Another EMC resource
from EMC Standards

Part 6b, Matched Transmission Lines continued

Helping you solve your EMC problems

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

compliance-club.com
Keep up to date. Free Membership.

Advanced PCB Design and Layout for EMC Part 6 - Transmission Lines - 2nd Part

Eur Ing Keith Armstrong C.Eng MIEE MIEEE, Cherry Clough Consultants

Due to its length we have had to split Part 6 between three issues. This issue contains Section 2 - Terminating transmission lines and Section 3 - Transmission line routing constraints.

This is the sixth in a series of eight articles on good-practice design techniques for electromagnetic compatibility (EMC) for printed circuit board (PCB) design and layout. This series is intended for the designers of any electronic circuits that are to be constructed on PCBs, and of course for the PCB designers themselves. All applications areas are covered, from household appliances; commercial, medical and industrial equipment; through automotive, rail and marine to aerospace and military.

These PCB techniques are helpful when it is desired to…

- Save cost by reducing (or eliminating) enclosure-level shielding
- Reduce time-to-market and compliance costs by reducing the number of design iterations
- Improve the range of co-located wireless datacomms (GSM, DECT, Bluetooth, IEEE 802.11, etc.)
- Use very high-speed devices, or high power digital signal processing (DSP)
- Use the latest IC technologies (130nm or 90nm chip processes, ‘chip scale’ packages, etc.)

The topics to be covered in this series are:

1. Saving time and cost overall
2. Segregation and interface suppression
3. PCB-chassis bonding
4. Reference planes for 0V and power
5. Decoupling, including buried capacitance technology
6. **Transmission lines**
7. Routing and layer stacking, including microvia technology
8. A number of miscellaneous final issues

A previous series by the same author in the EMC & Compliance Journal in 1999 “Design Techniques for EMC” [1] included a section on PCB design and layout (“*Part 5 – PCB Design and Layout*”, October 1999, pages 5 – 17), but only set out to cover the most basic PCB techniques for EMC – the ones that all PCBs should follow no matter how simple their circuits. That series is posted on the web and the web versions have been substantially improved over the intervening years [2]. Other articles and publications by this author (e.g. [3] [4] and Volume 3 of [5]) have also addressed basic PCB techniques for EMC.

Like the above articles, this series will not spend much time analysing why these techniques work, it

will focus on describing their practical applications and when they are appropriate. But these techniques are well-proven in practice by numerous designers world-wide, and the reasons why they work are understood by academics, so they can be used with confidence. There are few techniques described in this series that are relatively unproven, and this will be mentioned where appropriate.

Table of Contents, for this Part of the series

Section 1 is contained in Issue 57, March 2005

2 Terminating transmission lines

A transmission line must be terminated at its ends with resistors that match its characteristic impedance. Only then does it become a 'matched transmission line', and only then do all of the SI and EMC benefits arise.

Where the rise/fall time is long enough – or the highest frequency of concern low enough – the entire length of a trace with all of its low impedance discontinuities (e.g. due to capacitive loads) and high-impedance discontinuities (e.g. due to perforation of the reference plane) can be treated as one averaged impedance. This is then the impedance that must be matched by the terminating resistors.

But where the rise/fall time is short enough for one section of a trace to have a significantly different characteristic impedance from another section (see the 300ps curve on Figure 6E), it will be impossible to match the transmission line. The characteristics of the trace and its loads, reference planes and other physical features must be modified so that each section has the same characteristic impedance as is desired, then this is terminated in matched resistors at the ends of the line.

Variations in impedance that are considered acceptable for SI purposes may be too high for good EMC. So, for good EMC, take more care in matching the impedances along the length of a trace, and at its terminating resistors.

2.1A range of termination methods

RF and microwave designers use very narrow frequency ranges, so they can terminate their transmission lines in complex impedances that are carefully designed to do exactly what they want them to over that small range. But SI and EMC require termination to be effective over a wide range of frequencies – so resistors that maintain their resistive behaviour from d.c. to the highest frequency of concern are required. Ideally, this requires using very small surface mounted 'chip' resistors (but not MELF types, because they are too inductive).

Very short and direct connections between the terminating devices and the reference plane(s) are required. Ideally 'via-in-pad' techniques should be used, where practical. Crosstalk between termination resistors can occur, so where there are two or more close together they should be arranged in a neat line, with their 0V ends (or other reference plane connections) at the same end.

When choosing a resistor network or array for termination: to be effective at frequencies above 100MHz they must have reference plane connections at *both* ends of the device, at least. For effective termination above 1GHz they will probably need to have as many reference plane pins as signal pins. Some PCB manufacturers can laminate high-resistance metal foils as internal layers in PCBs, which are then etched to provide termination resistors.

'Classical RF' termination, where the transmission line is resistively matched at both ends, is the best method for both SI and EMC, but attenuates the signal voltage by 50%. This type of termination is often used for RF and datacommunications over cables, and for very high-speed backplanes, and it is the best termination method for both SI and EMC. But digital ICs will not work with half-height input signals, so for digital circuits on PCBs single-ended resistive termination is traditionally used instead [13]. Some types of ICs may soon be available with drivers that use internal voltage boosters so that classical RF termination can be used whilst providing the full signal voltage that the receiving ICs require.

There are a variety of single-ended termination methods available, divided into 'reflected wave switching' and 'incident wave switching' types, as shown in Figures 6L, 6M and 6N. Designing classical RF termination combines both reflected and incident wave methods – so will not be discussed here as a separate topic.

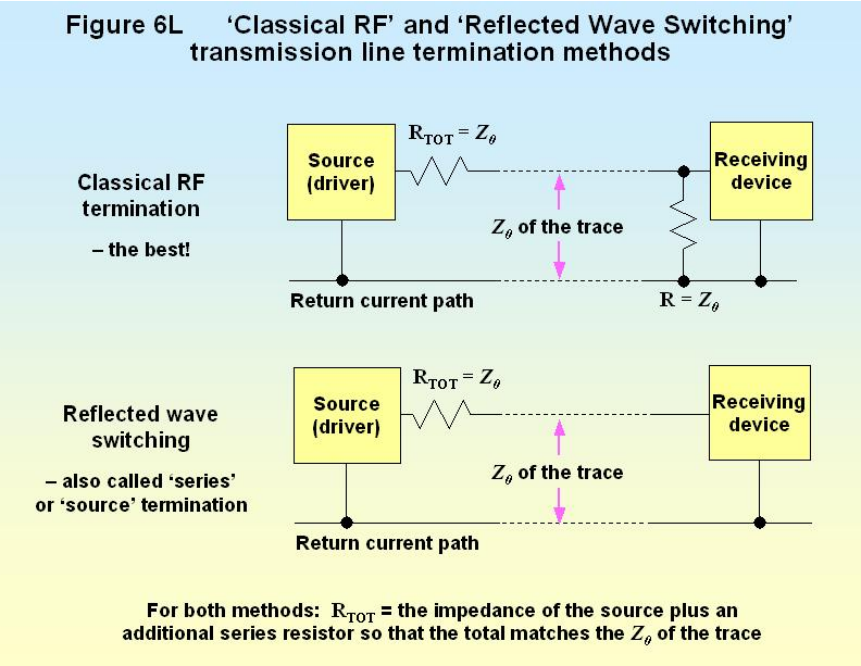

Reflected wave switching is sometimes called 'series' or 'source' termination, and is shown in Figure 6L. The driven end of the transmission line is terminated with an *overall* resistance equal to the line's Z_0, and – when perfectly matched – launches a half-height signal voltage transition into the trace. When the signal transition reaches the end of the trace, the very high impedance there creates a reflection coefficient of almost +1, reflecting the transition back down the line – but this time with a voltage that is almost the full height of the wanted signal. The reflected transition travels back along the trace, establishing the correct signal voltage along it as it goes. When it reaches the driver termination the reflection coefficient is zero (if the line is perfectly matched), so the signal voltage transition does not reflect again.

Reflected wave switching is good for point-to-point interconnections. If used for a multidrop bus there is the possibility that the devices along the line will be 'double clocked' because they will see half the signal amplitude some time before they see the full signal. Such devices need to be slow enough when responding, when compared with the propagation time along the line, to ensure that the reflections along the line will not confuse them. Alternatively, ensure that all of the load devices are located at the far end of the transmission line, separated from each other by traces for which $t_p \geq t_r/10$ or $t_p \geq 1/10\pi f$ for SI, or for good EMC: $t_p \geq t_r/40$ or $t_p \geq 1/40\pi f$

.

The series terminating resistors used in reflected wave switching must be located very close to the driver, and, as shown in Figure 6L, the value of the resistor is chosen so that the combination of driver impedance and external resistance equals the line's Z_0. With 'ordinary' digital ICs (not designed for transmission line matching) it may be hard to achieve good matching when pulling-up and when pulling down, and this is discussed later.

Figures 6M and 6N show the variety of incident wave switching methods. They all use a single-ended resistive termination at the furthest end of the transmission line from the driver. An idealised driver with a 0Ω source impedance launches a full-amplitude signal voltage transition onto the trace, which then travels along the trace until it reaches the very end, where a 'parallel' or 'shunt' terminating resistor is installed. Assuming matching of this resistance with the Z_0 of the trace, the reflection coefficient will be zero so there will be no reflections.

There will be reflections at the launching point, but these are all contained within the driver and not seen outside its IC. If the impedance of the load is less than 100 times Z_0, the value of the termination resistor should be chosen so that the parallel combination of load and the termination resistance equals Z_0. Parallel termination resistors are usually connected to the 0V plane, but some logic families use other reference planes (e.g. ECL uses the positive power plane) so these are where the termination resistors should connect to, instead.

Figure 6M 'Incident Wave Switching' using shunt and RC transmission line termination methods

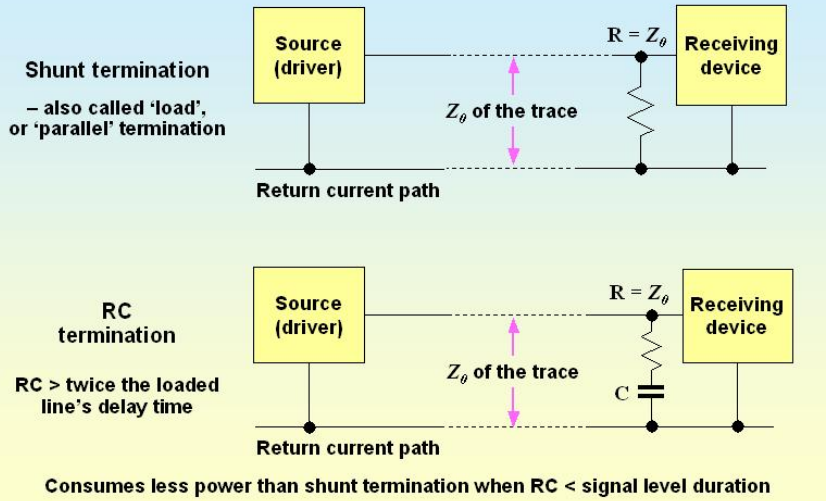

Incident wave switching allows multidrop busses to be used at much greater data rates than reflected wave switching. Using shunt resistor terminations consumes a lot of power, and may exceed the d.c. current limits of some drivers, so a variety of alternative termination circuits have been developed to reduce power consumption and d.c. current (hopefully without compromising the line matching too much at the highest frequencies). Alternative types of incident wave switching termination include 'RC', 'Thévenin' and 'Active'.

RC termination (Figure 6M) uses a resistor value equal to the line's Z_0, and capacitor values that are usually between 10 and 620pF. It only terminates the line at high frequencies, and the idea is that it will consume less power, and because it consumes no d.c. current it should be easier to drive with 'ordinary' devices.

But the effectiveness of RC termination depends upon the randomness of the data and the length of the trace, so it is not a universally applicable technique as the other three incident wave methods are. The power saving is only achieved when the RC time constant is less than the duration of the signal level. Also, the RC time constant must be larger than twice the loaded line's delay time for it not to cause SI problems. For more on this see "*AC Terminators*" in [12].

Because of the self-inductance of capacitors [16] it may be more difficult for an RC termination to equal the highest frequency performance of a purely resistive termination methods. Some component manufacturers offer single or array devices that package the R and C together to save PCB area, and these can be designed to have better high-frequency performance than discrete components.

Thévenin termination (Figure 6N) uses resistor values designed so that their parallel resistance is Z_0, and their ratio minimises the power consumption and the d.c. loading of the driver. Ideally, this means choosing the resistor ratio so that the d.c. voltage at the transmission line, with the driver

disconnected, would equal the average line voltage when it has data on it. Thévenin termination needs a properly decoupled power plane at all frequencies of concern, refer to [16] for how to achieve this.

Figure 6N 'Incident Wave Switching' using Thévenin and Active transmission line termination methods

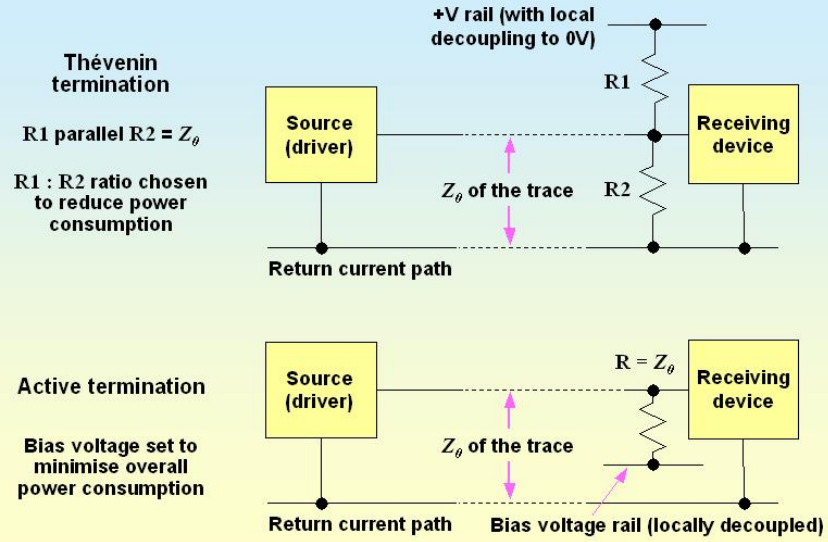

Active termination (Figure 6N) uses a voltage regulator to drive an additional power rail (or reference plane) at the nominal average value of the digital signals, or some other suitable voltage. A parallel terminating resistor with a value of Z_0 connects to this plane, which must be properly decoupled for the frequencies of concern (see [16]). Electrically equivalent to the Thévenin method, active termination can save overall power consumption by running the voltage regulator (which needs to be able to source as well as sink current) in Class AB mode.

Where a transmission line passes through a connector, a common problem is an intermittent or failed connection. The undriven line will still have crosstalk and other noises on it, which can lead to random data activity of the device the line connects to. Finding such faults can be made much easier when Thévenin or active termination methods are used, by biasing their ratios or voltages so that when the driver is disconnected the load device is guaranteed to go to a logic high or low. Choose the 'connector failure' logic level on the basis of which causes the fewest operational problems or best aids fault diagnosis. Where safety is a factor, this should influence the choice of logic level too.

Where a 'stub' trace connects to the trace that is the main transmission line, two consequences are possible, depending on the length of the stub and the capacitance or any load connected to it. Where the stub's propagation time is *very* short compared with the signal's rise/falltimes (or wavelength at the highest frequency of concern) it, and any device capacitance, will appear as a point capacitive load, and should be treated as discussed in the earlier section on capacitive loading.

But where the stub's propagation time is **not** insignificant (compared with the signal's rise/falltimes or wavelength at the highest frequency of concern) the effect is more complex. The Z_0 at the junction will be the parallel combination of the main line plus the stub (for example: if both have the same Z_0, the result will be $Z_0/2$), and reflections from both the stub junction and the end of the stub will propagate in the main transmission line and distort the signal.

Ideally, the end of such a stub would be terminated correctly in an accurately matched resistance, but this still leaves the reflections at the junction. Correctly terminating a transmission line with stubs

can be a complex task, and sometimes the only practical solutions may be less than ideal. Some of the references at the end of this article provide some guidance on this, especially the textbooks and the IPC standards.

2.2 Difficulties with drivers

Matching a trace's characteristic impedance at the driver end can be difficult for devices that are not carefully designed for driving transmission lines. For most 'ordinary' digital ICs, their pull-up impedance is higher than the trace's Z_0, and their pull-down impedance is lower than Z_0, so correct series termination at the source end – whether for classical RF or reflected wave switching – is impossible. [13] gives several worked examples of this situation.

When using incident wave switching, the non-zero driver impedance means that the full signal voltage can not be initially put on the line. The driver will eventually charge the line to the correct voltage, but this might take a while so could result in data timing problems. One common solution to this is to terminate the line in a value that is higher than Z_0, so that there is some reflection at the end that will compensate for the lack of initial driving voltage, and [13] gives a worked example of this. But this solution may mean that point-to-point connection is required instead of multidrop bussing (unless the devices along the line are slow in responding when compared with the trace's propagation time).

The above situation is further complicated by the fact that digital ICs have a dynamic output impedance that depends upon their output voltage. Calculating what voltages will result from driving transmission lines becomes iterative, which is why [13] describes the Bergeron Plot method for avoiding endless calculations.

Further complications arise from the fact that there can be wide variations in driver output impedances from one batch to another, of the same device from the same manufacturer. The variations between different batches of the 'same' device from different suppliers can be even greater.

Some drivers cannot source enough d.c. current to drive an incident wave terminator if it just a single shunt resistor. RC, Thévenin or active terminations may be required. If these cannot be used either, reflected wave switching using a series resistor at the driver end may be the only solution.

But many ICs are now designed with output drivers suitable for driving transmission lines properly, especially where high clock or data rates are required. These devices can be recognised by the transmission line matching data provided in their data sheets, or from the specifications for their output impedances. A driver designed for driving transmission lines will have a low output impedance, say around 10Ω;, when pulling up *or* down, making it much easier to correctly match to a transmission line. Such devices are much preferred over 'ordinary' or 'glue logic' digital ICs when transmission lines are to be driven. Even so, these devices will have small differences between their output impedances when pulling up or down, so some types may be better than others for EMC reasons.

Note that ECL devices have an output impedance of around 10Ω (pulling up *or* down) and have always been excellent for driving transmission lines.

2.3 Compromises in line matching

Termination techniques recommended by device manufacturers, such as [17], are usually based on SI requirements not EMC. Some techniques will be necessary because others won't work (e.g. Thévenin or active termination is usually required for CMOS 'glue-logic'), but others will often be a compromise between component cost and PCB area; power consumption and battery life; shielding,

filtering and number of PCB layers; use of 'terminating plugs'; EMC, etc.

Reflected wave switching and higher values of trace Z_0 both tend to cause lower levels of emissions (all else being equal) because their signal currents are lower, so they emit lower levels of magnetic fields.

Small impedance discontinuities along a trace and small mismatches in its resistive terminations might be considered unimportant for SI, and yet be *very* important for EMC. SI considerations usually ignore mismatches of ±10% or less, because they make very little difference to the voltage waveform. But measuring the trace current waveform can reveal large overshoots caused by the discontinuities or mismatches, significantly increasing the emissions from the PCB. For good EMC it is important to terminate a transmission line in an accurately-matched resistance, as shown by the brief article "*Reducing Emissions*" in [12].

Trace current can be easily measured without having to cut the trace or lift an IC leg – using current probes supplied by oscilloscope manufacturers or magnetic close-field probes supplied by spectrum analyser manufacturers (or made by yourself [18]) and used to measure trace current without contact. To detect individual traces they should be made very small, usually less than 5mm diameter (although square shapes are often easier to use). Don't forget that magnetic field probes respond to the differential of the current. The use of classical termination (at both ends) helps reduce the effect of reflections caused by changes in Z_0 along a trace, such as are caused by stubs and vias. As mentioned earlier, one effect of this is to attenuate the signal by 50% – so the receivers used should have an adequate input range.

2.4 ICs with 'smart' terminators

Some ICs incorporate 'smart' output drivers (for example: Vertex-II FPGAs and Rocket I/O, both from Xilinx) that can be programmed to various impedances so that they match the PCB transmission lines, with no need for any series resistors for reflected wave switching. Some ICs include on-chip voltage doublers for their output drivers, so they can use classical termination without losing 50% of the signal voltage. These are likely to give the best EMC (all else being equal). These smart driver technologies are usually available as IP for designing into FPGAs or ASICs.

But CMOS processes do not create very accurate resistors, so designers should always ask about the tolerance of the matching resistors built into their drivers or receivers, then take the worst-cases into account in their transmission-line design. Prototype ICs with integrated output drivers that automatically adjust their output and input impedances to optimise SI have been demonstrated – but it is not yet clear if they will achieve as good EMC as accurate matching with a discrete resistor.

In July 2004 a company called Adiabatic Logic Ltd claimed (EPD Magazine, page 43) to have developed "adiabatic" digital drivers. These were said to use reflected wave switching, but instead of terminating the reflected signal edge in resistance they put its energy back into the local power supply. The aim is to reduce power consumption, especially for battery-powered products. The author does not yet know of any EMC measurements made on devices using drivers based on this new technology.

2.5 Bi-directional terminations

Many modern busses don't have a single bus master and one or more slaves – some or all of the devices can be drivers or receivers, and they are called bi-directional busses. The PCI Bus used in PCs (until its replacement by PCI Express) is an example of a bi-directional bus.

Bi-directional transmission lines require reflected or incident wave termination at both of their ends,

and maybe at some/all of the stubs along their route as well. In the case of reflected wave termination all of the line's drivers are fitted with series-terminating resistors. The only problem that this causes is to create a small delay for a receiver that can also be a driver. But in the case of incident wave termination, each driver must now be powerful enough to drive the Z0 of the transmission line plus a nearby parallel termination resistor.

Terminating multidrop bi-directional buses can be very difficult, and is briefly described in "*Bi-directional Terminations*" in [12]. Compromises are often required – for example "*Bi-directional Alternatives*" in [12] describes the PCI Bus as requiring its signals to make three or four 'round trips' of the line before settling down to a stable enough value (see Figure 6B).

2.6 Non-linear termination techniques

Diode termination prevents device damage or latch-up due to overshoots beyond the power rails, but it doesn't terminate a transmission line's Z_0 properly. So the only significant EMC benefit from diode termination may be that they might help prevent transient damage as a result of a nearby electrostatic discharge (ESD). The diodes are not rated for handling ESD energies so should not be relied upon to protect an IC whose interconnections are directly exposed to ESD events. A correctly specified ESD protection device would be required for this.

Figure 6P Non-linear transmission line termination methods

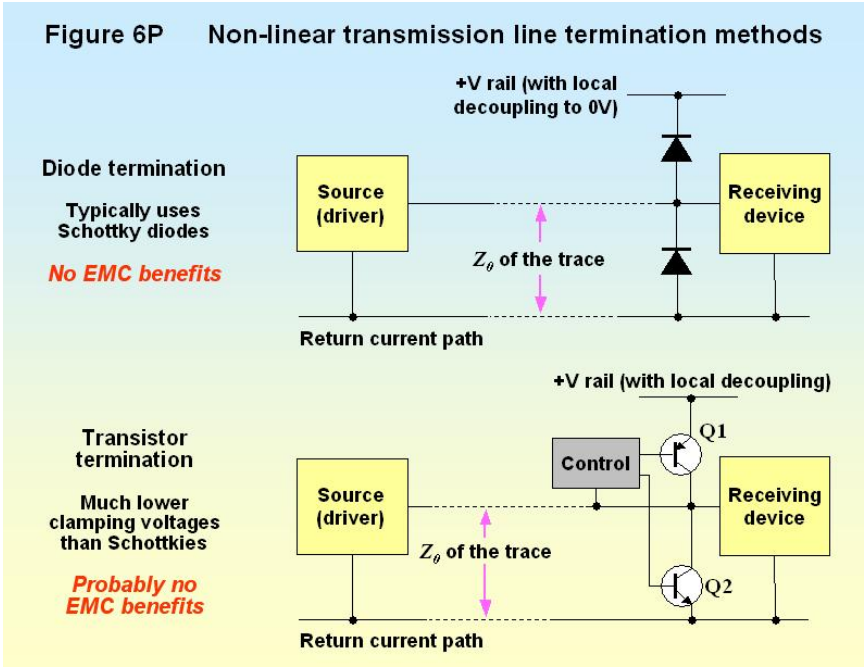

Diode termination techniques use Schottky diodes for the best clamping of overshoots, but even these are become ineffective at the low values of IC voltage rails used these days (e.g. 1.2V) because their diode voltage drop is too large in comparison. At least one company has developed 'active clamping' devices that use transistors to clamp an overshooting signal to the PWR or 0V rails with very small voltage offsets. It has been claimed that they achieved SI as good as resistor termination, but the author is not aware of any data on EMC performance and suspects that it would not be good. Where multi-drop and/or bi-directional busses are used with variable populations of receivers (and/or drivers), achieving correct resistive or RC termination of all possible loading situations can be very difficult indeed, especially if users cannot be trusted to fit 'terminating plugs' in the unused positions. Using non-linear terminators such as the above active clamp devices would be a solution as far as SI was concerned, because they would not all appear in parallel (as shunt terminating resistors would). They are only in circuit when the signal tries to overshoot one of the rails. However, as mentioned above, the EMC of such a system is not expected to be very good.

2.7 'Equalising' terminations

As mentioned earlier, for good SI and EMC (emissions and immunity) digital and analogue signals generally require resistive terminations for their transmission lines, using resistors that behave resistively up to the highest frequency of concern. And, also mentioned earlier, the complex impedance of a load (usually its capacitance) can cause a local reduction in *V*, hence a local reduction in $Z_{>0}$, possibly causing problems.

But when using incident-wave switching there are methods for making line terminations that use complex impedances (networks of resistors, inductors and capacitors) that compensate for the load's impedances. For more on these techniques read "*Constant-Resistance Termination*" and "*Constant-Resistance Equalizer*" in [12].

3 Transmission line routing constraints

3.1General routing guidelines

Transmission line routing should follow the guide below. These 'rules' are very important for traces carrying signals or noises with rise/fall times under 1ns, or carrying sinewave signals or noises above 300MHz. These rules are also useful for any high-speed signals.

- Route power supply decoupling first of all. These traces should be very short (see [16]) so won't restrict routing elsewhere on the PCB. This will probably require hand routing, not autorouting.
- Next, route the most 'aggressive' and most 'sensitive' signal nets. Aggressive nets include all clocks, write strobes, output enables and chip selects; sensitive nets include edge-triggered clocks, interrupts, asynchronous presets/resets/sets, and all resets.
- Make the most aggressive (or most sensitive) nets as short as possible, moving components as required (within component placement constraints). This will probably require hand routing, not autorouting.
- Route the most aggressive (or most sensitive) nets on the PCB layer that is closest to the 0V plane of a 0V/power plane pair.
- Route data busses and other transmission lines next, and keep them well away from the above lines to minimise crosstalk.
- Route all transmission lines (and any high-speed traces) on one PCB layer, where possible (see later). This is most important for any very aggressive (or sensitive) lines.
- Use striplines instead of microstrip (for EMC, if not for SI). This is most important for any very aggressive (or sensitive) lines.
- Don't route any traces over breaks in their reference planes. This is most important for any transmission lines, especially very aggressive (or sensitive) lines or other high-speed interconnections.
- Don't route transmission lines close to edges or gaps in their reference planes. Ideally route at least 5mm away, or five times the trace width, whichever is the greater.
- Don't use mesh or grid routing, unless all the resulting 'stubs' are so short that they don't create significant impedance discontinuities. Instead, it is best to route directly to use 'daisy chain' routing (as in Figures 6K and 6Q) or 'star' routing (see Figure 6R). When star connecting lines: all of the lines appear as parallel loads to their common driver (e.g. three 100Ω lines look like 33Ω), so ensure the driver can drive the resulting impedance.
- When all of the decoupling and transmission line (or high-speed) nets have been routed, the rest of the PCB can be routed.

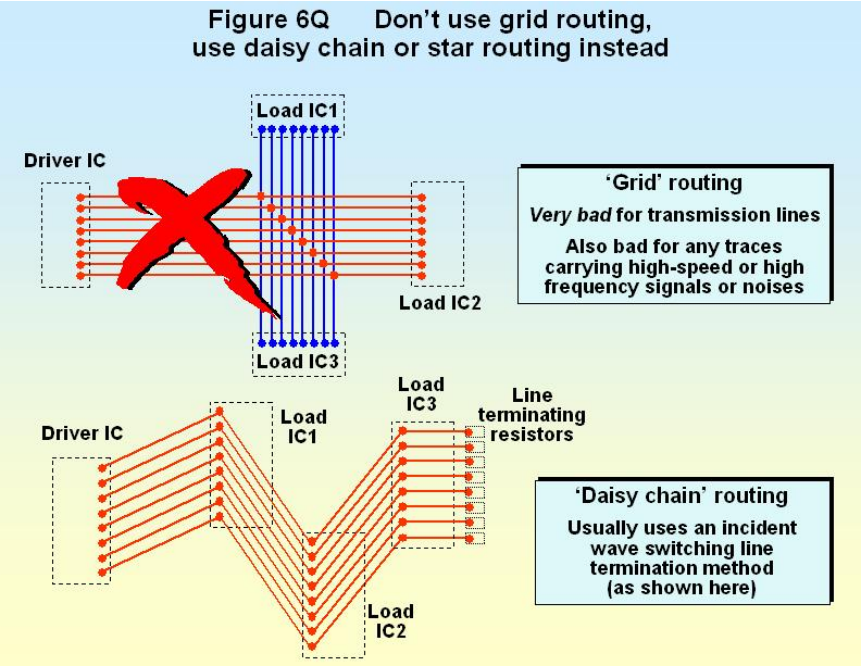

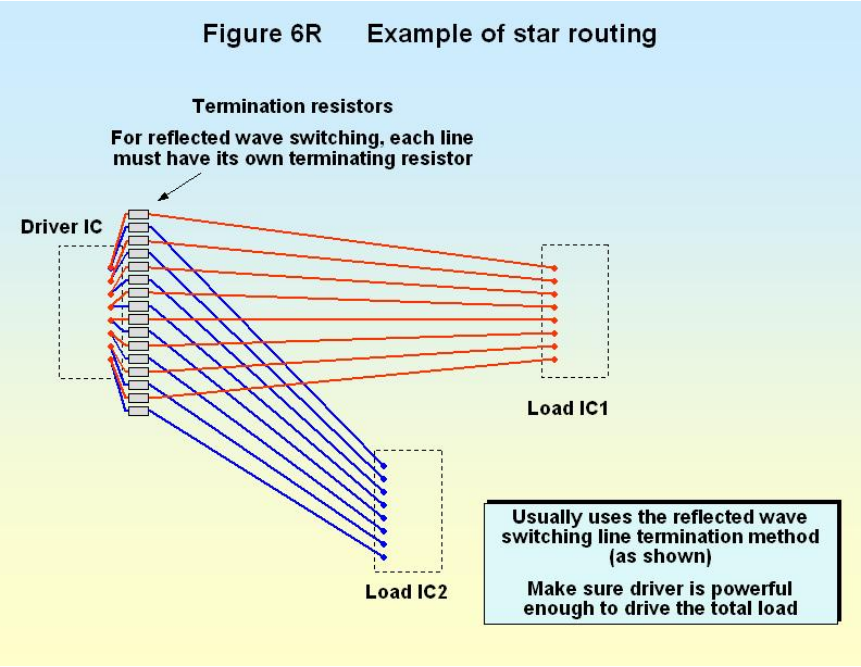

Ensure each load on the line is much greater than Z_0 – except for any parallel termination resistors, of course.

Many PCB designers with good EMC experience now refuse to use autorouters at all.

3.2 A transmission line exiting a product via a cable

Where a transmission line leaves a PCB (and the reference planes in the PCB) but is carried by a *shielded* controlled-impedance cable, it is important that the shield of the cable has 360° electrical bonding to the reference planes used for the return currents in the transmission line. 360° electrical bonding of shields is described in Part 2 of [1] or [2], and Volume 2 of [5]. For differential transmission lines, it may be important to use filtering at connectors even when shielded cables and connectors are used, see later.

Where a transmission line leaves a PCB (and the reference planes in the PCB) but is carried by an *unshielded* controlled-impedance cable, it is important for EMC that filters are located at the point where the signals leave their reference planes. Sections 6.2, 6.6, and 7, figure 2P and figures 2S – 2Y of [8] are particularly relevant. The filtering should remove all the signal and noise frequencies that are not required for the reliable communication of the wanted signals via the cables. The best filters employ common-mode chokes, and these are especially important for serial data communications in very noisy environments (e.g. CAN bus) or where high data rates are required (e.g. Ethernet, USB2.0, Firewire).

The connectors and cables used when routing a transmission line off-board must maintain exactly the same characteristic impedance as the trace(s) in the PCB. We are all familiar with the venerable BNC connector, available in 50Ω and 75Ω versions, and the 50Ω and 75Ω coaxial cables used to connect to them, but there are many more types of controlled-impedance connectors and cables available. In general, for good EMC, avoid coaxial cables and connectors and instead use types that route the return current path in its own conductor, so that there are always two wired conductors per transmission line, such as shielded twisted pairs. Cat 5, 5e, and 6 unshielded twisted pairs are good cables for SI, but their shielded versions are much better for EMC.

When using a shielded cable, the shield should always be 360° terminated at both ends wherever the signal carries RF signals or noises – as even low-frequency analogue signals do these days, due to the common-mode noise from switch-mode power converters and digital electronics. Also use 360° termination at both ends whenever there are RF electromagnetic fields in the environment – as there always are. Of course it is possible for low-frequency analogue signals to be heavily filtered, at their inputs and outputs from a product, so that they do not need cable shielding for good EMC.

Where a transmission line exits a metal (or metallised) enclosure (or chassis), there should be at least one very short electrical bond between the 0V plane of the PCB and the metalwork of the enclosure or chassis. This is normal good practice for EMC [19]. Where the shield of a shielded cable is bonded to the PCB's 0V plane, it can also provide the necessary 0V plane – chassis bond at the chassis by 360° electrical bonding in the chassis-mounted connector or cable gland.

3.3 Interconnections between PCBs inside a product

It is by far the best solution for EMC to use a single PCB for each product, so that there is no need for any interconnections between PCBs. Even mezzanine boards, daughter boards, and plugging boards into a backplane are not ideal from an EMC point of view, because of the difficulty of effectively routing the reference plane(s) through the connectors and cables between the PCBs.

Non-EMC advantages of the single-PCB-per-product solution include…

- Saving the costs of the connectors and cables
- Saving the costs of assembling the cables and connectors
- Saving the costs of rework (connectors are usually the most unreliable part of a product)
- Reducing warranty costs (connectors are usually the most unreliable part of a product)
- Improving product quality and customer perception (because of the increased reliability through using fewer connectors)

Flexi-rigid PCBs can sometimes be used in place of separate PCBs connected by connectors and cables. Their use is often not even considered, because of the higher costs of their bare boards, but this can be short-sighted because they may be able to reduce the overall cost-to-make, and the warranty costs, of a product.

Wherever a transmission line (or other conductor carrying high-speed data or high-frequency signals or noises) exits a PCB – even if it is only travelling a short distance to another PCB – it may need to

be treated as if it were leaving the product (as described in the section above). This is vital for products that are intended to be enclosed in unshielded enclosures. Ribbon cables can generally only be used for data communications in unshielded products if their signals are filtered to remove high frequencies from signals, and high-frequency noises.

Where a serial data bus connects between PCBs, a twisted-pair or shielded cable can be used. But it has been common practice in some areas to route parallel data busses between numbers of PCBs in a product, usually using ribbon cables. The EMC of such a design is very difficult without good shielding, even when the signals are filtered as described earlier. The SI and EMC of such product designs can be significantly improved by ensuring that each signal pin or conductor has an adjacent pin or conductor to carry its return currents. Where more than one reference plane is carrying return currents, more than one adjacent pin or conductor will be needed. In a ribbon cable, the outermost conductors should always be 0V or PWR, not signal.

The same guide applies to inter-board connectors where PCBs plug together. Figure 6S shows an example of a 'backplane bus' system, using a parallel data bus in which each conductor is a transmission line, and plug-in boards or modules. This is a very common type of product in some areas of industry, but the maximum data rate of the bus is limited by the lengths of the 'stubs' created by the connectors and traces of the plugged-in cards. A short stub looks like a capacitive load at a point on the line, whereas a long stub will create its own reflections if not terminated, and the point where it connects to the main line will have an impedance of half the line's Z_0.

High-speed busses in backplane systems like this might transfer data at up to a 50MHz clock rate. Some designers have managed to get such busses to run at 100MHz, and even 200MHz has been achieved with heroic design efforts. The problems of getting parallel busses to work at speeds as high as 200MHz or above is one reason why point-to-point serial busses are becoming so popular. A single 2.5Gb/s serial differential transmission line carries more data than a 16-bit parallel bus clocking at 100MHz, and the design of its transmission line is much, much easier.

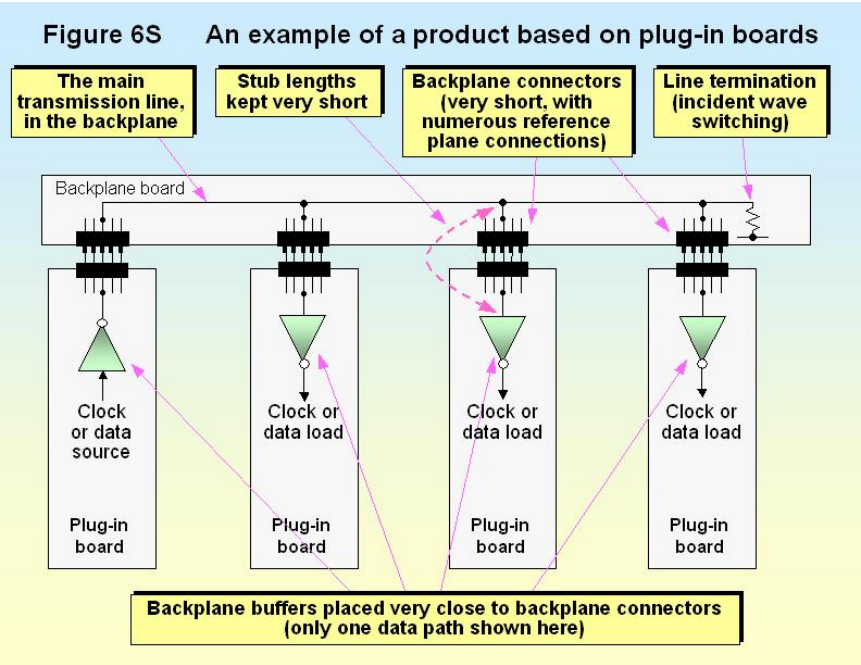

Figure 6S also shows how important it is to place the data buffers as close as possible to the backplane connectors, to minimise stub lengths. The backplane connectors will be controlled-impedance types that match the line's impedance, and they must have very short signal paths if high-speed is required. As was mentioned earlier, the connectors must also make numerous connections for all of the reference planes associated with the signals that pass between the boards. These

reference plane connections must be spread over the whole length of the connectors – and if the connectors are very wide (for example, some backplane connectors have five or more rows of pins), they should also be spread over the width of the connector.

3.4 Changing plane layers within one PCB

Ideally, transmission lines (and high-speed signals) should be routed on a single layer in the PCB, over solid (unbroken) reference planes that extend a very great distance on both sides of their traces and well beyond both ends. But sometimes there is no alternative but to change layers along a trace.

Above about 10MHz, the plane return currents associated with traces are forced by 'skin effect' to flow in the surfaces of the reference planes or traces that are closest to the path of the opposing current (send or return). When traces change layers, their surface RF return currents simply *cannot flow through the thickness of the plane* to follow the route of trace. It may seem odd that electricity cannot flow through a metal plate, from one side to the other, but at high frequencies this is just what happens. (It is the same effect that allows thin sheets of metal to be used as RF shields – so we can think of it as one side of a plane shielding its return currents from the other side.) This self-shielding creates difficulties for transmission lines (and for any traces carrying high-frequency signals or noise).

Where a trace is routed against the other side of the same plane, its return current has to flow to the other side of the plane *around the rim of the antipad in the plane* created by the via hole for the trace. This forces the return current to be concentrated at the hole's location, which increases the inductance and alters the characteristic impedance at that point.

Figure 6T An example of current flow when changing layers from one side of the reference plane to the other

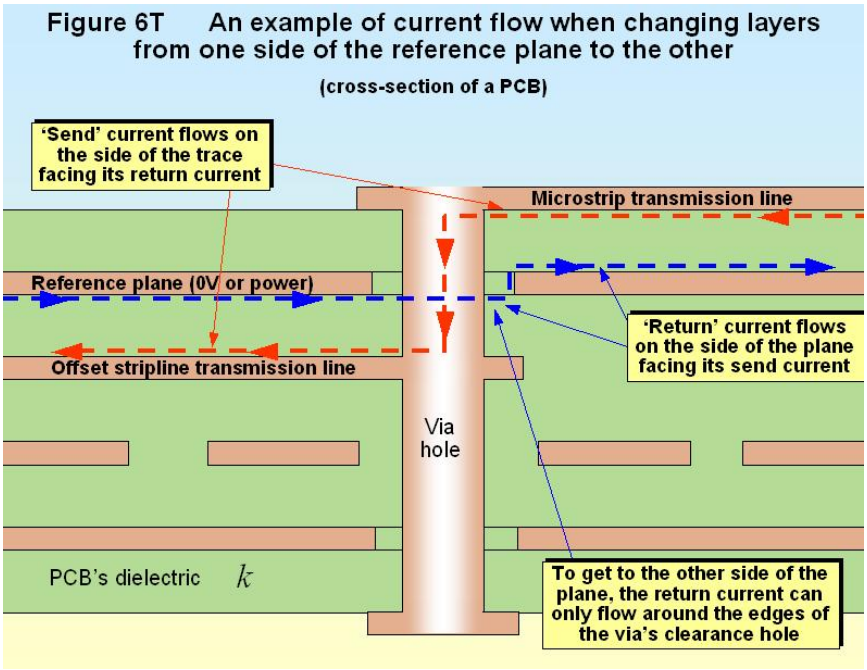

Where a trace is routed against a different plane at the same voltage (usually 0V) the interplane capacitance is insufficient to make a good connection – so two via holes, one on each side of the trace should be added to link the two planes together. These via holes should be in a line with the via carrying the signal, perpendicular to the trace direction, and within 2mm of the trace via. (Where this is not possible, make sure that there is at least one plane-linking via as close as possible to the trace via.)

The inductance associated with the return path current as it flows in the via(s) will affect the local value of Z_0, (see "*Via Inductance*" and "*Via Inductance II*" in [12]), and this can be important where

signals (or noises) have rise/fall times under 1ns or frequencies above 300MHz.

But where a trace becomes routed against a reference plane with a *different* voltage, it is important to connect a capacitor between the two planes, very close (i.e. < 3mm) to the trace's layer-change. Two capacitors, one on each side of the trace, are better than one. The return current will pass through the capacitor(s), so the value and type of capacitor and its self-resonant frequency, should be chosen to suit the part of the spectrum that needs to be well controlled.

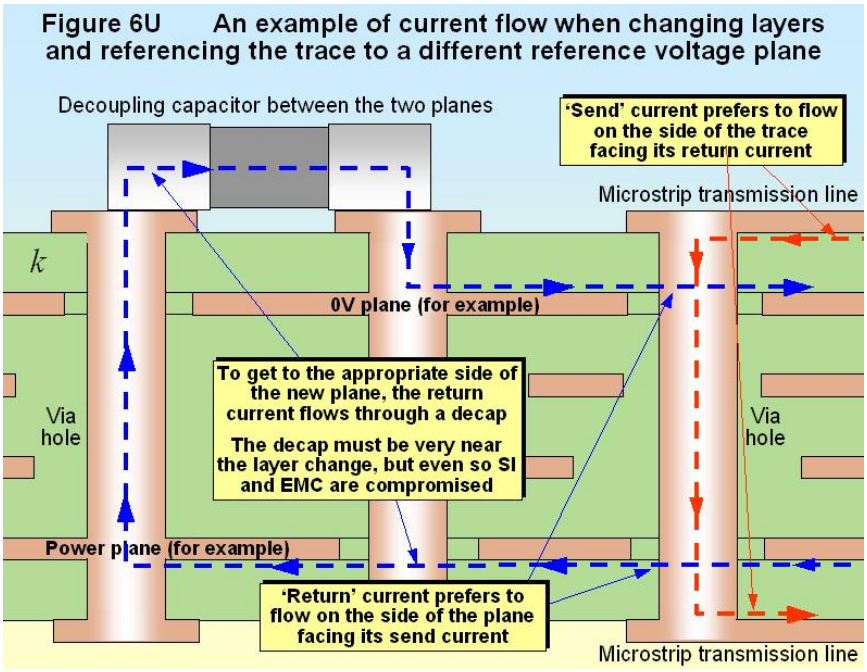

For transmission lines carrying signals with rise/fall times of 300ps or less (or sinewaves of 2GHz or more) it can be difficult to maintain the desired Z_0 over any portion of a transmission line for which the signal and its return currents are carried by vias. But it will be almost impossible to maintain the correct Z_0 where capacitors are used in the return path. For such signals it can be very important to have no layer changes *at all*, or to only allow layer changes that are either side of the same plane.

When changing layers near to a break or edge in a reference plane, the layer changes should be located at least 5mm from the nearest edge of the plane, preferably even further away.

Where there is more than one plane in a PCB, there are noise voltages between the planes. Via holes couple strongly with interplane voltage noise – so signals passing through via holes tend to add to the noise between the planes. The noise between the planes also tends to couple noise into the signals passing through the vias. Higher frequency signals and noise couple more strongly, so this is more of a problem for signals or noise with fast rise/fall times or high frequencies. This is another reason for referring all critical traces to either side of the same plane – and (better still) for not changing layers along their route.

For traces routed on internal layers in a PCB, there will always be a layer change (or a joint or 'spur' in the trace) at each point where the trace must connect to a device mounted on an outer layer. For signals with edges shorter than, say, 150ps there are significant problems in dealing with the mismatches caused by these layer changes or 'stubs' (see later). This is the main reason why high-data-rate PCB interconnections are now abandoning traditional parallel multidrop bus technology (e.g. PCI) – and using point-to-point LVDS serial data communications (e.g. PCI Express) instead.

These design constraints can lead to an increase in the number of plane layers – but the signal integrity requirements may leave the designer no choice but to increase the number of layers. Where achieving adequate EMC at PCB-level requires an increase in the number of PCB layers and hence the bare-board costs, the usual temptation is to hope for the best and use the lowest number of layers that will permit functionality. But this is usually a false economy – increasing the number of layers in a PCB is usually the lowest-cost and quickest way to deal with EMC.

When changing layers, don't forget to change the trace width as required to maintain the same Z_0 for each section of the trace. Particular care is needed when using both stripline with microstrip sections along the same trace (see later).

3.5 Crossing plane breaks or gaps within one PCB

Ideally, no traces should ever cross plane breaks or gaps unless some kind of common-mode isolation (or at least filtering, see figure 4P of [20]) is applied at the location of the break, to both the send *and return* paths of the signal. Transmission-line traces that cross plane breaks or gaps should be closely-coupled coplanar types, with return traces on either side of the signal trace. Closely-coupled coplanar differential transmission lines (see later) are the best type for this purpose.

If there is no choice but to cross a break or gap in a reference plane with a type of transmission line that employs a reference plane, fit a 'stitching capacitor' on both sides of the trace – within 3mm of the trace on each side. Take care to minimise the inductance of the capacitor (e.g. by using a very small surface-mounted type, such as an 0201) and also to minimise the inductance of its interconnections to the planes on either side of the split.

Where a reference plane break changes from a 0V plane to a power plane, the stitching capacitors are effectively part of the power decoupling system and could be called 'decaps' (see [16]) – but this does not mean they can be moved further away from the trace.

As above, it will be almost impossible to maintain the correct Z_0 where capacitors are used in the return path to cross a break in a reference plane. So, where signals (or noises) must cross a plane break, they should be ones with very long rise/fall times, or where the highest frequencies of concern are not very high.

As was mentioned above, when routing transmission lines close to edges or gaps in reference planes, or when changing layers near to an edge or a gap - the routing and/or layer changes should be located at least 5mm from the edge or gap (preferably even further away).

3.6 Avoid sharp corners in traces

A sharp corner in a trace creates an impedance discontinuity due to the stray capacitance from one part of the trace to another, around the corner. Some studies have shown that 90° corners in a trace only have a significant effect on SI where rise/fall times are 10ps or less. Nevertheless Intel and National Semiconductor recommend keeping the number of corners (and vias) in high-speed traces to a minimum, and 'chamfering' using two 45° bends to achieve a 90° bend (for example), or using smooth curves instead.

But PCB techniques for achieving adequate high-frequency SI are much less demanding than when those same techniques are used to achieve adequate EMC at PCB level. So even where sharp corners are acceptable for SI reasons – it is best for EMC to chamfer or curve them where real risetimes could be under 1ns.

Figure 6V Avoiding sharp corners in traces

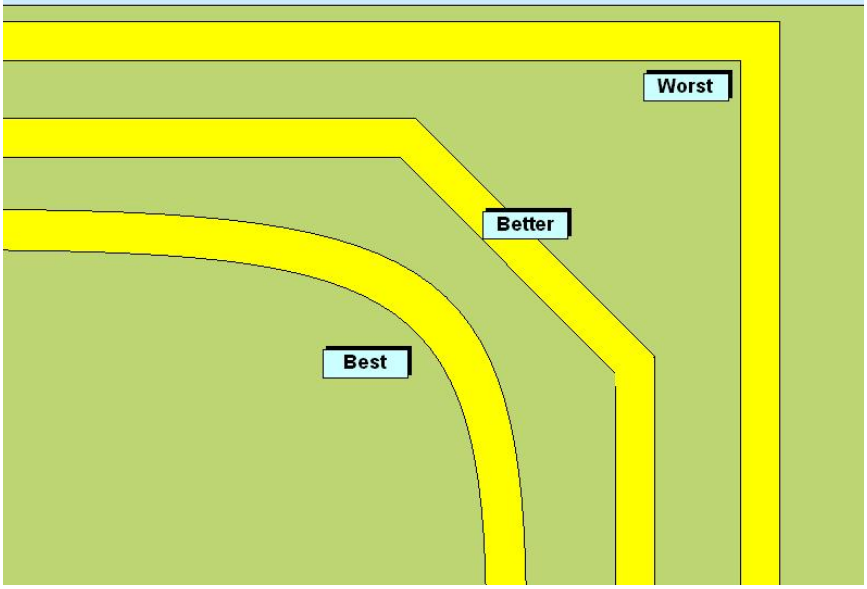

Of course, a via hole also represents two sharp 90° corners in a trace, which is another reason why transmission lines, and any traces carrying high-speed signals or noises, should ideally be routed on a single layer with no vias along its length, even if this means increasing the number of PCB layers. As has been mentioned many times in this series, increasing bare-board PCB costs to improve EMC is almost always much more cost-effective and quicker than any other EMC measures (other than those applied to the silicon die in the ICs).

3.7Linking return current planes with vias or decaps

To prevent cavity resonances from occurring in the frequency range of concern for EMC, all the 0V planes (including any small or broken plane areas) should be interconnected with via holes at least every λ?/10 at the highest frequency of concern, taking the dielectric constant of the PCB into account as described above. So, for example, to prevent cavity resonances in a plane pair in an FR4 PCB at frequencies up to 1GHz, plane-linking vias should be located at least every 15mm *all over* the planes' areas.

Where the planes forming a cavity have different potentials, they should be linked by decoupling capacitors (decaps) instead of vias. These issues were discussed in more depth in Parts 4 and 5 [18] and [16] of this series. Linking planes together with vias or decaps is especially important when striplines are used. For a stripline, the return current flows in both of its planes, above and below the trace, and it is important to link these currents together at distances less than λ/10. However, the *close* proximity of via holes or decaps along the route of the trace can add capacitive loading and affect its Z_0 and its *V*.

3.8 Effects of via stubs

The 'unused length' or stub of each via hole acts like a microwave designer's 'tuned stub filter' that looks like a short circuit at resonance when its length is λ/4, 3λ/4, etc. Signals with frequencies that are within about ±30% of the resonant frequency will also suffer significant attenuation [21] [22]. The extra capacitance of a via hole decreases the local propagation velocity when compared with a normal PCB trace, so in real-life the frequency of maximum attenuation will be significantly lower than that calculated from λ/4, 3λ/4, etc. using the normal values for FR4.

Considering the λ/4 stub resonance for example, via stubs with the following lengths can be expected

to resonate and strongly attenuate signals over the following frequency ranges:

- 5mm (typical of a 22 layer backplane): 3 - 6 GHz (maximum attenuation around 4.5GHz)
- 4mm: 4 - 7 GHz (maximum attenuation around 5.5GHz)
- 3mm: 5 - 9 GHz (maximum attenuation around 7.3GHz)
- 2mm: 8 - 13GHz (maximum attenuation around 10.5GHz)

For signals that have their third or fifth harmonics in the above ranges (especially near the frequency of maximum attenuation) the result is waveform distortion and a more closed eye pattern, in other words a degraded SI. At the same time emissions are increased due to the poor transmission line mismatching and increased reflections occurring over these frequency ranges.

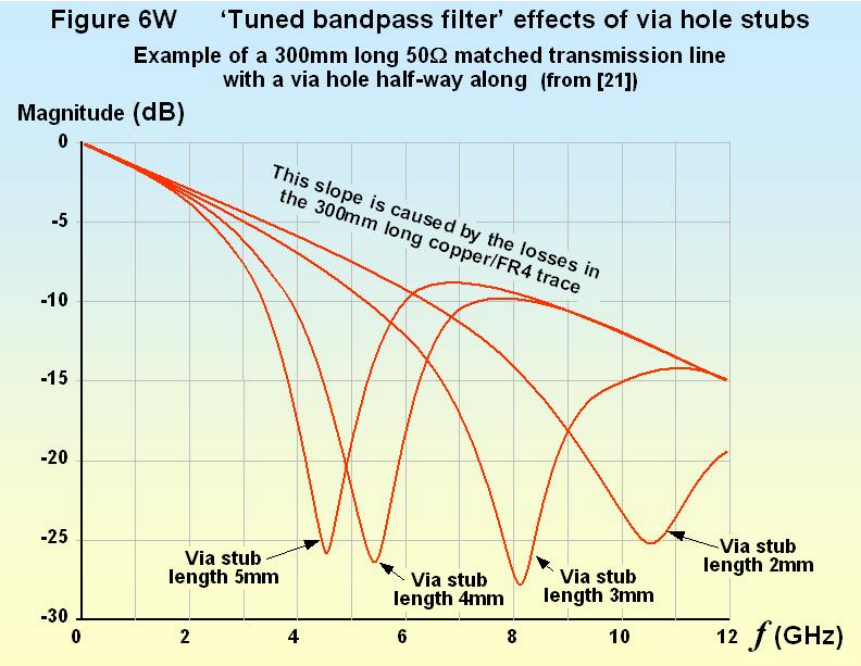

So, when using Through-Hole-Plate (THP) PCBs it is best to keep PCB thickness to a minimum, ideally so that the problems caused by the length of the stub occur at higher frequency than is cared about. Alternatively, counterbore the PCB to remove the unused lengths of their vias (their stubs). This sounds like a costly exercise, but it can use the same machines as are used to drill the PCB in the first place, and Teradyne engineers have estimated that the bare-board cost would increase by about 7% [22]. They also found that it was possible to back-drill vias without harming PCB quality.

But microvia/HDI PCB technology (see Part 7 of this series) uses via holes that are just as long as needed for a signal to change layers, so does not suffer from via stub problems like THP technology does [22].

3.9 Effects of routing around via fields

When using THP PCB technology, arrays or fields of via holes such as those associated with ball grid array (BGA) devices and multiway connectors cause high levels of perforation in reference planes. It is very important to ensure that at least a thin 'web' of copper remains around all the via holes' antipads in each plane layer, but even so an impedance error is introduced in transmission lines that pass over such perforated areas. As BGA pin pitches reduce below 1mm even retaining a thin copper web between via holes becomes increasingly difficult – ultimately impossible.

This plane perforation reduces the shielding provided by the planes – increasing crosstalk and degrading SI. Also, a perforated plane has a significantly higher inductance, which increases the Z_0

of traces routed over such areas and can also create SI and EMC problems due to reflections at the resulting impedance discontinuities [23] when rise/falltimes are short enough. From an EMC point of view a perforated plane does not behave as an image plane as well as a solid plane does, so emissions are increased and immunity worsened.

The best solution to this problem is to use microvia/HDI PCB technology. This does not require every via to pass through every layer on a PCB, and so it can enjoy unperforated reference planes (see Part 7 of this series). But when THP technology is used the problems of reference plane perforation in via fields should be taken into account in the design. If the signal and noise rise/fall times are much more than twice the propagation time it takes to cross the perforated area, the effect of the perforation will be 'smoothed out' and should be easier to design around. But where rise/fall times are shorter, more detailed design analyses and TDR measurements on prototype board assemblies will probably be required, for optimum SI and EMC.

3.10 Other effects of the PCB stack-up and routing

FR4 has a nominal dielectric constant, *k*, of 4.7 at the usual measurement frequency (100kHz), but it can vary between 4.0 and 5.5 at that frequency, and it reduces as frequency increases so that is nominally 4.2 at ≥ 1MHz. Grades of FR4 with controlled *k* are readily available (e.g. 4.7 ±0.1 measured at 100kHz), cost very little more than the regular material, and should always be used for FR4-based PCBs using matched transmission lines. When not using FR4, it is still important to use dielectrics with specified worst-case limits for their *k* values. See a later section for a discussion of PCB dielectrics.

The PCB manufacturer uses the following types of materials in a PCB's 'stack-up'…

- Caps. Dielectric layers laminated with single-sided copper foil ready for etching to create the top and bottom layers of the PCB.
- Cores. These are cured flexible dielectric layers laminated with double-sided copper foil ready for etching to create internal layers for traces and planes.
- Prepregs. These are simply layers of uncured dielectric. They are used as insulating spacers between the copper foil layers of the cores, and of course they glue the PCB together when cured.

Only a limited range of dielectric thicknesses are available to the PCB manufacturer for caps, cores and prepregs – and this places limits on the ranges of trace widths that a designer can use to achieve the desired Z_0s. So PCB designers should always work closely with their PCB manufacturers, to design trace geometries that they can make at low cost that will achieve the desired Z_0s.

The dielectric constant *k* of FR4 varies with resin content and thickness (as well as with frequency, temperature and humidity). Because it is a mixture of high-*k* glass fibre and lower-*k* epoxy resin, the usual assumption of 4.2 is only true for the bulk material. In the stack-up of a real PCB, resin-rich areas can be found close alongside each trace and in-between closely spaced traces on the same layer, such as differential pairs. The resin has a *k* of around 3.0, so these resin-rich areas tend to cause Z_0s to be lower than expected.

When designing a PCB stack-up, designers tend to choose round numbers for the thicknesses of the 'prepreg' layers. But the finished thicknesses of caps, cores and prepregs – and hence the actual characteristic impedance achieved – depends upon the PCB construction, lamination pressure, heat input rate and border damming (to stop resin being squeezed out) during lamination. It even depends on the copper patterns etched into the foils opposite the traces concerned.

For example: when a PCB is laminated but before it is cured, the 'core' layers are already cured and

the prepreg layers are not. So the etched copper layers on the cores press into the prepreg layers, displacing a small amount of their material and reducing the thickness of the dielectric over the copper traces. When thin prepreg layers are employed, the difference in their thickness can have a significant effect on Z_0. But the wider the area of copper, the less is this effect, so that for two large planes facing each other the reduction in prepreg thickness can be negligible even with the thinnest prepregs. This issue, and other similar ones, are discussed in "*Offset Stripline – Construction Matters*" in [24].

Figure 6X An example of a PCB stack-up and its impedance calculations (from [24])

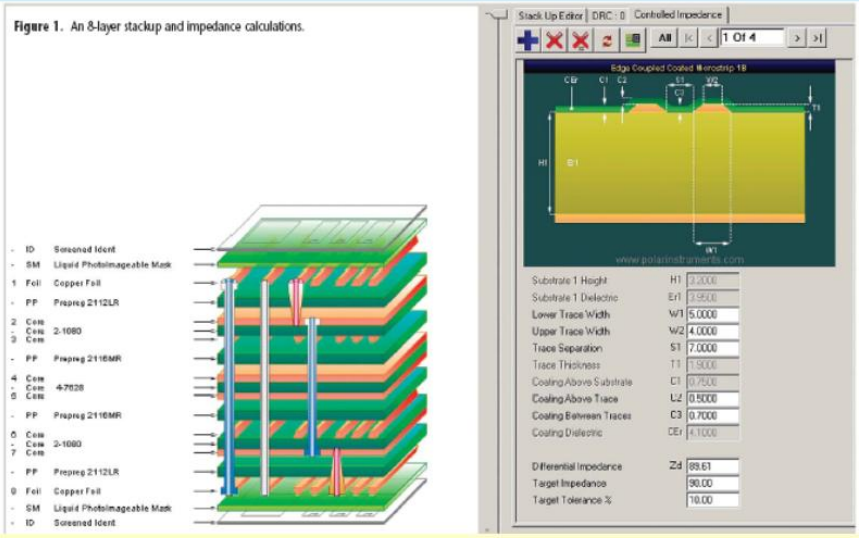

When all of these effects are considered alongside the need for very good impedance matching to improve EMC and avoid adding cost to the product by improving its shielding and filtering – we find that manual calculations are impossible and 'trace tuning' board iterations are too time-consuming. We need to use PCB design tools that can handle these issues, and tools such as the SB200 stack-up tool and SI8000 impedance solver from Polar Instruments [24] [25] can help.

A later section discusses the use of computer-aided design tools to help reduce unit manufacturing cost and time-to-market. TDR is also a very valuable tool for these purposes (see above).

An increasing number of PCB manufacturers are becoming very skilled with transmission line PCBs (also known as 'controlled-impedance' PCBs) and some of them are equipped with sophisticated simulators and measuring equipment. The simulators help them help designers achieve what they need, and the measuring gear confirms that what was required is actually been achieved. These manufacturers can be a great help to a novice transmission line designer, as well as a good 'technology partner' for experts.

But some PCB manufacturers, especially at the low-cost end of the market, may not be very helpful, and/or may not be very well equipped for reliably manufacturing controlled-impedance PCBs. Some may not have much of a clue. Their salespeople may nevertheless manage to create the impression that they will be able to work with you to easily get up their learning curve and still produce cost-effective and reliable controlled-impedance PCBs on time. Do not even *imagine* that you can employ such companies to make a PCB using modern sub-micron devices and transmission line technology. Since manufacturing buyers naturally prefer to use the cheapest possible PCB suppliers – the project manager (or someone else with sufficient authority) should always keep an eye on who they are buying from to prevent them 'saving money at any cost'.

3.11 Some issues with microstrip

Microstrip traces (those that are routed on the top or bottom layers a PCB) have FR4 or other PCB dielectric material on one side, and air on the other side. The result is that the effective *k* value they experience is less than that of the bulk PCB material itself. All electrical signals and noises are propagating electromagnetic waves, and in the case of microstrip traces some of the electric waves and half of the magnetic waves travel through the air, whilst the rest of the wave energy travels though the PCB dielectric. The wider the microstrip trace, the greater the proportion of electric waves travelling through the PCB dielectric – so for a microstrip, the effective *k* value depends on the trace width. As a result, the Z_0 and *V* of a microstrip line depend upon its width.

Another issue for microstrip traces is that they are usually coated with a solder resist layer, and sometimes conformally coated to prevent moisture ingress. Sometimes PCBs are encapsulated in silicone or resin or submerged in containers filled with special oils, to improve their reliability in harsh environments. Such layers, coatings or encapsulations have a *k* that is very different from air (*k* =1) and can have a significant impact on a trace's characteristic impedance and the accuracy of its matching (see later). Sometimes beads of condensation or even films of water can occur on a PCB's solder resist, and since the *k* of water is 80 this can have a very significant effect indeed on the Z_0 and *V* of microstrip traces. Where a trace changes layers so that part of its route is stripline and part is microstrip, these different parts will have different velocity factors and see different values of *k*. Unless care is taken over such mixed-type striplines, SI and EMC can be degraded by skew and mismatches.

12 References

[1] "Design Techniques for EMC" (in six parts), Keith Armstrong, UK EMC Journal, Jan – Dec 1999
[2] "Design Techniques for EMC" (in six parts), Keith Armstrong, http://www.compliance-club.com/keith_armstrong.asp
[3] "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 1", M K Armstrong, IEE Electronics and Communication Engineering Journal, Vol. 11 No. 4 August 1999, pages 185-194
[4] "PCB Design Techniques for Lowest-Cost EMC Compliance: Part 2", M K Armstrong, IEE Electronics and Communication Engineering Journal, Vol. 11 No. 5 October 1999, pages 219-226
[5] "EMC and Electrical Safety Design Manuals", a set of four volumes edited by Keith Armstrong and published by York EMC Services Ltd., 2002, ISBN 1-902009-07-X, sales@yorkemc.co.uk, phone: +44 (0)1904 434 440.
Volume 1 – What is EMC? ISBN 1-902009-05-3
Volume 2 – EMC Design Techniques – Part 1 ISBN 1-902009-06-1
Volume 3 – EMC Design Techniques – Part 2 ISBN 1-902009-07-X