



Another EMC resource
from EMC Standards

EMC techniques in electronic design Part 5 - Printed Circuit Board (PCB) Design and Layout

Helping you solve your EMC problems

Design Techniques for EMC

Part 5 — Printed Circuit Board (PCB) Design and Layout

Originally published in the EMC Compliance Journal in 2006-9,
and available from <http://www.compliance-club.com/KeithArmstrong.aspx>

Eur Ing Keith Armstrong CEng MIET MIEEE

Partner, Cherry Clough Consultants, www.cherryclough.com, Member EMCIA, EMCTLA

Phone/Fax: +44 (0)1785 660 247, Email: keith.armstrong@cherryclough.com

This is the **fifth** in a series of six articles on basic good-practice electromagnetic compatibility (EMC) techniques in electronic design, to be published during 2006-7. It is intended for designers of electronic modules, products and equipment, but to avoid having to write modules/products/equipment throughout – everything that is sold as the result of a design process will be called a ‘product’ here.

This series is an update of the series first published in the UK EMC Journal in 1999 [1], and includes basic good EMC practices relevant for electronic, printed-circuit-board (PCB) and mechanical designers in all applications areas (household, commercial, entertainment, industrial, medical and healthcare, automotive, railway, marine, aerospace, military, etc.). Safety risks caused by electromagnetic interference (EMI) are not covered here; see [2] for more on this issue.

These articles deal with the practical issues of what EMC techniques should generally be used and how they should generally be applied. Why they are needed or why they work is not covered (or, at least, not covered in any theoretical depth) – but they are well understood academically and well proven over decades of practice. A good understanding of the basics of EMC is a great benefit in helping to prevent under- or over-engineering, but goes beyond the scope of these articles.

The techniques covered in these six articles will be:

- 1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
- 2) Cables and connectors
- 3) Filtering and suppressing transients
- 4) Shielding (screening)
- 5) PCB layout (including transmission lines)**
- 6) ESD, surge, electromechanical devices, power factor correction, voltage fluctuations, supply dips and dropouts

Many textbooks and articles have been written about all of the above topics, so this magazine article format can do no more than introduce the various issues and point to the most important of the basic good-practice EMC design techniques. References are provided for further study and more in-depth EMC design techniques.

Table of contents for this article

5.	Printed Circuit Boards (PCBs)	2
5.1	Introduction	2
5.1.1	Real financial benefits	2
5.1.2	The scope of this article	2
5.1.3	Basic techniques are covered here	3
5.2	Segregation	3
5.3	Interface analysis, filtering, and suppression	6
5.4	0V and power planes	9
5.4.1	General plane design issues	9
5.4.2	Only use thermal break pads (thermal reliefs) when really necessary	11
5.4.3	RF-bonding planes to components, conductors and chassis	11
5.4.4	Don't split 0V planes any more (and what to do, if you do)	13
5.4.5	Traces routed close to plane edges, or across plane splits	14
5.4.6	Can't afford multilayer PCBs?	15
5.5	Power supply decoupling	15

5.5.1	General decoupling design rules.....	15
5.5.2	Decoupling with ferrites.....	17
5.5.3	Benefits of 0V/power plane pairs.....	17
5.5.4	Dealing with PDS resonances.....	18
5.6	Matched transmission line techniques.....	19
5.6.1	When to use matched transmission lines.....	19
5.6.2	Correcting for load capacitance.....	21
5.6.3	Choosing the dielectric materials for the stack-up.....	22
5.6.4	Terminating transmission lines in a <i>matching resistance</i>	22
5.6.5	Differential matched transmission lines.....	22
5.6.6	Transmission line routing.....	24
5.6.7	Stubs and branches.....	24
5.7	Layer stacking.....	26
5.8	References.....	28
5.9	Acknowledgements.....	28

5. Printed Circuit Boards (PCBs)

5.1 Introduction

5.1.1 Real financial benefits

As Part 0 of [3] made clear, the most cost-effective EMC techniques are those applied early in the design process, at the lowest level of assembly. Ideally, this means in the design of the integrated circuits (ICs) and other semiconductors – but (with a few exceptions) the semiconductor industry that provides standard parts ignores EMC completely and leaves the consequences of dealing with their design decisions to their customers.

FPGA and ASIC manufacturers usually make a range of EMC options available to their customers, and their use is recommended – but not all EMC issues can be dealt with in the semiconductors themselves. Techniques described in Part 1 of [3] help with the choice of semiconductors and passive components.

The PCB is the next lowest level of assembly, and has the advantage that any/all EMC issues can be dealt with in its design and construction. After the EMC design of the electronic circuits to be placed on the PCBs has been addressed (see Part 1 of [3]), and the components chosen (or designed, in the case of FPGAs and ASICs) the design and layout of the PCB is the most cost-effective level to deal with EMC.

Unfortunately, as discussed in section 0.1 of [3], many project managers seem to think that the PCB with the lowest bill of materials cost ('BOM cost') will result in the most profitable product, when a moment's thought will show that this is not necessarily the case. (In fact, there is almost *never* any *direct* relationship between the BOM cost and the financial success of a product.) For EMC, this misguided approach usually leads to boards that have poor EMC performance, discovered late in the project when the product fails EMC tests, and fixed in a desperate hurry using very costly components and manufacturing techniques whilst missing sales due to the inevitable delay in market introduction.

The consequences of this routine are painfully engraved on the minds of most electronic designers worldwide, and also show up in the poorer financial performance of their employers. On the other hand, EMC test laboratories enjoy having products back for retesting time and time again – as the design teams struggle to fix EMC problems, whilst their design freedom is seriously restricted by the perceived urgency and the large sums of money already spent on production tooling and software development.

EMC experts the world over have for decades recommended taking all necessary EMC precautions in the design and layout of the PCBs. Usually the extra design time is insignificant, and the extra unit manufacturing costs are negligible, certainly when compared with the financial benefits of timely market introduction and reduction in last-minute design changes. Indeed, there are many anecdotes of products that were redesigned for EMC with more sophisticated boards and a higher BOM cost – but which nevertheless enjoyed a lower overall cost-of-manufacture and much greater market success.

5.1.2 The scope of this article

Earlier parts of this series have tended to require two, or even three issues of The EMC Journal. Since I have recently published a book on basic and advanced EMC techniques for PCB design and layout [4], I could easily fill three or four issues just copying the basic material from my book. This doesn't seem a very sensible thing to do, so instead this Part of this series will be a text version of the PCB section of the EMC course I teach at the University of Manchester (UK) to graduate students on the Electronic Instrumentation Systems (EIS) MSc

Course (<http://www.eee.manchester.ac.uk/research/groups/sisp/postgraduate/taught/>). These are basic EMC techniques only, as discussed in 5.1.3.

5.1.3 Basic techniques are covered here

This article briefly describes the basic EMC techniques for PCB design and layout that are now generally needed for all types of PCB. As semiconductor technology continually advances, PCB technology also has to advance, and so the basic EMC techniques for PCBs are always advancing too. EMC techniques that were only needed for 'advanced' boards ten years ago, are now routinely needed for all boards (for cost-effective design, see 5.1.1). As a result, there are significant differences between this article and Part 5 of [1].

Engineering is all about compromise. This series covers a large number of good EMC design techniques, but in real projects some of them may be impractical, or too costly, or inappropriate. It is the engineering compromises necessary for real designs that makes EMC design really interesting. So good EMC practices should always be treated as recommendations – or as lists of issues to be considered. In general, where a good EMC design technique described in this article is not used for whatever reason – the risk of suffering EMC problems increases, unless alternative (effective) methods have been used instead (e.g. shielding and filtering the overall enclosure).

Because the material in this article covers basic PCB EMC techniques, the general advice is to only deviate from them for good technical reasons. If they are not applied for financial reasons (e.g. BOM cost), the financial argument will probably be incorrect (see 5.1.1.).

More advanced PCB EMC design techniques may be required for...

- Reducing or eliminating the cost or weight of enclosure shielding and/or filtering
- Achieving good sensitivity and range for integrated radiocommunications (e.g. GSM, 3G, DECT, IEEE 802.11x, Wi-Fi, Bluetooth, ZigBee, etc.)
- Achieving good sensitivity for GPS receivers with nearby antennas
- Very high-speed circuits (e.g. PC motherboards)
- Using the latest silicon technologies (e.g. ICs made on 90nm, 65nm, 45nm process lines)
- Using chip-scale packaging (e.g. to make very small and/or low-cost products)
- Reducing time-to-market without increasing risks of non-compliance

5.2 Segregation

The first and most cost-effective EMC technique (it is free, if done early in a project) is segregation. Firstly identify all the components and conductors (cables, connectors, PCB traces, etc.) that lie in the 'Inside World' and 'Outside World' EM zones.

As Figure 5A shows, the Outside World is where the designer has no control over some or all of the electromagnetic (EM) disturbances that could occur. The Inside World is where the designer has the ability to control all of the EM disturbances (even if some of them are not actually controlled).

It is not always obvious what should be included in the Inside World and what in the Outside World. EM disturbances are controlled by shielding [5], filtering [6], and other 'EM mitigation' techniques such as surge suppression [6] or power quality improvement methods (see Part 6 of this series, published later in 2007) – so where components and/or conductors are not protected from all of the external EM disturbances by appropriate EM mitigation measures – they are in the Outside World.

Where an equipment does not have a shielded and filtered overall enclosure, a 0V plane in the PCB (see 5.4) will provide some useful shielding for components and traces that are not too high, or too close to an edge. So low-profile components and traces not connected to external cables are treated as being Inside World as long as they are placed on – and surrounded by – a PCB 0V plane used as their circuit's 0V reference. Components and traces completely covered by a board-mounted shielding-can (see section 4.4 of [5]) are also Inside World.

However, ribbon cables, flexible jumpers, unshielded connectors, and traces that are not protected by a 0V plane should be treated as Outside World if they are within an inadequately shielded and/or filtered overall enclosure (see [5] [6]). They are all 'accidental antennas' as discussed in section 2.2 of [7], which shows, for example, that a flexible jumper, ribbon cable or other wires as short as 50mm are very efficient accidental transmitting/receiving antennas for GSM cellphones at 900MHz and 1.8GHz, and PCS at 1.9GHz.

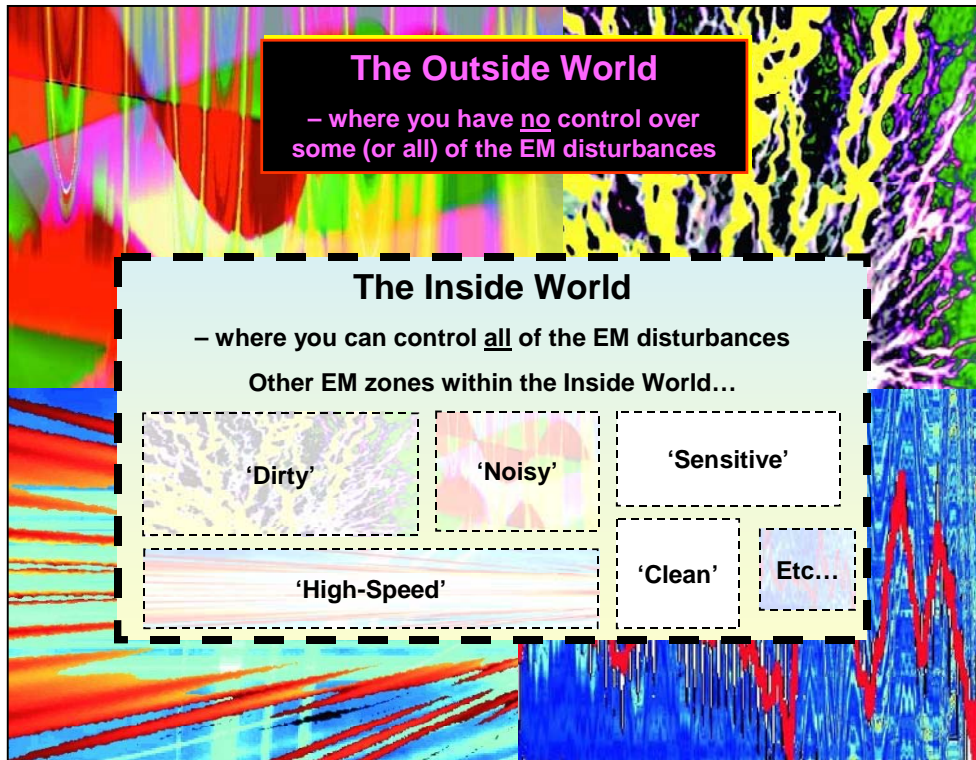


Figure 5A Segregating circuits into Inside and Outside worlds

It also shows that typical inter-board stake connectors and similar unshielded connectors are efficient accidental antennas for GSM at 1.8GHz, PCS at 1.9GHz, 3G (UMTS) around 2GHz, Wi-Fi, Bluetooth, ZigBee and microwave cookers at 2.45 GHz, and UWB (e.g. 'Wireless USB') at 3.1GHz and above.

The Inside World zone is also segregated into further zones, which are usually given names such as dirty, clean, high-speed, quiet, etc., depending on the circuits concerned.

- 'Clean' or 'quiet' circuits are especially sensitive to EM disturbances, for example analogue signal amplifiers and receivers (from DC through RF to microwave).
- 'High-speed' generally means digital signal processing, with its microprocessors, RAM and ROM, and clocks and data buses. It can also mean RF transmitters.
- 'Noisy' generally means switch-mode power converters (AC-DC, DC-AC, AC-AC or DC-DC) and all electromechanical contacts such as switches, relays, contactors, commutators and sliprings.

Each of the EM zones that have been identified is to be segregated from all of the other zones, both mechanically and electrically, starting at the earliest design phase (ideally the 'blank sheet of paper' stage). The segregation should be clearly shown on the schematic and all other relevant drawings, for example by drawing 'dotted boxes' around the zones and labelling them with their agreed names.

All of the components, traces, connectors and other conductors within a dotted box *must* remain totally within the area set aside for their EM zone on the PCB. Only the essential inter-zone connections are allowed to enter or exit an EM zone – and they might all need filtering or some other EM mitigation technique applying (see 5.3).

The segregation should be rigorously maintained throughout the design of the board layout; wiring harnesses; mechanical packaging, etc., taking the three-dimensional structure of the final product assembly fully into account. Because designers work with their products dismantled, they sometimes get caught out by zones that are physically very well segregated in the dismantled state, but not when finally assembled.

For example, when the product is finally assembled, a cable that belongs in one EM zone might lie too close to a circuit in a different zone; or two PCBs with different zones might end up being so close to each other that they couple excessive noise from one to the other through stray capacitances and mutual inductances (i.e. near-field EM coupling, often called 'crosstalk'). Such problems can generally be solved by shielding cables and areas of the board (section 4.4 of [5]), but it is quicker, easier, and less costly to avoid them in the first place by careful attention to segregating the EM zones, in the final assembly, in three dimensions.

Figure 5B shows an example of a segregated board, with its Inside and Outside Worlds, and shows its Inside World further segregated into further nested EM zones, which may be further subdivided into even smaller

zones. The short black double-headed arrows indicate the essential inter-zone connections, which are the only traces allowed to be routed between segregated zones.

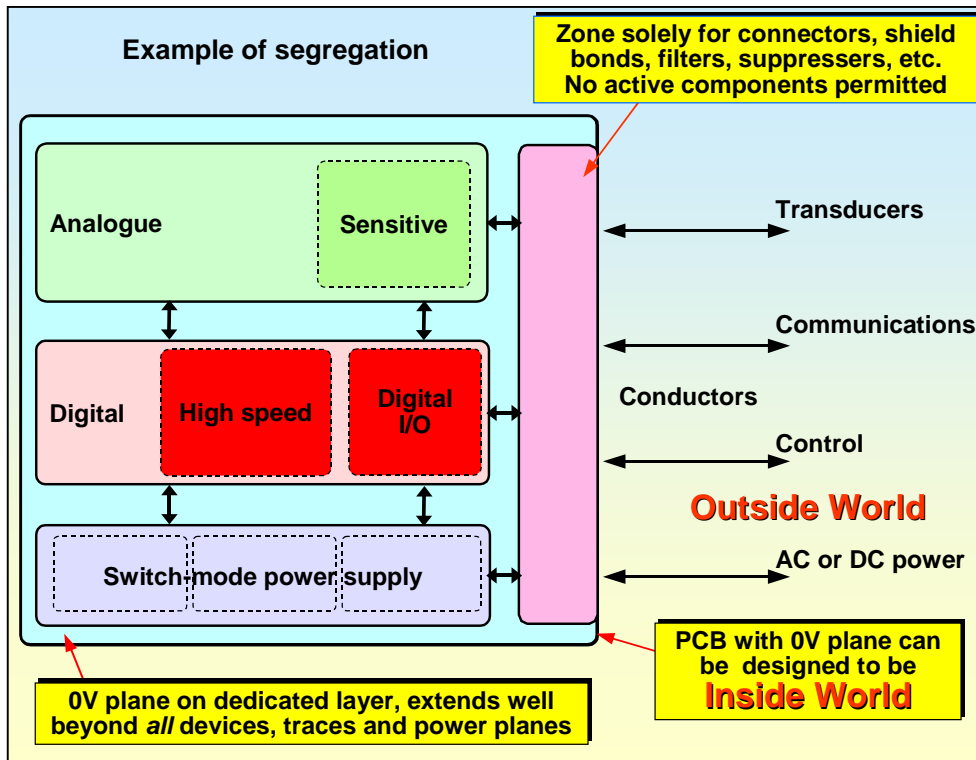


Figure 5B Example of a segregated single-PCB product

Unless well-shielded on the board [5], all high-speed and RF transmitter circuitry should be located close to the centre of the board's 0V plane, well away (at least 50mm) from any off-board connectors or wires, and at least 50mm away from routes taken by external cables or wires when the product is finally assembled.

Any voltage differences between off-board cables will drive them as radiating antennas, and so could cause problems with emissions. Even though a proper 0V plane (see 5.4) develops very low voltages in response to the currents flowing in it, they are not zero, so it is important to place all of the off-board (Inside-Outside World) connections along just one edge of the PCB – keeping them close together – and with no active circuitry located between them, to minimise the voltage difference between them.

All Outside World interconnections will need some sort of EM suppression applying, so it is important that there is a zone set aside for doing just that (see 5.3).

There should be a 0V plane (at least, see 5.4) that underlies all of the Inside World zones and extends beyond their traces and components by as far as possible. It can be cost-effective to make a board larger, where there is room available in the product, just to extend its 0V plane further beyond its components and traces.

Figure 5B shows an example of a product consisting of a single PCB. Where products are split into separate PCBs, some of the unshielded internal board-to-board interconnections might need to be treated as Outside World conductors as discussed above. 2.2 in [7] shows that even conductors as short as a few centimetres can be very effective accidental antennas at frequencies now commonplace, so the shielding effectiveness of the overall enclosure, that protects the internal conductors from such exposure, can be an important issue. Of course, shielding external enclosures, or shielded cables, see 2.6 in [7]) are costly, which is one of the reasons why the single-PCB product is usually the most cost-effective.

A single-PCB product might be the most cost-effective, even if more expensive flexi-rigid PCB technology has to be used to achieve it. Flexi-rigid PCBs with a 'solid' 0V plane (see 5.4) over their whole area (both flexi and rigid) have much better EM characteristics than a number of PCBs interconnected by unshielded connectors, flexible jumper strips, or unshielded cables.

Flexi-rigid assemblies are generally also much quicker to assemble, with fewer assembly errors (hence lower rework costs). Because they don't need board-to-board connectors they are generally more reliable in real life (hence fewer warranty claims). Products using single-PCBs using flexi-rigid PCB technology are often good examples of increasing the BOM cost to achieve a more cost-effective and profitable product.

5.3 Interface analysis, filtering, and suppression

After segregation has been achieved (see 5.2), each interface between each segregated zone should be analysed for all relevant electromagnetic (EM) disturbances, both conducted and radiated, and EM mitigation measures such as filtering (see [6]); shielding (see [5] and [7]); transient protection (see [6]); galvanic isolation such as transformers or optoisolators; etc., added as needed...

- At the zone boundaries between the segregated areas inside the product
- At the zone boundary between the Inside and Outside worlds

Digital control signals are not just clean 1 or 0 levels, they have quite high levels of ground bounce and other digital RF noises on them. So if used to control an analogue circuit in a different EM zone – their traces will generally need filters located at the boundary of the ‘quieter’ zone at the point where they enter that zone, to reduce the amount of digital noise coupled into the analogue signals.

Another common issue is the noise on shared power supplies. Because of the habit of not drawing power supplies as continuous lines on a schematic, the fact that a power rail is used by two different EM zones can be overlooked. If the power rail is not filtered at the boundary of the ‘quieter’ zone – with the filters located at the point where the trace or plane enters that zone – noises from one zone will couple into the signal path in the other zone.

Devices which interface *between* two zones – for example: A/D or D/A converters; filters; opto-isolators, etc. – should be positioned at the nearest edge of one of their zones, so that the traces that interconnect the two zones are not routed around inside *both* of the zones’ areas.

Inside any zone, even routing just a few millimetres of unsuppressed trace that has come from another zone can cause serious EMC problems. This is because the extremely fast switching edges of modern digital devices have such high frequency components that even very tiny stray capacitances and stray mutual inductances (e.g. the stray couplings between conductors that are just a few mm long), can have a very negative influence.

ICs have the smallest feature sizes of any mass-produced man-made items, which makes them very weak indeed. If they are to be connected to Outside World conductors they should always be protected by some filtering, transient protection, or other suppression or isolation, depending on the EM environment, and should never connect directly. Some serial I/O ICs are available with high levels of ESD protection, but this still leaves many other types of EM disturbance that they need to be protected from. Connecting an IC’s pin directly to an outside world conductor, without EM mitigation to protect it from all EM disturbances it is likely to encounter in its life, is rather like putting your five-year-old child on the bus to the city with a brown bag lunch and bus fare, and instructions not to return until they have earned some money. It is cruel. So never do it.

Figure 5C shows more detail of the recommended PCB layout in the area of the Inside-Outside world interface zone shown on Figure 5B, for a shielded off-board cable that is also filtered (see [6]) with a simple RC or LC layout. This example does not include any transient or surge overvoltage protection, which might be needed as well (or instead) and which should follow similar layout rules.

The shielded connector body should be soldered directly to the 0V plane at multiple points, and the local 0V plane should in turn make a multipoint connection to any shielding being used at that zone boundary. Any 0V plane to shield connection should achieve very low impedance at the highest frequency of concern; see [7] for more on this.

To minimise stray coupling around the series filter elements (R or L), in order to maximise filter performance, it is vital that the series elements are all aligned in a neat row. Where the selected filter components are not available in a small enough package style – place them on both sides of the board. Where even this is not enough, use arrays of resistors or ferrite beads instead of discretely. But never ‘stagger’ their placement on the board – they must *always* be laid out in straight lines.

In Figure 5C the series filter elements actually create the Inside World-Outside World boundary. Every filter, shield or other type of EM mitigation lies on a zone boundary, and if those boundaries get confused the degree of EM control that is achieved overall can be very poor indeed. Visualising the EM boundaries and making sure they are well maintained is one of the keys to professional EM design at the highest level – it is actually all about separating the flows of surface currents, to keep the inside and outside currents apart, as discussed in more detail in [6] and [7].

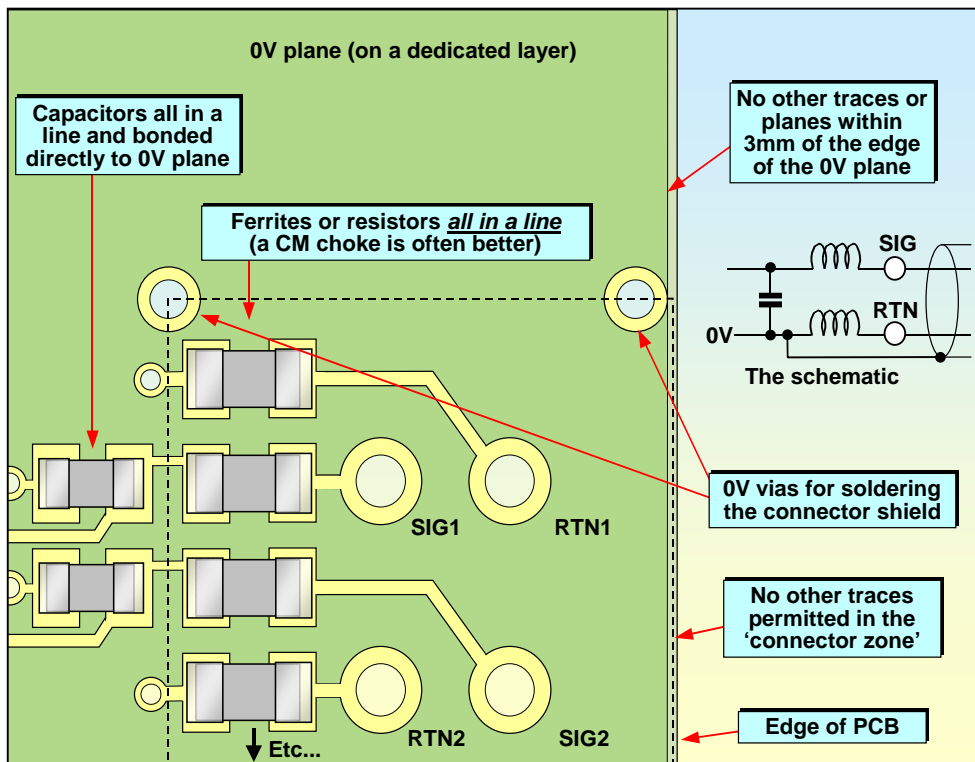


Figure 5C Example layout for a shielded off-board connector, using RC or LC filtering

Never route any traces down the edge of a PCB past an off-board connector, or sneak them between the traces and components associated with the filter and connector pins. The stray coupling between the traces and the unfiltered Outside World conductors (connector pins and/or attached cables) – even if only a tiny fraction of a picoFarad (pF) – can completely destroy the filter's attenuation at very high frequencies. This is why Figure 5C and Figure 5D show a 'Connector Zone' in which no other traces are permitted. This zone extends right to the edge of the PCB.

Figure 5D shows a similar example, to Figure 5C, but this time using a connector to an *unshielded* off-board cable (preferably using twisted-pair conductors, see [7]).

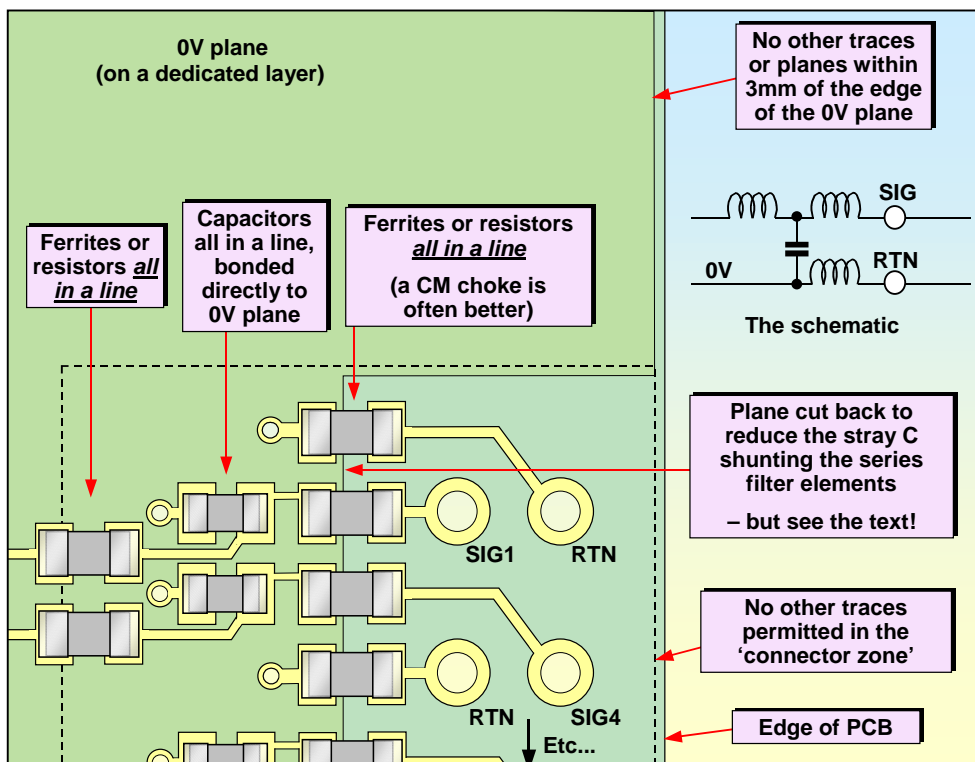


Figure 5D Example layout for an unshielded off-board connector, using Tee filtering

Apart from the differences in the filters (notice the neat lines, as in Figure 5C) and the lack of solder points for the cable connector's shield, the only significant difference from Figure 5C is that the 0V plane has been cut back. The series elements rely on achieving high impedance at high RF frequencies, but the proximity of both of their terminals to a metal plane increases their stray shunt capacitance and reduces their RF impedance. Cutting back the plane, as shown, reduces the stray capacitance shunting the series filter elements, and improves their high-frequency impedance.

However – as discussed in 5.4.1 – we don't like to have any holes in our 0V planes, so this technique always requires a 'judgement call'. If *all* of the EMC techniques described in this article have been correctly applied to a PCB – then above 100kHz or so the currents flowing in the 0V plane will remain very close indeed to their traces and components, mostly remaining within the 'dotted line' boundaries of their EM zones.

In such a happy situation, making a gap in the 0V plane in an EM zone at the edge of the PCB, should cause very little of the plane currents to be diverted from the paths that the laws of physics find most energy-efficient in a different zone. So – on balance – the gap should provide significant benefits for both emissions and immunity.

But where there could be significant levels of plane currents from other zones flowing in the connector zone, a gap in the 0V plane as shown in Figure 5D could – on balance – be counter-productive for EMC. Computer simulation might be able to provide the necessary information to tell whether a gap will or will not be beneficial, but in the absence of that technique, if there is any concern, test-bench experiments using close-field probes are recommended, as early as possible in a project.

Close-field probing techniques, including how to make your own probes and use them with oscilloscopes or spectrum analysers, are described in Parts 1 and 2 of [8], and can be used in an ordinary development setting to test two versions of a prototype, one with a gap in the plane and one without.

It is good practice to make provision for fitting shielding-cans over the most emissive (or most susceptible) ICs or circuit zones, at least, even if it is hoped to do without them, just in case they do turn out to be needed. Planning and designing for such flexibility from the start of a project is very worthwhile, and an example of what John R Barnes [9] calls "wobble room" and I call "anti-Murphy design", based on the well-known Murphy's Law. Since we do not (yet) have accurate computer simulators that will predict the emissions and immunity compliance of a real product from its design drawings and parts specifications alone, it makes good sense to add these little features during design.

Designers who try to anticipate the (unpleasant) surprises that Murphy might have in store for them, reach their design targets and timescales more reliably. They are much less likely to need major redesigns of their circuits and layouts of their boards at (what was supposed to be) the end of the project, when compliance tests were failed and modifications most costly (see Part 0 of [3]).

(Note that as a designer, your manager will criticise you for using anti-Murphy design measures that were not eventually required. But if you don't use them, Murphy will make the project fail, and your manager will criticise you for not thinking ahead to what might go wrong. Long experience with Murphy's law shows that whatever you do, Murphy will always achieve the maximum embarrassment for you, and the greatest criticism from your managers, that he can.)

(The choice is between whether we want to be criticised for getting a successful project to market on time but with a few things in it that turned out not to be necessary – or blamed for not designing thoughtfully enough and delaying a project and increasing costs hugely. Against this we should realise that designers who put their family, friends and health in second place, whilst working all hours to try to salvage a bad design, are often thought of very highly of by their managers (even where it was their bad design that caused the problem), often more so than an engineer who just quietly does a professional job that makes more money for the company.)

The walls of the shielding-can should follow the boundary of the segregated circuit zone it is shielding, so it is also good practice to base the shapes of the segregated zones on simple shapes that can easily be 'canned'. Some of the figures in [5] show examples of shielding-cans that have been used in real products.

If relying on the 0V plane to act as one of the shielding-can's walls, as is usually the case, provide multiple bonds between the shielding-can walls and the 0V plane under the segregated circuit zone. The maximum spacing between these bonds should be $\lambda/20$ (where the wavelength λ is measured in the air) at the highest frequency of concern, f_{\max} , or $15/f_{\max}$ (f_{\max} in GHz gives spacing in millimetres) for example the maximum bond spacing should be 15mm for up to 1GHz. With this maximum spacing the shielding effectiveness achieved at f_{\max} will not be very good, but at least resonances in the spaces between the bonds will be prevented. Much smaller spacings are recommended.

5.4 0V and power planes

5.4.1 General plane design issues

A well-designed 0V plane (sometimes called a 'ground plane' or 'RF Reference plane') on its own layer in a PCB is possibly the most cost-effective EMC design technique that has ever existed, or ever will. So it is always recommended to use a 0V plane wherever *possible*. Trying to reduce BOM costs by removing this plane layer is almost certainly a bad financial decision for the overall project.

It is also good practice to use well-decoupled (see 5.5) power planes too (especially where any rates of change of voltage in any circuits exceed $200\text{V}/\mu\text{s}$), but 5.5.2 shows an alternative technique that can be as good – in some types of circuits.

Planes are continuous 'solid' copper sheets on a dedicated PCB layer. They are *definitely not* 'ground fills' or 'ground meshes'. Any gaps, apertures, holes, splits, etc. in a plane reduce its effectiveness, and so should be avoided.

All 0V or power connections should bond directly to their respective planes using the shortest widest traces that can be cost-effectively achieved.

Figure 5E shows an example of a 0V plane underneath a through-hole connector, showing how the diameters of a plane's antipads (clearance holes) should be small enough to allow substantial 'webbing' of the plane between the pins. The aim is to maximise the EM characteristics of the plane at higher RF frequencies by preventing the necessary holes from joining up to create larger gaps in the plane.

The same minimisation of antipad diameter and 'plane webbing' should also be used at all via holes, and every other kind of hole in a plane, for the same reason. All modern board manufacturers should be able to get excellent yields with antipad diameters no greater than 0.36mm (14 thousandths of an inch) more than the hole diameter. Even so, some holes may need to be moved to prevent their antipads from joining together.

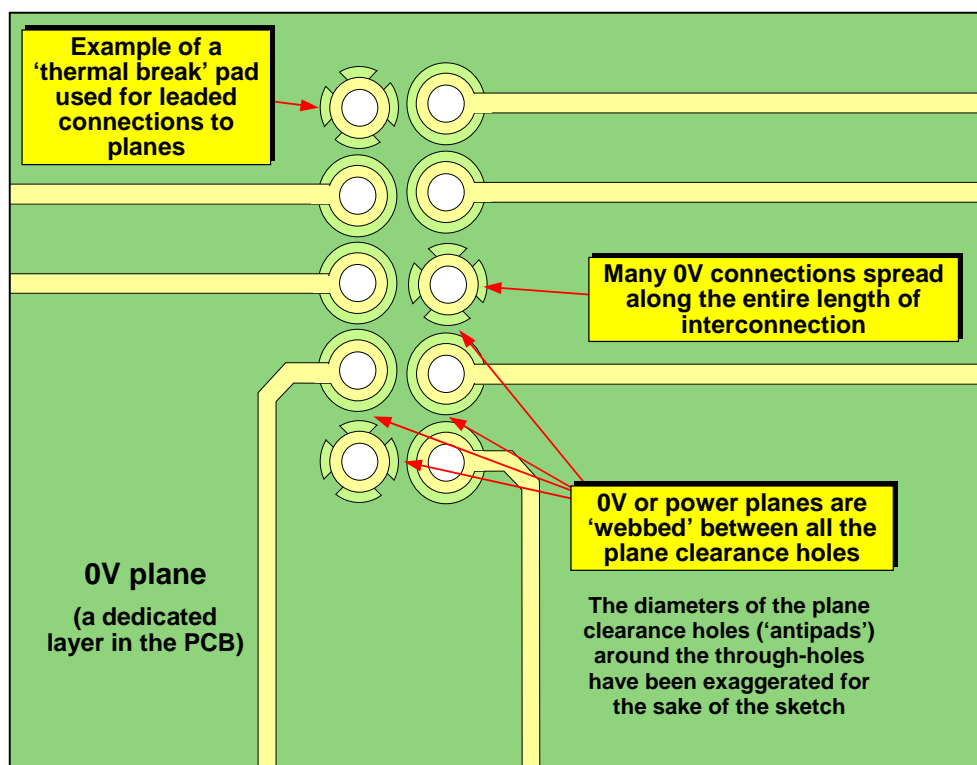


Figure 5E Example of a 0V plane under a through-hole connector

Some PCB design departments use the (by modern standards) huge via and/or antipad diameters, which were standard practices in the 1970s or 80s, to "be able to use the cheapest board manufacturers". Dealing with the resulting poor EMC performance almost certainly costs a great deal more overall than it would if better board manufacturers were used to achieve the much smaller diameters that are normal for 2007.

In Figure 5E, and in all of such figures in this article, the diameters of the plane clearance holes ('antipads') around the through-holes have been exaggerated for the sake of making the sketch clearer, they would normally be about the same diameter as the pads that 'cap' the barrels of the through-holes on the top and bottom layers of the board, and so they would be hard to see in practice just by looking at a completed board. A

well-planed board should look solid black, with just little pin-pricks of light showing through the via holes, when viewed against a light source.

Figure 5F shows an example of how not to design a 0V plane. This is a two-sided PCB with a 'ground fill' on one side, connected to the 0V. From an EMC point of view it doesn't satisfy any of the requirements of a plane, and in fact it is simply a mess of RF resonators and accidental antennas (see [7]). How best to deal with PCBs that have only one or two copper layers is dealt with in 5.4.6.




Figure 5F An example of how not to design a plane.

As mentioned in 5.2, the general rule is that a 0V plane should lie under all components, traces and power planes, and extend beyond them all around their perimeter by as far as is possible: at least 3mm, preferably 6mm or more. Using a larger 0V plane helps to reduce emissions and improve immunity.

Where possible, do not place components, or route traces very near to any plane edges, splits, holes, apertures, gaps, etc. If possible – never cross *any* plane splits, gaps, etc., with *any* traces or components (but see 5.4.5 for what to do when it is unavoidable).

It is very important to maintain the segregation of components and traces in their allotted zones, even when they share the same 0V plane. In almost all cases, the 0V plane can interconnect between any zones without needing any EM mitigation itself.

The general design rule for boards that are not very dense, is that there should be no plane gaps larger than 0.01λ at f_{max} . The value of λ that matters is the one *inside* the PCB's dielectric, which can be approximated quite well as $300/\{f_{max}\cdot\sqrt{\epsilon_r}\}$ metres, where f_{max} is in MHz and ϵ_r is the relative dielectric constant of the board material, for FR4 typically 4.2 above 1MHz. So for FR4 we can say that no plane perforations should exceed $1.5/f_{max}$. For example: for a plane in an average PCB to be reasonably effective up to 1GHz, no perforations in it should exceed 1.5mm.

1.5mm is not a problem for the vast majority of via holes and leaded semiconductors and other components, but there are leaded components that require larger hole diameters, and of course board fixing or mounting holes are larger than this. Knowing that larger holes create EMC problems, we try to keep them outside any noisy, dirty, high-speed, or especially sensitive EM Zones.

There can be other practical problems associated with trying to achieve a 'solid' 0V plane all over a board. It can be necessary to remove areas of 0V plane, when using: very high impedance circuits; very small currents; impedance-matching some types of RF devices or RF transmission lines, etc. But knowing that this can have dire consequences for EMC informs the design and allows alternative solutions to be used (e.g. using double-sided PCB shielding over the segregated area with poor 0V plane, RF-bonded to the 'solid' 0V plane all around its perimeter).

Parallel 0V planes should be bonded together at least every $\lambda/10$ at f_{max} by vias (or $15/f_{max}$, f_{max} in MHz gives the maximum spacing in metres, in GHz it gives it in mm).

5.4.2 Only use thermal break pads (thermal reliefs) when really necessary

Through-hole-plate (THP) PCB manufacturing technology, when used with leaded components (as was common in the 1980s), had a problem with the automated soldering of component leads to planes: the planes had such good thermal conductivity that they 'sucked' the heat out of the joints and dry joints were a common problem. This was solved by the use of 'thermal break' pads, sometimes called 'thermal reliefs' (or even 'wagon wheels' because of their superficial visual resemblance).

The downside is that thermal break pads perforate planes quite considerably, and so decrease their EMC benefits. So they should only be used where necessary for reliable automated soldering. They are generally not required for reflow-soldered surface-mounted components, because their break-out or pin-escape traces provide the necessary thermal relief from the via to the plane, for their soldered joints.

Unfortunately, some PCB departments apply thermal reliefs for every plane connection, and with the component density typical of modern PCBs this practice significantly increases the RF impedance of the planes, so is bad for EMC. Instead, thermal break pads should now only be used for leaded components, and only then when they are going to be automatically soldered.

5.4.3 RF-bonding planes to components, conductors and chassis

Figure 5G shows some example layouts for connecting decoupling capacitors to 0V and power planes. All other connections between components and planes should follow similar guidance. Just 1mm of trace can have an inductive impedance approaching 6Ω at 1GHz (60Ω for a 10mm trace), so it is clearly important for plane-bonding traces to be as short and wide as possible, to minimise their inductance.

There is a compromise to be made between the lengths of the plane-bonding traces and the production yield of the soldered PCB assembly. Traces that are too short, or used with too ineffective a solder resist, can result in the solder intended for the surface-mounted component being sucked into the via hole, resulting in a dry joint. Sometimes it is just a matter of specifying a better quality solder resist rather than trying to save pennies by using the cheapest resist available.




Figure 5G Examples of devices connected to planes (e.g. decoupling capacitors)

Figure 5G shows that where a device has connections to both 0V and power planes, there are some real advantages to placing their 0V and power plane via holes very close together (say 1mm or less) so that their mutual inductance and opposing directions of current cancels out some of the via holes' series inductance. Break-out and pin-escape traces should never be lengthened for this purpose. Firstly make the plane connections as short as possible, then place their 0V and power plane vias close together *without* lengthening their attached traces.

Traces crossing the edge of a 0V plane, and therefore entering or exiting an EM zone boundary, should be RF-bonded to the 0V plane near to that edge. Traces at 0V potential should be directly connected to the plane with a via. Other power and signal traces should be connected via a capacitor, the purpose of which is to provide a low-impedance return path for common-mode surface currents, but it is effectively just a capacitive filter. The value of the capacitor should not be so much as to cause a problem for the signal driver or signal quality.

As discussed in section 3.2.5 of [6], when a signal or noise source has low impedance at the frequencies concerned, using capacitive filtering on its own can sometimes increase emissions. In such cases it is usually better to use RC, LC or Tee filtering. Figure 5D shows an example of a filter circuit and layout that can be effective in a wide variety of circuits and applications. Where electronic units have well-shielded enclosures mounted directly onto metal chassis – as is common in military vehicles such as tanks and warships – π filters might be preferable to Tee, see [6].

0V planes should be RF-bonded to any metal chassis or enclosure shield, especially near high-speed devices (e.g. clock generators, clock buffers), and near any shielded (see [7]) and/or filtered (see [6]) I/O connectors – and then as frequently as possible all over the PCB's area. Ideally, the spacing between the bonds should be less than $\lambda/10$ at f_{\max} , or $30/f_{\max}$ metres, where f_{\max} is in MHz.

It is a good idea to make provision for these RF bonds, even if there is no metal chassis or shield, in case a chassis or shield has to be added later in the project, for EMC compliance or to solve actual interference problems. The chassis or shield required might even be as simple and low-cost as a sheet of metallised cardboard, the sort of thing most EMC engineers take to test labs with them to help solve customers' problems quickly. In 5.3 the term 'anti-Murphy design' was introduced for this sort of precautionary design measure.

A typical RF bond between 0V and chassis just uses a mounting pillar or screw to make the connection directly. But where there are many 0V-chassis bonds, assembly time can be reduced by using conductive gaskets or spring fingers to make automatic connections to the metal chassis or enclosure. Companies like Kitagawa, W.L Gore and others supply components intended for just that purpose.

Some designers, and some customers (e.g. automotive, rail, marine) don't like direct 0V-chassis bonds on PCBs, in case the large currents they allow to flow in the metal structures of their vehicles should decide to flow through a PCB instead, causing it to catch fire, or at least be damaged. The practice of using the chassis or other metalwork as a high-current return path, generally makes acceptable EMC performance much more costly to achieve.

Where direct 0V-chassis connections are forbidden, or where you are not sure what to do for the best, prototypes can be designed using pad patterns like that of Figure 5H, to have a range of chassis-bonding options...

- Zero-ohm links to provide direct (DC) bonds.
- Capacitors to provide RF bonds with high-voltage isolation. The lowest effective frequency depends on the capacitor's value, and the isolation voltage generally required in automotive applications is around 500Vdc, in some railway systems it is 2kV.
- Resistors to dampen RF resonances that could occur in the cavity between the 0V plane and the chassis or enclosure shield.
- Resistors in series with capacitors to damp RF resonances while maintaining high-voltage isolation.

It is important to understand that multi-point bonding is *always* required for RF, with the spacing between the bonds less than $\lambda/10$ at f_{\max} , or $30/f_{\max}$ metres, where f_{\max} is in MHz. Single-point bonding is incapable of being very effective above a few hundreds of kHz, and is generally completely useless above 30MHz.

Using the above 0V-chassis RF bonding components, we can achieve what is sometimes called 'hybrid bonding': one 0V-chassis bond is direct, whilst the others are via capacitors (or capacitors and resistors in series). This might satisfy the instrumentation and audio circuit designers who want to stick to their traditional single-point grounding practices, whilst also achieving EMC compliance for their products, or improve their EMC performance for other reasons.




Figure 5H Example flexible 0V plane – chassis bond layout

5.4.4 Don't split 0V planes any more (and what to do, if you do)

Of course, where galvanic isolation is necessary between two parts of a PCB, the 0V plane must be split between the two areas.

But never split a 0V plane just because a guideline, textbook, data sheet or application note says so. Articles, papers, guidelines, textbooks and application notes dating from before 2003 can be out of date as regards cost-effective PCB layout for EMC. An example of a textbook that contains good advice on breaks in 0V planes is [10], published in 2007. And many semiconductor manufacturers ignore EMC when they write their application notes, or else use 'traditional' practices that are, in fact, well out of date.

For example, it has been common, in the past, to split 0V planes between analogue and digital – but if the design recommendations in this article are implemented, you will generally achieve much better EMC – *and much better functional performance* (e.g. signal/noise ratio) if you use a single 0V plane over the entire PCB and all of its different EM zones.

The author learned to use unbroken 0V planes in the early 1980s, just to improve functional performance in the most demanding analogue applications using PCBs that mixed analogue and digital technologies. It was only in the early 1990s that I learned that this approach was also the best thing to do for EMC as well. The functional performance achieved for products and even large systems were well beyond what was thought possible with split analogue/digital 0V planes, and frequently amazed designers who thought splitting planes was some sort of law. The technique is very well proven (and not just by me), and [4] goes into much more detail.

These days, split 0V planes should only be used as part of a well thought-out EMC plan, which of course requires considerable EMC expertise to develop. If you are not sure whether to split 0V planes or not, prototype PCBs can provide both options, as sketched in Figure 5J...

- Split the planes between the segregated circuit zones (this is easy to do, because of the segregation discussed in 5.2)
- Place pads and 0V vias on both sides of the split so that it can be 'stitched' together *at least* every $\lambda/10$ at f_{\max} ($30/f_{\max}$ metres, where f_{\max} is in MHz) with small zero-ohm links or capacitors (an 'anti-Murphy' design technique).
- Test prototypes, using functional as well as EMC tests (the close field probe methods described in Parts 1 and 2 of [8] can be very helpful, but 'proper' emissions/immunity tests provide the best proof) to see whether split or stitched planes work best overall, and if stitched, what type of stitching is best. If you find that zero-ohm links across every stitching point is best, then removing the split altogether at the next iteration will probably improve performance even more.

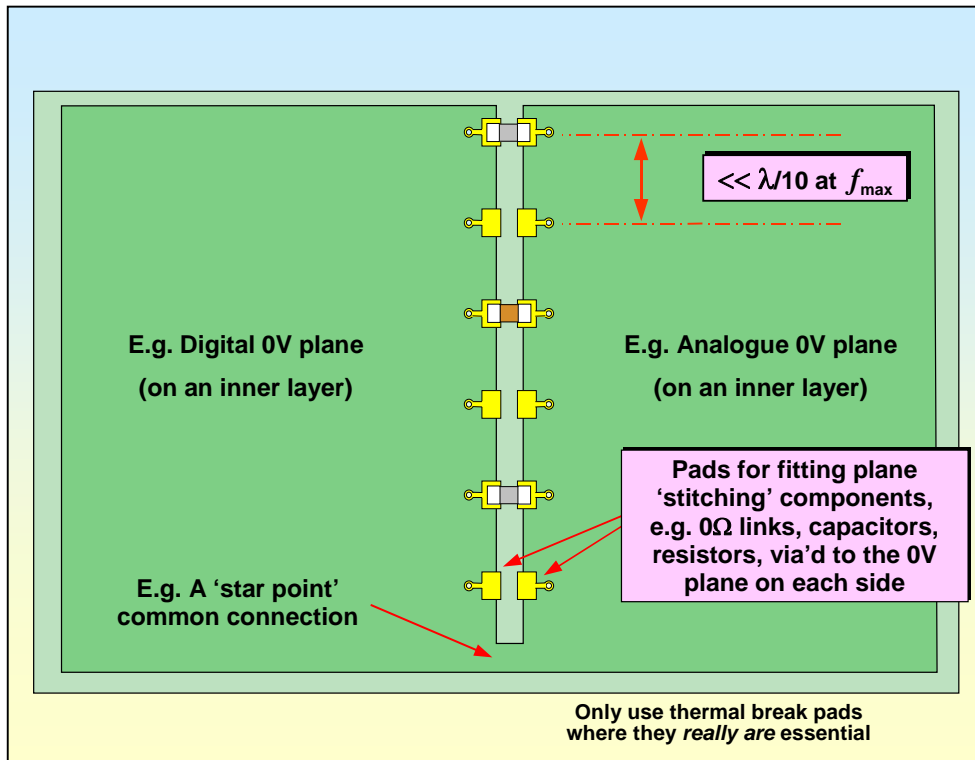


Figure 5J Example of 'stitching' across split planes

As was described in 5.3, this 'anti-Murphy' EMC design technique – used with a complete split between the planes – allows a direct connection to be made between the planes with a zero-ohm link at one of the stitching pads, with capacitors (or series resistor-capacitors) at the other stitching pads, to provide what is often called 'hybrid bonding'.

For galvanically isolated circuit zones, link the isolated 0V plane to the PCB's main 0V plane with a number of small capacitors spaced all around its perimeter. The capacitors should be rated to withstand the maximum voltage difference across the split. Where the galvanic isolation is for safety reasons, it is important to use no more than the maximum total value allowed by the relevant safety standards, and it is strongly recommended to use capacitors that have third-party safety-approvals to the relevant standards for the application. It is also recommended to check that the safety approvals are valid, and not counterfeit, by contacting the approvals body to confirm.

5.4.5 Traces routed close to plane edges, or across plane splits

Almost all EMC design can be seen as a process of controlling the return current paths so they are always in very close proximity to their send paths. So it is important indeed that no 'single-ended' (i.e. 0V referenced) signal or power traces cross any perforations or splits in their adjacent plane layers (whether they are 0V or power planes).

This is because the return currents naturally travel in the adjacent plane layers (whether 0V or power) and any perforation or split in those planes forces the return current away from the lowest-energy route preferred by the laws of nature, causing a great increase in EM fields, hence higher emissions and worse immunity.

In fact, traces should not go closer than 3mm to the edge of a plane (preferably more, see 5.4.1), because this also causes problems for the return currents.

But where a signal or power trace *has* to cross a plane split – it *must* have a return path provided in intimate proximity – *even if it means shorting-out the split at that point*. Keeping the return current path in intimate proximity to the send path is vital for EMC (and signal integrity, SI) – much more important than maintaining a split in the plane.

The best way to maintain the isolation of a plane split whilst crossing it with a power or signal trace, is to pass both the send and return currents for that power or signal through a common-mode (CM) choke that straddles the split. For single-ended power or signals, one winding of the choke connects to the plane on either side of the split.

A less effective alternative is to provide a plane 'stitching' capacitor very close to the power or signal trace, for the return current to take instead of diverting around the split. This method should be used where the planes on each side of the split crossed by the trace are at different voltages.

5.4.6 Can't afford multilayer PCBs?

Analogue and 'glue logic' boards can often achieve a good 0V plane with only two layers, but many modern boards have so many interconnections that they need at least four layers to be able to fit in a good 0V plane (see 5.4.1). There is often pressure from managers to remove dedicated 0V plane layers to reduce the BOM cost of the bare board, but this should be strongly resisted because:

- Adding a 0V layer does not add very much cost, providing you purchase from the correct manufacturer. Most corporate buyers, if asked to get a price for adding a couple of layers, will go to the suppliers they already use and be given a silly price. This is because all PCB manufacturers specialise in boards with a certain number of layers, although they often do not reveal this fact to customers. They will not give the best price for a board if it has a different number of layers from what their processes are optimised for.
In reasonable volumes, and for at least the last 17 years, the cost of adding two layers should not add more than 25% (typically 20%) to the bare-board unit cost. But the company buyer will have to do more work to find the suppliers who offer that price.
- Saving PCB layer cost by deleting a 0V plane usually adds to the overall cost of manufacturing a unit. This is because a continuous PCB 0V plane is probably the most powerful and cost-effective EMC technique of all (see 5.4.1), and if it is not used it is likely that more costly EM mitigation measures will be found to be necessary (with associated project delays) to pass compliance tests.

However, if using single layer PCBs, it is best to fill as much area as possible with 0V fills and traces, then fit 0V links that cross the gaps to create the smallest possible 0V mesh. A mesh is only effective at frequencies $\ll 30/L$ MHz (where L is the diagonal of the largest element in the mesh, in metres). Above frequencies of $50/L$ MHz, a mesh can *amplify* emissions and susceptibility problems, so it is important to get the mesh size small enough for a given f_{\max} .

If using 2-layer through-hole-plate (THP) PCBs, the best method is to move as many traces as possible onto one layer, then make the other layer as 'solid' a 0V plane as is practical. Any traces 'sneaked' into the 0V plane layer should be as short as possible. Repositioning of components is usually required to do this well. Using this method alone, the author has improved the immunity of a passive infra-red (PIR) detector from 5V/m at 900MHz (the burglar alarm system was being triggered by cellphones outside the building) to 50V/m.

The above 2-layer method is most appropriate for analogue or simple 'glue-logic' boards. Anything with a microprocessor with external memory is likely to have so many interconnections that both layers must be used for routing. (Digital processors with internal memories, and no high-data-rate buses on the PCB, are better for EMC).

These boards should route with horizontal traces on one side and vertical on the other, then '0V fill' both sides, then 'stitch' the fills together with vias or links to make the densest possible mesh. The $30/L$ guide for mesh effectiveness is the same as that given above. Unfortunately, the area where the smallest mesh is required for EMC, is closest to the microprocessor, where the trace density often makes it most difficult to get a close mesh. Where a 2-layer PCB fails an EMC test at RF by less than 10dB, it is often possible to use this 0V meshing method to achieve a pass – but it can be very time-consuming and require several iterations of the board before the optimum layout is found.

5.5 Power supply decoupling

5.5.1 General decoupling design rules

Power rail decoupling aims to keep the noisy RF currents drawn by semiconductors from exciting the product's power distribution system (PDS) as an accidental antenna (see [7]) and increasing emissions. The technique uses decoupling capacitors ('decaps') to provide very low impedances at the frequencies of concern, local to the noisy semiconductors, which encourage the noisy currents to remain local and not spread through the PDS. Similarly, decoupling maintains a low impedance to improve immunity to a range of conducted transient and continuous EM phenomena that can exist on the power rail.

Every power pin on an IC (or other type of semiconductor) should have a nearby decap to its local RF Reference (its nearby 0V plane). Figure 5K shows the general guidelines for an IC with only a 0V plane; and for an IC that has an adjacent pair of 0V and power planes in its PCB stack-up (strongly recommended, see 5.5.3).

For the IC in Figure 5K that has only a 0V plane, the decap should be placed very close to the power pin, and the power trace should pass through the decap's pad before connecting to the IC. The 0V pins of the IC and of the decap itself should connect directly to the 0V plane using short wide traces. Where 0V *and* power planes are both available, the IC to be decoupled and the decap itself should connect directly to both planes using short wide traces.

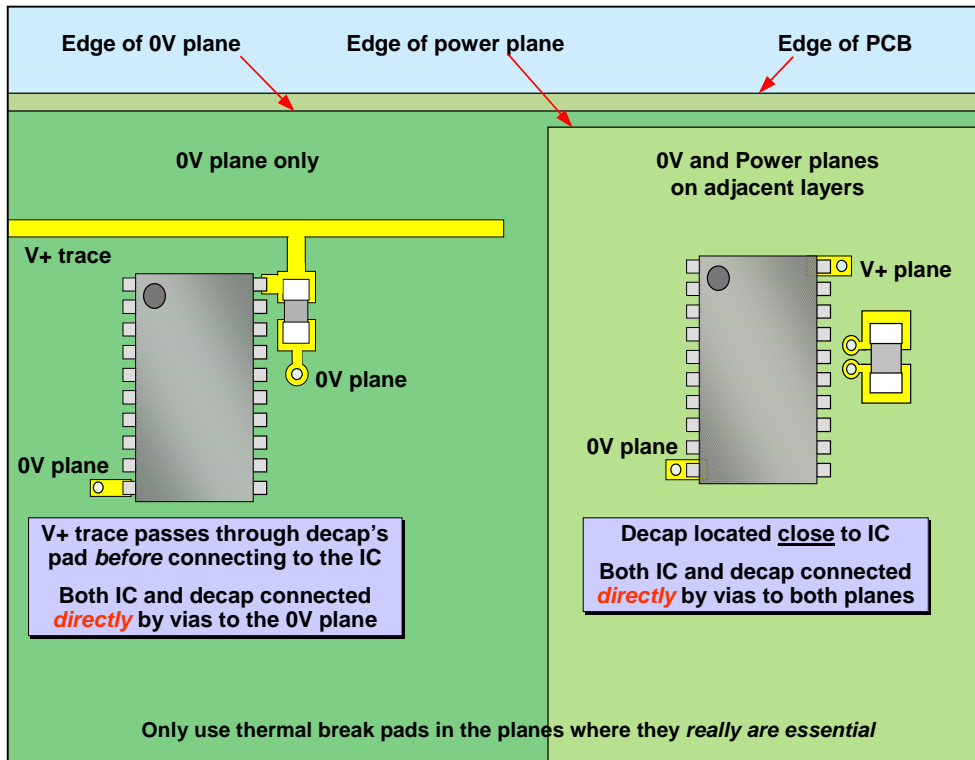


Figure 5K Examples of decoupling with planes

Figure 5L shows an example of an IC with multiple sets of 0V and power pins, each of which should have at least one local decap.

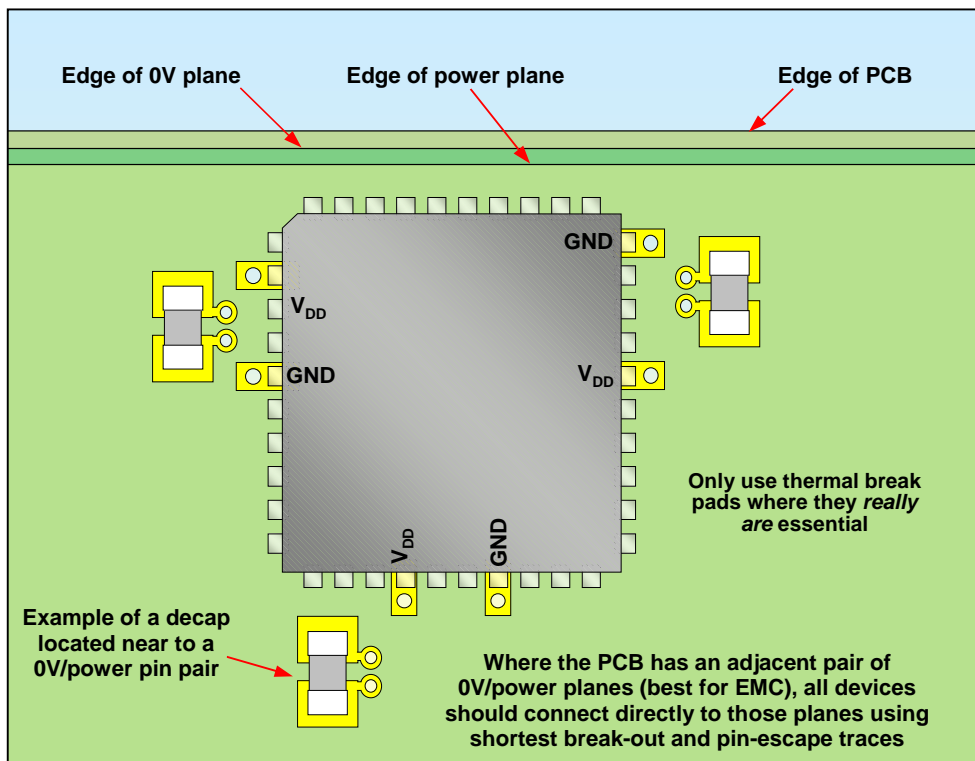


Figure 5L Another example of decoupling with planes

The RF impedance between two connections to a plane is orders of magnitude less than an equivalent length of trace, which is why we always use a plane instead of a trace where a plane is available. Some designers are in the habit of connecting IC pins to decaps using traces, and then connecting the decaps to the planes. Their intention is to 'keep the noise out of the planes' but, as described in 5.5.3, an adjacent pair of 0V and power planes is much better at decoupling noise frequencies above 300MHz than any discrete decaps possibly can – and modern digital ICs produce a great deal of power supply noise above 300MHz.

Also, the use of traces from IC to decap increases the total inductance of the decoupling, making the decaps much less effective at frequencies above 100MHz. So it is now best to use the layout techniques shown in Figure 5K and Figure 5L.

The best discrete decaps are surface mounted multilayer ceramic capacitors (SMD MLCCs). Decaps with COG and NPO dielectrics generally have the best performance at the highest frequencies, but the lower cost X7R dielectrics are often better for EMC overall because their higher series resistance provides some useful damping for the resonances that plague PDSs.

The self-resonant frequency of a decap of capacitance C , is given by $1/\{2\pi\sqrt{(L_{tot}C)}\}$, where L_{tot} is the total inductance associated with the decap (the decap itself plus its associated pads, traces and via holes), and is almost always well below 100MHz. Above this frequency the decoupling of the noise on the power supply rails is provided by the inductive impedance of the decaps. So smaller decaps (e.g. 0402, 0201) with lower profiles, used with the better layouts from Figure 5G, are best for EMC because their L_{tot} is the lowest.

5.5.2 Decoupling with ferrites

PCBs carrying mostly analogue ICs, or where the number of digital components is not very large, are sometimes designed with soft ferrite ('RF suppresser') beads in series with the power rails of their ICs, and with a 0V plane but no power plane. The ferrite beads help to restrict an IC's power supply noise currents to its nearby decap, and so provide better EMC performance than the '0V plane only' example in Figure 5K.

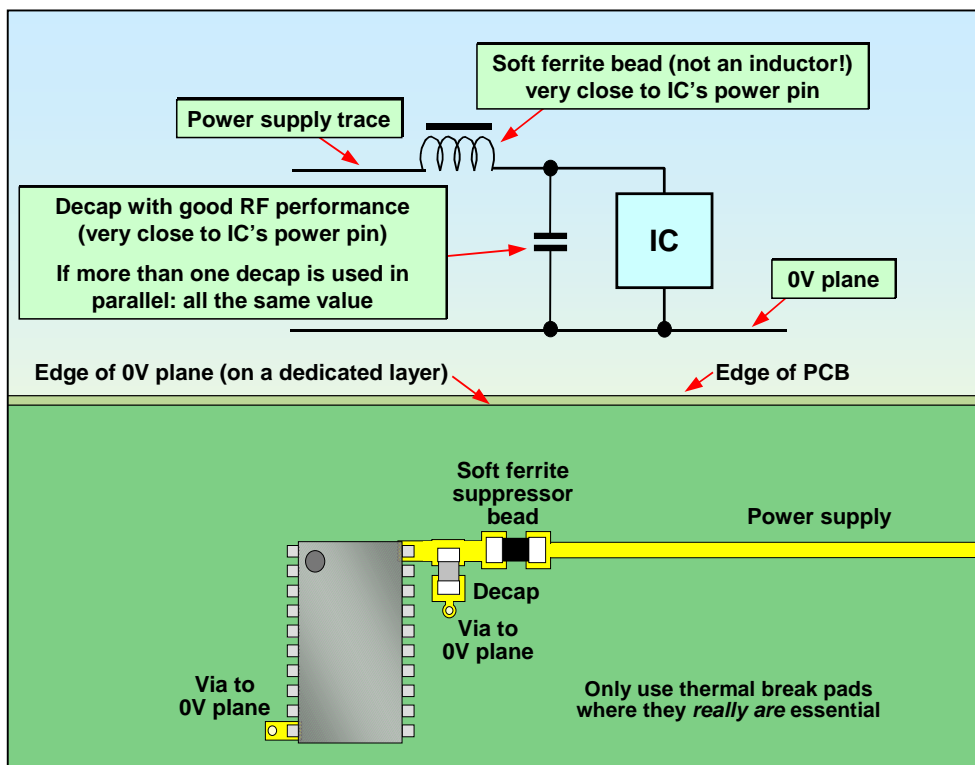


Figure 5M Example of decoupling with ferrites

For high-performance ICs, especially RF and analogue devices (e.g. ADSL drivers) the MLCCs used in this method may need to have better RF performance, which might mean using more costly types (e.g. the Murata ERB32 series). If more than one decap is used in parallel to achieve the desired total capacitance, they should all be the same value, to help avoid parallel resonances (which cause high impedances in the PDS).

Not all of the ICs on a board might need to use the ferrite beads, and some manufacturers spend longer in EMC test laboratories, replacing ferrites with zero-ohm links to reduce BOM costs by finding which ferrites are actually necessary.

5.5.3 Benefits of 0V/power plane pairs

As mentioned in 5.5.1, decoupling with discrete decaps relies on achieving a low inductance above 100MHz. But inductive impedance rises with frequency, so as frequencies exceed 300MHz decoupling schemes that rely on MLCCs (with/without ferrite beads, see 5.2.2) suffer higher impedances and hence emit more, and are also more susceptible to interference.

However, an adjacent pair of 0V and power planes in a PCB's layer stack provides an intrinsic, distributed capacitance that still behaves like a capacitor at well over 1GHz, due to its very low self-inductance. So these

days it is generally recommended to include a 0V/power plane adjacent pair inside a PCB, to help achieve good decoupling above 300MHz.

For maximum benefit from a 0V/power plane pair, connect 0V and power pins and terminals of all devices and components directly to their respective planes, as shown in Figure 5K and Figure 5L. Any traces used to connect planes to the pins or pads of devices should be very short and wide to minimise their inductance (see Figure 5G).

Achieving low PDS impedance over the whole frequency range of concern these days requires a combination of decoupling techniques, as shown diagrammatically by Figure 5N.

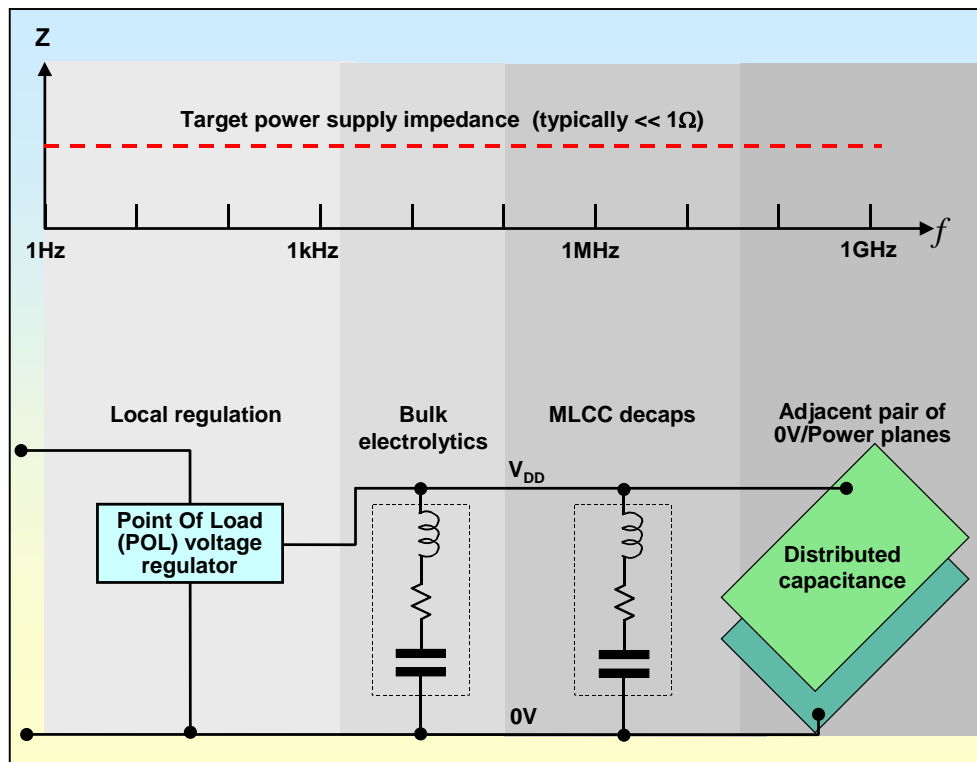


Figure 5N Achieving low PDS impedance to > 1GHz

5.5.4 Dealing with PDS resonances

When not using series ferrites in the PDS as in Figure 5M, decaps appear in parallel between the power and 0V – giving rise to parallel resonances, which cause high impedances that can cause EMC problems if any EM disturbances occur within their narrow ranges of frequencies. Also, 0V/power plane adjacent pairs can suffer cavity resonances due to their dimensions, which also cause high impedances.

Like all of the topics discussed in this article, [4] goes into a lot more detail on this issue, but it is enough for this basic guide to recommend the following approach to dealing with parallel resonances:

- If there are less than 10 decaps in total on the board – use same value (e.g. 10nF) for each
- If there are more than 10 decaps in total, use a range of values, e.g. 1, 2.2, 4.7, 10, 22nF, etc.
- Use the largest value capacitors available in the chosen package style (e.g. 0603).
- As well as placing decaps near to every power pin of every IC, also add more decaps all over the power plane, or along the power trace, so that they are no further apart from each other than $\lambda/10$ at f_{max} taking the board's dielectric constant into account ($15/f_{max}$ metres for FR4, where f_{max} is in MHz).
- Better still, use one of the several PDS simulators that have recently become available, to help decide on the above issues. If they are SI simulators rather than EMC, set 5 to 10 times lower limits for the PDS noise levels than are required just for the ICs to function correctly.
- Finally, verify the layout with 'proper' EMC tests.

The aim of these techniques is generally to ensure that whenever an IC's noisy power current cannot flow in its normal path, due to a parallel or cavity resonance at that frequency, there will be another path nearby that still has low impedance.

If resonances are still a problem, add series combinations of 4.7Ω resistors and 10nF MLCCs between the 0V and power to dampen down the peaks of the impedance resonances, at the locations on the board that provide

the most benefit. The necessary experimentation can be begun with a spectrum analyser and close-field probe (parts 1 and 2 of [8]) on an ordinary development bench.

5.6 Matched transmission line techniques

5.6.1 When to use matched transmission lines

The digital industry guidance for good signal integrity (SI) – to achieve low-enough overshoots and ringing on the signals for reliable functionality – is to use a matched transmission line when the propagation time from the source to the load exceeds $t_r/2$, where t_r is the shortest rise-time carried by the trace (use fall-time instead, if it is less). Some digital designers recommend $t_r/3$ instead.

If working with a spectrum analyser rather than an oscilloscope, the $t_r/2$ ‘rule’ is equivalent to when the trace length exceeds $\lambda/7$ at f_{max} , taking the board’s dielectric constant ϵ_r into account in the calculation of λ (approximately $24/f_{max}$ metres for an FR4 board, when f_{max} is in MHz).

The propagation time for a conductor in air is 3.3ps/mm, and for a conductor in a PCB dielectric it is $3.3\sqrt{\epsilon_r}$ ps/mm. The ϵ_r of FR4 is 4.2 (above 1MHz) so for an FR4 board we can assume a propagation time of about 6.6ps/mm (approximately 2ns per foot). So, for example, signals with 2ns risetimes would need to use transmission lines for trace lengths exceeding 150mm (about 6 inches), in an FR4 board.

But actual rise/fall-times are always much shorter than the data sheet specifications for an IC, and they get smaller each time the device undergoes a die shrink (which can be every couple of years, even for a mature part). **And** capacitive loading due to devices decreases the velocity of propagation and so increases the propagation time of a trace. So a 2ns (specified by the data sheet) IC driving a heavily loaded data bus might need to use matched transmission line techniques for traces longer than 40mm.

But the above guides are just for SI, not for EMC. To reduce cost-of-manufacture by improving EMC at PCB levels – reducing the need for costly and restrictive filtering and shielding – it will help to use matched transmission line techniques when traces are one-quarter of the lengths calculated above (preferably even less). In other words: use matched transmission lines when the source to load propagation time exceeds $t_r/8$, using the real rise-time value for t_r , not the data sheet value.

This ‘rule’ is equivalent to when the trace length exceeds $\lambda/28$ at f_{max} , taking the board’s ϵ_r into account (approximately $6/f_{max}$ metres for an FR4 board, when f_{max} is in MHz. f_{max} in GHz gives the answer in mm).

Don’t forget that capacitive loading by devices connected along a trace increase its propagation time (appropriate calculations are given in 5.6.2).

RF and microwave designers work with devices that must be connected to transmission lines with specified values of Z_0 , for those devices to function correctly.

Many designers apply transmission line techniques to the signals that they already know need to use them. But for good EMC all traces should be analysed in terms of rise-times (or high-frequency content). It can happen that even traces to test points might need to be treated as transmission lines for EMC purposes. The analysis should consider all the wanted signals and unwanted noises on the traces, equally. Where transmission line techniques turn out to be required for signals that do not need to have very high frequencies, or very small rise/fall times, an alternative is to low-pass filter their sources, using an RC or LC filter as discussed in [6].

More than 30 types of PCB transmission lines can easily be created for ‘single-ended’ power and signals, by controlling the geometry and stack-up of the traces and planes, and some of the more common types are shown in Figure 5P.

Figure 5Q focuses on the ‘surface microstrip’ type of transmission line, and gives the simplest equation that can be used for calculating its characteristic impedance, Z_0 . Microstrips are quite complicated structures, because a fraction of the wave energy that is the propagating signal travels in the air at about twice the velocity of the remainder that travels under the trace, inside the dielectric. So the simple equation is only valid over a range of trace geometries, which nevertheless meets most normal requirements.




Figure 5P Some types of PCB transmission lines for single-ended signals or power




Figure 5Q Example of a 'surface microstrip'

Figure 5R focuses on the example of a 'symmetrical stripline', and gives its simplest equation for calculating Z_0 . Striplines are traces that lie between two planes, and the planes do not have to be at the same potential. As described in 5.4 and 5.5, parallel 0V planes should be bonded together at least every $\lambda/10$ at f_{\max} by vias, and power planes should have decaps every $\lambda/10$ at f_{\max} to the main (most unbroken) 0V plane, and this is also important for striplines. λ should take the board's ϵ_r into account, so these spacings are equivalent to approximately $15/f_{\max}$ metres for an FR4 board, when f_{\max} is in MHz (e.g. 30mm for an f_{\max} of 500MHz).

Striplines have better EMC characteristics than microstrip.




Figure 5R Example of a 'symmetrical stripline'

Because there are so many different types of transmission line that can be designed on a PCB, and because they all have different equations, it is now much more cost-effective to use a computer simulator PCB transmission-line Z_0 'solver', some of which are available free on the Internet [11]. Of course, you always get what you pay for, and the better computer simulators (e.g. those from Polar Instruments) can take into account all of the PCB manufacturing issues that can influence the final Z_0 (e.g. areas of the PCB where the ϵ_r varies due to the effect of manufacturing processes on the dielectric materials in the stack-up).

Relying instead on formulae from textbooks and standards, it would be difficult to get through a dozen different stack-ups and geometries in a day (most designers seem to give up after the first six, and settle for the best of the few they have calculated).

With a fully-featured transmission line solver it is a simple matter to click on the style of transmission line to be calculated, fill in its spreadsheet with the few dimensions needed, and have the result calculated within 1 second. In this way it is possible to evaluate *many* dozens of different PCB stack-ups and trace geometries in an hour. Productivity is higher, and design quality, SI and EMC are much better.

In real PCBs, transmission lines might change their geometry along their length, and/or change layers in the stack-up. It is important to change the trace width (if required) to maintain the same value of Z_0 for each different segment of the line.

A great deal of useful information on transmission lines is given in IPC-2141 and IPC-D-317A (www.ipc.org), and also in IEC 61188-1-2 (www.iec.ch), all available to buy on-line. [12] is an excellent reference for the basics, and it includes a number of other useful references.

5.6.2 Correcting for load capacitance

The capacitances of the active devices (e.g. the input capacitances of gates, typically 3-5pF) connected to a trace cause the propagation velocity to reduce, and they also reduce the Z_0 of the trace. This is not calculated by transmission line solvers, which only take the PCB itself into account. The following formula can be used to estimate the effect of device loading:

$$Z_0 \text{ (loaded)} = \frac{Z_0}{\sqrt{(1 + C_d/C_0)}} \quad v \text{ (loaded)} = \frac{v_0}{\sqrt{(1 + C_d/C_0)}}$$

where: C_d is the load capacitance per unit length
 C_0 is the intrinsic (bare-board) line capacitance per unit length
 v_0 is the intrinsic (bare-board) line velocity

For example, if there were 6 devices, each with an input capacitance of 5pF, connected along a 200mm long trace being designed as a matched transmission line, then the value for C_d would be 150pF/metre, or 0.15pF/mm.

Except when using very thin, narrow traces, it is usually possible to reduce the width of the transmission line trace either side of the point of connection of a device, to compensate for the added capacitance at that point, so that the Z_0 of the trace can be maintained. The trace should be subdivided into sections having propagation times no longer than $t_r/10$ (for good SI) or $t_r/40$ (for good EMC) and the trace geometry adjusted to maintain the same Z_0 in each section.

IEC 61188-1-2:1998 gives formulae for C_0 and v_0 , for some common types of transmission line – but it is now much better, quicker, and more accurate to extract these parameters from actual PCB layouts by using field solvers running on personal computers.

5.6.3 Choosing the dielectric materials for the stack-up

FR4 has a nominal ϵ_r of 4.7 at the usual measurement frequency (100kHz). But different batches can vary between 4.0 and 5.5, and some poor quality board materials can apparently vary over this range over the length or width of an individual panel. ϵ_r reduces as f increases, and it is nominally 4.2 at frequencies above 1MHz (but varying in practice over the range 3.6 to 4.9).

Grades of FR4 with a more accurately controlled ϵ_r are readily available for the construction of PCBs with matched transmission lines (known as ‘controlled impedance’ boards), if you know to ask for them. The author’s favourite prototype board manufacturer buys a material specified as having an ϵ_r of 4.7 ± 0.1 , at 100kHz, and the extra cost is so low that they don’t bother to increase the price of the prototypes.

There are many other PCB dielectrics than FR4 that may be more suitable for a particular application, for example high voltage, high temperature, high vibration, etc. There are also specialist materials designed for boards carrying microwave signals, although as the volume manufacturers of cellphones and personal computers continue to raise their data rates, techniques are being developed to use the (low cost) FR4 substrate at ever higher frequencies.

Board manufacturers purchase dielectrics in sheets, either with no copper plating, plating on one side or both, and in various thicknesses, to stack up to make a PCB. These sheet materials are only available in a limited range of thicknesses, so it is important to discuss with the chosen board manufacturer what sheet thicknesses are available, then design the trace widths and geometry accordingly.

5.6.4 Terminating transmission lines in a *matching resistance*

To function correctly as a matched transmission line, a trace that is designed as a transmission line *must* be terminated in a resistance equal to its Z_0 , and for good EMC the termination resistors must maintain their overall impedance right up to f_{max} . This usually means using small SMD ‘chip’ resistors (not MELF types), surface-mounted resistor arrays, or integrated transmission-line termination devices, all with direct connections to the relevant planes as shown in Figure 5G.

There are a number of different transmission-line matching techniques available, all with different engineering compromises. These are described in section 2.7.3 of [7] and will not be repeated here.

5.6.5 Differential matched transmission lines

Differential signalling (also known as ‘symmetrical’ or ‘balanced’ signalling), uses two signal conductors driven in antiphase, as shown in Figure 5S. It is increasingly used on PCBs for microprocessor clocks and serial datacommunications, to improve signal integrity, and it can also – if carefully designed – provide better EMC (lower emissions, higher immunity).

Differential signalling requires the creation of differential matched transmission lines on the PCB, which are pairs of traces routed close together along their entire route, maintaining their specified characteristic impedance values for each segment along their entire length. Differential transmission lines have *three* types of Z_0 ...

- The single-ended Z_0 , when each trace is driven independently
- The CM Z_0 , when both traces are driven with a common signal
- The differential-mode (DM) Z_0 , when the two traces are driven with antiphase signals

More than 30 types of differential PCB transmission lines can easily be created for power or signals, by controlling the geometry and stack-up of the traces and planes, as shown in Figure 5T.

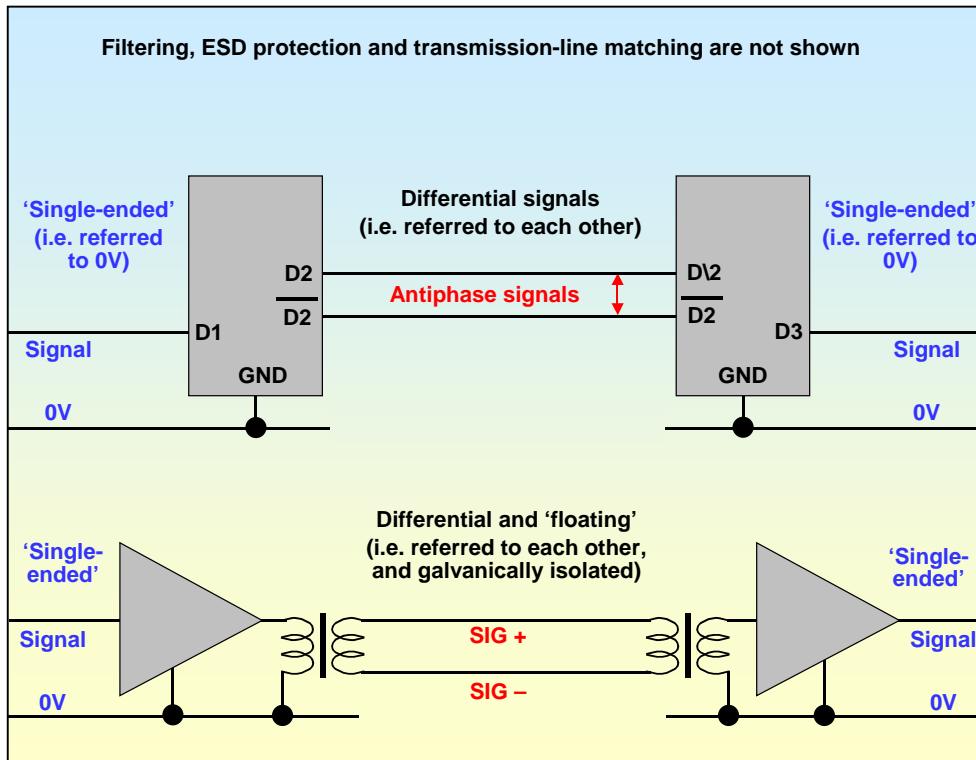


Figure 5S Examples of differential signalling

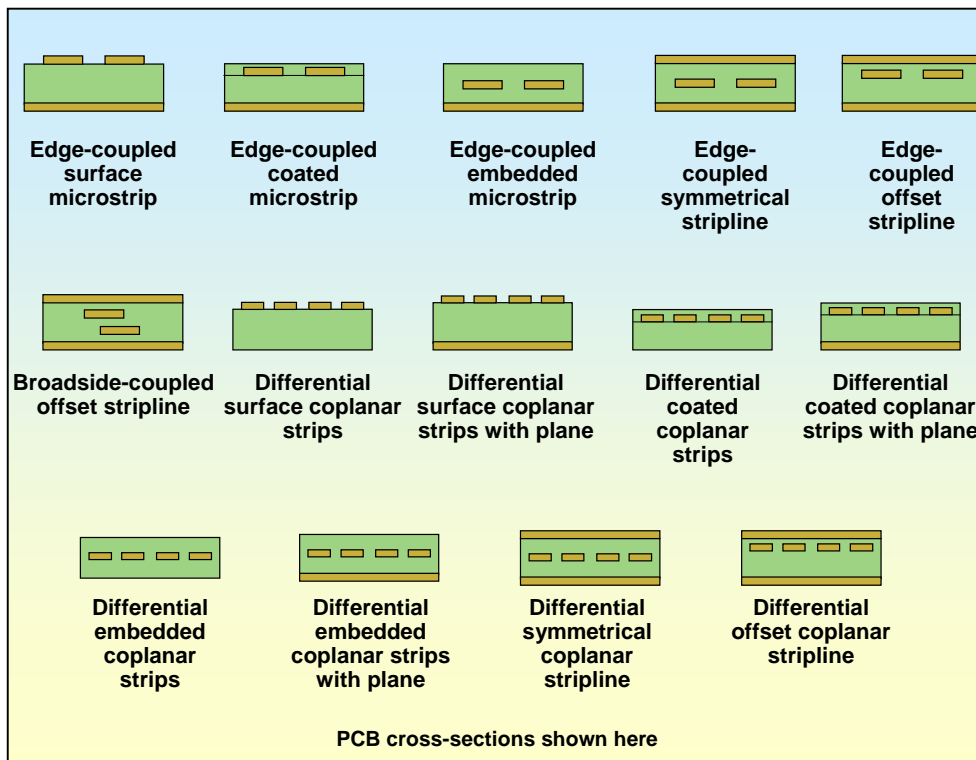


Figure 5T Some types of differential PCB transmission lines

Differential transmission lines suffer from imbalances that increase their emissions and worsen their immunity. Cables deal with similar imbalances by twisting their conductors, but of course this is almost never practical for PCB traces. Imbalances can be caused by:

- Routing too close to metalwork, edges of planes, other traces, etc., because one trace has more stray C or stray mutual inductance than the other. The use of stripline and/or coplanar techniques can help, by 'shielding' the traces above and below, and/or on both sides.
- Not maintaining the trace separation, widths, etc. so that all three types of Z_0 are maintained at the same time.

- Connectors. So choose pins that are symmetrical with respect to the shell and other pins, and so have identical stray capacitances.

It is common when terminating differential lines with matching resistors, simply to terminate the differential signals. But for good EMC all three types of Z_0 need to be correctly terminated, as shown in Figure 5U for an example of parallel (shunt) termination. Real differential signals can contain significant amounts of single trace, CM and DM signals, and unless all three modes are matched properly, emissions will be higher and immunity poorer than they need be.

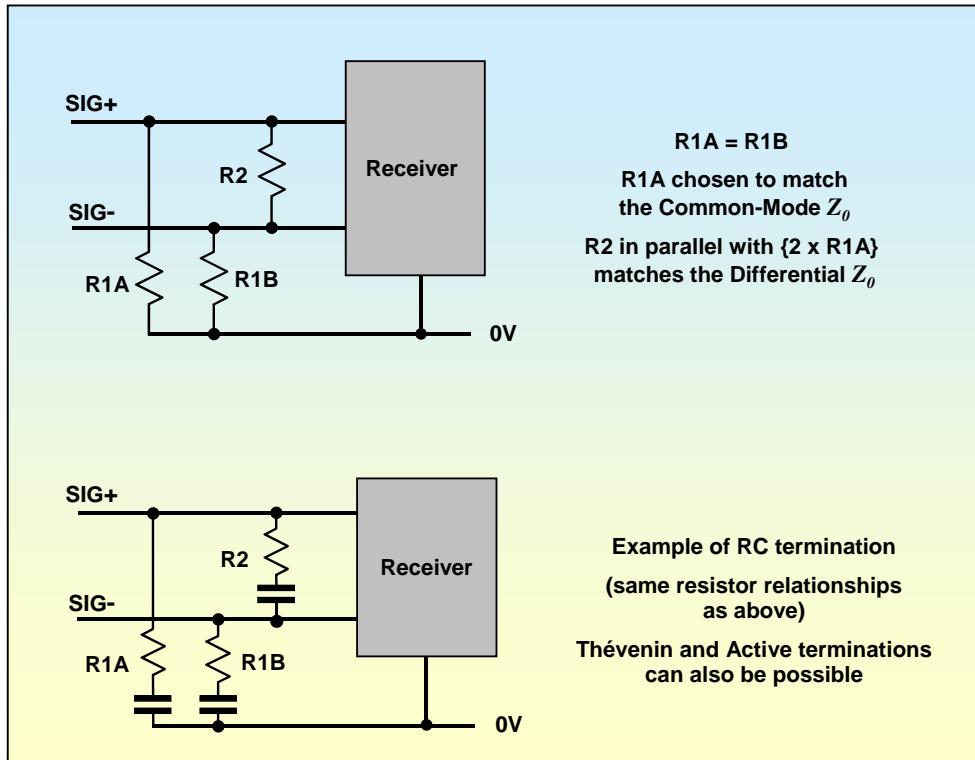


Figure 5U Examples of differential termination

5.6.6 Transmission line routing

These routing guidelines are also relevant for any traces carrying high-speed or RF signals or noises, and assume the board has at least one 0V/power plane pair in its layer stack.

- Firstly, route the 0V, power and decoupling. Because the correct use of 0V and power planes removes the need for all but the shortest traces, this will use up very little routing space and block off very little of the board area.
- Secondly, route the traces with the fastest digital edges or highest frequencies, using one PCB layer only, and keeping them short. Ideally, this would be a layer adjacent to a 0V plane that was part of a 0V/power plane adjacent pair. These traces typically include: microprocessor clocks; write strobes on SRAMs and FIFOs; output enables and chip enables and high-speed serial data buses. It can be good to route any very sensitive signal traces in a similar manner.
- Thirdly, route any parallel data busses.
- Lastly, fit all the other traces in somehow.

Design the traces and PCB layer stack to maintain the desired Z_0 over each segment of the entire length of each trace. The trace should be subdivided into sections with propagation delays no longer than $t_r/10$ (for good SI only) or $t_r/40$ (for good EMC), and the trace geometry adjusted to maintain the same Z_0 in each section.

If a transmission line has to change layers and/or crosses a plane split, so that its adjacent plane(s) are now different, a path must be provided for its return current, very close indeed to the point where the trace changes planes, as described in 5.4.5.

5.6.7 Stubs and branches

Ideally, transmission lines should have no joints along their lengths, because of the impedance discontinuities these cause. This is only possible for traces that are routed point-to-point, as shown in Figure 5V. Note that when two loads are to be driven point-to-point from one driver, they should fan out from the driver pins, not from

the series line matching resistors. Each load must have its own line matching resistor, and each driver output must be powerful enough to drive the parallel impedance resulting from all the lines connected to it.




Figure 5V Example of 'star routing' using series line terminations

Decades ago, arrays of memory ICs used to be connected to data buses using a 'grid' or 'mesh' routing as shown in Figure 5W. But this is not acceptable these days because the rise-times of the ICs are so much less that serious signal degradation would occur, causing problems for SI, and much worse problems for EMC. Instead, we use 'daisy chain' routing as shown in Figure 5W, where the transmission line is routed to each IC in turn.

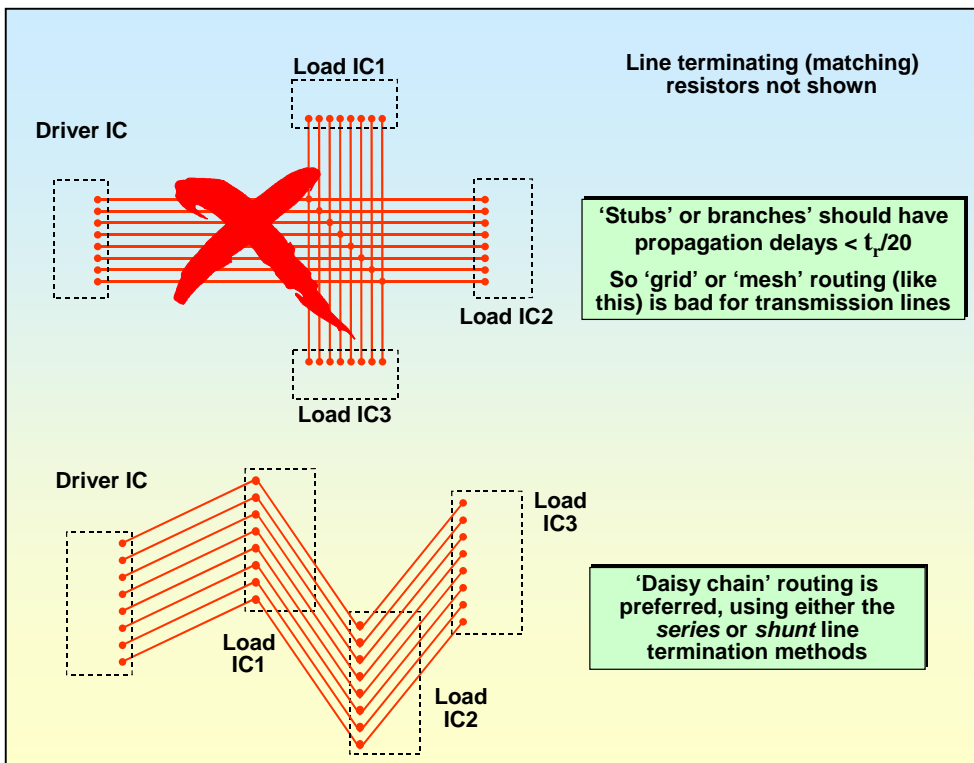


Figure 5W Example of 'daisy chain' routing

Of course, the vias, pins, lead frames and bond wires associated with the daisy-chained ICs act as short stubs or branches along the line, but as long as their source-to-load propagation delay is less than $t_r/20$, they can be treated as 'lumped' capacitances as described in 5.6.2. Longer stubs or branches cause serious problems for SI, and of course for EMC.

So far, this section has concerned itself with products consisting of a single PCB. But it is common for products to be made of plugged-together modules, using backplane boards, or daughterboards plugged into a motherboard. Figure 5X sketches the issues when a transmission line is used to interconnect a number of modules or daughterboards.

The lengths of the stubs or branches created by the modules or daughterboards limits the rise-time of the system and hence the maximum data rate. As mentioned above they need to have propagation times (sometimes called 'flight times') of less than $t_r/20$, and this is made more difficult for the example in Figure 5X because of the lengths of the pins in the module or daughterboard connectors. The backplane drivers and receivers are placed very close to the connectors to limit the overall stub lengths so that faster rise-times and hence faster data can be used.

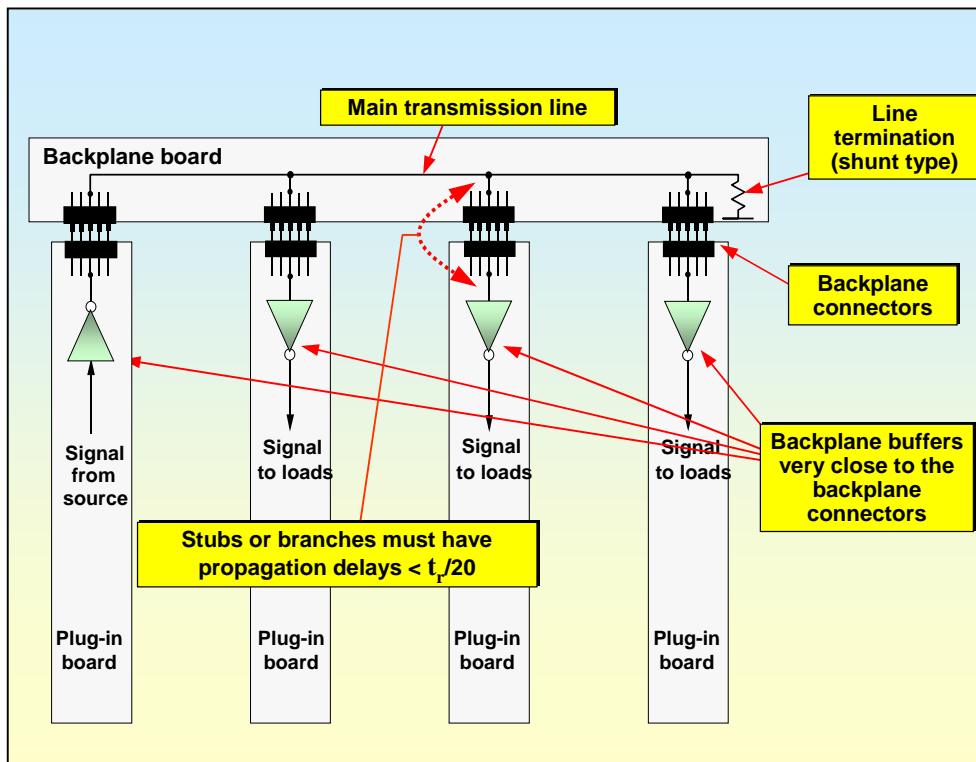


Figure 5X Example of transmission line techniques with plug-in modules

Where a stub or branch in a transmission line has a propagation delay longer than $t_r/20$, a buffer should be added at the point where it connects, so that instead of a long stub or branch it is instead a new transmission line. This technique could be applied to the example in Figure 5X, by fitting buffers on the backplane to drive the module connectors with new transmission lines.

For good EMC at board level, the "propagation delay no longer than $t_r/20$ " guide above should be replaced by a "propagation delay no longer than $t_r/80$ ".

5.7 Layer stacking

Due to successive die shrinks over the years it is now not uncommon for fairly simple HCMOS 'glue logic' boards clocking at 40MHz (say) to fail emissions tests at over 700MHz. As described in 5.5.3, an adjacent 0V/power plane pair is required in a board's the stack-up to improve the emissions and immunity of its PDS above 300MHz – so the use of such plane pairs is assumed in this section.

With the above assumption, a 4-layer board's stack-up would be as shown in Figure 5Y.

A 4-layer board with equally-spaced layers and a symmetrical stack-up, to help prevent board warp during automated soldering is not the best for EMC, because its 0V/power plane pair are not very close together, and not very close to the component mounting layers. Closer plane-pair spacings and closer spacings to the outer layers in the stack-up all help improve decoupling above 300MHz.

There are many possible permutations of unequal and/or unsymmetrical layer stacking in 4-layer boards. And there are also many possible permutations of layers for 6-layer boards. But only a board with 8 layers (or more) can satisfy all of the good practice EMC requirements and have a symmetrical stack-up to prevent board warp, and Figure 5Y gives an example.

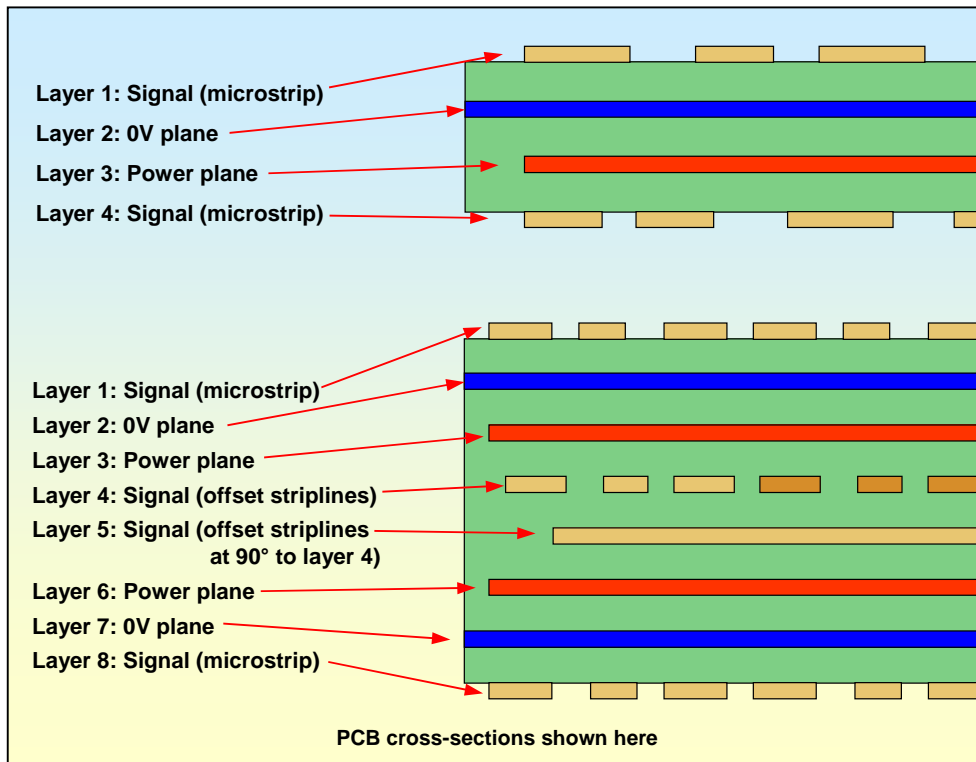


Figure 5Y Examples of 4 and 8 layer stack-ups

Closer layer spacing is better. Spacings of 0.15mm (6 thousands of an inch) or less between traces and planes in a layer stack, and between 0V/power plane pairs, can significantly improve decoupling, reduce emissions and improve immunity (and improve SI).

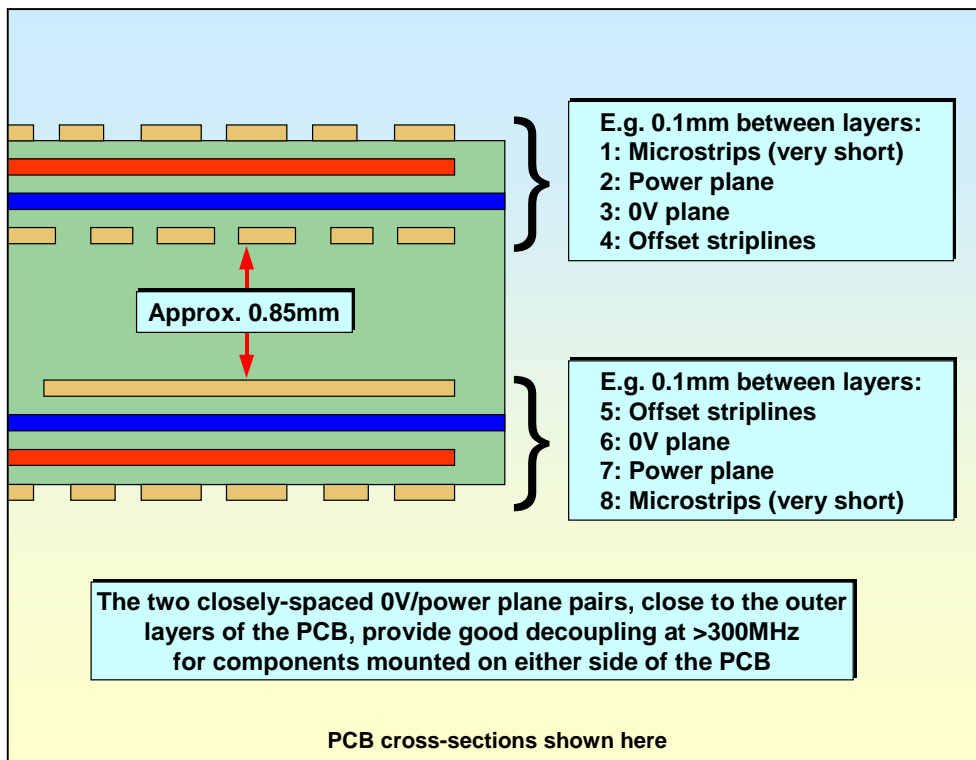


Figure 5Z Example of a close-layer-spacing but symmetrical 8 layer stack-up

So avoiding the traditional equally-spaced layer stacking can be very good for EMC, but we still want to have a symmetrical stack-up to help prevent board warp during automated soldering. Figure 5Z sketches an example of an 8-layer board with a stack-up designed in this way. Its two closely-spaced 0V/power plane pairs, close to the outer layers of the PCB, provide excellent high-frequency decoupling for components mounted on either side of the PCB.

The stack-up in Figure 5Z has proven to be very effective at solving difficult EMC problems at lowest cost in recent years. It is recommended for all new designs that need to mount active digital or high frequency analogue devices on both sides of the board.

5.8 References

- [1] Keith Armstrong, “*Design Techniques for EMC*”, UK EMC Journal, a 6-part series published bi-monthly over the period February – December 1999. An improved version of this original series is available from the “Publications & Downloads” page at www.cherryclough.com.
- [2] “*EMC for Functional Safety*”, The IET, 2008, a 180-page practical guide suitable for all industries and applications, available for free download from www.theiet.org/factfiles/emc/index.cfm or as a colour printed book for £27 from: <http://www.emcademy.org/books.asp>
- [3] Keith Armstrong, “*Design Techniques for EMC, Part 0 – Introduction, and Part 1 – Circuit Design and Choice of Components*”, The EMC Journal, January 2006 pp 29-41, plus March 2006 pp 30-37, available from <http://www.compliance-club.com>.
- [4] Keith Armstrong, “*EMC for Printed Circuit Boards – Basic and Advanced Design and Layout Techniques*”, February 2007, www.emcademy.org/books.asp, £47 plus p&p, perfect bound: ISBN 978-0-9555118-1-3, spiral bound (lays flat): ISBN 978-0-9555118-0-6
- [5] Keith Armstrong, “*Design Techniques for EMC, Part 4 – Shielding*”, The EMC Journal, Issues 69-71, March, May and July 2007, www.compliance-club.com.
- [6] Keith Armstrong, “*Design Techniques for EMC, Part 3 – Filtering and Suppressing Transients*”, The EMC Journal, September and November 2006, and January 2007, from the EMC Journal archives at www.compliance-club.com.
- [7] Keith Armstrong, “*Design Techniques for EMC, Part 2 – Cables and Connectors*”, The EMC Journal, Issues 64 and 65, May and July 2006, from the EMC Journal archives at www.compliance-club.com.
- [8] Tim Williams and Keith Armstrong, “*EMC Testing*”, a series in seven parts published in the EMC & Compliance Journal 2001-2, available from the ‘Publications & Downloads’ page at www.cherryclough.com.
- [9] John R Barnes, “*Robust Electronic Design Reference Book, Volume 1*”, Kluwer Academic Publishers, 2004, ISBN: 1-4020-7737-8.
- [10] Tim Williams, “*EMC for Product Designers, 4th Edition*”, Newnes, December 2006, ISBN: 0-750-68170-5.
- [11] Some free Transmission Line calculators:
 - a) www.emclab.umr.edu, click on ‘PCB Trace Impedance Calculator’;
 - b) www.amanogawa.com/index.html, click on ‘Transmission Lines’, then ‘Java Applets’;
 - c) Microwave Office TX-Line Transmission Line Calculator:
http://web.appwave.com/Products/Microwave_Office/Feature_Guide.php?bullet_id=9.
- [12] Freescale Application Note AN1051: “*Transmission Line Effects in PCB Applications*”, http://www.freescale.com/files/microcontrollers/doc/app_note/AN1051.pdf, or visit <http://www.freescale.com> and search its technical support documentation for AN1051.

5.9 Acknowledgements

The colourful graphics in Figure 5A were created by John Symonds, and many more of John’s artwork appears on the covers of the very useful EMC Guides available free at www.reo.co.uk/guides.

I am very grateful to the following people for suggesting a number of corrections, modifications and additions to the first series published in 1999 [1]: Feng Chen, Kevin Ellis, Neil Helsby, Alan Keenan, Mike Langrish, Tom Liszka, Tom Sato, and John Woodgate.