Extension MEDIA

Chip Design

Winter 2014

Tools, Technologies & Methodologies

EDA INDUSTRY PREDICTIONS FOR 2014

Design for Yield Trends

Shared Viewpoint for Integration

New Market Challenges

IDEM and ASIP Coverage - From 3D Transistors to Systems

Chip Design Online Network: • Chipdesignmag.com • SLDCommunity.com • LowPowerEngineering.com

call for participation

Deadline Early Registration fee:
FRIDAY, FEBRUARY 28, 2014

DATE 14

DATE 2014 ★ big ★ experienced ★ open minded ★ successful competent ★ eager to learn ★ business driven

ICC, Dresden, Germany

Event 24–28 March **Exhibition** 25–27 March

Design, Automation & Test in Europe

The European System Design Show

From Systems-on-Chip to Embedded Computing




www.date-conference.com

follow us on facebook

call for participation

DATE 14

DATE is a leading international event and unique networking opportunity for design and engineering of **Systems-on-Chip, Systems-on-Board and Embedded Systems Software.**

the conference

The **conference** addresses all aspects of research into technologies for electronic and (embedded) systems engineering. It covers the design process, test, and tools for design automation of electronic products ranging from integrated circuits to distributed large-scale systems. This includes both hardware and embedded software design issues. The conference scope also includes the elaboration of design requirements and new architectures for challenging application fields such as telecom, wireless communications, multimedia, healthcare and automotive systems. Persons involved in innovative industrial designs are particularly encouraged to submit papers to foster the feedback from design to research. Panels, hot-topic sessions and embedded tutorials highlight and inform about emerging topics.

Special Days in the programme will focus on two areas bringing new challenges to the system design community:

Advancing Electronics beyond CMOS As the issues associated with CMOS scaling become harder and more costly to solve, this special day will cover the latest trends towards alternative devices and paradigms for computing and data acquisition, storage and transport. While the potential for innovation is tremendous, the main challenge is to understand how to choose from such a broad spectrum of approaches. Important topics include nanoscale switching devices and computing nanofabrics, trends in memory technologies in hybrid 3D integration, flexible and organic electronics, new state variable vectors (spintronics, photonics) and interfaces to the natural world (sensor networks, data analysis, displays, MEMS).

System-Level Design to reflect current industrial practices as well as present recent advances in System-Level Design research. A particular emphasis will be on ultra-low power design and modeling at multiple abstraction levels, virtual platforms for software development and architecture design of MPSoCs and system integration through cosimulation. We will also take a look at high-level synthesis from different input languages as well as software code generation techniques. Important topics also include multi-core enablement for safety-critical and real-time embedded systems as well as accelerator-rich design for the fight against dark silicon.

On the first day of the DATE event, half-day in-depth technical **tutorials** are given by leading experts in their respective fields. The tutorials are well suited for researchers, tool developers and system designers.

Friday Workshops concentrate on specialised and novel topics.
The preliminary programm is available online at
www.date-conference.com

the exhibition

The conference is complemented by an exhibition and unique networking opportunity for vendors of tools and services for hardware and embedded software for the design, development and test of Systems-on-Chip, IPs, Embedded Systems, ASICs, FPGAs and PCBs including a broad range of design reuse technologies and services.

To inform attendees on commercial and design related topics, there will be a full programme in the Exhibition Theatre which will combine presentations by exhibiting companies, best-practice reports by industry leaders on their latest design projects and selected conference special sessions.

registration

The registration to the conference is solely possible online via the online registration platform.

On the homepage www.date-conference.com you will find an overview of the registration fees, the link to the online registration platform as well as hotel and travel offers.

Accommodation can be booked during the online registration process (recommended) or afterwards via the separate online reservation link.

Deadline Early Registration fee: Friday, February 28, 2014

ICC Dresden Germany

The International Congress Center (ICC) Dresden and the adjoining Maritim hotel are located on the banks of the river Elbe with a spectacular view at the historic city centre. The ICC with its modern architecture offers all facilities and technical requirements for hosting large international conferences. It is easily accessible from Dresden Airport by public transport and within walking distance to downtown Dresden and all famous sights.

unique Dresden

The Saxony state capital is characterised by its historic flair and the numerous cultural sights attracting visitors from all over the world every year. It is a modern and vibrant city with short distances and a cutting-edge technology sector. Silicon Saxony is a registered industry association of nearly 300 companies in the microelectronics and related sectors in Saxony, Germany, with around 40,000 employees. The majority of those firms are situated in the north of Dresden. DATE is cooperating with Silicon Saxony to maximise quality visitor attendance to the show and increase visibility of the conference in the region. Dresden therefore states an excellent venue for DATE 14.

sponsors

The event is sponsored by the European Design and Automation Association, the EDA Consortium, the IEEE Council on EDA, ECSI, ACM – SIGDA, and RAS. In cooperation with ACM – SIGBED, IEEE Solid-State Circuits Society (SSCS), IFIP and IET.

contact

Event Secretariat K.I.T. Group GmbH, Münzgasse 2, 01067 Dresden, Germany
Phone: +49 351 4967 312, Fax: + 49 351 4956 116, date14@kit-group.org

General Chair Gerhard P. Fettweis, TU Dresden, Germany
gerhard.fettweis@tu-dresden.de

Programme Chair Luca Fanucci, University of Pisa, Italy
luca.fanucci@iet.unipi.it

To add your name to the mailing list go to the website:
www.date-conference.com.

Chip Design
Tools, Technologies & Methodologies

IN THIS ISSUE

Cover Story

Features

Departments

Methods/EDA Tools

Semiconductor Technologies

Chip Design
Tools, Technologies & Methodologies

www.chipdesignmag.com

Publisher & Sales Director
Karen Popp (415) 255-0390 x19
kpopp@extensionmedia.com

EDITORIAL STAFF

Chief Content Officer
John Blyler (503) 614-1082
jblyler@extensionmedia.com

Associate Editor—China
Jane Lin-Li

Embedded Editor
Cheryl Coupe, Chris Ciufo

Contributing Editors
Caroline Hayes, Hamilton Carter, Sara Verbruggen, Gabe Moretti, Cheryl Ajluni, Dave Bursky, Cheryl Coupé, Chris Ciufo, David Lammers

Editorial Board
Cheryl Ajluni, Technical Consultant, Custom Media Solutions • Karen Bartleson, Standards Program Manager, Synopsys • Chuck Byers, Professor, Santa Clara University • Kathryn Kranen, CEO, Jasper Design Automation • Tom Moxon, Consultant, Moxon Design • Walter Ng, Senior Director, Design Services, Chartered Semiconductor • Scott Sandler, CEO, Novas Software • Steve Schulz, President, SI2 • Adam Traidman, Chip Estimate

CREATIVE/PRODUCTION/ONLINE

Production Manager
Spryte Heithecker

Media Coordinator
Yishian Yao

Graphic Designers
Nicky Jacobson & Jacob Ewing

Senior Web Developer
Slava Dotsenko
Mariam Moattari

SALES STAFF

Advertising and Reprints
Karen Popp (415) 255-0390 x19
kpopp@extensionmedia.com

Audience Development / Circulation

Extension MEDIA

Jenna Johnson • jjohnson@extensionmedia.com

President
Vince Ridley (415) 255-0390 x18
vridley@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.

By John Blyler, Vice President/Chief Content Officer

Sensor Swarm Proponent Predicted IoT Technical Barrier

Earlier warnings by Berkeley's Dr. Rabaey are echoed in ARM-Economist report.

Earlier warnings by Dr. Jan Rabaey of US Berkeley about complexity issues support a recent ARM-Economist report on barriers of immature standards and the high cost of networking infrastructures for IoT.

Years ago, I moderated a Cadence Design Systems panel where Dr. Rabaey talked about the coming age of wireless sensor swarms (a precursor to IoT). He noted the growing complexity of sensory networks, adding that, "if you ignore system-level design, you're toast."

Low-power panel, from left: John Blyler, Ted Vucurevich, Nikhil Jayaram, Juan Antonio Carballo, Jan Rabaey, and Carl Guardino

A similar warning has been put forth in an ARM sponsored report by the Economist Intelligence Unit, titled; "The Internet of Things Business Index: A quiet revolution gathers pace." Among the many findings of the report was a caution about connectivity: "The IoT will not flourish without genuine co-operation. Turning 50bn so-called smart things into a global network requires businesses to agree on standards for inter-connectivity and data sharing."

THE IOT WILL NOT FLOURISH WITHOUT GENUINE CO-OPERATION. TURNING 50BN SO CALLED SMART THINGS INTO A GLOBAL NETWORK REQUIRES BUSINESSES TO AGREE ON STANDARDS FOR INTER-CONNECTIVITY AND DATA SHARING.

Two of the top five barriers for companies hoping to increase the use of IoT were; 1) the immaturity of industry standards around IoT, and 2) the high costs of required investments of IoT infrastructure. Both of these barriers relate directly to networking/connectivity issues, including the design of the overall system – from the analog sensor to the digital MCU and up through the individual wired/wireless networks to the cloud.

Echoing Dr. Rabaey's words from many years ago, the technical success of IoT will depend upon the smooth interconnection between the many small sensor networks of individually connected things to the big network of connected things that extended across industries and organizations. This inter-connectivity will involved both wired and wireless protocols and standards.

Once this system-level infrastructure is secure, then the business success of IoT will depend upon the free flow of information across all of these networks. But part of the business allure of IoT is the services that companies can charge, e.g., the sale of data. This business model that supports both the free flow of information and the sale of a portion of that information has yet to be developed.

Both of these challenges can – and will – be met, but not without dealing with complex issues on both the technical and business fronts of IoT

John Blyler covers today's latest high-tech, R&D and even science fiction in blogs, magazine articles, books and videos. He is an experienced physicist, engineer, journalist, author and professor who continues to speak at major conferences and before the camera on. John is the Vice President, Chief Content Office for Extensionmedia, which includes the brands Chip Design, Solid State Technology, Embedded Intel and others. He holds a BS in Engineering Physics and a MSEE. John plays the piano and holds a black belt in TKD.

By Pete Singer, EiC, Solid State Technology

ISS-2014: Gadget Magic and Need for Innovation

Rick Wallace, president and CEO of KLA-Tencor, provided the keynote talk at the SEMI Industry Strategy Symposium (ISS) this year, held Jan 12-15 in Half Moon Bay, CA. He said he believes the semiconductor industry might be facing a "Concorde" moment, referring to the demise of supersonic passenger transport, the last flight of which was on 24 October 2003. "That failed not because of technology but because of economics," Wallace said. He sees a similar challenge coming down the road for continued scaling. "Moore's Law is much more likely to die in the boardroom than the laboratory," said.

Wallace also spoke about "The Road Less Traveled," seeming to indicate that the more traveled one is that of consolidation, which Wallace said leads to "losses in agility, flexibility and innovation." He said larger firms are not effective at driving innovation although they are effective at driving continuous improvement. "It's tough to see how a large scale merger makes a company better," he said. "Some firms will be too big to fail but my fear is that they will become too big to innovate."

THE MORE TRAVELED (ROAD) IS THAT OF CONSOLIDATION, WHICH WALLACE SAID LEADS TO "LOSSES IN AGILITY, FLEXIBILITY AND INNOVATION." HE SAID LARGER FIRMS ARE NOT EFFECTIVE AT DRIVING INNOVATION

The solution he said is young people. "We need to attract the young talent if we want real innovation. The longer you're around the more you see what can't be done," he said.

Wallace told a story about explaining to his 10 year old daughter what his company by using the iPad as an example. His daughter thought about it and said she understood: it was the magic behind the gadget.

Part of attracting young people to the semiconductor industry is through education. After Rick's presentation, Denny McGuirk, president of SEMI, presented an award to Rick and to L.T. Guttadauro, president of the Fab Owners Association, in recognition of their work on SEMI's High Tech University (HTU). HTU is a career exploration program that encourages student interest in science, technology, engineering and match. Since 2001, the SEMI Foundation has delivered 143 programs to 4800 students and teachers worldwide.

Although some view the semiconductors as a commodity, hopefully efforts such as that of the HTU will explain the magic behind the gadget. "Who doesn't want to work on magic?" Wallace asked.

Pete Singer has been involved in technical journalism for more than 30 years and has written well over 150 articles on all aspects of semiconductor manufacturing and related industries, Pete holds a degree in Electrical Engineering from the University of Illinois and is a member of IEEE, Electrochemical Society, American Vacuum Society and Materials Research Society

By Gabe Moretti, Senior Editor/EDA

EDA Industry Predictions for 2014 – Part I and II

Projected growth in the third-party semiconductor intellectual-property (IP) market through 2017 may affect the direction and evolution of subsystem designs.

I always ask for predictions for the coming year, and generally get good response. But this year the volume of responses was so high that I could not possibly cover all of the material in one article. So I will use two articles, one week apart, to record the opinions submitted. This first section details the contributions of Andrew Yang of ANSYS – Apache Design, Mick Tegethoff of Berkeley Design Automation, Michel Munsey of Dassault Systèmes, Oz Levia from Jasper Design Automation, Joe Sawicki from Mentor Graphics, Grant Pierce and Jim Hogan from Sonics, and Bob Smith of Uniquify.

ANDREW YANG – ANSYS APACHE DESIGN

For 2014 and beyond, we'll see increased connectivity of the electronic devices that are pervasive in our world today. This trend will continue to drive the existing mobile market growth as well as make an impact on upcoming automotive electronics. The mobile market will be dominated by a handful of chip manufacturers and those companies that support the mobile ecosystem. The automotive market is a big consumer of electronics components that are part of a complex system that help improve safety and reliability, as well as provide users with real-time interaction with their surroundings.

For semiconductor companies to remain competitive in these markets, they will need to take a "system" view for their design and verification. The traditional silo-based methodology, where each component of the system is designed and analyzed independently can result in products with higher cost, poor quality, and schedule delay. An adoption of system-level simulation will allow engineers to carry out early system prototyping, analyze the interaction of each of the components, and achieve optimal design tradeoffs.

MICK TEGETHOFF – BERKELEY DESIGN AUTOMATION

FinFET Technology will dominate the landscape in semiconductor design and verification as more companies adopt the technology. FinFET is a revolutionary change to device fabrication and modeling, requiring a more complex SPICE model and challenging the existing circuit behavior "rules of thumb" on which experienced designers have relied for years with planar devices.

Designers of complex analog/RF circuits, including PLLs, ADCs, SerDes, and transceivers, will need to relearn the device behavior in these applications and to explore alternative architectures. As a result, design teams will have to rely more than ever on accurate circuit verification tools that are foundry-certified for FinFET technology and have the performance and capacity to handle complex circuits including physical effects such as device noise, complex parasitics, and process variability.

In memory applications, FinFET technology will continue to drive change and challenge the status quo of "relaxed accuracy" simulation for IP characterization. Design teams are realizing that it is no longer acceptable to tolerate 2–5% inaccuracy in memory IP characterization. They are looking for verification tools that can deliver SPICE-like accuracy in a time frame on a par with their current solutions.

However, accurate circuit verification alone will not be sufficient. The impact of FinFET devices and new circuit architectures in analog, RF, mixed-signal, and memory applications demand full confidence from design teams that their circuits will meet specifications across all operational, environmental, and process conditions. As a result, designers will need to perform an increased amount of intelligent, efficient, and effective circuit characterization at the block level and at the project level to ensure that their designs meet rigorous requirements prior to silicon.

MICHAEL MUNSEY – DASSAULT SYSTÈMES

We at Dassault Systèmes see a few key trends coming to the semiconductor industry in 2014.

1) Extreme Design Collaboration: Complexity and cost in IC design and manufacturing now demand that semiconductor vendors engage an ever broader, more diverse pool of specialist designers and engineers.

At the same time, total costs for designing a cutting-edge integrated circuit can top $100 million for just one project. Respins can drive these costs even higher, adding huge profitability risks to new projects.

Technology-enabled extreme collaboration, over and above that in traditional PLM, will be required to assure manufacturable, profitable designs. Why? Because defects arise at the interchange between designers. And with more designers and more complex projects, the risk of misperceptions and miscommunications increases.

Pressure for design teams to interlock using highly specialized collaboration technology will increase in parallel with the financial risk of new semiconductor design projects.

2) Enterprise IP management: The move towards more platform-based designs in order to meet shortening time to market windows, application driven designs, and the increasing cost of producing new semiconductor devices, will explode the market for IP and create a new market for enterprise IP management.

The deeper insight is how that IP will be acquired, used, configured, validated and otherwise managed. The challenges will be (1) building a intelligent process that enables project managers to evaluate the lowest cost IP blocks quickly and effectively; (2) managing the licensed IP so that configuration, integration and validation know-how is captured and is easily reused; and (3) ensuring that licensing and export compliance attributes of each licensed block of IP are visible to design decision makers.

3) Flexible Design to Manufacturing: In 2011, the Japanese earthquake forced a leading semiconductor company to cease manufacturing operations because their foundry was located close to Fukushima. That earthquake and the floods in Thailand have awakened semiconductor vendors to the stark reality that global supply chains can be dramatically and unexpectedly disrupted without any prior notice.

At the same time, with increased fragmentation and specialization occurring within the design and supply chain for integrated circuit, cross chain information automation will be mission-critical.

Examples of issues that will require IT advances are (1) the increasing variations in how IP is transferred down the supply chain. It could be a file, a wafer, a die or a packaged IC – yet vendors will need to handle all options with equal efficiency to maximize profitability; and (2) the flexible packaging of an IC design for capture into ERP systems will become mandatory, in order to enable the necessary downstream supply chain flexibility.

OZ LEVIA – JASPER DESIGN AUTOMATION

There are a few points that we at Jasper consider important for 2014.

1) Low power design and verification will continue to be a main challenge for SoC designers.

2) Heterogeneous multi-processor designs will continue to grow. Issue such as fabric and NOC design and verification will dominate.

3) The segments that will drive the semiconductor markets will likely continue to be in the mobile space(s) – phones, tablets, etc. But the server segment will also continue to increase in importance.

4) Process will continue to evolve, but there is a lot of head room in current processes before we run out steam.

5) Consolidation will continue in the semiconductor market. More important, the strong will get stronger and the weak will get weaker. Increasingly this is a winner takes all market and we will see a big divide between innovators and leaders and laggards.

6) EDA will continue to see consolidation. Large EDA vendors will continue increasing investments in SIP, and Verification technologies. We will not see a new radically different technology or methodology. The total amount of investments In the EDA industry will continue to be low.

7) EDA will grow at slow pace, but Verification, Emulation and SIP will grow faster then other segments.

JOSEPH SAWICKI – MENTOR GRAPHICS

FinFETs will move from early technology development to early adopter designs. Over the last year, the major foundry

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

GSA

12400 Coit Road, Suite 650 | Dallas, TX 75251 | 888-322-5195 | T 972-866-7579 | F 972-239-2292 | www.gsaglobal.org

ecosystems moved from alpha to production status for 16/14nm with its dual challenges of double patterning and FinFET. Fabless customers are just beginning to implement their first test chip tape-outs for 16 /14 nm, and 2014 will see most of the 20 nm early-adopter customers also preparing their first 16 nm/14 nm test chips.

FinFETs are driving a need for more accurate extraction tools, and EDA vendors are turning to 3D field solver technology to provide it. The trick is to also provide high performance that can deliver quick turnaround time even as the number of required extraction corners jumps from 5 to 15 and the number of gates doubles or triples.

Test data and diagnosis of test fail data will play an increasingly important role in the ramp of new FinFET technologies. The industry will face new challenges as traditional approaches to failure analysis and defect isolation struggle to keep pace with changes in transistor structures. The opportunity is for software-based diagnosis techniques that leverage ATPG test fail data to pick up the slack and provide more accurate resolution for failure and yield analysis engineers.

16/14nm will also require more advanced litho hotspot checking and more complex and accurate fill structures to help ensure planarity and to also help deal with issues in etch, lithography, stress and rapid thermal annealing (RTA) processes.

In parallel with the production ramp at 20 nm, and 16 nm/14 nm test chips, 2014 will see the expansion of work across the ecosystem for 10 nm. Early process development and EDA tool development for 10 nm began in 2012, ramped up in intensity in 2013, and will be full speed ahead in 2014.

Hardware emulation has transitioned from the engineering lab to the datacenter where today's virtual lab enables peripheral devices such as PCIe, USB, and Ethernet to exist in virtual space without specialized hardware or a maze of I/O cables. A virtual environment permits instant reconfiguration of the emulator for any design or project team and access by more users, and access from anywhere in the world, resulting in higher utilization and lower overall costs.

The virtual lab is also enabling increased verification coverage of SoC software and hardware, supporting end-to-end validation of SW drivers, for example. Hardware emulation is now employed throughout the entire mobile device supply chain, including embedded processor and graphics IP suppliers, mobile chip developers, and mobile phone and tablet teams. Embedded SW validation and debug will be the real growth engine driving the emulation business.

The Internet of Things (IoT) will add an entirely new level of information sources, allowing us to interact with and pull data from the things around us. The ability to control the state of virtually anything will change how we manage and interact with the world. The home, the factory, transportation, energy, food and many other aspects of life will be impacted and could lead to a new era of productivity increases and wealth creation.

Accordingly, we'll see continued growth in the MEMS market driven by sensors for mobile phones, automobiles, and medical monitoring, and we'll see silicon photonics solutions being implemented in data and communications centers to provide higher bandwidth backplane connectivity in addition to their current use in fiber termination.

Semiconductor systems enabling the IoT trend will need to respond to difficult cost, size and energy constraints to drive real ubiquity. For example, we'll need 3D packaging implementations that are an order of magnitude cheaper than current offerings. We'll need a better ways to model complex system effects, putting a premium on tools that enable design and verification at the system level, and engineers that can use them. Cost constraints will also drive innovation in test to ensure that multi-die package test doesn't explode part cost. Moreover, once we move from data to actually interacting with the real world analog/mixed signal, MEMS and other sensors role in the semiconductor solution will become much greater.

GRANT PIERCE AND JIM HOGAN – SONICS

For a hint at what's to come in the technology sector as a whole and the EDA and IP industries specifically, let's first look at the global macro-economic situation. The single greatest macro-economic factor impacting the technology sector is energy. Electronic products need energy to work. Electronic designers and manufacturers need energy to do their jobs. In the recent past, energy has been expensive to produce, particularly in the US market due to our reliance on foreign oil imports. Today in the US, the cost of producing energy is falling while consumption is slowing. The US is on a path to energy self-sufficiency according to the Energy Department's annual outlook. By 2015, domestic oil output is on track to surpass its peak set in 1970.

What does cheaper energy imply for the technology industry? More investment. Less money spent on purchasing energy abroad means more capital available to fund new ventures at home and around the world. The recovery of US financial markets is also restoring investors' confidence in earning higher ROI through public offerings. As investors begin to take more risk and inject sorely needed capital into the technology sector, we expect to see a surge in new startups. EDA and IP industries will participate in this "re-birth" because they are critical to the success of technology sector as enabling technologies.

For an understanding of where the semiconductor IP business is going, let's look at consumer technology. Who are the leaders in the consumer technology business today? Apple, Google, Samsung, Amazon, and perhaps a few others. Why? Because they possess semiconductor knowledge coupled with software expertise. In the case of Apple, for example, they also own content and its distribution, which makes them extremely profitable with higher recurring revenues and better margins. Content is king and the world is becoming application-centric. Software apps are content. Semiconductor IP is content. Those who own content, its publication and distribution, will thrive.

In the near term, the semiconductor IP business will continue to consolidate as major players compete to build and acquire broader content portfolios. For example, witness the recent Avago/LSI and Intel/Mindspeed deals. App-happy consumers have an insatiable appetite for the latest and greatest content and devices. Consumer technology product lifecycles place immense pressure on chip and system designers when developing and verifying the flexible hardware platforms that run these apps. Among their many important considerations are functionality, performance, power, security, and cost. System architectures and software definition and control are becoming the dominant source of product differentiation rather than hardware. The need for semiconductor IP that addresses these trends and accelerates time-to-volume production is growing. The need for EDA tools that help designers successfully use and efficiently reuse IP is also growing.

So what are the market opportunities for new IP and tool companies in the coming years? These days, talk about the Internet of Things (IoT) is plentiful and there will be many different types of IP in this sensor-oriented market space. Perhaps, the most interesting and promising of these IoT IP technologies will address our growing concerns about health and quality of life. The rise of wearable technologies that help monitor our vital signs and treat chronic health conditions promises to extend our human survival rate beyond 100 years. As these technologies progress, surely the "Bionic Man" will become common place in the not-too-distant future. Personally, and being members of the aging "Baby Boomer" generation, we hope that it happens sooner rather than later!

BOB SMITH – UNIQUIFY

I spent a good deal of 2013 traveling around the globe doing a series of seminars on double data rate (DDR) synchronous dynamic random-access memory (SDRAM), the ubiquitous class of memory chips. The seminars were meant to promote the fastest, smallest and lowest power state-of-the-art adaptive DDR IP technology. They highlighted how it can be used to enhance design speed and configured to minimize the design footprint and hit increasingly smaller low-power targets.

While marketing and promotion was on the agenda, the seminars were a great way to check in with designers to better understand their current DDR challenges and identify a few trends that will emerge in 2014. What we learned may be a surprise to more than a few semiconductor industry watchers and offers some tantalizing predictions for next year.

The biggest surprise was hearing designers confirm plans to go directly to LPDDR4 (that is low-power DDR4, the latest JEDEC standard) and skip LPDDR3. The reasons are varied, but most noted that they're getting greater gains in performance and low power by jumping to LPDDR4, especially important for mobile applications. According to JEDEC, the LPDDR4 architecture was designed to be power neutral, offer 2X bandwidth performance over previous generations, with low pin-count and low cost. It's also backward compatible.

Even though many of the designers we heard from agreed that DDR3 is now mainstream, even more are starting projects based on DDR4. Some are motivated to move to DDR4 even without the need for extra performance for a practical and cost-effective reason. If they have a product with a long lifetime of five years or more, they are concerned that the DDR3 memory will cost more than DDR4 at some point. They have a choice: either build in the DDR4 now in anticipation or look for combination IP that handles both DDR3/4 in one IP. Many have chosen to do the former.

One final prediction I offer for 2014 is that 28nm is the technology node that will be around for a long time to come. Larger semiconductor companies, however, are starting new projects at 14/16 nm, taking advantage of the emerging FinFET technology.

According to my worldwide sources, memories and FinFET will dominate the discussion in 2014, which means it will be a lively year.

PART 2

BERNARD MURPHY – ATRENTA

"Smart" will be the dominant watchword for semiconductors in 2014. We'll see the maturing of biometric identification technologies, driven by security needs for smart payment on phones, and an increase in smart-home applications. An example of cool applications? Well, we'll toss our clunky 20th-century remote controls, and manage our smart TV with an app on our phone or tablet, which will, among a host of other functions, allow point / text input to your center of living entertainment system – your smart TV. We'll see indoor GPS products, enabling the mobile user to navigate shopping malls – an application with significant market potential. We'll see new opportunities for Bluetooth or WiFi positioning, 3D image recognition and other technologies.

In 2014 smart phones will be the dominant driver for semiconductor growth. The smart phone industry will grow but will be constrained by adoption costs and immaturity. But I foresee that one of the biggest emerging technologies will be smart cards. Although common for many years in Europe, this technology has been delayed in the US for lack of infrastructure and security concerns. Now check out businesses near you with new card readers. Chances are they have a slot at the bottom as well as one at the side. That bottom slot is for smart cards. Slated for widespread introduction in 2015, smart card technologies will explode due to high demand.

The EDA industry in 2014 will continue to see implementation tools challenged by conflicting requirements of technology advances against the shrinking customer-base that can afford the costs at these nodes. Only a fundamental breakthrough enabling affordability will affect significant change in these tools. Front-end design will continue to enjoy robust growth, especially around tools to manage, analyze and debug SoCs based on multi-sourced IPs – the dominant design platform today. Software-based analysis and verification of SoCs will be an upcoming trend, which will largely skip over traditional testbench-based verification. This will likely spur innovation around static hookup checking for the SoC assembly, and methods to connect software use-cases to implementation characteristics such as power and enhanced debug tools to bridge the gap between observed software behavior and underlying implementation problems.

THOMAS L. ANDERSON – BREKER

Electronic design automation (EDA) and embedded systems have long been sibling markets and technologies, but they are increasingly drawing closer and starting to merge. 2014 will see this trend continue and even accelerate. The catalyst is that almost all significant semiconductor designs user a system-on-chip (SoC) architecture, in which one or more embedded processors lie at the heart of the functionality. Embedded processors need embedded programs, and so the link between the two worlds is growing tighter every year.

One significant driver for this evolution is the gap between simulation testbenches and hardware/software co-verification using simulation, emulation or prototypes. The popular Universal Verification Methodology (UVM) standard provides no links between testbenches and code running in the embedded processors. The UVM has other limitations at the full-SoC level, but verification teams generally run at least some minimal testbench-based simulations to verify that the IP blocks are interconnected properly.

The next step is often running production code on the SoC processors, another link between EDA and embedded. It is usually impractical to boot an operating system in simulation, so usually the verification team moves on to simulation acceleration or emulation. The embedded team is more involved during emulation, and usually in the driver's seat by the time that the production code is running on FPGA prototypes. The line between the verification team (part of traditional EDA) and the embedded engineers becomes fuzzy.

When the actual silicon arrives from the foundry, most SoC suppliers have a dedicated validation team. This team has the explicit goal of booting the operating system and running production software, including end-user applications, in the lab. However, this rarely works when the chip is first powered up. The complexity and limited debug features of production code lead the validation team to hand-write diagnostics that incrementally validate and bring up sections of the chip. The goal is to find any lurking hardware bugs before trying to run production software.

Reuse Delivered

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**

OCP International Partnership

Closer alignment between EDA and embedded will lead to two important improvements in 2014. First, the simulation gap will be filled by automatically generated multi-threaded, multi-processor C test cases that leverage portions of the UVM testbench. These test cases stress the design far more effectively than UVM testbenches, hand-written tests, or even production software (which is not designed to find bugs). Tools exist today to generate such test cases from graph-based scenario models capturing the design and verification intent for the SoC.

Second, the validation team will be able to use these same scenario models to automatically generated multi-threaded, multi-processor C test cases to run on silicon and replace their hand-written diagnostics. This establishes a continuum between the domains of EDA, embedded systems, and silicon validation. Scenario models can generate test cases for simulation, simulation acceleration, emulation, FPGA prototyping, and actual silicon in the lab. These test cases will be the first embedded code to run at every one of these stages in 2014 SoC projects.

SHAWN MCCLOUD - CALYPTO

While verification now leverages high-level verification languages and techniques (i.e:, UVM/OVM and SystemVerilog) to boost productivity, design creation continues to rely on RTL methodologies originally deployed almost 20 years ago. The design flow needs to be geared toward creating bug-free RTL designs. This can be realized today by automating the generation of RTL from exhaustively verified C-based models. The C++/SystemC source code is essentially an executable spec. Because the C++/SystemC source code is more concise, it executes 1,000x–10,000x faster than RTL code, providing better coverage.

C and SystemC verification today is rudimentary, relying primarily on directed tests. These approaches lack the sophistication that hardware engineers employ at the RTL, including assertions, code coverage, functional coverage, and property-based verification. For a dependable HLS flow, you need to have a very robust verification methodology, and you need metrics and visibility. Fortunately, there is no need to re-invent the wheel when we can borrow concepts from the best practices of RTL verification.

Power analysis and optimization have evolved over the last two years, with more changes ahead. Even with conventional design flows there is still a lot more to be optimized on RTL designs. The reality is, when it comes to RTL power optimization, the scope of manual optimizations is relatively limited when factoring in time to market pressure and one's ability to predict the outcome of an RTL change for power. Designers have already started to embrace automated power optimization tools that analyze the sequential behavior of RTL designs to automatically shut down unused portions of a design through a technique called sequential clock gating. There's a lot more we can do by being smarter and by widening the scope of power analysis. Realizing this, companies will start to move away from the limitations of predefined power budgets targets toward a strategy that enables reducing power until the bell rings, and it's time for tape out.

BILL NEIFERT – CARBON

Any prediction of future advances in EDA has to include a discussion on meeting the needs of the software developer. This is hardly a new thing, of course. Software has been consuming a steadily increasing part of the design resources for a long time. EDA companies acknowledge this and discuss technologies as being "software-driven" or "enabling software development," but it seems that EDA companies have had a difficult time in delivering tools that enable software developers.

At the heart of this is the fundamental cost structure of how EDA tools have traditionally been sold and supported. An army of direct sales people and support staff can be easily supported when the average sales price of a tool is in the many tens or hundreds of thousands of dollars. This is the tried-and-true EDA model of selling to hardware engineers.

Software developoers, however, are accustomed to much lower cost, or even free tools. Furthermore, they expect these tools to work without multiple calls and hand-holding from their local AE.

In order to meet the needs of the software developers, EDA needs to change how it engages with them. It's not just a matter of price. Even the lowest-priced software won't be used if it doesn't meet the designer's needs or if it requires too much direct support. After all, unlike the hardware designers who need EDA tools to complete their job, a software programmer typically has multiple options to choose from. The platform of choice is generally the one that causes the least pain and that platform may be from an EDA provider. Or, it could just as likely be homegrown or even an older generation product.

Receive a Complimentary EDA & IP Industry Market Report

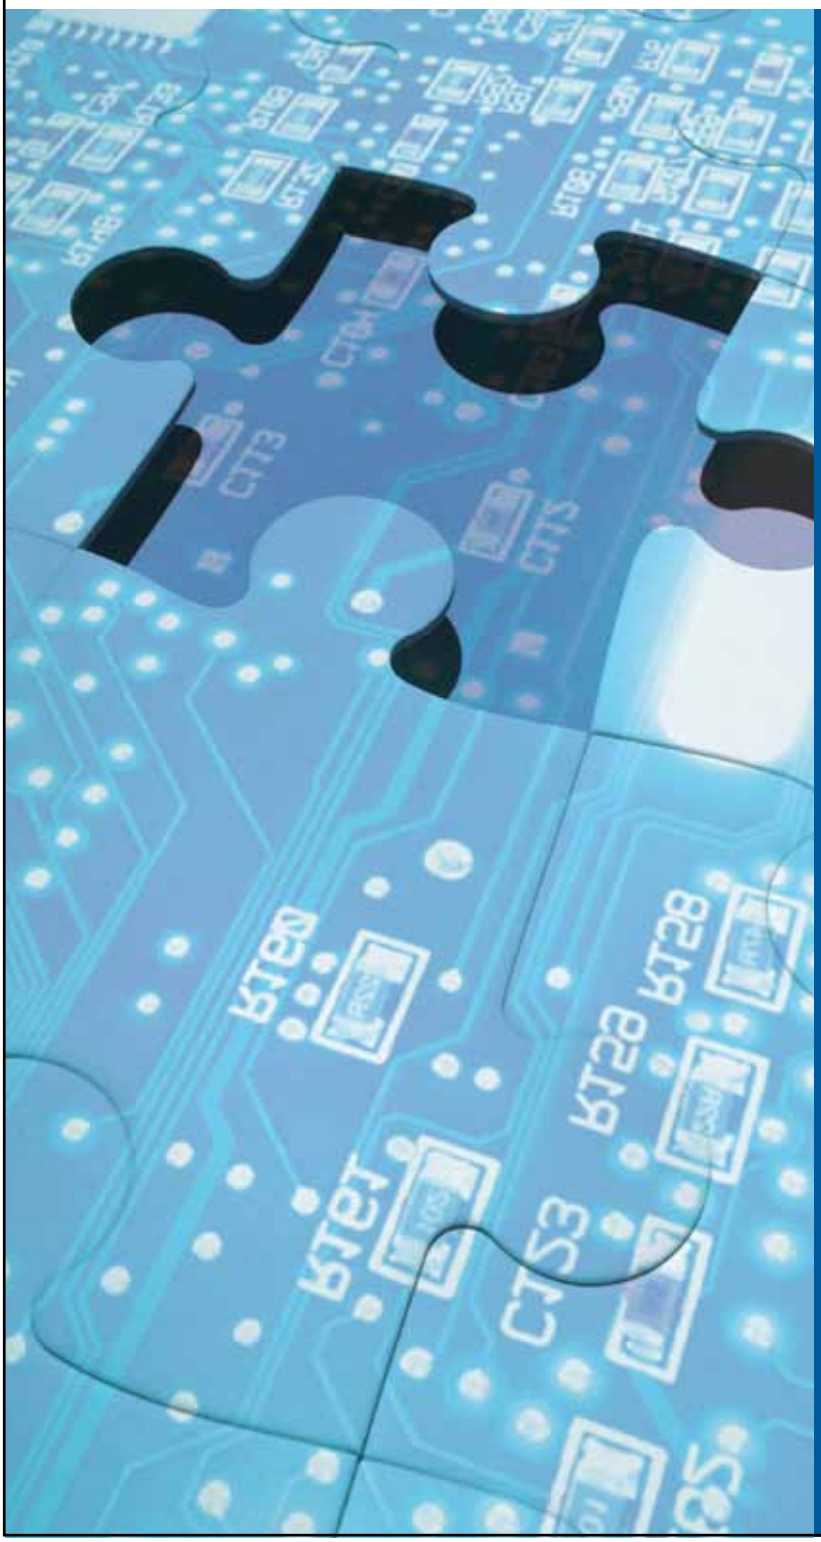

Is your company a missing piece in the EDA & IP industry?

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

EDA CONSORTIUM

If EDA is going to start bringing on more software users in 2014, it needs to come out with products that meet the needs of software developers at a price they can afford. In order to accomplish this, EDA products for programmers must be delivered in a much more "ready-to-consume" form. Platforms should be as prebuilt as possible while allowing for easy customization. Since support calls are barriers to productivity for the software engineer and costly to support for the EDA vendor, platforms for software engineers should be web-accessible. In some cases, they may reside fully in the cloud. This completely automates user access and simplifies support, if necessary.

Will 2014 be the year that EDA companies begin to meet the needs of the software engineer or will they keep trying to sell them a wolf in sheep's clothing? I think it will be the former because the opportunity's too great. Developing tools to support software engineers is an obvious and welcome growth path for the EDA market.

BRETT CLINE – FORTE

In the 19th century, prevailing opinion held that American settlers were destined to expand across North America. It was called Manifest Destiny.

In December 2014, we may look back on the previous 11 months and claim SystemC-Based Design Destiny. The Semiconductor industry is already starting to see more widespread adoption of SystemC-based design sweeping across the United States. In fact, it's the fastest growing worldwide region right now. Along with it comes SystemC-based High-level synthesis, gaining traction with more designers because it allows them to perform power tradeoffs that are difficult if not impossible in RTL due to time constraints. Of course, low power continues to be a major driver for design and will be throughout 2014.

Another trend that will be even more apparent in 2014 is the use of abstracted IP. RTL-based IP is losing traction for system design and validation due to simulation speed and because it's difficult to update, retarget and maintain. As a result, more small IP companies emerge with SystemC as the basis of their design instead of the long-used Verilog hardware design language.

SystemC-Based Design Destiny is for real in the U.S. and elsewhere as design teams struggle to contain the multitude of challenges in the time allotted.

DR. RAIK BRINKMANN – ONESPIN SOLUTION

Over the last few years, given the increase in silicon cost and slowdown in process advancement, we have witnessed the move toward standardized SoC platforms, leveraging IP from many sources, together with powerful, multicore processors.

This has driven a number of verification trends. Verification is diversifying, where the testing of IP blocks is evolving separately to SoC integration analysis, a different methodology from virtual platform software validation. In 2014, we will see this diversification extend with more advanced IP verification, a formalization of integration testing, and mainstream use of virtual platforms.

With IP being transferred from varied sources, ensuring thorough verification is absolutely essential. Ensuring IP block functionality has always been critical. Recently, this requirement has taken on an additional dimension where the IP must be signed off before usage elsewhere and designers must rely on it without running their own verification. This is true for IP from separate groups within a company or alternative organizations. This sign-off process requires a predictable metric, which may only be produced through verification coverage technology.

We predict that 2014 will be the year of coverage-driven verification. Effective coverage measurement is becoming more essential and, conversely, more difficult. Verification complexity is increasing along three dimensions: design architecture, tool combination, and somewhat unwieldy standards, such as UVM. These all affect the ability to collect, collate, and display coverage detail.

To read the entire story, please visit sldcommunity.com

Gabe Moretti has been in EDA for 45 years. First as an individual contributor with TRW Systems and Compucorp. Then as a manager with Intel and Signetics. He has been a member of the executive management team with EIS Modeling (a company he founded), HDL Systems, and Intergraph/Veribest. From 2000 to 2005 he was technical editor for EDA at EDN. Since then Gabe has run his own consulting company, GABEonEDA. He has a B.A. in Business Administration and a Master in Computer Sciences.

By Hamilton Carter, Contributing Editor/Power

Apache's Aveek Sarkar: Low Power Design

Low Power Engineering interviews Apache's Sarkar about low power trends.

LPE: Where do design engineers start when looking at different power efficiency design techniques?

Clock gating was one of the first things people started to look at for power optimization. In clock gating you shut off parts of your clock network such as a the portion of the clock network driving a particular bank of registers that do not have any data activity – this way you can save dynamic power. It's still one of the first places to look for savings as dynamic power is a big component of overall power. There's a pitfall however. A common mistake is not taking the design of the enable signals or their efficiency into account. If 99% of the design is clock gated, it doesn't mean that you have a low power design – it also depends on whether the enable signal is operating efficiently. Often times, projects which have clock gated the entire design can still lose sight of the fact that the efficiency of clock gating is controlled by the enable signal. If the enable signal is not designed properly it can limit the amount of power reduction you can achieve. So it's important to simulate the design and perform rigorous power exploration, tracking various metrics including clock gating efficiency especially at the RTL level to isolate and fix such scenarios.

LPE: What about power gating?

Power gating is commonly used to control leakage current or standby power. Its use started to become prominent around 2005 by mobile IC design teams, especially starting with the 65 nm process nodes. The leakage current was an increasing trend in those process nodes – to control that, they started adding power gating as a low power design technique. Power gating effectively breaks up power supply into two paths: an external path and an internal path by putting NMOS or PMOS transistor in series with the power rails. For example, if the video block of the device is not being accessed then the video processing section of the chip does not need to operate. By flipping off that particular switch on the power rail, it turns off or disconnects that block from the rest of the power supply. As a result, the

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

WWW.TRUECIRCUITS.COM/TIMING

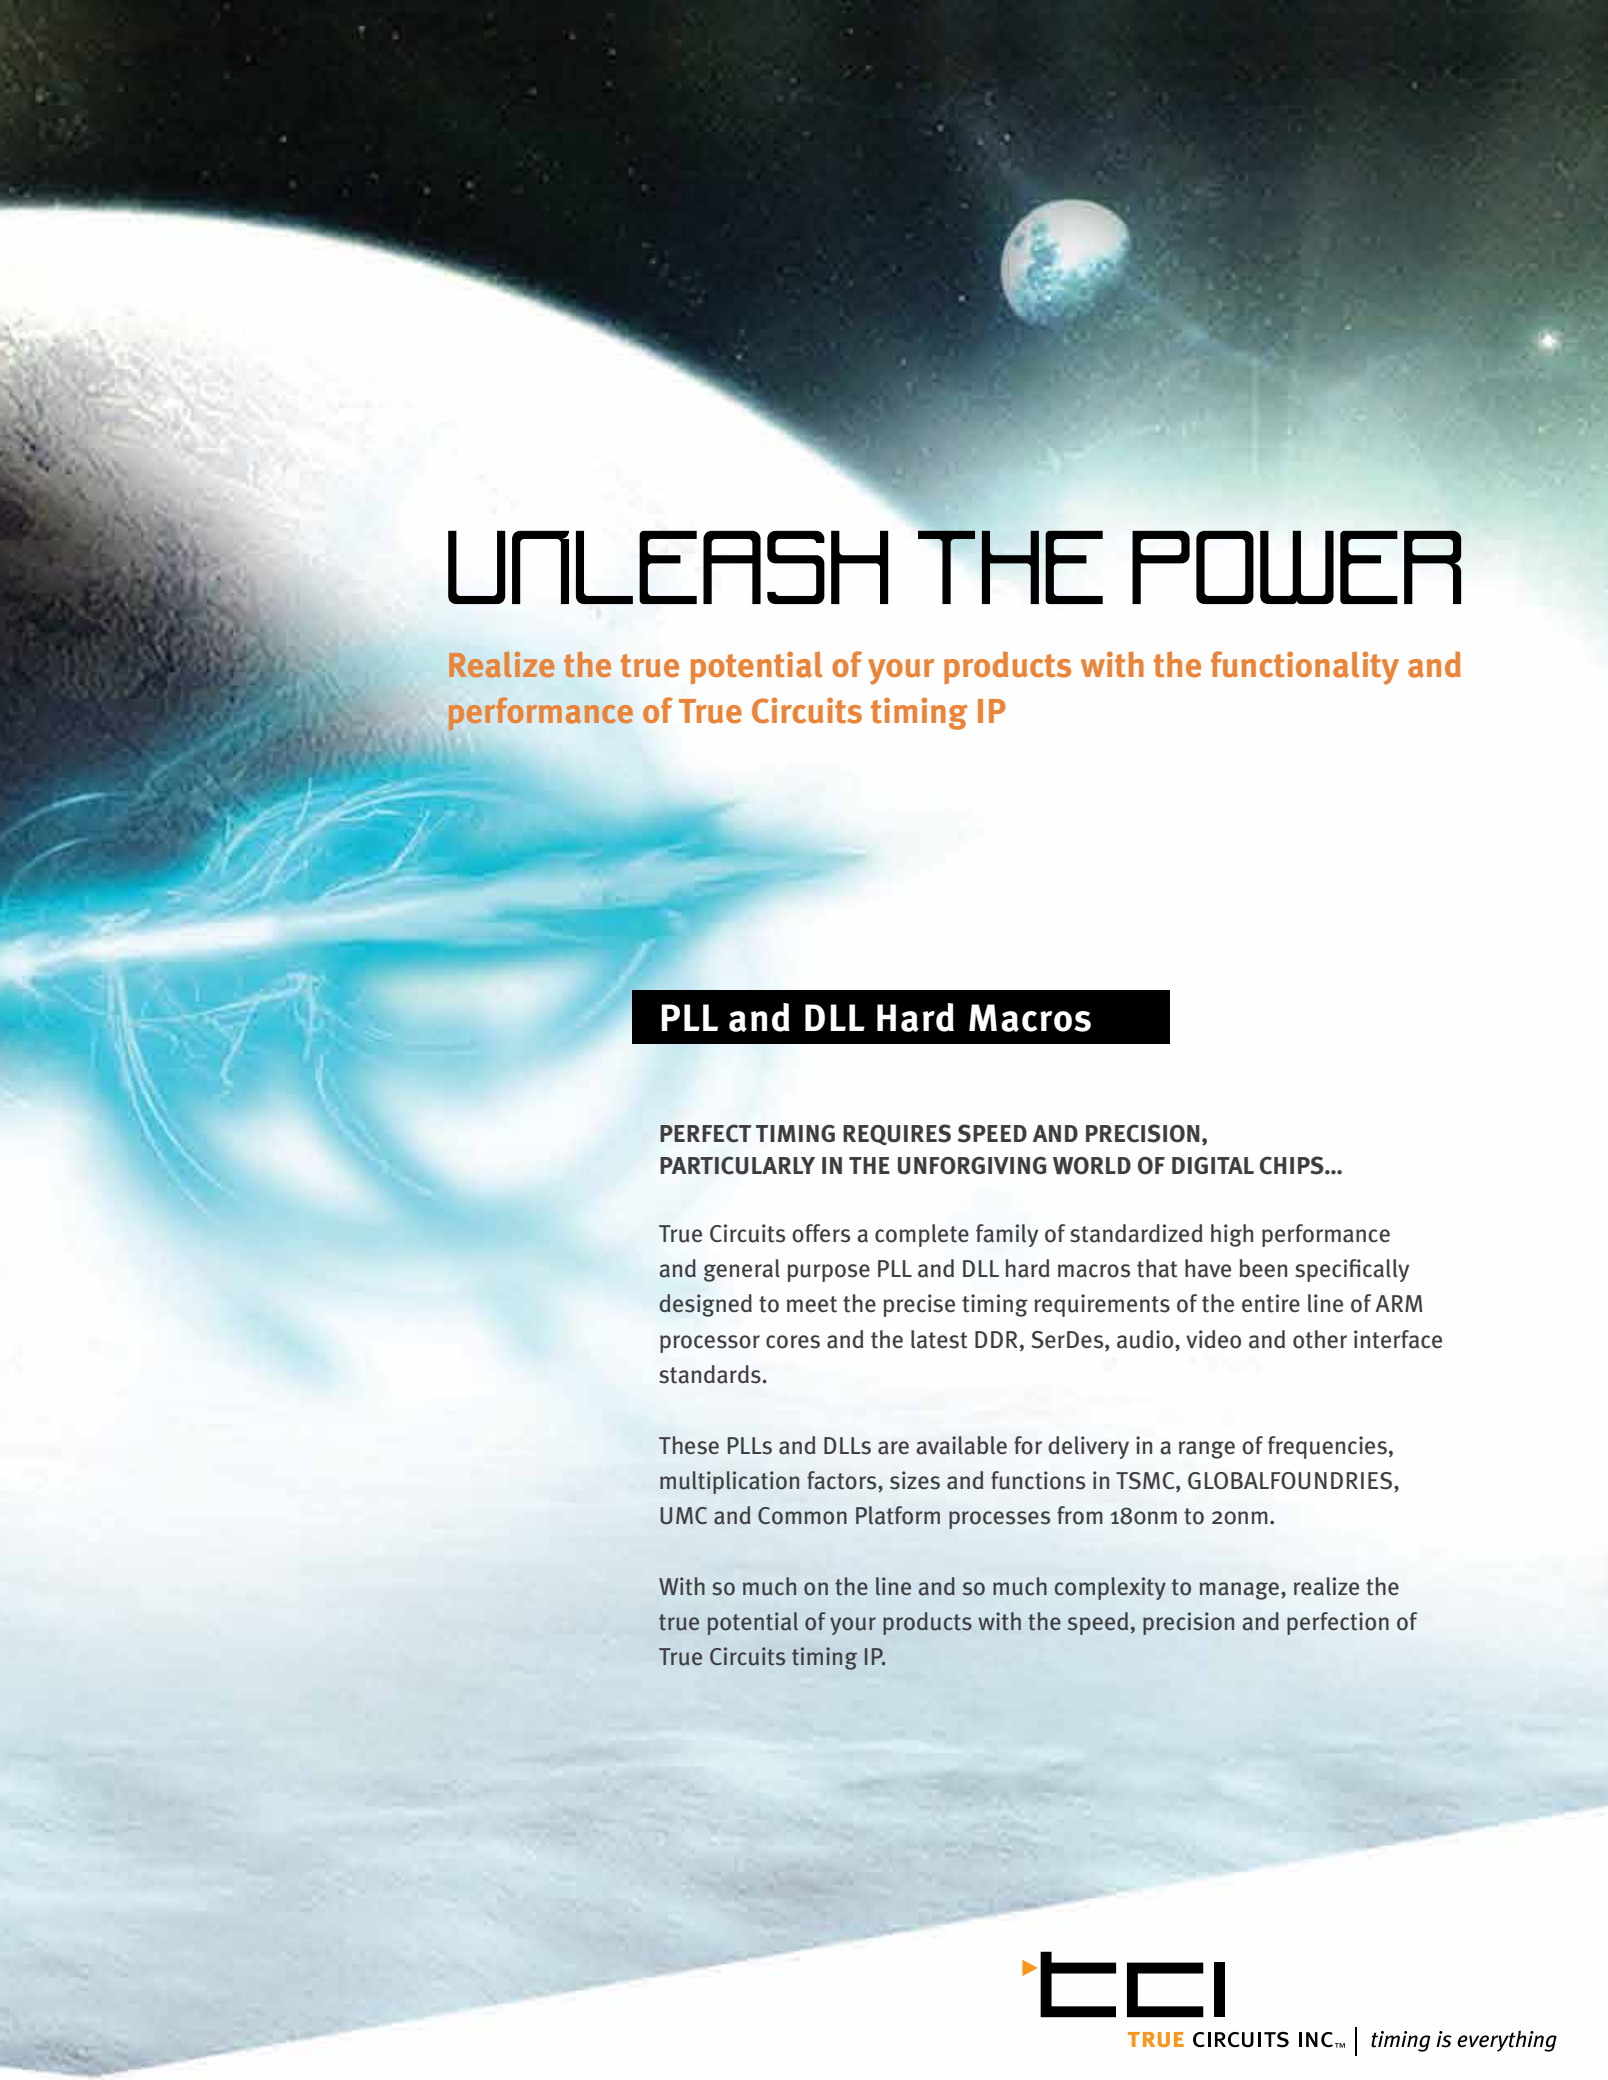

UNLEASH THE POWER

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

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, UMC and Common Platform processes from 180nm to 20nm.

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*

running. So, as you take charge from the neighboring CPU it can end up experiencing higher voltage drop or increased noise coupling, which can impact its performance or cause it to fail. This kind is scenario is fairly common and has to be watched out for, ideally with the use of a full-chip level package-aware power noise analysis flow that can model such a scenario.

LPE: How do dynamic voltage and frequency scaling factor in?

Depending on the type of application that is running, the chip can become very hot and to control the heat, the chip will need to be slowed down to extend the its lifetime or improve its power performance. But this technique, like others has the same set of challenges since it can introduce unknown behavior that can crop up when the device goes from one mode to another, or from one activity to another. To protect against unpredictable behavior, the design needs to be modeled at the full chip level with the package and the board to simulate the transient current changes that accompany these mode transitions.

LPE: What other power saving techniques might be of interest?

Forward and reverse biasing rate techniques are making a comeback. We are seeing designs where these techniques are being used. Another interesting technology that is becoming prevalent is the use of on-chip voltage regulators (LDO). When people want to control voltage for mission critical devices in an automotive IC or say for a sensor inside a pacemaker or other devices where you consume very little power over time, on-chip regulators will become increasingly critical. For these devices, it is very important to model the operation of the LDO in context of the design it is supplying power to. This ensures that the LDO can operate reliably across the entire range of operation of the chip.

Aveek Sarkar joined Apache Design, Inc., a subsidiary of ANSYS, as a senior applications engineer in 2003. Since then he has taken on different roles and responsibilities. Prior to joining Apache, Mr. Sarkar worked for Sun Microsystems on several generations of UltraSparc processors. Prior to Sun, he held engineering positions at Cadence Design Systems and National Semiconductor. Mr. Sarkar holds a B. Tech from the Indian Institute of Technology, Kanpur, a MSEE from Oregon State University, and MBA from Santa Clara University.

Hamilton Carter is the Chief Consultant at Pythagorean Productions. He has formerly worked as an engineer at Toshiba and Cadence Design Systems. Hamilton graduated from Ohio State University.

By Sara Verbruggen, Senior Editor/Manufacturing

Design for Yield Trends

Should foundries establish and share best practices to manage sub-nanometer effects to improve yield and also manufacturability?

TEAM EFFORT

Design for yield (DFY) has been referred to previously on this site as the gap between what the designers assume they need in order to guarantee a reliable design and what the manufacturer or foundry thinks they need from the designer to be able to manufacture the product in a reliable fashion. Achieving and managing this two-way flow of information becomes more challenging as devices in high volume manufacturing have 28 nm dimensions and the focus is on even smaller dimension next-generation technologies. So is the onus on the foundries to implement DFY and establish and share best practices and techniques to manage sub-nanometer effects to improve yield and also manufacturability?

'Certainly it is in the vital interest of foundries to do what it takes to enable their customers to be successful,' says **Mentor Graphics' Senior Marketing Director, Calibre Design Solutions, Michael Buehler**, adding, 'Since success requires addressing co-optimization issues during the design phase, they must reach out to all the ecosystem players that enable their customers.'

Mentor refers to the trend of DFY moving closer to the manufacturing/foundry side as 'design-manufacturing co-optimization', which entails improving the design both to achieve higher yield and to increase the performance of the devices that can be achieved for a given process.

But foundries can't do it alone. 'The electronic design automation (EDA) providers, especially ones that enable the critical customer-to-foundry interface, have a vital part in transferring knowledge and automating the co-optimization process,' says **Buehler**. IP suppliers must also have a greater appreciation for and involvement in co-optimization issues so their IP will implement the needed design enhancements required to achieve successful manufacturing in the context of a full chip design.

As they own the framework of DFY solutions, foundries that will work effectively with both the fabless and the equipment vendors will benefit from getting more tailored DFY solutions that can lead to shorter time-to-yield, says Amiad Conley, Applied Materials' Technical Marketing Manager, Process Diagnostics and Control. But according to Ya-Chieh Lai, Engineering Director, Silicon and Signoff Verification, at Cadence, the onus and responsibility is on the entire ecosystem to establish and share best practices and techniques. 'We will only achieve advanced nodes through a partnership between foundries, EDA, and the design community,' says Ya-Chieh.

But whereas foundries are still taking the lead when it comes to design for manufacturability (DFM), for DFY the designer is intimately involved so he is able to account for optimal trade-off in yield versus PPA that result in choices for specific design parameters, including transistor widths and lengths.

For DFM, foundries are driving design database adjustments required to make a particular design manufacturable with good yield. 'DFM modifications to a design database often happen at the end of a designer's task. DFM takes the "ideal" design database and manipulates it to account for the manufacturing process,' explains Dr Bruce McGaughy, Chief Technology Officer and Senior Vice President of Engineering at ProPlus Design Solutions.

The design database that a designer delivers must have DFY considerations to be able to yield. 'The practices and techniques used by different design teams based on heuristics related to their specific application are therefore less centralized. Foundries recommend DFY reference flows but these are only guidelines. DFY practices and techniques are often deeply ingrained within a design team and can be considered a core competence and, with time, a key requirement,' says McGaughy.

IN THE SPIRIT OF COLLABORATION

Ultimately, as the industry continues to progress requiring manufacturing solutions that increasingly tailored and more

and more device specific, this requires earlier and deeper collaboration between equipment vendors and foundry customers in defining and developing the tailored solutions that will maximize the performance of equipment in the fab. 'It will also potentially require more three-way collaboration between the designers from fabless companies, foundries, and equipment vendors with the appropriate IP protection,' says Conley.

A collaborative and open approach between the designer and the foundry is critical and beneficial for many reasons. 'Designers are under tight pressures schedule-wise and any new steps in the design flow will be under intense scrutiny. The advantages of any additional steps must be very clear in terms of the improvement in yield and manufacturability and these additional steps must be in a form that designers can act on,' says Ya-Chieh. The recent trend towards putting DFM/DFY directly into the design flow is a good example of this. 'Instead of purely a sign-off step, DFM/DFY is accounted for in the router during place and route. The router is able to find and fix hotspots during design and, critically, to account for DFM/DFY issues during timing closure,' he says. Similarly, Ya-Chieh refers to DFM/DFY flows that are now in place for custom design and library analysis. 'Cases of poor transistor matching due to DFM/DFY issues can be flagged along with corresponding fixing guidelines. In terms of library analysis, standard cells that exhibit too much variability can be systematically identified and the cost associated with using such a cell can be explicitly accounted for (or that cell removed entirely).'

'The ability to do "design-manufacturing co-optimization" is dependent on the quality of information available and an effective feedback loop that involves all the stakeholders in the entire supply chain: design customers, IP suppliers, foundries, EDA suppliers, test vendors, and so on,' says **Buehler**. 'This starts with test chips built during process development, but it must continue through risk manufacturing, early adopter experiences and volume production ramping. This means sharing design data, process data, test failure diagnosis data and field failure data,' he adds.

A pioneer of this type of collaboration was the Common Platform Consortium initiated by IBM. Over time, foundries have assumed more of the load for enabling and coordinating the ecosystem. 'GLOBALFOUNDRIES has identified collaboration as a key factor in its overall success since its inception and been particularly open about sharing foundry process data,' says **Buehler**.

TSMC has also been a leader in establishing a well-defined program among ecosystem players, starting with the design tool reference flows it established over a decade ago. Through its Open Innovation Platform program TSMC is helping to drive compatibility among design tools and provides interfaces from its core analysis engines and third party EDA providers.

In terms of standards Si2 organizes industry stakeholders to drive adoption of collaborative technology for silicon design integration and improved IC design capability. **Buehler** adds: 'Si2 working groups define and ratify standards related to design rule definitions, DFM specifications, design database facilities and process design kits.'

Open and trusting collaboration helps understand the thriving ecosystem programs that top-tier foundries have put together. McGaughy says: 'Foundry customers, EDA and IP partners closely align during early process development and integration of tools into workable flows. One clear example is the rollout of a new process technology. From early in the process lifecycle, foundries release 0.x versions of their PDK. Customers and partners expend significant amounts of time, effort and resources to ensure the design ecosystem is ready when the process is, so that design tapeouts can start as soon as possible.'

DFY is even more critically involved in this ramp-up phase, as only when there is confidence in hitting yield targets will a process volume ramp follow. 'As DFY directly ties into the foundation SPICE models, every new update in PDK means a new characterization or validation step. Only a close and sustained relationship can make the development and release of DFY methodologies a success,' he states.

Sara Verbruggen, based in Europe, is a technology journalist and editor, writing about the semiconductor, solar photovoltaic (PV), printed electronics, nanotech and energy storage industries. Before going freelance in 2010, she worked as a B2B editor for several years, specialising in printed electronics and nanotechnology.

By John Blyler, Vice President/Chief Content Officer

System Integration Requires a Shared Viewpoint

Qualcomm's Hexagon DSP chip benefits from Dassault Systemes' dashboarding tool.

Qualcomm uses Dassault Systemes' dashboarding tool in its Hexagon DSP chip to incorporate multiple design metrics from key EDA tools.

The EDA tool market has longed talked about its need to expand beyond the creation of silicon-based system-on-chips (SoCs) to provide packages that integrate the larger hardware and software system. Specifically, the major tool vendors emphasized the need to move beyond EDA-centric issues like electronic system level (ESL) design, functional verification, design-for-yield or any similar so-called crisis issues. The goal has been to move beyond chip creation to system integration to deal with both hardware and software at the chip, board, and end-user product levels.

"It begins with a shift from design creation to integration in the electronic systems industry," states the Cadence's EDA Vision 360 report. EDA tool companies have had to expand their coverage into the larger system market, thanks to changes in the semiconductor supply chain.

Regardless of the drivers, the expansion from creation to integration tools for the larger system has not been easy move for a variety of technical and cultural reasons. Consider but one aspect of the problem: How to provide higher-level integration when your customer uses a variety of internal and competitive tools? For example, most IDMs like Intel, Samsung and Apple, as well as fabless chip companies use a variety of EDA tools for synthesis, place and route (P&R), time and power closure, etc. Further, many use a mix of internal tools that have been tailored to the needs of the customer.

To become a system integrator – at least from the chip design space viewpoint – tool providers will need a mechanism to gather, analyze and display useful data metrics from a variety EDA packages. One of the few companies that come close to such an application is not an EDA company at all, but rather comes from a higher-level, project lifecycle management (PLM) provider.

Qualcomm recently shared their challenges in integrating the metrics from a mix of chip design tools. Their problem was how to put together all of the disjointed design pieces for development of its Hexagon DSP-based multithreaded CPU architecture. With a global design team (San Diego, India and Austin), the company had to communicate all of the traditional design metrics like timing and area, with secondary metrics like power and signal integrity. Adding to this technical complexity was the diversity of professionals that needed access to these metrics, from system architects, RTL coders to logical and physical designers.

The answer was simply to use dashboards to display data and metrics in such as way as to quickly show trends and trouble spots. Good dashboards highlight the metrics data in a graphical analysis format while also providing a transition from high-level to detailed low-level views. This abstraction-level zoom-in/zoom-out capability helps designers quickly spot trouble areas and then probe down into the details.

Dashboarding is nothing new. "Qualcomm has many internal dashboards," explained Dwight Galbi, Principle Manager of Qualcomm's physical design team at a recent Dassault Systemes's Customer Forum. "We have dashboards that cover some of the (design metrics) … but not one that incorporated all of them." What was needed was a dashboard to provide design metrics from a variety of EDA tools throughout the chip design process.

That's where Dassault Systemes's dashboarding tool called Pinpoint came to into play. In his presentation, Galbi listed the mix of life cycle tools (albeit from one vendor, i.e.,, Synopsys) used in his recent DSP project. The list included Design Compiler for synthsis; IC and Talis for P&R; and Prime Time for sign off.

"The beauty here is that these are four different tools but you can incorporate all of the reports into the same web-based server," said Galbi. Equally important (though not mentioned by Galbi) was that the tool provides a graphical visualization of physical design, timing paths, etc., without needing to reload the entire design block. This saves both time and money – since the user doesn't need to activate a license from the EDA tool vendors.

Further, using a dashboard can provide a way for geographically dispersed teams to communicate via a common view of the design. This is a key requirement for any system integration. For example, the chip's Register-Transfer-Level (RTL) codes are often developed by teams in different geographic locations. Complicating the geographic challenges is the need to incorporate third party-IP and reused internal design blocks with the various RTL designs before the implementation process even begins. This is a problem since the physical layout and design team requires the RTL synthesized code (with all the IP), design planning and place-and-route (P&R) data to decide if the primary chip design constraints can be met.

Getting the detailed RTL design team to work with the physical layout-design teams as soon as possible encourages communication and successful design practices. It helps mitigate the problems of siloed design activates. Also, a dashboard approach incorporates the essential data metrics from several different EDA tools into one place. This single, global view increases the likelihood of a successful SoC design as well as integrating that design – and the team – with the next level of system development.

John Blyler covers today's latest high-tech, R&D and even science fiction in blogs, magazine articles, books and videos. He is an experienced physicist, engineer, journalist, author and professor who continues to speak at major conferences and before the camera on. John is the Vice President, Chief Content Office for Extensionmedia, which includes the brands Chip Design, Solid State Technology, Embedded Intel and others. He holds a BS in Engineering Physics and a MSEE. John plays the piano and holds a black belt in TKD.

By Dave Bursky, Technology Editor

Design Challenges Continue as New Market Opportunities Arise

ARM, Sonics, Skyworks, Inside Secure and Infinitedge speak about market opportunities at Semico IP panel.

Designs, power and security for a new world of applications combine perspectives from ARM, Sonics, Skyworks, Inside Secure and Infinitedge (moderator) at recent Semico IP panel.

At last month's Semico IP Impact conference in San Jose, Calif., participants in a panel titled "Designing for New World Applications" discussed design challenges, security, and power-efficiency issues related to creating the "Internet of Things" (IoT) and other new applications.

The panel kicked off with a question from moderator Kent Shimasaki, a Managing Partner at Infinitedge. He asked what customers have been asking for from the various panelists and their companies.

Responding to that question, Grant Pierce, CEO of Sonics, stated that the customers he has talked to are mostly from large companies. They are chasing large market opportunities or going after new markets like the IoT. He explained, "They want us to make their business easier and less risky for them to pursue. We can do this by integrating complex solutions of tens, hundreds, or even thousands of IP cores, and exploit every opportunity to reduce power consumption to extend battery life. Another aspect that comes up is the need to raise the security levels appropriate to the actual application. Thus, different levels of security and different access permissions are key requirements."

Ron Moore, Director of Strategic Accounts and Marketing for Physical IP at ARM, explained that low power is by far the largest demand across the company's entire breadth of CPU IP. The availability of a wide range of manufacturing process technologies from 180 nm all the way down to 14 nm allows the company to offer multiple performance and power options. In addition, ARM can redesign the cores to take advantage of the latest design techniques to lower power and maintain or improve performance. Security is also a concern and ARM has been seeing more interest in its TrustZone IP.

John O'Neill, VP of Skyworks, agreed. He said that power and risk mitigation are key demands that Skyworks is addressing, thanks to its vertically integrated structure that covers design, fabrication, and system design. In doing so, it allows the company to tailor products to meet customer needs. Servicing a different type of customer, Steve Singer, Director of Systems Engineering for Embedded Security Solutions at Inside Secure, said that his customers were asking for software stacks, drivers, and semiconductor IP for embedded security applications to handle IPSEC, SSL, and other popular standards. Protecting data in motion and providing device protection to prevent hacking are key customer issues.

Another major question the panel focused on was how to determine the best technology for a given application. It's a mix of technologies and packaging options, according to Moore. "The IoT designers need processing power, sensors, and low-power radios—technologies that do not readily integrate on a single chip." Addressing the same question, Pierce indicated that his customers wanted chips with very power-efficient functions. In addition, they requested that those functions be enabled with very-fine-grained power management. To that end, Sonics has developed a design approach that creates multiple power domains, which can be switched on and off very quickly. However, such domains need knowledge from all of the blocks on the chip to know when to turn off or on the necessary blocks.

According to Singer, a lot of parallel processing will be needed to handle the high data rates now possible on WiFi, cable, and fiber-optic networks. Keep in mind that Crypto engines typically have to handle multiple protocols. Because only one protocol typically runs at a time, however, power consumption stays relatively low. But there could be thousands to a million tunnels running in cell systems, which means that a lot has to be done in parallel just to move the data. On top of that, there's the need for security to mitigate any threats.

The power budget for the radio portion of a wireless system is a key concern at Skyworks, explained O'Neill. A system may typically have half-a-dozen or more RF radios—multiple channels of WiFi, Bluetooth, LTE, GSM, GPS, and still other radios. Unlike the encryption processor, where only one algorithm runs at a time, multiple radios are often active simultaneously. This aspect increases power consumption. New techniques that modulate power using voltage scaling based on actual signal modulation can reduce instantaneous power consumption. Yet this approach requires a very complex control scheme. In sensor nodes, short-burst communications for small packets (sub-10-kbits) can have the power supplemented with the use of supercapacitors, which collect and store harvested energy from the environment to minimize the dependency on battery power. These supercaps can then power the communications circuits for short data bursts.

Keeping all of the radio and line-based devices is a challenge for customers in all areas. Simple protection schemes, such as ARM's TrustZone, provide a starting point, explained Moore. But additional, more secure solutions are needed as hackers get more sophisticated. Software-only schemes provide limited security, stated Singer, emphasizing that "many other aspects need additional hardware support to provide trusted solutions in the silicon." In some devices, such as pacemakers (which now can be updated through a wireless link), Singer asked: Where is the boundary to protect the device? Anyone can write an AES algorithm, he notes. In addition to the software, however, designers need to do a good hardware implementation. Pierce also felt that on-chip firewalling—perhaps coupled with the ARM Trustzone—will be needed to protect content.

These aspects of design, power, and security are all parts of the solution. Few companies have the resources to address them all. Thus, many companies have set up an ecosystem of multiple partners, which cooperate to support customers. As Singer put it: Companies must sell solutions, and that means they need partners with solutions that can plug into various infrastructures.

Dave Bursky is a contributing editor for Chip Design and Embedded Intel® Solutions magazines. He also is the technical editorial manager at Maxim Integrated Products Inc. in Sunnyvale, CA.

BEHIND THE GRAPHICS

By Caroline Hayes, Senior Editor

Graphics processing can bring color, depth and perspective to many end applications. It is no longer a case of monochrome or color, pixelated images or text, but shaded images that can add depth to anything from a computer game to digital signage or medical equipment.

Graphics are used for more than computer games today, although it could be argued that computing gaming has pioneered many graphics processing developments. For example, it was computer gaming that shifted from 2D to 3D graphics, as characters were required to move around, requiring their environments to adapt to the changes in real-time. This prompted the move to 3D graphics, with changing views required onscreen. Graphics accelerators were required to copy bitmaps onto surfaces, paying regard to where the camera or light source would be positioned in the scenario. Pixels had to be transformed to reflect the changes in shape and perspective that any movement or position change would bring. This is known as texture mapping.

Read the entire story on the "System Design Engineering" portal: **sldcommunity.com**

By Dave Bursky, Technology Editor

IDEM and ASIP Coverage - From 3-D Transistors to 2.5D or 3D Systems

From 3-D Transistors to 2.5D or 3D Systems

From the ultra-small 3D transistors described in papers at this month's International Electron Devices Meeting (IEDM) in Washington, D.C., to the 2.5D and 3D multichip structures described at the 3D Architectures for Semiconductor Integration and Packaging (ASIP) conference held in Burlingame, Calif., designers are finding more ways to pack more transistors on a chip and to pack more functions into a limited area on a printed-circuit board. For instance, at IEDM TSMC Shien-Yang Wu and his team of researchers described a 16-nm FinFET process in paper 9.1 that they feel is one of the world's most advanced semiconductor technologies.

The process is the first integrated technology platform to be announced below the 20 nm node, with key capabilities that include a 48-nm fin pitch and the smallest SRAM cell ever incorporated into an integrated process—a 128-Mb SRAM with a cell area of just 0.07 μm2 per bit. The process' short-channel effects were well-controlled, with DIBL <30 mV/V, saturation current of 520/525 μA/μm at 0.75V (NMOS and PMOS, respectively) and an off-current of 30 pA/μm. Depending on the designer's goal, the process delivers either a 35% speed gain or a 55% power reduction in comparison with

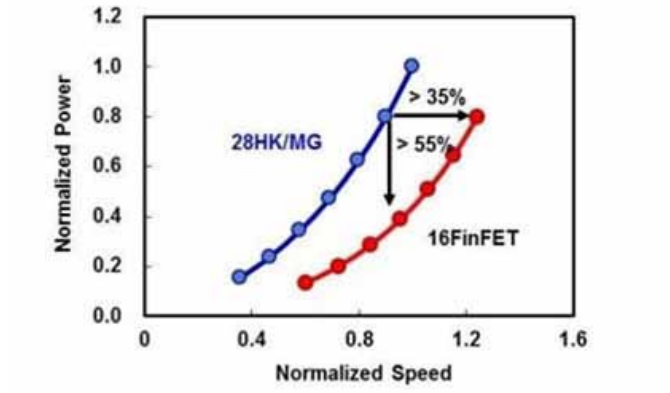

Figure 1: The 16 nm process platform developed by TSMC allows designers to get 55% reduction in operating power or a 35% improvement in operating speed vs the company's established 28 nm high-K/metal-gate process.

TSMC's existing 28-nm high-k/metal-gate planar process, and with twice the transistor density (Figure 1).

Creation of a "superchip" was the goal of researchers at the New Industry Creation Hatchery Center at Tohoku University in Sendai, Japan. The heterogeneous 3D integration described by the Professor Mitsumasa Koyanagi in a plenary presentation at IEDM allows various kinds of device chips with different sizes, different functions, and different materials to be stacked to form the superchip. A key technology developed to achieve this consists of self-assembly and electrostatic (SAE) temporary bonding. To demonstrate the technology, the university fabricated several prototype superchips—examples include stacking MEMS chips, spin memory chips and a photonic device chip on a CMOS logic chip; a 3D back-illuminated image sensor with through-silicon vias stacked on top of an image processing chip; and a 3D microprocessor with self-test and self-repair functions.

The assembly process to create superchips starts with known good die (KGD) that are sorted from several device wafers and simultaneously bonded as a batch onto a carrier wafer (Figure 2, left). High alignment accuracy is achieved using the self-assembly and electrostatic bonding. The process repeats with additional carrier wafers. Multiple carrier wafers with the KGDs are then stacked onto a target interposer wafer. This allows multiple superchips to simultaneously be fabricated. The surface tension of liquid is used in the self-assembly scheme to simultaneously align many dies in parallel. Hydrophilic areas and hydrophobic areas are formed on the surface of the wafer or chip to obtain high alignment accuracy. As many as 500 chips have been simultaneously aligned with an average alignment accuracy of 0.05 μm within 0.1 seconds.

The electrostatic temporary-bonding and de-bonding method for assembly of the multiple carrier wafers allows the stacking integration of multiple chips (Figure 2, right). Many chips are simultaneously bonded onto the electrostatic carrier wafer

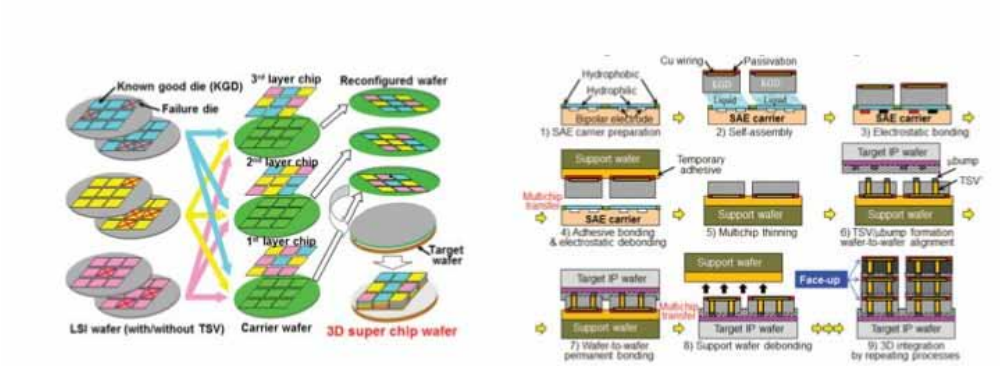

Figure 2: The assembly process to create superchips starts with known good die (KGD) that are sorted from several device wafers and simultaneously bonded as a batch onto a carrier wafer (left). The process repeats with additional carrier wafers and then multiple carrier wafers with the KGDs are then stacked onto a target interposer wafer using electrostatic bonding and debonding (right).

(e-carrier) by the electrostatic force after the simultaneous alignment by self-assembly. The electrostatic force for temporary bonding is generated by applying a high voltage to the electrodes embedded in the e-carrier wafer. A high voltage with opposite polarity is applied to the electrodes for de-bonding the chips.

These two presentations are just the proverbial tip of the iceberg representing several hundred paper presentations at IEDM that covered process and manufacturing, memory technology, nano-device technology, power and compound semiconductors, advanced CMOS technology, and many other subjects. For more information, go to www.ieee-iedm.org.

Running concurrently with IEDM but on the opposite coast, the 3D-ASIP Conference in Burlingame delved into many aspects of 2.5 and 3D integration, ranging from basic integration, to various interposer technologies, to wafer handling and thermal challenges to name a few. Many of the presentations examined the evolution of assembly techniques to move from 2D to 2.5 D to true 3D implementations. Doug Yu, the Director of the Integrated Interconnect and Packaging Division at TSMC provided an overview of wafer-level system integration technology, while Robert Patti, the CTO of Tezzaron Semiconductor described a combination of dis-integration and then integration to create a high-density and high-performance memory stack (See "Advances in DRAM and non-volatile memories keep upping system performance", Aug. 26, 2013). The architecture of the memory array provides 256 independent channels, each containing 256 Mbits of storage and capable of transferring data at 64 Gbits/s with a latency of just 9 ns

Another memory presentation by Eric Beyne, the Program Director for 3D System Design at IMEC examined high-bandwidth memory-logic 3D integration by either direct stacking or the use of interposers. One of the key aspects of leveraging the 3D integration is to reduce the power consumption of the chip-to-chip interconnects by lowering the voltage swing, widening the I/O to lower the transfer frequency, and use vertical interconnects in a chip stack to reduce the wiring length (Figure 3). Using 3D through-silicon vias and microbump interconnects designers at

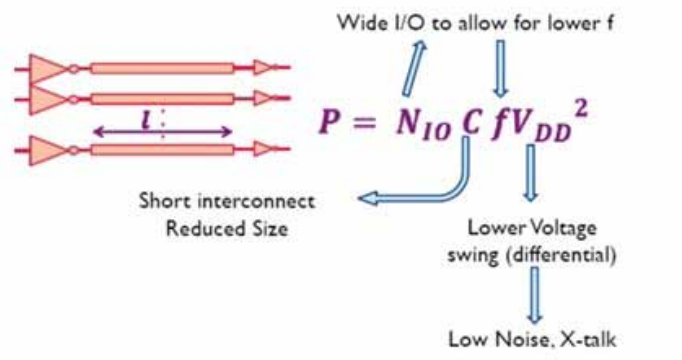

Figure 3: Multiple factors such as the I/O width, the load capacitance, the transfer frequency, and the operating voltage must be taken into account when estimating the power consumed in chip-to-chip interconnects. (Source, IMEC).

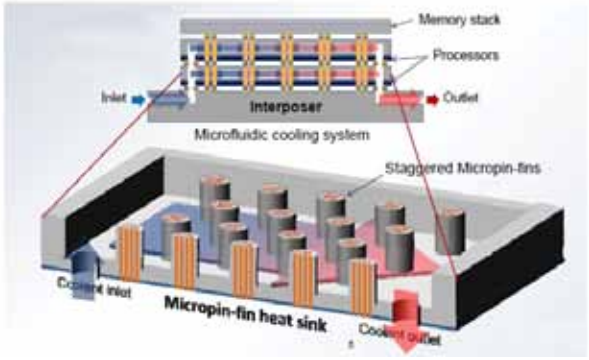

Figure 4: Microfluidic cooling between layers of chips, can pull out the heat, but there are concerns about the reliability of the microfluidic I/O technology as well as power-supply noise and the durability of TSVs due to the pressure of the liquid coolant as it flows through the package. (diagram courtesy of Georgia Tech).

IMEC were able to assemble high-density chip stacks, but encountered issues with the increased power density. The high power density can result in thermal issues (increased temperatures) and higher temperatures could affect DRAM data retention since retention time decreases as temperature increases.

Just such thermal issues were discusses by Joseph Maurer a support contractor to DARPA, and by Muhannad Bakir, Associate Professor, School of Electrical and Computer Engineering at the Georgia Institute of Technology. At DARPA, Maurer described multiple projects aimed at pulling out the heat and improving thermal conductivity. Techniques such as the use of copper nanosprings; near-junction thermal transport with liquid cooling and high-thermal conductivity diamond substrates; the use of a 3D vapor chamber with vibrating elements; the use of thin-film superlattice materials; and still other approaches are all being explored. Examining the use of microfluidic cooling on 3D ICs, Bakir showed a potential solution using coolant cycled through a multichip stack composed of two processor layers and a memory stack (Figure 4). With such a stack there are concerns about the reliability of circulation system used the microfluidic cooling, as well as the endurance of the TSVs since they are under pressure from the liquid flowing between the layers.

These few papers were just a few of the presentations at the 3D-ASIP conference. For more details, go to www.3dasip.org to view the program or purchase the proceedings.

Dave Bursky is a contributing editor for Chip Design and Embedded Intel® Solutions magazines. He also is the technical editorial manager at Maxim Integrated Products Inc. in Sunnyvale, CA.

By Jonah McLeod, blogger

The Year of the Wearables

When Humans Become Nodes on the Internet of Things

A good number of press releases coming out of the 2014 International Consumer Electronics Show the week of January 6th deal with the topic of "wearables," the wrist bands, belt clips, and other appendages that monitor a person's physical activity. Chip giant Intel spent time talking up the market segment in their sessions at CES. And health insurance giants contributed to joint announcements with software companies promoting wearables as a solution to the rising cost of medical care.

Wearables provide a potent new growth market for chip vendors, and it has created a whole new category of products for consumers to purchase. They will add to the collection of gadgets a person carries around rather than displacing a gadget every consumer already has—mobile phone, tablet, etc. Unlike the poor MP3 player that got gobbled up, the new wearable devices will not suffer the same fate. In fact, the number of wearables will multiply contributing to the proliferation of devices on the Internet of Things (IoT).

Neura Inc., a Sunnyvale, Calif. start-up that offers a software platform for apps development and cloud services to enhance the value of these apps, came to CES to engage with developers. Gilad Meiri, CEO of Neura declares IoT proliferation will occur in waves, beginning with wearables, following by myriad devices installed into every home, and finally filling the car with heretofore unavailable functions.

Neura is hoping to win a share of an IoT market worth $8.9 trillion by 2020, growing at a compound annual rate of 7.9 percent, according to market research firm IDC. IDC's prediction encompasses the entire IoT ecosystem, including the intelligent systems, connectivity services, platforms, analytics and vertical applications in addition to the security and professional services required to make the entire system work.

Wearable is where the action is at CES. In his keynote at CES Intel CEO Brian Krzanich showed off his own wearable wristband, which provides the runner with steps taken, distance run, active minutes—upright and moving versus sedentary in front of a screen—and calories consumed. If you're really engaged, you will accurately enter your weight, the amount of calories you consume and the amount of fluid you take in daily into the app on your phone or PC, which will record this data and provide you with feedback on your health.

These devices are rudimentary in that they consists of an ARM CPU core and the accelerometer and gyro that goes into most every mobile phone. The wearable detects steps taken and from this computes everything else. The shortcomings of this is that the distance covered is typically less than the physical distance as measured by Google and the device doesn't measure other vital signs—temperature, heart rate, respiration rate, and blood pressure. But, the next generation will most likely address this shortcoming.

Start-up Scanadu of Moffett Field, Calif. is taking the concept of measuring vital signs to the limit. Its prototype device called Scout aspires to emulate the Tricoder on Star Trek. Place the Scout on a patient's forehead and it reads all his vital signs that can be transferred to a smartphone for analysis and processing. The ambition is to help distinguish minor illnesses treatable at home with bed rest and over the counter medications from major problems that demand emergency room treatment. The device is a contender for the Tricorder X prize that Qualcomm has established for the invention that provides Tricoder functionality.

What provides these wearables their value is the "Big Data" behind them. The large computing platforms analyzing the endless amount of data being collected by wearables and the proliferating numbers of IoT devices is what will provide the value to the individual user.

Jonah McLeod, blogger

Verissimo SystemVerilog Testbench Linter

Thorough audit of your test benches

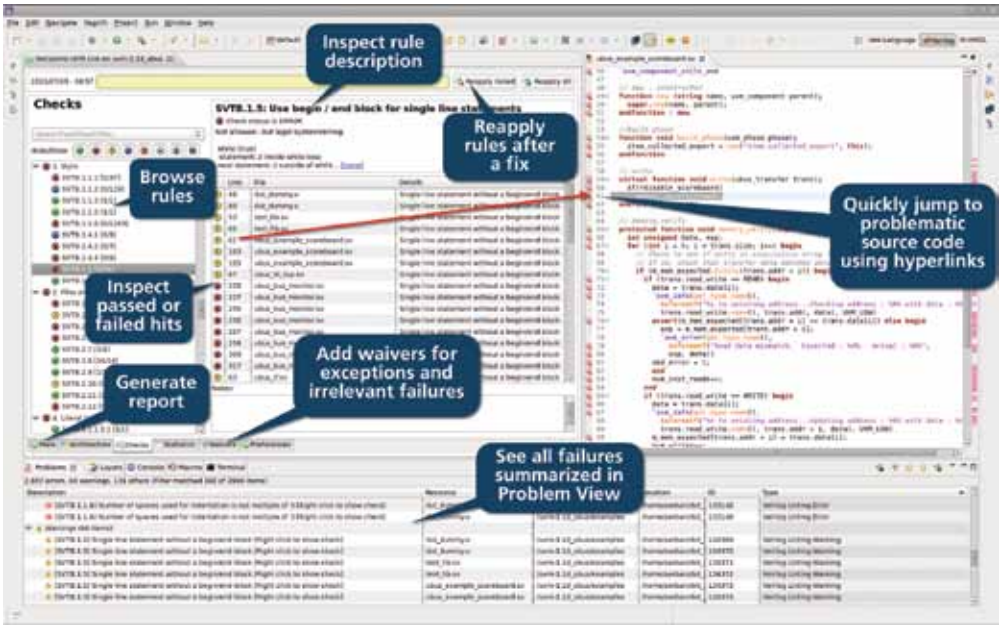

Verissimo SystemVerilog Testbench Linter is a static code analysis tool that allows engineers to perform a thorough audit of their testbenches. It enables users to easily identify improper SystemVerilog language, semantics, and styling usage, as well as verification methodology violations.

Verissimo offers a comprehensive library of generic SystemVerilog and Universal Verification Methodology (UVM) rules. In addition, users can customize the existing rules by tuning their parameters or can create custom rule sets by selecting those that correspond to their particular requirements. Verification groups that need their own rules can use the dedicated Java application programming interface (API) delivered with the linter. The API allows users to query the linter's internal database to find the relevant information and create new rules. A report generator is available to save the results of a linting session as a text or HTML file.

AMIQ's linter runs both in GUI and batch modes. In the GUI mode, the Verissimo linter integrates with the Design and Verification Tools (DVT) integrated development environment (IDE). Users can perform linting and then visualize the results on the DVT GUI, which offers an effective way to read and understand the error and warning messages.

With the DVT IDE, the messages can also be easily sorted and filtered by category, severity, and source location. In addition, DVT's code navigation features such as hyperlinks, allow users to jump directly to the problematic source line to fix the issue flagged by the linter. Users can also annotate rules and share the notes with the team.

In short, the Verissimo linter enables verification engineers to improve testbench code reliability and maintainability and implement best coding practices.

FEATURES

- Comprehensive library of generic SystemVerilog and UVM built-in checks
- Custom rule set configuration
- API for creating new custom checks
- Rule severity adjustment
- Waivers for exceptions and irrelevant failures
- Rule annotation
- Message filtering by source, category, and severity
- Text or HTML analysis reports
- Batch or GUI mode
- Integration with the DVT Eclipse IDE (GUI mode)

CONTACT INFORMATION

AMIQ EDA
Support & Evaluation:
support@amiq.com
sales@amiq.com
www.amiq.com; dvteclipse.com

True Circuits, Inc.

DDR 4/3 PHY

Organizations: GSA

The TCI DDR 4/3 PHY is a high-performance, scalable system using a radically new architecture that continuously and automatically adjusts each pin individually, correcting skew within byte lanes. This state-of-the-art tuning acts independently on each pin, data phase and chip select value. Read data eye and gate timing are also continuously adjusted. Automatic training is included for multi-cycle read gate timing and write leveling, write data eye timing, and internal and external (on DRAM) Vref setting.

Remarkable physical flexibility allows the PHY to adapt to each customer's die floorplan and package constraints, yet is delivered and verified as a single hard macro for easy timing closure with no assembly required.

The PHY is DFI 3.1 compliant, and when combined with the Northwest Logic DDR 4/3 memory controller, a complete and fully-automatic DDR 4/3 system is realized.

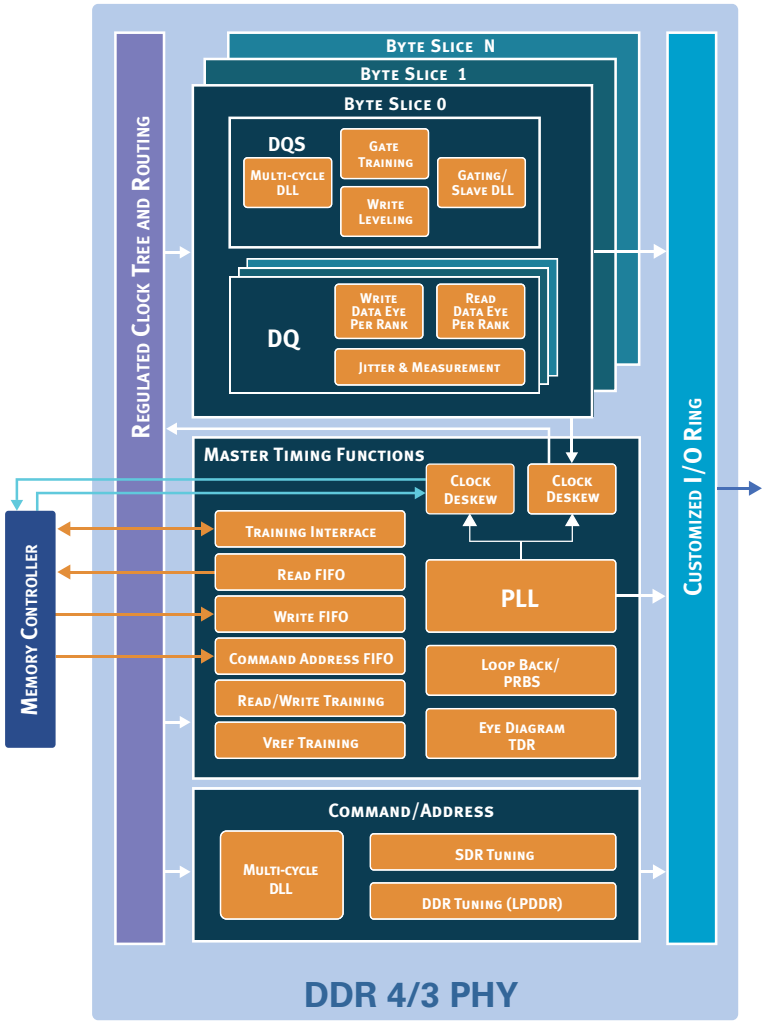

FEATURES

- Supports DDR4-2400, DDR3-2133 and LPDDR3-1066, simultaneously with one hard macro
- DFI 3.1 compliant
- Supports x4, x8 and x16 DRAMs
- Up to 144 bits wide
- Up to 4 chip selects, each with unique tuning
- Includes PLL, with frequency multiplication from low frequency reference
- Per-pin architecture, similar to a SerDes, automatically corrects skew, increases data eye and eliminates most parallel interface problems
- Continuous adjustment of read data eye and gate timing
- Automatic Training includes:
  - Multi-cycle read gate training and write leveling
  - Write data eye centering
  - Internal Vref adjustment
  - External Vref adjustment in each DRAM
- Localized, clock-deskewed PHY-to-memory controller interface to ease timing closure
- Full speed read/write BIST with pseudo-random data, mux-scan ATPG and 1149.1 Boundary Scan
- Circuitry in each pin able to measure the data eye and jitter, and calculate flight delays

Key Features	Benefits
Automatic Deskew	Skew among pins is automatically corrected; intentional skew can reduce SSO
Tuning	State-of-the-art tuning is the key to a high performance DDR system
Complete PHY	Completely assembled and validated hard PHY and I/O ring means no assembly is required and performance is guaranteed
Flexibility	Proprietary tools generate and validate a PHY fitted to the customer's die floorplan and package
Timing Closure	Memory controller to PHY timing closure is eased by a localized interface and clock deskew circuitry
Instrumentation	PHY resources can measure data eye and jitter per pin, speeding up board bring-up

http://www.truecircuits.com/ddr_phy.html

IP - Core

HIGH PERFORMANCE

TUNING FOR PERFORMANCE

The PHY has been designed from the ground up to provide extensive, automatic and continuous tuning. Each pin constantly adjusts separate read data eyes for even and odd data phases, taking jitter into account. Tuning is also done separately for each chip select value. Pervasive tuning is the key to performance.

TIMING CLOSURE

To make timing closure of the DDR 4/3 PHY to the memory controller faster, two elements are employed: 1) the interface from PHY to memory controller is localized and optimized for easy timing, and 2) the clock to the memory controller is driven by the PHY and is internally deskewed, minimizing clock-to-clock uncertainty.

AUTOMATIC TRAINING

DDR4 systems require a great deal of training to function properly. The TCI PHY, combined with an appropriate controller, does all of the required training with almost no user interaction. Low overhead, incremental training can be done at the user's discretion to achieve even higher data rates.

Automatic training includes multi-cycle read gate training and write leveling, and read and write data eye centering. DDR4 systems require internal Vref and external Vref (on DRAM) adjustment, which the TCI PHY performs automatically using very sophisticated algorithms.

TOOLS

TCI uses many proprietary tools to achieve a level of quality, flexibility and automation unseen in mixed-signal design, and not currently available in this type of hard IP.

EASY INTEGRATION

NO ASSEMBLY REQUIRED

The IP is delivered as a completely assembled hard macro. It is fully tested and verified with state-of-the-art timing analysis. Through a careful, joint process, the I/O ring and package are co-designed prior to PHY delivery, so that the PHY can be fully described, verified and delivered as a whole. Tremendous flexibility is allowed and no assembly is required.

LOWERING PACKAGE AND BOARD COSTS

Simpler and cheaper (fewer layer) chip packages and boards can be designed by eliminating the need for matched trace lengths, and by allowing for tremendous flexibility in the I/O ring/package co-development.

By intentionally skewing adjacent pins, lower cost power delivery systems can be used, and wire bond packages can be used at a higher speed.

MEASUREMENT RESOURCES FOR CHARACTERIZATION

The PHY contains many resources that can be set up to quickly characterize a new chip, a package or a customer's PCB board. Per pin measurements include: DQ switching jitter, read DQS jitter, read data eye, write data eye, Vref sensitivity and flight times. Pin and pattern weaknesses can be found quickly, without expensive lab equipment. Using an appropriate controller, the DDR interface can be fully characterized without CPU interaction.

TEST

The PHY includes a full speed read/write BIST, which tests the complete read and write paths of every pin simultaneously with pseudo-random data. The PHY design kits include industry-standard boundary scan chains and all the appropriate views for DFT.

CONTACT INFORMATION

True Circuits, Inc.
4300 El Camino Real, Suite 200
Los Altos, CA 94022
Phone: 650.949.3400
Fax: 650.949.3434
sales@truecircuits.com

DESIGN & VERIFICATION

DVCON 2014

CONFERENCE & EXHIBITION

PREMIER CONFERENCE FOR DESIGN & VERIFICATION ENGINEERS

Conference Sponsor: accellera SYSTEMS INITIATIVE

MARCH 3-6, 2014

DoubleTree, San Jose, CA

BOOST YOUR KNOWLEDGE BASE IN 2014

REGISTER FOR A HIGH-LEVEL TUTORIAL

MONDAY, MARCH 3

9:00am - 12:00pm

UVM – What's Now and What's Next
Using UPF for Low Power Design and Verification

2:00 - 5:00pm

Experience the Next ~Wave~ of Analog and Digital Signal Processing Using SystemC AMS 2.0

THURSDAY, MARCH 6

8:30am - 12:00pm

The How To's of Metric Driven Verification to Maximize Verification Productivity	cadence
Block to System Verification: Smooth Sailing from Simulation to Emulation	Mentor Graphics
SoC Verification Challenges Offer Opportunities to Take a New Look at Debug	SYNOPSYS
Formally Verifying Security Aspects of SoC Designs	JASPER design automation

2:00 - 5:30pm

Revolutionary Debug Techniques to Improve Verification Productivity	cadence
Formal Verification in Practice: Technology, Methods and Applications	Mentor Graphics
Optimizing Coverage-Driven Verification	SYNOPSYS

ADVANCE REGISTRATION RATES

DEADLINE: JANUARY 28, 2014

View full technical program & register:

DVCON.ORG

Register now and secure your ticket!
embedded-world.de

Nürnberg, Germany

February 25 – 27, 2014

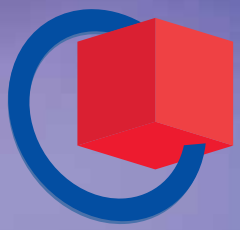

embedded world 2014

Exhibition&Conference

... it´s a smarter world

Be there to secure your head start!

Meet the Embedded-Community at the world's largest event for the Embedded-Technology. Obtain groundbreaking impulses for tomorrow's technology!

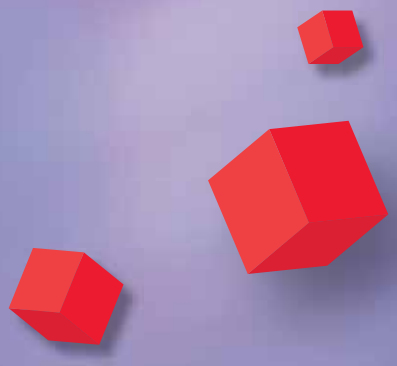

Media Partners

elektroniknet.de
computer-automation.de
energie-und-technik.de
karriere-ing.de
Markt&Technik
Elektronik
automotive
ENERGIE & TECHNIK
AUTOMATION
MEDIZIN+elektronik
DESIGN & ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER

Organizers

Trade Show
NürnbergMesse GmbH
Tel +49 (0) 9 11.86 06-49 12
visitorservice@nuernbergmesse.de

Conference
WEKA FACHMEDIEN GmbH
Tel +49 (0) 89.2 55 56-13 49
info@embedded-world.eu

NÜRNBERG MESSE

Run your design faster and get to market quicker.

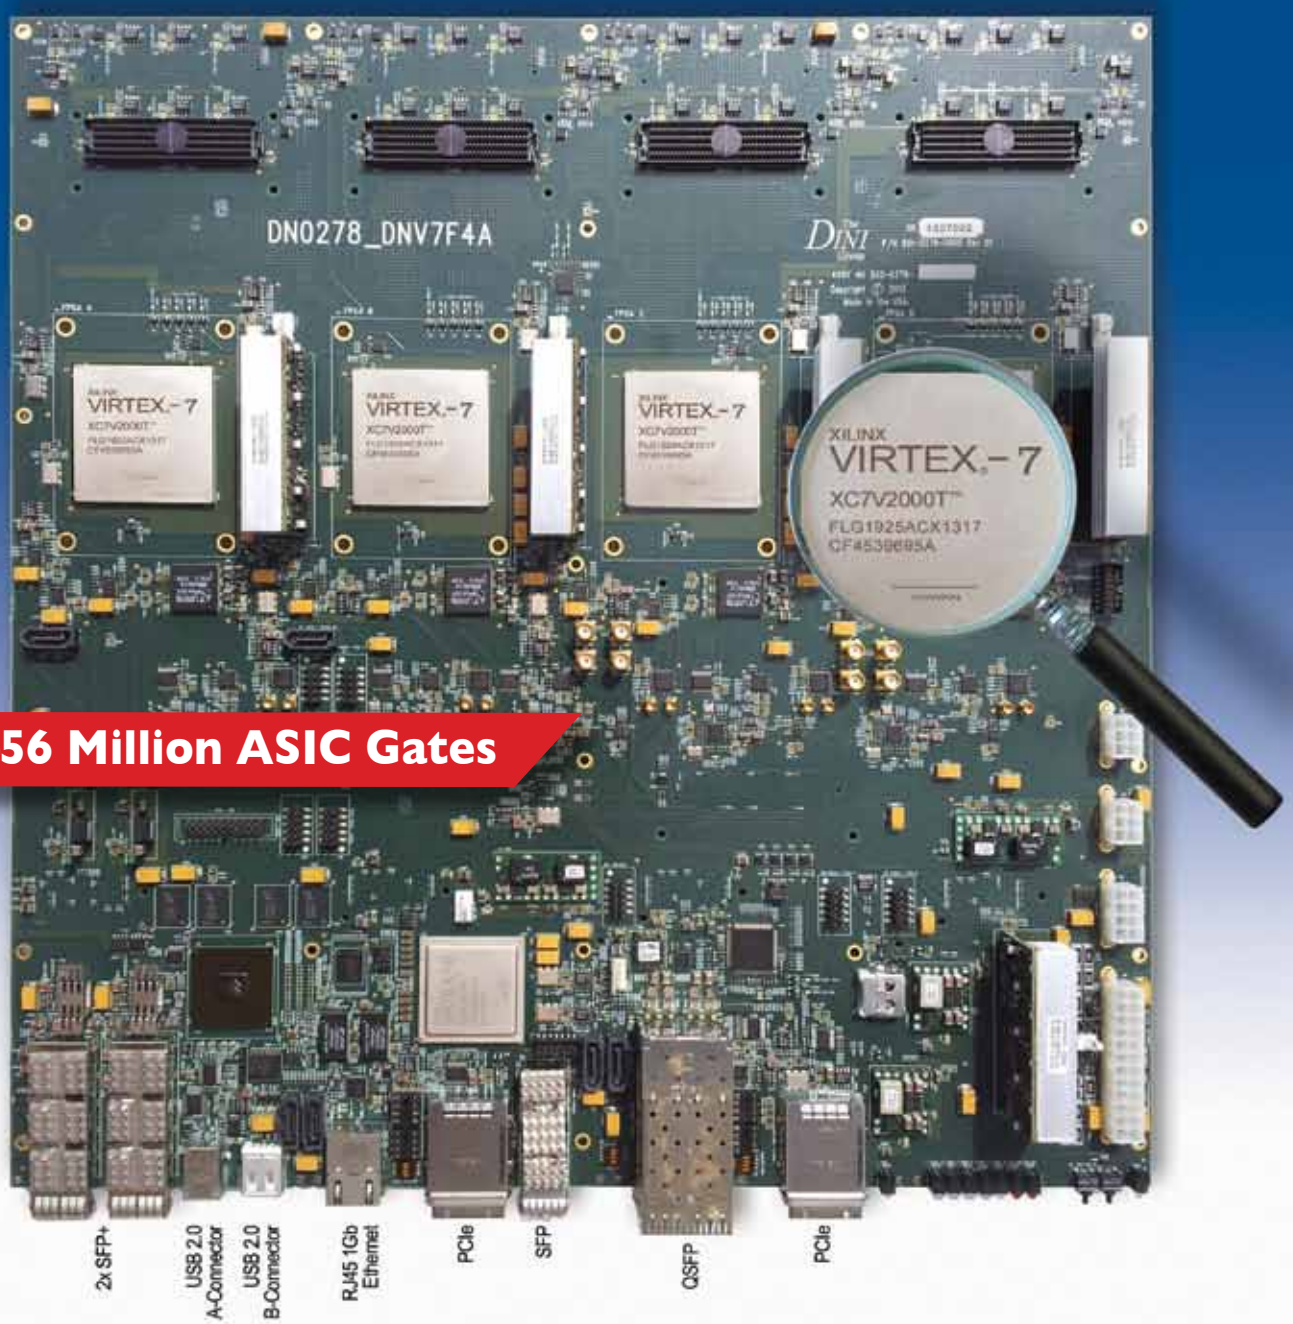

New—56 Million ASIC Gates

This new Virtex 7 based board is the biggest and fastest ASIC Prototyping platform—ever. All the features you need to run your big fast design are on this off-the-shelf logic prototyping system that captures all the capabilities of the new Xilinx 7V2000T. 14 million ASIC gates/chip, 1.4 Gb/sec LVDS, 1200 I/O per chip, and new software tools to get your RTL running faster, are all features of these 2.5D stacked silicon FPGAs from Xilinx. Four of these massive chips plus configuration FPGA and a Marvell Dual CPU for data handling, make it easier to get your design up to speed.

Is your design still too big? Seamlessly stack two boards together (a first for DINI) and 112 million gates are 100% user available. Need special functions, added memory, or signal processing? Eight daughter cards can support your unique requirements. There is an optional chassis available to ensure a robust and properly cooled system, because your design is going to run "hot".

Call DINI today, this surprisingly low priced board is now shipping.

DINI Group

7469 Draper Ave. • La Jolla, CA 92037-5026 • Phone: (858) 454-3419 • Fax: (858) 454-1728 • www.dinigroup.com