

Controlling the Beast

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

This Application Note discusses crosstalk - how to minimize and control it.

This beastie is a fact of life in digital systems. It can't be eradicated completely, so we need to minimize and control it and learn to live with it. Like flees on the dog.




Figure 1.

This insidious little creature is the dreaded *Microstripum crosstalkus radiarta* more commonly known as microstrip crosstalk radiation.

It thrives on the outer layers of printed circuit boards, is particularly prevalent where there are high speed traces in close proximity, it prefers to breed on extended piles of dielectric, it feeds on electromagnetic signals, regurgitating unpalatable spectrums, and is particularly difficult to eradicate prior to EMC testing. This beast is best suppressed at its source to stop it [crosstalk] multiplying in plague like proportions.

Laying the beast to rest between copper planes suppresses its activity somewhat but unfortunately it is still able to infect nearby signals with its relentless radiation. This beastie is a fact of life in digital systems. It can't be eradicated completely, so we need to minimize and control it and learn to live with it. Like flees on the dog.

First of all – what is crosstalk?

Crosstalk is the unintentional electromagnetic coupling between traces on a PCB.

In Figure 1. above, the red lines represent the magnetic field which couples voltage inductively to the nearby trace and also radiates. The blue lines are electric fields which 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)

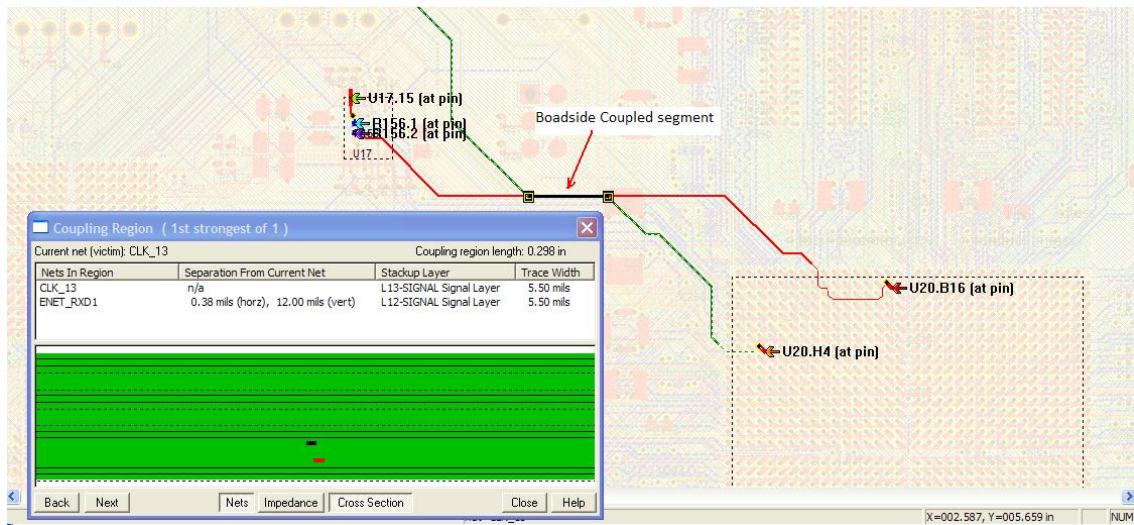


Figure 2. Boadside coupled crosstalk

Crosstalk can be coupled trace-to-trace on the same layer or can be boadside coupled by traces on adjacent layers. The coupling is three dimensional. Boadside 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 boadside cause greater amounts of crosstalk than those routed side by side. It is therefore good practice to route adjacent signal layers, in the stackup, orthogonally to each other to minimize the coupling region (not like in Figure 2). A better solution is to only have one signal layer between two planes to totally avoid boadside coupling altogether.

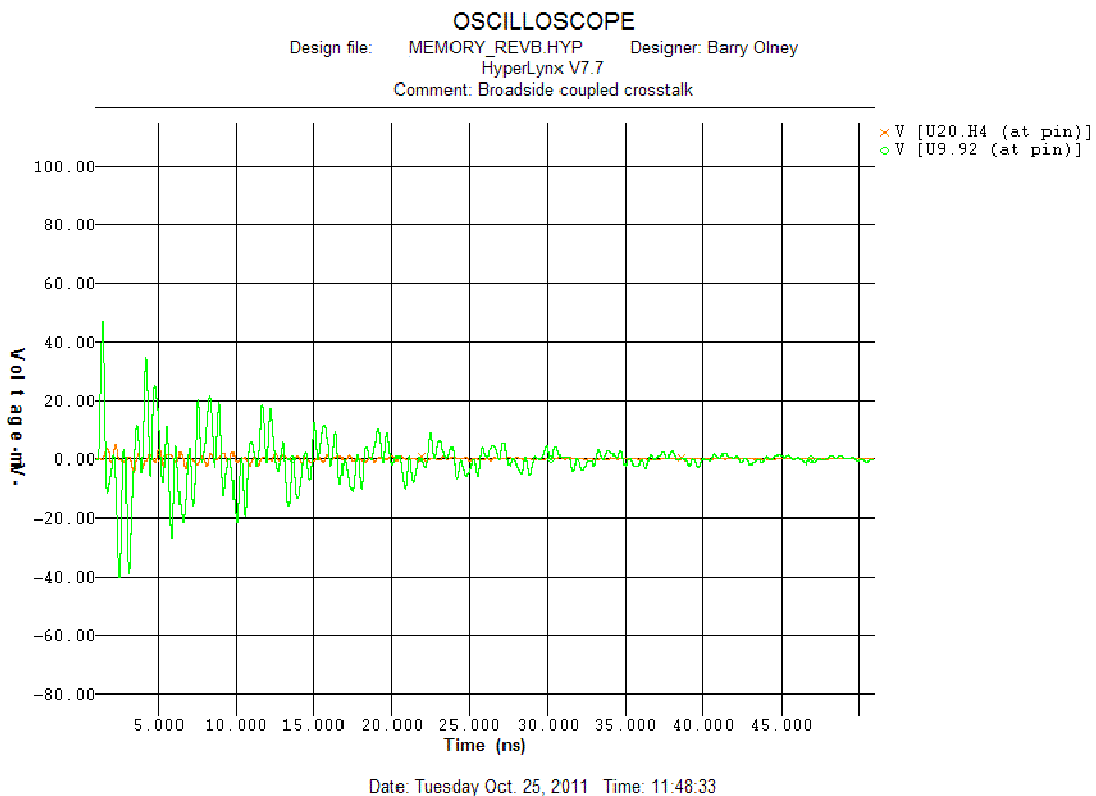


Figure 3. Boadside crosstalk on ENET_RXD1 (struck low) from CLK_13

In order to measure the crosstalk on the victim trace, we hold the victim signal low (zero volts) and pulse the aggressor.

Since crosstalk is induced by the aggressor onto the 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).

The ratio of crosstalk that is applied to a victim signal by an aggressor signal is:

$$Xtalk = 20 \log * \frac{V_{victim}}{V_{aggressor}} \text{ dB}$$

Or, looking at crosstalk from a PCB Designers point of view:

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

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.

A good rule of thumb for this is Gap = 3 x trace width. However, in today's complex designs it is not always possible to use up valuable real estate to satisfy the above.

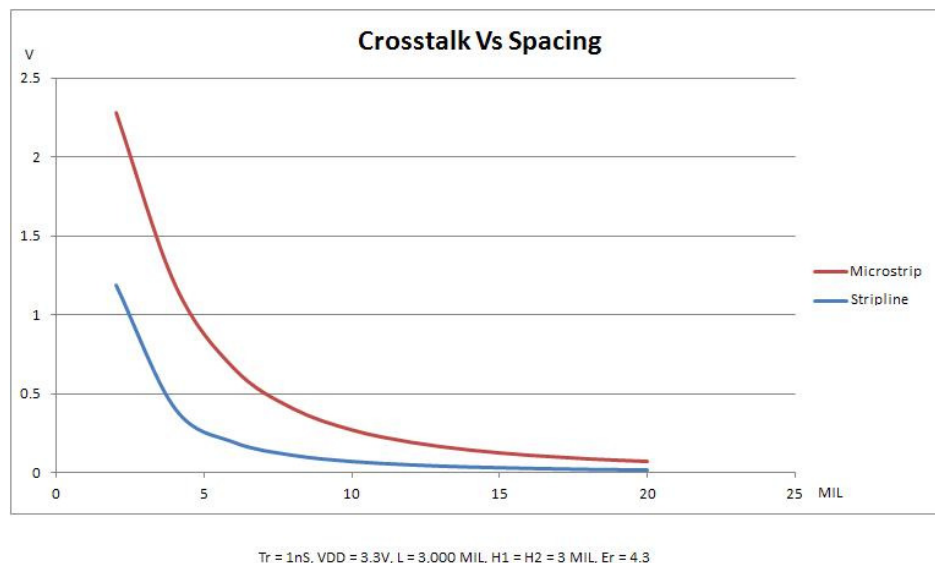


Figure 4. Crosstalk Vs Trace Spacing (edge to edge)

Figure 4 shows the effect of the tracing spacing (edge to edge) on the crosstalk for both microstrip (outer layers) and stripline (inner layers). Note that the stripline, because it is embedded between

the planes, has much less crosstalk than microstrip that radiates from the outer layers of a multilayer PCB.

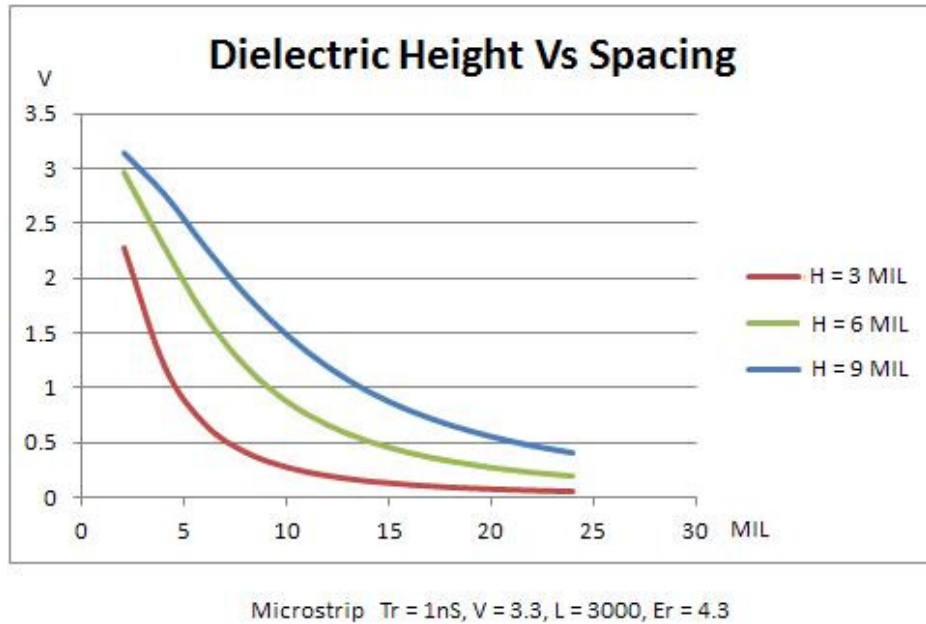


Figure 5. Affect of dielectric height on crosstalk

The effect of height over a reference plane on trace-to-trace coupling can be seen in Figure 5. The thickness of the dielectric material plays an important role in reducing the crosstalk. The 3 MIL thickness dielectric material reduces the crosstalk by approximately a quarter compared to the 6 MIL given the same trace spacing.

ICD STACKUP PLANNER – www.icd.com.au 25/10/2011 Total Board Thickness: 60.2													
Differential Pairs > 100 ohm													
Layer	Material	Dielectric	Copper	Trace	Current	Impedance	Edge Coupled	Broadside Coupled					Description
Number	Name	Type	Constant	Thickness	Thickness	Clearance	Width	(Amps)	Characteristic(Zo)	Differential(Zdiff)	Differential(Zdbs)		
1	Top	Dielectric	3.3	0.5	0.7	8	4	0.19	55.93	104.48			Soldermask
		Conductive											Signal
		Dielectric	4.3	3									Prepreg
2	GND	Conductive			1.4								Plane
		Dielectric	4.3	47									Core
3	VCC	Conductive			1.4								Plane
		Dielectric	4.3	6									Prepreg
4	Bottom	Conductive			0.7	8	8	0.31	57.72	97.03			Signal
		Dielectric	3.3	0.5									Soldermask

Figure 6. Affect of reduced dielectric thickness on impedance

However, if we reduce the dielectric thickness from 6 to 3 MIL (as in Figure 6) then we must also reduce the trace width by a corresponding amount from 8 to 4 MIL to maintain approximately same impedance. Also, assuming that the overall trace spacing remains fixed, the lowest crosstalk is obtained using the narrowest traces with the highest impedance.

When evaluating a stackup it is best to talk to your fabrication shop and make sure that the desired materials are available and that they can produce the trace widths/clearances selected without affecting yields.

With all of the above issues to take into account, how do we ever get high speed transmission lines to work effectively? Luckily synchronous buses, as typically used in DDR designs, benefit from an extraordinary immunity to crosstalk. Crosstalk only occurs when the signals are being switched and this crosstalk only has an affect within a small window around the moment of the clocking. So,

providing the receiver waits sufficiently long enough for the crosstalk to settle, before sampling the bus, the crosstalk has no affect on the signal quality at the receiver.

What is an acceptable level of crosstalk? That depends on the technology being used and has changed quite dramatically over the years going from TTL logic devices to today's high speed Gb/s devices. The amount of power a CPU uses, and thus the amount of heat it dissipates, is the product of the voltage and the current it draws. The trend is towards lower core voltages, which conserves power. But, reducing the core voltage also reduces maximum operating frequency and the level of acceptable crosstalk.

In a 3.3V system a driver going low might be expected to produce a signal no greater than 300mV low threshold (VOL), while the receiver is guaranteed to respond to any signal less than 800mV (VIL). The difference (the noise margin) between VOL (driver) and VIL (receiver) in the low state is 500mV.

For DDR3 memory devices for instance, the following values are taken from the Jedec Specification JESD79-3E:




Figure 7. Noise margin for DDR3 devices

The maximum crosstalk value is the difference between the expected voltage at the receiver and the receiver threshold. In this case the maximum crosstalk is 350mV. This is for single ended signals. Differential technologies do not have the noise margin concerns of single ended technologies. This is due to common mode rejection, which is the ability of the input to reject noise that appears coincident on both inputs. Although differential technologies are much better at rejecting input noise, they are not immune. Excessive noise is still an issue and can cause serious problems. Also, the crosstalk varies depending on the load which may vary considerably when driving banks of memory modules.

Keep in mind that the total crosstalk on each victim trace is the total crosstalk from each of several nearby aggressors, all of which must sum to maximum value. Setting the simulator to 150mV maximum crosstalk, may be low considering the above but it makes sure we pickup any coupling that may be detrimental to signal integrity.

Crosstalk creates noise that erodes the noise margin. This noise may not be so great that it alone will cause a bit failure, but it can be enough to push the total noise over the edge.

If we look into it further, the degree of crosstalk is also dependant of several other factors including driver strength (which can normally be adjusted by firmware), transmission line length, how far the segments run closely in parallel and signal rise time. In the case of long line lengths, a series

terminator slows the signal rise time and extinguishes reverse-coupled crosstalk at the near end improving crosstalk considerably.

Summary:

In order to minimize and control 'the beast' we need to:

1. Route traces with as much spacing as possible – 3 x trace width if possible.
2. Keep the signal to reference plane height as thin as possible.
3. Avoid long parallel segments > 500 MIL.
4. Avoid inadvertent broadside coupling on adjacent dual striplines layers.
5. Route adjacent dual striplines orthogonally to minimize coupled regions.
6. The lowest crosstalk is obtained using the narrowest traces with the highest impedance.
7. Use slow rise time signals or use a series terminator to slow the rise time.
8. Reduce the driver fanout – number of loads.
9. Reduce the driver strength – mid range is generally fine but should be checked by simulation.
10. Route on stripline (inner layers) rather than microstrip (outer layers). This also reduces EMI.
11. Segregate different technologies according to amplitude.
12. And finally, use all of the above techniques to ensure that you don't exceed the maximum crosstalk – 150 mV is a good rule of thumb.

References:

Advance Design for SMT – Barry Olney
PCB Design Techniques for DDR, DDR2 & DDR3 – Barry Olney
Design for EMC – Barry Olney
Embedded Signal Routing – Barry Olney
Controlling Emissions and Improving EMC – Barry Olney
Differential Pair Routing – Barry Olney
Defeat ground bounce, far end and near end crosstalk – Eric Bogatin
High Speed Signal Propagation – Howard Johnson
Jedec DDR3 Specification JESD79-3E
The ICD Stackup Planner can be downloaded from www.icd.com.au

All consideration has been taken to ensure that this Application Note is accurate, based on the information and data available. The liability of In-Circuit Design Pty Ltd is limited to correcting any unforeseen errors and revising the Application Note to meet the specified requirements. In no event shall In-Circuit Design Pty Ltd be liable for indirect, special, incidental, punitive or consequential damages including but not limited to whether occasioned by the act, breach, omission, default or negligence of In-Circuit Design Pty Ltd, its employees, contractors and subcontractors, and shall include without limitation, loss of business, revenue or profits, loss of use or data, loss of savings or anticipated savings, loss of investments, loss of goodwill, loss of reputation or cost of capital or loss of extra administrative cost, or economic loss, whether or not foreseeable, and arising out of or in connection with this Application Note. All trademarks are registered trademarks of their respective owners. E&OE