Design Techniques for EMC Part 3: Filters and surge protection devices

Design Techniques for EMC
Part 3: Filters and surge protection devices

Eur Ing Keith Armstrong CEng MIET MIEEE ACGI BSc(hons)
Partner, Cherry Clough Consultants, www.cherryclough.com, Associate of EMC-UK
Phone & fax: (+44) (0)1785 660247, Email: keith.armstrong@cherryclough.com

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 Part 3

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 that 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**
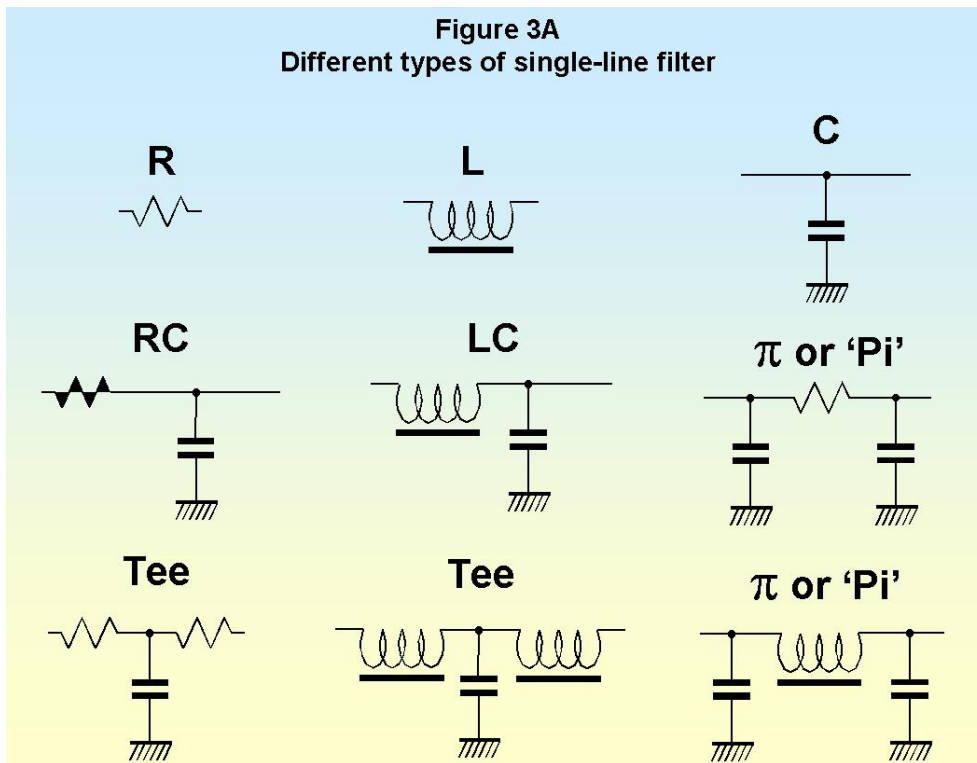
Resistors eventually lose their performance at high frequencies due to their stray shunt capacitance. Inductors also suffer from stray shunt capacitance, and this causes self-resonance and limits their high-frequency performance. The best inductors for EMC filter use have closed magnetic circuits (such 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 10kΩ) 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 that 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 3dB 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 versus common-mode "leakage"
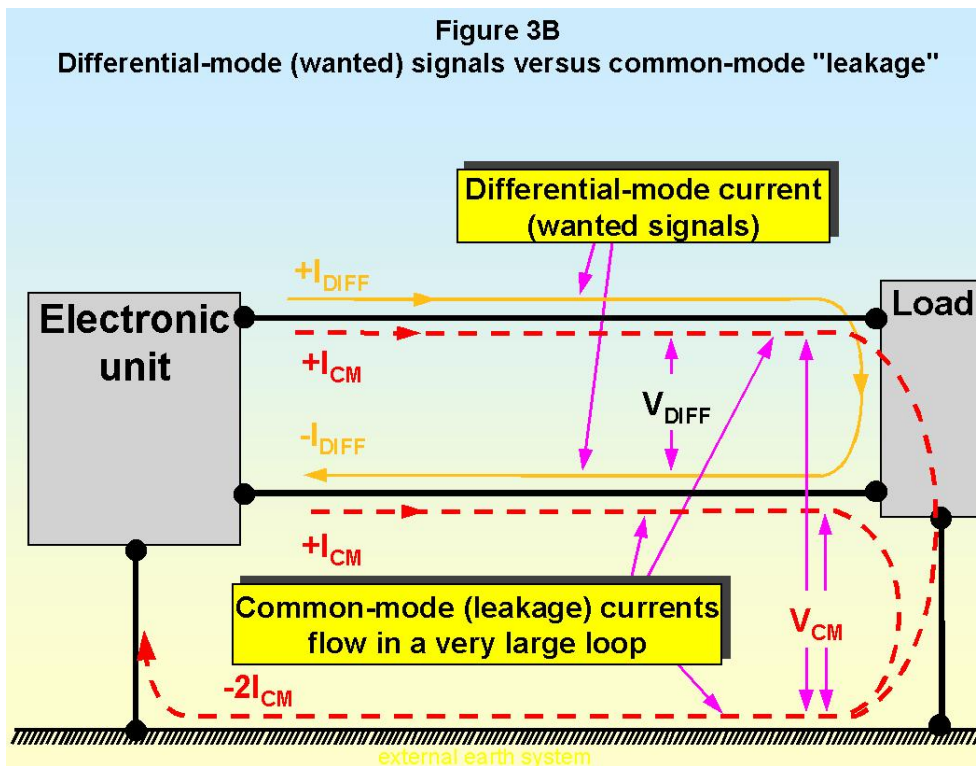
On the other hand, CM voltages and currents are identical for all the conductors in a cable (including any screens), and all CM currents flow in the same direction. The return path for CM currents is via the external earth system, other cables, metalwork, etc. CM currents and voltages are often very 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
CM filtering

A 2-circuit CM choke with ½ turn on it

A 3-circuit CM choke with 1½ turns on it

CM choke's attenuation

Cable-fitted chokes shown above
PCB-mounted chokes use same basic construction
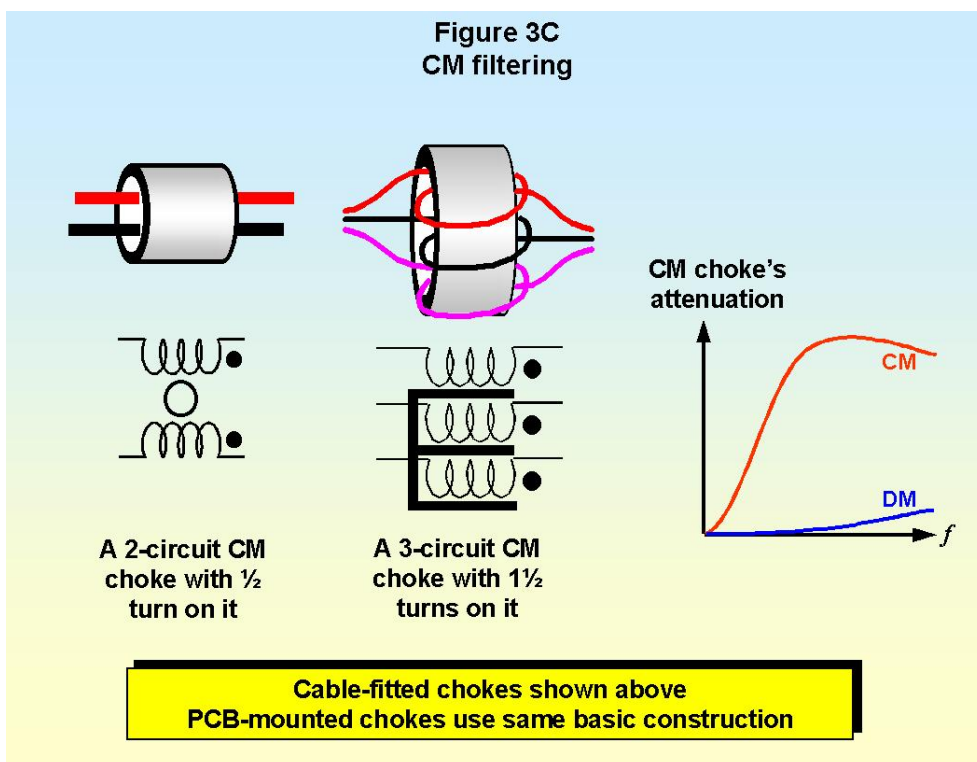
For most EMC work, common-mode chokes are made from soft ferrite. The cancellation of DM flux in CM chokes (often called a current-balanced choke) allows large inductance values (milliHenries) to be achieved with small components, whereas DM chokes the same size are measured in microHenries, and become 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 – Typical *single-stage* filters using CM and DM filtering techniques
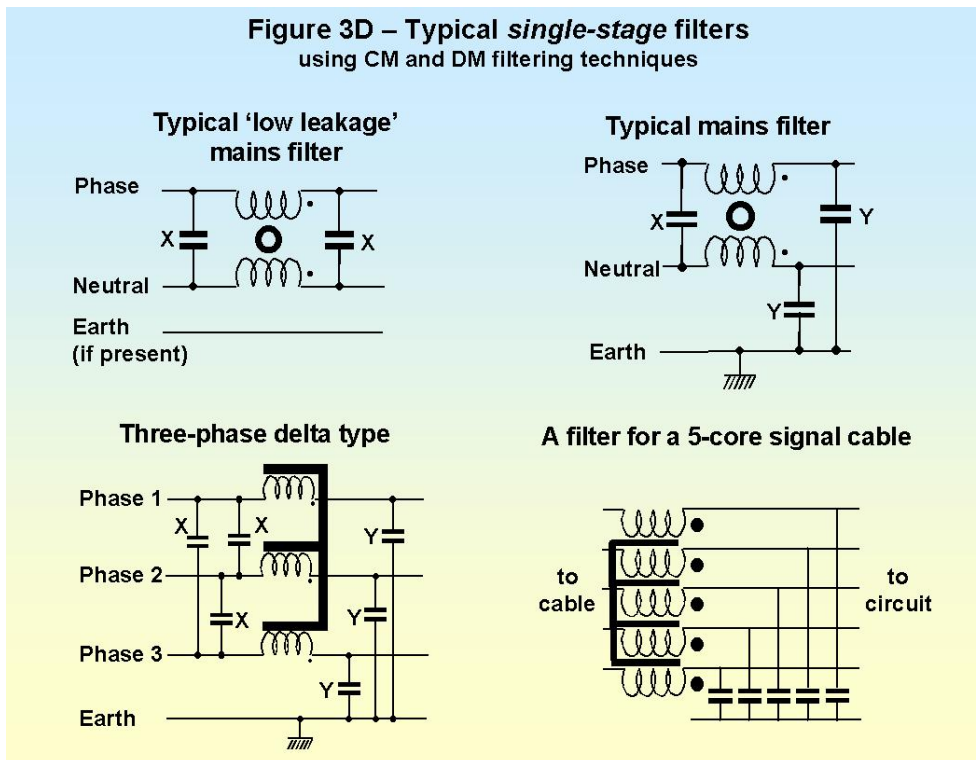
3.5 Crude filter rules-of-thumb

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 2Ω 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.
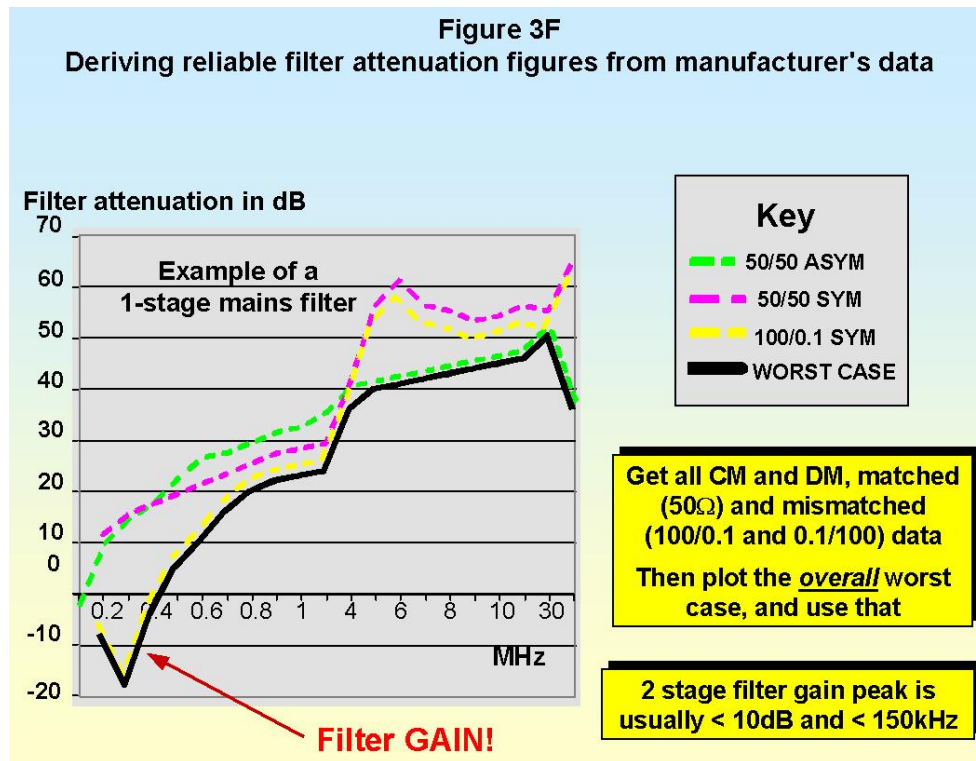
Filters with a single stage (such as those in Figure 3D) are very sensitive to source and load impedances. Such filters can easily give *gain*, rather than attenuation, when operated with source and load impedances other than 50Ω. This *filter gain* usually pops up in the 150kHz to 10MHz region and can be as bad as 10dB 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 50Ω/50Ω specifications. Of course, they are larger and cost more.

Figure 3E– A typical 2-stage mains filter
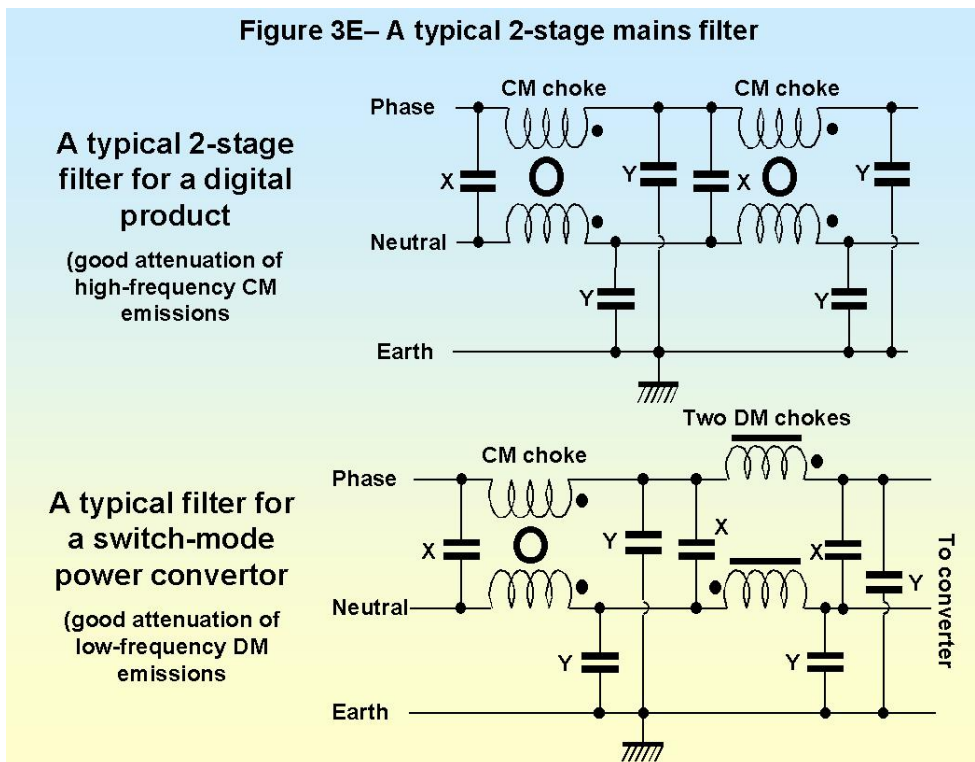
The best way to deal with the source/load impedance problem is to only purchase filters whose manufacturers specify both CM (sometimes called "asymmetrical") and DM (sometimes called "symmetrical") performance, for both matched 50Ω/50Ω and mismatched sources and loads.

Mismatched figures are taken with 0.1Ω source and 100Ω load, and vice versa, using the CISPR17 test standard that is also used for 50Ω/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

Filter attenuation in dB

Example of a 1-stage mains filter

Key
50/50 ASYM
50/50 SYM
100/0.1 SYM
WORST CASE

Get all CM and DM, matched (50Ω) and mismatched (100/0.1 and 0.1/100) data

Then plot the *overall* worst case, and use that

2 stage filter gain peak is usually < 10dB and < 150kHz

Filter GAIN!

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 that 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 that 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 surface, 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.

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.

Figure 3G   Mounting mains filters in shielded enclosures
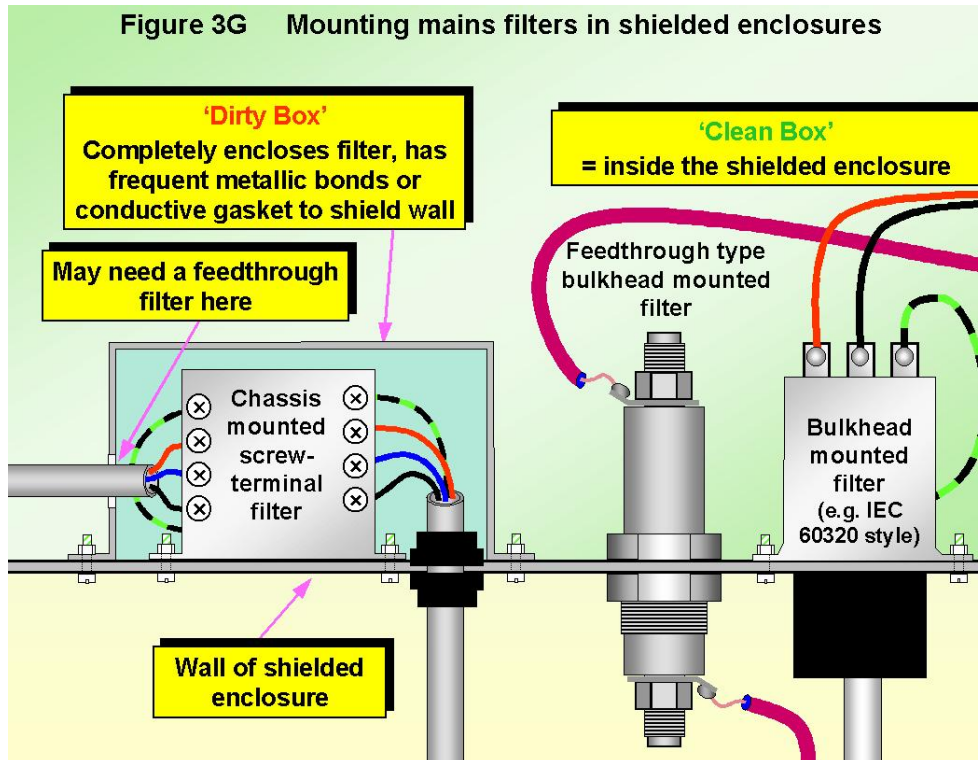
The filter input and output cables in the dirty box must be very short and far away from each other, but even so, very high frequencies may still couple between them and soft-ferrite cylinders may be needed on either (or both) cables.

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
Time-domain performance of SPDs

The SPD's 'let-through' voltage

Rising edge of surge without SPDs

Trigger voltage (for GDT and SCR types)

MOV

Avalanche

GDT, SCR

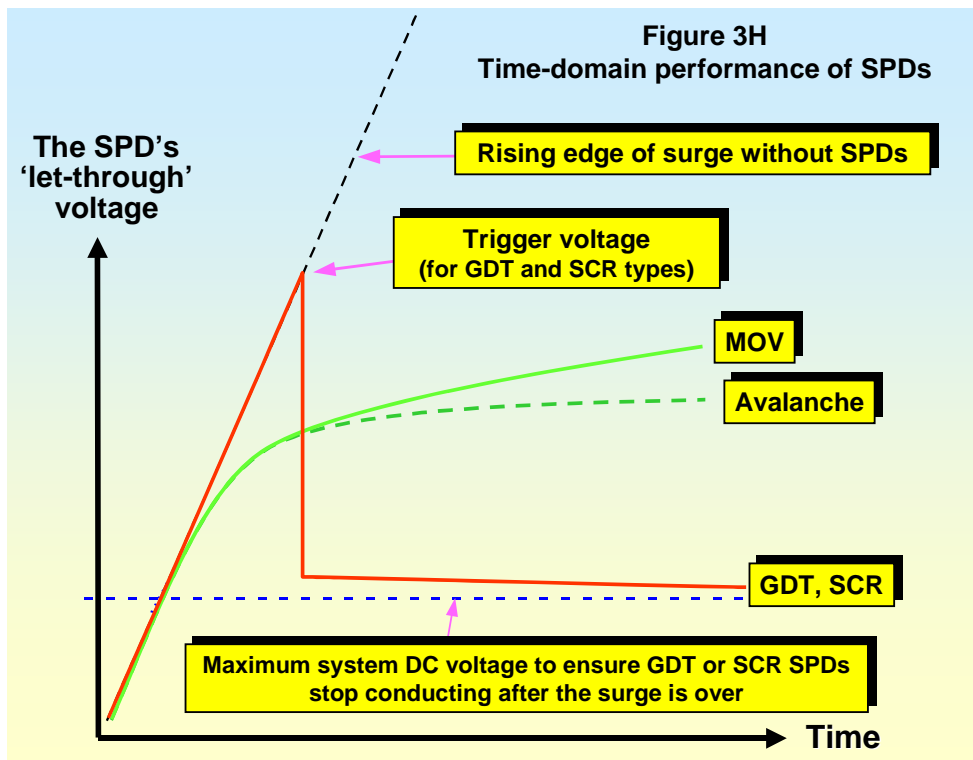Maximum system DC voltage to ensure GDT or SCR SPDs stop conducting after the surge is over

Time

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 that 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 forever.

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 that 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, 2,000 Amps
- 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, 5,000 Amps
- 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 that 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 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.

Figure 3J   Example of fitting SPDs to the switch-mode power supply of a portable product
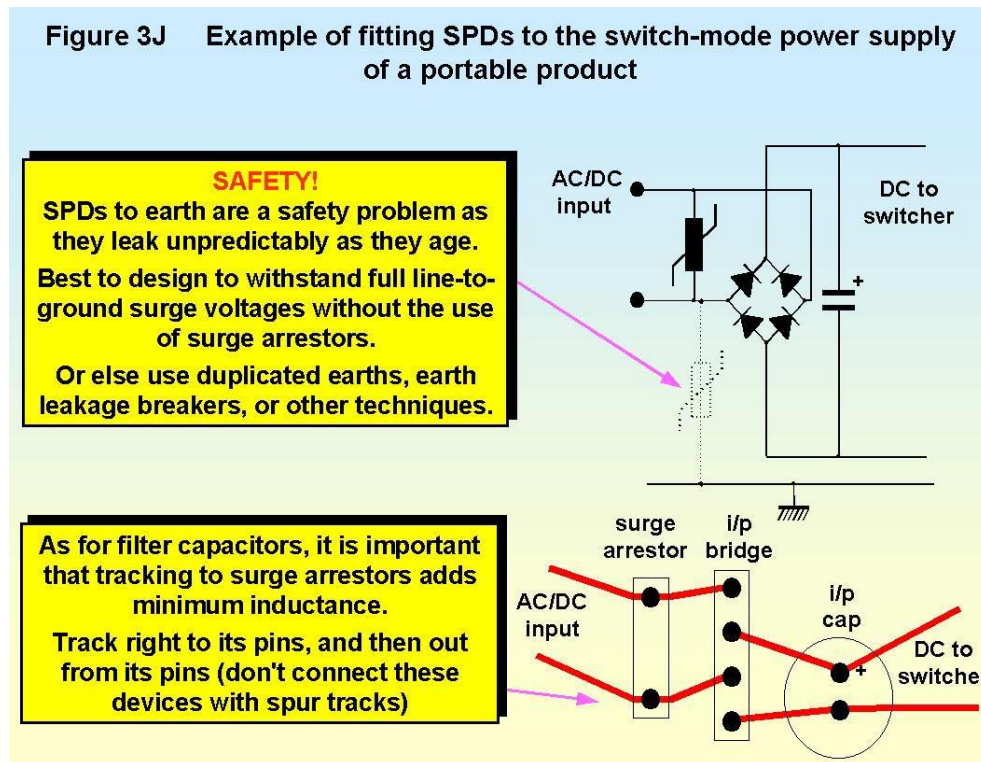
However, SPDs between supply phases and ground may be acceptable for permanently-wired equipment, especially if it has duplicated protective earth conductors, so that if one is faulty the other prevents electric shock from any SPD earth leakage. SPDs are commonly connected phase to ground in building 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:

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

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, April 2000, pp8 – 17. Available from the 'Publications & Downloads' page at www.cherryclough.com