


Automation frees high-speed design

UNTIL RECENTLY, TYPICAL DESIGN ENVIRONMENTS WERE JUST SOFTWARE TOOLS FOR TRANSLATING SCHEMATICS INTO PC BOARDS. BUT TODAY'S HIGH-SPEED CIRCUITS NOW FORCE SUCH TOOLS TO ADDRESS THE TIMING AND SIGNAL-INTEGRITY ISSUES THAT CAN CHALLENGE YOUR ENTIRE PRODUCT-DESIGN CYCLE.



IT'S 1992, AND YOU'RE A LEADING-EDGE DIGITAL DESIGNER, eagerly reading the data sheet for the soon-to-be released i486-DX2 microprocessor. You're keen to upgrade your company's i386-powered products just as soon as possible, so you go ahead and design

the schematics for a new version. With the promise of samples on their way, you pass your netlist to your layout engineer, and you soon have fully populated boards to evaluate. The problem is that one board works, another works most of the time, and the third doesn't even boot. Upon investigation, you find that your scope traces and your logic-analyser measurements agree. Those new Schottky-TTL bus drivers that you chose to guarantee timing margins present signal-integrity problems that prevent reliable data transmission. Just like many other contemporary designers, you realise too late that you've entered the world of high-speed design—an environment that ultimately challenges your entire product-design cycle.

Today, high-speed design is no longer the preserve of rarefied logic designers. Although the latest generation of PCs employs a fourfold 133-MHz front-side-bus architecture to yield 533-MHz-equivalent performance, 100-MHz designs are becoming commonplace in mainstream embedded applications. Unsurprisingly, a growing awareness exists among designers that you can no longer simply generate a schematic and expect your pc-board designer to arrange your netlist without consideration for physical-layout effects. Experience with pc-board respins leads to the acceptance that high-speed design is an iterative process, but the earlier you detect potential high-speed layout problems, the fewer your design iterations and the lower your over-

At a glance24
Physical layer determines speed26
For more information31

all design cost. Ideally, you should be able to practice a methodology that ensures correctness by design from the ground up; accordingly, a typical high-speed design flow now comprises a prelayout topological-exploration stage ahead of schematic generation. Meanwhile, silicon density and I/O-count growth force leading-edge designers to consider each IC's package as part of their system-level design, and it's only a matter of time before this approach becomes mainstream.

So, what constitutes a high-speed design, and what are the principal issues that routinely concern product designers? In general, you can classify an interconnect, such as a pc-board trace, as high-speed if the driving signal's rise time is small enough to change logic state in the time it takes for the signal to travel the length of the net and back to the driver. That is, any net that's long enough to delay the signal by more than half the signal's rise-time value exceeds the "critical length" that qualifies the net as high-speed. Notice that this definition comprises no frequency-dependent terms; the signal's slew rate—not the clock frequency—is the important factor. If you have a low-frequency signal with a 5-nsec rise time that commodity CMOS meets, the effective frequency component of the signal's leading edge is $1/\pi \times 5$ nsec, or ~63.6 MHz. Thus, "high-speed design" is a misnomer in the context of a signal's repetition rate, and basic issues must address timing robustness, signal integrity, and EMI/EMC concerns. To meet structured design methods, you need tools to analyse, implement, and manage the multiplicity of parameters that your designs spawn—that is, some form of EDA from vendors (see sidebar "For more information"). Experience with first-generation tools demonstrates that customers are rarely expert users and thus require tools that are easy to learn and use. It's also important that tools are flexible, allowing growth to accommodate ever more dense, complex designs without eternally becoming more complex to operate. And, because today's designs increasingly involve team members at remote locations, the tools must allow you to easily share and exchange data. Ultimately, such tools should encompass your entire design cycle and make it possible to reapply hard-won design lessons from previous work.

Before considering tool-set choices,

AT A GLANCE

- ▶ As many as 80% of today's nets are high-speed.

- ▶ Software tools derive and manage design-constraint data.

- ▶ Consider all interconnecting elements as components.

- ▶ Allow time for up-front analysis to reduce costs.

- ▶ Value your layout-design experience as proprietary IP.

you'll need some basic knowledge that's not so easily available, as the popularity of high-speed-design training classes shows. The two major components that colour high-speed-design practices are time and voltage, for which you need budgets to ensure that physical layout doesn't compromise reliable operation (see sidebar, "Physical layer determines speed"). "Constraint management" refers to assigning and managing values, such as maximum track lengths, and a balance invariably exists between reliable values and acceptable cost. John Berrie, a senior product consultant at Zuken, acknowledges that no high-speed design tool set can solve every problem: "A pc board generates a raging sea of electromagnetic fields that you must discipline. Treat every track as a component." Berrie says that, at first, designers find it easiest and most effective to employ defensive design strategies, such as controlling every high-speed interconnect with global or local ground planes and terminating all clock and control lines. As your experience grows, you can refine your models and reuse your IP (intellec-

tual property) in future designs.

Bob Williams, marketing director at Pulsonix, considers that constraint management divides into two levels: "At the base level, you have guys who simply wish to check that their designs meet predetermined criteria. At the top level, you have expert systems that calculate critical parameters and guide the routing process via interactive or automatic means." Williams explains that, in general, high-speed pc-board-design tools employ a range of net attributes to characterise track paths. Attributes can take the form of track-length constraints, the branch point where a connection divides to another track, and the pin-by-pin net order in which to daisy-chain connections. Other constraints typically include track-width and via dimensions and limits for the number of vias per connection after routing. You may also define maximum stub lengths to constrain track deviations to multiple receivers that aren't directly in the track's path. To provide shielding, you can often specify that a ground or another net surround a sensitive track. Differential-pair or bus-routing rules can apply to pairs of signals or whole buses that must route across a design in close proximity. Systems that include a field solver can accept crosstalk rules and perform physical-to-electrical transformations, such as track length to propagation-delay time.

Ideally, design tools should observe all rules during manual or autorouting runs and guide the router accordingly. Then, when you complete the layout, the tool should verify each rule against the schematic's constraint data. If you need to perform an electrical-signal analysis of your layout, you need an integrated tool that interprets the layout information and board-building data together with electrical models of each device on the pc board. Williams says, "Not many systems have realistically priced products that are totally integrated, but there are some add-on tools to do this specialised work. Pulsonix is able to hook into these tools using its netlister function." Williams contends that if you simulate your

schematic using a Spice tool before layout, the design is a long

way toward being electrically stable; in itself, this step can save many layout iterations and subsequent costs. Pulsonix Version 2 handles a subset of length- and

(continued on pg 28)

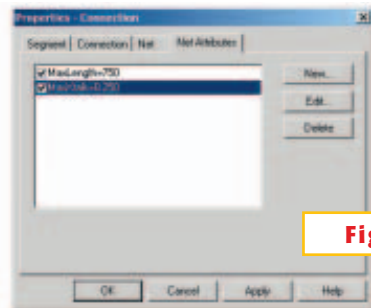


Figure 1

Pulsonix, a low-cost pc-board-design environment, handles constraint data using freeform "attributes" that reports can check.

PHYSICAL LAYER DETERMINES SPEED

To identify high-speed nets, you need to establish the net's critical length relative to driver rise time. Any pulse that travels along a conductor reflects upon reaching the end of the conductor back to the driver. To constrain reflection amplitude, the reflection must arrive back at the driver before the driver changes logic state. Thus, critical length is that value that occurs when the signal's round-trip propagation delay exceeds its rise time; above this value, you likely need termination to control transmission-line effects. Electrical signals travel at the speed of light in a vacuum and at similar velocity in free air. But when an insulator surrounds a conductor, the insulator's relative electric permittivity (ϵ_r) slows the signal by impeding the free flow of the electromagnetic field around the conductor. Assuming that the insulator is a homogenous material, you can approximate the velocity degradation by dividing the speed of light by the square root of ϵ_r .

The relationship between ϵ_r and propagation delay ($\sim 0.033366 * \sqrt{\epsilon_r}$ nsec/mm) yields some interesting results. In free space with unity ϵ_r value, signal speed is close to 298.7 mm/nsec. Ignoring corrections, FR4 pc-board material has an ϵ_r of approximately 4.1 that halves signal speed to yield a propaga-

tion delay of approximately 0.00675 nsec/mm. That doesn't sound like much delay until you estimate a net's maximum non-critical length by dividing signal rise time by twice the propagation delay per-unit-length. For FR4's raw ϵ_r value, the maximum critical length for a 5-nsec driver is approximately 370 mm, falling to approximately 74 mm for a 1-nsec driver. Don't simply use the fastest drivers to simplify timing calculations; rather, select drivers with the slowest slew rate that your application tolerates. Nonlinear equations describe the effective ϵ_r value, which the net's site within the board affects. The net's site relative to layer stacks also affects its impedance. For a pc board with multiple power/ground planes, each net is a transmission line that takes one of four forms (Figure A). A net on the board's surface with a plane below constructs an open microstripline with an effective ϵ_r between air's unity value and that of the pc-board substrate material. Assuming materials with ϵ_r values of 2 to 6, approximate effective ϵ_r with $\epsilon_{r(\text{eff})} = 0.475\epsilon_r + 0.67$ (Reference A). Thus, signals incur minimal propagation delay when you route them on the pc board's surface.

You're likely to route signals into a receiver from sources that are varying distances from that

receiver. If these distances are disparate, the mismatch between propagation delays creates skew between signals that can stimulate erratic operation. Methods of skew compensation include equalising net lengths, controlling driver delays, and balancing nets or buses. Vias complicate delay balance by introducing discontinuities in a net's LC characteristics; stubs off the main trace and the receivers that they serve have a similar effect. Methods for reducing such effects include adding a clearance between the power plane and the vias and limiting the number of vias and stubs per net. It's also likely that a practical layout compromises optimum electrical practices, such as siting clock drivers centrally to the devices that they control, to minimise skew. When possible, site memory modules close to a pc board's centre, where the ground plane works best and provides better high-frequency decoupling than do bypass capacitors. Site fast devices close to any edge connectors that they drive and move lower speed components farther in-board. Other good practices include separating analogue and digital circuit blocks

and providing each with individual ground planes and power-supply-decoupling schemes.

The trend from 5V to less-than-1.8V logic erodes system-noise margins and increases susceptibility to crosstalk. When two traces lie in close proximity, LC coupling transfers bidirectional signal energy between them. The current flowing toward the second receiver is forward crosstalk, and the current that flows into the second driver is backward crosstalk. For short traces, backward crosstalk typically dominates and is interesting in that its magnitude linearly increases with the coupling length until the coupled length equals the net's critical length; beyond this saturation point, magnitude does not increase. It's also interesting that striplines and transmission lines buried in a homogenous dielectric generate zero forward crosstalk. Also note—and contrary to popular belief—you can't suppress the inductive coupling element using a shield trace; it's far better simply to move coupled traces farther apart. Rigorous crosstalk synthesis requires 3-D geometric analysis and a field solver, but you can observe that effects increase with tighter

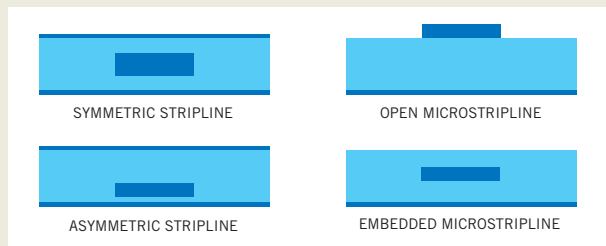


Figure A

A net's location within a pc-board structure influences both propagation delay and characteristic impedance (courtesy Per Viklund of DDE-EDA).

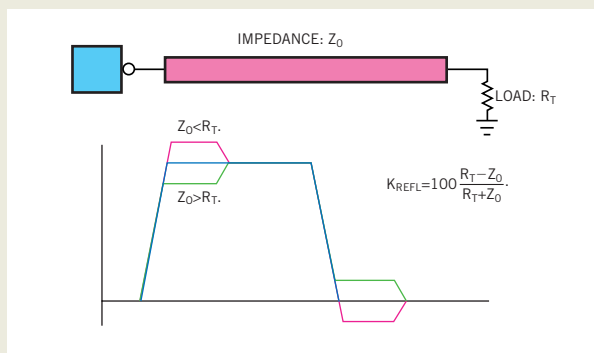


Figure B

If the termination exactly matches line impedance, no reflections result (courtesy Per Viklund of DDE-EDA).

tracking, wider tracking, or both; increasing distance from ground planes and power planes; and lower (shorter) logic rise times. Thus, to reduce crosstalk effects, separate and shorten tracks; use a thin lamination between signal planes and ground and power planes; and design low-impedance, properly terminated traces.

When you can't route nets to less than their critical length, consider transmission-line-termination techniques. If the impedance at the end of the line matches the driver-and-line impedance, pulse reflections and resulting signal distortions will not occur. As a rule of thumb, for nets shorter than their critical length—and regardless of line impedance—reflections are generally less than 15% of the driving signal's amplitude. For nets longer than the critical length, you estimate the percentage reflection amplitude by calculating the mismatch between line impedance and termination impedance, as follows:

$$K_{\text{refl}} = 100 \frac{R_t - Z_0}{R_t + Z_0}$$

(Figure B).

But how do you assess line impedance? You can find the driver's output impedance from data-sheet statistics, but each net is a series of distributed RLC circuits along the net's length. Because the resistance term is normally low, the line's impedance simplifies to $\sqrt{L/C}$, but these values are complex to derive; to solve for a net's impedance, expert signal-integrity tools typically apply Maxwell's electromagnetic-field equations within the structure's layout geometry. In general, pc-board

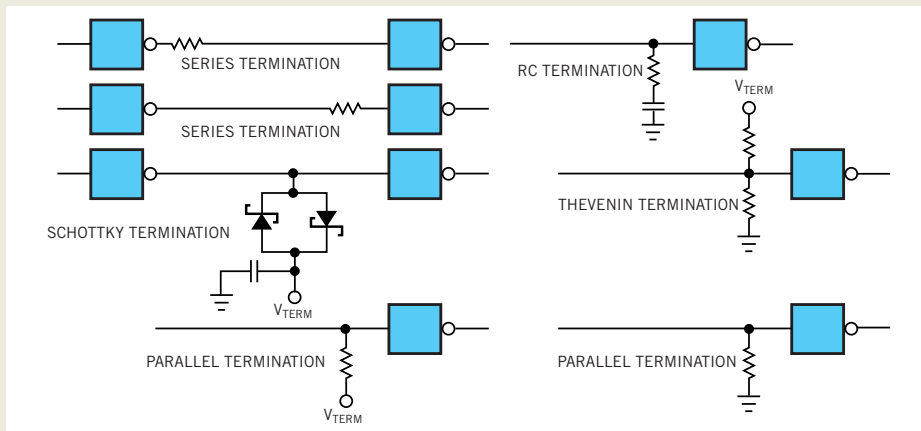


Figure C

Terminations cost components, power consumption, signal delay, or a combination (courtesy Per Viklund of DDE-EDA).

tracking has an inductance of approximately 1 nH/mm, but the stray capacitance is harder to quantify because it's proportional to trace thickness, trace length, and distance from ground and power planes. Math fans can see **Reference B** for numerical methods that derive a net's impedance and **Reference C** for real-world examples. But be sure to acquire Agilent's AppCAD RF design tool, which is freely available from www.agilent.com. AppCAD is a taste of the company's EEs of high-speed analogue-design environment that targets microwave and optical engineers working in the 10- to 40-Gbps spectrum. Among its capabilities, AppCAD calculates line impedance for configurations such as microstrip, stripline, and wire-over-groundplane. You may also explore free downloads from Ansoft, including Maxwell SV, a limited-function version of the Maxwell-2D modeller/field solver for electrical and magnetic fields.

All termination strategies increase component cost, power consumption, and signal delay. Hence, if you can avoid termi-

nating nets through smarter routing strategies, do so; if you can't, consider the trade-offs that each configuration inherits. The simplest termination is a parallel resistor to ground or to a voltage rail; the Thevenin connection uses a resistive divider to bias the line between V_{CC} and ground (Figure C). If you use a dedicated termination voltage, select a regulator that sinks and sources current; most voltage regulators are only current sources. Both approaches increase static power consumption; alternatively, a series-resistor termination trades no increase in power consumption for additional signal delay. The RC termination incurs a dynamic rather than static power-consumption increase, but this tuned-circuit approach is frequency-dependent; the purely resistive terminations work regardless of signal frequency. Be sure to place any termination as close to the receiver as possible.

A strategy that can suit bidirectional nets substitutes Schottky diode clamps to ground and V_{CC} in place of resistors. Another variation that suits

low-voltage signal swings arranges the diodes back to back to divert reflection energy into a termination voltage. If each diode-pair forms a perfect clamp with zero saturation voltage or propagation delay—and each incoming signal is full-amplitude—the arrangement works perfectly. But signals of less than full amplitude don't activate the clamps and remain undamped; also, their residual reflection products don't decay exponentially, as resistive terminations do. It's also important for the diodes to conduct immediately upon receiving a fast wave front; otherwise, the unabsorbed energy reflects rather than dissipates into the supplies.

REFERENCES

- A. IPC-D-317A, "Design guidelines for electronic packaging utilising high-speed techniques," www.ipc.org.
- B. Burkhardt AJ, CS Gregg, and JA Staniforth, "Calculation of PCB Track Impedance," www.mentor.com.
- C. "Transmission line effects in pcb applications," AN1051/D, Motorola Semiconductor, 1990, www.mot-sps.com.

(continued from pg 24)

net-based rules, allowing you to define such constraints during schematic entry (Figure 1). The system automatically forward-annotates constraint data to your layout and similarly handles ECO (engineering-change-order) data. You can check the rules that apply to nets using the integral DRC (design-rules-checking) facility at any stage of design and extract data with the user-configurable report generator. A synchronisation check between the schematic and the layout editors ensures that rules remain consistent between

the electrical designer and the layout engineer. Complete with its Spice-simulator option, an unlimited-capacity version of Pulsonix costs around €6000 per seat.

Familiar to a generation of pc-board designers, the UK-originated Cadstar suite now wears Japanese-owned Zuken branding. The company has just released Version 6, which targets desktop Windows users with core schematic-capture, library-management, and pc-board-layout and -routing tools. Options include an EMC-analysis module, and new signal-integrity and design-constraint tools that suit high-speed design tasks. The EMC adviser module integrates within the design environment and comprises a field solver with 20 predefined algorithms that check for parameters such as crosstalk, track impedance, and propagation delay. You can enter parameters within Cadstar's schematic editor that EMC adviser carries through to layout, enabling you to explore what-if scenarios. Two signal-integrity tool-set options suit postlayout analysis, extracting netlist data to analyse transmission-line behaviour. The base signal-integrity product offers a signal-screening tool that identifies high-speed nets, a reflection simulator and wave analyser that identify undershoot and overshoot, a termination adviser that allows what-if evaluation, and a macro-model editor that imports or builds component models; the enhanced signal-integrity module adds crosstalk-simulation capabilities to determine induced

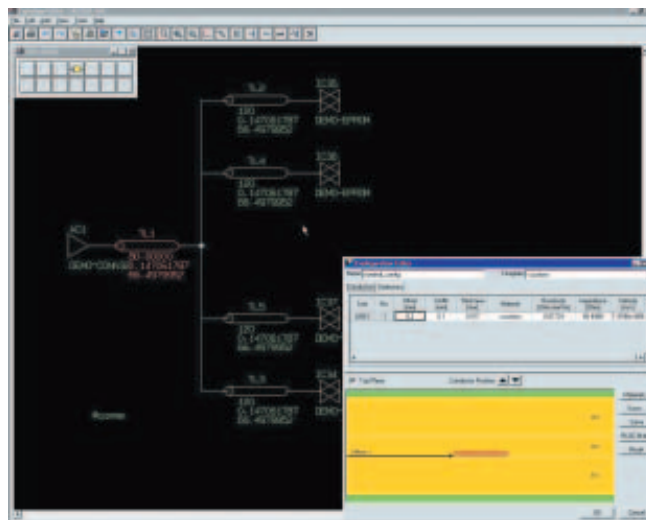


Figure 2

Zuken's Scenario Editor allows you to explore alternative topologies ahead of schematic entry.

voltages in adjacent tracks. You can import model data in the IBIS (I/O- buffer-information-specification) format that component vendors use to describe simulation data while obscuring the underlying IP that Spice models reveal. The constraints manager allows you to set electrical and physical design constraints within a spreadsheetlike GUI, which the system then maintains throughout the design flow. Not all constraints necessarily target high-speed design issues. For example, by setting limits for track-width values, you can improve routing flexibility by defining individual track widths during layout. Depending upon functions, guide pricing for Cadstar configurations spans €10,000 to €20,000.

For enterprisewide use, Zuken offers a modular board-integrity tool chain that caters to design-cycle needs from




Figure 3

Constraint-driven routers within environments such as Cadence Design Studio can automatically balance net length to minimise skew.

preschematic-entry planning through design-and-manufacturing data management. System Designer is the native schematic-entry package, and Board Designer, Visula, and PR Editor furnish a choice of layout and routing tools. To facilitate virtual prototyping, the Hot-Stage package hosts the master design-constraints database and facilitates design verification before the availability of a physical prototype. Hot-Stage integrates Zuken's autorouter technology with a simulator that permits two-way data exchanges, allowing you to quickly explore layout strategies and trade-offs

that compromise signal-integrity parameters (Figure 2). Hot-Stage comprises four design views, starting with a constraint manager that manually or automatically defines and verifies hierarchical design constraints. The scenario-editor facility provides a virtual scratchpad for exploring alternative device technologies, net topologies, and layer-stack arrangements. You can use the scenario editor to perform preschematic-entry what-if analyses and later to experiment with implementation issues, such as evaluating termination strategies. The physical-prototyping tool includes calculators for impedance and crosstalk that help you optimise high-speed net performance.

Zuken's Berrie notes: "Hot-Stage helps you derive generic design-constraint templates that you can view and reuse within the same development environment throughout your current and future design cycles." The tool comes with a set of standard net-topology patterns, or you can develop your own topologies to constrain the router. You can experiment with trial configurations, such as alternative layer-stack arrangements, and run simulations at an early stage to identify potential problem areas. Later, as you develop the physical placement, you can rerun simulations to view parameters such as impedance and propagation delay in real time as you arrange tracks. "Postlayout simulation alone is an old-fashioned approach that can't meet today's chal-

allenges, where frequencies double every 12 to 18 months,” says Berrie. “It’s becoming important to be able to simulate, assess, and then embed design-constraint data ahead of schematic capture.” Entry-level price for the board-integrity tool chain starts at approximately €10,000 per seat, rising to approximately €90,000 for a full system that includes the Hot-Stage tools.

Hemant Shah, director of high-speed systems design for Cadence’s pc-board systems division, agrees: “Postlayout simulation is better than board debugging, but you’re still too late.” Shah says that your first challenge is to find the optimum set of constraints that avoids high-speed-design problems but without incurring unacceptable cost or making your board impossible to route; you then need to implement your findings (Figure 3). This design flow requires the prelayout analysis of a tool such as Cadence’s SpectraQuest. A signal-integrity product family in its own right, SpectraQuest interfaces with the company’s Windows-compatible PCB Design Studio and its high-end Unix/Windows PCB Design Expert suites to suit single-seat to enterprisewide use. The SpectraQuest Signal Explorer edition provides Design Studio with transmission-line analysis that allows you to explore preroute and postroute instances of single-net topologies. The more capable SI Expert edition works with PCB Design Expert and with Mentor’s Graphics Board Station to benefit designers who need signal-integrity analysis and high-speed-design-rule derivation to constrain their component placement and routing processes. Along with the simulation engine and analysis tools, you get a topology editor that allows you to experiment with various component placements, a basic pc-board editor that supports critical-component placement and trial routing, and a hierarchical constraints manager that controls the design-rules database.

Shah also has an enlightening view on the qualification of high-speed design: “If you need to do more than route A to B, it’s high-speed. Most customers accept this realisation only when boards that should work fail.” As designers come to accept that high-speed is now the norm

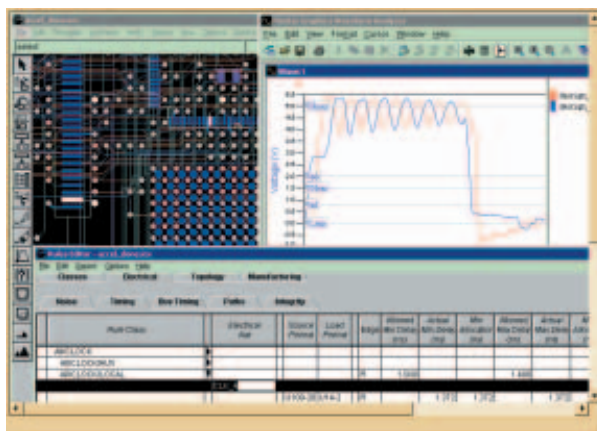


Figure 4

Mentor Graphics' ICX handles all constraint data directly in time and voltage units.

rather than the exception, Shah reports that there’s a migration from the lab-test-and-layout-iteration approach, through postplacement layout analysis, to the structured preschematic planning that modern practice requires. “Do your constraint management ahead of logic design,” he says. Shah also notes that any simulation environment is only as good as the models that it employs. “Although IBIS has been around for about a decade, model quality has until recently been poor,” he reflects. Driven by customer pressure, the semiconductor industry is now providing not just more representative models, but also tool sets that include sample topologies and test waveforms that you can run within your environment to correlate performance with reference criteria. This approach also provides you with a starting point to develop your own layout practices. A guide price for an entry-level Design Studio seat complete with SpectraQuest is €24,000.

Chris Swaim, European marketing manager for Mentor’s high-speed pc-board products, estimates that the latest generation of IBIS models are as much as 98% as accurate as a Spice equivalent. But, as much as accuracy, he says, “engineers need to be able to mix and match IBIS models with others that they already have or can generate, such as Spice and VHDL descriptors.” This multilingual-simulation capability is a key feature of Mentor’s ICX product family and embraces mixed-signal hardware-simulation languages—notably, IEEE 1076.1. Also known as VHDL-AMS (analogue and mixed-signal extensions to very-high-density logic), IEEE 1076.1 is a structured language that combines the

behavioural-modelling approach of IBIS with Spice’s freeform topological descriptors (see www.eda.org/vhdl-ams/). This approach especially suits the emerging generation of multigigabit-per-second serial buses, such as Yellowstone, in which edge-rate-programmable I/O structures and phase-adaptive receivers enable designers to match the silicon to its interconnection environment (see www.rambus.com).

Now in Version 3.0, Mentor’s ICX simulation environment integrates with the com-

pany’s Board Station and Expedition design flows to support designs with greater-than-500-MHz clock and subnanosecond edge rates. Uniquely, ICX employs an all-electrical approach to constraint rules that operates in direct engineering units (Figure 4). Swaim says, “Engineers prefer to think in terms of time rather than length and in millivolts of crosstalk rather than abstracts, such as capacitive and inductive coupling values.” The ADMS (analogue-digital-mixed-signal) single-kernel simulation engine at the core of ICX allows you to evaluate analogue and digital phenomena in one run, rather than setting up and managing separate scenarios.

Swaim believes that today’s biggest obstacle to effectively tackling high-speed designs is the need for engineers and their managers to change their work practices. “Accept that you’ll spend longer up-front analysing to avoid respinning your designs,” he advises. “This philosophy will save you time and money.” Consider the relationships between driver strength and receiver characteristics up-front rather than overconstraining designs with unnecessary rules. “Timing is almost always the most important issue, and establishing reliable margins is the key,” Swaim says. Accordingly, ICX 3.0 tightly couples with Mentor’s Tau board-level timing simulator to perform worst-case timing analysis and verification. Tau uses a symbolic timing methodology to eliminate the false paths that typical static-timing tools report. The program dispenses with the need for separate test vectors by taking advantage of models that contain both functional and timing information. Like ICX, Tau integrates

with your physical design by generating constraints that directly drive the placement and routing processes. Prices for ICX 3.0 and Tau 3.0 start at approximately €45,000 and €35,000, respectively; prices for Board Station seats start around €28,000.

But don't think that only billion-dollar multinationals can develop tools and environments that support correct-by-design methodologies. Incorporated as a separate company last year with headquarters in Denmark and offices in Sweden and the United States, DDE-EDA has its origins as Dansk Data Elektronik's EDA division. In development since 1975, the company's Supermax ECAD is a front-to-back design system with features that include constraint-driven placement-and-routing functions that can interpret electrical, signal-integrity, and manufacturability design rules. Today's trend toward silicon-package-board convergence drives Per Viklund, DDE-EDA's product and technical support manager, to recognise that a fuzzy boundary exists between pc-board, multichip-module, and device-packaging technologies: "Today, it is crucial that the designer can stay in one environment for the complete design flow," he says.

Accordingly, Supermax ECAD handles high-speed-, microvia-, and RF-design practices. Usefully, it allows you to view your entire design in 3-D to provide an intuitive visualisation of alternative pad- and layer-stack strategies (Figure 5). The program's single-file ASCII-format database facilitates and simplifies data ex-

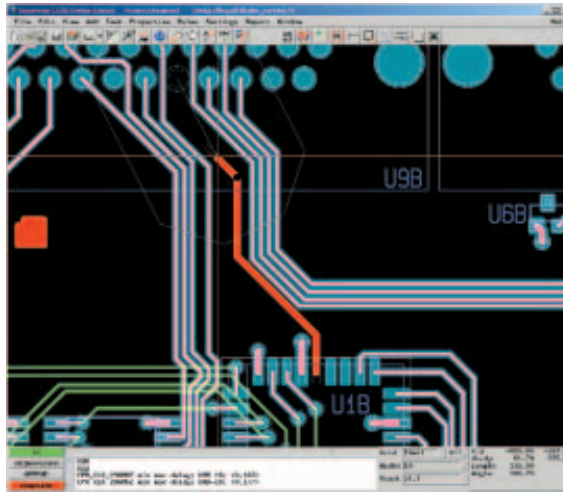


Figure 5 Supermax ECAD's white-octagon outline shows the remaining delay margin for a high-speed track during routing.

changes with third-party environments, such as thermal solvers, solid modellers, and other CAD tools. A comprehensive option list includes the Supermax ECAD high-speed-design package that adds facilities including controls for topology, timing and reflection, termination, crosstalk, and impedance. These functions operate in automatic, interactive, and batch-check modes. You can also select options such as Supermax ECAD RF Design, which provides two-way integration with microwave tools from Agilent and Mentor; an advanced editor that provides interactive 45° push-and-shove capabilities; and various viewer and librarian functions. Suitable for Windows and various Unix environments, including Linux, the pc-board-design environment and the high-speed-design package sell for €25,000 per seat.

In Santa Clara, CA, Sigrity develops software products that extend the capabilities of a multiplicity of silicon-package-board design flows, including Mentor Graphics' ICX tools, for which Sigrity is a Mentor Open-Door program partner. A component of the company's SpeedXP suite, PowerSI is a new, Windows-compatible application that also imports Cadence and Zuken file formats. Complementing Sigrity's Speed2000 time-domain electrical-analysis tool, PowerSI targets frequency-domain operation. Both tools share an analysis environment that supports high-performance prelayout and postlayout power- and signal-integrity analysis of IC packages and pc boards.

Ekkehard Miersch, consulting engineer at Sigrity's German representatives EFM Consulting, observes that the solvers within Speed2000 and PowerSI are the only tools of their kind that perform Spice-based and full-wave-electrodynamic analysis of complex high-speed pc boards and planar-chip packages as complete, 3-D entities. Miersch says, "Speed2000 and PowerSI let you observe expected and unexpected high-speed effects in great detail. This [feature] helps you to understand the physics behind your design, which reduces prototype builds and facilitates design improvements." Sigrity's newest and complementary product, Broadband Spice, converts frequency-domain network parameters into accurate Spice-circuit models over broadband frequencies at the touch of a button. SpeedXP applications include characterising and optimising the effectiveness of power and ground planes and distributed decoupling capacitors; the avoidance of self-resonance in or between circuits, packages, and boards; identifying and minimising the effects of signal-return-path discontinuities and simultaneous switching noise; and predicting electromagnetic radiation from packages and boards. Expect to pay approximately €95,000 for the Speed2000/PowerSI/Broadband Spice combination. □

You can reach Contributing Editor David Marsh at forncett@btinternet.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to www.edn-info.com and enter the reader service number. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN Europe*.

Agilent
www.agilent.com
Enter No. 350

DDE-EDA
www.dde-eda.com
Enter No. 353

Pulsonix
www.pulsonix.com
Enter No. 356

Ansoft
www.ansoft.com
Enter No. 351

EFM Consulting
miersch@attglobal.net
Enter No. 354

Sigrity
www.sigrity.com
Enter No. 357

Cadence
www.cadence.com
Enter No. 352

Mentor Graphics
www.mentor.com
Enter No. 355

Zuken
www.zuken.com
Enter No. 358

SUPER INFO NUMBER

For more information on the products available from all of the vendors listed in this box, enter no. 359 at www.edn-info.com.