Modern High-Speed Design Methodology Integrates EMC and Signal Integrity

RON MATTHEWS
Cadence Design Systems, Inc., Chelmsford, MA

It is now necessary to treat EMC and signal integrity concurrently.

The Challenge

As modern high-speed electronic system design evolves, the challenges which designers must deal with evolve as well. In the past, it was possible for various design functions, such as signal integrity, EMI and thermal, to be considered separately. Signal integrity, which is a feature that can be sold to customers, was given primary importance. As the system prototypes were being developed, mechanical engineers concerned with thermal issues made their contributions to the system design. Finally, when the designers were happy with the way the system worked, it was passed on to the EMI people who worked their magic to make it pass regulatory testing and get the necessary certifications.

Now, however, faster and faster clock and data speeds create a new class of challenges for system designers who are beginning to realize what members of the EMC community have maintained from the start—that EMC must be designed in from the very beginning and not added as an afterthought. It is now necessary to treat EMC and signal integrity concurrently. Add this to the fact that competitive pressures necessitate getting products to market as quickly as possible, and one can easily see that developing a successful design process—one that considers EMC and signal integrity from the beginning—is absolutely essential. This article will focus on superimposing a high-speed design methodology onto a typical schematic/layout oriented digital design process. It will further show how the high-speed process steps relate to the typical design process used by most companies today.

Why is High-Speed Design So Important Now?

There are many design processes that are aimed at achieving electrical integrity, but for the most part, these have been based on static, rather than dynamic operation. For instance, in slower systems running at 10 MHz, signals spend over 90% of their cycle duration in the static condition (Figure 1).

This being the case, parameters I_{OL} and V_{IH}, which describe this static state, are sufficient and form the basis of electrical integrity analysis processes used in the past. In contrast, signals in today's 66-MHz systems typically use about 1/3 of the cycle for switching. Often, these signals never even reach a "static" condition before they are required to switch again. Simulating and understanding these dynamic characteristics have become essential parts of determining how well a design operates.

To better illustrate how this has occurred, let's take a look at the way the design of the desktop PC has evolved over the last decade (Table 1).

Referring to the first row, we see that each of the older interfaces, and even the 386 and 486, topped out in the early '90s at 33 MHz. Interim solutions emerged in the mid-90s at higher frequencies to help performance continue to rise, but eventually gave place to the (even higher frequency) stable solutions primarily due to the problems stated.

It is even more interesting to consider how high-speed design and simulation was used to arrive at both the interim and stable solutions. The interim solutions *were* derived by "experts" on very advanced, but less user-friendly tools using a process that was still being developed. In contrast, the simulations of the stable solutions are now being performed by engineers everywhere on tools supplied by leading electronic design automation (EDA) vendors that are much easier to use. Further, the

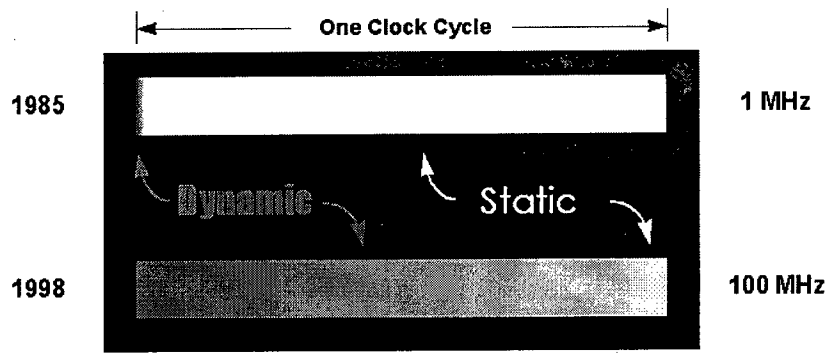

Figure 1. *Evolution of Signals.*

Subsystem ->	I/O	Processor	Cache	Graphics	Memory
Old Interface	XT - AT	CMOS	Async CMOS	ISA	FPM / EDO
frequency (MHz)	4	33	33	4	async 33
year problematic	'87 - '90	'90 - '92	'93 - '94	'90 - '93	'94 - '95
Interim Solution	EISA	CMOS w/rules	Burst CMOS	PCI	SDRAM
frequency (MHz)	8	66	66	33 / 66	66 / 100
timeframe	'90 - '92	'93 - '94	'94 - '95	'94 - '96	'95 - '97
problem	cumbersome	delicate	integration	too crowded	no future
Stable Solution	PCI	GTL	Into IC/Module	AGP	RDRAM
frequency (MHz)	33	66	100+	66 / 133	800
in place	1992	1995	1996	1997	1998

Table 1. *Evolution of the Desktop PC.*

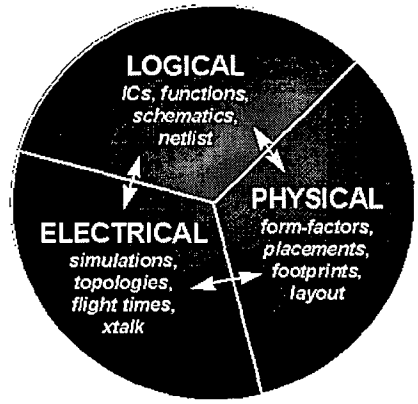

Figure 2. *Domains of High-speed Design.*

process used is becoming more crystallized into a common methodology—largely what we are attempting to illustrate in this article. Converging on an understandable and repeatable methodology will be the primary focus of high-speed design during this latter portion of the 1990s and into the next century.

Now that we've briefly reviewed how technology has migrated to higher performance, let's look at how design processes have changed in the 1990s to accommodate this migration. As high-speed design becomes more critical in today's mainstream designs, it must be more tightly integrated into the hardware development process. An ideal solution would integrate the bilateral electronic sharing of design information between each of the disciplines (Figure 2).

While yesterday's design success hinged on accurate passing of information between logical and physical domains (to/from schematic and layout), today's high-speed boards require close coupling of the electrical and physical domains (shown in the lower-most arrow).

Figure 3 presents a detailed process flow for the proper way to address high-speed issues and concurrently derive PCB topologies (shown at left). Successful high-speed engineering requires driving this process in conjunction with the larger system design process. The key challenge in making this process work effectively is using the data to drive and coordinate system-level tradeoffs and changes among the design disciplines (refer to the circle in the diagram). Usually the simplest design adjustments become obvious when viewed at the system level.

The key deliverable from this process is a set of "route topologies files" that are known to work well in order to ensure a rapid and effective process of physical layout. Deriving good topologies drastically reduces the risk of having high-speed problems arise when the design becomes physical and problems become

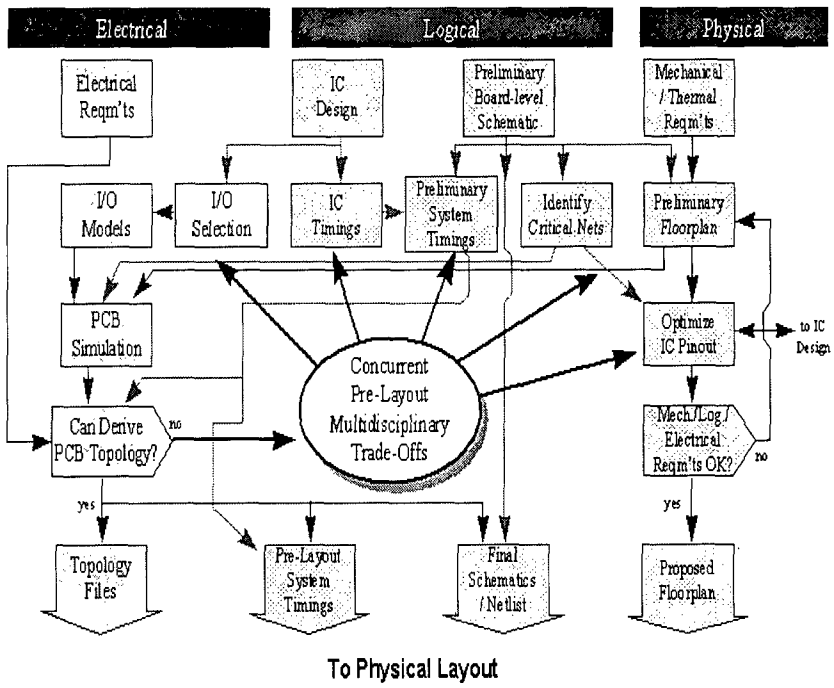

Figure 3. *Flow Chart for High-speed Issues.*

Continued on page 94

. . . *Continued from page 44*

harder to fix. Some design software can now read in these topology constraint files electronically to help guide the layout process. Automating this step helps to accelerate the design process, eliminate miscommunication, and enable electronic archiving of successful design solutions.

When done properly, the process shown correctly influences I/O buffers, IC timings, system timings, IC pinouts, final schematics, and the proposed floor plan to achieve proper high-speed operation and the mitigation of EMI. Significant production risk is reduced by coordinating these aspects during the design phase. Obviously, a methodology that waits until the implementation phase to derive topologies misses the opportunity to influence and optimize the larger system-level environment for high performance.

The proper functioning of lower-speed boards hinged only on routing the connectivity expressed in the final net list generated by the logical design. But readers should understand that logical connectivity expressed *what* was to be connected, not *how* the connection must be made (for example: in what order?; over what impedance?; and with what acceptable delay?).

In contrast, today's high-speed boards require *both* the connectivity requirements expressed in the logical net list *and also* the electrical and topological information (shown here in "topology files"). This is the primary change caused by higher speed devices: to successfully route a high-speed board, both sets of information must be fed into the physical layout process. Actually, that's been fairly obvious for awhile, but the fact that now *most of the board* has high-speed ICs and nets has made the automation of this task imperative. As such, this article focuses on exactly how this can be done.

One of the critical steps to ensure that the methodology is successful is to derive the constraints that will drive the routing of critical nets. Successful determination of the constraints requires several parameters to be varied. These simulations should determine several things about the topology. Most importantly is how the net should be routed (e.g., Star or T topology), microstrip or stripline medium and termination impedance and location.

More and more people involved in modern high-speed system design are becoming aware, at least to some degree, of the importance of sound design methodologies to ensure signal integrity. But these same methodologies must be extended to good EMC design as well. High-speed signal integrity design is often viewed as a difficult endeavor reserved for the specialist. If this is so, then design for the mitigation of EMI in high-speed systems is close to black magic. For example, in signal integrity the noise limit is often in the 50 mV range, whereas in EMI the noise limit is in the 500 mV range, a two-orders-of-magnitude difference! Further, common-mode current—one of the most vexing problems for EMC engineers—is created unintentionally and can cause a system to fail FCC class B when there is as little as 5 mA on cables. So it's easy to see how EMI mitigation is approached with trepidation and anxiety.

However, this need not be the case. EDA tools that are capable of extracting interconnect geometries from board files for use with signal integrity analysis and simulation software tools already have most of the information they require to be able to do a good job of simulating EMI, at least at the PCB level. This author takes the view that for high-speed digital systems, the PCB is really a collection of "offenders" that can cause EMI.

It must be noted here that EMI is not just one problem. Rather, EMI is a catch-phrase referring to a number of phenomena, any one of which can cause a system to fail to comply with published U.S. and international EMI limits. This being the case, it is most effective to simulate each of these phenomena individually in order to understand the contribution that each one makes (or could make) to a failing EMI profile.

Two of the phenomena which are being effectively simulated at present are differential-mode radiated emissions from PCB traces and common-mode voltage in the power/ground planes. The latter is one of the most insidious sources of radiated emissions from cables caused by switching noise voltages in the power/ground planes. Both of these sources of radiated emissions are interesting to different types of high-speed systems designers and manufacturers. Differential-mode radiated emissions from PCB traces are a source of great concern to personal computer manufacturers who sell motherboards and modules directly to the public. The FCC requires that these devices be tested in an enclosure with the cover off and two sides exposed. When so tested, the emissions from these modules may be no more than 3 dB over the FCC Class B limit. This precludes such companies from being able to "hide their sins" inside a shielded enclosure.

As with signal integrity, the correct time to address EMI issues is at the beginning of the design cycle. As drivers and interconnect structures are chosen, these choices can be scrutinized with an eye toward EMC by being able to simulate differential-mode radiated emissions even before routing begins. Further, the design constraints associated with each class of driver/receiver pairs and associated interconnect structure can be captured as a topology file which will guide the design through placement and routing to ensure that the constraints are correctly implemented. This process has been in place for some time for signal integrity, but it is now possible to address EMI-related concerns up front using this same methodology.

Figures 4, 5 and 6 show how state-of-the-art EDA tools can be

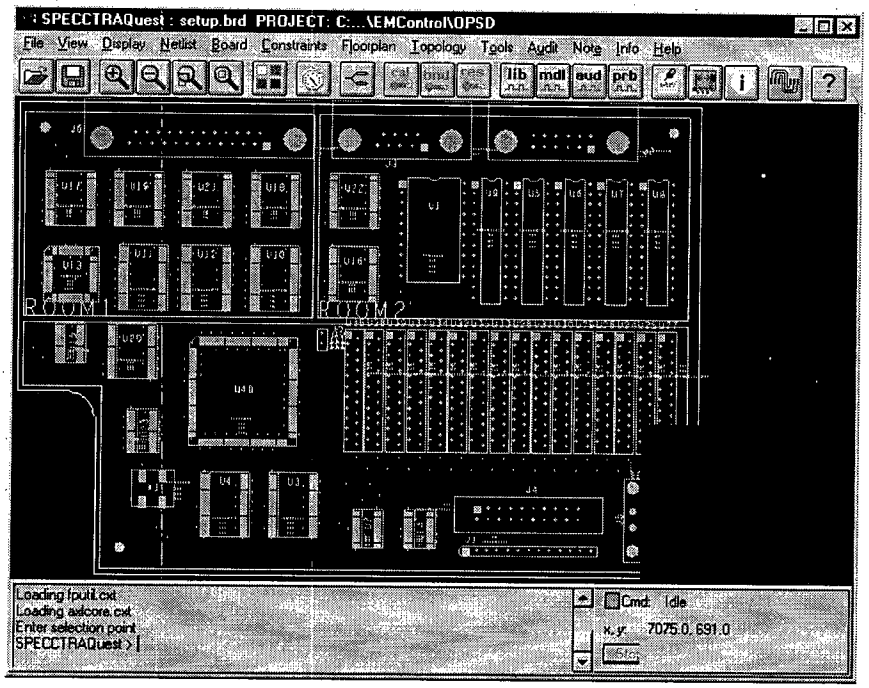

Figure 4. *Board File.*

used to extract a net from a PCB layout to a tool which displays the net and associated drivers/receivers logically. This logical display is then used to kick off signal integrity and EMI simulations. The screen shot in Figure 4 represents a graphic display of the board file. Using engineering design tools that push the envelope in state-of-the-art design practices, a net can be selected for further simulation and analysis. This user interface for this simulation environment is shown in Figure 5. In this environment, driver and receiver characteristics can be specified and what-if analyses can be carried out with respect to the interconnect. This type of work is well-developed for post-route signal integrity, but the figures shown are snaps of analyses done on a pre-route basis. The EMI field plot shown is developed in the same environment as the signal integrity tool.

This type of analysis is very useful for anyone involved in the design of high-speed systems but is especially interesting to those companies that sell motherboards and modules directly to the public.

As mentioned previously, one of the most vexing problems for the EMC engineer is common-mode radiated emissions from cables. Common-mode voltage exists due to the finite impedance in the power/ground structure in PCBs. There are two major sources of this common-mode voltage in power/ground planes in PCBs:

- High-speed power/ground switching at the device
- Signal and image return current in the power/ground planes

It is becoming imperative that design teams take problems like this into account as early in the design phase as possible. With the new generation of EDA tools, it is now possible to simulate the effects of common-mode voltage in the power/ground planes and to evaluate the effects of judiciously placing decoupling capacitors so as to mitigate the effects of these common-mode "hot spots."

By making the power and ground planes part of the circuit file, signal integrity engines can be used to depict the modal voltages in the planes. If a device like a clock or clock driver has spectral components higher in frequency than the natural resonant frequency of the power/ground planes to which it connects, like any electrically large structure, the power/ground plane pair behaves like a length of transmission line and can exhibit resonance modes. These modes are better known in the EMC community as "hot spots" in the PCB. This author has been faced with numerous EMI problems where, for example, the parallel port or serial port was causing the system to fail

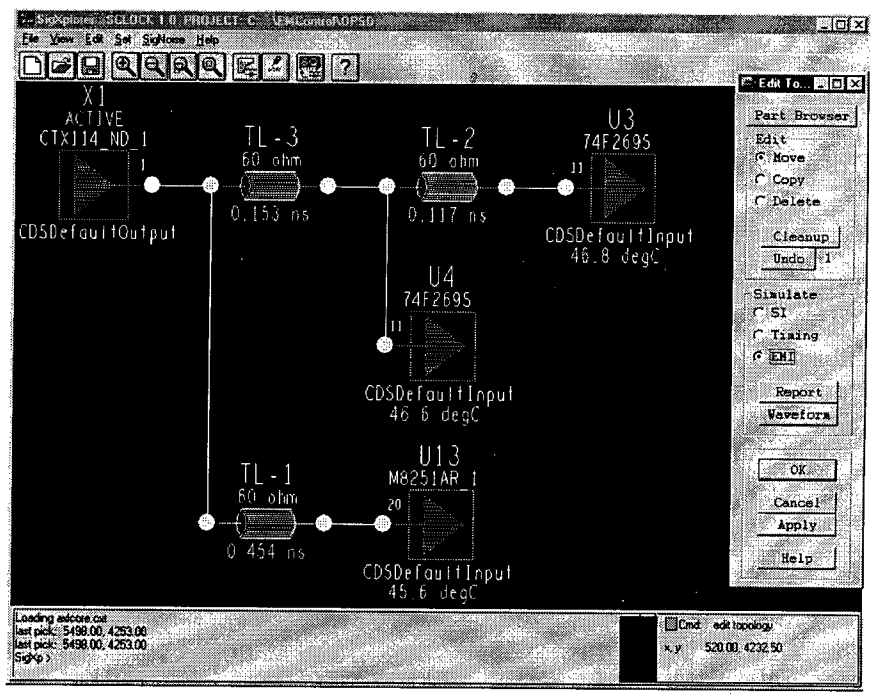

Figure 5. *Simulation Environment.*

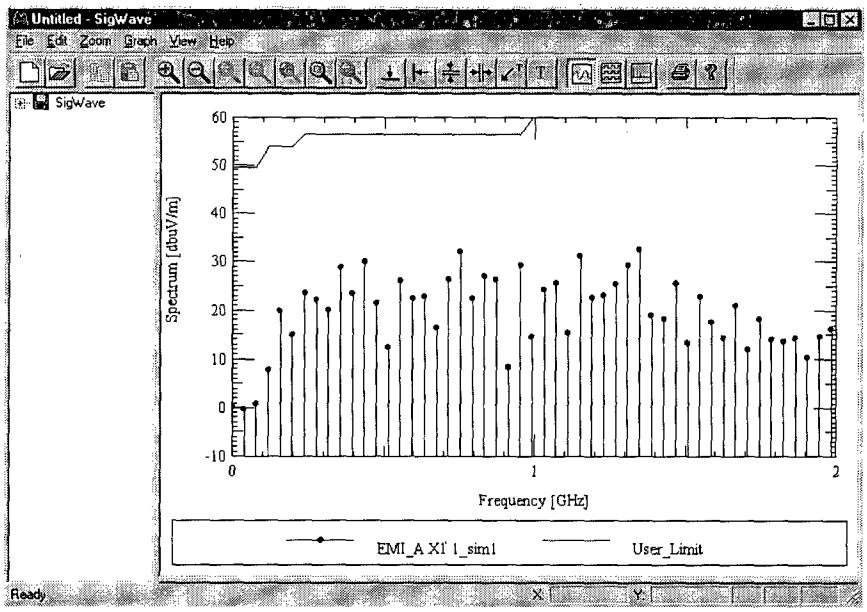

Figure 6. *EMI Field Plot.*

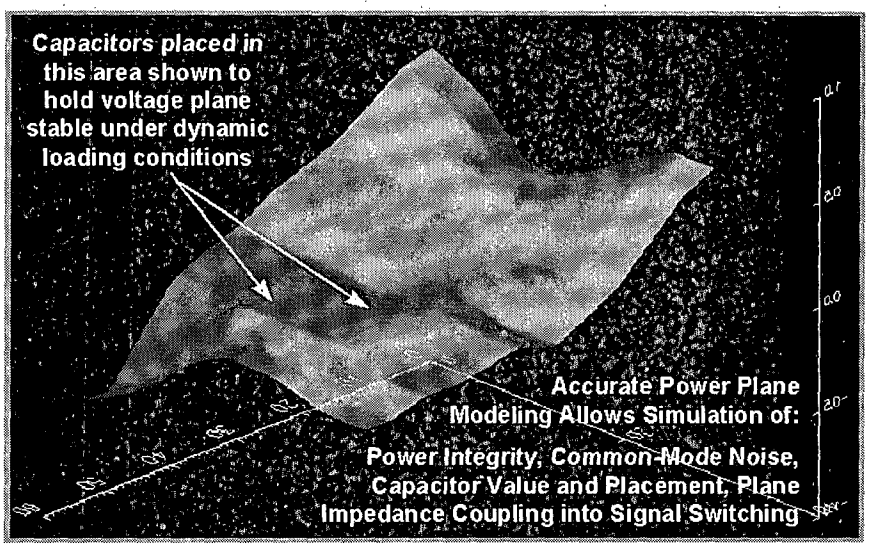

Figure 7. *Simulation of Common-mode Behavior.*

radiated emissions at a harmonic of the system clock, which was usually located as much as a foot away on the PCB. This type of problem is attributable to these common-mode hot spots caused by high-speed switching.

While it is extremely difficult to simulate this type of radiated emissions problem directly, the problem can be effectively dealt with by simulating this common-mode behavior in the power/ground planes and then attempting to use decoupling capacitors as "mode killers."

Figure 7 shows a graphic representation of this type of simulation. The user is able to see the effects of this common-mode voltage between the power and ground planes as color contours and/or as spatial fluctuations along the z-axis. Note the flat area which is where an array of capacitors was effectively employed to kill the traveling mode in that region.

What Does the Future Hold?

As we move forward we can expect to see even more progress in the area of EMI simulation both as a pre-route design and post-route checking. The ability to perform pre-route checking in batch mode will allow the user to evaluate entire buses or the traces routed in a specified functional "room" or area on the PCB, with the results being displayed in spreadsheet form. This will allow the user to tell at a glance which nets need to be managed from an EMI standpoint and which can be treated less aggressively. Incorporating such a tool early in the design allows us to treat EMI effectively without over-designing.

Post-route checking will employ full-wave techniques to give a more accurate prediction of the differential-mode radiated emissions and common-mode voltages and currents present on the board. Using full-wave techniques will not only allow for more accurate simulation of differential-mode radiated emissions, but will allow the effect of coupled nets to be evaluated. EMC engineers need to be concerned with crosstalk in a much more fundamental way than most signal integrity engineers. The type of crosstalk that concerns us in the EMC community is often multilevel in nature and is three orders of magnitude smaller than crosstalk associated with signal integrity. The term "multilevel crosstalk" is used to describe a situation where the signal currents flowing on a high-speed net (e.g., a clock net) couple to and an adjacent net which then couples to yet another adjacent net until finally the "noise" is coupled onto an I/O line like a keyboard or mouse line.

As our understanding of how to effectively simulate these phenomena using computing techniques increases, we will be able to decrease (and maybe someday even eliminate) EMI-related PCB re-spins.

RON MATTHEWS *is the EMI Core Competency Leader at Cadence Design Systems, Inc. Mr. Matthews has worked in EMI and RF design for 20 years, and is a NARTE-Certified EMC engineer. (978) 262-6000.*