Embedded COMPUTING DESIGN

Connecting Silicon, Software, and Strategies for Intelligent Systems

Available for download on iTunes App Store and Kindle.

FOLLOW US:

JUNE 2013
VOLUME 11 • NUMBER 4

www.embedded-computing.com

SPECIAL ISSUE: MOST INFLUENTIAL WOMEN IN EMBEDDED

Victoria Mitchell
Director, SoC Software Engineering, Altera Corporation

Weili Dai
Cofounder, VP, and GM, Communications & Consumer Business, Marvell Technology Group

Jane Donaldson
President, Annapolis Micro Systems, Inc.

EMBEDDED's TOP INNOVATORS

MOST INNOVATIVE PRODUCT NOMINATIONS

Pg 41

Winners Announced in August Edition

www.opensystemsmedia.com

Portwell

Embedded with Unlimited Applications

Small Form Factor

Portwell's flexible business model and optimized SFF boards & systems have helped numerous Medical, Networking, and Industrial OEMs reduce costs, save space, accelerate development schedules and extend system longevity. Contact us to learn how we can tailor our embedded solution offering to your exact needs!

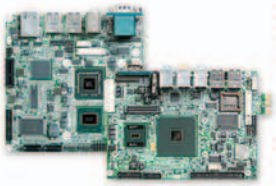

Small Form Factor Boards

- Small footprints to save space
- 120x120mm or 102x146mm
- Low power consumption
- Single-/ dual-core Intel® Atom™ and Intel® Core™ processors
- Extended temp: -40 to 80°C with Intel® Atom™ processor Z510PT Z520PT
- On board dual display, GbE LAN, PCIe, IDE or SATA and more

Small Form Factor Systems

- Compact turnkey solutions
- Fanless system with Intel® Atom™ or ULV Intel® Core™ processors
- All-in-one with LCD and touchscreen
- Ideal for digital signage, telematics, industrial automation, control and communications
- Expansion I/O options

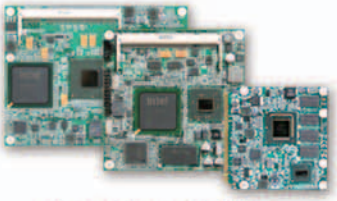

Computer-on-Modules

- Wide selection of COM Express and Qseven modules
- The latest Intel® Core™ i5/i7 and single-/ dual-core Intel® Atom™ processors
- Custom carrier board design and manufacturing
- Quick time-to-market solutions
- Extended temp and ECC memory support

Portwell's extensive product portfolio includes single-board computers, embedded computers, specialty computer platforms, rackmount computers, communication appliances, and human-machine interfaces.

We provides both off-the-shelf and versatile custom solutions for applications in the medical equipment, factory automation, retail automation, semiconductor equipment, financial automation, mission critical and network security markets.

American Portwell is both an ISO 9001:2008 and ISO 13485:2003 certified company.

intel Intelligent Systems Alliance Premier

Portwell

www.portwell.com

1-877-278-8899

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

Direct Seamless Connections with no Data Reduction
Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE
Between FPGAs and Standard Output Modules
Between FPGAs and Storage Arrays

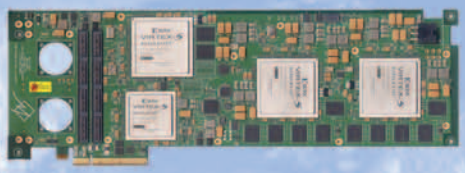

Ultimate Modularity
From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
Input/Output Modules Include:
Quad 130 MSPS thru Quad 550 MSPS A/D
1.5 GSps thru 5.0 GSps A/D, Quad 600 MSps D/A,
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10G, 40G or 100G Ethernet or SFPDP

VME/VXS/VPX, IBM Blade, PCI-X/PCI Express, PMC/XMC, MicroTCA

Annapolis
Micro Systems

No Other FPGA Board Vendor Streams This Volume of Data Real Time Straight Into the Heart of the Processing Elements and Then Straight Back Out Again

190 Admiral Cochrane Drive, Suite 130, Annapolis, Maryland USA 21401
wfinfo@annapmicro.com USA (410) 841-2514 www.annapmicro.com

Embedded COMPUTING DESIGN
www.embedded-computing.com
June 2013 | Volume 11 • Number 4

ON THE COVER

Embedded Computing Design editors have been on the lookout for this year's Top Embedded Innovators, and – for the first time this year – the Most Influential Women in Embedded. Our two contests pulled in many inspirational, highly qualified candidates who are forging new ideas and making a difference in the embedded industry. Read about the winners in this edition's exclusive Q&As, and check out the nominees for Most Innovative Product, winners to be announced in our August edition.

ISSN: Print 1542-6408, Online: 1542-6459

Introducing the AMD Embedded G-Series SOC.

AMD EMBEDDED SOLUTIONS

The Evolution is Complete

AMD's new Embedded G-Series SOC platform is a high-performance, low-power System-on-Chip design which includes even higher performance x86 CPU cores, integrated graphics capable of supporting beyond-HD resolutions across multiple displays and a fully-integrated I/O controller. As a result, AMD's new Embedded G-Series SOC enables you to deliver your next small form factor embedded design with an exceptional level of quality, reliability, longevity and innovative features.

Learn more at www.amd.com/g-series

Advertiser Information

Embedded COMPUTING DESIGN

OpenSystems media.

ECD Editorial/Production Staff

Warren Webb, Editorial Director
wwebb@opensystemsmedia.com

Sharon Hess, Managing Editor
sharon_hess@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Tom Varcie
Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker, Strategic Account Manager
rbarker@opensystemsmedia.com

Eric Henry, Strategic Account Manager
ehenry@opensystemsmedia.com

Ann Jesse, Strategic Account Manager
ajesse@opensystemsmedia.com

Christine Long
Vice President, Online Business
clong@opensystemsmedia.com

International Sales

Elvi Lee, Account Manager – Asia
elvi@aceforum.com.tw

Gerry Rhoades-Brown
Account Manager – Europe
gerry.rhoadesbrown@husonmedia.com

Christian Hoelscher
Account Manager – Europe
christian.hoelscher@husonmedia.com

Lauren Palmer
Account Manager – Europe
lauren.palmer@husonmedia.com

Regional Sales Managers

Barbara Quinlan, Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

Ron Taylor, East Coast/Mid Atlantic
rtaylor@opensystemsmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

OpenSystems Media Editorial/Production Staff

Warren Webb
Editorial Director
Embedded Computing Design
Industrial Embedded Systems
wwebb@opensystemsmedia.com

Sharon Hess
Managing Editor
Embedded Computing Design
Military Embedded Systems
Industrial Embedded Systems
sharon_hess@opensystemsmedia.com

John McHale, Editorial Director
Military Embedded Systems
jmchale@opensystemsmedia.com

Joe Pavlat, Editorial Director
xTCA and CompactPCI Systems
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VITA Technologies
jgipper@opensystemsmedia.com

Monique DeVoe
Assistant Managing Editor
VITA Technologies
EDA Digest
DSP-FPGA.com
mdevoe@opensystemsmedia.com

Brandon Lewis
Associate Editor
xTCA and CompactPCI Systems
PC/104 and Small Form Factors
blewis@opensystemsmedia.com

Curt Schwaderer
Technology Editor

Steph Sweet
Creative Director

David Diomede, Art Director

Joann Toth, Senior Designer

Konrad Witte, Senior Web Developer

Matt Jones, Web Developer

Editorial/Business Office

Patrick Hopper, Publisher
Tel: 586-415-6500 ■ Fax: 586-415-4882
phopper@opensystemsmedia.com

Subscriptions Updates

www.opensystemsmedia.com/subscriptions
Karen Layman, Business Manager
30233 Jefferson
St. Clair Shores, MI 48082

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

Wayne Kristoff, CTO

16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466

Tracking Trends in Embedded Technology

By Warren Webb

@warrenwebb wwebb@opensystemsmedia.com

Top Innovators streamline embedded technology

As the embedded industry transitions from stand-alone, fixed-function devices to networks of fully interconnected intelligent systems, the new product development process is also changing. For each new project, the embedded designer must constantly evaluate the latest ideas and innovations to deliver the highest-performance, most cost-effective products in the shortest possible timeframe. To help with this process, this issue of *Embedded Computing Design* presents our Second Annual roundup of the top innovative products and embedded innovators that together are responsible for dramatic changes in our industry. Our editorial team selected the top innovators from multiple entries based on creativity, technical guidance, and unique contributions to the embedded industry. The winners of the 2013 *Embedded Computing Design* award for Top Embedded Innovator are:

- Josh Lee, Cofounder and CEO at Uniquify
- Darren Humphrey, VP and CTO at DiSTI
- Adnan Hamid, CEO at Breker Verification Systems

You can read all the details about each winner and their contributions in the Top Embedded Innovator question and answer section of this issue. For the first time ever, this year we have also recognized the Most Influential Women in Embedded. These are women who are providing technical leadership, thriving in technology management roles, and making a difference in the embedded industry. The winners of the 2013 *Embedded Computing Design* award for Most Influential Women in Embedded are:

- Weili Dai, Cofounder and VP, Marvell Technology
- Vicki Mitchell, Director of Software Engineering, Altera
- Jane Donaldson, President, Annapolis Micro Systems

The three winners of our Top Innovative Products contest – one each in the Software, Silicon, and Strategies categories – will be announced in our August edition. (The nomination process is closed.)

In addition to the special embedded innovation coverage, this issue digs into the details of several topics including multicore processors, Electronic Design Automation (EDA) technology, operating system selection, and small form factor platforms. In the Silicon section, Markus Levy, Founder and President of the Embedded Microprocessor Benchmark Consortium (EEMBC), describes the technical challenges faced by multicore developers and gives his projections on the power and performance gains we can expect over the next few years. Extending the Silicon coverage, John Goodacre, Director of Technology and Systems at the ARM Processor Division, answers questions on portable multicore tech and explains how ARM's big.LITTLE architecture can deliver both higher performance and extended battery life.

In this edition's Strategies section, industry experts take a look at some of the latest technologies at work to expand the performance levels and lower the cost of small form factor modules. With design flexibility in mind, Kevin Roth, an Electrical Engineer at Alpha Data, reveals how the increased processing power and reconfigurable I/O of the latest generation of FPGAs have extended the design possibilities with VPX architecture. Along with an objective of increasing student interest in computer science and engineering, Pete Lomas, Director of Engineering at Norcott Technologies and cofounder of the Raspberry Pi foundation, presents a case study of the design challenges his team faced in developing a low-cost, credit card sized single board computer.

In this edition's Software section, a roundtable of embedded industry experts answered questions covering the technology trends, standards, and future projections in the EDA marketplace. Their responses included software and Intellectual Property (IP) for design verification, virtual prototyping, and a host of proprietary solutions. Participants in the roundtable are:

- Wally Rhines, Chairman and CEO, Mentor Graphics
- Brett Cline, Vice President, Forte Design Systems
- Marc Serughetti, Business Development Director, Synopsys
- Michał Siwiński, Director of Product Marketing at Cadence
- Bill Neifert, Cofounder and CTO, Carbon Design Systems

Extending the software coverage of this edition, my article outlines some of the basic requirements and evaluation steps necessary in the embedded operating system selection process.

If you have ideas for future articles and coverage in *Embedded Computing Design* that would help in your development efforts, please let us know. We are always interested in contributed technical articles, guest blogs, or videos that would appeal to other embedded designers.

Q & A

Moving target: EEMBC evolves its benchmark suites to keep pace with the multicore revolution

Q&A with Markus Levy, Founder and President of EEMBC

By Sharon Hess, Managing Editor

Constantly evolving, the EEMBC's focus has shifted from benchmarking microprocessors to SoCs and even entire systems, targeting the automotive, networking, and smartphone/tablets markets. Additionally, the emergence and constant expansion of the multicore processor has brought added challenges as EEMBC benchmarks such real-world, yet ever shifting, multicore targets.

Remind us briefly why the EEMBC was started, by whom, what its mission is, and what it provides to which industries.

LEVY: I started EEMBC in 1997 (with some guidance by guys such as Derek Meyer and Geoff Lees) because we realized that there was a clear need in the embedded industry for standardized benchmarks that all processor vendors could utilize. At the time, there were de facto standard benchmarks such as the infamous Dhrystone MIPS, but more often than not, processor vendors used their own proprietary benchmarks, which, of course, meant that there was no way to make apples-to-apples comparisons.

EEMBC has evolved considerably since its inception, with benchmarks going from a microprocessor focus to an SoC focus and even benchmarking the entire system. However, the mission has remained the same – to provide real-world benchmarks to help designers select the right embedded processors for their systems and to ensure that the benchmarking methodology will provide fair and reasonable comparisons. The EEMBC benchmarks are targeted at a variety of industries including automotive, networking, and smartphones/tablets.

Since multicore processors' emergence onto the embedded scene in recent years, what are the top 3 technical challenges faced by the embedded computing industry and the EEMBC?

LEVY: The top three technical challenges are software, software, and software. Seriously, with the emergence and continued expansion of multicore processors, system developers are challenged with either converting their legacy applications to take advantage of the parallel processing offered by multicore or by understanding how to create new applications that can utilize some of the extremely complex hardware accelerators included in many newer-generation SoCs. For example, even if a legacy application utilizes a threaded programming model, moving to multicore entails dealing with issues such as data races. Alternatively, writing an application from scratch to deal with complex hardware accelerators is a daunting task, and programmers are often unwilling to take this risk, especially considering that the SoC functions could change on the release of the next-generation device.

As this topic relates to EEMBC, you have to look back a decade ago when processor benchmarks were entirely focused on single-core devices. But over time, as the processor evolved into the multicore SoC, benchmarks have also had to evolve in order to adequately measure these application-specific and complex devices. In other words, you're not able to judge an SoC's capabilities just by running a standard benchmark on the processor core (although EEMBC's CoreMark has certainly grown in popularity). Hence, the benchmarks have evolved into system tests, such as BrowsingBench, which tests the page-load performance of a complete mobile phone. Likewise, EEMBC's networking benchmarks have evolved into DPIBench, which tests the sustained throughput of full-fledged routers while performing deep packet inspection searching for virus-infected files.

While BrowsingBench and DPIBench are system-level benchmarks, they are implicitly testing the underlying multicore hardware (as well as software stacks). Take BrowsingBench as an example. The primary function of this benchmark is to perform html page loads on a client device (i.e., smartphone, tablet). The page-load operation involves a series of steps that includes serial functions (i.e., enter or click URL, fetch initial HTML, parse the HTML, and determine the workload) and parallel functions that could take advantage of a multicore device (i.e., parsing, Javascript, image decoding, page rendering after all elements have been assembled, and animation). Inherently, if the client device's browser and operating system are designed appropriately, we will see significant performance improvements on these parallel functions.

How are these technical problems best solved, by industry and the EEMBC?

> Although the silver bullet for multicore programming will never arrive, programmers are learning to work with what's available and are changing their expectations, realizing that the move to multicore will require a fair amount of effort.

LEVY: EEMBC is solving these problems by continually evolving its benchmark suites to take advantage of (and stress) multicore devices, and at the same time the benchmarks are becoming much more real-world in nature, because that's the only way to accurately reflect the device's capabilities. However, the industry as a whole needs to continually strive for better development tools that will help analyze and automate the application's ability to take advantage of the multicore hardware.

Other industry organizations, such as the Multicore Association, are putting together Application Programming Interfaces (APIs) that will make it easier for programmers to develop portable and scalable applications. These APIs will also allow semiconductor vendors to enable their customers to take advantage of the (sometimes) extremely complex hardware accelerators, many of which don't even get used today because customers are unwilling or unable to make the investment to read a 1,000-page manual for programming these accelerators. PolyCore Software is an example of one company that has already commercialized tools based on these APIs; these tools allow a system developer to easily assign software functions to specific cores (or nodes) within the system. But more importantly, the tools make it much easier to move those functions to different cores to take advantage of a changing SoC topography.

Over the next 5 years, what characteristics and features will contribute most to multicore processor performance gains? Will it be enough to keep up with embedded industry demands?

LEVY: Certainly, vendors will continue to add more cores and accelerators to increase performance, and hopefully they will also provide the tools to support this. Another hardware-related area that will be beneficial will be increasing memory throughput and higher bandwidth – somehow you have to feed these hungry multicore processors. Software tools will also continue to improve: A tremendous amount of effort is going into this area from a multitude of vendors such as CriticalBlue and PolyCore Software. Although the silver bullet for multicore programming will never arrive, programmers are learning to work with what's available and are changing their expectations, realizing that the move to multicore will require a fair amount of effort.

Which types of multicore processors have the market edge now? Why? What about 2 years from now? 5 years from now and beyond?

LEVY: The types of processors with the market edge today are entirely dependent on the market that they play in. For example, in some areas, high-performance SMP devices are best utilized, especially for applications that have traditionally used programming frameworks such as Pthreads.

However, inevitably, the types of processors that will succeed in the future will be the SoCs that provide hardware-accelerated functions. It's the only way that applications will be able to meet their performance-power budgets. In other words, with homogeneous SMP devices, the performance gained by increased core count is not scalable. For example, the more cores that share a common bus structure, the more that each core must compete for memory bandwidth. This problem can be alleviated by designing chips that divide cores into clusters, where each cluster can operate autonomously if necessary.

What plans does the EEMBC have to expand its offerings in the future, and how can the industry get involved?

LEVY: As I mentioned earlier, EEMBC has already begun expanding its offerings to address the system-level performance measurement. I see this trend continuing (especially in the mobile, automotive, and networking areas), although there will always be a need for C-coded benchmarks that purely address the CPU + memory subsystem. As an industry association, EEMBC always encourages companies to get involved with the definition and development of new benchmarks. Our philosophy is "the more the merrier" because it helps ensure comprehensiveness and fairness. Companies interested in joining EEMBC or the Multicore Association can contact me directly for more information. ECD

Markus Levy is founder and President of EEMBC. He is also President of the Multicore Association and Chairman of the Multicore Developer's Conference. Markus began his career at Intel Corporation, where he served as both a Senior Applications Engineer and Customer Training Specialist for Intel's microprocessor and flash memory products. He was also previously a senior analyst at In-Stat/MDR and an editor at EDN *magazine, focusing in both roles on embedded processors. He is the coauthor of the book* Designing with Flash Memory *and holds several patents related to flash memory architecture and usage as a disk drive alternative. He is also a volunteer firefighter.*

EEMBC • markus.levy@eembc.org • www.EEMBC.org

Follow:

Q & A

ARM's big.LITTLE architecture aims to satisfy the hunger for power

Q&A with John Goodacre, Director, Technology and Systems, ARM Processor Division

By Sharon Hess, Managing Editor

As smartphone and tablet users continue to demand much higher performance to keep pace with an ever-more-connected lifestyle, the demand for extended battery life naturally follows. However, one way ARM is meeting the challenge is with its big.LITTLE multicore-based architecture, which aims to boost power savings by as much as 95 percent to help satisfy today's power-hungry computing paradigm. Additionally, John discusses how preserving embedded market technologies' safety-critical design features while maintaining predictability and amping up performance using multicore processors can present a daunting challenge.

Anyone who hasn't heard of ARM has clearly been living under a rock, at least in regard to smartphone offerings. But describe (briefly) your multicore offerings for automotive infotainment, smart meters, and embedded computing.

GOODACRE: ARM characterizes the market into three primary segments, those that require the support of rich operating systems, those that require true real-time and predictable execution, and those that will be embedded within microcontrollers, also known as the Application, Real-time, and Microcontroller profiles.

The same kind of technology that goes into smartphones is also applied to automotive infotainment. We provide Cortex-R and Cortex-M series processors for embedded computing markets and smart metering applications that are low power and can operate for extremely long periods of time on a single battery. For application and embedded computing markets alike, our partners utilize the Cortex-A9 with ARM graphics- and video-processing engines to deliver products like the Ford Sync and many of Samsung and LG's Smart TVs.

What have been the top technical challenges your customers have faced in those markets? How can those challenges be solved?

GOODACRE: As a whole, the processing needs for different markets can be clustered around common sets of requirements. The Application profile processors have been driven by the peak performance available to a system running on a platform operating system. This profile is typified by those of mobile, infotainment, and other consumer devices such as the Smart TV. This has necessitated multicore support for an SMP-based operating system where the tasks of the system can be automatically shared among multiple processors. Sharing tasks across the processors, especially when arranged in a big.LITTLE arrangement, allows the required performance to be shared across the cores, and as such deliver the performance also in a lower power envelope. In the highlighted markets, smartphones and auto infotainment are addressed by this Cortex profile.

The Real-time profile processors have also been driven by peak available performance; however, this segment utilizes an RTOS, and as such defines specifically where and when specific tasks will run. These systems must also maintain the predictability and safety characteristics required by many markets. In the highlighted markets, most embedded computing is addressed by this Cortex profile. For example, within a hard disk, the head must be positioned exactly at the right place, at the right time, so as to read the data as it spins past under the head. In automotive, when the driver presses the brake, that level of braking must be applied to the physical brakes.

Finally, the Microcontroller profile is similar to the Real-time profile in its use of an RTOS; however, it is more driven by the needs of embedded flash than the scalability of real-time performance. As such, these microcontroller parts, assuming they use more than one processor, will deploy them as independent systems within the multicore microcontroller. The smart meter is addressed by this profile of Cortex processor because of its extremely small size, and hence low cost, while also consuming little power – allowing battery life measured in years or scavenged from the environment around the sensor.

What is the biggest challenge right now for ARM in engineering multicore processors?

GOODACRE: One of today's most significant challenges is how to create an SoC that meets the conflicting consumer

demand for devices with both higher performance and extended battery life. Mobile usage has changed significantly and today's consumers are increasingly using their smartphone for the majority of their connected lives. Because of that, the performance demanded of current smartphones and tablets is increasing at a much faster rate than the capacity of batteries or the power savings from semiconductor process advances. At the same time, users are demanding longer battery life within roughly the same form factor.

Each profile of processors has its own technology challenges. The Application challenges are around how to support even more processors while maintaining support for the software models used by this market. The interest from the enterprise and server markets in using ARM technology also drives challenges around how to utilize many dozens, or significantly more, ARM processors in a single system. For embedded markets, maintaining the predictability and safety-critical design features while increasing performance through multicore also presents its own challenges.

Briefly describe your big.LITTLE technology, introduced in 2011 and recently described as gaining more momentum. How does it work, technically speaking?

GOODACRE: big.LITTLE processing is no more complex for software applications than today's SMP capable operating systems. However, under the hood, it is an energy savings method where high-performance CPUs and efficiency tuned CPUs are connected in a cache-coherent combination so the operating system can dynamically assign application tasks to the appropriate CPU based on performance needs.

The more powerful "big" core is responsible for handling computing intensive tasks, such as rendering a Web page, whereas the less powerful "LITTLE" core handles lesser demanding tasks, such as MP3 playback. Both the cores implement exactly the same processor architecture (ARMv7), and are capable of executing the same instructions. The only difference lies in the way the cores handle the execution. While the "big" core is designed with performance as its primary goal, the "LITTLE" core is designed with efficiency as its principal target. Thus, an application or program run on one core can also run on the other without knowing any difference except with different performance and power consumption levels.

Can big.LITTLE be accomplished with single-core processors, or just multicore?

GOODACRE: big.LITTLE is fundamentally based on a multicore system where a given processor has the ability to more power efficiently execute the same software as a higher peak performance core can. The big.LITTLE architecture is designed using the technique of employing separate cores with different computing powers within the same system.

This asymmetry in power efficiency in processors can be realized by running a specific CPU at a different voltage to another, or by implementing a CPU in a more power-efficient manner. The highest dynamic range and hence amount of power savings, however, is delivered when a CPU can use both the voltage and implementation aspects, but most importantly is built using a fundamentally more power efficient microarchitecture such as that realized by the Cortex-A7 alongside the Cortex-A15.

Does the big.LITTLE estimated 70 percent savings on processor energy consumption apply only to smartphones?

GOODACRE: big.LITTLE processing is currently targeted at the smartphone market. That said, the power savings is driven by the dynamic range in the required performance of an application. If the application only required high performance 5 percent of the time, then the savings would be around 95 percent. If an application only leaves the high-performance case 5 percent of the time, then the savings would be closer

to 5 percent. Thankfully, the big.LITTLE system can move tasks between these states very quickly, in the manner of a few 10's of a nanosecond; this means that even the most predictable high-performance systems also spend more than expected in the lower-performance state.

What are the emerging multicore trends for the aforementioned markets? How does ARM plan to keep pace with these trends?

GOODACRE: Each market has its own specific requirement when it comes to answering its multicore demands of peak performance, power efficiency, real-time nature, and embedded characteristics. ARM has structured its R&D to address these multiple markets while extending their capabilities into new markets, both with higher peak performance designs such as through the 64-bit capable Cortex-A57, but also in the efficiency requirements of the most embedded devices through the Cortex-M0+. ARM's road map contains devices that support each of these market's trends, and will continue to do so. The trend for performance in enterprise scale out drives need into the interconnect products such as CCI400 and CCN504 in utilizing the most power-efficient processor for the required single thread performance level.

Today, again, mobile devices are at 4 processors and moving to 6 or 8 utilizing big.LITTLE. Auto infotainment will follow this through its commonality with the Cortex Application profile. The embedded markets primarily still dedicate a specific processor to a specific task, so in smart meters it's more likely to still use a single processor, or potentially two: one for regular data capture and one for bursting the data across a network. The trend across many embedded markets is to make the intelligence of the system visible, whether a washing machine hosting an LCD or a meter emailing you the cost of your energy. This all leads to an increased use of processors, both individually and in multicore designs. ECD

__John Goodacre__ is Director, Technology and Systems, ARM Processor Division. He joined ARM in February 2002 with his current position focusing on various programs around the application processor's road map including the definition and market development of the ARM multicore technology. Prior to ARM, he specialized in enterprise software, having worked for Microsoft as Group Program Manager, delivering Exchange 2000 Server. Graduating from the University of York with a B.Sc. in Computer Science, John has more than 25 years of experience in realizing new technologies across various markets.

ARM • www.arm.com/products/processors/index.php

Follow:

ROUNDTABLE

EXPERT PANEL: Is EDA as easy as 1, 2, 3 these days?

By Sharon Hess, Managing Editor

QUESTION 1

Remind us briefly about your organization, when started, what its technical focus is, and what it provides to which industries.

No doubt about it. Today's ever-shifting SoC, IC, PCB, and electronic systems design paradigm is burdened with challenges including: soaring development costs/complexity, product differentiation dilemmas thanks to more commercial IP usage, the perpetual shift to finer-grain silicon geometries, and gaps between hardware design and software development timeframes and methodologies. These all make the engineer's job … complicated. Throw in the latest trends and it's even more so. The good news, though, is that the Electronic Design Automation (EDA) industry aims to smooth this rocky path for its customers. Accordingly, in the following virtual discussion, we asked 5 top EDA industry experts 6 questions to find out how they are helping their customers solve these challenges today. We also found out how the industry has evolved and what the hottest trends are now.

***Brett Cline** is VP of Marketing and Sales for Forte Design Systems. Previously, he was Director of Marketing at Summit Design and held positions in development, applications, and technical marketing at Cadence and General Electric. He holds a Bachelor of Science degree in Electrical Engineering from Northeastern University.*

***Marc Serughetti** is Director of Business Development at Synopsys, where he drives deployment of virtual prototypes and embedded software technologies. He has more than 18 years' experience in software development technologies, having led product marketing and business development teams at Integrated Systems, Wind River, and CoWare.*

***Walden (Wally) Rhines** is Chairman and Chief Executive Officer of Mentor Graphics. Prior to joining Mentor Graphics, he was Executive Vice President of Texas Instruments' Semiconductor Group, sharing responsibility for TI's Components Sector and having direct responsibility for the entire semiconductor business comprising more than 30,000 people.*

***Michał Siwiński** is the Senior Director of Product Marketing for System and Software Realization Group (SSG) at Cadence, responsible for strategy, products, partnerships, business development, and communications. He focuses on the realization of the EDA360 vision by concentrating on emerging and established offerings used by system and semiconductor companies.*

***Bill Neifert**, CTO and Cofounder of Carbon, has 13 years of electronics engineering experience and more than 18 years in EDA including C-level design and quickturn systems. Bill has designed high-performance verification and system integration solutions and developed an architecture and coding style for high-performance RTL simulation in C/C++. He has Bachelor of Science and Master of Science degrees in Computer Engineering from Boston University.*

Brett Cline
Vice President of Marketing and Sales
Forte Design Systems

CLINE: Forte Design Systems helps designers improve their productivity and quality of results with SystemC high-level synthesis. We've been doing this since the early 2000s. In addition, Forte has a number of sophisticated Intellectual Property (IP) offerings anchored by its floating-point implementations. Forte's software is used by virtually every market segment doing digital design and its IP is primarily used by 3D graphics and GPU applications.

Marc Serughetti
Director of Business Development
Synopsys

SERUGHETTI: Synopsys, Inc. provides EDA products and services that accelerate innovation in the global electronics market. Our system-level, IP, implementation, verification, manufacturing, optical, and FPGA solutions help address the key challenges designers face such as power and yield management, system-to-silicon verification, and time-to-results. Founded in 1986, Synopsys expanded into the prototyping space with the acquisition of the COSSAP product in 1994 and focuses on enabling companies to accelerate both the development and deployment of virtual prototypes via next-gen virtual prototyping technology.

Michał Siwiński
Senior Director of Product Marketing for System and Software Realization Group (SSG)
Cadence

SIWIŃSKI: Cadence Design Systems, Inc. was formed in 1988 by the merger of ECAD and SDA Systems. Today Cadence is a major provider of EDA software, hardware, and IP for the design of ICs, PCBs, and electronic systems. Cadence provides end-to-end solutions for custom/analog design, digital IC design, silicon-package-board, and functional verification, as well as design and verification IP. Major vertical market segments served by Cadence include mobile, consumer, and communications, as well as automotive, medical, mil/aero, and industrial. Main horizontal market segments served are semiconductor companies and system companies.

Walden (Wally) Rhines
Chairman and Chief Executive Officer
Mentor Graphics

RHINES: Mentor Graphics provides electronic hardware and software design solutions for worldwide electronic, semiconductor, and systems companies. We offer a broad portfolio of best-in-class hardware and software design solutions focused on IC design and physical verification, functional verification, FPGA/PLD, design-for-test, PCB design, and embedded software.

Bill Neifert
Chief Technology Officer and Cofounder
Carbon Design Systems

NEIFERT: Carbon Design Systems was founded in 2002. Our primary focus is to provide virtual prototypes that can be used throughout the design process, from IP selection and architectural exploration to O/S boot and software development. Instead of providing a range of different virtual prototypes to address the different speed and accuracy requirements through the design cycle, we provide a unified virtual prototype. It is capable of running at tens to hundreds of MIPS but also can run with 100 percent accuracy. Industry teams designing using virtual prototypes are focused primarily on developing leading-edge SoCs, especially in the mobile, consumer, and networking markets.

QUESTION 2

Which technology has been the most influential in forging the embedded industry's current EDA paradigm?

CLINE

Most likely, virtual system prototyping, such as the tools supplied by Carbon Design Systems. High-level models used in products like Carbon's can be utilized in Forte's HLS flow.

SIWIŃSKI

EDA enables the development of ICs with billions of transistors at advanced semiconductor process nodes, making it possible to build today's smartphones, tablets, servers, and other electronic products. One influential technology that made all this possible was the move from gate-level design to Register Transfer Level (RTL) design in the 1990s. RTL design is a much higher level of abstraction, and it permits design using specialized programming languages (Hardware Description Languages or HDLs) such as IEEE standard Verilog as opposed to wiring gates together. Tools that enable RTL design include logic synthesis and RTL verification.

Today the abstraction level is moving even higher to Transaction Level Modeling (TLM) using the IEEE standard SystemC programming language. This is made possible by such technologies as high-level synthesis, which can automatically convert SystemC into RTL, and virtual prototyping, which allows early software development using SystemC models. While EDA has historically focused on IC and System-on-Chip (SoC) development and verification, a new paradigm is emerging in which EDA tools are facilitating hardware/software codevelopment and coverification, due to the dominant role that software is starting to play in today's electronics. As a result, the gap between traditional EDA and embedded software development is closing.

NEIFERT

The embedded technology requirement driving EDA approaches now is more a realignment of design from being a hardware-focused task to being a more software-focused one. While there are fewer and fewer hardware starts, the number of software engineers continues to grow rapidly. Traditional EDA companies continue to make most of their revenue selling back-end implementation tools for the design cycle, which continues to grow in complexity. They are making attempts to develop a pricing model to sell tools to the large group of embedded software designers.

SERUGHETTI

Today's electronic systems – ranging from a System-on-Chip (SoC) to an electronic device/product – rely heavily on software executing on a chip. Software executing on a chip and prototyping are very closely linked to deliver to the end user the expected functionality and performance and today represent half the cost and half the time-to-market of an SoC design. They also have very important implications on the SoC development. Designers must get the architecture right, as no amount of downstream tools will compensate for a fundamentally wrong architecture.

Thus, prototyping has been the key to moving away from a standard project flow where architecture design, SoC hardware development, manufacturing, and software development, hardware/software integration, and system validation are done serially to a development process; this enables the parallelization of SoC hardware development with software development integration and system validation. Prototyping includes virtual and FPGA-based prototyping with the objectives of architecture design, early software development, hardware/software integration, and system validation.

RHINES

Codesign and verification of embedded software with hardware has been the most important technology advancement for embedded development in EDA. Issues almost always occur in the "system integration" phase, when the hardware and software are finally tested together, causing major program delays, functional problems, cost overruns, and unhappy customers.

Our company first tried addressing the system integration problem in 1995, when we rolled out our corporate "integrated system design" strategy. The goal was to compress the time and effort required to do system integration by having both the hardware and software work in the same integrated environment. What we failed to appreciate was the enormous gulf between the hardware and software engineers, who work in their native environments with different toolsets. To bridge this gap, we recently announced the Mentor Embedded Virtual Platform, suitable for both hardware and software teams.

QUESTION 3

What are the three hottest EDA technology trends right now?

SIWIŃSKI

One clear trend is the move to a higher level of abstraction, sometimes called Electronic System Level (ESL). Hardware/software codevelopment is made possible by such tools as virtual prototyping, simulation, acceleration and emulation, and FPGA-based prototyping. An integrated set of development platforms, such as the Cadence System Development Suite, for example, can serve these needs, enabling early software design ahead of the accurate hardware representation or with mixed abstractions, hardware-software integration, and system validation of the hardware running hardware-dependent software.

Another trend is support for semiconductor process nodes at 20 nm and below. Here, lithography is so difficult that a technology called "double patterning," which requires extra mask layers, is required. This is resulting in many changes to IC implementation tools. Additionally, a new transistor technology – the FinFET – will be supported by most foundries at 16 nm or 14 nm. While FinFETs promise huge power and performance advantages, IC design tools must adapt to support them.

A third new trend that's received a lot of notice is the emergence of 3D-ICs. This may involve multiple chips laid side-by-side on a silicon interposer layer (called "2.5D" IC), or true 3D stacking in which chips are placed on top of each other and connected with Through-Silicon Vias (TSVs).

RHINES

1) For EDA's core hardware design market, the semiconductor industry must continue to deliver increased functionality per dollar with lower power, as it has done so reliably over the past half century. However, shrinking feature dimensions can no longer shoulder the entire burden, because of the rising costs associated with reliably manufacturing 28 nm devices and below. Semiconductor manufacturers are now deploying new approaches to achieve the same results: multi-die packaging, 3D stacks, interposers, etc., which will become increasingly important and require a wide range of newer EDA solutions, many already available.

2) More exciting for EDA going forward is the growth of embedded software development. Software design now consumes more than half the effort required to create an SoC or system. And the rapid growth of hardware acceleration products (close to 100 percent per year the past two years) has created a verification platform for hardware and software that is replacing simulation for the leading-edge chip and system development programs.

3) There is promise of growth of EDA applications in system design beyond PCB layout and Electronic System Level (ESL), although both are expanding again in response to new challenges. The opportunity lies in addressing the enormous challenge of designing and verifying complex electronics for macro systems such as automobiles, aerospace, commercial vehicles, and other equipment. EDA for automotive and aerospace applications is growing faster than the market for IC design software.

SERUGHETTI

The hottest EDA technology trends, specific to embedded systems, relate to virtual prototyping. Among them are earlier and more robust architectural exploration and analysis. There are more functions, more software, and more resource sharing in systems that are dynamic. As a result, it makes performance hard to predict and the hot questions from architects become: "Will my SoC architecture meet the performance requirements … for the required combinations of application use cases? … without overdesign? This architectural exploration/analysis needs to be done before software is available, via industry tools that enable earlier analysis while avoiding both under- and over-design.

Another hot trend is early software development. From a software development perspective, the integration of increasingly complex hardware and software is a significant challenge for semiconductor and OEM companies developing next-generation wireless, consumer, and automotive devices. Traditional methods of serialized hardware and software development – where the vast majority of software is developed and verified after the silicon design is complete – often fail to meet aggressive product development schedules. Virtual prototyping enables software engineers to start development months before the hardware design is complete, enabling full system bring-up to occur within days of silicon availability.

CLINE

Electronic System Level (ESL) continues to be one of the hottest areas in EDA. Of course, high-level synthesis is only one aspect of ESL. Others, including virtual system prototyping, design modeling, and verification are also in the forefront.

IP reuse has been a trend for two decades. We are now seeing IP developed at a higher level of abstraction, making it much more usable. While Register Transfer Level (RTL) IP reuse has sustained momentum longer than most people thought it would, behavioral IP is starting to gain traction with several startups in this area appearing over the past year or two.

Third, but certainly not last, is anything that has to do with low power. Everyone wants power to be lower, end of story. As it turns out, the earlier that power is considered in the design process, the more control designers have over the power. ESL happens to help with that as well.

NEIFERT

Most EDA advances seem confined to back-end implements processes, far removed from the daily needs of most embedded designers, though necessary to enable many advanced designs. However, from the front-end design and software perspective where Carbon focuses, three hot trends are emerging:

1. The restructuring of EDA toward the design IP space: Synopsys is already [a major] IP company … The string of recent purchases by Cadence demonstrates that this is an important part of its strategy as well.
2. The consolidation of IP offerings as packaged subsystems: As designs get bigger, blocks being purchased and reused are moving from single design blocks to entire compute subsystems, complete with software, changing the rules on what companies need to do to differentiate their end products.
3. The strong push to introduce accuracy into virtual prototypes: As emulation becomes ubiquitous, there is a strong push to integrate it with virtual prototypes as a way to introduce the accuracy needed throughout the design cycle.

QUESTION 4

What are the biggest EDA engineering hurdles these days – faced by your organization and also by your customers? How are these challenges being solved?

NEIFERT

The biggest problem I see is how to differentiate products in an environment where more and more of the IP is coming from commercial sources. This means that there are typically fewer areas in which to differentiate, at least using traditional hardware approaches.

Companies are solving this by differentiating using varied approaches. They are doing so by configuring their IP uniquely so that various pieces of IP play together in an optimal fashion. This is accomplished using a virtual prototype to sweep through the matrix of configuration settings for various pieces of IP while running the real system software to find which combination yields the optimal result.

They're also optimizing the system power while running Register Transfer Level (RTL) code by instrumenting their virtual prototype using real system software to identify hot spots in the design. Finally, they're differentiating by getting software up and running early on a virtual prototype and using that to achieve a time-to-market advantage.

RHINES

The single biggest challenge is the rising cost of system development – whether it is for a complex SoC module or the electronics for a new car. Hardware and software subsystem complexity has not only increased recently, but also the interactions and dependencies between hardware and software subsystems create bottlenecks in the release process. Ideally, the embedded software development for an SoC or system would be completed ahead of the actual hardware. But that requires a virtual representation of the hardware that can run sufficiently fast and be robust enough for developers to thoroughly verify their code.

Although software development and hardware development are deeply intertwined, they are intensely unique disciplines and development tools must focus on a specific domain while providing insight into the other. Forcing software engineers to adopt hardware tools is a nonstarter as it not only forces an unfamiliar and inefficient experience on developers but also doesn't provide insight into software concerns at the firmware, operating system, and application software level. Emerging to address this is a cohesive embedded software development environment that is the same, whether the target is a simulation, emulation, prototype, or final product. This environment embeds the most advanced pre-prototype technology available from the hardware design tool flow deeply into the native software environment, resulting in a significant time-to-market advantage for software development. Now software developer teams can remain in their core development environment and develop, debug, and optimize their complete software stack on virtual prototypes and emulation platforms, before and after first prototype.

SERUGHETTI

SoC architects face big challenges as their chips move to finer-grain silicon geometries and the costs of system architecture inefficiencies become prohibitive. A good starting point to address these challenges is to use early architecture simulation to better understand and optimize multicore hardware-software partitioning and its relationship to performance, power, and cost of the SoC. This is more than optimizing the number of processors; it is optimizing how the software is mapped or assigned to each core. Another technique used during early architecture simulation and exploration is sensitivity analysis. This enables the architect to look at many factors to understand how performance and power metrics are sensitive to changes in various system parameters.

For virtual prototyping to enable these techniques, a complete solution needs to be available to users and must support the creation of virtual prototypes and their use by software developers. This creation requires model libraries for different IP used in an SoC (examples: ARM processor models, DesignWare IP models), tools to create customer-specific IP models, and tools to assemble and debug the virtual prototype model. Once created, the virtual prototype must be provided as a VDK to the end user who will use it as an embedded target. It must include integration with software debuggers, access to advance scripting, debug and analysis capabilities unique to a virtual prototype, and depending on the application, interface to the environment in which the SoC will reside (for example, the ability to connect to a physical Internet network, ability to cosimulate with analog or mechanical simulation tools).

CLINE

The "system" continues to become more complex. Integrating IP from multiple vendors, custom accelerators designed in-house, embedded software, and more is an extremely difficult task. Creating an implementation model that ties closely to embedded software and provides significant time-to-market advantages allows design teams to get more done with fewer resources.

SIWIŃSKI

The move to 16 nm/14 nm FinFETs demands major changes in tools such as parasitic (resistance/capacitance) extraction, as well as new foundry design rules. Cadence is making a major R&D investment to adapt its IC physical implementation tools to FinFET technology, and has successfully participated in several FinFET test chip efforts. Work is ongoing in both custom/analog design and digital design to ease FinFET implementation.

Further, today's large SoCs are not designed from scratch; they use semiconductor IP sourced from other design groups or outside providers. Accordingly, Cadence provides silicon-proven design IP and verification IP.

The growing role of software, with software teams now becoming larger than hardware teams in both system and semiconductor companies, is also driving an evolution of needs. In particular, the ability to continue to provide offerings to feed the "shift-left" notion of early software development, hardware-software integration, and system validation continues to be a big and growing focus.

Pulling it all together into overall end-product creation is another big challenge, requiring broad front-to-back solutions and ecosystem collaboration to address the needs of both semiconductor and system companies across the plethora of vertical markets with their distinct needs.

QUESTION 5

Which technologies or standards are needed in EDA now but not yet available?

NEIFERT

A real standard for Transaction Level Modeling (TLM) interchange and configuration is needed to enable higher-level system exploration and development. TLM-2.0 lays the framework for some of this, but there is no guarantee that the AMBA TLM IP model created by one vendor will work with the AMBA TLM IP model created by another vendor. Additionally, it's not clear how well they'll integrate with the virtual prototype tool tying it all together. Finally, once all models are talking to each other at a functional level, they still have different interfaces for control, configuration, and inspection.

CLINE

It's difficult to say. Design flows continue to evolve and project teams want access to the best-in-class technologies. To implement these sophisticated methodologies, standards must continue to evolve with the flows. IEEE 1666-2011 is a great example of a standard that has continued to evolve with the ESL design flow. The SystemC organizations continue to enhance the standard to move quickly with the ESL design flows being used today. (The latest IEEE 1666 information can be found at http://standards.ieee.org/findstds/standard/1666-2011.html.)

SERUGHETTI

Fortunately for designers in the embedded space, there are enough standards to support existing use models such as virtual prototyping; a number of system IP blocks and verification IP blocks are interacting with a combination of simulator, emulator, and virtual prototypes through SystemC (IEEE 1666) and TLM standards. The industry is working on IP-XACT (IEEE 1685), UPF (IEEE 1801), IP Encryption (IEEE 1785), and SystemRDL (Accellera) to name a few. Most of these standards are in their second or third incarnation, going through evolutionary enhancements as recommended by real users who are identifying the shortcomings.

SIWIŃSKI

With the evolving demands placed on system and semi companies, there is a need to enable faster and better end-product creation. This means that EDA will continue to evolve beyond its classical IC automation roots toward a greater focus on system needs. In turn, that will result in more collaboration across software, electrical, and mechanical domains, which will also fuel growth in standards to enable development of complex systems across many domains.

RHINES

One area ripe for standards is the automotive electronics market. More than 35-40 percent of the cost of a premium-class car today is software and electronics. Many of these embedded systems support safety-critical features such as automatic braking systems or airbag deployment. There is also significant development around in-vehicle infotainment systems that entertain and inform passengers.

In the past few years, the emergence of the AUTOSAR and GENIVI standards has provided a basic infrastructure for developing vehicle software within the automotive market. What is needed is a complete software design environment that unifies standards-based development between the real-time deterministic and safety-oriented systems with a richer infotainment and telematics system. With such an environment, embedded designers could support multiple operating systems on a single processor, or migrate applications from single core to multiple cores, and make use of hybrid or heterogeneous architecture SoCs. This will reduce overall development and Bill-of-Material (BOM) costs, while enabling development of customized, feature-rich automotive applications.

QUESTION 6

How can future EDA innovation best be fostered?

CLINE

Innovation in EDA comes from small companies. Large systems and semiconductor companies are best served by making sure these small, innovative companies thrive and continue to challenge conventional wisdom.

SIWIŃSKI

Traditionally a lot of EDA innovation has come from startups. In recent years there have been fewer startups, due to a lack of venture capital funding and a more difficult environment for IPOs. This, in turn, has rekindled in-house incubator and innovation investments in the EDA companies to make sure that daunting customer challenges can continue to be addressed.

The starting point, as always, has to be the market and the customers, posing new challenges. The good news is that we live at a time when electronics are evolving at an unprecedented rate, providing a very good driver for EDA innovation indeed.

RHINES

It is important to examine historic trends when predicting the future. Over the past decade, nearly all growth in EDA came from new design methodologies in response to emerging design challenges. The largest growth was in Design For Manufacturing (DFM) with 28 percent compounded annual growth, followed by formal verification, ESL, and IC/ASIC analysis (principally power estimation). In contrast, the rest of traditional EDA – including custom IC, ASIC, and PCB design – grew only 1 percent compounded over the entire decade.

With history as our guide, there is no doubt that the EDA industry should focus on solving new problems and finding new users for EDA technology, while continuing to invest heavily in maintaining and upgrading traditional design methodologies.

SERUGHETTI

Gaining and sharing experience constitute the best approach to foster innovation while ensuring that there is a viable industry on both the user and vendor side. Collaborative and open engagements are key, which are fostered by establishing one-to-one relationships between users and vendors as well as by participating in industry events, user forums, and technical communities such as OSCI and the Embedded Vision Alliance to promote an exchange of ideas and solutions.

NEIFERT

The best innovations are driven by close partnerships between EDA companies and their industry partners. This way, real-world designs are used to drive the innovation that can be rolled out to the broader market. Carbon has been using this type of partnership with Samsung to drive virtual prototyping innovations.

Choose the right embedded operating system

By Warren Webb

A multitasking operating system has become essential to untangle the complex layers of software needed to drive the latest generation of interactive and networked embedded products. But before you pick the operating system software to be at the heart of your next design, take a look at the basic requirements and how they might affect overall product performance and the software development process.

As embedded devices continue to increase in complexity, the software development task has become the largest element of the typical project budget. Graphical interfaces, network protocols, and data security are just a few of the new requirements that design teams can find added on top of their custom application software. With this growing software burden along with customer demand for faster response times and instant data access, operating systems have become an essential element to organize and prioritize the software and hardware interaction routines. Unlike the desktop environment where only a few operating systems prevail, embedded designers have hundreds of options and the right choice depends on the special needs and requirements of each project.

The basic functions of an operating system are to manage the system's peripherals and schedule software tasks to ensure that each program gets some processor time. A file system is also part of a standard OS to store software modules and boot instructions. Another big benefit of an OS is to provide networking software and drivers for common hardware peripherals, eliminating constant reinvention. However, an embedded OS is quite different from its desktop counterpart. Desktop systems assume a keyboard, a mouse, a display, a hard disk, and plenty of memory. However, there is no such standardization in embedded products. One embedded system might have no hard disk and limited memory while another has no user I/O at all. An embedded OS must also be modular, allowing components to be added or removed to adjust the memory footprint such as is possible with the Neutrino real-time OS from QNX (see Figure 1). Before settling on an OS, designers should understand scheduling algorithms, memory requirements, latencies, tool support, and pricing models.

Hard or soft

As designers piece together the requirements for each new project, real-time performance is one of the early considerations. OS vendors use terms like "hard" and "soft" to describe the real-time operation of their systems. Hard real-time systems are scheduled so tasks are guaranteed to start within a precise length of time from an external event. Hard real-time systems are deterministic. Soft real-time systems generally list the average length of time to start the routine, but there is a small probability that the maximum time can be much longer. Mission-critical applications must be deterministic. For example, industrial actuators, airbag controllers, and even arcade games must react in a known time. Soft real-time applications usually respond within a few seconds but an occasional slow response is not critical.

The basic architecture of a real-time operating system includes a program

interface, the kernel, device drivers, and optional service modules. The kernel is the core of the operating system and provides an interrupt handler, task scheduler, resource sharing flags, and memory management. The kernel's services are requested by calls to its Application Program Interface (API).

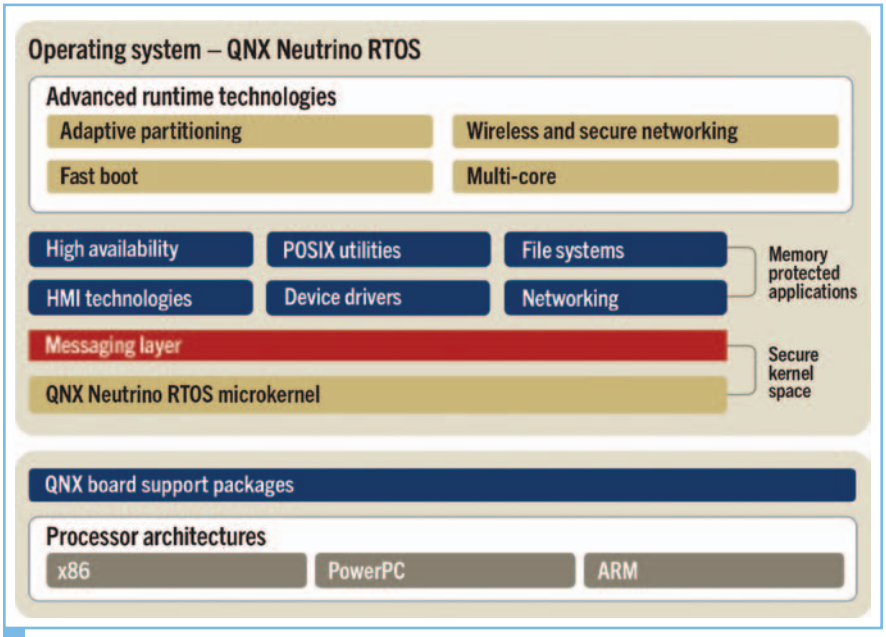

Figure 1 | The QNX Neutrino RTOS features a modular framework that enables designers to create optimized operating software for a variety of processor architectures.

The kernel is active continuously during real-time operation and must remain memory resident. One of the kernel's primary functions is to process interrupts caused by external or internal events. When an interrupt occurs, the processor transfers control to an interrupt service routine that logs in the interrupt, sends a message to the scheduler, and returns to the active code. The scheduler sets up the order of execution of individual tasks to ensure that a higher-priority task can preempt a lower-priority task to maintain a deterministic response. The most popular scheduling technique is preemptive, prioritized scheduling where tasks can interrupt a lower-priority task and continue execution until finished or until preempted by a higher-priority task.

The development tool chain is another big issue in the selection of an OS. Developers will spend most of their software design and debug efforts interacting with the Integrated Development Environment (IDE) to gain quick access to the editor, compiler, linker, downloader, and runtime tools. Most vendors provide a full IDE including the source code editor, the code manager, links to the compiler and linker, software to download code to the target platform, and one or more debuggers. Software vendors also supply software performance analysis tools to help developers profile and visualize the real-time activity in software. Many of these analysis tools are optional and will add to the overall tool cost. Green Hills Software offers the MULTI Integrated Development Environment (IDE) to support a variety of OSs. The IDE software tools include several C compiler options, editor, configuration manager, code browser, and debugger in a single package. MULTI also features DoubleCheck, an integrated static analyzer that isolates bugs caused by complex interactions between code segments that might not be in the same source file.

Cost savings

Cost is obviously a major deciding factor when selecting an embedded operating system. Since there is essentially no cost of goods with software, vendors can adjust their pricing model to generate revenue from various sources. Vendors can charge for an initial license fee, development seats per engineer, development tools, per-unit royalties, continuing support, or major upgrades. To fairly compare prices, the overall cost of ownership must be computed for each OS being considered. The calculation must include the number of units planned for delivery over the life of the product. Although a commercial OS can be expensive, cost savings is an important reason to purchase an off-the-shelf product. If you can purchase and therefore eliminate the coding, debug, and documentation of the most complicated portion of the software, you should give it careful consideration. Vendors promote product technical support as a major benefit of a commercial OS. They are able to provide continuous support for the operating system portion of the

software by spreading the cost over all customers.

After a life-cycle analysis and the "sticker shock" associated with the cost of a commercial OS, designers usually also consider "free" open source operating systems. Although open source software eliminates some initial cost problems, there are still significant issues that must be addressed. One problem voiced by potential open source users is the lack of a central resource to provide support similar to a commercial software vendor. Developers can often find answers to their questions through the Internet, but no one is on the hook to research and respond to a specific question. Open source products are generalized in order to fit the widest array of users and can force designers to modify the hardware configuration, resulting in higher recurring cost for the embedded device. Also, open source software is not in the public domain and users must adhere to specific rules set forth in individual licenses that may force designers to reveal the source code to proprietary software. Even with these hurdles, open source operating systems are widely used in embedded design.

Small footprint

An added consideration when selecting an OS is the trade-off between the initial hardware footprint required and the ability to add features when updates become necessary. The OS must be scalable so that users can select just those portions or features of the software system that they need. With constantly changing requirements and embedded technology, designers want to be able to download new software modules to add or revise features for future updates. Offering to reduce the bill of materials cost for embedded devices, the Nucleus real-time OS from Mentor Graphics allows developers to deploy a kernel as small as 2k on a wide selection of processor hardware. With the modular structure, designers can reduce the initial cost of the OS by selecting a subset of the available features and also the amount of ROM and RAM to support the system.

One of the favorite shortcuts of embedded designers is to start a project with a Commercial Off-the-Shelf (COTS) module to eliminate much of the hardware development and testing. This same shortcut can be extended to the software by selecting an operating system with a board support package for the target hardware. For example, LynuxWorks offers board support packages for more than 150 COTS processor boards for their LynxOS real-time operating system. These packages include ARM, PowerPC, x86, MicroBlaze, and MIPS processor boards. This approach can shorten the overall development schedule, allowing the design team to concentrate on the proprietary portion of the embedded project.

As customer demand for universal connectivity, faster response times, complex functionality, and instant data access spreads into more and more devices, designers can expect real-time operating system software as the norm in future embedded products. Fortunately, real-time vendors and open source collections offer plenty of off-the-shelf, ready-to-run software packages to fit most embedded configurations. The challenge will be to analyze your requirements to match an operating system to the application with ample room for growth while minimizing the cost and development effort. ECD

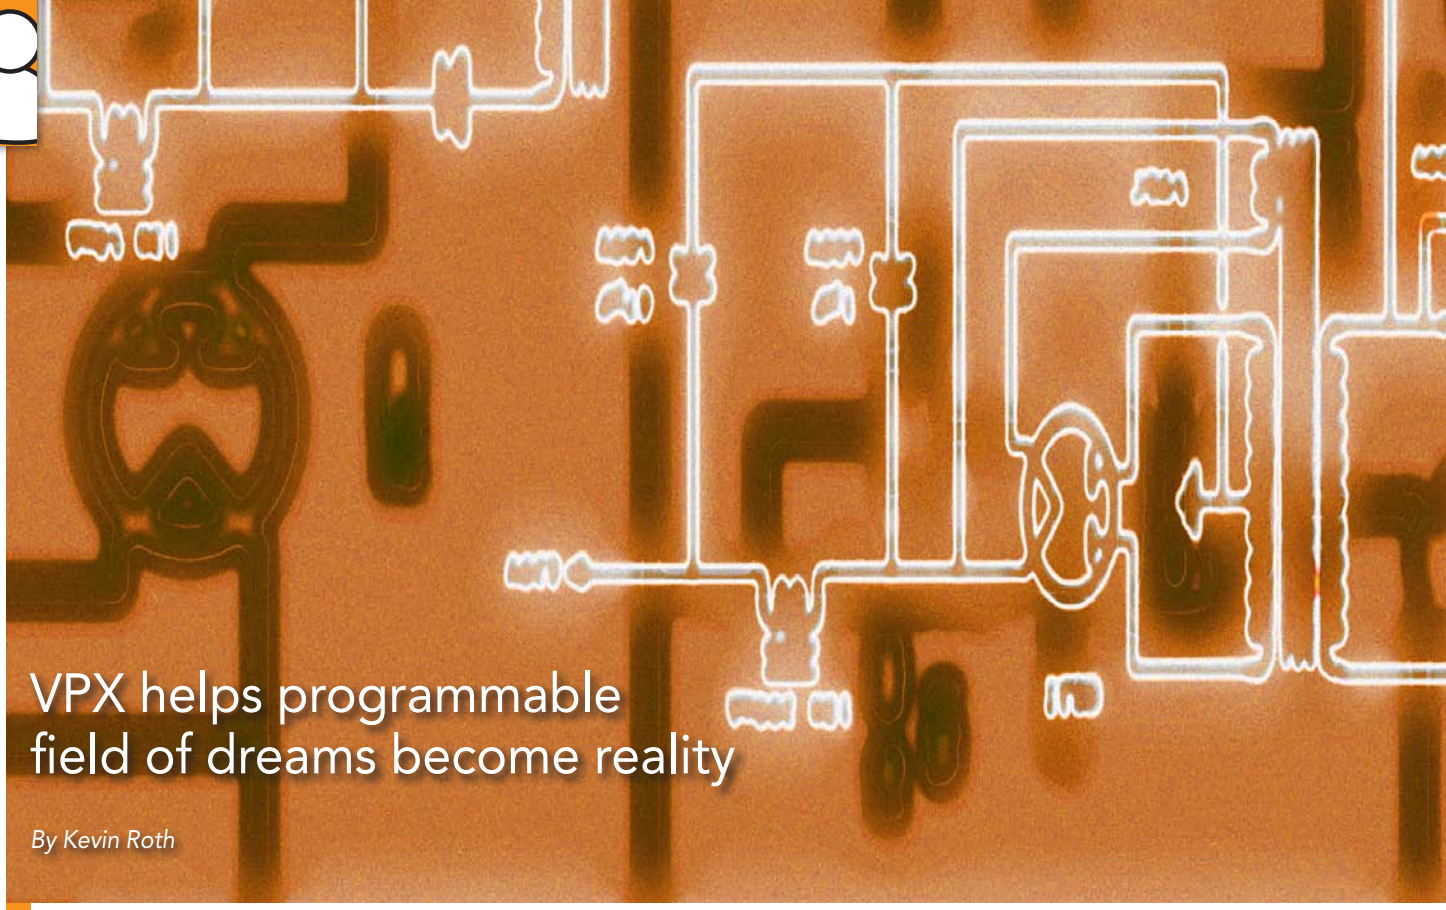

VPX helps programmable field of dreams become reality

By Kevin Roth

Thanks to the mind-boggling processing power available in the newest FPGAs and the organizational structure of 3U OpenVPX systems, a new generation of miniature ruggedized hardware designs is possible. Commercially available hardware combining these technologies allows engineers to quickly and easily build compact systems for the most demanding tasks.

Thirty-six multigigabit transceivers, 850 high-speed General-Purpose I/O (GPIO), and 2 million reprogrammable logic cells seem like a digital designer's dream come true. So why does architecting systems with the industry's newest FPGAs sometimes feel like a nightmare? Somewhere between trying to line up 10 Gbps transceiver lanes, Complementary Metal Oxide Semiconductor (CMOS) control signals, and Low-Voltage Differential Signaling (LVDS) sensor interfaces, designers may swear they've seen Freddy Krueger sneaking around.

Fortunately, finding a compact, open architecture that doesn't decimate an FPGA's outstanding I/O capability is easy to accomplish using 3U VPX cards, which offer an ultrarugged form factor loaded with I/O that can operate at 10 Gbps. Most importantly, the VPX universe utilizes a standardized control and data flow architecture supported by vendors around the world, letting designers utilize every square inch of the latest, most massive FPGAs on the market.

OpenVPX opens design possibilities

Having an interconnect scheme with a high pin count and outstanding signal integrity sounds powerful, but it isn't useful without an organized and widely accepted architecture behind it. The OpenVPX architecture has moved VPX from a highly customized embedded system into the realm of COTS rugged computing. The most common OpenVPX architecture is similar to modern PCI Express (PCIe) desktops, with a central processing/control card operating a PCIe switch (see Figure 1). A standard backplane carries the PCIe end points to an arbitrary number of peripheral cards, just like the motherboard in a desktop computer. The biggest difference is that VPX can withstand vibration, shock, and temperature extremes that would send a desktop PC to an early grave.

Beyond benefits provided by a fast data backplane, OpenVPX offers options for control over Ethernet and ultrafast SATA solid-state drives. Processor cards work best as system controllers, thanks to their built-in PCIe arbitration functionality. However, FPGAs have a distinct advantage for peripheral cards. They give OEMs the ability to use almost every standard configuration option available with a single piece of hardware, which translates directly into design flexibility for system integrators.

Massive I/O brings reconfigurable options

With organized data and control lines, a few rays of light are starting to filter through the initial design nightmare. Add the flexibility and speed of an FPGA, and the scare factor is reduced even further. However, a system can't be composed of just a few SATA devices, PCIe lanes, and an Ethernet connection alone.

Enter the knight in shining armor, reconfigurable I/O. This technology has been one of the biggest advantages in FPGAs since their conception, and it is making strides again with VPX. A single 3U VPX card slot has three interface connectors on the backplane: P0 is for power

and system management; P1 is for the OpenVPX architecture; and P2 is entirely reserved for user-defined I/O. That's a whopping 72 pins with a max speed above 10 Gbps.

VPX also defines a handful of standard I/O arrangements to help designers correctly line up multigigabit lanes and lower-speed GPIO. The FPGA goes a step further, allowing the multigigabit lanes to operate over a wide range of speeds and signaling standards such as XAUI, 1000BASE-X, Serial RapidIO, and many more. On the low-speed side, pin flexibility ranges from 3.3 V CMOS down to LVDS in a single COTS hardware design.

The VPX P2 connector offers several options, but designers can find even more configurability on the module's front-panel side. Due to wide industry adoption, modular front-panel I/O is common and supported by multiple vendors. FPGA Mezzanine Cards (FMCs) offer the widest range of industry support and deliver 160 GPIO and 10 multigigabit SERDES links that can run at 10 Gbps each. The market is filling up with front I/O modules that enable system architects to perform practically any task desired, including high-performance analog capture, video processing with Camera Link, copper and optical SERDES, 10BASE-T, RS-232/422 control, and simple GPIO. Where there are gaps in industry-wide offerings, various OEMs are delivering custom, low-cost products with short lead times (see Figure 2).

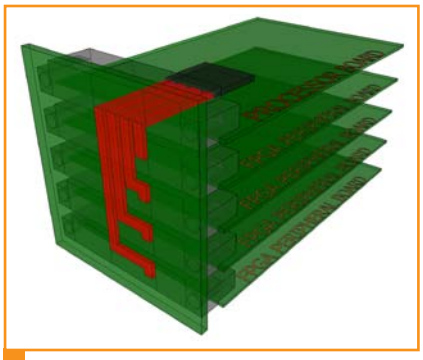

Figure 1 | In a five-slot OpenVPX system with a PCIe architecture, a standard backplane carries PCIe end points to the peripheral cards.

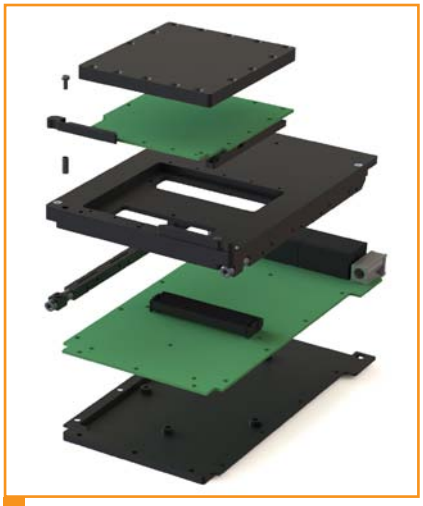

Figure 2 | Products such as the OpenCOTS 3U Conduction-Cooled VPX Development Kit offer designers support for developing 3U VPX modules with front I/O and an FMC site. Image courtesy of WaveTherm (www.wavetherm.com).

Ready for rugged deployment

To leverage these unlimited connectivity options, massive FPGA processing power, and an architecture that is practically plug-and-play, designers need the ability to package it all up and put it into an environment where only cockroaches could survive. VPX is designed to handle every environment imaginable. Simple forced-air lab and commercial deployment chassis are low cost and readily available. This is particularly useful when prototyping designs using rear I/O modules before integrating up to fully rugged deployment versions.

To help transition from the lab to deployment in a scorching dust-filled desert or at 40,000 feet, OEMs offer conduction-cooled versions of their hardware to get the heat out without airflow. The 3U form factor also helps reduce the length of metalwork between the ambient environment and the blazing-hot components in the middle of the board. For applications trying to defy the laws of thermodynamics, liquid flow-through chassis can actively suck the heat out of boards with incredible efficiency and without any degradation at altitude or fear of dust contamination.

VPX and FPGA technology in action

Add all of these elements together, and system integrators will find they can create compact, cost-effective, ultra-rugged systems for a plethora of challenging applications. Understanding the power of VPX takes a run through the ropes. The hardware can fulfill any designer's dream by providing the

following capabilities, as illustrated in the application in Figure 3:

- **Output:** A live video feed overlaid with directional vector graphics pointing out targeted radar signal sources
- **Inputs:** Full Camera Link feed of high-altitude video, GPS, and RF input tuned to suspect radio frequencies

This is not a simple task, but the hardware to accomplish this is already built and ready to order. The user interface could be a workstation or flight control unit with a 10 GbE link to view the live video stream.

A processor card utilizing a Zynq processor with integrated programmable logic could easily manage the system control. Mount this on a PCIe switch card, and the processor XMC gains four lanes of PCIe Gen2 access to each VPX slot. The second slot has Camera Link front I/O passing 5.4 Gbps of raw camera data to an FPGA for preprocessing. In the third slot, a VPX carrier with an FPGA XMC card and a dual ADC front I/O card takes in the DF radar signals and processes them into directional information regarding the radar source. Lastly, a GPS receiver decodes the information from a GPS antenna to deliver a high-resolution high-refresh rate for location, altitude, and velocity data to help align the RF and image data. OpenVPX even has a standardized 1 PPS distribution system built into the backplane to keep all boards in sync with current GPS information.

This is all completed using Xilinx's latest Virtex-7 FPGAs, which offer greatly reduced power consumption compared to previous generations. By using reconfigurable logic at every stage in the design, the entire internal communication architecture can be configured to achieve the highest possible performance. For example, data can be passed between slot 2 and slot 3 with a 10G or 40G SERDES link over the OpenVPX expansion plane using the four-lane PCIe link or user I/O on P2. Virtex-7 FPGAs also come with the Vivado Design Suite, which was written from the ground up to allow teams of engineers to work collaboratively on complex projects like this.

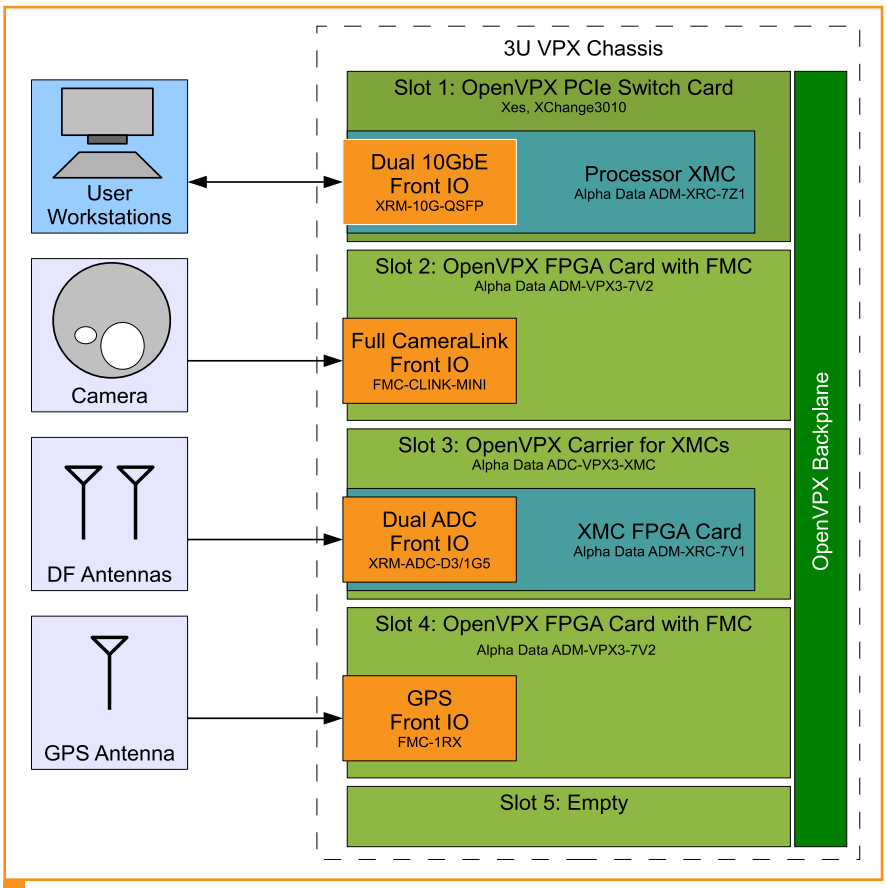

Figure 3 | OpenVPX hardware equipped with the latest FPGAs can provide the I/O functionality needed in a DF-Camera overlay system.

Another impressive feature of this hardware is a highly upgradable and expandable framework. Say a designer wants to add another feature that requires a whole new input device and some heavy processing power. This system supports an additional PCIe link to slot 5 allowing future expandability without a system redesign. The hardware itself is upgradable, as Alpha Data and other OEMs continue to release boards with the newest FPGAs year after year.

Take the design another step further and add a few terabytes of solid-state storage for video recall. Use the mSATA slots located under the XMCs, or populate the empty VPX slot with one of many 3U VPX storage drives available from various manufacturers. All of this hardware is currently available on the market today. No need to wait for hardware design cycles and board reruns; just submit a purchase order and find some good programmers.

Sleeping easy

It's time to say goodbye to those design sessions with your favorite horror movie character because the nightmare is finally over. The combination of 3U OpenVPX systems and the unreal processing power of modern FPGAs is creating a new generation of miniature, ruggedized hardware designs. Take advantage of these technologies and see how good it feels to start sleeping easy. ECD

Kevin Roth is an Electrical Engineer at Alpha Data.

Alpha Data
kevin.roth@alpha-data.com
www.alpha-data.com

Follow:

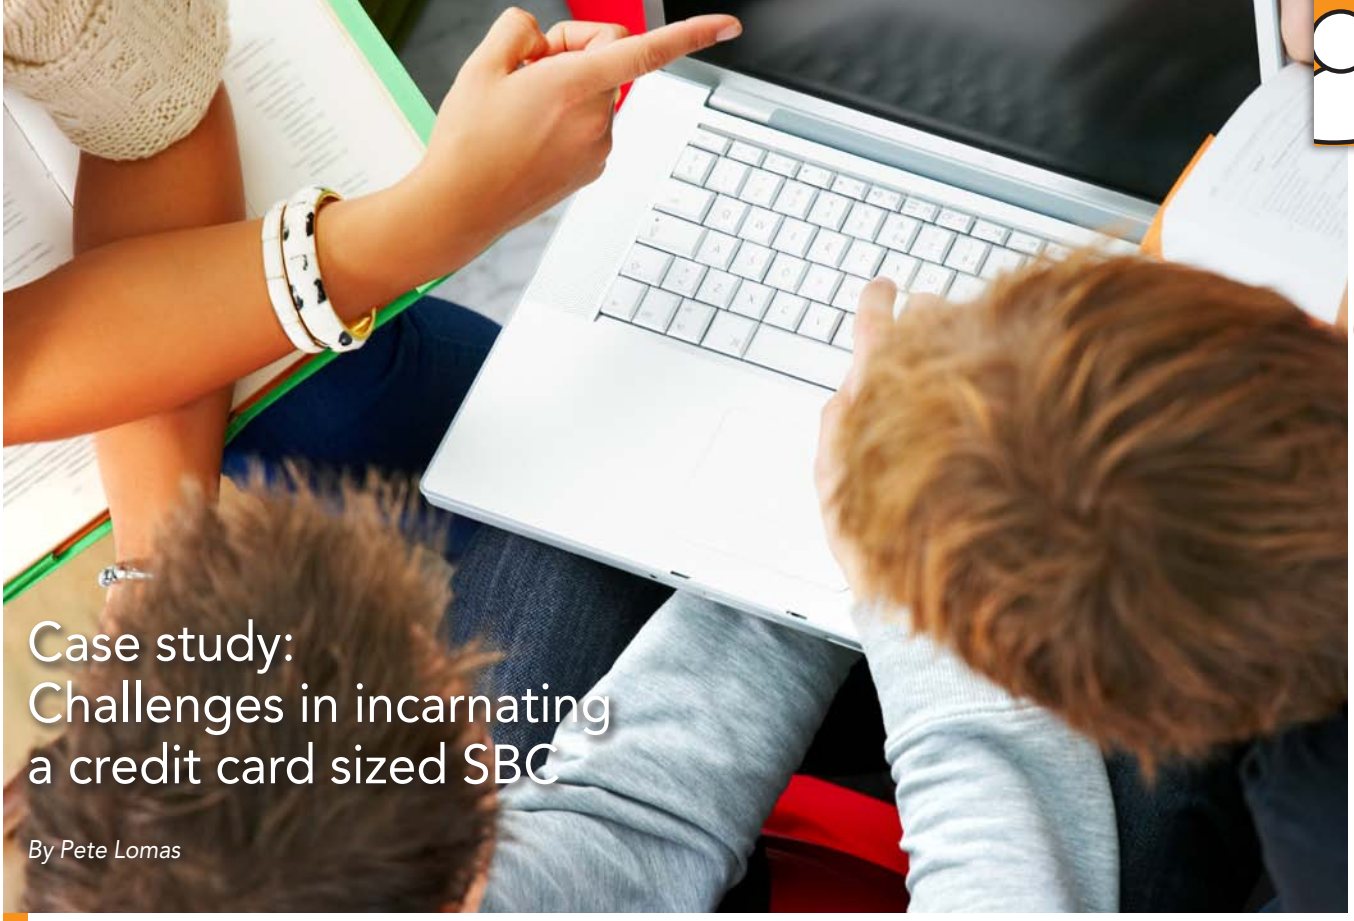

Case study: Challenges in incarnating a credit card sized SBC

By Pete Lomas

The initial goal in creating the Raspberry Pi credit card sized, Linux-based Single Board Computer (SBC) – targeted primarily at education – was to develop a response to the decline of students engaging with computer science and related engineering disciplines. Our desire was to reverse the trend of children becoming consumers rather than creators. The following case study follows the hardware development process from an early failure, initial prototypes, and through to the finished production design.

Over recent years there has been an increasing trend for children to be consumers of digital content rather than be future creators or engineers. This trend is driven by manufacturers looking to provide a seamless experience for target customers on a variety of electronic platforms, from gaming consoles to tablets and laptop computers.

As a result, access to raw I/O has become restricted. Similarly, any packaged provision of a programming environment is an anathema to the products' commercial goals. The knowledge required to create "hello world" or flash an external LED has become simply too vast and the opportunity to learn vital skills such as structuring/codifying ideas and debugging has been largely subsumed by a click-and-shoot world. Any motivation to get under the hood and see how these products work is largely dissipated by the impenetrable barriers presented by these "locked down" systems.

The challenge in developing the Raspberry Pi credit card sized, Linux-based SBC was to break down these barriers and provide access at a sufficiently low cost so any fear of breaking the hardware was effectively removed. Having the hardware is only half the story; the provision of a rich set of programming environments such as Scratch and Python with libraries to allow control of peripheral hardware provides an engaging toolset for learning through experimentation and play in either the formal classroom or at the many school and independent maker (hackspace) clubs. The following case study shows how Raspberry Pi was developed from the ground up.

Initial designs

The initial concept was a small form factor single board computer with good graphics performance that would boot directly into a Python IDE. There was never any doubt that an SoC solution was required; however, we needed to compromise to align target requirements with the feature set of any chosen SoC. Providing support for peripherals external to the SoC would typically add cost and complexity.

In late 2008 we pursued a design using a BCM2727 multimedia processor. It was apparent that the (23 x 18 array) 0.4 mm ball pitch of the BCM's BGA package would challenge the fabrication of a low-cost PCB (PWB). Coupled with the cost of space requirements of the supporting silicon we missed our target cost point by approximately 50 percent for the 10,000 unit volumes initially targeted.

In 2010, Broadcom developed the BCM2835 Media Applications Processor SoC with a substantial GPU capability and relatively modest, in terms of performance, ARM1176JZ-F application processor. Notably, several members of the design team were active supporters

of the Raspberry Pi foundation and had intimate knowledge of the BCM2835 SoC, significantly reducing the implementation risk involved with using the part, even when compared with one in commercial distribution. From a software viewpoint, our initial target had been a single language implementation of Python. Linux was now a realistic possibility providing access to wide range of preexisting languages and tools, which would have huge benefits in broadening the scope of applications for the SBC.

Development started in earnest in summer 2011. A team at Broadcom led by Gert van Loo produced an Alpha evaluation platform specifically targeted for the Raspberry Pi project with all the features that the final design might require. This was used to implement Linux while in parallel the foundation team started work on an optimized (pared back) design that became the final hardware launched in January 2012.

Memory

A key design goal throughout development of the Raspberry Pi was project transportability. Other systems such as PCs have the OS and core programming environment on a local hard disk. Often subtle setup and configuration changes on one system will frustrate the running of an application developed on another. In the Raspberry Pi SBC (Figure 1), all application, OS, and boot information is carried on the SD card. The absence of any user or OS NV storage onboard ensures a clean working environment. The Raspberry Pi can be completely repurposed, with a different OS or application just by the use of a different SD card. SD cards can be sized (4-64 GB) to suit the application and additional storage can be achieved through a USB stick or USB HDD. This flexibility provides a minimal entry cost with a massive range of memory expansion options.

The SoC utilizes a Package on Package (PoP) design for primary RAM. The dynamics of the memory market enabled an upgrade from 256 MB to 512 MB only five months following the launch. An important technical benefit of the PoP is the removal of the SDRAM connectivity from the SoC/PCB interface, reducing ball count, overall footprint, and routing and timing complexity for only a small increase in assembly complexity.

Physically the processor has a 0.65 mm pitch almost full 18 x 18 array (299 balls) in a 12 mm square package; the SDRAM interface is two peripheral rows 23 x 23 (168 pads) at 0.5 mm pitch.

Display

One design dilemma was VGA. We had integrated HDMI and composite video for default legacy support, but the lack of VGA was troubling as many surplus school and college monitors were typically VGA. In the home, however, the situation is reversed and HDMI is more common. The availability of external third-party HDMI-to-VGA converters would have to meet this need.

Thus, we implemented a Dual Videocore IV GPU, with full 1080p30 Full HD H.264 video encode/decode capabilities to allow for the creation of a low-cost media center – a cornerstone application for the board as it provides an important route for interactive teaching material deployment.

Internet

Internet connectivity is a critical feature that the SoC did not provide directly. This was resolved with the addition of an integrated USB 2.0-based triple hub and 10/100 Ethernet end point. The single SoC USB host supports Ethernet traffic and basic USB peripherals. Wi-Fi would have been more desirable but we would need 10/100 as a backstop. As USB Wi-Fi modules are freely available, the decision was 10/100.

Expandability

The strategy we adopted in all elements of the design was to provide as many opportunities for incremental expansion as possible. The idea was this expandability would track the increasing knowledge and enthusiasm of the user and provide avenues for further experimentation. One important area was the interfacing of peripheral electronics. The GPIO port was therefore designed to support key low-level peripheral interfaces including I^2C, SPI, RS-232, and PWMs. The provision of these interfaces has yielded a wide variety of third-party modules for digital, ADC/DAC, motor control, RTC, PoE, and prototyping boards to encourage users to create their own circuits.

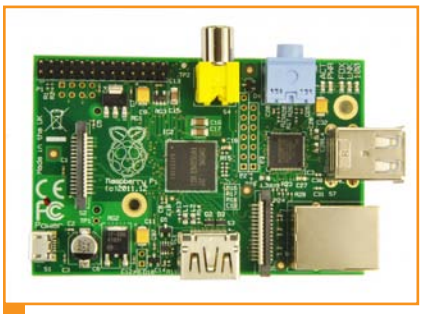

Figure 1 | In the Raspberry Pi SBC, all application, OS, and boot information is carried on the SD card.

The USB ports provide further, more standardized expansion, especially if connected to a powered hub. Originally these were intended for simple HID peripherals and limited to 100mA. Later production builds recognized users' expanded utilization of these ports and this restriction was removed.

A camera had always been on the road map as the processing of image data would open up a wide variety of learning and experimental opportunities. Thus, the Camera Serial Interface on the SoC provides video capture without overloading either the CPU or the USB as is the case with USB-based Webcams. With data rates of 1 Gbps per lane (2 lanes), it comfortably supports a 5 M Pixel sensor (2592 x 1944) at 15fps. Raw camera data is processed in the GPU to render an RGB image, and the GPU can also be used to provide compression for single frames or streamed video. Also included is a dual-lane DSI that will ultimately provide a low-cost interface to a flat panel display as an alternative to the HDMI or composite video.

Power tree

The initial alpha design and the early beta design utilized a 6-12 V input and a dual SMPS to generate +5 V and +3.3 V. The threatened European Directive for common (R&TTE) chargers created almost universal adoption of the micro USB. Evaluation of implementation

costs led us to design revisions to utilize this standard. This removed one SMPS and the use of a linear for the +3.3 V removed the second, giving a fourfold cost point improvement for this element of the design, at the expense of 700 mW additional dissipation onboard.

The result: With a full operational load exercising the GPU, ARM, Ethernet, and USBs, the power requirement of the board is 3.2 W. Approximately 1,100 mW is utilized by the SoC and memory and 760 mW by the LAN/USB hub with 850 mW being dissipated in the power tree itself. Certainly not ideal but necessary.

Form factor and PCB technology

Achieving a small form factor was an important design goal, as it reduced PCB cost and allowed multiple boards (6) to be built on the same manufacturing panel, thus saving production handling. With users providing interconnect cables, cost and availability were important, hence full-size HDMI, SD, and USB connectors and 0.1" pitch header for the GPIO were retained in the final design. Ultimately the connectors drove the form factor rather than the silicon.

The alpha board PWB design utilized full HDI technology with buried, blind, and micro via in pad. As part of the production engineering exercise, we were able to reduce this to a conventional six-layer design with one minor HDI concession. We introduced laser-drilled blind vias between the BGA placement layer and the underlying ground plane, providing space for critical decoupling on the rear directly under the processor/memory PoP stack.

We paired down the GPIO and removed unnecessary lanes from the CSI and DSI ports present on the original design to ease the routing from the SoC. The loss of these signals was not going to negatively impact our design goals but it allowed the use of just three of the six layers to carry the BGA escape avoiding compromising ground and power plane integrity.

Looking ahead

Going forward, the design team is evaluating incremental improvements in both the design and manufacturing techniques. These will translate to the addition of more features and extended performance at the same price point. The accessible nature of the Raspberry Pi SBC allows the curious to delve "under the hood" right down to the smallest detail of the systems operation. It encourages learning through exploration and experimentation. Importantly, it is also interesting, engaging, and fun. **ECD**

Pete Lomas *is a cofounder and trustee at the Raspberry Pi foundation, a UK registered charity. He was responsible for the hardware design and development of the Raspberry Pi SBC. He is also Director of Engineering at Norcott Technologies, an electronics design and contract manufacturer based in Warrington (UK). Pete holds a B.Sc. and M.Sc. in Computer Science from the University of Manchester (UK). He can be contacted at* pete@raspberrypi.org.

Raspberry Pi • www.raspberrypi.org

Follow:

Embedded COMPUTING DESIGN
TOP Embedded Innovator

Josh Lee

President and CEO, Uniquify

Josh Lee is President and CEO of Uniquify in San Jose, California. With two cofounders, he bootstrapped Uniquify into what is now a thriving IP and chip design and manufacturing services business collectively referred to as "ideas2silicon." The team has deep expertise in the specification, design, implementation, and manufacturing of complex SoCs. Lee and the team also built a comprehensive design management system, Perseus, used to manage all SoC projects at Uniquify. Lee holds a Bachelor of Science degree in Electrical Engineering and Computer Sciences from the University of California, Berkeley.

As you look ahead to the next few years, which embedded technologies and applications present the most interesting opportunities?

LEE: Adaptive or variation-aware IP is the wave of the future for embedded systems and will give IP suppliers and applications providers a clear path for growth.

As the industry moves to 28 nanometer, process variation effects and dynamic variations due to fluctuating operating conditions degrade system performance or cause system instability or render the system inoperative, which is why adaptive IP will become so critically important.

Development teams are looking for ways to manage static and dynamic variations in an SoC design that's full of third-party IP. Because adaptive IP automatically compensates for the effects of static and dynamic variations, a system can achieve its best performance, while maintaining a high degree of reliability.

Adaptive IP measures relevant parameters critical to performance and reliability, and automatically makes adjustments to ensure they are optimized. Such precise measurements and corrections are made during system initialization and at regular times during system operation.

Adaptive routines run quickly with little impact on system operation and throughput, and have enough latitude to correct for a range of variations. Because adaptive IP is in the chip, each system is optimized for static variations in each component and dynamic variations caused by the system environment. As a result, the system optimizes its operation to deliver the best performance with robustness and reliability.

What are the largest obstacles to innovation in the embedded realm, and how should those challenges be solved?

LEE: A study conducted last year found that the IP market is growing faster than the semiconductor market – twice as fast, in fact. While the opportunities are many in the IP market, the challenges might be even greater. For example, concerns about IP reliability and complexity remain, which is where adaptive or variation-aware IP can come to the rescue. DDR memory controller subsystem IP, found in almost all modern electronic products, illustrates the point. Uniquify sought a way to solve the fundamental problems of variations in the DDR memory subsystem that directly impact system yield and reliability. Our solution is circuitry embedded within the DDR PHY that precisely measures the DDR interface timing windows and automatically adjusts them. We refer to this patented circuitry as "Self-Calibrating Logic" [SCL] and "Dynamic Self-Calibrating Logic" [DSCL].

SCL and DSCL enhance device and system yield and reliability, reducing the effects of variation and maintaining DDR memory system performance as operating conditions fluctuate during system operation. The DDR memory subsystem timing calibration can be applied at system power-up (SCL) and during system operation (DSCL).

In which market segment and geographic area do you foresee the fastest growth for embedded products?

LEE: The fastest growing market segments range from networking and consumer electronics to image and media processing and digital TV. We see several geographic regions with explosive growth including Korea, Taiwan, Japan, and Israel. In the U.S., southern California and Texas are growing rapidly as well.

How does a company stay on the cusp of innovation, rather than just following the embedded crowd?

LEE: Listening carefully to the concerns and challenges of design teams helps identify various trends as well as opportunities. Building close partnerships with design teams helps develop the trust to get a firsthand perspective on what it will take for the next wave of innovation.

An expert team of technologists is a critical component to staying innovative. Also, reading publications such as *Embedded Computing Design* can give insights into new application areas, as can participating in industry events. ECD

Embedded COMPUTING DESIGN
TOP Embedded Innovator

Darren Humphrey

Sr. Vice President and Chief Technology Officer, DiSTI

Since 1994, Darren Humphrey has led research and development at DiSTI as Chief Technology Officer. He developed the flagship product, GL Studio, which has since grown to a suite of industry standard HMI tools, including products for safety-critical and embedded applications. He also advises DiSTI clients on integrating HMI technology and is the software architect for the company's commercial products. Darren holds a B.S. in Computer Science from the University of Central Florida.

As you look ahead to the next few years, which embedded technologies and applications present the most interesting opportunities?

HUMPHREY: At DiSTI, we are concerned with the software that drives the Human Machine Interface (HMI) experience. To that end, some of the most interesting opportunities we see emerging revolve around a higher end user interface experience in a broader range of consumer areas, based on the lower barriers to entry into the market.

The availability of low-cost Systems-on-Chip (SoCs) with highly capable OpenGL graphics hardware presents opportunities for exciting and superior user interfaces on a broad range of consumer devices. We've seen digital instrumentation and infotainment displays move from a luxury car feature to a standard feature of most new automobiles. More and more, consumers are demanding that the devices they interact with on a daily basis, from thermostats to home appliances, give them the same types of user experiences they are used to on their mobile phones.

> "We've only hit the tip of the iceberg for the automotive vertical market as price becomes less of a factor for digital displays."

What are the largest obstacles to innovation in the embedded realm, and how should those challenges be solved?

HUMPHREY: As advanced graphical displays flourish, the largest obstacle, from an HMI perspective, is developing good user interfaces. In many cases the quality and usability of the interface are either overlooked or undervalued, which can not only be a detriment to the product being delivered, but can also set back the advancement of the technology in the eyes of the consumer. The tools and skill sets used for HMI development are different from those normally used for embedded software development. One solution is leveraging commercially available HMI development tools such as GL Studio.

Developers also need to be concerned with handling change. Selecting a board support package is a moving target in the early stages of product development. Developers have to be cognizant of announcements during product development cycles for advancements in hardware/software and have the agility in their development tools to keep their options open. Choosing HMI tools based on industry standards (OpenGL) makes the difference. Developers easily shift from platform A to platform B because they want to take advantage of these advancements and their HMI development goes with it.

In which market segment and geographic area do you foresee the fastest growth for embedded products?

HUMPHREY: For us, as a B2B provider, we look toward automotive, medical, and space, and the geographic areas that support these industries, as markets that offer significant near-term opportunity. Specifically for user interfaces, embedded in the automotive industry is hot and we are concentrating on areas such as Detroit and Los Angeles in the U.S., Germany, France, the UK, Korea, and Japan. We've only hit the tip of the iceberg for the automotive vertical market as price becomes less of a factor for digital displays.

How does a company stay on the cusp of innovation, rather than just following the embedded crowd?

HUMPHREY: Well I think there are a lot of answers to that question. I personally read a lot to stay aware of trends, find and learn from fresh talent, and preview new product releases. But I think as a company, our biggest advantage is that we are willing to take chances and try new activities. As a small business, we are large enough to have the capital to take on the risks of trying new opportunities and the willingness to do so. On the flip side, we are small enough to be able to act quickly and move on new ideas. We are constantly evaluating new technologies to see if they are a good fit for us or maybe represent a new direction. An innovator can't be afraid to fail. ECD

Embedded COMPUTING DESIGN
TOP Embedded Innovator

Adnan Hamid

CEO of Breker Verification Systems

Adnan Hamid is CEO of Breker Verification Systems and the inventor of its core technology. He has more than 20 years of experience in functional verification automation. Prior to founding Breker in 2003, he managed AMD's System Logic Division and led its verification team. Hamid worked at Cadence Design Systems, serving a system-level verification expert. He holds nine patents in test case generation and synthesis. He received Bachelor of Science degrees in Electrical Engineering and Computer Science from Princeton University, and an MBA from the University of Texas at Austin.

As you look ahead to the next few years, which embedded technologies and applications present the most interesting opportunities?

HAMID: We focus on embedded systems within System-on-Chip (SoC) devices, and the industry is hitting a wall in its ability to assemble and verify such systems. Companies are discovering that it is no longer possible just to plug together well-verified IP blocks and produce a working SoC. Complete data flows and realistic user scenarios must be verified at the full-chip level. SoC verification products such as our TrekSoC product, for example, automatically generate self-verifying C test cases to run in the embedded processors before taping out the chip. The resultant test cases rigorously stress a design, more than traditional simulation test benches or handwritten embedded diagnostics. It is critical to get this verification done because the cost of finding and fixing a bug increases tenfold for each stage later in the development process. Worst case, the process of refabricating an SoC can run into the millions, and much more if lost market opportunity is taken into account. The ability to get an embedded SoC right the first time is both a big challenge for the industry and an opportunity for Breker.

What are the largest obstacles to innovation in the embedded realm, and how should those challenges be solved?

HAMID: The verification challenge is an obstacle to producing the size and complexity of chips demanded by the most innovative projects. But the sheer complexity of SoC designs is daunting as well. You can't innovate what you can't imagine. Some of these chips are so huge that it's hard for the development teams to understand all the different modes of operation, the interactions, side effects, and unintended behaviors. One way to manage embedded SoC complexity is to represent the design in the form of a graph-based scenario model. This looks like the sort of dataflow diagram that engineers draw all the time, but formalized enough so that TrekSoC can use them to generate the self-verifying test cases for verification. In this way, the developers can see behavior they didn't expect, and even discover parallelism they didn't know they had. Of course, this approach finds bugs, but it can also offer deep insight into designs so complex that it's hard for any one person to understand it all.

In which market segment and geographic area do you foresee the fastest growth for embedded products?

HAMID: Consumer products of all kinds are the biggest drivers of the electronics industry as a whole, and embedded products in particular. Smartphones and tablets contain incredibly complex SoC devices unimaginable just a few years ago. The electronic content of cars continues to grow dramatically. The "Internet of Things" will connect many more kinds of devices, and not necessarily with simple chips. Multiple heterogeneous processors are the norm in almost anything connected to the cellular network or wireless Ethernet. What's interesting is that this is happening everywhere, not just in places with traditionally better computer/Internet access. Cellular technology has become the choice for developing countries because it's cheaper to install towers than to wire every home and business. Our consumers' consumers are everywhere, but increasingly our consumers – the SoC developers – are everywhere too. Despite the challenges, the demands of the market and availability of engineering talent have made this a much smaller world than it used to be.

How does a company stay on the cusp of innovation, rather than just following the embedded crowd?

HAMID: Three things are essential for innovation. The first, and most obvious, is to hire the best people and set high expectations for them. Embedded developers today have to be really smart, moving seamlessly between hardware and software and from IP block to full SoC. Keeping them challenged and motivated comes naturally when pushing the boundaries of technology and products. The second is to keep up on academic and industrial research, leveraging it whenever possible. Today's exotic proposal may well be tomorrow's mainstream solution. Finally, find the best, smartest development teams and work closely with them. We work side by side with people developing the biggest, baddest chips in the world. That's a lot of fun, but it means that we are constantly being pushed and challenged to evolve our products at a faster pace. We have no choice: We must innovate or die. ECD

Weili Dai

Cofounder, Vice President, and General Manager of Communications and Consumer Business, Marvell Technology Group Ltd.

Weili Dai is one of the most successful women entrepreneurs in the world today. Widely considered a technology visionary, she is the only woman cofounder of a global semiconductor company, and since it began in 1995, she has helped Marvell rise to become one of the top semiconductor companies in the world. For her contributions to technology, Newsweek *named her among the "150 Women Who Shake the World,"* Forbes *named her one of the "World's 100 Most Powerful Women," and UC Berkeley selected her as the first woman commencement speaker at the College of Engineering graduation ceremony in May 2012.*

What are the biggest challenges you face on the job every day as a woman in high tech?

DAI: The world has changed dramatically and women have moved beyond the traditional roles of wives and mothers to also become leaders in every industry. Yet there are still many challenges. For example, women still play the traditional roles of wife and mother in today's society, and balancing those demands with work is a challenge. For me, I personally take pride in my responsibilities as a wife and mother raising two sons who both are Ph.D. students at UC Berkeley now while also working very hard to lead our global company over the past 18 years. Because of my passion, my determination, and my dedication for my family and my company, I make it work – and a "48-hour day" is good for me.

How do you overcome those challenges? What or who is your inspiration?

DAI: I was fortunate to have a great upbringing by my lovely parents and an education that provided me with a strong foundation. I believe the key elements of one's success start with passion, integrity, intelligence, and hard work. My philosophy in life has always been about achieving a "win-win" outcome. And that is the approach I bring to my customers and partners and employees. My parents are my source of inspiration. Growing up, my mom was the "glue" for our family, taking care of my dad, my brothers, and, of course, me. I believe that being a caretaker is a part of a woman's natural talent and should be embraced with pride. I am also thankful for the excellent education I received from my teachers and professors. In my everyday life and practice, I believe that "passion and integrity" are the foundation for all success; "fair and care" are essential principles of all successful leaders.

How can more women be prepared to enter traditionally male fields such as engineering?

DAI: I have always believed in the power of the female mind and the unique perspective we bring to problem solving in business and engineering. I believe if we can develop environments that engage more women and leverage their natural talents, we'd be able to practice business in a more beautiful way. For example, technology is embedded in the everyday lives of consumers globally. It is ideal for utilizing women's natural talent for designing the look and feel. I believe it is a passion of women to learn and turn "nerdy" technology into fashionable and user-friendly Smart solutions for the new era of the "always-on, always-connected lifestyle." By embracing this passion, a woman can create her own ecosystem to contribute and add great value to the high-tech industry.

How do you recognize when a new technology or application is one your company should invest/innovate in, versus a technology that will experience fast burnout?

DAI: Since I cofounded Marvell in 1995, our strategy has been to focus on investing and building a strong technology foundation throughout our invention and innovation for long-term impact that advances the high-tech industry and makes the world a better place for all. Today, Marvell has delivered numerous breakthrough semiconductor solutions to enable our global Tier 1 OEM customers' and partners' success. Marvell is a top five fabless semiconductor solution provider in the world, with broad technology diversity for storage, networking, computing, and mobility. In order to empower the always-on, always-connected Smart consumer, we have built depth and breadth of end-to-end semiconductor design and software capabilities.

In the next 5 to 10 years, which technologies will present the most viable development opportunities for your organization and for the embedded industry?

DAI: Today, we live in a new era where consumers are attached to mobile Smart devices, live content is shared seamlessly through the cloud globally, and "smart furnishings" transform the traditional tabletops and wall surfaces into interactive displays and live content command centers in our home and workplace. I believe, moving forward, more and more semiconductor solution development will be tailored to support and to be embedded toward "software-defined" solutions, which I refer to the term as "make to fit." In addition, I believe the aspects of developing solutions for ease of use, seamless connectivity, and beautiful look and feel are important. ECD

Embedded COMPUTING DESIGN
Most Influential WOMEN

Victoria Mitchell

Director of SoC Software Engineering, Altera Corporation

Victoria (Vicki) Mitchell manages the worldwide embedded software team at Altera. Her group is responsible for models and simulation, development tools, and OS enablement for Altera's microprocessor subsystems. She is responsible for the build-out of the software function at the Austin Technology Center. Vicki has more than 25 years of experience in low-level, systems software, mostly at semiconductor companies such as Cirrus Logic, SigmaTel, and IDT. Vicki holds a BS in Software Engineering from Colorado Technical University.

What are the biggest challenges you face on the job every day as a woman in high tech?

MITCHELL: The challenges I've faced have evolved over the years, from being a target of stereotyping to time management while raising a family. One of my biggest challenges is reaching optimal work-life balance. I tend to overcompensate for my gender, to aim for Super Woman, and the workload required for that goal eats into my personal life and interests.

How do you overcome those challenges? What or who is your inspiration?

MITCHELL: My inspiration is Sally Jewell, our next Secretary of the Interior. In her work as CEO of REI, as a banker, and as an oil company engineer, Sally has balanced obligations to her company, a practical obligation to engineering fundamentals, and a passion for the environment.

As women in engineering, we can use the inherent strengths of our gender to help achieve balance in our career, and also in product development. One example in embedded is Linux. We must safeguard developing IP, but at the same time actively promote our technology in the ecosystem. The analogy for women is nurturing and protecting our children, yet raising them in their community. I think women are uniquely comfortable with simultaneously holding close and reaching out.

How can more women be prepared to enter traditionally male fields such as engineering?

MITCHELL: For women in technology fields, success requires fostering two personal attributes that nullify stereotyping and demonstrate significant advantage to the organization:

1. **Leadership:** It cannot be taught, but it can be mentored. As Sheryl Sandberg writes, "Female leaders are key to the solution." It is up to today's leaders to set an example and to inspire. There are practical aspects of leadership applicable to anyone, but women can leverage our innate abilities to nurture and build community.

2. **Fearless creativity:** This is hard to foster when preparing for a tech career because applied science is not fantasy science. An active interest in literature, art, music, and dance helps inspire out-of-the-box thinking and provides a little bravado when voicing ideas.

How do you recognize when a new technology or application is one your company should invest/innovate in, versus a technology that will experience fast burnout?

MITCHELL: Every new idea that solves a problem in a unique way is worth some form of investment, but for how long, and how much? If a technology flames out, maybe the organization couldn't find a path to implement it. Or, it could be because the idea is ahead of its time. In embedded, the investment path often differs significantly from that of application or enterprise software. Agile methodology and object-oriented design took a long time to catch on in embedded, but we now wouldn't do without them. Consider how much simulation, modeling, and SystemC rely on OOD. Even if a technology does not appear sustainable or implementable, it still might be worth limited investment because it is a necessary step to the next "big thing."

In the next 5 to 10 years, which technologies will present the most viable development opportunities for your organization and for the embedded industry?

MITCHELL: The ever-expanding Internet of Things will continue to drive embedded development – specifically distributed tech. Back in the days of mainframe terminal controllers, we all shared resources to keep availability up and costs down. When PCs entered the market, everyone got their own resources, and it was all about owning bigger hard drives and more RAM. Now, we are moving toward everyone having thin clients and sharing resources again. It's all about mobility and always-on availability. This distributed evolution will drive opportunity for Altera and for embedded, in both infrastructure and end-user equipment.

Another factor is consumer trends. As an example: 3D printing is really hot right now. It brings robotics and automation to the people – makes it affordable and available just like PCs did for computing in the early '80s. Robotic technology helps with science, like medicine and mechanics, and also with art, productivity, and efficiency. As engineers, we should keep our eyes open for anything that melds the analog and digital, the human-machine interface that bridges organic and inorganic. ECD

Jane Donaldson

President of Annapolis Micro Systems, Inc.

Jane (Jenny) Donaldson started Annapolis Micro Systems, Inc. in 1982 with Bob Donaldson and Lawrence Marshall, Jr., serving as the company's first president. Annapolis performed custom engineering: software for ground stations for Comsat, point-of-sale terminals for Schlumberger, medical instruments, contract assembly for IBM, and ASIC design for Atmel, in addition to touch technology work with IBM fellow Evon Greanias. In 1994, Jenny guided the company's transition from custom engineering to FPGA-based products. She has a BA in Philosophy and a minor in English from the University of Washington and took Computer Science classes in the late 1970s at the University of Maryland.

What are the biggest challenges you face on the job every day as a woman in high tech?

DONALDSON: Oddly enough, I have never found any challenges caused by just being a woman. My main challenges are as a businessperson in the high-tech community: Pick the right product and develop and implement a tight plan for getting it developed, marketed, manufactured, and sold under budget and within a reasonable timeframe. Keep the customers happy. Keep the cash flowing in and out.

How do you overcome those challenges?

DONALDSON: I think hard all the time, work hard all the time, and constantly review everything and try to do better at everything. I work with the staff, particularly senior management, to help them do the same thing.

What or who is your inspiration?

DONALDSON: My parents, Ben and Jane Van Zwalenburg, taught me by example that:

1. Every human being has a unique and intrinsic value;
2. Work has an inherent worth and we should do every task to the best of our ability;
3. I could and should use my talents to do something significant, to be a useful member of society.

My husband, Bob Donaldson, has stood by me and worked with me through every challenge life has offered us.

Our mentor, Evon Greanias, IBM fellow, taught us that in technology, if you are not pushing the envelope then you are wasting everyone's time and money, and that doing your absolute best is not good enough. It must work and must be on time.

How can more women be prepared to enter traditionally male fields such as engineering?

DONALDSON: My advice for others, both male and female:

1. Figure out your strengths, and play to them.
2. Be brave. You need to make mistakes in order to learn. If you never make mistakes, then you are not pushing yourself hard enough.
3. Challenge yourself. Often and always. You will be surprised to see what you can accomplish.

> "If you never make mistakes, then you are not pushing yourself hard enough."

How do you recognize when a new technology or application is one your company should invest/innovate in, versus a technology that will experience fast burnout?

DONALDSON: We look for a need that is not being filled, something that we believe customers have the money for and will pay for, if it were available. We review the latest available technology and put that together with our technical abilities. If we think we can advance the state of the art, provide a good resolution to the need, and do it within the budget customers have for this problem and within the timeframe that this need will exist, then we go for it.

In the next 5 to 10 years, which technologies will present the most viable development opportunities for your organization and for the embedded industry?

DONALDSON: We have been using FPGAs for processing since 1992. Today we are designing our 14th modular family of FPGA processing boards. Our new line of products, the latest FPGA technologies – combined with the connectivity and speed possible with OpenVPX architectures – and high-performance A/D and D/A will meet the current and near-term needs of our customers in fields like radar, signal processing, SIGINT, ELINT, and communications. I foresee at least another four or five years of technology improvements in FPGAs, A/D, D/A, and SSDs – the fields we currently care most about. Ten years is too far out for me to predict. ECD

Most Influential Women in Embedded contest

Lisa Groeneveld
COO
Logic Supply, Inc.

Susan Wooley
President
Micro/sys

Joyce Tokar
President
Pyrrhus Software

Top Embedded Innovator contest

Raik Brinkmann
President and CEO
OneSpin Solutions

S. Tucker Taft,
VP and Director of Language Research
AdaCore

Charlie Cheng
CEO
Kilopass Technology Inc.

Jason Kridner
Cofounder
BeagleBoard.org

Gerald Coley
Cofounder
BeagleBoard.org

Innovative Products Index

And the Nominations for the

Most Innovative Embedded Products

Are...

Embedded COMPUTING DESIGN
TOP Embedded Innovations

AAEON Electronics, Inc. **POSEIDON CRS-300S-2R**

AAEON's 1U high density/VMS unit is one of the world's first 12 SATA disk (3.5" 4TB each) storage capacity servers. It is a very powerful and reliable system specialized for cloud computing, IP surveillance, and storage appliance applications. It is cost effective, has extended service/longevity (guaranteed up to 5 to 7 years product life) and is equipped with redundant power supply (550W 80 plus gold certified) and fans for efficiency and effectiveness. Other features include Intel® Xeon E3-1200/1200 v2 Series Processors, up to 32GB DDR3 UDIMM memory, 3 HDMI outputs, and Server Grade Intel 82574L Gigabit Ethernet x 2. Its unique design and 12 disk capacity saves not only costs but rack space and offers remote management capabilities.

http://aaeonusa.com/products/details/?item_id=1748 **embedded-computing.com/p9911573**

AAEON Electronics, Inc. **POSEIDON CRS-720X-2R**

AAEON's high performance 2U storage server is one of the world's first ever to offer 66 slim 2.5" SATA Disks in a 2U chassis (64 x Storage + 2 x OS). The CRS-720X-2R is powered by the Intel® Xeon E5-2600 Processor and is ideal for applications such as Object Storage Controller Node, High Performance Computing, High IOPS Storage Server Platforms/Hadoop, and as a cloud computing gateway. The unit features intelligent thermal management with smart fan detection, 9 expansion slots, and utilizes dedicated dual, hot swappable slim 2.5" disks for OS mirroring.

http://aaeonusa.com/products/details/?item_id=1759 **embedded-computing.com/p9911633**

ADI Engineering, Inc.

Seacliff Trail

The Intel Seacliff Trail is a new SDN/OpenFlow enabled reference platform that OEMs can use to quickly create and launch innovative SDN-based products such as switches, routers, and converged storage/compute/security appliances. Seacliff Trail is based on the Intel Alta FM6700 Ethernet Switch, and it includes an integrated Intel Xeon E3-1100 CPU providing control, management and higher-layer packet processing capabilities. ADI Engineering offers Seacliff Trail as a development kit and a private labeled OEM platform, and ADI can quickly tailor the design to meet unique customer requirements. ADI's "Open IP" supply chain model lets customers choose whether to use ADI as their ongoing hardware supplier, or directly build product at customer-preferred factories for the ultimate in control and supply chain efficiency.

www.adiengineering.com **embedded-computing.com/p9911927**

ADL Embedded Solutions Inc.

ADLMES-8200 – Modular Enclosure System

The ADLMES-8200 is a highly innovative embedded enclosure design. Its highly configurable modularity makes it possible to expand or reduce a system without replacing the entire enclosure. Sidewall modules may be added or removed as system requirements evolve. Three standard profiles provide quick-turn inventory availability. A broad portfolio of PC/104 SBC options ranging from Intel® Atom™ (Z510/Z530, D525 Pineview and N2600 Cedarview) to 2nd Generation Intel Core i7 processors is available.

Potential Applications: Military and Defense • Communications Applications • Mini-Routers and Other Network Appliances • Railway Train Control • Transportation • Rugged Industrial Applications • Imaging Applications ... and many more!

www.adl-usa.com **embedded-computing.com/p9911917**

Altia Inc.

User Interface Engineering – Altia Design 11.0

Altia's GUI builder is the centerpiece of Altia's concept-to-code suite for embedded graphics development. With a revolutionized new user interface, Altia's tool chain is the fastest way to get first class interfaces on cost-effective chips. Why has Altia become the industry choice for GUI development?

Reduced engineering time (average 50%) • Portability across multiple hardware platforms • Supports low end processors to high end systems • Works with any operating system or no operating system at all • Custom GUI development (no graphics programming required) • Generates complete, efficient C source code that is ready for production

www.altia.com **embedded-computing.com/p9911746**

AMD

AMD Embedded G-Series System-on-Chip (SOC)

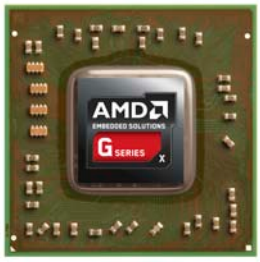

The AMD Embedded G-Series System-on-Chip (SOC) is a fully integrated processor with CPU, GPU and I/O controller on the same chip. It's the world's first embedded quad-core x86 SOC and the first available SOC based on AMD's "Jaguar" microprocessor design and AMD Radeon™ 8000 Series graphics. The AMD G-Series SOC uses 33 percent less board real estate than its predecessor while providing unparalleled performance and excellent energy efficiency with a thermal TDP below 9 watts, yet scaling to 25 watts enabling fan-less cooling designs. The platform offers critical features for embedded applications including ECC, Industrial Temperature Range and remote management.

www.amd.com/us/products/embedded/processors/Pages/g-series.aspx **embedded-computing.com/p9911833**

AMTELCO

AMTELCO PCI Express E&M Interface Boards

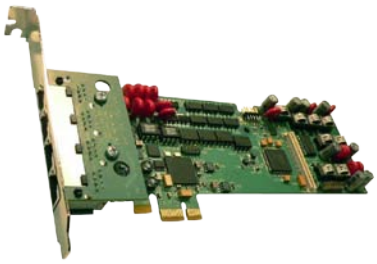

AMTELCO PCI Express E&M interface boards accommodate connections to radio dispatch equipment, with manual push-to-talk signaling control, echo suppression, and VOX resources. Standard Type I, Type IV, and Type V E&M protocols are also supported. Audio interfaces support both 2-wire and 4-wire modes and can be configured on a per-port basis. AMTELCO PCIe E&M interface boards are available as 4-port (x1) half-height boards for HMP environments, as well as 8-port (x1) full-length boards for H.100 TDM switching environments. Software drivers are available for most common operating systems, including channel drivers for use in Asterisk® environments. AMTELCO PCIe boards are also available for alternate line interface types, including station interface, loop start interface, and T1/E1/ISDN digital interfaces.

www.amtelco.com **embedded-computing.com/p9911907**

Annapolis Micro Systems, Inc.

WILD Data Storage Solution – 8 TB/Slot

8 TB of Storage Per Each 6U OpenVPX Slot • Up to Sixteen SSDs per Board in a Single 1" OpenVPX Slot • Up to 4 GB/s Write and Up to 8 GB/s Read Bandwidth • 8x PCIe Gen3 for almost 8 GB/s Full Duplex to each Adjacent Slot on Backplane (P2 and P5) • Up to 40Gb Ethernet or QDR InfiniBand on each of Four Fat Pipes on P1, for a total of 20GB/s on P1 • Hot Swappable Canister with up to Sixteen Pluggable 1.8" SSD SATA 3.x Drives, with 2, 4 or 8 Terabytes per Canister • 10,000 Insertion Cycle Rating • Canister Carrier Remains in Chassis • 4 GB DDR3 DRAM • System Management using IPMI • API for Command and Control of the Storage Process • Diagnostic Monitoring and Configuration over Backplane or Front Panel Status LEDs for 16 SSDs and Backplane Control, Data Plane Connections

www.annapmicro.com **embedded-computing.com/p9911930**

Annapolis Micro Systems, Inc. **WILDSTAR COTS OpenVPX Ecosystem**

WILDSTAR COTS OpenVPX Ecosystem begins with a 4 slot chassis with a custom designed high performance mesh backplane, a 6 slot chassis with a custom designed high speed switched backplane with room for 1 switch card, and a 12 slot chassis with a custom designed high speed backplane with accomodation for 2 switch cards. It also includes WILDSTAR 6 FPGA Processing boards, with 1, 2 or 3 Xilinx Virtex-6 FPGAs, and WILDSTAR A5 FPGA Processing boards, with 1, 2 or 3 Altera Stratix V FPGA Processing boards. Each of the processing cards can have up to 2 A/D or D/A I/O mezzanine cards, and up to 6 QSFP+ front panel I/O connectors, each capable of up to 40Gb Ethernet. We have a one slot 8 TB WILD Storage Board, and a WILD Switch Board that will carry data at up to 14 Gbps on its dataplanes.

www.annapmicro.com **embedded-computing.com/p9911916**

Annapolis Micro Systems, Inc. **WILD OpenVPX 40 Gb Ethernet and FDR InfiniBand Switch**

Multiprotocol Switch: SDR/DDR/QDR/FDR InfiniBand and 1/10/40/56 Gb Ethernet. Up to Four Tb/s Non-Blocking Switching Capacity with up to Eight Switch Partitions • 6U OpenVPX (VITA 65) Compliant • Supports OpenVPX Slot Profile: SLT6-SWH-16U20F-10.4.2-n • Backplane Ports: Sixteen 1Gb Ethernet Lanes, Twenty High Speed Four Lane Data Plane Connections • Front Panel Ports: Eight QSFP+, Two SFP+, RJ45 Managment Port, USB UART, Status LEDs • Each Backplane and Front Panel Port can be Configured for either InfiniBand or Ethernet • InfiniBand and IP Routing • Ethernet Gateways • Data Center Bridging (DCB): FPC, DCBX, ETS Layer 2 Protocols: VLANs (802.1q), DHCP Server, IGMP Snooping • Chassis Management Control using Intelligent Platform Management Interface (IPMI)

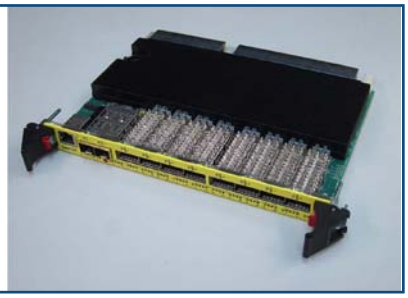

www.annapmicro.com **embedded-computing.com/p9911931**

Bitreactive AS **Reactive Blocks Software Development Kit**

Reactive Blocks Software Development Kit radically simplifies the way Java developers build mission-critical systems. It is a tool that automatically generates more than 60% of the required code. It lets designers see and modify their system graphically, and understand how their systems behave under pressure. Designers use verified building blocks for connecting to popular network services, including 4G data networks and M2M cloud services, and it also enables the Internet of Things. Reactive Blocks SDK also has many additional advantages.

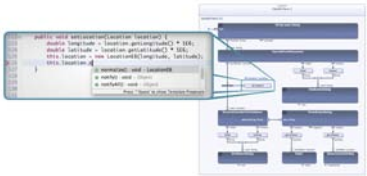

www.bitreactive.com **embedded-computing.com/p9911721**

Breker Verification Systems, Inc. **TrekSoC**

Functional verification of complex embedded systems is hitting a wall. Traditional simulation testbenches do not interact with embedded processors and hand-writing diagnostic code to run on the processors is hard and time-consuming. Jumping directly to hardware-software coverification is not the answer, since production software is inefficient at finding hardware bugs. Breker's TrekSoC product fills this gap by automatically generating C test cases to run on embedded processors in simulation. These test cases are multi-threaded and span multiple heterogeneous processors. Hardware teams can find bugs earlier and easier, while software teams can focus on developing production code rather than throwaway diagnostics.

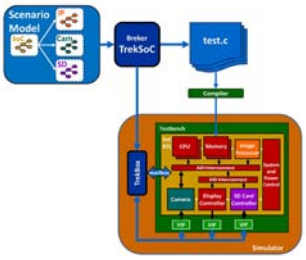

www.brekersystems.com **embedded-computing.com/p9911990**

Cadence Design Systems **Incisive Debug Analyzer (IDA)**

Companies today spend more than 50% of verification efforts in debug because bugs are hard to find at the HDL or testbench level. In 2012, Cadence solved this problem with its new Incisive Debug Analyzer (IDA) – a multi-language, "interactive" post-process debug solution that helps customers find bugs in minutes instead of hours. As the industry's only tool to deliver comprehensive and innovative debug functionality in a single, integrated and synchronized debug environment – IDA can cut debug time by 40-50% and by more than 2X on complex bugs. Intelligent handling of debug data • "Interactive" step by step debug in post process – even backwards • Guided debug providing clues where to look for errors • Collaborative debug - easily share data files globally • Sophisticated search/filter capability

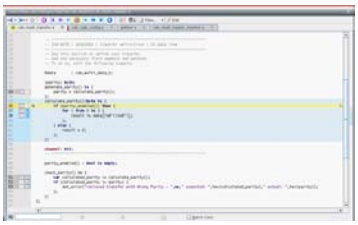

www.cadence.com/us/pages/default.aspx **embedded-computing.com/p9911637**

Cavium Networks **OCTEON Fusion Processors/CNF71XX**

OCTEON Fusion® Family, a small cell "Base Station-on-a-chip" family, is specifically designed for LTE & 3G multi-mode small cell base stations, including picocell and micro base stations. OCTEON Fusion processors combine OCTEON's widely successful multi-core architecture along with purpose-built Baseband DSP cores, extensive LTE/3G hardware accelerators and digital front end (DFE) features into a single chip, slashing the BOM cost and power envelope of small cell base stations. This product family consists of multiple SoCs (systems-on-chip) that enable small cells ranging from 32 users to 300+ users and up to dual 20MHz carriers, delivering industry-leading scalability. OCTEON Fusion is accompanied by FusionStack™ software, a comprehensive, interoperability tested (IOT), carrier-class L1 to L3 software suite.

www.cavium.com/OCTEON-Fusion.html **embedded-computing.com/p9911800**

Cavium Networks — **OCTEON II CN68XX family of Multi-Core MIPS64 Processors**

The OCTEON II CN68XX family of Multi-Core MIPS64 Processors targets high-performance, high-throughput, service-rich applications in secure data center, mobile Internet, and borderless enterprise applications. The family includes six software and pin-compatible processors, with 16 to 32 cnMIPS64 v2 cores, over 85 application acceleration engines, and real-time Power Optimizer features. Extremely high-bandwidth connectivity based on the latest standards-based SERDES I/Os including PCIe Gen2, XAUI, Double-speed XAUI, RXAUI, and Interlaken enable throughputs over 40Gbps using a single chip or scaling to over 100Gbps using multiple chips. The CN68XX offers a 3x compute advantage and 2x performance per Watt advantage over alternative solutions.

www.cavium.com/OCTEON_MIPS64.html — **embedded-computing.com/p9911852**

Cavium Networks — **NITROX® III Security Processor Family, CNN35xx**

The NITROX family supports cryptography and compression at up to 40 Gbps and up to 200K RSA Ops/sec. It is ideally suited for enterprise routers and switches, data center networking appliances, integrated security appliances and service provider infrastructure.

• Application Delivery Controller (ADC) • WAN Optimization • Cloud Computing • Server Offload – Web, Mail, Search • Integrated VPN/Firewall appliances • UTM Gateway, Routers • Load Balancers/L4+ Switches • Wireless LAN and WAN equipment • Security Telco gateways • SAN appliances • 3G & 4G/LTE infrastructure equipment

www.cavium.com/processor_security_nitrox-III.html — **embedded-computing.com/p9911908**

Connect Tech, Inc. — **Xtreme/GPU**

The Xtreme/GPU is the first high-end GPU bringing exceptional desktop-level graphics, outstanding multimedia features, and GPU processing power to the PCIe/104 form factor, with the ability to select either AMD or NVIDIA graphical solutions. Partnered with the AMD solution, the Xtreme/GPU can be used to drive up to four Mini DisplayPort connections for either independent use, or leveraging AMD Eyefinity technology, as a single large surface. On-board passive circuitry allows for the use of Dual-Mode DisplayPort, enabling the use of DVI or HDMI dongles if DisplayPort is not preferred. Combined with the NVIDIA solution, the Xtreme/GPU transforms into a processing powerhouse. With access to the Nvidia CUDA™ Cores, the GPU can become a parallel computational CPU for non-graphical applications.

www.connecttech.com — **embedded-computing.com/p9911926**

Creative Electronic Systems — **Multi-Format Video Processor**

With the VBP-2764, CES has succeeded to combine multiple video capture interfaces with multi-standard compression and custom processing into a rugged single-slot 3U VPX solution. Multiple digital and analog HD video channels can be acquired and processed in a powerful FPGA, and compressed with the H.264 and JPEG-2000 codecs. The VBP-2764 is intended for embedded systems with high-resolution cameras that require powerful video processing and compression at low power consumption. Through its PCIe x8 interface, the VBP-2764 is fully integrated with the full range of 3U VPX processors from CES, including Intel® and Freescale® SBCs and Xilinx FPGA boards, and is fully supported by the CES Video API under Linux® and VxWorks®.

www.ces.ch — **embedded-computing.com/p9911830**

Cypress Semiconductor — **PSoC 4 Programmable System-on-Chip**

The PSoC® 4 4100 and 4200 programmable system-on-chip families combine programmable precision analog circuitry, programmable digital blocks, and fully routable I/Os with Cypress' industry-leading CapSense® capacitive touch technology and a power-efficient ARM® Cortex™-M0 core. With PSoC 4, designers can replace an entire portfolio of proprietary MCU-based solutions and migrate legacy 8-bit and 16-bit designs to a single 32-bit platform. The truly scalable, cost-efficient families leverage the easy-to-use PSoC Creator™ integrated design environment, enabling designers to drag and drop dozens of free, pre-characterized, and production-ready analog and digital IP blocks – PSoC Components™ – into a single PSoC device to create customized, feature-rich products.

www.cypress.com — **embedded-computing.com/p9911928**

Express Logic, Inc. — **TraceX® Real-Time Event Analysis Application**

Express Logic's TraceX® real-time event analysis application gives developers an "instant-replay" of all RTOS and application events that occurred during a period of interest on their embedded system. TraceX paints a picture of each event, in a sequential or time-stamped format, broken down by application thread and RTOS subsystem. Interrupts, context switches, priority inversions, and all thread activities can be mapped with color-coded icons for easy view, with pan and zoom capability for zeroing in on a period of interest. TraceX aids in debugging, and in system optimization, by revealing exactly what the real-time system has done, step-by-step. TraceX is available for multicore systems, providing a breakdown of all system activity by core and by application thread within each core.

www.rtos.com — **embedded-computing.com/p9911859**

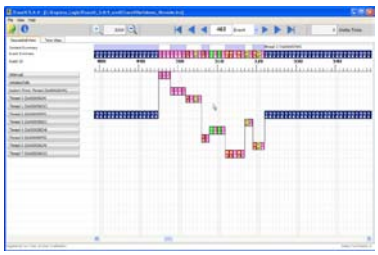

Freescale Semiconductor — **Healthcare Analog Front End (AFE) reference platform**

Freescale's Healthcare Analog Front End (AFE) reference platform is a comprehensive, easy-to-use solution that allows medical device designers to speed time to market for a wide variety of healthcare devices. Featuring the Kinetis K53 MCU with analog measurement engine and the Tower System modular development platform, it reduces system cost, board size and complexity. The highly integrated reference platform with hardware, software and schematics includes six plug-in boards that support all of the key vital signs including blood pressure, digital stethoscope, electrocardiograph, spirometry, pulse oximetry and glucometry. It is the most comprehensive healthcare-specific reference platform in the industry.

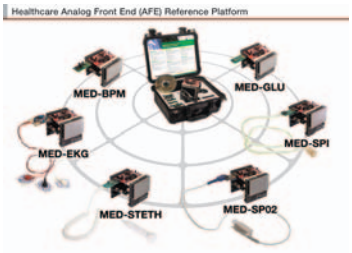

www.freescale.com — **embedded-computing.com/p9911823**

Freescale Semiconductor — **Kinetis L Series**

Kinetis L series is the first microcontroller build on the ARM® Cortex™-M0+ processor, featuring the latest low power innovations and high performance, high precision mixed-signal capability. It is supported by the Freescale Freedom development platform, along with third-party development resources from the extensive ARM ecosystem. Recently, with the proliferation of the Internet of Things (IoT), Freescale announced Kinetis KL02, the world's smallest ARM Powered® 32-bit MCU. It's small in size while retaining rich MCU feature integration, flash, analog, timers and communication interfaces – everything a body needs to be a basic tiny computer, ideal for space-constrained applications in the industrial, consumer and healthcare markets.

www.freescale.com — **embedded-computing.com/p9911822**

FTDI Chip — **Innovative Graphic Controller IC**

With display, audio and touch functionality on a single-chip, FTDI Chip's FT800 Embedded Video Engine (EVE) is optimized to provide high quality human machine interfaces (HMIs), while lowering overall system costs, minimizing board space utilization and shortening development times. Suitable for QVGA/WQVGA TFT display implementations, its unique object oriented approach (where objects can be images, fonts, or audio elements) simplifies the development process and dispenses with the need for frame buffer memory. The FT800's functionality includes: four-wire intelligent resistive-touch sensing, audio processing of midi-sounds with audio playback and the graphical capability to render images to 1/16th pixel resolution.

www.ftdichip.com/EVE.htm — **embedded-computing.com/p9911832**

HCC-Embedded — **HCC 'Rapid Customization' Tool for Embedded Software**

HCC's 'build-to-order' tool is a new approach to embedded software. HCC has created an online tool that allows engineers to customize key features of their project including MCU, development board, RTOS, compiler, peripherals, flash memory, etc. All projects are checked for correctness to create a fully customized project using HCC middleware. This removes the need for engineers to integrate 3rd party software into their target environment since the project described is delivered fully customized and tested. This is possible since HCC's approach to developing embedded software uses a strict methodology, independent of target specific dependencies such as RTOS or Compiler. MCUs including ARM Cortex and Renesas are supported with some projects available in as little as 48 hours.

www.hcc-embedded.com — **embedded-computing.com/p9911918**

Icon Labs — **Floodgate™ Defender**

Floodgate™ Defender is a small footprint security appliance providing drop-in security to unprotected legacy endpoint devices. There is a vast installed base of unprotected legacy devices operating our critical infrastructure, transportation systems, factories and homes. Many of these devices cannot be upgraded and lack security. Floodgate Defender fills a critical security hole by providing easy to install and manage protection for these devices. Floodgate Defender Security Appliance stops cyber-attacks against industrial, military and other fixed function devices. By controlling which packets are processed by endpoint devices, Floodgate Defender blocks attacks before a connection is established, providing protection from a wide array of attacks.

www.iconlabs.com — **embedded-computing.com/p9911846**

Impinj Inc. — **Indy RS500 UHF RFID Reader SiP**

The Indy RS500 RFID reader SiP (system in package) is a completely integrated UHF reader solution in a small, 30 x 32mm, surface mount package that enables OEM and device manufacturers to achieve fast time-to-market with low development risk. Designed to work as a surface mount technology (SMT) component in a standard PCB manufacturing process, the Indy RS500 SiP eliminates costly mechanical hardware and RF cables that are typically required with reader modules on the market today. Ideal for moderate read range of small tag populations, the Indy RS500 SiP's small form factor enables a diverse range of applications – such as consumables authentication, access control, process control, and anti-counterfeiting – that needs low-cost embedded UHF Gen 2 RFID reader capability.

www.impinj.com — **embedded-computing.com/p9911856**

Infineon Technologies Corporation — XMC1000 Microcontroller Family

Infineon's XMC1000 is the only ARM® Cortex™-MO based microcontroller family manufactured in a 65nm embedded flash process on 300mm wafers, delivering true 32-bit power at an 8-bit price. Flash sizes from 8-200 kB eliminate memory constraints faced with other low-end price point devices. Designers also leverage a pool of powerful integrated peripherals, including PWM units, 12-bit ADC, and user-configurable serial communications interfaces. Some family members include innovative application specific peripherals such as a patented LED brightness and color control unit and a powerful math coprocessor to support complex control loops. Use common tool chains or work solely in Infineon's DAVE™ IDE, bringing component-based programming to the low-end microcontroller segment.

www.infineon.com/xmc1000 — **embedded-computing.com/p9911848**

InForce Computing, LLC — IFC6410

IFC6410 offers 1.7 GHz Quad Core A15 class processing powered by Snapdragon. It's the first SBC featuring the Snapdragon S4 Pro APQ8064, and at U.S. $149 per unit, it's pretty close to a no-brainer for any serious embedded developer. It's small – 100mm x 70mm – with a Pico-ITX footprint that's easy to design around with powerful CPU and multimedia performance – 1.7 GHz Quad Core Krait CPUs, with advanced 1080p HD video, and graphics, with industry leading performance-to-power ratio. Full featured Pico-ITX Single Board Computer for embedded applications that is product ready • Rich multimedia capabilities with 1080p capable video and graphics accelerators • Android and a Linux software stack • Bringing the power of the leading edge mobile technologies, like gesture and Alljoyn, into the embedded world

www.inforcecomputing.com — **embedded-computing.com/p9911920**

InnoDisk — iSLC Technology for Solid State Drives

InnoDisk's new high-performance iSLC™ SSD Technology is available on SATA III & SATA II Series products. By using superior quality, preselected MLC NAND Flash and InnoDisk's patented Flash management algorithms, iSLC nearly approaches SLC performance and reliability at about half the cost. iSLC™ SSD Technology reprograms the two bits per cell of MLC into one bit per cell, which increases the sensitivity delta between each level. This practice enables the NAND Flash to perform similarly to an SLC Flash-based solution. The average endurance of iSLC can surpass 30,000 Program/Erase (P/E) cycles, which increases the lifespan of the drive over MLC Flash by a factor of 10. Additionally, write performance for iSLC is about 70% faster than MLC on SATA II.

www.InnoDisk.com — **embedded-computing.com/p9911825**

InnoDisk — SATA III Series – Solid State Drives

InnoDisk's latest SATA III Series of products includes our 2.5" SSD, SATADOM, SATA Slim, mSATA, and CFast. InnoDisk SATA III storage devices offer sequential read and write speeds up to 500/340MBps – about two times faster than SATA II. InnoDisk offers wider temperature ranges (-20°C to 85°C), customized firmware, power outage data protection and other advanced functions. The upgraded SATA III 2.5" SSDs provide a lower power design and reduced operating costs with enhanced power saving slumber and sleep modes. All SATA III devices come with InnoDisk's iSMART proprietary disk health monitoring tool. iSMART's one touch activate mode can save hours of setup and maintenance time, and gives users easy one touch access to features such as write-protection, quick erase and power saving settings.

www.InnoDisk.com — **embedded-computing.com/p9911824**

Integrated Device Technology, Inc. — Dual-Mode Wireless Power Receiver: IDTP9021

The IDTP9021 is an enhanced version of IDT's IDTP9020 wireless power receiver that complies with the PMA Type 1 Interoperability Specification in addition to the WPC "Qi" standard, allowing customers to consolidate multi-chip configurations into a single-chip solution. IDT's device is the industry's first to receive pre-certification from the PMA, and is the only solution with demonstrated compatibility between the two most popular magnetic induction-based wireless power standards. The IDTP9021 integrates a high-efficiency synchronous full bridge rectifier, high-efficiency synchronous buck converter, and control circuits to wirelessly receive an AC power signal from a compatible transmitter and convert it into a regulated 5 V output voltage for powering and charging portal electronics.

www.idt.com — **embedded-computing.com/p9911847**

Intel — 4th Generation Intel® Core™ processors

4th generation Intel® Core™ processors provide superior CPU/graphics/media performance and greater power efficiency over 3rd generation processors for intelligent systems in industrial, medical, retail, defense and aerospace environments. Intel® AVX 2.0 delivers improved integer/matrix-based calculation abilities for signal and image processing medical and defense applications. A 15-watt, one-chip package offers high performance and small size to enable medical device OEMs to design thinner, lighter, and sleeker portable ultrasound equipment and patient monitors. The next generation Intel graphics engine significantly improves graphics/media performance to enable compelling visual experiences for digital signs and kiosks that display videos, 2D/3D graphics, and interactive content.

www.intel.com/embedded — **embedded-computing.com/p9911906**

Intel — MinnowBoard

MinnowBoard is the first open hardware/open software, Intel® Atom™ 1.0 GHz 32-bit CPU with Hyper-Threading and Virtualization Technology based board which introduces the Intel® Architecture to the small and low cost embedded market for the developer and maker community. MinnowBoard combines great performance and a mature ecosystem via X86 compatibility/PC architecture standards (PCIe, SATA and USB) which provide generous I/O, video capabilities and Gigabit Ethernet. MinnowBoard also includes key embedded standards (SPI, I²C, CAN, GPIO) all with Angstrom, a Yocto Project Compatible Linux® Distribution and UEFI firmware with Fast Boot. The board is expandable via a flexible expansion connector which is chassis-friendly and stackable. The MinnowBoard is a powerful, low cost solution for embedded projects.

www.minnowboard.org **embedded-computing.com/p9912045**

Intel — Yocto Project

"It's not an embedded Linux distribution – it creates a custom one for you"

The Yocto Project is an open source collaboration project that provides templates, tools and methods to help you create custom Linux-based systems for embedded products – regardless of the hardware architecture. It provides a common environment to end the fragmented development process of having to use a different toolset for each architecture. The free tools are easy to get started with, powerful to work with (including emulation environments, debuggers, an Application Toolkit Generator, etc.). Community-tested images include the Yocto Project kernel and build profiles supporting multiple architectures including ARM, PPC, MIPS, x86, and x86-64. Specific platform support is via BSPs layers.

www.yoctoproject.org **embedded-computing.com/p9912046**

Klocwork — Klocwork Cahoots

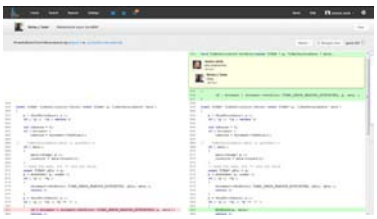

Klocwork Cahoots is a flexible and easy-to-use peer code review tool that simplifies the code review process. Language-agnostic and designed for development teams of all sizes, Klocwork Cahoots fits into the developer workflow to ensure code reviews are both effective and fast.

Klocwork Cahoots is: • Simple – Start a code review, participate, add reviews and more – with just one click. • Personal – Design your own space. Discover activity and participate based on your interests. • Social – Follow people, teams and components with customized feeds on an infinite wall. Join in threaded discussions. • Fast – Get real-time notifications, assign actions, leverage intuitive search and create instant custom reports.

www.klocwork.com **embedded-computing.com/p9911861**

Klocwork — Klocwork Insight™

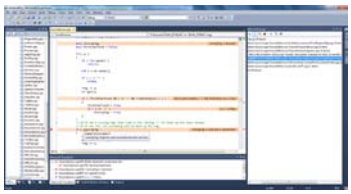

Klocwork® has reinvented source code analysis with Klocwork Insight™. It is the first C/C++ source code analysis tool to provide 'on-the-fly' analysis. Similar to a 'spell checker' found in popular word processors, Klocwork Insight instantly underlines defects as developers are writing their code. In the background, sophisticated syntax and dataflow analysis, build comprehension and incremental analysis are continuously running to ensure thorough and highly accurate defect detection. With Klocwork Insight's on-the-fly analysis, individual developers get all the benefits of complex code analysis at their desktop instantly, simply and with little training or configuration needed.

www.klocwork.com **embedded-computing.com/p9911860**

Lattice Semiconductor Corporation — Lattice iCE LP384 Field Programmable Gate Array

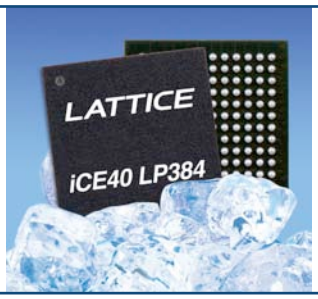

Lattice's iCE LP384 FPGA enables designers to rapidly add capabilities that differentiate cost-sensitive, very small, low power systems such as smartphones, portable medical equipment, cameras, and embedded systems. The tiny FPGA has 384 look-up tables (LUTs), consumes only 25 Microwatts, and is available in ultra-small packages (2.0 x 2.0 mm) with high-volume cost structure. The device offers 100 MHz hardware performance for high speed processing that turns incoming data into valuable information. Key functions include: Universal I/O Interface bridge to modules, apps processors, ASICs and ASSPs • Ultra-low power sensor manager and hub to increase information gathering • Timing critical control for coprocessing functions • Flexible I/O and port expander that enable processors and ASICs to easily add new functions

www.latticesemi.com **embedded-computing.com/p9911929**

LynuxWorks, Inc. — LynxSecure 5.2

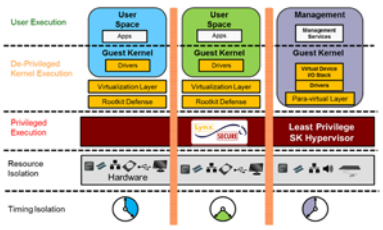

LynxSecure is a mature "Type-0" hypervisor that offers the combination of ultra security, based on a least privilege architecture, real-time performance using separation kernel technology and flexible virtualization married with processor virtualization extensions. Well suited to embedded systems due to its small footprint and real-time scheduler, LynxSecure 5.2 offers a new level of protection for the connected devices that will be running our critical infrastructure in an M2M environment by adding a rootkit defense layer beneath the virtualized operating systems. Rootkits are the most sophisticated and lethal type of malware-stealthy and extremely potent. When LynxSecure 5.2 is used on embedded devices, it can help detect these malicious infections as they inject their payload, long before the start of the actual cyber attack.

www.lynuxworks.com **embedded-computing.com/p9911991**

MEN Micro Elektronik GmbH

F75P CompactPCI PlusIO SBC

The F75P CompactPCI PlusIO SBC was specially designed for rail applications and brings functional safety to the board level, using redundant Intel Atom processors. The dual-redundant safe computer is certified according to SIL 4 – the highest safety level for railway applications. While two redundant CPUs with independent RAM and Flash memory increase system safety, a third processor controls the I/Os across an FPGA-based communication interface. Two supervisor units monitor the environmental conditions and an event logger allows tracing of any hardware and application events. The extended temperature range of -40 to +85°C, EN 50155 conformity, conformal coating and assembly of M12 connectors make the F75P a safe and at the same time robust solution.

www.men.de **embedded-computing.com/p9911784**

MicroMax Computer Intelligence

Rugged M-Max 810 – 1/2 ATR PC/104 system enclosure

The M-Max 810 high-performance rugged industrial computer provides reliable operation in tough environments. The fully ruggedized 1/2 ATR PC/104 aluminum chassis is fanless and uses natural convection and conduction cooling in accordance with MIL-STD-810 standards. COTS technology components allow configuring the M-Max 810 family to comply with a wide variety of airborne, marine and ground vehicle applications. Providing dust and humidity protection, the Max 810 can operate under extreme temperatures up to +65°C. Delivering excellent performance comparable to high-end desktop systems, it can dissipate up to 60W of power. Rugged construction and sealed case for tough environments – airborne, marine and ground vehicle applications.

www.micromax.com/catalog/product.shtml?i=2:1-0:15 **embedded-computing.com/p9911652**

Microsemi

SECURRE-Stor

Microsemi's SECURRE-Stor self-encrypting, high performance solid state drive (SSD) incorporates NIST certified AES 256 encryption core and is made in the United States. The product is targeted at banking, financial, medical, industrial, and secure laptop applications requiring a very high level of data security. The SSD can be used in a networked environment with full IT control of key management. Keys can be self-generated, download by IT/user, or be customized for the application. A unique feature also allows the network Host to erase the key, leaving the SSD with unbreakable encrypted info. The SSD can also be configured to then automatically erase 100% of the encrypted memory. To learn more visit www.microsemi.com or email sales.support@microsemi.com.

www.microsemi.com **embedded-computing.com/p9911849**

Microsemi

SmartFusion2 SoC FPGAs

Microsemi's SmartFusion®2 system-on-chip (SoC) field programmable gate array (FPGA) family features the only devices that address important fundamental requirements for advanced security, high reliability and low power in critical industrial, military, aviation, communications and medical applications. SmartFusion2 integrates inherently reliable flash-based FPGA fabric, a 166 megahertz (MHz) ARM® Cortex™-M3 processor, advanced security processing accelerators, DSP blocks, SRAM, eNVM and industry-required high-performance communication interfaces all on a single chip.

www.microsemi.com **embedded-computing.com/p9911850**

Parasoft

Development Testing Platform for Embedded Systems

The Development Testing Platform for Embedded Systems is the result of Parasoft's 25 years of experience creating software development solutions. The Development Testing Platform automates the application of software quality activities throughout the SDLC, which prevents errors from entering the code at the process level. The open platform is the only solution on the market to directly drive business objectives using a policy-driven approach, while providing 360-degree process intelligence and unprecedented visibility into development activities. The platform's business-driven architecture, unparalleled code analysis tools, and traceability reporting help organizations deliver compliant software on time and on budget.

www.parasoft.com **embedded-computing.com/p9911923**

Pentek, Inc.

Bandit Analog RF Slot Receivers

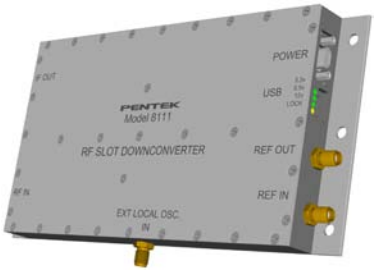

The Bandit™ Model 8111 provides a series of high-performance, low-cost RF slot receiver modules. Packaged in a shielded enclosure, the modules offer programmable gain, high dynamic range and low noise figure. Input options cover specific frequency bands of the RF spectrum, and have IF output optimized for A/D converters. Seven different input-frequency band options are offered, each tunable across a 400 MHz band with an overlap of 100 MHz between adjacent bands. These options accommodate signals from 800 MHz to 3.000 GHz. The 8111 is an ideal solution for receiving, amplifying, and downconverting antenna signals for satellite communications, radar, and signal intelligence systems.

http://pentek.com/go/ecd8111 **embedded-computing.com/p9911855**

Pico Technology

PicoScope 5000 Series

Most digital oscilloscopes gain their high sampling rates by interleaving multiple 8-bit ADCs. Despite careful design, the interleaving process introduces errors that always make the dynamic performance worse than the performance of the individual ADC cores. The new PicoScope 5000 Series scopes have a significantly different architecture in which multiple high-resolution ADCs can be applied to the input channels in different series and parallel combinations to boost either the sampling rate or the resolution. In series mode, the ADCs are interleaved to provide 1 GS/s at 8 bits. Interleaving reduces the performance of the ADCs, but the result (60 dB SFDR) is still much better than oscilloscopes that interleave 8-bit ADCs. This mode can also provide 500 MS/s at 12 bits resolution.

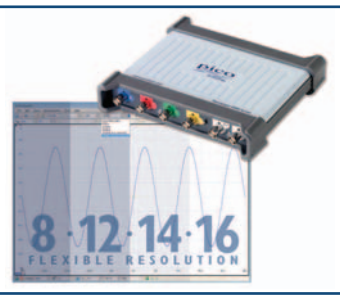

www.picotech.com **embedded-computing.com/p9911857**

PLDA

System-On-Module (SoM) Platforms

PLDA announced a project to produce system-on-module (SoM) platforms. Embedded software developers can use the SoM and its carrier board as a turnkey prototyping solution. This approach allows them to move to production with the same SoM architecture on a customized carrier board with the right set of interfaces and extensions, either designed internally or by PLDA's and ReFLEX's partners, optimizing time-to-market. The new line of PLDA SoM products is built by ReFLEX CES using PLDA's IP and will be marketed and supported by PLDA through its worldwide sales and distribution channels. The PLDA SoM line is expected to include products based on leading SOC FPGA platforms from Altera and Xilinx.

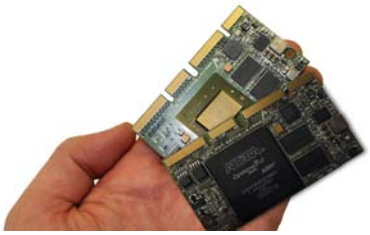

www.plda.com **embedded-computing.com/p9911922**

pls Development Tools

Universal Access Device 2pro

The UAD2pro is a high-performance target access device for on-chip debugging of a wide range of microcontroller architectures like TriCore, Power Architecture (e200), Cortex, ARM, XC2000/XE166 and others. Connection to the target takes place by means of architecture-specific adapters – available for JTAG, cJTAG, DAP, SWD, OnCE and other manufacturer-specific debug interfaces – via a universal interface. For debugging microcontroller boards with high-voltage components, as is common for instance with motor or inverter controls, target adapters with an electrical isolation can optionally also be used with the UAD2pro. The UAD2pro is offered together with a software license of the Universal Debug Engine (UDE).

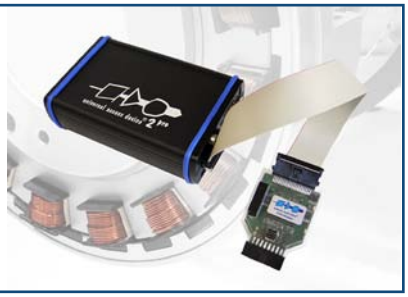

www.pls-mc.com **embedded-computing.com/p9911909**

Presagis

VAPS XT

VAPS XT is an easy-to-use yet powerful software tool from Presagis to rapidly develop Human-Machine Interface (HMI) displays for avionics and other embedded applications. Users can design sophisticated graphics for deployment to actual embedded systems through VAPS XT's support for advanced graphical features, including transparencies, gradients, texture mapping, and smooth shading. VAPS XT also helps designers and developers of embedded avionics displays meet the ARINC 661 standard easily and cost-effectively. Launched in April 2013, the latest release of VAPS XT includes design, performance, and usability enhancements that streamline HMI development. This includes access to an expanded library of pre-built and coded widgets that make it easier to develop ARINC 661 compliant avionic displays.

www.presagis.com/eg **embedded-computing.com/p9911921**

Rapita Systems Ltd.

DO-178B/C Tool Qualification Pack for RapiTime

With this qualification pack, RapiTime becomes the only tool that can be qualified to automate measurement-based worst-case execution time (WCET) for projects requiring DO-178B/C certification (the de facto standards governing the development of all avionics software). DO-178B/C forbids automation without evidence of the automation tool's reliability, something only achieved through qualification. Determining WCET is recommended under DO-178B/C guidelines. However, determining WCET through existing manual measurement and analysis techniques is effort-intensive and difficult to conduct. The pack therefore allows RapiTime to replace manual activities in projects requiring DO-178B/C certification, resulting in significant reductions in effort while maintaining (or even improving) quality levels.

www.rapitasystems.com **embedded-computing.com/p9911854**

Red Rapids

Signal Stream

The Signal Stream product family transforms a general purpose computer into a high speed signal acquisition and generation platform. A wide selection of sample rates and support for AC or DC coupled analog connections is offered, and different form factors target either an embedded chassis or traditional server/desktop environment. The Signal Stream hardware incorporates a rich set of software programmable features that includes selectable operating modes (continuous, snapshot, periodic), external or timed event triggers, timestamped data samples, and flexible data formatting. Each channel can stream raw samples or data packets defined by the VITA 49 specification. A unified code base allows application developers to transition between hardware options and across operating systems with a common API.

www.redrapids.com **embedded-computing.com/p9911718**

SafeNet, Inc.

SafeNet Sentinel Embedded Solutions

SafeNet helps intelligent device manufacturers evolve their business models to be more innovative in increasingly competitive markets. With SafeNet, intelligent device manufacturers not only can protect their business-critical IP from tampering, reverse engineering, and theft, but they can dramatically change how they package and therefore monetize their offerings. With Sentinel, intelligent device manufacturers (or embedded software vendors) leverage licensing and entitlement management technology to remotely deliver upgrades, bug fixes, and provide support to improve user experience while reducing operating costs. Remote feature management also enables a dramatic reduction in the number of separate hardware SKUs required to support a feature-rich product line.

www.safenet-inc.com/embedded **embedded-computing.com/p9911785**

Siemens AG

SIMATIC IPC 427D Microbox

Siemens embedded industrial computers are small in size but big on performance. The SIMATIC IPC 427D Microbox features a compact, fanless design that allows OEMs to reduce their IPC footprint by up to 20% – without sacrificing performance. In fact, with a 3rd Generation Intel Core i7 processor, the SIMATIC Microbox PC delivers the highest performance-to-size ratio in its class. The design is more user-friendly. With flexible mounting options and simple wiring connections, the IPC can be integrated with new or existing cabinet designs with minimal effort. Compact, flexible and powerful, the SIMATIC Microbox PC is a perfect fit for any industrial environment.

usa.siemens.com **embedded-computing.com/p9911831**

Swissbit

Swissbit X-500 Series Industrial SATA II 2.5" SSD

The rugged, high-performance Swissbit X-500 Series Industrial SATA II SSD 2.5" storage solution is designed for demanding applications. The data rate reaches up to 240 MB/sec and an impressive 14'500 IOPS with 4 KB random accesses. The high end architecture utilizes up to 8 channels with SLC NANC Flash. Special features such as ATA-8, NCQ and TRIM support enable higher IOPS (input/output per second) and sequential performance providing the best combination of performance and reliability for industrial applications. Additionally, the S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) protocol enables full control of mission critical data all the time. The BCH-ECC (error correction code) ensures data reliability with the power fail protected X-500 Series.

www.swissbit.com **embedded-computing.com/p9911755**

Tag-Connect

TC2030-CTX Plug-of-Nails™ JTAG Spring-Pin Cable for ARM Cortex

Tag-Connect's TC2030-CTX family of Plug-of-Nails™ Spring-Pin cables is designed for use with ARM Cortex Debuggers such as Keil's ULINK2. The patented Spring-Pin Connector plugs directly to a tiny footprint of pads and holes in the PCB thus eliminating all of the cost and most of the space of the traditional JTAG header. The PCB footprint can be about the size of an 0805 resistor (under one third of the PCB space required by the 10-pin Cortex 50-mil header). Since Tag-Connect footprints have no vertical height they are also ideal for use in low-profile or space-constrained products. As designs continue to shrink, Tag-Connect has tackled the problems of connector size and cost head-on and delivered a "No Header – No Brainer" solution!

www.tag-connect.com **embedded-computing.com/p9911919**

TE Connectivity

MULTIGIG RT 2-R Connector

TE Connectivity's MULTIGIG RT 2-R connector is a rugged, lightweight and high-speed board-to-board interconnect that is compliant to VITA 46, including backward compatibility with all existing VITA 46 daughter cards, and supports OpenVPX architecture. The connector system features the modularity and flexibility of the field-proven MULTIGIG RT 2 connector products, but with a new quad-redundant contact structure designed for extreme vibration levels. The MULTIGIG RT 2-R connectors have been tested to HALT vibration levels per VITA 72, and the robust "pinless" interface has been tested up to 10,000 mating/unmating cycles.

www.te.com/multigig **embedded-computing.com/p9911827**

Telco Systems

T-ATCA 404

Telco Systems' T-ATCA404 ATCA switching blade delivers unparalleled connectivity to ATCA platforms with support of multi-rate 10/40GE interfaces on the FI switch and 1/10GE interfaces on the BI switch. T-ATCA404 is one of the densest 40/10GE ATCA based solutions in the market – the only one that provides such an extensive feature set including OAM, Load balance, HQoS, MPLS/Ethernet enhanced security, synchronization and other L2-L4 features.

www.telco.com **embedded-computing.com/p9910362**

Texas Instruments — **BeagleBone Black**

A ready-to-use, 1 GHz computer for only $45. This credit-card-sized Linux computer is an open hardware and software development platform that makes it quick and easy to transform great ideas into products. BeagleBone Black allows developers to leverage the ideas and knowledge of the highly active and engaged users of the BeagleBoard.org community who support each other from concept through development. Opportunities for innovation are endless.

www.beagleboard.org — **embedded-computing.com/p9911590**

Texas Instruments — **KeyStone-based multicore SoCs**

On Nov. 13, 2012, Texas Instruments introduced six new KeyStone-based multicore SoCs. These SoCs were the industry's first implementation of quad-ARM® Cortex™-A15 MPCore™ processors in infrastructure systems, giving developers twice the capacity and performance at half the power relative to existing solutions for networking, HPC, gaming and media processing applications. The unmatched combination of ARM Cortex-A15 processors, C66x DSPs, security processing and Ethernet switching transforms the real-time cloud into an optimized high-performance, power-efficient processing platform. These processors, as part of HP's Project Moonshot for example, will give customers the performance, scalability and programmability needed for a variety of applications.

www.ti.com/Multicore — **embedded-computing.com/p9911845**

Toradex — **Apalis T30 ARM Computer Module**

Apalis T30 with NVIDIA® Tegra® 3 comes with a wide range of new features. Apalis supports a variety of industry standard interfaces, while at the same time providing advanced multimedia and high-speed connectivity making it suitable for an almost unlimited number of applications. Apalis brings to market various new technologies. Direct Breakout™ considerably simplifies routing of high speed signals on the carrier board. The module also comes with support for ready-made passive cooling solutions and industrial temperature support for the most demanding applications. Apalis T30 comes with pre-installed operating system and includes a Windows Embedded Compact 7 Core Runtime License. The available operating systems are Android, Linux and Windows Embedded Compact 7.

www.toradex.com — **embedded-computing.com/p9911828**

Wind River — **Wind River Intelligent Network Platform**

Wind River Intelligent Network Platform is a software platform for the development of network equipment that can accelerate and secure ever increasing traffic. It includes software engines delivering fast packet acceleration and DPI capabilities for greater network intelligence while maintaining high performance, up to 1100% improvement in IP forwarding, up to 500% improvement in throughput for UDP and performance for TCP. For customers wishing to avoid modifying any applications, the platform also includes patent-pending technology delivering up to 300% performance boost, ideal for massive traffic or high volume transactions. The platform can help customers build faster, smarter, and more secure networks and realize significant time and cost savings.

www.windriver.com — **embedded-computing.com/p9911924**

Zilog — **ZGATE™ Embedded Security Firewall**

The ZGATE™ Embedded Security Firewall combines Zilog's eZ80F91 MCU and full-featured ZTP Embedded Internet Software Suite and TCP/IP Stack with a world-class embedded firewall. This highly configurable firewall protects the ZTP networking layers from attack by discarding suspicious packets before they reach ZTP and your embedded application. The ZGATE Firewall includes a static packet filtering engine that filters packets according to user-defined configuration rules and a stateful packet inspection engine that can automatically filter suspicious packets based on unusual activity. Additionally, select ZGATE products include threshold-filtering mechanisms that can minimize the affect of packet floods. Three packages are available: Standard, Extended, and Premium.

www.zilog.com — **embedded-computing.com/p9911858**

Editor's Choice

OpenVPX module features user-programmable FPGA

A new 3U OpenVPX module from Creative Electronic Systems incorporates an FPGA enabling significant performance and I/O customization by the user. The FIOV-2310 allows the flexible configuration of the high-speed links between the FPGA Mezzanine Card (FMC) connector and various VPX backplane profiles, making it suitable for a variety of complex applications. PCI Express support is built-in, while the internal Xilinx Kintex-7 FPGA makes it possible to program the backplane interconnect to any of the popular serial fabric protocols including Serial RapidIO, Serial Front Panel Data Port (SFPDP), and 10 GbE. Onboard clock generators are also provided for precision timing applications, while clock inputs and outputs are available for external synchronization and to improve signal integrity. The FIOV-2310 module has options for both air-cooled and conduction-cooled operating environments. Software support packages are available for INTEGRITY, Linux, and VxWorks operating systems.

Creative Electronic Systems | www.ces.ch | www.embedded-computing.com/p9911883

Touch-screen controllers track gloved fingers

To simplify the interaction with portable electronic systems and smartphones in a variety of conditions, Cypress Semiconductor Corporation recently announced new features for the TrueTouch touch-screen controller family. The Gen4 family enables accurate navigation of capacitive touch screens even with thick gloves on, along with new functionality including passive stylus support to capture handwriting and signatures accurately. In addition, proximity detection senses an object approaching the touch screen and powers down the display before contact to avoid accidental touches. The Gen4 TMA467 controller tracks up to 10 fingers in thin gloves and 2 fingers with thicker gloves, such as those used by skiers. The Gen4 TMA468 controller supports passive styli with tips as small as 2 mm, enabling touch screens to capture characters as small as 7 mm. This capability is important for writing in languages that require enhanced character recognition for reliable text input, such as Chinese and Japanese. The Gen4 family can control capacitive touch screens up to 10.1 inches.

Cypress Semiconductor | www.cypress.com | www.embedded-computing.com/p9912445

Compact vision systems target embedded applications

Two new machine vision systems designed for integration into Original Equipment Manufacturer (OEM) devices and equipment were recently released by Cognex Corporation. The Advantage 100 vision system is a compact (23 mm x 44 mm x 54 mm), enclosed vision system with customizable optics, lighting, and Ethernet communications suitable for integration into large automated clinical diagnostic systems. The AE2 Advantage vision engine provides an even smaller (14.5 mm x 29.1 mm x 20.5 mm) form factor designed for integration inside embedded devices. The interface to the vision engine is via a 31-pin ZIF connector allowing access to all communication, power, and I/O signals. A standard Cognex I/O board is also available with an RJ-45 connection for Ethernet and a cable with DB-15 connector for power, I/O, and RS-232 communication. The built-in optics and illumination configuration is modular, offering designers a wide range of field of view and working distance alternatives. The Advantage 100 and AE2 Advantage engines include Cognex application software to test and develop the vision and identification tools needed for specific requirements in space-constrained and cost-sensitive applications such as those in the clinical diagnostic, printing, kiosk, or medical device industries.

Cognex Corporation | www.cognex.com | www.embedded-computing.com/p9912446

E-community Post

Joining the embedded conversation

By Sharon Hess

www.embedded-computing.com

BLOGS | MARKET STATS | INNOVATION | VIDEOS | SOCIAL MEDIA

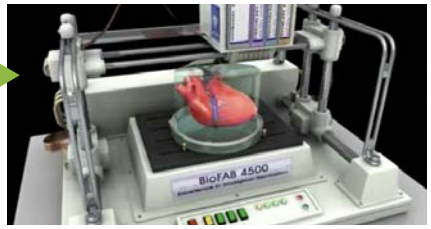

'Will 3D printing change the world?'

Will waiting for an organ transplant donor really become a thing of the past 10 to 20 years from now as the red hot 3D printing trend forges on? The new "Will 3D printing change the world?" video produced by the Off Book arm of PBS explores the disruptive technology, shedding light on a possible future where 3D blood vessels or skin tissues could be printed from living cells, where healthy replacement tissues for the heart could be printed, or entire organs could be printed using any individual patient's genetic makeup. 3D printing techniques from Japan aim to eliminate old-school black-and-white sonogram photos of a baby in-utero, by instead printing a 3D model of the mother's torso that shows the child's actual form. Countless other industries will also be affected by 3D printing, and MIT scientists are even working to develop a 3D printer capable of printing food. But in the meantime, issues such as 3D printing's effect on copyright and patent laws, economic scarcity, manufacturing, and the rules of consumer and brand engagement are evolving.

Watch the video: http://opsy.st/19XlONR

E-cast

Embedded industry upgrades M2M networks

Machine-to-Machine (M2M) communications strategies and cloud computing are transforming industrial interconnects from an assortment of fragmented, proprietary technologies to open standards easily integrated into new designs. In this webcast you will learn how M2M technology allows embedded design teams to contain costs, improve security, enable remote management, and maximize system availability.

Presented by: Eurotech, ThingWorx

Watch it on demand any time:
http://opensystemsmedia.com/events/e-cast/archive

$26 billion in mHealth app services predicted by 2017

As the smartphone/mobile connected device app trend continues to explode, the mHealth (mobile health) industry is following suit, with mHealth app services anticipated to climb to $26 billion worldwide by 2017, according to research2guidance's recent "Global Mobile Health Market Report 2013-2017." Top mHealth app creators garner more than 300,000 paid and 3 million free iOS app downloads within the United States. Other countries' and other platforms' statistics differ but also indicate that mHealth apps are on the rise, the company says.

Fifteen percent of mHealth apps are geared toward healthcare professionals, such as the Medscape app, pictured, published by WebMD and utilized by more than 3 million healthcare professionals using iPads, iPhones, iPod touch, and Android-based devices. Medical apps for healthcare pros typically include those for healthcare management, remote monitoring, and Continued Medical Education (CME).

Additionally, 42 percent of the 97,000 mHealth apps residing in major app stores currently require payment for downloading. However, this percentage is slated to drop to 9 percent within the next five years, when 84 percent of the mHealth application sector's income will result from products like sensors and mHealth-related services.

Roving Reporter blog: Hardening infrastructure against attack with Intel vPro and Intel Intelligent Systems Framework

By Warren Webb

As industrial designers incorporate remote, fully interconnected factory equipment to reduce costs and simplify maintenance, the exposure to data disruptions or system cyber attacks becomes apparent. According to the 2013 Internet Security Threat Report published by Symantec Corporation, manufacturing was the most-targeted sector in 2012 with 24 percent of all attacks, compared with 15 percent in 2011. Consequently, designers are devising techniques to protect vulnerable elements. The Intel Intelligent Systems Framework can be used to simplify connecting, managing, and securing embedded devices by leveraging system processors with Intel vPro.

Read the blog: http://opsy.st/19pB1L5

Atom™ Powered SBCs
High-Performance, Small and Fanless

For your next design, select rugged WinSystems' single board computers powered with single- or dual-core Intel® Atom™ processors. Our Industry Standards-based SBCs have a wealth of onboard I/O, plus expansion capabilities.

- Long-life Intel® Atom™ CPUs
- Simultaneous VGA and LVDS Video
- Gigabit Ethernet Port(s)
- Eight USB 2.0 Ports
- Four Serial Ports
- SATA and CompactFlash Interface
- Digital I/O with Event Sense
- -40°C to +85°C Operation
- Outstanding Technical Support
- Industry Standard Platforms
  - PC/104 – 3.6 x 3.8 inches
  - PC/104-*Plus* – 3.6 x 3.8 inches
  - SUMIT-ISM™ – 3.6 x 3.8 inches
  - EPIC – 4.5 x 6.5 inches
  - EBX – 5.75 x 8.00 inches
- Software Support: Windows®, Linux, and x86-compatible RTOS

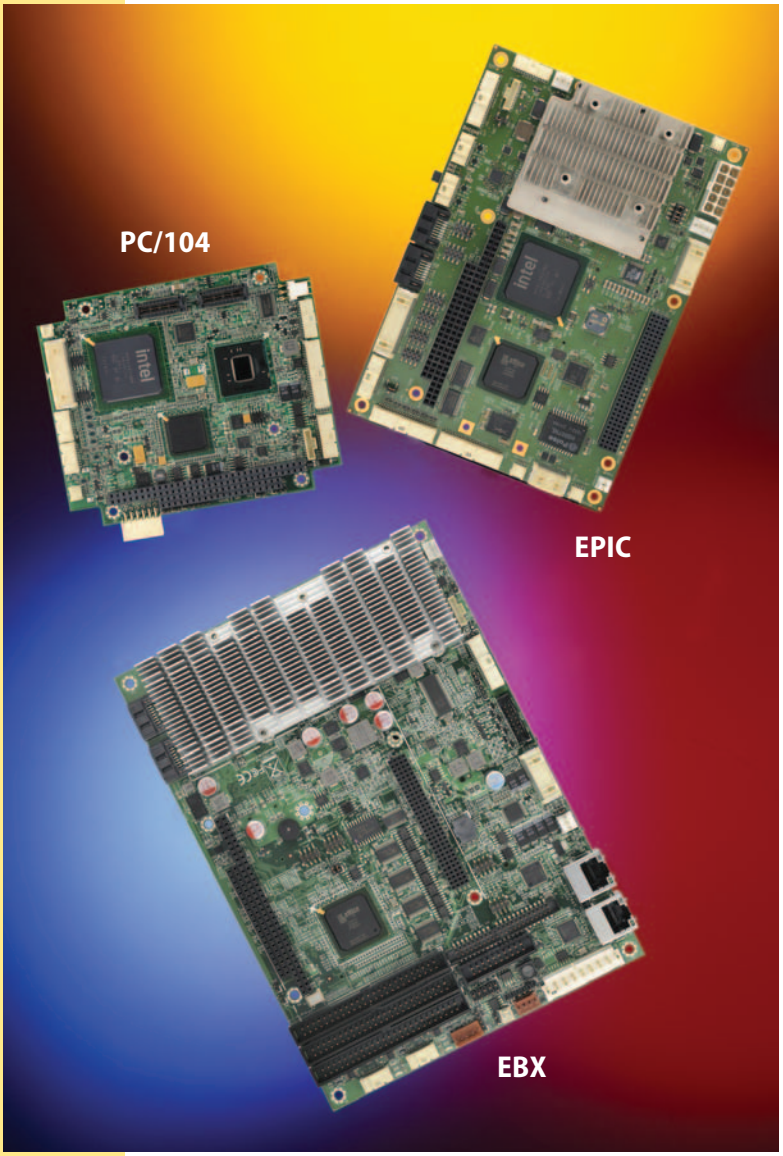

Our SBCs are the right choice for industrial, pipeline, communications, transportation, medical, instrumentation, and MIL/COTS applications.

Go to WinSystems' SBC Selection Guide at **www.WinSystems.com/SBCsPC**

Ask about our evaluation program.

Call 817-274-7553 or Visit WinSystems.com/AtomPC

WinSystems®
An Employee-Owned Company

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

Scan this tag to read more about our SBCs.