emc STANDARDS

Another EMC resource
from EMC Standards

EMC Techniques for Heatsinks

Helping you solve your EMC problems

9 Bracken View, Brocton, Stafford ST17 0TF T:+44 (0) 1785 660247 E:info@emcstandards.co.uk

EMC Techniques for Heatsinks

First published in The EMC Journal, Issue 89, July 2010, pp 31-38

Eur Ing **Keith Armstrong** CEng, FIET, Senior MIEEE, ACGI
phone/fax: +44 (0)1785 660 247, keith.armstrong@cherryclough.com

Contents

1 Controlling stray heatsink currents

Heatsinks can be a significant cause of emissions, that can cause "Internal EMI" problems (see Section 8 in Part 4 of [1]) and/or "External EMI" problems, including failure to comply with regulatory or customer emissions limits.

As operating frequencies rise and ICs and power transistors switch ever-faster, their heatsinks are becoming more important as a cause of emissions. This is especially so because the upper frequencies are now so high that they are more likely to excite electrical resonances in even quite small heatsink assemblies. Resonance can increase a heatsink's emissions by as much as 40dB.

An IC or transistor fitted with a heatsink has a significant stray capacitance to that heatsink, as sketched in Figure 1, because thermal interface washers are thin, and have high values of dielectric constant (ε_r).

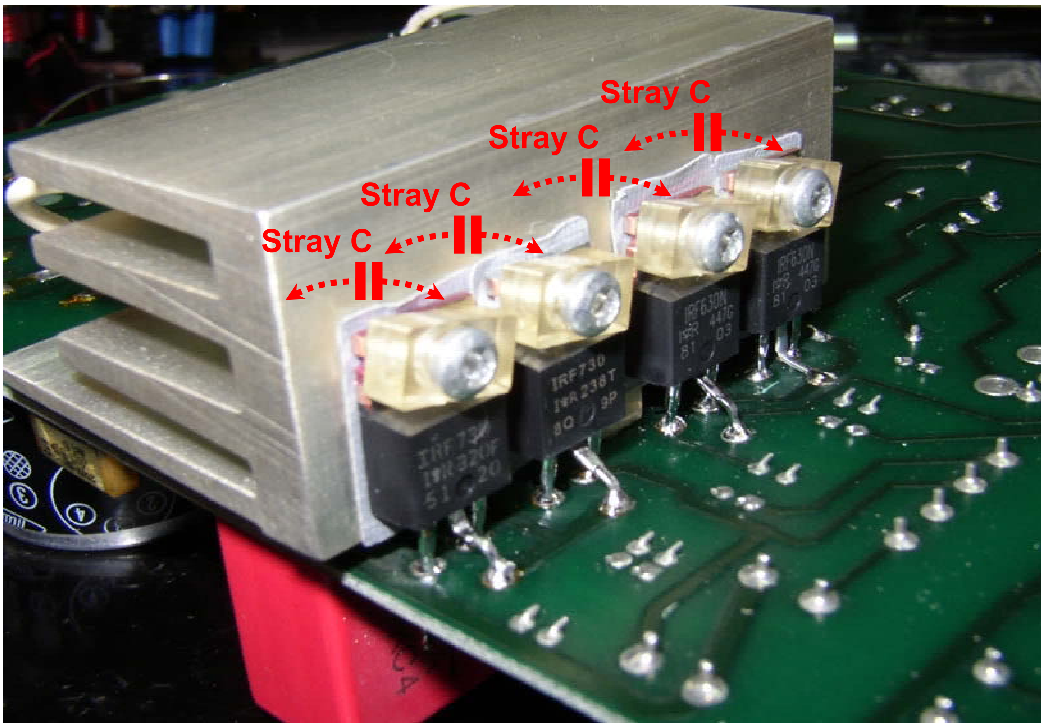

Figure 1 Example of transistors on a heatsink

Sometimes the heatsink tabs of devices are bonded directly to a heatsink. Such devices might have internal electrical isolation, meaning they have thermal washers inside that still cause stray capacitance coupling. Sometimes device heatsink tabs are directly connected to a drain or collector, so the heatsink experiences their full voltage excursions and is "live" (and often dangerous as well).

When the voltages of the IC or power transistor fluctuate, this stray capacitance (or direct connection) injects stray currents into conductive heatsinks, so returning these stray currents to their sources – the power supplies that power the ICs or transistors – is very important for controlling EMC emissions.

Part 1 of [1] discusses how fundamental it is to good EMC design (that saves time and cost and reduces field failures and warranty returns) to control the flow of return currents, including return currents caused by stray capacitive and inductive coupling – such as occurs with heatsinks.

Ceramic or plastic heatsinks (e.g. "CeramCool", "CoolPoly") are not conductive and so carry no stray currents. They do not need thermal insulating washers so can sometimes have better heat dissipation than metal, as well as having no EMC emissions. Some types can have metallised/etched surfaces, like a single-layer printed circuit board (PCB), allowing direct bonding of devices, originally designed for the power LED arrays for new LED-lighting products, but useful as a cost-saving measure in a wide variety of other electronic products.

It may well be that the best way to prevent emissions from heatsinks from causing internal or external EMI problems, or failure to comply with specified emissions limits, is to use a ceramic or plastic heatsink.

However, device heatsink 'tabs' will still emit, and although usually negligible they can sometimes cause serious emissions problems for certain types of designs, so some of the following techniques may still be required even when ceramic or plastic heatsinks are used.

The best (worst?) example of this that I have seen was a small low-cost inverter to drive a cold-cathode fluorescent lamp from a 12V supply. The PCB was only about 30mm square, most of it taken up by the base of the lamp. A single T0-220 power transistor did all the hard work of switching the flyback step-up voltage converter at about 70kHz, switching so quickly that a heatsink was not required, helping to save cost.

Its conducted emissions were more than 50dB above the Class B limit line, all the way up to a few MHz! We discovered quickly that they were all due to the emissions from the tiny heatsink tab of the TO-220 switching device. The TO-220 was standing up vertically from the tiny PCB, so its heatsink tab was as far away from the PCB's 0V plane as it could be – not that it was a very large 0V plane anyway, given the size of the product.

It was a huge surprise to see such high levels of emissions at such low frequencies, from such a small stray capacitance as could be possibly be caused by a TO-220 device, especially as there was no way that it could have been experiencing a structural resonance at such frequencies.

This story has a happy ending, with the problem being completely solved (emissions reduced by over 60dB) simply by controlling the stray current path using the techniques described below, for a negligible additional cost.

2 Returning stray heatsink currents to the power rails

The heatsink currents are sourced from both power rails feeding the IC or power switching device, so they must to be returned to the same power rails with an impedance that is much lower than free-space (377Ω), ideally << 1Ω, at all frequencies up to the highest frequency of concern (f_{MAX}).

This is called "RF-bonding" and the inductance (L) of the return current loop that is created must be << $1/2\pi f_{MAX}$ (f_{MAX} in MHz gives L in μH), and to prevent resonance, the loop's largest dimension (D, diagonal or diameter) must be << $15/f_{MAX}$ (f_{MAX} in MHz gives D in metres)

Generally, we avoid earthed/grounded heatsinks, because they usually have so many possible low-impedance paths for the heatsink's stray return current that it is difficult to ensure it mostly flows in the smallest local loop.

Designers often imagine that anything called earth, ground, or chassis is some sort of magical infinite sink for electrical current with a negligible impedance for currents of any size up to tens of GHz. This very widespread belief in a sort of electronic fairy-tale, almost certainly costs manufacturers of electronic products, and of all the applications that use them, billions of dollars in project delays and warranty costs worldwide every year (and costs their customers at least as much in downtime).

What really matters for EMC and signal integrity are the loops in which currents flow [1], including all of the stray currents, and if those loops do not include the earth, ground or chassis, then connecting to the earth, ground or chassis simply adds an extra path for common-mode currents with their powerful "accidental antenna" effects. Connecting to the safety earth or ground system usually makes most EM phenomena worse, for emissions and immunity, which is why a very highly respected EMC engineer once described the safety earthing/grounding system as an "Interference Distribution Network".

Where a heatsink *must* form part of a metal chassis or enclosure (not a good idea in general), either use shielded heatsink washers to divert most of the stray current away from the heatsink (see later) and/or shield the product (see later).

We usually return the stray heatsink current to just *one* of the power rails, as shown in Figure 2, and rely on good power decoupling on the PCB to carry the return current from it to the other rail with << 1Ω up to f_{MAX}. For digital and mixed-technology ICs this generally requires both power rails to be PCB planes that maintain a very low impedance up to at least the highest frequency of concern, as described in Part 5 of [2] and Chapter 4 of [3].

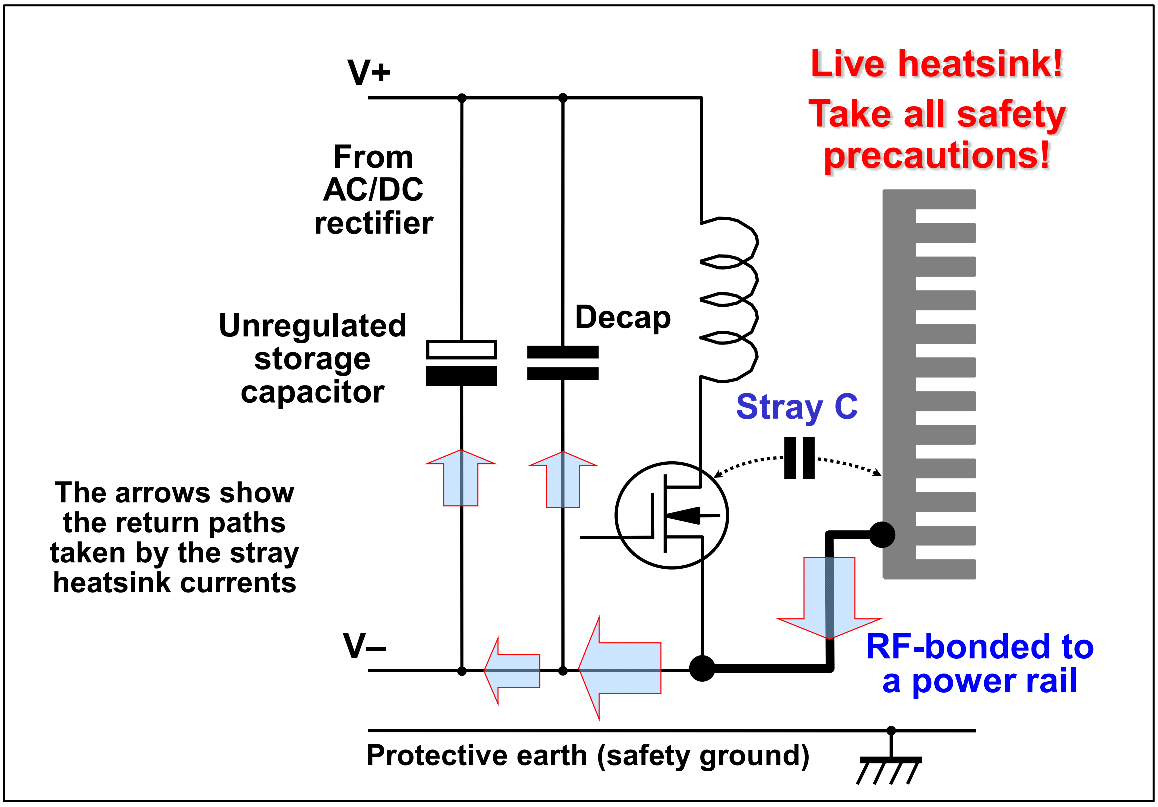

Figure 2 Example of an off-line switch-mode converter with direct RF-bonding

Directly connecting a heatsink to a power supply at a hazardous voltage, for example the rectified mains in an off-line switch-mode power converter, can increase safety risks to unacceptable levels. Of course, ICs and transistors powered from "Safety Extra Low Voltage" (SELV) can use direct RF-bonding for their heatsinks, without creating safety hazards.

Where power devices have their drains or collectors connected directly to a heatsink, it may not be possible to connect the heatsink directly to a power rail without stopping the circuit from operating.

In either of the above situations, instead of a direct connection to the power rail, a series capacitor may be used. The problem with a capacitor is achieving a low-enough impedance in the return path over a wide range of frequencies. Where the emissions are only troublesome over a narrow range of frequencies, it may be practical to use a series capacitor and choose its type and value so that so that it series-resonates with the inductance inherent in the current return path, at the frequency of the heatsink's highest emissions.

Such designs can easily be optimised ("tuned") at an ordinary electronics testbench by picking up the heatsink's E and H field emissions with close-field probes (see Parts 1 and 2 of [4]) – there is no need to visit an EMC test laboratory.

Existing metal parts, such as heatsink mounting components, are often pressed into service in the heatsink-return-current's RF-bonding path. Surfaces forming part of the current path should be conductively plated or passivated to prevent oxidation over time, especially aluminium and mild steel. The metals and plating used on mating parts should be chosen for low galvanic corrosion at their junctions – the importance depending on their exposure to condensation and other wetness in their environment.

Beware of manufacturing residues or coatings on metal parts (especially anodising or polymer passivation). Anodising is often used because it has much higher thermal emissivity than plain oxidised aluminium (0.8 instead of 0.2) and so makes for smaller heatsinks, but it makes RF-bonding very difficult. Tin plating is quite good (emissivity of 0.5) and makes RF-bonding very easy to do.

Because metal platers sometimes apply insulating polymer passivation coatings when they are not specified or requested, and sometimes in direct contravention of instructions, I strongly recommend checking the conductivity of all RF-bonding surfaces at Goods Receiving before accepting any batch of parts, using a simple ‘buzzer’ tester with low contact pressures for the probes (e.g. strips of “fabric-over-foam” conductive EMC gasket).

3 Using a PCB plane as a heatsink

For cost-effective EMC, most circuits need to use PCB layers dedicated to 0V and power planes over most/all of their areas. Using a plane as a heatsink injects stray currents from devices directly into that plane, so (as long as it is one of the rails that supplies the devices) there is no need for any RF-bonding.

This requires mounting the devices so that they are in thermal contact with the plane, via thermal grease and/or an insulating thermal washer. Some devices are designed for their heatsink tabs to be soldered to an area of PCB copper to act as a heatsink. If this copper area must be isolated from the 0V or power rails, then it needs RF-bonding to them as described in Section 2 above.

Where the PCB plane that is to be used as a heatsink is an internal layer, Figure 3 shows the two main techniques: either machine an aperture in the board material to allow direct access to the plane, or use a small plane on the component side, and (if necessary) thermally-bond it to a larger internal PCB heatsink plane with dozens of via holes spread all over its area – often called “thermal vias”.

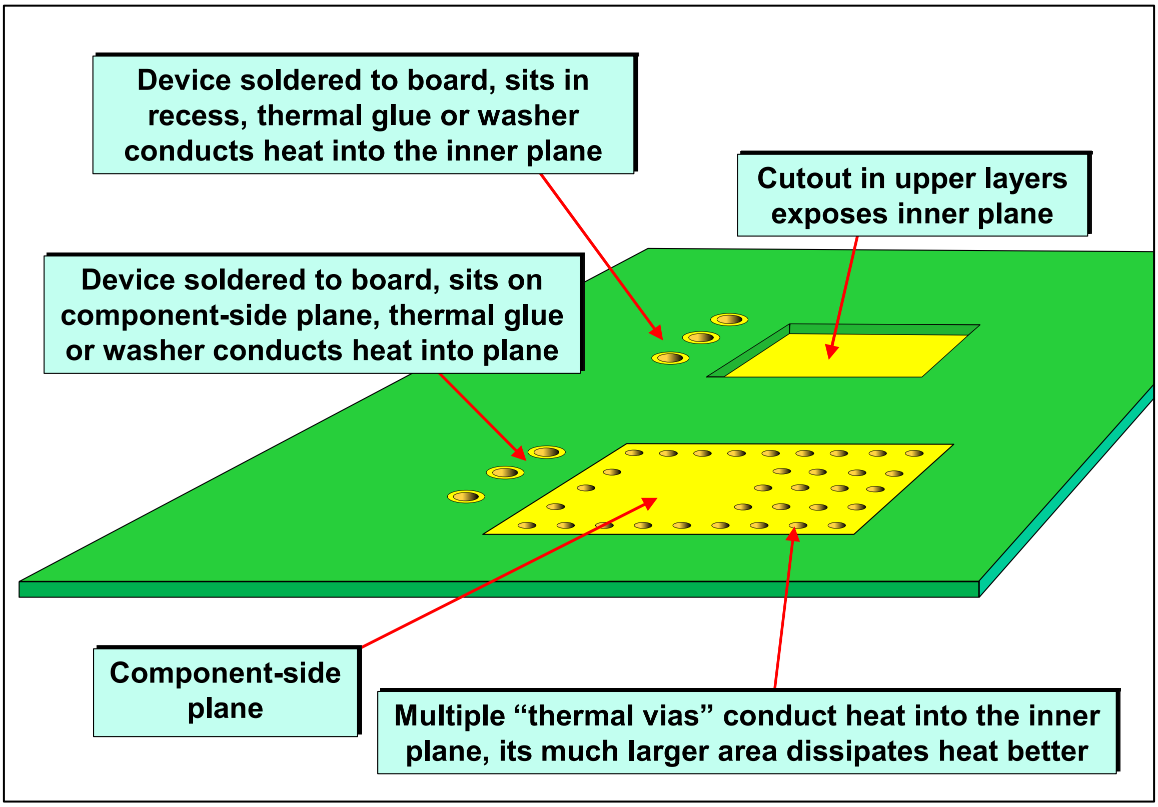

Figure 3 Examples of using a PCB plane as a heatsink

The thermal dissipation of a PCB plane can be improved by using thicker copper plating. Some PCB manufacturers can plate-up the normal “½ oz.” etched copper foil, to as much as “8 oz.”, and I have recently been offered as much as 12 oz from one Chinese PCB manufacturer.

Plane-chassis bonding with many short metal mounting pillars can help dissipate the heat, as well as providing many other important EMC benefits (see Part 5 of [2], Chapter 3 of [3], and [5]).

Some PCB manufacturers can laminate metal sheets or plates in the PCB and RF-bond them to a 0V or power plane. These sheets or plates can be extended beyond the PCB to form a bracket that can be clamped to a chassis or heatsink to help dissipate the heat from the PCB’s devices. If a PCB uses embedded plates as heatsinks, but they are isolated from the 0V or power planes, they need RF-bonding as described in Section 2 above.

Although embedding metal plates increases the bare-board cost, don’t forget that what really matters is the *overall cost of manufacture*, not the BOM cost [6]. Some manufacturers of high-frequency fluorescent lamp ballasts use this expensive PCB method to reduce the overall cost of their products and improve profitability by reducing the time/cost of assembly and improving yields.

4 Heat sink resonances

Where a heatsink is not itself a PCB plane that powers the heat-sunk devices, as discussed in section 3 above, it should be RF-bonded to such a plane, or at least to a power rail that supplies the devices.

If the impedance of the RF-bonding is inadequate, the heatsink stray return current flowing in it will cause the heatsink to suffer a significant dV/dt, making it radiate more as an 'accidental antenna'.

If the heatsink has a resonant frequency close to a signal frequency (or one of its harmonics), its emissions can increase by 20dB or more, because it is now a 'tuned' accidental antenna (see Part 4 of [1]).

Resonance effects only become significant when the longest heatsink dimension (D) exceeds $30/f_{MAX}$ (f_{MAX} in MHz gives D in metres, f_{MAX} in GHz gives D in millimetres).

The lowest resonance frequency f_{RES} = 150/D, where D is the longest 'three-dimensional' diagonal (D in metres gives f_{RES} in MHz, D in millimetres gives GHz). For example, e.g. a 54mm cube heatsink could have its first (lowest) resonance at around 1.6 GHz, so could cause big problems for nearby GPS antennas.

Accurate analysis requires computer-aided simulation, taking into account:

a) the heatsink's geometry / shape
b) types of semiconductors and their locations
c) the proximity of 0V planes and chassis
d) any RF-bonding

Square or cube shaped heatsinks tend to have the highest resonant frequencies, so are good only if their lowest resonant frequency is well above the highest frequency of concern. But such symmetrically-shaped heatsinks have a higher gain (their "Q") at their resonant frequency, and can increase emissions if their resonances are not well above the highest frequency of concern – in which case make them rectangular, and avoid simple ratios of length:width:height (such as 1.5, 2 etc.). The "Golden Mean" (approximately 1.618) is a good choice when there are no significant constraints. Figure 7D summarises these very basic guidelines.

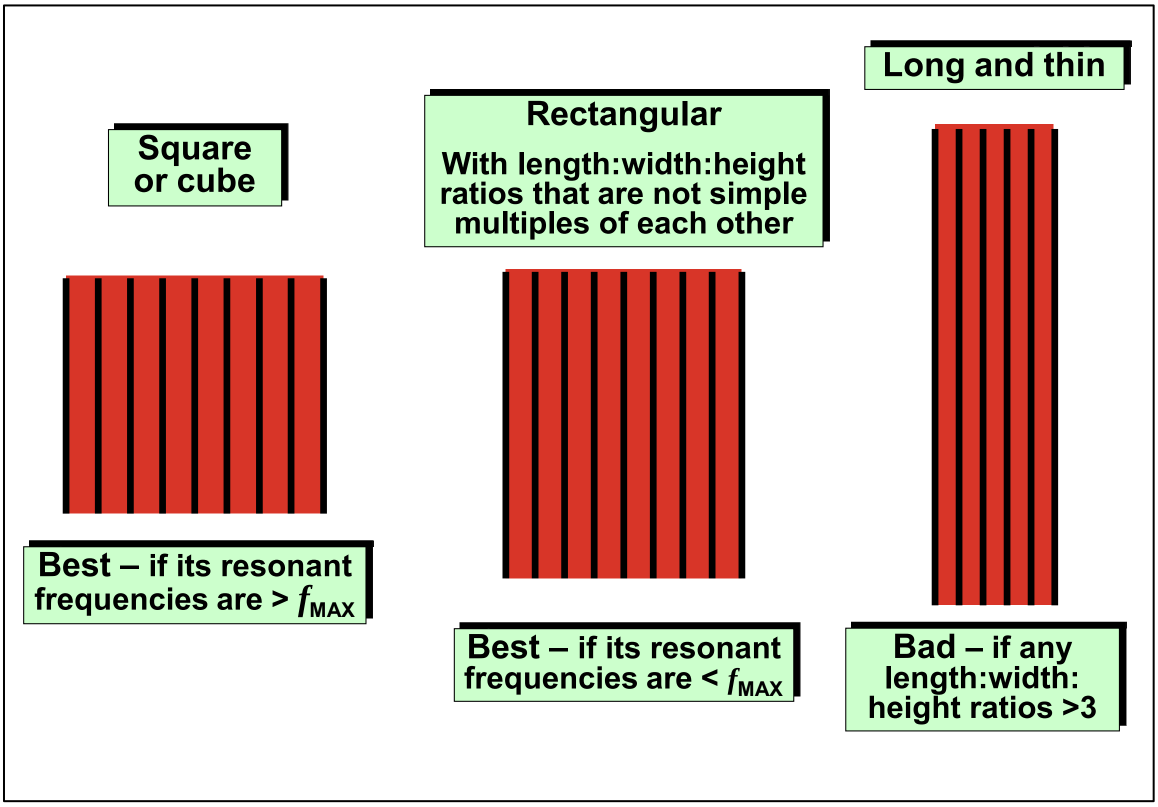

Figure 4 Heatsink shapes (viewed from above)

Where heatsink resonances could lie within (or near to) the frequency range concerned, the best place for the IC or power transistor is generally in the centre of the heatsink's base – usually the best position for good thermal performance too.

Edge locations cause greater resonant gain, and so cause higher levels of emissions.

Vertical fins reduce the resonant gain for the direction in which they run, so it is usually best for the fins to run along the longest heatsink dimension. Fins running perpendicular to a resonant dimension increase its resonant gain, which is not desirable. Figures 5, 6, 7 and 8 sketch these design guidelines.

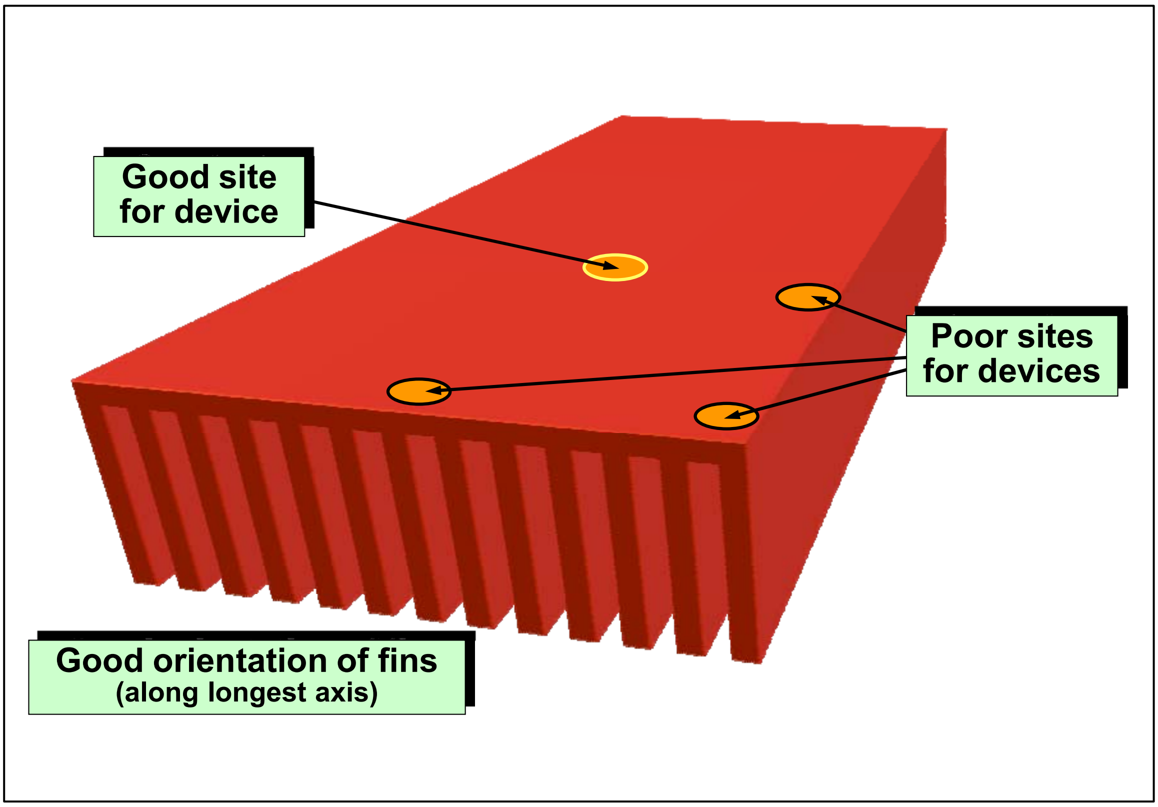

Figure 5 Some sketches of heatsinks

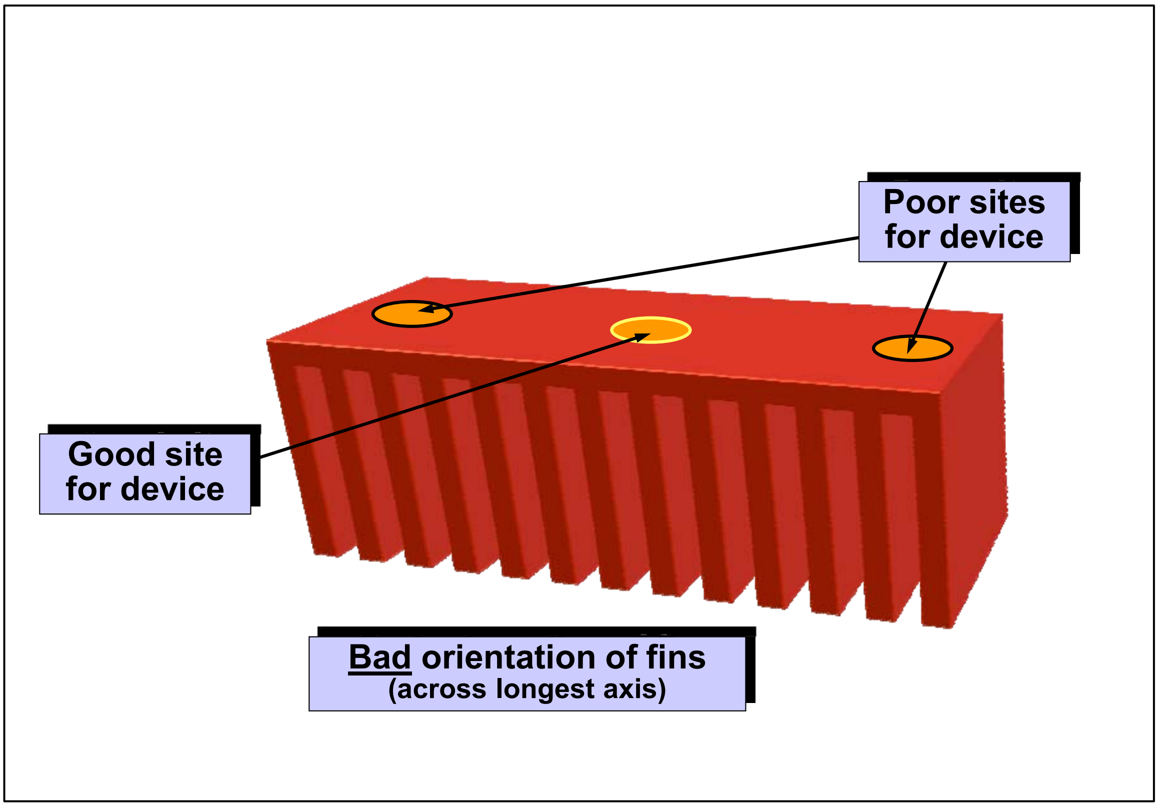

Figure 6 More sketches of heatsinks

Vertical pins increase the resonant gain of all resonant dimensions, which is a pity because "pin" heatsinks are often more efficient than finned types.

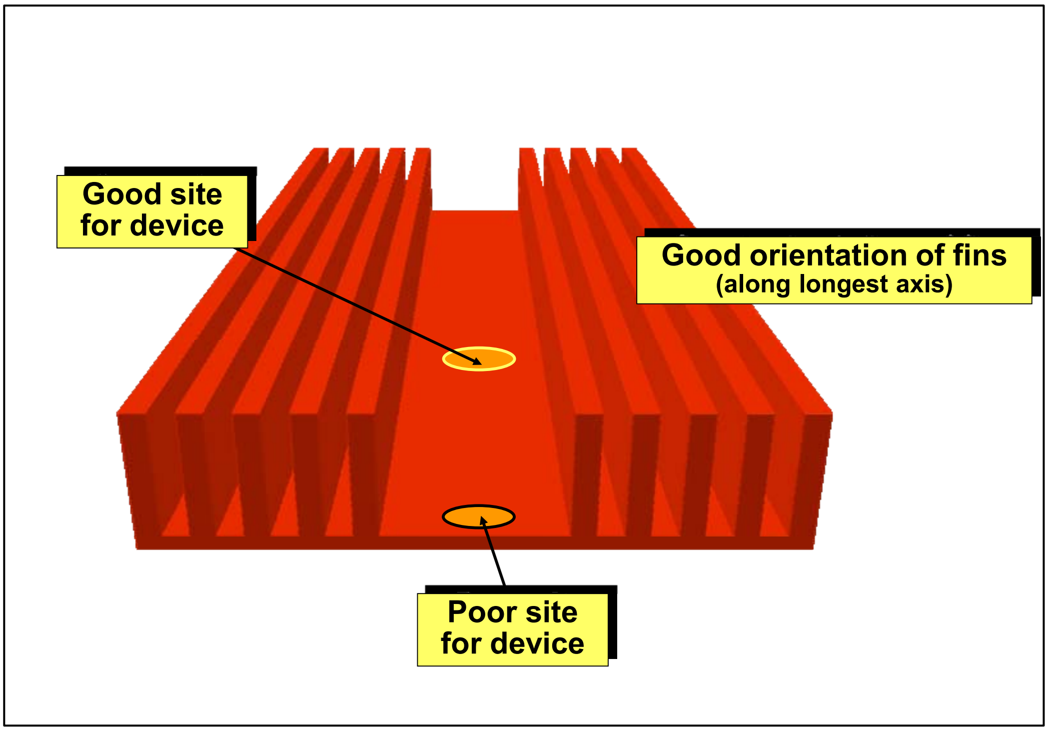

Figure 7 Yet more sketches of heatsinks

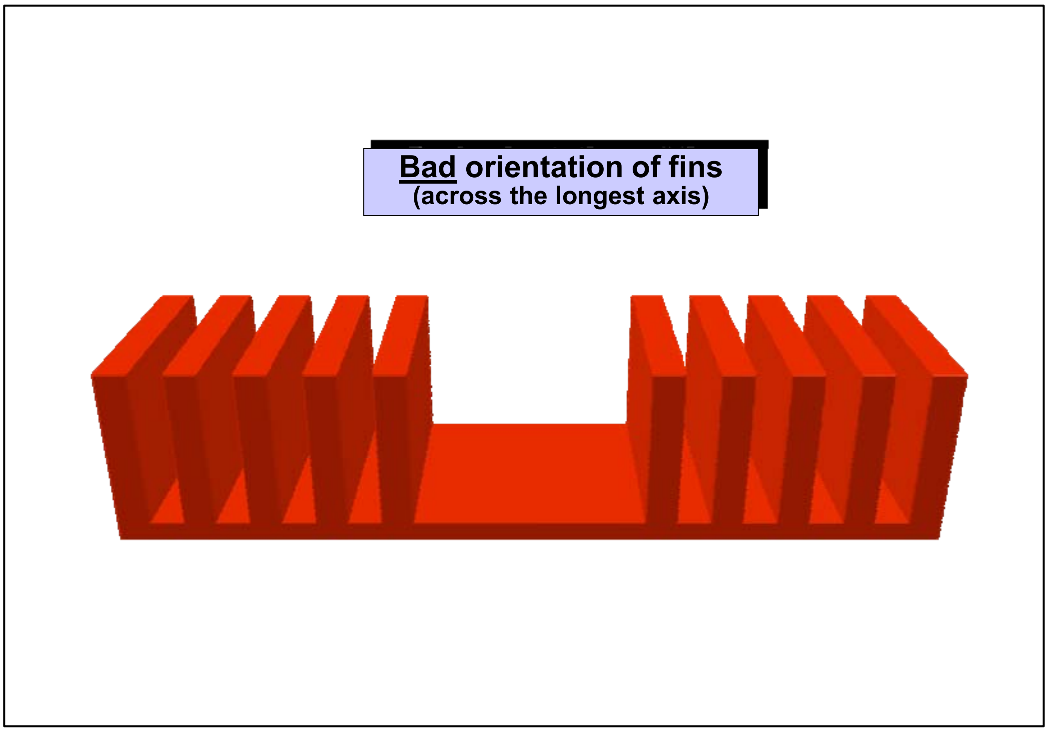

Figure 8 Even more sketches of heatsinks

5 Some useful tricks

Use thick aluminium oxide or gel washers to reduce the device-to-heatsink capacitance to reduce the stray current without increasing the thermal resistance too much.

Use smaller heatsinks, to make their first resonant frequency higher than f_{MAX}, by some combination of:

a) Increasing the airflow rate (e.g. using fans to blow the air, instead of relying on convection)

b) Use liquid instead of air for cooling

c) Use heatsinks with embedded heat lanes (a type of heat pipe) in their bases to spread the device heat uniformly all over the base of the heatsink

d) Use shielded thermal washers (two insulating layers with a metal shield inbetween), with the shield layer RF-bonded to a power rail (see Figures 9 and 10)

Shielded washers are never perfect, so when using them with a safety earthed/grounded heatsink, there will still be some stray current in the heatsink that will need recycling back to the power rail of the device that caused it in the first place. If nowhere else, it should be recirculated by Y capacitors at the power input connection.

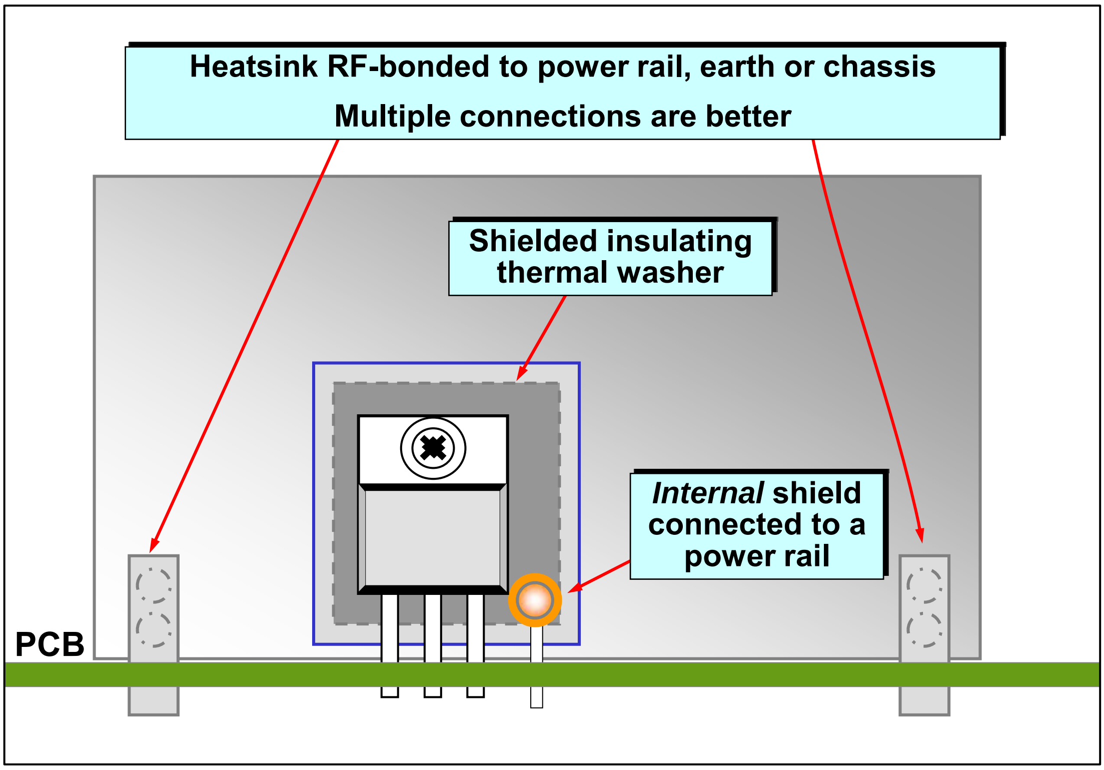

Figure 9 Example of using shielded thermal washers

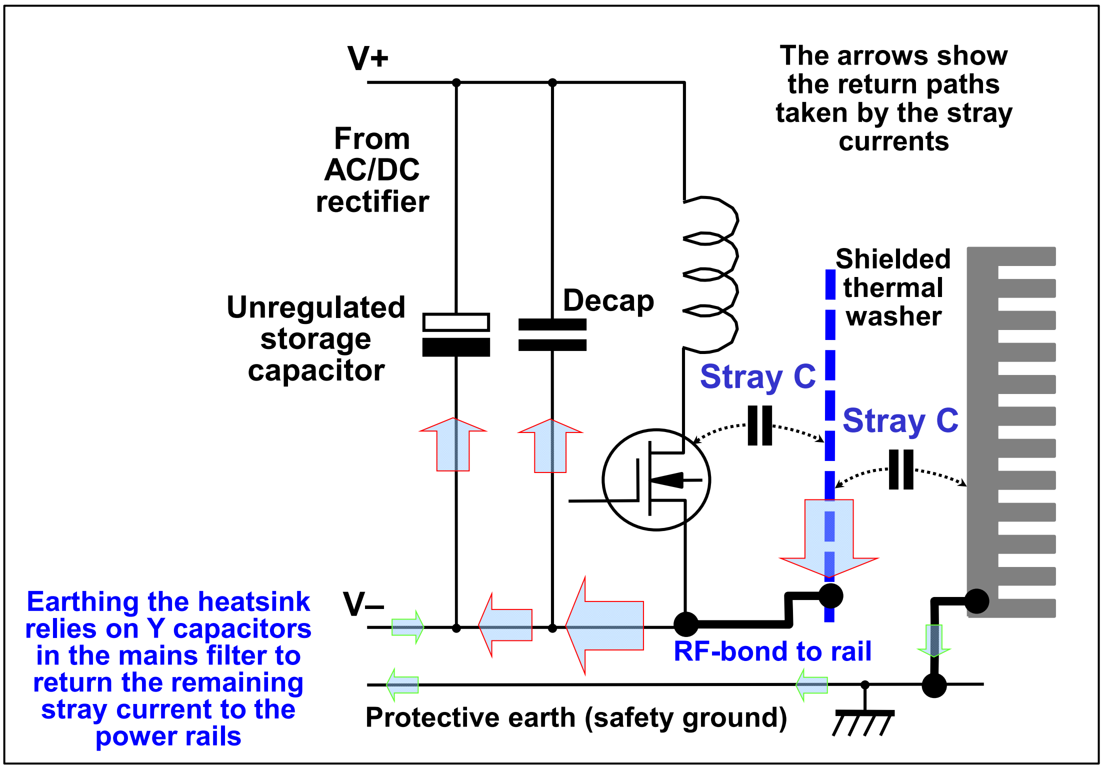

Figure 10 Where the stray current flows, when using a shielded thermal washer

6 Heat pipes

Heat pipes use a small part to collect the heat from a device, and because it is small it has very high resonant frequencies, ideally much higher than f_{MAX}. But the length of the heat pipe, plus the large area of heat dissipater at its far end, make its resonant frequencies very low overall.

It may be possible to fit a clip-on ferrite sleeve/clamp to the pipe close to the collector, to help reduce emissions from the pipe itself and its dissipater, as shown in Figures 11 and 12.

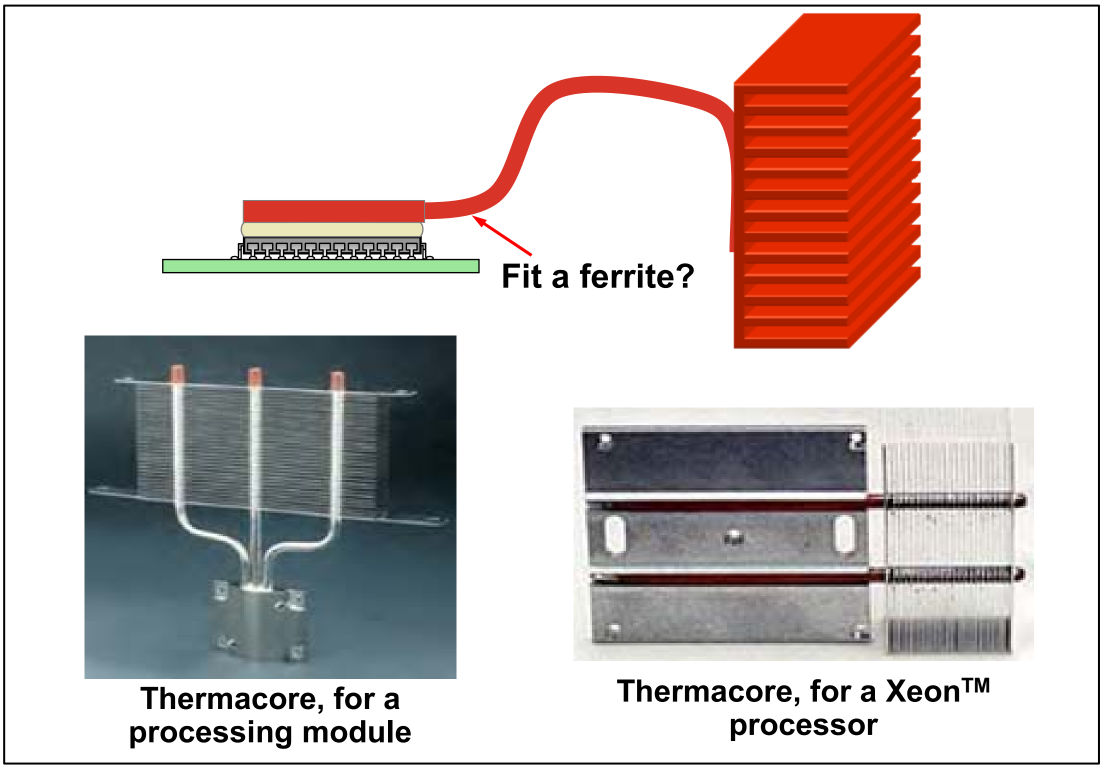

Figure 11 Examples of heat pipes

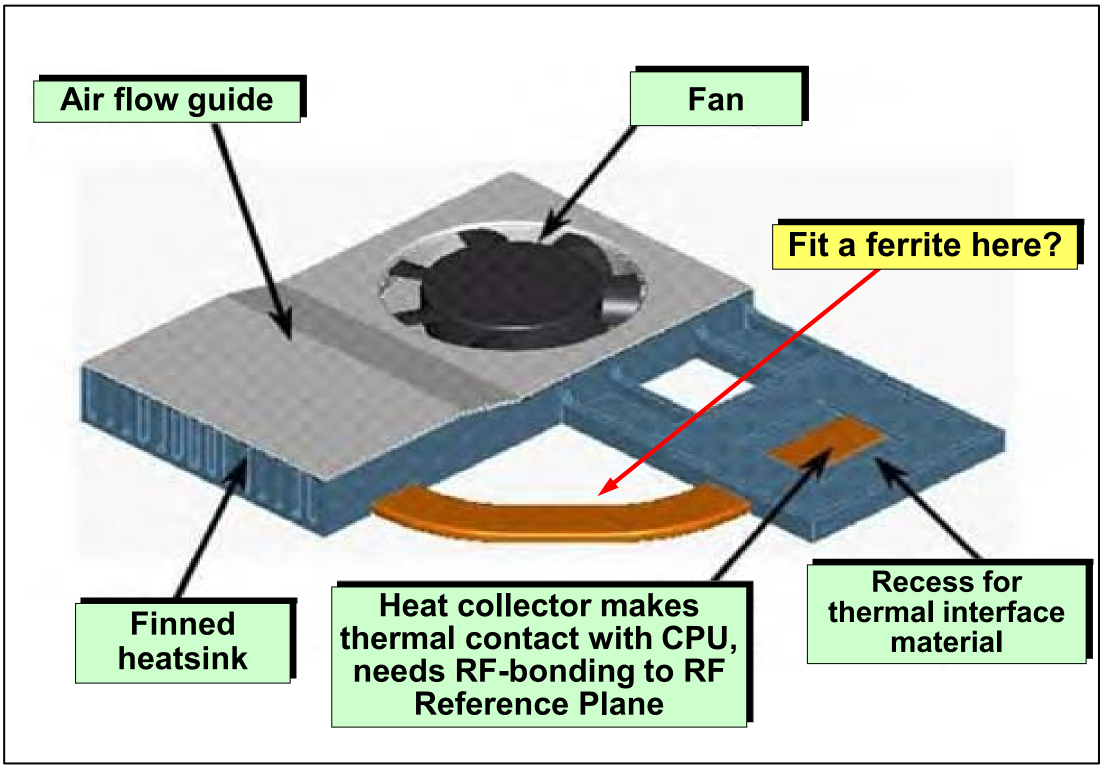

Figure 12 A Pentium III™ Mobile Module heat exchanger

7 RF-bonding heatsinks to reduce emissions

When heatsinking RF and digital ICs, a low bonding inductance is important – as mentioned earlier – and also to prevent the "heatsink-plus-RF-bonding" structure from resonating below f_{MAX}.

The radiated emissions from a heatsink increases with the bonding network inductance to the power of 3.5 (approx.), so for a transistor running at <1MHz, a single short bond to the closest point on its power rail is usually adequate. But, for higher frequencies, evenly distributed bonds are important, as they can have much lower emissions than if they are poorly distributed but achieve the same overall inductance.

Intel recommends RF-bond spacings < 75/f, where f is the 3rd harmonic of the processor core clock frequency. I recommend < 30/f_{MAX} instead (f and f_{MAX} in MHz gives bond-spacings in metres, GHz gives them in mm).

3-D field solvers can quickly reveal problems with heatsink bonding, and also help to find their solutions quickly. Figure 13 shows a simulation of the emissions from an example 'floating' heatsink, very kindly done for me by Paul Duxbury, using FLO/EMC, which is now called Microstripes and supplied by CST (www.cst.com).

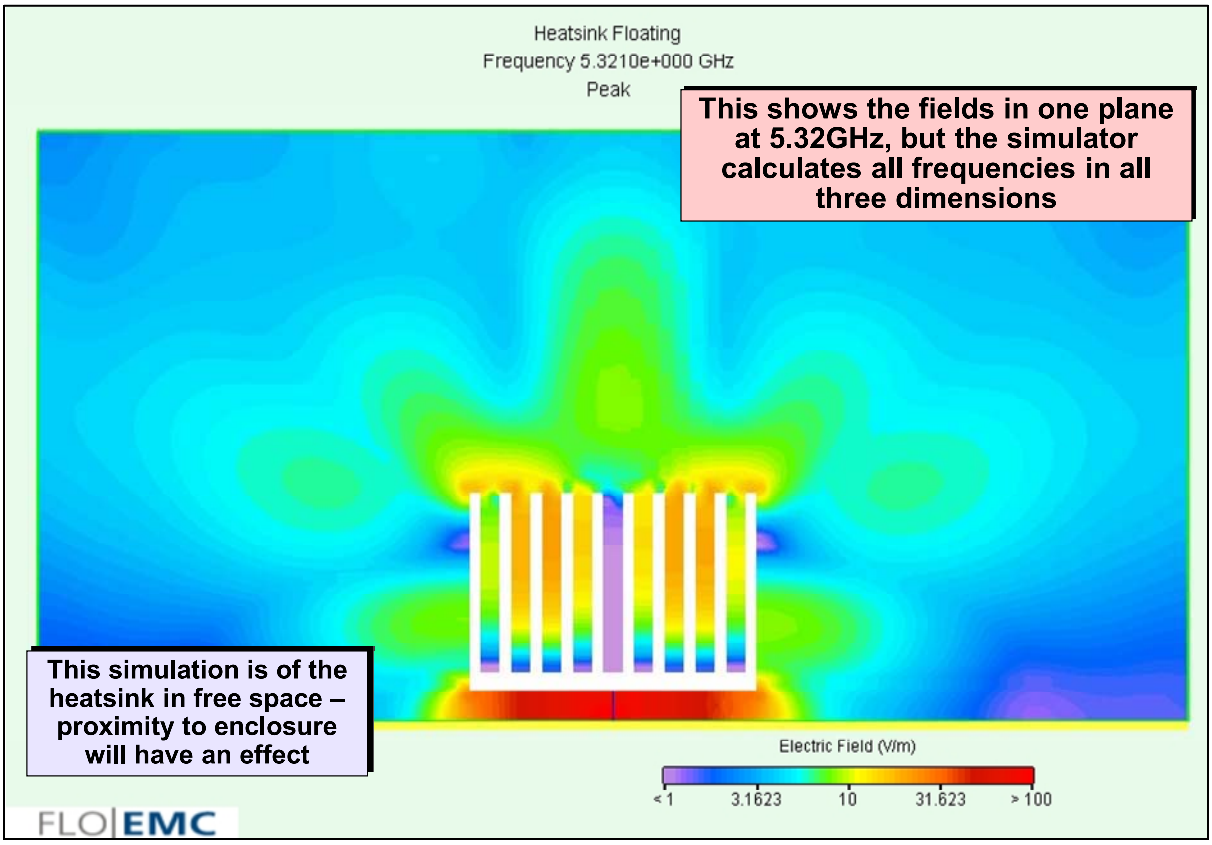

Figure 13 The electric near-fields around the 'floating' heatsink at its resonant frequency

One of the common output formats from 3-D field solvers like Microstripes is a three-dimensional view of the fields around an object, coloured to represent intensity, and fluctuating as the simulated frequency is ramped through the selected range. The resulting colour animations are very attractive, and help designers quickly get an overview of whether they might have problems at any frequency.

A "cutting plane" can be chosen, at any angle, to view the 3-D animation in 2-D, for more precise observations and measurements, and this is what has been used to create Figure 13, which has also chosen the video frame corresponding to the first resonant frequency that appeared.

In Figure 13, the heatsink is assumed to be floating above an RF Reference Plane, with the thin wire centred under it being a source of RF current. Because nothing else is nearby, the electric near-fields around the heatsink are symmetrical. In the far-field (see Part 2.4 of [1]) these complex near-fields will become a simple expanding spherical shell.

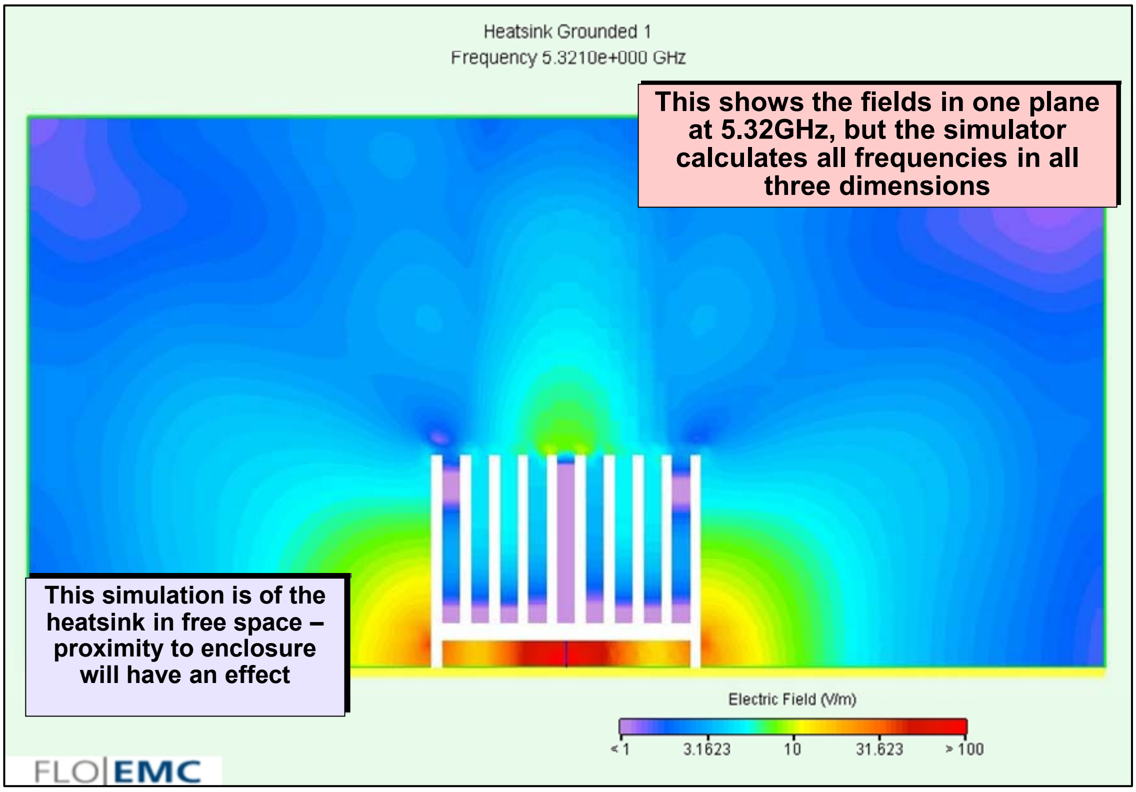

Figure 14 The Fig. 13 with heatsink RF-bonded directly to Reference Plane at each corner

Figure 14 shows the same heatsink as Figure 13, at the same frequency, but this time with zero-Ω bonds to the RF Reference Plane at each of its four corners. The near-fields are now noticeably closer to the RF Reference Plane – indicating lower far-field emissions – but of course from this one video frame we don't know whether the emissions have now worsened at a different frequency.

Microstripes can also plot the field strength at a specified measuring point, and Figure 15 shows a number of such graphs plotted on the same axis, for different combinations of heatsink bonding.

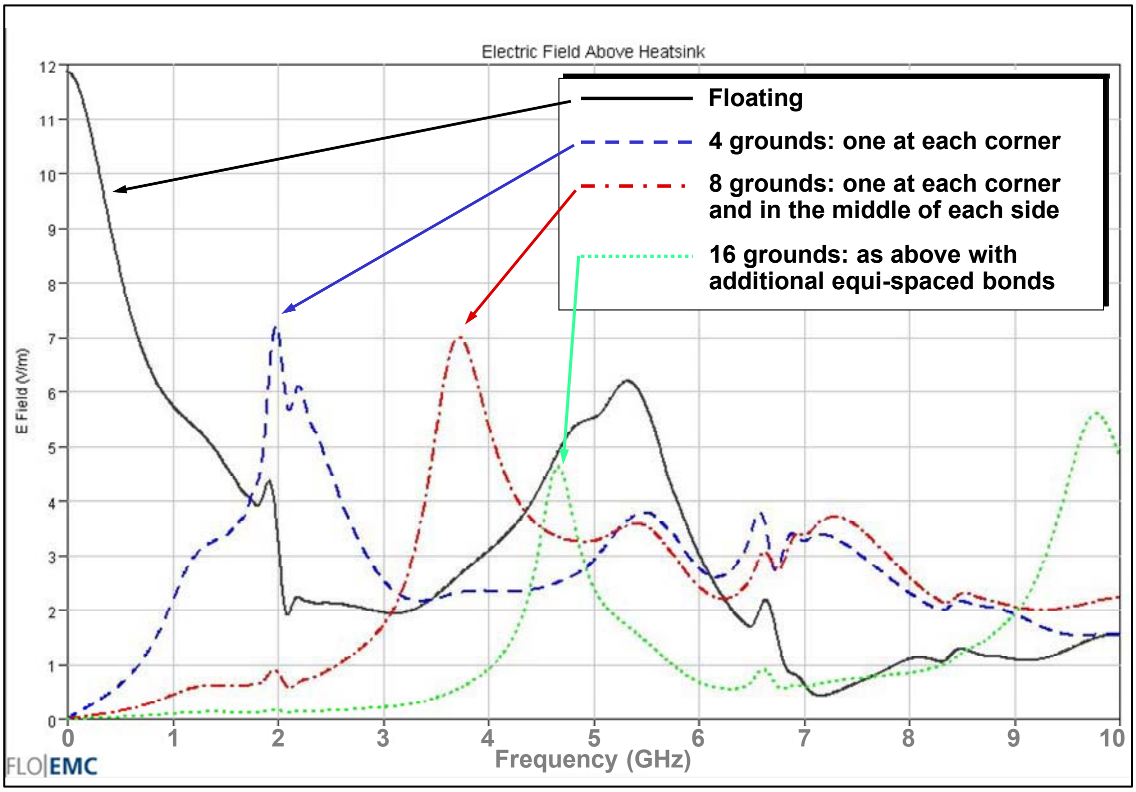

Figure 15 Variations in E-field emissions as heatsink RF-bonding is varied

The vertical axis of Figure 15 shows the far-field emissions at a specified point, but note that it uses a linear scale whereas most emissions measurements use a dB (i.e. logarithmic) scale. (I'm sure that dB scaling is a display option). It shows that the floating heatsink had high emissions at low frequencies, declining as frequency increased until peaking up again at its first resonance of 5.32GHz.

Direct bonding at all four corners had a very dramatic effect in reducing the emissions below 600MHz, but reduced the resonance frequency from 5.32 to 2.0GHz. Direct bonding at all four corners plus in the middle of each side of the heatsink – making 8 equally spaced RF-bonds in all – reduced the emissions significantly up to 2.5GHz, and caused the resonant frequency to increase to 3.75GHz.

Increasing the number of equally-spaced bonds to 16 reduced the emissions significantly all the way up to 4GHz, with the first resonance now shifted further up, to 4.7GHz. The level of emissions at frequencies below 2GHz is now very low indeed.

Figure 16 experiments with varying the bond resistance, starting with the direct corner-bonding of Figure 14 and then increasing the bond resistance.

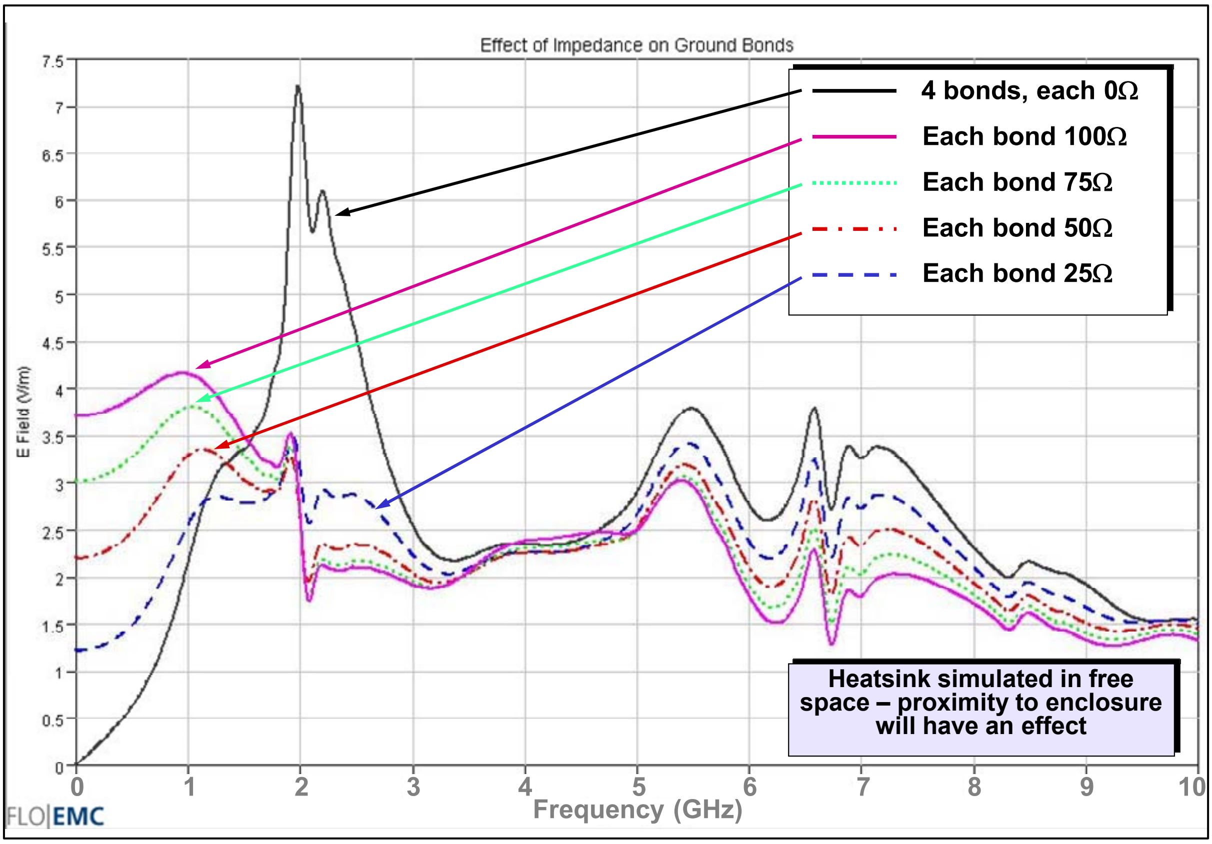

Figure 16 Variations in E-field emissions as the RF-bonding resistance is varied

Figure 16 shows that all of the resistance values used dampened down the first resonance, but the higher the resistance the worse the emissions became at low frequencies. Resistive bonding might be a useful technique when it was not practical to obtain a first (lowest) resonance above the highest frequency of concern – certain product designs may have to trade off between emission levels at resonant frequencies and the levels at low frequencies.

Paul created the above four graphics and made the 9 measurements, and emailed them to me while I was creating the EMC training course this article is based upon, over a couple of hours one evening. Imagine how long it would have taken to gain this kind of insight by making prototypes and testing them!

Modern electronic design is so complex and fast-paced that it requires the use of field solvers to save time and cost during design and development, delivering products to customers with less risk of time-to-market delays, and consequently lower corporate financial risks.

Figure 16 prompts the question of whether we could use "soft-ferrite RF suppressor beads" to our advantage, by choosing ferrites that have high (resistive) impedances above 1GHz to dampen down the first resonance, whilst having low impedances below 1GHz so as to minimise the emissions at the lower frequencies (but I didn't ask Paul to try this for me, because I felt I had already exceeded what I could ask as a favour).

8 Combining shielding with heatsinking

Shielding often restricts airflow and when devices have significant power dissipation can cause overheating. But shielding and heatsinking can be combined, on PCBs and in enclosures, as the examples in Figures 17 and 18 show...

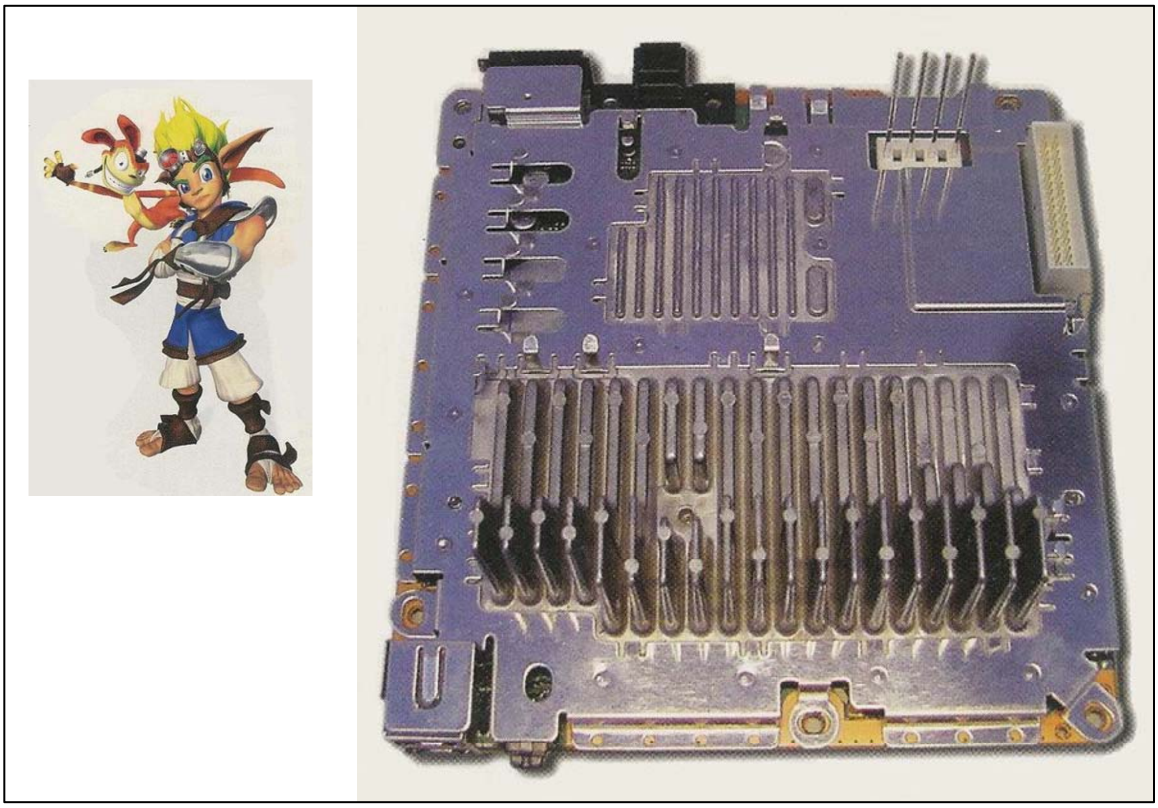

Figure 17 Example of Playstation 2, which uses a PCB-sized shield as its heatsink

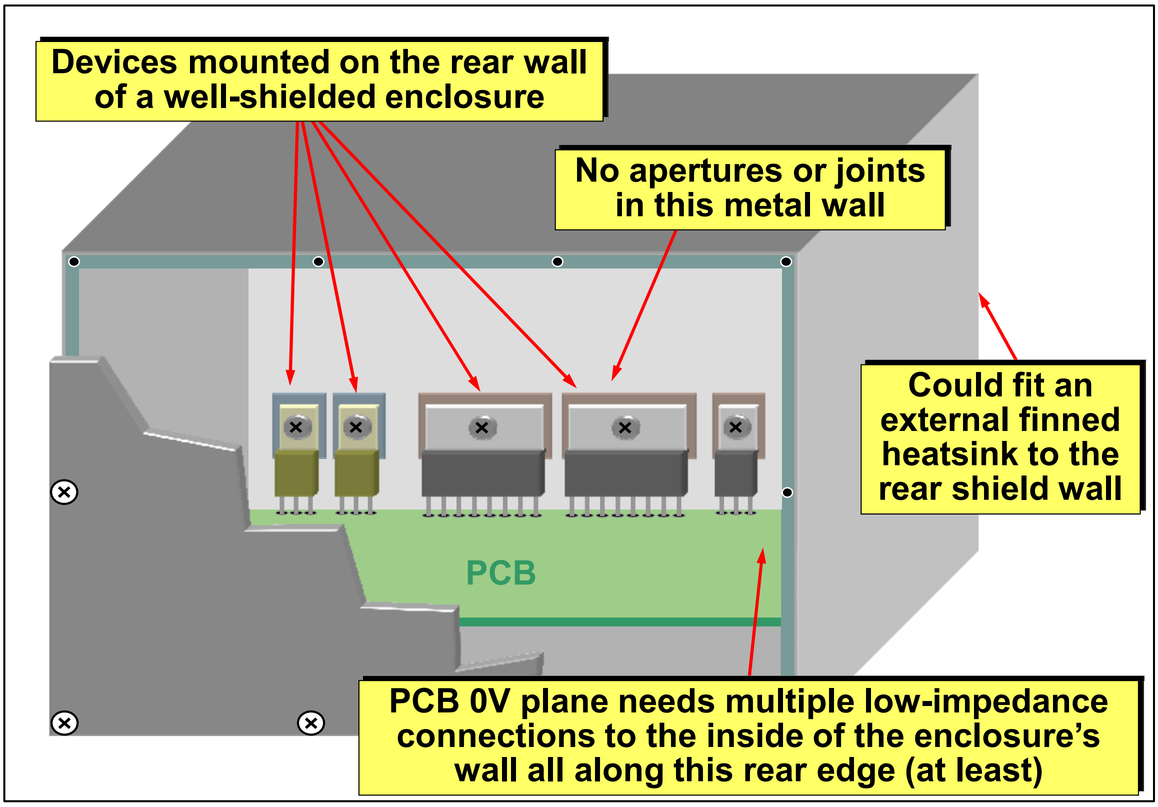

Figure 18 Example of combining shielding and heatsinking for an enclosure

The stray RF currents flowing on the inside surface of the rear panel of the example in Figure 18 can be very intense as they return via the edge of the PCB. Because these currents are so intense, it is best to ensure that there are no splits, joints, or other apertures in the rear panel, especially not around the ICs and between them and the RF-bonding to the 0V plane at the PCB's edge.

Good shielding effectiveness is required for Figure 18's overall enclosure, because there are many other possible paths for the stray heatsink currents that can have low impedances – especially at their series-resonant frequencies, or if the desired return path has any parallel resonances.

For this reason it is important to connect the 0V plane to the rear panel in a way that ensures a very short and very low impedance, non-resonant path up to f_{MAX}. One way to do this is to continue the 0V plane right to the rear edge of the PCB, then use the "edge-plating" technique to provide a continuous 0V connection all along the rear edge of the PCB. The rear panel must have a high-conductivity plated surface and connect to the edge-plated 0V strip on the PCB using a continuous strip of conductive EMC gasket all along the plated 0V edge of the PCB.

The plating of the PCB edge, inside surface or the rear panel and the gasket material, must all be galvanically compatible, given the anticipated lifecycle environment that the product will be used in.

It should be enough for the PCB's fixings to press the board's rear edge firmly into the gasket along its entire length, but a couple or three metal fixing brackets between the PCB (and its 0V plane) and the rear panel wouldn't go amiss (taking appropriate precautions against galvanic corrosion, as before).

Not every PCB manufacturer can do edge plating properly, however, and where there is some good reason why such manufacturers cannot be used, alternative means of achieving a reliable low-impedance RF-bond all along the rear edge of the PCB should be found. Remember, that it is the overall cost of manufacture that *really* matters, not the bare-board cost [6].

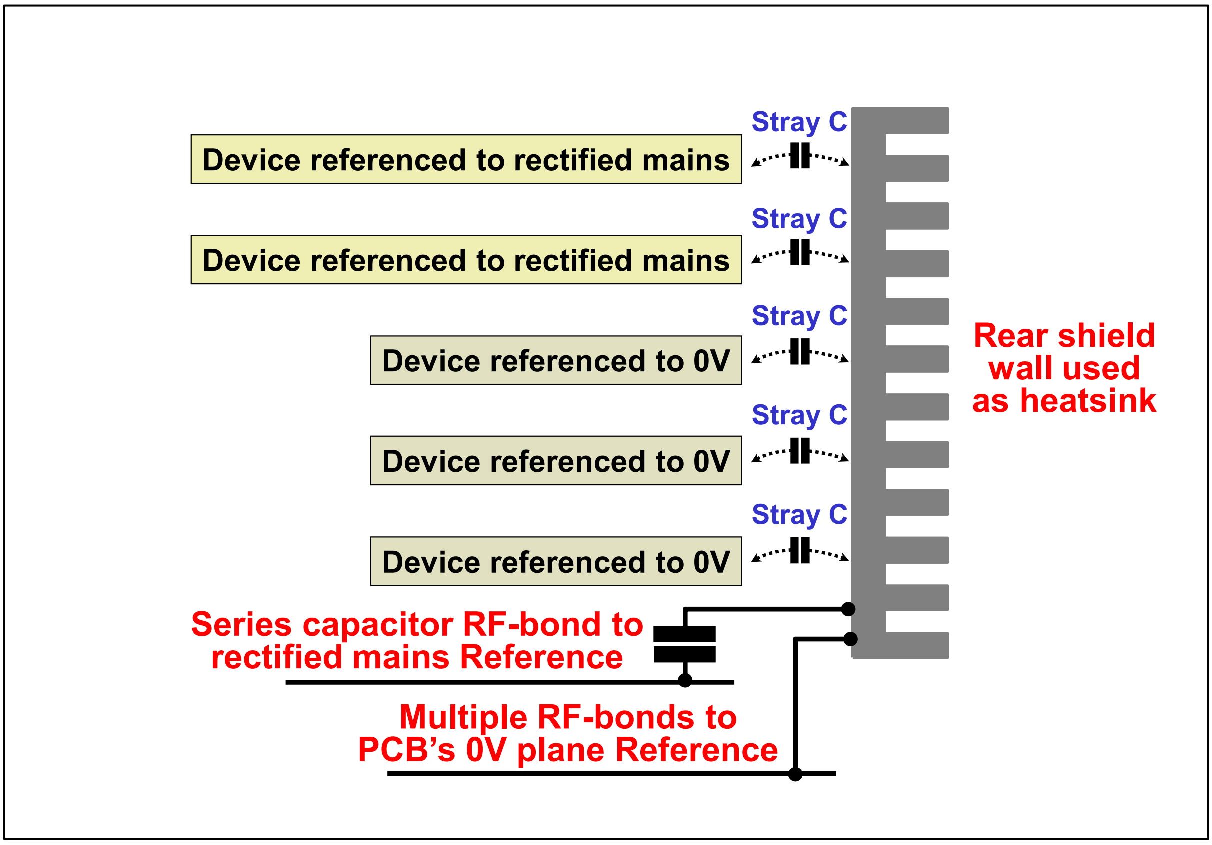

Figure 19 Where the stray heatsink currents flow in Figure 18

For the example shown in Figure 18, Figure 19 shows that the heatsunk ICs supplied by power rails referenced from the 0V plane have their stray currents returned directly to that plane.

The ICs that are powered from the AC mains via the bridge rectifier (typically the off-line mains switcher) have their stray currents pass through the 0V plane to the (safety-rated and (ideally) third-party safety agency approved) capacitors that connect to their hazardous voltage rails.

It is of course important to ensure that these capacitors don't increase the leakage current by too much and cause safety hazards (e.g. if a safety earth/ground wire breaks) or a failure to comply with the relevant safety standards.

Figure 20 shows how to combine shielding with heatsinking for a PCB-mounted part, along the lines of the PlayStation 2 in Figure 17. There are many different ways of creating the shielding wall around the IC, from soldered-in parts to conductively glued gaskets (possibly even form-in-place conductive elastomer, for a very skinny IC), to parts that are held in place by clips of housings and make contact by pressure alone.

It is important for the metal heatsink part to make a reliable electrical contact with the shielding wall all around, over the life of the product, so oxidation, galvanic and fretting corrosion at all the contact surfaces will all need to be taken into account.

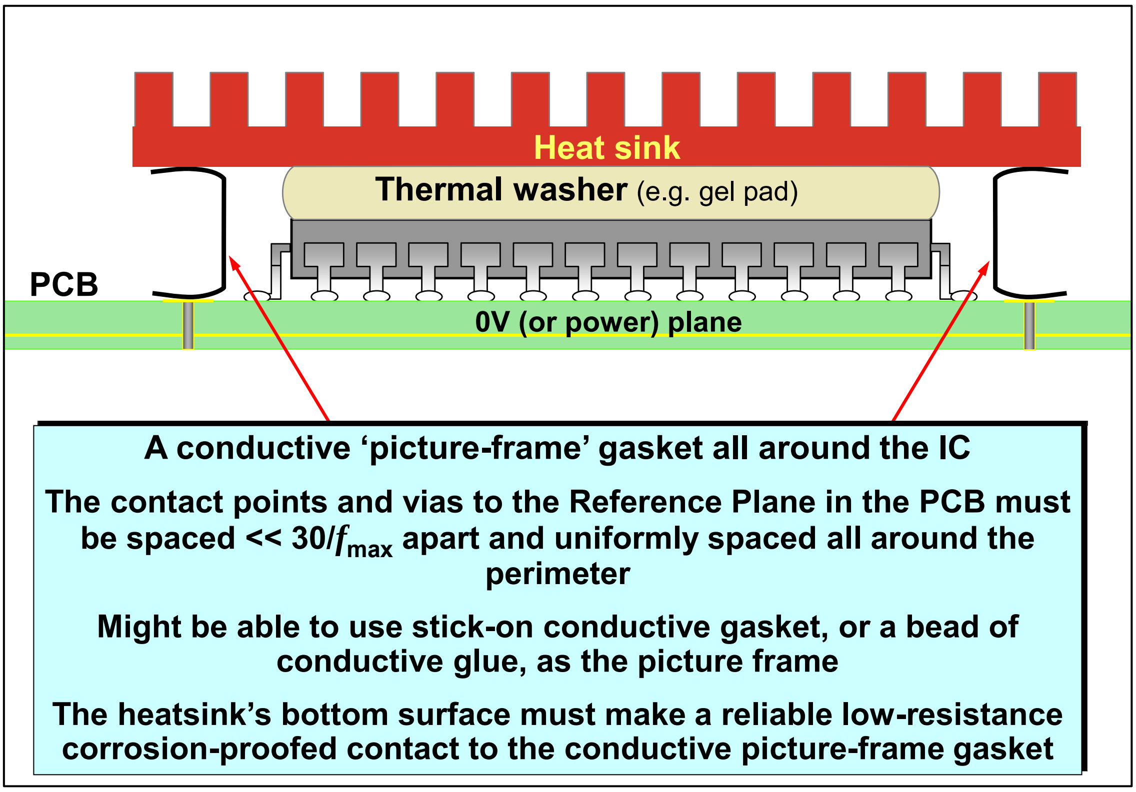

Figure 20 Example of combining shielding and heatsinking for an IC on a PCB

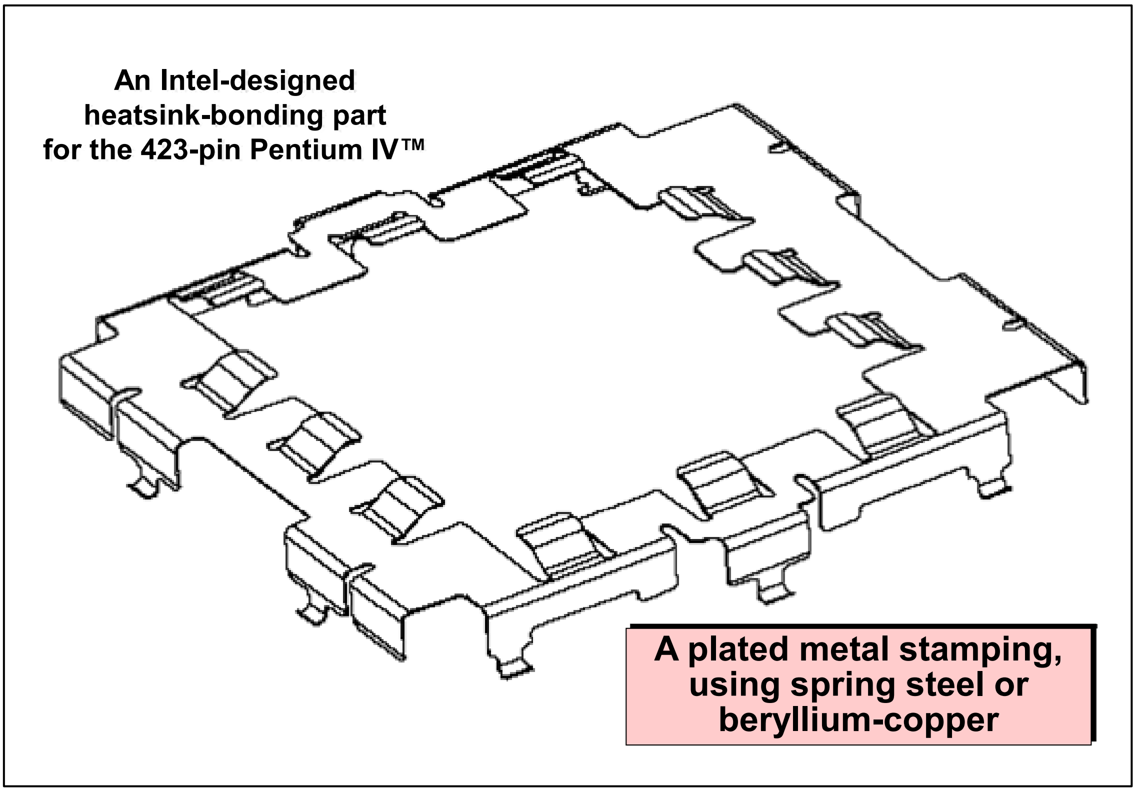

Figure 21 A heatsink-bonding part designed by Intel for the 423-pin Pentium IV™

Figure 21 shows a good example of a part designed as suggested by Figure 20. The part is stamped from a suitable springy metal and plated with a high conductivity metal (e.g. bright acid tin) if necessary. It is pick-and-placed over the IC and reflow soldered to the PCB at the same time as all the other components, then the heatsink is pressed down on top making contact with the spring fingers incorporated into the part.

As mentioned earlier, the metal base of the heatsink should have a high-conductivity plating/passivation with a metal that is galvanically compatible with the part, given the anticipated lifecycle environment of the product it is used in. A mechanical clip or other method of retention will be needed, to keep the heatsink pressed firmly down onto the part.

It is easy to design such parts yourself, and have them made and plated at low cost.

9 References

[1] "The Physical Basis of EMC", Keith Armstrong, The EMC Journal, in four Parts: Issues 85, 86, 87 and 88; November 2009, January, March and May 2010, www.theemcjournal.com

[2] "Design Techniques for EMC", Keith Armstrong, from the archives at www.theemcjournal.com:

Part 1) Circuit design (digital, analogue, switch-mode, comm's), and choosing components, Jan & March 2006

Part 2) Cables and connectors, May and July 2006

Part 3) Filtering and suppressing transients, Sept and Nov 2006, and Jan 2007

Part 4) Shielding (screening), March, May and July 2007

Part 5) PCB layout (including transmission lines), Sept and Nov 2007

Part 6) ESD, electromechanical devices, power factor correction, voltage fluctuations, immunity to power quality issues, Jan, March, May and July 2008

[3] "EMC for Printed Circuit Boards – basic and advanced design and layout techniques", Keith Armstrong, Nutwood UK February 2007, ISBN: 978-0-9555118-1-3 (perfect bound) or 978-0-9555118-0-6 (spiral bound), www.emcacademy.org/books.asp

[4] "EMC Testing", Tim Williams and Keith Armstrong, from the archives at www.theemcjournal.com:

Part 1) Radiated emissions, Feb 2001

Part 2) Conducted emissions, Apr 2001

Part 3) Fast transient burst, surge, electrostatic discharge, June 2001

Part 4) Radiated immunity, Aug and Oct 2001

Part 5) Conducted immunity, Oct 2001 and Jan 2002

Part 6) Low frequency magnetic fields emissions and immunity; plus mains dips, dropouts, interruptions, sags, brownouts and swells, Jan and March 2002

Part 7) Harmonics and flicker emissions, March 2002

[5] "Designing I/Os so they will not suffer from 'ground loop' currents in cable screens (shields)", Keith Armstrong, The EMC Journal, Issue 83, July 2009, pp 38-42

[6] "BOM cost and profitability", Keith Armstrong, The EMC Journal, May 2009, from the archives at www.theemcjournal.com