Embedded COMPUTING DESIGN

www.embedded-computing.com

@embedded_mag
Embedded Computing Design

NOVEMBER 2010
VOLUME 8 • NUMBER 7

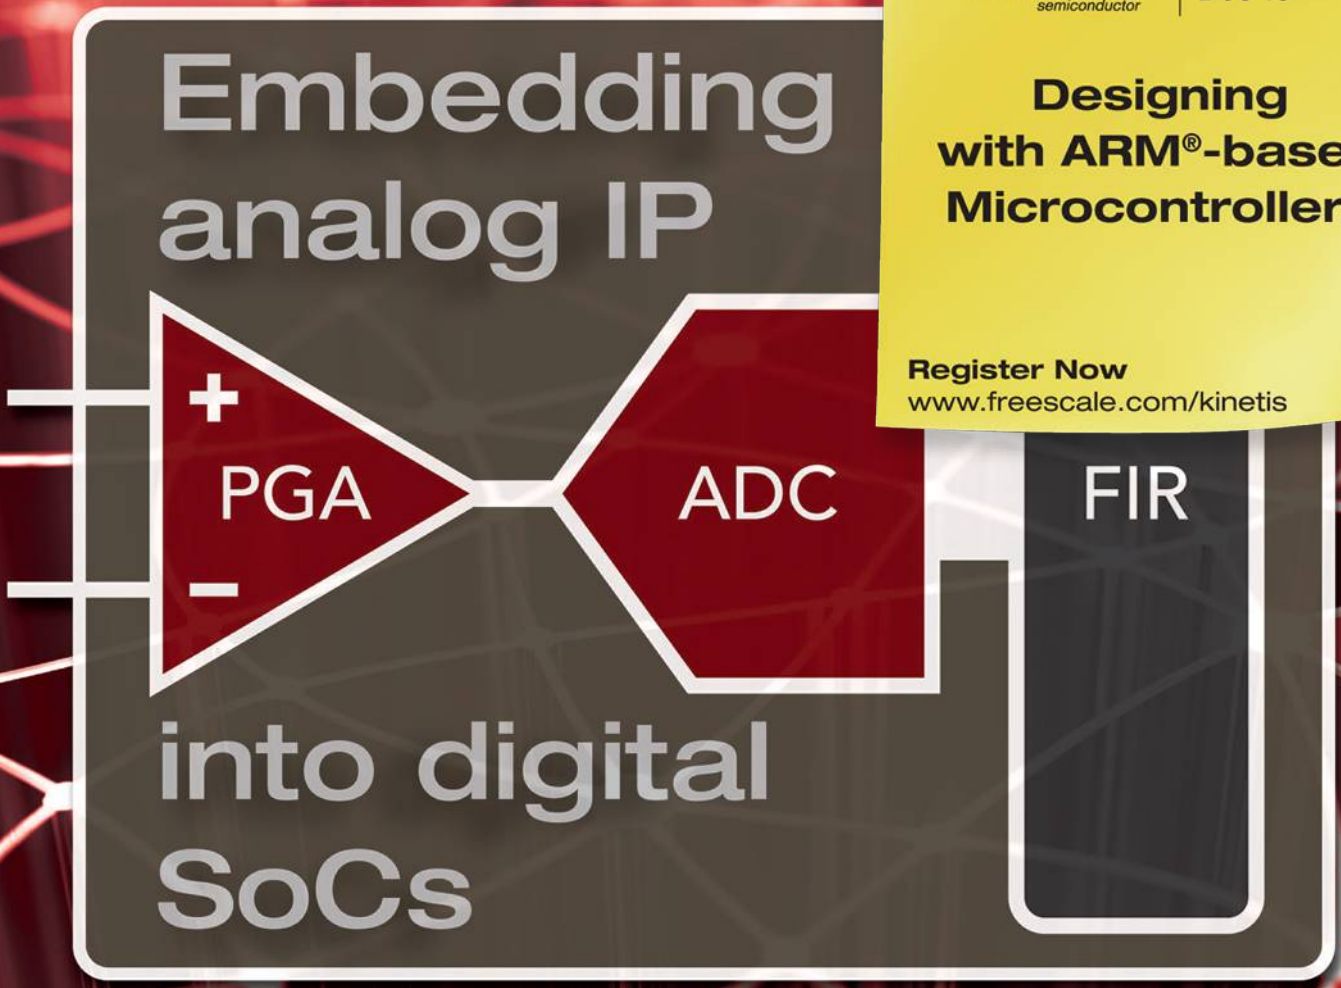

PC/104-*Plus* SBC Supports Industrial Control, HMI, Networking and I/O

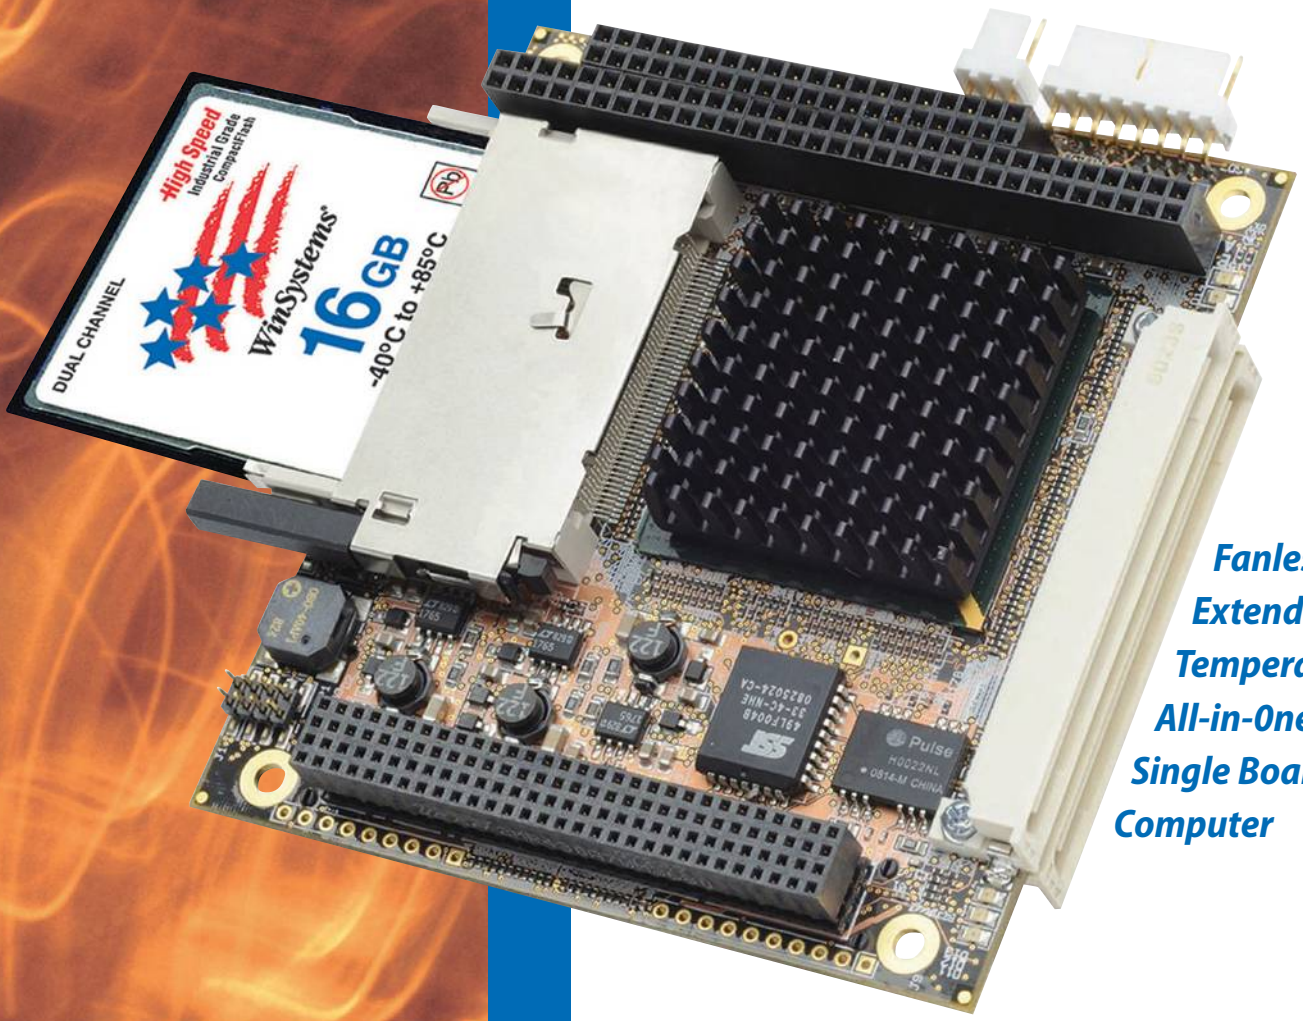

Fanless Extended Temperature All-in-One Single Board Computer

The PPM-LX800 SBC is available on a 6.5-inch Open Frame Panel PC.

If your application needs a compact, industrial PC with fanless -40°C to +85°C operation, check out WinSystems' PPM-LX800.

Designed for embedded, space-limited, and low-power applications requiring reliability and long-term availability, the PPM-LX800 runs Linux, Windows® XP embedded, and other x86 real-time operating systems. Plus it comes with a wealth of on-board I/O functions and free software drivers.

Features include:

- AMD LX800; x86-compatible CPU
- Compact size, 3.6" x 3.8" (90mm x 96mm)
- Socket for a bootable CompactFlash card
- High-resolution video controller provides
  - CRT resolutions up to 1920 x 1440
  - Panel resolutions up to 1600 x 1200
- Custom splash screen on start up
- 10/100 Mbps Intel Ethernet controller
- Four serial ports with FIFO
- Two USB 2.0 ports
- 16 Lines of digital I/O
- PS/2 keyboard and mouse controller plus LPT port
- +5V only operation
- Ultra DMA100 EIDE controller for 1 or 2 devices
- PC/104 and PC/104-*Plus* expansion connectors
- AC97 Audio with MIC, Line In, and Line Out
- -40°C to +85°C extended operational temperature
- Long-term product availability
- Responsive and knowledgeable technical support
- Quick Start Kits offered for easy software development

Contact our factory application engineers for additional product information, custom configurations, and pricing.

Call 817-274-7553 or
Visit www.winsystems.com/PPM-LX800

Ask about our 30-day product evaluation

WinSystems®
An Employee-Owned Company

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

Linux

Extended Temperature

RoHS compliant

Embedded COMPUTING DESIGN
www.embedded-computing.com
November 2010 | Volume 8 • Number 7

On the cover

Programmable gain amplifiers, analog/digital converters, and more mixed-signal circuits are moving onto the digital SoC, joining finite impulse response filters and other digital signal processing technologies. Choosing and integrating analog IP is the subject Synopsys explores starting on page 20 as part of our look at what's next in EDA.

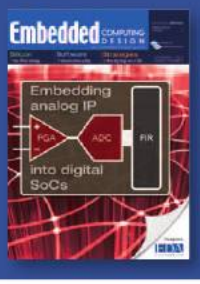

ENVIROINK

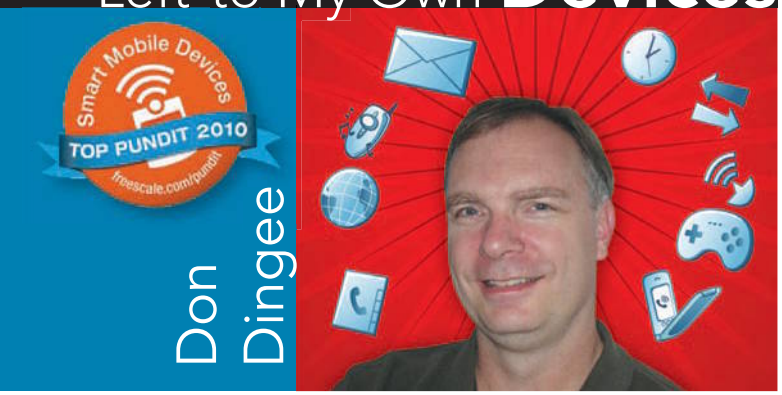

Mixed multicore messages

By Don Dingee

How do you see the multicore maze right now? Your experience can make all the difference in how you approach solutions and where your application goes next.

I opened my recent presentation at the Wind River Multicore Regional Conference with the idea that multicore is like a maze – it's easy to get into, but harder to get out with the results you want unless development is viewed differently.

Multicore hardware is plentiful and multicore operating systems are catching up, but application developers are just getting started. VDC just published a September 2010 report indicating that nearly 65 percent of engineers have less than one year of experience deploying software over multicore platforms. The questions of what multicore is capable of, what it will do for my application, and how my development has to change are on a lot of developers' minds lately.

For the latest ideas, visit

http://channels.opensystemsmedia.com/Multicore

There is a huge variety of multicore architectures with compute cores, graphics cores, DSP cores, networking cores, and I/O cores, and in many cases software designers don't actually know what's under the hood. That idea prompted an interesting audience remark: "We've found that it's easier to just start over." Taking that lump of existing C code and trying to spread it out over a new multicore architecture isn't the fastest path in some instances.

Another audience member posed this question: "What kind of parallel compiler technology is available?" This query is an indication that people tend to think of symmetric multiprocessing first. There are applications that benefit from technology like OpenMP and parallel programming. But often the application either has subsets or regions of intensity that lend to breaking things up instead of just spreading them out equally across cores.

Case in point: Wind River presenters said that for packet acceleration, they're discovering that two groups of four cores each running a suite of tasks significantly outperform one group of eight cores running the complete set. That's true even in an advanced multicore architecture with fabric connections between the cores and a hypervisor layer spanning independent guest operating systems running on each core.

ECD in 2D:

The Wind River Network Acceleration Platform leverages asymmetrical multiprocessing for ultra-fast packet throughput. Use your smartphone, scan this code, watch a video: http://bit.ly/9x77Ua.

Our lunch table challenged me for saying this concept is as big as the changeover from digital logic to microprocessors: "This isn't new and it's not that big a transition." They thought they had seen all these things before a couple of times, except the problem used to be spread across multiple boards instead of across cores within a single processor. I believe there are some parallels, but the world of pipelined microarchitecture, multilevel cache, fabric interconnect, and intelligent peripherals is very different from the world of shared parallel buses, register polling, and mailboxes.

Visibility inside multicore is another big issue – both getting in and keeping out. Wind River officials said they're seeing a resurgence in JTAG analyzer use cases, and Freescale representatives described their efforts in QorIQ trusted platform technology to make it much more difficult to see inside once an application is fielded.

The value-add on multicore platforms will move into optimizing applications and securing the IP. It'll take some relearning for developers to get the most out of the technology.

What are your thoughts on the multicore state of things? Are operating systems caught up to multicore processors? Is application development and optimization really easy or really difficult or somewhere in between? What's changing in your development practices for multicore? Share your ideas with us.

twitter @dondingee, @embedded_mag

ddingee@opensystemsmedia.com

Up next:

Consumer Electronics Show
January 6-9, 2011
Las Vegas, NV
www.cesweb.org

Reader resources

Subscribe to magazines and E-letters:
www.opensystemsmedia.com/subscriptions

E-casts:
ecast.opensystemsmedia.com

White Papers:
whitepapers.opensystemsmedia.com

Videos:
videos.opensystemsmedia.com

Get news via RSS:
www.embedded-computing.com/news/rss.xml

Advertiser resources

Submit information:
submit.opensystemsmedia.com

Media kits:
www.opensystemsmedia.com/publications

Join the discussion

Comment on blogs, articles, news, and more
Twitter: @embedded_mag
LinkedIn: www.linkedin.com/groups?gid=1802681

Advertiser Information

Page #	Advertiser	Ad title
11	American Portwell Technology	Your trusted provider for embedded applications
27	Annapolis Micro Systems, Inc.	High performance signal and data processing
26	Arbor Technology	Patient Infotainment & eHealthcare Station
19	CoreSolid Storage Solutions	XDOM mini PCI Express DiskOnModule
8	IEI Technology USA Corp.	Fanless Intel Atom D525 Processor Family
28	Kontron	Next generation Intel Atom E6xx based scalable platforms
26	Sanmina-SCI	Viking modular solutions
21	Sealevel Systems, Inc.	The right connection creates incredible power
24	Synopsys, Inc.	Virtual prototyping quad-core ARM with Android
9	Technologic Systems	TS-WIFIBOX-2
18	Themis Computer	Themis' rugged servers have speed to burn
13	Tri-M Systems Inc.	PC/104 Can-Tainer
25	Tri-M Systems Inc.	100 MHz PC/104 module
14, 15	Wind River Systems, Inc.	Innovation isn't making something new
2, 3	WinSystems, Inc.	PC/104-*Plus* SBC supports industrial control

Embedded COMPUTING DESIGN

OpenSystems media.

CompactPCI Advanced TCA Micro Systems | DSP-FPGA.com | INDUSTRIAL EMBEDDED SYSTEMS | Military EMBEDDED SYSTEMS | PC/104 and small form factors | VME and Critical Systems

Embedded Group

Don Dingee, Editorial Director
ddingee@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

Monique DeVoe, Web Content Editor

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Matt Jones, Web Developer

David Diomede, Art Director

Steph Sweet, Creative Director

Joann Toth, Senior Designer

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Reprints and PDFs

Nan Holliday
800-259-0470
republish@opensystemsmedia.com

Sales Group

Dennis Doyle, Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker, Strategic Account Manager
rbarker@opensystemsmedia.com

Bridget Hildebrand,
Strategic Account Manager
bhildebrand@opensystemsmedia.com

Christine Long, Digital Content Manager
clong@opensystemsmedia.com

International Sales
Dan Aronovic, Account Manager – Israel
daronovic@opensystemsmedia.com

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

Regional Sales Managers
Barbara Quinlan, Midwest/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

Editorial/Business Office

16626 E. Avenue of the Fountains, Ste. 203
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystemsmedia.com

Publishers: John Black, Michael Hopper, Wayne Kristoff
Vice President Editorial: Rosemary Kristoff
Vice President Marketing & Sales: Patrick Hopper
phopper@opensystemsmedia.com
Business Manager: Karen Layman

Squeezing the last penny of efficiency from embedded designs

By Kimkinyona Fox and Raviraj Mahatme

Image courtesy of Medtronic, Inc.

Not every processor is built on cutting-edge processes – many processors are being designed around mature nodes to reduce risk and cost. As with any embedded System-on-Chip, the challenge is to become more power-efficient. Physical IP can help extend battery life by reducing dynamic and leakage power.

The embedded microcontroller and mixed-signal System-on-Chip (SoC) markets have shown strong growth while manufacturing on mainstream technology process nodes. At almost $30 billion, the mixed-signal market constitutes about 10 percent of the $300 billion semiconductor industry, and this growth will continue with the advent of new consumer applications. For example, touch-screen technology and cell phones existed independently for many years. The emergence of better human user interfaces in the form of smartphones has established a market that did not exist three years ago.

Mixed-signal circuitry is what enables our analog "real" world to interact with its electronic digital counterparts. As embedded applications continue to reach higher volumes and more mobile implementations, designers are confronting the issues of lowering costs and reducing power.

A challenge in embedded and mixed-signal designs is the cost-sensitive nature of high-volume design. The lower the product price, the greater the number of potential applications that can incorporate that technology. One semiconductor cost-reduction approach is fabrication in small process geometries. But in the mixed-signal market, there are several reasons why developers stay on mainstream nodes:

- Analog components do not scale as well as digital components on a semiconductor process. As a result, the density benefits of scaling are much less than those obtained by digital components.
- Most mixed-signal devices do not have the high-speed requirement of digital devices. Many mixed-signal devices operate at speeds of less than 100 MHz, which eliminates the need to quickly move to the next leading process node.
- Older process technologies are stable with well-understood analog characteristics. The low risk and low cost make manufacturing at the older technology nodes more compelling.

Embedded SoC designers must look for innovative solutions to reduce die size while remaining on mainstream process technologies.

Another challenge for embedded designers is power reduction. At 250 nm, foundries are not investing in the process innovations necessary to reduce dynamic and leakage power. Some applications need additional power optimization to increase battery life to 2-5 days, as is the case with smartphones. Some medical applications such as pacemakers need low-power solutions to enable batteries to last 5-10 or more years. Nano- or pico-amp usage can minimize extensive surgeries and thereby increase patient quality of life.

Whether an application is on an extreme of the battery life continuum or somewhere in the middle, the one constant is that users need more functionality on a single charge. The embedded designer's challenge is to increase functionality while staying at a low risk – on a mature process node – and decrease area and power for added feature and battery life differentiation.

Feature innovation with embedded processors

Mixed-signal SoCs either incorporate a low-power microcontroller or work with a stand-alone microcontroller to execute commands from the user and address interrupts and provide readings of incoming data. ARM offers a family of modern 32-bit processors, the Cortex-M series, specifically designed for low-power microcontrollers.

To address the requirement for additional features, these processors offer more performance at a given footprint via performance efficiency – the ability to complete tasks faster and therefore reduce activity periods. The performance

advantage stems from Cortex-M performing single-cycle 32-bit arithmetic and logic operations (including single-cycle 32-bit multiplication) and performing 8-, 16-, or 32-bit data transfers with indexed addressing in a single instruction. This dramatically reduces the processor clock frequency required and increases the performance from a single instruction. Furthermore, it reduces the memory required for program storage and the power needed to fetch programs from memory.

The reduced clock frequencies mean lower-noise and higher-precision analog, boosting the device's analog sensor capability. RF applications also benefit from the reduction of electromagnetic interference. This efficiency increases performance and provides low-noise, high-precision analog operation ideally suited for mixed-signal applications.

Reducing manufacturing risk and silicon area

Increased functionality from an efficient processor design is a great start, but embedded designers also need to decrease SoC area and retain a low-cost, low-risk process node. To meet this need, the increasing trend is to migrate to the mainstream 180 nm node. Some of the most aggressive designs are even targeting the 110 nm technology node to shrink die size. Several factors are enabling this trend:

- Twelve-year-old 180 nm technology is stable enough that there is little to no risk in migrating from 250 nm technology to 180 nm.
- The availability of value-adding nonvolatile memory components such as flash and one-time programmable memory is not offered at the 250 nm node.
- Emerging niche process technologies like 180 Ultra-Low Leakage (ULL), 180 Bipolar CMOS DMOS (BCD), and 110 ULL are perfectly suited for the embedded market with their low dynamic and leakage power profiles.

Optimized Physical IP for reduced area and power

Area reduction leads directly to decreased die cost, and combined with lower power, can enable less expensive SoC packaging and cut overall system costs. Furthermore, lower dynamic and leakage power extends battery life. ARM enables reduced die size and power with Physical IP platforms implemented on ULL processes.

The ULL Physical IP platform comprises a range of logic products: a nine-track SC9 High-Density (HD) standard cell library, a tapless seven-track SC7 Ultra-High-Density (UHD) standard cell library, and a seven-track SC7 UHD Power Management Kit (PMK). The platform contains a full range of memory compilers including HD SRAMs, register files, and ROMs. The SC7 UHD library

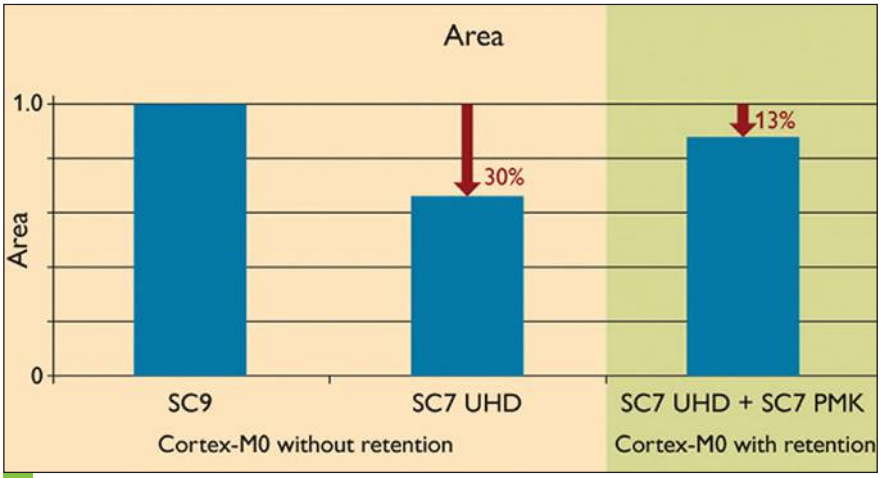

Figure 1 | The SC7 Ultra-High-Density library saves up to 30 percent of the area on the ARM Cortex-M0.

typically delivers up to 30 percent area savings compared to the SC9 HD library (see Figure 1).

The SC7 UHD library can be paired with the complementary SC7 UHD PMK, which can, for a small logic area increase, result in significant leakage savings. For example, when implementing the Cortex-M0, the leakage using SC7 is 12x lower compared to that of the SC9. When implementing with the SC7 UHD library and SC7 UHD PMK, the Cortex-M0 has up to 50x reduction in leakage (see Figure 2).

The future of embedded and mixed-signal design

The embedded and mixed-signal markets will continue to grow with added innovations to increase features and differentiation while reducing costs. Efficient embedded processors and Physical IP will aid in the growth of these markets by providing unparalleled design and efficiency in tomorrow's embedded device solutions.

ARM 32-bit processors bring added functionality and performance efficiency. Physical IP implemented on smaller, yet still low-risk process nodes will keep manufacturing costs low and reduce die size. Optimized solutions such as ARM's Physical IP bring dramatic area and leakage savings for lower total system costs and extended battery life. ECD

Kimkinyona Fox *is a senior platform marketing manager at ARM, Inc. Her experience includes senior-level marketing and engineering positions at Rambus, Cypress, PLX Technology, and C-Cube Microsystems. Kimkinyona has authored multiple articles on high-speed interfaces and system design. She earned her BSEE from California Polytechnic State University, San Luis Obispo.*

Raviraj Mahatme *is a platform marketing manager with the Physical IP Division at ARM. He started his career with ARM as a circuit design engineer and has worked in multiple ARM ofáces around the world. He holds an MSCE from North Carolina State University, Raleigh.*

ARM
kimkinyona.fox@arm.com
ravi.mahatme@arm.com
www.arm.com

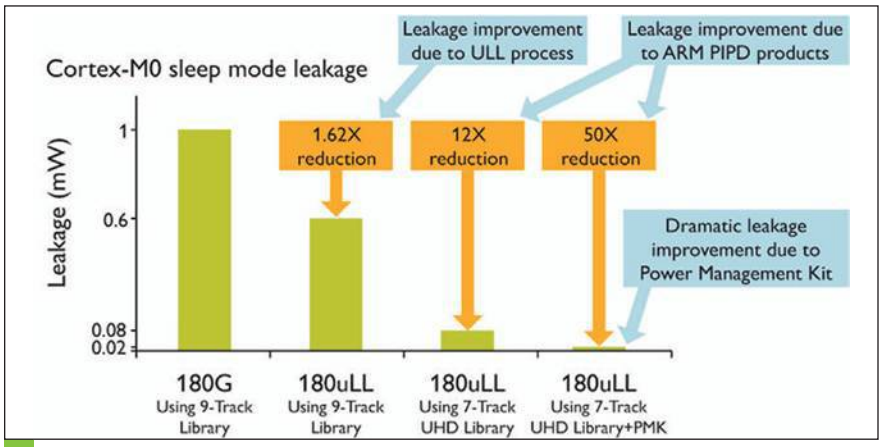

Figure 2 | Implemented with the SC7 Physical IP, the ARM Cortex-M has up to 50x reduction in sleep mode leakage.

ECD in 2D:

ARM CEO Warren East discusses the company's collaboration with Freescale on Kinetis microcontrollers based on ARM Cortex-M4. Use your smartphone, scan this code, watch a video: http://bit.ly/aKSMgK.

Instant replay for multicore systems

By John A. Carbone

Real-time event analysis is critical for multicore success, but gaining visibility into these types of events is difficult using traditional debugging tools. A new type of tool specifically made to facilitate multicore analysis provides new capabilities for seeing what's really going on.

Real-time systems must react quickly to external and internal demands. When a system uses a multicore architecture, the speed and number of interactions rise sharply. While this improves system performance, it complicates real-time sequencing of application events, given that multicore system events can occur simultaneously over multiple independent processors instead of sequentially over a single processor.

For the multicore developer, the increased complexity of managing the number of events and their simultaneous nature represents an exponentially more challenging system to design. Diagnosing the cause of a system failure or inefficiency is much more difficult with a multicore system than it is with a single-processor system. With few multicore-ready tools available, developers are left with primitive print statement techniques that leave bread crumbs throughout the system's operation indicating data about various events that have occurred. The developer must gather and make sense of the crumbs and infer the system's state, a process that often requires subsequent re-instrumentation to gain a finer degree of granularity and a repeat of the process.

To efficiently unravel the intricate sequence of operations on a multicore system, developers need an instant replay enabling them to examine the system's operations that immediately precede an area of interest. As shown in Figure 1, a new type of debugging tool shows exactly what is going on in a multicore system across a particular period of time.

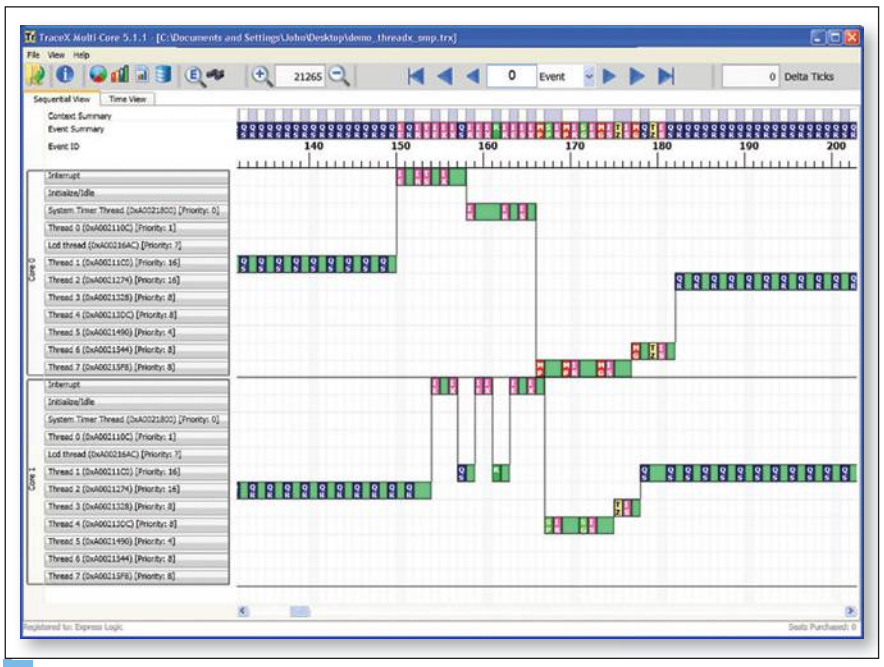

Figure 1 | TraceX offers a graphical view of real-time events in a multicore system. In this example, Core-0 and Core-1 are seen simultaneously executing different threads.

A graphical analysis of all system events is displayed across a single timescale organized by application thread and grouped by processor core.

Traditional approach to system-event analysis

Real-time programmers have long understood the importance of system behavior to the functionality and performance of their applications. The conventional approach addresses these issues by generating data on system behavior when the code reaches a certain stage by toggling an I/O pin, using printf, setting a variable, or writing a value to a file.

Inserting such responses requires a substantial amount of time, especially considering that the instrumentation code often doesn't work exactly as expected the first time around and must be debugged. Once that part of the application is verified, the instrumentation code needs to be removed, and its removal needs to be debugged. Most of the instrumentation process is manual and thus time-consuming and prone to additional errors.

Besides instrumenting the code, the developer also needs to find a way to interpret the data generated. The volume of information generated by the instrumentation code complicates the task of determining what system events took place in what sequence.

New approach offers advantages

In contrast to the conventional method, TraceX automatically analyzes and graphically depicts system and application events captured on the target system during runtime. Events such as thread context switches, preemptions, suspensions, terminations, and system interrupts each leave a bread crumb that the debugging tool recognizes and displays. These bread crumbs describe what event just happened, which thread was involved, which core that thread was running on, when it occurred, and other relevant information.

With this tool, the user can log any desired application events using an Application Programming Interface (API). Event information is stored (logged) in a circular buffer on the target system with buffer size determined by the application. A circular buffer enables the most recent "n" events to be stored at all times and available for inspection in the case of a system malfunction or other significant event.

Good multicore debugging tools allow event logging to be stopped and started dynamically by the application program at a specific time, such as when an area of interest is encountered. This avoids cluttering the database and consuming target memory when the system is performing correctly. The event log can be uploaded to the host for analysis when encountering a breakpoint or system crash or after the application has finished running.

Once the event log is uploaded from target memory to the host, the tool displays the events graphically on the horizontal axis,

which represents time (refer to Figure 1). The various application threads and system routines related to events are listed along the vertical axis, and the events themselves appear in the appropriate row. For multicore systems, the events are linked to their respective processor core and grouped together so that developers can easily see all the events for a core.

All events are also presented in the top summary row, regardless of core or thread, giving developers a handy way to obtain a complete picture of system events without scrolling down through all threads and cores. Events are represented by color-coded icons located at the point of occurrence along the horizontal timeline as well as to the right of the relevant thread or system routine. The axes can be expanded to show more event detail or collapsed to show more events. The timescale can be panned left (back) or right (ahead) to show any point in the trace buffer. When an individual event is selected, as shown in Figure 2, detailed information is provided for that event, including the core, context, event, thread pointer, new state, stack pointer, and next thread point.

Solving priority inversion problems

One of the most challenging real-time problems is priority inversions. Priority inversions arise because Real-Time Operating Systems (RTOSs) employ a priority-based preemptive scheduler to ensure the highest-priority thread that is ready to run actually runs. The scheduler can preempt a lower-priority thread in mid-execution to meet this objective.

Problems can occur when high- and low-priority threads share resources, such as a memory buffer. If the lower-priority thread is using the shared resource when the higher-priority thread is ready to run, the higher-priority thread must wait for the lower-priority thread to finish. If the higher-priority thread must meet a critical deadline, then the maximum time it might have to wait for all its shared resources must be calculated to determine its worst-case performance. Priority inversions occur when a high-priority thread is forced to wait while the CPU serves a lower-priority thread.

Priority inversions are difficult to identify and correct. Their symptom is normally poor performance, but poor performance stems from many potential causes. Compounding the challenge of identifying the cause is the fact that priority inversion can evade testing, which could mean the inversion is non-deterministic.

A systems event tool like TraceX makes it possible to easily and automatically identify priority inversions. The trace buffer clearly identifies which thread is running at any point in time and records any change in a thread's readiness. Thus, it is easy to go back in time to determine if a higher-level priority thread is ready to run but blocked by a lower-priority thread that holds a resource needed by the higher-priority thread. Figure 3 shows non-deterministic priority inversion.

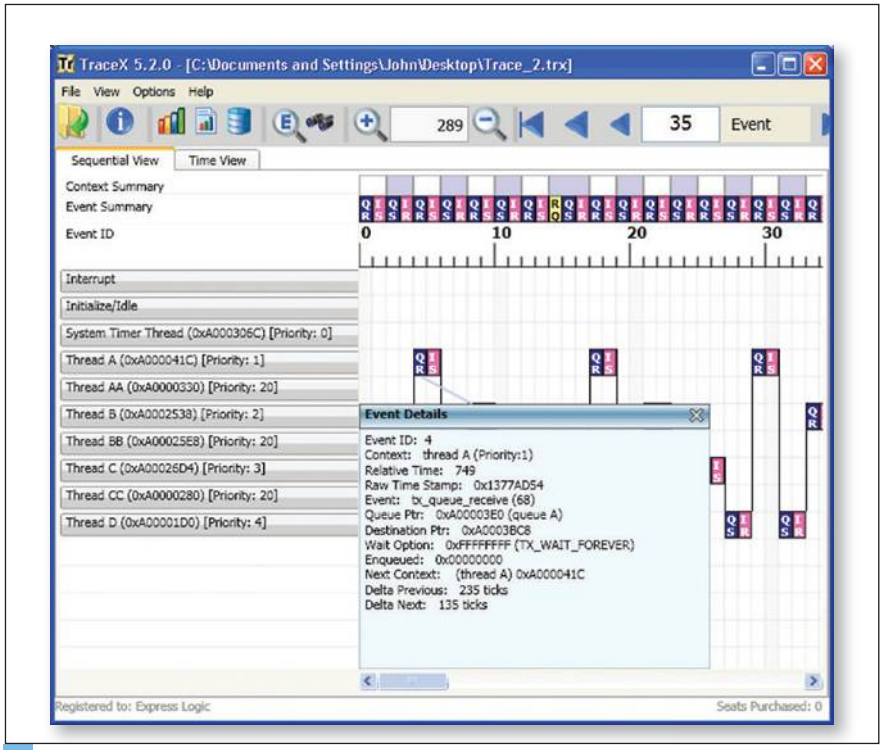

Figure 2 | Individual event details can be displayed by clicking on an event icon.

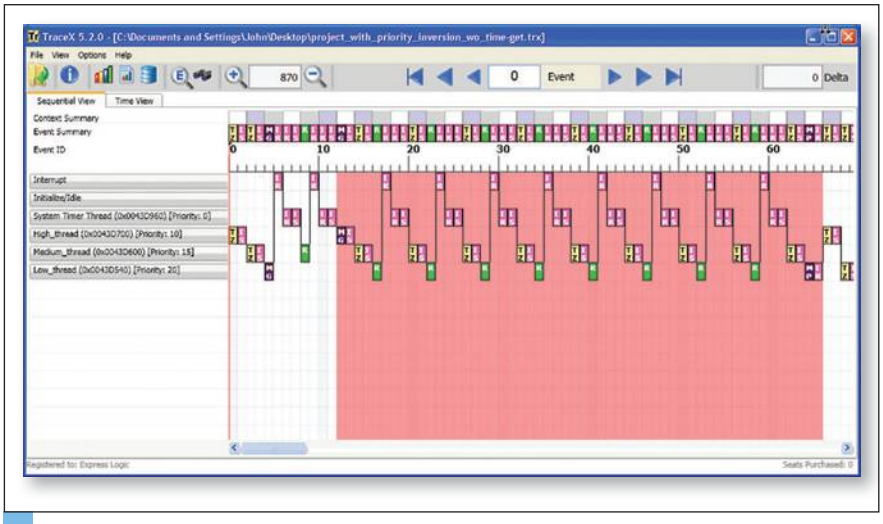

Figure 3 | The higher-priority thread must wait for the lower-priority thread to release a mutex in a non-deterministic priority inversion.

As shown in this graphic, Low_thread holds a mutex when it is preempted by High_thread. High_thread then seeks the same mutex, but must wait for Low_thread to release it. However, Medium_thread has intervened and can run for an indeterminate length of time, delaying not only Low_thread, but also High_thread. Only when Medium_thread yields enough time to Low_thread for it to complete its processing and release the mutex can High_thread resume.

Improving application performance

While most developers use multicore-enabled tools to understand and correct problems, the benefits don't end there. These tools offer an execution profile for analyzing and improving system-level application performance. Using an execution profile, developers see the amount of CPU time used by each thread and system services (see Figure 4). The developer can easily drill down on specific events for diagnostic purposes.

Even more relevant to multicore system operation, balancing the processing load across all available cores can achieve greater system throughput. If a system profile provides information about which cores have greater idle time, as shown in Figure 4, the developer gets a strong clue regarding how to shift processing to an otherwise idle core.

A multicore-enabled debugging tool paints a graphical picture of a system in a way that standard debuggers can't provide. It gives developers a clear view of interrupts, context switches, and other system events typically detected through time-consuming code instrumentation and tedious examination of the resulting data. Consequently, developers can find and fix bugs and optimize application performance in substantially less time than is required using standard debugging tools alone. With debugging taking up to 70 percent of application development, these tools significantly improve products while requiring less development time. ECD

John A. Carbone *is VP of marketing for Express Logic. He has 35 years of experience in real-time computer systems and software, ranging from embedded system developer and FAE to sales and marketing management roles. Prior to joining Express Logic, John was VP of marketing for Green Hills Software. He holds a BS in Mathematics from Boston College.*

Express Logic
858-613-6640
jcarbone@expresslogic.com
www.expresslogic.com

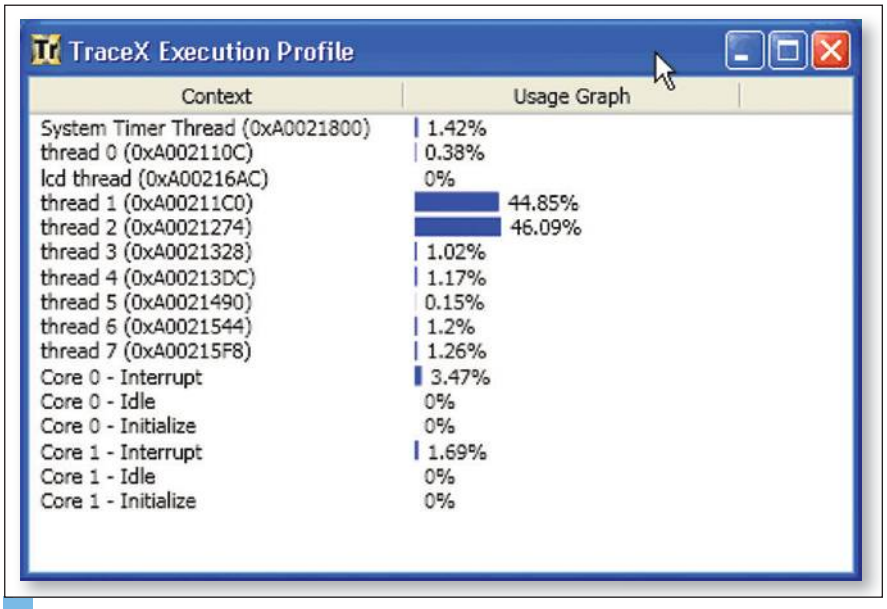

Figure 4 | An execution profile shows the CPU time used by each thread.


/* starts motor */

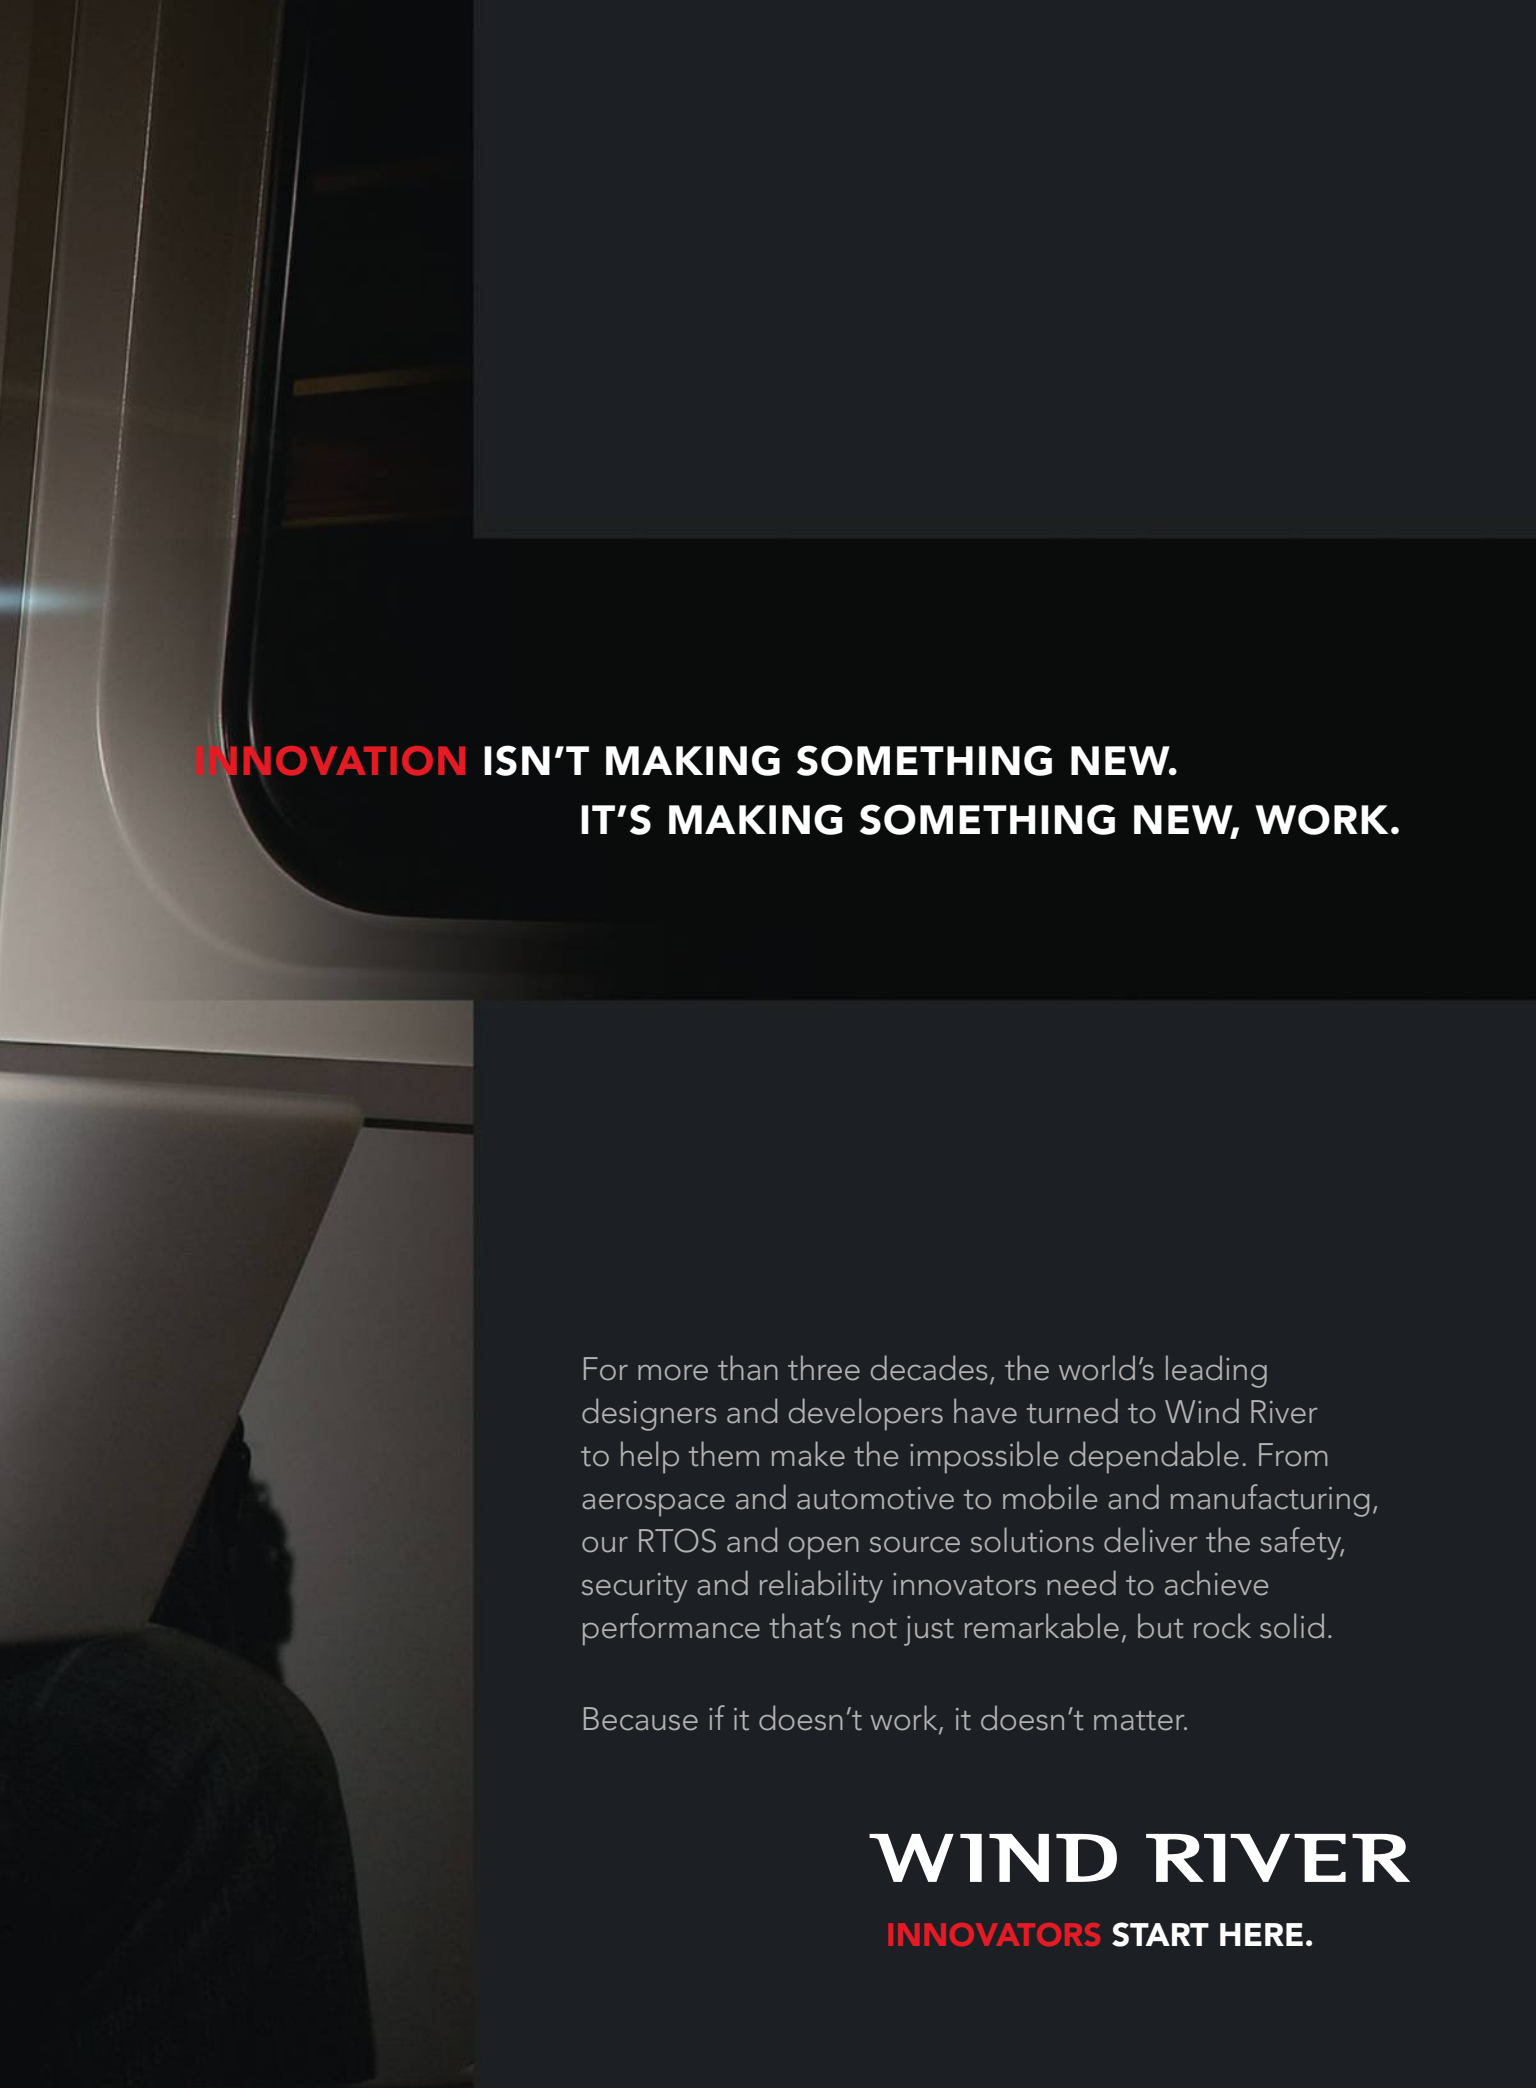

INNOVATION ISN'T MAKING SOMETHING NEW.
IT'S MAKING SOMETHING NEW, WORK.

For more than three decades, the world's leading designers and developers have turned to Wind River to help them make the impossible dependable. From aerospace and automotive to mobile and manufacturing, our RTOS and open source solutions deliver the safety, security and reliability innovators need to achieve performance that's not just remarkable, but rock solid.

Because if it doesn't work, it doesn't matter.

WIND RIVER

INNOVATORS START HERE.

EDA Consortium: Accomplishing together that which we cannot do as individual companies

By Robert Gardner

This quick overview of the EDA Consortium's main efforts illustrates key issues where cooperation helps both customers and vendors add value and drive the next big things in electronic design automation.

The Electronic Design Automation Consortium (EDAC) is the international association of companies developing EDA tools, software, services, and semiconductor intellectual property. Founded in 1989, EDAC addresses issues of common concern to EDA vendors. While initial efforts were focused on trade shows, EDAC efforts have expanded to include many areas where cooperation benefits both the EDA industry and customers.

The EDA industry includes three large companies, some medium-sized companies, and a number of smaller companies. Combined, the three large companies have about a 70 percent market share, with each having a dominant position in at least one subcategory of EDA tools. This suggests many designers are using tools from multiple vendors, creating a best-in-class design flow based on each vendor's strengths.

Much of EDAC's work is done by the operating committees, which consist of experts from member companies who dedicate part of their time to serve the overall good of the industry.

Working for fair competition

One of the more active EDAC committees is Anti-Piracy. Not long ago, it was the "collective wisdom" of the industry that EDA software was too complex and required too much support for software piracy to be a significant issue. More recently, the industry began to see evidence that this might no longer be true. The EDAC Anti-Piracy Committee worked with various anti-piracy vendors to evaluate the impact of piracy on the EDA industry. The committee selected a representative sample of EDA software from a number of members for an in-depth analysis in an attempt to gain a clearer understanding of the impact of EDA software piracy.

They found that EDA software piracy rates are about the same as other software products and represent a significant impact on EDA vendors and their customers. A design house using a pirated version of one vendor's software will not buy equivalent software from any of the vendors, regardless of the merits of the software. Furthermore, legitimate customers can find it difficult to compete against a design house with lower costs because they are not paying for their tools, resulting in decreased R&D budgets. Thus, pirating even a single EDA vendor's tools affects the entire industry. This is a major area where cooperation amongst vendors helps each company compete fairly based on the merits of their products.

Roadmaps and segments

Other EDAC committees include the Interoperability Committee, which publishes an industry OS roadmap. Agreeing to support, at a minimum, a core subset of the available OSs and versions reduces development costs within EDA companies and decreases support costs in the customer's development environment.

The EDAC Market Statistics Service (MSS) Committee publishes a quarterly report containing a detailed view of EDA industry revenue broken out by tool categories and geographic regions. Using revenue data collected confidentially from members and nonmembers, the MSS report (available by subscription) provides EDA companies, investment bankers, and analysts with a detailed look at trends in many areas of specialization within EDA.

These are just some examples of the work being done by EDAC to benefit the EDA industry, allowing members to focus on and deliver quality EDA tools. ECD

***Robert Gardner** has been an executive director of the EDA Consortium since 2006 and an ofácer/member of the board for more than 14 years. He has more than 40 years of management, engineering, operations, and sales experience, and has held executive management positions at several EDA and semiconductor companies. He holds a BSEE from California Polytechnic State University, Pomona.*

EDA Consortium
408-287-3322
www.edac.org

New paradigm speaks volumes for Knowles MEMS microphone design

By Pete Loeppert, PhD, and Nicolas Williams

Ever wonder what that microphone in your smartphone looks like? Or how it's designed? Electronic design automation tools have made microelectromechanical systems components like surface-mount microphones more efficient, compact, and innovative.

Jeremie Bouchaud and Richard Dixon, analysts at iSuppli, refer to Microelectromechanical Systems (MEMS) microphones as "one of the great success stories of MEMS." By outperforming Electret Condenser Microphone (ECM) technology in size, scalability, and ease of assembly characteristics, it's no wonder the MEMS microphone market is expanding from about $100 million in 2006 to a projected $300 million in 2013.

Traditional microphone technology has not kept up with market demand and is becoming a stumbling block as demand grows. Previously, microphone suppliers would stack up individual components and assemble microphones one at a time. Most are cylindrical, about 6 mm in diameter by 1-2 mm high, and inexpensive enough to be deployed in a range of applications from phones to toys. The problem is traditional microphones are heat-sensitive, which precludes the use of lead-free solder and the option for surface-mounting into circuits. So, to work around microphones' intolerability of high temperatures or reflow soldering, most manufacturers of high-volume products resort to an offline task, such as hand assembly or a special insertion machine, at the end of the mainstream assembly line.

Mounted for sound

Knowles has overcome the problems of ECM technology with the SiSonic MEMS microphone, which is batch-produced on silicon wafers and assembled like an integrated circuit except for an air pocket that allows sound waves to vibrate the diaphragm. An outline of the air pocket is shown in Figure 1. The SiSonic is reflowable, so an assembler can place it on a circuit board with a chip insertion machine, just like any other component. Microphones that can fit in with the normal assembly flow help reduce production cycle time, improve quality, and lower overall product cost.

The MEMS microphone is a bit of a renegade component in that it's neither a traditional microphone nor a conventional integrated circuit. The difference in the design paradigm is that in MEMS, there are no circuits *per se*. Knowles doesn't deal with schematic versus layout in the MEMS group, but instead draws complex polygonal and curved structures.

Electronic Design Automation (EDA) tools for designing these MEMS microphones have shortcomings. Most tools cannot handle complex geometries and are geared toward rectilinear design and layout schema. The Knowles microphone design, for instance, is largely circular and requires a tool that can create and manage toroidal elements (3D circles). And preparing to use a conventional EDA tool requires a significant amount of setup and configuration. While this up-front work might be needed for large circuit designs, where the engineer is working for weeks on a cell, it is hard to justify on smaller circuits.

When designing a cost-efficient MEMS microphone, it is critical to find a tool with a short ramp time and small up-front investment in setup and configuration. The different high-end tools the Knowles design team has used through the years to design microphones have two main drawbacks: an inability to handle complex geometries and too much overhead.

To better meet the challenges of this new paradigm, Knowles chose L-Edit from Tanner EDA for its MEMS design. This solution's hierarchical architecture provided flexibility to manipulate

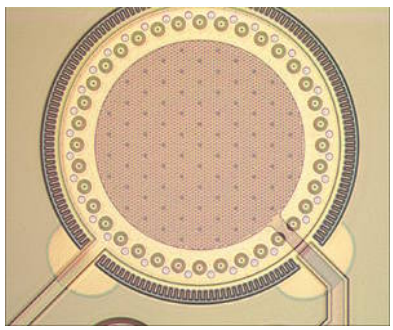

Figure 1 | An air pocket allows sound waves to vibrate the diaphragm in a Knowles MEMS microphone.

thousands of repeating elements. The team saved production time creating a variety of parametrically driven shapes. Circles, pie wedges, and tori were instanced slightly differently in every layer because Knowles used them as primitives. Although this was possible in a limited way with a tool like AutoCAD, designers wasted hours going back and forth between mechanical design and EDA tools. The ideal solution was to create and analyze designs in one environment and then send them out for photomask fabrication.

EDA scripting and fab rules

Knowles designs are not especially small – the die size is about 1 mm – but they are intricate. Drawings with 10-12 million objects are common. While MEMS design flow today does not lend itself to direct object generation through standardized libraries, scripting functions make creating and managing thousands of parametric objects extremely easy.

High-end tools have highly specialized scripting languages, but L-Edit users can write scripts using ordinary C/C++ code. The scripting function is flexible and often used to create primitives with a one- or two-page script. For example, in MEMS, Knowles often etches holes through the wafers and cannot have die intersecting the edge of the wafer. This means the dies must be arrayed in a circular pattern. Knowles now starts with an instance of a die and uses L-Edit to make it a rectangular array. The scripts clip the rectangular array to fit within the wafer extents, leaving a few millimeters for an exclusion zone around the edge, thus saving time.

The Knowles R&D team takes advantage of the scripting functions for other tasks as well, such as creating mapping programs for die bonding pick-and-place equipment. The script goes through the database of cells in the layout and automatically generates a wafer map, which is particularly important when working on a matrix design that has several different designs in it. The team can create the maps and assign different letters to different design styles, then tell the die-bonding engineers to pick a particular letter out of the map. The ability to generate the map automatically also helps save time.

Design Rule Checking (DRC) is different in the world of MEMS because there are few set rules. Designers have to create their own rules or work them out on-the-fly with the fab. The cross-sectioning tool in the editor helps with this by allowing Knowles to visualize designs in the third dimension of stacked layers.

Knowles exports to GDSII for handoff to their fabs. Using L-Edit helps the engineers work around two limitations peculiar to many GDS tools. First, the fab can have instances only at 90/180/270/360 degrees, but when Knowles engineers rotate things, they don't always end up with these particular angles. A script allows them to scan the database for any acute or obtuse angles and flush them out. Each one is then ungrouped, changed to be rectilinear, then later regrouped as it was originally.

Also supported is user-controllable fracturing of polygons, which lets Knowles set the limits for various mask fabrication vendors. Even when dealing with

limitations in Knowles' vendors' systems, L-Edit has helped overcome problems and interfaced smoothly with the fab.

Summing it all up

Knowles MEMS microphones require creation and manipulation of complex geometry and integration of mask layout data with advanced scripting. High-end EDA tools are not well suited to these types of requirements, and simple mechanical CAD tools can't perform the requisite tasks.

Knowles uses L-Edit tools to handle complex geometries and manipulate thousands of repeating elements at the heart of the company's MEMS designs. The results speak for themselves – 1 billion chips and counting. ECD

***Pete Loeppert** has been VP of engineering for the Knowles Acoustic Division of Knowles Electronics since 1999. His previous experience includes work at the Gould Electronics corporate research center on a variety of projects including high-speed A/D converters, chemical sensors, åber-optic systems, and GaAs circuits. Pete received his PhD in Electrical Engineering (specializing in solid-state devices) from Purdue University.*

Knowles Electronics
630-250-5930
www.knowles.com

***Nicolas Williams** is director of product management at Tanner EDA, where he leads product direction from concept to release. His åeld of specialty is analog EDA and analog, mixed-signal, and RFIC design. Nicolas received his BS and MS degrees in Electrical and Computer Engineering from the University of Tennessee.*

Tanner EDA
626-471-9700
www.tannereda.com

ECD in 2D:

EDA tools like Tanner EDA's L-Edit can handle complex geometries and accelerate analog IC layout. Use your smartphone, scan this code, watch a video: http://bit.ly/b8EjBF.

Embedding flexible analog interface IP into digital SoCs

By Manuel Mota, PhD

More and more analog and mixed-signal capability is being integrated into Systems-on-Chip, including baseband RF functions. Careful attention to detail is needed to maintain circuit performance and deliver functionality while keeping power consumption low. Designers can explore the criteria for a flexible analog interface to build a successful wireless communications device.

Today's consumers are getting more and more sophisticated. They are now creators of their own content, like HDTV videos or large sets of digital photographs. They expect to be able to share it anytime, anywhere, with anybody, and without delay; they do not want to be limited by slow wired connections.

These consumers are driving the implementation of broadband wireless network terminals in every conceivable device, beyond the traditional mobile phone and into the PC, TV, car, camcorder, and even picture frames.

To serve this consumer expectation, terminal devices must be cost-effective and simple to use. In addition, the wireless transceiver in these devices should be compatible with all broadband communication standards, including Long-Term Evolution (LTE), WiMAX, and Wi-Fi. These goals can be achieved by carefully defining the analog interface's characteristics and internal components, paying attention to several details.

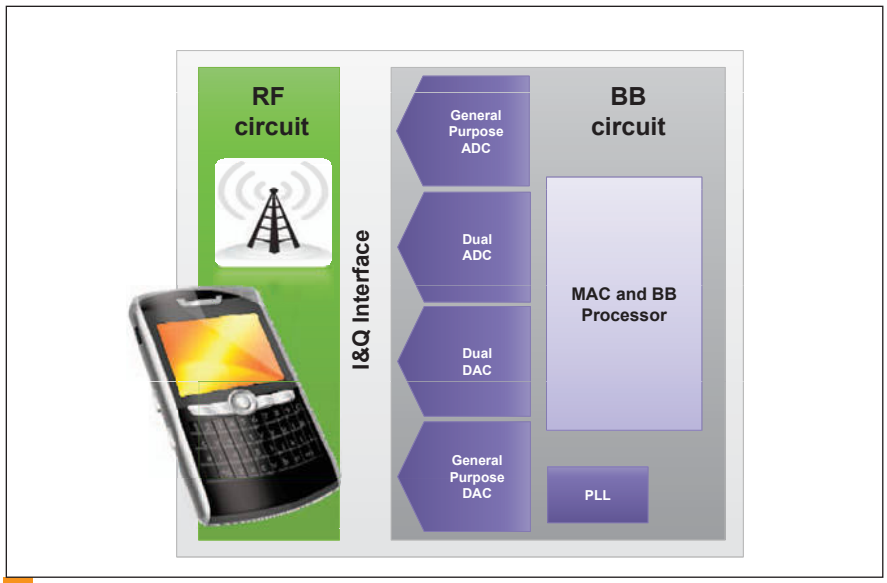

Figure 1 | An analog interface connects the analog RF block and digital baseband processor block in a broadband communications transceiver system.

The wireless baseband analog interface

A traditional communications system (Figure 1) comprises an analog RF block to translate over-the-air communication into baseband (BB) analog signals and a digital BB processor block to translate modulated signals into meaningful communication content. A wireless BB analog interface translates signals between the analog and digital domains.

To implement interchip communication between the digital System-on-Chip (SoC) with an external RF chip (RFIC), the analog interface is typically integrated with the digital BB processor in a complex digital SoC.

The analog interface comprises an IQ-Analog-to-Digital Converter (ADC) in the receive path and an IQ-Digital-to-Analog Converter (DAC) in the transmit path. It also includes an auxiliary ADC and an auxiliary DAC for measurement and control purposes. A Phase-Lock-Loop (PLL), which generates the sampling clock for all these converters, can also be considered part of the interface.

Most modern implementations of wireless communications transceivers are

multi-antenna array (Multiple In, Multiple Out or MIMO). In these cases, the receive and transmit paths use multiple instantiations of an IQ-ADC and IQ-DAC (one per antenna). These architectures implement advanced processing techniques such as diversity or special multiplexing to improve communication quality.

Embedding a flexible analog interface into a digital SoC

To implement multiple communication standards and enable the BB chip to be used with any RFIC, the analog interface must be very flexible. The definition of the interface and the internal data converters should take this into account to facilitate the analog interface's integration inside the digital SoC. The following ground rules should be followed:

- The interface should be compatible with multiple wireless communication standards while keeping power dissipation at a minimum.
- The interface should be flexible, allowing seamless connection between the BB chip and any external RFIC without requiring additional external circuitry.

The system designer should thus look beyond the basic characteristics of a data converter IP (area, power dissipation, throughput, dynamic range) and into the detailed characteristics that make it flexible and easy to integrate.

Compatibility with multiple standards

It is interesting to observe that most broadband wireless communication standards have defined a similar set of performance characteristics for the analog interface. In fact, most protocols define a maximum communication channel width at 20 MHz (or 40 MHz for Wi-Fi 802.11n). Furthermore, the conversion resolution required is typically 10- or 12-bit. This means a data converter with these characteristics is effectively compatible with multiple protocols. However, the optimal sampling rate for different communication modes within the same protocol (LTE mode 1 or mode 6, for example) may differ.

The communications system should be configured such that the same hardware can be used in all modes in the most power-efficient way. For the data converters in the analog interface, this means they should allow operation at a wide range of sampling rates without performance variation. Furthermore, operation at low sampling rates should not penalize the system's energy efficiency; the converter's power dissipation should be proportional to the sampling rate while maintaining performance, as shown in Figure 2.

Optimal interface with the RFIC (ADC/DAC)

For the highest system flexibility, the BB processor chip should be implemented in a way that does not impose restrictions on the RFIC connected to it.

Systems are typically implemented using a DC-coupled zero IF demodulation scheme. The analog signal level coming out of the RF block is not known and varies for different vendors.

To communicate seamlessly with all RFICs, the data converters in the analog interface need to support a wide range of input signal common mode voltages and perform the signal-level translation in its internal signaling levels without the need for extra external components (coupling capacitors or operational amplifiers).

For the receive path, dedicated level-shifting circuitry in the ADC sample and hold can implement all these functions and thus guarantee maximum interface flexibility. Figure 3 shows an example of an ADC input stage that can process signals with a wide range of common mode levels.

In the transmit path, a DAC using a current steering architecture enables the level translation function to be implemented without further overhead.

It is possible to alleviate the complexity of the analog reconstruction filter and thus simplify the RF block by using the transmit DAC at a higher sampling rate to increase signal oversampling. However, this solution is only viable if the DAC's power dissipation and performance are not affected by the higher sampling rate.

Support for 2G/3G/4G communications

A special category of wireless communication modems must support communications based on protocols for multiple generations (2G, 3G, and 4G), such as in the context of cellular applications. Modems for cellular applications must implement widely different requirements using the same hardware, for example:

- Narrowband Gaussian-filtered Minimum Shift Keying (GMSK) with 200 kHz channel bandwidth requires ADC performance in excess of 74 dB Signal to Noise and Distortion Ratio (SNDR).
- Broadband LTE with up to 20 MHz channel bandwidth requires ADC performance in excess of 63 dB SNDR.

The optimal solution for these modems is to implement the ADC using a wideband sigma-delta architecture. These converters are designed using a highly programmable analog sigma-delta modulator, followed by programmable digital filters. This embedded programmability allows trading off speed (signal bandwidth) and dynamic range (SNDR) while maintaining the power dissipation at a minimum. Additionally, due to the high oversampling rate, the analog anti-aliasing filters present in the RF block can be simplified.

Processing large bandwidth analog signals

The most common demodulation technique used for broadband wireless communications is based on a zero/near-zero IF implementation. However, in some cases, IF demodulation is preferred. In these situations, the IF signal is digitized and processed directly in the digital domain, thus simplifying the analog RF circuitry. In IF demodulation, the communication channel band to be digitized by the ADC is centered on a high frequency (the IF frequency).

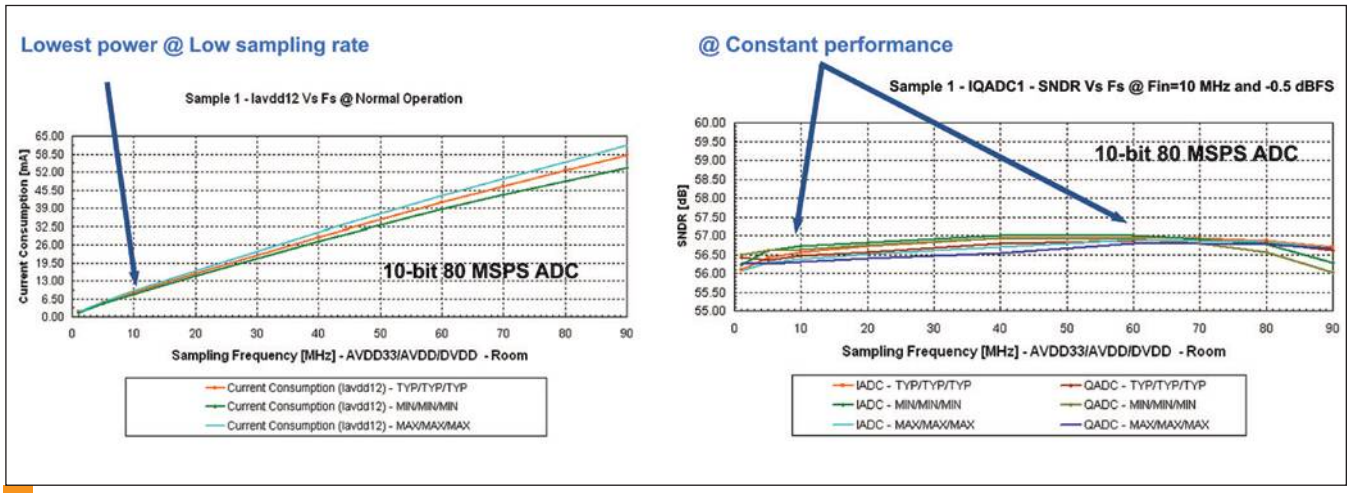

Figure 2 | For optimal implementation of the analog interface, the power dissipation scales proportionally with sampling rate while maintaining consistent performance.

To digitize these signals, the ADC input must be able to sample high-frequency signals while not jeopardizing the ADC's performance. Traditional converters for baseband applications exhibit performance degradation for high-frequency input signals. An ADC that integrates a dedicated high-frequency input stage targeting high-frequency signals can overcome this limitation (see Figure 4).

Ready-to-integrate IP

Achieving a successful design that integrates broadband wireless communications interfaces in a digital SoC depends on the definition of the analog interface and a careful selection of the blocks that make up the interface. Selection should look beyond the common criteria of resolution/performance and take into account the features embedded in these converters that help build a flexible solution. These converters should seamlessly communicate with any RFIC and accommodate all common broadband communication standards in the most power-effective way using the same hardware.

Synopsys DesignWare Data Converters were created in view of such needs. Synposys offers a complete portfolio of analog interface solutions that greatly simplifies its integration in a digital BB processor SoC. ECD

For more information on the Synopsys DesignWare Analog IP Data Converters portfolio, visit www.synopsys.com/IP/AnalogIP/DataConversion/Pages/default.aspx.

Manuel Mota *is technical marketing manager for data converter IP within the Solutions Group at Synopsys. He has worked in the semiconductor industry for more than eight years as analog IP designer and business developer for data conversion products at Chipidea Microelectronica. Manuel holds a PhD in Electronic Engineering from Lisbon Technical University, which he completed while working at CERN as a Research Fellow.*

Synopsys
650-584-5000
manuel.mota@synopsys.com
www.synopsys.com

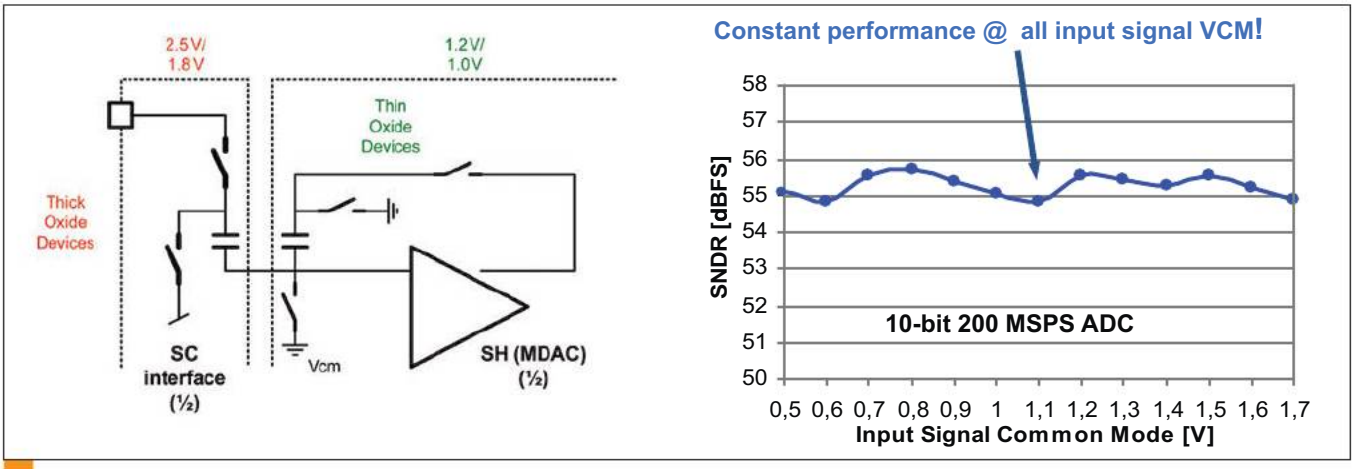

Figure 3 | An ADC input stage implementation compatible with all common RFIC signal levels can process signals with a wide range of common mode levels.

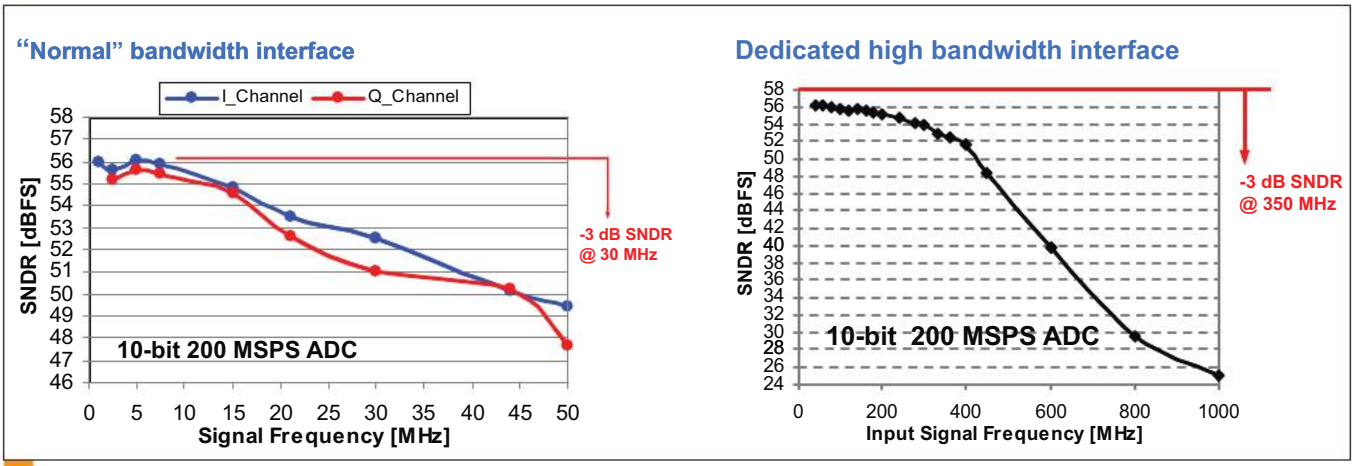

Figure 4 | Processing high-frequency input signals can impede the performance of ADCs unless they integrate a dedicated high-frequency input stage.

Virtual prototyping quad-core ARM with Android

By Achim Nohl

Applications and software services are the main differentiators for mobile devices like smartphones and tablets. Prototyping, development, debug, and validation of complex multifunction software applications is a major endeavor that is especially challenging when the software depends on hardware functions still under development and boards are not yet available.

Today's hardware is defined and configured based on software demands. The bring-up of applications based on a variety of software and hardware services such as audio/video, GPS, sensors, and radio poses major debugging challenges. Debugging does not end at the boundary of a single software module, library, or software layer; it involves analyzing the interactions between multiple software entities. Similarly, validation needs to cover functional performance as well as power requirements.

For prototyping, development, debug, and validation of complex software services, Synopsys provides a prototyping environment that addresses both software and hardware architecture. The core of this environment is a virtual prototype (Figure 1) of the ARM Cortex-A9x4 quad-core Versatile Express board, which is bundled with a complete Android 2.1 port. Software developers can take full advantage of the four-way Symmetric Multi-Processing (SMP) enabled by the Cortex-A9 CPU as well as Android and the Linux 2.6 kernel.

Integrating with Google's Eclipse-based software development kit via the Android Debug Bridge through a virtualized Ethernet connection, the virtual prototype can be used for general Android application development. It provides dedicated support to ease native code analysis and debugging inside Java applications. An increasing amount of performance-critical code as well as legacy C code is compiled into native libraries for the ARM CPU. The challenge is to debug those libraries along with the Java applications, Android runtime environment, middleware libraries, and OS. This is even harder with SMP Linux, as each process and thread needed to form a complex multifunction software service might execute on a different CPU and change the CPU when being rescheduled.

The virtual prototype is complemented with a software analysis environment that visualizes the scheduling of those processes per CPU over time. The ability to perform this analysis is based on dedicated OS-aware monitors inside the virtual prototype. Thus, embedded software instrumentation is not required. Each process can be analyzed down to the function and even instruction level.

Scripting enables fully automated control and inspection of all hardware and software aspects. This is crucial for the functional, performance, and power validation. Complex scenarios with environmental interaction such as touch-screen, keyboard, or radio events can be easily recorded and repeated in a deterministic fashion. A system-level software-centric assertion framework allows the software engineer to easily hook assertions into the replay scenario. Those assertions can validate if the actions taken by the user through the application are correctly propagated through the different software layers down to the hardware. Together with functional aspects, power states are analyzed over time, providing information about consumed energy.

The virtual prototype is based on Synopsys model libraries and Fast Models from ARM. Users can extend the virtual prototype with custom SystemC TLM-2.0 compliant models.

***Achim Nohl** is solution architect at Synopsys, driving the adoption of virtual prototypes for software development and veriåcation. Achim holds a degree in Electrical Engineering from the Institute for Integrated Signal Processing Systems in Aachen, Germany.*

Synopsys
650-584-5000
SLSinfo@synopsys.com
www.synopsys.com/virtualprototypes

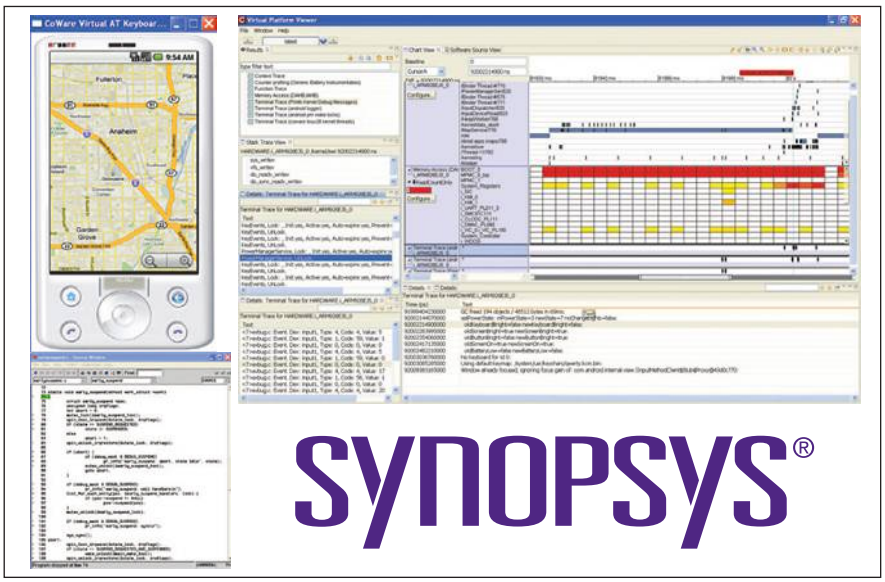

Editor's Choice

Finding bugs in concurrent apps

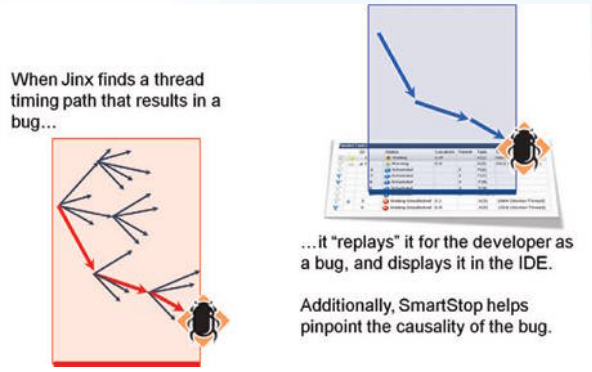

Debugging multicore code is one of the most interesting, nontrivial problems software designers face today. Problems that only show up on Mondays with a full moon and no silver bullets available can devastate humanity, and problems with concurrent code similarly devastate programmers because they walk right through most normal debugging tools.

Corensic's Jinx is designed specifically for multicore systems and is based on Deterministic Multi-Processing (DMP) technology. It forces bugs by setting up thread timing and simulating different paths of execution. Because it runs down different paths, DMP can make even a non-deterministic app behave in a predictable manner. Implemented in a very lightweight hypervisor sitting just above the operating system, Jinx is an interesting approach to finding potential multicore code problems where other tools might not.

Corensic | www.corensic.com

www.embedded-computing.com/p46088

Connecting with telehealth

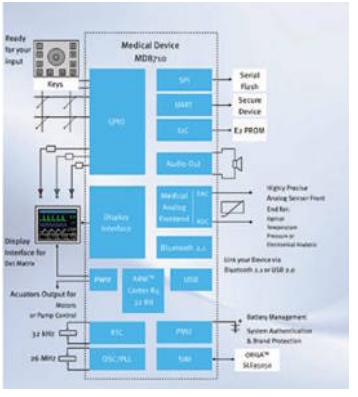

We've been looking for examples of approaches enabling telehealth devices in a growing market projected to be as big as $8 billion by 2015. Whether building a glucose monitor, blood pressure meter, or fitness device, a small but powerful and flexible System-on-Chip (SoC) platform enables both the processing and connectivity needed.

The MD8710 Medical Platform from Infineon shows how standard functions can fit a range of telehealth applications. The platform is more than just an SoC; it includes customer-specific analysis software that can be implemented on the ARM Cortex-R4 core. Other features include an analog front end, display interface, GPIO, audio, USB, Bluetooth, and a power management unit.

Infineon Technologies | www.infineon.com/medical

www.embedded-computing.com/p46500

Well-known RF name with new modules

Folks familiar with RF and microwave design know and trust the name Anaren for expertise and quality. So what if that expertise could be shrunk into an easy-to-use module for adding wireless capability to an embedded design? That's exactly why Anaren's foray into a new range of compact radio modules was interesting to see.

Anaren worked with Texas Instruments on developing low-power microcontroller and radio technology to make these modules simple to design with and simple to use. The modules are as small as 9 mm x 12 mm x 2.5 mm, come precertified to FCC Part 15 and Industry Canada RSS-210, and offer a choice of embedded or connectorized antenna. Several frequencies are available: 400-450 MHz, 868 MHz, 902-928 MHz, and 2,400-2,483.5 MHz.

Anaren | www.anaren.com
www.embedded-computing.com/p46498

New ref design for solar inverters

Solar panel installations are becoming much more sophisticated, with requirements for high-efficiency digitally controlled topologies supporting multiple panels, and are connecting the installation into the public power grid so that AC power can be fed back when not required locally. Solar panel designers are being challenged to get designs out quickly, but most aren't well versed in digital control.

To speed up the process of getting digital control into designs, Microchip has introduced a new solar inverter reference design based on the dsPIC33 GS family. The reference design converts a solar panel's DC output (up to 220 W) into AC power, with full digital control for improved efficiency up to 95 percent and power point tracking of 99.5 percent. The reference design itself consumes less than 1 W in nighttime mode, so it's quite power-efficient.

Microchip Technology | www.microchip.com
www.embedded-computing.com/p46499

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

Annapolis Micro Systems

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

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

» Where can I find the next generation Intel® Atom™ E6xx based scalable platforms? «

Kontron brings the next generation Atom™ E6xx processor to market on high performing, embedded platforms for use in extended temperature ranges of -40 °C to +85 °C.

- » The next generation Atom™ is a cost effective building block for future designs
- » Open standard interfaces deliver unprecedented levels of I/O flexibility
- » Extended graphics capabilities and low power enhance technologies like streaming video, video displays and audio
- » Increased software services including consultation, design, porting and validation

INTEL® ATOM™ E6XX BASED SCALABLE PLATFORMS

Learn more at **kontron.com/NextGen**

HMITR
EN 50155 Compliant HMI

- » Fanless and rugged design
- » Upgradable and scalable

nanoETXexpress-TT
Computer-On-Module

- » Ready for use in harsh environments
- » Optimized integrated resources
- » Simplified expansion options

CRITICAL QUESTIONS ... ANSWERED

Call, Email or Visit today.

Call: 1-888-294-4558
Email: info@us.kontron.com
Visit: kontron.com/NextGen

intel Embedded Alliance Premier

If it's embedded, it's Kontron.