# emc STANDARDS

Another EMC resource
from EMC Standards

# Part 8, Other miscellaneous issues

*Helping you solve your EMC problems*

# Advanced PCB Design and Layout for EMC Part 8 - A number of miscellaneous final issues

**By Eur Ing Keith Armstrong C.Eng MIEE MIEEE, Cherry Clough Consultants**

This is the last in a series of eight articles on good-practice design techniques for electromagnetic compatibility (EMC) for printed circuit board (PCB) design and layout. This series is intended for the designers of any electronic circuits that are to be constructed on PCBs, and of course for the PCB designers themselves. All applications areas are covered, from household appliances; commercial, medical and industrial equipment; through automotive, rail and marine to aerospace and military.

These PCB techniques are helpful when it is desired to.

- Save cost by reducing (or eliminating) enclosure-level shielding
- Reduce time-to-market and compliance costs by reducing the number of design iterations
- Improve the range of co-located wireless datacomms (GSM, DECT, Bluetooth, IEEE 802.11, etc.)
- Use very high-speed devices, or high power digital signal processing (DSP)
- Use the latest IC technologies (130nm or 90nm chip processes, 'chip scale' packages, etc.)

The topics to be covered in this series are:

1. Saving time and cost overall
2. Segregation and interface suppression
3. PCB-chassis bonding
4. Reference planes for 0V and power
5. Decoupling, including buried capacitance technology
6. Transmission lines
7. Routing and layer stacking, including microvia technology
8. **A number of miscellaneous final issues**

**This is the final part of this series, and I hope you have enjoyed reading it, or at least found some things in it that were interesting or useful.**

**A previous series by the same author in the EMC Compliance Journal in 1999 "Design Techniques for EMC" [1] included a section on PCB design and layout ("Part 5 - PCB Design and Layout", October 1999, pages 5 - 17), but only set out to cover the most basic PCB techniques for EMC - the ones that all PCBs should follow no matter how simple their circuits. That series is posted on the web and the web versions have been substantially improved over the intervening years [2]. Other articles and publications by this author (e.g. [3] [4] and Volume 3 of [5]) have also addressed basic PCB techniques for EMC.**

**Like the above articles, this series does not spend much time analysing why these techniques work, it focuses on describing their practical applications and when they are appropriate. But these techniques are well-proven in practice by numerous designers world-wide, and the reasons why they work are understood by academics, so they can be used with confidence. There are few techniques described in this series that are relatively unproven, and this will be**

mentioned where appropriate.

If you have missed any parts of this series, you can download them from the archive at http://www.compliance-club.com/keith_armstrong.asp or http://www.compliance-club.com/KeithArmstrongPortfolio. But please note that the most recent two issues of the EMC Compliance Journal are posted in full at http://www.compliance-club.com and articles are only placed in the archives when they are not in one of these - so if the part you want is the one previous to this, it might not yet have been transferred to the Journal's archives.
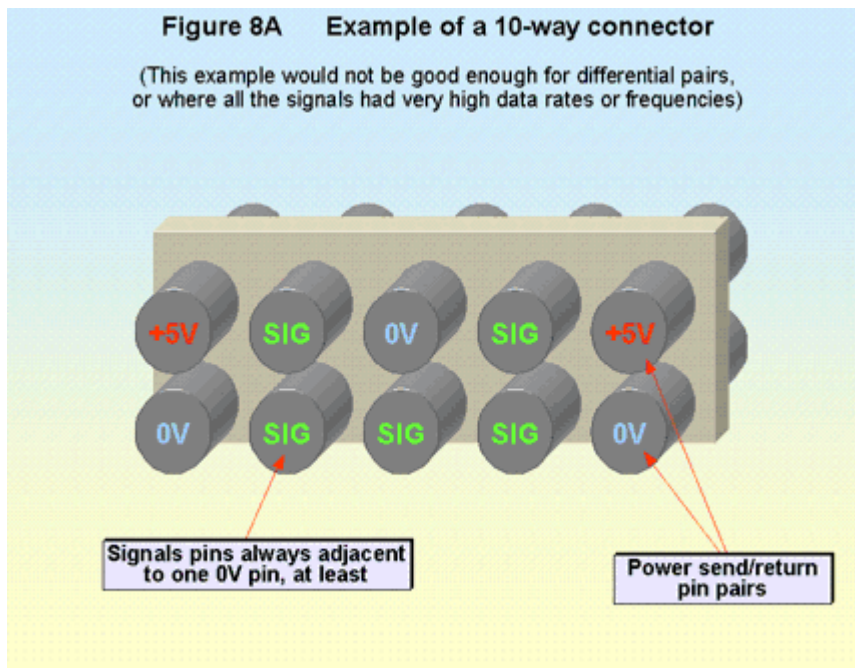
## Table of Contents, for this Part of the series

## 1 Power supply connections to PCBs

All PCB connectors carrying power and 0V should use adjacent pins for their power and 0V connections. If the connector is long, there should be a number of power/0V pin pairs spread along its whole length. If the connector is also wide there should be a number of power/0V pin pairs spread across its whole area.

Ideally there should be a power/0V pin pair adjacent to each signal pin, but given cost and space constraints this is usually only done for critical signals, such as high-speed (e.g. Gb/s) interconnections. Using differential pairs for Gb/s interconnections relaxes this constraint and allows one power/0V pair for every two or more signals. The lower the imbalance in the differential pairs (in their drivers, PCB traces, connector pins and external cabling, see [6]) the fewer power/0V pin pairs are required for a given number of signals, for a given level of EMC performance.

Figure 8A    Example of a 10-way connector

(This example would not be good enough for differential pairs,
or where all the signals had very high data rates or frequencies)

A decoupling capacitor (typically 10 - 100nF) should be placed between the power and ground very near to each power/0V pair of connector pins. Where the main power enters the PCB, a bulk decoupler should also be located. In the past bulk decoupling above 470nF was always done with electrolytics, but multilayer ceramics are now available with up to 100µF (with low voltage ratings) and they should be much more effective than electrolytics, as well as being smaller, more reliable, and reversible.

The above guide also applies to connections to cables, and also to inter-board connections. When used in a backplane type of assembly, spreading power/0V pin pairs across the whole area of each connector helps achieve a low impedance connection at radio frequencies (RF) between all of the planes in the assembly. In a mezzanine board type of assembly the same is also true, and it is also recommended to spread 0V connections all around the edges of the daughter/mezzanine boards (and maybe over their areas too). This also helps control cavity resonances (as discussed in section 6 of [7]).

Where there are multiple power supplies associated with a signal, as in some analogue designs (e.g. +/-15V), the above guide applies but instead of power/0V pin pairs they should be power/0V pin triples (e.g. +/-15V and 0V). For more on off-PCB interconnections and EMC see [8].

## 2 Low-K dielectrics

Homogeneous substrates (as opposed to epoxy-glass substrates such as FR4) generally have a lower dielectric constant ('$k$') than FR4, and also a lower loss tangent. Examples include pure polymer, and liquid crystal polymer (LCP). Teflon (and even foamed Teflon) substrates are sometimes used where very low $k$ is important, but the softer the substrate the harder it is to process in a PCB assembly. The lower $k$ makes the velocity of propagation higher than FR4, and the lower losses make it possible to send high frequency signals further and still maintain good SI.

The EMC benefits of low-$k$ substrates mainly resides in the lower levels of imbalance that can be achieved with differential pairs, see section 5.2 of [6]. Higher impedance lines use thinner traces, but when using thin layers of substrate the trace width can be too small for low-cost manufacture. Lower-$k$ dielectrics use wider traces for the same characteristic impedance (all

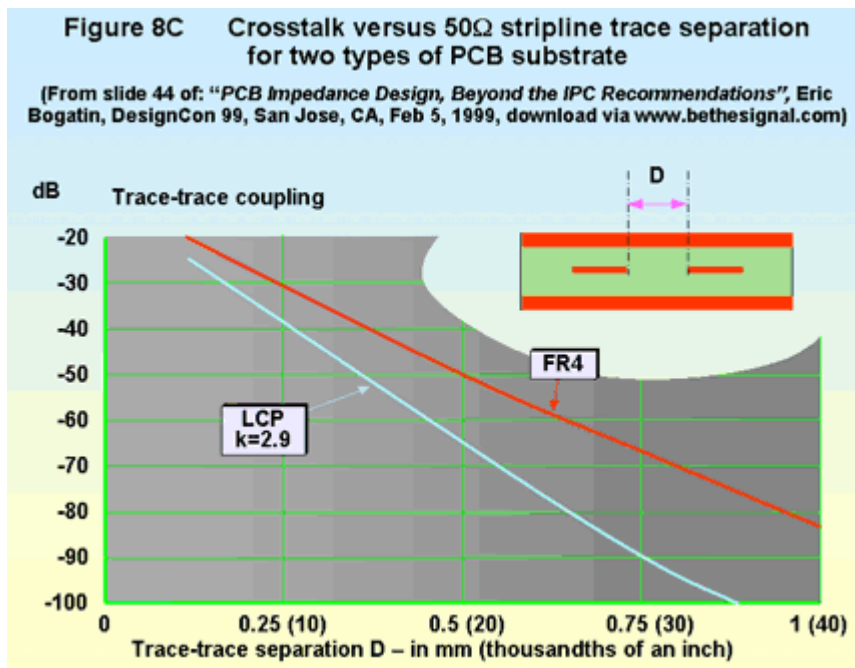else being equal) so can make it easier or less costly to use higher impedance transmission lines. As well as being useful in their own right, higher impedance transmission lines have lower currents flowing in them, so their emissions are lower.

The use of low-*k* substrates has been commonplace in microwave applications (e.g. satellite communications) for decades. As signal speeds increase it has been expected for a long time that the use of low-*k* substrates would migrate to more mainstream PCBs, particularly PC motherboards and cellphones, but so far designers have been clever enough to find ways to keep on using low-cost FR4 and its cousins. Figure 8B shows the relationship between SI, substrate loss tangent, and trace length.



Figure 8B    Signal loss versus type of PCB substrate
– for a trace width of 0.25mm (10 thousands of an inch)
(From slide 32 of: "Practical Characterization and Analysis of Lossy Transmission Lines"
Eric Bogatin, Gary Otonari, PC Design West, March 2001, download via www.bethesignal.com)
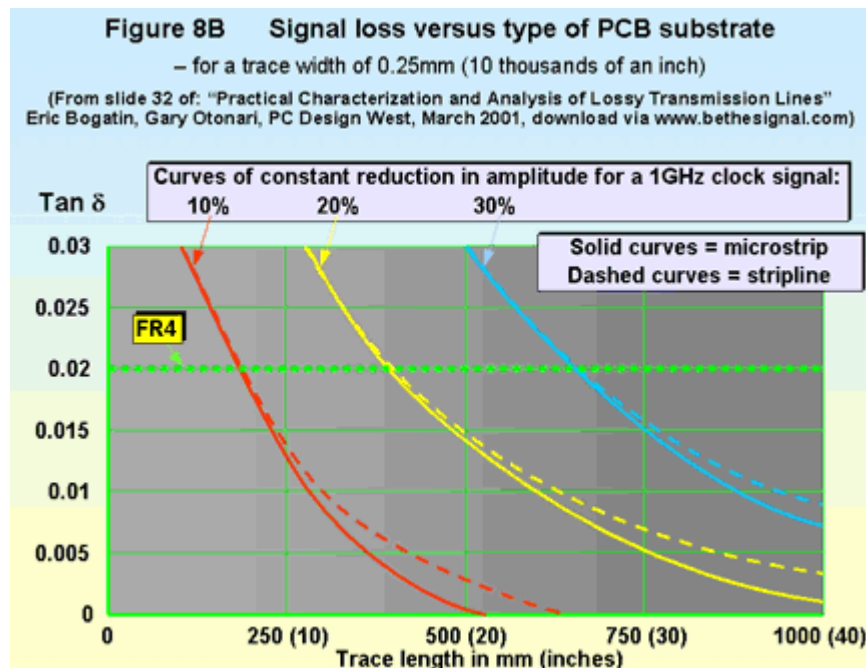
When signals are 10Gb/s or more, [9] shows that traces 600mm long or longer on glass-fibre substrates like FR4 will have serious SI and/or EMC difficulties. Section 5.2 of [6] and its Figure 6AK shows a method of using one or two layers of homogenous dielectric to gain their benefits in an otherwise low-cost PCB.

Although it is not the intention of this series to deal with SI issues, Figure 8C shows an interesting comparison of crosstalk versus trace-to-trace separation between FR4 and LCP.

Figure 8C     Crosstalk versus 50Ω stripline trace separation for two types of PCB substrate

(From slide 44 of: "PCB Impedance Design, Beyond the IPC Recommendations", Eric Bogatin, DesignCon 99, San Jose, CA, Feb 5, 1999, download via www.bethesignal.com)

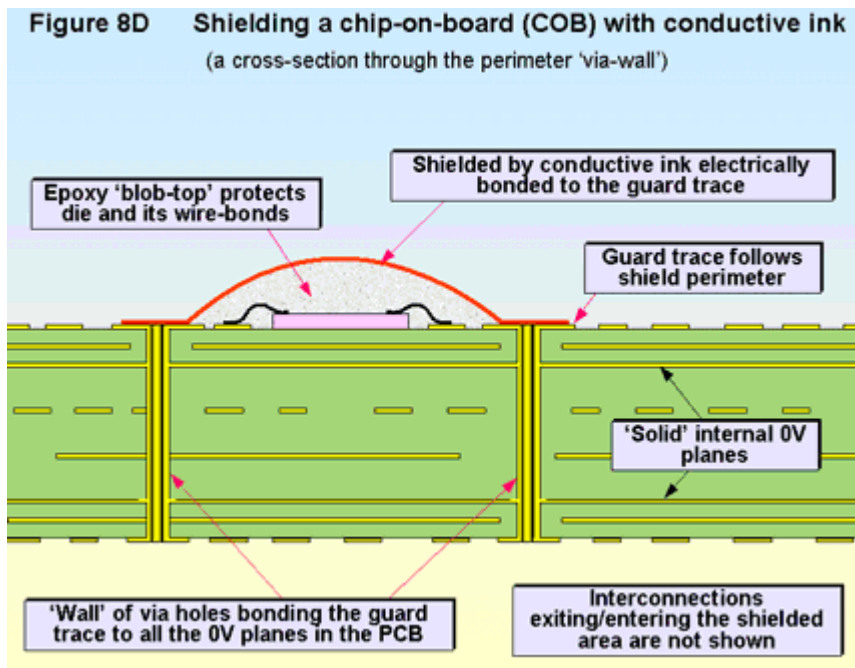## 3 Chip-scale packages (CSPs)

These very small IC packages [10] can generally be made to have excellent SI and EMC because their thinner packages place their silicon in closer proximity to the 0V plane in the PCB, so the image plane effect [11] is more powerful. Also, their much smaller internal interconnections means they are less effective as accidental antennas at frequencies below their first resonance, and their first resonance is at a much higher frequency, so these devices tend to emit less from their bodies.

However, a problem with these very small devices is that they allow much higher switching edges and higher-frequency noises to be conducted into the PCB's power distribution and signal traces, which can worsen emissions considerably. I know of someone using a CSP with a clock frequency of 1kHz, that exceeded the emissions limits all the way up to 1GHz (the one-millionth harmonic of its clock). But they can generally be made to have better EMC than the larger devices they replace, if all of the recommendations in this series of articles are followed.

## 4 Chip-on-board (COB)

In COB, bare silicon die (chips) are glued onto a PCB and wire-bonded to gold-plated PCB pads, then 'blob topped' with epoxy resin or silicone for protection. This is a very low-cost PCB assembly technique most often used on PCBs for high-volume low-cost consumer products, or for very rugged and reliable products. But despite its small PCB footprint, low cost (in volume) and ruggedness, COB seems to be ignored by most other product designers.

The small size of COB and its close proximity to the 0V plane in the PCB means it has better EMC, but if they still need to be shielded they can use small shielding cans as described in [12]. However, it may be easier and less costly to shield a COB by printing conductive ink over the top of its 'blob-top' protection. The size of the shield formed in this way is usually so small that internal resonances only start above 10GHz, and there are no apertures in it (above the surface of the PCB), so its shielding effectiveness can be very high up to many GHz. See [12] for more on shielding on PCBs.

**Figure 8D    Shielding a chip-on-board (COB) with conductive ink**
(a cross-section through the perimeter 'via-wall')

Epoxy 'blob-top' protects die and its wire-bonds

Shielded by conductive ink electrically bonded to the guard trace

Guard trace follows shield perimeter

'Solid' internal 0V planes

'Wall' of via holes bonding the guard trace to all the 0V planes in the PCB

Interconnections exiting/entering the shielded area are not shown

## 5 Heatsinks on PCBs

### 5.1 EMC effects of heat sinks

The stray capacitance between a metal heatsink and the IC or power semiconductor it is cooling injects stray noise currents into the heatsink, when the voltages in the IC or power transistor fluctuate. As a result, a floating heatsink experiences a fluctuating noise voltage, giving rise to electric field emissions. Since heatsinks can be quite large, and tend to be quite high above the 0V plane in their PCB, they can be very efficient radiators of emissions.

Inside the IC or power semiconductor, the stray capacitance to the heatsink arise from its bond wires and lead frames, maybe even (at above 1GHz) from its silicon metallisation patterns. Many types of power transistors connect one of their terminals (e.g. collector, drain, anode) to their metal tabs or metal body, and the stray capacitance between these relatively large areas of metal and the heatsink they are mounted on can be as much as 100pF.

If the heatsink is connected to a reference voltage, it will inject fluctuating currents into it, and also experience a fluctuating voltage depending on the impedance of the connection and magnitude of the current. If the reference voltage is not the correct one for the purpose, the noise currents injected into it in this way can cause self-interference or give rise to more emissions.

The electric fields from the heatsinks can radiate directly, causing problems with radiated emissions, and they can also couple with conductors and metalwork to cause problems for conducted emissions. If a heatsink causes emission problems, immunity problems might arise from external electromagnetic fields coupling to the heatsink, from where it might couple into the IC or power device. So reducing heatsink emissions can also improve immunity. Low-frequency analogue circuits that have no emissions above 150kHz might nevertheless need to use the heatsink techniques described here to improve immunity.

For good EMC (emissions and immunity) it is important to connect a heatsink to the reference voltage for the semiconductors that are actually causing the fluctuating voltages in the first place. In the case of an IC this is its 0V plane, and in the case of a power transistor it is one of the voltage rails that it draws power from, usually the one that is connected to the device via

the lowest impedance. It is also important that the connection method, and the plane or power rail, have low impedance at the highest frequency of concern. It will be assumed in the rest of this section that the heatsink is connected to a copper reference plane in the PCB, as this is generally what is required to control frequencies above a few MHz.

This technique can be thought of as reducing the loop area of the stray capacitive currents that originally flowed from a conductor in the IC or transistor into the heatsink. (For good EMC, it is always best to return stray currents to where they originated, using the smallest path length and smallest loop area. The path length should be much less than λ/10 at the highest frequency of concern - the shorter the better.)

For all but the simplest heatsinks (e.g. rectangular blocks of metal) accurate analysis at up to 1GHz (or more) requires computer-aided simulation, taking into account.

- The heatsink's geometry (shape)
- The types of semiconductors and their locations
- The proximity of any 0V planes and/or chassis
- The physical structure of any connection to a reference
- The physical characteristics of the reference it is connected to
- The proximity of the heatsink to the actual source of the emissions (e.g. a silicon chip and its bond wires and lead frame)
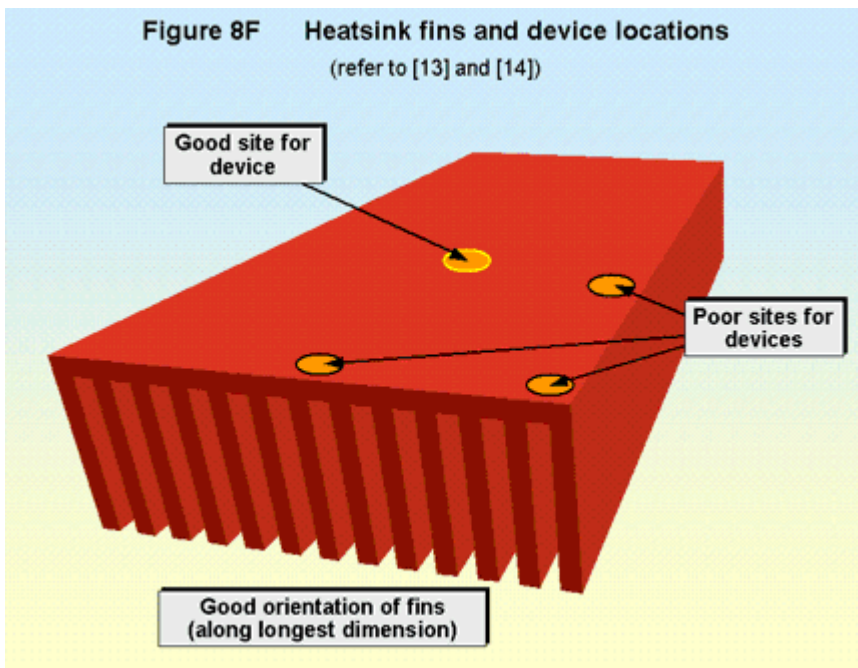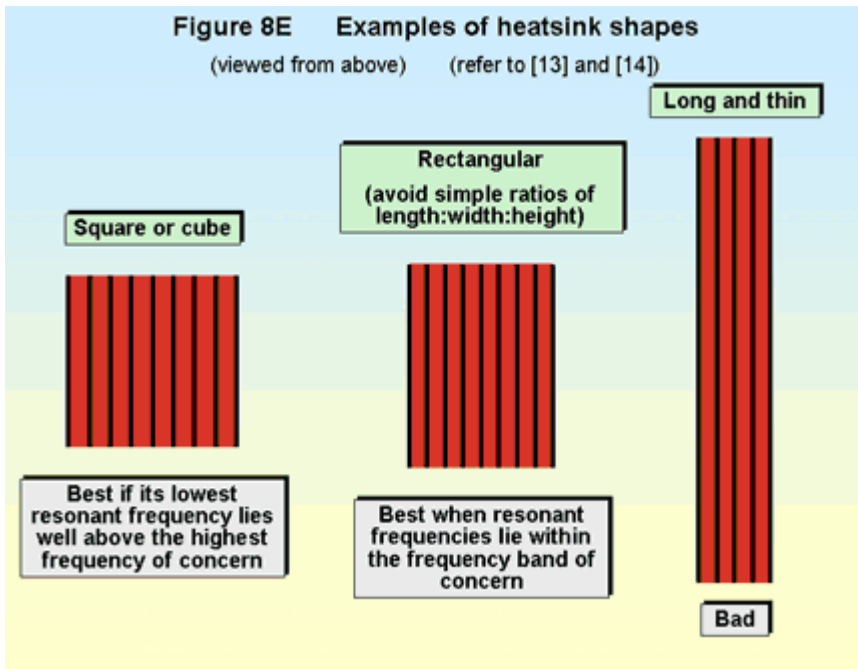
The rest of this section will discuss these design issues in more detail.
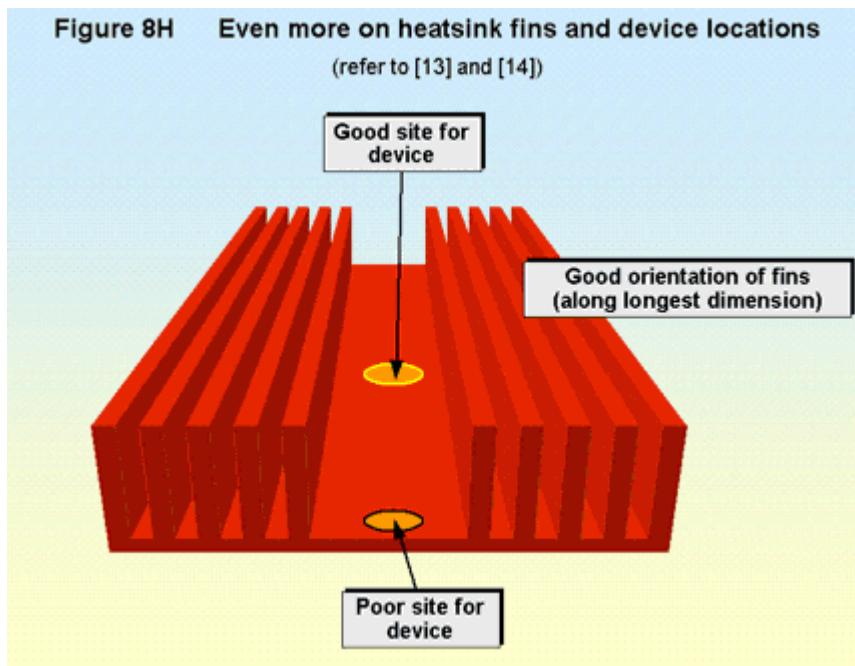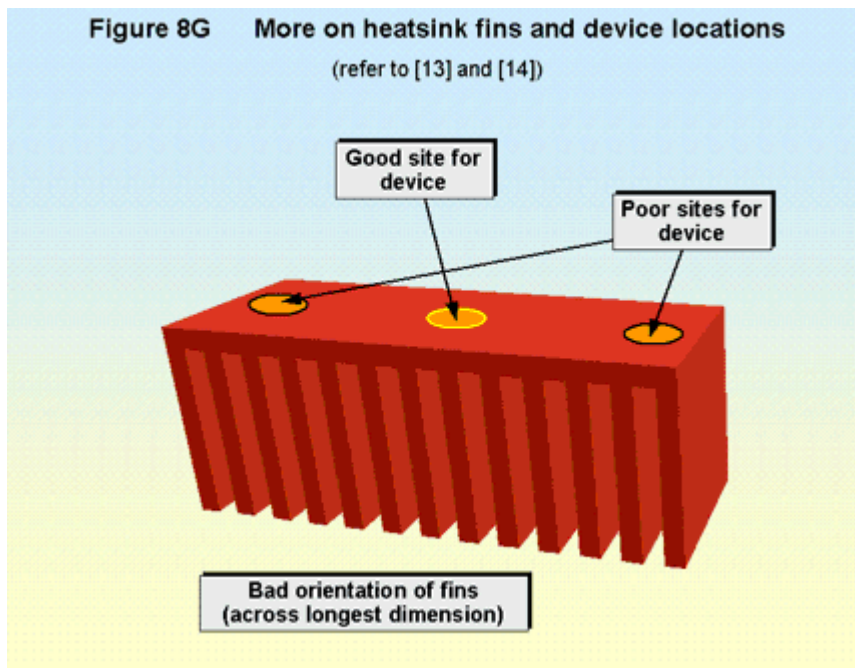
5.2 Heat sink RF resonances

[13] and [14] contain a great deal of useful information on heatsink resonances, including guidance on heatsink shapes and device locations. Resonance effects only begin to occur when any dimension of the heatsink, or of the cavity (or cavities) it creates, exceeds λ/10 at the highest frequency of concern. When a resonant frequency happens to coincide with a signal frequency or its harmonic, emissions from heatsinks can increase by 30dB or more. The lowest resonance frequency is given by the half-wave resonance of its longest 'three dimensional' diagonal. For example, a 60mm cube heatsink should have its first (lowest) resonance around 1.4 GHz. [7] discusses cavity resonances between PCBs and chassis, and between two PCBs, and its design guidelines can also be applied to cavities created by heatsinks.

Square or cube shaped heatsinks tend to have the highest resonant frequencies in their structure, so are good if their lowest resonant frequency is well above the highest frequency of concern. But such symmetrically-shaped heatsinks tend to have a higher Q at their lower resonant frequencies, so can be problematic if these frequencies fall within the frequency range of concern (especially when they fall on the same frequency as a clock harmonic). To avoid this problem, make them rectangular, but not too long and thin, and avoid simple ratios of length : width : height (e.g. 1 : 2 : 3), as shown in Figure 8E.

Where heatsink resonances exist within the frequency range of concern, the best place for the IC or power transistor is generally in the centre of the heatsink's base, which is usually best for thermal performance too. Edge locations cause greater resonant gain and higher emissions. In the centre of a channel between fins is best until the channel itself resonates (but this is generally a much higher frequency than the base resonances). Vertical fins reduce the resonant gain (the Q) for the direction in which they run, so it is usually best for the fins to run along the longest heatsink dimension. Vertical pins increase the resonant gain of most/all of the resonant modes. See Figures 8F, 8G and 8H.

**Figure 8E    Examples of heatsink shapes**

(viewed from above)      (refer to [13] and [14])

Long and thin

Rectangular

(avoid simple ratios of length:width:height)

Square or cube

Best if its lowest resonant frequency lies well above the highest frequency of concern

Best when resonant frequencies lie within the frequency band of concern

Bad

**Figure 8F    Heatsink fins and device locations**

(refer to [13] and [14])

Good site for device

Poor sites for devices

Good orientation of fins (along longest dimension)

Figure 8G     More on heatsink fins and device locations
(refer to [13] and [14])

Good site for device

Poor sites for device

Bad orientation of fins
(across longest dimension)

Figure 8H     Even more on heatsink fins and device locations
(refer to [13] and [14])

Good site for device

Good orientation of fins
(along longest dimension)

Poor site for device

## 5.3 Bonding heatsinks to a PCB plane

**Bonding heatsinks to a copper reference plane in the PCB (usually the 0V plane in the case of signal processing ICs), generally reduces emissions and increases the frequency of the lowest resonance [15], [16], [17], [18]. All such bonds should be very short and direct, and can usefully be combined with mounting hardware (e.g. short mounting pillars). Multiple plane bonds should be spread around the heatsink to reduce the bonding inductance and raise the heatsink's resonant frequencies.**

**Low bonding inductance between heatsink and plane is very important - [16] found that radiated emissions increase with the inductance of the heatsink's bonding network approximately to the power of 3.5. [16] also found that evenly distributed bonds are important - they can have 20dB lower emissions than if the same number of bonds are unevenly distributed but achieve the same overall inductance. Intel recommends [18] bond spacings of less than λ/4 at the 3rd harmonic of the processor core's clock frequency (e.g. <25mm for a**

**1GHz clock) but I recommend less than λ/10 at the highest frequency of concern for heatsink emissions or immunity.**

**Figure 8J shows a computer simulation of the fields around an example heatsink when it is 'floating' (not bonded to anything), and Figure 8K simulates the same heatsink bonded to an infinite plane with four direct connections - one at each corner. The simulator used was FLO/EMC running on a modern PC, and although it calculated the three-dimensional field patterns for all of the frequencies from d.c. to 10GHz, these two figures just show the field patterns in one plane at just one frequency.**



Figure 8J    Example of computer-aided simulation of the emissions from a 'floating' heatsink (using FLO/EMC, www.floemc.com)



Figure 8K    Simulation of the Figure 8J heatsink with a ground to a plane at each of its four corners (using FLO/EMC, www.floemc.com)
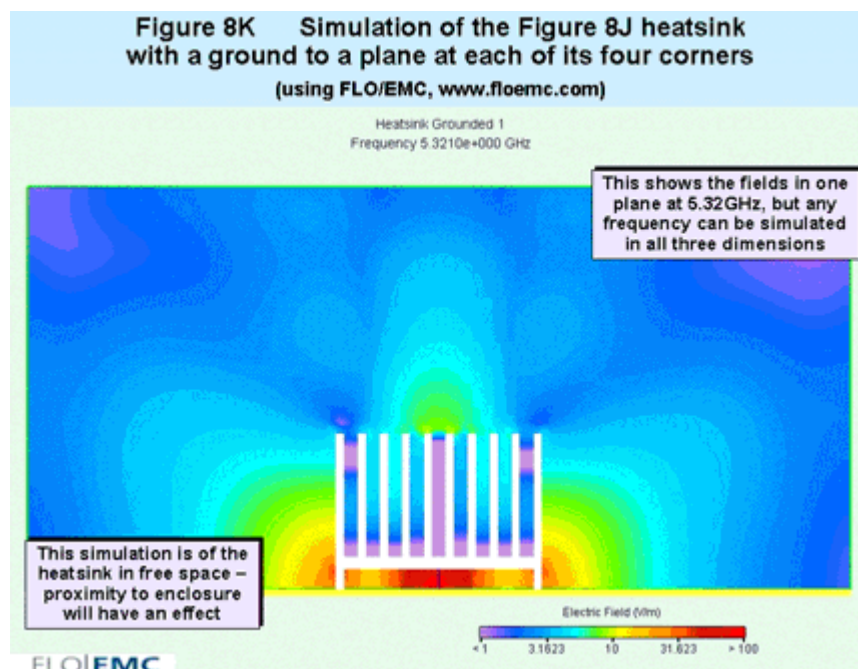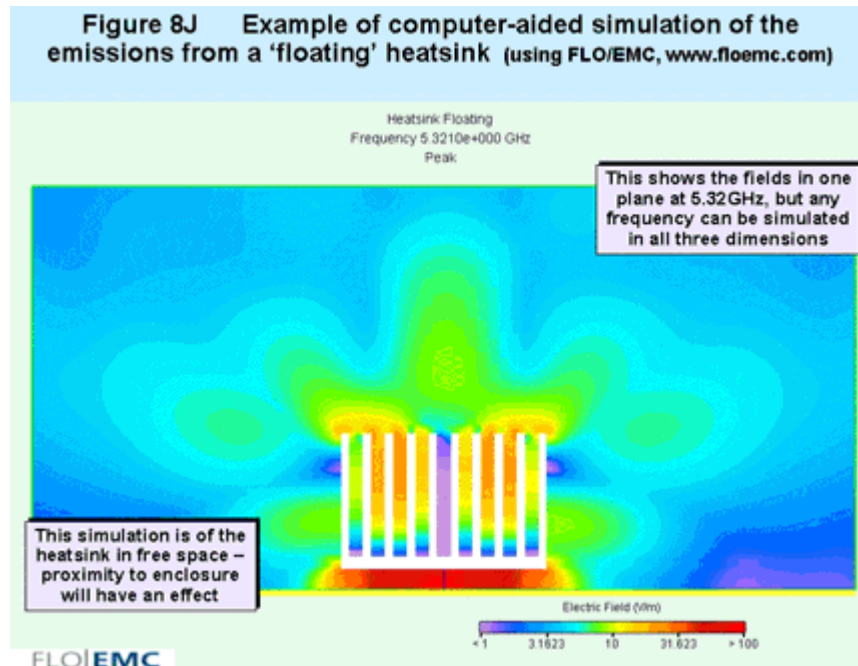
**Figure 8L graphs the field strength versus frequency at a point above the heatsink, for the situations described in Figures 8J and 8K. The floating heatsink is clearly a very good emitter of fields from d.c. to 2GHz, plus it has a resonance at 5.32GHz; whereas with four bonds at its corners it has significantly lower emissions from d.c. to 1GHz and its lowest resonance is at 2.0GHz.**

Figure 8L    Simulating the electric fields above the heatsinks in Figures 8J and 8K (using FLO/EMC, www.floemc.com)
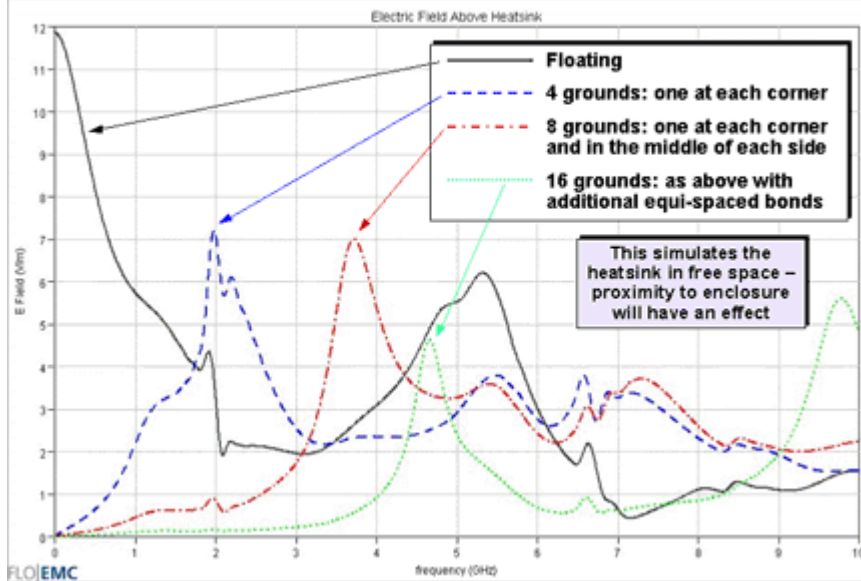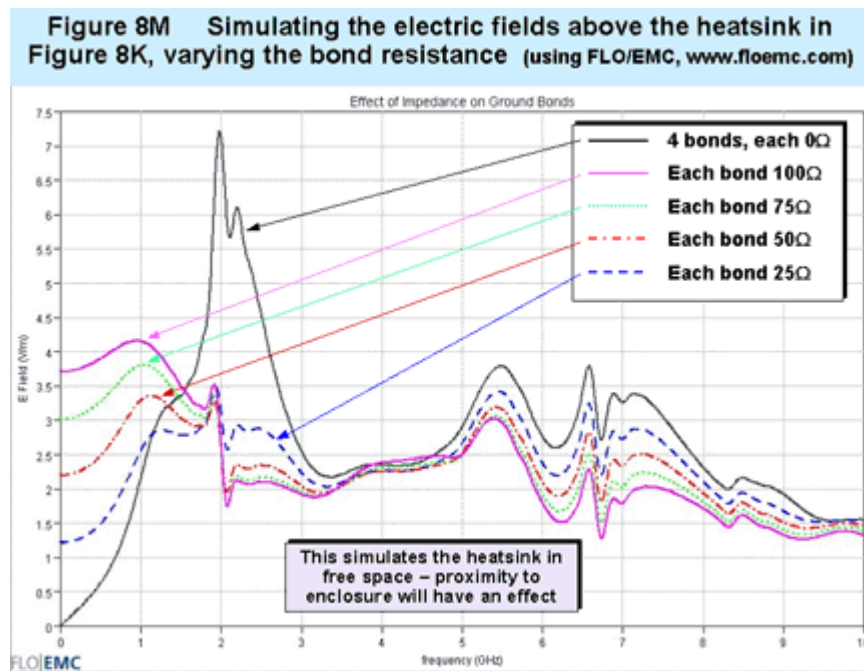
**Figure 8L then goes on to simulate the same heatsink with eight and sixteen equally-spaced bonds. With eight it has significantly lower emissions than the floating heatsink from d.c. to 2GHz and its lowest resonance is increased to 3.75GHz. With 16 it has very much lower emissions from d.c. to 3GHz and its lowest resonance is further increased to 4.6GHz. Adding even more bonds would continue the trend: increasing the frequency of the lowest resonance; and reducing the level of the emissions below about 75% of that frequency.**

**Note that the simulation with sixteen bonds has a new resonance at 9.75GHz, peaking at about +12dB, which was not present in the others. It is not yet known what the physical cause of this is, but it shows the value of simulation in revealing unexpected results that could potentially cause problems, early in a project before any hardware has been made, when changes cost the least.**

**The above simulations employed 'direct' heatsink bonds assumed to have a resistance of 0Ω. Section 5.9 of [7] showed that, for PCB plane-chassis bonding, using resistive (lossy) bonds instead of 0Ω bonds reduced the amplitude of the resonant peaks at the cost of increased emissions below the resonant frequency. Figure 8M shows a simulation that explores whether the same trade-off might work for heatsink-plane bonds. It shows that with bond resistances of 25Ω and above the lowest resonance is almost completely damped, at the cost of increased emissions below the resonant frequency. Higher values of bond resistance result in higher emissions from d.c. up to the lowest resonant frequency, and lower emissions above that frequency. Using 50Ω ferrite beads instead of resistors for the bonds should restore the good emissions performance at the lower frequencies where the ferrites have a very low impedance.**

Figure 8M    Simulating the electric fields above the heatsink in Figure 8K, varying the bond resistance (using FLO/EMC, www.floemc.com)

These simulation results point to some interesting and useful applications, especially where it proves impractical to move the lowest heatsink resonance above the frequency range of concern. However, I know of no one who has yet tried to use resistive or ferrite heatsink-plane bonding in a real design so do not know if it is really as effective in real-life as the simulation implies.

Some processor sockets may now be available with built-in heatsink grounding posts for heatsinks, but in the main it is up to each designer to work with his electrical and mechanical engineering colleagues to come up with a low-cost solution that is easy to assemble on the PCB and does not consume too much board area. Figure 8N shows an example designed by Intel Corporation [19] - simply a metal stamping that solders to the PCB and makes spring-finger contact with the heatsink above the device (the hole in the middle is for the thermal transfer medium between the IC and heatsink).

It is relatively easy to design and make similar metal parts, and they will work well if designed with the right kind of springy metal, with plating to make a reliable electrical connection to whatever the heatsink is made of. Manufacturers of PCB shielding cans or spring-finger gaskets should be able to recommend the best materials, and may also be the most appropriate manufacturers for them.

Figure 8N    Example of a heatsink-bonding part designed by Intel
for use with the 423-pin Pentium® IV    (from [19])
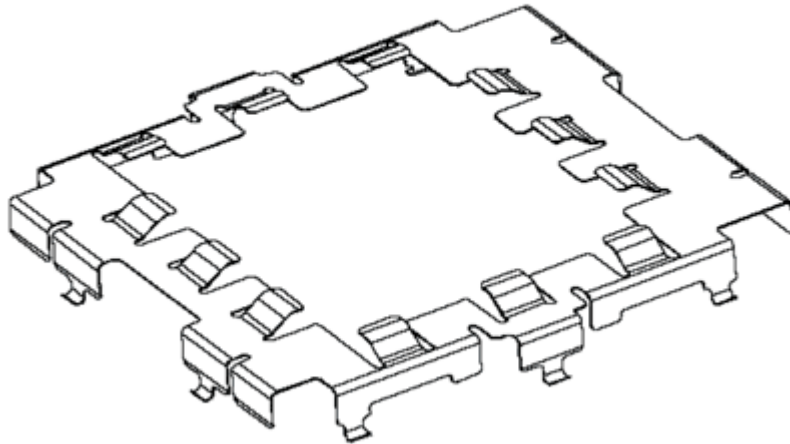
When designing methods of bonding PCB planes to heatsinks, be aware that the desired low impedance can be severely compromised by corrosion at any contacts between dissimilar metals, oxidation effects over time (especially with aluminium or steel), manufacturing residues or protective coatings on metal parts (especially anodising or passivation using insulating polymer coatings). Corrosion can be a big problem where condensation or other wetting by liquids can occur. Use appropriate techniques to ensure that all of the intended heatsink bonds maintain very low impedance contact despite the anticipated physical, chemical, biological and climatic exposures over the entire lifecycle.

In manufacturing always check that each batch of metal parts (heatsinks, fixings, etc.) have surfaces with very high conductivity, using smooth probes with low contact pressure. This is recommended because some manufacturers or platers of metal items do not understand surface conductivity specifications, and are liable to apply coatings of insulating clear polymer (which are invisible to the eye) on a whim. Also, some manufacturing buyers will purchase anodised aluminium instead of alochromed aluminium, or polymer-passivated coated items despite what is written on the purchasing specification, because "it looks just the same: grey metal".

### 5.4 Combining shielding with heatsinkingM/h4>

Looking at Figure 8N, the next step seems obvious - design the part that connects the heatsink to the PCB so that it encloses the IC in a 'Faraday Cage' - shielding the IC and making even better bonding between heatsink and plane to give the best EMC performance at PCB-level. Shielding ICs at PCB level (as discussed in [12]) restricts airflow and can cause overheating, and combining shielding with heatsinking can solve this problem. An example of a suitable construction is shown in Figure 8P.

Figure 8P      An example of combining
IC shielding with heatsinking

Heat sink

Thermal washer (e.g. gel pad)

PCB

A metal 'picture frame' all around the IC

Requires frequent low-impedance electrical contacts to heatsink and PCB
plane all around its perimeter

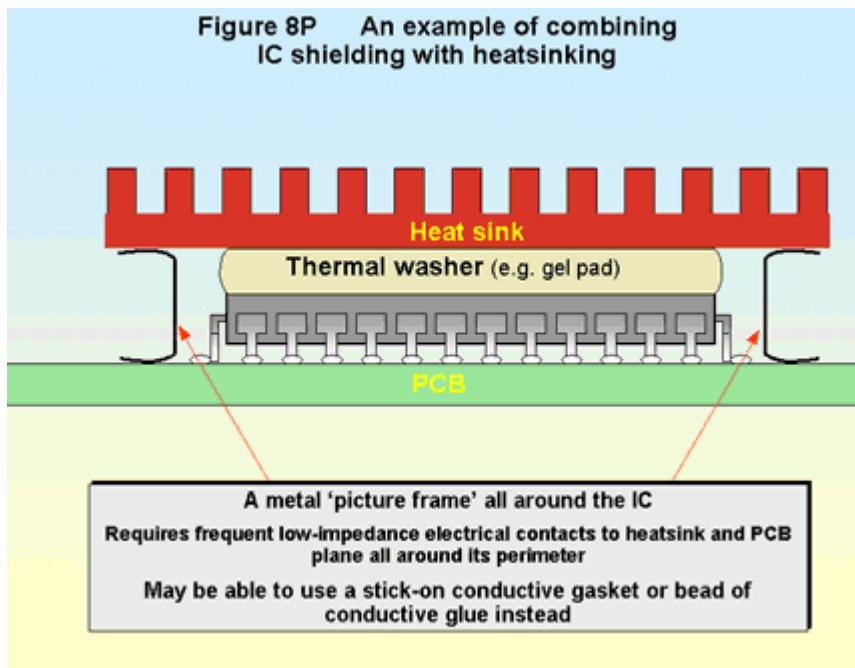May be able to use a stick-on conductive gasket or bead of
conductive glue instead

**Figure 8P shows a metal structure surrounding the IC, with spring finger contacts top and bottom to make electrical contact with a 0V guard trace on the component side of the PCB (tinned and not covered by solder resist), and to make electrical contact with the base of the heatsink. Such metal structures are sometimes called 'picket fences' or 'picture frames'. The figure does not show any alignment or mounting pins on the picket fence, or how the heatsink is held in place.**

**There are many design possibilities for bonding the heatsink to the 0V plane all around the IC to make a shield: it could use a picket fence made of metal or plated plastic, soldered to the PCB at multiple locations instead of using spring fingers. If the IC has a low enough profile it could use a compressible conductive gasket or even a bead of conductive glue. It could be a metal structure that is welded or soldered directly to the heatsink, or formed as an intrinsic part of it.**

**To be an effective shield the picket fence (or whatever) must be continuous all around the IC, and must make frequent low-impedance electrical contacts to both the heat sink and a 0V 'guard ring' or plane on top of the PCB all around its perimeter. The 0V guard ring or topside plane must be via'd very frequently to the main 0V plane, *all around its perimeter*. The design guidance for the spacing of the contact points and of the vias are the same as for the PCB shields in section 2 of [12].**

**5.5 Other heatsink techniques that may help**

**Where practical, increase the flow of air (or other cooling medium) so as to be able to reduce the size of the heatsinks to increase their lowest resonant frequencies.**
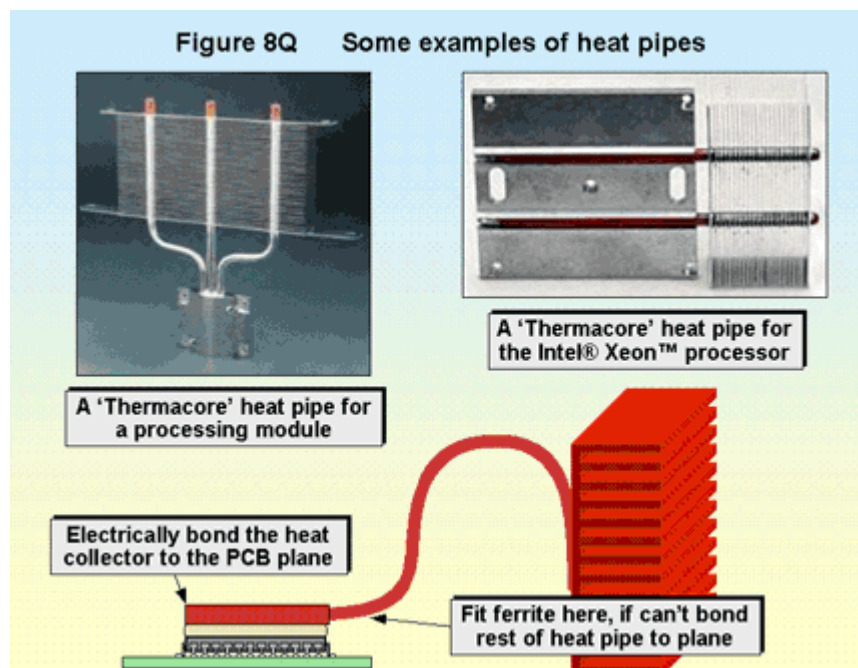
**Use insulating thermal interface materials that have a lower dielectric constant and/or are thicker to reduce the amount of stray capacitance between the heatsink and the semiconductors' silicon chips and bond wires [20].**

**Shielded thermal interfaces may be able to be used. These have two insulating thermally-conductive layers with a metal shielding layer sandwiched between them, and help to keep stray noise currents out of the heatsinks. The inner metal layer should be bonded directly to the appropriate plane, ideally at multiple places around its periphery to reduce emissions and**

increase its lowest resonant frequency, as discussed earlier for heatsinks.

When heat pipes are used, their length will generally give them lower resonant frequencies than equivalent heatsinks fitted directly to devices. Heat pipes should be bonded to the appropriate PCB plane as described earlier, along their length, to reduce their emissions at low frequency and place their resonant frequencies away from any clock harmonics or other strong signals in the PCB. Ideally, the heat pipe resonant frequency would be located higher than the highest frequency of concern, but this is often not possible.

The part of the heat pipe that collects the heat from the device is quite small, and so has much higher resonant frequencies - on its own - than the overall heat pipe (or an equivalent heatsink fitted directly to the device). So if it is difficult to bond the whole heat pipe assembly to the PCB plane or associated chassis (in turn well bonded to the PCB plane, see [7]) - it might be possible to bond the heat collector to a PCB plane as described earlier, then clip a ferrite around the heat pipe's 'pipe' element to reduce the emissions from the larger part of the heat pipe's assembly, as shown in Figure 8Q.



Figure 8Q    Some examples of heat pipes

A 'Thermacore' heat pipe for a processing module

A 'Thermacore' heat pipe for the Intel® Xeon™ processor

Electrically bond the heat collector to the PCB plane

Fit ferrite here, if can't bond rest of heat pipe to plane

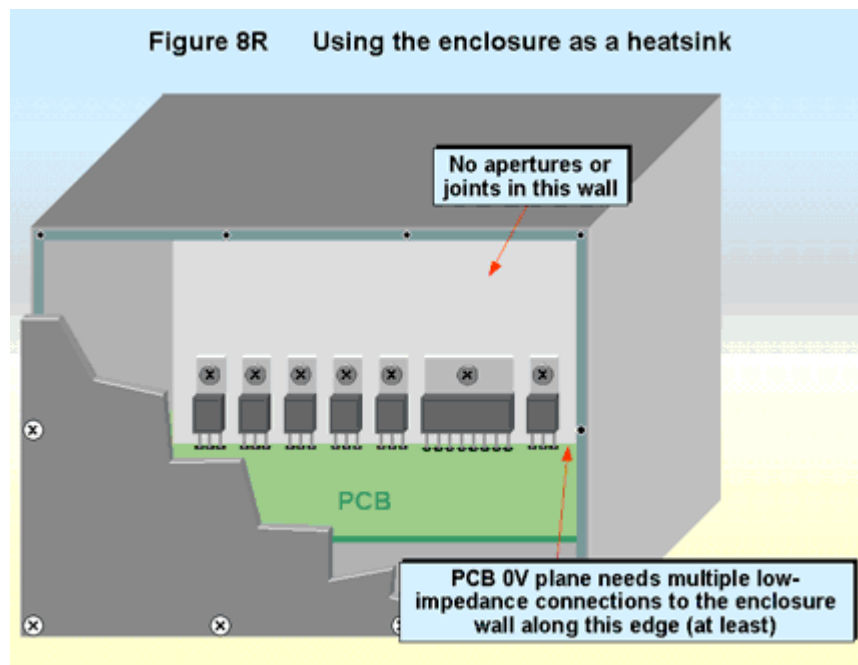## 5.6 Heatsinks for power devices

The methods described in 6.2 to 6.5 above, have centred on heatsinks for ICs, but the methods described can also be applied to heatsinks for power devices.

Where there is no electrical isolation between a device and its heatsink (a low-cost method commonly used) the heatsink may need to be connected to the relevant power plane via a suitably-rated capacitor that provides the necessary galvanic isolation. The type and number of capacitors used, their location on the PCB, and their trace routing will have a significant effect on the heatsink bonding at higher frequencies, similar to the issues discussed for decoupling capacitors in [21]. With such 'live' heatsinks, electrical safety issues should always be taken into account in the design, (e.g. the total value of capacitance used may be limited to prevent electric shock, the capacitor may need to be safety-approved to Class Y1 or Y2) see Volume 4 of [5].

The heatsink tabs of power transistors can also be effective radiators of emissions even when an external heatsink has not been fitted. I once saw a low-cost inverter on a PCB not much

more than 20mm square, with a single TO-220 power-switching device standing vertically above the PCB. The emissions from the TO-220's metal tab caused the inverter to fail the generic emissions limits (conducted and radiated) over a wide range of frequencies, but of course it could not be connected to the reference because it was connected internally to the drain of the switching semiconductor. A capacitor of a few pF, from the tab to a suitable reference, solved the problem.

Power devices are often mounted along one edge of a PCB so that they can use the metal enclosure as a heatsink to save cost, as shown in Figure 8R. In this case the metal enclosure should be bonded to the main 0V plane in the PCB at multiple points, especially in the vicinity of the power devices. The purpose is to encourage the stray currents caused by the power devices to be returned to the PCB (and thence to all relevant power rails) with very low impedance at the highest frequency of concern, so that (due to the skin effect) these currents remain mostly on the *inside* of the enclosure and do not cause external emissions. [7] deals with PCB-chassis bonding, so its design guidance is relevant here. The metal enclosure should ideally have no apertures or joints in the vicinity of the power devices and the nearby enclosure-PCB bonds, as the intense currents flowing in the enclosure in this area will make them 'leak' RF fields.



Figure 8R    Using the enclosure as a heatsink

No apertures or joints in this wall

PCB

PCB 0V plane needs multiple low-impedance connections to the enclosure wall along this edge (at least)

## 6 Package resonances

The frequencies that are now being used (e.g. radiocommunications, serial data, clocks and their harmonics at many GHz) plus the increasing requirements to test for emissions up to as much as 40GHz in some radiocommunication compliance, EMC compliance or military applications, means that semiconductor packages themselves can be large enough to resonate and become very effective 'accidental antennas' in their own right - even before a heatsink is attached (see earlier).

Eric Bogatin, in [22], says: "Know the resonant frequencies of all packages and change the package geometry if there is an overlap with a clock harmonic" and I don't think it is possible to improve on this advice. If it is not possible to avoid overlaps between the package resonances and the clock harmonics that are necessary for the functioning of the circuit, and if spread-spectrum clocking is not sufficient or not possible for some reason, then shielding is likely to be the only answer. [12] describes PCB shielding techniques - the lowest-cost way to provide

shielding.

A 3-D field solver is required to accurately determine the resonant frequencies of packages in real applications, but it may be possible to roughly estimate the lowest resonant frequency using simple mathematics.

Chip-scale packaging (see earlier) provides the smallest possible package size. But where there are hundreds of contacts to be made to a device, affordable PCB technology might not be able to connect to them all because their contact pitch is so small. So HDI/microvia PCB technology [23] might be required simply to connect to the device, even if not to provide good EMC.

## 7 Eliminate the test pads for bed-of-nails or flying probe testing

Testing bare or assembled PCBs using 'bed-of-nails' or 'flying probe' techniques is a traditional manufacturing process that became common in the 1970s but is getting more difficult to employ as devices get smaller and the number of device pins rises. The test pads required for these test methods can seriously damage EMC, and may even damage SI causing functionality or reliability problems. And in many cases these test methods are no longer the most cost-effective.

It is very important for EMC that test pads are *not* added to *every* circuit node. Adding such pads can cause huge difficulties for the routing of high pin count devices such as BGAs (which are difficult enough to route anyway). But, more importantly, test pads on the end of traces act as 'accidental antennas' just like any other conductor - but very few designers think to analyse them to see if transmission line techniques are needed.

BGAs often end up surrounded by a forest of test pad accidental antennas, some of them using very long traces when compared with the electrical lengths of the signals or noise present on them. The stray capacitance associated with a test pad itself considerably increases the effective electrical length of the trace that connects to it. The result is that test pads behave as unterminated stub transmission lines (making them very efficient at damaging SI) and as effective accidental antennas above 100MHz (compromising emissions and immunity). As a rough example, a test pad with a trace length longer than 15mm will probably cause significant SI and EMC problems below 1GHz.

Unless each test pad is analysed from a transmission line viewpoint for SI, the signals measured by the test probes may be distorted and perfectly good PCBs may be sent for rework. Unless each test pad is analysed and designed as a matched transmission line for reasons of good EMC performance, such a PCB will probably require a great deal of costly shielding to comply with EMC emissions requirements. But the cost of doing this and the PCB space lost to test pads and their traces and terminators (where it is even possible to use them), makes it much more likely that it will be more economical for products using modern ICs to scrap such 1970s techniques and instead invest in good quality board manufacture.

However, there are well-proven PCB testing techniques that can locate faults without using test pads. "Device-centric" boundary scan testing (JTAG, etc.) techniques [24a] now allow many/all device faults to be located at the level of an individual component, using a PCB's own address and data busses and testing using only PCB's functional connectors, even where some of the devices on the board are not JTAG compliant. The 'Built-In Self Test' (BIST) facilities provided in most modern VLSI devices and in the software of many systems can also be accessed with JTAG, extending the usefulness of the technique even further [24b] [24c].

Board rework is very costly and reduces the reliability of the PCB, leading to increased warranty costs. So even if JTAG is not available it can be more cost-effective to rely on built-in

self test (BIST) and/or functional testing to check PCBs - and improve the quality of the circuit design, PCB layout, components and assembly so that the number of faulty PCBs becomes so low that it costs less to scrap the faulty boards than pay for bed-of-nails or flying probe testing.

I often come across companies whose purchasing departments insist on using the lowest-cost board manufacturers and assemblers, and lowest-cost component suppliers, thinking that they are saving their company money. But they do not consider that the *overall* costs of implementing the resulting bed-of-nails testing and reworking may be greater than what they saved by using such low-quality suppliers. In effect, they are 'saving money at any cost'. Marconi Instruments discovered in the 1980's that despite suppliers' claims to the contrary, there were known bad batches of components being sold as full-spec in the UK, and they were sold to the buyers who showed the least loyalty to their suppliers and drove the hardest bargains (surprise!). As in so many other areas of life, you get what you pay for. Although it clearly makes sense to reduce the cost of the bill of materials (BOM) - if in doing this the *overall cost of manufacture* (including the costs of warranty returns) is increased by more than the BOM cost savings, it is not cost-effective after all.

When the EMC problems of advanced PCBs are added, the cost balance shifts even further in the direction of eliminating test pads by using higher quality suppliers and manufacturing, and/or using JTAG design/testing techniques.

If test pads or connectors, or programming connectors must be used, it is important to add them to the PCB at its initial design, so that all of the functional and EMC proving is done with them in place (page 31-072 of [25]). In some companies the manufacturing departments add the test pads to the layouts after the PCB design is (supposedly) complete, but with modern semiconductor devices this can easily completely ruin the EMC, even if it doesn't affect functionality.

Traces connecting to test or programming connectors have exactly the same EMC problems as traces to test pads, so should be treated exactly the same way. Any traces that were only needed during prototyping should either be removed before EMC testing begins, or treated with the same care for EMC as if they were functional traces. Such traces may need to be treated as stubs on transmission lines (see [6]).

Some companies design their prototype PCBs with huge numbers of additional traces that go off to an appendage whose sole purpose is to test most/all nodes on the prototypes. When judged ready for manufacture the 'test appendage' is simply cut off the PCB artwork, leaving large numbers of test traces reaching to the edge of the PCB. On a PCB that uses modern digital or RF devices, with their very fast signal edges and/or high frequencies, such a technique makes it impossible for the prototype PCBs to have SI or EMC that bears any resemblance to the manufactured PCB, and also makes it very difficult, if not impossible, to achieve good SI or EMC (for reasons that should be obvious to anyone who has been following this series of articles).

Even if the redundant test traces are completely removed from the layout, the layout that remains is generally less than ideal because its traces were originally laid out to fit in with the test traces. Usually at this stage in a project there is no time left to reroute the whole PCB and then retest its functionality and EMC. This whole technique is something that was possibly acceptable with the slow devices and low frequencies of the 1970s or 80s, but is wholly inappropriate for modern electronic devices.

If the above technique must still be used, the traces and planes that are required for the functionality and EMC of the final design should be designed using the layers and stack-up they will have in the final product - taking all the guidance in this series into account. The test

traces should then be added as additional PCB layers taking care to avoid crosstalk with the original traces. The 'test PCB' appendage should be cut off as early in the project as possible, and its associated test trace layers removed completely. Even if the 'test appendage' design is carefully done and transmission-line matching maintained (which it often is not) removing it will significantly alter the EMC performance, and could also alter the functional and SI performance. It is best not to use this technique at all with advanced PCBs (see Figure 1D in [26] for how to determine if a PCB requires treating as 'advanced').

Where a bus is provided for a cable (or a docking station), and when it is unterminated when the connector is not plugged in, or when the product is undocked - arrange for these clocks and data to be disabled when not required.

## 8 Unused I/O pins

Unused input pins should be pulled up or down, to prevent unintended oscillation, see page 31-28 of [25]. I have seen equipment behave like a Class C oscillator and radiate about 2W of RF power at 200MHz, simply because one section of an 'HC00' style inverter was unused and accidentally left with its input floating. Tying inputs directly to a rail might cause functional problems, or make fault testing difficult, so it is best to use a resistor of 10K$\Omega$ or so for pulling-up, and 200$\Omega$ or so for pulling down.

Some EMC engineers recommend programming any unused I/O ports to be outputs, with a 'strong' output driving capability, then programming them LOW and connecting them to the 0V plane to reduce ground bounce. Some could be programmed HIGH and connected to the power plane to reduce power supply bounce. Ground and power supply bounce is a problem for SI, and also a source of common-mode noise voltages and currents that increase emissions - so reducing them is a good thing.

But a problem can occur with this technique if, during booting-up or interference, the pin programming could be altered even momentarily, possibly causing high currents to be drawn from the power supply that could cause malfunction or damage.

## 9 Crystals and oscillators

Crystals and crystal oscillators should lie flat down on the PCB, close to their 0V plane and very far away from its edges (at least 15mm, unless they are placed inside a larger shielding can as discussed in [12]). Where the 0V plane layer in the PCB is not closest to the side on which the crystal is mounted, it helps to place an area of 0V plane on the component side underneath the crystal and extending at least 5mm around its outline on all sides. This small plane should be via'd to the main (solid) 0V plane at least every $\lambda/10$ at the highest frequency of concern, taking the dielectric constant of the PCB substrate into account (e.g. for up to 1GHz, vias in an FR4 PCB should be placed at least every 15mm over the whole area). It may help some devices if their metal can is soldered directly to the 0V plane by a very short wire or a U-shaped metal strap.

## 10 IC tricks

Although not PCB design techniques as such, it is worth mentioning that there are an increasing number of advanced EMC techniques becoming available in FPGAs and ASICs. These include.

**a) Output drivers with user-programmable slew rates.**

**b) Output drivers with user-programmable drive 'strength'.**

**c) Drivers that automatically adjust their output impedance to reduce the overshoot and ringing on their signals.**

**d) Drivers with built-in voltage-doubling and output impedances that allow 'classical' (both-ends) matching of transmission lines with no loss of signal amplitude.**

**e) Differential receivers with increased sensitivity that allow 'classical' (both-ends) matching of transmission lines with no problems caused by the 50% loss of signal amplitude.**

**f) 'Dithered' output switching that prevents outputs from changing state at exactly the same time (helping reduce ground bounce and rail collapse and CM noise in the PCB's power structure).**

**g) On-chip, in-package, or in-multi-chip-module decoupling capacitors.**

**IC designers are constantly developing more techniques for improving the EMC of their devices, so when beginning a new ASIC of FPGA design always ask what is available. The lowest-cost place to control EMC is generally in the semiconductor devices themselves, even if they cost more as a result. Controlling EMC at PCB level is generally the next most cost-effective [26].**

**11 Location of terminations at the ends of transmission-lines**

**With the complexity of modern PCB layouts, it is sometimes difficult to locate transmission-line terminators immediately adjacent to the input or output pins of their receivers or drivers. [27] describes how to calculate if the resulting stub trace length is too great for SI, taking into account the inductance of the bond wires and lead frame inside the devices themselves. For good EMC, the stub trace lengths should be at least one-third of the length that is required for good SI.**
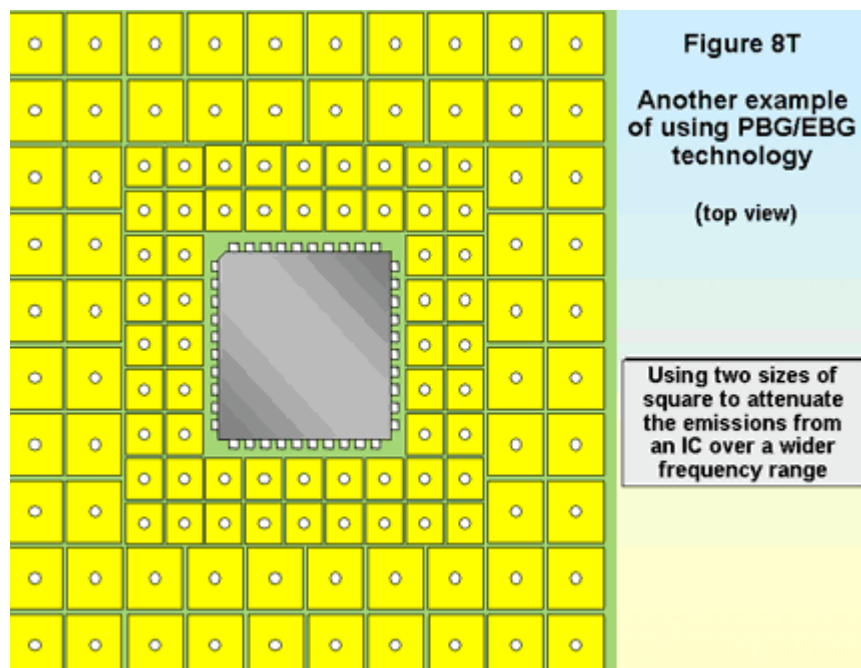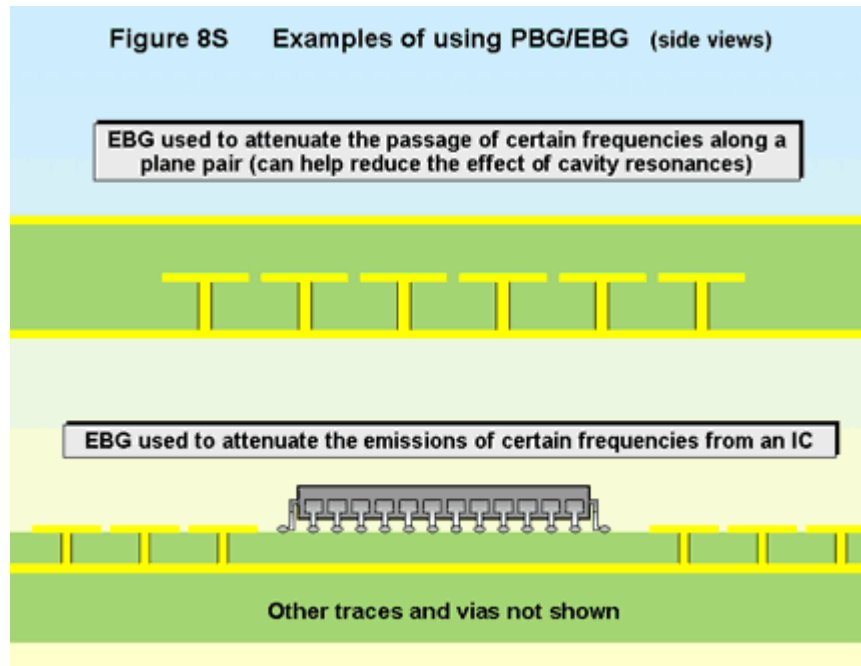
**Some PCB manufacturers offer embedded resistive metal layers (instead of copper), which can be etched to create transmission-line matching resistors inside the PCB stack-up itself. It is much easier to place such terminators exactly on the pins of their devices.**

**Some devices now operate so quickly that the length of their internal interconnections is sufficient - on its own - to prevent good enough line-matching for EMC reasons, even when the termination resistors are placed as close as possible to their pins. There is always some overshoot and ringing caused by the 'stubs' internal to the devices. This can be another reason for using the smallest packages available, such as the chip scale packages discussed earlier.**

**12 Electromagnetic Band Gap (EBG)**

**Arrays of small copper squares, each connected to a large plane by a via as shown in Figures 8S and 8T, can reduce the propagation of some frequencies along a PCB plane (or in a cavity between two planes) by as much as 30dB. This technology is generally either called photonic band gap (PBG) or electromagnetic band gap (EBG), and the earliest reference that I have is [28] (which itself references earlier**

**papers).**

Figure 8S    Examples of using PBG/EBG   (side views)

EBG used to attenuate the passage of certain frequencies along a plane pair (can help reduce the effect of cavity resonances)

EBG used to attenuate the emissions of certain frequencies from an IC

Other traces and vias not shown

Figure 8T

Another example of using PBG/EBG technology

(top view)

Using two sizes of square to attenuate the emissions from an IC over a wider frequency range

It requires a field solver or experiments to correctly assess the attenuation achieved by PBG/EBG structures, but I estimate that when using small square planes of side D as the elements of the PBG/EBG structure in a PCB employing an FR4 substrate, the useful attenuation occurs approximately over the frequency range 22/D to 42/D. D in millimetres gives frequencies in GHz, for example an array of 10mm squares creates useful attenuation over 2.2 - 4.2 GHz. Higher attenuation requires a larger area of PBG/EBG arrays.

PBG/EBG arrays can be used to reduce the noise or signals in one zone from flowing along a common plane to another zone, so might be able to be used all around the perimeter of a PCB zone, or around a 'connector zone', to help isolate

one area from another as described in [12], or even as an alternative to the guard trace technique described in [11]. Do not use PBG/EBG to attenuate waves that are important for the operation of the circuit, e.g. preventing the 0V/Power planes from providing current at certain necessary frequencies. For more information on this new technology with potentially very significant possibilities, see [29] [30] [31] [32] and [33]. [34] explores the possibility of combining PBG/EBG technology with high dielectric constant thin films to improve the performance of embedded decoupling capacitance (section 3.14 of [21] discusses embedded decoupling).

The number of vias required to use PGB/EBG in a PCB implies that it would be easiest to use with blind or buried microvia PCB manufacturing techniques (see Part 7 of the series [23]). PBG/EBG technology either uses more PCB area, or an extra PCB layer, so will add to costs, but it does provide a method of controlling propagation and resonances at frequencies above 1GHz that could be more difficult and/or more costly by other means. I do not know of anyone using this technique in anything other than experimental PCBs so far, and would be pleased to hear of any experiences of using it in a PCB in a real product.

### 13 Some final PCB design issues

Unused traces should have any clocks, data or high-frequency signals on them disabled to reduce emissions, and if they connect to receivers they may need suppressing to improve immunity.

If a trace has to route to the other side of a microprocessor, RAM or other noisy device, it should circle around the device instead of cutting across its footprint (see pafe 31-72 of [25])
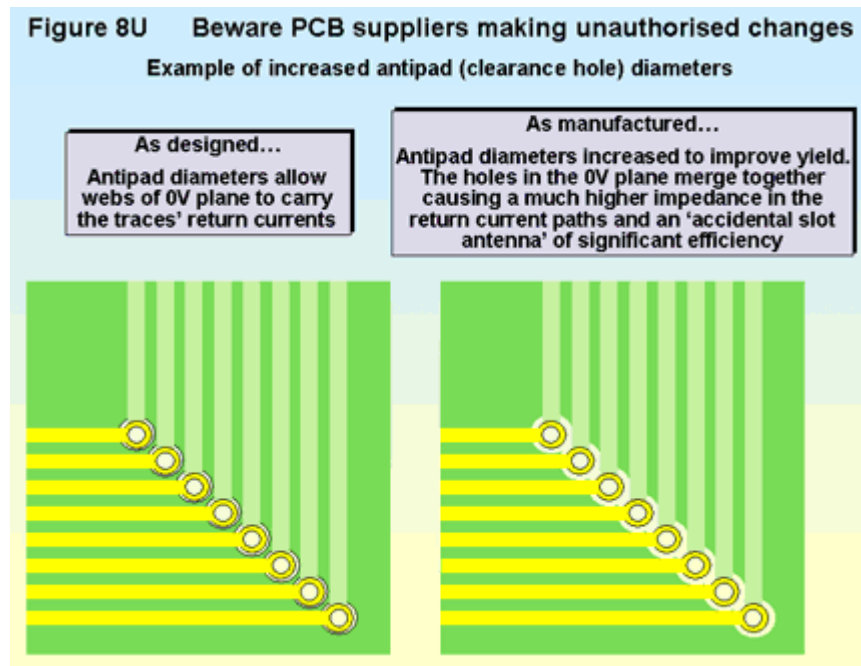
### 14 Beware board manufacturers changing layouts or stack-ups

Board manufacturers routinely adjust layouts, to make them easier to manufacture with good yields. This usually means increasing the diameter of the clearance holes (antipads) around vias and through-holes, even if this cuts into nearby traces or creates large holes in planes when the antipads merge with each other. The tendency of some manufacturing buyers to only buy PCBs from the supplier that quoted the lowest price only encourages this practice. So it is important to always check the PCB suppliers' final films after panellising [35], or check their preliminary *unlaminated* panels, to prevent the all-too-common situation depicted in Figure 8U or other unauthorised changes from causing problems.

I have seen PCBs modified without authorisation by low-cost high-volume PCB manufacturers to such a degree that they became unsafe. Of course, this was only discovered after a batch of 100,000 finished products using those unsafe PCBs had been delivered to the UK. The original designer's identification and version numbers were left unchanged on the PCB, but the rest of the layout bore little resemblance to his design! I have also seen PCBs where the merged antipads effectively split the main 0V plane, intended to be continuous over the whole PCB, into two large planes connected only by a whisker of copper no more than 0.25mm (10 thousands of an inch) wide. The EMC performance of the products built with

those batches of PCBs bore no resemblance at all to the ones that had been tested for EMC compliance, which had been made with the correct antipad diameters.

Oversized clearance holes or other unauthorised routing adjustments are more likely to occur on 'production' batches than on prototype PCBs. So if functional or EMC problems not present on the prototypes arise during serial manufacture - always suspect the PCBs. With PCBs of more than two layers, X-ray inspection might be required to see what is really going on inside a finished PCB.



Figure 8U   Beware PCB suppliers making unauthorised changes
Example of increased antipad (clearance hole) diameters

As designed…
Antipad diameters allow webs of 0V plane to carry the traces' return currents

As manufactured…
Antipad diameters increased to improve yield. The holes in the 0V plane merge together causing a much higher impedance in the return current paths and an 'accidental slot antenna' of significant efficiency

PCB manufacturers sometimes add copper pads or areas to assist with copper balancing (see sections 3.3 and 4 of [23]) or for solder 'thieving' when using wave-soldering (page G-23 of [36]). These can upset the impedance of controlled-impedance traces, or unbalance differential lines (see [6] and section 4 of [23]). When a PCB is revised, the manufacturer might add the copper pads or areas in a different way, causing different problems. Where such additional pads or areas are required, they should be applied by the PCB designer and their impact on EMC taken into account.

Every aspect of a finished PCB must be under the *complete* control of its designers - no PCB manufacturer must be allowed to change anything at all without their carefully considered approval in writing. But even if this is written in the purchasing contract for the PCBs it is best to assume it will not be followed, and insist on checking the final films or preliminary panels in every case.

15 Future-proofing the EMC design

Designs often suffer decreased EMC performance when they are no longer under the control of their original designers. This is especially a problem for 'value analysis' or 'cost reduction' projects - they should use people who are at least as knowledgeable and experienced as the original designers - but in practice usually use people with lesser skills who just delete anything they don't see the reason for.

This exact problem has cost at least one company their reputation for reliability, which took them decades to achieve.

In many companies, it does not help if the original designer writes technical reports about the EMC features of a design - because the people working on the next version or the cost-reduction either cannot understand them, or do not feel that they have the time to read them, or do not even bother to ask if such documents exist. We need to make sure that all of our good EMC work is 'future-proofed' so it is not wasted (at great cost to the company) when someone else becomes responsible for the design.

### 15.1 Marking EMC design features or critical parts on the design drawings

EMC features associated with the PCB design include such issues as multiple 0V return traces, additional 0V planes, routing constraints, PCB stack-up to achieve buried capacitance, etc. Components that are critical for EMC (but may not be for functionality) include multiple decoupling capacitors, transmission-line terminations, etc.

To help future-proof the EMC of a design it is very important to annotate all of the design features or components - that are needed for the EMC performance to be achieved - on the schematics, PCB design documents, assembly drawings, work instructions, and maybe on the parts lists (bills of material) too. In some cases all that will be needed is an adjacent 'EMC' symbol (maybe just an E), along with a footnote on each drawing page that this symbol means it is important for EMC. In other cases some instructions will be needed, for example: "These trace pairs must be routed symmetrically as striplines between two solid 0V planes along their entire length, on a single PCB layer."

At least this should warn the people who may be tempted to change those aspects of the design that they are messing with the EMC performance, so had better know what they are doing, or else expect to spend a lot of time testing and redesigning while they waste their employer's money in large amounts learning the hard way about how to do EMC design properly. If you feel that a stronger warning is required (e.g. over a subtle design aspect that took a long time to get right - add it!

### 15.2 A quality-controlled procedure for EMC design

If an organisation has an effective QC system (e.g. ISO 9000), adding an EMC Design Procedure should help ensure that designers and others (e.g. 'cost-reducers') who take over other people's projects or products do not destroy the good EMC performance they achieved. Such a design procedure will generally include requirements for identifying EMC design features and EMC-critical components as described above.

### 16 EMC-competent QA, change control, cost-reduction

See section 8 of [7] for a discussion of this important issue.

The electronic circuit designer, who should take full responsibility for ensuring the layout is good for EMC, should always review PCB layouts. The only exception to this rule is when the person doing the PCB layout is competent in understanding EMC in PCBs and applying that knowledge effectively. As discussed in [26] - good modern PCB design for SI and EMC requires a great deal of knowledge and technical expertise and the ability to understand and use computer applications such as field solvers. The level of expertise required can be equal to, or greater than, what is required to design the electronic circuits or write the software.

## 17 Compromises

It is easy to write an article like this and simply list all of the relevant good EMC design techniques - but in real-life there are a great many design compromises to be made, and this is where the circuit and PCB designers really earn their keep.

Designers are often put under cost or time pressure by managers who don't understand the technical trade-offs, and so don't understand that their actions could have the opposite effect to that which they intend and actually increase project costs and delays, as well as maybe increasing manufacturing and warranty costs. For more on this topic, please refer to [26], plus the final section of Part 2 [12].

## 18 References

[1] "Design Techniques for EMC" (in six parts), Keith Armstrong, UK EMC Journal, Jan - Dec 1999

[2] "Design Techniques for EMC" (in six parts), Keith Armstrong, http://www.compliance-club.com/keith_armstrong.asp

[3] "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 1", M K Armstrong, IEE Electronics and Communication Engineering Journal, Vol. 11 No. 4 August 1999, pages 185-194

[4] "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 2", M K Armstrong, IEE Electronics and Communication Engineering Journal, Vol. 11 No. 5 October 1999, pages 219-226

[5] "EMC and Electrical Safety Design Manuals", a set of four volumes edited by Keith Armstrong and published by York EMC Services Ltd., 2002, ISBN 1-902009-07-X, sales@yorkemc.co.uk, phone: +44 (0)1904 434 440.
Volume 1 - What is EMC? ISBN 1-902009-05-3
Volume 2 - EMC Design Techniques - Part 1 ISBN 1-902009-06-1
Volume 3 - EMC Design Techniques - Part 2 ISBN 1-902009-07-X
Volume 4 - Safety of Electrical Equipment ISBN 1-902009-08-8

[6] "Advanced PCB Design and Layout for EMC - Part 6: Transmission lines (and any traces carrying high-speed signals or noises)", Keith Armstrong, EMC & Compliance Journal, March 2005 pp 33-41, May 2005 pp 31-41 and July 2005 pp 30-41, available as one article from http://www.compliance-club.com/keith_armstrong.asp

[7] "Advanced PCB Design and Layout for EMC - Part 3: PCB-Chassis Bonding", Keith Armstrong, EMC & Compliance Journal, July 2004 pp 34-43, available from http://www.compliance-club.com/keith_armstrong.asp

[8] "Design Techniques for EMC - Part 2: Cables and Connectors", The UK EMC

Journal, April 1999, pp 7-16, available from http://www.compliance-club.com/keith_armstrong.asp [9] "The Impact of PCB Laminate Weave on the Electrical Performance of Differential Signaling at Multi-Gigabit Data Rates", Scott McMorrow and Chris Heard, DesignCon 05, January 2005, http://www.teradyne.com/prods/tcs/resource_center/whitepapers.html.

[10] "Chip Scale Packaging, Choices and Challenges", Joseph Fjeldstad, Electronic Engineering magazine, May 2000

[11] "Advanced PCB Design and Layout for EMC - Part 4: Reference Planes for 0V and Power", Keith Armstrong, EMC & Compliance Journal, July 2004, pp 34- 43, http://www.compliance-club.com/keith_armstrong.asp

[12] "Advanced PCB Design and Layout for EMC - Part 2: Segregation and interface suppression ", Keith Armstrong, EMC & Compliance Journal, May 2004, pp 32-42, http://www.compliance-club.com/keith_armstrong.asp

[13] "Minimising EMI from Heatsinks", N J Ryan, D A Stone and B Chambers, Compliance Engineering 2000 Annual Reference Guide, pp 52-56, from the archives at www.ce-mag.com or directly from: http://www.ce-mag.com/ARG/Ryan.html.

[14] "FDTD Modeling of Heatsink RF Characteristics for EMC Mitigation", N J Ryan, D A Stone and B Chambers, IEEE Transactions on EMC, Vol. 44 No. 3 August 2002, ISBN: 0018-9375, pp 458-465.

[15] "Grounding of Heatpipe/Heatspreader and Heatsink Structures for EMI Mitigation", C Wang, J L Drewniak, J L Knighten, D Wang, R Alexander and D M Hockanson, IEEE 2001 International EMC Symposium, Montreal, Canada, August 13-17, ISBN: 0-7803-6569-0, pp 916-920.

[16] "The Effect of Grounding on Radiated Emissions from Heatsinks", J F Dawson, A C Marvin, S J Porter, A Nothofer and J E Will, S Hopkins, IEEE 2001 International EMC Symposium, Montreal, Canada, August 13-17, ISBN: 0-7803-6569-0, pp 1248-1252.

[17] "Radiated Emission Analysis of a Heat Sink with Electrically-Long Noise Sources", Toshihiko Matsuura, IEEE 2001 International EMC Symposium, Montreal, Canada, August 13-17, ISBN: 0-7803-6569-0, pp 661-666.

[18] "Heatsink Grounding Effect on the Radiated Emissions of an Electronic Device", Li Rong and Zhang Lin-Chang, Proceedings of 3rd International Symposium on EMC, Beijing, China, May 21-24, 2002, pp 704-09.

[19] "Intel© Pentium© 4 Processor in the 423-Pin Package, EMI Guideline", Intel Corporation, Order Number: 249234-001, October 2000, section 2.4 on pages 6 and 7, http://developer.intel.com/design/pentium4/guides/24923401.pdf.

[20] "Design for EMI", for the Pentium© II, Intel Corporation Application Note AP-589, February 1999, Order No. 243334-002, section 2.3 on pages 6 and 7, http://www.intel.com/design/PentiumII/applnots/24333402.pdf.

[21] "Advanced PCB Design and Layout for EMC - Part 5: Decoupling, Including Buried Capacitance Technology", Keith Armstrong, EMC & Compliance Journal, July 2004, pp 34- 43, http://www.compliance-club.com/keith_armstrong.asp.

[22] "PCB Directions", Dr Eric Bogatin, Printed Circuit Design and Manufacture, October 2003, pp 36, download via http://wwwbethesignal.com. [23] "Advanced PCB Design and Layout for EMC - Part 7: Routing and Layer Stacking, Including Microvia Technology", Keith Armstrong, EMC & Compliance Journal, September 2005, pp 28-40, http://www.compliance-club.com/keith_armstrong.asp.

[24a]"Taking JTAG Testing Further", Dominic Plunkett, EPD magazine,

September 2005, pp 58-60.

[24b]"Testability beyond JTAG", Louis Y Ungar, Evaluation Engineering magazine, September 2005, pp 48-54, http://www.evaluationengineering.com

[24c] "Boundary Scan Spawns a Host of New Standards", Reg Waller, Micro Technology Europe, October 2005, pp 18-21, http://www.mtemag.com

[25] "Robust Electronic Design Reference Book, Volume I", John R Barnes, Kluwer Academic Publishers 2004, ISBN: 1-4020-7737-8.

[26] "Advanced PCB Design and Layout for EMC - Part 1: Saving Time and Cost Overall", Keith Armstrong, EMC & Compliance Journal, March 2004, pp 31-40, http://www.compliance-club.com/keith_armstrong.asp.

[27] "How Close is Close Enough", Howard Johnson, EDN Magazine, April 1998, http://www.sigcon.com/Pubs/edn/close.htm.

[28] "Photonic Bandgap Structures and Their Application to EMC Antennas", Vicente Rodriguez-Pereya, ITEM 2002, download via the 'article archives' at www.rbitem.com.

[29] "Simple PBG Structures Serve Microwave Designs", Jwo-Shiun Sun and Guan-Yu Chen, Microwaves and RF, November 2003, pp 51-57.

[30] "Electromagnetic Band Gap Material: Exotic Stuff for Not-So-Exotic EMI/EMC Applications", O M Ramanai, S Shahparnia and B Mohajer-Irvani, Interference Technology Annual EMC Guide 2004, pp 82-89, www.InterferenceTechnology.com.

[31] "Electromagnetic Interference (EMI) Reduction From Printed Circuit Boards (PCB) Using Electromagnetic Band Gap Structures", S Shahparnia and O M Ramanai, IEEE Transactions on EMC, Vol. 46, N0. 4, November 2004, ISBN: 0018-9375/04, pp 580-587, downloadable from http://www.glue.umd.edu/~oramahi/EIR-2004.pdf.

[32] "Towards New Paradigms in Noise Mitigation In and From Printed Circuit Baords Using Metallo-Dielectric Electromagnetic Band Gap Structures", O M Ramanai, S Shahparnia and B Mohajer-Irvani, EMC-Europe 2004 (6th International Symposium on EMC), Eindho