Design Techniques for EMC – Part 1
Circuit Design, and Choice of Components

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the first in a series of six articles on best-practice EMC techniques in electrical/electronic/mechanical hardware design, to be published in this journal over the following year. The series is intended for the designer of electronic products, from building block units such as power supplies, single-board computers, and "industrial components" such as motor drives, through to stand-alone or networked products such computers, audio/video/TV, instruments, etc.

These articles were first published in the EMC Journal as a series during 1999. This version includes a number of corrections, modifications, and additions, many of which have been made as a result of correspondence with the following, to whom I am very grateful: Feng Chen, Kevin Ellis, Neil Helsby, Mike Langrish, Tom Liszka, Alan Keenan, T Sato, and John Woodgate. I am also indebted to Tom Sato for translating these articles into Japanese and posting them on his website: http://member.nifty.ne.jp/tsato/, as well as suggesting a number of improvements.

The techniques covered in these six articles are:

1) **Circuit design (digital, analogue, switch-mode, communications), and choosing components**
2) Cables and connectors
3) Filters and transient suppressors
4) Shielding
5) PCB layout (including transmission lines)
6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), so this magazine article format can do no more than introduce the various issues and point to the most important of the best-practice techniques.

Before starting on the above list of topics it is useful see them in the context of the ideal EMC lifecycle of a new product design and development project.

The project EMC lifecycle

The EMC issues in a new project lifecycle are summarised below:

- Establishment of the target electromagnetic specifications for the new product, including:
  - The electromagnetic environment it must withstand (including continuous, high-probability, and low-probability disturbance events) and the degradation in performance to be allowed during disturbance events;
  - Its possible proximity to sensitive apparatus and allowable consequences, hence the emissions specifications;
  - Whether there are any safety issues requiring additional electromagnetic performance specifications. Safety compliance is covered by safety directives, not by EMC Directive;
  - All the EMC standards to be met, regulatory compliance documentation to be created, and how much "due diligence" to apply in each case (consider all markets, any customers' in-house specifications, etc.).
- System design:
  - Employ system-level best-practices ("bottom-up");
  - flow the "top-level" EMC specifications down into the various system blocks ("top-down").

- System block (electronic) designs:
  - Employ electrical/electronic hardware design best-practices ("bottom-up") (covered by these six articles);
  - Simulate EMC of designs prior to creating hardware, perform simple EMC tests on early prototypes, more standardised EMC tests on first production issue.
- Employ best-practice EMC techniques in software design.
- Achieve regulatory compliance for all target markets.
- Employ EMC techniques in QA to control:
  - All changes in assembly, including wiring routes and component substitutions;
  - All electrical/electronic/mechanical design modifications and software bug-fixes;
  - All variants.
- Sell only into the markets originally designed for;
  - To add new markets go through the initial electromagnetic specification stage again.
- Investigate all complaints of interference problems
  - Feed any resulting improvements to design back into existing designs and new products (a corrective action loop).

This may look quite daunting, but it is only what successful professional marketeers and engineers already know to do, so as not to expose their company to excessive commercial and/or legal risks.

As electronic technology becomes more advanced, more advanced management and design techniques (such as EMC) are required. There is no escaping the ratcheting effects of new electronic technologies if a company wants to remain profitable and competitive. But new electronics technologies are creating the worlds largest market, expected to exceed US$1 trillion annually in value (that's $1 million million) within a couple of years and continue to increase at 15% or so per annum after that. Rewards are there for those that can take the pace.

The following outlines a number of the most important best-EMC-practices. They deal with "what" and "how" issues, rather than with why they are needed or why they work. A good understanding of the basics of EMC is a great benefit in helping to prevent under or over-engineering, but goes beyond the scope of these articles.

Table of contents for Part 1

1. Circuit design and choice of components for EMC

Correct choice of active and passive components, and good circuit design techniques used from the beginning of a new design and development project, will help achieve EMC compliance in the most cost-effective way, reducing the cost, size, and weight of the eventual filtering and shielding required.

These techniques also improve digital signal integrity and analogue signal-to-noise, and can save at least one iteration of hardware and software. This will help new products achieve their functional specifications, and get to market, earlier. These EMC techniques should be seen as a part of a company’s competitive edge, for maximum commercial benefit.

1.1 Digital components and circuit design for EMC

1.1.1 Choosing components

Most digital IC manufacturers have at least one glue-logic range with low emissions, and a few versions of I/O chips with improved immunity to ESD. Some offer VLSI in “EMC friendly” versions (some “EMC” microprocessors have 40 dB lower emissions than regular versions).

Most digital circuits are clocked with squarewaves, which have a very high harmonic content, as shown by Figure 1.

Figure 1
The spectrum of an ideal 60 MHz squarewave with 1ns rise and fall times

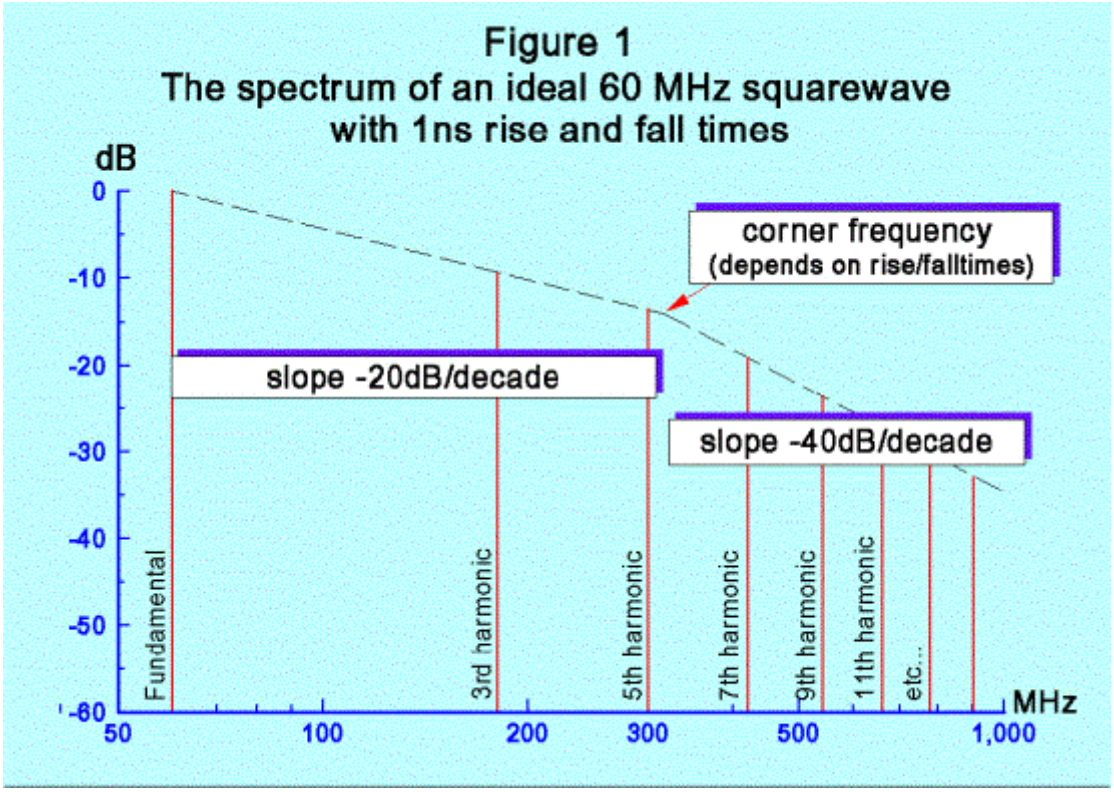

The faster the clock rate, and the sharper the edges, the higher the frequency and emissions levels of the harmonics.

So always choose the slowest clock rate, and the slowest edge rate that will still allow the product to achieve its specification. Never use AC when HC will do. Never use HC when CMOS 4000 will do.

Choose integrated circuits with advanced signal integrity and EMC features, such as:

- Adjacent, multiple, or centre-pinned power and ground.

  Adjacent ground and power pins, multiple ground and power pins, and centre-pinned power and ground all help maximise the mutual inductance between power and ground current paths, and minimise their self-inductance, reducing the current loop area of the power supply currents and helping decoupling to work more effectively. This reduces problems for EMC and ground-bounce.

- Reduced output voltage swing and controlled slew rates.

  Reduced output voltage swing and controlled slew rates both reduce the dV/dt and dI/dt of the signals and can reduce emissions by several dB. Although these techniques improve emissions, they could worsen immunity in some situations, so a compromise may be needed

- Transmission-line matching I/Os.

  ICs with outputs capable of matching to transmission-lines are needed when high-speed signals have to be sent down long conductors. E.g. bus drivers are available which will drive a 25Ω shunt-terminated load. These will drive 1 off 25Ω transmission line (e.g. RAMBUS); or will drive 2 off 50Ω lines, 4 off 100Ω lines, or 6 off 150Ω lines (when star-connected).

- Balanced signalling.

  Balanced signalling uses ± (differential) signals and does not use 0V as its signal return. Such ICs are very helpful when driving high-speed signals (e.g. clocks > 66MHz) because they help to preserve signal integrity and also can considerably improve common-mode emissions and immunity.

- Low ground bounce.

  ICs with low ground-bounce will generally be better for EMC too.

- Low levels of emissions.

  Most digital IC manufacturers offer glue-logic ranges with low emissions. For instance ACQ and ACTQ have lower emissions than AC and ACT. Some offer VLSI in “EMC friendly” versions, for example Philips have at least two 80C51 microprocessor models which are up to 40dB quieter than their other 80C51 products.

- Non-saturating logic preferred.

  Non-saturating logic is preferred, because its rise and fall times tend to be smoother (slew-rate controlled) and so contain lower levels of high-order harmonics than saturating logic such as TTL.

- High levels of immunity to ESD and other disturbing phenomena.

  Serial communications devices (e.g. RS232, RS 485) are available with high levels of immunity to ESD and other transients on their pins. If their immunity performance isn’t specified to at least the same standards and levels that you need for your product, additional suppression components will be needed.

- Low input capacitance.

  Low input capacitance devices help to reduce the current peaks which occur whenever a logic state changes, and hence reduce the magnetic field emissions and ground return currents (both prime causes of digital emissions).

- Low levels of power supply transient currents.

Totem-pole output stages in digital ICs go through a brief period when both devices are on, whenever they switch from one state to the other. During this brief period the supply rail is shorted to 0V, and the power supply current transient can exceed the signal's output current. Both the transient current (sometimes called the 'shoot-through' current) and the voltage noise it causes on the power rails are prime causes of emissions. Relevant parameters may include the transient current's peak value, its d*I*/d*t* (or frequency spectrum) and its total charge, any/all of which can be important for the correct design of the power supply's decoupling. ICs with *specified* low levels of power supply transients should be chosen where possible.

- Output drive capability no larger than need for the application.

  The output drive current of an IC (especially a bus driver) should be no larger than is needed. Drivers rated for a higher current have larger output transistors, which can mean considerably larger power supply transients. Their increased drive capability can also mean that the traces they drive can experience faster rise and falltimes than are needed, leading to increased overshoot and ringing problems for signal integrity as well as higher levels of RF emissions.

All of the above should have guaranteed minimum or maximum (as appropriate) specifications (or at least typical specifications) in their data sheets.

Second-sourced parts (with the same type number and specifications but from different manufacturers) can have significantly different EMC performance – something it is important to control in production to ensure continuing compliance in serial manufacture. If products haven't been EMC tested with the alternative ICs fitted, it will be best to stick with a single source.

Suppliers of high-technology ICs may provide detailed EMC design instructions, as Intel does for its Pentium MMO chips. Get them, and follow them closely. Detailed EMC design advice shows that the manufacturer cares about the real needs of his customers, and may tip the balance when choosing devices.

Some FPGAs (and maybe other ICs) now have the ability to program the slew rate, output drive capability and/or output impedance of their drive signals. Their drive characteristics can be adjusted to give better signal integrity and/or EMC performance and this should help save time in development by reducing the need to replace ICs, change the values of components on the PCB, or modify the PCB layout.

Where ICs' EMC performances are unknown, correct selection at an early design stage can be made by EMC testing a variety of contenders in a simple standard functional circuit that at least runs their clocks, preferably performs operations on high-rate data too.

Testing for emissions can easily be done in a few minutes on a standard test bench with a close-field magnetic loop probe connected to a spectrum analyser (or a wideband oscilloscope). Some devices will be obviously much quieter than others. Testing for immunity can use the same probe connected to the output of a signal generator (continuous RF or transient) – but if it is a proprietary probe (and not just a shorted turn of wire) first check that its power handling is adequate.

Close-field probes need to be held almost touching the devices or PCBs being probed. To locate the "hottest spots" and maximise probe orientation they should first be scanned in a horizontal and vertical matrix over the whole area (holding the probe in different orientations at 90^{o} to each other for each direction), then concentrating on the areas with the strongest signals.

1.1.2 Batch and mask-shrink problems

Some batches of ICs with the same type numbers and manufacturers can have different EMC performance.

Semiconductor manufacturers are always trying to improve the yields they get from a silicon wafer, and one way of doing this is to mask-shrink the ICs so they are smaller. Mask-shrunk ICs can have significantly different EMC performance, because smaller devices means:

- less energy is required (in terms of voltage, current, power or charge) to control the internal transistors, which can mean lowered levels of immunity

- thinner oxide layers, which can mean less immunity to damage from ESD, surge, or overvoltage

- lower thermal capacity of internal transistors can mean higher susceptibility to electrical overstress

- faster operation of transistors, which can mean higher levels of emissions and higher frequencies of emissions.

Large users can usually arrange to get advance warnings of mask-shrinks so they can buy enough of the 'old' ICs to keep them in production while they find out how to deal with the changed EMC from the new mask-shrunk IC.

It is possible to perform simple goods-in checks of IC EMC performance to see whether a new batch has different EMC performance, for whatever reason. This helps discover problems early on, and so save money.

Alternatively, sample-based EMC testing in serial manufacture is required to avoid shipping non-compliant or unreliable products, but it is much more costly to detect components with changed EMC performance this way than it is at goods-in.

1.1.3 IC sockets are bad

IC sockets are very bad for EMC, and directly soldered surface-mount chips (or chip and wire, or similar direct chip termination techniques) are preferred. Smaller ICs with smaller bondwires and leadframes are better, with BGA and similar styles of chip packaging being the best possible to date.

Often the emissions and susceptibility of non-volatile memory mounted on sockets (or, worse still, sockets containing battery backup) ruin the EMC of an otherwise good design. Field-programmable low-profile SMD non-volatile memory ICs soldered direct to the PCB are preferred.

Motherboards with ZIF sockets and spring-mounted heatsinks for their processors (to allow easy upgrading) are going to require additional costs on filtering and shielding, even so it will help to choose surface-mounted ZIF sockets with the shortest lengths of internal metalwork for their contacts.

1.1.4 Circuit techniques

- Level detection (rather than edge-detection) preferred for control inputs and keypresses.

  Use level detection ICs for all control inputs and keypresses. Edge detecting ICs are very sensitive to high-frequency interference such as ESD. (If control signals need to use such very high rates that they need to use edge-detecting devices, they should be treated for EMC as for any other high-speed communication link.)

- Use digital edge-rates that are as slow and smooth as possible should be used wherever possible, especially for long PCB traces and wired interconnections (without compromising skew limits).

  Where skew is not a problem very slow edges should be used (could be 'squared-up' with Schmitt gates where locally necessary).

- On prototype PCBs allow for control of logic edge speed or bandwidths (e.g. with soft ferrite beads, series resistors, RC or Tee filters at driven ends).

  Many IC data books don't specify their output rise or fall times at all (or only specify the maximum times, leaving typical rates unspecified). Because it is often necessary to control unwanted harmonics, it is advisable to make provision for control of logic edge speed or bandwidths, (on prototype PCBs at least).

  Series resistors or ferrite beads are usually the best way to control edge rates and unwanted harmonics, although R-C-R tee filters can also be used and may be able to give better control of harmonics where transmission lines are used. (simple capacitors to ground can increase output transient currents and increase emissions.)

- Keep load capacitance low.

This reduces the output current transient when the logic state changes over and helps to reduce magnetic field emissions, ground bounce, and transient voltage drops in the ground plane and power supply, all important issues for EMC.

- Fit pull-ups for open-collector drivers near to their output devices, using the highest resistor values that will work.

  This helps reduce the current loop area and the maximum current, and so helps to reduce magnetic field emissions. However, this could worsen immunity performance in some situations, so a compromise may be needed.

- Keep high speed devices far away from connectors and wires.

  Coupling (e.g. crosstalk) can occur between the metallisation, bond wires, and lead frame inside an IC and other conductors nearby. These coupled voltages and currents can greatly increase CM emissions at high frequencies. So keep high speed devices away from all connectors, wires, cables, and other conductors. The only exception is high-speed connectors dedicated to that IC (e.g. motherboard connectors).

  When a product is finally assembled, flexible wires and cables inside may lie in a variety of positions. Ensure that no wires or cables can lie near any high-speed devices. (Products without internal wires or cables are usually easier to make EMC compliant anyway.)

  A heatsink is an example of a conductor, and clearly can't be located a long way away from the IC it is to be cooling. But heatsinks can suffer from coupled signals from inside an IC just like any other conductor. The usual technique is to isolate the heatsink from the IC with a thermal conductor (the thicker the better as long as thermal dissipation targets are met), then 'ground' the heatsink to the local ground plane with many very short connections (the mechanical fixings can often be used).

- A good quality watchdog that 'keeps on barking' is required.

  Interference often occurs in bursts lasting for tens or hundreds of milliseconds. A watchdog which is supposed to restart a processor will be no good if it allows the processor to be crashed or hung permanently by later parts of the same burst that first triggered the watchdog. So it is best if the watchdog is an astable (not a monostable) that will keep on timing out and resetting the microprocessor until it detects a successful reboot. (Don't forget that the watchdog's timeout period must be longer than the processor's rebooting time.)

  AC-coupling of the watchdog input from a programmable port on the micro helps ensure reliable watchdog operation. For more on watchdogs, see section 7.2.3 in [1].

- An accurate power monitor is needed (sometimes called a 'brownout' monitor).

  Power supply dips, dropouts interruptions, sags, and brownouts can make the logic's DC rail drop below the voltage required for the correct operation of logic ICs, leading to incorrect functioning and sometimes over-writing areas of memory with corrupt instructions or data. So an accurate power monitor is required to protect memory and prevent erroneous control activity. Simple resistor-capacitor 'power-on reset' circuits are almost certainly inadequate.

- Never use programmable watchdogs or brownout monitors.

  Because programmable devices can have their programs corrupted by interference, programmable devices must not be used for watchdog or power monitor functions.

- Appropriate circuit and software techniques also required for power monitors and watchdogs so that they cope with most eventualities, depending on the criticality of the product, (not discussed further in this series of articles).

- High quality RF bypassing (decoupling) of power supplies is vital at every power or reference voltage pin of an IC (refer to Part 5 of this series).

- High quality RF reference potential and return-current planes (usually abbreviated to 'ground planes') are needed for all digital circuits (refer to Part 5 of this series).

- Use transmission line techniques wherever the rise/fall time of the logic signal edge is shorter than the “round trip time” of the signal in the PCB track (transmission lines are described in detail in the 5th article in this series).

  Rule of thumb: round trip time equals 13ps for every millimetre of track length. For best EMC it may be necessary to use transmission line techniques for tracks which are even shorter than this rule of thumb suggests.

- Asynchronous processing is preferred.

  Asynchronous (naturally clocked) techniques have much lower emissions than synchronous logic, and much lower power consumption too. ARM have been developing asynchronous processors for many years, and other manufacturers are now beginning to produce asynchronous products.

  One of the limitations on designing asynchronous ICs was the lack of suitable design tools (e.g. timing analysers). But at least one asynchronous IC design tool is now commercially available.

Some digital ICs emit high level fields from their own bodies, and often benefit from being shielded by their own little metal box soldered to the PCB ground plane. Shielding at PCB level is very low-cost, but can’t always be applied to devices that run hot and need free air circulation.

Clock circuits are usually the worst offenders for emissions, and their PCB tracks will be the most critical nets on a PCB, requiring component layout to be adjusted to minimise clock track length and keep each clock track on one layer with no via holes.

When a clock must travel a long distance to a number of loads, fit a clock buffer near the loads so the long track (or wire) has smaller currents in it. Where relative skew is not a problem clock edges in the long track should be well-rounded, even sine-waves, squared up by the buffer near the loads.

1.1.5 Spread-spectrum clocking

So-called "spread-spectrum clocking" is a recent technique that reduces the measured emissions, although it doesn't actually reduce the instantaneous emitted power so could still cause the same levels of interference with some fast-responding devices. It modulates the clock frequency by 1 or 2% to spread the harmonics and give a lower peak measurement on CISPR16 or FCC emissions tests. The reduction in measured emissions relies upon the bandwidths and integration time constants of the test receivers, so is a bit of a trick, but has been accepted by the FCC and is in common use in the US and EU. The modulation rates in the audio band so as not to compromise clock squareness specifications.

Figure 2 shows an example of an emission improvement for one clock harmonic.

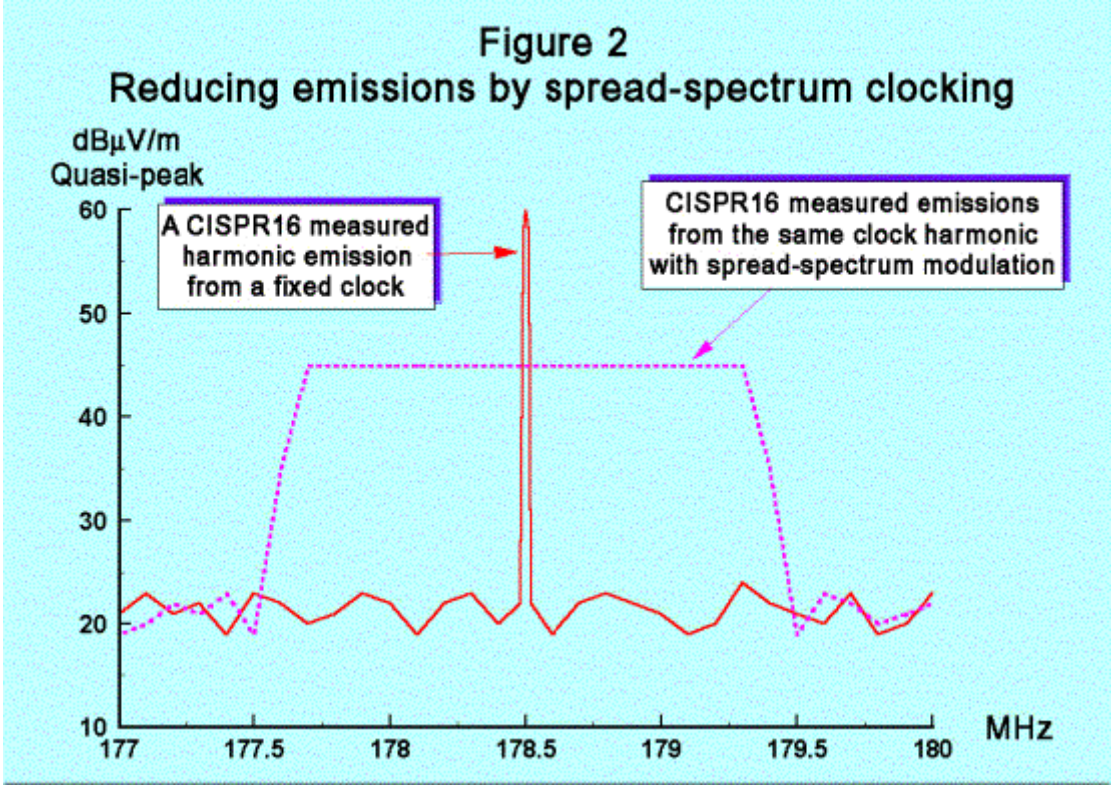

Debate continues about the possible effects of spread-spectrum clocking on complex digital ICs with the suppliers claiming no problems and some pundits still urging caution, but at least one major manufacturer of high-quality PC motherboards is using this technique as standard on new products.

Spread-spectrum clocking should not be used for timing-critical communications links, such as Ethernet, Fibre channel, FDDI, ATM, SONET, and ADSL.

Most of the problems with emissions from digital circuits are due to synchronous clocking. Asynchronous logic techniques (such as the AMULET microprocessors being developed by Prof. Steve Furber's group at UMIST) will dramatically reduce the total amount of emissions and also achieve a true spread-spectrum instead of concentrating emissions at narrow clock harmonics.

1.2 Analogue components and circuit design

1.2.1 Choosing analogue components

Choosing analogue components for EMC is not as straightforward as for digital because of the greater variety of output waveshapes. But as a general rule for low emissions in high-frequency analogue circuits: slew rates, voltage swings, and output drive current capability should be selected for the minimum necessary to achieve the function (given device and circuit tolerances, temperature, etc.).

But the biggest problem for most analogue ICs in low-frequency applications is their susceptibility to demodulating radio frequency signals which are outside their linear band of operation, and there are few if any data sheet specifications which can act as a guide for this. Specifications and standards for immunity testing of ICs are being developed, and in the future it may be possible to buy ICs which have EMC specifications on their data sheets.

Different batches, second-sourced, or mask-shrunk analogue ICs can have significantly different EMC performance for both emissions and immunity. It is important to control these issues by design, testing, or purchasing to ensure continuing compliance in serial manufacture, and some suitable techniques were described earlier (section on choosing digital ICs).

Manufacturers of sensitive or high-speed analogue parts (and data converters) often publish EMC or signal-to-noise application notes for circuit design and/or PCB layout. This usually shows they have

some care for the real needs of their customers, and may help tip the balance when making a purchasing decision.

1.2.2 Preventing demodulation problems

Most of the immunity problems with analogue devices are caused by RF demodulation.

Opamps are very sensitive to RF interference on all their pins, regardless of the feedback schemes employed (see Figure 3).

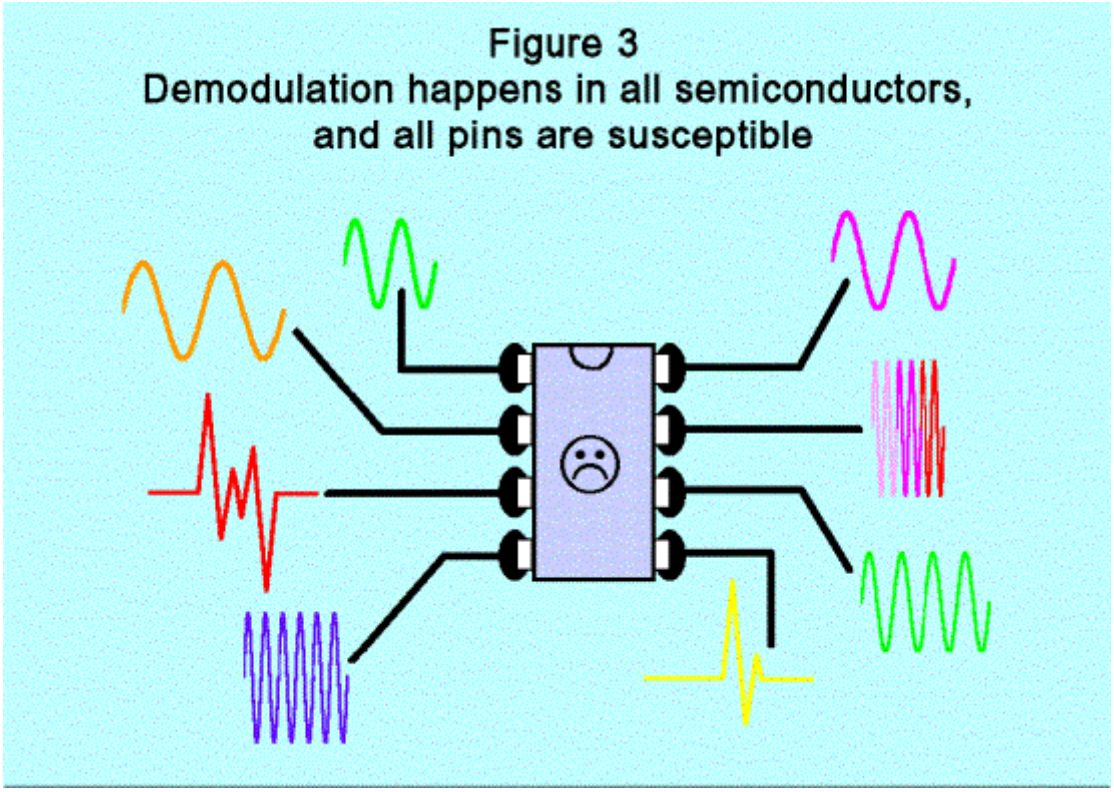

All semiconductors demodulate RF. Demodulation is more common problem for analogue circuits, but can produce more catastrophic effects in digital circuits (when software gets corrupted).

Even slow opamps will happily demodulate interference up to cellphone frequencies and beyond, as shown by the real product test results of Figure 4. To help prevent demodulation, analogue circuits need to remain linear and stable during interference. This is a particular problem for feedback circuits. Test the stability and linearity of the feedback circuit by removing all input and output loads and filters, then injecting very fast-edged (<1ns risetime) square waves into inputs (and possibly into outputs and power supplies, via small capacitors). The test signal amplitude is set so that the output pk-pk is about 30% maximum, to prevent clipping. The test signal's fundamental frequency should be near the centre of the intended passband of the circuit.

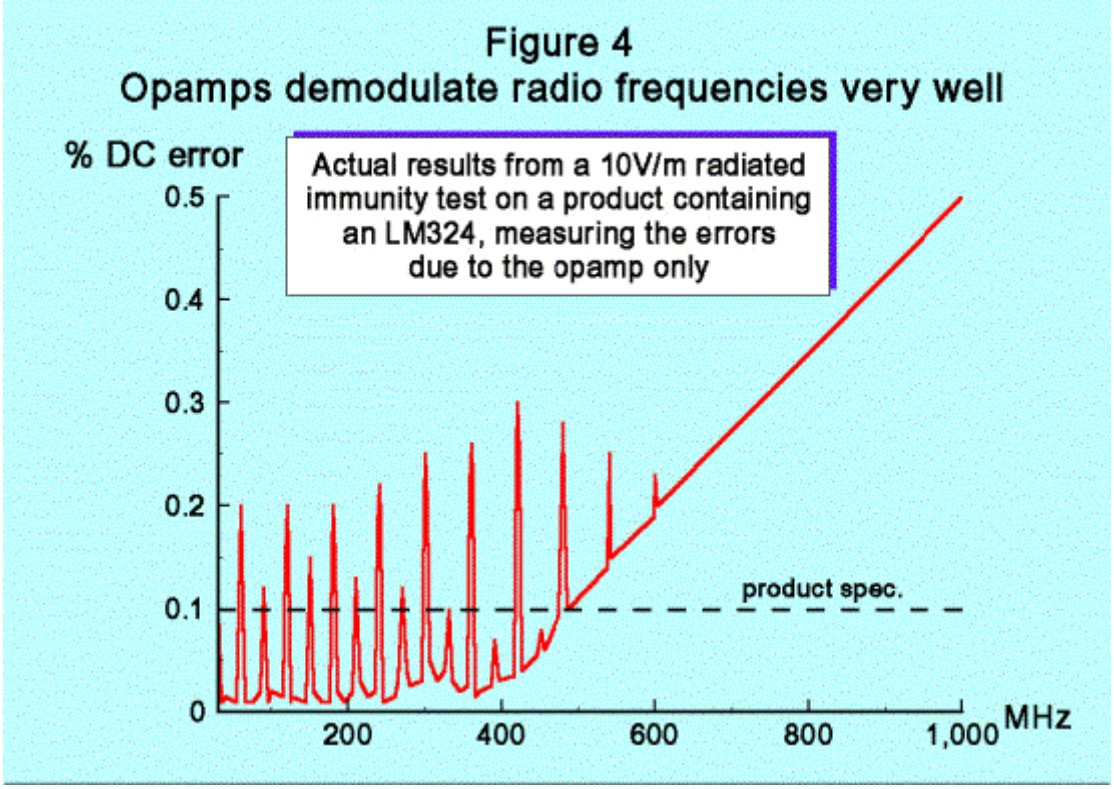

The circuit's output is observed with a 100MHz (at least) oscilloscope and probes for its slew rate, overshoot and ringing, even for audio or instrument circuits. For higher-speed analogue circuits use an appropriately faster 'scope and take great care to use appropriate high-speed probing techniques.

Feedback circuits should be adjusted so that slew rates are maximised, overshoots are low (heights of more than 50% of the signal's nominal height indicate instability). Any long periods of ringing (say, longer than two cycles) or bursts of oscillation also indicate instability.

Different batches of ICs can have very different stability performance, most easily simulated by cooling and heating the device under test over a wide range of temperatures (say: -30 to + 180oC) and ensuring the circuit is as fast and stable as it is possible to achieve over the whole temperature range.

Testing could use a swept frequency instead, with a spectrum analyser at the output. Take care not to overdrive the spectrum analyser's input.

1.2.3 Other analogue circuit techniques

Achieving good stability in feedback circuits usually requires that capacitive loads be buffered with a small resistance or choke which is outside the feedback loop.

Integrator feedback circuits usually need a small resistor (often around 560W) in series with every integrator capacitor larger than about 10pF.

Never try to filter or control RF bandwidth for EMC with active circuits – only use passive (preferably RC) filters outside any feedback loops. The integrator feedback method is only effective at frequencies where the opamp has considerably more open-loop gain than the closed-loop gain required by its circuit. It cannot control frequency response at higher frequencies.

Having achieved a stable and linear circuit, all of its connections might need protecting by passive filters or other suppression methods (e.g. opto-isolators). Any digital circuits in the same product will cause noise on all internal interconnections, and all external connections will suffer from the external electromagnetic environment.

Filtering is covered in Part 3 of this series, and the filters associated with an IC should connect to its local 0V plane. Filter design can be combined with galvanic isolation (e.g. a transformer) to provide

protection from DC to many GHz. Using balanced (differential) inputs and outputs can help reduce filter size while maintaining good rejection at lower frequencies.

Input or output filters are always needed where external cables are connected, but may not be necessary where opamps interconnect with other opamps by PCB traces over a dedicated 0V plane. Any wired interconnections inside unshielded enclosures might need filtering due to their antenna effect, as might wired interconnections inside shielded enclosures which also contain digital processing or switch-mode converters.

Analogue ICs need high-quality RF decoupling of all their power supplies and voltage reference pins, just as do digital ICs. RF decoupling techniques are described later in this volume.

But analogue ICs often need low-frequency power supply bypassing because the power supply noise rejection ratio (PSRR) of analogue parts are usually increasingly poor for frequencies above 1kHz. RC or LC filtering of each analogue power rail at each opamp, comparator, or data converter, may be needed. The corner frequency and slope of such power supply filters should compensate for the corner frequency and slope of device PSRR, to achieve the desired PSRR over the whole frequency range of interest.

Transmission line techniques may be essential for high-speed analogue signals (e.g. RF signals) depending on the length of their connection and the highest frequency to be communicated (see Part 5 of this series). Even for low-frequency signals, immunity will be improved by using transmission line techniques for interconnections, since correctly matched transmission lines of any length behave as very poor antennas and don't resonate.

Not many EMC design guides mention RF design. This is because RF designers are generally very good with most continuous EMC phenomena. However, local oscillators and IF frequencies often leak too much, so may need more attention to shielding and filtering.

Avoid the use of very high-impedance inputs or outputs. they are very sensitive to electric fields. Because the wave impedance of air is 377Ω, electric fields dominate outside of the near field of an emissions source.

Because most of the emissions from products are caused by common-mode voltages and currents, and because most environmental electromagnetic threats (simulated by immunity testing) are common-mode, using balanced send and receive techniques in analogue circuits has many advantages for EMC, as well as for reducing crosstalk. Balanced circuits drive antiphase (±) signals over two conductors, and does not use the 0V system for the return current path. Sometimes called differential signalling.

Comparators must have hysteresis (positive feedback) to prevent false output transitions due to noise and interference, also to prevent oscillation near to the trip point. Don't use faster output-slewing comparators than are really necessary (i.e. keep their dV/dt low).

Some analogue ICs themselves are particularly susceptible to radiated fields. They may benefit from being shielded by their own little metal box soldered to the PCB ground plane (take care to provide adequate heat dissipation too).

Figure 4B shows a simple opamp circuit (inverting amplifier) with some of the techniques described above applied. Even though the circuit uses single-ended signalling (i.e. uses 0V as the signal return) and is not balanced, common mode chokes will generally improve the EMC performance when used in the input and output filters.

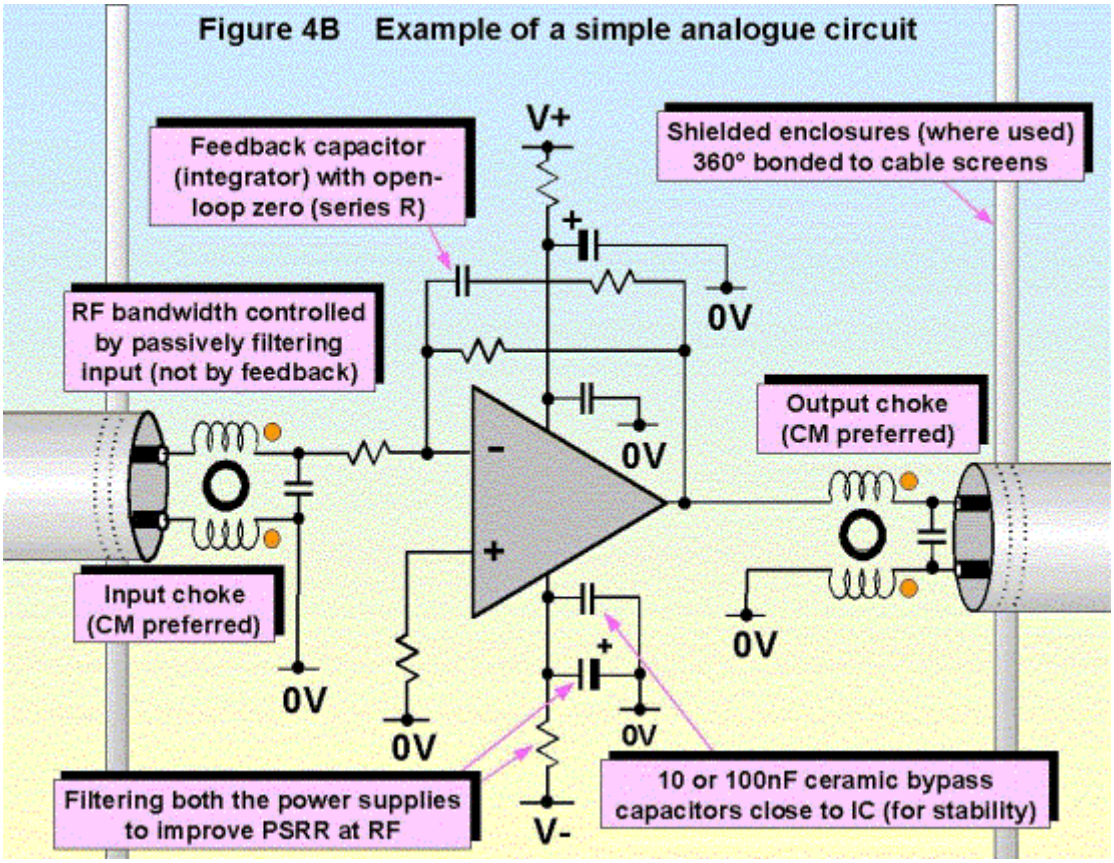

Input or output filters are always needed where external cables are connected, but may not be necessary where opamps interconnect with other opamps by PCB traces over a dedicated 0V plane. Any wired interconnections inside unshielded enclosures may also need filtering, as might wired interconnections inside shielded enclosures which also contain digital processing or switch-mode converters.

1.3 Switch-mode design

This technology is inherently electromagnetically noisy and will produce lots of interference if not firmly controlled, as outlined below. These techniques will also help make switch-mode power supplies low-noise enough to power sensitive analogue circuits.

1.3.1 Choice of topology and devices

Always switch power softly rather than abruptly, keeping both dV/dt and dI/dt low at all times. There are a number of circuit topologies which produce minimum emissions by reducing dV/dt and/or di/dt, whilst also reducing the stresses on the switching transistors. These include ZVS (zero-voltage switching), ZCS (zero current switching), resonant mode (a type of ZCS), SEPIC (single-ended primary inductance converter), Cük (an integrated magnetics topology, named after its inventor), etc.

In traditional (more noisy) topologies, where the power devices are not switched at zero volts or zero current, it is *not* true to say that reducing switching time *always* leads to efficiency improvements. All systems, circuits, and components (especially wound components) have natural resonant frequencies at radio frequencies. When the waveforms used by a circuit contain spectral components close to these natural resonant frequencies their resonances will become 'excited' and cause ringing, unwanted oscillations and emissions, and voltage overshoots that can increase the dissipation in power switching devices and even damage them.

Suppressing these resonances requires snubbing techniques which are usually lossy, as well as requiring costly components and PCB area. So switching at an ever-faster rate (which means increasingly high frequency content) eventually leads to diminishing efficiency and/or worsened reliability. For the most cost-effective design overall – soft-switching techniques trade a percentage point or two of device dissipation for much lower costs and sizes of filtering and shielding, minimum heatsink sizes and good reliability.

From an EMC point of view, faster switching edges means more energy in higher-frequency harmonics, hence larger and more complex filters and shielding. In poorly designed switch-mode power converters, harmonics of up to 1000 times the basic switching rate often cause failure to meet emissions tests.

One of the problems with switching power FETs is that their rate of change of drain voltage is a non-linear function of their gate voltage. Using the 'gate charge model' (which includes the 'Miller effect' from Cdg) provides much better accuracy when designing gate drive circuits so that they control the dV/dt at the drain.

1.3.2 Snubbing

Snubbing is usually required to protect the switching transistors from the peak voltages produced by the resonance of stray elements in the circuit components. Figure 5 shows the stray leakage inductance and inter-winding capacitance typical of an isolating transformer.

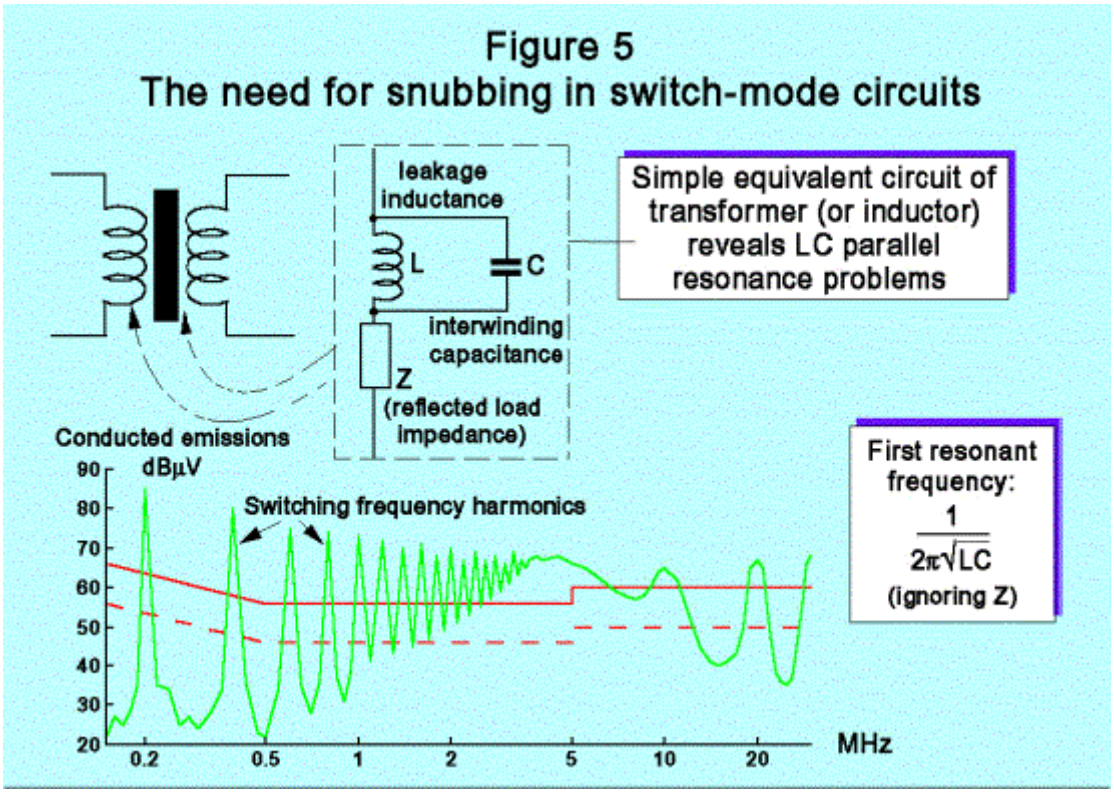

These form a resonant circuit which causes larger voltage overshoots the more abruptly its current is switched. On an emissions spectrum these resonances are often seen as a regular variation in the envelope of the emissions.

In the case of transformers, snubbers are connected across the winding whose overshoots are to be suppressed. Snubbers come in many types: A resistor and capacitor in series (RC type) is usually the best for EMC but can run hotter than other types.

Be prepared to compromise, and beware of using inductive components in snubbers. Inductance compromises snubber performance, so very low-inductance power resistors and pulse-rated capacitors should be used, with very short leads to the winding concerned.

1.3.3 Heatsinks

Heatsinks have around 50pF of capacitance to the collectors or drains of a TO247 power device, and similar capacitances to other package styles, so are strongly-coupled with the dV/dt of the collector or drain and can create strong emissions of electric fields through their own stray capacitances to other components either inside the product or the outside world. It is usually best to connect primary switching device heatsinks directly to one of the primary DC power rails – taking full account of all safety requirements, including a clear warning on or near the heatsink that it is live.

Heatsinks could be capacitively connected to the hazardous rail to improve safety, and it may even be possible to “tune” the capacitance with the length of its leads and traces to minimise troublesome frequencies.

It is important to return the RF current injected into the heatsink (via its 50pF or so capacitance) as quickly as possible back to its source whilst enclosing the smallest loop area, to avoid replacing an electric-field emissions problem with a magnetic field emissions problem. Always allow for some iteration on a prototype to find the best heatsink suppression method (for instance, which DC rail is the best to connect the heatsink to).

An alternative is to use shielded heat-sink thermal insulators. Their shielded inner layer is connected to the appropriate DC rail. The heatsink itself can remain isolated or else be connected to chassis. Although this is the safest, it is more costly.

Similar problems afflict the heatsinks of secondary rectifiers, but their heatsinks can usually be connected to their local 0V with no safety worries.

1.3.4 Rectifiers

The rectifiers used for primary flywheels and secondary rectifiers can cause a great deal of noise (hence emissions) due to their reverse current flow.

Faster-switching devices need less reverse charge (current x time) and can cause less noise. But if they are hard-switching types they can excite resonances in the switcher components (especially the isolation transformer) and cause excessive overshoots and emissions.

It is best for EMC to use rectifier types which have fast operation but soft-switching characteristics, as shown by Figure 6.

Figure 6
Improvements due to fast-recovery soft-switching rectifiers

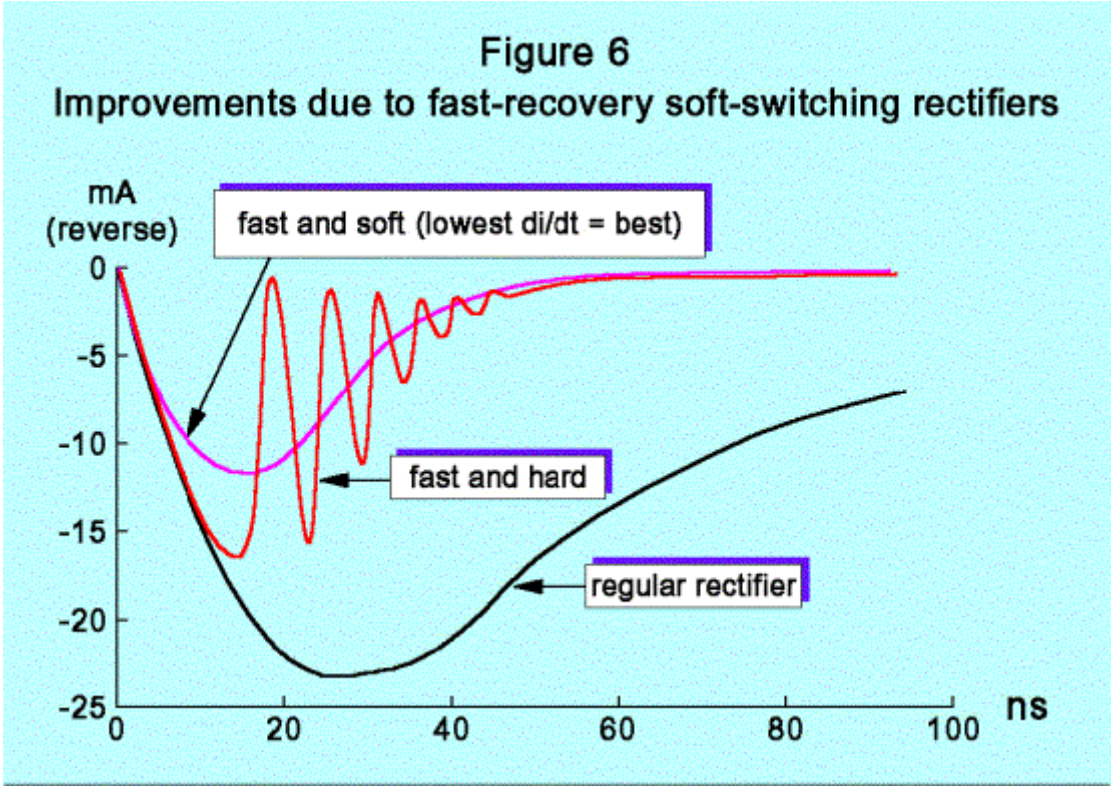

1.3.5 Problems and solutions relating to magnetic components

Pay particular attention to closing the magnetic circuits of inductors and transformers, e.g. using toroids or gapless cores. Iron powder toroidal cores are available for energy-storage magnetics, these effectively have a distributed air gap and so emit lower fields than gapped cores.

If air gaps have to be used, for instance in C, E or pot cores, an *overall* shorted turn may be needed to reduce the leakage fields. 'Overall' means that it goes around the *entire* body of the transformer, so it is only a shorted turn for the leakage fields.

Primary switching noise is injected via the interwinding capacitance of isolating transformers, creating common-mode noise in the secondaries. These noise currents are difficult to filter, and travel long distances, enclosing large loop areas (to keep Mr Kirchoff happy) thereby creating emissions problems.

Interwinding shields in an isolating transformer can suppress primary switching noise in the secondaries. One shield is a great help, and should be connected to a primary DC rail. Up the five shields is not unheard of, but three is more likely. When using three shields, the shield adjacent to the secondary windings usually connects to the common output ground (if there is one) and the shield in the middle usually connects to chassis. Be prepared to iterate a prototype to find their best connections.

PCB-transformers are becoming increasingly popular, and adding shields to these is simply a matter of adding more PCB layers (making sure that creepage and clearance distances are achieved despite tolerances in PCB manufacture).

Another powerful technique is to provide a local return path for these currents with small (safety approved!) capacitors connected between the secondary ground and one of the primary power rails.

Make sure that these capacitors don't cause the total earth leakage current to exceed the specification in the relevant safety standard.

These capacitors also help any filters on the secondaries to work much better, by reducing the source impedance of the emissions so that common-mode chokes can function effectively.

The above two techniques also reduce the secondary switching noise which appears at the input, via the isolating transformer's interwinding capacitance. The primary to secondary capacitor also makes filtering at the input more effective.

Figure 6B shows a simple switcher with a single interwinding shield and a primary-secondary bridging capacitor.

Figure 6B Two remedies for isolation transformer interwinding capacitance

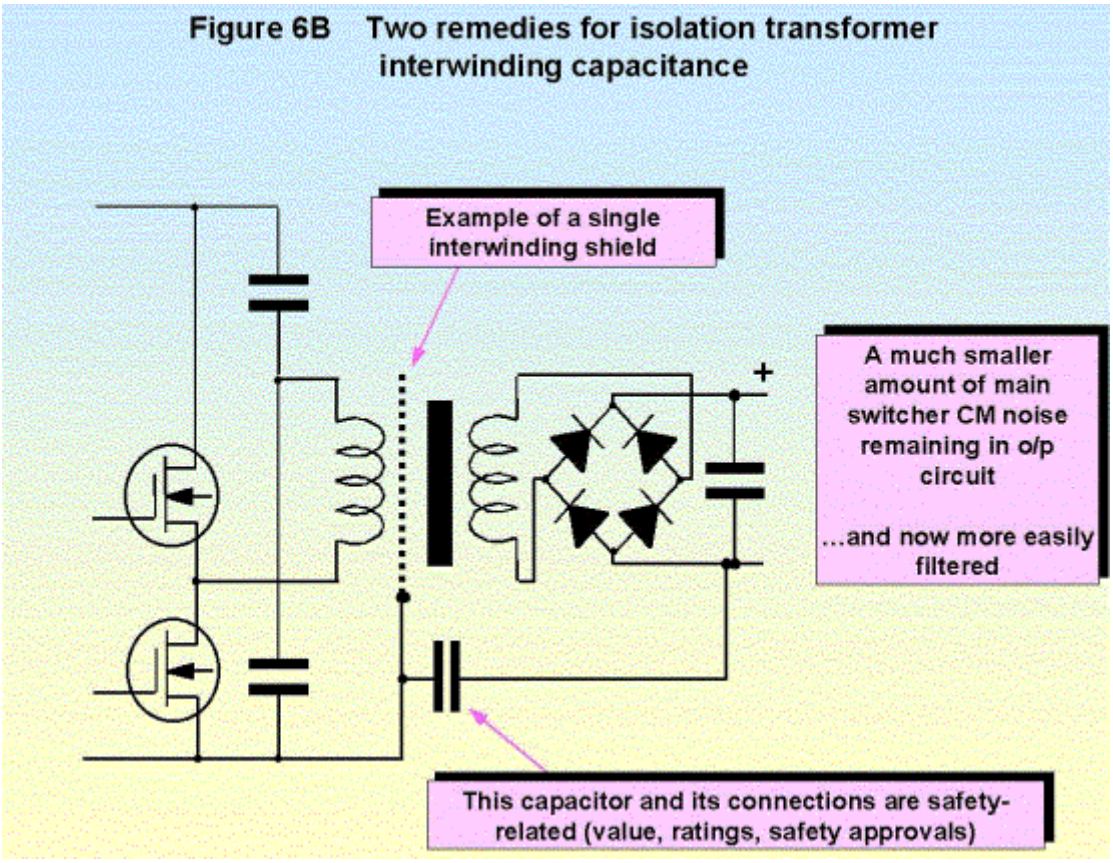

1.3.6 Spread-spectrum clocking for switch-mode

'Spread-spectrum clocking' techniques as described in 1.1.5 above can also be used with some switch-mode topologies to spread the emissions spectrum of the individual harmonics so that they measure less on an EMC test. Commercial and industrial conducted emissions tests use a 9kHz bandwidth from 150kHz to 30MHz, so spreading a harmonic by ±90kHz can give reductions of more than 10dB.

The spreading range can often be much larger than 1 or 2%, and some high-power converter manufacturers use almost white noise.

1.4 Signal communication components and circuit design

1.4.1 Non-metallic communications are best

The best communications for EMC purposes are infrared or optical, via free-space (e.g. IRDA) or fibre-optics. Their transmitters must not emit too much, and the receivers must be immune enough, but these are usually easier to control than the EMC of a long cable. Metal-can shielded transmitters and receivers are now readily available. It is often possible to bring metal-free fibre-optic cables right through the walls of shielded enclosures to PCBs or modules inside, without compromising the

enclosure shielding, whereas metallic wires and cables need to be filtered and/or 360° shield bonded at the points where they cross shielded enclosure boundaries.

Wireless communications are another alternative, but because they use the radio spectrum they sometimes cause interference with nearby electronics, and they can also be interfered with by electromagnetic disturbances.

Wires and cables may appear at first sight to be more cost-effective, but by the time their EMC problems have eventually been solved at the end of a project the non-metallic alternatives would often have been preferable for reasons of cost and timescale. Another reason for using non-metallic communications is that galvanic isolation to very high values is automatically achieved, improving product reliability and greatly easing the risks of failing EMC tests.

Wires and cables are usually cost-effective within a fully shielded product enclosure, but even then 'internal EMC' problems and the slow propagation velocity in cables can make infra-red or optical alternatives more attractive. (Don't forget to take account of the delays in the infra-red or optical transceivers themselves into account.)

1.4.2 Techniques for metallic communications

Single-ended signal communication techniques have very poor EMC performance for both emissions and immunity, and are best restricted to low frequency, low data rate, or short distance applications. They are usually all right as long as they remain on a PCB with a solid ground plane under all the tracks and don't go through any connectors or cables, which means that the single-PCB product is often the most cost-effective.

High-frequency or long-distance signals should be sent/ received as balanced signals (sometimes even on PCBs) for good signal integrity and EMC, and this is going to be a main issue in this sub-section.

Figure 8 shows examples of good and bad practices when connecting a millivolt output transducer to an amplifier via a cable.

Figure 8
Examples of good and bad practice in communications

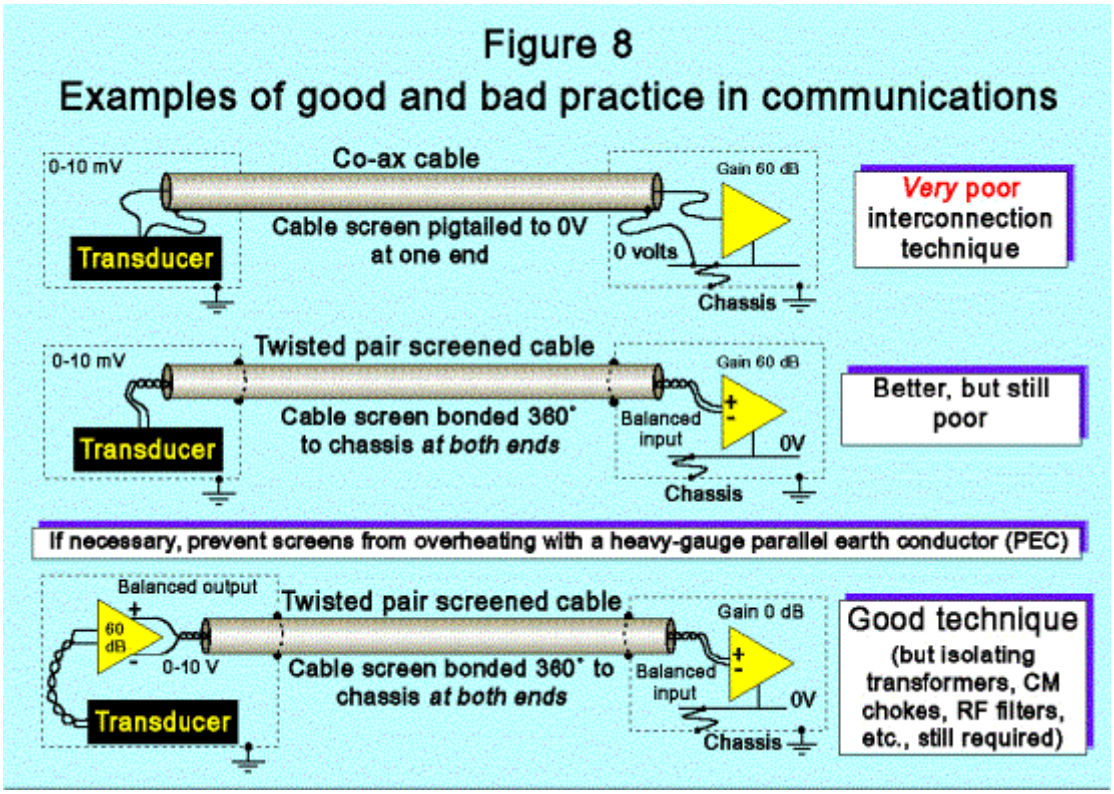

In general, connecting a cable shield to a circuit's 0V is very bad practice, as is the use of pigtails and grounding cable screens at one end only. Some older textbooks divide cables up into low and high frequency types, with different shield-bonding rules for each. But the electromagnetic environment is now so polluted with RF threats (and as was shown earlier, even 'slow' opamps will demodulate

>500MHz), and so many signals are polluted with RF common-mode noise from digital processors inside their products, that all cables should now be treated as high-frequency.

The three schemes in figure 8 show a hierarchy from a poor system for connecting to a transducer, through a better one, to a good system. Fitting an A/D converter in the transducer enclosure and sending high-level encoded data (with error-correction) over the cable to the product for decoding would be better than the best shown opposite. A perfect system would send the digital data over a fibre-optic instead of a metallic cable, and such systems are increasingly used in industry.

Concerns about cable shield heating in large or industrial premises are best dealt with by running the communications cable over a parallel earth conductor (PEC) to divert the majority of the heavy low-frequency currents (which will prefer to follow paths with lower *resistance*) and not by 'lifting' a shield connection at one end – which ruins the cable's shielding benefits at that end. Fitting a capacitor in series with the shield at one end is also not recommended as a design technique, although it may be useful as a remedial technique, because it is very difficult to make a capacitive bond work effectively over the full range of frequencies. PECs and other installation cabling and earthing techniques are discussed in detail in [2] [3] and [4].

For low frequency signals (say, under 100kHz) higher voltage levels in the communication link are better, for reasons of immunity. Where signal frequencies are above 10MHz (say) high voltages can lead to high levels of emissions – lower voltages are often preferred as the best compromise (e.g. as used by ECL, LVDS, USB). The signal frequency at which lower voltages are preferred depends on the length of cable and its type and EMC performance (especially its longitudinal conversion loss) and the design of the transmit and receive circuits.

Transmission line techniques may be essential for high-speed analogue or digital signals, depending on the length of their connection and the highest frequency to be communicated (see Part 5 of this series). Even for low-frequency signals, immunity will be improved by using transmission line techniques for their interconnections.

The best type of cable for EMC usually has a dedicated return conductor associated with each signal conductor, and any cable shields are used only to control interference. Co-axial cable is generally not preferred. Some cables need individually shielded signal pairs. It is very important to achieve a good *balance* over the whole frequency range, as this means a good common-mode rejection ratio (CMRR) and hence improved emissions and immunity. Balanced send / receive ICs are good, but isolation transformers have the benefit of adding galvanic isolation (up to the point where they flash-over) and also extending the common-mode range well beyond the DC supply rails.

Balanced construction twisted-pair or twinaxial cables usually give the best and most cost-effective emissions and immunity performance and very small differences in twist (and even the dielectric constants of the pigments used to colour their insulation) can be important. Balance is so important that in high-performance circuits even a physically balanced (mirror-image) PCB layout will be needed, using the same PCB layers.

Transformers and balanced send/receive ICs all suffer from degraded balance at RF. They generally require a common-mode choke in series to maintain good balance over the whole frequency range of interest. The CM choke always goes closest to the cable or connector at the boundary of the product.

Transformer isolation, balanced drive and receive, and CM chokes, all help to get the best EMC performance from a cable.

Figure 9 shows two examples, both equally applicable to providing good emissions and immunity for digital or analogue signalling (communications) of any speed or frequency range.

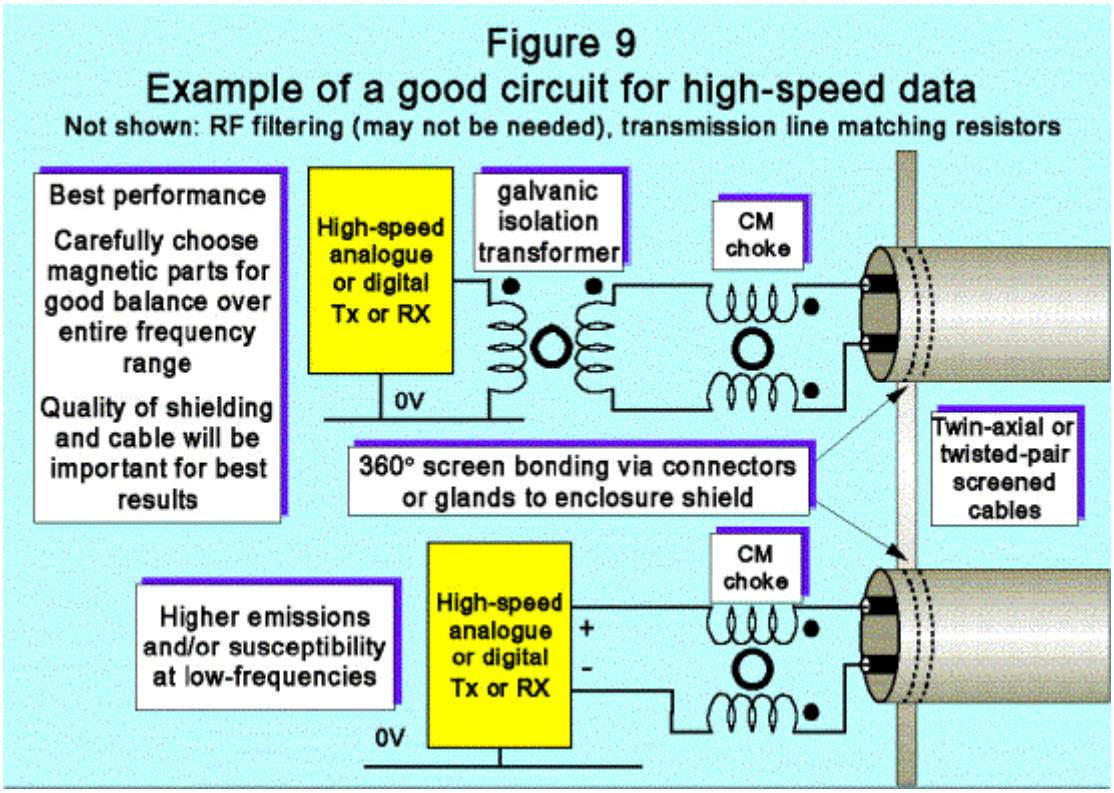

These circuits are ideal, in that a balanced send or receive circuit (in one case from a transformer, in the other an IC with balanced output or input) is connected to a balanced communications medium (the twin-axial or twisted-pair cable) via a CM choke.

Figure 10 shows how the CMRR of the choke is tailored to suit the transformer to give good balance over the whole frequency range, for a high-speed data example such as Ethernet. A similar design technique is used for the balanced IC.

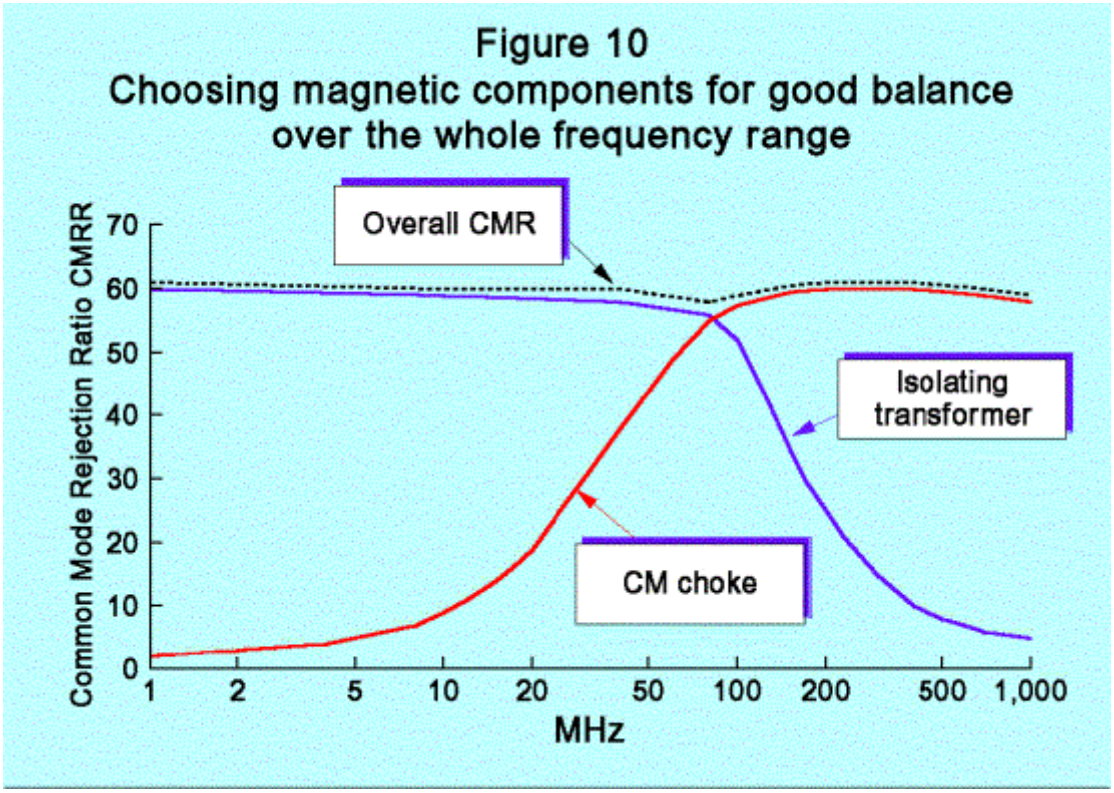

For a professional audio communication link the signal frequencies extend to 20Hz or less, so the isolating transformer will be large. Its large interwinding capacitance rolls its CMRR off to zero before 1MHz, so the CM choke then needs to be larger to provide CMRR down to 100 kHz or less. It is difficult to find a choke that has good CMRR from 100kHz to 1,000 MHz, so two chokes with different specifications may be needed in series to cover the range.

Where co-axial cables are used instead of twisted-pairs or twin-ax, EMC and signal integrity will suffer and the techniques shown in Figure 11 will help to achieve the best possible performance from the cables used.

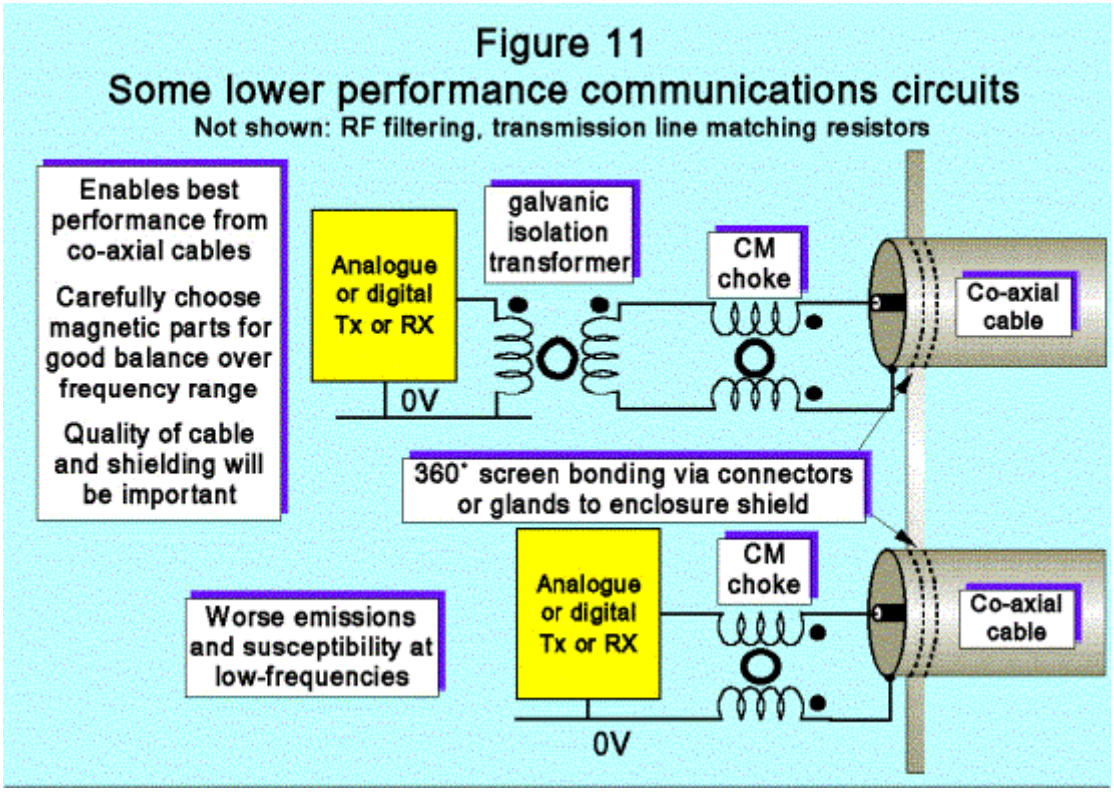

The circuit without the isolation transformer will generally suffer from poorer immunity at lower frequencies.

Many communications are still low frequency or low rate, and their signals are not particularly prone to causing emissions or suffering from interference. E.g. analogue to/from 8-bit converters will not be as sensitive as that from 12-bit converters, whereas 16 and higher number of bits will be very sensitive indeed.

Such signals are often sent down single wires in multiconductor cables to save cost, as shown by Figure 12 (an example of an RS232 application).

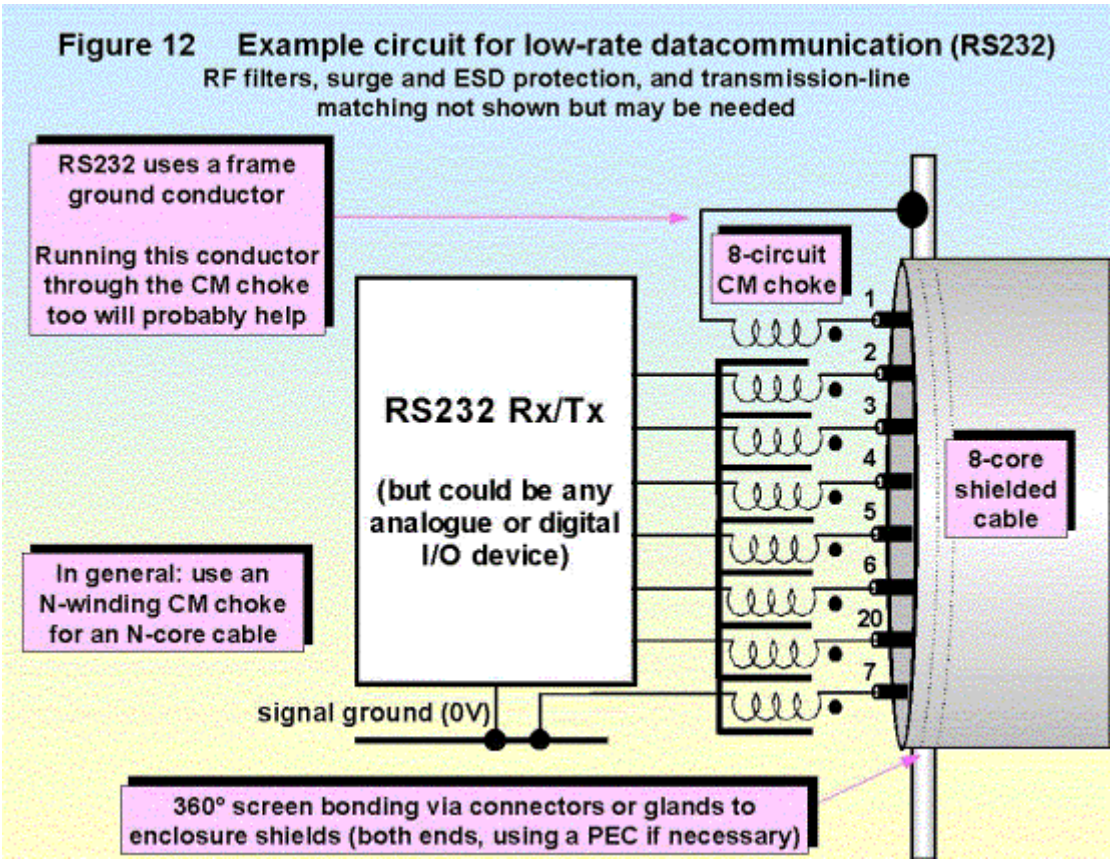

Where a conductor has N cores, it is best to connect it to the electronics at each end with a CM choke with N windings. Figure 12 shows a seven winding choke used for an eight-core cable, because one of the conductors is dedicated to “frame ground” according to the RS232 standard. (The frame ground lead is not likely to carry heavy currents and require a PEC because RS232 is only used for short-distances.)

RS232 only suits short distances because its single-ended signals lose their integrity rapidly as they radiate their energy as emissions. So although figure 12 (and the bottom circuit in figure 11) looks easy enough, the use of single-ended signals will require attention to CM choke and/or cable and/or connector quality. (Cable and connector types and qualities are discussed in the 2nd part of this series.)

Using drivers with very slow output edges (preferably slew-rate limited) can ease emissions problems significantly. Alternatively, standard drivers can be passively filtered to reduce their high-frequency content.

1.4.3 Opto-isolation

Opto-isolation is a common technique for digital signals, but the input-output capacitance of typical opto-coupler is around 1 pF – this creates a low enough impedance at frequencies above 10MHz to interact with the circuit impedances and destroy the balance of the signals in the cable.

As before, the selection of a suitable common-mode choke will restore the balance at high frequencies, allowing fast-edged signals to be communicated with fewer emissions or immunity problems.

Figure 13 shows an example of good EMC practices in a high-speed optically isolated link.

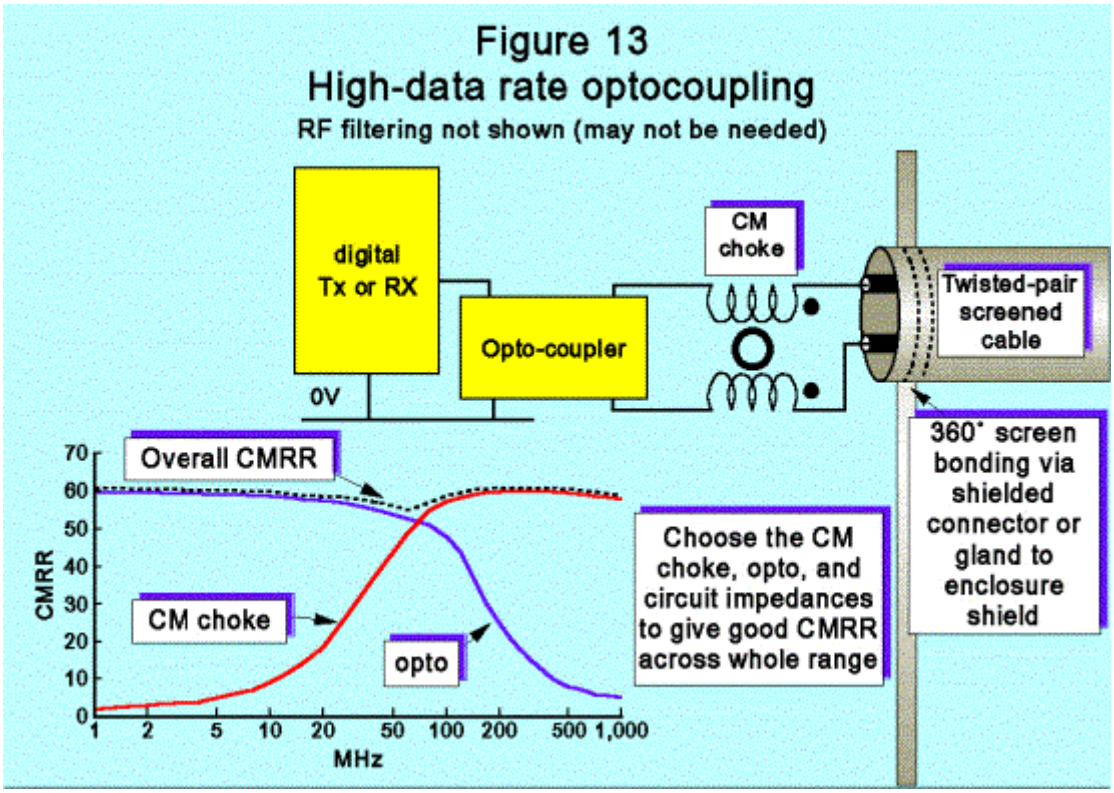

Similar to the previous examples, the CMRR of the CM choke is chosen to compensate for the fall-off in the balance of the opto-isolator at high frequencies, so that a good balance (equal to a good CMRR) is maintained across the full frequency range (DC to 1GHz in this example).

In many cases the CM choke can be replaced by two individual ferrite beads, and sometimes no choke or ferrites at all prove to be necessary.

But if they are not placed and routed on the PCB Murphy's Law predicts that they *will* be needed, and furthermore it is likely that there will be no room for them, no doubt making a wholesale redesign of the product necessary, including its plastic enclosure.

If the cable needs to be shielded, it must be 360° bonding via a shielded connector or gland to enclosure shield at both ends, using a PEC if necessary (see IEC 61000-5-2). But where galvanic isolation is needed bonding the shield at both ends may be forbidden. In this case a capacitive bond at one end may be used (the capacitor rated for the full voltage, and probably safety-approved too) - or the shield left unterminated at one end, which is liable to have poor EMC performance.

Analogue signals can now benefit also from opto-isolation with up to 0.1% linearity (e.g. using IL300 and the like). This can save having to use voltage-frequency converters (and vice-versa) in many opto-coupled applications.

Because of the common drawing practice of not showing power rails in full, it sometimes happens that both sides of an opto-isolator are powered form the same DC power rails, seriously compromising the isolation achieved and the RF performance. The RF performance of opto-isolators can only be as good as the RF isolation between their power supplies.

1.4.4 External I/O protection

External I/O is exposed to the full range of electromagnetic phenomena. The better circuits in the above figures should need less filtering or protection, for a given signal and semiconductors.

All the above communication circuits may need additional filtering for emissions or immunity with continuous EMC phenomena.

For ESD, transient, and surge phenomena the upper circuits of figures 9 and 11, and figure 13, are well-protected – providing their isolating transformers or opto-couplers will withstand the voltage stresses applied. RF filtering can also give some protection against ESD or fast transients.

The above circuits without isolating transformers or opto-couplers will almost certainly need overvoltage protection with diodes or transient suppressors, although heavy filtering might be adequate if data rates or frequencies are very low. For control signals a series 10k or 100k resistor closest to the connector followed by a 100 nF or 10 nF capacitor to the PCB ground plane makes a marvellous barrier against almost all EMC phenomena, but does not allow rapid changes in logic state.

Digital communications generally need a robust digital protocol (see below) to prevent data corruption, as protection devices only prevent actual damage to the semiconductors.

Allow for additional protection devices on a prototype board, and test it as early as possible to see which are needed.

1.4.5 "Earth – free" and "floating" communications

Another name for galvanic isolation is "earth free" or "floating", but these terms are often misunderstood or misused.

The above circuits using isolating transformers or opto-couplers are all "earth-free" and "floating", because no currents from the communications devices are assumed to flow between Tx and Rx via the 0V or chassis. This is true even though their cable screens are bonded at *both* ends to local chassis (enclosure shield). In fact, leakage currents flow through parasitic capacitances, and when CMRR is poor they can reach surprisingly large values.

The terms "earth-free" and "floating" are also sometimes applied to electronically balanced inputs or outputs, such as the lower circuit of figure 9. Although good CMRR performance will still give low leakage via 0V or chassis, such circuits are not galvanically isolated and are intrinsically more vulnerable to surges. Electronically balanced circuits also have a reputation for suffering from instability when one of the two lines is accidentally connected to ground.

Don't forget that the quality of the isolation achieved in practice is limited by the isolation performance of the power supplies supplying each side.

Never try to achieve "earth-free" operation by removing the protective earth from any equipment – this creates serious safety hazards and immediately contravenes several mandatory laws. If "ground loops" are a problem, use the proper circuit and installation techniques (e.g. PECs) and never compromise safety.

It is best to avoid jargon phrases like "earth-free" and "floating", instead state what is actually required or meant in plain circuit terms.

When screens *cannot* be connected at both ends

In some applications it is mandatory not to connect equipment grounds via cable screens or other conductors. The equipment concerned is still connected to main supply system's earth, but the earthing system is controlled in a special way. This does *not* help to achieve EMC at low cost. A screen connection at only one end will make the balance of the circuit and its conductors more important, and it will be more difficult and expensive to achieve a given EMC performance for a given signal.

Attention to creepage and clearances will also be important for safety reasons. In larger installations: when screens are not bonded at both ends, surges can cause arcing at the unconnected end possibly causing fire or toxic fumes. People can also receive shocks if they happen to be touching the screen and other equipment when a surge arrives. Clearly, not connecting the screens at both ends must place extra electrical and EMC stresses on some of the circuit components and cables, making surge, transient, and ESD damage more likely.

1.4.6 Hazardous area and intrinsically safe communications

Special barrier devices to limit the maximum power available in normal and fault conditions, and other restrictions, may be required. The EMC performance of these devices, which are made by specialist companies, is crucial. Further discussion is beyond the scope of this series.

1.4.7 Communication protocols

The data protocols used for digital communications are vital for both emissions and immunity, and it is much better to purchase chips that implement proven protocols than to try to develop them yourself. Simple protocols are easy, but they are very poor for EMC. Chips implementing CAN, MIL-STD-1553, LONWORKS, etc, have hundreds of man-years experience with interference control built into them, which no normal project team can ever hope to equal. Spend the extra few dollars on robust protocols, it will be worth it. Protocols are not discussed further in this series.

1.5 Choosing passive components

All passive components contain parasitic resistance, capacitance, and inductance. At the high frequencies at which many EMC problems occur these parasitic elements often dominate, making the components behave completely differently. E.g.: at high frequencies a film resistor becomes either a capacitor (due to its shunt C of around 0.2pF) or an inductor (due to its lead inductance and spiral tolerancing). These two can even resonate to give even more complex behaviour. Wire-wound resistors are useless above a few kHz, whereas film resistors under 1k usually remain resistive up to a few hundred MHz. A capacitor will resonate due to the effect of its internal and lead inductances, and above its first resonance it will have a predominantly inductive impedance.

Surface mounted components are preferred for good EMC because their parasitic elements are much lower and they provide their nominal value up to a much higher frequency. E.g. SMD resistors under 1k are usually still resistive at 1,000 MHz.

All components are also limited by their power handling capacity (especially for surges handling), dV/dt capacity (solid tantalum capacitors go short-circuit if their dV/dt is exceeded), dI/dt, etc. Passive components can also suffer severe temperature coefficients, or need de-rating. SMD parts have lower power ratings than leaded, but since most power occurs at lower frequencies it is often possible to use leaded parts in those areas, although taking care to minimise lead length.

For capacitors, ceramic dielectrics usually give the best high frequency performance, so SMD ceramics are often excellent. Some ceramic dielectrics have strong temperature or voltage coefficients, but COG or NPO dielectric materials have no tempco or voltco to speak of and make very stable and rugged high-quality high-frequency capacitors. They tend to be larger and cost more than other types, for values above 1nF.

Magnetic parts should have closed magnetic circuits, as has been described above. This is important for immunity as well as for emissions. Rod-cored chokes or inductors must be used with great care, if they cannot be avoided altogether (what shape is the ferrite antenna of a radio receiver?). Even the mains transformers used in linear power supplies can give better EMC performance if they have an interwinding screen connected to protective earth.

All these imperfections in passive components makes filter design very much more complicated than the circuits in textbooks and on simulator screens might suggest. Where a passive component is to be used with high frequencies (e.g. to decouple interfering currents up to 1,000MHz to a ground plane) it helps to know all about its parasitic elements and to do a few simple sums to work out their effects. Helpful manufacturers of quality components publish parasitic data, even sometimes impedance performance over a broad range of frequencies (often revealing their self-resonances).

Some passive components will need to be rated for safety, especially *all* those connected to hazardous voltages, of which the AC supply is often the worst case. It is best to only use parts here which have been approved to the correct safety standard(s) at the correct ratings by an accredited third-party laboratory and allowed to carry their distinguishing mark (SEMKO, DEMKO, VDE, UL, CSA, etc.). But the presence of the mark on the component means nothing. Much better is to get a copy of all the test labs' certificates for the safety approved parts and check they cover all they should.

The use of components with unknown parasitics for high-speed signals and/or EMC purposes makes it more likely that the number of product design iterations will be high and time-to-market delayed.

1.6 References:

[1] Tim Williams, *EMC for Product Designers* 3rd edition, Newnes 2001, ISBN: 0-7506-4930-5, www.newnespress.com

[2] IEC 61000-5-2:1997 Electromagnetic Compatibility (EMC) – Part 5: Installation and mitigation guidelines – Section 2: Earthing and cabling, www.iec.ch.

[3] Tim Williams and Keith Armstrong, *EMC for Systems and Installations*, Newnes 2000, ISBN 0 7506 4167 3 www.newnespress.com, RS Components Part No. 377-6463.

[4] Keith Armstrong, *EMC for Systems and Installations Part 2 – EMC techniques for installations*, EMC+Compliance Journal, April 2000, pp8 – 17. All UK EMC Journal and EMC+Compliance Journal articles are available electronically from the magazine archive at www.compliance-club.com .

Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, www.cherryclough.com
Phone: 01457 871 605, Fax: 01457 820 145, e-mail: keith.armstrong@cherryclough.com

Design Techniques for EMC – Part 2
Cables and Connectors

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the second in a series of six articles on best-practice EMC techniques in electrical/electronic/mechanical hardware design. The series is intended for the designer of electronic products, from building block units such as power supplies, single-board computers, and "industrial components" such as motor drives, through to stand-alone or networked products such computers, audio/video/TV, instruments, etc.

The techniques covered in the six articles are:

1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
2) **Cables and connectors**
3) Filters and transient suppressors
4) Shielding
5) PCB layout (including transmission lines)
6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), so this magazine article format can do no more than introduce the various issues and point to the most important of the best-practice techniques. Many of the techniques described in this series are also important for improving signal integrity.

Table of contents for this part

2. All cables are antennas

2.1 Spectrum use and the possibilities for interference

Figure 2A shows the frequencies in common use in civilian daily life, from AC powerlines through audio frequencies, long, medium, and short-wave radio, FM and TV broadcast, to 900MHz and 1.8GHz cellphones.

Figure 2A The frequencies we use

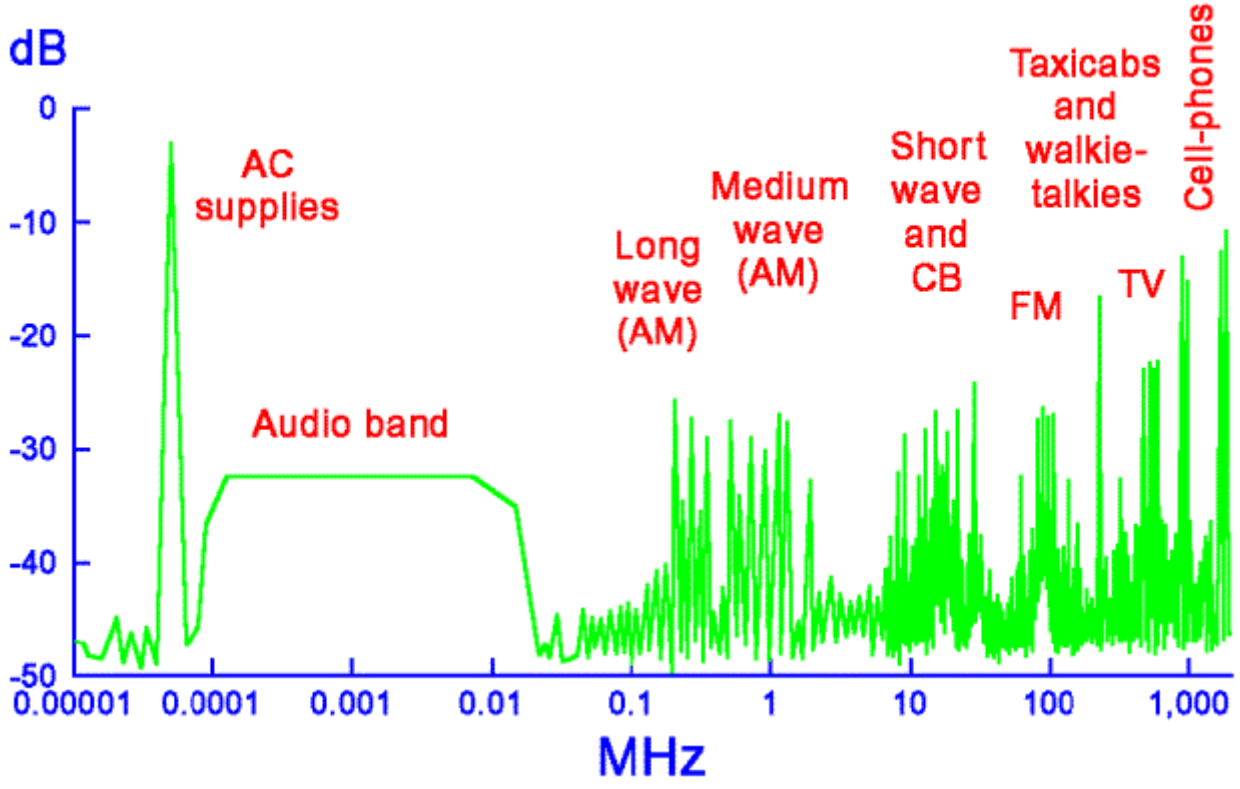

The real spectrum is busier than this – all of the range above 9kHz is used for something by someone.

This figure will soon need extending to 10 (or even 100GHz) as microwave techniques become more commonplace in ordinary life.

Figure 2B overlays the usage spectrum of Figure 2A with a less familiar spectrum showing the typical emissions from commonplace electrical and electronic equipment.

Figure 2B Plus the interference we create

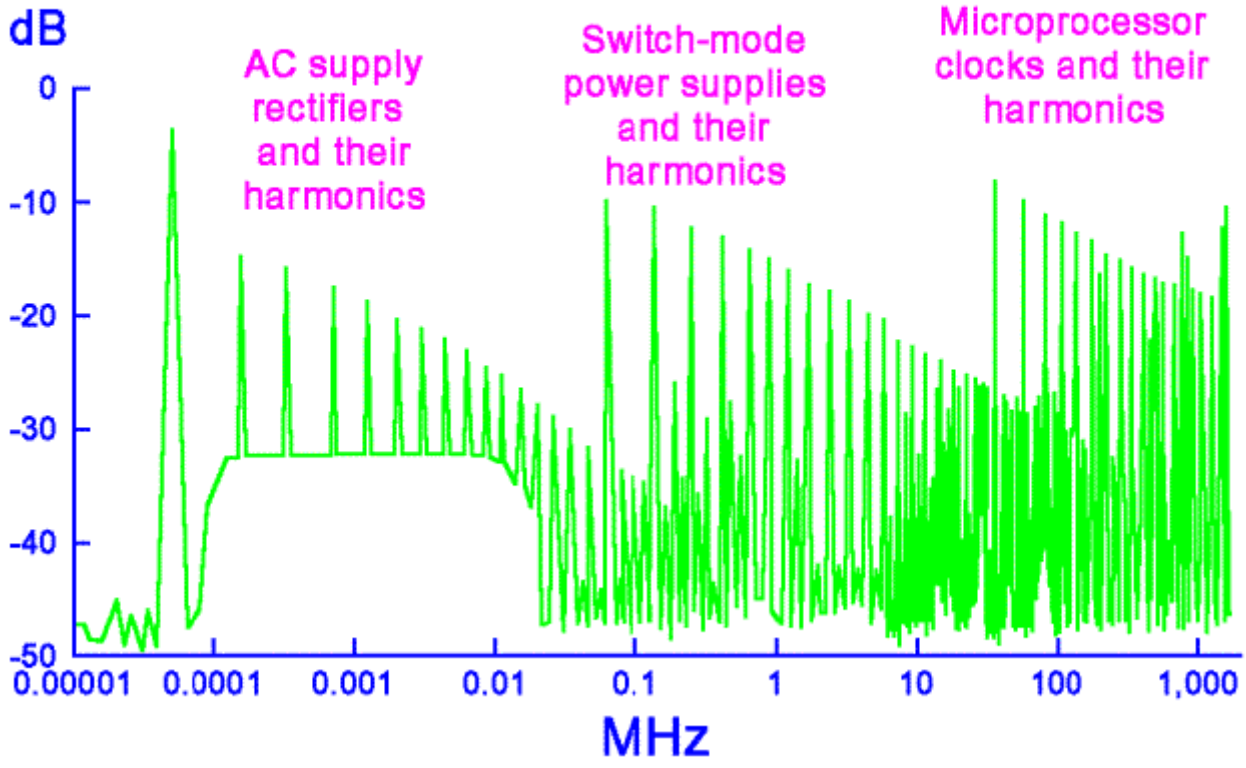

AC mains rectifiers emit switching noise at harmonics of the fundamental to considerable frequencies, depending on their power.

A 5kVa or so power supply (whether linear or switch-mode) can fail conducted emissions limits up to several MHz due to the switching noise of its 50 or 60Hz bridge rectifier.

Thyristor-based DC motor drives and phase-angle AC power control will have similar emissions. These emissions can easily interfere with long and medium wave broadcasting, and part of the short-wave band.

Switch-mode power convertors can operate at fundamental frequencies between 2 and 500kHz. It is not unusual for a switch-mode convertor to have significant levels of emissions at 1,000 times its switching frequency. Figure 2B shows the emissions from a 70kHz switching power supply typical of a personal computer. These emissions can interfere with radio communications up to and including the FM broadcast band.

Figure 2B next shows the typical emissions spectrum from a 16MHz clocked microprocessor or microcontroller. It is not unusual for these commonplace items to exceed emissions limits at frequencies of 200MHz or more. As personal computers are now using 400MHz clocks and heading for 1GHz, it is obvious that digital technology is capable of interfering with (and being interfered with) all the upper range of our spectrum.

The reason for mentioning this is that *all conductors are antennas*. They all convert conducted electricity into electromagnetic fields, which can then leak out into the wider environment. They all convert electromagnetic fields in their locality into conducted electrical signals. There are no exceptions to this rule in our universe.

Conductors are thus the principal means by which signals cause radiated emissions, and by which external fields contaminate signals (susceptibility and immunity).

2.2 Leakage and antenna effect of conductors

Electric (E) fields are created by voltages on conductor areas, and magnetic (M) fields are created by currents flowing (in loops, as they always do). All electrical signals create both types of field with their conductors, so all conductors leak their signals to their external environment, and allow external fields to leak into their signals.

At distances greater than one-sixth of the wavelength (λ) of the frequencies of concern, E and M fields develop into full electromagnetic (EM) fields with both electric and magnetic components.

For example: the transition to full EM fields occurs at 1.5 metres for 30MHz, 150mm for 300MHz, and 50mm for 900MHz.

So as frequencies increase, treating conductors as merely electric or magnetic field emitters and receivers becomes inadequate, as shown by figure 2C.

Figure 2C Electric and magnetic fields become electromagnetic fields with distance

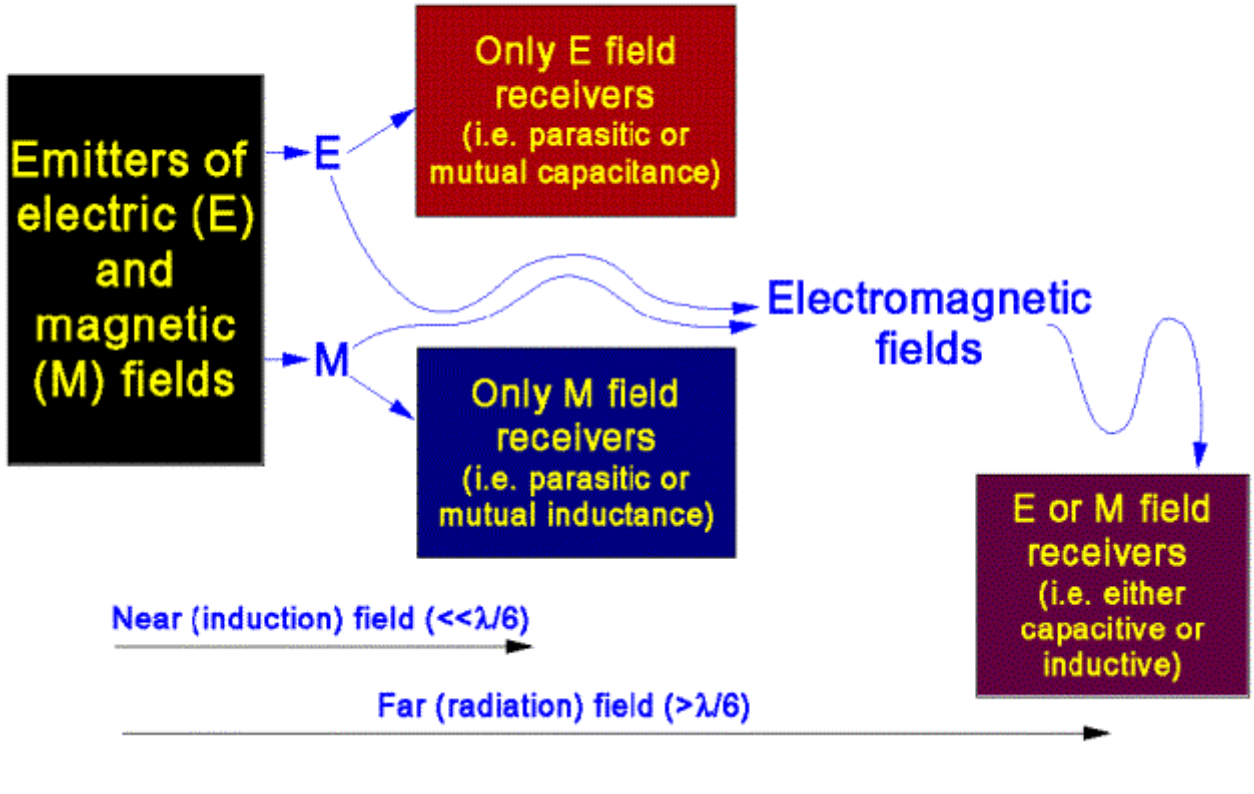

Another effect of increasing frequencies is that when λ is comparable with conductor length, resonances occur. At some of these the conversion of signals to fields (and vice-versa) can reach almost 100%. E.g. a standard whip antenna is merely a length of wire, and is a perfect convertor of signals to fields when its length equals one-quarter of λ.

This is a very simplistic description, but as far as the user of cables and connectors is concerned the important thing is that all conductors can behave as resonant antennae. Obviously we want them to be very poor antennas, and assuming that a conductor is like a whip antenna (good enough for our purposes) we can use Figure 2D to help guide us.

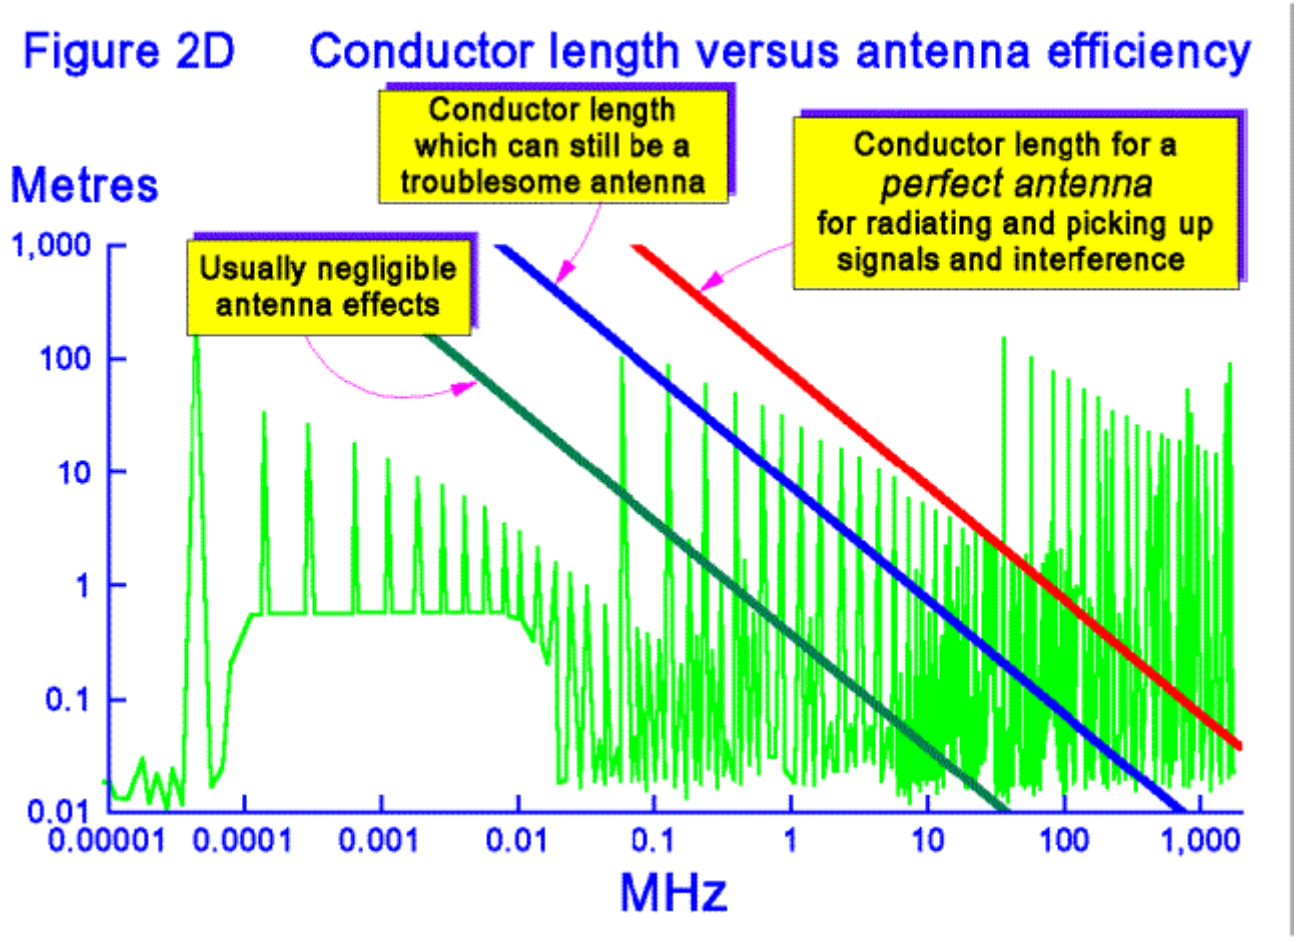

The vertical axis of Figure 2D is in metres of conductor length, and the spectrum of Figure 2B is retained as a visual guide. The right-hand-most (red) diagonal shows conductor length versus frequency for a perfect antenna.

Obviously, at frequencies in common use, even very short conductors can cause emissions and immunity problems. A signal or field at 100MHz finds a 1 metre long conductor to be a very efficient antenna, and at 1GHz 100mm conductors make good antennae. *This simple fact is responsible for a large number of "black magic" EMC problems.*

Not so many years ago, the frequencies in commonplace use were much lower and typical cable lengths were not very effective antennae, which is why electrical wiring "custom and practice" tends to be out of date.

The middle (blue) red diagonal in Figure 2D shows conductor lengths which do not make very efficient antennae, but can still cause problems. The left-hand (green) diagonal shows lengths which are so short that (for all except the most critical products) their antenna effects can usually be neglected.

How many times have you heard someone say: "It's OK, I've earthed it."? It is a standing joke in the EMC community that RF is colour blind, and so can't tell that the green/yellow striped conductor they are travelling in is supposed to be a perfect earth, and consequently all earth conductors are antennas too.

2.3 All cables suffer from intrinsic resistance, capacitance, and inductance

Forgetting fields and antennas for a moment: a few quick-and-dirty examples will show how even very tiny departures from the ideal cause problems for signals carried by conductors at commonplace modern frequencies.

- The resistance of a 1mm diameter wire at 160MHz is 50 times more than at DC, due to the skin effect forcing 67% of the current to flow in its outermost 5 microns at that frequency.

- A 25mm long 1mm diameter wire has an intrinsic space-charge capacitance of around 1pF, which does not sound much but loads it by around 1kΩ at 176MHz. If this 25mm long piece of wire alone was driven in free space by a perfect 5V peak-to-peak 16MHz square wave, the eleventh harmonic of the 16 MHz would take 0.45mA just to drive the wire.
- A connector pin 10mm long and 1mm diameter has an intrinsic inductance around 10nH, which does not sound like much. When driven with a perfect 16MHz square wave into a backplane bus impedance that draws 40mA, the voltage drop across this pin will be around 40mV, enough to cause significant problems for signal integrity and/or EMC.
- A 1 metre long wire has an intrinsic inductance of around 1μH, preventing surge protection devices from working properly when used to connect them a building's earth-bonding network.
- A 100mm long earth wire for a filter has so much intrinsic inductance (around 100nH) that it can ruin filter performance at > 5MHz or so.
- The inductance of a 25mm long "pigtail" termination for the screen of a 4 metre cable is enough to ruin the cable's screening effectiveness at >30MHz or so.

The rules of thumb for intrinsic capacitance and inductance for wires under 2mm diameter is 1pF per inch and 1nH per mm (sorry to mix units, but they stick in the mind better).Very simple maths such as $Z_C = \frac{1}{\sqrt{2\pi fC}}$ and ZL = $2\pi fL$ found in most basic electronic textbooks allows any engineer to discover whether the intrinsic imperfections in conductors are likely to be significant.

2.4 Avoiding the use of conductors

The above rather hand-wavy analysis shows that cables are increasingly problematic as frequencies increase: it is difficult to get them to carry signals properly, and difficult to stop them from leaking.

Even for low-frequency signals such as audio, cables present increasing problems. Since all semiconductors act like "crystal set" detectors up to many hundreds of MHz (typical even of slow op-amps like LM324), the antenna effects of cables pollutes the audio signals unpredictably.

So the best advice on signal and data cables and connectors for the most cost-effective EMC compliance may be not use metallic conductors at all. Non-metallic communications are preferred, and there are a lot of alternatives these days, including:

- Fibre-optics (preferably metal-free)
- Wireless (e.g. Bluetooth; wireless LANs)
- Infra-red (e.g. IrDA)
- Free-space microwave and laser links (e.g. between buildings)

2.4.1 Cost/benefit analyses of alternatives to conductors

Many designers feel they have to keep material costs down by using traditional cables and wires. But when the overall costs of completing a project and producing reliable, EMC compliant products, systems, and installations is taken into account it is often found that a fibre-optic or wireless link would have cost less overall. By then it is too late, of course.

For signal cables and connectors: *material cost is no longer related in any predictable way to the profitable selling price*, except for the simplest of electronic products. A proper cost/benefit analysis should take account of likely problems with signal integrity and EMC compliance, risks of incurring penalty charges, plus the risks of high levels of product returns, warranty claims, and lower levels of sales due to market perception.

Design engineers prefer not to consider the commercial risks of their designs, but they are the only people able to do this with any accuracy (usually with inputs from more commercial personnel). But as long as electronic designers insist on only considering the functional performance and material costs of their designs, their companies are missing competitive advantages and suffering commercial risks of unknown size.

2.5 Cable segregation and routing

Installation cabling rules are really outside the scope of this series, but the product designer needs to know what they are so as to design his product's external connections. Here is a quick summary of the main recommendations of IEC 61000-5-2:1997 and many other recent standards concerned with the installation of information technology and telecommunications:

a) All buildings to have a lightning protection system to BS6651 Appendix C or equivalent, bonded at ground level at least to their internal bonding network. All building steel, metalwork, cable ducts, conduits, equipment chassis, and earthing conductors in a building to be cross-bonded to create a 3-dimensional bonding network with mesh size no greater than 4 metres.

b) Segregate power and signal cables into at least four "classes", from very sensitive to very noisy.

c) Run all cables along a single route between items of equipment (which should therefore have a single connection panel each), whilst preserving at least minimum specified spacings between cable classes.

d) 360^{o} Bond cable screens (and any armouring) to the equipment enclosure shields *at both ends* (see later) unless specifically prohibited by the manufacturer of the (*proven* EMC compliant) transducer or equipment.

e) Prevent excessive screen currents by routing all cables (signal and power) very close to conductors or metalwork forming part of the meshed earth network.

f) Where meshed building earth is not available, use cable trays, ducts, conduits, or if these don't exist a heavy gauge earth conductor, as a Parallel Earth Conductor (PEC). A PEC must be bonded at both ends to the equipment chassis earths and the signal cable strapped to it along its entire length.

The needs for segregation, PECs, and (in general) screen bonding at both ends will have an impact on the design of interconnections panel layout, choice of connector types, and the provision of some means for bonding heavy-duty PECs.

Figure 2E gives an overview of the techniques involved in connecting screened enclosures together with both screened and unscreened cables.

Figure 2E Good practices when installing shielded enclosures

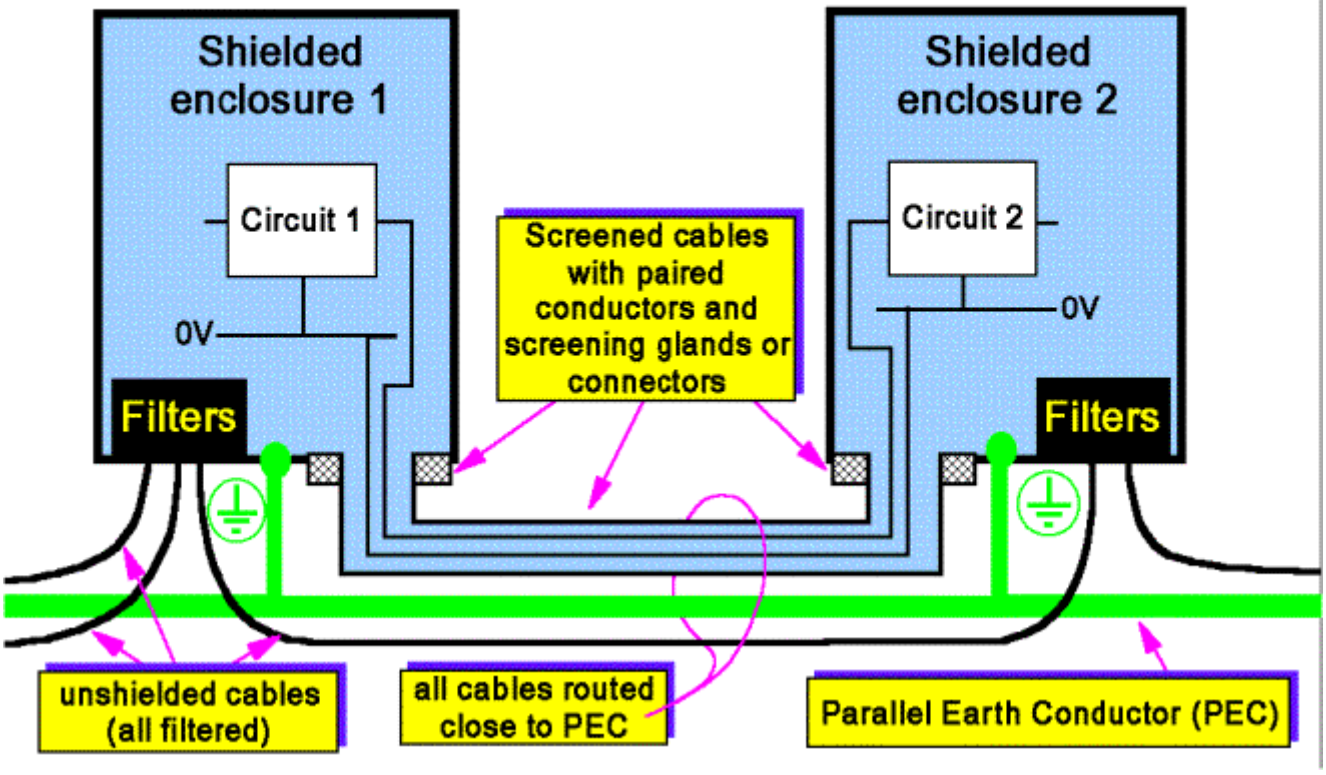

For short connections between items of equipment such as a PC and its VDU, dedicated printer, and modem, only d) above (360° cable screen bonding to enclosure shields at both ends) is needed – providing all the interconnected items are powered from the same short section of ring main, and all long cables to other parts of the building (e.g. network cables) are galvanically isolated (e.g. Ethernet). These screen bonding techniques are also needed for the EMC domestic hi-fi and home theatre systems. However, a) often comes in handy as well for protecting such equipment from damage during a thunderstorm.

2.6 Getting the best from cables

Open any signal cable manufacturer's catalogue and you will find a huge variety of cable types, even for similar tasks. This is a warning that cables are all imperfect. The best cable for a given application will be difficult to select, and then will probably be too expensive, too bulky, too stiff, and only available to special order on 26 week leadtime in 5km reels.

2.6.1 Transmission lines

Transmission line techniques prevent cables from acting as resonant antennas.

When the send and return conductors of a signal current loop are physically close together and so enjoy strong mutual coupling, the combination of their mutual capacitance and inductance results in a *characteristic impedance* $Z_0 = \sqrt{\frac{L}{C}}$, where L and C are the capacitance and inductance per unit length (a fraction of the λ of the highest frequency of concern). Z_0 can be calculated for cables and connectors (also for PCB tracks, see Part 5 of this series).

When Z_0 is kept constant over the entire length of an interconnection, and when drive and/or send (source or load) impedances are "matched" to Z_0, a controlled-impedance transmission line is created and *resonant effects do not happen*. The intrinsic inductance and capacitance of the conductors also

create far fewer problems. This is why RF and all EMC test equipment use 50Ω transmission line cables and connectors (see Figure 2F), and why high-speed and/or long distance data busses and serial communications also use transmission lines (usually in the range 50 to 120Ω).

Figure 2F The transmission line principle familiar from standard "50Ω" test gear

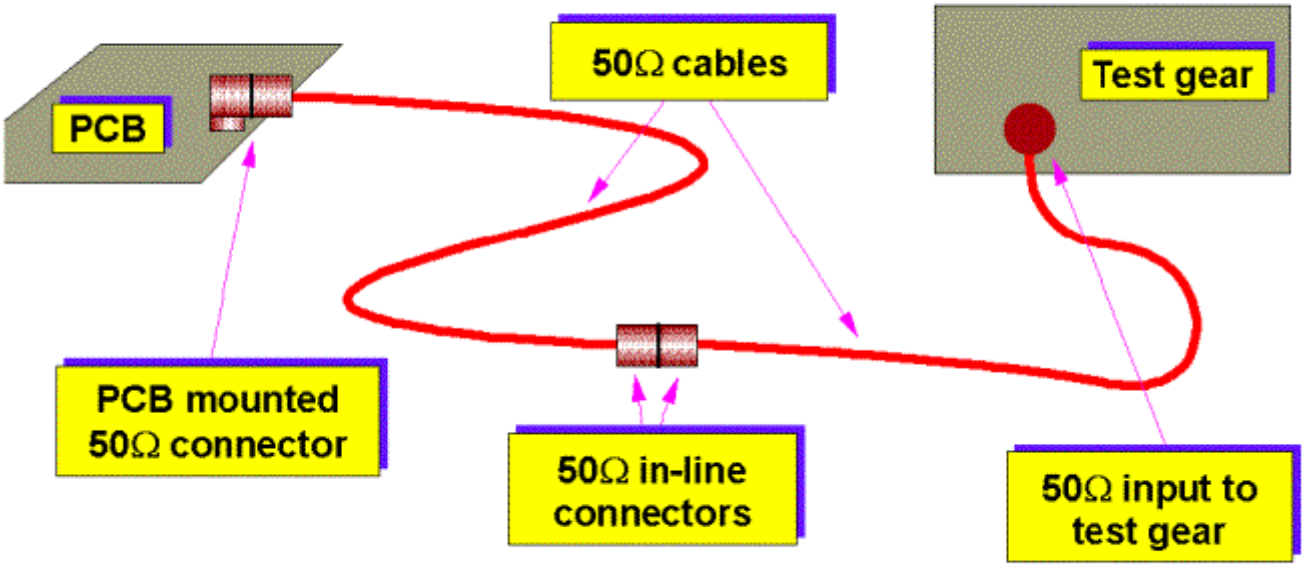

Impedances in the range 50 to 120Ω are available, especially for data communications. Manufacturers specify the impedances.

Lines must be matched, and the classical method is to match at both source and load. This provides maximum power transfer from source to load, but as it results in a 50% voltage loss for each interconnection, it is often not used for normal signal interconnections in non-RF equipment. Instead, transmission lines are often terminated at just one end, so as not to lose voltage, even though this is not ideal from either an EMC or signal integrity point of view. Terminating at one end only is a conscious decision to compromise on the engineering to save cost. Figures 2G, H, and J show the main termination methods.

Figure 2G Transmission line terminations

Classic RF termination

best, but loses 50% voltage

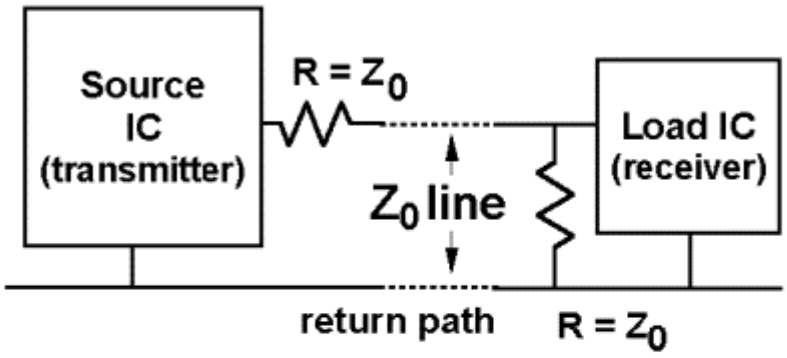

Series (source) termination

low power, only for reflected wave switching

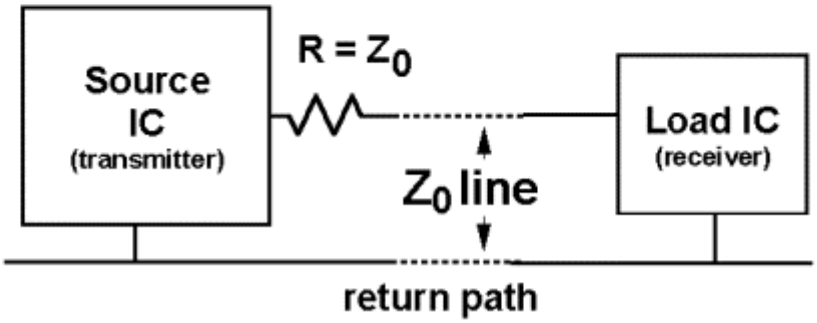

Figure 2H Transmission line terminations continued

for incident wave switching (fastest)

Parallel (shunt, or load) termination

best, but needs a lot of power

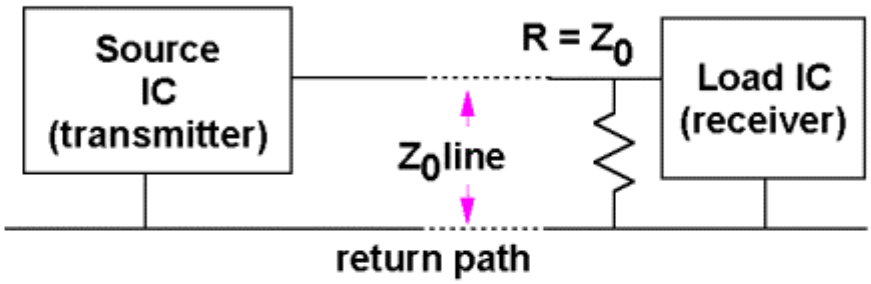

RC (parallel) termination

consumes less power

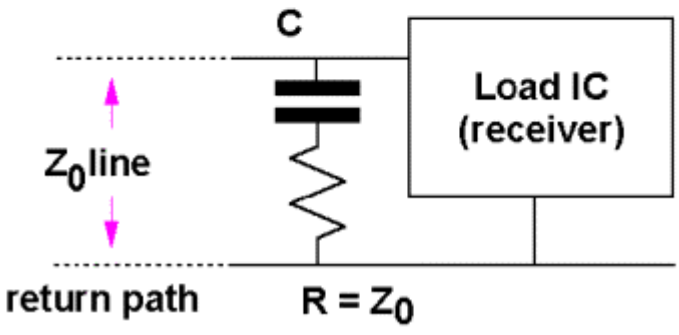

Figure 2J Thevenin (parallel) termination for incident wave switching

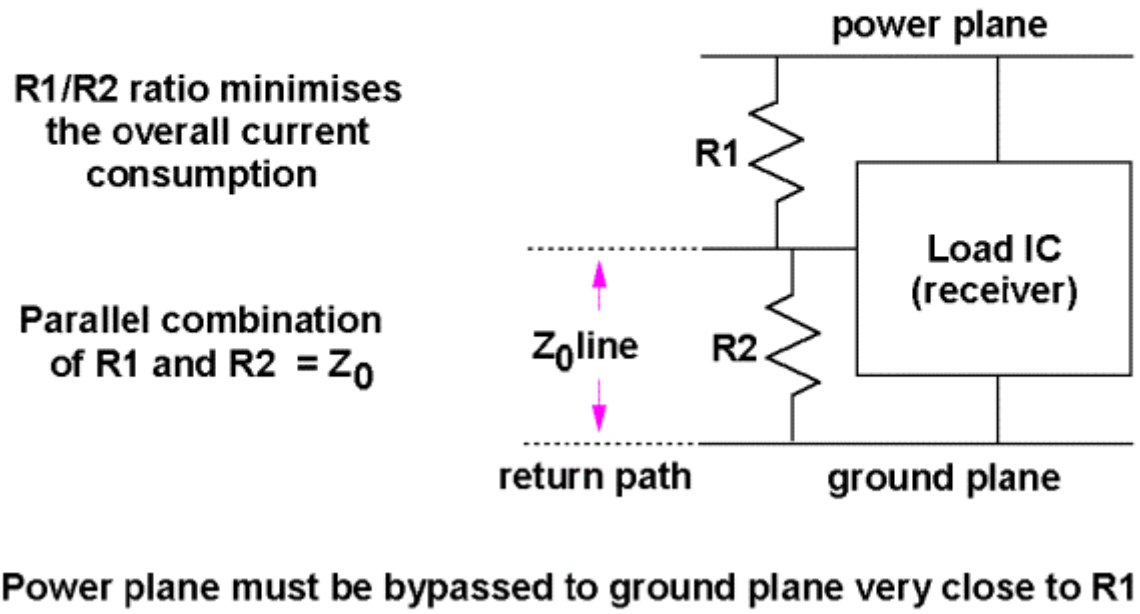

But nothing is perfect and even though not resonant, even the best practical transmission lines still leak a bit. Installation also reduces transmission line performance by causing variations in Z_0 (increasing leakage) when cables are bent sharply, crushed, strapped or clipped too tightly, repeatedly flexed, damaged, or fitted with inadequate connectors.

Unfortunately, the overall cost of creating transmission line cable interconnections with high enough quality at modern high frequencies can be very high. Flexible cables for microwave test equipment, for instance, can cost hundreds of pounds per metre. This is why, for GHz Ethernet to run on low-cost Cat 5 UTP (unscreened twisted pair), it has to use sophisticated DSP algorithms to reduce data rate and spread it randomly, and it still needs four pairs. So although transmission lines are very powerful, they are not a universal panacea for cable problems at high frequencies.

2.6.2 EMC considerations for conductors used inside and outside products

Inside a product – if the product's enclosure shields, and the screening and filtering of its external cables is good enough, almost any type of wire or cable can be used, although signal integrity will suffer. The problem here is that for high-performance digital or analogue electronics the cost of the enclosure shielding and filtering required can be so high that it would have been cheaper to use more expensive internal cables.

It is generally most cost-effective to *avoid all internal cables*, keeping all non-optical-fibre signals in the tracks of plugged-together PCBs (preferably a single PCB, even using flexi-rigid types). To make this work the PCBs need to be designed according to the Part 5 of this series, using a ground plane under all tracks. This generally reduces the cost of enclosure shielding and filtering to give the most cost-effective product, and because it also improves signal integrity it usually saves a couple of development iterations too.

Outside a product – unscreened cables with single-ended signals are now a serious liability whether the product is digital or analogue. Filtering digital signals does not help much to reduce emissions: single-ended drive produces copious common-mode currents *at the signal frequencies themselves*,

causing the product to fail conducted or radiated emissions tests depending on signal frequency. Any filtering would need to remove the signal, which does not help.

Filtering can work quite effectively for low-frequency analogue signals, but for precision beyond ±0.05% (12 bits) the cost of the filter and its board area increases rapidly. Of course filters have difficulty removing in-band interference (such as powerline hum) that a properly designed balanced communication system would easily reject.

2.6.3 Pairing send and return conductors

Even when not using transmission lines, always use paired conductors. Provide a dedicated return path for the return current as close as possible to the send path (and not via an earth or a screen). This works even when signals are single-ended and all their return conductors are bonded to a common reference potential. The fluxcompensation effect encourages return currents to flow in the path nearest to the send conductor, in preference to alternative current paths, and we can use this natural phenomenon to help keep the field patterns of our cables tight and reduce their E and M leakages. Figure 2K shows the general principle, which is of universal application.

Figure 2K Routing of the conductors in a cable
(this example is of switched circuits, but the principle is universally applicable)

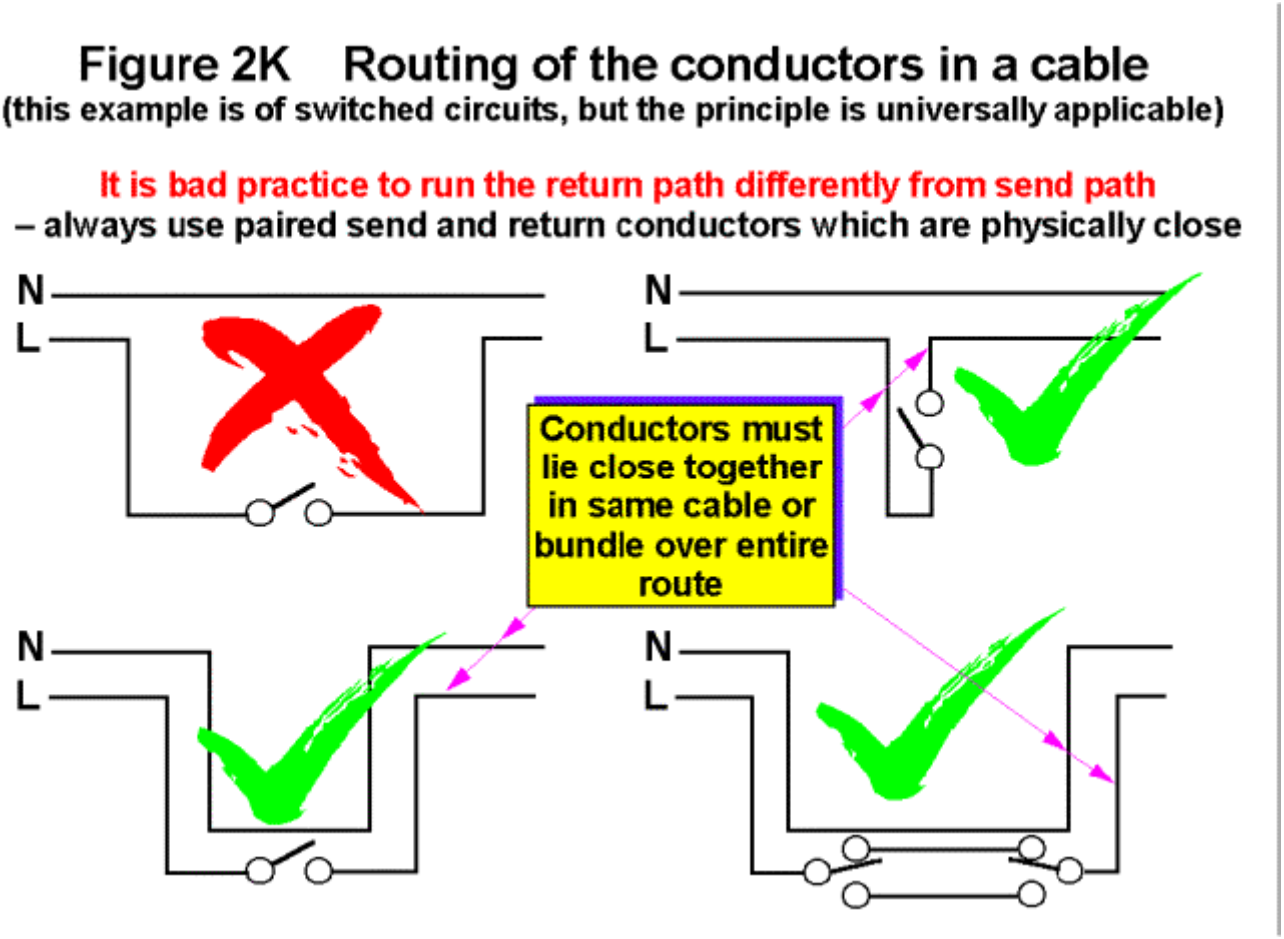

Figure 2K shows a mains supply with a switch in one line, but the same principle applies to signals.

The closeness of the send and return conductors over the entire current loop is absolutely crucial at the highest frequencies for circuits to work at all, never mind good EMC.

Ribbon cables carrying a number of single-ended (i.e. 0V referenced) signals are very poor indeed for EMC and signal integrity, but screening them results in stiff, bulky, expensive cable assemblies which is what flat cables were supposed to avoid.

Using the pairing technique for flat cables improves their EMC considerably and this conductor arrangement is the best:

return, signal, *return*, signal, *return*, etc.

A less effective alternative which is often recommended is:

return, signal, signal, *return,* signal, signal, *return*, etc.

Significant improvements can often be made by fitting flat cable ferrite clamps (common-mode chokes) at the source end(s), so that the conductor pairs behave as if they were driven from a balanced source at high frequencies, although proper balanced drive/receive circuits are better (see Part 1 of this series).

Twisted pairs are very much better than parallel pairs. Use twisted triples, quads, etc. where this is what it takes to get all the send and return paths of a signal in close proximity.

Twisted send and return conductors are strongly recommended for power cables: combining all phase and neutral conductors (two for single-phase, three for three-phase, four for three-phase plus neutral) in a single cable with a slow twist greatly reduces the emissions of powerline M field emissions. M fields from power busbars or individually routed phase and neutral cables can render whole areas of buildings unfit for CRT-based VDU monitors.

Twisted pairs using balanced circuitry (see Part 1 of this series) and common-mode chokes can be good for signals up to some tens of MHz, depending on the "balance" of the circuit, cable, and connectors. Any unbalance will convert some of the wanted signal into useless common-mode currents, which all leaks away as fields. Just a few micro-amps of common-mode can fail an emissions test. Tighter and more precisely regular twists make cables better for higher frequencies.

A great many types of twisted-pair cables are available, some intended for transmission lines (Z_0 will be specified). But twisted pair technology does not suit mass-termination. So-called "twist + flat" flat cable has multiple twisted pairs all formed into a ribbon, but has regular lengths of 100mm or so of parallel conductors for mass-terminating connectors – and the flat bits are so long they compromise EMC.

2.6.4 Getting the best from screened cables: the screen

There is no such thing as a cable that "complies with the EMC Directive", there are only cables with frequency-dependant screening performance.

Cable screens must cover the entire route with 360^o coverage. Making screening work effectively with low-cost these days is increasingly difficult, except for the least aggressive *and* least sensitive signals.

It is no longer best practice to use the shield of a cable as the signal return. The problem with co-axial cables is that the screen carries currents for both the signal return and external interference, and they use the skin effect to keep them on different sides of the screen (known as "tri-axial mode"). This works fine for solid copper screens (plumbing, to you and me), but flexible screened cables aren't very good at keeping the two currents apart, so return currents leak out, and interfering currents leak in.

But (I hear you say) all RF test equipment uses flexible co-ax, so it must be OK. Look carefully at these cables next time you are in an EMC test lab: the cables used for higher frequencies are very thick, stiff, and expensive, partly because they are double-screened at least. They use expensive screwed connectors (e.g. N-types), and are always used in matched (at both ends) 50Ω transmission lines. They are also treated reverentially and woe betide you if you tread on one. At higher frequencies than the average EMC test lab, semi-rigid or rigid co-axial cables have to be used, as stiff as automotive brake pipes.

The ability of a screened cable to prevent interference is measured in two ways, as shielding effectiveness (SE), and also as Z_T. SE seems obvious enough, and Z_T is simply the ratio of the voltage which appears on the centre conductor in response to an external RF current injected into the screen. For a high SE at a given frequency, we need a low Z_T. A flat Z_T of a few milliΩ over the whole frequency range would be ideal.

A very broad-brush summary of the screening qualities of typical types of screened cables follows, but remember that within each broad category there are many different makes and grades with different performances:

- Spiral wrapped foil is not terribly good at any frequency, and gets progressively worse > 1MHz.
- Longitudinal foil wrap is better than spiral foil.
- Single braid is better than foil at all frequencies, but still gets progressively worse > 10MHz.
- Braid over foil, double braid, or triple braid, are all better than single braid and all start to get progressively worse >100MHz
- Two or more insulated screens are better still, but only up to about 10MHz, at higher frequencies resonances between the screens can reduce their effectiveness to that of just one screen at some frequencies.
- Solid copper screens (e.g. semi-rigid, rigid, plumbing) are better than braid types and *their screening performance continually improves at higher frequencies*, unlike braid or foil, which always degrades above some frequency.
  Round metal conduit can be used to add a superb high-frequency performance screen. (Armouring is also useful as a screen but only at low frequencies, say up to a few MHz.)
- "Superscreened" cables use braid screens with a MuMetal or similar high-permeability wrap. These can be as good or even better than a solid copper screen, whilst still retaining some flexibility, but are expensive and suit applications where performance is more important than price (e.g. aerospace, military).
- I'm only aware of one manufacturer (Eupen) offering ferrite-loaded screened cables, which may offer improved high-frequency performance with good flexibility without the high cost of superscreened cables.

To reduce the bulk and cost of our screened cables and still get good EMC for high-performance modern products we need to use paired conductors for every signal and its return, preferably twisted pairs, just as described above for unscreened cables. Balanced drive/receive is also a great help.

2.6.5 Getting the best from screened cables: terminating the screen

Using co-axial cables and connecting their screens to a circuit 0V track is almost a guarantee of EMC disaster for high-performance digital and analogue products, for both emissions and immunity. Insulated BNC connectors on a product are usually a sign that all may not be well for EMC.

Cable screens should always be connected to their enclosure shield (even if they then go on to connect to circuit 0V), unless there are very good quantitative engineering and EMC reasons why not. "We've always done it this way" is not a reason.

Circuit development benches need to create the real structure of the product and the real interconnections with the outside world as closely as possible. Otherwise circuit designers may use various interconnection tricks to make their PCBs test well on the bench (I know, I used to do it too) – leaving it to someone else to sort out the resulting real-life application and EMC problems.

But even a high-quality screened cable is no good if the connection of the screen to the product is deficient. Cable screens need to be terminated in 360^{o} – a complete circumferential connection to the skin of the screened enclosure they are penetrating, so the connectors used are very important.

"Pigtails" should never be used, except where the screen is only needed up to a couple of MHz. Where pigtails are used they must be kept as short as assembly techniques allow, and splitting a pigtail into two on opposite sides also helps a bit. In the mid 1980's a company replaced all their pigtailed chassis-mount BNCs with crimped types for EMC reasons. Although the crimping tool cost around £600 they were surprised to find they quickly saved money because crimping was quicker and suffered fewer rejects. So pigtailing may be uneconomic as well as poor for EMC.

The "black magic" of cable screening is to understand that a cable's SE is compromised if its connectors, or the enclosure shields it is connected to, have a lower SE.

It is possible to use screened cables successfully with some unshielded products, if they use no internal wires and their PCBs are completely ground-planed with low-profile components. This is because the PCB ground plane, like any metal sheet, creates a zone of reduced field strength – a volumetric shield for a limited range of frequencies. Successful use of this technique depends upon

the electronic technology in the product, and is unlikely to be adequate for high-performance digital or analogue products. The cable screens would connect 360° to the PCB ground plane.

2.6.6 Terminating cable screens at both ends

This seems like heresy to some, but with the high frequencies in use these days leaving an end unterminated will usually leak too much. Screen termination at both ends also allows the screen to work on all orientations of magnetic fields.

Of course, connecting screens at both ends allows any earth potential differences to drive currents in the screen that might cause hum pickup, and even melt the cables. Where such earth currents exist it is an indication of poor facility earth-bonding which could allow earth faults or thunderstorm surges to destroy inadequately protected electronics. It is not unknown for screened cables to flash-over at unterminated ends during thunderstorms, creating obvious hazards.

It is sometimes recommended to electrically bond a cable screen at one end, and terminate the other with a small capacitor. The aim is to prevent excessive power frequency screen currents, and it does work to some extent although it is difficult to get capacitors to provide low enough inductance for very high frequencies and it does nothing for surge and flashover problems. Insulated BNC connectors are available with screen-to-chassis capacitors built-in, but the last price I had for types which worked well to 1GHz was nearly £20 each.

Galvanically isolated communications offer a way out of bonding screens at both ends, but such an approach needs careful attention to detail, especially the safety and reliability issues associated with installation earth-faults and surges. Metal-free fibre optics are the best kind of galvanically-isolated signal communications, and the easiest to use.

Section 2.5 above briefly lists the currently accepted best installation practices in achieving good EMC (e.g. using a PEC to divert heavy earth currents) *without* suffering hot cables *and* without compromising safety. Refer to IEC 61000-5-2:1997 for more detail.

2.7 Getting the best from connectors

Connectors suffer from many of the same EMC problems as cables, after all, they are just short lengths of conductor in a rigid body.

It is best to segregate connectors into those used for internal connections, and those used for outside connections, because of the possibility of flashover from outside to inside pins during surge or electrostatic discharge events. Such flashovers can bypass protective devices.

2.7.1 Unscreened connectors

Controlled-impedance transmission-line connectors are increasingly available for high-speed backplanes or cables.

When not using transmission lines, much improved EMC and signal integrity can be had from ordinary multiway connectors (e.g. DIN41612, screw-terminal strips) by making sure that each “send” pin has a return pin alongside. At the very least provide one return pin for every two signals. It is best when these are used with balanced signals, but the technique also helps when signals are single-ended.

2.7.2 Connectors between PCBs

Connectors between PCBs (e.g. daughterboard to motherboard) also benefit greatly from multiple 0V pins spread over their full length and width, and a similar arrangement of power pins also helps significantly. Optimum signal integrity and EMC is generally achieved (for signals sharing the same power rails) with a pin pattern that goes:

0V, signal, *+V, 0V*, signal, *+V, 0V*, signal, *+V, 0V*, etc.

The following pattern provides lower performance but is often adequate, and uses fewer pins:

0V, signal, *+V,* signal, *0V*, signal, *+V,* signal, *0V*, etc.

It is also best to extend the connector over the full length of the common edge between the two boards, and liberally sprinkle 0V and power connections over the full length. There is some evidence

that random pin allocations may provide better performance by breaking up standing wave patterns. Additional grounds between sensitive and noisy signals can help be a barrier to crosstalk.

We don't really like connectors on PCBs, and they should be avoided if possible (reliability will improve as a result) – as mentioned earlier, single or flexi-rigid PCBs with ground planes under all their tracks are better. Also: don't socket ICs (use field programmable PROMS). Each IC socket pin is a little antenna positioned right at the most vulnerable or noisy location possible.

2.7.3 Screened connectors

There is no such thing as a connector that "complies with the EMC Directive", there are only connectors with frequency-dependant screening performance. Screened cables must maintain 360° screen coverage over their whole length, including the backshells of their connectors at both ends. Connector backshells must make 360° electrical bonds to the enclosure shields they are mounted on, using iris springs or some other method. Saddle-clamps seem to make adequate screen bonds for most purposes in rectangular connector backshells, but avoid the ones that need the screen to made into a pigtail, however short. Figure 2L shows a typical D-type screen termination.

Figure 2L Example of 360° termination of cable screen in connector backshell (D-type shown)

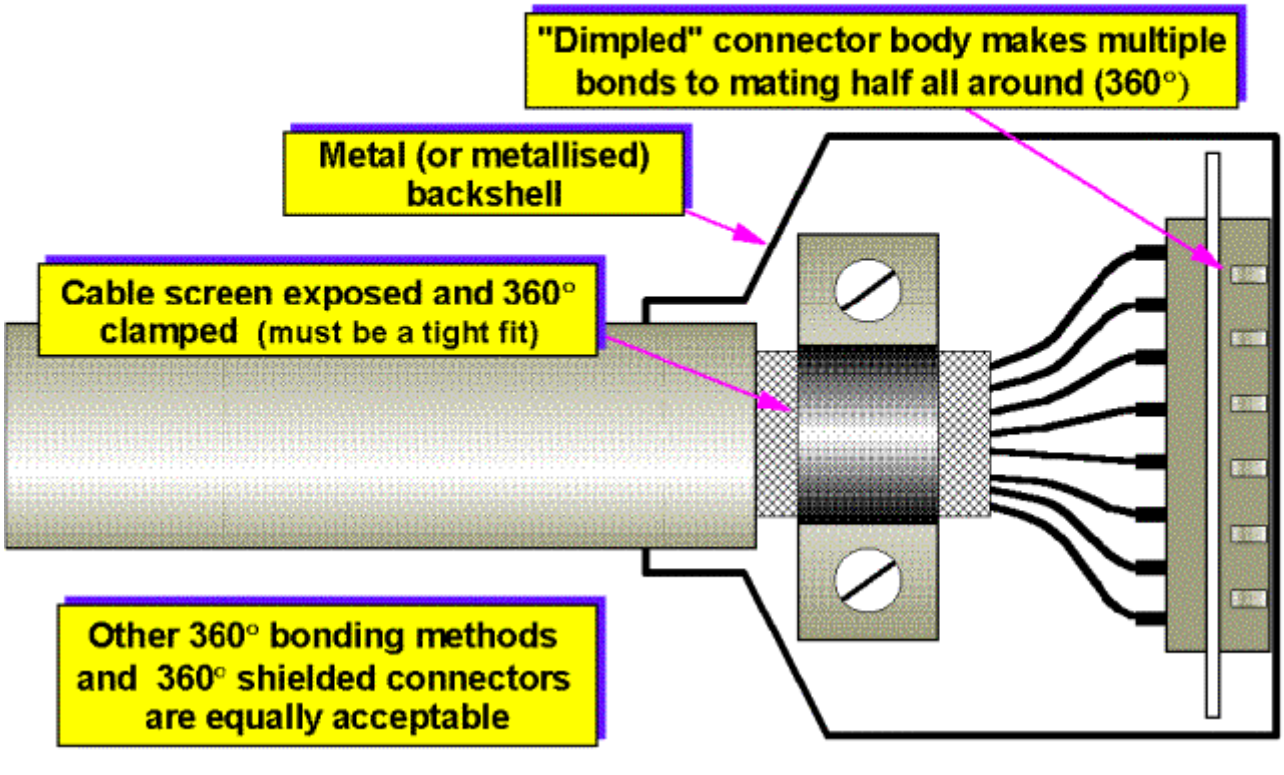

Co-axial and twinaxial screened cables benefit most from connectors with screwed metal backshells. These are better and more reliable at high frequencies than bayonet types (like BNC), which is why they are used on satellite TV convertor boxes.

Multiway screened cables are also best used with round connectors with screwed backshells, but are often specified with rectangular connectors such as D-types or larger. Making a 360° connection from cable screen to connector backshell, and from backshell to enclosure shield, seems a lot to ask of some connector manufacturers, so check this has been done effectively before committing to a particular make or type. Especially watch out for single (or long) springs, clips, and wires, when used to make a screen bond: these are just pigtails and will limit SE at high frequencies.

Unfortunately, many industry-standard connectors do not allow correct termination of cable shields (e.g. jack plugs, XLRs, phonos, and any number of proprietary connector models). Also unfortunately,

connector systems are still being designed without adequate thought being given to the need for simple 360° cable screen termination (e.g. RJ45, USB). The overall SE of a connector will be compromised if used with an enclosure shield or cable with a lower SE.

Figure 2M shows some of the important considerations when bonding connectors to shielded enclosures.

Figure 2M Cable screens entering a shielded cabinet

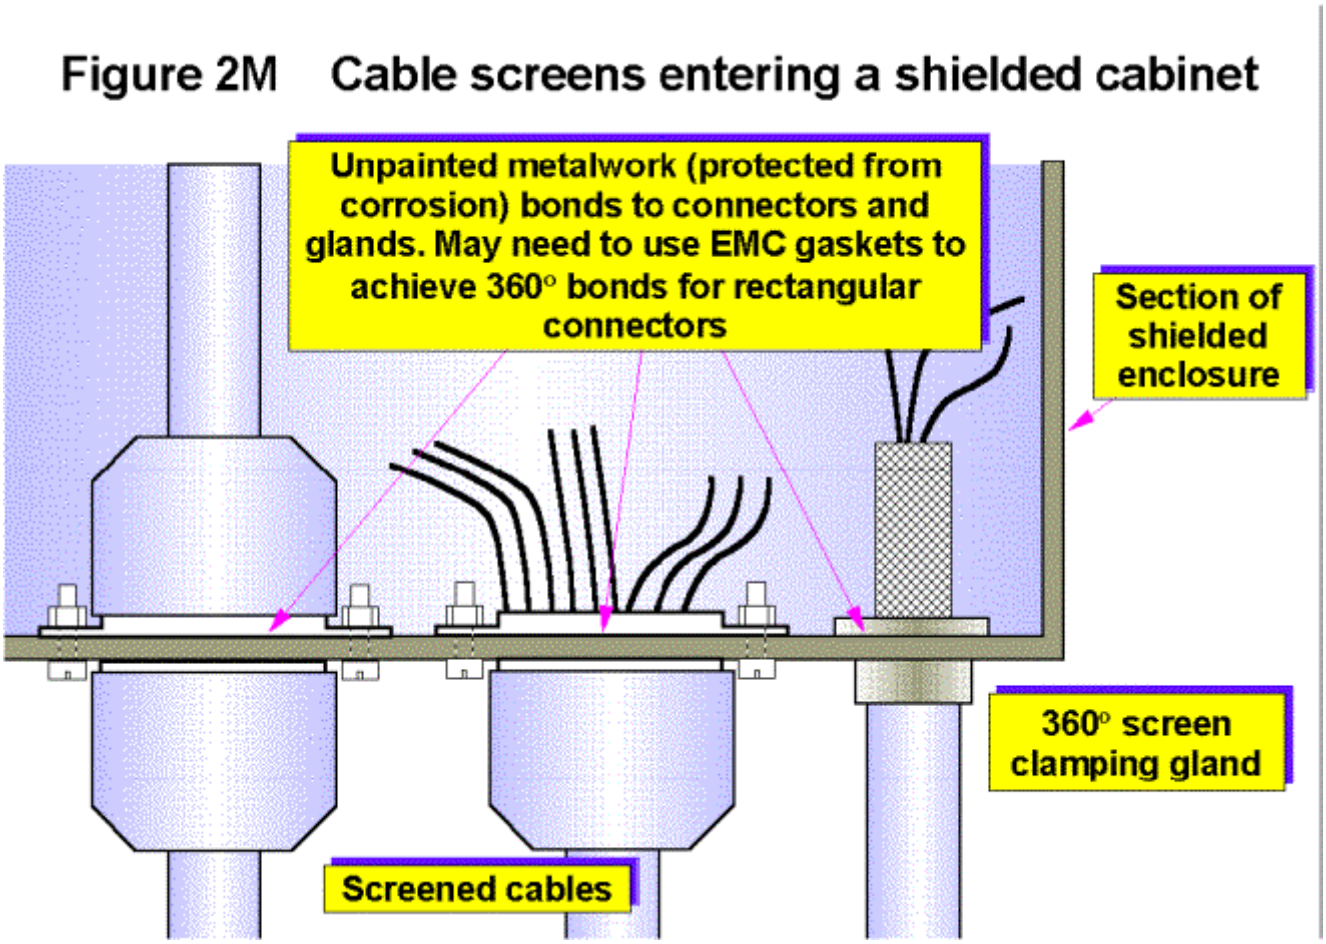

2.8 Further reading

There is a great deal more on cabling for EMC in the following:

Tim Williams, *EMC for Product Designers* 3rd edition, Newnes 2001, ISBN: 0-7506-4930-5, www.newnespress.com

IEC 61000-5-2:1997 Electromagnetic Compatibility (EMC) – Part 5: Installation and mitigation guidelines – Section 2: Earthing and cabling, www.iec.ch.

Tim Williams and Keith Armstrong, *EMC for Systems and Installations*, Newnes 2000, ISBN 0 7506 4167 3 www.newnespress.com, RS Components Part No. 377-6463.

Keith Armstrong, *EMC for Systems and Installations – Parts 1 through 6*, EMC+Compliance Journal, 2000. All UK EMC Journal and EMC+Compliance Journal articles are available electronically from the magazine archive at www.compliance-club.com .

Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone: 01457 871 605, Fax: 01457 820 145, Email: keith.armstrong@cherryclough.com

Design Techniques for EMC – Part 3
Filters and Surge Protection Devices

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the third in a series of six articles on best-practice EMC techniques in electrical/electronic/mechanical hardware design. The series is intended for the designer of electronic products, from building block units such as power supplies, single-board computers, and "industrial components" such as motor drives, through to stand-alone or networked products such computers, audio/video/TV, instruments, etc.

The techniques covered in the six articles are:

1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
2) Cables and connectors
3) **Filters and transient suppressors**
4) Shielding
5) PCB layout (including transmission lines)
6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), so this magazine article format can do no more than introduce the various issues and point to the most important of the best-practice techniques. Many of the techniques described in this series are also important for improving signal integrity, reducing the number of iterations during development, and reducing manufacturing costs.

Table of contents for this part

3. Filters and surge protection devices

Filters are used to attenuate unwanted frequencies travelling along conductors, and are characterised by attenuation versus frequency curves. Surge protection devices (SPDs) attenuate unwanted voltage surges travelling along conductors, and are characterised by graphs of voltage "let-through" versus time.

Incorrect use of filters or SPDs can make a product's emissions or immunity worse than if they were not used at all. More expensive filters or SPDs are not necessarily the best. You cannot in general choose a filter or SPD from a distributor's catalogue, only checking its ratings, number of circuits, and intended application, and expect it to provide the benefits you want.

3.1 Filter design, selection, and installation

The design and/or selection of filters can seem like a black art. It is not, but even when all the best efforts have been made along the lines described here, it is usually still necessary to try out a number of filters to find the optimum.

There are many books written on filter design. I have found Arthur B Williams book "Electronic Filter Design Handbook" very useful (McGraw Hill, 1981, ISBN 0-07-070430-9). No doubt there is a more modern edition available, but filter design has not changed much over the years. There are now a number of filter simulators which run happily on PCs.

I am not going to go into filter design in this article – but I will describe the things which need to be taken into account so that filters designed using textbooks, simulation programs, or chosen from catalogues, stand a chance of performing as required.

This article assumes that we are discussing filters to be fitted at the boundary between a product and its external environment. Filters for use inside a product, for example between a switch-mode power supply and a sensitive circuit, will share some if not all of the same considerations.

3.2 How filters work

A filter that works well in one situation can be poor in a different one, so we need at least a basic understanding of how filters work.

Filters work by creating a discontinuity in the impedance seen by a signal travelling along a conductor. The greater the discontinuity, the greater the attenuation. So if the impedance of an unwanted signal is 100Ω and we put a 1kΩ impedance in series with it, only about 10% of the signal gets through the high impedance – an attenuation of around 20dB. A similar effect can be created by instead shunting the 100Ω conductor with a low impedance: 10Ω would also provide an attenuation of around 20dB.

A variety of basic EMC (low-pass) filter types is shown in Figure 3A. Simple R-only and L-only filters create a series high impedance and are best used where the impedance of the unwanted signals is low. C-only filters create a shunt low impedance and so are best used where the impedance of the unwanted signals is high. Data sheet figures for C-only filters are rarely achieved in real life because they rely on the RF integrity of their earths (ideally zero impedance over all frequencies), and this is never achieved in reality.

Figure 3A
Different types of single-line filter

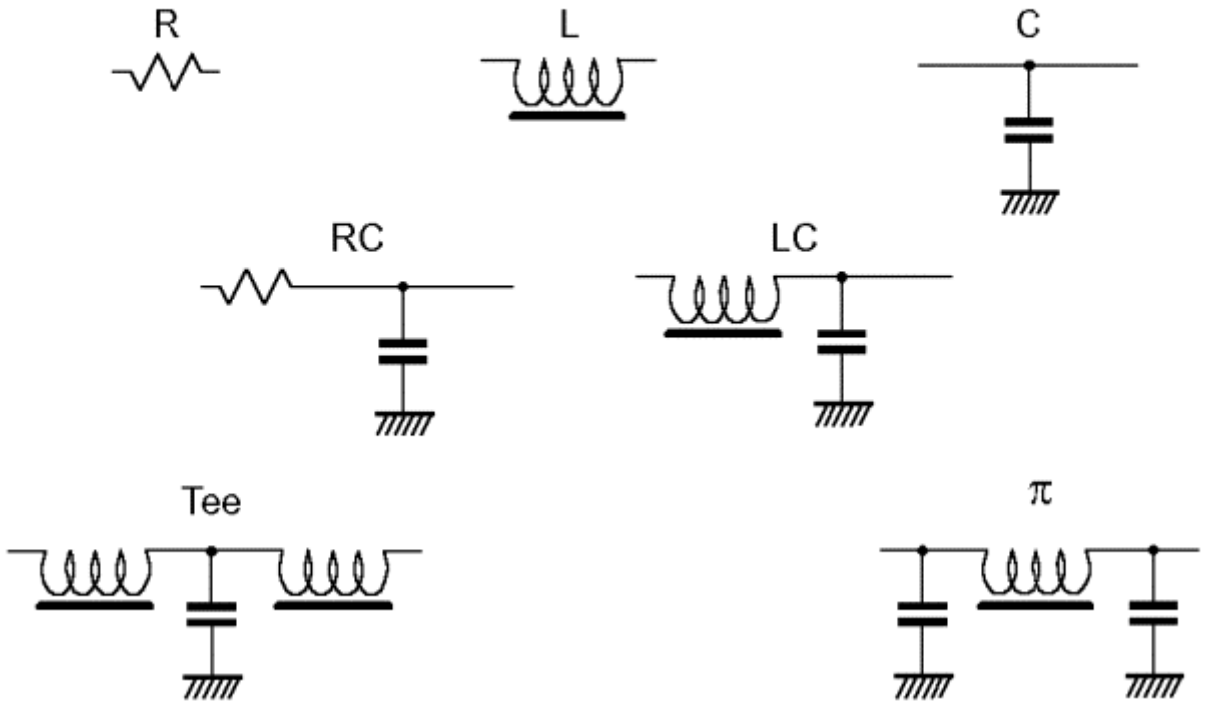

Resistors eventually lose their performance at high frequencies due to their stray shunt capacitance. Inductors also suffer from stray shunt capacitance, and this causes self-resonance and limits their high-frequency performance. The best inductors for EMC filter use have closed magnetic circuits (such as toroids, cylinders, etc., and other shapes with no air gaps), unfortunately this means that they are more likely to suffer from saturation effects at high currents (discussed later).

Capacitors suffer from intrinsic inductance, plus lead inductance, causing self-resonance and limiting high-frequency performance. The three-leaded capacitor has fewer problems with inductance (as long as its “earthy” lead is very short), but the ultimate elimination of inductive effects is only achieved by the feed-through capacitor, which has excellent high-frequency behaviour. Traditional feed-through capacitors are soldered or screwed into a shield wall, but some surface-mount devices now offer similar benefits when soldered to a PCB ground plane.

RC filters are the most predictable of all EMC filters because they do not resonate strongly. High values of R are best (say, from 1 to 10 kΩ) with low values of C (say 10nF or less). The most common use for RC filters is where a DC or low frequency signal from a low-ish source impedance is input to a high impedance circuit (the R connected to the source side, the C connected to the circuit side), and in these applications they can reliably provide very high attenuation at lowest cost.

LC, tee, and π filters can provide higher levels of attenuation with lower losses than filters using resistors, but are resonant circuits and are more sensitive to the impedances they are connected to.

Tee filters using resistors (RCR) are increasingly used to connect computer motherboards to displays and other peripherals which use a high data rate via ribbon cables or flexi-circuits. These are really C-only filters – the resistors are for terminating the transmission lines on either side of the capacitor to maintain signal integrity. The resistors are usually in the range 22 to 100Ω, and the filters are usually fabricated as arrays in small outline surface mount packages.

3.3 The advantages of soft ferrites

LC, tee, and π filters also suffer from the imperfections intrinsic to inductors and capacitors and their leads, which are a significant cause of the differences between simple calculations (or simulations) and real life. The use of soft ferrites for EMC filter inductors helps to ease these problems.

Soft ferrites are inductors at low frequencies, but become lossy (resistive) at higher frequencies. When used as an L-only filter soft ferrites usually provide between 3 and 20dB attenuation on typical cables, without significant resonances. When used in more complex filters, soft ferrites provide much more predictable performance (although some simulators may be unable to model their complex impedance-versus-frequency behaviour). Standard soft ferrite materials for EMC applications are still effective at 1,000MHz, with some recent components extending this to 2GHz.

A very wide range of soft ferrite RF suppression components exist, and are continually being added to, and recent offerings include SMD parts which provide 1kΩ or more around 100MHz, yet have DC resistances under 0.1Ω and are rated at 3A continuous. Choosing the right soft ferrite requires checking their impedance versus frequency graphs for the desired resistive impedance over the required frequency range. A true soft ferrite has no discontinuities in its impedance plot.

3.4 Common-mode (CM) and differential mode (DM)

To use filters effectively we need to know which conductors to connect our filter elements to, so we need to know the difference between common-mode and differential mode.

Figure 3B shows that wanted signals are almost always DM: the current flows *out* on one defined conductor, *back* on another, and these two conductors have a voltage between them. Sometimes more than two conductors are required for a signal and all of its current return paths (e.g. +5, signal, 0V; three-phase AC power), but the principle is the same.

Figure 3B
Differential-mode (wanted) signals and common-mode "leakage"

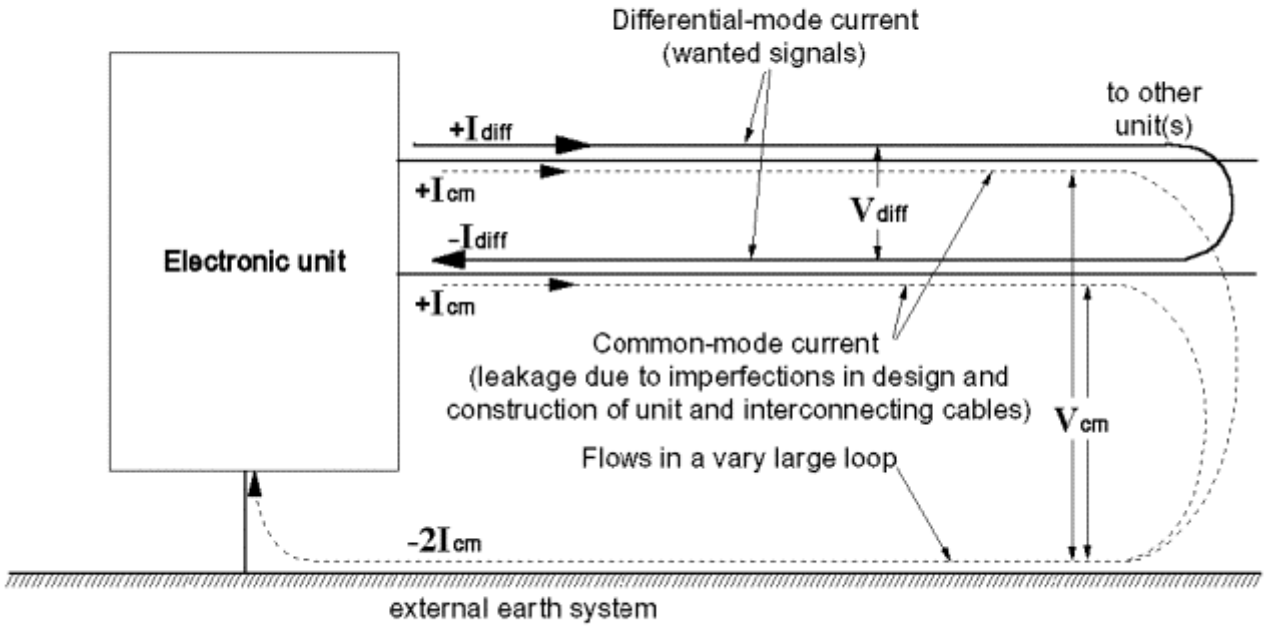

On the other hand, CM voltages and currents are identical for all the conductors in a cable (including any screens), and all CM currents flow in the same direction. The return path for CM currents is via the external earth system, other cables, metalwork, etc. CM currents and voltages are often very small in amplitude, but their large loop areas create worse EMC problems than equivalent DM signals.

CM is almost without exception caused by imperfections in design and construction which convert a fraction of the wanted (DM) signals to CM, causing problems with emissions. These same imperfections also convert a fraction of external CM interference into DM signals, worsening signal-to-noise ratios and/or digital signal integrity and causing problems with immunity. Cables are a significant contributor of DM to CM conversion, and the measure of this is called their "longitudinal conversion loss".

The powerful digital integrated circuits (ICs) used in many modern products create a lot of internal high-frequency switching noise due to ground and power "bounce". This leaks out of the ICs on all their connections and couples to PCB tracks, eventually appearing on I/O and all other cables as CM noise. This noise has a much higher frequency content and amplitude than would be expected from knowledge of the wanted signals, so is difficult to predict and is usually only quantified by measurement. (The techniques described in Part 1.1 and Part 5 will reduce this noise considerably.)

When a magnetic circuit (such as a ferrite toroid) is wrapped around a single conductor, it will attenuate all the currents in it, whether DM or CM. But when a magnetic circuit is wrapped around both (or all) the send and return conductors associated with a differential signal or power circuit, it will only attenuate CM currents. The magnetic fluxes created by the DM send and return current paths cancel out in a CM inductor, so (ideally) it has no effect on the wanted signal or power. In practice there is always some leakage inductance, hence always some attenuation of DM currents, and this can be used to provide both CM and DM filtering in one component. Figure 3C shows some aspects of common-mode choke filters.

Figure 3C
Common-mode filtering

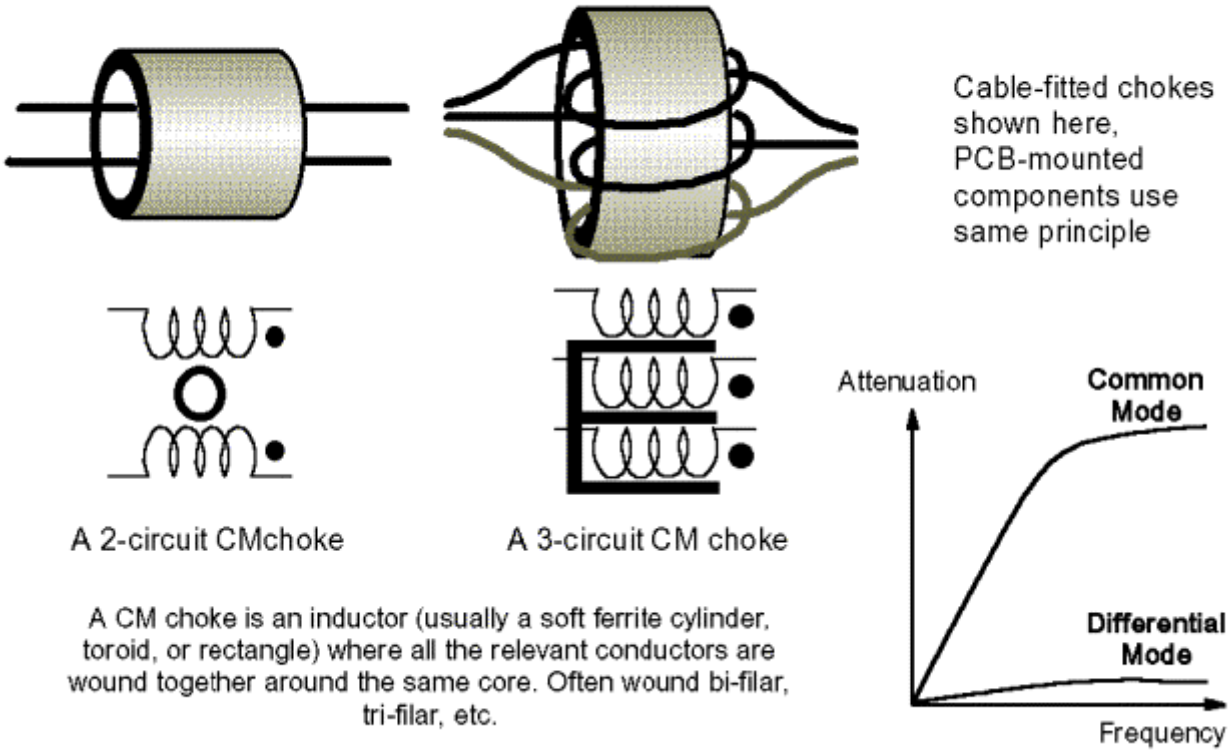

A CM choke is an inductor (usually a soft ferrite cylinder, toroid, or rectangle) where all the relevant conductors are wound together around the same core. Often wound bi-filar, tri-filar, etc.

For most EMC work, common-mode chokes are made from soft ferrite. The cancellation of DM flux in CM chokes (often called a current-balanced choke) allows large inductance values (milliHenries) to be achieved with small components, whereas DM chokes the same size are measured in microHenries, and become physically larger as currents increase.

Wrapping a set of conductors several times around the same magnetic circuit increases CM impedance, but is not effective at high frequencies because stray capacitances "short out" the choke. For higher attenuation at very high frequencies it is better to string a number of ferrite tubes or toroids along a cable, with only one pass of the cable through each. (Surface mounted ferrite beads can achieve high impedances at high frequencies because their parasitic capacitances are so small.)

A useful soft ferrite component is a cylinder split lengthways and held in a plastic "clip-on" housing. They are very easy to apply to cables (and to remove when found not to do much), and are available in a wide range of shapes and sizes, including flat type for ribbon cables. All EMC engineers' carry many of these with them, using them for diagnosis and isolation of EMC problems, as well as cures.

Figure 3D shows how common-mode and differential-mode filtering is combined in a number of simple filters.

Figure 3D
Simple single-stage filters which combine CM and DM techniques
Mains filters use X capacitors for DM filtering (safety-rated for line-to-line use) and Y capacitors for CM filtering (safety-rated for line-to-ground use)

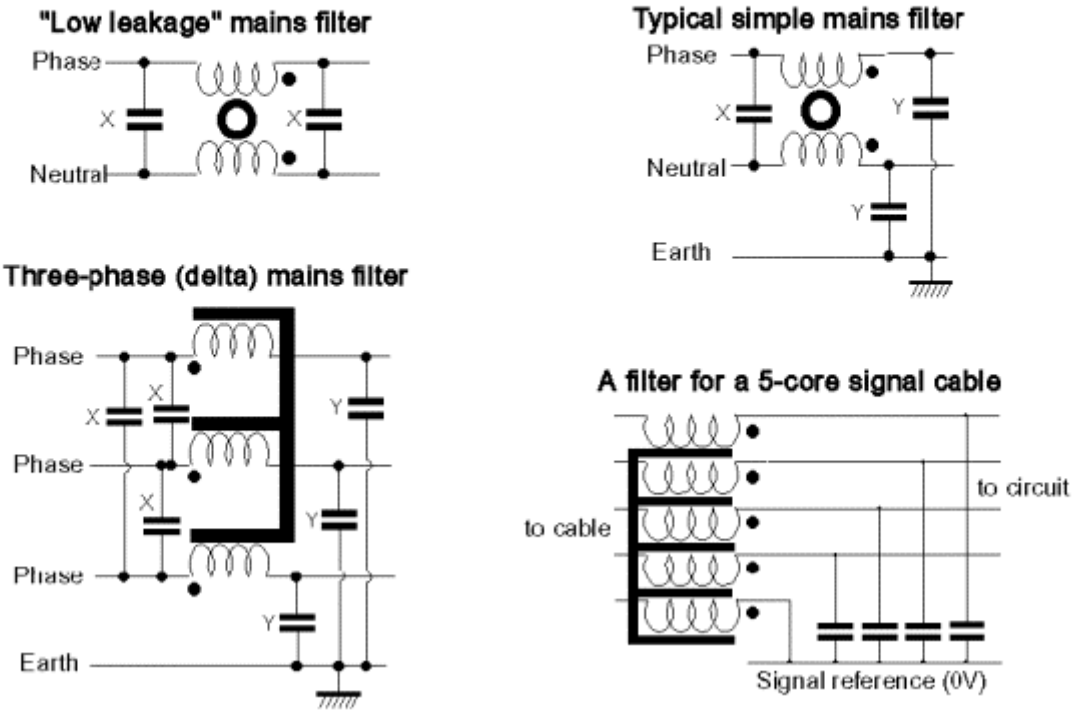

3.5 Crude filter rules-of-thumb

So we can now define our crude filter rules-of-thumb:

- If unwanted signal is high impedance and DM, attenuate with a shunt capacitor (connected between its send and return conductors)

- If the unwanted signal is low impedance and DM, attenuate with a series inductor (good high-frequency performance requires a symmetrical layout of identical inductors, one in the send conductor, and one in the return)
- If the unwanted signal is high impedance and CM, attenuate with identical shunt capacitors from each conductor to the local earth reference (usually the chassis)
- If the unwanted signal is low impedance and CM, attenuate with a common-mode inductor applied to all the signal conductors at the same time.

These rules are very crude, because the terms “low” and “high” impedances are relatively ill-defined, and depend upon the impedances of the suppression components available for handling the wanted signal, (and also on the components’ costs and availability, the quality of earth reference and ease of bonding to it, etc.).

3.6 Inductance variation with current

All inductors suffer a reduction in their inductance as current increases, up until they saturate (where they have no inductance at all). This is a common cause of the differences between simple calculations or simulations and real-life filter performance.

As a rule of thumb, this effect should be taken into account whenever the wanted current exceeds 20% of the rated current. Take account of the fact that power supplies that do not meet EN61000-3-2 draw their input currents as peaks many times higher than their rated RMS supply current.

3.7 Determining filter specifications

For controlling emissions: the filter performance required can be estimated by comparing the spectrum of the product’s emissions with the limits in the relevant EMC standard. The emissions may be either predicted or measured, and the limits are most often related to EN 55022 or EN 55011.

For controlling immunity: the filter performance required can be estimated by comparing the specification of the threats in the intended electromagnetic environment with the susceptibility of the electronic circuits to be protected. The functional performance degradation allowed should also need to be taken into account. The environment specification is usually taken from an EMC standard, often derived from EN 50082-1 (preferably the 1997 version) or EN 50082-2. But even EN 50082-2 may be inadequate in some industrial, scientific, or medical environments where high levels of 50Hz or radio-frequency (RF) power are being used, or when the user expects to use a portable radio transmitter whilst operating the product.

Where safety-critical systems are concerned no functional degradation is permitted during interference events, and a Safety Integrity Level (SIL) should be determined (for e.g. using the new IEC 61508) and used to increase the immunity test level accordingly to achieve the desired level of risk.

This all sounds very organised, but EMC should be designed-in from the start of any project for the greatest cost-effectiveness and we usually don’t know what the actual emissions or susceptibility are until we have built something and tested it, by which time it is fairly late in the project.

The answer to this is to assume that all conductors will need filtering to some degree. But we still need to know: what frequencies? and to what degree?

Sadly, most actual emissions are caused by unwanted CM voltages and currents. Immunity is a similar story: we can specify the frequency range and threat levels, but most problems are caused by CM interference being converted to DM and polluting the signal. Since the conversion from DM to CM, or CM to DM, is caused by imperfections we can’t easily predict filter specifications. (Most of the design techniques discussed in this series will reduce these imperfections and hence reduce the conversion between DM/CM and CM/DM.)

Murphy's Law ensures that when you have thought of everything, the expensive options will not be needed and you will be damned for over-engineering. But if you overlook any possibility Murphy will expose it and you will be damned for that. Since we are bound to be damned whatever we do, we may as well make our lives a lot easier by including a number of filter options in our initial designs.

When a product is first tested for EMC (long before production drawings are produced) some/all of the filters may be linked out at first, or simple inexpensive filters fitted. Anti-Murphy precautions then require having a wide range of alternative filter types and complexity handy, as well as the tools to fit them quickly. This is why all EMC engineers and test labs have stacks of sample boxes from filter manufacturers, overflowing toolboxes, with soldering and de-soldering irons already warmed-up.

Happily, experience with filtering various electronic technologies to meet various EMC standards is soon gained, and most engineers soon learn which filters usually work best for the different types of conductors in their products. Be aware, though, that every new product has its quirks (even related to its mechanical assembly) and a filter that worked on Model 1 might not be adequate for Model 2. So *always* make provision (at least on the early prototypes) for more expensive and larger filters than you hope to use, and only remove this insurance when everyone else's design is complete (including software, although this is probably a vain hope) and the product passes its EMC tests with a suitable margin.

Don't forget that there are inevitable unit-to-unit variations, so for a serially-manufactured product a prudent designer will aim for a 6dB "engineering margin" on emissions and immunity tests, at least.

3.8 Problems with real-life impedances

Most filter data comes from tests done with 50Ω source and load impedances, which leads us to a very important point – filter specifications are always hopelessly optimistic when compared with their performance in real life.

Consider a typical supply filter, installed at the AC power input to the DC power supply of an electronic apparatus. The CM and DM impedances of the AC supply can vary from 2 to 2,000Ω during the day depending on the loads that are connected to it and the frequency of interest. The DM impedance of the AC-DC converter circuitry looks like a short-circuit when the rectifiers are turned on at the peaks of the waveform, but otherwise looks like an open-circuit. The CM impedance of the DC power supply's AC input is very high indeed, due its isolation from earth for safety reasons (this is why most mains filters connect Y capacitors from line to earth on the equipment side of a mains filter: to create the maximum impedance discontinuity). This is clearly very far from being a matched 50Ω / 50Ω situation.

Because filters are made from inductors and capacitors they are resonant circuits, and their performance and resonance can depend critically on their source and load impedances. An expensive filter with excellent 50/50Ω performance may actually give *worse* results in practice than a cheaper one with a mediocre 50/50Ω specification.

Filters with a single stage (such as those in Figure 3D) are very sensitive to source and load impedances. Such filters can easily give *gain*, rather than attenuation, when operated with source and load impedances other than 50Ω. This *filter gain* usually pops up in the 150kHz to 10MHz region and can be as bad as 10 or 20dB, so it is possible that fitting an unsuitable mains filter can increase emissions and/or worsen susceptibility.

Filters with two or more stages, such as those in Figure 3E, maintain an internal circuit node at an impedance which does not depend very much on source or load impedances, so provide a performance at least vaguely in line with their 50/50Ω specifications. Of course, they are larger and cost more.

Figure 3E
Examples of 2-stage mains filters

A typical 2-stage filter for a digital product
(suffers from high levels of high-frequency CM emissions)

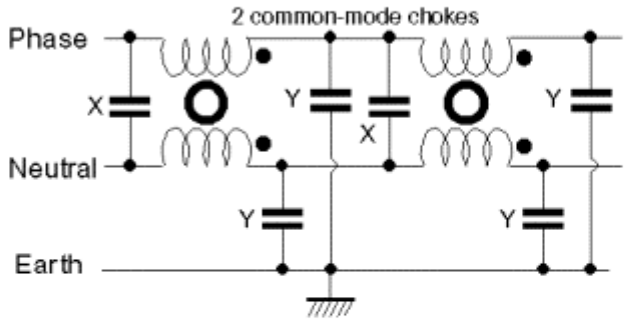

A typical filter for a switch-mode power converter
(suffers from high levels of low-frequency DM emissions)

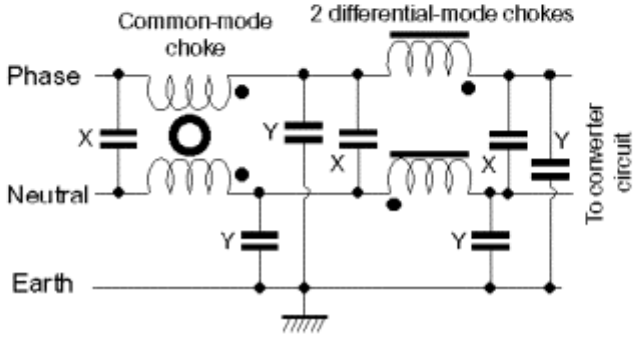

The best way to deal with the source/load impedance problem is to only purchase filters whose manufacturers specify both CM (sometimes called “asymmetrical”) and DM (sometimes called “symmetrical”) performance, for both matched 50/50Ω and mismatched sources and loads.

Mismatched figures are taken with 0.1Ω source and 100Ω load, and vice versa, using the CISPR17 test standard that is also used for 50/50Ω tests. Combining all the worst-cases of all the different curves results in a filter specification that may be relied upon, providing the filter is not overloaded with current (as discussed above), and earthed properly (as discussed below). An example of extracting the worst-case filter curve is sketched in figure 3F.

Figure 3F
Deriving reliable filter attenuation figures from manufacturer's data

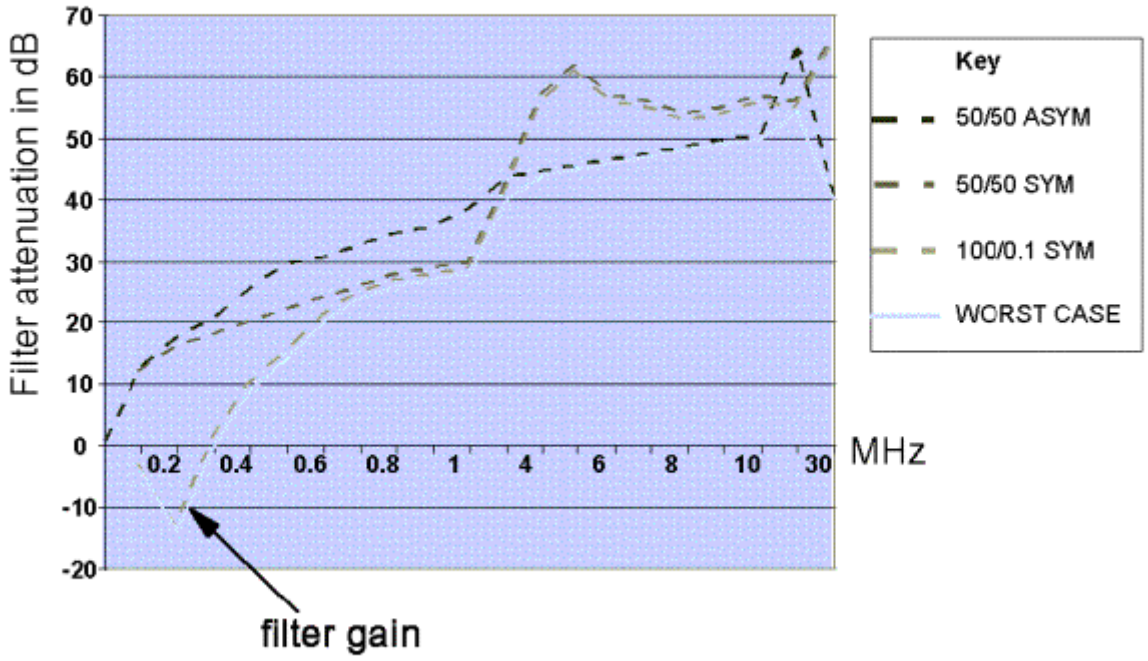

3.9 Earth leakage currents, and safety

Most supply filters use Y-rated capacitors between phases and earth, with values around a few nF not to exceed the earth leakage limits in the relevant safety standard. Fixed equipment permanently wired-in is allowed higher earth leakage currents, up to 5% of phase current in some cases (when appropriate warning labels are fitted). Industrial power conversion equipment can have very high levels of emissions, and often requires large filter capacitors and hence large earth leakage currents. This is one area where EMC and Safety considerations are unavoidably intertwined, and of course safety wins, so the relevant safety standards must always be referred to when designing mains filters, remembering that most filter capacitors have tolerances of ±20%.

For medical apparatus which may be connected to patients, earth-leakage currents may be limited by safety standards to such low levels that the use of any reasonable size of Y capacitor is impossible. Such filters tend to use better CM chokes to achieve the same performance without Y capacitors, and/or more stages, so tend to be larger and more expensive.

In systems, the earth leakages from numbers of Y capacitors (even small ones) can create large earth currents. These can cause earth voltage differences which impose hum and high levels of transients on cables between different equipments. Modern best-EMC-practices require equipotential three-dimensional meshed earth bonding, but many older installations do not have this so apparatus intended for systems in older buildings may benefit from the use of low leakage filters.

It is always best to use mains filters (or components) for which third-party safety approval certificates have been obtained *and* checked for authenticity, filter model and variant, temperature range, voltage and current ratings, and the application of the correct safety standard.

Filters sold for use on 50/60Hz may generally be used on power ranging from DC to 400Hz with the same performance, but it is best to check with the manufacturers beforehand. Also remember that

earth leakage currents will increase as the supply frequency increases, so filters which just meet safety standards at 50Hz may not meet them at 60Hz, and may be decidedly dangerous on 400Hz.

3.10 Issues of frequency and/or sensitivity of wanted signals

Most EMC filters are low-pass. Power supply filters have their design difficulties, but do benefit from their wanted signal (DC, or 50/60Hz) being very much lower than the frequencies of most types of interference. Where signals are digital or high-level analogue, and not very high frequency or very sensitive, simple R, L, C, RC, LC, tee, or π filters are often used, as shown by figure 3A.

But where emission/immunity frequencies overlap with, or are close to, the wanted signals, it is no good fitting DM filters such as those in Figure 3A – filtering out the unwanted signals will eliminate the interference, but will also eliminate the wanted signals. Screened cables and connectors will be required instead of/as well as filtering.

At high data rates single-ended signals can prove very difficult for EMC, even with very expensive (thick and inflexible) cables. The use of balanced drive/receive circuits (described in Part 1.4) with balanced cabling makes the filtering and screening of high-rate signals much easier, reduces cable costs (and thicknesses), and makes EMC compliance much easier. CM filters, rather than DM, may then be used *within* the spectrum of the wanted signals. This is an example of good thinking at an early design stage to minimise overall project timescales and manufacturing cost, even though the component cost of the functional circuit may not be minimised. Examples abound of cheap functional circuit designs that incur huge costs and delays when the time comes to make them EMC compliant or to fix an interference problem in a system.

Low-frequency instrumentation, audio, and other sensitive analogue signals may need to use multi-stage filters to achieve the desired immunity, unless adequate screening has been applied over the entire length of the conductor (unfortunately, good RF screening is not “traditional” in industries that still think cable shields should be bonded to earth at only one end, see Part 2.6.6).

Where an electronics module has a sensitive input, high-performance filters are often needed on *all* its inputs, outputs, and power conductors (unless the sensitive internal circuitry has been adequately protected with internal filtering and shielding from the rest of its circuits, discussed in Part 5).

3.11 Filter earthing

One of the secrets of RF filters which use capacitors connected to earth, is that they can never be a lot better than the RF performance of the reference (almost always earth or 0V) they are connected to. Most earths in domestic, commercial, and industrial applications have poor RF performance and are nothing like an ideal “infinite RF sink”.

The best place for mounting a filter for the purposes of this article is at the boundary between the product’s “inside world” and the cables in its “outside world”. For a shielded enclosure, filters should be RF-bonded (i.e. metal-to-metal) to an external surfaces, preferably using a through-bulkhead style of filter. For an unshielded enclosure, filters are generally best bonded to the printed-circuit board’s ground plane, at one edge of the PCB.

The connection between the capacitors in the filters and whatever is being used for its RF reference should be very short and direct, less than one-hundredth of a wavelength long at the highest frequency to be attenuated, and should also have a very low inductance. This means that wires cannot be used as filter grounds except for low frequencies (say, below 1MHz), even if they do have green/yellow insulation (electricity is colour-blind). For example, if a supply filter with 2.2nF Y capacitors is earthed solely by a 100mm long wire, its Y capacitors will be rendered completely ineffective at frequencies above 20MHz by the inductance of that wire.

1nH per millimetre is a good rule-of-thumb when calculating the effects of wired connections to earth. The only correct bonding for filters is at least one (preferably more) direct metal-metal connection(s) from the filter’s metal body to the earth reference. It is acceptable to fit green/yellow

wires to filters for safety reasons, as long as they are in parallel with at least one good RF earth bond.

Military signal filters tend to rely on C-only and π types, apparently because most traditional military equipment has a substantial and well-engineered RF earth reference (die-cast alochromed boxes bolted firmly to metal bulkheads in all-metal bodied machines and vehicles), so their earthed capacitors do not suffer from poor earth RF integrity.

Unfortunately, RF earth integrity is often a serious problem for domestic, commercial, and industrial products, which are constructed with low-cost in mind. The most predictable filters in these applications tend to be RC, LC, or tee types (using soft ferrites for the L components). These impose lower levels of RF currents on the earth reference than C-only or π filters. As military vehicles use new materials such as carbon fibre, their earth reference becomes less effective and they may find R, L, RC, LC, and tee filters more cost-effective than C or π.

D-type and some other connectors are available with a wide range of filter and shielding options, and because they are so easy to apply remedial EMC improvements to they are a good choice for low-current and signal cables. Most of these use capacitors from each signal/power pin to their earthed metal bodies. There are some D-types available with a soft-ferrite tube around each pin, and even some with LC, tee, or π filters for each pin. Pin filters treat CM and DM currents identically, so may not be suitable where wanted signals have a high frequency.

RJ45 and similar telecommunication or Ethernet connectors are available with built-in common-mode chokes. The baluns and pulse transformers often used in high-speed LANs to reject low-frequency common-mode noise and/or provide galvanic isolation, are sometimes available combined with common-mode chokes for better rejection of high frequency noise.

3.12 The synergy of filters and shielding

If a supply filter allows a very high frequency (say 300MHz, harmonics from digital processing) to pass through to the product's mains lead, the radiation of these frequencies from the lead will compromise the product's *radiated* emissions.

So another secret of cost-effective EMC is that the correct way to view filtering and shielding is as a synergy, with each one complementing the other. Incorrect filter design or mounting technique can easily compromise radiated emissions and immunity. Likewise, inadequate shielding can easily compromise conducted emissions and immunity.

Some manufacturers only make their filters to work over the frequency range of the conducted emissions tests (up to 30MHz), to keep costs low. Unfortunately such filters often compromise the shielding integrity of shielded enclosures and cause problems with radiated electromagnetic disturbances.

3.13 Filter construction, mounting, and cabling

Filters are easily compromised by RF coupling from the conductors on their unfiltered side to the conductors on their filtered side. Many engineers have been very surprised by the ease with which high frequencies will "leak around" a filter, given half a chance.

Where an external cable to be filtered enters a shielded enclosure or room, the filter should be fixed into the metal wall at the point of cable entry and RF bonded (metal-to-metal) to the metalwork of the wall. For better performance at high frequencies, and to avoid compromising enclosure shielding, a conductive gasket or spring fingers may be fitted to create a low-impedance electrical bond between the filter's metal body and the wall along the entire circumference of the filter's cut-out.

An IEC inlet filter installed in a shielded enclosure can only give a good account of itself at frequencies above a few tens of MHz *if* its body has a seamless construction *and* its body is RF bonded to the shielding metalwork, as shown in Figure 3G.

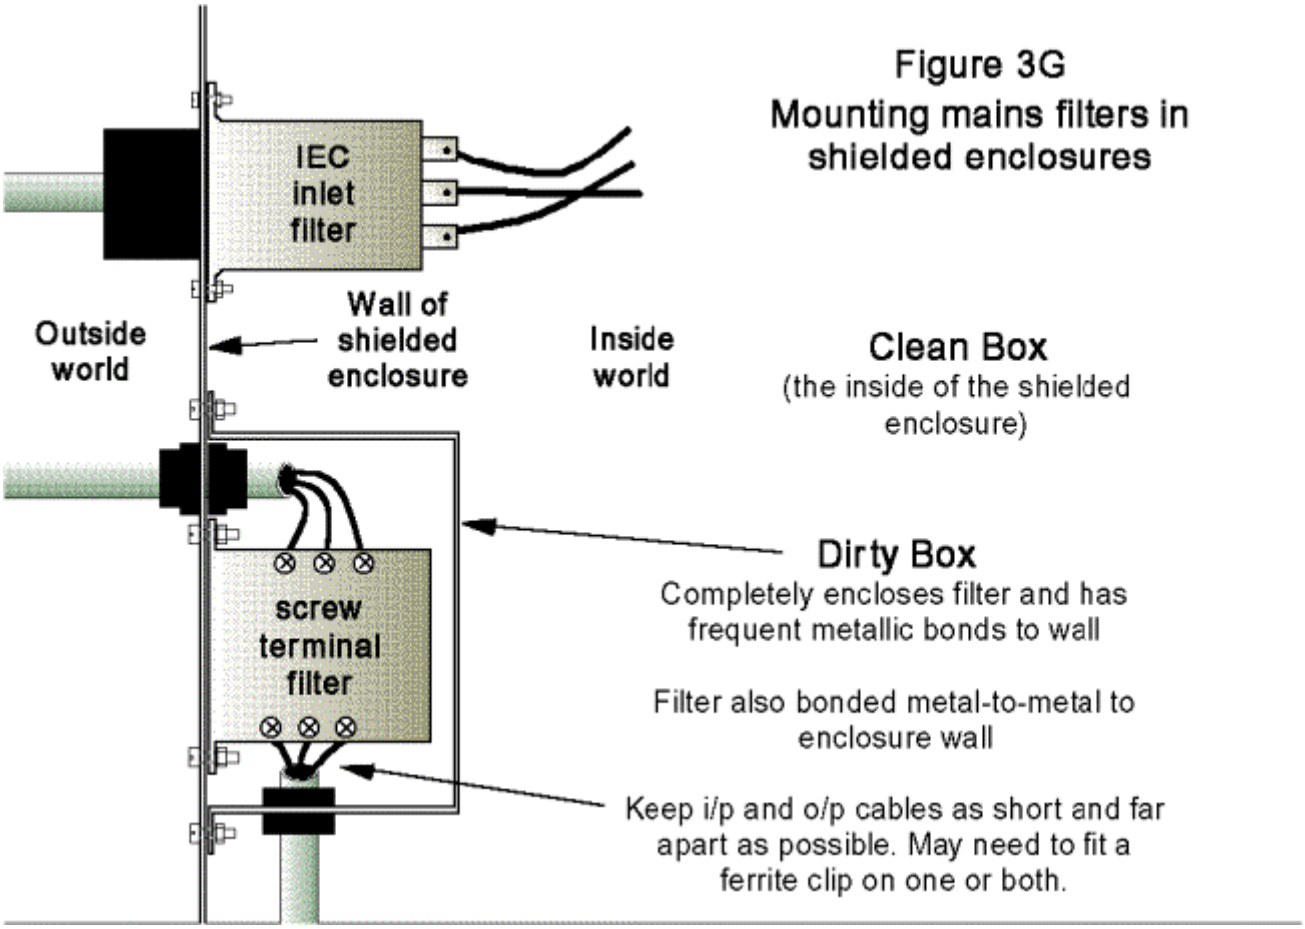

Figure 3G
Mounting mains filters in shielded enclosures

Clean Box
(the inside of the shielded enclosure)

Dirty Box
Completely encloses filter and has frequent metallic bonds to wall

Filter also bonded metal-to-metal to enclosure wall

Keep i/p and o/p cables as short and far apart as possible. May need to fit a ferrite clip on one or both.

Through-bulkhead filters are best, but may be too expensive for some applications, such as mains currents above 10A (the maximum rating of the IEC 320 style mains connector). For higher powers most commercially available mains filters are just rectangular units with screw-terminal connections. Figure 3G shows how to mount such filters using the "dirty-box" method, which encloses the filter in its own box within the main shielded enclosure (the “clean box”) to help achieve good high-frequency performance.

The filter input and output cables in the dirty box must be very short and far away from each other, but even so, very high frequencies may still couple between them and soft-ferrite cylinders may be needed on either (or both) cables.

Filters which are not through-bulkhead types but must have the highest performance often deal with cable coupling by enclosing their input and output terminals inside their metal enclosure (a “dirty box”) and bringing their cables out through standard circular galvanised conduit fittings. Running the cables in conduit effectively shields the unfiltered from the filtered cables, and the filter functions effectively up to the highest frequencies. This is the method used by most of the supply filters intended for EMC test chamber applications. Input and output cables may be screened instead of being run in conduit for the same effect.

3.14 Surge protection devices (SPDs)

3.14.1 Types of SPD

Surge arrestors are variable resistance devices, whose resistance is a function of the applied voltage. They are designed so that they provide a clamping effect when the voltage across them exceeds a certain level, rather like a zener diode.

There are four basic types of SPD:

- Gas discharge tube (GDT), essentially just a spark gap, slow but very high power
- Metal-oxide varistor (MOV), fast and available in a wide range of energy ratings
- Avalanche devices, semiconductors with a zener type action, very fast but not very high power
- SCR devices, another type of semiconductor device, slow but will handle high currents

Figure 3H sketches the voltage/time curves of these four types of SPD when exposed to the leading edge of a typical surge test waveform. It shows that GDT and SCR devices are slow to start suppressing. They have to reach a trigger voltage before they begin to conduct, and during this time they let through surge voltages which may be potentially damaging. They also have a foldback characteristic, which means that once they are triggered and are carrying current, the voltage across them drops to well below the voltage they were previously quite happy to block. Consequently, careful design is needed to make sure that when they are connected to a source of DC current they do not remain in conduction for ever.

Figure 3H
Typical performance of the four main types of SPD
(ignoring the effects of inductance in their wiring or earth bonding)

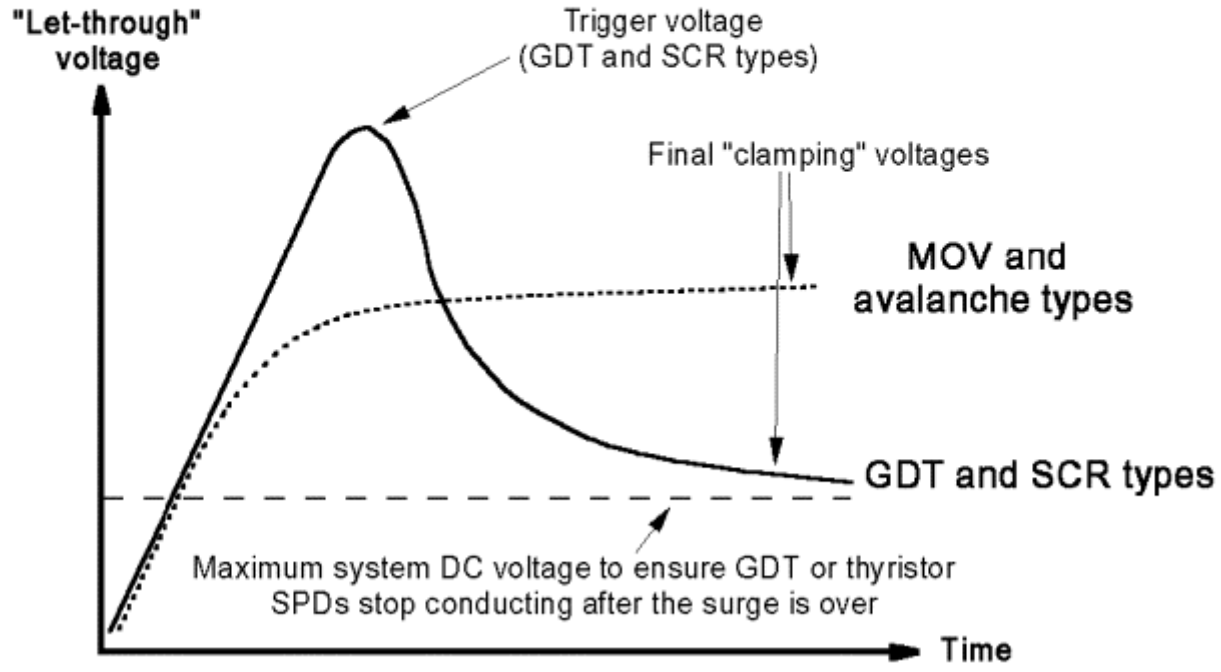

MOV and Avalanche devices act like zener diodes, with a knee voltage where they start to draw current. As their current increases their clamping voltage rises slowly.

The wired-in SPDs usually used for mains suppression may be unsuitable for use on signal or data lines, because of too much leakage, or too much capacitance, or impedance mismatch, or a variety of other problems. Special SPDs are available fitted with connectors intended to plug directly into a wide variety of digital and analogue signal cables and be mounted to earthed metalwork, like through-bulkhead filters. Radio antenna and high-speed data lines use matched transmission-line type SPDs, which use connectors such as BNCs.

3.14.2 Are SPDs needed on data lines?

SPDs are often advertised for use on data lines, but this is only really necessary on long cables which remain within a building, where there is a poor earth structure between the equipments at

each end such that lighting surges can cause a high voltage to appear between them. Where a product's cables exit a building (or connect to an external antenna) they should always have surge protection fitted.

For short interconnections (for example between keyboards or printers and PCs) the problems of earth voltage surge differences and fast transient bursts does not exist, and all that is left for an SPD to act on are electro-static discharge (ESD) transients. We shall see in Part 6 that ESD is best dealt with by using plastic to prevent discharges altogether, or else using good earth-bonding of metalwork.

3.14.3 SPDs and data integrity

Surge protection of analogue or digital data is usually not complete in itself. Even though the surges do not cause damage, they will have given a false value or bit. Where no memory or program is involved, as in simple analogue indicating instruments, a momentary glitch in the reading may be acceptable (depending on function), but for some analogue signals and all digital data (such as control signals) a momentarily incorrect signal can alter the stored data or operational mode, and this is usually unacceptable. Very slow data may be able to use filtering to reduce the "spike" to below detection thresholds.

Where glitched data is not acceptable and yet the only protection from surges is the use of SPDs, there needs to be some way of identifying and recovering from the incorrect data. Communication protocols are the usual answer, briefly mentioned Part 1.4.7. There are a number of these, ranging from simple to exotic, all with various overheads and penalties, and it is by far the best to buy the chips which implement protocols proven to be robust in real-life applications (e.g. CAN) rather than imagine that you can create a protocol of your own that will be good enough, even for the most cost-sensitive high-volume applications.

3.14.4 Ratings of SPDs

SPD ratings should really be chosen in conjunction with the design of a building and its lightning protection network. BS6651 Appendix C deals with this issue, and specifies the SPD ratings for equipment fitted in different parts of a building.

This standard, or others which deal with the lightning protection of electronic equipment (e.g. IEEE C62.41–1991) should be used where the generic or product EMC standards are lacking in surge requirements, or where their surge requirements are incomplete (for example, EN50082-1:1992 and EN 50082-2:1995 have no surge requirements at all, whereas EN 50082-1:1997 only applies surge requirements to mains and DC power ports).

BS6651 Appendix C specifies the following SPD ratings for the mains supply for equipment located more than 20 metres from the building's incoming mains connection (the most benign location, known as Category A):

- Low risk premises: 2kV, 167 Amps
- Medium risk premises: 4kV, 333 Amps
- High risk premises: 6kV, 500 Amps

Buildings not large enough to have a Category A, which includes most residences, are specified as Category B for internally connected equipment:

- Low risk premises: 2kV, 1,000 Amps
- Medium risk premises: 4kV, 2,000Amps
- High risk premises: 6kV, 3,000 Amps

External telecommunication and other signal/data cables (no matter how far they travel within a building), and the mains supplies to equipment mounted outside a building, are known as Category C:

- Low risk premises: 6kV, 3,000 Amps
- Medium risk premises: 10kV, 5,000Amps
- High risk premises: 20kV, 10,000 Amps

When a product is adequately protected against lightning surges, it is generally protected well enough against common surges generated by other means, such as switchgear. Where an application is known or suspected of suffering high levels of surges not of lightning origin, it is generally enough to protect the product to the next higher level of surge than is called for by the lightning protection standards.

Some superconducting magnet or power generation applications, or Nuclear Electromagnetic Pulse (NEMP) involve extreme or special types of surges, and are not addressed here.

3.14.5 Fusing of SPDs

All SPDs fail eventually, and since the majority of products use metal-oxide-varistor types (whose failure mode is to leak increasingly and finally to go short-circuit) in their mains inputs, they may need to be fused to prevent fire or shock hazards.

If the fuse is in the SPD circuit only, when it opens during the surge event that kills the SPD the protected equipment may be exposed to the remaining parts of the surge and damaged. Afterwards, even if the protected equipment is undamaged, it has lost its surge protection and so is very exposed to the next surge that comes along.

If the fuse is in series with the line that also goes to the protected equipment, the opening of the fuse due to SPD failure will disconnect the line to the equipment, which may not be acceptable in critical applications.

There is no easy answer to the problem of SPD fusing, but either of the above methods are generally acceptable providing the SPD is adequately rated for a goodly number of the maximum surges expected to be experienced. Once again, the lighting protection standards come to our aid, with risk assessments based on geography and application which allow the number and magnitude of surges caused by lightning to be assessed so that long-term reliability is likely.

3.14.6 Assembly of SPDs

SPDs, like filters, can be used to suppress unwanted DM or CM signals. In the jargon common to the surge community, DM surges are usually called line-to-line (or symmetrical), and CM surges become line-to-ground (or asymmetrical).

SPDs are always used in shunt mode, to “short out” the surge voltages. Also like filters, SPDs only function as expected when assembled correctly. One of the most important issues is to avoid lead inductance.

When surge currents flow in lead inductance they create voltages which increase the “let-through” voltage. As is shown in Figure 3J, it is best to track (or wire) the incoming power (or signal) directly to the terminals of its SPD, and then connect the protected circuitry to the SPD terminals too.

Figure 3J
Example of fitting SPDs at input of switch-mode power supply

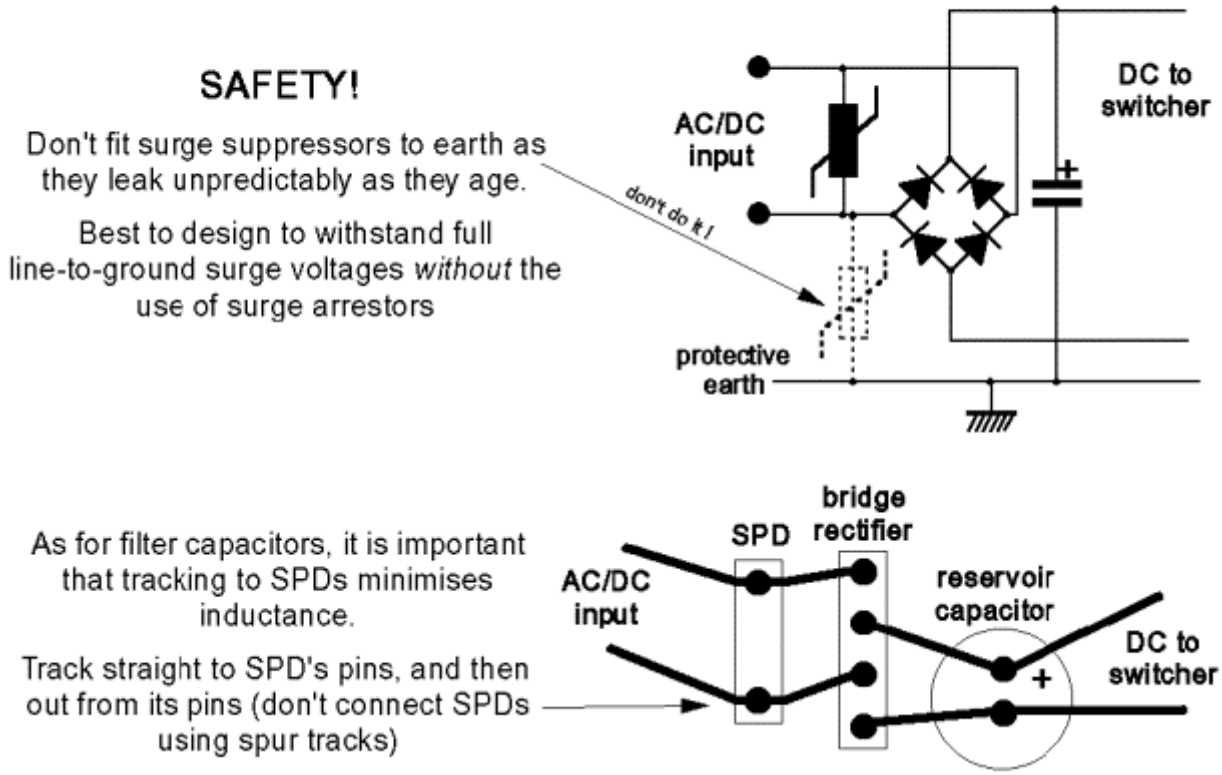

Figure 3J also shows that the use of SPDs from phase to earth is strongly discouraged, if not prohibited, for portable or pluggable equipment because of the unpredictable earth leakages of SPDs as they wear out, and the resulting safety problems. So it is best to use other means to deal with line-to-earth (common-mode) surges.

However, SPDs between supply phases and ground may be acceptable for permanently-wired equipment, especially if it has duplicated protective earth conductors, so that if one is faulty the other prevents electric shock from any SPD earth leakage. SPDs are commonly connected phase to ground in building installations, to help protect the equipment in the building from lightning surges (e.g. see BS6651 Appendix C), but in such instances the building’s earth structure (its common bonding network) should have a high degree of redundancy and not depend upon any single conductor for safety.

There are other types of surge limiters that are used in series with the circuit to be protected. Different types are used to prevent emissions of surge currents at switch-on, or to protect a circuit from current surges in its supply. These are not discussed here.

3.14.7 The problems of earth lift

Applying SPDs between phase and protective earth in a fixed product which is permanently wired may be acceptable under the relevant safety standards, given adequate protective earthing. But this can create new problems due to the inductance of the protective earth conductor. With ordinary protective conductor wire having an inductance of around 1μH/metre, and lightning surge currents of around 1kA with rise times of around 1μs – trying to suppress a power surge with an SPD connected to a length of earth wire does not actually suppress the beginning of the surge. Instead, it forces the voltage on the earth wire (and the chassis of the product) to follow the initial few microseconds of the surge voltage on the phase conductor. In the avionics world (at least) this phenomenon is called “earth lift”, which sums it up nicely.

This protects the mains input itself, but where there are signal cables connected to the product the surge voltage now present on the earthed chassis exposes the signal cables to the early part of the mains surge, which could damage their associated circuitry (even if opto-coupled, as most opto-couplers are only rated for 500V). So when SPDs are fitted to mains inputs to protect from line-to-ground surges, all signal cables may need surge protection too.

The same problem arises when an external telephone line or LAN is protected by SPDs connected to an equipment’s protective earth or chassis. This can allow surges on signal cables to damage the products’ power supply, necessitating the use of SPDs on the mains too (or else achieving the necessary protection by adequate line-to-ground voltage isolation).

Computer and telecommunication cabinets often deal with “earth-lift” surges by bonding their metal cabinets directly to the common-bonding network of their building with one or more heavy-duty cables each no more than 500mm long. Where they have “system blocks” of several cabinets passing signals between them, modern good practice is to construct a local earth mesh to reduce the inductance of the earth between the cabinets. This adds to the cost of the installation but removes the need for heavy-duty SPDs on all internal communication I/Os when SPDs are fitted on incoming mains and external LAN or telephone lines. Clearly, these earthing and bonding techniques are not appropriate for domestic or portable products.

More details on using SPDs within products to protect circuits may be found in Part 6. Although this is concerned with ESD, many of the circuit comments apply equally to using SPDs for surge protection.

3.15 Further reading

Other useful references on filters and surge protection are:

EMC for Product Designers 3rd *edition*, by Tim Williams, Newnes 2001, ISBN: 0-7506-4930-5, www.newnespress.com

EMC for Systems and Installations, by Tim Williams and Keith Armstrong, Newnes 2000, ISBN 0 7506 4167 3 www.newnespress.com, RS Components Part No. 377-6463.

EMC for Systems and Installations Parts 1 through 6, by Keith Armstrong, EMC+Compliance Journal, 2000. All UK EMC Journal and EMC+Compliance Journal articles are available electronically from the magazine archive at www.compliance-club.com.

Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone: 01457 871 605, Fax: 01457 820 145, Email: keith.armstrong@cherryclough.com

Design Techniques for EMC – Part 4
Shielding

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the fourth in a series of six articles on best-practice EMC techniques in electrical/electronic/mechanical hardware design. The series is intended for the designer of electronic products, from building block units such as power supplies, single-board computers, and "industrial components" such as motor drives, through to stand-alone or networked products such computers, audio/video/TV, instruments, etc.

The techniques covered in the six articles are:

1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
2) Cables and connectors
3) Filters and transient suppressors
4) **Shielding**
5) PCB layout (including transmission lines)
6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), so this magazine article format can only introduce the various issues and point to the most important best-practice techniques. Many of the techniques described in this series are also important for improving signal integrity: reducing the number of iterations during development and reducing manufacturing costs.

Table of contents for this part

4. Shielding

4.1 Shielding and the commercial imperative

A complete volumetric shield is often known as a "Faraday Cage", although this can give the impression that a cage full of holes (like Mr Faraday's original) is acceptable, which it generally isn't.

There is a cost hierarchy to shielding which makes it commercially very important to consider shielding early in the design process. Shields may be fitted around:

- Individual ICs — example cost 25p
- Segregated areas of PCB circuitry — example cost £1
- Whole PCBs — example cost £10
- Sub-assemblies and modules — example cost £15
- Complete products — example cost £100
- Assemblies (e.g. industrial control and instrumentation cubicles) — example cost £1,000
- Rooms — example cost £10,000
- Buildings — example cost £100,000

Please don't take these example costs as more than a very rough guide to orders of magnitude for domestic, commercial and industrial products. For example the shielding used in the new MI6 building in London will have cost very much more than £100,000, and that used on a flight-critical avionics computer will cost a lot more than £100. These costs do not take account of the re-engineering costs of adding shielding late in a project, which can be very much greater. They also don't take account of the lost sales and market position which can result from delays caused by shield re-engineering late in a project.

The important point is that shielding can be very low cost if it is designed-in carefully from the start, but can be extremely expensive indeed if it has to be applied at the last minute to make a product acceptable to a customer (or to an EMC enforcement agent).

Shielding always adds cost and weight, so it is always best to use the other techniques described in this series to improve EMC and reduce the need for shielding. Even where it is hoped to avoid shielding altogether it is best to allow for Murphy's Law and design from the start so that shielding can be added later if necessary. Mr Murphy can usually be discouraged from upsetting your product roll-out if you have a variety of shielding (and filtering) solutions ready to be dropped in when you first test your new product for EMC.

A degree of shielding can also be achieved by keeping all conductors and components very close to a solid metal sheet. Ground-planed PCBs populated entirely by low-profile surface mounted devices are therefore recommended for their EMC advantages. Even though such PCBs may require additional enclosure (volumetric) shielding, it should not need to have as high a shielding effectiveness (SE) and so will be easier to make and cost less.

A useful degree of shielding can be achieved in electronic assemblies by keeping their internal electronic units and cables very close to an earthed metal surface at all times, and bonding their earths directly to it instead of (or as well as) using a safety star earthing system based on green/yellow wires. This technique usually uses zinc-plated mounting plates or chassis, and can help avoid the need for high values of enclosure SE.

4.2 General concepts in shielding

Many textbooks have been written on the subject of how shields work, and it is not intended to repeat them here. However, a few broad concepts will help. A shield puts an impedance discontinuity in the path of a propagating radiated electromagnetic wave, reflecting it and/or absorbing it. This is conceptually very similar to the way in which filters work – they put an

impedance discontinuity in the path of an unwanted conducted signal. The greater the impedance ratio, the greater the SE.

At thicknesses of 0.5mm or over, most normal fabrication metals provide good SE above 1MHz and excellent SE above 100MHz. Problems with metal shields are mostly caused by thin materials, frequencies below 1MHz, and apertures, and this article focuses mainly on these.

4.3 Bigger and rectangular is better

It is generally best to allow a large distance between the circuits being shielded and the walls of their shield. The emitted fields outside the shield, and the fields that the devices are subjected to, will generally be more “diluted” the larger the shielded volume. This advice generally raises hollow laughter, since most product design seems to be about fitting a quart into a pint pot.

Where enclosures have parallel walls opposite each other standing waves can build up at resonant frequencies, and these can cause SE problems. Irregular shaped enclosures, or ones with curved or non-parallel walls (more hollow laughter) will help prevent resonances. Where opposing shield walls are parallel, try to prevent the resonances due to the width, height, or length from occurring at the same frequencies. So avoid cubic enclosures, use rectangular cross-sections instead of square, and try to avoid dimensions that are simple multiples of each other. E.g. if the length is 1.5 times the width, the second resonance of the width coincides with the third resonance of the length. Best to use irrationally ratio’d dimensions, such as that provided by the venerable Fibonacci series (which the Greeks knew as the Golden Mean). It is probably not worth worrying *too* much about this because the internal PCBs, components, and wiring alter its resonances unpredictably anyway.

4.4 Skin effect

Fields come in two flavours: electric (E) and magnetic (M). Electromagnetic fields consist of E and M fields in a given ratio (giving a wave impedance E/M of 377Ω in air). Electric fields are easily stopped by thin metal foils, since the mechanism for electric field shielding is one of charge re-distribution at a conductive boundary, so almost anything with a high conductivity (low resistance) will present a suitably low impedance. At high frequencies quite considerable displacement currents can result from the rapid rate of charge re-distribution, but even thin aluminium can cope with this quite nicely.

However, magnetic fields are much more difficult to stop. They need to generate eddy currents inside the shield material to create magnetic fields that oppose the impinging field. Thin aluminium is not going to be very suitable for this purpose, and the depth of current penetration required for a given SE depends on the frequency of the field, and on the characteristics of the metal used for the shield, and is known as the “skin effect”.

One skin depth is the depth in the shield material at which the “skin effect” causes the currents caused by the impinging magnetic field to be reduced by approximately 9dB. So a material which was as thick as 3 skin depths would have an approximately 27dB lower current on its opposite side and have an SE of approximately 27dB for that M field.

Skin effect is especially important at low frequencies, where the fields experienced are more likely to be predominantly magnetic with a lower wave impedance than 377Ω. The formula for skin depth is given in most textbooks, but requires knowledge of the shielding material’s conductivity and relative permeability. Figure 4A solves this for aluminium and steel, with copper thrown in for comparison. Pure zinc will have skin depths close to those of aluminium.

Figure 4A Graph of skin depths for copper, aluminium, and steel

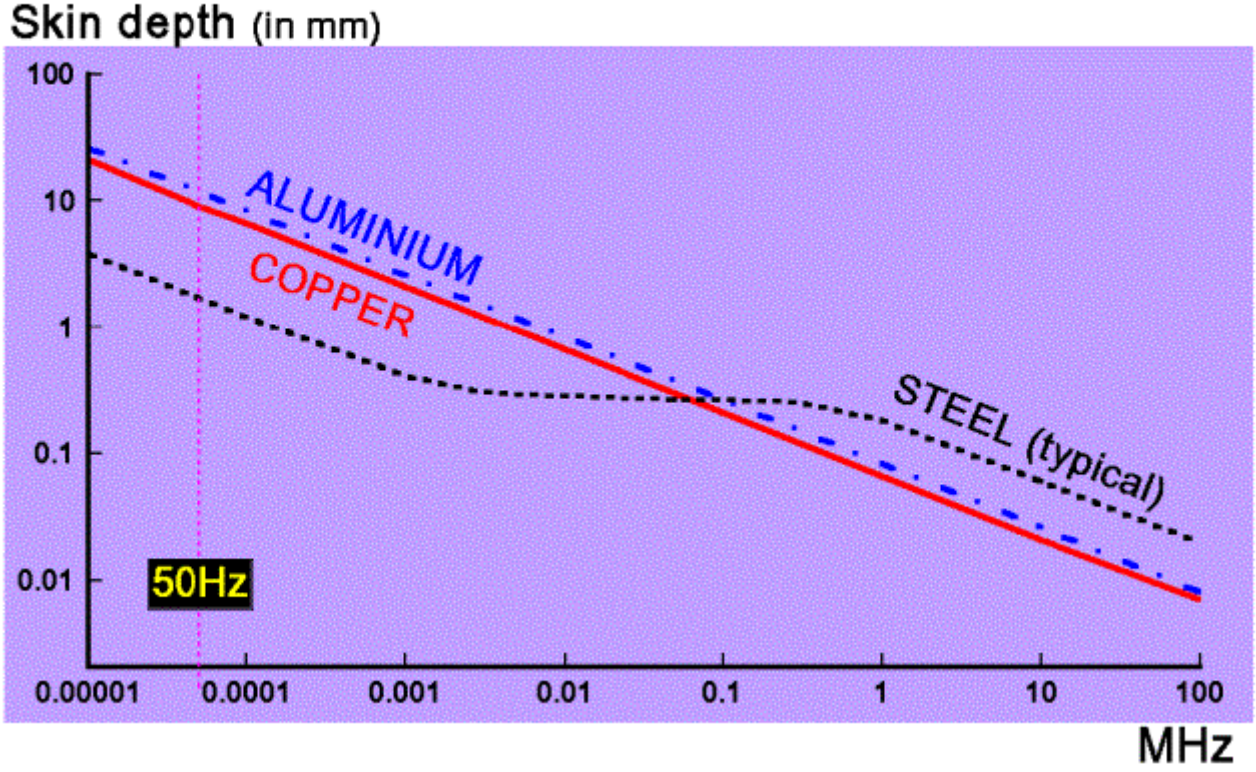

Copper and aluminium have over 5 times the conductivity of steel, so are very good at stopping electric fields, but have a relative permeability of 1 (the same as air). Typical mild steel has a relative permeability of around 300 at low frequencies, falling to 1 as frequencies increase above 100kHz, and its higher permeability gives it a reduced skin depth, making reasonable thicknesses of mild steel better than aluminium for shielding low frequencies. Different grades of steels (especially stainless) have different conductivities and permeabilities, and their skin depths will vary considerably as a result.

A good material for a shield will have high conductivity and high permeability, and sufficient thickness to achieve the required number of skin-depths at the lowest frequency of concern. 1mm thick mild steel plated with pure zinc (say, 10 microns or more) is just fine for many applications.

4.5 Apertures

It is easy to achieve SE figures of 100dB or more at frequencies above 30MHz with ordinary constructional metalwork. But this assumes a perfectly enclosing shield volume with no joints or gaps, which makes assembly of the product rather difficult unless you are prepared to seam-weld it all around and also have no external cables, antennae, or sensors (rather an unusual product).

In practice, whether shielding is being done to reduce emissions or improve immunity, most shield performance is limited by the apertures in it, as these two simple figures 4B and 4C try to show.

Figure 4B Shielding for immunity

Most shielding effectiveness is governed by apertures

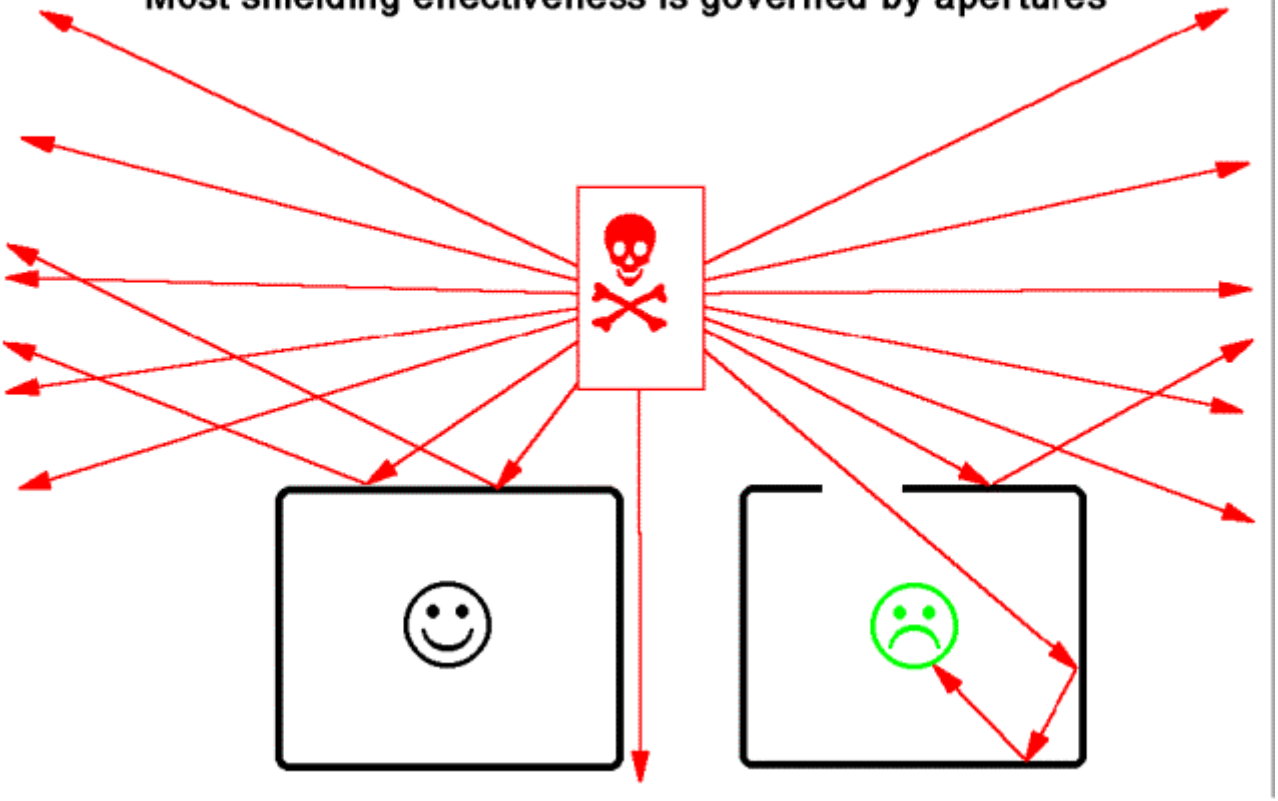

Figure 4C Shielding for emissions

Most shielding effectiveness is governed by apertures

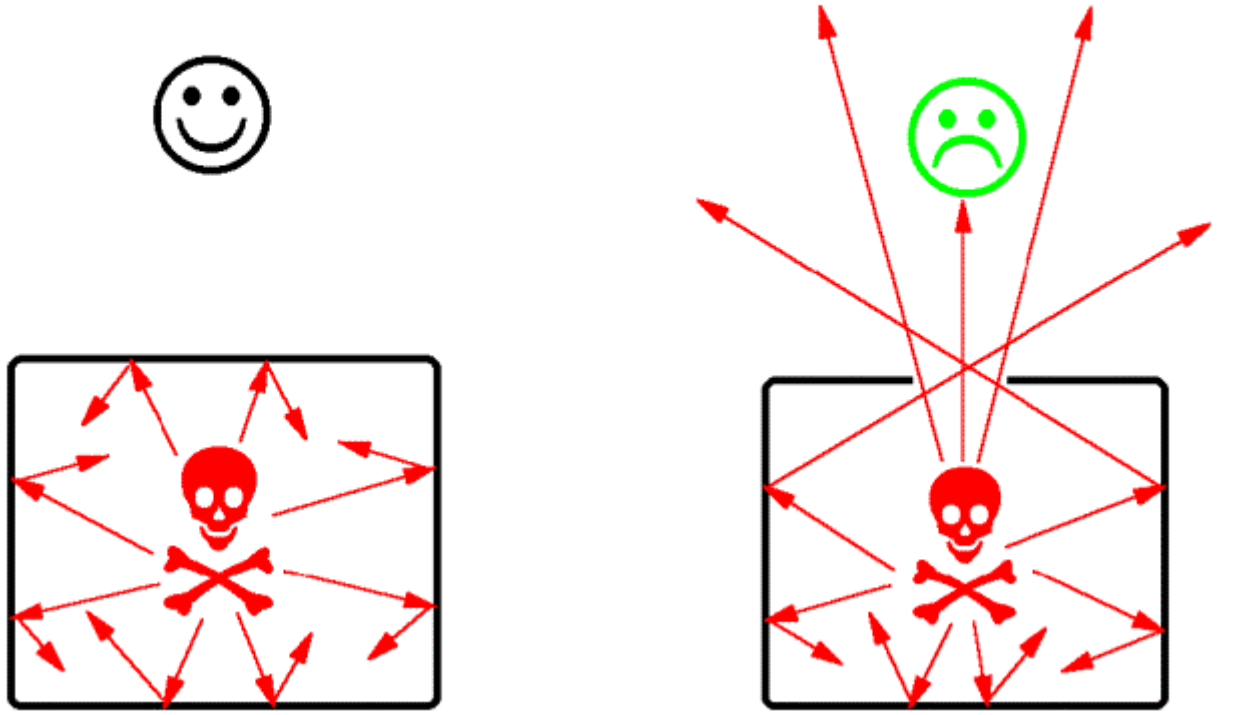

Considering apertures as holes in an otherwise perfect shield implies that they act as half-wave resonant “slot antennae”, and this allows us to make predictions about maximum aperture sizes for a given SE: for a single aperture, SE = 20log(λ/2d) where λ is the wavelength at the frequency of interest and d is the longest dimension of the aperture. In practice this assumption may not always be accurate, but it has the virtue of being an easy design tool that is better than doing nothing. Where its predications found to be are inaccurate it may be possible to refine it following practical experiences with the technologies and construction methods used on specific products.

The resonant frequency of a slot antenna is governed by its longest dimension – its diagonal. It makes little difference how wide or narrow an aperture is, or even whether there is a line-of-sight through the aperture.

Even apertures the thickness of a paint or oxide film, formed by overlapping metal sheets, still radiate (leak) at their resonant frequency as well as if they were wide enough to poke a finger through.

Figures 4D and 4E have been used before in this series to try to give a feel for the fact that the frequencies inside modern electronic products use the same range of frequencies as we rely on for communications and broadcasting.

Figure 4D The frequencies we use

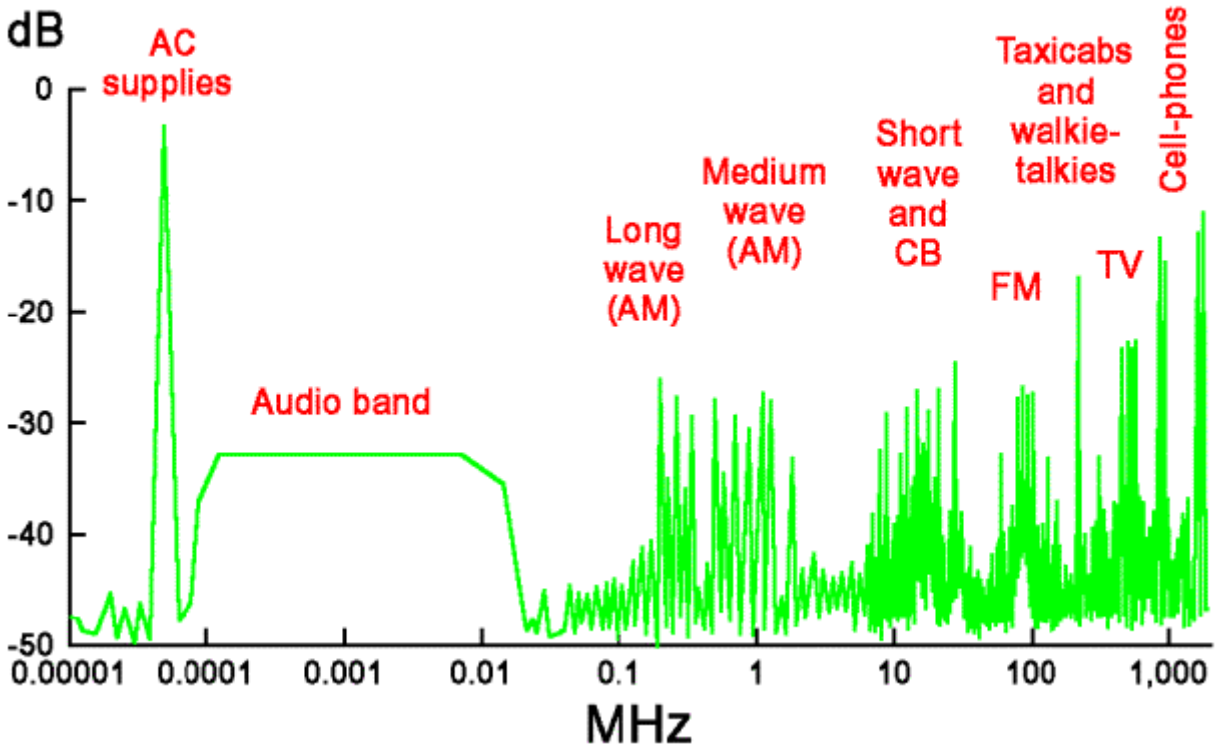

Figure 4E Plus the interference we create

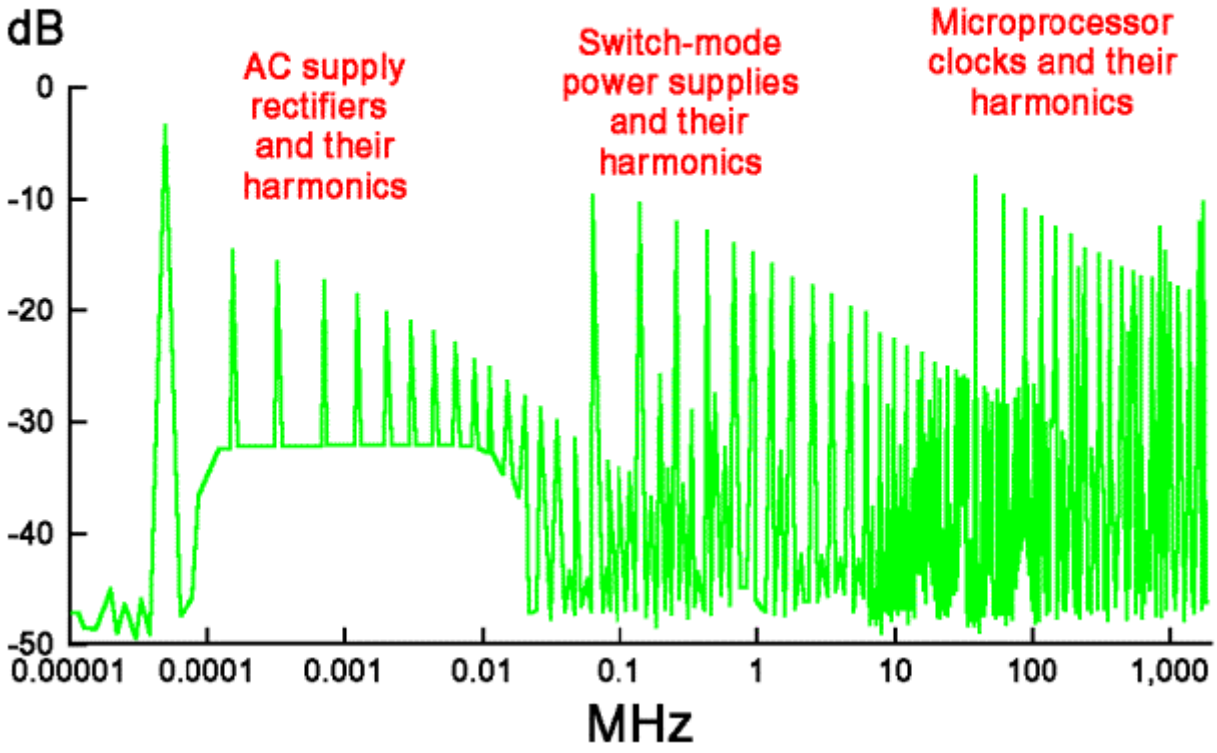

One of the most important EMC issues is keeping the products' internal frequencies inside, so they don't pollute the radio spectrum outside.

Figure 4F shows how effective apertures in shields can be at behaving like antennas and allowing the internal frequencies to disturb the radio spectrum and cause interference.

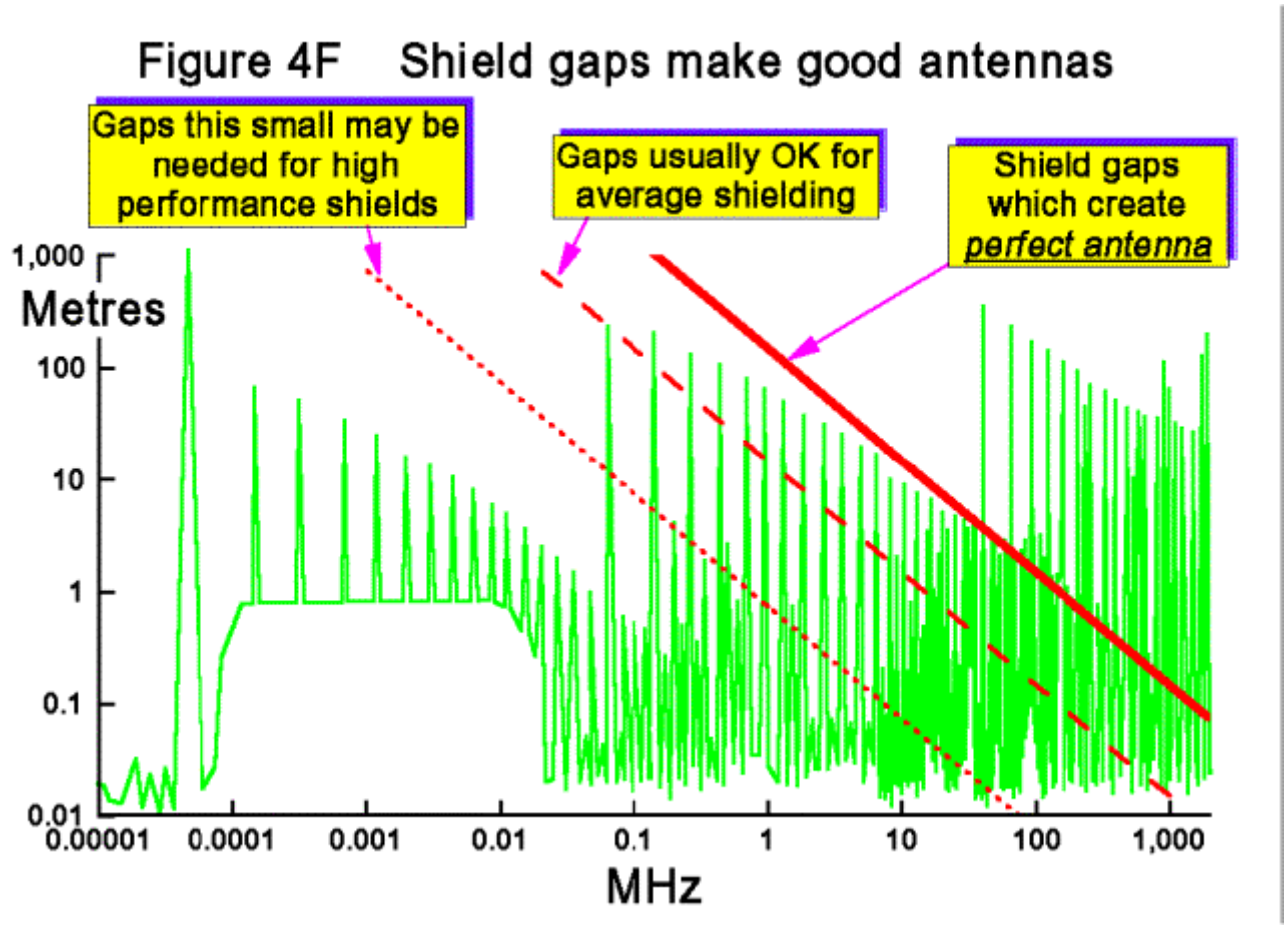

The half-wave resonance of slot antennae (expressed in the above rule of thumb: SE = 20log(λ/2d)) is the basis for the solid line in figure 4F (and for the rule-of-thumb of figure 4G) using the relationship $\nu = f\lambda$ (where ν is the speed of light: 3.10^8 metres/sec, f is the frequency in Hz, and λ is the wavelength in metres).

We find that a narrow 430mm long gap along the front edge of a 19-inch rack unit's front panel will be half-wave resonant at around 350MHz. At this frequency our example 19" front panel is no longer providing much shielding, and removing it entirely might not make much difference.

Figure 4G is useful when estimating the maximum size of an aperture for a given SE, and may be easily scaled to suit different dimensions.

Figure 4G Rule of thumb for the SE of apertures

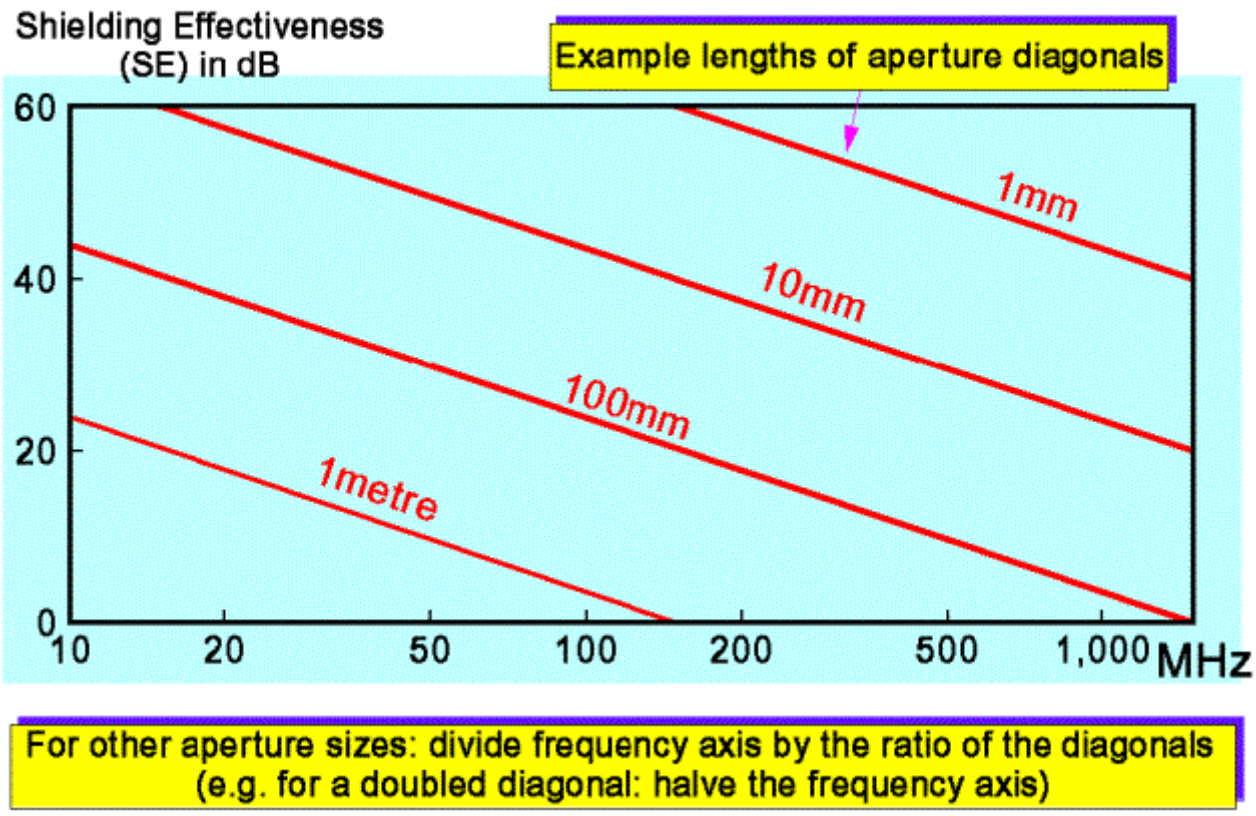

For an SE of 20dB at 1GHz (the present upper limit of testing in most standards) figure 4G suggests an aperture no larger than around 16 mm. For 40dB this would be only 1.6 mm, requiring gaskets to seal apertures and/or the use of the waveguide below cutoff techniques described later.

Actual SE in practice will depend on internal resonances between the walls of the enclosure itself, the proximity of components and conductors to apertures (keep noisy cables such as ribbon cables carrying digital busses well away from shield apertures and joints), the impedances of the fixings used to assemble the parts of the enclosure, etc.

Wherever possible, break all necessary or unavoidable apertures into a number of smaller ones. Unavoidably long apertures (covers, doors, etc) may need conductive gaskets or spring fingers (or other means of maintaining shield continuity). The SE of a number of small identical apertures nearby each other is (roughly) proportional to their number ($\Delta SE = 20\log n$, where n is the number of apertures), so two apertures will be worse by 6dB, four by 12dB, 8 by 18dB, and so on. But when the wavelength at the frequency of concern starts to become comparable with the overall size of the array of small apertures, or when apertures are not near to each other (compared with the wavelength), this crude '6dB per doubling' rule breaks down because of phase cancellation effects. However, at least this simple rule errs on the side of caution.

Apertures placed more than half a wavelength apart do not in general worsen the SEs that each achieves individually, but half a wavelength at 100MHz is 1.5 metres, so at such low frequencies on typical products smaller than this an increased number of apertures will tend to worsen the enclosure's SE.

Apertures don't merely behave as slot antennae. Currents flowing in a shield and forced to divert their path around an aperture will cause it to emit magnetic fields. Voltage differences across an aperture will cause the aperture to emit electric fields. The author has seen dramatic levels of

emissions at 130MHz from a hole no more than 4mm in diameter (intended for a click-in plastic mounting pillar) in a small PCB-mounted shield over a microcontroller. Figure 4G implies quite a good SE at that frequency from a 4mm hole, but the emissions from a particularly noisy microcontroller were causing significant currents to flow in the shielding can, and it appeared to be these that were causing the emissions, not the hole's antenna effect. Where long narrow apertures are concerned, it is sometimes possible to reduce such emissions by orienting the aperture's longer dimension appropriately with respect to the internal circuitry.

But be warned that the suggestions given are very very approximate and often confounded by proximity of cables, heavy loop currents flowing in the enclosure, etc., etc.

The only really sensible way to discover the SE of any complex enclosure with apertures is to model the structure, along with any PCBs and conductors (especially those that might be near any apertures) with a 3-dimensional field solver. Software packages that can do this now have more user-friendly interfaces and run on desktop PCs, alternatively you will be able to find a university or design consultancy that has the necessary software and the skills to drive it.

If you don't want to do field-solving, your best bet for any reasonable accuracy is to mock-up the construction as best you can, plus the internal electronics or whatever, and test its SE in an EMC test laboratory at the earliest stage in a project to avoid unpleasant surprises later on. Various methods exist for SE testing of enclosures, although few of them are standardised.

Since SE will vary strongly with the method and quality of assembly, materials, and internal PCBs and cables, it is always best to allow yourself an SE 'safety margin' of 20dB, or at least design-in features that will allow you to improve the SE by at least 20dB if you have problems with the final design's verification/qualification testing.

4.6 Low frequency (magnetic field) shielding

The frequency of 50Hz is highlighted on figure 4A, to show how difficult it is to achieve good SE at this frequency with any reasonable thickness of ordinary metals.

Special materials such as Mumetal and Radiometal have very high relative permeabilities, often in the region of 10,000. Their skin depth is correspondingly very small, but they are only effective up to a few tens of kHz. Care must be taken not to knock items made of these materials, as this ruins their permeability and they have to scrapped or else re-annealed in a hydrogen atmosphere. These exotic materials are used rather like channels to divert the magnetic fields away from the volume to be protected – a different concept to that used by ordinary shielding.

All metals shield materials with relative permeability greater than 1 can saturate in intense magnetic fields, and then don't work well as shields and often heat up. A steel or Mumetal shield box over a mains transformer to reduce its hum fields can saturate and fail to achieve the desired effect. Often, all that is necessary is to make the box larger so it does not experience such intense local fields.

Another shielding technique for low frequency shielding is active cancellation, and at least two companies have developed this technique specifically for stabilising the images of CRT VDUs in environments polluted by high levels of power frequency magnetic fields.

4.7 Waveguides below cutoff

Figure 4H shows that if we extend the distance that a wave leaking through an aperture has to travel between surrounding metal walls before it reaches freedom, we can achieve respectable SEs despite apertures large enough to put your fist through. This very powerful technique is called "waveguide below cutoff". Honeycomb metal constructions are really a number of waveguides below cutoff stacked side-by-side, and are often used as ventilation grilles for shielded rooms, and similar high-SE enclosures.

Figure 4H Waveguides can leak less

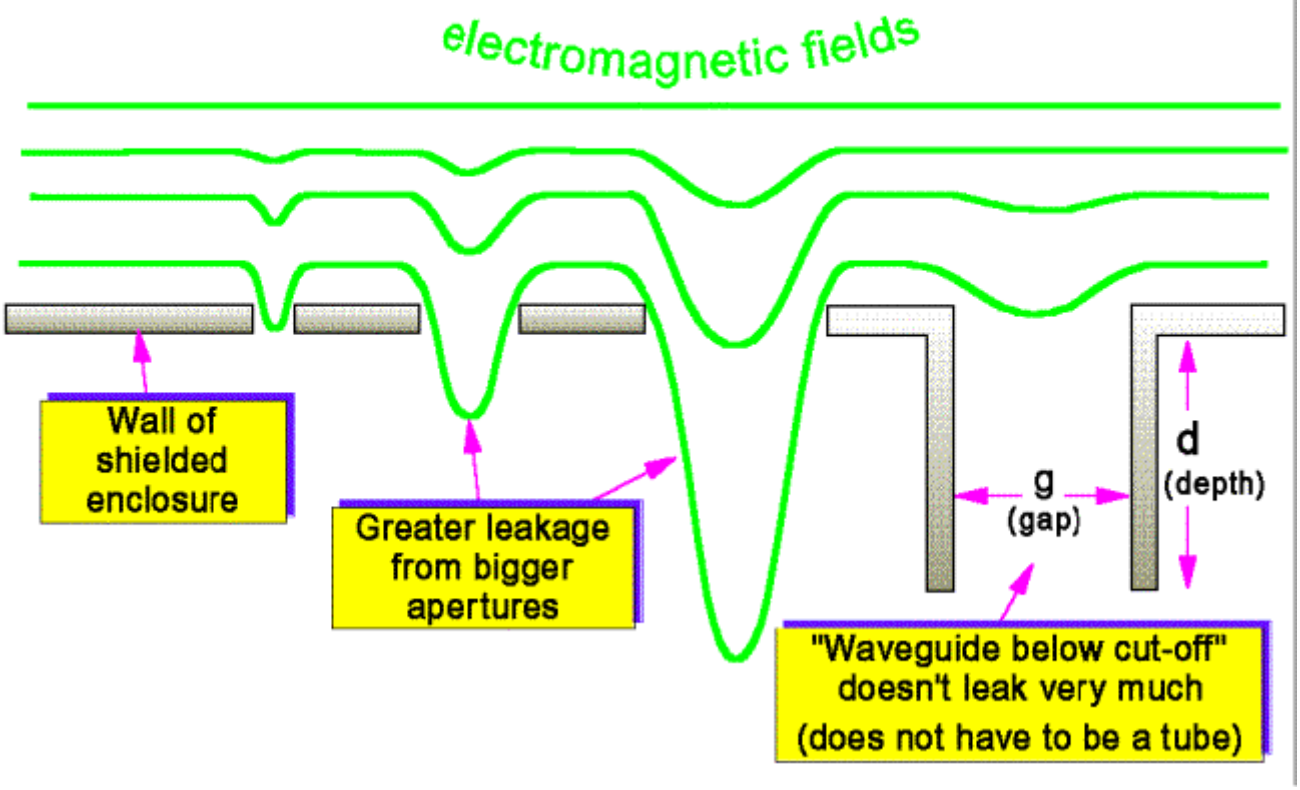

Like any aperture, a waveguide allows all its impinging fields to pass through when its internal diagonal (g) is half a wavelength, so the cutoff frequency of our waveguide is given by: f_{cutoff} = 150,000/g (answer in MHz when g is in mm). Below its cutoff frequency, a waveguide does not leak like an ordinary aperture (as shown by figure 4H) and can provide a great deal of shielding: for $f < 0.5 f_{cutoff}$ SE is approximately 27d/g where d is the distance through the waveguide the wave has to travel before it is free.

Figure 4J shows examples of the SE achieved by six different sizes of waveguide below cutoff. Smaller diameter (g) results in a higher cutoff frequency, with a 50mm (2 inch) diameter achieving full attenuation by 1GHz. Increased depth (d) results in increased SE, with very high values being readily achieved.

Figure 4J Some waveguides below cut-off

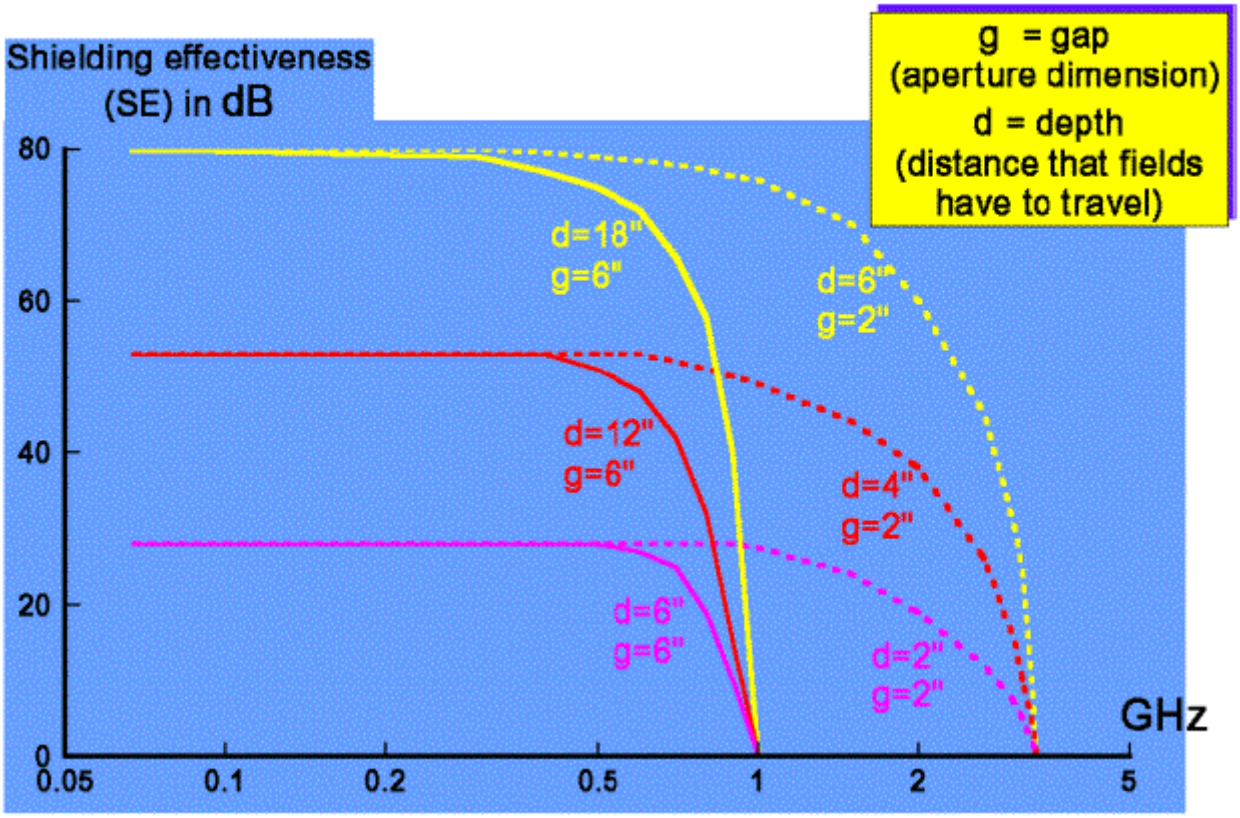

Waveguides below cutoff do not have to be made out of tubes, and can be realised using simple sheet metalwork which folds the depth (d) so as not to increase the size of the product by much. As a technique it is only limited by the imagination, but it must be taken into consideration early in a project as it is usually difficult to retrofit to a failing product not intended to use it.

Conductors should never be passed through waveguides below cutoff, as this compromises their effectiveness. Waveguides below cutoff can be usefully applied to plastic shafts (e.g. control knobs) so that they do not compromise SE where they exit an enclosure. The alternative is to use metal shafts with a circular conductive gasket and suffer the resulting friction and wear.

Waveguides below cutoff can avoid the need for continuous strips of gasket, and/or for multiple fixings, and thus save material costs and assembly times, but they appear to be rarely used: as a mechanical technique it is not of interest to electronic designers; and whoever saw a mechanical designer attend an EMC course?

4.8 Gasketting

Gaskets are used to prevent leaky apertures at joints, seams, doors and removable panels. For fit-and-forget assemblies gasket design is not too difficult, but doors, hatches, covers, and other removable panels create many problems for gaskets, as they must meet a number of conflicting mechanical and electrical requirements, not to mention chemical (to prevent corrosion). Shielding gaskets are sometimes required to be environmental seals too, adding to the compromise.

Figure 4K shows a typical gasket design for the door of an industrial cabinet, using a conductive rubber or silicone compound to provide an environmental seal as well as an EMC shield. Spring fingers are often used in such applications too.

Figure 4K Bonding the doors and removable panels of shielded cabinets
example of a door on industrial cabinet
(horizontal cross-section, viewed from above)

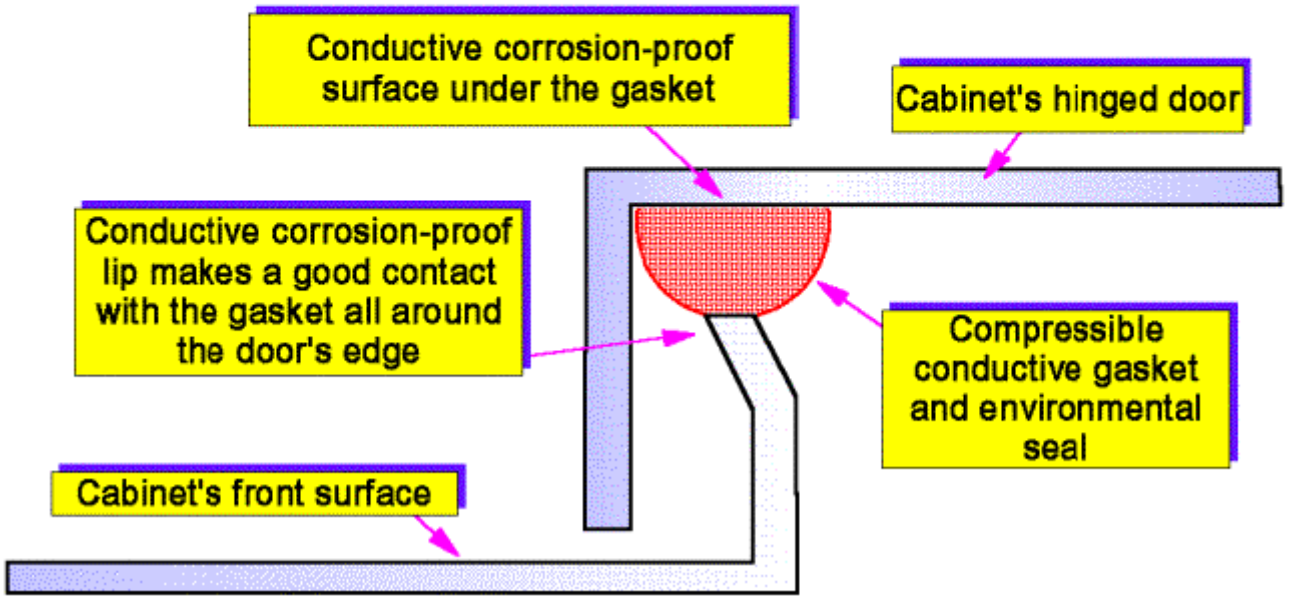

It is worth noting in passing that the green/yellow wire used for safety earthing of a door or panel has no benefits for EMC, above a few hundred kHz. This might be extended to a few MHz if a short wide earthing strap is used instead of a long wire.

A huge range of gasket types is available from a number of manufacturers, most of whom also offer customising services. This observation reveals that no one gasket is suitable for a wide range of applications. Considerations when designing or selecting gaskets include:

- Mechanical compliance
- Compression set
- Impedance over a wide range of frequencies
- Resistance to corrosion
  (low galvanic EMFs in relation to its mating materials, appropriate for the intended environment)
- Ability to withstand the expected rigours of normal use
- Shape and preparation of mounting surface
- Ease of assembly and dis-assembly
- Environmental sealing, smoke and fire requirements

There are four main types of shielding gaskets:

i. Conductive polymers (insulating polymers with metal particles in them). These double up as environmental seals, have low compression set but need significant contact pressure, making them difficult to use in manually-opened doors without lever assistance.

ii. Conductively wrapped polymers (polymer foam or tube with a conductive outer coating). These can be very soft and flexible, with low compression set and some only need low levels of contact pressure. However, they may not make the best environmental seals and their conductive layer may be vulnerable to wear.

iii. Metal meshes (random or knitted) are generally very stiff but match the impedance of metal enclosures better and so have better SEs than the above types. Poor environmental selaing performance, but some are now supplied bonded to an environmental seal, so that two types of gasket may be applied in one operation.

iv. Spring fingers ("finger stock"). Usually made of beryllium copper or stainless steel and can be very compliant. Their greatest use is on modules (and doors) which must be easy to manually extract (open), easy to insert (close), and which have high level of use. Their wiping contact action helps to achieve a good bond, and their impedance match to metal enclosures is good, but where they don't apply high pressures maintenance may be required (possibly a smear of petroleum jelly every few years). Spring fingers are also more vulnerable to accidental damage, such as getting caught in a coat sleeve and bending or snapping off. The dimensions of spring fingers and the gaps between them causes inductance, so for high frequencies or critical use a double row may be required, such as can be seen on the doors of most EMC test chambers.

Gaskets need appropriate mechanical provisions made on the product to be effective and easy to assemble. Gaskets simply stuck on to a surface and squashed between mating parts may not work as well as was hoped – the more their assembly screws are tightened in an effort to compress the gasket and make a good seal the more the gaps between the fixings can bow, opening up leaky gaps. This is because of inadequate stiffness in the mating parts, and it is difficult to make the mating parts rigid enough without a groove for the gasket to be squashed into, as shown by figure 4L. This groove also helps correctly position and retain the gasket during assembly.

Figure 4L Example of mechanical design to use a gasket

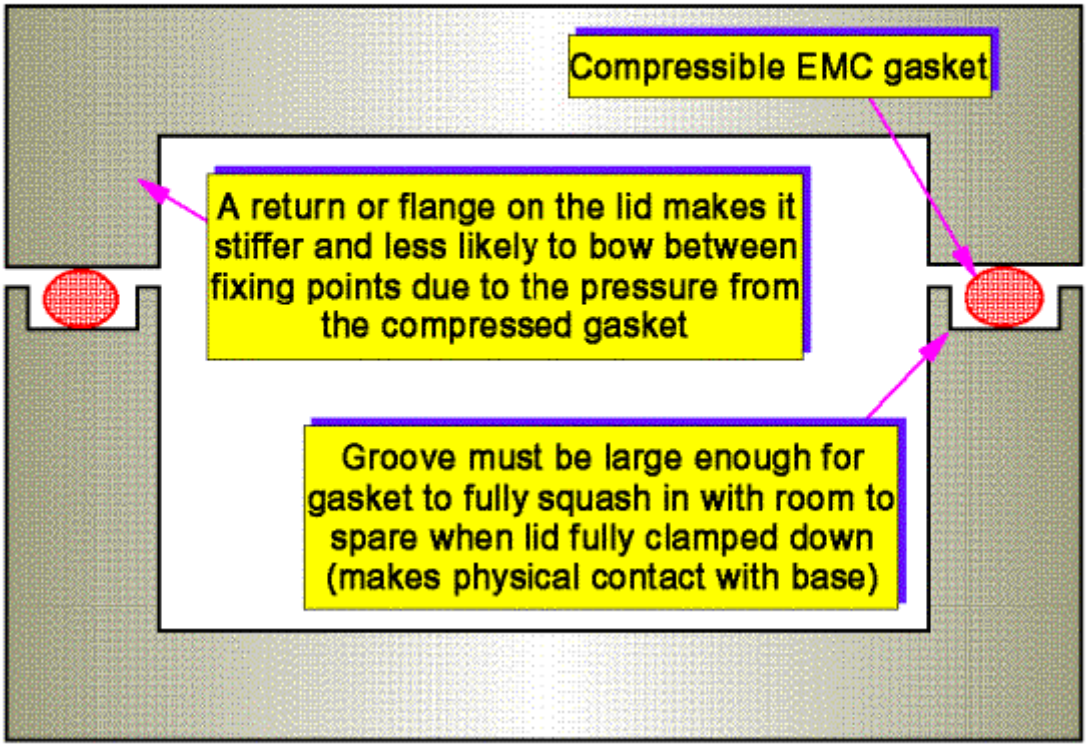

Gasket contact areas must not be painted (unless it is with conductive paint), and the materials used and their preparation and plating must be carefully considered from the point of view of galvanic corrosion.

All gasket details and measures must be shown on manufacturing drawings, and all proposed changes to them assessed for their impact on shielding and EMC. It is not uncommon, when painting work is transferred to a different supplier, for gaskets to be made useless because masking information was not put on the drawings. Changes in the painting processes used can also have a deleterious effect (as can different painting operatives) due to varying degrees of overspray into gasket mounting areas which are not masked off.

4.9 Shielding of displays (and the like)

Displays require apertures in enclosures, compromising shielding. A few small LEDs usually present few problems (although when using plastic enclosures they are often a weak spot for personnel ESD susceptibility). Mounting a display outside a shielded enclosure avoids the aperture, but removes the benefit of the shielding from the display and creates the new problem of what to do with the display's data and power cables and their penetration of the enclosure.

Figure 4M shows a display unit mounted in a large aperture in the wall of the shielded enclosure, using an internal "dirty box" to control the field leakage through the aperture. The joint between the dirty box and the inside of the enclosure wall must be treated the same as any other joint in the shield.

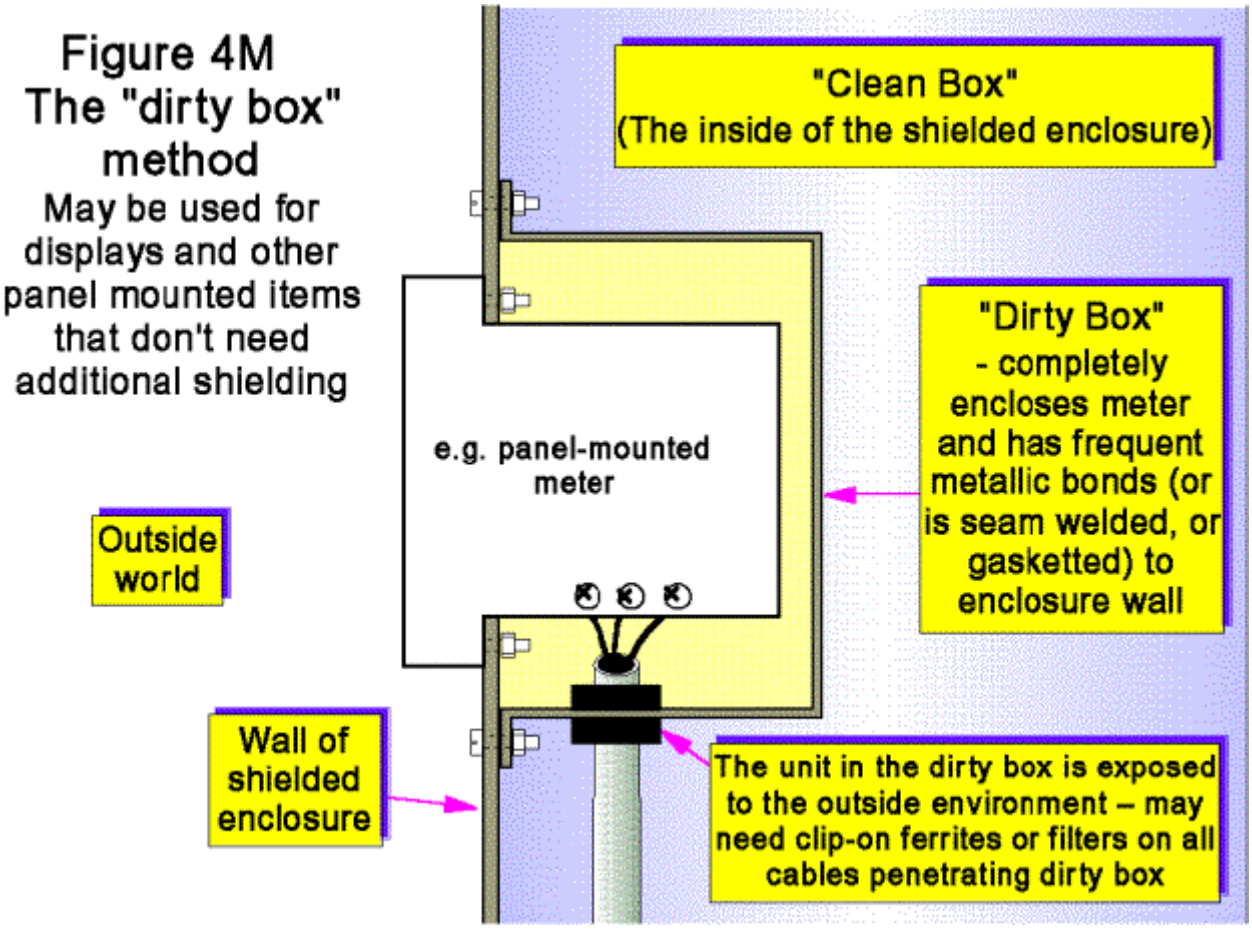

Shielded windows are needed where a display needs shielding by a product's enclosure. Some high-grade CRTs can provide a good shield when the metal frame around the front of their tube is electrically bonded to the front panel all around the aperture. Active matrix LCDs upgrades to products which had used high-grade CRTs have been known to be the cause of more emissions than the CRTs, and some have needed additional shielded windows where the CRTs had not.

A variety of shielded windows are available, based on two main technologies:

♦ Thin metal films on plastic sheets, usually indium-tin-oxide (ITO). At film thicknesses of 8 microns and above optical degradation starts to become unacceptable, and for battery-powered products the increased backlight power may prove too onerous. Referring to figure 4A, we see that the thickness of these films may be insufficient to provide good SEs below 100MHz.

♦ Embedded metal meshes, usually a fine mesh of blackened copper wires. For the same optical degradation as a metal film these provide much higher SEs, but they can suffer from Moiré fringing with the display pixels if the mesh is not sized correctly. One trick is to orient the mesh diagonally.

Honeycomb metal display screens are also available for the very highest shielding performance. These are large numbers of waveguides below cutoff, stacked side by side, and are mostly used in security or military applications where their extremely narrow viewing angle means that the operator's head prevents anyone else from sneaking a look at their displays.

A vital issue for screened windows is that their conducting layers (mesh, film, or honeycomb metal) must be bonded directly to the enclosure shield surface around all the edges of their cut -out.

Figure 4N shows one typical assembly method, which can use a conductive sealants / glues to avoid the need for mechanical fixings. The use of UV-curable conductive adhesives can make assembly times equal or better mechanical fixing methods.

Figure 4N Bonding shielded windows

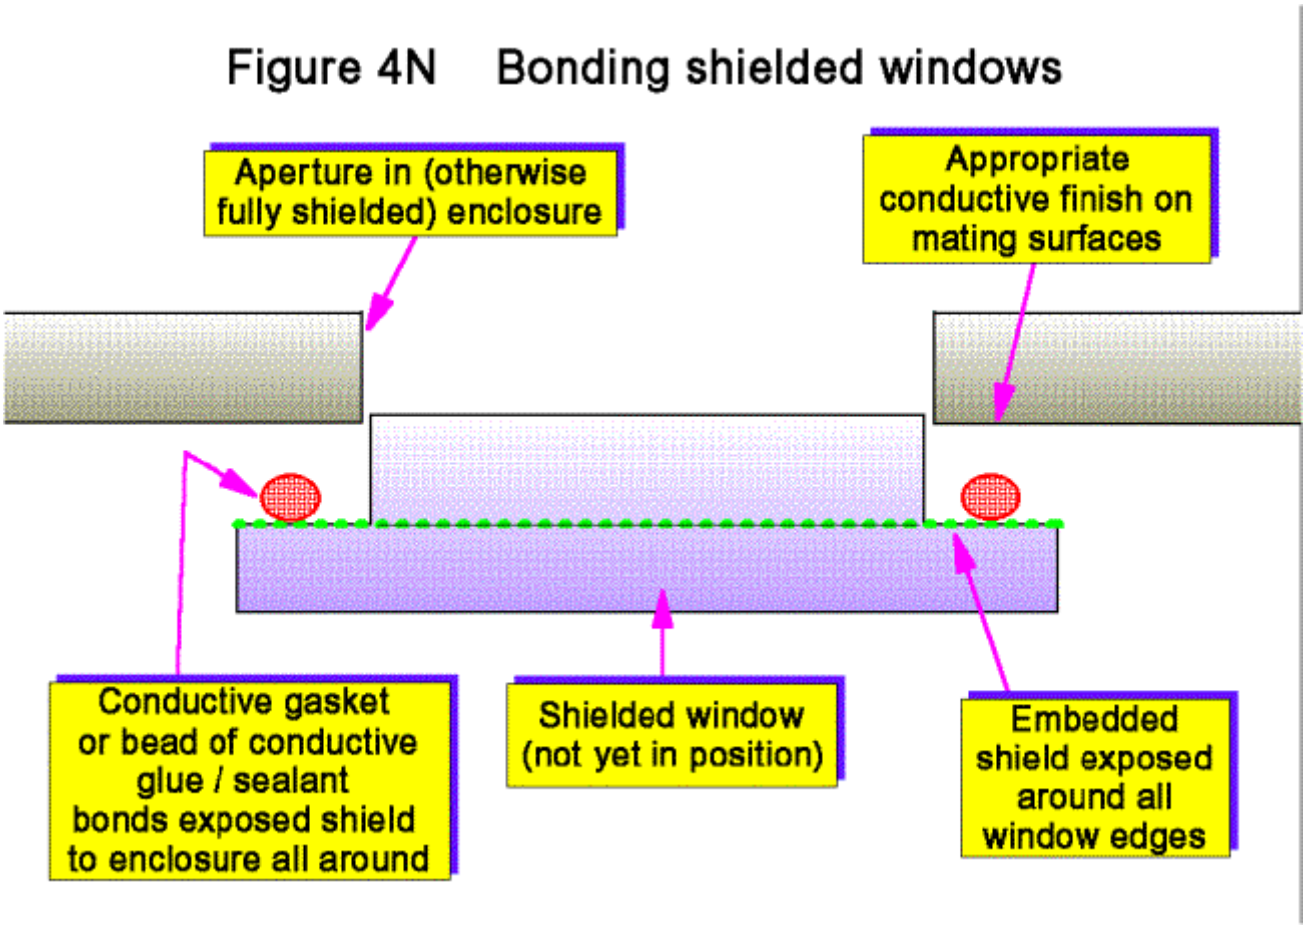

4.10 Shielding ventilation apertures

This presents similar issues to the shielding of displays, but only meshes or waveguides below cutoff can be used. As above, these must be bonded metal-to-metal (or with conductive gaskets) to the enclosure shield all around their ventilation aperture. Mesh size can of course be much coarser than for display applications, and expanded metal is sometimes used.

The mesh size must be small enough not to reduce the enclosure's SE too much. The SE of a number of small identical apertures near to each other is (roughly) proportional to their number, n, (ΔSE = 20logn), so two apertures will make SE worse by 6dB, four by 12dB, 8 by 18dB, and so on. For a large number of small apertures typical of a ventilation grille, mesh size will be considerably smaller than one aperture on its own would need to be for the same SE. At higher frequencies where the size of the ventilation aperture exceeds one-quarter of the wavelength, this crude "6dB per doubling" formula can lead to over-engineering, but no simple rule of thumb exists for this situation.

Waveguides below cutoff allow high air flow rates with high values of SE, and honeycomb metal ventilation shields (consisting of many long narrow hexagonal tubes bonded side-by-side) have been used for this purpose for many years. It is believed that at least one manufacturer of highly shielded 19" rack cabinets claims to use waveguide below cutoff shielding of the top and bottom ventilation apertures using ordinary sheet metalwork techniques.

The design of shielding for ventilation apertures can be complicated by the need to clean the shield of the dirt deposited on it from the air. Careful air filter design can allow ventilation shields to be welded or otherwise permanently fixed in place.

4.11 Shielding with painted or plated plastics

Plastic enclosures are often used for a pleasing feel and appearance, but can be difficult to shield. Coating the inside of the plastic enclosure with conductive materials such as metal particles in a binder (conductive paint), or with actual metal (plating), is technically demanding and requires attention to detail during design of the mould tooling if it is to stand any chance of working.

It is often found – when it is discovered that shielding is necessary – that the design of the plastic enclosure does not permit the required SE to be achieved by coating its inner surfaces. The weak points are usually the seams between the plastic parts: they often cannot ensure a leak-tight fit, and usually cannot easily be gasketted. Expensive new mould tools are often needed, with consequent delays to market introduction and to the start of income generation from the new product.

Whenever a plastic case is required for a new product, it is *financially vital* that consideration be given to achieving the necessary SE right from the start of the design process.

Paint or plating on plastic can never be very thick, so the number of skin-depths achieved can be quite small. Some clever coatings using nickel and other metals have been developed to take advantage of nickel's reasonably high permeability to reduce skin depth and achieve better SE.

Other practical problems with painting and plating include making them stick to the plastic substrate over the life of the product, in its intended environment. Not easy to do without expert knowledge of the materials and processes. Conductive paint or plating flaking off inside a product can do a lot more than compromise EMC – it can short conductors out, causing unreliable operation and risking fires and electrocution. Painting and plating plastics must be done by experts with long experience in that specialised field.

A special problem with painting or plating plastics is voltage isolation. For Class II products (double insulated) adding a conductive layer inside their plastic cases can reduce creepage and clearance distances and compromise electrical safety. Also, for any plastic-cased product, adding a conductive layer to the internal surface of the case can encourage personnel electrostatic discharge (ESD) to occur through seams and joints, possibly replacing a problem of radiated interference with one of susceptibility to ESD. For commercial reasons it is important that careful design of the plastic enclosure occurs from the beginning of the design process, if there is any possibility that shielding might eventually be required.

Some companies box clever (pun intended) by using thin and unattractive low-cost metal shields on printed circuit boards or around assemblies, making it unnecessary for their pretty plastic case to do double duty as a shield. This can save a great deal of cost and headache, but must be considered from the start of a project or else there will be no room available (or the wrong type of room) to fit such internal metalwork.

4.12 Shielding without metal

Volume-conductive plastics or resins generally use distributed conductive particles or threads in an insulating binder which provides the mechanical strength. Sometimes these suffer from forming a "skin" of the basic plastic or resin, making it difficult to achieve good RF bonds without helicoil inserts or similar means. These insulating skins make it difficult to prevent long apertures being created at joints, and also makes it difficult to provide good bonds to the bodies of connectors, glands, and filters. Problems with the consistency of mixing conductive particles and polymer can make enclosures weak in some areas, and lacking in shielding in others.

Materials based on carbon fibres (which are themselves conductive) and self-conductive polymers are starting to become available, but they do not have the high conductivity of metal and so do not give as good an SE for a given thickness.

4.13 Failing conducted tests due to inadequate shielding

Just because radiated phenomena are generally only tested above 30MHz, does not mean that shielding is unimportant below 30MHz. An enclosure that leaks excessively at low frequencies can cause a failure on a conducted test.

Small products with dimensions less than 0.5 metres usually make relatively inefficient antenna below 30MHz, and most problems are due to leaky cable screens (cables usually being long enough to be good antennae below 30 MHz). However, even a small product might need enclosure shielding that is effective at under 30MHz if it contains a powerful source of low-frequency fields.

4.14 Installation of shielded enclosures

A wire poked through an aperture in a shielded enclosure will completely destroy any SE pretensions. Figure 4P shows the main aspects of how to install an shielded enclosure without ruining it.

Figure 4P Installing shielded enclosures

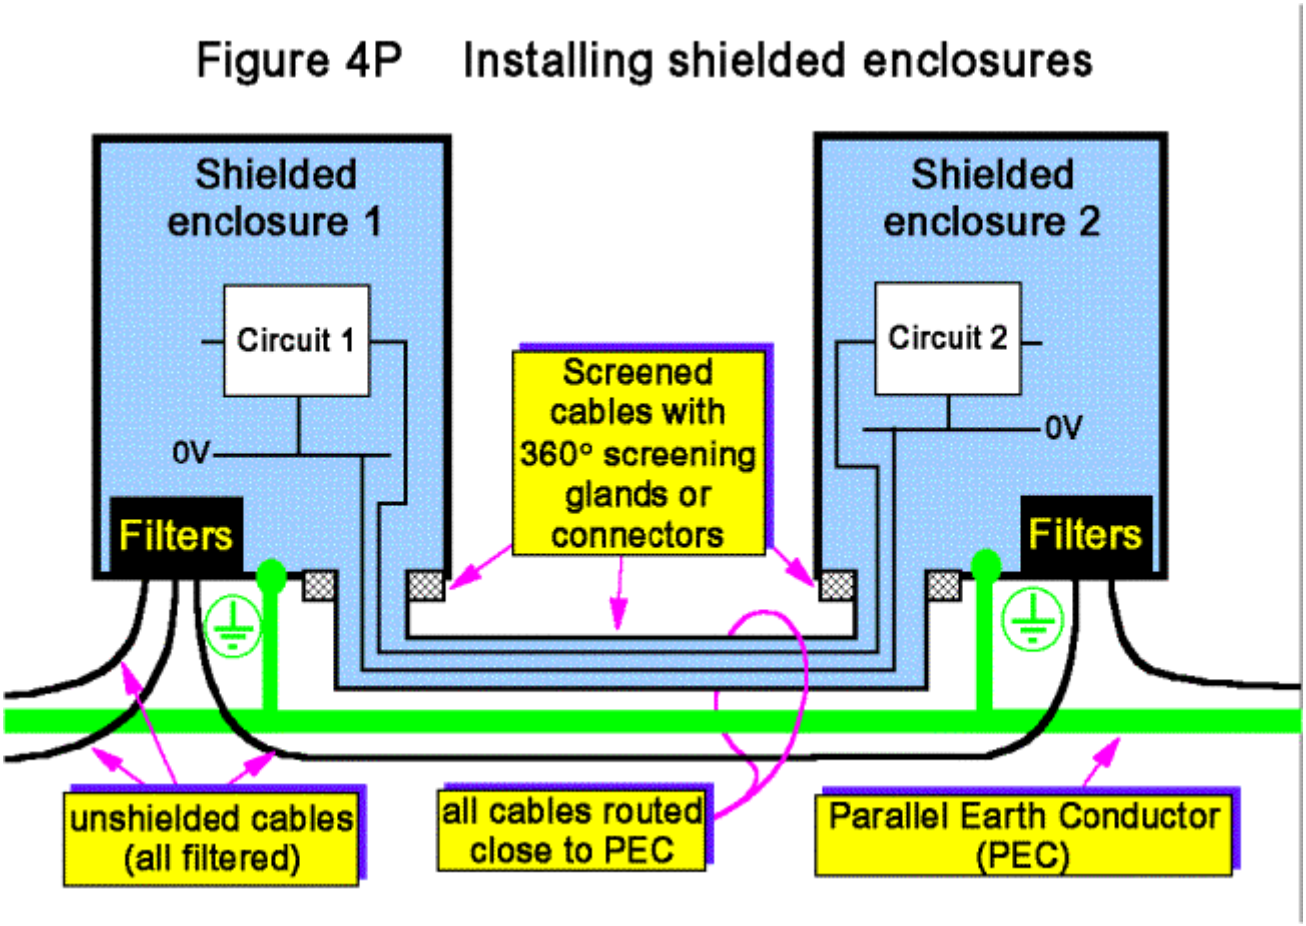

The screens and connectors (or glands) of all screened cables that penetrate a shielded enclosure, and their 360º bonding, are as vital a part of any "Faraday Cage" as the enclosure metalwork itself. The thoughtful assembly and installation of filters for unshielded external cables is also vital to achieve a good SE. These points were made in Parts 2 and 3 of this series, and they are worth making again.

In passing it is worth repeating that a cable screen pigtailed to an enclosure shield (instead of being 360º bonded) will ruin the SE of that enclosure from quite low frequencies upwards (say, above 10MHz for a short pigtail, lower if it is longer).

Refer to the draft IEC1000-5-6 (95/210789 DC from BSI) for best practices in industrial cabinet shielding (and filtering) and BS IEC 61000-5-2:1998 for best practices in cabling (and earthing) – including why pigtails are best consigned to the history books along with soldering irons that had to be heated in a fire. Figure 4P shows the main points of installing shielded cabinets according to these two best-practice standards, and is repeated from part 2 of this series. Refer to Part 2 for more details on this figure as regards the installation of shielded cabling, and Part 3 for details of installing filters on unshielded cables.

4.15 Using PCB-level shielding

Returning to our original theme of applying shielding at as low a level of assembly as possible to save costs, we should consider the issues of shielding at the level of the PCB.

The ideal PCB-level shield is a totally enclosing metal box with shielded connectors and feedthrough filters mounted in its walls, really just a miniature version of a product-level shielded enclosure as described above. The result is often called a module, can provide extremely high SEs, and is very often used in the RF and microwave worlds.

Lower cost PCB shields are possible, although their SE is not usually as good as a well-designed module. All depend upon a ground plane in a PCB being used to provide one side of the shield, so that a simple five-sided box can be assembled on the PCB like any other component. Soldering this five-sided box to the ground plane at a number of points around its circumference creates a "Faraday cage" around the desired area of circuitry. A variety of standard five-sided PCB-mounted shielding boxes are readily available, and companies who specialise in this kind of precision metalwork often make custom designs. Boxes are available with snap-on lids so that adjustments may easily be made, test points accessed, or chips replaced, with the lid off. Such removable lids are usually fitted with spring-fingers all around their circumference to achieve a good SE when they are snapped in place.

Weak points in this method of shielding are obviously the apertures created by the gaps between the ground-plane soldered connections, by any apertures in the ground plane (e.g. clearances around through-leads and via holes), and any other apertures in the five-sided box (e.g. ventilation, access to adjustable components, displays, etc.). Seam-soldering the edges of a five-sided box to a component-side ground plane can remove one set of apertures, at the cost of a time-consuming manual operation.

For the lowest cost, we want to bring all our signals and power into the shielded area of our PCB as tracks, avoiding wires and cables. This means we need to use the PCB equivalents of bulkhead-mounting shielded connectors, and bulkhead-mounting filters.

The PCB track equivalent of a shielded cable is a track run between two ground planes, often called a "stripline". Sometimes guard tracks are run on both sides of this "shielded track" on the same copper layer, these guard tracks having very frequent via holes bonding them to the top and bottom ground planes. The number of via holes per inch is the limiting factor here, as the gaps between them act as shield apertures (the guard tracks have too much inductance on their own to provide a good SE at high-frequencies). Since the dielectric constant of the PCB material is roughly four times that of air, when figures 4F and 4G are used to determine via spacing their frequency axes should be divided by two (the square root of the PCB's dielectric constant). Some designers don't bother with the guard tracks and just use via holes to "channel" the track in question. It may be a good idea to randomly vary the spacings of such rows of via holes around the desired spacing, to help avoid resonances.

Where striplines enter an area of circuitry enclosed by a shielded box, it is sufficient that their upper and lower ground planes (and any guard tracks) are bonded to the screening can's soldered joints on both sides, close to the stripline.

Track which only have a single ground plane layer in parallel, their other side being exposed to the air, are said to be of "microstrip" construction. When a microstrip enters a shielded PCB box it will suffer an impedance discontinuity due to the wall of the box. If the wavelength of the highest

frequency component of the signals in the microstrip is greater than 100 times the thickness of the box wall (or the width of box mounting flange), the discontinuity may be too brief to register. But where this is not the case some degradation in performance may occur, and such signals are best routed using striplines.

All unshielded tracks must be filtered as they enter a shielded PCB area. It is often possible to get valuable improvements using PCB shielding without such filtering, but this is difficult to predict so filtering should always be designed-in (at least on prototypes, only being removed from the PCB layout after successful EMC testing).

The best filters are feedthrough types, but to save cost we need to avoid wired types. Leaded PCB-mounting types are available which can be soldered to a PCB in the usual manner and then hand-soldered to the wall of the screening box when it is fitted at a later stage. Quicker assembly can be achieved by soldering the central contact of the filter to the underlying ground plane, making sure that solder joints between the shielding box and the same ground plane layer are close by, on both sides. This latter construction also suits surface-mounted "feedthrough" filters, further reducing assembly costs.

But feedthrough filters, even surface mounted types, are still more expensive than simple ferrite beads or capacitors. To allow the most cost-effective filters to be found during development EMC testing, whilst also minimising delay and avoiding PCB layout iterations, multipurpose pad patterns can easily be created to take any of the following filter configurations:

- zero-ohm link (no filtering, often used as the starting point when EMC testing a new design)
- a resistor or ferrite bead in series with the signal
- a capacitor to the ground plane
- common-mode chokes
- resistor/ferrite/capacitor combinations (tee, π, LC, etc. see Part 3 of this series for more details)
- feedthrough capacitor (i.e. centre-pin grounded, not truly feedthrough)
- feedthrough filter (tee, π, LC, etc., centre-pin grounded, not truly feedthrough)

Multipurpose padding also means we are not restricted to proprietary filters and can cook up our own to best suit the requirements of the circuit (and the product as a whole) at the lowest cost.

By now it should go without saying that all these PCB mounted filters should ideally be lined-up with their centres along the line of the wall of the shielding box, which will probably need a little cut-out in it to accommodate the components. Using surface-mounted devices rather than leaded allows the box cut-out size to be minimised, improving SE.

Where simple lines of filters are hard to achieve, take great care not to allow any unfiltered tracks to run close to any filtered tracks. Figure 4Q attempts to sketch what low-cost PCB shielding could look like.

Figure 4Q Sketch of shielding at PCB level

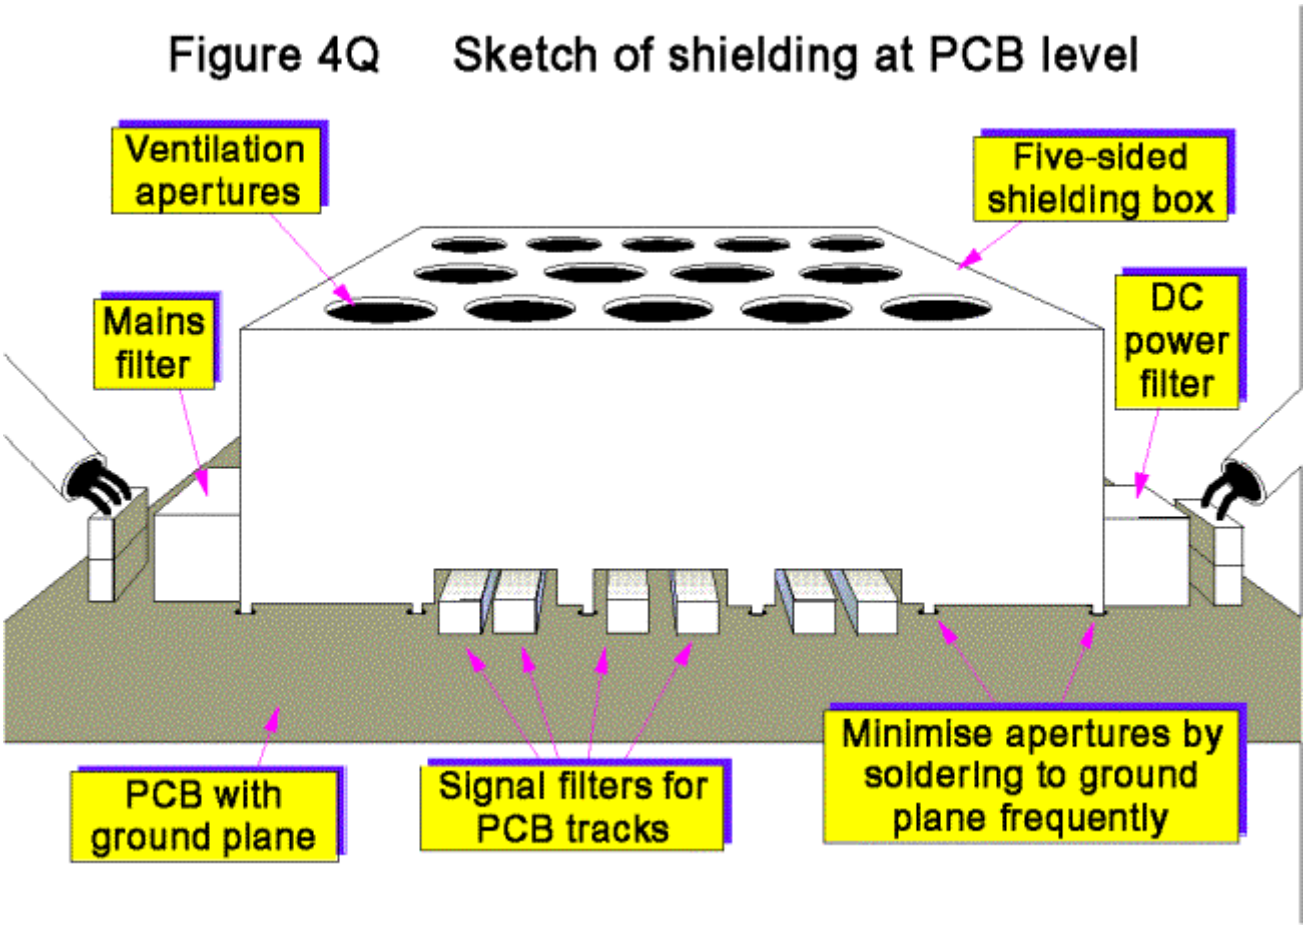

As with cables, it may be necessary to use shielding and filtering together, so it may be a wise precaution to provide for multi-padded filter layouts for all the “shielded” tracks entering a shielded PCB area, or at least make provision for a ferrite bead.

4.16 Further reading

EMC for Product Designers 3rd edition, by Tim Williams, Newnes 2001, ISBN: 0-7506-4930-5, www.newnespress.com

EMC for Systems and Installations, by Tim Williams and Keith Armstrong, Newnes 2000, ISBN 0 7506 4167 3 www.newnespress.com, RS Components Part No. 377-6463.

EMC for Systems and Installations Part 3 and Part 4, EMC+Compliance Journal, 2000. All UK EMC Journal and EMC+Compliance Journal articles are available electronically from the magazine’s archives at www.compliance-club.com

Eur Ing Keith Armstrong C.Eng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK, www.cherryclough.com
Phone: 01457 871 605, Fax: 01457 820 145, Email: keith.armstrong@cherryclough.com

Design Techniques for EMC & Signal Integrity – Part 5 PCB Design and Layout

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the fifth in a series of six articles on best-practice EMC and signal integrity techniques in electrical/electronic/mechanical hardware design. The series is intended for designers of electronic products, from building-block units such as power supplies, single-board computers, and 'industrial components' such as PLCs and motor drives, through to stand-alone or networked products such as computers, audio/video/TV, appliances, instrumentation and control, etc.

The techniques covered in these six articles are:

1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
2) Cables and connectors
3) Filters and transient suppressors
4) Shielding
5) **PCB layout (including transmission lines)**
6) ESD, electromechanical devices, and power factor correction

A textbook could be written about any one of the above topics (and many have), and this magazine article format merely introduces the various issues and points to the most important best-practice techniques. Signal integrity is treated as 'internal EMC'. Employing these well-proven techniques from the start of a new design generally reduces the number of iterations of hardware and software during development, and often reduces unit manufacturing costs too. EMC compliance is generally quicker, easier, with less risk of serious delays in time-to-market.

Table of contents for this part

5. PCB layout

These PCB-level design techniques are well proven to reduce the cost and effort of meeting “external” EMC requirements such as FCC, VCCI, and/or the EMC Directive. They also improve “internal EMC”, part of which is signal integrity, and help reduce the number of design iterations it takes to get a product to market. As electronic technology advances (clock speeds increase, A/D converter resolutions improve) internal EMC problems multiply, and the well-proven techniques described here become more important for commercial success.

The PCB techniques described here interact with each other to give improvements which are much greater than each can achieve on its own. They mostly improve the PCB's RF coupling mechanisms, and apply equally well to all types of analogue and digital circuits and to all the high-frequency emissions and immunity phenomena involved with both “internal” and “external” EMC.

Understanding ‘why’ these methods work helps extract their maximum benefits, but all we have room for here is a brief tour of these techniques – a few excellent references are provided at the end.

5.1 Circuit segregation

For cost-efficiency, this needs to be employed from the start of the real design process. The layout of the PCB should not begin until it is known where any shielding and filtering techniques need to be physically applied, so an overview of mechanical assembly and component placement should be done early in the product development lifecycle.

The following areas are first identified:

Outside-world: Total control of the electromagnetic (EM) environment is not practicable.

Inside-world: Where total control of the EM environment will be achieved.

5.1.1 The boundary between outside- and inside-worlds

This can be a hard boundary to draw. Conductors which run outside of a product's enclosure are clearly subject to the full outside-world EM environment, but cables which remain internal to a product may also suffer a subset of those phenomena if the product enclosure is not adequately shielded or external cables are not adequately filtered and/or suppressed. For example, a ribbon cable or jumper strip connecting two PCBs will not be protected from the outside-world’s high-frequency radiated RF environment unless there is to be an overall enclosure that provides adequate shielding over the whole frequency range of concern for both emissions and immunity.

The use of a single PCB for all the circuitry in a product is usually the most cost-effective way to meet EMC requirements. This is because it is easier to control the EM environment of a single PCB, with its obvious boundary between inside- and outside-worlds, than it is to control that of several PCBs and internal wires and cables. Many types of electronic products can avoid the need for a shielded enclosure if made using a single PCB (with no internal wires and cables) and the techniques described here. This can save costs in both materials and assembly, and allows a great deal more aesthetic freedom with plastic enclosure design.

5.1.2 Boundaries within an inside-world

When the inside-world circuitry has been determined, it should be further subdivided into dirty, high-speed, noisy, (etc.) potentially “aggressive” circuits, and clean, sensitive, quiet, (etc.) potential “victim” circuits. The likelihood of a circuit node being aggressive depends on its maximum dV/dt

and/or dI/dt. The likelihood of a circuit node being a victim of EM phenomena depends on its signal levels and noise margins (less = greater sensitivity).

5.1.3 Segregation

The various inside-world areas should be physically segregated from each other, and from the outside-world, both mechanically and electrically. Start at the earliest design phase by showing the segregated areas clearly on all drawings, usually done by drawing dotted lines around rectangular areas each covering one segregated portion of the circuit.

Ensure that this segregation is maintained throughout the rest of the design process including system design, PCB layout, wiring harness design, mechanical packaging, etc. Showing segregation clearly on all system, wiring, and circuit diagrams is of great help in communications between electronic designers, mechanical designers, and PCB layout persons – and is especially important where work is done by people on different sites, or by subcontractors.

Most design occurs in two dimensions. It is not uncommon to find that, in the final assembly, a PCB carrying a very sensitive circuit (such as a thermocouple or microphone amplifier) finds itself in close proximity to a noisy circuit (such as a switch-mode power converter), with consequent signal quality problems. Such unpleasant and time-consuming three-dimensional assembly problems should be avoided by detailed visualisation of the final assembly from the start, even before the circuits are designed and the PCBs laid out.

Figure 5A shows good segregation practices applied to a single-pcb product, whether it has an overall enclosure shield or not.

Figure 5A Example of circuit segregation for a single-PCB product

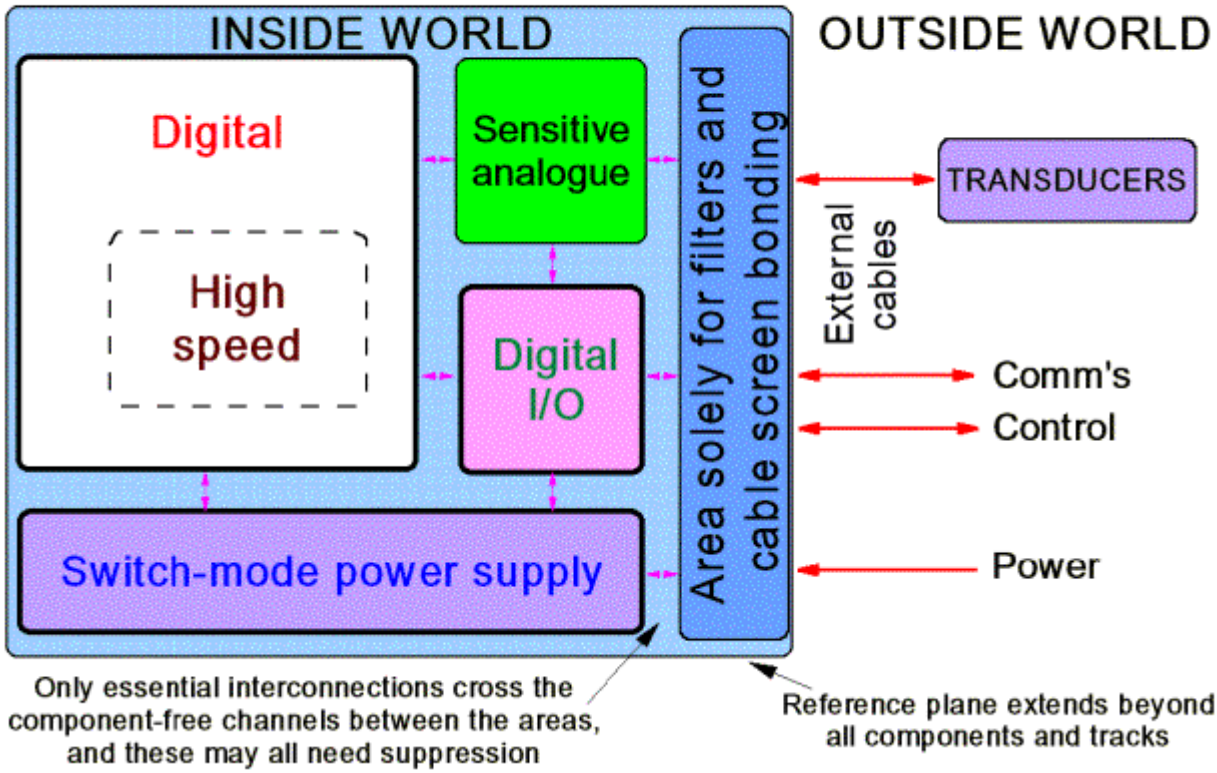

This example shows that the segregated area where the outside world interface suppression components are fitted, is along one edge of the PCB, as if it were a wall between outside and inside-worlds (which it is, in a way).

This area would only contain opto-isolators, isolating transformers, baluns, filters, transient absorbers, similar interfacing components but no ICs. It would also contain bonding points for the screens of any screened cables, and/or for any enclosure shielding. If this example PCB was part of a larger assembly, the segregation techniques employed for best EMC would be just the same.

The inside/outside-world interface components are restricted to one dedicated edge of the PCB to encourage all unwanted external currents (e.g. caused by voltage differences in protective earths) to restrict themselves to that area of the PCB, and discourage them from flowing through circuit areas.

Where an effective enclosure shield exists, the inside/outside-world boundary becomes the shielded wall of the enclosure. All of the associated filtering and suppression components, and cable screen bonding, must then use a connector panel set in the wall of the enclosure as their reference (as described in Part 4). A single area for all interconnections is still best. A wider range of PCB-mounted screened and/or filtered connectors that can also bond to a metal panel is now available. These parts would be soldered to the PCB reference plane, then electrically bonded metal-to-metal to the wall of a screened enclosure during final assembly, and can be very cost-effective.

Narrow channels free from components should be left between each of the segregated circuit areas on the PCB. These should be wide enough for the fitting of a PCB-mounted “tuner-can” shield, and provision should be made (at least on prototype boards) for bonding such screening cans to the 0V plane at frequent intervals (say, every 15mm) along all edges.

5.1.4 Component placement and routing of tracks

The most noisy or susceptible components in each area should be positioned first, as close to the centre of their areas and as far away from cables or wires as possible. Such components include clock generators and distribution (extremely noisy); bussed digital ICs (very noisy); microcontrollers (noisy); switch-mode power transistors and rectifiers and their chokes, transformers, and heatsinks (all very noisy), analogue ICs (sensitive), and millivolt level amplifiers (very sensitive). Remember (from Part 1) that even low-frequency operational amplifiers can be extremely susceptible to interference, even beyond 1GHz.

After the extremely short connections from components to reference planes, digital clock distributions (very aggressive signals) must be the next "nets" to be routed, and must be run on a single PCB layer adjacent to a 0V plane. These tracks must be as short as possible, and even so may need to use transmission-line techniques (described later). It may be necessary to experiment with component placement to achieve minimum track lengths. Where clock tracks are made longer than necessary to minimise skew, a "serpentine" layout is best.

Digital busses and high-speed I/O should be routed next, in a similar manner to clock tracks, deferring only to clock tracks and plane bonds where there is a conflict. Very susceptible tracks, such as those carrying millivolt transducer signals, should also be routed as if they were clock or data buss tracks, although they will always be in a different segregated area of the PCB. The later section on transmission lines describes what to do where critical tracks have to change layers.

All other types of analogue, digital, and power signals should also be routed according to how aggressive or sensitive they are. Where these characteristics are not obvious from a circuit analysis, probing a prototype with a wide-band oscilloscope (and/or spectrum analyser) with voltage or current probes will reveal which are the most aggressive, and injecting voltages or currents from a wideband sweep generator will reveal which are most sensitive. A loop probe can be most useful here, being able to inject signals into tracks without requiring connection of external equipment to potentially sensitive area of the circuit concerned.

All components and their tracks must be contained within their designated PCB areas. The only tracks to exit or enter an area are those that have to connect to other areas. If it has not proved possible to eliminate all the wires and cables inside a product, make sure that their routes are fixed so they can’t stray into the wrong PCB areas.

It is best to check that segregation instructions have been followed on *draft* PCB layouts, well before PCB manufacture. An easy check is to count the tracks and other conductors which cross the dotted lines showing the segregated areas on the circuit diagram – there should be exactly the same

number crossing the channels between areas on the draft PCB layout. Where PCBs have been autorouted it is usual to find additional tracks crossing area boundaries – these are often the source of much design heartache, so eliminate them right away by applying more skill to the track layout. Autorouting does not generally provide good layouts for EMC purposes.

5.2 Interface Suppression

EM disturbances can be radiated and/or conducted across interfaces between segregated areas, and shielding, filtering, or isolation techniques (such as opto-coupling) are used to reduce this to acceptable levels. To decide on the most cost-effective methods for each interface, they should be assessed for all the EM phenomena possible, given the operational EM environment and the emissions/immunity characteristics of the circuits concerned.

Don't ignore internal power supplies and other common connections such as 0Vs or grounds when considering interfaces between areas. Circuit designers abbreviate such connections on their circuit diagrams, often to invisibility, even though they provide the return current paths that are as important as the send path.

5.2.1 Suppressing outside/inside-world interfaces

Conductors passing from outside to inside-worlds may need the full range of suppression techniques – shielding, filters, isolating transformers, opto-isolators, surge protection devices, etc. As described above, best practice is to use a single PCB area or panel in the enclosure shield for all outside/inside-world interconnections and their suppression.

Visual displays (such as LCDs, LEDs, VDUs, moving-coil meters, etc.) and controls (such as pushbuttons, potentiometers, rotary knobs, etc.) are also interfaces between outside and inside-worlds, and are particularly exposed to personnel electro-static discharge (ESD), which will be covered by Part 6 of this series.

Shielding (See Part 4) may be applied to chips or areas of the PCB; the whole PCB; sub-assemblies of PCBs; entire assemblies of PCBs; or the entire product (listed in ascending order of cost and difficulty). The segregation methods described above help make low-cost shielding possible.

5.2.2 Interfaces between dirty/high speed/noisy and clean/sensitive/quiet areas

Determining the types/amounts of suppression to be applied to tracks and other conductors interconnecting different PCB areas needs an assessment of both the desired signals and the unwanted noise they may carry, plus the sensitivity of the circuits they connect to.

Digital clocks and data busses are aggressively noisy and should not be allowed in clean/sensitive/quiet PCB areas. Data intended for a sensitive area should be latched from its bus no closer than the boundary of that area, and the data busses themselves restricted to a noisier area.

Power distribution networks are often overlooked routes for conducted noise from one segregated area to another, as are "static" data lines, and other low-frequency signals. Digital control lines which remaining at logic 1 or 0 for long periods are often thought to be quiet, but they usually carry tens or even hundreds of millivolts of high-frequency noise generated by the electrical activity of their source ICs (e.g. by "ground bounce" and its corresponding "power bounce"). Many an analogue circuit has suffered from noise on its power supply rails from switch-mode power supplies or DC/DC converters, or from digital processing sharing the same rails or from noise injected into analogue switches and opamps from "static" logic control signals. It is often necessary to fit small filters to such inter-area connections, but sometimes more drastic measures are required, such as opto-isolation.

Components that interface between segregated areas, such as analogue-to-digital convertors, transformers, data bus latches, filters, isolators, and the like, should be positioned at an edge common to the areas they interconnect. They should usually remain wholly within one area or the other (so as to keep a component-free channel) and their tracks must route directly to their respective areas and not mingle with tracks associated with the "other side" of these components or

other areas. The purpose of keeping the channel component-free is to make it easier to fit shielding over the segregated areas of circuitry, should it be needed. Where interface components like ferrite beads, common-mode chokes, or opto-isolators are placed *in* one of these channels it can help to achieve good separation between the tracks associated with each circuit area, but the cut-outs they require in any PCB-mounted shield may compromise its shielding effectiveness.

This compromise between the need for good track segregation and for shielding effectiveness does not apply when 'feedthrough' filter components are fitted as interconnections between segregated areas. These are designed to fit into and actually penetrate the walls of screened enclosures, so when fitted in (otherwise component- and track-free channels) they encourage good track segregation and don't compromise shielding effectiveness. Traditionally, feedthrough filters are screwed or soldered into a hole in the shield, with wires connecting to their ends. This does not suit robotic surface-mounted assembly techniques, prompting some manufacturers to produce 'SMD feedthrough filters'. These generally have an earth electrode around their centre, which is intended to be soldered to the PCB reference plane (although some types may also be able to be hand-soldered to a cut-out in their shield). In general the small size and low-profile of these parts means that they only require a very small cut-out in the shield they penetrate, and so may be expected to have little effect on shielding effectiveness. Where SMD 'feedthroughs' are used, their performance will be improved if the shield they are associated with is soldered to the PCB reference plane as close as possible to the SMD feedthroughs, as frequently as can be achieved. Very stringent applications sometimes require PCB shields to be seam-soldered all around their circumference, and such assemblies would probably need to use the more traditional wired-in feedthrough devices.

Radiated interference between segregated areas is possible. The stray capacitance between components may only be a fraction of a pF, but at high frequencies can inject significant displacement currents into components and tracks in neighbouring areas. Combining small-sized low-profile components with PCB reference planes, and placing the noisiest devices (e.g. clocks, processors, switch-mode power devices) and signals in the centres of their areas, can help avoid the need to shield PCB areas from each other.

5.2.3 Details of interface suppression techniques

Suppression techniques include:

- common-mode and/or differential mode filtering
- galvanic isolation using opto-isolators or transformers
- communications protocols (to improve bit error rate in the presence of interference)
- surge protection devices
- the use of balanced drive and receive signals (instead of "single ended")
- the use of fibre-optic, infra-red, wireless, laser, or microwave instead of copper cables
- shielding of areas, volumes, cables, and connectors

All of these are covered by other parts of this series. It is important to realise that on a PCB only a plane (described next) can provide a good enough reference at high frequencies to enable the full performance of filters, cable screens, and internal shields to be achieved on a PCB.

5.3 Reference planes

Due to their intrinsic reactance and resonances, tracks, wires, "star grounding", area fills, guard rings, etc., cannot provide an adequate reference for a PCB except at low frequencies (usually below 1MHz). For example, the rule-of-thumb for the inductance of PCB tracks on their own or single wires, is 1nH/mm. This means that just 10mm of PCB track has an impedance of 6.3 Ω at 100MHz, and 63Ω at 1GHz. For this reason, only unbroken areas of metal conductor can provide an adequate reference up to 1GHz (and beyond), and these are called reference planes. In a PCB these are usually called power, ground, or 0V planes, but it is best to avoid the use of the words "ground" or "earth" in connection with EMC and circuits (reserving them for specific uses associated with safety bonding). As far as most EMC design techniques are concerned, a connection to the green/yellow protective earth conductor can often be more of a problem than a solution.

Reference plane techniques allow dramatic reductions in all unwanted EM coupling when used in conjunction with the other techniques described here. Reference planes are also essential for almost every other PCB EMC design technique to function properly.

5.3.1 Creating proper reference planes

A high-quality high-frequency reference must have a vanishingly small partial inductance, and can be created on a PCB by devoting one layer to an unbroken copper sheet, called a reference plane. A 0V reference plane would be used as the 0V (or “ground”) connection for all its associated circuits, so that all 0V return currents flow in the plane and not in tracks. Power planes are created and used in a similar manner for power connections and their return currents.

0V reference planes must lie under all their components and all their associated tracks, and extend a significant distance way beyond them. The segregation and interface suppression techniques described above must still be followed even where a common 0V plane is used for a number of circuit areas.

Perforations such as leads, pins, and via holes increase the inductance of a plane, making it less effective at higher frequencies. “Buried via” techniques have been developed for cellphones, allowing interconnections between tracking layers without perforating the reference plane. For less demanding products a rule-of-thumb is that any gaps must have dimensions of 0.01λ or less at the maximum frequency concerned. For a good plane at 1GHz, (e.g. to help meet most of the present EU harmonised EMC standards cost-effectively) this rule implies that plane gaps should have dimensions $\leq$ 1.5mm (remembering that the velocity of propagation in FR4 is approx. half of what it is in air). "Sneaking" tracks into a plane layer is not allowed.

Unavoidable gaps in a plane must not merge to create larger ones. PCB design rules should size clearance holes so that for regular hole spacings such as DIL packages, the plane "webs" between holes as shown by Figure 5B.

Figure 5B Example of a webbed 0V plane under a leaded IC

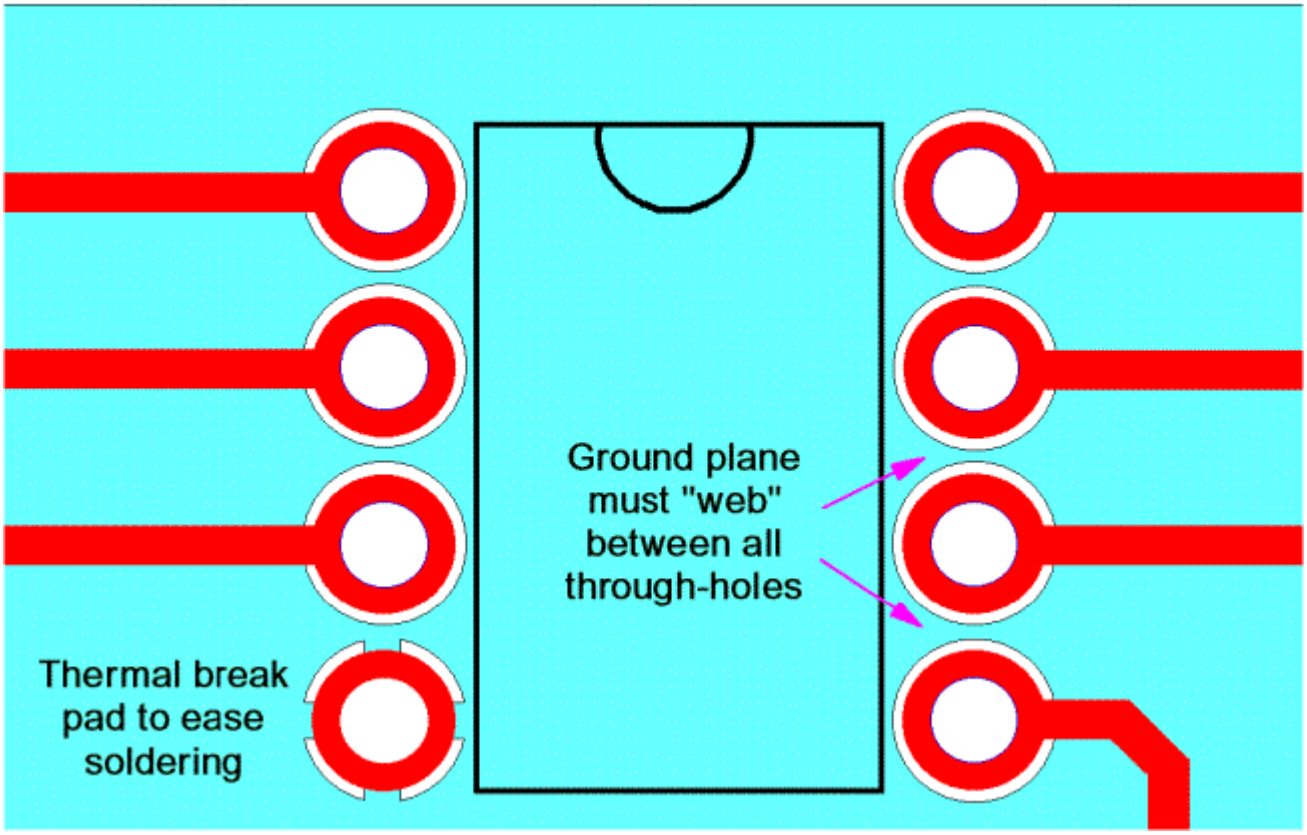

Tracks, area fills, guard rings, etc. forming part of the reference on signal layers *can* be used to good effect at high frequencies – but *only* when bonded to an underlying 0V plane with at least one via hole every 5 to15mm (using a random allocation of spacings).

0V planes should extend well beyond all components, tracks and power planes. [1] recommends "the 20H rule": 0V planes should extend by at least 20 times their layer spacing. High-speed components (such as digital clocks, processors, and memory) and their signal tracks should always be placed near the centres of their segregated areas, well away from plane edges.

All 0V and power connections must bond immediately to their respective planes to minimise their connection inductance. Leaded components must have their through-plated holes directly connected to planes using thermal-break pads as shown by figure 5B (sometimes called wagon-wheels) to help with soldering. Surface mounted devices (SMDs) for reflow soldering have to compromise the prevention of dry joints or "tomb-stoning" with the need to minimise inductance of plane connections.

Figure 5C shows various methods for connecting reflow-soldered SMDs to planes. Best is to use over-sized pads, tenting the solder-resist over a number of plane vias. Plane connections that do not need to be soldered (typical of the vias for reflow-soldered SMD components) may not need to use thermal-break pads – and using solid plane connections instead will reduce inductance.

Figure 5C A comparison of the partial inductances of various plane connection methods for SMDs

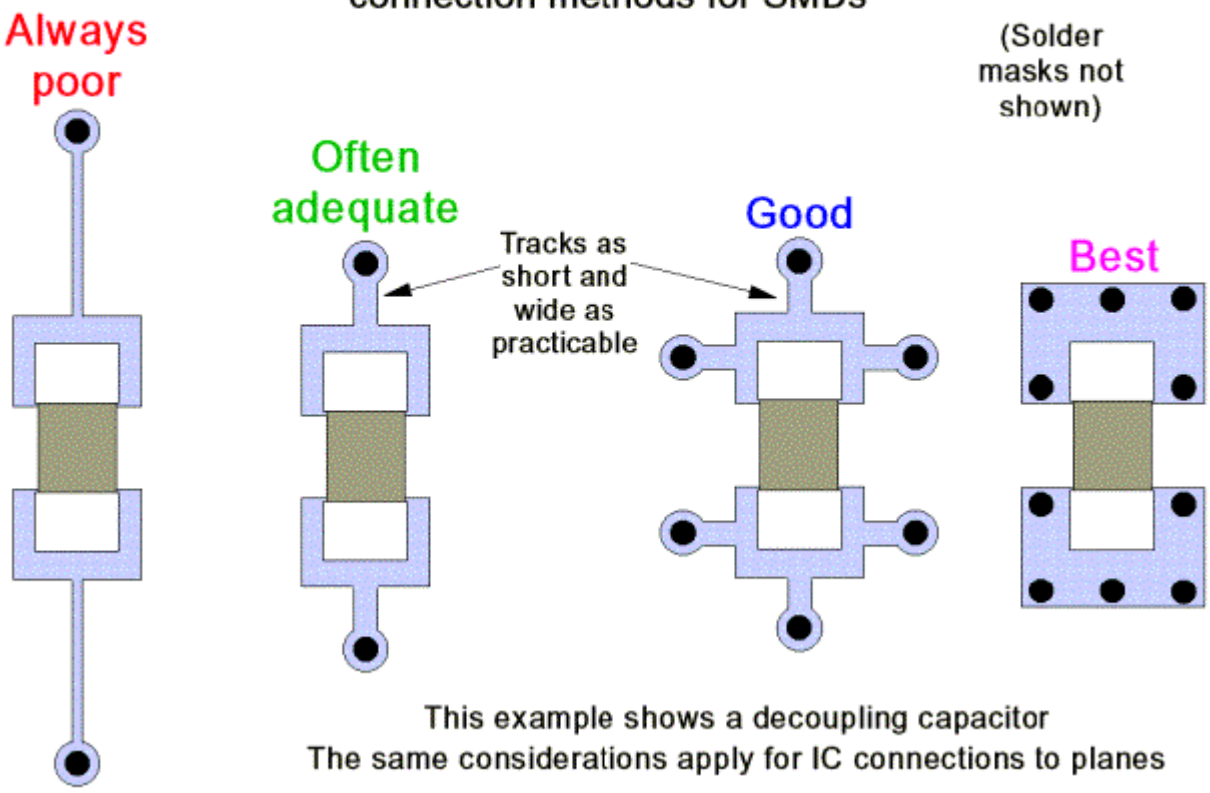

It is best to make reference planes rectangular (but not thin) to minimise their partial inductance, and also to make the fitting of PCB level shielding easier. Square planes, and planes with simple aspect ratios such as 1:2, should be avoided to help reduce possible problems with resonances. Where there are a number of different power supplies, there may need to be a number of different power planes. Segregation of circuit areas (see earlier) makes it easier to fit several broadly rectangular power planes on the same layer.

5.3.2 Connecting 0V planes to chassis

Components and tracks have weak capacitive coupling to everything else. Electrical activity causes displacement currents to flow in these “stray” capacitances, a cause of common-mode emissions. High-speed circuits usually need at least a nearby metal surface, and (increasingly) a fully shielded enclosure, to reduce the resulting emissions problems. The metal chassis or shields need to be connected to the reference plane of their PCB, preferably at a number of points spread over the PCB so that the high frequency displacement currents can be returned to their source within a fraction of their wavelength. PCB mechanical supports and fixings are often used for these chassis bonds, but should be very short (< 4mm). There should be at least one bond in the centre of each area of high-speed circuitry, especially clock generators and distribution. For high-speed digital boards, 0Vplane to chassis bonds every 50 to 100mm all over the PCB may not be overkill, and provision (at least) should be made for these on prototype PCBs. Even if it is not intended to have a metal chassis or shielded enclosure, it is still a good idea to include a number of potential chassis bonding points, just in case. Sometimes a sheet of aluminised cardboard or PVC is sufficient to overcome unexpected problems, providing it can be bonded to the right place(s).

To add flexibility, especially for mixed analogue/digital PCBs, each plane-to-chassis bond can have tracks and pads that allow the bond to be left open, or else fitted with direct links or capacitors of various types and values. Fitting a direct link at one chassis bond and capacitors at the others allows low frequencies (for which the inductance is not important) to be controlled with a “star ground” system, whilst high frequencies are controlled by the low inductance of the widely distributed capacitive links. Care should be taken to minimise the inductances of all these tracks, pads, and linking components (SMD preferred). Where reference planes must be galvanically isolated only capacitor bonds may be used, but care should be taken with safety approvals and earth-leakage requirements (especially for patient-coupled medical apparatus).

5.3.3 Shielding effect of planes

Antennas placed close to metal planes are less effective at radiating and receiving. Many advantages of planes are due to the way they allow the return currents to take the path of least inductance, but their “antenna shielding” effect is also important. For any significant advantage to be achieved from this effect, the tops of all the PCB components must be no more than one-twentieth of a wavelength above a PCB plane, at the highest frequency of concern for emissions and immunity, e.g. 15mm to give a degree of shielding to analogue circuits exposed to 1GHz immunity testing.

Even lower profiles will give improved shielding, one reason why SMD components are much preferred for EMC, with very low profile ball-grid-array and flip-chip technologies being better still. The plane needs to extend by considerably more distance around the components than their height above it.

5.3.4 Interconnecting planes in multi-PCB assemblies

Card cage, backplane, and mother/daughterboard structures will experience considerable signal integrity and EMC advantages from linking their reference planes together with very low inductance. This may be achieved with frequent low-inductance links between their planes, more-or-less uniformly distributed along the *full* length of all their common boundaries. Shielded backplane connectors are happily becoming more commonly available. Where shielded connectors aren’t used, using one 0V plane-linking pin alongside *every* signal or power pin in a connector may seem expensive, but sometimes it is the lowest-cost (or only practical) way to improve the EMC of a multiple-PCB product. Bonding planes via front panels and/or card guides is also very worthwhile.

5.3.5 To split or not to split?

Split reference planes may give better *or worse* EMC (and signal integrity) than unsplit planes, and this depends very much on the PCB layout and circuit design so it is often hard to decide which method to use. Note that where a 0V plane is to be split off from the main 0V plane, it may still need chassis bonds as described earlier. This is particularly true of “traditional” method of splitting off connector panel 0V plane areas (to try stop noise on the main board from exiting via the external

connectors), when the connector area on the PCB must have its local 0V plane bonded to any enclosure shielding. Also note that the inevitable stray capacitance across a split progressively "shorts it out" above 500MHz anyway.

To get any benefits from split planes with modern electronic technologies requires significant attention to detail, which is one reason why an increasing number of designers are now using common, unsplit 0V planes as a matter of course.

Allow for both split and unsplit options, on prototype PCBs at least, by splitting all planes at the natural boundaries between the segregated circuit areas, but also providing the means to "stitch" them together manually later on. Stitching requires pairs of via holes on each side of the split every 10mm or so (random spacing of 5 to 15mm preferred). These via pairs may be left open, or bridged with short wires or capacitors, and it is important to pitch the via pairs close together so that small capacitors or "zero-ohm links" can be used (preferably SMD). Linking planes with a single copper link and multiple capacitors can control lower frequencies (where inductance is not significant) by "star grounding", whilst also controlling higher frequencies by creating the effect of a single low-inductance plane.

Because a split in a plane is a slot antenna, it is best if no tracks cross the split (or even go near to it). Where tracks *have* to cross – they *must* have carefully-defined return current paths, and for high frequency currents these paths *must* be physically adjacent to their send tracks. These tend to defeat the purpose of the split, so should be limited to the bandwidth of the wanted signal (which should already have been restricted to just what is needed, as described in the section on interface analysis and suppression above). High-speed signals can usually be returned through a suitable size and type of capacitor, although some data streams with highly-variable content may need a more wideband return path than a single capacitor can easily provide (may need a direct link).

Balanced signals would ideally need no local return path, but in practice their balance always degrades at some frequency so a nearby return path is needed for the resulting common-mode "leakage" (usually a small-value capacitor). DC power and low-frequency signals that have been filtered to remove all high-frequency noises can use the star point between the split 0Vs for their return, as long as the inductance of the resulting current loop is negligible. Beware of assuming that a conductor is only carrying low frequencies just because that is what its signal name implies. In modern mixed digital/analogue products all the tracks and other conductors in a product usually carry significant levels of high frequency noise. A local return path for a low-frequency signal could be a ferrite bead.

Common-mode (CM) chokes fitted to any types of signals and their associated returns (e.g. a 4-circuit CM choke for a set of three related signals and their return) will probably help get the best performance from split planes, but cost more.

When all the above has been designed into the split-plane PCB, it will need testing and optimisation. Direct or capacitive links to/from the "stitching vias" should be added/subtracted to achieve the best EMC performance. If it is discovered that the best EMC is achieved when all the stitching points are directly linked, the next iteration of the PCB could remove the splits and their stitching points completely, saving manufacturing costs.

5.3.6 Galvanically isolated planes

The split planes described above are all ultimately powered from the same power rails (0V, at least), so there is a clear need for return current paths to be catered for every conductor (signal or power) that crosses from one plane area to another. It is often assumed that galvanically isolated areas have no return current requirements, but this is not so at high frequencies.

Galvanic isolation devices (opto-isolators, transformers, etc.) suffer from stray internal capacitance. A typical opto has 0.8pF internal capacitance, which provides a shunting impedance of only 2kΩ at 100MHz, or 200Ω at 1GHz, which will clearly prevent signal isolation from being maintained at high frequencies. Transformers (especially in DC/DC power converters) tend to have even larger stray internal capacitances. Common-mode chokes may be used to improve the isolation at high frequencies, but struggle to increase it by an order of magnitude at 1GHz. There are also many other stray capacitances around to compromise isolation. So there is a need, at high frequencies, to

provide a local return path for the displacement currents that flow in these stray capacitances, to prevent them from causing common-mode conducted and radiated emissions and immunity problems.

Because we usually only need isolation for low frequencies (usually only 50Hz) we can connect galvanically isolated planes to the main reference plane with a number of low-value capacitors (spread around the gap perimeter), so as to achieve the effect of a single reference plane for high frequencies and provide low-inductance local return paths for stray displacement currents.

Of course, great care may need to be taken with component approvals and leakage currents where safety is concerned.

5.3.7 What if multilayer PCBs are thought too costly?

In volume, four-layer PCBs now only cost between 20% and 50% more than two-layer. The use of planes usually turns out, in retrospect, to have been the most cost-effective EMC technique possible, especially when the *overall* financial break-even time and profitability of a product is considered.

An appropriate technique for low-density double-sided PCBs is to put all the tracks on one side, and a *solid* 0V plane on the other. For digital products, the lack of a power plane might require a number of ferrite beads in the power rails (see later), so it might not prove to be most cost-effective.

Where tracks must use both sides of a two-layer PCB, some EMC improvements may be had by "gridding" 0V tracks. This can be done by using a "maximum copper" or "area fill" on the 0V tracks of both PCB layers, which must run perpendicular to each other, "stitching" the resulting horizontal and vertical 0V areas and lines together with via holes wherever they cross to create a grid over the whole PCB area. Smaller grid sections are needed around the more sensitive or aggressive components, often difficult to achieve for leaded microprocessors but easier for SMD types. Time should be allowed for moving components and tracks around to achieve the best grid structure, but any grid will always be much less effective than a proper solid plane.

Single-sided PCBs are extremely difficult to make EMC compliant without enclosure shielding and filtering, except for circuits which naturally have very low emissions (low dV/dt and dI/dt) and also have naturally very high immunity (e.g. high signal levels and low impedances).

5.4 Power decoupling

The aim of power decoupling is to maintain the power supply impedance to each IC at 1Ω or less across the entire frequency range of interest (at least 150kHz to 1GHz for EMC). Some devices may need 0.1Ω or less over some frequency ranges for correct operation. Wires and PCB tracks have too much inductance to provide these low impedances, which require local capacitance of suitable quality and great attention to detail in PCB layout to minimise inductances.

Another aim is to reduce the size of the current loops in the power distribution, to reduce the emissions form this source. Happily, this is accomplished by the same techniques that lower the power supply impedance.

5.4.1 Power decoupling techniques

A large decoupling capacitor (typically 100μF, might be larger for power-hungry circuits) should be fitted where power supplies enter or leave a PCB, and some smaller ones (e.g. 10mF) should be ‘sprinkled’ around the PCB on a "μF per unit area" principle, as well as being positioned near to heavy power usage such as microprocessors, memory, and other powerful digital ICs. Using electrolytic technology these ‘bulk’ capacitors can provide a low impedance to about 3MHz.

Recently, several manufacturers have added high-capacitance multilayer ceramic capacitors to their surface-mounted product ranges. These are smaller or less costly or have lower ESR and/or better high frequency performance than electrolytics (such as solid tantalum), often several of these attributes at once. They also don’t suffer from reverse polarity or dV/dt problems, so should improve yields and reliability).

Next, the power supplies to every IC should be decoupled very nearby using appropriate capacitor sizes and types. Where an IC has a number of power pins, each pin should have an appropriate decoupling capacitor nearby, even if they are on the same supply (e.g. Vdd).

Achieving good decoupling above 10MHz gets more difficult as frequency increases, because the inductance of component leads, PCB tracks, via holes, and capacitor self-inductance, inevitably limit their performance. The achievement of good power supply decoupling at higher frequencies using capacitors mounted close to IC power pins is discussed next.

The total local decoupling capacitance required depends on the IC's transient power demands and the tolerances of its DC power rails. VLSI and RAM manufacturers should be able to specify the values (and maybe even the capacitor types and preferred layout patterns) for their products, but note that they will probably have assumed an accurate 5V power supply – usually not true of real life.

The formula $C(\Delta V) = I(\Delta t)$, using the units Farads, Volts, Amps, and seconds, covers what we want to know. ΔV is obtained by subtracting the IC's minimum operational voltage (from its data sheet) from the worst-case minimum power rail voltage (taking account of initial tolerances, regulation, temperature coefficients, ageing drift, and the voltage drops in the power conductors). ΔV often turns out to be a mere +100mV. I is the IC's transient current demand from its power rail, which lasts for Δt. I and Δt are almost never found in data sheets, and must be measured in some reasonably sensible way with an oscilloscope. An obvious component of I is the device's output (load) current, but this is often negligible in comparison with "shoot-through" currents, also known as "transient supply current". There is no point in measuring I or Δt with greater than ±20% accuracy.

Where ΔV is low it may be cost-effective to increase it by improving the regulation of the power supply, and/or reducing the resistance of the power rails, rather than fit larger capacitors with their lower performance at high frequencies. This is a common argument for local power regulation.

5.4.2 Self-resonance problems

Self-resonance in capacitors stops them providing low impedances at high frequencies, with higher values generally being worse. The first self-resonant frequency (SRF) of a capacitor is a series resonance, and a rule of thumb for this is: $f_{res} = \frac{1}{2\pi\sqrt{LC}}$, where *L* = ESL (internal to the capacitor) + the total inductance of any leads + the total inductance of any tracks and/or vias. 1nH/mm may be assumed for leads and/or tracks from a capacitor to the power pins of its IC. The inductance contributed by 0V and power planes may be neglected when the capacitor is near to its IC. Decoupling capacitors generally become ineffective at more than 3 times their SRF, as shown by Figure 5D.

Figure 5D Series resonances in decoupling capacitors (guide only)

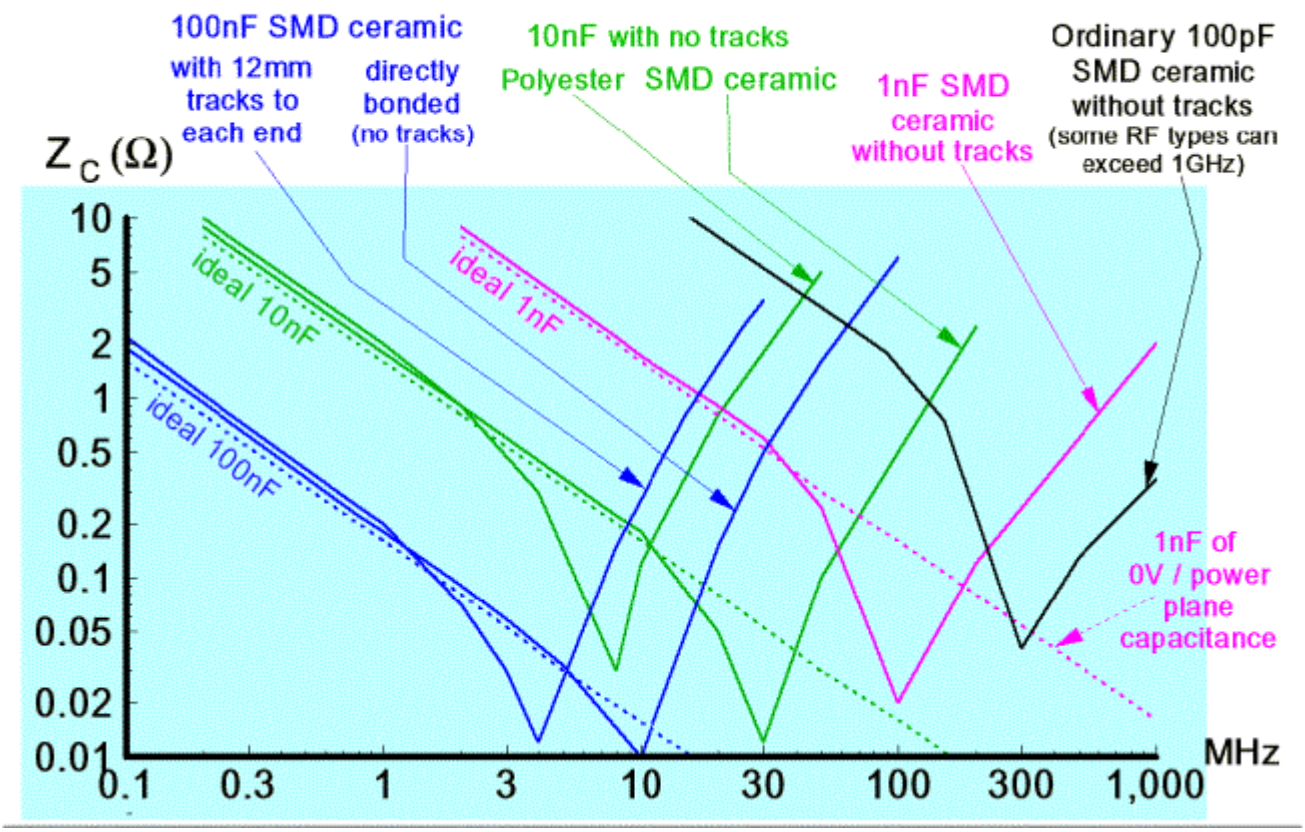

It is interesting to note that the favourite 100nF capacitor, even with no tracks at all, is effectively useless above 50MHz, yet it is still often seen in circuits with clocks of 50MHz or over, where it can do nothing to help control the fundamental clock frequency, never mind its harmonics.

Close proximity of adjacent 0V and power planes (with their low internal and connection inductances) can provide capacitance with no SRF below 1GHz. Two planes separated by 0.15mm in an FR4 PCB achieve approximately 23pF/sq.cm of high quality RF capacitor. Good decoupling from 10 to 1000MHz can be achieved by combining adjacent 0V and power planes with SMD ceramic capacitors (COG and NPO types are best). Sometimes two different values of capacitors (e.g. 100nF and 1nF) may be required. Low inductance bonds from IC power pins and decoupling capacitors to their planes is essential, and the capacitors must be positioned close to their IC. The common practice of tracking from IC power pin to decoupling capacitor, and only then connecting to the plane, does not make best use of plane capacitance.

Whenever two capacitors are connected in parallel, a high-Q (i.e. sharp) high impedance resonance is created which could compromise power impedance at that frequency. This is easily dealt with on PCBs which have a dozen or more decoupling capacitors, since for every sharp high-Z resonance there are a number of alternate current paths with low-Z, which will swamp it. It may be a good idea to fit decoupling capacitors of 10 to 100nF in large areas of planes which are devoid of ICs, to help this swamping process. Parallel resonances are very sharp and often don't correspond to any harmonics so have no effect, but unless it is *known* that this will also be the case for a new PCB (and that no-one will ever alter its clock frequencies) it is risky to ignore their potential for upset.

Parallel resonance problems are more likely to occur where only a few decoupling capacitors are used, for example where a small circuit area is powered from a dedicated power plane. It may be controlled by fitting a low-value resistor (say 1Ω) or small ferrite bead (preferably using SMD packages and short tracks) in series with one lead of the larger value capacitors. Alternatively, adding a number of additional capacitors with differing values should help.

The sudden change in characteristic impedance at the edges of parallel PCB planes creates resonances at integer multiples of half-wavelengths. For example, the first such resonance for a 150mm width or length would be around 500MHz for a bare PCB, moving down in frequency as the PCB was loaded with decoupling capacitors (which slows the velocity of propagation in the planes). This was the reason for recommending non-square rectangular plane shapes (and non-simple aspect ratios) in an earlier section. The resulting high impedances at various areas of the PCB can be controlled by fitting lots of decoupling capacitors, so it is only likely to be a problem for circuits operating at high frequencies with large planes and few decoupling capacitors. There is a suggestion that fitting 1 to 10nF decoupling capacitors around the edges of planes can reduce this effect.

Figure 5E shows a time domain view of how good power supply decoupling functions in an example situation. The first nanosecond or so of transient current can only be provided by local 0V/power plane capacitance, with from 1 to 3ns being provided by SMD ceramic capacitors up to 10nF located nearby.

Figure 5E A time domain view of good power supply decoupling

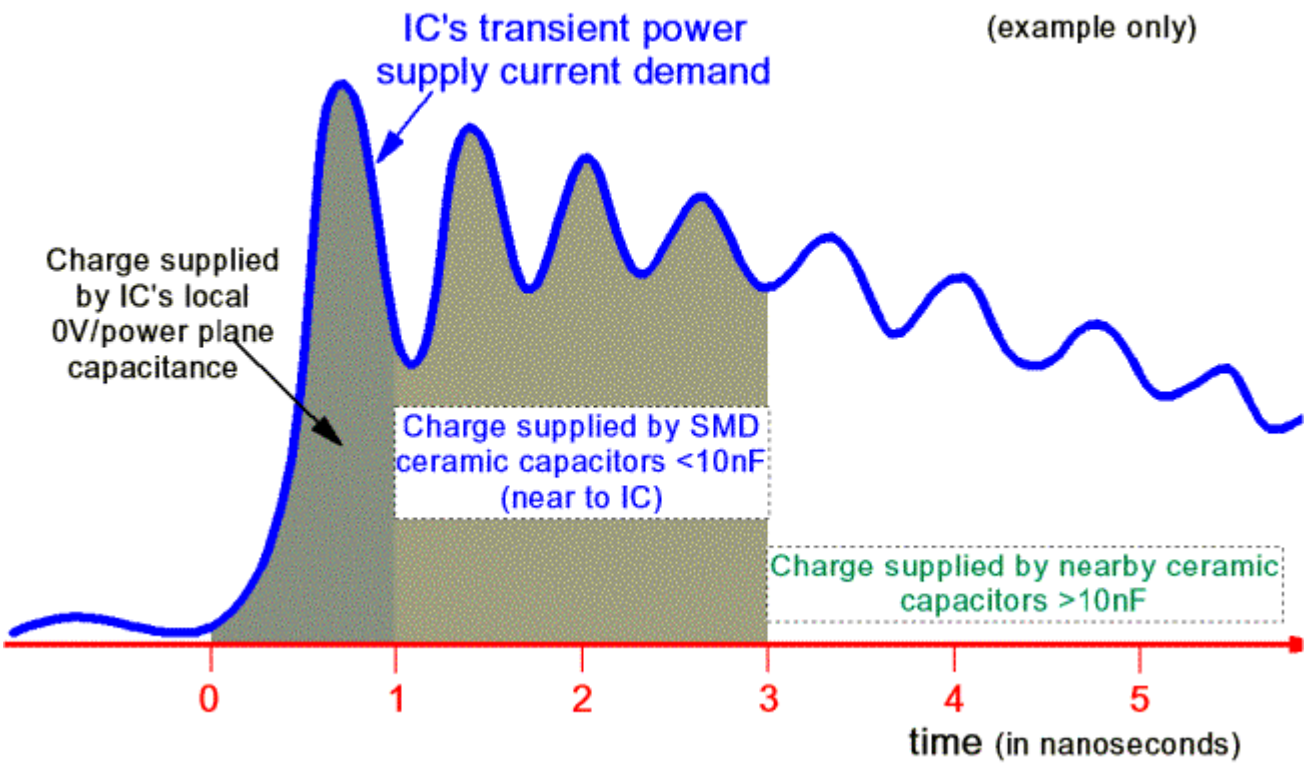

Larger (or further away) capacitors are only able to contribute to the current demand after at least 3ns. "Bulk" capacitors (e.g. tantalums) only provide significant current after 20ns or so, even if nearby (non-ceramic dielectrics and electrolytics are slow to respond to transient current demand due to dielectric absorption effects, also known as dielectric memory or dielectric relaxation).

A PCB process is available that uses a special dielectric between adjacent 0V and power planes to increase their capacitance and eliminate the need for most of the smaller values of decoupling capacitors. Three-terminal or “feedthrough” SMD decoupling capacitors have much higher SRFs than regular two-terminal capacitors, but are more expensive. There are also laminar capacitor components (such as the Micro/Q range) made to fit under leaded ICs, which are also expensive and perhaps best used in attempts to improve existing PCBs without relaying them.

5.4.3 Decoupling without power planes

One way to achieve possibly adequate decoupling without a power plane is to connect one end of an IC's decoupling capacitor to its power pin with very short fat track, then connect that end of the capacitor to the power distribution via a thin track (to create some inductance) or ferrite bead, rated for the IC's current. Three-terminal or high-specification capacitors may be used to advantage so that a high SRF is achieved with a single decoupler. This technique still requires a 0V plane. Where a large number of ferrite beads or expensive capacitors are required, multi-layer boards may prove to be more cost-effective and require less area.

5.5 Transmission Lines

Transmission lines maintain a chosen characteristic impedance, Z_0, from a signal's source to its load, and (as discussed in Part 2 of this series) unlike all other interconnections do not resonate however long they are. Transmission lines can easily be made on PCBs by controlling materials and dimensions and providing accurate termination resistances at source and/or load. They may also be extended off the PCB (if necessary) with appropriate controlled-impedance cables and connectors.

Comparing the length of a PCB track conductor with the wavelength of the highest frequencies of concern in the relevant medium (e.g. FR4), or with the rise- and fall-times of a signal, gives us what is called the 'electrical length' of the track. Electrical length may be expressed as a fraction of a wavelength or as a fraction of the rise- or fall-time. When a conductor is 'electrically long', transmission lines need to be used to maintain the frequency response (sometimes called 'flatness') or to prevent excessive distortion of the waveshape. For high-speed signals on PCBs, transmission line techniques are required for all electrically long tracks both for signal integrity and EMC.

The crude rule of thumb are that a conductor is electrically long when it exceeds one-seventh of the shortest wavelength of concern, or when the time that the leading edge of a signal takes to travel from the source to the furthest receiver exceeds half of its rise or fall times. Consider Fast TTL, which is *specified* as having 2ns risetimes. The dielectric constant of FR4 at high frequencies is around 4.0, which gives a signal velocity of 50% of *c*, or 1.5×10^8 m/s, equivalent to a track propagation time of 6.7ps/mm. In 2ns a signal in an FR4 PCB would have travelled about 300mm, so it appears that Fast TTL signals need only use transmission lines for tracks of 150mm or longer. Unfortunately, this answer is wrong. The 'half risetime' rule is very crude and can lead to problems if its shortcomings are not understood.

Databook specifications for output rise/fall times are *maximum* values, and devices almost always switch a lot faster (assume at least four times faster in the absence of actual data). It is best to measure a number of samples from different batches, and obtain an agreement from the device manufacturer that he will warn well in advance of mask-shrinks. Also, the inevitable capacitive loading from connected devices reduces the propagation velocity from what would be achieved on the bare board. So transmission lines should be used for much shorter lengths of track than suggested by the above rule, merely to achieve adequate digital signal integrity. Taking these two issues into account, we may find that specified 2ns rise/falltime signals should use transmission lines for tracks that are longer than 30mm (and possibly even less).

Transmission lines are often used for clock distribution and high-speed busses; for slower signals that have to travel further, such as SCSI and USB; and also for even slower communications such as 10base-T Ethernet and RS485, which have to travel very long distances.

Most transmission lines are used to preserve the waveshape of high-speed signals, and to reduce their emissions, but transmission-line techniques work in just the same way to reduce the amount of external fields *picked-up* by a track, so are valuable for EMC immunity reasons too. It may help to use transmission lines for low-bandwidth signals (e.g. analogue instrumentation) to prevent their contamination by high-frequency fields in their environment (which could be inside a product), since analogue devices are particularly prone to demodulating RF at hundreds of MHz (refer to Part 1). When designing a transmission line for immunity purposes, the 'shortest wavelength of concern' or the 'highest frequency of concern' is the important parameter.

IEC 1188-1-2 : 1998 [2], gives a wealth of details on constructing a wide variety of transmission lines with PCB tracks, plus how to specify their manufacture and check quality at goods-in. [1], [3],

[4], and [5] are also very helpful with this large and detailed topic, so only the two most common types are described below.

The first example is a *surface microstrip* (see figure 5F), and its Z_0 is given in ohms by:

$$Z_o = \frac{87}{\sqrt{\varepsilon_r + 1.41}} \times Ln\frac{5.98H}{0.8B + C}$$

where ε_r is the relative permeability of the substrate (typically 4.4 for FR4 at 100MHz), B is the track width, C is the thickness of the copper material used, and H is the substrate thickness.

Its propagation velocity in ns/metre is: $3.335\sqrt{0.475\varepsilon_r + 0.67}$.

Figure 5F A surface microstrip

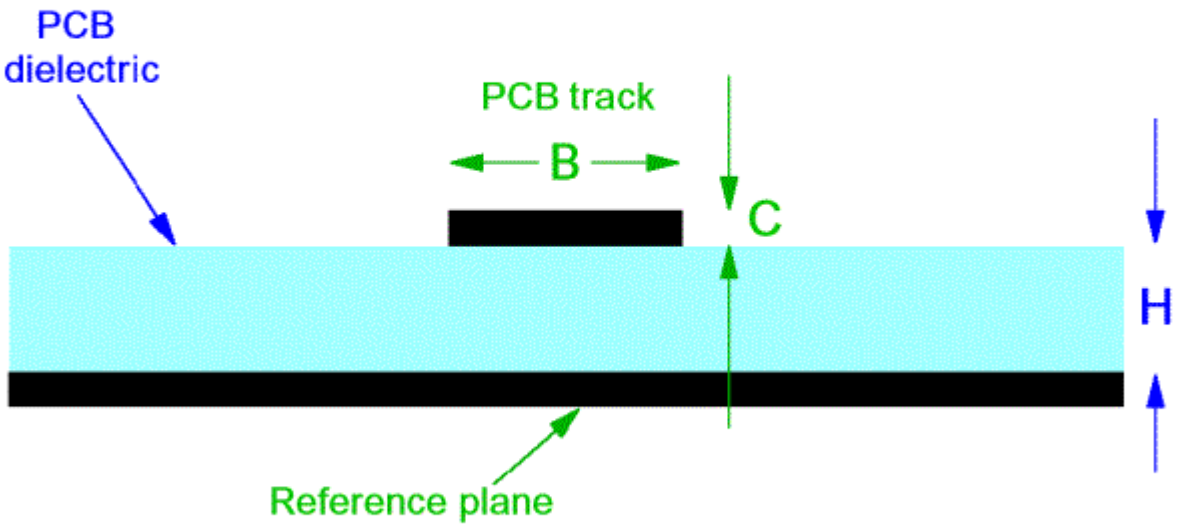

The second example is the symmetrical stripline (figure 5G), which uses two reference planes:

$$Z_o = \frac{60}{\sqrt{\varepsilon_r}} \times Ln\frac{1.9H}{0.8B + C}.$$

The propagation velocity for a symmetrical stripline in ns/metre is: $3.335\sqrt{\varepsilon_r}$.

Figure 5G A 'symmetrical stripline'

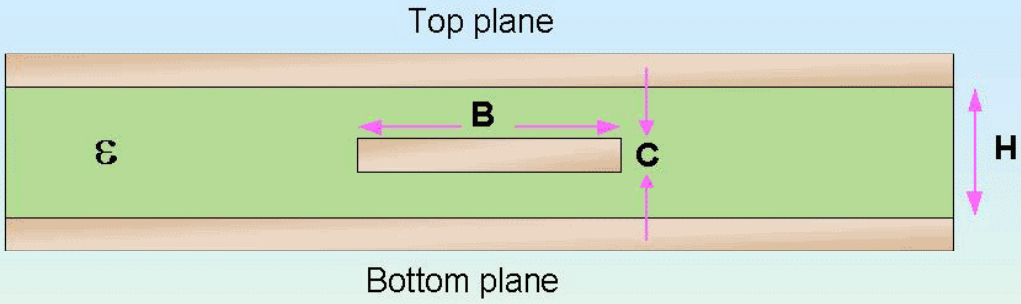

In a symmetrical stripline, the trace is mid-way between the two planes

$$Z_0 = \frac{60}{\sqrt{\varepsilon}} \times \mathrm{LOG_e}\ \frac{1.98H}{0.8B + C} \quad \Omega$$

Where:

- **ε = Relative dielectric constant** (typically 4.2 for FR4 at 100MHz)
- **H = Dielectric thickness**
- **B = Trace width**
- **C = Trace copper thickness** (e.g. 0.017mm for '½oz' and 0.034mm for '1oz' copper)

'Thou' (mil), inches, metres or millimetres may be used, as long as the same units are used throughout the formula

Striplines are slightly slower than microstrip, but have zero forward crosstalk and much less off-board leakage, so are best for EMC.

[1] gives correction factors for the above formulae to compensate for capacitive loading (typically a few pF per gate): $Z'_0 = \frac{Z_0}{\sqrt{1+\frac{C_d}{C_0}}}$ where C_d is the sum of all capacitive loads, Z_0 is the original characteristic impedance (unloaded) of the line, and C_0 is the characteristic capacitance of the (unloaded) line obtained from the basic formulae given in [2].

Velocity is slowed according to the formula: $V' = \frac{V_0}{\sqrt{1+\frac{C_d}{C_0}}}$ from [1], where V_0 is the original (unloaded) velocity. A constant "gates per unit length" is preferred for the layout of an array of load devices, rather than bunching them together, although it may be possible to adjust the line dimensions for different portions of the track so that the same Z_0 is maintained all along its length, even where load devices are bunched together.

The highest-speed (or most critical) signals should run adjacent to a 0V plane, preferably one paired with a power plane for decoupling. Less critical signals may be able to be routed against a power plane where the power plane has been adequately decoupled and is not too noisy (i.e. has been properly decoupled, see earlier). Any such power plane must be the one associated with the signal's ICs. Striplines routed between two 0V planes (one or both of which is paired with a power plane for decoupling) give the best signal integrity and EMC.

Transmission lines must not have any breaks, gaps, or splits in any of the planes they are routed over, as these cause sudden changes in Z_0. They should also stay as far away as possible from any breaks, gaps, splits, or plane edges, by at least ten times their track's width. Low crosstalk requires spacing adjacent transmission lines by at least three times their track widths. A very critical or aggressive signal (e.g. a radio antenna connection) may benefit from using a symmetrical stripline with a row of closely spaced vias between its two 0V planes all along each side, 'walling it off' from other tracks and creating a coaxial type of structure in the PCB. This requires a different Z_0 formula from those above.

The two transmission line types above require two or more PCB layers, so can be costly to achieve in high-volume low-cost products (although in volume a 4-layer PCB should cost no more than 20% more than a 2-layer). Balanced and co-planar line types can be constructed on a single PCB layer, so may be a solution where high-speed signals must use lowest-cost PCBs. Single-layer transmission lines will take between two and three times the area as microstrip or stripline, so be prepared for their real estate demands. Also beware of saving so much cost on the PCB that the cost of the enclosure shielding and filtering has to be increased. It is a general rule that solving an EMC problem at enclosure level costs between 10 and 100 times more than it would have if it was solved at PCB level. So when trying to pare costs to the bone by reducing the number of PCB layers, allow the time and cost for a couple of additional PCB iterations to get the EMC and signal integrity within specification and within budget, and also allow for additional PCB area.

5.5.1 Changing layers

High-speed or other critical transmission lines should not change layers. This means routing clock distribution first, moving components around to achieve the smallest area of highest-speed circuitry and tracks. High-speed busses, fast data communications, and the like are routed next, still sticking to one layer, and then everything else (less critical for signal integrity or EMC) is routed around them, changing layers as necessary. Where there is no reasonable alternative to changing the layers of a critical transmission line, a decoupling capacitor (with a suitable frequency response) should be fitted, with its vias linking all the relevant power and 0V planes, near to the point where the signal changes layers.

Keeping to the same layer is easy when using surface-mounted devices with microstrip transmission lines on the same side of the PCB. Stripline is less leaky than microstrip – but this would mean changing layers, which is generally a bad thing. (Microwave circuit designers often employ microstrip with surface-mount devices with leads that are exactly the same width as the PCB transmission line impedance (usually 50Ω), but they also generally fit each gain stage inside its own milled pocket in an aluminium housing. Such techniques are not generally suitable for computer and DSP boards.) Changing layers is usually necessary, if not to use striplines for their beneficial effects, then because of track densities, so how can we mitigate its effects?

We already have at least one decoupling capacitor associated with each IC (see 5.4 above) so we can change layers near to an IC, but we must consider the electrical lengths of the portions of the signal path that do not share the stripline layer. A crude rule of thumb is that these portions should not have an electrical length longer than one-eighth of the rise time (approximately one-thirtieth of the shortest wavelength of concern). Where very large changes in Z_0 can occur (e.g. when using a ZIF or other IC socket) it would be better to aim for less than one-tenth of the rise time. Use these rules to determine the *maximum* length, and keep well within this length wherever possible.

So for signals specified as having 2ns rise time signals we should probably change layers no further than 10mm from the centre of the IC's body or the centre of the line termination resistor. This includes a 'safety factor' of 4 to allow for the actual edge rate of the signal being faster than the data sheet maximum. At least one decoupling capacitor, which connects all the relevant powers and grounds together (respectively) should also be within a similar distance from any transmission line that changes layers. Such short lengths are often difficult to achieve with larger ICs, and this reveals some of the compromises inherent in modern high-speed PCB layout. This also reveals at least one reason why physically smaller ICs are preferred, and why bonding techniques such as BGA and flip-chip (which reduce the distance from the PCB track to the silicon itself) are continually being developed and improved upon.

5.5.2 Simulation and prototype testing

Because of the variations in the types of ICs, and the applications they find themselves in, some engineers will find these rules of thumb not tough enough, and some will wonder whether they are over-engineered, but that is the function of a rule of thumb, after all.

Computer-based circuit simulation techniques that calculate EMC and / or signal integrity based on parameters extracted from an actual PCB layout are becoming more capable, and their use is recommended instead of the crude rules-of-thumb expressed here. However, remember that device switching speed are almost always considerably faster than their data sheet specifications, so a simulation that uses data sheet figures will give a false sense of confidence.

Tests with a high-speed oscilloscope and probing system should be carried out on the first PCB prototypes to see whether the waveshape is good enough. A waveshape that does not distort as it travels around the PCB is the goal, and merely following these rules of thumb is unlikely to achieve such perfection, although the result may be good enough. Close-field probing with a single-turn loop, using a high-speed 'scope and/or a spectrum analyser, is another good way to detect signal integrity or EMC problems at prototype board level. The techniques involved in prototype testing are not discussed further here.

Even when using sophisticated modelling or simulation techniques, always perform signal integrity and EMC checks on early prototypes.

5.5.3 Manufacturing issues with transmission lines

Normal FR4 PCB material has a nominal relative dielectric constant (ε_r) of approximately 4.7 at 1MHz falling roughly in a linear fashion with increasing frequency to 4.2 at 1GHz. Actual values of ε_r, can vary by ±25%. Controlled ε_r grades of FR4 are available at little or no extra cost, but PCB manufacturers may not use these grades unless specifically requested.

PCB manufacturers work with standard thickness laminations ("prepregs"), and their thicknesses should be discovered (along with their manufacturing tolerances) before design starts. The track widths can then be chosen to achieve the required Z_0 for the available range of dielectric thicknesses. Track widths after PCB processing are usually about one thousandth of an inch less than those used on the photoplots. Ask what thickness to add to the drawn tracks to achieve the required finished track widths.

For signal frequencies greater than 1GHz it may be necessary to use other dielectric materials than FR4, such as those used for microwave applications (e.g. Duroid from Rogers Corporation Inc., or a number of more modern dielectrics).

5.5.4 Terminating transmission lines

"Classical" RF transmission lines are terminated both at signal source and load by impedances equal to their Z_0 (allowing for the internal impedances of source and load devices). Although an ideal and sometimes necessary technique, it halves the received voltage – so most ordinary analogue and digital circuits use low-Z sources and high-Z loads with the line only terminated at one end, to preserve signal levels.

RF engineers often use reactive components or even lengths of track as line terminations, but the terminations for wideband analogue and conventional digital signals require individual resistors, preferably SMD types for their excellent high-frequency performance. To get the best from SMD resistors they must be connected to the reference planes using low inductance techniques as shown by figure 5C.

Figure 5H shows the common termination techniques. Classical RF termination is still often used for high-speed signals such as fast backplane systems.

Figure 5H Various transmission line termination methods

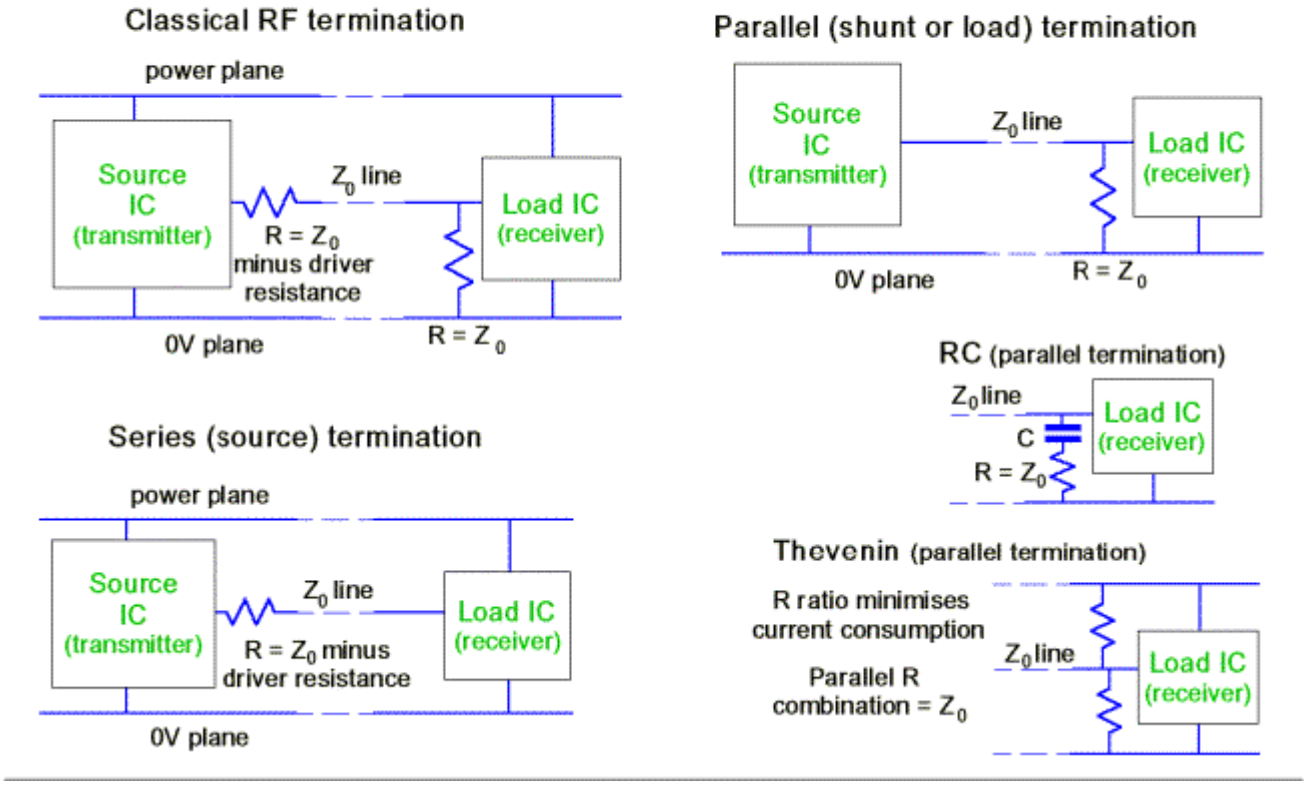

Where signals are restricted to a single PCB, series (source) termination may be used at the driver end of a transmission line, with the resistor chosen so that in series with the impedance of the output driver it matches the line's Z_0. This method has the advantage of consuming little power, and is most suitable for lines with a single load device at their far end. Where other loads exist along the length of the line they experience "reflected wave switching" and their response may need to be slowed to prevent false clocking.

Parallel (shunt, or load) termination at the very far end of a line is used where there are a number of devices spread along the length of the line and they need to respond most quickly, and achieves "incident wave switching". Figure 5H shows the termination resistor connected to the 0V plane, but some logic families use other reference voltages (e.g. the positive plane for ECL). Parallel termination dissipates a lot of power, and may also load some IC outputs too heavily.

Alternative types of parallel termination include "Thévenin" and RC. Thévenin uses resistor values designed so that their parallel resistance is Z_0 and they would provide a DC voltage at their junction equal to the average line voltage, to minimise power dissipation. Thévenin termination needs a properly decoupled power plane at all frequencies of concern so needs decoupling capacitors nearby. RC termination uses capacitor values between 10 and 620pF (typically) and only terminates the line for high frequencies. Because of the problems of capacitors (discussed earlier) it may be more difficult for an RC termination to equal the highest frequency performance of a parallel resistor or Thévenin termination.

"Active termination" uses a voltage regulator to drive an additional power plane at the nominal average value of the digital signals. A parallel line termination connects to this plane, which must be properly decoupled for the frequencies of concern. Electrically equivalent to the Thévenin method, this can save power by running the voltage regulator (which needs to be able to source as well as sink current) in Class AB.

Where a line is driven bi-directionally the compromise position for terminating resistors (series or parallel) is in the centre of the line, so such lines should always be kept very short and may not be able to run as fast as the device speeds may suggest. Series terminations at all possible drivers may be used instead of series termination at the centre of the line, but this may not give good signal integrity unless all the lines concerned are very short. Parallel termination at both ends of the line can give very good performance and allows the highest data rates, but drivers must capable of driving the resulting lower impedances, and power dissipation will also increase. Parallel (or Thévenin or active) termination at both ends is used for serial or parallel data cables such as SCSI and Ethernet.

When “star” connecting a number of individual series-terminated transmission lines, either use one termination resistor chosen so that the total source resistance equals the parallel combination of all the starred lines, or else use one resistor to match each line. The latter technique should be better. The star configuration may also be used to drive multiple parallel-terminated lines. In either case, the signal source must be capable of driving the parallel combination of all the lines' Z_0s.

It is generally better to choose higher values of Z_0 to reduce signal currents and reduce radiation from the tracks. Many ordinary CMOS or TTL ICs were never designed for driving transmission lines, and have neither the drive capability or an output impedance that is equal for both sourcing and sinking. Such devices may be able to use series, Thevenin, RC, or active terminations on high-impedance lines, but the best method to use, and the line impedance, may be difficult to predict for a given logic family.

However, an increasing number of devices are becoming available to drive transmission lines, and the increasing range of LVDS and similar devices is making clock and bus driving much easier and easing EMC problems. Backplane bus driver ICs are available with 25Ω output impedances, suitable for “star” driving four individual 100Ω, or six 150Ω lines. Some devices now have on-chip DC/DC converters which cause their unloaded outputs to achieve double the correct logic levels, so that when operated into a classically terminated line the received logic levels are correct.

5.5.5 Layer "stack up"

The above section on decoupling shows it is good EMC practice to provide 0V and power planes on adjacent layers and to maximise their capacitance by using a thin dielectric (say 0.15mm) between them. The above section on transmission lines shows that proximity to a reference plane is important for high-speed tracks. We can put this all together to decide how to stack up our PCB layers.

Four-layer PCBs often have their layers stacked as follows:

1) Microstrip transmission lines and other critical signals
2) 0V plane
3) +5V plane
4) Non-critical signals

Where more signal layers are required, a 0V and power plane "core" should be retained. Additional layers of high-speed signals may need additional 0V planes to be added, but high-speed clocks and data busses and similarly aggressive or very critical tracks should not swap layers.

Here is one of a number of possible stack-ups for an 8-layer computer motherboard:

1) 0Vplane
2) Most critical “offset striplines” and other signals, routed at 90° to layer 3) to reduce crosstalk
3) Most critical “offset striplines” and other signals, routed at 90° to layer 2) to reduce crosstalk
4) 0V plane
5) +5V plane
6) Non-critical signals routed at right angles to layer 7) to reduce crosstalk
7) Less critical “offset striplines” and other signals, routed at 90° to layer 6) to reduce crosstalk
8) 0V plane

5.5.6 Joints, stubs, and buffers

The above has treated transmission lines as if they were all point-to-point connections, so we need to address bussed systems such as RAM arrays, and situations where several cards interconnect at high speeds, such as backplane systems.

A length of track that springs off from a joint with the main track is called a ‘stub’. For memory arrays, the usual PCB layout technique is to route busses horizontally on one layer, ‘via-d’ through to another layer with vertically routed stubs to connect to the array devices. To preserve the transmission line the electrical length of the stubs must be kept to under one-eighth of the rise time (and preferably much less). Don’t forget that the important parameter is the real switching rate of the signals, not the data sheet specification of the drivers. If in doubt, assume the drivers switch four time faster than their data sheet maxima.

The stub length used for the calculation should include the distance from the end of the track (the IC’s soldered pin) to the centre of the IC itself. Where the permissible stub length is too short for traditional ‘horizontal and vertical’ array routing, daisy chain tracking should be used instead. Daisy chain tracking is generally better for high-speed signals in any case, especially when the tracks remain all on one layer. Daisy chain tracking means that the bussed tracks go from the source directly to each load in turn. Abrupt changes in track direction should be avoided, with gentle curves or large chamfers used instead. In an incident wave system, the daisy-chained tracks would end in the parallel, Thevenin, RC, or active termination resistors.

When electrically long stubs can’t be avoided, buffers should be fitted close to the main track to minimise the stub lengths. This is often used in backplane systems, where a number of plug-in cards must all run from the same clock lines and data busses, as shown in Figure 5J. The clock buffers must all be fitted very close to the backplane board connectors, and as signal speeds and data rates increase it is more common to find that matched-impedance backplane connectors are needed. Where a plug-in card only has one or two ICs that need to connect to the backplane clock and data lines, by placing them close the backplane connector it may be possible to do without the buffers.

Figure 5J An example of buffering to prevent long stubs

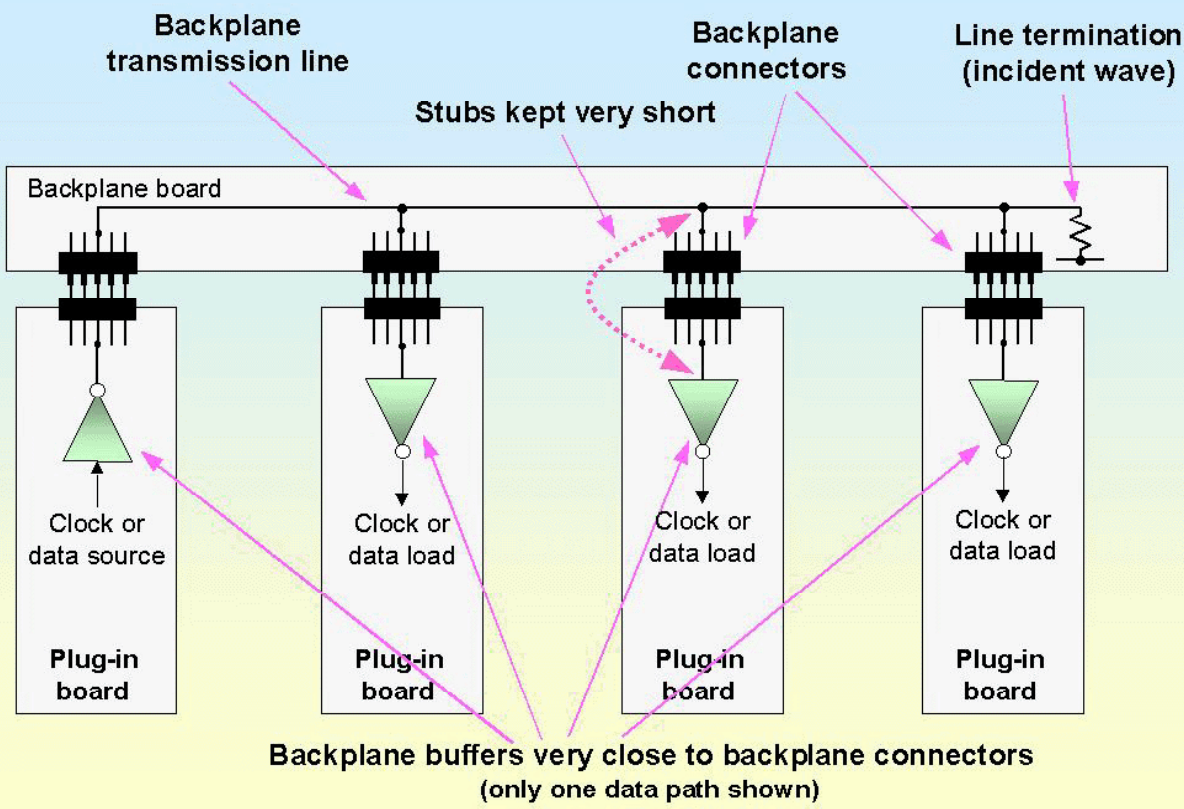

Planes must make low-impedance connections between backplane and plug-in PCBs, ideally multiple connections along the full height of the plug-in boards

Buffering is also a good technique for reducing the loading on a transmission line. For example, where there are ten plug-in cards each with ten ICs, all receiving one signal, their combined load capacitance can be around 400pF. The signal and return currents for this high value of capacitance have a long way to flow, increasing the likelihood that they will create EMC problems. Buffering the signal at each card means that the main line is only loaded by around 40pF, while the signal and return currents for the ten devices on each card now flow only in that card, improving signal integrity and reducing EMC problems.

Carrying high-speed signals through connectors and backplanes, it is important (vital for transmission lines) to maintain the same physical structure. For example, striplines in plug-in boards should be continued as striplines in the backplane, (although it is possible with some degradation in signal integrity to swap from one type of transmission line to another as long as the track dimensions maintain the same Z_0). Where transmission lines entering a backplane connector are routed against a power plane, that power plane should be continued through the connector into a power plane in the backplane and then to the associated power planes in the other cards using that signal. The interconnections between the power planes in the boards and the backplane should be designed in the same way as for the 0V return planes. Some boards may find that their optimum backplane connector pinning needs to be: 0V, signal 1, +5V, signal 2, 0V, signal 3, +5V, signal 4, 0V, ….etc.

5.5.7 Segregation in backplane systems

Section 5.1 above said high-speed devices should be kept in the middle of their segregated area, well away from any PCB or reference plane edges, or connectors. The backplane system described above and in Figure 5J places the fastest ICs close to the backplane connector but does not

compromise the earlier rules if the backplane is designed to be an extension of the plug-in boards at high frequencies.

This requires RF bonding the reference planes in the backplane to the corresponding reference planes in the plug-in boards, so that at the highest frequency of concern there appears to be no impedance discontinuity between them. Using shielded connectors helps - their shields should bond their mating halves to each other in 360^{o} (refer to Part 2 of this series), and also bond all along their length to the 0V reference plane on both sides. Whether shielded connectors are used or not, there should be one signal return pin for every one or two (at most) signal or power pins in the connector, and these return pins should be spaced fairly regularly along the entire length of the connector. Most designers would place the return pins according to a regular scheme, but there is some evidence that a randomised allocation can have benefits. Impedance-matched connectors will almost always have a return pin alongside every signal pin, in any case.

It is important to make sure that all the high-speed devices associated with the backplane connector are closer to the middle of the connector, and do not go near to the outer edges of the boards, near the ends of the connector. It is best if the backplane connector extends to occupy the entire length of the card edge, but if it can't it should still extend well to both sides of the area of the high-speed devices associated with the backplane connector.

5.6 Useful references

[1] M. Montrose, *EMC and the Printed Circuit Board, design, theory, and layout made simple*, IEEE Press, 1999, ISBN 0-7803-4703-X

[2] *IEC 61188-1-2 : 1998 Printed Boards and Printed Board Assemblies – Design and use. Part 1-2: Generic Requirements – Controlled Impedance, www.iec.ch.*

[3] IPC-2141, *Controlled Impedance Circuit Boards and High Speed Logic Design*, Institute for Interconnecting and Packaging Electronics Circuits, April 1996, www.ipc.org.

[4] T Williams, *EMC for Product Designers 3rd Edition*, Newnes, 1992, ISBN 0-7506-4930-5, www.newnespress.com.

[5] H W Johnson and M Graham, *High Speed Digital Design, a Handbook of Black Magic*, Prentice Hall, 1993, ISBN 0-13-39-5724-1.

Eur Ing Keith Armstrong C.Eng MIEE MIEEE
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone: 01457 871 605, Fax: 01457 820 145, Email: keith.armstrong@cherryclough.com

Design Techniques for EMC – Part 6

ESD, Dips, Flicker, Dropouts, Electromechanical Switching, and Power Factor Correction

By Eur Ing Keith Armstrong CEng MIEE MIEEE
Partner, Cherry Clough Consultants, Associate of EMC-UK

This is the final part of a series of six articles on best-practice EMC techniques in electrical/electronic/mechanical hardware design. The series is intended for the designer of electronic products, from building block units such as power supplies, single-board computers, and "industrial components" such as motor drives, through to stand-alone or networked products such as computers, audio/video/TV, instruments, etc.

The techniques covered in the six articles are:

1) Circuit design (digital, analogue, switch-mode, communications), and choosing components
2) Cables and connectors
3) Filters and transient suppressers
4) Shielding
5) PCB layout (including transmission lines)
6) **ESD, dips, flicker, dropouts, electromechanical switching, and power factor correction**

A textbook could be written about any one of the above topics (and many have), so this magazine article format can only introduce the various issues and point to the most important best-practice techniques. Many of the techniques described in this series are also important for improving signal integrity: reducing the number of iterations during development and reducing manufacturing costs.

Table of contents for this part

8. Conclusion to the Series

6. A number of specific issues

The previous parts of this series focussed on design techniques which will benefit a large number of emissions and immunity characteristics whilst also improving signal integrity. This article finishes off the series with a number of issues where specific techniques may be required, and signal integrity may not be a concern.

6.1 Electrostatic Discharge (ESD)

6.1.1 Different types of ESD

The high voltages that cause ESD arise through tribo-charging, the natural process by which electrons get transferred from one material to another of a different type when they are rubbed together. Man-made fabrics and plastic materials are often very good at tribo-charging, so ESD problems tend to be on the increase. ESD is a very fast phenomenon, and very intense while it lasts (usually just a few tens of nanoseconds overall).

Machinery ESD occurs when isolated metal parts rub against insulating materials, or have a flow of insulating liquids or gases over them. The metal parts tribo-charge until they discharge with a spark into something nearby which was not previously charged, equalising their potentials. Sparks created in this way by machinery can be very intense, especially when the metal part being charged is large and so has a large capacitance, which can store a large amount of charge.

Furniture ESD occurs when metal furniture (or parts of furniture) such as chairs, tables, cabinets, etc., become tribo-charged by friction against insulating materials. This may happen when the furniture is moved across a carpet or plastic floor covering, or because materials are rubbed against it, for instance when a person gets up from a chair.

Personnel ESD is caused by people becoming tribo-charged, usually by walking around. Walking on plastic floor coverings, synthetic carpets, etc., is the usual cause of personnel ESD. Few people can even notice sparks from their fingers which are under 2.5kV.

Spacecraft ESD is not covered here, although many of the techniques described will be applicable.

All these three types of ESD are very important in the manufacture of semiconductors and the assembly of electronic products, and in these areas great lengths are taken to prevent the three types of ESD from reducing yields. Machine ESD can be a big problem for process control automation. But personnel ESD is the only type of ESD which we find in EMC standards harmonised under the EMC Directive. ESD causes EMC problems in three main ways:

- The spark voltages which get into semiconductors can easily damage them. Modern semiconductors use internal insulation which can breakdown and permanently short out areas of the device at just a few tens of volts. This is known as a hard failure.
- Most ICs are made with built-in protective devices to help prevent them from damage by ESD during handling and assembly. However, these internal devices can't be made large enough to handle large amounts of power, and a significant ESD event can over-dissipate them, sometimes while leaving the semiconductor still functional. This is known as a soft failure, because the semiconductor usually fails a few weeks or months later.
- The intense transient electric and magnetic fields created in the vicinity of an ESD spark can induce voltages or currents into nearby circuitry and upset its operation. This does not usually cause direct damage, although the resulting malfunction can sometimes cause consequential damage of some sort.

6.1.2 The "Human Body Model" and ESD testing

The ESD simulator used for testing to EN 61000-4-2 is based upon the 150pF/330Ω human body model, and generates a waveform with a risetime of between 700ps and 1ns to reach a peak of several kV, which then decays to about 50% in 50ns. At a voltage of 8kV the peak current into a 50Ω calibration load is close to 20A. The frequency content of such an ESD waveform is flat to around 300MHz before it begins to roll off, so contains significant energy at 1GHz and above.

Some older test standards use an older human body model which only has a 5ns risetime, so its spectrum begins to roll off at 60MHz and it is not as aggressive a test as EN 61000-4-2. As high-speed measurement techniques improve, it appears that real ESD events may have risetimes faster than 700ps.

Testing to EN61000-4-2 (personnel discharge) involves the following:

- Air discharges of up to ±8kV (using an 8mm round tip to simulate a human finger) are applied to everything non-metallic which is normally accessible to the operator.
- Contact discharges of up to ±4kV (using a sharp tip which is touched against the product before the discharge) are applied to operator-accessible metal parts – and also to nearby vertical and horizontal metal planes.

Test voltages are increased gradually from low values, often using the settings 25%, 50%, 75%, and then 100% of the test voltage. This is because ESD failures are sometimes seen to occur at lower voltages but not at the maximum test level. The highest test level on an ESD test is not necessarily the one most likely to cause a failure (this is also true for other types of transients). Figure 6A is a sketch of the barest essentials of an ESD 'gun'.

Figure 6A

Figure 6A
Basic circuit of ESD 'gun' for personnel ESD testing to EN 61000-4-2

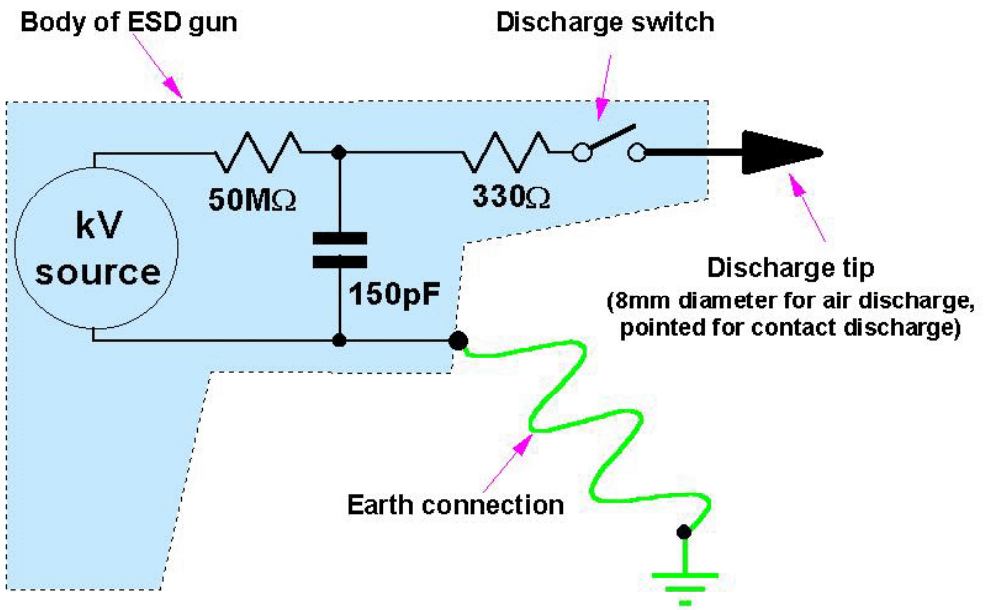

Bear in mind that in dry climates personnel ESD events can easily exceed 8kV. 15kV or even 20kV is not that unusual during freezing winter conditions when the air is very dry, especially in heated homes and buildings without humidity control. So, meeting an ESD test at ±8kV is not a guarantee of freedom from actual ESD problems in the field, and the environment and needs of the users should be taken into account when ESD testing to help produce reliable products.

6.1.3 Design techniques for personnel ESD

All the design techniques described in the previous parts of this series help a great deal in improving the immunity of circuits to electric and magnetic fields, and so help circuits cope with the brief but intense bursts of wideband disturbances from ESD events. However, they are not usually enough on their own. The two main techniques for preventing ESD sparks from upsetting products are:

- Dielectric protection (insulation)
- Shielding (metal or metallised enclosures)

  Dielectric protection is the preferred technique, but where it cannot be used for an entire product ESD problems can occur with both internal and external connections. These are discussed below. Apart from dielectric isolation, many of the techniques described below will also be useful for protection against other conducted transients and surges, which have not been dealt with in this series of articles as a separate topic.

6.1.4 Dielectric protection

This is the best ESD protection method. By not allowing an ESD spark to occur at all, not only are sparks prevented from getting into sensitive circuitry, but no bursts of electric and magnetic fields occur either.

Plastic enclosures, membrane keyboards, plastic knobs and control shafts, plastic switch caps, plastic lenses, etc., are all pressed into service to insulate the product (especially the operator-interface areas and controls). A 1mm thickness of common plastics such as PVC, polyester, polycarbonate, or ABS, is usually more than adequate to protect from 8kV of ESD (check the breakdown voltage rating of the material in kV/mm of thickness). But since no practical enclosure is without seams, joints, and ventilation, the achievement of adequate creepage and clearance distances becomes very important. Creepage is the shortest path that a current would have to take if it ‘crept’ along all available surfaces to reach the vulnerable part, while clearance is the shortest path to the vulnerable part through air (metal parts encountered along the way counting for zero distance regardless of their dimensions).

Clearance is the easiest to deal with, because the breakdown voltage of air is usually around 1kV/mm. So as long as the distance from the tip of the ESD gun to the vulnerable part is at least 8mm (preferably 10 or 12mm to give a design margin) an ESD spark can’t occur.

Creepage is more difficult, because the surfaces of plastics are always contaminated with mould-release chemicals, fingerprints, dust, etc., which attract moisture from the air and form a variable conductivity surface. Sparks from the tips of ESD guns are often seen to follow a random path over the surfaces of plastic enclosures, displays, keyboards, etc., sometimes for as long as 50mm as they follow the path of least resistance through the dirt on the surface of the plastic, eventually ending on a metal part. (Painted metal surfaces often show similar long random spark tracking, usually leading to a pinhole defect in the paint that it takes a microscope to see.) So it is very difficult to specify an adequate creepage distance which will protect from an ESD test, although more than 50mm is probably adequate except for polluted or wet environments.

Figure 6B shows a combined creepage and clearance design issue. A joint in a plastic enclosure could allow an ESD spark to travel along the surfaces of the plastic, then through the air inside the enclosure to terminate on a vulnerable PCB track. Figure 6B shows that it is usually a good idea not to line up PCBs with seams or joints in their plastic enclosure.

Figure 6B ESD sparks can enter via joints in plastic enclosures

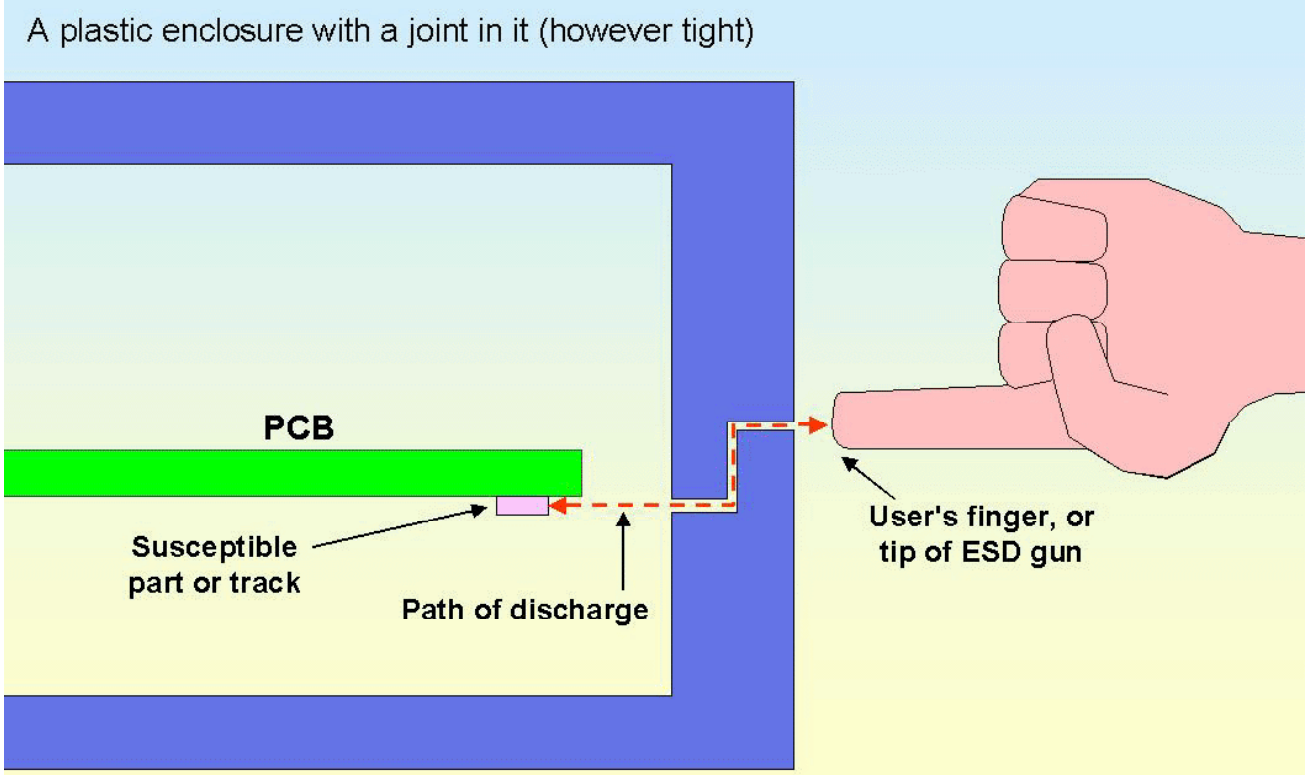

LCD displays, membrane panels, and tactile rubber keypads can be very good at preventing ESD if a few basic precautions are taken. Although their surfaces are ESD-proof at least to 15kV, they can have problems at their edges. ESD sparks can track along the dirt on their insulating surfaces, and go around their edges to reach vulnerable internal tracks.

LCDs often dealt with this problem by using large bezels which prevented fingers from getting too close to their edges. Insulating sealant and similar materials are now more likely to be used these days. Another method is to surround the LCD panel with a metal bracket that 'catches' the spark before it gets to any sensitive parts, but then something has to be done to remove the charge from the metal surround without it discharging itself into some sensitive part.

Membrane keypads and panels have internal conductive tracks, sandwiched between glued layers of plastic. If these tracks get too close to the edge of the panel, and if the glue has an airgap in it, sparks can track from the front surface (where the air discharge tip is applied), around the edge, through the void in the glue, and into the internal track, giving a false keypress if nothing worse. So whilst all attempts should be made to ensure there are no voids in the glue, it is still best to keep internal tracks at least 12mm from the edge of the panel (much more if possible).

Tactile rubber keypads also suffer from sparks that track through their surface dirt around the edges of their rubber mouldings and into the vulnerable keypad tracks behind. Unlike membrane panels, they usually don't have the benefit of glue to provide insulation, so it is important to extend the rubber edges of the tactile key moulding for far enough out, whilst keeping the tracks on the underlying PCB far enough in, so that any sparks have too far to go.

When a plastic enclosure has an internal shielding coating applied to meet RF emissions or immunity requirements, this can compromise dielectric isolation measures. For the conductive

layers to make a connection across enclosure seams they must extend at least a little way into the seams, and may even be fitted with a conductive gasket. This can compromise the creepage and clearance distances that had existed on the unshielded version, and ESD tests on such enclosures often find that when the tip of the ESD gun gets anywhere near seams and joints in the enclosure, a spark flies from the tip and disappears into the seam or joint to meet the internal shielding layer. In this situation it is usually very hard to achieve dielectric isolation for the whole enclosure, and the shielding method described next may have to be employed instead. Where possible it is a good idea to plan ahead so that plastic cases are designed to allow internal shielding to be added later without compromising the dielectric isolation ESD protection. This can be difficult to achieve, especially on small products.

6.1.5 Shielding

Shielding attempts to divert the (very large) ESD currents away from internal circuitry. In general it is not as good as dielectric isolation because it exposes all the external conductive connections (and possibly internal circuits too) to indirect ESD injection via 'ground lift'.

When an ESD spark occurs to a metal enclosure, for the first few microseconds the enclosure will be at a much higher voltage than any protective earth it is connected to. This local 'ground lift' decays as the charge on the enclosure leaks away through the inductance of any protective earth connections (usually several tens of μH). Where an enclosure is not connected to protective earth the charge on the shielded enclosure leaks away slowly through ionisation currents in the air around it, conduction through humid fabrics, and similar mechanisms.

During the beginning of a ground-lift event, internal circuits may still be at their previous voltages and sparking may occur between the enclosure and internal parts. This is known as 'secondary arcing' and it can be as bad for semiconductors and signals as the original ESD spark.

As long as the internal circuitry can cope with the sudden change in potential of its enclosure they don't care whether they are at protective earth potential, or 8kV relative to it. So it is quite practical to make battery powered or double-insulated products withstand ESD events (even though they are not connected to earth), although their isolation from earth may create greater problems for their external interconnections (discussed later).

So, on its own, having an external shield is not enough. One solution to secondary arcing is to bond the internal circuits to the enclosure shield, using connections that have a low enough inductance to maintain a low voltage between circuit and shield during an ESD event. These connections are often direct bonds from the 0V planes in the circuit boards to the enclosure, but they could be capacitive connections instead. Another solution is isolate the internal circuits from the enclosure using materials (or air spaces) that will not break down due to the voltage overstress.

Having dealt with secondary arcing we need to address the problem of transient current injection into the internal circuits.

The circuits within a shielded enclosure are all exposed to the internal fields created by the momentary ground-lift of the enclosure, before the internal potentials have had time to equalise. Because all the different PCB traces and components (even lead frames and bond wires) have different amounts of stray capacitance to the enclosure they each experience different amounts of injected stray current. Because the risetime of the ESD event is so fast and contains such a high frequency content, even very small stray capacitances can inject quite large currents. These different currents can create differential signals which can upset circuit operation. Figure 6C tries to show how this problem occurs.

Figure 6C ESD problems with shielded enclosures

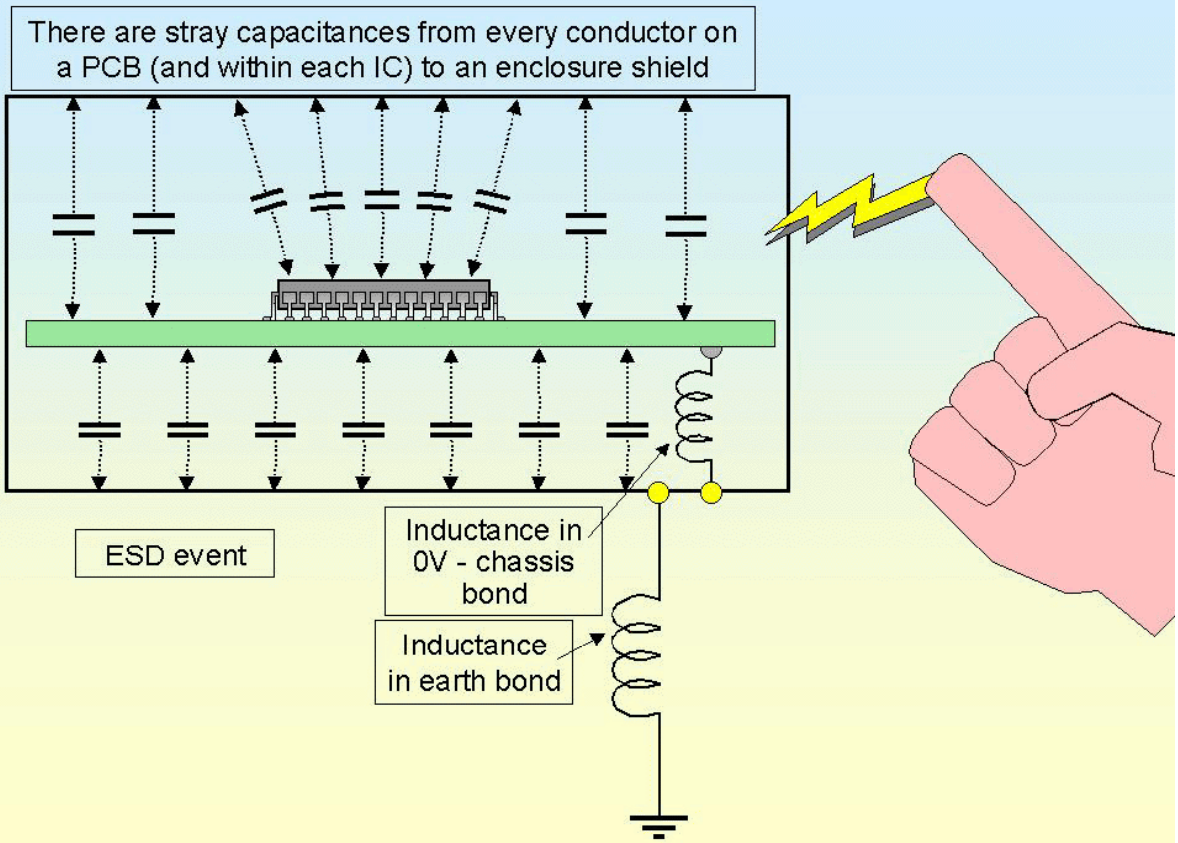

Solutions to the transient current injection include:

- Bonding the PCB ground plane(s) to the enclosure at very frequent intervals, to reduce the bonding inductance so that the ESD voltages equalise as quickly as possible. Quicker equalisation = lower internal voltage differences = lower values of transient current injection.
- Shielding sensitive circuits or ICs with PCB-mounted metal boxes bonded to the local PCB ground plane. These may be thought of as intercepting the stray capacitances from the components and PCB traces to the enclosure, diverting their transient currents into the ground plane where they will do less harm.

The problem illustrated by Figure 6C is often made worse where an enclosure contains a number of interconnected PCBs, as it is so difficult to ensure that during an ESD event they all charge up to the enclosure voltage at the same rate. If one circuit board has a low inductance connection to the enclosure, whilst another has a high inductance connection, then there can be a substantial transient voltage difference between them. This would inject a pulse of current into the boards' interconnections, causing cause signal corruption if not actual damage. So it is always a good idea to bond the 0V planes in different boards together using a number of conductors (one reference conductor for every one or two signal conductors may not be excessive) to help prevent large internal voltage differences.

Where galvanic isolation of circuits from their enclosure is used to protect sensitive signals from differences in earth potentials, usually where long external cables are involved, the isolation is usually only needed at mains frequencies so it is often possible to capacitively bond across the galvanic barrier (with suitably rated capacitors). This equalises transient voltages during an ESD

event without compromising functionality. Where this method is unsuitable, an internal shield over the isolated circuitry, bonded to its local isolated reference plane, may be required.

Even tiny gaps or joints in enclosure shields are weak spots, because they divert the very large fast currents from the ESD spark as they flow around the enclosure, causing locally intense pulses of electric and magnetic fields to be emitted through the shield and into the enclosure. Looked at from the frequency domain, we would say instead that the high frequency components of the ESD event find gaps and joints useful as slot antennae, radiating into the enclosure. So it is important to keep all gaps and joints in shields to a minimum size, as described in Part 4 of this series. Even if they are very small, sensitive circuitry should be kept well away from them.

6.1.6 Adding impedance to signal lines

Some interface devices (such as some RS232 ICs) are claimed to withstand EN61000-4-2 testing directly on their RX and TX connections. But for the vast majority of semiconductors it is by far the best to protect them from the direct effects of ESD by using the dielectric isolation (insulation) or shielding techniques, described above.

But if for some reason circuit conductors really *must* be exposed to ESD sparks: using a series resistor or choke (as shown in Figure 6D) may allow the ESD injection into the circuit to be handled by an IC's own internal protection devices. This is only likely to be possible for interface or 'glue logic' devices, and can't be generally recommended for VLSI devices such as microprocessors not intended to connect directly to external cables.

Figure 6D

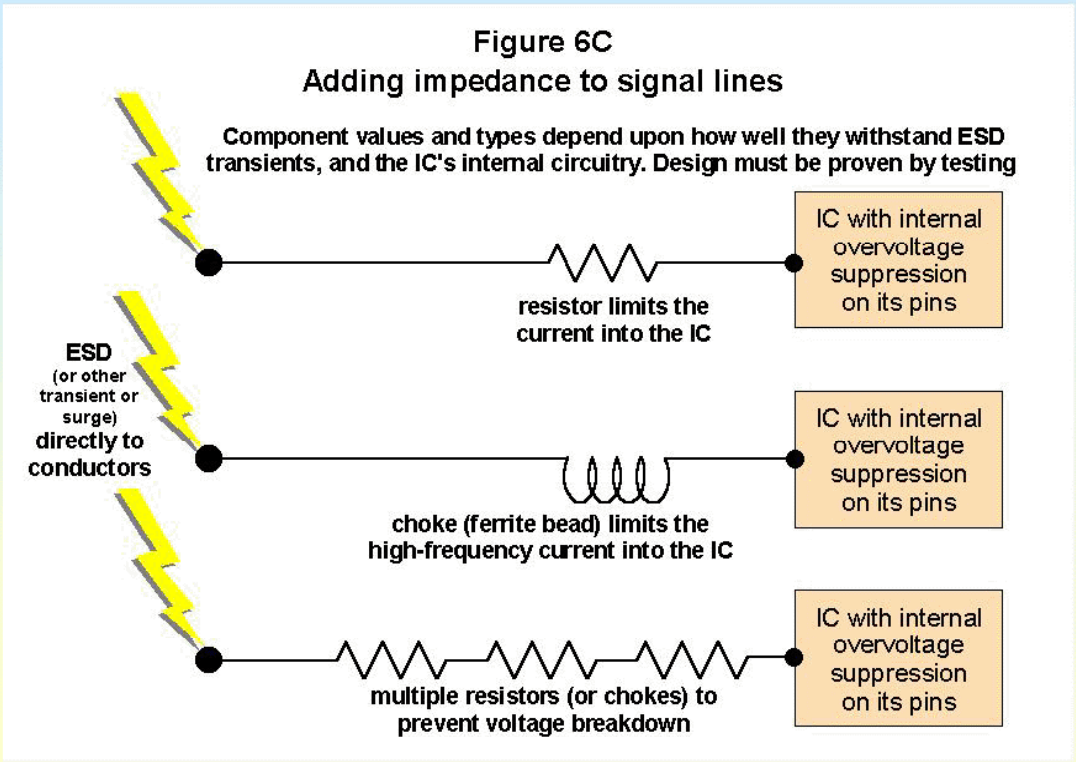

Most resistors or chokes aren't rated for ESD voltages, but may cope if they are physically large, have enough thermal inertia, and don't arc-over themselves. Since there is very little energy in a

personnel ESD event, a large thermal inertia will prevent the resistor from suffering damage, although the same may not be able to be said of some machine ESD events involving large stored charge. It may be necessary to put a number of resistors or ferrites in series so that they share the ESD voltage and don't arc-over or break down (close to each other and arranged in a straight line, to get maximum advantage from this technique).

Since resistor and choke manufacturers don't specify their products' specifications on EN 61000-4-2 tests, it is best to check a few in a representative circuit and PCB layout. It may be found that different values may be needed for different ICs. The chokes used should be carefully chosen to provide a high impedance over the range of frequencies encompassed by the ESD event. Always perform a number of tests, to make sure that 'soft' failures aren't happening.

Mask shrinks by semiconductor manufacturers can make such a protective circuit design inadequate (just as it can invalidate all other EMC measures), so it is always best to have an arrangement with your semiconductor suppliers that they will warn you in plenty of time of any impending mask shrinks, so you can purchase some buffer stock to continue production whilst you evaluate all the EMC effects of the mask shrink, including any ESD protection.

The resistors or chokes may need to have such high values that they prevent the circuit from operating correctly. This is a particular problem for high-speed signals. One solution is to add discrete transient voltage suppressers, as described later. These will handle much larger voltage and current transients than most devices' internal protection measures, allowing series resistors or chokes to be reduced in impedance, or removed altogether.

However, a potentially serious problem for all methods that merely prevent fatal damage to semiconductors from direct ESD injection, is that they do not prevent momentary corruption of signals. Signal corruption can cause an ESD test to be failed, even though devices are undamaged. This is discussed later.

6.1.7 Transient Voltage Suppressors (TVSs)

As asserted above, it is always best not to allow sparks to get into circuit conductors. But if, for some reason, circuit conductors really *must* be exposed to ESD sparks: it may be possible to arrange for the associated ICs to survive by fitting series resistors or chokes as described above. An alternative technique that is widely advertised by TVS manufacturers is to fit suitably-rated discrete TVSs between the vulnerable conductors and the local 0V plane, a shown in Figure 6D. Being very much higher rated than an IC's internal protection devices, series resistance or chokes can have much lower impedance, and are often dispensed with altogether, which is sometimes much better for functionality.

Low inductance plane bonds are required for TVSs to protect ICs correctly, so the techniques described in Part 5 for connecting decoupling capacitors to 0V planes should be followed. Also, the PCB track to be protected must pass through the other terminal of the TVS. 'Spurring' a track from the signal path to a nearby TVS creates inductance which can prevent the TVS from protecting the IC. Since the TVS is required to bypass the powerful discharge current from the signal conductor into its local 0V plane, care must be taken that the ESD transient discharge current – now flowing in the 0V circuit and looking for a route back to the external earth – does not cause problems elsewhere in the product.

TVS components are now available in a wide range of voltage and power ratings, in leaded and SMD styles. They are also available as space-saving PCB-mounted arrays, and also fitted into connectors where they shunt the spark current into the metal body of the connector (which therefore needs to have a good high-frequency bond to the metal enclosure or the PCB's 0V plane). Some manufacturers make very thin flexible circuits which fit over the solder pins of common connector types, allowing an easy TVS retrofit to every pin of the connector.

A big advantage of discrete TVSs over filtering is that they do not compromise high-speed signals as much. Most of them are based on zener technology so they do have some capacitance, and for very high speed signals the choice of an adequately low-capacitance TVS which still has the desired ratings can still be quite limited. Where extremely low capacitance ESD protection is needed,

reverse-biased diodes between the signal conductor and the 0V and power planes, as shown in Figure 6E, can be effective. Reverse-biased diodes have a significantly lower capacitance than zener diodes. The diodes need to be able to handle the ESD currents when forward-biased, and the local power plane needs sufficient high-frequency decoupling capacitance so that its voltage does not rise too much when it absorbs the ESD charge. Where the leakage currents of TVSs or diodes are the problem, rather than capacitance, it is possible to use FETs instead to get nanoamp leakage.

Figure 6E Adding transient voltage suppression

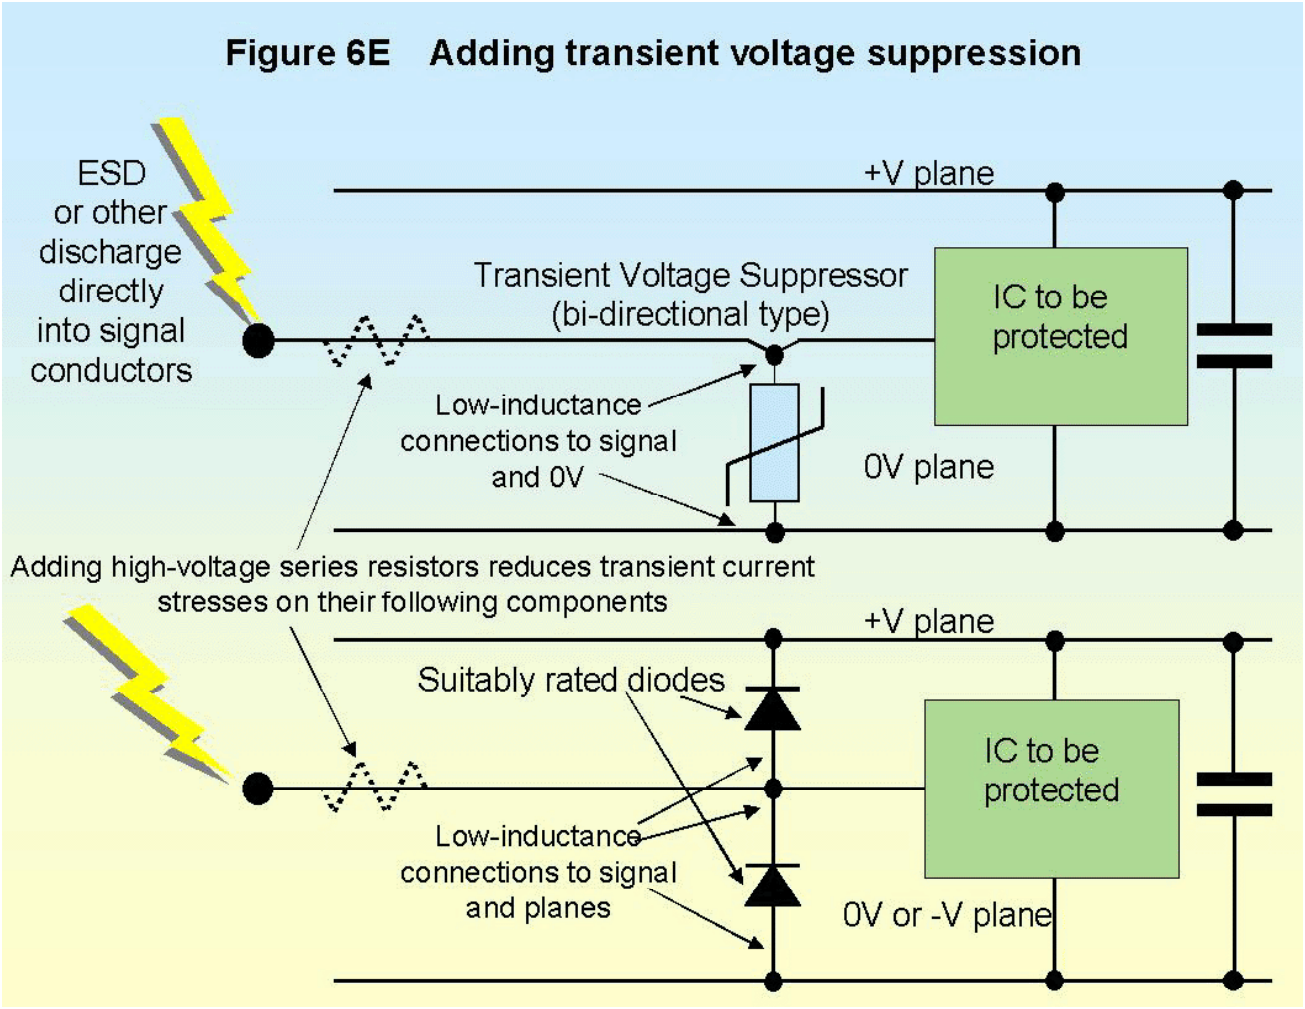

Unfortunately, although TVSs are often advertised as a complete solution to problems of EN 61000-4-2 ESD sparks getting into exposed circuitry, this is not the case in practice. Their problem, which they share with the ICs internal protection devices and the reverse-biased diodes above, is that they cannot prevent the ESD transient from corrupting the signals on the conductors they are protecting.

Techniques additional to the use of TVSs are thus required to prevent signal corruption, and these are described in a later section.

6.1.8 Low-pass filtering the signal lines

Following a series resistor or choke with a capacitor to the local 0V plane (as shown by Figure 6F) can provide excellent protection from direct ESD injection to a conductor. As well as helping prevent actual damage, it can reduce transient voltages to such low levels that excessive signal corruption is prevented. Unfortunately, this approach has its limitations for high-speed signals.

Figure 6F

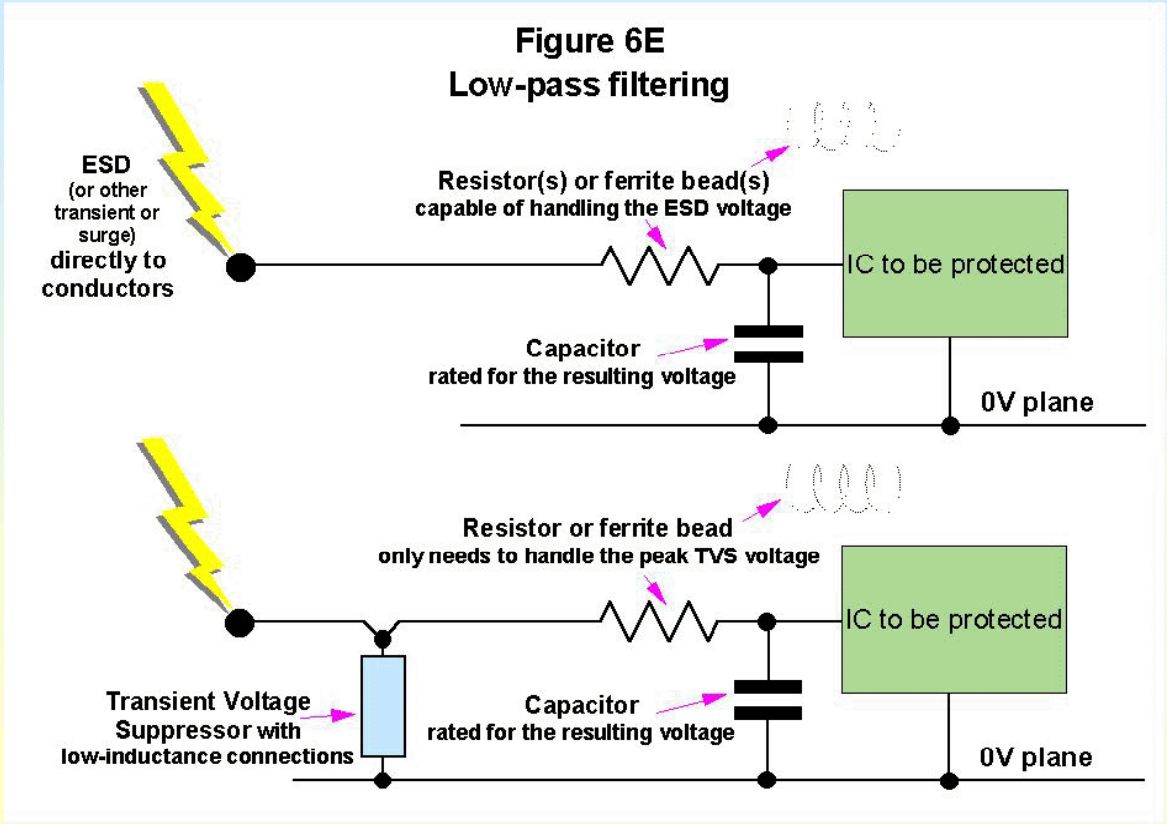

The capacitor in the circuit shares the charge with the capacitance in the source of the ESD event. Since the source of a personnel ESD event has a capacitance of 150pF (although some older human-body models used by older test standards use 330pF) a 1nF capacitor in the circuit would reduce an 8kV voltage to about 1000V, 10nF would reduce it to about 120V, and 100nF would reduce it to about 12V.

Adding a series impedance such as a resistor or choke between the ESD event and the capacitor reduces the peak currents, which if left unchecked could be tens of amps and could physically overstress the capacitor and might also cause secondary problems due to the intense magnetic fields resulting. A 1kΩ series resistor would limit the peak current of an 8kV discharge to 8A, 10kΩ to 800mA, and 100kΩ to 80mA.

So we can see that a 100kΩ resistor followed by a 100nF capacitor would tame a personnel ESD event very considerably, where the wanted signals concerned were slow enough not to mind a filter with a 1 second time-constant. Of course, the series resistor or choke would need to be able to withstand 8kV without damage, and this rules out the use of ordinary resistors (unless a number are used in series). Surge-rated resistors are available, but they are larger and more expensive than ordinary resistors. Some manufacturers of ferrite beads test their products to check that they are not damaged when ESD events of up to 25kV are applied across them, but I am not sure whether they allow them to flash-over during the test. Surface-mounted resistors and ferrite beads may be relied upon to flash-over at quite low voltages, due to the small spacings between their terminations. Unfortunately a number of SMDs in series may not voltage-share very well due to variations in surface contamination, so when using this technique rather than a guaranteed non-flash-over (probably leaded) component – it may be best to design-in a very large margin (e.g. use twice as many components in series as appears necessary).

When the load impedance is much greater than the series impedance in the protection circuit, the high voltages due to the ESD event can last a long time and the protection is achieved by charge balancing and peak current limitation. But when the load impedance of the IC or circuit to be protected is much lower than the impedance of the series resistor of ferrite in the RC circuit, the ESD event lasts a much short time and so the RC circuit also acts as a low-pass filter, giving even greater attenuation of the ESD event and hence greater protection.

Applying the 'low-pass filter' protection circuit after a TVS (as shown in figure 6E) can be useful where load impedances are less than the impedance of the series element (resistor or ferrite choke): the TVS reduces the ESD's 8kV (or whatever) to a few tens of volts, and the low-pass filter then attenuates it to less than the desired level. Of course, the peak currents into the TVS will be very large (unless limited by a series resistor) and these may cause problems due to their local magnetic fields, but the resistor or ferrite choke following the TVS will not have to be rated for the ESD voltage and can be an ordinary component.

Where fast signals are involved, filtering on individual lines may not be able to achieve useful ESD protection without negatively affecting the wanted signal, although common-mode filtering may well be possible (see later). Transient voltage suppressers may need to be used instead (see above), since they add only little capacitive loading, along with communication protocols or software techniques (see below), but it is still best not to let the spark get into the conductors in the first place.

6.1.9 Common-mode filtering at connectors

Because external signal and power cables are initially at their previous voltages, a 'ground lift' event due to a metal enclosure being struck by an ESD spark makes a transient high voltage appear across circuitry which interfaces with external cables.

All things being relative, it is as if the enclosure was at earth potential and all its external cables had been suddenly raised to a high voltage. For example, if an enclosure was suddenly charged to +4kV by an ESD spark, it is as if the enclosure remained at earth potential but all the external cables had been suddenly charged to -4kV instead.

Although filters and/or TVSs may be used on each individual power or signal line, to help protect the electronics interfacing with the external cables (as described in the previous two sections), filtering each signal line independently is usually not compatible with high-speed signals, whereas TVSs can't prevent signal corruption (see later).

Because the transient voltage experienced by the interface circuitry is identical for all the conductors in a given cable it is a common-mode (CM) transient, and CM suppression techniques may be used to help suppress the transient voltage without compromising high-speed signals.

Figure 6G shows a common-mode choke used for this purpose. It needs as many windings as there are conductors in the cable. Surface-mounted and leaded common-mode chokes are available with up to eight windings. If the CM choke has a high enough value of common-mode impedance over the frequencies of concern it may be able to prevent an ESD transient from damaging the interface electronics. In situations where signal corruption is the problem, a CM choke can greatly reduce the signal perturbation.

Figure 6G

Figure 6F
Common-mode filtering

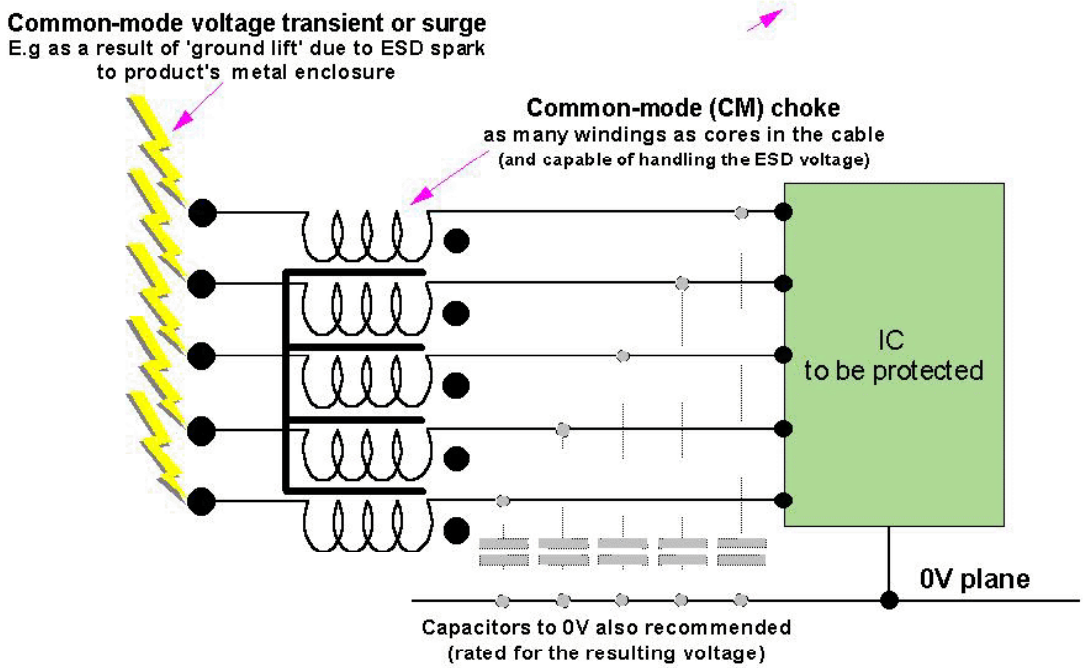

When a CM choke is used with filter capacitors even greater suppression can be achieved, although this is not as suitable for high-speed signals unless quite small capacitor values are used (say between 10pF and 1nF, depending on the application).

Where a large number of conductors must be accommodated specially-wound chokes may be required. A favourite technique is to pass the cable through a soft-ferrite cylinder or toroid. A single pass through a typical 32mm long ferrite cylinder creates a common-mode impedance of around 250Ω at 100MHz, not as much as some PCB-mounted components, but often enough to make a useful improvement to an EMC problem.

Where higher common-mode impedance is required, two or more turns of the cable could be made through the ferrite, to increase the impedance. Unfortunately, multiple turns also reduces the frequency at which the peak impedance is achieved, so beyond a certain number of turns no further benefit (at a given frequency) may be achieved by adding more turns. We would normally expect that the impedance presented by a choke would increase according to the square of the number of turns, but this is not always the case for soft ferrites, as shown by the following example (kindly provided by Alan Keenan of Steward Inc.). Tests on a particular soft-ferrite core found that with two turns its Z peaked at 692Ω at 322MHz; with three turns its Z peaked higher at 809Ω but at the lower frequency of 152MHz; and with four turns wound on it - its Z peaked even higher at 1300Ω but at the even lower frequency of 108MHz.

Multiple ferrites strung along a cable increase impedance proportionally to the number of ferrites, without reducing the frequency of the impedance peak, but can look a bit Heath-Robinson. Cables passed through ferrites should always finish on the other side of the ferrite from the side they entered, so that the ferrites are like beads strung on a length of string, for best high-frequency

performance. Any stray capacitance from one end of the ferrite (or chain of ferrites) to the other will compromise the high-frequency performance, making the routing of the cable or track very important.

6.1.10 Galvanic isolation techniques for ESD

Galvanically isolated external connections help a great deal with making products immune to ESD (and also to conducted transients and surges). PCB-mounted opto-isolators and transformers may be used, although many types will not have sufficient creepage and clearance and voltage withstand for ESD, and many types will have internal parasitic capacitances which may allow excessive transient currents to flow. Isolating mains transformers that meet their appropriate safety standards often seem quite adequate for protecting power supplies from personnel ESD, despite their shortcomings as ESD barriers.

The very best galvanic isolation for external signal communications is achieved with fibre-optic, wireless, or infra-red techniques, since these do not involve any conductors with dissimilar voltages being anywhere near the product concerned. Of course, the transmitting and receiving modules concerned can be very sensitive, and are often best fitted with their own local shielding, but even quite large PCB-mounted shields are now available taped and reeled for automatic placement and so can be much more cost-effective than when such items required manual assembly. Fibre optic cables sometimes use metal strengtheners, armour, or metal vapour barriers, and these can compromise the creepage and clearance distances required for ESD protection so should be stripped well back from any connections with products, where it is not possible to avoid the use of such cables entirely.

Five years ago I began expressing the opinion that by 2005 people who used copper cables to connect signals between items of equipment would be thought rather old-fashioned, and I still believe that this timescale is correct.

6.1.11 Dealing with signal corruption

When ESD testing to harmonised EMC standards using the EN 61000-4-2 test method, it is normal for the requirement to be that the operational state of the product being tested (and its displays and stored data) is exactly the same after the test as before. This is usually also the case for other transient tests such as Fast Transient Bursts (EN 61000-4-4), Surge (EN 61000-4-5), and the various automotive transients described by ISO 7637.

Using an IC's internal protection or a TVS to protect against ESD (or other types of conducted transient or surge) does nothing to prevent signal or data corruption, and can result in the product's operational state, displays, or memory being different after the test, leading to test failure. Additional techniques are often needed to prevent signal corruption.

For ordinary control lines such as keyboards, 'debouncing' techniques (whether hardware or software, such as are routinely used to debounce mechanical contacts) can work very well. Low-pass filtering (see above) is a good old-fashioned 'debouncing' technique often used after a TVS.

High-speed data can't use powerful debouncing techniques, and some type of error-protecting protocol is usually required, although of course this slows down the data rate too.

Where a single ESD event causes a single bit error in a serial communication, the use of simple parity check or Hamming coding can make a big improvement in immunity. However, these techniques will be less effective when the interference comes in bursts, as is usual for non-ESD transients, or when an ESD event causes secondary arcing. Some very comprehensive error-correcting protocols are now commercially available (see section 1.4.7), and it is usually much more cost-effective and time-saving to purchase the necessary chips or software licenses than it is to try to develop your own protocols. Unlike simple techniques such as parity checking or Hamming coding most such protocols have a high resistance for burst errors, and have been field proven in a wide range of applications and situations.

Momentary corruption of analogue signals is sometimes acceptable, when all that results is a brief click in a headphone or momentary flicker of a meter needle. But where decision thresholds may be crossed, or where the analogue signal is put through an averaging routine or stored in a memory after digitisation, even a brief error in an analogue signal may be unacceptable.

Since analogue signals can't easily use error-correcting protocols, debouncing techniques (such as a low-pass filter) are often used. Where the analogue signal must use high frequencies, it becomes very difficult to filter the transient event from the signal, the more so as the accuracy required (equivalent to number of bits of resolution) is increased, so dielectric isolation or shielding techniques are necessary after all, to prevent the ESD transient from getting into the signal conductor.

6.2 Dips, sags, brownouts, swells, dropouts, interruptions and power outages

According to some experts, the effects of poor supply quality on electronic equipment is one of the most significant causes of downtime and financial loss world-wide. Dips, sags, brownouts, swells, voltage variations, dropouts and interruptions are the main causes of poor supply quality. (In some areas waveform distortion is also becoming an important issue.)

EN 61000-4-11 is the basic test standard for dips, sags, brownouts, swells, voltage variations, dropouts and interruptions in the AC mains supply, but is not harmonised under the EMC directive so cannot be used when self-declaring conformity to standards. For self-declaration to standards it is necessary to use the harmonised generic or product standards - which now generally call-up the use of EN 61000-4-11 as the basic test method. It is always the generic or product standards that set the actual test limits, and they are often a lot less comprehensive than the limits in EN 61000-4-11.

When using the Technical Construction File route to conformity it is often permissible to use EN 61000-4-11 directly.

Of course there is no reason why a manufacturer should not apply EN 61000-4-11 to a product in any case, in addition to the requirements of any harmonised product or generic standards, and this may sometimes discover reliability problems with products that would otherwise have been overlooked.

Dips are short-term reductions in supply voltage caused by load switching and fault clearance in the AC supply network. They can also be caused by switching between the mains and alternative supplies in uninterruptible power supplies or emergency power back-up systems. Examples of dips: 30% dip for 10ms, 60% dip for 100ms. Figure 6H shows a 40% dip for 20ms (one mains cycle). A dip of 40% is equivalent to a reduction in supply voltage to 60% of its nominal value.

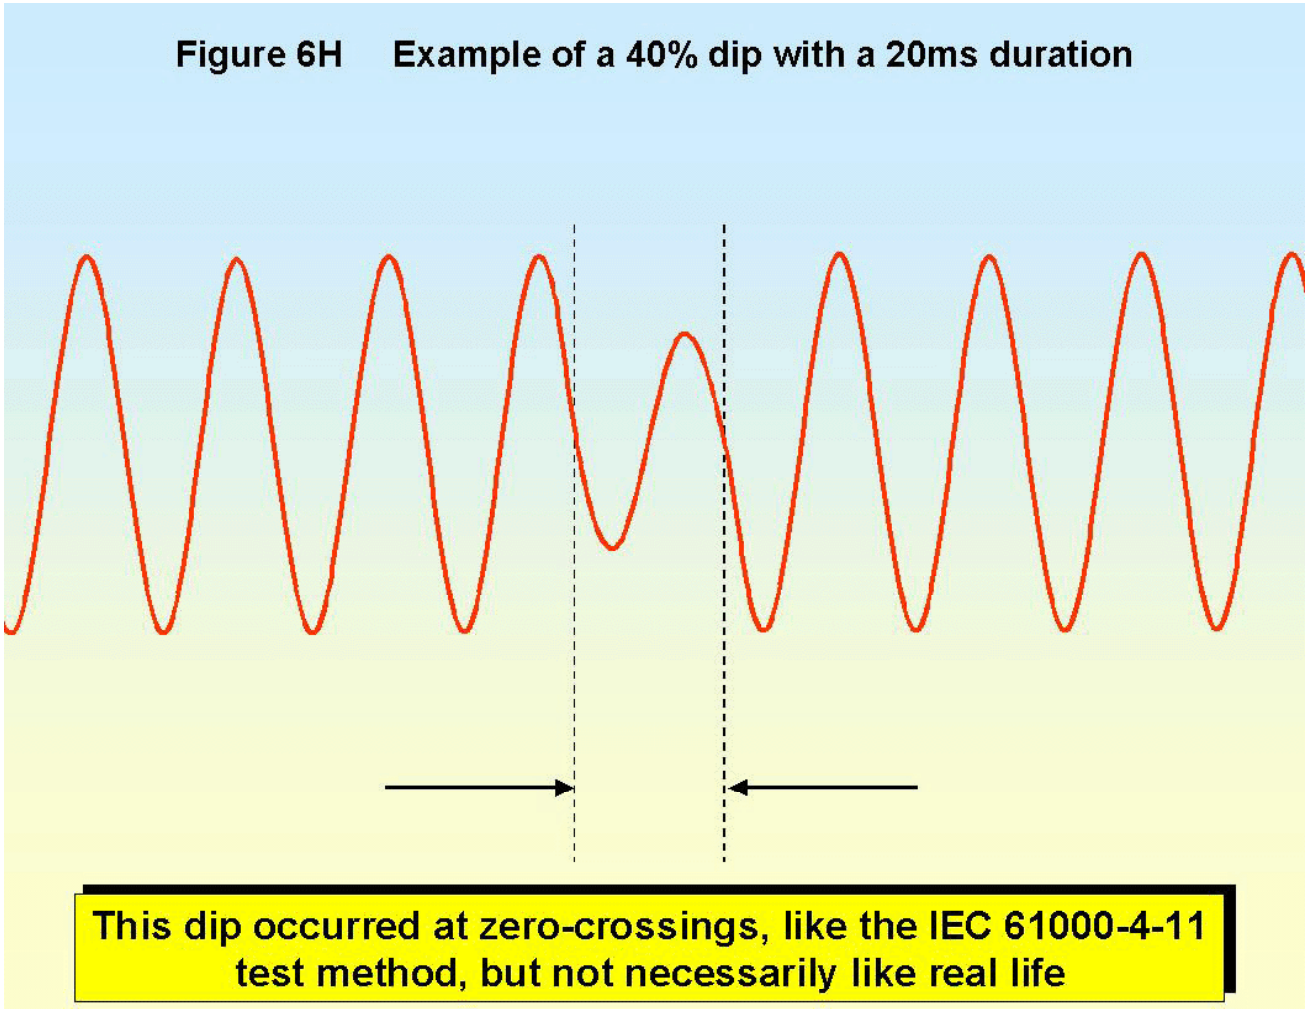

EN / IEC 61000-4-11 refers to a Unipede study of dips which covered public mains supplies.

Dip depth	Number of dips/yr with given durations			
	10-100ms	100-500ms	0.5-1s	1-3s
10-30%	61	66	12	6
30-60%	8	36	4	1
60-100%	2	17	3	2

Dips may not occur in isolation - sometimes a fast sequence of dips occurs. Most immunity tests apply dips abruptly, starting and finishing at a zero crossing, but in practice they can have gentle rates of change and start and stop at any point in the mains cycle.

EN61000-4-11 recommends testing with 30% and 60% dips (70% and 40% of nominal voltage respectively) for 0.5, 1, 5, 10, 25 and 50 cycles of the supply waveform in each case. Although the generic immunity standard for domestic, commercial, and light industrial environments EN 50082-1:1997 (which will be replaced by EN 61000-6-1 sometime around 2004) uses EN 61000-4-11 as its basic test method, it merely requires testing with 30% dips for 10ms and 60% dips for 100ms.

Likewise, EN 61000-6-2:1999 (which replaces EN 50082-2:1995 in April 2002) also uses EN 61000-4-11 as its basic test method but only requires testing with 30% dips for 10ms, 60% dips for 100ms and 1 second. Many of the other harmonised standards under the EMC and RTTE directives specify

dip test % and durations which are different from the above generics and from each other, and which are not as comprehensive as EN 61000-4-11.

Because no high frequencies are involved it is fairly easy to construct your own supply dips tester using a pair of solid-state zero-crossing relays, a variable transformer and a timing generator, as shown by Figure 6J. Safety issues must not be neglected ! A test laboratory would be more likely to use a programmable mains synthesiser with enough short-term capacity to handle the product's inrush current.

Figure 6J

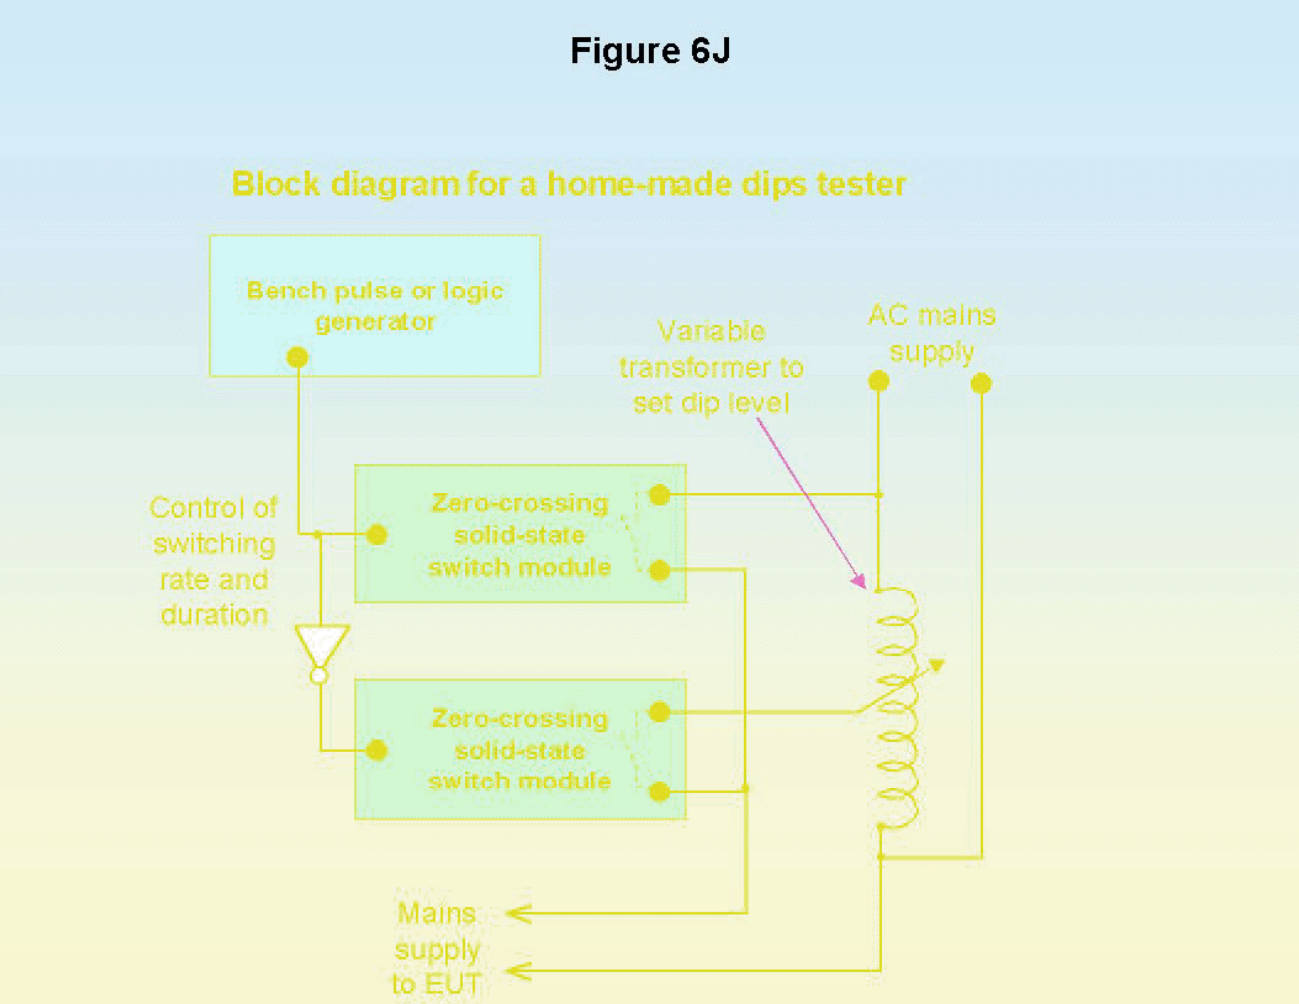

Sags, swells, and brownouts are slowly-varying changes in voltage, sometimes over periods of hours. They are sometimes called 'voltage variations' and some people don't consider them to have anything to do with EMC. A brownout is another name for a sag, and is a term most often heard in the USA and Canada. Sags and brownouts can go right down to zero volts. Swells are simply slow increases in voltage, as opposed to surges and transients, which are fast increases.

Typical mains supplies are specified at ±6% voltage tolerance (sometimes ±10%) and it is not unusual to experience these ranges daily. Safety tests apply ±15% for long periods of time, so voltages of these levels are clearly not impossible. It is often claimed that European mains supplies don't suffer from the traditional 'brownout' where the voltage falls by a large amount for minutes or even hours – but 60% nominal voltage for an hour or two every weekday is known to occur in parts of Spain – and 50% nominal for 8 hours has been measured by the author in the UK.

A sagging supply voltage is a particular problem for all AC motors and some types of DC motor because they can stall, overheat, and damage their insulation. This can lead to an increased risk of fire, fumes and electric shock, never mind the damage to the equipment. In industrial plant AC

motors are usually protected by undervoltage trips which are usually incorporated with other protective functions into motor control circuit breakers devices (MCCBs).

EN 61000-4-11:1994 suggests testing with 40% supply voltage for 1 second and also 0% supply for 1 second with the supply voltage ramped down and then up again over a period of 2 seconds in each case. It does not suggest testing for swells despite the fact that these can damage surge suppression devices.

Surge suppression devices are often designed to operate as close to the maximum expected mains voltage as possible, to protect the equipment better. But their leakage at voltages above nominal can cause overheating and damage if the higher voltages last for more than a few seconds. Maybe EN 61000-4-11 assumes that safety testing would discover such problems.

Most (if not all) harmonised generic and product standards don't specify any testing at all for sags, brownouts or swells. Testing to EN 61000-4-11 may be a good way to improve product reliability in the field. It may also be a good idea to test for anticipated swells, although EN 61000-4-11 doesn't recommend any test levels.

Figure 6K shows the voltage versus time profile for a typical sag or brownout test to 50% nominal voltage (in this case). It can easily be done manually by simply using a variable transformer ('Variac'), although a test laboratory would be more likely to use a programmable mains synthesiser with enough short-term capacity to handle the product's inrush current.

Figure 6K A typical EN 61000-4-11 sag test

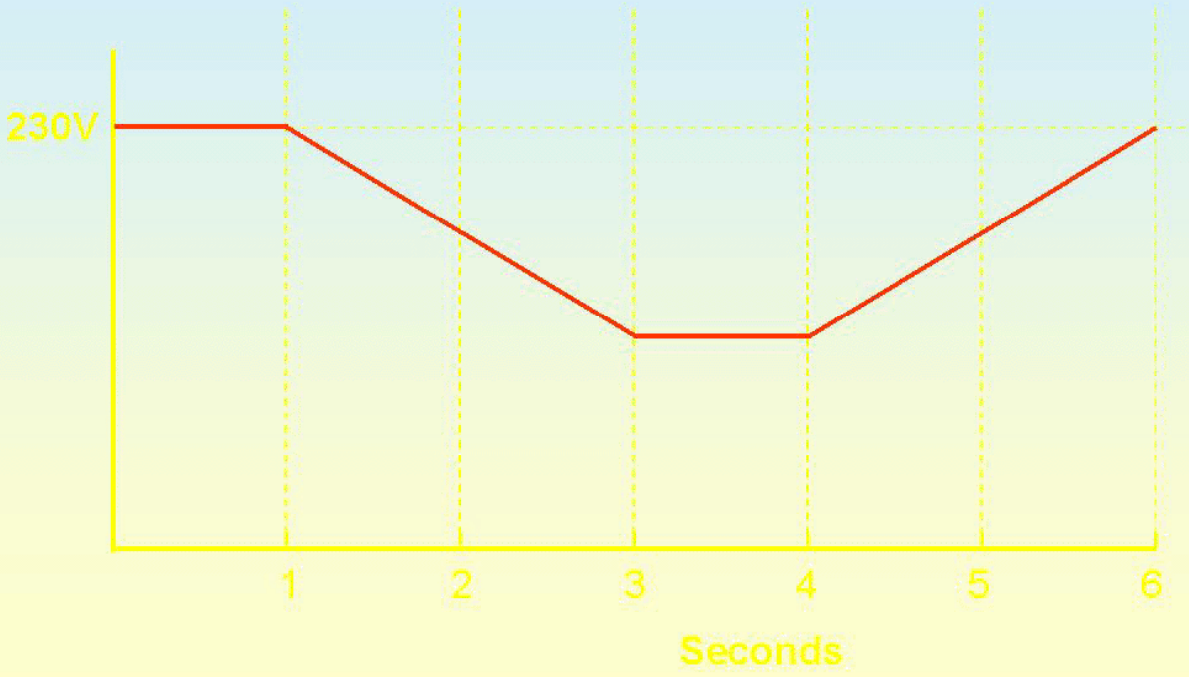

In real life brownouts can last for minutes at a time (8 hours at 50% voltage was once seen in the UK) so where motors or other vulnerable parts are concerned it may be best to test for a longer time than merely 1 second to ensure reliable operation in the field.

Swell testing can also use a variable transformer and manual control. A test laboratory would be more likely to use a programmable mains synthesiser with enough short-term capacity to handle the product's inrush current.

Dropouts are short-term 100% reductions in supply voltage (effectively 100% dips). Figure 6L shows a dropout 60ms long (3 mains cycles). Like dips, they are caused by load switching and fault clearance. Also like dips, they can be caused by switching between mains and alternative supply in uninterruptible power supplies or emergency power back-up systems.

Figure 6L Example of a dropout with a 60ms duration

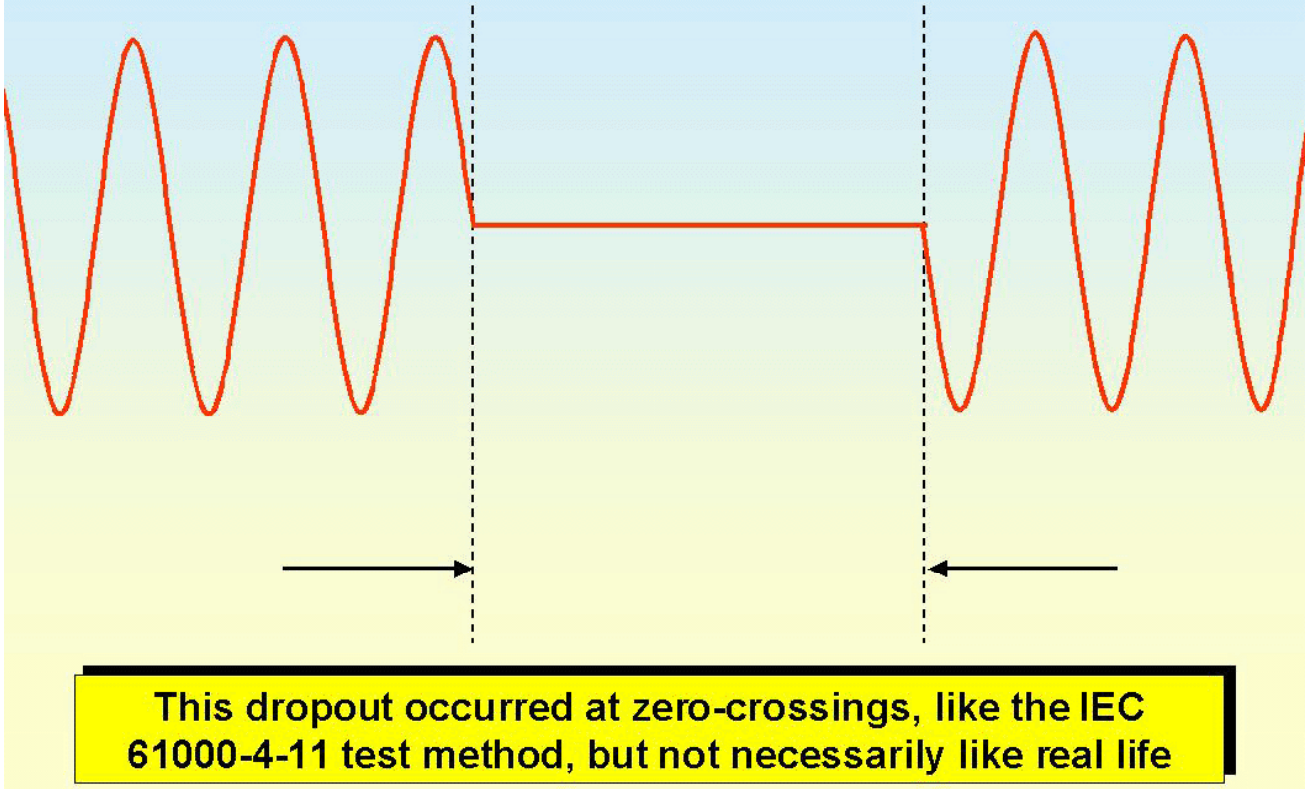

EN 61000-4-11 refers to a Unipede study of public mains supplies which recorded 12 dropouts per year with durations lasting between 100 and 500ms, and suggests testing with dropouts lasting for 0.5, 1, 5, 10 supply cycles.

But EN 50082-1:1997 (EN 61000-6-1 from 2004) and EN 61000-6-2:1999 (replaces EN 50082-2:1995 in April 2002) don't specify any dropout testing at all. It is clearly possible to declare EMC conformity without testing for dropouts, but your products might not be as reliable in the field as you would like. It is easy to make your own dropout tester using the method described for the dips tester but setting the alternative supply to zero.

Interruptions are supply voltage dropouts (= 100% dips) that last for a long time: seconds, minutes, or hours. They are really power outages. As for dips and interruptions, it is easy to make your own test equipment.

EN 61000-4-11 refers to a Unipede study of public mains supplies that showed...

Interruptions lasting: 0.5-1s 1-3s >3s

Number of interruptions/year: 24 5 0

(How did they manage to find no interruptions lasting more than 3 seconds? This doesn't agree with many people's experiences of UK mains supplies.)

EN 61000-4-11 suggests testing with interruptions lasting 0.5, 1, 5, 10, 25, and 50 supply cycles. But EN 50082-1:1997 (replaced by EN 61000-6-1 in 2004) and EN 61000-6-2:1999 (replaces EN 50082-2:1995 in April 2002) both specify testing with interruptions of 5 seconds duration only (= 250 cycles).

Equipment which has been designed to pass a 5 second power outage test might fail a 1 second outage test, so testing with all the durations suggested by EN 61000-4-11 might be a good idea to increase product reliability.

Waveform distortion is an increasing problem due to the proliferation of electronic loads on the power networks. 4% is getting to be the typical total harmonic distortion of European mains supplies, and this reach 8% in a few years. Most of the distortion consists of 'flat-topping', which means that the mains' peak voltage can be much lower than $\sqrt{2}$ times its RMS value.

In parts of China the mains can be an actual square wave. As Figure 6M shows, parts of Israel also suffer from significant waveform distortion and this is probably not untypical of a number of countries around the world. So although a meter might read 230Vrms, rectifier-capacitor power supplies could give as little as 75% output *before* any supply voltage tolerances are considered.

Figure 6M

Example of a domestic supply waveform in Israel, Oct 2000

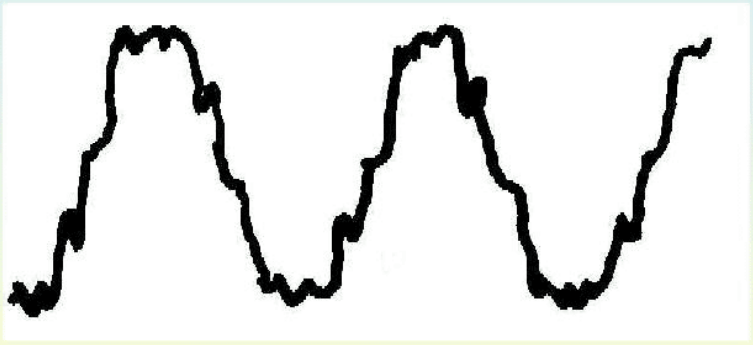

Privately-generated mains supplies can be much worse than the normal mains, although in the 'Internet Hotels' which now host so many of our web sites the mains supplies can be very much better (they need to be - a typical 'Cisco hotel' has an availability specification of 99.9999% and a power consumption of around 10MW).

An oil exploration platform had a 230V mains supply with a tolerance of ±100% caused by the effects of starting and stopping its huge drilling motor on its diesel generator. The voltage variations would last for several seconds each. Emergency 230/400V mains generators can have output quality much worse than a typical public mains supply, and much worse than the suggested test limits in EN 61000-4-11. So if your product is likely to be operated on a private mains supply, it is best to always find out what it is likely to be exposed to!

It is often thought that electro-mechanical components such as relays, contactors and solenoids are immune from EMC considerations. Certainly the designers of so-called hard-wired safety systems seem to believe this untruth. In fact dips, sags, brownouts, dropouts and interruptions can cause energised relays and contactors to 'drop out' and solenoids to lose power and move.

Although most designers would presumably consider the effects of a power outage (interruption) they would probably not consider a momentary dip in the supply that might cause some types of relays to drop out while other types remain energised. This situation is made more complicated because the 'pull-in' (must operate) and 'drop out' (minimum hold-in) voltages of relays and contactors rise significantly as they age.

Sometimes relays, contactors, solenoids, etc. are 'held in' at a reduced voltage, to save power consumption, and so might not pull back in again when the supply recovers. So we can see that 'relay logic' might not recover to its original state after the dip, dropout, sag, or interruption; and during the event itself they might operate in an unpredicted manner. Few designers of industrial automation seems to design or test for these problems, even though dips and dropouts are considered by experts to be a a major cause of lost production world-wide. (Few also seem aware of the fact that supply surges can arc unpredictably across mechanical contacts, also causing unpredictable operation of 'relay logic' systems.)

Figure 6N shows a typical problem for electronic circuits – the logic supply voltage going out of tolerance due to a momentary dip or dropout. The example is of a 40% dip for 40ms (two cycles of the supply), and shows how the unregulated DC rail droops to below the 5V regulator's drop-out voltage during the dip. The 5V supply to the logic ICs thus falls out of regulation and can fall outside of the tolerance band required for correct logical operation of the ICs.

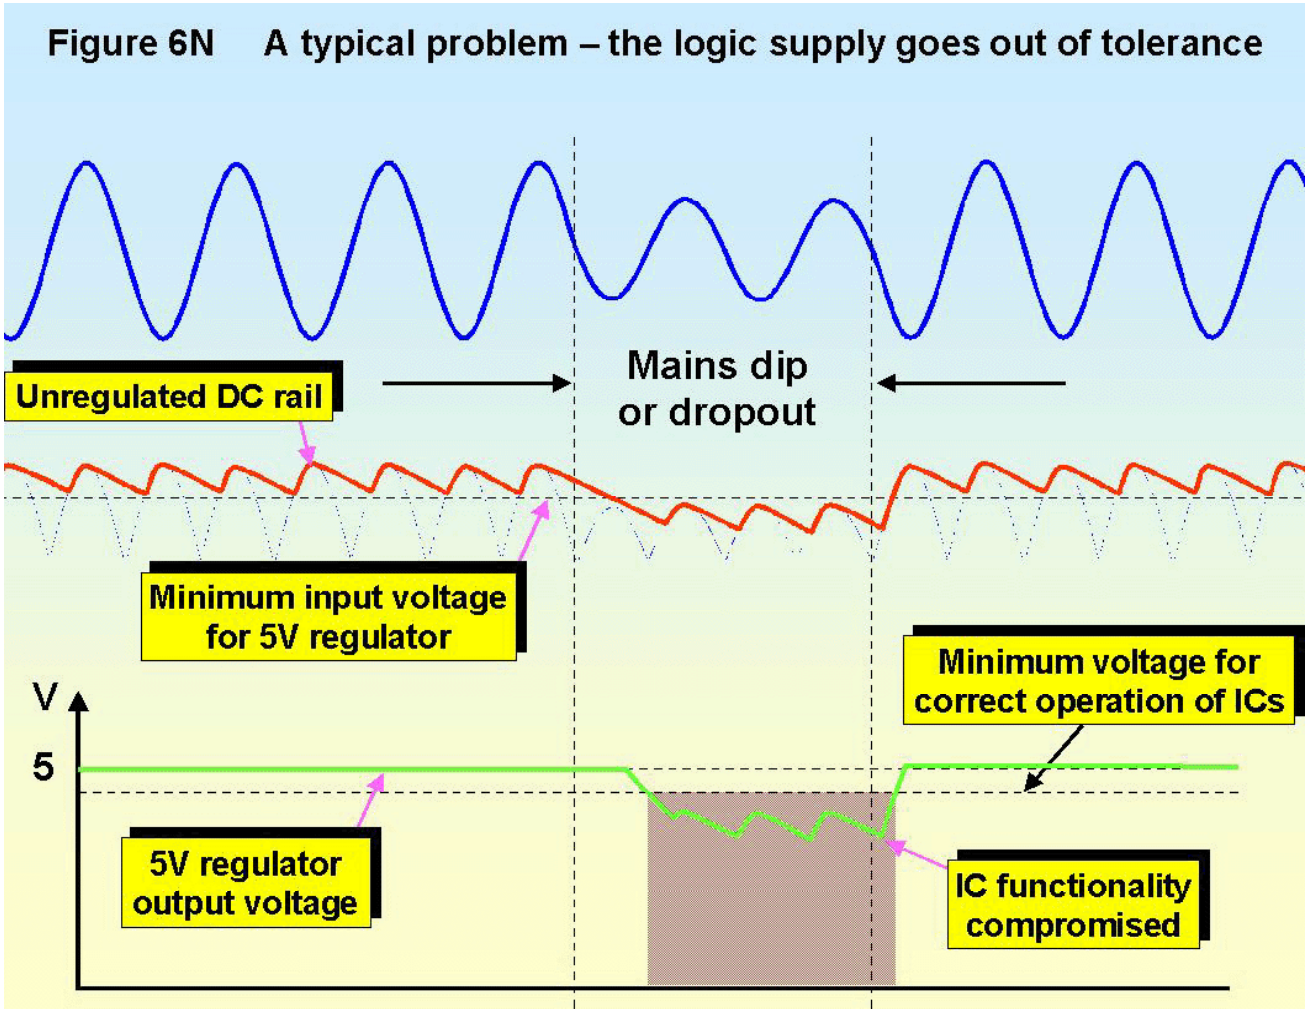

During this out of tolerance period anything can happen to logic systems and any software they are running - it is a very unpredictable situation and the results can be very damaging to product functionality. Operational malfunctions may be expected and it is not unknown for areas of memory to be wiped or over-written with garbage.

Because most designers only provide the minimum amount of energy storage on the unregulated side, many modern products are vulnerable to this problem. It is often found that products which are quite immune to their mains supply being switched on and off are vulnerable to dips and dropouts lasting just a few tens of milliseconds. It is difficult to manually switch mains faster than 200ms, while surveys (such as Unipede's) show that shorter dips and dropouts are more commonplace.

Immunity to dips, flicker, and dropouts is usually easiest to achieve by making sure that the unregulated DC rails of a product's power supply have long enough 'hold-up' times. This is usually achieved by using higher values of capacitance on the unregulated DC rail and also by increasing the unregulated voltage so as to allow for the increased DC ripple during the dips and dropouts. We need sufficient energy storage, given the power consumption of the product, to ride through the frequent short dips and dropouts. If non-essential functions can be switched off during a 'hold up' event, or microprocessors put into 'sleep' mode, it will help increase 'hold-up' time.

When dips or dropouts last so long that they become sags, brownouts, or interruptions and the unregulated energy storage cannot 'hold-up' any longer, it is important to ensure that before the logic (or other) DC supplies go out of tolerance the product 'cleanly' and safely ceases operation and shuts down. Then, when the event has passed, 'cleanly' and safely restoring operation (see later for more on the safety issues associated with power restoration). (In this context 'cleanly' means that no unintended operations occur and no memory is lost or corrupted.)

Figure 6P shows that because capacitive energy storage (in Joules) = ½CV^2, for two identical CV products the energy storage of the higher voltage part is greater. CV product is closely related to physical volume and cost, so the two capacitors in the example will have the same physical size and cost despite the higher voltage one holding twice the energy. Note that to use a much higher unregulated voltage usually requires appropriate circuit design for the regulator so that the higher voltage can be efficiently utilised.

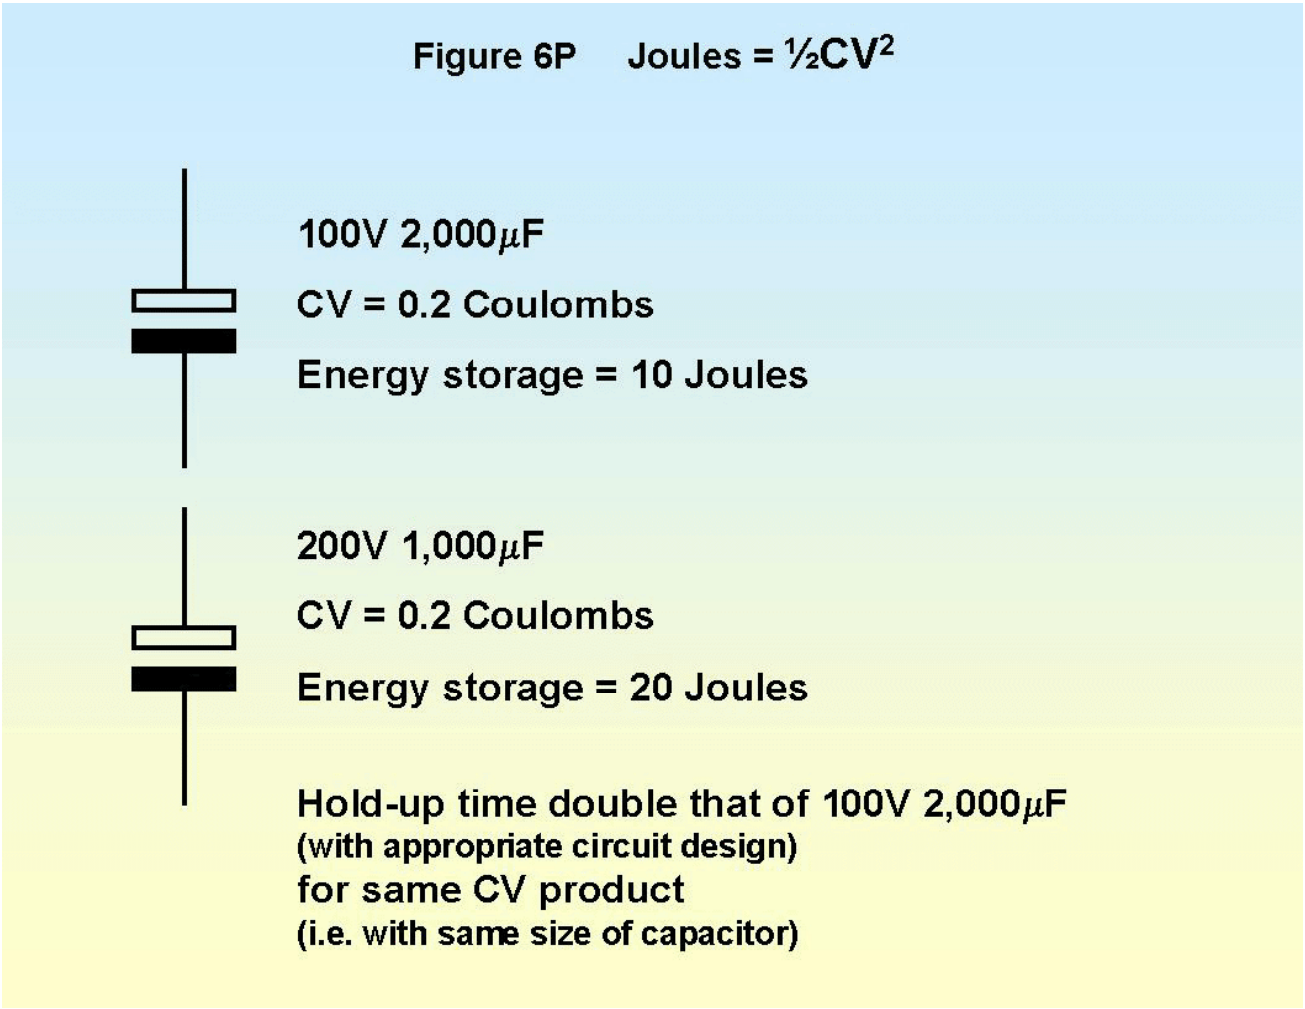

Continuous operation from a battery which is usually being charged up by mains-powered charger, as is normal for portable computers, is an excellent way of overcoming all dips, sags, dropouts; and interruptions lasting from minutes to hours (even months in some cases).

Most mains power supplies use unregulated DC rail voltages that exceed the dropout voltage of their regulators by the minimum amount possible, to improve efficiency and reduce size and cost. Unfortunately, this makes it difficult for the unregulated storage to provide much protection against dips and dropouts. (In the section on emissions of voltage fluctuations it is also shown that such designs make some products more likely to have emissions problems.)

Using unregulated DC rails with an increased voltage makes it easier to achieve significant energy storage to cope with dips and dropouts at reasonable cost, because capacitors store more energy per unit volume at higher voltages, as described earlier.

Operating the unregulated DC storage at a much higher voltage than is needed for circuit operation makes it possible for all the excess voltage to be used for DC ripple caused by dips and dropouts. When linear regulation is used higher unregulated voltages make efficiency worse, so this method is more suited to switch-mode regulators (which will probably need higher voltage ratings for their power devices).

Where an ‘Active PFC’ or similar boost converter precedes the unregulated DC storage capacitor (refer to the section on harmonic emissions) – designing the boost converter for a wide range of input voltages can ensure that the unregulated DC rail is kept fully charged even during quite severe dips and sags. A standard ‘universal input’ power supply (80 - 260V) used on 230V mains will be happy to operate continuously with only 80Vrms input, equivalent to a 65% dip or sag lasting for *any* length of time.

But ‘universal input’ power supplies won’t protect from dips and sags of more than 30% when used on 110V mains. To protect from larger dips and sags they would need to be re-designed for a 40 - 130V range. Alternatively they could be preceded by a selectable (110/230) voltage-doubler type rectifier circuit. In both these cases the benefit of ‘universal’ operation is lost unless the switch from 115 to 230V operation is done automatically.

Some products with linear power supplies have used triac-tap-switched mains transformers for many years. The triacs are switched automatically between a number of tappings on the primary of the transformer, to maintain the secondary AC voltage (and the unregulated DC voltage) within limits during dips, sags and swells. Hysteresis is used to prevent the triacs from ‘hunting’ between taps when the sensed voltage is close to a switching threshold. The unregulated DC needs to have enough hold-up for the response time of the tap-changer.

No boost converter or tap-changer can cope with a dropout or interruption – but at least they can help a little by making sure the unregulated DC capacitors are fully charged when the dropout begins. This won’t make any difference on formal EMC tests, but it may make a useful difference in real life, where dropouts can follow immediately on from a dip or a sag (brownout).

Designing equipment to cope with swells can just be a matter of rating components to cope with the overvoltage, and maybe with some extra heating too. Alternatively, it is possible to detect a dangerous overvoltage and shut down the equipment to protect it – although this should only be done where the user can cope with the resulting non-availability and no safety hazards are caused.

Surge protection devices (SPDs) are usually dimensioned to start conducting at just above normal mains tolerances. A swell which is higher than these tolerances could easily burn them out, since they aren’t rated for continuous dissipation. Some SPD types fail short-circuit and cause the equipment’s supply fuse to open, taking the equipment out of service. Some types fail open-circuit and can leave the equipment vulnerable to supply surges, with no warning. Overheating SPDs can cause fire and shock hazards if the designer has not considered the fact that they might glow red hot during a supply swell.

A number of ICs are now available for monitoring power supplies and effecting a controlled shut-down and reset when unregulated or regulated DC rails drop too low. These are usually sold as ‘brownout detectors’. Almost no equipment nowadays should rely on a simple resistor-capacitor-gate ‘power-on reset’ circuit, because they have always been inadequate for protecting circuits from the full range of perturbations of the AC supply.

Some power-monitoring devices monitor the unregulated rail so they are able to protect logic circuits before their DC rails goes out of tolerance by inhibiting read/write operations to memories of various types. This prevents memories from being overwritten with garbage during an undervoltage situation. In the case of brief events which don’t require a reset some devices can simply freeze RAM data momentarily enabling operation to continue as if no hiccup in the supply had ever happened.

Various grades of brownout detector are available, depending on the accuracy and tolerance of their voltage detecting functions. The more accurate and the tighter their tolerance, the more costly the tend to be. They are also available with a variety of functions to suit different applications.

Some circuits sample the mains voltage, usually to control heat or other parameters. These can often use a smoothing capacitor to ride out short-term disturbances. But some circuits take timing data from the mains, usually from the zero crossings. These can miscount, misfire, and go wrong in

a number of ways – so their design must ensure that all foreseeable timing glitches are prevented or else don't cause any malfunctions or damage.

The desired behaviour of equipment during supply dips, sags, swells, dropouts and interruptions depends on the application, and may be critical in safety-related applications. Some applications may need a controlled power-down, which may mean providing sufficient energy storage to manage the power down process. But some applications (e.g. medical life support) may not permit a power down at all, making their energy storage requirements very high.

Logic circuits should be prevented from causing malfunction by brownout detectors, but analogue circuits also need consideration. For example an audio amplifier must not produce loud instability, pops, clicks, or thumps during power down (or power up).

For motor drives the requirements depend strongly upon the application. They might be required to stop as quickly as possible, coast, or slow down in synchronism with a number of other motors controlled by other manufacturer's equipment as a part of a large installation. When power is restored motors might need to ramp their speed up slowly or quickly, or they may not be allowed to rotate until manually commanded.

Since the EMC Directive does *not* cover safety, where errors and malfunctions in safety-related systems can be caused by the effects of electromagnetic disturbances on electronic circuits, these must be dealt with under safety directives such as the LVD or Machinery Directives.

Uninterruptible power supplies (UPSs) can be used to help products 'ride through' dips, sags, dropouts, and interruptions, but are not a universal panacea. Their ability to cope with all AC power EMC disturbances without passing them on to the 'protected' equipment should be investigated carefully, as should their reliability performance.

It has been known for the reliability of electronics to be reduced because the UPSs they were run from had lower reliability than the mains supply they were supposed to be improving on. Having said that, properly dimensioned and reliable UPSs which use full-time double-conversion can be a great help in preventing mains power quality problems from disturbing the operation of equipment.

Continuous double-conversion types are preferable as they don't cause dips and dropouts by switching between their mains and alternative power sources. Many low-cost UPSs don't use continuous double conversion and may not be able to respond to short dips and dropouts – indeed they may actually cause dips and dropouts by their mode of operation. These may be able to be used if the equipment's power supply has sufficient hold-up time.

Some models of UPS provide good isolation from conducted RF interference, transients, surges, etc. – but not all of them do – so check the manufacturer's specification carefully! Also make sure that the UPS will withstand the types of dips, sags, swells, and dropouts that you anticipate for the application. Some UPSs do not themselves have low levels of harmonic emissions into their mains supply, and some might supply distorted output waveforms (including rotary types such as motor-generator sets).

Serious energy storage may be needed if equipment is to keep functioning during a long power outage. Where loads are light, batteries or 'supercapacitors' may provide sufficient energy for weeks, but at the other extreme serious power users may need rooms full of batteries just to keep them going for the 30 seconds or so it takes their diesel-powered generators to get up to speed.

One way of reducing energy storage needs and getting increased operational time from your energy storage is to shut down non-essential functions during an outage. Some lighting, heating, chilling, and some visual displays may be able to be treated as non-essential.

Some types of UPSs only support and preserve the mains waveform, rather than take over from the mains. A paper mill in South Africa uses a superconducting energy ring in such a UPS to provide enough storage just for dips and dropouts. At a cost US$10 million it cost less than the web-breaks they were getting every week.

Motor-generator sets with alternative power sources for the motor are the original continuous double-conversion UPS. Motor generator sets are also very useful for preventing harmonic emissions from getting into the public mains supply. If they have enough rotational inertia they can also help reduce the load current fluctuations that would otherwise cause emissions of voltage fluctuations and flicker.

Automatically controlled motorised variable transformers (variacs), ferroresonant constant voltage transformers (CVTs), etc. are all systematic ways of dealing with sags and swells.

6.3 Emissions of voltage fluctuations and flicker

EN 61000-3-3 assumes a typical value for mains supply impedance at 50Hz of (0.24 + *j*0.15) Ω for the phase conductors and (0.16 + *j* 0.1)Ω for the neutral. When testing single-phase equipment these two are combined into one overall source impedance of (0.4 = *j* 0.25)Ω, and it is worth noting that *j*0.25Ω at 50Hz corresponds to an inductance of 796μH. When testing balanced three-phase equipment which draws negligible neutral current the neutral impedance is neglected and the overall phase-to-phase impedance is (0.48 + *j*0.3)Ω.

When equipment draws a fluctuating current from its mains supply these impedances in the supply convert the fluctuating currents into fluctuating mains voltages, as shown by Figure 6Q. These can cause problems for other equipment and also cause lighting 'flicker'. Lightning flicker is very annoying to most people and can reduce their working efficiency. In some cases flicker can cause actual illnesses (especially those related to stress and/or eyestrain) and it may be able to encourage epileptic fits in susceptible people.

Figure 6Q How mains voltage fluctuations and flicker arise

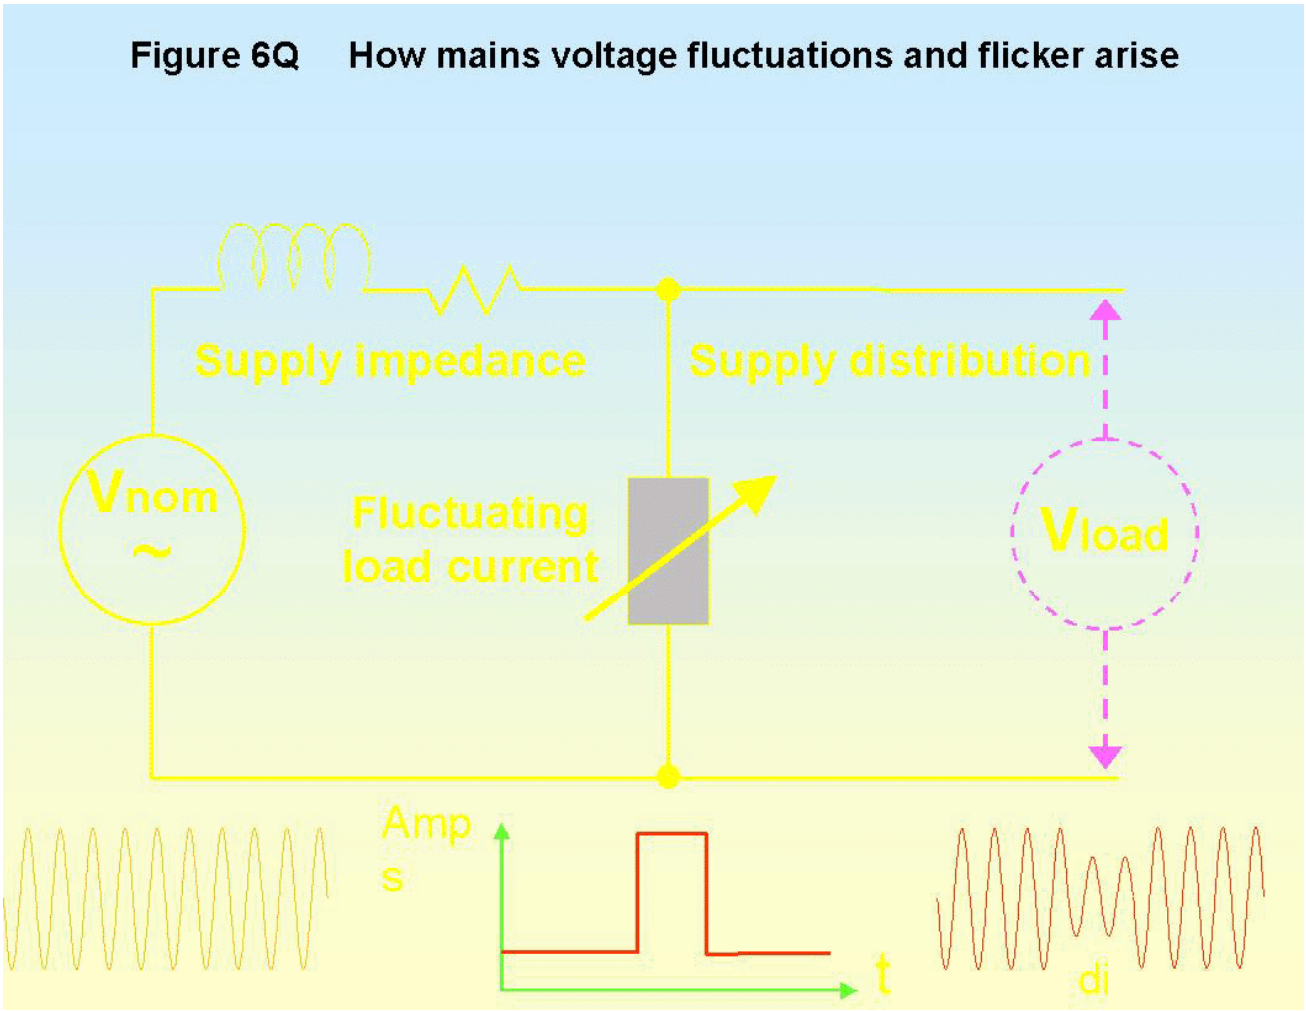

EN 61000-3-3 is now mandatory under the EMC Directive for all equipment that consumes up to 16A/phase from the public 230V supply. Test laboratories use an instrument called a 'flickermeter' to measure to this standard. EN 61000-3-3 limits the emissions of voltage fluctuations and flicker and also limits the inrush current at equipment switch-on. The inrush current requirement arises because after a power outage the very heavy inrush currents typical of electronic equipment (as they recharge their DC storage capacitors directly from their bridge rectifiers) can cause an overcurrent trip to occur in the high-voltage distribution, making restarting the network very difficult.

In the case of momentary power outages caused by arc-suppression protection devices on high-voltage overhead power lines when a lightning strike occurs, an overcurrent trip when the power is re-applied can extend what was a two-second outage to one that could last several hours.

Because so much of the load on the low-voltage power network is now electronic apparatus, high-voltage overcurrent tripping due to inrush is now a serious problem and the inrush currents of apparatus need to be controlled. Unfortunately, the issues of EN 61000-3-3 which exist at the time of writing do not mention the reason for its inrush current limit, making it easy for designers and test laboratories to accidentally overlook this requirement.

The requirements of EN 61000-3-3 are very complex because they are based on human perceptions of lighting flicker. The limits for voltage fluctuations and flicker follow a curve of amplitude versus rate (figure 4 in EN 61000-3-3) intended to correspond to the human perception of lighting flicker on a 60W filament light bulb, and are shown in Figure 6R. The peak human sensitivity to flicker is between 5 and 20Hz (300 to 1200 fluctuations per minute) so these rates have the toughest limits. At lower or higher rates greater levels of voltage fluctuation are permitted. For instance: at a rate of once per minute, the permitted voltage fluctuation level is approximately 10 times higher than the limit for 1000 fluctuations per minute.

Figure 6R Sketch of Figure 4 of EN 61000-3-3

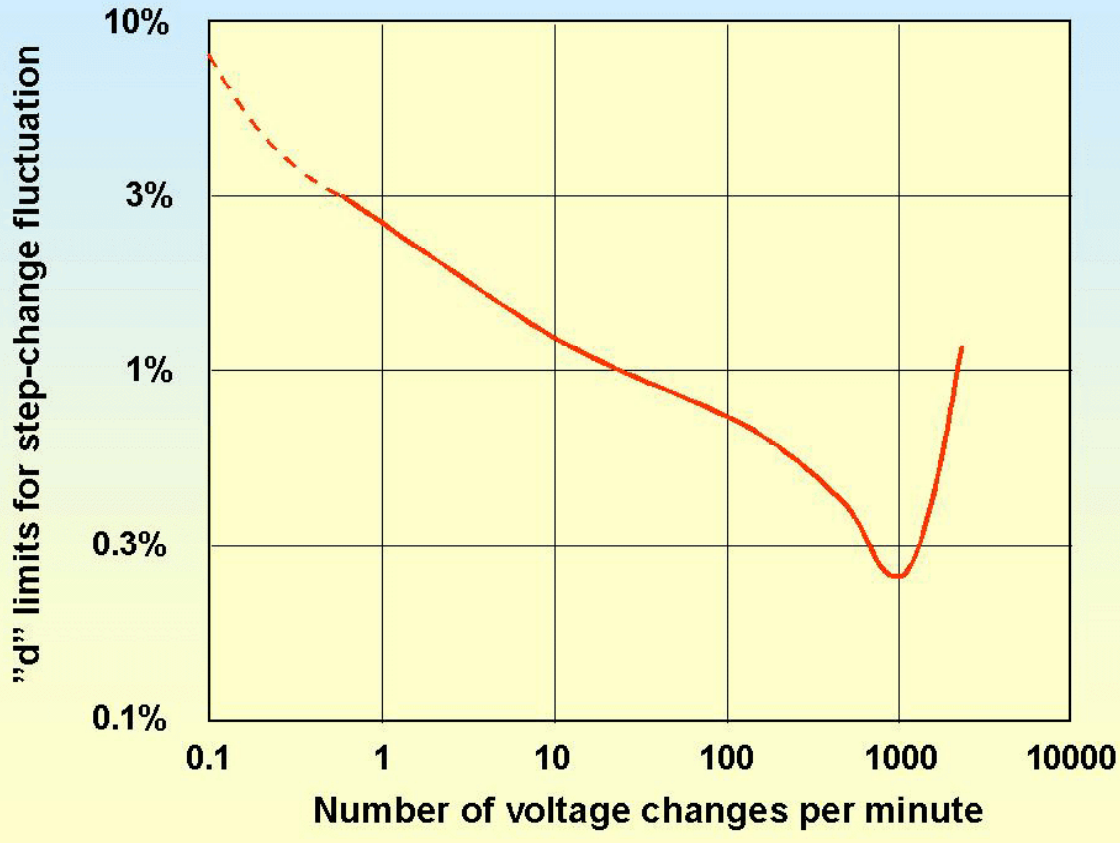

For voltage fluctuations that occur every 1.25 minutes, the steady-state limits for a step-change are 3% of the nominal voltage, corresponding to a change in resistive load current of 15A every two minutes or so. This allows typical domestic hotplates, heaters, etc., to be switched on and off without suppression as long as they don't switch more often than once every 1.25 minutes.

The peak limit for a step-change fluctuation every 1.25 minutes is 4%, one-third higher than for steady-state value, equivalent to a resistive load current change of about 20A. The steady-state limit must not be exceeded for more than 200ms.

During an initial switch-on inrush situation, or for load-switching events that occur less than once per hour, the values of the steady-state and peak voltage fluctuations are allowed to be 33% higher than the above. It may be that future versions of EN 61000-3-3 will permit equipment that delays its switch-on inrush current by more than 10ms to have even higher emissions, since delayed inrush places less of a burden on a mains network which is being restarted after a power interruption.

The basic limit curve (figure 4 in EN 61000-3-3) assumes a simple step change in mains voltage (whether an increase or a decrease) – but different voltage fluctuation waveshapes cause different flicker perceptions. Some waveforms will measure lower or higher values for their peak voltage fluctuations than others, even if their instantaneous peak values are the same. A complex mathematical transformation is required to determine whether a non-step type of waveform complies with the limits. This transformation is conducted automatically by the digital signal processing in 'flickermeters' (for which the defining standard is EN 61000-4-15).

Section 4.2.3 of EN 61000-3-3 gives some guidelines to designers for estimating the effects of waveshape on peak voltage fluctuation for a few commonly-encountered shapes – but only for fluctuations that occur less than once per second. This makes it easier to estimate the likely voltage fluctuation emissions from an equipment calculation, simulation, or by simple measurements using standard laboratory equipment (oscilloscopes, for example). The accuracy of these estimates is claimed to be no better than ±10%, so results which are within 20% of a limit should be checked with a flickermeter on the actual equipment to make sure it complies.

If you are using a test supply with a total harmonic distortion of under 10% and a supply impedance the same as that specified by EN 61000-3-3 you can measure the voltage fluctuation directly with an oscilloscope. If instead you measure the load current fluctuation with another supply impedance you would need to transform it mathematically into the voltage fluctuation that could be expected using the standard impedance. But beware – the load current fluctuations will themselves depend upon the supply impedance, so if measuring load current it is best to make sure your supply impedance is close (both in resistance and inductance) to the standard supply impedance. Synthesised sources of mains voltage are now available, either with programmed impedances (or zero) achieved by feedback techniques or with the standard impedance. As time goes on more manufacturers are entering this market and the cost of these sources is falling. Figure 6S shows an example of such a source.

Figure 6S Example of a mains source

If you know your equipment's load current waveshape – whether calculated, simulated, or measured with close to the standard source impedance – you can calculate or simulate the resulting mains voltage fluctuation waveshape. By referring to 4.2.3 and figures 5, 6 and 7 in EN 61000-3-3 you may find that altering the load current waveform can give useful reductions in the measured flicker value, even if the rate per minute, peak and steady-state amplitudes of the fluctuation remain unchanged.

EN 61000-3-3 requires the integration of each measured sample of voltage fluctuation for a period of 10 milliseconds. Voltage fluctuations occurring on shorter timescales are 'smoothed out' by this integration process.

6.4 The influence of the supply inductance

The supply impedance has an inductive component (L), so since $V = L\,di/dt$ the rise/fall times of the load's current waveform could have an effect on voltage fluctuations. The standard total single-phase L is around 796μH and its impedance exceeds the resistive component of the supply impedance above 80Hz. The rate of change of current required to create a 4% voltage fluctuation – due to the standard supply inductance alone – is 11,558 A/s.

80Hz is not a very high frequency, so equipment which draws fluctuating currents with spectral components above 80Hz may find that the inductive component of the standard supply impedance contributes more to the measured voltage fluctuation than the resistive component.

For example, high rates of current change are commonly exceeded by DC storage capacitors charging-up via rectifiers directly from the mains, especially when they are connected at the peak of

the mains cycle – the worst-case for capacitor inrush. But under such conditions typical capacitors of under 1,000μF require less than 1ms to charge to 370V and although the peak voltage fluctuation caused by the capacitive inrush would be very large indeed, the 10ms integration time 'smooths it out' to make the EN 61000-3-3 measured value very much less. Simplified analyses indicate that the rectifier-capacitor power supply input stages typical of switch-mode power supplies risk failing the inrush current limit when capacitor values exceed 600μF.

Where the DC current load on the storage capacitor is significant during the switch-on period, the capacitance values may need to be much less.

For many low-power consumption products the rise and fall times of their load current may not create significant emissions problems *due to the supply inductance*, at switch-on or during operation. But equipment which has unregulated 400VDC capacitors of over 600μF powered directly from mains rectifiers – and equipment with high levels of pulsed power (e.g. strobe lamps, powerful pulse generators) – *should* always consider the effects of L*d*i*/*d*t* on their emissions of voltage fluctuations.

Most DC capacitors aren't significantly discharged between one half-cycle and the next and only need 'top-ups' after their switch-on inrush. The fluctuating currents drawn by their DC load will then be the major contributor to their emissions.

Products such as computers, computer monitors, and TVs were traditionally not a cause of voltage fluctuations and flicker (except at switch-on, when CRT degaussing coils often caused the biggest problems). But they are increasingly adopting 'Energy Star' or other energy saving methods which can cause their DC load currents to change by 90% or more in under 1ms between standby and operational modes, making modern versions into significant sources of voltage fluctuations and flicker.

Reducing inrush currents

Where AC-DC power converters draw such high inrush currents at switch-on that their corresponding voltage fluctuation (integrated over 10ms) creates a problem with meeting EN 61000-3-3, steps should be taken to reduce their inrush current.

Inrush current into DC power supplies can be limited and/or have its rate-of-change slowed down by fitting resistors in series with their mains inputs. These resistors are usually shorted-out by relay contacts or triacs after the first second or two to permit normal operation of the equipment.

An alternative to an expensive resistor/relay combination is to use negative temperature coefficient (NTC) devices instead of a series resistor. NTCs have an initially high resistance, reducing to a low value as they heat up due to the passage of the equipment's mains current. Their benefit is that they do not need to be shorted-out by a relay or triac for full operation of the equipment – but their dimensioning can sometimes be tricky. Their thermal inertia and hot/cold resistances need to be taken into account with the operational current consumed by the equipment and the permitted maximum inrush current. NTC devices can get very hot, and must be positioned so they cannot melt plastics, damage cables or components, or burn operators or service technicians.

A useful technique is the soft-start, sometimes called a ramp start. This gradually increases a DC or AC voltage from zero to the full value over a second or two. As well as helping to reduce inrush current soft-starts also help reduce the stresses on components. (They also allow protection devices time to operate before components are damaged when switch-on occurs during a fault or misuse situation, such as an incorrect mains voltage.)

Many types of switch-mode controller ICs have soft-start features designed into them. These help inrush currents by ramping the DC load on the power supply. 'Active PFC' types of controllers don't have large DC storage capacitors after their mains rectifiers, so using soft-starts with these controller types can reduce switch-on inrush to negligible amounts.

Mains-powered motors, transformers and other inductive components can draw large inrush currents, which can vary over a 2:1 range depending on the phase angle of the AC mains voltage at the instant of switch-on. It is often found that switching on when the mains voltage is at a maximum gives the lowest inrush currents with inductive loads. When testing such parts it is important to find their worst-case switch-on phase angle.

Soft-starting/stopping mains-powered inductive loads in the past often used motorised variable transformers, but are now more likely to use phase-angle-controlled triacs. Because such triacs are only used intermittently they often need little or no heatsinking. The conducted and radiated RF emissions from the triacs during starting or stopping might not need to be suppressed if they are infrequent enough. An added advantage of soft-starting motors is that it helps prevent sudden application of torque.

Series resistors and NTCs can often be used with inductive and motor loads, but the reactive nature of the load might make them more difficult to design.

Reducing emissions in normal operation

Where time-proportioning on/off control (sometimes called bang-bang control) is used, it is sometimes practical to replace electro-mechanical switches or burst-fired triacs with electronic power converters. These can be designed to vary the load power in a continuous manner and so emit no voltage fluctuations or flicker. Microwave ovens have been known to use this technique.

If time-proportioning control *must* be used, load switching rates of between 5 and 20Hz are best avoided altogether. Reducing the switching rate of a load below 5Hz (or increasing it above 20Hz) allows it to benefit from the higher limits permitted by figure 4 of EN 61000-3-3. For heaters, chillers, and motors decreasing the switching rate may require additional thermal or mechanical inertia to achieve the desired degree of control.

A supplier of a system must meet EN 61000-3-3 for the whole system. In computer systems where the power-down times of the individual computers are set to be quite low, the system as a whole might create significant amounts of flicker at quite a high rate. It may be enough to merely set the power-down energy-saving times to be longer, to take advantage of the more relaxed limits allowed at lower rates of flicker.

Apparatus with switched loads can also benefit from soft-start/stop techniques – typically using phase-angle-controlled triacs. Figure 6T shows that a ramp up/down time of 100ms reduces the measured emissions to 70% of the corresponding step-change voltage fluctuation. Longer ramp times give correspondingly lower measured emissions, with 1 second reducing the measured value to just 20% of what would be measured on a step-change of the same magnitude.

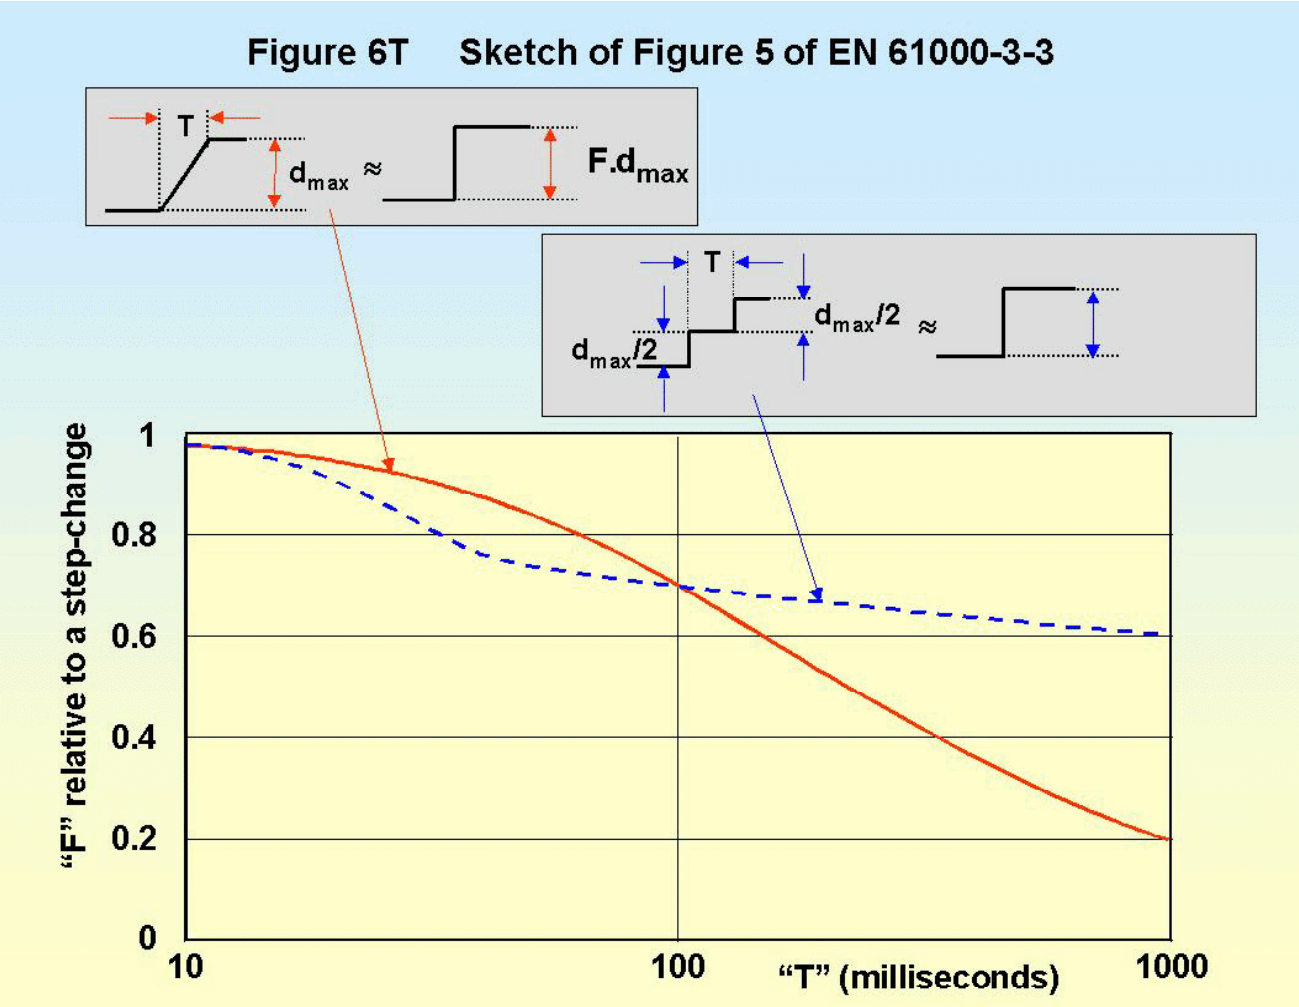

In equipment with multiple smaller heating elements or motors, sometimes all that is necessary is to make sure that more than one or two heaters or motors cannot be switched on or off at the same time. Splitting a large load into two or more smaller loads, each individually controlled so that they do not switch on or off at the same times can be a useful technique. Figure 6T shows that splitting a voltage fluctuation step change into two half-height step changes with a delay of more than 100ms between them reduces the measured emissions to about 70% of the single-step value. Longer delays do not result in very much lower emissions (a 1 second delay only reduces emissions to 50%). Splitting a load into three or more and delaying their switching by at least 100ms each should bring greater reductions in emissions.

6.4.1 Reducing emissions from fluctuating DC loads

Most power supply designers design unregulated rail voltages so that they exceed the dropout voltage of the following regulator by the minimum amount possible, to improve efficiency and reduce size and cost. Unfortunately, this approach makes it difficult for the unregulated DC storage to provide much benefit in reducing the effects of load current fluctuations. (In the section on immunity to supply dips and dropouts it is shown that such design also makes products more susceptible to supply quality issues such as dips and dropouts.)

Where the fluctuating load current is drawn from a DC rail, increasing the size of the unregulated DC storage capacitors can help by reducing the d*i*/d*t* of the load current. This can reduce the contribution to the fluctuations from the inductive part of the supply impedance and (if d*i*/d*t* is low enough) reduce the measured emissions by achieving a ramp start and stop instead of a step-change (see Figure 6T). This technique also requires some inductance or other current limiting

between the mains rectifiers and the storage capacitors – and often needs a higher unregulated voltage too – because the storage capacitors need to experience a large ripple voltage in order to 'smooth out' the energy demands of the fluctuating DC load.

'Smoothing out' the mains current variations by increasing the size of DC storage capacitors would need very much larger capacitance values than are generally used to control DC ripple, especially when the rate of fluctuations is not very high. However, modern developments in 'supercapacitors' (also known as 'ultracapacitors' and 'boostcaps') mean that such options may not now be impractical and novel solutions may now be possible to previously intractable problems.

Where larger storage capacitors follow immediately after an AC rectifier they can mean larger inrush currents at switch-on, and also an increase in the harmonic emissions. This is another reason for adding inductance or other current limiting devices between the rectifier and the storage capacitor.

Using unregulated rails with a very much increased voltage makes it easier to achieve enough energy storage to 'smooth out' DC load changes (and improve immunity to supply dips and dropouts) at reasonable cost. This is because the energy stored in a capacitor is proportional to the square of its voltage. So where large capacitive energy storage is required, using higher voltages can reduce size and cost. Also, operating the unregulated DC storage at a much higher voltage than is needed for circuit operation makes it possible for all the excess voltage to be used for DC ripple caused by the variable load current demands. When linear regulation is used the higher unregulated voltages make efficiency much worse, so this method is more suited to switch-mode regulators.

'Active PFC' switch-mode boost converters (see the section on reducing harmonic emissions) can also be used to reduce voltage fluctuations and flicker. They typically set the time-constant of their input current demand to 0.5s or so and will make step-changes in DC load appear as ramp-like changes in mains current, giving a lower measurement on a flickermeter (as shown by Figure 6T, useful benefits occur when ramp times >100ms).

Where the rate of occurrence of load fluctuations exceeds 120 per minute a 0.5 second time-constant in the boost circuit will 'smooth out' the fluctuations to some degree. With the same time-constant, greater reductions will be achieved for higher rates. Longer boost-circuit time-constants will give greater reductions at fluctuation rates above 120/min and/or help achieve some useful 'smoothing' at lower rates.

For 'smoothing' to work in an active PFC boost circuit the values of the storage capacitors and the unregulated voltage need to be dimensioned correctly for the size and rate of the DC load current fluctuations, and the time constant of the boost circuit.

Beware - most active PFC control ICs will suddenly switch off the input current completely when the maximum voltage on their storage capacitor is exceeded (this usually occurs just after a heavy DC load current has been removed). So if the unregulated capacitors don't have enough stored energy these active PFC circuits can sometimes make emissions of voltage fluctuations worse.

If circuit techniques fail or aren't appropriate for some reason, system-level approaches can help reduce flicker although they won't help the products concerned to meet EN 61000-3-3.

One (expensive) solution may be to run the problem equipment from its own low-voltage distribution transformer, so that their voltage fluctuations aren't applied to equipment powered from the public mains supply. The much lower impedance of the high-voltage distribution network attenuates the effect of their fluctuations considerably. Equipment run from a private LV supply is not covered by EN 61000-3-3 at all.

Motor-generator sets or continuous double-conversion on-line UPS can also be used to reduce voltage fluctuations, if they are dimensioned correctly. At least they can reduce the di/dt of the load current fluctuations. At most – with sufficient energy storage given the rate of the current fluctuations – they may be able to 'smooth out' the mains current so that the peak amplitudes of each fluctuation are reduced.

Some types of power factor correction equipment used in systems and installations may also be able to reduce the levels of voltage fluctuations and flicker caused by equipment.

6.5 Electromechanical switching

Every conductor stores energy in its intrinsic inductance, and inductive devices such as motors also store energy in their magnetic fields. When the flow of current is suddenly interrupted by breaking an electro-mechanical contact, such as a switch, relay, commutator, or slip-ring, the 'flyback' of this stored energy causes a spark due, to breakdown of the air as the circuit-interrupting contact first opens (or when it bounces after closing).

Sparks emit electromagnetic disturbances quite literally from DC to daylight, and many microprocessor circuit designers have been surprised by the ease with which their higher-frequency components can couple into their digital circuits (e.g. via coil-to-contact capacitance, or proximity of cables or PCB tracks) and crash their microprocessors.

It is best to avoid the generation of arcs and sparks by avoiding electromechanical switching completely. The use of solid-state relays, brushless DC motors, AC motors, and the like all help eliminate sparking, although some of these will add new EMC problems of their own.

6.5.1 Suppressing arcs and sparks at switches, relays, and contactors

Where sparks cannot be avoided, emissions standards will be easier to meet by making sure there are no more than 5 sparks per minute in the product, with a spark duration of 10ms or less (less than half a mains cycle, typical of a microswitch or fast-acting relay). In heavy industrial applications it also helps meet emissions standards if the total rate of spark production is less than five per minute, but spark durations of up to a second or two may be acceptable. Beware – although these rates and durations of sparks may be allowed by an emissions standard, they can still upset sensitive circuits so may not be desirable for operational reasons, especially where critical functions are being controlled or monitored.

Emissions from arcs and sparks are usually reduced by 'snubbing'. Simple snubbing involves connecting a series combination of R and C (sometimes just C) near the switching element to slow the rate of rise of inductive flyback voltage and so limit the size of the resulting spark. Connecting a snubber across the contact gap has the disadvantage, in AC circuits, of allowing a leakage current to flow which might shock a person who worked on a circuit expected it to be safe because its relay contacts or switch were open.

Connecting snubbers in parallel with the load's send and return conductors, close to the switching element, sometimes gives better results than connection across the switched contacts, and does not allow leakage past the contacts. Sometimes two sets of snubbers may be required, one to deal with the flyback of the load's inductance, and one to deal with the flyback of the supply's inductance. Figure 6U shows the alternatives for snubbing switch and relay contacts.

Figure 6U RC snubbing techniques at switched contacts

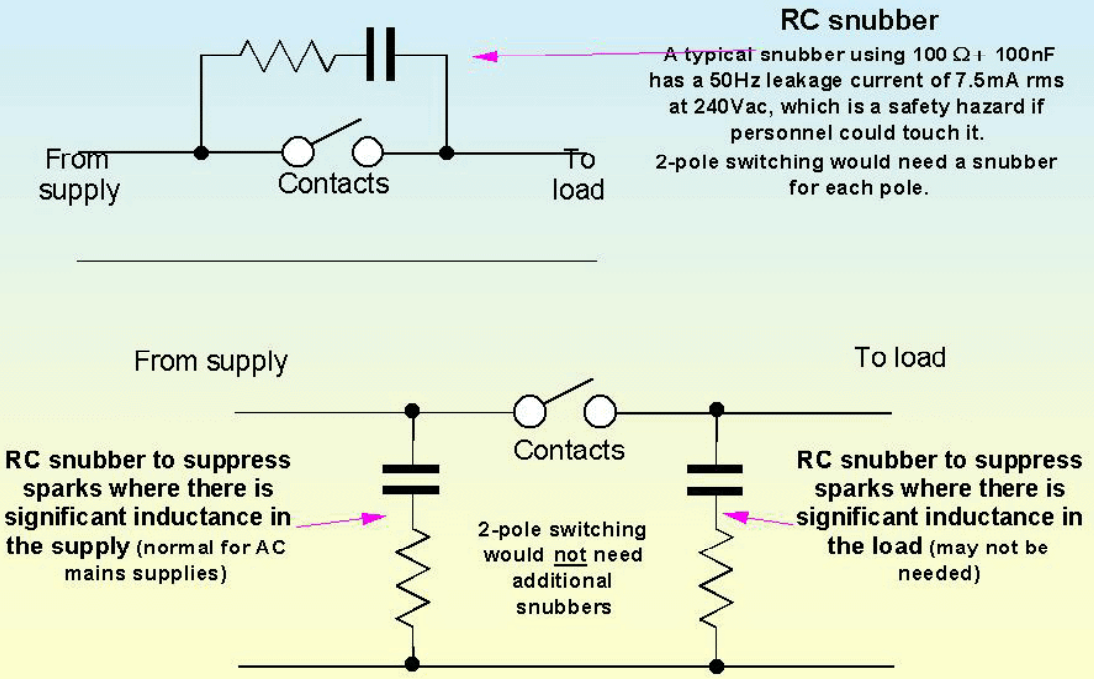

Snubbers can also use non-linear devices such as diodes, rectifiers, zeners, and a variety of surge protection devices (see Part 3 of this series) to provide an alternative path for the flyback currents, either on their own or in conjunction with RC snubbers. The higher the turn-on voltage of the device, the faster the stored energy will collapse and the quicker can be the rate of cycling of the load. Unfortunately, the higher the turn-on voltage, the greater will be the spark at the contacts, so this can lead to a compromise between rate of operation and emissions. DC circuits can use unidirectional semiconductor snubbers, remembering that the flyback voltage has the opposite polarity to the applied voltage.

A side-benefit of all spark suppression techniques is that they generally increase contact life.

6.5.2 Suppressing arcs and sparks in DC motors

In general, DC motors are a very serious source of conducted and radiated emissions, and are very difficult to suppress. The filters and other suppression devices required for them to meet emissions standards can cost more than (and sometimes bulk as large as) the motor or bell itself.

Some ‘pancake’ DC motors don’t spark because their brushes connect to a number of rotor windings at once, so there is always one of them in circuit to provide a path for the flyback currents in the windings.

Larger, more industrial DC motors with fully-enclosed metal bodies tend to emit less and be easier to suppress than lower-cost motors. Larger DC motors connected by many metres of cable to their controls or drives are able to reduce their emissions by using good quality screened cable, as long as its screen is 360° bonded at the motor’s metal terminal box (and probably to the earthed cabinet enclosure at the controlling end).

Where this technique is not sufficient, or impossible to apply (as in many motorised toys or domestic equipment such as CD players), it is best to use a motor with transient suppression fitted to its *rotor*. The rotor is where the energy is stored in a brushed DC motor, and is best dealt with before it causes sparks in the commutator. A 'varistor disc' can easily be fitted to most low-voltage DC motors, essentially connecting a varistor (voltage-dependant resistor, described in Part 3 of this series) between each pair of contacts on the commutator. For a 24V motor the varistor disc may be designed to conduct at 30V or so, and only conducts current when flyback occurs. During flyback, it conducts the energy into the neighbouring winding and limits the resulting overvoltage at the commutator to under 45V or so. This still causes sparking, but only small ones with much lower emissions.

Where a varistor-disc motor cannot be obtained, it is usually necessary to shield the motor and filter after the commutator, not always very easy to do at low cost. Metal shielded motor bodies are preferred to (crudely) catch the radiated emissions from the sparks and return them back into the motor where they came from, via the filter. The filter is also needed to reduce the conducted emissions. Since DC motor emissions are still going strong at 1GHz (and also, in fact, at 10GHz), motor shielding needs to have very few very small gaps. Motors with metal end caps and metal bodies may appear well-shielded, but the bonds between the metal parts may be poor due to paint or anodising.

A filtering technique which works well is to bond one of the commutator terminals directly to the metal motor body (the shield). The other terminal is decoupled to the motor's metal body by a capacitor with very good high-frequency characteristics, such as an 820pF multilayer ceramic with a COG or NPO dielectric and very short leads. The capacitor must be rated to cope with the transient voltages caused by commutation. Where it is not possible to bond one of the brushes directly to the metalwork, it should be decoupled in the same way as the other brush. A low self-inductance is very important for these bonds and decoupling, and even 5mm of length or distance can be crucial. Feedthrough capacitors of around 1nF, screwed into the body of a fully metal enclosed motor and used as the brush terminals to the motor cable, often work very well indeed, although they are not inexpensive.

Correct application of shielding, bonding, and decoupling, may make the motor's emissions low enough. If not, the next step is to add chokes to the brush leads, as close as possible to the decoupling capacitors but immediately outside the motor body. Differential chokes and across-the-line capacitors may be needed to reduce low-frequency emissions, whereas common-mode chokes and line-to-chassis capacitors are usually best at suppressing high-frequency emissions. A multi-stage filter using both types of choke may be needed in difficult cases, and is often best implemented with a PCB mounted directly on the motor end-cap at the commutator end, to keep all lead lengths low and to permit low-inductance bonding of capacitors to the motor body.

A varistor-disc motor with only very tiny sparks on its commutator should last longer before it commutator wears out, whereas a shielded and filtered motor will not benefit in this way because its sparks have not been made any smaller.

6.5.3 Suppressing arcs and sparks in electric bells

Like commutator motors, electric bells create emissions from DC to daylight. The best technique is to remove their spark gap and use an oscillator (astable) circuit to pulse current through the hammer solenoid at the hammer's natural frequency. This is usually very much cheaper than any filtering methods. Such an electric bell could be much more reliable, and of course would require no adjustment to its spark-gap during manufacture. It may be that this is the first significant improvement in the design of electric bells since the 1880s.

6.6 Power factor correction

EN 61000-3-2 came into force on 1/1/2001 under the EMC Directive for all equipment consuming up to 16Amps/phase and connected to public low-voltage mains supplies. It limits the harmonic (non-sine wave) currents drawn by products, for all lighting equipment consuming above 25W and all

other products consuming above 75W. Professional equipment rated at over 1kW has no limits to meet at the time of writing.

The problem for typical rectifier-capacitor AC-DC power converters is that they appear to the power distribution as non-linear loads because they only top up their DC storage capacitors at the peaks of the AC supply waveform. Their supply currents are discontinuous, non-sine wave, and rich in harmonics (as shown by Figure 6V).

Figure 6V Harmonic emissions from single-phase power supplies

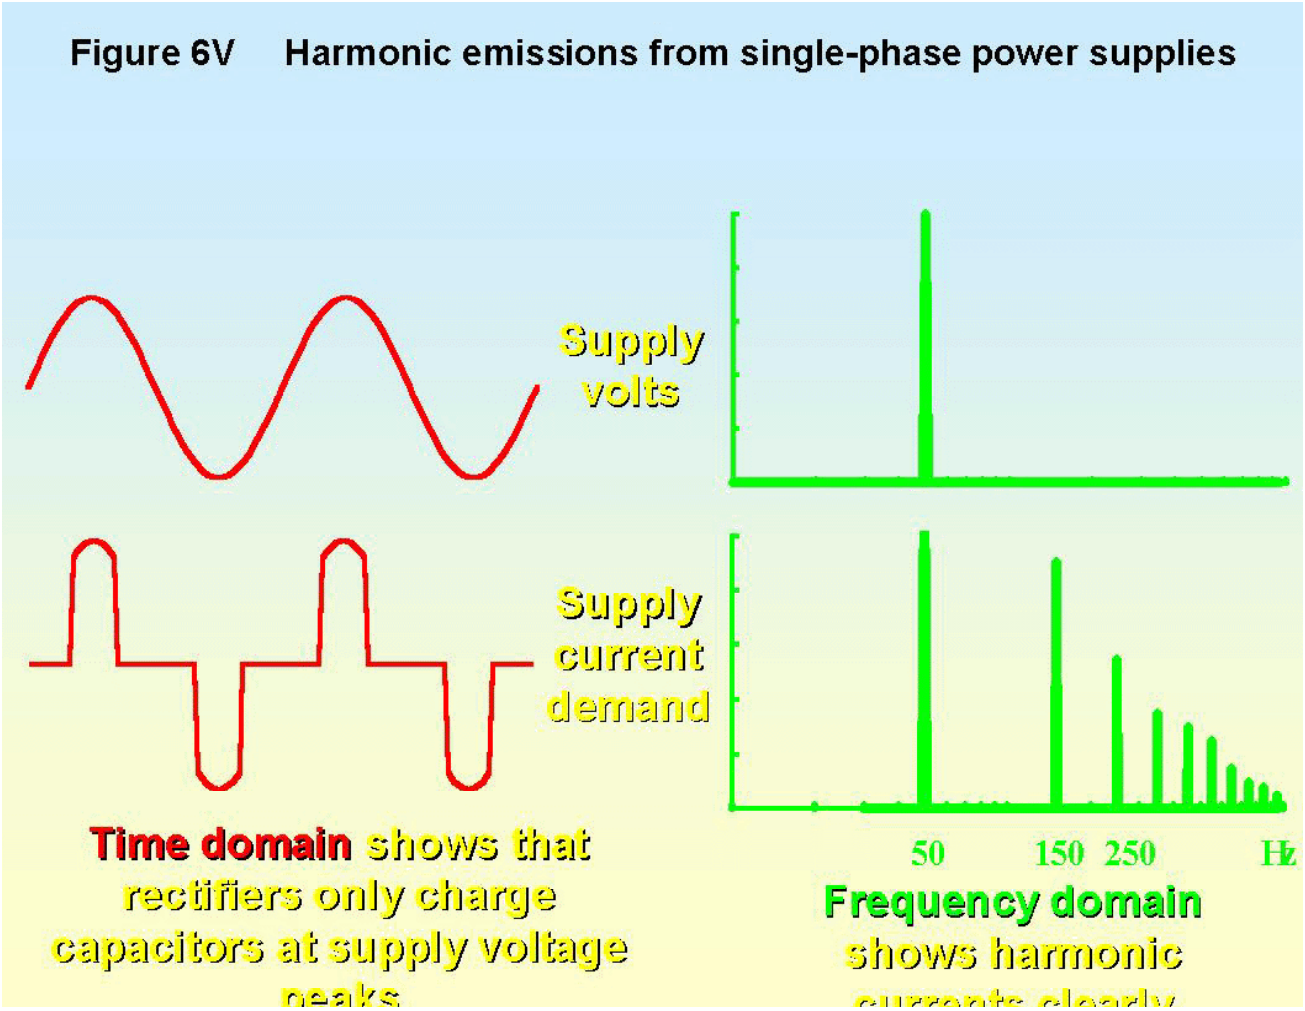

The special problem for single-phase power supplies is that they emit triple (or triplen) harmonics (3rd, 9th, 15th, etc.), which are a particular nuisance since they add linearly in neutral conductors (no cancellation) and are a major cause of cable and transformer overheating.

In a larger installation with a lot of single-phase electronic loads (typical of a modern office) the neutral currents can reach over 1.7 times the size of the phase currents. Since many older buildings are wired with half-size neutrals, and since building neutrals aren't fused, the fire hazard is clear.

Harmonic emissions create a number of problems for power generation and distribution, not least of which is overheating and fire (something that fire insurers are becoming increasingly aware of). There are a number of ways of dealing with this problem at the equipment and installation levels. Electronic solutions at the equipment level are the main concern here.

There are many other non-linear loads which also cause harmonic currents in the supply, such as transformers and motors; arc furnaces and welding equipment. Fluorescent lamps with magnetic ballasts have harmonic emissions too, and although they include even-order harmonics they usually don't extend to very high frequencies. High-frequency ballasts for fluorescent lamps (including the

popular ‘low energy’ filament bulb replacement products) are simply single-phase AC-DC switch-mode power supplies – with all their harmonic problems. Three-phase power converters (sometimes called 6-pulse converters) are also a source of harmonic emissions, but if operated with balanced loads they produce low triplen levels.

When an item of equipment draws (‘emits’ in EMC terminology) harmonic currents from a sine-wave AC supply, the harmonic currents are reactive and increase the VA consumption of the equipment without affecting its consumption when measured in Watts. The ratio of Watts to VA consumed by a load is known as its Power Factor (PF), so where an equipment has significant emissions of harmonics it also has a poor power factor.

A PF of 1 means that the Watts consumed equals the VA of the equipment, in which case it looks like a pure resistive load and has no harmonic emissions. AC-DC power converters with no harmonic reduction techniques tend to have PFs of around 0.6. Techniques which reduce the emissions of harmonic currents into the AC supply also improve the equipment’s PF, so they are usually called Power Factor Correction (PFC) techniques.

Don’t confuse real Power Factor (= W / VA) with the power factor traditionally used by electrical generation and distribution engineers, which is the cosine of the angle between the sine-wave supply voltage and the load current and can be adjusted by either adding capacitance or inductance to a power line. The electrical engineers’ traditional PF is based on sine wave voltages and linear loads (resistive, inductive, or capacitive) and so is actually a special case of real PF. Few, if any, power distributions these days have linear loads, and you cannot correct the PF of an electronic AC-DC power supply using the traditional methods for linear loads.

There are a number of techniques for reducing harmonic emissions (improving PF) for items of electronic equipment:

- Filtering
- Passive PFC using an inductor between bridge rectifier and DC storage capacitor
- Passive PFC using a charge pump with a suitable SMPS controller
- Active PFC using a boost regulator after the bridge rectifier
- Increasing 3ϕ rectifiers to 6ϕ

Filtering means simply connecting filters at the AC input to the power converter to limit the emissions of some or all of the harmonics. Because the frequencies are so low, and the currents involved are often measured in Amps, these filters can be physically large, heavy, and expensive.

Small linear power supplies have relatively high impedances in their mains transformers, which spreads their pulses of supply current in time and so reduces their harmonic content. They sometimes meet the harmonic limits without modification. As power transformers get bigger their impedance drops and the resulting current pulses into their bridge rectifiers are sharper and contain more troublesome harmonics. Larger linear supplies therefore emit harmonic currents as readily as do switch-mode power supplies, which no transformer between their bridge rectifier and their unregulated DC storage capacitor.

One solution is to add a series inductance either before or after the bridge rectifier, as in Figure 6W. This widens the conduction angle of the rectifiers and so reduces their harmonic emissions. The lowest harmonics are realised when the choke has a very large inductance, but these can be comparable in size with a mains transformer rated for the product’s full power.

Figure 6W Passive PFC with series DC inductor

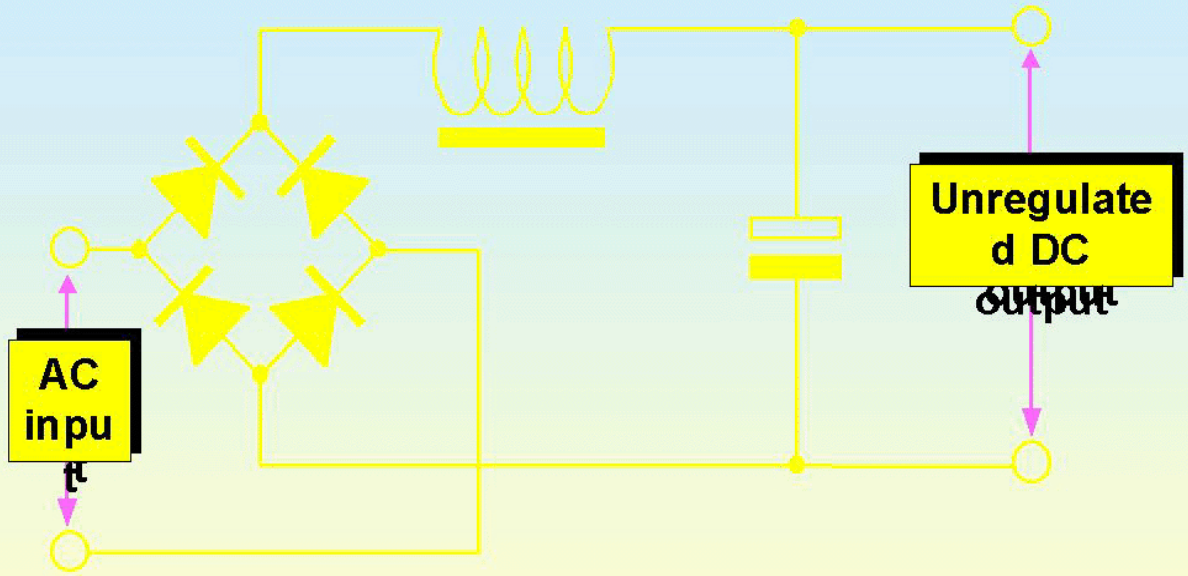

The inductor makes the rectifiers conduct for longer, reducing harmonic currents. The inductor must be sized so its flux does not become discontinuous. As the inductor is used on DC, care must be taken to ensure it does not saturate

It appears that single-phase rectifiers with constant-inductance choke input filters can be designed to meet the toughest harmonic limits in EN 61000-3-2 for power ratings <1500W. Choke values of 7mH to 70mH are approximately required for direct-on-line voltage-doubler rectifiers, the higher values applying at lower powers and currents. Transformers with the requisite inductance could be made, to save adding a separate component for applications where low-voltage high-current supplies are required and switch-mode techniques are not favoured.

The inductor in series with the unregulated DC capacitor resists rapid changes in current, so makes the rectifiers conduct for longer – reducing harmonic currents. The inductor must be sized so its flux does not become discontinuous at any point during a cycle. As the inductor is used on DC, care must be taken to ensure it does not saturate. The large air-gaps used to prevent saturation will emit quite strong ‘hum’ magnetic fields locally and if the inductor is not shielded these may affect the placement of other devices and the routing of cables and PCB traces nearby. It may be possible to adapt this circuit to put the inductor on the AC side, so that it sees an AC current and saturation is less of a problem.

With the right kind of switch-mode power supply (SMPS) controller it is possible to use passive ‘charge pump’ circuitry to correct power factor and reduce harmonics, as shown in Figure 6X.

Figure 6X Example of an SMPS with a PFC charge-pump

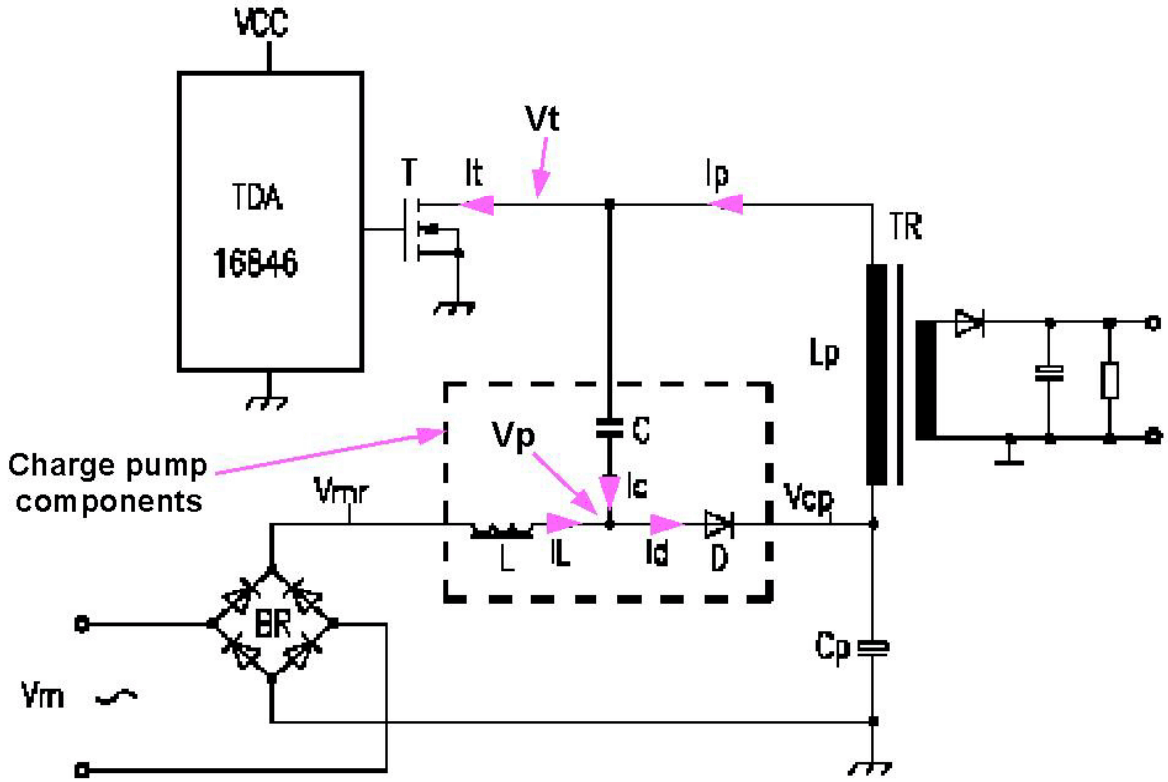

An example is the Infineon TDA 1684X family (e.g. TDA 16846, TDA16847). Their load dependent frequency response for free running SMPS's allows them to control sinusoidal switched mode power supplies featuring a PFC charge pump circuit (as well as provide a low power standby operation). For more on this technique refer (for example) to the Infineon Application Note: AN- TDA 1684X (version 1.2, dated June 2000).

Traditional SMPS converters do not have the charge-pump circuitry shown in the dotted box of Figure 6W. With the charge-pump circuitry, the current (IL) in the charge pump choke (L) becomes higher with a longer switch-on time of the switching device (T). The controller operates so as to make T's switch-on time increase with a decreasing mains voltage, thereby increasing the current through the PFC charge pump circuit. With suitably dimensioned components this makes the load on the AC mains appear resistive. The additional stored energy in capacitor Cp makes the SMPS capable of stabilising the ripple voltage during mains zero crossings.

Adding the charge-pump circuitry does not worsen any other EMC emissions (unlike the 'Active PFC' technique described later). But it does increase the dissipation in T by a significant amount and so can worsen efficiency and make larger devices and/or heatsinks necessary. The waveforms of Figure 6Y show the various currents and voltages in the charge pump and SMPS circuit for an example cycle, to help understand the operation of the circuit.

Figure 6Y One cycle of the example SMPS

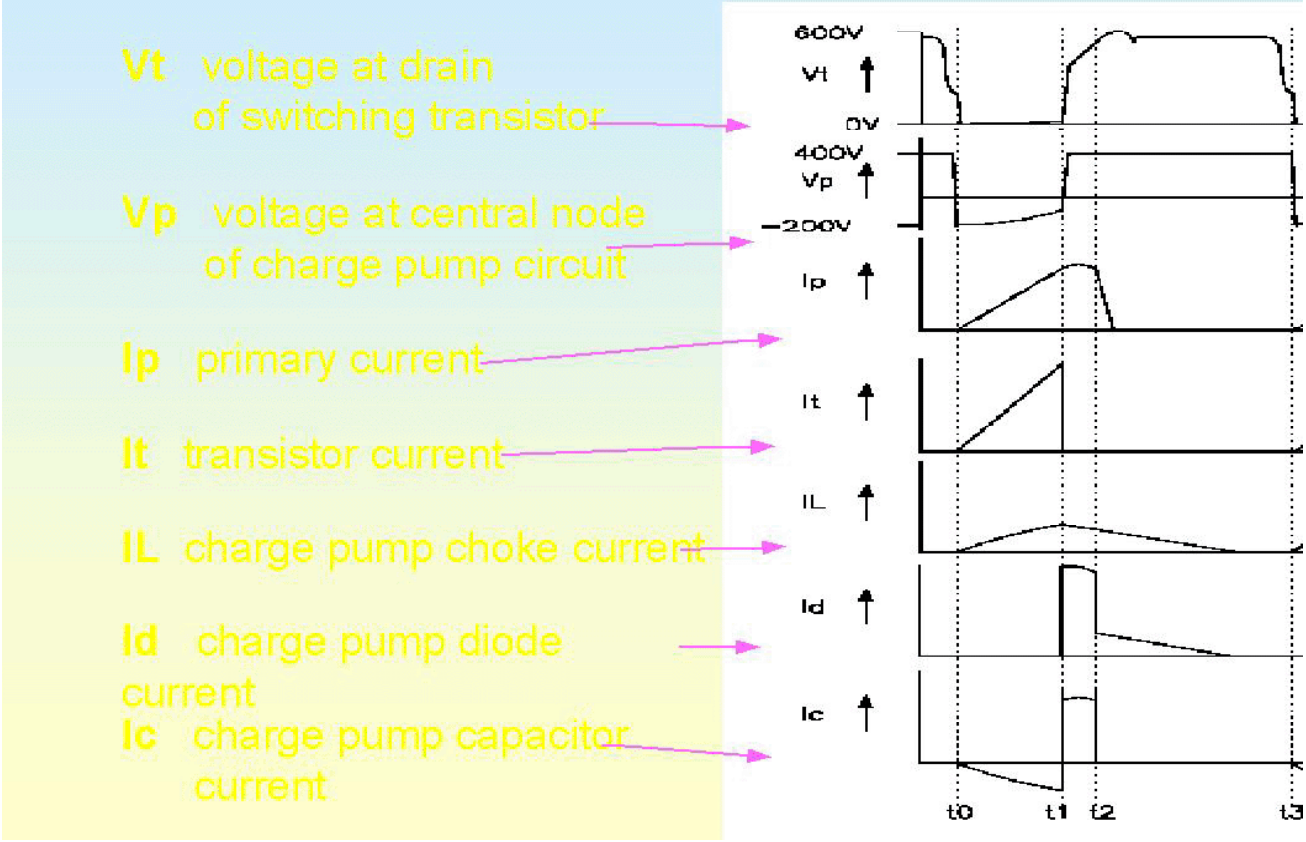

Different waveforms obtain under different conditions of supply voltage and load current, although they still retain the same principle (for more on this refer to the Infineon Application Note: AN- TDA 1684X (version 1.2, dated June 2000).

The charge pump's capacitor current peaks when the voltage across T is high. This considerably increases the dissipation in the switching transistor T, when compared with the same converter without the charge pump, requiring larger devices and/or larger heat sinks for reliable operation. However, it is possible to feed the charge pump's capacitor from a separate winding on the transformer instead of from the switched primary winding. This can be designed in such a way as to reduce the dissipation in T to close to what it would be without a charge pump – saving cost and PCB area.

The 'Active PFC' technique interposes a switch-mode boost converter between the bridge rectifier and the storage capacitor, as shown in Figure 6Z. It boosts the full-wave rectified supply when it is lower than the voltage on the storage capacitor.

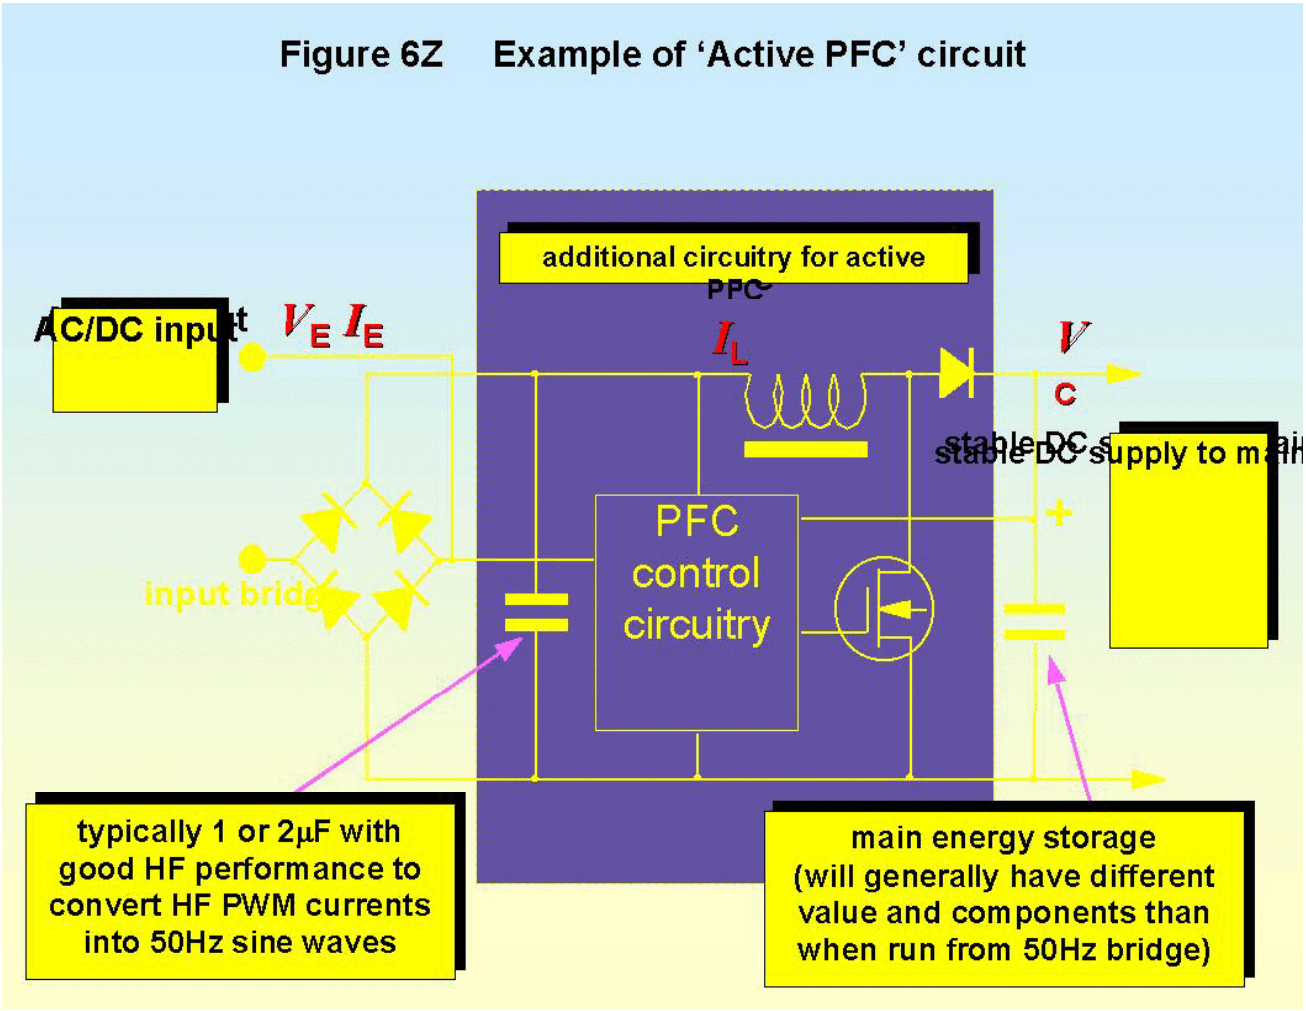

The rectified mains voltage is boosted under the control of an IC that causes the current into the storage capacitor to approximate a full-wave rectified sine wave. Thus the storage capacitor plus active PFC circuit looks like a resistive load to the bridge rectifier, so the whole lot appears to the AC supply like a resistor (although with little notches around zero-crossings due to the rectifiers). The boost circuit typically operates at a high frequency, even several MHz, so a filter capacitor (usually around 1μF) is required to convert the fast-switching current pulses into a reasonably-looking rectified sine wave.

Active PFC circuits have a time constant of around 0.5 seconds to smooth out load current fluctuations so that their AC supply's simulated resistance appears to change in value slowly, not to cause harmonic emissions. This needs to be taken into account when designing the size of the DC storage capacitor. The PFC boost circuit used generates high-frequency conducted and radiated emissions which need suppressing. Usually these boost circuits are added to existing switch-mode power supplies where some filtering and shielding will already be in place, although they are likely to need modifying.

Figure 6AA is taken from an EPCOS application note on their active PFC products. The waveforms shown for VC, IL, VE and IE correspond to the points marked on the overview circuit on Figure 6Z. Figure 6AA shows how the inductor current is pulse-width-modulated and smoothed by the input capacitor to simulate a sine-wave input current. The voltage on the DC storage capacitor will have 100Hz ripple on it (or 120Hz).

Figure 6AA Active PFC example waveforms

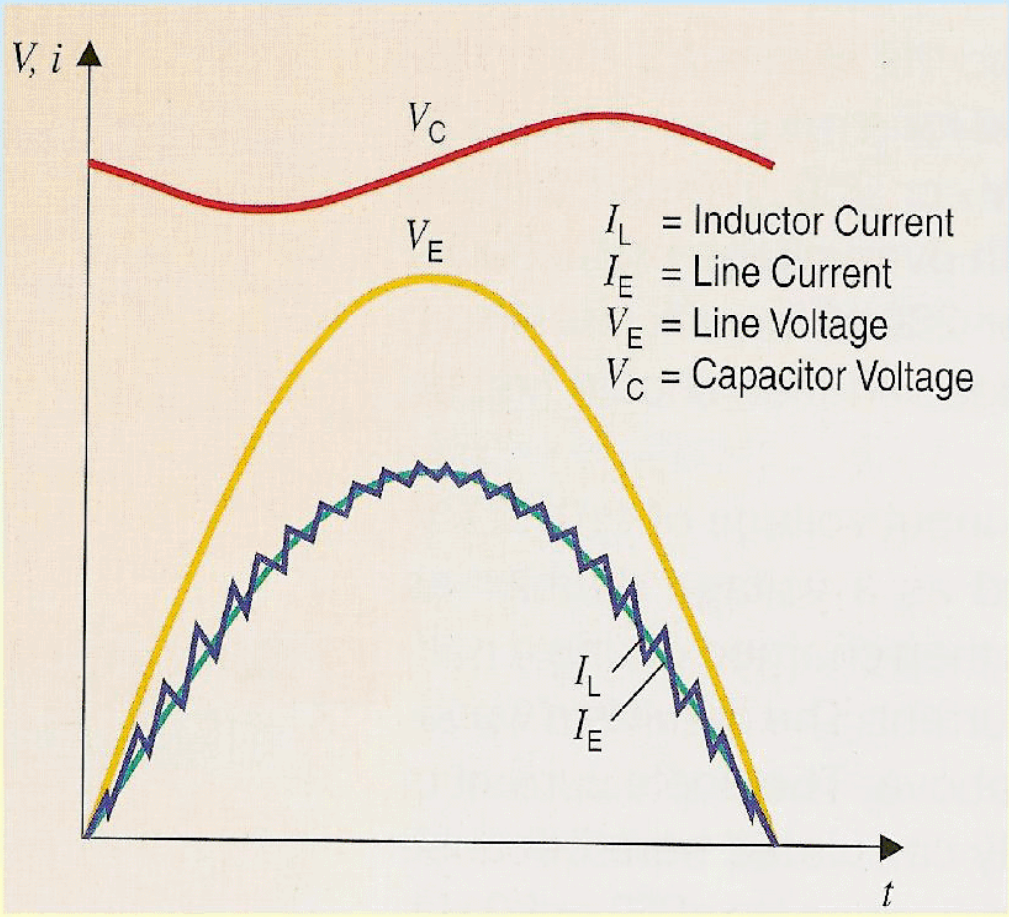

Most PFC circuits will be followed by voltage regulation, usually another switch-mode converter. Control ICs are now becoming available that will manage both functions: PFC and regulation.

Figure 6BB shows an example application circuit for an EPCOS PFC controller, complete with input filtering and surge protection.

Figure 6BB Example 'Active PFC' circuit from EPCOS

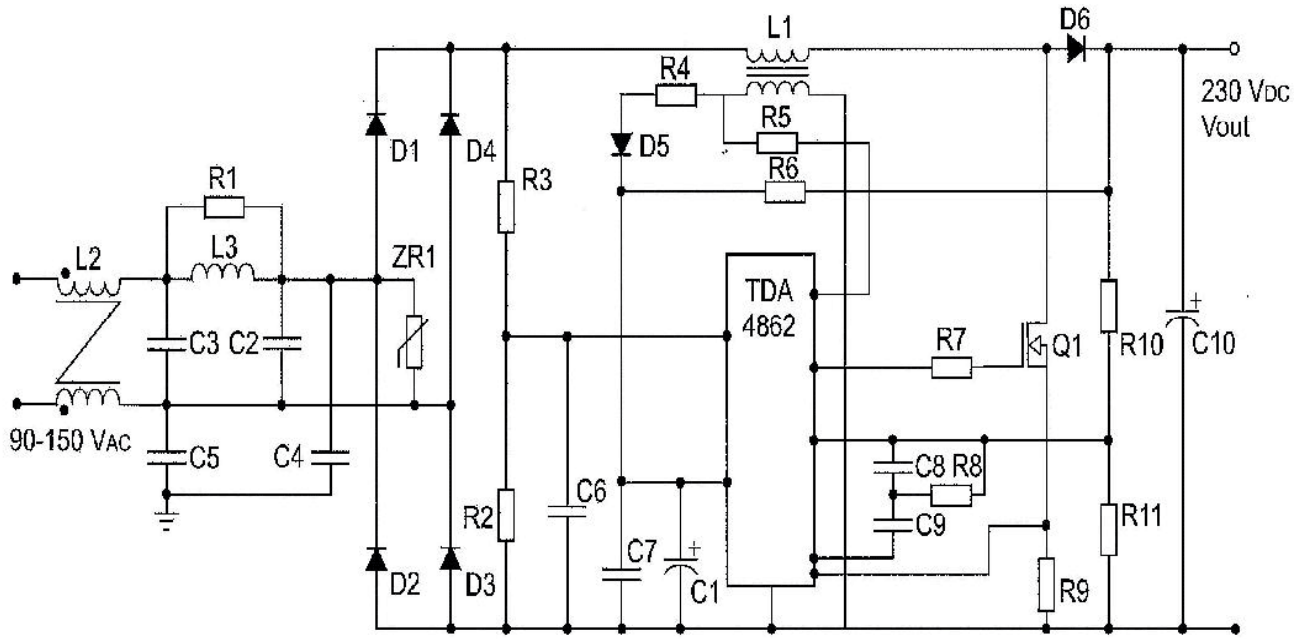

Active PFC has many advantages apart from helping to meet harmonic emissions standards, including...

Universal input

Having active PFC boost circuitry makes it easier to design power supplies which will cope with a wide range of supply voltages, such as 85 to 264V, DC to 400Hz, allowing operation from the public low-voltage mains (and some aircraft supplies) anywhere in the world.

This reduces the need for 'country variants', reduces stockholding, and allows faster order fulfilment.

Full power available from wall sockets

Active PFC can draw full power from wall sockets (e.g. 3kW from a UK 230V 13A plug, 1700W from a US 110V 15A plug). Previously they were typically limited to < 50% of the possible power before the fuse in the plug or the building's over-current protection operated.

Waveform insensitivity

Some countries, especially in the developing world, can have very badly distorted supply waveforms. 'Traditional' power supplies are sensitive to crest factor (the ratio of peak to RMS) so do badly on distorted supplies. Active PFC helps cope with these situations.

Universal-input PFC supplies help cope with dips, sags, brownouts, and dropouts

Mains supplies everywhere suffer from dips, sags, brownouts, and dropouts. These can cause microprocessors to reset frequently, causing annoyance to users even if they don't cause data to be lost or control systems to go out of control. A universal-input power converter running on a nominal 230V will run perfectly well, keeping its storage capacitors fully charged, on -50% mains.

The harmonic emissions from three-phase power rectifiers (AC-DC converters) can be reduced by making them six-phase. Six-phase rectifiers are often called 12-pulse converters and use the fact that the phase outputs from star and delta wound 3ϕ transformer secondaries are 60^o apart, so use both star and delta secondaries each feeding a three-phase rectifier which feeds into the DC rail, as shown by Figure 6CC.

Figure 6CC A 6ϕ bridge (12-pulse) rectifier

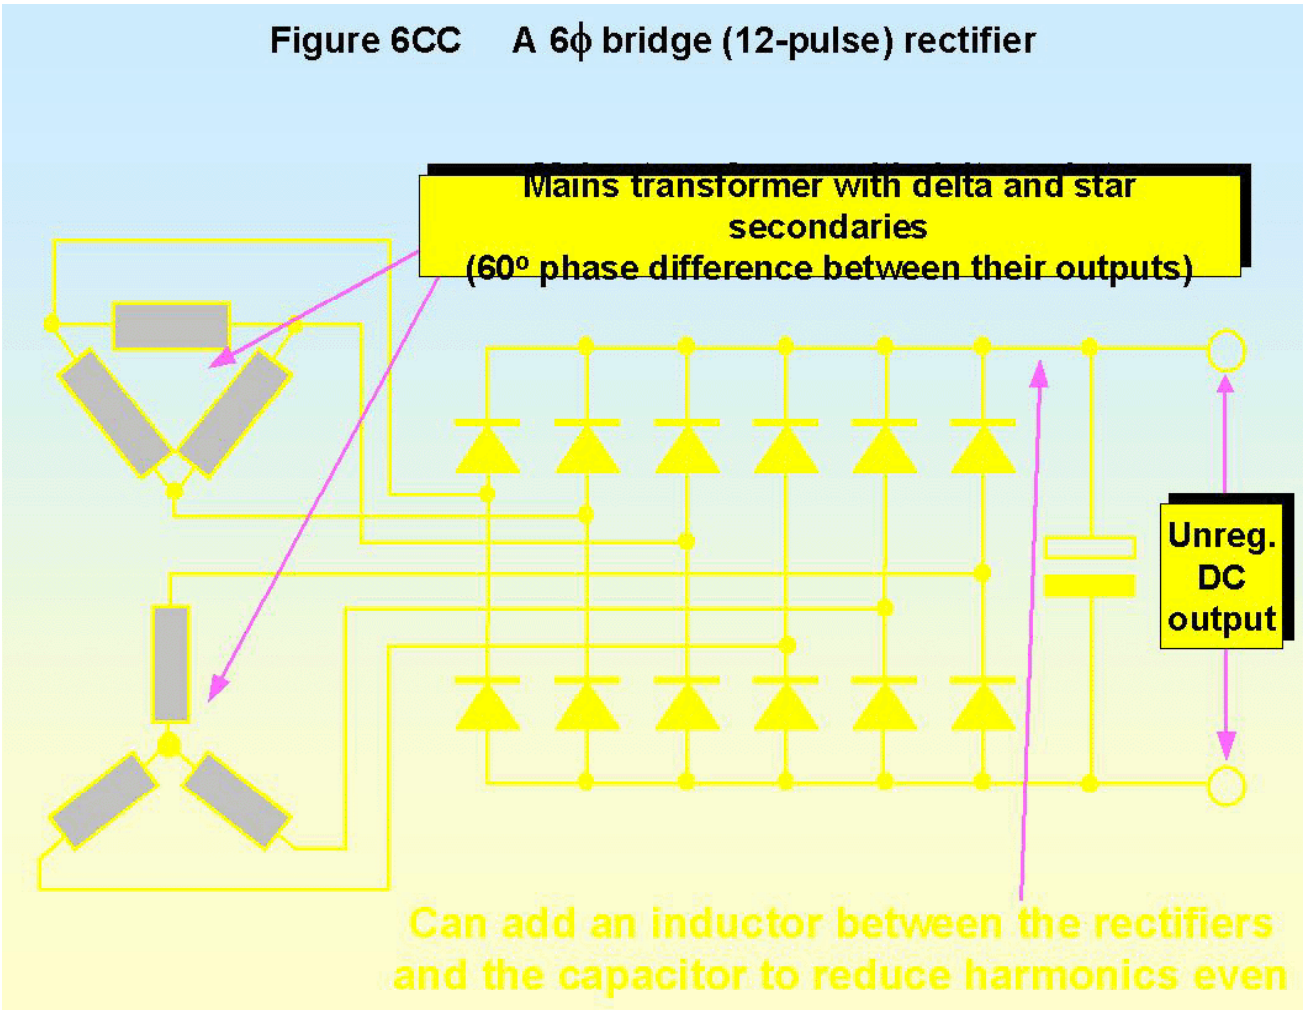

Like three-phase converters, six-phase converters naturally have low levels of triplen harmonics (into balanced load) but the conversion to six-phase reduces their emissions of 5th and 7th harmonics too. A series inductor can be added, in the AC or DC sides, to reduce the harmonics even more (these are often called 'line reactors' or 'DC link reactors') and operate in the same way as the series inductors described earlier, but as they do not have to deal with very high levels of current ripple their design constraints are eased.

Three-phase 'boost' rectifiers use PWM control of power transistors (usually IGBTs) instead of ordinary rectifiers, using 'rectifier' circuits similar to that of Figure 6DD. When the switching rate of the PWM is much faster than 50Hz the switching functions can be replaced by their average values in the switching period.

Figure 6DD A 3 phase 'boost' rectifier

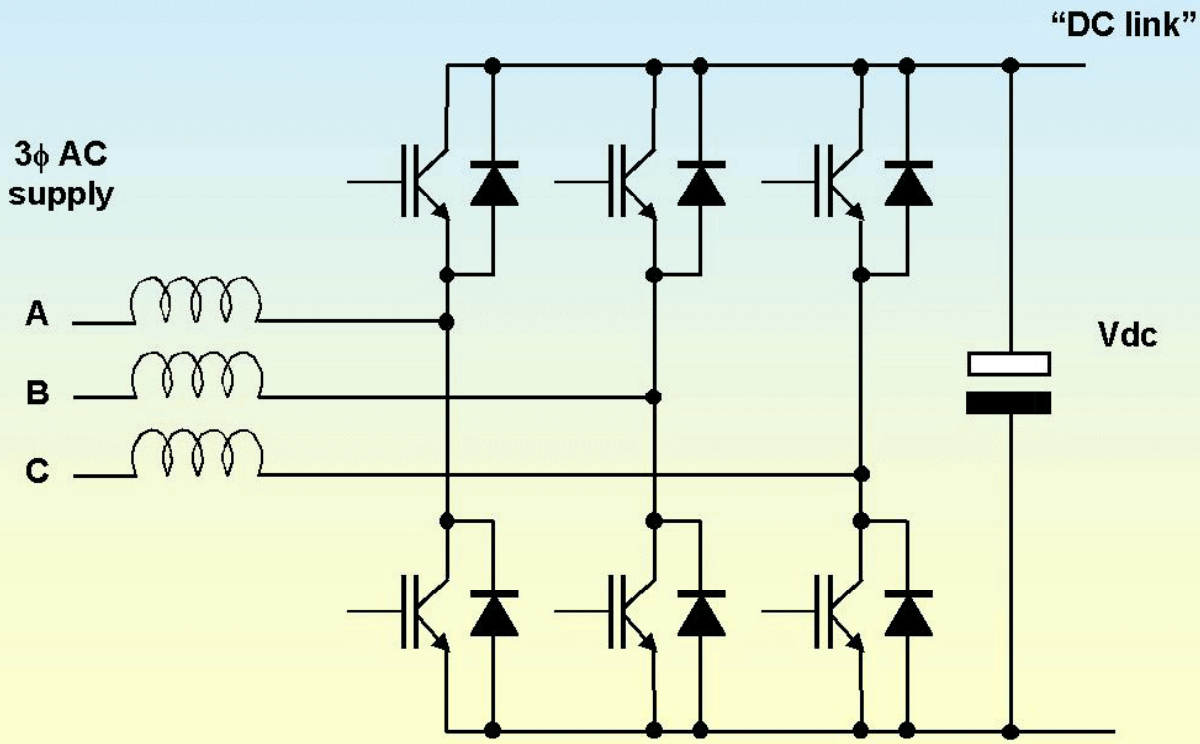

By appropriate control of the PWM switching patterns these boost rectifiers can achieve a near sine-wave input current from the mains supply, hence PFC. They can also achieve bi-directional power transfer. They have a tendency to become unstable which is usually overcome by using a large value for the DC capacitor.

7. System level techniques

There are mains harmonic reduction techniques which can be used at the level of the system or installation, but of course they don't help equipment meet EN 61000-3-2.

Some harmonic problems in systems and installations are simply dealt with by up-rating conductors (neutrals may need double the cross-sectional area of the phase conductors) and transformers.

Star-delta transformers help remove the triplen harmonics typical of single-phase mains distributions. Their loads must be well balanced and they must be rated for the appropriate levels of zero-phase magnetic flux in their delta windings

Filters (series and parallel types) can restrict the flow of harmonic currents in areas of the mains distribution. Filters are 'tuned' to each problem frequency, and can interact unpredictably with existing network resonances. Filters should be left to power system harmonics experts.

'Active filters' are now available which are much easier to use than passive filters. A better name for these types of equipment is 'active harmonic conditioners' as they are not really filters but 'anti-harmonic generators' that take energy from the mains supply and add it back in as necessary to preserve its sine waveform. They only use the energy associated with their internal inefficiencies,

since on over periods of one second or more the power they need to extract to preserve the waveshape is exactly equal to the power they need to inject.

Uninterruptible power supplies (UPSs) can be used providing they themselves have sufficiently low levels of harmonic emissions into their mains supply (check manufacturers' data carefully).

Motor-generator sets are a traditional way of confining harmonic currents to mains conductors which are isolated form the public supply.

- Operating equipment from its own low-voltage distribution transformer is another traditional technique. The low impedance at the common point of connection to the high voltage supply reduces the effect of the harmonic currents on the public low-voltage network.

8. Conclusion to the Series

A textbook could be written about any one of the topics covered above, and I am keenly aware of the topics I have left out or dealt with in a rather brief manner. I am sure many specialists are wincing at the treatment I have given to their subjects. Many people have written to me with useful comments and suggestions, for which I am most grateful. I especially acknowledge the huge contribution made by Tomonori Sato since I wrote the first versions in 1999.

The aim of this series was to introduce the EMC non-specialist to the various issues and point to the most important and commonly-used best-practice design techniques, and I hope I have achieved this.

Theoretical background has been minimal, where it existed at all, despite the fact that an understanding of the theory behind EMC helps designers to deal with the unexpected (of which EMC often appears to have more than its fair share). The books recommended in Part 5 will be found useful in obtaining this knowledge. Many more EMC textbooks and guides exist for the interested reader.

Many of the techniques described in this series are also important for improving signal integrity. They are thus powerful competitive weapons in reducing the number of iterations during development, reducing manufacturing costs, and improving reliability in the field and enjoying lower levels of product returns and claims under warranty. They may also be helpful in reducing a company's risk under product liability legislation. Even if the EMC Directive did not exist in the EU, or the FCC in the USA, the techniques described in these six articles would still be worth using for the real engineering and financial benefits they bring.

Eur Ing Keith Armstrong C.Eng MIEE MIEEE
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone: 01457 871 605, Fax: 01457 820 145, Email: keith.armstrong@cherryclough.com