Another EMC resource
from EMC Standards

PCB Layout Techniques for Low Cost EMC

Helping you solve your EMC problems

PCB design and layout techniques
for lowest-cost EMC compliance, and signal integrity

(Printed in the IEE's Electronics and Communications Engineering Journal in August and October 1999)

Eur Ing M. K. Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, UK
Phone: +44 (0)1457 871 605 Fax: +44 (0)1457 820 145 Email: karmstrong@iee.org

Abstract

The requirement to comply with the Electromagnetic Compatibility (EMC) Directive 89/336/EEC [1] [2] for almost every electrical and/or electronic equipment or system sold in the European Union is focusing many minds on the best ways to achieve this in the most cost-effective manner.

Printed circuit board (PCB) technology is very widely used in such equipment and systems, and there is a very large body of experience which shows that it is possible to incorporate good EMC practices during their design.

These PCB-level EMC practices usually help achieve the required EMC performance at much lower cost than alternative EMC measures at higher levels of integration, such as whole-product shielding.

These EMC best-practices also improve signal integrity, so help reduce the number of iterations in product development and improve time-to-market. As modern electronic circuits increase their digital clock or analogue oscillator frequencies (especially above 50 MHz), these EMC practices become increasingly necessary for functionality.

EMC design is a complex topic, but the proven best EMC practices for generalised PCB layout can be fairly simply stated, and this is the purpose of this paper.

The practical "what" and "how" of these techniques is the subject of this paper, rather than the "why" (although understanding the "why" will of course help any practising engineer to use these techniques to the full, and compromise where necessary).

Contents

1.   Introduction

In the context of this paper, EMC includes all relevant electromagnetic (EM) emissions and immunity phenomena as they concern all types of analogue and digital circuits.

The best-practice EMC techniques described in this paper interact with each other to give dramatic improvements in the EMC of electronic circuits which use PCBs. Good control of EMC up to 1GHz will only be realised on PCBs by fully implementing all these techniques, and the cost-effectiveness of using all of them together generally exceeds that of other methods. Note that there are also requirements in some market areas for EMC at frequencies exceeding 1GHz.

For EM interference to occur (i.e. inadequate EMC) there must be a source of EM disturbances (an aggressor), a coupling path, and a circuit which is susceptible to EM disturbances (the victim). Coupling out of or into a circuit can occur through:

- The finite impedance of shared circuits (most often AC mains, DC power rails, earth and 0V references)

- Capacitive coupling, where the voltages on one conductor give rise to currents in the victim conductor

- Inductive coupling, where the currents flowing in one circuit give rise to voltages in the victim circuit

- EM coupling (mostly at high frequencies) where the voltages *or* currents in one circuit give rise to EM fields which in turn give rise to currents *and* voltages in the victim circuit.

The techniques described here mostly improve the PCB's EM coupling mechanisms and do little for the actual sources, or for the susceptible circuits themselves (although techniques do exist for improving these too).

Because they operate on the coupling between the circuits on a PCB, and between the PCB circuits and their external environment, these techniques are equally effective for improving both EM emissions and immunity. A valuable side-effect is that they help to maintain signal integrity in both analogue and digital circuits, meaning that it will take fewer design iterations of hardware and software before a product functions adequately, saving significant engineering costs and reducing time-to-market.

To create competitive products it is necessary to use these proven good EMC techniques from the start of a project, continuing all through it. If it is not thought commercially sensible to fully implement *all* these techniques, additional project costs and lengthened timescales should be allowed for extra EMC testing and extra design iterations.

2.   Circuit segregation

In a cost-efficient design process this is the first technique to be applied, and it needs to be employed right at the beginning of the real design process. The layout of the PCB cannot begin until it is known where (physically) any shielding and filtering techniques need to be applied, so this overview of the mechanical assembly should precede PCB design, and should be done as early in the product development lifecycle as possible.

To employ this technique the following areas are first identified:

Outside World: Any part of the circuit for which total control of the EM environment is not available to the product designers.

Inside World: Any part of the circuit for which total control of the EM environment is achieved by the product designers.

2.1   The boundary between outside and inside worlds

Where a product has an adequately shielded enclosure, the outside world is easily identified as being everything outside that enclosure. But where a shielded enclosure is not used it can be quite difficult to identify the boundary between the outside and inside worlds.

It is generally undesirable for all but the most benign circuits to be exposed to the outside world environment. It is often possible to fit PCBs inside screened "module" enclosures of their own, to create their own inside worlds. PCBs designed using all the practices described here will have a good degree of control over their local EM environment, and some types of products will be able to create their own inside worlds even without module or enclosure shielding.

Conductors which run outside of a product's enclosure are clearly subject to the full outside world EM environment, but cables which run inside a product may also suffer a subset of those phenomena if the product enclosure is not adequately shielded. For example, a ribbon cable or jumper strip connecting two shielded PCBs will not be protected from the outside-world radiated RF environment unless there is an enclosure that provides adequate shielding. The frequencies at which internal cables start to act as antennas for outside world radiated fields depend upon their type and length, and upon the circuitry and filters they are connected to.

The use of a single-PCB for all the circuitry in a product is usually the most cost-effective way to meet EMC requirements. This is because it is easier and least costly to control the EM environment of a single PCB, with its obvious boundary between "inside" and "outside" worlds, than it is to control that of a product with several PCBs and internal wires and cables. Many types of electronic products will not need a shielded enclosure if made using a single PCB with no internal wires and cables and the techniques described here. This promises to save considerable costs in both materials and assembly, and to allow a great deal more aesthetic freedom in styling.

Considerable manufacturing cost savings and improvements in reliability result from a single PCB approach, even where EMC is not an issue. Where a single flat PCB is impractical, flexi-rigid techniques are increasingly cost-effective, for EMC purposes, at realising complex three-dimensional circuit arrangements as if they were a single PCB.

2.2   Boundaries within an inside world

When the circuitry which is in the inside world has been determined, it should be further subdivided into dirty, high-speed, noisy, (etc) potentially aggressive areas, and clean, sensitive, quiet, (etc) potential victim areas.

All of the above terms are poorly defined jargon used by electronic engineers to describe the likelihood of one circuit emitting EM disturbances, or being interfered with by EM disturbances. A better rule of thumb for ranking the likelihood of a circuit node to *emit* "noise" is to refer to its maximum dV/dt and dI/dt (rates of change of voltage and current). Another rule of thumb for the likelihood of a circuit node to *suffer from* interference is to refer to its signal levels and noise margins (lower level signals and/or lower noise margins implies greater potential for susceptibility). This analysis should result in identification of the worst-case sources and victims of EM disturbance.
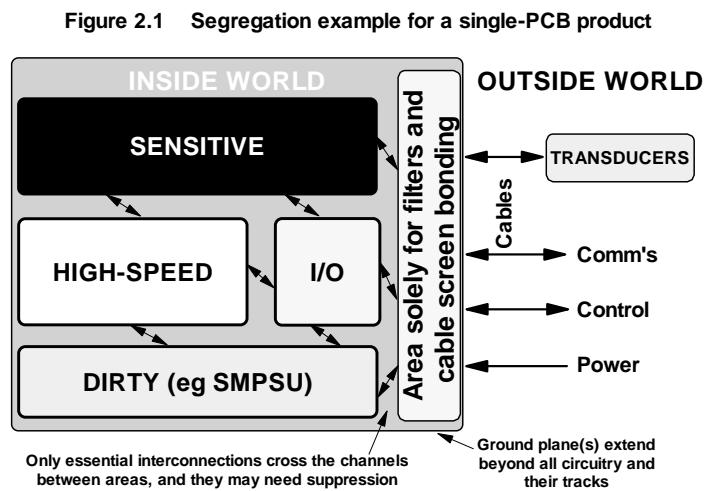
2.3   Segregation

The various inside world areas identified are to be physically segregated from each other, and from the outside world. This activity should segregate all these areas mechanically and electrically, starting at the earliest design phase by showing the various areas clearly on all drawings. This is usually done by drawing dotted lines around rectangular areas each consisting solely of a segregated portion of the overall circuit.

It is necessary to ensure that this segregation is maintained throughout the rest of the design process including system design, PCB layout, wiring harness design, mechanical packaging, etc. Showing the segregation clearly on all system, wiring, and circuit diagrams will prove to be a great help to communications between the electronic designers and others who may not have the necessary EM skills – such as mechanical designers and PCB layout persons – especially where such work is contracted to industrial designers or PCB layout companies.

Most design occurs in two dimensions. It is not uncommon to find that in the final assembly, a PCB carrying a very sensitive circuit (such as a thermocouple or microphone amplifier) turns out to be in close proximity to a PCB carrying a very noisy circuit (such as a switch-mode power convertor), with consequent signal quality problems. Such unpleasant and time-consuming three-dimensional assembly issues should be avoided by detailed visualisation of the final assembly from the start, even before the circuits are designed and the PCBs laid out.

Figure 2.1 is typical of a good segregation practice applied to a product, which uses a single PCB, whether it has an overall enclosure shield, or not. It is generally not too difficult to make such products EMC compliant without enclosure shielding as long as they use analogue and digital technologies that are not too noisy (e.g. use clocks under 32MHz) and not too sensitive (e.g. no more sensitive than 12-bit accuracy, or ±0.012%, for analogue signals).

This example shows that the segregated area, where the outside world interface suppression components are fitted, is positioned along one edge of the PCB forming a physical boundary between the outside and inside worlds. This area would only contain opto-isolators, isolating transformers, baluns, filters, transient absorbers, and similar interfacing components. It would also contain bonding points for the screens of any screened cables, and/or for any enclosure shielding. If this example PCB was part of a larger assembly, the segregation techniques employed for best EMC would be just the same.

Figure 2.1   Segregation example for a single-PCB product

Why restrict all interface components to one edge of the PCB? All cables and metalwork in the outside world carry unwanted, potentially interfering, currents. Some are due to their antenna-like propensity to pick up radiated EM fields, while others are caused by voltage differences between the earths of the units the cables connect to. Some of these unwanted external cable currents wish to flow from one cable into another, whilst others wish to flow from their cables to their local earth. So, having all the outside world screen connections and filters along a single edge of the PCB encourages these unwanted currents to restrict themselves to that area of the PCB, and discourages them from flowing through the areas dedicated to other circuitry, helping to prevent these external currents from upsetting circuit operation.

Where an effective enclosure shield is used, the boundary line between inside and outside worlds moves from the PCB to the wall of the enclosure. All of the related interface suppression components and cable screen bonding must consequently move from using the PCB as their ground reference to using a connector panel which is set in the wall of

the enclosure. If this is not done the shielding integrity of the enclosure will be compromised by leakage of unwanted currents through the enclosure shield by the cables. A single area for all the interconnections is still best practice.

There is now available a wide range of PCB-mounted screened and/or filtered connectors that can also bond to a metal panel. This makes it possible for many products to use the same PCB layout as shown above even when using an enclosure shield. Instead of mounting the connectors on a panel and wiring them to the PCB, they are soldered onto the PCB and then bonded to the connector panel during final assembly. The bodies of the screened connectors will bond the PCB ground plane in the interconnection area to the connector panel (enclosure shield) at multiple points. A compliant conductive gasket may be used between the connector bodies and the metal of the connector panel to avoid using multiple screw fixings. This is a most cost-effective way to design and assemble the interconnection area for a product with a shielded enclosure.

Within the PCB itself, a narrow "channel" free of components should be created between each of the segregated circuit areas on the PCB, to help avoid coupling between the different circuit areas. This channel should be wide enough for the fitting of a PCB-mounted shield including bonding it to a PCB ground plane at frequent intervals (say, every 12mm) along all its edges. Prototype PCBs should provide ground-plane bonding points in these channels to facilitate fitting local "tuner can" shields to noisy or susceptible circuit areas, should they turn out to be required by functional or EMC testing.

2.4   Component placement and routing of tracks

All relevant components and their tracks must be contained within their designated PCB area. The only tracks that may exit or enter any area are those that have no choice but to connect to other areas of the PCB. Wires which are free to move around, or which may be routed differently by assembly staff, can cause EMC problems in serial manufacture when they fall close to a circuit area they are not connected to. If it is not possible to eliminate all use of wires inside a product, care is needed to ensure that the lengths and routes of any remaining wires are fixed.

The most noisy or susceptible components in each area should be positioned first, as close to the centre of their areas as possible and as far away from any cables or wires as possible. Examples of such components include clock generators and distribution (extremely noisy); bussed digital ICs (very noisy); microcontrollers (noisy); switch-mode power transistors and rectifiers and their chokes, transformers, and heatsinks (all very noisy), analogue ICs (sensitive), and millivolt level amplifiers (very sensitive).

It is not commonly realised that analogue devices such as operational amplifiers can be extremely susceptible to interference, even up to 1GHz when their fastest slew-rate may be as low as $1V/\mu s$. This susceptibility can easily occur even in integrator circuits with very slow (e.g. several seconds) time constants.

Digital clock distributions (being very aggressive signals) must be the first "nets" to be routed, and must be run on a single PCB layer adjacent to a ground plane (see section 4). Clock tracks must be kept as short as possible and even then may require the use of transmission-line techniques (described later). It may be necessary to experiment with a variety of component placements to achieve the minimum lengths of clock tracks. Where clock tracks have to be lengthened to minimise skew between different devices, a "serpentine" layout is preferred.

Digital busses and high-speed I/O should be routed next, and should be treated exactly the same as clock tracks, deferring to clock tracks where there is any conflict of design rules. Very susceptible tracks, such as those carrying millivolt transducer signals, should also be routed as if they were clock or data buss tracks, although they will always be in a different segregated area of the PCB.

All other types of analogue, digital, and power signals should also be routed according to how aggressive they are at emitting EM disturbances, or how sensitive they are to being interfered with by them. Where these characteristics are not obvious from a circuit analysis, probing a prototype with a wide-band oscilloscope (and/or spectrum analyser) with both voltage or current probes will reveal which are the most aggressive, and injecting voltages or currents from a wideband sweep generator will reveal which are most sensitive.

There should be no layer changes for any digital clock or other very aggressive tracks, and also for very sensitive tracks (i.e. these tracks should not have any via holes). Where there is no reasonable alternative to changing the layers of such tracks, the techniques described in the last paragraph of section 6.1 should be used.

2.5   Communicating with PCB layout personnel

PCB layout personnel need to understand that the segregation applies not only to the components, but also to all tracks and other conductors. Only tracks or conductors that must interconnect segregated areas are allowed to cross the channels between them.

It is best to check that segregation instructions have been followed on *draft* PCB layouts, well before PCB manufacture. An easy check is to count the tracks and other conductors which cross the dotted lines showing the segregated areas on

the circuit diagram – there should be exactly the same number crossing the channels between areas on the draft PCB layout. Where PCBs have been autorouted it is usual to find additional tracks crossing area boundaries – these are often the source of much design heartache, so eliminate them right away by applying more skill to the track layout.

3. Interface Suppression

This is the second technique which should be employed during a cost-efficient EMC design procedure.

3.1 Overview of the technique

EM disturbances can be radiated and/or conducted across interfaces between segregated areas. Shielding, filtering, and isolation techniques (such as opto-coupling) are used to reduce this coupling to acceptable levels.

To decide on the most cost-effective methods to apply at each interface, it is necessary to analyse them for all the EMC phenomena possible, given the operational EM environment, circuits and components.

It is important not to ignore internal power supplies and other common connections such as 0Vs or grounds. Circuit designers usually abbreviate such connections on their circuit diagrams, often to the point of invisibility. This leads to inevitable problems when EM disturbances coupling via overlooked common connections is not anticipated and so not suppressed, and is a frequent cause of additional PCB layout iterations. A common error of this type is the powering of both sides of an opto-isolator from the same power supply rails.

3.2 Suppressing outside – inside world interfaces

Conductors passing from outside to inside worlds may need to employ the full range of suppression techniques – shielding, filters, isolating transformers, opto-isolators, surge protection devices, etc. – to cope with all the EMC threats and susceptibilities appropriate to the EM environment, the EMC standards it is hoped to declare compliance to, and the product's internal design. As described above, best practice is to use a single PCB area for all outside-inside world interconnections.

Visual displays (such as LCDs, LEDs, VDUs, moving-coil meters, etc.) and controls (such as pushbuttons, potentiometers, rotary knobs, etc.) are usually exposed at an interface between outside and inside worlds, because they need to be connected to the circuit and also be seen and touched by operators. They experience a different balance of EM phenomena from the cables discussed above, and personnel electro-static discharge (ESD) is often an important consideration. Once again, numerous suppression techniques may need to be employed depending on the circumstances. For example, a small LED mounted in a hole in a plastic front panel provides an easy route for ESD to enter a circuit by sparking to the LED's leads and into vulnerable circuitry. Solutions could include covering the LED with a clear plastic window, using an earthed metal front panel to divert the ESD, or fitting surge protection devices to each of the LED's leads.

It may be necessary to shield certain of the areas of the PCB, the whole PCB, a whole sub-assembly of PCBs, the entire assembly of PCBs, or the entire product (in ascending order of cost and difficulty). Following the segregation methods described above makes lower-cost and easier shielding possible, because the component-free channels between the segregated circuit areas make it relatively easy to fit PCB-level shielding (such as "tuner cans") to especially noisy or sensitive areas.

It is not unusual, when testing a finished product, to find that a small part of its circuit needs shielding, but the layout of the PCBs is such that a complete redesign would be needed to be able to use low-cost shielding methods. Planning ahead for low-cost shielding, as suggested above, can save a great deal.

3.3 Interfaces between dirty/high speed/noisy and clean/sensitive/quiet areas

The type and amount of suppression required to be applied to tracks and other conductors between different PCB areas should be determined by an analysis of the desired signals. This includes the unwanted noises that may be inadvertently communicated along them and the emissions and sensitivity of the circuits they connect to.

Digital clocks and data busses are very aggressive noise emitters and should not be allowed into clean/sensitive/quiet areas at all. Data intended for a sensitive area should be latched from the bus no closer than the boundary of that area, and the data busses themselves restricted to a noisier area. (Clocks and data busses should also be as short as possible.)
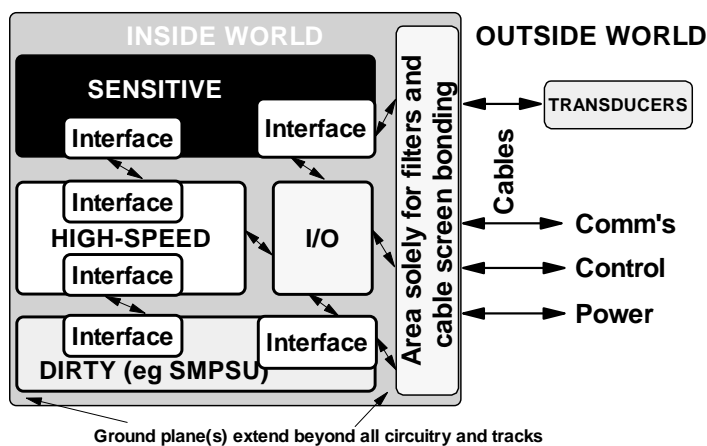
Often overlooked routes for conducted noise carried from one segregated area to another are power distribution networks, "static" data lines, and other low-frequency signals. Many an analogue circuit has suffered from noise carried on its power supply rails from a switch-mode power supply. Many an analogue circuit has also suffered increased noise levels carried by digital control lines – despite remaining static at logic 1 or 0 they are usually polluted by high-frequency noise

originating from internal switching operations in digital ICs. It may be necessary to fit filters or isolators to such inter-area connections.

Components that interface between segregated areas, such as analogue-to-digital convertors, transformers, data bus latches, filters, isolators, and the like, should be positioned at a common edge of the areas they interconnect. They should generally remain wholly within one area or the other to keep the inter-area channels free of components so that shielding can easily be fitted. Their tracks must be routed directly into their respective areas, and not intermingle with tracks that remain within the area they exit from. Figure 3.1 shows the previous example (Figure 2.1) with its internal interfacing component areas.

The purpose of keeping the channel component-free is to make it easier to fit shielding over the segregated areas of circuitry, should it be needed. Where interface components like ferrite beads, common-mode chokes, or opto-isolators are placed *in* one of these channels it can help to achieve good separation between the tracks associated with each circuit area, but the cut-outs they require in any PCB-mounted shield may compromise its shielding effectiveness.

Figure 3.1 Siting the interfaces between the segregated areas

This compromise between the need for good track segregation and for shielding effectiveness does not apply when 'feedthrough' filter components are fitted as interconnections between segregated areas. These are designed to fit into and actually penetrate the walls of screened enclosures, so when fitted in (otherwise component- and track-free channels) they encourage good track segregation and don't compromise shielding effectiveness. Traditionally, feedthrough filters are screwed or soldered into a hole in the shield, with wires connecting to their ends. This does not suit robotic surface-mounted assembly techniques, prompting some manufacturers to produce 'SMD feedthrough filters'. These generally have an earth electrode around their centre, which is intended to be soldered to the PCB reference plane (although some types may also be able to be hand-soldered to a cut-out in their shield). In general the small size and low-profile of these parts means that they only require a very small cut-out in the shield they penetrate, and so may be expected to have little effect on shielding effectiveness. Where SMD 'feedthroughs' are used, their performance will be improved if the shield they are associated with is soldered to the PCB reference plane as close as possible to the SMD feedthroughs, as frequently as can be achieved. Very stringent applications sometimes require PCB shields to be seam-soldered all around their circumference, and such assemblies would probably need to use the more traditional wired-in feedthrough devices.

Radiated interference between segregated areas is also possible. The placement within an area of the most aggressive or sensitive components, or larger components such as large-value capacitors, is critical. If positioned near the edge of an area, the local EM fields unavoidably present around components can more easily couple with components and tracks in their neighbouring area, possibly causing a range of problems. The use of small sized, low-profile components, sensibly placed where they have the least opportunity to couple to neighbouring areas, can often avoid the costs of shielding.

Nevertheless, separating shields such as guard tracks or screening boxes around (or over) some areas may be necessary. This is made relatively easy without major design iteration by the presence the component-free channels between segregated areas discussed earlier. Where the channels in prototype PCBs are already fitted with via holes making frequent connections to the ground plane, low-cost shielding can easily be fitted if needed.

3.4   Details of suppression techniques

Available suppression techniques include:

- common-mode and/or differential mode filtering
- opto-isolation
- transformer isolation
- communications protocols (to improve bit error rate in the presence of interference)
- surge protection devices
- the use of balanced drive and receive signals (instead of "single ended")
- the use of fibre-optic, infra-red, wireless, laser, or microwave communications instead of copper cables.

All of these are circuit or system design issues which are outside the scope of this paper. There are many excellent references for these suppression topics, but it is important to realise that on a PCB only a ground plane (as described

next) can provide a good enough RF reference ground to allow the full performance of filters, cable screens, and internal shields to be achieved. Providing a good RF connection to an effective local RF reference plane is one of the secrets of EMC "black magic", without which a perfectly designed, simulated, and expensive filter may provide no benefits, or even make EMC worse.

4. Ground and power planes

It is not possible to use a length of track or wire (or star grounding) as an equipotential ground except for the very lowest frequencies (usually below 1MHz). Only a solid area of conductor can provide a good ground up to 1GHz (and beyond), and these are called ground planes.

Ground and power plane techniques allow dramatic reductions in all unwanted EM coupling when used in conjunction with the other techniques described here. Ground planes are also essential for all the other techniques to function properly.

4.1 Ground and power plane techniques

A high-quality high-frequency ground plane can easily be created by devoting one layer of the PCB to a solid copper sheet, and using it as the 0V and RF reference ground for all the circuits. Ground planes must lie under all their components and all their associated tracks. Power planes may also be created in the same way as ground planes to provide similar benefits for other common connections. The segregation and interface suppression techniques of sections 2 and 3 above must be followed even where a common ground plane is used for a number (or all) of the circuit areas.

High-performance circuits using PCBs with more than four layers may need to devote a number of layers to ground planes to use with transmissions lines and for other EMC reasons described later. PCB areas can also be divided up into different ground plane areas, each serving a segregated circuit area as described in section 2. In these cases the interconnections between the ground planes can become quite complicated if high-frequency performance is to be maintained, and this is described later.
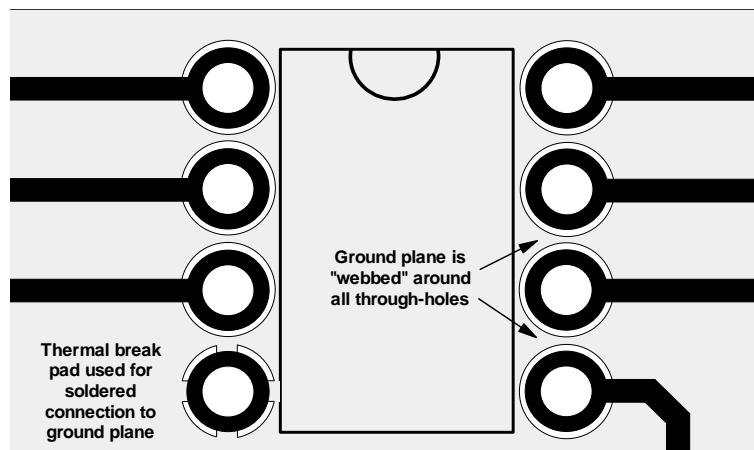
Most practical PCBs are perforated with via holes, even if they are not also perforated with through-holes for leaded components. These perforations increase the inductance of a ground plane, making it less effective at higher frequencies. "Buried via" techniques have had to be developed for the PCBs used in cellphones, allowing interconnections to occur between tracking layers without the via needing to also perforate the ground plane. The general rule is that any gaps must have dimensions of 0.01λ or less at the maximum frequency concerned. For a good ground plane at 1GHz, (e.g. to help meet EU harmonised EMC standards cost-effectively) no ground-plane gaps should have any dimensions greater than 1.5mm. "Sneaking" tracks into a ground plane layer to make layout easier is therefore not allowed.

Where there are a number of closely-spaced vias or through-holes there is danger that their clearance holes through a plane will merge together, creating a gap longer than 1.5mm. PCB design rules should size clearance holes so that for regular hole spacings such as DIL packages the plane's copper forms "webs" in between each pair of holes, as shown by Figure 4.1.

The minimum thickness of these webs generally needs to be between 7 and 10 thousands of an inch to suit low-cost PCB manufacturing processes. Of course, it would be better to use SMD rather than leaded components, both to reduce the perforation of the planes and so that the IC is closer to the ground plane and benefits more from the shielding effect it creates. Where hole spacings are not fixed by leaded components, but are manually or automatically placed, they should not be placed so close together as to create gaps in planes by preventing "webbing" between the holes.

Figure 4.1 Example of a webbed ground plane under a leaded IC

Ground plane is "webbed" around all through-holes

Thermal break pad used for soldered connection to ground plane

PCB tracks or wires cannot provide any sort of ground integrity at high frequencies, e.g. a 5mm length of a 1mm diameter wire, or of a 1mm wide track on a PCB, has an impedance at 1GHz close to 30Ω (using the 1nH per millimetre rule of thumb for inductance). Every doubling of the width of a track will reduce its impedance by only around 20%. Consequently, "ground fills", "ground tracks", "guard rings", and all the other traditional ground techniques used on PCB layers which carry signal or power tracks are all completely ineffective at high frequencies, and at high frequencies might

even resonate and make EMC and signal integrity worse. Some PCB CAD packages claim to automate the creation of planes, but actually only create useless "area fills".

Tracks, area fills, guard rings, etc. used as grounds on signal layers *can* be used to good effect at high frequencies – but *only* when bonded to an underlying ground plane with at least one via hole every 5 to15mm (using a random allocation of spacings). Signal tracks will get the best benefits from the proximity of a plane if they have a width of at least 0.7 times (preferably 1 or 2 times) the thickness of the PCB material between them and the plane.

Ground planes should extend as far as possible beyond the boundaries of components and their tracks and power planes. [6] recommends using a "20H rule": ground planes should extend beyond power planes and any tracks by at least 20 times their layer spacing (H). High-speed components (such as digital clocks, processors, and memory) and their tracks should always be placed near the centre of their segregated areas, well away from the edges of their ground and power planes. Where it is essential to route a high-speed track close to the edge of its respective ground plane, a guard track should be routed just outside it and on the same layer (bonded frequently with via holes to the edge of the ground plane as described above.

. Spacing between power and ground planes should be minimised, so they are best run on adjacent PCB layers. Section 5 will show that this has benefits for power supply decoupling. [7] shows that a spacing of 0.2mm can reduce emissions by 15dB compared with a spacing of 1mm.

All 0V and power connections must bond directly to their respective ground and power planes. The 0V and power leads of through-hole components must have their through-plated holes directly connected into their planes using thermal-break pads (sometimes called wagon-wheel pads) to allow their automated soldering, as shown by Figure 4.1. The 0V and power pads of surface mounted devices (SMDs) must be connected to via holes which connect directly into the respective ground and power planes. This is easy for flow-soldered SMD, but when reflow soldering the vias can suck the solder away from the pad, leaving a dry joint, so plane vias are not usually included in the SMD pads but are positioned nearby and tracked to the pad. The inductance of these tracks has a deleterious effect, so it is vital to use the shortest and fattest tracks possible, whilst taking account of reflow soldering requirements.

Figure 4.2 compares alternative methods for connecting SMD components to planes. Where very high-frequency performance is required multiple plane vias and short fat tracks will help by reducing the inevitable inductance of the tracks and vias. Using an over-sized SMD pad with multiple vias to the plane actually in the pad itself (tented over by solder resist to avoid removing solder from the joint) makes a very significant improvement to the high-frequency performance of plane connections [3].

Where ground plane connections do not need to be soldered (typical of reflow-soldered SMD components), solid connections to planes should be used. This will give better EMC performance than the use of thermal break connections.

Figure 4.2 Comparison of ground and power plane connection methods for reflow soldered SMD components

(Flow soldered SMDs may use several vias to plane *underneath the component terminations themselves* – giving the best possible high-frequency performance)
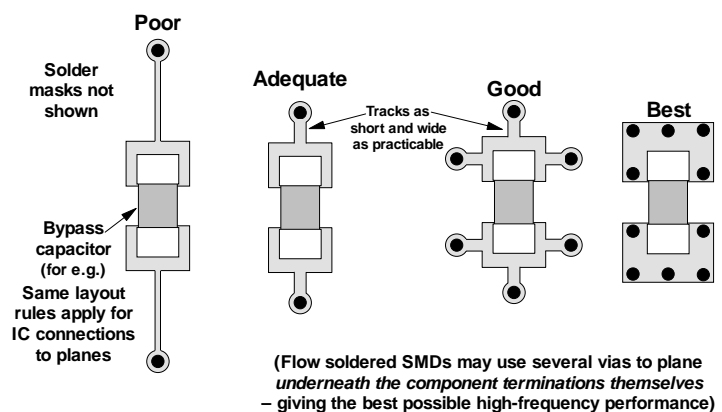
Rectangular ground planes are preferred, as long as their aspect ratio is not extreme (although square plane areas should generally be avoided to minimise the effects of resonances in the plane). Rectangular shapes perform better than exotic shaped ground plane areas or thin, extended, areas, and also aid the fitting of PCB-level shielding if it is found to be needed. Sometimes several power planes are placed on one PCB layer, in which case they may need to have complex shapes (avoiding any long or thin areas).

4.2   Benefit of using low-profile components with ground planes

Ground planes also provide a substantial degree of volumetric shielding close to their surface and combining them with low-profile components (e.g. SMD, BGA, TAB, chip and wire, etc.) is strongly recommended for the best EMC performance (emissions and immunity).

For example, axial electrolytic capacitors lying close to the PCB ground plane are better than taller radial types. To benefit from a useful shielding effect up to 1GHz, no part of any PCB-mounted components should exceed 10mm above the ground plane layer, with ≤3mm for especially high-speed (i.e. most digital clocks, processing, and memory) or

sensitive components. No IC sockets should be used, and this often means using field-programmable memory ICs in place of socketed PROMs or ROMs.

Increasingly, modern digital products are only able to function at all by using ground planes and SMD devices. Larger digital ICs are best used in flip-chip, tape automated bonding (TAB), or ball-grid array (BGA) packages. Quad flat-pack (QFP), J-lead, and pin-grid-array (PGA) devices are not too much of a liability for EMC if they are soldered directly to the PCB.

4.3   Splitting planes

Split planes may give either better or worse performance than unsplit planes, depending strongly on the details of the PCB layout and circuit design [8], [9]. It may be hard to determine in advance which will give the best performance, so it is recommended here that both options are allowed for, on prototype PCBs at least, by gapping all planes at the natural boundaries between the segregated circuit areas, but also providing the means to "stitch" them together with pairs of via holes placed along the split every 10mm or so (random spacing allocation of 5 to 15mm is preferred). The inevitable stray capacitance "shorts out" the split at very high frequencies, so the benefits of splitting only exist below 1GHz.

The pairs of vias across plane splits can be linked directly together with short wires, or with small capacitors, so it helps to pitch the via pairs so that small capacitors or "zero-ohm links" may readily be fitted (preferably SMD types). The use of capacitors instead of direct bonds across splits allows the achievement of "star grounding" techniques for wanted low-frequency signals, while achieving the effect of a single overall plane for the best high-frequency performance. Capacitors all have self-resonances (described in section 5) which can worsen performance, but taking their self-resonances into account can be used to improve their performance.

The author has had considerable success with improving signal integrity and noise in mixed technology PCBs using powerful digital processing and sensitive analogue circuits by using a common ground plane for all circuits, but most EMC debate revolves around whether to split the ground plane for the I/O interconnection area (see Figures 2.1 and 3.1). Common-mode noise often exists on ground planes, especially for PCBs with digital processing. This is not an issue for the circuits on that PCB, but it is an EMC problem for any wires and cables exiting the PCB. Providing a "clean ground plane" for the interconnection devices and filters is one way of dealing with this. An I/O ground plane split off from the main ground plane must be directly connected to the chassis or shielded enclosure of the connector panel area at multiple points spread over its area (or else by a conductive gasket).
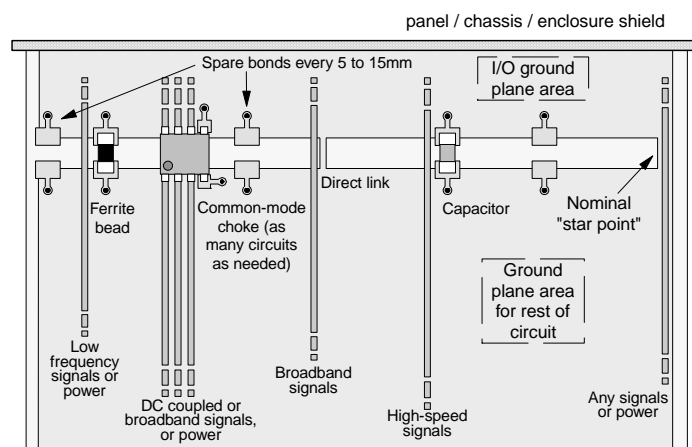
For a split I/O area ground plane to provide the required benefits, all tracks crossing the split *must* have carefully-defined return current paths physically adjacent to them, as shown by Figure 4.3.

The return path should, if practical, be limited to the bandwidth of the wanted signal. For high-speed signals with no low frequency content the return path can be a capacitor (chosen to have a suitably low impedance over the frequency range of the signal) connected across the split between the ground planes, adjacent to the signal track. For low-speed signals with no high-frequency content the return path can be a ferrite bead, also bridging the planes and adjacent to the signal track. Some high-rate data streams can also need DC coupling because their data can include low-frequency patterns, and such broadband signals

Figure 4.3   Crossing a split between ground planes

(and Dc power supplies) should use an adjacent direct link across the split planes. Balanced signals should need no ground return path, but this depends on the degree of their balance and some adjacent ground bridging may prove necessary. Common-mode chokes fitted to signals and their ground returns may also be used, and may provide advantages over direct links for DC power or broadband signals.

Galvanically isolated planes (e.g. associated with an isolation transformer or opto-coupler) will usually benefit from being bridged by a number of low-value capacitors to the main ground plane. Where the isolation is for safety reasons, only suitably rated safety approved capacitors should be used, and care should be taken not to exceed earth leakage safety requirements.

Testing a prototype PCB for emissions (or immunity) with a close-field EM probe helps obtain the best EMC performance from a split I/O ground plane, although only testing the (almost) finished product in an EMC test laboratory can give definitive results. The first part of the exercise is to make sure that the ground plane bridges associated with the

tracks that cross the split are of the optimum type for both the signal and EMC. When this is completed, additional capacitors or links should be added at the spare bond positions to optimise EMC further. Clearly, this can be a time-consuming process if the product needs to be dismantled to make the changes to the PCB.

4.4   Connecting ground planes to chassis

PCB components and tracks have rapidly fluctuating voltages on them. These will experience weak capacitive coupling to everything in their vicinity (if they are not already fitted with shields soldered to their ground planes). Weak currents inevitably flow in these parasitic capacitances, one cause of common-mode emissions. Very noisy circuits (such as modern computer motherboards), or especially susceptible circuits (e.g. >12 bit analogue resolution), generally need at least a nearby metal plate, and preferably a fully shielded enclosure, to catch and return these leaked currents before they can interact with the outside world.

When such metalwork (or plated plastic) has been provided, it must be connected to the ground plane of the PCB at many points to encourage as much as possible of the leaked currents to return quickly to their origin on the PCB. Mechanical mountings are usually used for these "chassis bonds", but they must be as short as possible (stand-off pillars taller than 4mm are not recommended). There should be one chassis bond in the centre of each area of high-speed circuitry, especially clock generators and their distribution. [6] recommends a $\lambda/20$ rule for the spacing of chassis bonds: maximum spacing between bonds should not exceed $^1/_{20}$ of the wavelength at the highest frequency of interest. For 1 GHz this rule would result in chassis bonds every 15mm, which would clearly create significant track layout and assembly difficulties. For high-speed digital boards a ground plane to chassis bond every 75 to 100mm is recommended.

It is usual to bond the ground planes of digital circuits to the local chassis at each mechanical mounting using an M3 or larger screw in an appropriately-sized plated-through via hole. Where there are concerns about this technique, or where the ground plane must be isolated from chassis for functional reasons (e.g. galvanic isolation), the large vias for the mounting screws should be isolated from the ground planes and provision made for a connection between them to be made using a passive component. Then each individual chassis bond may be connected to its local ground plane with
a link, or a capacitor of appropriate type and value, or even left unconnected.

On PCBs which just carry digital processing, direct links at every mounting point are likely to be best. But for analogue (or mixed analogue/digital) boards the ability to vary the type of ground plane-chassis bond can be useful during product development. This technique makes it possible to fit a link at one bonding point whilst fitting the others with appropriate capacitors, so that at low frequencies there is only a single-point connection between ground plane and chassis (to reduce the flow of DC or 50Hz "earth loop" currents in the ground plane),whilst at very high frequencies the multiple bonds created by the capacitors help improve EMC. Where ground planes must be galvanically isolated from chassis, capacitors to chassis at each generally provide significant EMC benefits, but care should be taken not to exceed any earth-leakage requirements (e.g. for patient-coupled medical apparatus).

A single link or other component can have enough self-inductance to limit its EMC benefits at very high frequencies. For this reason surface mount components (being smaller and hence less inductive) are preferred to leaded types, and even so it may be found that an arrangement of three or more components with equal radial spacing around the screw's via hole is required to reduce their inductance enough to achieve the necessary EMC performance. The track lengths associated with such components should be minimal (refer to Figure 4.2).

4.5   Interconnecting planes in multi-PCB assemblies

Where two or more PCBs are in close proximity and sharing signals (e.g. card cage systems) there will generally be considerable functional performance and EMC advantages in linking their ground planes together at RF. This is achieved by making frequent short interconnections between the cards' planes, uniformly distributed along the *full* length of their common boundary (e.g. via a full-height card connector to a ground plane in the backplane card).

For better RF performance, ground plane interconnecting pins or wires should be shorter, have reduced pitch, and be spread over more of the interconnecting boundary. Using one ground plane link pin alongside every signal or power pin is not excessive, and although it seems expensive it is often the lowest-cost way to improve the emissions and immunity

of a multiple-PCB product. Card-cage and similar systems can also benefit from bonding ground planes together at all card edges, e.g. by means of front panels and/or card guides.

4.6   What if multilayer PCBs are too costly?

It is important to realise that, in volume, four-layer PCBs now only cost between 20% and 50% more than two-layer. Where a ground plane is not used, an enclosure shield (or higher-specification shield) is more likely to be needed, resulting in a higher overall manufacturing cost. Also, the project timescale improvements afforded by the improved signal integrity of these EMC best practices will not materialise. The use of ground planes usually turns out in retrospect to have been the most cost-effective EMC technique, especially when the *overall* financial break-even time and profitability of a product is considered.

Where the scale of silicon integration used is low, and/or where SMD technology is used, it is possible to get all signal and power tracks of a double-sided PCB onto one side, leaving the other side free for a solid ground plane. This is an appropriate technique for two-layer PCBs. When used for digital products, the lack of a power plane might requires a number of ferrite beads (see section 5) and it may not prove to be the most cost-effective construction.

Where restricting signal and power tracks to one layer of a double-sided PCB is not feasible, signal integrity and EMC become more difficult to achieve. However, useful improvements in the EMC performance of two-layer PCBs without ground planes can be had by "gridding" ground tracks between two tracking layers to reduce inductance. This is often done most easily by using a "maximum copper" or "area fill" on the ground tracks of two layers whose tracks run perpendicular to each other. This results in many horizontal ground runs on one side, and many vertical runs on the other. As it is, this is useless – the EMC improvements are achieved by "stitching" these horizontal and vertical runs together with vias wherever they cross, to create a grid over the whole area of the components and their tracks. This gridding technique depends upon keeping the longest diagonals of all the grid openings as short as possible, especially underneath the most sensitive or aggressive components. This is often difficult to achieve for through-hole microprocessors and support chips, but less difficult for SMD types. It is usually necessary to spend considerable time shuffling components and tracks around to achieve the required gridding.

Gridded grounds only provide barely adequate ground integrity up to the frequency where the longest grid diagonal equals $^1/_{20}$ of a wavelength. Remembering that velocity of propagation is roughly 60% of free-space in an FR4 PCB, this implies that a square grid with a 50mm diagonal could only provide a barely effective ground up to 150MHz. Gridded grounds also suffer from high resistive impedance, making them significantly less effective than ground planes at any frequency.

Single-sided PCBs are extremely difficult to make EMC compliant for all emissions and immunity phenomena without expensive external shielding, with the exception of circuits which naturally have low emissions (signals with low values of dV/dt and dI/dt) and which also have naturally high susceptibility (usually due to high signal levels and low impedances).

5.   Power decoupling

This technique must be considered as a part of the circuit design process itself, as well as part of the PCB layout. The aim of power decoupling, crudely put, is to maintain the power supply impedance to each IC on the PCB to 1Ω or less across the entire frequency range of interest for EMC (at least 150kHz to 1GHz). Some high-speed or powerful ICs may require a power supply impedance of 0.1Ω or less over certain frequency ranges, for correct operation. Wires and PCB tracks all have much too much inductance to provide such low impedances.

Decoupling a power rail means connecting capacitors between it and its respective ground (0V) reference. Correct power decoupling keeps current loops in power supply circuits small, thereby reducing inductive and common impedance coupling. It also keeps the voltage fluctuations on power rails low, reducing capacitive coupling. Also, the functional operations of most digital ICs (especially those using higher clock speeds) can depend critically upon the high-frequency performance of their power supplies.

5.1   Power decoupling techniques

A large decoupling capacitor (e.g. 100µF) should be fitted at every point where power supplies enter (or leave) the PCB, and some smaller capacitors (e.g. 10µF tantalum or solid aluminium types) should be "sprinkled" around the PCB on a "µF per unit area" principle, e.g. one such tantalum for every nine square inches of PCB, as well as being positioned near to heavy power usage such as microprocessors, memory, and other VLSI ICs, plus being fitted at every point where power enters or leaves the PCB. These capacitors will help to control lower frequencies, say up to 10MHz.

The power supplies to every IC should be decoupled very close to the device concerned (and to every transistor connected to a power rail), using appropriate capacitor sizes and types (see later). Where an IC has a number of power pins of the same type (e.g. V_{CC}), each should have an appropriate decoupling capacitor mounted as close as possible to it.

Achieving good decoupling performance above 10MHz is more difficult as the frequency increases, because the inductance of component leads, PCB tracks, via holes, and capacitor self-inductance, inevitably limit the RF performance. The achievement of good power supply decoupling at higher frequencies using capacitors mounted close to IC power pins is discussed next.

5.2   Determining the values of the local decoupling capacitors for each IC

The total local decoupling capacitance required for an IC depends on its transient current demands and the power supply's voltage tolerances. VLSI and RAM manufacturers should be able to specify the values (and maybe even the capacitor types and preferred layout patterns) for their products, but note that they will probably have assumed an accurate 5V power supply, not always achieved in real life.

To calculate the value of decoupling capacitance required for an IC, use the formula: $C(\Delta V) = I(\Delta t)$, using the units Farads, Volts, Amps, and seconds.

In this formula, ΔV is the available tolerance on the IC's power supply voltage, derived by subtracting the IC's minimum supply voltage specification for correct operation (from its data sheet) from the minimum power supply voltage (taking account of all the tolerances, regulation parameters, and the voltage drops due to the resistance of the power conductors). This calculation should be done for the worst case operational temperatures, and should take into account temperature and ageing co-efficients of the power supply regulators and their load regulation curves. ΔV often turns out to be a mere +100mV. (A negative figure would indicate potential problems with reliability in the field.)

I is the IC's transient current demand from its power supplies, which lasts for the time Δt. For many modern digital ICs, Δt is 2ns or even less and may be assumed to be the same as the specified low-high or high-low transition times for the digital outputs. I is a parameter almost never provided by device manufacturers. It will usually be estimated or measured in some reasonably sensible way. For many digital devices, the output current cannot be used for I because their "shoot-through" currents (caused by overlapping conduction in totem-pole output stages) are more significant. (Shoot-through currents are known by other names, such as "transient supply current"). These can be an order of magnitude higher than the IC's transient output currents into its capacitive load, although they usually only last for a fraction of the output's transition times.

In the absence of an IC manufacturer's data giving detailed power supply current versus time data, or detailed power supply decoupling specifications, a very high-speed two-channel oscilloscope may be used to measure both I and Δt. Two probes should be used in A-B mode to measure the voltage across a low value SMD resistor inserted into the power supply of the IC concerned (SMD so that it behaves as a resistor up to 1GHz at least). The value of the resistor should be such that the peak voltage measured across the resistor by the A-B probes is no more than ±100mV, and the resistor should be positioned between the IC's power pins and its PCB mounted decoupling capacitors. Some larger ICs such as the Pentium™ may not operate properly if the resistor or its lead lengths are not kept so short that a magnifying glass is needed to solder them in place. For the Pentium™ good quality data on its transient current demands is provided by its manufacturer, Intel. There is no point in trying to measure I and Δt with greater than ±10% accuracy, since this tolerance is less than the typical value tolerance of decoupling capacitor components.

Where ΔV values are low, and decoupling capacitance needs to be correspondingly high, it can be cost-effective to improve the regulation of the power supply, and the resistance of its PCB connections, to increase ΔV and allow the use of smaller decoupling capacitance near to each IC. This can also help achieve the desired EMC, due to the higher self-resonance frequencies of smaller capacitors.

5.3   Self-resonance limitations in capacitors

The typical capacitor value resulting from the above calculation is often in the 1 to 10nF region for small or medium-scale integrated "glue logic" ICs and ECL, and often 100nF or more for VLSI devices such as microprocessors, memory ICs, gate arrays, and the like.

Unfortunately, self-resonance in capacitors prevents them from providing a low impedance at high frequencies. As capacitance value increases, the highest frequency at which they are remain useful decreases. The first (series) resonant frequency (SRF) of a capacitor depends upon its internal inductance (equivalent series inductance, or ESL), the inductance of its leads, and the inductance of any tracks and vias it is connected to. A first-order approximation for

capacitor SRF is given by: $f_{res} = \dfrac{1}{2\pi\sqrt{LC}}$

where: L = ESL (internal to the capacitor) + the total inductance of any leads + the total inductance of any tracks and/or vias. The rule-of-thumb is to assume 1nH/mm for any leads and/or tracks between the capacitor's body and the power pins of the IC to be decoupled. Via holes are assumed to have 1nH per millimetre of depth. The inductance contributed by ground and power planes may be neglected (but only when the decoupling capacitor is close to its IC). When a via hole connects a track to a plane, only the depth from track to plane is counted.

Decoupling capacitors generally become ineffective at more than 3 times their SRF, as shown by Figure 5.1. [3] shows that the best cost-effectiveness is achieved when very careful attention is paid to the inductance of the decoupling capacitors and their connections to the planes. The use of SMD capacitors with low inductance connections, as shown by Figure 4.2, is clearly very important for effective (and cost-effective) decoupling.

The capacitance created by the close proximity of ground and power planes (on adjacent PCB layers) has low internal and connection inductance and can provide low decoupling impedance to above 1GHz. For example: ground + power planes in PCBs made from FR4 glass-fibre material, with 0.15mm separation, creates approximately 23pF/sq.cm of high quality RF capacitor.

Figure 5.1 Series resonances in decoupling capacitors (guide only)

Good decoupling performance at frequencies from 10 to 1000MHz is therefore best achieved by combining the intrinsic capacitance of adjacent ground and power planes with short-leaded multilayer ceramic capacitors (preferably SMD types) mounted close to each of the power pins of every IC (or transistor). At least one company offers a PCB process that uses special dielectric materials between adjacent ground and power planes to increase the inter-plane capacitance and allow many of the smaller decoupling capacitors to be eliminated.

Three-terminal and "feedthrough" types of SMD decoupling capacitor exist, and have much higher SRF than the regular two-terminal capacitors. However, they are expensive and may not allow the lowest-inductance connection to the power plane, so should be used with care. There are also laminar capacitors devices (such as the Micro/Q range) made to fit under leaded ICs, which are also expensive and seem to be used mostly for improving the EMC of existing PCBs without resorting to redesign.

5.4   Combining decoupling capacitors – problems with parallel resonances

When two decoupling capacitors are provided on a PCB, a parallel (high impedance) resonance is created in between their individual series resonant frequencies, as shown by Figure 5.2. This high impedance prevents the achievement of good EMC at that frequency (and close to it).

This problem is more likely to occur on PCBs with only a few ICs and hence only a few decoupling capacitors, and in these cases should be controlled by fitting a small value resistor, such as 1Ω, or a small ferrite bead (preferably using SMD packages and short tracks) in series with one of the leads on the larger value capacitors.

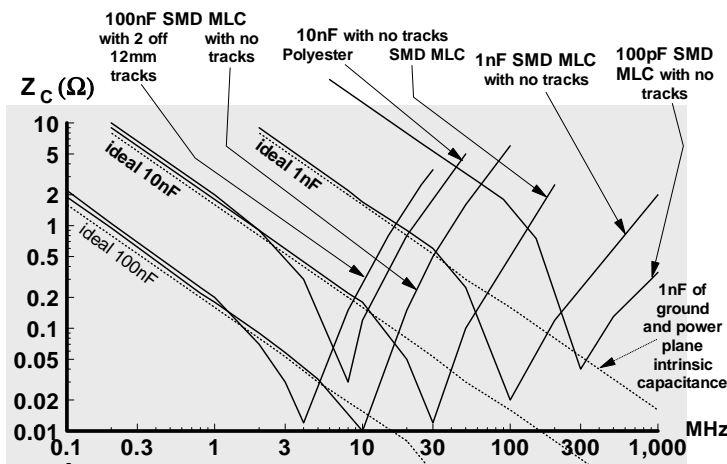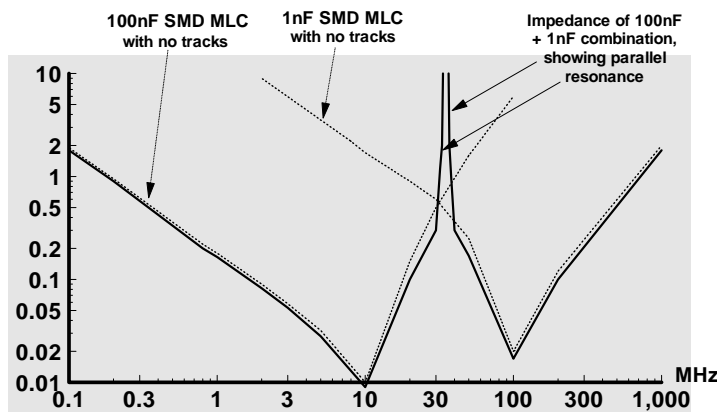When a large number of decoupling capacitors (say over 30) are spread over a ground/power

Figure 5.2   Example of a parallel resonance between one 100nF and one 1nF decoupling capacitors (guide only)

plane PCB, the resulting multiple parallel resonances are swamped by the multiple series resonances. The impedance of the supply is effectively maintained at a low level across a very wide frequency range [3].

Although the capacitance between close proximity ground/power planes is intrinsically non-resonant (at least to 1 GHz), ground and power planes do in fact exhibit resonance due to their dimensions. High-frequency currents in the planes
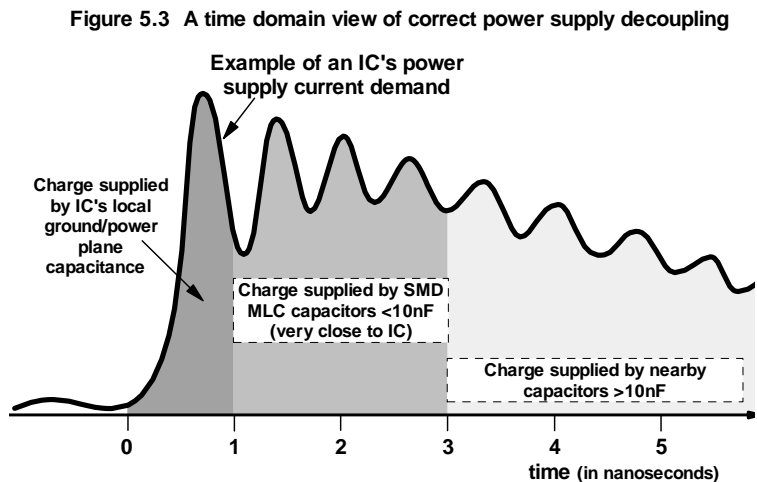
suffer strong reflections at their edges, and their ricocheting from side to side of the PCB makes the planes resonant, with effects in the frequency domain similar to series and parallel resonances in capacitors. However, this is not a reason for avoiding ground/power planes, as the time domain analysis below shows.

Discrete component capacitors are limited in the rate at which they can provide current (di/dt) by their internal inductance, and it is very difficult to get any current even from low-value SMD MLC types during the first nanosecond or two of a transient current demand. But ground/power plane combinations have extremely low internal inductances, so can provide current very rapidly, say for the first nanosecond of a transient demand.

Figure 5.3 shows a time domain view of effective power supply decoupling. The first nanosecond or so of an IC's transient power current demand can only be provided by its local ground/power plane capacitance, the current demand from 1 to 3 ns being provided by the small values of MLC capacitors (up to 10nF) which are located within 20mm or so of the IC's power pins, with larger or further away capacitors (>100nF) only being able to contribute significantly to the current demand from the IC after at least 3ns.

Figure 5.3 A time domain view of correct power supply decoupling

"Bulk" capacitors (such as tantalums) can only begin to provide significant current after 20ns or more, even if they are nearby. Non-ceramic dielectrics and aluminium electrolytics can be slower than expected in their response to a current demand than might be expected from their frequency-domain behaviour, due to a phenomenon known as dielectric absorption, or dielectric memory (another reason for using ceramic types for high frequency decoupling).

Ground/power plane capacitance is essential to maintain low power supply impedance, even though it does create resonances in the frequency domain which need to be controlled for good EMC. For low power-supply impedance up to 1GHz, which is increasingly necessary as clock speeds exceed 100MHz, the addition of very low-value decoupling capacitors, typically in the 100s of pF, may be necessary. (100pF capacitors with SRFs of 1GHz are available from some manufacturers.)

5.5   PCB layer "stack up"

Considering the above, it is best EMC practice to provide ground and power planes on adjacent layers in a PCB, and to maximise their capacitance by using only a thin dielectric (say 0.15mm) between them. In multilayer boards with more than one ground plane, at least one ground/power plane combination should be provided.

Where there are a number of different power supplies, there may need to be a number of different power planes, each on an adjacent layer to a ground plane. Due to the careful segregation of circuits previously described, it is usually possible to have several power planes on the same PCB layer, e.g. for digital +5V and +3.3V supplies, and analogue ±12V supplies.

A four-layer PCB thus has its layers arranged as follows: 1: signal, 2: ground, 3: power, 4: signal

– with the ground and power planes usually separated by only 0.15mm thickness of glass-fibre. Highest-speed or most sensitive tracks should be routed on the layer closest to the ground plane, whereas more indifferent tracks should be routed on the side closest the power plane. High-speed clocks and data busses, and similarly aggressive tracks, should not swap layers. Highest-speed signals may need to use transmission-line techniques, described in section 6. Where more signal layers are required, the ground and power plane "core" should be retained and additional signal layers added. If additional ground planes are not added, the extra tracking layers should only carry signals which are unlikely to cause EMC problems.

In modern high-power computing products almost all the digital tracks are considered aggressive, and there are so many tracks that two tracking layers are often insufficient. For these it is often best to go straight from four layers to an eight-layer board, making sure that every signal tracking layer runs close to a ground or power plane. With a total of three ground plane layers needed to ensure that each signal layer has an adjacent plane, and retaining a ground/power pair for decoupling, an 8-layer stack-up for a computer motherboard may be made as follows:

1: ground   2: signal   3: signal   4: ground   5: power   6: signal   7: signal   8: ground

1: signal   2: ground   3: signal   4: ground   5: power   6: signal   7: ground   8: signal

1: ground   2: signal   3: signal   4: ground   5: power   6: signal   7: ground   8: signal

To help prevent cross-talk between adjacent signal layers, they should be routed with their tracks at 90° to each other (horizontal and vertical). Having ground planes on the outer layers (as in the first stack-up above) can provide the greatest benefits for EMC by reducing external coupling to inner tracks. [12] makes a different case for the benefits of having a ground plane layer on the active component side of high-speed PCBs.

5.6   Further improving decoupling performance and cost-effectiveness

To achieve a low power supply impedance over the whole frequency range of interest, it is necessary to combine various values of decoupling capacitance with the power/ground plane capacitance, according to the needs of the circuit. Where two decoupling capacitors are required at an IC's power pin, the smallest value should be mounted closest to the pin.

When a product with digital clocks or other frequencies of under 100MHz is tested for EMC all the decoupling capacitors with values below 1nF could be left out at first. If they are needed at all will be revealed by high levels of emissions above 300MHz, and they may then be added to the PCB. This is a much better approach than discovering that emissions are too high and only then trying to squeeze a number of additional decoupling capacitors onto the PCB, close to IC power pins.

Ceramic capacitors which employ a COG or NPO dielectric (rather than X7R, Z5U, Y5V, or any of the other standard ceramic dielectric materials) will generally have superior performance at higher frequencies, and should always (at the very least) be used for decoupling capacitors of less than 1nF. Capacitors for which SRF and/or ESL data is provided are preferred over those which do not have these parameters specified.

Devices which have to operate at very high rates or frequencies (say above 200MHz) may benefit from having their own small area of power plane separated by a "moat" from the rest of the power plane for their segregated area. This dedicated power plane is supplied from the main power plane through a ferrite bead, forming a segregated high-speed area within an existing circuit area. The high-speed ICs' local decoupling capacitors are connected between this small plane and the ground plane (which is common with the rest of their segregated circuit area) and a tantalum decoupling capacitor of 1 to 10μF should also be provided. The aim is to increase the "dimensional" resonant frequency of the IC's local ground/power plane combination until they are so high as to have no deleterious effects. Using this technique in practice may not be easy. It will be necessary to ensure that the small power plane does not suffer parallel resonances at a few 100MHz due to the low number or variety of decoupling capacitors. Also, all the tracks interconnecting the small plane area to other areas should be routed on the PCB layer adjacent to the continuous ground plane, and any routed next to the power layers may need filtering by ferrite beads or series resistors at the point where they cross the break.

To get the very best from decoupling, the latest research shows that positioning the relevant ground and power planes closer to the capacitors (and the ICs they decouple) reduces their connection inductance and raises their resonant frequencies [10], [11]. Improving emissions by intentionally tailoring the SRF of decoupling capacitors to correspond to the worst-case frequencies is reported in [11]. [7] recommends that, once each IC (or other device) is decoupled locally, additional benefits may be realised by positioning additional low-value decoupling capacitors around the perimeter of the power plane.

5.7   If not using planes

One way to achieve possibly adequate decoupling without using a power plane is to connect the power end of each IC's decoupling capacitor(s) directly to the power pin(s) of the IC, and then to power the IC from its power rail through a (suitably rated) ferrite bead connected to the power end of the decoupling capacitor. This technique requires a ground plane, so it is suitable for use with double-sided PCBs where all the signal and power tracks have been routed onto one side leaving the other side for a ground plane. Where a large number of ferrite beads are required, multi-layer boards may prove to be more cost-effective and require less area.

6.   Transmission Lines

This technique must be considered as a part of the circuit design process itself, as well as part of the PCB layout.

6.1   Introduction to transmission lines

Transmission lines are interconnections which maintain a chosen characteristic impedance (Z_0) all the way from signal source to load. They have the lowest electric and magnetic field coupling of any interconnections, and unlike all other interconnections do not resonate however long they are. Transmission lines can easily be made on PCBs by controlling their materials and dimensions and providing accurate termination resistances at source and/or load, and may be extended off the PCB (if necessary) by commonly available controlled-impedance cables and connectors.

Transmission lines are needed for signal integrity when a PCB track is so long that a signal travelling along it cannot, due to its finite velocity of propagation, maintain the same potential at all points along its path. In FR4 the velocity of propagation of a signal is roughly 60% of the free space value of 3.10^8m/s (the speed of light in vacuum). A rule of thumb has grown up in the digital community that PCB tracks need only be treated using transmission line techniques when the connection length is greater than $\lambda/7$ at the maximum frequency of interest, f, calculated from:

$$f = \frac{0.35}{t}$$

– where t is the faster of the signal's rise and fall times. From a time domain point of view this is roughly the same as saying that transmission line techniques are required when the time that a digital signal takes to reach the furthest end of its track exceeds half its rise or fall times.

For example: to maintain digital signal integrity for Schottky TTL with 3ns risetimes, the maximum frequency of interest is calculated to be 120MHz, and $\lambda/7$ at this frequency in FR4 is roughly 180mm. Using the time domain view: in 3ns a signal in an FR4 PCB could travel roughly 450mm, and half of this distance is 225mm. The frequency domain estimate of 180mm compares reasonably well with the time domain estimate of 225mm. Likewise, for 1 ns rise and fall time signals, the corresponding track length (beyond which transmission-line techniques would be required) is 66mm.

Databook specifications for output rise or fall times are maximum values, and most of devices will switch a lot faster. Another problem is that the capacitive loading created by other devices connected to a signal track slows the propagation velocity below what would be expected just from the dielectric constant of the PCB material (e.g. FR4). Both these effects mean that transmission lines should be used for even shorter lengths of track than the above examples imply, merely to achieve adequate digital signal integrity. To improve EMC, and to achieve flat analogue signal frequency response, it is recommended that transmission line techniques are used when tracks are at most one-half of the length which is found necessary to achieve digital signal integrity. Further improvements in EMC for aggressive PCB tracks (such as digital clocks and data busses) may be possible by using transmission-line techniques on even the shortest tracks.

In most digital designs, transmission lines are generally used for clock distribution and high-speed data busses. They are also used for slower signals that have to travel further, such as SCSI and USB, and also for even slower communications such as 10base-T Ethernet and RS485, which have to be able to remain sufficiently un-degraded for very long distances.

IEC 1188-1-2 : 1998 [4], gives a wealth of details on constructing a variety of transmission lines using PCB tracks, plus how to specify their manufacture and check their received quality on manufactured boards, [5] and [6] will also be found to be very detailed. Transmission lines are a very large and detailed topic, so this paper will only briefly describe two of the most common types.

The Z_0 for a surface microstrip (figure 6.1) is given in ohms by: $Z_o = \dfrac{87}{\sqrt{\varepsilon_r + 1.41}} \times Ln \dfrac{5.98H}{0.8B + C}$ [4]

– where ε_r is the relative permeability of the substrate (typically 4.4 for FR4 at 100MHz), B is the track width, C is the thickness of the copper material used, and H is the substrate thickness. The propagation velocity for a surface microstrip is given in nanoseconds per metre by: $3.335\sqrt{0.475\varepsilon_r + 0.67}$ [6].

The Z_0 of a symmetrical stripline (figure 6.2) is given in ohms by: $Z_o = \dfrac{60}{\sqrt{\varepsilon_r}} \times Ln \dfrac{1.9H}{0.8B + C}$ [4]

Figure 6.1   Surface Microstrip

Figure 6.2   Symmetrical Stripline
(track is centred between ground planes)

The propagation velocity for a symmetrical stripline is given in ns/metre by: $3.335\sqrt{\varepsilon_r}$ [6]

Striplines are slightly slower than microstrip, but have zero forward crosstalk and much less leakage (ie: are better for EMC). Consequently, computer boards and similar circuits with clock rates exceeding 60MHz or logic transition times of under 5ns will benefit significantly from the use of stripline, rather than microstrip, for these high-speed signals.

Although striplines are best run between two ground planes, they may also be run between a ground plane and a power plane if the power plane has been correctly decoupled (see section 5). Breaks in such power planes must not exist as this would cause sudden changes in Z_0.

To compensate the above formulae for the capacitive loading (typically a few pF per gate):

$$Z'_0 = \frac{Z_0}{\sqrt{1+\dfrac{C_d}{C_0}}}$$ from [6] where Z'_0 is the new (loaded) characteristic impedance, C_d is the sum of all capacitive (gate)

loads, Z_0 is the original characteristic impedance (unloaded) of the line, and C_0 is the original characteristic capacitance of the (unloaded) line obtained from the formulae given in [4]. Velocity is slowed according to the formula:

$$V' = \frac{V_0}{\sqrt{1+\dfrac{C_d}{C_0}}}$$ from [6], where V_0 is the original (unloaded), and V' is the new (loaded) velocity.

It helps transmission line performance if a uniform rate of "gates per unit length" is maintained, rather than bunching ICs together, although the use of sophisticated CAD will allow the line width to be adjusted to maintain the same impedance along its length even when ICs are bunched together.

Different formulas from the ones given here for microstrip and stripline sometimes appear in articles and databooks, and even in microwave substrate manufacturers' application notes. The equations given above have been proven over many years to give the correct results, and are used by major manufacturers of microwave measuring instruments.

Formulae for coated (or embedded) microstrip, offset striplines, balanced striplines (for balanced signals), and others may be found in [4] and [5], as well as in many modern PCB CAD packages. Some CAD packages use field solver methods to calculate Z_0 from first principles for every point along a line – which is much better than relying on choosing the correct formula and hoping that Z_0 is not affected too much by proximity to plane edge, proximity to other tracks, use of via holes, capacitive loading by ICs, and accidental errors in layout.

The highest-speed (or most susceptible) signals should be run adjacent to a ground plane, preferably one that is paired with a power plane to provide power supply decoupling, rather than a power plane. It is not unusual for computer system backplanes (e.g. in VME systems) to have more ground plane layers in them than signal layers, plus more than one power plane layer, merely to maintain signal integrity.

For low inter-track crosstalk, spacing between adjacent transmission lines should be greater than 3B (three times the track width). It is possible to "channelise" an individual stripline to reduce its crosstalk and coupling even further, by running rows of close-spaced vias between its two ground planes, all along both sides of its track. This technique is used for very high frequency signals (e.g. in cellphones) and requires a modified stripline formula.

High-speed or other critical transmission lines must not use any via holes (i.e. should not change layers), and as described above this means routing clock distribution first, moving components as necessary to achieve the smallest area of highest-speed circuitry and tracks. High-speed busses, fast data communications, and the like are routed next, and then everything else (less critical for signal integrity or EMC) is routed around them.

Where there is no reasonable alternative to changing the layers of a transmission line, it is important to realise that its return currents also have to change layers to be able to flow on the adjacent surface of the new plane. High-frequency currents can't easily penetrate planes (they act as shields) so a path must be provided for them. Where the new plane is a ground plane, a via hole adjacent to the via in the track will be needed for the return current path. Even where the new ground plane is merely the other side of the same plane, a via hole will still be needed in the plane. These return current vias must not be so small that they have no inside diameter, or be filled with solder, in fact the larger the inside diameter the better (within reason). Where the new plane is a power plane, a decoupling capacitor with a suitable frequency response should be provided at the point where the signal changes layers.

6.2   Manufacturing issues with transmission lines

Normal FR4 PCB material has a nominal relative dielectric constant (ε_r) of approximately 4.7 at 1MHz falling roughly in a linear fashion with increasing frequency to 4.2 at 1GHz. Actual values of ε_r, can vary from nominal by up to $\pm 25\%$ depending on fibre/epoxy ratio and the epoxy's composition, and can even vary across a single PCB. Controlled ε_r grades of FR4 are available from most PCB substrate material suppliers at little or no extra cost, but PCB manufacturers may not stock this grade of material unless it is specifically requested.

Using these formulae for the first time usually results in dimensions for dielectric thicknesses that PCB manufacturers find difficult to achieve, since they work with standard thickness laminations. It will be best to discuss lamination thicknesses (and manufacturing tolerances) with PCB manufacturers, and then design the track widths to suit the range of dielectric thickness options available.

Some low-cost PCB manufacturers may not have the necessary quality control to be able to reliably manufacture transmission line PCBs. They should be asked what experience they have had with controlled impedance (transmission line) PCBs. Any references should be taken up, rather than accepted at face value. PCB manufacturers who are experienced and equipped to manufacturer controlled-impedance boards will be able to provide test results for agreed specimen tracks that show what impedance has actually been achieved.

The actual track widths achieved after processing often differ by one or even two thousandths of an inch from those shown on the manufacturing photoplots or designer's CAD station. A good PCB manufacturer will be able to suggest what size the tracks should be drawn as in order to achieve the desired finished thickness and hence impedance.

For signal frequencies greater than 1GHz it may be necessary to use other dielectric materials than FR4, such as those used for microwave applications.
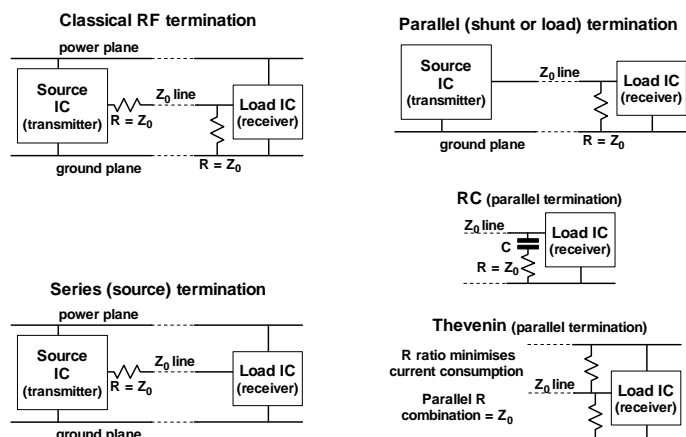
6.3   Terminating Transmission Lines

Textbook radio-frequency transmission lines are terminated both at signal source and load by impedances equal to their line's Z_0 (allowing for internal impedances of source and load devices). This is fine for RF applications, but it halves the received voltage, and so most analogue and digital circuits use low-Z sources and high-Z loads and use only one line termination (at either the source or the load) to achieve adequate signal integrity and EMC.

Signals that travel long distances, especially in cables will generally need to use the "classical" termination method. Standardised or proprietary high-speed or long-distance digital communications such as RS485, 10base-T and 100base-T Ethernet, SCSI, USB, Firewire, will already have their termination methods and values specified, and often the type of cable and connectors (and even circuitry) too.

RF engineers are usually concerned only with extremely narrow frequency bands, so can use complex terminations involving reactive components or even carefully-dimensioned lengths of track. For wideband analogue or digital signals transmission line terminations must work over a very wide range of frequencies, so it is best to use individual resistors. Leaded resistors, SIP and DIP resistor arrays may be used, but SMD resistors are preferred for their higher frequency performance (e.g. for digital signals with transition times shorter than 5ns). To realise the speed advantages of SMD resistor line termination they must be connected using the lowest possible inductances using layout techniques (refer to Figure 4.2).

Figure 6.3 shows a number of common transmission line termination techniques. Source (series) termination at the driver end of a transmission line consumes little or no power. The source IC has an intrinsic output impedance, so the value of the source termination resistor should be chosen so that their series combination matches the transmission line. Series termination is most suitable for transmission lines with a single load device at their far end. Where there are other loads distributed along the length of the line they experience what is called "reflected wave switching" and may need to be slower devices in order not to suffer from false clocking.

Figure 6.3 Transmission line terminations

Classical RF termination

Series (source) termination

Parallel (shunt or load) termination

RC (parallel termination)

Thevenin (parallel termination)

Parallel (load) termination at the very far end of a transmission line is the best method where *very* high speed signals are concerned, especially where there are a number of load ICs devices spread along the length of the transmission line. The load termination resistor is shown in the above figure connected to the ground plane. Some logic families require the use of other reference voltages (such as the positive power plane for ECL). Unfortunately parallel terminations dissipate a lot of power, and may also be found to load some IC outputs too heavily.

Alternative types of parallel termination include "Thevenin" and RC. Thevenin uses resistor values calculated to terminate the line when in parallel, and to provide a DC voltage at their junction which is equal to the long-term average voltage on the line to minimise power consumption and driver loading. Thevenin termination requires that the power plane be properly decoupled to the ground plane at all frequencies of concern, and so may require additional decoupling capacitors nearby.

RC termination typically uses capacitor values between 10 and 620pF and only terminates the line at high frequencies (where transmission line effects are most important). Because of the resonance effects in capacitors and the deleterious effects of long tracks and via holes (discussed in section 5) it is more difficult for an RC termination to equal the highest frequency performance of a single parallel resistor or Thevenin termination.

Another alternative parallel termination technique, often called "active termination" (e.g. often used for SCSI busses), uses a voltage regulator to drive an additional power plane at a voltage equal to the nominal average value of the digital signals. Line terminations use a single resistor to this plane (which of course must be properly decoupled for the frequencies of interest). This is electrically equivalent to the Thevenin method, but reduces the number of components, and can also save power since the voltage regulator (which needs to be able to source as well as sink current) may operate in Class AB and not consume standing currents as high as the equivalent number of Thevenin terminations.

Where a line is driven bi-directionally the compromise position for terminating resistors (series or parallel) is in the centre of the line, so such lines should always be kept very short and may not be able to run as fast as the device speeds may suggest. Series terminations at all possible drivers may be used instead of series termination at the centre of the line, but this may not give good signal integrity unless all the lines concerned are very short. Parallel termination at both ends of the line can give very good performance and allows the highest data rates, but drivers must capable of driving the resulting lower impedances, and power dissipation will also increase. Parallel (or Thévenin or active) termination at both ends is used for serial or parallel data cables such as SCSI and Ethernet.

When using a transmission line to bus a signal to several ICs, when using the series termination method the best scheme may be to "star" connect individual transmission lines to each load from the series termination resistor. The star configuration may also be used for bussing parallel terminated lines, but "daisy chain" tracking is more common, with all unavoidable "stubs" on the line kept to less than 0.01λ in length (λ being the highest frequency of interest) and the termination resistors at the furthest end from the driver.

When using a star configuration for driving a number of transmission lines from one source, the source should be capable of driving the resulting lower impedance, which will be the parallel combination of all the lines' Z_0s. Backplane driver ICs are available with 25Ω output impedances, suitable for "star" driving four individual 100Ω transmission lines.

The interested reader is referred to [6], [13], [14], and [15] for more information on transmission lines. These references will also be found to contain other useful information on EMC, and on EMC techniques for PCBs.

References

[1]   P Ridley, *Implications of the Electromagnetic Compatibility Directive*, IEE Electronics and Communications Engineering Journal, October 1995, pp 195-199.

[2]   G A Jackson*, EMC Directives and Standards*, IEE Engineering Science and Education Journal, October 1993, pp 225 - 232.

[3]   Todd H Hubing, et al, *Power Bus Decoupling on Multilayer Printed Circuit Boards*, IEEE Transactions on Electromagnetic Compatibility, Vol. 37, No. 2, May 1995, pp 155-166

[4]   *IEC 61188-1-2 : 1998 Printed Boards and Printed Board Assemblies – Design and use. Part 1-2: Generic Requirements – Controlled Impedance.*

[5]   IPC-2141, *Controlled Impedance Circuit Boards and High Speed Logic Design*, Institute for Interconnecting and Packaging Electronics Circuits, April 1996.

[6]   M. Montrose, *EMC and the Printed Circuit Board, design, theory, and layout made simple*, IEEE Press, 1999, ISBN 0-7803-4703-X

[7]   Osamu Oeno et al, *Three-dimensional noise current distribution on power and ground planes in printed circuit boards*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.2 pp 1136-1141

[8]   Bruce Archambeault, *Proper design of intentional splits in the ground reference plane of PC boards to minimise emissions on I/O wires and cables*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.2 pp 768-773

[9]   Dheena Moongilan, *Grounding optimisation techniques for controlling radiation and crosstalk in mixed signal PCBs*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.1, pp 495-500

[10]   George Tang, *Surface mount capacitor loop inductance calculation and minimisation*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.1 pp 505-510

[11]   Cornelius B O'Sullivan, *Developing a decoupling methodology with SPICE for multilayer printed circuit boards*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.1

[12]   S.Radu et al, *An impact of layer stack-up on EMI*, IEEE EMC Symposium, Denver, August 24-28 1998, Vol.1

[13]   C Christopoulos, *Principles and Techniques of Electromagnetic Compatibility*, CRC Press, 1995.

[14]   T Williams, *EMC for Product Designers*, Newnes, Oxford, 1992, ISBN 0-7506-1264-9.

[15]   H W Johnson and M Graham, *High Speed Digital Design, a Handbook of Black Magic*, Prentice Hall, 1993, ISBN 0-13-39-5724-1.