

Barry Olney's High-Speed Simulation Primer

Feature Interview by the I-Connect007 Editorial Team

The I-Connect007 editorial team recently spoke with Barry Olney of iCD about simulation. Barry, a columnist for *Design007 Magazine*, explains why simulation tools can have such a steep learning curve, and why many design engineers are still not using simulation on complex high-speed designs.

Barry also highlights common mistakes that design engineers make using simulation tools, and he offers a variety of tips and techniques for anyone dealing with simulation challenges. Among them: Don't trust reference designs and datasheets.

Andy Shaughnessy: What are some of the biggest problems in simulation? In our surveys, engineers say that they have trouble doing simulation and analysis. What is so tough about it?

Barry Olney: I think the biggest problem, Andy, is time, and that's the same with PCB design in general. The PCB design is the last process in

the design flow, and when I get a job for a board layout, it's already behind schedule. I've never ever had a job where it was on schedule and everything was running smoothly. So, you're pushed for time during the whole process, and to add simulation on top of that, that pushes it back another week or so. Management is reluctant to do it because they think, "Well, it may work and then we can get it through just a little bit behind schedule, but if we leave it another week then it delays things even further." It seems they don't have time to do it right the first time, but they've got the time for a re-spin.

A lot of engineering managers actually schedule in a re-spin because they believe they need at least two iterations before they get a working product. So, time is the biggest factor. That's why they generally tend to skip simulation. There's also the learning curve associated with the high-end tools that requires experience—not just with the tools, but with high-speed design rules. Sourcing IBIS models is another big issue. Maybe you can't find the model, so you have to compromise. IC vendors are now supplying most IBIS models, but for FPGAs,

in particular, if you get the default IBIS model from the vendor's website, it has a default pin assignment, but once the EE places and routes the actual FPGA chip, you need to redefine the pin assignments for each signal.

Now, the pin assignment of the FPGA that someone designs isn't the same as the one on the IBIS model, and that's where it all goes haywire. You think, "That simulation is simple. You just have to import the IBIS models into the transmission line model and click Go." That should happen, but, it doesn't. Where you don't have pin assignments matching, you have to manually select the required FPGA sub-models from the thousands of pins with 50 or so sub-models. And these have all got cryptic names that are different for each IC vendor. So, you actually have to find the model or driver model that matches the transmission line that you're trying to simulate. Or maybe you cannot find the model at all. Good luck sourcing a connector model. So, again, this all takes time. It took me years to figure out how to do it properly, quickly, and efficiently.

The other thing is the cost of ownership, of course. The high-end simulation tools are expensive, as you know, and are beyond the budget of the average punter. So, it's better to outsource simulation. From a cost point of view, it's not all that expensive. Alternatively, EDA vendors have professional packages that include basic simulation tools. OrCAD has Signal Explorer, that I am now using; PADS has BoardSim EXT, a cut-down version of HyperLynx; and Zuken has tools like Lightning in their CR-5000. These are all entry level simulation tools, but they still require a lot of background knowledge to get them up and running. They are not that simple to use.

Shaughnessy: Do you basically have to be an SI or PI engineer to use simulation tools? Are there any simpler tools that fall under the simulation umbrella that, say, a PCB designer could use?



Barry Olney

Olney: Simulation tools encompass many sub-tools. For instance, stackup planning and PDN planning is all part of simulation and that can certainly be done by the average PCB designer. The stackup needs to be assigned before you even start placing a chip on the board. You need to define the differential pairs and accommodate different technologies on the same layers of the PCB. You may have a 50-ohm digital signal with 100-ohm differential pair. You may have a DDR3, which has 40-ohm single ended and 80-ohm differential impedance. You may have a USB that's 90 ohms. All these have to work on one substrate, so you need to account for all these different technologies and work out the impedance for all those before you even start your design. PCB designers can certainly do that. It's not difficult. It's pretty straightforward.

PDN planning can also be done by the PCB designer. The optimization of the PDN is a trial-and-error process that needs to be done in conjunction with the stackup materials, planar capacitance, and decoupling, to fully

exploit all avenues. The idea is to lower the AC impedance of the planes.

Shaughnessy: Barry, what's the cutoff? How should a company know when it's time to either invest in a simulation tool or just start farming out simulation? Do they look at it by the speed or the rise and fall rates, edge rates, or if it's serial links? Who would you recommend?

Oiney: When it's too close for comfort, basically it's time to do signal integrity analysis. There's two ways of looking at it. There's the old, lumped-element method working in the time domain. That's static timing or the relationship between clock, data, address, and command signals. But what you must appreciate is that that static timing rides on an electromagnetic carrier wave. You have your static timing, which is critical, that rides on this wave of electromagnetic energy through the transmission line. We are used to visualizing the static waveforms on the oscilloscope in the time domain, but we also need to think in the frequency domain. It's a different world working in the frequency domain compared to the time domain. The frequency domain is particularly suited to analyzing the PDN.

This electromagnetic field that transports the static signals can also vary depending on whether it's running on a microstrip or a stripline configuration. If it's on the outer microstrip layer, because there's a mixture of solder mask and air in the dielectric, it tends to speed up the signal. It reduces the dielectric constant because the velocity of propagation is the speed of light divided by the square root of the dielectric constant. If you have a lower dielectric constant, then you have a faster speed. So, you must take into account the actual static timing of the devices, their relationship between the clock and the data and address, plus the timing of the propagating signal through the transmission lines.

Shaughnessy: When you have customers come to you, what are some of the typical jobs where they should have done a simulation and they didn't?

Oiney: Generally, memory analysis. DDR is new territory to a lot of designers. They want to feel confident in what they've done and make sure they haven't made a mistake. That's one reason why people get simulation done, and the other, of course, is they have a board that doesn't work, they've been trying to fix it and they need help. That's when they're desperate, and it shouldn't come to that; that's plan B. Simulation should be done up front, as I've mentioned before, you need to do it before you even start placing a chip on the board.

Happy Holden: Are there situations where the board might work but you fail FCC? What other external things can you fail even though you connected the point, and it seems that the board works?

Oiney: Yes, that happens a lot with reference designs. A reference design is done by smart engineers who design chips, and they know how to do a basic PCB layout, but in general they're not highly experienced PCB designers, so they don't know a lot about design for manufacturability. These new reference designs work great in the lab with a few wires hanging off here and there. Great, we've got it working. They make the reference design available, and everyone copies it thinking it's a golden board. But if you were to temperature cycle the board for the equivalent of 10 years, there are reliability issues. So that's generally what happens—intermittent problems in the field.

Shaughnessy: Do you find datasheets to be helpful?

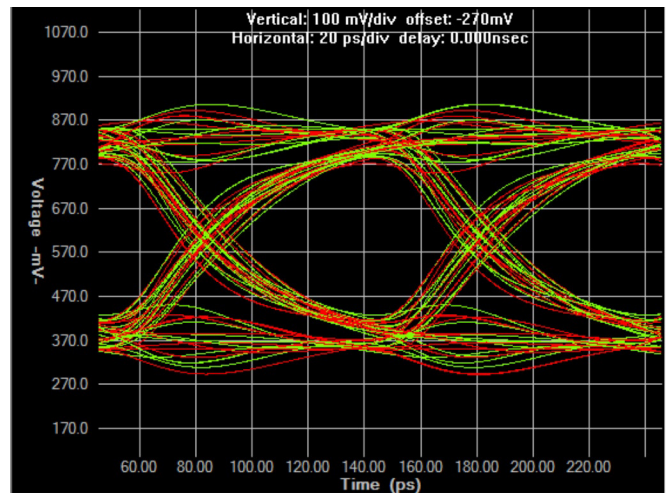
Oiney: I always say to assume the datasheets are wrong, especially libraries of physical components because you look at a mechanical draw-

ing and there's always one number on there that just doesn't fit, and you've got to work it all out. You can generally go back to the actual physical component and measure it. There's always something wrong with the datasheet.

Shaughnessy: I remember the big three EDA companies were talking about having simulation tools that would run concurrent with the design, like they would be running during schematic and all through. Is that how it is now or is it something that is right at the very end?

Olney: Ideally you need to do that, yes. You need to simulate the schematic, which you can do. You can create freeform schematic models and then add an estimate of the lengths of the transmission lines. For example, for memory you may say, "I have two and a half inches here. So, I'll put a two-and-a-half-inch, 50Ω, transmission line, attach an IBIS model at either end and see how it goes." That's really good for building up the maximum length you have to play with because before you do the placement you need to know the maximum timing you can have on a chip. If you place the chips too close together it's difficult to route, and if they are too far apart then it may mess up your transmission lines with long delays and too many reflections.

So ideally you need to work out the placement before you put the chip on the board, and that also helps with your routing topology. You may want to do flyby topology—that's how you generally do DDR3 and DDR4—but there's no need to do a flyby if you've only got a couple of memory chips. You can still do a T-topology route, which means branching out with equal delay into each chip. All the professional EDA tools now have a way of extracting the topology, so once you have routed a particular section you can extract the topology of that from the PCB, so you get the physical transmission line. That includes whether it's running on microstrip or asymmetric stripline, for example. It builds the actual physical structure of the PCB, including layer transitions (vias), into



Much of high-speed simulation is centered on achieving a nice eye diagram.

the simulation model. Then you can virtually simulate the board itself, so that's the final process—making sure all the timing is right on the physical simulation.

Holden: Does anybody publish or sell a good dielectric material library?

Olney: I do (laughs). Actually, our stackup planner has 35,330 individual materials at different frequencies. There is a choice of over 700 series of rigid and flexible dielectric materials from over 60 manufacturers. There is every material you would ever need. If you can't find a new material, I'll add it free of charge. It's very comprehensive.

Shaughnessy: What good sources of information on simulation analysis are out there? If someone is new to the industry, what would you recommend to them as far as getting some more information?

Olney: Well, I think reading your magazine, Andy, is a good start, and other industry magazines. The SI Forum is a really good source of information with many experts on there that can answer your questions, as is attending any of the technical seminars that Rick Hartley and people like him run. They are a good start. I

started off attending conferences and technical sessions. Happy's "Design for Manufacturability" was probably one of my first back in 1990. You can learn about the engineering process, but design for manufacturability is very important, and it's not taught in universities. There are a few courses on signal integrity that are targeted at EEs but not so much PCB designers. I developed the "Advanced Design for SMT" course for the Australian Electronics Development Centre back in 1994 to teach high-speed design techniques.

Holden: I guess we should qualify that a lot of what we're talking about in terms of problems apply to analog and digital designers, but if people are involved in RF design and doing it correctly then they have learned a lot of these things early on. It's just that the percentage who are really good RF designers is so small.

Olney: RF and microwave design is a different world altogether; the circuits are very high frequency with fast rise time and short wavelengths, so you need to be very careful. Every stub is an antenna. That's why coplanar waveguides are typically used for RF, because they tend to stabilize the signal. For a microstrip signal, for instance, you've got coplanar waveguides with ground pours on either side of the signal which tends to dampen the radiation.

Shaughnessy: A lot of designers and engineers agree that simulation is a great idea, but many of them don't use it. Why don't more people use simulation?

Olney: I'm sure it's mainly to conserve time. Managers always strive to reduce their design cycle. But simulation isn't that expensive or time consuming if you outsource it. For instance, I can analyze an average complex board for about US\$5,000 and find the problem. If I were a manager and I had a choice of a small investment in analysis or taking the risk of a re-spin, I know which I'd choose. It's a no-

brainer as it verifies that the product will perform reliably!

Dan Feinberg: One of the things that you just raised and Happy just mentioned is that people often look at price, but it's really cost, and they're different.

Olney: Absolutely. The cost of a re-spin is a lot more than the price of the tools or outsourcing analysis. It's not just engineering time you are wasting, but time to market and lost opportunity which can be extremely costly.

Shaughnessy: Lee Ritchey says a lot of these conflicts stem from a managerial viewpoint versus the engineering viewpoint. The manager thinks, "Great, the design was done quickly," even though there may be a costly re-spin down the road, because for now, the EE can get on to another project.

Olney: That's right, you'll always have conflict. Engineering and management require different skill sets. The manager is responsible for getting the project completed on schedule and to budget. And, he has to look good to his boss. If he schedules a re-spin, then it is expected. It's in the timeline. Whereas the engineer or PCB designer is more interested in the completion of their work and learning new techniques to improve their skills.

Holden: Is PCB design growing in Australia?

Olney: I guess I am a little out of touch with the Australian electronics industry as I mainly do business overseas now. We have small groups of designers, but there is not a large electronics industry except for defense contractors. It was always difficult to sell VeriBest PCB and Mentor (Siemens) tools in Australia because it's the home of Altium, of course. Everyone uses Altium and going against tradition is very difficult. I guess you could say that the Altium Designer's market is growing. There isn't much

of a market for simulation tools either. In fact, in the past, anyone who did well ended up going to the U.S., so you lose a lot of talent when a company becomes successful; it just disappears offshore.

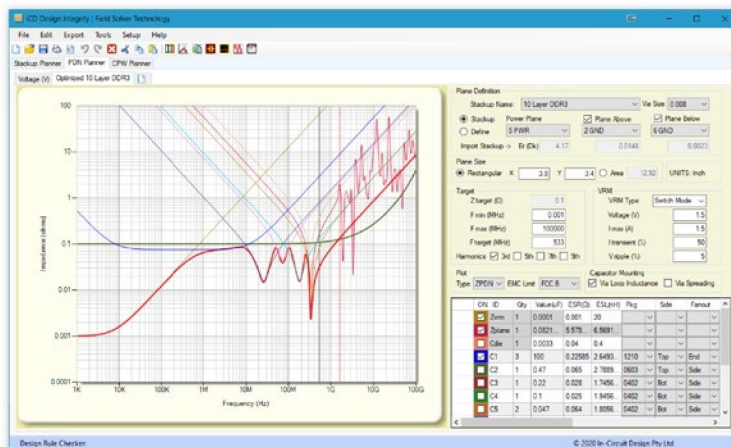
Holden: Where do your customers come from?

Olney: Lately, the U.S. Probably 70% of my customer base is in the U.S., 25% in Europe, and 5% from other countries.

Shaughnessy: Do you have any recommendations for somebody dealing with simulation challenges right now? Any tips?

Olney: Sure. The biggest thing is forgetting what you learned at university. Circuit theory is all DC, but when you get into the frequency domain, with high frequency, fast rise time signals, and a distributed system, you need to think in the way of the electromagnetic fields, coupling, and radiation. It's a different world altogether. You really need to educate yourself before you can become a good high-speed designer. The other point I'd like to make is that simulation tools don't complete the analysis toolset. My eyeball is just as good as any signal integrity tool. That's because I have (embedded in my skull) 30 or 40 years of IPC design rules, design for manufacturability, and SI and PI requirements, so I can see a lot of issues before I really need to simulate. Just looking at the topology of the routing, for instance, gives me an idea of whether a design will perform to expectations. On the other hand, simulation gives you another set of eyes and allows you to visualize other issues.

When I do an analysis, I think outside the square. I drill down through the board; I do a teardown of the substrate. First, I look at the impedance and different technologies that must run on each layer, the materials used, and at the minimum cost, of course. It must provide adequate performance, but it doesn't have to be the best material, and it also has



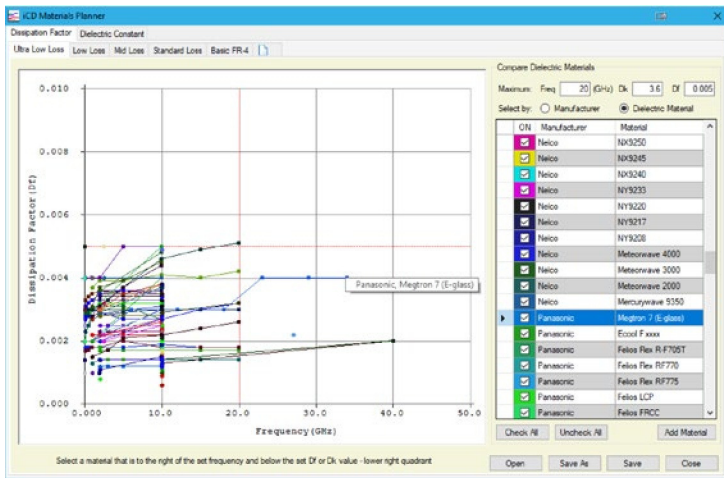
An In-Circuit Design field solver simulates a DDR3 stackup.

to be stocked by your fab shop. So, you must choose the right material for the right job.

These days, one of the largest problems is the amount of power supplies you have on a board. I think the last design I did had 35 power supplies.

Some believe, "Routing is simple. You just route between two planes and keep the critical signals isolated, etc." But until you get in amongst it (as a PCB designer, which I am) you come across the real problems of space. You just don't have the space to put in everything, while working out where these 35 copper pours go, and not interfering with anything else is a major problem in board layout. If you have a trace that crosses from one power plane to another, you are creating return path discontinuities and it can cause untold problems with EMI.

So, you really need to plan how you're going to place these power supplies to get the most effective route. If you can't route a board it may not be because you don't have enough layers. People often say to me, "How do you know when you've got enough layers? Is eight layers enough, or do I need 12 layers?" After a while you get a feel for it. But a good way to test it is just to let your autorouter go loose, and if you don't get 85% completion, then you've got something wrong. If the board is difficult to route, then it may be because of placement. You need to open channels for the router.



A materials planner tool allows the user to compare Df and Dk of various materials.

One of the things I talked about was time. A lot of people start designs when they are not ready for layout. They say, “We started the design, and we’ll give you the main processor, the memory, and the power supplies,” for instance. You start with these, move things around, and change them from one side of board to the other. Spin them and rotate them and whatever to get the best placement to prevent the rat’s nest from crossing over, then route that section. Eventually the EE finishes the schematic, and you add all the other chips. “But the processor would have been better on the other side of the board.” Turn it the other way around and all the signals would not be going through the center of the board from one side to the other. So, my recommendation is to never start a board until the schematic has been finished, approved, and it’s ready to go.

The other thing I do when I’m analyzing a design is to look at the PDN, which is very important from a stability point of view. The key points for stability in the design are stackup impedance and the AC impedance in the PDN. Adding planar capacitance by using very closely coupled power/ground planes pairs, positioned close, in the stack, to the top and bottom ICs, has a dramatic effect on reducing the AC impedance. This is where the stackup configuration needs to be adjusted in conjunction with the PDN plot.

Minimizing the reflections on high-speed signals is another thing. You may have a driver, and typically the source impedance of the driver is between about 10–30 ohms. But, you need to have a transmission line of say 50–60 ohms. Now, that won’t match to a driver of 10–30 ohms. In high-speed design you need to have either a series terminator or a parallel terminator at the end of a long line. Parallel terminations simple match the transmission line impedance of 50Ω. But a series termination must be calculated. The iCD Termination Planner, for instance, extracts the attributes required to determine the source impedance of the

driver from an IBIS models IV curves. Then the series termination resistance is calculated, based on a distributed system and load, to match the transmission line for the selected layer in the stackup. In a typical digital design, you’ve only got the rise time to worry about, but with DDR you’ve got the data being clocked from both the rising and falling edges of the signal. So, you need to also look at the falling edge of the waveform, and that’s quite a lot faster than the rising time of the device. This is due to the design of typical CMOS output drivers. For the same feature size transistor, an n transistor can turn on faster than a p transistor. So, the fall time is always faster than the rise time—which you need to consider.

Simple things like that will get your reflections under control. Now, reflections create a lot of crosstalk because if you have close coupling on the signals, which you normally do with tightly routed boards, you don’t have a lot of room. You’ve got to route things close, and so you get coupling and crosstalk, which creates radiation. EMI is created from all these reflections. It’s stabilizing the impedance and the power supply that stops these reflections and radiation allowing the product to pass EMC tests. All critical signals should be simulated. I don’t actually simulate every signal within a bus. If you’ve got a huge bus with 128 bits, you’re not going to simulate each one, but one

from each byte lane is enough because they're all routed much the same. Pick the worst-case signals to simulate then the rest will be fine.

When you're routing the data group of signals, it should all be done on the one layer of the substrate. A lot of reference designs that I have seen have address lines and data lines routed on the outer microstrip layer and some on the inner stripline layers. As I mentioned before, that all impacts timing, because of the different velocity of propagation of the outer and inner layers. If you run the entire bus on one layer, then you're matching the timing automatically, so that saves a lot of simulation time and risk.

Crosstalk is very difficult to predict. When I do a teardown of the board, I look at every layer with respect to every other layer. I'll have a look at layer one, which may be a signal layer with respect to the layer two, which may be a ground, and look for any signals outside that ground plane, because once the signal leaves the plane area that increases its impedance, and it will radiate and cause problems.

When you get to stripline, you may have dual asymmetric stripline, and you will have two traces between the planes. You need to look at the planes with reference to each signal layer and see where the return paths will flow. We tend to route signals from one point to another from a driver to a receiver, but what you must understand is the return current will flow with high frequency directly underneath that signal trace, and if there's a split in the plane, or it has to change layers in some way, then you need to facilitate that.

If you have two ground reference planes, for instance, you can place a ground stitching via where the layer transition is. However, if the reference plane changes from ground to power, traditionally we use one decap across the planes to provide a return path for the current. If you have two different power supplies—for instance, a 5V and a 3V—and you have traces going across the split, you should put in a decoupling capacitor going from the 5V to ground and then from the 3V to ground.

This stops the noise coupling between the two different power supplies that are going through ground rather than just from supply to supply.

There are a lot of things to look at about overlapping your signals. Broadside coupling is very, very difficult to spot because when we are routing we generally turn off the other layers or just dim them in the background. We are just looking at the one layer we are routing and pushing things around, and we don't understand or see what's on the other signal layer near it. In the stripline situation where you have broadside coupling, you have the width of the trace coupling to the width of the trace below, which is a lot more coupling area than you would have with an edge coupling because you only have the very thin edge of the trace with edge coupling. With broadside coupling, you have a wide trace coupling to another wide trace, and generally you've got a very thin dielectric—like a 3-mil—in between. That makes a beautiful coupling, maybe a good RF coupling, but it's not what you want to do in digital design. Also, be careful of broadside coupling on built-up microstrip layers.


Shaughnessy: It sounds like a lot of this comes down to following solid design and engineering practices, but also having the experience to know what to look for. You can't be a part-timer with simulation tools, I guess.

Olney: It's all about experience. Experience is the greatest teacher of all! That's why I say you can't just put a tool in front of an EE who's just out of college. They don't know what to do—it doesn't make sense to them. They don't understand the principles or concepts, and someone needs to give them direction in how to analyze high-speed designs quickly and efficiently.

Shaughnessy: This has been really great. Thanks for speaking with us.

Olney: Thank you. DESIGN007

We **DREAM** Impedance!



Did you know that two seemingly unrelated concepts are the foundation of a product's performance and reliability?

- Transmission line impedance and
- Power Distribution Network impedance

DISCOVER MORE

iCD software quickly and accurately analyzes impedance so you can sleep at night.

iCD Design Integrity: Intuitive software for high-speed PCB design.

"iCD Design Integrity software features a myriad of functionality specifically developed for PCB designers."

– Barry Olney

