advanced layout solutions

Tel: 0118 971 1930
Email: info@alspcb.com

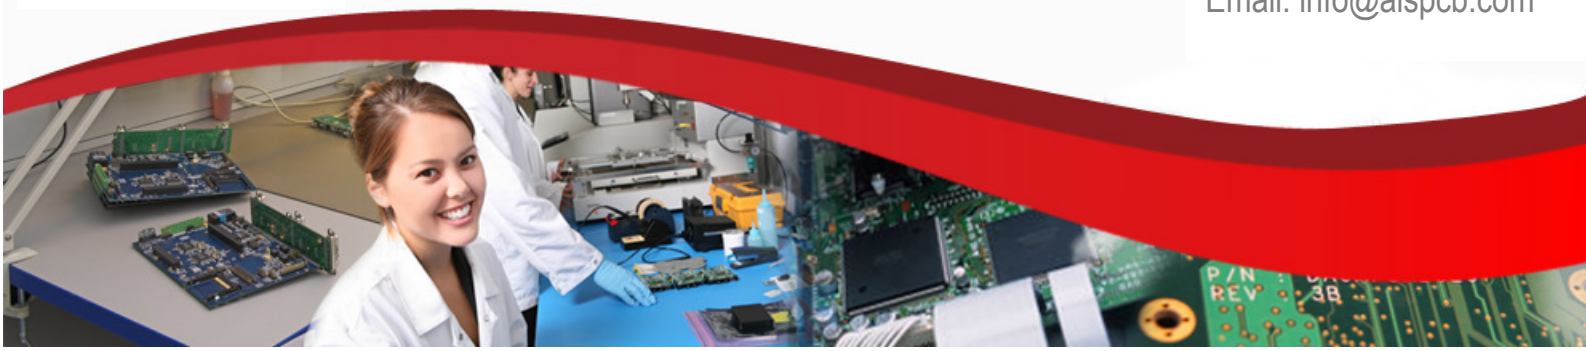

Controlled Impedance

What is it? Why is it important? How is it controlled through the PCB design flow?

by chris halford
Signal Integrity Engineer
Advanced Layout Solutions Ltd
chris.halford@alspcb.com

At Advanced Layout Solutions Ltd (ALS) we control impedance throughout the board layout process by utilising an accurately defined cross section, a trusted field solver and a constraint driven design flow.

But what is trace impedance?

If you measure the impedance of the 65 Ohm trace in Figure 1, your multi-meter will report ~10 Ohms. Faulty board? Of course not, just a slow meter. The meter is reporting the (negligible) resistance of the PCB traces plus the 10 Ohm resistor.

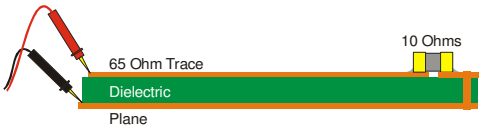

Figure 1. Test Circuit

However, it takes a finite time for a voltage step to propagate along a PCB trace, and if you could measure the initial instantaneous impedance before the test voltage reached the resistor, then it would read ~65 Ohms.

During this initial time we are not driving the 10 Ohm load at the end of the trace - the load being driven ***is*** the trace. The current is flowing into, and returning from, the microstrip structure itself.

It's not unlike using your hosepipe: prior to anything squirting out of the end, water flows steadily for a short time to fill the capacity of the empty pipe.

Whilst the signal (in this case the step voltage from the meter) travels down the trace, the driver sources current as the instantaneous impedance is overcome and the trace is, in essence, charged up. If this impedance is constant then it is *characteristic* of this line, and is denoted as 'Zo'. If we design a line to have a specific characteristic impedance throughout its length, then it is referred to as a *controlled impedance* line.

Consider Figure 2. In this diagram a signal is travelling along a trace with a return current path through the reference plane directly beneath it (a transmission line). As the signal propagates, current only flows through the dielectric at the signal's wavefront, where the voltage is changing. Ahead of this wavefront the trace has no part to play in the circuit; it is blissfully unaware of the approaching signal. In its wake the wavefront leaves a length of trace charged up to the new signal voltage, with no current (other than leakage) flowing through the dielectric. The (ideal) transmission line can be modelled as the long LC ladder shown, in which each element represents the L and C of the structure per-unit-length.

The unit-length inductance (L) of a printed circuit board trace is a function of the physical geometry of the conductors and dielectrics that make up the transmission line. In addition to geometry, the capacitance (C) is also a function of the relative permitivity (εr, DK, Er) of the dielectric material. Therefore, to control the impedance we must construct a trace of uniform cross-sectional geometry and consistent εr along its length for a given routing layer.

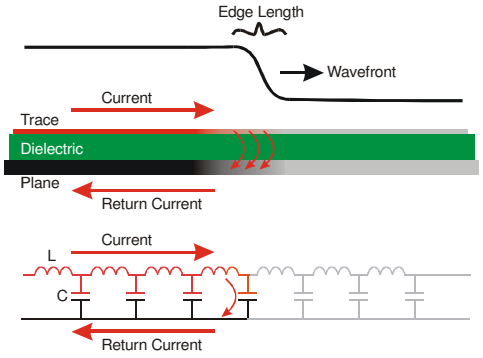

Figure 2. Transmission line propagation on PCB microstrip and equivalent ideal model

Who cares about characteristic impedance?

In general, it is an important issue for anyone designing systems in which the physical length of a switching edge is short in comparison to the length of the transmission line along which it is propagating. With today's fast rise-time outputs, this situation applies to almost all designs that we see from our clients.

Fast rise-times mean we are not dealing with a *lumped* system, we are dealing with a *distributed* system, in which the voltage seen on a particular trace may vary significantly along the trace length. In this mode an impedance discontinuity will cause a reflection whereby some of the signal's energy will reflect back down the line whilst the remaining signal will continue onwards, distorted. Examples of impedance discontinuities are variations in trace width, and the presence of vias, connectors, or the receiving device itself.

What impedance value should be used?

Within reason, the absolute impedance value chosen is not normally important, providing it is controlled along the entire length of the line. Other constraints in a design often dictate the impedance for us; it may be chosen based on a design specification (e.g. 65 Ohms for PCI) or chosen to reduce current (a high impedance). It will generally be between 45 and 80 Ohms due to typical material geometries, and if the signal changes layer then the trace geometry should be adjusted as necessary to maintain a consistent Zo.

Does Zo calculation have to be accurate?

Yes. There is generally a large tolerance on finished trace impedance (±10% is common), the purpose of which is to allow manufactures to achieve an acceptable yield. This tolerance is not there so that designers can approximate the ***nominal*** value. If a finished trace is to be 50 Ohms ±10% then, although a nominal 54 Ohm geometry will keep the layout tools happy, it does not give the fabricator much room to move (and one way or another, you'll pay for that lower yield).

Secondly, if you send the same design to several PCB fabricators you will find that they all want to 'tweak' your design in different ways. One supplier may want to use a different trace width to achieve a specific Zo whilst another may want to change to a different dielectric material. Producing a design with accurate impedance control makes the board more portable to multiple suppliers and reduces the number of impedance DFM issues when handing the design over to them. Finally if, like ALS, you run a constraint-driven simulation-based design flow, then it is essential that your stack-up is representative and that your nominal impedance calculation is accurate. This ensures that Z constraints are being correctly met and that interconnect models are representative when running simulations.

Accurate stack-up

In order to design for controlled impedance, the design's stack-up must be well defined - '8 layer in FR4' is not sufficient! At ALS we assign a realistic stackup: if your design includes a ply of Isola IS410 2116 pre-preg, then this is what will go in the cross section, along with the εr for that material's resin content at an appropriate operating frequency. We work closely with material suppliers and PCB fabricators to maintain a rapidly growing library of base materials from which designs are stacked.

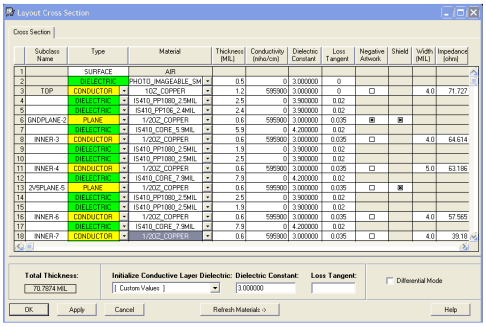

Figure 3. Allegro cross section editor

Field solver

With the stack-up defined, the Allegro 2.5D field solver is able to accurately calculate the characteristic impedance of traces by breaking down the area that surrounds the conductors into small 'mesh' elements and solving Maxwell's equations for each one. The accuracy of the tool can be increased by reducing the size of these mesh elements, though this comes with a penalty in the form of increased computation time. Figure 4 illustrates some of the advanced features of the field solver we use.

Trapezoidal Etch Angle can be defined to take into account the over-etched top surface of the conductors. This parameter, between 0 and 90 degrees, specifies the anticipated angle on the side-wall of the trace.

Conductive Layer Dielectric allows local εr variation to be taken into account by allowing a separate value to used between traces. When traces are closely spaced the gap between them tends to be devoid of glass fibres, making them extremely resin-rich and of a lower dielectric constant, increasing the impedance of edge-coupled structures. Use of CLD allows the yellow areas in Figure 4 to have a different permitivity value to the surrounding pre-preg into which the traces are embedded.

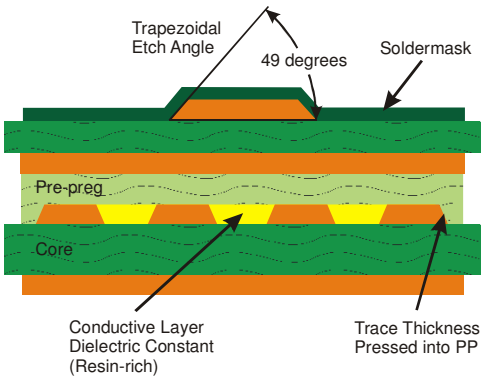

Figure 4. Allegro cross-section features

Solder-Mask Layers have the effect of lowering the impedance of the traces on the outer layers of the board. It is important that this coating is taken into account when calculating stripline impedances.

The tools most commonly used by PCB fabricators to calculate impedance are the offerings from Polar Instruments Ltd. Not only do Polar provide software solutions to predict impedance, they also provide hardware test equipment to measure impedance on the finished PCB. To create as few DFM issues as possible when handing your design over for manufacture, it is advantageous if your layout and simulation tools agree with the Polar impedance results. We compared three design tools with the Polar Si8000 field solver to see how closely the results correlate. Figure 5 shows the results of this test. In this stripline example the Allegro PCB SI and Polar Si8000 results are virtually identical.

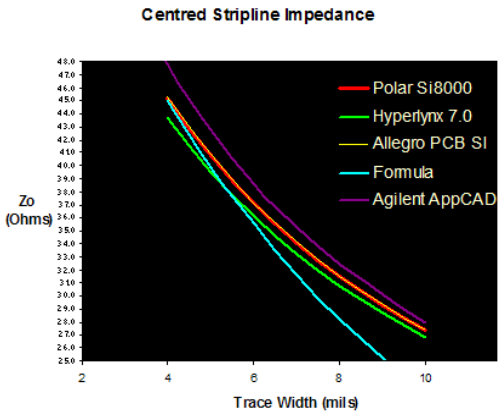

Figure 5. Comparison of field solvers
Details of this comparison can be seen at www.alspcb.com.

Constraint driven design

When using a constraint-driven design flow, accurate field solver results allow impedance constraints to automatically select the optimum trace width for a particular routing layer. This approach also allows any controlled impedance violations to be clearly seen as DRCs in the design, as shown in Figure 6.

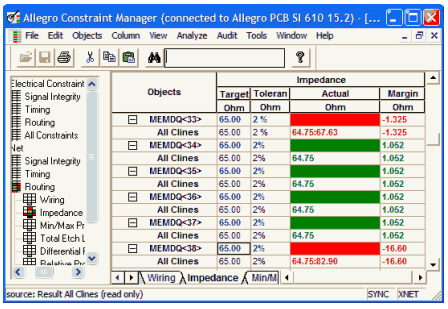

Figure 6. Allegro constraint manager - impedance violations are clear to see.

And now the bad news…

When designing trace geometries up front, field solvers will make the assumption that your trace has the luxury of being the only trace in the entire design. In this ideal view of the world there is no electromagnetic coupling between traces, except for the intentional coupling in the case of the differential pair. Figure 7 shows the electric and magnetic field lines of three structures when routed sufficiently far apart so as not to interact.

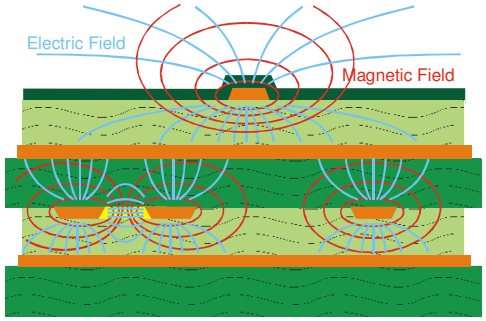

Figure 7. Field lines associated with single-ended and differential traces.

In reality there are usually many neighbouring traces routed in close proximity. The presence of these neighbours will both raise and lower the impedance of the surrounding traces, as they switch both in the same and opposite directions.

A reduction in impedance will be sustained when the signal on a neighbouring trace switches with opposite polarity (Zoo, odd mode). Conversely, an increase in impedance will result if neighbours switch with the same polarity (Zoe, even mode). Figure 8 shows the magnetic (red) and electric (blue) field lines associated with the three traces during odd and even mode drive.

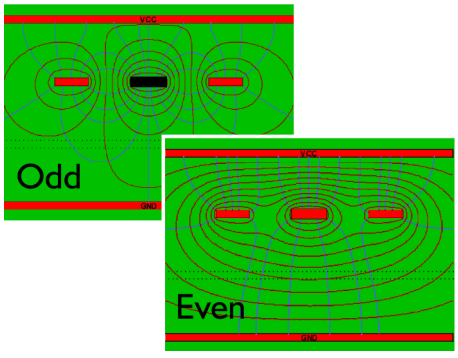

Figure 8. Odd and even mode field lines

From these cross sections it is clear to see why the impedance would shift. In odd mode the centre trace's electric field couples not only into the power planes but also into the neighbouring traces (of opposite polarity). The trace behaves as if it had more reference plane surrounding it, reducing its impedance. In the even mode case there is no coupling between the conductors and less capacitive coupling to the planes, causing the impedance to rise. The net result is that trace impedance will suffer dynamic variation as random data signals propagate through the traces. As these proximity effects are proportional to the field strength of neighbouring traces they can be substantially reduced by increasing trace-to-trace spacing.

Differential impedance.

Coupling can, however, be used to advantage in the form of differential impedance. When a differential pair is routed on a printed circuit board we assume that the signals will be of equal magnitude and opposite polarity. When these traces are routed together in close proximity (usually the minimum spacing possible) then the impedance seen on *each* conductor will not be Zo but Zoo (odd mode impedance). The differential impedance will be twice the value of Zoo and will be attributable to the proximity of both the reference plane(s) and the neighbouring compliment signal.

In a cable situation the differential pair is superb, allowing the linking together of pieces of equipment with different ground voltages and offering immunity to common-mode noise. In the PCB setting, however, their use is questionable, as they do not actually route together at all. A noisy aggressor signal will nearly always be closer to one signal than its compliment and, despite some coupling between the pair, most of the current still returns through the reference plane(s). To compound this, as mentioned earlier, the local εr between an edge-coupled pair is lower than the εr between the conductors and plane(s), further reducing the coupling effect. One of the real benefits of diff-pairs in PCB layout is that at last we have some signals that we are encouraged to squash up as closely together as possible. If a pair has to be separated for a portion of its length then this is not a problem, we can simply route the signals of equal length, and with each conductor maintaining Zoo; it will still function well as a diff-pair. This is why when design constraints start asking for a pair to be 'loosely coupled' and set wide apart then it's time to rethink our objectives – is it really a differential pair?

Case Study: The Heavily Constrained Design.

Perhaps the toughest board requirement that we often see comes from the following type of requirements and constraints:

- Thickness constraint of 1.6mm (62 mils)
- 4 plane layers / 8 signal layers required
- 65 Ohm single ended / 100 Ohm differential
- Cost effective FR-4 materials with 1oz (1.22mils / 35um) copper throughout

If we could lose just one of these items then we could easily lay out the board with all good design practices in place. However, twice as many signal layers as plane layers dictates adjascent signal layers in the stack. A 12 layer board 1.6mm thick implies ~5mil dielectrics. Cost-effective standard FR4 materials imply that in order to achieve the high impedances required, traces will have to be the minimum width possible. This arrangement can work well if we use orthogonal routing on adjacent signal layers, one routed as north-south and the other east-west. However, where this falls down is when we need every signal layer to head in the same direction, most commonly from dense edge-connectors to ICs. The resulting cross section that drops out from this is something like that shown in Figure 9.

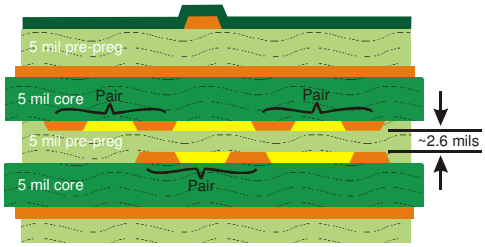

Figure 9. Cross section for the brave! Crosstalk simulation is the only safe way forward.

In the lamination process the narrow copper traces press into the semi-cured pre-preg (a fact taken into consideration in Z calculation). For 1oz copper built on 5mil dielectrics the traces could be as close a 2.6mils from one another. When this structure is viewed as a cross-section it becomes apparent that the resulting '*diff-pairs*' are a world away from the ideal case used to calculate the impedance. In fact, it's not even clear which traces are paired, or whether the intention was broadside or edge coupling!

There is a way forward from even this scenario, in the form of crosstalk simulation. On these occasions when we are forced to go against our rules of thumb then crosstalk simulation will allow us to quantitatively analyse the damage. During analysis, Allegro PCB SI will build multi-trace models and will then stimulate all valid aggressor nets whilst measuring coupled noise and impedance shift on the surrounding victim traces.

Conclusions

As rise times continue to reduce, it is a certainty that the number of traces requiring impedance control will continue to increase. Where impedance control is needed it is important to control it accurately, calculating it with the most representative cross-section you can create.

Board design, like all engineering challenges, is all about compromise. It is very easy to get too focused on one aspect of a design, and by attempting to perfect one detail, we can easily harm another aspect of the design. Controlling impedance is no exception to this. Before spending days in front of the field solver and specifying trace widths to several decimal places, ask yourself:

- Is this manufacturable?
- Will other signals nearby affect this figure?
- What's the tolerance on the dimensions (especially in the Z-axis)?
- At what frequency was the εr measured?
- What does my PCB fabricator say about my calculations?

Definitions (Polar conventions)…

Zo	Impedance of a single-ended line.
Zoo	Impedance of one of a pair of lines being driven by equal and opposite polarity signals.
Zoe	Impedance of one of a pair of lines being driven by equal signals.
Zdiff	Impedance between a pair of lines being driven by equal and opposite polarity signals.
Zcm	Impedance between a pair of lines being driven by equal signals.

If you would like more information on PCB design or signal integrity analysis, please contact:

Advanced Layout Solutions Ltd
Fronds Park, Frouds Lane, Aldermaston,
Reading, Berkshire, RG7 4LH, UK
Tel: +44 (0) 118 971 1930
Fax: +44 (0) 118 971 1931
http://www.alspcb.com
mailto:info@alspcb.com