sponsored by

Gabe on EDA • EDAmarket

Assembling the Future

A Newsletter About the Design and Production of Electronics

ISSUE 010 • DECEMBER 2011

SUBSCRIBE

PDF and Archives

twitter

In this issue:

PCB Design Can Be As Complex As IC Design

Gabe Moretti

About a month ago I attended Z-DAC, the yearly conference held by Zuken for both customers and prospects. I had not seen in person an advanced Printed Circuit Board (PCB) design or analysis tool for quite sometime. At the conference I was reminded how advanced PCB design really is. Most of the problems that are discussed daily in the press and on the web regarding IC design also exist, and in most cases have been solved, for PCB design.

The major difference in the development of an IC versus a PCB seems to me to be how engineers deal with the problem. IC designers approach the problem in an abstract way, using tools like higher level languages that have no relationship with the physical implementation of the design, while PCB designers deal with the physical realization of the design from the very beginning of the project.

A component or a functional block of components assumes a physical reality immediately as its location on the board is chosen. The physical characteristics of the board also must be dealt with form the beginning. Its shape, its dimensions, in fact its physical surroundings all come into consideration much sooner than what happens for IC designs.

PCB designers have dealt with three dimensional aspects of their products longer than the majority of EDA vendors have been in business, using tools from the pioneers of the EDA industry that have

now disappeared through mergers and acquisitions.
IC designers do not concern themselves with how their logic design is actually fabricated until their design becomes a circuit netlist that must be validated according to rules defined by the foundry. This is a source of additional development cost and, I suspect, even larger manufacturing costs than it would be necessary.

On the other hand, PCB designers must deal with how portions of the circuit are positioned on the board. The limited size on the surface creates the need to add layers. This, in turn, creates advanced physical requirements that are solved by using buried and blind vias and embedded components. PCB designers are quite aware of the implications on manufacturing costs of their choices.

While 3D issues in IC design are the result of trying to achieve a smaller footprint for a transistor, 3D in PCB design is the result of physical size restrictions on the design. On the other hand IC designers do not consider manufacturing issues until they must address the requirements of the mask set needed to fabricate the circuit. In PCB design, partitioning and placing functional blocks is a job that must be dealt with during the entire design process. Adding or removing one layer in a PCB has significant fabrication cost implications, as well as it impacts the dynamic characteristics of the system.

The router employed to connect the components, both active and passive, of the circuit must, of course, be able to work in 3D, something that is not required of an IC router. Thus the router used in PCB design is much more sophisticated than the IC routers. It is true that an IC router must deal with much larger designs, but it does so on a flat surface without having to deal with connecting two non-adjacent layers without disturbing the intermediate ones.

As the article from EVE points out the technological requirements for PCB design are just as demanding as those for IC design. The complexity is increased by the demands that the fabrication of a PCB product makes on the engineers. While the burden of producing a successful IC design is split between architects, logic developers, and circuit implementers, PCB design efforts are more concentrated on one type of developer. The development team must be, by necessity, architect, developer and physical implementer. Only the actual circuit analysis for power distribution and consumption, as well as noise reduction and EMI characteristics can be handed over to specialists.

It is clear that the often heard refrain that IC designers are a class above PCB designers is a myth without any basis in reality. The mystique of developing something that one does not get to see or touch, gives an IC its privileged status, but today's PCBs almost always contain internal layers that are quite sophisticated and include not just connections but also active components.

The Market

As in IC design, three vendors, Mentor, Zuken and Cadence, represent about 75% of the market. While both Mentor and Cadence are in the top three vendors in PCB and IC design market, Synopsys has chosen not to be in the PCB market. Zuken, on the other hand, is playing a significant role in PCB design and analysis but has no presence in the IC design market.

While the IC design market is served by at least another one hundred small companies besides the largest three, the PCB market is served by at most a dozen EDA tool vendors. This is spite of the

fact that Gary Smith's Wallchart (http://www.garysmitheda.com/paper/CADCAM_Wall11.pdf) shows 32 companies offering PCB design tools. Some companies on the list no longer offer PCB design tools, like Aldec, while others have very small yearly revenue numbers and are really equivalent to family businesses. Familiar names, like Orcad and PADS are no longer independent companies. Cadence owns and distributes Orcad products, while Mentor does the same for PADS.

While the IC design market has seen growth and differentiation with the introduction of new methods, such as formal verification and High Level Synthesis, the PCB market in general as seen only incremental improvements of existing tools. The release by Zuken of the CR-8000 product has been a major event in this market. The product is a totally new tool that improves functionality, human interface, and integration, something not seen since the introduction of Expedition PCB by Veribest in 1997.

The PCB tools market which includes the support for the development and analysis of boards, as well as wiring and harnesses, is growing. As IC grow in sophistication so must the boards that support them. But it must be pointed out that PCB designers are very conservative when it comes to adopting new tools. Mentor knows something about this. it acquired Veribest in December of 1999. At the time it marketed its own PCB desing tool called BoardStation. It is clear that Expedition PCB is a more modern tool with capabilities superior to those in BoardStation. Yet, almost eleven years after the acquisition, Mentor still markets and sells both products. The fact that Mentor has not been able to lower its development, marketing, and sales costs by unifying the product line is either a strong signal that BoardStation users do not want to change, or that Mentor does not possess the engineering capabilities to unify the products. Of course the System division is highly profitable, mostly through acquisitions, so it may be possible that lowering costs is not as important as keeping the status quo and not risking a comparative study by its customers when they are asked to change.

Big Boards With Lots Of Everything

Francois Douezy, Hardware Manager and Helena Krupnova Product Manager, EVE

EVE makes emulator systems that allow system-on-chip (SoC) designers to validate both hardware and software on billion-gate devices by exercising billions of clock cycles prior to committing their chips to silicon. These systems comprise many boards, each of which is made up of many interconnected FPGAs and other support components that can be configured to implement the user's design.

In order to have enough logic capacity for such large designs, these boards must pack as much FPGA content as possible. They're high-performance, complex boards with as many as 32 layers, roughly half of which are power and ground planes.

But these systems have an even more fundamental requirement that most other boards don't have: they must be able to implement a wide range of designs. They aren't geared to one specific design, so they must be over-provisioned with resources; any given design will leave unused nets or FPGA logic or clocks or something. Those unused resources have to be there for some other customer's design that will need them.

This creates significant challenges for the board design team. Getting it right requires careful planning both at the architecture and design stages.

Before any design is done, the architecture must accommodate a roadmap that will ensure a long life for the board. This is particularly challenging because we need to keep up with the latest FPGA generations, which are revised every 1½ to 2 years – limiting the amount of time we have between generations. So a good plan will minimize the rework required.

This also means that board planning and design must start as early in the system design process as possible, long before details are available. This is only possible through a close FPGA vendor relationship that can point us in the right direction as early as possible. Such relationships take years to develop.

The specific high-level things we must plan for include:

- the details of the FPGA to be used;
- the memory interfaces;
- all of the interconnect interfaces to be supported;
- all of the power supplies: how many and how to distribute;
- debug access – particularly important on a board of this complexity;
- and future generations of the FPGAs.

Because of the high performance requirements our users place on us, we must pay particular attention to clock resource planning to ensure availability, good quality, and low skew.

Reliability is also critical: we have to plan and design the board in a way that ensures that we identify any issues before delivery to our customers. They'll be using our boards to debug their designs – they'll be very angry if they end up debugging our board instead.

Once we get to the design phase, the biggest obvious issue is routing a mix of different high-performance signals. These include noisy single-ended signals, differential pairs, and multi-GHz serial signals. These all have different design rules and require experience for a successful design. Good simulation prior to implementation is critical to ensuring that we're starting in the right direction.

The problem gets harder because automatic routing doesn't provide good results for us. It tends to introduce too many vias, which hurts signal quality. That means that a high percentage of the nets must be routed manually – an enormous task that takes time and requires experts that understand the needs of the different types of signals.

Unfortunately, for the most part, the layouts around each FPGA can't be stepped and repeated for all FPGAs. Careful planning can provide some regularity, but it takes a very experienced layout engineer to do this. Busses can be routed together – say, 20 signals at a time, meaning that they

don’t need to be individually routed. But the engineer must ensure that the nets in the bus are electrically matched to minimize skews. This also ties into the mapping of I/Os from the FPGA – a change in mapping can improve the board layout. We also use tools to tune the signal quality.

Numerous power supplies create yet another challenge. They arise from the various electrical and I/O standards, and some of them must be programmable. It’s a lot of work to implement these in a way that provides good power integrity without going overboard on the number of power planes required.

These FPGAs also draw a lot of power, around 20-25 W per chip, so we have to make sure the boards stay cool. We do that primarily through the use of heat sinks and fans. But we also add a temperature monitoring scheme that provides several levels of warning so that users can take action if things get too hot – and so that, if necessary, the system can shut itself down gracefully before going up in smoke.

Finally, we need to pay attention to the capabilities of the board and assembly subcontractors. If we place too many constraints on the board, it won’t yield well. In addition, signals above 1 GHz aren’t built on FR-4 – they need more exotic materials that not every subcontractor can produce.

We can also make life difficult for the assembly guy, given that the board will have thousands of components, all of which need to be placed and spaced in a way that works for automated assembly. Making matters more difficult are the numerous I/O standards we have to support, each of which has its own physical plug requirements, and all of which must fit and play nicely with each other.

Roll all of this together, and the board design process becomes no less challenging than the overall system design itself. Add to that the fact that we’ll have to do another one every couple years, and we’ve got our work cut out for us.

About the Authors:

Francois Douezy has been hardware manager at EVE since 2001. He graduated from ECE in Paris with a degree in Electronics Engineering ECE and began working as an image processing hardware designer for NOESIS. Douezy then joined META-System/Mentor Graphics as a hardware designer, than hardware manager for Mentor’s emulation division.

Helena Krupnova is product manager at EVE. Previously, she was responsible for FPGA prototyping within STMicroelectronics central verification/emulation team. She holds a PhD on FPGA prototyping and multi-FPGA partitioning techniques from Institut National Polytechnique in Grenoble.