sponsored by

Gabe on EDA • EDAmarket

Assembling the Future

A Newsletter About the Design and Production of Electronics

ISSUE 021 • NOVEMBER 2012

SUBSCRIBE

PDF and Archives

twitter

In this issue:

THIS IS THE LAST ISSUE OF THE NEWSLETTER FOR 2012. WE WILL RESUME PUBLISHING WITH THE JANUARY 2013 ISSUE. THE ENTIRE PRODUCTCION TEAM WISHES ALL OF YOU A RELAXING HOLIDAYS SEASON AND A PROSPEROUS NEW YEAR.

The Approaching Discontinuity

Gabe Moretti

The world of EDA is about to change. The subtle signs are there for all to see, and the coming reality is so different to be scary to some. Thus better not to talk about it. The changes will include how ICs are designed, developed, and verified. They will involve designers, tools developers, and manufacturers, and force an integration that the EDA industry has not experienced so far.

I have followed with great interest the various press releases from TSMC, Cadence, Mentor, and Synopsys describing the work, and the progress, toward finalizing a commercial grade 20 nm process. It is interesting that the vast majority of the news is about TSMC. There is a perplexing lack of news from other foundries about their work on the 20 nm process. Thus the question: are they already done or are they lagging behind?

I tend toward the second explanation. Accustomed to moving from one processing node to the next with regularity, I believe that most commercial foundries have been caught by surprise by the increased difficulty that the 20 nm process holds. it is not just a matter of developing a cell library, or to create and calibrate a new nanolithography process.

What the EDA companies and TSMC have found is that the processes of design and of manufacturing can no longer be considered as two separate methods. The traditional development of a new process was sequential and stand alone, while the new one reveals itself to require non-trivial feedback loops among the foundry and the EDA companies. So far it looks like we have managed to avoid getting the designer involved, although not for long.

Traditionally the foundry developed a cell library and a set of design rules. Of course the design rules have become more numerous and more complex with each new process. Then the EDA companies, using these as input, went to work to either improve existing algorithms or invent new ones that would allow designers to develop circuitry to be manufactured with the new process. The goal of developers was to start with an architectural concept driven by market requirements and end with a layout that could be manufactured using both new circuitry built with the cell library and re-used IP.

The responsibilities were well defined and delimited. Designers had the responsibility to design and develop circuitry that performed the desired functions within specified physical characteristics. EDA companies had the task to develop and support computer based tools that allowed the synthesis of circuitry so that it could be laid out and verified, again using EDA tools, ready for manufacturing, and that provided a way to verify that the tasks were done without errors. The foundry had the task to manufacture the devices with acceptable yields. In the end all made a profit while working in well defined areas of expertise.

Overflowing the Boundaries

As the traditional methods began to be employed to support the 20 nm process, it soon became obvious that such clean demarcation of responsibilities could not be maintained. For one thing the design rules became so complex using the established method to significantly impact productivity, and thus profitability. Transforming a design into a manufacturable layout now requires more knowledge of the manufacturing process than ever before. And, on the other hand, developing a cell library and a set of design rules requires more knowledge of the development process than before. Thus the work of the foundry and the EDA companies had to be more integrated than ever before.

The issue is no longer can we implement algorithms that respect the design rules without asking designers to know more than they are prepared to know about semiconductors physics and manufacturing without changes to the manufacturing process? The answer turns out to be negative. The manufacturing process is no longer a given. It must be developed with the knowledge of what it is possible to do with EDA tools in a profit making environment.

Obviously, given infinite time and resources, the problem is solvable. But the semiconductors industry is now fueled by consumers products. These are notoriously low margin products that require relatively short development times and large volume of product in order to be successful. Short development time means that the process of transforming an idea into a layout must be performed by average engineers in the shortest amount of time, generally 9 to 12 months. And the

resulting circuitry must be manufacturable with yields that will result in a profit margin sufficient to assure a return on investment that is competitive.

The result is that professional staffs from EDA companies and foundries have had to work together more than ever in order to explore and test tradeoffs necessary to make sure that the entire method, from product architecture to IC packaging, is reliable. This has been a learning process, thus it is taking longer than usual and it is costing more than expected.

The result is that we will have products manufactured with the 20 nm process, but they will cost more to design, develop, and manufacture. End users will have to pay more for the additional functionality, breaking a pattern of more functionality for the same price that has become the norm with electronics products.

Although a first reading all this seems like bad news, it is actually very good news. The reason is that this new development method is a reliable introduction to what will be required to develop and support the 14 nm process. The task of using UV photolithography at the required resolution may in fact be so complicated to finally force the industry to invest in new manufacturing technologies. If not, given the staggering costs of such transformation, new methods that integrate design, development, and manufacturing will have to be developed.

Integrated Verification

The interaction between development, tools providers, and foundries is becoming a requirement. To some extent this is true at 20 nm, but it will be the norm at 14 nm. Verification is the aspects of product development that will see this impact the most. But using traditional verification methods will not be enough. They will generate tasks that are either too complex or too time consuming to be practical.

It is a well known and widely discussed fact that Verification is the most expensive portion of a design project. And, in many respects, it is also the most difficult to achieve correctly. Traditionally we have taken a cops and robbers approach to verification. We employ to different teams, one for development and one for verification. Generally, although not always, junior staff is assigned the verification tasks. This goes to the point that in some companies development is done in the US while verification is done in a developing country in order to lower costs.

Verification must be integrated in the development tool itself. Synopsys, in my opinion, has shown itself to be a visionary when it invested in both Springsoft and EVE, two companies that know a lot about verification. Development tools must become smarter, must understand the sources of errors. Just like formal analysis techniques make use of assertions generated by engineers, tools will have to be able to infer their own set of assertions and build correct by construction circuitry. The identification of errors will be directed by the set of design rules from the foundry, thus every tool will have to use a set of assertions that are foundry specific.

Whether an EDA company chooses to have one configurable tool or a set of foundry specific tools will depend mostly on the run time profile of the tool. The version that takes less memory and executes fastest will be the chosen solution.

Financial consideration may transform emulation boxes into intelligent peripherals that are integrated

into a simulation environment that can handle multiple levels of abstraction. We must achieve a general purpose verification environment that avoids special purpose hardware. The latter is just to costly and has lower profit margins than software.

Stand alone verification tools will continue to exists for a time while the integration process evolves, but the history of EDA has shown that the tendency is toward more intelligent tools. The purpose is not only to make the work of a designer shorter by improving execution speed, but to make the designers more productive by guiding the process in error avoidance. The goal is not to help an engineer create the largest possible amount of errors in the shortest time possible, but to avoid errors in the shortest feasible amount of time.

The 14 nm Team

All development done at 14 nm will involve a team that includes development engineers, tools providers, process and manufacturing engineers. There is no other possible approach. The manufacturing flow and characteristics will dictate what is possible to build.

It will be too costly to design a circuit that cannot be built and try to fix it at the back end. The product must be architected knowing what can be built. That means the choice of third party IP will be restricted, in fact internal re-use may become more expensive, requiring a stricter selection and at times a significant re-design.

As I just wrote, EDA tools will have to be adapted to the specific constraints of the foundry and may be to the methods used by the development team. The key words "design for manufacturing" and "design for yield" will actually mean something more than developing a robust design. They will mean developing a design that can be successfully manufactured with the proprietary dialect of the 14 nm process "spoken" by the specific foundry.

Process development will have to take into account common circuitry and insure that those collections of transistors and connections will execute without generating parasitic effects that inhibit proper execution. The cell library will look more like a collection of macros not primitives. Proximity effects will be important and generate new layout considerations that will impact both timing and power features.

Clearly companies will try to minimize the cost of new hardware by using software as the "tailoring" ingredient in the recipe. EDA companies have already put increased resources in supporting software integration, but more needs to be done. Specifically the way to measure architectural tradeoff results must be improved.

Although each company will still be focused in performing its own mission, the final device will be the product of a consortium of specific companies. Grand alliances will be formed and, I fear, smaller companies will be the one suffering from this. Leading edge technology in EDA will be the province of companies with access to large amount of investments and established business partners. But, it must also be said, the number of electronics vendors who will be able to use the 14 nm process will be limited both by cost and by business reasons. There will be the need to develop a new business model that assures the proper return to EDA vendors, not just electronics companies and foundries. EDA cannot continue to be squeezed in the middle.

Smaller EDA companies will find good business opportunity serving the vast majority of producers who will continue to use 65 to 32 nm processes. Productivity improvements that lower costs at proven process nodes will offer opportunities to both established and new small EDA companies.

From GDSII to Good Silicon: Integrated Yield Analysis Tools are Key

Dr. Zhihong Liu, Executive Chairman, ProPlus Design Solutions, Inc.

While the semiconductor industry's attention was focused on front-end functional verification— and for good reason because it takes up the bulk of a project schedule — process variations at 45nm and lower geometries were becoming unpredictable, but manageable. At 28nm, resources continue to be poured into front-end verification, while silicon yield has reached crisis proportions for development engineers and circuit designers and without much industry fanfare. As a result, going from GDSII to good silicon has become a perilous endeavor. Circuit designers: Try this at your own risk!

Yield is only now getting the attention it deserves. In short order, both the semiconductor and EDA industries need to refocus attention on achieving good silicon, not the aspect of the design flow causing the most problem or the one with the best publicity. Without question, a few key components of said design flow need careful consideration to make yield analysis more reliable.

Let's start with the challenges that impact chip yield — increased random variations and layout-dependent effects — and identify solutions to offer a bridge between different design stages of the design flow, something that isn't well addressed today. A good example would be a modeling strategy to bridge the foundry process and the design completed by a fabless semiconductor company.

Design for Yield (DFY) technology needs a bridge between a design from a fabless semiconductor company and the fabrication facility or foundry. Another bridge must be extended to link modeling to simulation to yield for performance with tradeoff. These bridges must be flexible, able to handle problems at different levels and be able to scale from small to large, from device level to circuit level to wafer level.

This all sounds great and logical, but as a savvy reader of GABEonEDA's Assembling the Future, you may want more details. Process effects must be understood physically and accurately modeled in SPICE models at the outset. These models must be usable later by circuit designers at various circuit design stages, including simulation and verification. As a result, they need accurate statistical models and specialized design tools with high prediction accuracy and superior simulation

performance.

Currently, circuit designers separately run process, voltage and temperature (PVT) corner analysis and Monte Carlo analysis, with overly optimistic or conservative information from the foundry models. In turn, these models may be used incorrectly on a specific application. This disparate approach with selective corner models and Monte Carlo analysis produces limited information about yield prediction and design optimization. The result can be inclusive information on the overall design effects, leaving the circuit designer less than confident about his or her design.

As design teams reassess their needs, they have concluded that good model knowledge and properly applying it in the design flow are important considerations. The solution they want is a transistor-level design strategy to bridge accurate yield prediction and realistic design optimization with performance and yield. Devising such a strategy should include a DFY toolkit for device modeling, parallel SPICE simulation and statistical analysis.

The toolkit would need to manage process variations with accurate SPICE models for process variations, a fast and reliable statistical simulation engine and hardware-validated sampling technologies. Ideally, these three tools would be tightly integrated to ensure greater accuracy and less degradation in simulation performance. The software would determine the efficiency of yield analysis for performance through regular 3-sigma Monte Carlo runs for analog circuit designs and special High Sigma runs, such as 5~6 sigma, or more, if necessary, for memory designs.

In chip design, nothing seems more of a constant than new challenges and greater complexity. At the risk of oversimplification, process variation issues were set aide to focus on front-end functional verification challenges and other, more pressing concerns. It's now time to focus attention toward the back end to ensure GDSII to good silicon through the use of more practical, reliable, faster and integrated statistical yield analysis software toolkit that bridges different design stages.

About Zhihong Liu

Dr. Zhihong Liu is executive chairman of ProPlus Design Solutions, Inc. He was most recently corporate vice president for CSV R&D at Cadence Design Systems Inc. Dr. Liu co-founded BTA Technology Inc. in 1993 and invented BSIMPro, the leading SPICE modeling product. He also served as the president and chief executive officer of BTA Technology Inc. and later Celestry Design Technology Inc., acquired by Cadence in 2003. Dr. Liu holds a Ph.D. degree in Electrical Engineering from the University of Hong Kong and co-developed the industry's first standard model (BSIM3) for IC designs as one of the main contributors at the University of California at Berkeley.