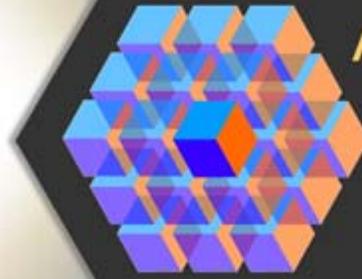


sponsored by

Gabe on EDA • EDAMarket • EDA Confidential



Assembling the Future

A Newsletter About the Design
and Production of Electronics

ISSUE 005 • JUNE 2011

[SUBSCRIBE](#)

[PDF and Archives](#)

[twitter](#)

In this issue:

- [EDA Gets Ready For DAC](#)
by Gabe Moretti
- [Power is a Chip-Level Problem](#)
by Paul van Besouw
- [What's new at DAC? Moving Beyond
Textbook High-Level Synthesis](#)
by Brett Cline
- [Hardware/Software Co-Verification Rules This
DAC](#)
by Lauro Rizzatti
- [Silicon One: The Next Technology Solutions
to Address Semiconductor Companies'
Business Objectives](#)
by Behrooz Zahiri
- [Verification Engineers Develop Triage Plan](#)
by Andreas Veneris
- [New at DAC: A Perl Interface](#)
by Rob Dekker
- [What's New at DAC: Multi-Time
Programmable, Non-Volatile Memory Core in
40nm](#)
by Linh Hong

EDA Gets Ready For DAC

Gabe Moretti

Another DAC is fast approaching and the organizing committee had the task to pull another rabbit from the proverbial top hat in order to lure attendees to San Diego. Hotel prices certainly did not do the trick. It seems that in San Diego the hotel managers have not heard about the recession, or may be they have and figured on DAC attendees to provide needed profits. Either way, room prices, given that DAC has negotiated a block of rooms, are over twice what I will pay at a good hotel with convenient public transportation to the convention center.

So Embedded Software has become the rabbit. And it was about time that EDA recognized the fact that software is an important, and more than often the principal, component of today's systems. Today's chips are truly mainframe computers with sophisticated operating systems and application programs that need to be developed, debugged, verified, and maintained with the same level of efforts we dedicated to mainframes not so long ago.

So I must give credit where credit is due. To the Program Chairs and their committee, to the General Chair, and the Panels Committee for working hard to provide visibility to embedded software problems and solutions.

But the rabbit is already being oversold. The DAC website still needs work, and the use of a good editor is paramount. Paul McLellan, enlisted as the editor of the "DAC Knowledge Center" has no say on what else is being published on the site, and thus has no blame for the hype displayed at <http://www.dac.com/dac+news.aspx?article=241&category=1>.

Six, or almost 40%, of the sixteen companies listed in the Exhibitors ESS zone would have been at DAC whether or not embedded systems had been singled out as a new feature of the conference. And to state that the keynote session "Up Close and Personal with Steve Wozniak" is about embedded systems and software is just an exaggeration. Hopefully Steve, who did not trust any EDA person as his interviewer and demanded to be interviewed by an editor from the San Jose Mercury as his condition to appear at DAC, will talk about systems. Note that the session is sponsored by Magma Design Automation, a company that is not directly involved with embedded software. My problem is that a good thing sells itself, obviously with the help of competent professional that can "spread the word", and does not need to receive accolades that really do not belong to it. If the session was really about embedded systems, then one of the embedded systems company would have sponsored the event, don't you think?

Promising and not delivering will only hurt future conferences. But given that DAC lives for only ten months of the year, future events are of little concern to the present "administration".

Finally, has anyone noticed that ESC20 is taking place in Chicago during DAC? Once again the EDA industry falls victim of parochialism, and the greed to make money at the expense of both users and vendors. People at UBM certainly knew the DAC schedule. So couldn't they schedule ESC20 around DAC? Is this a case of mine is bigger than yours? It is not like UBM is winning the hearts and minds of the electronics community lately, as it continues what seems to be a perennial shuffle of editors.

Seven companies contributed to this month's issue, and all have products or services that are related to the topic of embedded systems and software. Oasys Design Systems writes about power consumption. Software consumes more power than dedicated hardware, but system complexity

requires the use of software, so back to the hardware engineers goes the burden of solving the problem.

Forte Design Systems talks about High Level Synthesis (HLS), a family of tools that makes it possible to simulate at a higher level of abstraction and deal with the hardware/software interactions and then synthesize the hardware description without the need to re-write it. A very powerful tool in the hands of hardware designers. By the way Brett Cline of Forte is the organizer of a DAC Pavilion panel about HLS to be held on Monday June 6 at 11:30AM.

EVE writes about emulation, or to use their term "Hardware Assisted Verification". It is all about being able to simulate the hardware while executing the software in a reasonable amount of time and with a reasonable volume of output.

Magma Design Automation points out that the flow is as important as the individual EDA tool. Integration of tools from different vendors has always been a significant problem for users. Magma thinks it is time to change that.

Vennsa, a Canadian startup, looks at the problem of tracing the origins of software and hardware bugs. It offers its own solution to the problem borrowing the approach from the medical world of the emergency room.

Verific acknowledges the work required by engineers to properly use EDA tools and offers some relief in the form of a tool that makes scripting much simpler.

Finally Kilopass talks about Itera, the industry's first CMOS logic embedded multi-time programmable (MTP) non-volatile memory (NVM) in 40nm. After all one needs a secure storage space to keep its software, right?

Power is a Chip-Level Problem

Paul van Besouw

President and CEO
Oasys Design Systems

It wasn't all that long ago that timing was the toughest design constraint facing designers. Today, power is the big limiter. Power constraints add another dimension of complexity, and traditional synthesis tools can't handle power in ways that are effective for project teams. That's because power is a chip-level problem, not a block-level problem.

Every SoC has a power challenge today. Chips for battery-powered devices such as cell-phones need to dissipate extremely low power when they are not being used, and cannot be wasteful under

any circumstances. Battery life is something that shows through to the consumer and one which dictates buying decisions. Even chips for tethered devices have challenges. Devices for the living room rarely have fans. Server farms are often limited by getting power in and heat out, so even the highest performance chips do not escape the power challenge.

Block-level tools cannot handle chip-level issues. However, power is a problem to be solved at the chip level. Obviously, the power dissipated on an SoC is power generated by all the various blocks. Unfortunately, there is no good way to budget power among blocks, especially when there are complex power architectures involving multiple voltage islands that may only be finalized late in the design cycle.

Battery consumption or heat dissipation, rather than local temperature hot-spots, is a problem confronting SoC designers, more easily managed at the chip level. Dividing the design into blocks each with a power budget in the 28nm environment is an impossible task. Add multiple voltage domains, libraries with multiple power-performance points, libraries with high and low leakage cells and the challenge becomes acute, with complex interactions between performance, dynamic power and static power.

Unquestionably, a change in design methodology is required. The communication of intent by formats such as Common Power Format (CPF) from Si2 and IEEE Standard 1801-2009, based on Accellera's Unified Power Format (UPF) is a start. Chip-level power challenges, however, need to be handled at the chip level and not the block level. Designers need a way to experiment with voltage levels and power tradeoffs at the architectural level for maximum impact, while taking all power measurements from a fully placed netlist.

As we head into DAC, imagine the power headaches we're going to be dealing with inside the San Diego Convention Center. I say it's time to face the power challenge and admit it's a chip-level problem not solvable at the block level. Oasys Design Systems, noted for its Chip Synthesis platform, offers chip-level power analysis and the ability to re-synthesize a design from RTL with power constraints. Demonstrations of RealTime Designer, a tool to manage power at the chip level, will be demonstrated in **Booth #2031** during DAC.

What's new at DAC? Moving Beyond Textbook High-Level Synthesis

Brett Cline

Vice President of Marketing and Sales
Forte Design Systems

Ah, June. Burning coffee, overworked and stressed out marketing people, and far too many

Powerpoint slides. It must be time for DAC!

Over the last 10 years at DAC, we've witnessed the growth of electronic system level (ESL) design and verification to where it now makes up more than 10% of the worldwide CAE revenue. At the heart of ESL is high-level synthesis (HLS), which promises designers the ability to take highly abstracted models and automatically turn them into register transfer level (RTL) design implementations.

It further claims to improve designer productivity, quality of results ... and may even have you have you golfing nine months out of the year if you use it, at least according to some companies. Let's face it, there are a number of tools available and the bold claims in the marketing material can confuse things enough where it's difficult to tell what's true and what's simply hype.

At this point, a number of tools on the market can do the basics found in any HLS textbook, making it *seem* like they are viable solutions for your production design flow. Are they? Only a few of these tools have really been around the block and have documented successes and tapeouts from the leading systems and semiconductor companies around the world. There are some key features that differentiate the leaders and, with a couple of pointers, you can quickly become versed in what it takes to find the leading products.

What's new this year at DAC, however, is that there's a real focus on getting beyond textbook HLS to the practical features needed for hardcore behavioral design with HLS. You'll find a focus on:

- The abstraction level of the input code
- The deep feature set to handle any type of design
- Getting to area, power, and performance numbers that beat hand coded RTL and the supporting technologies for ESL flow automation, verification, 3rd party integrations, and so on.

Of course there are many other product features that seem (and may actually be) useful, but in a 60-minute presentation, there is only so many things you can dig into.

And, while the language debate has been largely decided with the leading tools supporting SystemC and untimed C++ code, there are still some stragglers pushing their own odd languages and the "benefits" of going with a non-standard language or proprietary extensions.

As always, if you'd like to get the truth about HLS, we're happy to serve it up at the Forte **booth #3417**. And, on Tuesday from 4-6pm, we'll also be serving up some ice-cold beer for cocktail hour. We'll help you cut through the hype and understand the basics of high-level synthesis and demonstrate the advanced features in Cynthesizer for real production success. And, while there's no burnt coffee here, you might see a Powerpoint or two and bagpipers. Yes, Forte will close DAC with a bagpiper's performance Wednesday, June 8, at 5:50 p.m. in **Booth #3417**. It is DAC after all.

Hardware/Software Co-Verification Rules This DAC

Lauro Rizzatti

General Manager of EVE-USA

DAC's emphasis this year seems to be on embedded software, an area of design and verification well suited for today's hardware-assisted verification platforms and a term that's been part of the verification lexicon for many years now. In fact, a variety of hardware/software co-verification tools have been available to design teams for at least five years and in used to great success.

But first, a definition: Hardware-assisted verification platforms combine traditional emulation and rapid prototyping systems into one environment for ASIC and SoC debugging and embedded software validation ... or hardware/software co-verification. They have a means of concurrently debugging hardware and embedded software, and are good at handling large, complex designs in excess of hundred million ASIC-equivalent gates that are now consuming upwards of one-trillion verification cycles per month.

Think of it! One-trillion verification cycles per month to validate that a hundred-million gate SoC design will work as intended. It wasn't all that long ago that we were talking billions of cycles and that seemed incomprehensible. Even so, thorough and exhaustive functional verification at this stage continues to be the most cost-effective debugging approach available. And, hardware-assisted verification platforms differentiate themselves with high performance, due in large measure to users' demand for performance that, in turn, is driven by software requirements.

Driven by the massive complexity of today designs, it is becoming a necessary replacement for the venerable event-based HDL simulator, a 35-year old verification tool often thought of as the industry's mainstay that's fast running out of juice. In fact, hardware-assisted verification supersedes large simulation farms not adapt at processing very long sequences of verification cycles inherent in the execution of embedded software. Further, it adds the significant benefit of saving enormous quantities of electrical power by lowering direct consumption and alleviating air cooling requirements.

Hardware-assisted verification allows engineering teams to plan more strategically and implement a debugging approach based on multiple abstraction levels. Teams can track concurrently a bug between the hardware and embedded software to identify where the problem lies.

The advent of transaction level modeling (TLM) and availability of transactors can turn hardware-assisted verification tools into virtual platform test environments for a range of vertical markets. A transactor, part of verification IP, is a high-level abstraction model of a peripheral function or protocol. Transactors, often provided as off-the-shelf IP, are available for a variety of different protocols. A typical catalog includes PCIe, USB, FireWire, Ethernet, Digital Video, RGB, HDMI, I2C, UART and JTAG components.

With DAC's program focused on embedded software, I predict hardware-software co-verification and other embedded software application tools will rule the DAC exhibit floor. EVE will be at DAC in **Booth #2836**. Don't miss an opportunity to meet the EVE team, here about our hardware/software co-verification solutions and see a demonstration of ZeBu, a cost-effective hardware-assisted verification platform.

Silicon One: The Next Technology Solutions to Address Semiconductor Companies' Business Objectives

Behrooz Zahiri

Vice President of Business Development
Magma Design Automation Inc.

The semiconductor industry has overcome a range of technology issues and market pressures to reach a level of innovation and success that resulted in ... even greater issues and pressures. EDA can no longer meet these by producing better tools alone, but needs to address the full range of technological requirements of semiconductor companies and reconcile them with their business constraints.

Let's look at the Semiconductor industry. Chip makers can be viewed as victims of their own success. Semiconductor manufacturers have delivered more functionality, along with better and faster chips at an increasing rate. Visionary applications of these chips have resulted in new products that do more, which in turn, results in demand for any number of things.

Lifestyle-driven technology and new products must continue to do even more. As always, the availability of the "next great thing" drives demand for more great things, such as more connectivity, more performance and more integration.

Silicon technology continues to advance at a fast clip to accommodate the "better faster cheaper" phenomena of a market created by lifestyle enhancements. Integration of more functionality at lower cost and higher performance has led the convergence of silicon to mixed signal. From increasing analog IP in SoC designs, to emergence of more digital content in analog/mixed-signal chips to complete memory subsystems, silicon is emerging as the next great thing.

The "next great thing" is an extremely competitive market of short market windows, aggressive pricing and similar functionalities. It's a furious race to get to market first for margins, cost reduction across the entire design cycle and silicon manufacturing and product differentiation in a market flooded with the same IP, manufacturing capabilities, and commodity functions.

All this calls for the need for better design tools, but it's not enough. No semiconductor company in the world can hope to achieve its goals if it is armed with only the latest software from an EDA vendor. No one EDA vendor can supply 100 percent of what any semiconductor company needs.

Software is just part of the solution. EDA companies must recognize the range of requirements that semiconductor companies face and innovate technologies accordingly. EDA needs to help them assemble the necessary differentiated solutions within their ever-tightening time windows and cost requirements.

Magma is using DAC to unveil Silicon One, a means of bringing together superior technology, design expertise and creativity to deliver profit-driving differentiated silicon. It includes new innovations to deliver fast, integrated solutions for design team and a way to collaborate with others. Visit Magma's **booth #1743** to learn more.

Verification Engineers Develop Triage Plan

Andreas Veneris

President and Chief Executive Officer
Vennsa Technologies Inc.

As the verification community struggles with the manual error localization and debugging effort that consumes more than 60% of the verification effort, it advocates the use of new automated tools and methodologies to slash this productivity burden. One new tool encompasses the community's feedback and addresses a wide set of debugging applications native to modern verification flows.

Among its features is a powerful "triage" engine. Triage is a term that originates in the medical community referring to the quick diagnostic process performed when patients enter the emergency rooms. Based on the triage, medical professionals determine the urgency of the situation and the doctor type to examine the patient.

An analogous process exists in verification where engineers must quickly diagnose simulation failures at a high level to understand the cause of the failure and identify which engineers should further analyze and fix them. Today, failure triage is performed in an ad hoc manner based only on the messages available in the simulation log files. Design teams waste valuable resources as problems are handed to the wrong engineers, failures are incorrectly dismissed as duplicates or duplicate bugs are treated as distinct bugs.

Triage is hard because it is a Catch-22 problem: In order to do accurate triage, the verification engineer needs to do low-level debug to determine where the bug is. But he or she cannot do low-level debug effectively without good triage.

With a tool to automate triage, the debugging effort will run much faster and more efficiently. OnPoint from Vennsa is the first such verification tool with a triage engine that can differentiate automatically between the sources of failures. The process starts with a diagnostic step that performs automated root cause analysis to generate a list of bug sources, called suspects. To continue the medical analogy, it automatically screens the failures and generates an "X-ray" picture of what is going on in the design and where the error source is originating from.

This X-ray is used to do intelligent binning. Suspects are analyzed by the embedded triage engine

to generate a signature based on the error source. Signatures determine which failures contain similar or distinct bug sources. The end result of the triage process is a set of bins, where each bin contains failures with the same root cause. This allows engineers to simply assign a bin to an engineer to make a fix. The entire process can be naturally automated within nightly regression failures to accelerate the debug and triage tasks.

Verification engineers at DAC don't need medical gear to perform triage on their designs, but they do indeed need a good diagnostic tool. I encourage them to stop by the Vennsa Technologies' DAC **Booth #2912** for a demonstration of the first automated functional verification debugging and error localization tool performing triage.

New at DAC: A Perl Interface

Rob Dekker

Founder and CTO

Verific Design Automation, Inc.

EDA and FPGA companies, as well as semiconductor corporations with in-house CAD development, have invested many hours in developing parsers and elaborators for VHDL and Verilog over the years. Many of them were ready to give up when they faced the daunting task of supporting SystemVerilog on top of that.

Fortunately, by that time affordable off-the-shelf parsers and elaborators, written and maintained by (System)Verilog and VHDL parser specialists had become available. It takes the burden of full IEEE language support of EDA developers, although still requiring a fair knowledge of C++ to interface with the commercial parser and elaborator packages.

At the same time, a different class of engineers working in the semiconductor industry also had made occasional attempts in writing parsers to solve a specific problem. Such in-house parsers tended to be single-problem solutions that withered rapidly once a project was finished, lacking maintenance and full IEEE standard support.

This group of engineers is not necessarily C++ programmers. Most of them are proficient in csh, Tcl, Perl, or even Python, and would much rather put a solution together using a scripting language. If they encounter a project for which a language parser is required, they prefer to use their scripting skills instead of brushing up their C++.

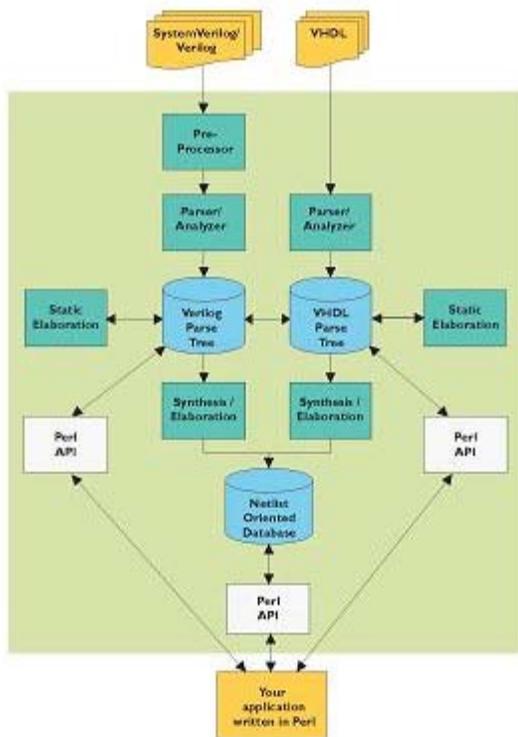
To accommodate this, one could imagine creating a Perl command for each C++ API inside the available parser / elaborator packages, but that sounds like a lot of work. Verific's software, for instance, contains well over 2,000 APIs.

As it turns out, that is not really necessary. The good news is that the majority of these APIs can be developed with relatively little effort through the use of SWIG (www.swig.com). SWIG is an interface generator that connects programs written in C and C++ with scripting languages such as Perl. It works by taking the declarations found in C/C++ header files and using them to generate the wrapper code that scripting languages need to access the underlying C/C++ code. The Simplified Wrapper and Interface Generator (SWIG) is free software and the code that SWIG generates is compatible with both commercial and non-commercial projects.

At Verific, we undertook this exercise and created a Perl interface to our industry standard (System)Verilog and VHDL parsers, analyzers and elaborators (figure 1). Most of Verific's APIs translate rather easily using SWIG. But, we also found plenty of areas where we needed to step in. For instance, SWIG does not support Perl callbacks. Once we took care of the anomalies, we ended up with a comprehensive Verific Perl package.

It should be noted that this is a Perl API to a system defined in C++. All the underlying objects are written in C++ and there is a strong correlation between Perl and C++ APIs.

Verific will demo the SystemVerilog and VHDL Perl APIs in **Booth #2733** at DAC June 6-8.



[Figure 1]

What's New at DAC: Multi-Time Programmable, Non-Volatile Memory Core in 40nm

Linh Hong

Vice President of Marketing
Kilopass Technology Inc.

This year's Design Automation Conference is coping with the transition from 65nm to 40nm as the mainstream, high-volume silicon production process with 28nm beginning to ramp at leading edge foundries such as TSMC and GLOBALFOUNDRIES.

According to Michael Splinter, president and CEO of Applied Materials, smart phones and tablet performance demands are driving the increased production at 40nm and below. For an embedded non-volatile memory IP company, like Kilopass Technology, this represents major upside for its latest product offering Itera, the industry first and only multi-time programmable (MTP), non-volatile memory (NVM) in 40-nm.

Built using any standard logic CMOS process, this new memory core employs antifuse technology to record data. Unlike a metal fuse that burns away metal to store a bit of data, antifuse technology causes a breakdown in the gate-oxide of a standard CMOS transistor, which converts the transistor in a quiescent off-state from an open to a short. Unlike metal fuses in which the debris created when the fuse was blown can over time re-grow and reconnect an open circuit, the antifuse oxide breakdown is permanent.

How is it possible to create an MTP memory using OTP technology? The answer lies in increasingly smaller process geometries making it economically feasible to overpopulate a memory and simply discard sectors that would ordinarily be overwritten and replace them with blanks sectors. The memory core can be reprogrammed up to 1,000 times and can accommodate a storage capacity of up to a 1-Mbit.

The advantage that comes with storing program code on-chip is the increase in performance and reduction in system costs that result. Without it, SoC vendors would have to store program code in an off-chip EEPROM or flash memory chip. Connecting this off-chip memory to an embedded processor on board the SoC means using a four-wide serial interface bus. By having the program data on chip in a memory core, SoC designers can achieve bus interface speeds 24-times faster than the off-chip memory, thus enabling the on-chip CPU to execute directly from it. Eliminating the off-chip memory device from the bill of material cost can be significant for large volume production. For example, an SoC in production for three years and shipping 10 million units a year will achieve a \$6 million savings by eliminating this one component.

To eliminate the challenge of integrating it into an SoC design, the new MTP memory comes with the Open Core Protocol interface. This makes the memory core look to the computing devices accessing the memory like a standard re-programmable memory. Finally, it comes with built-in self-test (BIST) to eliminate the need for ATE testing at wafer and package. The memory tests itself.

Need more information? Visit ChipEstimate.com's **Booth (#1731)** or the GLOBALFOUNDRIES' **Booth (#1517)** during the 48th Design Automation Conference to learn more about Kilopass and Itera.

