

Printed Circuit Board Routing at the Threshold

by DAVID WIENS

During the past several years, we have witnessed an unparalleled wave of technological innovation and rapid market adoption. Spurred on by global competitive pressures, new product introductions have come fast and furiously. Sustaining this blistering pace means decreasing both the time and cost of product design cycles. To successfully collapse design cycles in this way, manufacturers must achieve a significant increase in productivity: the Holy Grail of the technology-driven marketplace. It follows that design tools must also become more productive to facilitate attainment of this goal. In today's environment, you get your PCB done right the first time, or pay a heavy price.

Development schedules are shrinking at a dizzying rate—and so are the sizes of the products themselves. Suddenly, everything is shrinking, from cell phones, laptops, and digital cameras to the integrated circuits and printed circuit boards that drive them. The push for smaller form factors is accompanied by equally insistent demands for more functionality and better performance. Make it smaller and faster, more powerful, feature-rich and reliable—all at the same time. No wonder the challenges for PCB designers are growing. Now more than ever, their software tools must help them deliver measurably better products within the framework of an efficient and cost-effective design process.

In routing terms, this translates into faster routes, higher completion rates, and fewer iterations. It also means producing smaller, denser boards with fewer layers, manufactured at less cost. In addition, it implies the ability to handle the newest improvements

in chip, packaging, and fabrication technologies as they emerge. Finally, the router itself must be seamlessly integrated into the core design environment, as well as easy to use. All this may sound straightforward, but only the most sophisticated PCB routers are truly up to the task.

THE SILICON CONNECTION

The breathtaking rate of technological progress in the integrated circuit arena is responsible for the present advances—and pressures—in printed circuit board design. Foreseen by Moore's Law, die sizes are shrinking as capacity increases, aided by the use of ever-finer process technologies. In conjunction, signal-switching times are accelerating. These combined effects translate directly into added complexity for the PCB designer downstream. Faster clock speeds and device edge rates have introduced a host of high-speed issues into the board routing process—issues which threaten to lengthen design cycles, increase costs, and adversely affect product quality if not properly addressed.

The corresponding improvements in IC packaging technology have proven to be just as challenging. High-density, fine-pitch packages like the ball grid array (BGA) house devices with pin counts higher than 1,000 while occupying minimal board space. The benefits these packages provide (the ability to create smaller, denser, higher-performance boards) come with a price, however.

The reduced space between a larger number of package pins dramatically increases routing challenges, and can require the use of advanced fabrication techniques to com-

pensate. The truth is, the newest packages are testing the limits and revealing the weaknesses of some entrenched PCB routers.

CRUNCH TIME FOR PCB ROUTERS

For many years, the only way to pack more functionality onto a printed circuit board was to increase the board size and the number of layers, which, in turn, raised production costs. Those days are over. Now the preferred method is to reduce both board and package size, while maintaining or increasing pin counts. This is a tall order and it's only going to get harder as package counts rise.

Perhaps not surprisingly, those designing more complex boards today anticipate that their next designs will take them even longer. Because layout and routing comprise a major portion of PCB design, it seems clear that performing those activities more efficiently will have a direct impact on shortening the entire cycle and enhancing overall productivity. The changing times require a robust printed circuit board routing solution, one that addresses high-speed issues, handles new packaging and fabrication requirements, is automatic and interactive, integrated, and easy to use. Only then will the community of printed circuit board designers be well equipped to manage the growing technical challenges in the months and years ahead.

ADVANCES IN ROUTING TECHNOLOGY The Packaging Effect

What makes an autorouter productive? Speed, completion rate, high route quality, ability to apply and maintain high-speed rules, and the support of current packaging

and fabrication technologies are all important factors. Yet historically, changes in packaging and fabrication technology have caused the most important improvements in router performance and productivity. Routers that could not keep pace with those technological advances eventually faded away, to be replaced by the next standard.

A BRIEF HISTORY OF ROUTING

In the days of through-hole components, gridded routers dominated. Because through-holes were relatively large and pin-to-pin spacing was wide, the task of routing traces between pads was generally a straightforward exercise. With the advent of surface mount technology, pin pitches began to shrink. The big advantages surface mount offered at the time were smaller footprints and higher pin counts. While the first surface-mount components featured pin pitches of 25 mils, they decreased over time to around 11 mils. Minimum trace widths and clearances decreased accordingly, putting a tremendous strain on existing routers. The finer the grid, the slower the router ran until, eventually, the gridded router was replaced by the gridless or shape-based variety. Shape-based routers took hold in the mid-1990s, through the era of the pin grid array and the leadless chip carrier, and into the present day. They proved to be better than their predecessors at handling a variety of components and pitch sizes, especially finer-pitched devices. The most prevalent shape-based routers—with algorithms based on 90-degree angles—performed well until faced with high-density, fine-pitch attach packaging, at which point certain fundamental shortcomings of the 90-degree approach became abundantly clear.

THE KEY: NATIVE 45-DEGREE ROUTING

The large number of fine-pitch pins on a given device, the growing number of devices on a board, and decreasing board sizes overall, make today's routing challenges more difficult than ever. Smaller pads, shrinking trace widths, and tighter clearances create problems for the existing 90-degree-angle routers. The difficulty is amplified by the presence of numerous staggered pin connectors. In addition, 90-degree routing fails to adequately address the critical BGA escape problem, unable to provide the necessary 45-degree-angle via fanout from tightly packed solder ball pins.

Some 90-degree routers attempt to include 45-degree angles as a post-process,

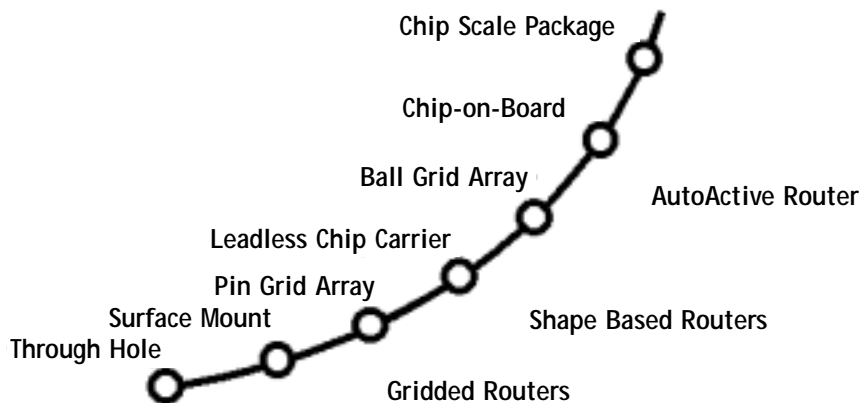


Figure 1. Packaging Changes Affect Router Evolution. Next-generation routing technology is necessary to handle the newest device packages, including successive generations of the ball grid array, chip scale packaging, and chip-on-board, among others.

which is, in essence, a productivity killer. Because the entire board is first routed with 90-degree-angle traces, the designer is not able to make the best use of already-limited board space. Certain areas are simply not available, crippling the router performance. Other capabilities useful for some high-speed designs—such as trace corner chamfering—are also virtually impossible to perform with a 90-degree router. In summary, when vendors attempt to retrofit 90-degree routers to meet current challenges, they end up compromising speed, performance, and adherence to high-speed design rules. In the end, it's the designer who loses.

The ideal solution is a fully integrated, shape-based, native 45-degree PCB design system, where placement, interactive and automatic routing all run within a unified environment, using the same editor. In this situation, the designer can actually build the via fanout into the BGA footprint during layout and then rely on the automatic router to rapidly complete the task, optimizing escape routes according to pre-established design constraints. If interactive intervention is subsequently required, any changes the designer makes will automatically incorporate 45-degree angles as needed.

ROUTERS AND THE MICROVIA REVOLUTION

To accommodate shrinking packaging, PCB interconnect (trace and via) dimensions must also change. Through-vias are too large and unwieldy to work with BGAs and other




Figure 2. Native 45-degree Routing. Only true 45-degree routing gives the designer the flexibility to deal with the twin problems of growing density and scarce real estate. Here, the router automatically runs multiple traces between staggered pin connectors.

miniature, high I/O components. Adding board layers is not an option either; on the contrary, designers are seeking to decrease layer counts, in order to fulfill corporate directives calling for lower production costs and reduced manufacturing times. Advanced, finer-geometry interconnect appears to be the answer for achieving denser routing and, at the same time, lowering the number of layers. The challenge for the PCB designer is managing the higher level of complexity this technological solution brings with it. Once again, it is the router that lies at the heart of the effort.

In response to the current dilemma, microvia technology has finally come into its own. Long considered promising but expensive, microvias are now becoming more viable, as burgeoning demand and process




Figure 3. Buildup Technology Enables Microvias. In this eight-layer board, the middle four layers contain through-vias and blind-vias manufactured utilizing the traditional laminate process. In contrast, the top and bottom two layers are created using the buildup process, introducing new via structures and rules to the routing problem.



Figure 4. Rules by Area. Another approach to routing in congested device pad areas is to define different, smaller rules within a limited region.

advances drive the associated manufacturing costs down. Tiny microvias have become the method of choice for routing designs containing BGA and CSP components. The proliferation of microvias has also revived the popularity of blind and buried vias, as a way to facilitate signal escape from dense, attach-technology arrays. The microvia revolution has been made possible by several noteworthy manufacturing developments: buildup fabrication processes pioneered in Europe and Japan, and enhancement of laser and plasma drilling techniques.

Clearly, the PCB router must incorporate advanced interconnect features to ensure the correct functioning of designs utilizing microvias. New requirements include allowing vias between any two layers, with rules and delay values per via span. Support for blind and buried microvias is essential: a blind via connects the surface layer with one

or more internal layers, and a buried via interconnects internal layers only. Microvia support must be included in the automated routing capabilities, to further decrease route times and, hence, the entire design cycle.

To achieve the tightest wiring densities possible, vias can also be placed within pads. Called via-under-pad or via-in-pad, this feature is necessary when routing pin-outs located in the center or interior of a high I/O array. Surrounding pin density can often make it impossible to run traces through these congested areas, resulting in an incomplete route. To access an interior pin, the router must be able to automatically drop a microvia directly through the corresponding pad, allowing the signal to exit to another layer, thus sidestepping the top-level density problem.

HIGH-SPEED DESIGN CHALLENGES

High-speed issues—including both timing and signal integrity—are now taking center stage for designers of digital printed circuit boards. Customarily, high speed has referred to boards with clock speeds of 50 MHz or above. By that measure, somewhere between 50 and 60 percent of board designers are presently coping with high-speed effects. But the magnitude of the problem is actually much greater still. High-speed issues are actually caused primarily by fast edge rates. Simply replacing an old component can

compromise signal integrity on traditionally “slow” boards.

A related indicator of the growing challenge for designers is the rise in the number of critical or high-speed nets on a board. In the past, only two to five percent of nets were considered critical; today that number commonly surpasses 50 percent. In some high-end applications, the number of critical nets per board can reach 90 percent. With the average number of nets per board around 5,000, today’s designers have their hands full managing the consequences of high speed.

A CLOSER LOOK

Conditions are ripe for a wave of signal integrity and timing problems to strike the board design world. Why? Because three critical factors are rapidly converging: board densities are increasing, clock frequencies are climbing, and device-switching speeds are dipping into the 0.5-nanosecond range. Resulting signal integrity issues can include increased noise, ringing, reflections, coupling, ground bounce, and crosstalk; if unresolved, they cause serious signal degradation. In turn, timing issues such as excessive gate delays, excessive interconnect delays, clock skew, and signal instability can give rise to troubling switching errors. Boards suffering from such high-speed effects exhibit intermittent symptoms that can be difficult to diagnose, but ultimately compromise reliability and may eventually lead to product failure. Obviously, such an outcome is as undesirable as it is costly. The question for board designers—many of whom are unfamiliar with electrical engineering principles—is how to successfully navigate this potential minefield, delivering a high-quality end product on schedule and within budget.

THE RIGHT TOOLS

“Rule-of-thumb” or overly conservative design rules are no match for the new era of high speed. Nor are many present-day printed circuit board routers. In order to successfully handle high-speed boards, a router must be able to apply an extensive set of complex design constraints, both automatically and interactively.

Important high-speed design constraints include:

Timing

- Maximum net length or delay
- Matched length or delay
- Skew control




Figure 5. The Importance of Interconnect to Signal Timing and Integrity The net in Example 1 is unconstrained, resulting in a signal with poor quality and skew problems. The net in Example 2 utilized a star topology with matched driver-to-load delays, resulting in a clean signal with minimal skew.




Figure 6. High-Speed Constrained Routing. Constraints for net topology, differential pairing, and matched delays ensure this design will operate with clean signals and correct timing.

- Delay formulas
- Pin connection order

Signal integrity

- Differential pairs
- Layer restrictions
- Trace width range
- Routing topology and priority
- Termination scheme
- Controlled impedance

The router must be able to perform customizable, automatic net tuning to ensure that the physical design is always meeting all predefined constraints. The same capability should also be available interactively; this makes it possible for engineers to place and

route critical components and nets up front, before handing the board over to the designer. This type of process flow is common practice where high-speed designs are concerned. Ideally, all activity takes place within a single, integrated editing environment, so that design rules are instantly available at every stage in the process, from early constraint definition—again, often done initially by an engineer—through logic design and physical layout.

The ability to auto-route and tune differential pairs is particularly important for high-speed design. It's not unusual to encounter complex, high-performance boards that contain 1,000–2,000 differential pairs, particularly in the areas of computers, networking, and telecommunications. Differential pairs are typically used to reduce signal degradation over long distances, sometimes between several boards. For that reason, boards intended for telecommunications base stations often make heavy use of differential pairs.

Verification of high-speed routing typically requires tedious and time-consuming iterations between transmission line simulation and routing. A new breed of routers is integrating the traditional strengths of routing technology with simulation into a single automated step known as interconnect synthesis. This integration eliminates the need for back-end verification, producing higher quality designs in a shorter time.

NEXT-GENERATION DESIGN ENVIRONMENT

The market and technology changes now affecting printed circuit board design are fundamental, fast moving, and irreversible. Routers that were built to address the issues of another era are no longer adequate to manage today's challenges, as many PCB designers are discovering. PCB routers for the next generation of designs must deliver

uniformly higher performance; handle the latest developments in IC, device packaging, and board fabrication technology; include advanced high-speed design features; provide both automatic and interactive capabilities; and, above all, be easy to learn and use.

PRODUCTIVITY MAKES THE DIFFERENCE

In the final analysis, the PCB router that can deliver the greatest gains in productivity will win the designer's confidence, especially now, as market pressures intensify to reduce both the time and cost of product design cycles.

PCB design tools make a concrete contribution to the designer's bottom line by reducing routing times, improving route quality and manufacturability, and decreasing production costs. These benefits will only become more important as we head into a new millennium, one that promises ever-greater technological advances to come. ■