



Another EMC resource
from EMC Standards

Design Techniques for EMC Part 5: PCB design and layout

Helping you solve your EMC problems

Design Techniques for EMC and signal integrity

Part 5: PCB design and layout

Originally published in the EMC Compliance Journal in 1999, significantly improved since then and available from <http://www.compliance-club.com/KeithArmstrong.aspx>

Eur Ing Keith Armstrong CEng MIET MIEEE ACGI BSc(hons)
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone & fax: (+44) (0)1785 660247, Email: keith.armstrong@cherryclough.com

This is the fifth in a series of six articles on best-practice EMC and signal integrity techniques in electrical/electronic/mechanical hardware design. The series is intended for designers of electronic products, from building-block units such as power supplies, single-board computers, and 'industrial components' such as PLCs and motor drives, through to stand-alone or networked products such as computers, audio/video/TV, appliances, instrumentation and control, etc.

The techniques covered in these six articles are:

- 1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
- 2) Cables and connectors
- 3) Filters and transient suppressors
- 4) Shielding
- 5) PCB layout (including transmission lines)**
- 6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), and this magazine article format merely introduces the various issues and points to the most important best-practice techniques. Signal integrity is treated as 'internal EMC'. Employing these well-proven techniques from the start of a new design generally reduces the number of iterations of hardware and software during development, and often reduces unit-manufacturing costs too. EMC compliance is generally quicker, easier, with less risk of serious delays in time-to-market.

Table of contents for Part 5

5.	PCB layout	2
5.1	Circuit segregation	3
5.1.1	The boundary between outside- and inside-worlds	3
5.1.2	Boundaries within an inside-world	3
5.1.3	Segregation	3
5.1.4	Component placement and routing of traces	5
5.2	Interface Suppression	5
5.2.1	Suppressing outside/inside-world interfaces	6
5.2.2	Interfaces between dirty/high speed/noisy and clean/sensitive/quiet areas	6
5.2.3	Details of interface suppression techniques	8
5.3	Reference planes	8
5.3.1	Creating proper reference planes	9
5.3.2	Connecting 0V planes to chassis	12
5.3.3	Shielding effect of planes	12
5.3.4	Interconnecting planes in multi-PCB assemblies	12
5.3.5	To split or not to split?	12
5.3.6	Galvanically isolated planes	13
5.3.7	What if multilayer PCBs are thought too costly?	14
5.4	Power decoupling	14
5.4.1	Power decoupling techniques	14
5.4.2	Self-resonance problems	15

5.4.3	Decoupling without power planes	17
5.5	Transmission Lines	17
5.5.1	Changing layers	20
5.5.2	Simulation and prototype testing	21
5.5.3	Manufacturing issues with transmission lines	21
5.5.4	Terminating transmission lines.....	21
5.5.5	Layer "stack up"	23
5.5.6	Joints, stubs, and buffers.....	23
5.5.7	Segregation in backplane systems	25
5.6	Useful references.....	25

5. PCB layout

5.1 Introduction

These PCB design techniques are well-proven to reduce the cost and effort of meeting “external” EMC requirements such as FCC, VCCI, and/or the EMC Directive. They also improve ‘internal EMC’ and signal integrity, and help reduce the number of design iterations it takes to get a product to market.

As electronic technology advances (clock speeds increase, A/D converter resolutions improve) internal EMC problems multiply, and the well-proven techniques described here become more important for commercial success.

5.1.1 The synergy and cost-effectiveness of using these PCB techniques

The PCB techniques described here interact with each other to give improvements which are much greater than each can achieve on its own. Solving EMC problems at PCB level always costs much less than solving them at a higher level of assembly. Even if EMC problems can only be reduced at PCB level, the costs of the higher-level EMC measures required will be considerably reduced.

The PCBs resulting from the use of these techniques may take longer to design initially (and need higher skill levels) and may also have higher material costs, but...

- Unit cost-of-manufacture should be reduced overall
- The number of development iterations will be reduced
- Time to market will be reduced
- Reliability will be improved, leading to improved market perception and lower warranty costs
- Engineering support for early production reduced

Understanding ‘why’ these methods work helps extract their maximum benefits, but all we have room for here is a brief tour of these techniques – a few excellent references are provided at the end.

A major survey conducted in 2000 found that 43% of rework/redesign was caused by EMC and signal integrity-related problems. As electronics technology advances its EMC emissions and susceptibility both worsen and as a result many companies are finding that the ‘EMC fixes’ they used to be able to apply to their products to get them through EMC compliance tests are just not good enough any more. The fixes that would make their products compliant would also make them far too expensive compared with competitors and what the market will stand. So EMC must be now designed-in from the first to avoid serious financially embarrassing delays.

The techniques described in this section improve signal integrity (e.g. ground bounce) as well as improving EMC, so they are very valuable techniques and should always be used. If any of the following PCB techniques cannot be applied correctly for any reason, this is an indication that problems with signal integrity and/or EMC will probably arise during development, so a suitable amount of time and resource should be pre-allocated to deal with this by the project planner.

5.1.2 Profitable products

It is possible, by using the PCB techniques described here, to make a PCB-based product that meets commercial/industrial EMC requirements yet that only needs a low-cost plain plastic enclosure to protect it physically and from electro-static discharges. This allows almost unlimited freedom with the plastic enclosure design.

The lowest-cost way to do this is to use a single PCB for all the circuits and interconnections within a product. There will be no wires or cables inside the product. Although the PCB might be quite sophisticated (e.g. a multilayer stack-up with some ICs covered by PCB-mounted shielding cans) it will generally have a low overall cost-of-manufacture because it will be so easy and quick to assemble, and it will tend to be more reliable (reduced problems with poor connections in wires and cables, plus good EMC performance).

Remember, it is low cost-of-manufacture and low warranty costs that are important for a profitable product. Material costs are *totally irrelevant* as long as the overall cost-of-manufacture is low and the product is reliable enough. Unfortunately many companies do not have a good way to determine assembly costs during design (much less reliability), so tend to focus on material costs because these are easy to establish. Some companies even assume that the selling price is a multiple of the material cost, whereas for the vast majority of electronic products the selling price has no direct (i.e. causal) relationship with material cost at all. In fact, reducing the material cost can lead to more expensive cost-to-make, not to mention delays in time-to-market and high warranty costs through unreliability.

5.1.3 Advanced PCB techniques

This part of the series covers basic PCB design techniques for EMC. Products with PCBs that use high-speed signals (e.g. clocks >66MHz, risetimes of under 1ns, analogue >1GHz) may need to use advanced PCB techniques as well as those described here. Even low-frequency signals may need advanced PCB techniques, if they are susceptible to interference and in an aggressive electromagnetic environment without high-specification enclosure shielding and filtering. The basic techniques described here are a pre-requisite for the use of more advanced techniques.

5.2 Circuit segregation

For cost-efficiency, segregation needs to be employed from the start of the real design process. The various segregated zones may need filtering and/or shielding from each other, so it is clear that the layout of the PCB cannot begin until it is known where these are to be put.

The two main zones are:

The Outside World: Where total control of the electromagnetic (EM) environment is not practicable.

The Inside World: Where total control of the EM environment will be achieved.

5.2.1 The boundary between the Outside and Inside worlds

Outside and inside zones are easy to establish for a product with a fully-shielded enclosure (they are segregated by the shield) but can be hard to draw for unshielded products.

Conductors completely within an unshielded product will couple with the external EMC environment. They can be a cause of radiated emissions, and the radiated RF environment will cause them to carry noise current and voltage and inject it into the circuits, depending on how efficient the conductor is as an unintentional antenna at the frequencies concerned (see Volumes 1 and 2).

5.2.2 Boundaries within an inside-world

When the zone which is the inside-world has been determined (there may be two or more 'inside zones' in a multi-PCB product where the PCBs are connected by wires), each zone should be further subdivided into dirty, high-speed, noisy, (etc.) potentially "aggressive" circuits, and clean, sensitive, quiet, (etc.) potential "victim" circuits. The likelihood of a circuit node being aggressive depends on

its maximum dV/dt and/or dI/dt . The likelihood of a circuit node being a victim of EM phenomena depends on its signal levels and noise margins (less = greater sensitivity).

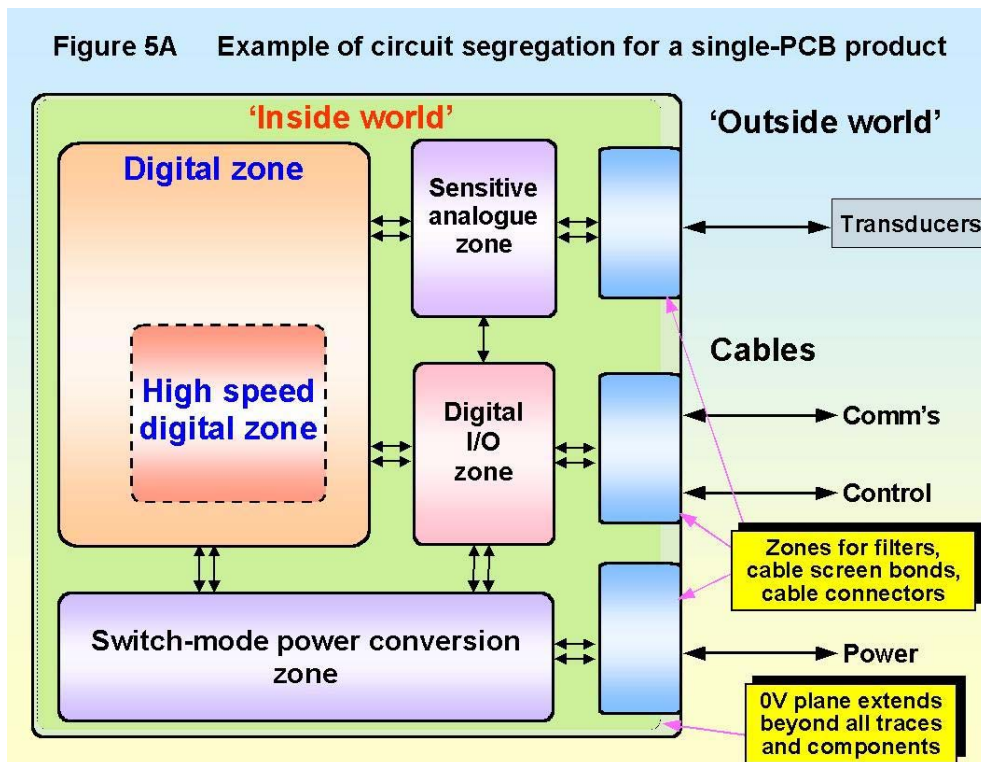
5.2.3 Segregating the zones

The various inside-world areas should be physically segregated from each other, and from the outside-world, both mechanically and electrically. Start at the earliest design phase by showing the segregated areas clearly on all drawings, usually done by drawing dotted lines around rectangular areas each covering one segregated portion of the circuit.

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 segregation clearly on all system, wiring, and circuit diagrams is of great help in communications between electronic designers, mechanical designers, and PCB layout persons – and is especially important where work is done by people on different sites, or by subcontractors.

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) finds itself in close proximity to a noisy circuit (such as a switch-mode power converter), with consequent signal quality problems. Such unpleasant and time-consuming three-dimensional assembly problems 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 5A shows good segregation practices applied to a single-pcb product, whether it has an overall enclosure shield or not.



This example shows that the segregated area where the outside world interface suppression components are fitted, is along one edge of the PCB, as if it were a wall between outside and inside-worlds (which it is, in a way).

This area would only contain opto-isolators, isolating transformers, baluns, filters, transient absorbers, similar interfacing components but no ICs. 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.

The inside/outside-world interface components are restricted to one dedicated edge of the PCB to encourage all unwanted external currents (e.g. caused by voltage differences in protective earths) to restrict themselves to that area of the PCB, and discourage them from flowing through circuit areas.

Where an effective enclosure shield exists, the inside/outside-world boundary becomes the shielded wall of the enclosure. All of the associated filtering and suppression components, and cable screen bonding, must then use a connector panel set in the wall of the enclosure as their reference (as described in Part 4). A single area for all interconnections is still best. A wider range of PCB-mounted screened and/or filtered connectors that can also bond to a metal panel is now available. These parts would be soldered to the PCB reference plane, then electrically bonded metal-to-metal to the wall of a screened enclosure during final assembly, and can be very cost-effective.

Narrow channels free from components should be left between each of the segregated circuit areas on the PCB. These should be wide enough for the fitting of a PCB-mounted "tuner-can" shield, and provision should be made (at least on prototype boards) for bonding such screening cans to the 0V plane at frequent intervals (say, every 15mm) along all edges.

5.2.4 Component placement and routing of traces

The most noisy or susceptible components in each area should be positioned first, as close to the centre of their areas and as far away from cables or wires as possible. 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). Remember (from Part 1) that even low-frequency operational amplifiers can be extremely susceptible to interference, even beyond 1GHz.

After the extremely short connections from components to reference planes, digital clock distributions (very aggressive signals) must be the next "nets" to be routed, and must be run on a single PCB layer adjacent to a 0V plane. These traces must be as short as possible, and even so may need to use transmission-line techniques (described later). It may be necessary to experiment with component placement to achieve minimum trace lengths. Where clock traces are made longer than necessary to minimise skew, a "serpentine" layout is best.

Digital busses and high-speed I/O should be routed next, in a similar manner to clock traces, deferring only to clock traces and plane bonds where there is a conflict. Very susceptible traces, such as those carrying millivolt transducer signals, should also be routed as if they were clock or data buss traces, although they will always be in a different segregated area of the PCB. The later section on transmission lines describes what to do where critical traces have to change layers.

All other types of analogue, digital, and power signals should also be routed according to how aggressive or sensitive they are. Where these characteristics are not obvious from a circuit analysis, probing a prototype with a wide-band oscilloscope (and/or spectrum analyser) with 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. A loop probe can be most useful here, being able to inject signals into traces without requiring connection of external equipment to potentially sensitive area of the circuit concerned.

All components and their traces must be contained within their designated PCB zones (areas). The only traces which may exit or enter a zone are those that have to interconnect with other zones – these traces might all need filtering or suppressing in some other manner, discussed later. Where it has not proved possible to eliminate all the wires and cables inside a product, make sure that their routes are fixed so they are as short as possible and can't stray into the wrong PCB zones.

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 traces and other conductors which cross the dotted lines showing the segregated zones 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 unexpected traces crossing zone boundaries – these are often the source of much design heartache, so eliminate them right away by applying more skill to the layout. Autorouting does not generally provide good layouts for EMC purposes.

Devices which connect directly to other zones (e.g. A/D and D/A converters, filters, etc.) are positioned right at the nearest edge of their zone. All external interconnections should be restricted to one edge of the PCB only, to help prevent external circulating (noise) currents from flowing through the circuitry on the PCB.

5.3 Interface Suppression

EM disturbances can be radiated and/or conducted across interfaces between segregated areas, and shielding, filtering, or isolation techniques (such as opto-coupling) are used to reduce this to acceptable levels. To decide on the most cost-effective methods for each interface, they should be assessed for all the EM phenomena possible, given the operational EM environment and the emissions/immunity characteristics of the circuits concerned.

Don't ignore internal power supplies and other common connections such as 0Vs or grounds when considering interfaces between areas. Circuit designers abbreviate such connections on their circuit diagrams, often to invisibility, even though they provide the return current paths that are as important as the send path.

5.3.1 Suppressing outside/inside-world interfaces

Conductors passing from outside to inside-worlds may need the full range of suppression techniques – shielding, filters, isolating transformers, opto-isolators, surge protection devices, etc. As described above, best practice is to use a single PCB area or panel in the enclosure shield for all outside/inside-world interconnections and their suppression.

Visual displays (such as LCDs, LEDs, VDUs, moving-coil meters, etc.) and controls (such as pushbuttons, potentiometers, rotary knobs, etc.) are also interfaces between outside and inside-worlds, and are particularly exposed to personnel electro-static discharge (ESD), which will be covered by Part 6 of this series.

Shielding (See Part 4) may be applied to chips or areas of the PCB; the whole PCB; sub-assemblies of PCBs; entire assemblies of PCBs; or the entire product (listed in ascending order of cost and difficulty). The segregation methods described above help make low-cost shielding possible.

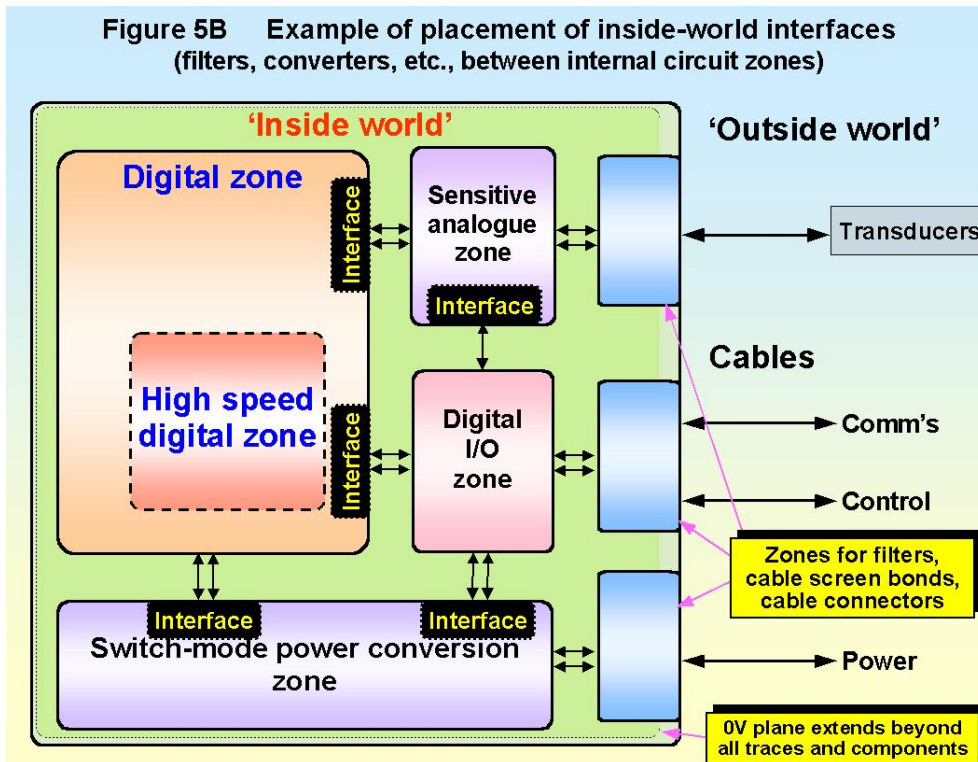
5.3.2 Interfaces between dirty/high speed/noisy and clean/sensitive/quiet areas

Determining the types/amounts of suppression to be applied to traces and other conductors interconnecting different PCB areas needs an assessment of both the desired signals and the unwanted noise they may carry, plus the sensitivity of the circuits they connect to.

Digital clocks and data busses are aggressively noisy and should not be allowed in clean/sensitive/quiet PCB areas. Data intended for a sensitive area should be latched from its bus no closer than the boundary of that area, and the data busses themselves restricted to a noisier area.

Power distribution networks are often overlooked routes for conducted noise from one segregated area to another, as are "static" data lines, and other low-frequency signals. Digital control lines which remaining at logic 1 or 0 for long periods are often thought to be quiet, but they usually carry tens or even hundreds of millivolts of high-frequency noise generated by the electrical activity of their source ICs (e.g. by "ground bounce" and its corresponding "power bounce"). Many an analogue circuit has suffered from noise on its power supply rails from switch-mode power supplies or DC/DC converters, or from digital processing sharing the same rails or from noise injected into analogue switches and opamps from "static" logic control signals. It is often necessary to fit small filters to such inter-area connections, but sometimes more drastic measures are required, such as opto-isolation.

Components that interface between segregated areas, such as analogue-to-digital converters, transformers, data bus latches, filters, isolators, and the like, should be positioned at an edge common to the areas they interconnect, as shown by Figure 5B. They should usually remain wholly within one area or the other (so as to keep a component-free channel between segregated circuit areas) and their traces must route directly to their respective areas and not mingle with traces associated with the “other side” of these components or other areas.



One purpose in keeping the channel component-free is to reduce crosstalk, but it also makes 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 traces associated with each circuit area, but the cut-outs they require in any PCB-mounted shield may compromise its shielding effectiveness.

This compromise between the need for good trace 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 trace-free channels) they encourage good trace 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 (at least as often as 50mm times the risetime of the fastest signal in nanoseconds). 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 possible. The stray capacitance between components may only be a fraction of a pF, but at high frequencies can inject significant displacement currents into components and traces in neighbouring areas. Combining small-sized low-profile

components with PCB reference planes, and placing the noisiest devices (e.g. clocks, processors, switch-mode power devices) and signals in the centres of their areas, can help avoid the need to shield PCB areas from each other.

5.3.3 Details of interface suppression techniques

Suppression techniques include:

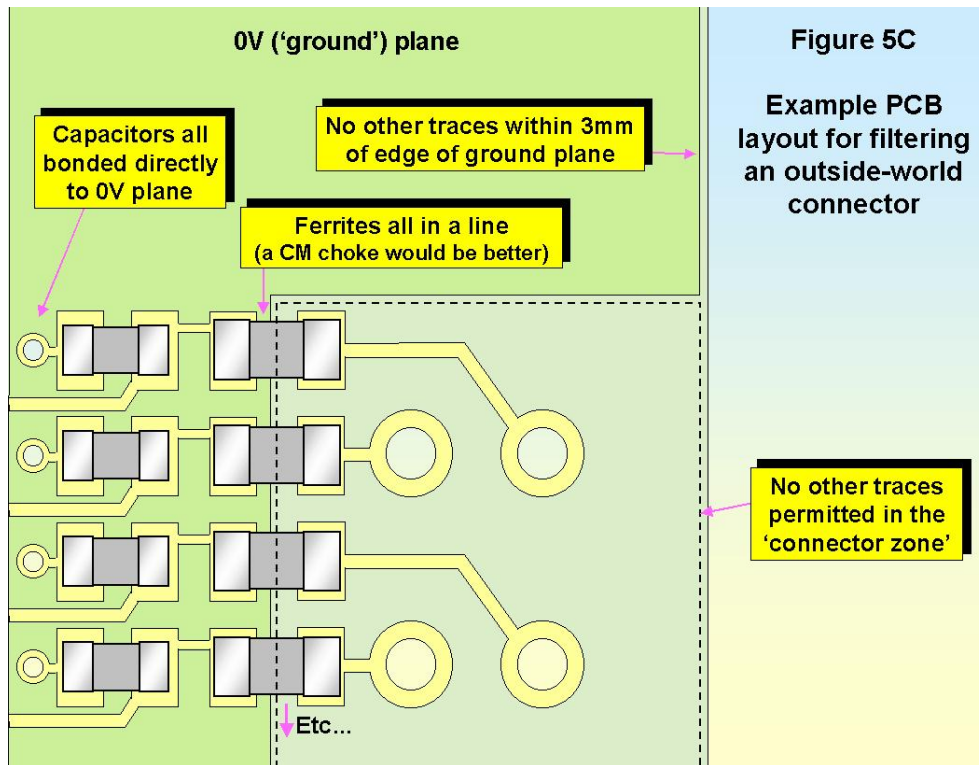
- common-mode and/or differential mode filtering
- galvanic isolation using opto-isolators or transformers
- communications protocols (to improve bit error rate in the presence of interference)
- surge protection devices
- the use of balanced drive and receive signals (instead of "single ended")
- the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

All of these are covered by other parts of this series. It is important to realise that on a PCB only a plane (described next) can provide a good enough reference at high frequencies to enable the full performance of filters, cable screens, and internal shields to be achieved on a PCB.

Figure 5C shows an example of a PCB layout for filtering a through-hole connector connected to an outside-world signal cable. The filter is simple LC type (see Part 3) using individual RF-suppressing ferrite beads. A common-mode choke is often a better solution, especially for high-speed signals for which individual line filters might affect the signals, and it is recommended to place and route the ferrites so that their pads can be used to fit individual beads (as shown) or a common-mode choke, at least on prototype PCBs.

The PCB's ground plane is cut away around the connector to reduce capacitive coupling to the unfiltered conductors. The connector area only contains the PCB traces between the connector pins and the row of ferrites. All the other traces are routed over the ground plane, keeping 3mm or so away from its edges at all times.

The ferrites are arranged in a row so that the trace and pad associated with the 'connector side' of any one ferrite is always as far away as possible from the trace and pad associated with the 'capacitor side' of all other ferrites. If this is not done, RF can couple across (leak around) the filter through stray capacitance paths. The capacitors are placed in a row and close to their associated ferrites for a similar reason.



The bonds between the capacitors and the ground plane must be very short and direct, to provide the lowest possible inductance and so help the capacitors to filter out higher frequencies (see the section on decoupling, later in this Part).

5.4 Reference planes

Due to their intrinsic reactance and resonances, traces, wires, “star grounding”, area fills, guard rings, etc., cannot provide an adequate reference for a PCB except at low frequencies (usually below 1MHz). For example, the rule-of-thumb for the inductance of PCB traces on their own or single wires, is 1nH/mm. This means that just 10mm of PCB trace has an impedance of 6.3Ω at 100MHz, and 63Ω at 1GHz. For this reason, only unbroken areas of metal conductor can provide an adequate reference up to 1GHz (and beyond), and these are called reference planes. In a PCB these are usually called power, ground, or 0V planes, but it is best to avoid the use of the words “ground” or “earth” in connection with EMC and circuits (reserving them for specific uses associated with safety bonding). As far as most EMC design techniques are concerned, a connection to the green/yellow protective earth conductor can often be more of a problem than a solution.

Reference plane techniques allow dramatic reductions in all unwanted EM coupling when used in conjunction with the other techniques described here. Reference planes are also essential for almost every other PCB EMC design technique to function properly.

5.4.1 Creating proper reference planes

As frequencies and component densities increase and rise-times decrease, PCB planes are increasingly necessary. They are the only way to provide a known voltage, or a minimum-loop-area current return path, at RF. They should be called Reference Potential Planes and/or Return Planes, but are usually called simply ground, 0V, or power planes.

Unless a circuit is within a well-shielded and filtered enclosure with no digital or switch-mode or other noisy circuits, it is nowadays exposed to an environment which contains a great deal of RF. So almost all circuits should now use 0V planes. Where 0-5V signals on a PCB are >4MHz digital or >100MHz sine-wave, power planes should also be used.

To function as intended and provide a plane needs a vanishingly small inductance between any two connections made to it. A plane can be achieved on a PCB by devoting one copper layer to an *unbroken* copper sheet and making all connections to it *direct*. A ‘copper fill’ is not a plane.

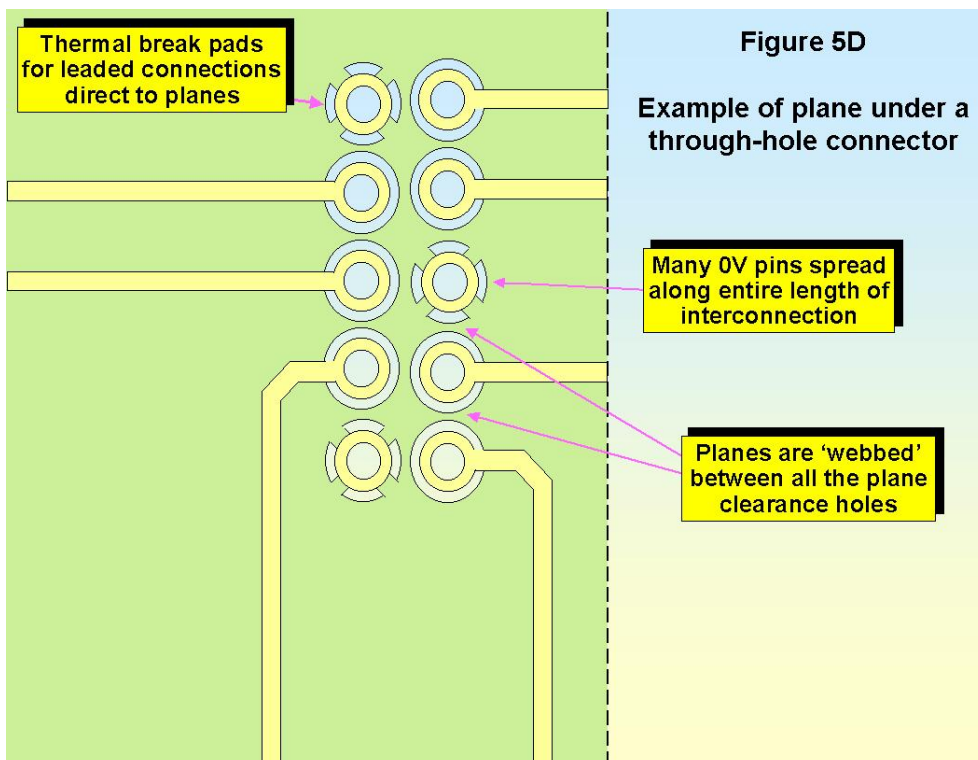
A 0V plane must be used as the 0V (or ground) connection for all its associated circuits, so that no 0V return currents flow in any traces. To help minimise inductance in the return path, only use ‘thermal break’ pads where necessary for soldering, and minimise the lengths of any ‘pin-escape’ or ‘break out’ traces from surface-mounted components to the vias that connect to a plane (see later).

Power planes are created and used in a similar manner for power connections and their return currents.

0V planes must lie under all their components and all their associated traces, and extend a significant distance beyond them and their power planes (a distance of 20 times their height above the plane is often recommended). The segregation and interface suppression techniques described earlier must still be followed even where a 0V plane is common to a number of different circuit areas.

Perforations such as leads, pins, and via holes increase the inductance of a plane, making it less effective at higher frequencies. “Buried via” techniques have been developed for cellphones, allowing interconnections between routing layers without perforating the reference plane. For less demanding products a rule-of-thumb is that any gaps must have dimensions of 0.01λ or less at the maximum frequency concerned. For a good plane at 1GHz, (e.g. to help meet most of the present EU harmonised EMC standards cost-effectively) this rule implies that plane gaps should have dimensions $\leq 1.5\text{mm}$ (remembering that the velocity of propagation in FR4 is approx. half of what it is in air). “Sneaking” traces into a plane layer is not allowed.

Unavoidable gaps in a plane must not merge to create larger ones. PCB design rules should size clearance holes so that for regular hole spacings such as through-hole connectors or DIL packages, the plane “webs” between holes as shown by Figure 5D.




Traces, area fills, guard rings, etc. forming part of the reference on signal layers *can* be used to good effect at high frequencies – but *only* when bonded to an underlying 0V plane with at least one via hole every 5mm to 15mm (using a random allocation of spacings).

0V planes should extend well beyond all components, traces and power planes. [1] recommends “the 20H rule”: 0V planes should extend by at least 20 times their layer spacing. High-speed components

(such as digital clocks, processors, and memory) and their signal traces should always be placed near the centres of their segregated areas, well away from plane edges.

All 0V and power connections must bond immediately to their respective planes to minimise their connection inductance. Leaded components must have their through-plated holes directly connected to planes using thermal-break pads as shown by Figure 5D (sometimes called wagon-wheels) to help with soldering. Surface mounted devices (SMDs) for reflow soldering have to compromise the prevention of dry joints (e.g. tomb-stoning) with the need to minimise the inductance of plane connections.

Figure 5E shows various methods for connecting reflow-soldered decoupling capacitors to their ground and power planes. The most important issues in this case are to get the two plane connections as close together as possible and to keep any 'break out' traces as short and as wide as is possible without risking dry joints. Other components connecting to planes should follow similar rules, especially IC power pins and transmission-line matching resistors.



The very best method is to use 'via-in-pad'. This can be done with the standard through-hole processes by using over-sized pads and tenting a good-quality solder-resist over a number of plane vias. Via-in-pad is automatic when using Microvia (high density interconnect or HDI) PCB technology

High-quality solder resist, tented solder resist, minimum size soldered areas, and similar techniques should be used where they allow the plane-connecting vias to be closer to the component pads without causing dry joints. Plane connections that do not need to be soldered (typical of the vias for reflow-soldered SMD components) may not need to use thermal-break pads – using solid plane connections instead will reduce inductance still further (but never extend a trace to avoid having to use a thermal-break pad).

It is best to make reference planes rectangular (but not thin) to minimise their inductance, and also to make the fitting of PCB level shielding easier. Square planes, and planes with simple aspect ratios such as 1:2, should be avoided to help reduce possible problems with resonances. Where there are a number of different power supplies, there may need to be a number of different power planes. Segregation of circuit areas (see earlier) makes it easier to fit several broadly rectangular power planes on the same layer.

5.4.2 Connecting 0V planes to chassis

Components and traces have weak capacitive coupling to everything else. Electrical activity causes displacement currents to flow in these “stray” capacitances, a cause of common-mode emissions. High-speed circuits usually need at least a nearby metal surface, and (increasingly) a fully shielded enclosure, to reduce the resulting emissions problems. The metal chassis or shields need to be connected to the reference plane of their PCB, preferably at a number of points spread over the PCB so that the high frequency displacement currents can be returned to their source within a fraction of their wavelength. PCB mechanical supports and fixings are often used for these chassis bonds, but should be very short (<4mm). There should be at least one bond in the centre of each area of high-speed circuitry, especially clock generators and distribution. For high-speed digital boards, 0V plane to chassis bonds every 50mm to 100mm all over the PCB may not be overkill, and provision (at least) should be made for these on prototype PCBs. Even if it is not intended to have a metal chassis or shielded enclosure, it is still a good idea to include a number of potential chassis bonding points, just in case. Sometimes a sheet of aluminised cardboard or PVC is sufficient to overcome unexpected problems, providing it can be bonded to the right place(s).

To add flexibility, especially for mixed analogue/digital PCBs, each plane-to-chassis bond can have traces and pads that allow the bond to be left open, or else fitted with direct links or capacitors of various types and values. Fitting a direct link at one chassis bond and capacitors at the others allows low frequencies (for which the inductance is not important) to be controlled with a “star ground” system, whilst high frequencies are controlled by the low inductance of the widely distributed capacitive links. Care should be taken to minimise the inductances of all these traces, pads, and linking components (SMD preferred). Where reference planes must be galvanically isolated only capacitor bonds may be used, but care should be taken with safety approvals and earth-leakage requirements (especially for patient-coupled medical apparatus).

5.4.3 Shielding effect of planes

Antennas placed close to metal planes are less effective at radiating and receiving. Many advantages of planes are due to the way they allow the return currents to take the path of least inductance, but their “antenna shielding” effect is also important. For any significant advantage to be achieved from this effect, the tops of all the PCB components must be no more than one-twentieth of a wavelength above a PCB plane, at the highest frequency of concern for emissions and immunity, e.g. 15mm to give a degree of shielding to analogue circuits exposed to 1GHz immunity testing.

Even lower profiles will give improved shielding, one reason why SMD components are much preferred for EMC, with very low profile ball-grid-array and flip-chip technologies being better still. The plane needs to extend by considerably more distance around the components than their height above it.

5.4.4 Interconnecting planes in multi-PCB assemblies

Card cage, backplane, and mother/daughterboard structures will experience considerable signal integrity and EMC advantages from linking their reference planes together with very low inductance. This may be achieved with frequent low-inductance links between their planes, more-or-less uniformly distributed along the *full* length of all their common boundaries. Shielded backplane connectors are happily becoming more commonly available. Where shielded connectors aren't used, using one 0V plane-linking pin alongside *every* signal or power pin in a connector may seem expensive, but sometimes it is the lowest-cost (or only practical) way to improve the EMC of a multiple-PCB product. Bonding planes via front panels and/or card guides is also very worthwhile.

5.4.5 To split or not to split?

Split reference planes may give better *or worse* EMC (and signal integrity) than unsplit planes, and this depends very much on the PCB layout and circuit design so it is often hard to decide which method to use. Note that where a 0V plane is to be split off from the main 0V plane, it may still need chassis bonds as described earlier. This is particularly true of “traditional” method of splitting off

connector panel 0V plane areas (to try stop noise on the main board from exiting via the external connectors), when the connector area on the PCB must have its local 0V plane bonded to any enclosure shielding. Also note that the inevitable stray capacitance across a split progressively “shorts it out” above 500MHz anyway.

To get any benefits from split planes with modern electronic technologies requires significant attention to detail, which is one reason why an increasing number of designers are now using common, unsplit 0V planes as a matter of course.

Allow for both split and unsplit options, on prototype PCBs at least, by splitting all planes at the natural boundaries between the segregated circuit areas, but also providing the means to “stitch” them together manually later on. Stitching requires pairs of via holes on each side of the split every 10mm or so (random spacing of 5mm to 15mm preferred). These via pairs may be left open, or bridged with short wires or capacitors, and it is important to pitch the via pairs close together so that small capacitors or “zero-ohm links” can be used (preferably SMD). Linking planes with a single copper link and multiple capacitors can control lower frequencies (where inductance is not significant) by “star grounding”, whilst also controlling higher frequencies by creating the effect of a single low-inductance plane.

Because a split in a plane is a slot antenna, it is best if no traces cross the split (or even go near to it). Where traces *have* to cross – they *must* have carefully-defined return current paths, and for high frequency currents these paths *must* be physically adjacent to their send traces. These tend to defeat the purpose of the split, so should be limited to the bandwidth of the wanted signal (which should already have been restricted to just what is needed, as described in the section on interface analysis and suppression above). High-speed signals can usually be returned through a suitable size and type of capacitor, although some data streams with highly-variable content may need a more wideband return path than a single capacitor can easily provide (may need a direct link).

Balanced signals would ideally need no local return path, but in practice their balance always degrades at some frequency so a nearby return path is needed for the resulting common-mode “leakage” (usually a small-value capacitor). DC power and low-frequency signals that have been filtered to remove all high-frequency noises can use the star point between the split 0Vs for their return, as long as the inductance of the resulting current loop is negligible. Beware of assuming that a conductor is only carrying low frequencies just because that is what its signal name implies. In modern mixed digital/analogue products all the traces and other conductors in a product usually carry significant levels of high frequency noise. A local return path for a low-frequency signal could be a ferrite bead.

Common-mode (CM) chokes fitted to any types of signals and their associated returns (e.g. a 4-circuit CM choke for a set of three related signals and their return) will probably help get the best performance from split planes, but cost more.

When all the above has been designed into the split-plane PCB, it will need testing and optimisation. Direct or capacitive links to/from the “stitching vias” should be added/subtracted to achieve the best EMC performance. If it is discovered that the best EMC is achieved when all the stitching points are directly linked, the next iteration of the PCB could remove the splits and their stitching points completely, saving manufacturing costs.

5.4.6 Galvanically isolated planes

The split planes described above are all ultimately powered from the same power rails (0V, at least), so there is a clear need for return current paths to be catered for every conductor (signal or power) that crosses from one plane area to another. It is often assumed that galvanically isolated areas have no return current requirements, but this is not so at high frequencies.

Galvanic isolation devices (opto-isolators, transformers, etc.) suffer from stray internal capacitance. A typical opto has 0.8pF internal capacitance, which provides a shunting impedance of only 2k Ω at 100MHz, or 200 Ω at 1GHz, which will clearly prevent signal isolation from being maintained at high frequencies. Transformers (especially in DC/DC power converters) tend to have even larger stray

internal capacitances. Common-mode chokes may be used to improve the isolation at high frequencies, but struggle to increase it by an order of magnitude at 1GHz. There are also many other stray capacitances around to compromise isolation. So there is a need, at high frequencies, to provide a local return path for the displacement currents that flow in these stray capacitances, to prevent them from causing common-mode conducted and radiated emissions and immunity problems.

Because we usually only need isolation for low frequencies (usually only 50Hz) we can connect galvanically isolated planes to the main reference plane with a number of low-value capacitors (spread around the gap perimeter), so as to achieve the effect of a single reference plane for high frequencies and provide low-inductance local return paths for stray displacement currents.

Of course, great care may need to be taken with component approvals and leakage currents where safety is concerned.

5.4.7 What if multilayer PCBs are thought too costly?

In volume, four-layer PCBs now only cost between 20% and 50% more than two-layer. The use of planes usually turns out, in retrospect, to have been the most cost-effective EMC technique possible, especially when the *overall* financial break-even time and profitability of a product is considered.

An appropriate technique for low-density double-sided PCBs is to put all the traces on one side, and a *solid* 0V plane on the other. For digital products, the lack of a power plane might require a number of ferrite beads in the power rails (see later), so it might not prove to be most cost-effective.

Where traces must use both sides of a two-layer PCB, some EMC improvements may be had by "gridding" 0V traces. This can be done by using a "maximum copper" or "area fill" on the 0V traces of both PCB layers, which must run perpendicular to each other, "stitching" the resulting horizontal and vertical 0V areas and lines together with via holes wherever they cross to create a grid over the whole PCB area. Smaller grid sections are needed around the more sensitive or aggressive components, often difficult to achieve for leaded microprocessors but easier for SMD types. Time should be allowed for moving components and traces around to achieve the best grid structure, but any grid will always be much less effective than a proper solid plane.

Single-sided PCBs are extremely difficult to make EMC compliant without enclosure shielding and filtering, except for circuits which naturally have very low emissions (low dV/dt and dI/dt) and also have naturally very high immunity (e.g. high signal levels and low impedances).

5.5 Power decoupling

The aim of power decoupling is to maintain the power supply impedance to each IC at 1Ω or less across the entire frequency range of interest (at least 150kHz to 1GHz for EMC). Some devices may need 0.1Ω or less over some frequency ranges for correct operation. Wires and PCB traces have too much inductance to provide these low impedances, which require local capacitance of suitable quality.

5.5.1 Power decoupling techniques

A large decoupling capacitor (e.g. 100μF) should be fitted where power supplies enter or leave a PCB, and some smaller ones (e.g. 10μF tantalum or solid aluminium) should be "sprinkled" around the PCB on a "μF per unit area" principle, as well as being positioned near to heavy power usage such as microprocessors, memory, and other powerful digital ICs. These capacitors provide low impedance up to, say, 10MHz. Next, the power supplies to every IC should be decoupled very nearby using appropriate capacitor sizes and types. Where an IC has a number of power pins, each pin should have an appropriate decoupling capacitor nearby, even if they are on the same supply (e.g. Vdd).

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

The total local decoupling capacitance required depends on the IC's transient power demands and the tolerances of its DC power rails. 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 – usually not true of real life.

The formula $C(\Delta V) = I(\Delta t)$, using the units Farads, Volts, Amps, and seconds, covers what we want to know. ΔV is obtained by subtracting the IC's minimum operational voltage (from its data sheet) from the worst-case minimum power rail voltage (taking account of initial tolerances, regulation, temperature coefficients, ageing drift, and the voltage drops in the power conductors). ΔV often turns out to be a mere +100mV. I is the IC's transient current demand from its power rail, which lasts for Δt . I and Δt are almost never found in data sheets, and must be measured in some reasonably sensible way with an oscilloscope. An obvious component of I is the device's output (load) current, but this is often negligible in comparison with "shoot-through" currents, also known as "transient supply current". There is no point in measuring I or Δt with greater than $\pm 20\%$ accuracy.

Where ΔV is low it may be cost-effective to increase it by improving the regulation of the power supply, and/or reducing the resistance of the power rails, rather than fit larger capacitors with their lower performance at high frequencies. This is a common argument for local power regulation.

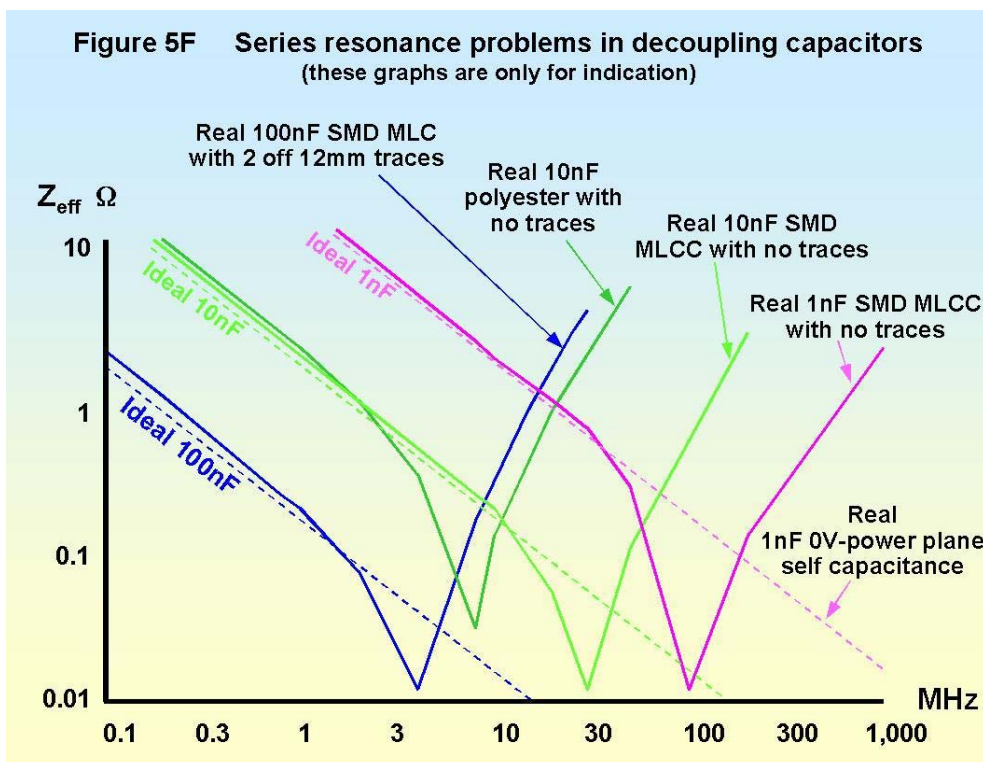
5.5.2 Self-resonance problems

Self-resonance in capacitors stops them providing low impedances at high frequencies, with higher values generally being worse. The first self-resonant frequency (SRF) of a capacitor is a series resonance, and a rule of thumb for this is:

$$f_{res} = \frac{1}{2\pi\sqrt{LC}}$$

where $L = ESL$ (internal to the capacitor)

+ the total inductance of any leads + the total inductance of any traces and/or vias. 1nH/mm may be assumed for leads and/or traces from a capacitor to the power pins of its IC. The inductance contributed by 0V and power planes may be neglected when the capacitor is near to its IC. Decoupling capacitors generally become ineffective at more than 3 times their SRF, as shown by Figure 5F.



It is interesting to note that the favourite 100nF capacitor, even with no traces at all, is effectively useless above 50MHz, yet it is still often seen in circuits with clocks of 50MHz or over, where it can do nothing to help control the fundamental clock frequency, never mind its harmonics.

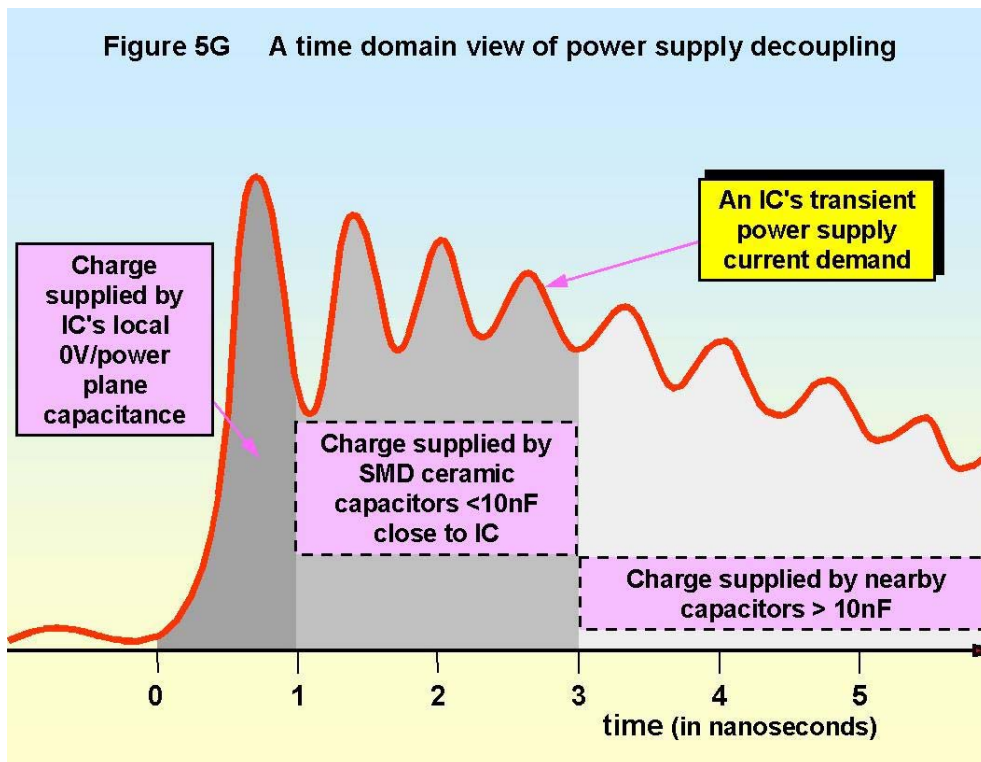
Close proximity of adjacent 0V and power planes (with their low internal and connection inductances) can provide capacitance with no SRF below 1GHz. Two planes separated by 0.15mm in an FR4 PCB achieve approximately 23pF/sq.cm of high quality RF capacitor. Good decoupling from 10 to 1000MHz can be achieved by combining adjacent 0V and power planes with SMD ceramic capacitors (COG and NPO types are best). Sometimes two different values of capacitors (e.g. 100nF and 1nF) may be required. Low inductance bonds from IC power pins and decoupling capacitors to their planes is essential, and the capacitors must be positioned close to their IC. The common practice of routing from IC power pin to decoupling capacitor, and only then connecting to the plane, does not make best use of plane capacitance.

Whenever two capacitors are connected in parallel, a high-Q (i.e. sharp) high impedance resonance is created which could compromise power impedance at that frequency. This is easily dealt with on PCBs which have a dozen or more decoupling capacitors, since for every sharp high-Z resonance there are a number of alternate current paths with low-Z, which will swamp it. It may be a good idea to fit decoupling capacitors of 10nF to 100nF in large areas of planes which are devoid of ICs, to help this swamping process. Parallel resonances are very sharp and often don't correspond to any harmonics so have no effect, but unless it is *known* that this will also be the case for a new PCB (and that no-one will ever alter its clock frequencies) it is risky to ignore their potential for upset.

Parallel resonance problems are more likely to occur where only a few decoupling capacitors are used, for example where a small circuit area is powered from a dedicated power plane. It may be controlled by fitting a low-value resistor (say 1 Ω) or small ferrite bead (preferably using SMD packages and short traces) in series with one lead of the larger value capacitors. Alternatively, adding a number of additional capacitors with differing values should help.

The sudden change in characteristic impedance at the edges of parallel PCB planes creates resonances at integer multiples of half-wavelengths. For example, the first such resonance for a 150mm width or length would be around 500MHz for a bare PCB, moving down in frequency as the PCB was loaded with decoupling capacitors (which slows the velocity of propagation in the planes). This was the reason for recommending non-square rectangular plane shapes (and non-simple aspect ratios) in an earlier section. The resulting high impedances at various areas of the PCB can be controlled by fitting lots of decoupling capacitors, so it is only likely to be a problem for circuits operating at high frequencies with large planes and few decoupling capacitors. There is a suggestion that fitting 1 to 10nF decoupling capacitors around the edges of planes can reduce this effect.

Figure 5G shows a time domain view of how good power supply decoupling functions in an example situation. The first nanosecond or so of transient current can only be provided by local 0V/power plane capacitance, with from 1ns to 3ns being provided by SMD ceramic capacitors up to 10nF located



nearby.

Larger (or further away) capacitors are only able to contribute to the current demand after at least 3ns. "Bulk" capacitors (e.g. tantalums) only provide significant current after 20ns or so, even if nearby (non-ceramic dielectrics and electrolytics are slow to respond to transient current demand due to dielectric absorption effects, also known as dielectric memory or dielectric relaxation).

A PCB process is available that uses a special dielectric between adjacent 0V and power planes to increase their capacitance and eliminate the need for most of the smaller values of decoupling capacitors. Three-terminal or "feedthrough" SMD decoupling capacitors have much higher SRFs than regular two-terminal capacitors, but are more expensive. There are also laminar capacitor components (such as the Micro/Q range) made to fit under leaded ICs, which are also expensive and perhaps best used in attempts to improve existing PCBs without relaying them.

5.5.3 Decoupling without power planes

One way to achieve possibly adequate decoupling without a power plane is to connect one end of an IC's decoupling capacitor to its power pin with very short fat trace, then connect that end of the capacitor to the power distribution via a thin trace (to create some inductance) or ferrite bead, rated for the IC's current. Three-terminal or high-specification capacitors may be used to advantage so that a high SRF is achieved with a single decoupler. This technique still requires a 0V plane. Where a large number of ferrite beads or expensive capacitors are required, multi-layer boards may prove to be more cost-effective and require less area.

5.6 Transmission Lines

Transmission lines maintain a chosen characteristic impedance, Z_0 , from a signal's source to its load, and (as discussed in Part 2 of this series) unlike all other interconnections do not resonate however long they are. Transmission lines can easily be made on PCBs by controlling materials and dimensions and providing accurate termination resistances at source and/or load. They may also be extended off the PCB (if necessary) with appropriate controlled-impedance cables and connectors.

Comparing the length of a PCB trace conductor with the wavelength of the highest frequencies of concern in the relevant medium (e.g. FR4), or with the rise- and fall-times of a signal, gives us what is called the ‘electrical length’ of the trace. Electrical length may be expressed as a fraction of a wavelength or as a fraction of the rise- or fall-time. When a conductor is ‘electrically long’, transmission lines need to be used to maintain the frequency response (sometimes called ‘flatness’) or to prevent excessive distortion of the waveshape. For high-speed signals on PCBs, transmission line techniques are required for all electrically long traces both for signal integrity and EMC.

The crude rule of thumb are that a conductor is electrically long when it exceeds one-seventh of the shortest wavelength of concern, or when the time that the leading edge of a signal takes to travel from the source to the furthest receiver exceeds half of its rise or fall times. Consider Fast TTL, which is *specified* as having 2ns risetimes. The dielectric constant of FR4 at high frequencies is around 4.0, which gives a signal velocity of 50% of c , or 1.5×10^8 m/s, equivalent to a trace propagation time of 6.7ps/mm. In 2ns a signal in an FR4 PCB would have travelled about 300mm, so it appears that Fast TTL signals need only use transmission lines for traces of 150mm or longer. Unfortunately, this answer is wrong. The ‘half risetime’ rule is very crude and can lead to problems if its shortcomings are not understood.

Databook specifications for output rise/fall times are *maximum* values, and devices almost always switch a lot faster (assume at least four times faster in the absence of actual data). It is best to measure a number of samples from different batches, and obtain an agreement from the device manufacturer that he will warn well in advance of mask-shrinks. Also, the inevitable capacitive loading from connected devices reduces the propagation velocity from what would be achieved on the bare board. So transmission lines should be used for much shorter lengths of trace than suggested by the above rule, merely to achieve adequate digital signal integrity. Taking these two issues into account, we may find that specified 2ns rise/falltime signals should use transmission lines for traces that are longer than 30mm (and possibly even less).

Transmission lines are often used for clock distribution and high-speed busses; 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 travel very long distances.

Most transmission lines are used to preserve the waveshape of high-speed signals, and to reduce their emissions, but transmission-line techniques work in just the same way to reduce the amount of external fields *picked-up* by a trace, so are valuable for EMC immunity reasons too. It may help to use transmission lines for low-bandwidth signals (e.g. analogue instrumentation) to prevent their contamination by high-frequency fields in their environment (which could be inside a product), since analogue devices are particularly prone to demodulating RF at hundreds of MHz (refer to Part 1). When designing a transmission line for immunity purposes, the ‘shortest wavelength of concern’ or the ‘highest frequency of concern’ is the important parameter.


IEC 1188-1-2: 1998 [2], gives a wealth of details on constructing a wide variety of transmission lines with PCB traces, plus how to specify their manufacture and check quality at goods-in. [1], [3], [4], and [5] are also very helpful with this large and detailed topic, so only the two most common types are described below.

The first example is a *surface microstrip* (see Figure 5H), and its Z_0 is given in ohms by:

$$Z_0 = \frac{87}{\sqrt{\epsilon_r + 1.41}} \times Ln \frac{5.98H}{0.8B + C}$$

where ϵ_r is the dielectric constant of the substrate (typically 4.4 for FR4 at 100MHz), B is the trace width, C is the thickness of the copper material used, and H is the substrate thickness.

Its propagation velocity in ns/metre is: $3.335\sqrt{0.475\epsilon_r + 0.67}$.



The second example is the symmetrical stripline (Figure 5J), which uses two reference planes:

$$Z_0 = \frac{60}{\sqrt{\epsilon_r}} \times \text{Ln} \frac{1.9H}{0.8B + C}$$

The propagation velocity for a symmetrical stripline in ns/metre is: $3.335\sqrt{\epsilon_r}$.

Striplines are slightly slower than microstrip, but have zero forward crosstalk and much less off-board leakage, so are best for EMC.

[1] gives correction factors for the above formulae to compensate for capacitive loading (typically a few pF per gate):

$$Z'_0 = \frac{Z_0}{\sqrt{1 + \frac{C_d}{C_0}}}$$

where C_d is the sum of all capacitive loads, Z_0 is the original

characteristic impedance (unloaded) of the line, and C_0 is the characteristic capacitance of the (unloaded) line obtained from the basic formulae given in [2].

Velocity is slowed according to the formula: $V' = \frac{V_0}{\sqrt{1 + \frac{C_d}{C_0}}}$ from [1], where V_0 is the original

(unloaded) velocity. A constant “gates per unit length” is preferred for the layout of an array of load devices, rather than bunching them together, although it may be possible to adjust the line dimensions for different portions of the trace so that the same Z_0 is maintained all along its length, even where load devices are bunched together.

The highest-speed (or most critical) signals should run adjacent to a 0V plane, preferably one paired with a power plane for decoupling. Less critical signals may be able to be routed against a power plane where the power plane has been adequately decoupled and is not too noisy (i.e. has been properly decoupled, see earlier). Any such power plane must be the one associated with the signal’s ICs. Striplines routed between two 0V planes (one or both of which is paired with a power plane for decoupling) give the best signal integrity and EMC.

Transmission lines must not have any breaks, gaps, or splits in any of the planes they are routed over, as these cause sudden changes in Z_0 . They should also stay as far away as possible from any breaks, gaps, splits, or plane edges, by at least ten times their trace’s width. Low crosstalk requires spacing

adjacent transmission lines by at least three times their trace widths. A very critical or aggressive signal (e.g. a radio antenna connection) may benefit from using a symmetrical stripline with a row of closely spaced vias between its two 0V planes all along each side, 'walling it off' from other traces and creating a coaxial type of structure in the PCB. This requires a different Z_0 formula from those above.

The two transmission line types above require two or more PCB layers, so can be costly to achieve in high-volume low-cost products (although in volume a 4-layer PCB should cost no more than 20% more than a 2-layer). Balanced and co-planar line types can be constructed on a single PCB layer, so may be a solution where high-speed signals must use lowest-cost PCBs. Single-layer transmission lines will take between two and three times the area as microstrip or stripline, so be prepared for their real estate demands. Also beware of saving so much cost on the PCB that the cost of the enclosure shielding and filtering has to be increased. It is a general rule that solving an EMC problem at enclosure level costs between 10 and 100 times more than it would have if it were solved at PCB level. So when trying to pare costs to the bone by reducing the number of PCB layers, allow the time and cost for a couple of additional PCB iterations to get the EMC and signal integrity within specification and within budget, and also allow for additional PCB area.

5.6.1 Changing layers

High-speed or other critical transmission lines should not change layers. This means routing clock distribution first, moving components around to achieve the smallest area of highest-speed circuitry and traces. High-speed busses, fast data communications, and the like are routed next, still sticking to one layer, and then everything else (less critical for signal integrity or EMC) is routed around them, changing layers as necessary. Where there is no reasonable alternative to changing the layers of a critical transmission line, a decoupling capacitor (with a suitable frequency response) should be fitted, with its vias linking all the relevant power and 0V planes, near to the point where the signal changes layers.

Keeping to the same layer is easy when using surface-mounted devices with microstrip transmission lines on the same side of the PCB. Stripline is less leaky than microstrip – but this would mean changing layers, which is generally a bad thing. (Microwave circuit designers often employ microstrip with surface-mount devices with leads that are exactly the same width as the PCB transmission line impedance (usually 50Ω), but they also generally fit each gain stage inside its own milled pocket in an aluminium housing. Such techniques are not generally suitable for computer and DSP boards.) Changing layers is usually necessary, if not to use striplines for their beneficial effects, then because of trace densities, so how can we mitigate its effects?

We already have at least one decoupling capacitor associated with each IC (see 5.4 above) so we can change layers near to an IC, but we must consider the electrical lengths of the portions of the signal path that do not share the stripline layer. A crude rule of thumb is that these portions should not have an electrical length longer than one-eighth of the rise time (approximately one-thirtieth of the shortest wavelength of concern). Where very large changes in Z_0 can occur (e.g. when using a ZIF or other IC socket) it would be better to aim for less than one-tenth of the rise time. Use these rules to determine the *maximum* length, and keep well within this length wherever possible.

So for signals specified as having 2ns rise time signals we should probably change layers no further than 10mm from the centre of the IC's body or the centre of the line termination resistor. This includes a 'safety factor' of 4 to allow for the actual edge rate of the signal being faster than the data sheet maximum. At least one decoupling capacitor, which connects all the relevant powers and grounds together (respectively) should also be within a similar distance from any transmission line that changes layers. Such short lengths are often difficult to achieve with larger ICs, and this reveals some of the compromises inherent in modern high-speed PCB layout. This also reveals at least one reason why physically smaller ICs are preferred, and why bonding techniques such as BGA and flip-chip (which reduce the distance from the PCB trace to the silicon itself) are continually being developed and improved upon.

5.6.2 Simulation and prototype testing

Because of the variations in the types of ICs, and the applications they find themselves in, some engineers will find these rules of thumb not tough enough, and some will wonder whether they are over-engineered, but that is the function of a rule of thumb, after all.

Computer-based circuit simulation techniques that calculate EMC and/or signal integrity based on parameters extracted from an actual PCB layout are becoming more capable, and their use is recommended instead of the crude rules-of-thumb expressed here. However, remember that device switching speed are almost always considerably faster than their data sheet specifications, so a simulation that uses data sheet figures will give a false sense of confidence.

Tests with a high-speed oscilloscope and probing system should be carried out on the first PCB prototypes to see whether the waveshape is good enough. A waveshape that does not distort as it travels around the PCB is the goal, and merely following these rules of thumb is unlikely to achieve such perfection, although the result may be good enough. Close-field probing with a single-turn loop, using a high-speed 'scope and/or a spectrum analyser, is another good way to detect signal integrity or EMC problems at prototype board level. The techniques involved in prototype testing are not discussed further here.

Even when using sophisticated modelling or simulation techniques, always perform signal integrity and EMC checks on early prototypes.

5.6.3 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 by $\pm 25\%$. Controlled ϵ_r grades of FR4 are available at little or no extra cost, but PCB manufacturers may not use these grades unless specifically requested.

PCB manufacturers work with standard thickness laminations ("prepregs"), and their thicknesses should be discovered (along with their manufacturing tolerances) before design starts. The trace widths can then be chosen to achieve the required Z_0 for the available range of dielectric thicknesses. Trace widths after PCB processing are usually about one thousandth of an inch less than those used on the photoplots. Ask what thickness to add to the drawn traces to achieve the required finished trace widths.

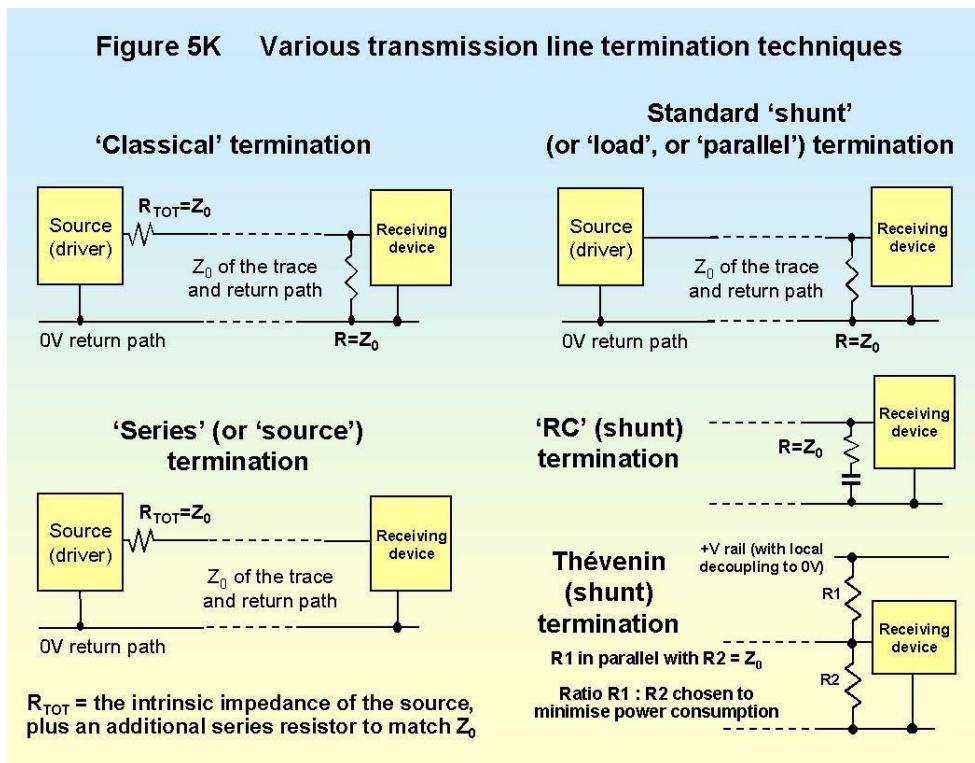
For signal frequencies greater than 1GHz it may be necessary to use other dielectric materials than FR4, such as those used for microwave applications (e.g. Duroid from Rogers Corporation Inc., or a number of more modern dielectrics).

5.6.4 Terminating transmission lines

"Classical" RF transmission lines are terminated both at signal source and load by impedances equal to their Z_0 (allowing for the internal impedances of source and load devices). Although an ideal and sometimes necessary technique, it halves the received voltage – so most ordinary analogue and digital circuits use low-Z sources and high-Z loads with the line only terminated at one end, to preserve signal levels.

RF engineers often use reactive components or even lengths of trace as line terminations, but the terminations for wideband analogue and conventional digital signals require individual resistors, preferably SMD types for their excellent high-frequency performance. To get the best from SMD resistors they must be connected to the reference planes using low inductance techniques as shown by Figure 5C.

Figure 5K shows the common termination techniques. Classical RF termination is still often used for high-speed signals such as fast backplane systems.



Where signals are restricted to a single PCB, series (source) termination may be used at the driver end of a transmission line, with the resistor chosen so that in series with the impedance of the output driver it matches the line's Z_0 . This method has the advantage of consuming little power, and is most suitable for lines with a single load device at their far end. Where other loads exist along the length of the line they experience "reflected wave switching" and their response may need to be slowed to prevent false clocking.

Parallel (shunt, or load) termination at the very far end of a line is used where there are a number of devices spread along the length of the line and they need to respond most quickly, and achieves "incident wave switching". Figure 5K shows the termination resistor connected to the 0V plane, but some logic families use other reference voltages (e.g. the positive plane for ECL). Parallel termination dissipates a lot of power, and may also load some IC outputs too heavily.

Alternative types of parallel termination include "Thévenin" and RC. Thévenin uses resistor values designed so that their parallel resistance is Z_0 and they would provide a DC voltage at their junction equal to the average line voltage, to minimise power dissipation. Thévenin termination needs a properly decoupled power plane at all frequencies of concern so needs decoupling capacitors nearby. RC termination uses capacitor values between 10pF and 620pF (typically) and only terminates the line for high frequencies. Because of the problems of capacitors (discussed earlier) it may be more difficult for an RC termination to equal the highest frequency performance of a parallel resistor or Thévenin termination.

"Active termination" uses a voltage regulator to drive an additional power plane at the nominal average value of the digital signals. A parallel line termination connects to this plane, which must be properly decoupled for the frequencies of concern. Electrically equivalent to the Thévenin method, this can save power by running the voltage regulator (which needs to be able to source as well as sink current) in Class AB.

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 be 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 “star” connecting a number of individual series-terminated transmission lines, either use one termination resistor chosen so that the total source resistance equals the parallel combination of all the starred lines, or else use one resistor to match each line. The latter technique should be better. The star configuration may also be used to drive multiple parallel-terminated lines. In either case, the signal source must be capable of driving the parallel combination of all the lines' Z_0 s.

It is generally better to choose higher values of Z_0 to reduce signal currents and reduce radiation from the traces. Many ordinary CMOS or TTL ICs were never designed for driving transmission lines, and have neither the drive capability or an output impedance that is equal for both sourcing and sinking. Such devices may be able to use series, Thevenin, RC, or active terminations on high-impedance lines, but the best method to use, and the line impedance, may be difficult to predict for a given logic family.

However, an increasing number of devices are becoming available to drive transmission lines, and the increasing range of LVDS and similar devices is making clock and bus driving much easier and easing EMC problems. Backplane bus driver ICs are available with 25Ω output impedances, suitable for “star” driving four individual 100Ω , or six 150Ω lines. Some devices now have on-chip DC/DC converters which cause their unloaded outputs to achieve double the correct logic levels, so that when operated into a classically terminated line the received logic levels are correct.

5.6.5 Layer "stack up"

The above section on decoupling shows it is good EMC practice to provide 0V and power planes on adjacent layers and to maximise their capacitance by using a thin dielectric (say 0.15mm) between them. The above section on transmission lines shows that proximity to a reference plane is important for high-speed traces. We can put this all together to decide how to stack up our PCB layers.

Four-layer PCBs often have their layers stacked as follows:

- 1) Microstrip transmission lines and other critical signals
- 2) 0V plane
- 3) +5V plane
- 4) Non-critical signals

Where more signal layers are required, a 0V and power plane "core" should be retained. Additional layers of high-speed signals may need additional 0V planes to be added, but high-speed clocks and data busses and similarly aggressive or very critical traces should not swap layers.

Here is one of a number of possible stack-ups for an 8-layer computer motherboard:

- 1) 0V plane
- 2) Most critical “offset striplines” and other signals, routed at 90° to layer 3) to reduce crosstalk
- 3) Most critical “offset striplines” and other signals, routed at 90° to layer 2) to reduce crosstalk
- 4) 0V plane
- 5) +5V plane
- 6) Non-critical signals routed at right angles to layer 7) to reduce crosstalk
- 7) Less critical “offset striplines” and other signals, routed at 90° to layer 6) to reduce crosstalk
- 8) 0V plane

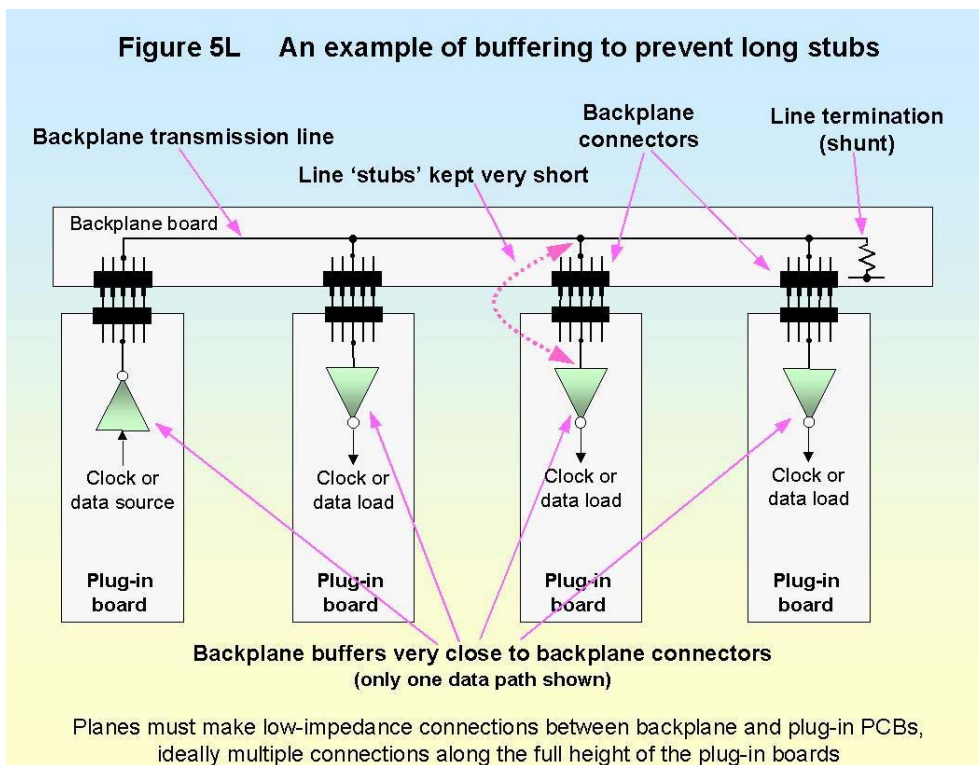
5.6.6 Joints, stubs, and buffers

The above has treated transmission lines as if they were all point-to-point connections, so we need to address bussed systems such as RAM arrays, and situations where several cards interconnect at high speeds, such as backplane systems.

A length of trace that springs off from a joint with the main trace is called a 'stub'. For memory arrays, the usual PCB layout technique is to route busses horizontally on one layer, 'via-d' through to another layer with vertically routed stubs to connect to the array devices. To preserve the transmission line the electrical length of the stubs must be kept to under one-eighth of the rise time (and preferably much less). Don't forget that the important parameter is the real switching rate of the signals, not the data sheet specification of the drivers. If in doubt, assume the drivers switch four times faster than their data sheet maxima.

The stub length used for the calculation should include the distance from the end of the trace (the IC's soldered pin) to the centre of the IC itself. Where the permissible stub length is too short for traditional 'horizontal and vertical' array routing, daisy chain routing should be used instead. Daisy chain routing is generally better for high-speed signals in any case, especially when the traces remain all on one layer. Daisy chain routing means that the bussed traces go from the source directly to each load in turn. Abrupt changes in trace direction should be avoided, with gentle curves or large chamfers used instead. In an incident wave system, the daisy-chained traces would end in the parallel, Thevenin, RC, or active termination resistors.

When electrically long stubs can't be avoided, buffers should be fitted close to the main trace to minimise the stub lengths. This is often used in backplane systems, where a number of plug-in cards must all run from the same clock lines and data busses, as shown in Figure 5L. The clock buffers must all be fitted very close to the backplane board connectors, and as signal speeds and data rates increase it is more common to find that matched-impedance backplane connectors are needed. Where a plug-in card only has one or two ICs that need to connect to the backplane clock and data lines, by placing them close the backplane connector it may be possible to do without the buffers.



Buffering is also a good technique for reducing the loading on a transmission line. For example, where there are ten plug-in cards each with ten ICs, all receiving one signal, their combined load capacitance can be around 400pF. The signal and return currents for this high value of capacitance have a long way to flow, increasing the likelihood that they will create EMC problems. Buffering the signal at each card means that the main line is only loaded by around 40pF, while the signal and return currents for the ten devices on each card now flow only in that card, improving signal integrity and reducing EMC problems.

Carrying high-speed signals through connectors and backplanes, it is important (vital for transmission lines) to maintain the same physical structure. For example, striplines in plug-in boards should be continued as striplines in the backplane, (although it is possible with some degradation in signal integrity to swap from one type of transmission line to another as long as the trace dimensions maintain the same Z_0). Where transmission lines entering a backplane connector are routed against a power plane, that power plane should be continued through the connector into a power plane in the backplane and then to the associated power planes in the other cards using that signal. The interconnections between the power planes in the boards and the backplane should be designed in the same way as for the 0V return planes. Some boards may find that their optimum backplane connector pinning needs to be: 0V, signal 1, +5V, signal 2, 0V, signal 3, +5V, signal 4, 0V,etc.

5.6.7 Segregation in backplane systems

Section 5.1 above said high-speed devices should be kept in the middle of their segregated area, well away from any PCB or reference plane edges, or connectors. The backplane system described above and in Figure 5L places the fastest ICs close to the backplane connector but does not compromise the earlier rules if the backplane is designed to be an extension of the plug-in boards at high frequencies.

This requires RF bonding the reference planes in the backplane to the corresponding reference planes in the plug-in boards, so that at the highest frequency of concern there appears to be no impedance discontinuity between them. Using shielded connectors helps - their shields should bond their mating halves to each other in 360° (refer to Part 2 of this series), and also bond all along their length to the 0V reference plane on both sides. Whether shielded connectors are used or not, there should be one signal return pin for every one or two (at most) signal or power pins in the connector, and these return pins should be spaced fairly regularly along the entire length of the connector. Most designers would place the return pins according to a regular scheme, but there is some evidence that a randomised allocation can have benefits. Impedance-matched connectors will almost always have a return pin alongside every signal pin, in any case.

It is important to make sure that all the high-speed devices associated with the backplane connector are closer to the middle of the connector, and do not go near to the outer edges of the boards, near the ends of the connector. It is best if the backplane connector extends to occupy the entire length of the card edge, but if it can't it should still extend well to both sides of the area of the high-speed devices associated with the backplane connector.

5.7 Useful references

- [1] M. Montrose, *EMC and the Printed Circuit Board, design, theory, and layout made simple*, IEEE Press, 1999, ISBN 0-7803-4703-X
- [2] *IEC 61188-1-2 : 1998 Printed Boards and Printed Board Assemblies – Design and use. Part 1-2: Generic Requirements – Controlled Impedance*, www.iec.ch.
- [3] IPC-2141, *Controlled Impedance Circuit Boards and High Speed Logic Design*, Institute for Interconnecting and Packaging Electronics Circuits, April 1996, www.ipc.org.
- [4] Tim Williams, *EMC for Product Designers* 3rd edition, Newnes 2001, ISBN: 0-7506-4930-5, www.newnespress.com.
- [5] H W Johnson and M Graham, *High Speed Digital Design, a Handbook of Black Magic*, Prentice Hall, 1993, ISBN 0-13-39-5724-1.