Characterisation of emissions due to power electronics heatsinks

Tim Williams
Elmac Services
Wareham, UK
timw@elmac.co.uk

***Abstract*— Heatsinks create a hazardous path for RF emissions from power electronic devices in switching converters. This paper describes a simple demonstration of the problem and discusses how the equivalent circuit can be modelled and quantified. Different heatsink configurations are seen to have a dramatic effect on the amplitude of conducted emissions.**

Keywords—conducted emissions, heatsinks, circuit simulation, power electronics

I. INTRODUCTION

Power electronics applications are widespread; they include mains-to-DC and DC-to-DC power supplies, variable speed drives, battery chargers and renewable energy power converters among many others. The importance of efficient power conversion is unquestioned, and for this application the high frequency switching converter is king. But it has a penalty in terms of generation of high frequency disturbance which must be addressed, both to satisfy regulatory requirements and to prevent degradation of the function of the system in which the converter is embedded. One route for emission of these disturbances is via the metalwork structures of the converter, and of these, the heatsink is a significant part.

This paper describes the problem posed by the simultaneous requirements to extract heat dissipation from a power converter and to prevent the emission of HF disturbances. It goes on to look at a simple demonstration of the problem and discusses how the equivalent circuit can be modelled and quantified, so that the disturbances can be minimised in the most effective way.

The heatsink's relationship to EMC has already been studied and modelled in some depth with respect to radiated emissions from VLSI ICs [1][2][3][4][5] and from power electronics [6]. However, these studies concentrate on frequencies for which the heatsink structure is resonant, i.e. hundreds of MHz up to several GHz. For many power electronics applications the problems start much lower, tens to hundreds of kHz up to several MHz, where there is no question of structural resonance. Even so, the effect of the heatsink is still dramatic and it can be modelled in a more simplified fashion using equivalent circuits, which makes understanding the problem from the point of view of the product designer that much easier. This is the subject of this paper.

II. THE HEATSINK PROBLEM

Switchmode power converters have one or more switching devices, such as MOSFETs or IGBTs, which are usually mounted on heatsinks to deal with their dissipation and exhaust the heat created to their near environment. Unfortunately their heatsink tab will typically be closely coupled to a node which carries the full supply voltage at the switching frequency, such as the drain terminal of a MOSFET – indeed the tab often *is* the drain. The external heatsink is in turn mechanically and thermally coupled to this tab. Even if there is no intentional electrical connection, capacitance between the device tab and its heatsink will cause the heatsink to carry some of the switching voltage. But, being a large metal component, its capacitance to its environment then creates an effective path out of the product for the switching noise.

Unfortunately again, as Sochoux *et al* point out, "there is little or no design input from EMC engineers during the initial design phase of heatsinks. Most of the EMC designs for heatsinks are left to pure luck. Mechanical engineers design heat sinks based on thermal and mechanical constraints. The EMC team only becomes involved ... at the end of the design cycle". [1]

It is possible, though, to include the heatsink's electrical properties at the start of the design. We can create a simplified equivalent circuit of the path for conducted emissions as in Figure 1.

The principal components are the switching device itself, its load, the path back to the incoming power supply including any filtering, and the line impedance stabilising network

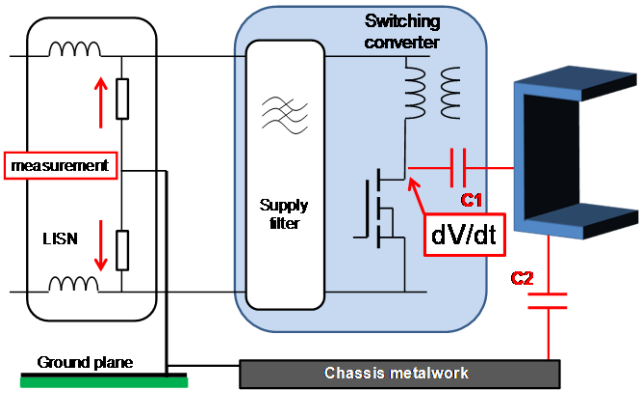

Figure 1 Equivalent circuit for conducted emissions

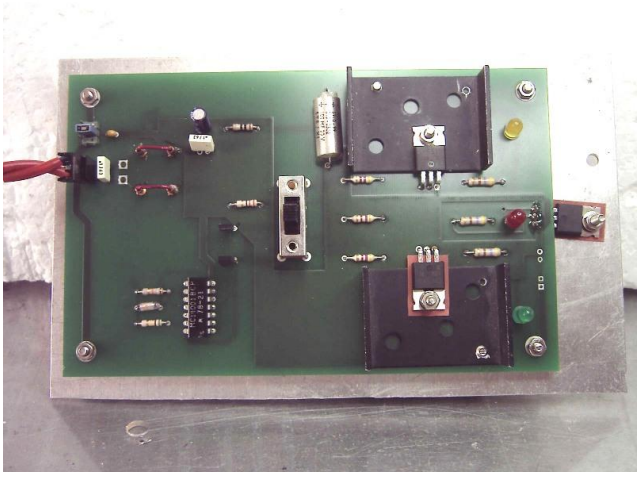

Figure 2 Photo of Model circuit

(LISN) across which a conducted emissions measurement is made. A further crucial item is the metal chassis, which is typically connected via a safety earth wire to the measurement ground reference. If there is no chassis, stray capacitance directly to the measurement ground plane must be included.

The circuit node which typically carries the highest dV/dt is the drain of the switching device. Any capacitive coupling to this point will create a severe emissions threat. But to extract heat from the device, the design will mount it on a heatsink, and this enhances parasitic capacitive coupling, from the device tab to the heatsink (C1 in Figure 1) and from the heatsink to the chassis (C2 in Figure 1). This makes for an effective path for emissions into the power supply connection.

We can deal with this path by controlling C1 and C2, and by returning the current they create to the right place.

III. THE MODEL CIRCUITS

To demonstrate the effect of different heatsink connections, and to create a simplified circuit for modelling purposes, a pair of circuits have been constructed (Figure 2 shows one of these).

The switching circuit is reduced to its essentials, which are an oscillator running at around 80kHz, driving into a simulated power switcher. The switching device is an ordinary TO220 package bipolar transistor which has a resistive load, for one version of the model, or an ordinary MOSFET with the same resistive load for the other. Both types of device have their collector or drain terminal directly connected to the package tab. The switching frequency can be applied to one of three configurations of heatsink. In case 1, the device tab is connected to the heatsink which is isolated from everything else. In case 2, the tab is isolated from the heatsink which is connected to 0V, not measurement ground. In case 3, the device is mounted directly on the chassis, which is connected to measurement ground as would be the case in a safety class I product, but the device tab is of course electrically isolated from it. Simplified circuits of each of the three cases are shown in Figure 3.

The 10V DC supply to the circuit is taken through a measurement LISN so that a conducted emissions measurement can be done in exactly the same way as on a

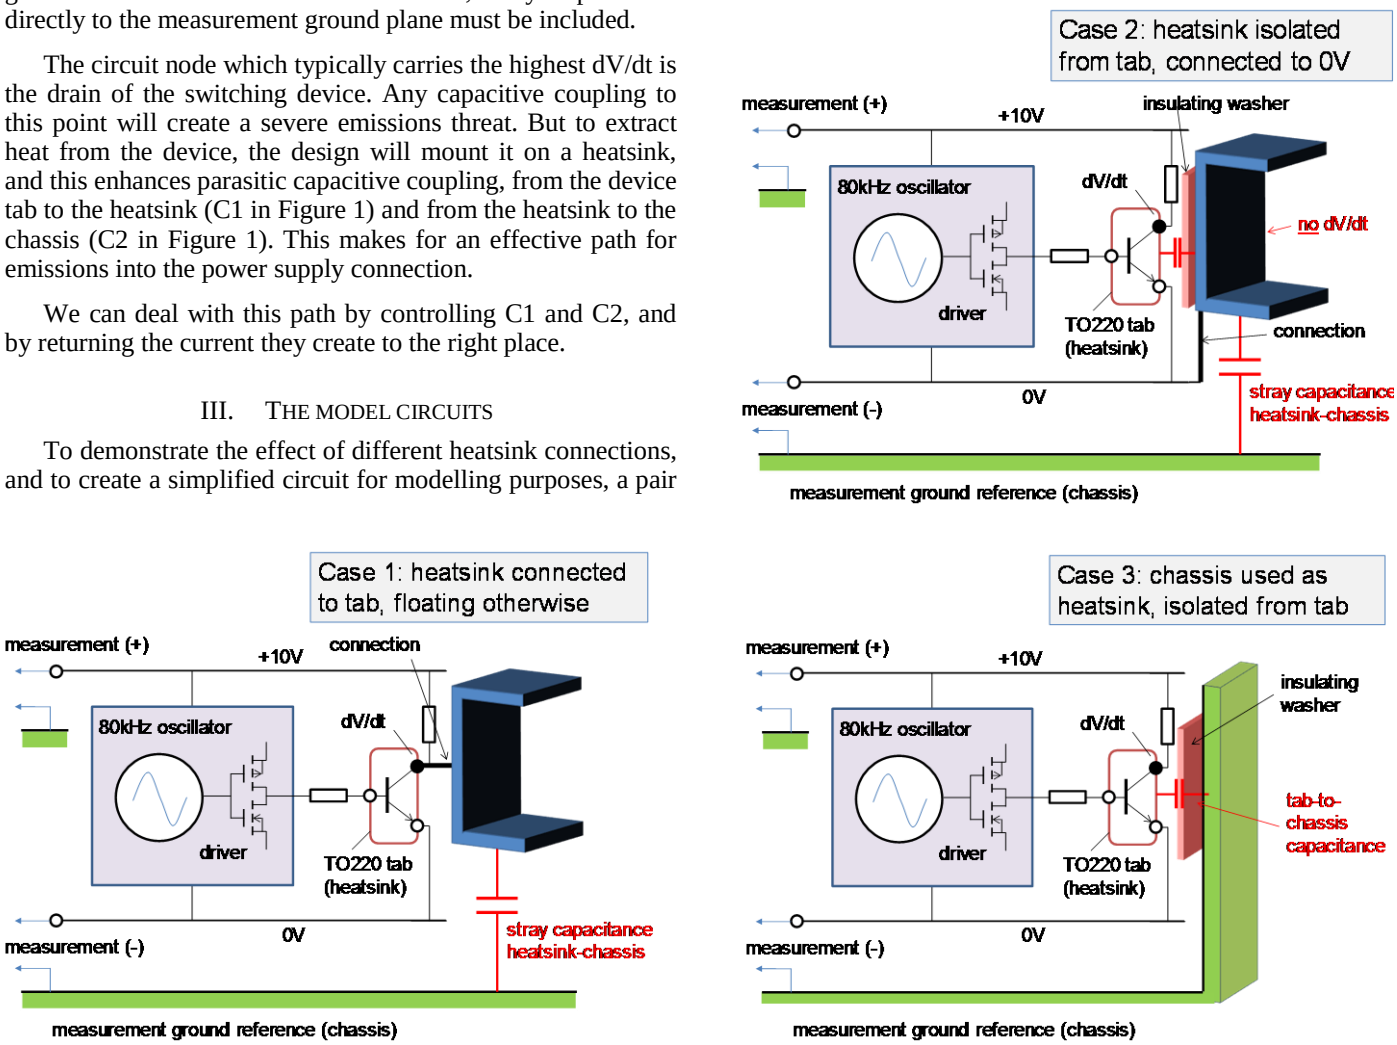

Figure 3 Model circuit configurations (bipolar version shown)

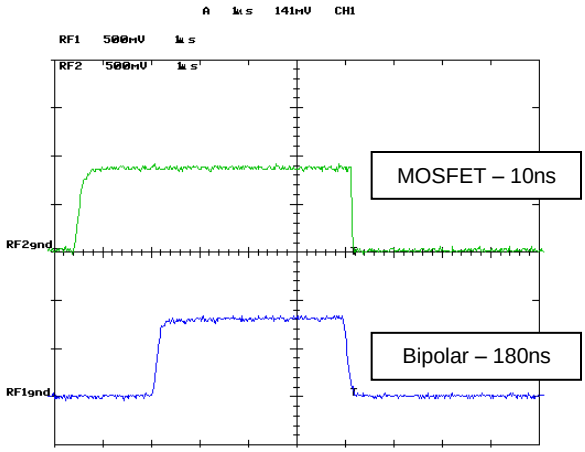

Figure 4 Switching waveforms

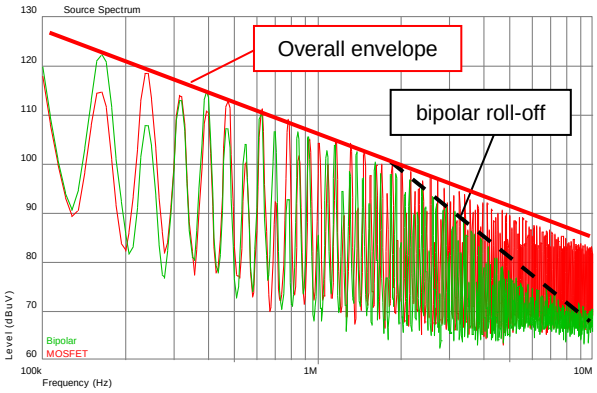

Figure 5 Measured spectrum of switching waveform

power supply input; in the examples presented here the LISN is the CISPR 25 5µH/50Ω version. The same principles of course would apply to a mains 50µH/50Ω LISN, where the DC +10V would be equivalent to live, and the DC 0V equivalent to neutral. Good differential mode filtering is applied between +10V and 0V to minimise the signal that is measured in differential mode; initially, no common mode filter is used.

The waveform of the signal at the collector or drain terminal of each transistor is shown in Figure 4; the peak voltage swing is 8V and the measured fall times (rise times are slower) are 180ns and 10ns respectively. This waveform appears directly on the TO220 tab which is mounted on the heatsink.

IV. PREDICTED COMMON MODE EMISSIONS

The spectrum of the switching waveform can be measured directly, with a calibrated CISPR-16 voltage probe across the switching transistor, as presented in Figure 5 up to 10MHz. It could also be derived from the waveform, of course, by Fourier analysis. Note how the low-frequency content is largely identical for both devices, but the MOSFET's spectrum extends at a much higher level above a few MHz. The

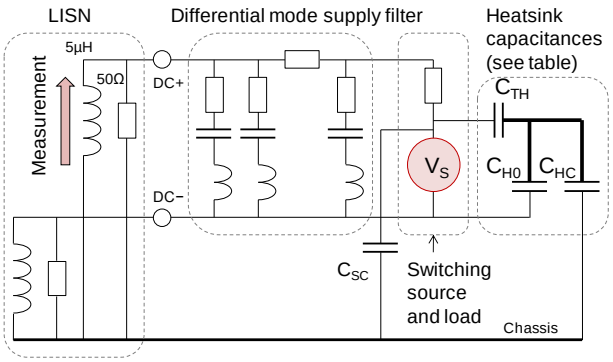

Figure 6 Simulation circuit

MOSFET spectrum envelope (red line) decays at a constant 20dB/decade across this range from a maximum at 80kHz of 130dBµV. The bipolar spectrum envelope (black dotted) has the same 80 kHz amplitude but decays at −40dB/decade above 2MHz.

The coupling circuits of Figure 3 require a knowledge of the relevant parasitic capacitances. For cases 2 and 3 where the TO220 device is mounted to heatsink or chassis via a silicone washer, the capacitance through this washer can be easily calculated or measured, and for the parts used in this demonstration it is 18pF. The capacitance of the heatsink to chassis in cases 1 and 2 can be estimated from the formula for overlapping plates in air: 0.0885*A/d pF, where A is the area of the overlap and d is the plate separation. In this case the area of the metal is 12cm^2 and its separation from the chassis is 7mm, which yields 1.5pF. We can ignore the dielectric constant of the intervening PCB, since most of the separation distance is through air.

We can simulate the overall common mode circuit, including these parasitic capacitances and the LISN impedance, and this simulation circuit is shown in Figure 6. Additional parasitic capacitances are relevant, the first is shown as C_{H0} in Figure 6; this is the capacitance from the heatsink to circuit 0V (DC-), which is a function of the layout of each of the three cases. In Case 1, the overlapping area between the heatsink metal and the circuit 0V trace on the other side of the board is 1.2cm^2, and with the board thickness of 1.6mm and ε_r of 4.5 this gives around 3pF. In Case 3, the chassis metalwork is the heatsink and this couples with the whole of the PCB 0V trace, but through the separation distance board-to-chassis of 7mm, which yields around 6pF.

In the practical circuit, there is another capacitance which is modeled as C_{SC}, the capacitance from source to chassis. This is shown as appearing from the switching device tab to the chassis in Figure 6. In fact this is the amalgamation of all the stray capacitances from the driving circuit and switching device to the chassis metalwork, *excluding* the heatsink path. It is responsible for limiting the variations due to heatsink configuration, and is estimated at 1pF.

The capacitances are tabulated in table 1 for the three configurations, with reference to the figure.

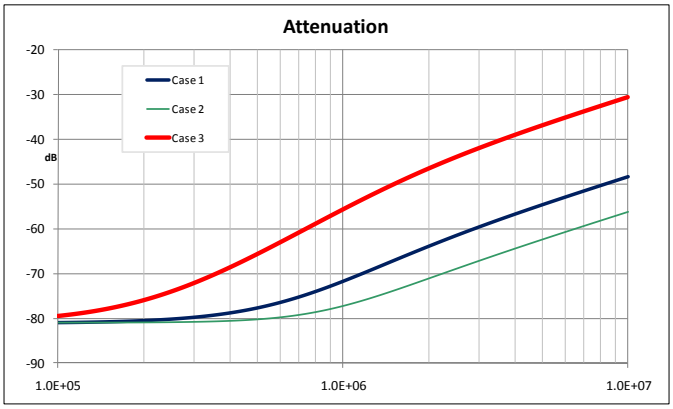

Figure 7 Simulated attenuation from tab to LISN

Table 1 Parasitic capacitances

	Case 1	Case 2	Case 3
C_{TH} tab-heatsink (C1)	Link	18pF	18pF
C_{H0} heatsink-DC 0V	3pF	Link	6pF
C_{HC} heatsink-chassis (C2)	1.5pF	1.5pF	Link
C_{SC} source-chassis	1pF	1pF	1pF

The practical circuit has a 470 ohm resistor plus LED in series with the drain or collector of the switching device; this is shown as R_L in Figure 6. The switching device itself is regarded as a pure voltage source with the waveform as given in Figure 4, shown as V_S.

The WinSpice [7] simulated attenuation in dB between the heatsink tab (V_S in Figure 6) and the measurement point at the LISN, for the different configurations is shown in Figure 7. To arrive at the expected emissions voltages for the three cases, we can combine this by simple addition with the envelope of the driving voltage spectrum (red and dotted black lines, derived from Figure 5). The resulting envelopes of the expected emissions are shown in Figure 8.

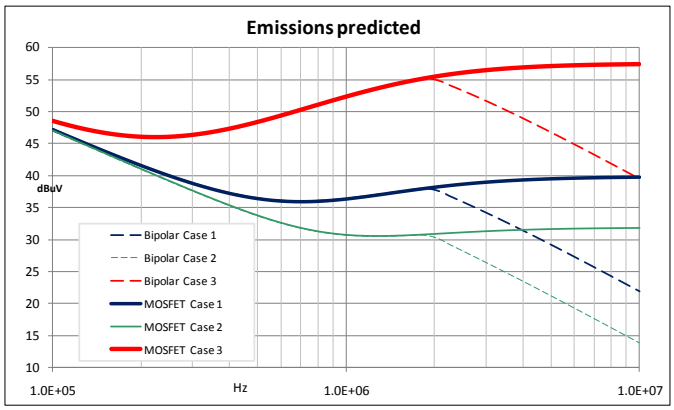

Figure 8 Predicted emissions envelope

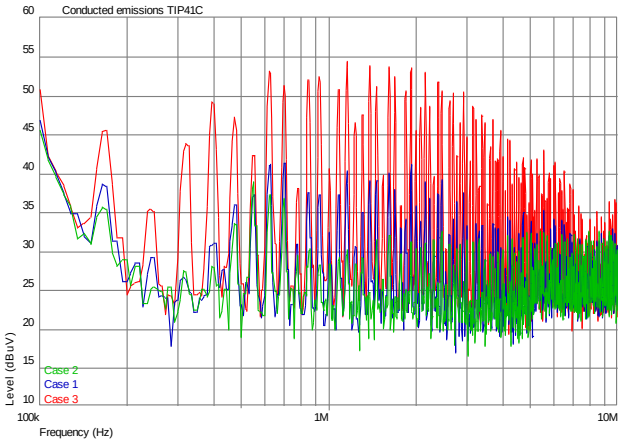

Figure 9 Bipolar measured emissions

V. MEASURED EMISSIONS

The measured emissions from the bipolar circuit are shown in Figure 9. Although the bipolar transistor is little used in power applications, its spectrum is similar to that of the IGBT, which has switching speeds of the same order, and is very common in many power applications. Configuration 1 (blue) is with heatsink connected to tab but floating. The capacitance of the heatsink is doing most of the coupling, as can be demonstrated by temporarily touching the heatsink; the levels go up, since the capacitance has gone up. With configuration 2 (green), the heatsink is now isolated from the tab but connected to 0V. Although its capacitance to the chassis hasn't changed, it no longer carries the switching voltage, and the levels have dropped significantly, by 6dB or more. Touching the heatsink makes no difference.

Configuration 3 (red) uses the chassis, which in this case is connected to measurement reference ground, as the heatsink. There is now a high capacitance (as per Table 1) between the noise source, that is the tab, and the measurement ground reference. The emissions have increased dramatically. The difference between this worst case, and the best case, is over 20dB across most of the spectrum.

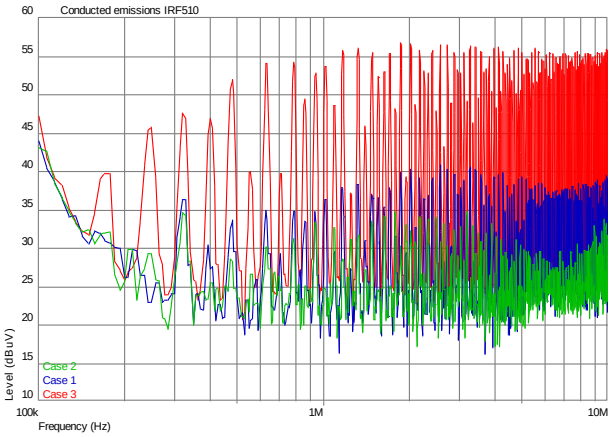

Figure 10 MOSFET measured emissions

Turning to the MOSFET circuit (Figure 10) – which is more representative of the majority of medium-power applications – we can see that the high frequency spectrum is greatly extended, due to the faster fall times of the switch. But the general relationship between the different heatsink configurations still holds, since the layout parasitics are the same. The consequence is simply that emissions above around 3MHz become substantially more significant.

For both devices, we can see that the correlation between measured emissions and those predicted by WinSpice circuit simulation is very good. This circuit is of course simple by comparison with real products; for instance, the output circuitry has been completely ignored. The interaction between common mode and differential mode components of the emissions has been deliberately suppressed; the heatsink doesn't contribute to differential mode. The demonstration shown here becomes more complex above 10MHz where additional factors, particularly contributions from other parts of the circuit as well as the effects of component inductance and PCB layout, become significant. Nevertheless it shows that predicting a conducted emissions level is fairly straightforward if the relevant parasitics are known or can be reliably estimated.

VI. Filtering the Supply

The model so far has not used the conventional common mode filter Y-capacitor on the DC terminals to chassis, but its effect can easily be predicted by adding this in (with its attendant parasitic inductance) to the simulation circuit. The practical units both have this capability and so the effect of a filter capacitor can be demonstrated both by the simulation and by measurement. Figure 11 repeats Figure 6 with a Y-capacitor added between DC 0V and chassis, at the terminals of the unit. Two values were used; a 3.3nF ceramic, and a 22nF polyester in parallel with the 3.3nF giving 25nF. The smaller capacitor had an estimated self-inductance of 30nH while the larger was 40nH, figures derived from the measured total lead and track length at 10nH/cm.

At frequencies where the LISN impedance is less than the capacitor impedance there is little effect; but at higher frequencies the capacitive divider formed by the Y-capacitor and the combined source capacitance reduces the emissions in exactly the expected manner. The simulation results for case 3 (Figures 12 and 13) again show very good agreement with the

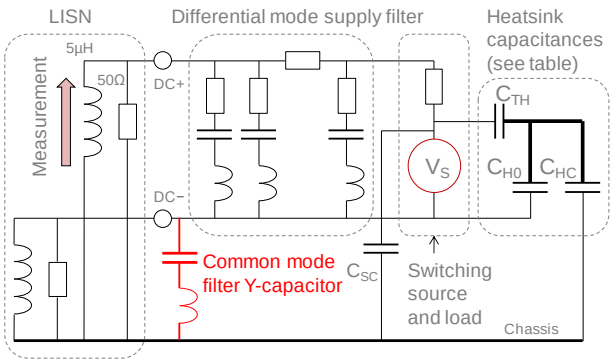

Figure 11 Simulation circuit with Y-capacitor

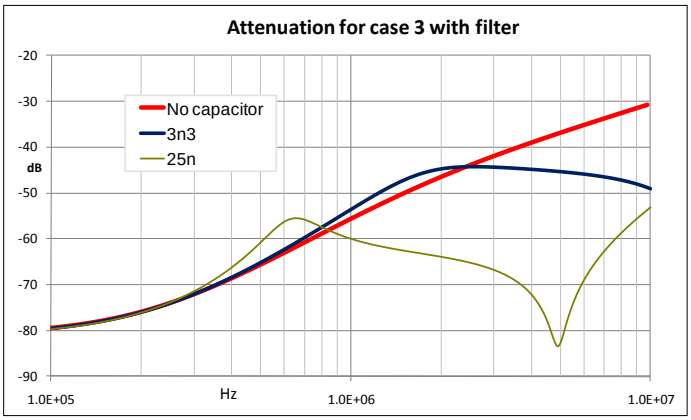

Figure 12 Modelled Y-cap filter attenuation – Case 3

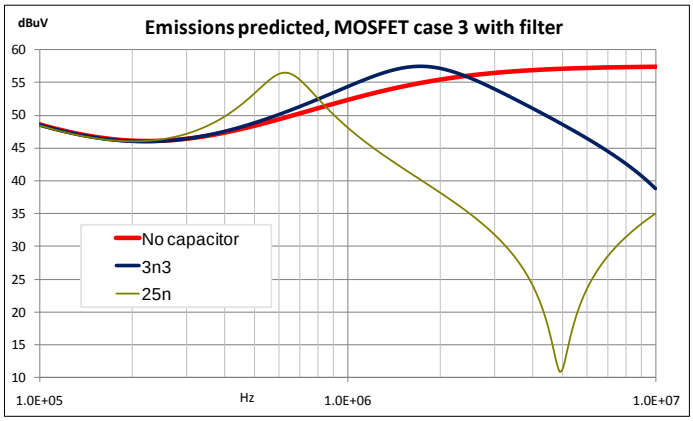

Figure 13 Predicted emissions envelope – filtered, MOSFET Case 3

measurements (Figure 14); it is noteworthy that for the 25nF capacitor the increase in the 600kHz area due to resonance with the LISN's 5µH inductor, and the drop in the 5MHz area due to the series resonance of the capacitor with its lead inductance, are both clearly modeled and clearly shown in the measurements.

The Y-capacitor is an important component, but its value is

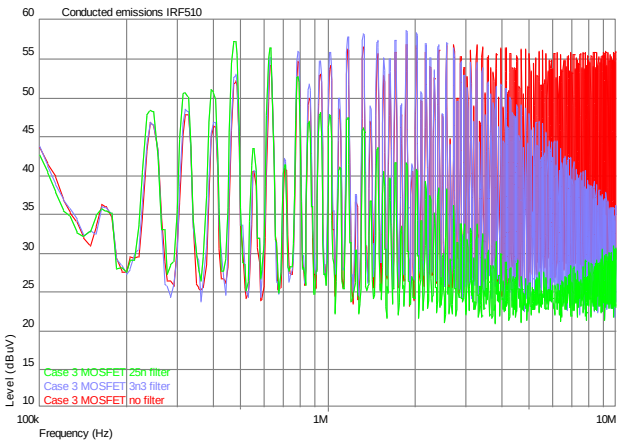

Figure 14 Measured emissions levels – Filtered, MOSFET Case 3

usually limited at least in mains-powered applications by the maximum allowed safety leakage current to earth. Therefore keeping the source capacitance to a minimum effectively increases the attenuation offered by the highest available Y-cap value.

VII. CONCLUSION

The product designer needs to be careful to map out the coupling paths between any areas of high dV/dt in the circuit, and the emissions measurement ground reference. Very often, in a power conversion circuit, a heatsink will be found in this path. If the heatsink is electrically separate from the chassis, it should be connected to a benign DC voltage, typically circuit 0V, and not left floating or directly coupled to the switching device. And a crucial further rule is, never use the chassis as a heatsink with a switching device directly mounted to it; or if you must, and after all for good thermal reasons a lot of designs do, then expect to have to take substantial extra measures in filtering. Or, find a way, such as an additional E-field screening foil, to neutralize the capacitive coupling to the chassis without compromising the thermal coupling.

To help with the design, it is possible to make a reasonably accurate estimation of the expected conducted emissions using circuit simulation techniques, provided that the relevant parasitic components are known. This allows other methods that can minimise the impact of the unwanted capacitance to be investigated: draw out the equivalent circuit, and this should provide clues as to the possibilities.

References

[1] Philippe Sochoux, Jinghan Yu, Alpesh Bhobe, Federico Centola, "Heat sink design flow for EMC", Design Con 2008, IEC publications, pp: 1 - 27.

[2] Colin E. Brench, "Heat sink radiation as a function of geometry", IEEE Int Symp EMC 1994, pp: 105-109.

[3] "Optimization of Heat Sink EMI Using Design of Experiments with Numerical Computational Investigation and Experimental Validation", S.Manivannan, R.Arumugam, S.Prasanna Devi, S.Paramasivam, P.Salil, B.Subbarao, IEEE Int Symp EMC 2010 pp 295-300

[4] "An Overview of Chip Level EMC Problems", Sergiu Radu, Sun Microsystems, various IEEE presentations 2010

[5] Richard Georgerian, Mark Montrose, "Product Safety and the Heat Sink – Dilemma of Minimizing Radiated Emissions and Maximizing Thermal Cooling", IEEE Int Symp EMC 2003 pp 134-137

[6] "FDTD Modeling of Heatsink RF Characteristics for EMC Mitigation", Nick J. Ryan, Barry Chambers, D. A. Stone, IEEE Trans EMC, Vol. 44, NO. 3, Aug 2002 pp 458-465

[7] www.winspice.com