

Evening Panel Discussion*

Monday March 26

7:00pm-8:30pm

The 50-Million Transistor Chip: The Quality Challenge for 2001



Moderator

Rick Merritt

Editor in Chief, EETimes



Organizer

Richard Goering

Managing Editor, EETimes

Description

The panel will examine the core issues that leading edge semiconductor makers face in making quality and time-to-market decisions about how they design, verify and test high-end silicon devices.

One provocative question that can be posed is, "How much verification is enough?" Every design team must come to a point where a chip design is completed and ready for fabrication — but knowing when you've hit that point is very difficult. There are few guidelines or metrics. It's pretty much a subjective, seat-of-the-pants kind of decision, although there are a few tools (such as code coverage tools) that might offer some help.

In the manufacturing test area, a similar question arises — "How much fault coverage is enough?" Just a slight increase in coverage can make a tremendous difference in yield, but those last few percentage points of fault coverage can be painfully difficult to get. A panel representing users, EDA vendors, and semiconductor executives will share their experiences and concerns about verification and test.

Panelists Statements:



Thomas Daniel, VP ASIC Technology, LSI Logic, Milpitas, CA

There are many facets of verification. The designer is concerned about functional verification at the RTL level and timing verification at the gate level after synthesis. In the last few years, concepts such as elaborating behavioral constructs at a higher level or semi-formal verification techniques have been introduced to alleviate the verification bottleneck. Unfortunately, verification is and will be needed because designs cannot be coded right by construction.

The ASIC vendor is concerned about timing and physical verification. Timing verification after place and route is done using static timing analysis (STA) tools, but it requires robust and high quality timing constraint files. Chip designers are still coming up with right methodologies for timing constraints as different tools accept different constructs.

* Reception starts at 6:30pm, Sponsored by EETimes

With the advent of UDSM technologies, instead of asking how much fault coverage is enough, we should be asking which fault models are used to achieve the fault coverage. Single stuck-at fault models are not adequate to address timing defects caused by low voltage, high density, and multi-level metal designs.



Bryan Hoyer, Senior Director of System Level Solutions, Altera, San Jose, CA

As chips become larger and achieve greater complexity and higher integration there is an understandable concern over verification. Today's high tech marketplace no longer purchases on brand loyalty, the first *working* product in the marketplace wins. Many products and even entire companies have devoted themselves to the science of verification, making sure that the design is right. For all of their noble efforts they fail to deal with a fundamental issue, that of making the wrong thing right. Many a product has failed in the marketplace not by being wrong but by being the wrong thing. What tools deal with this in a changing world?

Certainly solutions which are in system re-programmable provide the ability to deal not only with the engineering error that slips through verification, but also with changes in the marketplace.

The other side of the fast time to market coin is product life. Increasingly, products are being designed which recover their costs not only through initial purchase but also by service fees. For these marketplaces maintaining in system flexibility is key to extending product life as well as reducing design risk and verification time.



Chris Malachowsky, VP Engineering, Nvidia, Santa Clara, CA

For NVIDIA, getting a 50M transistor chip right was last years problem. The solution emphasized both tools and methodology while involving diverse groups from across the company. Getting the next 'big one' logically correct will be an even bigger challenge. It will inevitably draw from the approach used last year but will need to evolve in some significant ways.



Janusz Rajski, Chief Scientist, Design for Test, Mentor Graphics, Portland Oregon

Any future solution to high quality manufacturing test should consider its impact on time-to-market, and cost-of-test. Commercial EDA tools measure test quality by fault coverage of several types of fault models, such as stuck-at, transition, Iddq, and path delay faults. Often tests with even a high coverage of stuck-at and Iddq faults are not sufficient to achieve the required quality objectives and subsequently need to be supplemented by functional tests that require expensive testers.

To guarantee a high degree of automation, future solutions to manufacturing test will continue using structural test methodology based on scan. These solutions, however, will require additional fault models to cover new deep submicron defects affecting logic functions and timing.

In addition, manufacturing test, and specifically fault diagnostics, will play a significantly increased role in design verification as the ultimate solution that can find "the truth hidden in silicon" and help debugging the design process. The new solution should also address the soaring cost of test. While moving entirely to structural testing will reduce, or even eliminate, complex and expensive functional testing, the new fault models are likely to push the volume of test data well above one kilo-bit per gate for high-end silicon devices. Even for low cost structural testers this amount of test data could be overwhelming. The new solution should also drastically reduce the volume of test data to reduce the cost of test. A solution with these characteristics will offer high quality structural test, with very short time to market, and very low cost of manufacturing test.




Greg Spirakis, VP Microprocessor Products Group, GM MPG Design Technology, Intel, Santa Clara, CA

In any given design, the designer inherently cannot know what fraction of the design's defects has been caught or missed in the verification effort.

The end-point of verification is therefore usually determined by pragmatic time or resource constraints, rather than theoretical defect targets. The question of "how much verification is enough" is not central to improving design quality.

It is more important instead to ask how the heavy reliance on verification can be reduced. Verification contributes to design quality only by catching design defects after they have been introduced. Verification resources will not keep pace with the increased design complexities of 0.13 micron technologies and beyond. Efforts to maintain or improve design quality must therefore address defect prevention during the design process. This applies to all aspects of the product design, but the leverage and need for progress are greatest in functional/logic design, to match the product functionality afforded by 0.13 micron transistor budgets.

Intel's experience suggests that proactive defect prevention is a practical and an essential complement to its verification efforts. The correct-by-construction methodologies and tools span the entire design spectrum, ranging from timing and signal integrity through high-level logic design.



Tom Williams, Chief Scientist and Director Test, Synopsys, Mountain View, CA

As the number of transistors increases so do the number of possible defect mechanisms increase (e.g. Cross Talk). Tests are created and evaluated against abstractions of defects called faults. The most popular fault model of today is the stuck-at fault model with the high nineties being the typical coverage achieved by test patterns created and evaluated against this model. At this level of coverage the variation caused by the abstraction of defects to stuck-at faults is so large that any incremental stuck-at fault coverage may or may-not correlate well with defect coverage. At this point, additional coverage should be achieved by switching fault models to cover an orthogonal

set of defects to those already covered by the stuck-at faults.

The test industry has researched a number of fault models such as pseudo-stuck-at faults, toggle faults, bridging faults, open faults, transition and path delay faults. The industry is at a stage where the votes are counted (automated counts and manual counts are in) and the second place winner is quite clear. Transition faults are the most practical set of faults that cover delay defects that are not covered by stuck-at faults. In this panel my position is to discuss aspects about process yield, fault coverage, fault models, and the tools to support this position. In addition to quality a fundamental balance must be made with the Cost of Test. The demands of Low Cost Testing will also be discussed.