Embedded COMPUTING DESIGN

www.embedded-computing.com

@embedded_mag
Embedded Computing Design

JUNE 2011
VOLUME 9 • NUMBER 4

Embedded testing: A real lifesaver

1 GHz PC/104 SBC
Supports Networking and Communications

Hardened for harsh, industrial environments, the low power PCM-VDX-2 PC/104 single board computer is designed for medical, transportation, communication, and control applications. It is packed with serial, USB, and Ethernet ports, plus GPIO; and offers -40°C to 85°C extended temperature operation.

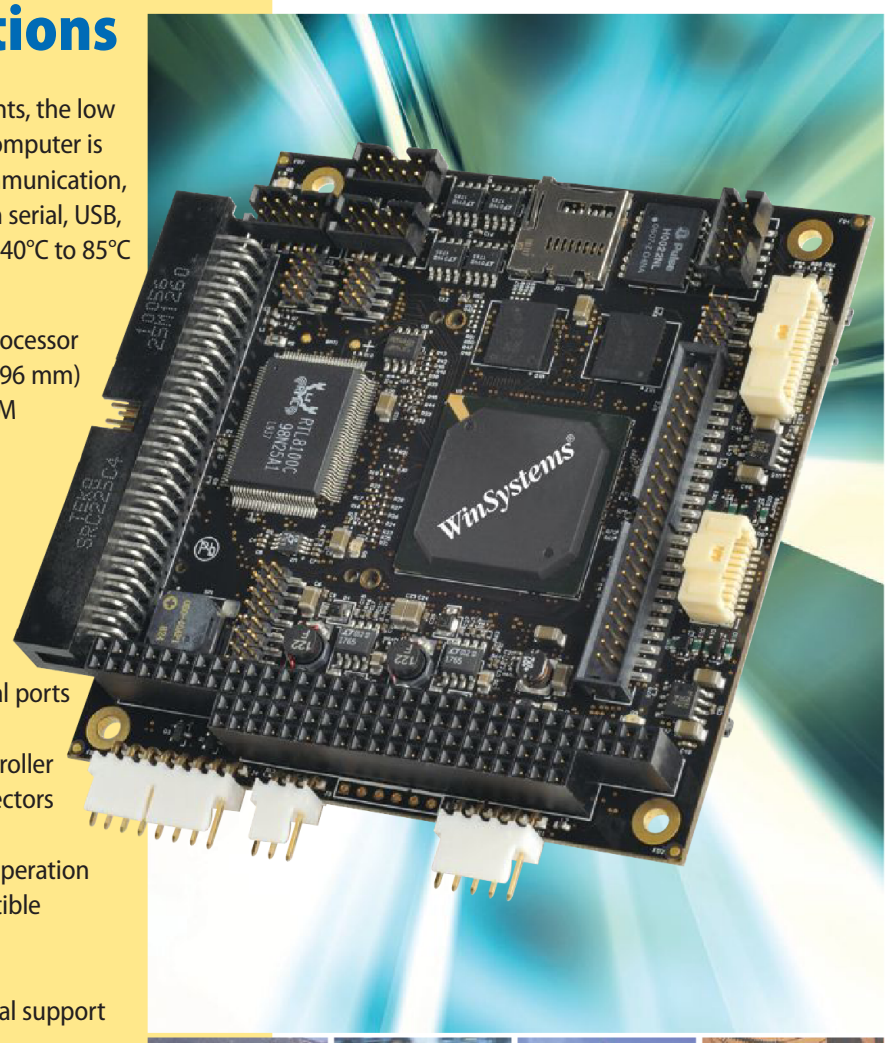

- Fanless, low power 1GHz Vortex86DX processor
- PC/104 Bus compliant form factor (90 x 96 mm)
- 512MB of soldered-down onboard DRAM
- 1MB of battery-backed SRAM
- CompactFlash socket
- Optional 512MB onboard SSD flashdisk
- Full-featured I/O includes:
  - Two 10/100 Mbps Ethernet ports
  - Four USB 2.0 ports
  - Four serial RS-232/422/485 ports
  - ESD protection on LAN, USB, and serial ports
  - 16 lines of general purpose I/O
  - PATA, LPT, PS/2 KYBD and Mouse controller
  - Mini PCI and PC/104 expansion connectors
  - WDT, RTC, status LEDs, and beeper
- Extended temperature -40°C to +85°C operation
- Runs Linux, DOS, and other x86-compatible operating systems
- Downloadable drivers available
- Responsive and knowledgeable technical support
- Long-term product availability

The PCM-VDX-2 can be customized by depopulating certain features or adding soldered-down flash memory, CompactFlash card retention clip, and/or a Mini PCI video card.




Understanding long-term product availability is a critical issue for our customers, the PCM-VDX-2 is offered beyond 2017. Contact us for additional product information and pricing. Our factory application engineers look forward to working with you.

Ask about our 30-day product evaluation.

Call 817-274-7553 or Visit
www.winsystems.com/PCM-VDX

WinSystems®
An Employee-Owned Company

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

AMD Embedded G-Series Processor: One Design, Fewer Watts, Massive Capability

AMD is ushering in a new era of embedded computing. The AMD Embedded G-Series processor is the world's first integrated circuit to combine a low-power CPU and a discrete-level GPU into a single embedded Accelerated Processing Unit (APU).

- Unprecedented level of graphics integration
- High performance multi-media content delivery
- Small form factor and power efficient platform

Learn more about new levels of performance in a compact BGA package at: www.amd.com/G-series

AMD **Fusion**
FAMILY OF APUs
EMBEDDED TECHNOLOGY

AMD

Embedded COMPUTING DESIGN

www.embedded-computing.com

June 2011 | Volume 9 • Number 4

On the cover

Medical devices and other critical applications require software verification tools to authenticate quality and ensure reliable operation. This issue's sections on medical device development and static analysis tools offer ideas for addressing the challenges of designing these vital systems.

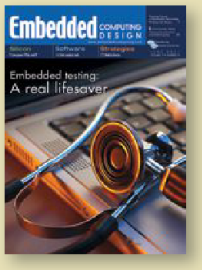

ISSN: Print 1542-6408, Online: 1542-6459

ZQ 7000

Introducing ZYNQ, the new element in processing.

Finally, the processor comes together with the FPGA in a fully extensible processing platform called Zynq.™ More intuitive to program in the way you already know. Fully customizable to your requirements. Faster to implement and get to market. As a software engineer, if you know ARM® Cortex,™ you already know Zynq. And if you know Xilinx, you already know this is innovation you can count on.

Visit us at www.xilinx.com

XILINX ZYNQ™

Reader resources

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

E-casts:
ecast.opensystemsmedia.com

TechChannels:
tech.opensystemsmedia.com

White Papers:
whitepapers.opensystemsmedia.com

Videos:
video.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/Description
3	**AMD** – AMD Embedded G-Series processor
32	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
23	**Arbor Technology** – Robust vehicle computer
21	**BittWare, Inc.** – Looking to maximize your FPGA processing power?
15	**CoreSolid Storage Solutions** – XDOM PCIe mini card
17	**EMAC, Inc.** – ARM9 module
31	**InnoDisk** – PCIeDOM II PCIe disk on module
13	**Sealevel Systems, Inc.** – RISC computing power
9	**Technologic Systems** – TS-SOCKET macrocontrollers
2	**WinSystems, Inc.** – 1 GHz PC/104 SBC supports networking and communications
5	**XILINX, Inc.** – Introducing ZYNQ, the new element in processing

Embedded COMPUTING DESIGN

OpenSystems media.

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

Embedded Group

Warren Webb
Editorial Director
wwebb@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

Christine Capuano, Web/Editorial Assistant

Brandon Lewis, Copy Editor

Steph Sweet, Creative Director

David Diomede, Art Director

Joann Toth, Senior Designer

Phyllis Thompson
Circulation/Office Manager
subscriptions@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

Christine Long, Online Manager
clong@opensystemsmedia.com

Laura Arsenault, Media Assistant
larsenault@opensystemsmedia.com

International Sales
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

Reprints and PDFs

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

Editorial/Business Office

16626 E. Avenue of the Fountains, Ste. 201 Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ **Fax:** 480-837-64661 ■ **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

Tracking Trends in Embedded Technology

By Warren Webb

@warrenwebb wwebb@opensystemsmedia.com

Verifying vital designs

Despite a weakened economy, the medical device industry continues to grow and provide plenty of embedded design challenges. According to AARP, the number of Americans over age 65 will double in the next 30 years. This increase will spark continued development of portable medical devices to monitor and analyze the complex medical conditions of an aging population. At this year's Embedded Systems Conference (ESC) in Silicon Valley, I met with several vendors developing hardware and software for the next generation of medical device applications, and everyone agreed that embedded designers must prepare for increased software complexity and ubiquitous connectivity.

Medical device design

In this issue's Strategies section, we cover some of the techniques and tools designers can employ for medical device development, manufacturing, and support. Chiman Patel of WIN Enterprises describes several medical case studies where off-the-shelf Computers-On-Modules (COMs) were used to simplify device development, improve time to market, ease product updates, and extend device life expectancy. His examples include a blood-gas analyzer, a DNA research system, and a laser cosmetic treatment device.

Another important part of medical device manufacturing is a calibration program to certify the accuracy of production measurements. Voler Systems' Walt Maclay outlines the requirements for a low-cost verification program to calibrate the temperature and pressure sensors contained in disposable catheters. Walt shows how manufacturers can conform to FDA regulations for measurement device validation and traceability to the National Institute of Standards and Technology (NIST). He also presents techniques to determine the test method, measurement accuracy, and record-keeping requirements of an approved product calibration program.

Medical devices are among a wide range of critical applications with escalating software complexity that can benefit from software verification tools to authenticate quality and ensure safe and reliable operation. Our Software section reviews static code analysis tools and methods designers can use to help locate and correct potentially vulnerable code. Rutul Dave from Coverity introduces static analysis as "the most cost-effective, automated, and repeatable way to meet the challenge of ensuring the quality of complex software."

The migration to multicore processors is increasing the complexity of today's embedded software. Developers must write code that takes advantage of multiple cores to gain the promised performance improvements. Multithreaded programming is difficult and subject to new defect types such as race conditions, deadlocks, and starvation. Paul Anderson of GrammaTech examines these multicore software defects and describes the tools needed to isolate them. I was fortunate to spend some time with Paul at ESC to uncover a few of his ideas for what's next in static analysis for multicore and multithreaded applications.

Low-power strategies

To reduce component count and decrease power requirements, designers turn to silicon technology to enable low-cost platforms that fit multiple embedded applications. This technology includes multicore processors, FPGAs, and System-on-Chip (SoC) devices, along with IP cores from device vendors, third-party suppliers, and the open-source community. The common goal among these technologies is to achieve higher performance with less power.

In our Silicon section, we provide updates on the latest techniques to extract performance from low-power hardware. Pete Hardee of Cadence Design Systems reminds designers that even if hardware is optimized for minimum power, the software and integration teams must properly utilize the power-saving features. Pete points out that the greatest power consumption in a typical mobile multimedia device comes from peripherals, memory architecture, and the digital processor subsystem. He also provides power-saving tips and cautions for system bring-up on a hardware prototype based on FPGAs rather than the final SoC.

Low-power requirements are not unique to portable embedded devices. Joseph Spisak of Sigma Designs describes the power consumption rules for a set-top box prescribed by regulatory bodies such as Energy Star and the European Commission. Although these set-top boxes must support advanced applications like HD video, 3D graphics, Web browsing, gaming, and telepresence, they must do it with less power. The advantages to a low-power design include quiet operation and a lower bill of materials. Joseph remarks that fans are the most fragile consumer electronic component and usually the first to break.

These articles tell us that no matter whether your design objective is a medical instrument or a set-top box, you face similar challenges: product development, time to market, device certification, software analysis, and reduced power consumption. We plan to scour the embedded community to uncover tips and techniques to keep you ahead of the competition. In particular, we're keeping a close eye on multicore updates and associated software. Please give us your ideas for articles and updates that we can provide to support your design efforts. If you have a suggestion for a technical article or video that would be of interest to you or other designers, please let me know.

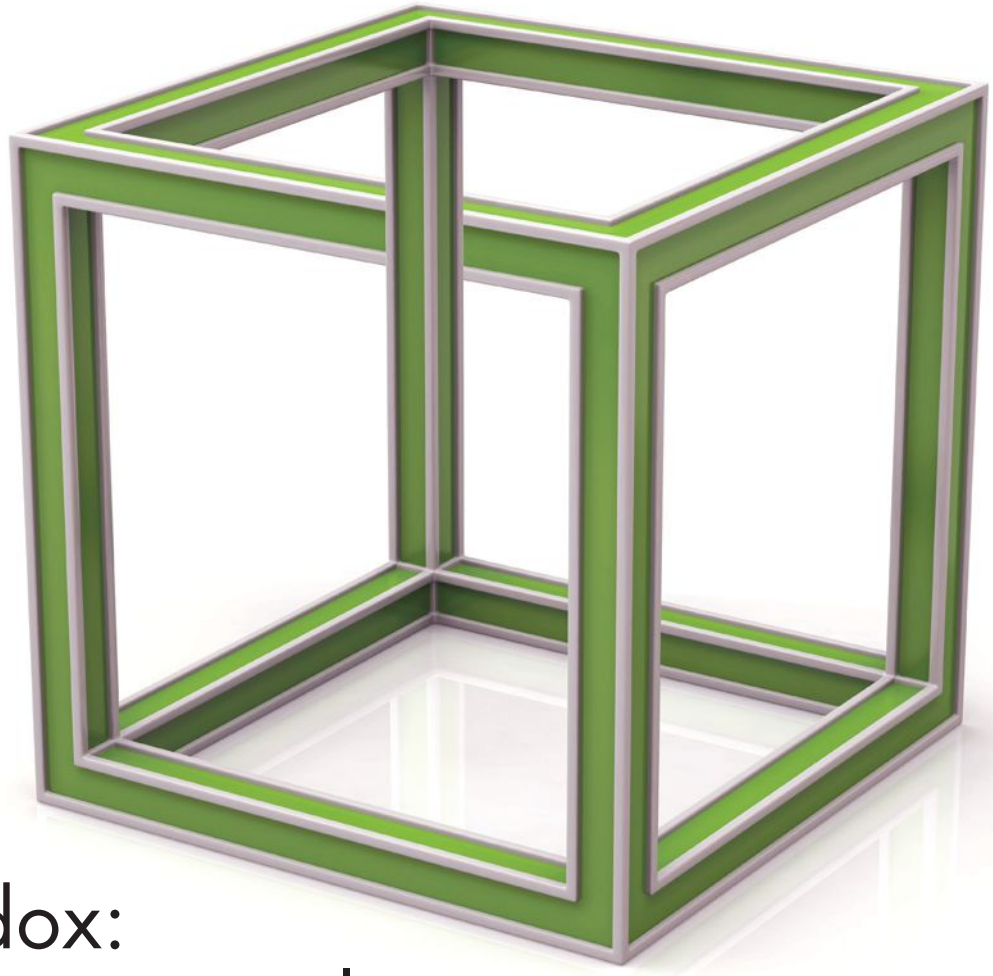

Set-top box paradox: High performance, low power

By Joseph Spisak

Today's set-top boxes and all other consumer electronic devices face the high-performance/low-power challenge. As the market demands more functionality, regulatory bodies are stipulating lower levels of power consumption, and businesses are seeing the advantages of low power, including greater reliability and lower bills of materials. Designers can use different techniques to increase performance when it's needed and save power when it's not.

The challenge: Low power, high functionality

Much is expected of today's Set-Top Boxes (STBs, see example in Figure 1). They must support new and sophisticated applications like HD video, 3D graphics and video, widgets, fast Web browsing, telepresence, gaming, and more. Moreover, they must support these new power-hungry applications with less power.

This is the high-performance/low-power challenge. All consumer electronic devices face it these days, STBs included. Even as general market forces demand more functionality, regulatory bodies are prescribing lower and lower levels of power consumption. Energy Star and the European Commission are the most prominent regulatory bodies driving the low-power trend. But private businesses also see an advantage to low power. Low-power devices often come with lower bills of materials. They're smaller, simpler, quieter, and also more reliable, in large part because they don't need a fan or heat sink. Fans are the most fragile consumer electronic component and usually the first to break.

It is generally accepted that devices consuming less than 5 W can avoid using fans and heat sinks. But on a chip level, how do low-power devices achieve unprecedented functionality? Chip manufacturers can use different methods to pack a greater punch with less power.

Less is more

There are two ways chips use power: dynamically and through leakage. Dynamic power consumption is the direct result of the quantity and speed of the digital processing a chip performs in operation. In STBs, these chips might be decoding audio and video and running a GPU to display graphics and video. They might be browsing the Web, supporting Skype, and so on. These activities result in dynamic power consumption.

Figure 1 | Low-power set-top boxes are smaller, simpler, quieter, and more reliable than other consumer electronic devices that consume a great deal of power.

Leakage, on the other hand, is power consumption unrelated to activity. It refers to energy that is consumed even when chips are idle. The only way to prevent leakage is to entirely shut off power to all or part of a chip.

Chip geometry also plays into power consumption. Oftentimes, smaller chips consume more power because smaller chips have smaller feature sizes. Smaller feature sizes, in turn, enable more transistors to fit on a single chip. Although this increases the functionality of a chip – enabling snappier responses, faster Web browsing, and better 3D and multiprocessing – it also increases dynamic power consumption and leakage. This brings us once again to the seemingly unavoidable trade-off between low power and high functionality. Is there any way to achieve the best of both worlds: low power and high performance?

Different ways of achieving standby mode

One popular way to create low-power STBs is implementing a standby mode. There are different ways of doing this. *Frequency scaling* involves reducing the frequency to a level that greatly decreases power but still enables fast wake-up. While this method reduces dynamic power consumption, it does not address the issue of leakage. *Power islanding* is another method that uses on-chip mechanisms to shut off certain blocks or cores when they are not being used. This method completely eliminates both dynamic power as well as leakage at the cost of chip complexity and increased wake-up time. *Dynamic frequency scaling* and *dynamic voltage scaling* use on-chip or discrete power management blocks that tailor the chip's operation to concurrently running applications, thereby providing performance where it's needed and saving power where it's not.

Sigma Designs uses frequency scaling and power islanding techniques to make its chips not only low power, but also low cost (see Figure 2). The SMP8652 chip supports HD video decoding, 3D video, a variety of conditional access solutions, and digital rights management solutions. At full speed, the Secure Media Processor requires a mere 2.2 W. Yet power consumption drops even lower to only 0.3 W in standby mode. The reduction is achieved by lowering the frequency of the chip's CPU from 500 MHz to 3.6 MHz and putting the DRAM in self-refresh mode. To save additional energy, the DRAM refreshes once every four cycles, not every cycle. In addition to the aforementioned frequency scaling, the chip achieves lower power with clock gating, a software-controlled process that shuts off the clock to different blocks on the chip but leaves the power on to enable a speedy wake-up from standby mode. Although

Figure 2 | Sigma Designs uses frequency scaling and power islanding techniques with its chips to achieve low power.

this also saves dynamic power, it does not affect leakage.

Another example is Sigma's SMP8910, a powerful 55 nm multiprocessor chip capable of providing 3D graphics, Blu-ray 3D video, rapid Web browsing, Over-The-Top (OTT) video, Skype, and the latest TV-viewing experiences. With all that functionality, the chip needs an energy management solution like power islanding to address both dynamic power consumption and leakage.

To achieve power islanding on the SMP8910, engineers created a chip with the following blocks: 3D graphics, video processing, and video display. These blocks constitute approximately 50 percent of the chip, so the ability to cut power to them significantly reduces leakage. (The mechanism for turning off these blocks when they are not in use is on-chip FET power switches.)

In deciding which blocks to put on a power island, engineers generally examine usage models. They ask themselves, "When certain applications are in use, which other functionalities aren't needed?" For example, if a person is simply watching a movie on TV, the chip's GPU (used mostly for video games and graphics) doesn't need to run, too. The GPU is a big block with lots of leakage. Using the power islanding method for leakage control, the chip's software will determine that the GPU is idle and turn it off.

Architecture's role in power savings

Chip architecture is also relevant to power savings. So as not to waste clock cycles, it is important to verify that clocks assigned to each block are running at optimal speeds. It is also important to make sure that clocks are turned off completely during standby mode. Whether a clock is actually off isn't readily apparent without special design tools enabling engineers to eliminate power bugs. Power bugs don't necessarily affect chip functionality; however, they invisibly increase power consumption. Design methodologies that eliminate these bugs help low-power chips live up to their name.

Hardware- vs. software-accelerated solutions

Engineers have found yet another way of saving power – up to a full 2 W on tasks like dual decode functionality (required by Blu-ray players). Instead of implementing a software-accelerated solution, they use a hardware-accelerated solution. This method offers a tremendous power advantage because it offloads work from the general CPU onto the discrete blocks within the chip. A hardware-accelerated solution enables decoding of HD video with a mere 200 MHz clock frequency. This leaves plenty of headroom for other applications. A software-accelerated solution, on the other hand, requires five times that amount – 1 GHz – putting a heavy burden on the system's CPU, requiring a much larger CPU and usually necessitating fans and heat sinks.

From a thermal engineering perspective, a hardware-accelerated solution is crucial to eliminating the need for fans and heat sinks, meeting current and future regulatory requirements, and lowering bills of materials. As stated earlier, chips that run cooler are known to be more reliable.

Dynamic Voltage and Frequency Scaling (DVFS) is another way to lower both dynamic power consumption and leakage. It has proven to be helpful in controlling the heat produced by today's smaller chips. DVFS offers the ability to adjust speed and voltage to the requirements of specific applications. In other words, if an application does not require a high frequency, DVFS enables it to run at a lower frequency while setting the chip's regulator at a lower voltage, reducing power consumption even further. Because idle blocks are dimmed rather than completely shut off, there is less wake-up time and blocks don't lose their state when they go into standby mode. However, unlike power islanding, DVFS doesn't halt leakage completely.

Lower STB energy consumption

Moving forward, chip engineers face the challenge of reducing STB energy consumption to levels even lower than those prescribed today. They also face the challenge of creating STBs that can wake up from standby mode in split seconds and from a variety of sources, like IR remote controls, the Ethernet, HDMI cables, wireless signals, and more. Indeed, STBs are now a part of a whole home network of digital devices, and they need to respond to commands from different devices within that network.

The high-performance/low-power challenge discussed in this article stands to benefit chip makers, device manufacturers, and service providers, as well as end users. Creating and implementing green technologies that allow STBs to become more powerful but less power-hungry will enable the STB industry to reduce its costs while increasing its value to consumers. **ECD**

Joseph Spisak *is the product marketing manager at Sigma Designs, where he manages the company's line of advanced media processors. An expert in video codecs and video processing, Joseph has also worked in engineering and marketing roles at Motorola, Freescale Semiconductor, Conexant Systems, and Gennum. He has a BSEE from Michigan State University as well as an MSc in Finance and an MBA in International Business from the University of Denver.*

Sigma Designs
Joseph_spisak@sigmadesigns.com
www.sigmadesigns.com

ECD in 2D: Check out a demo application designed for a TV or set-top box running on the Sigma Designs 8656 development board. With OpenGL/ES 2.0 and the FancyPants engine from Fluffy Spider Technologies, the application renders accelerated 3D graphics at 1080p. Use your smartphone, scan this code, watch a video: http://opsy.st/jp6Dcv.

Low power: The key issue for system integration in mobile devices

By Pete Hardee

The hardware and software worlds are colliding, and integration is providing a spark that will require system developers to keep a sharp eye on overall power demands. Innovative techniques can enable developers to validate whether software is correctly controlling the power-saving capabilities in the hardware platform and verify that the device can meet the power requirements in real system conditions.

As the lines blur between hardware and software development and integration, engineers tend to overlook the importance of developing these systems with power in mind. Even if a hardware design is optimized, the embedded software delivered within these systems must correctly and efficiently use the power-saving capabilities built into the hardware.

Teams developing these latest and greatest electronic systems need techniques focused on relieving the pressure of hardware/software integration. Techniques such as Power Shut-Off (PSO, also known as power gating), Multi-Supply Voltages (MSV), and Dynamic Voltage and Frequency Scaling (DVFS) can leverage platforms and advanced system-level verification capabilities such as emulation. Engineers require new forms to measure and dynamically analyze power requirements for integrating embedded software with hardware, which must be tested in real-world system modes by leveraging virtual platforms.

Yes, software burns power

Designers tend to think of power as a hardware issue, and indeed, it is the electronic hardware in an embedded mobile device that dissipates the power. Yet many recent instances highlighted throughout the blogosphere demonstrate how the latest release of a mobile device Operating System (OS) resulted in consumers complaining that their battery life was suddenly drastically worse. How so? While the power is dissipated in hardware, what the hardware is doing at any given moment depends on the user activity and system modes of operation, all under software control. For all the power-saving techniques that can be added to today's devices, the software needs to use them correctly to get the desired result.

The OS updates that have caused the greatest problems invariably changed the behavior of the system to leave one or many real-world interfaces in the "on" state, either by default or for longer periods of time compared to the previous release. In order of importance, here are

the areas where the greatest amount of power is consumed in a typical mobile multimedia device:

1. Peripherals and modems
2. Memory architecture
3. Digital subsystem and processor cores

Peripherals and modems

Pick up your favorite mobile device and press any key. The display lights up to full brilliance, and the keyboard probably lights up as well. How long before the display dims? When does the keyboard cease to be illuminated? We spend a lot more time interfacing with our mobile devices – browsing the Web, watching movies, texting, e-mailing – than we used to.

The LCD display driver in particular has become a big power sink, especially with the increased demand for bigger, brighter, and higher-definition screens. Furthermore, a modern global cellular device can have up to four modems. Voice, SMS, e-mail, and Web browsing all use a cellular modem. Alternatively, data may use an 802.11 Wi-Fi modem. Additionally, the device may use a Bluetooth modem or a satellite modem in a GPS subsystem. What is the device doing, not only while you are using it, but when you are not? Even then, the modems may be on, polling and synchronizing with various cellular, local, or personal networks.

The power amplifiers in these RF modems are particularly power hungry, but usually necessarily so, and the suppliers of these components make them as efficient as possible. Depending on the device usage profile, the modems may use more energy over time than the peripherals. However, software defaults such as always connecting the device to a Wi-Fi network if one is available or leaving the LCD display powered may have a detrimental effect on battery life, and sometimes unnecessarily.

Memory architecture

After the interfaces to the outside world, the movement of data within the device is responsible for the next most significant power usage. The memory architecture is structured in layers, usually referred to as L1, L2, and L3. L1 is cache memory, and usually caches exist for both processor instructions and data. L2 is on-chip memory and L3 is off-chip, usually in the form of DRAM and flash.

When software executes, memory transactions are the inevitable consequence. Transactions in L1 are quickest and cost the least in terms of power, then L2, then L3. There may be an order of magnitude difference in the power consumed by transactions successfully serviced in L1 versus L2 and L2 versus L3. Different cache policies and sizes are selected. Cache algorithms such as Last-In, First-Out (LIFO), First-In, First-Out (FIFO), or random can be more or less efficient depending on the regularity of the application being run. Once the policy is set, cache misses are subject to great attention, as the system will need to go to L2 or even L3 memory to retrieve the missing data. DDR or flash controllers manage access to L3. For memory transactions being serviced, most of the power is dissipated in the memory itself, rather than the controller or PHY interface. However, in most systems, especially ones processing video or graphical data, the memory controller is inevitably the bottleneck.

More power can be consumed in the system by memory transactions that are not granted access to memory compared to those that are processed; thus, depending on another set of priorities and policies, these waiting transactions must either be stored in FIFO until they can be processed or cancelled and resent at a later time. Depending on the current memory usage, memory pages can be switched off or put in standby mode when not being accessed. The efficiency of the policies and algorithms for this memory gating is another area that is highly dependent on software application and usage.

Digital subsystem

The rest of the digital platform including processor cores, hardware accelerators, random logic, and distributed registers account for the remaining power dissipation. It may be counterintuitive to some that the processor core itself would account for less of the power dissipated when software runs on a hardware platform, compared to memory or the peripherals the processor subsystem controls, but in most mobile devices with typical usage profiles it is true. Another rule of thumb is that typically in a digital subsystem, the clock network may account for 40 to 50 percent of the dynamic power.

Techniques to minimize power

Today's mobile devices are multifunction communication, entertainment, and productivity systems that pack an incredible array of functionality and performance. These devices are powered by the same size lithium-ion battery as the one-trick-pony voice-only cell phone of a few years ago – a battery that upsets us if it doesn't last at least a day without needing to recharge.

So what techniques were introduced in recent years to control the rate of energy used? At first, low-power design usually meant no more than clock gating – controlling dynamic power by disabling the clock to parts of the circuitry that could be idle and choosing a frequency of operation no faster than what was needed to meet performance. Then came processes and library cells with Multiple threshold Voltages (MVt). A low-threshold voltage switches quickly but uses more power, while a higher-threshold voltage uses less power at the expense of performance. These methods helped with both dynamic and leakage power.

As technology progressed down successive process nodes, this increased frequency (good for performance but not for power) but reduced supply voltage. Because energy is proportional to the square of voltage, this made up for the increase in frequency. Engineers could implement processes and libraries with MSV using different levels to balance performance needs versus power dictated.

However, the reduction in supply voltage has its limits. Reduction in V_{dd} too close to V_t leads to greatly increased leakage. Despite the availability of other techniques such as body biasing to control leakage, technology reached the point somewhere around the 45 nm node where leakage became a bigger issue than dynamic power, and the only sure way

to deal with it was to turn the circuit off when not in use. Hence, one advanced low-power design technique, PSO with or without state retention, has become common in today's designs. This technique is also known as power gating, or State Retention Power Gating (SRPG).

Meanwhile, developers continued to control the trade-off of dynamic power versus performance needs with techniques like DVFS. This technique works well on self-contained blocks whose performance supply and demand can be somewhat easily measured. DVFS has been successfully implemented on processor cores, for example.

These techniques have significantly increased the complexity of both design and verification. There are many different power states, all under software control. Certain protocols must be followed to successfully switch various parts of the design between power states, as well as switch them off and bring them back up when required.

Besides the functional design and verification needs all of that implies, there are many other structural needs. Each group of circuitry whose supply voltage may be separately switched, known as a power domain, must be fully separated from other power domains using appropriate isolators or level shifters on every signal that crosses the domains. The hardware team can exhaustively test and ensure that these techniques are implemented correctly, and all the necessary methods and solutions are available to do that. The purpose here is twofold: First, to validate during system integration that the software can correctly and efficiently control the power-saving capabilities provided in the hardware platform; and second, to analyze the resultant power savings and verify that the device can meet the overall power specification in real system conditions.

Validating power management under software control

System bring-up is usually executed on prototypes, often a board-based representation of the hardware implemented in FPGAs rather than the final System-on-Chip (SoC). However, it is very rare for power-management capabilities such as those described in the preceding section to

be successfully represented in a hardware prototype. Hence, the power-management aspects of system integration are all too often tested piecemeal, separately on the hardware side and the software side, and the two typically don't get integrated until the real hardware is available.

Can simulation help? Yes, but execution times are often prohibitively long to thoroughly check out power management with software. The problem is that once power domains are introduced to a design, many different power modes need very complex (and much longer) vectors to place the chip into that power mode and provide traffic representative of the system mode or combination of modes to which each power mode corresponds. Some companies painstakingly work out the data bandwidths for various parts of a chip to come up with vectors for maybe 30 different modes for power analysis. Even this might be missing a lot.

Bear in mind that these power-saving techniques do not come free. There is overhead associated with switching power domains off and bringing them back online. Unless a power mode endures for a certain time, turning idle circuitry off may waste power, not save it. Computation can be sped up so the power domain can be turned off for a longer amount of time. Hence, all the combinations of moving between those 30 modes must be considered to obtain an accurate picture. This becomes challenging for any simulation-based technique.

One answer is to use an emulation tool such as Cadence's Palladium Verification Computing Platform. Palladium was recently extended for power-aware execution, so vectors can be run thousands of times faster than logic simulation. This allows significant software to be executed, enabling engineers to check that power domains can be powered down and restored under software control. Virtual platforms – high-level simulation models of the hardware platform executing fast enough to allow significant software testing – also show promise as a solution in this area. However, more work must be done to successfully represent the necessary power states in these high-level models. Some of the advantages of emulation versus simulation can be seen in Figure 1.

System-level power analysis

Fundamentally, dynamic power boils down to a function of two things: characterization and switching activity. Characterization means accurately measuring and modeling what happens when a transistor switches – a function of V_{dd}^2, R, C, and increasingly, L. Switching activity means the frequency and duty cycle at which the switching happens for each of the transistors in the circuit of interest. That activity can be reduced to some extent by using the lowest clock frequency that gets the job done and turning the clock off when not needed.

To date, characterization has dominated thinking in power analysis. This has led to the idea that acceptable accuracy can be ensured only when there is a placed and routed netlist, all the transistors and wires are known, and all the RC values are extracted. However, characterization no longer seems to be the problem designers are struggling with; it's the activity. What vectors can run on the transistor netlist? What are all the system modes to generate realistic activity in today's multifunction devices? Are the vectors replicating those modes, or just running test patterns, or using statistical methods, which bear scant relation to real-life device operation? Given the complexity of power modes, and hence the quantity of activity vectors, this approach cannot cope with the activity, and the so-called accurate power analysis tools operating at gate level or physical netlist level may not be representative of the circuit used in the system. At the other end of the design abstraction spectrum, it may be easier to analyze real activity driven by executing system software, but the accurate characterization is missing. These issues are illustrated in Figure 2.

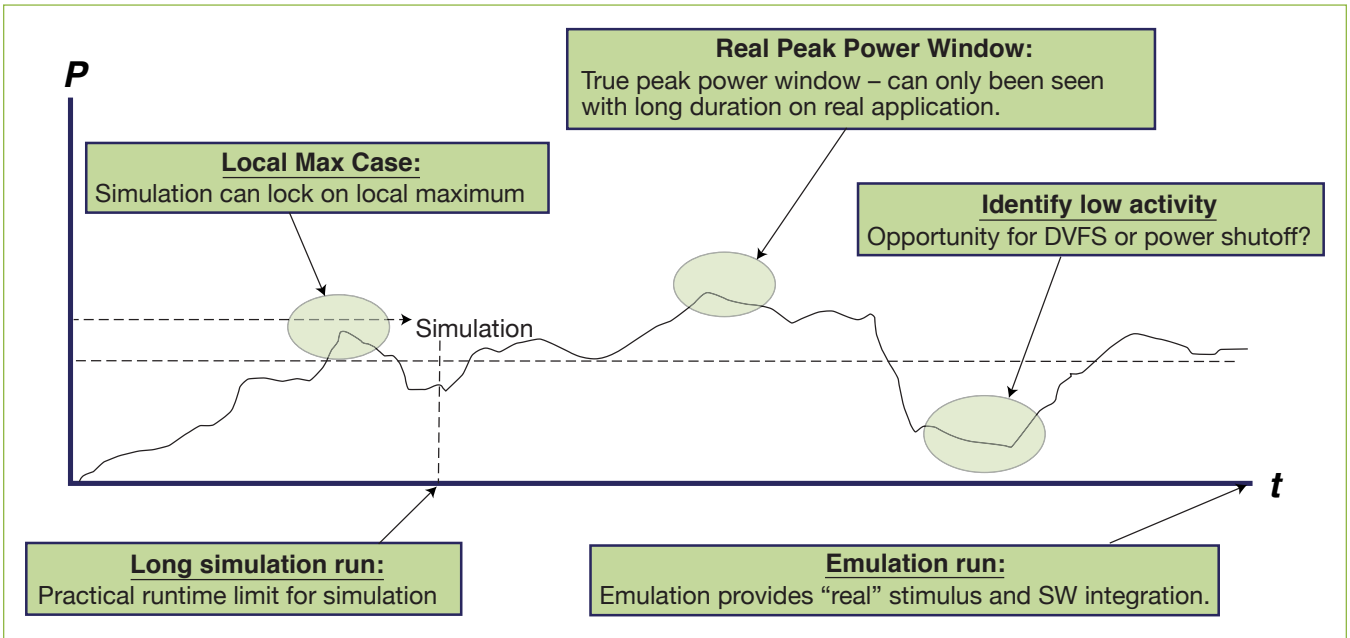

Figure 1 | Validating power management with emulation provides many advantages over simulation.

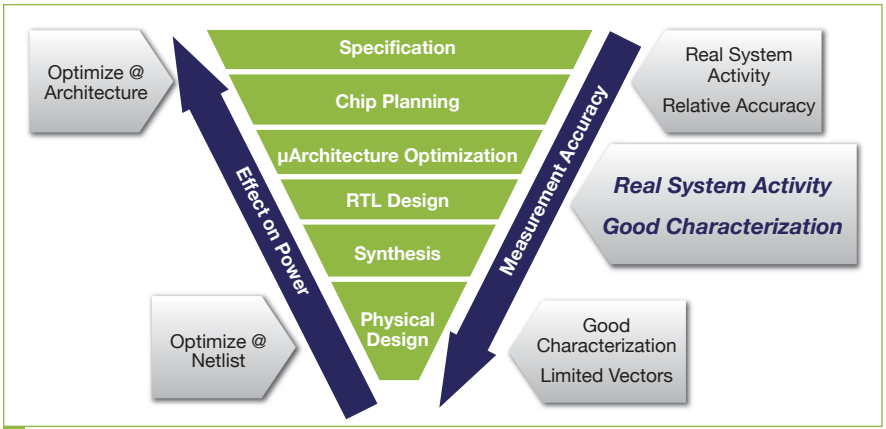

Figure 2 | The complexity of power modes makes system-level power analysis more difficult to accomplish accurately.

Palladium has a capability called Dynamic Performance Analysis (DPA) that allows designers to run as much real system-level activity as necessary, running real system modes under software control. The resultant activity is measured on all relevant points in the design mapped onto the emulator. This is very efficient, as this is exactly what an emulator is designed to do.

Besides the hardware design being mapped to the emulator box, the design is characterized using Cadence logic synthesis technology under the hood to map the design to the real cell library. This renders full representative system activity and characterization from the actual silicon process. The implementation is not exactly the same as the real chip would be, so it's still an estimate, but it may be the closest designers will get until the actual silicon running the actual application software is available. Which of course is probably too late. **ECD**

***Pete Hardee** is a director of solutions marketing at Cadence Design Systems. He is a 16-year veteran of the EDA and silicon IP industries. His experience previous to Cadence includes positions at Synopsys, CoWare, and Silistix. While at CoWare, he was a founder and first cochair of the Open SystemC Initiative. He has a BSc in Electrical Engineering from Imperial College, London, and an MBA from Warwick Business School.*

Cadence Design Systems
phardee@cadence.com
www.cadence.com

ECD in 2D: See a smartphone simulation running on the Cadence Virtual System Platform demonstrated at this year's ESC Silicon Valley. Use your smartphone, scan this code, watch a video: http://opsy.st/kyt6fA.

Static analysis: Beyond a simple list of defects

By Rutul Dave

You've been diligent and used static code analysis to identify defects early during development. Great. So now what? Finding defects is just the first step in the process of ensuring software integrity. Contextual information on every identified defect is essential to prioritize fixes and maintain bug-free software.

Embedded software is ubiquitous and provides critical functionality in a variety of devices, from the latest smartphones and gaming gadgets to life-saving medical devices. Engineering organizations creating embedded software understand that ensuring quality of code is a key differentiator and competitive advantage. Along with other methods of testing and verification, many companies have taken advantage of the benefits of code testing with modern static analysis to identify defects early in development. During the past few years, various reports by embedded market research firm VDC Research indicate strong growth in companies adopting static analysis as a critical test automation tool. Modern static analysis is arguably the most cost-effective, automated, and repeatable way to meet the challenge of ensuring the quality of complex software.

A strong reason driving this growth is that the technology used to identify critical defects such as memory corruptions, resource leaks, null pointer dereferences, and invalid memory accesses has matured to a point where finding large numbers of hard-to-find defects that traverse function and file boundaries can now be accomplished accurately, resulting in a very small number of false positives. However, the real innovation lies in providing contextual information for every defect identified. A developer needs to know why the defect exists, what impact it will have, and where it needs to be fixed.

The answer to the question of where it needs to be fixed is not as simple as knowing the file name and line number. Code branching and merging for versioning, reuse of code, and reuse of code components for development productivity allow a defect to make its way into multiple versions and products.

Consider the case of a software team with multiple branches for various versions of the product. A bug in one of these branches might exist in one or more other branches due to code replication. In another case, consider a team creating the framework to support applications for smartphones. Because they might be porting the framework onto various platforms like Windows, Android, or

iPhone, it is critical that the static analysis results clearly indicate whether identified defects exist in just one place or on multiple platforms. Similarly, when software is created by aggregating from multiple sources, it is a nightmare when a particular component is used in various products, as a defect in one third-party component might end up affecting all the different products that include it.

Multiple branches for different versions of an operating system

Imagine a software development team responsible for creating a new Operating System (OS) for mobile smartphones. Because multiple mobile phone vendors (OEMs) must be supported, every vendor in the source control management system needs a development branch. In addition, each vendor typically has multiple branches for different releases and product generations. The picture starts to get complex very quickly.

Static analysis performed on every branch of the code produces a list of defects. However, depending on when a defect is introduced, it could exist in all versions or a subset. When looking at a single defect in isolation in a single branch, the challenge for developers is that they can't gauge the severity of the defect without knowing where else it is present. A defect that is not limited to a single version or one OEM client would be severe, and fixing it would need to be prioritized over anything else. Additionally, a developer writing code to fix the defect needs to know exactly which branches in the source control management system the fix needs to be checked in. Analysis results that pinpoint the exact location where the defect exists and provide information such as the branches where defects occur are highly valuable to developers (see Figure 1).

A single framework for multiple platforms

On the flip side of branching, there is often a need to write code designed to run on multiple platforms. A software component such as a framework for mobile applications is usually built to run on various types of mobile phone platforms. For embedded devices, a common requirement is to build 32- and 64-bit

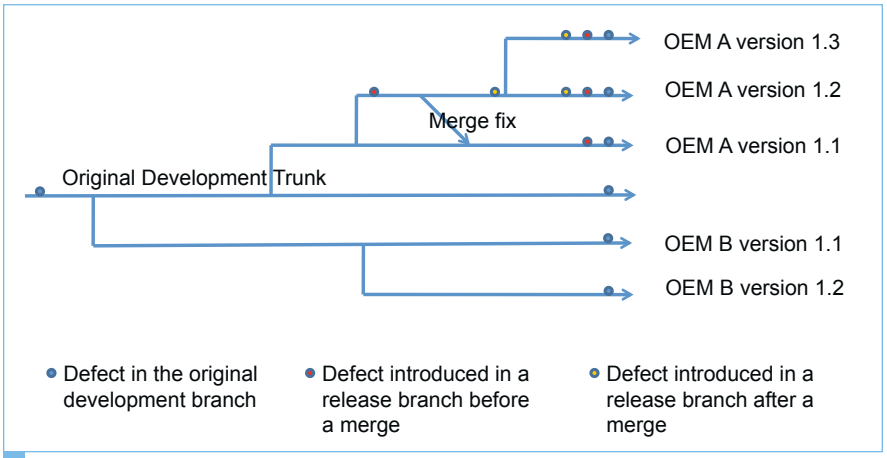

Figure 1 | Defects duplicate due to code branching and merging.

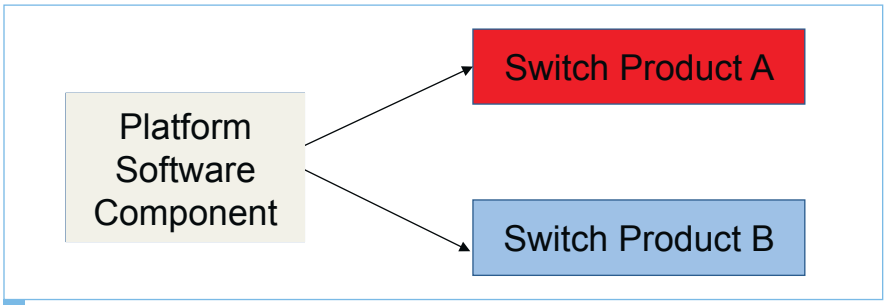

Figure 2 | A single software component is reused in multiple products.

versions of the same code base. Let's take a simple example:

```
gcc --m32 -c foo.c
```

// 32-bit compile. Contains a null pointer dereference defect.

```
gcc -c foo.c
```

// 64-bit compile. Contains the same null pointer dereference defect.

A defect in `foo.c` that gets triggered in both 32- and 64-bit binaries will be detected and reported as a single defect. However, because the source code is the same, a sophisticated analysis would not report it as a duplicate defect. Duplicates are as harmful as false positives in losing the developer's trust in the static analysis solution.

Sharing common code components

In this final example, consider a team developing the platform software for a family of networking switches. Because the functionality provided by the platform software must be implemented in all products, this code component will be shared (see Figure 2). For developers working on this team, the best assessment of the severity of a defect reported by static analysis is not only the impact it will have on one switch product, but also information on all the products that use this platform software component.

A product is usually created by combining many such shared components. Each component is not only a project itself, but

also a part of various other projects using it. The analysis result needs to identify that a defect in this shared component has an impact on the various projects using it.

Taking the guesswork out of code testing

The adoption of modern developer-side testing methods such as static analysis is a positive trend in the embedded software industry. The technology has matured to a point where it is a strong weapon in the software engineer's arsenal. Without needing to create elaborate test cases and testing infrastructures, static analysis automatically finds critical defects as code is written and compiled. However, for static analysis to become a developer's most valuable tool, the analysis must provide answers to questions such as "What is the impact of this defect?" and "Where do I need to check in the fix?" to help prioritize fixing the identified defects and take the guesswork out of ensuring that software is as bug-free as possible. ECD

Rutul Dave *is senior development manager at Coverity. He has several years of software development experience in embedded and real-time systems, including work developing bleeding-edge technology systems at Procket Networks, Topspin Communications, and Cisco Systems. When not evangelizing about the benefits of software integrity, Rutul scratches the coding itch by developing mobile apps and understanding the Linux kernel. He received his Master's in Computer Science with a focus on networking and communications systems from the University of Southern California.*

Coverity
info@coverity.com • www.coverity.com

ECD in 2D: Coverity CTO Ben Chelf explains how the software analysis capabilities offered by Coverity Integrity Center can help avoid debugging nightmares. Use your smartphone, scan this code, watch a video: http://opsy.st/ikAgqg.

More static analysis coverage online:

White paper: Safety-critical software integrity
http://opsy.st/jXxygH

Article: Better efficiency and lower costs for integration
http://opsy.st/kYDTEK

Finding concurrency errors with static analysis

By Paul Anderson

Due to the rise of multicore processors, programmers increasingly need to write multithreaded code. Pitfalls such as race conditions, starvation, and deadlock make it hard to get such code right, and traditional testing can only help so much. New approaches based on static analysis are proving effective at quickly finding difficult bugs.

Although decades of advances in miniaturization have yielded enormous performance gains for single processors, it appears this era is coming to a close. The best bet for achieving significant additional performance with single chips is through multiple cores, but only if software can be programmed to take advantage of them.

Unfortunately, concurrent programming is difficult. Even expert-level programmers familiar only with single-threaded programming often fail to appreciate that concurrent programs are susceptible to entirely new classes of defects such as race conditions, deadlocks, and starvation. It is difficult for humans to reason about concurrent programs, and some aspects of programming languages themselves are ill-suited to concurrency. Consequently, experts frequently stumble over these hazards. The following discussion describes common concurrency pitfalls and explains how static analysis tools can find such defects without executing the program.

Consequences of race conditions

A race condition arises when multiple threads of execution access a shared piece of data and at least one of them changes the value of that data without an explicit synchronization operation to separate the accesses. Depending on the interleaving of the two threads, the system can be left in an inconsistent state.

Race conditions are especially insidious because they can lurk undetected indefinitely, and only show up in rare circumstances exhibiting mysterious symptoms that are difficult to diagnose and reproduce. In particular, they are likely to survive through testing into deployed software. At best, this means increased development times; at worst, the consequences can be devastating.

One reason the Northeast Blackout of 2003 was so widespread was that a race condition in a computerized energy management system caused misleading information to be communicated to the operators. As Kevin Poulsen noted in a 2004 article (www.securityfocus.com/news/8412), "the bug had a window of opportunity measured in milliseconds." The chances of a problem like this manifesting during testing are infinitesimal. In another case, a race condition in iOS 4.0 through 4.1 (now fixed) meant that any person with physical access to an iPhone 3G or later could bypass its passcode lock under certain conditions.

An example of a simple race condition is shown in Figure 1. A manufacturing assembly line with entry and exit sensors maintains a running count of the items currently on the line. This count is incremented every time an item enters the line and decremented every time an item reaches the end of the line and exits. If an item enters the line at the same time that another item exits, the count should be incremented and then decremented (or vice versa) for a net change of zero. However, normal increment and decrement are not atomic operations; they are composed of a sequence of individual instructions that first loads the value from memory, then modifies it locally, and finally stores it back in memory. If the updating transactions are processed in a multithreaded system without sufficient safeguards, a race condition can arise because the sensors read and write a shared piece of data: the count. The interleaving in Figure 1 results in an incorrect count of 69. There are also interleavings that can result in an incorrect count of 71, as well as some that correctly result in a count of 70.

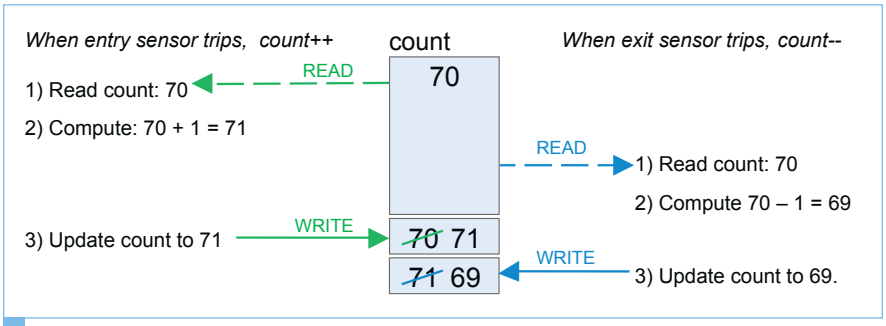

Figure 1 | A race condition leads to an incorrect count of items on an assembly line.

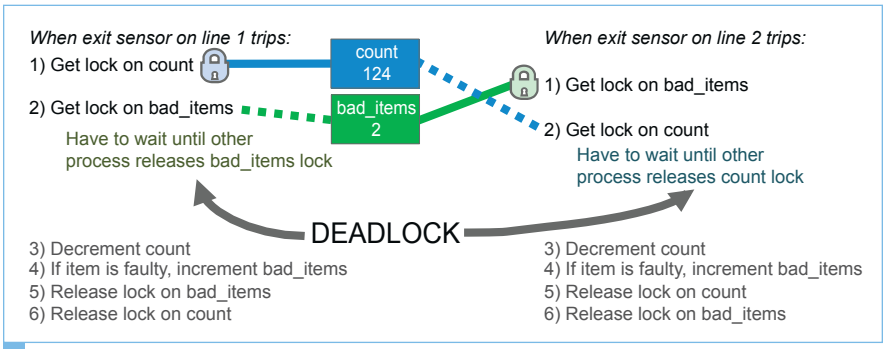

Figure 2 | In a deadlock between two threads, neither thread can progress.

For this example and for race condition bugs in general, standard debugging techniques may be ineffective for several reasons.

Rare occurrence means a reduced chance of noticing there is a problem. If the problem manifests infrequently, it may never show up during testing. The issue is twofold. Firstly, the number of possible interleavings of the instructions in two threads can be huge and increases enormously as the number of instructions grows. This phenomenon is known as *combinatorial explosion*. If thread A executes M instructions and thread B executes N instructions, the possible interleavings of the two threads are:

$$\binom{N+M}{N}$$

For example, given two trivial threads with 10 instructions each, there are 184,756 possible interleavings of the instructions. Real-world software is large and complex; testing every interleaving is simply impossible. Secondly, even if testers can identify a few interleavings that merit inspection, it is difficult to set up test cases to ensure they actually occur because thread scheduling can be highly nondeterministic.

If exhaustive testing is intractable, then what can developers do? One extremely useful approach is static analysis. Advanced static analysis tools such as CodeSonar use highly sophisticated symbolic execution techniques to consider many possible execution paths and interleavings at once. These techniques can find race conditions and other concurrency errors without needing to execute the program.

Several factors make race condition diagnosis difficult. Firstly, the symptoms can be perplexing. In the Figure 1 example, the running count will usually be correct, but sometimes too high and other times too low. Secondly, programmers unaccustomed to considering the particular pitfalls of multithreaded programming may spend a lot of time puzzling over the code before the possibility of a race condition occurs to them. Advanced static analysis tools are especially helpful in this regard. They identify race conditions by examining patterns of access to shared memory locations; that is, they focus on the race itself, not its symptoms. When a race condition is identified, an advanced static analysis tool will report it along with supporting information to aid the user in evaluation and debugging. The onus on the programmer is substantially reduced.

More complexity, more bugs

Race conditions are typically avoided by using locks to protect shared resources. However, locks can introduce performance bottlenecks that might prevent the program from taking advantage of the full potential of multiple cores, so programmers must exercise care in using them. It can be tricky to write code that uses locks effectively, and this complexity can lead to a different set of problems, namely deadlock and starvation.

In a deadlock, two or more threads prevent each other from making progress because each holds a lock needed by another. Figure 2 shows how a deadlock can arise with two locks used to protect two shared variables. In this example,

multiple assembly lines share a count of the total number of items currently under assembly, and a second bad_items value records how many finished items failed quality control. One thread acquires the lock on count, another acquires the lock on bad_items. Neither thread can obtain the second lock it needs; thus neither can carry out its operations, nor can it get to the point where it will release its lock. As neither update can be completed, both threads are completely stuck.

Static analysis tools can identify software at risk of deadlock by flagging situations where the same locks can be acquired in different orders by different threads, such as the threads shown in Figure 2. Eliminating all such cases is sufficient to ensure that the system cannot become deadlocked.

Starvation is another problem that occurs in multithreaded programs that use locks. A thread can starve if it is waiting for a resource currently held by another thread that takes a very long time. For example, suppose the aforementioned manufacturing automation system includes a regular audit thread that examines all entry and exit records to ensure that the running count matches total items entering less total items exiting. The audit thread would need to hold locks on the count and on all sensors, so all updates must wait for the audit to finish. If the audit runs for a long time, updates can be significantly delayed. If it runs too long, the next audit may manage to acquire all the locks and start running before the outstanding thread can make any progress. In the worst case, some or all of the updates may never have the opportunity to run.

Static analysis can provide significant value here by posing questions such as, "Is there a call to a long-running library function while a lock is held?" Tools such as CodeSonar also provide mechanisms for users to add their own checks. If an in-house function f() is known to have a long running time, engineers could add a custom check that triggers a warning whenever f() is called by a thread that holds one or more locks.

Multithreading adds entirely new classes of potential bugs to those that embedded developers must consider, making it significantly more difficult to find bugs of all kinds. The latest generation of static analysis tools can help with both of these issues. ECD

Paul Anderson *is VP of engineering at GrammaTech, where he manages the engineering team and architects static analysis tools. He has worked in the software industry for 20 years, helping organizations including NASA, the FDA, the FAA, MITRE, Draper Laboratory, GE, Lockheed Martin, and Boeing apply automated code analysis to critical projects. He received his BSc from King's College London and his PhD in Computer Science from City University London.*

GrammaTech
paul@grammatech.com
www.grammatech.com

ECD in 2D: A demo of CodeSonar shows how the tool can identify bugs in each version of the software being analyzed. Use your smartphone, scan this code, watch a video: http://opsy.st/kSRTWh.

Calibration verifies accuracy in disposable medical devices

By Walt Maclay

FDA regulations for medical devices require manufacturers to develop and monitor production practices to ensure that the end device conforms to its specifications. A calibration program is a major part of this quality system to verify the accuracy and precision of measurements. Walt describes the techniques, tools, and equipment used to calibrate a disposable medical device.

Medical devices that rely on one or more embedded sensors not only need testing, but also sensor calibration. Digital circuits can be subject to pass/fail testing. However, sensors require precise testing or calibration to be accurate. Calibration ensures that different individual devices give the same result. Sensors needing calibration include temperature, pressure, acceleration, strain, and displacement.

Medical devices must meet FDA regulations regarding device traceability and validation. The calibration process must ensure traceable measurements where the start and end of a calibration chain can be followed. Disposable devices add another constraint because the cost must be kept low without sacrificing accuracy and traceability.

The following discussion explains how the temperature and pressure sensors in a disposable catheter met accuracy requirements while maintaining low cost and meeting FDA requirements for traceability to the National Institute of Standards and Technology (NIST).

Achieving accuracy in temperature and pressure measurements

Depending on the required accuracy of the sensor, calibration may be simple or complex. A sensor that meets the accuracy requirements according to a manufacturer data sheet may only need to be tested to show that it is functioning, if the manufacturer is qualified by the device manufacturer's quality system. The pressure sensor in the catheter fell into this category.

A sensor that nearly meets the accuracy requirements may only need a single-point calibration to remove the offset error. For many sensors the largest error is the offset. A single-point calibration can significantly improve accuracy.

A less accurate sensor or a more demanding application may need a two-point calibration to remove offset and gain errors. Calibration requires subjecting the sensor to two different levels of temperature, pressure, or whatever is being calibrated, usually near the extremes of the range. For a temperature sensor, there is a time delay to reach a stable temperature, which adds cost to

the calibration. This type of calibration is often impractical for a disposable device.

A sensor that is far from meeting the accuracy requirements or one that is very nonlinear may need to be calibrated at multiple points to compensate for offset, gain, and linearity errors. This is rarely done on disposable devices.

When a device is calibrated, the results must be used to modify the device output. Sometimes a potentiometer is turned to set the calibration. Potentiometers are subject to change with temperature, vibration, and aging. More often, the calibration is stored in nonvolatile memory on the device. A processor reads the calibration and adjusts the readings. Storage in nonvolatile memory lends itself well to low-cost disposable devices with automated calibration.

Meeting FDA traceability requires NIST traceable calibration and tracking of devices by serial number, which is common for devices in other fields besides medical devices. The record-keeping requirements are more stringent for medical devices, and automated calibration can efficiently generate the electronic records used by the record-keeping system.

Calibrating a disposable catheter

Voler Systems built a tester for disposable catheters (Figure 1) that required calibration of the temperature and pressure sensors as well as functional tests. It was critical that the catheter temperature and pressure sensors were verified to meet accuracy specifications to assure that the catheter was safe for the operating temperature range of -80 °C to +37 °C.

The catheter used thermocouples to measure temperature. Thermocouples, although very common, are one of the most difficult sensors to use. They need a separate cold junction sensor, which gives a reference temperature since the thermocouple itself only reports temperature relative to where its wires are connected. Thermocouples are nonlinear and have an output of only a few microvolts per °C, requiring high gain. In modern digital systems the linearization is easily performed with a processor using a calculation or lookup table. The high-gain amplifier and the cold junction sensor each add errors that need to be considered in the calibration.

In this case the thermocouple calibration needed to be more precise than a one-point calibration. The thermocouples were calibrated at +37 °C. With a one-point calibration, an offset is added to all temperature readings. However, in this case the offset could actually increase the error at -80 °C. So it was important to not apply any offset at -80 °C.

This was done with a calculation. It was assumed there was no error at -80 °C, which was used as the second point in a two-point calibration without actually making a measurement at -80 °C. The result is an offset at +37 °C and no offset at -80 °C. Between these extremes, the offset decreases linearly from a maximum at +37 °C to no offset at -80 °C.

The test system used a commercial dry bath metrology well for temperature calibration, as placing the catheter into water could increase the bio-burden, making it difficult to sterilize. The calibration was made traceable to NIST by calibrating two temperature probes: one in the metrology well and another that measured the cold junction temperature near the catheter handle.

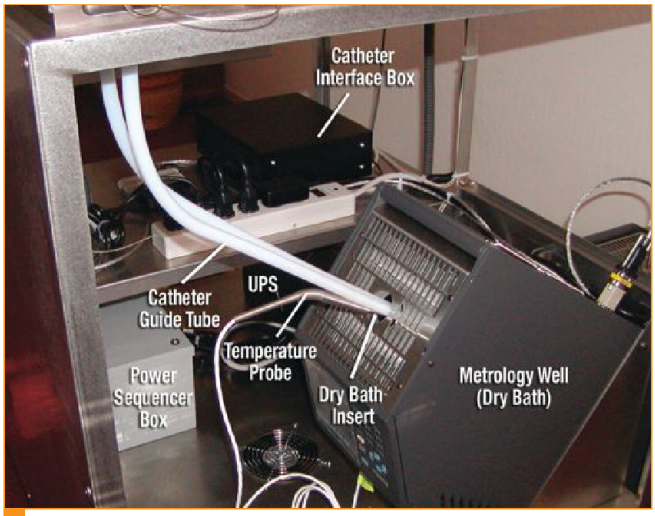

Figure 1 | A catheter is calibrated in an air bath (metrology well on the right) using a custom-designed guide tube.

The pressure sensors met the accuracy requirements according to vendor specifications, and the electronics added little to this error. The vendor was going to be approved through the manufacturer's quality system, so it was deemed sufficient to check the pressure sensors' outputs at atmospheric pressure, without applying

a calibrated pressure or performing a calibration. Because the pressure sensors are absolute devices, one point was sufficient to verify the pressure sensors were working accurately.

Some of the errors in Table 1 were in the console that the catheter plugs into (see Figure 2). Catheters must be interchangeable with any console, so the console errors could not be removed by calibrating the catheter. Because the console is not disposable, the electronics in it were designed to have very small errors, and almost all the errors were in the catheter.

REQUIRED ACCURACY			
Temperature		Pressure	
±2.0 °C at +37 °C		±1.1% of full scale	
±5.0 °C at -80 °C			
ERROR SOURCES			
Temperature		Pressure	
Sensor – all errors at +37 °C	0.5 °C	Sensor – all errors	0.9%
Sensor – all errors at -80 °C	1.9 °C	ADC offset	0.003%
Amplifier offset	0.7 °C	ADC gain	0.046%
Amplifier gain	0.75 °C	ADC resolution	0.0015%
Cold junction sensor offset	0.5 °C	Total error	0.95%
Cold junction sensor gain	0.25 °C		
Cold junction gradient	0.5 °C		
Cold junction amp offset	0.00 °C		
Cold junction amp gain	0.02 °C		
ADC offset	0.00 °C		
ADC gain	0.15 °C		
ADC resolution	0.00 °C		
Total error at +37 °C	3.37 °C		
Total error at -80 °C	4.77 °C		

Table 1 | Error calculation without calibration.

The settling time for the temperature calibration proved to be a major consideration. In the calibration bath the thermocouples stabilized rapidly in about 30 seconds. The cold junction sensor, located in the handle of the catheter, stabilized more slowly, taking 10 minutes. The long stabilization was caused by the electronics in the handle heating the cold junction sensor. The electronics dissipated only a few milliwatts, but that was enough to create a temperature change of more than a degree after power was applied to the catheter.

Because calibration could not be done with the catheter disassembled, there was no choice but to wait for 10 minutes for the catheters to stabilize. The calibration system was designed to hold two to four catheters simultaneously. The power was applied to all of them at once, so up to four catheters stabilized in 10 minutes, yielding a test time just longer than 2.5 minutes per catheter, which was acceptable. Once the temperature was stable, the entire automatic calibration process took about two seconds.

The calibration of the thermocouples was repeatable to ±0.05 °C including all error sources. By calibrating two temperature sensors, the system could be maintained traceable to NIST, meeting the FDA requirements for traceability. The pressure sensors did not need to be calibrated.

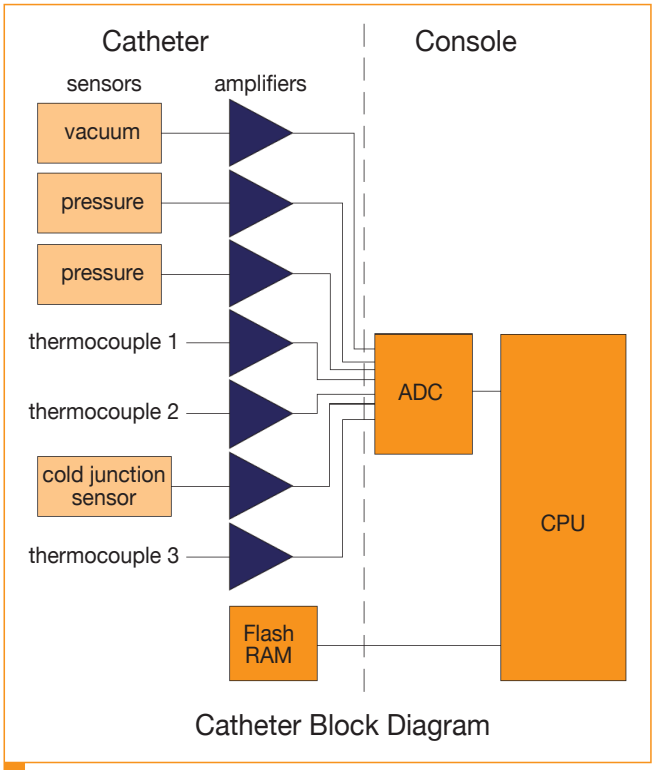

Figure 2 | The electronics that contribute to errors are in both the catheter and the console to which it connects.

Ensuring traceable measurements

In the real world, all measurements are subject to some error. Calibration must achieve the accuracy required with high reliability and repeatability. Calibration of disposable devices is challenging due to accuracy and cost trade-offs. Medical devices must also meet government regulations regarding device traceability and validation. The calibration process must ensure traceable measurements where the start and end of a calibration chain can be followed. ECD

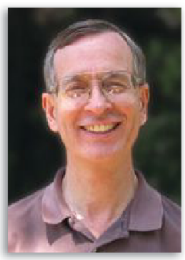

Walt Maclay *is president and chief engineer of Voler Systems and Strawberry Tree, Inc. In 2008, he was elected president of the Professional and Technical Consultants Association. In 2010 he taught at Foothill College as an instructor for the Product Realization Certificate Program. Walt holds a BSEE from Syracuse University.*

Voler Systems
A division of Strawberry Tree, Inc.
walt@volersystems.com • www.volersystems.com

ECD in 2D: Experts discuss the critical issues surrounding embedded medical software design in the webcast "Security in eHealth: Designing Software to Connect Devices Safely." Use your smartphone, scan this code, watch a webcast: http://opsy.st/mKiMgs.

OpenSystems media®

Leadership in the embedded industry

Join the leading print and online publications in the embedded industry

www.opensystemsmedia.com/subscriptions

@embedded_mag

Embedded Computing Design

E-letters

embedded-computing.com/eletter

Print and Digital magazines

embedded-computing.com/emag

TechChannels

tech.opensystemsmedia.com

E-casts and Virtual Events

ecast.opensystemsmedia.com

COM cases show longer medical device life expectancy

By Chiman Patel

Like many other industries, health care has become much more agile as practitioners have tried to keep expenses under control. Smaller, more portable devices now enable more efficient and economical diagnostic and treatment procedures to be delivered at the patient's bedside. Computers-On-Modules (COMs) are helping provide the portability, computing power, and integration needed to increase the longevity of medical device designs.

Electronic device manufacturers are increasingly using Computers-On-Modules (COMs) to produce compact, portable, and easily modified solutions for health care. In addition to smaller solutions, COMs enable longer product life, which is particularly desirable in medical equipment. One reason for this is the stringent and expensive certification process medical devices go through as mandated by the FDA. Furthermore, successful products often become industry standards that medical technicians must learn and adhere to.

By definition, COM designs are modular. The COM approach for embedded technology enables a medical solution's processing power to be easily updated or upgraded while the physical dimensions and mounting system remain exactly the same. This allows the appearance of the industrial design to stay consistent with market expectations while providing for any needed upgrades. If a particular processor is reaching End-Of-Life (EOL), a module with a current processor can be introduced.

With COMs, the processor and supporting features are placed on a compact module that comprises a complete computer. The I/O for the system and the unique features that differentiate the customer's product are included on a separate baseboard. Processor upgrades can be made by utilizing different processing modules that mate with the baseboard. This enables easy processor upgrades, plus various modules can be offered spanning a range of performance levels.

Other benefits of COM use in medical device designs include:

- Fast time time to market
- Simplified development
- Easy modification of popular products
- Easier life-cycle extension

Many medical devices are cart-based to provide point-of-care diagnosis and treatment. COMs with the latest processors and

features such as onboard graphics and GbE LAN capability enable fast transmission to remote terminals or centralized storage. LAN capabilities and standard x86-based solutions help meet the need for information integration with centralized patient records, a key concern in health care.

Medical application examples

COM suppliers can meet the technological needs of medical devices by designing from the ground up, modifying reference designs, and retrofitting newer technology into existing embedded designs to prolong their life. Generally, OEM customers maintain responsibility for the industry certification process. This is because health care certification requires special expertise and is more of a lifeblood issue for health care companies than for those in other industries.

Case study #1: Blood-gas analyzer

A large medical device manufacturer wanted to go to market with a comprehensive yet easy-to-operate blood-gas analyzer. One key goal was long product life. The company selected the COM approach as a way to enable easy processor upgrades and virtually future-proof the product.

The manufacturer chose WIN Enterprises' MB-90140 module (Figure 1) as the core of the hardware. The module features the embedded version of the Intel Celeron M processor to meet production requirements long into the future. The board also offers CRT, LVDS, LAN, and audio support.

Figure 1 | The MB-90140 SOM ETX module, shown here mounted on its baseboard, is equipped with the low-voltage Intel Pentium M/Celeron M processor and offers support for CRT, LVDS, LAN, and audio capabilities.

Case study #2: DNA duplication and research

Another medical device manufacturer needed to update a market-proven solution for DNA duplication and study. Due to the product's existing market acceptance, the manufacturer chose not to redesign its external packaging, but instead, opted to retrofit its core electronics to increase its longevity. In a custom design, WIN Enterprises introduced a new processor and redesigned the circuitry of the old board, reducing the number of boards to two to provide better manufacturing economy.

As in the previous example, the product life cycle was managed through the pre-purchase of the key components over the solution's planned life. Along with the hardware improvements, the manufacturer has continually evolved the product's user interface and software features to maintain its popularity and market position.

> *COM suppliers can meet the technological needs of medical devices by designing from the ground up, modifying reference designs, and retrofitting newer technology into existing embedded designs to prolong their life.*

Case study #3: Laser cosmetic treatment system

A medical OEM marketing a broad line of laser systems used for a range of non-invasive cosmetic treatments needed a custom ETX module that mated with the company's existing baseboard. The ETX module is a variant of the same module used in the blood-gas analyzer application, but with a Pentium M processor. The module supports VGA, LVDS, LAN, SATA, audio, and solid-state disk. As in the first case, a custom baseboard was designed to provide easy future upgrades.

The OEM is currently transitioning to a new design based on the Intel Atom Pineview processor. The goal is to reduce the design impact on the baseboard design during the transition.

Savvy future-proof designs

Although it is one of the few steadily growing industries, health care is still budget-conscious, as are its suppliers. This requires portable, economical, and proven solutions. Once accepted by the health care industry, designs enjoy very long life spans. Medical equipment OEMs must, therefore, select partners that can carefully manage the product life cycle. OEMs and their ODM partners must be savvy enough to future-proof products from the design perspective. The COM design approach helps meet these concerns by enabling existing, proven medical devices to be easily updated. **ECD**

Chiman Patel *is the founder, CEO, and CTO of WIN Enterprises and a pioneer in the design and manufacture of x86-based products for the embedded industry. He holds two patents and is a charter member of TiE-Boston, a nonprofit Massachusetts organization with a mission to foster and support entrepreneurship.*

WIN Enterprises
chimanp@win-ent.com
www.win-ent.com

ECD in 2D: WIN Enterprises' MB-80100 board can be used in medical devices to power high-end graphics and other intense applications. Use your smartphone, scan this code, watch a video: http://opsy.st/mJtmXt.

Editor's Choice

Pico-ITX SBC targets graphics-intensive projects

Digital signage, infotainment systems, and multimedia sales kiosks are examples of projects that require designers to package control functionality into a small form factor, reduced-power embedded device that supports HD video displays. With these applications in mind, Kontron recently announced the new KTA55/pITX, a 2.5" Pico-ITX embedded SBC based on the AMD Embedded G-Series processor. The AMD accelerated processing unit integrates one or two 64-bit processor cores, a programmable graphics unit that supports DirectX 11 and dual-screen 1080p video displays, plus a PCI Express Gen2 controller and a DDR3 RAM controller.

The module features a 100 mm x 72 mm footprint and includes six USB 2.0 ports, an integrated 10/100/1000 Ethernet port, three 6.0 Gbps SATA connections, a low pin count bus for legacy peripherals, and eight configurable GPIO. A trusted platform module provides data security for safety-critical applications. The KTA55/pITX supports Windows XP, Windows XP Embedded, Windows Embedded Standard 7, and Linux.

Kontron | www.kontron.com | www.embedded-computing.com/p52208

Dev kits enable Android developers to accessorize

With Android-based devices now representing a sizable segment of the smartphone universe, there is a huge potential market for hardware accessories, including automotive, home, fitness/health, and business applications. Addressing this new market, Microchip Technology recently announced Accessory Development Starter Kits for Android. Specifically, Android versions 2.3.4, 3.1, and later include a new framework that allows applications to communicate directly with an accessory connected to a smartphone or tablet via USB.

Each kit consists of a development board and software library that enables the development of Android smartphone and tablet accessories based on Microchip's 16- and 32-bit PIC microcontrollers. Example accessory applications include automotive (car kits, audio, GPS), home (audio docks, remote controls, data backup), fitness/health (glucose meters, fitness equipment), and business (credit card terminals, projectors). Microchip's PIC24F Accessory Development Starter Kit for Android is available today for $79.99. The PIC32 Accessory Development Starter Kit for Android is expected to be available in Q3 2011 also for $79.99. The software library, which works with both kits, is currently available via a free download from Microchip.

Microchip Technology | www.microchip.com | www.embedded-computing.com/p52209

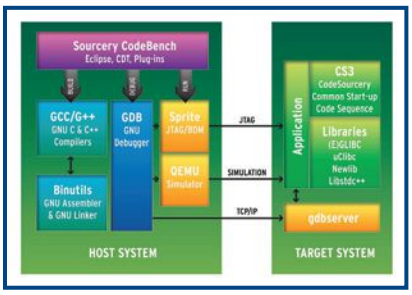

Tool suite optimizes multicore performance

As multicore technology invades the embedded space, development tools vendors are updating their products to enhance the performance of multithreaded code. For example, Mentor Graphics recently announced Mentor Embedded Sourcery CodeBench, a next-generation Integrated Development Environment (IDE) based on the open-source GNU toolchain. Sourcery CodeBench contains all the tools needed to build and debug embedded applications and features an IDE based on Eclipse, the Eclipse C/C++ development tools and compilers, and GNU toolchain, including an assembler, linker, runtime libraries, and source- and assembly-level debuggers.

Sourcery CodeBench includes the Mentor Embedded Sourcery System Analyzer, a specialized tool that helps developers visualize and analyze data to easily locate problem areas and improve design performance. The tool can work with large data sets, integrating and correlating multiple sets of information and creating compelling visualizations. The Mentor Embedded Sourcery CodeBench Personal Edition is available starting at $199, and the Professional Version with System Analyzer is available starting at $2,799.

Mentor Graphics | www.mentor.com | www.embedded-computing.com/p52210

Dual-core processors enhance embedded designs

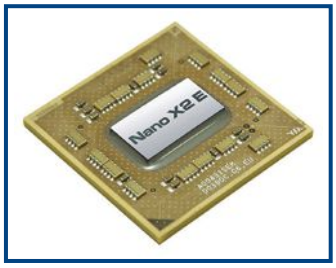

Designers are integrating multicore processors into their embedded devices for improved performance, reduced component count, and lower power requirements. Supporting this trend, VIA Technologies announced a new family of dual-core processors at this year's Embedded Systems Conference Silicon Valley. Targeting x86 embedded system design applications, VIA Nano X2 E-Series processors combine two 64-bit superscalar VIA Nano cores onto one 40 nm die.

Available in two models running at speeds of 1.2+ GHz and 1.6+ GHz, the VIA Nano X2 E-Series processors come with a seven-year component longevity guarantee. These processors bring additional features: VIA VT virtualization to combine applications in virtual scenarios without affecting performance, and the VIA AES Security Engine to deliver hardware-based data encryption for content protection and system security. VIA Nano X2 E-Series processors are compatible with Windows Embedded Standard 7, Windows CE, and Linux.

VIA Technologies | www.via.com.tw | www.embedded-computing.com/p52214

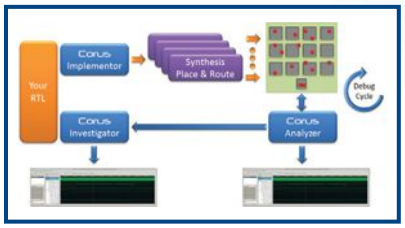

Debug suite validates complex FPGA designs

As FPGA systems increase in both complexity and bandwidth, the test and debug phase has become a development bottleneck. With software that promises to redefine how engineers design for debug, Veridae Systems introduced the Corus validation and debug software suite for complex FPGA-based systems. Corus integrates with popular test equipment and software debuggers to provide a synchronized view of the entire system, including serial I/O, buses, software code, and FPGA hardware, giving designers and validation engineers access to any observed signal at any time.

The validation and debug suite is based on Veridae's set of existing software tools, including the Implementor for designing on-chip signal capture probes, the Analyzer for managing capture data, and the Investigator for interpolating and displaying the signals. The combination of these tools enables faster FPGA-based system validation and debug without continuous resynthesis. Because FPGA products are often embedded into complex systems in the field, remote validation and debug is also supported.

Veridae Systems | www.veridae.com | www.embedded-computing.com/p52212

Evaluation kits simplify ZigBee development

The new ZigBee RF4CE technology specifies a low-data-rate, ultra-low-power wireless network for consumer electronics applications, while ZigBee PRO includes an expanded feature set for application profiles such as automatic meter reading, commercial building automation, and home automation. Supporting this market, Atmel recently announced the ZigBee RF4CE wireless evaluation kit designed for ZigBee PRO and ZigBee RF4CE application development, and prototyping. Atmel is also launching the REB231ED-EK evaluation kit with antenna diversity support for the Atmel AT86RF231 device.

The kits support ZigBee application profiles such as Smart Energy, Home and Building Automation, and Remote Control, as well as IEEE 802.15.4 and proprietary wireless applications. The Atmel ZigBee evaluation kits come with reference boards based on the Atmel ATmega128RFA1 and AT86RF231 wireless low-power transceivers and ZigBee-compliant software stacks ready for download at Atmel's website. The Atmel RF4CE Remote Control Evaluation Kit is available now for $439. The REB231ED-EK kit will be available in Q2 2011 for $239.

Atmel | www.atmel.com | www.embedded-computing.com/p52213

Embedded Linux platform fulfills security criteria

Embedded designs destined for defense applications require secure platforms for military communications such as software-defined radios, command/control ground stations, and combat systems. As security requirements expand into broader market segments, demand increases for security-certified embedded Linux platforms in networking infrastructure, industrial, energy, and medical systems. To participate in these new markets, Wind River announced Linux Secure, asserted to be the first commercial embedded Linux platform to achieve Common Criteria Evaluation Assurance Level (EAL)4+ certification by the National Information Assurance Partnership using the General-Purpose Operating System Protection Profile.

This platform provides organizations facing strict security and cryptography certification requirements with a secure COTS embedded Linux solution. Wind River Linux Secure features Common Criteria EAL4+ certification on ARM including hardware from Texas Instruments, Intel, and Power architectures. Organizations looking to extend this security certification to their own hardware and software environments can also create incremental certifications of Wind River Linux Secure on new COTS or custom hardware and applications. Wind River Linux Secure is based on the stable Linux 2.6.27 kernel and GNU Compiler Collection version 4.3.2.

Wind River | www.windriver.com | www.embedded-computing.com/p52211

Tiny receiver chip eases wireless power charging

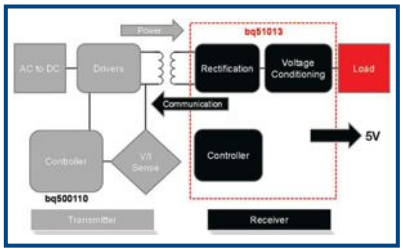

Coming as a relief to portable device manufacturers wanting to eliminate that nightly plug-in ritual of recharging batteries, Texas Instruments (TI) announced its next generation of wireless power technology, the bq51013 wireless power receiver. The new receiver performs all three functions of rectification, voltage conditioning, and digital control and is contained in one small 1.9 mm x 3 mm package, reducing the board area by 80 percent compared to TI's first-generation receiver.

The bq51013 is Wireless Power Consortium Qi compliant, which assures interoperability among various charging pads and portable electronic devices. The receiver features built-in protection against voltage, current, and temperature fault conditions and provides 93 percent peak efficiency, allowing charge rates comparable to an AC adapter. Samples of the bq51013 wireless power receiver are available online. TI also offers the bqTESLA150LP wireless power development kit, which includes the bq51013 wireless power receiver, and the bq500110 wireless power transmitter controller for $299.

Texas Instruments | www.ti.com | www.embedded-computing.com/p52215

E-community Beat

Keeping the embedded conversation going

By Jennifer Hesse

www.embedded-computing.com

BLOGS | MARKET STATS | SURVEYS | VIDEOS | SOCIAL NETWORKING

MCUs lead embedded processor growth

With the proliferation of embedded devices in markets such as data management and storage (think: cloud computing), smart grid/smart home, and the ever-demanding consumer space, embedded CPUs and GPUs will enjoy a unit volume shipment increase CAGR of 21 percent over the 2009-2014 forecast period, according to VDC Research. Embedded MCUs will continue to score the largest unit volume, thanks to their low cost and deployment in automotive and consumer electronics. While DSPs will show the lowest five-year CAGR of 15 percent, they'll retain their position as second most popular embedded processor type.

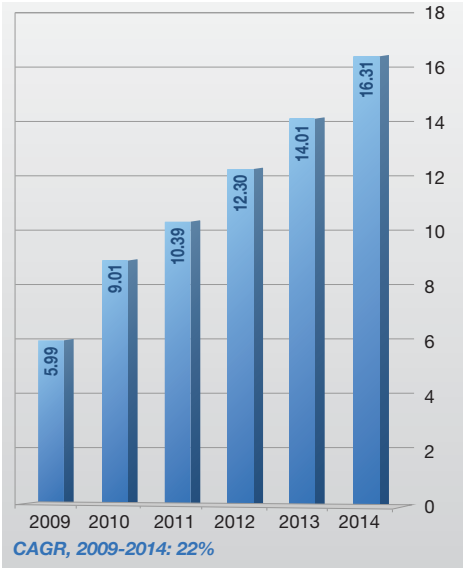

Image courtesy of VDC Research and Power.org.

Total market, embedded processors, unit volume shipments, 2009-2014 (billions of units).

ESC demo video: Blood glucose meter

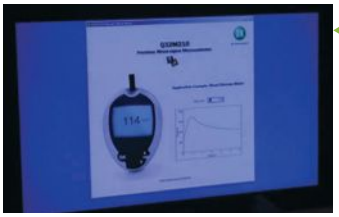

ON Semiconductor demonstrates how the Q32M210 precision mixed-signal microcontroller based on the ARM Cortex-M3 calculates measurements from a blood glucose meter.

See more videos at our ESC TechChannel: http://tech.opensystemsmedia.com/esc.

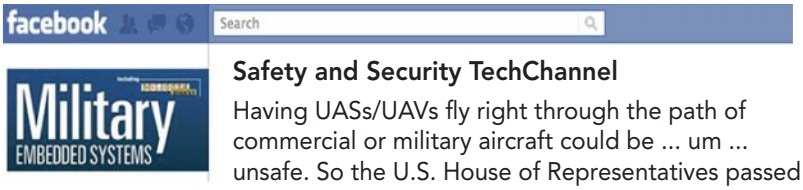

Safety and Security TechChannel

Having UASs/UAVs fly right through the path of commercial or military aircraft could be ... um ... unsafe. So the U.S. House of Representatives passed the FAA Reauthorization and Reform Act (H.R. 658). Among several other imperatives, the Act specifies UAS operation at least five miles from the nearest airport, and daylight operation. Another must: Operation in the operator's line of sight: a necessary no-brainer?

See more ideas at our Safety and Security TechChannel: http://tech.opensystemsmedia.com/safety-and-security.

GROUP: Connected Health Community
DISCUSSION: Incredible growth predicted in global telemedicine

Analysis from BCC Research shows that the telemedicine market will jump from $9.8 billion in 2010 to $23 billion in 2015 for a CAGR of 18.6 percent. Also predicted: Global telehospitals/clinics will notch nearly 17 percent CAGR and telehome will grab 22.5 percent CAGR over that five-year period. Companies like GlobalMedia attending the American Telemedicine Association 2011 meeting and expo confirm that this forecast is on target, given the interest at the show.

Check out other Embedded Computing connections: www.linkedin.com/groups?gid=1802681.

Roving Reporter blog: Virtualization shrinks embedded platforms

By Warren Webb

Embedded designers face the challenge of combining slower legacy interface circuitry with the latest high-speed control devices and multiple displays. The resulting system includes the original hardware with its operating system and application software plus a separate controller with software to handle the newer requirements. This approach increases component count and power requirements and does nothing to increase the performance of the legacy application. With the enhanced processing power of the 2nd generation Intel Core processor family, designers can employ recently updated virtualization software to integrate new requirements with existing embedded applications.

Read more at http://opsy.st/iq3U4M.

INNODISK®
Beyond your imagination

pciedom® II
pcie disk on module

- Smallest PCI Express disk module
- PCI Express X1 & Mini PCIe form factor
- PCI Express Gen.2
- Excellent performance
- Plug & Play, easy installation (patented mechanical design)
- Reliability and optimized lifespan
- Anti-Shock & Vibration
- Built in ECC & Wear leveling architecture

BC AWARD BEST CHOICE 2011

The total solution for Memory and Flash Storage

SATADOM® D150 SATA Module

SATADOM D150QV
- 2GB~ 32GB
- SATA II 3.0G
- Read: 125MB/sec. (max.)
- Write: 120MB/sec. (max.)

SATADOM D150QH
- 2GB~ 32GB
- SATA II 3.0G
- Read: 125MB/sec. (max.)
- Write: 120MB/sec. (max.)

SATA Slim D150
- 2GB~32GB
- Mini PCIe (SATA II)
- Read: 125MB/sec. (max.)
- Write: 120MB/sec. (max.)

mSATA D150
- 2GB~32GB
- SATA II 3.0G
- Read: 125MB/sec. (max.)
- Write:120MB/sec. (max.)

CFast D150Q
- 2GB~ 32GB
- SATA II 3.0G
- Read: 125MB/sec. (max.)
- Write: 120MB/sec. (max.)

SSD Industrial Solid State Disk

2.5" SATA 20000H
- 256GB
- Read:230MB/sec. (max.)
- Write:190MB/sec.(max.)

1.8" SATA 10000-D
- 2GB~32GB
- Read:120MB/sec. (max.)
- Write:100MB/sec.(max.)

i-DIMM® Wide Temp

- Support -40°C ~ 85°C
- 30um golden finger
- ESD and steam proof
- 50% power saving under 85°C ~ 95°C

EverGreen Series Industrial Flash Storage

EverGreen 2.5" SATA SSD
- 16GB~128GB
- Read:220MB/sec. (max.)
- Write:150MB/sec. (max.)

Industrial SD card

www.innodisk.com

InnoDisk USA Corporation
43130 Osgood Road, Fremont, CA 94539 TEL: 510-314-8393 FAX: 510-314-8394 Email:usasales@innodisk.com

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