Filtering and suppression

Section 9

1

Outline

- filter configuration and layout
- components
- ferrites
- I/O filtering
- mains filtering
- transient and motor suppression

2

Noise being conducted out of or into equipment along connecting leads will bypass and render ineffective any attempts at screening the circuit. It is unfortunately impossible to completely eliminate it. The purpose of filtering at interfaces is to attenuate such noise to a level either at which it meets a given specification, for exported noise, or at which it does not result in malfunction of the system, for imported noise.

If a filter contains lossy elements, such as a resistor, a ferrite component or a transient suppressor, then the noise energy may be absorbed and dissipated within the filter. If it does not, then the energy is reflected and must be dissipated elsewhere in the system, typically in the impedance of the source or as re-radiation from the incoming cables. Absorptive filters are generally to be preferred for EMC applications.

Filtering is not confined just to interfaces. It is often helpful to limit the bandwidth or the induced voltages at particular circuit nodes. Suppression of motors and inductive loads is one example where interference is prevented at source.

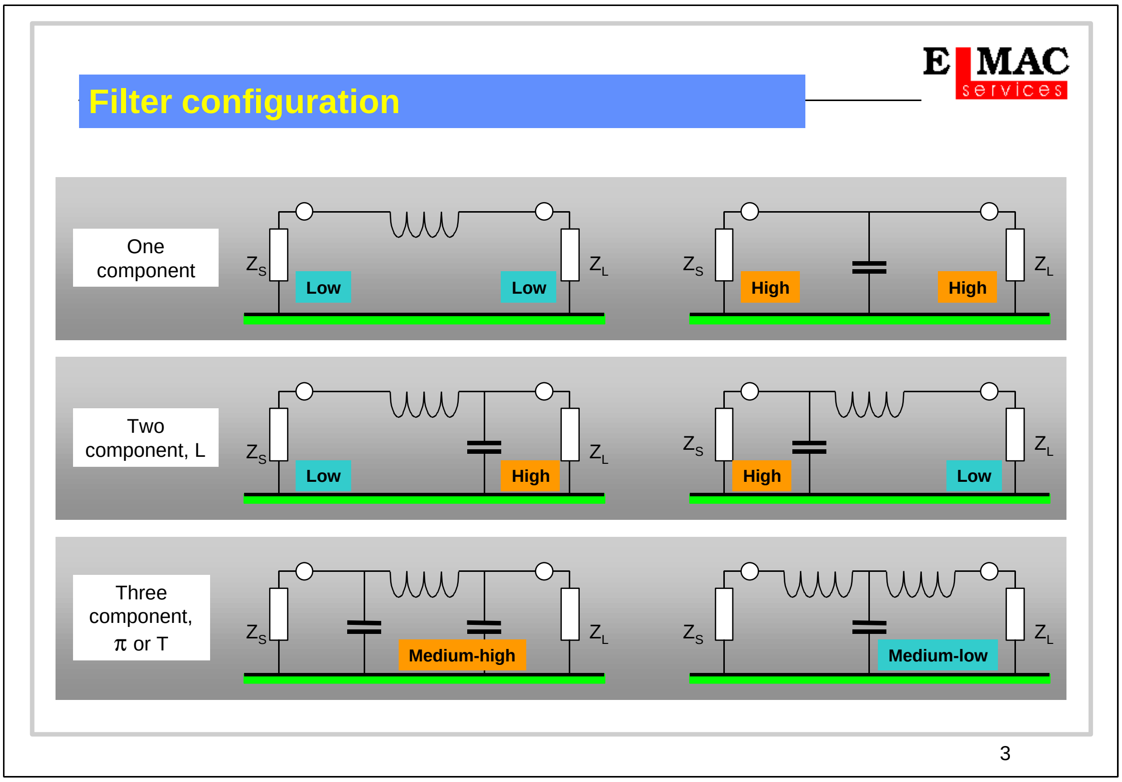

In EMC work, “filtering” almost always means low-pass filtering. The purpose is normally to attenuate high-frequency components while passing low-frequency ones. Various simple low-pass configurations are shown above, and filter circuits are normally made up from a combination of these.

The effectiveness of the filter configuration depends on the impedances seen at either end of the filter network.The simple inductor circuit will give good results - better than 40dB attenuation - in a low impedance circuit but will be quite useless at high impedances. The simple capacitor will give good results at high impedances but will be useless at low ones. The multi-component filters will give better results provided that they are configured correctly; the capacitor should face a high impedance and the inductor a low one.

Conventionally, manufactured filter assemblies have their attenuation specified for terminating impedances of 50Ω at each end. In the real application, Z_S and Z_L are complex and perhaps unknown at the frequencies of interest for suppression, which makes accurate design of filter properties largely academic. There is usually a limitation on maximum component values as a result of size or circuit performance, and the actual L and C components will be chosen to be some figure less than this.

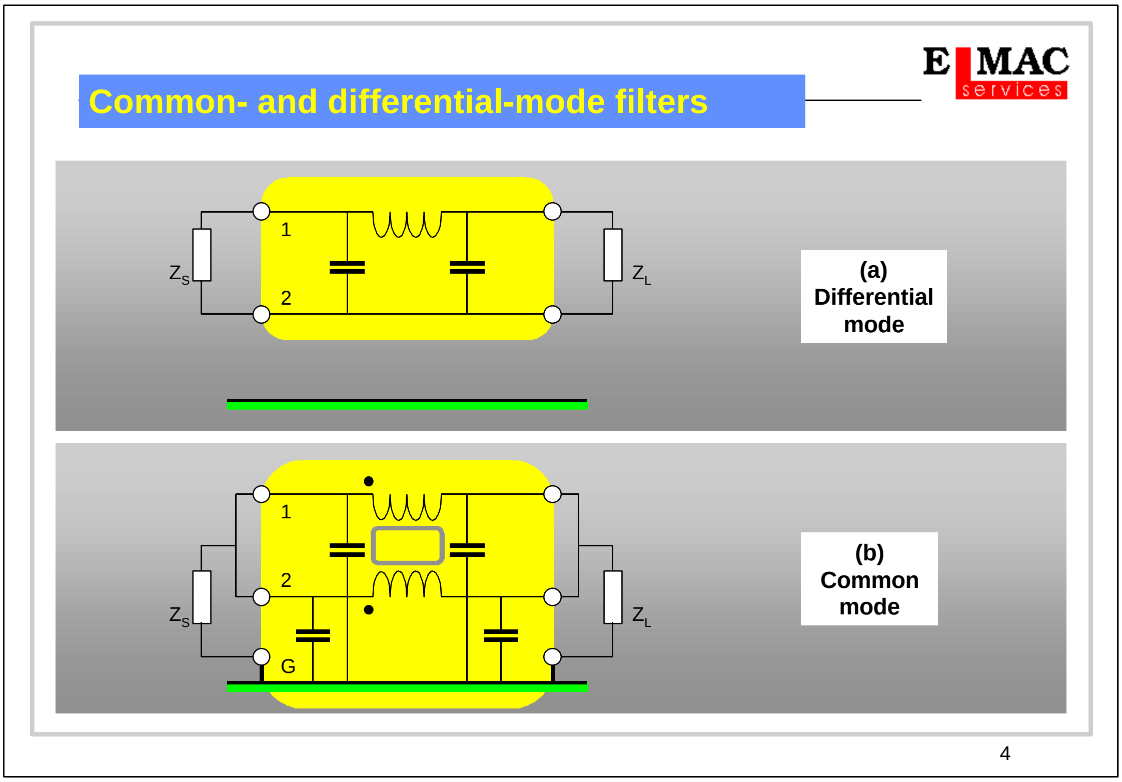

The distinction between the two modes of interference coupling has been covered in section 3. Filters to attenuate these modes must be configured appropriately.

The filter shown at a) will attenuate interference which appears between terminals 1 and 2, that is, differentially. It will have no effect on interference which appears in common mode between terminals 1 or 2 and ground, since there is no parallel capacitance to ground, and there is a straight-through path via terminal 2.

The common mode filter shown at b) will attenuate interference appearing between terminals 1 + 2 together, and ground. It may also have a lesser effect on differential mode interference; the differential circuit effectively sees a balanced π-filter with capacitance values of 0.5C and an inductance value equal to the leakage inductance of the common mode choke. This may be replaced by two separate chokes in order to increase differential attenuation, at the expense of limiting the bandwidth of the differential circuit.

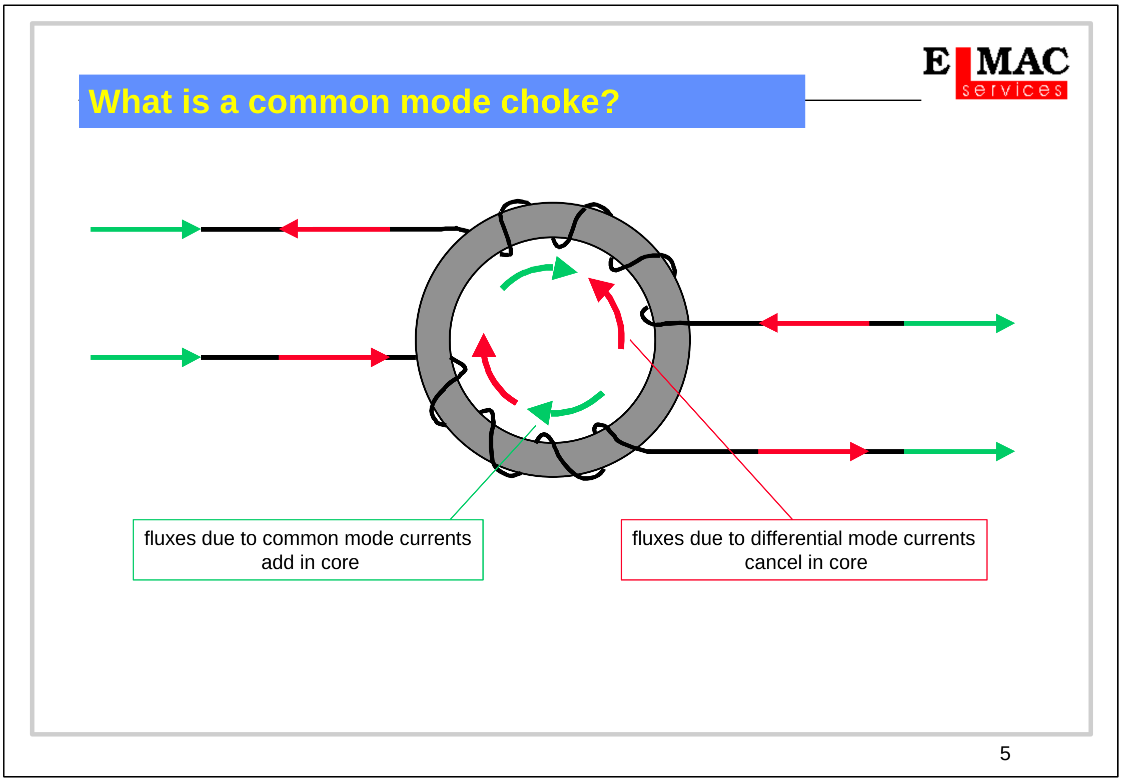

A common mode choke is most easily described when it has two windings, one in each leg of a two wire circuit. The two windings are nominally identical and on the same core, which is often though not invariably toroidal, to minimise magnetic flux leakage.

The sense of the windings is such that differential currents, in which the “go” current in one wire is equal and opposite to the “return” current in the other, each create a magnetic flux in the core, but because they are equal and opposite the two fluxes cancel, leaving no net magnetic flux. Thus since the core is invisible the differential mode inductance is very small, being dominated by the residual difference between the windings, known as the leakage inductance.

By contrast the flux from common mode currents in the wires adds in the core, and therefore the full inductance of the choke is presented to common mode signals.

Chokes used in this way are sometimes known as “current-compensated” chokes, since as well as being invisible to differential signals, they can carry large values of low frequency or DC current without fear of core saturation and loss of inductance. Alternatively, by designing in a suitable amount of imbalance, the differential mode inductance can be tailored to provide some DM attenuation as well as CM, at the cost of a reduced current rating.

Common mode chokes can be constructed with multiple windings for filtering interfaces with many lines. As long as the fluxes due to the differential mode currents sum to zero in the core, the choke will operate correctly. Thus a winding must be inserted in series even with the 0V rail, if this carries differential return currents. Crosstalk between the circuits will be determined by the balance of the windings, and their inter-winding capacitance.

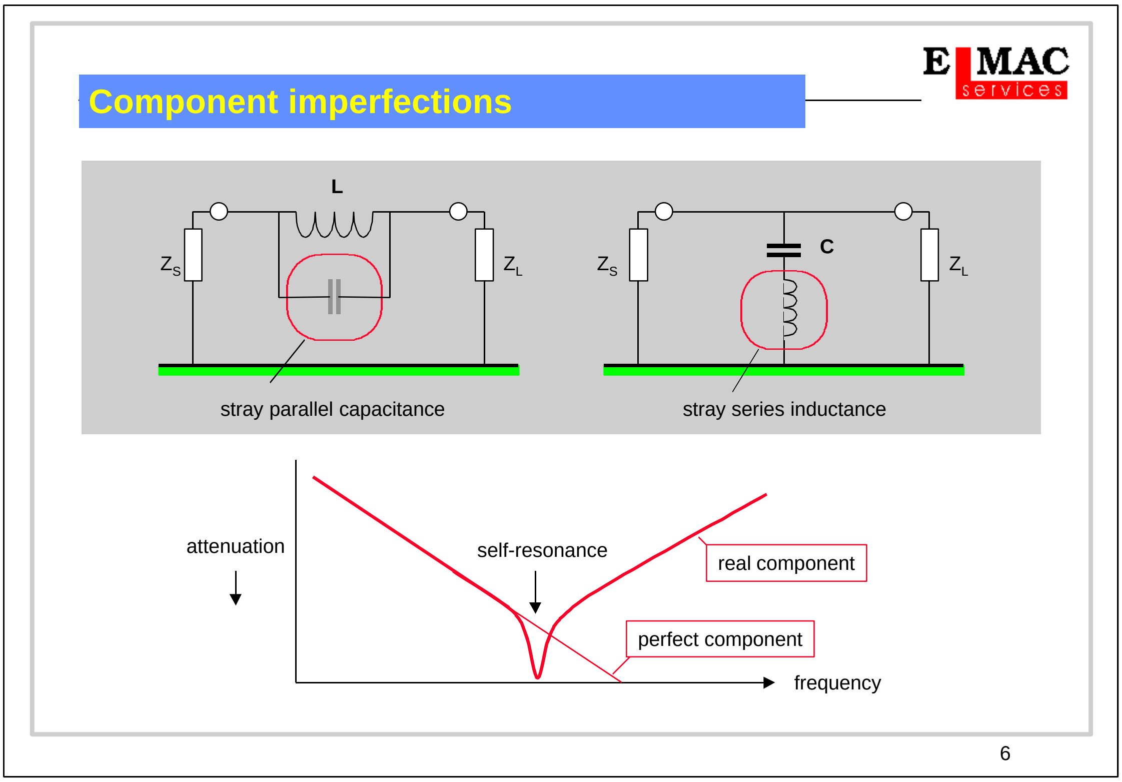

Filter components, like all others, are imperfect. Inductors have self-capacitance, capacitors have self-inductance. This complicates the equivalent circuit at high frequencies and means that a typical filter using discrete components will start to lose its performance above about 10MHz. The complementary stray reactance resonates with the main component to cause “self resonance”. The larger the components are physically, the lower will be the self resonant frequency (SRF).

For capacitors, as the frequency increases beyond capacitor self-resonance the impedance of the capacitors in the circuit actually rises, so that the insertion loss begins to fall. This can be countered by using special construction for the capacitors. Similarly, inductors have a SRF beyond which their impedance starts to fall. Particular forms of winding geometry will reduce the internal winding capacitance and raise the SRF: the most important aspect is that the start and finish of the winding should be physically separate. To obtain good performance at high frequencies, it is important to choose components carefully.

Sizing filter components

- **Capacitor value depends on:**
  - differential circuit bandwidth
  - acceptable self-resonance
  - circuit voltage and leakage current
  - allowable physical size
- **Inductor value depends on:**
  - allowable physical size
  - current rating and series resistance
  - differential circuit bandwidth
  - acceptable self-resonance

7

Accurate design of filter components, as discussed earlier, is usually not reasonable. Typical limitations on the filter components are as shown above. For capacitors, usually the most important consideration is whether the operating circuit can support extra capacitance, either across circuit nodes (differentially) or from a circuit node to ground (common mode). This is determined by the desired bandwidth and driving impedance, bearing in mind that some analogue circuits will oscillate if they are asked to drive a capacitive load. This will set a maximum value, which may then be affected by other considerations: is this value in too large a package, so that self-inductance is unacceptably high and the SRF is too low, or the package will not fit in the available space? Is the value available in the required voltage rating - sometimes an isolation of several hundred volts from earth is needed, and if there is AC in the circuit, will the leakage created by the capacitor cause a circuit malfunction? In precision analogue circuits, even the DC leakage of the capacitor may be too great.

For inductors, similar issues of bandwidth and physical size arise. Here there is usually a choice between wound components on ferrite cores, which are effective at low frequencies but have quite low SRFs, and ferrite multilayer chip components, which are smaller but have a lower inductance and therefore impedance, and a SRF in the hundreds of MHz range. Because the inductor is a series component it must carry the circuit current. Therefore the current rating is important for power circuits, as is the winding resistance, which causes a series voltage drop and may well be the most important limiting factor. For signal circuits these are usually not relevant.

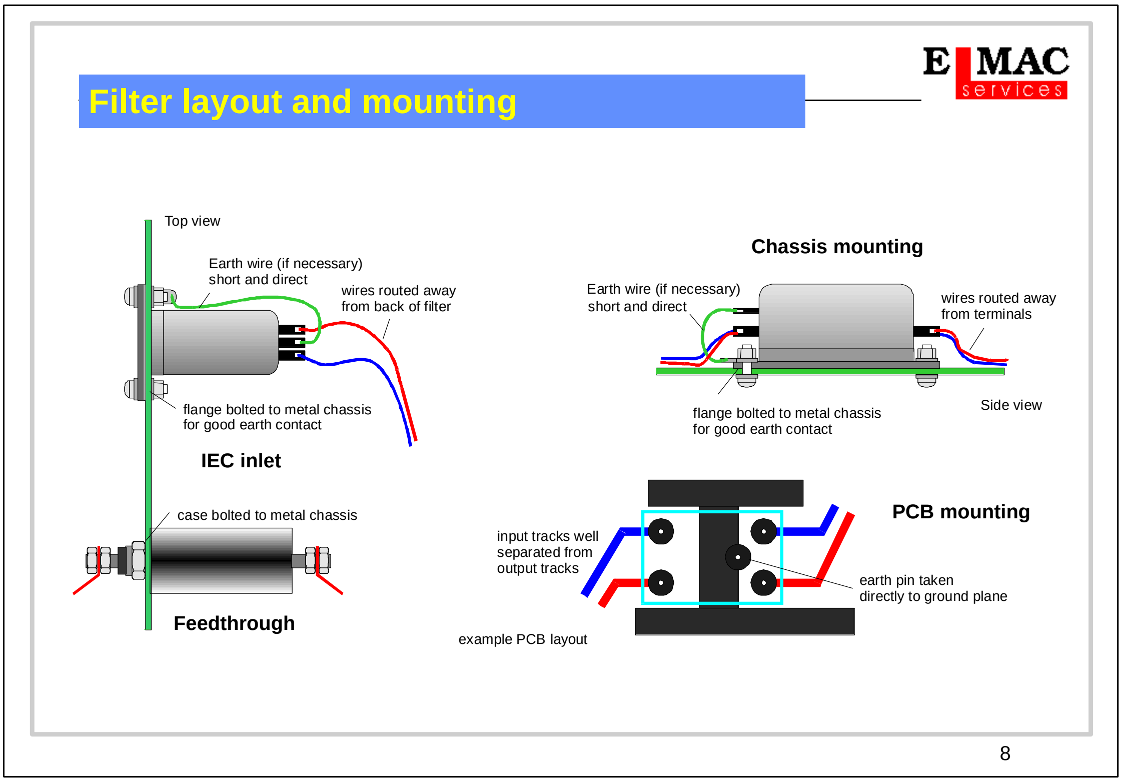

Lead inductance and stray capacitance can degrade filter performance markedly at high frequency. Two common faults in filter applications are not to provide a low-inductance ground connection, and to wire the input and output leads in the same loom or at least close to each other.

A poor ground offers a common impedance which rises with frequency and couples HF interference straight through via the filter's local ground path. Common input-output wiring does the same thing through stray capacitance or mutual inductance. The cures are to mount the filter so that its ground node is directly coupled to the lowest inductance ground of the equipment, preferably the chassis, and to keep the I/O leads separate, preferably screened from each other. The best solution is to position the filter so that it straddles the equipment shielding, where this exists.

Component layout within the filter itself is also important. Input and output components should be well separated from each other.

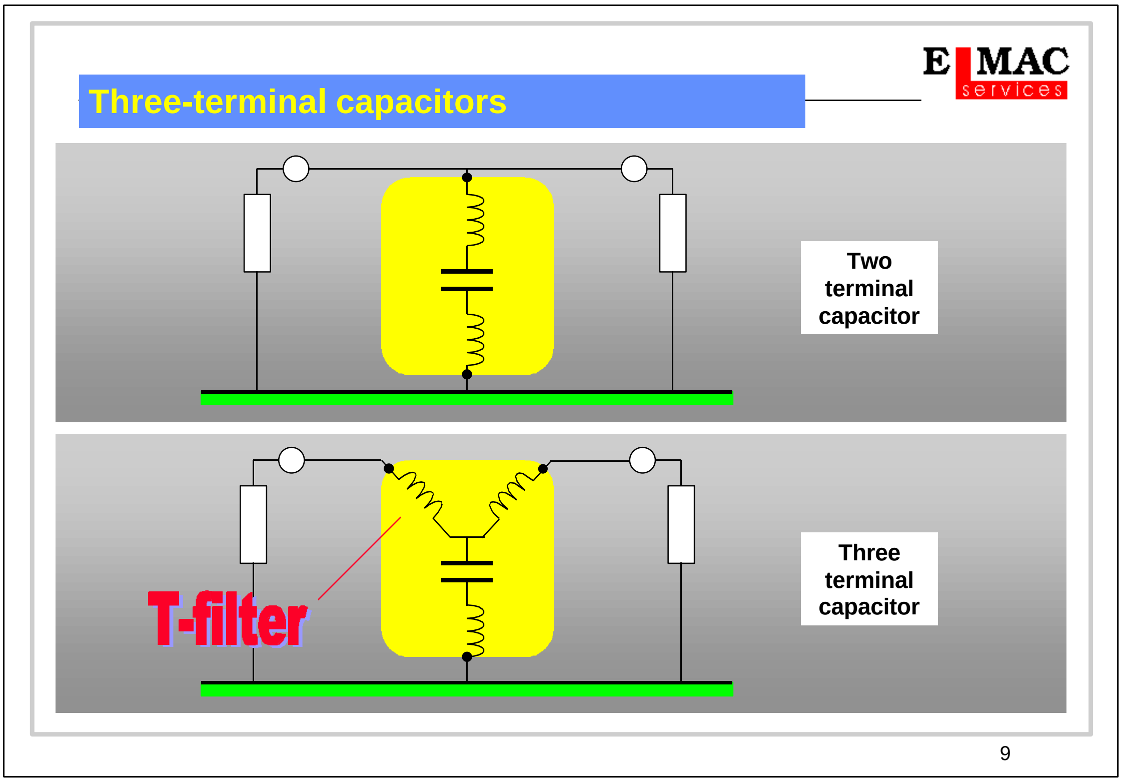

Any low-pass filter configuration except for the simple inductor uses a capacitor in parallel with the signal path. A perfect capacitor would give an attenuation increasing at a constant 20dB per decade as the frequency increased, but a practical wire-ended capacitor has some inherent lead inductance which in the conventional configuration puts a limit to its high frequency performance as a filter. The impedance characteristics show a minimum at some frequency and rise with frequency above this minimum. This lead inductance can be put to some use if the capacitor is given a three-terminal construction.

The lead inductance now forms a T-filter with the capacitor, greatly improving its high-frequency performance. Lead inductance can be enhanced by incorporating a ferrite bead on each of the upper leads. The ground terminal of the capacitor still has some residual inductance which limits its performance, and the ground connection must be made within the shortest possible difference. Even so, the 3-terminal configuration can extend the effectiveness of a small ceramic capacitor from below 50MHz to upwards of 200MHz, which is particularly useful for interference in the vhf band.

Feedthrough capacitors

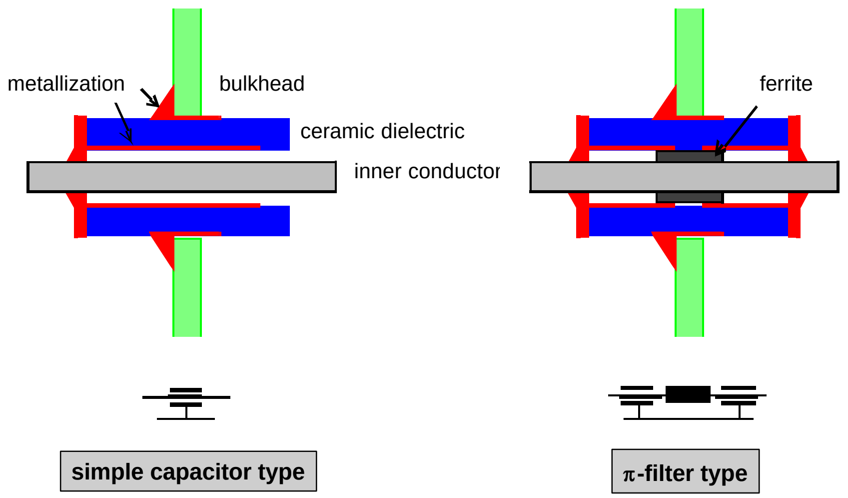

10

Any leaded capacitor is still limited in effectiveness by the inductance of the connection to the ground point. For the ultimate performance, and especially where penetration of a screened enclosure must be protected at uhf and above then a feedthrough construction is essential.

Here, the ground connection is made by the outer body of the capacitor being screwed or soldered directly to the metal screening or bulkhead. Because the current can flow out from the central conductor over 360º around it, there is effectively no inductance associated with this terminal and the capacitor performance is maintained well into the GHz region. The inductance of the through lead can be increased, thereby creating a π-section filter, by separating the ceramic metallization into two parts and incorporating a ferrite bead within the construction.

Feedthrough capacitors are available in a wide range of voltage and capacitance ratings but they are generally expensive and their cost increases with size. They are rarely used for EMC purposes in low-cost commercial products.

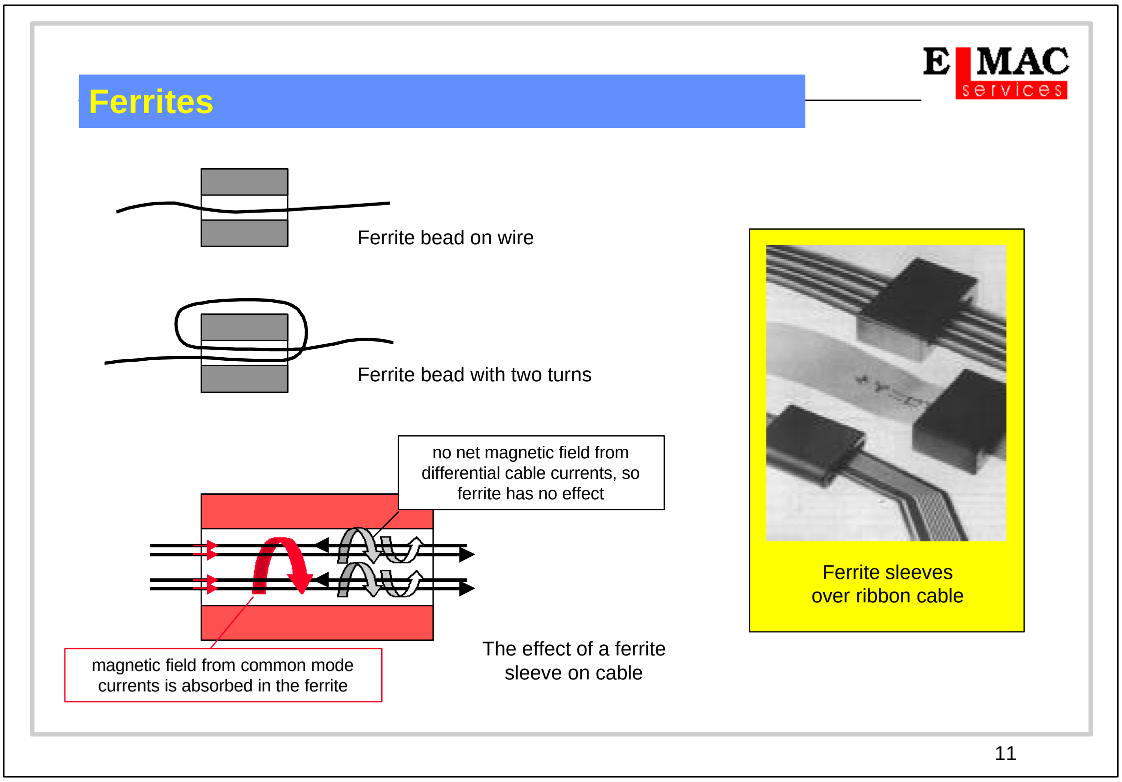

A very simple, inexpensive and easily-fitted filter is obtained by slipping a ferrite sleeve around a wire or cable. Any wire carrying a current has a magnetic field around it. The effect of the ferrite is to concentrate this field and hence to increase the wire’s inductance by several hundred times.

The attractiveness of the ferrite choke is that it involves no circuit redesign, and often no mechanical redesign either. It is therefore very popular for retro-fit applications. Several manufacturers offer kits which include halved ferrites, which can be applied to cable looms immediately to check for improvement.

If a ferrite is put over a cable which includes both signal and return lines, it will have no effect on the signal (differential-mode) current but it will increase the impedance to common-mode currents. This is because the *differential* currents, by definition, sum to zero in each wire pair and therefore there is no net magnetic field. If there is no field, the ferrite is invisible. But the *common mode* currents do produce a net magnetic flux and this flux is concentrated in the bulk of the ferrite, leading to an increased impedance for these currents only.

The effectiveness can be increased by looping the cable several times through the core, or by using several cores in series. The ferrite is at its most effective when it fits snugly over the wire or cable and if it is long and narrow rather than short and fat.

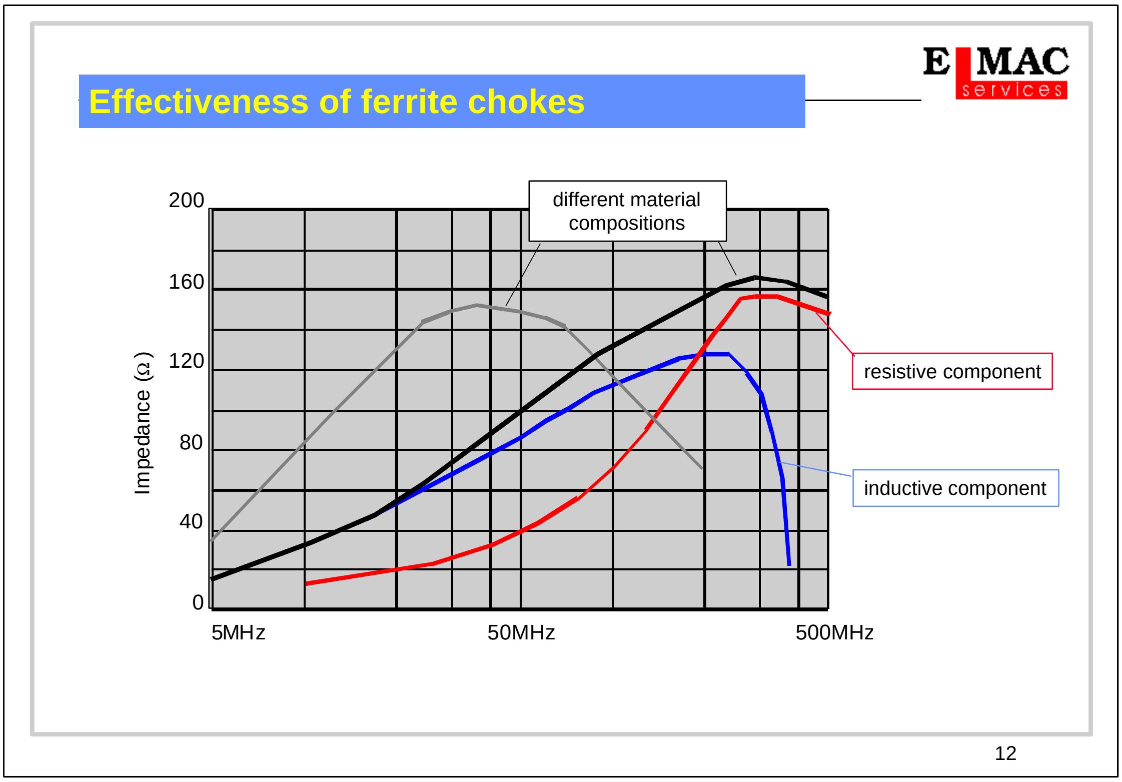

Ferrite effectiveness increases with frequency. The impedance of a ferrite choke is typically around 50Ω at 30MHz, rising to hundreds of ohms above 100MHz (the actual value depends on shape, size and material composition). Usually a ferrite is ineffective at frequencies lower than 30MHz, becomes most effective above 100MHz and falls off in performance as the frequency approaches 1GHz. A useful property of ferrites is that their impedance becomes resistive at the higher frequencies, so that interference energy tends to be absorbed rather than reflected. This property is deliberately enhamced in parts that are designed for suppression purposes, whereas for other applications it is usually undesirable. Therefore ferrite components that are not intended for suppression should not be used for this.

Because a ferrite choke is effectively a lossy inductor, it only functions usefully between low impedances. A ferrite included in a high-impedance line will offer little or no attenuation. Most circuits, and especially cables, show impedances that vary with frequency in a complex fashion but normally stay within the bounds of 10 - 1000Ω, so a ferrite will give modest attenuation factors averaging around 10dB and rarely better than 20dB.

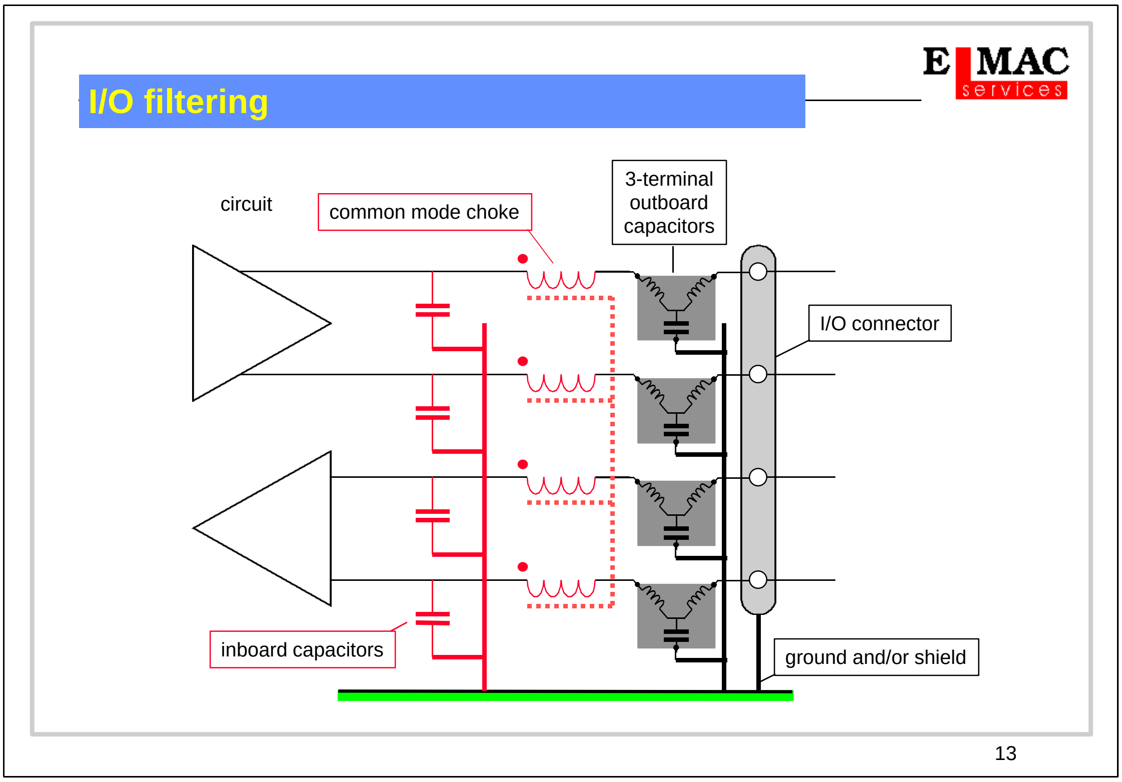

If I/O connections carry only low bandwidth signals it is possible to filter them using simple RC low-pass networks. This is not possible with high-speed data links, but it is possible to attenuate common-mode currents entering or leaving the equipment without affecting the signal frequencies by using a discrete common-mode choke arrangement. Such units are available commercially or can be custom designed. If they are mounted on a pcb, some degradation in the high frequency response is to be expected because of stray capacitance feedthrough and the method of choke construction. C-M chokes are only effective when faced with low impedances and are therefore often implemented with low value parallel capacitors, which have an increasing effect at RF but do not significantly restrict circuit bandwidth.

The capacitors must be decoupled to the clean I/O ground, not to the circuit 0V, which is often the prime carrier of interference. This defines a preferential route to ground for both incoming and outgoing interference. The clean ground may be the case metalwork, a grounding plate for the I/O connectors or a designated ground plane area on the pcb.

Three terminal capacitors are often a good choice for these components since the circuit configuration lends itself to their use. Capacitors inboard of the CM choke may be useful for increased performance, especially if the circuit impedance is too high for the choke to be effective on its own.

Mains filters

- **Reasons for popularity of mains filter modules**
  - Mandatory conducted emission standards concentrate on mains port, hence established market for filter units
  - Add-on “fit and forget” filters can be retro-fitted
  - Safety approvals already achieved
  - Many equipment designers not familiar with RF filter design

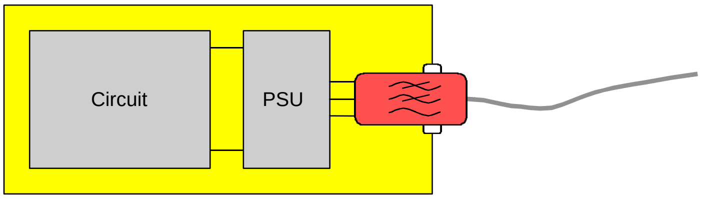

14

RFI filters for mains supply inputs have developed as a separate species and are available in many physical and electrical forms from several specialist manufacturers. A typical "block" filter for European mains supplies with average insertion loss might cost around £5.

Some of the reasons for the development and use of block mains filters are summarized above. Merely adding a block filter to a mains input will improve low frequency emissions such as the low harmonics of a switching power supply. But HF emissions (above 1MHz) require attention to the layout of the circuitry around the filter. Treating it like any other power supply component will not give good HF attenuation and may actually worsen the coupling, through the addition of spurious resonances and coupling paths.

The combined filter and CEE22 inlet connector modules are a good method of ensuring correct layout, providing they are used within a grounded conducting enclosure.

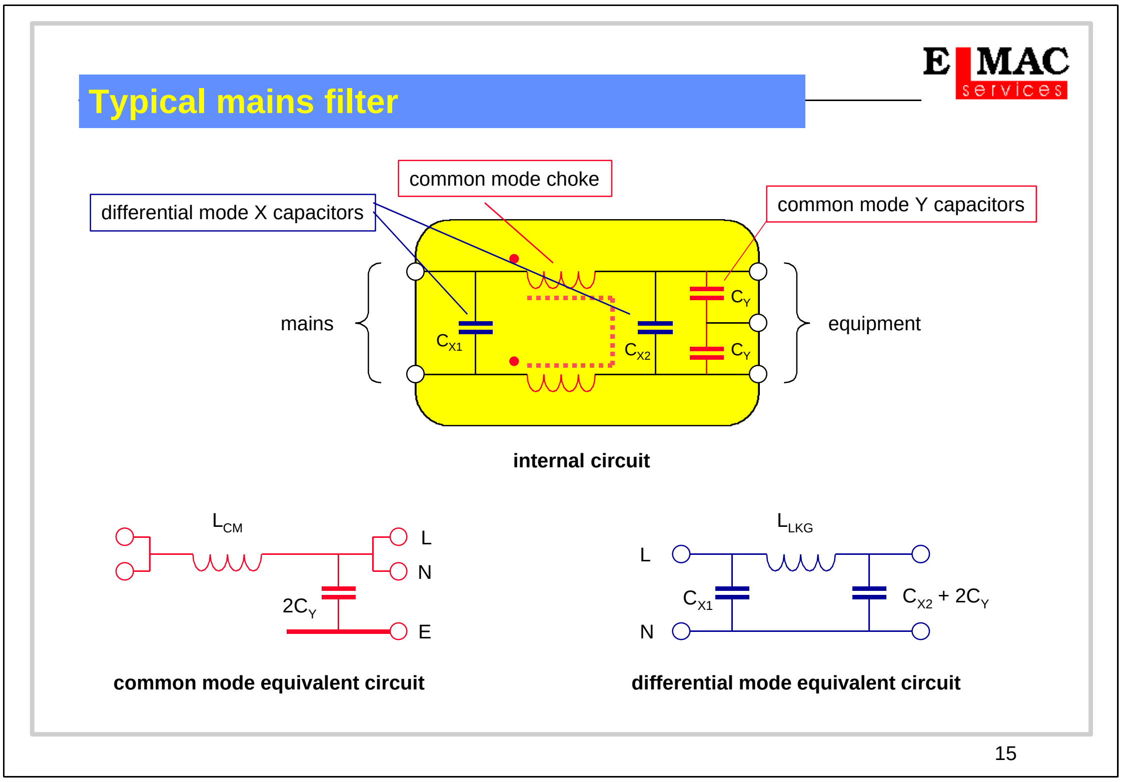

A typical filter includes components to block both common-mode and differential-mode components. The common-mode choke L consists of two identical windings on a single high permeability, usually toroidal, core, configured so that differential (line-to-neutral) currents cancel each other. This allows high inductance values, typically 1-10mH, in a small volume without fear of choke saturation caused by the mains frequency supply current. The full inductance of each winding (L_{CM}) is available to attenuate common-mode currents with respect to earth, but only the leakage inductance (L_{LKG}, which is determined by winding imbalance) will attenuate differential-mode interference.

Capacitors C_{X1} and C_{X2} attenuate differential-mode only but can have fairly high values, 0.1 to 0.47µF being typical. Either may be omitted depending on the detailed performance required, remembering that the source and load impedances may be too low for the capacitor to be useful. Capacitors C_{Y1} and C_{Y2} attenuate common-mode interference and if C_{X2} is large, have no significant effect on differential mode.

Extended performance

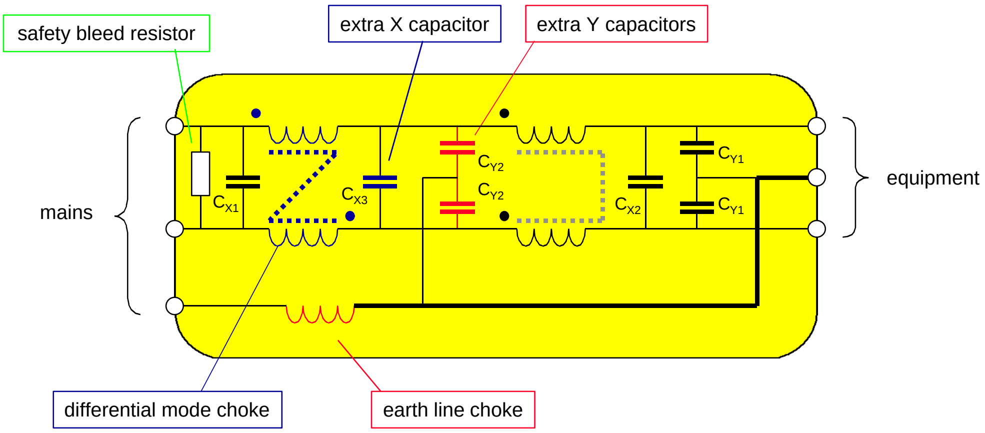

Internal circuit of high-performance mains filter

16

Suppression of common mode interference using a common mode choke and capacitors on the live and neutral lines can be supplemented by increasing the RF impedance of the earth return, using an earth line choke. This is especially useful if the case of the equipment itself, being an imperfect screen, is found to be “live” at RF. It assumes that a protective earth connection exists and is the main return path for CM currents; the technique is not suitable for Class II double insulated equipment.

The choke must be sized to safely carry the full predicted earth fault current and have a low DC resistance, and is therefore large. Also, it cannot be used when any other ground connection is made to the equipment, such as via signal or control lines to other grounded instruments, or via a separate earth bonding strap. This would short circuit the choke, or at least radically alter the overall inductance of the path. Similarly, if a large stray capacitance to earth exists it will appear across the choke and cause unpredictable resonant effects.

Extra differential mode attenuation can be provided by adding a differential mode choke and/or a further X-capacitor. The differential choke is wound in the opposite direction to a common mode choke and therefore does not compensate the power current, so cannot have a high inductance value. Similarly, extra common mode attenuation is provided by adding a further set of Y capacitors and/or another common mode choke.

Many safety standards require a bleed resistor to discharge the X capacitors when the connection to the supply is broken. In general, this is needed if the total capacitance value is more than 0.1µF.

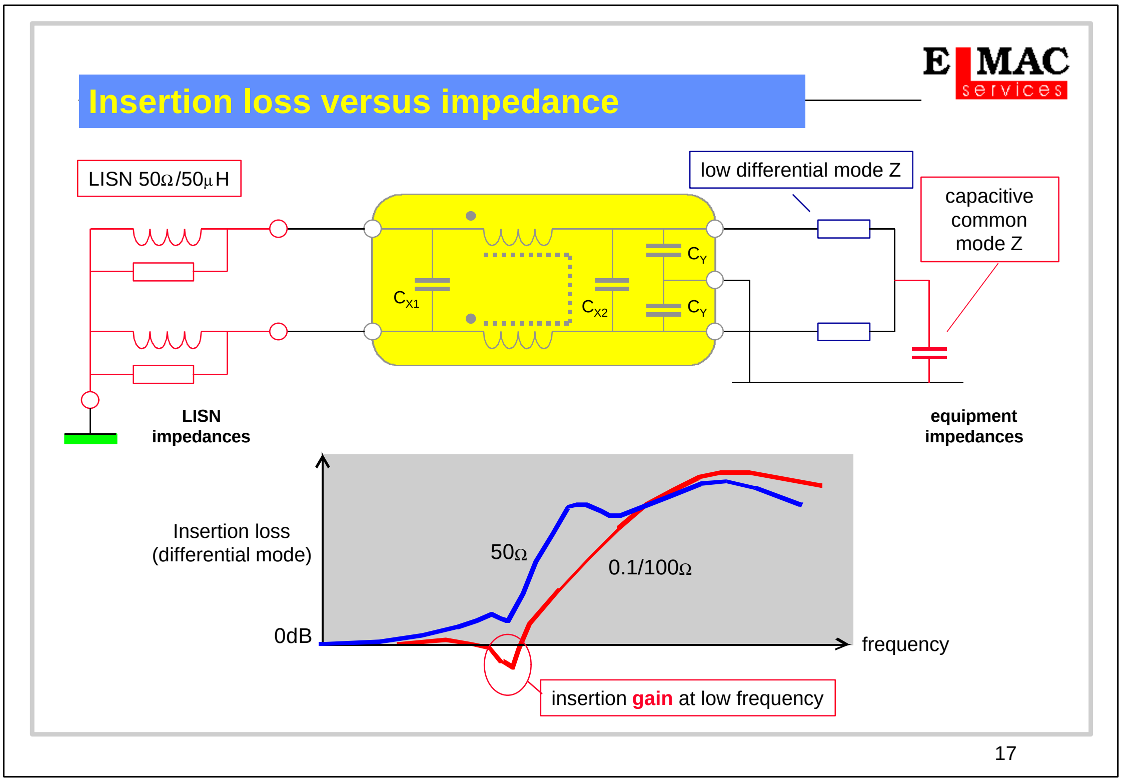

Ready-made filters are universally specified between 50Ω source and load impedances. This does not reflect the real situation.

The mains port HF impedance can be generalized for both common- and differential-mode by a 50Ω//50µH network as provided by a CISPR-16 LISN. The equipment port impedance will vary substantially depending on load and on the HF characteristics of the input components. Common-mode impedance is a function of stray capacitance coupling to earth and can normally be approximated by a capacitive reactance of around 1000pF.

The effect of these load impedances differing from the nominal may be to enhance resonances within the filter and thus to achieve insertion *gain* at some frequencies, typically at or around 150kHz. Some filter manufacturers provide extra curves for attenuation with 0.1/100Ω differential mode load impedances, measured according to CISPR 17, and these figures are more realistic for real situations.

The problem of poor low frequency performance is usually worst for single-stage filters and can be dealt with by adding extra differential mode components, but this makes the filter more expensive.

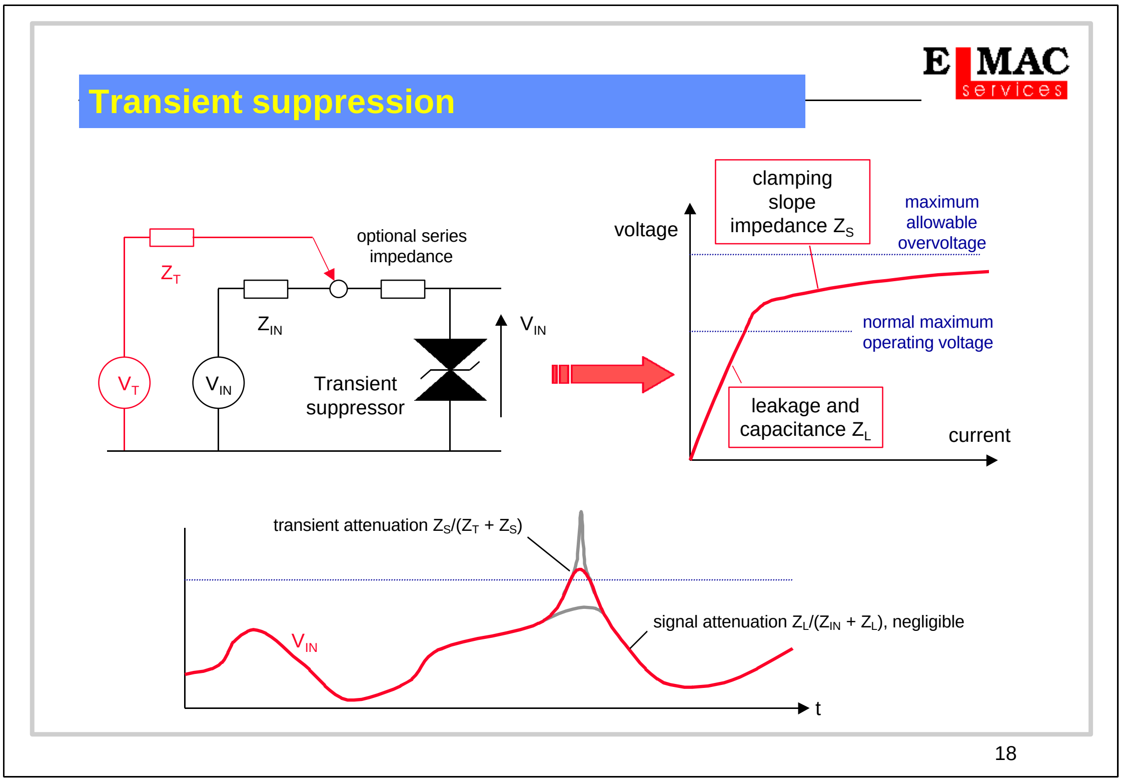

Incoming overvoltage transients are clamped at the circuit interface by using non-linear components: varistors (MOVs), zeners (including foldback types) and spark gaps (gas discharge tubes). These are placed in parallel with the transient source and their effectiveness depends on the ratio of their dynamic slope impedance Z_S to the transient source impedance Z_T.

The choice of device is dictated by circuit operating parameters: leakage current, capacitance and threshold voltage are important with respect to normal circuit operation; clamp voltage, follow-on current, energy capability and response time are important when the device is faced with a transient. Response time is particularly critical for transients with fast rising edges, as a slow device will let through the leading edge which may well be the most damaging part of the transient. Device inductance is as important as slope resistance and short leads or chip packages are preferred. A high transient current flows through the device when it operates and this dictates careful layout.

Use of transient suppressors

Device	Leakage	Follow-on current	Clamp voltage	Energy capability	Capacitance	Response	Cost
ZnO Varistor	Moderate	No	Medium-high but degrades	High	High	Medium	Low
Zener	Low	No	Low to medium (can foldback)	Low	Medium	Fast	Moderate
Spark gap	Zero	Yes	High ignition, low clamp	High	Very low	Slow	Moderate

19

The device must be sized to withstand the continuous operating voltage of the circuit, with a safety margin, and to be able to absorb the energy from any expected transient. The first requirement is fairly simple to design to, although it means that the transient clamping voltage is usually 1.5 - 2 times the continuous voltage, and circuits that are protected by the suppressor must be able to withstand this. The second requirement calls for a knowledge of the source impedance and likely amplitude of the transients, which is often difficult to predict accurately especially for external connections. Some research has been done on mains transient amplitudes and in the telecom and automotive environments. See e.g. Transient Voltage Suppression Manual, General Electric, 1986; for guidance on expected transients in mains-supplied environments, see “IEEE Recommended Practice on Surge Voltages in Low-Voltage AC Power Circuits”, IEEE/ANSI C62.41-1991.

Devices can be combined with each other or with capacitors and ferrites to offer a total suppression performance that would be unobtainable from any single device.

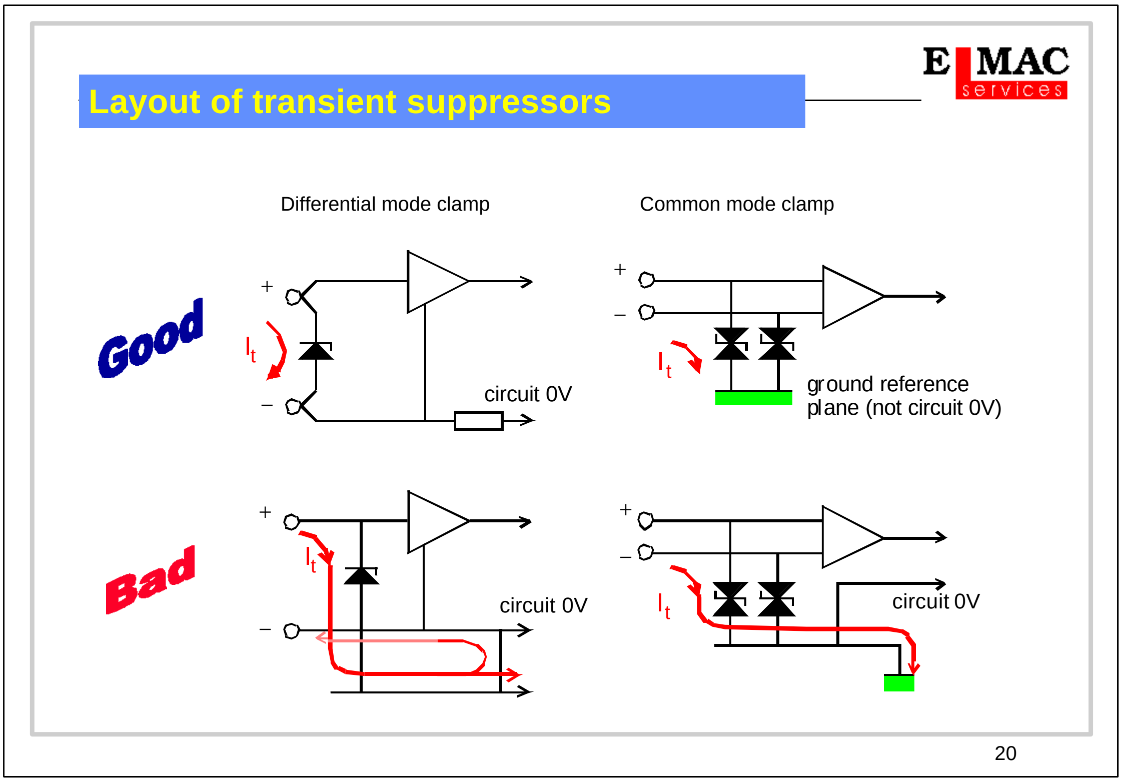

Short and direct connections to the suppressor (including the ground return path) are vital to avoid compromising the high-speed performance by undesired extra inductance. Transient edges have very fast risetimes (a few nanoseconds for switching-induced interference down to sub-nanosecond for ESD) and any inductance in the clamping circuit will generate a high differential voltage during the transient plus ringing after it, which will defeat the purpose of the suppressor.

The component leads must be short (suppressors are available in SM chip form) and they must be connected locally to the circuit that is to be clamped. Any common impedance coupling, via ground or otherwise, must be avoided. When the expected transient source impedance is low (less than a few ohms), it is worthwhile raising the RF impedance of the input circuit with a lossy component such as a ferrite bead. Where suppressors are to be combined with I/O filtering it may be possible to use the 3-terminal varistor/capacitor devices that are now available.

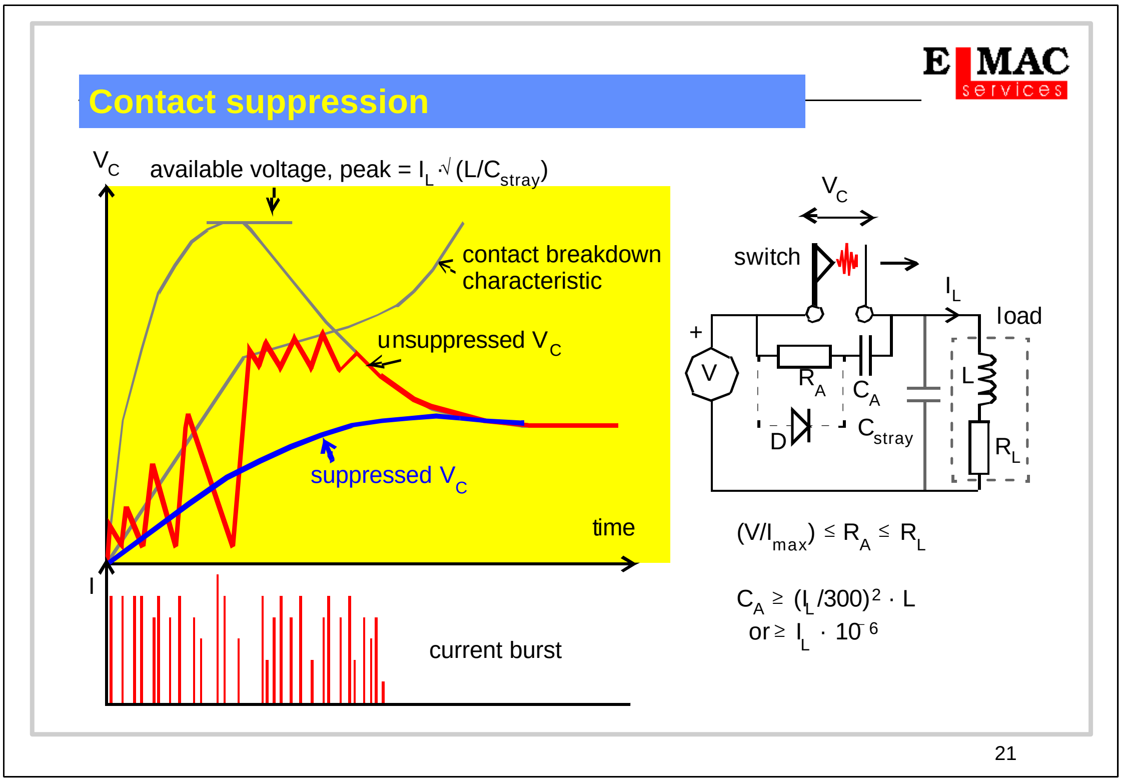

An opening contact which interrupts a flow of current - typically a switch or relay - will initiate an arc across the contact gap. The arc will continue until the available current is not enough to sustain a voltage across the gap. The stray capacitance and inductance associated with the contacts and their circuit will in practice cause a repetitive discharge until their energy is exhausted, and this is responsible for considerable broadband interference. A closure can also cause interference because of contact bounce.

Any spark-capable contact should be suppressed. The conventional suppression circuit is an RC network connected directly across the contacts. The capacitor is sized to limit the rate-of-rise of voltage across the gap to below that which initiates an arc, typically 1V/µs for most contacts. The resistor limits the capacitor discharge current on contact closure; its value is a compromise between maximum rated contact current and limiting the effectiveness of the capacitor. A parallel diode can be added in DC circuits if this compromise cannot be met.

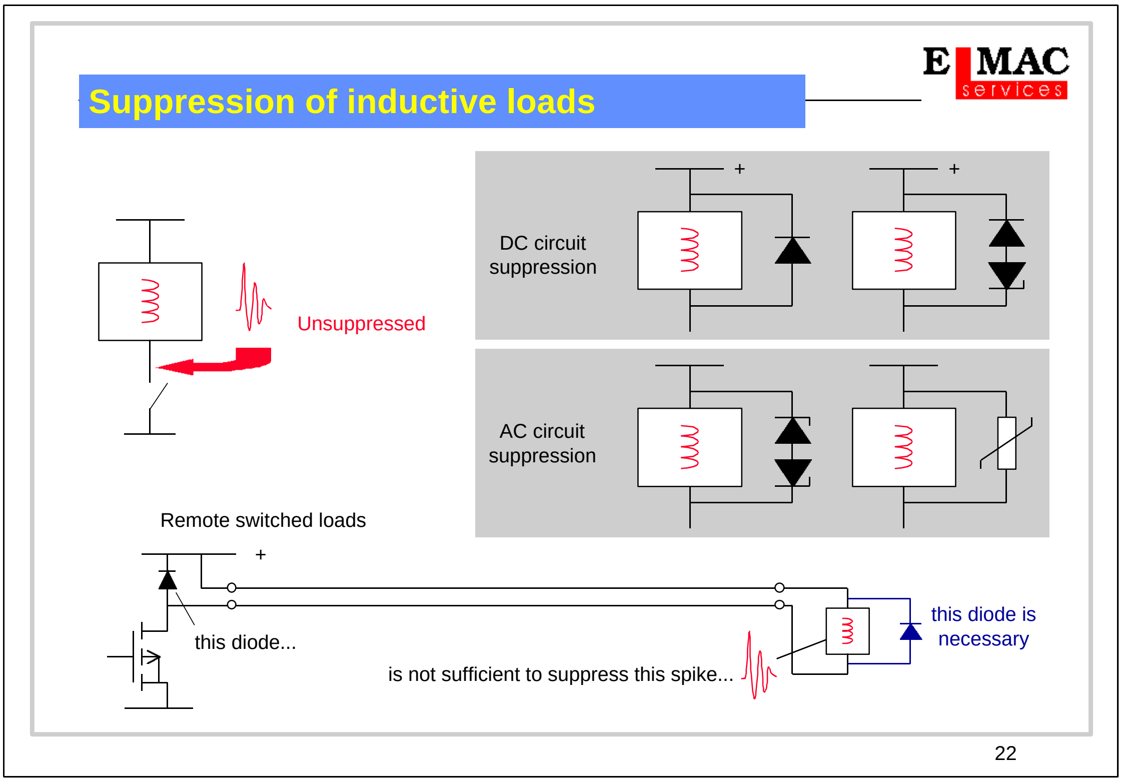

When current through an inductance is interrupted a large transient voltage is generated, governed by $V = -L \cdot di/dt$. Theoretically if di/dt is infinite then the voltage is infinite too; in practice it is limited by stray capacitance if no other measures are taken. Typical examples of switched inductive loads are motors and relay coils. Switching can either be via an electromechanical contact or a semiconductor, and the latter can easily suffer avalanche breakdown if the transient is unsuppressed. RF interference is generated in both cases at frequencies determined by stray circuit resonances and is usually coupled by the wiring between switch and load.

The RC snubber circuit can be used in some cases to damp an inductive transient. Other circuits use diode, Zener or varistor clamps as shown. In all cases the suppression components must be mounted immediately next to the load terminals, otherwise a radiating current loop is formed by the intervening wiring. Protection of a remote drive semiconductor must be considered separately.

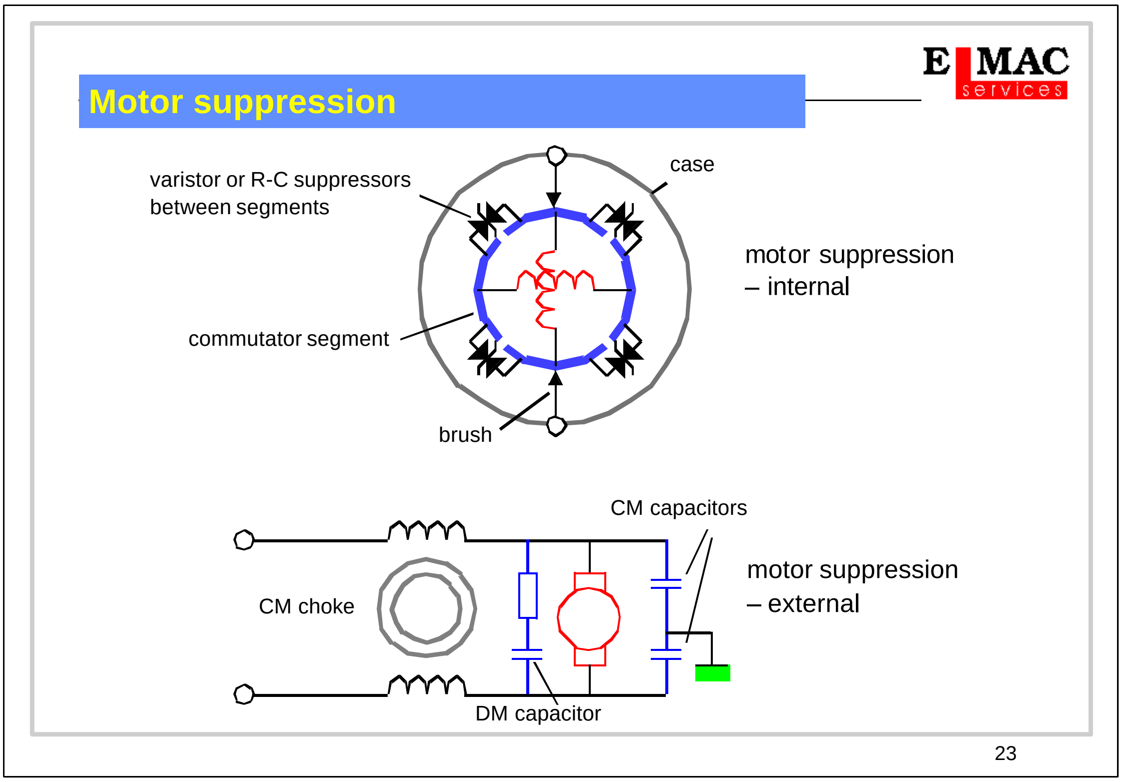

DC motor noise is particularly aggressive, since it consists of impulsive and hence wideband transients repeated at a rate determined by the commutation speed – in other words, several hundred to several thousand times a second. The spectral composition of this noise may extend up to several hundred MHz. This appears both as differential mode noise across the terminals and common mode with respect to the housing, coupled through stray capacitance.

The best method of suppression is to prevent the motor from generating impulsive voltages across the commutator segments. This can only be achieved by the motor manufacturer, by incorporating varistor or RC components between each commutator segment, but ensures that the motor is quiet without further suppression being necessary. Otherwise, capacitors across the terminals (for differential mode) and from the terminals to the local earth (for common mode) are required; if the local earth is not available, then a common-mode choke at the motor terminals may be needed.

End of this section

24