P.c.b. layout for high-speed Schottky t.t.l.

Requirements of printed-board design for low inductance and effective decoupling

by D. Walton, B.Sc. (Hons), Ph.D.

A great deal has been written on the subject of logic design and quite comprehensive books appear almost monthly. In general, however, the published material neglects an extremely important area and one which probably gives the most trouble to practising engineers. This area, which is dealt with in the present article, is concerned with the layout of logic on printed circuit boards in order to ensure reliable operation. The impetus for writing this article comes from the author's own experience of the lamentable lack of understanding of these basic considerations.

IT SHOULD not be concluded from the preamble that the subject is a difficult one; indeed the mathematics employed in the present paper is extremely elementary. The problems are caused rather by the historical progression from analogue to digital techniques with the consequent carrying out of well-tried analogue practices into the digital environment. Unfortunately, the requirements for digital circuitry are frequently opposite to those needed by the analogue variety and hence there is a need for a complete reconsideration of the requirements.

Low inductance bussing

To understand the criteria which determine how the supply and GND lines should be distributed to the t.t.l., first take the case of a t.t.l. gate driving its output line from low to high. For the gate to drive the output line high it must pass current into it. The output line must be considered as a transmission line of impedance Z_0 if its length exceeds 10cm. In practice, Z_0 will be in the region of 100Ω and for a single logic signal changing from low to high the instantaneous output current will be given by $I_0 = 5/100 = 50\text{mA}$. This current must be obtained from the supply rails in a time comparable to the risetime of the signal. If, for Schottky t.t.l., $t_{r(\min)} \simeq 1.5\text{ns}$, then charge must be transferred from the decoupling capacitor to the gate and hence to the output line in this time. Remember that charge is obstructed from flowing into the gate by the inductance, L, of the loop ABCD in Fig. 1. If this is approximately 2cm

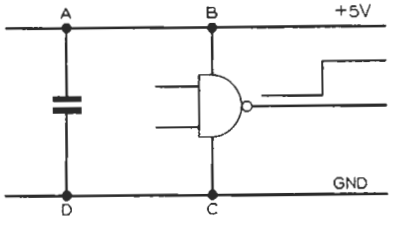

Fig 1. Example of gate, with decoupling, producing a low-to-high transition.

square with reasonable track width then, using the formula for parallel wires, $L = \ln(a/r)\,\mu_0/4\pi \approx 30\text{nH}$. The e.m.f. dropped across L will then be given by $E = -L\,di/dt$. Therefore,

$$E = \frac{30\times10^{-9}\times50\times10^{-3}}{1.5\times10^{-9}}$$

$$= 1 \text{ volt}$$

This is a considerable voltage and it should be remembered that it is the result of a single gate switching. If all four gates in a pack switch together the currents will be additive and the rail will fall by 4 volts.

The first requirement of a power distribution system must therefore be low inductance between the i.c. and the decoupling capacitor. This is achieved by the track layout shown in Fig. 2(b), where a low inductance path from C to the i.c. is provided by keeping the V_{CC} and GND tracks close together.

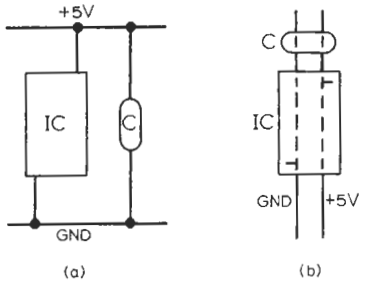

Fig. 2. Two ways of laying out supply lines. Preferred method, giving lower inductance, is at (b).

Manufacturers of i.cs usually specify one decoupling capacitor for every 5–10 i.cs which, with the track layout of Fig 2(a) results in prohibitively high inductance between the capacitor and the worst-case positioned i.c. The safest course is to provide the track layout as in Fig. 2(b) but also to put one capacitor adjacent to each i.c. Clearly, this can be achieved by having one capacitor for each pair of i.cs.

Decoupling capacitors

The foregoing argument shows that the capacitor is better thought of as a reservoir capacitor which supplies the local, instantaneous current demands as i.cs switch. This means that the important parameter for such a capacitor is the instantaneous current which it can supply. Some manufacturers specify capacitors for i.c. decoupling by giving the maximum pulse risetime, which corresponds to a maximum current for a given size of capacitor. For instance, a 47nF capactior specified at 50V/μs can supply a current given by

$$i = C.\frac{dv}{dt} = 47\times10^{-9}\times\frac{50}{10^{-6}}$$

$$= 2.5\text{A},$$

which is adequate in the context of the previous calculation.

The other check to make is that the current drawn from the capacitor does not cause its voltage and hence the rail voltage to fall excessively. If the local demand is equal to 10 gates switching, the current demand will be 500mA; to be safe, assume that this demand lasts for 10ns, and design for a voltage drop at the capacitor of 50mV.

Thus,

$$i = C.\frac{dv}{dt}$$

$$0.5 = C\frac{50\times10^{-3}}{10\times10^{-9}}$$

$$C = 100\text{nF}.$$

This suggests that we should provide approximately 100nF for each pair of packages.

It might be thought that radio frequency type capacitors are necessary for t.t.l. decoupling, but this is not so. To show why requires more space than can be spared in an article of this type but essentially it is because the frequently adopted model of a capacitor, which

proposes that it possesses a lumped series inductance, breaks down in the case of a single applied step. There is therefore no reason for the designer to be afraid to employ non-ceramic capacitors provided they have adequate V/μs ability. In the author's experience 1μF tantalum beads perform well as decoupling capacitors.

Transmission-line model

The best way to think of the power distribution system is as a transmission line, with each package connected to an ideal voltage source via an impedance equal to the transmission line impedance*. This impedance must be sufficiently low for negligible voltage transients to be produced on the line by gates switching within the package. The impedance of a transmission line is given by $Z_o=\sqrt{L/C}$, where L and C are the inductance and capacitance per unit length respectively. To calculate Z_o for the case of two tracks close together:

$$L=\frac{\mu_0}{4\pi}\ln\frac{a}{r}$$

where μ_0 is 5. A and r are taken as 2mm and 0.5 mm. Therefore

$$L=0.6\mu H/m.$$

If a 100nF capacitor is placed every 5cm along this line, then:

$$C=100\times 20\ nF\ m^{-1}=2\mu F\ m^{-1}$$

Therefore $Z\approx 0.5\Omega$.

An instantaneous current demand of 200mA – corresponding to 4 gates switching – will produce a voltage transient of 100mV. This is only just acceptable and suggests that the value of C should be increased. Note however, that laying out the tracks with wider spacing and using smaller capacitors – 10nF for every few i.cs, which is not uncommon, will create a situation much worse than this.

Auto-decoupling in t.t.l.

In the context of the preceding remarks some readers may wonder how systems which they have seen or have worked with managed to function at all, since it is common to see most or all of the above design guidelines violated. To see the answer to this, consider the structure of the t.t.l. gate output circuit, when this is driving the following gate input low, as in Fig. 3.

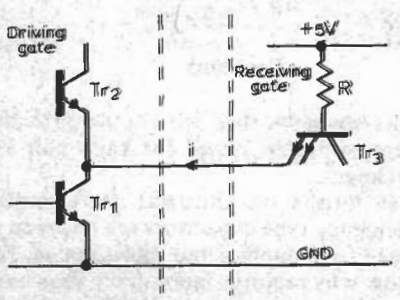

Fig. 3. Totem-pole t.t.l. output stage, driving succeeding gate low.

According to the specification for, say, a 7400 the typical values of i and R are 1.0mA and 4kΩ respectively. When the gate output is low it sinks a current i, given by $i=(V_{cc}-V_{be}-V_{CE(sat)})/R$, where V_{be} is the base-emitter voltage of Tr_3 and $V_{CE(sat)}$ is the collector saturation voltage of Tr_1.

If V_{be} and $V_{CE(sat)}=0.7$ volts, to take a worst-case example, and $V_{cc}=5$ volts

$$\text{then } i=\frac{3.6}{R}$$

Now consider what happens if the rail voltage drops, due to a transient load imposed by the output of another gate switching. When V_{CC} drops there is no change (to a good approximation) in the V_{be} drops. Suppose the rail drops by 10% then:

$$i_1=\frac{5-1.4}{R}$$

$$i_2=\frac{4.5-1.4}{R}.$$

Therefore

$$\frac{i_1-i_2}{i_1}=\frac{0.5}{3.6}=14\%.$$

In other words a 10% change in V_{CC} produces a 14% change in the current load placed on the rail. In effect what is happening is that each gate output which is holding another input low acts as a 'reservoir' of current and when the rail voltage drops as another gate drives its output high all the other gates give up some of their current to assist. This is what I would call the 'good neighbourliness effect' in t.t.l. In general, some gates on a voltage bus will be low and so act as current supplies. The problem arises when none or only a few are in this state – a critical situation for a badly designed system and one which could cause a failure. It should be remembered that a logic system should work for all possible combinations of states which can occur in practice and a hazard of this type could have serious consequences. It is therefore insufficient to demonstrate that a system 'works' because if the power distribution system is badly designed there is always the chance of an untested situation bringing about a failure of the system. It is assumed that in a logic system of reasonable size it is impossible to test all possible combinational situations, and doubly impossible to test all possible changes of situation!

The problem with Schottky t.t.l. is that the increase in speed does not allow time for the "good neighbourliness effect" to act, consequently one is many times worse off with Schottky than with ordinary t.t.l. Schottky is a less forgiving family than conventional t.t.l. and much more care must therefore be taken with power distribution to ensure reliable performance.

* A package at the centre of a power bus will see two lines in parallel and hence half the impedance. We will adopt the worse figure for the purpose of this argument.

The current spike

As just described, the main cause of transient current demands in a Schottky t.t.l. system is the initial current surge when a gate switches into its transmission line load. The manufacturers' data overlooks the mechanism entirely. There is another cause of transient current demand which results from the 'push-pull' design of the t.t.l. output stage shown in Fig. 4. The current spike is produced because, on the 0 to 1 transition, the upper transistor turns on while the lower transistor is still turning off. This leads to a current surge of 10mA with duration of about 10ns[1]. Provided the design guidelines laid down in the earlier sections with regard to power supply bussing and decoupling have been followed, this small additional hazard will be taken care of. In fact, since a logic gate is driving a transmission line which is a resistive rather than a capacitive load, there is no need to provide a totem pole output and this must be regarded as one of the bad features of the t.t.l. family.

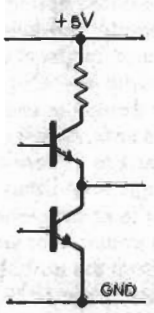

Fig. 4. T.t.l. output configuration leads to current spike at transmission.

Interconnexions

To implement a system successfully using the t.t.l. family it is necessary to interconnect correctly between logic gates.

Transmission lines. The correct model to use for interconnexion between logic gates is a two-wire transmission line. It is impossible to understand how a signal travels from gate to gate without taking the return path into consideration. Indeed it is impossible for a signal to travel without a return path! Consider the two-wire transmission line shown in Fig. 5, in which a zero rise-time is pro-

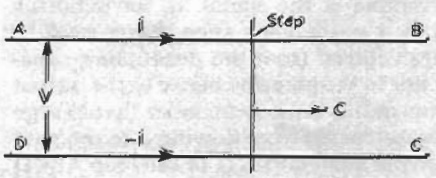

Fig. 5. Two-wire transmission line.

pagating to the right with velocity c. Ahead of the step there is no current in the wires and no voltage differences between them. Behind the step there is a current i in the direction of AB and a current -i in the direction of DC with a voltage difference V between the wires. It can be shown[2] that $V = iZ_0$, where $Z_0 = \sqrt{L/C} = \sqrt{\mu/\epsilon}$ where Z_0 = characteristic impedance of line, L = inductance per unit length of line, C = capacitance per unit length of line, μ = permeability of medium between wires, ϵ permittivity of medium between wires. The velocity of propagation $c = 1/\sqrt{LC} = 1/\sqrt{\mu\epsilon}$.

These equations are true for any two-conductor system where the resistance of the conductors can be neglected and the medium between the conductors is well-behaved. These conditions are met by tracks on a printed circuit board for any track width which can be manufactured. The step which we have just described is a transverse electromagnetic disturbance. Since the equation relating current and voltage on a transmission line is $V = iZ_0$, it follows that the effect of a transmission line on the driving circuit can be considered in terms of a resistance $R = Z_0$ connected in place of the line. This was the procedure followed earlier in calculating the current drawn from the supply rail by a gate as it switches.

The impedance Z_0 depends on the cross-sectional geometry of the conductors employed and its calculation is extremely difficult except for very simple cases. It is, however, a relatively slowly varying function of the geometry[3] (usually logarithmic) and therefore this need not worry us too much. For a track on a printed circuit board laid out according to the design rules evolved in this paper a value of Z_0 of around 150Ω can be assumed.

One key feature of a board of logic which distinguishes it from most analogue systems is that there are a multiplicity of signal paths from various points scattered about the board to various other similar points. It is essential that each of these signal routes has an adjacent return path. The simplest way, conceptually, to achieve this is to provide a ground plane on one side of the board. In practice this is difficult since it usually requires multi-layer construction, with the increased cost and complexity which this entails, in order to accommodate the signal interconnexions. With Schottky t.t.l. it is not necessary to go to this extreme; all that is required is a ground grid laid out so that a signal line is never more than one inch away from its return path.

Ground loops. It might be argued that this scheme leads to ground loops which, from our experience with analogue systems (e.g. audio equipment) are to be avoided. The plain fact is though, that on a logic board, ground loops are of no importance. The reasons for this are somewhat complex but it is probably useful to note one simple argument. In a high-gain amplifier, induction of a few millivolts at the input due to ground loop pickup can lead to an output of the same order as the signal. In logic this is not the case; a few millivolts into a gate input make no difference whatsoever. Hundreds of millivolts of noise are required before we will significantly degrade the noise immunity of a t.t.l. system.

It is probably valuable to examine a situation where a logic board has been laid out in order to avoid ground loops. A possible layout of power and ground connexions, which is quite commonly adopted in the industry, is shown in Fig. 6. Now, if circuit A sends a step to circuit B there is no adjacent return path. In practice, since a fast step requires a return path it will simply use adjacent signal lines as returns, resulting in the induction of transient noise on these other signal lines. A further consequence is that the input to B will take a longer time to settle with a consequent reduction in the speed of the system. As was explained earlier, the layout of Fig 6 is also bad from the point of view of placing excessive inductance in the way of charge travelling between i.cs and decoupling capacitors.

Recommended layout

A recommended scheme for laying out a printed circuit board is shown in Fig. 7. The power rails are run as close together as possible along the columns of integrated circuit packages and are interconnected at the top and bottom of the board. These provide return paths for logic signals travelling parallel to them. To provide return paths for signals travelling across the board the ground pins of the packages are connected together from left to right. Thin track, of the same thickness used for signal interconnexions can be used for this. A tantulum bead 10μF decoupling capacitor is provided between each pair of i.cs. Notice also that ground connexions are brought out at regular intervals across the edge connector. These provide return paths for signals travelling on and off the board.

If all these design rules are followed a reliable system will result and the consequent savings in servicing and testing will amply repay a little consideration given to board layout at the design stage.

References

1. Bonham, B. 'Schottky t.t.l.', T.I. application report B93.
2. Catt, I. 'Crosstalk (noise) in digital systems', I.E.E.E. Transactions on Electric Components', EC-16, 743-763 (1967).
3. Gunston, M. A. R. 'Microwave transmission live impedance data'. Van Nostrand Reinhold.

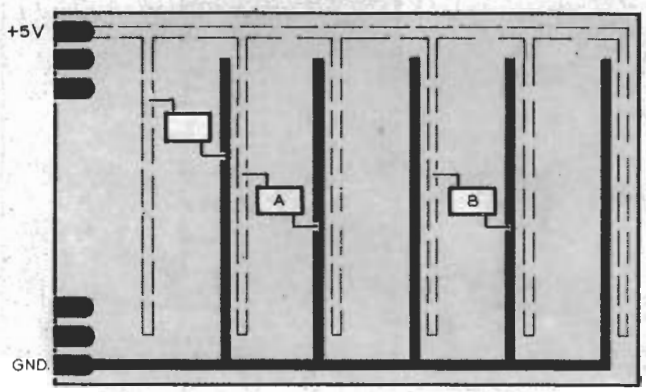

Fig. 6. A bad layout giving high inductance and few adjacent signal return paths, which leads to cross-talk.

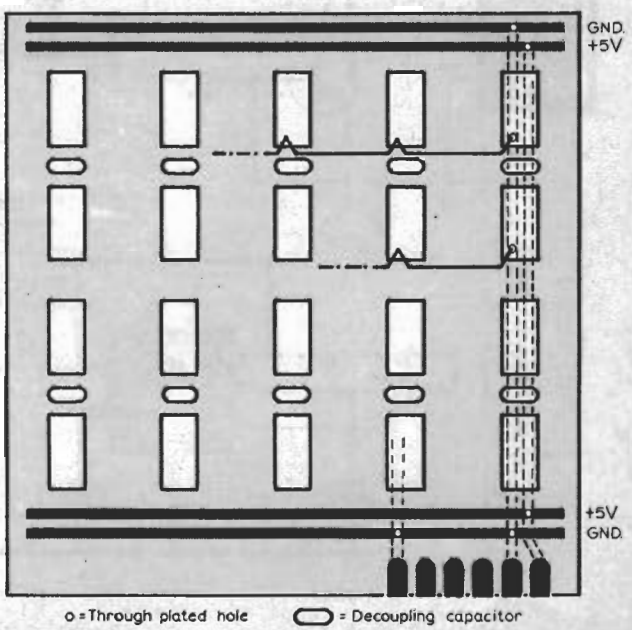

Fig. 7. Recommended layout.