Practical Signal Integrity

by Barry Olney | In-Circuit Design Pty Ltd |Australia

"There are two types of designers: Those that have signal integrity problems – and those that will." - *Sun Microsystems*. So if you are a digital designer – you will eventually have SI problems whether you like it or not. But all is not lost. If you learn to work with these issues, then you will soon become proficient with high-speed design.

Advances in semiconductor lithography enables IC manufacturers to ship smaller and smaller dies. However, Moore's law (1965) is still in effect–the number of transistors on ICs doubles every two years and will continue for at least ten years. Arguably the predictions about the law were short-sighted, and the paradigm will continue to apply as chip sizes continue to scale down. But keeping up with it is becoming more challenging. Intel for instance, changed transistor structure into 3D form, by placing transistors on top of each other, on the latest 22nm process to enable them to continue shrinking silicon.

Each new generation of semiconductor process technology delivers greater levels of integration, higher performance and lower cost. However, these benefits are offset by increases in power consumption that seem to unavoidably accompany each reduction in feature size. In order to reduce power consumption, IC manufacturers have moved to lower core voltages and higher operating frequencies which of course mean high current requirements and faster edge rates.

Faster edge rates mean reflections and signal quality problems. So even when the package hasn't changed and your clock speed hasn't changed a problem may exist for legacy designs. The enhancements in driver edge rates have a significant impact on signal quality, crosstalk, timing and EMI. So, whether you like it or not – **Welcome to the domain of high-speed design**.

Impedance Control

Impedance is the key factor that controls the stability of a design – it is the core issue of the Signal Integrity methodology. At low frequencies, a PCB trace is almost an ideal circuit with little resistance, and without capacitance or inductance. Current follows the path of least resistance. But at high frequencies, alternating current circuit characteristics dominate causing impedances, inductances and capacitances to become prevalent. Current then follows the path of least inductance. The impedance of an ideal lossless transmission line is related to the capacitance and inductance:

$$Zo = \sqrt{L/C}$$

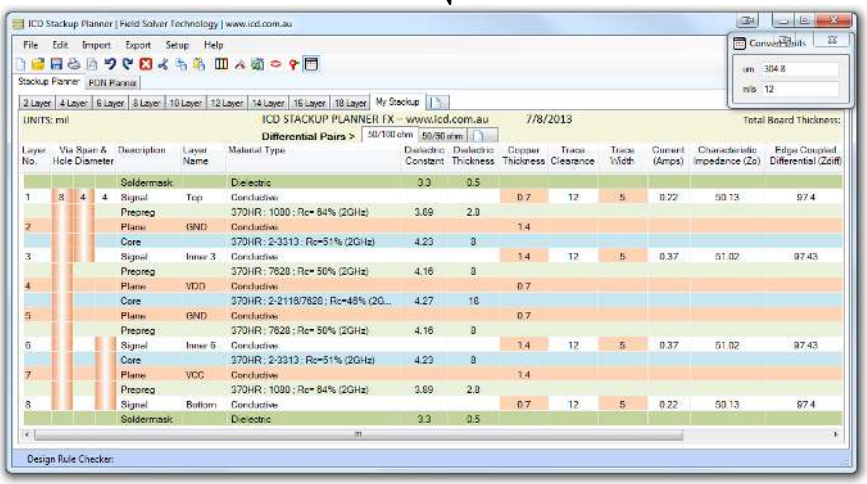

Figure 1 – Impedance of the transmission lines simulated by a BEM field solver

But this is very simplistic and the impedance should be simulated by a field solver (as in Figure 1) to obtain accurate values of impedance for each signal layer of the substrate. The impedance of the trace is extremely important, as any mismatch along the transmission path will result in a reduction in quality of the signal and possibly radiation of noise. For perfect transfer of energy, the impedance at the source must equal the impedance at the load. However, this is not usually the case and terminations are generally required at fast edge rates to limit ringing.

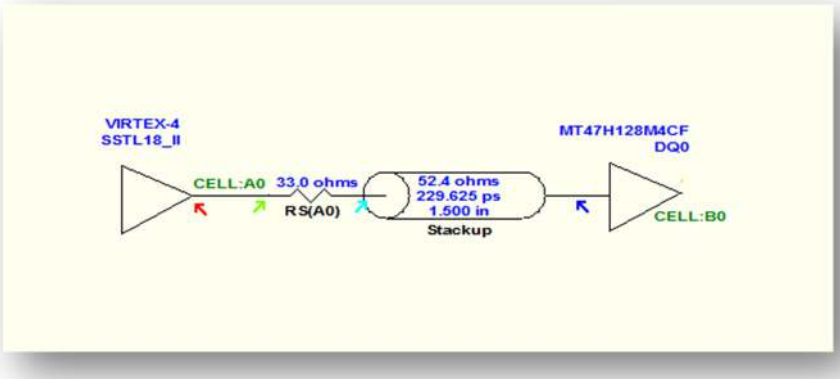

Figure 2 – Model of a Virtex-4 driver, a 1.5" transmission line and a DDR2 Receiver

Figure 2 shows a typical scenario where there is a Virtex-4 driving into a 1.5" transmission line and then to the DDR2 receiver. The impedance of the driver is 20 ohms – its value is embedded in the devices IBIS model. Now obviously a 20 ohm driver does not match a 52.4 ohm transmission line.

The red waveform in Figure 3 shows the ringing of the un-terminated trace. However, once a 33 ohm series terminator is placed close to the driver, the impedances are matched resulting in the blue waveform. Notice that the rise time is slower now. There are, of course, different types of termination strategies but series is the best for point to point terminations because it slows down the edge rate without drawing extra current. Parallel terminations are typically used as end terminators on the address busses pulling the signal up to VTT.

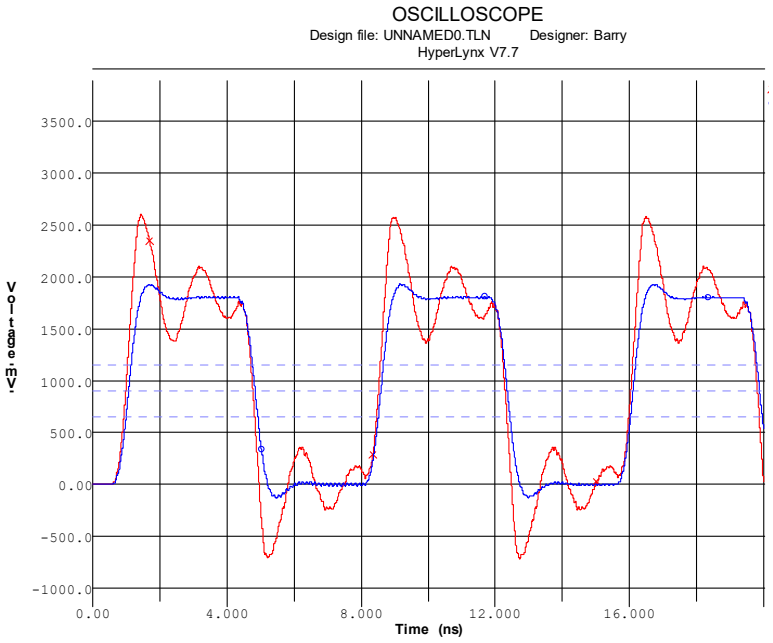

Figure 3 – Comparison of an un-terminated (red) and a terminated (blue) trace

In this case, the rise time of the un-terminated signal is 235ps whereas it is slowed down to just 350ps once terminated. It is the rise time rather than the frequency that is of concern.

Rule of Thumb: All drivers whose trace length (in inches) is equal to or greater than the rise time (in ns) must have provision for termination.

Crosstalk

Crosstalk is the unintentional electromagnetic coupling between traces on a PCB. But crosstalk can also be induced in the return path–which often gets overlooked. The *insidious little creature* pictured in Figure 4, is the crosstalk associated with two parallel trace segments on the outer (microstrip) layer of a PCB.

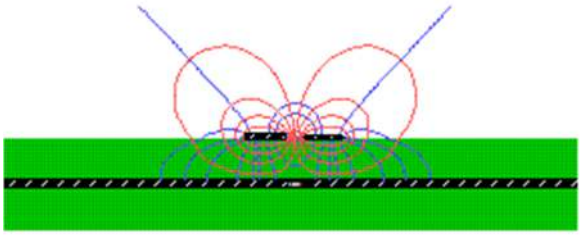

Figure 4 – Crosstalk on the outer (microstrip) layer

The red lines represent the magnetic field that couples voltage inductively to the nearby trace and also radiates electromagnetic emissions. The blue lines are electric fields that capacitively couple current into the nearby trace and are somewhat absorbed by the plane but still tend to radiate noise outward.

Crosstalk is caused by capacitive and inductive coupling:

- Capacitive coupling causes signal voltages to couple current into nearby nets. This is also referred to as forward or far end crosstalk (FEXT)
- Inductive coupling causes signal currents to couple voltage into nearby nets. This is also referred to as backward or near end crosstalk (NEXT)

In the case of inductive coupling, return currents can overlap also causing *Ground Bounce* to occur. The return currents follow the path of least inductance with the streams of electrons crossing over each other. Although this is probably immeasurable, one would assume that there would be some sort of detrimental interaction (crosstalk) between these streams. When the stackup is planned, be aware of which plane(s)–either power or ground–will be the return path for the critical signals and make sure there is an unobstructed return path.

Crosstalk can be coupled trace-to-trace, on the same layer, or can be broadside coupled by traces on adjacent layers. The coupling is three dimensional. Broadside coupling is difficult to spot as generally we look for trace clearances when evaluating crosstalk but a simulator will pick this up. Traces routed in parallel and broadside cause greater amounts of crosstalk than those routed side by side. This is due to the width of the trace being much larger that the thickness – so more coupling occurs in the broadside configuration. It is therefore good practice to route adjacent signal layers, in the stackup, orthogonally to each other to minimize the coupling region. A better solution is to only have one signal layer between two planes to totally avoid broadside coupling altogether.

Since crosstalk is induced by one or more aggressors onto a victim trace, it is obvious that the higher the aggressor voltage the more crosstalk will be induced. It is therefore best to segregate groups of

nets according to their signal amplitude. This strategy prevents larger voltage nets (3.3V) from affecting smaller voltage nets (1.5V).

Crosstalk is defined by:

$$Xtalk = \frac{1}{1 + (\frac{D}{H})^2}$$

The above equation clearly shows that in order to reduce crosstalk, we need to minimize H (height above the plane) and maximize D (distance between traces). The easiest way to reduce crosstalk, from a nearby aggressor signal, is of course by increasing the spacing between the signals in question. Crosstalk falls off very rapidly with distance. Crosstalk plummets roughly quadratically with increased separation. Doubling the spacing cuts the crosstalk to roughly a quarter of its original level.

Rule of Thumb: Gap = 3 x trace width.

However in today's complex, dense designs, it is not always possible to use up valuable real estate to satisfy the above. An alternative is to set up parallel segment rules to prevent traces running in parallel for more than 500mils. Also, the effect of dielectric height above a reference plane on trace-to-trace coupling plays an important role in reducing the crosstalk. A 3mil thickness dielectric material reduces the crosstalk by approximately a quarter compared to the 6mil given the same trace spacing.

Rule of Thumb: Couple the signal traces closely to the plane.

Rail Collapse in the Power Distribution Network

Now that the quality of signal paths is sorted out, we need to look at other sources of noise. In order to reduce power consumption, IC manufacturers have moved to lower core voltages and higher operating frequencies. But as the voltage drops – the current increases and as the frequency increases – more energy is required more often. And then, if a 64 bit wide bus is switching simultaneously, then one hell of a lot of current is required–up to 10A–instantaneously. This tends to collapse the power supply rails creating further SI and also EMI problems. The Power Distribution Network (PDN) must be designed to maintain constant voltage levels under maximum switching load.

The goal of PDN analysis is to maintain low AC impedance, on the supply voltage planes, from DC to the maximum required frequency. Excessive electromagnetic radiation typically occurs where there is a peak in the AC impedance. This maximum frequency or bandwidth should also take into consideration the odd harmonics of the clock. The bandwidth increases as the rise time gets faster. For DDR2 memory running at 400MHz, the fifth harmonic is 2GHz. So what was assumed to be just a 400MHz design now has to be stable up to 2GHz. On second thought maybe you should leave this to the experts? But it is actually quite simple if you have the right tools.

In Figure 5, the plane data has been extracted from the ICD Stackup Planner into the ICD PDN Planner, and the AC impedance of a 1.8V DDR2 supply is analyzed up to 10GHz. The AC impedance is below the target impedance of 60mΩ up to 400MHz–the fundamental frequency–and plane resonance is clear at the 5th harmonic of 2GHz.

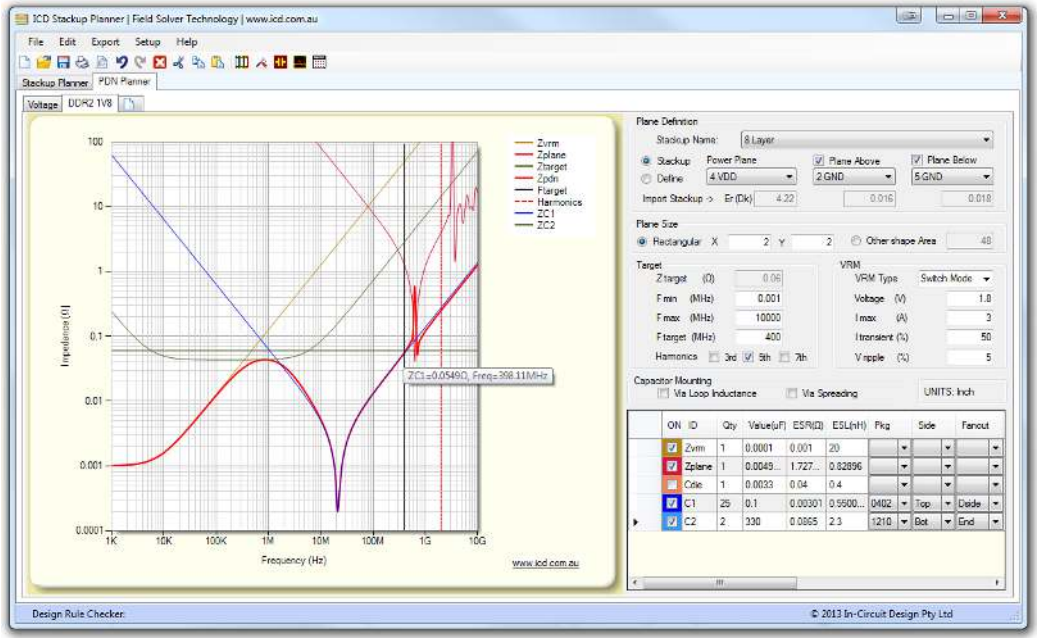

Figure 5 – The 1.8V Power Distribution Network of a DDR2 memory

Timing and Skew

Skew refers to the time difference between any two single-ended signals. For instance, if a clock signal arrives at the receiver before an address signal then the skew may cause false triggering. Clock signals should always have the longest delay, of the group, as the data, address, control and command signals need to settle down before the clock arrives at the chip sampling the bus.

Skew can also be the time difference of the two signals in a differential pair. Any mismatch in delay will result in changing part of the differential signal power into common-mode power. While skew is a timing problem, it is often caused by mismatched interconnects.

Ideally, the delay should be confirmed by simulation rather than by matched length. The length of the signal is not necessarily directly proportional to the delay. For instance for serpentine traces, forward crosstalk can increase the signal propagation even if the lengths are matched.

There are many other factors that can influence Signal Integrity but basically the Stackup Planning and the PDN Analysis, of a PCB, are the two main factors that control the stability of a design. Getting these two factors right, helps ensure the long term reliability and performance of any high-speed digital design.

Points to remember:

- "There are two types of designers: Those that have signal integrity problems – and those that will." - Sun Microsystems.
- In order to reduce power consumption, IC manufacturers have moved to lower core voltages and higher operating frequencies which of course mean high current requirements and faster edge rates. Faster edge rates mean reflections and signal quality problems.
- The impedance of the trace is extremely important, as any mismatch along the transmission path will result in a reduction in quality of the signal and possibly radiation of noise.

- Rule of Thumb: All drivers whose trace length (in inches) is equal to or greater than the rise time (in ns) must have provision for termination. It is the rise time rather than the frequency that is of concern.
- Crosstalk can be coupled trace-to-trace on the same layer or can be broadside coupled by traces on adjacent layers. The coupling is three dimensional.
- It is best to segregate groups of nets according to their signal amplitude to reduce crosstalk.
- Rule of Thumb: Gap = 3 x trace width. Crosstalk plummets roughly quadratically with increased separation.
- Set up parallel segment rules to prevent traces running in parallel for more than 500mils.
- Rule of Thumb: Couple the signal traces closely to the plane. A 3mil thickness dielectric material reduces the crosstalk by approximately a quarter compared to the 6mil given the same trace spacing.
- The goal of PDN analysis is to maintain low AC impedance, on the supply voltage planes, from DC to the maximum required frequency – including harmonics.
- Clock signals should always have the longest delay, of the group, as the data, address, control and command signals need to settle down before the clock arrives at the chip sampling the bus.

References:

Advanced Design for SMT – Barry Olney
Intel: Keeping up with Moore's Law is becoming a challenge - pcworld.com May 8, 2013
Intro to Board-Level Simulation and the PCB Design Process – Barry Olney
Board Level Simulation and the Design Process: Plan B: Post Layout Simulation – Barry Olney
Beyond Design: The Dumping Ground – Barry Olney
Beyond Design: Controlling the Beast – Barry Olney
Beyond Design: Power Distribution Network Planning – Barry Olney
Beyond Design: Skewed Again – Barry Olney
The ICD Stackup and PDN Planner can be downloaded from www.icd.com.au

Bio - Barry Olney is Managing Director of In-Circuit Design Pty Ltd (ICD), Australia. The company developed the ICD Stackup Planner and ICD PDN Planner software, is a PCB Design Service Bureau and specializes in board level simulation.