6

Distribution of D.C. Power to Logic

In a digital system, sudden massive changes of load current in the power supply serving the logic can easily occur. For instance, it is not uncommon for a 32-bit word of data to be gated on to 32 parallel signal lines. If the word happens to be 'all ones', the amount of current suddenly switched out of the +5 V supply is quite large. For instance, if the lines are shunt terminated in their characteristic impedance of 100 Ω, the total current switched, at 50 mA per line for a 5 V signal, is 1.6 A. We see what a massive event this is when we assume a reasonable risetime for the signal, say 5 ns, and find that the rate of change of current out of the supply

$$\frac{di}{dt} = \frac{1.6}{5 \times 10^{-9}} \text{ A/s} = 320 \times 10^6 \text{ A/s}$$

For proper functioning of the logic it is important that such a massive, rapid current demand should not cause a significant drop in the voltage bus in the vicinity of the demand at any time after that demand has occurred. (A 50 mV drop could be regarded as the maximum acceptable.)

A hierarchy of energy reservoirs, or capacitors, sustains the d.c. voltage at the critical point. During the first few nanoseconds, a small capacitor, near to the point of the suddenly changing load, supplies the current. Later, other similar capacitors distributed around the logic a little further away help out. Later still, when all the small local capacitors would have been depleted, the larger energy reservoir of, say, 20 000 μF in the power supply itself sustains the voltage.

As this final reservoir begins to be depleted, the series pass transistors in the power supply alter their current levels to supply the deficiency.

Figure 6.1 shows a typical hierarchy of energy stores such as might be found in a system in use today. At every stage it is deficient. Starting at the right, each capacitor has to have become seriously depleted, its voltage dropping by a volt or two (compared with the maximum reasonable drop of 50 mV) before there is enough voltage and time developed across the next inductor to the left to cause it to pass the required increase of 1.6 A. The situation is unsatisfactory by orders of magnitude, showing that it is not merely some systems that are deficient. In fact, most systems in service today are unable to supply a perfectly normal change in current demand at one point in the system while keeping the d.c. voltage bus near to its nominal value.

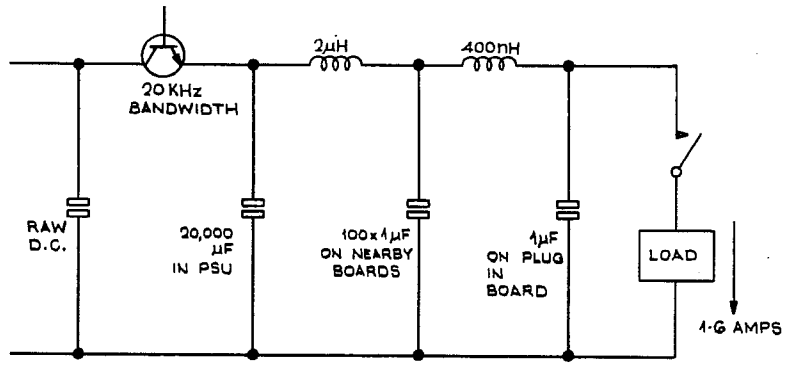

Figure 6.1

Why do systems continue to function if the above assertions are true? The answer is that an important element in the total picture was omitted from the diagram, and that was the other steady loads. When the 1.6 A is switched on, the d.c. voltage supply begins to sag. The result is that, with a reduced voltage across them, the other circuits demand less current. If the d.c. voltage supply drops by 5 per cent, the current taken by an integrated circuit falls by even more than 5 per cent. (The current drop is more than linear with voltage because some of the available voltage is 'lost' in fixed transistor V_{be} drops. This means that if there is a 5 per cent drop in the full voltage, the percentage drop across resistors in the circuit is more than 5 per cent.)

If the full system load is 30 A, a further load of 1.6 A only calls for the removal of about 5 per cent from this standing load, which is probably achieved by a drop of only some 3 per cent in the d.c. voltage supply; a mere 150 mV in the case of a 5 V supply. So we see that in a traditionally designed system, the voltage supply decoupling capacitors are to quite a degree merely a gesture of intent, and the real voltage decoupling is by the steady loads. This is unsatisfactory, because any change in d.c. voltage levels leads to a degradation of performance and possible failure. Timing circuits are upset, and margins against other types of noise are reduced.

When the 1.6 A load is suddenly switched off, a similar problem arises. This time, the voltage bus rises above its correct value in order to repel the now unwanted current which continues to arrive. The amount of over-voltage developed is about equal to the amount of under-voltage which occurred in the previous case. As before, the real voltage decoupling (or stabilisation) in traditionally built systems is caused by the other steady loads, which start to take more current when the voltage across them increases.

The designer should make a conservative (that is, large) estimate of the maximum change of load that can be expected in the logic, and design a hierarchy of energy stores that can sustain such a load.

7

Local Decoupling of Voltage Supplies by Printed Circuit Voltage Planes

When a logic gate switches on, a signal is launched down the signal line, the current step i being taken from the positive supply line, and an equal and opposite current $-i$ being dumped into the 0 V line at the point where the logic signal originated (see figure 7.1).

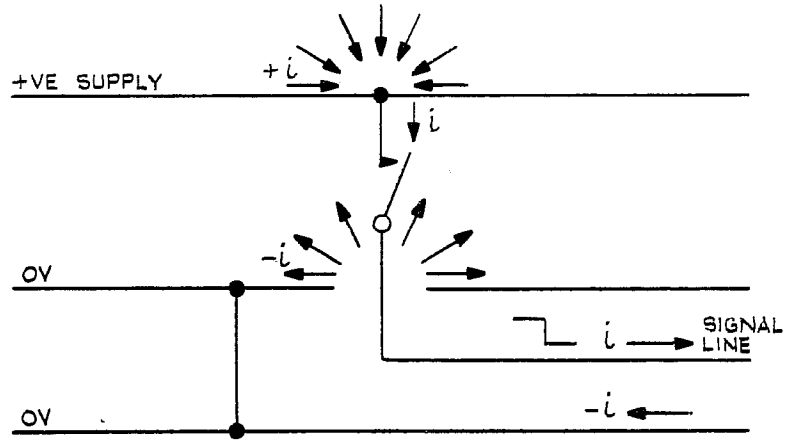

Figure 7.1

As a result, in order to satisfy Kirchhoff's law that the total current at any point must be zero, a circular wave of current and voltage – like the ripple caused by a stone dropped into a pond – flows out between the positive supply voltage plane and the 0 V plane. It is important that the impedance seen by such a signal

flowing out between the voltage planes, which can be regarded as a source impedance of the logic signal, should be small. Otherwise the positive voltage difference between the two planes at the signal source will temporarily collapse.

The Pie-shaped Transmission Line

In order to understand the nature of the decoupling action at a point between parallel voltage planes, it is easiest to approach it by way of the parallel-plate transmission line and the pie-shaped transmission line (see figures 7.2 and 7.3).

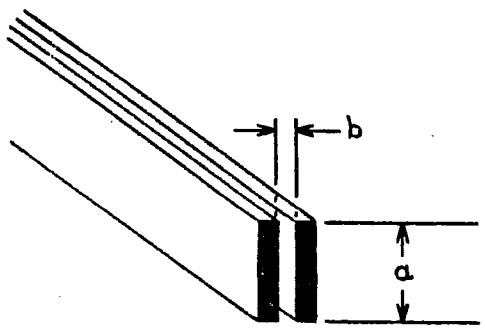

Figure 7.2

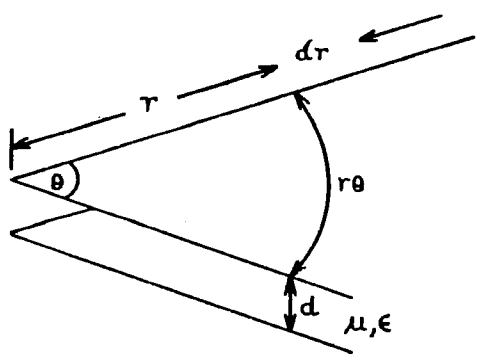

Figure 7.3

The formula for the characteristic impedance of a parallel-plate transmission line of width a and separation b with negligible fringing at the edge (because a is much bigger than b) is

$$Z_0 = \frac{b}{a}\sqrt{\left(\frac{\mu}{\epsilon}\right)}$$

This formula still applies for each small section of a transmission line where the width a is varying, in particular for a pie-shaped transmission line (figure 7.3), where over a distance δr it becomes

$$Z_0 = \frac{d}{r\theta}\sqrt{\left(\frac{\mu}{\epsilon}\right)}$$

Now if $\theta = 2\pi r$, that is, we are considering a complete plane (figure 7.4), we get

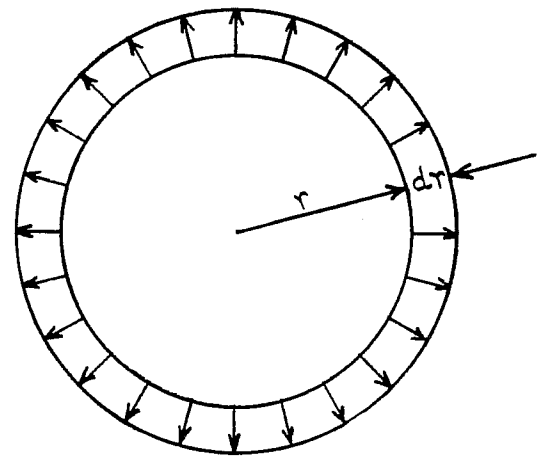

Figure 7.4

$$Z_0 = \frac{d}{2\pi r}\sqrt{\left(\frac{\mu}{\epsilon}\right)}$$

Now we know that, in a medium with permittivity ϵ and permeability μ, the outwards velocity of the signal must be $1/\sqrt{(\mu\epsilon)}$, therefore the velocity is

$$c = \frac{r}{t} = \frac{1}{\sqrt{(\mu\epsilon)}}$$

where time t is the time since the signal was introduced at the centre.

So using this last equation to substitute for the distance r we get

$$Z_0 = \frac{d\mu}{2\pi t} \ \Omega$$

where d is in metres.

A reflection related to Z_0 arrives back at the centre at time $2t$. If Z_0 is small when $2t = t_r$ (the risetime of the output of the logic gate), then natural decoupling between planes is satisfactory.

As an example, if $2t = 1$ ns, $d = 0.5$ mm, $\mu = 4\pi \times 10^{-7}$

$$Z_0 = \frac{0.5 \times 10^{-3} \times 4\pi \times 10^{-7}}{2\pi \times \frac{1}{2} \times 10^{-9}} = 0.2\ \Omega$$

This calculation shows that it is possible to keep the decoupling between voltage planes satisfactory by the addition of extra decoupling capacitors distributed at intervals of a few centimetres, to prevent superposition of signals from generating unacceptably large voltage transients between planes.

The proportion of the surface area occupied by the miniature tantalum capacitors will be insignificant.

8

Printed Circuit Board Layout for High Speed Schottky TTL

This subject is not a difficult one; indeed, the mathematics used is extremely elementary. The problems are caused rather by the historical progression from analog to digital techniques, with the consequent carrying over of well tried analog techniques into the digital environment. Unfortunately, the requirements for digital circuitry are frequently opposite to those needed by the analog 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 +5 V and 0 V lines should be distributed to the TTL, first take the case of a TTL 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 100 mm. In practice, Z_0 will be in the region of 100 Ω, and for a single logic signal changing from low to high the instantaneous output current will be given by $I_0 = 5/100 = 50$ mA. This current must be obtained from the supply rails in a time comparable to the risetime of the signal. If, for Schottky TTL, $t_{r(\min)} \approx 1.5$ 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 figure 8.1. If this is approximately 20 mm square with reasonable track width, then, using the formula for parallel wires

$$L = \frac{\mu_0}{4\pi} \ln \frac{a}{r} \text{ H/m} \approx 30 \text{ nH}$$

The e.m.f. dropped across L will then be given by $E = -L\, di/dt$. Therefore

$$E = \frac{30 \times 10^{-9} \times 50 \times 10^{-3}}{1.5 \times 10^{-9}} = 1 \text{ V}$$

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 V.

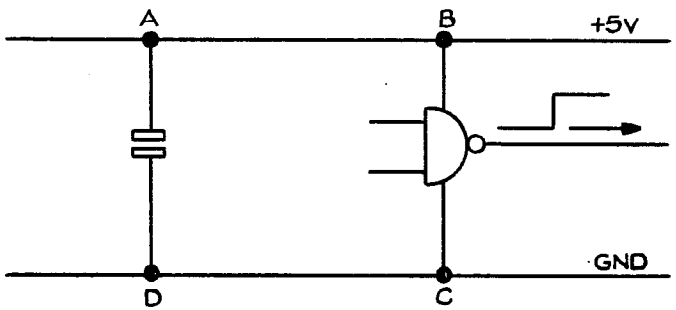

Figure 8.1

The first requirement of a power distribution system must therefore be low inductance between the IC and the decoupling capacitor. This is achieved by the track layout shown in figure 8.2, where a low inductance path from C to the IC is provided by keeping the +5 V and 0 V tracks close together.

Manufacturers of ICs usually specify one decoupling capacitor for every 5 to 10 ICs, which, with the track layout shown in figure 8.2a, results in prohibitively high inductance between the capacitor and the worst-case positioned IC. The safest course is to provide the track layout as in figure 8.2b, but also to put one capacitor adjacent to each IC. Clearly, this can be achieved by having one capacitor for each pair of ICs.

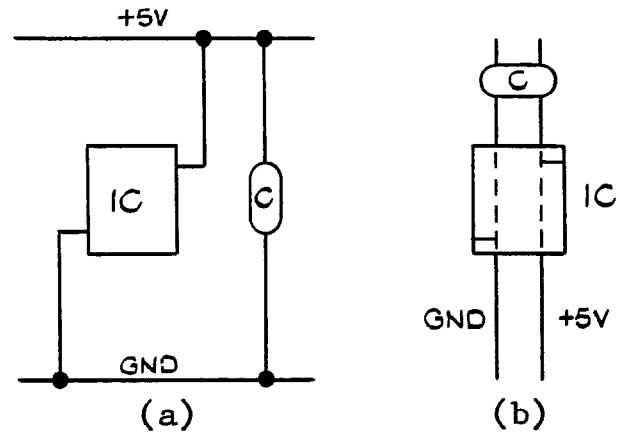

Figure 8.2

Decoupling Capacitors

The foregoing argument shows that the capacitor is better thought of as a reservoir capacitor which supplies the local, instantaneous current. Some manufacturers specify capacitors for IC decoupling by giving the maximum pulse risetime, which corresponds to a maximum current for a given size of capacitor. For instance, a 47 nF capacitor specified at 50 V/μs can supply a current given by

$$i = C\frac{dv}{dt} = 47 \times 10^{-9} \times \frac{50}{10^{-6}} = 2.5 \text{ A}$$

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 500 mA; to be safe, assume that this demand lasts for 10 ns, and design for a voltage drop at the capacitor of 50 mV. Thus

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

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

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

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

It might be thought that radio frequency type capacitors are necessary for TTL decoupling, but this is not so. To show why will be explained in a later volume. Briefly, 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 use non-ceramic capacitors. 1 μF 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. (A package at the centre of a power bus will see two transmission lines in parallel and hence half the impedance. We shall adopt the worse figure for the purpose of this argument.) 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 $Z_0 = \sqrt{(L/C)}$ where L and C are the inductance and capacitance per unit length respectively. To calculate Z_0 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 2 mm and 0.5 mm. Therefore

$$L = 0.6\ \mu\text{H/m}$$

If a 100 nF capacitor is placed every 50 mm along this line, then

$$C = 100 \times 20\ \text{nF/m} = 2\ \mu\text{F/m}$$

therefore

$$Z \approx 0.5\ \Omega$$

An instantaneous current demand of 200 mA – corresponding to 4 gates switching – will produce a voltage transient of 100 mV. 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 – 10 nF for every few ICs, which is not uncommon – will create a situation much worse than this.

Auto-decoupling in TTL

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 paradox, consider the structure of the TTL gate output circuit, when this is driving the gate input low shown in figure 8.3.

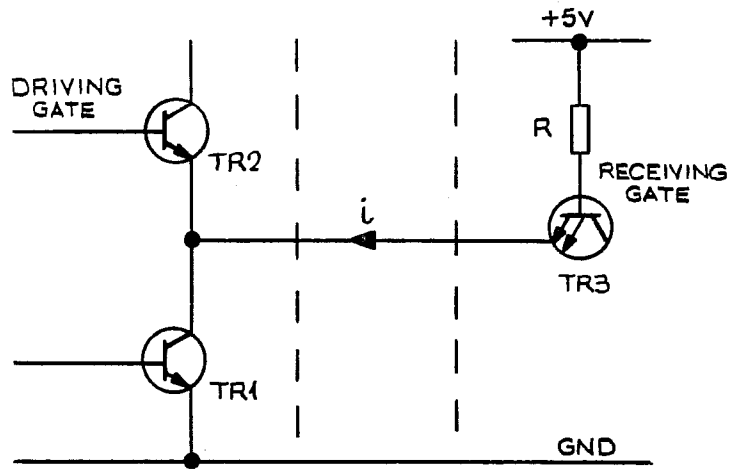

Figure 8.3

According to the specification for, say, a 7400, the typical values of i and R are 1.0 mA and 4 kΩ respectively. When the gate output is low it sinks a current i, given by

$$i = \frac{V_{cc} - V_{be} - V_{ce\ sat}}{R}$$

where V_{be} is the base-emitter voltage of TR3 and $V_{ce_{sat}}$ is the collector saturation voltage of TR1.