Frequency synthesizers — 3

The generation of wanted frequencies from other frequencies

by R. Thompson, M.I.E.E.

Parts 1 and 2 of this three-part series described the synthesis of frequencies by means of addition, multiplication and division. This final part covers the principles of phase-locked loops as synthesizers and presents three solutions to specific problems.

Phase-locked loops

It has already been mentioned that the major difficulty with synthesizers is not the generation of the wanted frequency so much as the rejection of unwanted products. The phase-locked loop, p.l.l., is a circuit capable of providing highly selective amplification, its unique feature being its precise and automatic tuning. It is because of this feature that the p.l.l. is used extensively in frequency synthesis.

Figure 14 shows the block diagram of a basic p.l.l. It is simply a feedback circuit which controls the phase, and hence frequency, of the output with respect to the phase of an input signal. For steady state conditions there can be no frequency error between f_{in} and f_{out}. This follows from the action of the phase detector which integrates any frequency error. There will in general be a steady state phase error, though this can be made arbitrarily small by increasing the loop gain.

The term 'type 1 loop' is sometimes applied to low gain loops and 'type 2 loop' to high gain loops. Strictly speaking type 2 loops have zero steady state phase error and employ a second integrator in the loop. However, it should be appreciated that a range of performance is available between types 1 and 2 by varying gain.

The simplest type of loop has no low pass filter. The open loop gain characteristics therefore take the forms shown in Fig. 15(a). The closed loop frequency response is also shown and the bandwidth is approximately equal to the frequency at which the open loop gains equals unity. This follows from the action of the feedback which is able to stabilise gain at all frequencies where the loop gain is greater than unity. The 6dB/octave slope is produced by the integral relationship of phase and frequency.

The loop will be able to track changes in input frequency, or compensate for voltage controlled oscillator (v.c.o.) errors, provided the frequency change does not exceed $\pi/2K_0$. This assumes that the phase detector saturates at phase errors of $\pm\pi/2$, as many detectors do. Increasing the loop gain K_o therefore increases the "hold in" range. However, this increases the bandwidth as shown in Fig. 15(a) and there is therefore a conflict between bandwidth and hold in range in the choice of K_o.

The transient response shown in Fig. 15(a) is that of a simple CR filter, as of course is the frequency response. This type of loop is in fact referred to as a first order loop since the Laplace transfer function has a first order denominator (giving a simple pole in the s plane).

If the loop low pass filter consists of a simple CR filter the transfer function becomes second order (giving two poles in the s plane). The advantage of introducing a filter can be seen from Fig. 15(b). To a first order, the closed loop bandwidth will always be equal to the open loop unity gain frequency ω_o. With the added filter therefore it is possible to increase K_o without increasing the bandwidth. The frequency responses of the loop illustrate this, although the exact shape of the response depends on the ratio ω_1/ω_0. The transient response, like the frequency response, can exhibit overshoots which are characteristic of multiple pole networks.

The time response is characterised by a 'natural' frequency, ω_n, and a damping factor ζ. In general, responses with large overshoots are highly undesirable and we may therefore still have conflicting design requirements. Loop gain (for accuracy and 'hold in'), bandwidth (for filtering) and ζ (for transient response) cannot be independently controlled by the two variables K_o and ω_1. The solution to such a design conflict is to use a modified CR filter, normally referred to as a lead/lag filter. Fig. 15(c) shows the filter. The purpose of R_2 is to take out the effect of C as the frequency increases. The filter therefore has an initial 'break' frequency ω_1 when $R_1 = X_c$ and 'breaks back' towards a simple resistive attenuator when $X_c \doteq R_2$. A loop with such a filter is still 'second order'.

The transient response of the loop is determined primarily by the open loop phase shift around its unity gain frequency. Lowering ω_1 below ω_o increases this phase shift towards 180° from the 90° shift present with no filter.

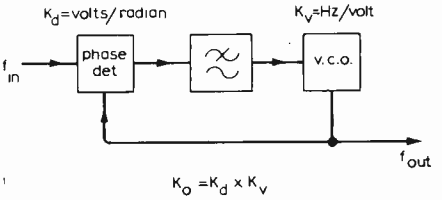

Fig. 14. *Basic phase-locked loop, in which no frequency error can exist between f_{out} and f_{in} in the steady condition.*

It has been seen that this results in decreased damping. Adding the break back at ω_2 tends to cancel the increased phase shift, its effect obviously increasing as ω_2 approaches ω_1. The effect of varying ω_2 is shown in Fig. 15(c).

The most common type of loop found in synthesizers is in fact a high gain second order loop. High gain is often not required specifically for accuracy but to reduce the level of noise generated by the phase detector. This detector is often a logic circuit producing a variable width square wave at the comparison frequency. The components of this, even after filtering in the low pass filter, will produce phase noise in the output signal. A high gain loop will have a very small steady state phase error producing very narrow low energy pulses at the detector output. Use of such high gain means that the filter must use a lead/lag network. The filter is in fact often in the form of an operational amplifier integrator which provides increased loop gain.

The loop bandwidth is controlled by a simple low pass filter and bandwidths of only a few Hertz can easily be obtained. When it is remembered that the input to the loop can be at high radio frequency the unique filtering characteristics of the p.l.l. can be appreciated. It can be used to separate wanted from unwanted signals even when the percentage frequency separation is very small.

The cut-off rate of the loop frequency response is only 6 dB/octave when a lead/lag filter is used. This can be increased at frequencies substantially greater than ω_o without altering the basic characteristics of the loop. In synthesizers this is often done to reduce comparison frequency noise generated in the phase detector which can modulate the v.c.o.

The 'hold in' range of the loop represents the range of frequency

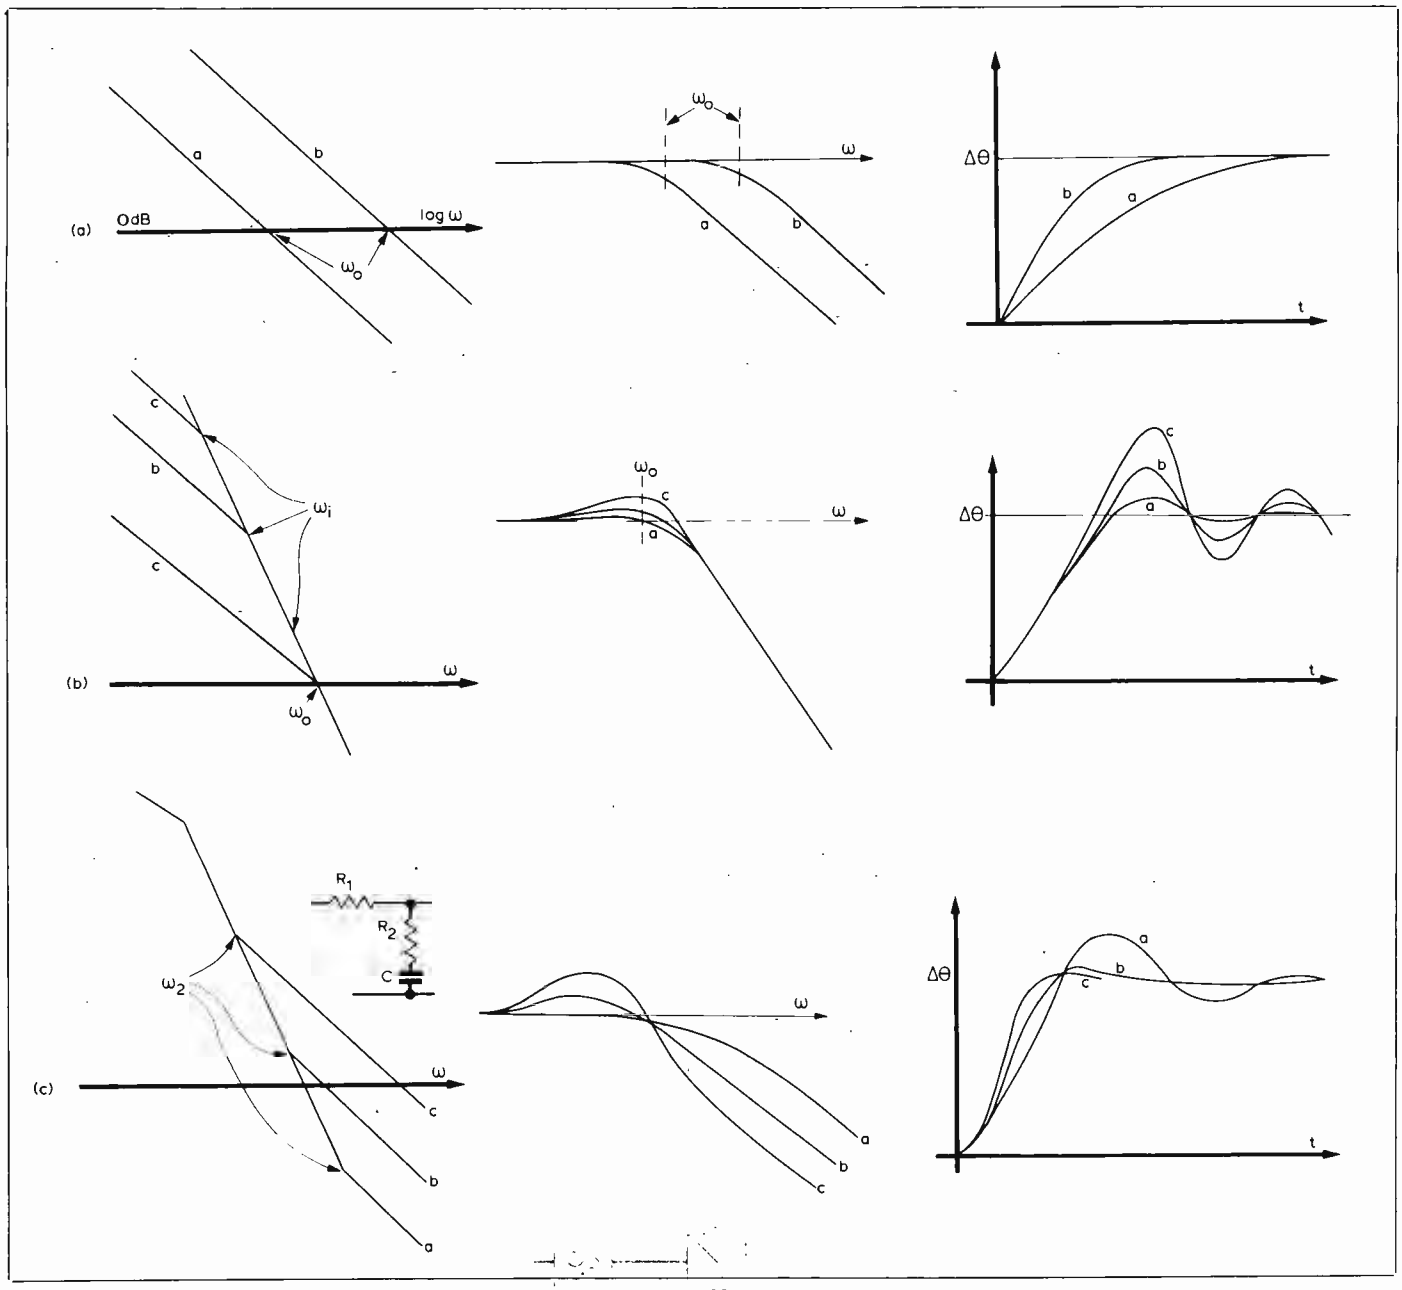

Fig. 15. *Characteristics exhibited by loops with no filter (a), simple CR filter (b) and lead/lag filter (c). Loop gains are shown in the left-hand column and the associated frequency responses in the centre. Corresponding transient respose is in the right-hand column.*

which the loop can track once phase lock has been established. The maximum frequency error for which the loop can acquire lock is called the 'pull in' range. This is never greater than the 'hold in' range and even with this range acquisition can take very long times. High gain second order loops are particularly bad in this respect. Because of this type of acquisition problem p.l.l.s sometimes use a subsidiary frequency discrimination to provide coarse tuning correction. Once phase lock is established the frequency discriminator action can be ignored.

Synthesizer circuits

Having looked at four basic types of "building block"; addition, multiplication, division and the p.l.l., we will now consider their combined application in frequency synthesizers. The basic requirement of a synthesizer is the generation of one frequency from another and this can be stated as:

$$f_2/f_1 = X/Y.$$

where X and Y are rational numbers.

Synthesis can therefore be achieved by the simple cascade of an X-times multiplier followed by a Y-times divider. Practical realisation in this form is limited by the values of X and Y.

The divider techniques described earlier allow a wide range of division ratios to be achieved with proprietory devices operationg at frequencies up to about 1 GHz. In general therefore it will be the realisation of high multiplication factors which will impose practical difficulties. Where p.l.l. techniques are not used multiplication factors of greater than 5 in one stage are difficult if reasonable spectral purity is to be maintained. The procedure adopted to avoid excessive multiplication is to split the requirement into the form:

$$X/Y = (x/y)\,(X_1 \pm X_2/Y_2)$$

For instance, if we have a specific requirement to generate 2.35 MHz from a reference to 3.52 MHz, then:

$$(f_2/f_1 = 2.35/3.52 = 235/352 = X/Y$$

this can be split into:

$$235/32 = (5/32) \times (47/11) = (5/32)\,(4 + 3/11)$$

This synthesis involves a mixer to carry out the subsidiary addition as shown in Fig. 16. The multiplication terms, 5, 4 and 3 are all reasonable, but the addition of 4 + 3/11 represents a ratio of 44:3 in the input frequencies to the mixer. As explained earlier unwanted frequencies generated in the mixer give upper and lower bounds on the ratio for practical filtering. A better factoring is therefore:

$$(5/11)(3 + 14/11)$$

giving a mixing ratio of 33:14. The multiplication factor of 14 can be reduced

by a second mixing loop as shown in Fig. 16, providing synthesis in the form:

$$(5/32)(3 + (1 + 3)/11)$$

The requirement has thus been factored into a form readily realised with practical circuit elements. It can be seen that the complexity of the synthesizer is dependent on the practical bounds set primarily by filtering problems.

Phase-locked loops are widely used in synthesizers because their excellent filtering characteristics allow very wide bounds on multiplication and mixing ratios. Before looking in more detail at the application of p.l.ls, an alternative form of selective filtering is worth mentioning. This is the "triple mixing" system shown in Fig. 17.

Highly selective fixed-tune filtering is provided and the auxiliary oscillator is used to "scan" the incoming spectrum with these filters. Filtering having been accomplished, the third mixer cancels out the auxiliary oscillator frequency together with any associated frequency drift. Mixer 2 is used to introduce an interpolation frequency, normally an incremental frequency which can be used to span the separation of the harmonics in the input signal f_1.

The p.l.l. can also be considered as a form of heterodyne/fixed filter system, the signal being converted down to zero frequency and fixed filtering provided by the low pass filter. In its basic form it can be used simply as a high grade filter in synthesizing loops, allowing constraints on multiplication and mixing ratios to be relaxed. More frequently, the p.l.l. is modified by the inclusion of a divider in the feedback loop as shown in Fig. 18(a). When modified in this way, the loop stabilises the v.c.o. frequency at n times the input frequency. The p.l.l. loop gain is reduced by a factor of n and the characteristics of the loop are modified accordingly. This arrangement operates as a multiplier, having the very important feature that its multiplying ratio is readily controllable over a wide range by proprietory divider circuits. It is also "fail safe" on its operation, simple harmonic multipliers have the problem that the wrong harmonic may be selected, particularly where high ratios are involved. In synthesisers, it is normally an advantage to have maximum bandwidth, as opposed to the p.l.l. being used as a narrow band filter of *input* signals. The wider bandwidth allows cancellation of noise generated *in* the loop, in particular v.c.o. noise. Wide bandwidth also reduces tuning time. Since the bandwidth must be subtantially less than the comparison frequency we have a design conflict between low comparison frequency for fine tuning increments and wide bandwidth for noise and tuning time performance.

The solution of the synthesizer requirement $F_1/F_1 = X/Y$ can be realised in a simple two-step operation, using a multiplying p.l.l. Figure 18(b) illustrates the solution to our previous example of 235/352.

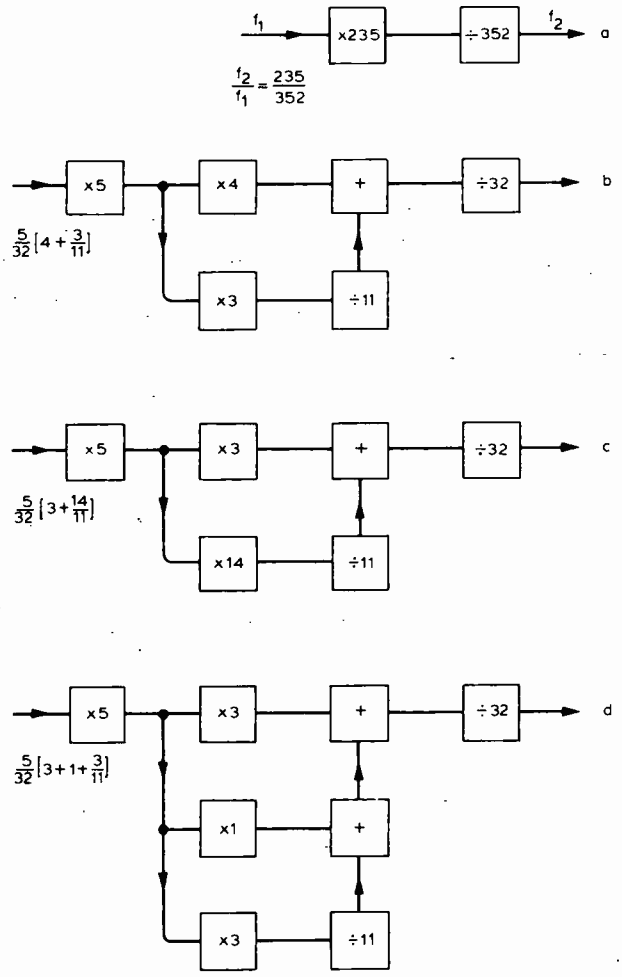

Fig. 16. *Evolution of synthesizer to perform basic functions shown at (a).*

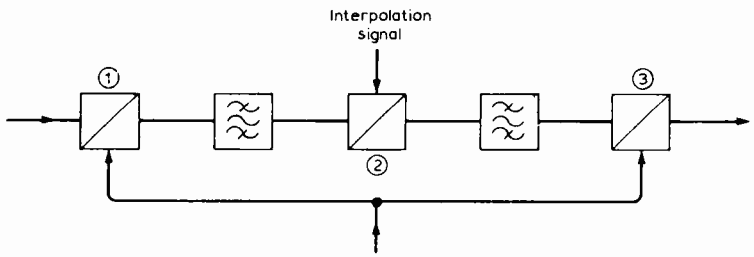

Fig. 17. *Triple mixing synthesizer.*

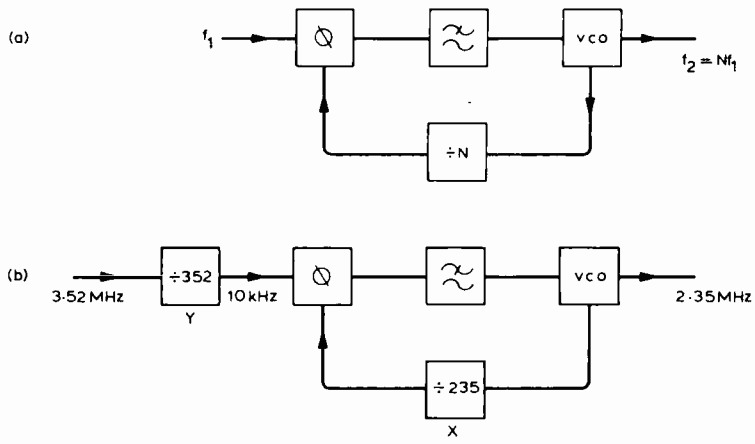

Fig. 18. *Phase-locked loop multiplier (a) and a p.l.l. multiplier solution to the 2.35/3.52 problem (b).*

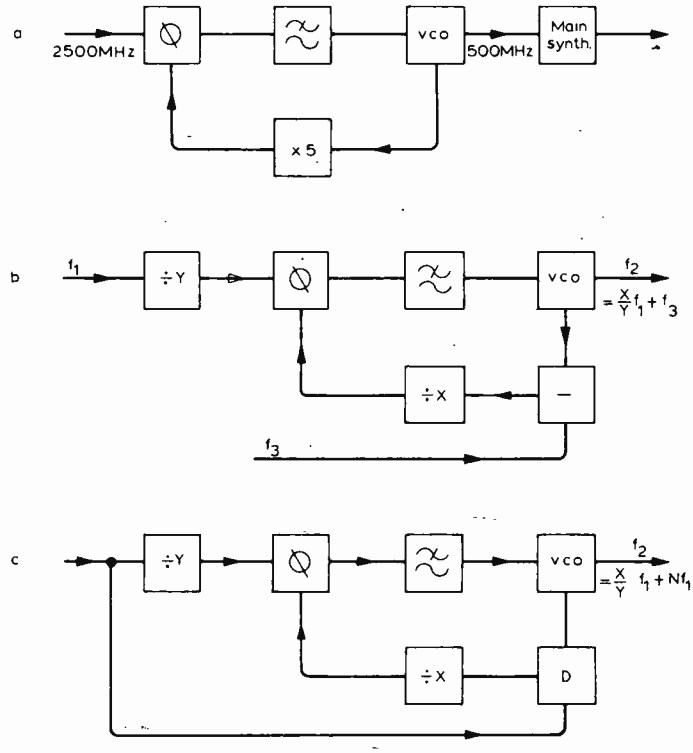

Fig. 19. *P.l.ls can be used as high-speed dividers by inclusion of a multiplier in the loop. P.l.l. at (a) divides by 5. Mixer at (b) reduces frequency requirements of divider and (c) shows alternative method.*

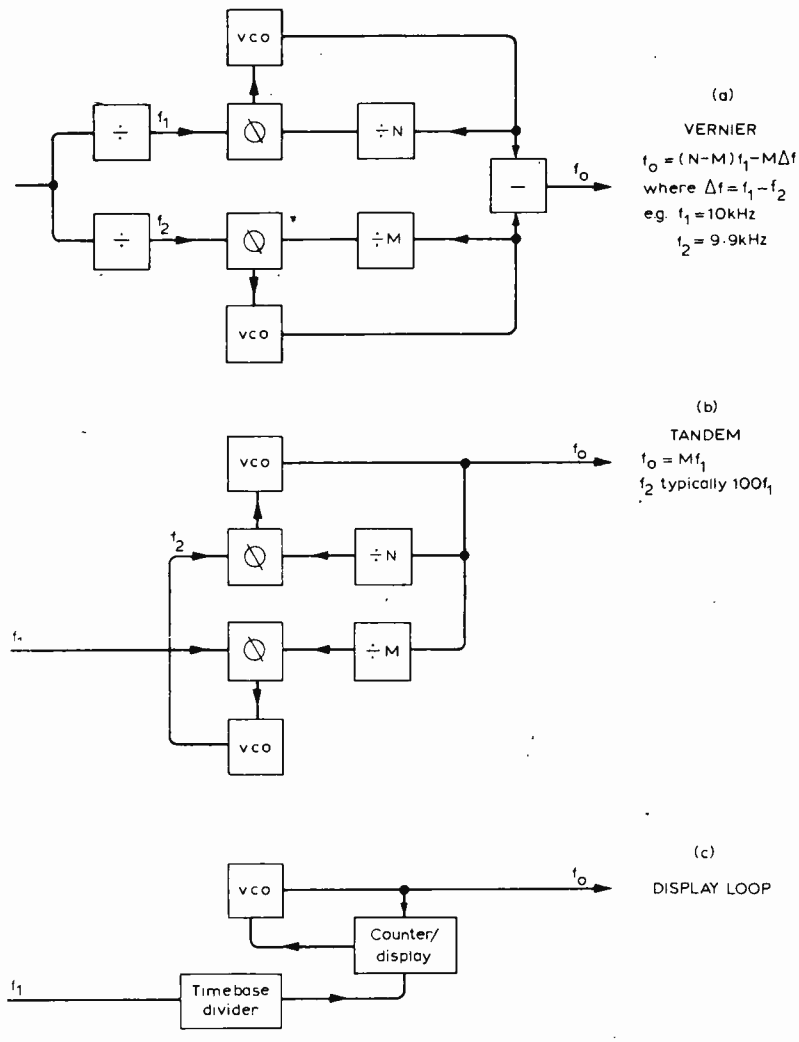

Fig. 20. *Methods of achieving wide bandwidth and small increments.*

The operating frequency of the divider circuits is often an important constraint on the synthesizer design. For this reason, Fig. 18(b) shows the division followed by multiplication; reversal of this would require the divider to operate at over 850 MHz. Currently-available dividers in integrated circuit form have the following approximate frequency limits:

direct variable dividers	– 10 MHz
variable division with "early decode"	– 25 MHz
variable modulus prescaler	– 500 MHz
fixed prescaling	–1500 MHz

At frequencies higher than this, or as an alternative to prescaling techniques, the p.l.l. can be used as a high-frequency divider by employing a multiplier in the feedback path. Figure 19(a) shows an arrangement where the reference frequency is very high.

Another alternative used to cope with high frequency requirements is to introduce a mixer into the p.l.l. This reduces the frequency handled by the divider, such an arrangement being shown in Fig. 19(b). The synthesis is $f_2 = f_1X / Y + f_3$, with a maximum frequency at the divider of $f_2 - f_3$. A harmonic of the reference frequency can be used and the digital mixer described earlier can be used as a combined mixer and harmonic multiplier. Figure 19(c) shows such an arrangement. The synthesis is $f_2=f_1X / Y+ Nf_1$ where N is the harmonic of f_1 selected by the loop. For example f_1 might be 5 MHz and f_2 required to be in the region of 56 MHz. The 11th harmonic of f_1 is subtracted from f_2 in the mixer, giving a frequency into the variable divider in the region of 1 MHz (the exact frequency depending on X and Y).

As described earlier we have problems in achieving wide bandwidth and low frequency increments. The vernier system uses two synthesizers, offset in their reference frequencies so that increments in frequency are set by the difference, as shown in Fig. 20(a).

In the Tandem system a wideband synthesizer produces a low-noise output while the fine frequency increments are achieved via a second synthesizer providing the reference frequency for the output loop. The secondary loop is slow because of the low comparison frequency, so tuning is slow when a change to that loop is involved.

The display loop system uses an accurate timebase counter to display the frequency. When the required frequency is obtained the loop is completed with correction being generated by any variation of the last digit of the counter.

The next article will conclude the series with examples of specific designs.