

Chip Design

Winter 2015

Tools, Technologies & Methodologies

“SNEAK PATH” BREAKTHROUGH

p.8

Hot Trends for 2015
p.12

New Markets for EDA
p.14

**Pushing the
Performance
Boundaries of ARM
Cortex-M Processors**
p.22

www.chipdesignmag.com

Sponsors:



Seeing What Others Can't

The Key to Unlocking New Insights



Mark Wallace
Vice President and General Manager
Keysight Technologies, Inc.

You've known us as Hewlett-Packard, Agilent Technologies and, now, Keysight Technologies. For more than 75 years we have been helping you unlock measurement insights.

There have always been two sides to the story. One is the work we do, creating innovative instrumentation and software. The other is the work you do: design, develop, debug, troubleshoot, manufacture, test, install and maintain components, devices and systems.

Those seemingly separate activities are connected by something profound: the "A-ha!" that comes with a moment of insight. When those happen for us, the results are innovations that enable breakthroughs for you.

Enabling the right idea at the right time

This is our legacy. Keysight is a company built on a history of firsts, dating back to the days when Bill Hewlett and Dave Packard worked in the garage on 367 Addison Avenue in Palo

Alto, California. Our firsts began with U.S. patent number 2,268,872 for a "variable-frequency oscillation generator." Appropriately, the centerpiece of Bill's design was a light bulb, which is often used to symbolize a new idea.

Our future depends on your success, and our vision is simple: by helping engineers find the right idea at the right time, we enable them to bring next-generation technologies to their customers — faster.

Offering expertise you can leverage

This is happening in aerospace and defense applications where increasingly realistic signal simulations are accelerating the development of advanced systems that protect those who go in harm's way. It's happening in research labs where our tools help turn scientific discovery into the discovery of new sciences. It's taking place with DDR memory, where our line of end-to-end solutions ranges from simulation software to protocol-analysis hardware. And in wireless communications we're providing leading-edge measurement tools and sophisticated, future-friendly software that support the development and deployment of LTE-Advanced.

Within those systems, there are more standards than a single engineer can keep up with. That's why so many of our engineers are involved in standards bodies around the world. We're helping shape those standards while creating the tools needed to meet the toughest performance goals.

Enabling your next breakthrough

To help Keysight customers continue to open new doors, we're concentrating our effort and experience on what comes next in test and measurement. Our unique combination of hardware, software and people will help enable your next "A-ha!" moment, whether you're working on mobile devices, cloud computing, semiconductors, renewable energy, or the latest glimmer in your imagination. Keysight is here to help you see what others can't — and then make it reality.





A NEW DIMENSION IN TRANSLATION FOR TEST

POWERFUL TRANSLATIONS | LONG HISTORY OF SUCCESS | COST EFFECTIVE | CUSTOMER DRIVEN

SOURCE III ACHIEVES A QUANTUM LEAP IN VECTOR TRANSLATION TOOLS

Our products combine methods, software and efficiencies to bring translation for ATE into new dimensions—providing efficient Translation to Link Simulation and ATPG Test Patterns!

VTRAN

is the most powerful vector translation solution available. The new "VUI" Visual User Interface makes it more powerful and intuitive than ever and integrates all three products under one easy-to-use interface.

DFTView

Conceptualizes test vectors with timing-accurate waveform visualization of popular test programs in WGL or STIL; and VTRAN-generated ATE. Allows understanding of test patterns' impact on the device across all signals

VCAP

Vector comparison and analysis pushes VTRAN into new dimensions, especially when used with the new VUI visual user interface to support vcd/evcd translations to ATE

Contact us today for free evaluation ►



SOURCE III

CUSTOMER DRIVEN TEST SOLUTIONS

SOURCEiii.COM | 916 941 9403

Exponentially Rising Costs Will Bring Changes



Scott McGregor, President and CEO of Broadcom, sees some major changes for the semiconductor industry moving forward, brought about by rising design and manufacturing costs.

Speaking at the SEMI Industry Strategy Symposium (ISS) in January, McGregor said the cost per transistor was rising after the 28nm, which he described as “one of the most significant challenges we as an industry have faced.”

He said that in the past, it was a “no brainer” for a design company to move its entire set of products always forward to the next generation. “Every generation would be better than the previous one. It would be faster, it would be lower power, it would be more cost effective,” he said. “We think we’re now seeing this come to a bottom.” The reason for increasing transistor cost is the complexity of the devices, and the cost of the equipment required to produce them. These costs are going up exponentially, McGregor said.

Chip design cost is also increasing exponentially. McGregor showed a chart with dramatically increasing cost for each process node for such things as software, prototyping, validation, verification and IP qualification per process node.

McGregor also pointed out that the semiconductor industry as a whole is maturing. He said we have moved from a new market phase with double digit growth into an evolving market with high single digit growth to now a stabilizing market with mid-single digit growth year-on-year.

Although mature, the industry will still see some volatility, although less than in the past. “supply is easily overridden and it takes a long time to build some of these devices like scanners so that’s going to create volatility,” he said.

He also predicted that SoCs would become even more pervasive. He used a set top boxes as an example, noting that they used to be full of boards crammed with lots of discrete parts. “Today, if you open up a set top box, you’ll see a relatively small circuit board inside with a large chip that integrates almost all the

functionality,” he said. “The box is still the same size because consumers perceive value in the size of the box, but it’s mostly air inside.”

McGregor said the tapeout costs to do a single device are very high. “You have to put \$100 million into a semiconductor startup today to be able to get to productization.” This means that big companies will be getting bigger. “There will still be some small companies – but I think the mid-sized company in our industry, in devices, is going to dramatically go away because of the scale and other things required,” he said.

A big impact of these changes is that the process node selection is going to change. “Instead of immediately going to the next node, you’re going to stay in nodes longer. That means, for example, that 28nm is going to be a very long-lived node. There are a lot of things that probably will not make sense to move beyond 28nm for a long time. It will not automatically mean you should go to 16 or 14nm, or 10nm. There will be relatively few devices that economically make sense to do that,” he said.

He noted that Broadcom crossed the threshold where software engineers outnumbered hardware engineers a number of years ago and now has “significantly more” software engineers than hardware engineers. “That’s an interesting transition because we’re now delivering systems instead of just chips. The value just doesn’t come from the transistors -- It comes from all the other pieces put together. One of the challenges for us as an industry is getting paid for that. Unfortunately, for many of us, the software is the ‘gift wrap’ for the chip rather than something we can monetize,” he said. ♦

Pete Singer is the Editor-in-Chief of Chip Design, Semiconductor Manufacturing & Design (SemiMD.com) and Solid State Technology. He is also the conference chair of The ConFab, an annual networking event and conference focused on the economics of semiconductor manufacturing.



Why Reinvent The Wheel?

Maximize the reuse of your IP Cores with Open Core Protocol, the only open, industry-wide socket standard. Utilizing OCP-IP's free infrastructure eliminates the need to design, document, train and evolve a proprietary standard and support tools. This frees up valuable resources and maximizes ROI.



The OCP Specification describes a robust, continuously updated, IP core interface, or socket, that is supported and built upon by the biggest names in the industry. The many benefits of OCP-IP membership include support, tools, training, marketing opportunities and much more.

Visit our Web site today and see how membership with industry leaders can save you time, money and resources. www.ocpip.org

IN THIS ISSUE

Cover Story

8

Focus Report—"Sneak Path" Breakthrough Heralds Arrival of Ultra-High Density Resistive Memory

By Sung Hyun Jo, Ph.D., Senior Fellow,
Advanced Technology Development, Crossbar,
Inc.

Features

12

Hot Trends for 2015

The new system-design imperative

By Chi-Ping Hsu, Senior Vice President, Chief
Strategy Officer, EDA and Chief of Staff to the
CEO at Cadence

14

The Coming Year in EDA: What Will Shape 2015

By Gabe Moretti, Senior Editor

22

Pushing the Performance Boundaries of ARM Cortex-M Processors for Future Embedded Design

By Ravi Andrew and Madhuparna Datta,
Cadence Design Systems

27

IoT Cookbook: Analog and Digital Fusion Bus Recipe

Experts from ARM, Mathworks, Cadence,
Synopsys, Analog Devices, Atrienta, Hillcrest
Labs and STMicroelectronics cook up ways to
integrate analog with IoT buses.

By John Blyler, JB Systems

Departments

2 Editor's Note—Exponentially Rising Costs Will Bring Changes

By Pete Singer, Editor-in-Chief

6 In the News—Product News

32 Viewpoint—From Data to Information to Knowledge

By Gabe Moretti, Senior EDA-IP
Editor

METHODS/EDA TOOLS


30 Source III, Inc.

Industry Leading Tools Linking
Simulation and ATPG to Test

SEMICONDUCTOR TECHNOLOGIES

31 True Circuits, Inc.

PLL and DLL Hard Macros



Cover image courtesy of Crossbar

READ ONLINE
www.chipdesignmag.com

www.chipdesignmag.com

EDITORIAL STAFF

Editor-in-Chief

Pete Singer
psinger@extensionmedia.com

Senior Editor

Gabe Moretti

Web Editor

Shannon Davis

Embedded Editor

Chris Ciuffo, Anne Fisher

Contributing Editors

Caroline Hayes, Cheryl Ajluni, Dave Bursky, Brian Fuller,
David Lammers, Craig Szydlowski, Ed Korczynski, Jeff Dorsch

Editorial Board

Thomas L. Anderson, Vice President of Marketing, Breker
Verification Systems, Inc. • Cheryl Ajluni, Technical
Consultant, Custom Media Solutions • Karen Bartleson,
Senior Director, Corporate Programs and Initiatives, Synopsys
• Chuck Byers, Director Communications, TSMC • Lisa
Hebert, PR Manager, Agilent • Tom Moxon, Consultant,
Moxon Design • Walter Ng, Senior Director, Design Services,
Chartered Semiconductor • Steve Schulz, President, S12

CREATIVE/PRODUCTION/ONLINE

Production Manager

Spryde Heithecker

Media Coordinator

Yishian Yao

Graphic Designers

Nicky Jacobson & Caldin Seides

Senior Web Developer

Mariam Moattari

SALES STAFF

Advertising and Reprints

Jenna Johnson
Sales Director
jjohnson@extensionmedia.com
(612) 598-3446

Michael Cloward

Sales Manager

mcloward@extensionmedia.com

Audience Development / Circulation
Jenna Johnson • jjohnson@extensionmedia.com

Extension MEDIA

President

Vince Ridley (415) 255-0390 x18
vr Ridley@extensionmedia.com

Vice President, Business Development
Melissa Sterling • msterling@extensionmedia.com

Vice President, Sales
Embedded Systems Media Group
Clair Bright • cbright@extensionmedia.com

TO SUBSCRIBE OR UPDATE YOUR PROFILE
www.chipdesignmag.com/subscribe

SPECIAL THANKS TO OUR SPONSORS



Chip Design is sent free to design engineers and engineering managers in
the U.S. and Canada developing advanced semiconductor designs. Price for
international subscriptions is \$125, US is \$95 and Canada is \$105.

Chip Design is published bimonthly by Extension Media LLC, 1786 18th
Street, San Francisco, CA 94107. Copyright © 2015 by Extension Media
LLC. All rights reserved. Printed in the U.S.

Receive a Complimentary EDA & IP Industry Market Report

Is your company a missing piece in the EDA & IP industry?

The EDA Consortium Market Statistics Service (MSS) is the timeliest and most detailed market data available on EDA & IP. Top companies have been participating in the MSS since 1996, but no report is complete without ALL the pieces.

If your company is not contributing to the MSS, it is a missing piece in the puzzle. Contribute your company's data to the MSS and receive a complimentary Annual MSS Executive Summary.

What's in it for your company?

- Worldwide statistics on EDA & IP revenue
- Ability to size your market
- Six years of market data for trends analysis
- The opportunity to make a positive difference in the industry

For details contact:

mss@edac.org

(408) 287-3322

www.edac.org



Product News

MENTOR GRAPHICS ACQUIRES FLEXRAS TECHNOLOGIES

Mentor Graphics Corp. acquired Flexras Technologies, a developer of proprietary technologies that reduce time required for prototyping, validation, and debug of integrated circuits (ICs) and systems on chip (SoCs). Flexras timing-driven partitioning technology will expand and strengthen the portfolio of tools available from Mentor to help engineers overcome the challenges of increasingly complex design prototyping. Terms of the deal were not disclosed.

Flexras Technologies was founded in August 2009 by a group of CAD and FPGA experts, established to build upon the knowledge from ten years of research at the University of Pierre et Marie Curie and LIP6 Lab. Flexras's products solve hardware/software validation bottlenecks by complementing rapid prototyping platforms with timing-driven partitioning, offering a material increase in clock frequency for highly complex design prototyping.

LATEST TENSILICA PROCESSORS DELIVER UP TO 75% MEMORY POWER AND AREA SAVINGS

Cadence Design Systems, Inc. announced the 11th generation of the Tensilica® Xtensa® processors. The new Xtensa LX6 and Xtensa 11 processors enable users to create innovative custom processor instruction sets with up to 25 percent less processor logic power consumption and up to 75 percent better local memory area and power efficiency.

The new Xtensa 11 and Xtensa LX6 processors feature several architectural improvements, including:

- Enhancements in flexible length instruction extensions (FLIX) for Xtensa LX6 that allow for very long instruction word (VLIW) instructions of any length from 4 to 16 bytes, resulting in code size savings of up to 25 percent compared to prior Xtensa versions, thus enabling local memory and cache size reductions of up to 25 percent for the same performance level.
- An option for run-time power-down of portions of cache memories, yielding up to 75 percent local memory power savings in select operating scenarios with dynamic cache-way control.
- More efficient data cache block prefetch lowers system power and boosts system performance by speeding functions such as

MemCpy by 6.5 times faster and reducing the total number of system bus read operations by up to 23 percent.

- Reduced dynamic switching power of the processor logic gates by up to 25 percent.

MENTOR GRAPHICS ANNOUNCES NEW VERIFICATION IP FOR PCIE 4.0

Mentor Graphics Corp. announced the availability of its new Mentor® EZ-VIP PCI Express Verification IP. The new Verification IP (VIP) reduces testbench assembly time for ASIC (application-specific integrated circuit) and FPGA (field-programmable gate array) design verification by a factor of up to 10X.

Verification IP is intended to help engineers reduce the time spent building testbenches by providing re-usable building blocks for common protocols and architectures. However, even standard protocols and common architectures can be configured and implemented differently from design to design. As a result, traditional VIP components can take days, or even weeks, to prepare for a simulation or emulation testbench.

GLOBALFOUNDRIES AND LINEAR DIMENSIONS TO OFFER JOINT ANALOG SOLUTION FOR WEARABLES AND MEMS SENSORS MARKETS

GLOBALFOUNDRIES, a provider of advanced semiconductor manufacturing technology, and Linear Dimensions Semiconductor Inc., a semiconductor company specializing in low power analog and mixed signal integrated circuits, today announced that they are working together to manufacture a 14-channel programmable reference from Linear Dimensions for multiple markets including IoT (Internet of Things) sensor and wearable device applications.

The LND1114 is a 14-channel reference designed to meet the tuning needs of emerging IoT sensors and Wearable applications. The LND1114 is available in QFN-3x2.2mm form factor, and is the world's smallest programmable multi-channel reference product. With a typical drift of only 13uV after 10 years at 70C, low temperature drift and an initial accuracy of 0.2%, the LND1114 is ideally suited for precision sensor biasing.

SYNOPSYS SHIPS OVER 5000 HAPS PROTOTYPING SYSTEMS

Synopsys, Inc. shipped over 5,000 HAPS® FPGA-based prototyping systems to more than 400 companies. These companies have selected HAPS systems to accelerate software development, hardware/software integration and system validation. Prototyping

teams seeking the highest performance ASIC prototypes select HAPS systems for the scalable architecture, integrated prototyping software and rich catalog of real world I/O interfaces.

CADENCE ANNOUNCES FOURTH GENERATION TENSILICA HIFI DSP ARCHITECTURE

Cadence Design Systems, Inc. announced the Cadence® Tensilica® HiFi 4 audio/voice digital signal processor (DSP) intellectual property (IP) core for system-on-chip (SoC) designs, which offers the industry's highest performance licensable digital signal processing (DSP) core for 32-bit audio/voice processing. This fourth generation HiFi architecture enables emerging multi-channel object-based audio standards and offers 2X the performance versus the HiFi 3 DSP, making it ideal for DSP intensive applications including digital TV, set-top box (STB), Blu-ray Disc and automotive infotainment.

With emerging object-based audio standards, individual sounds become objects that can be placed anywhere in a room. Instead of being mixed at a recording studio, these sounds are mixed on the fly at each location where they are played back. The object-based audio system calculates where the sound should emanate so that it appears in roughly the same space, no matter where speakers are located. This provides a more natural and immersive sound experience, and requires much more DSP processing power. Instead of having to use multiple DSPs to accomplish this, designers can now get the DSP power they need from one core – the HiFi 4 DSP.

ON SEMICONDUCTOR DEMONSTRATES HIGH EFFICIENCY 3D SENSOR STACKING TECHNOLOGY

ON Semiconductor has successfully characterized and demonstrated its first fully-functional stacked CMOS imaging sensor featuring a smaller die footprint, higher pixel performance and better power consumption compared to traditional monolithic non-stacked designs. The technology has been successfully implemented and characterized on a test chip with 1.1 µm pixels and will be introduced in a product later this year.

NEW LPDDR4 VERIFICATION IP FROM SYNOPSYS

Synopsys, Inc. announced the availability of verification IP (VIP) for LPDDR4. Synopsys VIP for LPDDR4 is based on a 100 percent native SystemVerilog Universal Verification Methodology (UVM) architecture to enable ease of use, ease of integration and performance. Complete with verification plans, built-in coverage and a protocol-aware memory debug environment, Verdi® Protocol Analyzer, Synopsys VIP for LPDDR4 is a complete VIP solution that accelerates verification closure for designers of low power memory controllers and systems on chips (SoCs).

HLMC DEVELOPS SPECIALTY TECHNOLOGY IN PURSUIT OF GROWING IOT MARKET

Chinese IC manufacturer Shanghai Huali Microelectronics Corporation gave a presentation on its outlook for the Internet of Things (IoT) market and the wide application of its specialty technology at the 2014 China Semiconductor Industry Association IC Design Branch Annual Conference (“ICCAD”), which was held at Hong Kong Science Park.

Henry Liu, senior director of marketing at HLMC, said, “With the development of smart automotive, smart grid, smart home and smart medical services, among other sectors, coupled with the pursuit among the general population of a simpler lifestyle and more efficient management of one’s day to day affairs, IoT has become the new hot topic of the market. The development of the market is set to further promote the prosperity of the semiconductor industry as semiconductor components are the basic core and data gateway of IoT equipment.”

According to Cisco IBSG, IoT connections worldwide are expected to reach 50 billion units, a milestone that is expected to have a profound impact on both consumers and vendors around the world. Currently, many of the world’s leading IC producers are accelerating expansion into the IoT sector in preparation for building their own ecosystem.

As one of the most advanced 12-inch wafer foundries in mainland China, HLMC’s technology starts from 55nm technology node and mainly covers 55nm LP, 40nm LP and 28nm LP as well as 55nm HV, 55nm eFlash and specialty technology. HLMC provides customers with low-cost wafer foundry solutions.

PLDA AND M31 TECHNOLOGY COMBINE PCIE 3.0 CONTROLLER AND PHY IP

PLDA, the company that designs and sells intellectual property (IP) cores and prototyping tools, and M31 Technology, a global silicon intellectual property (SIP) provider, have developed a comprehensive controller-plus-PHY solution for ASIC design projects. The combined solution -- PLDA’s PCIe 3.0 controller in Gen2 XpressRICH3 with M31 Technology’s PHY IP – received certification from PCI-SIG which was validated on PLDA Kintex-7-based platform XpressK7 and M31 Technology’s daughter card. The complete solution is optimized for storage applications.

“Sneak Path” Breakthrough Heralds Arrival of Ultra-High Density Resistive Memory

As floating-gate flash memory technologies used in the majority of products on the market today quickly approach the limits of their ability to scale to higher densities, it has become widely recognized that a new non-volatile memory technology is needed to replace them. Resistive random-access memory (RRAM) is widely hailed as the “most likely to succeed” in the race to replace today’s flash memory with a new, more scalable, higher-capacity, higher-performance and more reliable non-volatile memory. While many companies are actively pursuing RRAM technologies, however, the road to creating commercially viable RRAM products has not been easy.

One of the greatest challenges facing developers in achieving ultra-high density RRAM (>1Tb) has been overcoming the leakage (sneak) current problem in crossbar arrays that interferes with the reliable reading of data from individual memory cells and increases power consumption. Three dimensional (3D) stackable 1TnR (1 Transistor per “n” Resistive memory cells) or 1S1R (1 Selector per 1 Resistive memory cell) crossbar structures are necessary to achieve these density levels, and various selector devices, such as tunneling diodes, bidirectional varistors, and ovonic threshold switches, have been tested as solutions to the sneak path issue with only limited success.

Among the key requirements for a suitable selector are high selectivity ($\Delta I @VR, 1/2VR$), a steep turn-on slope, high current density, fast turn-on and recovery, and high endurance. Previous attempts at developing suitable selector devices have been reported with selectivity ranging from 150 to about 10^5 , but circuit simulations have shown that a selectivity far larger than 10^5 is required to design megabit-level passive crossbar arrays.

Engineers at Crossbar, Inc., a start-up company developing 3D-stackable RRAM technology, have solved the sneak path problem with the development of the Field Assisted Superlinear Threshold (FAST) selector, a new design which has demonstrated the highest reported selectivity of 10^{10} , as well as an extremely sharp turn-on slope of less than 5mV/dec, fast turn-on and recovery (<50ns), an endurance greater than 100M cycles, and a processing temperature less than 300°C ensuring commercial viability.

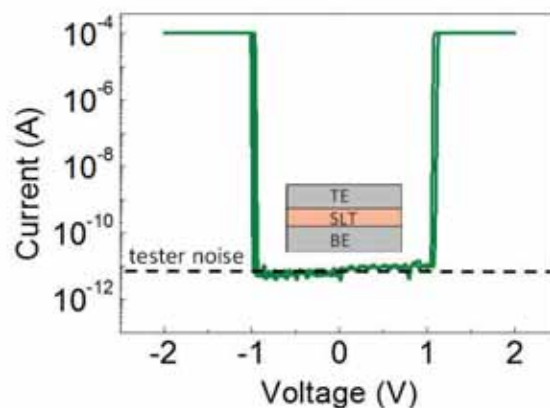


Figure 1a. I-V characteristics of FAST selectors (100nm x 100nm). The device exhibits bidirectional threshold switching with an on/off ratio (test limited) greater than 10^7 .




Figure 1b. Zoomed-in plot showing the extremely sharp FAST selector turn-on slope of less than 5mV/dec.

Measurements on a 4Mb 1S1R crossbar array with integrated FAST selectors show the sneak current suppressed to below 0.1nA, while maintaining a 10^2 memory on/off ratio and greater than 10^6 selectivity during cycling, making it ideal for ultra-high density memory applications.

The Global Semiconductor Alliance (GSA) mission is to accelerate the growth and increase the return on invested capital of the global semiconductor industry by fostering a more effective ecosystem through collaboration, integration and innovation. It addresses the challenges within the supply chain including IP, EDA/design, wafer manufacturing, test and packaging to enable industry-wide solutions. Providing a platform for meaningful global collaboration, the Alliance identifies and articulates market opportunities, encourages and supports entrepreneurship, and provides members with comprehensive and unique market intelligence. Members include companies throughout the supply chain representing 25 countries across the globe.

GSA Member Benefits Include:

- Access to Profile Directories
- Ecosystem Portals
- IP ROI Calculator
- IP Ecosystem Tool Suite
- Global Semiconductor Funding, IPO and M&A Update
- Global Semiconductor Financial Tracker
- End Markets Tool
- MS/RF PDK Checklist
- AMS/RF Process Checklist
- MS/RF SPICE Model Checklist
- Collaborative Innovation Study with the Wharton School

Discounts on various reports and publications including:

- Wafer Fabrication & Back-End Pricing Reports
- *Understanding Fabless IC Technology Book*
- IC Foundry Almanac, 2011 Edition

Global Exposure Opportunities:

- Advertising
- Sponsorships
- Member Press Section on GSA Web site
- Company Index Listing on Ecosystem Portals
- Member Spotlight on GSA Web site






Figure 1c. The threshold voltages can be tuned by controlling the SLT layer thickness.

THE FAST SELECTOR DEVICE

The patented FAST selector device utilizes a superlinear threshold layer (STL), which forms a conduction path at the specified threshold voltage (V_{TH}). The device provides extremely fast, bidirectional switching with a large resistance ratio, high turn-on current and steep turn-on slope, as shown in Figure 1a.

The measured selectivity for the 100nm x 100nm device is greater than 10^7 and limited by the test setup. The switching slope is extremely sharp – less than 5mV/dec. (Figure 1b), which is advantageous for array level operations (e.g. larger read voltage margin, faster read time).

The threshold voltage (V_{TH}) of the selector can be tuned by controlling either the STL thickness or the device structure (Figures 1c and 1d).

The FAST selector can switch reliably for more than 100M cycles (test limited), and reliable switching can also be maintained in 15nm x 100nm devices, with current density greater than 5×10^6 A/cm². The FAST selector was DC stress tested at a constant $0.5V_{TH}$




Figure 2. Hold voltage (V_H) characteristics of FAST selectors. (a) Once a device switches on, a small V_H is required for passing a specific target current (passing current I_P). (b) I_P vs. V_H . (c) Median on resistances vs. I_P . (d) V_{TH} vs. V_H .




Figure 1d. Asymmetric threshold voltages can be achieved by controlling the device structure (by modulating the electric field).


applied for two hours, and the device did not turn on, confirming good immunity to program/read disturb.

Switching speeds faster than 50ns can be achieved with voltages above V_{TH} . The FAST selector can be turned on within 30ns. The off-to-on transition time is about 5ns (test limited) for over 300 μ A of passing current through the selector, and the device can quickly recover to the off state once the voltage is removed, with a recovery time of less than 50ns.

Once a FAST device switches to the on state, a target pass current, I_P , can be delivered with a much smaller hold voltage, V_H (Figure 2a). V_H increases as I_P increases, but it is independent of V_{TH} (Figures 2b-d). The small V_H (< 0.3V for 200 μ A) and very large off-state resistance minimize the voltage overhead when integrated with RRAM.

INTEGRATING FAST SELECTORS IN CROSSBAR ARRAYS

The FAST selector has been successfully integrated into a 4Mb passive crossbar array, with a sneak current below 0.1nA at both 25°C and 125°C, demonstrating very high device yield and low leakage current.






Figure 3. Passive crossbar integration of RRAM devices with FAST selectors. (a) I-V characteristics of a single cell level RRAM, (b) selector, and (c) integrated 1S1R device. (d) I-V characteristics of a 4Mb passive crossbar array based on 1S1R.

For the integration, forming-free, low current ($\leq 20\mu\text{A}$) RRAM cells were developed to minimize IR drop and power consumption in large arrays (Figure 3a). FAST selectors with V_{TH} larger than $0.5V_{\text{PRG}}$, but smaller than the V_{PRG} of the RRAM, were designed to suppress sneak currents during both program and read operations (Figure 3b).

The resulting integrated 1S1R device exhibits a memory on/off ratio greater than 10^2 and selectivity greater than 10^6 (Figure 3c), and device operations have been successfully maintained for 4Mb 1S1R crossbar arrays (Figure 3d), which is the largest array size demonstrated to date for a passive crossbar structure.

The integrated 1S1R device switches reliably for more than 100K cycles while maintaining the large memory on/off ratio and selectivity (Figure 4).

When the leakage current through an entire 40Kb selector array was measured to extract the intrinsic leakage current of an individual selector (Figure 5a), the extracted selectivity was found to be 10^{10} (in a 100nm device). Figure 5a also shows that there is no single shorted selector device within the 40Kb selector array. The selectivity for different device areas was calculated using the same test method (Figure 5b). Circuit simulations show that a selectivity greater than 10^5 is required to design megabit-level passive crossbars, and the FAST selector clearly surpasses this requirement.




Figure 5. Leakage current test of selectors and the projected selectivity. (a) Leakage current through entire 40Kb devices and the projected 1-bit (100nm x 100nm) leakage current. Inset: Typical I-V of a selector on the same wafer. (b) The selectivity vs. device area based on the leakage current measurement.




Figure 4. Integrated 1S1R device cycling demonstration. On and off states and half-selected currents are shown. The integrated 1S1R device maintained $> 10^2$ memory on/off ratio and $> 10^6$ selectivity during the cycling.

CONCLUSION

The FAST selectors developed by Crossbar offer the largest reported selectivity to date (10^{10}), as well as excellent performance metrics for other characteristics required for high density memory applications, including steep slope and fast turn on/recovery. The high selectivity of the FAST device, and its ability to be integrated directly into each RRAM memory cell, make it possible to move beyond the density limitations of 1T1R array structures and implement commercial memory products, based on 3D stackable 1TnR memory architectures, for ultra-high density nonvolatile memory application. ♦

Hot Trends for 2015

The new system-design imperative

We're at a tipping point in system design. In the past, the consumer hung on every word from technology wizards, looking longingly at what was to come. But today, the consumer calls the shots and drives the pace and specifications of future technology directions. This has fostered, in part, a new breed of system design companies that has taken direct control over the semiconductor content.

These systems companies are reaping business (pricing, availability), technical (broader scope of optimization) and strategic (IP protection, secrecy) benefits. This is clearly a trend in which the winning systems companies are partaking.

They're less interested in plucking components from shelves and soldering them to boards and much more interested in conceiving, implementing and verifying their systems holistically, from application software down to chip, board and package. To this end, they are embracing the marriage of EDA and IP as a speedy and efficient means of enabling their system visions. For companies positioned with the proper products and services, the growth opportunities in 2015 are enormous.

THE SHIFT LEFT

Time-to-market pressures and system complexity force another reconsideration in how systems are designed. Take verification for example. Systems design companies are increasingly designing at higher levels, which requires understanding and validating software earlier in the process. This has led to the "shift left" phenomenon.

The simple way to think about this trend is that everything that was done "later" in the design flow is now being started "earlier" (e.g., software development begins before hardware is completed). Another way to visualize this macroscopic change is to think about the familiar system development "V-diagram" (Figure 1). The essence of this evolution is the examination of any and all dependencies in the product



Figure 1. Everything that was done "later" in the design flow is now being started "earlier."

planning and development process to understand how they can be made to overlap in time.

This overlap creates the complication of "more moving parts" but it also enables co-optimization across domains. Thus, the right side of the "V" shifts left (Figure 2) to form more of an accelerated flow. (Note: for all of the engineers in the room, don't be too literal or precise; it is meant to be thematic of the trend).

Prime examples of the shift left are the efforts in software development that are early enough to contemplate hardware changes (i.e., hardware optimization and hardware dependent software optimization), while at the other end of the spectrum we see early collaboration between the foundry, EDA tool makers and IP suppliers to co-optimize the overall enablement offering to maximize the value proposition of the new node.




Figure 2. The right side of the "V" shifts left to form more of an accelerated flow.

A by-product of the early software development is the enablement of software-driven verification methodologies that can be used to verify that the integration of sub-systems does not break the design. Another benefit is that performance and energy can be optimized in the system context with both hardware

and software optimizations possible. And, it is no longer just performance and power – quality, security and safety are also moving to the top level of concerns.

CHIP-PACKAGE-BOARD INTERDEPENDENCIES

Another design area being revolutionized is packaging. Form factors, price points, performance and power are drivers behind squeezing out new ideas. The lines between PCB, package, interposer and chip are being blurred.

Having design environments that are familiar to the principle in the system interconnect creation, regardless of being PCB, package or die centric by nature, provides a cockpit from which the cross fabric structures can be created, and optimized. Being able to provide all of the environments also means that data interoperable data sharing is smooth between the domains. Possessing analysis tools that operate independent of the design environment offers the consistent results for all parties incorporating the cross fabric interface data. In particular power and signal integrity are critical analyses to ensure design tolerances without risking the cost penalties of overdesign.

THE RISE OF MIXED-SIGNAL DESIGN

In general, but especially driven by the rise of Internet of Things (IoT) applications, mixed-signal design has soared in recent years. Some experts estimate that as much as 85% of all designs have at least some mixed-signal elements on board.

Being able to leverage high quality, high performance mixed signal IP is a very powerful solution to the complexity of mixed signal design in advanced nodes. Energy-efficient design features are also pervasive. Standards support for power reduction strategies (from multi-supply voltage, voltage/frequency scaling, and power shut-down to multi-threshold cells) can be applied across the array of analysis, verification and optimization technologies.

To verify these designs, the industry has been a little slower to migrate. The reality is that there is only so much tool and methodology change that can be digested by a design team while it remains immersed in the machine that cranks out new designs. So, offering a step-by-step progression that lends itself to incremental progress is what has been devised. “Beginning with the end in mind” has been the mantra of the legions of SoC verification teams that start with a sketch of the outcome desired in the planning and management phase at the beginning of the program. The industry best practices are summarized as: MD-UVM-MS – that is, metrics-driven unified verification methodology with mixed signal. ♦




Figure 3. IBS Mixed-signal design start forecast (source: IBS)

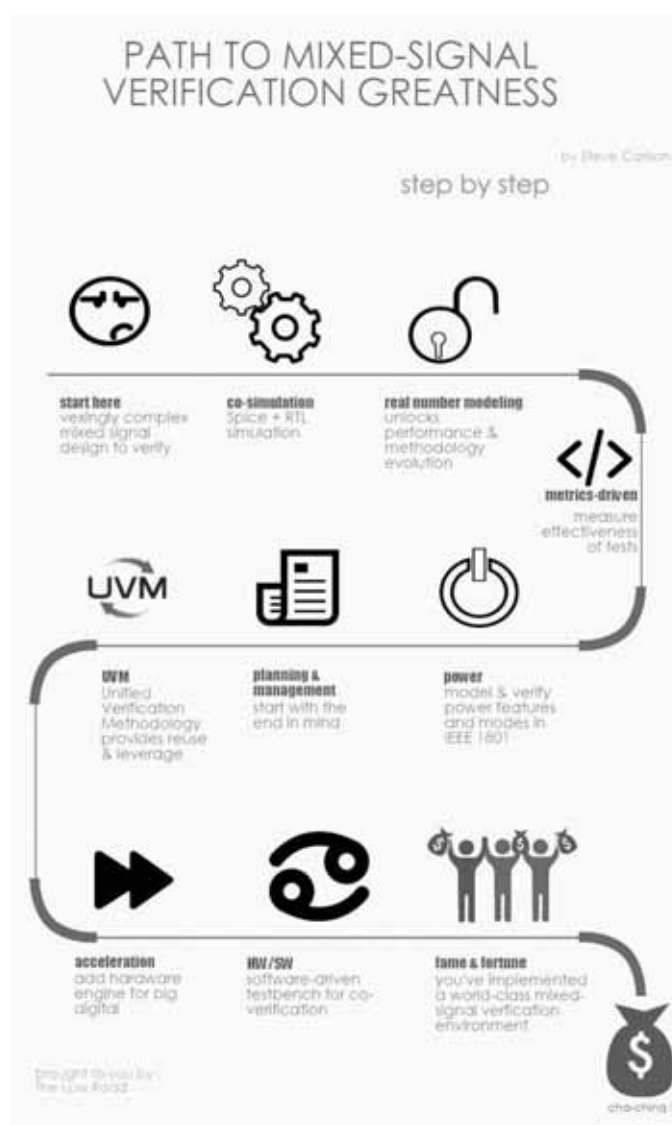


Figure 4. Path to MS Verification Greatness

The Coming Year in EDA: What Will Shape 2015

In spite of the financial importance of the “big three” – Cadence, Mentor Graphics and Synopsys -- a significant amount of innovation comes from smaller companies focused on one or just a few sectors of the market.

Piyush Sancheti, VP of Marketing at Atrenta pointed out that users drive the market and that users worry about time to market. In the companion article Chi-Ping Hsu of Cadence stated the same. To meet the market window they need predictability in product development, and therefore must manage design size and complexity, handle IP quality and integration risks, and avoid surprises during development. He observed that “The EDA industry as a whole is still growing in the single digits. However, certain front-end domains like RTL design and verification are growing much faster. The key drivers for growth are emulation, static and assertion-based verification, and power intent verification.

As the industry matures, consolidation around the big three vendors will continue to be a theme. Innovation is still fueled by startups, but EDA startup activity is not quite as robust. In 2014 Synopsys, Cadence and Mentor continued to drive growth with their investment and acquisitions in the semiconductor IP space, which is a good trend.”

Hamhua Ng, CEO of Plunify said: “There is much truth in the saying, ‘Those who don’t learn from history are doomed to repeat it,’ especially in the data-driven world that we live in today. It seems like every retailer, social network and financial institution is analyzing and finding patterns in the data that we generate. To businesses, being able to pick out trends from consumer behavior and quickly adapt products and services to better address customer requirements will result in significant cost savings and quality improvements.

Intuitively, chip design is an ideal area to apply these data analysis techniques because of an abundance of data generated in the process and the sheer cost and expertise required in realizing a design from the drawing board all the way to silicon. If only engineers can learn useful lessons – For instance, what worked well and what didn’t work as well – from all the chips that have ever been designed in history, what insights we would have today. Many companies




Figure 1: A new tool for DDR System analysis can check bit-level margins on a DDR read.

already have processes in place for reviewing past projects and extracting information.”

While talking about design complexity Bill Neifert, CTO at Carbon Design Systems noted that: “Initially targeted more at the server market, ARM’s 64-bit V8 architecture has been thrust into mobile since Apple announced that it was using it for the iPhone 5. Since then, we’ve seen a mad dash as semiconductor companies start developing mobile SoC designs containing multiple clusters of multicore processors.

Mobile processors have a large amount of complexity in both hardware and software. Coping with this move to 64 bits has placed a huge amount of stress on the hardware, software and systems teams.

With the first generation of 64-bit designs, many companies are handling this migration by changing as few variables as possible. They’ll take reference implementations and heavily leverage third-party IP in order to get something out the door. For this next generation of designs though, teams are starting to add more of their own differentiating IP and software. This raises a host of new verification and validation issues especially when addressing the complications being introduced with hardware cache coherency.”

INTERNET OF THINGS (IOT)

IoT is expected to drive much of the growth in the electronics industry and therefore in EDA. One can begin to see a few examples of products designed to work in the IoT architecture, even if the architecture is not yet completely finalized. There are wearable products that at the moment work only locally but have the potential to be connected via a cell phone to a central data processing system that generates information. Intelligent cars are at the moment self-contained IoT architectures that collect data, generate information, and in some cases, act on the information in real time.

David Kelf, VP of Marketing at OneSpin Solutions talked about the IoT in the automotive area. “2015 is yet again destined to be an exciting year. We see some specific IoT applications taking off, particularly in the automotive space and with other designs of a safety critical nature. It is clear that automotive electronics is accelerating. In particular is the concept of various automotive “apps” running on the central computer that interfaces with sensors around the car. This leads to a complex level of interaction, which must be fully verified from a safety critical and security point of view, and this will drive the leading edge of verification technology in 2015. Safety Critical standards will continue to be key verification drivers, not just in this industry sector but for others as well.”

Drew Wingard, CTO of Sonics said that: “The IoT market opportunity is top-of-mind for every company in the electronics industry supply chain including EDA tool vendors. For low-cost IoT devices, systems companies cannot afford to staff 1000-person SoC design teams. Furthermore, why do system design companies need two verification engineers for every designer? From an EDA tools and methodology perspective, today’s approach doesn’t work well for IoT designs.

SoC designers need to view their IoT designs in a more modular way and accept that components are “known good” across levels of abstraction. EDA tools and the verification environments that they support must eliminate the need to re-verify components whenever they are integrated into the next level up. It boils down to verification reuse. Agile design methodologies have a focus on automated component testing that SoC designers should consider carefully. IoT will drive the EDA industry trend toward a more agile methodology that delivers faster time-to-market. EDA’s role in IoT is to help lower the cost of design and verification to meet the requirements of this new market.”

VERIFICATION

Verification continues to be a hot topic. The emphasis has shifted from logic verification to system verification, where system is

understood to contain both hardware and software components. As the level of abstraction of design under test (DUT) has increased, the task of verification has become more demanding.


Michael Sanie, Senior Director of Verification Marketing at Synopsys talked about the drivers that will influence verification progress in 2015.

“SoCs are growing in unprecedented complexity, employing a variety of advanced low power techniques and an increasing amount of embedded software. While both SoC verification and software development/validation traditionally have been the long-poles of project timelines, they are now inseparable and together have a significant impact on time-to-market. Advanced SoC verification teams are now driven by not only reducing functional bugs, but also by how early they can achieve software bring-up for the SoCs. The now-combined process of finding/debugging functional bugs and software bring-up is often comprised of complex flows with several major steps including virtual platforms, static and formal verification, simulation, emulation and FPGA-based prototyping, with tedious and lengthy transitions between each step taking as long as weeks. Further complicating matters, each step requires a different technology/methodology for debug, coverage, verification IP, etc.

In 2015, the industry will continue its journey into new levels of verification productivity and early software bring-up by looking at how these steps can be approached universally with the introduction of larger platforms built from the industry’s fastest engines for each of these steps, further integration and unification compile, set up, debug, verification IP and coverage. Such an approach creates a continuum of technologies leveraging virtual platforms, static and formal verification, simulation, emulation and FPGA-based prototyping, enabling a much shorter transition time between each step. It further creates a unified debug solution across all domains and abstraction levels. The emergence of such platforms will then enable dramatic increases in SoC verification productivity and earlier software bring-up/development.”

Bill Neifert of Carbon says that: “In order to enable system differentiation, design teams need to take a more system-oriented approach. Verification techniques that work well at the block level start falling apart when applied to complex SoC systems. There needs to be a greater reliance upon system-level validation methodologies to keep up with the next generation of differentiated 64-bit designs. Accurate virtual prototypes can play a huge role in this validation task and we’ve seen an enormous upswing in the adoption of our Carbon Performance Analysis Kits (CPAKs) to perform exactly this task. A CPAK from System Exchange,

Register and
get your tickets now!
embedded-world.de



Media partners

elektroniknet.de

computer-automation.de

energie-und-technik.de

MEDIZIN-UND-elektronik.DE

Markt&Technik
DIE UNABHÄNGIGE WOCHENZEITUNG FÜR ELEKTRONIK

**DESIGN &
ELEKTRONIK**
KNOW-HOW FÜR ENTWICKLER

elektroniknet.de
Elektronik
Fachmedium für industrielle Anwender und Entwickler

**Elektronik
automotive**
Fachmedium für professionelle Automobilelektronik


**ENERGIE
& TECHNIK**
— Fachmedium für Energieeffizienz —

**Computer &
AUTOMATION**
Fachmedium der Automatisierungstechnik

MEDIZIN  **elektronik**
Fachmedium für Elektronik in der Medizintechnik

Nürnberg, Germany

February 24 – 26, 2015



embedded world 2015


Exhibition & Conference

... it's a smarter world

THE gathering for the embedded community!

The world's largest event for embedded technologies brings together all the players in the embedded sector for a productive conversation.

Make sure you are there when contacts and networks are being cemented on the international level and new trends are being set.



Trade fair organizer

NürnbergMesse GmbH

Tel+49 (0) 9 11.86 06-49 12

visitorservice@nuernbergmesse.de

Congress organizer

WEKA FACHMEDIEN GmbH

Tel+49 (0) 89.2 55 56-13 49

info@embedded-world.eu

NÜRNBERG / MESSE



for example, can be customized quickly to accurately model the behavior of the SoC design and then exercised using system-level benchmarks or verification software. This approach enables teams to spend far less time developing their validation solution and a lot more time extracting value from it.”

We hear a lot about design reuse, especially in terms of IP use. Drew Wingard of Sonics points to a lack of reuse in verification. “One of the biggest barriers to design reuse is the lack of verification reuse. Verification remains the largest and most time-consuming task in SoC design, in large part due to the popularity of constrained-random simulation techniques and the lack of true, component-based verification reuse. Today, designers verify a component at a very small unit level, then re-verify it at the IP core level, then re-verify the IP core at the IP subsystem level, then re-verify the IP subsystem at the SoC level and then, re-verify the SoC in the context of the system.

They don’t always use the same techniques at every one of those levels, but there is significant effort spent and test code developed at every level to check the design. Designers run and re-write the tests at every level of abstraction because when they capture the test the first time, they don’t abstract the tests so that they could be reused.

SoC designers need to view their IoT designs in a more modular way and accept that components are “known good” across levels of abstraction. EDA tools and the verification environments that they support must eliminate the need to re-verify components whenever they are integrated into the next level up. It boils down to verification reuse. Agile design methodologies have a focus on automated component testing that SoC designers should consider carefully. IoT will drive the EDA industry trend toward a more agile methodology that delivers faster time-to-market. EDA’s role in IoT is to help lower the cost of design and verification to meet the requirements of this new market.”

Piyush Sancheti of Atrenta acknowledges that front-end design and verification tools are growing driven by more complex designs and shorter time-to-market. But design verification difficulty continues to increase with shrinking time to completion reality. Companies are turning more and more to static verification, formal techniques and emulation. The goal is RTL debug and signoff aiming at more automatic, or knowledge based place and route functions.

Regarding formal tools David Kelf of OneSpin noted that: “Formal techniques, in general, continue to proliferate through many verification flows. We see an increase in the use of the technology by designers to perform initial design investigation, and

greater linkage into the simulation flow. Combining the advances in FPGA, the safety critical driver for high-reliability verification and increases in formal technology, we believe that this year fundamental shifts.”

Jin Zhang, Senior Director of Marketing at Oski Technology had an interesting input to the subject of formal verification because it was based on the feedback she received recently from the Decoding Formal Club. Here is what he said: “In October, Oski Technology hosted the quarterly Decoding Formal Club where more than 40 formal enthusiasts gathered to talk about Formal Sign-off and processor verification using formal technology. The sheer energy and enthusiasm of Silicon Valley engineers speaks to the growing adoption of formal verification.

Several experts on formal technology who attended the event view the future of formal verification similarly. They echoed the trends we have been seeing — formal adoption is in full bloom and could soon replace simulation in verification sign-off.

What’s encouraging is not just the adoption of formal technology in simple use models, such as formal lint or formal apps, but in an expert use model as well. For years, expert-level use has been regarded as academic and not applicable to solving real-world challenges without the aid of a doctoral degree. Today, End-to-End formal verification, as the most advanced formal methodology, leads to complete verification of design blocks with no bugs left behind. With ever-increasing complexity and daunting verification tasks, the promise and realization of signing off a design block using formal alone is the core driver of the formal adoption trend.

The trend is global. Semiconductor companies worldwide are recognizing the value of End-to-End formal verification and working to apply it to critical designs, as well as staffing in-house formal teams. Formal verification has never been so well regarded.

While 2015 may not be the year when every semiconductor company has adopted formal verification, it won’t be long before formal becomes as normal as simulation, and shoulders as much responsibility in verification sign off.”

ADVANCED PROCESSES CHALLENGES

Although the number of system companies that can afford to use advanced processes is diminishing, their challenges are an important indicator of future requirements for a larger set of users.

Mary Ann White, Director of Product Marketing, Galaxy Design Platform at Synopsys points out how timing and power requirement analysis are critical elements of design flows.

“The endurance of Moore’s law drives design and EDA trends where consumer appetites for all things new and shiny continue to be insatiable. More functional consolidation into a single SoC pushes ultra-large designs more into the norm, propelling the need for more hierarchically oriented implementation and signoff methodologies. While 16- and 14-nm FinFET technologies become a reality by moving into production for high-performance applications, the popularity of the 28-nm node will persevere, especially for mobile and IoT (Internet of Things) devices.

Approximately 25% of all designs today are ≥ 50 million gates in size according to Synopsys’ latest Global User Survey. Nearly half of those devices are easily approaching one billion transistors. The sheer size of these designs compels adoption of the latest hierarchical implementation techniques with either black boxes or block abstract models that contain timing information with interfaces that can be further optimized. The Galaxy Design Platform has achieved several successful tapeouts of designs with hundreds of millions of instances, and newer technologies such as IC Compiler II have been architected to handle even more. In addition, utilization of a sign-off based hierarchical approach, such as PrimeTime HyperScale technology, saves time to closure, allowing STA completion of 100+ million instances in hours vs. days while also providing rapid ECO turnaround time.

The density of FinFET processes is quite attractive, especially for high-performance designs which tend to be very large multi-core devices. FinFET transistors have brought dynamic, rather than leakage (static) power to the forefront as the main concern. Thanks to 20-nm, handling of double patterning is now well established for FinFET. However, the next-generation process nodes are now introduced at a much faster pace than ever before, and test chips for the next 10-nm node are already occurring. Handling the varying multi-patterning requirements for these next-generation technologies will be a huge focus over the next year, with early access and ecosystem partnerships between EDA vendors, foundries and customers.

Meanwhile, as mobile devices continue their popularity and IoT devices (such as wearables) become more prevalent, extending battery life and power conservation remain primary requirements. Galaxy has a plethora of different optimization techniques to help mitigate power consumption. Along with the need for more silicon efficiency to lower costs, the 28-nm process node is ideal for these types of applications. Already accounting for more than a third of revenue for foundries, 28-nm (planar and FD-SOI) is poised to last a while even as FinFET processes come online.”

Dr. Bruce McGaughy, CTO and VP of Engineering at ProPlus was kind enough to provide his point of view on the subject. “The challenges of moving to sub-20nm process technologies are forcing designers to look far more closely at their carefully constructed design flows. The trend in 2015 could be a widespread retooling effort to stave off these challenges as the most leading-edge designs start using FinFET technology, introducing a complication at every step in the design flow.

Observers point to the obvious: Challenges facing circuit designers are mounting as the tried-and-true methodologies and design tools fall farther and farther behind. It’s especially apparent with conventional SPICE and FastSPICE simulators, the must-have tools for circuit design.

FastSPICE is faltering and the necessity of using Giga-scale SPICE is emerging. At the sub-28nm process node, for example, designers need to consider layout dependent effects (LDE) and process variations. FastSPICE tricks and techniques, such as isomorphism and table models, do not work.

As we move further into the realm of FinFET at the sub-20nm nodes, FastSPICE’s limitations become even more pronounced. FinFET design requires a new SPICE model called BSIM-CMG, more complicated than BSIM3 and BSIM4 models, industry-standard models for CMOS technology used by the industry for 20 years. New FinFET physical effects include strong Miller Capacitance effects, break FastSPICE’s partitioning and event-driven schemes. Typically, FinFET models have over 1,000 parameters per transistor, and more than 20,000 lines of C code, posing a tremendous computational challenge to SPICE simulators.

Furthermore, the latest advanced processes pose new and previously undetected challenges. With reduced supply voltage and increased process variations, circuits now are more sensitive to small currents and charges, such as SRAM read cycles and leakage currents. FastSPICE focuses on event-driven piecewise linear (PWL) voltage approximations rather than continuous waveforms of currents, charges and voltages. More delay and noise issues are appearing in the interconnect, requiring post layout simulation with high accuracy, and multiple power domains and ramp-up/ramp-down cycles are more common.

All are challenging for FastSPICE, but can be managed by Giga-scale SPICE simulators. FastSpice simulators are lacking in accuracy for sensitive currents, voltage regulators and leakage, which is where Giga-scale SPICE simulators can really shine.

Time to market and high mask costs demand tools that always give accurate and reliable results, and catch problems before tapeout. Accordingly, designers and verification engineers are demanding a tool that does not require the “tweaking” of options to suit their situations, such as different sets of simulation options for different circuit types, and assigned accuracy where the simulator thinks it’s needed. Rather, they should get accuracy everywhere without option tweaks.

Often, verification engineers are not as familiar with the circuits as the designers, and may inadvertently choose a set of options that causes the FastSPICE simulator to ignore important effects, such as shorting out voltage regulator power nets. Weak points could be lurking in those overlooked areas of the chip. With Giga-scale SPICE, such approximations are not used and unnecessary.

Here’s where Giga-scale SPICE simulation takes over, being perfectly suited for the new process technologies including 16/14nm FinFET. They offer pure SPICE accuracy and deliver comparable FastSPICE simulator capacity and performance.

For the first time, Giga-scale SPICE makes it possible for designers to use one simulation engine to design both small and large circuit blocks, and simultaneously use the same simulator for full-chip verification with SPICE accuracy, eliminating glitches or inconsistencies. We are at the point that retooling the simulation tools, including making investments in parallel processing hardware, is the right investment to make to improve time to market and reduce the risk or respins. At the same time, tighter margins can be achieved in design, resulting in better performance and yield.”

MEMORY AND FPGA

With the increase use of embedded software in SoC designs memories and memory controllers are gaining in importance. Bob Smith, Senior VP of Marketing and Business Development at Uniquify presents a compelling argument.

”The term ‘EDA’ encompasses tools that both help in automating the design process (think synthesis or place and route) as well as automating the analysis process (such as timing analysis or signal integrity analysis). A new area of opportunity for EDA analysis is emerging to support the understanding and characterization of high-speed DDR memory subsystems.

The high-performance demands placed on DDR memory systems requires that design teams thoroughly analyze and understand the system margins and variations encountered during system operation. At rated DDR4 speeds, there is only about 300ps of allowable timing margin across the entire system (ASIC or SoC,

package, PCB or other interconnect media, and the DDR SDRAM itself). Both static (process-related) and dynamic variations (due to environmental variables such as temperature) must also be factored into this tight margin. The goal is straightforward: optimize the complete DDR system to achieve the highest possible performance while minimizing any negative impacts on memory performance due to anticipated variations and leave enough margin such that unanticipated variations don’t cause the memory subsystem to fail.

However, understanding the available margin and how it will be impacted by static and dynamic variation is challenging. Historically, there is no visibility into the DDR subsystem itself and the JEDEC specifications only address the behavior of the DDR-SDRAM devices and not the components outside of the device. Device characterization helps, but only accounts for part-part variation. Ideally, DDR system designers would like to be able to measure timing margins in-situ, with variations present, to fully and accurately understand system behavior and gain visibility into possible issues.

A new tool for DDR System analysis provides this visibility (Figure 1). A special interface in the DDR PHY allows it to run numerous different analyses to check the robustness of the entire DDR system including the board and DDR-SDRAM device(s). The tool can be used to determine DDR system margins, identify board or DDR component peculiarities and be used to help tune various parameters to compensate for issues discovered and maximize DDR performance in a given system. Since it is essentially a window into the DDR subsystem, it can also be used to characterize and compare the performance of different boards and board layouts and even compare the performance and margins of different DDR-SDRAM components.

Charlie Cheng, CEO of Kilopass talks about the need for new memory technology. “For the last few years, the driver for chip and EDA tool development has been the multicore SoC that is central to all smartphones and tablets. Memory dominates these devices with 85% of the chip area and an even larger percentage of the leakage power consumption. Overlay the strong friction that is slowing down the transition to 14/16nm from 28nm — the most effective node. It becomes quickly obvious that a new high-density, power-thrifty memory technology is needed at the 28nm node. Memory design has been sorely lacking in innovation for the last 20 years with all the resources getting invested on the process side. 2015 will be the year of major changes in this area as the industry begins to take a better look at support for low-power, security and mobile applications.”

The use of FPGA devices in SoC has also increased in 2014. David Kelf thinks that: “The significant advancement in FPGA technology will lead to a new wave of FPGA designs in 2015. New device geometries from the leading FPGA vendors are changing the ASIC to FPGA cost/volume curve and this will have an affect in the size and complexity of these devices. In addition, we will see more specialized synthesis tools from all the vendors, which provide for greater, device-targeted optimizations. This in turn drives further advancement of new verification flows for larger FPGAs and it is our prediction that most of the larger devices will make use of a formal based verification flow to increase overall QoR (Quality of Results).”

The need for better tools to support the use of FPGA is also acknowledged by Harnhua NG at Plurify. “FPGA software tools contain switches and parameters that influence synthesis optimizations and place-and-route quality of results. Combined with user-specified timing and location constraints, timing, area and power results can vary by as much as 70% without even modifying the design’s source code. Experienced FPGA designers

intuitively know good switches and parameters through years of experience, but have to manually refine this intuition as design techniques, chip architectures and software tools rapidly evolve. Continuing refinement and improvement are better managed using data analysis and machine learning algorithms.”

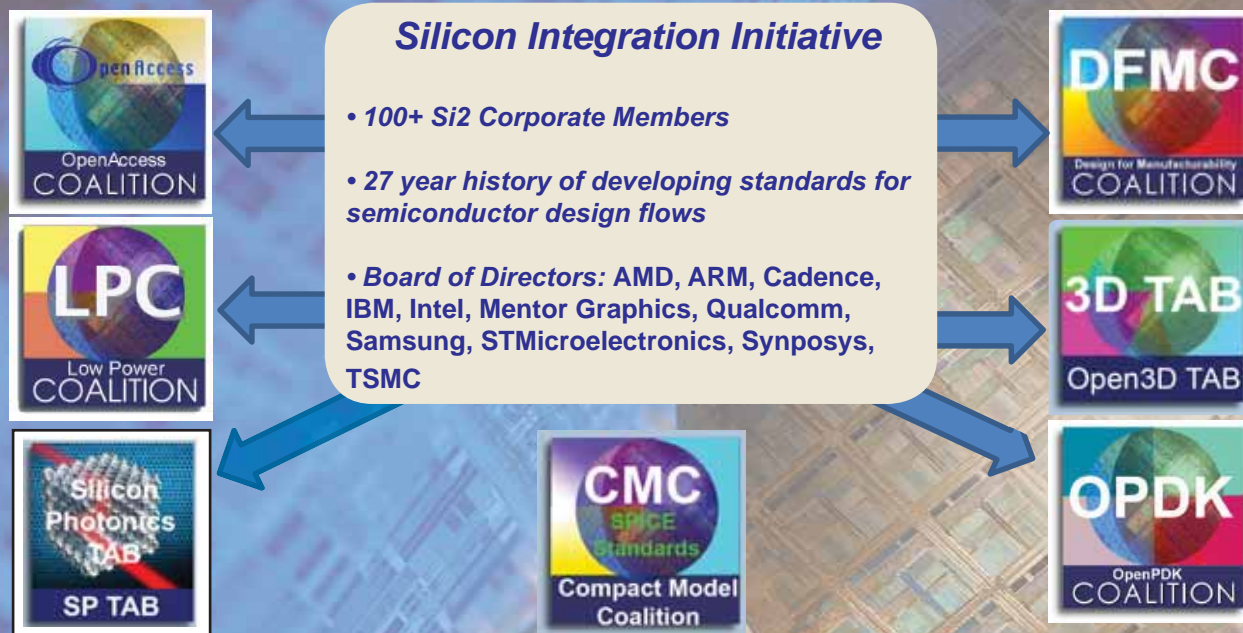
CONCLUSION

It should not be surprising that EDA vendors see many financial and technical opportunities available in 2015. Consumers’ appetite for electronic gadgets, together with the growth of cloud computing, and new IoT implementations provide markets for new EDA tools. How many vendors will hit the proper market windows is still to be seen and timing to market will be the fundamental characteristic of the 2015 EDA industry. ♦

Gabe Moretti is Senior Editor at Chip Design.



Far more than just Standards....



.....It's about Adoption and ROI

Featured Downloads

At: www.si2.org/?page=1038

www.si2.org



Pushing the Performance Boundaries of ARM Cortex-M Processors for Future Embedded Design

One of the toughest challenges in the implementation of any processors is balancing the need for the highest performance with the conflicting demands for lowest possible power and area. Inevitably, there is a tradeoff between power, performance, and area (PPA). This paper examines two unique challenges for design automation methodologies in the new ARM®Cortex®-M processor: how to get maximum performance while designing for a set power budget and how to get maximum power savings while optimizing for a set target frequency.

The ARM®Cortex®-M7 processor is the latest embedded processor by ARM specifically developed to address digital signal control markets that demand an efficient, easy-to-use blend of control and signal processing capabilities (Figure 1). The ARM Cortex-M7 processor has been designed with a large variety of highly efficient signal processing features, which demands very power- efficient design.

The energy-efficient, easy-to-use microprocessors in the ARM Cortex-M series have received a large amount of attention recently as portable and wireless / embedded applications have gained market share. In high-performance designs, power has become an issue since at those frequencies power dissipation can easily reach several tens of watts. The efficient handling of these power levels requires complex heat dissipation techniques at the system level, ultimately resulting in higher costs and potential reliability issues. In this section, we will isolate the different components of power consumption on a chip to demonstrate why power has become a significant issue. The remaining sections will discuss how we approached this problem and resolved it using Cadence® implementation tools, along with other design techniques.

We began the project with the objective of addressing two simultaneous challenges:

1. Reach, as fast as possible, a performance level with optimal power (AFAP)




Figure 1: ARM Cortex-M7 Block Diagram

2. Reduce power to the minimum for a lower frequency scenario (MinPower)

Before getting into the details of how we achieved the desired frequency and power numbers, let's first examine the components which contribute to dynamic power and the factors which gate the frequency push. This experiment has been conducted on the ARM Cortex-M7 processor. The ARM Cortex-M7 processor has achieved 5 CoreMark / MHz – 2000 CoreMark* in 40LP and typical 2X digital signal processing (DSP) performance of the ARM Cortex-M4 processor.

DYNAMIC POWER COMPONENTS

In high-performance microprocessors, there are several key reasons which are causing a rise in power dissipation. First, the presence of a large number of devices and wires integrated on a big chip results in an overall increase in the total capacitance of the design. Second, the drive for higher performance leads to increasing clock frequencies, and dynamic power is directly proportional to the rate of charging capacitances (in other words, the clock frequency). A third reason that may lead to higher power consumption is an inefficient use of gates. The total switching device capacitance consists of gate oxide capacitance, overlap capacitance, and junction capacitance. In

In addition, we consider the impact of internal nodes of a complex logic gate. For example, the junction capacitance of the series-connected NMOS transistors in a NAND gate contributes to the total switching capacitance, although it does not appear at the output node. Dynamic power is consumed when a gate switches. However, interest has risen in the physical design area, to make better use of the available gates by increasing the ratio of clock cycles when a gate actually switches. This increased device activity would also lead to rising power consumption. Dynamic power is the largest component of total chip power consumption (the other components are short-circuit power and leakage power). It occurs as a result of charging capacitive loads at the output of gates. These capacitive loads are in the form of wiring capacitance, junction capacitance, and the input (gate) capacitance of the fan-out gates. Since leakage is <2% of total power, the focus of this collaboration was only on dynamic power.

The expression for dynamic power is:

$$P_{\text{dynamic}} = \alpha C V_{\text{dd}}^2 f \dots\dots\dots(1)$$

In (1), C denotes the capacitance being charged /discharged, V_{dd} is the supply voltage, f is the frequency of operation, and α is the switching activity factor. This expression assumes that the output load experiences a full voltage swing of V_{dd}. If this is not the case, and there are circuits that take advantage of this fact, (1) becomes proportional to (V_{dd} * V_{swing}). A brief discussion of the switching factor α is in order at this point. The switching factor is defined in this model as the probability of a gate experiencing an output low-to-high transition in an arbitrary clock cycle. For instance, a clock buffer sees both a low-to-high and a high-to-low transition in each clock cycle. Therefore, α for a clock signal is 1, as there is unity probability that the buffer will have an energy-consuming transition in a given cycle. Fortunately, most circuits have activity factors much smaller than 1. Some typical values for logic might be about 0.5 for data path logic and 0.03 to 0.05 for control logic. In most instances we will use a default value of 0.15 for α, which is in keeping with values reported in the literature for static CMOS designs [1,2,3]. Notable exceptions to this assumption will be in cache memories, where read /write operations take place nearly every cycle, and clock-related circuits.

Here are five key components of dynamic power consumption and how we addressed a few of these components:

- Standard cell logic and local wiring

- Global interconnect (mainly busses, inter-modular routing, and other control)
- Global clock distribution (drivers + interconnect + sequential elements)
- Memory (on-chip caches) — this is constant in our case
- I/Os (drivers + off-chip capacitive loads) — this is constant in our case

TIMING CLOSURE COMPONENTS

One fundamental issue of timing closure is the modeling of physical overcrowding. The problem involves, among other factors, the representation and the handling of layout issues. These issues include placement congestion, overlapping of arbitrary-shaped components, routing congestion due to power/ground, clock distribution, signal interconnect, prefixed wires over components, and forbidden regions of engineering concerns. While a clean and universal mathematical model of physical constraint remains open, we tend to formulate the layout problem using multiple constraints with sophisticated details that complicate the implementation. We need to consider multiple constraints with a unified objective function for a timing-closure design process. This is essential because many constraints are mutually conflicting if we view and handle their effects only on the surface. For example, to ease the routing congestion of a local area, we tend to distribute components out of the area to leave more room for routing. However, for multi-layer routing technology, eliminating components does not save much on routing area. The spreading of components actually increases the wire length and demands more routing space. The resultant effect can have a negative impact on the goals of the original design. In fact, the timing can become much worse. Consequently, we need an intelligent operation that identifies both the component to move out and the component to move in to improve the design.

Accurately predicting the detail routed signal-integrity (SI) effects, before the detail routing happens, and its impact to timing is of key interest. This is because a reasonable misprediction of timing before the detail route would create timing jumps after the routing is done. Historically, designs for which it is tough to close timing have relied solely on post-route optimization to salvage setup /hold timing. With the advent of “in-route optimization”, timing closure has been bridged earlier during the routing step itself using track assignment. In addition, if we can reduce the wire lengths and make good judgment calls based on the timing profiles, we can




Figure 2: Soft-Blocked Floorplan



Figure 3: "GigaPlace" Placement Engine

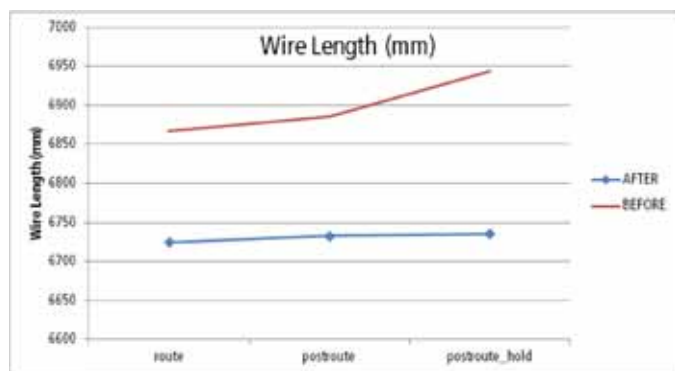


Figure 4: Wirelength Reduced with "GigaPlace and Soft-Blocked" Placement

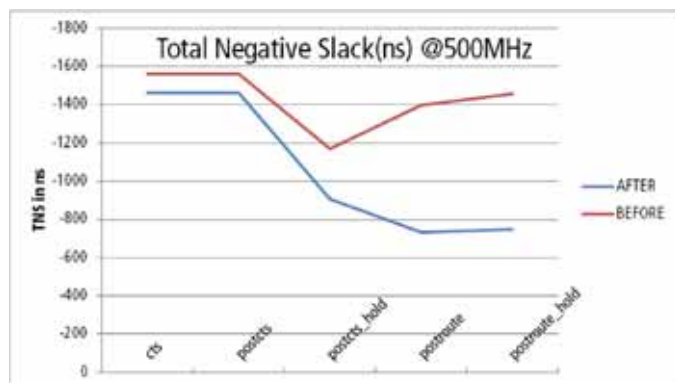


Figure 5: Total Negative Slack (ns) Chart

find opportunities to further reduce power. This paper will walk through the Cadence digital implementation flow and new tool options used to generate performance benefits for the design. The paper will also discuss the flow and tool changes that were done to get the best performance and power efficiency out of the ARM Cortex-M7 processor implementation.

BETTER PLACEMENT AND REDUCED WIRELENGTH

As discussed in the introduction, wire capacitance and gate capacitance are among the key factors that impact dynamic power, while also affecting wire delays. While evaluating the floorplan and cell placement, it was noticed that the floorplan size was bigger than needed and the cell placement density was uniform. These two aspects could lead to spreading out of cells, resulting in longer wirelength and higher clock latencies. In order to improve the placement densities, certain portions of the design were soft-blocked, and the standard cell densities were kept above 75% (Figure 2).

Standard cell placement plays a vital role. If the placement is done right, it will eventually pay off in terms of better Quality of Results (QoR) and wirelength reduction. If the placement algorithms can take into account some of the power dissipation-related issues, like reducing the wirelength and considering overall slack profile of the design, and also make the right moves during placement, this would tremendously improve the above mentioned aspect. This is the core principle behind the "Giga Place" placement engine. The Giga Place engine, available in Cadence Encounter® Digital Implementation System 14.1, helps place the cells in a timing-driven mode by building up the slack profile of the paths and performing the placement adjustments based on these timing slacks. We have introduced this new placement engine on the ARM Cortex-M7 design and seen good improvements on the overall wirelength and Total Negative Slack (TNS).

With a reduced floorplan and by removing the uniform placement and utilizing the new GigaPlace technologies (Figure 3), we were able to reduce the wirelength significantly (Figure 4). This helped push the frequency as well as reduce the power (Figure 5). But, there were still more opportunities available to further benefit the frequency and dynamic power targets.

IN-ROUTE OPTIMIZATION

"In-route optimization" for timing optimization happens before routing begins. This is a very close representation of the real routes, which does not account for the DRC fixes and the leaf-cell pin access. This enables us to get an accurate view of timing /SI and make bigger changes without disrupting the routes.

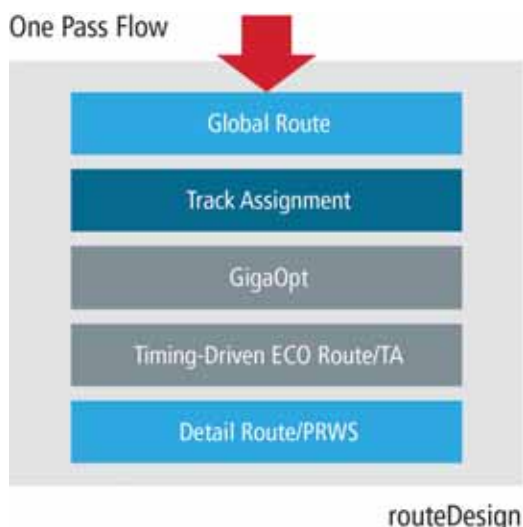


Figure 6: In-Route Optimization Flow Chart

These changes are then committed to a full detail route. In-route optimization technology utilizes an internal extraction engine for more effective RC modeling. The timing QoR improvement observed after post-route optimization was significant at the expense of a slight runtime increase (currently observed at only 2%). A successful usage of an internal extraction model during in-route optimization (Figure 6) helped reduce the timing divergence seen as we go from the pre-route to the post-route stage. This optimization technology pushed the design to achieve the targeted frequency.

DESIGN CHANGES AND FURTHER DYNAMIC POWER REDUCTION

In the majority of present-day electronic design automation (EDA) tools, timing closure is the top priority and, hence, many of these tools make the trade-off to give priority to timing. However, opportunities exist to reduce area and gate capacitance by swapping cells to lower gate cap cells and by reducing the wirelength. To address the dynamic power reduction in the design, three major sets of experiments were done to examine the above aspects.

In the first set of experiments, two main tool features were used in the process of reducing dynamic power (Figure 7). These were the introduction of the “dynamic power optimization engine” along with the “area reclaim” feature in the post-route stage. These options helped save 5% of dynamic power @400MHz and enabled us to nearly halve the gap that earlier existed between the actual and desired power target.

In the second set of experiments, the floorplan was soft-blocked by 100 microns to reduce the wirelength. This was discussed in detail in an earlier section. This floorplan shrink resulted in:



Figure 7: Example of Power Optimization

Run Details	Relative Dynamic Power Reduction
400MHz - RC/EDI 13.2	100
400MHz - RC/EDI 14.1	96
400MHz - 14.1 + PowerOpt	91
200MHz - 14.1	75
200MHz - 14.1 + PowerOpt	71
200MHz - 14.1 + PowerOpt + GigaPlace	67
200MHz - 14.1 + PowerOpt + GigaPlace + Relax Clock Skew	62

Table 1: Dynamic Power Reduction Results

- Increasing the density from ~76% to 85%
- Wirelength reduction by 5.1% – post route
- Area (with combo of #1 and shrink) shrinkage by ~4% – post route

This helped saved an additional 2% @400MHz, and the impact was similar across the frequency sweep.

The third set of experiments was related to design changes where flop sizes were downsized to a minimum at pre_ cts opt and the remaining flops of higher drive strengths were set to “don’t use”. This helped to further reduce the sequential power. An important point to note is that the combinational power did not increase significantly. After we introduced the above technique, we were able to reduce power significantly, as shown in the charts below.

RESULTS

By using these latest tool technologies and design techniques, we were able to achieve 10% better frequency and reduced the dynamic power by 10%. Results are shown in Table 1 based on the 400MHz and 200MHz for the dynamic power reduction.

The joint ARM /Cadence work started with addressing challenges at two points /scenarios on the PPA curve:

1. Frequency focus with optimal power (400MHz)
2. Lowest power at reduced frequency (200MHz)

For scenario #1, out of box 14.1 allowed us to reach 400MHz. With the use of PowerOpt technology, available in Encounter

Digital Implementation System 14.1, we were able to reduce power to an optimal number. For scenario #2, additional use of GigaPlace technology and inherently better SI management allowing relaxed clock slew, and much higher power reduction at 200MHz was possible. With the combination of ARM design techniques and Cadence tool features, we were able to show 38% dynamic power reduction (for standard cells) going from

400MHz – 13.2-based run to 200MHz – 14.2 best power recipe run.

SUMMARY

Reducing the wirelength and slack profile-based placement, and predicting the detailed routing impact in the early phase of the design, are important aspects to improve the performance and reduce the dynamic power consumption in designs. Tools perform better when given the right floorplan along with the proper directives at appropriate places. With a combination of design changes, advanced tools, and engineering expertise, today's physical design engineers have the means to thoroughly address the challenges associated with timing closure while keeping the dynamic power consumption of the designs low (Figure 8).

Several months of collaborative work between ARM and Cadence, driven by many trials, have led to optimized PPA results. Cadence tools – Encounter RTL Compiler/ Encounter Digital Implementation System 14.1 – have produced better results out of box compared to Encounter RTL Compiler/ Encounter Digital Implementation System 13.x. The continuous refinement of the flow along with design techniques such as floorplan reduction and clock slew relaxation allowed a 38% dynamic power reduction. The ARM/Cadence implementation Reference Methodology (iRM) flow uses a similar recipe for both scenarios: lowest power (MinP) and highest frequency (AFAP).




Figure 8: Dynamic Power (Normalized) for Logic

REFERENCES

1. D. Liu and C. Svensson, "Power consumption estimation in CMOS VLSI chips," *IEEE Journal of Solid-State Circuits*, vol. 29, pp. 663-670, June 1994.
2. A.P. Chandrakasan and R.W. Broderon, "Minimizing power consumption in digital CMOS circuits," *Proc. of the IEEE*, vol. 83, pp. 498-523, April 1995.
3. G. Gerosa, et al., "250 MHz 5-W PowerPC microprocessor," *IEEE Journal of Solid-State Circuits*, vol. 32, pp. 1635-1649, Nov. 1997. ♦

IoT Cookbook: Analog and Digital Fusion Bus Recipe

Experts from ARM, Mathworks, Cadence, Synopsys, Analog Devices, Atrenta, Hillcrest Labs and STMicroelectronics cook up ways to integrate analog with IoT buses.

Many embedded engineers approach the development of Internet-of-Things (IoT) devices like a cookbook. By following previous embedded recipes, they hope to create new and deliciously innovative applications.

While the recipes may be similar, today's IoT uses strong concentration of analog, sensors and wireless ingredients. How will these parts combine with the available high-end bus structures like ARM's AMBA?

To find out, "IoT Embedded Systems" talked with the head technical cooks including Paul Williamson, Senior Marketing Manager, ARM; Rob O'Reilly, Senior Member Technical Staff at Analog Devices; Mladen Nizic, Engineering Director, Mixed Signal Solution, Cadence; Ron Lowman, Strategic Marketing Manager for IoT, Synopsys; Corey Mathis, Industry Marketing Manager - Communications, Electronics and Semiconductors, MathWorks; Daniel Chaitow, Marketing Manager, Hillcrest Labs; Bernard Murphy, CTO, Atrenta; and Sean Newton, Field Applications Engineering Manager, STMicroelectronics. What follows is a portion of their responses.

KEY POINTS:

- System-level design is needed so that the bus interface can control the analog peripheral through a variety of modes and power-efficient scenarios.
- One industry challenge is to sort the various sensor data streams in sequence, in types, and include the ability to do sample or rate conversion.
- To ensure the correct sampling of analog sensor signals and the proper timing of all control and data signals, cycle accurate simulations must be performed.
- Control system and sensor subsystems are needed to help reduce digital bus cycles by tightly integrating the necessary components.

- Hardware design and software design have inherently different workflows, and as a result, use different design tools and methodologies.
- For low-power IoT sensors, the analog-digital converter (ADC) power supply must be designed to minimize noise. Attention must also be paid to the routing of analog signals between the sensors and the ADC.
- Beyond basic sensor interfacing, designer should consider digitally assisted analog (DAA) – or digital logic embedded in analog circuitry that functions as a digital signal processor.

Blyler: What challenges do designers face when integrating analog sensor and wireless IP with digital buses like ARM's AMBA and others?

Williamson (ARM): Designers need to consider system-level performance when designing the interface between the processor core and the analog peripherals. For example a sensor peripheral might be running continuously, providing data to the CPU only when event thresholds are reached. Alternatively the analog sensor may be passing bursts of sampled data to the CPU for processing. These different scenarios may require that the designer develop a digital interface that offers simple register control, or more advanced memory access. The design of the interface needs to enable control of the peripheral through a broad range of modes and in a manner that optimizes power efficiency at a system and application level.

O'Reilly (Analog Devices): One challenge is ultra-low power designs to enable management of the overall system power consumption. In IoT systems, typically there is one main SoC connected with multiple sensors running at different Output Data Rates (ODR) using asynchronous clocking. The application processor SoC collects the data from multiple sensors and completes the processing. To keep power consumption low, the SoC generally isn't active all of the time. The SoC will collect data at certain intervals.

To support the needs of sensor fusion it's necessary that the sensor data includes time information. This highlights the second challenge, the ability to align a variety of different data types in a time sequence required for fusion processing. This raises the question "How can an entire industry adequately sort the various sensor data streams in sequence, in types, and include the ability to do sample or rate conversion.?"

Nizic (Cadence): Typically a sensor will generate a small (low voltage/current) analog signal which needs to be properly conditioned and amplified before converting it to digital signal sent over a bus to memory register for further processing by a DSP or a controller. Sometimes, to save area, multiple sensor signals are multiplexed (sampled) to reduce the number of A2D converters.

From the design methodology aspect, the biggest design challenge is verification. To ensure analog sensor signals are sampled correctly and all control and data signals are timed properly, cycle-accurate simulations must be performed. Since these systems now contain analog, in addition to digital and bus protocol verification, a mixed-signal simulation must cover both hardware and software. To effectively apply mixed-signal simulation, designers must model and abstract behavior of sensors, analog multiplexers, A2D converters and other analog components. On the physical implementation side, busses will require increased routing resources, which in turn mean more careful floor-planning and routing of bus and analog signals to keep chip area at minimum and avoid signal interference.

Lowman (Synopsys): For an IC designer, the digital bus provides a very easy way to snap together an IC by hanging interface controllers such as I2C, SPI, and UARTs to connect to sensors and wireless controllers. It's also an easy method to hang USB and Ethernet, as well as analog interfaces, memories and processing engines. However, things are a bit more complicated on the system level. For example, the sensor in a control system helps some actuator know what to do and when to do it. The challenge is that there is a delay in bus cycles from sensing to calculating a response to actually delivering a response that ultimately optimizes the control and efficiency of the system. Examples include motor control, vision systems and power conversion applications. Ideally, you'd want a sensor and control subsystem that has optimized 9D Sensor Fusion application. This subsystem significantly reduces cycles spent traveling over a digital bus by essentially removing the bus and tightly integrating the necessary components needed to sense and process the algorithms. This technique will be critical to

reducing power and increasing performance of IoT control systems and sensor applications in a deeply embedded world.

Mathis (Mathworks): It is no surprise that mathematical and signal processing algorithms of increasing complexity are driving many of the innovations in embedded IoT. This trend is partly enabled by the increasing capability of SoC hardware being deployed for the IoT. These SoCs provide embedded engineers greater flexibility regarding where the algorithms get implemented. The greater flexibility, however, leads to new questions in early stage design exploration. Where should the (analog and mixed) signal processing of that data occur? Should it occur in a hardware implementation, which is natively faster but more costly in on-chip resources? Or in software, where inherent latency issues may exist? One key challenge we see is that hardware design and software design have inherently different workflows, and as a result, use different design tools and methodologies. This means SoC architects need to be fluent in both C and HDL, and the hardware/software co-design environments needed for both. Another key challenge is that this integration further exacerbates the functional, gate- or circuit-level, and final sign-off verification problems that have dogged designers for decades. Interestingly, designers facing either or both of these key challenges could benefit significantly from top-down design and verification methodologies.

Chaitow (Hillcrest Labs): In most sensor-based applications, data is ultimately processed in the digital realm so an analog to digital conversion has to occur somewhere in the system before the processing occurs. MEMS sensors measure tiny variations in capacitance, and amplification of that signal is necessary to allow sufficient swing in the signal to ensure a reasonable resolution. Typically the analog to digital conversion is performed at the sensor to allow for reduction of error in the measurement. Errors are generally present because of the presence of noise in the system, but the design of the sensing element and amplifiers have attributes that contribute to error. For a given sensing system minimizing the noise is therefore paramount. The power supply of the ADC needs to be carefully designed to minimize noise and the routing of analog signals between the sensors and the ADC requires careful layout. If the ADC is part of an MCU, then the power regulation of the ADC and the isolation of the analog front end from the digital side of the system is vital to ensure an effective sampling system.

As always with design there are many tradeoffs. A given analog MEMS supplier may be able to provide a superior

Source III, Inc.

INDUSTRY LEADING TOOLS LINKING SIMULATION AND ATPG TO TEST

Source III provides the industry's most comprehensive and cost-effective vector translation product (**VTRAN**®) which links simulation/ATPG vector data to ATE, a powerful vector analysis tool (**VCAP**®) and WGL/STIL waveform viewing/editing tool (**DFTView**®). With over 25 years of providing vector translation solutions to hundreds of electronics and semiconductor companies, Source III has gained a reputation for outstanding customer service and high-quality tools.

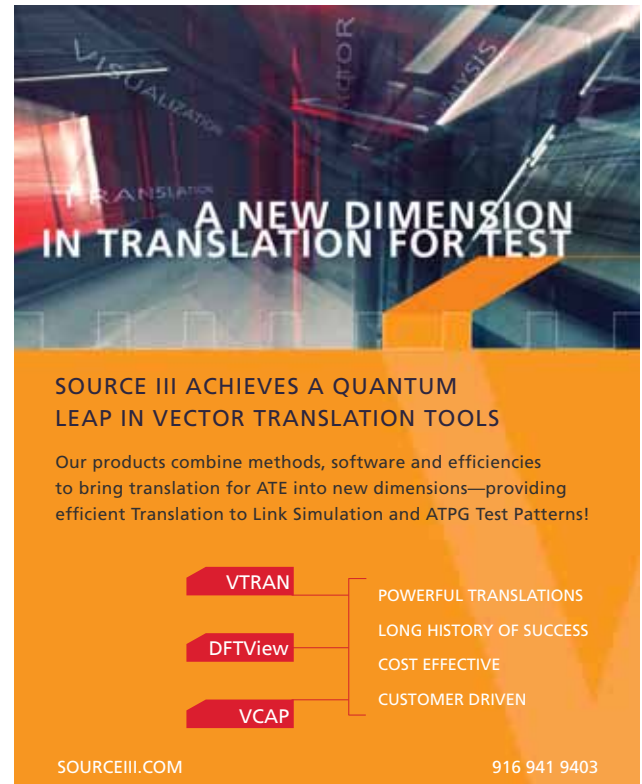
LINKING SIMULATION/ATPG TO TEST

Translating event-based simulation data (VCD and EVCD) to cycle-based ATE device testers can be complex and error-prone. Running under the Vtran User Interface (VUI), VTRAN, VCAP and DFTView provide a powerful set of tools to efficiently and flexibly accomplish this task. Simulation or ATPG files are automatically read, optionally have various conditioning and editing processes applied to them, and finally are cyclized and formatted for a target ATE. The resulting test programs can then be viewed graphically and optionally translated back into a Verilog or VHDL testbench for a validation simulation prior to being loaded onto the tester.

POWERFUL VECTOR PROCESSING

VTRAN, VCAP and DFTView gives test engineers an unparalleled level of power and flexibility during the vector translation process – often a huge benefit when trying to get working test programs. Some of the key processing features include;

- ◆ Analysis of timing and waveforms in VCD/EVCD files
- ◆ Vector masking using numerous masking algorithms
- ◆ Adding new signals with state values a function of other signals
- ◆ Re-arranging signal ordering and name aliasing
- ◆ Inserting vector repeat functions and statement insertions
- ◆ Tracking of vector number, cycle number and scan bit counts
- ◆ Waveform viewing of STIL, WGL and VTRAN-generated ATE test programs



**A NEW DIMENSION
IN TRANSLATION FOR TEST**

**SOURCE III ACHIEVES A QUANTUM
LEAP IN VECTOR TRANSLATION TOOLS**

Our products combine methods, software and efficiencies to bring translation for ATE into new dimensions—providing efficient Translation to Link Simulation and ATPG Test Patterns!

VTRAN — POWERFUL TRANSLATIONS
DFTView — LONG HISTORY OF SUCCESS
VCAP — COST EFFECTIVE
CUSTOMER DRIVEN

SOURCEIII.COM 916 941 9403

Visit us at www.sourceiii.com to access more information, review technical documents and download evaluation software to try it out for yourself. Source III's technical support staff becomes an extension to your design and test groups to ensure your success. You get proven, reliable products with a full-time technical staff to support you. No other tools on the market can offer the level of scope, features, performance, support and cost that we offer. Learn more today!

CONTACT INFORMATION



CUSTOMER DRIVEN TEST SOLUTIONS

Source III, Inc.
3941 Park Drive, Suite #20-342,
El Dorado Hills, CA 95762
United States
916.941.9403 Telephone
916.941.9403 Fax
corps@sourceiii.com
www.sourceiii.com

PLL and DLL Hard Macros

Organizations: GSA

True Circuits' complete family of standardized, silicon-proven, low-jitter PLL and DLL hard macros spans nearly all performance points and features typically requested by ASIC, FPGA and SoC designers. Using robust state-of-the-art circuits, a methodical and proven design strategy, and close associations with the world's leading fabs, IDMs and design services companies, True Circuits is able to quickly and reliably create new and innovative designs in a variety of advanced process technologies.

True Circuits' clock generator, spread spectrum, deskew and general purpose PLLs support wide operating frequency ranges and multiplication factors over which they deliver optimal jitter performance, avoiding the cost and complexity of licensing multiple point-solution PLLs from other vendors.

The clock generator PLL supports fractional-N frequency multiplication and is available with multi-phase outputs. The spread spectrum PLL allows precise adjustment of spreading rate and depth, with 4 bits of precise fractional-N control, optional multi-phase outputs, and very low long-term jitter. The general purpose PLL is ideal for low power, area sensitive and low cost applications.


True Circuits' multi-slave DLL delays a set of signals by precise and adjustable fractions of a reference clock cycle independent of voltage and temperature. The multi-phase DLL generates precise multi-phase clocks. These DLLs have flexible form factors for easier integration and are ideal for high-speed DDR and other interface applications.

FEATURES & BENEFITS


- ◆ Clock Generator, Spread Spectrum and Deskew PLLs
- ◆ General Purpose PLLs
- ◆ Multi Slave and Multi Phase DLLs
- ◆ Support the latest ARM processor, DDR, SerDes and audio/video standards
- ◆ Optimized for size, power and performance
- ◆ Low-jitter and very process tolerant
- ◆ Large multiplication factors (1-4096)
- ◆ Wide output frequency ranges
- ◆ Pin Programmable
- ◆ Customization available upon request




CLOCK GENERATOR, SPREAD SPECTRUM AND GENERAL PURPOSE PLLS




DESKEW PLL



DDR DLL



MULTI PHASE DLL



TECHNICAL SPECS

- ◆ Available in TSMC, GLOBALFOUNDRIES and UMC processes from 180nm to 16nm
- ◆ GDSII and LVS Spice netlists
- ◆ Behavioral, synthesis and LEF models
- ◆ Extensive user documentation
- ◆ Integration support to ensure successful tape outs

INDUSTRIES SERVED

Audio, Automotive, Biometric, Computer Services, Consumer Products, Defense, Industrial, Entertainment, Healthcare, Networking, Security, Telephony, Telecommunications, Wireless

CONTACT INFORMATION



True Circuits, Inc.
 4300 El Camino Real
 Suite 200
 Los Altos, CA 94022
 United States
 650.949.3400 Telephone
 650.949.3434 Fax
 timing@truecircuits.com
 www.truecircuits.com

From Data to Information to Knowledge

During my conversation with Luco Lanza late last year, we exchanged observations on how computing power, both at the hardware and the system levels, had progressed since 1968, the year when we both started working full time in the electronics field. And how much was left to do in order to achieve knowledge based computing.

Observing what computers do, we recognized that computers are now very good at collecting and processing data. In fact the concept of the IoT is based on the capability to collect and process a variety of data types in a distributed manner. The plan is, of course, to turn the data into information.

Definitely there are examples of computer systems that process data and generate information. Financial systems provide profit and loss reports using income and expense data for example, and logic simulators provide the behavior of a system by using inputs and output values. But information is not knowledge.

Knowledge requires understanding, and computers do not understand information. The achievement of knowledge requires the correlation and synthesis of information, sometimes from disparate sources, in order to generate understanding of the information and thus abstract knowledge. Knowledge is also a derivative of previous knowledge and cannot always be generated simply by digesting the information presented. The human brain associates the information to be processed with knowledge already available in a learning process that assigns the proper value and quality to the information presented in order to construct a value judgment and associative analysis that generates the new knowledge. What follows are some of my thoughts since that dialogue.

As an aside I note that unfortunately many education systems give students information but not the tools to generate knowledge. Students are thought to give the correct answer to a question, but not to derive the correct answer from a set of information and their own existing knowledge.

We have not been able to recreate the knowledge generating processes successfully with a computer, but nothing says that it will not be possible in the future. As computing power increases and new computer architectures are created, I know we will be able to automate the generation of knowledge. As far as EDA is concerned I will offer just one example.

A knowledge based EDA tool would develop a SoC from a set of specifications. Clearly if the specifications are erroneous the resulting SoC will behave differently than expected, but even this eventuality would help in improving the next design because it would provide information to those that develop a knowledge based verification system.

When we achieve this goal humanity will finally be free to dedicate its intellectual and emotional resources to address those problems that derive directly from what humans are, and prevent most of them, instead of having to solve them.

At this moment I still see a lot of indeterminism with respect of the most popular topic in our field: the IoT. Are we on the right track in the development of the IoT? Are we generating the correct environment to learn to handle information in a knowledge generating manner? To even attempt such a task we need to solve not just technical problems, but financial, behavioral, and political issues as well. The communication links to arrive to a solution are either non-existent or weak. Just think of the existing debate regarding "free internet". Financial requirements demand a redefinition of the internet. From a communication mechanism akin to a utility, to a special purpose device used in selective manners defined by price. How would a hierarchical internet defined by financial parameters modify the IoT architecture? Politicians and some business interests do not seem to care and engineers will have to patch things up later. In this case we do not have knowledge. ♦

Gabe Moretti is Senior Editor at Chip Design.



The ConFab®

MAY 19-22, 2015
ENCORE AT THE WYNN
LAS VEGAS
www.theconfab.com

SAVE 20%!
Sign up by Feb. 27th and Save

THE SEMICONDUCTOR MANUFACTURING & DESIGN INDUSTRY'S PREMIER ANNUAL EVENT

The ConFab, presented by Solid State Technology, draws influential executives & esteemed thought leaders from organizations such as:

GLOBALFOUNDRIES
TSMC
G450C

INTEL
AMD
QUALCOMM

TEXAS INSTRUMENTS
SAMSUNG
PLUS MANY OTHERS

Confirm your Participation for the 11th Annual ConFab Today

SIGN UP TODAY AND SAVE!
Sign up by February 27, 2015 and Save up to 20%

For more information, contact:
Sabrina Straub
Sales Manager
+1.603.770.6569
sstraub@extensionmedia.com



Presented by:

Solid State
TECHNOLOGY®

www.theconfab.com

Owned and Produced by:

Extension
MEDIA

UNLEASH THE POWER

Realize the true potential of your products with the functionality and performance of True Circuits timing IP

SPEED

Wide output frequency ranges
Large multiplication factors (1-4096)
GHz frequencies for Gb/s data rates

PRECISION

Low jitter
Pin programmable
Process independent

PERFECTION

High quality, silicon proven IP
Comprehensive documentation
Easy to integrate

PLL and DLL Hard Macros

PERFECT TIMING REQUIRES SPEED AND PRECISION, PARTICULARLY IN THE UNFORGIVING WORLD OF DIGITAL CHIPS...

True Circuits offers a complete family of standardized high performance and general purpose PLL and DLL hard macros that have been specifically designed to meet the precise timing requirements of the entire line of ARM processor cores and the latest DDR, SerDes, audio, video and other interface standards.

These PLLs and DLLs are available for delivery in a range of frequencies, multiplication factors, sizes and functions in TSMC, GLOBALFOUNDRIES, and UMC processes from 180nm to 16nm.

With so much on the line and so much complexity to manage, realize the true potential of your products with the speed, precision and perfection of True Circuits timing IP.



TRUE CIRCUITS INC.™ | *timing is everything*

WWW.TRUECIRCUITS.COM/TIMING