Advanced PCB design and layout for EMC. Part 3 – PCB-to-chassis bonding

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

This is the third 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 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

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. This series will not repeat the basic design information in these articles – it will build upon it.

Like the above articles, this series will not spend much time analysing why these techniques work, they will focus on describing their practical application 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.

Table of Contents, for this Part of the series

1. Introduction to PCB-to-chassis bonding

Electronic equipment is typically constructed using PCBs fixed to a conductive chassis. Most PCBs will have an unbroken 0V plane on one layer (and if they don’t, they probably should, see [1] – [5]). There is usually at least one electrical bond, usually at a fixing, between the 0V plane and the chassis, usually near to the external cable connectors if nowhere else, see Figure 3A.

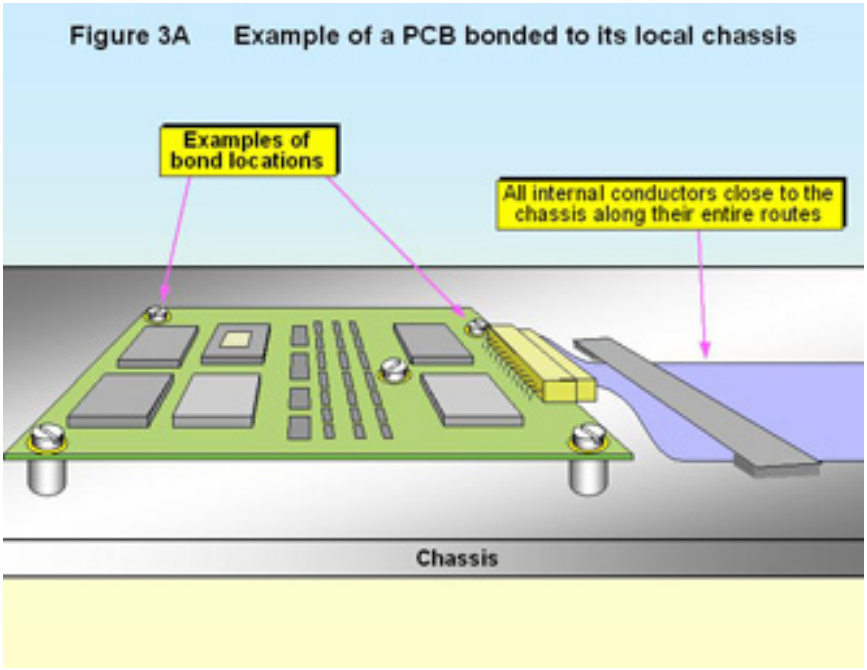

Sometimes galvanic isolation is required between a PCB and its chassis, in which case the bonds are made via suitably rated capacitors. This article discusses the EMC design issues associated with bonding PCB 0V planes to their local chassis.

1.1 What do we mean by 'chassis'?

The word 'chassis' in this article refers to a metal support structure for a PCB, e.g. the wall or floor of a metal enclosure. A shielded metal enclosure makes a very good chassis for a PCB, but a chassis need not be a shield, or even an enclosure – it could just be a piece of metal. For instance, some PCBs with power semiconductors are supported entirely by their heatsink, in which case the heatsink would be considered the PCB's chassis. In the personal computer (PC) industry, some companies (such as Intel) recommend bonding a motherboard to a sheet metal chassis, which they sometimes call a 'basepan' (or even a 'diaper', for reasons best known to themselves).

Where plastic housings, enclosures or other structures are used to support a PCB they can be used as a 'chassis' in the context of this article if they are made conductive (e.g. by metallisation, conductive painting, or using a conductive filler).

Bonding the PCB 0V plane to its local chassis generally has benefits for EMC; with cast or sheet metal chassis providing the best benefits. Some manufacturers use plain plastic housings and add sheets of metallised cardboard or PVC, or conductively paint or metallise some of their housing's internal surfaces, as needed, to help pass EMC tests. It is difficult to call a sheet of metallised cardboard or similar a chassis, when it seems more like a shield of some sort, but unless it wraps all around the PCB it is best to consider it as if it is a chassis.

When designing an entire item of equipment, the chassis we are discussing here forms part of what is sometimes called the RF Reference Plane – created by bonding all of the different metal (or conductive) parts together using connections that have a very low impedance at the highest frequency of concern. A totally shielded enclosure makes the best RF reference plane, but shielding is not essential.

1.2 What do we mean by 'bonding'?

For the purposes of this article, the type of bonding that concerns us is 'RF bonding' – a low impedance connection at the frequencies of concern for EMC. The word 'low' here generally means

less than 1?, although how low it needs to be will depend upon the particular application. For EMC purposes, we often do not care whether a PCB-to-chassis bond provides a low impedance at frequencies below 10kHz, or even at DC, as long as it has a low impedance over the frequency range of concern for EMC.

The lowest impedance bonds at RF are achieved by pressing one metal surface against another. Pointed metal contacts tend to have higher impedance at RF than is achieved by pressing one area of metal against another, so EMC designers generally try to avoid using point contact bonds. It is usually recommended that each bonding location employs a PCB pad of at least 3mm diameter, not covered by solder resist, pressed against a metal spacer which in turn presses against the chassis. It is normal to ensure a good bond between these pads on the outer layers of the PCB and an inner 0V plane by using a ring of via holes, as shown in Figure 3B.

Figure 3B Some examples of capacitor RF impedances

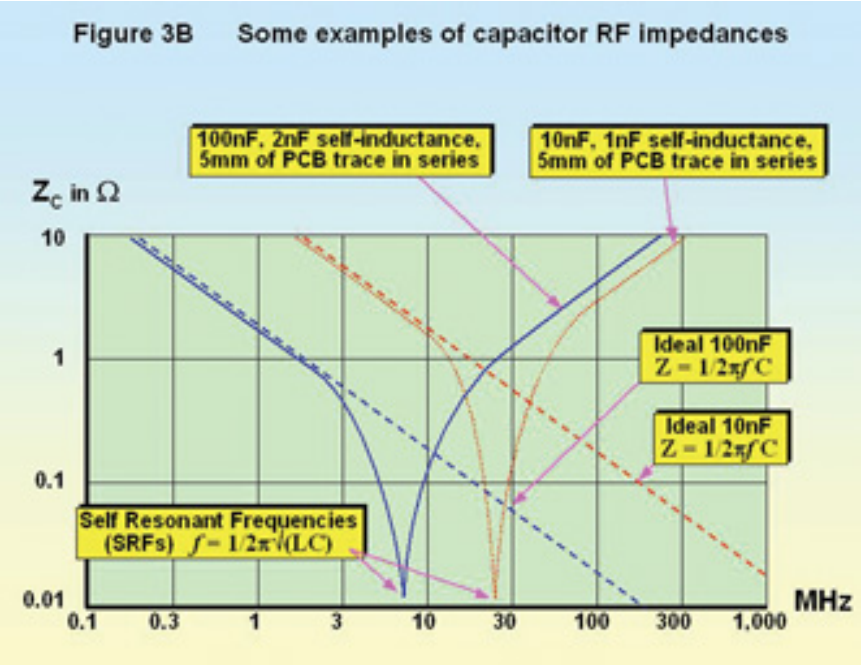

Good surface conductivity is required for all the metal-to-metal contact areas, and they should not corrode over the life of the product – so it is best to plate each part with the same metal, such as tin. Fixings such as screws or bolts should not be relied upon as bonds, they should simply ensure that the contact areas of the bonding components are pressed together firmly. Captive metal spacers with threaded holes are often used, but the means by which the spacer is held captive in the chassis must achieve a cold-weld all around the spacer, add no impedance, and not corrode.

A variety of components are available to help achieve 0V plane-chassis bonding at low cost, such as the conductive snap-in PCB spacers shown in Figure 3P. Some other components and bonding techniques are mentioned later.

The minimum requirements are usually for RF bonds between the 0V plane and the chassis or frame near every input or output connector, and at least one RF bond close to the highest-speed or 'noisiest' (most emissive) devices.

Where a shielded cable is connected to a PCB, its shield should usually be RF bonded (sometimes called 360° bonded) to the chassis or frame of the unit at the place where it enters (see Parts 2 and 4 of [2]), which requires a shielded connector (or a shielded shell on a connector). To ease assembly and reduce assembly time we prefer to use PCB-mounted connectors that locate with cutouts in connector panels.

360° bonding can easily be achieved for PCB-mounted connectors by sandwiching a conductive gasket between their metal bodies or shells and their connector panel. If their metal bodies or shells are also soldered at multiple points to the 0V plane (usually at the lugs that provide their mechanical support) — then 0V plane-chassis bonding is easily achieved too. This technique can sometimes be used without any mechanical fixings between the connector panel and the connectors, further reducing assembly time whilst helping to achieve good EMC performance. An example of such an assembly is shown in Figure 3Q.

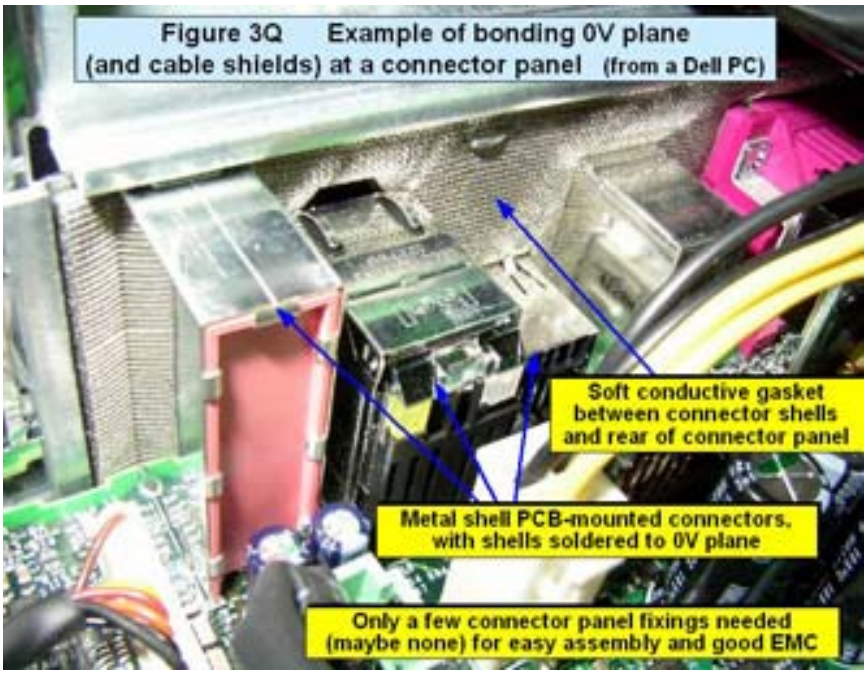

Even if the cable is not a shielded type, it is a good idea to use shielded PCB-mounted connectors as described here, to improve 0V plane — chassis bonding near the input and output connectors. Soft conductive gasket material is usually used, die-cut to suit the arrangement of the connectors (as in Figure 3Q). The typical gasket material consists of a plastic foam core covered with metallised fabric, but types are now becoming available that use a conductive foam to give improved bonding (so-called "Z-axis conductive").

Spring finger gaskets are available for individual D-types, and custom spring-finger gaskets can also be created. Examples of custom-designed spring finger gaskets can be found at the expansion card slots of most modern PCs. Unfortunately, some types of PCs do not use a stiff enough connector panel in the expansion card area, and the combined pressure on all the spring fingers causes the panel to bow outwards in the middle, opening up large gaps and reducing the performance of the PC case's shielding whilst also degrading the 0V-chassis bonding for the expansion cards.

Because for EMC we often don't care whether a low bond impedance is achieved all the way down to DC, we sometimes choose to make our bonds through capacitors. When galvanic isolation is required between a PCB and its chassis, capacitive bonding is the only solution (using suitably rated, and maybe safety-approved capacitors). Galvanically isolated medical circuits may have to meet very stringent leakage current specifications, and this could place an upper limit on the total value of capacitance that may be used between a PCB and its local chassis. In this case some of the techniques described later might not be able to be used.

An important consideration with the use of bonding capacitors is that, in conjunction with the inductive impedance that inevitably appears in series with them (e.g. due to traces and via holes), they self-resonate and only provide low impedances over a limited range of frequencies, as Figure 3C shows.

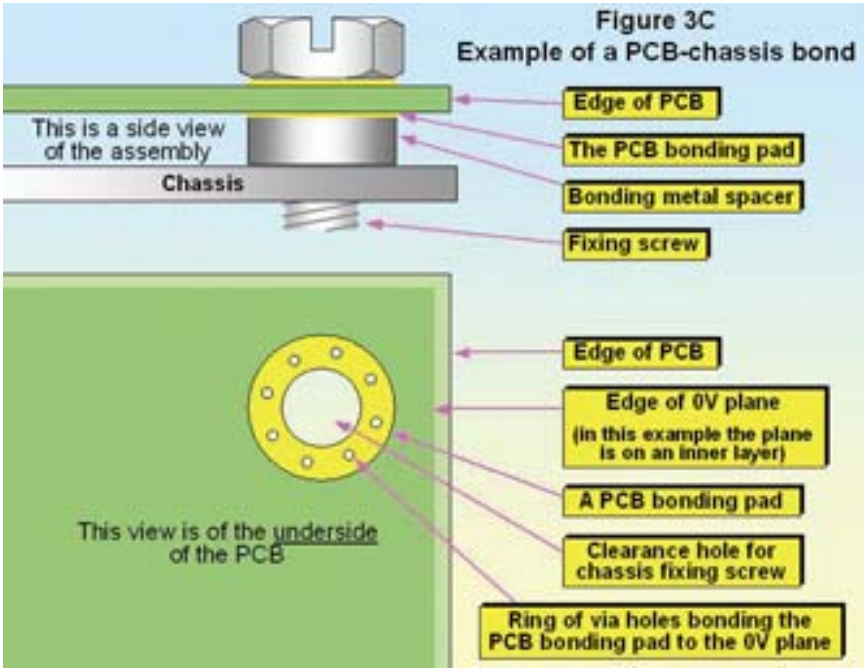

Above its self-resonant frequency (SRF), the impedance of a capacitor increases with frequency. In fact, this slope is the overall inductive impedance of the capacitor and its pads, traces and via hole(s). All capacitors inevitable have some internal inductance, typically between 1 and 2 nanoHenries (nH) for small multi-layer ceramics rated up to 100V. The pads, traces and via holes that connect the capacitors to the PCB accumulate inductance at the rate of about 1nH per millimetre of their length (as do leads, if using leaded capacitors instead of surface mounted). At 1GHz each additional nH of inductance (L) adds 6.3Ω additional inductive impedance to a PCB-chassis bond ($X_L = 2\pi f L$).

1.3 Hybrid bonding

This technique uses a variety of types of bonds between the PCB's 0V plane and the chassis…

- direct (as shown in Figure 3B);
- via a capacitor (see Figure 3C);
- via a resistor, for damping structural resonances (see later).

Each application is different, and there are sometimes good technical reasons for using hybrid bonding despite the fact that capacitive bonds are generally less than ideal for broadband RF bonding (see Figure 3C). It is often difficult to know in advance of EMC testing what type of bond will give the best overall EMC performance, for each location, so prototype PCBs can benefit from the use of a multi-purpose bonding pad pattern, like that shown in Figure 3D.

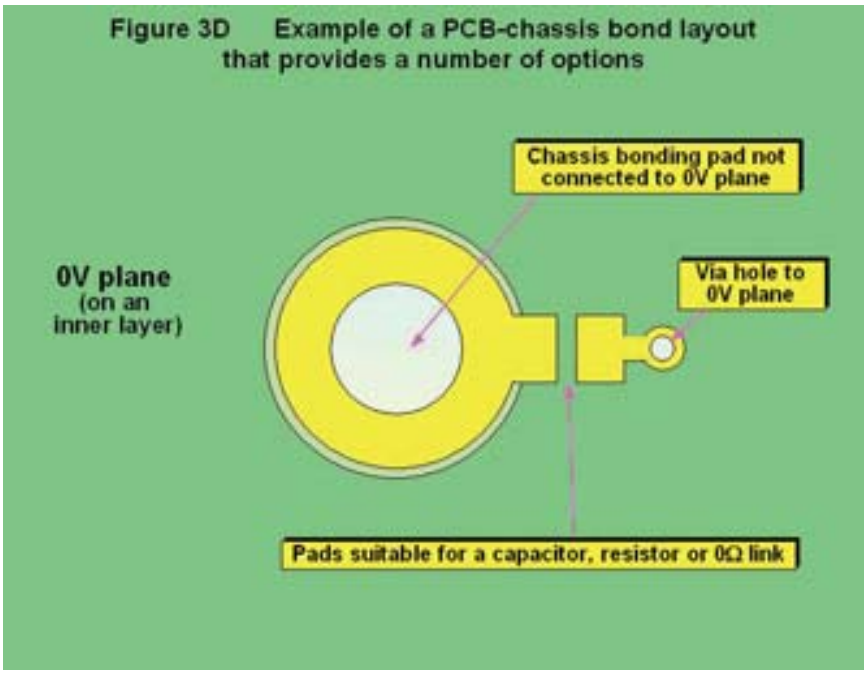

Experimental EMC tests (e.g. so-called 'pre-compliance' tests) are recommended to discover which is the best configuration of DC, capacitive, and resistive (see later) bonds, and which are the best values to use.

1.4 'Ground loops' and tradition

A tradition has grown up amongst designers of DC and low-frequency (LF) instrumentation and amplification (e.g. audio) of avoiding 'ground loops' by only having a single bond between the 0V reference for their circuit and the chassis (which is itself often connected to the protective earthing safety conductor in the AC power lead). But at the same time, designers of RF amplifiers and the like established a tradition of multi-point bonding between their 0V planes and their chassis, with the spacing between the bonds being very small compared with the wavelength at their highest frequency of concern.

The DC/LF camp have in the past employed what they called single-point grounding (or 'star' grounding) almost as if it was an article of faith, and seemingly never stopped to ask why their RF counterparts could design perfectly good DC/LF circuits using multipoint bonding which creates numerous ground loops.

These days, many DC/LF circuit designers are learning that to achieve good RF immunity and pass their EMC Directive compliance tests they need to employ multipoint bonding after all. Some of them are learning that giving their PCBs a good 'solid' unbroken 0V plane and using that for all the 'grounding', and bonding that 0V plane to their chassis at multiple locations, generally is a big help in meeting immunity standards and also improves their circuit's functional performance.

In some professional audio equipment, changing from traditional single-point grounding to unbroken 0V planes with multiple PCB-chassis bonds has reduced audio-band noise levels by 10dB below what had for decades been considered the minimum physically achievable (see [6] – [8]). A similar multi-point bonding approach can work very well in large audio installations too, see [6] - [11].

But we still find DC/LF circuit designers using 'hybrid' bonding to try to preserve their ancient tradition of single-point chassis bonding whilst achieving multi-point bonding via capacitors for EMC. Unfortunately, the result can be a less than optimal design for signal-to-noise, cost and/or EMC. For cost-effective engineering, traditional design techniques should never be accepted without question.

2. Why bond PCB 0V planes to chassis anyway?

2.1 Reduced transfer impedance

As signal return currents flow through the inevitable impedances in a PCB's 0V structure, they cause different parts of the PCB to experience different voltages from each other. These voltages are common-mode (CM), and CM voltages and currents are the major cause of EMC problems. One of the many EMC and signal integrity benefits of an unbroken 0V plane in a PCB is that it has lower impedance at RF, so CM voltage differences at RF are reduced and emissions of electric fields are reduced. An impedance that converts a wanted current (e.g. signal return) into unwanted CM voltage (or a wanted voltage into an unwanted CM current) is called a 'transfer impedance' – an important concept in EMC engineering.

0V planes are especially valuable where off-board conductors (e.g. cables) are attached to different parts of the PCB. The length of these conductors often makes them very efficient 'accidental antennas' and reducing the CM voltage difference between the different parts of the PCB they are attached to is an excellent way to reduce their emissions.

When RF CM currents are injected into a PCB by coupling from the external electromagnetic environment (with off-board conductors being major sources of injection) – having a lower transfer impedance in the PCB's 0V structure means that the resulting signal voltage noise is lower and much less likely to interfere with circuit operation.

The transfer impedance of a well-designed 0V plane (see the next part of this series) is several orders of magnitude less than the transfer impedance of a PCB trace or a wire. To take advantage of the low transfer impedance of the plane, all traces, wires or cables that exit a plane's area must be RF bonded to the plane, either by '360° bonding' of their shields or by filters with a capacitor connected to the 0V plane.

Bonding the 0V plane to the chassis at multiple points helps reduce the transfer impedance even more than can be achieved with a plane alone, and so helps improve EMC performance. Metal chassis have much lower *resistance* than can be achieved in a layer of copper in a PCB, so they also help reduce the transfer impedance at much lower frequencies than the plane can achieve on its own, even down to DC.

Because of the above considerations, it is normal EMC practice to bond 0V planes to chassis at least at each of the corners of the PCB's 0V plane and also near to the entry/exit point of each off-board conductor. Figure 3E shows one of the benefits of reducing a PCB's transfer impedance by bonding its 0V plane to its chassis.

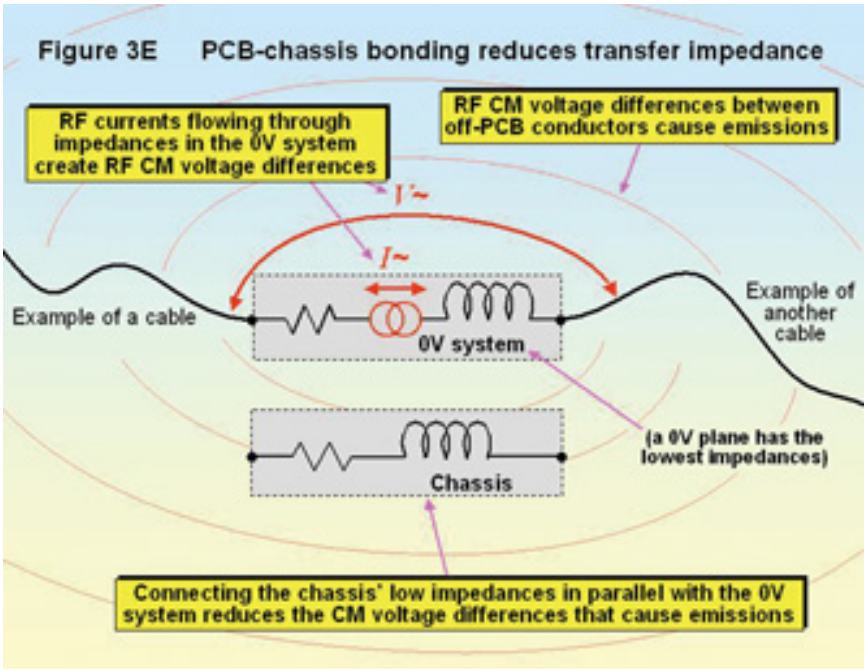

2.2 Better control of 'fringing fields'

The operation of the circuit on the PCB causes RF voltages that differ from the voltages on nearby conductive items, such as the chassis, and these RF voltage differences give rise to emissions, as shown by Figure 3F. Where the 0V system is a plane, all the emissions are from its edges, sometimes called fringing fields.

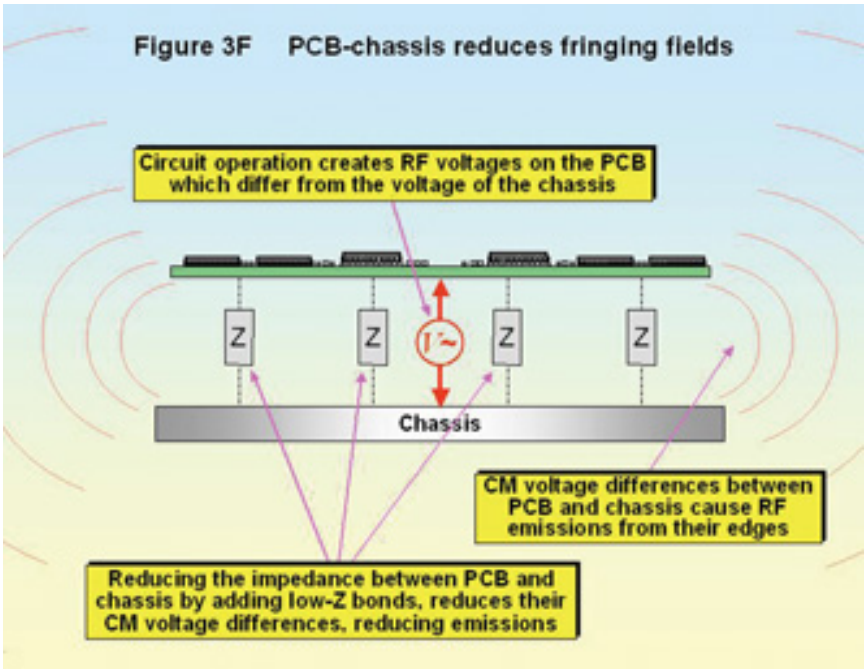

There are a number of sources of RF voltage on a PCB. Signal return currents flowing in the 0V system are one source of RF voltages (see above and Figure 3E). RF currents flowing in the impedances of the circuit's external power supply (e.g. the AC supply) are yet another source of RF voltages. The signals in the circuit may have some RF content in their spectrum (digital signals always do). The RF voltages on transistors and ICs themselves (especially those fitted with heatsinks) are becoming more significant as operating frequencies increase.

So another reason for bonding a PCB's 0V plane to the chassis is to try to equalise the RF voltages

between them, to reduce the emissions from their fringing fields. A reciprocal argument can be made to show that improving the RF bonding between PCB and its nearby chassis helps improve immunity.

Reducing the effects of CM emissions is the reason why PCBs often have at least one bond between their 0V planes and local chassis, located close to the source of the PCB's highest-frequency noise emissions. This is often a clock oscillator or clock buffer, or a VLSI integrated circuit (IC) such as a powerful microprocessor, gate array, or digital signal processor.

Basic PCB-chassis bonding guidelines thus require PCB-chassis bonds at least at each corner of the 0V plane, at least one near to each cable port, and at least one near to each high-speed device.

Now that we understand the basics of PCB-chassis bonding, we can move to discussing the design issues that are arising due to the continual increase in the highest frequency of concern, due to the technology issues discussed in Part 1 of this series [12].

3. The 'highest frequency of concern'

Throughout this series of articles, an important issue is the 'highest frequency of concern' because this governs a great many of EMC design issues. The choice of the highest frequency is up to the head of the equipment's design team, who might choose it on the basis of…

- the minimum regulatory requirements it is hoped to 'get away with' in the countries being marketed to;
- what could cause annoyance to customers and/or poor quality performance or unreliability in real life;
- the highest frequency the technology is capable of emitting or being susceptible to;
- future developments in regulations, test standards, the electromagnetic environment or the devices used (e.g. die shrinks) to avoid having to redesign the equipment too soon.

It is worth pointing out here that although the EMC Directive's notified emissions standards may be limited in the upper frequency they test to, compliance with the EMC Directive's essential Protection Requirements mean that emissions that could cause a nuisance to others are not permitted at any frequency up to 400GHz.

It is a similar issue for immunity. Even though the most relevant notified immunity standard is limited in its range of disturbances and their levels and frequencies, compliance with the EMC Protection Requirements requires that the equipment be sufficiently immune to its real electromagnetic environment up to 400GHz. (The 400GHz limit of the EMC Directive comes about because signals above that frequency are considered to be infra-red, then visible light, ultra-violet, X-rays, etc. as the frequency increases still further. Although these are all electromagnetic phenomena, they are not covered by the EMC Directive.)

So, simply testing to the current versions of the most relevant EN or IEC EMC emissions and immunity standards is not enough to ensure EMC Directive compliance, and not enough to ensure happy customers either. Where the electronic technology used could cause emissions above the highest frequency covered by the standard it is recommended that quick tests (at least) be done to see if the emissions are excessive and need suppressing to prevent interference.

In the case of immunity, the issue is whether there could be significant sources of ambient noise present in the equipment's operational environment, at frequencies higher than those covered by the most relevant immunity standard. If there are, it is recommended that quick tests (at least) be done to see if the susceptibility to these frequencies is significant, and whether the equipment needs modifying to prevent it from being interfered with in real-life operation.

4. Benefits of closer spacing between a PCB and its chassis

Ideally, we would like zero coupling between our digital, analogue and switch-mode circuits and their external electromagnetic environment, because then we would have no emissions and perfect immunity. We can never achieve zero coupling, but a great many EMC design techniques are associated with reducing it.

The fact that circuit operation causes voltage differences between PCB and chassis, causing fringing fields, was mentioned earlier. Some EMC engineers analyse the structure consisting of the PCB's 0V plane and its nearby chassis using transmission-line methods. They find that the closer they are together; the lower is their coupling to the external electromagnetic environment [13]. Figure 3G shows this graphically.

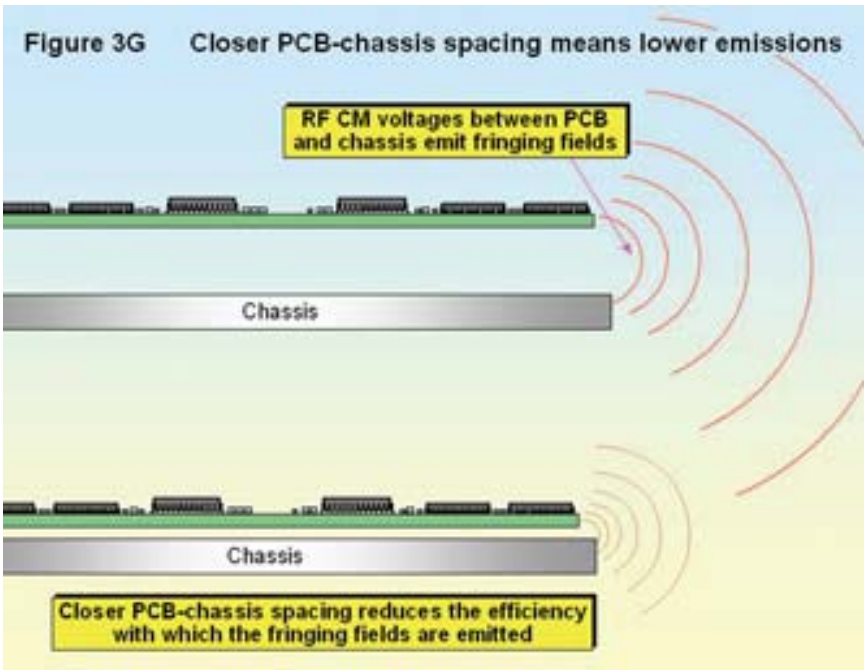

Where the PCB-chassis spacing is greater than half a wavelength, reducing the spacing could possibly make the coupling between the transmission-line structure and the external environment more efficient, increasing emissions and worsening immunity. One solution to this problem is to ensure that the spacing is reduced to much less than half a wavelength.

The closer the PCB is to its chassis, the lower is the impedance of the bonds between them. Inductance scales linearly with length, so halving the PCB's spacing from the chassis will half the length (and hence halve the partial inductance) of the bonds. Reduced inductance in PCB-chassis bonds has benefits for the overall transfer impedance, and helps return CM currents more quickly to the PCB, improving emissions and immunity in two ways.

Finally, closer spacing will increase the resonant frequencies of the cavities between the PCB and the chassis (see later) – not by much, but it all helps.

5. Controlling resonances in the PCB-chassis cavity

5.1 Why and how the cavity resonates

As the 'highest frequency of concern' keeps getting higher (see above, and [12]) many of the physical structures in electronic equipment become resonant, creating new problems for both signal integrity and EMC. The resonance problem is caused when the wavelengths (λ) of the frequencies

become comparable with the physical dimensions of the structure. This is why EMC design engineers are often as concerned with millimetres as they are with MHz.

The black curve in Figure 3L below is a measurement of an example PCB-chassis structure. At resonance, a structure couples very well indeed with its external electromagnetic environment, becoming a very efficient 'accidental antenna'. This is not what we want because it increases emissions and worsens immunity. At resonance, emissions from physical structures can be 20dB (or more) higher than at nearby non-resonant frequencies (and by a reciprocal argument, their immunity can be 20dB or more worse).

Also, at resonance, coupling between circuits on the PCB and elsewhere in the equipment is increased. This is often called crosstalk, and can cause problems for signal-to-noise and signal integrity.

Part 2 of this series [14] discussed the structural resonances of PCB-mounted shielding cans, and its Figures 2H and 2J shows the effects of their structural (cavity) resonances on crosstalk and emissions respectively. Exactly the same issues arise due to the structural resonances of the cavities formed between the PCB and its chassis [15] [16], the topic covered by this part of the series.

If there were only PCB-chassis bonds at the four corners of a rectangular PCB, then there would be just the one cavity to analyse, but often there are more bonds and so more (and smaller) cavities. In most cases, what we are mostly concerned with is the first (lowest) resonant frequency the PCB-chassis bonded structure, which is associated with the longest diagonal of its cavities.

A crude analysis of the likely resonant frequency for a structure that has PCB-chassis bonds only at its corners, and where PCB-chassis spacing is small (as is usually the situation) is easy enough, uses the formula…

$$f_{lowest} = 150 \sqrt{\{(L^2 + W^2)-1\}} \text{ (in GHz, when L and W are in mm)}$$

…where L and W are the PCB 0V plane's length and width respectively. For example, a structure consisting of a PCB 0V plane 160 x 120mm spaced 5mm above a chassis, and bonded to it at its four corners, would have its lowest resonant frequency at around 0.75GHz.

The above formula is only a crude approximation because it is based on the resonances inside a totally metal-sided cavity, whereas our structure is simply two metal plates with open edges connected together at a few locations with non-zero impedances. Such a structure requires a three-dimensional field solver to simulate what is really going on at any frequency of interest.

Where a cavity has PCB-chassis bonds on only one side (e.g. at only two corners), or on two adjacent sides (e.g. at only three corners), a crude approximation for its lowest resonant frequency is…

$$f_{lowest} = 75 \sqrt{\{(L^2 + W^2)-1\}} \text{ (in GHz, when L and W are in mm)}$$

5.2 Wavelength rules

To help avoid structural resonances, EMC engineers often employ general guidelines ('rules of thumb') for physical dimensions based on λ/10. The dimensions concerned could be spacings between RF bonds, or a number of other issues, but the idea is that as long as they are less than λ/10 resonance cannot occur. λ/20 or λ/100 'rules' are sometimes used instead, for better EMC performance.

Some designers know the rise and fall times of their signals, but not their associated highest frequency of concern. In this case the dimension given by the λ/10 guide is equivalent to 100mm times the real risetime in ns (or the real falltime, if it is shorter). It is very important to use the real rise/falltimes achieved at the output pins of the ICs, not their datasheet figures, since die-shrinks usually result in ICs that have actual rise and fall times that are much faster than the maximum values in their datasheets [12]. If you do not know the real rise and falltimes for saturating logic like CMOS and TTL, divide their datasheet figures by 10 to be on the safe side. For non-saturating logic (such as ECL), divide the datasheet values by 4. But it is much better to measure them than to use these estimates.

When measuring rise and fall times: use an oscilloscope and probes that have much faster rise and fall times than the measured signals (ideally more than twice as fast); and use correct high-frequency probing techniques as described in the oscilloscope manufacturer's application notes. Always measure at the actual pins of the IC or transistor that is the source of the signal.

5.3 Increasing the number of bonds to increase resonant frequencies

To help prevent resonances from occurring in the PCB-chassis structure the general guidance is to use PCB-chassis bonds that are no further than λ/10 from any other bond (preferably closer), where the λ used corresponds to the highest frequency of concern. In air $\lambda = c/f$, where c = the speed of light (3.10^8 m/s), and when f is in Hz, λ is in metres.

Mark Montrose [17] recommends using bonds that are no more than λ/20 apart from each other, at the highest frequency of concern, as shown in Figure 3H. He uses an argument based on the efficiency of dipole antennas, rather than cavity resonances. λ/20 will give better performance than ?/10, but will quadruple the number of bonds for a given λ.

Figure 3H The λ/20 PCB-chassis bonding rule
(a copy of Figure 2.27 from Chapter 2 of reference [11])

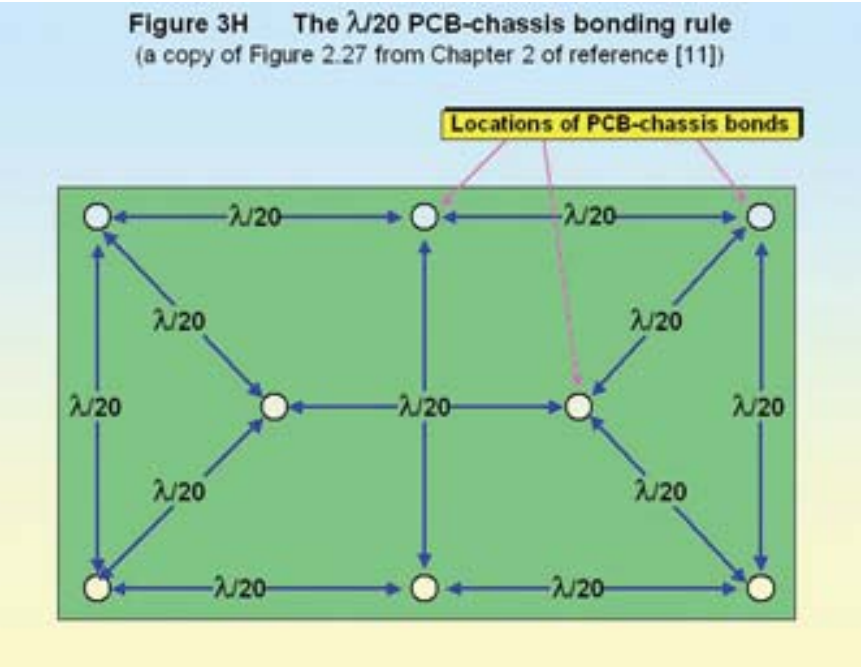

Once a PCB has sufficient fixings to control its movement during shock and vibration, adding extra PCB-chassis bonds that required fixings would add to assembly time. However, a number of techniques exist that allow PCB-chassis bonds to be added without increasing assembly times. One of them is shown in Figure 3J, a (slightly out of focus) photograph of a PCB-chassis bonding location between a PC motherboard and its chassis (basepan).

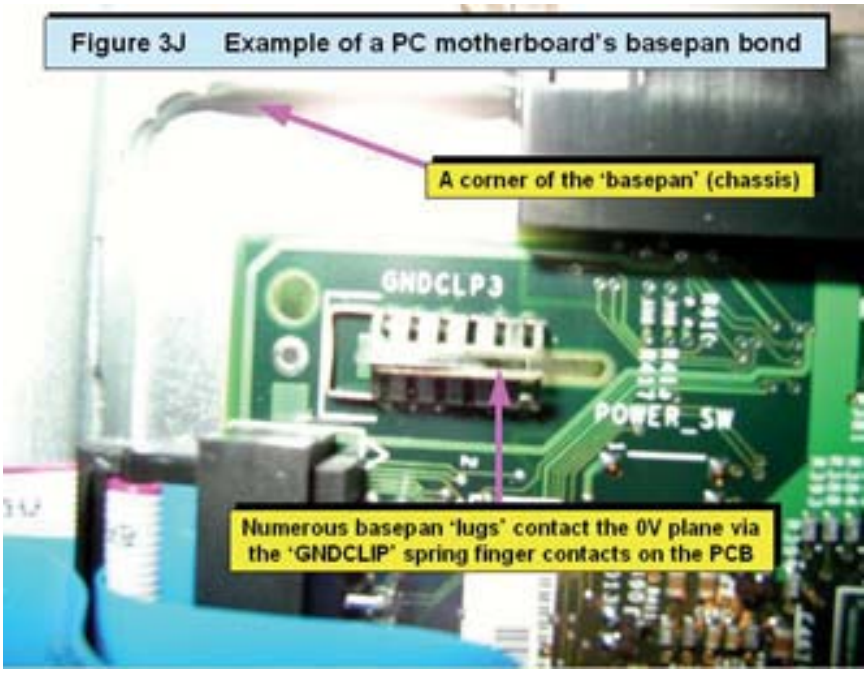

This technique uses a sheet metal chassis that is 'semi-punched' to create a number of vertical protrusions ('lugs') that align with slots cut out of the PCB. The PCB is pushed down over the lugs and metal spring finger clips make contact with the lugs as they protrude through the slots. The example shown in Figure 3J is especially notable because the lugs were shaped so that once pressed down, the PCB was then slid sideways into a feature in the lugs to lock it into position at each lug. Only a single screw-fixing was required to prevent the PCB from sliding sideways and becoming disengaged from the lugs, so assembling this motherboard to its basepan was a very quick and easy operation. It is very rewarding when good EMC design techniques can also be combined with savings in the overall cost of manufacture, as they often can when the natural urge to spend the least cost on the PCB is resisted (see [12]).

Another technique is to provide PCB pads on the bottom of the PCB so that spring fingers or conductive 'bumps' can make contact between the pad and the chassis, as shown in Figure 3K. A wide variety of alternatives are available for use as conductive bumps, usually small pieces of compressible conductive gasket.

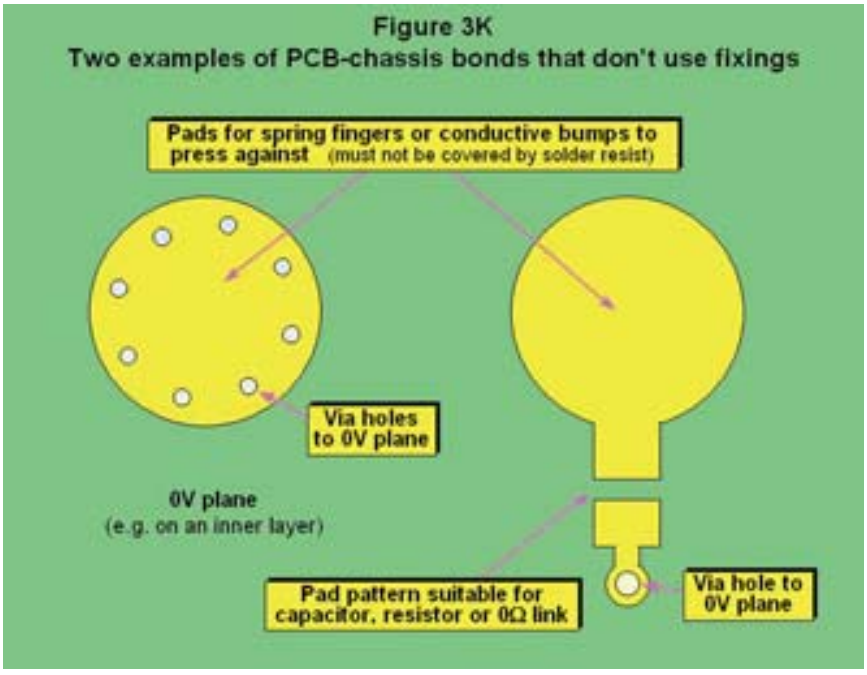

The spring fingers or conductive gaskets could either be fitted to the PCB, or to the chassis. Spring fingers are available as surface-mounted components that can be soldered to the PCB pads. Many types of spring finger and conductive gasket are also available with self-adhesive backing, making their assembly to the chassis or PCB relatively easy. Some shielding and other manufacturers make components designed especially for bonding PCB 0V planes to chassis, but anything based on a coiled metal spring should be avoided unless the highest frequency of concern is not very high (say, below 50MHz, depending on the application). For example, Kitagawa makes some very small surface-mounted spring finger components, initially for use in cellphones.

In high volume serial manufacture, and where PCB-chassis spacings are 2.5mm or less, it may be worthwhile considering the robotic application of blobs of form-in-place (FIP) conductive gasket. Types of FIP gasket material are now available that foam up after application, to create larger and/or softer gasket bumps.

One of the problems with PCB-chassis bonds is that PCBs are often changed late in a design and development project – when problems are found during functional or compliance testing – and this can mean that the locations of some of the bonds need to be changed too. Where the PCB bonds rely on the chassis metalwork (e.g. captive metal spacers, semi-punched lugs, etc.) the changes in the PCB layout have a knock-on effect that adds to costs and timescales. But bonds that use spring fingers or conductive bumps applied to the PCB don't cause the same problems.

Mark Montrose describes (in [17]) using conductive polymer parts for PCB-chassis bonding, retained between the PCB and the chassis by punched holes in a plastic sheet. When the PCB changes, the punching pattern for the plastic sheet is changed to suit.

Good PCB-chassis bonding performance relies on good (high) surface conductivity being achieved over the life of the equipment. This is especially important where spring finger or conductive bump techniques are used, because their contact pressures are much less than at a screwed fixing, so thin films of oxide or corrosion will have a much worse effect. Where metal parts are concerned, it is best to plate each part with the same highly-conductive metal, such as tin or gold. Conductive gasket materials are harder to choose to prevent corrosion, but any professional EMC gasket manufacturer should be able to provide highly detailed application guides and test results showing which gasket materials should be used with which metals.

5.4 What if we can't use enough bonds?

Ideally, we would reduce the spacing of the PCB-chassis bonds until the lowest resonance frequency was higher than the highest frequency of concern. Then – as long as the highest frequency of concern does not increase – we should have design for which the PCB-chassis bonding structure is not likely to be the cause of significant EMC problems despite design changes, alternative components, and manufacturing tolerances.

But to achieve a lowest resonant frequency of 3.0GHz (for example) – so that no resonant effects should cause problems below about 2.5GHz – following the $\lambda/10$ rule would require PCB-chassis bonds no further apart than 10mm (5mm if using the $\lambda/20$ rule).

Such very frequent bond spacing would increase PCB layout difficulties, and the total force required to compress all the spring fingers or conductive bumps might make the PCB and/or the chassis bend, unless extra fixings or stiffening components are used.

So to control PCB-chassis cavity resonances above about 500MHz, we may need to employ other techniques, described below.

5.5 Spreading the resonances more widely to reduce peak amplitude

It might prove impractical to add sufficient PCB-chassis bonds to make the lowest resonant frequency higher than the highest frequency of concern. In this situation there may be an advantage in avoiding too regular an arrangement of bond locations. An irregular bonding pattern has two advantages…

- The resonances in the length and width directions of different cavities are not the same
- The resonant frequencies of the multiple smaller cavities created by the multiple bonds do not coincide

Breaking up the resonances in this way should reduce the worst-case peak amplitudes of the resonances, trading them for broader resonant regions (lower Q values). If using this technique to reduce emissions or improve immunity at a particular frequency, be aware that it might worsen EMC performance at other frequencies.

This approach should be reasonably 'robust' as far as design changes and component variations are concerned, but EMC-competent QA and change control are always recommended (see later).

5.6 Designing resonances to miss problem frequencies

Each cavity resonance covers a range of frequencies (see the black line in Figure 3L below for a typical example), and generally only cause problems when the frequencies emitted by the PCB's circuit fall within this range. Most circuits have their highest emissions at their clocks' fundamentals and their harmonics, and careful design of the PCB-chassis bonds may be able to ensure that these do not fall into any resonant frequency ranges.

Figure 3L Example of damping, using resistive bonds
(a copy of Figure 10 from reference [9])

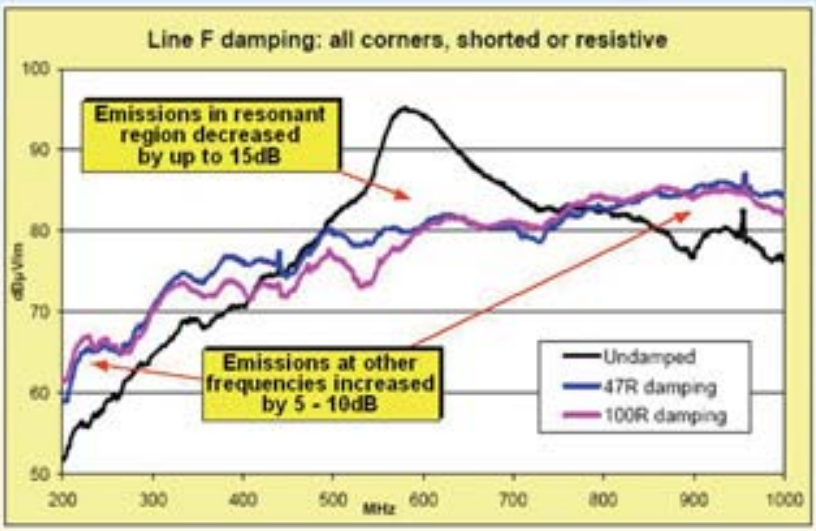

The use of high-frequency clocks can make this technique easier to apply, if their harmonic spacing is great enough that they 'bracket' the resonant regions rather than fall into any of them. It can also be an advantage to increase the Q of each resonance, by making all of the resonant cavities created by the multiple PCB-chassis bonds identical in shape and size wherever practical. This is the exact opposite effect to the previous technique. A higher Q means the frequency range over which emissions are increased is narrower, and easier to avoid with clock harmonics.

Simple experiments with plain copper sheets may not be sufficiently representative of real life, because the PCB's devices, trace routings, perforations and gaps in the 0V plane – plus any apertures, shape changes and the proximity of cables and other components – will modify the

resonant frequencies. But we do not want to wait until we can EMC test a structure that is reasonably representative of the final design – because by this time changes to the design will be much more costly than if they were done earlier (see [12], especially Figure 1A).

This clever technique can help reduce unit manufacturing costs, but investing in a three-dimensional field solver (and in the training in how to use it properly) is almost a necessity if is to be used cost-effectively. The author's experience with such computer-aided engineering tools, over many years, has been that they seem to take a very long time to learn how to use effectively, but then they suddenly allow design iterations that would have taken days or weeks to be done in minutes or hours. A dozen design iterations in a day is not unreasonable at an early stage in the project, and at that point the true value of the investment becomes apparent to all. Of course, the field solver will need to be provided with a reasonably accurate description of the final design, to give useful predictions that reduce project timescale risk. The garbage in garbage out rule always applies.

But this is not a design technique that is very 'robust' as far as design changes and component variations is concerned. Even quite small changes, such as altering the clock frequency to use the latest ICs, could completely alter the EMC performance of the equipment, so EMC-competent QA and change control is very important (see later).

5.7 Being clever with capacitors

The values of the capacitors used in capacitive or hybrid bonding might affect the EMC performance. Unlike the 'designing resonances to miss problem frequencies' technique above, this approach lends itself to last-minute modifications – as long as the PCB has already been laid out using appropriate pad patterns at the bonding locations (see Figures 3D, 3K and 3M).

Usually, the equipment is subjected to pre-compliance EMC tests and the types and values of capacitors (or zero-Ω links, or resistors) fitted at each bond are varied until the optimum is found. An assembly bench equipped with appropriate soldering/desoldering tools and a complete set of all likely components needs to be provided just outside the test chamber. The number of possible alternatives is huge, so most people stop iterating when further improvement is proving too time-consuming.

If zero-Ω links provide the best EMC, on a future revision of the PCB they could be replaced by direct bonds between the chassis and the 0V plane (e.g. as shown in Figure 3B). However, a zero-Ω link plus its pads, traces and via hole will have a significant overall series inductance – replacing this with a direct bond might affect EMC, so retesting is recommended.

If it is found that carefully chosen values of capacitance are necessary, then – like the "designing resonances to miss problem frequencies" technique above – small changes in ICs, circuits or assembly parts or methods could dramatically worsen the EMC performance, so EMC-competent QA and change control is very important (see later).

5.8 Using resistors to 'dampen' cavity resonances

[15] and [16] describe using resistive PCB-chassis bonds to 'dampen' cavity resonances and reduce the amplitude of the resonant peaks. In practice, this technique is like the hybrid bonding technique described earlier, using resistors instead of capacitors (or resistors in series with capacitors) with appropriate pad patterns (see Figures 3D and 3K, and 3M below).

The higher the inductive impedance of a PCB-chassis bond, the lower the benefits for EMC. But the higher the resistive impedance of a bond, the lower the amplitude of the PCB-chassis cavity resonances. At the resonant frequencies of a PCB-chassis cavity, much higher currents flow through the bonding points. Resistance is lossy, converting current into heat, so any resistance in the bonds

will cause greater losses at the resonant frequencies – reducing the peak amplitude of the resonances (reducing the cavity's "Q").

Increasing PCB-chassis bonding impedances with resistors has the downside of decreasing EMC performance at non-resonant frequencies. So, when using this technique to reduce emissions or improve immunity at a cavity resonance frequency, be aware that it might make a marginal performance at other frequencies non-compliant.

[15] and [16] found that resistors between 47 and 100Ω worked best, but maybe this was related to the 50Ω source impedance used for the experiments. Experimental EMC tests (e.g. pre-compliance tests) are always recommended to discover which is the best configuration of DC, capacitive, and resistive bonds, and which are the best values of capacitors and resistors to use.

A well-damped PCB-chassis structure is not likely to be the cause of significant EMC problems despite design changes, alternative components, and manufacturing tolerances, but EMC-competent QA and change control is always recommended.

5.9 Using absorber to 'dampen' cavity resonances

Carbon and/or ferrite-loaded materials (usually elastomers) are available that absorb RF energy. These are similar to the material used to line EMC test chambers to dampen their cavity resonances, and they can also be used inside shielded boxes to dampen their resonances (see [14]).

When they are placed in a cavity they convert electromagnetic energy into heat, and (like the resistive bonds described above) they reduce the resonant peakiness (the cavity's "Q") thereby reducing emissions at the resonant frequencies (and, by a reciprocal argument, increasing immunity).

A number of manufacturers now offer suitable ferrite-loaded materials, usually based on flexible sheets with self-adhesive backing. The thicker sheets usually achieve greater damping at lower frequencies. Ferrite-loaded sheets should work best when located near to the PCB-chassis bonds, where the magnetic fields at resonance should be the highest. Carbon-loaded materials (usually blocks of foamed plastic) should provide their best damping when located in-between the bonds, where the electric fields are the highest at resonance. However, experimental EMC tests (e.g. pre-compliance tests) are always recommended to discover which is the best and lowest-cost material to use, where best to locate it and how best to support it.

There seems to be no reason why ordinary ferrite material, such as has been used for many years to suppress CM currents on cables, could not be used instead, apart from the difficulty of mounting such hard, brittle and dense materials. Maybe standard ferrite cylinders or toroids could simply be slipped over the metal spacers used for bonding (perhaps with a blob of silicone glue to help prevent vibration and wear)?

Ferrite cylinders and toroids have for many years been available at low cost in a very wide variety of dimensions, in a material that gave a wide choice of resistive impedances all with their peak lossiness around 300MHz. Recently new materials have been developed that achieve their peak lossinesses around 700MHz or 2.45GHz. As far as the author knows, solid ferrite has never been used in this way as a means of damping a PCB-chassis resonance, but it seems an idea worth trying.

As for resistive damping above, absorber damping should continue to be effective despite design changes, alternative components, and manufacturing tolerances, but EMC-competent QA and change control is always recommended.

5.10 Reducing the impedance of capacitive bonds

The extra inductance associated with using capacitive bonds (see above, and Figure 3D and 3K) can be reduced by using multiple capacitors arranged radially around each bonding location, as shown by Figure 3M.

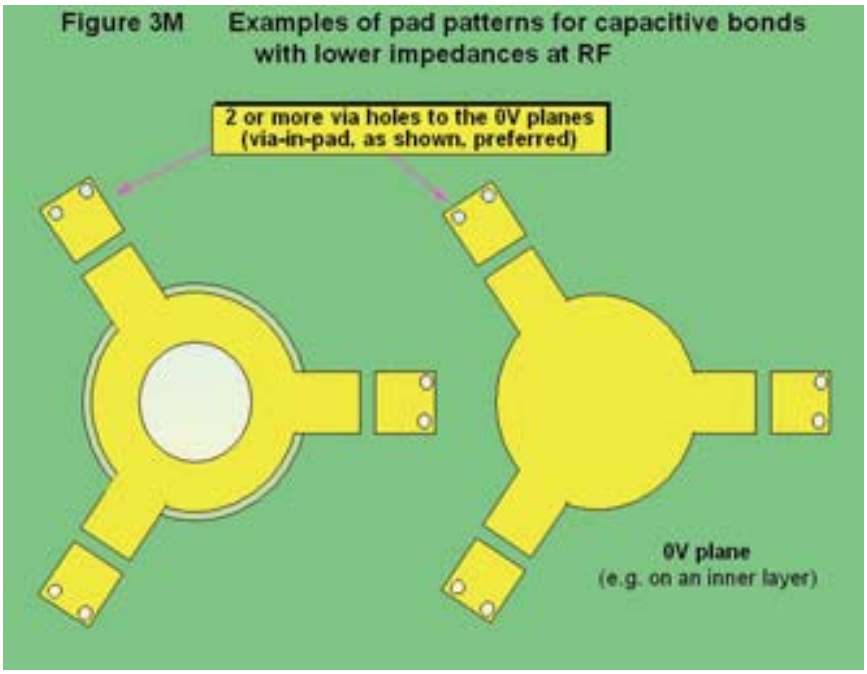

The inductances associated with each capacitor and its pads, traces and via holes all appear in parallel. Using a radial arrangement, such as that as shown in Figure 3M, ensures that the mutual inductance between them cancels out, so that their overall inductance is simply the value of one capacitive bond divided by the number of capacitors. Arranging the capacitors in a parallel array would not cancel out their mutual inductances and the overall inductance achieved would not be as low as with the radial arrangement shown.

Three capacitors were chosen for the sketch in Figure 3M, but just two capacitors, or four or more, could (of course) be used instead.

5.11 Using shielding techniques

We could surround our PCB with spring fingers or conductive gasket that makes a continuous perimeter bond between the edges of the 0V plane and the chassis. This would make the PCB-chassis cavity into a fully shielded enclosure, reducing its emissions considerably (and improving the immunity of the PCB's circuit).

Just like the PCB shielding cans shown in Figures 2H and 2J in [14], such shielding could increase crosstalk between different circuits on the PCB, and at its resonant frequencies its emissions might not be reduced by as much as might be expected. The same remedies for this problem as were described in [14] also apply to the PCB-chassis cavity…

- Break the cavity up into many smaller cavities each with a its lowest resonant frequency higher than the highest frequency of concern
- Add absorber to the cavity to dampen its resonant peaks

Breaking the cavity up into many smaller cavities could in some applications simply mean using a mesh-shaped conductive gasket, making contact with exposed traces on the bottom of the PCB. These bonding traces replace the bonding pads described earlier in this article, and should be via'd to the internal unbroken 0V plane using the PCB shielding rules given in [14].

5.12 Using fully shielded PCB assemblies

If fully-shielded PCB assemblies as described in [14] are used, the need for PCB-chassis bonding to improve EMC is reduced, although bonding to a solid metal or sheet metal chassis might still provide a useful reduction in transfer impedance (see above), especially at lower frequencies.

6. Daughter and mezzanine boards

Everything that has been written above also applies to mezzanine and daughter boards. For such boards, the 0V plane in their motherboard can provide some of the benefits of a local chassis. The cavity created between them can resonate, creating problems for emissions and immunity, especially where the mezzanine/daughter board is connected to external cables (as numerous designers have discovered at great cost to their projects). For these reasons, all of the PCB-chassis bonding benefits and techniques discussed so far also apply to a mezzanine/daughter board-motherboard structure.

In some circumstances it can be sufficient to bond the mezzanine/daughter board frequently enough to the motherboard. The inter-board connectors can help with the bonding, by using numerous 0V pins to bond the two boards' 0V planes all along the length of the connector.
A chassis can also be used to improve the EMC of a mezzanine/daughter board, using all of the techniques described above. It will generally be best if the bonds between the mezzanine/daughter board and its motherboard are carried straight through the motherboard to bond to the chassis.

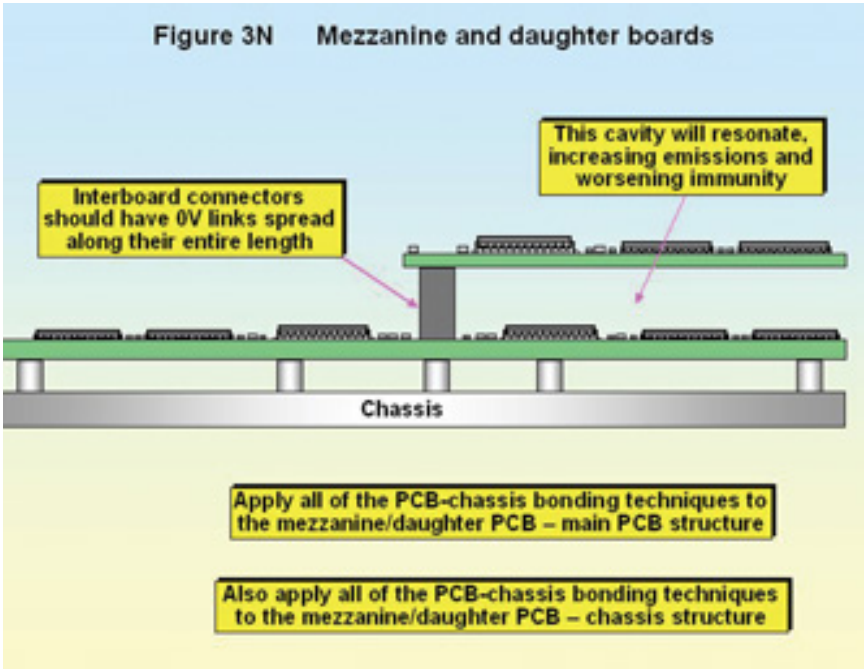

7. EMC-competent QA, change control, cost-reduction

Most types of equipment these days rely on advanced silicon ICs, although they are often such commodity items that we don't tend to notice the amazingly sophisticated technologies they routinely employ. High-tech silicon has many EMC problems; and EMC regulations are widespread world-wide; and as a result EMC-competent QA and change control is now a standard requirement in almost all electronics companies.

EMC-competent QA would ensure that serial manufacture always resulted in goods with similar enough EMC performance. This requires that all the EMC-critical aspects of the design are identified and controlled in manufacture, for instance by using appropriate Work Instructions. Good EMC QA also employs custom EMC tests when goods are delivered, plus custom EMC tests at various stages

during assembly, and at final test.

Custom EMC tests can be surprisingly low-cost to set-up and easy to apply, but they are often application-specific. Good EMC QA will also randomly select a product from time-to-time and subject it to full compliance tests, with records kept and trends analysed to prevent a bad batch of products from being made, or worse still – shipped.

EMC-competent change control is vital. EMC-competent people and EMC test facilities (ideally in-house, to save cost) are used to assess every request for a change or a concession, no matter how trivial they might seem, for their possible EMC implications. Of course, change control cannot be done if it is not known what EMC design techniques an equipment relies upon, so it relies on designers documenting important EMC issues in such a way that future engineers who need to know what consequences a change might have, are fully informed.

Products in serial manufacture are often subjected to cost-reduction exercises to reduce manufacturing costs and improve profitability. The author knows of several instances where such cost-reduction exercises resulted in such severe problems with interference in operation that the overall cost to the company was increased dramatically, rather than reduced. So the EMC-competent change control described above must also be applied to any cost-reduction exercises, no matter how trivial the issues seem to be.

For instance, the production manager in one company decided to employ self-tapping screws that were pre-lubricated with a blob of wax, to save time and tool-wear in assembly. But the wax increased the resistance at all the metal bonds that they were relying on to create a reasonable RF Reference Plane to help achieve EMC compliance (although they did not realise this at the time). Numerous similar very costly real-life examples could be described to emphasise how very important it is to fully understand the EMC design of an equipment, and to employ EMC-competent control on every change to it.

8. 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 trade-offs (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 part 1 of this series [12], plus the final section of part 2 [14].

9. References

[1] Keith Armstrong, "Design Techniques for EMC" (in six parts), UK EMC Journal, Jan – Dec 1999
[2] Keith Armstrong, "Design Techniques for EMC" (in six parts), http://www.compliance-club.com/KeithArmstrongPortfolio
[3] M K Armstrong, "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 1", IEE Electronics and Communication Engineering Journal, Vol. 11 No. 4 August 1999, pages 185-194
[4] M K Armstrong, "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 2", 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] Tony Waldron and Keith Armstrong, "Interference-free pro-audio systems and installations", EMC York 2002 Conference, Conference Proceedings: ISBN 1-902009-09-6.
[7] Tony Waldron and Keith Armstrong, "AES-X13 and Interference-free Pro-Audio Systems", IEEE International Symposium on EMC, Minneapolis, August 2002, Symposium Record: ISBN 0-7803-72643-6, pp 434-439.
A similar paper with the same title was presented at the IEE's Seminar: "EMC – its (nearly) all about the cabling", Wednesday 22 January 2003, sales@iee.org.uk.
[8] Tony Waldron, "Designing for Interference-Free Audio System Components", EMC Compliance Journal, July 2002, pp 22-33, from the archives at http://www.compliance-club.com
[9] Tony Waldron, "A Practical Interference-Free Audio System", EMC Compliance Journal, Part 1 September 2002 pp 18-28, Part 2 November 2002 pp 22-31, from the archives at http://www.compliance-club.com
[10] Tony Waldron and Keith Armstrong, "Bonding Cable Shields at Both Ends to Reduce Noise", EMC Compliance Journal May 2002, pp 14-27, from the archives at: http://www.compliance-club.com
[11] Keith Armstrong, "Audio-Frequency Shield Current Induced Noise is Negligible", Audio Engineering Society 114th Convention, Amsterdam, March 2003, tutorial session S5 on "Grounding and Shielding".
[12] Keith Armstrong, "Advanced PCB Design and Layout for EMC – Part 1: Saving Time and Cost Overall", EMC & Compliance Journal, March 2004, http://www.compliance-club.com.
[13] A A Smith, "Coupling of External Electromagnetic Fields to Transmission Lines", Wiley, 1977.
[14] Keith Armstrong, "Advanced PCB Design and Layout for EMC – Part 2:Segregation and interface suppression ", EMC & Compliance Journal, May 2004, http://www.compliance-club.com
[15] T Williams, "Controlling Resonances in PCB-Chassis Structures", EMC Europe 2002, International Symposium on EMC, September 9-13 2002, Sorrento, paper OD1, pp 305-310. Download at: http://www.elmac.co.uk.
[16] T Williams, "Resonant Circuit Board Effects", Approval magazine, Jan/Feb 2002, pp 28-31.
[17] Mark I. Montrose, "Printed Circuit Board Design Techniques for EMC Compliance, Second Edition", IEEE Press, 2000, ISBN 0-7803-5376-5

I would like to reference all of the academic studies that back-up the practical techniques described in this series, but the reference list would take longer to write than this series! But I must mention the papers presented at the annual IEEE International EMC Symposia organised by the IEEE's EMC Society (http://www.ewh.ieee.org/soc/emcs), especially the dozens of wonderful papers by Todd Hubing's staff and students at the University of Missoura-Rolla EMC Lab (http://www.emclab.umr.edu), and papers by Bruce Archambeault of IBM and the EMC experts at Sun Microsystems.

Many other contributors to the IEEE EMC Symposia, and other conferences and symposia and private correspondence are far too numerous to be named here, but the following stand out: Tim Williams of Elmac Services, http://www.elmac.co.uk; Mark Montrose of Montrose Compliance Services, http://www.montrosecompliance.com; John Howard, http://www.emcguru.com; Tim Jarvis of RadioCAD, http://www.radiocad.com; Eric Bogatin of Giga-Test Labs, http://www.gigatest.com; and dozens of application notes from National Semiconductor; Rambus Corp.; California Micro Devices; Intel; IBM; Cypress Semiconductor; Xilinx; Sun; Motorola; AVX; X2Y Attenuators; Giga-Test Labs; Ansoft and Flomerics. I apologise to the very many excellent people and companies that I have had to leave out of this list.

Some useful textbooks and other references are:
[A] Mark I. Montrose, "EMC and the Printed Circuit Board, design, theory, and layout made simple", IEEE Press, 1999, ISBN 0-7803-4703-X

[B] Howard W Johnson and M Graham, “High Speed Digital Design, a Handbook of Black Magic”, Prentice Hall, 1993, ISBN 0-1339-5724-1
[C] Howard W. Johnson, “High Speed Signal Propagation: Advanced Black Magic”, Prentice Hall, 2003, ISBN 0-1308-4408-X

Eur Ing Keith Armstrong C.Eng MIEE MIEEE
Cherry Clough Consultants,
Member of: EMCTLA, EMCUK, EMCIA keith.armstrong@cherryclough.com