FIFTH EDITION

Electric Motors and Drives

Fundamentals, Types and Applications

Austin Hughes
Bill Drury

Newnes

Electric Motors and Drives
Fundamentals, Types and Applications

Electric Motors and Drives

Fundamentals, Types and Applications

Fifth Edition

Austin Hughes

Bill Drury

ELSEVIER

Newnes
An imprint of Elsevier

Newnes is an imprint of Elsevier
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, United Kingdom
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-08-102615-1

For information on all Newnes publications
visit our website at https://www.elsevier.com/books-and-journals

Publisher: Joe Hayton
Acquisition Editor: Lisa Reading
Editorial Project Manager: Thomas Van Der Ploeg
Production Project Manager: Kamesh Ramajogi
Cover Designer: Matthew Limbert

Typeset by SPi Global, India

Preface

The fifth edition, like its predecessors, is intended primarily for non-specialist users or students of electric motors and drives, but we are pleased to have learned that many researchers and specialist industrialists have also acknowledged the value of the book in providing a clear understanding of the fundamentals of the subject.

Our aim is to bridge the gap between specialist textbooks (which are pitched at a level which is too academic/analytical for the average user) and the more prosaic handbooks which are full of detailed information but provide little opportunity for the development of any real insight. We intend to continue what has been a successful formula by providing the reader with an understanding of how each motor and drive system works, in the belief that it is only by knowing what should happen (and why) that informed judgements and sound comparisons can be made.

Given that the book is aimed at readers from a range of disciplines, sections of the book are of necessity devoted to introductory material. The first and second chapters therefore provide a gentle introduction to electromagnetic energy conversion and power electronics respectively. Many of the basic ideas introduced here crop up frequently throughout the book (and indeed are deliberately repeated to emphasise their importance), so unless the reader is already well-versed in the fundamentals it would be wise to absorb the first two chapters before tackling the later material. At various points later in the book we include more tutorial material, e.g. in Chapter 8 where we unravel the mysteries of field-oriented control.

The book covers all of the most important types of motor and drive, including conventional and brushless d.c., induction motors, synchronous motors of all types (including synchronous reluctance motors and salient Permanent Magnet motors), switched reluctance, and stepping motors. Induction motors, synchronous motors and their drives are given most attention, reflecting their dominant market position in terms of numbers. Conventional d.c. machines and drives are deliberately introduced early on, despite their much-reduced importance: this is partly because understanding is relatively easy, but primarily because the fundamental principles that emerge carry forward to the other types. Experience shows that readers who manage to grasp the principles of the d.c. drive will find this know-how invaluable in dealing with other more challenging types.

The fifth edition has been completely revised, updated and expanded. Modest changes to Chapter 2 now allow all of the most important converter topologies to be brought together, and the treatment of inverters now includes more detail of the PWM switching waveforms. The middle chapters dealing with the fundamentals of d.c. and induction motors and their drives have required only a few minor additions, whereas major changes and additions have been made to Chapters 8 and 9.

The new Chapter 8 is devoted exclusively to the treatment of Field Oriented control, reflecting its increasing importance for both induction and synchronous motor drives. The principles of control are explained in a unique physically-based way that builds on the understanding of motor behaviour developed earlier in the book: we believe that the largely non-mathematical treatment will dispel much of the mystique surrounding what is often regarded as a difficult topic.

The discussion of synchronous, permanent magnet and reluctance motors and drives has been greatly expanded in the new Chapter 9. There has been significant innovation in this area since the fourth edition, particularly in the automotive, aircraft and industrial sectors, with novel motor topologies emerging, including hybrid designs that combine permanent magnet and reluctance effects. We now include a physical basis for understanding and quantifying torque production in these machines, and this leads to simple pictures that illuminate the control conditions required to optimise torque.

We have responded to requests by providing Review Questions at the end of each chapter, together with fully-worked solutions intended to guide the reader through a logical approach to the question, thereby reinforcing knowledge and understanding.

Younger readers may be unaware of the radical changes that have taken place over the past 60 years, so a couple of paragraphs are appropriate to put the current scene into perspective. For more than a century, many different types of motor were developed, and each became closely associated with a particular application. Traction, for example, was seen as the exclusive preserve of the series d.c. motor, whereas the shunt d.c. motor, though outwardly indistinguishable, was seen as being quite unsuited to traction applications. The cage induction motor was (and still is) the most numerous type but was judged as being suited only to applications which called for constant speed. The reason for the plethora of motor types was that there was no easy way of varying the supply voltage and/or frequency to obtain speed control, and designers were therefore forced to seek ways of providing for control of speed within the motor itself. All sorts of ingenious arrangements and interconnections of motor windings were invented, but even the best motors had a limited operating range, and they all required bulky electromechanical control gear.

All this changed from the early 1960s, when power electronics began to make an impact. The first major breakthrough came with the thyristor, which provided a relatively cheap, compact, and easily controlled variable-speed drive using the d.c. motor. In the 1970s the second major breakthrough resulted from the development of power electronic inverters, providing a three-phase variable-frequency supply for the cage induction motor and thereby enabling its speed to be controlled. These major developments resulted in the demise of many of the special motors, leaving the majority of applications in the hands of comparatively few types. The switch from analogue to digital control also represented significant progress, but it was the availability of cheap digital processors that sparked the most recent leap forward. Real time modelling and simulation are now incorporated as standard into induction and synchronous motor drives, thereby allowing them to achieve levels of dynamic performance that had long been considered impossible.

The informal style of the book reflects our belief that the difficulty of coming to grips with new ideas should not be disguised. The level at which to pitch the material was based on feedback from previous editions which supported our view that a mainly descriptive approach with physical explanations would be most appropriate, with mathematics kept to a minimum to assist digestion. The most important concepts (such as the inherent e.m.f. feedback in motors, the need for a switching strategy in converters, and the importance of stored energy) are deliberately reiterated to reinforce understanding, but should not prove too tiresome for readers who have already 'got the message'. We have deliberately not included any computed magnetic field plots, nor any results from the excellent motor simulation packages that are now available because experience suggests that simplified diagrams are actually better as learning vehicles.

At the time of writing, the $150 Bn Electric Motors and Drives market is seeing tremendous innovation and growth driven by the electrification of automotives, aeroplanes and the fourth industrial revolution 'Industry 4.0', all of which are enabled by motors and drives. At such an exciting time it is important to understand the fundamental, underpinning technology—that is the mission of this book.

Finally, we welcome feedback, either via the publisher, or using the email addresses below.

Bill Drury (w.drury@btinternet.com)
Austin Hughes (a.hughes@leeds.ac.uk)

Chapter 1

Electric motors—The basics

1.1 Introduction

Electric motors are so much a part of everyday life that we seldom give them a second thought. When we switch on an ancient electric drill, for example, we confidently expect it to run rapidly up to the correct speed, and we don't question how it knows what speed to run at, nor how it is that once enough energy has been drawn from the supply to bring it up to speed, the power drawn falls to a very low level. When we put the drill to work it draws more power, and when we finish the power drawn from the supply reduces automatically, without intervention on our part.

The humble motor, consisting of nothing more than an arrangement of copper coils and steel laminations, is clearly rather a clever energy converter, which warrants serious consideration. By gaining a basic understanding of how the motor works, we will be able to appreciate its potential and its limitations, and (in later chapters) see how its already remarkable performance is dramatically enhanced by the addition of external electronic controls.

The great majority of electric motors have a shaft which rotates, but linear electric motors have niche applications, and whilst they appear very different from their rotating sister, their principle of operation is the same.

This chapter deals with the basic mechanisms of motor operation, so readers who are already familiar with such matters as magnetic flux, magnetic and electric circuits, torque, and motional e.m.f. (electromotive force) can probably afford to skim over much of it. In the course of the discussion, however, several very important general principles and guidelines emerge. These apply to all types of motor and are summarised in Section 1.9. Experience shows that anyone who has a good grasp of these basic principles will be well equipped to weigh the pros and cons of the different types of motor, so all readers are urged to absorb them before tackling other parts of the book.

1.2 Producing rotation

Nearly all motors exploit the force which is exerted on a current-carrying conductor placed in a magnetic field. The force can be demonstrated by placing a

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00001-5

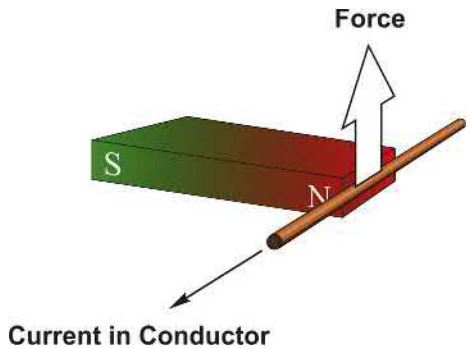

FIG. 1.1 Mechanical force produced on a current-carrying wire in a magnetic field.

bar magnet near a wire carrying current (Fig. 1.1), but anyone trying the experiment will probably be disappointed to discover how feeble the force is, and will doubtless be left wondering how such an unpromising effect can be used to make effective motors.

We will see that in order to make the most of the mechanism, we need to arrange for there to be a very strong magnetic field, and for it to interact with many conductors, each carrying as much current as possible. We will also see later that although the magnetic field (or 'excitation') is essential to the working of the motor, it acts only as a catalyst, and all of the mechanical output power comes from the electrical supply to the conductors on which the force is developed.

It will emerge later that in some motors the parts of the machine responsible for the excitation and for the energy converting functions are distinct and self-evident. In the d.c. motor, for example, the excitation is provided either by permanent magnets or by field coils wrapped around clearly-defined projecting field poles on the stationary part, while the conductors on which force is developed are on the rotor and supplied with current via sliding contacts known as brushes. In many motors, however, there is no such clear-cut physical distinction between the 'excitation' and the 'energy-converting' parts of the machine, and a single stationary winding serves both purposes. Nevertheless, we will find that identifying and separating the excitation and energy-converting functions is always helpful to understanding both how motors of all types operate, and their performance characteristics.

Returning to the matter of force on a single conductor, we will look first at what determines the magnitude and direction of the force, before turning to ways in which the mechanism is exploited to produce rotation. The concept of the magnetic circuit will have to be explored, since this is central to understanding why motors have the shapes they do. Before that a brief introduction to the magnetic field and magnetic flux and flux density is included for those who are not already familiar with the ideas involved.

1.2.1 Magnetic field and magnetic flux

When a current-carrying conductor is placed in a magnetic field, it experiences a force. Experiment shows that the magnitude of the force depends directly on the current in the wire, and the strength of the magnetic field, and that the force is greatest when the magnetic field is perpendicular to the conductor.

In the set-up shown in Fig. 1.1, the source of the magnetic field is a bar magnet, which produces a magnetic field as shown in Fig. 1.2.

The notion of a ‘magnetic field’ surrounding a magnet is an abstract idea that helps us to come to grips with the mysterious phenomenon of magnetism: it not only provides us with a convenient pictorial way of picturing the directional effects, but it also allows us to quantify the ‘strength’ of the magnetism and hence permits us to predict the various effects produced by it.

The dotted lines in Fig. 1.2 are referred to as magnetic flux lines, or simply flux lines. They indicate the direction along which iron filings (or small steel pins) would align themselves when placed in the field of the bar magnet. Steel pins have no initial magnetic field of their own, so there is no reason why one end or the other of the pins should point to a particular pole of the bar magnet.

However, when we put a compass needle (which is itself a permanent magnet) in the field we find that it aligns itself as shown in Fig. 1.2. In the upper half of the figure, the S end of the diamond-shaped compass settles closest to the N pole of the magnet, while in the lower half of the figure, the N end of the compass seeks the S of the magnet. This immediately suggests that there is a direction associated with the lines of flux, as shown by the arrows on the flux lines, which conventionally are taken as positively directed from the N to the S pole of the bar magnet.

The sketch in Fig. 1.2 might suggest that there is a ‘source’ near the top of the bar magnet, from which flux lines emanate before making their way to a corresponding ‘sink’ at the bottom. However, if we were to look at the flux lines

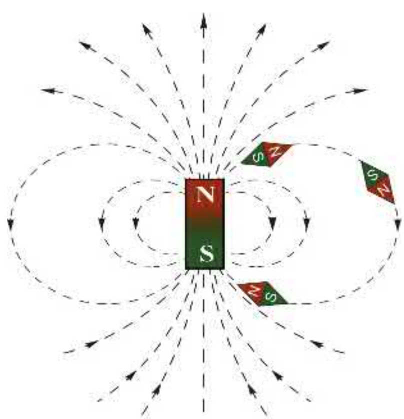

FIG. 1.2 Magnetic flux lines produced by a permanent magnet.

inside the magnet, we would find that they were continuous, with no 'start' or 'finish'. (In Fig. 1.2 the internal flux lines have been omitted for the sake of clarity, but a very similar field pattern is produced by a circular coil of wire carrying a direct current—see Fig. 1.7 where the continuity of the flux lines is clear.) Magnetic flux lines always form closed paths, as we will see when we look at the 'magnetic circuit', and draw a parallel with the electric circuit, in which the current is also a continuous quantity. (There must be a 'cause' of the magnetic flux, of course, and in a permanent magnet this is usually pictured in terms of atomic-level circulating currents within the magnet material. Fortunately, discussion at this physical level is not necessary for our purposes.)

1.2.2 Magnetic flux density

As well as showing direction, the flux plots also convey information about the intensity of the magnetic field. To achieve this, we introduce the idea that between every pair of flux lines (and for a given depth into the paper) there is the same 'quantity' of magnetic flux. Some people have no difficulty with such a concept, while others find that the notion of quantifying something so abstract represents a serious intellectual challenge. But whether the approach seems obvious or not, there is no denying the practical utility of quantifying the mysterious stuff we call magnetic flux, and it leads us next to the very important idea of magnetic flux density ($\boldsymbol{B}$).

When the flux lines are close together, the 'tube' of flux is squashed into a smaller space, whereas when the lines are further apart the same tube of flux has more breathing space. The flux density ($\boldsymbol{B}$) is simply the flux in the 'tube' (Φ) divided by the cross-sectional area (A) of the tube, i.e.

$$\boldsymbol{B} = \frac{\Phi}{A} \tag{1.1}$$

The flux density is a vector quantity, and is therefore often written in bold type: its magnitude is given by Eq. (1.1), and its direction is that of the prevailing flux lines at each point. Near the top of the magnet in Fig. 1.2, for example, the flux density will be large (because the flux is squashed into a small area), and pointing upwards, whereas on the equator and far out from the body of the magnet the flux density will be small and directed downwards.

We will see later that in order to create high flux densities in motors, the flux spends most of its life inside well-defined 'magnetic circuits' made of iron or steel, within which the flux lines spread out uniformly to take full advantage of the available area. In the case shown in Fig. 1.3, for example, the cross-sectional area of the iron at bb' is twice that at aa', but the flux is constant so the flux density at bb' is half that at aa'.

It remains to specify units for quantity of flux, and flux density. In the SI system, the unit of magnetic flux is the Weber (Wb). If one Weber of flux is distributed uniformly across an area of one square metre perpendicular to the

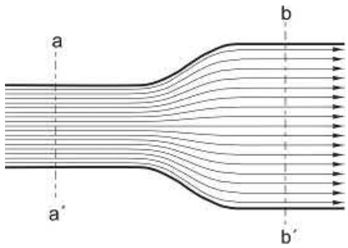

FIG. 1.3 Magnetic flux lines inside part of an iron magnetic circuit.

flux, the flux density is clearly one Weber per square metre (Wb/m^2). This was the unit of $\boldsymbol{B}$ until about 60 years ago, when it was decided that one Weber per square metre would henceforth be known as one Tesla (T), in honour of Nikola Tesla whose is generally credited with inventing the induction motor. The widespread use of B (measured in Tesla) in the design stage of all types of electromagnetic apparatus means that we are constantly reminded of the importance of Tesla; but at the same time one has to acknowledge that the outdated unit did have the advantage of conveying directly what flux density is, i.e. flux divided by area.

The flux in a 1 kW motor will be perhaps a few tens of milliwebers, and a small bar magnet would probably only produce a few microwebers. On the other hand, values of flux density are typically around 1 Tesla in most motors (regardless of type and rating), which is a reflection of the fact that although the quantity of flux in the 1 kW motor is small, it is also spread over a small area.

1.2.3 Force on a conductor

We now return to the production of force on a current-carrying wire placed in a magnetic field, as revealed by the set-up shown in Fig. 1.1.

The force is shown in Fig. 1.1: it is at right angles to both the current and the magnetic flux density, and its direction can be found using Fleming's left hand rule. If we picture the thumb, first and middle fingers held mutually perpendicular, then the **f**irst finger represents the **f**ield or **f**lux density ($\boldsymbol{B}$), the m**i**ddle finger represents the current ($\boldsymbol{I}$), and the thu**m**b then indicates the direction of **m**otion, as shown in Fig. 1.4.

Clearly, if either the field or the current is reversed, the force acts downwards, and if both are reversed, the direction of the force remains the same.

We find by experiment that if we double either the current or the flux density, we double the force, while doubling both causes the force to increase by a factor of four. But how about quantifying the force? We need to express the force in terms of the product of the current and the magnetic flux density, and this turns out to be very straightforward when we work in SI units.

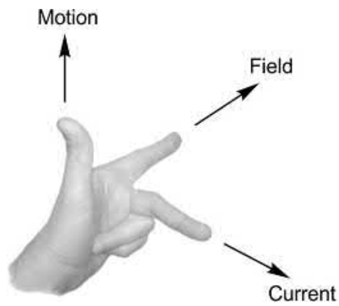

FIG. 1.4 Fleming's Left Hand rule for finding direction of force.

The force $\boldsymbol{F}$ on a wire of length $\boldsymbol{l}$, carrying a current $\boldsymbol{I}$ and exposed to a uniform magnetic flux density $\boldsymbol{B}$ throughout its length is given by the simple expression

$$\boldsymbol{F} = \boldsymbol{B}\boldsymbol{I}\boldsymbol{l} \tag{1.2}$$

In Eq. (1.2), F is in Newtons when B is in Tesla, I in Amps, and $\boldsymbol{l}$ in metres.

This is a delightfully simple formula, and it may come as a surprise to some readers that there are no constants of proportionality involved in Eq. (1.2). The simplicity is not a coincidence, but stems from the fact that the unit of current (the Ampere) is actually defined in terms of force.

Eq. (1.2) only applies when the current is perpendicular to the field. If this condition is not met, the force on the conductor will be less; and in the extreme case where the current was in the same direction as the field, the force would fall to zero. However, every sensible motor designer knows that to get the best out of the magnetic field it has to be perpendicular to the conductors, and so it is safe to assume in the subsequent discussion that B and I are always perpendicular. In the remainder of this book, it will be assumed that the flux density and current are mutually perpendicular, and this is why, although $\boldsymbol{B}$ is a vector quantity (and would usually be denoted by bold type), we can drop the bold notation because the direction is implicit and we are only interested in the magnitude.

The reason for the very low force detected in the experiment with the bar magnet is revealed by Eq. (1.2). To obtain a high force, we must have a high flux density, and a lot of current. The flux density at the ends of a bar magnet is low, perhaps 0.1 Tesla, so a wire carrying 1 Amp will experience a force of only 0.1 N (approximately 10 g wt) per metre. Since the flux density will be confined to perhaps 1 cm across the end face of the magnet, the total force on the wire will be only 0.1 g wt. This would be barely detectable, and is too low to be of any use in a decent motor. So how is more force obtained?

The first step is to obtain the highest possible flux density. This is achieved by designing a 'good' magnetic circuit, and is discussed next. Secondly, as many conductors as possible must be packed in the space where the magnetic field exists, and each conductor must carry as much current as it can without heating up to a dangerous temperature. In this way, impressive forces can be obtained from modestly sized devices, as anyone who has tried to stop an electric drill by grasping the chuck will testify.

1.3 Magnetic circuits

So far we have assumed that the source of the magnetic field is a permanent magnet. This is a convenient starting point as all of us are familiar with magnets. But in the majority of motors, the magnetic field is produced by coils of wire carrying current, so it is appropriate that we look at how we arrange the coils and their associated 'magnetic circuit' so as to produce high magnetic fields which then interact with other current-carrying conductors to produce force, and hence rotation.

First, we look at the simplest possible case of the magnetic field surrounding an isolated long straight wire carrying a steady current (Fig. 1.5). (In the figure, the + sign indicates that current is flowing into the paper, while a dot is used to signify current out of the paper: these symbols can perhaps be remembered by picturing an arrow or dart, with the cross being the rear view of the fletch, and the dot being the approaching point.) The flux lines form circles concentric with the wire, the field strength being greatest close to the wire. As might be expected, the field strength at any point is directly proportional to the current. The convention for determining the direction of the field is that the positive direction is taken to be the direction that a right-handed corkscrew must be rotated to move in the direction of the current.

Fig. 1.5 is somewhat artificial as current can only flow in a complete circuit, so there must always be a return path. If we imagine a parallel 'go' and 'return' circuit, for example, the field can be obtained by superimposing the field produced by the positive current in the go side with the field produced by the negative current in the return side, as shown in Fig. 1.6.

We note how the field is increased in the region between the conductors, and reduced in the regions outside. Although Fig. 1.6 strictly only applies to an infinitely long pair of straight conductors, it will probably not come as a surprise to learn that the field produced by a single turn of wire of rectangular, square or round form is very much the same as that shown in Fig. 1.6. This enables us to

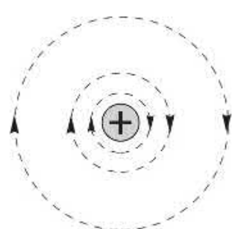

FIG. 1.5 Magnetic flux lines produced by a straight, current-carrying wire.

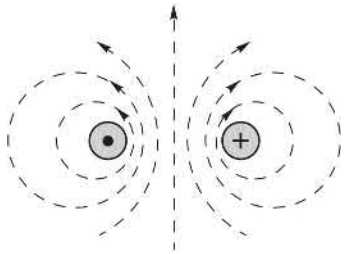

FIG. 1.6 Magnetic flux lines produced by current in a parallel go and return circuit.

build up a picture of the field that would be produced—in air—by the sort of coils used in motors, which typically have many turns, as shown for example in Fig. 1.7.

The coil itself is shown on the left in Fig. 1.7 while the flux pattern produced is shown on the right. Each turn in the coil produces a field pattern, and when all the individual field components are superimposed we see that the field inside the coil is substantially increased and that the closed flux paths closely resemble those of the bar magnet that we looked at earlier. The air surrounding the sources of the field offers a homogeneous path for the flux, so once the tubes of flux escape from the concentrating influence of the source, they are free to spread out into the whole of the surrounding space. Recalling that between each pair of flux lines there is an equal amount of flux, we see that because the flux lines spread out as they leave the confines of the coil, the flux density is much lower outside than inside: for example, if the distance 'b' is say four times 'a' the flux density B_b is a quarter of B_a.

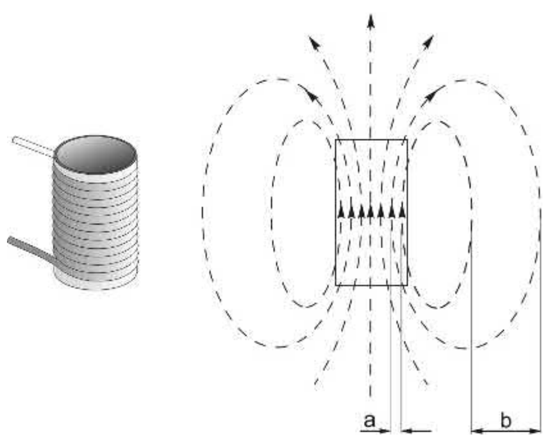

FIG. 1.7 Multi-turn cylindrical coil and pattern of magnetic flux produced by current in the coil. (For the sake of clarity, only the outline of the coil is shown on the right.)

Although the flux density inside the coil is higher than outside, we would find that the flux densities which we could achieve are still too low to be of use in a motor. What is needed is firstly a way of increasing the flux density, and secondly a means for concentrating the flux and preventing it from spreading out into the surrounding space.

1.3.1 Magnetomotive force (m.m.f.)

One obvious way to increase the flux density is to increase the current in the coil, or to add more turns. We find that if we double the current, or the number of turns, we double the total flux, thereby doubling the flux density everywhere.

We quantify the ability of the coil to produce flux in terms of its Magnetomotive Force (m.m.f.). The m.m.f. of the coil is simply the product of the number of turns (N) and the current (I), and is thus expressed in Ampere-turns. A given m.m.f. can be obtained with a large number of turns of thin wire carrying a low current, or a few turns of thick wire carrying a high current: as long as the product NI is constant, the m.m.f. is the same.

1.3.2 Electric circuit analogy

We have seen that the magnetic flux which is set up is proportional to the m.m.f. driving it. This points to a parallel with the electric circuit, where the current (Amps) which flows is proportional to the electromotive force (e.m.f., in Volts) driving it.

In the electric circuit, current and e.m.f. are related by Ohm's Law, which is

$$\text{Current} = \frac{\text{e.m.f.}}{\text{Resistance}}, \quad \text{i.e. } I = \frac{V}{R} \tag{1.3}$$

For a given source e.m.f. (Volts), the current depends inversely on the resistance of the circuit, so to obtain more current, the resistance of the circuit has to be reduced.

We can make use of an equivalent 'magnetic Ohm's law' by introducing the idea of Reluctance ($\mathcal{R}$). The reluctance gives a measure of how difficult it is for the magnetic flux to complete its circuit, in the same way that resistance indicates how much opposition the current encounters in the electric circuit.

The magnetic Ohm's law is then

$$\text{Flux} = \frac{\text{m.m.f.}}{\text{Reluctance}}, \quad \text{i.e. } \Phi = \frac{NI}{\mathcal{R}} \tag{1.4}$$

We see from Eq. (1.4) that to increase the flux for a given m.m.f. we need to reduce the reluctance of the magnetic circuit. In the case of the example (Fig. 1.7), this means we must replace as much as possible of the air path (which is a 'poor' magnetic material, and therefore constitutes a high reluctance) with a

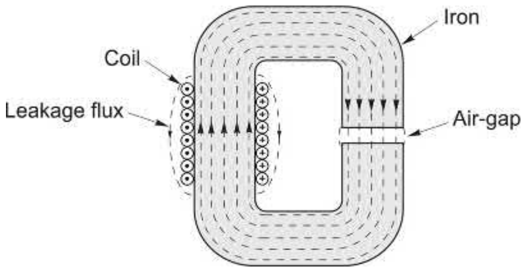

FIG. 1.8 Flux lines inside low-reluctance magnetic circuit with air-gap.

'good' magnetic material, thereby reducing the reluctance and resulting in a higher flux for a given m.m.f.

The material which we usually choose is good quality magnetic steel, which for historical reasons is often referred to as 'iron'. This brings several very dramatic and desirable benefits, as shown in Fig. 1.8.

Firstly, the reluctance of the iron paths is very much less than the air paths which they have replaced, so the total flux produced for a given m.m.f. is very much greater. (Strictly speaking therefore, if the m.m.f.'s and cross-sections of the coils in Figs. 1.7 and 1.8 are the same, many more flux lines should be shown in Fig. 1.8 than in Fig. 1.7, but for the sake of clarity a similar number are indicated.) Secondly, almost all the flux is confined within the iron, rather than spreading out into the surrounding air. We can therefore shape the iron parts of the magnetic circuit as shown in Fig. 1.8 in order to guide the flux to wherever it is needed. And finally, we see that inside the iron, the flux density remains constant over the whole uniform cross-section, there being so little reluctance that there is no noticeable tendency for the flux to crowd to one side or another.

Before moving on to the matter of the air-gap, a question which is often asked is whether it is important for the coils to be wound tightly onto the magnetic circuit, and whether, if there is a multi-layer winding, the outer turns are as effective as the inner ones. The answer, happily, is that the total m.m.f. is determined solely by the number of turns and the current, and therefore every complete turn makes the same contribution to the total m.m.f., regardless of whether it happens to be tightly or loosely wound. Of course it does make sense for the coils to be wound as tightly as is practicable, since this not only minimises the resistance of the coil (and thereby reduces the heat loss) but also makes it easier for the heat generated to be conducted away to the frame of the machine.

1.3.3 The air-gap

In motors, we intend to use the high flux density to develop force on current-carrying conductors, which must then move to produce useful work. We have now seen how to create a high flux density in a magnetic circuit, but, of course, it is not possible to put current-carrying conductors inside the iron. We therefore arrange for an air-gap in the magnetic circuit, as shown in Fig. 1.8. We will see shortly (see for example Fig. 1.12) that the conductors on which the force is to be produced will be placed in this air-gap region. So although we will find that the reluctance of the air-gap is unwelcome from the magnetic circuit viewpoint, it is obviously necessary for there to be mechanical clearance to allow the rotor to rotate.

If the air-gap is relatively small, as in motors, we find that the flux jumps across the air-gap as shown in Fig. 1.8, with very little tendency to balloon out into the surrounding air. With most of the flux lines going straight across the air-gap, the flux density in the gap region has the same high value as it does inside the iron.

In the majority of magnetic circuits with one or more air-gaps, the reluctance of the iron parts is very much less than the reluctance of the gaps. At first sight this can seem surprising, since the distance across the gap is so much less than the rest of the path through the iron. The fact that the air-gap dominates the reluctance is simply a reflection of how poor air is as a magnetic medium, compared with iron. To put the comparison in perspective, if we calculate the reluctances of two paths of equal length and cross-sectional area, one being in iron and the other in air, the reluctance of the air path will typically be 1000 times greater than the reluctance of the iron path.

Returning to the analogy with the electric circuit, the role of the iron parts of the magnetic circuit can be likened to that of the copper wires in the electric circuit. Both offer little opposition to flow (so that a negligible fraction of the driving force (m.m.f. or e.m.f.) is wasted in conveying the flow to where it is usefully exploited) and both can be shaped to guide the flow to its destination. There is one important difference, however. In the electric circuit, no current will flow until the circuit is completed, after which all the current is confined inside the wires. With an iron magnetic circuit, some flux can flow (in the surrounding air) even before the iron is installed. And although most of the flux will subsequently take the easy route through the iron, some will still leak into the air, as shown in Fig. 1.8. We will not pursue leakage flux here, though it is sometimes important, as will be seen later.

1.3.4 Reluctance and air-gap flux densities

If we neglect the reluctance of the iron parts of a magnetic circuit, it is easy to estimate the flux density in the air-gap. Since the iron parts are then in effect ‘perfect conductors’ of flux, none of the source m.m.f. (NI) is used in driving the flux through the iron parts, and all of it is available to push the flux across the air-gap. The situation depicted in Fig. 1.8 therefore reduces to that shown in Fig. 1.9, where an m.m.f. of NI is applied directly across an air-gap of length g.

To determine how much flux will cross the gap, we need to know its reluctance. As might be expected the reluctance of any part of the magnetic circuit

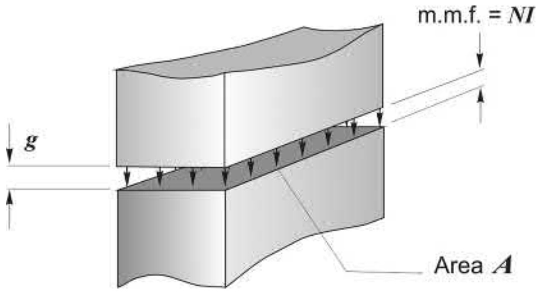

FIG. 1.9 Air-gap region, with m.m.f. acting cross opposing pole-faces.

depends on its dimensions, and on its magnetic properties, and the reluctance of a rectangular 'prism' of air, of cross-sectional area A and length g is given by

$$\mathcal{R}_g = \frac{g}{A\mu_0} \tag{1.5}$$

where μ_0 is the so-called 'primary magnetic constant' or 'permeability of free space'. Strictly, as its name implies, μ_o quantifies the magnetic properties of a vacuum, but for all engineering purposes the permeability of air is also μ_o. The value of the primary magnetic constant (μ_o) in the SI system is $4\pi \times 10^{-7}$ Henry/m: rather surprisingly, there is no widely-used name for the unit of reluctance.

(In passing, we should note that if we want to include the reluctance of the iron part of the magnetic circuit in our calculation, its reluctance would be given by

$$\mathcal{R}_{iron} = \frac{l_{iron}}{A\mu_{iron}}$$

and we would have to add this to the reluctance of the air-gap to obtain the total reluctance. However, as stated earlier, because the permeability of iron (μ_{iron}) is so much higher than μ_0, its reluctance will be very much less than the gap reluctance, despite the path length l_{iron} being considerably longer than the path length (g) in the air.)

Eq. (1.5) reveals the expected result that doubling the air-gap would double the reluctance (because the flux has twice as far to go), while doubling the area would halve the reluctance (because the flux has two equally appealing paths in parallel). To calculate the flux, Φ, we use the magnetic Ohm's law (Eq. 1.4), which gives

$$\Phi = \frac{\text{m.m.f.}}{\mathcal{R}}, \quad \text{i.e.} \quad \Phi = \frac{NIA\mu_0}{g} \tag{1.6}$$

We are usually interested in the flux density in the gap, rather than the total flux, so we use Eq. (1.1) to yield

$$B = \frac{\Phi}{A} = \frac{\mu_0 NI}{g} \tag{1.7}$$

Eq. (1.7) is delightfully simple, and from it we can calculate the air-gap flux density once we know the m.m.f. of the coil (*NI*) and the length of the gap (*g*). We do not need to know the details of the coil-winding as long as we know the product of the turns and the current, and neither do we need to know the cross-sectional area of the magnetic circuit in order to obtain the flux density (though we do if we want to know the total flux, see Eq. 1.6).

For example, suppose the magnetising coil has 250 turns, the current is 2 A, and the gap is 1 mm. The flux density is then given by

$$B = \frac{4\pi \times 10^{-7} \times 250 \times 2}{1 \times 10^{-3}} = 0.63\,\text{Tesla}$$

(We could of course create the same flux density with a coil of 50 turns carrying a current of 10 A, or any other combination of turns and current giving an m.m.f. of 500 Ampere-turns.)

If the cross-sectional area of the iron was constant at all points, the flux density would be 0.63 T everywhere. Sometimes, as has already been mentioned, the cross-section of the iron reduces at points away from the air-gap, as shown for example in Fig. 1.3. Because the flux is compressed in the narrower sections, the flux density is higher, and in Fig. 1.3 if the flux density at the air-gap and in the adjacent pole-faces is once again taken to be 0.63 T, then at the section aa' (where the area is only half that at the air-gap) the flux density will be $2 \times 0.63 = 1.26$ T.

1.3.5 Saturation

It would be reasonable to ask whether there is any limit to the flux density at which the iron can be operated. We can anticipate that there must be a limit, or else it would be possible to squash the flux into a vanishingly small cross-section, which we know is not the case. In fact there is a limit, though not a very sharply defined one.

Earlier we noted that the 'iron' has very little reluctance, at least not in comparison with air. Unfortunately this happy state of affairs is only true as long as the flux density remains below about 1.6–1.8 T, depending on the particular magnetic steel in question: if we try to work at higher flux densities, it begins to exhibit significant reluctance, and no longer behaves like an ideal conductor of flux. At these higher flux densities a significant proportion of the source m.m.f. is used in driving the flux through the iron. This situation is obviously undesirable, since less m.m.f. remains to drive the flux across the air-gap. So, just as we would not recommend the use of high-resistance supply leads to the load in an electric circuit, we must avoid overloading the iron parts of the magnetic circuit.

The emergence of significant reluctance as the flux density is raised is illustrated qualitatively in Fig. 1.10. When the reluctance begins to be appreciable,

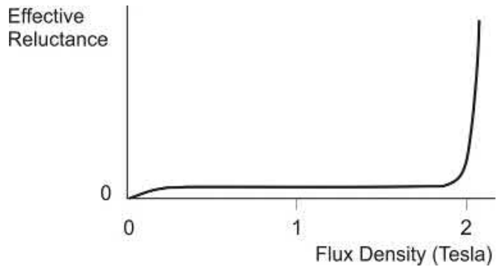

FIG. 1.10 Sketch showing how the effective reluctance of iron increases rapidly as the flux density approaches saturation.

the iron is said to be beginning to 'saturate'. The term is apt, because if we continue increasing the m.m.f. or reducing the area of the iron, we will eventually reach an almost constant flux density, typically around 2 T. To avoid the undesirable effects of saturation, the size of the iron parts of the magnetic circuit are usually chosen so that the flux density does not exceed about 1.5 T. At this level of flux density, the reluctance of the iron parts will remain small in comparison with the air-gap.

1.3.6 Magnetic circuits in motors

The reader may be wondering why so much attention has been focused on the gapped C-core magnetic circuit, when it appears to bear little resemblance to the magnetic circuits found in motors. We will now see that it is actually a short step from the C-core to a typical motor magnetic circuit, and that no fundamentally new ideas are involved.

The evolution from C-core to motor geometry is shown in Fig. 1.11, which should be largely self-explanatory, and relates to the field system of a traditional d.c. motor.

We note that the first stage of evolution (Fig. 1.11, left) results in the original single gap of length g being split into two gaps of length $g/2$, reflecting the requirement for the rotor to be able to turn. At the same time the single magnetising coil is split into two to preserve symmetry. (Relocating the magnetising

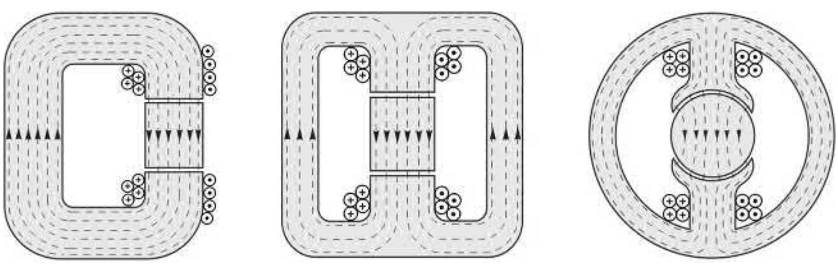

FIG. 1.11 Evolution of d.c. motor magnetic circuit from gapped C-core.

coil at a different position around the magnetic circuit is of course in order, just as a battery can be placed anywhere in an electric circuit.) Next, (Fig. 1.11, centre) the single magnetic path is split into two parallel paths of half the original cross-section, each of which carries half of the flux: and finally (Fig. 1.11, right), the flux paths and pole faces are curved to match the cylindrical rotor. The coil now has several layers in order to fit the available space, but as discussed earlier this has no adverse effect on the m.m.f. The air-gap is still small, so the flux crosses radially to the rotor.

1.4 Torque production

Having designed the magnetic circuit to give a high flux density under the poles, we must obtain maximum benefit from it. We therefore need to arrange a set of conductors, fixed to the rotor, as shown in Fig. 1.12, and to ensure that conductors under a N-pole (on the left) carry positive current (into the paper), while those under the S-pole carry negative current. The tangential electromagnetic ('*BIl*') force (see Eq. 1.2) on all the positive conductors will be downwards (towards the bottom of the page), while the force on the negative ones will be upwards (towards the top of the page): a torque will therefore be exerted on the rotor, which will be caused to rotate.

(The observant reader spotting that some of the conductors appear to have no current in them will find the explanation later, in Chapter 3.)

At this point we should pause and address three questions that often crop up when these ideas are being developed. The first is to ask why we have made no reference to the magnetic field produced by the current-carrying conductors on the rotor. Surely they too will produce a magnetic field, which will presumably interfere with the original field in the air-gap, in which case perhaps the expression used to calculate the force on the conductor will no longer be valid.

The answer to this very perceptive question is that the field produced by the current-carrying conductors on the rotor certainly will modify the original field (i.e. the field that was present when there was no current in the rotor conductors.) But in the majority of motors, the force on the conductor can be calculated correctly from the product of the current and the 'original' field. This is very

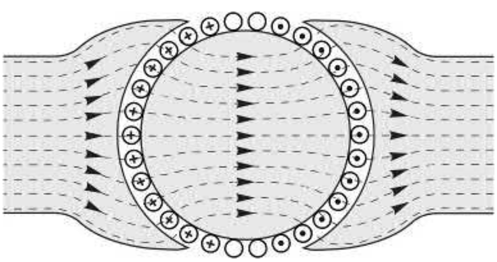

FIG. 1.12 Current-carrying conductors on rotor, positioned to maximise torque. (The source of the magnetic flux lines (arrowed) is not shown.)

fortunate from the point of view of calculating the force, but also has a logical feel to it. For example in Fig. 1.1, we would not expect any force on the current-carrying conductor if there was no externally applied field, even though the current in the conductor will produce its own field (upwards on one side of the conductor and downwards on the other). So it seems right that since we only obtain a force when there is an external field, all of the force must be due to that field alone. (In Chapter 3 we will discover that the field produced by the rotor conductors is known as 'armature reaction', and that, especially when the magnetic circuit becomes saturated, its undesirable effects may be combatted by fitting additional windings designed to nullify the armature field.)

The second question arises when we think about the action and reaction principle. When there is a torque on the rotor, there is presumably an equal and opposite torque on the stator; and therefore we might wonder if the mechanism of torque production could be pictured using the same ideas as we used for obtaining the rotor torque. The answer is yes, there is always an equal and opposite torque on the stator, which is why it is usually important to bolt a motor down securely. In some machines (e.g. the induction motor) it is easy to see that torque is produced on the stator by the interaction of the air-gap flux density and the stator currents, in exactly the same way that the flux density interacts with the rotor currents to produce torque on the rotor. In other motors (e.g. the d.c. motor we have been looking at), there is no simple physical argument which can be advanced to derive the torque on the stator, but nevertheless it is equal and opposite to the torque on the rotor.

The final question relates to the similarity between the set-up shown in Fig. 1.11 and the field patterns produced for example by the electromagnets used to lift car bodies in a scrap yard. From what we know of the large force of attraction that lifting magnets can produce, might we not expect there to be a large radial force between the stator pole and the iron body of the rotor? And if there is, what is to prevent the rotor from being pulled across to the stator?

Again the affirmative answer is that there is indeed a radial force due to magnetic attraction, exactly as in a lifting magnet or relay, although the mechanism whereby the magnetic field exerts a pull as it enters iron or steel is different from the '*BIl* 'force we have been looking at so far.

It turns out that the force of attraction per unit area of pole-face is proportional to the square of the radial flux density, and with typical air-gap flux densities of up to 1 T in motors, the force per unit area of rotor surface works out to be about 40 N/cm^2. This indicates that the total radial force can be very large: for example the force of attraction on a small pole face of only 5 cm $\times$ 10 cm is 2000 N, or about 200 kg. This force contributes nothing to the torque of the motor, and is merely an unwelcome by-product of the '*BIl*' mechanism we employ to produce tangential force on the rotor conductors.

In most machines the radial magnetic force under each pole is actually a good deal bigger than the tangential electromagnetic force on the rotor conductors, and as the question implies, it tends to pull the rotor onto the pole. However, the majority of motors are constructed with an even number of poles equally spaced around the rotor, and the flux density in each pole is the same, so that—in theory at least—the resultant force on the complete rotor is zero. In practice, even a small eccentricity will cause the field to be stronger under the poles where the air-gap is smaller, and this will give rise to an unbalanced pull, resulting in noisy running and rapid bearing wear.

In the majority of motors we can quantify the torque via the *'BIl'* approach. The source of the magnetic flux density B may be a winding, as in Fig. 1.11, or a permanent magnet. The source (or 'excitation') may be located on the stator (as implied in Fig. 1.12) or on the rotor. If the source of B is on the stator, the current carrying conductors on which the force is developed are located on the rotor, whereas if the excitation is on the rotor, the active conductors are on the stator. In all of these the large radial magnetic forces discussed above are an unwanted by-product.

However, in some motors the stator and/or rotor geometry is arranged so that some of the flux crossing the air-gap to the rotor produces tangential forces (and thus torque) directly on the rotor iron, without any rotor currents. We will see in later chapters that in some of these 'reluctance' machines, we can still quantify the torque using the 'BIl' method, while in others we have to employ an alternative to the '*BIl*' method to obtain the turning forces.

1.4.1 Magnitude of torque

Returning to our original discussion, the force on each conductor is given by Eq. (1.2), and it follows that the total tangential force F depends on the flux density produced by the field winding, the number of conductors on the rotor, the current in each, and the length of the rotor. The resultant torque (T) depends on the radius of the rotor (r), and is given by

$$T = Fr \tag{1.8}$$

We will return to this after we examine the remarkable benefits gained by putting the rotor conductors into slots.

1.4.2 The beauty of slotting

If the conductors were mounted on the surface of the rotor iron, as in Fig. 1.12, the air-gap would have to be at least equal to the wire diameter, and the conductors would have to be secured to the rotor in order to transmit their turning force to it. The earliest motors were made like this, with string or tape to bind the conductors to the rotor.

Unfortunately, a large air-gap results in an unwelcome high reluctance in the magnetic circuit, and the field winding therefore needs many turns and a high current to produce the desired flux density in the air-gap. This means that the field winding becomes very bulky and consumes a lot of power. The early (nineteenth-century) pioneers soon hit upon the idea of partially sinking the conductors on the rotor into grooves machined parallel to the shaft, the intention being to allow the air-gap to be reduced so that the exciting windings could be smaller. This worked extremely well as it also provided a more positive location for the rotor conductors, and thus allowed the force on them to be transmitted more directly to the body of the rotor. Before long the conductors began to be recessed into ever deeper slots until finally (see Fig. 1.13) they no longer stood proud of the rotor surface and the air-gap could be made as small as was consistent with the need for mechanical clearances between the rotor and the stator. The new 'slotted' machines worked very well, and their pragmatic makers were unconcerned by rumblings of discontent from sceptical theorists.

The theorists of the time accepted that sinking conductors into slots allowed the air-gap to be made small, but argued that, as can be seen from Fig. 1.13, almost all the flux would now pass down the attractive low-reluctance path through the teeth, leaving the conductors exposed to the very low leakage flux density in the slots. Surely, they argued, little or no '*BIl*' force would be developed on the conductors, since they would only be exposed to a very low flux density.

The sceptics were right in that the flux does indeed flow down the teeth; but there was no denying that motors with slotted rotors produced the same torque as those with the conductors in the air-gap, provided that the average flux densities at the rotor surface were the same. So what could explain this seemingly too good to be true situation?

The search for an explanation preoccupied some of the leading thinkers long after slotting became the norm, but finally it became possible to show that what happens is that the total force remains the same as it would have been if the conductors were actually in the flux, but almost all of the tangential force now acts on the rotor teeth, rather than on the conductors themselves.

This is remarkably good news. By putting the conductors in slots, we simultaneously enable the reluctance of the magnetic circuit to be reduced, and

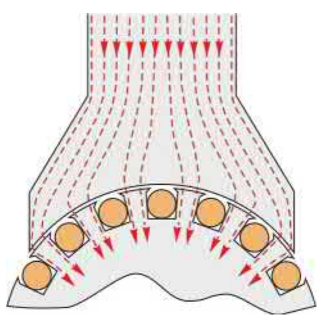

FIG. 1.13 Influence on flux paths when the rotor is slotted to accommodate conductors.

transfer the force from the conductors themselves to the sides of the iron teeth, which are robust and well able to transfer the resulting torque to the shaft. A further benefit is that the insulation around the conductors no longer has to transmit the tangential forces to the rotor, and its mechanical properties are thus less critical. Seldom can tentative experiments with one aim have yielded rewarding outcomes in almost every other relevant direction.

There are some snags, however. To maximise the torque, we will want as much current as possible in the rotor conductors. Naturally we will work the copper at the highest practicable current density (typically between 2 and 8 A/mm^2), but we will also want to maximise the cross-sectional area of the slots to accommodate as much copper as possible. This will push us in the direction of wide slots, and hence narrow teeth. But we recall that the flux has to pass radially down the teeth, so if we make the teeth too narrow, the iron in the teeth will saturate, and lead to a poor magnetic circuit. There is also the possibility of increasing the depth of the slots, but this cannot be taken too far or the centre region of the rotor iron—which has to carry the flux from one pole to another—will become so depleted that it too will saturate. Finally, an unwelcome mechanical effect of slotting it that it increases the frictional drag and acoustic noise, effects which are often minimised by filling the tops of the slot openings so that the rotor surface becomes smooth.

1.5 Torque and motor volume

In this section we look at what determines the torque that can be obtained from a rotor of a given size, and see how speed plays a key role in determining the power output.

The universal adoption of slotting to accommodate conductors means that a compromise is inevitable in the crucial air-gap region, and designers constantly have to exercise their skills to achieve the best balance between the conflicting demands on space made by the flux (radial) and the current (axial).

As in most engineering design, guidelines emerge as to what can be achieved in relation to particular sizes and types of machine, and motor designers usually work in terms of two parameters, the specific magnetic loading, and the specific electric loading. These parameters will seldom be made available to the user, but, together with the volume of the rotor, they define the torque that can be produced, and are therefore of fundamental importance. An awareness of the existence and significance of these parameters therefore helps the user to challenge any seemingly extravagant claims that may be encountered.

1.5.1 Specific loadings

The specific magnetic loading $(\overline{B})$ is the average of the magnitude of the radial flux density over the entire cylindrical surface of the rotor. Because of the slotting, the average flux density is always less than the flux density in the teeth, but in order to calculate the magnetic loading we picture the rotor as being smooth, and calculate the average flux density by dividing the total radial flux from each 'pole' by the surface area under the pole.

The specific electric loading (usually denoted by the symbol $(\overline{A})$, the $\overline{A}$ standing for Amperes) is the axial current per metre of circumference on the rotor. In a slotted rotor, the axial current is concentrated in the conductors within each slot, but to calculate $\overline{A}$ we picture the total current to be spread uniformly over the circumference (in a manner similar to that shown in Fig. 1.13, but with the individual conductors under each pole being represented by a uniformly distributed 'current sheet'). For example, if under a pole with a circumferential width of 10 cm we find that there are five slots, each carrying a current of 40 A, the electric loading is

$$\frac{5 \times 40}{0.1} = 2000\,\text{A/m}.$$

The discussion in Section 1.4 referred to the conflicting demands of flux and current, so it should be clear that if we seek to increase the electric loading, for example by widening the slots to accommodate more copper, we must be aware that the magnetic loading may have to be reduced because the narrower teeth will mean there is less area for the flux, and therefore a danger of saturating the iron.

Many factors influence the values which can be employed in motor design, but in essence the specific magnetic and electric loadings are limited by the properties of the materials (iron for the flux, and copper for the current), and by the cooling system employed to remove heat losses.

The specific magnetic loading does not vary greatly from one machine to another, because the saturation properties of most core steels are similar, so there is an upper limit to the flux density that can be achieved. On the other hand, quite wide variations occur in the specific electric loadings, depending on the type of cooling used.

Despite the low resistivity of the copper conductors, heat is generated by the flow of current, and the current must therefore be limited to a value such that the insulation is not damaged by an excessive temperature rise. The more effective the cooling system, the higher the electric loading can be. For example, if the motor is totally enclosed and has no internal fan, the current density in the copper has to be much lower than in a similar motor which has a fan to provide a continuous flow of ventilating air. Similarly, windings which are fully impregnated with varnish can be worked much harder than those which are surrounded by air, because the solid body of encapsulating varnish not only gives mechanical rigidity but also provides a much better thermal path along which the heat can flow to the stator body. Overall size also plays a part in determining permissible electric loading, with larger motors generally having higher values than small ones.

In practice, the important point to be borne in mind is that unless an exotic cooling system is employed, most motors (induction, d.c., etc.) of a particular size have more or less the same specific loadings, regardless of type. As we will now see, this in turn means that motors of similar size have broadly similar torque capabilities, regardless of the specific motor type/technology. This fact is not widely appreciated by users, but is always worth bearing in mind.

1.5.2 Torque and rotor volume

In the light of the earlier discussion, we can obtain the total tangential force by first considering an area of the rotor surface of width w and length L. The axial current flowing in the width w is given by $I = w\overline{A}$, and on average all of this current is exposed to radial flux density $\overline{B}$, so the tangential force is given (from Eq. 1.2) by $\overline{B} \times w\overline{A} \times L$. The area of the surface is wL so the force per unit area is $\overline{B} \times \overline{A}$. We see that the product of the two specific loadings expresses the average tangential stress over the rotor surface.

To obtain the total tangential force we must multiply by the area of the curved surface of the rotor, and to obtain the total torque we multiply the total force by the radius of the rotor. Hence for a rotor of diameter D and length L, the total torque is given by

$$T = \left(\overline{B}\overline{A}\right) \times (\pi D L) \times \frac{D}{2} = \frac{\pi}{2}\left(\overline{B}\overline{A}\right) D^2 L \tag{1.9}$$

What this equation tells us is extremely important. The term D^2L is proportional to the rotor volume, so we see that **for given values of the specific magnetic and electric loadings, the torque from any motor is proportional to the rotor volume.** We are at liberty to choose a long thin rotor or a short fat one, but once the rotor volume and specific loadings are specified, we have effectively determined the torque.

It is worth stressing that we have not focused on any particular type of motor, but have approached the question of torque production from a completely general viewpoint. In essence our conclusions reflect the fact that all motors are made from iron and copper, and differ only in the way these materials are disposed, and how hard they are worked.

We should also acknowledge that in practice it is the overall volume of the motor which is important, rather than the volume of the rotor. But again we find that, regardless of the type of motor, there is a fairly close relationship between the overall volume and the rotor volume, for motors of similar torque. We can therefore make the bold but generally accurate statement that the overall volume of a motor is determined by the torque it has to produce. There are of course exceptions to this rule, but as a general guideline for motor selection, it is extremely useful.

Having seen that torque depends on rotor volume, we must now turn our attention to the question of power output.

1.5.3 Output power—Importance of speed

Before deriving an expression for power a brief digression may be helpful for those who are more familiar with linear rather than rotary systems.

In the SI system, the unit of work or energy is the Joule (J). One Joule represents the work done by a force of 1 Newton moving 1 m in its own direction. Hence the work done (W) by a force F which moves a distance d is given by

$$W = F \times d$$

With F in Newtons and d in metres, W is clearly in Newton-metres (Nm), from which we see that a Newton-metre is the same as a Joule.

In rotary systems, it is more convenient to work in terms of torque and angular distance, rather than force and linear distance, but these are closely linked as we can see by considering what happens when a tangential force F is applied at a radius r from the centre of rotation. The torque is simply given by

$$T = F \times r$$

Now suppose that the arm turns through an angle θ, so that the circumferential distance travelled by the force is $r \times \theta$. The work done by the force is then given by

$$W = F \times (r \times \theta) = (F \times r) \times \theta = T \times \theta \tag{1.10}$$

We note that whereas in a linear system work is force times distance, in rotary terms work is torque times angle. The units of torque are Newton-metres, and the angle is measured in radians (which is dimensionless), so the units of work done are Nm, or Joules, as expected. (The fact that torque and work (or energy) are measured in the same units does not seem self-evident to the authors!)

To find the power, or the rate of working, we divide the work done by the time taken. In a linear system, and assuming that the velocity remains constant, power is therefore given by

$$P = \frac{W}{t} = \frac{F \times d}{t} = F \times v \tag{1.11}$$

where v is the linear velocity. The angular equivalent of this is

$$P = \frac{W}{t} = \frac{T \times \theta}{t} = T \times \omega \tag{1.12}$$

where ω is the (constant) angular velocity, in radians per second.

We can now express the power output in terms of the rotor dimensions and the specific loadings, using Eq. 1.9 which yields

$$P = T\omega = \frac{\pi}{2}\left(\overline{BA}\right) D^2 L \omega \tag{1.13}$$

Eq. (1.13) emphasises the importance of speed (ω) in determining power output. For given specific and magnetic loadings, if we want a motor of a given power we can choose between a large (and therefore expensive) low-speed motor or a small (and generally cheaper) high-speed one. The latter choice is preferred for most applications, even if some form of speed reduction (using belts or gears, for example) is needed, because the smaller motor is cheaper. Familiar examples include portable electric tools, where rotor speeds of 12,000 rev/min or more allow powers of hundreds of Watts to be obtained, and electric traction: in both the high motor speed is geared down for the final drive. In these examples, where volume and weight are at a premium, a direct drive would be out of the question.

1.5.4 Power density (specific output power)

By dividing Eq. (1.13) by the rotor volume, we obtain an expression for the specific power output (power per unit rotor volume), Q, given by

$$Q = 2\overline{B}\,\overline{A}\,\omega \tag{1.14}$$

The importance of this simple equation cannot be overemphasised. It is the fundamental design equation that governs the output of any 'BIl' machine, and thus applies to almost all motors.

To obtain the highest possible power from a given volume for given values of the specific magnetic and electric loadings, we must clearly operate the motor at the highest practicable speed. The one obvious disadvantage of a small high-speed motor and gearbox is that the acoustic noise (both from the motor itself and the from the power transmission) is higher than it would be from a larger direct drive motor. When noise must be minimised (for example in ceiling fans), a direct drive motor is therefore preferred, despite its larger size.

In this section, we began by exploring and quantifying the mechanism of torque production, so not surprisingly it was tacitly assumed that the rotor was at rest, with no work being done. We then moved on to assume that the torque was maintained when the speed was constant and useful power was delivered, i.e. that electrical energy was being converted into mechanical energy. The aim was to establish what factors determine the output of a rotor of given dimensions, and this was possible without reference to any particular type of motor.

In complete contrast, the approach in the next section focuses on a generic 'primitive' motor, and we begin to look in detail at what we have to do at the terminals in order to control the speed and torque.

1.6 Energy conversion—Motional e.m.f

We now examine the behaviour of a primitive linear machine which, despite its obvious simplicity, encapsulates all the key electromagnetic energy conversion processes that take place in electric motors. We will see how the process of conversion of energy from electrical to mechanical form is elegantly represented in an ‘equivalent circuit’ from which all the key aspects of motor behaviour can be predicted. This circuit will provide answers to such questions as ‘how does the motor automatically draw in more power when it is required to work’, and ‘what determines the steady speed and current’. Central to such questions is the matter of motional e.m.f., which is explored next.

We have already seen that force (and hence torque) is produced on current-carrying conductors exposed to a magnetic field. The force is given by Eq. (1.2), which shows that as long as the flux density and current remain constant, the force will be constant. In particular we see that the force does not depend on whether the conductor is stationary or moving. On the other hand relative movement is an essential requirement in the production of mechanical output power (as distinct from torque), and we have seen that output power is given by the equation $P = T\omega$. We will now see that the presence of relative motion between the conductors and the field always brings ‘motional e.m.f.’ into play; and we will find that this motional e.m.f. plays a key role in quantifying the energy conversion process.

1.6.1 Elementary motor—Stationary conditions

The primitive linear machine is shown pictorially in Fig. 1.14 and in diagrammatic form in Fig. 1.15. It consists of a conductor of active[1] length l which can move horizontally perpendicular to a magnetic flux density B.

It is assumed that the conductor has a resistance (R), that it carries a d.c. current (I), and that it moves with a velocity (v) in a direction perpendicular to the field and the current (see Fig. 1.15). Attached to the conductor is a string which passes over a pulley and supports a weight: the tension in the string acting as a mechanical ‘load’ on the rod. Friction is assumed to be zero.

We need not worry about the many difficult practicalities of making such a machine, not least how we might manage to maintain electrical connections to a moving conductor. The important point is that although this is a hypothetical

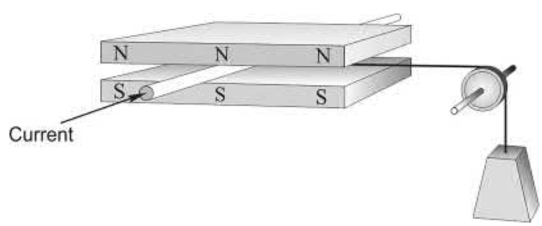

FIG. 1.14 Primitive linear d.c. motor.

1. The active length is that part of the conductor exposed to the magnetic flux density—in most motors this corresponds to the length of the rotor and stator iron cores.

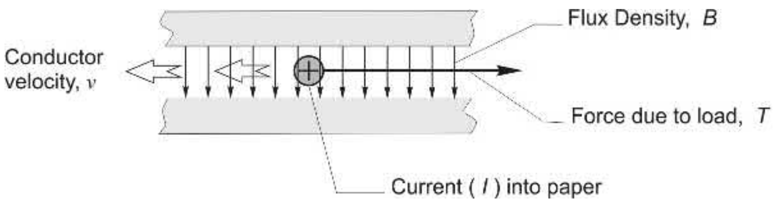

FIG. 1.15 Diagrammatic sketch of primitive linear d.c. motor.

set-up, it represents what happens in a real motor, and it allows us to gain a clear understanding of how real machines behave before we come to grips with more complex structures.

We begin by considering the electrical input power with the conductor stationary (i.e. $v = 0$). For the purpose of this discussion we can suppose that the magnetic field (B) is provided by permanent magnets. Once the field has been established (when the magnet was first magnetised and placed in position), no further energy will be needed to sustain the field, which is just as well since it is obvious that an inert magnet is incapable of continuously supplying energy. It follows that when we obtain mechanical output from this primitive 'motor', none of the energy involved comes from the magnet. This is an extremely important point: the field system, whether provided from permanent magnets or 'exciting' windings, acts only as a catalyst in the energy conversion process, and contributes nothing to the mechanical output power.

When the conductor is held stationary the force produced on it (BIl) does no work, so there is no mechanical output power, and the only electrical input power required is that needed to drive the current through the conductor.

The resistance of the conductor is R, the current through it is I, so the voltage which must be applied to the ends of the rod from an external source will be given by $V_1 = IR$ and the electrical input power will be $V_1 I$ or $I^2 R$. Under these conditions, all the electrical input power will appear as heat inside the conductor, and the power balance can be expressed by the equation

$$\text{Electrical input power } (V_1 I) = \text{Rate of production of heat in conductor } \left(I^2 R\right) \quad (1.15)$$

Although no work is being done because there is no movement, the stationary condition can only be sustained if there is equilibrium of forces. The tension in the string (T) must equal the gravitational force on the mass (mg), and this in turn must be balanced by the electromagnetic force on the conductor (BIl). Hence under stationary conditions the current must be given by

$$T = mg = BIl, \text{ or } I = \frac{mg}{Bl} \quad (1.16)$$

This is our first indication of the essential link that always exists (in the steady state) between the mechanical and electric worlds, because we see that in order to maintain the stationary condition, the current in the conductor is determined by the mass of the mechanical load. We will return to this interdependence later.

1.6.2 Power relationships—Conductor moving at constant speed

Now let us imagine the situation where the conductor is moving at a constant velocity (v) in the direction of the electromagnetic driving force that is propelling it. What current must there be in the conductor, and what voltage will have to be applied across its ends?

We start by recognising that constant velocity of the conductor means that the mass (m) is moving upwards at a constant speed, i.e. it is not accelerating. Hence from Newton's law, there must be no resultant force acting on the mass, so the tension in the string (T) must equal the weight (mg).

Similarly, the conductor is not accelerating, so its net force must also be zero. The string is exerting a braking force (T), so the electromagnetic force (BIl) must be equal to T. Combining these conditions yields

$$T = mg = BIl, \text{ or } I = \frac{mg}{Bl} \tag{1.17}$$

This is exactly the same equation that we obtained under stationary conditions, and it underlines the fact that the steady-state current is determined by the mechanical load. When we develop the electrical equivalent circuit, we will have to get used to the idea that in the steady-state one of the electrical variables (the current) is determined by the mechanical load.

With the mass rising at a constant rate, mechanical work is being done because the potential energy of the mass is increasing. This work is coming from the moving conductor. The mechanical output power is equal to the rate of work, i.e. the force ($T = BIl$) times the velocity (v). The power lost as heat in the conductor is the same as it was when stationary, since it has the same resistance, and the same current. The electrical input power supplied to the conductor must continue to furnish this heat loss, but in addition it must now supply the mechanical output power. As yet we do not know what voltage will have to be applied, so we will denote it by V_2. The power balance equation now becomes

$$\text{Electrical input power} = \begin{array}{l}\text{Rate of production of heat in conductor} + \\ \text{Mechanical output power}\end{array}$$

$$\text{i.e. } V_2 I = I^2 R + (BIl)v \tag{1.18}$$

We note that the first term on the right hand side of Eq. (1.18) represents the heating effect, which is the same as when the conductor was stationary, while the second term corresponds to the additional power that must be supplied to provide the mechanical output. Since the current is the same but the input power is now greater, the new voltage V_2 must be higher than V_1.

By subtracting Eq. (1.15) from Eq. (1.18) we obtain

$$V_2 I - V_1 I = (BIl)v,$$

and thus

$$V_2 - V_1 = Blv = E \tag{1.19}$$

Eq. (1.19) quantifies the extra voltage to be provided by the source to keep the current constant when the conductor is moving. This increase in source voltage is a reflection of the fact that whenever a conductor moves through a magnetic field, an electromotive force or voltage (E) is induced in it.

We see from Eq. (1.19) that the e.m.f. is directly proportional to the flux density, to the velocity of the conductor relative to the flux, and to the active length of the conductor. The source voltage has to overcome this additional voltage in order to keep the same current flowing: if the source voltage was not increased, the current would fall as soon as the conductor began to move because of the opposing effect of the induced e.m.f.

We have deduced that there must be an e.m.f. caused by the motion, and have derived an expression for it by using the principle of the conservation of energy, but the result we have obtained, i.e.

$$E = Blv \tag{1.20}$$

is often introduced as the 'flux-cutting' form of Faraday's law, which states that when a conductor moves through a magnetic field an e.m.f., given by Eq. (1.20), is induced in it. Because motion is an essential part of this mechanism, the e.m.f. induced is often referred to as a 'motional e.m.f.'. The 'flux-cutting' terminology arises from attributing the origin of the e.m.f. to the cutting or slicing of the lines of flux by the passage of the conductor. This is a useful mental picture, though it must not be pushed too far: after all the flux lines are merely inventions which we find helpful in coming to grips with magnetic matters.

Before turning to the equivalent circuit of the primitive motor two general points are worth noting. Firstly, whenever energy is being converted from electrical to mechanical form, as here, the induced e.m.f. always acts in opposition to the applied (source) voltage. This is reflected in the use of the term 'back e.m.f.' to describe motional e.m.f. in motors. Secondly, although we have discussed a particular situation in which the conductor carries current, it is certainly not necessary for any current to be flowing in order to produce an e.m.f.: all that is needed is relative motion between the conductor and the magnetic field.

1.7 Equivalent circuit

We can represent the electrical relationships in the primitive machine in an equivalent circuit as shown in Fig. 1.16.

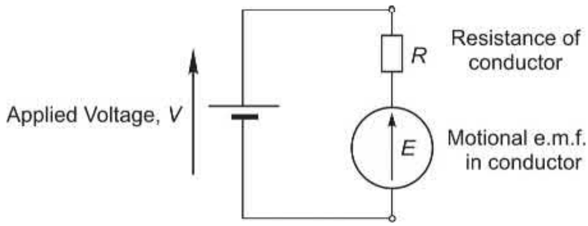

FIG. 1.16 Equivalent circuit of primitive d.c. motor.

The resistance of the conductor and the motional e.m.f. together represent in circuit terms what is happening in the conductor (though in reality the e.m.f. and the resistance are distributed, not lumped as separate items). The externally applied source that drives the current is represented by the voltage V on the left (the old-fashioned battery symbol being deliberately used to differentiate the applied voltage V from the induced e.m.f. E). We note that the induced motional e.m.f. is shown as opposing the applied voltage, which applies in the 'motoring' condition we have been discussing. Applying Kirchoff's law we obtain the voltage equation as

$$V = E + IR \text{ or } I = \frac{V - E}{R} \tag{1.21}$$

Multiplying Eq. (1.21) by the current gives the power equation as

$$\text{Electrical input power } (VI) = \begin{array}{l}\text{Mechanical output power } (EI) + \\ \text{Copper loss } \left(I^2R\right)\end{array} \tag{1.22}$$

(Note that the term 'copper loss' used in Eq. (1.22) refers to the heat generated by the current in the windings: all such losses in electric motors are referred to in this way, even when the conductors are made of aluminium or bronze!)

It is worth seeing what can be learned from these equations because, as noted earlier, this simple elementary 'motor' encapsulates all the essential features of real motors. Lessons which emerge at this stage will be invaluable later, when we look at the way actual motors behave.

1.7.1 Motoring and generating

If the e.m.f. E is less than the applied voltage V, the current will be positive, and electrical power will flow from the source, resulting in motoring action in which energy is converted from electrical to mechanical form. The first term on the right hand side of Eq. (1.22), which is the product of the motional e.m.f. and the current, represents the mechanical output power developed by the primitive linear motor, but the same simple and elegant result applies to real motors. We may sometimes have to be a bit careful if the e.m.f. and the current are not simple d.c. quantities, but the basic idea will always hold good.

Now let us imagine that we push the conductor along at a steady speed that makes the motional e.m.f. greater than the applied voltage. We can see from the equivalent circuit that the current will now be negative (i.e. anticlockwise), flowing back into the supply and thus returning energy to the supply. And if we look at Eq. (1.22), we see that with a negative current, the first term ($-VI$) represents the power being returned to the source, the second term ($-EI$) corresponds to the mechanical power being supplied by us pushing the rod along, and the third term is the heat loss in the conductor.

For readers who prefer to argue from the mechanical standpoint, rather than the equivalent circuit, we can say that when we are generating a negative current ($-I$), the electromagnetic force on the conductor is ($-BIl$), i.e. it is directed in the opposite direction to the motion. The mechanical power is given by the product of force and velocity, i.e. ($-BIlv$), or $-EI$, as above.

The fact that exactly the same kit has the inherent ability to switch from motoring to generating without any interference by the user is an extremely desirable property of all electromagnetic energy converters. Our primitive set-up is simply a machine which is equally at home acting as a motor or a generator.

Finally, it is obvious that in a motor we want as much as possible of the electrical input power to be converted to mechanical output power, and as little as possible to be converted to heat in the conductor. Since the output power is EI, and the heat loss is I^2R, we see that ideally we want EI to be much greater than I^2R, or in other words E should be much greater than IR. In the equivalent circuit (Fig. 1.16) this means that the majority of the applied voltage V is accounted for by the motional e.m.f. (E), and only a little of the applied voltage is used in overcoming the resistance.

1.8 Constant voltage operation

Up to now, we have studied behaviour under 'steady-state' conditions, which in the context of motors means that the load is constant and conditions have settled to a steady speed. We saw that with a constant load, the current was the same at all steady speeds, the voltage being increased with speed to take account of the rising motional e.m.f. This was a helpful approach to take in order to illuminate the energy conversion process, but is seldom typical of normal operation. We therefore turn to how the moving conductor will behave under conditions where the applied voltage V is constant, since this corresponds more closely with normal operation of a real motor.

Matters inevitably become more complicated because we consider how the motor gets from one speed to another, as well as what happens under steady-state conditions. As in all areas of dynamics, study of the transient behaviour of our primitive linear motor brings into play additional parameters such as the mass of the conductor (equivalent to the inertia of a rotary motor) which are absent from steady-state considerations.

1.8.1 Behaviour with no mechanical load

In this section we assume that the hanging weight has been removed, and that the only force on the conductor is its own electromagnetically generated one. Our primary interest will be in what determines the steady speed of the primitive motor, but we begin by considering what happens when we first apply the voltage.

With the conductor stationary when the voltage V is applied, there is no motional e.m.f. and the current will immediately rise to a value of V/R, since the only thing which limits the current is the resistance. (Strictly we should allow for the effect of inductance in delaying the rise of current, but we choose to ignore it here in the interests of simplicity.) The resistance will be small, so the current will be large, and a high 'BIl' force will therefore be developed on the conductor. The conductor will therefore accelerate at a rate governed by Newton's law, i.e. acceleration = the force acting on it divided by its mass.

As the speed (v) increases, the motional e.m.f. (Eq. 1.20) will grow in proportion to the speed. Since the motional e.m.f. opposes the applied voltage, the current will fall (Eq. 1.21), so the force and hence the acceleration will reduce, though the speed will continue to rise. The speed will increase as long as there is an accelerating force, i.e. as long as there is a current in the conductor. We can see from Eq. (1.21) that the current will finally fall to zero when the speed reaches a level at which the motional e.m.f. is equal to the applied voltage. The speed and current therefore vary as shown in Fig. 1.17, both curves having the exponential shape which characterises the response of systems governed by a first-order differential equation. The fact that the steady-state current is zero is in line with our earlier observation that the mechanical load (in this case zero) determines the steady-state current.

We note that in this idealised situation (in which there is no load applied, and no friction forces), the conductor will continue to travel at a constant speed, because with no net force acting on it there is no acceleration. Of course, no mechanical power is being produced, since we have assumed that there is no opposing force on the conductor, and there is no input power because the current

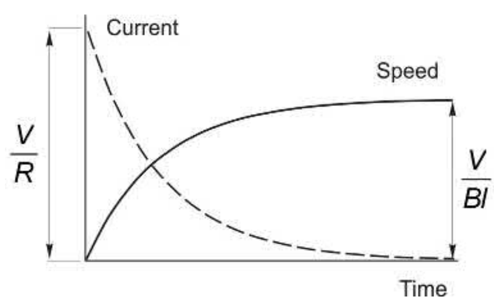

FIG. 1.17 Dynamic (run-up) behaviour of primitive d.c. motor with no mechanical load.

is zero. This hypothetical situation nevertheless corresponds closely to the so-called 'no-load' condition in a motor, the only difference being that a motor will have some friction (and therefore it will draw a small current), whereas we have assumed no friction in order to simplify the discussion.

An elegant self-regulating mechanism is evidently at work here. When the conductor is stationary, it has a high force acting on it, but this force tapers-off as the speed rises to its target value, which corresponds to the back e.m.f. being equal to the applied voltage. Looking back at the expression for motional e.m.f. (Eq. 1.18), we can obtain an expression for the no-load speed v_0 by equating the applied voltage and the back e.m.f., which gives

$$E = V = Blv_0,\ i.e.\ v_0 = \frac{V}{Bl} \tag{1.23}$$

Eq. (1.23) shows that the steady-state no-load speed is directly proportional to the applied voltage, which indicates that speed control can be achieved by means of the applied voltage. We will see later that one of the main reasons why d.c. motors held sway in the speed-control arena for so long is that their speed can be controlled via the applied voltage.

Rather more surprisingly, Eq. (1.23) reveals that the speed is inversely proportional to the magnetic flux density, which means that the weaker the field, the higher the steady-state speed. This result can cause raised eyebrows, and with good reason. Surely, it is argued, since the force is produced by the action of the field, the conductor will not go as fast if the field is weaker. This view is wrong, but understandable.

The flaw in the argument is to equate force with speed. When the voltage is first applied, the force on the conductor certainly will be less if the field is weaker, and the initial acceleration will be lower. But in both cases the acceleration will continue until the current has fallen to zero, and this will only happen when the induced e.m.f. has risen to equal the applied voltage. With a weaker field, the speed needed to generate this e.m.f. will be higher than with a strong field: there is 'less flux', so what there is has to be cut at a higher speed to generate a given e.m.f. The matter is summarised in Fig. 1.18, which shows how the speed will rise for a given applied voltage, for 'full' and 'half' fields

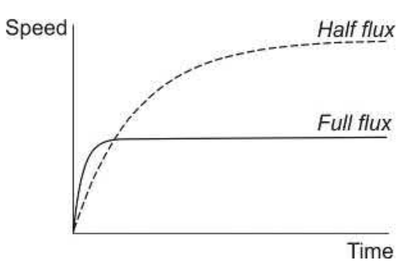

FIG. 1.18 Effect of flux density on the acceleration and steady running speed of primitive d.c. motor with no mechanical load.

respectively. Note that the initial acceleration (i.e. the slope of the speed-time curve) in the half-flux case is half that of the full flux case, but the final steady speed is twice as high. In motors the technique of reducing the flux density in order to increase speed is known as 'field weakening'.

1.8.2 Behaviour with a mechanical load

Suppose that, with the primitive linear motor up to its no-load speed we suddenly attach the string carrying the weight, so that we now have a steady force T (= mg), opposing the motion of the conductor. At this stage there is no current in the conductor and thus the only force on it will be T. The conductor will therefore begin to decelerate. But as soon as the speed falls, the back e.m.f. will become less than V, and current will begin to flow into the conductor, producing an electromagnetic driving force. The more the speed drops, the bigger the current, and hence the larger the force developed by the conductor. When the force developed by the conductor becomes equal to the load (T), the deceleration will cease, and a new equilibrium condition will be reached. The speed will be lower than at no-load, and the conductor will now be producing continuous mechanical output power, i.e. acting as a motor.

We recall that the electromagnetic force on the conductor is directly proportional to the current, so it follows that the steady-state current is directly proportional to the load which is applied, as we saw earlier. If we were to explore the transient behaviour mathematically, we would find that the drop in speed followed the same first-order exponential response that we saw in the run-up period. Once again the self-regulating property is evident, in that when load is applied the speed drops just enough to allow sufficient current to flow to produce the force required to balance the load. We could hardly wish for anything better in terms of performance, yet the conductor does it without any external intervention on our part.

(Readers who are familiar with closed-loop control systems will probably recognise that the reason for this excellent performance is that the primitive motor possesses inherent negative speed feedback via the motional e.m.f.).

Returning to Eq. (1.21), we note that the current depends directly on the difference between V and E, and inversely on the resistance. Hence for a given resistance, the larger the load (and hence the steady-state current), the greater the required difference between V and E, and hence the lower the steady running speed, as shown in Fig. 1.19.

We can also see from Eq. (1.21) that the higher the resistance of the conductor, the more it slows down when a given load is applied. Conversely, the lower the resistance, the more the conductor is able to hold its no-load speed in the face of applied load, as also shown in Fig. 1.19. We can deduce that the only way we could obtain an absolutely constant speed with this type of motor is for the resistance of the conductor to be zero, which is of course not possible. Nevertheless, real d.c. motors generally have resistances which are small, and

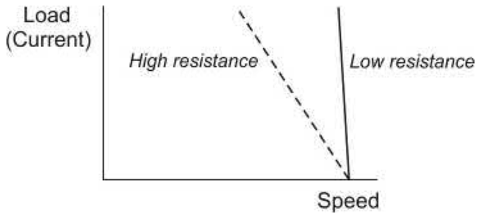

FIG. 1.19 Influence of resistance on the ability of the motor to maintain speed when load is applied.

their speed does not fall much when load is applied—a characteristic which is highly desirable for most applications.

We complete our exploration of the performance when loaded by asking how the flux density influences behaviour. Recalling that the electromagnetic force is proportional to the flux density as well as the current, we can deduce that to develop a given force, the current required will be higher with a weak flux than with a strong one. Hence in view of the fact that there will always be an upper limit to the current which the conductor can safely carry, the maximum force which can be developed will vary in direct proportion to the flux density, with a weak flux leading to a low maximum force and vice-versa. This underlines the importance of operating with maximum flux density whenever possible.

We can also see another disadvantage of having a low flux density by noting that to achieve a given force, the drop in speed will be disproportionately high when we go to a lower flux density. We can see this by imagining that we want a particular force, and considering how we achieve it firstly with full flux, and secondly with half flux. With full flux, there will be a certain drop in speed which causes the motional e.m.f. to fall enough to admit the required current. But with half the flux, for example, twice as much current will be needed to develop the same force. Hence the motional e.m.f. must fall by twice as much as it did with full flux. However, since the flux density is now only half, the drop in speed will have to be four times as great as it was with full flux. The half-flux ‘motor’ therefore has a load characteristic with a load/speed gradient four times more droopy than the full-flux one. This is shown in Fig. 1.20 the applied

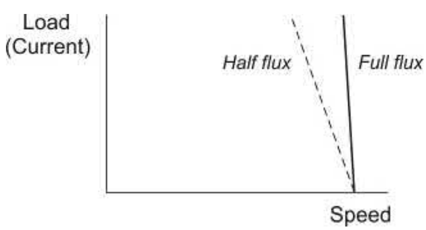

FIG. 1.20 Influence of flux on the drop in steady running speed when load is applied.

voltages having been adjusted so that in both cases the no-load speed is the same. The half-flux motor is clearly inferior in terms of its ability to hold the set speed when the load is applied.

We may have been tempted to think that the higher speed which we can obtain by reducing the flux somehow makes for better performance, but we can now see that this is not so. By halving the flux, for example, the no load speed for a given voltage is doubled, but when the load is raised until rated current is flowing in the conductor, the force developed is only half, so the mechanical power is the same. We are in effect trading speed against force, and there is no suggestion of getting something for nothing.

1.8.3 Relative magnitudes of *V* and *E*, and efficiency

Invariably we want machines which have high efficiency. From Eq. (1.20), we see that to achieve high efficiency, the copper loss (I^2R) must be small compared with the mechanical power (EI), which means that the resistive volt-drop in the conductor (IR) must be small compared with either the induced e.m.f. (E) or the applied voltage (V). In other words we want most of the applied voltage to be accounted for by the 'useful' motional e.m.f., rather than the wasteful volt drop in the wire. Since the motional e.m.f. is proportional to speed, and the resistive volt drop depends on the conductor resistance, we see that a good energy converter requires the conductor resistance to be as low as possible, and the speed to be as high as possible.

To provide a feel for the sorts of numbers likely to be encountered, we can consider a conductor with resistance of 0.5 Ω, capable of carrying a current of 4 A without overheating, and moving at a speed such that the motional e.m.f. is 8 V. From Eq. (1.19), the supply voltage is given by

$$V = E + IR = 8 + (4 \times 0.5) = 10\,\text{Volts}$$

Hence the electrical input power (VI) is 40 W, the mechanical output power (EI) is 32 W, and the copper loss (I^2R) is 8 W, giving an efficiency of 80%.

If the supply voltage was doubled (i.e. $V = 20$ Volts), however, and the resisting force is assumed to remain the same (so that the steady-state current is still 4 A), the motional e.m.f. is given by Eq. (1.21) as

$$E = 20 - (4 \times 0.5) = 18\,\text{Volts}$$

which shows that the speed will have more than doubled, as expected. The electrical input power is now 80 W, the mechanical output power is 72 W, and the copper loss is still 8 W. The efficiency has now risen to 90%, underlining the fact that the energy conversion process gets better at higher speeds.

When we operate the machine as a generator, we again benefit from the higher speeds. For example, with the battery voltage maintained at 10 V, and the conductor being propelled by an external force so that its e.m.f. was 12 V, the allowable current of 4 A would now be flowing *into* the battery, with energy being converted from mechanical to electrical form. The power into the battery (VI) is 40 W, the mechanical input power (EI) is 48 W and the heat loss 8 W. In this case efficiency is defined as the ratio of useful electrical power divided by mechanical input power, i.e. 40/48, or 83.3%.

If we double the battery voltage to 20 V and increase the driven speed so that the motional e.m.f. rises to 22 V, we will again supply the battery with 4 A, but the efficiency will now be 80/88, or 90.9%.

The ideal situation is clearly one where the term IR in Eq. (1.22) is negligible, so that the back e.m.f. is equal to the applied voltage. We would then have an ideal machine with an efficiency of 100%, in which the steady-state speed would be directly proportional to the applied voltage and independent of the load.

In practice the extent to which we can approach the ideal situation discussed above depends on the size of the machine. Tiny motors, such as those used in wrist watches, are awful, in that most of the applied voltage is used up in overcoming the resistance of the conductors, and the motional e.m.f. is very small: these motors are much better at producing heat than they are at producing mechanical output power! Small machines, such as those used in hand tools, are a good deal better with the motional e.m.f. accounting for perhaps 70–80% of the applied voltage. Industrial machines are very much better: the largest ones (of many hundreds of kW) use only 1 or 2% of the applied voltage in overcoming resistance, and therefore have very high efficiencies.

1.8.4 Analysis of primitive machine—Conclusions

All of the lessons learned from looking at the primitive machine will find direct parallels in almost all of the motors we look at in the rest of this book, so it is worth reminding ourselves of the key points.

Although this book is primarily about motors, perhaps the most important conclusion so far is that electrical machines are inherently bi-directional energy converters, and any motor can be made to generate, or vice-versa. We also saw that the efficiency of the energy-conversion process improves at high speeds, which explains why direct-drive slow-speed motors are not widely used.

In terms of the theoretical underpinning, we will make frequent reference to the formula for the force on a conductor in a magnetic field, i.e.

$$\text{Force}, F = BIl \tag{1.24}$$

and to the formula for the motional induced e.m.f, i.e.

$$\text{Motional e.m.f.}, E = Blv \tag{1.25}$$

where B is the magnetic flux density, I is the current, l is the length of conductor and v is the velocity perpendicular to the field.

Specifically in relation to d.c. machines, we have seen that the *speed* at which the primitive motor runs unloaded is determined by the *applied voltage*, while the steady-state *current* that the motor draws is determined by the *mechanical load*. Exactly the same results will hold when we examine real d.c. motors, and very similar relationships will also emerge when we look at a.c. machines.

1.9 General properties of electric motors

All electric motors are governed by the laws of electromagnetism, and are subject to essentially the same constraints imposed by the materials (copper or aluminium for the electric circuit, magnetic steel, and insulation) from which they are made. We should therefore not be surprised to find that, at the fundamental level, all motors—regardless of type—have a great deal in common.

These common properties, most of which have been touched on in this chapter, are not usually given prominence. Books tend to concentrate on the differences between types of motor, and manufacturers are usually interested in promoting the virtues of their particular motor at the expense of the competition. This divisive emphasis can cause the underlying unity to be obscured, leaving users with little opportunity to absorb the sort of knowledge which will equip them to make informed judgements.

The most useful ideas worth bearing in mind are therefore given below, with brief notes accompanying each. Experience indicates that users who have these basic ideas firmly in mind will find themselves better able to understand why one motor is better than another, and will feel more confident when faced with the difficult task of weighing the pros and cons of competing types.

1.9.1 Operating temperature and cooling

The cooling arrangement is the single most important factor in determining the permissible output from any given motor.

Any motor will give out more power if its electric circuit is worked harder (i.e. if the current is allowed to increase). The limiting factor is normally the allowable temperature rise of the windings, which depends on the class of insulation.

For class F insulation (the most widely used) the permissible temperature rise is 100 K, whereas for class H it is 125 K. Thus if the cooling remains the same, more output can be obtained simply by using the higher-grade insulation. Alternatively, with a given insulation the output can be increased if the cooling system is improved. A through-ventilated motor, for example, might give perhaps twice the output power of an otherwise identical but totally enclosed machine.

1.9.2 Torque per unit volume

For motors with similar cooling systems, the rated torque is approximately proportional to the rotor volume, which in turn is roughly proportional to the overall motor volume.

This stems from the fact that for a given cooling arrangement, the electric and magnetic loadings of machines of different types will be more or less the same. The torque per unit length therefore depends first and foremost on the square of the diameter, so motors of roughly the same diameter and length can be expected to produce roughly the same torque.

1.9.3 Power per unit volume and efficiency—Importance of speed

Output power per unit volume is directly proportional to speed.

Low-speed motors are unattractive for most applications because they are large, and therefore expensive. It is usually better to use a high-speed motor with mechanical speed reduction. For example, a direct drive motor for a portable electric screwdriver would be an absurd proposition. On the other hand the reliability and inefficiency of gearboxes may sometimes outweigh the size argument, especially in high-power applications.

The efficiency of a motor improves with speed.

For a given torque, power output usually rises in direct proportion to speed, while electrical losses tend to rise less rapidly, so that efficiency rises with speed.

1.9.4 Size effects—Specific torque and efficiency

Large motors have a higher specific torque (torque per unit volume) and are more efficient than small ones.

In large motors the specific electric loading is normally much higher than in small ones, and the specific magnetic loading is somewhat higher. These two factors combine to give the higher specific torque.

Very small motors are inherently very inefficient (e.g. 1% in a wrist-watch), whereas motors of over say 100 kW have efficiencies above 96%. The reasons for this scale effect are complex, but stem from the fact that the resistance volt-drop term can be made relatively small in large electromagnetic devices, whereas in small ones the resistance becomes the dominant term.

1.9.5 Rated voltage

A motor can be provided to suit any voltage.

Within limits it is possible to rewind a motor for a different voltage without affecting its performance. A 200 V, 10 A motor could be rewound for 100 V, 20 A simply by using half as many turns per coil of wire having twice the cross-sectional area, the m.m.f. remaining the same. The total amounts of active material, and hence the performance, would be the same. This argument breaks down if pushed too far of course: a very small motor originally wound for 100 V would almost certainly require a larger frame if required to operate at 690 V, because of the additional space required for the higher-grade insulation.

1.9.6 Short-term overload

Most motors can be overloaded for short periods without damage.

The continuous electric loading (i.e. the current) cannot be exceeded without overheating and damaging the insulation, but if the motor has been running with reduced current for some time, it is permissible for the current (and hence the torque) to be much greater than normal for a short period of time. The principal factors which influence the magnitude and duration of the permissible overload are the thermal time-constant (which governs the rate of rise of temperature) and the previous pattern of operation. Thermal time constants range from a few seconds for small motors to many minutes or even hours for large ones. Operating patterns are obviously very variable, so rather than rely on a particular pattern being followed, it is usual for motors to be provided with over-temperature protective devices (e.g. thermistors) which trigger an alarm and/or trip the supply if the safe temperature is exceeded.

Motors can of course be designed for specific applications with known duty cycles, where the thermal and electric loading can be accommodated. This is discussed further in Chapter 11.

1.10 Review questions

(1) The current in a coil with 250 turns is 8 A. Calculate the m.m.f.

(2) The coil in Q1 is used in a magnetic circuit with a uniform cross-section made of good-quality magnetic steel and with a 2 mm air-gap. Estimate the flux density in the air-gap, and in the iron. ($\mu_0 = 4 \times 10^{-7}$ Hm^{-1})

How would the answers change if the cross-sectional area of the magnetic circuit was doubled, with all other parameters remaining the same?

(3) A magnetic circuit of uniform cross-sectional area has two air-gaps of 0.5 and 1 mm respectively in series. The exciting winding provides an m.m.f. of 1200 Amp-turns. Estimate the m.m.f. across each of the air-gaps, and the flux density.

(4) The rotor of a d.c. motor had an original diameter of 30 cm and an air-gap under the poles of 2 mm. During refurbishment the rotor diameter was accidentally reground and was then undersize by 0.5 mm. Estimate by how much the field m.m.f. would have to be increased to restore normal performance. How might the extra m.m.f. be provided?

(5) Calculate the electromagnetic force on:

- **(a)** a single conductor of length 25 cm, carrying a current of 4A, exposed to a magnetic flux density of 0.8 T perpendicular to its length.
- **(b)** a coil-side consisting of 20 wires of length 25 cm, each carrying a current of 2 A, exposed to a magnetic flux density of 0.8 T perpendicular to its length.

(6) Estimate the torque produced by one of the early machines illustrated in Fig. 1.12 given the following:- Mean air-gap flux density under pole-face = 0.4 T; pole-arc as a percentage of total circumference = 75%; active length of rotor = 50 cm; rotor diameter = 30 cm; inner diameter of stator pole = 32 cm; total number of rotor conductors = 120; current in each rotor conductor = 50 A.

(7) If the field coils of a motor are rewound to operate from 220 V instead of 110 V, how will the new winding compare with the old in terms of number of turns, wire diameter, power consumption and physical size?

(8) A catalogue of DIY power tools indicates that most of them are available in 240 V or 110 V versions. What differences would you expect in terms of appearance, size, weight and performance?

(9) Given that the field windings of a motor do not contribute to the mechanical output power, why do they consume power continuously?

(10) Explain briefly why low-speed electrical drives often employ a high-speed motor and some form of mechanical speed reduction, rather than a direct-drive motor.

Answers to the review questions are given in the *Appendix*.

Chapter 2

Power electronic converters for motor drives

2.1 Introduction

In this chapter we look at the power converter circuits which are widely used with motor drives, providing either d.c. or a.c. outputs, and working from either a d.c. (e.g. battery) supply, or from the conventional (50 or 60 Hz) utility supply. All of the principal converter types are introduced, but the coverage is not intended to be exhaustive, with attention being concentrated on the most important topologies, their features and aspects of behaviour which recur in many types of drive converter.

All except very low power converters are today based on power electronic switching. The need to adopt a switching strategy is emphasised in the first example, where the consequences are explored in some depth. We will see that switching is essential in order to achieve high-efficiency power conversion, but that the resulting waveforms are inevitably less than ideal from the point of view of both the motor and the power supply.

The examples have been chosen to illustrate typical practice, so the most commonly used switching devices (e.g. thyristor, MOSFET, and IGBT) are shown. In many cases, several different switching devices may be suitable (see later), so we should not identify a particular circuit as being the exclusive preserve of a particular device.

Power semiconductor devices are the subject of continuous innovation, with faster switching, lower on-state resistance, higher operating temperatures, higher blocking voltage capability, improved robustness to fault conditions and of course cost/competitive advantage being the main drivers. These developments rarely impact on either the topology of power converters or their principal characteristics, and so detailed review of power semiconductor switches is beyond the scope of this book.[1]

1. Readers who want to learn more about different power semiconductor devices should have a look at "Power Semiconductor Devices" by Lutz, Schlangenotto, Scheuermann, and De Donker, Springer, 2018, ISBN-10: 331970916x, ISBN-13: 978-3319709161.

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00002-7

This chapter covers all the most significant types of power converter circuits used in d.c. and a.c. drives. There are a small number of further topologies which are specific to the control of a particular motor type, and these will be discussed when we come to look at those motors.

Before discussing particular circuits it will be useful to take an overall look at a typical drive system, so that the role of the converter can be seen in its proper context.

2.1.1 General arrangement of drive

A complete drive system is shown in block diagram form in Fig. 2.1.

The job of the power converter is to draw electrical energy from the utility supply (at constant voltage and frequency) and supply electrical energy to the motor at whatever voltage and frequency is necessary to achieve the desired mechanical output. In Fig. 2.1, the 'demanded' output is the speed of the motor, but equally it could be the torque, the position of the motor shaft, or some other system variable.

Except in the very simplest converter (such as a basic diode rectifier), there are usually two distinct parts to the converter. The first is the power stage, through which the energy flows to the motor, and the second is the control section, which regulates the power flow. Low-power control signals tell the converter what it is supposed to be doing, while other low-power feedback signals are used to measure what is actually happening. By comparing the demand and feedback signals, and adjusting the output accordingly, the target output is maintained.

The basic arrangement shown in Fig. 2.1 is clearly a speed control system, because the signal representing the demand or reference quantity is speed, and we note that it is a 'closed-loop' system because the quantity that is to be controlled is measured and fed back to the controller so that action can be taken if they do not correspond. All drives employ some form of closed-loop (feedback) control, so readers who are unfamiliar with the basic principles will probably be well advised to do some background reading.[2]

In later chapters we will explore the internal workings and control arrangements at greater length, but it is worth mentioning that almost all drives employ current feedback, not only to facilitate over-current protection circuits, but in order to control the motor torque, and that in all except high-performance drives it is unusual to find external transducers, which can account for a significant fraction of the total cost of the drive system. Instead of measuring the actual speed using, for example, a shaft-mounted tachogenerator as shown in Fig. 2.1, speed is more likely to be derived from sampled measurements of

2. Readers who are not familiar with feedback and control systems would find this book helpful: Schaum's Outline of Feedback and Control Systems, 2nd Edition (2013) by Joseph Distefano. McGraw-Hill Education; ISBN-10: 9780071829489.

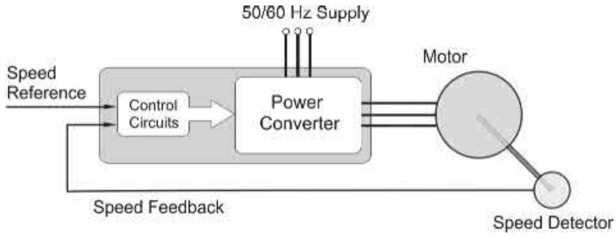

FIG. 2.1 General arrangement of speed-controlled drive.

motor voltages, currents, and frequency, used in conjunction with a stored mathematical model of the motor.

A characteristic of power electronic converters, which is shared with most electrical systems, is that they usually have very little capacity for storing energy. This means that any sudden change in the power supplied by the converter to the motor must be reflected in a sudden increase in the power drawn from the supply. In most cases this is not a serious problem, but it does have two drawbacks. Firstly, sudden increases in the current drawn from the supply will cause momentary drops in the supply voltage, because of the effect of the supply impedance. These voltage 'dips' will appear as unwelcome distortion to other users on the same supply. And secondly, there may be an enforced delay before the supply can furnish extra power. For example, with a single-phase utility supply, there can be no sudden increase in the power supply at the instant where the utility voltage is zero, because instantaneous power is necessarily zero at this point in the cycle because the voltage is itself zero.

It would be ideal if the converter could store at least enough energy to supply the motor for several cycles of the 50/60 Hz supply, so that short-term energy demands could be met instantly, thereby reducing rapid fluctuations in the power drawn from the utility supply. But unfortunately this is rarely economic: most converters do have a small store of energy in their smoothing inductors and capacitors, but the amount is not sufficient to buffer the supply sufficiently to shield it from anything more than very short-term fluctuations.

2.2 Voltage control—D.C. output from d.c. supply

In Chapter 1 we saw that control of the basic d.c. machine is achieved by controlling the current in the conductor, which is readily achieved by variation of the voltage. A controllable voltage source is therefore a key element of a motor drive, as we will see in later chapters.

However, for the sake of simplicity we will begin by exploring the problem of controlling the voltage across a 2 Ω resistive load, fed from a 12 V constant-voltage source such as a battery. Three different methods are shown in Fig. 2.2, in which the circle on the left represents an ideal 12 V d.c. source, the tip of the arrow indicating the positive terminal. Although this set-up is not quite the same

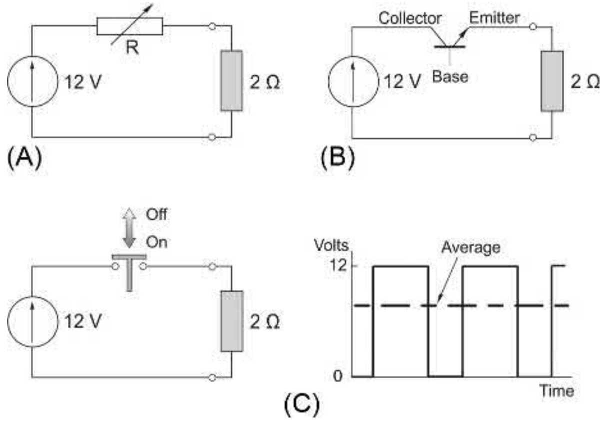

FIG. 2.2 Methods of obtaining a variable-voltage output from a constant-voltage source.

as if the load was a d.c. motor the conclusions which we draw are more or less the same.

Method (A) uses a variable resistor (R) to absorb whatever fraction of the battery voltage is not required at the load. It provides smooth (albeit manual) control over the full range from 0 to 12 V, but the snag is that power is wasted in the control resistor. For example, if the load voltage is to be reduced to 6 V, the resistor R must be set to 2 Ω, so that half of the battery voltage is dropped across R. The current will be 3 A, the load power will be 18 W, and the power dissipated in R will also be 18 W. In terms of overall power conversion efficiency (i.e. useful power delivered to the load divided by total power from the source) the efficiency is a very poor 50%. If R is increased further, the efficiency falls still lower, approaching zero as the load voltage tends to zero. Whilst this method of control was commonplace before the advent of power semiconductors in the 1960s, and some examples will still be seen in use, the cost of energy, cooling and safety considerations have led to its demise, except perhaps in applications such as low cost, low power hand tools and toy racing cars.

Method (B) is much the same as (A) except that a transistor is used instead of a manually-operated variable resistor. The transistor in Fig. 2.2B is connected with its collector and emitter terminals in series with the voltage source and the load resistor. The transistor is a variable resistor, of course, but a rather special one in which the effective collector-emitter resistance can be controlled over a wide range by means of the base-emitter current. The base-emitter current is usually very small, so it can be varied by means of a low-power electronic circuit (not shown in Fig. 2.2) whose losses are negligible in comparison with the power in the main (collector-emitter) circuit.

Method (B) shares the drawback of method (A) above, i.e. the efficiency is very low. But even more seriously, the ‘wasted’ power (up to a maximum of

18 W in this case) is burned-off inside the transistor, which therefore has to be large, well-cooled, and hence expensive. Transistors are almost never operated in this 'linear' way when used in power electronics, but are widely used as ON/OFF switches, as discussed below.

2.2.1 Switching control

The basic ideas underlying a switching power regulator are shown by the arrangement in Fig. 2.2C, which uses a mechanical switch. By operating the switch repetitively and varying the ratio of ON to OFF time, the average load voltage can be varied continuously between 0 V (switch OFF all the time) through 6 V (switch ON and OFF for half of each cycle) to 12 V (switch ON all the time).

The circuit shown in Fig. 2.2C is often referred to as a 'chopper', because the d.c. battery supply is 'chopped' ON and OFF. A constant repetition frequency (switching frequency) is normally used, and the width of the ON pulse is varied to control the mean output voltage (see the waveform in Fig. 2.2): this is known as 'pulse width modulation' (PWM).

The main advantage of the chopper circuit is that no power is wasted, and the efficiency is thus 100%. When the switch is ON, current flows through it, but the voltage across it is zero because its resistance is negligible. The power dissipated in the switch is therefore zero. Likewise, when OFF the current through it is zero, so although the voltage across the switch is 12 V, the power dissipated in it is again zero.

The obvious disadvantage is that by no stretch of the imagination could the load voltage be seen as 'good' d.c.: instead it consists of a mean or 'd.c.' level, with a superimposed 'a.c.' component. Bearing in mind that we really want the load to be a d.c. motor, rather than a resistor, we are bound to ask whether the pulsating voltage will be acceptable. Fortunately, the answer is yes, provided that the chopping frequency is high enough. We will see later that the inductance of the motor causes the current to be much smoother than the voltage, which means that the motor torque fluctuates much less than we might suppose; and the mechanical inertia of the motor filters the torque ripples so that the speed remains almost constant, at a value governed by the mean (or d.c.) level of the chopped waveform.

Obviously a mechanical switch would be unsuitable, and could not be expected to last long when pulsed at high frequency. So a power electronic switch is used instead. The first of many devices to be used for switching was the bipolar junction transistor (BJT), so we will begin by examining how such devices are employed in chopper circuits. If we choose a different device, such as a metal oxide semiconductor field effect transistor (MOSFET) or an insulated gate bipolar transistor (IGBT), where the gate drive circuits are simpler and lower power, the main conclusions we draw will be much the same.

2.2.2 Transistor chopper

As noted earlier, a transistor is effectively a controllable resistor, i.e. the resistance between collector and emitter depends on the current in the base-emitter junction. In order to mimic the operation of a mechanical switch, the transistor would have to be able to provide infinite resistance (corresponding to an open switch) or zero resistance (corresponding to a closed switch). Neither of these ideal states can be reached with a real transistor, but both can be approximated closely.

The typical relationship between the collector-emitter voltage and the collector current for a range of base currents rising from zero is shown in Fig. 2.3. The bulk of the diagram represents the so-called 'linear' region, where the transistor exhibits the remarkable property of the collector current remaining more or less constant for a wide range of collector-emitter voltages: when the transistor operates in this region there will be significant power loss. For power electronic applications, we want the device to behave like a switch, so we operate on the margins of the diagram, where either the voltage or the current is close to zero, and the heat released inside the device is therefore very low.

The transistor will be OFF when the base-emitter current (I_b) is zero. Viewed from the main (collector-emitter) circuit, its resistance will be very high, as shown by the region Oa in Fig. 2.3.

Under this 'cut-off' condition, only a tiny current (I_c) can flow from the collector to the emitter, regardless of the voltage (V_{ce}) between the collector and emitter. The power dissipated in the device will therefore be minimal, giving an excellent approximation to an open switch.

To turn the transistor fully ON, a base-emitter current must be provided. The base current required will depend on the prospective collector-emitter current, i.e. the current in the load. The aim is to keep the transistor 'saturated' so that it

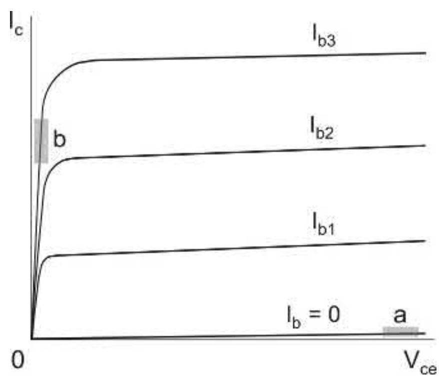

FIG. 2.3 Transistor characteristics showing high-resistance (cut-off) region Oa and low-resistance (saturation) region Ob. Typical 'off' and 'on' operating states are shown by the shaded areas a and b respectively.

has a very low resistance, corresponding to the region Ob in Fig. 2.3. In the example shown in Fig. 2.2, if the resistance of the transistor is very low, the current in the circuit will be almost 6 A, so we must make sure that the base-emitter current is sufficiently large to ensure that the transistor remains in the saturated condition when $I_c = 6$ A.

Typically in a bipolar transistor (BJT) the base current will need to be around 5–10% of the collector current to keep the transistor in the saturation region: in the example (Fig. 2.2), with the full load current of 6 A flowing, the base current might be 400 mA, the collector-emitter voltage might be say 0.33 V, giving an on-state dissipation of 2 W in the transistor when the load power is very nearly 72 W. The power conversion efficiency is not 100%, as it would be with an ideal switch, but it is acceptable.

We should note that the on-state base-emitter voltage is very low, which, coupled with the small base current, means that the power required to drive the transistor is very much less than the power being switched in the collector-emitter circuit. Nevertheless, to switch the transistor in the regular pattern shown in Fig. 2.2, we obviously need a base current which goes on and off periodically, and we might wonder how we obtain this 'control' signal. In most modern drives the signal originates from a microprocessor, many of which are designed with PWM auxiliary functions, which can be used for this purpose. Depending on the base circuit current requirements of the main switching transistor, it may be possible to feed it directly from the microprocessor, but it is more usual to see additional transistors interposed between the signal source and the main device to provide the required power amplification.

Just as we have to select mechanical switches with regard to their duty, we must be careful to use the right power transistor for the job in hand. In particular, we need to ensure that when the transistor is ON, we don't exceed the safe current, or else the active semiconductor region of the device will be destroyed by overheating. And we must make sure that the transistor is able to withstand whatever voltage appears across the collector-emitter junction when it is in the OFF condition. If the safe voltage is exceeded, even for a very short period, the transistor will breakdown, and be permanently ON.

A suitable heatsink will be a necessity in all but the smallest and most efficient drives. We have already seen that some heat is generated when the transistor is ON, and at low switching rates this is the main source of unwanted heat. But at high switching rates, 'switching loss' can also be very important.

Switching loss is the heat generated in the finite time it takes for the transistor to go from ON to OFF or vice-versa. The base-drive circuitry will be arranged so that the switching takes place as fast as possible, but in practice, with silicon-based power semiconductors, it will seldom take less than a few microseconds. During the switch-on period, for example, the current will be building up, while the collector-emitter voltage will be falling towards zero. The peak power reached can therefore be large, before falling to the relatively low on-state value. Of course the total energy released as heat each time the device switches is modest because the whole process happens so quickly. Hence if the switching rate is low (say once every second) the switching power loss will be insignificant in comparison with the on-state power. But at high switching rates, when the time taken to complete the switching becomes comparable with the ON time, the switching power loss can easily become dominant. In drives, switching rates vary somewhat depending upon the power rating. In general, the higher the power the lower the switching frequency. Many commercial drives allow the user to select the switching frequency, but by selecting high switching frequencies to give smoother current waveforms (and lower audible noise) the higher losses mean that the rating of the drive must be reduced. New, so called Wide Band Gap (WBG) power devices, notably Gallium Nitride and Silicon Carbide are now available which offer the prospect of much faster switching times, and therefore much lower switching losses, than is possible with silicon. Such devices allow efficient switching frequencies into the MHz region, and will become more prevalent in future as their costs fall.

2.2.3 Chopper with inductive load—Overvoltage protection

So far we have looked at chopper control of a resistive load, but in a drives context the load will usually mean the winding of a machine, which will invariably be inductive.

Chopper control of inductive loads is much the same as for resistive loads, but we have to be careful to prevent the appearance of dangerously high voltages each time the inductive load is switched OFF. The root of the problem lies with the energy stored in the magnetic field of the inductor. When an inductance L carries a current I, the energy stored in the magnetic field (W) is given by

$$W = \frac{1}{2} L I^2 \tag{2.1}$$

If the inductor is supplied via a mechanical switch, and we open the switch with the intention of reducing the current to zero instantaneously, we are in effect seeking to destroy the stored energy. This is not possible (because it would imply an energy pulse of zero duration and hence infinite power), and what happens instead is that the energy is dissipated in the form of a spark (or arc) across the contacts of the switch.

The appearance of a spark indicates that there is a very high voltage, sufficient to break down the surrounding air. We can anticipate this by remembering that the voltage and current in an inductance are related by the equation

$$V_L = L \frac{di}{dt} \tag{2.2}$$

which shows that the self-induced voltage is proportional to the rate of change of current, so when we open the switch in order to force the current to zero quickly, a very large voltage is created in the inductance. This voltage appears across the terminals of the switch and, if sufficient to break down the air, the resulting arc allows the current to continue to flow until the stored magnetic energy is dissipated as heat in the arc.

Sparking across a mechanical switch is unlikely to cause immediate destruction, but when a transistor is used its immediate demise is certain unless steps are taken to tame the stored energy, and in particular ensure that the voltage seen by the transistor does not exceed its rated peak voltage. The usual remedy lies in the use of a 'freewheel diode' (sometimes called a flywheel diode), as shown in Fig. 2.4.

A diode is a one-way valve as far as current is concerned: it offers very little resistance to current flowing from anode to cathode (i.e. in the direction of the broad arrow in the symbol for a diode), but blocks current flow from cathode to anode. Actually, when a power diode conducts in the forward direction, the voltage drop across it is usually not all that dependent on the current flowing through it, so the reference above to the diode 'offering little resistance' is not strictly accurate because it does not obey Ohm's law. In practice the volt-drop of power diodes (most of which are made from silicon) is around 0.7 V, regardless of the current rating.

In the circuit of Fig. 2.4A, when the transistor is ON, current (I) flows through the load, but not through the diode, which is said to be reverse-biased (i.e. the applied voltage is trying—unsuccessfully—to push current down through the diode). During this period the voltage across the inductance is positive, so the current increases, thereby increasing the stored energy.

When the transistor is turned OFF, the current through it and the battery drops very quickly to zero. But the stored energy in the inductance means that

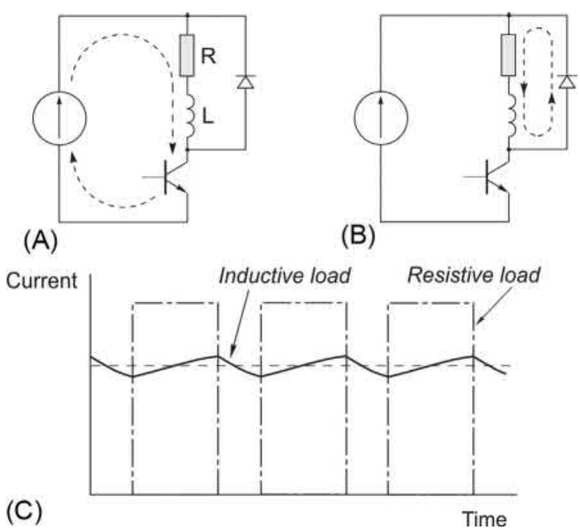

FIG. 2.4 Operation of chopper-type voltage regulator.

its current cannot suddenly disappear. So since there is no longer a path through the transistor, the current diverts into the only other route available, and flows upwards through the low-resistance path offered by the diode, as shown in Fig. 2.4B.

Obviously the current no longer has a battery to drive it, so it cannot continue to flow indefinitely. During this period the voltage across the inductance is negative, and the current reduces. If the transistor were left OFF for a long period, the current would continue to 'freewheel' only until the energy originally stored in the inductance is dissipated as heat, mainly in the load resistance but also in the diode's own (low) resistance. In normal chopping, however, the cycle restarts long before the current has fallen to zero, giving a current waveform as shown in Fig. 2.4C. Note that the current rises and falls exponentially with a time-constant of L/R, though it never reaches anywhere near its steady-state value in Fig. 2.4. The sketch corresponds to the case where the time-constant is much longer than one switching period, in which case the current becomes almost smooth, with only a small ripple. In a d.c. motor drive this is just what we want, since any fluctuation in the current gives rise to torque pulsations and consequent mechanical vibrations. (The current waveform that would be obtained with no inductance is also shown in Fig. 2.4: the mean current is the same but the rectangular current waveform is clearly much less desirable, given that ideally we would like constant d.c.).

The freewheel (or flywheel) diode was introduced to prevent dangerously high voltages from appearing across the transistor when it switches off an inductive load, so we should check that this has been achieved. When the diode conducts the forward-bias volt-drop across it is small—typically 0.7 Volts. Hence while the current is freewheeling, the voltage at the collector of the transistor is only 0.7 V above the battery voltage. This 'clamping' action therefore limits the voltage across the transistor to a safe value, and allows inductive loads to be switched without damage to the switching element.

We should acknowledge that in this example the discussion has focused on steady-state operation, when the current at the end of every cycle is the same, and it never falls to zero. We have therefore sidestepped the more complex matter of how we get from start-up to the steady state, and we have also ignored the so-called 'discontinuous current' mode, when the current in the load does fall to zero during the OFF period. We will touch on the significant consequences of discontinuous operation in drives in later chapters.

We can draw some important conclusions which are valid for all power electronic converters from this simple example. Firstly, efficient control of voltage (and hence power) is only feasible if a switching strategy is adopted. The load is alternately connected and disconnected from the supply by means of an electronic switch, and any average voltage up to the supply voltage can be obtained by varying the mark/space (ON/OFF) ratio. Secondly, the output voltage is not smooth d.c., but contains unwanted a.c. components which, though undesirable, are tolerable in motor drives. And finally, the load current waveform will be smoother than the voltage waveform if—as is the case with motor windings—the load is inductive.

2.2.4 Boost converter

The previous section dealt with the so-called step-down or buck converter, which provides an output voltage less than the input. However, if the motor voltage is higher than the supply (for example in an electric vehicle with a 240 V motor driven from a 48 V battery), a step-up or boost converter is required. Intuitively this seems a tougher challenge altogether, and we might expect to have to first convert the d.c. to a.c. so that a transformer could be used, but in fact the basic principle of transferring 'packets' of energy to a higher voltage is very simple and elegant, using the circuit shown in Fig. 2.5. Operation of this circuit is worth discussing because it again illustrates features common to many power electronic converters.

As is usual, the converter operates repetitively at a rate determined by the frequency of switching ON and OFF the transistor (T). During the ON period (Fig. 2.5A), the input voltage (V_{in}) is applied across the inductor (L), causing the current in the inductor to rise linearly, thereby increasing the energy stored in its magnetic field. Meanwhile, the motor current is supplied by the storage capacitor (C), the voltage of which falls only a little during this discharge period. In Fig. 2.5A, the input current is drawn to appear larger than the output (motor) current, for reasons that will soon become apparent. Recalling that the aim is to produce an output voltage greater than the input, it should be clear that the voltage across the diode (D) is negative (i.e. the potential is higher on the right than on the left in Fig. 2.5), so the diode does not conduct and the input and output circuits are effectively isolated from each other.

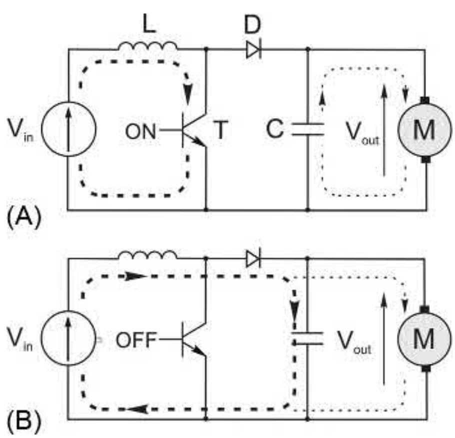

FIG. 2.5 Boost converter. The transistor is switched ON in (A) and OFF in (B).

When the transistor turns OFF, the current through it falls rapidly to zero, and the situation is much the same as in the step-down converter where we saw that because of the stored energy in the inductor, any attempt to reduce its current results in a self-induced voltage trying to keep the current going. So during the OFF period, the inductor voltage rises extremely rapidly until the potential at the left side of the diode is slightly greater than V_{out}, the diode then conducts and the inductor current flows into the parallel circuit consisting of the capacitor (C) and the motor, the latter continuing to draw a steady current (as the voltage across the large storage capacitor and hence motor cannot change instantaneously), while the major share of the inductor current goes to recharge the capacitor. The resulting voltage across the inductor is negative, and of magnitude $V_{out} - V_{in}$, so the inductor current begins to reduce and the extra energy that was stored in the inductor during the ON time is transferred into the capacitor. Assuming that we are in the steady-state (i.e. power is being supplied to the motor at a constant voltage and current), the capacitor voltage will return to its starting value at the start of the next ON time.

If the losses in the transistor and the storage elements are neglected, it is easy to show that the converter functions like an ideal transformer, with output power equal to input power, i.e.

$$V_{in} \times I_{in} = V_{out} \times I_{out}, \quad or \quad \frac{V_{out}}{V_{in}} = \frac{I_{in}}{I_{out}}$$

We know that the motor voltage V_{out} is higher than $V_{in,}$ so not surprisingly the quid pro quo is that the motor current is less than the input current, as indicated by line thickness in Fig. 2.5.

Also if the capacitor is large enough to hold the output voltage very nearly constant throughout, it is easy to show that the step-up ratio is given by

$$\frac{V_{out}}{V_{in}} = \frac{1}{1-D}$$

where D is the duty ratio, i.e. the proportion of each cycle for which the transistor is ON. Thus for the example at the beginning of this section, where $V_{out}/V_{in} = 240/48 = 5$, the duty ratio is 0.8, i.e. the transistor has to be ON for 80% of each cycle.

The advantage of switching at a high frequency becomes clear when we recall that the capacitor has to store enough energy to supply the output during the ON period, so if, as is often the case, we want the output voltage to remain almost constant, it should be clear that the capacitor must store a good deal more energy than it gives out each cycle. Given that the size and cost of capacitors depends on the energy they have to store, it is clearly better to supply small packets of energy at a high rate, rather than use a lower frequency that requires more energy to be stored. Conversely, the switching and other losses increase with frequency, so a compromise is inevitable.

2.3 D.C. from a.c.—Controlled rectification

The majority of drives of all types draw their power from a constant voltage 50 Hz or 60 Hz utility supply, and so in many drive converters the first stage consists of a rectifier which converts the a.c. to a crude form of d.c. Where a constant-voltage (i.e. unvarying average) 'd.c.' output is required, a simple (uncontrolled) diode rectifier is usually sufficient. But where the mean d.c. voltage has to be controllable (as in a d.c. motor drive to obtain varying speeds), a controlled rectifier is used.

Many different converter configurations based on combinations of diodes and thyristors are possible, but we will focus on 'fully-controlled' converters in which all the rectifying devices are thyristors, because they are predominant in modern motor drives. Half-controlled converters are used less frequently, so will not be covered here.

From the user's viewpoint, interest centres on the following questions:

- How is the output voltage controlled?
- What does the converter output voltage look like? Will there be any problems if the voltage is not pure d.c.?
- How does the range of the output voltage relate to the a.c. supply voltage?
- How is the converter and motor drive 'seen' by the supply system? What is the power factor, and is there distortion of the supply voltage waveform, with consequent interference to other users?

We can answer these questions without going too thoroughly into the detailed workings of the converter. This is just as well, because understanding all the ins and outs of converter operation is beyond our scope. On the other hand it is well worth studying the essence of the controlled rectification process, because it assists in understanding the limitations which the converter puts on drive performance (see Chapter 4, etc.). Before tackling the questions posed above, however, it is obviously necessary to introduce the thyristor.

2.3.1 The thyristor

The thyristor is a power electronic switch, with two main terminals (anode and cathode) and a 'switch-on' terminal (gate), as shown in Fig. 2.6.

Like a diode, current can only flow in the forward direction, from anode to cathode. But unlike a diode, which will conduct in the forward direction as soon as forward voltage is applied, the thyristor will continue to block forward current until a small current pulse is injected into the gate-cathode circuit, to turn it

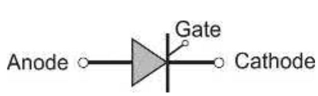

FIG. 2.6 Circuit diagram of thyristor.

ON or 'fire' it. After the gate pulse is applied, the main anode-cathode current builds up rapidly, and as soon as it reaches the 'latching' level, the gate pulse can be removed and the device will remain ON.

Once established, the anode-cathode current cannot be interrupted by any gate signal. The non-conducting state can only be restored after the anode-cathode current has reduced to zero, and has remained at zero for the turn-off time (typically 100–200 μs).

When a thyristor is conducting it approximates to a closed switch, with a forward drop of only one or two volts over a wide range of current. Despite the low volt drop in the ON state, heat is dissipated, and heatsinks must usually be provided, perhaps with fan or other forms of cooling. Devices must be selected with regard to the voltages to be blocked and the r.m.s. and peak currents to be carried. Their overcurrent capability is very limited, and it is usual in drives for devices to have to withstand perhaps twice full-load current for a few seconds only, though this is dependent on the application. Special fuses must be fitted to protect against large fault currents.

The reader may be wondering why we need the thyristor, since in the previous section we discussed how a transistor could be used as an electronic switch. On the face of it the transistor appears even better than the thyristor because it can be switched OFF while the current is flowing, whereas the thyristor will remain ON until the current through it has been reduced to zero by external means. The primary reason for the use of thyristors is that they are cheaper and their voltage and current ratings extend to higher levels than in power transistors. In addition, the circuit configuration in rectifiers is such that there is no need for the thyristor to be able to interrupt the flow of current, so its inability to do so is no disadvantage. Of course there are other circuits (see for example the next section dealing with inverters) where the devices need to be able to switch OFF on demand, in which case the transistor then has the edge over the thyristor.

2.3.2 Single pulse rectifier

The simplest phase-controlled rectifier circuit is shown in Fig. 2.7. When the supply voltage is positive, the thyristor blocks forward current until the gate

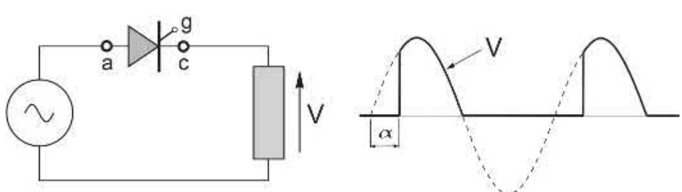

FIG. 2.7 Simple single-pulse thyristor-controlled rectifier, with resistive load and firing angle delay α.

pulse arrives, and up to this point the voltage across the resistive load is zero. As soon as a firing pulse is delivered to the gate-cathode circuit (not shown in Fig. 2.7) the device turns ON, the voltage across it falls to near zero, and the load voltage becomes equal to the supply voltage. When the supply voltage reaches zero, so does the current. At this point the thyristor regains its blocking ability, and no current flows during the negative half-cycle.

If we neglect the small on-state volt-drop across the thyristor, the load voltage (Fig. 2.7) will consists of part of the positive half-cycles of the a.c. supply voltage. It is obviously not smooth, but is 'd.c' in the sense that it has a positive mean value; and by varying the delay angle (α) of the firing pulses the mean voltage can be controlled. With a purely resistive load, the current waveform will simply be a scaled version of the voltage.

This arrangement gives only one peak in the rectified output for each complete cycle of the supply, and is therefore known as a 'single-pulse' or half-wave circuit. The output voltage (which ideally we would like to be steady d.c.) is so poor that this circuit is almost never used in drives. Instead, drive converters use four or six thyristors, and produce much superior output waveforms with two or six pulses per cycle, as will be seen in the following sections.

2.3.3 Single-phase fully-controlled converter—Output voltage and control

The main elements of the converter circuit are shown in Fig. 2.8. It comprises four thyristors, connected in bridge formation. (The term bridge stems from early four-arm measuring circuits which presumably suggested a bridge-like structure to their inventors.)

The conventional way of drawing the circuit is shown on the left side in Fig. 2.8, while on the right side it has been redrawn to assist understanding: the load is connected to the output terminals, where the voltage is V_{dc}. The top of the load can be connected (via T1) to terminal A of the supply, or (via T2) to terminal B of the supply, and likewise the bottom of the load can be connected either to A or to B via T3 or T4 respectively.

We are naturally interested to find what the output voltage waveform on the d.c. side will look like, and in particular to discover how the mean d.c. level can

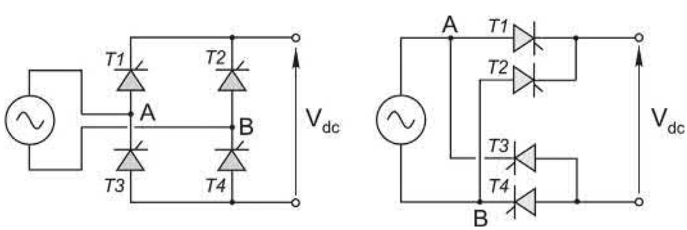

FIG. 2.8 Single-phase 2-pulse (full-wave) fully-controlled rectifier.

be controlled by varying the firing delay angle α. The angle α is measured from the point on the waveform when a diode in the same circuit position would start to conduct, i.e. when the anode becomes positive with respect to the cathode.

This is not such a simple matter as we might have expected, because it turns out that the mean voltage for a given α depends on the nature of the load. We will therefore look first at the case where the load is resistive, and explore the basic mechanism of phase control. Later, we will see how the converter behaves with a typical inductive motor load.

Resistive load

Thyristors T1 and T4 are fired together when terminal A of the supply is positive, while on the other half-cycle, when B is positive, thyristors T2 and T3 are fired simultaneously. The output voltage and current waveform are shown in Fig. 2.9A and B, respectively, the current simply being a replica of the voltage. There are two pulses per cycle of the supply, hence the description '2-pulse' or full-wave.

The load is either connected to the supply by the pair of switches T1 and T4, or it is connected the other way up by the pair of switches T2 and T3, or it is disconnected. The load voltage therefore consists of rectified chunks of the incoming supply voltage. It is much smoother than in the single-pulse circuit, though again it is far from pure d.c.

The waveforms in Fig. 2.9A and B correspond to delay angles $\alpha = 45°$ and $\alpha = 135°$ respectively. The mean value, V_{dc} is shown in each case. It is clear that the larger the delay angle, the lower the output voltage. The maximum output voltage (V_{do}) is obtained with $\alpha = 0°$: this is the same as would be obtained if the thyristors were replaced by diodes, and is given by

$$V_{d0} = \frac{2}{\pi}\sqrt{2}V_{rms} \tag{2.3}$$

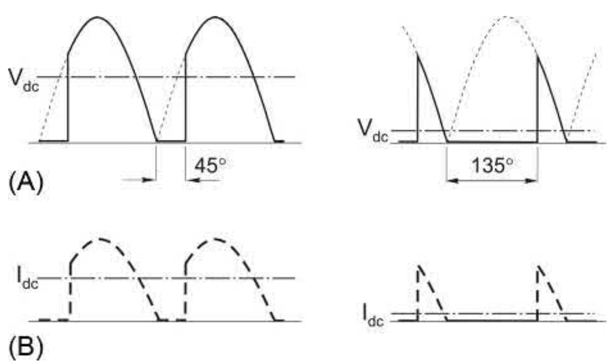

FIG. 2.9 Output voltage waveform (A) and current (B) of single-phase fully-controlled rectifier with resistive load, for firing angle delays of 45° and 135°.

where V_{rms} is the r.m.s. voltage of the incoming supply.

The variation of the mean d.c. voltage with α is given by

$$V_{dc} = \left\{ \frac{1 + cos\alpha}{2} \right\} V_{d0} \tag{2.4}$$

from which we see that *with a resistive load* the d.c. voltage can be varied from a maximum of V_{do} down to zero by varying α from 0° to 180°.

Inductive (motor) load

As mentioned above, motor loads are inductive, and we have seen earlier that the current cannot change instantaneously in an inductive load. We must therefore expect the behaviour of the converter with an inductive load to differ from that with a resistive load, in which the current was seen to change instantaneously.

The realisation that the mean voltage for a given firing angle might depend on the nature of the load is a most unwelcome prospect. What we would like is to be able to say that, regardless of the load, we can specify the output voltage waveform once we have fixed the delay angle α. We would then know what value of α to select to achieve any desired mean output voltage. What we find in practice is that once we have fixed α, the mean output voltage with a resistive-inductive load is not the same as with a purely resistive load, and therefore we cannot give a simple general formula for the mean output voltage in terms of α. This is of course very undesirable: if for example we had set the speed of our unloaded d.c. motor to the target value by adjusting the firing angle of the converter to produce the correct mean voltage, the last thing we would want is for the voltage to fall when the load current drawn by the motor increased, as this would cause the speed to fall below the target.

Fortunately however, it turns out that the output voltage waveform for a given α does become independent of the load inductance once there is sufficient inductance to prevent the load current from ever falling to zero. This condition is known as 'continuous current'; and, happily, many motor circuits do have sufficient self-inductance to ensure that we achieve continuous current. Under continuous current conditions, the output voltage waveform only depends on the firing angle, and not on the actual inductance present. This makes things much more straightforward, and typical output voltage waveforms for this continuous current condition are shown in Fig. 2.10.

The waveforms in Fig. 2.10 show that, as with the resistive load, the larger the delay angle the lower the mean output voltage. However with the resistive load the output voltage was never negative, whereas we see that, although the mean voltage is positive for values of α below 90° there are brief periods when the output voltage becomes negative. This is because the inductance smoothes out the current (see Fig. 4.2, for example) so that at no time does it fall to zero. As a result, one or other pair of thyristors is always conducting, so at every

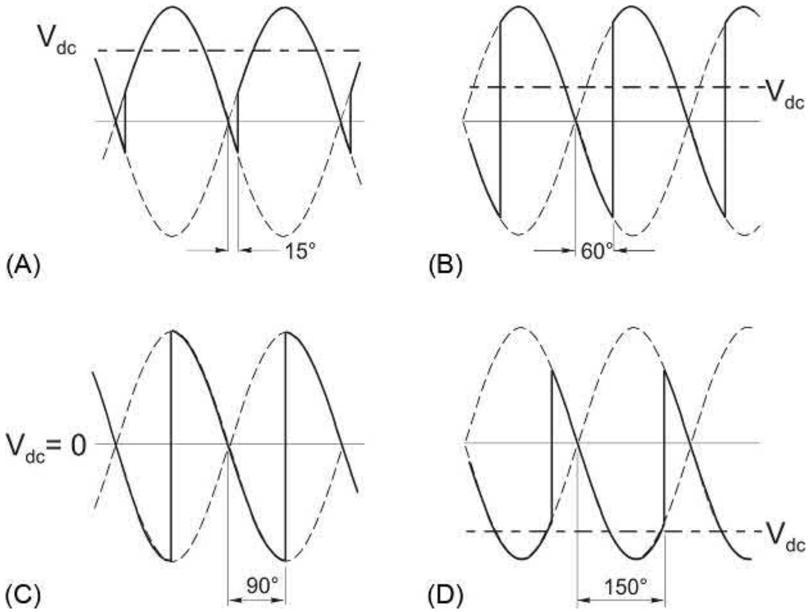

FIG. 2.10 Output voltage waveforms of single-phase fully-controlled rectifier supplying an inductive (motor) load, for various firing angles.

instant the load is connected directly to the supply, and therefore the load voltage always consists of chunks of the supply voltage.

Rather surprisingly, we see that when α is $>90°$, the average voltage is negative (though of course, the current is still positive). The fact that we can obtain a net negative output voltage with an inductive load contrasts sharply with the resistive load case, where the output voltage could never be negative. The combination of negative voltage and positive current means that the power flow is reversed, and energy is fed back to the supply system. We will see later that this facility allows the converter to return energy from the load to the supply, and this is important when we want to use the converter with a d.c. motor in the regenerating mode.

It is not immediately obvious why the current switches (or 'commutates') from the first pair of thyristors to the second pair when the latter are fired, so a brief look at the behaviour of diodes in a similar circuit configuration may be helpful at this point. Consider the set-up shown in Fig. 2.11, with two voltage sources (each time-varying) supplying a load via two diodes. The question is what determines which diode conducts, and how does this influence the load voltage?

We can consider two instants as shown in the diagram. On the left, V_1 is 250 V, V_2 is 240 V, and D1 is conducting, as shown by the heavy line. If we

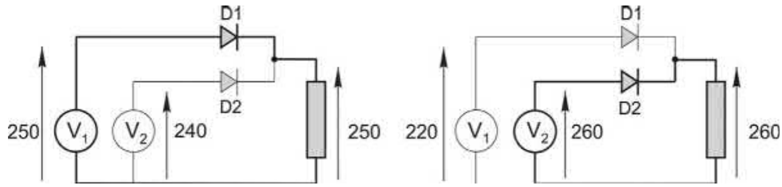

FIG. 2.11 Diagram illustrating commutation between diodes: the current flows through the diode with the higher anode potential.

ignore the volt-drop across the diode, the load voltage will be 250 V, and the voltage across diode D2 will be 240–250 = −10 V, i.e. it is reverse biased and hence in the non-conducting state. At some other instant (on the right of the diagram), V_1 has fallen to 220 V while V_2 has increased to 260 V: now D2 is conducting instead of D1, again shown by the heavy line, and D1 is reverse-biased by −40 V. The simple pattern is that the diode with the highest anode potential will conduct, and as soon as it does so it automatically reverse-biases its predecessor. In a single-phase diode bridge, for example, the commutation occurs at the point where the supply voltage passes through zero: at this instant the anode voltage on one pair goes from positive to negative, while on the other pair the anode voltage goes from negative to positive.

The situation in controlled thyristor bridges is very similar, except that before a new device can take over conduction, it must not only have a higher anode potential, but it must also receive a firing pulse. This allows the changeover to be delayed beyond the point of natural (diode) commutation by the angle α, as shown in Fig. 2.10. Note that the maximum mean voltage (V_{d0}) is again obtained when α is zero, and is the same as for the resistive load (Eq. 2.3). It is easy to show that the mean d.c. voltage is now related to α by

$$V_{dc} = V_{d0} \cos \alpha \tag{2.5}$$

This equation indicates that we can control the mean output voltage by controlling α, though Eq. (2.5) shows that the variation of mean voltage with α is different from that for a resistive load (Eq. 2.4), not least because when α is >90° the mean output voltage is negative.

It is sometimes suggested (particularly by those with a light-current background) that a capacitor could be used to smooth the output voltage, this being common practice in some low-power d.c. supplies. However, the power levels in most d.c. drives are such that in order to store enough energy to smooth the voltage waveform over the half-cycle of the utility supply (20 ms at 50 Hz), very bulky and expensive capacitors would be required. Fortunately, as will be seen later, it is not necessary for the voltage to be smooth as it is the current which directly determines the torque, and as already pointed out the current is always much smoother than the voltage because of inductance.

2.3.4 Three-phase fully-controlled converter

The main power elements are shown in Fig. 2.12. The three-phase bridge has only two more thyristors than the single-phase bridge, but the output voltage waveform is vastly better, as shown in Fig. 2.13. There are now six pulses of the output voltage per cycle, hence the description '6-pulse'. The thyristors are again fired in pairs (one in the top half of the bridge and one—from a different leg—in the bottom half), and each thyristor carries the output current for one third of the time. As in the single-phase converter, the delay angle controls the output voltage, but now $\alpha=0$ corresponds to the point at which the phase voltages are equal (see Fig. 2.13).

The enormous improvement in the smoothness of the output voltage waveform is clear when we compare Figs. 2.13 and 2.10, and it underlines the benefit of choosing a 3-phase converter whenever possible. The very much better voltage waveform also means that the desirable 'continuous current' condition is much more likely to be met, and the waveforms in Fig. 2.13 have therefore been drawn with the assumption that the load current is in fact continuous.

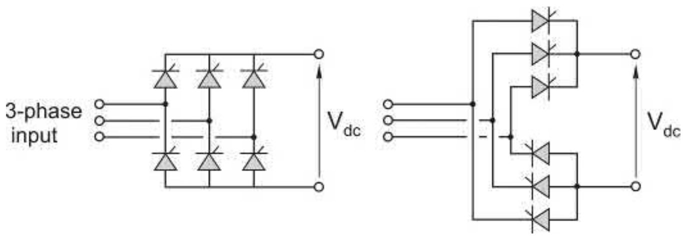

FIG. 2.12 Three-phase fully-controlled thyristor converter. (The alternative diagram (right) is intended to assist understanding.)

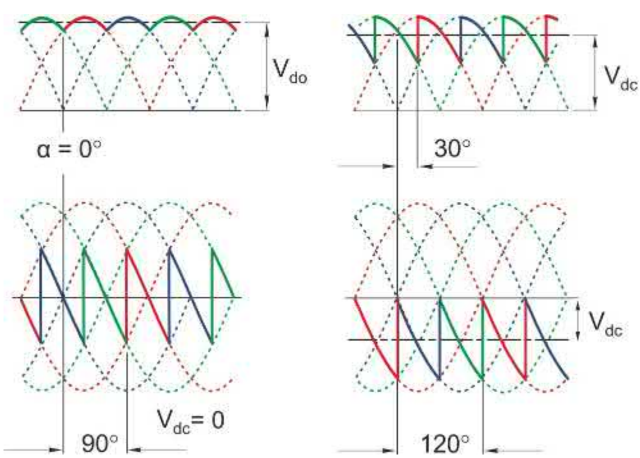

FIG. 2.13 Output voltage waveforms for three-phase fully-controlled thyristor converter supplying an inductive (motor) load, for various firing angles from 0° to 120°. The mean d.c. voltage is shown by the horizontal line, except for $\alpha=90°$ where the mean d.c. voltage is zero.

Occasionally, even a six-pulse waveform is not sufficiently smooth, and some very large drive converters therefore consist of two six-pulse converters with their outputs in series. A phase-shifting transformer is used to insert a 30° shift between the a.c. supplies to the two 3-phase bridges. The resultant ripple voltage is then 12-pulse.

Readers with prior knowledge of three-phase systems will recall that the three line-line voltages differ in phase by 120°, and they may therefore wonder why the output voltage waveforms in Fig. 2.13 show a transition every 60°, rather than 120° as might be expected. The explanation is simply that the top thyristors switch every 120°, as do the bottom ones, but the transitions are 60° out of step, so there are 6 transitions per cycle.

Returning to the 6-pulse converter, as long as the d.c. current remains continuous, the mean output voltage can be shown to be given by

$$V_{dc} = V_{d0}\cos\alpha = \frac{3}{\pi}\sqrt{2}V_{rms}\cos\alpha \tag{2.6}$$

We note that we can obtain the full range of output voltages from $+V_{d0}$ to close to $-V_{d0}$, so that, as with the single-phase converter, regenerative operation will be possible.

It is probably a good idea at this point to remind the reader that, in the context of this book, our first application of the controlled rectifier will be to supply a d.c. motor. When we examine the d.c. motor drive in Chapter 4, we will see that it is the average or mean value of the output voltage from the controlled rectifier that determines the speed, and it is this mean voltage that we refer to when we talk of 'the' voltage from the converter. We must not forget the unwanted a.c. or ripple element, however, as this can be large. For example, we see from Fig. 2.13 that to obtain a very low d.c. voltage (to make the motor run very slowly) α will be close to 90°; but if we were to connect an a.c. voltmeter to the output terminals it could register several hundred volts, depending on the incoming supply voltage!

2.3.5 Output voltage range

In Chapter 4 we will discuss the use of the fully-controlled converter to drive a d.c. motor, so it is appropriate at this stage to look briefly at the typical voltages we can expect. Utility supply voltages vary, but single-phase supplies are usually 200–240 V, and we see from Eq. (2.3) that the maximum mean d.c. voltage available from a single phase 230 V supply is 207 V. This is suitable for 180–200 V motors. If a higher voltage is needed (for say a 300 V motor), a transformer must be used to step up the incoming supply.

Turning now to typical three-phase supplies, the lowest three-phase industrial voltages are usually around 380–480 V. (Higher voltages of up to 11 kV are used for large drives, but these will not be discussed here). So with $V_{rms} = 400$ V for example, the maximum d.c. output voltage (Eq. 2.6) is 540 Volts. After allowances have been made for supply variations and impedance drops, we could not rely on obtaining much more than 500–520 V, and it is usual for the motors used with 6-pulse drives fed from 400 V, 3-phase supplies to be rated in the range 430–500 V. (Often the motor's field winding will be supplied from single phase 230 V, and field voltage ratings are then around 180–200 V, to allow a margin in hand from the theoretical maximum of 207 V referred to earlier.)

2.3.6 Firing circuits

Since the gate pulses are only of low power, the gate drive circuitry is simple and cheap. Often a single electronic integrated circuit contains all the circuitry for generating the gate pulses, and for synchronising them with the appropriate delay angle α with respect to the supply voltage. To avoid direct electrical connection between the high voltages in the main power circuit and the low voltages used in the control circuits, the gate pulses are coupled to the thyristor either by small pulse transformers, or opto-couplers.

2.4 A.C. from d.c.—Inversion

The business of getting a.c. from d.c. is known as inversion, and nine times out of ten we would ideally like to be able to produce sinusoidal output voltages of whatever frequency and amplitude we choose, to achieve speed control of a.c. motors. Unfortunately the constraints imposed by the necessity to use a switching strategy mean that we always have to settle for a voltage waveform which is composed of 'chunks', and is thus far from ideal. Nevertheless it turns out that a.c. motors are remarkably tolerant, and will operate satisfactorily despite the inferior waveforms produced by the inverter.

2.4.1 Single-phase inverter

We can illustrate the basis of inverter operation by considering the single-phase example shown in Fig. 2.14. This inverter uses IGBT's (see later) as the switching devices, with diodes to provide the freewheel paths needed when the load is inductive.

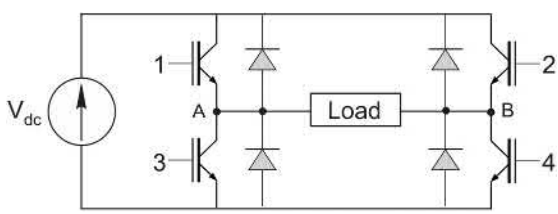

FIG. 2.14 Single-phase inverter.

The input or d.c. side of the inverter (on the left in Fig. 2.14) is usually referred to as the 'd.c. link', reflecting the fact that in the majority of cases the d.c. is obtained by rectifying an incoming constant-frequency utility supply as we have discussed earlier. The output or a.c. side is taken from terminals A and B in Fig. 2.14.

When transistors 1 and 4 are switched ON, the load voltage is positive, and equal to the d.c. link voltage, while when 2 and 3 are ON it is negative. If no devices are switched ON, the output/load voltage is zero. These 3 distinct possible output voltage states lead us to describe such an inverter as a 3 level inverter. Typical output voltage waveforms at low and high switching frequencies are shown in Fig. 2.15A and B, respectively.

Here each pair of devices is ON for one-half of a cycle. Careful inspection of Fig. 2.15 reveals that there is a very short zero voltage period at the transition point. This is because when switching from one pair of transistors to another it is essential that both transistors in the same arm are never switched ON at the same time as this would provide a short circuit or 'shoot-through' path for the d.c. supply: a safety angle is therefore included in the switching logic. The output waveform is clearly not a sine wave, but at least it is alternating and symmetrical. The fundamental component is shown dotted in Fig. 2.15.

We will make frequent references throughout this book to the term 'fundamental component', so it is worth a brief digression to explain what it means for any readers who are not familiar with Fourier analysis. In essence, any periodic waveform can be represented as the sum of an infinite series of sinusoidal waves or 'harmonics', the frequency of the lowest harmonic (the 'fundamental') corresponding to the frequency of the original waveform: Fourier analysis provides the means for calculating the magnitude, phase and order of the harmonic components.

For example, the fundamental, third and fifth harmonics of a square wave are shown in Fig. 2.16.

In this particular case, the symmetry of the original waveform is such that only odd harmonics are present, and it turns out that the amplitude of each harmonic is inversely proportional to its order. We can see from Fig. 2.16 that the sum of the fundamental, third and fifth harmonics makes a reasonably good approximation to the original square wave, and if we were to include higher harmonics, the approximation would be even better.

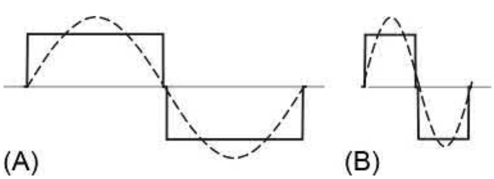

FIG. 2.15 Single-phase inverter output voltage waveforms—resistive load.

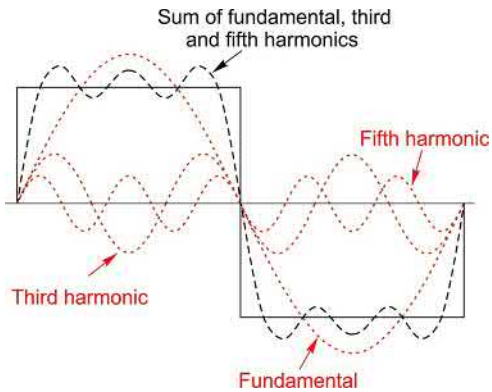

FIG. 2.16 Harmonic components of a square wave.

Readers who have studied a.c. circuits under sinusoidal conditions will recall that there are powerful methods based on complex numbers for the steady-state analysis of such circuits, in which inductors and capacitors are represented in terms of their reactances, which depend on the supply frequency. These techniques clearly cannot be applied directly when the waveforms are not sinusoidal, but, because most electrical circuits are linear (i.e. they obey the principle of superposition), we can solve the response to each harmonic separately, and add the responses to determine the final output.

For example, if we want to find the steady-state current in an inductor when a square wave voltage is applied, we find the currents due to the fundamental and each harmonic, and sum the separate currents to obtain the result. Hence the current is given by

$$I = \frac{V_1}{X_1} + \frac{V_3}{X_3} + \frac{V_5}{X_5} + \cdots,$$

where X_n is the reactance at the nth harmonic frequency.

Recalling that under sinusoidal conditions the reactance of an inductance (L) at frequency f is given by $X = 2\pi fL = \omega L$, we note that reactance is directly proportional to frequency. Hence if the reactance at the fundamental frequency is X_1, the current is given by

$$I = \frac{V_1}{X_1} + \frac{V_3}{3X_1} + \frac{V_5}{5X_1} + \cdots,$$

It should be clear that for the square wave, where the amplitude of the third harmonic voltage is one third that of the fundamental, the third harmonic current will be only one ninth of the fundamental, and the fifth harmonic current will be only one twenty-fifth. In other words, the current harmonics decrease very rapidly with order, which means that the current waveform is more sinusoidal than

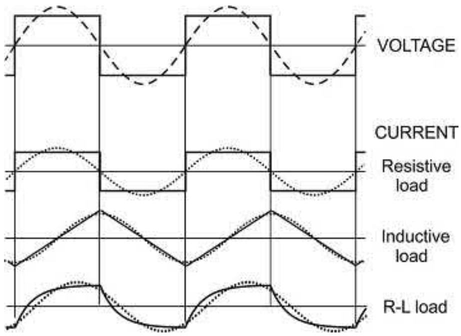

FIG. 2.17 Effect of the load on the current waveform of a square wave voltage.

the voltage waveform. (In contrast, a resistive load looks the same to each harmonic, so the current waveform is the same as the voltage.)

These conclusions are confirmed by Fig. 2.17, which shows the steady state current waveforms when a square wave of voltage is applied to various loads, the fundamental being shown by the dotted line.

The most noteworthy feature as far as we are concerned is how much smoother the current wave is for the purely inductive case, and how, as a result of its reduced harmonic content, the actual current is closely approximated by the fundamental component. We note also that, as expected, the fundamental current is in phase with the voltage for the resistive load; it lags by 90° for the purely inductive load; and it lies in between for the mixed R-L load.

All this is very fortunate for our purposes. Later in the book we will be studying the behaviour of a.c. motors (that were originally designed for use from the sinusoidal utility supply) when fed instead from inverters whose voltage waveforms are not sinusoidal. We will discover that, because the motors have high reactances at the harmonic frequencies, the currents are much more sinusoidal than the voltages. We can therefore apply our knowledge of how motors perform with sinusoidal waveforms by ignoring all except the fundamental components of the inverter waveforms, and nevertheless obtain a close approximation to the actual behaviour.

2.4.2 Output voltage control

There are two ways in which the amplitude of the output voltage can be controlled. First, if the d.c. link is provided from the utility supply via a controlled rectifier or from a battery via a chopper, the d.c. link voltage can be varied. We can then set the amplitude of the output voltage to any value within the range of the d.c. link. For a.c. motor drives (see Chapter 7) we can arrange for the link voltage to track the output frequency of the inverter, so that at high output frequency we obtain a high output voltage and vice-versa. This method of voltage control results in a simple inverter, but requires a controlled (and thus relatively expensive) rectifier for the d.c. link.

The second method, the principle of which predominates in all sizes, achieves voltage control by pulse-width-modulation (PWM, see Section 2.2.1) within the inverter itself i.e. it does not require any additional power semiconductor switches. A cheaper uncontrolled rectifier can then be used to provide a constant-voltage d.c. link.

The principle of voltage control by PWM is illustrated in Fig. 2.18, which assumes a resistive load for the sake of simplicity.

There are many methods for generating PWM waveforms, but originally this was done by comparing a triangular modulating wave, the frequency of which is often referred to as "the switching frequency", that sets the pulse frequency, shown at the top left, with a reference sinusoidal wave that determines the frequency and amplitude of the desired sinusoidal output voltage.

The inverter (on the right) is the same 3-level one discussed in Section 2.4.1, Fig. 2.14. When devices 1 and 4 are ON, the load voltage is positive and current takes the red path shown, and when devices 2 and 3 are ON the load voltage is negative and the current takes the blue path. When no devices are ON the voltage and current are both zero. With a purely resistive load, the freewheel diodes never conduct current.

The intersections of the reference and modulating wave define the switching times. Taking the positive half-cycle for example, if the amplitude of the reference wave is greater than that of the triangular wave, devices 1 and 3 are turned ON, and they remain ON until the amplitude of the reference wave becomes less than that of the triangular wave, whereupon they are turned OFF.

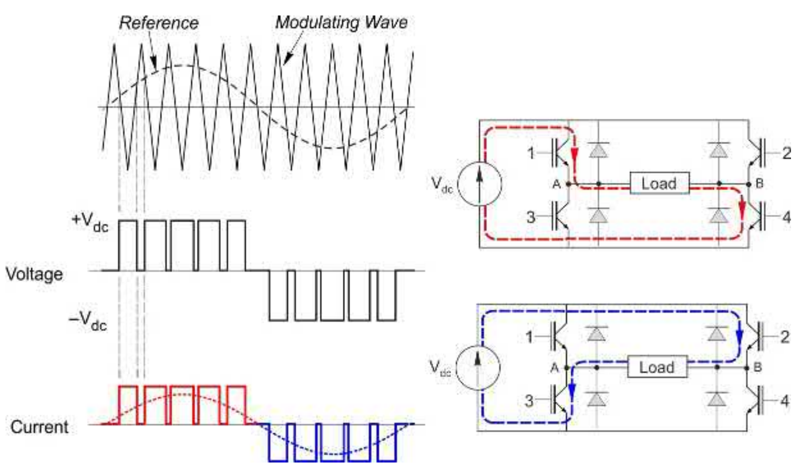

FIG. 2.18 Inverter output voltage and frequency control with pulse-width modulation—resistive load.

Readers who do not find it self-evident that this strategy generates a good waveform (i.e. one with the desired fundamental component and low harmonic content) can rest assured that it is a well-proven approach. We should also acknowledge that in a drive converter, the frequency of the modulating wave would be very much higher than that shown (for the sake of clarity) in Fig. 2.18, which in turn means that the harmonic components originating from the modulating wave will be at such a high frequency that they will be barely visible in the current waveform.

The introductory discussion above applies to a resistive load, but in the drives context the load will always include an inductive element. As a result, as we have seen earlier, when all four devices are turned OFF, the current does not fall to zero as it does with a resistive load. Instead it continues by taking the path through a pair of freewheel diodes, and within each cycle, the current will normally still be flowing until the next pair of devices turns ON, i.e. we have what we earlier called a 'continuous current' condition.

Typical voltage and current waveforms for an inductive load are shown in Fig. 2.19.

The output voltage no longer has any zero intervals (see Fig. 2.20 later for details), but our interest here is to show that the amplitude of the fundamental component of the output voltage is determined by the amplitude of the reference wave (dotted).

The left-hand waveform represents a medium amplitude condition in which every cycle of the modulating wave results in a switch ON and switch OFF, and

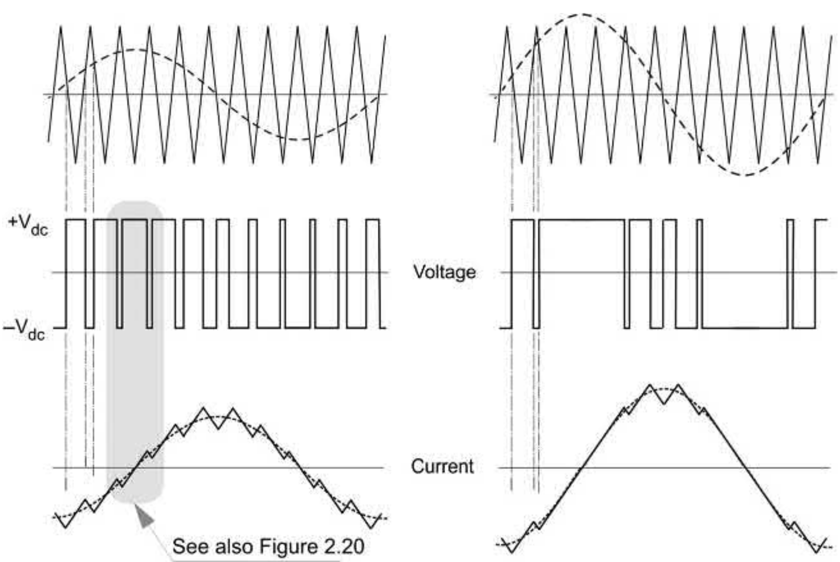

FIG. 2.19 Amplitude control by PWM—inductive load.

in this condition the output harmonic content is modest. However when a very high amplitude is called for (on the right in Fig. 2.19), the amplitude of the reference signal may exceed that of the modulating wave for an extended period. This is the so-called 'over-modulated' condition, the final limit of which takes us back to the square wave, with a consequent unwelcome increase in harmonic content.

An enlarged view of the shaded region in Fig. 2.19 is shown in Fig. 2.20. This period includes the transition in the load current from the negative half-cycle (on the left of the vertical dotted line) to the positive half-cycle. The four sketches (two each on either side of the zero current crossover) represent all possible modes of operation of an inverter when there are no discontinuities in the load current. For the sake of simplicity, only the devices that are conducting are shown in each sketch.

Mode A

Ignoring the low voltage drop across each conducting device, we can see in sketch A that the load is connected to the d.c. link via the two transistors (2 and 3 in Fig. 2.18). The right-hand side of the load is connected to the positive terminal of the d.c. link, so, using the polarity convention adopted previously,

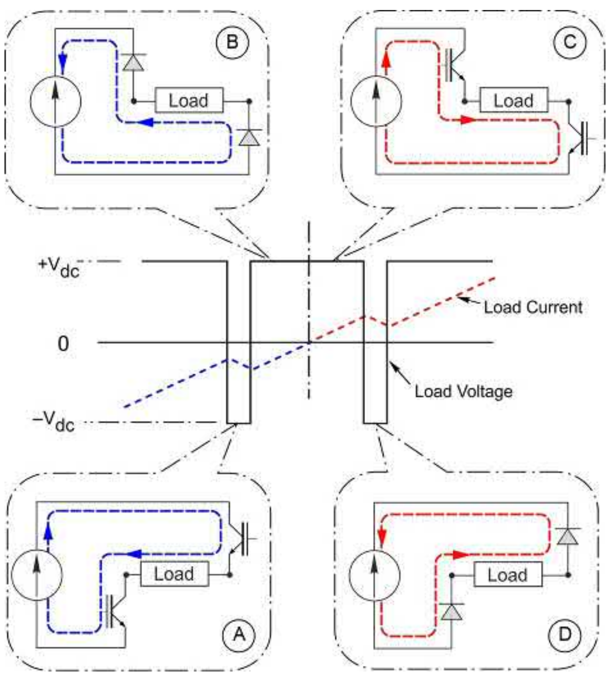

FIG. 2.20 The commutation process in a single-phase bridge.

the voltage across the load is negative. The current is also negative (blue), so the power is positive and energy is flowing into the load from the d.c. link.

In this particular case the load is a pure inductance, so the negative voltage causes the already negative current to increase, the additional energy being stored in the magnetic field. More generally, if the inverter is supplying a motor, this condition would correspond to motoring, with energy being supplied to the motor from the d.c. link.

Mode B

If the transistors in sketch A are switched OFF, the current that is already flowing in sketch A will have to go somewhere, because the load is inductive and the current cannot change instantaneously. Even if they were switched ON, the current cannot flow through the transistors (1 and 4 in Fig. 2.18), because they will only conduct current in one direction (down the page as we look at the diagram). The freewheel diodes shown in sketch B will therefore conduct the negative load current (blue), but because the load is now connected to the d.c. link via the two diodes, the load voltage has reversed, and is now positive.

This changeover or ‘commutation’ happens almost instantaneously: the power is now negative, and energy is flowing from the load to the d.c. link. With a pure inductive load (as here) energy is supplied from the magnetic field of the inductor, but if the inverter is supplying a motor this period would represent the recovery of excess kinetic energy from the motor and its load, an efficient technique known as ‘regenerative braking’ that we will encounter later in the book.

Mode C

If we switch transistors 1 and 4 (see Fig. 2.18) ON at any time during period B, they will not conduct until the polarity of the current becomes positive (red), whereupon the current will cease flowing in the freewheel diodes and commutate naturally to the transistors, as shown in sketch C. The load voltage and current are now both positive, so the power is positive and energy is flowing from the d.c. link to the load. The current in the load will increase.

Mode D

If we now switch OFF the transistors in sketch C the positive current must find a new path, so it flows through the freewheel diodes. The load current is positive but the voltage negative and so energy is flowing from the load to the d.c. link.

This example highlights the fact that the inverter is inherently capable of supplying or accepting energy during both positive and negative half-cycles. It will emerge later that this facility is extremely important when the inverter drives a motor.

It is also interesting, and from a design perspective important, to note that when energy is flowing from the d.c. link to the inductive load then the transistors are conducting and when energy is flowing from the load to the d.c. link the freewheel diodes are conducting. It follows that the power factor of the load therefore has a significant impact on the necessary rating of the freewheel diodes in the converter. With a low power factor load, energy flows in and out during each cycle, thereby stressing the diodes. In contrast, if the power factor is unity, there are no energy reversals, and only the transistors carry the load current.

Having digressed to look in detail at the internal modes of operation of a PWM inverter, we now return to the reason why PWM is employed, namely to control the amplitude of the output voltage wave.

It should be clear from the discussion based on Fig. 2.19 that when the frequency reaches the point at which PWM ceases to be effective in supplying the required voltage, the output waveform degenerates to a square wave whose fundamental harmonic has an amplitude of $4/\pi$ times the d.c. link voltage: the d.c. link voltage therefore determines the maximum a.c. voltage that can be produced by the inverter.

The number, width and spacing of the pulses of PWM inverters are optimised to keep the harmonic content as low as possible and fundamental voltage as high as possible. There is an obvious advantage in using a high switching frequency since there are more pulses to play with. Ultrasonic switching frequencies (generally understood to be $>16\,\text{kHz}$, but for we older folk switching frequencies above 10 kHz are actually inaudible!) are now widely used, and as devices improve, frequencies continue to rise. The latest Wide Band Gap devices, mentioned earlier, allow switching frequencies in the MHz range. Most manufacturers claim their particular PWM switching strategy is better than the competition, but it is not clear which, if any, is best for motor operation. Some early schemes used comparatively few pulses per cycle, and changed the number of pulses in discrete steps with output frequency rather than smoothly, which earned them the nick-name ‘gear-changers’: their sometimes irritating sound can still be heard in older traction applications. Others have irregular switching patterns, which claim to develop “white noise” rather than a single irritating tone. As a rule of thumb, the higher the switching frequency the lower the noise, but the higher the losses.

2.4.3 Three-phase inverter

Single-phase inverters are seldom used for drives, the vast majority of which use 3-phase motors because of their superior characteristics. A 3-phase output can be obtained by adding only two more switches to the four needed for a single-phase inverter, giving the typical power-circuit configuration shown in Fig. 2.21.

As usual, a freewheel diode is required in parallel with each switching device to protect against overvoltages caused by an inductive (motor) load. We note that the circuit configuration in Fig. 2.21 is essentially the same as for the 3-phase controlled rectifier discussed earlier. We mentioned then that

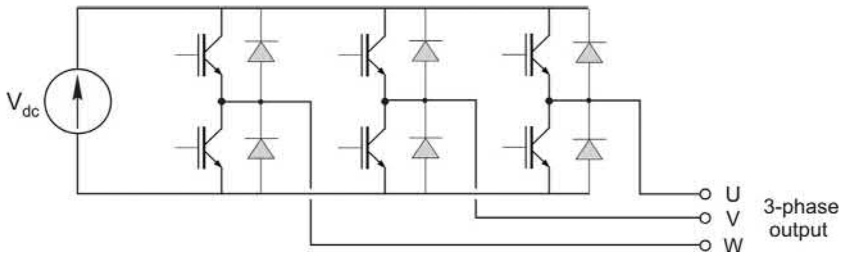

FIG. 2.21 Three-phase inverter power circuit.

the controlled rectifier could be used to regenerate, i.e. to convert power from d.c. to a.c., and this is, of course, 'inversion' as we now understand it.

This circuit forms the basis of the majority of converters for motor drives, and will be explored more fully in Chapter 7. In essence, the output voltage and frequency are controlled in much the same way as for the single-phase inverter discussed in the previous section, to produce the line-to-line output voltages (V_{uv}, V_{vw}, and V_{wu}) which are identical waveforms but displaced by 120° from each other. For reasons of simplicity, let us consider square wave operation, as shown in the lower part of Fig. 2.22.

The derivation of the line-line voltages can best be understood by considering the potentials of the mid-points of each leg (U, V, and W), as shown in the upper part of Fig. 2.21. The zero reference is taken as the negative terminal of the d.c. link, and the sketches show that, as expected, each mid-point is either connected to the positive terminal of the d.c. link, in which case the potential is V_{dc}, or to the negative (reference) in which case the potential is zero. To obtain the line voltages, for example $V_{uv,}$ we subtract the potential of V from the potential of U, and this yields the three lower waveforms.

Not surprisingly, there are more constraints on switching here than in the single-phase case, because we are producing three output waveforms rather than one, and we only have two more switches to play with. In addition, it is easy to see that applying PWM to each phase may well throw up an unacceptable requirement for the upper and lower switches in one leg to be on simultaneously. Fortunately, these potential problems are limited to well-defined regions of the cycle and switching patterns or associated protection circuits include built-in delays to prevent such 'shoot-through' faults from occurring.

Numerous 3 phase PWM switching strategies have been developed to maximise the available voltage whilst retaining the waveform enhancing effects of PWM. The most common of these is to add third harmonic (usually 1/6 the magnitude of the fundamental) to the reference waveform. This boosts the available fundamental voltage and has no impact on torque in a symmetrically wound 3 phase a.c. motor.

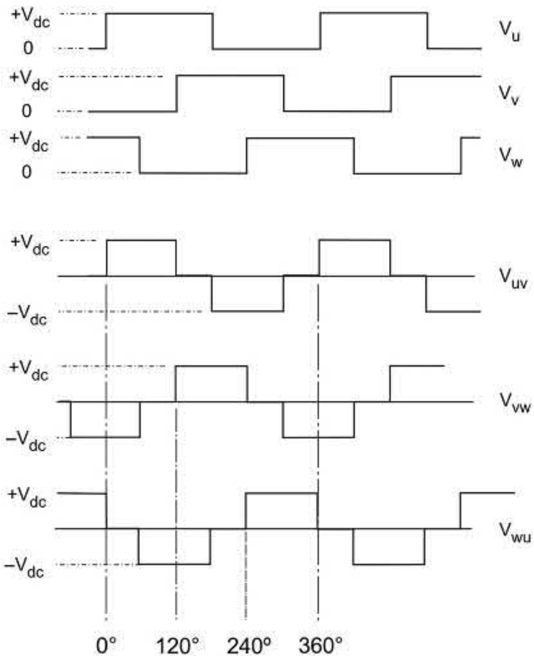

FIG. 2.22 Three-phase square wave output voltage waveforms.

It is worth mentioning at this stage that in a.c. motor drives, it is usually desirable to keep the ratio of voltage/frequency constant in order to keep the motor flux at its optimum level. This means that as the frequency of the reference sine wave is increased, its amplitude must increase in direct proportion. The frequency corresponding to the maximum available output voltage is often taken to define the 'base speed' of the motor, and the frequency range up to this point is often referred to as the 'constant torque' region, as it is only in this region that the motor is operated with its designed flux, and can therefore provide continuous rated shaft torque. Above the frequency corresponding to base speed, the inverter can no longer match voltage to frequency, the inverter effectively having run out of steam as far as voltage is concerned. The motor will then be operating at a level of flux below the design level, and the available continuous torque will reduce with frequency/speed. This region is often referred to as the 'Constant Power' region.

As mentioned earlier, the switching nature of these converter circuits results in waveforms which contain not only the required fundamental component but also unwanted harmonic voltages. It is particularly important to limit the magnitude of the low-order harmonics because these are most likely to provoke an unwanted torque response from the motor, but the high-order harmonics can lead to acoustic noise particularly if they happen to excite a mechanical resonance.

2.4.4 Multi-level inverter

The three-phase, three level inverter power circuit shown in Fig. 2.21, where the d.c. supply is derived from a three phase diode bridge, is widely used in drive systems rated to 2 MW and above at voltages from 400 to 690 V. At ratings above 1 MW, products designed for direct connection to medium voltage supplies (2 kV … 11 kV) are also available. At higher powers, switching the current in the devices proves more problematic in terms of the switching losses, and so switching frequencies have to be reduced. At higher voltages the impact of the rate of change of voltage on the motor insulation requires switching times to be extended deliberately, thereby further increasing losses, and limiting switching frequencies. In addition, whilst higher voltage power semiconductors are available, they tend to be relatively expensive and so commercial consideration is given to the series connection of devices. However, voltage sharing between devices then becomes a problem due to the disproportionate impact of any small difference between switching characteristics.

Multi-level converters have been developed which address these issues. One example from among many different topologies is shown in Fig. 2.23.

In this example four capacitors act as potential dividers to provide four discrete voltage levels, and each arm of the bridge has 4 series connected IGBT's with anti-parallel diodes. The four intermediate voltage levels are connected by means of clamping diodes to the link between the series IGBT's. To obtain full voltage between the outgoing lines all four devices in one of the upper arms are switched ON, together with all four in a different lower arm, while for say half voltage only the bottom two in an upper arm are switched ON. The quarter and three quarter levels can be selected in a similar fashion, and in this way a good 'stepped' approximation to a sine wave can be achieved. This can then be further enhanced by PWM.

Clearly there are more devices turned ON simultaneously than in a basic inverter, and this gives rise to an unwelcome increase in the total conduction loss. But this is offset by the fact that because the VA rating of each device is smaller than the equivalent single device, the switching loss is much reduced.

For the inverter in Fig. 2.23 the stepped waveform will have four positive, four negative and a zero level. The oscilloscope trace in Fig. 2.24 is from a 13-level inverter, and clearly displays the 13 discrete levels. A sophisticated modulation strategy is required in order not only to achieve a close approximation to a sine wave, but also to ensure that the reduction of charge (and voltage) across each d.c. link capacitor during its periods of discharge period is compensated by a subsequent charging current.

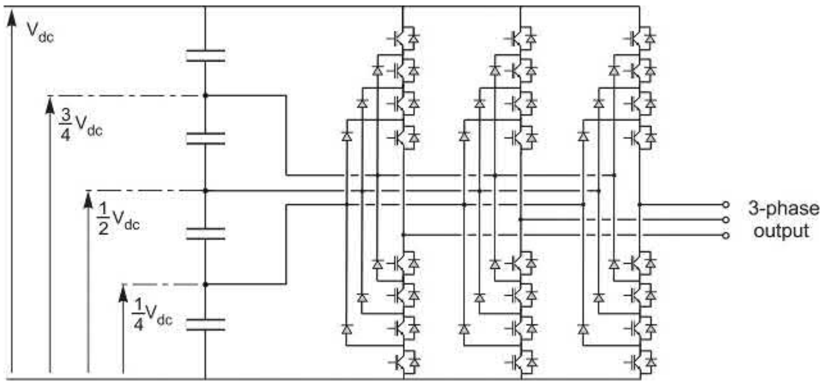

FIG. 2.23 A multi-level inverter.

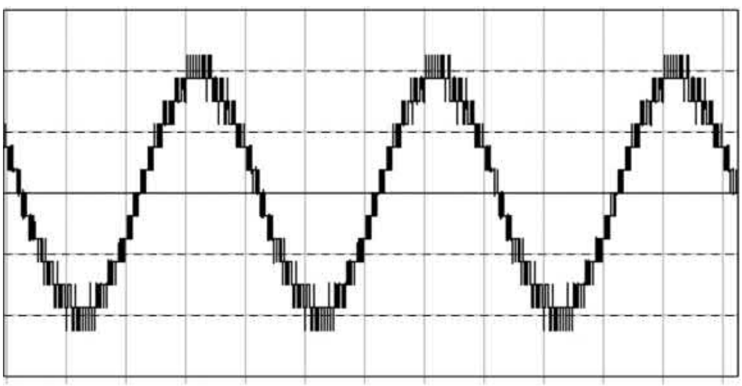

FIG. 2.24 A 13-level inverter voltage waveform.

Multi-level inverters have advantages over the conventional PWM inverter:

- A higher effective output frequency is obtained for a given PWM frequency, and smaller filter components are required.
- EMC (Electromagnetic compatibility) is improved due to the lower dV/dt at output terminals (lower voltage steps in the same switching times), which also reduces stress on the motor insulation.
- Higher DC link voltages are achievable for medium voltage applications due to voltage sharing of power devices within each inverter leg.

However, there are drawbacks:

- The number of power devices is increased by at least a factor of two; each IGBT requires a floating (i.e. electrically isolated) gate drive and power supply; and additional voltage clamping diodes are required.

- The number of DC bus capacitors may increase, but this is unlikely to be a practical problem as lower voltage capacitors in series are likely to offer a cost-effective solution.
- The balancing of the d.c. link capacitor voltages requires careful management/control.
- Control/modulation schemes are more complicated.

For low-voltage drives (690 V and below) the disadvantages of using a multi-level inverter tend to outweigh the advantages, with even five - level converters being significantly more expensive than conventional topologies. However multi-level converters are entering this market, so the situation may change.

2.4.5 Braking

As we have seen, the three-phase inverter power stage shown in Fig. 2.21 inherently allows power to flow in either direction between the d.c. link and the motor. However, the simple diode rectifier that is often used between the utility supply and the d.c. terminals of the inverter does not allow power to flow back into the supply. Therefore an a.c. motor drive based on this configuration cannot be used where power is required to flow from the motor to the utility supply. On the face of it, this limitation might be expected to cause a problem in almost all applications when shut down of a process is involved and the machine being driven has to be braked by the motor. The kinetic energy has to be dissipated, and during active deceleration power flows from the motor to the d.c. link, thereby causing the voltage across the d.c. link capacitor to rise to reflect the extra stored energy.

To prevent the power circuit from being damaged, d.c. link over-voltage protection is included in most industrial drive control systems. This shuts down the inverter if the d.c. link voltage exceeds a trip threshold, but this is at best a last resort. It may be possible to limit the voltage across the d.c. link capacitor by restricting the energy flow from the motor to the d.c. link by controlling the slow down ramp, but this is not always acceptable to the user.

To overcome this problem, the diode rectifier can be supplemented with a d.c. link braking resistor circuit as shown in Fig. 2.25.

In this arrangement, when the d.c. link voltage exceeds the braking threshold voltage, the switching device is turned ON. Provided the braking resistor has a low enough resistance to absorb more power than the power flowing from the inverter, the d.c. link voltage will begin to fall and the switching device will be turned OFF again. In this way the ON and OFF times are automatically set depending on the power from the inverter, and the d.c. link voltage is limited to the braking threshold. To limit the switching frequency of the braking circuit, appropriate hysteresis is included in its control. Note the freewheel diode across the braking resistor, which is required because even braking resistors have inductance!

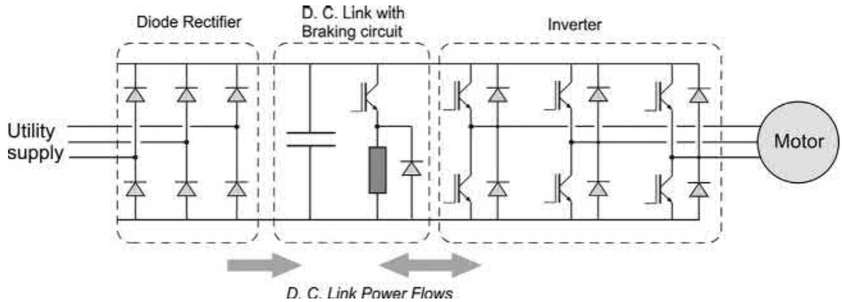

FIG. 2.25 A.C. motor drive with a braking resistor circuit.

A braking resistor circuit is used in many applications where it is practical and acceptable to dissipate stored kinetic energy in a resistor, and hence there is no longer a limit on the dynamic performance of the system.

2.4.6 Active front end

An alternative way to deal with energy flow from the inverter to the utility supply is the active rectifier shown in Fig. 2.26, in which the diode rectifier is replaced with an IGBT inverter.

The labelling of the two converters shown in Fig. 2.26 reflects their functions when the drive is operating in its normal 'motoring' sense, but during active braking (or even continuous generation) the converter on the left will be inverting power from the d.c. link to the utility supply, while that on the right will be rectifying the output from the induction generator.

This arrangement if often referred to as an 'Active Front End'. Additional input inductors are usually required to limit the unwanted switching frequency currents generated by the switching action of the active rectifier, but by using PWM control the supply converter can be controlled to give near sinusoidal supply current waveforms with a power factor close to unity, so that the complete system presents a near perfect load to the utility supply.

An active rectifier is used where motoring and braking/generating operation and/or good quality input waveforms are required. Cranes and elevators, engine test rigs and cable laying ships are some applications where an active rectifier

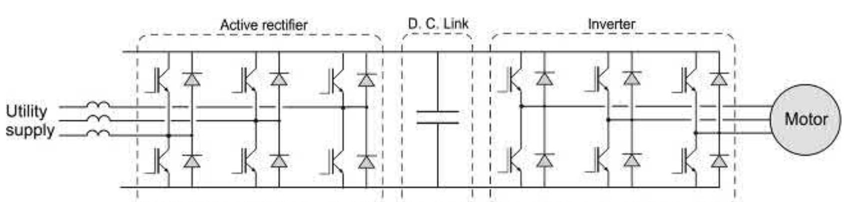

FIG. 2.26 A.C. motor drive with active front-end rectifier.

may be appropriate. The performance characteristics of this configuration are summarised below:

- Power flow between the motor and the utility supply is possible in both directions, and so this makes the drive more efficient than when a braking resistor is used.
- Good quality input current waveforms i.e. low harmonic distortion of the utility supply.
- Supply power factor can be controlled to near unity (in fact it is technically possible to control this sort of bridge to operate at any power factor from approximately 0.8 lagging to 0.8 leading).

Clearly, however, an active rectifier is significantly more expensive than a simple diode rectifier and a brake resistor.

2.5 A.C. from a.c.

The converters we have looked at so far have involved a d.c. stage, but the ideal power electronic converter would allow power conversion in either direction between two systems of any voltage and frequency (including d.c.), and would not involve any intermediate stage, such as a d.c. link. In principle this can be achieved by means of an array of switches that allow any one of a set of input terminals to be connected to any one of a set of output terminals, at any desired instant. Whilst direct a.c. to a.c. converters have had only limited niche areas of application, two forms are worthy of mentioning: the Cycloconverter and the Matrix Converter.

2.5.1 The cycloconverter

The cycloconverter has always been considered a niche topology, but is still used for very large (e.g. 1 MW and above) low-speed induction motor or synchronous motor drives. The reason for its limited use lies primarily in the fact that the cycloconverter is only capable of producing acceptable output waveforms at frequencies well below the utility frequency, but this, coupled with the fact that it is feasible to make large motors with high pole-numbers (e.g. 20—see later) means that a very low-speed direct (gearless) drive becomes practicable. A 20-pole motor, for example, will have a synchronous speed of only 30 rev/min at 5 Hz, making it suitable for mine winders, kilns, crushers, etc.

The principal advantage of the cycloconverter is that naturally commutated devices (thyristors) can be used instead of self-commutating devices, which means that the cost of each device is low and high powers can be readily achieved.

The power conversion circuit for each of the three output phases is the same, so we can consider the simpler question of how to obtain a single variable-frequency supply from a three-phase supply of fixed frequency and constant voltage. We will see that in essence, the output voltage is synthesised by switching the load directly across whichever phase of the utility supply gives the best approximation to the desired load voltage at each instant of time.

Assuming that the load is an induction motor, we will discover later in the book that the power-factor varies with load but never reaches unity, i.e. that the current is never in phase with the stator voltage. So during the positive half-cycle of the voltage waveform the current will be positive for some of the time, but negative for the remainder; while during the negative voltage half-cycle the current will be negative for some of the time and positive for the rest of the time. This means that the supply to each phase has to be able to handle any combination of both positive and negative voltage and current.

We have already explored how to achieve a variable-voltage d.c. supply using a thyristor converter, which can handle positive currents, but here we need to handle negative current as well, so for each of the three motor windings we will need two converters connected in parallel, as shown in Fig. 2.27, making a total of 36 thyristors. To avoid short-circuits, between phases, isolating transformers, also shown in Fig. 2.27, are used.

Previously the discussion focused on the mean or d.c. level of the output voltages, but here we want to provide a low-frequency (preferably sinusoidal) output voltage for an induction motor, and the means for doing this should now be becoming clear. Once we have a double thyristor converter, we can generate a low-frequency sinusoidal output voltage simply by varying the firing angle of the positive-current bridge so that its output voltage increases from zero in a sinusoidal manner with respect to time. Then, when we have completed the positive half-cycle and arrived back at zero voltage, we bring the negative bridge into play and use it to generate the negative half-cycle, and so on.

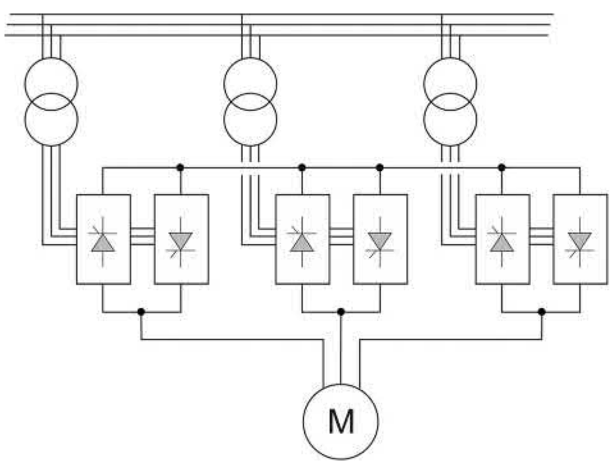

FIG. 2.27 Cycloconverter power circuit.

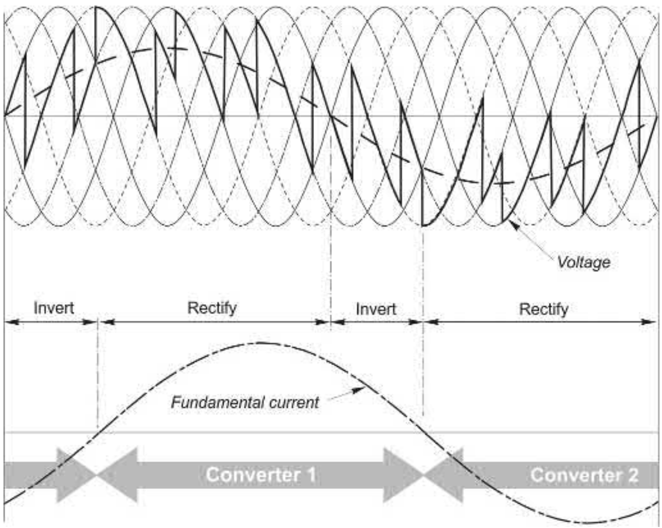

FIG. 2.28 Typical output voltage waveform for one phase of 6-pulse cycloconverter supplying an inductive (motor) load. The output frequency shown in the figure is one-third of the utility frequency, and the amplitude of the fundamental component of the output voltage (shown by the dotted line) is about 75% of the maximum that could be obtained. The fundamental component of the load current is shown in order to define the modes of operation of the converters.

Hence the output voltage consists of chunks of the incoming supply voltage, and it offers a reasonable approximation to the fundamental-frequency sine wave shown by the dotted line in Fig. 2.28.

The output voltage waveform is not significantly worse than the voltage waveform from a d.c. link inverter, and as we saw in that context, the current waveform in the motor will be much smoother than the voltage, because of the filtering action of the motor self leakage inductance. The motor performance will therefore be acceptable, despite the extra losses, which arise from the unwanted harmonic components. We should note that because each phase is supplied from a double converter, the motor can regenerate when required (e.g. to restrain an overhauling load, or to return kinetic energy to the supply when the frequency is lowered to reduce speed).

2.5.2 The matrix converter

Much attention has been focused in technical journals on the matrix converter, the principle of which is shown in Fig. 2.29.

The matrix converter operates in much the same way as a cycloconverter. For example, if we assume that we wish to synthesise a 3-phase sinusoidal output of a known voltage and frequency, we know at every instant what voltage we want, say between the lines A and B, and we know what the voltages are

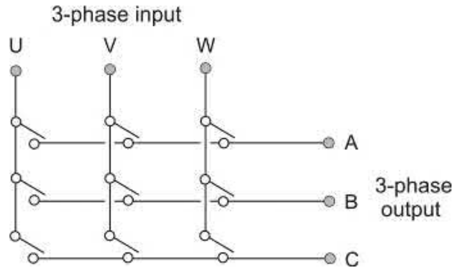

FIG. 2.29 Matrix converter power circuit.

between the three incoming lines. So we switch ON whichever pair of the A and B switches connects us to the two incoming lines whose voltage at the time is closest to the desired output line-to-line voltage, and we stay with it while ever it offers the best approximation to what we want. As soon as a different combination of switches would allow us to hook onto a more appropriate pair of input lines, the switching pattern changes.

Because there are only three different incoming line-to-line voltages to choose from, we cannot expect to synthesise a decent sinusoidal waveform with simple switching, so in order to obtain a better approximation to a sinusoidal waveform we must use chopping. This means that the switches have to be capable of switching OFF as well as ON and operating at much higher switching frequencies than the fundamental output frequency, so that switching loss becomes an important consideration.

Whilst the basic circuit is not new, recent advances in power devices offer the potential to overcome many of the drawbacks inherent in the early implementations that used discrete IGBT's due to the lack of suitable packaged modules. On the other hand, the fact that the output voltage is limited to a maximum of 86% of the utility supply voltage (without over-modulation and subsequent loss of voltage waveform quality) means that applications in the industrial market, where standard motors predominate, remain largely problematic. However, there appear to be good prospects in some aerospace applications and possibly in some specific areas of the industrial drives market, notably integrated motors (where the drive is built into the motor, the windings of which can therefore be designed to suit the voltage available).

2.6 Inverter switching devices

As far as the user is concerned, it does not really matter what type of switching device is used inside the inverter, but it is probably helpful to mention the three most important[3] families of devices in current use so that the terminology is familiar and the symbols used for each device can be recognised. The common

3. The gate turn-off thyristor is now seldom used, so is not discussed here.

feature of all three devices is that they can be switched ON and OFF by means of a low-power control signal, i.e. they are self-commutating. We have seen earlier that this ability to be turned ON or OFF on demand is essential in any inverter which feeds an active load, such as an induction motor.

Each device is discussed briefly below, with a broad indication of its most likely range of application. Because there is considerable overlap between competing devices, it is not possible to be dogmatic and specify which device is best, and the reader should not be surprised to find that one manufacturer may offer a 5 kW inverter which uses MOSFETs while another chooses to use IGBTs.

Power electronics is still subject to substantial innovation, and there are significant developments in the area of power semiconductor devices which are starting to impact drives. A lot of publicity and anticipation has surrounded the birth of Wide Band Gap (WBG) devices, such as Silicon Carbide and Gallium Nitride, which switch much more quickly than silicon-based devices and can operate at higher temperatures. This opens up opportunities for lower switching losses, significantly better power densities and allowing drive converters into environments which had hitherto been considered too hostile e.g. within motor windings or a jet engine. These devices therefore open up many new opportunities, but the fundamental principles of operation remain the same.

2.6.1 Bipolar junction transistor (BJT)

Historically the bipolar junction transistor was the first to be used for power switching. Of the two versions (npn and pnp) only the npn has been widely used in inverters for drives, mainly in applications ranging up to a few kW and several hundred volts.

The npn version is shown in Fig. 2.30: the main (load) current flows into the collector (C) and out of the emitter (E), as shown by the arrow on the device symbol. To switch the device ON (i.e. to make the resistance of the collector-emitter circuit low, so that load current can flow) a small current must be caused to flow from the base (B) to the emitter. When the base-emitter current is zero, the resistance of the collector-emitter circuit is very high, and the device is switched OFF.

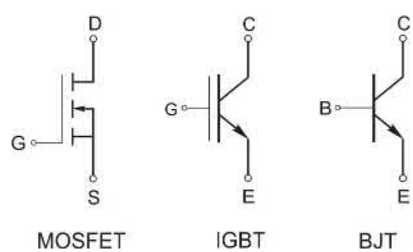

FIG. 2.30 Circuit symbols for switching devices.

The advantage of the bipolar transistor is that when it is turned ON, the collector-emitter voltage is low (see point b in Fig. 2.3) and hence the power dissipation is small in comparison with the load power, i.e. the device is an efficient power switch. The disadvantage is that although the power required in the base-emitter circuit is tiny in comparison with the load power, it is not insignificant and in the largest power transistors can amount to several tens of watts.

2.6.2 Metal oxide semiconductor field effect transistor (MOSFET)

Since the 1980s the power MOSFET has superseded the BJT in inverters for drives. Like the BJT, the MOSFET is a three-terminal device and is available in two versions, the n-channel and the p-channel. The n-channel is the most widely used, and is shown in Fig. 2.30. The main (load) current flows into the drain (D) and out of the source (S). (Confusingly, the load current in this case flows in the *opposite* direction to the arrow on the symbol.) Unlike the BJT, which is controlled by the base *current,* the MOSFET is controlled by the gate-source *voltage*.

To turn the device ON, the gate-source voltage must be comfortably above a threshold of a few volts. When the voltage is first applied to the gate, currents flow in the parasitic gate-source and gate-drain capacitances, but once these capacitances have been charged the input current to the gate is negligible, so the steady-state gate drive power is minimal. To turn the device OFF, the parasitic capacitances must be discharged and the gate-source voltage must be held below the threshold level.

The principal advantage of the MOSFET is that it is a voltage-controlled device which requires negligible power to hold it in the ON state. The gate drive circuitry is thus less complex and costly than the base drive circuitry of an equivalent bipolar device. The disadvantage of the MOSFET is that in the ON state the effective resistance of the drain-source is higher than an equivalent bipolar device, so the power dissipation is higher and the device is rather less efficient as a power switch. MOSFETs are used in low and medium power inverters up to a few kW, with voltages generally not exceeding 700 V.

2.6.3 Insulated gate bipolar transistor (IGBT)

The IGBT (Fig. 2.30) is a hybrid device which combines the best features of the MOSFET (i.e. ease of gate turn ON and turn OFF from low-power logic circuits) and the BJT (relatively low power dissipation in the main collector-emitter circuit). These obvious advantages give the IGBT the edge over the MOSFET and BJT, and account for their dominance in all but small drives. They are particularly well suited to the medium power, medium voltage range (up to several hundred kW). The path for the main (load) current is from collector to emitter, as in the npn bipolar device.

2.7 Converter waveforms, acoustic noise, and cooling

In common with most textbooks, the waveforms shown in this chapter (and later in the book) are what we would hope to see under ideal conditions. It makes sense to concentrate on these ideal waveforms from the point of view of gaining a basic understanding, but we ought to be warned that what we see on an oscilloscope may well look rather different, and is often not easy to interpret.

We have seen that the essence of power electronics is the switching process, so it should not come as much of a surprise to learn that in practice the switching is seldom achieved in such a clear-cut fashion as we have assumed. Usually, there will be some sort of high-frequency oscillation or 'ringing' evident, particularly on the voltage waveforms following each transition due to switching. This is due to the effects of stray capacitance and inductance: it will have been anticipated at the design stage, and steps will have been taken to minimise it by fitting 'snubbing' circuits at the appropriate places in the converter. However complete suppression of all these transient phenomena is seldom economically worthwhile so the user should not be too alarmed to see remnants of the transient phenomena in the output waveforms.

Acoustic noise is also a matter that can worry newcomers. Most power electronic converters emit whining or humming sounds at frequencies corresponding to the fundamental and harmonics of the switching frequency, though when the converter is used to feed a motor, the sound from the motor is usually a good deal louder than the sound from the converter itself. These sounds are difficult to describe in words, but typically range from a high-pitched hum through a whine to a piercing whistle. In a.c. drives it is often possible to vary the switching frequency, which may allow specific mechanical resonances to be avoided. If taken above the audible range (which is inversely proportional to age) the sound clearly disappears, but at the expense of increased switching losses, so a compromise has to be sought.

2.7.1 Cooling of switching devices—Thermal resistance

We have seen that by adopting a switching strategy the power loss in the switching devices is small in comparison with the power throughput, so the converter has a high efficiency. Nevertheless almost all the heat which is produced in the switching devices is released in the active region of the semiconductor, which is itself very small and will overheat and breakdown unless it is adequately cooled. It is therefore essential to ensure that even under the most onerous operating conditions, the temperature of the active switching junction inside the device does not exceed the safe value.

Consider what happens to the temperature of the junction region of the device when we start from cold (i.e. ambient) temperature and operate the device so that its average power dissipation remains constant. At first, the junction temperature begins to rise, so some of the heat generated is conducted to the metal case, which stores some heat as its temperature rises. Heat then flows into the heatsink (if fitted) which begins to warm up, and heat begins to flow to the surrounding air, at ambient temperature. The temperatures of the junction, case and heatsink continue to rise until eventually an equilibrium is reached when the total rate of loss of heat to ambient temperature is equal to the power dissipation inside the device.

The final steady-state junction temperature thus depends on how difficult it is for the heat to escape down the temperature gradient to ambient, or in other words on the total 'thermal resistance' between the junction inside the device and the surrounding medium (usually air). Thermal resistance is usually expressed in °C/W, which directly indicates how much temperature rise will occur in the steady state for every Watt of dissipated power. It follows that for a given power dissipation, the higher the thermal resistance, the higher the temperature rise, so in order to minimise the temperature rise of the device, the total thermal resistance between it and the surrounding air must be made as small as possible.

The device designer aims to minimise the thermal resistance between the semiconductor junction and the case of the device, and provides a large and flat metal mounting surface to minimise the thermal resistance between the case and the heatsink. The converter designer must ensure good thermal contact between the device and the heatsink, usually by a bolted joint smeared with heat-conducting compound to fill any microscopic voids, and must design the heatsink to minimise the thermal resistance to air. Heatsink design offers the only real scope for appreciably reducing the total resistance, and involves careful selection of the material, size, shape and orientation of the heatsink, and the associated air-moving system (see below).

In some applications, air cooled heatsinks are not the best choice, and alternatives such as liquid cooling (de-ionised water, glycol or oil) is used. An automotive application is a typical example of this, where a drive can be designed to take advantage of an existing liquid cooling system.

One drawback of the good thermal path between the semiconductor junction and the case of the device is that the metal mounting surface (or surfaces in the case of the popular high power hockey-puck package) can be electrically 'live'. This poses a difficulty for the converter designer, because mounting the device directly on the heatsink causes the latter to be potentially dangerous. In addition, several separate isolated heatsinks may be required in order to avoid short-circuits. The alternative is for the devices to be electrically isolated from the heatsink using thin mica spacers, but then the thermal resistance is increased appreciably.

Increasingly devices come in packaged 'modules' with an electrically isolated metal interface to get round the 'live' problem. The packages contain combinations of transistors, diodes or thyristors, from which various converter circuits can be built up. Several modules can be mounted on a single heatsink, which does not have to be isolated from the enclosure or cabinet. They are available in ratings suitable for converters up to hundreds of kW, and the range is expanding.

2.7.2 Arrangement of heatsinks and forced-air cooling

The principal factors which govern the thermal resistance of a heatsink are the total surface area, the condition of the surface and the air flow. Many converters use extruded aluminium heatsinks, with multiple fins to increase the effective cooling surface area and lower the resistance, and with a machined face or faces for mounting the devices. Heatsinks are usually mounted vertically to improve natural air convection. Surface finish is important, with black anodised aluminium being typically 30% better than bright. The cooling performance of heatsinks is however a complex technical area, with turbulence being very beneficial in forced air-cooled heatsinks.

A typical layout for a medium-power (say 200 kW) converter is shown in Fig. 2.31. The fan(s) are positioned either at the top or bottom of the heatsink, and draw external air upwards, assisting natural convection. Even a modest airflow is very beneficial: with an air velocity of only 2 m/s, for example, the thermal resistance is halved as compared with the naturally-cooled set-up, which means that for a given temperature rise the heatsink can be half the size of the naturally-cooled one. However, large increases in the air velocity bring diminishing returns and also introduce additional noise.

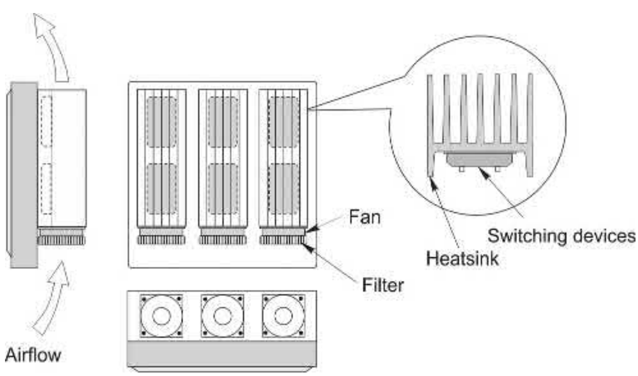

FIG. 2.31 Layout of converter showing heatsink and cooling fans.

2.8 Review questions

(1) In the circuit of Fig. Q1 both voltage sources and the diodes can be treated as ideal, and the load is a resistor. (Note: this question is specifically aimed at reinforcing the understanding of how diodes behave: it is not representative of any practical circuit.)

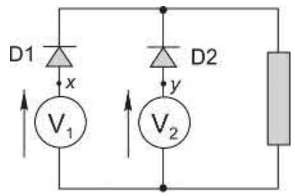

FIG. Q1

Sketch the voltage across the load under the following conditions:

(a) V_1 is a sinusoid of amplitude 20 V, and V_2 is a constant voltage of +10 V;

(b) V_1 is a sinusoid of amplitude 20 V, and V_2 is a constant voltage of −10 V;

(c) The same as (a), except that the diode D1 is placed below, rather above, V_1.

(2) What is the maximum d.c. voltage available from a fully-controlled bridge converter supplying a motor and operating from a low-impedance 230 V utility supply?

(3) Estimate the firing angle required to produce a mean output voltage of 300 V from a fully-controlled three-phase converter supplied from stiff 415 V, 50 Hz utility supply. Assume that the load current is continuous. How would the firing angle have to change if the supply frequency was 60 Hz rather than 50 Hz?

(4) Sketch the waveform of voltage across one of the thyristors in a fully-controlled single-phase converter with a firing angle delay of 60°. Assume that the d.c. load current is continuous. Fig. 2.10 may prove helpful.

(5) Sketch the current waveform in the a.c. supply when a single-phase fully-controlled converter with $\alpha = 45°$ is supplying a highly inductive load which draws a smooth current of 25 A. If the a.c. supply is 240 V, 50 Hz, and losses in the devices are neglected, calculate the peak and average supply power per cycle.

(6) A d.c. chopper circuit is often said to be like an a.c. transformer. Explain what this means by considering the input and output power relationships for a chopper-fed inductive motor load supplied with an average voltage of 20 V from a 100 V battery. Assume that the motor current remains constant throughout at 10 A.

(7) A 5 kHz step-down transistor chopper operating from a 150 V battery supplies an R/L load which draws an almost-constant current of 5 A. The resistance of the load is 8 Ω.

Treating all devices as ideal, estimate:

(i) the mark:space ratio of the chopper;

(ii) the average power in the load;

(iii) the average power from the source.

(8) This question relates to the switching circuit of Fig. Q8, and in particular to the function of the diode.

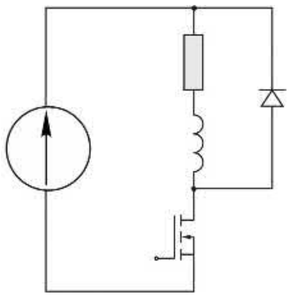

FIG. Q8

Some possible answers to the question 'what is the purpose of the diode' are given below:

- to prevent reverse current in the switch
- to protect the inductor from high voltages
- to limit the rate of change of current in the supply
- to limit the voltage across the MOSFET
- to dissipate the stored energy in the inductance

Discuss these answers and identify which one(s) are correct.

(9) In the circuit of Fig. Q8, assume that the supply voltage is 100 V, and that the forward volt-drop across the diode is 0.7 V. Some common answers to the question 'when the current is freewheeling, what is the voltage across the MOSFET' are given below:

- 99.3 V
- 0.7 V
- zero
- depends on the inductance
- 100.7 V

Discuss these answers and identify which one is correct.

Answers to the review questions are given in the Appendix.

Chapter 3

D.C. motors

3.1 Introduction

Until the 1980s the conventional (brushed) d.c. machine was the automatic choice where speed or torque control is called for, and large numbers remain in service despite a declining market share that reflects the general move to inverter-fed a.c. motors. D.C. motor drives do remain competitive in some larger ratings (several hundred kW) particularly where drip-proof motors are acceptable, with applications ranging from steel rolling mills, railway traction, through a very wide range of industrial drives.

Given the reduced importance of the d.c. motor, the reader may wonder why a whole chapter is devoted to it. The answer is that, despite its relatively complex construction, the d.c. machine is relatively simple to understand, not least because of the clear physical distinction between its separate 'flux' and 'torque producing' parts. We will find that its performance can be predicted with the aid of a simple equivalent circuit, and that many aspects of its behaviour are reflected in other types of motor, where it may be more difficult to identify the sources of flux and torque. The d.c. motor is therefore an ideal learning vehicle, and time spent assimilating the material in this chapter should therefore be richly rewarded later.

Over a very wide power range from several megawatts at the top end down to a only a few watts, all d.c. machines have the same basic structure, as shown in Fig. 3.1.

The motor has two separate electrical circuits. The smaller pair of terminals (designated E1, E2, with E for excitation—see later Fig. 3.6) connect to the field windings, which surround each pole and are normally in series: these windings provide the m.m.f. to set up the flux in the air-gap under the poles. In the steady state all the input power to the field windings is dissipated as heat—none of it is converted to mechanical output power.

The main terminals (designated A1, A2, with A for 'armature') convey the 'torque-producing' or 'work' current to the brushes which make sliding contact (via the commutator) with the conductors that form the so-called 'armature' winding on the rotor. The supply to the field (the flux-producing part of the

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00003-9

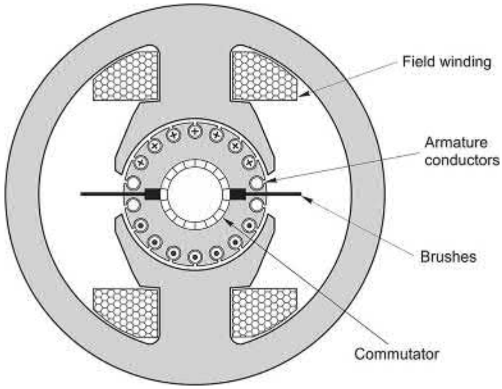

FIG. 3.1 Conventional (brushed) d.c. motor.

motor) is separate from that for the armature, hence the description 'separately excited'.

As in any electrical machine it is possible to design a d.c. motor for any desired supply voltage, but for several reasons it is unusual to find rated voltages lower than about 6 V or much higher than 700 V. The lower limit arises because the brushes (see below) give rise to an unavoidable volt-drop of perhaps 0.5–1 V, and it is clearly not good practice to let this 'wasted' voltage became a large fraction of the supply voltage. At the other end of the scale it becomes prohibitively expensive to insulate the commutator segments to withstand higher voltages. The function and operation of the commutator is discussed later, but it is appropriate to mention here that brushes and commutators are troublesome at very high speeds. Small d.c. motors, say up to hundreds of Watts output, can run at perhaps 12,000 rev/min, but the majority of medium and large motors are usually designed for speeds below 3000 rev/min.

Motors are usually supplied with power electronic drives, which draw power from the a.c. utility supply and convert it to d.c. for the motor. Since the utility voltages tend to be standardised (e.g. 110 V, 200–240 V, or 380–480 V, 50 or 60 Hz), motors are made with rated voltages which match the range of d.c. outputs from the converter (see Chapter 2).

As mentioned above, it is quite normal for a motor of a given power, speed and size to be available in a range of different voltages. In principle all that has to be done is to alter the number of turns and the size of wire making up the coils in the machine. A 12 V, 4 A motor, for example, could easily be made to operate from 24 V instead, by winding its coils with twice as many turns of wire having only half the cross-sectional area of the original. The full speed would be the same at 24 V as the original was at 12 V, and the rated current would be 2 A, rather than 4 A. The input power and output power would be unchanged, and externally there would be no change in appearance, except that the terminals might be a bit smaller.

Traditionally d.c. motors were classified as shunt, series, or separately excited. In addition it was common to see motors referred to as 'compound-wound'. These descriptions date from the period before the advent of power electronics: they reflect the way in which the field and armature circuits are interconnected, which in turn determines the operating characteristics. For example, the series motor has a high starting torque when switched directly on line, so it became the natural choice for traction applications, while applications requiring constant speed would use the shunt connected motor.

However, at the fundamental level there is really no difference between the various types, so we focus attention on the separately-excited machine, before taking a brief look at shunt and series motors. Later, in Chapter 4, we will see how the operating characteristics of the separately-excited machine with power electronic supplies equip it to suit any application, and thereby displace the various historic predecessors.

We should make clear at this point that whereas in an a.c. machine the number of poles is of prime importance in determining the speed, the pole-number in a d.c. machine is of little consequence as far as the user is concerned. It turns out to be more economical to use two or four poles (perhaps with a square stator frame) in small or medium size d.c. motors, and more (e.g. 10 or 12 or even more) in large ones, but the only difference to the user is that the two-pole type will have two brushes displaced by 180°, the four-pole will have four brushes displaced by 90°, and so on. Most of our discussion centres on the two-pole version in the interests of simplicity, but there is no essential difference as far as operating characteristics are concerned.

Finally, before we move on, we should point out that confusion can be caused when people describe a type of a.c. motor which is fed from a power electronic converter with a three phase square wave voltage waveform as a 'Brushless d.c. motor'. We suggest that the reader should forget about that at the moment—we will pick it up in Chapter 9.

3.2 Torque production

Torque is produced by interaction between the axial current-carrying conductors on the rotor and the radial magnetic flux produced by the stator. The flux or 'excitation' can be furnished by permanent magnets (Fig. 3.2A) or by means of field windings (Figs. 3.1 and 3.2B).

Permanent magnet versions are available in motors with outputs from a few watts up to a few kilowatts, while wound-field machines begin at about 100 watts and extend to the largest (MW) outputs. The advantages of the permanent magnet type are that no electrical supply is required for the field, and the overall size of the motor can be smaller. On the other hand the strength of the field cannot be varied so one possible option for control is ruled out.

Ferrite magnets have been used for many years. They are relatively cheap and easy to manufacture but their energy product (a measure of their

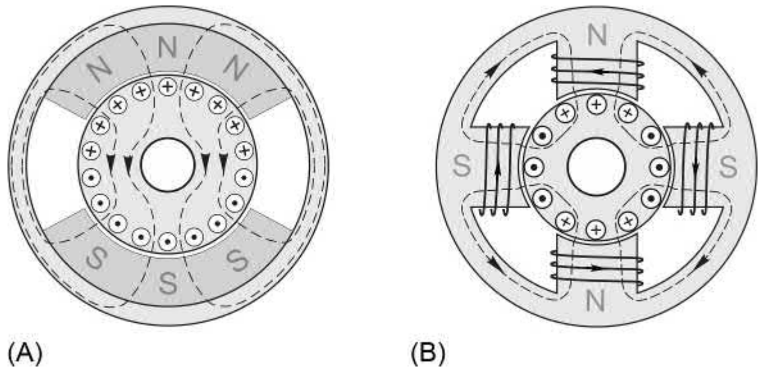

FIG. 3.2 Excitation (field) systems for d.c. motors (A) two-pole permanent magnet; (B) four-pole wound field.

effectiveness as a source of excitation) is poor. Rare earth magnets (e.g. Neodynium-iron-boron or Samarium cobalt) provide much higher energy products, and yield high torque/volume ratios: they are used in high performance servo motors, but are relatively expensive and difficult to manufacture and handle. Nd-Fe-B magnets have the highest energy product but have a modest Curie point, above which demagnetisation occurs, which may be a consideration for some demanding applications.

Although the magnetic field is essential to the operation of the motor, we should recall that in Chapter 1 we saw that none of the mechanical output power actually comes from the field system. The excitation acts like a catalyst in a chemical reaction, making the energy conversion possible but not contributing to the output.

The main (power) circuit consists of a set of identical coils wound in slots on the rotor, and known as the armature. Current is fed into and out of the rotor via carbon 'brushes' which make sliding contact with the commutator, which consists of insulated copper segments mounted on a cylindrical former. (The term 'brush' stems from the early attempts to make sliding contacts using bundles of wires bound together in much the same way as the willow twigs in a witch's broomstick. Not surprisingly these primitive brushes soon wore grooves in the commutator.)

The function of the commutator is discussed below, but it is worth stressing here that all the electrical energy which is to be converted into mechanical output has to be fed into the motor through the brushes and commutator. Given that a high-speed sliding electrical contact is involved, it is not surprising that to ensure trouble-free operation the commutator needs to be kept clean, and the brushes and their associated springs need to be regularly serviced. Brushes wear away, of course, though if the correct 'grade' of brush is used and they are properly 'bedded in' they can last for many thousands of hours. As a very broad rule of thumb, a commutator brush could be expected to wear at a rate of 3–4000 h per centimetre. All being well, the brush debris (in the form of graphite particles) will be carried out of harm's way by the ventilating air: any build up of dust on the insulation of the windings of a high-voltage motor risks the danger of short-circuits, while debris on the commutator itself is dangerous and can lead to disastrous 'flashover' faults.

The axial length of the commutator depends on the current it has to handle. Small motors usually have one brush on each side of the commutator, so the commutator is quite short, but larger heavy-current motors may well have many brushes mounted on a common arm, each with its own brushbox (in which it is free to slide) and with all the brushes on one arm connected in parallel via their flexible copper leads or 'pigtails'. The length of the commutator can then be comparable with the 'active' length of the armature (i.e. the part carrying the conductors exposed to the radial flux).

3.2.1 Function of the commutator

Many different winding arrangements are used for d.c. armatures, and it is neither helpful or necessary for us to delve into the nitty-gritty of winding and commutator design. These are matters which are best left to motor designers and repairers. What we need to do is to focus on what a well designed commutator winding actually achieves, and despite the apparent complexity, this can be stated quite simply.

The purpose of the commutator is to ensure that regardless of the position of the rotor, the pattern of current flow in the rotor is always as shown in Fig. 3.3.

Current enters the rotor via one brush, flows through the rotor coils in the directions shown in Fig. 3.3, and leaves via the other brush. The first point of contact with the armature is via the commutator segment or segments on which the brush is pressing at the time (the brush is usually wider than a single segment), but since the interconnections between the individual coils are made at each commutator segment, the current actually passes through all the coils via all the commutator segments in its path through the armature.

We can see from Fig. 3.3 that all the conductors lying under the N pole carry current in one direction, while all those under the S pole carry current in the opposite direction. All the conductors under the N pole will therefore experience a downward force (which is proportional to the radial flux density B and the armature current I) while all the conductors under the S pole will

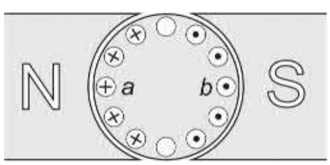

FIG. 3.3 Pattern of rotor (armature) currents in two-pole d.c. motor.

experience an equal upward force—remember Flemings left hand rule—Fig. 1.4. A torque is thus produced on the rotor, the magnitude of the torque being proportional to the product of the flux density and the current. In practice the flux density will not be completely uniform under the pole, so the force on some of the armature conductors will be greater than on others. However, it is straightforward to show that the total torque developed is given by

$$T = K_T \Phi I \tag{3.1}$$

where Φ is the total flux produced by the field, and K_T is constant for a given motor. In the majority of motors the flux remains constant, so we see that the motor torque is directly proportional to the armature current. This extremely simple result means that if a motor is required to produce constant torque at all speeds, we simply have to arrange to keep the armature current constant. Standard drive packages usually include provision for doing this, as will be seen later. We can also see from Eq. (3.1) that the direction of the torque can be reversed by reversing either the armature current (I) or the flux (Φ). We obviously make use of this when we want the motor to run in reverse, and sometimes when we want regenerative braking.

The alert reader might rightly challenge the claim—made above—that the torque will be constant regardless of rotor position. Looking at Fig. 3.3, it should be clear that if the rotor turned just a few degrees, one of the five conductors shown as being under the pole will move out into the region where there is no radial flux, before the next one moves under the pole. Instead of five conductors producing force, there will then be only four, so won't the torque be reduced accordingly?

The answer to this question is yes, and it is to limit this unwelcome variation of torque that most motors have many more coils than are shown in Fig. 3.3. Smooth torque is of course desirable in most applications in order to avoid vibrations and resonances in the transmission and load, and is essential in machine tool drives where the quality of finish can be marred by uneven cutting if the torque and speed are not steady.

Broadly speaking the higher the number of coils (and commutator segments) the better, because the ideal armature would be one in which the pattern of current on the rotor corresponded to a 'current sheet', rather than a series of discrete packets of current. If the number of coils was infinite, the rotor would look identical at every position, and the torque would therefore be absolutely smooth. Obviously this is not practicable, but it is closely approximated in most d.c. motors. For practical and economic reasons the number of slots is higher in large motors, which may well have a hundred or more coils and hence very little ripple in their output torque.

3.2.2 Operation of the commutator—interpoles

Returning now to the operation of the commutator, and focusing on a particular coil (e.g. the one shown as *ab* in Fig. 3.3) we note that for half a revolution—while side *a* is under the N pole and side *b* is under the S pole, the current needs to be positive in side *a* and negative in side *b* in order to produce a positive torque. For the other half revolution, while side *a* is under the S pole and side *b* is under the N pole, the current must flow in the opposite direction through the coil for it to continue to produce positive torque. This reversal of current takes place in each coil as it passes through the interpolar axis, the coil being 'switched-round' by the action of the commutator sliding under the brush. Each time a coil reaches this position it is said to be undergoing commutation, and the relevant coil in Fig. 3.3 has therefore been shown as having no current to indicate that its current is in the process of changing from positive to negative.

The essence of the current-reversal mechanism is revealed by the simplified sketch shown in Fig. 3.4. This diagram shows a single coil fed via the commutator and brushes with current that always flows in at the top brush.

In the left-hand sketch, coil-side *a* is under the N pole and carries positive current because it is connected to the shaded commutator segment which in turn is fed from the top brush. Side *a* is therefore exposed to a flux density directed from left (N) to right (S) in the sketch, and will therefore experience a downward force. This force will remain constant while the coil-side remains under the N pole. Conversely, side *b* has negative current but it also lies in a flux density directed from right to left, so it experiences an upward force. There is thus an anti-clockwise torque on the rotor.

When the rotor turns to the position shown in the sketch on the right, the current in both sides is reversed, because side *b* is now fed with positive current via the unshaded commutator segment. The direction of force on each coil side is reversed, which is exactly what we want in order for the torque to remain clockwise. Apart from the short period when the coil is outside the influence of the flux, and undergoing commutation (current-reversal) the torque is constant.

It should be stressed that the discussion above is intended to illustrate the principle involved, and the sketch should not be taken too literally. In a real multi-coil armature, the commutator arc is much smaller than that shown in Fig. 3.4 and only one of the many coils coil is reversed at a time, so the torque remains very nearly constant regardless of the position of the rotor.

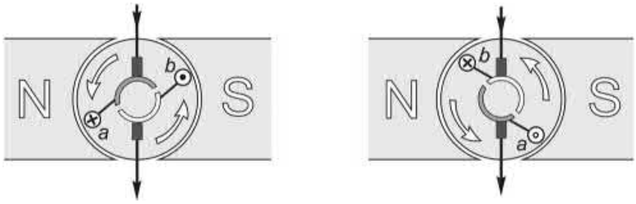

FIG. 3.4 Simplified diagram of single-coil motor to illustrate the current-reversing function of the commutator.

The main difficulty in achieving good commutation arises because of the self-inductance of the armature coils, and the associated stored energy. As we have seen earlier, inductive circuits tend to resist change of current, and if the current reversal has not been fully completed by the time the brush slides off the commutator segment in question there will be a spark at the trailing edge of the brush.

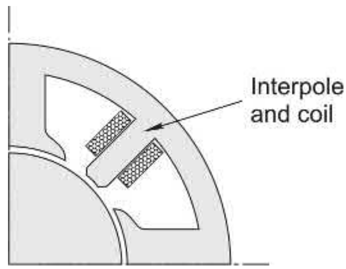

FIG. 3.5 Sketch showing location of interpole and interpole winding. (The main field windings have been omitted for the sake of clarity.)

In small motors some sparking is considered tolerable, but in medium and large wound-field motors small additional stator poles known as interpoles (or compoles) are provided to improve commutation and hence minimise sparking. These extra poles are located midway between the main field poles, as shown in Fig. 3.5. Interpoles are not normally required in permanent magnet motors because the absence of stator iron close to the rotor coils results in much lower armature coil inductance.

The purpose of the interpoles is to induce a motional e.m.f. in the coil undergoing commutation, in such a direction as to speed-up the desired reversal of current, and thereby prevent sparking. The e.m.f. needed is proportional to the current which has to be commutated, i.e. the armature current, and to the speed of rotation. The correct e.m.f. is therefore achieved by passing the armature current through the coils on the interpoles, thereby making the flux from the interpoles proportional to the armature current. The interpole coils therefore consist of a few turns of thick conductor, connected permanently in series with the armature.

3.3 Motional e.m.f.

Readers who have skipped Chapter 1 are advised to check that they are familiar with the material covered in Section 1.7 before reading the rest of this chapter, as not all of the lessons drawn in Chapter 1 are repeated explicitly here.

When the armature is stationary, no motional e.m.f. is induced in it. But when the rotor turns, the armature conductors cut the radial magnetic flux and an e.m.f. is induced in them.

As far as each individual coil on the armature is concerned, an alternating e.m.f. will be induced in it when the rotor rotates. For the coil *ab* in Fig. 3.3, for example, side *a* will be moving upward through the flux if the rotation is clockwise, and an e.m.f. directed out of the plane of the paper will be generated. At the same time the 'return' side of the coil (*b*) will be moving downwards, so the same magnitude of e.m.f. will be generated, but directed into the paper. The resultant e.m.f. in the coil will therefore be twice that in the coil-side, and this e.m.f. will remain constant for almost half a revolution, during which time the coil sides are cutting a constant flux density. For the comparatively short time when the coil is not cutting any flux the e.m.f. will be zero, and then the coil will begin to cut through the flux again, but now each side is under the other pole, so the e.m.f. is in the opposite direction. The resultant e.m.f. waveform in each coil is therefore a rectangular alternating wave, with magnitude and frequency proportional to the speed of rotation.

The coils on the rotor are connected in series, so if we were to look at the e.m.f. across any given pair of diametrically opposite commutator segments, we would see a large alternating e.m.f. (We would have to station ourselves on the rotor to do this, or else make sliding contacts using slip-rings).

The fact that the induced voltage in the rotor is alternating may come as a surprise, since we are talking about a d.c. motor rather than an a.c. one. But any worries we may have should be dispelled when we ask what we will see by way of induced e.m.f. when we 'look in' at the brushes. We will see that the brushes and commutator effect a remarkable transformation, bringing us back into the reassuring world of d.c.

The first point to note is that the brushes are stationary. This means that although a particular segment under each brush is continually being replaced by its neighbour, the circuit lying between the two brushes always consists of the same number of coils, with the same orientation with respect to the poles. As a result the e.m.f. at the brushes is direct (i.e. constant), rather than alternating.

The magnitude of the e.m.f. depends on the position of the brushes around the commutator, but they are invariably placed at the point where they continually 'see' the peak value of the alternating e.m.f. induced in the armature. In effect, the commutator and brushes can be regarded as a mechanical rectifier which converts the alternating e.m.f. in the rotating reference frame to a direct e.m.f. in the stationary reference frame. It is a remarkably clever and effective device, its only real drawback being that it is a mechanical system, and therefore subject to wear and tear.

We saw earlier that to obtain smooth torque it was necessary for there to be a large number of coils and commutator segments, and we find that much the same considerations apply to the smoothness of the generated e.m.f. If there are only a few armature coils the e.m.f. will have a noticeable ripple superimposed on the mean d.c. level. The higher we make the number of coils, the smaller the ripple, and the better the d.c. we produce. The small ripple we inevitably get with a finite number of segments is seldom any problem with motors used in drives, but can sometimes give rise to difficulties when a d.c. machine is used to provide a speed feedback signal in a closed-loop system (see Chapter 4).

In Chapter 1 we saw that when a conductor of length l moves at velocity v through a flux density B, the motional e.m.f. induced is given by $e = Blv$. In the complete machine we have many series-connected conductors; the linear velocity (v) of the primitive machine examined in Chapter 1 is replaced by the tangential velocity of the rotor conductors, which is clearly proportional to the speed of rotation (n); and the average flux density cut by each conductor (B) is directly related to the total flux (Φ). If we roll together the other influential factors (number of conductors, radius, active length of rotor) into a single constant (K_E), it follows that the magnitude of the resultant e.m.f. (E) which is generated at the brushes is given by

$$E = K_E \Phi n \tag{3.2}$$

This equation reminds us of the key role of the flux, in that until we switch on the field no voltage will be generated, no matter how fast the rotor turns. Once the field is energised, the generated voltage is directly proportional to the speed of rotation, so if we reverse the direction of rotation, we will also reverse the polarity of the generated e.m.f. We should also remember that the e.m.f. depends only on the flux and the speed, and is the same regardless of whether the rotation is provided by some external source (i.e. when the machine is being driven as a generator) or when the rotation is produced by the machine itself (i.e. when it is acting as a motor).

It has already been mentioned that the flux is usually constant at its full value, in which case Eqs (3.1) and (3.2) can be written in the form

$$T = k_t I \tag{3.3}$$

$$E = k_e \omega \tag{3.4}$$

where k_t is the motor torque constant, k_e is the e.m.f. constant, and ω is the angular speed in rad/s.

In this book the international standard (SI) system of units is used throughout. In the SI system, the units for k_t are the units of torque (Newton metre) divided by the unit of current (Ampere), i.e. Nm/A; and the units of k_e are Volts/rad/s. (Note, however, that k_e is more often given in Volts/1000 rev/min.)

It is not at all clear that the units for the torque constant (Nm/A) and the e.m.f. constant (V/rad/s), which on the face of it measure very different physical phenomena, are in fact the same, i.e. 1 Nm/A = 1 Volt/rad/s. Some readers will be content simply to accept it, others may be puzzled, a few may even find it obvious. Those who are surprised and puzzled may feel more comfortable by progressively replacing one set of units by their equivalent, to lead us in the direction of the other, e.g.

$$\frac{(Newton)(metre)}{Amp} = \frac{Joule}{Amp} = \frac{(Watt)(sec)}{Amp} = \frac{(Volt)(Amp)(sec)}{Amp} = (Volt)(sec)$$

This still leaves us to ponder what happened to the 'radians' in k_e, but at least the underlying unity is demonstrated, and after all a radian is a dimensionless quantity. Delving deeper, we note that 1 Volt.second = 1 Weber, the unit of magnetic flux. This is hardly surprising because the production of torque and the generation of motional e.m.f. are both brought about by the catalytic action of the magnetic flux.

Returning to more pragmatic issues, we have now discovered the extremely convenient fact that in SI units, the torque and e.m.f. constants are equal, i.e. $k_t = k_e = k$. The torque and e.m.f. equations can thus be further simplified as

$$T = kI \tag{3.5}$$

$$E = k\omega \tag{3.6}$$

We will make use of these two delightfully simple equations time and again in the subsequent discussion. Together with the armature voltage equation (see below), they allow us to predict all aspects of behaviour of a d.c. motor. There can be few such versatile machines for which the fundamentals can be expressed so simply.

Though attention has been focused on the motional e.m.f. in the conductors, we must not overlook the fact that motional e.m.f.s are also induced in the body of the rotor. If we consider a rotor tooth, for example, it should be clear that it will have an alternating e.m.f. induced in it as the rotor turns, in just the same way as the e.m.f. induced in the adjacent conductor. In the machine shown in Fig. 3.1, for example, when the e.m.f. in a tooth under a N pole is positive, the e.m.f. in the diametrically opposite tooth (under a S pole) will be negative. Given that the rotor steel conducts electricity, these e.m.f.s will tend to set up circulating currents in the body of the rotor, so to prevent this happening, the rotor is made not from a solid piece but from thin steel laminations (typically less than 1 mm thick) which have an insulated coating to prevent the flow of unwanted currents. If the rotor was not laminated the induced current would not only produce large quantities of waste heat, but also exert a substantial braking torque.

3.3.1 Equivalent circuit

The equivalent circuit can now be drawn on the same basis as we used for the primitive machine in Chapter 1, and is shown in Fig. 3.6.

The voltage V is the voltage applied to the armature terminals (i.e. across the brushes), and E is the internally developed motional e.m.f. The resistance and inductance of the complete armature are represented by R and L in Fig. 3.6. The sign convention adopted is the usual one when the machine is operating as a motor. Under motoring conditions, the motional e.m.f. E always opposes the applied voltage V, and for this reason it is referred to as 'back e.m.f.'. For current to be forced into the motor, V must be greater than E, the armature circuit voltage equation being given by

$$V = E + IR + L\frac{dI}{dt} \tag{3.7}$$

The last term in Eq. (3.7) represents the inductive volt-drop due to the armature self-inductance. This voltage is proportional to the rate of change of current, so under steady-state conditions (when the current is constant), the term will be zero and can be ignored. We will see later that the armature inductance has an unwelcome effect under transient conditions, but is also very beneficial in smoothing the current waveform when the motor is supplied by a controlled rectifier.

3.4 D.C. motor—steady-state characteristics

From the user's viewpoint the extent to which speed falls when load is applied, and the variation in speed with applied voltage are usually the first questions which need to be answered in order to assess the suitability of the motor for the job in hand. The information is usually conveyed in the form of the steady-state characteristics, which indicate how the motor behaves when any transient effects (caused for example by a sudden change in the load) have died away and conditions have once again become steady. Steady-state characteristics are usually much easier to predict than transient characteristics, and for the d.c. machine they can all be deduced from the simple equivalent circuit in Fig. 3.6.

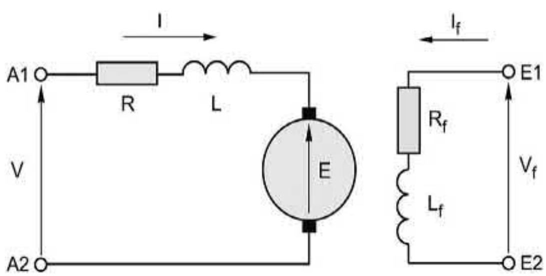

FIG. 3.6 Equivalent circuit of a d.c. motor.

Under steady conditions, the armature current I is constant and Eq. (3.7) simplifies to

$$V = E + IR, \text{ or } I = \frac{V - E}{R} \tag{3.8}$$

This equation allows us to find the current if we know the applied voltage, the speed (from which we get E via Eq. 3.6) and the armature resistance, and we can then obtain the torque from Eq. (3.5). Alternatively, we may begin with torque and speed, and work out what voltage will be needed.

We will derive the steady-state torque-speed characteristics for any given armature voltage V shortly, but first we begin by establishing the relationship between the no-load speed and the armature voltage, since this is the foundation on which the speed control philosophy is based.

3.4.1 No-load speed

By 'no-load' we mean that the motor is running light, so that the only mechanical resistance is that due to its own friction. In any sensible motor the frictional torque will be small, and only a small driving torque will therefore be needed to keep the motor running. Since motor torque is proportional to current (Eq. 3.5), the no-load current will also be small. If we assume that the no-load current is in fact zero, the calculation of no-load speed becomes very simple. We note from Eq. (3.8) that zero current implies that the back e.m.f. is equal to the applied voltage, while Eq. (3.2) shows that the back e.m.f. is proportional to speed. Hence under true no-load (zero torque) conditions, we obtain

$$V = E = K_E \Phi n, \ \ or \ \ n = \frac{V}{K_E \Phi} \tag{3.9}$$

where n is the speed. (We have used Eq. (3.8) for the e.m.f., rather than the simpler Eq. (3.4) because the latter only applies when the flux is at its full value, and in the present context it is important for us to see what happens when the flux is reduced.)

At this stage we are concentrating on the steady-state running speeds, but we are bound to wonder how it is that the motor reaches speed from rest. We will return to this when we look at transient behaviour, so for the moment it is sufficient to recall that we came across an equation identical to Eq. (3.9) when we looked at the primitive linear motor in Chapter 1. We saw that if there was no resisting force opposing the motion, the speed would rise until the back e.m.f. equalled the supply voltage. The same result clearly applies to the frictionless and unloaded d.c. motor here.

We see from Eq. (3.9) that the no-load speed is directly proportional to armature voltage, and inversely proportional to field flux. For the moment we will continue to consider the case where the flux is constant, and demonstrate by means of an example that the approximations used in arriving at Eq. (3.9) are justified in practice. Later, we can use the same example to study the torque-speed characteristic.

3.4.2 Performance calculation—example

Consider a 500 V, 9.1 kW, 20 A, permanent-magnet motor with an armature resistance of 1 Ω. (These values tell us that the normal operating voltage is 500 V, the current when the motor is fully loaded is 20 A, and the mechanical output power under these full-load conditions is 9.1 kW.) When supplied at 500 V, the unloaded motor is found to run at 1040 rev/min, drawing a current of 0.8 A.

Whenever the motor is running at a steady speed, the torque it produces must be equal (and opposite) to the total opposing or load torque: if the motor torque was less than the load torque, it would decelerate, and if the motor torque was higher than the load torque it would accelerate. From Eq. (3.3), we see that the motor torque is determined by its current, so we can make the important statement that, in the steady-state, the motor current will be determined by the mechanical load torque. When we make use of the equivalent circuit (Fig. 3.6) under steady-state conditions we will need to get used to the idea that the current is determined by the load torque—i.e. one of the principal 'inputs' which will allow us to solve the circuit equations is the mechanical load torque, which is not even shown on the diagram. For those who are not used to electromechanical interactions this can be a source of difficulty.

Returning to our example, we note that because it is a real motor, it draws a small current (and therefore produces some torque) even when unloaded. The fact that it needs to produce torque, even though no load torque has been applied and it is not accelerating, is attributable to the inevitable friction in the cooling fan, bearings and brushgear.

If we want to estimate the no-load speed at a different armature voltage (say 250 V), we would ignore the small no-load current and use Eq. (3.9), giving

$$\text{No-load speed at } 250\,\text{V} = (250/500) \times 1040 = 520\,\text{rev/min}$$

Since Eq. (3.9) is based on the assumption that the no-load current is zero, this result is only approximate.

If we insist on being more precise, we must first calculate the original value of the back e.m.f., using Eq. (3.8), which gives

$$E = 500 - (0.8 \times 1) = 499.2\,\text{Volts}$$

As expected the back e.m.f. is almost equal to the applied voltage. The corresponding speed is 1040 rev/min, so the e.m.f. constant must be 499.2/1040 or 480 Volts/1000 rev/min. To calculate the no-load speed for V = 250 Volts, we first need to know the current. We are not told anything about how the friction torque varies with speed so all we can do is to assume that the friction torque is constant, in which case the motor current will be 0.8 A regardless of speed. With this assumption, the back e.m.f. will be given by

$$E = 250 - 0.8 \times 1 = 249.2\,\text{V}$$

And hence the speed will be given by

$$\text{No-load speed at } 250\,\text{V} = \frac{249.2}{480} \times 1000 = 519.2\,\text{rev/min}$$

The difference between the approximate and true no-load speeds is very small, and is unlikely to be significant. Hence we can safely use Eq. (3.9) to predict the no-load speed at any armature voltage, and obtain the set of no-load speeds shown in Fig. 3.7.

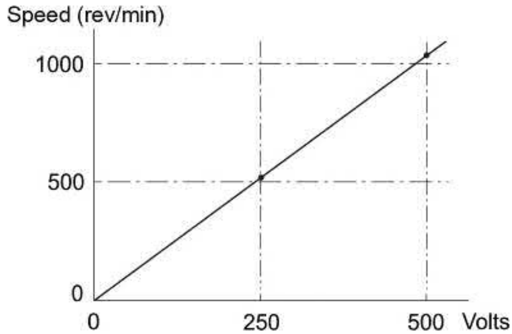

FIG. 3.7 No-load speed of d.c. motor as a function of armature voltage.

This diagram illustrates the very simple linear relationship between the speed of an unloaded d.c. motor and the armature voltage.

3.4.3 Behaviour when loaded

Having seen that the no-load speed of the motor is directly proportional to the armature voltage, we need to explore how the speed will vary when we change the load on the shaft.

The usual way we quantify 'load' is to specify the torque needed to drive the load at a particular speed. Some loads, such as a simple drum-type hoist with a constant weight on the hook, require the same torque regardless of speed, but for most loads the torque needed varies with the speed. For a fan, for example, the torque needed varies roughly with the square of the speed. If we know the torque/speed characteristic of the load, and the torque/speed characteristic of the motor (see below), we can find the steady-state speed simply by finding the intersection of the two curves in the torque-speed plane. An example (not specific to a d.c. motor) is shown in Fig. 3.8.

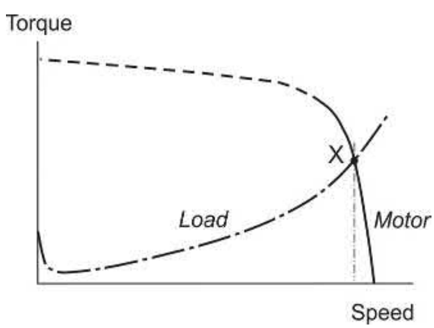

FIG. 3.8 Steady-state torque-speed curves for motor and load showing location (X) of steady-state operating condition.

At point X the torque produced by the motor is exactly equal to the torque needed to keep the load turning, so the motor and load are in equilibrium and the speed remains steady. At all lower speeds, the motor torque will be higher than the load torque, so the net torque will be positive, leading to an acceleration of the motor. As the speed rises towards X the acceleration reduces until the speed stabilises at X. Conversely, at speeds above X the motor's driving torque is less than the braking torque exerted by the load, so the net torque is negative and the system will decelerate until it reaches equilibrium at X. This example is one which is inherently stable, so that if the speed is disturbed for some reason from the point X, it will always return there when the disturbance is removed.

Turning now to the derivation of the torque/speed characteristics of the d.c. motor, we can profitably extend the previous example to illustrate matters. We can obtain the full-load speed for V = 500 Volts by first calculating the back e.m.f. at full load (i.e. when the current is 20 A). From Eq. (3.8) we obtain

$$E = 500 - 20 \times 1 = 480\,\text{Volts}$$

We have already seen that the e.m.f. constant is 480 Volts/1000 rev/min, so the full load speed is clearly 1000 rev/min. From no-load to full-load the speed falls linearly, giving the torque-speed curve for V = 500 Volts shown in Fig. 3.9. Note that from no-load to full-load the speed falls from 1040 rev/min to 1000 rev/min, a drop of only 4%. Over the same range the back e.m.f. falls from very nearly 500 Volts to 480 Volts, which of course also represents a drop of 4%.

We can check the power balance using the same approach as in Section 1.7 of Chapter 1. At full load the electrical input power is given by VI, i.e. 500 V × 20 A = 10 kW. The power loss in the armature resistance is $I^2 R = 400 \times 1 = 400$ W. The power converted from electrical to mechanical form is given by EI, i.e. 480 V × 20 A = 9600 W. We can see from the no-load data that the power required to overcome friction and iron losses (eddy currents and hysteresis, mainly in the rotor) at no-load is approximately 500 V × 0.8 A = 400 W, so this leaves about 9.2 kW. The rated output power

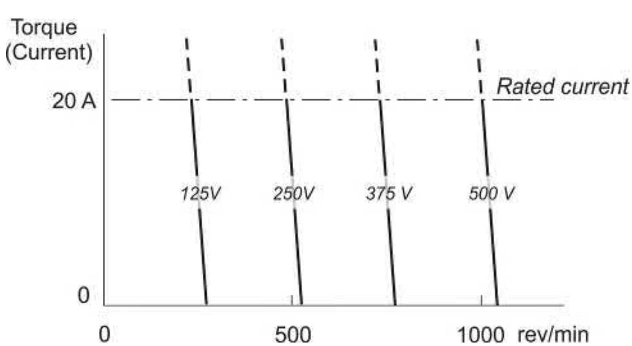

FIG. 3.9 Family of steady-state torque-speed curves for a range of armature voltages.

of 9.1 kW indicates that 100 W of additional losses (which we will not attempt to explore here) can be expected under full-load conditions.

Two important observations follow from these calculations. Firstly, the drop in speed when load is applied (the 'droop') is very small. This is very desirable for most applications, since all we have to do to maintain almost constant speed is to set the appropriate armature voltage and keep it constant. Secondly, a delicate balance between V and E is revealed. The current is in fact proportional to the difference between V and E (Eq. 3.8), so that quite small changes in either V or E give rise to disproportionately large changes in the current. In the example, a 4% reduction in E causes the current to rise to its rated value. Hence to avoid excessive currents (which cannot be tolerated in a thyristor supply, for example), the difference between V and E must be limited. This point will be taken up again when transient performance is explored.

A representative family of torque-speed characteristics for the motor discussed above is shown in Fig. 3.9. As already explained, the no-load speeds are directly proportional to the applied voltage, while the slope of each curve is the same, being determined by the armature resistance: the smaller the resistance the less the speed falls with load. These operating characteristics are very attractive because the speed can be set simply by applying the correct voltage.

The upper region of each characteristic in Fig. 3.9 is shown dotted because in this region the armature current is above its rated value, and therefore the motor cannot be operated continuously without overheating. Motors can and do operate for short periods above rated current, and the fact that the d.c. machine can continue to provide torque in proportion to current well into the overload region makes it particularly well-suited to applications requiring the occasional boost of excess torque.

A cooling problem might be expected when motors are run continuously at full current (i.e. full torque) even at very low speed, where the natural ventilation is poor. This operating condition is considered quite normal in converter-fed motor drive systems, and motors are accordingly fitted with a small air-blower motor as standard.

This book is about motors, which convert electrical power into mechanical power. But, in common with all electrical machines, the d.c. motor is inherently capable of operating as a generator, converting mechanical power into electrical power. And although the overwhelming majority of d.c. machines will spend most of their time in motoring mode, there are applications such as rolling mills where frequent reversal is called for, and others where rapid braking is required. In the former, the motor is controlled so that it returns the stored kinetic energy to the supply system each time the rolls have to be reversed, while in the latter case the energy may also be returned to the supply, or dumped as heat in a resistor. These transient modes of operation during which the machine acts as a generator may better be described as 'regeneration' since they only involve recovery of mechanical energy that was originally provided by the motor.

Continuous generation is of course possible using a d.c. machine provided we have a source of mechanical power, such as an internal combustion engine. In the example discussed above we saw that when connected to a 500 V supply, the unloaded machine ran at 1040 rev/min, at which point the back e.m.f. was very nearly 500 V and only a tiny positive current was flowing. As we applied mechanical load to the shaft the steady-state speed fell, thereby reducing the back e.m.f. and increasing the armature current until the motor torque was equal to the opposing load torque and equilibrium returned.

Conversely, if instead of applying an opposing (load) torque, we use the IC engine to supply torque in the opposite direction, i.e. trying to increase the speed of the motor, the increase in speed will cause the motional e.m.f. to be greater than the supply voltage (500 V). This means that the current will flow from the d.c. machine to the supply, resulting in a negative torque and reversal of electrical power flow back into the supply. Stable generating conditions will be achieved when the motor torque (current) is equal and opposite to the torque provided by the IC engine. In the example, the full-load current is 20 A, so in order to drive this current through its own resistance and overcome the supply voltage the e.m.f. must be given by

$$\mathrm{E = IR + V = 20 \times 1 + 500 = 520\,Volts}$$

The corresponding speed can be calculated by reference to the no-load e.m.f. (499.2 Volts at 1040 rev/min) from which the steady generating speed is given by

$$\frac{\mathrm{N_{gen}}}{1040} = \frac{520}{499.2} \quad \text{i.e. } \mathrm{N_{gen}} = 1083\,\mathrm{rev/min}$$

On the torque-speed plot (Fig. 3.9) this condition lies on the downward projection of the 500 V characteristic at a current of −20 A. We note that the full range of operation, from full-load motoring to full-load generating is accomplished with only a modest change in speed from 1000 to 1083 rev/min.

It is worth emphasising that in order to make the unloaded motor move into the generating mode, all that we had to do was to start supplying mechanical power to the motor shaft. No physical changes had to be made to the motor to make it into a generator—the hardware is equally at home functioning as a motor or as a generator—which is why it is best referred to as a 'machine'. An electric vehicle takes full advantage of this inherent reversibility, recharging the battery when we decelerate or descend a hill. (How nice it would be if the internal combustion engine could do the same; whenever we slowed down, we could watch the rising gauge as our kinetic energy was converted back into hydrocarbon fuel in the tank!)

To complete this section we will derive the analytical expression for the steady-state speed as a function of the two variables that we can control, i.e. the applied voltage (V), and the load torque (T_L). Under steady-state conditions the armature current is constant and we can therefore ignore the armature inductance term in Eq. (3.7); and because there is no acceleration, the motor torque is equal to the load torque. Hence by eliminating the current I between Eqs (3.5) and (3.7), and substituting for E from Eq. (3.6) the speed is given by

$$\omega = \frac{V}{k} - \frac{R}{k^2} T_L \tag{3.10}$$

This equation represents a straight line in the speed/torque plane, as we saw with our previous worked example. The first term shows that the no-load speed is directly proportional to the armature voltage, while the second term gives the drop in speed for a given load torque. The gradient or slope of the torque-speed curve is $\frac{R}{k^2}$, showing again that the smaller the armature resistance, the smaller the drop in speed when load is applied.

3.4.4 Base speed and field weakening

Returning to our consideration of motor operating characteristics, when the field flux is at its full value the speed corresponding to full armature voltage and full current (i.e. the rated full-load condition) is known as base speed (see Fig. 3.10). The upper part of the figure shows the regions of the torque speed plane within which the motor can operate without exceeding its maximum or rated current. The lower diagram shows the maximum output power as a function of speed.

The motor can operate at any speed up to base speed, and any torque (current) up to rated value by appropriate choice of armature voltage. This full flux

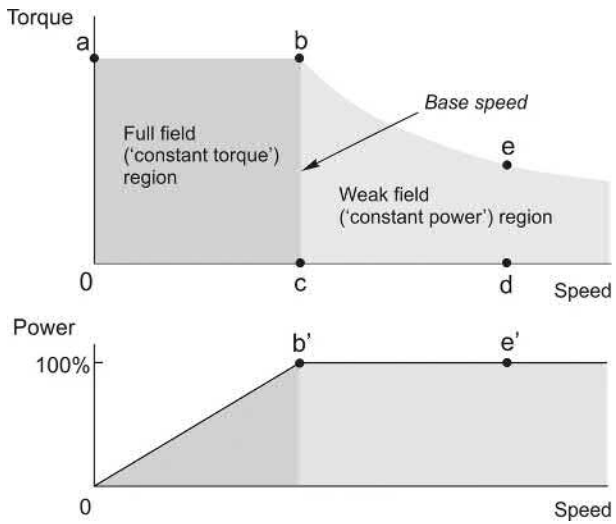

FIG. 3.10 Regions of continuous operation in the torque-speed and power-speed planes.

region of operation is indicated by the shaded area Oabc in Fig. 3.10, and is often referred to as the 'constant torque' region of the torque-speed characteristic. In this context 'constant torque' signifies that at any speed below base speed the motor is capable of producing its full rated torque. Note that the term constant torque does not mean that the motor *will* produce constant torque, but rather it signifies that the motor *can* produce constant torque if required: as we have already seen, it is the mechanical load we apply to the shaft that determines the steady-state torque produced by the motor.

When the current is at maximum (i.e. along the line ab in Fig. 3.10), the torque is at its maximum (rated) value. Since mechanical power is given by torque times speed, the power output along ab is proportional to the speed, as shown in the lower part of the figure, and the maximum power thus corresponds to the point b in Fig. 3.10. At point b both the voltage and current have their full rated values.

To run faster than base speed the field flux must be reduced, as indicated by Eq. (3.9). Operation with reduced flux is known as 'field weakening', and we have already discussed this perhaps surprising mode in connection with the primitive linear motor in Chapter 1. For example by halving the flux (and keeping the armature voltage at its full value), the no-load speed is doubled (point d in Fig. 3.10). The increase in speed is however obtained at the expense of available torque, which is proportional to flux times current (see Eq. 3.1). The current is limited to rated value, so if the flux is halved, the speed will double but the maximum torque which can be developed is only half the rated value (point e in Fig. 3.10). Note that at the point e both the armature voltage and the armature current again have their full rated values, so the power is at maximum, as it was at point b. The power is constant along the curve through b and e, and for this reason the shaded area to the right of the line bc is referred to as the 'constant power' region. Obviously, field weakening is only satisfactory for applications which do not demand full torque at high speeds, such as electric traction.

The maximum allowable speed under weak field conditions must be limited (to avoid excessive sparking at the commutator), and is usually indicated on the motor rating plate. A marking of 1200/1750 rev/min, for example, would indicate a base speed of 1200 rev/min, and a maximum speed with field weakening of 1750 rev/min. The field weakening range varies widely depending on the motor design, but maximum speed rarely exceeds three or four times base speed.

To sum up, the speed is controlled as follows:

- Below base speed, the flux is maximum, and the speed is set by the armature voltage. Full torque is available at any speed.
- Above base speed the armature voltage is at (or close to) maximum, and the flux is reduced in order to raise the speed. The maximum torque available reduces in proportion to the flux.

To judge the suitability of a motor for a particular application we need to compare the torque-speed characteristic of the prospective load with the operating diagram for the motor: if the load torque calls for operation outside the shaded areas of Fig. 3.10, a larger motor is clearly called for.

Finally, we should note that according to Eq. (3.9), the no-load speed will become infinite if the flux is reduced to zero. This seems an unlikely state of affairs: after all, we have seen that the field is essential for the motor to operate, so it seems unreasonable to imagine that if we removed the field altogether, the speed would rise to infinity. In fact, the explanation lies in the assumption that 'no-load' for a real motor means zero torque. If we could make a motor which had no friction torque whatsoever, the speed would indeed continue to rise as we reduced the field flux towards zero. But as we reduced the flux, the torque per ampere of armature current would become smaller and smaller, so in a real machine with friction, there will come a point where the torque being produced by the motor is equal to the friction torque, and the speed will therefore be limited. Nevertheless, it is quite dangerous to open-circuit the field winding, especially in a large unloaded motor. There may be sufficient 'residual' magnetism left in the poles to produce significant accelerating torque to lead to a run-away situation. Usually, field and armature circuits are interlocked so that if the field is lost, the armature circuit is switched off automatically.

3.4.5 Armature reaction

In addition to deliberate field-weakening, as discussed above, the flux in a d.c. machine can be weakened by an effect known as 'armature reaction'. As its name implies, armature reaction relates to the influence that the armature m.m.f. has on the flux in the machine: in small machines it is negligible, but in large machines the unwelcome field weakening caused by armature reaction is sufficient to warrant extra design features to combat it. A full discussion would be well beyond the needs of most users, but a brief explanation is included for the sake of completeness.

The way armature reaction occurs can best be appreciated by looking at Fig. 3.1 and noting that the m.m.f. of the armature conductors acts along the axis defined by the brushes, i.e. the armature m.m.f. acts in quadrature to the main flux axis which lies along the stator poles. The reluctance in the quadrature direction is high because of the large air spaces that the flux has to cross, so despite the fact that the rotor m.m.f. at full current can be very large, the quadrature flux is relatively small; and because it is perpendicular to the main flux, the average value of the latter would not be expected to be affected by the quadrature flux, even though part of the path of the reaction flux is shared with the main flux as it passes (horizontally in Fig. 3.1) through the main pole-pieces.

A similar matter was addressed in relation to the primitive machine in Chapter 1. There it was explained that it was not necessary to take account of the flux produced by the conductor itself when calculating the electromagnetic force on it. And if it were not for the non-linear phenomenon of magnetic saturation, the armature reaction flux would have no effect on the average value of the main flux in the machine shown in Fig. 3.1: the flux density on one edge of the pole-pieces would be increased by the presence of the reaction flux, but decreased by the same amount on the other edge, leaving the average of the main flux unchanged. However if the iron in the main magnetic circuit is already some way into saturation, the presence of the rotor m.m.f. will cause less of an increase on the one edge than it causes by way of decrease on the other, and there will be a net reduction in main flux.

We know that reducing the flux leads to an increase in speed, so we can now see that in a machine with pronounced armature reaction, when the load on the shaft is increased and the armature current increases to produce more torque, the field is simultaneously reduced and the motor speeds up. Though this behaviour is not a true case of instability, it is not generally regarded as desirable!

Large motors often carry additional windings fitted into slots in the pole-faces and connected in series with the armature. These 'compensating' windings produce an m.m.f. in opposition to the armature m.m.f., thereby reducing or eliminating the armature reaction effect.

3.4.6 Maximum output power

We have seen that if the mechanical load on the shaft of the motor increases, the speed falls and the armature current automatically increases until equilibrium of torque is reached and the speed again becomes steady. If the armature voltage is at its maximum (rated) value, and we increase the mechanical load until the current reaches its rated value, we are clearly at full-load, i.e. we are operating at the full speed (determined by voltage) and the full torque (determined by current). The maximum current is set at the design stage, and reflects the tolerable level of heating of the armature conductors.

Clearly if we increase the load on the shaft still more, the current will exceed the safe value, and the motor will begin to overheat. But the question which this prompts is 'if it were not for the problem of overheating, could the motor deliver more and more power output, or is there a limit'?

We can see straightaway that there will be a maximum by looking at the torque-speed curves in Fig. 3.9. The mechanical output power is the product of torque and speed, and we see that the power will be zero when either the load torque is zero (i.e. the motor is running light) or the speed is zero (i.e. the motor is stationary). There must be maximum between these two zeroes, and it is easy to show that the peak mechanical power occurs when the speed is half of the no-load speed. However, this operating condition is only practicable in very small motors: in the majority of motors, the supply would simply not be able to supply the very high current required.

Turning to the question of what determines the theoretical maximum power, we can apply the maximum power transfer theorem (from circuit theory) to the equivalent circuit in Fig. 3.6. The inductance can be ignored because we assume d.c. conditions. If we regard the armature resistance R as if it were the resistance of the source V, the theorem tells us that in order to transfer maximum power to the load (represented by the motional e.m.f. on the right-hand side of Fig. 3.6) we must make the load 'look like' a resistance equal to the source resistance, R. This condition is obtained when the applied voltage V divides equally so that half of it is dropped across R and the other half is equal to the e.m.f., E. (We note that the condition E = V/2 corresponds to the motor running at half the no-load speed, as stated above.) At the maximum power point, the current is V/2R, and the mechanical output power (EI) is given by $V^2/4R$.

The expression for the maximum output power is delightfully simple. We might have expected the maximum power to depend on other motor parameters, but in fact it is determined solely by the armature voltage and the armature resistance. For example, we can say immediately that a 12 V motor with an armature resistance of 1 Ω cannot possibly produce more than 36 W of mechanical output power.

We should of course observe that under maximum power conditions the overall efficiency is only 50% (because an equal power is burned off as heat in the armature resistance); and emphasise again that only very small motors can ever be operated continuously in this condition. For the vast majority of motors, it is of academic interest only, because the current (V/2R) will be far too high for the supply.

3.5 Transient behaviour

It has already been pointed out that the steady-state armature current depends on the small difference between the back e.m.f. E and the applied voltage V. In a converter-fed drive it is vital that the current is kept within safe bounds, otherwise the thyristors or transistors (which have very limited overcurrent capacity) will be destroyed, and it follows from Eq. (3.8) that in order to prevent the current from exceeding its rated value we cannot afford to let V and E differ by more than IR, where I is the rated current.

It would be unacceptable, for example, to attempt to bring all but the smallest of d.c. motors up to speed simply by switching on rated voltage. In the example studied earlier, rated voltage is 500 V, and the armature resistance is 1 Ω. At standstill the back e.m.f. is zero, and hence the initial current would be 500/1 = 500 A, or 25 times rated current! This would destroy the thyristors in the supply converter (and/or blow the fuses). Clearly the initial voltage we must apply is much less than 500 V; and if we want to limit the current to rated value (20 A in the example) the voltage needed will be 20×1, i.e. only 20 Volts. As the speed picks up, the back e.m.f. rises, and to maintain the full current V must also be ramped up so that the difference between V and E remains constant at 20 V. Of course, the motor will not accelerate nearly so rapidly when the current is kept in check as it would if we had switched on full voltage, and allowed the current to do as it pleased. But this is the price we must pay in order to protect the converter.

Similar current-surge difficulties occur if the load on the motor is suddenly increased, because this will result in the motor slowing down, with a consequent fall in E. In a sense we welcome the fall in E because this is what brings about the increase in current needed to supply the extra load, but of course we only want the current to rise to its rated value: beyond that point we must be ready to reduce V, to prevent an excessive current.

The solution to the problem of overcurrents lies in providing closed-loop current-limiting as an integral feature of the motor/drive package. The motor current is sensed, and the voltage V is automatically adjusted so that rated current is not exceeded continuously, although typically it is allowed to reach 1.5 times rated current for up to 60 s. We will discuss the current control loop in Chapter 4.

3.5.1 Dynamic behaviour and time-constants

The use of the terms 'surge' and 'sudden' in the discussion above will doubtless have created the impression that changes in the motor current or speed can take place instantaneously, whereas in fact a finite time is always necessary to effect changes in either. (If the current changes, then so does the stored energy in the armature inductance; and if speed changes, so does the rotary kinetic energy stored in the inertia. For either of these changes to take place in zero time it would be necessary for there to be a pulse of infinite power, which is clearly impossible.)

The theoretical treatment of the transient dynamics of the d.c. machine is easier than for any other type of electric motor but is nevertheless beyond our scope. However it is worth summarising the principal features of the dynamic behaviour, and highlighting the fact that all the transient changes that occur are determined by only two time-constants. The first (and most important from the user's viewpoint) is the electromechanical time-constant, which governs the way the speed settles to a new level following a disturbance such as a change in armature voltage or load torque. The second is the electrical (or armature) time-constant, which is usually much shorter and governs the rate of change of armature current immediately following a change in armature voltage.

When the motor is running, there are two 'inputs' that we can change suddenly, namely the applied voltage and the load torque. When either of these is changed, the motor enters a transient period before settling to its new steady state. It turns out that if we ignore the armature inductance (i.e. we take the armature time-constant to be zero), the transient period is characterised by first-order exponential responses in the speed and current. This assumption is valid for all but the very largest motors. We obtained a similar result when we looked at the primitive linear motor in Chapter 1 (see Fig. 1.16).

For example, if we suddenly increased the armature voltage of a frictionless and unloaded motor from V_1 to V_2, its speed and current would vary as shown in Fig. 3.11.

There is an immediate increase in the current (because we have ignored the inductance), reflecting the fact that the applied voltage is suddenly more than the back e.m.f.; the increased current produces more torque and hence the motor accelerates; the rising speed is accompanied by an increase in back e.m.f., so the current begins to fall; and the process continues until a new steady speed is reached corresponding to the new voltage. In this particular case the steady-state current is zero because we have assumed that there is no friction or load torque, but the shape of the dynamic response would be the same if there had been an initial load, or if we had suddenly changed the load.

The expression describing the current as a function of time (t) is:-

$$i = \left\{ \frac{V_2 - V_1}{R} \right\} e^{\frac{-t}{\tau}} \tag{3.11}$$

The expression for the change in speed is similar, the time dependence again featuring the exponential transient term, $e^{\frac{-t}{\tau}}$. The significance of the time-constant (τ) is shown in Fig. 3.11. If the initial gradient of the current-time graph is projected it intersects the final value after one time-constant. In theory, it takes an infinite time for the response to settle, but in practice the transient is usually regarded as over after about 4 or 5 time-constants. We note that the transient response is very satisfactory: as soon as the voltage is increased the current immediately increases to provide more torque and begin the acceleration, but the accelerating torque is reduced progressively to ensure that the new target speed is approached smoothly. Happily, because the system is first-order, there is no suggestion of an oscillatory response with overshoots.

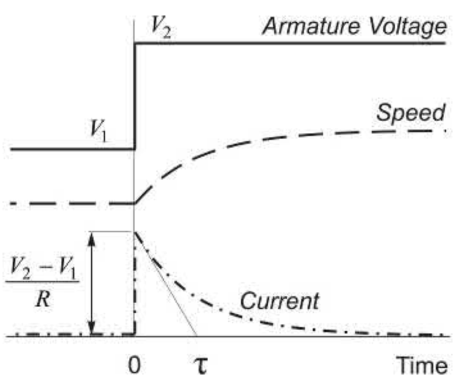

FIG. 3.11 Response of d.c. motor to step increase in armature voltage.

Analysis yields the relationship between the time-constant and the motor/system parameters as

$$\tau = \frac{RJ}{k^2} \tag{3.12}$$

where R is the armature resistance, J is the total rotary inertia of motor plus load, and k is the motor constant (Eqs 3.3 and 3.4). The appropriateness of the term 'electromechanical time-constant' should be clear from Eq. (3.12), because τ depends on the electrical parameters (R and k) and the mechanical parameter, J. The fact that if the inertia was doubled, the time-constant would double and transients would take twice as long is perhaps to be expected, but the influence of the motor parameters R and k is probably not so obvious.

The electrical or armature time-constant is defined in the usual way for a series L, R circuit, i.e.

$$\tau_a = \frac{L}{R} \tag{3.13}$$

If we were to hold the rotor of a d.c. motor stationary and apply a step voltage V to the armature, the current would climb exponentially to a final value of V/R with a time-constant τ_a.

If we always applied pure d.c. voltage to the motor we would probably want τ_a to be as short as possible, so that there was no delay in the build-up of current when the voltage is changed.

But given that most motors are fed with voltage waveforms which are far from smooth (see Chapter 2), we are actually rather pleased to find that because of the inductance and associated time-constant, the current waveform (and hence the torque) are smoother than the voltage waveform. So the unavoidable presence of armature inductance turns out (in most cases) to be a blessing in disguise.

So far we have looked at the two time-constants as if they were unrelated in the influence they have on the current. We began with the electromechanical time-constant, assuming that the armature time-constant was zero, and saw that the dominant influence on the current during the transient was the motional e.m.f. We then examined the current when the rotor was stationary (so that the motional e.m.f. is zero), and saw that the growth or decay of current is governed by the armature inductance, manifested via the armature time-constant.

In reality, both time-constants influence the current simultaneously, and the picture is more complicated than we have implied, as the system is in fact a second-order one. However, the good news is that for most motors, and most purposes, we can take advantage of the fact that the armature time-constant is much shorter than the electromechanical time-constant. This allows us to approximate the behaviour by decoupling the relatively fast 'electrical transients' in the armature circuit from the much slower 'electromechanical transients' which are apparent to the user. From the latter's point of view, only the electromechanical transient is likely to be of interest.

3.6 Four quadrant operation and regenerative braking

We have seen that the great beauty of the separately-excited d.c. motor is the ease with which it can be controlled. The steady-state speed is determined by the applied voltage, so we can make the motor run at any desired speed in either direction simply by applying the appropriate magnitude and polarity of the armature voltage; and the torque is directly proportional to the armature current, which in turn depends on the difference between the applied voltage V and the back e.m.f. E. We can therefore make the machine develop positive (motoring) or negative (generating) torque simply by controlling the extent to which the applied voltage is greater or less than the back e.m.f. An armature voltage controlled d.c. machine is therefore inherently capable of what is known as 'four-quadrant' operation, by reference to the numbered quadrants of the torque-speed plane shown in Fig. 3.12.

Fig. 3.12 looks straightforward but experience shows that to draw the diagram correctly calls for a clear head, so it is worth spelling out the key points in detail. A proper understanding of this diagram is invaluable as an aid to understanding how controlled-speed drives operate.

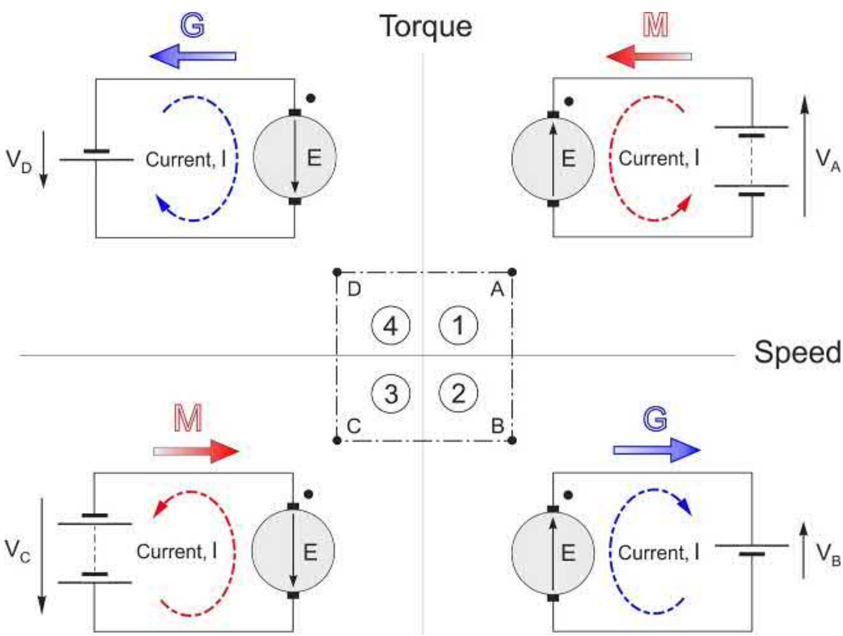

FIG. 3.12 Operation of d.c. motor in the four quadrants of the torque-speed plane.

Firstly, one of the motor terminals is shown with a dot, and in all four quadrants the dot is uppermost. The purpose of this convention is to indicate the sign of the torque: if current flows into the dot, the machine produces positive torque, and if current flows out of the dot, the torque is negative.

Secondly, the supply voltage is shown by the old-fashioned battery symbol, as use of the more modern circle symbol for a voltage source would make it more difficult to differentiate between the source and the circle representing the machine armature. The relative magnitudes of applied voltage and motional e.m.f. are emphasised by the use of two battery cells when $V > E$ and one when $V < E$.

We have seen that in a d.c. machine speed is determined by applied voltage and torque is determined by current. Hence on the right-hand side of the diagram the supply voltage is positive (upwards), while on the left-hand side the supply voltage is negative (downwards). And in the upper half of the diagram current is positive (into the dot), while in the lower half it is negative (out of the dot). For the sake of convenience, each of the four operating conditions (A, B, C, D) have the same magnitude of speed and the same magnitude of torque: these translate to equal magnitudes of motional e.m.f. and current for each condition.

When the machine is operating as a motor and running in the forward direction, it is operating in quadrant 1. The applied voltage V_A is positive and greater than the back e.m.f. E, and positive current therefore flows into the motor: in Fig. 3.12, the arrow representing V_A has accordingly been drawn larger than E. The power drawn from the supply ($V_A I$) is positive in this quadrant, as shown by the shaded arrow labelled M to represent motoring. The power converted to mechanical form is given by EI, and an amount I^2R is lost as heat in the armature. If E is much greater than IR (which is true in all but small motors), most of the input power is converted to mechanical power, i.e. the conversion process is efficient.

If, with the motor running at position A, we suddenly reduce the supply voltage to a value V_B which is less than the back e.m.f., the current (and hence torque) will reverse direction, shifting the operating point to B in Fig. 3.12. There can be no sudden change in speed, so the e.m.f. will remain the same. If the new voltage is chosen so that $E - V_B = V_A - E$, the new current will have the same amplitude as at position A, so the new (negative) torque will be the same as the original positive torque, as shown in Fig. 3.12. But now power is supplied from the machine to the supply, i.e. the machine is acting as a generator, as shown by the shaded arrow.

We should be quite clear that all that was necessary to accomplish this remarkable reversal of power flow was a modest reduction of the voltage applied to the machine. At position A, the applied voltage was E + IR, while at position B it is E – IR. Since IR will be small compared with E, the change (2IR) is also small.

Needless to say the motor will not remain at point B if left to its own devices. The combined effect of the load torque and the negative machine torque will cause the speed to fall, so that the back e.m.f. again falls below the applied voltage V_B, the current and torque become positive again, and the motor settles back into quadrant 1, at a lower speed corresponding to the new (lower) supply voltage. During the deceleration phase, kinetic energy from the motor and load inertias is returned to the supply. This is therefore an example of regenerative braking, and it occurs naturally every time we reduce the voltage in order to lower the speed.

If we want to operate continuously at position B, the machine will have to be driven by a mechanical source. We have seen above that the natural tendency of the machine is to run at a lower speed than that corresponding to point B, so we must force it to run faster, and create an e.m.f. greater than V_B, if we wish it to generate continuously.

It should be obvious that similar arguments to those set out above apply when the motor is running in reverse (i.e. V is negative). Motoring then takes place in quadrant 3 (point C), with brief excursions into quadrant 4 (point D, accompanied by regenerative braking) whenever the voltage is reduced in order to lower the speed.

3.6.1 Full speed regenerative reversal

To illustrate more fully how the voltage has to be varied during sustained regenerative braking, we can consider how to change the speed of an unloaded motor from full speed in one direction to full speed in the other, in the shortest possible time.

At full forward speed the applied armature voltage is taken to be +V (shown as 100% in Fig. 3.13), and since the motor is unloaded the no-load current will be very small and the back e.m.f. will be almost equal to V. Ultimately, we will clearly need an armature voltage of −V to make the motor run at full speed in reverse. But we cannot simply reverse the applied voltage: if we did, the armature current immediately afterwards would be given by (−V −E)/R, which would be disastrously high. (The motor might tolerate it for the short period for which it would last, but the supply certainly could not!).

What we need to do is adjust the voltage so that the current is always limited to rated value, and in the right direction. Since we want to decelerate as fast as possible, we must aim to keep the current negative, and at rated value (i.e. −100%) throughout the period of deceleration and for the run up to full speed in reverse. This will give us constant torque throughout, so the deceleration (and subsequent acceleration) will be constant and the speed will change at a uniform rate, as shown in Fig. 3.13.

We note that to begin with the applied voltage has to be reduced to less than the back e.m.f., and then ramped down linearly with time so that the difference between V and E is kept constant, thereby keeping the current constant at its rated value. During the reverse run-up, V has to be numerically greater than

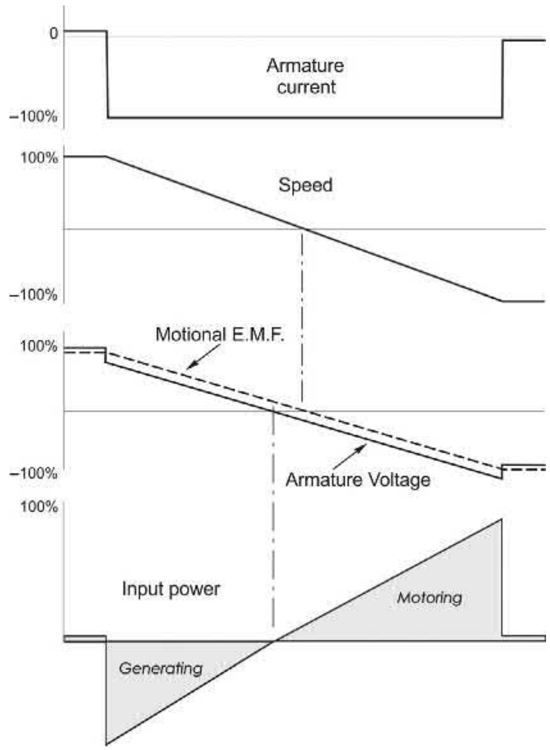

FIG. 3.13 Regenerative reversal of d.c. motor from full-speed forward to full-speed reverse, at maximum allowable torque (current).

E, as shown in Fig. 3.13. (The difference between V and E has been exaggerated in Fig. 3.13 for clarity: in a large motor, the difference may only be 1 or 2% at full speed.)

The power to and from the supply is shown in the bottom plot in Fig. 3.13, the energy being represented by the shaded areas. During the deceleration period most of the kinetic energy of the motor (lower shaded area) is progressively returned to the supply, the motor acting as a generator for the whole of this time. The total energy recovered in this way can be appreciable in the case of a large drive such as a steel rolling mill. A similar quantity of energy (upper shaded area) is supplied and stored as kinetic energy as the motor picks up speed in the reverse sense.

Three final points need to be emphasised. Firstly, we have assumed throughout the discussion that the supply can provide positive or negative voltages, and can accept positive or negative currents. A note of caution is therefore appropriate, because many simple power electronic converters do not have this flexibility. Users need to be aware that if full four-quadrant operation (or even two-quadrant regeneration) is called for, a basic converter will probably not be adequate. This point is taken up again in Chapter 4. Secondly, we should not run away with the idea that in order to carry out the reversal in Fig. 3.13 we would have to work out in advance how to profile the applied voltage as a function of time. Our drive system will normally have the facility for automatically operating the motor in constant-current mode, and all we will have to do is to tell it the new target speed. This is also taken up in Chapter 4.

3.6.2 Dynamic braking

A simpler and cheaper but less effective method of braking can be achieved by dissipating the kinetic energy of the motor and load in a resistor, rather than returning it to the supply. A version of this technique is employed in the cheaper power electronic converter drives, which have no facility for returning power to the utility supply.

When the motor is to be stopped, the supply to the armature is removed and a resistor is switched across the armature brushes. The motor e.m.f. drives a (negative) current through the resistor, and the negative torque results in deceleration. As the speed falls, so does the e.m.f., the current, and the braking torque. At low speeds the braking torque is therefore very small. Ultimately, all the kinetic energy is converted to heat in the motor's own armature resistance and the external resistance. Very rapid initial braking is obtained by using a low resistance (or even simply short-circuiting the armature).

3.7 Shunt and series motors

Before variable-voltage supplies became readily available, most d.c. motors were obliged to operate from a single d.c. supply, usually of constant voltage. The armature and field circuits were therefore designed either for connection in parallel (shunt), or in series. The operating characteristics of shunt and series machines differ widely, and hence each type tended to claim its particular niche: shunt motors were judged to be good for constant-speed applications, while series motors were widely used for traction applications.

In a way it is unfortunate that these historical patterns of association became so deep-rooted. The fact is that a converter-fed separately-excited d.c. motor, freed of any constraint between field and armature, can do everything that a shunt or series d.c. motor can, and more; and it is doubtful if shunt and series motors would ever have become widespread if variable-voltage supplies had always been around.

For a given continuous output power rating at a given speed, shunt and series motors are the same size, with the same rotor diameter, the same poles, and the same quantities of copper in the armature and field windings. This is to be expected when we recall that the power output depends on the specific magnetic and electric loadings, so we anticipate that to do a given job, we will need the same amounts of active material. However differences emerge when we look at the details of the windings, especially the field system, and they can best be illustrated by means of an example which contrasts shunt and series motors for the same output power.

Suppose that for the shunt version the supply voltage is 500 V, the rated armature (work) current is 50 A, and the field coils are required to provide an m.m.f. of 500 Amp-turns. The field might typically consist of say 200 turns of wire with a total resistance of 200 Ω. When connected across the supply (500 V), the field current will be 2.5 A, and the m.m.f. will be 500 AT, as required. The power dissipated as heat in the field will be 500 V × 2.5 A = 1.25 kW, and the total power input at rated load will be 500 V × 52.5 A = 26.25 kW.

To convert the machine into the equivalent series version, the field coils need to be made from much thicker conductor, since they have to carry the armature current of 50 A, rather than the 2.5 A of the shunt motor. So, working at the same current density, the cross-section of each turn of the series field winding needs to be 20 times that of the shunt field wires, but conversely only one-twentieth of the turns (i.e. 10) are required for the same ampere-turns. The resistance of a wire of length l and cross sectional area A, made from material of resistivity ρ is given by $R = \frac{\rho l}{A}$, so we can use this formula to show that the resistance of the new field winding will be much lower, at 0.5 Ω.

We can now calculate the power dissipated as heat in the series field. The current is 50 A, the resistance is 0.5 Ω, so the volt-drop across the series field is 25 V, and the power wasted as heat is 1.25 kW. This is the same as for the shunt machine, which is to be expected since both sets of field coils are intended to do the same job.

In order to allow for the 25 V dropped across the series field, and still meet the requirement for 500 V at the armature, the supply voltage must now be 525 V. The rated current is 50 A, so the total power input is 525 V × 50 A = 26.25 kW, the same as for the shunt machine.

This example illustrates that in terms of their energy-converting capabilities, shunt and series motors are fundamentally no different. Shunt machines usually have field windings with a large number of turns of fine wire, while series machines have a few turns of thick conductor. But the total amount and disposition of copper is the same, so the energy-converting abilities of both types are identical. In terms of their operating characteristics, however, the two types differ widely, as we will now see.

3.7.1 Shunt motor—steady-state operating characteristics

A basic shunt-connected motor has its armature and field in parallel across a single d.c. supply, as shown in Fig. 3.14A. Normally, the voltage will be constant and at the rated value for the motor, in which case the steady-state torque/speed curve will be similar to that of a separately-excited motor at rated field flux, i.e. the speed will drop slightly with load, as shown by the line ab in Fig. 3.14B. Over the normal operating region the torque-speed characteristic is similar to that

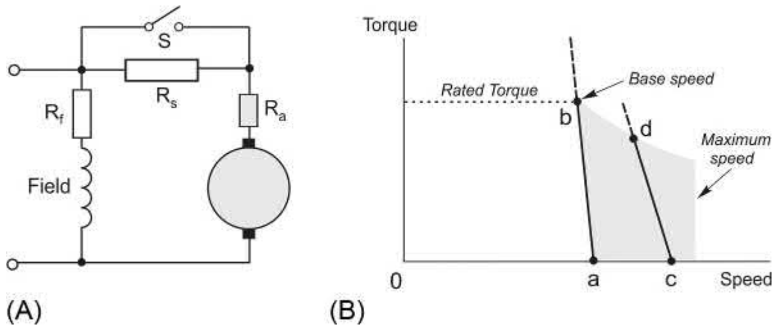

FIG. 3.14 Shunt-connected d.c. motor and steady-state torque-speed curve.

of the induction motor (see Chapter 6), so shunt motors are suited to the same duties, i.e. what are usually referred to as 'constant speed' applications.

Except for small motors (say less than about 1 kW), it will be necessary to provide an external 'starting resistance' (R_S in Fig. 3.14) in series with the armature, to limit the heavy current which would flow if the motor was simply switched directly onto the supply. This starting resistance is progressively reduced as the motor picks up speed, the current falling as the back e.m.f. rises from its initial value of zero.

We should ask what happens if the supply voltage varies for any reason, and as usual the easiest thing to look at is the case where the motor is running light, in which case the back e.m.f. will almost equal the supply voltage. If we reduce the supply voltage, intuition might lead us to anticipate a fall in speed, but in fact two contrary effects occur which leave the speed almost unchanged.

If the voltage is halved, for example, both the field current and the armature voltage will be halved, and if the magnetic circuit is not saturated the flux will also halve. The new steady value of back e.m.f. will have to be half its original value, but since we now have only half as much flux, the speed will be the same. The maximum output power will of course be reduced, since at full load (i.e. full current) the power available is proportional to the armature voltage. Of course if the magnetic circuit is saturated, a modest reduction in applied voltage may cause very little drop in flux, in which case the speed will fall in proportion to the drop in voltage. We can see from this discussion why, broadly speaking, the shunt motor is not suitable for operation below base speed.

Some measure of speed control is possible by weakening the field (by means of the resistance (R_f) in series with the field winding), and this allows the speed to be raised above base value, but only at the expense of torque. A typical torque-speed characteristic in the field-weakening region is shown by the line cd in Fig. 3.14B.

Reverse rotation is achieved by reversing the connections to either the field or the armature. The field is usually preferred since the current rating of the switch or contactor will be lower than for the armature.

3.7.2 Series motor—steady-state operating characteristics

The series connection of armature and field windings (Fig. 3.15A) means that the field flux is directly proportional to the armature current, and the torque is therefore proportional to the square of the current. Reversing the direction of the applied voltage (and hence current) therefore leaves the direction of torque unchanged. This unusual property is put to good use in the universal motor, but is a handicap when negative (braking) torque is required, since either the field or armature connections must then be reversed.

If the armature and field resistance volt-drops are neglected, and the applied voltage (V) is constant, the current varies inversely with the speed, hence the torque (T) and speed (n) are related by

$$T \propto \left\{ \frac{V}{n} \right\}^2 \tag{3.14}$$

A typical torque-speed characteristic is shown in Fig. 3.15B. The torque at zero speed is not infinite of course, because of the effects of saturation and resistance, both of which are ignored in Eq. (3.14).

As with the shunt motor, under starting conditions the back e.m.f. is zero, and if the full voltage was applied the current would be excessive, being limited only by the armature and field resistances. Hence for all but small motors a starting resistance is required to limit the current to a safe value.

Returning to Fig. 3.15B, we note that the series motor differs from most other motors in having no clearly defined no-load speed, i.e. no speed (other than infinity) at which the torque produced by the motor falls to zero. This means that when running light, the speed of the motor depends on the windage and friction torques, equilibrium being reached when the motor torque equals the total mechanical resisting torque. In large motors, the windage and friction torque is often relatively small, and the no-load speed is then too high for mechanical safety. Large series motors should therefore never be run uncoupled

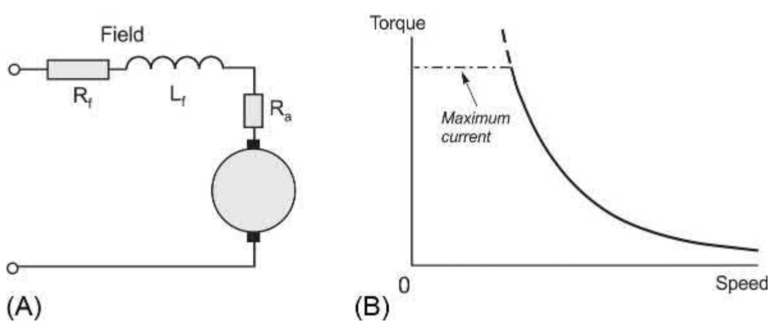

FIG. 3.15 Series-connected d.c. motor and steady-state torque-speed curve.

from their loads. As with shunt motors, the connections to either the field or armature must be reversed in order to reverse the direction of rotation.

The reason for the historic use of series motors for traction is that under the simplest possible supply arrangement (i.e. constant voltage) the overall shape of the torque-speed curve fits well with what is needed in traction applications, i.e. high starting torque to provide acceleration from rest and an acceptable reduction in acceleration as the target speed is approached. Early systems achieved coarse speed control via resistors switched in parallel with either the field or the armature in order to divert some of the current and thereby alter the torque-speed curve, but this inherently inefficient approach became obsolete when power electronic converters became available to provide an efficient variable-voltage supply.

3.7.3 Universal motors

In terms of numbers the main application area for the series commutator motor is in portable power tools, food mixers, vacuum cleaners, etc., where paradoxically the supply is a.c. rather than d.c. Such motors are often referred to as 'universal' motors because they can run from either a d.c. or an a.c. supply.

At first sight the fact that a d.c. machine will work on a.c. is hard to believe. But when we recall that in a series motor the field flux is set up by the current which also flows in the armature, it can be seen that reversal of the current will be accompanied by reversal of the direction of the magnetic flux, thereby ensuring that the torque remains positive. When the motor is connected to a 50 Hz supply for example, the (sinusoidal) current will change direction every 10 ms, and there will be a peak in the torque 100 times per second. But the torque will always remain unidirectional, and the speed fluctuations will not be noticeable because of the smoothing effect of the armature inertia.

Series motors for use on a.c. supplies are always designed with fully laminated construction (to limit eddy current losses produced by the pulsating flux in the magnetic circuit), and are intended to run at high speeds, say 8–12,000 rev/min. at rated voltage. Commutation and sparking are worse than when operating from d.c., and output powers are seldom greater than 1 kW. The advantage of high speed in terms of power output per unit volume was emphasised in Chapter 1, and the universal motor is a good example which demonstrates how a high power can be obtained with small size by designing for a high speed.

Until recently the universal motor offered the only relatively cheap way of reaping the benefit of high speed from single-phase a.c. supplies. Other small a.c. machines, such as induction motors and synchronous motors, were limited to maximum speeds of 3000 rev/min at 50 Hz (or 3600 rev/min at 60 Hz), and therefore could not compete in terms of power per unit volume. The availability of high-frequency inverters (see Chapter 7) has opened up the prospect of higher specific outputs from induction motors and PM motors in many appliances. The transition away from universal motor is underway, but because of the huge investment in high volume manufacturing that has been made over many years, they are still being produced in significant numbers.

Speed control of small universal motors is straightforward using a triac (in effect a pair of thyristors connected back to back) in series with the a. c. supply. By varying the firing angle, and hence the proportion of each cycle for which the triac conducts, the voltage applied to the motor can be varied to provide speed control. This approach is widely used for electric drills, fans etc. If torque control is required (as in hand power tools, for example), the current is controlled rather than the voltage, and the speed is determined by the load.

3.8 Self-excited d.c. machine

At various points in this chapter, we saw that the d.c. machine can change from motoring to generating depending on whether its induced or generated e.m.f. was less than or greater than the voltage applied to the armature terminals. We assumed that the machine was connected to a voltage source, from which it could draw both its excitation (field) and armature currents. But if we want to use the machine as a generator in a location isolated from any voltage source, we have to consider how to set up the flux to enable the machine to work as an energy converter.

When the field is provided by permanent magnets, there is no problem. As soon as the prime-mover (an internal combustion engine, perhaps) starts to turn the rotor, a speed-dependent e.m.f. (voltage) is produced in the armature and we can connect our load to the terminals and start to generate power. To ensure a more or less constant voltage, the engine will be governed to keep the speed constant.

If the machine has a field winding that is intended to be connected in parallel (shunt) with the armature, it is easy to imagine that once the machine is generating an armature voltage, there will be a field current and flux. But before we get a motional e.m.f. we need flux, and we can't have flux without a field current, and we won't get field current until we have generated an e.m.f.

The answer to this apparent difficulty usually lies in taking advantage of the residual magnetic flux remaining in the magnetic circuit from the time the machine was last used. If the residual flux is sufficient we will begin to generate as soon as the rotor is turned, and provided the field is connected the right way round across the armature, the first few milliamps of field current will increase the flux and we will then have positive feedback which we might expect to result in a runaway situation with ever-increasing generated voltage. Fortunately, provided that the driven speed is constant, the growth of flux is stabilised by saturation of the magnetic circuit, as we will discuss next. This procedure is known as self-excitation, for obvious reasons. (Should we be

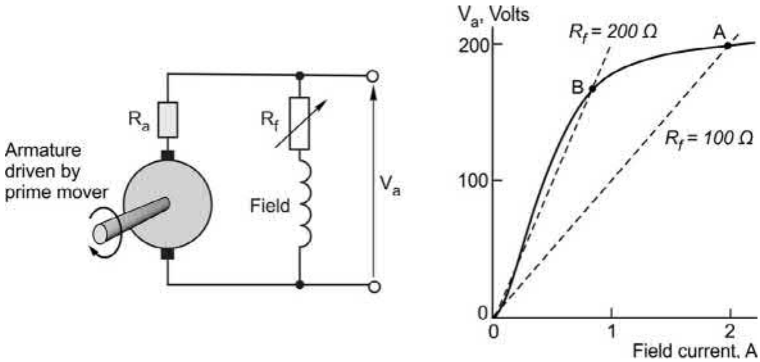

FIG. 3.16 Self-excited d.c. generator.

unfortunate enough to connect the field the wrong way round, we will be disappointed: but if the connections were undisturbed from the previous time, we should be OK.)

The stabilising effect of saturation is shown in Fig. 3.16. The solid line in the graph on the right is a typical 'magnetisation curve' of the machine. This is a plot of the generated e.m.f. (E) as a function of field current when the rotor is driven at a constant speed. This test is best done with a separate supply to the field, and (usually) at the rated speed. For low and medium values of the field current, the flux is proportional to the current, so the motional e.m.f. is linear, but at higher field currents the magnetic circuit starts to saturate, the flux no longer increases in proportion, and consequently the generated voltage flattens out. The armature circuit is open, so the terminal voltage (V_a) is equal to the induced e.m.f. (E).

When we want the machine to self-excite, we connect the field in parallel with the armature as shown in Fig. 3.16. The current in the field is then given by $I_f = V_a/R_f$ where R_f is the resistance of the field winding, and we have neglected the small armature resistance. The two dotted lines relate to field circuit resistances of 100 and 200 Ω, and the intersection of these with the magnetisation curve give the stable operating points. For example, when the field resistance is 100 Ω, and the armature voltage is 200 V, the field current of 2 A is just sufficient to generate the motional e.m.f. of 200 V (point A). At lower values of field current, the generated voltage exceeds the value needed to maintain the field current, so the excess voltage cause the current to increase in the inductive circuit until the steady state is reached.

Increasing the resistance of the field circuit to 200 Ω results in a lower generated voltage (point B), but this is a less stable situation, with small changes in resistance leading to large changes in voltage. It should also be clear that if the resistance is increased above 200 Ω, the resistance line will be above the magnetisation curve, and self-excitation will be impossible.

3.9 Toy motors

The motors used in model cars, trains etc. are rather different in construction from those discussed so far, primarily because they are designed almost entirely with cost as the primary consideration. They also run at high speeds, so it is not important for the torque to be smooth. A typical arrangement used for rotor diameters from 0.5 cm to perhaps 3 cm is shown in Fig. 3.17.

The rotor, made from laminations with a small number (typically three or five) of multi-turn coils in very large 'slots,' is simple to manufacture, and because the commutator has few segments, it too is cheap to make. The field system (stator) consists of radially-magnetised ceramic magnets with a steel backplate to complete the magnetic circuit.

The rotor clearly has very pronounced saliency, with three very large projections which are in marked contrast to the rotors we looked at earlier where the surface was basically cylindrical. It is easy to imagine that even when there is no current in the rotor coils, there is a strong tendency for the stator magnets to pull one or other of the rotor saliencies into alignment with a stator pole, so that the rotor would tend to lock in any one of six positions. This cyclic 'detent' torque, is due to the variation of reluctance with rotor position, an effect which is exploited in a.c. reluctance motors (see Chapter 9), but is unwanted here. To combat the problem the rotor laminations are skewed before the coils are fitted, as shown in the lower sketch on the left.

Each of the three rotor poles carries a multi-turn coil, the start of which is connected to a commutator segment, as shown in the cross-section on the right of Fig. 3.17. The three ends are joined together. The brushes are wider than the inter-segment space, so in some rotor positions the current from say the positive brush will divide and flow first through two coils in parallel until it reaches the common point, then through the third coil to the negative brush, while in other positions the current flows through only two coils.

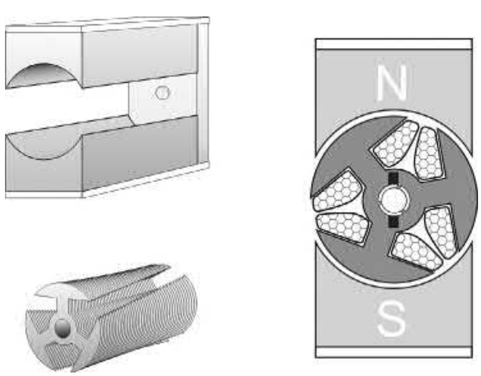

FIG. 3.17 Miniature d.c. motor for use in models and toys.

A real stretch of imagination is required to picture the mechanism of torque production using the '*BIl*' approach we have followed previously, as the geometry is so different. But we can take a more intuitive approach by considering, for each position, the polarity and strength of the magnetisation of each of the three rotor poles, which depend of course on the direction and magnitude of the respective currents. Following the discussion in the previous paragraph, we can see that in some positions there will be say a strong N and two relatively weak S poles, while at others there will be one strong N, one strong S, and one unexcited pole.

Thus although the rotor appears to be a three-pole device, the magnetisation pattern is always two-pole (because two adjacent weak S poles function as a single stronger S pole). When the rotor is rotating, the stator N pole will attract the nearest rotor S pole, pulling it round towards alignment. As the force of attraction diminishes to zero, the commutator reverses the rotor current in that pole so that it now becomes a N and is pushed away towards the S stator pole.

There is a variation of current with angular position as a result of the different resistances seen by the brushes as they alternately make contact with either one or two segments, and the torque is far from uniform, but operating speeds are typically several thousand rev/min and the torque pulsations are smoothed out by the rotor and load inertia.

3.10 Review questions

(1) **(a)** What is the primary (external) parameter that determines the speed of an unloaded d.c. motor?

(b) What is the primary external factor that determines the steady-state running current of a d.c. motor, for any given armature voltage?

(c) What determines the small current drawn by a d.c. motor when running without any applied mechanical load?

(d) What determines how much the speed of a d.c. motor reduces when the load on its shaft is increased? Why do little motors slow down more than large ones?

(2) What has to be done to reverse the direction of rotation of:-
(a) a separately-excited motor; (b) a shunt motor; (c) a series motor?

(3) Most d.c. motors can produce much more than their continuously-rated torque. Why is it necessary to limit the continuous torque?

(4) Why do d.c. motors run faster when their field flux is reduced?

(5) A separately-excited d.c. motor runs from a 220 V d.c. supply and draws an armature current of 15 A. The armature resistance is 0.8 Ω. Calculate the generated voltage (back e.m.f).

If the field current is suddenly reduced by 10%, calculate (a) the value to which the armature current rises momentarily, and (b) the percentage increase in torque when the current reaches the value in (a). Ignore armature inductance, neglect saturation, and assume that the field flux is directly proportional to the field current.

(6) **(a)** When driven at 1500 rev/min the open-circuit armature voltage of a d.c. machine is 110 V. Calculate the e.m.f. constant in volts per radian/s. Calculate also the machine torque when the armature current is 10 A.

(b) Suppose the machine was at rest, and a weight of 5 kg was suspended from a horizontal bar of length 80 cm attached to the shaft, as shown in Fig. Q6.

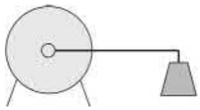

FIG. Q6

What current must be applied to the armature to keep the arm horizontal? Will the equilibrium be a stable one? (Neglect the mass of the bar; $g=9.81\,m/s^2$.)

(c) When the machine runs as a motor drawing 25 A from a 110 V supply, the speed is 1430 rev/min. Calculate the armature resistance. Hence find the voltage needed to keep the bar in (b) horizontal.

(d) At what speed must the machine be driven to supply 3.5 kW to a 110 V system?
Calculate the corresponding torque. If the field circuit consumes 100 W and the friction losses are 200 W, calculate the efficiency of the generator.

(7) The equations expressing torque in terms of current ($T=kI$) and motional e.m.f. in terms of speed ($E=k\omega$) are central in understanding the operation of a d.c. machine. Using only these equations, show that the mechanical output power is given by $W=EI$.

(8) Explain briefly why:-

(a) large d.c. motors cannot normally be started by applying full voltage;

(b) the no-load speed of a permanent-magnet motor is almost proportional to the armature voltage.

(c) a d.c. motor draws more current from the supply when the load on the shaft is increased.

(d) the field windings of a d.c. motor consume energy continuously even though they do not contribute to the mechanical output power.

(e) the field poles of a d.c. machine are not always laminated.

(9) It is claimed in this book that a motor of a given size and power can be made available for operation at any voltage. But it is clear that, when it comes to battery-powered hand tools of a given size and speed, the higher-voltage versions are more powerful. What accounts for this contradiction?

(10) This question relates to a permanent-magnet d.c. machine with an armature resistance of 0.5 Ω.

When the rotor is driven at 1500 rev/min by an external mechanical source, its open-circuit armature voltage is 220 V. All parts of the question relate to steady-state conditions, i.e. after all transients have died away.

The machine is to be used as a generator and act as a dynamic brake to restrain a lowering load, as shown in Fig. Q10A. The hanging mass of 14.27 kg is suspended by a rope from a 20 cm diameter winding drum on the motor shaft.

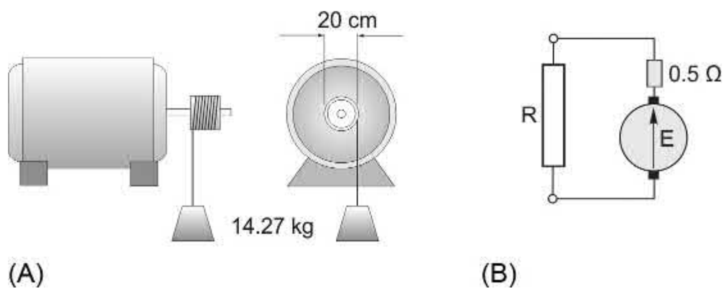

FIG. Q10

The majority of the generated power is to be dissipated in an external resistor (R) connected across the armature, as shown in Fig. Q10B.

(a) Calculate the value of resistance required so that the mass descends at a steady speed of 15 m/s. Take $g = 9.81\ m/s^2$.

(b) What is the power dissipated in (i) the external resistor, (ii) the armature? Where does the energy dissipated in the resistors come from?

Answers to the review questions are given in the Appendix.

Chapter 4

D.C. motor drives

4.1 Introduction

Until the 1960s, the only satisfactory way of obtaining the variable-voltage d.c. supply needed for speed control of an industrial d.c. motor was to generate it with a d.c. generator. The generator was driven at fixed speed by an induction motor, and the field of the generator was varied in order to vary the generated voltage. For a brief period in the 1950s these 'Ward-Leonard' sets were superseded by grid-controlled Mercury Arc rectifiers, but these were soon replaced by thyristor converters which offered lower cost, higher efficiency (typically over 95%), smaller size, reduced maintenance and faster response to changes in set speed. The disadvantages of these rectified supplies are that the waveforms are not pure d.c., that the overload capacity of the converter is very limited, and that a single converter is not inherently capable of regeneration.

Though no longer pre-eminent, study of the d.c. drive is valuable for two reasons:-

- The structure and operation of the d.c. drive are reflected in almost all other drives, and lessons learned from the study of the d.c. drive therefore have close parallels in other types.
- Under constant-flux conditions the behaviour is governed by a relatively simple set of linear equations, so predicting both steady-state and transient behaviour is not difficult. When we turn to the successors of the d.c. drive, notably the induction motor drive, we will find that things are much more complex, and that in order to overcome the poor transient behaviour, the control strategies adopted are based on emulating the inherent characteristics of the d.c. drive.

The first and major part of this chapter is devoted to thyristor-fed drives, after which we will look briefly at chopper-fed drives that are used mainly in the medium and small sizes, and finally turn attention to small servo-type drives.

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00004-0

4.2 Thyristor d.c. drives—general

For motors up to a few kilowatts the armature converter draws power from either a single-phase or three-phase utility supply. For larger motors three-phase is preferred because the waveforms are much smoother. Traction applications, where only a single-phase supply is available, require a series inductor to smooth the current. A separate thyristor or diode rectifier is used to supply the field of the motor: the power is much less than the armature power, the inductance is much higher, and so the supply is often single-phase, as shown in Fig. 4.1.

The arrangement shown in Fig. 4.1 is typical of most d.c. drives and provides for closed-loop speed control. The function of the two control loops will be explored later, so readers who are not familiar with the basics of feedback and closed-loop systems may find it necessary to consult an introductory text first.[1]

The main power circuit consists of a six-thyristor bridge circuit (as discussed in Chapter 2) which rectifies the incoming a.c. supply to produce a d.c. supply to the motor armature. By altering the firing angle of the thyristors the mean value of the rectified voltage can be varied, thereby allowing the motor speed to be controlled.

We saw in Chapter 2 that the controlled rectifier produces a crude form of d.c. with a pronounced ripple in the output voltage. This ripple component gives rise to pulsating currents and fluxes in the motor, and in order to avoid excessive eddy-current losses and commutation problems the poles and frame should be

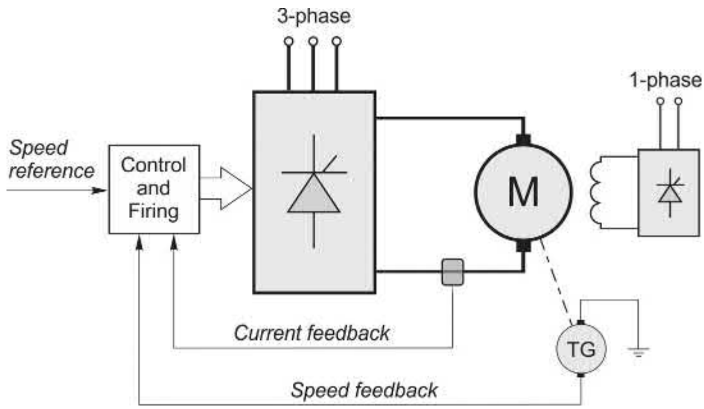

FIG. 4.1 Schematic diagram of speed-controlled d.c. motor drive.

1. Readers who are not familiar with feedback and control systems would find this book helpful: Schaum's Outline of Feedback and Control Systems, 2nd Edition (2013) by Joseph Distefano. McGraw-Hill Education; ISBN-10: 9780071829489

FIG. 4.2 High performance force-ventilated d.c. motor. The motor is of all-laminated construction and designed for use with a thyristor converter, and has a tacho mounted at the non-drive end. The small blower motor is an induction machine that runs continuously, thereby allowing the main motor to maintain full torque at low speed without overheating. *(Courtesy of Nidec Leroy Somer.)*

of laminated construction. It is accepted practice for motors supplied for use with thyristor drives to have laminated construction, but older motors may have solid poles and/or frames, and these will not always work satisfactorily with a rectifier supply. It is also the norm for d.c. motors for variable speed operation to be supplied with an attached 'blower' motor as standard (Fig. 4.2). This provides continuous through ventilation and allows the motor to operate continuously at full torque without overheating, even down to the lowest speeds.

Low power control circuits are used to monitor the principal variables of interest (usually motor current and speed), and to generate appropriate firing pulses so that the motor maintains constant speed despite variations in the load. The 'Speed Reference' (Fig. 4.1), was historically an analogue voltage varying from 0 to 10 V, obtained from a manual speed-setting potentiometer or from elsewhere in the plant, but now typically comes in digital form.

The combination of power, control and protective circuits constitutes the converter. Standard modular converters are available as off-the-shelf items in sizes from 100 W up to several hundred kW, while larger drives will be tailored to individual requirements. Individual converters may be mounted in enclosures with isolators, fuses, etc. or groups of converters may be mounted together to form a multi-motor drive.

4.2.1 Motor operation with converter supply

The basic operation of the rectifying bridge has been discussed in Chapter 2, and we now turn to the matter of how the d.c. motor behaves when supplied with 'd.c.' from a controlled rectifier.

By no stretch of the imagination could the waveforms of armature voltage looked at in Chapter 2 (e.g. Fig. 2.13) be thought of as *good* d.c., and it would not be unreasonable to question the wisdom of feeding such an unpleasant-looking waveform to a d.c. motor. In fact it turns out that the motor works almost as well as it would if fed with pure d.c., for two main reasons. Firstly, the armature inductance of the motor causes the waveform of armature current to be much smoother than the waveform of armature voltage, which in turn means that the torque ripple is much less than might have been feared. And secondly, the inertia of the armature (and load) is sufficiently large for the speed to remain almost steady despite the torque ripple. It is indeed fortunate that such a simple arrangement works so well, because any attempt to smooth-out the voltage waveform (perhaps by adding smoothing capacitors) would prove to be prohibitively expensive in the power ranges of interest.

4.2.2 Motor current waveforms

For the sake of simplicity we will look at operation from a single-phase (two-pulse) converter, but similar conclusions apply to the six-pulse one. The voltage (V_a) applied to the motor armature is typically as shown in Fig. 4.3A: as we saw in Chapter 2, it consists of rectified 'chunks' of the incoming utility supply voltage, the precise shape and average value depending on the firing angle.

The voltage waveform can be considered to consist of a mean d.c. level (V_{dc}), and a superimposed pulsating or ripple component which we can denote loosely as v_{ac}. If the supply is at 50 Hz, the fundamental frequency of the ripple is 100 Hz. The mean voltage V_{dc} can be altered by varying the firing angle, which also incidentally alters the ripple (i.e. v_{ac}).

The ripple voltage causes a ripple current to flow in the armature, but because of the armature inductance, the amplitude of the ripple current is small. In other words, the armature presents a high impedance to a.c. voltages. This

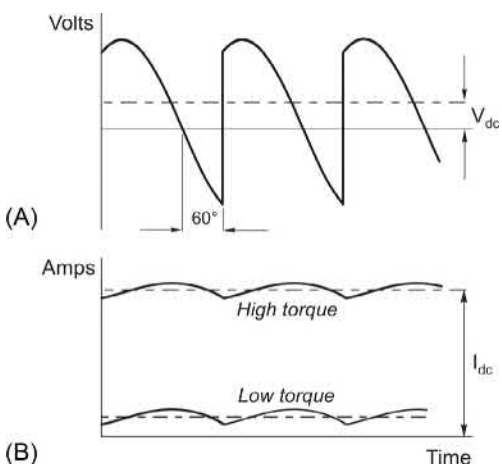

FIG. 4.3 Armature voltage (A) and armature current (B) waveforms for continuous-current operation of a d.c. motor supplied from a single-phase fully-controlled thyristor converter, with firing angle of 60 degrees.

smoothing effect of the armature inductance is shown in Fig. 4.3B, from which it can be seen that the current ripple is relatively small in comparison with the corresponding voltage ripple. The average value of the ripple current is of course zero, so it has no effect on the average torque of the motor. There is nevertheless a variation in torque every half-cycle of the mains, but because it is of small amplitude and high frequency (100 or 120 Hz for the 1-phase case here, but 300 or 360 Hz for 3-phase, (both for 50 or 60 Hz supplies respectively)) the variation in speed (and hence back e.m.f., E) will not usually be noticeable.

The current at the end of each pulse is the same as at the beginning, so it follows that the average voltage across the armature inductance (L) is zero. We can therefore equate the average applied voltage to the sum of the back e.m.f. (assumed pure d.c. because we are ignoring speed fluctuations) and the average voltage across the armature resistance, to yield

$$V_{dc} = E + I_{dc}R \tag{4.1}$$

which is exactly the same as for operation from a pure d.c. supply. This is very important, as it underlines the fact that we can control the mean motor voltage, and hence the speed, simply by varying the converter delay angle.

The smoothing effect of the armature inductance is important in achieving successful motor operation: the armature acts as a low-pass filter, blocking most of the ripple, and leading to a more or less constant armature current. For the smoothing to be effective, the armature time-constant needs to be long compared with the pulse duration (half a cycle with a two-pulse drive, but only one sixth of a cycle in a six-pulse drive). This condition is met in all almost all six-pulse drives, and in many two-pulse ones. Overall, the motor then behaves much as it would if it was supplied from an ideal d.c. source (though the I^2R loss is higher than it would be if the current was perfectly smooth).

The no-load speed is determined by the applied voltage (which depends on the firing angle of the converter); there is a small drop in speed with load; and, as we have previously noted, the average current is determined by the load. In Fig. 4.3, for example, the voltage waveform in (A) applies equally for the two load conditions represented in (B), where the upper current waveform corresponds to a high value of load torque while the lower is for a much lighter load, the speed being almost the same in both cases. (The small difference in speed is due to the IR term, as explained in Chapter 3.) We should note that the current ripple remains the same—only the average current changes with load. Broadly speaking, therefore, we can say that the speed is determined by the converter firing angle, which represents a very satisfactory state of affairs because we can control the firing angle by low-power control circuits and thereby regulate the speed of the drive.

The current waveforms in Fig. 4.3B are referred to as 'continuous', because there is never any time during which the current is not flowing. This 'continuous current' condition is the norm in most drives, and it is highly desirable because it is only under continuous current conditions that the average voltage from the converter is determined solely by the firing angle, and is independent of the load current. We can see why this is so with the aid of Fig. 2.8, imagining that the motor is connected to the output terminals and that it is drawing a continuous current. For half of a complete cycle, the current will flow into the motor from T1 and return to the supply via T4, so the armature is effectively switched across the supply and the armature voltage is equal to the supply voltage, which is assumed to be ideal, i.e. it is independent of the current drawn. For the other half of the time, the motor current flows from T2 and returns to the supply via T3, so the motor is again hooked-up to the supply, but this time the connections are reversed. Hence the average armature voltage—and hence to a first approximation the speed—are defined once the firing angle is set.

4.2.3 Discontinuous current

We can see from Fig. 4.3B that as the load torque is reduced, there will come a point where the minima of the current ripple touch the zero-current line, i.e. the current reaches the boundary between continuous and discontinuous current. The load at which this occurs will also depend on the armature inductance, because the higher the inductance the smoother the current (i.e. the less the ripple). Discontinuous current mode is therefore most likely to be encountered in small machines with low inductance (particularly when fed from two-pulse converters) and under light-load or no-load conditions.

Typical armature voltage and current waveforms in the discontinuous mode are shown in Fig. 4.4, the armature current consisting of discrete pulses of current that occur only while the armature is connected to the supply, with zero current for the period (represented by θ in Fig. 4.3) when none of the thyristors are conducting and the motor is coasting free from the supply.

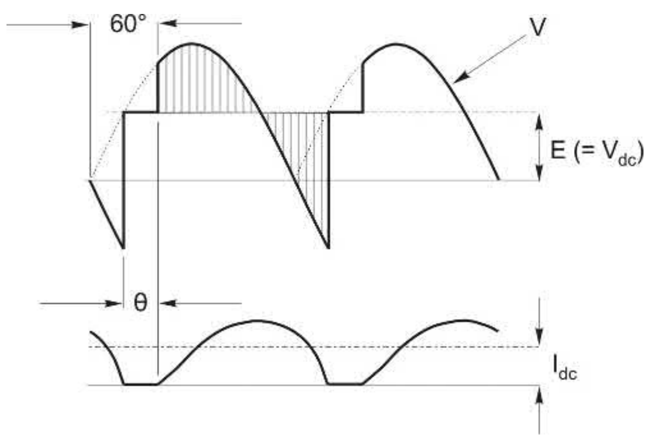

FIG. 4.4 Armature voltage and current waveforms for discontinuous-current operation of a d.c. motor supplied from a single-phase fully-controlled thyristor converter, with firing angle of 60 degrees.

The shape of the current waveform can be understood by noting that with resistance neglected, Eq. 3.7 can be rearranged as

$$\frac{di}{dt} = \frac{1}{L}(V - E) \tag{4.2}$$

which shows that the rate of change of current (i.e. the gradient of the lower graph in Fig. 4.4) is determined by the instantaneous difference between the applied voltage V and the motional e.m.f. E. Values of (V–E) are shown by the vertical hatchings in Fig. 4.4, from which it can be seen that if $V > E$, the current is increasing, while if $V < E$, the current is falling. The peak current is thus determined by the area of the upper or lower shaded areas of the upper graph.

The firing angle in Figs. 4.3 and 4.4 is the same, at 60 degrees, but the load is less in Fig. 4.4 and hence the average current is lower (though, for the sake of the explanation offered below, the current axis in Fig. 4.4 is expanded as compared with that in Fig. 4.2). It should be clear by comparing these figures that the armature voltage waveforms (solid lines) differ because, in Fig. 4.4, the current falls to zero before the next firing pulse arrives and during the period shown as θ the motor floats free, its terminal voltage during this time being simply the motional e.m.f. (E). To simplify Fig. 4.4 it has been assumed that the armature resistance is small and that the corresponding volt-drop ($I_a R_a$) can be ignored. In this case, the average armature voltage (V_{dc}) must be equal to the motional e.m.f., because there can be no average voltage across the armature inductance when there is no net change in the current over one pulse: the hatched areas—representing the volt-seconds in the inductor—are therefore equal.

The most important difference between Figs. 4.3 and 4.4 is that the average voltage is higher when the current is discontinuous, and hence the speed corresponding to the conditions in Fig. 4.4 is higher than in Fig. 4.3 despite both having the same firing angle. And whereas in continuous mode a load increase can be met by an increased armature current without affecting the voltage (and hence speed), the situation is very different when the current is discontinuous. In the latter case, the only way that the average current can increase is for the speed (and hence E) to fall so that the shaded areas in Fig. 4.4 become larger.

This means that the behaviour of the motor in discontinuous mode is much worse than in the continuous current mode, because as the load torque is increased, there is a serious drop in speed. The resulting torque-speed curve therefore has a very unwelcome 'droopy' characteristic in the discontinuous current region, as shown in Fig. 4.5, and in addition the I^2R loss is much higher than it would be with pure d.c.

Under very light or no-load conditions, the pulses of current become virtually non-existent, the shaded areas in Fig. 4.4 become very small and the motor speed approaches that at which the back e.m.f. is equal to the peak of the supply voltage (point (C) in Fig. 4.5).

It is easy to see that inherent torque-speed curves with sudden discontinuities of the form shown in Fig. 4.5 are very undesirable. If, for example, the

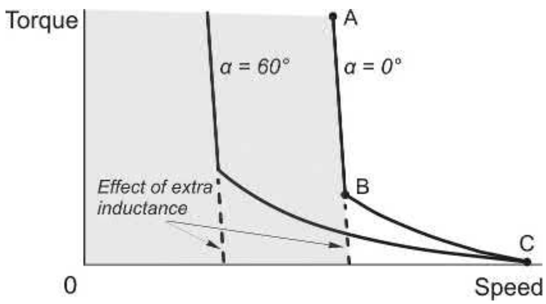

FIG. 4.5 Torque-speed curves illustrating the undesirable 'droopy' characteristic associated with discontinuous current. The improved characteristic (shown dotted) corresponds to operation with continuous current.

firing angle is set to zero and the motor is fully loaded, its speed will settle at point A, its average armature voltage and current having their full (rated) values. As the load is reduced, current remaining continuous, there is the expected slight rise in speed, until point B is reached. This is the point at which the current is about to enter the discontinuous phase. Any further reduction in the load torque then produces a wholly disproportionate—not to say frightening—increase in speed, especially if the load is reduced to zero, when the speed reaches point C.

There are two ways by which we can improve these inherently poor characteristics. Firstly, we can add extra inductance in series with the armature to further smooth the current waveform and lessen the likelihood of discontinuous current. The effect of adding inductance is shown by the dotted lines in Fig. 4.5. Secondly, we can switch from a single-phase converter to a three-phase one which produces smoother voltage and current waveforms, as discussed in Chapter 2.

When the converter and motor are incorporated in a closed-loop drive system the user should be unaware of any shortcomings in the inherent motor/converter characteristics because the control system automatically alters the firing angle to achieve the target speed at all loads. In relation to Fig. 4.5, for example, as far as the user is concerned the control system will confine operation to the shaded region, and the fact that the motor is theoretically capable of running unloaded at the high speed corresponding to point C is only of academic interest.

It is worth mentioning that discontinuous current operation is not restricted to the thyristor converter, but occurs in many other types of power electronic system. Broadly speaking, converter operation is more easily understood and analysed when in continuous current mode, and the operating characteristics are more desirable, as we have seen here. We will not dwell on the discontinuous mode in the rest of the book, as it is beyond our scope, and unlikely to be of concern to the drive user.

4.2.4 Converter output impedance: Overlap

So far we have tacitly assumed that the output voltage from the converter was independent of the current drawn by the motor, and depended only on the delay angle α. In other words we have treated the converter as an ideal voltage source.

In practice the a.c. supply has a finite impedance, and we must therefore expect a volt-drop which depends on the current being drawn by the motor. Perhaps surprisingly, the supply impedance (which is mainly due to inductive leakage reactances in transformers) manifests itself at the output stage of the converter as a supply resistance, so the supply volt-drop (or regulation) is directly proportional to the motor armature current.

It is not appropriate to go into more detail here, but we should note that the effect of the inductive reactance of the supply is to delay the transfer (or commutation) of the current between thyristors, a phenomenon known as overlap. The consequence of overlap is that instead of the output voltage making an abrupt jump at the start of each pulse, there is a short period where two thyristors are conducting simultaneously. (The reader should not be alarmed at the mention of two devices conducting at the same time: they are not in the same arm, so there is not a short circuit path across the d.c. that we warned of in Chapter 2 when discussing the d.c. to a.c. inverter.) During this interval the output voltage is the mean of the voltages of the incoming and outgoing voltages, as shown typically in Fig. 4.6. When the drive is connected to a 'stiff' (i.e. low impedance) industrial supply the overlap will only last for perhaps a few microseconds, so the 'notch' shown in Fig. 4.6 would be barely visible on an oscilloscope. Books always exaggerate the width of the overlap for the sake of clarity, as in Fig. 4.6: with a 50 or 60 Hz supply, if the overlap lasts for more than say 1 ms, the implication is that the supply system impedance is too high for the size of converter in question, or conversely, the converter is too big for the supply.

Returning to the practical consequences of supply impedance, we simply have to allow for the presence of an extra 'source resistance' in series with the output voltage of the converter. This source resistance is in series with

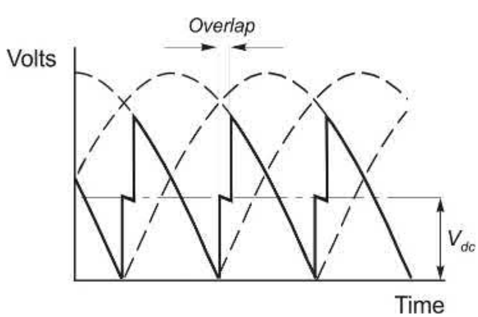

FIG. 4.6 Distortion of converter output voltage waveform caused by rectifier overlap.

the motor armature resistance, and hence the motor torque-speed curves for each value of α have a somewhat steeper droop than they would if the supply impedance was zero. However, as part of a closed loop control system, the drive would automatically compensate for any speed droop resulting from overlap.

4.2.5 Four-quadrant operation and inversion

So far we have looked at the converter as a rectifier, supplying power from the a.c. utility supply to a d.c. machine running in the positive direction and acting as a motor. As explained in Chapter 3, this is known as 1-quadrant operation, by reference to quadrant 1 of the complete torque-speed plane shown in Fig. 3.12.

But suppose we want to run as a motor in the opposite direction, with negative speed and negative torque, i.e. in quadrant 3. How do we do it? And what about operating the machine as a generator, so that power is returned to the a.c. supply, the converter then 'inverting' power rather than rectifying, and the system operating in quadrant 2 or quadrant 4. We need to do this if we want to achieve regenerative braking. Is it possible, and if so how?

The good news is that as we saw in Chapter 3 the d.c. machine is inherently a bi-directional energy converter. If we apply a positive voltage V greater than E, a current flows into the armature and the machine runs as a motor. If we reduce V so that it is less than E, the current, torque and power automatically reverse direction, and the machine acts as a generator, converting mechanical energy (its own kinetic energy in the case of regenerative braking) into electrical energy. And if we want to motor or generate with the reverse direction of rotation, all we have to do is to reverse the polarity of the armature supply. The d.c. machine is inherently a four-quadrant device, but needs a supply which can provide positive or negative voltage, and simultaneously handle either positive or negative current.

This is where we meet a snag: a single thyristor converter can only handle current in one direction, because the thyristors are unidirectional devices. This does not mean that the converter is incapable of returning power to the supply however. The d.c. current can only be positive, but (provided it is a fully-controlled converter) the d.c. output voltage can be either positive or negative (see Chapter 2). The power flow can therefore be positive (rectification) or negative (inversion).

For normal motoring where the output voltage is positive (and assuming a fully-controlled converter), the delay angle (α) will be up to 90 degrees. (It is common practice for the firing angle corresponding to rated d.c. voltage to be around 20 degrees when the incoming a.c. voltage is normal: then if the a.c. voltage falls for any reason, the firing angle can be further reduced to compensate and still allow full "rated" d.c. voltage to be maintained.)

When α is greater than 90 degrees, however, the average d.c. output voltage is negative, as indicated by Eq. (2.5), and shown in Fig. 4.7. A single

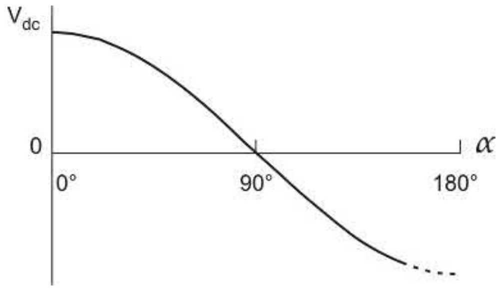

FIG. 4.7 Average d.c. output voltage from a fully-controlled thyristor converter as a function of the firing angle delay α.

fully-controlled converter therefore has the potential for two-quadrant operation, though it has to be admitted that this capability is not easily exploited unless we are prepared to employ reversing switches in the armature or field circuits. This is discussed next.

4.2.6 Single-converter reversing drives

We will consider a fully-controlled converter supplying a permanent-magnet motor, and see how the motor can be regeneratively braked from full speed in one direction, and then accelerated up to full speed in reverse. We looked at this procedure in principle at the end of Chapter 3, but here we explore the practicalities of achieving it with a converter-fed drive. We should be clear from the outset that in practice, all the user has to do is to change the speed reference signal from full forward to full reverse: the control system in the drive converter takes care of matters from then on. What it does, and how, is discussed below.

When the motor is running at full speed forward, the converter delay angle will be small, and the converter output voltage V and current I will both be positive. This condition is shown in Fig. 4.8A, and corresponds to operation in quadrant 1.

In order to brake the motor, the torque has to be reversed. The only way this can be done is by reversing the direction of armature current. The converter can

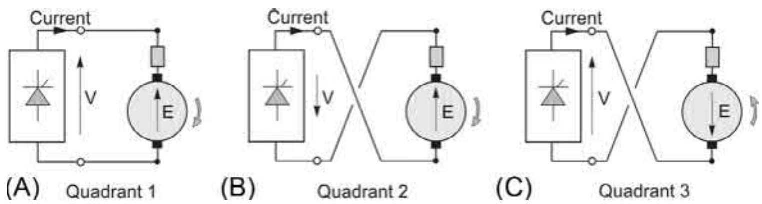

FIG. 4.8 Stages in motor reversal using a single-converter drive and mechanical reversing switch.

only supply positive current, however, so to reverse the motor torque we have to reverse the armature connections, using a mechanical switch or contactor, as shown in Fig. 4.8B. (Before operating the contactor, the armature current would be reduced to zero by lowering the converter voltage, so that the contactor is not required to interrupt current.) Note that because the motor is still rotating in the positive direction, the back e.m.f. remains in its original sense; but now the motional e.m.f. is seen to be assisting the current and so to keep the current within bounds the converter must produce a negative voltage V which is just a little less than E. This is achieved by setting the delay angle at the appropriate point between 90 and 180degrees. (The dotted line in Fig. 4.7 indicates that the maximum acceptable negative voltage will generally be somewhat less than the maximum positive voltage: this restriction arises because of the need to preserve a margin for commutation of current between thyristors.) Note that the converter current is still positive (i.e. upwards in Fig. 4.8B), but the converter voltage is negative, and power is thus flowing back to the supply system. In this condition the system is operating in quadrant 2, and the motor is decelerating because of the negative torque. As the speed falls, E reduces, and so V must be reduced progressively to keep the current at full value. This is achieved automatically by the action of the current-control loop, which is discussed later.

The current (i.e. torque) needs to be kept negative in order to run up to speed in the reverse direction, but after the back e.m.f. changes sign (as the motor reverses) the converter voltage again becomes positive and greater than E, as shown in Fig. 4.8C. The converter is then rectifying, with power being fed into the motor, and the system is operating in quadrant 3.

Schemes using reversing contactors are not suitable where the reversing time is critical or where periods of zero torque are unacceptable, because of the delay caused by the mechanical reversing switch, which may easily amount to 200–400 ms. Field reversal schemes operate in a similar way, but reverse the field current instead of the armature current. They are even slower, up to 5 s, because of the relatively long time-constant of the field winding.

4.2.7 Double-converter reversing drives

Where full four-quadrant operation and rapid reversal is called for, two converters connected in anti-parallel are used, as shown in Fig. 4.9. One converter supplies positive current to the motor, while the other supplies negative current.

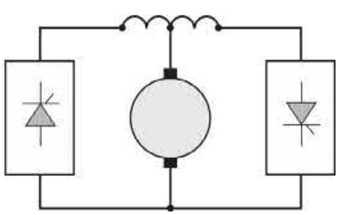

FIG. 4.9 Double-converter reversing drive.

The bridges are operated so that their d.c. voltages are almost equal thereby ensuring that any d.c. circulating current is small—When the bridges are operated together in this way, a reactor is almost always placed between the bridges to limit the flow of ripple currents which result from the unequal instantaneous voltages of the two converters.

In most applications, the reactor can be dispensed with and the converters operated one at a time. The changeover from one converter to the other can only take place after the firing pulses have been removed from one converter, and the armature current has decayed to zero. Appropriate zero-current detection circuitry is provided as an integral part of the drive, so that as far as the user is concerned, the two converters behave as if they were a single ideal bi-directional d.c. source. There is a dead (torque-free) time typically of only 10 ms or so during the changeover period from one bridge to the other.

Prospective users need to be aware of the fact that a basic single converter can only provide for operation in one quadrant. If regenerative braking is required, either field or armature reversing contactors will be needed; and if rapid reversal is essential, a double converter has to be used. All these extras naturally push up the purchase price.

4.2.8 Power factor and supply effects

One of the drawbacks of a converter-fed d.c. drive is that the supply power-factor is very low when the motor is operating at low speed (i.e. low armature voltage), and is less than unity even at base speed and full load. This is because the supply current waveform lags the supply voltage waveform by the delay angle α, as shown (for a three-phase converter) in Fig. 4.10, and also the supply current is approximately rectangular (rather than sinusoidal).

It is important to emphasise that the supply power-factor is always lagging, even when the converter is inverting. There is no way of avoiding the low power-factor, so users of large drives need to be prepared to augment their existing power-factor correcting equipment if necessary.

The harmonics in the utility supply current waveform can give rise to a variety of disturbance problems, and supply authorities generally impose statutory

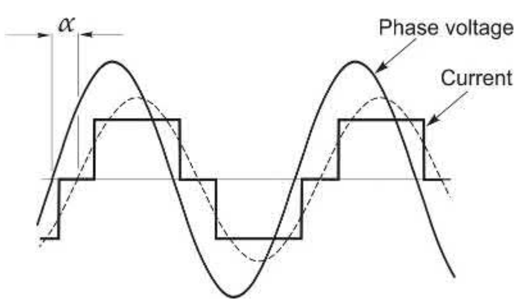

FIG. 4.10 Supply voltage and current waveforms for single-phase converter-fed d.c. motor drive.

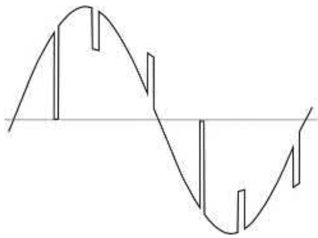

FIG. 4.11 Distortion of line voltage waveform caused by overlap in three-phase fully-controlled converter. (The width of the notches has been exaggerated for the sake of clarity.)

limits. For large drives (say hundreds of kW), filters may have to be provided to prevent these limits from being exceeded.

Since the supply impedance is never zero, there is also inevitably some distortion of the utility supply voltage waveform, as shown in Fig. 4.11, which indicates the effect of a six-pulse converter on the supply line-to-line voltage waveform. The spikes and notches arise because the utility supply is momentarily short-circuited each time the current commutates from one thyristor to the next, i.e. during the overlap period discussed earlier. For the majority of small and medium drives, connected to stiff industrial supplies, these notches are too small to be noticed (they are greatly exaggerated for the sake of clarity in Fig. 4.11); but they can cause serious disturbance to other consumers when a large drive is connected to a weak supply.

4.3 Control arrangements for d.c. drives

The most common arrangement, which is used with only minor variations from small drives of say 0.5 kW up to the largest industrial drives of several MW, is the so-called two-loop control. This has an inner feedback loop to control the current (and hence torque) and an outer loop to control speed. When position control is called for, a further outer position loop is added. A two-loop scheme for a thyristor d.c. drive is discussed first, but the essential features are the same in a chopper-fed drive. Later the simpler arrangements used in low-cost small drives are mentioned.

In order to simplify the discussion we will assume that the control signals are analogue, although in all modern versions the implementation will be digital, and we will limit consideration to those aspects which will be beneficial for the user to know something about. In practice, once a drive has been commissioned, there are only a few adjustments to which the user has access. Whilst most of them are self-explanatory (e.g. maximum speed, minimum speed, acceleration and deceleration rates), some are less obvious (e.g. ‘current stability’, ‘speed stability’, ‘IR comp’) so these are explained.

To appreciate the overall operation of a two-loop scheme we can consider what we would do if we were controlling the motor manually. For example, if we found by observing the tacho-generator output (usually referred to as a tacho) that the speed was below target, we would want to provide more current (and hence torque) in order to produce acceleration, so we would raise the armature voltage. We would have to do this gingerly however, being mindful of the danger of creating an excessive current because of the delicate balance that exists between the back e.m.f., E and applied voltage, V. We would doubtless wish to keep our eye on the ammeter at all times to avoid blowing-up the thyristors; and as the speed approached the target, we would trim back the current (by lowering the applied voltage) so as to avoid overshooting the set speed. Actions of this sort are carried out automatically by the drive system, which we will now explore.

A standard d.c. drive system with speed and current control is shown in Fig. 4.12. The primary purpose of the control system is to provide speed control, so the 'input' to the system is the speed reference signal on the left, and the output is the speed of the motor (as measured by a tacho) on the right. As with any closed-loop system, the overall performance is heavily dependent on the quality of the feedback signal, in this case the speed-proportional voltage provided by the tacho. It is therefore important to ensure that the tacho is of high quality (so that, for example, its output voltage does not vary with ambient temperature, and is ripple-free) and as a result the cost of the tacho often represents a significant fraction of the total cost.

We will take an overview of how the scheme operates first, and then examine the function of the two loops in more detail.

To get an idea of the operation of the system we will consider what will happen if, with the motor running light at a set speed, the speed reference signal is suddenly increased. Because the set (reference) speed is now greater than the actual speed there will be a speed error signal (see also Fig. 4.13), represented by the output of the left-hand summing junction in Fig. 4.12. A speed error indicates that acceleration is required, which in turn means more torque, i.e.

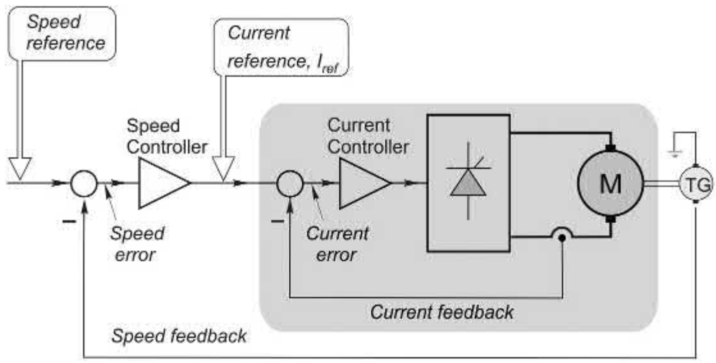

FIG. 4.12 Schematic diagram of analogue controlled-speed drive with current and speed feedback control loops.

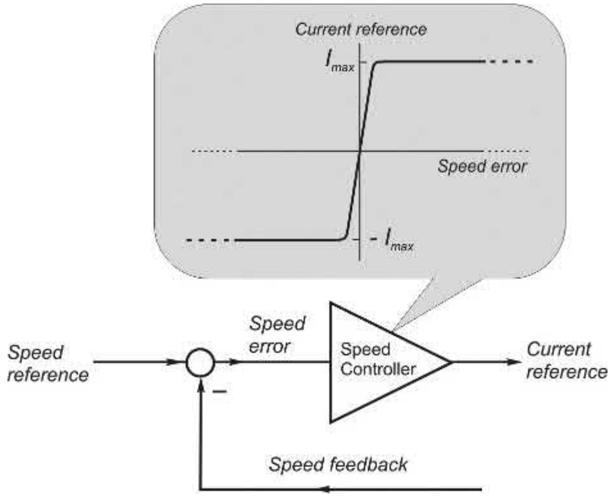

FIG. 4.13 Detail showing characteristic of speed error amplifier.

more current. The speed error is amplified by the speed controller (which is more accurately described as a speed-error amplifier) and the output serves as the reference or input signal to the inner control system. The inner feedback loop is a current-control loop, so when the current reference increases, so does the motor armature current, thereby providing extra torque and initiating acceleration. As the speed rises the speed error reduces and the current and torque therefore reduce to obtain a smooth approach to the target speed.

We will now look in more detail at the inner (current control) loop, as its correct operation is vital to ensure that the thyristors are protected against excessive current.

4.3.1 Current limits and protection

The closed-loop current controller, or current loop, is at the heart of the drive system and is indicated by the shaded region in Fig. 4.12. The purpose of the current loop is to make the actual motor current follow the current reference signal (I_{ref}) shown in Fig. 4.12. It does this by comparing a feedback signal of actual motor armature current with the current reference signal, amplifying the difference (the current error), and using the resulting current error signal to control the firing angle α—and hence the output voltage—of the converter. The current feedback signal is obtained either from a d.c. current transformer (which gives an isolated analogue voltage output), or from a.c. current transformer/rectifiers in the utility supply lines.

The job of comparing the reference (demand) and actual current signals and amplifying the error signal is carried out by the current-error amplifier. By giving the current-error amplifier a high gain, the actual motor current will always correspond closely to the current reference signal, i.e. the current-error will be small, regardless of motor speed. In other words, we can expect the actual motor current to follow the 'current reference' signal at all times, the armature voltage being automatically adjusted by the controller so that, regardless of the speed of the motor, the current has the correct value.

Of course no control system can be perfect, but it is usual for the current-error amplifier to be of the proportional plus integral (PI) type (see below), in which case the actual and demanded currents will be exactly equal under steady-state conditions.

The importance of preventing excessive converter currents from flowing has been emphasised previously, and the current control loop provides the means to this end. As long as the current control loop functions properly, the motor current can never exceed the reference value. Hence by limiting the magnitude of the current reference signal (by means of a clamping circuit), the motor current can never exceed the specified value. This is shown in Fig. 4.13, which represents a small portion of Fig. 4.12. The characteristics of the speed controller are shown in the shaded panel, from which we can see that for small errors in speed, the current reference increases in proportion to the speed error, thereby ensuring 'linear system' behaviour with a smooth approach to the target speed. However, once the speed error exceeds a limit, the output of the speed error amplifier saturates and there is thus no further increase in the current reference. By arranging for this maximum current reference to correspond to the full (rated) current of the system there is no possibility of the current in the motor and converter exceeding its rated value, no matter how large the speed error becomes.

This 'electronic current limiting' is by far the most important protective feature of any drive. It means that if for example the motor suddenly stalls because the load seizes (so that the back e.m.f. falls dramatically), the armature voltage will automatically reduce to a very low value, thereby limiting the current to its maximum allowable level.

The inner loop is critical in a two-loop control system, so the current loop must guarantee that the steady-state motor current corresponds exactly with the reference, and the transient response to step changes in the current reference should be fast and well damped. The first of these requirements is satisfied by the integral term in the current-error amplifier, while the second is obtained by judicious choice of the amplifier proportional gain and time-constant. As far as the user is concerned, a 'current stability' adjustment may be provided to allow him to optimise the transient response of the current loop.

On a point of jargon, it should perhaps be mentioned that the current-error amplifier is more often than not called either the 'current controller' (as in Fig. 4.12) or the 'current amplifier'. The first of these terms is quite sensible, but the second can be very misleading: there is after all no question of the motor current itself being amplified.

4.3.2 Torque control

For applications requiring the motor to operate with a specified torque regardless of speed (e.g. in line tensioning), we can dispense with the outer (speed) loop, and simply feed a current reference signal directly to the current controller. This is because torque is directly proportional to current, so the current controller is in effect also a torque controller. We may have to make an allowance for accelerating torque, by means of a transient 'inertia compensating' signal, which could simply be added to the torque demand.

In the current-control mode the current remains constant at the set value, and the steady running speed is determined by the load. If the torque reference signal was set at 50%, for example, and the motor was initially at rest, it would accelerate with a constant current of half rated value until the load torque was equal to the motor torque. Of course, if the motor was running without any load, it would accelerate quickly, the applied voltage ramping up so that it always remained higher than the back e.m.f. by the amount needed to drive the specified current into the armature. Eventually the motor would reach a speed (a little above normal 'full' speed) at which the converter output voltage had reached its upper limit, and it was therefore no longer possible to maintain the set current: thereafter, the motor speed would remain steady.

This discussion assumes that torque is proportional to armature current, which is true only if the flux is held constant, which in turn requires the field current to be constant. Hence in all but small drives the field will be supplied from a thyristor converter with current feedback. Variations in field circuit resistance due to temperature changes, and/or changes in the utility supply voltage are thereby automatically compensated and the flux is maintained at its rated value.

4.3.3 Speed control

The outer loop in Fig. 4.12 provides speed control. Speed feedback is typically provided by a d.c. tacho, and the actual and required speeds are fed into the speed-error amplifier (often known simply as the speed amplifier or the speed controller).

Any difference between the actual and desired speed is amplified, and the output serves as the input to the current loop. Hence if for example the actual motor speed is less than the desired speed, the speed amplifier will demand current in proportion to the speed error, and the motor will therefore accelerate in an attempt to minimise the speed error.

When the load increases, there is an immediate deceleration and the speed error signal increases, thereby calling on the inner loop for more current. The increased torque results in acceleration and a progressive reduction of the speed error until equilibrium is reached at the point where the current reference (I_{ref}) produces a motor current that gives a torque equal and opposite to the load torque. Looking at Fig. 4.13, where the speed controller is shown as a simple proportional amplifier (P control), it will be readily appreciated that in order for there to be a steady-state value of I_{ref}, there would have to be a finite speed error, i.e. a P controller would not allow us to reach exactly the target speed. (We could approach the ideal by increasing the gain of the amplifier, but that might lead us to instability.)

To eliminate the steady-state speed error we can easily arrange for the speed controller to have an integral (I) term as well as a proportional (P) term. A PI controller can have a finite output even when the input is zero, which means that we can achieve zero steady-state error if we employ PI control.

The speed will be held at the value set by the speed reference signal for all loads up to the point where full armature current is needed. If the load torque increases any more the speed will drop because the current-loop will not allow any more armature current to flow. Conversely, if the load attempted to force the speed above the set value, the motor current will be reversed automatically, so that the motor acts as a brake and regenerates power to the utility supply.

To emphasise further the vitally important protective role of the inner loop, we can see what happens when, with the motor at rest (and unloaded for the sake of simplicity), we suddenly increase the speed reference from zero to full value, i.e. we apply a step demand for full speed. The speed error will be 100%, so the output from the speed-error amplifier (I_{ref}) will immediately saturate at its maximum value, which has been deliberately clamped so as to correspond to a demand for the maximum (rated) current in the motor. The motor current will therefore be at rated value, and the motor will accelerate at full torque. Speed and back e.m.f. (E) will therefore rise at a constant rate, the applied voltage (V) increasing steadily so that the difference (V–E) is sufficient to drive rated current (I) through the armature resistance. A very similar sequence of events was discussed in Chapter 3, and illustrated by the second half of Fig. 3.13. (In some drives the current reference is allowed to reach 150% or even 200% of rated value for a few seconds, in order to provide a short torque-boost. This is particularly valuable in starting loads with high static friction, and is known as a 'two-stage current limit').

The output of the speed amplifier will remain saturated until the actual speed is quite close to the target speed, and for all this time the motor current will therefore be held at full value. Only when the speed is within a few per cent of target will the speed error amplifier come out of saturation. Thereafter, as the speed continues to rise, and the speed error falls, the output of the speed error amplifier falls below the clamped level. Speed control then enters a linear regime, in which the correcting current (and hence the torque) is proportional to speed error, giving a smooth approach to final speed.

A 'good' speed controller will result in zero steady-state error, and have a well-damped response to step changes in the demanded speed. The integral term in the PI control caters for the requirement of zero steady-state error, while the transient response depends on the setting of the proportional gain and time-constant. The 'speed stability' setting (traditionally a potentiometer) is provided to allow fine tuning of the transient response. It should be mentioned that in some high-performance drives the controller will be of the PID form, i.e. it will also include a differential (D) term. The D term gives a bit of a kick to the controllers when a step change is called for—in effect advanced warning that we need a bit of smart action to change the current in the inductive circuit.

It is important to remember that it is much easier to obtain a good transient response with a regenerative drive, which has the ability to supply negative current (i.e. braking torque) should the motor overshoot the desired speed. A non-regenerative drive cannot furnish negative current (unless fitted with reversing contactors), so if the speed overshoots the target the best that can be done is to reduce the armature current to zero and wait for the motor to decelerate naturally. This is not satisfactory, and every effort therefore has to be made to avoid controller settings which lead to an overshoot of the target speed.

As with any closed-loop scheme, problems occur if the feedback signal is lost when the system is in operation. If the tacho feedback became disconnected, the speed amplifier would immediately saturate, causing full torque to be applied. The speed would then rise until the converter output reached its maximum output voltage. To guard against this many drives incorporate tacho-loss detection circuitry, and in some cases armature voltage feedback (see later section) automatically takes over in the event of tacho failure.

Drives which use field-weakening to extend the speed range include automatic provision for controlling both armature voltage and field current when running above base speed. Typically, the field current is kept at full value until the armature voltage reaches about 95% of rated value. When a higher speed is demanded, the extra armature voltage applied is accompanied by a simultaneous reduction in the field current, in such a way that when the armature voltage reaches 100% the field current is at the minimum safe value.

4.3.4 Overall operating region

A standard drive with field-weakening provides armature voltage control of speed up to base speed, and field-weakening control of speed thereafter. Any torque up to the rated value can be obtained at any speed below base speed, and as explained in Chapter 3 this region is known as the 'constant torque' region. Above base speed, the maximum available torque reduces inversely with speed, so this is known as the 'constant power' region. For a converter-fed drive the operating region in quadrant 1 of the torque-speed plane is therefore as shown in Fig. 3.10. (If the drive is equipped for regenerative and reversing operation, the operating area is mirrored in the other three quadrants, of course.)

4.3.5 Armature voltage feedback and IR compensation

In low-power drives where precision speed-holding is not essential, and cost must be kept to a minimum, the tacho may be dispensed with and the armature voltage used as a 'speed feedback' instead. Performance is clearly not as good as with tacho feedback, since whilst the steady-state no-load speed is proportional to armature voltage, the speed falls as the load (and hence armature current) increases.

We saw in Chapter 3 that the drop in speed with load was attributable to the armature resistance volt-drop (IR), and the steady-state drop in speed can therefore be compensated by boosting the applied voltage in proportion to the current. 'IR compensation' would in such cases be provided on the drive circuit for the user to adjust to suit the particular motor. The compensation is far from perfect, since it cannot cope well with temperature variation of resistance, nor with load transients.

4.3.6 Drives without current control

Very low cost, low power drives may dispense with a full current control loop, and incorporate a crude 'current-limit' which only operates when the maximum set current would otherwise be exceeded. These drives usually have an in-built ramp circuit which limits the rate of rise of the set speed signal so that under normal conditions the current limit is not activated. They are however prone to tripping in all but the most controlled of applications and environments.

4.4 Chopper-fed d.c. motor drives

If the source of supply is d.c. (e.g. in a battery vehicle or a rapid transit system) a chopper-type converter is usually employed. The basic operation of a single-switch chopper was discussed in Chapter 2, where it was shown that the average output voltage could be varied by periodically switching the battery voltage on and off for varying intervals. The principal difference between the thyristor-controlled rectifier and the chopper is that in the former the motor current always flows through the supply, whereas in the latter, the motor current only flows from the supply terminals for part of each cycle.

A single-switch chopper using a transistor, MOSFET or IGBT, can only supply positive voltage and current to a d.c. motor, and is therefore restricted to one-quadrant motoring operation. When regenerative and/or rapid speed reversal is called for, more complex circuitry is required, involving two or more power switches, and consequently leading to increased cost. Many different circuits are used and it is not possible to go into detail here, but it will be recalled that the two most important types were described in Section 2.2.2: the simplest or 'buck' converter provides an output voltage in the range $0 < E$, where E is the battery voltage, while the slightly more complex ‘boost’ converter, Section 2.2.3, provides output voltages greater than that of the supply.

4.4.1 Performance of chopper-fed d.c. motor drives

We saw earlier that the d.c. motor performed almost as well when fed from a phase-controlled rectifier as it does when supplied with pure d.c. The chopper-fed motor is, if anything, rather better than the phase-controlled, because the armature current ripple can be less if a high chopping frequency is used.

A typical circuit and waveforms of armature voltage and current are shown in Fig. 4.14: these are drawn with the assumption that the switch is ideal. The rather modest chopping frequency of 100 Hz shown in Fig. 4.14 has been deliberately chosen so that the current ripple is significant in relation to the average current (I_{dc}), and so that the relationship to the voltage waveform is clearly shown. In practice, most chopper drives operate at much higher frequencies, and the ripple current is therefore much less significant. As usual, we have assumed that the speed remains constant despite the slightly pulsating torque, and that the armature current is continuous.

The shape of the armature voltage waveform reminds us that when the transistor is switched on, the battery voltage V is applied directly to the armature, and during this period the path of the armature current is indicated by the dotted line in Fig. 4.14A. For the remainder of the cycle the transistor is turned ‘off’ and the current freewheels through the diode, as shown by the dotted line in Fig. 4.14B. When the current is freewheeling through the diode, the armature voltage is clamped at (almost) zero.

The speed of the motor is determined by the average armature voltage, (V_{dc}), which in turn depends on the proportion of the total cycle time (T) for which the transistor is ‘on’. If the on and off times are defined as $T_{on}=kT$ and $T_{off}=(1-kT)$, where $0<k<1$, then the average voltage is simply given by

$$V_{dc}=kV \tag{4.3}$$

from which we see that speed control is effected via the on time ratio, k.

Turning now to the current waveforms shown in Fig. 4.14C, the upper waveform corresponds to full load, i.e. the average current (I_{dc}) produces the full rated torque of the motor. If now the load torque on the motor shaft is reduced to half rated torque, and assuming that the resistance is negligible, the steady-state speed will remain the same but the new mean steady-state current will be halved, as shown by the lower dotted curve. We note however that although, as expected, the mean current is determined by the load, the ripple current is unchanged, and this is explained below.

If we ignore resistance, the equation governing the current during the ‘on’ period is

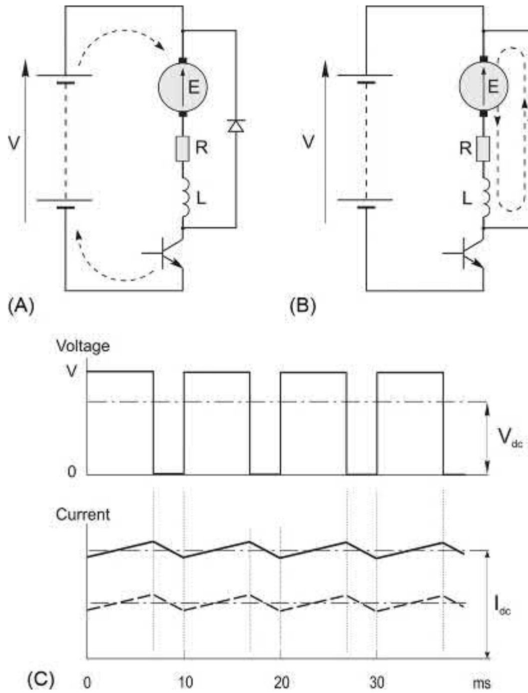

FIG. 4.14 Chopper-fed d.c. motor. In (A) the transistor is 'ON' and armature current is flowing through the voltage source; in (B) the transistor is 'OFF' and the armature current freewheels through the diode. Typical armature voltage and current waveforms are shown at (C), with the dotted line representing the current waveform when the load torque is reduced by half.

$$V = E + L\frac{di}{dt}, \quad \text{or} \; \frac{di}{dt} = \frac{1}{L}(V - E) \tag{4.4}$$

Since V is greater than E, the gradient of the current (di/dt) is positive, as can be seen in Fig. 4.14C. During this 'on' period the battery is supplying power to the motor. Some of the energy is converted to mechanical output power, but some is also stored in the magnetic field associated with the inductance. The latter is given by $\frac{1}{2}Li^2$, and so as the current (i) rises, more energy is stored.

During the 'off' period, the equation governing the current is

$$0 = E + L\frac{di}{dt}, \quad \text{or} \quad \frac{di}{dt} = \frac{-E}{L} \tag{4.5}$$

We note that during the 'off' time the gradient of the current is negative (as shown in Fig. 4.14C) and it is determined by the motional e.m.f. E. During this period, the motor is producing mechanical output power which is supplied from the energy stored in the inductance, so not surprisingly the current falls as the energy previously stored in the on period is now given up.

We note that the rise and fall of the current (i.e. the current ripple) is inversely proportional to the inductance, but is independent of the mean d.c. current, i.e. the ripple does not depend on the load.

To study the input/output power relationship, we note that the battery current only flows during the on period, and its average value is therefore kI_{dc}. Since the battery voltage is constant, the power supplied is simply given by $V(kI_{dc})=kVI_{dc}$. Looking at the motor side, the average voltage is given by $V_{dc}=kV$, and the average current (assumed constant) is I_{dc}, so the power input to the motor is again kVI_{dc} i.e. there is no loss of power in the ideal chopper. Given that k is <1, we see that the input (battery) voltage is higher than the output (motor) voltage, but conversely the input current is less than the output current, and in this respect we see that the chopper behaves in much the same way for d.c. as a conventional transformer does for a.c.

4.4.2 Torque-speed characteristics and control arrangements

Under open-loop conditions (i.e. where the mark-space ratio of the chopper is fixed at a particular value) the behaviour of the chopper-fed motor is similar to the converter-fed motor discussed earlier (see Fig. 4.4). When the armature current is continuous the speed falls only slightly with load, because the mean armature voltage remains constant. But when the armature current is discontinuous (which is most likely at high speeds and light load) the speed falls off rapidly when the load increases, because the mean armature voltage falls as the load increases. Discontinuous current can be avoided by adding an inductor in series with the armature, or by raising the chopping frequency, but when closed-loop speed control is employed, the undesirable effects of discontinuous current are masked by the control loop.

The control philosophy and arrangements for a chopper-fed motor are the same as for the converter-fed motor, with the obvious exception that the mark-space ratio of the chopper is used to vary the output voltage, rather than the firing angle.

4.5 D.C. servo drives

The precise meaning of the term 'servo' in the context of motors and drives is difficult to pin down. Broadly speaking, if a drive incorporates 'servo' in its description, the implication is that it is intended specifically for a high performance application and employs closed-loop or feedback control, usually of shaft torque, speed and often position. Early servomechanisms were developed primarily for military applications, and it quickly became apparent that standard d.c. motors were not always suited to precision control. In particular high torque to inertia ratios were needed, together with smooth ripple-free torque. Motors were therefore developed to meet these exacting requirements, and not surprisingly they were, and still are, much more expensive than their industrial counterparts. Whether the extra expense of a servo motor can be justified depends on the specification, but prospective users should always be on their guard to ensure they are not pressed into an expensive purchase when a conventional industrial drive could cope perfectly well.

The majority of servo drives are sold in modular form, consisting of a high-performance permanent magnet motor, often with an integral tacho, and a chopper-type power amplifier module. The drive amplifier normally requires a separate regulated d.c. power supply, if, as is normally the case, the power is to be drawn from the utility supply. Continuous output powers range from a few Watts up to perhaps 2–5 kW, with voltages of 12, 24, 48 and multiples of 50 V being standard: higher powers are available, but they are not so common.

There has been an even more pronounced movement in the servo market than in industrial drives away from d.c. in favour of the a.c. permanent magnet or induction motors, although d.c. servos do retain some niche applications.

4.5.1 Servo motors

Although there is no sharp dividing line between servo motors and ordinary motors, the servo type will usually be intended for use in applications which require rapid acceleration and deceleration. The design of the motor will reflect this by catering for intermittent currents (and hence torques) of many times the continuously rated value. Because most servo motors are small, their armature resistances are relatively high: the short-circuit (locked-rotor) current at full armature voltage is therefore perhaps only a few times the continuously rated current, and the drive amplifier will normally be selected so that it can cope with this condition, giving the motor a very rapid acceleration from rest. The even more arduous condition in which the full armature voltage is suddenly reversed with the motor running at full speed is also quite normal. (Both of these modes of operation would of course be quite unthinkable with a large d.c. motor, because of the huge currents which would flow as a result of the disproportionately low armature resistance.)

In order to maximise acceleration, the rotor inertia must be minimised, and one obvious way to achieve this is to construct a motor in which only the electric circuit (conductors) on the rotor move, the magnetic part (either iron or permanent magnet) remaining stationary. This principle is adopted in 'ironless rotor' and 'printed armature' motors.

In the ironless rotor or moving-coil type (Fig. 4.15) the armature conductors are formed as a thin-walled cylinder consisting essentially of nothing more than varnished wires wound in skewed form together with the disc-type commutator

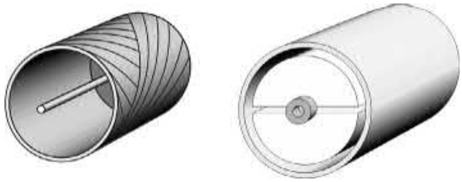

FIG. 4.15 Ironless rotor d.c. motor. The commutator (not shown) is usually of the disc type.

(not shown). Inside the armature sits a two-pole (upper N, lower S) permanent magnet which provides the radial flux, and outside it is a steel cylindrical shell which completes the magnetic circuit.

Needless to say the absence of any slots to support the armature winding results in a relatively fragile structure, which is therefore limited to diameters of not much over 1 cm. Because of their small size they are often known as micromotors, and are very widely used in cameras, video systems, card readers, medical instruments etc.

The printed armature type is altogether more robust, and is made in sizes up to a few kW. They are generally made in disc or pancake form, with the direction of flux axial and the armature current radial. The armature conductors resemble spokes on a wheel, the conductors themselves being formed on a lightweight disc. Early versions were made by using printed-circuit techniques, but pressed fabrication is now more common. Since there are usually at least a hundred armature conductors, the torque remains almost constant as the rotor turns, which allows them to produce very smooth rotation at low speed. Inertia and armature inductance are low, giving a good dynamic response, and the short and fat shape makes them suitable for applications such as machine tools and disc drives where axial space is at a premium.

It is worth briefly mentioning that there is another type of servo motor, where extremely smooth rotation is the primary consideration. Here, rotor inertia is usually maximised. Such an applications would be a machine tool spindle, where any variation in the rotational speed during a cutting process could be reflected in surface imperfections in the finished product.

4.5.2 Position control

As mentioned earlier, many servo motors are used in closed-loop position control applications, so it is appropriate to look briefly at how this is achieved. Later (in Chapter 10) we will see that the stepping motor provides an alternative open-loop method of position control, which can be cheaper for some less demanding applications.

In the example shown in Fig. 4.16, the angular position of the output shaft is intended to follow the reference voltage (θ_{ref}), but it should be clear that if the

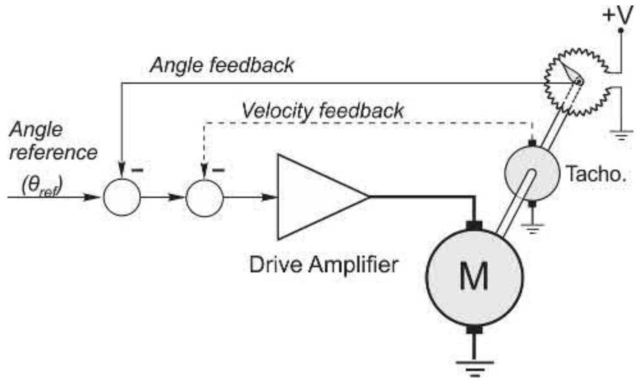

FIG. 4.16 Closed-loop angular position control using d.c. motor and angle feedback from a servo-type potentiometer.

motor drives a toothed belt linear outputs can also be obtained. The potentiometer mounted on the output shaft provides a feedback voltage proportional to the actual position of the output shaft. The voltage from this potentiometer must be a linear function of angle, and must not vary with temperature, otherwise the accuracy of the system will be in doubt.

The feedback voltage (representing the actual angle of the shaft) is subtracted from the reference voltage (representing the desired position) and the resulting position error signal is amplified and used to drive the motor so as to rotate the output shaft in the desired direction. When the output shaft reaches the target position, the position error becomes zero, no voltage is applied to the motor, and the output shaft remains at rest. Any attempt to physically move the output shaft from its target position immediately creates a position error and a restoring torque is applied by the motor.

The dynamic performance of the simple scheme described above is very unsatisfactory as it stands. In order to achieve a fast response and to minimise position errors caused by static friction, the gain of the amplifier needs to be high, but this in turn leads to a highly oscillatory response which is usually unacceptable. For some fixed-load applications matters can be improved by adding a compensating network at the input to the amplifier, but the best solution is to use 'tacho' (speed) feedback (shown dotted in Fig. 4.16) in addition to the main position feedback loop. Tacho feedback clearly has no effect on the static behaviour (since the voltage from the tacho is proportional to the speed of the motor), but has the effect of increasing the damping of the transient response. The gain of the amplifier can therefore be made high in order to give a fast response, and the degree of tacho feedback can then be adjusted to provide the required damping (see Fig. 4.17). Many servo motors have an integral tacho for this purpose. (This is a particular example of the general principle by which the response can be

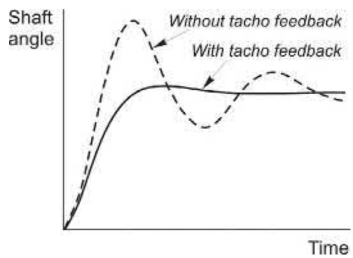

FIG. 4.17 Typical step responses for a closed-loop position control system, showing the improved damping obtained by the addition of tacho feedback.

improved by adding 'derivative of output' feedback: in this case the speed signal is the rate of change (or derivative) of the angular position.)

The example above dealt with an analogue scheme in the interests of simplicity, but digital position control schemes (with an encoder used to provide both position and speed information) are now much more common, especially when brushless motors (see Chapter 9) are used. Complete 'controllers on a card' are available as off-the-shelf items, and these offer ease of interface to other systems as well as providing for improved flexibility in shaping the dynamic response.

4.6 Digitally-controlled drives

As in all forms of industrial and precision control, digital implementations have replaced analogue circuitry in the vast majority of electric drive systems, but there are few instances where this has resulted in any real change to the basic structure of the drive with respect the motor control, and in most cases understanding how the drive functions is still best approached in the first instance by studying the analogue version. Digital control electronics has brought with it a considerable advance in the auxiliary control and protection functions which are now routinely found on a drive system. Digital control electronics has also facilitated the commercial implementation of advanced a.c. motor control strategies which will be discussed in Chapters 7–10. However, as far as understanding d.c. drives is concerned, users who have developed a sound understanding of how the analogue version operates will find little to trouble when considering the digital equivalent. Accordingly this section is limited to the consideration of a few of the advantages offered by digital implementations, and readers seeking more are recommended to consult a book such as the *Control Techniques Drives and Controls Handbook*, 2nd Edition, by W Drury.

Many drives use digital speed feedback, in which a pulse train generated from a shaft-mounted encoder is compared (using a phase-locked-loop) with a reference pulse train whose frequency corresponds to the desired speed, or where the reference is transmitted to the drive in the form of a synchronous serial word. Consequently, the feedback is more accurate and drift-free and noise in the encoder signal is easily rejected, so that very precise speed holding can be guaranteed. This is especially important when a number of independent motors must all be driven at identical speed. Phase-locked loops are also used in the firing-pulse synchronising circuits, to overcome the problems caused by noise on the utility supply waveform.

Digital controllers offer freedom from drift (the bugbear of analogue amplifier circuits), added flexibility (e.g. programmable ramp-up, ramp-down, maximum and minimum speeds etc.), ease of interfacing and linking to other drives and host computers and controllers, and self-tuning. User-friendly diagnostics represents another benefit, providing the local or remote user with current and historical data on the state of all the key drive variables. Digital drives also offer many more functions, including user programmable functions as are found on PLCs as well as a host of communications interfaces to allow incorporation into industrial automation systems.

4.7 Review questions

(1) A speed-controlled d.c. motor drive is running light at 50% of full speed. If the speed reference was raised to 100%, and the motor was allowed to settle, how would you expect the new steady-state values of armature voltage, tacho voltage and armature current to compare with the corresponding values when the motor was running at 50% speed?

(2) A d.c. motor drive has a PI speed controller. The drive is initially running at 50% speed with the motor unloaded. A load torque of 100% is then applied to the shaft. How would you expect the new steady-state values of armature voltage, tacho voltage and armature current to compare with the corresponding values before the load was applied?

(3) An unloaded d.c. motor drive is started from rest by applying a sudden 100% speed demand. How would you expect the armature voltage and current to vary as the motor runs up to speed?

(4) What would you expect to happen to a d.c. drive running with 50% torque at 50% speed if:-

- **(a)** the utility supply voltage fell by 10%;
- **(b)** the tacho wires were inadvertently pulled off;
- **(c)** the motor seized solid;
- **(d)** a short-circuit was placed across the armature terminals;
- **(e)** the current feedback signal was removed.

(5) Why is discontinuous operation generally undesirable in a d.c. motor?

(6) What is the difference between dynamic braking and regenerative braking?

(7) Explain why, in the drives context, it is often said that the higher the armature circuit inductance of d.c. machine, the better. In what sense is high armature inductance not desirable?

(8) The torque-speed characteristics shown in Fig. Q8 relate to a d.c. motor supplied from a fully-controlled thyristor converter.

Identify the axes. Indicate which parts of the characteristics display ‘good’ performance and which parts indicate ‘bad’ performance, and

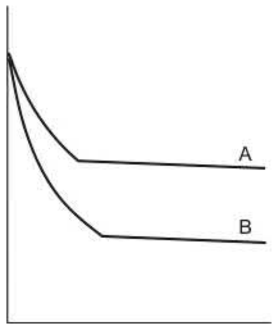

FIG. Q8

explain briefly what accounts for the abrupt change in behaviour. If curve A corresponds to a firing angle of 5 degrees, estimate the firing angle for curve B. How might the shape of the curves change if a substantial additional inductance were added in series with the armature of the motor?

(9) A 250 kW drive for a tube-mill drawbench had to be designed using two motors rated at 150 kW, 1200 rev/min and 100 kW, 1200 rev/min respectively, and coupled to a common shaft. Each motor was provided with its own speed-controlled drive. The specification called for the motors to share the load in proportion to their rating, so the controls were arranged as shown in Fig. Q9 (the load is not shown).

(a) Explain briefly why this scheme is referred to as a master/slave arrangement.

(b) How is load sharing achieved?

(c) Discuss why this arrangement is preferable to one in which both drives have active outer speed loops.

(d) Would there be any advantage in feeding the current reference for the smaller drive from the current feedback signal of the larger drive?

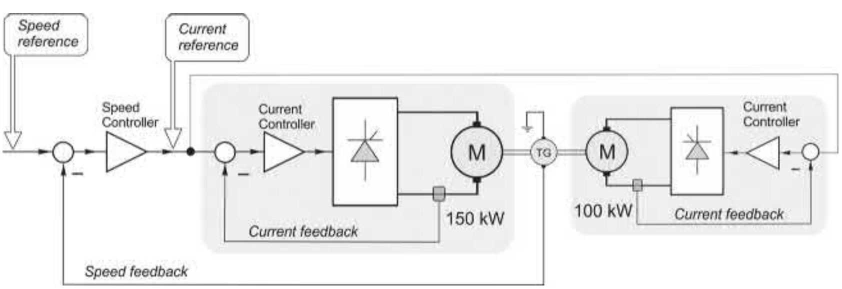

FIG. Q9

Answers to the review questions are given in the Appendix.

Chapter 5

Induction motors—Rotating field, slip and torque

5.1 Introduction

Judged in terms of fitness for purpose coupled with simplicity, the induction motor must rank alongside the screwthread as one of mankind's best inventions. It is not only supremely elegant as an electromechanical energy converter, but is also by far the most important, with something like half of all the electricity generated being converted back to mechanical energy in induction motors. Despite playing a key role in industrial society, it remains largely unnoticed because of its everyday role driving machinery, pumps, fans, compressors, conveyors, hoists, and a host of other routine but vital tasks. It will doubtless continue to dominate fixed-speed applications, but, thanks to the availability of reliable variable-frequency inverters, it is now also a leader in the controlled-speed arena.

Like the d.c. motor, the induction motor develops torque by the interaction of axial currents on the rotor and a radial magnetic field produced by the stator. But whereas in the d.c. motor the 'work' current has to be fed into the rotor by means of brushes and a commutator, the torque-producing currents in the rotor of the induction motor are induced by electromagnetic action, hence the name 'induction' motor. The stator winding therefore not only produces the magnetic field (the 'excitation'), but also supplies the energy which is converted to mechanical output. The absence of any sliding mechanical contacts and the consequent saving in terms of maintenance is a major advantage of the induction motor over the d.c. machine.

Other differences between the induction motor and the d.c. motor are firstly that the supply to the induction motor is a.c. (usually three-phase, but in smaller sizes single-phase); secondly that the magnetic field in the induction motor rotates relative to the stator, while in the d.c. motor it is stationary; and thirdly that both stator and rotor in the induction motor are non-salient (i.e. effectively smooth) whereas the d.c. motor stator has projecting poles or saliencies which define the position of the field windings.

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00005-2

Given these differences we might expect to find major contrasts between the performance of the two types of motor, and it is true that their inherent characteristics exhibit distinctive features. But there are also many aspects of behaviour which are similar, as we shall see. Perhaps most important from the user's point of view is that there is no dramatic difference in size or weight between an induction motor and a d.c. motor giving the same power at the same base speed, though the induction motor will usually be cheaper. The similarity in size is a reflection of the fact that both types employ similar amounts of copper and iron, while the difference in price stems from the simpler construction and production volume of the induction motor.

5.1.1 Outline of approach

Throughout this chapter we will be concerned with how the induction motor behaves in the steady state, i.e. the supply voltage and frequency are constant, the load is steady, and any transients have died away. We will aim to develop a sound qualitative understanding of the steady-state behaviour, based on the ideas we have discussed so far (magnetic flux, m.m.f., reluctance, electromagnetic force, motional e.m.f.). But despite many similarities with the d.c. motor, most readers will probably find that the induction motor is more difficult to understand. This is because we are now dealing with alternating rather than steady quantities (so, for example, inductive reactance becomes very significant), and also because (as mentioned earlier) a single winding acts simultaneously as the producer of the flux and the supplier of the converted energy.

In the next chapter, we will extend our qualitative understanding to look at how motor performance depends on design parameters: we will again be following an approach that has served us well since the early days of the induction motor, and was developed to reflect the fact that motors were operated at a fixed voltage and frequency. It turns out that under these 'utility supply' conditions, the transient performance is poor, fast control of torque is not possible, and so the induction motor was considered unable to compete with the d.c. motor in controlled speed drives.

All this changed rapidly beginning in the 1970s. The full set of governing equations (describing not only the steady state but also the much more complex dynamic behaviour) had become tractable with computer simulation, which in turn led the way to understanding how the stator currents would have to be manipulated to obtain fast control of torque. The hardware for implementing rapid current control became available with the development of PWM inverters, but it was not until digital signal processing finally became cheap and fast enough to deal with the complex control algorithms that so-called 'field-oriented' or 'vector' control emerged as a practicable commercial proposition. We will defer consideration of this spectacularly successful system until later, because experience has shown that a solid grounding based on the classical approach is invaluable before we get to grips with the more demanding ideas introduced in Chapter 7.

5.2 The rotating magnetic field

To understand how an induction motor operates, we must first unravel the mysteries of the rotating magnetic field. We will see later that the rotor is effectively dragged along by the rotating field, but that it can never run quite as fast as the field.

Our look at the mechanism of the rotating field will focus on the stator windings because they act as the source of the flux. In this part of the discussion we will ignore the presence of the rotor conductors. This makes it much easier to understand what governs the speed of rotation and the magnitude of the field, which are the two factors that most influence motor behaviour.

Having established how the rotating field is set up, and what its speed and strength depend on, we move on to examine the rotor, concentrating on how it behaves when exposed to the rotating field, and discovering how the induced rotor currents and torque vary with rotor speed. In this section we assume—again for the sake of simplicity—that the rotating flux set up by the stator is not influenced by the rotor.

Finally we turn attention to the interaction between the rotor and stator, verifying that our earlier assumptions are well justified. Having done this we are in a position to examine the 'external characteristics' of the motor, i.e. the variation of motor torque and stator current with speed. These are the most important characteristics from the point of view of the user.

Readers who are unfamiliar with routine a.c. circuit theory, including reactance, impedance, phasor diagrams (but not, at this stage, 'j' notation) and basic ideas about three-phase systems will have to do some preparatory work[1] before tackling the later sections of this chapter.

Before we investigate how the rotating magnetic field is produced, we should be clear what it actually is. Because both the rotor and stator iron surfaces are smooth (apart from the regular slotting), and are separated by a small air-gap, the flux produced by the stator windings crosses the air-gap radially. The behaviour of the motor is dictated by this radial flux, so we will concentrate first on establishing a mental picture of what is meant by the 'flux wave' in an induction motor.

The pattern of flux in an ideal four-pole motor supplied from a balanced three-phase source is shown in Fig. 5.1A. The top sketch corresponds to time $t = 0$; the middle one shows the flux pattern one quarter of a cycle of the supply later (i.e. 5 ms if the frequency is 50 Hz); and the lower one corresponds to a further quarter-cycle later. We note that the pattern of flux lines is repeated

1. 'Electrical and Electronic Technology, 12th Edition' by Edward Hughes (no relation) is a tried and tested favourite.

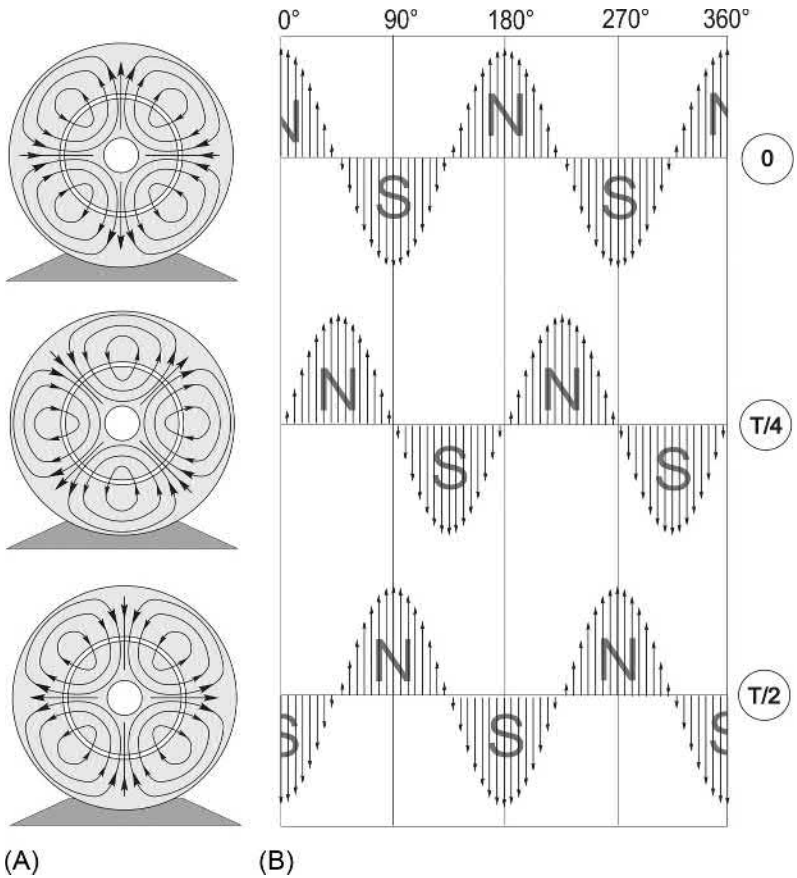

FIG. 5.1 (A) Flux pattern in a four-pole induction motor at three successive instants of time, each one-quarter of a cycle apart; (B) radial flux density distribution in the air-gap at the three instants shown in Fig. 5.1A.

in each case, except that the middle and lower ones are rotated by 45 and 90 degrees respectively with respect to the top sketch.

The term 'four-pole' reflects the fact that flux leaves the stator from two N poles, and returns at two S poles. Note however that there are no physical features of the stator iron that mark it out as being four-pole, rather than say two-pole or six-pole. As we will see, it is the layout and interconnection of the stator coils which sets the pole number.

If we plot the variation of the radial air-gap flux density with respect to distance round the stator, at each of the three instants of time, we get the patterns shown in Fig. 5.1B. The first feature to note is that the radial flux density varies sinusoidally in space. There are two N peaks and two S peaks, but the transition from N to S occurs in a smooth sinusoidal way, giving rise to the term 'flux wave'. The distance from the centre of one N pole to the centre of the adjacent S pole is called the pole-pitch, for obvious reasons.

Staying with Fig. 5.1B, we note that after one quarter of a cycle of the utility frequency, the flux wave retains its original shape, but has moved round the stator by half a pole-pitch, while after half a cycle it has moved round by a full pole-pitch. If we had plotted the patterns at intermediate times, we would have discovered that the wave maintained a constant shape, and progressed smoothly, advancing at a uniform rate of two pole-pitches per cycle of the mains. The term 'travelling flux wave' is thus an appropriate one to describe the air-gap field.

For the four-pole wave here, one complete revolution takes two cycles of the supply, so the speed is 25 revs/s (1500 rev/min) with a 50 Hz supply, or 30 rev/s (1800 rev/min) at 60 Hz. The general expression for the speed of the field (which is known as the synchronous speed) N_s, in rev/min is

$$N_s = \frac{120f}{p} \tag{5.1}$$

where p is the pole-number. The pole-number must be an even integer, since for every N pole there must be a S pole. Synchronous speeds for commonly-used pole-numbers are given in Table 5.1 below.

We can see from the table that if we want the field to rotate at intermediate speeds, we will have to be able to vary the supply frequency, and this is what happens in inverter-fed motors, which are dealt with in Chapters 7 and 8.

5.2.1 Production of a rotating magnetic field

Now that we have a picture of the field, we turn to how it is produced. If we inspect the stator winding of an induction motor we find that it consists of a uniform array of identical coils, located in slots. The coils are in fact connected to form three identical groups or phase-windings, distributed around the stator,

TABLE 5.1 Synchronous speeds in rev/min

Pole number	50 Hz utility supply	60 Hz utility supply
2	3000	3600
4	1500	1800
6	1000	1200
8	750	900
10	600	720
12	500	600

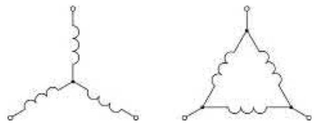

FIG. 5.2 Star (wye) and delta (mesh) connection of the three phase-windings of a three-phase induction motor.

and symmetrically displaced with respect to one another. The three phase-windings are connected either in star (wye) or delta (mesh), as shown in Fig. 5.2.

The three phase-windings are connected to a balanced three-phase a.c. supply, and so the currents (which produce the m.m.f. that sets up the flux) are of equal amplitude but differ in time-phase by one third of a cycle (120 degrees), forming a balanced three-phase set.

5.2.2 Field produced by each phase-winding

The aim of the winding designer is to arrange the layout of the coils so that each phase-winding, acting alone, produces an m.m.f. wave (and hence an air-gap flux wave) of the desired pole-number, and with a sinusoidal variation of amplitude with angle. Getting the desired pole-number is not difficult: we simply have to choose the right number and pitch of coils, as shown by the diagrams of an elementary four-pole winding in Fig. 5.3.

In Fig. 5.3A we see that by positioning two coils (each of which spans one pole-pitch) 180 degrees apart we obtain the correct number of poles (i.e. 4). However, the air gap field—shown by only two flux lines per pole for the sake of clarity—is uniform between each go and return coil side, not sinusoidal.

A clearer picture of the air-gap flux wave is presented in the developed view in Fig. 5.3B, where the equally-spaced flux lines emphasise the uniformity of the flux density between the go and return sides of the coils. Finally, the plot of the air-gap flux density underlines the fact that this very basic arrangement of coils produces a rectangular flux density wave, whereas what we are seeking is a sinusoidal wave.

We can improve matters by adding more coils in the adjacent slots, as shown in Fig. 5.4. All the coils have the same number of turns, and carry the same current. The addition of the extra slightly-displaced coils gives rise to the stepped waveform of m.m.f. and air-gap flux density shown in Fig. 5.4. It is still not sinusoidal, but is much better than the original rectangular shape.

It turns out that if we were to insist on having a perfect sinusoidal flux density waveform, we would have to distribute the coils of one phase in a smoothly-varying sinusoidal pattern over the whole periphery of the stator. This is not a practicable proposition, firstly because we would also have to vary the number of turns per coil from point to point, and secondly because we want the coils to be in slots, so it is impossible to avoid some measure of discretisation in the

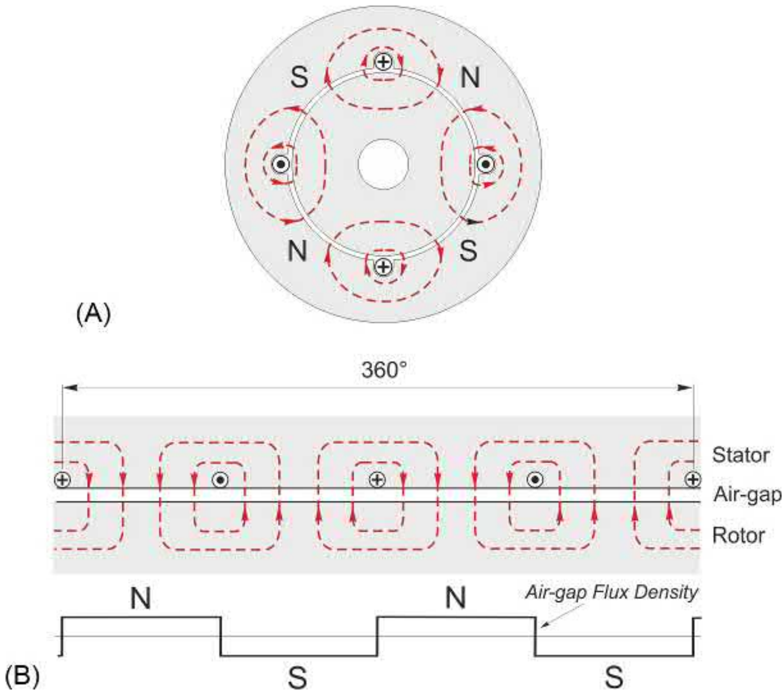

FIG. 5.3 Arrangement (A) and developed diagram (B) showing elementary four-pole, single-layer stator winding consisting of four conductors spaced by 90 degrees. The 'go' side of each coil (shown by the plus symbol) carries current into the paper at the instant shown, while the 'return' side (shown by the dot) carries current out of the paper.

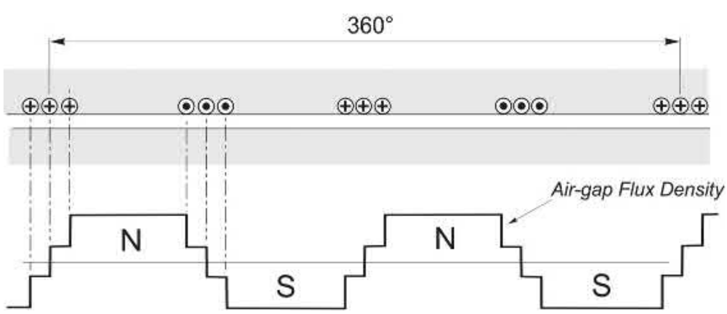

FIG. 5.4 Developed diagram showing flux density produced by one phase of a single-layer winding having three slots per pole per phase.

layout. For economy of manufacture we are also obliged to settle for all the coils being identical, and we must make sure that the three identical phase-windings fit together in such a way that all the slots are utilised.

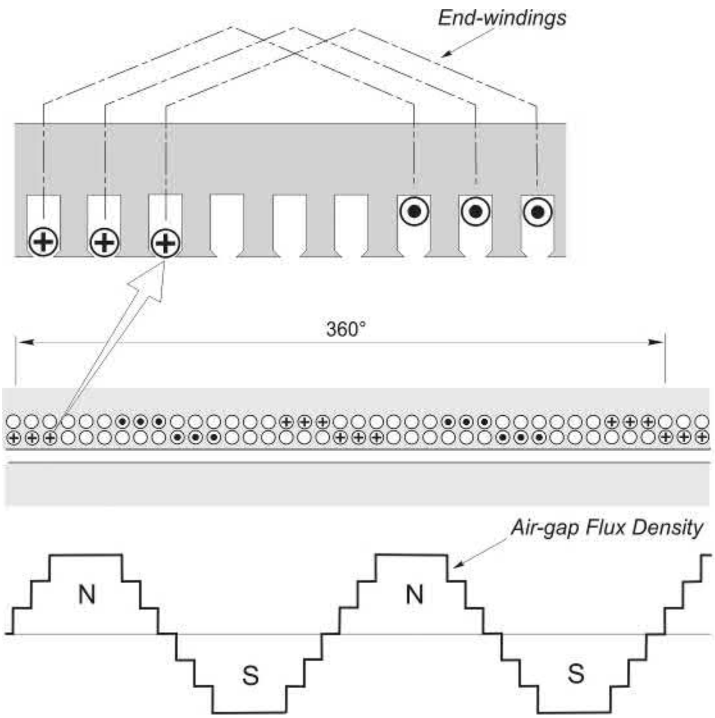

FIG. 5.5 Developed diagram showing layout of windings in a three-phase, four-pole, two-layer induction motor winding, together with the flux density wave produced by one phase acting alone. The upper detail shows how the coil-sides form upper and lower layers in the slots.

Despite these constraints we can get remarkably close to the ideal sinusoidal pattern, especially when we use a 'two-layer' winding (in which case the stator slots may contain turns from more than one phase winding). A typical arrangement of one phase is shown in Fig. 5.5. The upper expanded sketch shows how each coil sits with its go side in the top of a slot while the return side occupies the bottom of a slot rather less than one pole-pitch away. Coils which span less than a full pole-pitch are known as short-pitch or short-chorded: in this particular case the coil pitch is six slots, the pole-pitch is nine slots, so the coils are short-pitched by three slots.

This type of winding is widely used in induction motors, the coils in each phase being grouped together to form 'phase-bands' or 'phase-belts'. Since we are concentrating on the field produced by only one of the phase-windings (or 'phases'), only one third of the coils in Fig. 5.5 are shown carrying current. The remaining two-thirds of the coils form the other two phase-windings, as discussed in the next section.

Returning to the flux density plot in Fig. 5.5 we see that the effect of short-pitching is to increase the number of steps in the waveform, and that as a result the field produced by one phase is a fair approximation to a sinusoid.

FIG. 5.6 The photo (left) shows the slotted stator core and slot liners, with the first of the three phase windings inserted. It is an eight-pole fully pitched two-layer winding, and we can see that one coil-side is in the top of a slot and the return side is in the bottom. The photo (right) shows the same motor with all three phase windings inserted. These pictures are taken at the stage of manufacture prior to the end windings being shaped into their final form which is usually done by pressing, after which the stator is usually impregnated with a resin to improve electrical insulation and protect against vibration of the windings. *(Courtesy of Nidec Control Techniques Dynamics Ltd.)*

The current in each phase pulsates at the supply frequency, so the field produced by, say, phase A, pulsates in sympathy with the current in phase A, the axis of each 'pole' remaining fixed in space, but its polarity changing from N to S and back once per cycle. There is no hint of any rotation in the field of one phase, but when the fields produced by each of the three phases are combined, matters change dramatically, as we will see shortly.

We should mention that although it is desirable to minimise the harmonic content of the phase m.m.f. by striving for a sinusoidal m.m.f., it is not always feasible, particularly with high pole numbers, and in smaller motor sizes. For example, an eight-pole two-layer winding in 24 slots is shown in Fig. 5.6. In this case there is only one slot per pole per phase, so the phase m.m.f. is rectangular, with a high harmonic content. Nevertheless, because the overall performance of an induction motor is dominated by the fundamental component, the simple winding is entirely satisfactory.

5.2.3 Resultant three-phase field

The layout of coils for the complete four-pole winding under discussion is shown in shown in Fig. 5.7A. The go sides of each coil are represented by the capital letters (A, B, C) and the return sides are identified by bars over the letters ($\overline{A}, \overline{B}, \overline{C}$). (For the sake of comparison, a six-pole winding layout that

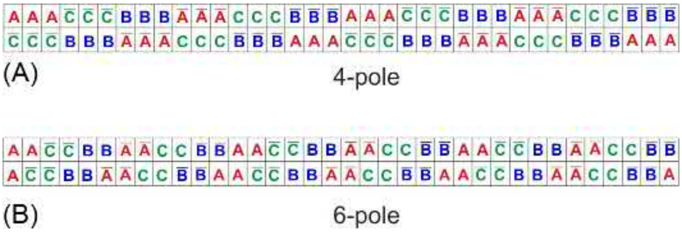

FIG. 5.7 Developed diagram showing arrangement of three-phase, two-layer windings in a 36-slot stator. A four-pole winding with three slots/pole/phase is shown in (A), and a six-pole winding with two slots/pole/phase is shown in (B).

uses the same stator slotting is shown in Fig. 5.7B: here the pole-pitch is six slots and the coils are short-pitched by one slot.)

Returning to the four-pole winding, we can see that the windings of phases B and C are identical with that of phase A apart from the fact that they are displaced in space by plus and minus two thirds of a pole-pitch respectively. Phases B and C therefore also produce pulsating fields, along their own fixed axes in space. But the currents in phases B and C also differ in time-phase from the current in phase A, lagging by one third and two thirds of a cycle respectively. To find the resultant field we must therefore superimpose the fields of the three phases, taking account not only of the spatial differences between windings, but also the time differences between the currents. This is a tedious process, so the intermediate steps have been omitted and instead we move straight to the plot of the resultant field for the complete four-pole machine, for three discrete times during one complete cycle, as shown in Fig. 5.8.

We see that the three pulsating fields combine beautifully and lead to a resultant four-pole field which rotates at a uniform rate, advancing by two pole-pitches for every cycle of the supply. The resultant field is not exactly sinusoidal in shape (though it is actually more sinusoidal than the field produced by the individual phase-windings), and its shape varies a little from instant to instant; but these are minor worries. The resultant field is amazingly close to the ideal travelling wave and yet the winding layout is simple and easy to manufacture. This is an elegant engineering achievement, however one looks at it.

Before leaving our look at windings, it should be pointed out that the windings we have illustrated are called 'integral-slot' windings, where the number of slots per pole per phase is an integer: for example, the four-pole three-phase winding in 36 slots shown in Fig. 5.7A has 3 s/p/ph, while the six-pole version in Fig. 5.7B has 2 s/p/ph.

However, suppose we wanted to use a 54-slot stator for a four-pole, three-phase winding, for which the number of slots per pole per phase would be 4.5. We might guess that such a winding would not be possible, but in fact such so-called 'fractional-slot' windings are widely used where a standard lamination is to be used for several different pole-numbers. In the four-pole example above,

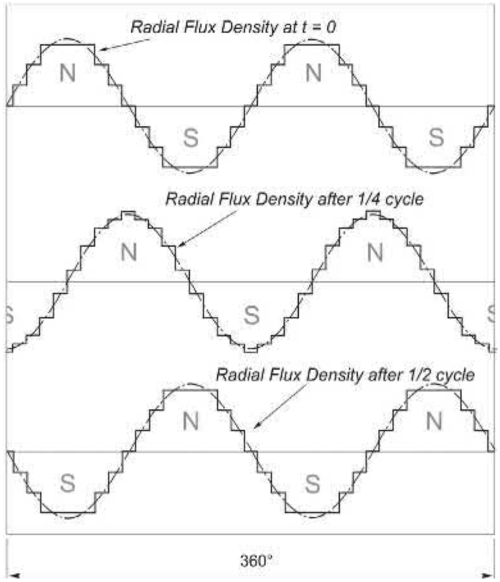

FIG. 5.8 Resultant air-gap flux density wave produced by a complete three-phase, four-pole winding at three successive instants in time.

the first and third phase-bands would comprise four adjacent slots, while the second and fourth would comprise five adjacent slots, an average of 4.5. These windings produce perfectly satisfactory sinusoidally-distributed travelling fields when supplied with sinusoidal currents from a utility supply.

We will come across an extreme version of fractional-slot windings later when we look at developments aimed at minimising production costs for mass-market permanent-magnet motors, particularly for use in electric vehicles. These motors are inverter-fed and their stators have relatively few slots, but with relatively high pole-numbers: consequently, the slots per pole per phase is normally fractional, and, surprisingly, is often less than unity.

5.2.4 Direction of rotation

The direction of rotation depends on the order in which the currents reach their maxima, i.e. on the phase-sequence of the supply. Reversal of direction is therefore simply a matter of interchanging any two of the lines connecting the windings to the supply.

5.2.5 Main (air-gap) flux and leakage flux

Broadly speaking the motor designer shapes the stator and rotor teeth to encourage as much as possible of the flux produced by the stator windings to pass right down the rotor teeth, so that before completing its path back to the stator it is fully linked with the rotor conductors (see later) which are located in the rotor slots. We will see later that this tight magnetic coupling between stator and rotor windings is necessary for good running performance, and the field which provides the coupling is of course the main or air-gap field, which we are in the midst of discussing.

In practice the vast majority of the flux produced by the stator is indeed main or 'mutual' flux. But there is some flux which bypasses the rotor conductors, linking only with the stator winding, and known as stator leakage flux. Similarly not all the flux produced by the rotor currents links the stator, but some (the rotor leakage flux) links only the rotor conductors.

The use of the perjorative-sounding term 'leakage' suggests that these leakage fluxes are unwelcome imperfections, which we should go out of our way to minimise. However whilst the majority of aspects of performance are certainly enhanced if the leakage is as small as possible, others (notably the large and unwelcome current drawn when the motor is started from rest directly on the utility supply) are made much worse if the coupling is too good. So we have the somewhat paradoxical situation in which the designer finds it comparatively easy to lay out the windings to produce a good main flux, but is then obliged to juggle the detailed design of the slots in order to obtain just the right amount of leakage flux to give acceptable all-round performance. (In contrast, as we will see later, an inverter-fed induction motor can avoid such issues as excessive starting current and, ideally, could be designed with much lower leakage than its utility-fed counterpart. It has to be said however that the majority of induction motors are still designed for general-purpose use, and in this respect they lose out, often unfairly, in comparison with other forms of motor that are specifically designed for operation with a drive.)

The weight which attaches to the matter of leakage flux is reflected in the prominent part played by the associated leakage reactance in equivalent circuit models of the induction motor, which we have chosen not to cover. However, such niceties are of limited importance to the user, so in this and the next chapters we will limit references to leakage reactance to well-defined contexts, and in general, where the term 'flux' is used, it will refer to the main air-gap field.

5.2.6 Magnitude of rotating flux wave

We have already seen that the speed of the flux wave is set by the pole number of the winding and the frequency of the supply. But what is it that determines the amplitude of the field?

To answer this question we can continue to neglect the fact that under normal conditions there will be induced currents in the rotor. We might even find it easier to imagine that the rotor conductors have been removed altogether: this may seem a drastic assumption, but will prove justified later. The stator windings are assumed to be connected to a balanced three-phase a.c. supply so that a balanced set of currents flows in the windings. We denote the phase voltage by V, and the current in each phase by I_m, where the subscript m denotes 'magnetising' or flux-producing current.

From the discussion in Chapter 1 we know that the magnitude of the flux wave (B_m) is proportional to the winding m.m.f., and is thus proportional to I_m. But what we really want to know is how the flux density depends on the supply voltage and frequency, since these are the only two parameters over which we have control.

To guide us to the answer, we must first ask what effect the travelling flux wave will have on the stator winding. Every stator conductor will of course be cut by the rotating flux wave, and will therefore have an e.m.f. induced in it. Since the flux wave varies sinusoidally in space, and cuts each conductor at a constant velocity, a sinusoidal e.m.f. is induced in each conductor. The magnitude of the e.m.f. is proportional to the magnitude of the flux wave (B_m), and to the speed of the wave (i.e. to the supply frequency f). The frequency of the induced e.m.f. depends on the time taken for one N pole and one S pole to cut the conductor. We have already seen that the higher the pole-number, the slower the field rotates, but we found that the field always advances by two pole-pitches for every cycle of the supply. The frequency of the e.m.f. induced in the stator conductors is therefore the same as the supply frequency, regardless of the pole-number. (This conclusion is what we would have reached intuitively, since we would expect any linear system to react at the same frequency at which we excited it.)

The e.m.f. in each complete phase winding (E) is the sum of the e.m.f's in the phase coils, and thus will also be at supply frequency. (The alert reader will realise that whilst the e.m.f. in each coil has the same magnitude, it will differ in time phase, depending on the geometrical position of the coil. Most of the coils in each phase-band are close together, however, so their e.m.f's—though slightly out of phase—will more or less add up directly.)

If we were to compare the e.m.f's in the three complete phase windings, we would find that they were of equal amplitude, but out of phase by one third of a cycle (120°), thereby forming a balanced three-phase set. This result could have been anticipated from the overall symmetry. This is very helpful, as it means that we need only consider one of the phases in the rest of the discussion.

So we find that when an alternating voltage V is applied, an alternating e.m.f., E, is induced. We can represent this state of affairs by the primitive a.c. equivalent circuit for one phase shown in Fig. 5.9.

The resistance shown in Fig. 5.9 is the resistance of one complete phase-winding. Note that the e.m.f. E is shown as opposing the applied voltage V. This must be so, otherwise we would have a runaway situation in which the voltage V produced the magnetising current I_m which in turn set up an e.m.f. E, which added to V, which further increased I_m and so on ad infinitum.

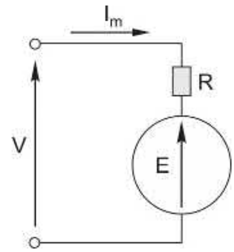

FIG. 5.9 Simple equivalent circuit for the induction motor under no-load conditions.

Applying Kirchoff's voltage law to the a.c. circuit in Fig. 5.8 yields

$$V = I_m R + E \tag{5.2}$$

We find in practice that the term $I_m R$ (which represents the volt drop due to winding resistance) is usually very much less than the applied voltage V. In other words most of the applied voltage is accounted for by the opposing e.m.f., E. Hence we can make the approximation

$$V \approx E \tag{5.3}$$

But we have already seen that the e.m.f. is proportional to B_m and to f, i.e.

$$E \propto B_m f \tag{5.4}$$

So by combining Eqs (5.3) and (5.4) we obtain

$$B_m = k\frac{V}{f} \tag{5.5}$$

where the constant k depends on the number of turns per coil, the number of coils per phase and the distribution of the coils.

Eq. (5.5) is of fundamental importance in induction motor operation. It shows that if the supply frequency is constant, the flux in the air-gap is directly proportional to the applied voltage, or in other words the voltage sets the flux. We can also see that if we raise or lower the frequency (in order to increase or reduce the speed of rotation of the field), we will have to raise or lower the voltage in proportion if, as is usually the case, we want the magnitude of the flux to remain constant. (We will see in Chapters 7 and 8 that the early inverter drives used this so-called 'V/f control' to keep the flux *constant* at all speeds.)

It may seem a paradox that having originally homed-in on the magnetising current I_m as being the source of the m.m.f. which in turn produces the flux, we find that the actual value of the flux is governed only by the applied voltage and frequency, and I_m does not appear at all in Eq. (5.5). We can see why this is by looking again at Fig. 5.9 and asking what would happen if, for some reason, the e.m.f. (E) were to reduce. We would find that I_m would increase, which in turn would lead to a higher m.m.f., more flux, and hence to an increase in E. There is clearly a negative feedback effect taking place, which continually tries to keep

E equal to V. It is rather like the d.c. motor (Chapter 3) where the speed of the unloaded motor always adjusted itself so that the back e.m.f. almost equalled the applied voltage. Here, the magnetising current always adjusts itself so that the induced e.m.f. is almost equal to the applied voltage.

Needless to say this does not mean that the magnetising current is arbitrary, but to calculate it we would have to know the number of turns in the winding, the length of the air-gap (from which we could calculate the gap reluctance) and the reluctance of the iron paths. From a user point of view there is no need to delve further in this direction. We should however recognise that the reluctance will be dominated by the air-gap, and that the magnitude of the magnetising current will therefore depend mainly on the size of the gap: the larger the gap, the bigger the magnetising current. Since the magnetising current contributes to stator copper loss, but not to useful output power, we would like it to be as small as possible, so we find that induction motors usually have the smallest air-gap which is consistent with providing the necessary mechanical clearances. Despite the small air-gap the magnetising current can be appreciable: in a four-pole motor, it may be typically 50% of the full-load current, and even higher in six-pole and eight-pole designs.

5.2.7 Excitation power and VA

The setting up of the travelling wave by the magnetising current amounts to the provision of 'excitation' for the motor. Some energy is stored in the magnetic field, but since the amplitude remains constant once the field has been established, no net power input is needed to sustain the field. We therefore find that under the conditions discussed so far, i.e. in the absence of any rotor currents, the power input to the motor is very small. (We should perhaps note that the rotor currents in a real motor are very small when it is running light, so the hypothetical situation we are looking at is not so far removed from reality as we may have supposed.)

Ideally the only source of power losses would be the copper losses in the stator windings, but to this must be added the 'iron losses' which arise from eddy currents and hysteresis in the laminated steel cores of rotor and stator. However we have seen that the magnetising current can be quite large, its value being largely determined by the air-gap, so we can expect an unloaded induction motor to draw appreciable current from the supply, but very little real power. The VA will therefore be substantial, but the power-factor will be very low, the magnetising current lagging the supply voltage by almost 90°, as shown in the time phasor diagram (Fig. 5.10).

Viewed from the supply the stator looks more or less like a pure inductance, a fact which we would expect intuitively given that—having ignored the rotor circuit—we are left with only an arrangement of flux-producing coils surrounded by a good magnetic circuit.

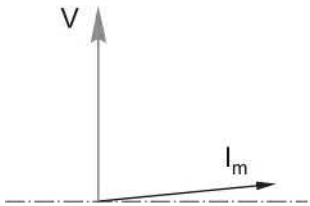

FIG. 5.10 Phasor diagram for the induction motor under no-load conditions, showing magnetising current I_m.

5.2.8 Summary

When the stator is connected to a three-phase supply, a sinusoidally distributed, radially-directed rotating magnetic flux density wave is set up in the air-gap. The speed of rotation of the field is directly proportional to the frequency of the supply, and inversely proportional to the pole-number of the winding. The magnitude of the flux wave is proportional to the applied voltage, and inversely proportional to the frequency.

When the rotor circuits are ignored (i.e. under no-load conditions), the real power drawn is small, but the magnetising current itself can be quite large, giving rise to a significant reactive power demand from the utility supply.

5.3 Torque production

In this section we begin with a brief description of rotor types, and introduce the notion of 'slip', before moving on to explore how the torque is produced, and investigate the variation of torque with speed. We will find that the behaviour of the rotor varies widely according to the slip, and we therefore look separately at low and high values of slip. Throughout this section we will assume that the rotating magnetic field is unaffected by anything which happens on the rotor side of the air-gap. Later, we will see that this assumption is pretty well justified.

5.3.1 Rotor construction

Two types of rotor are used in induction motors. In both the rotor 'iron' consists of a stack of silicon steel laminations with evenly-spaced slots punched around the circumference. As with the stator laminations, the surface is coated with an oxide layer which acts as an insulator, preventing unwanted axial eddy-currents from flowing in the iron.

The cage rotor is by far the most common: each rotor slot contains a solid conductor bar and all the conductors are physically and electrically joined together at each end of the rotor by conducting 'end-rings', as shown in Fig. 5.11. The term squirrel cage was widely used at one time and the origin should be clear from Fig. 5.11. The rotor bars and end-rings are reminiscent

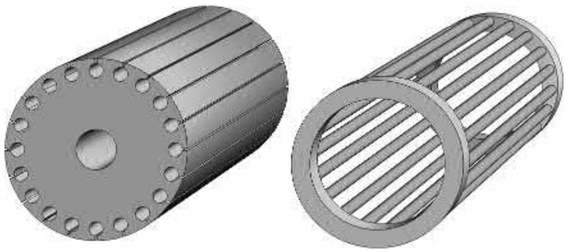

FIG. 5.11 Cage rotor construction. The stack of pre-punched laminations is shown on the left, with the copper or aluminium rotor bars and end-rings on the right.

FIG. 5.12 Cage rotor for induction motor. The rotor conductor bars and end rings are cast in aluminium, and the blades attached to the end rings serve as a fan for circulating internal air. An external fan will be mounted on the nondrive end to cool the finned stator casing (as shown in Fig. 8.4).

of the rotating cages used in bygone days to exercise small rodents (or rather to amuse their human captors).

In the larger sizes the conductors will be of copper, in which case the end-rings are brazed-on. In small and medium sizes, the rotor conductors and end rings may be of copper (particularly in high efficiency motors) or most commonly die-cast in aluminium, as shown in Fig. 5.12.

The absence of any means for making direct electrical connection to the rotor underlines the fact that in the induction motor the rotor currents are induced by the air-gap field. It is equally clear that because the rotor cage comprises permanently short-circuited conductor bars, no external control can be exercised over the resistance of the rotor circuit once the rotor has been made. This is a significant drawback which can be avoided in the second type of rotor, which is known as the 'wound-rotor' or 'slipring' type.

In the wound rotor, the slots accommodate a set of three phase-windings very much like those on the stator. The windings are connected in star, with the three ends brought-out to three slip-rings (Fig. 5.13). The rotor circuit is thus open, and connection can be made via brushes bearing on the sliprings.

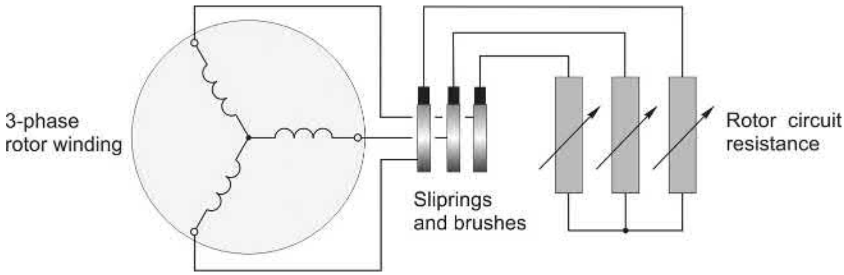

FIG. 5.13 Schematic diagram of wound rotor for induction motor, showing sliprings and brushes to provide connection to the external (stationary) three-phase resistance.

In particular, the resistance of each phase of the rotor circuit can be increased by adding external resistances, as indicated in Fig. 5.13. Adding resistance can be beneficial in some circumstances, as we will see.

Cage-rotors are usually cheaper to manufacture, and are very robust and reliable. Until the advent of variable-frequency inverter supplies, however, the superior control which was possible from the slip-ring type meant that the extra expense of the wound rotor and its associated control gear were frequently justified, especially for high-power machines. Nowadays comparatively few are made, and then only in large sizes. But many old motors remain in service, so they are included in Chapter 6.

5.3.2 Slip

A little thought will show that the behaviour of the rotor depends very much on its relative velocity with respect to the rotating field. If the rotor is stationary, for example, the rotating field will cut the rotor conductors at synchronous speed, thereby inducing a high e.m.f. in them. On the other hand, if the rotor was running at the synchronous speed, its relative velocity with respect to the field would be zero, and no e.m.f.'s would be induced in the rotor conductors.

The relative velocity between the rotor and the field is known as the slip speed. If the speed of the rotor is N, the slip speed is N_s–N, where N_s is the synchronous speed of the field, usually expressed in rev/min. The slip (as distinct from slip speed) is the normalised quantity defined by

$$s = \frac{N_S - N}{N_S} \tag{5.6}$$

and is usually expressed either as a ratio as in Eq. (5.6), or as a percentage. A slip of 0 therefore indicates that the rotor speed is equal to the synchronous speed, while a slip of 1 corresponds to zero speed. (When tests are performed on induction motors with their rotor deliberately held stationary so that the slip is 1, the test is said to be under 'locked-rotor' conditions. The same expression is often used loosely to mean zero speed, even when the rotor is free to move, e.g. when it is started from rest.)

5.3.3 Rotor induced e.m.f. and current

The rate at which the rotor conductors are cut by the flux, and hence their induced e.m.f., is directly proportional to the slip, with no induced e.m.f. at synchronous speed ($s=0$) and maximum induced e.m.f. when the rotor is stationary ($s=1$).

The frequency of the rotor e.m.f. (the slip frequency) is also directly proportional to slip, since the rotor effectively slides with respect to the flux-wave, and the higher the relative speed, the more times in a second each rotor conductor is cut by a N and a S pole. At synchronous speed (slip = 0) the slip frequency is zero, while at standstill (slip = 1), the slip frequency is equal to the supply frequency. These relationships are shown in Fig. 5.14.

Although the e.m.f. induced in every rotor bar will have the same magnitude and frequency, they will not be in phase. At any particular instant, bars under the peak of the N poles of the field will have maximum positive voltage in them, those under the peak of the S poles will have maximum negative voltage, (i.e. 180° phase shift) and those in between will have varying degrees of phase shift. The pattern of instantaneous voltages in the rotor is thus a replica of the flux density wave, and the rotor induced ‘voltage wave’ therefore moves relative to the rotor at slip speed, as shown in Fig. 5.15.

All the rotor bars are short-circuited by the end-rings, so the induced voltages will drive currents along the rotor bars, the currents forming closed paths through the end-rings, as shown in the developed diagram (Fig. 5.16).

In Fig. 5.16 the variation of instantaneous e.m.f. in the rotor bars is shown in the upper sketch, while the corresponding instantaneous currents flowing in the rotor bars and end-rings are shown in the lower sketch. The lines representing the currents in the rotor bars have been drawn so that their width is proportional to the instantaneous currents in the bars.

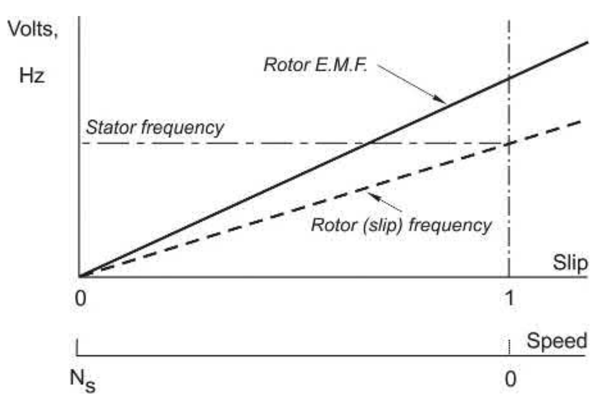

FIG. 5.14 Variation of rotor induced e.m.f. and frequency with speed and slip.

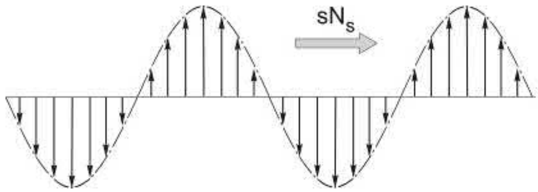

FIG. 5.15 Pattern of induced e.m.f's in rotor conductors. The rotor 'voltage wave' moves at a speed of sN_s with respect to the rotor surface.

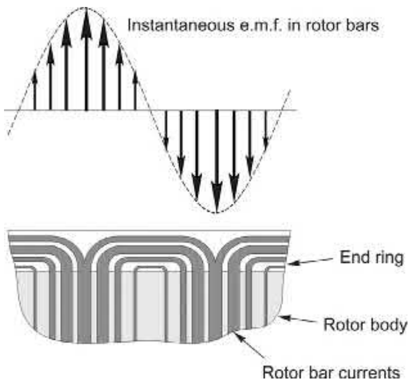

FIG. 5.16 Instantaneous sinusoidal pattern of rotor currents in rotor bars and end-rings. Only one pole-pitch is shown, but the pattern is repeated.

5.3.4 Torque

The axial currents in the rotor bars will interact with the radial flux wave to produce the driving torque of the motor, which will act in the same direction as the rotating field, the rotor being dragged along by the field. We note that slip is essential to this mechanism, so that it is never possible for the rotor to catch-up with the field, as there would then be no rotor e.m.f., no current, and no torque. The fact that motor action is only possible if the speed is less than the synchronous speed explains why the induction machine is described as 'asynchronous'. Finally, we can see that the cage rotor will automatically adapt to whatever pole-number is impressed by the stator winding, so that the same rotor can be used for a range of different stator pole numbers.

5.3.5 Rotor currents and torque—small slip

When the slip is small (say between 0 and 10%), the frequency of induced e.m.f. is also very low (between 0 and 5 Hz if the supply frequency is 50 Hz). At these low frequencies the impedance of the rotor circuits is predominantly resistive, the inductive reactance being small because the rotor frequency is low.

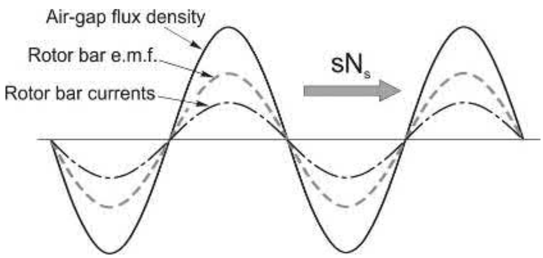

FIG. 5.17 Pattern of air-gap flux density, induced e.m.f. and current in cage rotor bars at low values of slip.

The current in each rotor conductor is therefore in time-phase with the e.m.f. in that conductor, and the rotor current-wave is therefore in space-phase with the rotor e.m.f. wave, which in turn is in space-phase with the flux wave. This situation was assumed in the previous discussion, and is represented by the space waveforms shown in Fig. 5.17.

To calculate the torque we first need to evaluate the 'BI_rl_r' product (see Eq. 1.2) in order to obtain the tangential force on each rotor conductor. The torque is then given by the total force multiplied by the rotor radius. We can see from Fig. 5.17 that where the flux density has a positive peak, so does the rotor current, so that particular bar will contribute a high tangential force to the total torque. Similarly, where the flux has its maximum negative peak, the induced current is maximum and negative, so the tangential force is again positive. We don't need to work out the torque in detail, but it should be clear that the resultant will be given by an equation of the form

$$T = KBI_r \tag{5.7}$$

where B and I_r denote the amplitudes of the flux density wave and the rotor current wave respectively. Provided that there are a large number of rotor bars (which is a safe bet in practice), the waves shown in Fig. 5.17 will remain the same at all instants of time, so the torque remains constant as the rotor rotates.

If the supply voltage and frequency are constant, the flux will be constant (See Eq. 5.5). The rotor e.m.f. (and hence I_r) is then proportional to slip, so we can see from Eq. (5.7) that the torque is directly proportional to slip. We must remember that this discussion relates to low values of slip only, but since this is the normal running condition, it is extremely important.

The torque-speed (and torque-slip) relationship for small slips is thus approximately a straight-line, as shown by the section of line AB in Fig. 5.18.

If the motor is unloaded, it will need very little torque to keep running—only enough to overcome friction in fact—so an unloaded motor will run with a very small slip at just below the synchronous speed, as shown at A in Fig. 5.18.

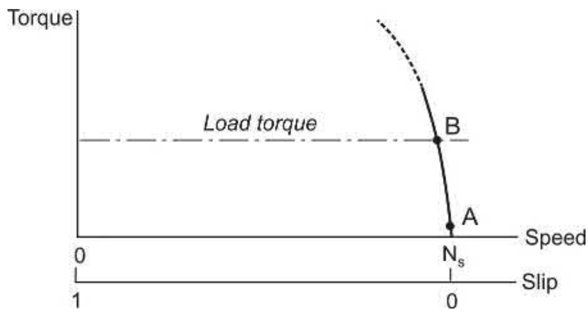

FIG. 5.18 Torque-speed relationship for low values of slip.

When the load is increased, the rotor slows down, and the slip increases, thereby inducing more rotor e.m.f. and current, and thus more torque. The speed will settle when the slip has increased to the point where the developed torque equals the load torque—e.g. point B in Fig. 5.18.

Induction motors are usually designed so that their full-load torque is developed for small values of slip. Small ones typically have a full-load slip of 8%, large ones around 1%. At the full-load slip, the rotor conductors will be carrying their safe maximum continuous current, and if the slip is any higher, the rotor will begin to overheat. This overload region is shown by the dotted line in Fig. 5.18.

The torque-slip (or torque-speed) characteristic shown in Fig. 5.18 is a good one for most applications, because the speed only falls a little when the load is raised from zero to its full value. We note that, in this normal operating region, the torque-speed curve is very similar to that of a d.c. motor (see Fig. 3.9).

5.3.6 Rotor currents and torque—large slip

As the slip increases, the rotor e.m.f. and rotor frequency both increase in direct proportion to the slip. At the same time the rotor inductive reactance, which was negligible at low slip (low rotor frequency) begins to be appreciable in comparison with the rotor resistance. Hence although the induced current continues to increase with slip, it does so more slowly than at low values of slip, as shown in Fig. 5.19.

At high values of slip, the rotor current also lags behind the rotor e.m.f. because of the inductive reactance. The alternating current in each bar reaches its peak well after the induced voltage, and this in turn means that the rotor current wave has a space-lag with respect to the rotor e.m.f. wave (which is in space-phase with the flux wave). This space-lag is shown by the angle ϕ_r in Fig. 5.20.

The space-lag means that the peak radial flux density and peak rotor currents no longer coincide, which is bad news from the point of view of torque

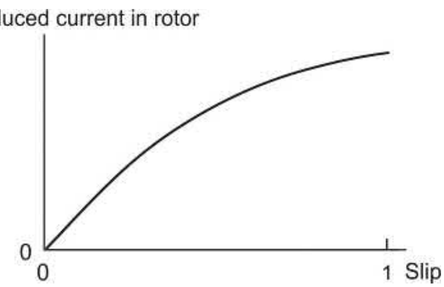

FIG. 5.19 Magnitude of current induced in rotor over the full (motoring) range of slip.

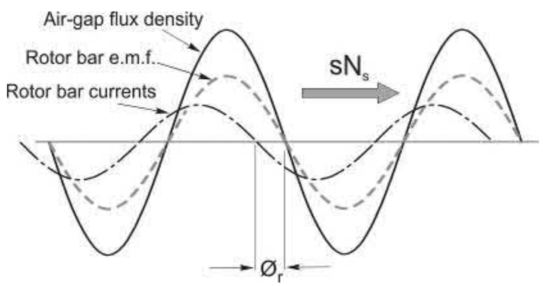

FIG. 5.20 Pattern of air-gap flux density, induced e.m.f. and current in cage rotor bars at high values of slip (These waveforms should be compared with the corresponding ones when the slip is small, see Fig. 5.17).

production, because although we have high values of both flux density and current, they do not occur simultaneously at any point around the periphery. What is worse is that at some points we even have flux density and currents of opposite sign, so over those regions of the rotor surface the torque contributed will actually be negative. The overall torque will still be positive, but is much less than it would be if the flux and current waves were in phase. We can allow for the unwelcome space-lag by modifying Eq. (5.7), to obtain a more general expression for torque as

$$T = KBI_r \cos\phi_r \tag{5.8}$$

Eq. (5.7) is merely a special case of Eq. (5.8), which only applies under low-slip conditions where $\cos\phi_r \approx 1$.

For most cage rotors, it turns out that as the slip increases the term $cos\ \phi_r$ reduces more quickly than the current (Ir) increases, so that at some slip between 0 and 1 the developed torque reaches a maximum value. This is illustrated in the typical torque-speed characteristic shown in Fig. 5.21. The peak torque actually occurs at a slip at which the rotor inductive reactance is equal to the rotor resistance, so the motor designer can position the peak torque at any slip by varying the reactance to resistance ratio.

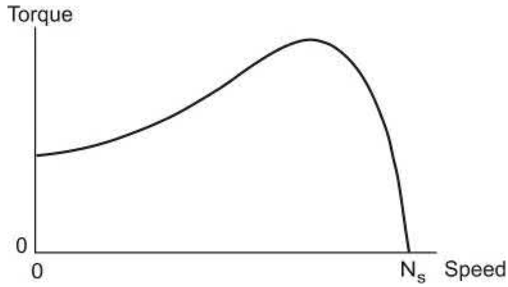

FIG. 5.21 Typical complete torque-speed characteristic for motoring region of cage induction motor.

5.3.7 Generating—Negative slip

When we explored the steady-state characteristics of the d.c. machine (see Section 3.4) we saw that at speeds less than that at which it runs when unloaded the machine acts as a motor, converting electrical energy into mechanical energy. But if the speed is above the no-load speed (e.g. when driven by a prime-mover), the machine generates and converts mechanical energy into electrical form.

The inherently bi-directional energy converting property of the d.c. machine seems to be widely recognised. But in the experience of the authors the fact that the induction machine behaves in the same way is far less well accepted, and indeed it is not uncommon to find users expressing profound scepticism at the thought that their 'motor' could possibly generate.

In fact, the induction machine behaves in essentially the same way as the d.c. machine, and if the rotor is driven by an external torque such that its speed is above the synchronous speed (i.e. the slip becomes negative), the electromagnetic torque reverses direction, and the power becomes negative, with energy fed back to the utility supply. It is important to note that, just as with the d.c. machine, this transition from motoring to generating takes place naturally, without intervention on our part.

When the speed is greater than synchronous, we can see from Eq. (5.6) that the slip is negative, and in this negative slip region the torque is also negative, the torque-speed curve broadly mirroring that in the motoring region, as shown in Fig. 5.22. We will discuss this further in Chapter 6, but it is worth noting that for both motoring and generating continuous operation will be confined to low values of slip, as indicated by the heavy line in Fig. 5.22.

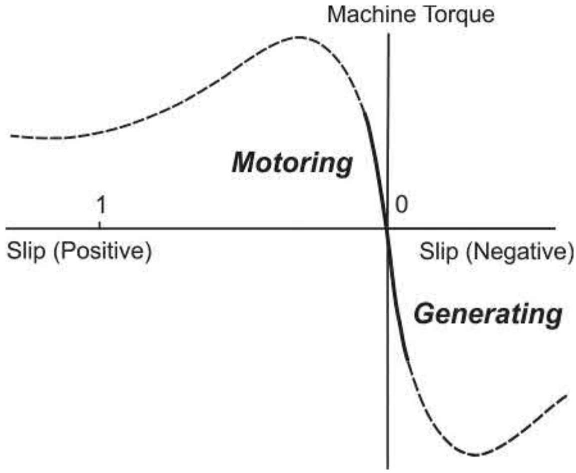

FIG. 5.22 Typical torque-speed characteristic showing stable motoring and generating regions.

5.4 Influence of rotor current on flux

Up to now all our discussion has been based on the assumption that the rotating magnetic field remains constant, regardless of what happens on the rotor. We have seen how torque is developed, and that mechanical output power is produced. We have focused attention on the rotor, but the output power must be provided from the stator winding, so we must turn attention to the behaviour of the whole motor, rather than just the rotor. Several questions spring to mind.

Firstly, what happens to the rotating magnetic field when the motor is working? Won't the m.m.f. of the rotor currents cause it to change? Secondly, how does the stator know when to start supplying real power across the air-gap to allow the rotor to do useful mechanical work? And finally, how will the currents drawn by the stator vary as the slip is changed?

These are demanding questions, for which full treatment is beyond our scope. But we can deal with the essence of the matter without too much difficulty.

5.4.1 Reduction of flux by rotor current

We should begin by recalling that we have already noted that when the rotor currents are negligible ($s=0$), the e.m.f. which the rotating field induces in the stator winding is very nearly equal to the applied voltage. Under these conditions a reactive current (which we termed the magnetising current) flows into the windings, to set up the rotating flux. Any slight tendency for the flux to fall is immediately detected by a corresponding slight reduction in e.m.f. which is reflected in a disproportionately large increase in magnetising current, which thus opposes the tendency for the flux to fall.

Exactly the same feedback mechanism comes into play when the slip increases from zero, and rotor currents are induced. The rotor currents are at slip frequency, and they give rise to a rotor m.m.f. wave, which therefore rotates at slip speed (sN_s) relative to the rotor. But the rotor is rotating at a speed of $(1-s)N_s$, so that when viewed from the stator, the rotor m.m.f. wave always rotates at synchronous speed, regardless of the speed of the rotor.

The rotor m.m.f. wave would, if unchecked, cause its own 'rotor flux wave', rotating at synchronous speed in the air-gap, in much the same way that the stator magnetising current originally set up the flux wave. The rotor flux wave would oppose the original flux wave, causing the resultant flux wave to reduce.

However, as soon as the resultant flux begins to fall, the stator e.m.f. reduces, thereby admitting more current to the stator winding, and increasing its m.m.f. A very small drop in the e.m.f. induced in the stator is sufficient to cause a large increase in the current drawn from the supply because the e.m.f. E (see Fig. 5.9) and the supply voltage V are both very large in comparison with the stator resistance volt-drop, IR. The 'extra' stator m.m.f. produced by the large increase in stator current effectively 'cancels' the m.m.f. produced by the rotor currents, leaving the resultant m.m.f. (and hence the rotating flux wave) virtually unchanged.

There must be a small drop in the resultant m.m.f. (and flux) of course, to alert the stator to the presence of rotor currents. But because of the delicate balance between the applied voltage and the induced e.m.f. in the stator the change in flux with load is very small, at least over the normal operating speed-range, where the slip is small. In large motors, the drop in flux over the normal operating region is typically less than 1%, rising to perhaps 10% in a small motor.

The discussion above should have answered the question as to how the stator knows when to supply mechanical power across the air-gap. When a mechanical load is applied to the shaft, the rotor slows down, the slip increases, rotor currents are induced and their m.m.f. results in a modest (but vitally important) reduction in the air-gap flux wave. This in turn causes a reduction in the e.m.f. induced in the stator windings and therefore an increase in the stator current drawn from the supply. We can anticipate that this is a stable process (at least over the normal operating range) and that the speed will settle when the slip has increased sufficiently that the motor torque equals the load torque.

As far as our conclusions regarding torque are concerned, we see that our original assumption that the flux was constant is near enough correct when the slip is small. We will find it helpful and convenient to continue to treat the flux as constant (for given stator voltage and frequency) when we turn later to methods of controlling the normal running speed.

It has to be admitted, however, that at high values of slip (i.e. low rotor speeds), we cannot expect the main flux to remain constant, and in fact we would find in practice that when the motor was first switched on to the utility supply (50 or 60 Hz), with the rotor stationary, the main flux might typically be only half what it was when the motor was at full speed. This is because at high slips, the leakage fluxes assume a much greater importance than under normal low-slip conditions. The simple arguments we have advanced to predict torque would therefore need to be modified to take account of the reduction of main flux if we wanted to use them quantitatively at high slips. There is no need for us to do this explicitly, but it will be reflected in any subsequent curves portraying typical torque-speed curves for real motors. Such curves are of course used when selecting a motor to run directly from the utility supply, since they provide the easiest means of checking whether the starting and run-up torque is adequate for the job in hand. Fortunately, we will see in Chapter 7 that when the motor is fed from an inverter, we can avoid the undesirable effects of high-slip operation, and guarantee that the flux is at its optimum value at all times.

5.5 Stator current-speed characteristics

To conclude this chapter we will look at how the stator current behaves, remembering that we are assuming that the machine is directly connected to a utility supply of fixed voltage and frequency. Under these conditions the maximum current likely to be demanded and the power factor at various loads are important matters that influence the running cost.

In the previous section, we argued that as the slip increased, and the rotor did more mechanical work, the stator current increased. Since the extra current is associated with the supply of real (i.e. mechanical output) power (as distinct from the original magnetising current which was seen to be reactive), this additional 'work' component of current is more or less in phase with the supply voltage, as shown in the phasor diagrams, Fig. 5.23.

The resultant stator current is the sum of the magnetising current, which is present all the time, and the load component, which increases with the slip.

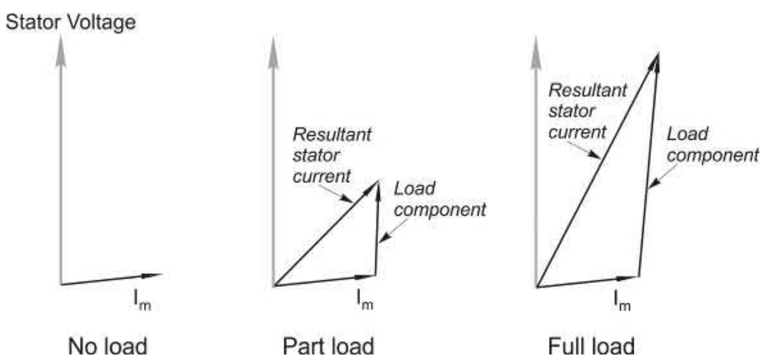

FIG. 5.23 Phasor diagrams showing stator current at no-load, part-load and full-load. The resultant current in each case is the sum of the no-load (magnetising) current and the load component.

We can see that as the load increases, the resultant stator current also increases, and moves more nearly into phase with the voltage. But because the magnetising current is appreciable, the difference in magnitude between no-load and full-load currents may not be all that great. (This is in sharp contrast to the d.c. motor, where the no-load current in the armature is very small in comparison with the full-load current. Note however that in the d.c. motor, the excitation (flux) is provided by a separate field circuit, whereas in the induction motor the stator winding furnishes both the excitation and the work currents. If we consider the behaviour of the work components of current only, both types of machine look very similar.)

The simple ideas behind Fig. 5.23 are based on an approximation, so we cannot push them too far: they are fairly close to the truth for the normal operating region, but break down at higher slips, where the rotor and stator leakage reactances become significant. A typical current locus over the whole range of slips for a cage motor is shown in Fig. 5.24. We note that the power factor is poor when the motor is lightly loaded, and becomes worse again at high slips, and also that the current at standstill (i.e. the 'starting' current) is perhaps five times the full-load value.

Very high currents when started direct-on-line are one of the worst features of the cage induction motor. They not only cause unwelcome volt-drops in the supply system, but also call for heavier switchgear than would be needed to cope with full-load conditions. Unfortunately, for reasons discussed earlier, the high starting currents are not accompanied by high starting torques, as we can see from Fig. 5.25, which shows current and torque as functions of slip for a general-purpose cage motor.

We note that the torque per ampere of current drawn from the utility supply is typically very low at start up, and only reaches a respectable value in the normal operating region, i.e. when the slip is small. This matter is explored further in Chapter 6.

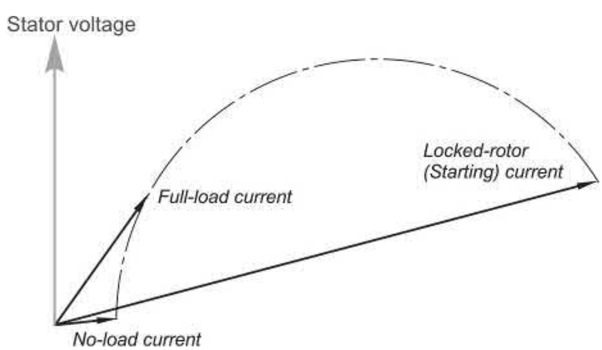

FIG. 5.24 Phasor diagram showing the locus of stator current over the full range of speeds from no-load (full speed) down to the locked-rotor (starting) condition.

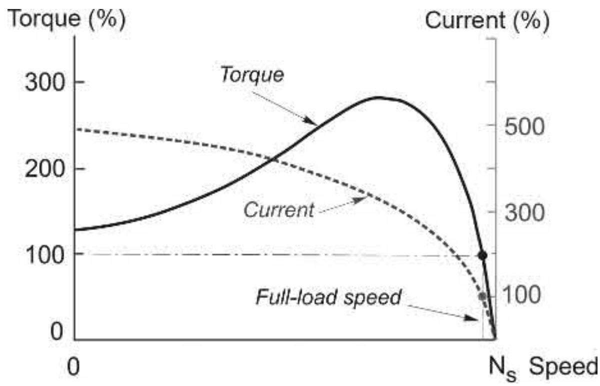

FIG. 5.25 Typical torque-speed and current-speed curves for a cage induction motor. The torque and current axes are scaled so that 100% represents the continuously-rated (full-load) value.

5.6 Review questions

(1) The nameplate of a standard 50 Hz induction motor quotes full-load speed as 2950 rev/min. Find the pole-number and the rated slip.

(2) A four-pole, 60 Hz induction motor runs with a slip of 4%. Find:
- **(a)** the speed;
- **(b)** the rotor frequency;
- **(c)** the speed of rotation of the rotor current wave relative to the rotor surface;
- **(d)** the speed of rotation of the rotor current wave relative to the stator surface.

(3) An induction motor designed for operation from 440 V is supplied at 380 V instead.

What effect will the reduced voltage have on the following:
- **(a)** the synchronous speed;
- **(b)** the magnitude of the air-gap flux;
- **(c)** the induced current in the rotor when running at rated speed;
- **(d)** the torque produced at rated speed.

(4) The rotor from a six-pole induction motor is to be used with a four-pole stator having the same bore and length as the six-pole stator. What modifications would be required to the rotor if it was
- **(a)** a squirrel-cage type, and
- **(b)** a wound-rotor type?

(5) A 440 V, 60 Hz induction motor is to be used on a 50 Hz supply. What voltage should be used?

(6) The stator of a 220 V induction motor is to be rewound for operation from a 440 V supply. The original coils each had 15 turns of 1 mm diameter wire. Estimate the number of turns and diameter of wire for the new stator coils.

(7) The slip of an induction motor driving a constant-torque load is 2.0%. If the voltage is reduced by 5%, estimate:

(a) the new steady-state rotor current, expressed in terms of its original value;

(b) the new steady-state slip.

(8) As the slip of an induction motor increases, the current in the rotor increases, but beyond a certain slip the torque begins to fall. Why is this?

(9) For a given rotor diameter, the stator diameter of a two-pole motor is much greater than the stator diameter of, say, an eight-pole motor. By sketching and comparing the magnetic flux patterns for machines with low and high pole-numbers, explain why more stator iron is required as the pole-number reduces.

(10) The layout of coils for four-pole and six-pole windings in a 36-slot stator are shown in Fig. 5.7. All the coils have the same number of turns of the same wire, the only real difference being that the six-pole coils are of shorter pitch.

Sketch the m.m.f. wave produced by one phase, for each pole-number, assuming that the same current flows in every coil. Hence show that to achieve the same amplitude of flux wave, the current in the six-pole winding would have to be about 50% larger than in the four-pole.

How does this exercise relate to the claim that the power-factor of a low-speed induction motor is usually lower than a high-speed one?

Answers to the review questions are given in the Appendix.

Chapter 6

Induction motor—Operation from 50/60 Hz supply

6.1 Introduction

This chapter is concerned with how the induction motor behaves when connected to a supply of constant voltage and frequency. Despite the widespread use of inverter-fed motors, direct connection to the utility supply remains the most widely used in many application areas.

The key operating characteristics are considered, and we look at how these can be modified to meet the needs of some applications through detailed design. The limits of operation are investigated for the induction machine operating as both a motor and a generator. Methods of speed control, which are not dependent on changing the frequency of the stator supply, are also explored. Finally, whilst the majority of industrial applications utilise the 3 phase induction motor, the role played by single phase motors is acknowledged with a review of the types and characteristics of this variant.

6.2 Methods of starting cage motors

6.2.1 Direct starting—Problems

Our everyday domestic experience is likely to lead us to believe that there is nothing more to starting a motor than closing a switch, and indeed for most low-power machines (up to a few kW)—of whatever type—that is indeed the case. By simply connecting the motor to the supply we set in train a sequence of events which sees the motor draw power from the supply while it accelerates to its target speed. When it has absorbed and converted sufficient energy from electrical to kinetic form, the speed stabilises and the power drawn falls to a low level until the motor is required to do useful mechanical work. In these low-power applications acceleration to full speed may take less than a second, and we are seldom aware of the fact that the current drawn during the acceleration phase is often higher than the continuous rated current.

For motors over a few kW, however, it is necessary to assess the effect on the supply system before deciding whether or not the motor can be started simply by

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00006-4

switching directly onto the supply. If supply systems were ideal (i.e. the supply voltage remained unaffected regardless of how much current was drawn) there would be no problem starting any induction motor, no matter how large. The problem is that the heavy current drawn while the motor is running up to speed may cause a large drop in the system supply voltage, annoying other customers on the same supply and perhaps taking it outside statutory limits.

It is worthwhile reminding ourselves about the influence of supply impedance at this point, as this is at the root of the matter, so we begin by noting that any supply system, no matter how complicated, can be modelled by means of the delightfully simple Thévenin equivalent circuit shown in Fig. 6.1. (We are assuming balanced 3-phase operation, so a 1-phase equivalent circuit will suffice.)

The supply is represented by an ideal voltage source (V_s) in series with the supply impedance Z_s. When no load is connected to the supply, and the current is zero, the terminal voltage is V_s; but as soon as a load is connected the load current (I) flowing through the source impedance results in a volt-drop, and the output voltage falls from V_s to V, where

$$V = V_s - IZ_s \tag{6.1}$$

For most industrial supplies the source impedance is predominantly inductive, so that Z_s is simply an inductive reactance, X_s. Typical phasor diagrams relating to a supply with a purely inductive reactance are shown in Fig. 6.2: in (a) the load is also taken to be purely reactive, while the load current in (b) has the same magnitude as in (a) but the load is resistive. The output (terminal) voltage in each case is represented by the phasor labelled V.

For the inductive load (a) the current lags the terminal voltage by 90° while for the resistive load (b) the current is in phase with the terminal voltage. In both cases the volt-drop across the supply reactance (IX_s) leads the current by 90°.

The first point to note is that, for a given magnitude of load current, the volt-drop is in phase with V_s when the load is inductive, whereas with a resistive load the volt-drop is almost at 90° to V_s. This results in a much greater fall in the magnitude of the output voltage when the load is inductive than when it is

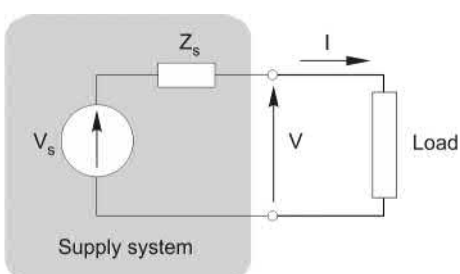

FIG. 6.1 Equivalent circuit of supply system.

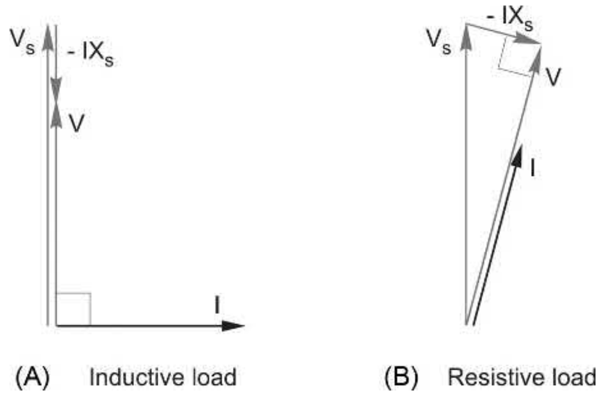

FIG. 6.2 Phasor diagrams showing the effect of supply-system impedance on the output voltage with (A) inductive load and (B) resistive load.

resistive. The second—obvious—point is that the larger the current, the more the drop in voltage.

Unfortunately, when we try to start a large cage induction motor we face a double-whammy because not only is the starting current typically five or six times rated current, but it is also at a low power-factor, i.e. the motor looks predominantly inductive when the slip is high. (In contrast, when the machine is up to speed and fully loaded, its current is perhaps only one fifth of its starting current and it presents a predominantly resistive appearance as seen by the supply. Under these conditions the supply voltage is hardly any different from at no-load.)

Since the drop in voltage is attributable to the supply impedance, it follows that if we want to be able to draw a large starting current without upsetting other consumers, it would clearly be best for the supply impedance to be as low as possible, and preferably zero. But from the supply authority viewpoint a very low supply impedance brings the problem of how to cope in the event of an accidental short-circuit across the terminals. The short circuit current is inversely proportional to the supply impedance, and tends to infinity as Z_s approaches zero. The cost of providing the switchgear to clear such a large fault current would be prohibitive, so a compromise always has to be reached, with values of supply impedances being set by the supply authority to suit the anticipated demands.

Systems with low internal impedance are known as 'stiff' supplies, because the voltage is almost constant regardless of the current drawn. (An alternative way of specifying the nature of the supply is to consider the fault current that would flow if the terminals were short-circuited: a system with a low impedance would have a high fault current or 'fault level'.) Starting on a stiff supply requires no special arrangements and the three motor leads are simply switched directly onto the utility supply terminals. This is known as 'direct-on-line'

(DOL) or 'direct-to-line' (DTL) starting. The switching will usually be done by means of a contactor, incorporating fuses and other overload/thermal protection devices, and operated manually by local or remote pushbuttons, or interfaced to permit operation from a programmable controller or computer.

In contrast, if the supply impedance is high (i.e. a low fault level) an appreciable volt-drop will occur every time the motor is started, causing lights to dim and interfering with other apparatus on the same supply. With this 'weak' supply, some form of starter is called for to limit the current at starting and during the run-up phase, thereby reducing the magnitude of the volt-drop imposed on the supply system. As the motor picks up speed, the current falls, so the starter is removed as the motor approaches full speed. Naturally enough the price to be paid for the reduction in current is a lower starting torque, and a longer run-up time.

Whether or not a starter is required depends on the size of the motor in relation to the capacity or fault-level of the supply, the prevailing regulations imposed by the supply authority, and the nature of the load.

The references above to 'low' and 'high' supply impedances must therefore be interpreted in relation to the impedance of the motor when it is stationary. A large (and therefore low impedance) motor could well be started quite happily direct-on-line in a major industrial plant, where the supply is 'stiff' i.e. the supply impedance is very much less than the motor impedance. But the same motor would need a starter when used in a rural setting remote from the main power system, and fed by a relatively high impedance or 'weak' supply. Needless to say, the stricter the rules governing permissible volt-drop, the more likely it is that a starter will be needed.

Motors which start without significant load torque or inertia can accelerate very quickly, so the high starting current is only drawn for a short period. A 10 kW motor would be up to speed in a second or so, and the volt-drop may therefore be judged as acceptable. Clutches are sometimes fitted to permit 'off-load' starting, the load being applied after the motor has reached full speed. Conversely, if the load torque and/or inertia are high, the run-up may take many seconds, in which case a starter may prove essential. No strict rules can be laid down, but obviously the bigger the motor, the more likely it is to require a starter.

6.2.2 Star/delta (wye/mesh) starter

This is the simplest and most widely used method of starting. It provides for the windings of the motor to be connected in star (wye) to begin with, thereby reducing the voltage applied to each phase to 58% (i.e. $1/\sqrt{3}$) of its direct-on-line value. Then, when the motor speed approaches its running value, the windings are switched to delta (mesh) connection. The main advantage of the method is its simplicity, while its main drawbacks are that the starting torque is reduced (see below), and the sudden transition from star to delta gives rise to a second shock—albeit of lesser severity—to the supply system and to the load. For star/delta switching to be possible both ends of each phase of the motor windings must be brought out to the terminal box. This requirement is met in the majority of motors, except small ones, which are usually permanently connected in delta.

With a star/delta starter the current drawn from the supply is approximately one third of that drawn in a direct-on-line start, which is very welcome, but at the same time the starting torque is also reduced to one third of its direct-on-line value. Naturally we need to ensure that the reduced torque will be sufficient to accelerate the load, and bring it up to a speed at which it can be switched to delta without an excessive jump in the current.

Various methods are used to detect when to switch from star to delta. Historically, in manual starters, the changeover is determined by the operator watching the ammeter until the current has dropped to a low level, or listening to the sound of the motor until the speed becomes steady. Automatic versions are similar in that they detect either falling current, or speed rising to a threshold level, or, where the load is always the same, they operate after a pre-set time.

6.2.3 Autotransformer starter

A three-phase autotransformer is usually used where star-delta starting provides insufficient starting torque. Each phase of an autotransformer consists of a single winding on a laminated core. The incoming supply is connected across the ends of the coils, and one or more tapping points (or a sliding contact) provide a reduced voltage output, as shown in Fig. 6.3.

The motor is first connected to the reduced voltage output, and when the current has fallen to the running value, the motor leads are switched over to the full voltage.

If the reduced voltage is chosen so that a fraction α of the line voltage is used to start the motor, the starting torque is reduced to approximately α^2 times its direct-on-line value, and the current drawn from the supply is also reduced to α^2

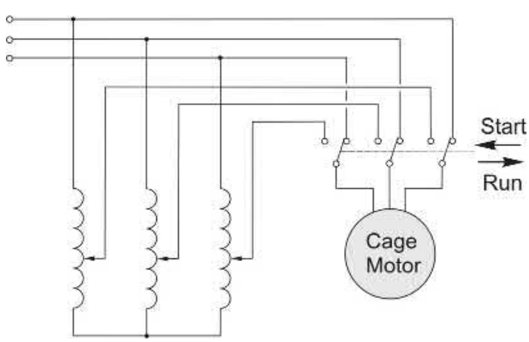

FIG. 6.3 Autotransformer starter for cage induction motor.

times its direct value. As with the star/delta starter, the torque per ampere of supply current is the same as for a direct start.

6.2.4 Resistance or reactance starter

By inserting three resistors or inductors of appropriate value in series with the motor, the starting current can be reduced by any desired extent, but only at the expense of a disproportionate reduction in starting torque.

For example if the current is reduced to half its direct-on-line value, the motor voltage will be halved, so the torque (which is proportional to the square of the voltage—see later) will be reduced to only 25% of its direct-on-line value. This approach is thus less attractive in terms of torque per ampere of supply current than the star/delta method. One attractive feature, however, is that as the motor speed increases and its effective impedance rises, the volt-drop across the extra impedance reduces, so the motor voltage rises progressively with the speed, thereby giving more torque. When the motor is up to speed, the added impedance is shorted-out by means of a contactor.

6.2.5 Solid-state soft starting

This method is now widely used. It provides a smooth build-up of current and torque, the maximum current and acceleration time are easily adjusted, and it is particularly valuable where the load must not be subjected to sudden jerks. The only real drawback over conventional starters is that the utility supply currents during run-up are not sinusoidal, which can lead to interference with other equipment on the same supply.

The most widely-used arrangement comprises three pairs of back-to-back thyristors (or triacs) connected in series with the three supply lines, as shown in Fig. 6.4A.

Each thyristor is fired once per half-cycle, the firing being synchronised with the utility supply and the firing angle being variable so that each pair conducts for a controllable proportion of a cycle. Typical current waveforms are shown in Fig. 6.4B: they are clearly not sinusoidal but the motor will tolerate them quite happily.

A wide variety of control philosophies can be found, with the degree of complexity and sophistication being reflected in the price. The cheapest open-loop systems simply alter the firing angle linearly with time, so that the voltage applied to the motor increases as it accelerates. The ‘ramp-time’ can be set by trial and error to give an acceptable start, i.e. one in which the maximum allowable current from the supply is not exceeded at any stage. This approach is reasonably satisfactory when the load remains the same, but requires resetting each time the load changes. Loads with high static friction are a problem because nothing happens for the first part of the ramp, during which time the motor torque is insufficient to move the load. When the load finally moves,

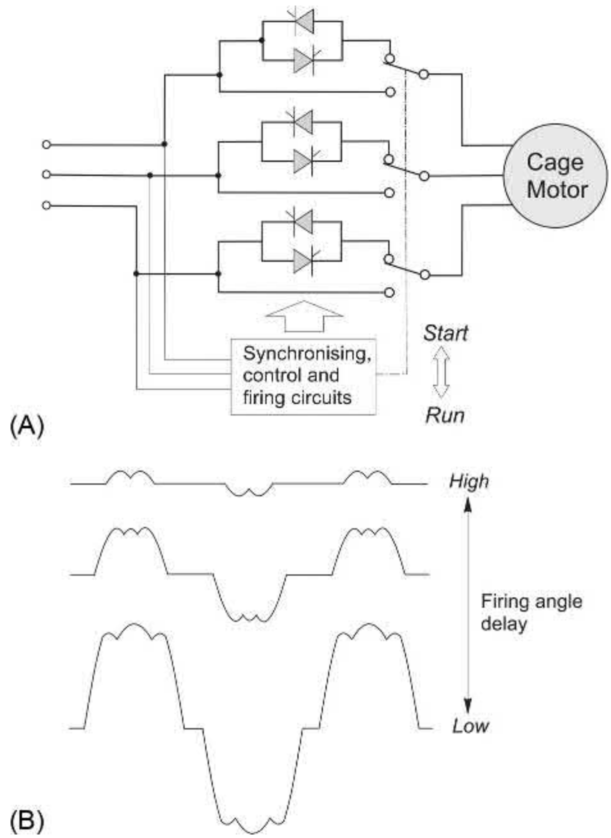

FIG. 6.4 (A) Thyristor soft-starter, (B) typical motor current waveforms.

its acceleration is often too rapid. The more advanced open-loop versions allow the level of current at the start of the ramp to be chosen, and this is helpful with 'sticky' loads.

More sophisticated systems—usually with on-board digital controllers—provide for tighter control over the acceleration profile by incorporating an inner current-control loop. After an initial ramping up to the start level (over the first few cycles), the current is held constant at the desired level throughout the accelerating period, the firing angle of the thyristors being continually adjusted to compensate for the changing effective impedance of the motor. By keeping the current at the highest value which the supply can tolerate the run-up time is minimised. As with the open-loop systems the velocity-time profile is not necessarily ideal, since with constant current the motor torque exhibits a very sharp rise as the pull-out slip is reached, resulting in a sudden surge in speed. Some systems also include a motor model, which estimates speed and allows the controller to follow a ramp or other speed-time profile.

Prospective users need to be wary of some of the promotional literature: claims are sometimes made that massive reductions in starting current can be achieved without corresponding reductions in starting torque. This is nonsense: the current can certainly be limited, but as far as torque per line amp is concerned soft-start systems are no better than series reactor systems, and not as good as the autotransformer and star/delta methods. Caution should also be exercised in relation to systems that use only one or two triacs: these are fairly common in smaller sizes (<50 kW). Although they do limit the current in one or two phases as compared with direct-on-line starting, the unbalanced currents distort the air-gap flux and this gives rise to uneven (pulsating) torque.

6.2.6 Starting using a variable-frequency inverter

Operation of induction motors from variable-frequency inverters is discussed in Chapters 7 and 8, but it is appropriate to mention here that one of the advantages of inverter-fed operation is that starting is not a problem because it is usually possible to obtain rated torque from standstill up to rated speed without drawing an excessive current from the utility supply. None of the other starting methods we have looked at have this ability, so in some applications it may be that the comparatively high cost of the inverter is justified solely on the grounds of its starting and run-up potential.

6.3 Run-up and stable operating regions

In addition to having sufficient torque to start the load it is obviously necessary for the motor to bring the load up to full speed. To predict how the speed will rise after switching-on we need the torque-speed curves of the motor and the load, and the total inertia.

By way of example, we can look at the case of a motor with two different loads (Fig. 6.5). The solid line is the torque-speed curve of the motor, while the dotted lines represent two different load characteristics. Load (A) is typical of a

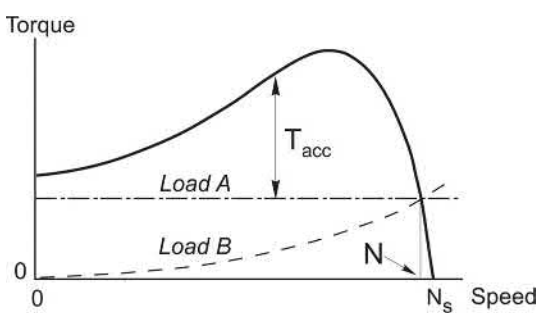

FIG. 6.5 Typical torque-speed curve showing two different loads which have the same steady running speed (N).

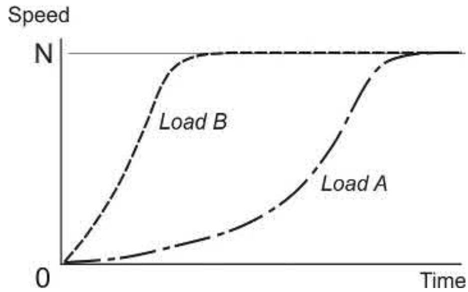

FIG. 6.6 Speed-time curves during run-up, for motor and loads shown in Fig. 6.5.

simple hoist, which applies constant torque to the motor at all speeds, while load (B) might represent a fan. For the sake of simplicity, we will assume that the load inertias (as seen at the motor shaft) are the same.

The speed-time curves for run-up are shown in Fig. 6.6. Note that the gradient of the speed-time curve (i.e. the acceleration) is obtained by dividing the accelerating torque T_{acc} (which is the difference between the torque developed by the motor and the torque required to run the load at that speed) by the total inertia.

In this example, both loads ultimately reach the same steady speed, N, (i.e. the speed at which motor torque equals load torque), but B reaches full speed much more quickly because the accelerating torque is higher during most of the run-up. Load A picks up speed slowly at first, but then accelerates hard (often with a characteristic 'whoosh' produced by the ventilating fan) as it passes through the peak torque speed and approaches equilibrium conditions.

It should be clear that the higher the total inertia, the slower the acceleration, and vice-versa. The total inertia means the inertia as seen at the motor shaft, so if gearboxes or belts are employed the inertia must be 'referred' as discussed in Chapter 11.

An important qualification ought to be mentioned in the context of the motor torque-speed curves shown by the solid line in Fig. 6.5. This is that curves like this represent the torque developed by the motor when it has settled down at the speed in question, i.e. they are the true steady-state curves. In reality, a motor will generally only be in a steady-state condition when it settles at its normal running speed, so for most of the speed range the motor will be accelerating.

In particular, when the motor is first switched on, there will be a transient period as the three currents gradually move towards a balanced three-phase pattern. During this period the torque can fluctuate wildly, with brief negative excursions, typically as shown in Fig. 6.7 which relates to a small unloaded motor. During the transient period the average torque may be very low (as in Fig. 6.7) in which case acceleration only begins in earnest after the first few cycles. In this particular example the transient persists long enough to cause an overshoot of the steady-state speed and an oscillation before settling.

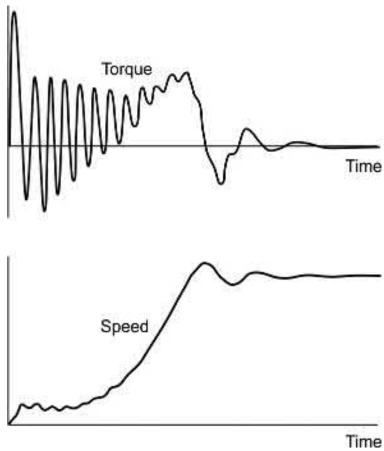

FIG. 6.7 Run-up of unloaded motor, showing torque transients persisting for the first few cycles of the supply.

Fortunately, the average torque during run-up can be fairly reliably obtained from the steady-state curves (usually available from the manufacturer), particularly if the inertia is high and the motor takes many cycles to reach full speed, in which case we would consider the torque-speed curve as being 'quasi-steady-state'.

6.3.1 Harmonic effects—Skewing

A further cautionary note in connection with the torque-speed curves shown in this and most other books relates to the effects of harmonic air-gap fields. In Chapter 5 it was explained that despite the limitations imposed by slotting, the stator winding m.m.f. is remarkably close to the ideal of a pure sinusoid. Unfortunately, because it is not a perfect sinusoid, Fourier analysis reveals that in addition to the predominant fundamental component, there are always additional unwanted 'space harmonic' fields. These harmonic fields have synchronous speeds that are inversely proportional to their order. For example a 4-pole, 50 Hz motor will have a main field rotating at 1500 rev/min, but in addition there may be a 5th harmonic (20-pole) field rotating in the reverse direction at 300 rev/min, a 7th harmonic (28-pole) field rotating forwards at 214 rev/min, etc. These space harmonics are minimised by stator winding design, but not all can be eliminated.

If the rotor has a very large number of bars it will react to the harmonic field in much the same way as to the fundamental, producing additional induction-motor torques centred on the synchronous speed of the harmonic, and leading to unwanted dips in the torque speed, typically as shown in Fig. 6.8.

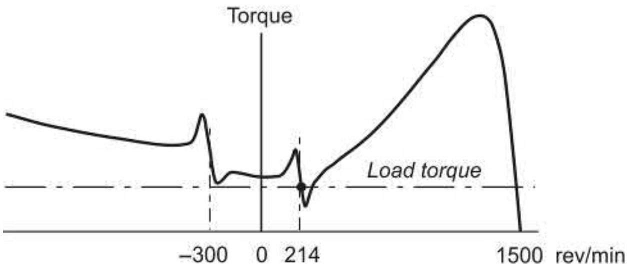

FIG. 6.8 Torque-speed curve showing the effect of space harmonics, and illustrating the possibility of a motor 'crawling' on the 7th harmonic.

Users should not be too alarmed as in most cases the motor will ride-through the harmonic during acceleration, but in extreme cases a motor might for example stabilise on the 7th harmonic, and 'crawl' at about 214 rev/min, rather than running up to 4-pole speed (1500 rev/min at 50 Hz), as shown by the dot in Fig. 6.8.

To minimise the undesirable effects of space harmonics the rotor bars in the majority of induction motors are not parallel to the axis of rotation, but instead they are skewed (typically by around one or two slot-pitches) along the rotor length, as shown in Fig. 6.9. This has very little effect as far as the fundamental field is concerned, but can greatly reduce the response of the rotor to harmonic fields.

FIG. 6.9 Cutaway view of cage induction motor showing skewed rotor bars. The bars are of cast aluminium, and the paddles attached to the end rings serve as a fan for circulating internal air. The external fan mounted on the non-drive end cools the finned stator casing. *(Courtesy of Siemens.)*

Because the overall influence of the harmonics on the steady-state curve is barely noticeable, and their presence might worry users, they are rarely shown, the accepted custom being that 'the' torque-speed curve represents the behaviour due to the fundamental component only.

6.3.2 High inertia loads—Overheating

Apart from accelerating slowly, high inertia loads pose a particular problem of rotor heating which can easily be overlooked by the unwary user. Every time an induction motor is started from rest and brought up to speed, the total energy dissipated as heat in the motor windings is equal to the stored kinetic energy of the motor plus load. Hence with high inertia loads, very large amounts of energy are released as heat in the windings during run-up, even if the load torque is negligible when the motor is up to speed. With totally-enclosed motors the heat ultimately has to find its way to the finned outer casing of the motor, which is cooled by air from the shaft-mounted external fan. Cooling of the rotor is therefore usually much worse than the stator, and the rotor is thus most likely to overheat during high inertia run-ups.

No hard and fast rules can be laid down, but manufacturers usually work to standards which specify how many starts per hour can be tolerated. Actually, this information is useless unless coupled with reference to the total inertia, since doubling the inertia makes the problem twice as bad. However, it is usually assumed that the total inertia is not likely to be more than twice the motor inertia, and this is certainly the case for most loads. If in doubt, the user should consult the manufacturer who may recommend a larger motor than might seem necessary if it was simply a matter of meeting the full-load power requirements.

6.3.3 Steady-state rotor losses and efficiency

The discussion above is a special case, which highlights one of the less attractive features of induction machines. This is that it is never possible for all the power crossing the air-gap from the stator to be converted to mechanical output, because some is always lost as heat in the rotor circuit resistance. In fact, it turns out that at slip s the total power (P_r) crossing the air-gap always divides so that a fraction sP_r is lost as heat, while the remainder $(1 - s)P_r$ is converted to useful mechanical output.

Hence when the motor is operating in the steady-state the energy-conversion efficiency *of the rotor* is given by

$$\eta_r = \frac{\text{Mechanical output power}}{\text{Power into rotor}} = (1 - s) \tag{6.2}$$

This result is very important, and shows us immediately why operating at small values of slip is desirable. With a slip of 5% (or 0.05) for example, 95% of the air-gap power is put to good use. But if the motor was run at half the synchronous speed ($s = 0.5$), 50% of the air-gap power would be wasted as heat in the rotor.

We can also see that the overall efficiency of the motor must always be significantly less than $(1 - s)$, because in addition to the rotor copper losses there are stator copper losses, iron losses and windage and friction losses. This fact is sometimes forgotten, leading to conflicting claims such as 'full-load slip = 5%, overall efficiency = 96%', which is clearly impossible.

6.3.4 Steady-state stability—Pull-out torque and stalling

We can check stability by asking what happens if the load torque suddenly changes for some reason. The load shown by the dotted line in Fig. 6.10 is stable at speed X, for example: if the load torque increased from T_a to T_b, the load torque would be greater than the motor torque, so the motor would decelerate. As the speed dropped, the motor torque would rise, until a new equilibrium was reached, at the slightly lower speed (Y). The converse would happen if the load torque reduced, leading to a higher stable running speed.

But what happens if the load torque is increased more and more? We can see that as the load torque increases, beginning at point X, we eventually reach point Z, at which the motor develops its maximum torque. Quite apart from the fact that the motor is now well into its overload region, and will be in danger of overheating, it has also reached the limit of stable operation. If the load torque is further increased, the speed falls (because the load torque is more than the motor torque), and as it does so the shortfall between motor torque and load torque becomes greater and greater. The speed therefore falls faster and faster, and the motor is said to be 'stalling'. With loads such as machine tools (a drilling machine, for example), as soon as the maximum or 'pull-out' torque is exceeded, the motor rapidly comes to a halt, making an angry humming sound. With a hoist, however, the excess load would cause the rotor to be accelerated in the reverse direction, unless it was prevented from doing so by a mechanical brake.

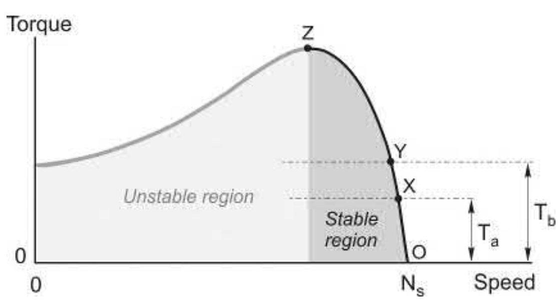

FIG. 6.10 Torque-speed curve illustrating stable operating region (0XYZ).

6.4 Torque-speed curves—Influence of rotor parameters

We saw earlier that the rotor resistance and reactance influenced the shape of the torque-speed curve. Both of these parameters can be varied by the designer, and we will explore the pros and cons of the various alternatives. To limit the mathematics the discussion will be mainly qualitative, but it is worth mentioning that the whole matter can be dealt with rigorously using the equivalent circuit approach.[1]

We will deal with the cage rotor first because it is the most important, but the wound rotor allows a wider variation of resistance to be obtained, so it is discussed later.

6.4.1 Cage rotor

For small values of slip, i.e. in the normal running region, the lower we make the rotor resistance the steeper the slope of the torque-speed curve becomes, as shown in Fig. 6.11. We can see that at the rated torque (shown by the horizontal dotted line in Fig. 6.11) the full-load slip of the low-resistance cage is much lower than that of the high-resistance cage. But we saw earlier that the rotor efficiency is equal to $(1 - s)$, where s is the slip. So we conclude that the low resistance rotor not only gives better speed holding, but is also much more efficient. There is of course a limit to how low we can make the resistance: copper allows us to achieve a lower resistance than aluminium, but we can't do any better than fill the slots with solid copper bars.

As we might expect there are drawbacks with a low resistance rotor. The direct-on-line starting torque is reduced (see Fig. 6.11), and worse still the

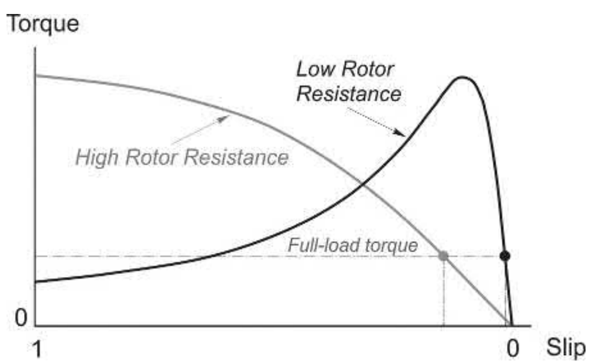

FIG. 6.11 Influence of rotor resistance on torque-speed curve of cage motor. The full-load running speeds are indicated by the vertical dotted lines.

1. See for example The Control Techniques Drives and Controls Handbook, 2nd Edition, by W Drury, pages 38–43.

starting current is increased. The lower starting torque may prove insufficient to accelerate the load, while increased starting current may lead to unacceptable volt-drops in the supply.

Altering the rotor resistance has little or no effect on the value of the peak (pull-out) torque, but the slip at which the peak torque occurs is directly proportional to the rotor resistance. By opting for a high enough resistance (by making the cage from bronze, brass or other relatively high-resistivity material) we can if we wish arrange for the peak torque to occur at or close to starting, as shown in Fig. 6.11. The snag in doing this is that the full-load efficiency is inevitably low because the full-load slip will be high (see Fig. 6.11).

There are some applications for which high-resistance motors were traditionally well suited, an example being for metal punching presses, where the motor accelerates a flywheel which is used to store energy. In order to release a significant amount of energy, the flywheel slows down appreciably during impact, and the motor then has to accelerate it back up to full speed. The motor needs a high torque over a comparatively wide speed range, and does most of its work during acceleration. Once up to speed the motor is effectively running light, so its low efficiency is of little consequence. (We should note that this type of application is now often met by drives, but because the induction motor is so robust and long-lived, significant numbers of 'heritage' installations will persist, particularly in the developing world.)

High-resistance motors are sometimes used for speed control of fan-type loads, and this is taken up again later when we explore speed control.

To sum up, a high rotor resistance is desirable when starting and at low speeds, while a low resistance is preferred under normal running conditions. To get the best of both worlds, we need to be able to alter the resistance from a high value at starting to a lower value at full speed. Obviously we can't change the actual resistance of the cage once it has been manufactured, but it is possible to achieve the desired effect with either a 'double cage' or a 'deep bar' rotor.

6.4.2 Double cage and deep bar rotors

Double cage rotors have an outer cage made of relatively high resistivity material such as bronze, and an inner cage of low resistivity, usually copper, as shown on the left in Fig. 6.12.

The inner cage of low resistance copper is sunk deep into the rotor, so that it is almost completely surrounded by iron. This causes the inner bars to have a much higher leakage inductance than if they were near the rotor surface, so that

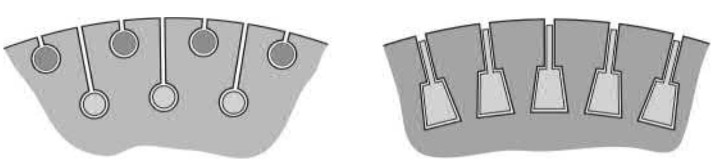

FIG. 6.12 Double cage (left) and deep bar (right) rotors.

under starting conditions (when the induced rotor frequency is high) their inductive reactance is very high and little current flows in them. In contrast, the bars of the outer cage (of higher resistance bronze) are placed so that their leakage fluxes face a much higher reluctance path, leading to a low leakage inductance. Hence under starting conditions, rotor current is concentrated in the outer cage, which, because of its high resistance, produces a high starting torque.

At the normal running speed the roles are reversed. The rotor frequency is low, so both cages have low reactance and most of the current therefore flows in the low-resistance inner cage. The torque-speed curve is therefore steep, and the efficiency is high.

Considerable variation in detailed design is possible in order to shape the torque-speed curve to particular requirements. In comparison with a single cage rotor, the double cage gives much higher starting torque, substantially less starting current, and marginally worse running performance.

The deep bar rotor has a single cage, usually of copper, formed in slots which are deeper and narrower than in a conventional single-cage design. Construction is simpler and therefore cheaper than in a double-cage rotor, as shown on the right in Fig. 6.12.

The deep bar approach ingeniously exploits the fact that the effective resistance of a conductor is higher under a.c. conditions than under d.c. conditions. With a typical copper bar of the size used in an induction motor rotor, the difference in effective resistance between d.c. and say 50 or 60 Hz (the so-called 'skin-effect') would be negligible if the conductor was entirely surrounded by air. But when it is almost completely surrounded by iron, as in the rotor slots, its effective resistance at supply frequency may be two or three times its d.c. value.

Under starting conditions, when the rotor frequency is equal to the supply frequency, the skin effect is very pronounced, and the rotor current is concentrated towards the top of the slots. The effective resistance is therefore increased, resulting in a high starting torque from a low starting current. When the speed rises and the rotor frequency falls, the effective resistance reduces towards its d.c. value, and the current distributes itself more uniformly across the cross-section of the bars. The normal running performance thus approaches that of a low-resistance single-cage rotor, giving a high efficiency and stiff torque-speed curve. The pull-out torque is however somewhat lower than for an equivalent single-cage motor because of the rather higher leakage reactance.

Most small and medium motors are designed to exploit the deep bar effect to some extent, reflecting the view that for most applications the slightly inferior running performance is more than outweighed by the much better starting behaviour. A typical torque-speed curve for a general-purpose medium-size (55 kW) motor is shown in Fig. 6.13. Such motors are unlikely to be described by the maker specifically as 'deep-bar' but they nevertheless incorporate a measure of the skin effect and consequently achieve the 'good' torque-speed characteristic shown by the solid line in Fig. 6.13.

The current-speed relationship is shown by the dotted line in Fig. 6.13, both torque and current scales being expressed in per-unit (p.u.). This notation is

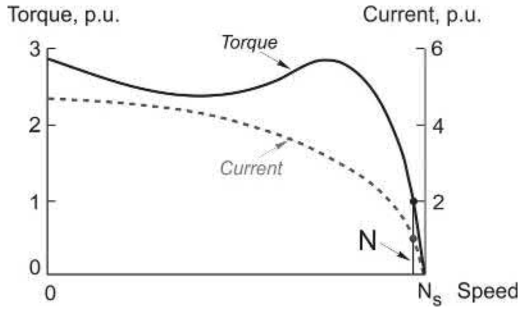

FIG. 6.13 Typical torque-speed and current-speed curves for a general-purpose industrial double cage motor.

widely used as a shorthand, with 1 p.u. (or 100%) representing rated value. For example a torque of 1.5 p.u. simply means one and a half times the rated value, while a current of 400% means a current of four times the rated value.

6.4.3 Starting and run-up of slipring motors

By adding external resistance in series with the rotor windings the starting current can be kept low but at the same time the starting torque is high. This was the major advantage of the wound-rotor or slipring motor, which made it well suited for loads with heavy starting duties such as stone-crushers, cranes and conveyor drives, for most of which the inverter-fed cage motor is now preferred.

The influence of rotor resistance is shown by the set of torque-speed curves in Fig. 6.14.

Typically the resistance at starting would be selected to give full-load torque together with rated current from the utility supply. The starting torque is then as indicated by point A in Fig. 6.14.

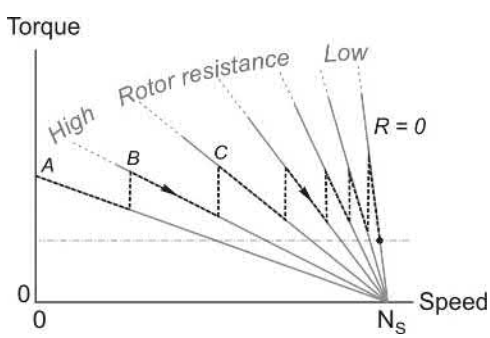

FIG. 6.14 Torque-speed curves for a wound-rotor (slipring) motor showing how the external rotor-circuit resistance (R) can be varied in steps to provide an approximately constant torque during acceleration.

As the speed rises, the torque would fall more or less linearly if the resistance remained constant, so in order to keep close to full torque the resistance is gradually reduced, either in steps, in which case the trajectory ABC, etc. is followed (Fig. 6.14), or continuously so that maximum torque is obtained throughout. Ultimately the external resistance is made zero by shorting-out the sliprings, and thereafter the motor behaves like a low-resistance cage motor, with a high running efficiency.

6.5 Influence of supply voltage on torque-speed curve

We established earlier that at any given slip, the air-gap flux density is proportional to the applied voltage, and the induced current in the rotor is proportional to the flux density. The torque—which depends on the product of the flux and the rotor current - therefore depends on the square of the applied voltage. This means that a comparatively modest fall in the voltage will result in a much larger reduction in torque capability, with adverse effects which may not be apparent to the unwary until too late.

To illustrate the problem, consider the torque-speed curves for a cage motor shown in Fig. 6.15. The curves (which have been expanded to focus attention on the low-slip region) are drawn for full voltage (100%), and for a modestly reduced voltage of 90%. With full voltage and full-load torque the motor will run at point X, with a slip of say 5%. Since this is the normal full-load condition, the rotor and stator currents will be at their rated values.

Now suppose that the voltage falls to 90%. The load torque is assumed to be constant so the new operating point will be at Y. Because the air-gap flux density is now only 0.9 of its rated value, the rotor current will have to be about 1.1 times rated value to develop the same torque, so the rotor e.m.f. is required to increase by 10%. But the flux density has fallen by 10%, so an increase in slip of 20% is called for. The new slip is therefore 6%.

The drop in speed from 95% of synchronous to 94% may well not be noticed, and the motor will apparently continue to operate quite happily. But

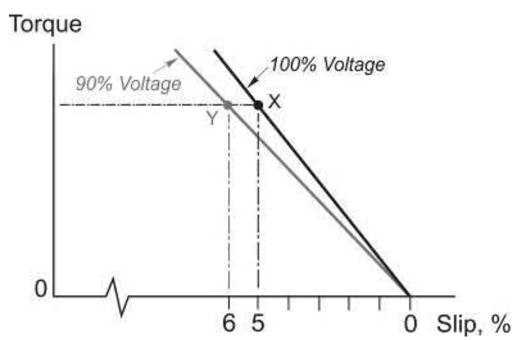

FIG. 6.15 Influence of stator supply voltage on torque-speed curves.

the rotor current is now 10% above its rated value, so the rotor heating will be 21% more than is allowable for continuous running. The stator current will also be above rated value, so if the motor is allowed to run continuously, it will overheat. This is one reason why all large motors are fitted with over-temperature protection. Many small and medium motors do not have such protection, so it is important to guard against the possibility of undervoltage operation.

Another potential danger arises if the supply voltage is unbalanced, i.e. the three line-to-line voltages are unequal. This is most likely where the supply impedance is high and the line currents are unequal due to unbalanced loads elsewhere on the system.

Electrical engineers use the technique of 'symmetrical components' to analyse unbalanced three-phase voltages. It is a method whereby the effect of unbalanced voltages on, say, a motor is quantified by finding how the motor behaves when subjected, independently, to three sets of balanced voltages that together simulate the actual unbalanced voltages.

The first balanced set is the positive sequence component: its phase-sequence is the normal one, e.g. UVW or ABC; the second is the negative sequence, having the phase sequence WVU; and the third is the zero sequence, in which the three phase voltage components are co-phasal. There are simple analytic formulae for finding the components from the original three unbalanced voltages.

If the supply is balanced, the negative sequence and zero sequence components are both zero. Any unbalance gives rise to a negative sequence component, which will set up a rotating magnetic field travelling in the opposite direction to that of the main (positive sequence) component, and thus cause a braking torque and increase the losses, especially in the rotor. The zero sequence components produce a stationary field with three times the pole-number of the main field, but this can only happen when the motor has a star point connection.

Perhaps the most illuminating example of the application of symmetrical components as far as we are concerned is in the case of single-phase machines, where a single winding produces a pulsating field. This is an extreme case of unbalance because two of the phases are non-existent! It turns out that the positive sequence and negative sequence components are equal, which leads naturally to the idea that the pulsating field can be resolved into two counter-rotating fields. We will see later in this chapter that we can extend this picture to understand how the single-phase induction motor works.

Returning to three phase motors, a quite modest unbalance can be serious in terms of overheating. For example, to meet international standards motors are expected to tolerate only 1% negative sequence voltage continuously. Large motors often have negative sequence protection fitted, while small ones will rely on their thermal protection device to prevent overheating. Alternatively, if continuous operation under unbalanced conditions is required the motor must be de-rated significantly, e.g. to perhaps 80% of full load with a voltage unbalance of 4%.

6.6 Generating

Having explored the torque-speed curve for the normal motoring region, where the speed lies between zero and just below synchronous, we must ask what happens if the speed is above the synchronous speed, i.e. the slip is negative. The unambiguous answer to this question is that the machine will switch from motoring to generating. Strangely, as mentioned previously, the authors have often encountered users who express deep scepticism, or even outright disbelief at the prospect of induction machines generating, so it is important that we attempt to counter what is clearly a widely-held misconception.

A typical torque-speed curve for a cage motor covering the full range of speeds which are likely in practice is shown in Fig. 6.16.

We can see from Fig. 6.16 that the decisive factor as far as the direction of the torque is concerned is the slip, rather than the speed. When the slip is positive the torque is positive, and vice-versa. The torque therefore always acts so as to urge the rotor to run at zero slip, i.e. at the synchronous speed. If the rotor is tempted to run faster than the field it will be slowed down, whilst if it is running below synchronous speed it will be urged to accelerate forwards. In particular, we note that for slips >1, i.e. when the rotor is running backwards (i.e. in the opposite direction to the field), the torque will remain positive, so that if the rotor is unrestrained it will first slow down and then change direction and accelerate in the direction of the field.

6.6.1 Generating region

For negative slips, i.e. when the rotor is turning in the same direction, but at a higher speed than the travelling field, the 'motor' torque is in fact negative. In other words the machine develops a torque that opposes the rotation, which can therefore only be maintained by applying a driving torque to the shaft. In this

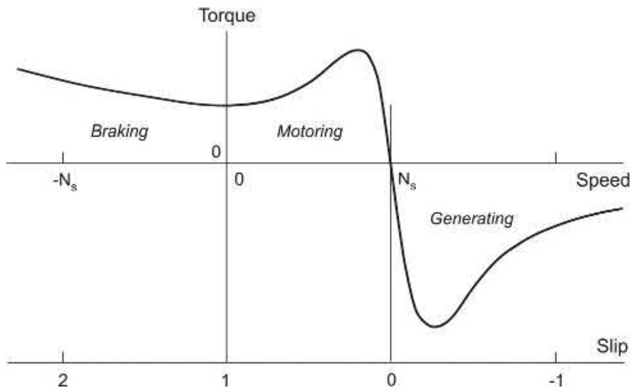

FIG. 6.16 Torque-speed curve over motoring region (slip between 0 and 1), braking region (slip >1) and generating region (negative slip).

region the machine acts as an induction generator, converting mechanical power from the shaft into electrical power into the supply system. Cage induction machines are used in this way in wind-power generation schemes, as described in a later section.

It is worth stressing that, just as with the d.c. machine, we do not have to make any changes to an induction motor to turn it into an induction generator. In both cases, all that is needed is a source of mechanical power to turn the rotor faster than it would run if there was zero load or friction torque. For the d.c. motor, its ideal no-load speed is that at which its back e.m.f. equals the supply voltage, whereas for the induction motor, it is the synchronous speed.

On the other hand, we should be clear that, unlike the d.c. machine, the induction machine can only generate when it is connected to the supply. If we disconnect an induction motor from the utility supply and try to make it generate simply by turning the rotor we will not get any output because there is nothing to set up the working flux: the flux (excitation) is not present until the motor is supplied with magnetising current from the supply. It seems likely that this apparent inability to generate in isolation is what gave rise to the myth that induction machines cannot generate at all: a widely held view, but wholly incorrect!

There are comparatively few applications in which utility-supplied motors find themselves in the generating region, though as we will see later it is quite common in inverter-fed drives. We will however look at one example of a utility supply fed motor in the so-called 'regenerative' mode to underline the value of the motor's inherent ability to switch from motoring to generating automatically, without the need for any external intervention.

Consider a cage motor driving a simple hoist through a reduction gearbox, and suppose that the hook (unloaded) is to be lowered. Because of the static friction in the system, the hook will not descend on its own, even after the brake is lifted, so on pressing the 'down' button the brake is lifted and power is applied to the motor so that it rotates in the lowering direction. The motor quickly reaches full speed and the hook descends. As more and more rope winds off the drum, a point is reached where the lowering torque exerted by the hook and rope is greater than the running friction, and a restraining torque is then needed to prevent a runaway. The necessary stabilising torque is automatically provided by the motor acting as a generator as soon as the synchronous speed is exceeded, as shown in Fig. 6.16. The speed will therefore be held at just above the synchronous speed, provided of course that the peak generating torque (see Fig. 6.16) is not exceeded.

6.6.2 Self-excited induction generator

In previous sections we have stressed that the rotating magnetic field or excitation is provided by the magnetising current drawn from the supply, so it would seem obvious that the motor could not generate unless a supply was provided to furnish the magnetising current. However, it is possible to make the machine ‘self-excite’ if the conditions are right, and given the robustness of the cage motor this can make it an attractive proposition, especially for small-scale isolated installations.

We saw in Chapter 5 that when the induction motor is running at its normal speed, the rotating magnetic field that produces the currents and torque on the rotor also induced balanced 3-phase induced e.m.f’s in the stator windings, the magnitude of the e.m.f’s being not a great deal less than the voltage of the utility supply. So to act as an independent generator what we want to do is to set up the rotating magnetic field without having to connect to an active voltage source.

We discussed a similar matter in Chapter 3, in connection with self-excitation of the shunt d.c. machine. We saw that if enough residual magnetic flux remained in the field poles after the machine had been switched off, the e.m.f. produced when the shaft was rotated could begin to supply current to the field winding, thereby increasing the flux, further raising the e.m.f. and initiating a positive feedback (or bootstrap) process which was ultimately stabilised by the saturation characteristic of the iron in the magnetic circuit.

Happily, much the same can be achieved with an isolated induction motor. We aim to capitalise on the residual magnetism in the rotor iron, and by turning the rotor, generate an initial voltage in the stator to kick-start the process. The e.m.f. induced must then drive current to reinforce the residual field and promote the positive feedback to build up the travelling flux field. Unlike the d.c. machine, however, the induction motor has only one winding that provides both excitation and energy converting functions, so given that we want to get the terminal voltage to its rated level before we connect whatever electrical load we plan to supply, it is clearly necessary to provide a closed path for the would-be excitation current. This path should encourage the build-up of magnetising current—and hence terminal voltage.

‘Encouraging’ the current means providing a very low impedance path, so that a small voltage drives a large current, and since we are dealing with a.c. quantities, we naturally seek to exploit the phenomenon of resonance, by placing a set of capacitors in parallel with the (inductive) windings of the machine, as shown in Fig. 6.17.

The reactance of a parallel circuit consisting of pure inductance (L) and capacitance (C) at angular frequency ω is given by $X=\omega L-{}^{1}/_{\omega C}$, so at low and high frequencies the reactance is very large, but at the so-called resonant frequency $\left(\omega_0={}^{1}/_{\sqrt{LC}}\right)$, the reactance becomes zero. Here the inductance is the magnetising inductance of each phase of the induction machine, and C is the added capacitance, the value being chosen to give resonance at the desired frequency of generation. Of course the circuit is not ideal because there is resistance in the windings, but nevertheless the inductive reactance can be ‘tuned out’ by choice of capacitance, leaving a circulating path of very low resistance. Hence by turning the rotor at the speed at which the desired frequency is

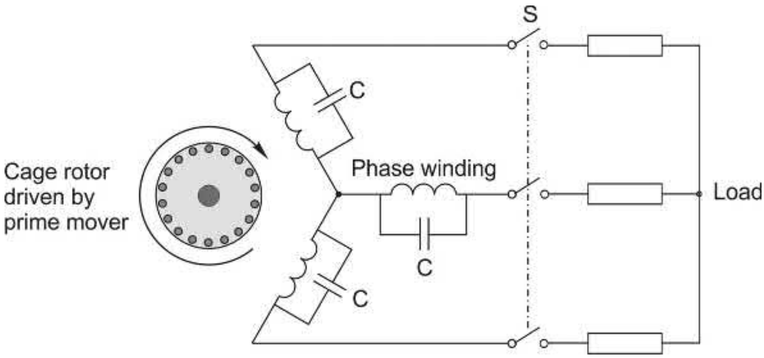

FIG. 6.17 Self-excited induction generator. The load is connected only after the stator voltage has built up.

produced by the residual magnetism (e.g. 1800 rev/min for a 4-pole motor to generate 60 Hz), the initial modest e.m.f. produces a disproportionately high current and the flux builds up until limited by the non-linear saturation characteristic of the iron magnetic circuit. We then get balanced 3-phase voltages at the terminals, and the load can be applied by closing switch S (Fig. 6.17).

The description above gives only a basic outline of the self-excitation mechanism. Such a scheme would only be satisfactory for a very limited range of driven speeds and loads, and in practice further control features are required to vary the effective capacitance (typically using triac control) in order to keep the voltage constant when the load and/or speed vary widely.

6.6.3 Doubly-fed induction machine for wind power generation

The term doubly-fed refers to an induction machine in which both stator and rotor windings are connected to an a.c. power source: we are therefore talking about wound rotor (or slipring) motors, where the rotor windings are accessed through insulated rotating sliprings.

Traditionally, large slipring machines were used in Kramer[2] drives to recover the slip energy in the rotor and return it to the supply, so that efficient operation was possible at much higher slips than would otherwise have been possible. Some Kramer drives remain (see also Section 6.8.4), but in the 21st century by far the major application for the doubly-fed induction machine is in wind-power generation where the wind turbine feeds directly into the utility grid.

Before we see why the doubly-fed motor is favoured, we should first acknowledge that in principle we could take a cage motor, connect it directly

2. See for example The Control Techniques Drives and Controls Handbook, 2nd Edition, by W Drury.

to the utility grid, and drive its rotor from the wind turbine (via a gearbox with an output speed a little above synchronous speed) so that it supplied electrical power to the grid. However, the range of speeds over which stable generation is possible would be only a few percent above synchronous, and this is a poor match as far as the turbine characteristics are concerned. Ideally, in order to extract the maximum power from the wind, the speed of rotation of the blades must vary according to the conditions, so being forced to remain at a more-or-less constant speed by being connected to a cage motor is not good news. In addition, when there are rapid fluctuations in wind speed that produce bursts of power, the fact that the speed is constant means that there are rapid changes in torque which produce unwelcome fatigue loading in the gearbox. What is really wanted is a generator in which the generated frequency can be maintained constant over a wider speed range, and this is where the doubly-fed system scores.

The stator is connected directly to the utility grid, while the rotor windings are also linked to the grid, but via a pair of ac/dc converters, as shown in Fig. 6.18. The converters - connected via a d.c. link - allow power to flow to or from the fixed-frequency grid into or out of the rotor circuit, the frequency of which varies according to the shaft speed (see below). The rating of these converters will be substantially less than the rated power output of the induction machine, depending on the speed range that is to be accepted. For example, if the speed range is to be $\pm 1.3N_S$, where N_s is the synchronous speed of the induction motor connected to the grid, and full torque is to be available over the full speed range, the rating of the converters will be only 30% of the rating of the machine. (This is a major advantage over the alternative of having a conventional synchronous generator and a frequency converter of 100% rating.)

Understanding all the details of how the doubly-fed induction generator operates is not easy, but we can get to the essence by picturing the relationship between the rotating magnetic field and the stator and rotor.

Given that the stator is permanently connected to the utility grid, we know that the magnitude and speed of the rotating magnetic field cannot vary, so the synchronous speed of a 4-pole machine connected to the 60 Hz supply will be

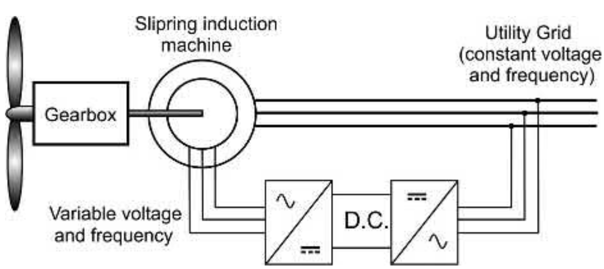

FIG. 6.18 Doubly-fed wound-rotor induction machine for wind power generation.

1800 rev/min. If we want to generate power to the grid at this speed, we feed d.c. (i.e. zero frequency) into the rotor, and we then have a synchronous machine (see chapter 9), which can motor or generate, and its power factor can be controlled via the rotor current. The 4-pole field produced by the rotor is stationary with respect to the rotor, but because the rotor is turning at 1800 rev/min, it rotates at the same speed as the travelling field produced by the stator windings, the two effectively being locked together so that torque can be transmitted and power can be converted. Energy conversion is only possible at the speed of exactly 1800 rev/min.

Now suppose that the wind turbine speed falls so that the speed at the machine shaft is only 1500 rev/min. For the rotor to be able to lock-on to the 1800 rev/min stator field, the field that it produces must rotate at 300 rev/min (in a positive sense) relative to the rotor, so that its speed relative to the stator is $1500+300=1800$ rev/min. This is achieved by supplying the rotor with 3-phase current at a frequency of 10 Hz. Conversely, if the turbine drives the machine shaft speed at 2100 rev/min, the rotor field must rotate at 300 rev/min in a negative sense relative to the rotor, i.e. the rotor currents must again be at 10 Hz, but with reversed phase sequence.

It turns out that if the input speed is below the 1800 rev/min synchronous speed (e.g. 1500 rev/min in the example above), electrical power has to be fed into the rotor circuit. This power is taken from the utility grid, but (neglecting losses, which are small) it then emerges from the stator, together with the mechanical power supplied by the turbine. Thus we can think of the electrical input power to the rotor as merely 'borrowed' from the grid to allow the energy conversion to take place: of course, the net power supplied to the grid all comes from the turbine.

In the example above, if the turbine torque is at rated value, the overall power output will be $\frac{1500}{1800}P$, i.e. $5P/6$, where P is the rated power at normal (synchronous) speed. This will be made up of an output power of P from the stator, from which the rotor converters take $\frac{300}{1800}P$, i.e. P/6 into the rotor.

When the driven speed is above synchronous (e.g.2100 rev/min), both stator and rotor circuits export power to the grid. In this case, with rated turbine torque, the power to the grid will be $\frac{2100}{1800}P$, i.e. 7P/6, comprising P from the stator and P/6, from the rotor converter.

We should be clear that there is no magic in being able to exceed the original rated power of our machine. At full torque, the electric and magnetic loadings will be at their rated values, and the increase in power is therefore entirely due to the higher speed. We discussed this towards the end of Chapter 1.

In Chapter 9, we discuss how, in a conventional (single-speed) synchronous machine, we can control the extent to which the stator and rotor sides contribute to the setting up of the resultant flux, and thereby control the grid power factor via the rotor circuit. The ability of the doubly-fed induction motor to do the same can be very advantageous where there is a system requirement to export or absorb reactive volt-amperes

6.7 Braking

6.7.1 Plug reversal and plug braking

Because the rotor always tries to catch up with the rotating field, it can be reversed rapidly simply by interchanging any two of the supply leads. The changeover is usually obtained by having two separate three-pole contactors, one for forward and one for reverse. This procedure is known as plug reversal or plugging, and is illustrated in Fig. 6.19.

The motor is initially assumed to be running light (and therefore with a very small positive slip) as indicated by point A on the dotted torque-speed curve in Fig. 6.19A. Two of the supply leads are then reversed, thereby reversing the direction of the field, and bringing the mirror-image torque-speed curve shown by the solid line into play. The slip of the motor immediately after reversal is approximately 2, as shown by point B on the solid curve. The torque is thus negative, and the motor decelerates, the speed passing through zero at point C and then rising in the reverse direction before settling at point D, just below the synchronous speed.

The speed-time curve is shown in Fig. 6.19B. We can see that the deceleration (i.e. the gradient of the speed-time graph) reaches a maximum as the motor passes through the peak torque (pull-out) point, but thereafter the final speed is approached gradually, as the torque tapers down to point D.

Very rapid reversal is possible using plugging; for example a 1 kW motor will typically reverse from full speed in under one second. But large cage motors can only be plugged if the supply can withstand the very high currents involved, which are even larger than when starting from rest. Frequent plugging will also cause serious overheating, because each reversal involves the 'dumping' of four times the stored kinetic energy as heat in the windings.

Plugging can also be used to stop the rotor quickly, but obviously it is then necessary to disconnect the supply when the rotor comes to rest, otherwise it

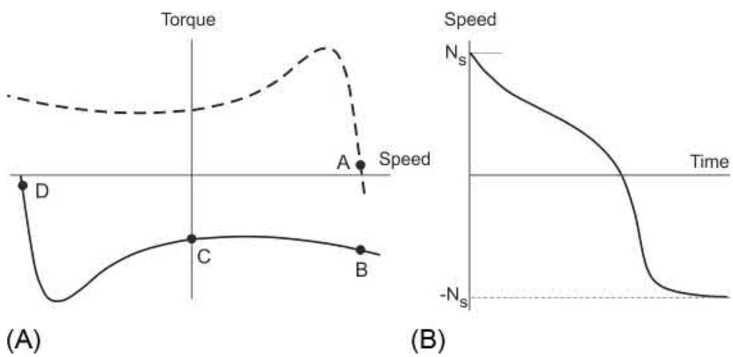

FIG. 6.19 Torque-speed and speed-time curves for plug reversal of cage motor.

will run up to speed in reverse. A shaft-mounted reverse-rotation detector is therefore used to trip out the reverse contactor when the speed reaches zero.

We should note that whereas in the regenerative mode (discussed in the previous section) the slip was negative, allowing mechanical energy from the load to be converted to electrical energy and fed back to the utility supply, plugging is a wholly dissipative process in which all the kinetic energy ends up as heat in the motor.

6.7.2 Injection braking

This is the most widely-used method of electrical braking. When the 'stop' signal occurs the three-phase supply is interrupted, and a d.c. current is fed into the stator via two of its terminals. The d.c. supply is usually obtained from a rectifier fed via a low-voltage high-current transformer.

We saw earlier that the speed of rotation of the air-gap field is directly proportional to the supply frequency, so it should be clear that since d.c. is effectively zero frequency, the air-gap field will be stationary. We also saw that the rotor always tries to run at the same speed as the field. So if the field is stationary, and the rotor is not, a braking torque will be exerted. A typical torque-speed curve for braking a cage motor is shown in Fig. 6.20, from which we see that the braking (negative) torque falls to zero as the rotor comes to rest.

This is in line with what we would expect, since there will only be induced currents in the rotor (and hence torque) when the rotor is 'cutting' the flux. As with plugging, injection (or dynamic) braking is a dissipative process, all the kinetic energy being turned into heat inside the motor.

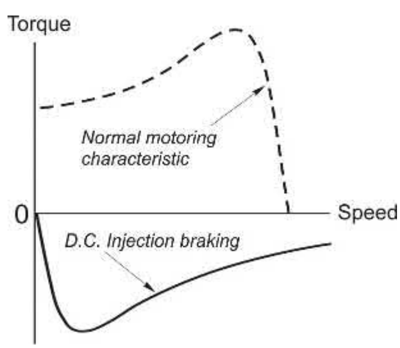

FIG. 6.20 Torque-speed curve for d.c. injection braking of cage motor.

6.8 Speed control (without varying the stator supply frequency)

We have seen that to operate efficiently an induction motor must run with a small slip. It follows that any efficient method of speed control must be based on varying the synchronous speed of the field, rather than the slip. The two factors which determine the speed of the field are the supply frequency and the pole-number (See Eq. 5.1).

The pole-number has to be an even integer, so where continuously adjustable speed control over a wide range is called for, the best approach by far is to provide a variable-frequency supply. This method is very important, and is dealt with separately in Chapters 7 and 8. In this Chapter we are concerned with constant frequency (utility-connected) operation, so we are limited to either pole-changing, which can provide discrete speeds only, or slip-control which can provide continuous speed control, but is inherently inefficient.

6.8.1 Pole-changing motors

For some applications continuous speed control may be an unnecessary luxury, and it may be sufficient to be able to run at two discrete speeds. Among many instances where this can be acceptable and economic are pumps, lifts and hoists, fans and some machine tool drives.

We established in Chapter 5 that the pole-number of the field was determined by the layout and interconnection of the stator coils, and that once the winding has been designed, and the frequency specified, the synchronous speed of the field is fixed. If we wanted to make a motor which could run at either of two different speeds, we could construct it with two separate stator windings (say 4-pole and 6-pole), and energise the appropriate one. There is no need to change the cage rotor since the pattern of induced currents can readily adapt to suit the stator pole-number. Early two-speed motors did have two distinct stator windings, but were bulky and inefficient.

It was soon realised that if half of the phase-belts within each phase-winding could be reversed in polarity, the effective pole-number could be halved. For example a 4-pole m.m.f. pattern (N-S-N-S) would become (N-N-S-S), i.e. effectively a 2-pole pattern with one large N and one large S pole. By bringing out six leads instead of three, and providing switching contactors to effect the reversal, two discrete speeds in the ratio 2:1 are therefore possible from a single winding. The performance at the high (e.g. 2-pole) speed is relatively poor, which is not surprising in view of the fact that the winding was originally optimised for 4-pole operation.

It was not until the advent of the more sophisticated Pole Amplitude Modulation (PAM) method in the 1960s that two-speed single-winding high-performance motors with more or less any ratio of speeds became available from manufacturers. This subtle technique allows close ratios such as 4/6, 6/8, 8/10 or wide ratios such as 2/24 to be achieved. The beauty of the PAM method is that it is not expensive. The stator winding has more leads brought out, and the coils are connected to form non-uniform phase-belts, but otherwise construction is the same as for a single-speed motor. Typically six leads will be needed, three of which are supplied for one speed, and three for the other, the switching being done by contactors. The method of connection (star or delta) and the number of parallel paths within the winding are arranged so that the air-gap flux at each speed matches the load requirement. For example if constant torque is needed at both speeds, the flux needs to be made the same, whereas if reduced torque is acceptable at the higher speed the flux can obviously be lower.

6.8.2 Voltage control of high-resistance cage motors

Where efficiency is not of paramount importance, the torque (and hence the running speed) of a cage motor can be controlled simply by altering the supply voltage. The torque at any slip is approximately proportional to the square of the voltage, so we can reduce the speed of the load by reducing the voltage. The method is not suitable for standard low-resistance cage motors, because their stable operating speed range is very restricted, as shown in Fig. 6.21A. But if special high-rotor-resistance motors are used, the slope of the torque-speed curve in the stable region is much less, and a somewhat wider range of steady-state operating speeds is available, as shown in Fig. 6.21B.

The most unattractive feature of this method is the low efficiency which is inherent in any form of slip-control. We recall that the rotor efficiency at slip s is (1 - s), so if we run at say 70% of synchronous speed (i.e. $s = 0.3$), 30% of the power crossing the air-gap is wasted as heat in the rotor conductors. The approach is therefore only practicable where the load torque is low at low speeds, so a fan-type characteristic is suitable, as shown in Fig. 6.21B. Voltage control became feasible only when relatively cheap thyristor a.c. voltage regulators arrived on the scene during the 1970's, and although it enjoyed some success it is now seldom seen. The hardware required is essentially the same as discussed earlier for soft starting, and a single piece of kit can therefore serve for both starting and speed control.

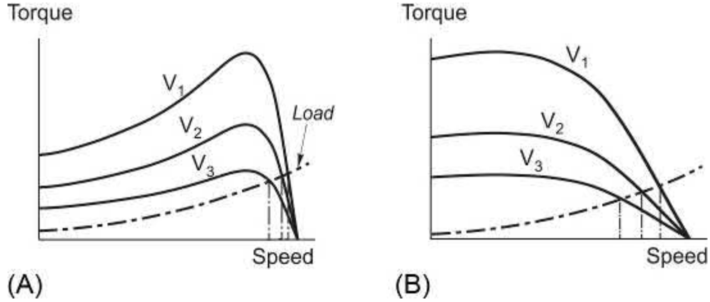

FIG. 6.21 Speed control of cage motor by stator voltage variation; (A) low resistance rotor, (B) high resistance rotor.

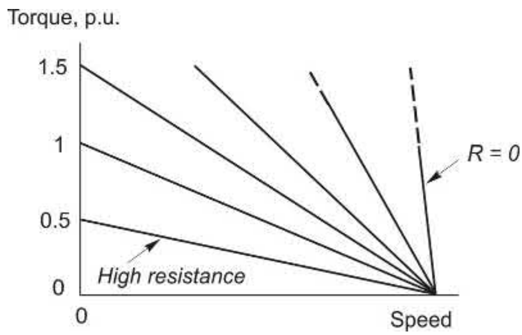

FIG. 6.22 Influence of external rotor resistance (R) on torque-speed curve of wound-rotor motor.

6.8.3 Speed control of wound-rotor motors

The fact that the rotor resistance can be varied easily allows us to control the slip from the rotor side, with the stator supply voltage and frequency constant. Although the method is inherently inefficient it is still sometimes used because of its simplicity and comparatively low cost.

A set of torque-speed characteristics is shown in Fig. 6.22, from which it should be clear that by appropriate selection of the rotor circuit resistance, any torque up to typically 1.5 times full-load torque can be achieved at any speed.

6.8.4 Slip energy recovery

Instead of wasting rotor-circuit power in an external resistance, it can be converted and returned to the utility supply. Frequency conversion is necessary because the rotor circuit operates at slip frequency, so it cannot be connected directly to the supply. These systems are known as static Kramer drives (see Section 6.6.3), and have largely been superseded by inverter-fed cage motors, but a brief mention is in order.

In a slip energy recovery system, the slip-frequency a.c. from the rotor is first rectified in a three-phase diode bridge and smoothed before being returned to the utility supply via a three-phase thyristor bridge converter operating in the inverting mode (see Chapter 4). A transformer is usually required to match the output from the controlled bridge to the supply voltage. Since the cost of both converters depends on the slip power they have to handle, this system was most often used where only a modest range of speeds (say from 70% of synchronous and above) is required, such as in large pump and compressor drives.

6.9 Power-factor control and energy optimisation

In addition to their use for soft-start and speed control, thyristor voltage regulators provide a means for limited control of power-factor for cage motors, and this can allow a measure of energy economy. However, the fact is that there are comparatively few situations where considerations of power factor and/or energy economy alone are sufficient to justify the expense of a voltage controller. Only when the motor operates for very long periods running light or at low load can sufficient savings be made to cover the outlay. There is certainly no point in providing energy-economy when the motor spends most of its time working at or near full-load.

Both power-factor control and energy optimisation rely on the fact that the air-gap flux is proportional to the supply voltage, so that by varying the voltage, the flux can be set at the best level to cope with the prevailing load. We can see straightaway that nothing can be achieved at full load, since the motor needs full flux (and hence full voltage) to operate as intended. Some modest savings in losses can be achieved at reduced load, as we will see.

If we imagine the motor to be running with a low load torque and full voltage, the flux will be at its full value, and the magnetising component of the stator current will be larger than the work component, so the input power-factor ($\cos\phi_a$) will be very low, as shown in Fig. 6.23A.

Now suppose that the voltage is reduced to say half (by phasing back the thyristors), thereby halving the air-gap flux and reducing the magnetising current by at least a factor of two. With only half the flux, the rotor current must double to produce the same torque, so the work current reflected in the stator will also double. The input power-factor ($\cos\phi_b$) will therefore improve considerably (See Fig. 6.23B). Of course the slip with 'half-flux' operation will be higher (by a factor of four), but with a low resistance cage it will still be small, and the drop in speed will therefore be slight.

The success (or otherwise) of the energy economy obtained depends on the balance between the iron losses and the copper losses in the motor. Reducing the voltage reduces the flux, and hence reduces the eddy current and hysteresis losses in the iron core. But as we have seen above, the rotor current has to increase to produce the same torque, so the rotor copper loss rises. The stator copper loss will reduce if (as in Fig. 6.23) the magnitude of the stator current

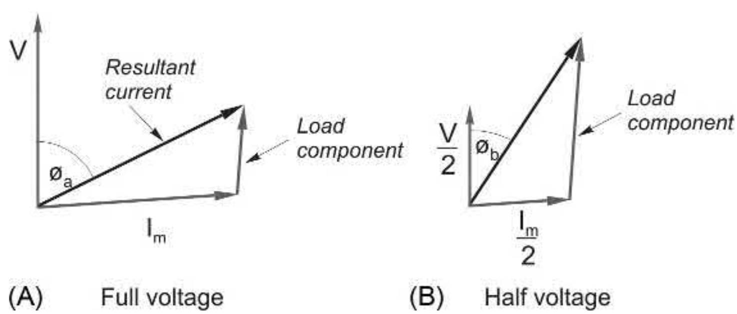

FIG. 6.23 Phasor diagram showing improvement of power-factor by reduction of stator voltage.

falls. In practice, with average general purpose motors, a net saving in losses only occurs for light loads, say at or below 25% of full load, though the power-factor will always increase.

6.10 Single-phase induction motors

Single-phase induction motors are simple, robust and reliable, and continue to be used in large numbers especially in domestic and commercial applications where three-phase supplies are not available. Although outputs of up to a few kW are possible, the majority are below 0.5 kW, and are used in straightforward applications such as refrigeration compressors, dryers, pumps and fans, small machine tools, etc. However, these traditional applications can now benefit from the superior control and low cost provided by the simple inverter and three-phase induction or permanent magnet motor, so the single-phase motor looks set for a niche role in future.

6.10.1 Principle of operation

If one of the leads of a 3-phase motor is disconnected while it is running light, it will continue to run with a barely perceptible drop in speed, and a somewhat louder hum. With only two leads remaining there can only be one current, so the motor must be operating as a single-phase machine. If load is applied the slip increases more quickly than under three-phase operation, and the stall torque is much less, perhaps one-third. When the motor stalls and comes to rest it will not restart if the load is removed, but remains at rest drawing a heavy current and emitting an angry hum. It will burn-out if not disconnected rapidly.

It is not surprising that a truly single-phase cage induction motor will not start from rest, because as we saw in Chapter 5 the single winding, fed with a.c., simply produces a pulsating flux in the air-gap, without any suggestion of rotation. It is however surprising to find that if the motor is given a push in either direction it will pick up speed, slowly at first but then with more vigour, until it settles with a small slip, ready to take-up load. Once turning, a rotating field is evidently brought into play to continue propelling the rotor.

We can understand how this comes about by first picturing the pulsating m.m.f. set up by the current in the stator winding as being the resultant of two identical travelling waves of m.m.f., one in the forward direction and the other in reverse. (This equivalence is not self-evident, but can be demonstrated by applying the method of symmetrical components, discussed earlier in this chapter.)

When the rotor is stationary, it reacts equally to both travelling waves, and no torque is developed. When the rotor is turning, however, the induced rotor currents are such that their m.m.f. opposes the reverse stator m.m.f. to a greater extent than they oppose the forward stator m.m.f. The result is that the forward flux wave (which is what develops the forward torque) is bigger than the reverse flux wave (which exerts a drag). The difference widens as the speed increases, the forward flux wave becoming progressively bigger as the speed rises while the reverse flux wave simultaneously reduces. This ‘positive feedback’ effect explains why the speed builds slowly at first, but later zooms up to just below synchronous speed. At the normal running speed (i.e. small slip), the forward flux is many times larger than the backward flux, and the drag torque is only a small percentage of the forward torque.

As far as normal running is concerned, a single winding is therefore sufficient. But all motors must be able to self-start, so some mechanism has to be provided to produce a rotating field even when the rotor is at rest. Several methods are employed, all of them using an additional winding.

The second winding usually has less copper than the main winding, and is located in the slots which are not occupied by the main winding, so that its m.m.f. is displaced in space relative to that of the main winding. The current in the second winding is supplied from the same single-phase source as the main winding current, but is caused to have a phase-lag, by a variety of means which are discussed later. The combination of a space displacement between the two windings together with a time displacement between the currents produces a two-phase machine. If the two windings were identical, displaced by 90°, and fed with currents with 90° phase-shift, an ideal rotating field would be produced. In practice we can never achieve a 90° phase-shift between the currents, and it turns out to be more economic not to make the windings identical. Nevertheless, a decent rotating field is set up, and entirely satisfactory starting torque can be obtained. Reversal is simply a matter of reversing the polarity of one of the windings, and performance is identical in both directions.

The most widely used methods are described below. At one time it was common practice for the second or auxiliary winding to be energised only during start and run-up, and for it to be disconnected by means of a centrifugal switch mounted on the rotor, or sometimes by a time-switch. This practice gave rise to the term ‘starting winding’. Nowadays it is more common to find both windings in use all the time.

6.10.2 Capacitor run motors

A capacitor is used in series with the auxiliary winding (Fig. 6.24) to provide a phase-shift between the main and auxiliary winding currents. The capacitor (usually of a few μF, and with a voltage rating which may well be higher than the supply voltage) may be mounted piggyback fashion on the motor, or located elsewhere: its value represents a compromise between the conflicting requirements of high starting torque and good running performance.

A typical torque-speed curve is also shown in Fig. 6.24; the modest starting torque indicates that capacitor run motors are generally best suited to fan type-loads. Where higher starting torque is needed, two capacitors can be used, one being switched out when the motor is up to speed.

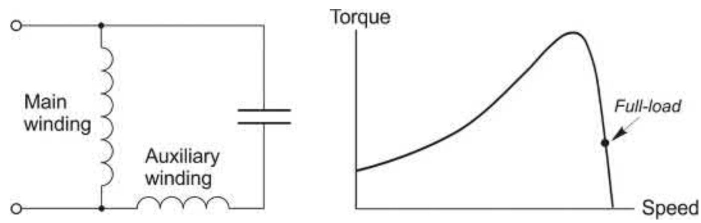

FIG. 6.24 Single-phase capacitor-run induction motor.

As mentioned above, the practice of switching out the starting winding altogether is no longer favoured, but many old ones remain, and where a capacitor is used they are known as 'capacitor start' motors.

6.10.3 Split-phase motors

The main winding is of thick wire, with a low resistance and high reactance, while the auxiliary winding is made of fewer turns of thinner wire with a higher resistance and lower reactance (Fig. 6.25). The inherent difference in impedance is sufficient to give the required phase-shift between the two currents without needing any external elements in series. Starting torque is good at typically 1.5 times full-load torque, as also shown in Fig. 6.25. As with the capacitor type, reversal is accomplished by changing the connections to one of the windings.

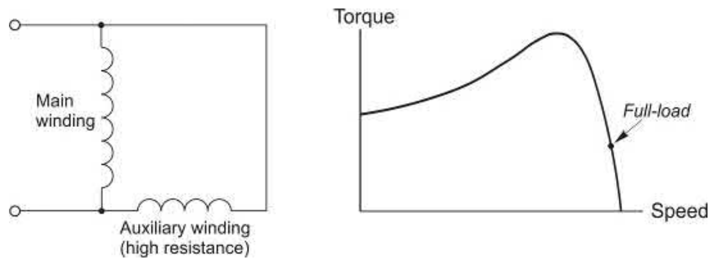

FIG. 6.25 Single-phase split-phase induction motor.

6.10.4 Shaded pole motors

There are several variants of this extremely simple, robust and reliable cage motor, which is used for low-power applications such as hair-dryers, oven fans, office equipment, display drives, etc. A 2-pole version from the low cost end of the market is shown in Fig. 6.26.

The rotor, typically between 1 and 4 cm diameter, has a die-cast aluminium cage, while the stator winding is a simple concentrated coil wound round the laminated core. The stator pole is slotted to receive the 'shading ring' which is a single short-circuited turn of thick copper or aluminium.

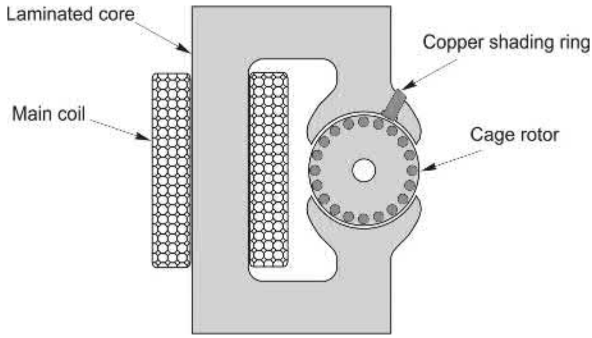

FIG. 6.26 Shaded-pole induction motor.

Most of the pulsating flux produced by the stator winding by-passes the shading ring and crosses the air-gap to the rotor. But some of the flux passes through the shading ring, and because it is alternating it induces an e.m.f. and current in the ring. The opposing m.m.f. of the ring current diminishes and retards the phase of the flux through the ring, so that the flux through the ring reaches a peak after the main flux, thereby giving what amounts to a rotation of the flux across the face of the pole. This far from perfect travelling wave of flux produces the motor torque by interaction with the rotor cage. Efficiencies are low because of the rather poor magnetic circuit and the losses caused by the induced currents in the shading ring, but this is generally acceptable when the aim is to minimise first cost. Series resistance can be used to obtain a crude speed control, but this is only suitable for fan-type loads. The direction of rotation depends on whether the shading ring is located on the right or left side of the pole, so shaded pole motors are only suitable for unidirectional loads.

6.11 Power range

Having praised the simplicity and elegance of the induction motor, it is perhaps not surprising that they are successful over an extraordinarily wide power range from multi MW down to a few tens of W. (Indeed we cannot think of any other non-electric energy converter that can exceed a span of six orders of magnitude!)

At the upper end the limit is largely a question of low demand, there being few applications that need a shaft that delivers many tens of megawatts. But at the lower end we may wonder why there are no very small ones. Industrial (3-phase) induction motors are rarely found below about 200 W, and single-phase versions, rarely extend below about 50 W, yet applications in this range are legion.

We will see that when we scale down a successful design, the excitation or flux-producing function of the windings becomes more and more demanding until eventually the heat produced in the windings by the excitation current causes the permissible temperature to be reached. There is then no spare capacity for the vital function of supplying the mechanical output power, so the machine is of no use.

6.11.1 Scaling down—The excitation problem

We can get to the essence of the matter by imagining that we take a successful design and scale all the linear dimensions by half. We know that in order to fully utilise the iron of the magnetic circuit we would want the air-gap flux density to be the same as in the original design, so because the air-gap length has been halved the stator m.m.f. needs to be half of what it was. The number of coils and the turns in each coil remain as before, so if the original magnetising current was I_m, the magnetising current of the half-scale motor will be $I_m/2$.

Turning now to what happens to the resistance of the winding, we will assume that the resistance of the original winding was R. In the half-scale motor, the total length of wire is half of what it was, but the cross-sectional area of the wire is only a quarter of the original. As a result the new resistance is twice as great, i.e. $2R$.

The power dissipated in providing the air-gap flux in the original motor is given by I_m^2R, while the corresponding excitation power in the half-scale motor is given by

$$\left(\frac{I_m}{2}\right)^2 \times 2R = \frac{1}{2} I_m^2 R.$$

When we consider what determines the steady temperature rise of a body in which heat is dissipated, we find that the equilibrium condition is reached when the rate of loss of heat to the surroundings is equal to the rate of production of heat inside the body. And, not surprisingly, the rate of loss of heat to the surroundings depends on the temperature difference between the body and its surroundings, and on the surface area through which the heat escapes. In the case of the copper windings in a motor, the permissible temperature rise depends on the quality of insulation, so we will make the reasonable assumption that the same insulation is used for the scaled motor as for the original.

We have worked out that the power dissipation in the new motor is half of that in the original. However, the surface area of the new winding is only one quarter, so clearly the temperature rise will be higher, and if all other things were equal, it will double. We might aim to ease matters by providing bigger slots so that the current density in the copper could be reduced, but as explained in Chapter 1 this means that there is less iron in the teeth to carry the working flux. A further problem arises because it is simply not practicable to go on making the air-gap smaller because the need to maintain clearances between the moving parts would require unacceptably tight manufacturing tolerances.

Obviously, there are other factors that need to be considered, not least that a motor is designed to reach its working temperature when the full current (not just the magnetising current) is flowing. But the fact is that the magnetisation problem we have highlighted is the main obstacle in small sizes, not only in induction motors but also in any motor that derives its excitation from the stator windings. Permanent magnets therefore become attractive for small motors, because they provide the working flux without producing unwelcome heat. Permanent magnet motors are discussed in Chapter 9.

6.12 Review questions

(1) **(a)** Why do large induction motors sometimes cause a dip in the supply system voltage when they are switched direct on line?
(b) Why, for a given induction motor, might it be possible to employ direct-on-line starting in one application, but be necessary to employ a starter in another application?
(c) What is meant by the term ‘stiff’ in relation to an industrial electrical supply?
(d) Why would a given motor take longer to run up to speed when started from a weak supply, as compared with the time it would take to run up to speed on a stiff supply?

(2) The voltage applied to each phase of a motor when it is star-connected is $1/\sqrt{3}$ times the voltage applied when in delta connection. Using this information, explain briefly why the line current and starting torque in star are both 1/3 of their values in delta.

(3) Explain briefly why, in many cage rotors, the conductor bars are not insulated from the laminated core.

(4) How could the pole-number of an induction motor be determined by inspection of the stator windings?

(5) Choose suitable pole-numbers of cage induction motors for the following applications:
(a) a grindstone to run at about 3500 rev/min when the supply is 60 Hz;
(b) a pump drive to run at approximately 700 rev/min from a 50 Hz supply;
(c) a turbo-compressor to run at 8000 rev/min from 60 Hz;

(6) The full-load speed of a 4-pole, 60 Hz induction motor is 1700 rev/min. Why is it unlikely that the full-load efficiency could be as high as 94%?

(7) Sketch a typical cage motor torque-speed curve and indicate:-
(a) the synchronous speed;
(b) the starting torque;
(c) the stable operating region;
(d) the stall speed.

(8) The full-load speed of a 4-pole, 60 Hz, low-resistance cage induction motor is 1740 rev/min.

Estimate the speed under the following conditions:-

(a) Half rated torque, full voltage

(b) Full torque, 85% voltage.

Why would prolonged operation in condition (b) be unwise?

(9) Why might the rotor of an induction motor become very hot if it was switched on and off repeatedly, even though it was not connected to any mechanical load?

(10) The book explains that the space harmonics of the air-gap field in an induction motor rotate at a speed that is inversely proportional to their order. For example the fifth harmonic rotates forward at one-fifth of synchronous speed, while the seventh rotates backward at one seventh of the synchronous speed. Calculate the frequencies of the e.m.f's induced by these two harmonic fields in the stator winding.

Answers to the review questions are given in the Appendix.

Chapter 7

Variable frequency operation of induction motors

7.1 Introduction

We saw in Chapter 6 the many attributes of the induction motor which have made it the preferred workhorse of industry. These include simple low cost construction, which lends itself to totally enclosed designs suitable for dirty or even hazardous environments; limited routine maintenance with no brushes; only three power connections; and good full load efficiency. We have also seen that when operated from the utility supply there are a number of undesirable characteristics, the most notable being that there is only one speed of operation (or more precisely a narrow load-dependent speed range). In addition, starting equipment is often required to avoid excessive currents of up to six times rated current when starting direct on line; reversal requires two of the power cables to be interchanged; and the instantaneous torque cannot be controlled, so the transient performance is poor.

We will see in this chapter that all the good features of the utility-operated induction motor are retained and all the bad characteristics detailed above can be avoided when the induction motor is supplied from a variable-frequency source, i.e. its supply comes from an inverter.

The first part of this chapter deals with the steady-state behaviour when the operating frequency is solely determined by the inverter, and is independent of what is happening at the motor. We will refer to this set-up as 'inverter-fed', and in the early days of converter-driven induction motors this was the norm—the frequency being set by an oscillator that controlled the sequential periodic switching of the devices in the inverter. We will see that by appropriate control of the frequency and voltage we are able to operate over a very wide range of the torque-speed plane, but we will also identify the limits on what can be achieved.

On a steady-state basis, the inverter-fed arrangement proved able to compete with the d.c. drive, but even when incorporated into a closed-loop speed control scheme the transient performance remained inferior.

The reason for the superior inherent transient performance of the d.c. drive is that the torque is directly proportional to the main armature (rotor) current,

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00007-6

which can be measured easily, and directly controlled with a high-gain current control loop. In complete contrast, the rotor current in the cage induction motor obviously cannot be measured directly, and any changes have to be induced from the stator side.

We will see in Chapter 8 that the field-oriented technique allows rapid and precise control of the rotor current (and hence torque), which it achieves by sophisticated control of the inverter switching according to what is happening in the motor. This technique (which was only made possible when cheap real-time digital processing became available) results in outstanding dynamic performance, but understanding how it works is relatively challenging, and can be confusing even for engineers with considerable experience of drive systems.

For some readers, understanding these relatively complex control schemes will not be necessary, while for others, getting to grips with and understanding advanced control strategies is critically important. The authors therefore decided that although it will be appropriate to make reference to the remarkable capabilities of field-oriented systems in this chapter, it would be best to defer detailed consideration to Chapter 8.

At this stage it is important to stress that there is no real difference between the traditional inverter-fed drive and the field oriented drives under true steady-state running conditions—a fact that is often not appreciated. It is therefore well worth absorbing the main messages from Section 7.2 (variable-frequency steady-state operation) because it covers most of the fundamental aspects applicable to all induction motor drives.

For most of this chapter we will assume that the motor is supplied from an ideal balanced sinusoidal voltage source. Our justification for doing this is that although the pulse-width-modulated voltage waveform supplied by an inverter will not be sinusoidal (see Fig. 7.1), the motor performance depends principally on the fundamental (sinusoidal) component of the applied voltage. This is a somewhat surprising but extremely welcome simplification, because it allows us to make use of our knowledge of how the induction motor behaves with a sinusoidal supply to anticipate how it performs when fed from an inverter.

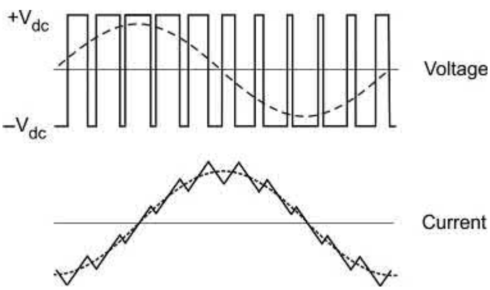

FIG. 7.1 Typical voltage and current waveforms for PWM inverter-fed induction motor. (The fundamental-frequency component is shown by the dotted line.)

In Section 7.3, we look at some of the practical aspects of inverter-fed induction motor drives and consider the impressive performance of commercially available drives. Having alluded to the excellent performance offered by advanced control strategies (see Chapter 8), for the sake of balance, we will look at an application where Field Orientation struggles and Direct Torque Control simply doesn't work.

The adoption of the standard induction motor in a variable speed drive is not without potential problems, so in Section 7.4 we look at the more important practical issues which affect the motor, and in Section 7.5 we consider the pros and cons from the point of view of the utility supply. Finally, in Section 7.6, we look briefly at inverter and motor protection.

7.2 Variable frequency operation

It was explained in Chapter 6 that the induction motor can only run efficiently at low slips, i.e. close to the synchronous speed of the rotating field. The best method of speed control must therefore provide for continuous smooth variation of the synchronous speed, which in turn calls for variation of the supply frequency. This is readily achieved using a power electronic inverter (as discussed in Chapter 2) to supply the motor. A complete speed control scheme, which is illustrated with speed feedback, is shown in simplified block diagram form in Fig. 7.2.

The arrangement shown in Fig. 7.2 shows the motor with a speed sensor/incremental encoder attached to the motor shaft. For all but the most demanding dynamic applications, or where full torque at standstill is a requirement, a speed sensor would not normally be required. This is good news because fitting a speed sensor to a standard induction motor involves extra cost and additional cabling.

We should recall that the function of the converter (i.e. rectifier and variable-frequency inverter) is to draw power from the fixed-frequency constant-voltage supply, rectify it and then convert it to variable-frequency,

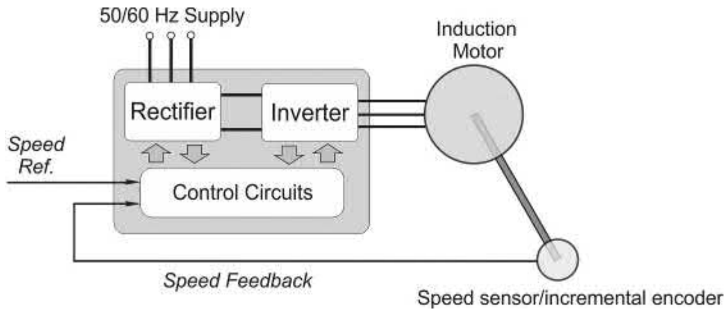

FIG. 7.2 General arrangement of inverter-fed variable-frequency induction motor controlled-speed drive.

variable-voltage for driving the induction motor. Both the rectifier and the inverter employ switching strategies (see Chapter 2), so the power conversions are accomplished efficiently and the converter can be compact.

7.2.1 Steady-state operation—Importance of achieving full flux

Three simple relationships need to be borne in mind in order to simplify understanding of how the inverter-fed induction motor behaves. Firstly, we established in Chapter 5 that for a given induction motor, the torque developed depends on the magnitude of the rotating flux density wave, and on the slip speed of the rotor, i.e. on the relative velocity of the rotor with respect to the flux wave. Secondly, the strength or amplitude of the flux wave depends directly on the supply voltage to the stator windings, and inversely on the supply frequency. Thirdly, the absolute speed of the flux wave depends directly on the supply frequency.

Recalling that the motor can only operate efficiently when the slip is small, we see that the basic method of speed control rests on the control of the speed of rotation of the flux wave (i.e. the synchronous speed), by control of the supply frequency. If the motor is a 4-pole one, for example, the synchronous speed will be 1500 rev/min when supplied at 50 Hz, 1800 rev/min at 60 Hz, 750 rev/min at 25 Hz, and so on. The no-load speed will therefore be almost exactly proportional to the supply frequency, because the torque at no load is small and the corresponding slip is also very small.

Turning now to what happens on load, we know that when a load is applied the rotor slows down, the slip increases, more current is induced in the rotor, and more torque is produced. When the speed has reduced to the point where the motor torque equals the load torque, the speed becomes steady. We normally want the drop in speed with load to be as small as possible, not only to minimise the drop in speed, but also to maximise efficiency: in short, we want to minimise the slip for a given load.

We saw in Chapter 5 that the slip for a given torque depends on the amplitude of the rotating flux wave: the higher the flux, the smaller the slip needed for a given torque. It follows that having set the desired speed of rotation of the flux wave by controlling the output frequency of the inverter we must also ensure that the magnitude of the flux is adjusted so that it is at its full (rated) value,[1] regardless of the speed of rotation. This is achieved, in principle, by making the output voltage from the inverter vary in the appropriate way in relation to the frequency.

1. In general, operating at rated flux gives the best performance and on most loads the highest efficiency. Some commercial drives offer a mode of control in which the flux is reduced typically with the square of the speed: this can provide some benefit at low speeds for fan and pump type loads where the magnetising current accounts for a significant proportion of the motor losses.

Given that the amplitude of the flux wave is proportional to the supply voltage and inversely proportional to the frequency, it follows that if we arrange that the voltage supplied by the inverter varies in direct proportion to the frequency, the flux wave will have a constant amplitude. This simple mode of operation—where the V/f ratio is constant—was for many years the basis of the control strategy applied to most inverter fed induction motors, and it can still be found in some commercial products.

Many inverters are designed for direct connection to the utility supply, without a transformer, and as a result the maximum inverter output voltage is limited to a value similar to that of the supply system. Since the inverter will normally be used to supply a standard induction motor designed, for example for 400 V, 50 Hz operation, it is obvious that when the inverter is set to deliver 50 Hz, the voltage should be 400 V, which is within the inverter's voltage range. But when the frequency was raised to say 100 Hz, the voltage should—ideally—be increased to 800 V in order to obtain full flux. The inverter cannot supply voltages above 400 V, and it follows that in this case full flux can only be maintained up to base speed. Established practice is for the inverter to be capable of maintaining the "V/f ratio", or rather the flux, constant up to the base speed (frequently 50 Hz or 60 Hz), but to accept that at higher frequencies the voltage will be constant at its maximum value. This means that the flux is maintained constant at speeds up to base speed, but beyond that the flux reduces inversely with frequency. Needless to say the performance above base speed is adversely affected, as we will see.

Users are sometimes alarmed to discover that both voltage and frequency change when a new speed is demanded. Particular concern is expressed when the voltage is seen to reduce when a lower speed is called for. Surely, it is argued, it can't be right to operate say a 400 V induction motor at anything less than 400 V. The fallacy in this view should now be apparent: the figure of 400 V is simply the correct voltage for the motor when run directly from the utility supply, at say 50 Hz. If this full voltage was applied when the frequency was reduced to say 25 Hz, the implication would be that the flux would rise to twice its rated value. This would greatly overload the magnetic circuit of the machine, giving rise to excessive saturation of the iron, an enormous magnetising current, and wholly unacceptable iron and copper losses. To prevent this from happening, and keep the flux at its rated value, it is essential to reduce the voltage in proportion to frequency. In the case above, for example, the correct voltage at 25 Hz would be 200 V.

It is worth stressing here that when considering a motor to be fed from an inverter there is no longer any special significance about the utility network frequency, and the motor can be wound for almost any base frequency. For example a motor wound for 400 V, 100 Hz could, in the above example, operate with constant flux right up to 100 Hz.

7.2.2 Torque-speed characteristics

When the voltage at each frequency is adjusted so that the ratio of voltage to frequency (V/f) is kept constant up to base speed, a family of torque speed curves as shown in Fig. 7.3 is obtained. These curves are typical for a standard induction motor of several kW output.

As expected, the no-load speeds are directly proportional to the frequency, and if the frequency is held constant e.g. at 25 Hz in Fig. 7.3, the speed drops only modestly from no-load (point a) to full-load (point b). These are therefore good, useful open-loop characteristics, because the speed is held fairly well from no-load to full-load. If the application calls for the speed to be held precisely, this can clearly be achieved by raising the frequency so that the full-load operating point moves to point (c).

We note also that the pull-out torque and the torque stiffness (i.e. the slope of the torque-speed curve in the normal operating region) is more or less the same at all points below base speed, except at low frequencies where the voltage drop due to the stator resistance becomes very significant as the applied voltage is reduced. A simple V/f control system would therefore suffer from significantly reduced flux and hence less torque at low speeds, as indicated in Fig. 7.3.

The low-frequency performance can be improved by increasing the V/f ratio at low frequencies in order to restore full flux, a technique which is referred to as ‘voltage boost’. In modern drive control schemes which calculate flux from a motor model (see Chapter 8), the voltage is automatically boosted from the linear V/f characteristic that the approximate theory leads us to expect. A typical set of torque-speed curves for a drive with the improved low-speed torque characteristics obtained with voltage boost is shown in Fig. 7.4.

The characteristics in Fig. 7.4 have an obvious appeal because they indicate that the motor is capable of producing practically the same maximum torque at all speeds from zero up to the base (50 Hz) speed. This region of the characteristics is known as the ‘constant torque’ region, which means that for frequencies up to base speed, the maximum possible torque which the motor can deliver is independent of the set speed. Continuous operation at peak torque will not be allowable because the motor will overheat, so an upper limit will be imposed by the controller, as discussed shortly. With this imposed limit, operation below

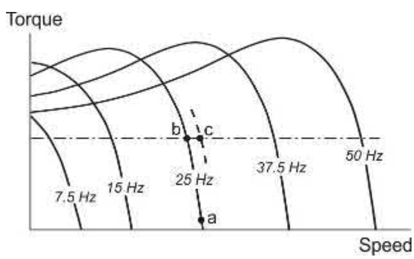

FIG. 7.3 Torque-speed curves for inverter-fed induction motor with constant V/f ratio.

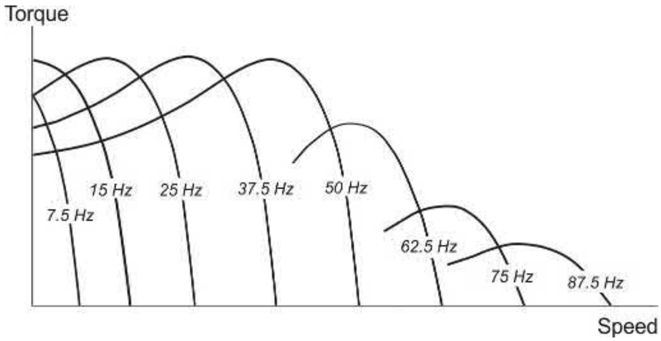

FIG. 7.4 Typical torque-speed curves for inverter-fed induction motor with constant flux up to base speed (50 Hz) and constant voltage at higher frequencies.

base speed corresponds to the armature-voltage control region of a d.c. drive, as exemplified in Fig. 3.9.

We should note that the availability of high torque at low speeds (especially at zero speed) means that we can avoid all the 'starting' problems associated with fixed-frequency operation (see Chapter 6). By starting off with a low frequency which is then gradually raised, the slip speed of the rotor is always small, i.e. the rotor operates in the optimum condition for torque production all the time, thereby avoiding all the disadvantages of high-slip (low torque and high current) that are associated with utility-frequency/Direct-on-Line (DOL) starting. This means that not only can the inverter-fed motor provide rated torque at low speeds, but—perhaps more importantly—it does so without drawing any more current from the utility supply than under full-load conditions, which means that we can safely operate from a weak supply without causing excessive voltage dips. For some essentially fixed-speed applications, the superior starting ability of the inverter-fed system alone may justify its cost.

Beyond the base frequency, the flux ("V/f ratio") reduces because V remains constant. The amplitude of the flux wave therefore reduces inversely with the frequency. Under constant flux operation, the pull-out torque always occurs at the same absolute value of slip, but in the constant-voltage region the peak torque reduces inversely with the square of the frequency and the torque-speed curve becomes less steep, as shown in Fig. 7.4.

Although the curves in Fig. 7.4 show what torque the motor can produce for each frequency and speed, they give no indication of whether continuous operation is possible at each point, yet this matter is extremely important from the users viewpoint, and is discussed next.

7.2.3 Limitations imposed by the inverter—Constant torque and constant power regions

A primary concern in the inverter is to limit the currents to a safe value as far as the main switching devices and the motor are concerned. The current limit will be typically set to the rated current of the motor, and the inverter control circuits will be arranged so that no matter what the user does the output current cannot exceed this safe (thermal) value, other than for clearly defined overload (e.g. 120% for 60 s) for which the motor and inverter will have been specified and rated. (For some applications involving a large number of starts and stops, the motor and drive must be specially designed for the specific duty.)

In modern control schemes (see Chapter 8) it is possible to have independent control of the flux and torque producing components of the current, and in this way the current limit imposes an upper limit on the permissible torque. In the region below base speed, this will normally correspond to the rated torque of the motor, which is typically about half the pull-out torque, as indicated by the shaded region in Fig. 7.5.

Above base speed, it is not possible to increase the voltage and so the flux reduces inversely with the frequency. Since the inverter current is thermally limited (as we saw in the constant torque region), the maximum permissible torque also reduces inversely with the speed, as shown in Fig. 7.5. This region is consequently known as the 'constant power' region: in this region, the flux is reduced and so the motor has to operate with higher slips than below base speed to develop the full (rated) rotor current and correspondingly reduced torque. There is of course a close parallel with the d.c. drive here, both systems operating with reduced or weak field in the constant power region. Note that if an inverter with a higher current rating were used, the motor would still operate with high slip, meaning high rotor losses and the thermal rating of the motor would become the critical factor, which we will explore next.

At all speeds in the constant power region, the maximum torque available is limited by the inverter current limit, the motor itself having some reserve before it reaches its pull-out torque. However, with constant voltage, the pull-out torque is inversely proportional to the square of the frequency, so the upper bound on torque ultimately becomes limited by the motor itself, rather than the inverter. This is shown by the hatched area in Fig. 7.5: the transition to this 'high-speed region' typically occurs at about twice base speed.

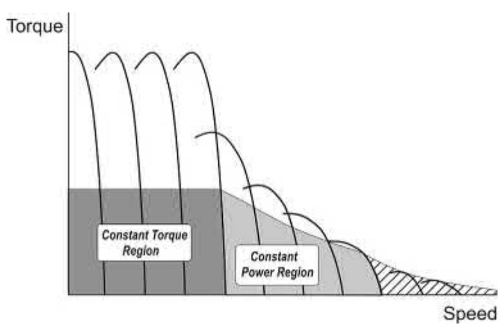

FIG. 7.5 Constant torque, constant power, and high-speed motoring regions.

7.2.4 Limitations imposed by the motor

The traditional practice in d.c. drives is to use a motor specifically designed for operation from a thyristor converter. The motor will have a laminated frame, will probably come complete with a tacho, and—most important of all—will have been designed for through ventilation and equipped with an auxiliary air blower. Adequate ventilation is guaranteed at all speeds, and continuous operation with full torque (i.e. full current) at even the lowest speed is therefore in order.

By contrast, it is still common for inverter-fed systems to use a standard industrial induction motor. These motors are usually totally enclosed, with an external shaft-mounted fan which blows air over the finned outer case (and an internal stirring fan to circulate air inside the motor to avoid spot heating). They are designed first and foremost for continuous operation from the fixed frequency utility supply, running at base speed.

As we have mentioned earlier, when such a motor is operated at a low frequency (e.g. 7.5 Hz), the speed is much lower than base speed and the efficiency of the cooling fan is greatly reduced. At the lower speed the motor will be able to produce as much torque as at base speed (see Fig. 7.4) but in doing so the losses in both stator and rotor will also be more or less the same as at base speed, so it will overheat if operated for any length of time.

However, induction motors bearing the name 'inverter grade' or similar are today readily available. As well as having reinforced insulation systems (see Section 7.4.5), they have been designed to offer a constant torque operating range below rated speed, typically down to 30% of base speed, without the need for an external cooling fan. In addition they may be offered with a separate external cooling fan to allow operation at constant (rated) torque down to standstill.

7.2.5 Four quadrant capability

So far in this chapter it is natural that we have concentrated on motoring in quadrant 1 of the torque-speed plane (see Fig. 3.12), because this is how the machine will spend most of its time, but it is important to remind ourselves that the induction motor is equally at home as a generator, a role that it will frequently perform, even with an ordinary load, when a reduction in speed is called for. We should also recall that in this part of the chapter we are studying variable-frequency operation at the fundamental level, so we should bear in mind that in practice details of the control strategy will vary from drive to drive.

We can see how intermittent generation occurs with the aid of the torque-speed curves shown in Fig. 7.6. These have been extended into quadrant 2, i.e. the negative-slip region, where the rotor speed is higher than synchronous, and a braking torque is exerted.

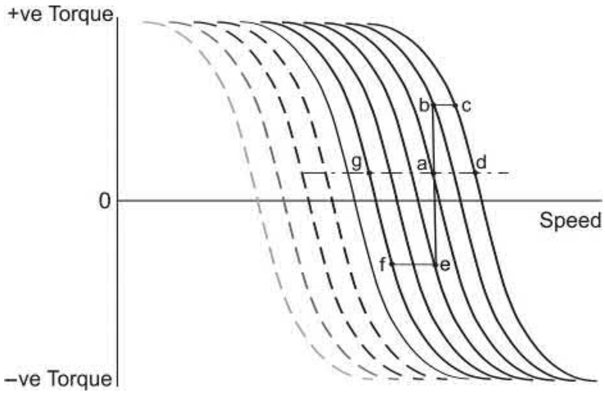

FIG. 7.6 Acceleration and deceleration trajectories in the torque-speed plane.

The family of curves indicate that for each set speed (i.e. each frequency) the speed remains reasonably constant because of the relatively steep torque-slip characteristic of the cage motor. If the load is increased beyond rated torque, an internal current limit comes into play to prevent the motor from reaching the unstable region beyond pull-out. Instead, the frequency and speed are reduced, and so the system behaves in a similar way to a d.c. drive.

Sudden changes in the speed reference are buffered so that the frequency is gradually increased or decreased. If the load inertia is low and/or the ramp time sufficiently long, the acceleration will be accomplished without the motor entering the current-limit region. On the other hand if the inertia is large and/or the ramp time was very short, the acceleration will take place as discussed below.

Suppose the motor is operating in the steady state with a constant load torque at point (a), when a new higher speed corresponding to point (d) is demanded. The frequency is increased, causing the motor torque to rise to point (b), where the current has reached the allowable limit. The rate of increase of frequency is then automatically reduced so that the motor accelerates under constant-current conditions to point (c), where the current falls below the limit: the frequency then remains constant and the trajectory follows the curve from (c) to settle finally at point (d).

A typical deceleration trajectory is shown by the path aefg in Fig. 7.6. The torque is negative for much of the time, the motor operating in quadrant 2 and regenerating kinetic energy. Because we have assumed that the motor is supplied from an ideal voltage source, this excess energy will return to the supply automatically. In practice however we should note that many drives do not have the capability to return power to the a.c. supply, and the excess energy therefore has to be dissipated in a resistor inside the converter. (The resistor is usually connected across the d.c. link, and controlled by a chopper. When the level of the d.c. link voltage rises, because of the regenerated energy, the chopper switches the resistor on to absorb the energy. High inertia loads which are subjected to frequent deceleration can therefore pose problems of excessive power dissipation in this 'dump' resistor.)

To operate as a motor in quadrant 3 all that is required is for the phase sequence of the supply to be reversed, say from ABC to ACB. Unlike the utility-fed motor, there is no need to swap two of the power leads because the phase sequencing can be changed easily at the low-power logic level in the inverter. With reverse phase sequence, a mirror image set of 'motoring' characteristics are available, as shown in Fig. 7.7. The shaded regions are as described for Fig. 7.5, and the dotted lines indicate either short-term overload operation (quadrants 1 and 3) or regeneration during deceleration (quadrants 2 and 4).

Note that unlike the d.c. motor control strategies we examined in Chapter 4, neither the motor current, nor indeed any representation of torque, play a role in the motor control strategy discussed so far (except when the current hits a limit, as discussed above).

We have seen that the inverter-fed induction motor is a very versatile variable speed drive. The control systems presented so far are rather simple, all based in principle on keeping the ratio of stator voltage to fundamental frequency constant, in an attempt to keep the air-gap flux constant and give the motor the potential for providing rated torque at the desired frequency. This form of control has many drawbacks, and in recent years a large number of commercial drive systems employ either Field-Oriented (sometimes referred to as Vector) Control or Direct Torque Control. These are discussed in Chapter 8. Before we move on to that however, there are a number of practical issues associated with inverter-fed induction motors, which are important to understand regardless of the particular type of drive in question.

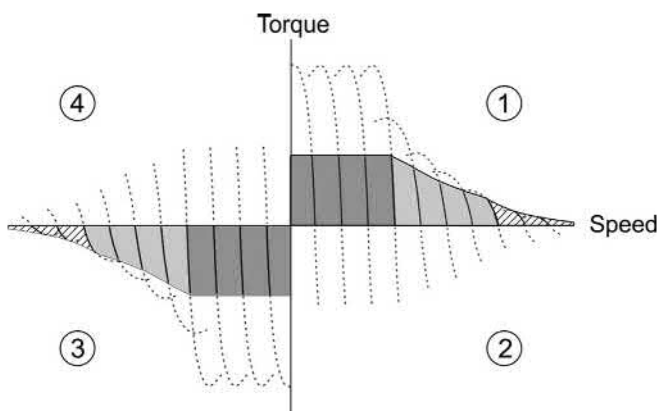

FIG. 7.7 Operating regions in all for quadrants of the torque-speed plane.

7.3 Practical aspects of inverter-fed drives

In this section we look at some of the practical aspects of inverter-fed induction motor drives and briefly consider the impressive performance of commercially available drives. By using Field-Oriented or Direct Torque Control (which we will cover in Chapter 8) it is possible to achieve not only steady-state speed control but also dynamic performance superior to that of a thyristor d.c. drive: such performance is dependent on the ability to perform very fast/real-time modelling of the motor and very rapid control of the motor voltage magnitude and phase.

Whilst it comes as no surprise that the inverter-fed induction motor is now the best-selling industrial drive, the adoption of the standard induction motor in a variable speed drive is not without potential problems, so it is important to be aware of their existence and learn something of the methods of mitigation. We will therefore consider some of the more important practical issues which result from operating standard (utility supply) motors from a variable frequency inverter.

7.3.1 PWM voltage source inverter

Several alternative drive topologies are applied to induction motors and the most relevant of these have been discussed in Chapter 2. There is also one seldom-used topology that is unique to induction motors, which we discuss briefly in Section 7.3.2. However, by far the most important for most industrial applications has a diode bridge rectifier (which only allows energy flow from the supply to the d.c. link) and a Pulse Width Modulated (PWM) Voltage Source Inverter (VSI), as shown in Fig. 7.8, and this arrangement will now be the focus of our attention.

Most low power inverters use MOSFET switching devices in the inverter bridge, and may switch at ultrasonic frequencies, which naturally results in

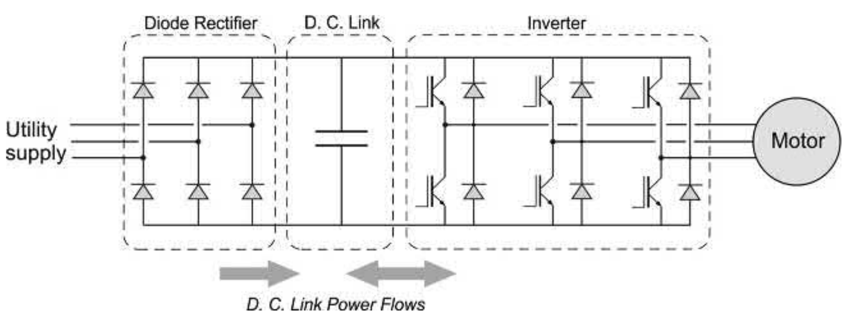

FIG. 7.8 Pulse-width modulated (PWM) voltage source inverter (VSI).

quiet operation. Medium and larger power inverters use IGBT's which can be switched at high enough frequencies to be ultrasonic for most of the population. It should be remembered however that the higher the switching frequency[2] the higher the inverter losses, and hence the lower the efficiency, and so a compromise must be reached.

Inverter-fed induction motor drives are used in ratings up to many megawatts. Standard 50 Hz or 60 Hz motors are often used, though the use of a variable frequency inverter means that motors of almost any rated/base frequency can be employed. We saw earlier in this chapter that operation above base frequency limits performance and so this needs to be carefully considered when specifying a drive system. Commercially available inverters operate with output frequencies typically from 0 Hz up to perhaps several hundred Hz, and in some cases to much higher frequencies. The low frequency limit is generally determined by the form of control, whilst the higher frequency depends on the control and the physical dimensioning of the power electronic circuits (where stray inductance can be a problem if (internal to the inverter) interconnections become too long).

The majority of inverters are three-phase input and three-phase output, but single-phase input versions are available up to about 7.5 kW. Some inverters (usually less than 3 kW) are specifically designed for use with single-phase motors, but these are unusual and will not be considered further. The upper operating frequency is generally limited by the mechanical stresses in the rotor. Very high speed motors for applications such as centrifuges and wood working machines can be designed, with special rotor construction and bearings, for speeds up to 50,000 rev/min, or occasionally even higher.

A fundamental aspect of any converter, which is often overlooked, is instantaneous energy balance. In principle, for any balanced three-phase load, the total load power remains constant from instant to instant, so if it was possible to build an ideal three-phase input, three-phase output converter, there would be no need for the converter to include any energy storage elements. In practice, all converters require some energy storage (in capacitors or inductors), but these are relatively small when the input is three-phase because the energy balance is good. However, as mentioned above, many low power (and some high power, rail traction) converters, are supplied from a single-phase source. In this case, the instantaneous input power is zero at least four times per cycle of the supply (because the voltage and current each go through zero every half-cycle). If the motor is three-phase and draws power at a constant rate from the d.c. link, it is obviously necessary to store sufficient energy in the converter to supply the motor during the brief intervals when the load power is greater than the input power. This explains why the most bulky components in many power inverters

2. Remember we are talking switching *frequency*. Faster switching *times* of compound power semiconductor devices such as SiC and GaN reduces switching losses.

are electrolytic[3] capacitors in the d.c. link. (Some drive manufacturers are now designing products, for connection to a 3 phase supply, with low values of d.c. capacitance for undemanding applications where the subsequent reduction in control/performance is acceptable).

The output waveform produced by the PWM inverter in an a.c. drive also brings with it challenges for the motor, which we will consider later. When we looked at the converter-fed d.c. motor we saw that the behaviour was governed primarily by the mean d.c. voltage, and that for most purposes we could safely ignore the ripple components. A similar approximation is useful when looking at how the inverter-fed induction motor performs: we assume that although the voltage waveform supplied by the inverter will not be sinusoidal, the motor behaviour depends principally on the fundamental (sinusoidal) component of the applied voltage. This allows us to make use of our knowledge of how the induction motor behaves with a sinusoidal supply to anticipate how it will behave when fed from an inverter.

In essence, the reason why the harmonic components of the applied voltage are much less significant than the fundamental is that the impedance of the motor at the harmonic frequencies is much higher than at the fundamental frequency. This causes the current to be much more sinusoidal than the voltage (as previously shown in principle in Fig. 7.1), and this means that we can expect a sinusoidal travelling field to be set up in much the same way as discussed in Chapter 5.

In commercial inverters the switching frequency is high and the measurement and interpretation of the actual waveforms is not straightforward. For example, the voltage and current waveforms in Fig. 7.9 relate to an industrial drive with a 3 kHz switching frequency. Note the blurring of the individual voltage pulses (a result of sampling limitations of the oscilloscope), and the near-sinusoidal fundamental component of current. We might be concerned at what appear to be spikes of current, but consideration of the motor leakage inductance and the limited forcing voltage will confirm that such rapid rates of change of current are impossible: the spikes are in fact the result of noise on the signal from the current transducer.

Measurement of almost all quantities associated with power electronic converters is difficult, and great care must be taken in the selection of instrumentation and interpretation of the results. A clear understanding of grounding is also important when reviewing inverter d.c. link and output waveforms since unlike a utility supply there is no clear, or at least simple, ground reference.

3. In some applications including aerospace and automotive, electrolytic capacitors are not allowed due to safety concerns should such capacitors dry out, and fail. In these cases physically larger film type capacitors are used.

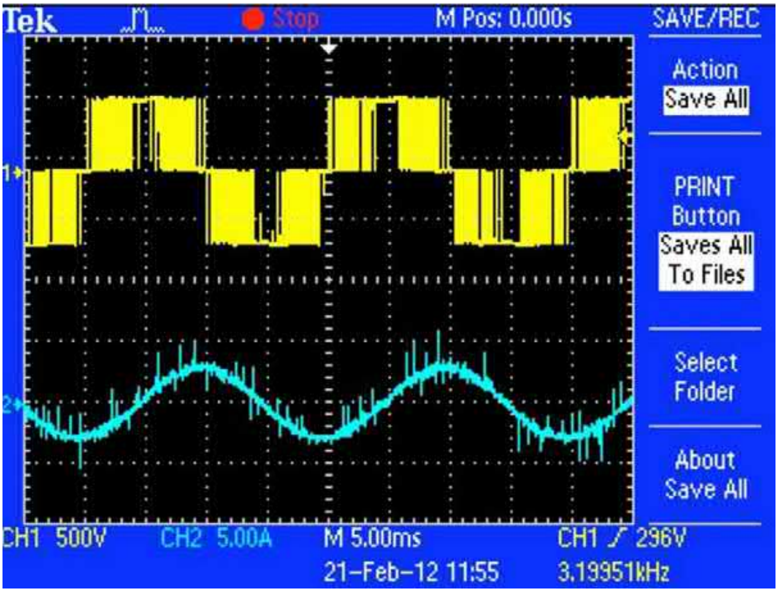

FIG. 7.9 Actual voltage and current waveforms for a star-connected, PWM-fed induction motor. Upper trace—voltage across U and V terminals. Lower trace—U phase motor current.

7.3.2 Current source induction motor drives

Although the majority of inverters used in motor drives are Voltage Source Inverters (VSI), described in Chapter 2 and discussed in the previous section, Current Source Inverters (CSI) are still sometimes used, particularly for high power applications, and warrant a brief mention.

The forced-commutated current-fed induction motor drive, shown in Fig. 7.10, was strongly favoured for single induction motor applications for a

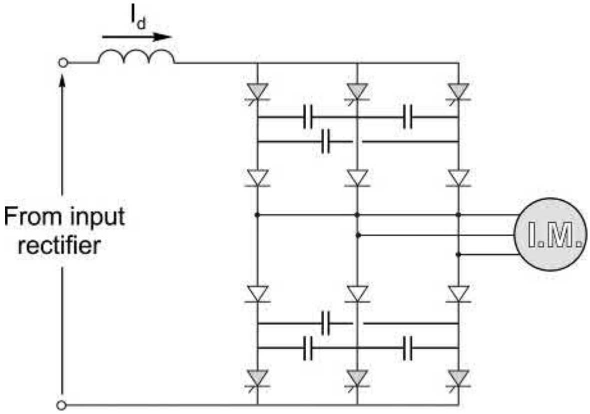

FIG. 7.10 Forced-commutated current-fed induction motor drive (motor converter only shown).

long period, and was available at power levels in the range 50–3500 kW at voltages normally up to 690 V. High voltage versions at 3.3 kV/6.6 kV were also developed but they have not proved to be economically attractive. Today it is not seen as having merit and has virtually disappeared from the portfolios of most companies. A brief description is included here for interest only.

The d.c. link current I_d, taken from a 'stiff' current source (usually in the form of a thyristor bridge and a series inductor in the d.c. link), is sequentially switched at the required frequency into the stator windings of the induction motor. The capacitors and extra series diodes provide the mechanism for commutating the thyristors by cleverly exploiting the reversal of voltage resulting from resonance between the capacitor and the motor leakage reactance. The resultant motor voltage waveform is, perhaps somewhat surprisingly, approximately sinusoidal apart from the superposition of voltage spikes caused by the rise and fall of machine current at each commutation.

The operating frequency range is typically 5–60 Hz, the upper limit being set by the relatively slow commutation process. Below 5 Hz, torque pulsations can be problematic but PWM control of the current can be used at low frequencies to ease the problem.

This system was most commonly used for single motor applications such as fans, pumps, extruders, compressors, etc. where very good dynamic performance is not necessary and a supply power factor which decreases with speed is acceptable.

7.3.3 Performance of inverter-fed drives

It has often been said that the steady-state performance of the inverter-fed induction motor is broadly comparable with that of an industrial d.c. drive, but in fact the performance of contemporary inverter-fed induction motor is better in almost all respects.

To illustrate this, we can consider how quickly an induction motor drive, with field-oriented control (see Chapter 8), can change the motor shaft torque. Remarkably, the torque can be stepped from zero to rated value and held there in less than 1 ms, and this can now be achieved by a motor even without a speed/position sensor. For comparison, a thyristor-fed d.c. motor could take up to one sixth of a 50/60 Hz supply cycle i.e. around 3 ms before the next firing pulse can even initiate the process of increasing the torque, and clearly considerably longer to complete the task.

The induction motor is also clearly more robust and better suited to hazardous environments, and can run at higher speeds than the d.c. motor, which is limited by the performance of its commutator.

As we will see in Chapter 8, field-oriented control, coupled with the ability, through a PWM inverter, to change the stator voltage phasor in magnitude, phase and frequency very rapidly is at the heart of this exceptional motor shaft performance. The majority of commercial inverter systems now embody such control strategies, but the quantification of shaft performance is subject to a large number of variables and manufacturer's data in this respect needs to be interpreted with care. Users are interested in how quickly the speed of the motor shaft can be changed, and often manufacturers quote the speed loop response, but many other factors contribute significantly to the overall performance. Some of the most important, considering a spectrum of applications, are:

- Torque Response: The time needed for the system to respond to a step change in torque demand and settle to the new demanded level.
- Speed Recovery Time: The time needed for the system to respond to a step change in the load torque and recover to the demanded speed.
- Minimum Supply Frequency at which 100% Torque can be achieved
- Maximum Torque at 1 Hz
- Speed Loop Response: This is defined in a number of different ways but a useful measure is determined by running the drive at a non-zero speed and applying a square wave speed reference and looking at the overshoot of speed on the leading edge of the square wave: an overshoot of 15% is—for most applications—considered practically acceptable. For the user, it is always a good idea to seek clarification from the manufacturer whenever figures are quoted for the speed (or current/torque) loop bandwidth of a digital drive, because this can be defined in a number of ways (often to the advantage of the supplier and not to the benefit of the application).

Note that the above measures of system performance should be obtained under conditions which avoid the drive hitting current limits, as this obviously limits the performance. Tests should typically be undertaken on a representative motor with a load inertia approximately equal to the motor inertia.

Indications of the performance of the open loop and closed loop field oriented induction motor control schemes are shown below:

	Open-loop (Without position feedback)	Closed loop (With position feedback)
Torque response (ms)	<0.5	<0.5
Speed recovery time (ms)	<20	<10
Min speed with 100% torque (Hz)	0.8	Standstill
Max torque at 1 Hz (%)	>175	>175
Speed loop response (Hz)	75	125

The performance of a closed loop inverter-fed induction motor is comparable to that of a closed loop permanent magnet motor, which we discuss in Chapter 9. This comes as a great surprise to many people (including some who have spent a lifetime in drives), possibly because the majority of induction motor drives use standard motors which were designed for fixed speed operation and broad application, whereas permanent magnet motors tend to be customised and many have been designed with relatively low inertias (long length and small diameter rotor), which facilitate rapid speed changes, or high inertias (short shaft and large diameter rotor) which promotes smooth rotation in the presence of load changes. Special induction motor designs are available, however, and are sometimes the preferred solution.

In the remainder of this section we give broad indications of the applicability of the various drive configurations that should prove helpful when looking at specific applications.

Open-loop (without speed/position feedback) induction motor drives

Open-loop induction motor drives are used in applications that require moderate performance (i.e. fans and pumps, conveyors, centrifuges, etc.). The performance characteristics of these drives are summarised below:

- Moderate transient performance with full torque production down to approximately 2% of rated speed.
- A good estimate of stator resistance improves torque production at low speeds, but the control system will work with an inaccurate estimate, albeit with reduced torque.
- A good estimate of motor slip improves the ability of the drive to hold the reference speed, but the control system will work with an inaccurate estimate, albeit with poorer speed holding.

The performance of open loop induction motor drives continues to improve. Techniques for sensing the rotational speed of an induction motor without the need for a shaft-mounted speed or position sensor pervade the technical literature, and will in time find their way into some commercial drives.

Closed-loop (with speed/position feedback) induction motor drives

Induction motor drives with closed-loop control are used in similar applications to d.c. motor drives (i.e. cranes and hoists, winders and un-winders, paper and pulp processing, metal rolling, etc.). These drives are also particularly suited to applications that must operate at very high speeds with a high level of field weakening, for example spindle motors. The performance characteristics of these drives are summarised below:

- Good dynamic performance at speeds down to standstill when position feedback is used.
- Only incremental position feedback is required. This can be provided with a position sensor or alternatively a sensorless scheme can be used. The transient performance of a sensorless scheme will be lower than when a position sensor is used and lower torque is produced at very low speeds.

- The robustness of the rotor makes induction motors particularly well suited to high speed applications that require field weakening. The motor current reduces as the speed is increased and the flux is reduced.
- Induction motors are generally less efficient than permanent magnet motors because of their additional rotor losses.

Applications when field orientation or Direct Torque Control cannot be used

Field orientation and Direct Torque Control both rely upon modelling the flux in the motor. If a single inverter is being used to feed more than one motor, as in Fig. 7.11, neither control strategy can be used.

In such systems individual motor control is not possible and the only practical form of control is to feed the motor group with an appropriate voltage source, of which the magnitude and frequency can be controlled. In fact, this is exactly the traditional form of V/f control which predominated in early inverters. The output frequency, and hence the no-load speed of the motors, is set by the speed reference signal, (traditionally 0–10 V or 4–20 mA). In most automated applications the speed demand comes from a remote controller such as a PLC, while in simpler applications there will be a digital user interface on a control panel, or on the drive itself. The drive would have facility to adjust the V/f ratio, and to boost the voltage at low frequencies to compensate for the dominating influence of the stator winding resistance.

The fact that field orientation schemes use a vector modulator/PWM controller (see later Fig. 8.19) indicates that to adapt the field oriented scheme for multi-motor drives (i.e. a number of motors connected in parallel to a single inverter) is relatively simple: all that is required is to provide a sawtooth waveform of the appropriate frequency as input to the vector modulator (see Fig. 8.18). However, for Direct Torque Control schemes (see Chapter 8) there is no such controller, and in order to provide for multi-motor operation manufacturers employing this control strategy are obliged to provide an additional PWM controller.

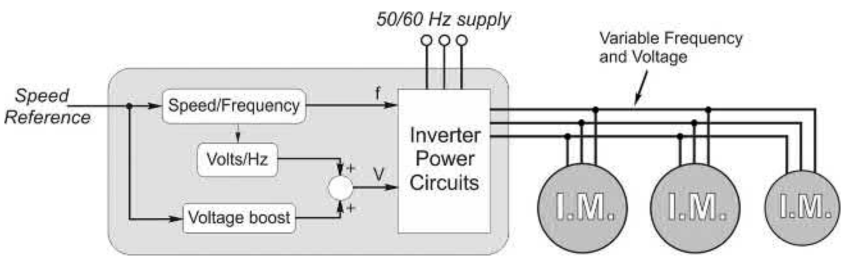

FIG. 7.11 Simple V/f induction motor control strategy applied to a multi-motor system.

7.4 Effect of inverter on the induction motor

It is often stated that standard 'off-the-shelf' a.c. motors can be used without problems with modern PWM inverters, and whilst such claims may be largely justified, inverters do have some impact and limitations are inevitable. In particular, the harmonic components of the voltages and currents create acoustic noise; they always give rise to additional iron and copper losses; and they have other effects which are perhaps less obvious. In addition, the operation of a standard motor—with its cooling system designed to suit fixed-speed operation—can be a significant limitation, and this will also be considered here.

7.4.1 Acoustic noise

Acoustic noise can usually be reduced by selecting a higher switching frequency (at the cost of higher inverter losses). It is interesting to note that not all motors exhibit the same characteristic when connected to identical inverters. The differences are usually small, and relate primarily to the clamping of the core iron in the stator of the motor. Certain switching frequencies may excite resonances in some motors, these often being related to the tie bars between the end frames: the vibrations can be alarming but easily remedied by changing the switching frequency, or, if the tie bar is external to the motor frame, by simply adding a wedge to change the natural frequency of the bar.

7.4.2 Motor insulation and the impact of long inverter-motor cables[4]

The PWM waveform has another very significant, but perhaps less obvious effect, related to the very high rates of change of voltage (dV/dt) which can result in damage to the winding insulation. In a modern 400 V (Silicon) IGBT power converter the d.c. link voltage is around 540 V, the voltage switches typically in 100 ns, and so *at the terminal of the drive* there is a very high dV/dt of over 5000 V/μs. (For Wide Band Gap devices using materials such as SiC and GaN, switching speeds are faster still.)

Recalling that the equation linking the current through and voltage across a capacitor is $i = C\frac{dv}{dt}$, it becomes clear that it is possible for appreciable current to flow in even a very small capacitance if the rate of change of voltage is high enough. In our context there are inevitably unwanted 'stray' capacitances within and between the phase-windings, and between the conductors in the supply cable and the cable screen or armouring. For example the capacitance of the cable to ground might be perhaps 100 pF/m, so a long cable of say 30 m would have a charging current pulse of 15A when the dV/dt was 5000 V/μs. However, as far as the charging current pulses are concerned, the issue is only likely to be a

4. For more detailed information on this refer to The Control Techniques Drives and Controls Handbook, 2nd Edition, by W. Drury, pages 337–351.

practical problem in small drives with very long motor cables, where the charging current pulses may, in extreme cases, exceed the rated current of the motor, and determine the rating of the required drive!

Very fast pulses take us from our familiar 'low-frequency' world into the territory of the communications or high-frequency engineer, where effects that we are usually unaware of begin to assert themselves. So whereas we can happily ignore the finite time taken for electrical effects to be transmitted from one part of a circuit to another, and hence are able to use simple lumped-parameter models (i.e. consisting of R, L, and C), these simplifications are inadequate at high frequencies. In essence, when the physical dimensions of the hardware become comparable with the wavelength of the electromagnetic phenomena, we have to resort to more elaborate distributed-parameter representations. A pulse travelling along our 50 m long drive to motor cable sees the cable as a transmission line along which energy travels at perhaps 60% of the speed of light, but nevertheless takes almost 300 ns to reach the motor. The so-called surge impedance of the cable is usually smaller than that of the motor, so a reflected pulse will be created, which in an extreme case could result in a doubling of the motor terminal voltage. The fast risetime pulses can also result in uneven voltage distribution within the motor windings, and consequent additional stressing of the insulation.

In case all this sounds alarming the fact is that such problems are extremely unusual and usually associated with systems employing old or very low cost motors with poor insulation systems, and/or with drive systems with rated voltages over 690 V. Naturally enough, the problem is more pronounced on medium voltage drives where it is not uncommon for dV/dt filters to be fitted between the inverter and the motor.

These phenomena are now very well understood by reputable motor manufacturers. International standards on appropriate insulation systems have also been published, notably IEC 34-17 and NEMA MG1pt31.

7.4.3 Losses and impact on motor rating

Operation of induction motors on an inverter supply inevitably results in additional losses in the machine as compared with a sinusoidal utility supply. These losses fall into three main categories:

(a) Stator Copper Loss—This is proportional to the square of the r.m.s. current although additional losses due to skin effect associated with the high frequency components also contribute. We have seen in Fig. 7.1 that the motor current is reasonably sinusoidal and hence, as we would expect, the increase in copper loss is seldom significant.

(b) Rotor Copper Loss—The rotor resistance is different for each harmonic current present in the rotor due to skin effect (and is particularly pronounced in deep bar rotors). Since the rotor resistance is a function of frequency, the rotor copper loss must be calculated independently for each harmonic. Whilst these additional losses used to be significant in the early days of PWM inverters with low switching frequencies, in modern drives with switching frequencies above 3 kHz the additional losses are minimal.

(c) Iron Loss—This is increased by the harmonic components in the motor voltage.

For PWM voltage source inverters using sinusoidal modulation and switching frequencies of 3 kHz or higher, the additional losses are therefore primarily iron losses and are generally small, resulting in a loss of motor efficiency by 1–2%. Motors designed for enhanced efficiency, e.g. to meet the IEC IE2/IE3 requirements or NEMA EPACT and Premium Efficiency requirements, also experience a proportionately lower increase in losses with inverter supplies because of the use of reduced-loss magnetic steels.

However, the increase in losses does not directly relate to a de-rating factor for standard machines since the harmonic losses are not evenly distributed through the machine. The harmonic losses mostly occur in the rotor and have the effect of raising the rotor temperature. Whether or not the machine was designed to be stator critical (stator temperature defining the thermal limit) or rotor critical, clearly has a significant impact on the need for, or magnitude of, any de-rating. The cooling system (see below) is at least as important, however, and in practice it emerges that a standard motor may have to be de-rated by 5 or even 10% for use on an inverter supply.

Whereas a d.c. motor was invariably supplied with through ventilation provided by an auxiliary blower, to allow it to operate continuously at low speeds without overheating, the standard induction motor has no such provision. Having been designed primarily for fixed-frequency/full-speed operation, most induction motors tend to be totally enclosed (IP44 or IP54) with a shaft-mounted fan at the non-drive end running within a cowl to duct the cooling air over a finned motor body as shown in Fig. 7.12. Note also the cast ‘paddles’ on the rotor endrings which provide internal air circulation and turbulence to assist with transmitting the heat from the rotor to the stator housing and from there to the atmosphere.

Thus although the inverter is capable of driving the induction motor with full torque at low speeds, continuous operation *at rated torque* is unlikely to be possible because the standard shaft mounted cooling fan will be less effective at reduced speed and the motor will overheat. We should say, however, that for applications such as fans and pumps where the load torque is proportional to the cube of the speed, no such problems exist, but for many applications it is a significant consideration.

7.4.4 Bearing currents

Scare stories periodically appear in the trade press and Journals relating to motor bearing failures in inverter-fed a.c. motors. It should be said immediately that such failures are rare, and mainly associated with medium voltage systems.

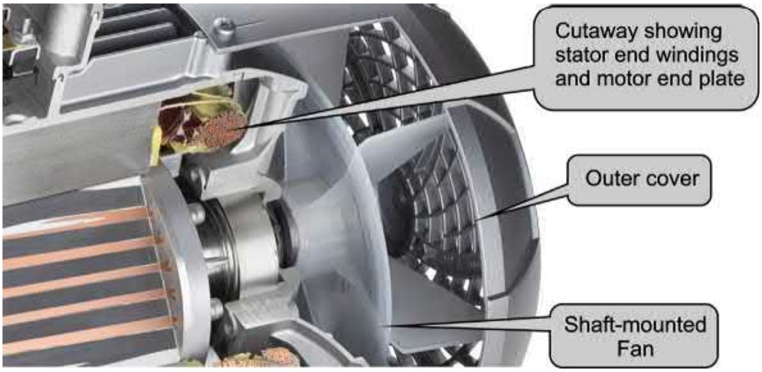

FIG. 7.12 Typical shaft-mounted external cooling fan on an a.c. induction motor. *(Courtesy of Siemens.)*

With a balanced three phase sinusoidal supply the sum of the three stator currents in an a.c. motor is zero and there is no further current flow outside the motor. In practice however there are conditions which may result in currents flowing through the bearings of a.c. motors even when fed with a sinusoidal 50 Hz or 60 Hz supply, and the risk is further increased when using an inverter supply. Any asymmetric flux distribution within an electrical machine can result in an induced voltage from one end of the rotor shaft to the other. If the bearing 'breakover voltage' is exceeded (the electrical strength of the lubricant film being of the order of 50 V) or if electrical contact is made between the moving and fixed parts of the bearing this will result in a current flowing through both bearings. The current is of low (often slip) frequency and its amplitude is limited only by the resistance of the shaft and bearings, so it can be destructive. In some large machines it is common practice to fit an insulated bearing, usually on the non-drive end, to stop such currents flowing.

Any motor may also be subject to bearing currents if its shaft is connected to machinery at a different ground potential from the motor frame. It is therefore important to ensure that the motor frame is connected through a low-inductance route to the structure of the driven machinery. This issue is well understood and with modern motors such problems are rare.

7.4.5 'Inverter grade' induction motors

Addressing the above potential hazards, induction motors carrying the name 'inverter grade' or similar are readily available. They would typically have reinforced winding insulation systems and have a thermal capacity for a constant torque operating range, often down to 30% of base speed, without the need for additional external cooling. Further they would have options to fit thermocouples, a separate cooling fan (for very low speed operation) and a speed/position feedback device.

International standards exist to help users and suppliers in this complex area. NEMA MG1-2016, Part 31 gives guidance on operation of squirrel cage induction motors with adjustable-voltage and adjustable-frequency controls. IEC 60034-17 and IEC 60034-25 give guidance on the operation of induction motors with converter supplies, and design of motors specifically intended for converter supplies, respectively.

7.5 Utility supply effects

It is a common misconception to believe that the harmonic content of the motor current waveform and the motor power factor are directly reflected on the utility supply, but this is not the case. The presence of the inverter, with its energy-buffering d.c. link capacitor results in near unity power factor as seen by the utility supply regardless of load or speed of operation, which is of course highly desirable. It is not all good news however, so we now look at the adverse impact of an inverter- fed drive on the utility.

7.5.1 Harmonic currents

Harmonic current is generated by the input rectifier of an a.c. drive shown in Fig. 7.8. The utility supply is rectified by the diode bridge, and the resulting d.c. voltage is smoothed by the d.c. link capacitor and, for drives rated typically at over 2.2 kW, the d.c. current is smoothed by an inductor in the d.c. circuit. The d.c. voltage is then chopped up in the inverter stage which uses PWM to create a sinusoidal output voltage of adjustable voltage and frequency.

Whilst small drive ratings may have a single phase supply, we will consider a three phase supply. We see from Fig. 7.13 that current flows into the rectifier as a series of pulses that occur whenever the supply voltage exceeds that of the d.c. link, which is when the diodes start to conduct. The amplitude of these pulses is much larger than the fundamental component, which is shown by the dotted line.

Fig. 7.14 shows the spectral analysis of the current waveform in Fig. 7.13.

Note that all currents shown in spectra comprise lines at multiples of the 50 Hz utility frequency. Because the waveform is symmetrical in the positive and negative half-cycles, apart from imperfections, even order harmonics are present only at a very low level. The odd order harmonics are quite high, but they diminish with increasing harmonic number. For the three phase input bridge there are no triplen (triple-frequency) harmonics, and by the 25th harmonic the level is negligible. The frequency of this harmonic for a 50 Hz supply is 1250 Hz which is in the audio frequency region of the electromagnetic spectrum and well below the radio frequency part which is generally considered to

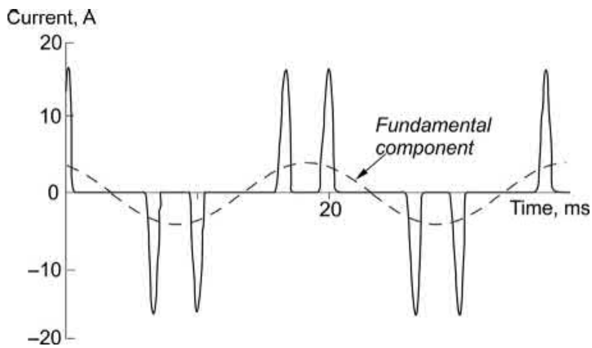

FIG. 7.13 Typical current from utility supply for a 1.5kW 3-phase drive.

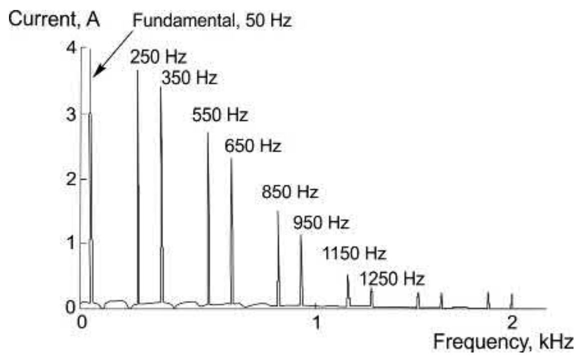

FIG. 7.14 Harmonic spectrum of the current waveform shown in Fig. 7.13.

begin at 150kHz. This is important, because it shows that supply harmonics are low frequency effects, which are quite different from radio frequency electromagnetic compatibility (EMC) effects. They are not sensitive to fine details of layout and screening of circuits, and any remedial measures which are required use conventional electrical power techniques such as tuned power factor capacitors and phase-shifting transformers. This should not be confused with the various techniques used to control the radio-frequency interference from fast switching devices, sparking electrical contacts, etc.—all matters that relate to the 'high-frequency world' referred to in Section 7.4.2.

The actual magnitudes of the current harmonics depend on the detailed design of the drive, specifically the values of d.c. link capacitance and, where used, d.c. link inductance, as well as the impedance of the utility system to which it is connected, and the other non-linear loads on the system.

We should make clear that industrial problems due to harmonics are unusual, although with the steady increase in the use of electronic equipment, they will become more common in future. Problems have occurred most frequently in office buildings with a very high density of personal computers, and in cases where most of the supply capacity is used by electronic equipment such as drives, converters and UPS.

As a general rule, if the total rectifier loading (i.e. drives, UPS, PCs, etc.) on a power system comprises less than 20% of its current capacity then harmonics are unlikely to be a limiting factor. In many industrial installations the capacity of the supply considerably exceeds the installed load, and a large proportion of the load such as uncontrolled (direct on line) induction motors and resistive heating elements generate minimal harmonics.

If rectifier loading exceeds 20% then a harmonic control plan should be in place. This requires some experience and guidance can often be sought from equipment suppliers. The good news is that if it is considered that a problem will exist with the estimated level of harmonics then there are a number of options available to reduce the distortion to acceptable levels.

A.C. drives rated over 2.2 kW tend to be designed with inductance built in to the d.c. link and/or the a.c. input circuit. This gives the much better supply current waveform and its dramatically improved spectrum as shown in Figs. 7.15 and 7.16 respectively, which are again for a 1.5 kW drive for ease of comparison with the previous illustrations. (In this case the inductance in each line is specified as '2%', which means that when rated fundamental current flows in the line, the volt-drop across the inductor is equal to 2% of the supply voltage.) Note the change of vertical scale between Figs. 7.13 and 7.15, which may tend to obscure the fact that the pulses of current now reach about 5 A, rather than the 17 A or so previously, but the fundamental component remains at 4 A because the load is the same. (Remember that whilst we have just demonstrated the tremendous improvement in supply harmonics achieved by adding d.c. link inductance to a 1.5 kW drive, standard drives would rarely be manufactured

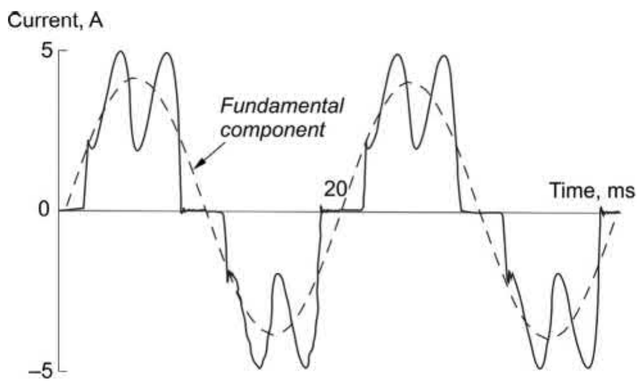

FIG. 7.15 Utility supply current waveform for a 1.5 kW 3-phase drive with d.c. and 2% a.c. inductors.

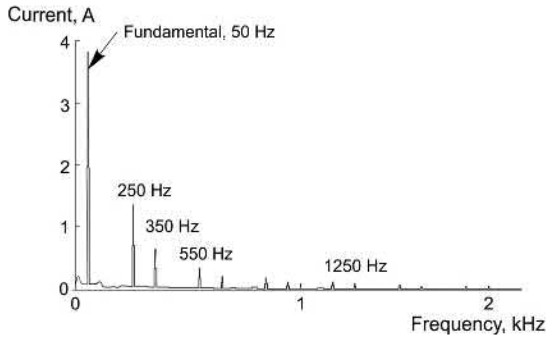

FIG. 7.16 Harmonic spectrum of the improved current waveform shown in Fig. 7.15.

with any inductance because whilst the harmonic spectrum looks worrying, the currents are at such a low level that they would rarely cause practical problems.)

Standard three phase drives rated up to about 200 kW tend to use conventional 6 pulse rectifiers. At higher powers, it may be necessary to increase the pulse number to improve the supply-side waveform, and this involves a special transformer with two separate secondary windings, as shown for a 12-pulse rectifier in Fig. 7.17.

The voltages in the transformer secondary star and delta windings have the same magnitude but a relative phase shift of 30°. Each winding has its own set of six diodes, and each produces a six-pulse output voltage. The two outputs are generally connected in parallel, and because of the phase shift, the resultant voltage consists of twelve pulses of 30° per cycle, rather than the six pulses of 60° shown for example in Fig. 2.13.

The phase shift of 30° is equivalent to 180° at the fifth and seventh harmonics (as well as 17, 19, 29, 31, etc.), so that flux and hence primary current at these harmonics cancels in the transformer, and the resultant primary

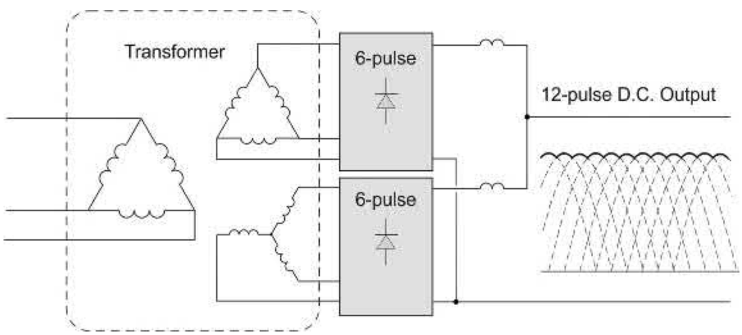

FIG. 7.17 Basic 12-pulse rectifier arrangement.

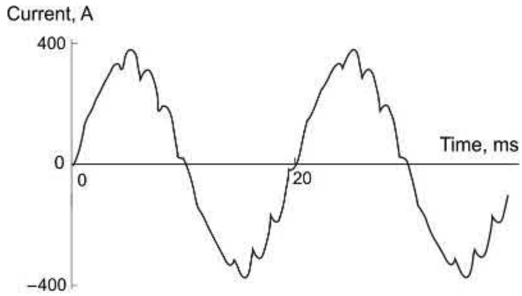

FIG. 7.18 Utility supply current waveform for 150 kW drive with 12-pulse rectifier.

waveform therefore approximates very well to a sinusoid, as shown for a 150 kW drive in Fig. 7.18.

The use of drive systems with an input rectifier/converter using PWM which generates negligible harmonic current in the utility supply, as described in Section 2.4.6, is becoming increasingly common. This also permits the return of power from the load to the supply.

7.5.2 Power factor

The power factor of an a.c. load is a measure of the ratio of the average power to the product of r.m.s. current and voltage, and is given by:

$$\text{Power Factor} = \frac{\text{Average Power (W)}}{\text{r.m.s. Volts} \times \text{r.m.s. Amps}}$$

With a sinusoidal supply voltage and a linear load the current will also be sinusoidal, with a phase-shift of ϕ with respect to the voltage. The power is then given by the simple expression

$$W = VIcos\phi,$$

where V and I are r.m.s. values (which are equal to the peak of the sinusoid divided by $\sqrt{}2$), and so in this case the power factor is equal to cos ϕ. Clearly the maximum possible power factor is 1.

Unfortunately, in power electronic circuits either the voltage or the current or both are non-sinusoidal, so there is no simple formula for the r.m.s. values or the mean power, all of which have to be found by integration of the waveforms. There is therefore no simple formula for the power factor, but frequent use is made of a related quantity known as the fundamental power factor, given by

$$\text{Fundamental Power Factor} = \frac{\text{Average Power}}{\text{Fundamental r.m.s. Volts} \times \text{Fundamental r.m.s. Amps}}$$

The influence of the harmonics in the non-sinusoidal waveforms causes the actual power factor to be lower than the fundamental power factor, so users should be aware that when suppliers quote the power factor of a drive they are usually ignoring the harmonic currents, and quoting cos φ, the fundamental power factor.

It may be worth reminding readers who are not familiar with industrial energy tariffs why maximising the power factor is important. All industrial users pay primarily for the energy used, which depends on the integrated total of the product of power and time, but most are also penalised for drawing the power at a low power factor (because the currents are higher and therefore switchgear and cables have to be larger than would otherwise be necessary). In addition, there may be a penalty related to the maximum VoltAmpere product in a specified period, so again a high power factor is desirable.

Fortunately, for the diode bridge, which is the most common form of rectifier in a commercial a.c. drive, cos ϕ is close to unity for all speed and load conditions. To illustrate this, we can consider a typical 11 kW induction motor operating at full load, connected either directly to the utility supply or through an a.c. variable speed drive. Comparative figures are given in the table below.

At supply terminals	Direct on line motor	Motor via a.c. drive	Notes on drive parameters
Voltage (V)	400	400	
r.m.s. current (A)	21.1	21.4	No significant change.
Fundamental current (A)	21.1	18.8	Reduced because magnetising current is not drawn directly from the utility supply.
Fundamental power factor (cosφ)	0.85	0.99	Improved because input rectifier current is in phase with supply voltage.
Power (W)	12,440	12,700	Slight increase at full load due to drive losses.

A typical PWM induction motor drive improves the power factor as compared with a direct on line motor because it reduces the requirement of the supply to provide the magnetising current for the motor, but in return generates harmonics. Power consumption at full load is slightly increased due to losses of the drive.

7.6 Inverter and motor protection

We have stressed before that power semiconductors are notoriously intolerant of excess current, and so even in the earliest drives of this type, current was measured in order to trip the drive when a simple current threshold was exceeded and before damage could be done to the inverter. Some protection schemes would also sense high currents and reduce the applied frequency and thereby reduce the current.

The stored energy in the drive and motor inductances and capacitances also needs to be handled without inducing voltages or currents which can damage the system components. As previously mentioned, the basic power circuit is not inherently capable of regenerating energy back into the supply, and when a braking duty results in energy flow into the d.c. link then a correctly-rated 'dump resistor' (see Section 2.4.5) must be provided in order to limit the circuit voltages.

Motor protection also requires current measurement, but here it is thermal protection of the motor that is of concern. A very approximate indication of the losses or heating effect in the motor is obtained by monitoring the product of the square of the motor current times time. This so-called 'i^2t' protection is still referred to as motor thermal protection in drives, though many of the thermal algorithms now employed are very much more complex and accurate than their primitive predecessors.

Modern commercial drives include extensive internal protection systems as well as thermal motor modelling systems, but such drives are designed for a multiplicity of applications and motor designs and so must be configured during installation. Where multiple motors are fed from a single inverter (as described in Section 7.3.3) each motor must have its own individual thermal trip, because the fault current of any individual motor alone may not be significant when a large number of motors are connected to the same inverter.

7.7 Review questions

(1) Choose a suitable pole-number for an induction motor to cover the speed range from 400 rev/min to 800 rev/min when supplied from a 30–75 Hz variable-frequency source.

(2) A 2-pole, 440 V, 50 Hz induction motor develops rated torque at a speed of 2960 rev/min; the corresponding stator and rotor currents are 60 A and 150 A, respectively. If the stator voltage and frequency are adjusted so that the flux remains constant, calculate the speed at which full torque is developed when the supply frequency is (a) 30 Hz, (b) 3 Hz.

(3) Estimate the stator and rotor currents and the rotor frequency for the motor in question 3 at 30 Hz and at 3 Hz.

(4) What is 'voltage boosting' in an open loop voltage-source inverter, and why is it necessary?

(5) An induction motor with a synchronous speed of N_s is driving a constant-torque load at base frequency, and the slip is 5%. If the frequency of the supply is then doubled, but the voltage remains the same, estimate the new slip speed and the new percentage slip.

(6) Approximately how would the efficiency of an inverter-fed motor be expected to vary between full (base) speed, 50% speed and 10% speed, assuming that the load torque was constant at 100% at all speeds and that the efficiency at base speed was 80%.

(7) Why is it unwise to expect a standard induction motor driving a high-torque load to run continuously at low speed?

(8) Explain briefly why an inverter-fed induction motor will probably be able to produce more starting torque per ampere of supply current than the same motor would if connected directly to the utility supply. Why is this likely to be particularly important if the supply impedance is high?

(9) Why is the harmonic content of an inverter-fed induction motor current waveform less than the harmonic content of the voltage waveforms?

Answers to the review questions are given in Appendix.

Chapter 8

Field oriented control of induction motors

8.1 Introduction

In this chapter, we explore the contemporary approach to control of the inverter/induction motor combination. Field-oriented (or vector) control allows the induction motor/inverter combination to outperform conventional industrial d.c. drives, and its progressive refinement since the 1980s represents a major landmark in the history of electrical drives. It is therefore appropriate that its importance is properly reflected in this book, because one of our aims is to equip readers with sufficient understanding to allow them to converse intelligently with manufacturers and suppliers. It is also all too easy for designers of these control systems to get so absorbed in the mathematical equations describing field orientation, that the fundamental understanding of what is actually happening can be lost: for those readers this chapter might offer a welcome re-orientation.

We prefer not to use the term ‘inverter-fed’ in these circumstances because although the motor derives its supply from an inverter, the switching of the inverter devices is determined at every instant by the state of the flux and currents in the motor, the switching being continuously optimised to achieve the torque required. Both field oriented and its close relative direct torque control methods only became possible with the development of fast, cheap, digital processors that can implement the high-speed calculations that are necessary to model and control the motor in real time.

Understanding field-oriented control is usually regarded as challenging, even for experienced drives personnel, not least because the subject tends to be highly mathematical. Anyone who has consulted an article or textbook on the subject of field-oriented control will quickly become aware that most treatments involve extensive use of matrices and transform theory, and that many of the terms used will not be familiar to anyone not already schooled in the analysis of electrical machines.

Our aim is to continue to cover new topics without recourse to any very demanding mathematics, and we believe that it is possible to understand the underlying basis of field-oriented control via a relatively simple graphical

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00008-8

approach. However, even for this we have to make use of several disparate ideas that we have not discussed previously, so Section 8.2 is included to acquaint the non-specialist reader with the new tools and insights that we employ later to explain how field orientation works. It consists of three parts, dealing with space phasors; transformation of reference frames; and transient and steady-state conditions in electric circuits. Readers who are already comfortable with these matters may wish to skip this section.

Section 8.3 covers the modelling of the induction motor in terms of a set of magnetically coupled circuits, rather than the physical field-based approach that we have used so far. We include an outline of this approach for completeness, but our interest is in interpreting the outcomes of the analysis, so we do not discuss the solution of the equations in detail.

In Section 8.4 we study the steady-state behaviour of an induction motor when its stator *currents*, rather than *voltages* are prescribed. Precise and rapid control of the stator currents is achieved via high-bandwidth closed-loop current controllers, and not surprisingly, it emerges that the behaviour of the motor is very different from what we have seen hitherto. In particular, by controlling the stator currents so that the rotor flux linkage remains constant regardless of slip, the torque is directly proportional to slip. This leads to a delightfully simple approach to torque control, the motor and inverter combination behaving in a similar (but superior) manner to the thyristor d.c. drive that we studied in Chapter 4.

The real benefit of the current-driven system is revealed in Section 8.5, where we see that steady-state torque control can equally well be applied under dynamic conditions, so that, for example, an almost instantaneous and transient-free step of torque is readily achievable. The practical implementation of this remarkable technique is dealt with in Section 8.6, while Section 8.7 provides an introduction to an alternative approach via so-called direct torque control.

8.2 Essential preliminaries

8.2.1 Space phasor representation of m.m.f. waves

The space phasor (or space vector) provides a shorthand graphical way of representing sinusoidally distributed spatial quantities such as the m.m.f. and flux waves that we explored in Chapter 5. It avoids us having to consider individual currents by focusing on their combined effects, and thus makes things easier to understand.

We begin by taking a fresh look at the rotating stator m.m.f., making the reasonable assumption that each of the three phase windings produces a sinusoidally distributed m.m.f. with respect to distance around the air-gap, which in turn implies that the winding itself is sinusoidally distributed (rather than sitting in clearly defined groups of coils as in the real machine discussed previously). For convenience, we will consider a 2-pole winding.

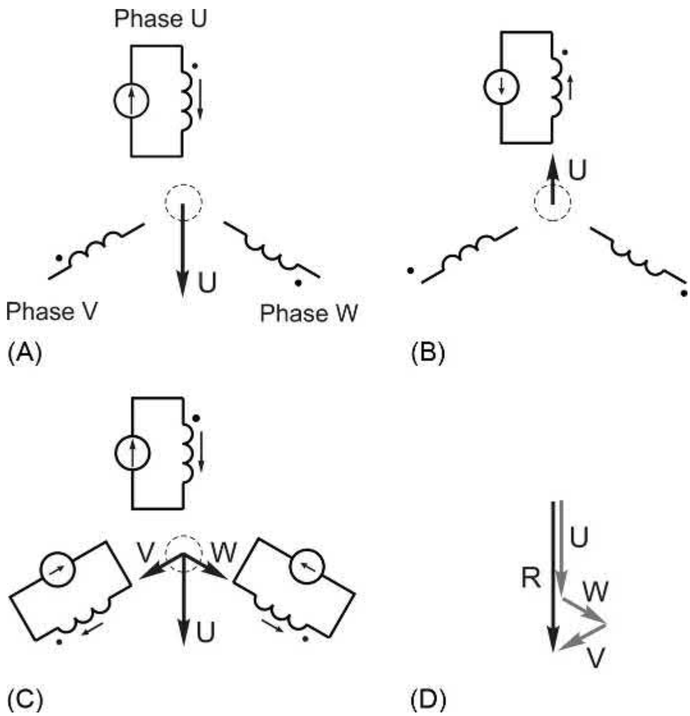

FIG. 8.1 Space phasor representation of m.m.f. waves.

We can represent the relative position of the windings *in space* as shown in Fig. 8.1. We will use the standard notation (UVW) from here onwards, although the notation ABC is still often used.

In Fig. 8.1A phases V and W are on open-circuit so that we can focus on how the m.m.f. of phase U is represented. When the current in phase U is positive (i.e. current flows into the dotted end), we have chosen to represent its sinusoidal m.m.f. pattern by a vector along the axis of the winding and pointing away from it (Fig. 8.1A), and so when the current is negative the vector points towards the coil (Fig. 8.1B). The length of the vector is directly proportional to the instantaneous value of the current, as indicated by the relative sizes of the two vectors.

In Fig. 8.1C, phase U has its maximum positive current, while phases V and W both have negative currents of half the maximum value. Because each m.m.f. is distributed sinusoidally in space, we can find their resultant (R) using the approach that is probably more familiar in the context of a.c. circuits, i.e. by adding the three components vectorially. In this particular example, the resultant m.m.f., R (Fig. 8.1D) is co-phasal with the m.m.f. of phase U, but one and half times larger.

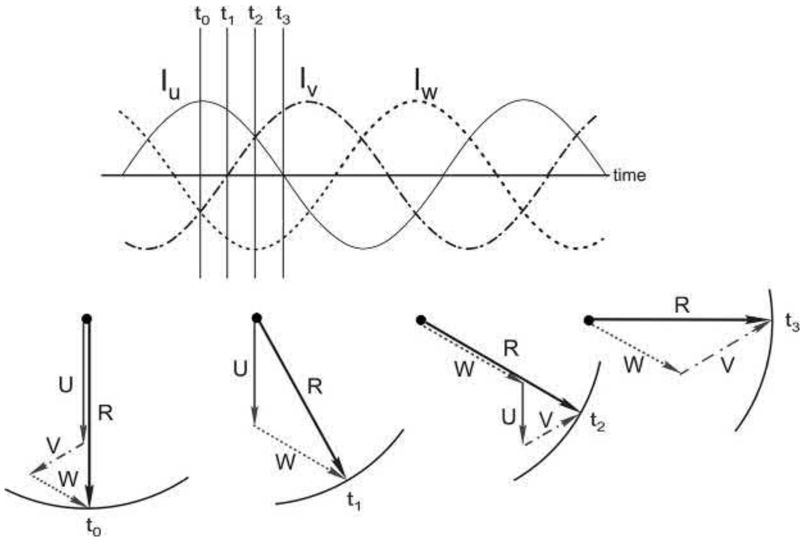

FIG. 8.2 Resultant m.m.f. space phasor for balanced 3-phase operation at four discrete times, each separated by one-twelfth of a cycle (i.e. 30°).

We can now use the approach outlined above to find the resultant m.m.f. when the windings are supplied with balanced 3-phase currents of equal amplitude but displaced in time by one third of a cycle (i.e. 120°). The axes of the phases are displaced in *space* as shown in Fig. 8.1, and the three currents are shown as functions of *time* in the upper part of Fig. 8.2. Four consecutive times are identified, separated by one twelfth of a complete cycle, or 30° in angle terms.

The lower part of the diagram represents the m.m.f.'s in a space phasor diagram. At each instant the three individual phase m.m.f.'s are shown in magnitude and position, together with the resultant m.m.f. (R). At time t_0 for example, the situation is the same as in Fig. 8.1, with phase U at maximum positive current and phases V and W having equal currents of half the magnitude of that in phase U; at t_1 phase V is zero while phases U and W have equal but opposite currents; and so on.

The four sketches suggest that the resultant m.m.f. describes an arc of constant radius, and it can easily be shown analytically that this is true. So although each phase produces a pulsating m.m.f. along its axis, the overall m.m.f. is constant in amplitude (with a value equal to 1.5 times the phase peak), and it rotates at a uniform rate, completing one mechanical revolution per cycle if the field is 2-pole (as here), half a mechanical revolution if 4-pole, etc. This is in line with our findings in Chapter 5.

We should note that although we have developed the idea of space phasors by focusing on steady-state sinusoidal operation, the approach is equally valid for any set of instantaneous currents, and is therefore applicable under transient conditions, for example during acceleration when the instantaneous frequency of the currents may change continuously.

An alternative way of representing the *resultant m.m.f. pattern* produced by a set of balanced 3-phase windings follows naturally from the discussion above. We imagine a hypothetical *single m.m.f. vector* that has a constant magnitude but *rotates relative to the stator* at the synchronous speed. This turns out to be an exceptionally useful mental picture when we come to study the behaviour of the inverter-fed induction motor, because the currents will be under our control and we are able to specify precisely the magnitude, speed and angular position of the stator m.m.f. vector in order to achieve precise control of torque.

8.2.2 Transformation of reference frames

In the previous section we saw that the resultant m.m.f. was of constant amplitude and rotated at a constant angular velocity with respect to a reference frame fixed to the stator. As far as an observer in the stationary reference frame is concerned, the same m.m.f. could equally well be produced by a sinusoidally distributed winding fed with constant (d.c.) current and mounted on a structure that rotated at the same angular velocity as the actual m.m.f. wave. On the other hand, if we were attached to a reference frame rotating with the m.m.f., the space phasor would clearly appear to us to be constant.

Transformations between reference frames have long been used to simplify the analysis of electrical machines, especially under dynamic conditions, but until fast signal-processing became available it was seldom used for live control purposes. We will see later in this chapter that the method is used in field-oriented control schemes to transform the stator currents into a rotating reference frame locked to the rotating rotor flux space phasor, thereby making them amenable for control purposes.

Transformation is usually accomplished in two stages, as shown in Fig. 8.3.

The first stage involves replacing the three windings by two in quadrature, with balanced sinusoidal currents of the same frequency but having a 90° phase shift. In this case the ‘α β’ stationary reference frame has phase α aligned with phase U. To produce the same m.m.f., either the two windings need more turns, or more current, or a combination of both. This is known as the Clarke transformation. The resultant space phasor (at the bottom of the diagram) is of course identical with the three-phase one on the left.

The second stage (the Park transformation) is more radical as the new variables I_d and I_q are in a rotating reference frame, and they remain constant under steady state conditions, as shown in Fig. 8.3. Again we need to specify the turns ratio and/or the current scaling. (Strictly speaking there is no need for the intermediate (2-phase) transformation, because we can transform directly from 3-phase to two-axis, but we have included it because it is often mentioned in

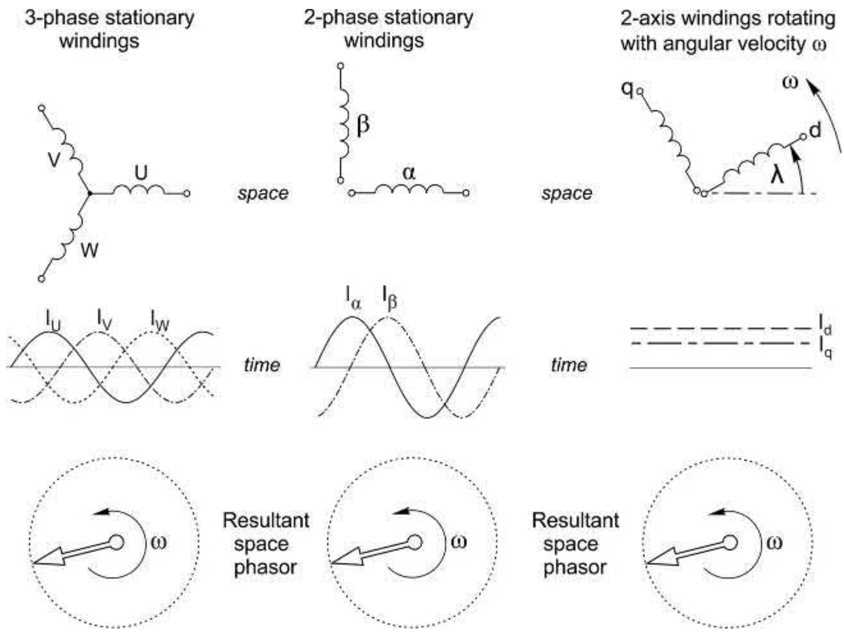

FIG. 8.3 Transformation from 3-phase stationary axis reference frame to two-axis (d-q) rotating reference frame.

the literature.) The resultant space phasor is again identical to the three-phase one.

It should be clear that the magnitude of the currents I_d and I_q will depend on the angle λ, which is the angle between the two reference frames at a specified instant, typically at $t = 0$. As far as we are concerned, it is sufficient to note that there are well-established formulae relating the input and output variables, both for the forward transformation (U, V, W to d, q) and for the inverse transformation, so it is straightforward to construct algorithms to perform the transformations. We should also note that whilst we have considered the transformation of sinusoidal currents, the technique is equally valid for instantaneous values.

8.2.3 Transient and steady-states in electric circuits

Field-oriented control allows us to obtain (almost) instantaneous (step) changes in torque on demand, and in essence it does this by *jumping directly from one steady-state condition to another*, without any unwelcome transient period of adjustment.

Given the very poor inherent transient response of the induction motor to sudden changes in load or utility supply (see for example Fig. 6.7), it will come as no surprise when we learn later that this sudden transition between steady states can only be achieved by precise control of the magnitude, frequency

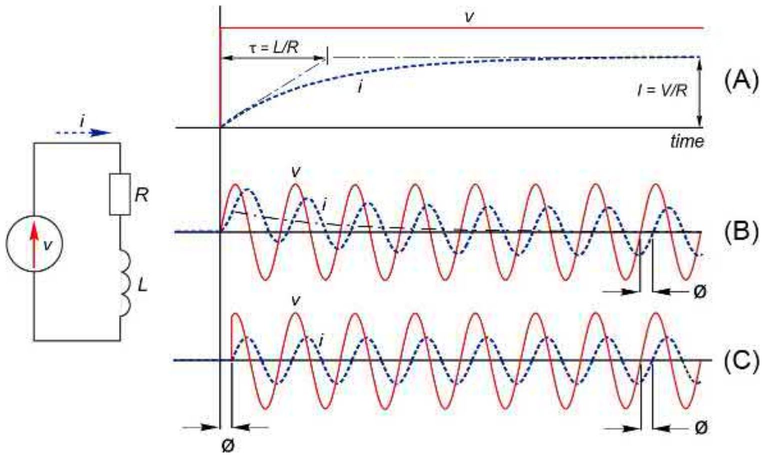

FIG. 8.4 Transition between steady states in series R-L circuit.

and instantaneous position of the stator current space vector. As will emerge, the key requirement for a successful sudden transition is that it must not involve a step change in the stored energy of the system.

As an introduction to the underlying principle of changing from one steady state to another without any transient, we can look at the behaviour of a series resistor and inductor circuit fed by an ideal voltage source (Fig. 8.4). This is much simpler than the induction motor, (it only has one energy storage element—the inductor) but it demonstrates the key requirement to be satisfied for transient-free switching.

First, we look at the current when the voltage is a step at $t=0$ (Fig. 8.4A). The steady-state current is simply V/R, but the current cannot rise instantaneously because that would require the energy stored in the inductor ($\frac{1}{2}Li^2$) to be supplied in zero time, which corresponds to an impulse of infinite power. So in addition to the steady state term $i_{ss}=\frac{V}{R}$, there is a transient term given by $i_{tr}=-\frac{V}{R}e^{-\frac{t}{\tau}}$, where the time-constant, $\tau={}^{L}/_{R}$. The total current is the sum of the steady-state and transient components, as shown by the dotted blue line in Fig. 8.4A.

Now consider a more relevant situation, where we wish the current to jump suddenly from a steady state at one frequency (in this case zero amplitude at zero frequency (or 'd.c.') for $t<0$ to a sinusoidal steady state for $t>0$.

Fig. 8.4B shows what happens if we make the sudden transition in the applied voltage (from zero d.c.) at a point where the new voltage waveform is zero but rising, i.e. at $t=0$. We note that the current does not immediately assume its steady state, but displays the characteristic decaying transient, lasting for several cycles before the steady-state is reached, with the current finally lagging the voltage by an angle ϕ. Examination of the steady state current waveform shows that the current is negative as the voltage rises through zero, so this particular attempt to jump straight into the steady state is clearly doomed from the outset because it would have required the circuit to anticipate the arrival of the voltage by having a negative current already in existence!

The fundamental reason for the transient adjustment in Fig. 8.4B is that we are seeking an instantaneous increase in the energy stored in the inductor from its initial value of zero, which is clearly impossible. It turns out that if we want to avoid the transient, we must make the jump without requiring a change in the stored energy, which in this example means at the point when the current passes through zero, as shown in Fig. 8.4C. The voltage that has to be applied therefore starts abruptly at a value $V\sin\phi$, as shown, and the current immediately enters its steady state, with no transient term.

We will see later that the principle of not disturbing the stored energy is essential to obtain sudden step changes in torque from an induction motor.

8.3 Circuit modelling of the induction motor

Up to now in this book we have developed our understanding by starting with a physical picture of the interactions between the magnetic field and current-carrying conductors, but we quickly realised that in the case of both the d.c. machine (and the utility-fed induction motor) there was a lot to be gained by making use of an 'equivalent circuit' model, particularly in terms of performance prediction. In so-doing we were representing all the distributed interactions of the motor by way of their ultimate effect as manifested at the electrical terminals and the mechanical 'terminal', i.e. the output shaft.

As long ago as the early nineteenth century it was known that the a.c. transformer could be analysed as a pair of magnetically linked coils, and it did not take long to show that all of the important types of a.c. electrical machine can also be analysed by regarding them as a set of circuits, the electrical parameters (resistance, inductance) being either measured or calculated. The vital difference compared with the static transformer is that in the machine, the coils on the rotor move with respect to those on the stator, thereby causing a variation in the extent of the magnetic interaction between the rotor and stator. This variation turns out to be the essential requirement for the machine to produce torque and to be capable of energy conversion.

8.3.1 Coupled circuits, induced EMF, and flux linkage

By 'coupled circuits' we mean two or more circuits, often in the form of multi-turn coils sharing a magnetic circuit, where the magnetic flux produced by the current in one coil not only links with its own winding, but also with those of the other coils. The coupling medium is the magnetic field, and as we will see the electrical effect of the coupling is manifested when the flux changes.

We know from Faraday's law that when the magnetic flux (ϕ) linking a coil of N turns changes, an e.m.f. (e) is induced in the coil, given by

$$e = -N\frac{d\phi}{dt},$$

i.e. the e.m.f. is proportional to the rate of change of the flux. (The minus sign indicates that if the induced e.m.f. is allowed to drive a current, the m.m.f. produced by the current will be in opposition to that which produced the original changing flux.) This equation only applies if all the flux links all N turns of the coil, the situation most commonly approached in transformer windings that share a common magnetic circuit, and are thus very tightly coupled.

We have seen that windings for induction motors are distributed, and the flux wave produced by the current in the winding is also distributed around the air-gap. As a result not all of the flux produced by one winding links with all of its turns, and we have to perform a summation (integration) of all the 'turns times flux that links them' contributions to find the 'total effective self flux linkage' which we denote by the symbol ψ (psi). The e.m.f. induced when the self-produced flux linkage changes in, say, a stator winding (subscript s) is then given by

$$e_S = \frac{d\psi_S}{dt}.$$

In an induction motor there are three distributed windings on the stator, and either a cage or three more distributed windings on the rotor, and some of the flux produced by current in any one of the windings will link all of the others. We term this 'mutual flux linkage', and often denote it by a double suffix: for example the symbol ψ_{SR} is the mutual flux linkage between a stator winding and a rotor winding.

In the same way that an e.m.f. is induced in a winding when its self-produced flux changes, so also are e.m.f.'s induced in all other windings that are mutually coupled to it. For example if the flux produced by the stator winding changes, the e.m.f. in the rotor (subscript R) is given by

$$e_R = \frac{d\psi_{SR}}{dt}.$$

8.3.2 Self and mutual inductance

The self and mutual flux linkages produced by a winding are proportional to the current in the winding: the ratio of flux linkage to the current that produces it is therefore a constant, and is defined as the inductance of the winding. The self inductance (L), is given by

$$L = \frac{Self\,flux\,linkage}{Current} = \frac{\psi_S}{i_S},$$

while the mutual inductance (M) is defined as

$$M_{SR} = \frac{Mutual\ flux\ linkage}{Current} = \frac{\psi_{SR}}{i_S}.$$

The self and mutual inductances therefore depend on the design of the magnetic circuit and the layout of the windings. In an induction motor the self inductances are constant, but the mutual inductance between a stator and a rotor winding varies with the angular position of the rotor.

We can now re-cast the expressions for e.m.f. derived above so that they involve the rates of change of the currents and the inductances, rather than the fluxes. This is a very important simplification because it allows us to represent the distributed effects of the magnetic coupling in single lumped-parameter electric circuit terms. The self-induced and mutually-induced e.m.f.'s are now given by

$$e_S = L\frac{di_S}{dt},$$

and

$$e_R = M_{SR}\frac{di_S}{dt}.$$

As mentioned above, the mutual inductance between stator and rotor windings varies with the rotor position, so M_{SR} is a function of θ, and the rotor e.m.f. has to be expressed as

$$e_R = M_{SR}\frac{di_S}{dt} + i_S\frac{dM_{SR}}{dt} = M_{SR}\frac{di_S}{dt} + i_s\frac{dM_{SR}}{d\theta}\left(\frac{d\theta}{dt}\right)$$

The first term of the equation is the 'transformer e.m.f.' that results from changes in the stator current, while the second is present even when the stator current is constant, and it is proportional to the speed of the rotor. We have already seen the vital role that this 'motional' e.m.f. plays in the energy conversion process. When the term occurs in a circuit model, it is often referred to as a 'speed' voltage.

8.3.3 Obtaining torque from a circuit model

We represent the two sets of 3-phase distributed windings of the induction motor by means of the six fictitious 'equivalent' coils shown in Fig. 8.5. (We are using the well-proven fact that a cage rotor behaves in essentially the same way as one with a wound rotor, as explained in Chapter 5.) The three stator coils remain stationary, while those on the rotor obviously move when the angle θ changes.

Because the air-gap is smooth, and the rotor is assumed to be magnetically homogeneous, all the self inductances are independent of the rotor position, as

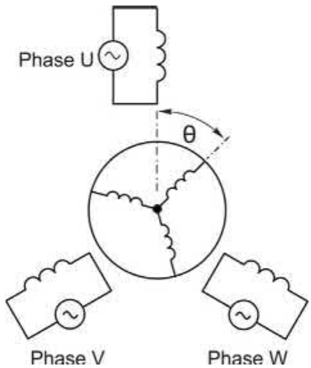

FIG. 8.5 Coupled-circuit model of 3-phase induction motor.

are the mutual inductances between pairs of stator coils and between pairs of rotor coils. Symmetry also means that the mutuals between any two stator or rotor phases are the same.

However, it is obvious that the mutual inductance between a stator and a rotor winding will vary with the position of the rotor: when stator and rotor windings are aligned, the flux linkage will be maximum, and when they are positioned at right angles, the flux linkage will be zero. With windings that are distributed so as to produce sinusoidal flux waves, the mutual inductances vary sinusoidally with the angle θ.

To a circuit theory practitioner, it is this variation of mutual inductance with position that immediately signals that torque production is possible. In fact it is straightforward (if somewhat intellectually demanding) to show that the torque produced when the sets of windings in Fig. 8.5 carry currents is given by the rather fearsome-looking expression

$$T = \sum i_S i_R \frac{dM_{SR}}{d\theta}$$

What this means is that to find the total torque we have to find the summation of nine terms, each of which represents a contribution to the total torque from one of the nine stator-rotor pairs. So we need the instantaneous value of each of the six currents, and the rate of change of inductance with rotor position for each stator-rotor pair. For example the term representing the contribution to torque made by stator coil U interacting with rotor coil V is given by

$$T_{SURV} = i_{SU} i_{RV} \frac{dM_{SURV}}{d\theta}$$

In practice we can use various expedients to simplify the torque expression, for example we know that mutual inductance is a reciprocal property, i.e. $M_{UV} = M_{VU}$, and we can exploit symmetry, but the important thing to note here is that it is a straightforward business to find the torque from the circuit model, provided that we know the currents and the angle-dependancy of the inductances.

8.3.4 Finding the rotor currents

In the induction motor, the rotor currents are induced, and to find them we have to solve the set of six equations relating them to the applied stator voltages, using Kirchoff's voltage law.

So for example the voltage equation below relating to rotor phase U includes a term representing the resistive volt drop, another representing the self-induced e.m.f. and five others representing the mutual coupling with the other windings. There are two more rotor equations and three similar ones for the stator windings.

$$v_{RU} = i_{RU}R_R + L_{RU}\frac{di_{RU}}{dt} + M_{RURV}\frac{di_{RV}}{dt} + M_{RURW}\frac{di_{RW}}{dt} + M_{RUSU}\frac{di_{SU}}{dt} + M_{RUSV}\frac{di_{SV}}{dt} + M_{RUSW}\frac{di_{SW}}{dt}$$

In the induction motor the rotor windings are usually short-circuited, so there is no applied voltage and the left-hand side of each rotor equation is zero.

If we have to solve these six simultaneous differential equations when the stator terminal *voltages* are specified (typical of utility-fed constant-frequency conditions), we have a very challenging task that demands computer assistance, even under steady-state conditions. However, when the stator *currents* are specified (as we will see is the norm in an inverter-fed motor under vector control), the equations can be solved much more readily. Indeed under steady-state locked rotor conditions we can employ an armoury of techniques such as j notation and phasor diagrams to solve the equations by hand.

We have now seen in principle how to predict the torque, and how to solve for the rotor currents, when the stator currents are specified. So we are now in a position to see what can be learned from a study of the known outcomes under two specific conditions.

In the next section, we look at how the torque varies when the stator windings are fed with a balanced set of 3-phase a.c. currents of constant amplitude but variable frequency, and the rotor is stationary. Although this is not of practical importance, it is very illuminating, and it points the way to the second and much more significant mode of operation, in which the net rotor flux is kept constant at all frequencies: this forms the basis for field-oriented control.

8.4 Steady-state torque under current-fed conditions

Historically there was little interest in analysis under current-fed conditions because we had no means of direct control over the stator currents, but the inverter-fed drive allows the stator currents to be forced very rapidly to whatever value we want, regardless of the induced e.m.f.'s in the windings. Fortunately, knowing the currents from the outset makes quantifying the torque very much easier, and it also allows us to derive simple quantitative expressions that indicate how the machine should be controlled to achieve precise torque control, even under dynamic conditions.

To simplify our mental picture we will begin with the rotor at rest, and we will assume that we have a wound rotor with balanced 3-phase windings that for the moment are open-circuited, i.e. that no current can flow in them. With balanced 3-phase stator currents of amplitude I_s we know from the discussion above that the travelling stator m.m.f. wave can be represented by a single space phasor that rotates at the synchronous speed, and that in the absence of any currents in the rotor (and neglecting saturation of the iron) the flux wave will be in phase with the m.m.f. and proportional to it. This is shown Fig. 8.6A: in this sketch the rotor and stator are stationary, but all the patterns rotate at synchronous speed.

On the left is a graphical representation of the sinusoidal distribution of resultant current around the stator at a given instant, and the corresponding flux pattern (dotted lines). Note that there is no rotor current. On the right of Fig. 8.6A is a phasor that can represent both the stator m.m.f. and what

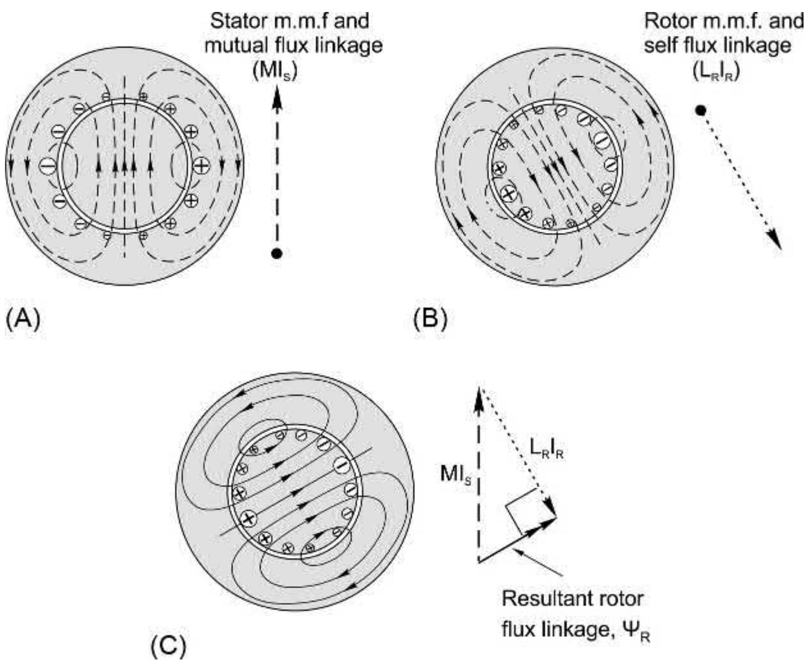

FIG. 8.6 Space phasors of m.m.f. and flux linkage under locked-rotor conditions.

we will call the resultant mutual flux linkage, both of which are proportional to the stator current. The expression 'mutual flux linkage' in Fig. 8.6A thus represents the total effective flux linkages with the rotor due to the stator travelling flux wave. In circuit terms, this flux linkage is proportional to the mutual inductance between the stator and rotor windings (M), and to the stator current (I_s), i.e. MI_s.

Now we short-circuit the rotor windings, and solve the set of equations for the rotor currents in the steady state. In view of the symmetry it comes as no surprise to find that they also form a balanced 3-phase set, at the same frequency as those of the stator, but displaced in time-phase. The resultant pattern of currents in the rotor is shown on the left in Fig. 8.6B, together with the flux pattern (dotted lines) that they would set up if they acted alone. Note that the stator currents that are responsible for inducing the rotor currents have been deliberately suppressed in Fig. 8.6B, because we want to highlight the rotor's reaction separately.

The m.m.f. due to the rotor currents is represented by the phasor shown on the right in Fig. 8.6B, and again this can also serve to represent the rotor self flux linkages (L_RI_R) attributable to the induced currents. It is clear that the time phase shift between stator and rotor currents causes a space phase shift between stator and rotor m.m.f.'s, with the rotor m.m.f. broadly tending to oppose the stator m.m.f.. If the rotor had zero resistance, the rotor m.m.f. would directly oppose that of the stator. The finite rotor resistance displaces the angle, as shown in Fig. 8.6B. We will see shortly that this phase angle varies widely and is determined by the frequency.

To find the resultant m.m.f. acting on the rotor we simply add the stator and rotor m.m.f. vectors, as shown in Fig. 8.6C. It is this m.m.f. that produces the resultant flux at the rotor, and we can therefore also use it to represent the net rotor flux linkage (ψ_R). The flux pattern at the rotor is shown by the solid lines in Fig. 8.6C. (But we should note that the number of flux lines shown in Fig. 8.6 are not intended to reflect the relative magnitudes of flux densities, which, if saturation was not present, would be higher in the two upper sketches.) We should also note that, as expected, the behaviour is independent of the rotor position angle θ, because the rotor symmetry means that when viewed from the stator, the rotor always looks the same overall. This is a feature that makes life easier when we come to look at the practical implementation of field-oriented control of induction motors.

Close examination of the lower sketch reveals an extremely important fact. The resultant rotor flux vector (ψ_R) is perpendicular to the rotor current vector. This means that the rotor current wave (shown in the left hand sketch) is oriented in the ideal position in space to maximise the torque production, because the largest current is coincident with the maximum flux density. If we look back to Figs. 3.1 and 3.2, we will see that this is exactly how the flux and current are disposed in the d.c. machine, the N pole facing the positive currents and the S pole opposite the negative currents.

When we evaluate the torque under these conditions, a very simple analytical result emerges: the torque turns out to be given by the product of the rotor flux linkage and the rotor current, i.e.

$$T = \psi_R I_R$$

The similarity of this expression and the torque expression for a d.c. machine is self-evident, and further underlines the fundamental unity of machines exploiting the '*BIl*' mechanism discussed in Chapter 1.

We should also note that in Fig. 8.6C one side of the right-angle triangle is ψ_R, while the adjacent side is proportional to the rotor current, I_R. Hence the area of the triangle is proportional to the torque, which provides an easy visualisation of how torque varies with frequency, which we look at shortly. (When we reach synchronous machines in Chapter 9, we will encounter a similar (though not right-angled) torque triangle, with adjacent sides proportional to the stator and rotor currents. In that case we will find it more useful to express the area of the triangle (i.e. torque) in terms of the product of the two currents and the sine of the angle between them.)

As an aside, the keen reader may recall that the mental pictures we employed in Chapter 5 were based on the calculation of torque from the product of the air-gap flux wave and the rotor current wave, and that these were not in phase, except at very low slip frequency. The much simpler picture which has now been revealed—in which the flux and current waves are always ideally disposed as far as torque production is concerned—arises because we have chosen to focus on the resultant rotor flux linkage, not the air-gap flux: we are discussing the same mechanism as in Chapter 5, but the new viewpoint has thrown up a much simpler picture of torque production.

We will see later that the rotor flux linkage is the central player in field-oriented (and direct torque) control methods that now dominate in inverter-fed drives. To make full use of the flux-carrying capacity of the rotor iron, and to achieve step changes in torque, we will keep the amplitude of ψ_R constant, and we will explore this shortly. But first we will look at how the torque depends on slip when the amplitude of the stator current is kept constant.

8.4.1 Torque vs slip frequency—Constant stator current

An alert reader might question why the title of this section includes reference to slip frequency, when we have specified so far that the rotor is stationary, in which case the effective slip is 1 and the frequency induced in the rotor will always be the same as the stator frequency. The reason for referring to slip frequency is that, as far as the reaction of the rotor is concerned, the only thing that matters is the relative speed of the travelling stator field with respect to the rotor.

So if we study the static model with an induced rotor frequency of 2 Hz, the torque that we predict can represent locked rotor conditions with 2 Hz on the

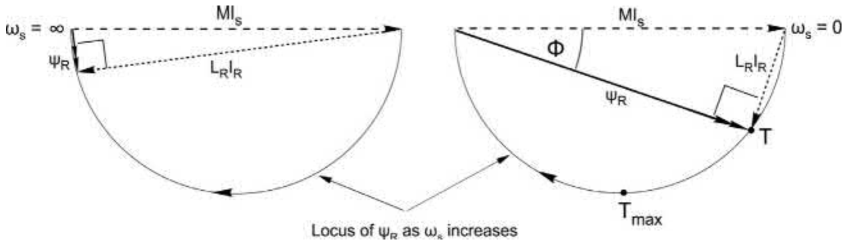

FIG. 8.7 Locus of rotor flux linkage space phasor as slip frequency varies.

stator; or the rotor running with a slip of 0.1 with 20 Hz on the stator; or a slip of 0.04 with 50 Hz on the stator, and so on. In short, under current-fed conditions, our model correctly predicts the rotor behaviour, including the torque, when we supply the stator windings at the slip frequency. (Note that all other aspects of behaviour on the stator side are not represented in this model, notably the fact that the stator voltage has to vary appropriately in order to keep the current constant.)

The variation in the flux linkage triangle with slip frequency, assuming that the amplitude of the stator current is constant, is shown in Fig. 8.7. The locus of the resultant rotor flux linkage as the slip is varied is shown by the semi-circles; the rotor current increases progressively with slip, but at the same time it moves out of space phase with the stator phasor.

The right hand side relates to low values of slip frequency, where the rotor self flux linkage is much less than the stator mutual flux linkage, so the resultant rotor flux linkage (ψ_R) is not much less than when the slip is zero. In other words, at low slips the presence of the rotor currents has little effect on the magnitude of the resultant flux, as we saw in Chapter 5. Low-slip operation is the norm in controlled drives.

The left hand figure relates to high values of slip, where the large induced currents in the rotor lead to a rotor m.m.f. that is almost able to wipe out the stator m.m.f., leaving a very small resultant flux in the rotor. We will not be concerned with this end of the diagram in an inverter drive.

There is a simple formula for the angle ϕ, which is given by

$$\tan\phi = \omega_S\tau \tag{8.1}$$

where $\tau = \frac{L_R}{R_R}$, the rotor time-constant.

We noted earlier that the torque is proportional to the area of the triangle, so it should be clear that the peak torque is reached when the slip increases from the point T and moves to $T_{max.}$ At this point, $\phi = \pi/4$ and the slip frequency is given by $\omega_S = \frac{1}{\tau} = \frac{R_R}{L_R}$. Under these constant-current conditions, the slip at which maximum torque occurs is much less than under constant-voltage conditions, because the rotor self inductance is much larger than the rotor leakage inductance.

8.4.2 Torque vs slip frequency—Constant rotor flux linkage

As already mentioned, it is clear that to make full use of the flux-carrying capacity of the rotor iron, we will want to keep the amplitude of the rotor flux ψ_R constant. Given that the majority of the rotor flux links the stator (see Fig. 8.6C), keeping the rotor flux constant also means that for most operating conditions, the stator flux is also more or less constant, as we assumed in Chapter 5.

In this section we explore how steady-state torque varies with slip when the rotor flux is maintained constant: this is illuminating, but much more importantly it prepares us for the final section, which deals with how we are able to achieve precise control of torque even under dynamic conditions.

We can see from Fig. 8.7 that to keep the rotor flux constant we will have to increase the stator current with slip. This is illustrated graphically in Fig. 8.8, in which the rotor flux linkage ψ_R is shown vertically to make it easier to see that it remains constant. In the left hand sketch, the slip is very small, so the induced rotor current and the torque (which is proportional to the area of the triangle) are both small. The rotor flux is more or less in phase with the applied stator flux linkage because the 'opposing' influence of the rotor m.m.f. is small.

In the middle and right-hand diagrams the slip is progressively higher, so the induced rotor current is larger and the stator current has to increase in order to keep the rotor flux constant.

There is a simple analytical relationship that gives the value of stator current required to keep ψ_R constant as slip varies, but of more interest is the relationship between the induced rotor current and the slip. From Fig. 8.8, we can see that the tangent of the angle ϕ is given by

$$\tan\phi = \frac{L_R I_R}{\psi_R}$$

Combining this with Eq. (8.1) we find that the rotor current is given by

$$I_R = \left(\frac{\psi_R}{R_R}\right)\omega_{slip} \tag{8.2}$$

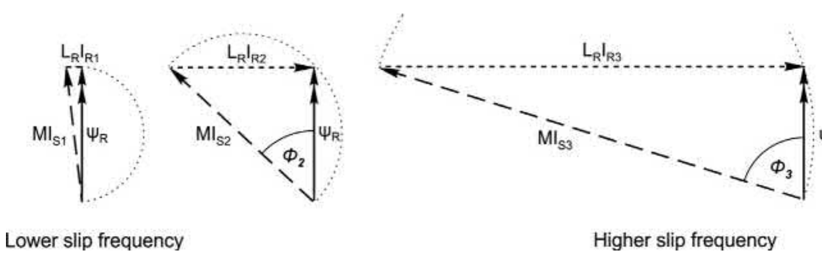

FIG. 8.8 Constant rotor flux linkage space phasors at low, medium and high values of slip, showing variation of stator current required to keep rotor flux constant.

The bracketed term is constant, therefore the rotor current is directly proportional to the slip. Hence the horizontal sides of the triangles in Fig. 8.8 are proportional to slip, and since the vertical side is constant, the area of each triangle (and torque) is also proportional to slip. To emphasise this simple relationship, the right-hand diagram in Fig. 8.8 has been drawn to correspond to a slip three times higher than that of the middle one, so the horizontal side of the right-hand sketch is three times as long, and the area of the triangle (and torque) is trebled.

We note that when the rotor flux is maintained constant, the torque-speed curve becomes identical to that of the d.c. motor. In this respect the behaviour differs markedly from that under both constant-voltage and constant-current conditions, where a peak or pull-out torque is reached at some value of slip. With constant rotor flux there is no theoretical limit to the torque, but in practice the maximum will be governed by thermal limits on the rotor and stator currents.

For those who prefer the physical viewpoint it is worth noting that the results discussed in this section could have been deduced directly from Fig. 8.6C, which indicates that the resultant rotor flux and rotor current waves are always aligned (i.e. the peak flux density coincides with the peak current density) so that if the flux is held constant, the torque is proportional to the rotor current. The rotor current is proportional to the motionally induced e.m.f., which in turn depends on the velocity of the flux wave relative to the rotor, i.e. the slip speed.

8.4.3 Flux and torque components of stator current

If we resolve the stator flux-linkage phasor MI_s into its components, parallel and perpendicular to the rotor flux, the significance of the terms ‘flux component’ and ‘torque component’ of the stator current becomes obvious (Fig. 8.9).

We can view the ‘flux’ component as being responsible for setting up the rotor flux, and this is the component that we must keep constant in order to maintain the working flux of the machine at a constant value for all slips. It is clearly analogous to the field current that sets up the flux in a d.c. motor.

The other (‘torque’) component (which is proportional to the rotor current) can be thought of as being responsible for nullifying the opposing effect of the

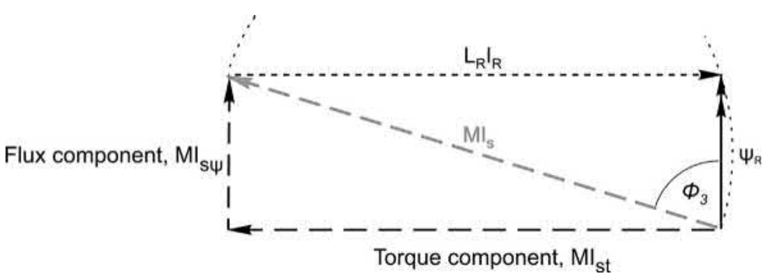

FIG. 8.9 Flux and Torque components of stator current.

rotor current that results when the rotor conductors are 'cut' by the travelling flux wave. This current component is therefore seen as the counterpart of the armature or work current in the d.c. motor.

Looking back to the left-hand diagram in Fig. 8.8, we see that at small slips (light load) the stator current is small and practically in phase with the flux; this is what we referred to as the magnetising current in previous chapters. At higher slips, the stator current is larger, reflecting that it now has a torque or 'work' component in addition to its magnetising component, which again accords with our findings in previous chapters.

8.5 Dynamic torque control

'Field-oriented control allows us to obtain (almost) instantaneous (step) changes in torque on demand, and it does this by jumping directly from one steady-state condition to another'. This simple statement is seldom given the prominence it deserves, but it is a simple truth, to be recalled whenever there is a danger of being bamboozled by a surfeit of technospeak.

If we want to obtain a step increase in torque, we have to change the rotor flux or the rotor current instantaneously, so as to jump instantaneously from one steady-state operating condition to another. But we have stressed many times that because a magnetic field has stored energy associated with it, it is not possible to change flux linkage instantaneously. In the case of the induction motor, any change in the rotor flux is governed by the rotor time-constant, which will be as much as 0.25 s for even a modest motor of a few kW rating, and much longer for large motors. This is not acceptable when we are seeking instantaneous changes in torque.

The alternative is to keep the flux constant, and make the rotor current change as quickly as possible: this is how dynamic control of torque is achieved in field-oriented systems, and the means whereby rapid changes in rotor current is achieved is discussed in the following section.

8.5.1 Special property of closely-coupled circuits

In Chapter 5 we saw that, when looked at from its terminals, the induction motor under steady-state conditions always looks—to a greater or lesser extent—inductive. At no-load, for example, the (magnetising) current was relatively small, but it lagged the voltage by almost 90°, so that the motor looked more or less like a high inductance. But as the load torque and slip increased, a load component of stator current came into play, so that the total current became much larger and moved closer in phase with the voltage, making the motor overall look more like a resistor, but still with a significant inductive element.

On the face of it therefore, the presence of stator inductance makes the prospect of making rapid changes to the rotor current look daunting. However, the important difference is that we are no longer talking about the steady-state (i.e.

when all the initial transients have settled down) but instead we must focus on how the stator-rotor combination reacts immediately after sudden changes in the voltage applied to the stator. The full treatment is complex and beyond our scope, but fortunately we can illustrate the essence of the matter by looking at the behaviour of a pair of coupled coils, one representing the stator winding and the other notionally representing the short-circuited cage rotor winding.

We will begin with a recap of the relationship between voltage and current in a pure inductor, the aim being to emphasise the difficulty of making a rapid change in the current. The differential equation linking voltage and current for a pure inductance is

$$v = L\frac{di}{dt}, \ or\ di = \frac{1}{L}(v\,dt)$$

It follows that unlike in a resistor, where the current follows the voltage immediately, in an inductor the current is determined by the time-integral of the applied voltage, so that to increase the current by *di* we have to apply a fixed volt-second product *vdt*.

To increase the current (and hence increase the stored energy in the inductor) we can choose to apply a high voltage (and high power) for a short time or a lower voltage for longer, but, whatever the inductance, we can never obtain a perfect step change in current because that would require an impulse of infinite voltage and zero duration, i.e. a pulse of infinite power.

For example, Fig. 8.10 shows the result of applying successive voltage-time pulses of the same area (shown shaded) but of voltages V, 2 V and 4 V, respectively, to a pure inductance, L. Each pulse raises the current by the same amount (VT/L), but the rate of rise of the current increases in direct proportion to the voltage. (Finally, the current is brought back down to zero by applying a

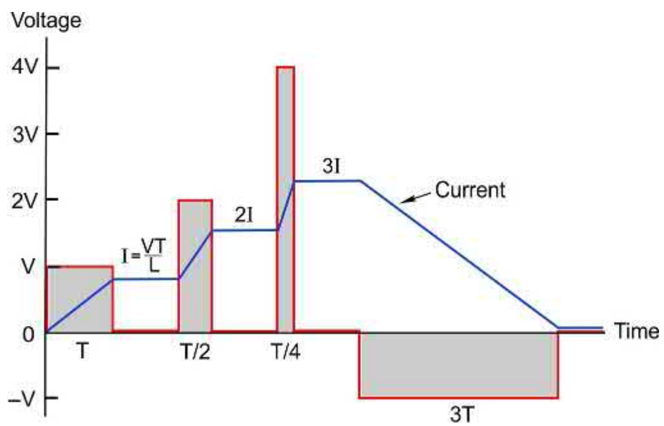

FIG. 8.10 Current (*blue*) when step voltages (*red*) are applied to a pure inductor.

negative pulse of area 3VT.) It should be clear that to effect a rapid change in current, the higher the voltage and the lower the inductance, the better.

Some readers may feel uneasy when observing Fig. 8.10 to note that the current remains constant when the voltage is zero, but they will be reassured when reminded that this is only true when the inductance is ideal, i.e. it has zero resistance. In practice there will always be resistance, and then, as soon as the voltage is reduced to zero, the current will begin to fall with a time-constant of L/R. The reason that we neglect resistance is not only that the analysis is very much simpler, but also we are interested in what happens immediately after a step change in voltage, in which case the effect of resistance is negligible, all the voltage then appearing across the inductive element.

Having looked at how a single inductor behaves, we return to the matter of what the stator 'looks like' to the supply under transient conditions, by considering a pair of magnetically coupled coils, one representing the stator and the other the rotor. We will assume that the coils are stationary (i.e. the rotor is at rest), but the arguments are equally valid when the rotor is moving because at every instant the cage rotor looks the same when viewed from the stator.

Let the self inductances of the coils representing stator and rotor be L_s and L_r, respectively, let their mutual inductance be M, and ignore their (very low) resistances. Because of the good magnetic circuit, all three inductances will be large. If we were able to open-circuit the rotor cage, the stator would therefore look like a high inductance (and its reactance at the utility supply frequency would correspond to what we earlier called the magnetising reactance). However, the rotor circuits are short-circuited, so we model this by setting the rotor voltage term to zero in the Kirchoff's voltage equations:-

$$v_s = L_s\frac{di_s}{dt} + M\frac{di_r}{dt}$$

$$v_r = 0 = L_r\frac{di_r}{dt} + M\frac{di_s}{dt}$$

Eliminating the rotor current we obtain

$$v_s = L_s\left(1 - \frac{M^2}{L_sL_r}\right)\frac{di_s}{dt}$$

The effective inductance at the stator is therefore given by

$$L_s' = L_s\left(1 - k^2\right),$$

where k is the coupling coefficient, defined as

$$k = \frac{M}{\sqrt{L_sL_r}}$$

The coupling coefficient always lies between 0 (no mutual flux) and 1 (no leakage flux), so we see the very welcome news that the inductance looking in at the stator is reduced from its open-circuit value by the factor $1-k^2$, and that if the coils were perfectly coupled (i.e. $k=1$), the effective inductance becomes zero.

At first sight this is a very unexpected (and welcome) result given that both windings taken separately have high self inductances. (But for those familiar with the idea of referred impedance in the context of a transformer, it is perhaps not surprising!)

In an induction motor the good magnetic circuit means that the coupling coefficient is high, perhaps 0.95, in which case the effective transient inductance looking in at the stator is barely 10% of its self-inductance (and it corresponds loosely to the so-called leakage inductance). It is this remarkably fortunate outcome that allows us to achieve rapid changes in the stator currents without requiring extremely high voltages from the inverter.

It remains for us to see how the rotor currents react when the stator current is changed so, knowing that we can come close to effecting a sudden change in practice, we will assume that the current in the stator suddenly increases from zero to I_s. Under this condition, it emerges that the rotor current also suddenly increases, its magnitude being given by

$$I_r = -\left(\frac{M}{L_r}\right) I_s$$

Rearranging this expression to highlight the flux linkage produced by the rotor yields

$$L_r I_r = -M I_s \,.$$

Hence the rotor reacts by producing self flux linkages that exactly cancel those reaching the rotor from the stator, leaving the rotor flux linkage at its initial value of zero. As mentioned earlier, in practice the rotor current will decay because of its finite resistance, but this will not be significant over the short timescale that we will be interested in for torque control.

For those who prefer to argue from a physical basis, we can see why the effective inductance is reduced by the closely-coupled and short-circuited secondary coil by invoking Faraday's and Lenz's laws. If the primary current changes, the flux linking the secondary also changes, inducing an e.m.f. in the secondary. The secondary is short-circuited, so a secondary current is induced in a direction that makes its associated flux oppose the changing flux that was responsible for producing the e.m.f.. If the circuits are perfectly coupled via zero-reluctance paths, the opposing flux completely neutralises the primary flux, and there is never any resultant flux, and the effective primary inductance is therefore zero. In the non-ideal case, cancellation is incomplete, but the effective primary inductance is always reduced.

To conclude this section, it is appropriate to say that in order to cover a potentially tricky matter we have deliberately made use of a very simple model, and we have been somewhat loose in suggesting that the stator and rotor can be represented by single coils. Our experience suggests that non specialist readers will accept such gross simplifications where the main message emerges without too much difficulty, and we will make use of similar arguments in the next section, that deals with establishing the rotor flux.

8.5.2 Establishing the flux

In Section 8.4 we assumed that steady-state conditions prevailed, with the rotor flux linkage remaining of constant magnitude and rotating relative to the rotor at the slip speed. We now look at how the rotor flux wave was first established.

We start with the rotor at rest, no current in any of the windings, and hence no flux. With reference to Fig. 8.1, we suppose that we supply a step (d.c.) current into phase U, which will split with half exiting from each of phases V and W, and producing a stationary sinusoidally distributed m.m.f. that, ultimately, will produce the flux pattern labelled 'final state' in Fig. 8.11.

But of course the rotor windings are short-circuited, with no flux through them, and as we showed in the previous section, closed electrical circuits behave like many things in the physical world in that they react to change by opposing

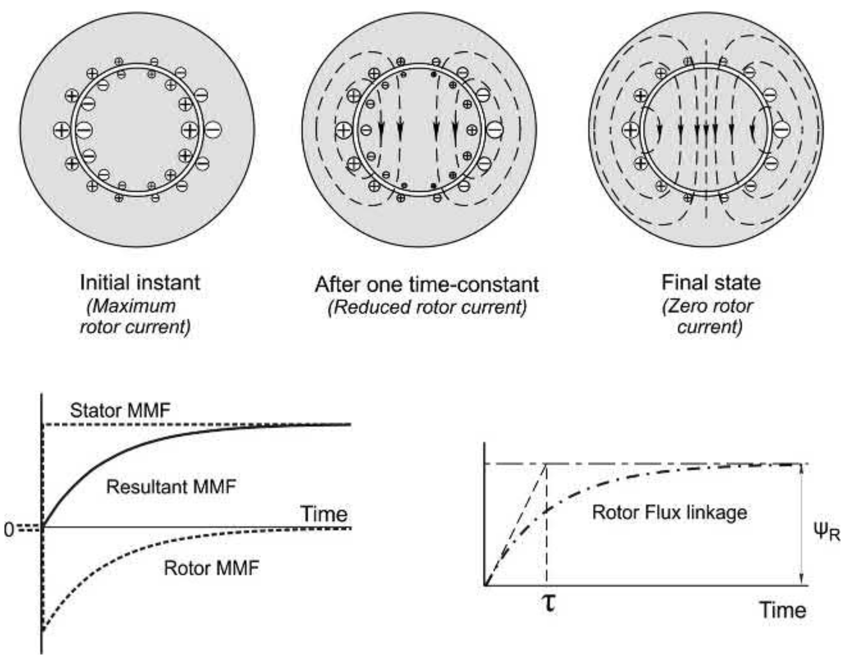

FIG. 8.11 Diagrams illustrating the build-up of rotor flux when a step of stator m.m.f. occurs.

it. In this context, if the flux linking a winding changes, there will be an induced e.m.f. that produces a current, the m.m.f. of which will be in opposition to the 'incoming' m.m.f..

So, when the stator m.m.f. phasor suddenly comes into existence, the immediate reaction of the rotor is the production of a negative stationary rotor m.m.f. pattern, i.e. in direct opposition to the stator m.m.f.: this is labelled 'Initial instant' in Fig. 8.11. Instantaneously, the magnitude of the rotor m.m.f. is such as to keep the rotor flux linkage at zero, as it was previously. This outcome agrees with what we found from a study of two coupled coils.

However, because of the rotor resistance, the rotor current needs a voltage to sustain it, and the voltage can only be induced if the flux changes. So the rotor flux begins to increase, rising rapidly at first (high e.m.f.) then with ever-decreasing gradient leading to lower current and lower rotor m.m.f. The response is a first-order one, governed by the rotor time-constant, so after one time constant (middle sketch) the flux linking the rotor reaches about 63% of its final value, while the rotor current has fallen to 37% of its initial value. Finally the rotor's struggle to prevent the flux changing comes to an end and the rotor flux linkage reaches a steady value determined by the stator current. If the resultant rotor flux linkage is the target value for steady-state running (ψ_R), the corresponding stator current is what we previously referred to as the 'flux component'.

The physical reason why it takes time to build the flux is that energy is stored in a magnetic field, so we cannot suddenly produce a field because that would require an impulse of infinite power. If we want to build up the flux more rapidly, we can put in a bigger step of stator current at first, so that the flux heads for a higher final value than we really need, then reduce the stator current when we get close to the flux we are seeking.

We began this section with d.c. current in the stator, which in effect corresponds to zero slip frequency, all the field patterns being stationary in space. Because there is no relative motion involved, there is no motional e.m.f. and hence no torque. The 'torque component' only comes into play when there is relative motion between the rotor and the rotor flux wave, i.e. when there is slip. Obviously, to cause rotation the frequency must be increased, and as we have seen in the previous section the stator current then has to be adjusted with slip and torque to keep the rotor flux linkage constant.

Finally, it is worth revisiting Fig. 8.9 briefly to reconcile what we have discussed in this section with our picture of steady-state operation, where the rotor currents are at slip frequency. On the left we have the fictitious 'flux component' of stator current, which remains constant in magnitude and aligned with the rotor flux linkage, $\psi_{R,}$ along the so-called 'direct' axis. When we first established this flux, the rotor reacted as we have discussed above, but after a few time-constants the flux settled to a constant value along the direct axis in the direction of the flux component of the stator m.m.f.. This is why the arrows on ψ_R point in the same direction as the stator flux producing component.

However, we note from Fig. 8.9 that the rotor flux linkage phasor ($L_R I_R$) is always equal in magnitude to the torque component of the stator mutual flux linkage phasor, but, as shown by the arrows, it is in the opposite direction. There is therefore no resultant m.m.f. or tendency for flux to develop along this, the so-called 'quadrature' axis. This is what we would expect in the light of the previous discussion, where we saw that the reaction of mutually coupled windings to any suggestion of change is for currents to spring up so as to oppose the change. In the literature when, as here, the 'torque' current does not affect the flux, the axes are said to be 'decoupled'.

8.5.3 Mechanism of torque control

In the previous section, our aim was to grow the rotor flux, which, because of its stored energy, took a while to reach the steady state. However if, subsequently, we keep the rotor flux linkage constant (by ensuring that the flux component of the stator current is constant and aligned with the flux) we can cause sudden changes to the motionally-induced rotor current by making sudden changes in the torque component of the stator current.

We achieve sudden step changes in the stator currents by means of a fast-acting closed-loop current controller. Fortunately, we have seen that under transient conditions the effective inductance looking in at the stator is quite small (it is equal to the leakage inductance), so it is possible to obtain very rapid changes in the stator currents by applying high, short-duration impulsive voltages to the stator windings. In this respect the stator current controller closely resembles the armature current controller used in the d.c. drive.

When a step change in torque is required the magnitude, frequency, and phase of the stator currents are changed (almost) instantaneously in such a way that the rotor current jumps suddenly from one steady-state to another. But in this transition it is only the torque component of stator current that is changed, leaving the flux component aligned with the rotor flux. There is therefore no change in the magnitude of the rotor flux wave and no change in the stored energy in the field, so the change can be accomplished almost instantaneously.

We can picture what happens by asking what we would see if we were able to observe the stator m.m.f. wave at the instant that a step increase in torque was demanded. For the sake of simplicity, we will assume that the rotor speed remains constant, and consider an increase in torque by a factor of three (as between the middle and right-hand sketches in Fig. 8.8), in which case we would find that:

(a) the stator m.m.f. wave suddenly increases its amplitude;
(b) the frequency of the stator m.m.f. wave suddenly increases, defining a new synchronous speed such that the slip increases by a factor of three (as the rotor speed has not changed);

(c) the stator m.m.f. wave jumps forward to retain its correct relative phase with respect to the rotor flux, i.e. the angle between the stator m.m.f. and the rotor flux increases from ϕ_2 to ϕ_3 as shown in Fig. 8.8.

Thereafter the stator m.m.f. retains its new amplitude, and rotates at its new speed.

The rotor current experiences a step change from a steady state at its initial slip frequency to a new steady state with three times the amplitude and frequency, and there is a step increase in torque by a factor of three, as shown in Fig. 8.12A. The new current is maintained by the new (higher) stator currents and slip frequency.

We should note particularly that it is the jump in the *angular position (*i.e. *space phase angle)* that accompanies the step changes in magnitude and frequency of the stator m.m.f. phasor that allows the very rapid and transient-free control of torque. Given that the definition of a vector quantity is one which has magnitude and direction, and that the angular position of the phasor defines the direction in which it is pointing, it is clear why this technique is sometimes known by the name *'vector control'*.[1]

To underline the importance of the sudden change in *phase* of the stator current (i.e. the sudden jump in angular position of the stator m.m.f. in achieving a step of torque), Fig. 8.12B shows what happens typically if only the magnitude and frequency, but not the position of the stator m.m.f. phasor are suddenly changed. The steady state conditions are ultimately reached, but only after

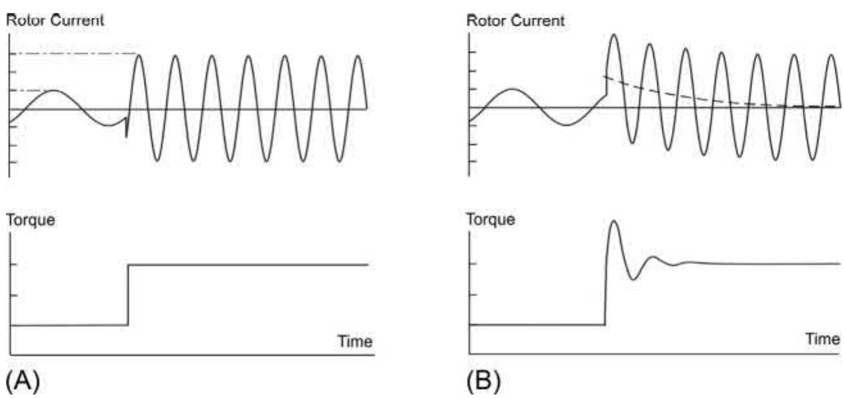

FIG. 8.12 Step changes in rotor current. (A) Transient-free transition with correct changes to magnitude, frequency and instantaneous position of stator m.m.f. wave (i.e. vector control). (B) Same changes to magnitude and frequency, but not phase.

1. The term vector control has sometimes been misused to refer to drives that do not include field-orientation. However the term is so ubiquitous that we cannot avoid it, so when we refer to 'vector control' we mean a proper field-oriented system.

an undesirable transient governed by the (long) rotor time-constant, which may persist for several cycles at the slip frequency. The fundamental reason for the transient is that if the magnitude of the stator current is suddenly increased without a change of position, the flux and torque components both increase proportionately. The change in the flux component portends a change in the rotor flux (and associated stored energy), which in turn is resisted by induced rotor currents until they decay and the steady-state is reached.

This section has described the underlying principles by which very rapid and precise torque control can be achieved from an induction motor, but we should remember that until sophisticated power electronic control became possible the approach outlined here was only of academic interest. The fact that the modern inverter-fed drive is able to implement torque control and achieve such outstandingly impressive performance from a motor whose inherent transient behaviour is poor, represents a major milestone in the already impressive history of the induction motor. The way in which such drives achieve field oriented control is discussed next.

8.6 Implementation of field-oriented control

As explained in Chapter 7, early inverter-fed systems used an internal oscillator to determine the frequency supplied from the inverter to the motor, the latter being left to its own devices to react to any changes in the inverter output voltage waveform. In particular, the motor current was free to do what comes naturally, so that, as we mentioned in Chapter 5, the inherent transient response is poor, and adjustments to changes in frequency, for example, typically take several cycles to reach the new steady state, with unwanted fluctuations in torque over which we have no control.

In complete contrast, with field-oriented control, the motor flux and current are continuously monitored, and rapid current control is employed to ensure that the instantaneous torque follows the demanded value. The switching of the devices in the inverter is thus determined by what is happening inside the motor, rather than being imposed by an external oscillator. So it will probably be good to remind the reader that a non-trivial change of mind-set is called for at this point to reflect the radically different inverter control philosophy which we are now beginning to examine.

An essential requirement if we are to unravel the workings of the overall scheme for field-oriented control is an understanding of the PWM vector modulator/inverter combination that is a feature of all such schemes, so this is covered first.

8.6.1 PWM controller/vector modulator

In the inverters we have looked at so far (see Section 2.4) we have supposed that the periodic switching required to approximate a sinusoidal output was provided from a master oscillator. The frequency of the oscillator determined the frequency of the a.c. voltage applied to the motor, and the amplitude was controlled separately. In terms of space phasors this allows control of the amplitude and frequency, but not the instantaneous angular position of the voltage and current phasors. As we have seen, it is the additional ability to make instantaneous changes to the *angular position* of the output phasor that is the key to dynamic torque control, and this is the key feature provided by the 'vector modulator'.

We now explore what the inverter can produce in terms of its output voltage phasor. We recall that there are six devices (switches) in three legs (see Fig. 2.21), and to avoid a short-circuit across the inverter d.c. link both switches in one leg must not be turned on at the same time. If we make the further restriction that each phase winding must at all times be connected to one or other of the d.c. link terminals, there are only eight possible combinations, as shown in Fig. 8.13.

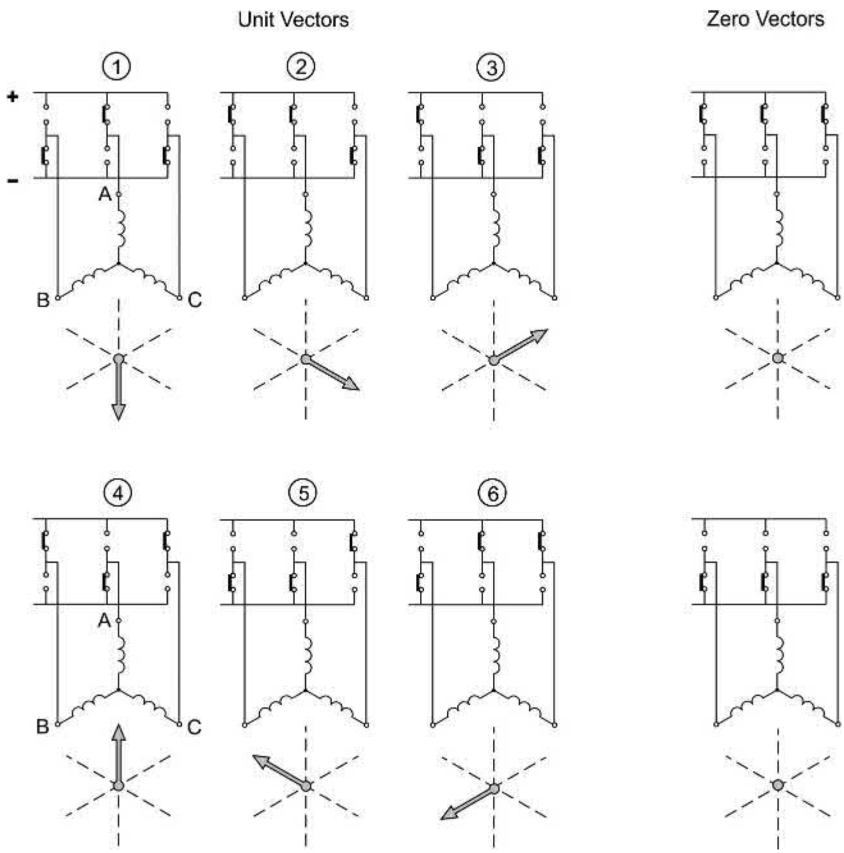

FIG. 8.13 Voltage phasors for all acceptable combinations of switching for a 3-phase inverter.

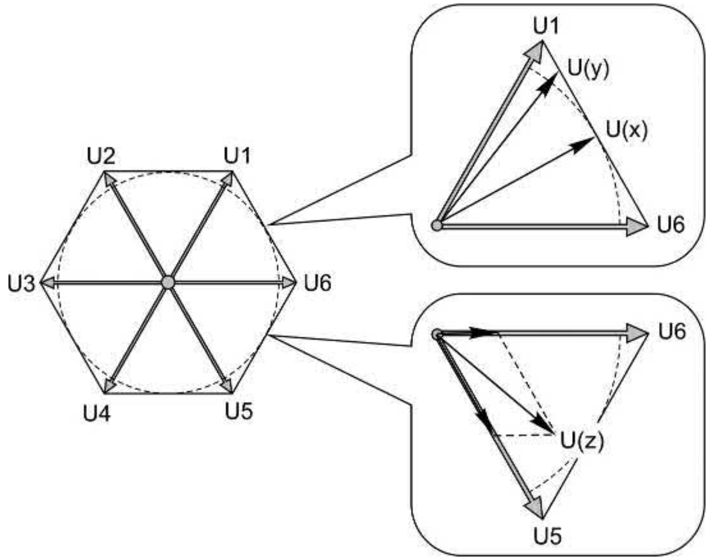

FIG. 8.14 Synthesis of intermediate voltage phasors in vector modulator.

The six switching combinations labelled 1–6 each produce an output voltage phasor of equal amplitude but displaced in phase by 60° as shown in the lower part of each sketch, while the final two combinations have all three terminals joined together, so the line voltage is zero. The six unit vectors are shown with their correct relative phase, but rotated so as to bring U6 horizontal, in Fig. 8.14.

Having only six states of the voltage phasor at our disposal is clearly not satisfactory, because we need to exert precise control over the magnitude and position of the voltage phasor at any instant, so this is where the 'time modulation' aspect comes into play. For example, if we switch rapidly between states U1 and U6, spending the same time with each, we will effectively have synthesised a voltage phasor lying half way between them, and of magnitude U1cos 30° (or 86.7% of U1), as shown by the vector U(x) in the upper part of Fig. 8.14. If we spend a higher proportion of the time on U1 and the remainder on U2, we could produce the vector U(y). As long as we spend the whole of the sample time on either U1 or U6, we will end up somewhere along the line joining U1 to U6.

We have used the terms 'switch rapidly' and 'the time' without specifying what they mean. In practice, we would expect the switching or modulating frequency to be perhaps a few kHz up to the low tens of kHz, so 'the time' means one cycle at this frequency, say 100 microseconds at 10 kHz. So for as long as we wished the voltage phasor to remain at U(x), we would spend 50 μs of each sample period alternately connected to U1 and U6.

Recalling that ideally we want to be able to choose both magnitude and position it is clearly not satisfactory to be constrained to the outer edges of the hexagon. So now we bring the zero vector into play. For example, suppose we wish

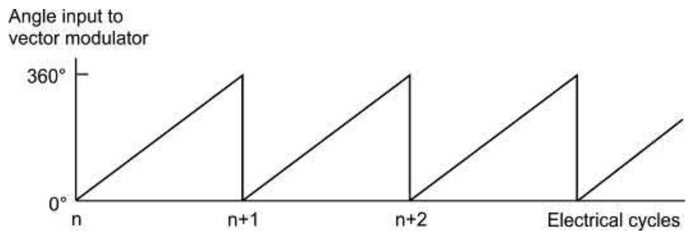

FIG. 8.15 Angle reference to vector modulator corresponding to constant frequency operation of inverter. (Note that the angle resets to zero at the end of each cycle.)

the voltage phasor to be U(z), as in the lower sketch. This is composed of (0.5) U5 plus (0.3)U6. Hence in each modulating cycle of 100 μs, we will spend times of 50 μs on U5, 30 μs on U6 and 20 μs with one of the zero states.

The precise way in which these periods are divided within one cycle of the modulating frequency is a matter of important detail in relation to the distribution and minimisation of losses between the six switching devices, but need not concern us here. Suffice it to say that it is a straightforward matter to arrange for digital software/hardware that has input signals representing the magnitude and instantaneous position of the output voltage phasor, and which selects and modulates the six switches appropriately to create the desired output until told to move to a new location.

When we introduced the idea of space phasors earlier in this chapter, we saw that if we begin with balanced three-phase sinusoidal voltages, the voltage phasor is of constant length and rotates at a uniform rate. Looking at it the other way round, it should be clear that if we arrange for the output of the inverter to be a voltage phasor of constant length, rotating at a constant rate, then the corresponding phase voltages must form a balanced sinusoidal set, which is what we want for steady state running.

We conclude that in the steady state, the magnitude of the input signal to the vector modulator would have a constant amplitude and its angle would increase at a linear rate corresponding to the desired angular velocity of the output. Clearly in order to avoid having to deal with ever-increasing angles it will reset each time a full cycle of 360° is reached, as shown in Fig. 8.15.

Away from the steady-state condition, for example during acceleration, we should recall that to preserve the linear relation between torque and the stator current component (I_t), the flux component of the stator current phasor (I_F) must remain aligned with the rotor flux. As we will see in the next section, this is achieved by deriving the angle input to the vector modulator directly from the absolute angular position of the rotor flux.

8.6.2 Torque control scheme

A simplified block diagram of a typical field-oriented torque control system is shown in Fig. 8.16.

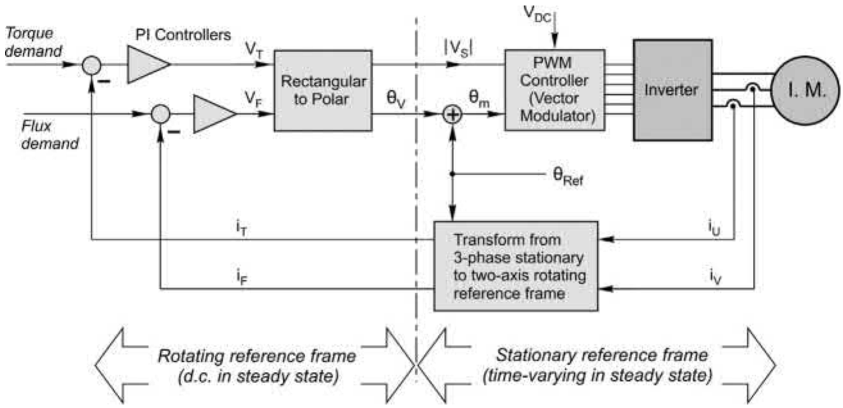

FIG. 8.16 Simplified block diagram of a typical field-oriented torque control system.

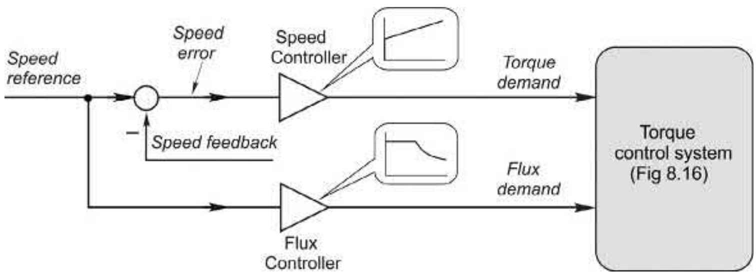

FIG. 8.17 Schematic diagram of closed-loop speed control system.

The first and most important fact to bear in mind in the discussion that follows is that Fig. 8.16 represents a *torque* control scheme, and that for applications that require speed control, it will form the 'inner loop' of a closed-loop speed control scheme. The torque and flux inputs will therefore be outputs from the speed controller, as indicated in Fig. 8.17. Note that the flux demand will be constant up to the base speed of the motor, and then falls. This is field-weakening operation (Section 7.3), as a consequence of the drive not being able to produce any more volts, and the V/f ratio therefore reducing. In Field oriented control, we still have the same practical limitations of maximum available voltage and so the flux must be reduced accordingly. As in any good control scheme, in order to facilitate good dynamic operation, it is necessary to keep a margin of volts in order to change the current quickly. Nonetheless, the transient performance in the field-weakening region remains impressive (though obviously not as good as at full flux).

Returning to Fig. 8.16 it has to be acknowledged that it looks rather daunting, and getting to grips with it is not for the faint-hearted. However, if we examine it a bit at a time, it should be possible to grasp the essential features of its operation. To simplify matters, we will focus on steady-state conditions, despite the fact that the real merit of the system lies in its ability to provide precise torque control even under transient conditions.

Taking the broad overview first, we can see that there are similarities with the d.c. drive with its inner current (torque) control loop (see Chapter 4), notably the stator current feedback and the use of proportional and integral (PI) controllers to control the torque and flux components of the stator current. It would be good if we could measure the flux and torque components directly, but of course the current components do not have separate existences: they are merely components of the stator current, which is what we can measure. The motor has three phases, but because we are assuming that there is no neutral connection, it suffices to measure only two of the line currents (because the sum of the three is zero). The information from these two currents allows us to keep track of the angular position of the stator current phasor with respect to the stationary reference frame (θ_S) as shown in Fig. 8.18.

However, the stator current feedback signals are alternating at the frequency supplied by the inverter, and the corresponding stator space phasor is rotating at the supply frequency with respect to a stationary reference frame. Before the flux and torque components of these signals (I_F and I_T) can be identified (and subsequently fed back to the PI controllers) they must first be transformed (see Section 8.2.2) into a reference frame that rotates with the rotor flux. As explained previously, the rotor flux angle θ_{Ref} is therefore an essential input to the transformation algorithm, as shown in Fig. 8.16.

The vertical dotted line in the middle of Fig. 8.16 separates quantities defined in the stationary reference frame (on the right) from those in the rotating reference frame (on the left). In the steady state, all those on the left are d.c., while all those on the right are time-varying.

The reader might wonder why, when we follow the signal path of the current control loops, beginning on the right with the phase current transducers, there is no matching 'inverse transform' to get us back from the rotating reference frame

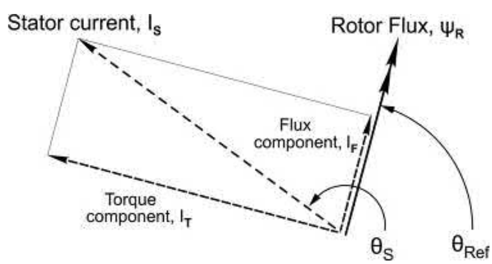

FIG. 8.18 Stator current and rotor flux reference angles.

on the left to the stationary reference frame on the right. The answer lies in the nature of the input signal to the PWM/vector modulator and inverter, which we discussed above. Let us suppose that the motor is running in the steady state, so that the output voltage phasor rotates at a constant rate with angular frequency ω. Under these conditions the rotor flux phasor also rotates with constant angular velocity ω, so the angle of the flux vector with respect to the stationary reference frame (θ_{Ref}) increases linearly with time. Also, in the steady state, the output from the PI controllers is constant, so the angle θ_V (Fig. 8.16) is constant. Hence the input angle to the modulator (θ_m in Fig. 8.16) which is the sum of θ_{Ref} and θ_V is also a ramp in time, and this is what provides the rotation of the output voltage phasor. In effect, the system is self-sustaining: the primary time-varying input angle to the modulator comes from the flux position signal (which is already in the stationary reference frame), and the PI controller provides the required magnitude signal ($|Vs|$) and the additional angle (θ_V).

Turning now to the action of the PI controllers, we see from Fig. 8.16 that the outputs are voltage commands in response to the differences between the feedback (actual) values of the transformed currents and their demanded values. The flux demand will usually be constant up to base speed, while the torque demand will usually be the output from the speed or position controller, as shown in Fig. 8.17. The proportional term gives an immediate response to an error, while the integral term ensures that the steady-state error is zero. The outputs from the two PI controllers (which are in the form of quadrature voltage demands, V_F and V_T) are then converted from rectangular to polar form, to produce amplitude and phase signals, $|V_s|$ and θ_V, where

$$|V_S| = \sqrt{V_F^2 + V_T^2}, \text{ and } \theta_V = \tan^{-1}\frac{V_T}{V_F},$$

as shown in Fig. 8.19.

The amplitude term specifies the magnitude of the output voltage phasor (and thus the three phase voltages applied to the motor), with any variation of the dc link voltage (V_{dc}) being compensated in the PWM controller. The phase angle (θ_V) represents the desired angle between the stator voltage phasor and the rotor flux phasor, both of which are measured in the stationary reference

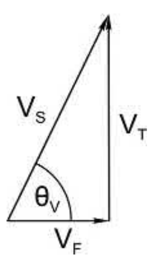

FIG. 8.19 Derivation of voltage phasor from flux and torque components.

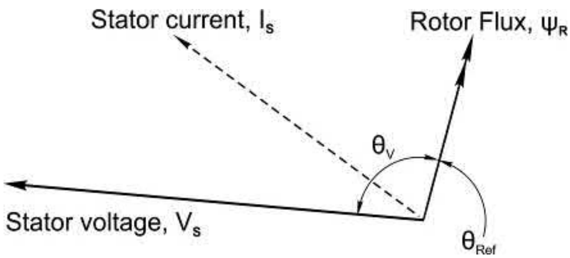

FIG. 8.20 Time phasor diagram showing stator voltage and current under steady-state conditions.

frame. The angle of the rotor flux phasor is θ_{Ref}, so θ_V is added at the input to the vector modulator to yield the stator voltage phasor angle, θ_m, as shown in Fig. 8.16.

We can usefully conclude our look at the steady state by adding the stator voltage phasor to Fig. 8.18 to produce Fig. 8.20, to provide reassurance that, in the steady state, the rather different approach we have taken in this section is consistent with the classical approach taken earlier.

8.6.3 Transient operation

We concluded earlier that for the motor torque to be directly proportional to the torque component of stator current, it is necessary to keep the magnitude of the rotor flux constant and to ensure that the flux component of stator current is aligned with the rotor flux. This is achieved automatically because the principal angle input to the vector modulator comes directly from the rotor flux angle (θ_{Ref}), as shown in Fig. 8.16. So during acceleration, for example, the instantaneous angular velocity of the rotor flux wave will remain in step with that of the stator current phasor, so that there is no possibility of the two waves falling out of synchronism with one another.

In Section 8.5.3 we discussed a specific example of how to obtain a step change in torque by making near-instantaneous changes to the magnitude, speed and position of the stator m.m.f. wave, and we are now in a position to see how this particular strategy is effected using the control scheme shown in Fig. 8.16.

A step demand for torque causes a step increase in $|V_s|$ and θ_V at the output of the rectangular to polar converter in order to effect a very rapid increase in the magnitude and instantaneous position of the stator current phasor. At the same time, the algorithm that calculates the slip velocity of the flux wave (see later, Eq. 8.3) yields a step increase because of the sudden increase in the torque component of stator current. The principal angular input to the vector modulator (the flux angle (θ_{Ref})) therefore changes gradient abruptly, as shown in Fig. 8.21.

Recalling that the steady-state stator frequency is governed by the angular velocity of the flux (i.e. $\frac{d\theta_{Ref}}{dt}$), this lines up with our expectation that (assuming the rotor velocity is constant) the stator frequency will increase in order to increase the slip and provide the new higher torque.

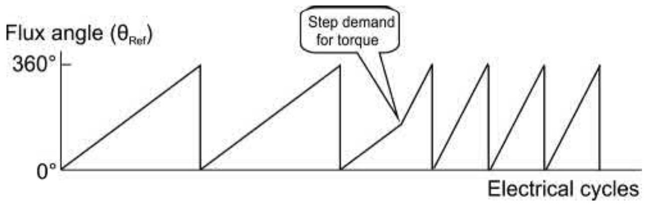

FIG. 8.21 Flux angle reference showing response to a sudden step demand for increased torque.

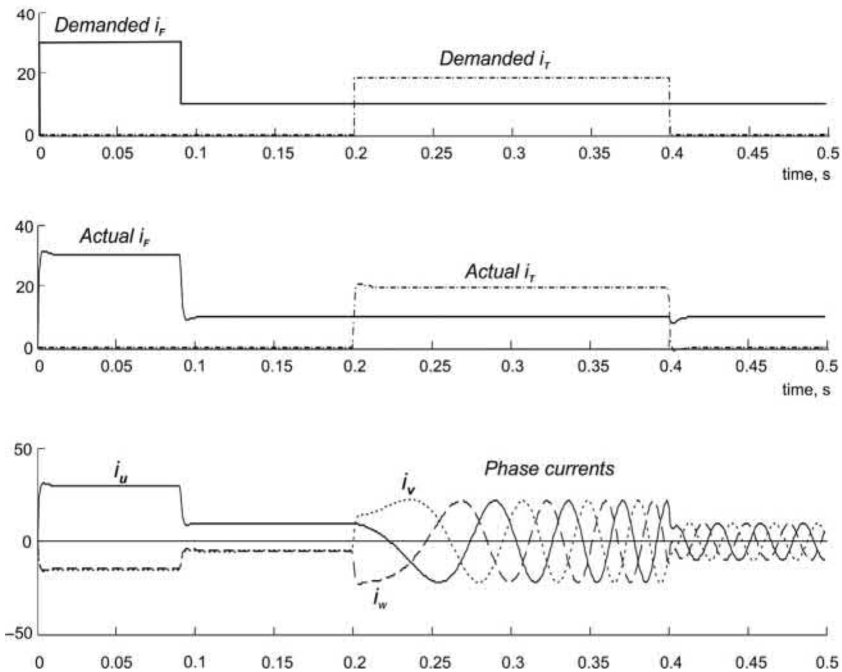

FIG. 8.22 Experimental results showing build-up of flux followed by sudden demand for step increase in torque until motor reaches its target speed. *(Courtesy of Nidec—Control Techniques.)*

8.6.4 Acceleration from rest

The practical results in Fig. 8.22 ably demonstrate the impressive performance of a field-oriented torque control system. These relate to a motor whose rotor time-constant is approximately 0.1 s, and covers an overall time period of 0.5 s. The motor is initially unexcited, rotor flux is then established and it is then accelerated up to a steady speed.

The upper diagram shows the demanded values for the transformed flux and torque components of stator current; the middle diagram shows the measured (actual) flux and torque components; and the lower shows the three phase currents.

This particular motor requires a stator current flux component (I_F) of 10 A to maintain full rotor flux linkage, but by applying an initial demand of 30 A, the initial rate of rise of the flux is trebled. After about one rotor time-constant the demand is reduced to 10 A to maintain the flux. Without this short-term 'forcing' it would have taken about five time-constants to establish the flux.

The measured value of the transformed flux component of stator current is shown in the middle plot, and it is seen to follow the demanded signal closely, with only slight overshoot. This signal is the transformed version of the actual stator winding currents, so the fact that it is on target demonstrates that the phase currents are established rapidly, and held while the flux builds up, as we can see in the lower figure. After 0.2 s phase U carries a positive d.c. current of 10 A while phases V and W each carry a negative d.c. current of 5 A. The rotor remains at rest, with full rotor flux now established, and at this time there is no demand for torque, so the motor remains stationary, the slip being zero.

At 0.2 s, a step demand signal (I_T) equivalent to a torque component of stator current of 20 A is applied, in order to accelerate the motor. The flux demand remains at 10 A during the acceleration, in order to keep the rotor flux linkage constant, thereby ensuring that torque is proportional to slip. The torque demand is maintained until 0.4 s, when the torque producing reference is reduced to zero, and the motor stops accelerating.

We note the almost immediate and transient-free transition of the three-phase currents from their initial steady (d.c.) values immediately prior to 0.2, into constant amplitude, 'smoothly increasing frequency' a.c. currents over the next 0.2 s. And then there is a similarly near-perfect transition to reduced amplitude steady-state conditions (at about 40 Hz) after 0.4 s. In the steady state, the torque component is negligible because the motor is unloaded, and the stator current consists only of the flux component, which traditionally would be referred to as the magnetising current.

During the acceleration the controller keeps the stator phase currents at the amplitude corresponding to the vector sum of the demanded flux and torque components (Fig. 8.19), and it continuously estimates the rotor flux position in order to keep the stator flux component aligned with the rotor flux. Hence while the rotor is accelerating, the instantaneous angular velocity of the rotor flux is greater than that of the rotor by an amount equal to the slip, as shown in the Fig. 8.23.

Younger readers will doubtless not require convincing of the validity of these remarkable results, but they might find it salutary to know that until the 1970's it was widely believed that such performance would never be possible.

To conclude this section we can draw a further parallel between the field-oriented induction motor and the d.c. motor. We see from Fig. 8.22 that as the motor accelerates, the frequency of the stator currents increases with the speed. If we stationed ourselves on the rotor of a d.c. motor as it accelerated,

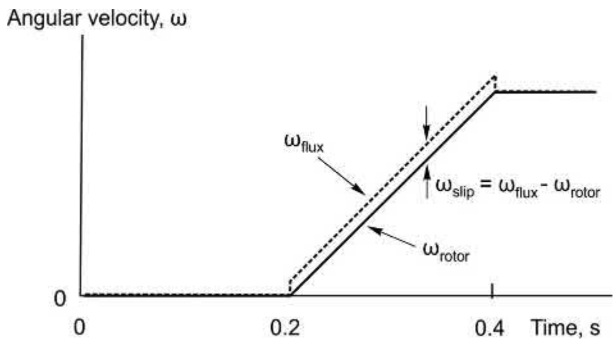

FIG. 8.23 Angular velocity during acceleration.

the rate at which the current in each rotor coil reversed as it was commutated would also increase in proportion to the speed, though of course we are not aware of it when we are in the stationary reference frame.

8.6.5 Deriving the rotor flux angle

By now, the key role played by the rotor flux angle should have become clear, so finally we look at how it is obtained. It is not practical nor economic to fit a flux sensor to the motor, so industrial control schemes invariably estimate the position of the flux.

We will first establish an expression for absolute rotor flux angle (θ_{Ref}) in the stationary reference frame in terms of quantities that can either be measured or estimated. Readers who find the derivation indigestible need not worry as it is the conclusions that are important, not the analytical detail.

If we let the angle of the rotor body with respect to the stationary reference frame be θ, then the instantaneous angular velocities of the rotor flux wave and the rotor itself are given by

$$\omega_{flux} = \frac{d\theta_{Ref}}{dt}$$

$$\omega_{rotor} = \frac{d\theta}{dt}$$

The rotor motional e.m.f. is directly proportional to the rotor flux linkage and the slip velocity, i.e.

$$V_R = \psi_R\left(\omega_{flux} - \omega_{rotor}\right),$$

and the rotor current is therefore given by

$$I_R = \frac{\psi_R\left(\omega_{flux} - \omega_{rotor}\right)}{R_R}.$$

The corresponding component of stator current is given (see Fig. 8.12) by

$$I_{ST} = \frac{L_R}{M} I_R$$

Combining these equations and rearranging gives

$$\frac{d\theta_{Ref}}{dt} = \frac{MR_R}{\psi_R L_R} I_{ST} + \omega_{rotor} = \left(\frac{M}{\tau\psi_R}\right) I_{ST} + \omega_{rotor} \tag{8.3}$$

where τ is the rotor time-constant. Hence to find the rotor flux angle at time t we must integrate the expression above.

The mutual inductance M is a constant, and although the time-constant will vary because the rotor resistance varies with temperature, it will change relatively slowly, so we can treat it as constant, in which case the rotor flux angle is given by

$$\theta_{Ref} = \int_0^t \omega_{rotor} dt + \frac{M}{\tau} \int_0^t \frac{I_{ST}}{\psi_R} dt = \theta + \frac{M}{\tau} \int_0^t \frac{I_{ST}}{\psi_R} dt$$

Note that because of the symmetry of the rotor, we only need the time-varying element of the rotor body angle (θ), not the absolute position, so the constant of integration is not required. (In contrast, for vector control of permanent-magnet motors, the absolute position is important, because the rotor has saliency.)

The various methods that are used to keep track of the flux angle are what differentiate the various practical and commercial implementations of field-oriented control, as we will now see.

If we have a shaft encoder we can measure the rotor position (θ), or if we have a measured speed signal, we can derive θ by direct integration. This approach involves the fewest estimations, and therefore will normally offer superior performance, especially at low speeds, but is more costly because it requires extra transducers. We will refer to systems that use shaft feedback as 'closed-loop', but in the literature they may be also referred to as 'direct vector control.' In common with all schemes, the second term has to be estimated.

Many different methods of estimating the instantaneous parameter values are employed, but all employ a digital simulation or mathematical model of the motor/inverter system. The model runs in real time and is subjected to the same inputs as the actual motor, the model then being continuously fine-tuned so that the predicted and actual outputs match. Modern drives measure the circuit parameters automatically at the commissioning stage, and refine them on a near-continual basis to capture parameter variations.

The majority of vector control schemes eliminate the need for measurement of rotor position, and instead the rotor position term in Eq. (8.3) is also estimated from a motor model, based on the known motor voltage and currents. Rather confusingly, in order to differentiate them from schemes that do have shaft transducers, these systems are known as 'open-loop' or 'indirect' vector control. The term 'open loop' is a misleading one because at its heart is the closed-loop torque control shown in Fig. 8.16, but it is widely used: what it really means is 'no shaft position or speed feedback'.

The main problems of the open-loop approach occur at low speeds where motor voltages become very small and measurement noise can render the algorithms unreliable. Techniques such as the injection of high frequency "diagnostic" voltage signals exist, but are yet to find widespread acceptance in the market. Open loop inverter-fed induction motors are usually unsuitable for continuous operation at frequencies below 0.75 Hz, and struggle to produce full torque in this region.

An additional difficulty is that the significant variation of rotor resistance with temperature is reflected in the value of the all-important rotor time-constant, τ. Any difference between the real rotor time constant and the value used by the model causes an error in the calculation of the flux position and so the reference frame becomes misaligned. If this happens, the flux and torque control are no longer completely decoupled which results in sub-optimum performance and possible instability. To avoid this, routines are included in the drive to provide on-going estimates of the rotor time constant.

8.7 Direct torque control

Direct torque control is an alternative high-performance strategy to vector/field-orientation, and warrants a brief discussion to conclude our look at contemporary schemes. Developed from work first published in 1985 it theoretically provides the fastest possible torque response by employing a 'bang-bang' approach to maintain flux and torque within defined hysteresis bands. Like field-oriented control, it only became practicable with the emergence of relatively cheap and powerful digital signal processing.

Direct torque control avoids co-ordinate transformations because all the control actions take place in the stator reference frame. In addition there are no PI controllers, and a switching table determines the switching of devices in the inverter. These apparent advantages are offset by the need for a higher sampling rate (up to 40 kHz as compared with 6–15 kHz for field-orientation) leading to higher switching loss in the inverter; a more complex motor model; and inferior torque ripple. Because a hysteresis method is used the inverter has a continuously variable switching frequency, which may be seen as an advantage in spreading the spectrum of acoustic noise from the motor.

We saw in the previous sections that in field-oriented control, the torque was obtained from the product of the rotor flux and the torque component of stator current. But (as discussed in Chapter 9), there are other ways in which the torque can be derived, for example in terms of the product of the rotor and stator fluxes and the sine of the angle between, or the stator flux and current and the sine of the angle between them. The latter is the approach discussed in the next section, but first a word about hysteresis control.

A good example of hysteresis control is discussed later in this book, in relation to ‘chopper drives’ for stepping motors in Chapter 11. Another more familiar example is the control of temperature in a domestic oven. Both are characterised by a simple approach in which full corrective action is applied whenever the quantity to be controlled falls below a set threshold, and when the target is reached, the power is switched off until the controlled quantity again drops below the threshold. The frequency of the switching depends on the time-constant of the process and the width of the hysteresis band: the narrower the band and the shorter the time-constant, the higher the switching frequency.

In the domestic oven, for example, the ‘on’ and ‘off’ temperatures can be a few degrees apart because the cooking process is not that critical and the time-constant is many minutes. As a result the switching on and off is not so frequent as to be irritating and wear out the relay contacts. If the hysteresis band were to be narrowed to a fraction of a degree to get tighter control of the cooking temperature, the price to be paid would be incessant clicking on and off, and shortened life of the relay.

8.7.1 Outline of operation

The block diagram of a typical direct torque control scheme is shown in Fig. 8.24. There are several similarities with the scheme shown in Fig. 8.16, notably the inverter, the phase current feedback, and the separate flux and torque demands, which may be generated by the speed controller, as in Fig. 8.17.

However, there are substantial differences. Earlier we discovered that the inverter output voltage space phasor has only six active positions, and two zero states (see Fig. 8.14), corresponding to the eight possible combinations of the six switching devices. This means that at every instant there are only eight options in regard to the voltage that we can apply to the motor terminals. In the field-oriented approach, PWM techniques are employed to alternate

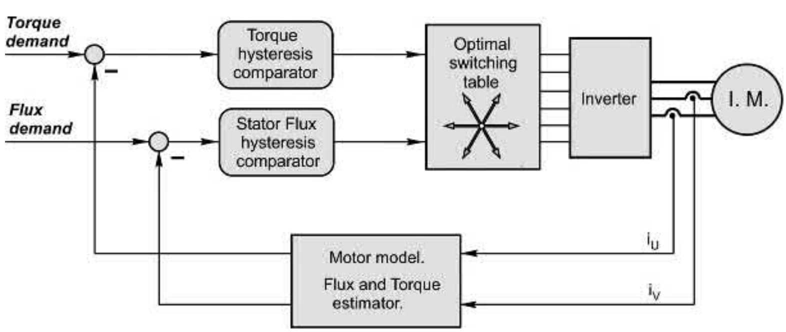

FIG. 8.24 Block diagram of typical direct torque control scheme.

between adjacent unit vectors to produce an effective voltage phasor of any desired magnitude and instantaneous position. However, with direct torque control, only one of the eight intrinsic vectors is used for the duration of each sample, during which the estimated stator flux and torque are monitored.

The motor model is exposed to the same inputs as the real motor, and from it the software continuously provides updated estimates of the stator flux and torque. These are compared with the demanded values and as soon as either strays outside its target hysteresis band, a logical decision is taken as to which of the six voltage phasors is best placed to drive the flux and/or torque back onto target. At that instant the switching is changed to bring the desired voltage phasor into play. The duration of each sample therefore varies according to the rate of change of the two parameters being monitored: if they vary slowly it will take a long time before they hit the upper or lower hysteresis limit and the sample will be relatively long, whereas if they change rapidly, the sample time will be shortened and the sample frequency will increase. Occasionally, the best bet will be to apply zero voltage, so one of the two zero states then takes over.

8.7.2 Control of stator flux and torque

We will restrict ourselves to operation below base speed, so we should always bear in mind that although we will talk about controlling the stator flux, what we really mean is keeping its magnitude close to its normal (rated) value, at which the magnetic circuit is fully utilised. We should also recall that when the stator flux is at its rated value and in the steady state, so is the rotor flux.

It is probably easiest to grasp the essence of the direct torque method by focusing on the stator flux linkage, and in particular on how (a) the magnitude of the stator flux is kept within its target limits and (b) how its phase angle with respect to the current is used to control the torque.

The reason for using stator flux linkage as a reference quantity is primarily the ease with which it can be controlled. When we discussed the basic operation of the induction motor in Chapter 5, we concluded that the stator voltage and frequency determined the flux, and we can remind ourselves why this is by writing the voltage equation for the stator as

$$V_S = I_S R_S + \frac{d\psi_S}{dt}$$

(We are being rather loose here, by treating space phasor quantities as real variables, but there is nothing to be gained by being pedantic when the message we take away will be valid.) In the interests of clarity we will make a further simplification by ignoring the resistance voltage term, which will usually be small compared with V_S. This yields

$$V_S = \frac{d\psi_S}{dt} \quad \text{or, in integral form,} \quad \psi_S = \int V_S dt$$

The differential form shows us that the rate of change of stator flux is determined by the stator voltage, while the integral form reminds us that to build the flux (e.g. from zero) we have to apply a fixed volt-second product, with either a high voltage for a short time, or a low voltage for a long time. We will limit ourselves to the fine-tuning of the flux after it has been established, so we will only be talking about very short sample intervals of time (Δt) during which the change in flux linkage that results ($\Delta\psi_S$) is given by

$$\Delta\psi_S = V_S \Delta t$$

As far as we are concerned, ψ_S represents the stator flux linkage space phasor, which has magnitude and direction relative to the stator reference frame, and V_S represents one of the six possible stator voltage space phasors that the inverter can deliver. So if we consider an initial flux linkage vector ψ_S as shown in Fig. 8.25A, and assume that we apply, over time Δt, each of the six possible options, we will produce six new flux-linkage vectors. The tips of the new vectors are labelled ψ_1 to ψ_6 in the figure, but only one (ψ_4) is fully drawn (dotted) to avoid congestion. There is also the option of applying zero voltage, which would of course leave the initial flux-linkage unchanged.

In (a) option 4 results in a reduction in amplitude and (assuming anticlockwise rotation) a retardation in phase of the original flux, but if the original flux linkage had a different phase, as shown in diagram (b), option 4 results in an increase in magnitude and an advance in phase. It should be clear that outcomes vary according to initial conditions, and therefore an extensive look-up table will be needed to store all this information.

Having seen how we can alter the magnitude and phase of the stator flux, we now consider the flux linkage phasor during steady-state operation with

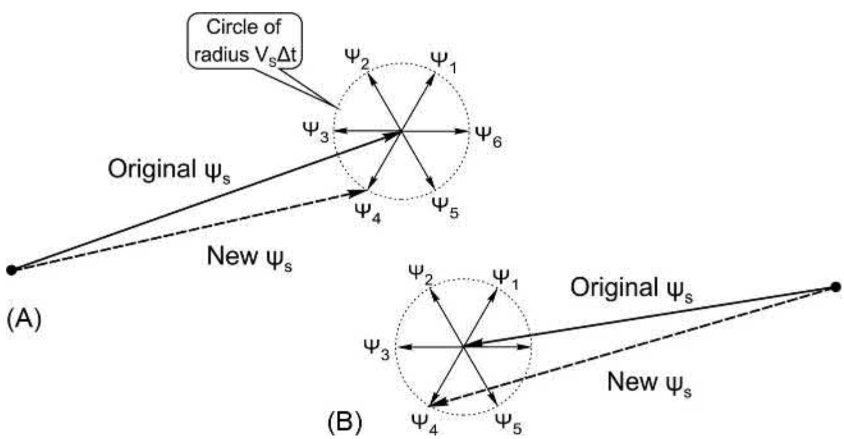

FIG. 8.25 Space phasor diagram of stator flux linkage showing how the outcome of applying a given volt-second product depends on the original phase angle. In (A) the magnitude is reduced and the phase is retarded, while in (B) the magnitude is increased and the phase is advanced.

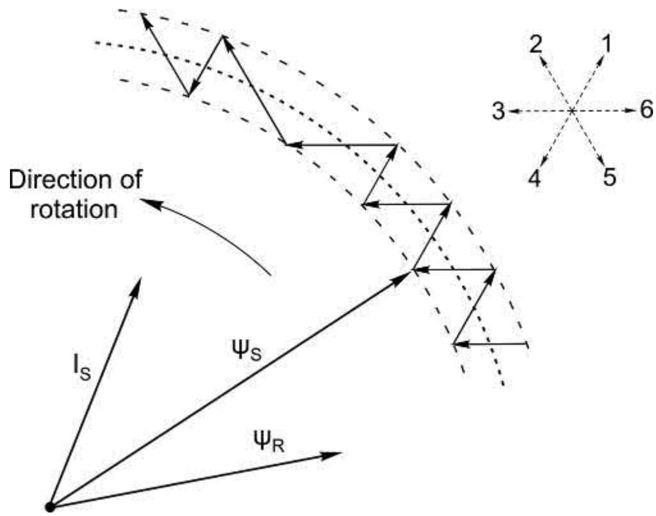

FIG. 8.26 Trajectory of stator flux linkage under steady-state conditions.

constant speed and torque, in which case we know that ideally all the space phasors will be rotating at a constant angular velocity.

The locus of the stator flux linkage space phasor (ψ_S) is shown in Fig. 8.26. In this diagram the spacing of hysteresis bands indicated by the innermost and outermost dotted lines have been greatly exaggerated in order to show the trajectory of the flux linkage phasor more clearly. Ideally, the phasor should rotate smoothly along the centre dotted line.

In this example, the initial position shown has the flux linkage at the lower bound, so the first switching brings voltage vector 1 into play to drive the amplitude up and the phase forward. When the upper bound is reached, vector 3 is used, followed by vector 1 again and then vector 3. Recalling that the change in the flux linkage depends on the time for which the voltage is applied, we can see from the diagram that the second application of vector 3 lasts longer than the first. (We should also reiterate that in this example only a few switchings take place while the flux rotates through sixty degrees: in practice the hysteresis band is very much narrower, and there may be many hundreds of transitions.)

We are considering steady-state operation, and so we would wish to keep the torque constant. Given that the flux is practically constant, this means we need to keep the angle between the flux and the stator current constant. This is where the torque hysteresis controller shown in Fig. 8.24 comes in. It has to decide what switching will best keep the phase on target, so it runs in parallel with the magnitude controller we have looked at here. Each controller will output a signal for either an increase or decrease in its respective variable (i.e. magnitude or phase) and these are then passed to the optimal switching table to determine the best switching strategy in the prevailing circumstances (see Fig. 8.24).

As we saw when discussing field-oriented control, it is not possible to make very rapid changes to the rotor flux because of the associated stored energy. Because the rotor and stator are tightly coupled it follows that the magnitude of the stator flux linkage cannot change very rapidly either. However, just as with field-oriented control, sudden changes in torque can be achieved by making sudden changes to the phase of the flux linkage, i.e. to the tangential component of the phasor shown in Fig. 8.26.

8.8 Review questions

(1) At full load, the rotor current in a motor with vector control is 30 A. Estimate the rotor current when the load torque reduces to 50%.

(2) In Fig. 8.17, a balloon attached to the flux controller shows a sketch graph. What do the axes represent? What does the shape of the graph indicate?

(3) In the field-oriented control literature, reference is often made to d-axis and q-axis currents. Which of them would you expect to stay more or less constant as the load changes, and why? What would you expect to happen to the other if the load torque were to double?

(4) When an induction motor with field-oriented control is first switched on, the alert onlooker may detect a brief delay before the rotor begins to turn. What is happening in this period?

(5) An induction motor with field oriented control drives an inertial load that has negligible friction. The duty typically requires the motor to track the symmetrical continuous speed-time demand shown in Fig. Q5, which it does with minimum speed error throughout.

At the maximum speed, the steady-state frequency is 40 Hz, and during acceleration the slip frequency is 2 Hz. The dotted lines are at half of the maximum speed.

Explain with reasons what you would expect the instantaneous frequency to be at each of the marked points (a to f).

If the speed -time profile also included periods of zero speed, how would you expect the motor to react if someone grasped its shaft and started to turn it?

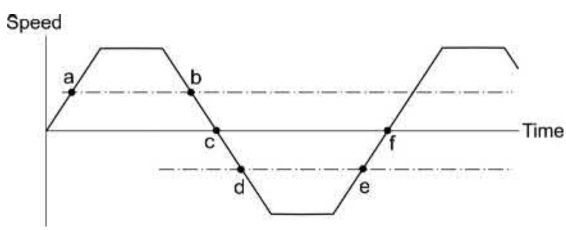

FIG. Q5

(6) When running in the steady state at full (base) speed and rated load, the angle between the stator flux linkage phasor (MI_s) and the rotor flux phasor (ψ_R) of a particular induction motor with field-oriented control is 70°.
Estimate the angle when the motor is running at:-
(a) 100% speed and 50% rated torque;
(b) 50% speed and full torque.

(7) Two identical induction motors drive identical loads. The supply to one is what the book refers to as 'inverter-fed', while the other has vector control. Both are running in the steady state at exactly the same speed and both are producing the same torque. How would the user know which was which?

Answers to the review questions are given in the Appendix.

Chapter 9

Synchronous, permanent magnet and reluctance motors and drives

9.1 Introduction

Chapters 6–8 have described the virtues of the induction motor and how, when combined with power electronic control, it is capable of meeting the performance and efficiency requirements of many of the most demanding applications. In this chapter another group of a.c. motors is described. In all of them the electrical power that is converted to mechanical power is fed into the stator, so, as with the cage induction motor, there are no sliding contacts in the main power circuits. The majority have stators that are identical (or very similar) to the induction motor, but some new constructional and winding techniques involving segmented construction are being applied at the lower power end: these will be introduced later.

It has to be admitted that the industrial and academic communities have served to make life confusing in this area by giving an array of different names to essentially the same machine, so we begin by looking at the terminology. The names we will encounter include:

(a) Conventional *Synchronous Machine* with its rotor field winding (excited-rotor). This is the only machine that may have brushes, but even then they will only carry the rotor excitation current, not the main a.c. power input.
(b) *Permanent Magnet Synchronous Machine* with permanent magnets replacing the rotor field winding.
(c) *Brushless Permanent Magnet Synchronous Motor* (same as (b)). The prefix 'brushless' is superfluous.
(d) *Brushless a.c. Motor* (same as (b).
(e) *Brushless d.c. Motor* (same as (b) except for detailed differences in the field patterns). This name was coined in the 1970s to describe 'inside out' motors that were intended as direct replacements for conventional d.c. motors, and in this sense it has some justification.
(f) *Permanent Magnet Servo Motor* (same as (b)).

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00009-X

The Reluctance motor is fundamentally different from all the above types. Whilst it rotates at synchronous speed, it does not have any form of excitation on the rotor and the method of torque production is somewhat different. It is nonetheless an increasingly important type of electrical machine.

Traditionally, as their name implies, synchronous motors were designed for operation directly off the utility supply, usually at either 50 Hz or 60 Hz. They then operate at a specific and constant speed (determined by the pole-number of the winding) over a wide range of loads, and therefore can be used in preference to induction motors when precise (within the tolerance of the utility frequency) constant speed operation is essential: there is no load-dependent slip as is unavoidable with the induction motor. These machines are available over a very wide range from tiny single-phase versions in domestic clocks and timers to multi-megawatt machines in large industrial applications such as gas compressors. (The clock application means that utility companies have a responsibility to ensure that the average frequency over a 24 h period always has to be precisely the rated frequency of the supply in order to keep us all on time. Ironically, in order to do this, they control the speed of very large, turbine-driven, synchronous machines that generate the vast majority of the electrical power throughout the world.)

To overcome the fixed-speed limitation that results from the constant frequency of the utility supply, inverter-fed synchronous motor drives are widely used. We will see that all forms of this generic technology use a variable-frequency inverter to provide for variation of the synchronous speed, but that in almost all cases, the switching pattern of the inverter (and hence the frequency) is determined by the rotor position and not by an external oscillator. In such so called 'self-synchronous' drives, the rotor is incapable of losing synchronism and stalling (which is one of the main drawbacks of the utility-fed machine). Field oriented control can be applied to synchronous machines to achieve the highest levels of performance and efficiency with machines which have higher inherent power densities than the equivalent induction motors.

9.2 Synchronous motor types

In the synchronous motor, the stator windings are essentially the same as in the induction motor, so when connected to the 3-phase supply, a rotating magnetic field is produced. However, instead of having a cylindrical rotor with a cage winding which automatically adapts to the pole number of the stator field, the synchronous motor has a rotor with either a d.c. excited winding (supplied via sliprings, or on larger machines an auxiliary exciter[1]), or permanent

1. An auxiliary exciter is simply a second, smaller machine with a 3-phase stator and rotor winding, mounted on the same shaft. A 3-phase supply on the stator is 'transformer coupled' to the rotor winding. The induced e.m.f. in the rotor is rectified and fed to the main motor field winding. The phase rotation of the supply to the auxiliary stator is opposite to that of the main motor so that when the motor comes to high speed the induced rotor e.m.f. remains high.

magnets, designed with the same pole number as the stator. The rotor is thus able to 'lock-on' or 'synchronise with' the rotating magnetic field produced by the stator. Once the rotor is synchronised, it will run at exactly the same speed as the rotating field despite load variation, so under constant-frequency operation the speed will remain constant as long as the supply frequency is stable.

As previously shown, the synchronous speed (in rev/min) is given by the expression

$$N_s = \frac{120f}{p}$$

where f is the supply frequency and p is the pole-number of the winding. Hence for 2, 4 and 6-pole industrial motors the running speeds on a 50 Hz supply are 3000, 1500, and 1000 rev/min, while on a 60 Hz supply they become 3600, 1800, and 1200 rev/min respectively. At the other extreme, the little motor in a time switch with its cup-shaped rotor with 20 axially projecting fingers and a circular coil in the middle is a 20-pole reluctance (synchronous) motor that will run at 300 rev/min when fed from a 50 Hz supply. Users who want speeds different from these discrete values will be disappointed, unless they are prepared to invest in a variable-frequency inverter.

With the synchronous machine, we find that there is a limit to the maximum (pull-out) torque (see Section 9.3) which can be developed before the rotor is forced out of synchronism with the rotating field. This 'pull-out' torque will typically be 1.5 times the continuous rated torque but can be designed to be as high as 4 or even 6 times higher in the case of high performance PM motors where, for example, high accelerating torques are needed for relatively short periods. For all torques below pull-out the steady running speed will be absolutely constant. The torque-speed curve is therefore simply a vertical line at the synchronous speed, as shown in Fig. 9.1. We can see that the vertical line extends into quadrant 2, which indicates that if we try to force the speed above the synchronous speed the machine will act as a generator.

Traditionally, utility-fed synchronous motors were used where a constant speed is required, high efficiency desirable, and power factor controllable. They were also used in some applications where a number of motors were required to run at precisely the same speed. However, a group of utility-fed synchronous motors could not always replace mechanical shafting[2] because whilst their rotational speed would always be matched, the precise relative rotor angle of each motor would vary depending on the load on the individual motor shafts.

2. Drive shafts were used in early textile factories and the like. A single mechanical shaft would be fed through the factory or machine and various functions would be connected to the same shaft via belts. The connected equipment would then run up and down together. With toothed belts, position synchronism could be achieved.

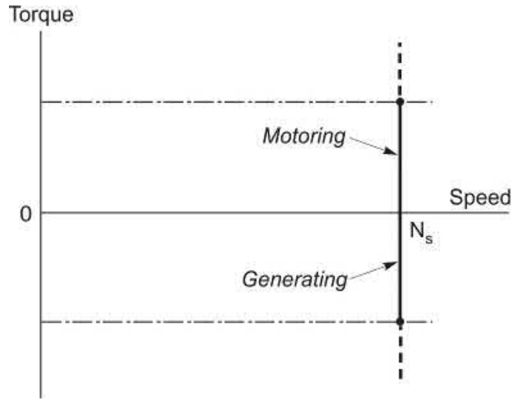

FIG. 9.1 Steady-state torque-speed curve for a synchronous motor supplied at constant frequency.

9.2.1 Excited-rotor motors

The rotor of a conventional synchronous machine carries a 'field' or 'excitation' winding which is supplied with direct current either via a pair of sliprings on the shaft, or via an auxiliary brushless exciter on the same shaft. The field winding is designed to produce an air-gap field of the same pole-number and spatial distribution (usually sinusoidal) as that produced by the stator winding. The rotor may be more-or-less cylindrical, with the field winding distributed in slots (Fig. 9.2A), or it may have projecting ('salient') poles around which the winding is concentrated (Fig. 9.2B).

A cylindrical-rotor motor has little or no reluctance (self-aligning) torque (discussed later), so it can only produce torque when current is fed into the rotor. On the other hand, the salient-pole type also produces some reluctance torque even when the rotor winding has no current. In both cases, however, the rotor 'excitation' power is relatively small, since all the mechanical output power is supplied from the stator side.

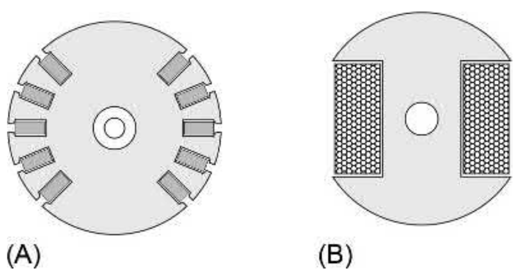

FIG. 9.2 Rotors for synchronous motors. (A) 2-pole cylindrical, with field coils distributed in slots, (B) 2-pole salient pole with concentrated field winding.

Excited rotor motors are used in sizes ranging from a few kW up to many MW. The large ones are effectively alternators (as used for power generation) but used as motors. Wound rotor induction motors (see Chapter 6) can also be made to operate synchronously by supplying the rotor with d.c. through the sliprings.

As we will see later, an advantage of the excited rotor type is that the power factor can be controlled over a wide range by varying the rotor excitation current.

9.2.2 Permanent magnet motors

The synchronous machines considered so far require two electrical supplies, the first to feed the field/excitation and the second to supply the stator. Brushless permanent magnet (Brushless PM) machines have magnets attached to the rotor to provide the field, and so only a stator supply is required. The principle is illustrated for 4-pole surface-mounted and 10-pole buried/interior types in Fig. 9.3. Motors of this sort have typical output ranging from about 100 W up to perhaps 500 kW, though substantially higher ratings have been made.

The advantages of the permanent magnet type are that no supply is needed for the rotor and the rotor construction can be robust and reliable. The disadvantage is that the excitation is inherently fixed, so the designer must either choose the shape and disposition of the magnets to match the requirements of one specific load, or seek a general-purpose compromise. Control of power-factor via excitation is of course no longer possible. Within these constraints the Brushless PM synchronous motor behaves in very much the same way as its excited-rotor sister.

9.2.3 Reluctance motors

The reluctance motor is arguably the simplest synchronous motor of all, the rotor consisting simply of a set of laminations shaped so that it tends to align itself with the field produced by the stator. This 'reluctance torque' action is discussed in Section 9.3.3.

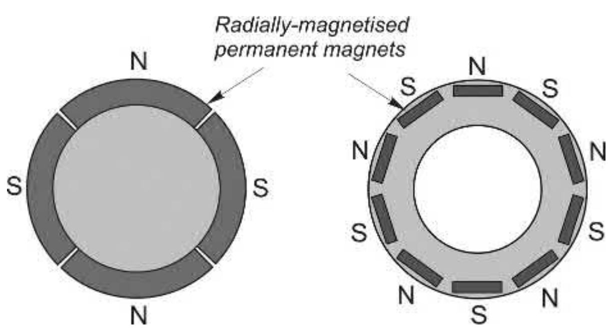

FIG. 9.3 Permanent magnet synchronous motor rotors; 4-pole (left); 10-pole (right).

Until recently reluctance motors were intended for use from the utility supply, so in order to get the rotor up to a speed close enough to the synchronous speed for the rotor to make the final leap and lock onto the rotating stator field, the rotor also carries a cage winding which provides accelerating torque for the run-up, but thereafter carries no current. The rotor therefore resembles a cage induction motor, with parts of the periphery cut away in order to force the flux from the stator to enter the rotor in the remaining regions where the air-gap is small, as shown in Fig. 9.4A. Alternatively, the 'preferred flux paths' can be imposed by removing iron inside the rotor so that the flux is guided along the desired path, as shown in Fig. 9.4B and C, the latter being for an inverter-fed motor, where a starting cage is not required. All these rotor types can be seen to have "salient" poles.

The rotor will tend to align itself with the field, and hence is able to remain synchronised with the travelling field set up by the 3-phase winding on the stator in much the same way as a permanent-magnet rotor. Early reluctance motors were invariably one or two frame sizes bigger than an induction motor for a given power and speed, and had low power-factor and poor pull-in performance. As a result they fell from favour except for some special applications such as textile machinery where large numbers of cheap synchronised motors driven from a single variable frequency inverter were used. Understanding of reluctance motors is now much more advanced, though their fundamental performance still lags the induction motor as regards power-output, power factor and efficiency. Recently, there is renewed interest in rotors that exploit both permanent magnet and reluctance torque, and this is discussed in Section 9.3.5.

9.2.4 Hysteresis motors

Whereas most motors can be readily identified by inspection when they are dismantled, the hysteresis motor is likely to baffle anyone who has not come across it before. The rotor consists simply of a thin-walled cylinder of what looks like steel, while the stator has a conventional single-phase or three-phase winding. Evidence of very weak magnetism may just be detectable on the rotor, but there

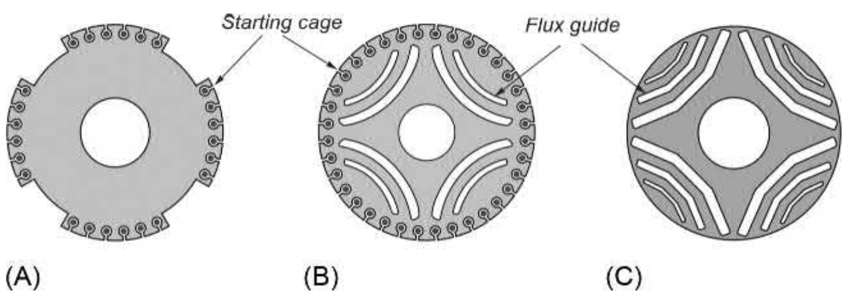

FIG. 9.4 Reluctance motor rotors (4-pole).

is no hint of any hidden magnets as such, and certainly no sign of a cage. Yet the motor runs up to speed very sweetly and settles at exactly synchronous speed with no sign of a sudden transition from induction to synchronous operation.

These motors (the operation of which is quite complex) rely mainly on the special properties of the rotor sleeve, which is made from a hard steel which exhibits pronounced magnetic hysteresis. Normally in machines we aim to minimise hysteresis in the magnetic materials, but in these motors the effect (which arises from the fact that the magnetic flux density B depends on the previous 'history' of the m.m.f.) is deliberately accentuated to produce torque. There is actually also some induction motor action during the run-up phase, and the net result is that the torque remains roughly constant at all speeds.

Small hysteresis motors were once used extensively in office equipment, fans, etc. The near constant torque during run-up and the very modest starting current (of perhaps 1.5 times rated current) means that they are also suited to high inertia loads such as gyro compasses and small centrifuges.

Hysteresis motors are used in niche areas, and so we will not consider them any further.

9.3 Torque production

The aim of this section is to present physical pictures of the torque production mechanism in the various types of synchronous motor. We deliberately concentrate on the same 'BIl' approach that we have followed in relation to the d.c. and the induction motor, in order to highlight the fundamental similarities between the three types of motor.

We start with the excited-rotor motor in order to establish a general approach to the interaction between stator and rotor currents and fields, and to develop qualitative expressions for torque. A relatively simple adaptation is then made to deal with the permanent-magnet type, and with a further twist, we can throw light on the torque in a reluctance motor (which has neither winding nor magnets on the rotor). Finally we look briefly at salient pole synchronous motors, which develop both excitation torque and reluctance torque.

The spatial images developed in this section will help later in the chapter when we link them to the voltages and currents in the steady-state time phasor diagram. Readers who are familiar with the theory of electrical machines will be aware that books often omit the material in this section in favour of the coupled-circuit approach that we discussed briefly in Section 8.3; however, we believe that the physical approach will be preferred by our target readership.

9.3.1 Excited rotor motor

In explaining the mechanism of torque production in the induction motor (Chapter 5) we chose to focus on the 'BIl' force (and hence the torque) resulting from the interaction between the rotor current wave and the resultant radial flux density wave (i.e. the resultant flux due to the combined effect of the stator and rotor m.m.f.'s), both of which are sinusoidally distributed in space.

For the synchronous machine we could use the same approach for the excited rotor type, but the PM rotor and reluctance versions don't have current on the rotor when running synchronously, so we consider the torque on the stator instead. We assume here that the reader is happy to accept that electromagnetic torque on the rotor will always be accompanied by an equal and opposite torque on the stator.

We will explore the torque mechanism by looking first at the static condition (which is equivalent to taking a snapshot of the running condition), as shown in Fig. 9.5. This represents a model of a smooth rotor machine with sinusoidally distributed currents on both stator and rotor: the stator current is fixed in all five sketches, and the rotor current is fixed relative to the rotor. The blue flux lines represent in simplified form the sinusoidal m.m.f. and flux density produced by the rotor current, and so when the rotor turns, they follow; only a few lines are shown for the sake of clarity. The drum on the rotor carries a rope to which weights can be attached to apply mechanical torque to the rotor.

The 'rule of thumb' relating to the torque produced by the interaction of rotor and stator fields is that the torque always acts so as to move the two fields into alignment. The flux pattern produced by the stator winding is not shown in Fig. 9.5 in order to avoid overcrowding, but it would be similar to the rotor flux, and in all three sketches it would be directed upwards.

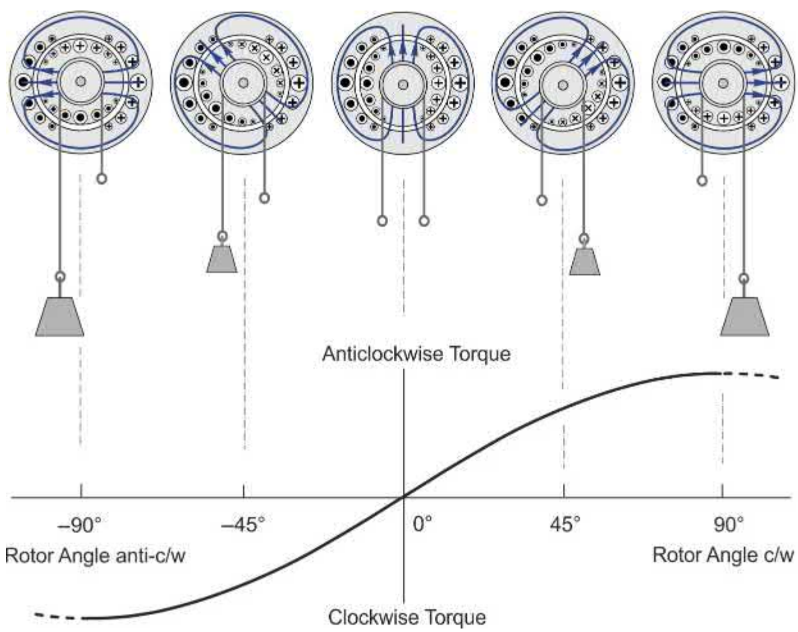

FIG. 9.5 Static torque in an excited rotor motor.

Hence if the rotor is free to move, and there is no external torque applied, it will come to rest as shown in the centre sketch, with zero torque. We can confirm that the resultant torque is zero in this position by considering the 'BIl' forces on the stator conductors caused by interaction with the rotor flux. Because of the symmetry there is no resultant force on the group of stator conductors carrying positive current, because any in the upper half exposed to positive (outward) flux will experience a clockwise force, while those in the lower half will be exposed to negative (inward) flux, and therefore will experience an equal anticlockwise force. By the same token, there is no resultant force on the conductors carrying negative current, and there is therefore no torque on the stator, and hence none on the rotor.

However, when the rotor is turned in a clockwise direction by a load, as in the two right-hand sketches, more stator conductors carrying positive current are exposed to positive rotor flux, and those carrying negative current are exposed to negative flux. Application of Fleming's left hand rule shows that there is a clockwise torque on the stator, and therefore an equal anti-clockwise torque on the rotor, the torque increasing with angle. When the rotor torque equals the load torque, the rotor is at rest, in a stable equilibrium: If it is displaced in either direction, and then released, it will settle back at the same angle.

The unloaded case shown in the centre sketch is also a stable equilibrium, with any clockwise displacement leading to an anti-clockwise restoring torque, and vice-versa. Beyond 90°, however, the rotor is unstable, and so although there is theoretically another zero torque at 180°, the unloaded rotor could not remain at rest there because the slightest nudge would cause it to flip round in the direction it had been disturbed, and come to rest as in the centre sketch.

The right hand sketch shows the position of maximum torque, the rotor flux being horizontal and aligned with the stator current distribution. Recalling that the flux produced by the stator current is vertical, we see that the condition for maximum torque is that the two fields are perpendicular.

The theoretical torque-angle curve is shown below the three sketches in Fig. 9.5, and is a sinusoidal function of rotor angle. The beginnings of the unstable regions are indicated by dotted lines. The peak torque is, as expected, proportional to the product of the stator current and the rotor current (or flux). For obvious reasons, the angle between the rotor and stator fields is known as the torque angle.

We have considered the stator current to be constant in order to explore the mechanism of torque production, and the mental picture of the two fields always tending to align themselves is a useful one, which we can easily extend to the running condition. We now imagine that the rotor is running in the steady-state at the synchronous speed, with the amplitude and frequency of the stator current kept constant, the rotor being 'locked onto' or dragged along by the rotating stator field. The variation of the 'torque angle' with load on the shaft is then the same as we have seen in the static case, the angle increasing as the load torque increases. Clearly, if the load torque exceeds the point at which the torque angle is 90°, the rotor will lose synchronism and stall. However it is important to stress that in practice 'constant stator current' is not a normal operating condition. When the motor is operated from the utility supply, for example, it is the stator voltage (and frequency) that are constant, while in an inverter-fed drive with field-oriented control (see later), the system will always control the magnitude, speed and instantaneous position of the stator current distribution relative to the rotor position in order to provide the torque required.

Torque—Excited rotor motor

Hitherto, when we have invoked the 'BIl' picture to explain the production of torque we have always taken the B to be the resultant flux density, resulting from the combined effect of both stator and rotor currents, but in the discussion above we only considered the torque produced on the stator by the rotor flux. We will now see that if the windings are sinusoidally distributed, and the rotor surface is smooth (i.e. it may have slots, but no major saliencies), there are several ways of expressing the torque, not all involving the resultant flux, and we can choose which suits best according to the circumstances.

We saw in Section 8.2.1 that sinusoidally space-distributed quantities can be represent by space phasors or vectors, so we can represent an m.m.f. (or its corresponding flux density distribution) by a vector whose length is proportional to the m.m.f. (and hence to the current), and whose direction is determined by the instantaneous angular position of the current wave.

The general case is shown in Fig. 9.6, where F_S, F_R, and F are the stator, rotor and resultant m.m.f.'s respectively. The angle between stator and rotor m.m.f's (the torque angle), is λ, and the angle between the rotor m.m.f. and the resultant m.m.f. (δ) is the load angle (see later). The rotor m.m.f. in this space phasor diagram has been chosen deliberately to be horizontal in order to be consistent with the time phasor diagrams of pm motors that we will discuss later in this chapter.

We have already established that the torque is proportional to the product of the stator and rotor m.m.f.'s, and to the sine of λ, i.e.

$$T \propto F_S F_R \sin\lambda$$

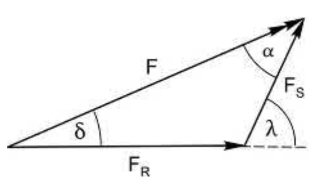

FIG. 9.6 Space phasors showing rotor, stator, and resultant m.m.f.'s.

The expression $F_S F_R \sin\lambda$ is the area of the triangle defined by the sides F_S and F_R, so we see that the area provides an immediate visual indication of the torque. We can also see that when the stator and rotor m.m.f.'s are aligned, the torque is zero, and when they are perpendicular, we get maximum torque.

Readers who are familiar with vector calculus will recognise the expression for the area of a triangle as the amplitude term of the so-called 'cross product' of two vectors (i.e. if the vectors are of magnitudes A and B and the angle between them is γ, their cross product is AB sin γ).

The triangles defined by the sides F and F_R, and F and F_S, both have the same area as the triangle defined by F_S and F_R, so the torque can equally well be expressed in two further ways, leading to the three equivalent cross-product formulations shown below:

$$\text{Stator m.m.f. with Rotor m.m.f.} \ (F_S F_R \sin\lambda)$$

$$\text{Resultant m.m.f. with Rotor m.m.f.} \ (F F_R \sin\delta)$$

$$\text{Resultant m.m.f. with Stator m.m.f.} \ (F F_S \sin\alpha)$$

We have used the first of these in slightly modified form in previous chapters, where instead of resultant m.m.f. we talked of resultant flux density, and instead of rotor m.m.f. we used rotor current distribution, but as both are proportional to their respective m.m.f.'s, and we are not being quantitative, there is no inconsistency. In fact, we will often substitute 'flux density' or (for a given machine) 'flux' in place of m.m.f. later in this chapter when we are discussing what determines the torque.

In the excited rotor case it turns out that the first version is the simplest way of picturing the torque mechanism when the currents are specified (as in an inverter-fed drive), and in particular it defines what we mean by the torque angle, λ. We should also note what many readers may think is obvious, which is that if the rotor m.m.f. is zero, and the only source of excitation is on the stator, the stator flux will not produce any torque. However, this is only true for smooth rotor machines, and things are different for rotors with salient poles, as we will see later.

When we move on to consider steady-state operation from the utility supply, the second torque formulation will prove more fruitful for our discussions, and we will make use of the load-angle (δ) rather than the torque angle.

It is worth mentioning that an alternative way of looking at a cross product is that it can be obtained by taking the product of the first vector with the component of the second that is perpendicular to the first, and we will make use of this later in this chapter, where we refer to the axis of the rotor flux as the 'direct axis' and to the perpendicular component as the 'quadrature axis' component.

Physically, we picture maximum torque when stator and rotor m.m.f.'s or fluxes are in quadrature, and zero torque when they are aligned. At other angles, we recognise that it is only the quadrature component of the second vector that contributes to torque.

9.3.2 Permanent magnet motor

The permanent magnet (PM) motor (Fig. 9.3) behaves in a similar way as the excited rotor one, with the obvious exception that the 'strength' of the flux produced by the magnet cannot be varied once it has been magnetised.

To avoid going deeply into the properties of permanent-magnets, we can picture the magnet as an m.m.f. source, whose external magnetic circuit comprises three reluctances in series, viz. the 'iron' part of the rotor body; the stator iron; and the air-gap between rotor and stator iron. The latter term is dominant in any machine, but it is even more so here because the "airgap" is at least equal to the radial thickness of the magnets, and so is much larger than in the excited rotor version. Despite this, the high magnet m.m.f. produces the required flux density at the stator. There is no variation in the reluctance 'seen' by the magnet as the rotor turns, so we can picture a rotor flux wave that remains constant and whose instantaneous position is determined by the rotor angle: in future, we will denote this by the symbol φ_{mag}.

Looked at from the stator side we might wonder how the magnet material influences the reluctance seen by the stator m.m.f. Again, we can take a simplified view and, as far as an externally applied field is concerned, we treat the magnet material as if it had the same permeability as air. The stator m.m.f. therefore sees a high reluctance because the effective air-gap is much larger than usual due to the radial depth of the magnets. This means the stator self-inductance is much lower than that of a similar excited rotor motor, which is advantageous as far as rapid current control is concerned.

Torque—Permanent magnet motor

In line with our discussion above, if we denote the stator (or, to use the traditional word 'armature') flux wave by φ_{arm} and the angle between the magnet and stator flux waves (the torque angle) by λ, the torque is expressed by

$$T \propto \left(\Phi_{mag}\right)\left(\Phi_{arm}\right)\sin\lambda$$

i.e. the torque depends on product of the rotor and stator fluxes and the sine of the angle between them. Note that we could equally have chosen to use the respective m.m.f.'s in the torque expression rather than the fluxes, in which case the torque expression would be identical to that for the excited rotor case.

9.3.3 Reluctance motor

The reader who absorbed Chapter 1 will recall that the word reluctance is used in the context of magnetic circuits to define the analogous quantity to resistance in an electric circuit, i.e. the ratio of m.m.f. to flux. Given that all motors have magnetic circuits, and involve reluctance, it may seem odd that the same word also describes a particular type of motor, but the justification for doing so should soon become clearer.

Up to now in this chapter we have seen how to picture torque from the 'BIl' interaction between the rotor component of the resultant flux density and the stator current. However, the rotor of a reluctance motor has no current-carrying conductors (apart from, perhaps, a starting cage, which only functions below synchronous speed), so clearly the only source of excitation is the stator winding. Nevertheless, torque is produced, so it would not be unreasonable to suppose that an entirely different mechanism is responsible, and that our trusted friend 'BIl' will not provide an explanation. (This impression would be reinforced if we looked at the literature on reluctance motors, which concentrates on the circuit modelling approach.) However, as we will see, the 'BIl' approach not only illuminates the physical basis for understanding reluctance torque, but it also allows us to obtain a simple expression for torque in terms of stator current.

We begin with a hypothetical set up of no practical use but which will help our later discussion.

Fig. 9.7 shows an idealised machine with a smooth air-gap and infinitely permeable stator and rotor cores, and an ideal sinusoidal distribution of stator current. As we have seen previously, this idealised model leads to a close approximation of the resultant field produced by a real machine with three-phase sinusoidally distributed windings fed from a three-phase sinusoidal supply.

The rotor is non-salient, i.e. a uniform homogeneous cylinder, and has no current carrying conductors.

The flux produced by the stator current is shown by the hand-sketched red lines. In this instance, sketching is aided by the fact that for this idealised situation, an analytical solution is available that shows that the flux density inside the rotor is uniform, say B. The sketch is too small to show the air-gap flux density in any detail, but we can find the radial flux density at any angle θ measured from the vertical axis by resolving B, which yields the radial air-gap flux density as

$$B_r = B\cos\theta$$

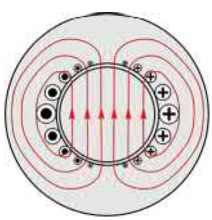

FIG. 9.7 Flux pattern produced by a 2-pole sinusoidally-distributed stator winding.

This is what is expected: the gap is uniform, so the radial flux density wave is proportional to, and in phase with the stator m.m.f., which in turn is in quadrature with the stator current distribution. Maximum air-gap flux density is at the top, and maximum current is on the horizontal axis, and there is therefore no resultant torque.

If the previous paragraph fails to convince, we can look directly at the forces on the stator conductors, using 'BIl'. First, consider the topmost conductor on the left: it carries a current out of the paper, and is exposed to a radially outward air-gap flux density, so it experiences a force to the left, and the rotor therefore experiences an equal and opposite force to the right. The corresponding conductor on the right carries current into of the paper, and is exposed to the same outward flux density, so its force is to the right, and the reaction on the rotor is to the left. Because of the symmetry, there is no resultant torque on the rotor, regardless of its angular position. (If we had taken the circuit viewpoint, where torque is linked to the change of inductance with position, we would have deduced that the torque was zero because the magnetic circuit 'seen' by the stator winding does not vary as the rotor turns, so the inductance is constant.)

Now we turn to the 'salient' rotor shown in the four sketches in Fig. 9.8. Two segments of the original cylindrical rotor have been removed to leave two projecting poles or 'saliencies', thereby forming the rotor of a 2-pole reluctance motor. (We have deliberately chosen a simple shape for illustrating the principle: a real rotor would probably look more like a 2-pole version of those shown in Fig. 9.4. The grey dot on the rotor is there to allow us to correlate sketches (a) to (d) with the torque plot below.

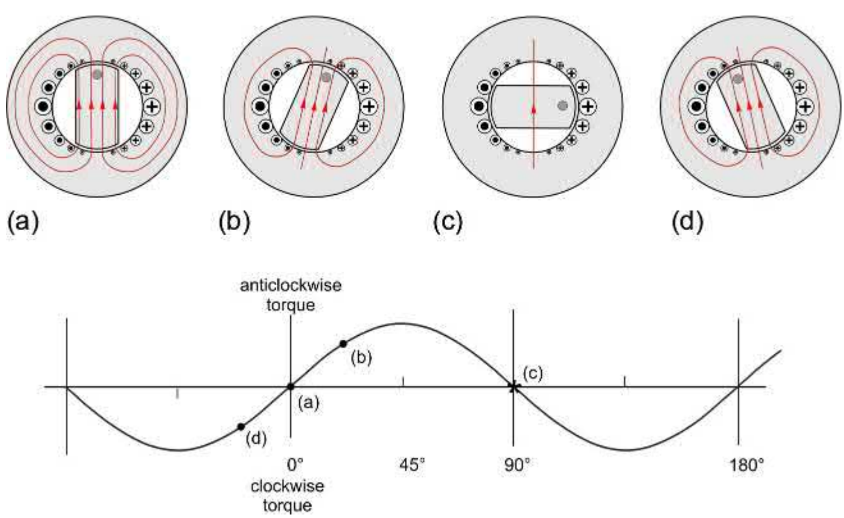

FIG. 9.8 Static torque in a reluctance motor.

When the rotor is aligned as in (a) the sinusoidal stator m.m.f. acts on a relatively low-reluctance path, so the flux is relatively large, (but not as high as it was when the rotor was cylindrical because the overall reluctance is now higher). When the flux takes the low-reluctance path as in (a) it is said to be flowing along the 'direct axis'. Recalling that inductance is the ratio of flux linking the stator winding to the current producing it, the inductance in case (a) is known as the direct-axis inductance, L_d: we will return to this later.

At the other extreme, when the rotor is aligned as in (c), the same m.m.f. now acts on a relatively high reluctance path (the 'quadrature axis') and the flux is low: the inductance in this situation is the quadrature-axis inductance (L_q), which is of course much less than L_d.

The terms direct axis and quadrature axis define mutually perpendicular axes fixed to the rotor, and will be used extensively later in this chapter, especially when we talk about the control of these machines with a variable frequency inverter.

It should be clear that because of the non-uniform air gap, the flux density wave will no longer be sinusoidal, which might suggest that our previous simple approach to torque (via the product of two displaced sine waves) would no longer apply. In fact, as we will see next, as long as we focus on the fundamental component of the flux density wave, we can obtain torque as before.

An important property of any winding that we have not mentioned previously is the reciprocal relationship between the m.m.f. and flux density produced by a winding, and the reaction of that winding to its own internal or another external field. By 'reaction to' we mean both the e.m.f. induced in the winding by either the self-produced field or the external field, and/or the torque produced on it when it is carrying current. For example, if the winding m.m.f. consisted of only a fundamental and fifth space harmonic, the winding would only react to external fields of fundamental and fifth space harmonic. (To emphasis the point, if we were to put a 4-pole permanent magnet rotor into a 2-pole stator carrying current, there would be 'BIl' forces on the individual conductors of the stator, some positive and some negative, but the resultant force would be zero.)

In the present context, we have a sinusoidally distributed winding but a non-sinusoidal flux density wave, so as long as we restrict consideration to the fundamental component of the flux density wave, we can continue to explore the mechanism of torque production as we have done so far.

When we apply 'BIl' to find the forces on the stator windings, we find that in both (a) and (c), the symmetry results in zero net torque on the stator windings, and thus zero torque on the rotor, as shown in the graph in Fig. 9.8.

When we displace the rotor as shown in (b), the stator conductors that overlap the top of the rotor all carry current into the paper, and thus experience forces to the right, while those at the bottom carry current out of the paper and experience a force to the left: the reaction forces on the rotor therefore give rise to anticlockwise torque which tends to return the rotor to position (a). Not surprisingly, if we displace the rotor anticlockwise instead (see (d)), the rotor torque becomes clockwise, again tending to restore the rotor to position (a). Position (a) is therefore one of stable equilibrium, i.e. any displacement in either direction brings a restoring torque into play, and the greater the displacement, the greater the restoring torque.

In practice we find that there is a maximum torque, and it occurs at 45°, as shown. Beyond that point the torque diminishes to zero at 90°, before reversing and reaching a peak again at 135°. At 90°, the torque is zero, but any displacement of the rotor not only results in a torque that tries to increase the angle further, but also the torque increases with the displacement. The 90° position is therefore an unstable equilibrium (marked by a star): if we put the rotor there, any tiny disturbance will cause it to flip away, and settle (after oscillating) at a stable equilibrium point such as (a), marked with a dot.

We have considered the stator current wave to be stationary in order to simplify the discussion, but in practice the field will be rotating synchronously at a speed determined by the frequency of the stator currents. Under such steady-state conditions the rotor torque must match the load torque, so under ideal no-load conditions the rotor direct axis will be aligned with the applied field (as in Fig. 9.8a) and no torque will be produced. When the load torque is increased, causing a momentary deceleration, the rotor axis drops back relative to the field, thereby producing a motor torque that increases with angle until it equals the load torque, whereupon the speed is again synchronous. It is also worth noting that, once synchronised, a reluctance machine will operate as a generator, in which case the load angle becomes negative.

Torque—Reluctance motor

For the non-salient rotors looked at in the previous two sections we were able to obtain very simple expressions for the torque because the amplitudes of the stator and rotor fields remained constant as the rotor angle varied. Here, things are more complex, because the amplitude of the stator field varies with rotor position (the variation being governed by the change in reluctance with angle), and as a result it is a bit more tricky to obtain an expression for torque in terms of the stator current. The exercise is useful however, because it introduces some ideas that are taken up again later when we discuss torque in salient-pole excited motors.

The sketches in Fig. 9.9 relate to the basic reluctance motor shown in Fig. 9.8. Fig. 9.9A is a reminder that we are dealing with a salient-pole rotor and a stator winding carrying a sinusoidally distributed current, while diagrams (B), (C), and (D) are space vector diagrams corresponding to various rotor positions. (We explained in Chapter 8 that it is very convenient to represent currents, m.m.f.'s and fluxes that are sinusoidally distributed in space by vectors, and the following discussion will underline the value of this approach.)

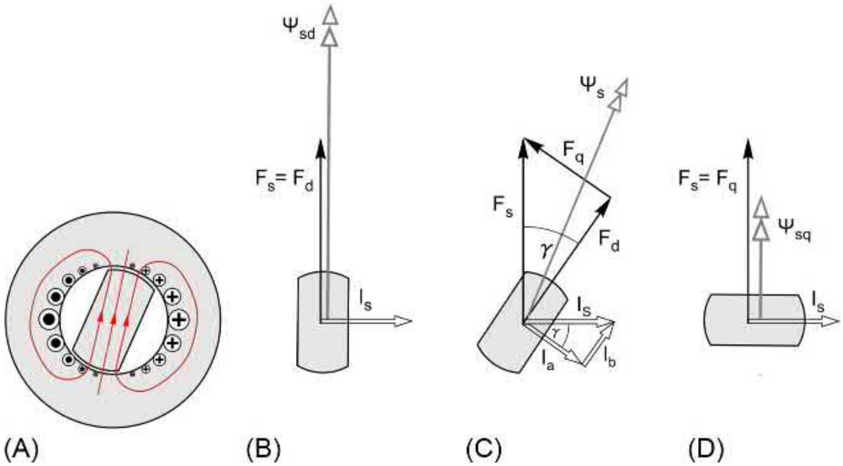

FIG. 9.9 Space phasors for a reluctance motor.

What we want to find is how the magnitude and phase of the stator flux linkage varies with the position of the rotor: if we know this, we can find the torque by applying the familiar 'BIl' approach used hitherto. The general position is shown in sketch (C), but it is best for us to establish some general ideas first by considering the limiting cases shown in sketches (B) and (D).

The peak of the stator current distribution (I_s) lies on the horizontal axis, so the vector that represents the stator current distribution is horizontal and pointing to the right: this vector remains constant in all three sketches. The corresponding m.m.f. wave (F_s) is directed vertically upwards, and again it is constant in all three sketches.[3] The suffices d and q refer to the direct and quadrature axes of the rotor, d being the low-reluctance axis, and q being the high reluctance path.

In sketch (B), for example, the stator m.m.f. F_s is directed along the rotor direct axis, and so it has been labelled F_d. The reluctance along the direct axis is low, so the stator flux linkage ψ_{sd} will be large. In contrast, in sketch (D) the same stator m.m.f. F_s is directed along the rotor's quadrature axis: it is therefore labelled $F_{q;}$ and because the reluctance is relatively high, the corresponding stator flux linkage ψ_{sq} is relatively small.

In sketch (C), the stator m.m.f. vector has been resolved into direct and quadrature axis components. If each m.m.f. component acted on the same reluctance, the resultant flux linkage would be in phase with F_s. But the d-axis reluctance is much lower than the q-axis reluctance, so the resultant stator flux linkage (ψ_s) is

3. The 90° shift in space between the current and m.m.f. vectors occurs because the m.m.f. is the integral with respect to distance of the current distribution. This can be seen in Fig. 5.4, where each time a positive current-carrying conductor (a current impulse in mathematical terms) is reached, the m.m.f. increases by a fixed step.

shifted towards the direct axis. There are some similarities here with our previous discussions, where we saw that torque production required two fields to be displaced by an angle. This condition is met to a varying degree for all rotor positions between sketches (B) and (D), with m.m.f. and flux misaligned; but at the extremes both fields become co-phasal and the torque falls to zero.

In order to obtain the variation of torque with rotor angle and stator current, we need to form an expression equivalent to the 'BIl' product at the stator winding. Hitherto we have worked in terms of the flux density (B), but we can equally well use the flux linkage vector because we have specified that all quantities are sinusoidally distributed.

The sketch on the left in Fig. 9.10 is a repeat of the space vector diagram (C) in the previous sketch, and it shows the stator current distribution (I_s) resolved into components I_a and I_b that are responsible for producing the m.m.f. components F_d and F_q, respectively. (It is important to point out that the resolved components of the current distribution vector in Fig. 9.10 should not be confused with the direct and quadrature axis currents that we will meet later in this chapter: the former are spatial quantities, represented here by outline arrows, while the latter are sinusoidal time-varying quantities, and are represented in time phasor diagrams by solid grey arrows.)

We need to find the component of the resultant flux linkage that is co-phasal with the current so that we can apply 'BIl'. That component is represented by the distance x, but we do not know the angle α, so instead we find the difference between z and y, yielding the flux linkage component (equivalent to the 'B' in BIl) as ($L_d I_a \sin\gamma - L_q I_b \cos\gamma$). The 'I' part of BIl is of course I_s. So we find that the torque (or strictly the force, but here we are seeking a general qualitative expression applicable to a generic motor, so dimensions are not considered) is given by

$$T \propto \left(L_d I_a \sin\gamma - L_q I_b \cos\gamma\right) I_s$$

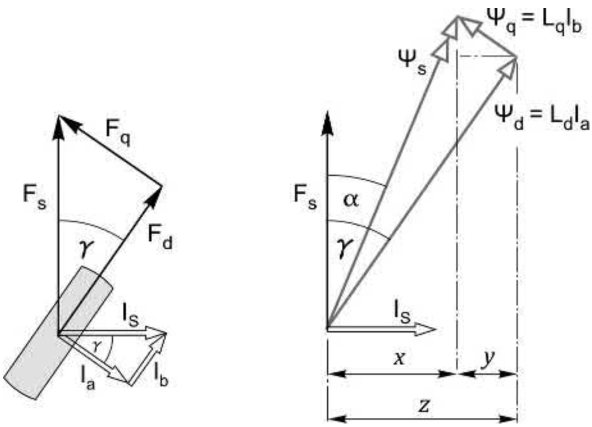

FIG. 9.10 Reluctance motor—axis m.m.f's and flux linkages.

If we now express the current distributions I_a and I_b in terms of Is, i.e. $I_a = I_s \cos\gamma$ and $I_b = I_s \sin\gamma$, the torque is given by

$$T \propto I_s^2 \left(L_d - L_q\right) \sin 2\gamma$$

This simple expression shows that the torque depends on the square of the current, and is therefore the same for both positive and negative current: the torque is also a double-frequency function of the rotor angle, as we saw earlier. We already knew that in order to work at all, the rotor had to have saliency, but we now see that the torque is directly proportional to the difference between the direct and quadrature axis inductances, so that, broadly speaking, the greater the difference, the better. However, we will see later that this is not the only criteria, and that the ratio of the inductances also has important consequences for steady-state running.

In practical reluctance machines, as the current is increased, the iron begins to saturate and the torque assumes a more linear relationship with the current. In view of the importance of saturation, reluctance torque is predicted at the design stage using computer-based finite element analysis of the magnetic field distribution at each rotor position, and then using 'BIl' or the Maxwell stress method to find the torque. Finite element analysis also enables inductance variation to be obtained, so that the 'coupled circuit' approach (see Section 8.3) can then be used to predict all aspects of performance.

9.3.4 Salient pole synchronous motor

Looking back at Fig. 9.2, it is obvious that both rotors exhibit saliency, although it is much more pronounced in the one on the right. It is therefore to be expected that in addition to the excitation torque that depends on the rotor current (Section 9.3.1), there will also be reluctance torque (Section 9.3.3) which will be there even when there is no rotor current.

If we ignore saturation, the torque-angle relationship is obtained by superposing the excitation torque term and the reluctance torque term, giving the typical torque-angle curve shown in Fig. 9.11. The region to the right of peak torque in Fig. 9.11 is of no practical interest because it represents an unstable condition, but it is included to show one full cycle of the reluctance torque.

The most noticeable effect of the reluctance component is to increase the 'stiffness' (i.e. the gradient of the torque-angle curve) about the zero torque position, but it also reduces the stable operating region lying to the left of the (in this case, slightly increased) peak torque. The increased stiffness is generally beneficial, while the reduction in the stable region is unlikely to be significant because operation at such large load angles would not be possible without exceeding the rated current.

The relative magnitude of the excitation torque and the reluctance torque is a critical design consideration. Today substantial research is being undertaken to

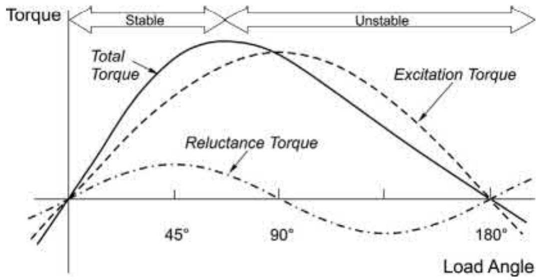

FIG. 9.11 Torque-angle curve for a salient pole excited rotor motor.

reduce the use of permanent magnets, resulting in motor designs with an increasing proportion of the total torque being reluctance torque: this is discussed further in Section 9.8.

9.3.5 Salient permanent magnet motor ('PM/Rel' motor)

In Section 9.3.4 we saw that for an excited rotor motor with rotor saliency, the reluctance torque acting alone caused the unloaded rotor to come to rest with the rotor direct axis aligned with stator m.m.f., i.e. at the same position as it would if the excitation torque acted alone. This is because the low-reluctance axis is the same as the excitation direct axis. The 'stiffness' of the torque-angle characteristic is increased by the presence of the reluctance torque, and depending on the relative magnitudes of the two components, the peak torque may also be increased, as shown in Fig. 9.11, so the combination is an attractive proposition.

The idea of replacing the rotor excitation circuit with simpler permanent magnets while continuing to exploit reluctance torque is clear in principle, but in practice is not as straightforward as might be expected. In order for the magnet flux to flow along the rotor direct axis (i.e. along a salient pole), a gap has to be inserted to accommodate the magnet, and the stronger the magnet, the longer the gap. This greatly increases the reluctance of the direct axis, which is the opposite of what we want in order to maximise the reluctance torque.

However, we have already talked about the move by many industries to reduce their dependency on rare earth magnets because of concerns over the global security of supply. This uncertainty, together with the incentive to provide low-cost motors for the burgeoning mass market (notably in hybrid electric vehicles), has led to renewed interest in motors that combine PM and reluctance torque. Compared with a purely PM motor, the aim is to achieve comparable performance with less magnet material: the ratio of PM torque to reluctance torque varies considerably (typically from 4:1 to as low as 1:1) depending on the detailed motor design and application, but in very simplistic terms, and stating the obvious, the less magnet material the higher the proportion of reluctance torque.

A typical six-pole rotor is shown in Fig. 9.12: it is basically a flux-guided reluctance motor rotor with buried permanent magnets sitting in the flux guides. Looking at the topmost N pole in Fig. 9.12 for example, its two magnets are effectively in series, and their direct (flux) axis is vertical. Apart from the air-gap, the main magnetic circuit external to each pair of magnets is of low reluctance through the core 'iron', so in this regard there is little compromise compared with PM-only design. However, for structural reasons there has to be a bridge of magnetic core material at the outer extremities of the flux guides, and this inevitably offers an attractive short-circuit for some of the magnet flux, which is thus diverted away from its useful path via the stator. This area remains saturated and unproductive in terms of torque.

As far as the reluctance aspect goes, the direct (low inductance) axis is shown by the chain-dotted lines, and the bridge referred to above again represents an unwanted short-circuit path for stator-produced flux, but in essence this is exactly as it would be in a flux-guided reluctance motor. If the reluctance torque acted alone, the unloaded rotor would come to rest with the stator m.m.f. aligned with the chain-dotted line, but if the PM acted alone, the unloaded rotor would come to rest with a N pole aligned with the stator m.m.f. So unlike the salient excited rotor motor, where the equilibrium positions coincided, we now have two distinct zero-torque positions, separated by 90° (elec.).

There is clearly potential confusion over which is the direct axis. On the face of it we have two rival contenders with conflicting claims: the reluctance camp would claim it was the chain-dotted line in Fig. 9.12, while the PM fans would argue it was an axis through the centre of the magnet poles. In practice the latter is usually preferred, i.e. the direct axis is defined in the same way as for a purely PM machine.

We can get a general idea of the shape of the overall torque-angle characteristic by superposing the separate reluctance and PM curves, as shown in

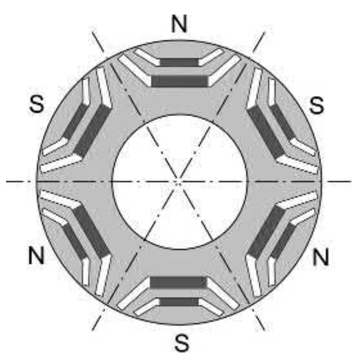

FIG. 9.12 Six-pole PM/Rel motor.

Fig. 9.13, but we have to accept that this is only an approximation because it ignores the effects of saturation in the magnetic circuits.

The excited rotor case is included in Fig. 9.13 for comparison with the PM/Rel motor, and as we have already seen in Section 9.3.4, the resultant torque-angle curve for the excited rotor is stiffer about the stable zero-torque position, and the motoring and braking regions are symmetrically disposed, with equal maximum torque angles for motoring and braking of γ_m and γ_b, respectively.

Our aim is to highlight the fundamental differences between the torque-angle characteristics of the excited rotor motor and the PM/Rel motor, so we have arbitrarily chosen the reluctance and torque components to have the same amplitude. (In practice, an excited rotor motor would have much less reluctance torque, whereas the ratio of torque components for a PM/Rel motor could be higher or lower.)

The shift of 90° between the reluctance and PM curves results in different stable operating regions for the PM case, and also new zero-torque rest positions. The peak motoring and braking torques remain the same as for the excited rotor case, but they are no longer symmetrical about a single equilibrium rest position. The maximum motoring torque angle is indicated by γ_m, while the

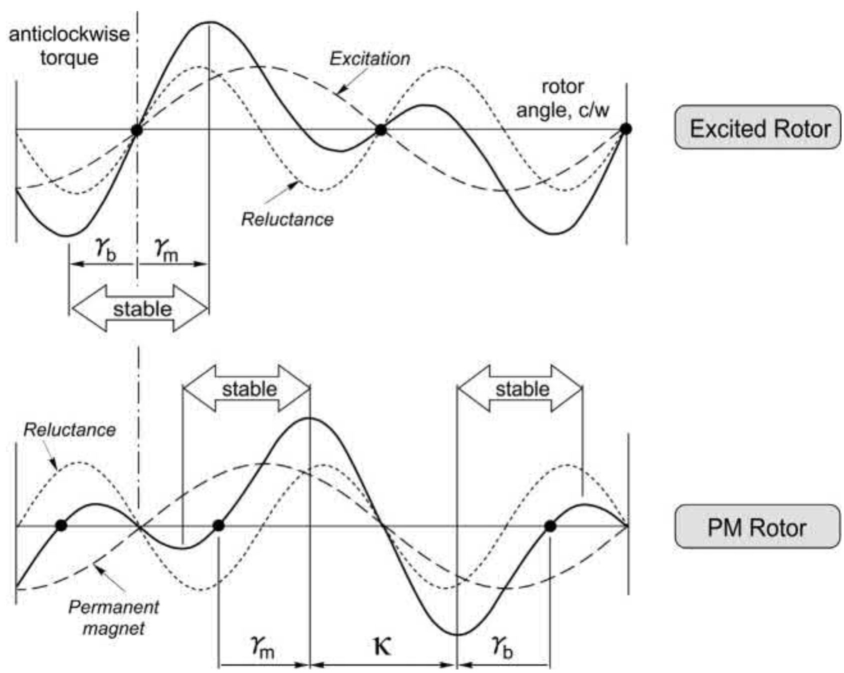

FIG. 9.13 Combined reluctance and PM torques.

maximum braking torque is shown as γ_b. Hence when the drive requires the torque to change from maximum motoring torque to maximum braking torque, the control system (see Section 9.6) will reposition the stator current vector relative to the rotor by the angle κ shown in the lower diagram.

There is still a great deal of ongoing work and interest in this emerging technology, and it will be some time before optimised solutions for the various application areas finally emerge.

9.4 Utility-fed synchronous motors

Historically most synchronous machines were operated directly from the utility supply, so it is appropriate that we look first at the operation of synchronous machines assuming that the supply voltage and frequency are constant.

The physical picture involving spatially distributed m.m.f's, fluxes and current distributions in the previous section provided an explanation for the mechanism of torque production. We now shift attention into the time domain, and explore the steady-state behaviour when a synchronous motor is supplied from a balanced, constant voltage, sinusoidal supply. The stator current, which we held constant in the previous section, will now be one of the principal dependent variables, while the independent variables are the rotor excitation (if a wound rotor type) and the load on the shaft. The speed will be constant, so the power will be directly proportional to the torque.

Fortunately, predicting the current and power-factor drawn from the supply by a cylindrical-rotor or PM synchronous motor supplied from a balanced three-phase supply is possible by means of the per-phase a.c. equivalent circuit shown in Fig. 9.14. To arrive at such a simple circuit inevitably means that approximations have to be made (notably in relation to saturation in the magnetic circuit), but we are seeking only a broad-brush picture, so the circuit is perfectly adequate.

In this circuit X_s (known as the synchronous reactance, or simply the reactance) represents the effective inductive reactance of the stator phase winding; R is the stator winding resistance; *V* is the applied voltage; and *E* is the e.m.f.

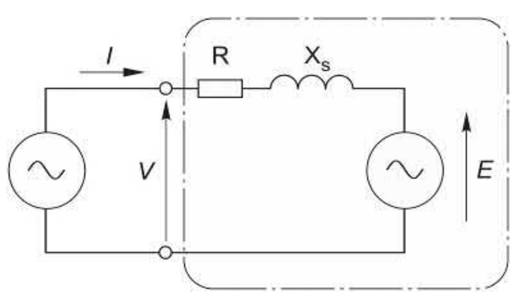

FIG. 9.14 Equivalent circuit for the excited-rotor and PM synchronous machine.

induced in the stator winding by the rotating flux produced either by the d.c. current on the rotor or the permanent magnet.

The term 'effective' reactance applied to X_s reflects the fact that the magnitude of the flux wave produced by balanced currents in each of the three phase windings is 1.5 times larger than the flux that would be produced by the same current in only the one winding. The effective inductance of each phase (i.e. the ratio of flux linkage to current) is thus 1.5 times the self inductance of each phase winding, and consequently the synchronous reactance $X_s = 1.5X$, where X is the reactance of one phase. For the benefit of readers who are familiar with the parameters of the induction motor, it should be pointed out that X_s is equal to the sum of the magnetising and leakage reactances, but because the effective air-gap in synchronous machines is usually larger than in induction motors, their per-unit synchronous reactance is usually lower than that of an induction machine with the same stator winding.

It may seem strange that in the previous section we were talking about both the rotor flux and the stator flux, but now, we refer to the e.m.f. induced by the rotor flux, and seem to have ignored the e.m.f. due to the rotating stator (armature) flux. Needless to say, we have not forgotten the stator flux, because we are following the conventional approach in which the self-induced e.m.f. due to the resultant stator flux (which is proportional to the stator current) is represented by the voltage (IX_s) across the inductive reactance, X_s.

At this point, readers who are not familiar with a.c. circuits and phasor diagrams will inevitably be disadvantaged, because discussion of the equivalent circuit and the associated phasor diagram greatly assists the understanding of motor behaviour. We have included a brief resume of phasors in Section 9.4.3, but we have also summed up the lessons learned in each case, so that readers who have been unable to absorb the theoretical underpinning will not be seriously handicapped.

9.4.1 Excited rotor motor

Our aim is to find what determines the current drawn from the supply, which from Fig. 9.14 clearly depends on all the parameters therein. But for a given machine operating from a constant-voltage, constant-frequency supply, the only independent variables are the load on the shaft and the d.c. current (the excitation) fed into the rotor, so we will look at the influence of both, beginning with the effect of the load on the shaft.

The speed is constant and therefore the mechanical output power (torque times speed) is directly proportional to the torque being produced, which in the steady-state is equal and opposite to the load torque. Hence if we neglect all the losses in the motor (and in particular we assume that the resistance R is negligible), the electrical input power is also determined by the load on the shaft. The input power per phase is given by $VI\cos\phi$ where I is the current and the power-factor angle is ϕ. But V is fixed, so the in-phase (or real) component of input current ($I\cos\phi$) is determined by the mechanical load on the shaft. We recall that, in the same way, the current in the d.c. motor (Fig. 3.6) was determined by the load. This discussion reminds us that although the equivalent circuits in Figs. 9.14 and 3.6 are very informative, they should perhaps carry a 'health warning' to the effect that one of the two independent variables (the load torque) does not actually appear explicitly on the diagrams.

Turning now to the influence of the d.c. excitation current, at a given supply frequency (i.e. speed) the utility-frequency e.m.f. (E) induced in the stator is proportional to the d.c. field current fed into the rotor. If we wanted to measure this e.m.f. we could disconnect the stator windings from the supply, drive the rotor at synchronous speed by an external means, and measure the voltage at the stator terminals, performing the so-called 'open-circuit' test. If we were to vary the speed at which we drove the rotor, keeping the field current constant, we would of course find that E was proportional to the speed. We discovered a very similar state of affairs when we studied the d.c. machine (Chapter 3): its induced motional or 'back' e.m.f. (E) turned out to be proportional to the field current, and to the speed of rotation of the armature. The main difference between the d.c. machine and the synchronous machine is that in the d.c. machine the field is stationary and the armature rotates, whereas in the synchronous machine the field system rotates while the stator windings are at rest: in other words, one could describe the synchronous machine, loosely, as an 'inside-out' d.c. machine.

We also saw in Chapter 3 that when the unloaded d.c. machine was connected to a constant voltage d.c. supply, it ran at a speed such that the induced e.m.f. was (almost) equal to the supply voltage, so that the no-load current was almost zero. When a load was applied to the shaft, the speed fell, thereby reducing E and increasing the current drawn from the supply until the motoring torque produced was equal to the load torque. Our overall conclusion was the simple statement that if E is less than V, the d.c. machine acts as a motor, while if E is greater than V, it acts as a generator.

The situation with the synchronous motor is similar, but now the speed is constant and we can control E independently via control of the d.c. excitation current fed to the rotor. We might again expect that if E was less than V the machine would draw-in current and act as a motor, and vice-versa if E was greater than V. But we are no longer dealing with simple d.c. circuits in which phrases such as 'draw in current' have a clear meaning in terms of what it tells us about power flow. In the synchronous machine equivalent circuit the voltages and currents are a.c., so we have to be more careful with our language and pay due regard to the phase of the current, as well as its magnitude. Things turn out to be rather different from what we found in the d.c. motor, but there are also similarities.

Phasor diagram and power-factor control

We will begin by looking at how the e.m.f. (E) influences the behaviour of the motor when it is running without any shaft load. By excluding one of the two independent variables (i.e. the load torque), we can highlight the role of the other independent variable (the rotor current) in relation to the excitation or flux producing requirement.

If we neglect winding resistance, iron loss and frictional losses, the input electrical power is equal to the mechanical output power, so in this case the input power will be zero, which means that the 'real' component of the phase current (i.e. the component in phase with V) will be zero, and the machine will therefore always be 'reactive' as viewed from the utility supply. The four sketches in Fig. 9.15 show how the input current (and thus the effective reactance) varies with the induced e.m.f. (E): they embody the result of applying Kirchoff's voltage law to the equivalent circuit in Fig. 9.14, i.e. $V = E + IR + jIX_s$, but with R neglected, so the phasor diagram simply consists of the volt-drop IX_s (which leads the current (I) by 90°) added to E to yield V.

In sketch (A), the rotor current (and hence E) is zero, so according to Fig. 9.14 the motor will look like a pure inductance when viewed from the supply, and the current (I_o, the 'o' denoting no-load) will be given by V/X_s, where X_s is the synchronous reactance at the frequency of the supply. The large current (perhaps of the order of 60% of the full-load current in a large motor) is lagging by 90°, so the motor is therefore consuming reactive power only. In this extreme case the motor might be described as un-excited because there is no rotor current, but in fact a rotating field must be set up to induce an internal e.m.f. equal to the terminal voltage, and in this case all of the necessary m.m.f. is provided by the stator current in each phase. We discussed a similar state of affairs when we looked at the unloaded induction motor, where it emerged that the no load current lagged the voltage by almost 90°, and was called the 'magnetising' current, because it was responsible for setting up the rotating flux wave. Here the use of the term 'excitation' harks back to the d.c. machine, but the word 'magnetising' would be equally suitable.

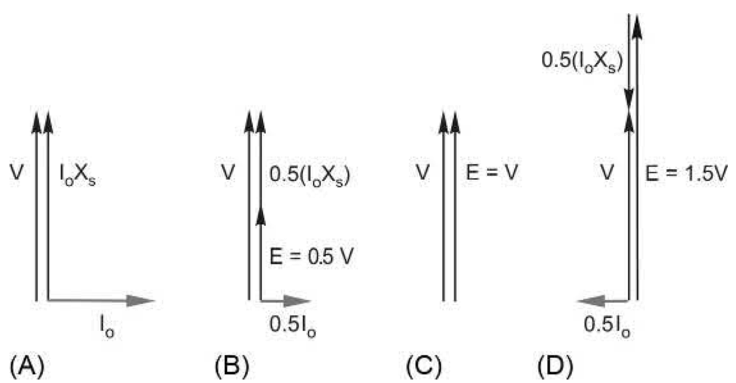

FIG. 9.15 No-load time phasor diagrams showing effect of excitation.

In case (B), the rotor current is sufficient to produce a flux wave that in turn induces an e.m.f. of half the terminal voltage, so now the stator current is only half of that in (A) because it only has to make up the shortfall of excitation. The reactive power is now reduced, but it is still lagging. This condition, where the induced e.m.f. E is less that V, is usually referred to as 'under-excited'.

In the special case shown in Fig. 9.15C, the induced e.m.f. is equal to the applied voltage, the current is zero, and the motor appears as an open circuit to the supply.

Finally, the 'over-excited' case is shown in sketch (D). The rotor excitation is now 50% larger than required to balance the voltage V, so the stator current reverses in phase and now opposes the rotor excitation. The current now leads by 90°, and the motor can equally well be viewed as exporting lagging VArs or consuming leading VArs, i.e. it looks capacitive. (When synchronous machines were operated with both ends of the rotor shaft sealed off, so that no mechanical power was involved, they were known as 'compensators', with the ability to 'look like' either an inductor or a capacitor, and until the advent of power electronics they were used for regulating purposes in power systems.)

We now turn to the more important practical case where the motor is supplying mechanical output power, and again we will investigate how the magnitude of the e.m.f. influences behaviour, with the aid of the phasor diagrams in Fig. 9.16.

The first point to clarify is that our sign convention is that motoring corresponds to positive electrical input power to the machine. The power is given by $VI\cos\phi$, so when the machine is motoring (positive power) the angle ϕ lies in the range $\pm 90^{\circ}$. If the current lags or leads the voltage by more than 90° the machine will be generating.

The sketches in Fig. 9.16 correspond to low, medium and high values of the induced e.m.f. (E), the shaft load (i.e. mechanical power) being constant. As discussed above, if the mechanical power is constant, so is $I\cos\phi$, and the locus

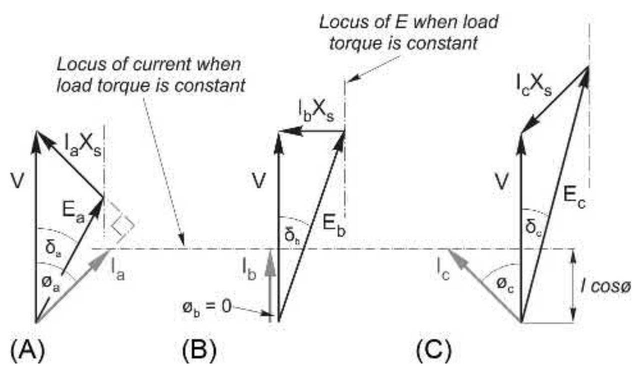

FIG. 9.16 Phasor diagrams for excited rotor synchronous motor operating with constant load torque, for three different values of rotor (excitation) current.

of the current is therefore shown by the horizontal dashed line. The load angle (δ), discussed earlier, is the angle between V and E in the phasor diagram.

Fig. 9.16A represents an under-excited condition where the field current has been set so that the magnitude of (E) is less than V, which leads to a relatively large lagging reactive current component. When the field current is increased (increasing the magnitude of E) the magnitude of the input current reduces and it moves more into phase with V: the special case shown in Fig. 9.16B shows that the motor can be operated at unity power-factor if the field current is suitably chosen.

A brief digression is appropriate at this point to relate the time phasor diagram to the space phasors in Fig. 9.6. We imagine the space phasors to be rotating at the synchronous speed, in which case each of them gives rise to an induced motional e.m.f. in the stator. The rotor space phasor (F_R) produces E (proportional to the d.c. current in the rotor), and the stator space phasor (F_S) produces an e.m.f. (which is proportional to the armature current) that we represent by the voltage IX_s. The resultant of these two space phasors (F) produces the resultant flux linkage at the stator winding (ψ_s), which must induce the terminal voltage, V, so, depending on the magnitude of E, the armature current adjusts accordingly.

The unity power factor case shown in (B) represents the minimum current for the given power (or torque), when the terminal voltage V and frequency are fixed. The corresponding space phasor triangle will have an area determined by the torque, a resultant (F) that is fixed, and F_R adjusted so as to minimise F_S, which will be achieved when the angle α in Fig. 9.6 is 90° and F_S is perpendicular to F. This represents the optimum condition for maximising torque when the resultant and rotor fluxes are specified, and the space phasor diagram will then be a scaled version of Fig. 9.16B.

Returning to Fig. 9.16C, the field current is considerably higher (the 'over-excited' case) which causes the current to increase again but this time the current leads the voltage and the power-factor is $\cos\phi_c$, leading. We see that we can obtain any desired power-factor by appropriate choice of rotor excitation, and in particular we can operate with a leading power-factor, a freedom not afforded to users of induction motors.

When we studied the induction motor we discovered that the magnitude and frequency of the supply voltage V governed the magnitude of the resultant flux density wave in the machine, and that the current drawn by the motor could be considered to consist of two components. The real (in-phase) component represented the real power being converted from electrical to mechanical form, so this component varied with the load. On the other hand the lagging reactive (quadrature) component represented the 'magnetising' current that was responsible for producing the flux, and it remained constant regardless of load.

The stator winding of the synchronous motor is essentially the same as the induction motor, so, as discussed above, it is to be expected that the resultant flux will be determined by the magnitude and frequency of the applied voltage.

This flux will therefore remain constant regardless of the load, and there will be an associated requirement for magnetising m.m.f. But as we have already seen, we have two possible means of providing the excitation m.m.f., namely the d.c. current fed into the rotor and the lagging component of current in the stator.

When the rotor is under-excited, i.e. the induced e.m.f. E is less than V (Fig. 9.16A), the stator current has a lagging component to make up for the shortfall in excitation needed to yield the resultant field that must be present as determined by the terminal voltage, V. With more field current (Fig. 9.16B), however, the rotor excitation alone is sufficient and no lagging current is drawn by the stator. And in the over-excited case (Fig. 9.16C), there is so much rotor excitation that there is effectively some reactive power to spare and the leading power-factor represents the export of lagging reactive power that could be used to provide excitation for induction motors elsewhere on the same system, thereby raising the overall system power factor. As might have been expected, these observations about the role of the excitation line up nicely with what we saw when we looked at the no-load behaviour.

To conclude our look at the excited rotor motor we can now quantify the torque. From Fig. 9.16, the real power is given by

$$W = VI\cos\phi = \left(\frac{V}{Xs}\right) IX_s\cos\phi = \frac{V}{Xs}E\sin\delta = \frac{EV}{Xs}\sin\delta.$$

The speed is constant, so the torque is also given by an expression of the form

$$T \propto \frac{EV}{Xs}\sin\delta.$$

This agrees with the conclusion we reached in Section 9.3.1, where we saw that the torque depended on the product of the resultant field (here represented by V), the rotor field (here represented by E) and the sine of the load angle between them (δ). We note that if the load torque is constant, the variation of the load-angle (δ) with E is such that $E\sin\delta$ remains constant. As the rotor excitation is reduced, and E becomes smaller, the load angle increases until it eventually reaches 90°, at which point the rotor will lose synchronism and stall. This means that there will always be a lower limit to the excitation required for the machine to be able to transmit the specified torque. This is just what our simple mental picture of torque being developed between two magnetic fields, one of which becomes very weak, would lead us to expect.

9.4.2 Permanent magnet motor

Although the majority of permanent magnet motors are supplied from variable-frequency inverters, some are directly connected to the utility supply, and we can again explore their behaviour using the equivalent circuit shown in Fig. 9.14. Because the permanent magnet acts as source of constant excitation, we no longer have control over the magnitude of the induced e.m.f. (E), which now depends on the magnet strength and the speed, the latter being fixed by the utility frequency. So now we only have the load torque as an independent variable, and, as we saw earlier, because the supply voltage is constant, the load torque determines the in-phase or work component of the stator current ($I\cos\phi$) as indicated in the phasor diagrams in Fig. 9.16.

In order to identify which of the three diagrams in Fig. 9.16 applies to a particular motor we need to know the motional e.m.f. (E) with the rotor spinning at synchronous speed and the stator open-circuited. If E is less than the utility voltage, diagram (A) applies; the motor is said to be under-excited; and it will have a lagging power-factor that worsens with load. Conversely, if E is greater than V, (the over-excited case) diagrams (B) or (C) are typical, and the power factor will be leading.

9.4.3 Reluctance motor

Time phasor diagrams for one phase winding under no-load and loaded conditions are shown in Fig. 9.17A and B, respectively, and in each case the time phasors of the flux linkage have been shifted to the right to avoid congestion with the voltages and current. These flux phasors can be compared with the space vectors shown in Fig. 9.10. Once again, resistance has been neglected in the interests of simplicity.

Two novel features of Fig. 9.17 are firstly the resolved current components I_d and I_q, and secondly the incongruity of finding a sketch of a salient rotor in a diagram that supposedly represents time-varying electric circuit quantities. We will therefore break off from the specifics of the reluctance motor, firstly to refresh our understanding of the properties of time phasors, and secondly to justify the presence of the rotor in Fig. 9.17.

The purpose of a phasor diagram is to provide an efficient graphical way of representing the steady-state inter-relationship between quantities that vary sinusoidally in time. We picture all phasors to be rotating anticlockwise at a constant speed and completing one revolution per cycle of the supply. The length of each phasor is proportional to the amplitude (or more usually r.m.s.

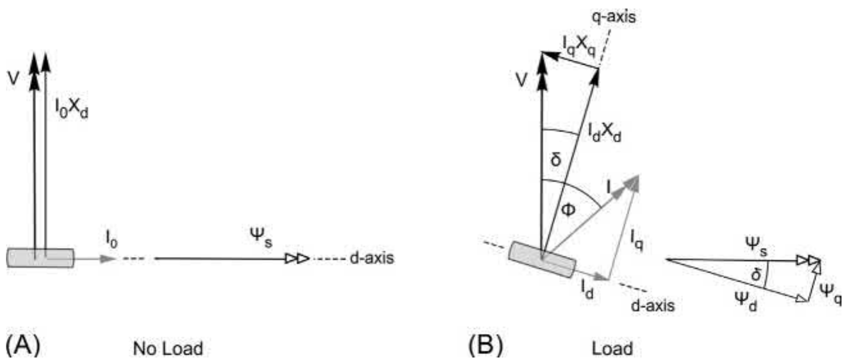

FIG. 9.17 Time phasor diagrams for the reluctance motor.

value) of the quantity represented, and its angular position represents its phase with respect to the other quantities. The projection of the tip of each phasor onto the vertical axis of the diagram then indicates its instantaneous value, which will, of course, vary sinusoidally in time: and if we were to arrange for a pen to be fixed to the tip of each phasor, and for the pen to bear on an endless strip of paper moving at a steady speed from left to right behind the rotating phasor, the trace on the paper would be a sinewave.

In this book, time phasor diagrams represent what is happening inside one of the stator windings, which by definition means that we are in a stationary reference frame, in which two distinct types of quantity may be represented. Firstly, there are the terminal voltages and currents (that we could measure with a voltmeter and ammeter), and the various 'internal to the equivalent circuit' voltages and currents that together make up the terminal quantities. These are single valued *time* functions which have no spatial meaning associated with them. But we also display quantities such as flux linkages that are sinusoidally distributed in *space:* this is justified because the rotating flux linkage space phasor manifests itself in the winding as a sinusoidally time-varying effect, the rate of change of which results in an induced voltage.

For the diagrams to be meaningful, it hardly needs saying that all quantities of the same kind (e.g. voltages) must be drawn to the same scale, while physically different quantities (e.g. currents) can be drawn to a different but consistent scale.

In the light of the above discussion, we would not expect to see the ghostly outline of a salient pole rotor superimposed on a diagram such as that in Fig. 9.17, and we are not about to argue that a rotor is a time-varying quantity. But the rotor does rotate by one pole-pair for each electrical cycle, and we have seen in the discussion of torque production (in Section 9.3.3) that in the steady state the rotor has a fixed angular relationship with the resultant flux phasor, so including it is not so fanciful after all, and as we will now see, it helps us to understand the phasor diagram by locating the direct and quadrature axes, and hence throwing light on the currents I_d and $I_{q.}$

In Section 9.3.3 we defined the low reluctance path through the rotor as the direct axis, and the higher reluctance path as the quadrature axis. Hence when we include the rotor outline in a phasor diagram we implicitly define the direct axis (for example it is horizontal in Fig. 9.17A) and the quadrature axis (vertical). (We adopt the most widely-used convention, with the quadrature axis leading the direct axis by 90°, but readers should not be surprised to find textbooks that use alternatives.)

Returning now to the discussion of the reluctance motor phasor diagram in Fig. 9.17A that represents the unloaded motor, we choose the constant supply voltage V as reference, and follow the usual practice by drawing it vertical.

As we have seen for the induction and synchronous machines, the resultant flux linkage ψ_s is determined by the supply voltage and frequency, and the in-phase component of current ($I\cos\phi$), is solely determined by the load torque.

The reluctance motor has no rotor excitation, so the stator current always has to draw an excitation or magnetising component.

At no load, (Fig. 9.17A, there is no in-phase component of current, and so the no-load current I_o is all excitation or magnetising current, and it produces the flux linkage ψ_s which is in time phase with the current, and in turn induces the e.m.f. which must equal V. We have encountered these flux-e.m.f. relations several times already.

Where there is no saliency, we chose to represent the self-induced e.m.f. by means of the synchronous reactance voltage drop IX_s. But with saliency, we have seen that for a given stator current, the self flux linkage (and hence the inductance) depends on the angular position of the rotor. We therefore introduced two new inductances L_d and L_q to assist analysis. L_d represents the inductance when the direct axis of the rotor is aligned with the m.m.f. axis of the stator phase, and L_q is the inductance when the rotor direct axis is perpendicular to the phase axis. The corresponding steady-state reactances are X_d and X_q, but it is not immediately obvious how these are to be reflected in the phasor diagram. This is where the rotor outline becomes invaluable.

Referring back to Fig. 9.10, we learned that torque is proportional to the sine of (twice) the space angle (δ) between the direct axis and the resultant flux, so we know that at no-load, the flux is along the direct axis. As the flux wave rotates, the rotor remains in synchronism with the flux, so we could legitimately add another 'phasor arrow' to the time diagram labelled 'rotor direct axis': in practice, it is more graphic to put a rotor outline, as shown in Fig. 9.17.

We now know that at no-load, the direct axis of the rotor is aligned with the flux wave, so this determines the rotor outline under no-load conditions as being in time phase with the flux, i.e. horizontal in Fig. 9.17A. The stator m.m.f. is thus entirely directed along the direct axis, and we therefore refer to the corresponding current as the 'direct axis' current. In this special (no-load) case all of the current is d-axis current ($I_0 = I_d$).

The flux produced by the d-axis current is determined by the d-axis inductance, and the e.m.f. induced by the flux is therefore related to the current by the d-axis reactance. The volt-drop that we choose to represent this e.m.f. is the phasor I_dX_d in Fig. 9.17A.

The phasor diagram when the motor is on load is shown in Fig. 9.17B, again with the voltage V as reference, and as at no load, the resultant flux remains the same because it has to induce an e.m.f. equal to V. The resultant flux is no longer on the direct axis, and the rotor is lagging the flux by the load angle (δ). The load torque determines the component of current that is in phase with V (i.e. $I\cos\Phi$—not shown), but the reactive component that finally determines the phase of the current depends on the two reactances, as explained below.

The significance of the current components I_d and I_q should now be clear. The m.m.f. associated with I_d acts along the direct axis of the rotor, and produces a flux linkage ψ_d that is given by L_dI_d, and a corresponding self-induced e.m.f. given by $I_d X_d$. We refer to I_d as the flux producing, magnetising or excitation component of the current, and as we have seen, at no load, all of the input current is direct axis current.

Conversely, I_q is the torque-producing component of current: its associated m.m.f acts along the quadrature axis, producing a flux linkage component ψ_q given by L_qI_q and a corresponding component of self-induced e.m.f. given by I_qX_q.

It is important to reiterate that the time phasor diagram represents voltages and currents that are varying sinusoidally in time, in a stationary reference frame. In particular, the direct and quadrature components I_d and I_q of the phase current that we have just discussed are not to be confused with the transformed variables i_d and i_q in the rotating reference frame that we discussed in Chapter 8. The latter are, under steady-state conditions, constant ('d.c.') currents that produce at the rotor the same m.m.f. that is produced by the a.c. currents in the stator windings. We will return to the transformed currents later when we deal with field-oriented control.

To obtain the output power per phase from the phasor diagram is a straightforward but tedious process, the result of which is that the power is given by

$$P = \frac{V_S^2}{2X_dX_q}\left(X_d - X_q\right)\sin 2\delta.$$

The speed is constant, so this expression also represents the relationship between the torque and the load angle, δ. Physically, as we have seen, δ is the angle between the resultant rotating field and the direct axis of the rotor, i.e. the 'lag' of the rotor with respect to the resultant field, and it therefore bears some similarity with the load angle of motors that have rotor excitation.

The static analysis in Section 9.3.3 showed that the torque was proportional to $(L_d - L_q)$ and sin 2γ, where γ is the torque angle, whereas here we find that the torque is proportional to $\left(\frac{1}{L_q} - \frac{1}{L_d}\right)$ and $\sin 2\delta$, where δ is the load angle. The differences stem from our choices of the contrasting constraints between the static and steady running cases: in the former, the current was kept constant, and the resultant flux varied according to the rotor angle, while in the latter, the voltage and frequency force the resultant flux to be constant, and the current has to adjust accordingly.

9.4.4 Salient pole motor

Modelling the salient pole motor is based on an equivalent circuit similar to that for the smooth rotor motor in Fig. 9.14, with an induced e.m.f. E due to the rotor excitation. But because of the saliency and the consequent reluctance component of torque, the single synchronous reactance (X_s) has to sub-divided into a direct-axis reactance (X_d) and a quadrature axis reactance (X_q). These reactances are the same as we discussed in the previous section.

The time phasor diagram is essentially a combination of those shown in Figs. 9.16 and 9.17, and again resistance has been neglected. Constructing

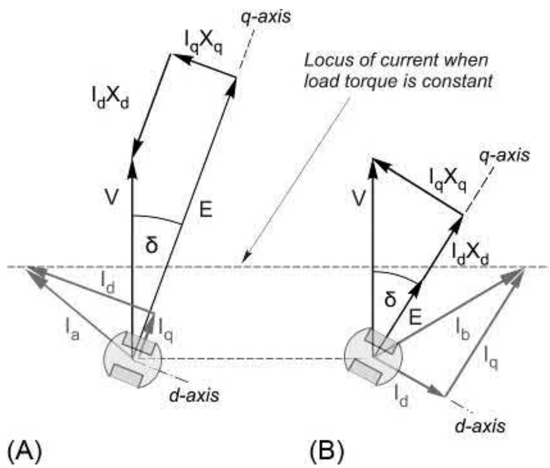

FIG. 9.18 Time phasor diagrams for salient pole excited rotor motor.

the diagram involves resolving the current into direct and quadrature axis components, before the reactive volt drops I_dX_d and I_qX_q can be identified. We will not go into detail, but typical phasor diagrams for over-excited and under-excited conditions are shown in Fig. 9.18A and B, respectively, both sketches being for the same output power or torque.

In the over-excited case, the current is leading, i.e. lagging Vars are being exported, whereas in the under-excited case additional lagging Vars are needed to supplement the excitation provided by the rotor current.

The power and torque can be derived from Fig 9.18 in terms of the controllable variables (V, E, and the load torque), but the lengthy manipulations are not necessary and we will simply quote the well-known result that the torque is given by the expression

$$T \propto \frac{EV}{X_d}\sin\delta + \frac{V^2}{2X_dX_q}\{X_d - X_q\}\sin 2\delta.$$

The first term is the same as that for the round-rotor excited machine that we saw earlier, but with X_d replacing X, and the second term is the same as we found for the reluctance motor. We note that, even when the rotor excitation is zero (i.e. $E=0$), the salient pole motor can produce torque by reluctance action alone, but needless to say, the reluctance torque depends on the degree of saliency: a motor with a rotor such as that in Fig. 9.2A might produce 5% of its torque via reluctance action, whereas the reluctance term for the type of rotor in Fig. 9.3B could easily contribute 30% or more.

9.4.5 Starting on utility supply

It should be clear from the discussion of how torque is produced, that unless the rotor is running at the same speed as the rotating field, no steady torque can be produced. If the rotor is running at a different speed, the two fields will be sliding past each other, giving rise to a pulsating torque with an average value of zero. Hence a basic synchronous machine is not self-starting, and some alternative method of producing a run-up torque is required.

Most line start synchronous motors, designed for direct connection to the utility supply, are therefore equipped with some form of rotor cage, similar to that of an induction motor, in addition to the main field winding. When the motor is switched onto the supply, it operates as an induction motor during the run-up phase, until the speed is just below synchronous. The excitation is then switched on, and as long as the load is not too high, the rotor is able to make the final acceleration and 'pull-in' to synchronism with the rotating field. Because the cage is only required during starting, it can be short-time rated, and therefore comparatively small. Once the rotor is synchronised, and the load is steady, no currents are induced in the cage, because the slip is zero. The cage does however come into play when the load changes, when it provides an effective method for damping out the oscillations of the rotor as it settles at its new steady-state load angle.

Large motors will tend to draw a very heavy current during run-up, perhaps 6 or more times the rated current, for many tens of seconds, or longer, so some form of reduced voltage starter is often required (see Chapter 6). Sometimes, a separate small or 'pony' motor, is used simply to run-up the main motor prior to synchronisation, but this is only feasible where the load is not applied until after the main motor has been synchronised.

9.5 Variable frequency operation of synchronous motors

Just as we have seen in Chapters 7 and 8 for the induction motor, once we interpose a power electronic converter between the utility supply and the machine we introduce new levels of performance and lose most of the inherent drawbacks which we find when the motor is directly connected to the utility supply.

Most obviously, a variable frequency converter frees the synchronous machine from the fixed-speed constraint imposed by utility-frequency operation. The obvious advantage over the inverter-fed induction motor is that the speed of the synchronous motor is exactly determined by the supply frequency whereas the induction motor always has to run with a finite slip. On the down side, we lose the ability of the excited rotor motor to vary the power factor as seen by the utility supply.

In principle, a precision frequency source (oscillator) controlling the inverter switching is all that is necessary to give precise speed control with a synchronous motor, while speed feedback is essential to achieve such accuracy with an induction motor. In practice however, we seldom use open-loop control, where the voltage and frequency are generated within the inverter, and are independent of what the motor does. Instead, field-oriented control, almost identical to that described for the inverter-fed induction motor, predominates (see

Section 9.6). The principal advantage of FOC is that it allows us to control the torque and flux components of the stator current independently, and in the case of the PM motor it prevents the motor from losing synchronism with the travelling field by locking the supply frequency to the speed of the rotor.

However, in the steady state, an observer looking at the stator voltage and current would see steady state sinusoidal waveforms, and would be unaware of the underlying control mechanism. We can therefore study the steady-state behaviour using the equivalent circuit in much the same way as we did with the utility-fed PM motor. We will continue to ignore resistance because this makes the phasor diagrams much simpler to understand without seriously compromising our conclusions.

In this section we again use the phasor diagram showing voltages and current, but we also include the time varying fluxes in order to emphasise the link between the two sets of variables.

As before, we imagine the flux produced by the magnet and the flux produced by the stator as if they existed independently, although in reality there is only one resultant flux. We have seen previously that because the fluxes rotate in synchronism, the magnitude of the torque depends on the product of the two field strengths and the angle between them: when aligned, the torque is zero, and when perpendicular, it is maximum. This is equivalent to saying that the torque is maximum when the stator current wave is aligned with the magnet flux wave, which is the traditional '*BIl*' picture.

When we discussed the various utility-connected motors in Section 9.4, we saw that because the voltage and frequency were fixed, the resultant (stator) flux was constant. For the excited-rotor motor, the rotor current could be adjusted to achieve a power factor of unity, but for the PM motor, in which the rotor excitation is constant, the stator current adjusted itself to satisfy the requirement for the resultant flux to be constant, and as a result we had no control over the power factor.

With an inverter-fed motor we gain control of both the stator voltage and frequency, so that together with the load torque we now have three independent variables in the case of the PM motor, or four for the excited rotor machine. Despite the recent re-emergence of the inverter fed reluctance motor, the majority of inverter-fed synchronous motor drives, employ PM motors, so we will concentrate on their behaviour for the remainder of this section.[4]

9.5.1 Phasor diagram of PM motor

The general diagram (Fig. 9.19) is for an under-excited case, i.e. at the speed in question, the open circuit e.m.f. (E) is less than the terminal voltage. The red and

4. Those readers who want to consider variable frequency operation of other synchronous motor types will be able to develop the necessary understanding by referring to the appropriate phasor diagrams described in Section 9.4.

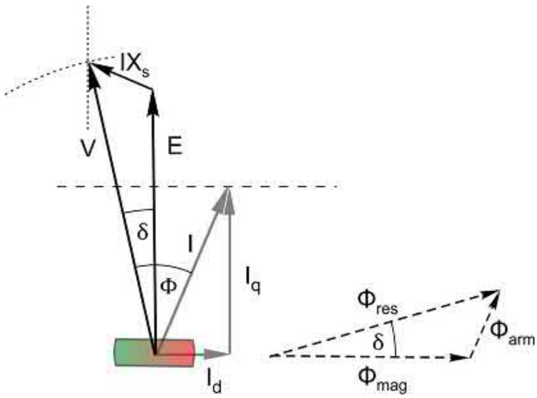

FIG. 9.19 Phasor diagram for the permanent magnet synchronous motor.

green regions within the stylised rotor outline indicate the poles of the magnet, and thus define the direct axis (horizontal) in the diagram. The magnet flux (Φ_{mag}) remains constant, so we choose it as our horizontal reference throughout this and the following sections.

Fig. 9.19 is similar to Fig. 9.16, but we have chosen to add the flux phasors in order to lay the groundwork for the later discussion of field-oriented control. As explained below, the flux and voltage triangles are similar, and each flux is proportional to the corresponding m.m.f. (see Fig. 9.6): for example, the rotor flux (φ_{mag}) corresponds with the rotor m.m.f. (F_R in Fig. 9.6).

E is the open circuit e.m.f. produced by the magnet flux (φ_{mag}): it is proportional to the magnet flux and the speed, which is proportional to the stator frequency, ω. Once we specify the frequency the magnitude of E is known, so we can start the left hand side of the diagram with E drawn vertically because, as previously discussed, the induced e.m.f. leads the magnet flux by 90°.

φ_{arm} is the flux that would be set up if the armature (stator) current existed alone. We are assuming no saliency, so the magnitude and direction of this flux depend only on the magnitude and phase of the stator current (I).

φ_{res} is the resultant ('the') flux. With resistance ignored, this flux is effectively determined by the applied voltage and frequency, as we have seen several times earlier in this book, so in the phasor diagram, the magnitude of this flux is proportional to, and in quadrature with, V.

As usual, the self-induced e.m.f. from the stator flux is modelled by means of the voltage across the effective stator inductive reactance, X_s, so in the phasor diagram the voltage IX_s leads the current by 90°.

V is the applied stator voltage, here regarded as an independent variable, and I is the resulting stator current. The stator phasor equation is $V=E+IX_s$ (the terms being vectors, of course), or $V=E+jIX_s$ for those who prefer the complex notation.

The other independent variable is the load torque, which is our next consideration in developing the diagram. We dealt with the same matter in relation to Fig. 9.16, where the utility supply voltage was constant. But now the voltage is not fixed, so we will take a different approach that focuses instead on how we need to control the current in order to achieve the required torque in the most efficient way. This will assist us in understanding the strategy employed in the torque control loop of a PM motor drive, which we look at later in this chapter.

In the steady state, the motor torque must equal the load torque. The motor torque is proportional to the product of the current (I), the magnet flux (φ_{mag}), and the sine of the angle between them (δ), or in other words to the product of the magnet strength and the component of current at right angles to the magnet flux. Hence the load torque determines the torque component of the stator current (I_q), as shown in Fig. 9.19. The locus of the current I is then the horizontal dashed line, and the locus of V is the vertical dotted line as shown in Fig. 9.19. So when we finally specify the magnitude of V (shown by the dotted arc), the intersection of the arc and vertical line fixes the phase of the stator current, thereby specifying the direct axis component I_d and finalising the diagram.

The current component in phase with E is the useful or torque component I_q, while the component in phase with the magnet flux is the flux component I_d. The area of the flux triangle is proportional to the product of I_q and φ_{mag}, and thus the area provides an immediate visual indication of the torque.

If we adjust V so that the stator current is in phase with E, (i.e. the flux component I_d is zero), we get the maximum torque per ampere of stator current, and thus the minimum stator copper loss for that particular torque: if iron loss is ignored, this would be the most efficient condition.

It is important to stress that the conclusion that we have reached—that the torque is proportional to the quadrature component of current - is not restricted to the steady-state condition that we have assumed so far, and is in fact the basis of the field-oriented control system that we will discuss in Section 9.6.

Alternative expressions for the power per phase can be derived from the geometry of the phasor diagram with the aid of a few construction lines. The first is identical to that obtained for the excited rotor motor, i.e. $P = \frac{EV}{Xs}\sin\delta$, while the second, which is more appropriate when the current is being controlled directly, is $P = EI_q$. The second equation is of exactly the same form as we found for the d.c. machine in Chapter 3.

Before we move on, we should note that because we looked at an example where the applied voltage is just a bit larger than E, the current turned out to be of modest amplitude and at a reasonable power-factor angle. However, if we had specified a much higher V, we would find that the current would have had a much larger lagging flux component (with the same torque component, determined by the load), and a much worse power factor. And conversely, a much smaller applied voltage would result in a large leading power-factor current. This behaviour is in line with what we have already seen, and neither condition is desirable because of the increased copper loss.

9.5.2 Variable speed and load conditions

In Section 9.7 we will look at the torque-speed and performance capabilities of the inverter-fed PM synchronous motor, and we will find that, as with the d.c. drive and the induction motor drive there is a so-called constant torque region that extends up to base speed, within which full rated torque is available on a continuous basis. And in common with the other drives, at higher speeds there is a field weakening region where the maximum available torque is reduced. However, the usual constraints imposed by the maximum supply voltage and allowable continuous stator current may result in more serious restrictions on operation in this region than we have seen with other drives, including the likelihood that only intermittent operation may be possible.

Given the number of parameters involved and their variation between motor designs, it is only possible for us paint a broad picture, so we will look at one hypothetical machine and use it to provide an insight into some of the issues involved. We will focus on what can be learned from the phasor diagrams for three steady-state conditions, two in the constant-torque region and one in field weakening. The lessons learned will prove invaluable when we move on to study field-oriented controlled drives.

For convenience in calculations in relation to the steady-state phasor diagrams we will take the open-circuit e.m.f. E at the base speed (ω_B) to be 1 p.u. (In practice, the rated applied voltage is usually taken as 1 p.u., but there is no reason why any other value should not be chosen.) We will assume that the reactance of the winding (X) at the base speed is 0.3 p.u. (which means that at rated current (and rated frequency) the voltage across the reactance is 0.3 times the rated voltage).

Full load (full torque at base speed)

By full load we mean that the machine is running at base speed and delivering its full rated torque. We saw above that the current depended on the applied voltage, and that in particular if we apply the right voltage we can minimise the current. This is what we do in field oriented control, so the diagram (Fig. 9.20) represents this condition.

We should note that the stator current has zero flux component, and so the armature flux is perpendicular to the magnet flux, i.e. in the optimum torque-producing position. We will define the current in this situation as 1 p.u., so the volt-drop across the reactance is $1 \times 0.3 = 0.3$ p.u. From the property of the right-angled triangle, V turns out to be 1.04 p.u., and because we have defined this as the rated power condition, we will regard 1.04 V as the maximum voltage that the inverter can produce. If iron loss is neglected, this condition is the most efficient state for the given current and torque, because the stator copper loss is minimised.

Full torque at half base speed (half power)

The phasor diagram is shown on the right hand side of Fig. 9.21, with the voltage optimised again to give the most efficient stator current: the full load diagram. (Fig. 9.20 is repeated on the left hand side to make comparison easy.)

The magnet flux is the same as in Fig. 9.20 of course, and so for the same (rated) torque the stator current has to be 1 p.u., as before. However the frequency is now only half ($\omega_B/2$) so the open-circuit e.m.f. is now reduced to 0.5. The stator reactance is proportional to frequency, so it has also halved, to 0.15 p.u., and the volt-drop IX becomes 0.15 p.u. also. In view of the similarity between the diagrams it is no surprise that the applied voltage turns out to be half, i.e. 0.52 p.u., and given the emphasis we have previously placed on the fact that the flux in an a.c. machine is determined by the V/f ratio, we will not be surprised to see that this is again confirmed by these results.

The angle θ between V and I (the power-factor angle) is the same at full and half speed, and since the input power is given by VI cos θ, it is clear that the input power at half speed is half of that at full speed. This is what we would expect because with resistance neglected the input power equals the mechanical power, which is half in Fig. 9.21 because the torques are the same but the speed is half of that in Fig. 9.20.

The two previous examples have shown how the motor can be operated to produce full rated torque up to a speed at which the full available voltage is applied, i.e. this is what we have referred to previously as the 'constant torque region'. In this region, the PM motor with field oriented control behaves, in the steady state, in a very similar manner to a dc motor drive, in that the applied voltage (and frequency) are proportional to the speed and the stator current is proportional to the torque.

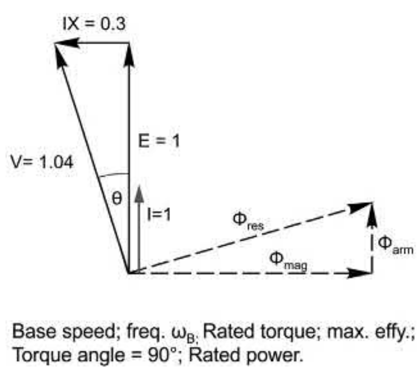

FIG. 9.20 PM motor phasor diagram—field oriented control—full (rated) torque at base speed.

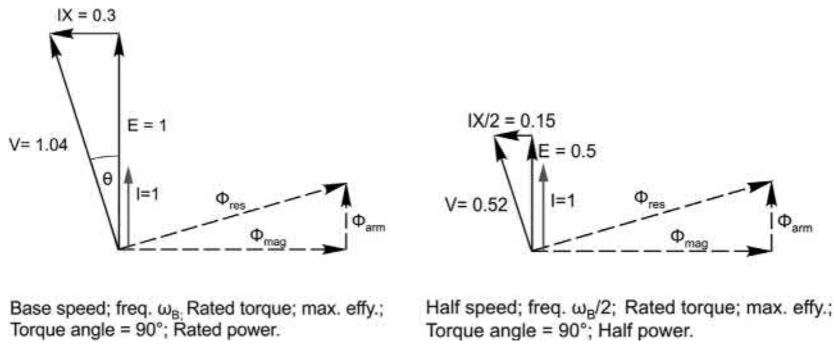

FIG. 9.21 PM motor phasor diagram—field oriented control. Full rated torque at base speed (left), Full rated torque at half base speed (right).

We saw with the dc drive and the induction motor drive that once we had reached full voltage and current, any further increase in speed could only be achieved at the expense of a corresponding reduction in torque, because the input power was already at its maximum or rated value. In both cases, higher speeds were obtained by entering the aptly named 'field weakening'region.

For the dc motor, the field flux is under our direct control, so we reduce the current in the field circuit. In the induction motor, the field is determined indirectly by the V/f ratio, so if f increases while V remains constant, the field flux reduces. The PM motor behaves in a somewhat similar fashion to the induction motor, in that if the voltage is constant at speeds above base speed, the V/f ratio reduces as the frequency (speed) is increased, so the resultant flux also reduces. However, whereas the only source of excitation in the induction motor is the stator winding, the magnet in the PM motor remains a constant (and potent!) source of excitation at all times, and so we can anticipate that in order to arrive at a reduced flux to satisfy the V/f condition, the stator current will have to nullify some of the influence of the magnet. We must therefore expect less than ideal behaviour in the field-weakening region, which we now examine.

Field weakening—Operation at half torque, twice base speed (full power)

We will consider a situation well into the field weakening region, i.e. we will assume that in line with other drives we can expect full power, and so settle for operation at twice base speed and half rated torque. We will see that while this is achievable on an intermittent basis, it is not possible without exceeding the rated current, a conclusion that we must expect to apply to the whole of the field-weakening region with a PM drive.

The condition is represented by the phasor diagram shown on the right hand side of Fig. 9.22: the full load, base speed diagram is repeated on the left hand side to make comparison easy.

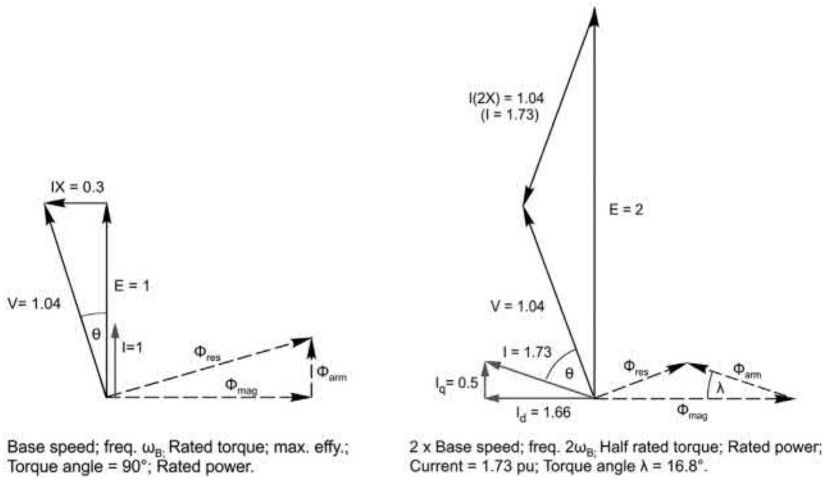

FIG. 9.22 PM motor phasor diagram—field oriented control. Full rated torque at base speed (left), half rated torque at twice base speed (right).

The frequency is twice the base value, i.e. $2\omega_B$ and so the open circuit e.m.f. is 2 p.u. As in other drives, above base speed we apply the maximum possible stator voltage, in this case 1.04 p.u. Because the voltage cannot be greater, it is not possible for us to arrange for the current to be in phase with E, and we are obliged to settle for a rather poor second-best.

The load torque is half rated, which means that the torque component of the current is 0.5 p.u. However, we note that the very large difference in voltage between E and V results in a very large stator volt-drop of $IX_s = 1.04$ p.u. The reactance is twice as large as at base speed because the frequency is doubled, so the current is given by $1.04/0.6 = 1.74$ p.u., which is 74% above its continuously rated value. The copper loss will therefore be increased by a factor of 1.74^2, i.e. to three times the rated value. Clearly continuous operation will not be possible because the stator will overheat, so in this field-weakening condition only intermittent operation at half torque will be possible.

By comparing the flux triangles in Fig. 9.22 we can see why the condition shown on the right side is referred to as field weakening. At base speed (the left hand figure) our freedom to adjust V allowed us to ensure that the stator or armature flux is in quadrature with the magnet flux (i.e. the torque angle λ is 90°), leading to a slightly higher resultant flux, and maximising efficiency by minimising the current for the given torque.

In contrast the constraint on V means that at twice base speed the armature flux has a large component which is in opposition to the magnet flux, leading to a resultant flux that is much less than the magnet flux, and a very low torque angle (λ) of only 16.8°. In a sense therefore, most of the stator current is 'wasted' in being used to oppose the magnet flux. Clearly this is not a desirable operating condition, but it is the best that we can get with the limited voltage at our disposal.

As a final check we can calculate the input power. The angle θ is given by $90° - 2\lambda = 56.4°$, so the input power (VIcosθ) is $1.04 \times 1.74 \times 0.553 = 1$ p.u., as expected since we assumed that the mechanical power was at rated value (twice base speed, half rated torque) and we ignored resistance in our calculations.

At this point we should recall that the aim throughout Section 9.5 has been to discover how the motor parameters determine the steady-state behaviour when the voltage, frequency and load vary over a range that is representative of a typical inverter-fed drive. Few readers will find it necessary to retain all of the detail (although a general awareness of the broad picture is always helpful) so those who have found it hard going can take comfort from the fact that in practice the drive will take care of everything for them. This is discussed in the next section.

9.6 Synchronous motor drives

Inverter fed operation of synchronous machines plays a very important and growing role in the overall drives market, as customers seek higher efficiency and higher power density than can be achieved with induction motors. We will see later, when we consider some of the available PM motor designs and control strategies, that very high dynamic performance can be achieved, making this the motor of choice in many of the most demanding applications.

9.6.1 Introduction

In Sections 9.4 and 9.5 we looked at the steady-state behaviour of the various types of synchronous motor when supplied with balanced sinusoidal voltages, and over a range of frequencies, when both voltage and frequency were independent of what was happening in the motor.

We saw that, as expected in a motor that has to compete with the induction motor, the synchronous motor has the inherent ability to accommodate to load changes: if the shaft load is increased, the rotor slows momentarily until the load angle and in-phase component of current have increased to produce more torque, and allow it to regain a new steady-state at the original speed. However, we also saw that under these ‘voltage fed’ conditions the maximum torque occurs when the load angle reaches 90°, and if it exceeds this value, the rotor will drop out of synchronism, the average motor torque will then be zero, and the motor will stall. It is easy to imagine that loss of synchronism may also occur if the frequency is increased too quickly: as the resultant field begins to accelerate, the load angle (between the resultant field and the rotor direct axis) initially increases, so the torque rises and the rotor begins to accelerate: but, depending on the inertia, if the load angle exceeds the point of maximum torque (90°), the rotor will ‘pole slip’, and stall. Loss of synchronism is clearly not acceptable for a general-purpose drive, which explains why early inverter-fed synchronous motor drives were not widely used.

Fortunately all the shortcomings of the utility-fed motor can be completely avoided by making the supply to the motor terminals respond almost instantaneously to what is happening inside the motor. As with the induction motor drive, this only became possible when fast digital processing could be harnessed to allow rapid control of the Inverter.

We will see that in these so-called ‘self-synchronous’ drives, the rotor is incapable of losing synchronism and stalling because the switching pattern of the inverter (and hence the frequency and speed) is determined by the rotor position and not by an external oscillator. We will also see that field oriented control can be readily applied to synchronous machines to achieve the highest levels of performance and efficiency with machines which have higher inherent power densities than the equivalent induction motors.

The reader may have noted that the stator current, which had been central to the discussion of torque production in Section 9.3, was relegated to being a dependent variable in Sections 9.4 and 9.5. This was because, having specified the voltage, frequency, and load, we had no direct control of the current, the magnitude and phase of which were obliged to assume the unique values that produced the required torque. In contrast, all of the drives that we look at in this chapter have current controllers at the heart of their torque control system, so we will see that current re-asserts its importance, and we will then be able to make use of many of the ideas relating to torque production that were introduced in Section 9.3. (We can continue to make use of the lessons we learned from the phasor diagrams in the previous sections, but they will only apply after the drive has reached a steady state in which the terminal voltage and frequency are constant.)

The reader will also be pleased to know that we will find many comforting similarities with d.c. and induction motor drives, but in relation to the sign of the torque, the synchronous machine is very different from the d.c. or induction motor. In the d.c. machine, the torque depends on the polarity of the armature current; in the induction motor the deciding factor is the whether the slip is positive or negative; but in the synchronous motor, it is the sign of the torque angle that matters, as shown in Fig. 9.23.

Fig. 9.23 shows in simplified form the rotor, the stator and rotor m.m.f. space waves (F_s and F_r, respectively), and the corresponding sinusoidally distributed stator current distribution (the latter being very crudely represented by a single-turn coil). The sketch is equally valid as a freeze-frame view of synchronous operation, when all of these elements are rotating anticlockwise at the synchronous speed (in which case the rotating stator current distribution is the resultant of three stationary windings with alternating currents) or with the machine at standstill with each (d.c.) phase current frozen at its appropriate value.

Recalling that the two m.m.f. waves always try to align themselves (or by applying ‘BIl’), we see that in the sketch on the left, the rotor torque is anticlockwise, because the stator m.m.f is displaced by a positive torque angle (γ_m) with

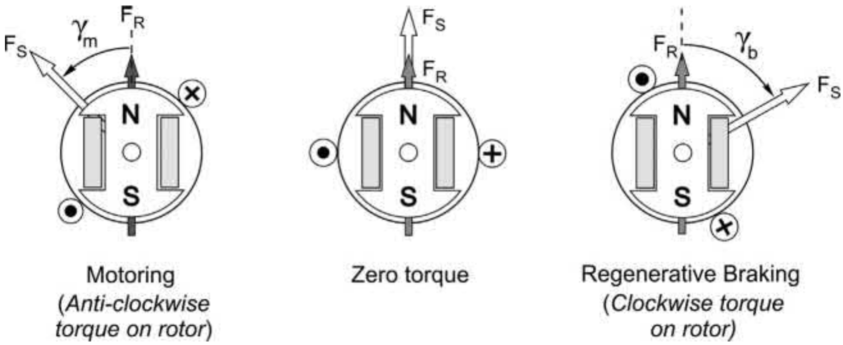

FIG. 9.23 Synchronous motor showing (left) torque angle (γ_m) when motoring, and (right) when braking (γ_b).

respect to the rotor axis, while if the stator m.m.f. is to the right of the rotor axis the torque angle (γ_b) is negative and the torque on the rotor is clockwise. It should be clear that when the motor current (rather than the voltage) is the independent variable, the torque angle (between the stator and rotor m.m.f.'s) becomes of central interest, rather than the load angle.

Under running conditions, with the rotor turning anticlockwise, the left hand sketch represents motoring, the rotor lagging behind the stator m.m.f./flux wave, with energy being converted from electrical to mechanical form. Our concern in what follows will be to explain how we supply and control the magnitude and position of the stator current wave with respect to the rotor axis in order to achieve the desired torque. In most cases where full torque is required (for example during acceleration) the torque control system will supply each phase with rated current, and the torque angle will be maintained at or close to 90°, so that as the motor picks up speed, the period of each complete electrical cycle progressively reduces and the rotor therefore remains synchronised throughout.

We mentioned above that Fig. 9.23 applies equally when the motor is running at a steady speed, or is accelerating or decelerating, or is at rest, so it will probably be a helpful image to bear in mind for the remainder of this chapter.

We will now look at the power electronic converters and control strategies used with the various types of synchronous machine, and in the same order as in previous sections of this chapter. The last three all use pulse width modulated voltage source inverters (VSI), but the first - the excited rotor machine – is often supplied from a current source inverter (CSI) because the latter is better suited to the very high (multi-megawatt) power levels involved.

9.6.2 Excited rotor motor

Applications for this type of drive fall into two main categories.

Firstly, as a starting mechanism for very large synchronous machines, to bring them up to speed prior to synchronising to the utility supply, the converter then being rated for only a fraction of the machine rating. The main motor is started off load, synchronised to the utility supply and the load is then applied.

Secondly, as large high power (and sometimes high speed) variable speed drives for a variety of applications. Power ratings, typically from 2 to 100 MW at speeds up to 8000 rpm are available. At these high powers, it is advantageous to increase the operating voltage in order to minimise the current and thereby make the windings and interconnecting power cables more manageable: supply voltages up to 12 kV are typical, but systems over 25 kV are in service in which case the high voltage converter technology is similar to that used for HVDC power-system converters.

Some manufacturers of synchronous machines of much more modest ratings (e.g. a few tens of kW), also offer low voltage, thyristor based converter technologies, but they tend to be niche products.

The detailed design of high power drives, and the detailed consideration of the impact on the utility supply are beyond our scope, but we can cover the main principles without too much difficulty.

The power circuit and basic operation

The basic components of the high power, high voltage drive system are shown in Fig. 9.24.

The two fully-controlled converters are connected via a d.c. link that includes a large inductance to ensure continuous current operation. Current in the top of the d.c. link always flows from left to right, (as shown by the direction of the thyristors), but, as we saw in Chapter 2, the link voltage can be positive or negative (see below) so that energy can flow in either direction. The direction of rotation of the machine is determined electronically by the switching sequence within the machine converter, and hence full four-quadrant operation is available without extra hardware. The labels 'rectifier' and 'inverter' in Fig. 9.24 indicate how each converter operates when the machine is operating as a motor, their roles being reversed when the machine is braking or generating.

In view of what we have read so far, we might have expected the switching devices in the motor inverter to be IGBT's, and to see PWM control of the output voltage waveform, but in these large sizes, IGBT devices are expensive, and

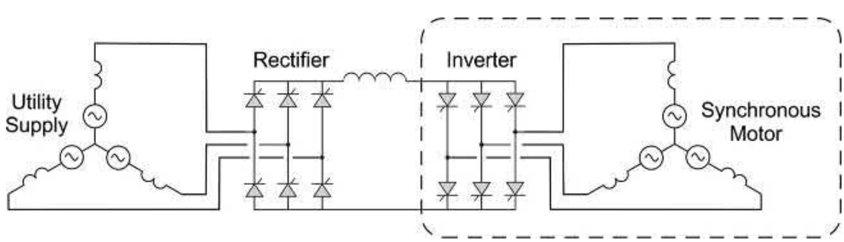

FIG. 9.24 Converter-fed large excited-rotor synchronous machine drive.

this explains why the motor converter is exactly like that used on the supply side of a d.c. motor drive.

The reason why we can use this cheaper inverter lies in the fact that, once rotating, a synchronous machine generates a.c. voltages in each phase winding that facilitate the commutation of a converter connected to its terminals. This is why it is frequently referred to as a Load Commutated Inverter (LCI). In effect the motor converter behaves in the same way as it does when connected to the utility supply.

There are close parallels between the combination of the motor side inverter and synchronous motor (enclosed by the dashed line in Fig. 9.24) and the conventional brushed d.c. motor, which explains why the combination is sometimes referred to as an 'inside-out' d.c. machine. We will explore the similarity, because we will then find it instructive to help us to understand how this particular drive operates.

In the synchronous motor the moving rotor carries the d.c. field winding. The rotating flux induces a sinusoidal e.m.f. in the three stator windings, the magnitude and frequency being proportional to the rotor speed and field current. The firing of the devices in the inverter is synchronised to these induced e.m.f.'s (and hence to the rotor position) so that the link current is fed sequentially into each phase at the optimum rotor angle to maximise torque, the inverter effectively acting as an 'electronic commutator'. The d.c. link therefore 'sees' a rectified 3-phase voltage that, although unidirectional, is not smooth d.c. (see Fig. 9.26): it can be thought of as the equivalent of the 'back e.m.f.' in a brushed d.c. motor.

Conversely, in the conventional d.c. motor, the field is stationary, and its flux induces a motional alternating e.m.f. in the (many) armature coils on the rotor. The mechanical commutator and its sliding brushes rectify the e.m.f. so that at the armature terminals we get a very smooth d.c. induced voltage that we refer to as the back e.m.f, and denote by the symbol E.

So when viewed from the d.c. link, the two are essentially the same. It is therefore not surprising that, as with a d.c. motor, the no-load speed of the synchronous motor depends on the d.c. link voltage provided by the supply-side converter, while when load is applied and speed tends to fall, reducing the induced e.m.f., the d.c. link current automatically increases, thereby producing more torque until the steady-state is reached. The effects of varying the d.c. excitation on the rotor of the synchronous machine also mirror those in the d.c. motor, so that field weakening will lead to higher speed but reduced torque for a given link current.

Returning to the synchronous motor itself, and assuming that the link current is sensibly constant (see below), the switching strategy of the motor converter is synchronised to the induced e.m.f. of the machine and hence to the rotor position. In the steady state, each phase conducts the link current (I_{dc}) for one third of a cycle of the generated e.m.f., i.e. 120^o (elect); it is zero for the next 60°; and the pattern is repeated in the negative half-cycle, giving the current waveforms as shown in Fig. 9.25. (These are idealised waveforms, in which the d.c. current

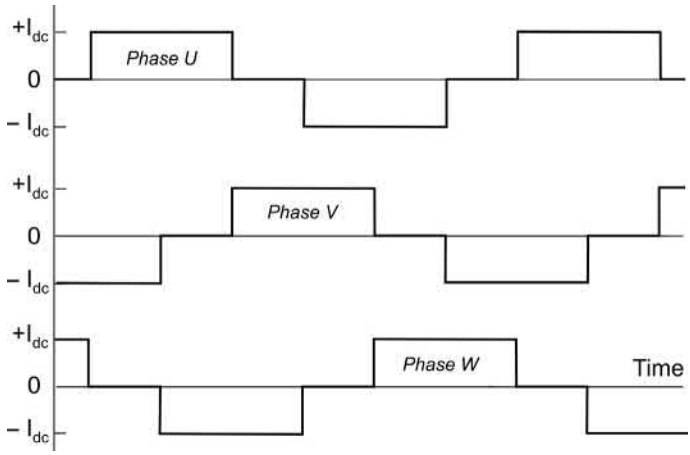

FIG. 9.25 Idealised steady state motor phase current waveforms.

is assumed to be constant, and the current transitions from one value to another instantly, in the interest of simplicity: in practice there will of course be some current ripple and finite rise and fall times, the latter determined by the so-called 'transient reactance' of each phase, which fortunately is relatively low.)

If as is usual the windings are sinusoidally distributed, the resultant stator m.m.f. will always be sinusoidally distributed around the machine despite the non-sinusoidal current waveforms. But instead of progressing smoothly around the stator (as it would if the phase currents were sinusoidal in time), the m.m.f. jumps forward in space by 60° (elec) (i.e. one third of a pole-pitch) every time the motor converter commutates the link current to the incoming phase. Hence in the steady state there will inevitably be a torque ripple, which is the price to be paid for having a relatively unsophisticated inverter. Happily, resonances excited by the torque ripple are rare and can usually be overcome by preventing continuous operation at particular speeds associated with any mechanical resonance.

We have illustrated the basic operation by focusing on the steady-state, but it is important to remind ourselves that this is very different from the utility-fed case we looked at previously. In the utility-fed situation, the motor will lose synchronism and stall if the pull-out torque is exceeded. Here, the motor is self-synchronising, and if the load on the shaft is greater than the torque that can be developed when the link current is at its maximum, the motor will simply slow down, and pick up speed again when the shaft load is reduced, just as in a d.c. motor.

Current source inverter (CSI)

The term 'Current Source Inverter' has already been used to describe the power circuit shown in Fig. 9.24, so it is now time to explain what the term means.

It may be unnecessary, but we will start by making the point that the term current source inverter does not mean that the link current never changes, which is what a reader who is familiar with current sources in other contexts, especially in low power electronics, might think. In the present context, it means that under normal operating conditions, the link current cannot be changed rapidly, i.e. not significantly during one complete period of the motor current waveform, even at the lowest operating speed. The reader will not be surprised to learn that the link inductor is central to achieving this state of affairs.

We have said many times before in this book that inductance in a circuit results in the current waveform being much smoother than the voltage waveform (see for example Fig. 8.10), and that the bigger the inductance, the smoother the current. We need to recall that the voltage across an inductor is related to the current through it by the equation

$$v = L\frac{di}{dt} \text{ or } \frac{di}{dt} = \frac{v}{L}$$

i.e. the rate of change of current is proportional to the voltage difference and inversely proportional to the inductance.

The rectified output voltage waveform of the supply-side rectifier will typically be as shown (left) in Fig. 9.26, which shows the potential of the top of the converter (i.e. at left hand end of the inductance) with respect to the bottom of the d.c. link. It a has a substantial ripple at six times the utility frequency, and the average (d.c.) voltage is (V_s).

At the same time the motor side converter (which is connected upside down) is inverting and the potential of the right hand end of the inductor will be as shown on the right of Fig. 9.26; the average (d.c.) voltage is (V_m). Note that whenever we want the link current to be constant, the first requirement is that the average voltage across the inductor is zero, which means that V_s must be equal to V_m, i.e. the d.c. voltage is the same for both converters. The current controller will adjust the firing angle of the supply-side rectifier in order to achieve this. (In practice there will be a small voltage difference because of the resistance of the inductor.)

The instantaneous voltage across the inductance is the difference between the two waveforms in Fig. 9.26. Finding the difference would be difficult because the two waveforms are not synchronised in time, but we can see that

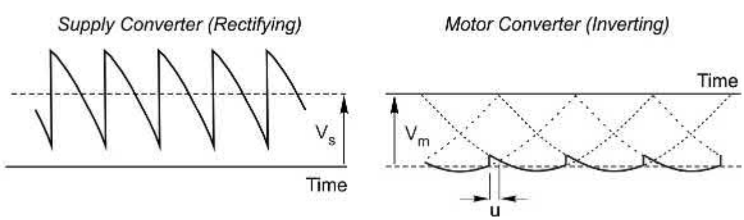

FIG. 9.26 Typical DC link voltages of the current source inverter.

there will be substantial voltages across the inductor, not least the sudden step changes as a result of each supply converter commutation. If no inductance was present, there would consequently be huge step changes in the link current, and wild fluctuations in the motor torque. Hence we need to decide how much 'ripple' current we can tolerate, and choose the inductor accordingly. In practice a peak ripple of say 5% of rated current is typical for most applications.

Having chosen the inductor to suppress the current ripples, it is inevitable that when we want to raise or lower the mean current in order to vary the torque, the inductor will impede our efforts, and the response of the current control loop will be more sluggish. Fortunately, in large motors, we are not usually seeking high bandwidth torque control, so the compromise is acceptable.

The reason for the description 'current source' should now be clearer. Despite the switching of the link current from one phase to another in which the instantaneous induced e.m.f. is very different, the link current remains more or less unchanged, so that the impression we get is that the link current is independent of the load that we present to it.

Starting

So far we have assumed that the motor is running in the steady state, so that the motor terminal voltage (or better the back e.m.f.) can be used to determine the rotor position, but at start up and at very low speeds the magnitude of the back e.m.f. is too small to be used for control purposes or to commutate the current in the motor converter. The commutation is usually achieved by momentarily switching off the d.c. link current (by phase control of the utility side converter). When the current has reduced to zero the next pair of thyristors can be fired and the current built up again. A shaft mounted absolute position sensor provides position information to determine when the current should be commutated from one switch to the next. Whilst this sounds a slow and laborious process, it is not, and large machines can be accelerated in a few seconds. Above approximately 5% of rated speed the machine generates sufficient voltage for natural commutation, and subsequent control is undertaken in a similar manner to that of a d.c. drive.

Where the load is highly predictable some systems actually impose a predetermined sequence of current pulses (applied sequentially to the motor phases) to "crank" the motor up to a speed at which the back e.m.f. becomes of sufficient magnitude to be used for position sensing and commutation. Such systems can avoid the need for a shaft mounted position sensor.

As in the d.c. drive, the a.c. supply power factor is poor at low speeds, but on the plus side, full four-quadrant operation is possible without any additional equipment.

Control

The fact that the synchronous motor/self-commutating thyristor converter combination behaves in much the same way as a conventional d.c. motor means that the control philosophy that we discussed in Chapter 5 can be employed, as shown in Fig. 9.27.

An inner current control loop provides torque control, and the torque reference is obtained from the error signal from the outer speed control loop. The maximum current is limited by clamping the current reference signal, I_{ref}.

The items inside the chain-dotted line in Fig. 9.27 form the essential elements of a self-synchronous motor, and all that is necessary for it to function is a d.c. supply, shown as V_m in Fig. 9.27. The control scheme is the same as we saw for the conventional d.c. motor (Fig. 4.12), and it operates in much the same way as previously described, including field weakening, the rotor field current being reduced progressively as the link voltage gets close to its maximum value.

Looking back to Fig. 9.26, the frequency and amplitude of the motor-side voltage waveforms are of course proportional to the speed, and in motoring mode the converter firing angle delay is kept constant at slightly less than 180°, so that the torque angle is maintained at +90°, thereby maximising the torque for the given current. The safety angle ‘u’ (see Fig. 9.26) is necessary to ensure successful commutation.

When regeneration is required, the firing angle delay of the motor converter is reduced to zero, the torque angle then becomes −90° to give maximum braking torque and the motor converter then rectifies with maximum voltage (and

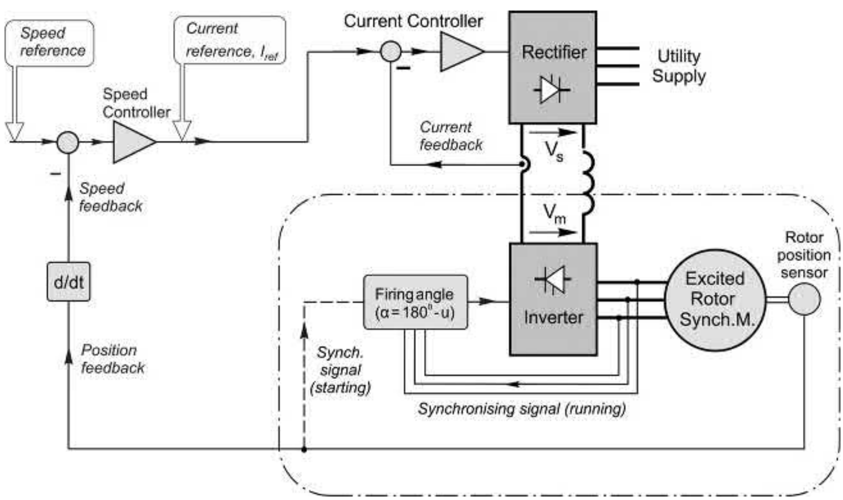

FIG. 9.27 Control scheme for converter fed excited rotor synchronous motor.

hence maximum power and torque). Energy then flows back to the supply converter, which will have moved into the inverting mode in order to control the current to its demanded value.

Finally, we should note that current source inverters are now being challenged at powers up to 5 MW by voltage source inverters, the basic circuit and control of which is the same as for the PM motor to which we now turn.

9.6.3 Permanent magnet motor

Permanent magnet motor drives are a very important and rapidly growing sector of the drives market, and we will be looking in Section 9.7 at the excellent performance characteristics of these drives in more detail to explain why. In terms of the practical implementation, we have little new to introduce because the converter circuits used in commercial drives are exactly the same as we have already discussed in Chapters 7 and 8 for the induction motor, and the control is also very similar.

The power circuit

The general arrangement used in the control of PM motors (and all synchronous motors apart from the excited-rotor type) is as shown in Fig. 7.2.

The inverter bridge is shown in Fig. 9.28 (it is the same Fig. 2.21, but repeated here to avoid turning back).

In Chapter 2 we discovered that during motoring operation (power flowing from the d.c. link to the motor) the power flows through the main power switch, usually an IGBT, and the reactive power flows through the anti-parallel diode. When the bridge feeds an induction motor, the magnetising (reactive) component of current is high, but in a PM motor there is no need for magnetising current when the motor is running below base speed, so there is the possibility of savings in terms of the diode rating, and also in improved inverter efficiency. We will consider the impact of operation above base speed (in the field weakening region) later.

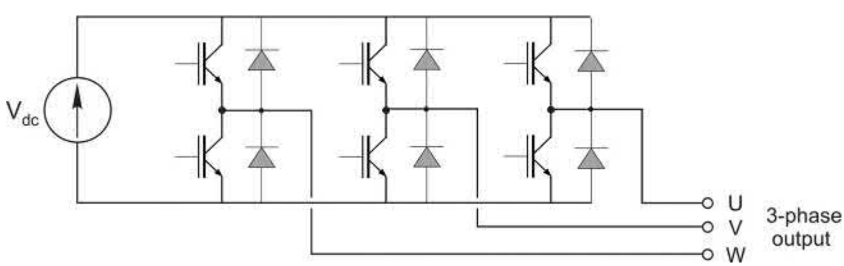

FIG. 9.28 Three-phase inverter power circuit.

Control

PM motors are almost always fed from a voltage source inverter, and so the control strategy is different from that for the current source inverter shown in Fig. 9.27. The arrangement of a typical field oriented control system for a PM motor is shown in Fig 9.29, in which the asterisk (*) is used to denote demanded quantities. As mentioned previously, the scheme is very similar to that for the induction motor (see Fig. 8.16).

The flux demand has been set to zero because the flux is provided by the magnets. However, if the field current demand was not set to zero, the control system would provide a flux component of stator current that could either increase or decrease the effect of the magnet's flux, depending on the polarity of the reference signal. Clearly there will be an upper limit on the flux because of saturation of the magnetic circuit, and in practice, reducing the flux is a more attractive proposition because it allows us to operate in a field weakening mode and so extend the speed range into a constant power region, as discussed at the end of Section 9.5.

It would be possible to apply field weakening control by applying a speed-dependent term ($-i_d$), but this would require a good understanding of the machine characteristics, which is far from easy. A simpler approach can be applied if we think back to how we control a d.c. motor in field weakening (see Chapter 2). Steady-state operation up to base speed requires the stator voltage to increase with speed, but once base speed has been reached, the voltage, by definition, cannot increase any further. For the PM motor drive the situation is no different and so negative i_d must be imposed via the stator current, and a simple way to do this is shown by the additional "Voltage Control (Field Weakening)" loop at the top of Fig. 9.29.

Field weakening operation of PM motors presents an interesting practical problem that may occur if the control system should fail when running at high speed. For example, suppose the motor is running at three times base speed and for some reason the control system loses control. The component of stator current that is opposing the magnet flux disappears and so the rotor is spinning at

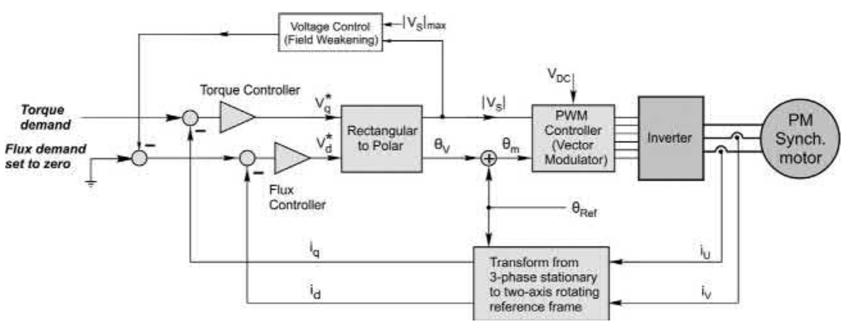

FIG. 9.29 Field oriented control scheme for permanent magnet synchronous motor.

three times base speed with full airgap flux, and the terminal voltage will rise to three times its rated value. The motor insulation systems and the power converter components are not normally rated to withstand such voltages and so catastrophic failure will result. Rating all components for such a situation would usually be cost prohibitive, and so other means of protection need to be sought. A common solution is to put a simple crowbar circuit near to the motor terminals, such that if such a fault does occur, and the terminal voltage rises, it is limited to an acceptable level.

9.6.4 Reluctance motor

The reluctance motor was once the motor of choice in the early days of power semiconductor based a.c. drives, but until recently had become all but forgotten, in a market dominated by the induction and PM motors. In recent years however, it has re-emerged as a commercial competitor of the induction motor. Whilst the motor is, in general, a little larger than the equivalent induction motor, some suppliers have put together credible motor/drive packages such as that shown in Fig 9.30.

The power circuit

The voltage source power converter (Fig 9.28) predominates for synchronous reluctance drives, and because it is common to induction motor and PM synchronous motor drives, a standard hardware platform can be utilised for the control of many different forms of a.c. motor. It turns out that the control strategy for the reluctance motor also has much in common with the induction and PM drives, as we will now see.

FIG. 9.30 Reluctance motor and drive package. *(Courtesy Siemens.)*

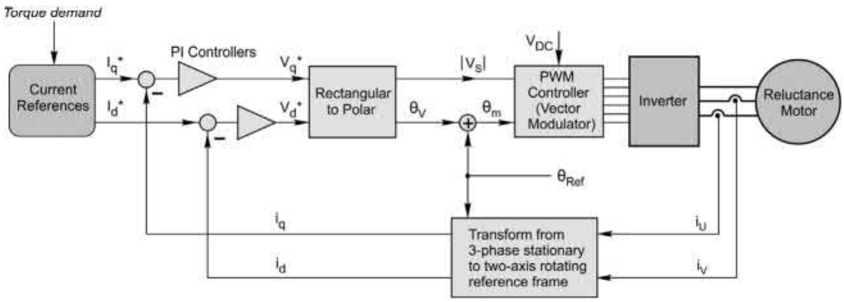

FIG. 9.31 Field oriented control scheme for reluctance motor.

Control

Whilst reluctance motors can be operated with very simple control strategies, field orientation control is used in most commercial drives. We have seen in section 9.3.3 that as long as we are able to control the angle of the stator current phasor in relation to the rotor saliency we are able to directly control the motor torque. We have also seen that the torque depends on the square of the current, and that it is a double frequency function of the rotor angle.

The rotor design and modelling of reluctance motors has received a great deal of academic and industrial attention in recent years and the optimisation of the stator current phasor angle, taking into account the effects of saturation, is a subject area which is well beyond the scope of this book. To simplify matters, we will introduce an undefined 'Current References' function block into the control diagram, Fig. 9.31. In practice, the reality in many cases is that the function simply creates a fixed ratio of i_q/i_d to give an angle close to the 45° associated with the peak torque as shown in Fig 9.8.

9.6.5 Salient permanent magnet motor

Control of the Salient PM ("PM/Rel") motor is undertaken in much the same way as described above for the reluctance motor, but with more complex algorithms in the "Current References" box. Again relatively simple implementations based on a fixed ratio of I_q/I_d are possible, but for optimised control, the impact of saturation, common in such machines, needs to be taken into account.

9.7 Performance of permanent magnet motors

Throughout this section we will use the terms 'brushless' and 'permanent magnet' to be synonymous when they are applied to an electric motor. We will also avoid the confusing term 'brushless d.c. motor' because, as already mentioned in the introduction, such motors are always supplied with alternating currents.

However we should mention that the inherent electromagnetic properties of a PM motor can be quantified by its 'motor constant' in much the same way as a d.c. motor (Chapter 3). If we spin the rotor of a PM machine at angular velocity ω, the r.m.s. value of the sinusoidal e.m.f. induced in each phase is given by $E = k\omega$, and if we supply balanced currents of r.m.s value I_a to the three phase windings in quadrature to the field, the torque is given by $T = kI_a$, where k is the machine constant, expressed in the SI units of Volts per radian per second or the equivalent Newton-metres per Ampere. (In practice, manufacturers usually quote k in terms of Volts per thousand revs per min.) These relationships are identical to those we discussed earlier for the DC machine, and once again they underline the unity of machines that operate on the '*BIl*' principle.

We have hinted previously that PM motors offer outstanding performance in terms of power density and performance in comparison with induction and d.c. machines, and in this section we look briefly at the underlying reasons. We then discuss the limitations that govern performance and finally give an example that illustrates the impressive results that can be obtained.

9.7.1 Advantages of PM motors

The stator windings of PM motors do not have to carry the excitation or magnetising current required by the induction motor, so a given winding can carry a higher work current without generating more heat, thereby increasing the electric loading and the specific power output (as discussed in Chapter 1).

Cooling the rotor is difficult in any enclosed machine because ultimately the heat has to get to the stator, so the absence of current on the rotor not only improves efficiency by reducing the total copper loss, but also eases the cooling problem.

Historically, brushless PM motors only became practicable with the advent of power electronics, so it became normal for them to be supplied via power electronics with the associated expectation that they would operate in a speed-controlled drive. The majority were therefore not expected to operate directly off the utility supply, and as a result their designers had much greater freedom to produce bespoke designs, tailored to a specific purpose. (The designers were also less hamstrung by standards developed over many years dictating shaft size and height, mounting arrangements and cooling arrangements.)

For example suppose we require a motor that can accelerate very quickly, which implies that the ratio of torque to inertia should be maximised. We saw in Chapter 1 that with given values of the specific magnetic and electric loadings, the torque is broadly dependent on the volume of the rotor, so we are free to choose long and thin or short and fat. The inertia of a homogeneous rotor is proportional to the fourth power of its radius, so clearly for this application we want to minimise the rotor radius, so a long thin design is required. Fortunately, there is considerable flexibility in regard to the shape and size of the rotor magnets, so no serious constraint applies in relation to rotor diameter. Many so-called servo motors have this profile, as illustrated in Fig. 9.32.

A different application with the same continuous torque and power requirements would require the same rotor volume, but if, for example, the application requires that unwanted changes in speed caused by step changes in load torque must be minimised, the inertia should clearly be maximised, with a short rotor of larger diameter.

A section through a high inertia PM motor typically rated up to 100 Nm and 3000 rev/min is shown in Fig. 9.33. It is interesting to note how little of the volume of the motor is actually taken up with active material, namely the stator and rotor laminations and the surface-mounted magnets on the rotor. The end

FIG. 9.32 Permanent magnet servo motors. *(Courtesy Nidec—Control Techniques Dynamics Ltd.)*

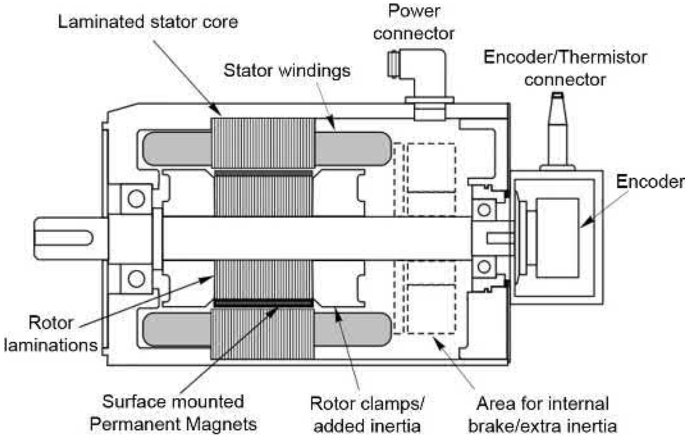

FIG. 9.33 Typical high inertia permanent magnet synchronous motor. *(Courtesy Nidec—Control Techniques Dynamics Ltd.)*

windings of the stator windings are seen to contribute significantly to the overall volume of the motor illustrated. (The move to using segmented stator windings reduces the impact of the end winding and can lead to significant reductions in total motor volume.) Most PM motors employ rare-earth magnets which have much higher energy product (in effect a measure of their magnetising ‘power’) than traditional materials such as Alnico, so they are very small, as shown in Fig. 9.33.

As in other motors the heat dissipated in the stator diffuses into the air through the frame to the finned case and hence to the surrounding air. However, the design of some motors of this type is based on the requirement that a substantial proportion (perhaps 40%) of the loss is conducted through the mounting flange to a suitable heatsink, so this is an area where great care needs to be taken with the thermal properties of the mounting.

9.7.2 Industrial PM motors

In the preceding section we mentioned that bespoke design of PM motors has long been considered unexceptional, but recent years have seen the emergence of PM motors packaged in the same industrial motor (IEC or NEMA) housings that are been used for induction motors, as shown in Fig. 9.34.

This type of motor is now marketed as a direct competitor of the induction motor in variable speed applications. They are targeted at general applications where the higher initial cost is offset by their higher efficiency and high power

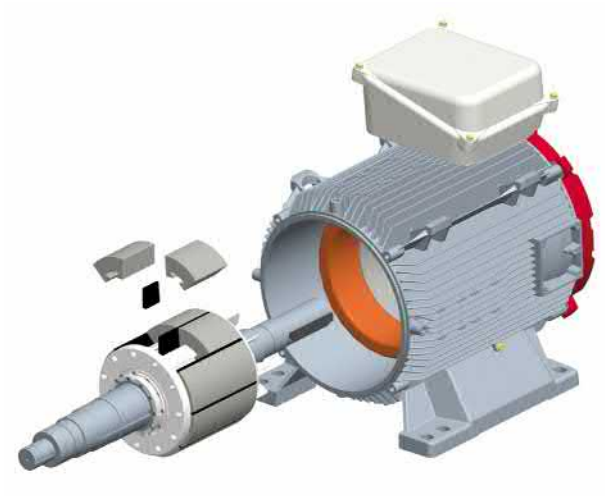

FIG. 9.34 Permanent magnet synchronous machine in a standard IEC frame. *(Courtesy Nidec—Leroy Somer.)*

density. The heat loss in the permanent magnet rotor is much less than in the corresponding induction motor, so the rotor runs cooler which may also be an advantage in aspects such as bearing life.

9.7.3 Summary of performance characteristics

Permanent magnet motors with low inertia rotors are used in high performance servo applications such as machine tools or pick and place applications where fast precise movements are required, and motors with high inertia rotors (and high pole numbers) suit low speed applications such as gearless lift systems.

The performance characteristics of these drives are summarised below:

- Excellent dynamic performance at speeds down to standstill when position feedback is used.
- For precision positioning the position feedback must define the absolute position uniquely within an electrical revolution of the motor. This can be provided with a position sensor or alternatively a sensorless scheme can be used. The performance of a sensorless scheme will be lower than when a position sensor is used.
- Field weakening of permanent magnet motors is possible to extend their speed range, but (as shown in Section 9.5) this requires additional motor current, and so the motor becomes less efficient in the field weakening range. This form of control also increases the rotor losses and raises the temperature of the magnet material, thereby increasing the risk of demagnetisation. Care is also needed in such applications to avoid overvoltage at the motor and drive terminals in the event of a loss of control: at high speeds, the open-circuit voltage will exceed the rated value.
- Permanent magnet motors exhibit an effect called cogging that results in torque ripple. It is caused by magnetic reluctance forces acting mainly in the teeth of the stator, and can be minimised by good motor design, but can still be a problem in sensitive applications.
- Permanent magnet motors can be very efficient as the rotor losses are very small.

To give an impression of the outstanding performance that can be achieved by a brushless PM motor, Fig. 9.35 shows the results from a bench test in which the speed reference begins with a linear ramp from zero to 6000 rev/min in 0.06 s, followed shortly by a demand for the speed to reverse to 6000 rev/min, then back to full forward speed and finally to rest, the whole process lasting less than 1 s.

The motor was coupled to a high-inertia load of 78 times the rotor inertia, which makes the fact that the speed reversal is accomplished in only 120 ms even more remarkable. It takes only three revolutions to come to rest and a further three to accelerate in reverse. In common with many high performance applications, the drive control is actually implemented in the form of position control, with the shaft angle being incremented at a rate equivalent to the

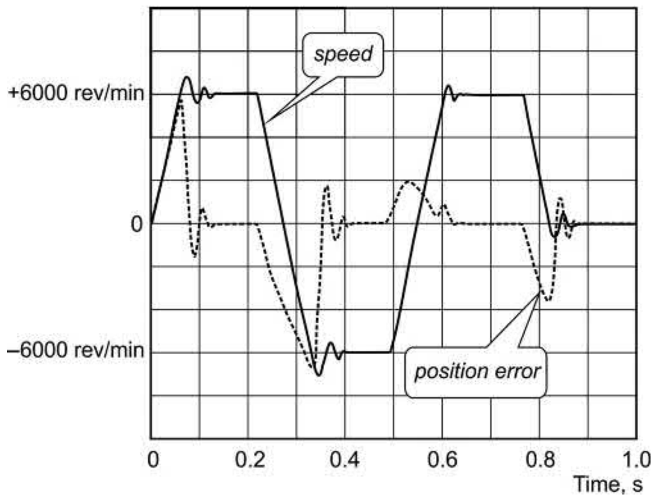

FIG. 9.35 Permanent magnet motor drive performance under rapid reversal test.

required speed. The dotted trace in Fig. 9.35 shows the position error of the motor shaft throughout the speed reversal. The maximum error is less than 0.05°, so by any standard this is truly impressive. It is no wonder that the brushless PM motor is frequently chosen for applications where closely co-ordinated motion control is called for.

Finally, on a matter of terminology, it is worth pointing out that brushless PM motors are sometimes referred to as 'servo' motors. The name "servo" originates from 'servomechanism', defined as a mechanical or electrical system for control of speed or position. The term tends to be used loosely, but broadly speaking when it is applied to a motor it implies superior levels of performance.

9.7.4 Limits of operation of a brushless PM motor

We have previously talked about the limits of operation that determine the rating and operating envelope of other types of electrical machine, so we will conclude this section by taking a closer look at the limits of operation for a Brushless PM servo motor (which usually has no external cooling fins or fan). A typical torque-speed characteristic is shown in Fig. 9.36.

The individual limits shown in Fig. 9.36 are discussed below, but the most striking feature is clearly the very large area where operation above rated torque is possible (albeit on an intermittent basis). This provision clearly reflects the potential application areas, such as rapid positioning systems, where high acceleration is needed for relatively short periods.

The continuously-rated region is, as usual, limited by the allowable temperature rise of the motor. At standstill the predominant source of loss is the stator

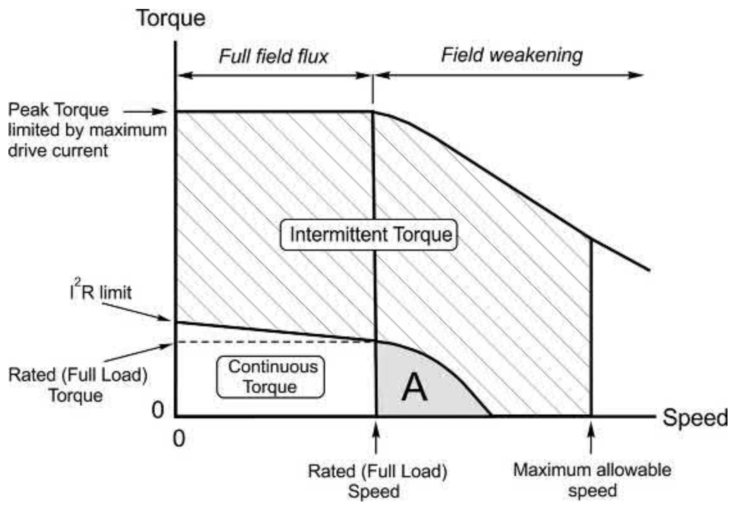

FIG. 9.36 Limits of operation of a typical brushless PM motor.

copper loss (shown as I^2R limit in Fig. 9.36), but at higher speeds the iron loss becomes significant and the full load torque at rated speed is therefore less than the standstill torque.

The upper boundary of the intermittent torque region is usually determined by the maximum current that the drive converter can supply: only intermittent operation is possible otherwise the motor overheats.

Continuous operation is not possible when field weakening either, because, as explained in Section 9.5, the stator current is large in order to oppose the magnet flux. For this thermal reason Region A shown in Fig. 9.36 can be somewhat truncated compared with the characteristic field weakening region of the induction motor.

Operation of PM motors at excessive temperatures can lead to demagnetisation of the rare earth magnets in the motor, as well as the hazards common to all other electrical machine such as degradation of insulation. Good thermal protection is therefore necessary. Below base speed, for applications involving few excursions outside the continuous operating region, a relatively simple motor thermal model in the drive control scheme, may be adequate. For applications involving significant operation in the intermittent torque region, and certainly where field weakening operation is used, more complex thermal models would be needed and these are usually supplemented by thermistors embedded in the stator windings.

9.7.5 Brushless PM generators

As with almost all forms of electrical machines the PM motor can be operated as a generator, with mechanical energy supplied to the shaft being converted to electrical energy. The advantages of high power density and efficiency offered by the PM motor, of course apply equally when the machine is being used as a generator.

Many commercial wind generators up to 75 kW use PM synchronous machines. Much larger wind generators are also in service with some multi-pole multi-MW motors being applied in utility scale turbines with direct drive systems, i.e. systems which do not employ a gearbox between the wind turbine and the generator.

9.8 Emerging developments in permanent magnet motors

So far in the book, our discussions of a.c. machines (both induction and permanent magnet) have concentrated on machines with low pole-numbers (e.g. 2, 4, 6….), because they are by far the most numerous. The corresponding speeds when fed from the 50 or 60 Hz utility supplies cover the range from 1500 rev/min at 50 Hz to 3600 rev/min at 60 Hz, and are therefore well-suited to the majority of industrial applications.

We have also seen that for motors with similar magnetic and electric loadings, and similar cooling arrangements, the specific power output is proportional to the speed. A 4-pole, 50 Hz, 1500 rev/min, totally enclosed fan ventilated induction motor is therefore smaller than a 12-pole, 50 Hz, 500 rev/min version producing the same power output.

The availability of inverters that are highly efficient at frequencies in the kHz range has largely removed the frequency constraint, so that, regardless of pole number, the synchronous speed can now be varied smoothly up to much higher frequencies than that of the utility supply, so it is no longer necessary to have a low pole number to achieve the high speed required for a high specific output.

Superficially, the high pole-number stator windings of recently developed PM motors (see Fig. 9.39) often appear radically different from their predecessors, so we might expect to need a new approach in order to explain how they work. In fact, we will see that their novel stator windings represent limiting versions of the conventional multi-slotted, 2-layer stator windings, with coils spanning several slots (see Chapter 5), which means that their behaviour can be explained using the same ideas as we have used previously.

9.8.1 Advantages of high pole number

For a PM motor, where the main excitation flux is produced by magnets on the rotor, a higher pole-number brings advantages in terms of economy of copper and iron, and often simplicity of manufacture, and comparatively few disadvantages, as we will see shortly. This explains why higher pole numbers now predominate in PM motors, particularly those aimed at high-volume emerging markets such as electric vehicle drives, where investment in mass production is justified. (By contrast, it is worth pointing out that when the excitation is provided by the stator winding, as in an induction motor, a higher pole numbers is not inherently desirable (apart from yielding a low speed at 50 or 60 Hz). This is because to achieve a given magnetic loading the magnetising m.m.f. per pole must remain the same, but the available slot space per pole reduces as the pole number increases. Hence the proportion of the slot space to be devoted to the magnetising function increases, leaving less for the useful or work component of current. Induction motors with high pole numbers therefore have an unattractive low power factor.)

We are discussing permanent magnet (PM) motors, where the main flux is provided by the rotor magnets, so (as explained in Section 9.3) we now choose to picture torque production as the result of the magnet flux interacting with the currents on the stator, rather than stator flux interacting with rotor current. Thus when we refer to electric loading (see Section 1.5.1), we are referring to the stator, rather than the rotor.

The first and most obvious advantage of a high pole number (as compared with say a 2-pole) is that, for a given magnetic loading, the flux per pole is less and consequently the depth of iron core behind the slots (necessary to carry the circumferential flux from north pole back to south pole without saturating) is reduced. This results in a cost saving in expensive core steel on both stator and rotor, the latter being annular: a sketch showing 2-pole and 8-pole stators for the same rotor diameter, and similar magnetic loadings, is shown diagrammatically in Fig. 9.37.

The second advantage relates to the progressive reduction in the length of the end windings as the pole number is increased. The function of the end windings is simply to connect the 'go' and 'return' sides of the coils: they contribute nothing to the output power, but their resistance represents a source of unwanted 'copper loss', so there are savings in both materials and losses when the pole number is increased.

9.8.2 Segmented core and concentrated windings

For high pole-number stators, splitting the stator core into segments has obvious attractions both in terms of automated coil assembly and reduction of waste in comparison with conventional (full circle) laminations: a typical set of interlocking segments is shown in Fig. 9.38.

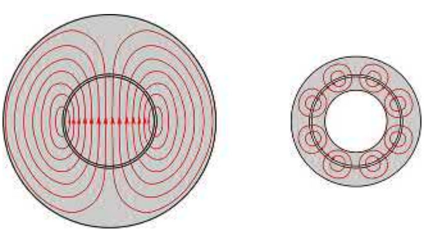

FIG. 9.37 Comparison of 2-pole and 8-pole stator cores.

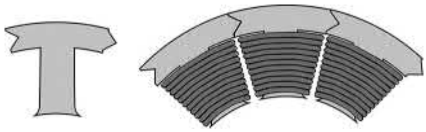

FIG. 9.38 Segmented stator.

Single-tooth versions that are pre-wound before assembly are now in widespread use for PM motors. They can have a relatively high fill factor, but the disadvantage of this 'concentrated coil' winding is that there is no longer any freedom (as there is with a 2-layer winding) to select the coil-pitch: in a concentrated winding spanning one tooth only, the coil pitch is of necessity equal to the slot pitch. We will see in the example in the next section that this leads to a relatively poor m.m.f. waveform.

9.8.3 Fractional slot windings

In order to understand how multi-pole stator windings with few slots derive from their conventional predecessors, we must briefly recap and then extend the discussion of fractional slot windings in Section 5.2.3. Readers who want to skip this detail will doubtless be happy to accept that it is not essential to understanding the remainder of the chapter.

Historically, the majority of a.c. machines (synchronous and induction) had stators with a large number of slots with 3-phase, 2-layer windings of the type discussed in Section 5.2 of Chapter 5. (By 'a large number of slots' we mean a range from perhaps 24 to around 100). This type of winding remains dominant, especially in the larger sizes.

In Section 5.2, we emphasised that conventional stator windings were laid out in order to produce a sinusoidal rotating magnetic field of a particular pole number. It became clear that by using a two-layer winding with coil pitch less than one pole-pitch, and with several slots per pole per phase (s/p/p), a set of identical coils could be distributed so that the resultant space wave of m.m.f. was a very close approximation to a sinewave, thereby minimising the undesirable effects of harmonic fields.

However, even with a large number of slots, there are occasions where the s/p/p is not an integer, for example a 4-pole, 3-phase, two-layer winding in a stator with 54 slots, for which the s/p/p = 4.5. Where the s/p/p is not an integer, the winding is described as 'fractional slot', and in such windings, some phase bands occupy more slots than others: in the example above, the first and third phase bands will occupy four adjacent slots, while the second and fourth will occupy 5 adjacent slots. As usual, all three phases will be identical but displaced by one third of a pole-pitch. Such windings have long proved to be entirely satisfactory.

Example: 10-pole, 3-phase winding in 12 slots

Returning for a moment to the 2-layer winding, we can see by comparing Figs. 5.4 and 5.3 that when the s/p/p reduces from 3 to 1, the m.m.f. waveform degenerates from a reasonable approximation to sinewave to a rectangular waveform, which is of course heavily polluted with harmonics. Many contemporary permanent magnet motor stators with concentrated windings take the move towards fewer slots per pole well beyond that shown in Fig. 5.3, and typically have values of s/p/p less than one. For example, for the 3-phase, 10-pole winding in 12 slots shown in Fig. 9.39, the s/p/$p = 0.4$.

The rotor in Fig. 9.39 clearly has 10 permanent magnets, which immediately indicates that it is a 10-pole machine. However, it is by no means self-evident (even to an experienced observer) that the 12-slot stator winding, fed with balanced 3-phase currents, will produce a 10-pole rotating field. (Indeed, given only the stator to examine, it would not be at all surprising if a non-expert concluded that it was for a 12-pole machine, or a switched reluctance motor, or a stepping motor (see Chapter 10).

To see that the winding will indeed produce a 10-pole field to cooperate with the rotor, we must examine the stator m.m.f., using the approach followed in Chapter 5.

Given that the aim of each phase winding is to produce a 10-pole m.m.f., the fact that there are only 4 slots for each phase clearly indicates that the outcome will inevitably be less than ideal. In fact, by sharing some of the slots with a coil from another phase (see Fig. 9.39), we manage to obtain 8 'half height steps' in the phase m.m.f. waveform, as shown in Fig. 9.40: in this figure, the current in phase U is at a positive maximum, while each of the other two phases carry a negative current of half peak amplitude. The phase m.m.f. is clearly a very poor approximation to a 10-pole sinusoid, it probably being best described as a 10-pole square wave with six of the ten poles missing.

The other two phase-windings are identical to phase U, but displaced in space by 120° and 240° electrical (i.e. with respect to the 10-pole wave), and when their m.m.f.'s are added the resulting m.m.f. wave is much better, as shown by the resultant waveform in Fig. 9.40.

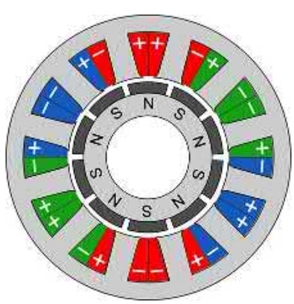

FIG. 9.39 10-pole, 3-phase winding in 12 slots.

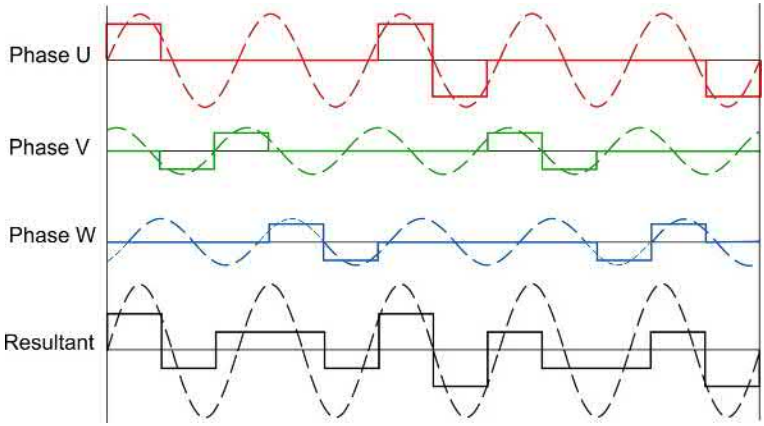

FIG. 9.40 m.m.f. waveform of 10-pole, 3-phase winding in 12 slots.

This example illustrates the extent to which expediency and cost of manufacturing outweigh considerations of purity of m.m.f. waveform for applications such as hybrid vehicles, where the motor will probably be integrated within the power train and cooled by engine oil. 'Fault tolerance' is at a premium in such duties, meaning that the motor is expected to continue to produce torque despite failure of one or more of its parts. In the case of the motor in Fig. 9.39, the four coils in each phase may be supplied from separate inverters, so that healthy coils can continue in the event of a failure elsewhere.

9.9 Review questions

(1) What voltage should be used to allow a 420 V, 60 Hz, 4-pole synchronous motor to be used on a 50 Hz supply?

(2) What purpose, might be served by a pair of 3-phase synchronous machines (one of which has ten polar projections on its rotor and the other twelve) mounted on a bedplate with their shafts coupled together, but with no shaft projections at their outer ends?

(3) The book explains that in excited-rotor synchronous machines the field winding can be supplied with d.c. current via sliprings. Given that the field winding rotates, why is there no mention of any motional e.m.f. in the rotor circuit?

(4) A large synchronous motor is running without any load on its shaft, and it is found that when the d.c. excitation on the rotor is set to either maximum or minimum, the a.c. current in the stator is large, but that at an intermediate level the stator current becomes almost zero. The stator power seems to remain low regardless of the rotor current.

Explain these observations by reference to the equivalent circuit and phasor diagram. Under what conditions does the motor look like a capacitor when viewed from the supply side?

(5) This question refers to Fig. 9.5. In the motor shown on the right of Fig. 9.5, what would happen if the polarity of the stator current was suddenly reversed?

(6) How many discrete stable equilibrium positions would you expect to find in a 4-pole reluctance motor? (Hint: imagine the stator current to be d.c. as in the 2-pole example in Fig. 9.8.)

(7) By clever re-design of the rotor of a current-fed reluctance motor, the quadrature axis reactance is reduced while the direct axis reactance is unchanged, and as a result the ratio of direct to quadrature reactance is increased from 4 to 5. If the motor is current-fed, by how much will the peak torque increase?

(8) The stator of a 10-pole PM motor is supplied with a steady (d.c.) current from a battery, and the unloaded rotor is at rest at an angle of 0°.
   (i) The rotor is then turned by hand before being released. At what angle will the rotor eventually come to rest if the angle it is initially turned through is (a) 15°, (b) 45°.
   (ii) At what angle(s) will the maximum torque be detected?
   (iii) If the polarity of the current is reversed, at what angle will the rotor settle?

(9) Explain briefly why, for a PM motor, it is possible for a given torque to be produced for a range of values of stator current.

(10) A PM motor driven from a voltage source inverter is running light at its base speed and the current is negligible. The speed is increased to twice base speed, and despite having no load on the shaft, the current is now much larger. Explain, preferably with the aid of phasor diagrams.

Answers to the review questions are given in Appendix.

Chapter 10

Stepping and switched reluctance motors

10.1 Introduction

We have grouped stepping and switched reluctance (SR) motors because despite substantial differences between the modes of operation (and the power ratings), the fundamental mechanism, i.e. reluctance torque, is the same in both. However, unlike the (**s**ynchronous) reluctance motors considered in Chapter 9, stepping and switched reluctance motors are 'doubly-salient' with projecting poles or saliencies on both the stator and the rotor (see Fig. 10.5).

In Chapter 9 we explained that reluctance torque referred to the natural tendency of the projecting poles on a rotor (without any windings or magnets) to align themselves with the axes of the multi-polar field produced by the windings on the smooth stator. Intuitively, therefore, it should not be difficult to see that the reluctance torque is further enhanced when we replace the smooth bore stator by one with projecting poles carrying discrete exciting windings. In effect, we are 'sharpening up the focus' of the stator field, and thus increasing the stiffness of the rotor about its equilibrium alignment position.

We will see that doubly salient machines require their discrete stator windings to be switched sequentially across a d.c. supply, and it was the difficulty of doing this that proved the stumbling block for the early nineteenth century pioneers, who saw these inherently simple machines as being suitable for electric traction. The mechanical switching devices at the time were inadequate, and it was not until power semiconductor devices arrived in the 1960s that interest in the technology was renewed.

The theoretical treatment of doubly-salient machines is not as straightforward as that for motors with smooth stators and rotors such as the induction motor or the non-salient synchronous motor, which we have studied in previous chapters. For those motors it was relatively easy to explain the mechanism of torque production using the 'force on a conductor' formula ('BIl'), and because the windings were sinusoidally distributed in space, and the currents were sinusoidal in time, we were able to make extensive use of space and time phasors to

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00010-6

illuminate performance. We even managed to explain and quantify reluctance torque in a singly-salient motor by the same approach.

Unfortunately the 'BIl' approach does not lend itself easily to the analysis of doubly-salient machines. Instead, most analysis and design is based on computer-aided modelling of the complex flux patterns (often involving high levels of saturation), in order to establish how the flux linking each winding varies with the position of the rotor and the currents in the windings. The torque is then predicted using the circuit-based approach outlined in Section 3.4 of Chapter 8. As a result there are few if any analytic results that are helpful in terms of understanding, and we will therefore follow a largely descriptive approach in this chapter, beginning with stepping motors because they provide useful groundwork for the SR material that follows.

10.2 Stepping motors

Stepping (or stepper) motors historically became attractive because they can be controlled directly by computers, microcontrollers and Programmable Logic Controllers (PLC's).[1] Their unique feature is that the output shaft rotates in a series of discrete angular intervals, or steps, one step being taken each time a command pulse is received. When a definite number of pulses has been supplied, the shaft will have turned through a known angle, and this makes the motor well suited for open-loop position control.

The idea of a shaft progressing in a series of steps might conjure up visions of a ponderous device laboriously indexing until the target number of steps has been reached, but this would be quite wrong. Each step is completed very quickly, often in less than a millisecond; and when a large number of steps is called for the step command pulses can be delivered rapidly, sometimes as fast as several thousand steps per second. At these high stepping rates the shaft rotation becomes smooth, and the behaviour resembles that of an ordinary motor. Typical applications include disc head drives, and small numerically-controlled machine tool slides, where the motor would drive a lead screw; and print feeds, where the motor might drive directly, or via a belt.

Most stepping motors look much like conventional motors, and as a general guide we can assume that the torque and power of a stepping motor will be similar to the torque and power of a conventional totally-enclosed motor of the same dimensions and speed range. Step angles are mostly in the range 1.8–90°, with torques ranging from 1 μNm (in a tiny wristwatch motor of 3 mm diameter) up to perhaps 40 Nm in a motor of 15 cm diameter suitable for a machine tool application where speeds of 500 rev/min might be called for, but the majority of applications use motors which can be held comfortably in the hand.

1. Many servo drives now have digital pulse train inputs, so the uniqueness of the stepper has somewhat diminished.

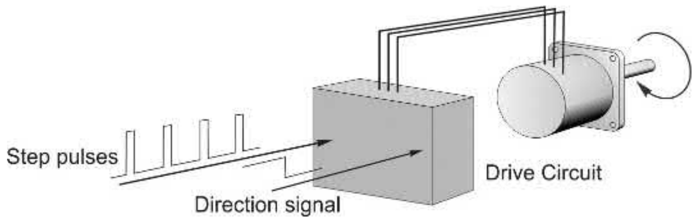

FIG. 10.1 Open-loop position control using a stepping motor.

10.2.1 Open-loop position control

A basic stepping motor system is shown in Fig. 10.1.

The drive contains the electronic switching circuits which supply the motor, and is discussed later. The output is the angular position of the motor shaft, while the input consists of two low-power digital signals. Every time a pulse occurs on the step input line, the motor takes one step, the shaft remaining at its new position until the next step pulse is supplied. The state of the direction line ('high' or 'low') determines whether the motor steps clockwise or anti-clockwise. A given number of step pulses will therefore cause the output shaft to rotate through a definite angle.

This one to one correspondence between pulses and steps is the great attraction of the stepping motor: it provides *position* control, because the output is the angular position of the output shaft. It is a *digital* system, because the total angle turned through is determined by the *number* of pulses supplied; and it is *open-loop* because no feedback need be taken from the output shaft.

10.2.2 Generation of step pulses and motor response

The step pulses may be produced by a digital controller or microprocessor (or even an oscillator controlled by an analogue voltage). When a given number of steps is to be taken, the step pulses are gated to the drive and the pulses are counted, until the required number of steps is reached, when the pulse train is gated off. This is illustrated in Fig. 10.2, for a six-step sequence. There are six step command pulses, equally spaced in time, and the motor takes one step following each pulse.

Three important general features can be identified with reference to Fig. 10.2. Firstly, although the total angle turned through (6 steps) is governed only by the number of pulses, the average speed of the shaft (which is shown by the slope of the broken line in Fig. 10.2) depends on the frequency. The higher the frequency, the shorter the time taken to complete the six steps.

Secondly, the stepping action is not perfect. The rotor takes a finite time to move from one position to the next, and then overshoots and oscillates before finally coming to rest at the new position. Overall single-step times vary with motor size, step angle and the nature of the load, but are commonly within the

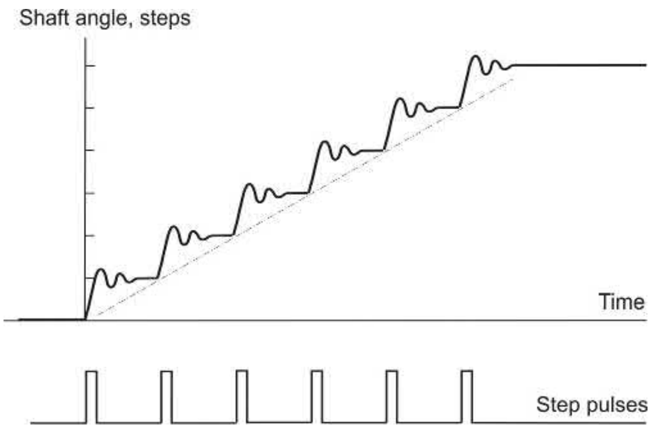

FIG. 10.2 Typical step response to low-frequency train of step command pulses.

range 5–100 ms. This is often fast enough not to be seen by the unwary newcomer, though individual steps can usually be heard; small motors 'tick' when they step, and larger ones make a satisfying 'click' or 'clunk'.

Thirdly, in order to be sure of the absolute position at the end of a stepping sequence, we must know the absolute position at the beginning. This is because a stepping motor is an incremental device. As long as it is not abused, it will always take one step when a drive pulse is supplied, but in order to keep track of absolute position simply by counting the number of drive pulses (and this is after all the main virtue of the system) we must always start the count from a known datum position. Normally the step counter will be 'zeroed' with the motor shaft at the datum position, and will then count up for clockwise direction, and down for anticlockwise rotation. Provided no steps are lost (see later) the number in the step counter will then always indicate the absolute position.

10.2.3 High speed running and ramping

The discussion so far has been restricted to operation when the step command pulses are supplied at a constant rate, and with sufficiently long intervals between the pulses to allow the rotor to come to rest between steps. Very large numbers of small stepping motors in watches and clocks do operate continuously in this way, stepping perhaps once every second, but most commercial and industrial applications call for a more exacting and varied performance.

To illustrate the variety of operations which might be involved, and to introduce high-speed running, we can look briefly at a typical industrial application. A stepping motor-driven table feed on a numerically-controlled milling machine nicely illustrates both of the key operational features discussed earlier. These are the ability to control position (by supplying the desired number of steps) and velocity (by controlling the stepping rate).

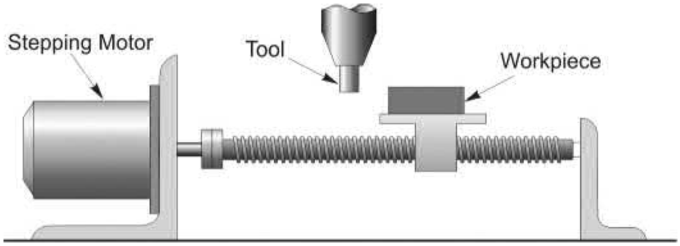

FIG. 10.3 Application of stepping motor for open-loop position control.

The arrangement is shown diagrammatically in Fig. 10.3. The motor turns a leadscrew connected to the worktable, so that each motor step causes a precise incremental movement of the workpiece relative to the cutting tool. By making the increment small enough, the fact that the motion is discrete rather than continuous will not cause any difficulties in the machining process in most applications. We will assume that we have selected the step angle, the pitch of the leadscrew, and any necessary gearing so as to give a table movement of 0.01 mm per motor step. We will also assume that the necessary step command pulses will be generated by a digital controller or computer, programmed to supply the right number of pulses, at the right speed for the work in hand.

If the machine is a general-purpose one, many different operations will be required. When taking heavy cuts, or working with hard material, the work will have to be offered to the cutting tool slowly, at say, 0.02 mm/s. The stepping rate will then have to be set to 2 steps/s. If we wish to mill out a slot 1 cm long, we will therefore programme the controller to put out 1000 steps, at a uniform rate of 2 steps/s, and then stop. On the other hand, the cutting speed in softer material could be much higher, with stepping rates in the range 10–100 steps/s being in order. At the completion of a cut, it will be necessary to traverse the work back to its original position, before starting another cut. This operation needs to be done as quickly as possible to minimise unproductive time, and a stepping rate of perhaps 2000 steps/s (or even higher), may be called for.

It was mentioned earlier that a single step (from rest) takes upwards of several milliseconds. It should therefore be clear that if the motor is to run at 2000 steps/s (i.e. 0.5 ms/step), it cannot possibly come to rest between successive steps, as it does at low stepping rates. Instead, we find in practice that at these high stepping rates, the rotor velocity becomes quite smooth, with hardly any outward hint of its stepwise origins. Nevertheless, the vital one-to-one correspondence between step command pulses and steps taken by the motor is maintained throughout, and the open-loop position control feature is preserved. This extraordinary ability to operate at very high stepping rates (e.g. at 20,000 steps/s), and yet to remain in synchronism with the command pulses, is the most striking feature of stepping motor systems.

Operation at high speeds is referred to as 'slewing'. The transition from single-stepping (as shown in Fig. 10.2) to high-speed slewing is a gradual

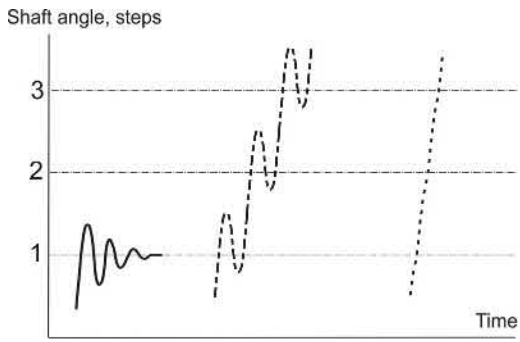

FIG. 10.4 Position-time responses at low, medium and high stepping rates.

one and is indicated by the sketches in Fig. 10.4. Roughly speaking, the motor will 'slew' if its stepping rate is above the frequency of its single-step oscillations. When motors are in the slewing range, they generally emit an audible whine, with a fundamental frequency equal to the stepping rate.

Naturally, a motor cannot be started from rest and expected to 'lock on' directly to a train of command pulses at, say, 2000 steps/s, which is well into the slewing range. Instead, it has to be started at a more modest stepping rate, before being accelerated (or 'ramped') up to speed: this is discussed more fully in Section 10.7. In undemanding applications, the ramping can be done slowly, and spread over a large number of steps; but if the high stepping rate has to be reached quickly, the timings of individual step pulses must be very precise.

We may wonder what will happen if the stepping rate is increased too quickly. The answer is simply that the motor will not be able to remain 'in step' and will stall. The step command pulses will still be being delivered, and the step counter will be accumulating what it believes are motor steps, but, by then, the system will have failed completely. A similar failure mode will occur if, when the motor is slewing, the train of step pulses is suddenly stopped, instead of being progressively slowed. The stored kinetic energy of the motor (and load) will cause it to overrun, so that the number of motor steps will be greater than the number of command pulses. Failures of this sort are prevented by the use of closed-loop control, as discussed later.

Finally, it is worth mentioning that stepping motors are designed to operate for long periods with their rotor held in a fixed (step) position, and with rated current in the winding (or windings). We can therefore anticipate that overheating after stalling is generally not a problem for a stepping motor.

10.3 Principle of motor operation

The principle on which stepping motors are based is very simple: when a bar of iron or steel is suspended so that it is free to rotate in a magnetic field, it will align itself with the field. If the direction of the field is changed, the bar will turn until it is again aligned, by the action of the so-called reluctance torque, which is the same as we have seen for the (synchronous) reluctance motor in Chapter 9.

The two most important types of stepping motor are the variable-reluctance (VR) type and the hybrid type. Both types utilise the reluctance principle, the difference between them lying in the method by which the magnetic fields are produced. In the VR type the fields are produced solely by sets of stationary current-carrying windings. The hybrid type also has sets of windings, but the addition of a permanent magnet (on the rotor) gives rise to the description 'hybrid' for this type of motor. Although both types of motor work on the same basic principle, it turns out in practice that the VR type is attractive for the larger step angles (e.g. 15°, 30°, 45°), while the hybrid tends to be best-suited when small angles (e.g. 1.8°, 2.5°) are required. (We should acknowledge that small versions of what we described in Chapter 9 as PM synchronous motors are also used as stepping motors, particularly where a large step angle is required, but we will not be discussing them in this book.)

10.3.1 Variable reluctance motor

A simplified diagram of a 30°/step VR stepping motor is shown in Fig. 10.5. The stator is made from a stack of steel laminations, and has six equally-spaced projecting poles, or teeth, each carrying a separate coil. The rotor, which may be solid or laminated, has four projecting teeth, of the same width as the stator teeth. There is a very small air gap—typically between 0.02 mm and 0.2 mm—between rotor and stator teeth. When no current is flowing in any of the stator coils, the rotor will therefore be completely free to rotate.

Diametrically opposite pairs of stator coils are connected in series, such that when one of them acts as a North pole, the other acts as a South pole. There are thus three independent stator circuits, or phases, and each one can be supplied with direct current from the drive circuit (not shown in Fig. 10.5).

When phase A is energised (as indicated by the thick lines in Fig. 10.5A), a magnetic field with its axis along the stator poles of phase A is created. The rotor is therefore attracted into a position where the pair of rotor poles distinguished

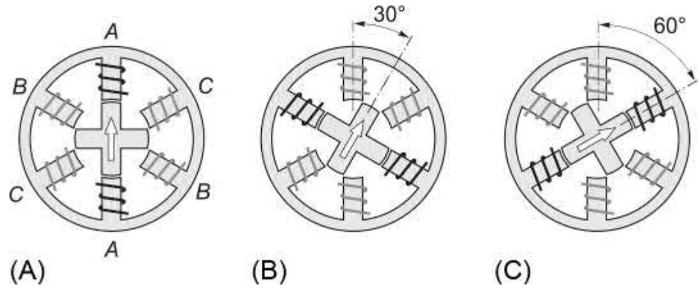

FIG. 10.5 Principle of operation of 30°/step variable-reluctance stepping motor.

by the marker arrow line up with the field, i.e. in line with the phase-A pole, as shown in Fig. 10.5A. When phase A is switched off, and phase B is switched on instead, the second pair of rotor poles will be pulled into alignment with the stator poles of phase B, the rotor moving through 30° clockwise to its new step position, as shown in Fig. 10.5B. A further clockwise step of 30° will occur when phase B is switched off and phase C is switched on. At this stage the original pair of rotor poles come into play again, but this time they are attracted to stator poles C, as shown in Fig. 10.5C. By repetitively switching-on the stator phases in the sequence ABCA, etc. the rotor will rotate clockwise in 30° steps, while if the sequence is ACBA, etc. it will rotate anticlockwise. This mode of operation is known as '1-phase-on', and is the simplest way of making the motor step. Note that the polarity of the energising current is not significant: the motor will be aligned equally well regardless of the direction of current.

This example demonstrates an interesting difference between doubly-salient reluctance motors and the singly-salient types discussed in Chapter 9. In a singly-salient machine, the pole-number of the field produced by the stator is determined by the winding layout and is the same as the number of rotor saliences, and both rotate together in the same direction. In most doubly-salient types, however, the rotor rotates in the opposite direction from the progression of the stator excitation: in Fig. 10.5 for example, the stator (2-pole) excitation axis rotates 60° anticlockwise each step, while the (4-pole) rotor turns 30° clockwise each step.

10.3.2 Hybrid motor

A cross-sectional view of a typical 1.8° hybrid motor is shown in Fig. 10.6. The stator has 8 main poles, each with 5 teeth, and each main pole carries a simple coil. The rotor has two steel end-caps, each with 50 teeth, and separated by a permanent magnet.

The rotor teeth have the same pitch as the teeth on the stator poles, and are offset so that the centreline of a tooth at one end-cap coincides with a slot at the other end-cap. The permanent-magnet is axially magnetised, so that one set of rotor teeth is given a North polarity, and the other a South polarity. Extra torque is obtained by adding more stacks, and stretching the stator, as shown in Fig. 10.7.

When no current is flowing in the windings, the only source of magnetic flux across the air-gap is the permanent magnet. The magnet flux crosses the air-gap from the N end-cap into the stator poles, flows axially along the body of the stator, and returns to the magnet by crossing the air-gap to the S end-cap. If there were no offset between the two sets of rotor teeth, there would be a strong periodic alignment torque when the rotor was turned, and every time a set of stator teeth was in line with the rotor teeth we would obtain a stable equilibrium position. However there is an offset, and this causes the alignment torque due to the magnet to be almost eliminated. In practice a small 'detent' torque remains, and

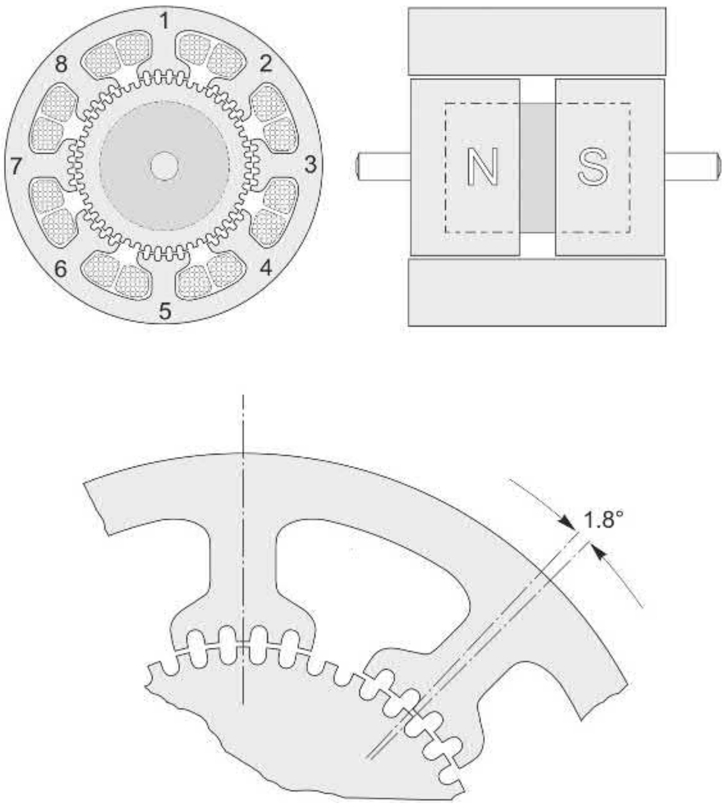

FIG. 10.6 Hybrid (200 step/rev) stepping motor. The detail shows the rotor and stator tooth alignments, and indicates the step angle of 1.8°.

this can be felt if the shaft is turned when the motor is de-energised: the motor tends to be held in its step positions by the detent torque. This is sometimes very useful: for example it is usually enough to hold the rotor stationary when the power is switched-off, so the motor can be left without fear of it being accidentally nudged into to a new position.

The 8 coils are connected to form two phase-windings. The coils on poles 1, 3, 5, and 7 form phase A, while those on 2, 4, 6, and 8 form phase B. When phase A carries positive current stator poles 1 and 5 are magnetised as South, and poles 3 and 7 become North. The teeth on the North end of the rotor are attracted to poles 1 and 5 while the offset teeth at the South end of the rotor are attracted into line with the teeth on poles 3 and 7. To make the rotor step, phase A is switched off, and phase B is energised with either positive current or negative current, depending on the sense of rotation required. This will cause the rotor to move by one quarter of a tooth pitch (1.8°) to a new equilibrium (step) position.

FIG. 10.7 Rotor of size 34 (3.4 in. or 8 cm diameter) 3-stack hybrid 1.8° stepping motor. The dimensions of the rotor end-caps and the associated axially-magnetised permanent magnet are optimised for the single-stack version. Extra torque is obtained by adding a second or third stack, the stator simply being stretched to accommodate the longer rotor. *(Courtesy of Astrosyn International Technology Ltd.)*

The motor is continuously stepped by energising the phases in the sequence +A, −B, −A, +B, +A (clockwise) or + A, +B, −A, −B, +A (anticlockwise). It will be clear from this that a bipolar supply is needed (i.e. one which can furnish +ve or −ve current). When the motor is operated in this way it is referred to as '2-phase, with bipolar supply'.

If a bipolar supply is not available, the same pattern of pole energisation may be achieved in a different way, as long as the motor windings consist of two identical ('bifilar wound') coils. To magnetise pole 1 North, a positive current is fed into one set of phase-A coils. But to magnetise pole 1 South, the same positive current is fed into the other set of phase-A coils, which have the opposite winding sense. In total, there are then four separate windings, and when the motor is operated in this way it is referred to as '4-phase, with unipolar supply'. Since each winding only occupies half of the space, the m.m.f. of each winding is only half of that of the full coil, so the thermally-rated output is clearly reduced as compared with bipolar operation (for which the whole winding is used).

The 200 step/rev hybrid is the most widely-used general-purpose stepper, and is available in a range of sizes, as shown in Fig. 10.8.

We round-off this section on hybrid motors with a comment on identifying windings, and a warning. If the motor details are not known, it is usually possible to identify bifilar windings by measuring the resistance from the common to the two ends. If the motor is intended for unipolar drive only, one end of each winding may be commoned inside the casing; for example a 4-phase unipolar motor may have only five leads, one for each phase and one common. Wires are also usually colour-coded to indicate the location of the windings; for example a

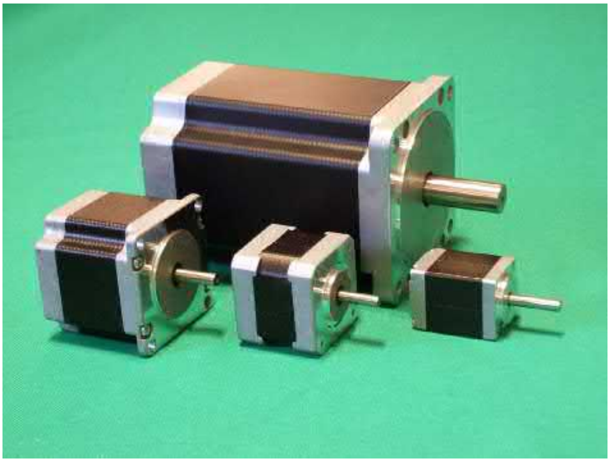

FIG. 10.8 Hybrid 1.8° stepping motors, of sizes 34 (3.4 in. diameter), 23 and 17 and 11. *(Courtesy of Astrosyn International Technology Ltd.)*

bifilar winding on one set of poles will have one end red, the other end red and white, and the common white. Finally, it is not advisable to remove the rotor of a hybrid motor because they are magnetised in-situ: removal typically causes a 5–10% reduction in magnet flux, with a corresponding reduction in static torque at rated current.

10.3.3 Summary

The construction of stepping motors is simple and robust, the only moving part being the rotor, which has no windings, commutator or brushes. The rotor is held at its step position solely by the action of the magnetic flux between stator and rotor. The step angle is a property of the tooth geometry and the arrangement of the stator windings, and accurate punching and assembly of the stator and rotor laminations is therefore necessary to ensure that adjacent step positions are exactly equally spaced. Any errors due to inaccurate punching will be non-cumulative, however.

The step angle is obtained from the expression

$$step\ angle = \frac{360^\circ}{(rotor\ teeth) \times (stator\ phases)}$$

The VR motor in Fig. 10.5 has 4 rotor teeth, 3 stator phase-windings, and the step angle is therefore 30°, as already shown. It should also be clear from the equation why small angle motors always have to have a large number of rotor teeth: the 200 step/rev hybrid type (see Fig. 10.6) has a 50-tooth rotor, 4-phase stator, and hence a step angle of 1.8° (= 360°/(50 × 4).

The magnitude of the aligning torque clearly depends on the magnitude of the current in the phase-winding. However, the equilibrium positions itself does not depend on the magnitude of the current, because it is simply the position where the rotor and stator teeth are in line. This property underlines the digital nature of the stepping motor.

10.4 Motor characteristics

10.4.1 Static torque–displacement curves

From the previous discussion, it should be clear that the shape of the torque–displacement curve, and the peak static torque, will depend on the internal electromagnetic design of the rotor. In particular the shapes of the rotor and stator teeth, and the disposition of the stator windings (and permanent magnet(s)) all have to be optimised to obtain the maximum static torque.

We now turn to a typical static torque–displacement curve, and look at how it determines motor behaviour. Several aspects will be discussed, including the explanation of basic stepping (which has already been looked at in a qualitative way); the influence of load torque on step position accuracy; the effect of the amplitude of the winding current; and half-step and mini-stepping operation. For the sake of simplicity, the discussion will be based on the 30°/step 3-phase VR motor introduced earlier, but the conclusions reached apply to any stepping motor.

Typical static torque–displacement curves for a 3-phase 30°/step VR motor are shown in Fig. 10.9. These show the torque that has to be applied to move the rotor away from its aligned position. Because of the rotor/stator symmetry, the magnitude of the restoring torque when the rotor is displaced by a given angle in one direction is the same as the magnitude of the restoring torque when it is displaced by the same angle in the other direction, but of opposite sign.

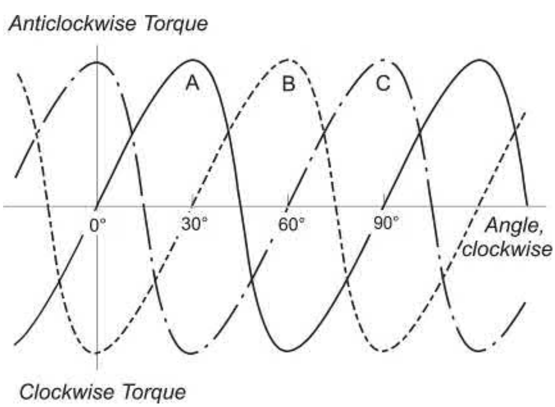

FIG. 10.9 Static torque–displacement curves for 30°/step variable reluctance stepping motor.

There are 3 curves in Fig. 10.9, one for each of the three phases, and for each curve we assume that the relevant phase winding carries its full (rated) current. If the current is less than rated, the peak torque will be reduced, and the shape of the curve is likely to be somewhat different. The convention used in Fig. 10.9 is that a clockwise displacement of the rotor corresponds to a movement to the right, while a positive torque tends to move the rotor anticlockwise.

When only one phase, say A, is energised, the other two phases exert no torque, so their curves can be ignored and we can focus attention on the solid line in Fig. 10.9. Stable equilibrium positions (for phase A excited) exist at $\theta = 0°$, 90°, 180° and 270°. They are stable (step) positions because any attempt to move the rotor away from them is resisted by a counteracting or restoring torque. These points correspond to positions where successive rotor poles (which are 90° apart) are aligned with the stator poles of phase A, as shown in Fig. 10.5A. There are also four unstable equilibrium positions, (at $\theta = 45°$, 135°, 225° and 315°) at which the torque is also zero. These correspond to rotor positions where the stator poles are mid-way between two rotor poles, and they are unstable because if the rotor is deflected slightly in either direction, it will be accelerated in the same direction until it reaches the next stable position. If the rotor is free to turn, it will therefore always settle in one of the four stable positions.

10.4.2 Single-stepping

If we assume that phase A is energised, and the rotor is at rest in the position $\theta = 0°$ (Fig. 10.9) we know that if we want to step in a clockwise direction, the phases must be energised in the sequence ABCA, etc., so we can now imagine that phase A is switched off, and phase B is energised instead. We will also assume that the decay of current in phase A and the build-up in phase B take place very rapidly, before the rotor moves significantly.

The rotor will find itself at $\theta = 0°$, but it will now experience a clockwise torque (see Fig. 10.9) produced by phase B. The rotor will therefore accelerate clockwise, and will continue to experience clockwise torque, until it has turned through 30°. The rotor will be accelerating all the time, and it will therefore overshoot the 30° position, which is of course its target (step) position for phase B. As soon as it overshoots, however, the torque reverses, and the rotor experiences a braking torque, which brings it to rest before accelerating it back towards the 30° position. If there was no friction or other cause of damping, the rotor would continue to oscillate; but in practice it comes to rest at its new position quite quickly in much the same way as a damped second-order system. The next 30° step is achieved in the same way, by switching-off the current in phase B, and switching-on phase C.

In the discussion above, we have recognised that the rotor is acted on sequentially by each of the three separate torque curves shown in Fig. 10.9. Alternatively, since the three curves have the same shape, we can think of the rotor being influenced by a single torque curve which ‘jumps’ by one step

(30° in this case) each time the current is switched from one phase to the next. This is often the most convenient way of visualising what is happening in the motor.

10.4.3 Step position error, and holding torque

In the previous discussion the load torque was assumed to be zero, and the rotor was therefore able to come to rest with its poles exactly in line with the excited stator poles. When load torque is present, however, the rotor will not be able to pull fully into alignment, and a 'step position error' will be unavoidable.

The origin and extent of the step position error can be appreciated with the aid of the typical torque–displacement curve shown in Fig. 10.10. The true step position is at the origin in the figure, and this is where the rotor would come to rest in the absence of load torque. If we imagine the rotor is initially at this position, and then consider that a clockwise load (T_L) is applied, the rotor will move clockwise, and as it does so it will develop progressively more anticlockwise torque. The equilibrium position will be reached when the motor torque is equal and opposite to the load torque, i.e. at point A in Fig. 10.10. The corresponding angular displacement from the step position (θ_e in Fig. 10.10) is the step position error.

The existence of a step position error is one of the drawbacks of the stepping motor. The motor designer attempts to combat the problem by aiming to produce a steep torque–angle curve around the step position, and the user has to be aware of the problem and choose a motor with a sufficiently steep curve to keep the error within acceptable limits. In some cases this may mean selecting a motor with a higher peak torque than would otherwise be necessary, simply to obtain a steep enough torque curve around the step position.

As long as the load torque is less than T_{max} (see Fig. 10.10), a stable rest position is obtained, but if the load torque exceeds T_{max}, the rotor will be unable to hold its step position. T_{max} is therefore known as the 'holding' torque. The

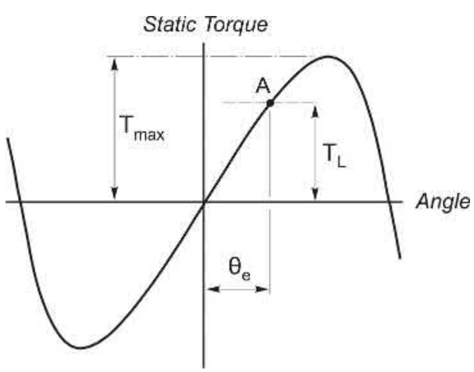

FIG. 10.10 Static torque–angle curve showing step position error (θ_e) resulting from load torque T_L.

value of the holding torque immediately conveys an idea of the overall capability of any motor, and it is—after step angle—the most important single parameter which is looked for in selecting a motor. Often, the adjective ‘holding’ is dropped altogether: for example ‘a 1 Nm motor’ is understood to be one with a peak static torque (holding torque) of 1 Nm.

10.4.4 Half stepping

We have already seen how to step the motor in 30° increments by energising the phases one at a time in the sequence ABCA, etc. Although this ‘1-phase-on’ mode is the simplest and most widely used, there are two other modes which are also frequently employed. These are referred to as the ‘2-phase-on’ mode and the ‘half-stepping’ mode. The 2-phase-on can provide greater holding torque and a much better damped single-step response than the 1-phase-on mode; and the half stepping mode permits the effective step angle to be halved—thereby doubling the resolution - and produces a smoother shaft rotation.

In the 2-phase-on mode, two phases are excited simultaneously. When phases A and B are energised, for example, the rotor experiences torques from both phases, and comes to rest at a point midway between the two adjacent full step positions. If the phases are switched in the sequence AB, BC, CA, AB, etc., the motor will take full (30°) steps, as in the 1-phase-on mode, but its equilibrium positions will be interleaved between the full-step positions.

To obtain ‘half-stepping’ the phases are excited in the sequence A, AB, B, BC, etc., i.e. alternately in the 1-phase-on and 2-phase-on modes. This is sometimes known as ‘wave’ excitation, and it causes the rotor to advance in steps of 15°, or half the full step angle. As might be expected, continuous half-stepping usually produces a smoother shaft rotation than full-stepping, and it also doubles the resolution.

We can see what the static torque curve looks like when two phases are excited by superposition of the individual phase curves. An example is shown in Fig. 10.11, from which it can be seen that for this machine, the holding torque (i.e. the peak static torque) is higher with two phases excited than with only one excited. The stable equilibrium (half-step) position is at 15°, as expected. The price to be paid for the increased holding torque is the increased power dissipation in the windings, which is doubled as compared with the 1-phase-on mode. The holding torque increases by a factor less than two, so the torque per watt (which is a useful figure of merit) is reduced.

A word of caution is needed in regard to the addition of the two separate 1-phase-on torque curves to obtain the 2-phase-on curve. Strictly, such a procedure is only valid where the two phases are magnetically independent, or the common parts of the magnetic circuits are unsaturated. This is not the case in most motors, in which the phases share a common magnetic circuit which operates under highly saturated conditions. Direct addition of the 1-phase-on

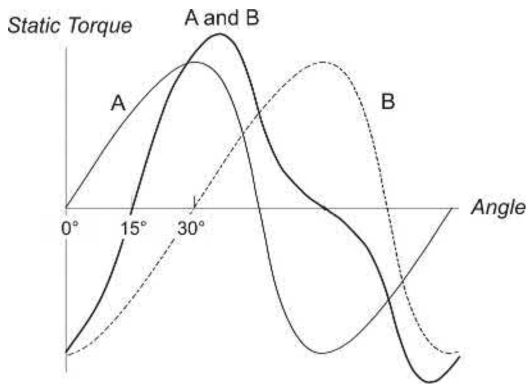

FIG. 10.11 Static torque curve (solid line) corresponding to two-phase-on excitation.

curves cannot therefore be expected to give an accurate result for the 2-phase-on curve, but it is easy to do, and provides a reasonable estimate.

Apart from the higher holding torque in the 2-phase-on mode, there is another important difference which distinguishes the static behaviour from that of the 1-phase-on mode. In the 1-phase-on mode, the equilibrium or step positions are determined solely by the geometry of the rotor and stator: they are the positions where the rotor and stator are in line. In the 2-phase-on mode, however, the rotor is intended to come to rest at points where the rotor poles are lined-up midway between the stator poles. This position is not sharply defined by the 'edges' of opposing poles, as in the 1-phase-on case; and the rest position will only be exactly midway if (a) there is exact geometrical symmetry and, more importantly (b) the two currents are identical. If one of the phase currents is larger than the other, the rotor will come to rest closer to the phase with the higher current, instead of half-way between the two. The need to balance the currents to obtain precise half stepping is clearly a drawback to this scheme. Paradoxically, however, the properties of the machine with unequal phase currents can sometimes be turned to good effect, as we now see.

10.4.5 Step division—Mini-stepping

There are some applications where very fine resolution is called for, and a motor with a very small step angle—perhaps only a fraction of a degree—is required. We have already seen that the step angle can only be made small by increasing the number of rotor teeth and/or the number of phases, but in practice it is inconvenient to have more than four or five phases, and it is difficult to manufacture rotors with more than 50–100 teeth. This means it is rare for motors to have step angles below about 1°. When a smaller step angle is required a technique known as mini-stepping (or step division) is used.

Mini-stepping is a technique based on 2-phase-on operation which provides for the sub-division of each full motor step into a number of ‘substeps’ of equal size. In contrast with half-stepping, where the two currents have to be kept equal, the currents are deliberately made unequal. By correctly choosing and controlling the relative amplitudes of the currents, the rotor equilibrium position can be made to lie anywhere between the step positions for each of the two separate phases.

Closed-loop current control is needed to prevent the current from changing as a result of temperature changes in the windings, or variations in the supply voltage; and if it is necessary to ensure that the holding torque stays constant for each mini-step both currents must be changed according to a prescribed algorithm. Despite the difficulties referred to above, mini-stepping is used extensively, especially in photographic and printing applications where a high resolution is needed. Schemes involving between 3 and 10 mini-steps for a 1.8° step motor are numerous, and there are instances where over 100 mini-steps (20,000 mini-steps/rev) have been achieved.

So far, we have concentrated on those aspects of behaviour which depend only on the motor itself, i.e. the static performance. The shape of the static torque curve, the holding torque, and the slope of the torque curve about the step position have all been shown to be important pointers to the way the motor can be expected to perform. All of these characteristics depend on the current(s) in the windings, however, and when the motor is running the instantaneous currents will depend on the type of drive circuit employed.

10.5 Steady-state characteristics—Ideal (constant-current) drive

In this section we will look first at how the motor would perform if it were supplied by an ideal drive circuit, which turns out to be one that is capable of supplying rectangular pulses of current to each winding when required, and regardless of the stepping rate. Because of the inductance of the windings, no real drive circuit will be able to achieve this, but the most sophisticated (and expensive) ones achieve near-ideal operation up to very high stepping rates.

10.5.1 Requirements of drive

The basic function of the complete drive is to convert the step command input signals into appropriate patterns of currents in the motor windings. This is achieved in two distinct stages, as shown in Fig. 10.12, which relates to a 3-phase motor.

The ‘translator’ stage converts the incoming train of step command pulses into a sequence of on/off commands to each of the three power stages. In the 1 phase-on mode, for example, the first step command pulse will be routed to turn

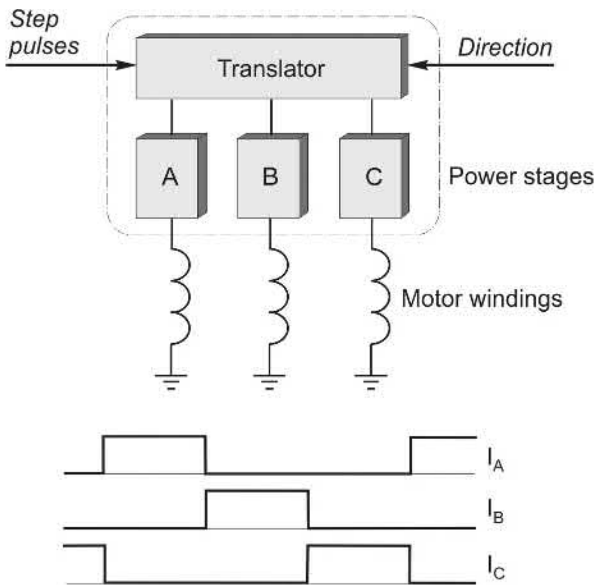

FIG. 10.12 General arrangement of drive system for 3-phase motor, and winding currents corresponding to an 'ideal' drive.

on phase A, the second will turn on phase B, and so on. In a very simple drive, the translator will probably provide for only one mode of operation (e.g. 1-phase-on), but most commercial drives provide the option of 1-phase-on, 2-phase on and half-stepping. Single-chip I.C.'s with these 3 operating modes and with both 3-phase and 4-phase outputs are readily available.

The power stages (one per phase) supply the current to the windings. An enormous diversity of types are in use, ranging from simple ones with one switching transistor per phase, to elaborate chopper-type circuits with four transistors per phase, and some of these are discussed in Section 10.6. At this point however, it is helpful to list the functions required of the 'ideal' power stage. These are firstly that when the translator calls for a phase to be energised, the full (rated) current should be established immediately; secondly, the current should be maintained constant (at its rated value) for the duration of the 'on' period; and finally, when the translator calls for the current to be turned off, it should be reduced to zero immediately.

The ideal current waveforms for continuous stepping with 1-phase-on operation are shown in the lower part of Fig. 10.12. The currents have a square profile because this leads to the optimum value of running torque from the motor. But because of the inductance of the windings, no real drive will achieve the ideal current wave-forms, though many drives come close to the ideal, even at quite high stepping rates. Drives which produce such rectangular current waveforms are (not surprisingly) called constant-current drives. We now look at the running torque produced by a motor when operated from an ideal constant current drive. This will act as a yardstick for assessing the performance of other drives, all of which will be seen to have inferior performance.

10.5.2 Pull-out torque under constant-current conditions

If the phase currents are taken to be ideal, i.e. they are switched on and off instantaneously, and remain at their full rated value during each 'on' period, we can picture the axis of the magnetic field to be advancing around the machine in a series of steps, the rotor being urged to follow it by the reluctance torque. If we assume that the inertia is high enough for fluctuations in rotor velocity to be very small, the rotor will be rotating at a constant rate which corresponds exactly to the stepping rate.

Now if we consider a situation where the position of the rotor axis is, on average, lagging behind the advancing field axis, it should be clear that, on average, the rotor will experience a driving torque. The more it lags behind, the higher will be the average forward torque acting on it, but only up to a point. We already know that if the rotor axis is displaced too far from the field axis, the torque will begin to diminish, and eventually reverse, so we conclude that although more torque will be developed by increasing the rotor lag angle, there will be a limit to how far this can be taken.

Turning now to a quantitative examination of the torque on the rotor, we will make use of the static torque–displacement curves discussed earlier, and look at what happens when the load on the shaft is varied, the stepping rate being kept constant. Clockwise rotation will be studied, so the phases will be energised in the sequence ABC. The instantaneous torque on the rotor can be arrived at by recognising (a) that the rotor speed is constant, and it covers one step angle (30°) between step command pulses, and (b) the rotor will be 'acted on' sequentially by each of the set of torque curves.

When the load torque is zero, the net torque developed by the rotor must be zero (apart from a very small torque required to overcome friction). This condition is shown in Fig. 10.13A. The instantaneous torque is shown by the thick line, and it is clear that each phase in turn exerts first a clockwise torque, then an anticlockwise torque while the rotor angle turns through 30°. The average torque is zero, the same as the load torque, because the average rotor lag angle is zero.

When the load torque on the shaft is increased, the immediate effect is to cause the rotor to fall back in relation to the field. This causes the clockwise torque to increase, and the anticlockwise torque to decrease. Equilibrium is reached when the lag angle has increased sufficiently for the motor torque to equal the load torque. The torque developed at an intermediate load condition like this is shown by the thick line in Fig. 10.13B. The highest average torque that can possibly be developed is shown by the thick line in Fig. 10.13C: if the

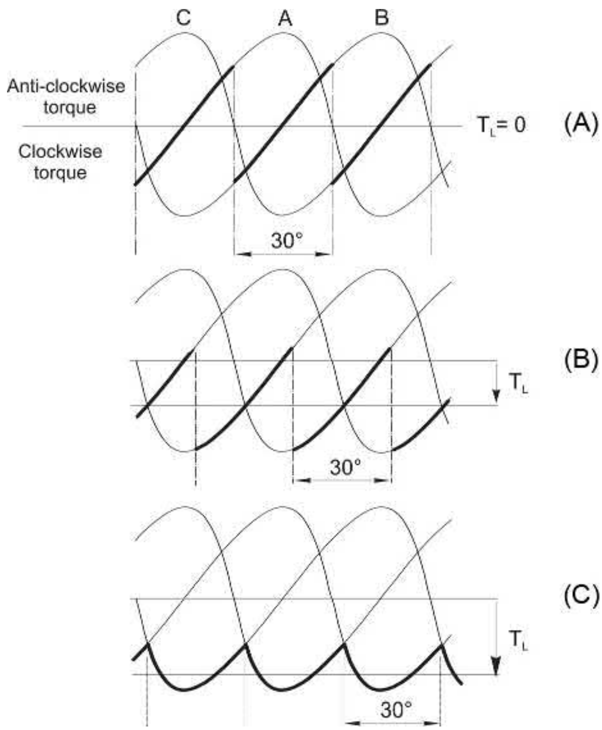

FIG. 10.13 Static torque curves indicating how the average steady-state torque (T_L) is developed during constant-frequency operation.

load torque exceeds this value (which is known as the pull-out torque) the motor loses synchronism and stalls, and the vital one-to-one correspondences between pulses and steps is lost.

Since we have assumed an ideal constant-current drive, the pull-out torque will be independent of the stepping rate, and the pull-out torque–speed curve under ideal conditions is therefore as shown in Fig. 10.14. The shaded area represents the permissible operating region: at any particular speed (stepping rate) the load torque can have any value up to the pull-out torque, and the motor will continue to run at the same speed. But if the load torque exceeds the pull-out torque, the motor will suddenly pull out of synchronism and stall.

As mentioned earlier, no real drive will be able to provide the ideal current waveforms, so we now turn to look briefly at the types of drives in common use, and at their pull-out torque–speed characteristics.

10.6 Drive circuits and pull-out torque–speed curves

Users often find difficulty in coming to terms with the fact that the running performance of a stepping motor depends so heavily on the type of drive circuit

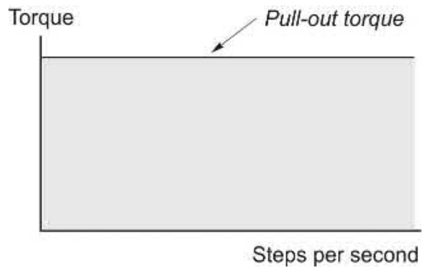

FIG. 10.14 Steady-state operating region with ideal constant-current drive. (In such idealised circumstances there would be no limit to the stepping rate, but as shown in Fig. 10.18, a real drive circuit imposes an upper limit.)

being used. It is therefore important to emphasise that in order to meet a specification, it will always be necessary to consider the motor and drive together, as a package.

There are three commonly-used types of drive. All use transistors which are operated as switches, i.e. they are either turned fully on, or they are cut-off. A brief description of each is given below, and the pros and cons of each type are indicated. In order to simplify the discussion, we will consider one phase of a 3-phase VR motor and assume that it can be represented by a simple series R-L circuit in which R and L are the resistance and self-inductance of the winding respectively. (In practice the inductance will vary with rotor position, giving rise to motional e.m.f. in the windings, which, as we have seen previously in this book, is an inescapable manifestation of an electromechanical energy-conversion process. If we needed to analyse stepping motor behaviour fully we would have to include the motional e.m.f. terms. Fortunately, we can gain a pretty good appreciation of how the motor behaves if we model each winding simply in terms of its resistance and self-inductance.)

10.6.1 Constant voltage drive

This is the simplest possible drive: the circuit for one of the three phases is shown in the upper part of Fig. 10.15, and the current waveforms at low and high stepping rates are shown in the lower part of the figure. The d.c. voltage V is chosen so that when the switching device (usually a MOSFET although a BJT is shown) turns on, the steady current is the rated current as specified by the motor manufacturer.

The current waveforms display the familiar rising exponential shape that characterises a first-order system: the time-constant is L/R, the current reaching its steady state after several time-constants. When the transistor switches off, the stored energy in the inductance cannot instantaneously reduce to zero, so although the current through the transistor suddenly becomes zero, the current in the winding is diverted into the closed path formed by the winding and the

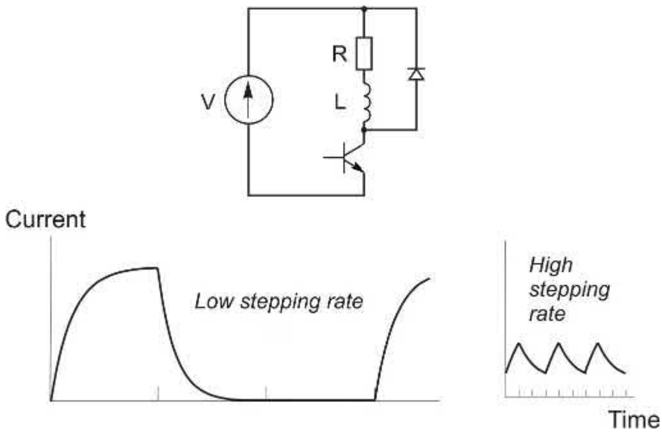

FIG. 10.15 Basic constant-voltage drive circuit and typical current waveforms.

freewheel diode, and it then decays exponentially to zero, again with time-constant L/R: in this period the stored energy in the magnetic field is dissipated as heat in the resistance of the winding and diode.

At low stepping rates (low speed), the drive provides a reasonably good approximation to the ideal rectangular current waveform. (We are considering a 3-phase motor, so ideally one phase should be on for one step pulse and off for the next two, as in Fig. 10.12.) But at higher frequencies (right-hand waveform in Fig. 10.15), where the 'on' period is short compared with the winding time-constant, the current waveform degenerates, and is nothing like the ideal rectangular shape. In particular the current never gets anywhere near its full value during the on pulse, so the torque over this period is reduced; and even worse, a substantial current persists when the phase is supposed to be off, so during this period the phase will contribute a negative torque to the rotor. Not surprisingly all this results in a very rapid fall-off of pull-out torque with speed, as shown later in Fig. 10.18A.

Curve (A) in Fig. 10.18 should be compared with the pull-out torque under ideal constant-current conditions shown in Fig. 10.14 in order to appreciate the severely limited performance of the simple constant-voltage drive.

10.6.2 Current-forced drive

The initial rate of rise of current in a series R-L circuit is directly proportional to the applied voltage, so in order to establish the current more quickly at switch-on, a higher supply voltage (V_f) is needed. But if we simply increased the voltage, the steady-state current (V_f/R) would exceed the rated current and the winding would overheat.

To prevent the current from exceeding the rated value, an additional 'forcing' resistor has to be added in series with the winding. The value of this

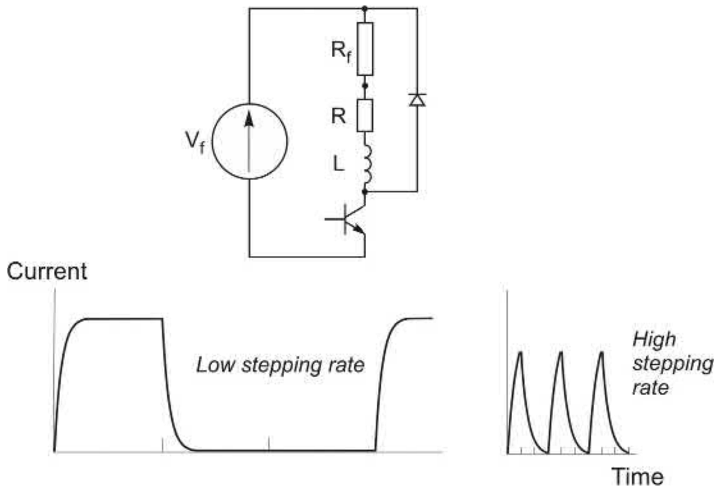

FIG. 10.16 Current-forced (L/R) drive and typical current waveforms.

resistance (R_f) must be chosen so that ($V_f/(R+R_f)=I$,where I is the rated current. This is shown in the upper part of Fig. 10.16, together with the current waveforms at low and high stepping rates. Because the rates of rise and fall of current are higher, the current waveforms approximate more closely to the ideal rectangular shape, especially at low stepping rates, though at higher rates they are still far from ideal, as shown in Fig. 10.16. The low-frequency pull-out torque is therefore maintained to a higher stepping rate, as shown in Fig. 10.18B. Values of R_f from 2 to 10 times the motor resistance (R) are common. Broadly speaking, if $R_f=10R$, for example, a given pull-out torque will be available at ten-times the stepping rate, compared with an unforced constant-voltage drive.

Manufacturers sometimes call this type of drive an 'R/L' drive, or an 'L/R' drive, or even simply a 'Resistor Drive'. Sets of pull-out torque speed curves in catalogues may be labelled with values R/L (or L/R)$=5$, 10, etc. This means that the curves apply to drives where the forcing resistor is five (or ten) times the winding resistance, the implication being that the drive voltage has also been adjusted to keep the static current at its rated value. Obviously, it follows that the higher R_f is made, the higher the power rating of the supply; and it is the higher power rating which is the principal reason for the improved torque–speed performance.

The major disadvantage of this drive is its inefficiency, and the consequent need for a high power-supply rating. Large amounts of heat are dissipated in the forcing resistors, especially when the motor is at rest and the phase-current is continuous, and disposing of the heat can lead to awkward problems in the siting of the forcing resistors. These drives are therefore only used at the low-power end of the scale.

It was mentioned earlier that the influence of the motional e.m.f. in the winding would be ignored. In practice, however, the motional e.m.f. always has a pronounced influence on the current, especially at high stepping rates, so it must be borne in mind that the waveforms shown in Figs. 10.15 and 10.16 are only approximate. Not surprisingly, it turns out that the motional e.m.f. tends to make the current waveforms worse (and the torque less) than the discussion above suggests. Ideally therefore, we need a drive which will keep the current constant throughout the on period, regardless of the motional e.m.f. The closed-loop chopper-type drive (below) provides the closest approximation to this, and it avoids the waste of power which is a feature of R/L drives, so this type is now the most widely used.

10.6.3 Constant current (chopper) drive

The basic circuit for one phase of a VR motor is shown in the upper part of Fig. 10.17, with the current wave-forms shown below. A high-voltage power supply is used in order to obtain very rapid changes in current when the phase is switched on or off.

The lower transistor is turned on for the whole period during which current is required. The upper transistor turns on whenever the actual current falls below the lower threshold of the hysteresis band (shown dotted in Fig. 10.17) and it turns off when the current exceeds the upper threshold. The chopping action leads to a current waveform which is a good approximation to the ideal (see Fig. 10.12). At the end of the on period both transistors switch off and the

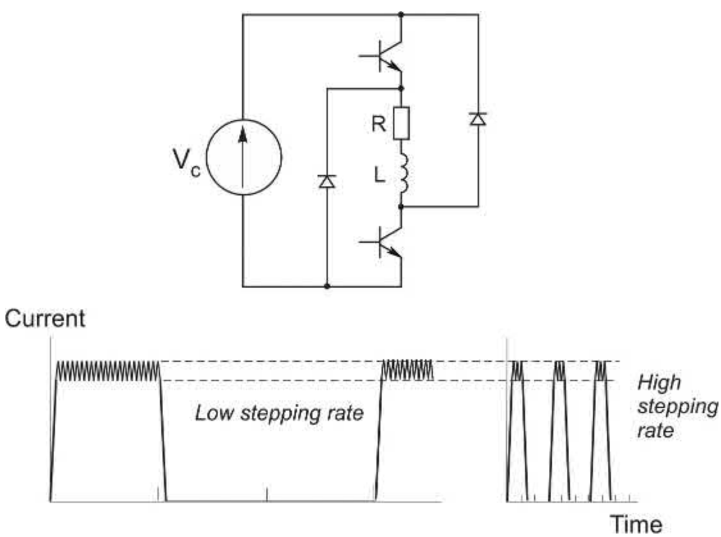

FIG. 10.17 Constant-current chopper drive and typical current waveforms.

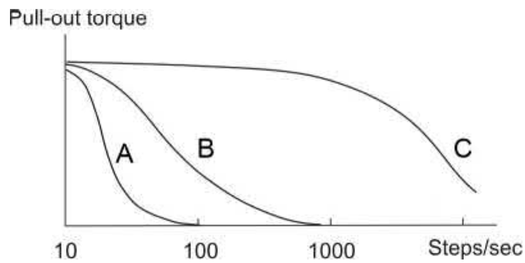

FIG. 10.18 Typical pull-out torque–speed curves for a given motor with different types of drive circuit. (A) Constant voltage drive; (B) Current-forced drive; (C) chopper drive.

current freewheels through both diodes and back to the supply. During this period the stored energy in the inductance is returned to the supply, and because the winding terminal voltage is then $-V_c$, the current decays as rapidly as it built up.

Because the current-control system is a closed-loop one, distortion of the current waveform by the motional e.m.f. is minimised, and this means that the ideal (constant-current) torque–speed curve is closely followed up to high stepping rates. Eventually, however, the 'on' period reduces to the point where it is less than the current rise time, and the full current is never reached. Chopping action then ceases, the drive reverts essentially to a constant-voltage one, and the torque falls rapidly as the stepping rate is raised even higher, as in Fig. 10.18C. There is no doubt of the overall superiority of the chopper-type drive, and it is now the standard drive. Single-chip chopper modules can be bought for small (say 1–2 A) motors: complete plug-in chopper cards rated up to 10 A or more are available for larger motors; and increasingly, motors may have the drive integrated with the motor housing.

The discussion in this section relates to a VR motor, for which unipolar current pulses are sufficient. If we have a hybrid or other permanent-magnet motor we will need a bipolar current source (i.e. one that can provide positive or negative current), and for this we will find that each phase is supplied from a 4-transistor H-bridge, as discussed in Chapter 2. A typical bipolar drive for a hybrid motor is shown in Fig. 10.19.

10.6.4 Resonances and instability

In practice, measured torque–speed curves frequently display severe dips at or around certain stepping rates: a typical measured characteristic for a hybrid motor with a voltage-forced drive is shown as (a) in Fig. 10.20. Manufacturers are not keen to stress this feature, so it is important for the user to be aware of the potential difficulty.

FIG. 10.19 Bipolar drive with integral heatsink providing a range of step divisions up to 25,600 steps/rev. It allows drive current to be set from 0.25 A to 2.0 A with supply voltage from 14 V to 40 V. *(Courtesy of Astrosyn International Technology Ltd.)*

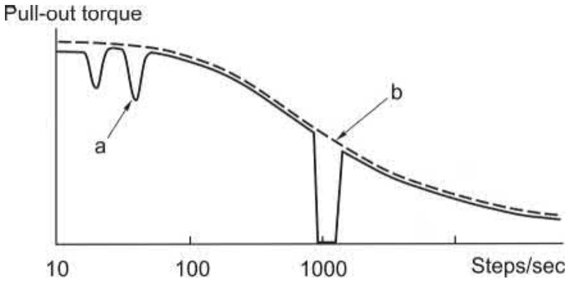

FIG. 10.20 Pull-out torque–speed curves for a hybrid stepping motor showing (curve a) low-speed resonance dips and mid-frequency instability at around 1000 steps/s; and improvement brought about by fitting an inertial damper (curve b).

The magnitude and location of the torque dips depend in a complex way on the characteristics of the motor, the drive, the operating mode and the load. We will not go into detail here, apart from mentioning the underlying causes and remedies.

There are two distinct mechanisms which cause the dips. The first is a straightforward 'resonance-type' problem which manifests itself at low stepping rates, and originates from the oscillatory nature of the single-step response. Whenever the stepping rate coincides with the natural frequency of rotor oscillations, the oscillations can be enhanced, and this in turn makes it more likely that the rotor will fail to keep in step with the advancing field.

The second phenomenon occurs because at certain stepping rates it is possible for the complete motor/drive system to exhibit positive feedback, and become unstable. This instability usually occurs at relatively high stepping rates, well above the 'resonance' regions discussed above. The resulting dips in the torque–speed curve are extremely sensitive to the degree of viscous damping present (mainly in the bearings), and it is not uncommon to find that a severe dip which is apparent on a warm day (such as that shown at around 1000 steps/s in Fig. 10.20) will disappear on a cold day.

The dips are most pronounced during steady-state operation, and it may be that their presence is not serious provided that continuous operation at the relevant speeds is not required. In this case, it is often possible to accelerate through the dips without adverse effect. Various special drive techniques exist for eliminating resonances by smoothing out the step-wise nature of the stator field, or by modulating the supply frequency to damp out the instability, but the simplest remedy in open-loop operation is to fit a damper to the motor shaft. Dampers of the Lanchester type or of the viscously-coupled inertia (VCID) type are used. These consist of a lightweight housing which is fixed rigidly to the motor shaft, and an inertia which can rotate relative to the housing. The inertia and the housing are separated either by a viscous fluid (VCID type) or by a friction disc (Lanchester type). Whenever the motor speed is changing, the assembly exerts a damping torque, but once the motor speed is steady, there is no drag torque from the damper. By selecting the appropriate damper, the dips in the torque speed curve can be eliminated, as shown in Fig. 10.20(b). Dampers are also often essential to damp the single-step response, particularly with VR motors, many of which have a highly oscillatory step response. Their only real drawback is that they increase the effective inertia of the system, and thus reduces the maximum acceleration.

10.7 Transient performance

10.7.1 Step response

It was pointed out earlier that the single-step response is similar to that of a damped second-order system. We can easily estimate the natural frequency ω_n in rad/s from the equation

$$\omega_n^2 = \frac{slope\ of\ torque - angle\ curve}{total\ inertia}$$

Knowing ω_n, we can judge what the oscillatory part of the response will look like, by assuming the system is undamped. To refine the estimate, and to obtain the settling time, however, we need to estimate the damping ratio, which is much more difficult to determine as it depends on the type of drive circuit and mode of operation as well as on the mechanical friction. In VR motors the damping ratio can be as low as 0.1, but in hybrid types it is typically 0.3–0.4. These values are too low for many applications where rapid settling is called for.

Two remedies are available, the simplest being to fit a mechanical damper of the type mentioned above. Alternatively, a special sequence of timed command pulses can be used to brake the rotor so that it reaches its new step position with zero velocity and does not overshoot. This procedure is variously referred to as 'electronic damping', 'electronic braking' or 'back phasing'. It involves re-energising the previous phase for a precise period before the rotor has reached the next step position, in order to exert just the right degree of braking. It can only be used successfully when the load torque and inertia are predictable and not subject to change. Because it is an open-loop scheme it is extremely sensitive to apparently minor changes such as day-to-day variation in friction, which can make it unworkable in many instances.

10.7.2 Starting from rest

The rate at which the motor can be started from rest without losing steps is known as the 'starting' or 'pull-in' rate. The starting rate for a given motor depends on the type of drive, and the parameters of the load. This is entirely as expected since the starting rate is a measure of the motor's ability to accelerate its rotor and load and pull into synchronism with the field. The starting rate thus reduces if either the load torque, or the load inertia are increased. Typical pull-in torque–speed curves, for various inertias, are shown in Fig. 10.21. The pull-out torque speed curve is also shown, and it can be seen that for a given load torque, the maximum steady (slewing) speed at which the motor can run is much higher than the corresponding starting rate. (Note that only one pull-out torque is usually shown, and is taken to apply for all inertia values. This is because the inertia is not significant when the speed is constant.)

It will normally be necessary to consult the manufacturer's data to obtain the pull-in rate, which will apply only to a particular drive. However, a rough assessment is easily made: we simply assume that the motor is producing its

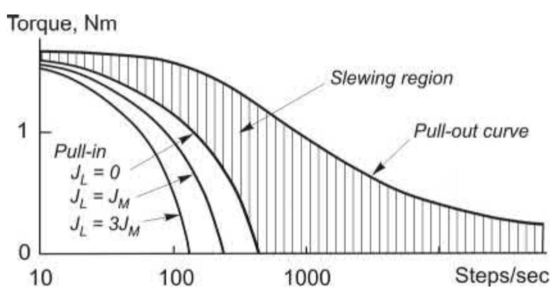

FIG. 10.21 Typical pull-in and pull-out curves showing the effect of load inertia on the pull-in torque. J_M, motor inertia; J_L, load inertia.

pull-out torque, and calculate the acceleration that this would produce, making due allowance for the load torque and inertia. If, with the acceleration as calculated, the motor is able to reach the steady speed in one step or less, it will be able to pull in; if not, a lower pull-in rate is indicated.

10.7.3 Optimum acceleration and closed-loop control

There are some applications where the maximum possible accelerations and decelerations are demanded, in order to minimise point-to-point times. If the load parameters are stable and well-defined, an open-loop approach is feasible, and this is discussed first. Where the load is unpredictable, however, a closed-loop strategy is essential, and this is dealt with later.

To achieve maximum possible acceleration calls for every step command pulse to be delivered at precisely optimised intervals during the acceleration period. For maximum torque, each phase must be on whenever it can produce positive torque, and off when its torque would be negative. Since the torque depends on the rotor position, the optimum switching times have to be calculated from a full dynamic analysis. This can usually be accomplished by making use of the static torque–angle curves (provided appropriate allowance is made for the rise and fall times of the stator currents), together with the torque–speed characteristic and inertia of the load. A series of computations is required to predict the rotor angle-time relationship, from which the switchover points from one phase to the next are deduced. The train of accelerating pulses is then pre-programmed into the controller, for subsequent feeding to the drive in an open-loop fashion. It is obvious that this approach is only practicable if the load parameters do not vary, since any change will invalidate the computed optimum stepping intervals.

When the load is unpredictable, a much more satisfactory arrangement is obtained by employing a closed-loop scheme, using position feedback from a shaft-mounted encoder. The feedback signals indicate the instantaneous position of the rotor, and are used to ensure that the phase-windings are switched at precisely the right rotor position for maximising the developed torque. Motion is initiated by a single command pulse, and subsequent step-command pulses are effectively self-generated by the encoder. The motor continues to accelerate until its load torque equals the load torque, and then runs at this (maximum) speed until the deceleration sequence is initiated. During all this time, the step counter continues to record the number of steps taken. This approach is essentially the same as we discussed in relation to self synchronous motors in Chapter 9.

Closed-loop operation ensures that the optimum acceleration is achieved, but at the expense of more complex control circuitry, and the need to fit a shaft encoder. Relatively cheap encoders are however now available for direct fitting to some ranges of motors, and single chip microcontrollers are available which provide all the necessary facilities for closed-loop control.

An appealing approach aimed at eliminating an encoder is to detect the position of the rotor by on-line analysis of the signals (principally the rates of change of currents) in the motor windings themselves: in other words, to use the motor as its own encoder. A variety of approaches have been tried, including the addition of high-frequency alternating voltages superimposed on the excited phase, so that as the rotor moves the variation of inductance results in a modulation of the alternating current component. Some success has been achieved with particular motors, but the approach has not achieved widespread commercial exploitation, perhaps because of competition from the PM synchronous drives discussed in Chapter 9.

To return finally to encoders, we should note that they are also used in open-loop schemes when an absolute check on the number of steps taken is required. In this context the encoder simply provides a tally of the total steps taken, and normally plays no part in the generation of the step pulses. At some stage, however, the actual number of steps taken will be compared with the number of step command pulses issued by the controller. When a disparity is detected, indicating a loss (or gain of steps), the appropriate additional forward or backward pulses can be added.

10.8 Switched reluctance motor drives

The switched reluctance drive was developed in the 1980s with the promise of offering advantages in terms of efficiency, power per unit weight and volume, robustness and operational flexibility compared with the then dominant induction motor. Its market position has been compromised by the rise of PM motor technologies and improved control strategies for induction motors. Nevertheless, the switched reluctance motor has been applied to a range of applications which can benefit from its performance characteristics, notably offering very high torque at standstill and very low speed. The motor (Fig. 10.22) and its associated power-electronic drive must be designed as an integrated package, and optimised for a particular specification, e.g. for maximum overall efficiency with a specific load, or maximum speed range, peak short-term torque or torque ripple (and associated acoustic noise).

10.8.1 Principle of operation

Like the stepping motor, the switched reluctance (SR) motor (Fig. 10.23) is 'doubly-salient', i.e. it has projecting poles on both rotor and stator. However, most SR motors are of much higher power than the largest stepper, and it turns out that in the higher power ranges (where the winding resistances become much less significant), the doubly-salient arrangement is very effective in terms of efficient electromagnetic energy conversion.

A cross section through a typical SR motor is shown in Fig. 10.23: this example has twelve stator poles and eight rotor poles, and represents a widely

FIG. 10.22 Switched reluctance motor. The toothed rotor does not require windings or magnets, and is therefore exceptionally robust. *(Courtesy of Nidec SR Drives Ltd.)*

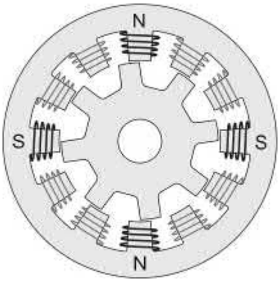

FIG. 10.23 Typical switched reluctance motor. Each of the twelve stator poles carries a concentrated winding, while the eight-pole rotor has no windings or magnets.

used arrangement, but other pole combinations are used to suit different applications. The stator carries coils on each pole, while the rotor, which is made from laminations in the usual way, has no windings or magnets.

In Fig. 10.23 the twelve coils are grouped to form three phases, which are independently energised from a three-phase converter.

The motor rotates by exciting the phases sequentially in the sequence A, B, C for anticlockwise rotation or A, C, B for clockwise rotation, the 'nearest' pair of rotor poles being pulled into alignment with the appropriate stator poles by reluctance torque action. In Fig. 10.23 the four coils forming phase A are shown in black, the polarities of the coil m.m.f.'s being indicated by the letters N and S on the back of core. Each time a new phase is excited the equilibrium position of the rotor advances by 15°, so after one complete cycle (i.e. each of the three phases has been excited once) the angle turned through is 45°. The machine therefore rotates once for eight fundamental cycles of supply to the stator windings, so in terms of the relationship between the fundamental supply frequency and the speed of rotation, the machine in Fig. 10.23 behaves as a 16-pole conventional machine.

The structure is clearly the same as the variable reluctance stepping motor discussed earlier in this chapter, but there are important design differences which reflect the different objectives (continuous rotation for the SR, stepwise progression for the stepper), but otherwise the mechanisms of torque production are identical. However whilst the stepper is designed first and foremost for open-loop operation, the SR motor is designed for self-synchronous operation, the phases being switched by signals derived from a shaft-mounted rotor position detector. In terms of performance, at all speeds below the base speed continuous operation at full torque is possible. Above the base speed, the flux can no longer be maintained at full amplitude and the available torque reduces with speed. The operating characteristics are thus very similar to those of the other most important controlled-speed drives.

10.8.2 Torque prediction and control

If the iron in the magnetic circuit is treated as ideal, analytical expressions can be derived to express the torque of a reluctance motor in terms of the rotor position and the current in the windings. In practice, however, this analysis is of little real use, not only because switched reluctance motors are designed to operate with high levels of magnetic saturation in parts of the magnetic circuit, but also because, except at low speeds, it is not easy to achieve specified current profiles.

The fact that high levels of saturation are involved makes the problem of predicting torque at the design stage challenging, but despite the highly nonlinear relationships it is possible to compute the flux, current and torque as functions of rotor position, so that optimum control strategies can be devised to meet particular performance specifications. Unfortunately this complexity means that there is no simple equivalent circuit available to illuminate behaviour.

As we saw when we discussed the stepping motor, to maximise the average torque it would (in principle) be desirable to establish the full current in each phase instantaneously, and to remove it instantaneously at the end of each positive torque period. But, as illustrated in Fig. 10.15, this is not possible even with a small stepping motor, and certainly out of the question for switched reluctance motors (which have much higher inductance) except at low speeds where current-chopping (see Fig. 10.17) is employed. For most of the speed range, the best that can be done is to apply the full voltage available from the converter at the start of the 'on' period, and (using a circuit such as that shown in Fig. 10.17) apply full negative voltage at the end of the pulse by opening both of the switches.

Operation using full positive voltage at the beginning and full negative voltage at the end of the 'on' period is referred to as 'single-pulse' operation. For all but small motors (of less than say 1 kW) the phase resistance is negligible and consequently the magnitude of the phase flux-linkage is determined by the applied voltage and frequency, as we have seen many times previously with other types of motor.

The relationship between the flux-linkage (ψ) and the voltage is embodied in Faraday's law, i.e. $v = \frac{d\psi}{dt}$, so with the rectangular voltage waveform of single-pulse operation the phase flux linkage waveforms have a very simple triangular shape, as in Fig. 10.24 which shows the waveforms for phase A of a 3-phase motor. (The waveforms for phases B are C are identical, but are not shown: they are displaced by one third and two thirds of a cycle, as indicated by the arrows.) The upper half of the diagram represents the situation at speed N, while the lower half corresponds to a speed of 2N. As can be seen, at the higher speed (high frequency) the 'on' period halves, so the amplitude of the flux halves, leading to a reduction in available torque. The same limitation was seen in the case of the inverter-fed induction motor drive, the only difference being that the waveforms in that case were sinusoidal rather than triangular.

It is important to note that these flux waveforms do not depend on the rotor position, but the corresponding current waveforms do because the m.m.f.

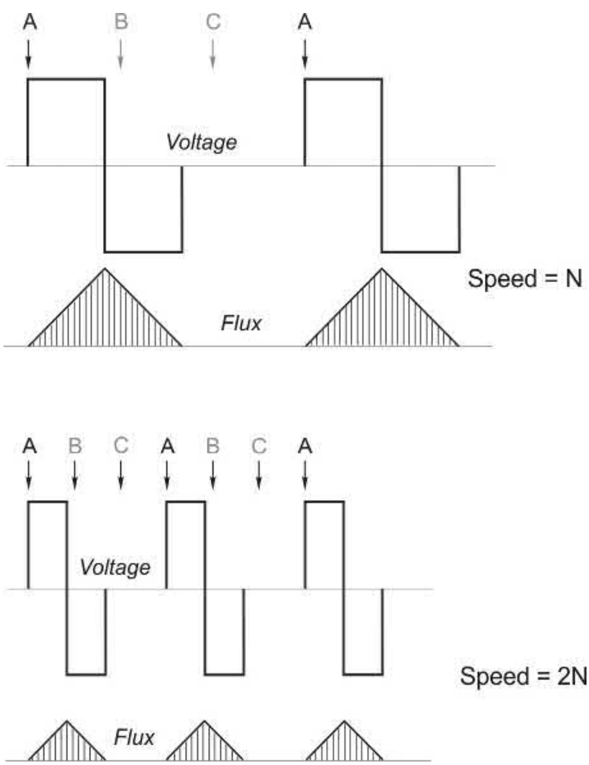

FIG. 10.24 Voltage and flux waveforms for switched-reluctance motor in 'single-pulse' mode.

needed for a given flux depends on the effective reluctance of the magnetic circuit, and this of course varies with the position of the rotor.

To get the most motoring torque for any given phase flux waveform, it is obvious that the rise and fall of the flux must be timed to coincide with the rotor position: ideally, the flux should only be present when it produces positive torque, and be zero whenever it would produce negative torque, but given the delay in build-up of the flux it may be better to switch on early so that the flux reaches a decent level at the point when it can produce the most torque, even if this does lead to some negative torque at the start and finish of the cycle.

The job of the torque control system is to switch each phase on and off at the optimum rotor position in relation to the torque being demanded at the time, and this is done by keeping track of the rotor position, possibly using a rotor position sensor.[2] Just what angles constitute the optimum depends on what is to be optimised (e.g. average torque, overall efficiency), and this in turn is decided by reference to the data stored digitally in the controller 'memory map' that relates current, flux, rotor position and torque for the particular machine. Torque control is thus considerably less straightforward than in d.c. drives, where torque is directly proportional to armature current, or induction motor drives where torque is proportional to slip.

10.8.3 Power converter and overall drive characteristics

An important difference between the SR motor and all other self-synchronous motors is that its full torque capability can be achieved without having to provide for both positive and negative currents in the phases. This is because the torque does not depend on the direction of current in the phase winding. The advantage of such 'unipolar' drives is that because each of the main switching devices is permanently connected in series with one of the motor windings (as in Fig. 10.17), there is no possibility of the 'shoot-through' fault (see Chapter 2) which necessitates the inclusion of 'dead time' in the switching strategies of the conventional inverter.

Overall closed-loop speed control is obtained in the conventional way with the speed error acting as a torque demand to the torque control system described above. However, if a position sensor is fitted the speed can be derived from the position signal.

In common with other self-synchronous drives, a wide range of operating characteristics are available. If the input converter is fully controlled, continuous regeneration and full 4-quadrant operation is possible, and the usual constant torque, constant power and series type characteristic is regarded as standard. Low speed torque can be uneven unless special measures are taken

2. As in all drives the use of a position sensor is avoided wherever possible because of the significant extra cost and complexity. Sophisticated software models of the motor are therefore employed to deduce rotor position from analysis of the current and voltage waveforms.

to profile the current pulses, but the particular merit of the SR drive is that continuous low speed high torque operation is usually better than for most competing systems in terms of overall efficiency.

10.9 Review questions

(1) Why are the step positions likely to be less well-defined when a motor is operated in '2-phase-on' mode as compared with 1-phase-on mode?

(2) What is meant by detent torque, and in what type of motors does detent torque occur?

(3) What is meant by the 'holding torque' of a stepping motor?

(4) The static torque curve of a 3-phase VR stepper is approximately sinusoidal, the peak torque at rated current being 0.8 Nm. Find the step position error when a steady load torque of 0.25 Nm is present.

(5) For the motor in question (4), estimate the low-speed pull-out torque when the motor is driven by a constant-current drive.

(6) The static torque–angle curve of a particular 1.8° hybrid step motor can be approximated about the equilibrium position by a straight line with a slope of 2 Nm/degree, and the total inertia (motor plus load) is 1.8×10^{-3} kg m^2. Estimate the frequency of oscillation of the rotor following a single step.

(7) Find the step angle of the following stepping motors:
   (a) 3-phase, VR, 12 stator teeth and 8 rotor teeth;
   (b) 4-phase unipolar, hybrid, 50 rotor teeth

(8) What simple tests could be done on an unmarked stepping motor to decide whether it was a VR motor or a hybrid motor?

(9) At what speed would a 1.8° hybrid steeping motor run if its two phases were each supplied from the 60 Hz utility supply, the current in one of the phases being phase-shifted by 90° with respect to the other?

(10) An experimental scientist read that stepping motors typically complete each single step in a few milliseconds. He decided to use one for a display aid, so he mounted a size 18 (approximately 4 cm diameter) 15°/step VR motor so that its shaft was vertical, and fixed a lightweight (30 g) aluminium pointer about 40 cm long onto the shaft. When he operated the motor he was very disappointed to discover that after every step the pointer oscillated wildly and took almost two seconds before coming to rest. Use some judicious estimates to suggest why he should not have been surprised.

Answers to the review questions are given in the Appendix.

Chapter 11

Motor/drive selection

11.1 Introduction

The selection process often highlights difficulties in three areas. Firstly, as we have discovered in the preceding chapters, there is a good deal of overlap between the characteristics of the major types of motor and drive. This makes it impossible to lay down a set of hard and fast rules to guide the user straight to the best solution for a particular application. Secondly, users tend to underestimate the importance of starting with a comprehensive specification of what they really want, and they seldom realise how much weight attaches to such things as the steady-state torque-speed curve, the dynamic performance required, the inertia of the load, the pattern of operation (continuous or intermittent), the environment and the question of whether or not the drive needs to be capable of regeneration. And thirdly they may be unaware of the existence of standards and legislation, and hence can be baffled by questions from any potential supplier.

The aim in this chapter is to assist the user in taking the first steps by giving some of these matters an airing. We begin by laying out the availability, ratings and speeds, and then the main characteristics of the various motor and drive types that we have highlighted earlier in the book. We then move on to the questions which need to be asked about the load and pattern of operation, and finally look briefly at the matter of standards. The whole business of selection is so broad that it really warrants a book to itself, but the cursory treatment here should at least help the user to specify the drive rating and arrive at a short-list of possibilities.

11.2 Power ratings and capabilities

The four tables in this section represent an attempt to condense the most important aspects that will be of interest to the user into a readily accessible form. It will be evident from the weight given to a.c. drives in the book that they now dominate the scene, and this is reflected in the fact that only the first table deals with d.c. drives. The scope of a.c. drives is too great for a single table, so to aid digestion we have split them into three.

Electric Motors and Drives. https://doi.org/10.1016/B978-0-08-102615-1.00011-8

DC Motor (Separately Excited or Permanent Magnet)			
Drive Type	**Single Quadrant**	**Four Quadrant**	**DC Chopper**
Converter	Single or three-phase fully (or half) controlled thyristor bridge.	Dual single/three phase fully controlled thyristor bridge.	DC Chopper
Torque/Speed range	Motoring in one direction only (Braking in other direction.)	Motoring and Braking in both directions.	2 or 4 quadrant versions.
Speed Control	Closed loop control of armature voltage with inner current control loop.		
Torque Control	Closed loop control of armature current.		
Ratings	10 W to 5 MW (Fractional HP to 7000 HP)		0.5 to 5 kW (0.7 to 7 HP). Traction > 500 kW.
Max Power	Available to multi-MW ratings, but motor limitations restrict the product of Power and Speed to 3×10^6 kW.rev/min.		
Min Speed	Good control down to standstill.		
Notable features	Separately-excited motors often used above base speed in constant-power mode.		DC to DC conversion.
		Fast torque reversals	Smooth torque possible.
Market Position	DC drives remain a significant part of the overall drives market. Popularity is diminishing and is focused mainly on applications where a DC motor already exists; on very simple drives; and high power drives where the DC motor can still be competitive.		
	Popular low-cost solution for low power drives in simple applications or for retrofit to existing motors.	Popular 4 quadrant solution for retrofit to existing motors. Motor and drive can be competitive at higher powers.	Once very popular as servo drive with brushed motor. Gradually losing markets to AC equivalents. Remains important for some retro-fits/upgrades in traction applications.

FIG. 11.1 Conventional (brushed) d.c. motor drives.

The first table (Fig. 11.1) covers conventional (brushed) d.c. motor drives, which remain significant despite their diminishing market share. However, the second table (Fig. 11.2—a.c. drives) is by far the most important because it includes the inverter-fed induction motor and brushless PM motor drives which are now the most widely used industrial drives. The (synchronous) reluctance motor is not separately identified here but can be considered as an alternative to the induction motor in most applications (though the motor tends to be a little larger). The emergence of the salient PM synchronous machine in recent years is having a growing impact in many markets including automotive traction: its operational characteristics closely follow those of the nonsalient PM motor and for this reason they have been discussed together in this summary. Fig. 11.3 deals with (predominantly large) a.c. drives, which have limited and/or niche application and are therefore less common: for completeness the static Scherbius drive is included although it was not discussed. The fourth table (Fig. 11.4) contains three unlikely bedfellows whose conjunction reflects the fact that they could not be fitted in elsewhere.

	Induction Motor			Brushless PM Motor (non-salient or salient)		Excited Rotor Synch. Motor
Drive /Converter	PWM Inverter	Multi-level PWM	Current source converter	PWM Inverter	Multi-level PWM	3-phase fully-controlled bridges
Torque/Speed Quadrants	Motoring and Braking in both directions with either dual converter on supply side, or d.c. link chopper and brake resistor.		Motoring and Braking in both directions	Motoring and Braking in both directions with either dual converter on supply side, or d.c. link chopper and brake resistor.		Motoring and Braking in both directions
Method of Torque/Speed Control	Field Orientation (Vector) (or Direct Torque Control). Simpler Scalar (V/f) variants available.			Field Orientation (Vector) (or Direct Torque Control). Position feedback is standard, but sensorless options are available.		Same as brushless PM
Typical Power Ratings	Up to 3 MW (400/690 V) >10 MW (6600 V)		4 MW (400/690 V) 10 MW (6600 V)	10 W to 500 kW, with higher powers available as custom designs.		2 to > 20 MW (MV only).
Maximum Speed (Typical)	> 40,000 rev/min		> 6000 rev/min	> 70,000 rev/min	>10,000 rev/min	>10,000 rev/min
Minimum Speed (Typical)	Control dependent. With position feedback to standstill with full torque: without, to 1 Hz.		5Hz standard: lower possible.	To standstill with full torque. (Without feedback to 1Hz.)		Open loop to 3 Hz; closed loop to standstill with full torque.
Notable Features	Simple, robust, readily available induction motor.			Very high power density (somewhat lower for salient motor). Excellent control dynamics. Motor inertia can be specified from many suppliers. Motors with integral position sensors available.		Simple converter available at very high voltages High speed motors available.
	Low torque ripple. Excellent control dynamics.		High efficiency			
		Low motor dV/dt (important at MV).			Low motor dV/dt (important at MV).	High efficiency. Good dynamic control.
Market Position	Voltage source PWM inverters are popular because they are readily applied to different types of motor and load.		Once the topology of choice for single motor drives, but has largely been superseded by the PWM inverter.	Most popular high performance/servo drive for precision motion applications where smooth rotation or rapid accln/deceln is required. Salient motors are growing in high volume markets in response to concerns about the supply chain of rare earth magnets.	Limited application. Most are below 700V where multi-level advantages are less clear.	The most popular solution for very high power (> 5 MW) and /or very high speed applications.
	The industrial workhorse. Most widely used drive.	Popular in MV drives.				

FIG. 11.2 Induction, brushless PM and excited rotor synchronous motor drives.

11.3 Drive characteristics

The successful integration of electronic variable speed drives into a system depends on knowledge of the key characteristics of the application and the site where the system will be used. The correct drive also has to be selected, but this is often simpler than it used to be because the voltage source PWM inverter has become universally adopted for almost all applications, and its control characteristics can be adapted to suit most loads with relative ease. It is still important to understand the general characteristics of the various types of drive, however, and these are listed in summary form in the table (Fig. 11.5). (Note that the data on overloads relates to typical industrial equipment, but users will inevitably encounter a range of different figures from the various suppliers.) Finally, a summary of the pros and cons of d.c. and a.c. drives is given in the table

	Static Kramer	Static Scherbius	Cycloconverter	Matrix Converter
Motor	Slipring Induction		Synchronous or Induction	
Converter	Diode bridge and 3-phase full controlled bridge	Cycloconverter or back to back PWM inverters	Cycloconverter	Matrix converter
Torque/Speed Quadrants	Motoring in one direction only	Motoring and braking in both directions		
Speed and Torque Control	Field Orientation control (or Direct Torque Control) Simple scalar V/f control variants available.			
Ratings	500 kW to 20 MW	500 kW to 20 MW	500 kW to > 10 MW	50 kW to 2 MW
Typical Maximum Speed	1470 rev/min	2000 rev/min	< 30 Hz	< supply frequency
Min Speed	900 rev/min (wider speed range impacts competitiveness)		0 Hz	
Features	Economic solution at high powers		No DC link components.	
Market Position	Not prominent but still very economic for limited speed range applications.	Has much of the benefit of the Kramer but with ability to operate above synchronous speed.	The topology of choice for very low speed, very high power applications	Has niche areas of application where DC link components troublesome. New topology variants may expand market.

FIG. 11.3 Slip energy recovery and direct converter drives.

	Soft Starter	Switched Reluctance Drive	Stepper Motor
Motor	Induction Motor	Switched Reluctance Motor	Stepper Motor
Converter Type	Inverse parallel thyristors in supply lines	Specific topologies	Specific topologies
Torque/Speed Quadrants	Motoring in one direction only	Motoring and Braking in both directions with either a supply-side dual converter or a d.c. link chopper and brake resistor	
Speed Control	Starting duty only	Specific Control Strategies	
Typical Industrial Power Ratings	1 kW to > 1 MW	1 kW to > 1 MW	10 W to > 5 kW
Typical Maximum Speed	Rated DOL speed	> 10,000 rev/min	> 10,000 rev/min
Typical Minimum Speed	Not Applicable	Standstill	Standstill
Notable Features		Very simple and robust motor	Simple motor and control strategy
Market Position	The soft start has very specific application in reducing the starting current and/or controlling the starting torque of a DOL induction motor	High starting torque applications benefit from SR technology	Open loop positioning systems

FIG. 11.4 Soft start, Switched Reluctance and stepping motor drives.

	DC Drives (Separately excited)		AC Drives		
	Phase Controlled	Chopper	Induction Motor (Field Orientation Control without speed feedback)	Induction Motor (Field Orientation Control with speed feedback)	Brushless PM Motor (non-salient and salient). (Field Orientation Control with position feedback)
Operating Speed Range	0 to Base Speed at Constant Torque. Above base speed at Constant Power. Max approx. 4 × base speed.		3-100% Base speed at constant torque. Above base speed at constant power. Max approx. 20 × base.	0-100% Base speed at constant torque. Above base speed at constant power. Max approx. 20 × base.	0-100% Base speed at constant torque. Above base speed at constant power. Max approx. 8 × base.
Braking Capability	150% (4Q drive)	150% (4Q drive)	150%		> 200%
Speed Loop Response	10 Hz	50 Hz	50 Hz	150 Hz	> 150 Hz
Speed Holding (100% load change)	0.1%–0.05% with good speed feedback	0.05%–0.001% with good speed feedback	0.1%	0.001%	0.0005%
Torque / Speed Capability	Constant Torque + Field Weakening	Constant Torque + Field Weakening	Constant Torque + Field Weakening	Constant Torque + Field Weakening	Constant Torque
Starting Torque	150% (60s)	150% (60s)	150% (60s)	150% (60s)	200% (4s)
Min speed with 100% torque	Standstill	Standstill	2% base speed	Standstill	Standstill
Motor IP Rating	IP23	IP54 / IP23	IP54	IP54	IP65
Motor Inertia	High	High (Low avail.)	Medium (Low avail.)	Medium (Low avail.)	Low (High avail.)
Motor size	Large	Large	Medium	Medium	Small
Cooling	Forced ventilated	FV or natural	Natural Cooling	Natural Cooling	Natural Cooling
Typical f/b device	Tachogenerator or Encoder	Tachogenerator or Encoder	Not Applicable	Encoder	Encoder

FIG. 11.5 Characteristics of the most popular types of drive.

(Fig. 11.6). The wealth of information in these tables should be of great value in assisting prospective users to narrow down their options.

11.3.1 Maximum speed and speed range

This is an important consideration in many drives, so a few words are appropriate. We saw in Chapter 1 that as a general rule, for a given power the higher the base speed the smaller the motor. In practice there are only a few applications where motors with base speeds below a few hundred rev/min are attractive, and it is usually best to obtain low 'maximum' speeds by means of mechanical speed reduction often in the form of a gearbox.

Motors speeds above 10,000 rev/min are also unusual except in small universal motors and special-purpose inverter-fed motors for applications such as aluminium machining. The majority of motors have base speeds between 1500 and 3000 rev/min: this range is attractive as far as motor design is concerned since acceptable power/weight ratios are obtained, and these speeds are also

	DC Drives (Separately excited)		A C Drives		
	Phase Controlled	Chopper	Induction Motor (Field Orientation Control without speed feedback)	Induction Motor (Field Orientation Control with speed feedback)	Brushless PM Motor (non-salient and salient). (Field Orientation Control with position feedback)
Principal Advantages	• Low Cost Controller • Relatively simple technology	• Good Dynamic performance • Relatively simple technology	• Good dynamic performance • Full torque down to very low speeds • Good starting torque	• Very good dynamic performance • Full torque available down to standstill • Standard motor (but with feedback added) • No zero torque dead band	• Excellent dynamic performance • Low (or High) inertia motors • High Motor IP ratings • Very smooth rotation possible
Principal Disadvantages	• Expensive motor <100 kW • Brush gear Maintenance • Note: Low loads reduce brush life • Zero Torque dead band • Possible failure on supply loss/dip • Possible instability on Fan/Pump type loads • Low Motor IP rating • Chopper at modest powers only (<10 kW) • Additional converter required for regeneration		• Additional converter required for regeneration to the supply • Limited Torque and Speed Loop response	• Additional converter required for regeneration to the supply	• Additional converter required for regeneration to the supply • Field weakening range difficult to facilitate

FIG. 11.6 Advantages and disadvantages of d.c. and a.c. drives.

satisfactory as far as mechanical transmission is concerned. Note, however, that as explained in Chapter 1, the higher the rated speed the more compact the motor, so there will be instances where a high speed motor is preferred when minimising motor size is of paramount importance.

In controlled-speed applications, the range over which the steady-state speed must be controlled, the required torque/speed characteristic, the duty cycle and the accuracy of the speed-holding, are significant factors in the selection process. In Chapter 7, we looked at some of these aspects, notably in relation to motor cooling. For constant torque loads which require operation at all speeds, the inverter-fed induction motor, the inverter-fed synchronous machine and the d.c. drive are possibilities, but only the d.c. drive would come as standard with a force-ventilated motor capable of continuous operation with full torque at very low speeds.

Fan or Pump type loads (see below), with a wide operating speed range, are a somewhat easier proposition because the torque is low at low speeds. In most circumstances, the inverter-fed induction motor (using a standard motor) is the natural choice.

Figures for accuracy of speed holding can sometimes cause confusion, as they are usually given as a percentage of the base speed. Hence with a drive claiming a decent speed holding accuracy of 0.2% and a base speed of 2000 rev/min, the user can expect the actual speed to be between 1996 rev/min and 2004 rev/min when the speed reference is 2000 rev/min. But if the speed reference is set for 100 rev/min, the actual speed can be anywhere between 96 rev/min and 104 rev/min, and still be within the specification.

11.4 Load requirements—torque-speed characteristics

The most important things we need to know about the load are the steady-state torque-speed characteristic, the effective inertia as seen by the motor and what dynamic performance is required. At one extreme, for example in a steel rolling mill, it may be necessary for the speed to be set at any value over a wide range, and for the mill to react very quickly when a new target speed is demanded. Having reached the set speed, it may be essential that it is held very precisely even when subjected to sudden load changes. At the other extreme, for example a large ventilating fan, the range of set speed may be quite limited (perhaps from 80% to 100%); it may not be important to hold the set speed very precisely; and the times taken to change speeds, or to run-up from rest, are unlikely to be critical.

At full speed both of these examples may demand the same power, and at first sight might therefore be satisfied by the same drive system. But the ventilating fan is obviously an easier proposition, and it would be overkill to use the same system for both. The rolling mill would call for a regenerative drive with speed or position feedback, while the fan could quite happily manage with a simple open-loop, i.e. sensorless inverter-fed induction motor drive.

Although loads can vary enormously, it is customary to classify them into two major categories, referred to as 'constant-torque' or 'fan or pump' types. We will use the example of a constant-torque load to illustrate in detail what needs to be done to arrive at a specification for the torque-speed curve. An extensive treatment is warranted because this is often the stage at which users come unstuck.

11.4.1 Constant-torque load

A constant torque load implies that the torque required to keep the load running is the same at all speeds. A good example is a drum-type hoist, where the torque required varies with the load on the hook, but not with the speed of hoisting. An example is shown in Fig. 11.7.

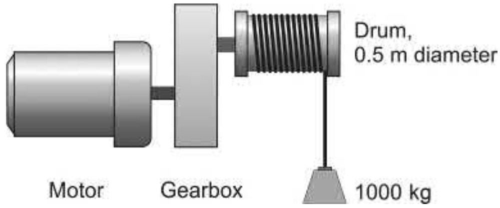

FIG. 11.7 Motor driven hoist—a constant-torque load.

The drum diameter is 0.5 m, so if the maximum load (including the cable) is say 1000 kg, the tension in the cable (mg) will be 9810 N, and the torque applied by the load at the drum will be given by $Force \times radius = 9810 \times 0.25 \approx 2500$ Nm. When the speed is constant (i.e. the load is not accelerating), the torque provided by the motor at the drum must be equal and opposite to that exerted at the drum by the load. (The word 'opposite' in the last sentence is often omitted, it being understood that steady-state motor and load torque must necessarily act in opposition.)

Suppose that the hoisting speed is to be controllable at any value up to a maximum of 0.5 m/s, and that we want this to correspond with a maximum motor speed of around 1500 rev/min, which is a reasonable speed for a wide range of motors. A hoisting speed of 0.5 m/s corresponds to a drum speed of 19 rev/min, so a suitable gear ratio would be say 80:1, giving a maximum motor speed of 1520 rev/min.

The load torque, as seen at the motor side of the gearbox, will be reduced by a factor of 80, from 2500 Nm to 31 Nm at the motor. We must also allow for friction in the gearbox, equivalent to perhaps 20% of the full load torque, so the maximum motor torque required for hoisting will be 37 Nm, and this torque must be available at all speeds up to the maximum of 1520 rev/min.

We can now draw the steady-state torque-speed curve of the load as seen by the motor, as shown in Fig. 11.8.

The steady-state motor power is obtained from the product of torque (Nm) and angular velocity (rad/s). The maximum continuous motor power for hoisting is therefore given by

$$P_{max} = 37 \times 1520 \times \frac{2\pi}{60} = 5.9\,\text{kW} \tag{11.1}$$

At this stage it is always a good idea to check that we would obtain roughly the same answer for the power by considering the work done per second at the

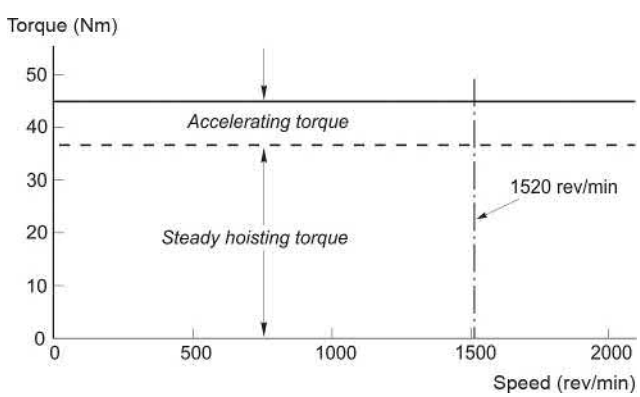

FIG. 11.8 Torque requirements for motor in hoist application (Fig. 11.7).

load. The force (F) on the load is 9810 N, the velocity (v) is 0.5 m/s so the power (Fv) is 4.9 kW. This is 20% less than we obtained above, because here we have ignored the power lost in the gearbox.

So far we have established that we need a motor capable of continuously delivering 5.9 kW at 1520 rev/min in order to lift the heaviest load at the maximum required speed. However we have not yet addressed the question of how the load is accelerated from rest and brought up to the maximum speed. During the acceleration phase the motor must produce a torque greater than the load torque, or else the load will descend as soon as the brake is lifted. The greater the difference between the motor torque and the load torque, the higher the acceleration. Suppose we want the heaviest load to reach full speed from rest in say 1 s, and suppose we decide that the acceleration is to be constant. We can calculate the required accelerating torque from the equation of motion, i.e.

$$\textit{Torque}\,(\text{Nm}) = \textit{Inertia}\,\left(\text{kg}\,\text{m}^2\right) \times \textit{Angular acceleration}\,\left(\frac{\text{rad}}{\text{sec}^2}\right). \qquad (11.2)$$

We usually find it best to work in terms of the variables as seen by the motor, and therefore we first need to find the effective total inertia as seen at the motor shaft, then calculate the motor acceleration, and finally use Eq. (11.2) to obtain the accelerating torque.

The effective inertia consists of the inertia of the motor itself, the referred inertia of the drum and gearbox, and the referred inertia of the load on the hook. The term 'referred inertia' means the apparent inertia, viewed from the motor side of the gearbox. If the gearbox has a ratio of n:1 (where n is greater than 1), an inertia of J on the low-speed side appears to be an inertia of J/n^2 at the high-speed side. In this example the load actually moves in a straight line, so we need to ask what the effective inertia of the load is, as 'seen' at the drum. The geometry here is simple, and it is not difficult to see that as far as the inertia seen by the drum is concerned the load appears to be fixed to the surface of the drum. The load inertia at the drum is then obtained by using the formula for the inertia of a mass m located at radius r, i.e. $J = mr^2$, yielding the effective load inertia at the drum as $1000\,\text{kg} \times (0.25\,\text{m})^2 = 62.5$ kgm^2.

The effective inertia of the load as seen by the motor is $1/(80)^2 \times 62.5 \approx 0.01$ kgm^2. To this must be added firstly the motor inertia which we can obtain by consulting the manufacturer's catalogue for a 5.9 kW, 1520 rev/min motor. This will be straightforward for a d.c. motor, but a.c. motor catalogues tend to give ratings at utility frequencies only, and here a motor with the right torque needs to be selected, and the possible torque speed curve for the type of drive considered. For simplicity let us assume we have found a motor of precisely the required rating which has a rotor inertia of 0.02 kgm^2. The referred inertia of the drum and gearbox, must be added and this again we have to calculate or look up. Suppose this yields a further 0.02 kgm^2. The total effective inertia is thus 0.05 kgm^2, of which 40% is due to the motor itself.

The acceleration is straightforward to obtain, since we know the motor speed is required to rise from zero to 1520 rev/min in 1 s. The angular acceleration is given by the increase in speed divided by the time taken, i.e.

$$\left(1520 \times \frac{2\pi}{60}\right) \div 1 = 160\,\text{rad}/\sec^2$$

We can now calculate the accelerating torque from Eq. (11.2) as

$$T = 0.05 \times 160 = 8\,\text{Nm}$$

Hence in order to meet both the steady-state and dynamic torque requirements, a drive capable of delivering a torque of 45 Nm (= 37 + 8) at all speeds up to 1520 rev/min is required, as indicated in Fig. 11.8.

In the case of a hoist, the anticipated pattern of operation may not be known, but it is likely that the motor will spend most of its time hoisting rather than accelerating. Hence although the peak torque of 45 Nm must be available at all speeds, this will not be a continuous demand, and will probably be within the short-term overload capability of a drive which is continuously rated at 5.9 kW.

We should also consider what happens if it is necessary to lower the fully-loaded hook. We allowed for friction of 20% of the load torque (31 Nm), so during descent we can expect the friction to exert a braking torque equivalent to 6.2 Nm. But in order to prevent the hook from running-away, we will need a total torque of 31 Nm, so to restrain the load, the motor will have to produce a torque of 24.8 Nm. We would naturally refer to this as a braking torque because it is necessary to prevent the load on the hook from running away, but in fact the torque remains in the same direction as when hoisting. The speed is however negative, and in terms of a 'four-quadrant' diagram (e.g. Fig. 3.12) we have moved from quadrant 1 into quadrant 4, and thus the power flow is reversed and the motor is regenerating, the loss of potential energy of the descending load being converted back into electrical form (and losses). Hence if we wish to cater for this situation we must go for a drive that is capable of continuous regeneration: such a drive would also have the facility for operating in quadrant 3 to produce negative torque to drive down the empty hook if its weight was insufficient to lower itself.

In this example the torque is dominated by the steady-state requirement, and the inertia-dependent accelerating torque is comparatively modest. Of course if we had specified that the load was to be accelerated in one fifth of a second rather than 1 s, we would require an accelerating torque of 40 Nm rather than 8 Nm, and as far as torque requirements are concerned the acceleration torque would be more or less the same as the steady-state running torque. In this case it would be necessary to consult the drive manufacturer to determine the drive rating, which would depend on the frequency of the start/stop sequence.

The question of how to rate the motor when the loading is intermittent is explored more fully later, but it is worth noting that if the inertia is appreciable the stored rotational kinetic energy ($\frac{1}{2}J\omega^2$) may become very significant, especially when the drive is required to bring the load to rest. Any stored energy either has to be dissipated in the motor and drive itself, or returned to the supply. All motors are inherently capable of regenerating, so the arrangement whereby the kinetic energy is recovered and dumped as heat in a resistor within the drive enclosure is the cheaper option, but is only practicable when the energy to be absorbed is modest. If the stored kinetic energy is large, the drive must be capable of returning energy to the supply, and this inevitably pushes up the cost of the converter.

In the case of our hoist, the stored kinetic energy is only

$$\frac{1}{2} \times 0.05\left(1520 \times \frac{2\pi}{60}\right)^2 = 633\,\text{Joules}$$

or about 1% of the energy needed to heat up a mug of water for a cup of coffee. Such modest energies could easily be absorbed by a resistor, but given that in this instance we are providing a regenerative drive, this energy would also be returned to the supply.

11.4.2 Inertia matching

There are some applications where the inertia dominates the torque requirement, and the question of selecting the right gearbox ratio has to be addressed. In this context the term 'inertia matching' often causes confusion, so it is worth explaining what it means.

Suppose we have a motor with a given torque, and we want to drive an inertial load via a gearbox. As discussed previously, the gear ratio determines the effective inertia as 'seen' by the motor: a high step-down ratio (i.e. load speed much less than motor speed) leads to a very low referred inertia, and vice-versa.

If the specification calls for the acceleration of the load to be maximised, it turns out that the optimum gear ratio is that which causes the referred inertia of the load to be equal to the inertia of the motor. Applications in which load acceleration is important include all types of positioning drives, e.g. in machine tools and phototypesetting. This explains why brushless PM synchronous motors, discussed in Chapter 9, are available from a wide range of manufacturers with different inertias. (There is an electrical parallel here—to get the most power into a load from a source with internal resistance R, the load resistance must be made equal to R.)

It is important to note however that inertia matching only maximises the *acceleration* of the load. Frequently it turns out that some other aspect of the specification (e.g. the maximum required load speed) cannot be met if the gearing is chosen to satisfy the inertia matching criterion, and it then becomes necessary to accept reduced acceleration of the load in favour of higher speed.

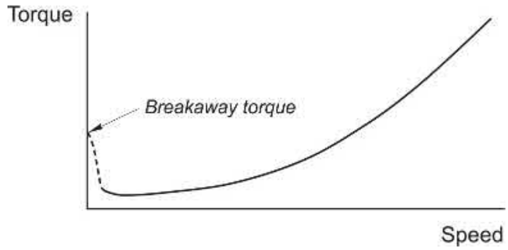

FIG. 11.9 Torque-speed characteristics for fan-type load.

11.4.3 Fan and pump loads

Fans and centrifugal pumps have steady-state torque-speed characteristics which generally have the shape shown in Fig. 11.9.

These characteristics are often approximately represented by assuming that the torque required is proportional to the square of the speed, and thus power proportional to the cube of the speed. We should note however that the approximation is seldom valid at low speeds because most real fans or pumps have a significant static friction or breakaway torque (as shown in Fig. 11.9) which must be overcome when starting.

When we consider the power-speed relationships the striking difference between the constant-torque and fan-type load is underlined. If the motor is rated for continuous operation at the full speed, it will be very lightly loaded (typically around 12% rated power) at half-speed, whereas with the constant torque load the power rating will be 50% at half speed. Fan-type loads which require speed control can therefore be handled by drives such as the inverter-fed cage induction motor, which can run happily at reduced speed and very much reduced torque without the need for additional cooling. If we assume that the rate of acceleration required is modest, the motor will require a torque-speed characteristic which is just a little greater than the load torque at all speeds. This defines the operating region in the torque-speed plane, from which the drive can be selected.

Many fans, which are required to operate almost continuously at rated speed, do not require speed control of course, and are well served by utility-frequency induction motors.

11.5 General application considerations

11.5.1 Regenerative operation and braking

All motors are inherently capable of regenerative operation, but in drives the basic power converter as used for the 'bottom of the range' version will not normally be capable of continuous regenerative operation. The cost of providing for fully regenerative operation is usually considerable, and users should always ask the question 'do I really need it?'

In most cases it is not the recovery of energy for its own sake which is of prime concern, but rather the need to achieve a specified dynamic performance. Where rapid reversal is called for, for example, kinetic energy has to be removed quickly, and, as discussed in the previous section, this implies that the energy is either returned to the supply (regenerative operation) or dissipated (usually in a braking resistor). An important point to bear in mind is that a non-regenerative drive will have an asymmetrical transient speed response, so that when a higher speed is demanded, the extra kinetic energy can be provided quickly, but if a lower speed is demanded, the drive can do no better than reduce the torque to zero and allow the speed to coast down.

11.5.2 Duty cycle and rating

This is a complex matter, which in essence reflects the fact that whereas all motors are governed by a thermal (temperature rise) limitation, there are different patterns of operation which can lead to the same ultimate temperature rise.

Broadly speaking the procedure is to choose the motor on the basis of the root mean square of the power cycle, on the assumption that the losses (and therefore the temperature rise) vary with the square of the load. This is a reasonable approximation for most motors, especially if the variation in power is due to variations in load torque at an essentially constant speed, as is often the case, and the thermal time-constant of the motor is long compared with the period of the loading cycle. (The thermal time-constant has the same significance as it does in relation to any first-order linear system, e.g. an R/C circuit. If the motor is started from ambient temperature and run at a constant load, it takes typically four or five time-constants to reach its steady operating temperature.) Thermal time-constants vary from more than an hour for the largest motors (e.g. in a steel mill) through tens of minutes for medium power machines down to minutes for fractional-horsepower motors and seconds for small stepping motors.

To illustrate the estimation of rating when the load varies periodically, suppose a constant-frequency cage induction motor is required to run at a power of 4 kW for 2 min, followed by 2 min running light, then 2 min at 2 kW, then 2 min running light, this 8-min pattern being repeated continuously. To choose an appropriate power rating we need to find the r.m.s. power, which means exactly what it says, i.e. it is the square root of the mean (average) of the square of the power. The variation of power is shown in the upper part of Fig. 11.10, which has been drawn on the basis that when running light the power is negligible. The 'power squared' is shown in the lower part of the figure.

The average power is 1.5 kW, the average of the power squared is $5\,\text{kW}^2$, and the r.m.s. power is therefore $\sqrt{5}\,\text{kW}$, i.e. 2.24 kW. A motor that is continuously

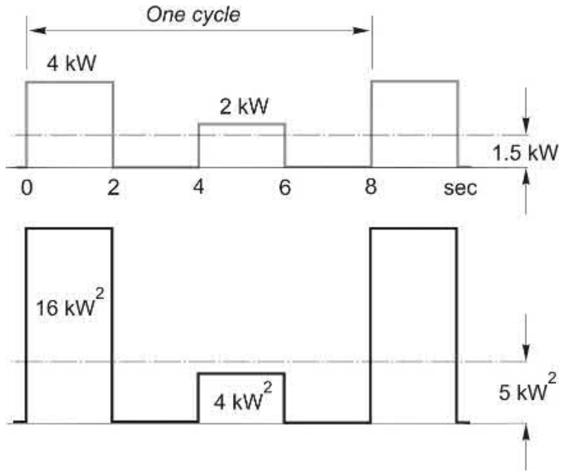

FIG. 11.10 Calculation of r.m.s. power rating for periodically varying load.

rated at 2.24 kW would therefore be suitable for this application, provided of course that it is capable of meeting the overload torque associated with the 4 kW period. The motor must therefore be able to deliver a torque that is greater than the continuous rated torque by a factor of 4/2.25, i.e. 178%: this would be within the capability of most general-purpose induction motors.

Motor suppliers are accustomed to recommending the best type of motor for a given pattern of operation, and they will typically classify the duty type in one of eight standard categories which cover the most commonly encountered modes of operation. As far as rating is concerned the most common classifications are maximum continuous rating, where the motor is capable of operating for an unlimited period, and short time rating, where the motor can only be operated for a limited time (typically 10, 30 or 60 min) starting from ambient temperature. NEMA, in the United States, refer to motor duty cycles in terms of continuous, intermittent or special duty. IEC 60034-1 defines duty cycles in eight ratings S1 … S8.

It is important to add that in controlled speed applications, the duty cycle also impacts the power electronic converter and here the thermal time constants are VERY much shorter than in the motor. The duty cycle can be very important in some applications, and should be clearly stated in a torque and speed-time profile to ensure that there is no ambiguity that could lead to subsequent disappointment.

11.5.3 Enclosures and cooling

There is clearly a world of difference between the harsh environment faced by a winch motor on the deck of an ocean-going ship and the comparative comfort enjoyed by a motor driving the drum of an office photocopier. The former must be protected against the ingress of rain and sea water, while the latter can rely on a dry and largely dust-free atmosphere.

Classifying the extremely diverse range of environments poses a potential problem, but fortunately this is one area where international standards have been agreed and are widely used. The International Electrotechnical Committee (IEC) standards for motor enclosures are now almost universal and take the form of a classification number prefixed by the letters IP, and followed by two digits. The first digit indicates the protection level against ingress of solid particles ranging from 1 (solid bodies greater than 50 mm diameter) to 5 (dust), while the second relates to the level of protection against ingress of water ranging from 1 (dripping water) through 5 (jets of water) to 8 (submersible). A zero in either the first or second digit indicates no protection.

Methods of motor cooling have also been classified and the more common arrangements are indicated by the letters IC followed by two digits, the first of which indicates the cooling arrangement (e.g. 4 indicates cooling through the surface of the frame of the motor) while the second shows how the cooling circuit power is provided (e.g. 1 indicates motor driven fan).

The most common motor enclosures for standard industrial electrical machines are IP 21 and 23 for d.c. motors; IP 44 and 54 for induction motors; and IP 65 and 66 for brushless PM synchronous motors. These classifications are known in the United States, but it is common practice for manufacturers there to adopt less formal designations, for example the following:

Open Drip Proof (ODP): Motors with ventilating openings, which permit passage of external cooling air over and around the windings. (Similar to IP23)

Totally Enclosed Fan Cooled (TEFC): A fan is attached to the shaft to push air over the frame during operation to enhance the cooling process. (Similar to IP54)

Washdown Duty (W): An enclosure designed for use in the food processing industry and other applications routinely exposed to washdown, chemicals, humidity and other severe environments. (Similar to IP65)

11.5.4 Dimensional standards

Standardisation is improving in this area, though it remains far from universal. Such matters as shaft diameter, centre height, mounting arrangements, terminal box position, and overall dimensions are fairly closely defined for the mainstream motors (induction, d.c.) over a wide size range, but standardisation is relatively poor at the low-power end because so many motors are tailor-made for specific applications.

Standardisation is also poor in relation to brushless PM motors. Paradoxically, the lack of standardisation, particularly in terms of frame size, mounting arrangements and shaft dimensions, has been cited as a significant driver of innovation in these machines, whereas the induction motor has been handicapped by its universal success and resulting standardisation!

11.5.5 Supply interaction and harmonics

As we have seen in Section 8.5, most converter-fed drives cause distortion of the utility voltage which can upset other sensitive equipment, particularly in the immediate vicinity of the installation. We have also discussed some of the options available to mitigate such effects. This is however a complex subject area and one which, in most cases, requires a systems overview rather than consideration of a single drive or even group of drives. It requires an intimate knowledge of the supply system, particularly its impedances.

With more and larger drives being installed, the problem of distortion is increasing, and supply authorities therefore react by imposing increasingly stringent statutory limits governing what is allowable.

The usual pattern is for the supply authority to specify the maximum amplitude and spectrum of the harmonic currents at the point of supply to a particular customer. If the proposed installation exceeds these limits, appropriate filter circuits must be connected in parallel with the installation. These can be costly, and their design is far from simple because the electrical characteristics of the supply system need to be known in advance in order to avoid unwanted resonance phenomena. Users need to be alert to the potential problem and, where appropriate, ensure that the drive supplier takes responsibility for handling it.

Measurement of harmonic currents is not straightforward and here the standards help, with IEC 61000-4-7:2002 providing valuable guidance.

11.6 Review questions

(1) The speed-holding accuracy of a 1500 rev/min drive is specified as 0.5% at all speeds below base speed. If the speed reference is set at 75 rev/min, what are the maximum and minimum speeds between which the drive can claim to meet the specification?

(2) A servo motor drives an inertial load via a toothed belt. The motor carries a 12-tooth pulley, and the total inertia of motor and pulley is 0.001 kgm^2. The load inertia (including the load pulley) is 0.009 kgm^2. Find the number of teeth on the load pulley that will maximise the acceleration of the load.

(3) Assuming that the temperature rise of a motor follows an exponential curve, that the final temperature rise is proportional to the total losses and that the losses are proportional to the square of the load power, for how long could a motor with a 30-min thermal time-constant be started from cold and overloaded by 60%?

(4) A special-purpose numerically-controlled machine-tool spindle has a maximum speed of 10,000 rev/min, and requires full torque at all speeds. Peak steady-state power is of the order of 1200 W. The manufacturer wishes to use a direct-drive motor. Discuss options, and suggest what additional information might be required in order to seek the best solution.

(5) When planning to purchase a variable-frequency inverter to provide speed control of an erstwhile fixed-frequency induction motor driving a hoist, what problems should be anticipated if low-speed operation is envisaged?

(6) List two controlled-speed applications for which conventional d.c. motors are not suitable, and suggest an alternative for each.

(7) A standard induction motor can produce twice full load for short periods, but a standard inverter is limited to its rated output. What accounts for the difference?

(8) A pump drive supplied from the 50 Hz utility supply requires a torque of 60 Nm at approximately 1400 rev/min for 1 min, followed by 5 min during which the motor runs unloaded. This cyclic pattern is repeated continuously.

The motor is to be selected from a range of general-purpose cage motors with continuous ratings of 2.2, 3, 4, 5.5, 7.5, 11 and 15 kW. The motors all have full-load slips of approximately 5% and pull-out torques of 200% at slips of approximately 15%.

Select the pole number and power rating, and estimate the running speed when the motor drives the pump.

(9) A speed-controlled drive rated 50 kW at its base speed of 1200 rev/min drives a large circular stone-cutting saw. When the drive is started from rest with the speed reference set to base speed, it accelerates to 1180 rev/min in 4 s, during which time the acceleration is more-or-less uniform. It takes a further second to settle at full speed, after which time the saw engages with the workpiece.

Estimate the total effective inertia of motor and saw. Make clear what assumptions you have had to make. Estimate the stored kinetic energy at full speed, and compare it with the energy supplied during the first 4 s.

(10) When the saw in question 9 is running light at base speed, and the power is switched-off, the speed falls approximately linearly, taking 20 s to reach 90% of base speed. Estimate the friction torque as a percentage of the full-load torque of the motor. Explain how this result justifies any approximations that had to be made in order to answer question 9.

Answers to review questions are given in the Appendix.

Appendix: Solutions to review questions

Chapter 1

(1) The magnetomotive force or m.m.f. is simply the product of the number of turns and the current, i.e. m.m.f. $= 250 \times 8 = 2000$ AT. The 'AT' stands for ampere-turns, but strictly the unit of m.m.f. is the Ampere.

(2) We are told that the magnetic circuit is made of good-quality magnetic steel, which is a coded way of saying that the reluctance of the steel part of the magnetic circuit is negligible in comparison with the air-gap. Under these circumstances all of the m.m.f. provided by the coil (2000 A) is available across the gap, and the flux density in the air-gap is given by Eq. (1.7) as

$$B = \frac{\mu_o NI}{g} = \frac{4\pi \times 10^{-7} \times 2000}{2 \times 10^{-3}} = 1.26\,\mathrm{T}.$$

The question then asks about the flux density in the iron, implying that it is different from that in the air-gap. But as we have seen in the sketches (e.g. Fig. 1.7) the flux density remains the same unless the cross-sectional area changes, so the answer is that the flux density in the iron is the same, i.e. 1.26 T.

If the cross-sectional area were doubled, it would make no difference to the flux density because, as revealed in Eq. (1.7), the air-gap flux density only depends on the m.m.f. and the length of the gap. However, with twice the area, and the same flux density, the total flux would increase by a factor of two.

An alternative way of reaching the same conclusion would be to say that if the cross-sectional area were doubled, the reluctance of the air-gap would be halved, so for a given m.m.f. the flux would double.

(3) Because no information is given about the iron part of the magnetic circuit, we are expected to assume that we can ignore its reluctance, and assume that the only significant reluctances are those due to the air-gaps. This hidden message is reinforced by the use of the word 'estimate' rather than calculate: if the (small) reluctance of the iron parts is neglected, our values of flux density are inevitably going to be slightly on the high side.

If we denote the reluctance of the 0.5 mm air-gap by R, the reluctance of the 1 mm gap will be $2R$, because reluctance is directly proportional to length. The two air-gaps are in series, so the total reluctance is $3R$. The flux through both is the same, and from the Magnetic Ohm's law is given by

$$\Phi = \frac{\text{m.m.f.}}{\text{Reluctance}} = \frac{1200}{3R} = \frac{400}{R}.$$

To find the m.m.f. across each air-gap we apply the magnetic Ohm's law again: so the m.m.f. across the 0.5 mm gap is given by

$$\text{m.m.f.} = \text{Flux} \times \text{Reluctance} = \frac{400}{R} \times R = 400\,\text{A}.$$

Similarly the m.m.f. across the 1 mm gap is given by

$$\text{m.m.f.} = \text{Flux} \times \text{Reluctance} = \frac{400}{R} \times 2R = 800\,\text{A}.$$

To find the flux density we can proceed as in question 2 and apply Eq. (1.7) to either gap. For the 1 mm gap, this yields

$$B = \frac{\mu_o NI}{g} = \frac{4\pi \times 10^{-7} \times 800}{1 \times 10^{-3}} = 1.0\,\text{T}.$$

(4) The new rotor diameter is 299.5 mm instead of 300 mm, so radius has decreased by 0.25 mm, meaning that the new air-gap is 2.25 mm, instead of its original 2 mm. To the uninitiated, this might appear to be of little consequence. However, the reluctance of the air-gap has increased by 12.5%, so, assuming that we can neglect the reluctance of the iron parts, the m.m.f. will have to increase by 12.5% in order to maintain the same flux density.

The power dissipated in the field windings depends on the square of the current, so it has to increase by a factor of $(1.125)^2$, i.e. 1.27. Such a large increase in heating will almost certainly be unacceptable, because unless the original cooling system had been overdesigned (in which case the field windings were running well-below their allowable temperature rise), the permissible rise will certainly be exceeded under the new conditions.

So what are the options? In the short run, the only course open is to tolerate the reduced flux density, which will be 89% of its original value. At rated armature voltage the motor will then run 12.5% above its rated speed, which can be restored by reducing the armature voltage to approximately 89% of its original value. With the same full-load armature current the full-load torque and power will then be 11% below their original values.

In the longer term it might be possible to add some extra turns to the field winding (if there is space!), but even then the field voltage will have to be raised to maintain the same field current.

(5) This is an exercise in applying Eq. (1.2), i.e. $F = BIl$. So for (a) the force is $0.8 \times 4 \times 0.25 = 0.8$ N, while for (b) the total force is $20 \times 0.8 \times 2 \times 0.25 = 8$ N. The result for (b) is 10 times as great as that for (a) because the total current in the coil-side is 40 A.

(6) To estimate the torque we first need to calculate the total tangential electromagnetic force, and then multiply by the radius at which the force acts. We are told that there are 120 conductors, but that only 75% of the circumference is covered by the poles: this means that only 75% of the conductors will be exposed to the radial flux density at any instant, i.e. we can assume that 90 of the 120 conductors make a contribution to the torque. We also make the very important assumption that if the conductors lying under a N pole carry positive current, those under a S-pole will be carrying negative current, so that they all contribute to torque.

The electromagnetic force on one conductor is given by $F = BIl = 0.4 \times 50 \times 0.5 = 10$ N. Note that we use the mean flux density under the pole-face in this calculation: some conductors may be exposed to a slightly higher flux density than others, but we do not need the details as long as we are given the average flux density.

The total tangential force is thus $90 \times 10 = 900$ N. We assume that the force acts at the centre of the conductors, but all we know about the diameter of the conductors is that it must be less than 1 cm in order to fit in the air-gap. If we take the diameter as 0.8 cm, the radius at which the electromagnetic force acts will be $15 + 0.4 = 15.4$ cm or 0.154 m. The torque is therefore given by $T = \text{Force} \times \text{radius} = 900 \times 0.154 = 139$ Nm.

(7) The rewound (220 V) field winding must produce the same m.m.f. as the original winding did when it was supplied at 110 V. If we are not sure how to proceed further, we can begin by arguing that, to do the same job as the old winding, the new one would probably consume the same power, in which case if the current in the 110 V field was I, the current in the 220 V one would be $I/2$. By progressing further with this approach, it should soon become clear whether or not we are on the right lines.

To produce the same m.m.f. with half the current the number of turns must be 2N, where N is the original number of turns. Because the current in the 2N turns is only I/2, the cross-sectional area of the original conductor can be halved, leaving the current density in the wire the same as it was in the original wire. This gives the new winding twice as many turns, but the cross-sectional area of each wire is halved, so the new winding should fit in the same space as the original one. We should now check that our initial assumption—that the power of the new one would be the same as the old one—is correct.

Let the resistance of the original winding be R: it was supplied at 110 V, so the current I was given by $I = 110/R$ and the power consumption was $(110)^2/R$.

The new winding has twice as many turns, so if it were made of the same wire its resistance would be $2R$. But the cross-sectional area of the new wire is only half of the original, so each turn of the new coil has twice the resistance of a turn of the original wire. The total resistance of the new winding is therefore $4R$. The new winding is supplied at 220 V, so the current is $220/4R = 55/R$, i.e. the new current is, as expected, half of the original. The power consumption is $220 \times 55/R = (110)^2/R$, the same as the original. So our original belief that to produce the same m.m.f., the same power would be required, is seen to be correct.

The new winding provides the same m.m.f., and contains the same volume of copper, occupies the same space, and dissipates the same power. We can conclude from this that what really matters is the amount of copper and how hard we work it (i.e. the current density): the number of turns and the cross-sectional area of wire can be chosen to suit any desired operating voltage.

(8) The discussion in question 7 related to the field windings, but the same argument can be applied to all the windings in the machine. It should therefore be clear that there will be very little external difference, except that the cable for the 110 V version would have to be thicker to carry just over twice as much current as the 240 V version.

(9) The windings inevitably have resistance, say R. Hence when they carry a steady current (I), there is a continuous power dissipation of $I^2 R$, and this is equal to the power supplied by the voltage source. If the field windings were made of superconducting wire having zero resistance, the steady-state power dissipation would be zero.

The energy supplied during the transient period (when the winding is first switched on, and the build-up of current is influenced by the inductance of the winding) is divided between that dissipated in the resistance as heat, and that stored in the magnetic field. Once the field is established, and the current becomes steady, no further energy is required to maintain the magnetic field.

(10) Mechanical power is given by the product of torque and speed, so a low-speed motor needs more torque than a high-speed one to produce the same power. The torque of an electrical machine (with a given sophistication of cooling system) is broadly determined by the volume of its rotor, which in turn is closely related to its overall volume. Hence for a given power a motor that runs say 10 times as fast will be 10 times smaller than its low-speed equivalent, and thus will be cheaper to manufacture.

The high-speed motor will have to be geared down to provide a low-speed drive, but this still turns out to be cheaper than a direct-drive motor in most cases.

Chapter 2

(1) (a) Taking the bottom rail as the reference for voltages, the potential of point x oscillates sinusoidally with an amplitude of 20 V, while the potential of point y remains constant at +10 V. The diode with the higher anode potential will conduct, thereby reverse-biassing the other one. The voltage at the load will therefore follow the sinewave while the voltage is greater than +10 V, and be held at +10 V at all other times, as shown by the thick line in part (a) of the figure below.

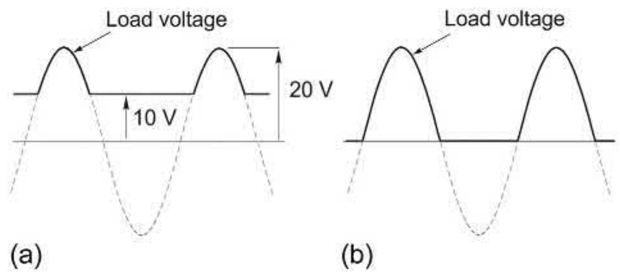

FIG. Q1

(b) This part probably needs a little more thought. If we imagine that V_2 reduces gradually from the +10 V that applied in part (a), we can see that the horizontal line in the load voltage waveform simply moves progressively downwards until it reaches zero. As this happens, the diode D2 conducts for shorter and shorter intervals. When V_2 reaches zero we obtain the waveform shown at (b). We observe that whichever diode is conducting, its current is flowing upwards, and returning by flowing downwards through the load.

When V_2 becomes negative, we might be tempted to think that the output voltage could become negative for part of each cycle, but this is not possible because the current would then be trying to flow up through the load and down through a diode. The diodes can only conduct from anode to cathode (i.e. in the direction of the broad arrow) so we conclude that neither diode will conduct and the waveform will be as shown in (b) regardless of the value of the negative voltage V_2.

(c) Whether the diode is above or below the voltage source makes no difference since they are in series.

(2) The reference to a motor load means that there is inductance present. If, as is most likely, we can assume that the current is continuous, the mean d.c. voltage from a fully-controlled single-phase converter is given by combining Eqs. (2.3), (2.5) to yield

$$V_{dc} = \frac{2\sqrt{2}}{\pi} V_{rms} \cos\alpha.$$

(If the current is discontinuous we cannot determine the voltage without knowing the load and the motor parameters.)

The maximum voltage is obtained when $\alpha = 0$, giving a value of 207 V. This ignores the volt-drop across the diodes, so in practice the true figure will be nearer to 205 V. The reference to a low-impedance utility supply signals that we can neglect the volt-drop due to the supply system impedance.

(3) The mean output voltage is given by Eq. (2.6), i.e.

$$V_{dc} = V_{d0}\cos\alpha = \frac{3}{\pi}\sqrt{2}V_{rms}\cos\alpha.$$

Substituting $V_{rms} = 415$ and $V_{dc} = 300$ gives $\alpha = 57.6°$.

The frequency does not affect the formula for the average d.c. voltage, so there would be no change.

(4) This is the sort of deceptively simple question that can easily trip up the unwary. Whatever approach is taken, experience shows that it is essential to define terms clearly and proceed systematically.

We will begin by noting that the circuit is as shown in Fig. 2.8, and the output voltage waveforms for continuous-current operation are shown in Fig. 2.10. We should also recall that devices T1 and T4 conduct in one half cycle, while T2 and T3 conduct in the other half cycle.

Let us focus attention on the waveform of voltage across T1, i.e. the potential difference between its anode and its cathode. We can see that the anode of T1 is permanently connected to terminal A of the supply, so its anode is always at the potential of terminal A. So we need to see what happens to the potential of the cathode of T1 in order to sketch the voltage across it.

For the half-cycle when T1 and T4 are conducting (which we conventionally take as predominantly during the 'positive' half-cycle of the supply), the forward volt-drop across T1 when it is conducting will be very small, so we will assume for the purposes of sketching that we can take the 'on' voltage across T1 as zero for this 180° of conduction.

During the other 180° conduction period, T2 and T3 are conducting (with negligible volt-drop), so the cathode of T1 is connected—via T2—to terminal B of the supply, while its anode remains as before, i.e. connected to terminal A of the supply. The potential difference between the anode and cathode of T1 is therefore the potential difference between terminals A and B, which is of course the supply voltage, V_{AB}. So for the second period the voltage across T1 is the supply voltage. The complete waveform of the voltage across T1 is therefore as shown in the sketch below.

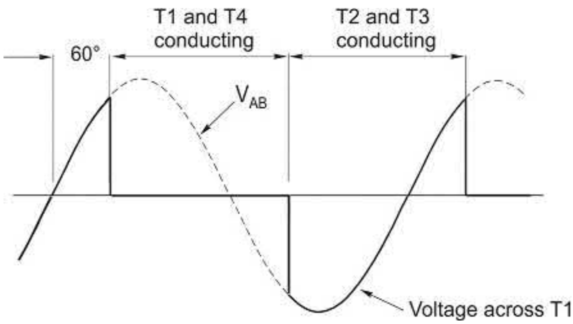

FIG. Q4

(5) We are told that the d.c. load current is smooth at 25 A. Referring to Fig. 2.7, we recall that for one period of 180° the load current flows out from T1 and returns through T4, i.e. upwards through the supply in Fig. 2.7, while during the other period of 180° the load current flows out from T2 and returns through T3, i.e. downwards through the supply. Since the load current is constant at 25 A, the supply current is a 25 A square-wave, as shown in the figure below.

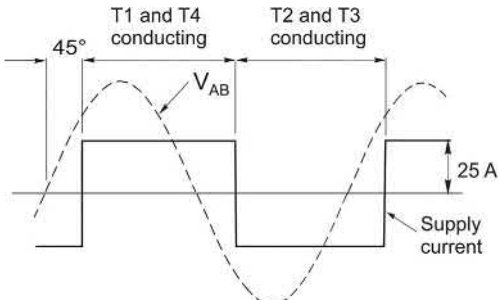

FIG. Q5

We should note that because we have ignored supply inductance (which causes the 'overlap' effect discussed in this chapter), the current commutates instantaneously, and the current waveform is rectangular. In practice the current takes a finite time to commutate, and the waveform is trapezoidal, but still very much non-sinusoidal and far from ideal from the point of view of the supply authorities, who do not welcome harmonic components of current!

To calculate the peak supply power we can see from the sketch that this occurs when the voltage reaches its peak, so the peak power is given by

$$P_{\max} = V_{\max} \times 25 \; = 230\sqrt{2} \times 25 = 8.13\text{kW}.$$

Perhaps the easiest way to obtain the average power is to take advantage of the fact that we are told to ignore the losses in the devices. This means that the average a.c. input power is equal to the mean d.c. power, which is easy to calculate in this instance because the current is constant. When the current is constant, the mean power is simply the mean d.c. voltage (which we can obtain from Eqs. 2.3, 2.5) multiplied by the current.

Hence the mean power is given by

$$P_{av} = V_{dc} \times I_{dc} = \frac{2\sqrt{2}}{\pi}(230)\cos 45^\circ \times 25 = 3.66\text{kW}.$$

It is important to note that the mean power can only be obtained by multiplying the mean voltage by the current if the current is constant. If both the voltage and current vary with time, it is necessary to integrate the instantaneous power (i.e. the product of instantaneous voltage and instantaneous current) to obtain the total energy per cycle, and then divide by the period to obtain the mean power.

An alternative approach to find the mean power directly from the a.c. side exploits the fact that if, as here, the voltage is sinusoidal, the average power can be obtained by finding the r.m.s. value of the fundamental-frequency component of the current, I_1; then taking the product $V_{rms}I_1\cos\varphi$, where φ is the phase angle between the fundamental components of current and voltage.

The amplitude of the fundamental component of a square wave of 25 A can readily be shown to be $\frac{4}{\pi} \times 25 = 31.83$ A, so the r.m.s. of the fundamental component of current is $\frac{31.83}{\sqrt{2}} = 22.51$ A. Hence the average power is given by $230 \times 22.51 \times \cos 45^\circ = 3.66$ kW, as above.

(6) The average load voltage is 20 V and the source voltage is 100 V, so it follows that the transistor is 'on' for one-fifth of the time. The load (motor) current is constant at 5 A, so the waveforms of source and load are as shown below.

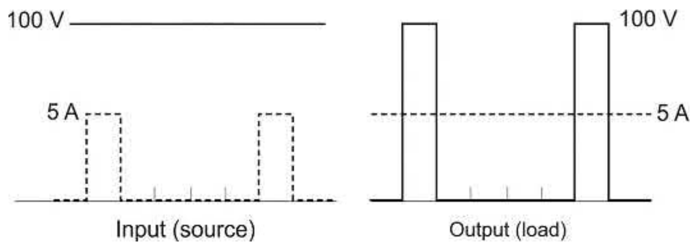

FIG. Q6

The average input current is 1 A, so the input power is 100 V × 1 A = 100 W. The average load voltage is 20 V, so the load power is 20 V × 5 A = 100 W.

To draw a parallel with a transformer we focus on the average values of the input and output voltages and currents, in which case the chopper appears to function as an ideal step-down transformer with a turns ratio of 5 to 1: the secondary voltage is one-fifth of the primary (input) voltage, while the secondary current is five times the primary current. And like an ideal transformer, the input and output powers are equal.

(7) The circuit and the load voltage and current waveforms are shown in Fig. 2.4. The current is described as being 'almost constant', which means that we treat it as constant for the purposes of calculation.

(i) The average load voltage is given by the product of the resistance and the average current, i.e. 8 × 5 = 40 V. The voltage waveform itself is rectangular, being 150 V when the transistor is 'on' and zero when the transistor is 'off' and the current is freewheeling through the diode. If we denote the on time (the mark) by T_{on} and the off time (the space) by T_{off}, the average voltage is given by $150 \times \frac{T_{on}}{T_{on}+T_{off}} = 40, \ \therefore \frac{T_{on}}{T_{on}+T_{off}} = 0.267, \ \therefore \frac{T_{on}}{T_{off}} = 0.364.$

(ii) Because the load current is constant, we can obtain the average load power from the product of the mean load voltage and the current, i.e. 40 × 5 = 200 W.

(iii) We are told to treat all the devices as ideal, so there are no losses and the input power must equal the output power, i.e. 200 W. As a check, we can easily calculate the average source power because the input voltage is constant (at 150 V), so the average power is simply the product of 150 V and the average source current.

The source current is 5 A when the transistor is on, and zero when it is off. So the average source current is $5 \times \frac{T_{on}}{T_{on}+T_{off}} = 5(0.267) = 1.33\,A$. Hence the mean source power is 150 × 1.33 = 200 W, as above.

(8) Taking the possibilities in the order given:

As there is no tendency for current to flow upwards through the switch there is no need to prevent it—and in any event a MOSFET contains an intrinsic diode that is reverse-biased during normal operation, but would allow reverse current to flow if necessary.

Inductors do not need protection from high voltages: any 'high voltages' present are likely to be generated by high rates of change of current through inductors anyway.

Voltage supplies are very unlikely to be subject to any restriction on the rate of change of current, but this answer may have come about as a result of confusion over the need to limit the rate of rise of current in some semiconductor devices. Where such restriction applies, an inductor would serve to limit the rate of rise of current, not a diode.

Limiting the voltage across the MOSFET is the real reason for the diode. There will be a maximum allowable voltage across the MOSFET (i.e. between Drain and Source), and if this is exceeded the device will break down. When no current is flowing in the load (Fig. Q8), the voltage across the MOSFET is equal to the supply voltage, so clearly the device must be capable of withstanding supply voltage. However, the real danger arises because of the load inductance, as discussed below.

The voltage across, and the current through, an inductor are related by the equation $v = L\frac{di}{dt}$, i.e. the voltage is determined by the rate of change of current, or vice-versa. So if we attempt to open a switch in a circuit containing an inductor through which current is flowing, thereby forcing a very rapid rate of change of current (i.e. $\frac{di}{dt}$ tends to infinity) a very high voltage will be developed across the terminals.

For example suppose that there was no diode in Fig. Q8, and, with the MOSFET switched on, a steady current of 2 A was flowing through an inductance of 50 mH. If we then switched off the MOSFET such that the current was reduced to zero at a uniform rate in 1 μs, the voltage across the inductor would be 100 kV! We know from Kirchoff's voltage law that the sum of the voltages around a circuit must equal the supply voltage, so it follows that almost all of the 100 kV would appear across the MOSFET, leading to its immediate destruction.

The real problem lies in the energy stored in the inductance, which is given by $E = \frac{1}{2}Li^2$. If we attempt to reduce the current to zero instantaneously, we are trying to destroy energy in zero time, which corresponds to infinite power, which is clearly impossible. The alternative is to provide a mechanism whereby the stored energy in the inductor can be released in a more measured fashion, and this is where the 'freewheel' diode comes in. When the switch opens, the diode provides an alternative path for the current that is flowing down through the load, so the current can continue by flowing in the closed path shown at (b) in Fig. 2.4. At first the current is undiminished, so the stored energy is unchanged, but because heat is dissipated in the resistance of the load, the current decays exponentially as the energy stored reduces. The term 'freewheeling' arises by analogy with riding a bicycle, where, having built-up kinetic energy by hard footwork, we can rest the pedals and let the stored energy sustain our motion until friction brings us to rest.

The last answer suggesting that the diode is there to dissipate stored energy is partially correct in that some energy is indeed dissipated in a real diode (though not in an ideal one). But most of the stored energy will be dissipated in the load resistance, and the resistance of the inductor itself.

(9) We should identify the right answer (100.7 V) and explain why before speculating on the origins of the remainder.

As explained in the answer to question 8, the freewheel diode conducts when the MOSFET switches off, the load current flowing upwards through the diode. To find the voltage across the MOSFET in this condition we need to find the potential of the anode of the diode with respect to the ground reference (i.e. the bottom of the supply). We know that the forward volt-drop across the diode is 0.7 V, so the potential of the anode is 0.7 V higher than the potential of the cathode. The cathode is connected to the top of the supply, so its potential with respect to ground is 100 V. Hence the potential difference (voltage) across the MOSFET is 100.7 V. This is therefore the maximum voltage that the MOSFET has to withstand.

The figure of 99.3 V clearly comes about because of confusion over the sign of the voltage across the diode. The suggestion 'zero' presumably arises because it is believed that since there is no current through the MOSFET during freewheeling, its voltage must also be zero. The voltage across the load (resistance and inductance) is 0.7 V during freewheeling, because the load is (almost) short-circuited by the diode. And the suggestion that the voltage depends on the inductance is understandable but wrong because while the value of inductance determines the stored energy and therefore the duration of freewheeling, the voltage across the MOSFET will be 100.7 V regardless of the inductance.

Chapter 3

The equivalent circuit shown here should be in our minds when we tackle any d.c. machine questions:

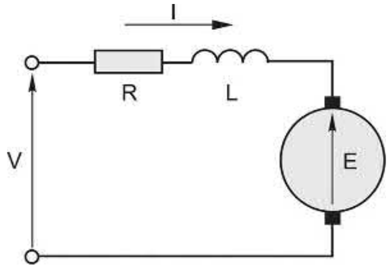

In the majority of cases we will be considering steady-state operation, so the current will be constant and therefore we can ignore the armature inductance in our calculations.

Unless told to the contrary, we will assume that the volt-drop across the brushes can be ignored.

(1) (a) As we are not told otherwise, we are expected to assume that the question refers to the steady-state running speed, in which case the answer is that the speed is determined by the armature voltage. Justification, if required, would be along the lines below.

Whenever the speed is steady, the motor torque must be equal and opposite to the load torque. Except for very tiny d.c. motors it is safe to assume that when the motor is unloaded, the friction torque is very small, so the motor torque would also be very small. Motor torque is proportional to armature current, so we can expect the armature current of an unloaded motor to be very small, hence the term IR in the armature voltage equation $V = E + IR$ is negligible and we can say that at no-load V is approximately equal to V. E is the motional e.m.f. induced in the armature, and is directly proportional to the angular velocity (speed), i.e. $E = k\omega$, hence the speed is given by $\omega = \frac{V}{k}$, i.e. the speed is determined by the applied voltage.

Note that in the majority of d.c. motors the term IR will be small compared with the armature voltage V even when the motor is on load and the current I is not small, so to a first approximation we can say that the on-load speed will also be determined by the applied voltage, the speed when loaded only being slightly less than that of the unloaded motor.

(b) As discussed in the answer to question 1, the steady running current must be such as to produce a torque equal and opposite to the load torque, so in the steady state it is the load torque that determines the armature current.

(c) The answer to part (a) indicates that the steady running current is always determined by the load torque. When no 'real' load torque is applied, we are left with friction, due to bearings, fan, and (especially in a d.c. machine) brush friction. The friction torque is therefore reflected in the no-load current.

(d) The answers are that the drop in speed from no-load depends directly on the load torque and the armature resistance.

First, let us consider the effect of load torque. For any given load, the speed will settle when the motor torque T_m equals the load torque T_L. Motor torque is proportional to armature current, i.e. $T_m = kI$, hence the steady current is given by $I = \frac{T_L}{k}$, i.e. the steady armature current is proportional to the load torque. Combining the armature voltage equation $V = E + IR$ and the e.m.f. equation $E = k\omega$, and substituting for I from above gives the speed as

$$\omega = \frac{V}{k} - \frac{R}{k^2} T_L$$

This equation shows that the no-load speed (i.e. when $T_L = 0$) is given by $\omega_0 = \frac{V}{k}$, and the drop in speed that is attributable to load is given by $\left(\frac{R}{k^2}\right) T_L$. The drop in speed is therefore directly proportional to the load torque and to the armature resistance.

We note also that the drop in speed for a given load is inversely proportional to the square of the motor constant. So if we were to reduce the field current so that, say, the flux was halved to double the no-load speed, we would find that because k had been halved, the drop in speed for a given load torque would be four times as great as with full flux. This matter was discussed in Chapter 3 and illustrated in Fig. 3.12.

What the question means when it refers to 'little ones slow down more than large ones' really means that the percentage drop in speed between no-load and full-load is usually higher in small motors than in large ones. The reason is simply that in small machines the term IR represents a higher fraction of the applied voltage than it does in large machines.

Alternatively we could say that the reason is that 'the per-unit resistance is higher in small machines', meaning that the ratio $\left(\text{Armature resistance} \div \frac{\text{Rated voltage}}{\text{Rated current}}\right)$ is larger in small machines than in large ones.

(2) To reverse the direction of rotation we must reverse the direction of current in the armature or the direction of current in the field. In a separately-excited motor or a shunt motor it is usually easiest to reverse the connections to the field, because the field current is less and the wires are thinner. In a series motor, the field and armature carry the same current, so either can be reversed.

(3) If the motor is producing more than its continuously-rated torque its armature current will be above the continuously-rated value and therefore it will overheat. If the cooling of an existing motor is improved it should be possible to increase the continuous rating without overheating, but other problems due to commutation and brush wear must be anticipated.

(4) We can consider the no-load condition, when the motional e.m.f. E is very nearly equal to the applied voltage V. If we reduce the flux that is cut by the armature conductors, they will have to cut through the weakened flux faster to achieve the same e.m.f., so the weaker the field, the higher the no-load speed.

Alternatively we can use the result from the solution to question 1(d), i.e. that the no-load speed is given by $\omega_0 = \frac{V}{k}$, where k is the e.m.f. constant, which is proportional to the field flux. If the flux is reduced, so is k, leading to a higher no-load speed.

(5)

$$E = V - IR = 220 - 15(0.8) = 208\,\text{V}.$$

We are told that the field current is 'suddenly' reduced by 10%, and that the flux is proportional to the field current. (We know that the current in an inductance (the field circuit) cannot change instantaneously, so we suppose that what the question means is 'very quickly, compared with any subsequent changes that may be initiated by the reduction in flux'.)

A reduction of flux by 10% will cause the motional e.m.f. to reduce by 10%, so the new e.m.f. is $0.9 \times 208 = 187.2\,\text{V}$. So the new current will be given by $I = \frac{220-187.2}{0.8} = 41\,\text{A}$. Note that a modest reduction of only 10% in the flux causes a dramatic increase in the armature current, which jumps from 15 to 41 A.

The increased current will lead to more torque, but not quite in proportion to the increase in current because there has been a reduction in the flux. In most of the calculations in the book the flux has remained constant, in which case the torque is proportional to the current. But the torque depends on the product of the flux and the current, so if we denote the original flux by Φ, the 'new' and 'original' torques are in the ratio $\frac{0.9\Phi \times 41}{\Phi \times 15} = 2.46$. The surge of torque will lead to a rapid acceleration to the new (higher) steady speed.

(6) (a) $k = \frac{\text{e.m.f}}{\text{speed}} = \frac{110}{1500 \times \frac{2\pi}{60}} = 0.70\,\text{V/rad/s} = 0.70\,\text{Nm/A}$. Hence when $I = 10\,\text{A}$, Torque $= 7\,\text{Nm}$.

(b) The gravitational force on the mass is given by $F = mg = 5 \times 9.81 = 49.05\text{N}$. Hence the torque exerted at the motor shaft is $49.05 \times 0.8 = 39.24\,\text{Nm}$.

The motor must exert an equal and opposite torque to achieve equilibrium, so the motor current is given by $39.24/0.70 = 56.06\,\text{A}$.

The stability question can be addressed by considering that, with the arm horizontal and zero net torque, we make a small change to one of the parameters and see if the system takes up a new equilibrium. If we slightly reduced the current in the motor, the load torque would then exceed the motor torque and the weight would move downwards. But as it did so the torque it exerts reduces because the line of action of the force moves closer to the axis of the motor. So when it has moved down to the point where the load torque again equals the motor torque, it will find a stable equilibrium.

However, if we slightly increase the current, the motor torque will be greater than the load torque and the weight will begin to move upwards. In so doing its line of force moves closer to the axis and the torque it exerts gets less. The amount by which the motor torque exceeds the load torque therefore increases with movement, and we have an unstable equilibrium.

So there isn't a simple answer to the question 'is it stable', because the stability depends on how the equilibrium is disturbed.

(c) This is another straightforward exercise using the armature voltage equation. First we need to find the back e.m.f. which is given by $E = k\omega = 0.70 \times \left(1430 \times \frac{2\pi}{60}\right) = 104.8\,\text{V}$. Then apply the armature voltage equation $V = E + IR$ to obtain $IR = 110 - 104.8 = 5.2\,\text{V}$. Hence since $I = 25\,\text{A}$, $R = 0.2\,\Omega$.

To drive a current of 56 A through $0.2\,\Omega$ when the motor is at rest (i.e. $E = 0$) requires a voltage of $56 \times 0.2 = 11.5\,\text{V}$.

(d) The machine is now acting as a generator, supplying power to a system at 110 V. The generated e.m.f. E is greater than the system voltage by IR.

If the power supplied to the system is 3500 W at 110 V, the current is $3500/110 = 31.82\,\text{A}$. Hence the generated e.m.f. is given by $E = 110 + 31.82(0.2) = 116.4\,\text{V}$. The speed is given by $\omega = \frac{E}{k} = \frac{116.4}{0.70} = 166.28\,\text{rad/s} = 1588\,\text{rev/min}$.

The corresponding torque is given by $31.82 \times 0.70 = 22.27\,\text{Nm}$.

The electromechanical power is $EI = 116.4 \times 31.82 = 3704\,\text{W}$, to which we must add 200 W to find the additional mechanical input power, and 100 W for the input power to the field, making a total of 4004 W. The useful output power, supplied to the system, is 3500 W, so the efficiency is $(3500/4004) \times 100\%$, i.e. 87.4%.

(7) In a linear system work is force times distance: in a rotary system force is replaced by torque and linear distance becomes rotary distance, i.e. angle. So in a rotary system, work is torque times angle.

Mechanical power is the rate of doing work, i.e. work/time. So in a rotary system mechanical power is torque times angle over time. But, assuming that speed is constant, angle over time is angular velocity, and power is thus given by

$$\text{Power} = \text{Torque} \times \text{Angular velocity, i.e. } P = T\omega.$$

We have the equations $T = kI$ and $E = k\omega$. Hence $P = T\omega = kI \times \frac{E}{k} = EI$.

(8) (a) When the motor is at rest the back e.m.f. is zero so if rated voltage (V) is applied the current will be V/R_a, where R_a is the armature resistance. In large d.c. motors the current V/R_a is very much greater than the rated

current. The motor will almost certainly be supplied from a thyristor converter, in which the thyristors would not be able to withstand such a large current. So the control scheme would automatically limit the voltage applied to the motor in order to restrict the current to an acceptable level.
(b) The torque required to maintain a steady speed when a motor is unloaded is very small. The torque produced by the motor is proportional to the current, so the no-load current is very small.

The current is given by $I = \frac{V-E}{R}$, where V is the applied voltage, R is the armature resistance and E is the motional or back e.m.f induced in the motor. As explained above, the no-load current is very small, which indicates that the back e.m.f. E is almost equal to the applied voltage. The motional e.m.f. is proportional to the speed, so the no-load speed is almost proportional to the applied voltage.
(c) When the motor is running at a steady speed, the torque it produces is equal to the load torque. When the load torque increases the previous state of equilibrium is disturbed because the load torque now exceeds the motor torque, so the net torque is negative and the system decelerates. The motional e.m.f. (E) is proportional to speed, so E reduces.

The armature current is given by $I = \frac{V-E}{R}$, where V is the applied voltage, R is the armature resistance and E is the back e.m.f. As E reduces, the current increases, and so does the motor torque. The net decelerating torque then reduces, so the deceleration is reduced but will continue until the speed has fallen to the point where the motor torque equals the load torque, at which point equilibrium will be restored, but at a new (lower) speed. The smaller the armature resistance, the less the speed has to drop in order for the current to reach the new load level.
(d) The voltage equation for a field winding is $v = ri + L\frac{di}{dt}$, so the instantaneous power is given by $vi = i^2 R_f + L_f i \frac{di}{dt}$. The first term in the power equation represents the loss of heat due to the field winding resistance, while the second term represents the rate of change of stored energy in the magnetic field.

Under d.c. conditions the second term is zero because $\frac{di}{dt}$ is zero, indicating that once the magnetic field has been established the energy stored remains constant. The first term ($I_{dc}^2 R_f$) represents the heat loss per second due to resistance (copper loss) and this has to be supplied continuously, even though none of it appears as mechanical output power. If the resistance could be made zero (e.g. with a superconducting winding) the power input would be zero once the field current had been established.
(e) If the supply to the field is pure d.c., then apart from the very short periods when the field flux is changing, the flux in the magnetic circuit is constant, so there is no danger of induced eddy currents in the body of the pole and therefore no need for it to be laminated.

When the supply is from a thyristor converter, however, there will be an additional alternating component of flux in the poles, which must therefore be laminated to minimise eddy-current losses.

(9) At the fundamental level it is true that in principle any electrical machine with rated voltage V and rated current I could be rewound to operate at voltage kV and rated current I/k, and that the rewound motor would contain the same amounts of active materials (copper and iron) and have the same performance (in particular the same power (VI) as the original.

However, in the case of the low-voltage d.c. motor, there are several additional factors which complicate matters.

The first relates to the size of the commutator. For a given power, the current is inversely proportional to the voltage, so a low voltage motor obviously has a higher current than a high-voltage one. The area of brush in contact with the commutator is determined by the current it has to carry, so the lower the voltage, the bigger the brushgear/commutator. In hand tools space and weight are at a premium so the high-voltage motor is at an advantage.

The second matter stems from the fact that the voltage/current characteristic of the carbon brushes is non-linear, so that under normal operation the volt-drop across the brushes contains a more-or-less fixed component that is of the order of 1 V, regardless of current. In a 110 V motor the loss of 1 V is not serious: but in battery-powered tools the supply voltage is only a few volts, in which case the loss of 1 V is serious, but becomes less so the higher the supply voltage. It is therefore desirable to avoid low voltages from the point of view of efficient use of energy.

The third factor relates to the properties of the semiconductor switches used in the chopper drive that provides speed control. The on-state volt-drop in transistors and diodes is (rather like the brush-drop referred to above) largely independent of the current, so that the on-state power loss is more-or-less proportional to the current. So when efficiency is important it is preferable to handle a given power at a high voltage and low current, rather than at a low voltage and high current.

Taken together these factors indicate that for a given output power the designer should aim to minimise the current, so that the higher the power, the higher the voltage.

(10) (a) We can find the machine constant from the data given in the first paragraph. When the machine is on open circuit there is no volt-drop across the armature resistances and the terminal voltage is therefore the same as the induced e.m.f. Hence using the relationship $E = k\omega$,

$$k = \frac{220}{1500 \times \frac{2\pi}{60}} = 1.40\,\mathrm{V/rad/s} = 1.40\,\mathrm{Nm/A}.$$

The question is all about steady-state conditions, so we must expect to make use of the fact that if a linear (or rotary) system is not accelerating, the resultant force (or torque) must be zero. We can make use of this knowledge to find the tension in the rope (Fig. Q10A), which we need to know in order to work out the torque exerted by the load.

The two forces acting on the mass m are the gravitational force (downwards), equal to mg, and the tension in the rope (F) upwards. Since the descent velocity is to be constant, the net force must be zero, i.e. $F = mg = 14.27 \times 9.81 = 140\,\text{N}$.

At the drum, this (downwards) tension acts at a radius of 10 cm, so the torque exerted by the load is $140 \times 0.1 = 14\,\text{Nm}$. We are not told anything about friction torque so all we can do is to assume it is negligible, so the total load torque is 14 Nm.

The linear speed of the rope at the drum is given as 15 m/s, the circumference of the drum is 0.2π, and the speed of rotation of the drum and machine is therefore $15/0.2\pi$ rev/s or 150 rad/s.

Because the speed is steady there is no acceleration, and the machine torque must be equal and opposite to the load torque, i.e. the machine torque must be 14 Nm at a speed of 150 rad/s.

We keep referring to the 'machine' rather than the 'motor' because in this application we are using the machine to restrain the descending load, not to drive it down. We need the machine torque to act in the opposite direction to the load torque, which it will do automatically if we complete the armature circuit with a resistance (as in Fig. Q10B), thereby allowing the generated e.m.f. to drive a current in the same direction as the e.m.f. (In contrast, if for some reason we wanted to operate as a motor rotating in the same direction, i.e. to drive down the load, we would apply a voltage greater than the e.m.f. and force the current to flow in the opposite direction to the e.m.f., yielding a driving torque rather than a braking torque.)

Now that we know the speed is 150 rad/s we can calculate the generated e.m.f. as $E = k\omega = 1.4 \times 150 = 210\,\text{V}$. The machine torque has to be 14 Nm, so, using $T_m = kI$, the armature current must be 10 A.

The total resistance in the armature circuit is therefore given by $210/10 = 21\,\Omega$. The armature resistance itself is $0.5\,\Omega$, so the external resistance required is $20.5\,\Omega$.

(b) The power dissipated in the external resistor is given by $I^2R = 100 \times 20.5 = 2050\,\text{W}$, while the power dissipated in the machine armature is $100 \times 0.5 = 50\,\text{W}$.

The electrical generated power is provided from its mechanical input power, which is derived from the steady reduction in potential energy of the lowering mass. We can check the power by considering the power (force times speed) of the falling mass, i.e. P_{mech} = Force × speed = 140 × 15 = 2100 W.

In this question we have ignored all but the armature copper loss in order to simplify the calculations, but nevertheless the situation is representative of many real-life applications, such as dynamic braking of railway vehicles where kinetic energy is dumped in large resistors, often mounted on the roof to assist cooling.

Chapter 4

(1) Under no-load conditions the speed of a d.c. motor is almost exactly proportional to the armature voltage, so when the speed reference is increased from 50% to 100% the armature voltage will double. Assuming the control scheme is good, the actual speed should precisely track the reference, so the new tacho voltage will be exactly twice what it was.

To compare the armature currents we would need to know how the friction torque varied with speed. As we have no information all we can say is that if the friction torque was independent of speed (a reasonable assumption for a separately-ventilated motor of the type usually employed), the no-load current will be independent of speed. Ideally the no-load current should be zero, and in practice it is seldom more than a percent or two of full-load current, except perhaps in very small machines.

(2) With a PI speed controller there will be no steady-state error in the speed, so after the transient has settled the on-load speed will be exactly 50%: the tacho voltage will therefore be the same as before the load was applied.

The torque of a d.c. motor is proportional to the armature current, so a load torque of 100% means that the armature current must be at its rated value.

If we denote the induced e.m.f. in the motor before the load was applied by E_1, the corresponding armature voltage is given by $V_1 = E_1 + I_1R$, where I_1 is the no-load current. When the motor is loaded, the control system will adjust the armature voltage to achieve the same speed, so the induced e.m.f will be exactly the same and the on-load voltage will be given by $V_2 = E_1 + I_2R$, where I_2 is the full-load current. The increase in armature voltage is thus $V_2 - V_1 = (I_2 - I_1)R$. In practice, because the armature resistance is small, the increase in voltage from no-load to full-load will also be only a few percent of rated voltage.

(3) At first there will be a 100% speed error, so the speed controller output will saturate and demand full (rated) current to provide maximum acceleration. The current will be maintained at 100% during acceleration, so the motor torque will be constant and with negligible friction torque the acceleration will be constant so the speed will increase at a uniform rate.

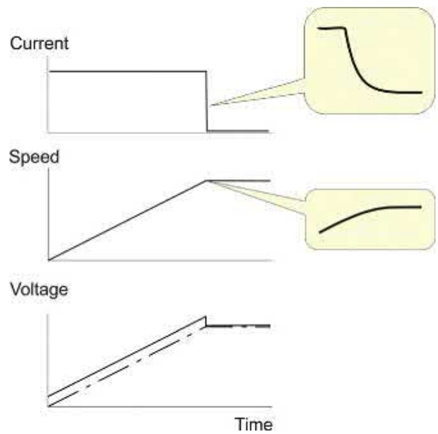

FIG. Q3

The output from the speed error amplifier will remain constant, and the current will therefore remain at full value, until the speed rises to within a few percent of the target. At this stage, the speed controller comes out of saturation and enters its linear regime, the demanded current (and the torque) reducing as the final speed is approached smoothly.

While the acceleration is constant the speed increases linearly, and so therefore does the induced e.m.f, E. The armature voltage is given by $V=E+IR$, so since I is held constant during most of the run-up phase, the armature voltage also increases linearly with time, as shown in Fig. Q3.

(4) (a) The drive will react to the drop in voltage (which will lead to a drop in armature voltage, torque, and speed) by reducing the firing angle a little so that the armature voltage is returned to its correct level to restore the speed to target.

A well-engineered drive will include provision for variation in the supply voltage of at least 10%, by arranging that under normal conditions the full-load armature voltage is obtained with a converter firing angle of say 15–20°. In this way, if the utility voltage should fall, the firing angle can be reduced towards zero in order to maintain the output voltage.

(b) If the tacho feedback disappears, the drive will act as if the speed was zero, so with a target speed of 50% the large speed error would cause the speed controller output to saturate and demand full current. The motor will then accelerate at full current (see answer to Q3 above) until the firing angle of the converter has been reduced to zero and the armature voltage is at its maximum possible value. The motor will then run at somewhat above base speed, drawing a small no-load current.

If the drive includes tacho loss detection circuitry it may shut-down automatically, or switch-over to armature voltage feedback.

(c) If the motor was stopped by some mechanical means, there would be a large speed error and the speed controller output would saturate and demand full current. The current controller would reduce the output voltage of the converter to a very low level because there would be no back e.m.f. with the motor stalled. A sophisticated drive would recognise that full current and no motion indicates trouble, and time-out after a few seconds at most.

(d) This would be much the same as (c), except that the output voltage from the converter would be even lower, depending on the resistance of the 'short-circuit'. (A large stationary d.c. motor is almost a short-circuit anyway!)

(e) This is very serious, since a principal function of the inner current-feedback loop is to protect the thyristors from the danger of excess current. With no current feedback, the current controller will sense a large error, and immediately increase the output voltage from the converter in an attempt to raise the current. Given the very delicate balance that has to be maintained between V and E to avoid excessive currents, the current will inevitably shoot up sufficiently rapidly to blow the expensive fuses that are the last line of defence, but in all probability some of the thyristors will be lost.

(5) When the armature current is discontinuous, the torque-speed curve for a given converter firing angle is very poor: a modest increase in torque causes a very large drop in speed. This occurs because when the current is discontinuous, the output voltage of the converter falls substantially as the current (i.e. the load) is increased. When the load increases and the current becomes continuous, the output voltage from the converter is almost independent of the load, and the speed therefore remains almost constant over a wide range of load.

Although the undesirable effects of discontinuous current can a largely be masked by the operation of the closed-loop speed control system, it is more difficult to optimise the control system (especially the transient response) when the inherent behaviour of the motor itself varies so markedly according to the load.

(6) In dynamic braking mechanical energy is converted to electrical form and then dissipated in resistors. Regenerative braking involves converting mechanical energy into electrical energy and returning it to the supply system.

(7) The output voltage waveform from a thyristor converter consist of rectified chunks of the incoming a.c. mains supply, as shown for example in Fig. Q7A. Ideally, since this is to be the armature voltage for a d.c. motor, we would like pure d.c. (shown by the dotted lines) but as we can see the actual waveform contains a lot of 'a.c' as well.

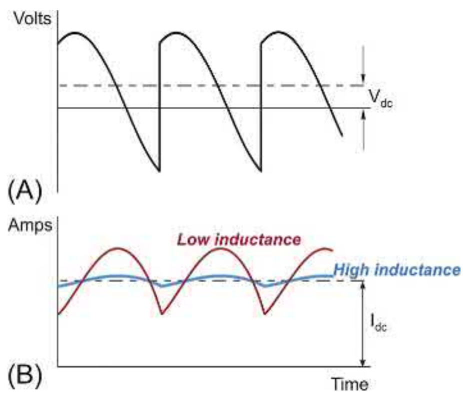

FIG. Q7

The corresponding armature current waveforms are shown in Fig. Q7B. The effect of the armature inductance is to make the current waveform a great deal smoother than the voltage waveforms, which is very desirable because the torque is proportional to the current, and we want to minimise torque pulsations. The higher the inductance, the smoother the current, as shown in Fig. Q7B. (Those who are familiar with a.c. circuit theory will explain the smoothing effect of the armature inductance in terms of inductive reactance ωL. The 'a.c' component of the voltage contains a series of harmonic terms (the lowest of which is 100 Hz when the supply is 50 Hz). The reactance is proportional to frequency, so higher-frequency voltage harmonics produce very little harmonic current, so the current waveform looks much smoother than the voltage waveform.)

The disadvantage of a high armature inductance is that to cause the current in an inductance L to change by an amount ΔI requires a voltage V for a time Δt such that $V\Delta t = L\Delta I$.

In drives, we usually want the transient response of the inner (current-control) loop to be as fast as possible, which means that we apply the highest available voltage in order to maximise the rate of change of current and minimise the time taken to achieve a given change in current. From the equation above we can see that the higher the inductance, the longer we have to apply the voltage, so in terms of transient response, the lower the inductance the better.

(8) The axes are as shown in Fig. Q8, which is the same as in the question but rotated through 90°. (The diagram in the question has the speed axis vertical, which seems to be the preference of mechanical engineers.)

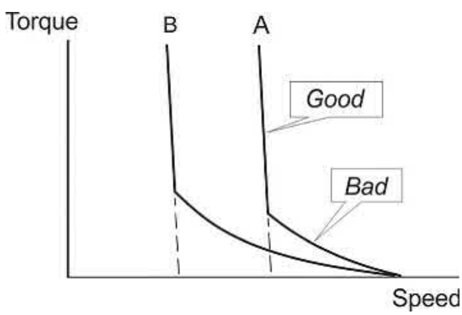

FIG. Q8

The good and bad sections are labelled in Fig. Q8: a good characteristic has only a small change in speed with load, and a bad characteristic is one in which the speed falls significantly when load is increased.

The abrupt change in character occurs at the point where the armature current changes from discontinuous to continuous. When current is discontinuous the converter output voltage depends on the current: the higher the current (i.e. the higher the torque) the lower the voltage, and hence the lower the speed. When the average current has become large enough that the current is continuous (i.e. it never falls to zero), the output voltage of the converter no longer depends on the current, and the speed is therefore almost constant regardless of the load.

We are told that the converter is fully-controlled, so the mean d.c. voltage with $\alpha = 5°$ is given by $V_{dc} = V_{do} \cos 5° \approx V_{do}$. There is no scale on the diagram, but we note that in the continuous current region, curve B corresponds to roughly half of the speed of curve A. The mean d.c. voltage for curve B must therefore be half of that for A, so if the firing angle is α_B, then $\cos\alpha_B = 0.5$, hence $\alpha_B = 60°$.

Adding additional armature circuit inductance makes the armature current smoother, and therefore reduces the likelihood of discontinuous current. The effect on the torque-speed characteristic is to extend the straight portion as shown by the dotted lines.

(9) (a) Since both motors are coupled to the same shaft, their speeds will always be the same, so to share the mechanical work we need to arrange that the torque provided by each motor is proportional to its power rating: when the original motor is at half rated torque, we want the other one to be at half rated torque, and so on.

The 150 kW drive has been chosen as the master, i.e. it has its outer (speed control) loop operational, as well as its inner current-control loop. The 100 kW drive only has its inner loop operational: its current reference is derived from that of the master.

(b) The current reference signal is fed to both drives so that when, for example, 50% of rated current is demanded in the master, 50% of the (different) rated current is also demanded from the slave.

(c) It would not be a good idea to ask both drives to operate in the speed control mode, as unless they were precisely matched, there would be a tendency for them to end up fighting one another.

(d) It could be argued that by making the slave machine current track the actual current in the master, rather than the current reference of the master, the slave current would more faithfully follow the master current. On the other hand, this would mean that if the master current control went astray, so would the slave. So on balance, it's probably not a good idea.

Chapter 5

(1) The maximum synchronous speed for a 50 Hz motor is 3000 rev/min, which is obtained with a 2-pole winding. A speed of 2950 rev/min is clearly appropriate for a 2-pole machine, as the slip of 1.67% is in line with our expectation that full-load slip will not exceed a few percent.

(2) (a) The synchronous speed is given by $N_S = \frac{120f}{p} = \frac{120\times60}{4} = 1800\,\text{rev/min}$. The actual speed is given by $N = N_S(1-s) = 1800(1-0.04) = 1728$ rev/min.

(b) The rotor frequency is given by $f_r = s\,f = 0.04 \times 60 = 2.4$ Hz.

(c) The speed of the 4-pole rotor induced current wave *relative to the rotor* is given by

$$n_r = \frac{120 f_r}{p} = \frac{120 \times 2.4}{4} = 72\,\text{rev/min}.$$

(d) The speed of the rotor induced current wave relative to the stator is 1728 + 72 = 1800 rev/min. This example illustrates the fact that the induced rotor current wave is perceived at the stator as a synchronous wave, so although the rotor currents are at slip frequency, their influence as seen at the stator takes place at the supply frequency.

(3) (a) The synchronous speed depends only on pole-number and frequency, so it is unaffected by the voltage.
(b) The air-gap flux is proportional to voltage and inversely proportional to frequency. Hence the air-gap flux will reduce in the ratio 380/440, i.e. the new flux will be 86.4% of the original flux.
(c) At all speeds, the magnitude of the current induced in the rotor depends on the magnitude of the air-gap flux wave and the slip. Hence for any given slip, if the air-gap flux wave is reduced to 86.4% the induced current will also be reduced to 86.4%.
(d) The torque is proportional to the induced rotor current and the air-gap flux, both of which have reduced to 86.4% of their original values. The torque is therefore reduced by a factor of $0.864^2 = 0.746$, or 74.6% of its original value.

This question illustrates the sensitivity of torque to supply voltage: because the torque is proportional to the square of the applied voltage, a modest voltage reduction produces a much more significant drop in torque.

(4) (a) A cage rotor reacts to all pole-numbers, so no modification is required.
(b) A 6-pole wound rotor reacts only to a 6-pole air-gap flux wave, so if it was placed in a 4-pole stator, there would be zero resultant e.m.f. in each rotor phase, zero induced current and no torque. (It is important to acknowledge that there will be induced e.m.f.'s in the individual coils of the rotor winding, but because of the progressive phase-shift between them the resultant e.m.f. in the complete phase-winding will be zero.) The 4-pole winding would have to be replaced by a 6-pole one in the same slots.

(5) It is almost always desirable that the magnetic circuit is fully utilised, which means that the magnitude of the air-gap flux wave should be kept at its rated value. The magnitude of the flux depends on the voltage/frequency ratio, which should therefore be kept constant.

Hence the optimum voltage for 50 Hz is given by

$$\frac{V_{50}}{50} = \frac{V_{60}}{60} = \frac{440}{60}, \text{ i.e. } V_{50} = 440 \times \frac{50}{60} = 367\,\text{V}.$$

(6) Consider no-load operation, where the induced e.m.f. in the stator phase-winding is virtually equal to the applied voltage, i.e. in the original machine the e.m.f. induced in the winding with 15 turns per coil was 220 V.

In the rewound machine, we want the same air-gap flux wave to induce 440 V, so we will need twice as many (i.e. 30 turns) in each coil. However, for the same power, the current of the 440 V machine will be only half of the current in the 220 V machine, so the wire can be thinner. Assuming that we work the copper at the same current density, the new wire will need only half the cross-sectional are of the original, so the new diameter is $1/\sqrt{2} = 0.71$ mm. The total cross-sectional area of the new 30-turn coil is therefore the same as the original 15-turn coil, so it should fit in the same slot.

(7) The load torque is constant, so the motor will run at a speed such that its torque is equal and opposite to the load torque.

We must make the reasonable assumption that the motor is operating with a small slip in which case the induced rotor current is proportional to the magnitude of the air-gap flux wave and to the slip. The torque is proportional to the product of the induced rotor current and the air-gap flux.

The air-gap flux is proportional to the applied stator voltage, so if the voltage is reduced by a factor of 0.95, the flux will reduce to 0.95 of its original value. To produce the same torque the current will therefore have to increase by a factor of 1/0.95 or 105.3%.

The induced current in the rotor is proportional to the air-gap flux and to the slip. We have discovered that in order to produce the same torque when we reduce the voltage, we need the rotor current to increase by a factor of 1.053. If the flux had remained the same, this would have called for an increase of slip by a factor of 1.053. But the flux is now only 0.95 of what it was, so the slip has to increase yet more, by a factor of $1.053 \times 1/0.95 = 1.108$. The new slip is therefore $2 \times 1.108 = 2.22\%$.

(8) The answer to this question is given in Section 5.3.5 in the book. Torque is produced by interaction of the air-gap flux wave and the rotor induced current wave. At low slips, these two waves are in space-phase with one another, i.e. the peak rotor current occurs at the same point on the rotor surface as the peak flux density. This is the optimum set-up for torque production, and if the two waves move out of space phase the torque reduces, reaching zero when the waves are out of space phase by 90°.

As the slip frequency increases, the rotor leakage inductance ($X_r = \omega L_r$), which is negligible at low slips, begins to be the dominant impedance, and the rotor current therefore falls more and more out of time phase with the induced rotor e.m.f. as the slip increases. This is reflected in the rotor current wave moving out of space phase with the flux wave, to such an extent that although the amplitude of the current wave continues to increase with slip, the torque actually reduces because the extra phase-lag more than cancels out the influence of the increasing current.

(9) Flux patterns for 2-pole and 6-pole machines with identical rotor diameters are shown in Fig. Q9. If we assume that the maximum radial air-gap flux density (at the centre of each pole) is the same, it should be clear that the total flux crossing the air-gap under each pole is three times greater in the 2-pole machine than in the 6-pole machine. After crossing the air-gap from the rotor, the flux splits, with half going clockwise around the stator core and the other half going anticlockwise. The sketches show that the peak circumferential flux density in the stator iron core occurs at a point mid-way between the poles. The radial depth of the stator must be sufficient to carry this flux without the iron saturating, so given that the 2-pole stator has three times as much flux to carry it clearly has to be much deeper than the 6-pole one. (The ratio of depths is not 3:1 because of the slots, which have not been included in the sketches.)

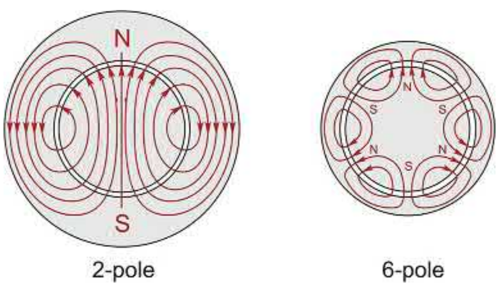

FIG. Q9

(10) The m.m.f. waveforms are shown in Fig. Q10, and are based on the assumption that all coils consist of N turns and carry current I.

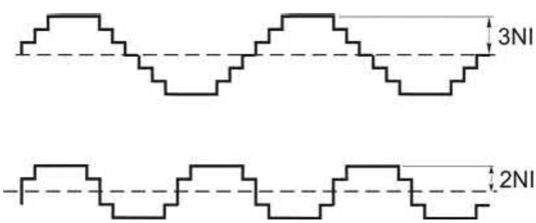

FIG. Q10

Clearly the amplitude of the 6-pole wave is only two-thirds of the amplitude of the 4-pole version, so if—as is usually the case—we want the amplitude of the air-gap flux density wave to be the same, the magnetising current (I) in the 6-pole case will have to be 50% larger than in the 4-pole case.

This example illustrates that the higher the pole-number the higher the magnetising current. We know that the magnetising current lags the voltage by 90°, whereas the 'power' component is in phase with the applied voltage. Hence for machines of a given power and voltage, the higher the pole-number the lower the power-factor.

Chapter 6

(1) (a) Any supply system can be represented by an equivalent circuit consisting of an ideal voltage source V_S in series with the supply system impedance, Z_S, as shown in Fig. 6.1. The supply impedance is usually predominantly inductive. The terminal (system) voltage, V, is generally less than V_S because of the volt-drop across the supply impedance, which is predominantly inductive. The volt-drop increases with current, but for a given current, the volt-drop is greatest when the load is inductive.

A large induction motor at rest has a very low, and predominantly inductive, impedance. When connected to the supply the current drawn by the motor will be several times the full-load current. And because the motor and supply impedances are both inductive, the fall in voltage when the load is applied will be much greater than if the load was resistive.

Other customers on the same system will experience a dip in voltage until the motor speed rises and the current it draws reduces and moves more into phase with the system voltage.

(b) If the supply system impedance is relatively high (a weak system), the volt-drop when the motor is started direct-on-line may be unacceptable to other consumers, or in extreme cases the voltage may fall so much that the motor has insufficient torque to start and/or accelerate to its normal speed. The same motor may however be started quite happily on a low-impedance (stiff) supply, where there is little or no dip in voltage even when a very large current is drawn.

(c) See the answers to questions 1 and 2.

(d) The torque developed by an induction motor at any speed is proportional to the square of the applied voltage. On a weak supply, for the reasons given in the answer to part (a), the voltage during starting will be less than it would be with a stiff supply, so the torque at all speeds will be less and the motor will therefore accelerate less rapidly and take longer to run up to speed.

(2) First, some background on 3-phase. The line-to-line voltage (referred to as 'line' voltage, V_L) in a three-phase system is the voltage between any pair of lines. The magnitudes of all three line-line voltages are the same, but they differ in phase by 120°.

When the three windings of the motor are connected in delta (see Fig. 5.2) the voltage across each phase is the relevant line voltage. When the load (i.e. the motor) phases are balanced, the currents in all three phases have the same magnitude, but differ in phase by 120°, and as a result the magnitude of the current in the supply line is $\sqrt{3}$ times the current in each phase. To sum up, for delta connection:

$$V_L = V_{ph}$$
$$I_L = \sqrt{3} I_{ph}$$

When the three motor windings are connected in star (see Fig. 5.2), then provided that (a) the three windings are balanced, or (b) the star point is connected to the neutral of the supply, the voltage across each phase is given by $\frac{V_L}{\sqrt{3}}$, and the line current is clearly the same as the phase current. To sum up, for star connection:

$$V_L = \sqrt{3} V_{ph}$$
$$I_L = I_{ph}$$

Turning now to the question, and thinking first about the line current, suppose that the impedance of each phase of the motor is Z.

When the motor phases are connected in delta to the supply, the current in each phase is given by $I_{ph} = \frac{V_L}{Z}$, and the current in each line is therefore given by $I_L = \sqrt{3} I_{ph} = \frac{\sqrt{3} V_L}{Z}$ (a).

When the motor phases are connected in star to the supply, the current in each phase is given by $I_{ph} = \frac{V_{ph}}{Z} = \frac{V_L}{\sqrt{3} Z}$, and the current in each line is therefore given by $I_L = I_{ph} = \frac{V_L}{\sqrt{3} Z}$ (b).

Comparing expressions (a) and (b) shows that the line current when the motor windings are connected in star is 1/3 of the line current when the motor windings are connected in delta.

As far as torque is concerned, we know that torque is proportional to the square of the applied voltage across each phase. The phase voltage in star is $\frac{1}{\sqrt{3}}$ times the line voltage, so the torque in star is one-third of the torque in delta.

The torque per line ampere is thus the same regardless of whether the motor is connected in star of in delta.

(3) The motionally induced e.m.f.'s in the rotor bars are directed axially, and currents flow along the low-resistance copper rotor bars, the circuit being completed via the circumferential path provided by the copper end-rings (see Fig. 5.11). Axial current flow in the iron core is prevented because the core is made from a stack of laminations that are insulated from one another. It is however possible for circumferential currents to flow in the laminations, but the currents will be small because of the relatively high resistance of the core material.

(4) It is usually possible to examine the end-windings to deduce the pole-number, provided that they are not completely obscured by insulating tapes. For example, in the most common (double-layer) winding the pitch (in slots) can usually be estimated by tracing the path of a top coil side from where it leaves the end of the stator core to the point at which it enters the bottom of a slot. For example if in a 48-slot stator the coil appears to span 8 or 9 slots, it is almost certain that the winding is 4-pole (full-pitch = 12 slots) with short-pitched coils of 2/3 or 3/4 pitch. On the other hand if the coil pitch was say 18 slots it would clearly be 2-pole (full-pitch = 24 slots), or if the pitch was 6 slots the winding would be 6-pole (full pitch = 8 slots).

(5) For this question we make use of the expression for the synchronous speed, N_S in terms of the pole-number of the machine (p) and the supply frequency (f) i.e. $N_S = \frac{120f}{p}$.

(a) If we choose a 2-pole motor the synchronous speed is $\frac{120 \times 60}{2} = 3600\,\text{rev/min}$. Allowing for a modest slip of say 4% the running speed will be about 3450 rev/min, which is fine.

(b) At 50 Hz the synchronous speed of an 8-pole motor is 750 rev/min, so allowing for modest slip the running speed will be about 700 rev/min.

(c) The pole-number must be an even integer, so the lowest is 2 and therefore the highest synchronous speed with a 60 Hz supply is 3600 rev/min.

If we want to reach 8000 rev/min with a 2-pole motor we would need to feed the motor via an inverter that could provide a frequency of $\frac{8000}{3600} \times 60 = 133\,\text{Hz}$. It would be unwise to run a standard 60 Hz motor at this speed without checking with the manufacturer that it is safe mechanically. Electrically, the voltage would have to be increased in proportion to the frequency if full torque is required: this will result in high iron losses and increased stress on the insulation, so again these aspects should first be checked.

(6) The synchronous speed is 1800 rev/min and the full-load slip is $\frac{1800-1700}{1800} = 0.056$.

The efficiency of the rotor is given by (1 − slip) × 100%, i.e. 94.4%, i.e. the rotor losses amount to 5.6%. The overall efficiency must be less that the rotor efficiency, and typically the stator losses are of similar magnitude to the rotor losses, which would suggest in this case an overall efficiency of around 89%. It is certainly highly unlikely that the stator losses are as low as 0.4%, which would be necessary to achieve an overall efficiency of 94%.

(7)

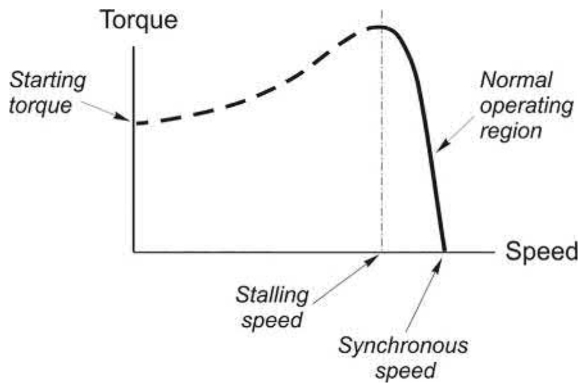

FIG. Q7

(8) The full-load slip is $\frac{1800-1740}{1800} = 0.033$.

(a) For low values of slip, the torque is proportional to slip, so at half rated torque, the slip is 0.033/2 = 0.0167 and the speed is 1800 (1 − 0.0167) = 1770 rev/min.

(b) For low values of slip, torque is proportional to slip and to the square of the applied voltage. With rated voltage, the slip for rated torque is 0.033. Hence if the voltage is reduced by a factor of 0.85, the slip must increase by a factor of $\left(\frac{1}{0.85}\right)^2 = 1.384$, so the new slip is 1.384 × 0.0333 = 0.046, corresponding to a running speed of (1 − 0.046)1800 = 1717 rev/min.

When the voltage is reduced by a factor of 0.85, so is the magnitude of the rotating flux wave. In order to develop full rated torque, the induced current in the rotor must increase by a factor of 1/0.85 = 1.176 in order to compensate for the reduced flux. This means that the rotor current is 11.8% higher than its rated value, so the rotor copper loss will be increased by a factor of $(1.18)^2$ i.e. 1.38. A 38% increase in rotor losses will cause overheating of the rotor.

(9) The stator and rotor currents are very large when the slip is large, as it is during most of the run-up period. Consequently the I^2R losses are high and a substantial amount of heat is released in the windings each time the motor runs up to speed. Repetitive starting therefore runs the risk of overheating the motor, particularly if it is coupled to a high-inertia load.

(10) The space harmonics are the unwanted by-products that are produced because the stator windings of real machines are at best an approximation to the ideal of a sinusoidally-distributed winding. When we refer to the fifth harmonic of say a 4-pole field, we mean the unwanted 20-pole field produced when we aim to produce a pure 4-pole field, and so on.

Let us suppose that the fundamental pole-number is p, and that the supply frequency is f. The speed of rotation of the fundamental field is given (in engineering units of rev/min) by the familiar formula $N_1 = \frac{120f}{p}$.

We are told that the fifth harmonic rotates backwards at one-fifth of the speed of the fundamental, i.e. the speed of the fifth harmonic is given by $N_5 = \frac{120f}{5p}$. To find the frequency induced by the fifth harmonic flux wave we can use the familiar formula again, to yield $f_5 = \frac{N_5 \times 5p}{120} = \frac{\frac{120f}{5p} \times 5p}{120} = f$. We see that the fifth harmonic reacts by inducing a fundamental-frequency e.m.f. in the stator, as indeed do all the space harmonics.

This is a result which we could have anticipated by noting that although the fifth harmonic flux wave has five times as many poles as the fundamental flux wave, it rotates at only a fifth of the speed. All points on the stator see one complete cycle of the fifth harmonic flux in the same time that they see a complete cycle of the fundamental, so both waves induce the same frequency.

Chapter 7

(1) A speed of 400 rev/min at 30 Hz would imply that the motor had 9 poles(!): but of course the pole number must be even so we could try 10 poles, for which the speed at 30 Hz is 360 rev/min. At 75 Hz, the synchronous speed of a 10-pole motor is 900 Hz, so the 10-pole will cover the speed range comfortably.

(2) The key here is that when the supply frequency is reduced from 50 Hz, the voltage is also reduced so that the magnitude of the air-gap flux density wave remains constant at its rated value. Under these conditions full torque is developed when the slip speed of the rotor (i.e. its speed relative to the synchronous travelling field) is the same as under rated conditions.

At 50 Hz, the synchronous speed is 3000 rev/min and the rated speed is 2960 rev/min, so the slip speed at full torque is 40 rev/min. At 30 Hz the synchronous speed is 1800 rev/min, and the slip speed for rated torque is again 40 rev/min, so the rotor speed is 1760 rev/min.

Similarly when the frequency is 3 Hz the synchronous speed is 180 rev/min, so with a slip speed of 40 rev/min the rotor speed is 140 rev/min.

(3) Because conditions on the rotor are the same as at 50 Hz, the rotor current is the same, i.e. 150 A. The rotor frequency is given by sf, where s is the slip and f is the supply frequency. The slip is given by: $\text{slip} = \dfrac{\text{slip speed}}{\text{synchronous speed}}$. Hence the slips at 50 Hz, 30 Hz and 3 Hz are 0.01333, 0.02222 and 0.22222 respectively, yielding rotor frequencies of 0.666 Hz in each case.

Since the rotor current at rated torque is independent of the supply frequency, it follows that the corresponding referred (stator) current is also the same. In addition, the magnetising current will be the same because the flux is the same, so the total stator current will also be the same in all three cases, i.e. 60 A.

We note that the question says estimate, and this is just as well because we have chosen to ignore the iron losses, which will not be independent of the supply frequency, and which will have an in-phase component of current associated with the loss. Because we are not told anything about how the loss varies with frequency, and because the loss component of the input current will be much smaller than the load component, we can legitimately ignore it in arriving at an estimate.

(4) To a first approximation, the magnitude of the flux-density wave in an induction motor is directly proportional to the stator voltage and inversely proportional to the frequency. When the frequency is reduced to lower the running speed, it is generally desirable to maintain the flux at its full value in order to exploit the magnetic circuit to the full, and maximise the torque per ampere of rotor current. So in order to keep the flux constant, the voltage must be reduced in proportion to the frequency, i.e. the ratio V/f is kept constant.

The approximation that flux is proportional to V/f is valid provided that the volt-drop due to the stator resistance is negligible in comparison with the applied stator voltage. This is true at high frequencies, where the stator voltage is high. But at low frequencies (e.g. a few Hz with a 50 Hz or 60 Hz motor), it turns out that the stator resistance drop is comparable with the applied voltage, and therefore merely keeping the ratio V/f constant causes the flux density to be below its full value. In order to bring the flux up to rated value, extra voltage (i.e. greater than that given by the ratio V/f) is called for, and this is referred to as 'low-speed voltage boost'.

(5) Except at very low frequencies, the magnitude of the flux density wave in an induction motor is directly proportional to the applied voltage and inversely proportional to the frequency. When the frequency is doubled, but the voltage remains the same, the magnitude of the flux density wave will be halved.

We are told that the load torque is constant, and that with the original flux, the slip speed at which the motor torque equals the load torque is $0.05N_S$.

The new flux is only half, so given that the induced current in the rotor is proportional to the flux and to the slip speed, we would have to double the original slip speed in order to induce the same current. However, we need twice as much current to develop the same torque because the flux is only half, so in fact the new slip speed will have to be four times the original slip speed, i.e. $0.2N_S$.

The new synchronous speed is $2N_S$, so the new slip is $0.2N_S/2N_S = 0.1$ or 10%.

(6) We will have to be prepared to make some 'guesstimates' to hazard an answer to this one, and not expect to place much confidence in the answers, although the main message (that at full torque, efficiency reduces with speed) should be clear enough.

If we denote the output power at base speed as 1 p.u., we know that the input must be $1/0.8 = 1.25$ p.u. so the losses at base speed are 0.25 p.u. But we have no data as to the breakdown of losses, so we will have to make some guesses.

We will begin by considering operation at full torque but half speed, and make the sensible assumption that in line with normal practice the voltage is reduced with the frequency when the speed is lowered, in order to maintain constant air-gap flux. The load torque is constant at all speeds, so, because the flux is constant, the induced rotor current will be the same and so will the referred or load component of the stator current. The magnetising component of the stator current will also be the same, so the total stator current will be the same at all three speeds. Conditions in the motor will therefore be essentially the same as at base speed and full torque, the only real difference being that the synchronous speed of the field is halved. Hence all the copper losses will be the same as at base speed, and the rotor iron losses will be the same because the slip frequency in the rotor is unchanged.

The stator iron losses will be less because the supply frequency has halved, and the windage and friction losses will have reduced because the speed has halved.

If we invoke the rule-of-thumb that says that in a well-designed motor at full load the load-dependent losses (i.e. the copper losses) are equal to the other losses (iron loss, windage and friction and stray loss), we might expect the copper loss at full torque to be 0.125 p.u. Then if we take the iron loss to be proportional to the square of the frequency, and the windage to be proportional to the square of the speed, the other losses will become 0.125/4 = 0.031 p.u. at half speed. The total loss is then 0.125 + 0.031 = 0.156 p.u. The output power is 0.5 p.u. and the efficiency is therefore 0.5/0.656 = 0.76 or 76%.

At 10% of base speed and full torque the power output will only be 0.1 p.u. The total copper loss remains 0.125 p.u. and the other losses will be very small and so we can neglect them and estimate the efficiency as 0.1/0.225 = 0.44 or 44%.

Although we cannot place too much trust in these estimates, this example nevertheless underlines the rapid fall-off in efficiency at low speeds when driving a constant-torque load.

(7) A standard totally-enclosed induction motor is designed so that at full load and base speed the shaft-driven fan provides sufficient cooling to prevent an excessive temperature rise. At low speeds, the fan is much less effective and therefore if the motor torque is high (and the stator and rotor copper losses are at their full values) there is a danger that the permissible temperature rise of the winding insulation will be exceeded.

(8) Taking the last part first, a high supply impedance means that when current is drawn from the supply, the consequent drop in voltage at the supply terminals is not negligible, and may prove unacceptable to other consumers sharing the same supply, or exceed the statutory limits imposed by the supply authority.

Starting an induction motor by direct connection to the utility supply results in a heavy current being drawn because the impedance of the motor is very low when the slip is 1, but unfortunately, the corresponding torque may be modest, particularly if the motor has a low-resistance rotor, so the motor may take a long time to accelerate up to speed, or may even not develop enough torque to start the load. A typical high-efficiency (low resistance) motor may draw perhaps four times full-load current yet produce a starting torque less than its full-load torque.

The starting torque is low despite the high current in the rotor because the rotor current wave is out of phase (in space) with the air-gap flux wave. This happens because, when the rotor is stationary, the rotor frequency is equal to the supply frequency, and therefore the rotor inductive leakage reactance is large, which makes the rotor current wave lag the voltage wave by almost 90°. (In contrast, under normal operating conditions, i.e. at low slip frequency, the leakage reactance is negligible and the rotor current wave is in phase with the flux/voltage wave.)

In complete contrast, when the same low-resistance motor is started from an inverter, the initial stator frequency will be much less than the utility frequency, so that conditions on the rotor are ideal in terms of optimum torque production. The flux, stator and rotor currents will all be at their rated values, so full-load torque will be produced. Depending on the arrangement of the inverter, the current drawn from the supply by the inverter will be the full-load current, so there will be no overloading of the supply but the torque will probably be greater than for a direct-on-line start.

(9) When an induction motor is supplied with a non-sinusoidal periodic waveform (such as that from an inverter) we know that the behaviour of the motor is governed primarily by the fundamental-frequency component of the applied voltage. Under normal operating conditions the motor runs with a small slip with respect to the fundamental rotating field and presents a predominantly resistive impedance at the fundamental supply frequency.

As far as the higher-frequency harmonics in the applied voltage waveform are concerned, however, the slip is very high and the effective motor impedance becomes dominated by the reactances arising from the leakage and magnetising inductances: these reactances are proportional to frequency, and therefore the amplitude of the harmonic current is inversely proportional to the harmonic order. A voltage waveform with a high harmonic content thus gives rise to a current waveform with a much lower harmonic content.

In more straightforward language, the inductive nature of the motor at higher frequencies causes the current waveform to be much smoother than the voltage waveform, as shown for example in Fig. 7.1.

Chapter 8

(1) Torque is directly proportional to rotor current when the rotor flux linkage is constant (as in vector control) so the new rotor current will be 15 A.

(2) The vertical axis is maximum possible torque and the horizontal axis is speed. Full torque is available at all speeds up to the base speed, at which the inverter output voltage is at its maximum. At higher speeds the V/f ratio reduces with frequency, thereby reducing the flux. In this field-weakening region the maximum available torque therefore reduces inversely with frequency.

(3) The d-axis current produces the rotor flux, which should be kept constant. The q-axis represents the torque component of stator current, so if the load torque doubles, so will the torque component.

(4) The rotor flux linkage has to be established, which takes time because of the inherent inductance.

(5) The acceleration is constant, so because the load is inertial it follows that during acceleration or deceleration the magnitude of the motor torque is constant; this requires the slip frequency to be maintained at 2 Hz.

Point a: the rotor speed is half of the full speed, so if the rotor was running in the steady state at this speed it would not require any torque because the load is inertial, so the steady state frequency would be 20 Hz. But because it is accelerating, and we are told that the slip frequency is 2 Hz, the instantaneous stator frequency will be 22 Hz.

Point b: negative torque is required for deceleration, so the rotor will be running faster than the synchronous speed, thereby producing negative slip and torque: the frequency at point b is thus 18 Hz.

Point c: the rotor is instantaneously at rest (i.e. zero speed) at point c, but the torque must continue to be negative so the stator frequency will be −2 Hz at point c, i.e. the phase-sequence will be reversed. (In fact the phase sequence will have to reverse at the point where the rotor speed corresponds to a synchronous speed of 2 Hz, at which point the instantaneous frequency passes through zero.)

Point d: similar to point a, i.e. frequency 22 Hz, negative phase sequence.

Point e: similar to point b, i.e. frequency 18 Hz, negative phase sequence.

Point f: similar to point c, i.e. frequency 2 Hz, positive phase sequence.

If the motor was at rest because the speed demand was zero, any attempt to cause acceleration would be result in a demand for an opposing torque, so the person grasping the shaft would feel substantial resistance to movement.

(6) (a) The motor torque is directly proportional to the torque component of stator current, so when the torque is halved, the torque component is also halved. The flux component will be kept constant. By reference to Fig. 8.9, the original angle φ_3 was 70°, so by simple trigonometry, the new angle is 53.9°.

(b) The quantities in Fig. 8.9 are not dependent on speed, so the angle will remain at 70°.

(7) The user would not normally be able to tell any difference.

Chapter 9

(1) We should always aim to keep the working flux at its designed value. In synchronous or induction motors above 1 kW connected to a utility supply we can ignore the effect of resistance as far as flux is concerned, and safely assume that the ratio of the applied voltage to the frequency (V/f) determines the flux. For the motor in question, V/f = 420/60 = 7. Hence when we want to operate at 50 Hz, the voltage should be $7 \times 50 = 350$ V. Note that if we did not reduce the voltage, the flux would increase by 20%, causing the magnetic circuit to become highly saturated, with a substantial increase in magnetising current, stator copper loss, and iron loss.

(2) With no shaft projecting at either end the set-up is clearly not intended to provide any mechanical output power, so we can deduce that one of the machines is intended to run as a motor and drive the other as a generator. The arrangement is in fact a bi-directional frequency changer, for allowing 50 Hz apparatus to be supplied from a 60 Hz system, or vice-versa.

The 10-pole machine has a synchronous speed of 600 rev/min at 50 Hz, while the 12-pole machine also has synchronous speed of 600 rev/min when it is supplied at 60 Hz. Both machines can work as motor of generator, so power can be transferred in either direction between the 50 and 60 Hz systems.

The machines could be separated and used as synchronous motors. With a 50 Hz supply, the speeds available are 500 and 600 rev/min, while these increase to 600 and 720 rev/min respectively with a 60 Hz supply.

(3) In a synchronous machine operating in the steady state the magnetic field produced by the stator windings rotates at the same speed as the rotor, so that from the point of view of an observer on the rotor there is no 'flux cutting' and therefore no motional e.m.f. The only time that there will be motional e.m.f. in the rotor is during run-up, or when there is a transient load change. This is important in only a limited number of cases, for example under fault conditions in very large synchronous machines connected to the power system.

(4) The first thing that we should note is that because there is no mechanical load on the shaft of the motor, the only mechanical power being produced is that associated with any frictional or windage losses, which must be very small. So we should not be surprised to be told that the input power is always low: after all, throughout the book we have stressed that the input power is determined by the load on the shaft.

To go further we must do as suggested and make use of the equivalent circuit (Fig. 9.14) and phasor diagram. We are told that the machine is a large one, so we can make the assumption that the stator resistance is negligible without having much effect on the result.

The phasor diagram is shown in Fig. Q6. This is based on Fig. 9.16, but with the load torque reduced to zero and the losses neglected, in which case the input power becomes zero. The locus of the stator current as the rotor excitation is varied is shown by the dotted line.

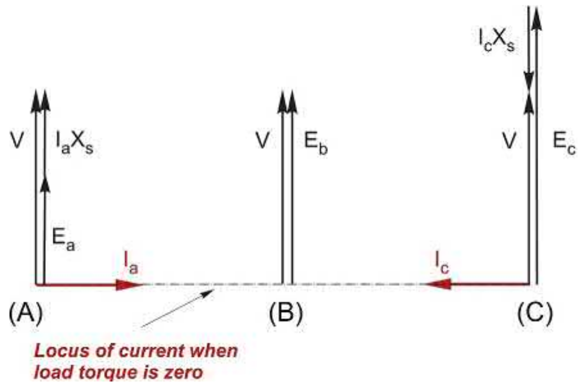

FIG. Q4

Because the power is zero (or at least very small), the real or in-phase component of the stator current is always almost zero: it has been taken as zero in Fig. Q4 for the sake of simplicity.

However, the stator current (I) must be such as to satisfy Kirchoff's law (that the applied voltage V must equal the sum of the induced e.m.f. E and the volt-drop across the synchronous reactance. Hence we see that when E is small (Fig. Q4A), the stator current is large and lagging the voltage by 90°, so that the machine looks like an inductance when viewed from the supply.

When the rotor excitation (E) is increased the stator current reduces (i.e. the apparent inductance increases) until, when E = V, the stator current is zero. Further increase in E causes the current to increase again, but his time it leads the voltage by 90°, and the machine therefore looks like a capacitor when viewed from the system. (We explained in the text that synchronous machines operating in this mode (i.e. without any mechanical power output) were widely used at one time as 'synchronous compensators' in power systems, allowing the power-factor of a region to be optimised by being set to have the desired value of inductive or capacitive reactance.)

(5) This is a rather vague and open-ended question, and it would be reasonable to challenge it on the grounds that we are not given enough information. However, we can answer what the questioner probably had in mind easily, by noting that initially, the motor on the right in Fig. 9.5 is in a state of equilibrium, with the rotor flux acting on the stator conductors to produce a negative (anticlockwise) torque that is equal and opposite to the positive (clockwise) torque provided by the weight.

If the stator current is then reversed, the motor torque will become positive, and will add to the torque provided by the weight, so the rotor will accelerate in a clockwise direction.

To say anything more we would have to make assumptions, perhaps beginning by assuming that the torque angle curve of the motor is independent of its velocity, so that the motor torque depends only on its position, and also that friction is negligible. In that case, the average motor torque for every complete revolution will be zero, whereas—as long as the rope remains on the drum—the load torque will always be clockwise. It follows that the average acceleration will always be positive, so the mean speed will increase until the rope finally leaves the drum, at which stage the kinetic energy of the rotor is equal to the potential energy lost be the falling weight.

Once the weight has dropped off, if friction is negligible the rotor will continue with the final mean speed being harmonically modulated by the motor toque. In a real motor, friction torque will produce deceleration, and the rotor will eventually end up with a pendulum-like damped oscillation about one of the stable zero-torque equilibrium positions.

(6) The stator winding will produce a 4-pole m.m.f. and flux distribution (N-S-N-S), and the rotor has four saliencies, so there are four equally spaced positions at which the rotor poles can align with the stator field.

(7) The expression for torque for a given stator current is given in Section 9.3.3, i.e.

$$T \; \alpha \; I_s^2 \left(L_d - L_q\right) \sin 2\gamma.$$

Assuming that the current is the same in both cases, the ratio of torques is thus given by

$$\frac{T_{new}}{T_{original}} = \frac{1-0.2}{1-0.25} = 1.067.$$

The new motor therefore produces almost just under 7% more torque.

(8) Part of a typical static torque-angle curve for a 10-pole pm motor will look like Fig. Q8, with each pole spanning 36°. Stable equilibrium positions are marked by dots, and unstable ones by a star shape.

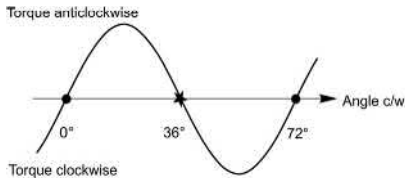

FIG. Q8

(i) When released at any angle less than 36°, the rotor experiences an anticlockwise torque, and therefore accelerates towards the stable position at 0°, overshoots and oscillates before finally coming to rest at 0°. Conversely, if it is released anywhere between 36° and 72°, it will experience a clockwise torque and thus eventually settle at 72°.

(ii) From the figure, maximum anticlockwise torque will be at ±18°, and maximum clockwise torque will be at ±54°.

(iii) If the current is reversed, the torque-angle curve changes sign, so the stable points will now be where the unstable ones were previously. The rotor will therefore settle at ±36°.

(9) The torque in a pm motor depends on the strength of the magnet and the magnitude and position of the stator current relative to the magnet. Assuming that the distributions of magnet and stator (armature) fluxes are sinusoidal in space, we showed in Section 9.3.2 that the torque expression is

$$T \propto \left(\Phi_{mag}\right)\left(\Phi_{arm}\right)\sin\lambda,$$

where λ is the angle between the flux distributions. The magnet flux is fixed, and the armature flux is proportional to the stator current, so in order to minimise the current for a given torque, the stator current/flux should be positioned relative to the rotor so that the torque angle (λ) is 90°. If the torque angle is less than 90°, the torque could still be obtained using a larger current; with $\lambda = 30°$, for example, twice as much current would be required to produce the same torque.

(10) The reference to base speed means that the motor voltage has reached its maximum (rated) value, and we can therefore assume that when the motor runs faster than base speed, the voltage remains at rated value. (At speeds below base speed, we would expect the V/f ratio to be constant to keep the resultant flux at its designed value.)

The motor is unloaded so once it is up to speed we can neglect losses and assume that the real input power is zero, which means that the in-phase component of current is zero. But in this case we are told that the current is negligible, so we know that the reactive component of current is also zero. Hence by reference to Fig. 9.14, the induced e.m.f. is equal to the applied voltage, as shown on the left had side of Fig. Q10. There is no armature current, (and thus no armature flux) and the resultant flux is simply that due to the magnet, and at the base frequency in this particular motor, the flux induces the e.m.f. E that is equal to the applied voltage, V.

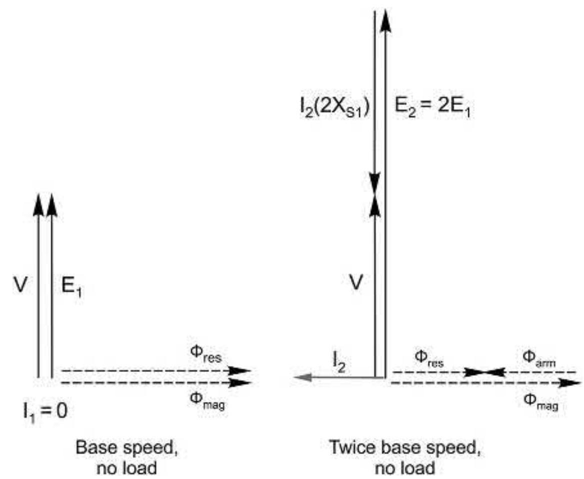

FIG. Q10

The two things that remain the same when the speed is doubled are the magnet flux and the armature voltage, and we know that the resultant flux must again induce the voltage V. But because the induced voltage is proportional to speed (i.e. frequency), we deduce that the resultant flux will be half of what it was at base speed. The armature flux must therefore oppose the magnet flux, and effectively nullify half of it, as shown in the right-hand sketch.

Turning to the time phasor diagram, we can see how the magnet flux and armature flux are reflected in the time phasors. At twice base speed, the magnet flux will induce twice as much motional e.m.f. as at base speed, so we know that in the right hand sketch, $E_2 = 2E_1$. The effect of the armature flux is represented by the voltage across the armature reactance, and we see that in this particular case the phase and magnitude of the current (I_2) have to adjust themselves so that the voltage drop ($I_2(2X_{S1})$) is antiphase to E_2 and of half the amplitude. (Note that the reactance is proportional to frequency so it doubles with the increase in speed.)

This example illustrates the less than ideal inherent behaviour of the pm motor in the 'field-weakening' region above base speed. Because the motor effectively becomes over-excited (i.e. the induced e.m.f. is much greater than the supply voltage) the flux produced by the armature current has to oppose the magnet flux: this demagnetising component of the total armature current produces no useful output power, but inevitably contributes to the stator copper loss.

Chapter 10

(1) In one-phase-on operation the equilibrium (rest) positions are where several rotor and stator teeth (which are usually of the same width) are directly aligned with each other. The alignment positions are therefore well-defined at the design/manufacture stage, and do not depend on the magnitude of current in the excited phase.

When two phases are excited simultaneously, the rotor comes to rest at a position intermediate between the two positions it would occupy if only one of the phases were excited. If the two exciting windings are identical and the currents are equal the rest position will be mid-way between the two one-phase-on positions. In practice, however, even if the windings are identical it is unlikely that the two currents will be exactly equal, so the final equilibrium position will tend to be slightly closer to the phase that has the higher current. The rest positions are thus less well-defined than during one-phase-on operation.

(2) Detent torque is the (small) torque that exists in hybrid motors when none of the phase-windings carries any current. It can easily be detected by turning the rotor slowly by hand with the motor unexcited. The rotor equilibrium (or rest) positions that result from the detent torque are the same as the normal step positions: this is an advantage if the motor is left unexcited, because the existence of the detent torque helps to prevent the rotor from losing registration by being inadvertently moved to another position.

(3) The holding torque is the maximum static torque that the motor can develop, i.e. the peak of the torque/angle curve, i.e. T_{max} in the typical torque/angle curve shown in Fig. 10.10. When a stepping motor is referred to as, say, 'a 2 Nm motor', it is understood to mean that the holding torque is 2 Nm.

(4) We are not told the step angle so clearly the answer to this question will have to be expressed as a fraction of a step. However we know that the motor is a 3-phase one, so one step corresponds to one-third of the angle between successive stable alignment positions for one phase. Hence if we use the sinusoidal approximation suggested, the static torque can be written as $T=0.8\ \sin\theta$ Nm, where $\theta=360^\circ$ corresponds to three steps and one step is represented by 120°.

With a steady load torque of 0.25 Nm the equilibrium position will therefore be given by $0.25=0.8\ \sin\theta$, i.e. $\theta=18.2^\circ$. Ideally, of course the equilibrium position should be at $\theta=0^\circ$: expressed as a fraction of step, the step position error is therefore $18.2/120=0.15$ step.

(5) When driven from an ideal constant-current source, the output torque effectively jumps instantaneously from the static torque curve of one phase onto the static torque curve of the incoming phase, the lag angle (see Section 10.5.2) adjusting such that the average motor torque is equal and opposite to the load torque. With the assumption that the static torque/angle relationship is sinusoidal, the maximum torque will be obtained when the rotor lag angle is such that the instantaneous torque follows the solid line in Fig. Q5

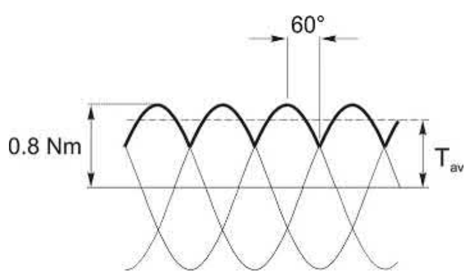

FIG. Q5

The question says estimate, so we note that the minimum torque occurs at 30° on the curve $0.8\sin\theta$, i.e. at 0.4, so by eye we could guess that the average is between 0.6 and 0.7: by integration the result is 0.66. The pull-out torque is therefore 0.66 Nm, and if the load torque exceeds this value, the motor will stall.

(6) If the torque/angle curve is assumed to be linear about the step position, the rotor will oscillate in much the same way as a simple pendulum as it comes to rest. Mathematically, if the restoring torque is given by $T = -k\theta$, where θ is the angle of displacement from the equilibrium position, the equation of motion is $-k\theta = J\frac{d\omega}{dt} = J\frac{d^2\theta}{dt^2}$, the solution of which is $\theta = A \sin \omega_n t$, where $\omega_n = \sqrt{\frac{k}{J}}$.

We can therefore calculate the angular frequency of oscillation provided we know k, the gradient of the torque/angle curve, and the total inertia, J. We must use SI units, so we first convert 2 Nm/degree to $2 \times 180/\pi = 114.6$ Nm/rad. Substituting in the expression above yields $\omega_n^2 = \frac{114.6}{1.8\times10^{-3}}, \omega_n = 252.3\,\text{rad/s}$, or just over 40 Hz.

(7) The step angle is given by: Step Angle $= \dfrac{360°}{\text{Rotor teeth} \times \text{Stator phases}}$.

(a) Applying this formula the step angle of the VR motor is given by $\frac{360°}{8\times3} = 15°$.

(b) Using the same formula, the step angle is given by $\frac{360°}{50\times4} = 1.8°$.

(8) Probably the easiest thing to do is to spin the rotor with the stator windings open circuited. If there is negligible mechanical resistance apart from friction, it is almost certainly a VR type. If a cyclic detent torque is evident, it must be a hybrid or a permanent-magnet type. (We have not dealt with the pm type in this book.) To be absolutely sure, the windings can be short circuited and if when the rotor is turned the mechanical resistance has increased we can deduce that the machine is acting as a generator, the currents in the windings being due to motional e.m.f. This can only happen if the motor includes a permanent-magnet to provide the excitation: the VR type has no excitation unless it is connected to the drive circuit, so it cannot generate in isolation.

(9) When operating in the stepping mode, a hybrid motor with two phases would normally run from a bipolar supply, the rectangular one-phase-on pattern of excitation typically being $+A$, $-B$, $-A$, $+B$, $+A$, etc,i.e. one complete cycle of excitation of each phase corresponds to four steps.

When supplied from a sinusoidal source rather than a rectangular one, the torque will be smoother and rather than an incremental motion the rotor velocity is effectively smooth. In one cycle of the supply, the rotor will therefore move four steps, so at 60 Hz the effective stepping rate is 240 steps/s. The step angle is 1.8°, so the motor speed will be $240 \times 1.8° = 432°$/s or 72 rev/min. (It is interesting to note that the 1.8° hybrid motor was originally conceived to produce a low-speed utility-fed direct drive motor, offering 72 rev/min from 60 Hz or 60 rev/min from 50 Hz. The utility-fed version has more turns of thinner wire to suit the higher voltage, but is otherwise the same as the stepping version.)

(10) Perhaps it should have been pointed out to the scientist that the claim that stepping motors 'typically complete each single step in a few milliseconds' does presuppose that the motor was being used in a 'typical' application, which normally means that the load inertia is of similar order to the motor inertia.

To estimate the inertia of the pointer (which we will assume is fixed to the rotor at one end) we use the formula for the inertia of a uniformly-distributed linear mass, i.e. $\frac{ML^2}{3} = 16{,}000$ g cm^2.

We will have to guess that the rotor diameter is at most 1.5 cm, and its length perhaps 6 cm. Assuming it is made of steel with a relative density of 7.8 its mass is approximately 85 g. The rotor inertia is therefore given by $\frac{85\times(0.75)^2}{2} = 24\,\text{g}\,\text{cm}^2$.

We see that the load inertia is 660 times the rotor inertia, so in no sense is this a typical application.

The acceleration at the beginning of the step will be only 1/660 of what it is for the unloaded motor, and if this massive reduction applied throughout the whole step trajectory, the transit time would therefore be increased by a factor of $\sqrt{660}$ or approximately 26 times. In practice however, the frictional damping in VR motor is low and therefore because of the high inertia the damping factor will be very low and the pointer must be expected to oscillate several times before coming to rest.

A viscous-coupled inertia damper will make matters much better, though the load inertia is so high that it may be difficult to find an 'off the shelf' damper that is optimum.

Chapter 11

(1) When speed accuracy is specified as a percentage, it invariably means that the speed will be held to within that percentage of the base speed. In this case the base speed is 1500 rev/min, so 0.5% of base speed is 7.5 rev/min. So when the speed reference is set to 75 rev/min, the speed can be between 82.5 and 67.5 rev/min and yet the drive will still be within specification.

(2) The condition that must be satisfied in order for the acceleration of the load to be maximised is that the effective inertia of the load (as seen at the motor) must be equal to the motor inertia.

In this question the actual load inertia is nine times larger than the motor inertia, so in order to reduce this as seen by the motor we must use gearing (or a toothed belt) so that the load speed is less than the motor speed. It is explained in Section 11.4 that if the gear ratio between motor (high speed) and load (low speed) is n:1, a load of inertia J at the low-speed side appears at the motor as if it were an inertia of $\frac{J}{n^2}$. We want an inertia of 0.009 kg m^2 to appear as an inertia of 0.001 kg m^2, i.e. we need

$\frac{1}{n^2} = \frac{1}{9}$ i.e. $n = 3$. Since the motor pulley has 12 teeth, it follows that the load pulley must have 36 teeth.

(3) First, consider operation at normal full load, i.e. 1 p.u. We are told to assume that the temperature rise is exponential (loose language—it really means that the temperature is given by an expression of the form $(1 - e^{-t/T})$), so if we denote the final temperature rise by Θ_{max}, the equation governing the temperature rise as a function of time will be

$$\Theta = \Theta_{max}\left(1 - e^{\frac{-t}{\tau}}\right)$$

where τ is the thermal time-constant. This is sketched as the lower curve in Fig. Q3.

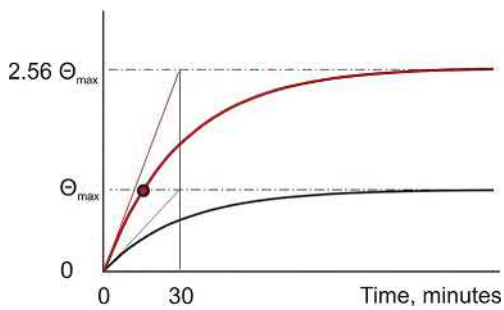

FIG. Q3

Now consider operation from cold with the motor overloaded by 60%, i.e. at 1.6 p.u. We are told that the losses are proportional to the square of the load, so the losses are now $(1.6)^2$ or 2.56 times greater than at full load. The final temperature rise is proportional to the losses, so the final temperature rise that would be reached at 60% overload is $2.56\Theta_{max}$. The equation governing the temperature rise in this case it therefore

$$\Theta = 2.56\Theta_{max}\left(1 - e^{\frac{-t}{\tau}}\right),$$

and this is sketched as the upper curve in Fig. Q3. Putting $\Theta = \Theta_{max}$ and solving this equation gives $t = 14.9$ min as the time at which the temperature rise will reach the allowable limit.

(4) The scant information provided can at best allow us to suggest some possibilities. The specification will have to be much more detailed before we can home in on the best solution. But there is information that provides useful guidance, particularly as a direct drive (i.e. no gearbox) is specified. Full torque is required at all speeds, and since we know that the maximum power is about 1200 W at 10,000 rev/min, the torque is a modest 1.1 Nm.

The torque of a machine will usually give an idea of its volume, but if we have little experience we can still get some idea by noting that at the more commonplace speed of say 2000 rev/min, the output power will only be 240 W. So we might expect the size of the motor to be comparable to say a medium-range hand-held power drill.

The top speed of 10,000 rev/min is on the high side, and immediately limits the choice of motor/drive. Conventional (brushed) d.c. machines are seldom expected to run at speeds as high as this because of difficulties with commutation. A standard 2-pole induction motor could not be assumed to run happily at 10,000 rev/min when its bearings and rotor design were intended for a maximum of 3600 rev/min (on a 60 Hz supply). An off the shelf induction motor aimed at the inverter-fed market would also probably have an upper limit of something like twice the utility-frequency speed, so a special would be required to allow speeds of 10,000 rev/min.

In short, we should not expect to be able to do this job with a run-of-the-mill motor, and instead we must expect to move into the specialist area and face the inevitable higher cost. The motor manufacturer/distributer should be consulted and they would be able to offer either a specially designed induction motor or PM synchronous motor. In either case the power electronic controller would be similar, or a common design capable of controlling either type of motor. The operating frequency would not be a limiting factor.

Further matters that will have a bearing on the selection include the space available (a short fat motor or a long thin one?); the required accuracy of speed holding; the dynamic performance (i.e. what is the required bandwidth of the speed control loop); the stall protection requirements; the operating environment (will the motor have to withstand the presence of cutting fluids?, what is the ambient temperature range?); the mounting and coupling arrangements (which can be surprisingly tricky for a direct drive); and maybe several more!

(5) As far as the hoist motor is concerned, the load on the hook directly determines the steady-state torque. The maximum weight that may be attached to the hook is specified on the hoist, and it is understood that this limit is not related to the hoisting speed. It follows that full motor torque must be available at all speeds including during starting. It is often a requirement to prove that the required torque is present at the motor shaft before the holding brake is released and so full torque is also required at zero speed.

When speed control using an inverter becomes available, the operator may set the drive to lift the maximum load at a slow speed, the motor then drawing full current in order to produce full torque. In this condition the motor losses will be high but because of the low speed the cooling fan will be much less effective than at full speed and prolonged operation will therefore result in the motor overheating. The remedy will be to fit a cooling blower, or an oversized motor.

(6) Two examples of controlled-speed applications for which conventional (brushed) d.c. motors are not suitable are:

(a) Any application in an environment where there is a danger of explosive gases being present, because of the danger of ignition caused by sparking of the commutator. The obvious area where this applies is in mining, where inverter-fed induction motors are inherently less dangerous (although they will also have to have flameproof enclosures). (D.C. motors in flameproof enclosures are available but they are prohibitively expensive.)

(b) Any application where it is essential to minimise periodic maintenance, such as unmanned off-shore installations. The conventional d.c. motor requires regular inspection and renewal of brushes, so again an inverter-fed induction motor is likely to be preferred.

(7) The rating of the induction motor is primarily determined by the temperature rise of the insulation material, the limit reflecting a compromise between the lifetime of the insulation and the complexity of the cooling system. After a prolonged period of operation at the full rated load, the final temperature will settle at a value that is acceptable for an indefinite period thereafter.

But because the motor is inherently capable of supplying further torque (in which case the cooling system will no longer match the increased losses), permissible patterns of short-term overload can be specified at the design stage so that, despite the temperature rising above the steady-state full-load value, no serious harm will be done to the insulation. Fortunately, the high thermal capacity of the copper and iron result in a thermal time-constant of anything from several seconds in a small motor to many minutes in a large one. Hence overloads that last for only a fraction of the time-constant can be tolerated.

The thermal issues in the inverter are similar, but with very different physical dimensions and time-scales. The vulnerable areas in any power electronic converter are the semiconductors, particular the active junction regions in the switching devices and diodes. We have seen that the on-state voltage drops are low (typically around 1 V), but the internal power dissipation can be tens or hundreds of Watts, depending on the current. This loss is produced in a tiny volume of semiconductor with very low specific heat, so the power density is very high and adequate heat removal (typically using heat sinks) is essential in order to limit the full-load temperature rise of the active material to a safe level. However, the associated thermal time-constant is very much shorter than that of the motor, so that if the current suddenly increased to provide overload power to the motor, the semiconductors would self-destruct in perhaps a few milliseconds.

The importance of discussing the potential application with the drive supplier should be clear, and the user who expects his N kW induction motor to deliver twice full load (even for short periods) should not be surprised to have to purchase an inverter rated at 2 N kW.

(8) This question relates to the rating of motors performing an intermittent cyclic duty, which was discussed in Section 11.5.2. The general approach is to rate the motor according to the r.m.s. of its power cycle, on the assumption that the losses (and therefore the temperature rise) vary with the square of the load.

We begin by estimating the r.m.s. output power. For the 1 min that the motor runs on load the torque is 60 Nm and the speed 1400 rev/min, so the power output is given by $P_{out} = 60 \times 1400 \times \frac{2\pi}{60} = 8{,}800\,\text{W}$ or 8.8 kW. For the following 5 min the motor runs light so the output power is zero. The power-squared plot is therefore as shown in Fig. Q8.

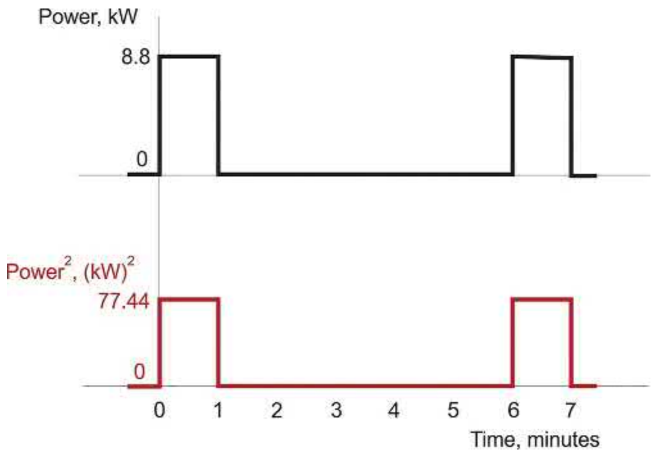

FIG. Q8

The mean of the power squared plot is given by: Mean of $(\text{Power})^2 = \frac{77.44 \times 1 + 0 \times 5}{6} = 12.9\,(\text{kW})^2$, so the root mean square is 3.6 kW.

The periods of operation are certainly short in comparison with the thermal time-constant of a motor of this output, so on the basis of the continuous power rating we could choose the 4 kW, motor.

But we need to check that it would have sufficient torque during the 1-min periods i.e. 60 Nm. We are told that the pull-out torque will be 200% or 2 p.u., i.e. the pull-out torque is twice the full-load torque. We can calculate the full-load torque by dividing the output power by the running speed. Clearly, since the supply is 50 Hz, we will choose a 4-pole motor (synchronous speed = 1500 rev/min): given that the full-load slip is about 5%, we can assume that the full-load speed is 1500 × 0.95, i.e. 1425 rev/min, close enough to the desired speed. So the rated torque of the 4 kW motor is given by $T_{f-l} = \dfrac{4000}{1425 \times \dfrac{2\pi}{60}} = 26.8\,\text{Nm}$. So now we know that the pull-out torque of the 4 kW motor is $2 \times 26.8 = 53.6\,\text{Nm}$. This is insufficient to meet the requirement of 60 Nm, so a 4 kW motor will not be suitable and we must go for the next one up, i.e. 5.5 kW.

The full-load torque of the 5.5 kW motor turns out to be 36.8 Nm, so its peak torque is almost 74 Nm, which provides a good margin over the 60 Nm we are seeking to drive the pump.

The rated torque of 36.8 Nm is obtained at a speed of 1425 rev/min, i.e. a slip of 75 rev/min. We are asked to estimate the running speed at 60 Nm, so first we can assume that the torque remains proportional to slip, in which case the slip at 60 Nm will be given by $s = \frac{60}{36.9} \times 75 = 122\,\text{rev/min}$. On this basis, the running speed would be 1500–122 = 1378 rev/min: but we know that between rated torque and pull-out torque the gradient of the torque slip curve reduces (see Fig. 5.25), so in practice the slip at 60 Nm will be rather more than 122 rev/min, say 140 rev/min, which means that the speed will be 1360 rev/min, which should be close enough to meet the specification.

(9) When the speed reference is set at 100% (i.e. not a controlled ramp) and the drive is at rest, the drive will apply full torque to accelerate as fast as possible. In most drives, full torque will be maintained until the speed comes within a few percent of the target, and only then will the torque be reduced to give a smooth approach to the final speed.

The information we are given is in line with normal behaviour, because we are told that the acceleration up to 1180 rev/min is more-or-less uniform, which means that the net torque remains constant. However, we are asked to estimate the system inertia, so we need to make use of the dynamic equation, i.e. *torque = inertia × angular acceleration*. But although we can calculate the rated torque of the motor, and assume that this remains constant during run-up, we have no information about friction torque, so all we can do is to make the reasonable assumption that it is small in comparison with the motor torque, and ignore it.

The rated torque of the motor is given by $T = \frac{50 \times 10^3}{1200 \times \frac{2\pi}{60}} = 398\,\text{Nm}$.

The angular acceleration is given by $\frac{d\omega}{dt} = \frac{(1180 - 0) \times \frac{2\pi}{60}}{4}$ $= 30.9\,\text{rad/s}^2$.

Hence the inertia is given by $J = \frac{398}{30.9} = 12.9\,\text{kgm}^2$.

To find the stored kinetic energy at full speed we use the expression $E = \frac{1}{2} J\omega^2$, i.e. $E = \frac{1}{2} \times 12.9 \times \left(\frac{1200 \times 2\pi}{60}\right)^2 = 101.85\,\text{kJ}$.

We could find the energy supplied in the first 4 s by repeating the calculation above with a speed of 1180 rev/min, but the point of this part of the question is to encourage us to recognise that when acceleration takes place at constant torque, the speed increases linearly with time and so therefore does the output power of the motor (because power = torque times speed, and torque is constant). The average power in this case is therefore half of the power at 4 s, which in turn is $\frac{1180}{1200} \times 50 = 49.17\,\text{kW}$. So the average power is $0.5 \times 49.17 = 24.58\,\text{kW}$. The energy supplied is thus given by the average power times the time, i.e. $24.58\,\text{kW} \times 4\,\text{s} = 98.33\,\text{kJ}$.

The extra 3.5 kJ represents the additional kinetic energy gained during the final part of the run-up from 1180 to 1200 rev/min. During this time the speed controller will not be saturated, and the torque will taper-off as the final speed is approached.

(10) When the motor (see question 9) is switched off completely, the only torque will be that due to friction, and because we know the total inertia we can estimate the friction torque if we know the deceleration.

The speed falls from 1200 to 1080 rev/min in 20 s so the angular deceleration is given by $\frac{(1200 - 1080) \times \frac{2\pi}{60}}{20} = 0.63\,\text{rad/s}^2$. We know that the corresponding acceleration under full motor torque is $30.9\,\text{rad/s}^2$, so we can express the magnitude of the friction torque as a fraction of the full-load torque as $0.63/30.9 = 0.02$. In other words the friction torque is about 2% of the full-load torque, which justifies our neglect of friction in answering question 9.

Further reading

A General motors and drives books – Delving a little deeper

- Bose, B.K. (1996) "Power Electronics and Variable Frequency Drives: Technology and Applications", Wiley-Blackwell, ISBN-10: 0780310845, ISBN-13: 978-0780310841
  Authoritative book covering many key areas of motor drive analysis, performance and design. Contains many references to key books and technical papers on the subject.

- Hindmarsh, J. and Renfrew, A. (1996) "Electrical Machines and Drives: Worked Examples" (3rd Edition), Butterworth-Heinemann, ISBN-10: 0750627247, ISBN-13: 978-0750627245 | Edition: 3
  Worked examples, which put theory in context. Good discussion of engineering implications. Numerous problems are also provided, with answers supplied!

B Control and modelling

- Vas, P. (1992) "Electrical Machines and Drives" Oxford University Press, ISBN 0-19-859378-3
  A definitive, highly analytical, book on space vector theory of electrical machines

- Vas, P. (1998) "Sensorless Vector and Direct Torque Control", Oxford University Press, ISBN: 0-19-856465-1
  Comprehensive analytical treatment of Field orientation and Direct Torque Control of a.c. machines

- Distefano, J., Schaum's Outline of Feedback and Control Systems, 2nd Edition (2013), McGraw-Hill Education; ISBN-10: 9780071829489
  This is a good introduction to feedback and control systems including lots of tutorial questions.

- Chaisson, J. (2005), "Modeling and High-Performance Control of Electrical Machines", John Wiley & Sons, IEEE Press series on power engineering, ISBN 0-471-68449-X
  A good book for anyone wanting to model an electrical machine as part of a simulation package or in a control system. Good treatment of a number of practical issues including noise on feedback signals.

C Practical aspects of design and application of motors and drives

- Drury, W. (2009), "The Control Techniques Drives and Controls Handbook", (2nd Edition), IET, London, ISBN: 978-84919-013-8
  A practical guide to the technology underlying drives and motors and consideration of many of the design issues faced with their use in many applications and environments.

- Mohan, N., Undeland, T.M. and Robbins W.P., (2002) "Power Electronics: Converters, Applications, and Design" (3rd Edition); John Wiley & Sons; ISBN-10: 0471226939, ISBN-13: 978-0471226932
  Cohesive presentation of power electronics fundamentals for applications and design. Lots of worked examples.

- Neale, M., Needham, P. and Horrell, R. (1991) "Couplings and shaft alignment", Professional Engineering Publishing Limited, ISBN 1-86058-170-6
  A practical guide to coupling selection and problems of shaft alignment

- Wu, B. (2006) "High Power Converters and AC Drives", John Wiley & Sons, ISBN-13 978-0-471-73171-9/ISBN-10 0-471-73171-4
  Good practical design book on high power a.c. drive topologies

D Reliability

- Chung, H.S., Wang, H., Blaabjerg, F. and Pecht, M. (2016), "Reliability of Power Electronic Converter Systems", IET, London, ISBN: 987-1-84919-901-8
  An excellent advanced book on all aspects of converter and drive reliability. Covers the majority of failure mechanisms and provides context for the importance in practical applications.

E Synchronous reluctance and permanent magnet motors

- Hendershot, J. R. and Miller T. J. E., (2010) "Design of Brushless Permanent-Magnet Motors" Motor Design Books, LLC, ISBN: 0984068708, 9780984068708
  A comprehensive guide to the design and performance of brushless permanent-magnet motors.

- Pellegrino, G., Jahns, T.M., Bianchi, N, Soong, W. and Cupertino, F., (2016), "The Rediscovery of Synchronous Reluctance and Ferrite Permanent Magnet Motors: Tutorial Course Notes", Springer, ISBN:-10: 3319322001
  This is a good step into better understanding these motors which are becoming more important in many applications including automotive traction.

F Energy efficient electric motors

- Jordan, H.E. (1994) Energy-Efficient Motors and their Applications (2nd Edition). New York: Springer, ISBN-10: 0306446987, ISBN-13: 978-0306446986
  Clearly written specialist text.

G Power semiconductor devices

- Lutz, J., Schlangenotto, H., Scheuermann, U. and De Donker, R., "Power Semiconductor Devices" by, Springer, 2018, ISBN-10: 331970916X, ISBN-13: 978-3319709161
  Readers who want to learn more about different power semiconductor devices should have a look at this book. It is a specialist text covering all key aspects.

Index

Note: Page numbers followed by *f* indicate figures and *t* indicate tables.

F

G

H

I

K

L

Q

R

S

T

U

V

W

FIFTH EDITION

Electric Motors and Drives

Fundamentals, Types and Applications

Austin Hughes and Bill Drury

Electric Motors and Drives, Fifth Edition, is intended for non-specialist users of electric motors and drives, filling the gap between mathematically-based academic textbooks and the more prosaic 'handbooks', which provide useful detail but little opportunity for the development of real insight and understanding.

The book explores all of the widely-used modern types of motor and drive, including conventional and brushless D.C., Induction Motors, Permanent Magnet and wound rotor Synchronous Machines and Reluctance Motors, providing readers with the knowledge to select the right technology for a given job.

The fifth edition provides a physical basis for understanding the control algorithms employed in Field Orientation and Direct Torque Control, and a comprehensive explanation of reluctance torque, which is becoming increasingly important when combined with permanent-magnet torque in applications ranging from domestic appliances to automotive.

This book assists readers to acquire knowledge and understanding of the capabilities and limitations of drive systems without having to struggle through unnecessary mathematics and theory. It is targeted at electrical and mechanical engineers and technicians requiring a working knowledge of motors and drives, and students of electrical and mechanical engineering seeking an accessible introduction and understanding. Practical review questions are provided at the end of each chapter, with model answers and commentary in an Appendix.

Austin Hughes, now retired, was a long-time member of the innovative motors and drives research team at the University of Leeds, UK, and has established a reputation for an informal style that opens up complex subjects to a wide readership, including students and managers as well as technicians and engineers.

Bill Drury is an independent consultant in power electronics, electrical machines and drives (PEMD). He has 45 years industrial PEMD experience at Siemens, Rolls-Royce and for 20 years was Executive VP of Technology of Control Techniques. He is a Visiting Professor at University of Bristol and Newcastle University.

ISBN 978-0-08-102615-1

9 780081 026151

ELSEVIER

Newnes
An imprint of Elsevier
elsevier.com/books-and-journals