MIL-EMBEDDED.COM

Military EMBEDDED SYSTEMS

VOLUME 5 NUMBER 4
JUNE 2009

INCLUDING:

Good, good vibrations

Also:
Visual computing made easy

Supplement inside:
FPGAs & Reconfigurable Computing

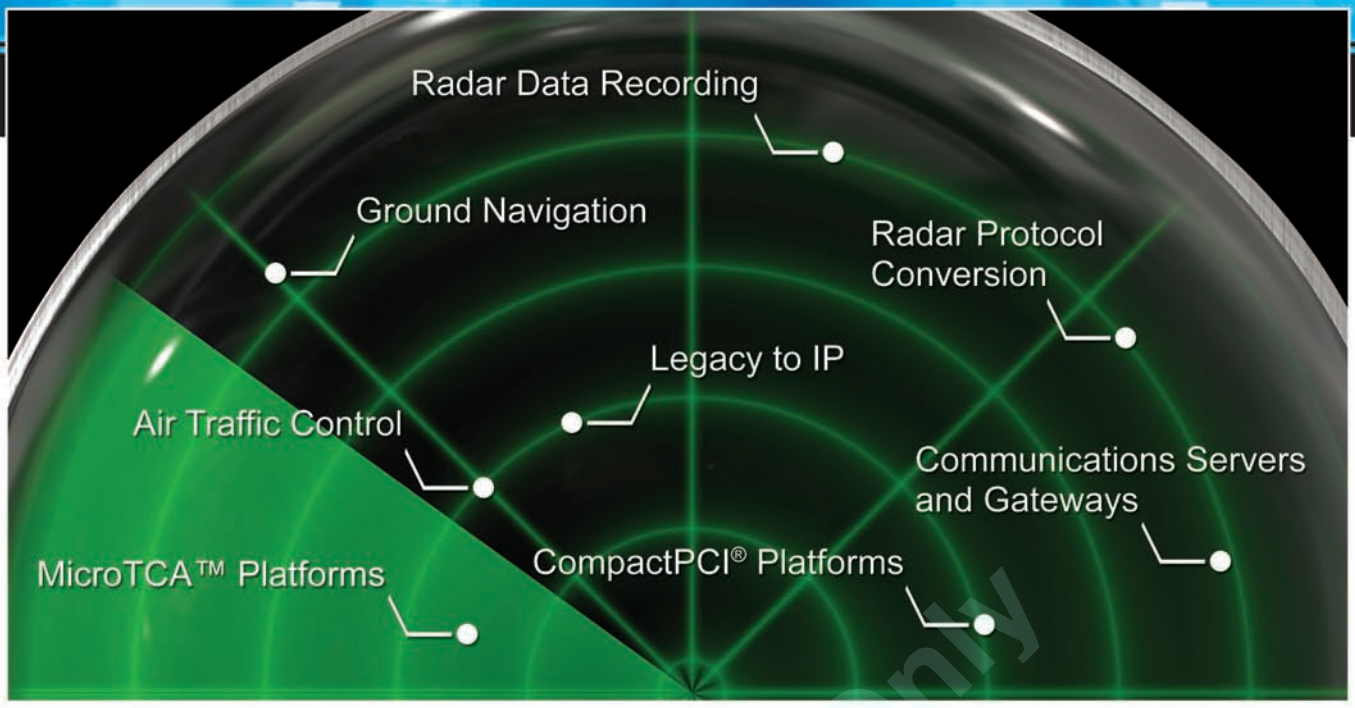

What's on your Radar?

Designing today's aerospace and defense communication systems requires highly reliable and scalable solutions that make the most of next-generation technologies, while protecting investments in legacy infrastructure.

Performance Technologies' COTS-based communication systems are built on our 28-year history of tightly integrated hardware AND software that gets your communications projects up and running fast. From gateways and servers, legacy-to-IP, radar protocols, and more – we provide the platforms and software that allow you to focus on applications development to meet your mission objectives.

When a new communications platform is on your radar, call Performance Technologies.

PERFORMANCE TECHNOLOGIES

Phone: 585.256.0200 • www.pt.com • E-mail: info-request@pt.com

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

Fully Integrated into IBM Blade Management System
Abundant Power and Cooling Ensure Maximum Performance

Made in the USA

Ultimate Modularity

From 2 to 8 Virtex 5 FPGA/Memory Modules
Input / Output Modules Include:
Quad 130 MSps thru Quad 500 MSps A/D
1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10 G Ethernet, FC4, SFPDP

Direct Seamless Connections with no Data Reduction

Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE Backplane
Between FPGAs and Standard Output Modules

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

Military EMBEDDED SYSTEMS

June 2009 Volume 5 Number 4

COLUMNS

DEPARTMENTS

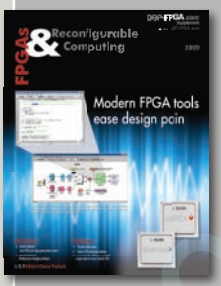

Supplement:
FPGAs & Reconfigurable Computing

ON THE COVER:
The famous Richter scale records seismic activity and is the first thing we think about when quantifying an earthquake or other geologic event. Similarly, shaker tables create earthquake-like vibrations for testing electronics such as those in the 140 M1A1 Abrams tanks destined for the Iraqi Army. But performing X,Y,Z testing one axis at a time isn't the same as simulating a real-world environment. Mechanical testing expert Wayne Tustin tells us why on page 20. (Abrams image courtesy of Pvt. 1st Class Evan Loyd, 2nd BCT PAO, 1st Armd. Div., MND-B)

EVENTS

www.opensystemsmedia.com/events

Sensors Expo & Conference
June 8-10 • Rosemont, IL
www.sensorsexpo.com/sensorsexpo/v42/index.cvn

Flash Memory Summit
Aug. 11-13 • Santa Clara, CA
www.flashmemorysummit.com

Intel Developer Forum
Sept. 22-24 • San Francisco, CA
www.intel.com/idf/us/fall2008/index.htm

WEB RESOURCES

Subscribe to the magazine or E-letter
Live industry news • Submit new products
http://submit.opensystemsmedia.com

White papers:
Read: http://whitepapers.opensystemsmedia.com
Submit: http://submit.opensystemsmedia.com

Published by: OpenSystems media

ISSN: Print 1557-3222

Dream Big...

Think Small.

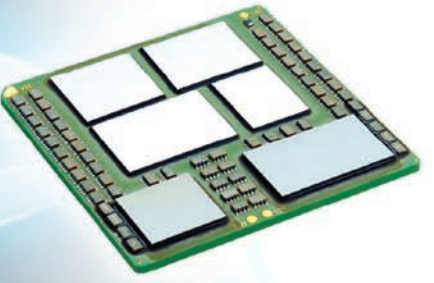

Bring your big ideas down to size with our space-saving semiconductor packaging that makes your board designs possible. Our new DDR2 SDRAM packs 1GB into a SnPb ball attach BGA package less than 1/2 in^2.

WEDC lean multi-chip and system-in-a-package solutions are designed to solve component incompatibilities, reduce design complexity, and extend product life and environmental performance in military applications.

Expand the possibilities;
change your perspective at
www.whiteedc.com/dream3.

602.437.1520 TEL | 602.437.9120 FAX

WHITE ELECTRONIC DESIGNS

www.whiteedc.com

Customer First. Quality Always.

ADVERTISER INFORMATION

Page	Advertiser/Ad title
3	**Annapolis Micro Systems, Inc.** – WILDSTAR 5
22	**Concurrent Technologies, Inc.** – x86 Processor Boards
48	**Curtiss-Wright** – Your data is critical
42	**ELMA Electronic** – Dish us the unusual
11	**Excalibur Systems, Inc.** – Ready for the unexpected?
47	**GE Fanuc Intelligent Platforms, Inc.** – Which of these platforms
24	**General Dynamics C4 Systems** – Airborne network security optimized for the mission
15	**Great River Technology, Inc.** – ARINC 818
17	**Green Hills Software, Inc.** – If it's not INTEGRITY
7	**Jacyl Technology Inc.** – Python X
23	**Parvus Corporation** – Rugged IPv6
2	**Performance Technologies** – What's on your radar?
16	**Phoenix International** – DATA STORAGE
44	**Technobox, Inc.** – XMC and PCIe
16	**TEWS Technologies LLC** – COTS I/O solutions
18	**Themis Computer** – Introducing the first SPARC-Based Blade
23	**Trident Space and Defense** – Protect your data
21	**Tri-M Systems Inc.** – PC/104 Can-Tainer
45	**Tri-M Systems Inc.** – 100Mhz PC/104 Module
10	**VersaLogic Corp.** – All the features you need
5	**White Electronic Designs** – Dream big ...

E-LETTER

June: www.mil-embedded.com/eletter

- Top 5 technologies for the warfighter
  By Chris A. Ciufo
- Ruggedized MicroTCA gears up for deployment
  By Bob Sullivan, Hybricon, and James Doyle, Emerson Network Power
- Portable power management for soldiers: Fuel cell hybrid system is lighter, safer
  By Dr. Peter Podesser, SFC Smart Fuel Cell AG

Military EMBEDDED SYSTEMS

OpenSystems media

Military EMBEDDED SYSTEMS | DSP-FPGA.com | VME Critical Systems | PC/104 and small form factors | INDUSTRIAL EMBEDDED SYSTEMS | CompactPCI and AdvancedTCA Systems | Embedded COMPUTING DESIGN

Military & Aerospace Group

Chris Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

Don Dingee, Contributing Editor
ddingee@opensystemsmedia.com

Jennifer Hesse, Senior Associate Editor
jhesse@opensystemsmedia.com

Sharon Schnakenburg, Associate Editor
sschnakenburg@opensystemsmedia.com

Terri Thorson, Senior Editor (columns)
tthorson@opensystemsmedia.com

Monique DeVoe, Copy Editor

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Steph Sweet, Creative Director

Joann Toth, Senior Designer

David Diomede, Art Director

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

Andrea Stabile
Advertising/Marketing Coordinator
astabile@opensystemsmedia.com

Christine Long, Digital Content Manager
clong@opensystemsmedia.com

International Sales

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

Sam Fan, Account Manager – Asia
sfan@opensystemsmedia.com

Regional Sales Managers

Ernest Godsey, Central and Mountain States
egodsey@opensystemsmedia.com

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 Lamade
800-259-0470
license@opensystemsmedia.com

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

Vice President Marketing & Sales:
Patrick Hopper
phopper@opensystemsmedia.com

Business Manager: Karen Layman

A Ruggedized Computer Designed for Severe Environment Applications

PYTHON X

Designed to Survive

Now Available with a Core i7 Processor

- 19" Rackmount, 1U size
- Available Processors:
  Intel Core 2 Duo
  Intel Core 2 Quad
  Intel Core i7
- Configurable up to 24GB RAM
- All drives are hot swappable
- Designed for minimal weight
- Multiple internal PCIe and PCI slots
- Can be configured for an AC or DC power source
- Wide functional temperature range
- Specialized testing to MIL-STD 810
- Can be custom configured for unique applications
- Also available in a non-ruggedized version

JACYL TECHNOLOGY

Advancing Today's Technology into Tomorrow

www.jacyltechnology.com

DSP library portability speeds application development

By Duncan Young

Many of today's military DSP applications such as radar, sonar, imaging, and signals intelligence make use of high-performance, multicomputing DSP configurations. Often combined with FPGA front-end processing, these applications will typically benefit from 128-bit vector processing engines such as Freescale Semiconductor's AltiVec found in their e600 core-based processing products. DSP libraries are well established for the AltiVec, but as new architectures are explored as potential hosts for newer DSP applications, there is a need for more portability of tools and libraries to enable better evaluation and implementation of these alternatives.

Libraries are an essential tool often delivered as part of a comprehensive DSP development environment. In addition to libraries, such a multicomputing development environment will typically include graphical tools, middleware, and debugging tools. The graphical tool is used for the visualization of software, hardware, and interconnect configurations, enabling performance modeling and the allocation of tasks and processes to processing nodes. Middleware abstracts intimate hardware details provided by the developer, offering a common set of function calls for booting, DMA engines, memory control, and interprocessor communication. Debugging tools can be used during design and development for instrumentation, analysis, and performance evaluation of processing nodes or complete DSP systems.

Understanding open-standard libraries

Vendor-supplied libraries are traditionally targeted to the specific hardware and real-time operating system configurations offered by that vendor. These libraries may be based on an open standard such as the Vector Signal Image Processing Library (VSIPL), or where a vendor is offering unique expertise, libraries might be highly optimized to suit a preferred application area or skill set, such as imaging or software radio. VSIPL provides many generic DSP library functions: matrix operations, FFT, convolution, Finite Impulse Response (FIR) filtering, and correlation for a broad range of commercial and military applications typically including radar and signals intelligence. VSIPL is a library definition that requires vendors to write a port for their particular hardware products. It must also be instrumented and integrated into a development environment to simulate algorithmic performance and to aid in system debugging and verification. Graphical representation and simulation are used to assist in early design tasks including partitioning and assigning functions to processing nodes.

Portability is key

COTS vendors and systems integrators are recognizing that many more processing architectures are candidates for military DSP applications. For example, a number of sonar systems deployed in submarines are being implemented on ruggedized, networked servers or on embedded, off-the-shelf computers based on Intel architecture with Linux or Windows. Having invested in the software development of such DSP systems, integrators are also looking for maximum reuse through portability and scalability for different platform types. COTS vendors are responding by taking the long view, offering generic, processor-agnostic C libraries integrated with their development environments. These libraries support early system architecture design and evaluation independently of the need to choose a particular processor type.

To make meaningful comparative analyses, these C libraries must be ported to specific hardware implementations and operating systems. Additionally, the development environment must graphically represent these new implementations and support simulated performance and assignment of processes between multiple, possibly disparate, nodes. These capabilities are provided by GE Fanuc Intelligent Platforms' AXIS DSP development environment, which offers highly optimized runtime libraries for AltiVec-based architectures and generic library support for Intel architectures (Figure 1). It is offered in desktop, laptop, or embedded form and supports non-AltiVec PowerPC devices and cores.

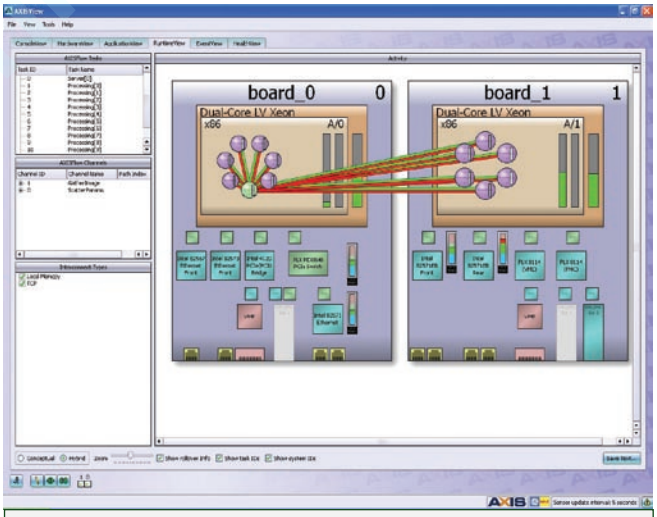

FIGURE 1: A screen capture shows the AXIS development environment, graphically depicting two Intel Xeon-based processing nodes.

Enabling the evaluation and migration of more portable DSP applications to next-gen processing devices is becoming a key part of COTS vendors' strategies. Providing generic libraries allows alternative hardware platforms to be compared, leaving the final choice until much later in the design process. Further performance optimizations, such as Intel's SSE4 extensions, can be incorporated later at runtime. Being able to start the design process earlier in the cycle will speed overall project development by delaying the all-important hardware platform choice until the system has been fully architected and partitioned.

To learn more, e-mail Duncan Young at young.duncan1@btinternet.com.

COTS steps up to embedded computing security

By John Wemekamp

Armed forces worldwide continue to be concerned for their security as they become increasingly dependent on the use of high technology in critical war-fighting arenas. The network-enabled battlefield and Global Information Grid (GIG) have introduced the need for more vigilant network security. This, in turn, has motivated the move to IPv6, with all implementations required to support Authentication Headers (AHs) and Encapsulating Security Payloads (ESPs). However, this is only one piece of the puzzle. Any situation where high-technology equipment, from manpack radios to armored vehicles or combat aircraft, is deployed on the front line poses a risk if the technology unintentionally falls into the wrong hands. Technology can be reverse engineered to discover its operation or to replicate it. In addition, unraveling its firmware and software code to reveal its algorithmic and performance data might also expose potentially compromising information such as orders, codes, and keys of much wider significance.

Protection of critical technology is now considered to be an essential element in the overall system design and procurement processes of new equipment. Identification of Critical Program Information (CPI), whether it is data or technology based, starts at the system level. It is then decomposed to application software, operating system, firmware, communications, subsystems, chassis, modules, deployment, operation, and maintenance down to the lowest level. New programs must identify CPI and develop a comprehensive protection plan to mitigate every aspect that could compromise its technical superiority if it were to be reverse engineered – or if details of its performance or operation were to be revealed.

Layered approach is best

Typical good practice means taking a layered approach to providing protection. In the context of an embedded system, this would start at the chassis or enclosure level with volume protection. Volume protection is designed to prevent physical intrusion, or if prevention is not possible, then detection of intrusion must trigger irreversible indications. Equipment is often recovered or returned at a later date, so understanding the depth of intrusion is of vital interest. The simplest example of this form of detection is a *warranty seal* to indicate that the enclosure has been opened. However, this is obviously not enough, and further active and passive measures must be taken for greater assurance. External I/O connections must also be considered, for example, debug ports, network connections, or special-to-type I/O. These could reveal critical performance parameters or provide access via potentially vulnerable information.

Additional stages of hardware protection might be necessary to protect FPGA or nonvolatile memory contents. Often electronic assemblies in transit for maintenance or repair, or in an emergency, will have had their nonvolatile memory erased. This basic feature has been incorporated by many COTS vendors into their products offered for military applications. However, erasure is not always practical for FPGAs or other types of programmable devices. Thus, other techniques might warrant consideration, such as encrypting or obfuscating the contents. Similarly, software and firmware require protection from the lowest levels of code storage, retrieval, and execution through I/O, communications, and the application layers.

Custom versus COTS in the military

The traditional method for the military to protect their critical technology leadership has been to employ custom hardware and software design. Each project has its own specific security requirements; therefore, off-the-shelf embedded computing equipment is perceived as unlikely to incorporate entirely. But as technology becomes more widely diffused across the battlefield, even as far as fleets of trucks or handheld devices for individual soldiers, the cost of this purely custom approach is escalating. Through much greater understanding of the issues, plus dialog with end users and integrators, COTS vendors are able to bridge the cost/capability gap between fully custom and off-the-shelf products. This gap is narrowed when COTS vendors introduce key protection capabilities into their products that meet many critical program requirements. Features such as active perimeter defense, a secure computing environment, and standards-based, NIST-approved cryptographic engines are examples of how COTS vendors like Curtiss-Wright Controls Embedded Computing (CWCEC) are addressing these security needs (Figure 1).

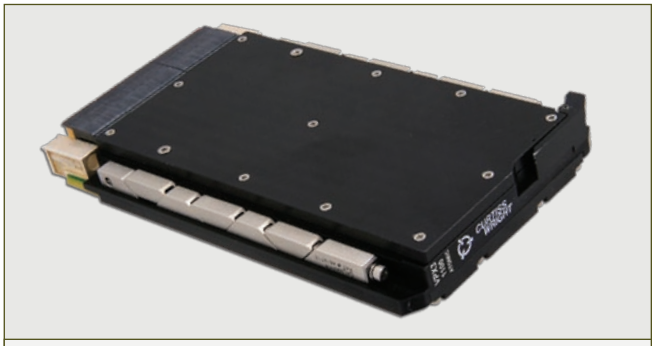

FIGURE 1: Curtiss-Wright Controls Embedded Computing's VPX3-1100 Atomic SBC is designed to address today's security needs via its active perimeter defense, secure computing environment, and standards-based, NIST-approved cryptographic engines.

Including security features early in systems engineering and development will reduce overall costs of implementation and eventual retrofit following deployment. At the higher layers, requirements and solutions to those requirements will vary significantly from project to project. However, for embedded computing hardware, it is possible to take a more generic approach, one that provides a comprehensive set of secure enabling technologies that can be used intelligently to create a safer environment to support those very specific outer layers.

To learn more, e-mail John at john.wemekamp@curtisswright.com.

Legacy Software Migration

By Jacques Brygier

Extending military software life expectancy through safe and secure virtualization

Building new software is extremely expensive and time consuming, so extending the useful lifetime of existing software is critical in times of tight budgets and lengthy budget cycles. The spectrum of available techniques for software life extension includes reuse at the code, design, or specification level where at least some aspects of the existing system can be salvaged and adapted to new hardware platforms and new requirements. Yet each of these paths requires costly manual adaptation and might not adequately retain hard-fought assurances for real-time safety and security management. Even traditional virtualization might not be adequate when real-time behavior is a concern. However, emerging Safe and Secure Virtualization (SSV) technology is meeting the challenge and providing the best of all worlds.

An ideal life extension technique would be one where entire subsystems are retained from the original platform and inserted, unchanged, into the new platform. Such an approach is now possible with an emerging technology called Safe and Secure Virtualization (SSV). SSV is now finding favor in the planning and development of next-generation systems where safety and security are as important as economy and timeliness.

SSV and legacy subsystems

SSV as a method of legacy reuse is not a fantasy of future technology, but rather it is supported with COTS technology already in use. An example of SSV adoption is the recent selection of SYSGO's PikeOS by Airbus for their A350 XWB aircraft. Among the many requirements for the new Airbus architecture is the ability to develop certifiably safe software in a native partition, while easily coexisting with legacy software subsystems in POSIX partitions.

Accordingly, SSV provides the ability to comingle software subsystems culled from a variety of legacy platforms, even including off-the-shelf open source subsystems. Because real-time behaviors can sometimes present problems when utilizing classic virtualization techniques, SSV builds upon proven ideas of virtualization,

such as the use of paravirtualization to adapt a hosted operating system to interact with the hypervisor layer of the hosting operating system. However, SSV goes further in providing deterministic behavior and resource partitioning to enable the development of embedded applications with predictable real-time response and to provide the basis for certifying independent levels of both safe and secure operations. More specifically, SSV integrates the separation kernel model that includes a low-level scheduler that guarantees the overall system preemption ability while still allowing the different scheduling policies in partitions to behave as originally designed.

Therefore, by definition with SSV, subsystems can feature different levels of safety and security without conflict or risk. The SSV real-time OS is based on a MILS-conformant microkernel that supervises every hardware access. A unique feature of the SSV architecture is the ability to classify all resources according to requirement subsets. Different operating systems, Runtime Environments (RTEs), and APIs are able to run simultaneously within protected software partitions. This is made possible through the use of accurate, manageable communication channels that allow both safety-critical and noncritical applications to coexist within a single hardware environment.

Reusing legacy code

SSV is an ultimate resolution for integrating otherwise obsolete software in a modern, high-capacity embedded system by isolating separate partitions. The advantage of this approach is the ability to reapply already-existing legacy code that does not have to undergo costly, time-consuming, and error-prone redevelopment phases. The established software is able to operate on a new hardware platform intermingled with other, newer software components such as a modern Linux OS. The modularity and independence of the separate systems allow this peaceful coexistence and cooperation to occur.

Isolating and encapsulating different software packages enable resource partitioning and time partitioning: A static assignment of all available and temporary resources takes place. Each application obtains guaranteed access to assigned resources but does not have any access to other partitions' resources. Strict separation enforcement guarantees that failures in one partition will not affect other partitions, thus ensuring safe and secure operation.

As an example, SSV allows PikeOS to run a Linux-based subsystem and a safety-critical application with its own proprietary operating system on a single CPU platform. All partitions run in user mode and do not influence the stable kernel mode. Many OS or RTE personalities are available in this environment. These include those based on POSIX, Ada, and Linux, providing the developer the ability to cleanly adopt legacy code into next-generation systems.

Many techniques have been defined and implemented to support code reuse in order to reduce development cost and hopefully minimize risks. Intuitively, traditional virtualization provides a very valid concept to mix legacy software with new software, but falls short for many embedded applications, particularly when real-time behavior is involved. Furthermore, when safety and security are mandatory requirements, then innovative technology based on the SSV concept is essential.

Jacques Brygier *is vice president of marketing at SYSGO. He has spent more than 20 years in the business of high technology and computer science, where he has acquired in-depth knowledge of the software industry, its evolution, and its main application fields. He holds a Ph.D. in Computer Science. He can be reached at jacques.brygier@sysgo.fr.*

Daily Briefing: News Snippets

By Sharon Schnakenburg, Associate Editor

www.mil-embedded.com/dailybriefing

Semiconductor sales: A matter of perspective?

The Semiconductor Industry Association's (SIA's) most recent financial report ostensibly serves to validate both the optimist and the skeptic. The good news is that March 2009 worldwide semiconductor sales climbed 3.3 percent as compared to February 2009 (see Table 1). Additionally, except for Japan, all worldwide regions reported semiconductor sales gains month-to-month (see again Table 1). However, first quarter 2009 chip sales of $44.0 billion reflect a 29.9 percent drop versus first quarter 2008's sales of $62.8 billion. According to George Scalise, SIA president, "While all major product sectors showed month-to-month growth, there continues to be limited visibility in end markets. There are some bright spots such as 'smart phones' and 'netbook' PCs, but there are no clear signs of early firming of demand … The global chip industry continues to reflect the influence of the worldwide economic slowdown." Scalise anticipates that future economic stimulus measures will "impact sales" toward the beginning of 2010.

Month-to-month sales in billions of dollars			
Market	**February 2009**	**March 2009**	**% Change**
Americas	2.46	2.58	5.1%
Europe	2.12	2.18	3.1%
Japan	2.79	2.52	-9.4%
Asia Pacific	6.85	7.38	7.8%
Total	**14.21**	**14.67**	**3.3%**

Table 1: February vs. March 2009 semiconductor sales. Source: SIA

U.S. Navy steers former Detroit autoworkers in new direction

Reports of mass layoffs in the U.S. and around the world hit airwaves and newspapers nearly every day, and Detroit autoworkers are certainly among the hardest-hit. However, Naval Sea Systems Command (NAVSEA) is now steering laid-off "mid-career automotive engineering professionals" into NAVSEA jobs located at NAVSEA headquarters, Navy-sponsored University Affiliated Research Centers (UARCs) nationwide, naval warfare centers, and program executive offices. "They come from diverse backgrounds and can help us think in different ways in terms of manufacturing: Lean initiatives, production line techniques, and modern design," NAVSEA Commander Vice Adm. Kevin McCoy explained in a statement to the media. A recent Livonia, Michigan NAVSEA career fair yielded about 240 qualified resumes from those wanting to make the transition from engineering autos to maintaining and developing systems and ships for the U.S. Navy.

VPX remains poised for call to duty

It's only *human* to falter under pressure. However, VPX technology, thankfully, does not suffer from the same limitations as humankind. Accordingly, a recent contract stipulates that a VPX-centric rugged signal processor subsystem be provided by GE Fanuc to the Turkish Armed Forces for inclusion in its Synthetic Aperture Radio (SAR) system. VPX was reportedly the architecture of choice because of its extreme rugged reliability and high-bandwidth processing capabilities. VPX progeny within GE's unnamed subsystem include the XMCV5 VPX FPGA-based processor, along with the DSP230 VPX multiprocessor computer. Because of immediate availability of the XMCV5 and DSP230, prime contractor Aselsan is slated to start immediate prototyping and test and validation. XMCV5 utilizes Xilinx's Virtex-5 FPGA family, while the 6U DSP230 includes four Freescale 8641D dual-core processor nodes.

Built-in technology for ground soldiers

Knowledge is power, and the recent first-phase Ground Soldier Ensemble contract between the U.S. Army and General Dynamics equips soldiers with a SWaP-efficient, built-in information system. The Ground Soldier Ensemble facilitates Army leaders' rapid access to critical information relative to their surroundings and each other – and provides small groups the ability to manipulate and view mission-planning information. A headset for digital and audio communications, radios, and the software and power source for the Army's computer applications are included. Delivery of 10 prototypes is anticipated by September.

Wind River seeks NSA's validation

It's highly coveted amongst embedded software vendors … Green Hills Software recently reached it … Wind River is striving to achieve it. What is it? EAL 6+ certification for High Robustness in conformance with Common Criteria. Case in point: Wind River's VxWorks MILS Platform 2.0, a development environment for designing Multilevel Secure (MLS) systems, was recently announced as "in evaluation" for EAL 6+. Headed by the NSA and in conjunction with the National Information Assurance Partnership (NIAP), EAL 6+ is designed to ensure an extremely high level of trust among product users, who then know that the high-value or classified information in their networks and information systems is secure against hackers, intruders, terrorists, and the like. The highest possible level of EAL certification is 7.

If the shoe fits, track it

Taking wearable technology to a new level, a recent DARPA contract has motion technology provider InterSense collaborating with Case Western Reserve University on DARPA's Micro Inertial Navigation Technology (MINT) program. MINT's objective is production of a tracking system geared toward commercial and military use that can render extremely accurate information on positioning, whether or not GPS is additionally available. MINT combines Case Western Reserve's unnamed high-resolution ground reaction sensor cluster and InterSense's unnamed micro-inertial navigation system based on its NavShoe boot-mounted sensor technology. NavShoe's filter algorithms constantly correct for drift, resulting in accuracy of less than 1 percent of traveling distance. Meanwhile, dead reckoning sensors typically only provide accuracy within 2 to 5 percent of traveling distance, InterSense reports.

Phantom Ray technology is … reality

Though its government funding could be described as "phantom-like," the technology behind Boeing's internally funded Phantom Ray is anything but elusive or illusion. The rationale: The Phantom Ray unmanned flying test bed, designed for air system technologies, will be based on the prototype vehicle originally developed by Boeing for the U.S. Air Force/DARPA/U.S. Navy Joint-Unmanned Combat Air System (J-UCAS) program. Resuming where UCAS stopped in 2006, Boeing will base Phantom Ray on the X-45C's airframe design. The Phantom Ray demonstrator's inaugural flight is slated for December 2010, followed by 10 additional flights within about six months. Possible mission genres for Phantom Ray include ISR, electronic attack, enemy air defense suppression, and autonomous aerial refueling, among others.

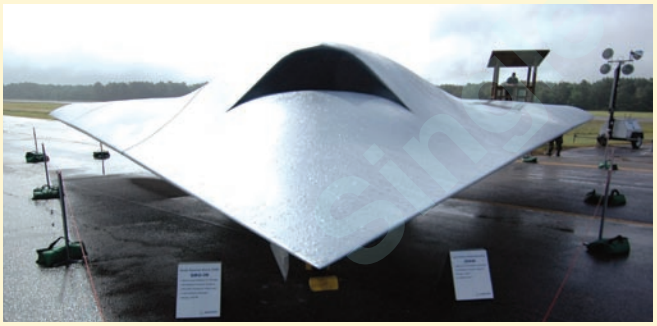

U.S. Navy photo of the X-45C by Photographer's Mate 2nd Class Daniel J. McLain

Have ultrasound, will travel

Apparently the DoD took notice when ultrasound imaging systems developer/manufacturer Imperium claimed its Acoustocam is "the world's only portable ultrasound imaging camera." Under the DoD's Small Business Innovation Research (SBIR) program and with the goal of meeting NAVAIR's objective to rapidly evaluate flight-critical structures for damage that's "non-visible," Imperium will develop two versions of Acoustocam: a through-transmission iteration that identifies skin-to-core disbonds that can occur in "honeycomb" structures like helicopter rotor blades; and a pulse-echo version that finds defects including disbonds, delaminations, and corrosion in metallic plate- or sheet-like structures in cases where both sides of the structure are not visible.

New U.S. Army helicopter: Duplicity?

Some might say that the U.S. Army's recently announced Armed Scout 645 helicopter bears a striking resemblance to the Army's UH-72A Lakota – but the similarity is primarily skin deep according to prime EADS North America. Though based on the Eurocopter EC145 commercial airframe that spawned the UH-72A, Armed Scout 645 will offer its own distinctive features: twin-engine performance geared for high-altitude and extended-temp environments, along with a modular weapons system and "modern design." Armed Scout 645's small footprint will also enable easy transport on C-17s. Meawhile, Armed Scout 645 is touted to provide "maximum mission flexibility": the byproduct of its unobstructed main cabin, which can be easily reconfigured.

Artist rendering of the Armed Scout 645 courtesy of EADS North America

Mission communications – no limits?

Apparently the *sky's* the limit – or maybe that's the *stratosphere* – for Quintron Systems, Inc. Within the past three months, the communications/VoIP producer was awarded *four* contract expansions, one of which calls for a new Atlas 5 rocket communications system and was awarded by the United Launch Alliance (ULA). The ULA expansion stipulates that a DICES VoIP mission voice system renders ancillary audio to ULA managers and engineers not stationed at the launch control center, which already utilizes DICES. The new expansion-provided DICES will interface with existing LAN/WAN corporate infrastructure. Meanwhile, one of Quintron's other contract expansions includes providing an additional DICES system to support U.S. Army and NATO training in Germany; and two DISA/USSTRATCOM-authorized expansions were also received: a 33 percent system capacity increase for a DICES IV system located at Offutt AFB in Nebraska, along with a 50 percent increase in audio connections at an Air National Guard locale in Idaho.

Atlas 5 photo courtesy of Lockheed Martin

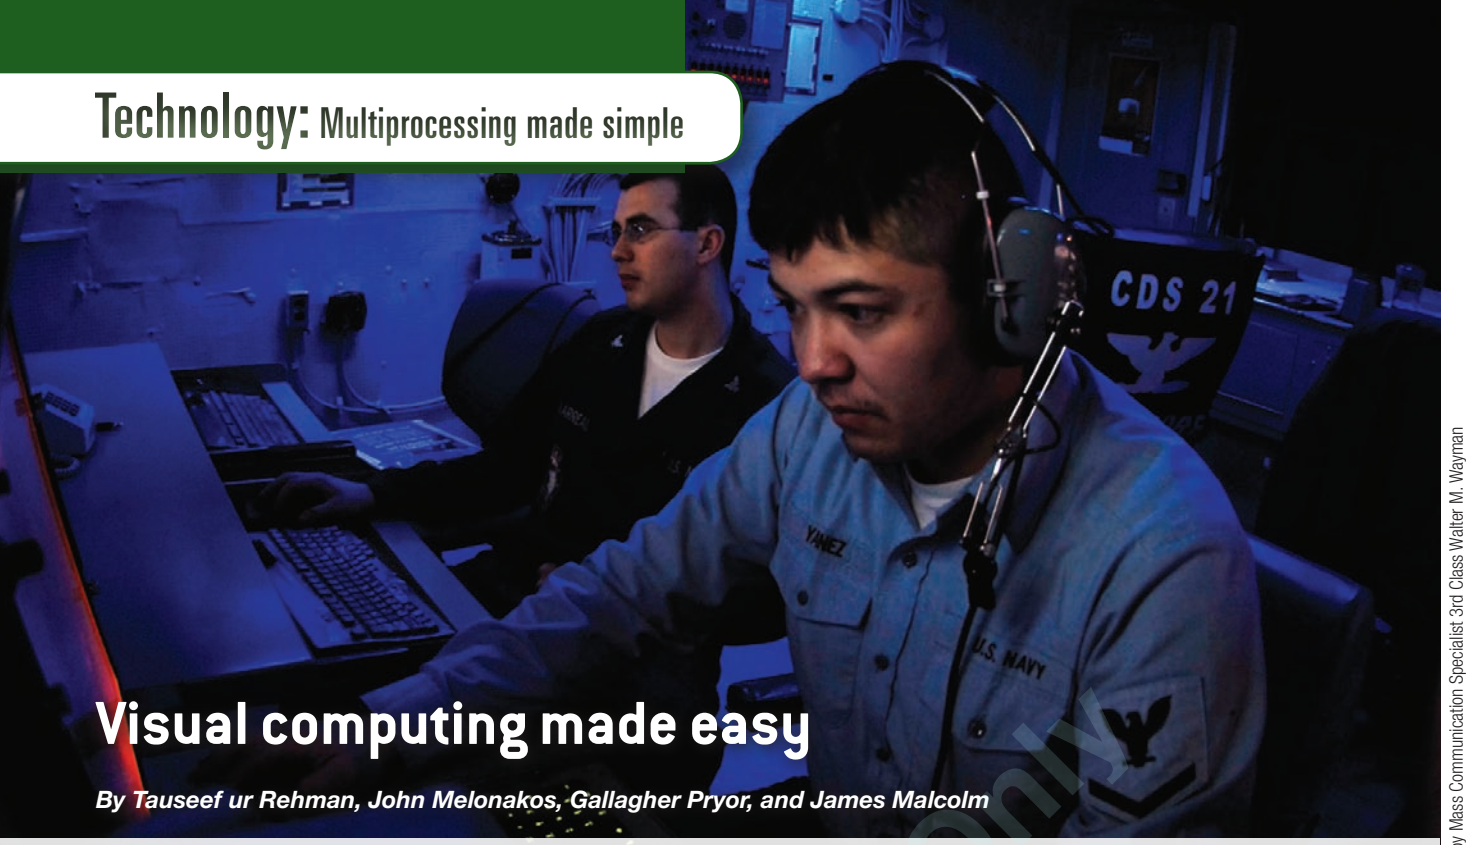

U.S. Navy photo by Mass Communication Specialist 3rd Class Walter M. Wayman

Visual computing made easy

By Tauseef ur Rehman, John Melonakos, Gallagher Pryor, and James Malcolm

We are at the dawn of a new and exciting era in computing called "visual computing." It has the potential to fundamentally change the way data is processed and visualized in design and engineering modern military systems. With the availability of Jacket, a new programming tool, it is surprisingly easy to make use of this new technology in complex computer simulations.

Recent developments associated with the emergence of multi-core architectures are transforming the field of High Performance Computing (HPC) in a disruptive way. The forerunner technology leading this transformation is the Graphics Processing Unit (GPU). These massively parallel processors were originally developed to perform complex graphics calculations, but they can also be used for general-purpose computations. GPUs offer up to 240 computational cores, in contrast to the 4 cores currently available on CPUs. Due to this performance discrepancy, a new era of "visual computing" is emerging in which GPUs are being used to increase application productivity and quality through the merger of computation with visualization.

These recent developments in the HPC market are not going unnoticed by the defense industry, which has always been an early adopter of emerging technologies. Major companies and governmental organizations in the industry have already started working with GPUs and are discovering the potential for significant acceleration in a wide range of applications ranging from signal and image processing to distributed control of multi-agent systems. However, as is common for new technologies, the development tools for GPU computing have not yet matured. GPU manufacturers and several third-party providers are now offering a variety of Application Programming Interfaces (APIs) to program these devices. Although these tools improve the accessibility to these new computing platforms, they still inherit the fundamental difficulties associated with parallel programming: switching from the sequential to parallel way of writing programs, as well as the steep learning curve associated with available software development tools.

One particular programming language that has the potential to bridge this gap between parallelizable applications and GPUs is the M-language. It is available in MATLAB, a very popular Integrated Development Environment (IDE) for prototyping algorithms and complex simulations. Available since earlier this year, Jacket, a new add-on toolbox to MATLAB, now enables standard M code to run on GPU-based accelerators from within MATLAB. A user can get started with GPU computing almost instantly with a negligible learning curve involved. The authors explore the nuances of M-language, then present a volumetric control simulation example to illustrate the ease of use and code reusability afforded with Jacket.

Visual computing and new challenges

Visual computing is the union of visualization and mathematics. It allows users to interact with data and visually explore results. Visual computing is important to numerous fields ranging from defense and industrial applications to games, sports, and more (Figure 1). Engineering problems are increasingly becoming dependent on larger amounts of data. Consequently, new

FIGURE 1: Visual computing is important to numerous fields ranging from defense and industrial applications to games, sports, and more.

challenges arise both in terms of computational power required to process in reasonable time as well as effective visualization of results. In performing computer simulations of systems and processes, the results are generally visualized after the computation has completed to avoid considerable computational overhead of visualization code. Now with the GPU as the computational workhorse, seamless integration of visualization and computation becomes possible where users can visually explore data and results in real time as they get computed in a simulation. Therefore, large-scale visualization is an ideal fit for GPU computing for several reasons:

- Visualization is a data-intensive application, and GPUs are well suited for data-intensive tasks.
- Visualization computations exhibit substantial parallelism.
- Visualization tasks should be closely coupled to the graphics system.
- Integrating computations with visualization on the GPU eliminates huge overhead of memory transfer between CPU and GPU, permitting entire applications to run faster.

GPU manufacturers are focused on developing programming tools to harness the incredible power of these processors such as NVIDIA's CUDA (Compute Unified Device Architecture) and AMD's CTM (Close To the Metal). These tools, though powerful, require learning new APIs as well as mastering the ability to write parallel code. Moreover, optimizing code for performance is also nontrivial, requiring a deeper knowledge of underlying hardware architectures. Therefore, despite the promise of significant performance benefits being reported in research literature, adoption of GPU computing has been below expectations.

M-language to the rescue

M-language allows complex systems to be represented in a concise manner and provides sophisticated visualization and analysis functions as compared to C/C++ with OpenGL visualization. The key to its growing popularity is that nonprogrammers find it relatively easy to learn. It is a data-parallel language and offers an easy way for scientists working with vectors and matrices to express element-wise parallelism on aggregate data. This feature combined with the ubiquity of MATLAB makes it an ideal candidate for mapping to fine-grained, data-parallel systems like GPUs. Moreover, from a scientist's perspective, minimal effort is needed to achieve basic visualizations. For example, the following code snippet shows M code that visualizes a signal as a surface while it is being manipulated within a loop.

```
A = zeros(1000,1000); % Allocate CPU memory
for  t = 1:100
    A  =  exp(-t*i) * A;  % manipulate
    B  =  ifft(A);        % inverse FFT
    surf(abs(B));         % visualize (transfer to video memory)
end
```

In contrast, the same function if written in C/C++ with OpenGL visualization would span, at minimum, hundreds of lines. Though concise and easy to learn, M-language also has its restrictions that become quickly noticeable in the form of slower computational speeds as the data size increases. There are solutions to this in the form of distributed/cluster computing, but their cost and complexity do not make sense for most applications.

Jacket, an add-on toolbox to MATLAB, can make up for these deficiencies by providing a trivial way to offload computations to the GPU, resulting in code speedups of 10-100x for many applications. Users can simply mark data, via casting operations, to indicate that computations should occur on the GPU. The following code shows how simple it is to achieve this.

```
A = gzeros(1000,1000);      % Allocate GPU memory
for  t = gsingle(1:100)     % Casting of iterator vector
   A  =  exp(-t*i) * A;     % manipulate
   B  =  ifft(A);           % inverse FFT
   surf(abs(B));            % visualize (no memory transfer)
end
```

Jacket also includes a Graphics Toolbox that enables seamless integration of computation with visualization, making difficult-to-program, multithreaded, and real-time graphical displays simple to achieve. For example, by placing a single visualization command at the end of a loop as shown in the code snippet above, data may be viewed as it is processed in place on the GPU. Load-balancing decisions are automatically made to optimally use GPU resources for compute as well as display. Further, the Graphics Toolbox exposes the entire OpenGL API and allows for interactive scene creation and rapid prototyping.

CPU Code

```
%% Swarm Simulation Loop

tic %Start Timer
initial_conditions = single(initial_conditions);
for i=1:length(tt)-1
    [initial_conditions] = discrete_swarm(initial_conditions,param);
    initial_conditions(:,3:4) = initial_conditions(:,3:4)*0.980;
end
toc %Stop Timer
```

GPU Code

```
%% Swarm Simulation Loop

tic %Start Timer
initial_conditions = gsingle(initial_conditions);
for i=1:length(tt)-1
    [initial_conditions] = discrete_swarm(initial_conditions,param);
    initial_conditions(:,3:4) = initial_conditions(:,3:4)*0.980;
end
toc %Stop Timer
```

FIGURE 2: Jacket is used to offload the computation and visualization to the GPU by maximally reusing existing CPU M code.

A volumetric control simulation example

Although the applications are numerous, the code reuseability and overall ease-of-use benefits of Jacket for a complex simulation scenario can be appreciated by an example from volumetric control. "Volumetric control" means monitoring, detecting, tracking, reporting, and responding to environmental conditions within a specified physical region. This is done in a distributed manner by deploying numerous vehicles, each carrying one or more sensors, to collect, aggregate, and fuse distributed data into a tactical assessment. Examples of this include surveillance and perimeter defense using multiple agents. Jacket is used in the code segment shown in Figure 2 to offload the computation and

If it's not INTEGRITY®, it's not secure

The only operating system certified to EAL6+ High Robustness.

INTEGRITY
The only secure operating system

Green Hills® SOFTWARE, INC.

www.ghs.com

visualization to the GPU by maximally reusing existing CPU M code.

Notice that a simple casting operation (outlined in red) is the only change between the two code segments. The result was a speedup of 8x using an NVIDIA Tesla C1060 GPU when compared to a 2.2 GHz dual core Intel CPU for a swarm (group of agents) size of 400. Another point to note is that the function `discrete_swarm()`, which models the swarm dynamics, is completely reused in the GPU version with no change at all. Moreover, with Jacket the simulation updates can be visualized live without incurring the huge overhead of memory transfers to video RAM as in the case of CPU code.

Simple software, powerful visual computing

Application developers designing high-performance computational systems are increasingly embracing GPU computing to enhance application speed and data visualization. Visualization is a critical step in the process of transforming data into knowledge. An integrated visualization capability alongside compute provides scientists and engineers a crucial capability when running complex simulations. Jacket makes this capability available to the technical computing community already familiar with the MATLAB programming language without any pains of rewrites or learning new APIs. For widespread multicore adoption, it is imperative that both scientists and engineers who do not want to be computer scientists continue to use familiar languages and constructs and focus on their algorithms rather than on implementation details.

Tauseef ur Rehman *is cofounder of AccelerEyes LLC, working on applications of GPU computing for scientific and engineering problems using Jacket. He is an expert on numerical methods for computer vision and holds an MS in Electrical and Computer Engineering from Georgia Institute of Technology. He can be reached at tauseef.rehman@accelereyes.com.*

John Melonakos *is cofounder and CEO of AccelerEyes LLC. He is an expert on medical image processing and holds a PhD in Electrical and Computer Engineering from Georgia Institute of Technology. He can be reached at john.melonakos@accelereyes.com.*

Gallagher Pryor *is cofounder of AccelerEyes LLC, providing strategic technological vision for current and future products. He currently heads the visualization division, which is developing tools for scientific visualization. He holds a BS in Computer Science from Georgia Institute of Technology. He can be reached at gallagher.pryor@accelereyes.com.*

James Malcolm *is cofounder of AccelerEyes LLC and chief architect of the Jacket project, where he is responsible for software strategy and development. He holds degrees in Mathematics (BS), Computer Science (BS, MS), and Electrical Engineering (MS) from Georgia Institute of Technology. He can be reached at malcolm@accelereyes.com.*

AccelerEyes LLC
770-500-2101
www.accelereyes.com

Multicore, Multiprocessor, and Multichip product showcase

Editor's note: *Multicore, multiprocessor, and multichip applications are becoming ever-more prevalent on the battlefield because of their well-known, distinctive SWaP advantages. Consequently, they're also appearing more often in magazines like the one you're holding now. Check out our inaugural "multi" product showcase, in addition to these related articles:*

- ***Top 5 technologies for the warfighter,*** by Chris A. Ciufo, Editor
  www.mil-embedded.com/articles/id/?3960
- ***Migrating legacy applications to multicore: Not as scary as it sounds,*** by Bill Graham, QNX Software Systems
  www.mil-embedded.com/articles/id/?3820
- ***MicroTCA's role expands in modern battlefields,*** by David Pursley, Kontron
  www.mil-embedded.com/articles/id/?3829
- ***Maximizing GPGPU computing for embedded systems,*** by Alan Commike, Quantum3D, Inc.
  www.mil-embedded.com/articles/id/?3742

CURTISS WRIGHT Controls Embedded Computing

VPX6-1952 6U VPX Single Board Computer

- Intel® Core™2 Duo T9400 at 2.5GHz
- Intel® GM45 chipset
- XMC site
- 4-8GB DDR3 SDRAM
- I/O including DVI Video, RGB, Ethernet, SATA, USB, serial (RS-232, -422), JTAG and GPIO
- 6U VPX air- and conduction-cooled
- VPX-REDI 2 level maintenance compatible

Curtiss-Wright Controls Embedded Computing

Tel: 613-254-5112
sales@cwcembedded.com
www.cwcembedded.com

CURTISS WRIGHT Controls Embedded Computing

CHAMP-AV6 6U VPX Digital Signal Processor

- 4x Freescale MPC8641/D CPUs at 1.0GHz (8 cores)
- XMC site with x8 PCIe interface
- 1GB DDR2 SDRAM per CPU
- Serial RapidIO architecture
- 6U VPX air- and conduction-cooled options

Curtiss-Wright Controls Embedded Computing

Tel: 613-254-5112
sales@cwcembedded.com
www.cwcembedded.com

CURTISS WRIGHT Controls Embedded Computing

VPX6-185 6U VPX Single Board Computer

- Freescale MPC8640/1D up to 1.33GHz
- 2x XMC/PMC sites
- Up to 2GB DDR2 SDRAM
- I/O including Ethernet, serial (RS-232, -422 sync), USB, 1553, PCIe, SRIO, TTL DIO and DIFFIO
- 6U VPX air- and conduction-cooled options

Curtiss-Wright Controls Embedded Computing

Tel: 613-254-5112
sales@cwcembedded.com
www.cwcembedded.com

CURTISS WRIGHT Controls Embedded Computing

HPE720 6U VPX Digital Signal Processor

- Single or dual-core Freescale MPC864x at 1GHz
- 2x Xilinx Virtex-5 LX220T, LX330T or SX240T FPGAs
- 2x FMC sites or 1 FMC and 1 XMC/PMC site
- I/O including Ethernet, serial (RS-232, -422), PCIe, Serial RapidIO, TTL DIO and DIFFIO
- 6U VPX air-cooled
- FusionXF development package

Curtiss-Wright Controls Embedded Computing

Tel: 613-254-5112
sales@cwcembedded.com
www.cwcembedded.com

EMBEDDED
Products you can count on!
A Xycom/VME Company

VME-6250

- Intel® Core™2 Duo 1.5GHz or 2.16GHz
- Dual 32/64-bit 33/66MHz PMC/XMC sites
- 2GB or 4GB ECC DDR2 400MHz memory
- Four Gigabit Ethernet ports
- Two SATA external devices
- Two serial ports
- SVGA, LVDT, keyboard, mouse, three USB
- Available in conduction- and non-conduction-cooled versions

Xembedded, LLC

Tel: 734-975-0577
Fax: 734-975-0588
sales@xembedded.com
www.xembedded.com

EMERSON Network Power

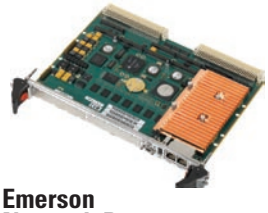

MVME7100 Multicore VMEbus SBC

The Emerson Network Power MVME7100 single-board computer is one of the first VMEbus products with the system-on-chip MPC8641D processor and 2eSST VMEbus interface. The MVME7100 delivers competitive advantage and investment protection to industrial, medical, defense and aerospace OEMs. VME users can now have significantly improved performance and extended product life cycles.

Emerson Network Power

Tel: 800-759-1107 or 602-438-5720
EmbeddedComputingSales@Emerson.com
www. EmersonNetworkPower.com/EmbeddedComputing

EMERSON Network Power

CPCI7200 Multicore CompactPCI SBC

The Emerson Network Power CPCI7200 SBC, with the Intel® Core™2 Duo processor and Intel® E7520 + 6300ESB chipset, is a low-power, high-performance SBC ideal for applications in thermally constrained environments. Dual GbE interfaces and dual channel 3.2GB/s high speed, double data rate DDR2 provide a combined maximum bandwidth of 6.4GB/s.

Emerson Network Power

Tel: 800-759-1107 or 602-438-5720
EmbeddedComputingSales@Emerson.com
www. EmersonNetworkPower.com/EmbeddedComputing

Photo courtesy of U.S. Air Force

Movers and shakers: Simultaneous multiaxis testing tops sequential axis methods

By Wayne Tustin

Simultaneous multiaxis vibration testing is shaking up the testing world – and beating out the more common sequential axis shaking method.

Most engineers would readily agree that "real-world" vibration is likely to be multi-axis. They would also agree that aircraft, missiles, rockets, land vehicles, and ships must simultaneously shake all onboard electronic and other hardware in multiple axes X, Y, and Z. In their own experience driving an automobile or riding in an airplane, bus, or train, they have felt multiaxis inputs. When they are first exposed to military and other vibration test procedures, most engineers are surprised by the ancient laboratory practice of sequential or single-axis-at-a-time shaking (see Sidebar 1).

However, with the appearance of new Test Method 527 on Multi Exciter Testing – which appears in the latest "G" revision to the venerable Military Standard 810 Department of Defense Test Method Standard for Environmental Engineering Considerations – engineers will at last find some incentive for simultaneous (Figure 1) rather than sequential multiaxis shaking of military embedded systems.

FIGURE 1: An array of three electrodynamic shakers simultaneously vibrates test items north-south, east-west, and up-down. (Courtesy Spectrum Technologies, Inc., Redford, Michigan)

While this discussion emphasizes automobiles and automotive testing, mainly because of likely reader familiarity with automobiles, the concepts presented also apply to testing a wide variety of military and commercial land, sea, and air vehicles. Single-axis-at-a-time or sequential shaking is reviewed, and the benefits of multiaxis shaking – servohydraulically at relatively low test frequencies and electrodynamically at higher test frequencies – are explored. These benefits apply not only to the usual "frequency domain" testing in terms of power spectral density but also to "time domain" testing in terms of force or acceleration versus time.

Servohydraulic single and multiaxis shaking

Mechanical shakers were stroke and frequency limited. Thus, the need for several-inch strokes shaking frequencies prompted the development of servohydraulic (electrically valved high-pressure,

Sequential axis testing: History

Way back in vibration testing history, circa 1950, all shake testing was done mechanically. Vibration testing of large structures such as bridges or buildings was accomplished by attaching a rotating variable-speed unbalanced mass. Nominally sinusoidal, single-frequency-at-a-time vibration testing of smaller hardware involved bolting it (via an adapter or fixture) onto the vibrating platform of the only shaker type then available: mechanical. Test engineers routinely positioned their hardware on that single-axis platform and shook it sequentially along its X axis, then its Y axis, and finally along its Z axis … three tests. Sequential axis testing, unfortunately, is still with us in 2009.

SIDEBAR 1: Sequential axis testing began in the 1950s and unfortunately is still used today despite its limitations.

oil-driven) shakers, various-sized actuators for transportation vibration testing of cartons or pallet loads of packaged goods. The frequency range is somewhat limited; however, servohydraulic shakers are rarely used above 500 Hz. Many are never used above 200 Hz. But automotive test engineers have long used four such actuators to vertically shake a test platform that represents a railcar or a motor truck carrying new cars from factories to dealerships.

Pioneering users of such systems were not satisfied with single-axis shaking, though. They could not replicate in the lab all the damage (such as assemblies loosening and parts losses) that occurred when shipping completed automobiles. Fortunately, a few engineers recognized that they could vary the phase between the electrical signals driving the servovalves on the four shakers (see Sidebar 2). Their test platforms could then be made to roll and pitch as well as shake vertically. Adding more shakers (lateral and fore-and-aft shakers) enabled two more translations as well as yaw: six platform motions. Such a multiaxis vibration laboratory test platform can represent a familiar sight: special-purpose railroad cars carrying completed automobiles from factory to dealership. The lab can identify potential "in transit" failures, leading to automobile redesign and/or manufacturing changes such as stiffening of suspension attachment bushings.

So what about road and off-road terrain inputs to vehicles in use? For these tests, land vehicle developers shake the individual wheels. Typically, a total of 12 servohydraulic shakers is used. Note the vast amount of lab area required for such testing. One manufacturer of servohydraulic shakers has put six such shakers inside a hollow steel cube, causing the cube to simultaneously shake in six axes. An automobile resting on four such cubes uses little lab space. Or one such cube can multiaxis shake a number of automotive assemblies and/or sections of automobiles in which electronics are embedded.

Military electrodynamic simultaneous multiaxis shaking

In order to shake microelectronic and other small assemblies at higher test frequencies, typically 1,000 or 2,000 Hz, electrodynamic (alternating current in coil located in strong magnetic field – similar to a loudspeaker) shakers such as those in Figure 1 had to be developed.

Common test lab interpretation of Military Standard 810 Test Method 514 has been sequential single-axis-at-a-time vibration. In most testing labs, engineers have never seen or even considered simultaneous multiaxis shaking. However, there have been several incidents in which military hardware has failed in the field, but those failures could not be duplicated in the laboratory. Very few military labs have obtained funding to add two more electrodynamic shakers:

- The U.S. Army at Adelphi, Maryland and at White Sands Proving Ground, New Mexico
- The U.S. Navy at Keyport, Washington
- The U.S. Air Force near Ogden, Utah

Consequently, axes X, Y, and Z can now be excited *simultaneously*. In some of these situations, the hidden failure modes have appeared. That is the greatest motivation

Random vibration: Making the grade

Other pioneers recognized that road inputs were not the classical sinusoidal, single-frequency-at-a-time sine vibration of mechanical shakers. The need for random vibration was another factor promoting servohydraulic shakers over mechanical shakers.

Such tests are described by continuous all-frequencies-at-the-same-time spectra. They typically are graphed as Power Spectral Density (PSD) versus frequency in hertz. The units of PSD (or alternate term is ASD: Acceleration Spectral Density or Auto-Spectral Density) are given in g^2/Hz units (vibration power per unit bandwidth). Constant 0.1 g^2/Hz was about 1,955 required by aircraft and missile manufacturers over a continuous band 20 to 2,000 Hz. (Such continuous spectrum testing is impossible with single-frequency-at-a-time mechanical shakers.) The area of that rectangle was (0.1 g^2/Hz) (1,980 Hz) or a mean square acceleration of 198 g^2. Then the Root of the Mean Square or RMS acceleration was 14 g. Today's spectra have g^2/Hz steps and slopes.

SIDEBAR 2: The need for random vibration was another factor promoting servohydraulic shakers over mechanical shakers.

for Military Standard 810G's new Test Method 527 on Multi Exciter Testing.

The other motivation is economic: faster, one test instead of three, and one attachment fixture instead of three to be designed and fabricated. Interestingly, automotive multiaxis shaking is common practice in the test labs of Japanese automobile manufacturers. Several Japanese shaker manufacturers offer multiple electrodynamic shaker arrays somewhat similar to Figure 1. They differ in their manner of connecting three shakers to a common load, Figure 2.

Only one independent environmental testing laboratory in North America provides simultaneous multiaxis vibration testing to 2,000 Hz: Spectrum Technologies, Inc. of Redford, Michigan. STI offers these services to military, automotive, and other hardware makers. Occasional training held there uses the system shown in Figure 1 to demonstrate simultaneous multiaxis vibration testing.

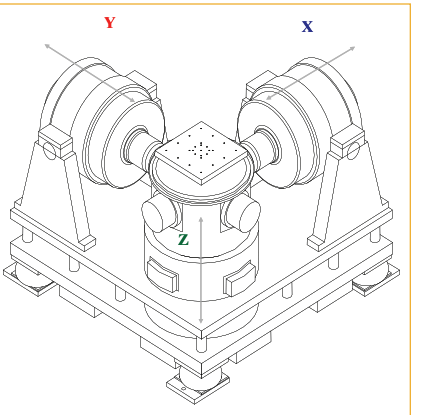

FIGURE 2: Payload adapter accepts mutually perpendicular inputs from three electrodynamic shakers. (Courtesy IMV Corporation, Osaka, Japan)

Time Domain Replication: An alternative to continuous spectrum testing

Military and other testing standards here and abroad commonly tell testers to gather and use their own data rather than use the provided spectra. Accordingly, some automotive STI clients gather their own over-the-road or off-road vibration data. They might mount "triaxial" accelerometers or they might appropriately mount three single-axis accelerometers at their locations of interest. They gather acceleration data and save it in the time domain, ready for Time Domain Replication (TDR) testing.

Under the previously described continuous spectrum testing, their "time domain" data (millivolts versus time) would have been time-averaged, then Fourier transformed into a spectrum in the frequency domain. Spectra would have been keyboard input by operators into digital shaker controllers. Digital-to-analog conversion would have been needed to develop a signal for each power amplifier to drive its electrodynamic shaker. Each shaker's motion would have spectrally matched its requirement. Unfortunately, that time averaging greatly reduces shaker reproduction of brief severe events such as the occasional highway bump or chuckhole. Those brief events were always under-represented in shaker motion.

TDR avoids that averaging and is believed to better reproduce shock events. The three acceleration signals remain in the time domain and pass without time averaging to each power amplifier and

each shaker. Another term for TDR is RoadLoad signal replication.

The future of simultaneous multiaxis shaking

In coming years, more and more test labs will "go simultaneous multiaxis," though there will be complaints about the required investment. Back in the 1950s, similar complaints were voiced about users' first electrodynamic shaker purchase. Once installed and used, that first shaker found so many weaknesses in the products tested that the shaker was soon used 24 hours per day. Then additional shakers were quickly purchased. That pattern will repeat with simultaneous multiaxis shaking.

Wayne Tustin *is founder and president of the Equipment Reliability Institute in Santa Barbara, California, a specialized engineering school.*

Equipment Reliability Institute
805-564-1260
www.equipment-reliability.com

Airborne network security optimized for the mission.

With over 40 years of high-grade and Type 1 system development experience, General Dynamics is the only independent supplier of advanced cryptographic and system solutions for airborne security.

Go to the source.

- Largest library of legacy and modern cryptographic algorithms for airborne networking and ground-based mission planning
- Proven Type 1 certified solutions for both avionics and ground applications
- One-stop system security engineering, certification and accreditation support for platform requirements
- Full suite of multi-level network and crew station solutions

For complete turnkey airborne security solutions or component-level embeddable security, contact General Dynamics today.

GENERAL DYNAMICS
C4 Systems

480-441-5448
866-400-0195 (toll-free)
IASystems@gdc4s.com
www.gdc4s.com/highassurance

DSP-FPGA.com
Supplement
www.DSP-FPGA.com

FPGAs & Reconfigurable Computing

2009

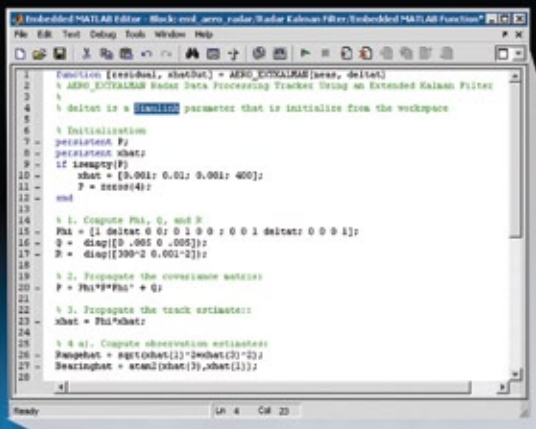

Modern FPGA tools ease design pain

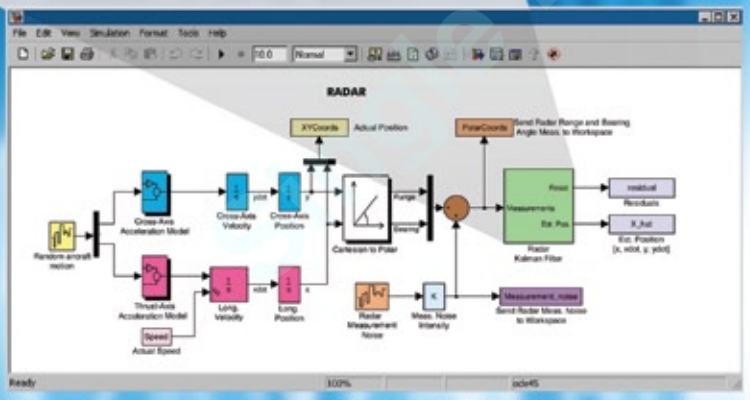

The MathWorks' MATLAB and Simulink tools make designing DSP functions, such as Kalman filters, practically drag-and-drop. FPGA implementations rely heavily on improved EDA and system tools. Image courtesy of The MathWorks.

Columns

Features

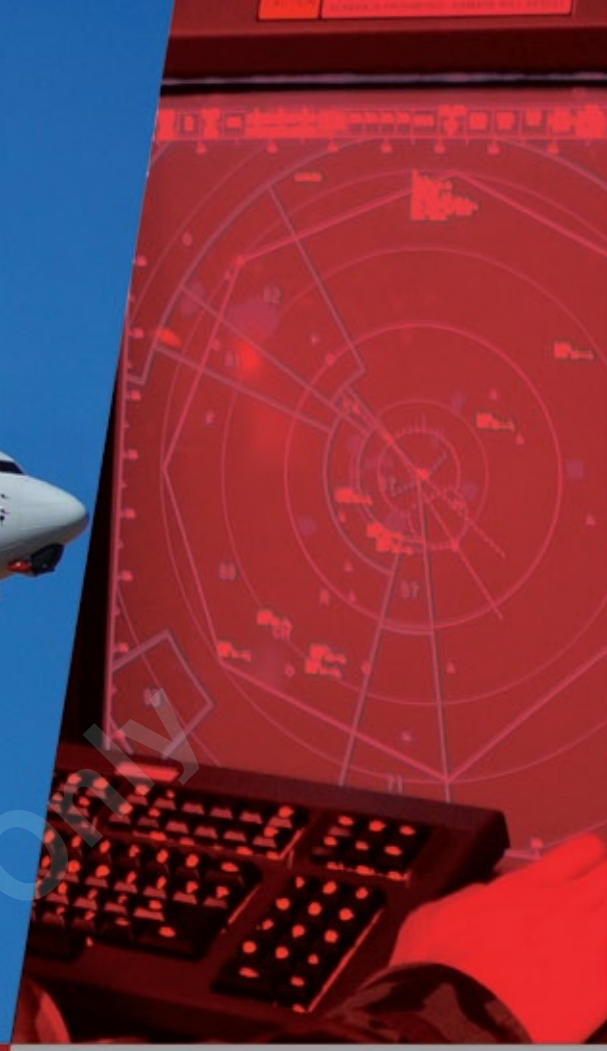

CHANGING I/O AND PROCESSING

FMC (VITA 57) is a new standard for I/O mezzanine modules that work closely with FPGA processors. FMCs provide a rugged I/O platform with high-speed signaling directly to FPGA processors on host cards. This new I/O standard is changing the battlefield today.

Curtiss-Wright delivers everything you need to integrate this new technology into commercial and rugged systems:

- FMC modules
- FMC carriers with Xilinx® Virtex®-5 FPGAs
- FPGA Development Kits.

High-performance FPGA-based processing is reaching new heights.

ADC510	ADC511	ADC512	ADC513
2x 500MSPS 12-bit ADC	2x 400MSPS 14-bit ADC	2x 3GSPS 8-bit ADC	4x 1.5GSPS 8-bit ADC

CURTISS WRIGHT Controls
Embedded Computing

www.cwcembedded.com
sales@cwcembedded.com

FPGA-BASED PROCESSING & I/O....ABOVE & BEYOND

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

Fully Integrated into IBM Blade Management System
Abundant Power and Cooling Ensure Maximum Performance

Made in the USA

Ultimate Modularity

From 2 to 8 Virtex 5 FPGA/Memory Modules
Input / Output Modules Include:
Quad 130 MSps thru Quad 500 MSps A/D
1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10 G Ethernet, FC4, SFPDP

Direct Seamless Connections with no Data Reduction

Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE Backplane
Between FPGAs and Standard Output Modules

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

Latest FPGA tools keep getting better, smarter

Tightening the grip on faster and easier development

By CHRIS A. CIUFO, EDITOR

More intuitive integration and the realization of multiple embedded processors, plus safety certification, highlight tools' new features.

FPGAs are no longer just glue logic. Let's face it, designing an FPGA requires a multidisciplined designer who's part RTL coder, system engineer, and software writer. The latest FPGA development tools from Altera and Xilinx are designed to simultaneously solve more real-world design problems, while dramatically improving ease-of-use. If you or your company implement FPGA designs – especially those used for DSP systems – you'll want to make note of these latest tools and development systems.

Xilinx: The end of *a la carte*

Xilinx finally recognizes that its battle is not against gate arrays, it's against companies designing ASSPs such as those for handheld music players, medical devices, or deployed military platforms. As long as the volume isn't in the tens of millions, chances are an FPGA will be used instead of a custom IC. But those target systems usually fall into several easily identified domains with known-in-advance IP such as MPEG4 decoders or 1024-point FFT algorithms. As well, the company's smorgasbord of design tools was a dog's breakfast of synthesis, floor-planning, probe, IP core, and reference designs that mostly represented a closed system – and even then, most tools didn't play well together.

By the time you read this, Xilinx will have gotten religion and announced ISE Design Suite 11.1. The new "Domain-specific Methodology for Targeted Design Platforms" creates four pre-packaged toolboxes with individual products designed for interoperability, improved performance (both in the software and in the final FPGA design), and tailored for the types of design problems engineers are actually battling. Figure 1 shows the various editions and their major tool components.

You might be thinking that on the surface this is just a repackaging exercise, especially since the lower end Logic platform is the basis for the others, and more tool components just get tacked on as you add Embedded, DSP, and finally the *full monty* System platform. But you'd be wrong. Now that Xilinx has defined the types of customers using all this stuff, ISE is better suited to solving problems. As well, new features have been added to simplify what engineers are actually designing. In the Embedded platform, for instance, Xilinx acknowledges that a full 20 percent of its customers are using an embedded CPU such as the PowerPC, soft MicroBlaze, or an ARM (in select silicon versions). So the Embedded platform includes multiprocessor support (but not cache coherency) and cross-optimization with a multiport memory controller that uses fewer resources and less silicon.

Figure 1 | Xilinx's ISE Design Suite 11.1 now has four "Design Platform" flavors.

In the DSP platform, System Generator automatically launches the SDK to modify software when using an embedded processor, allowing designers to code algorithms in the CPU, in ExtremeDSP slices, or in logic. The AccelDSP tool is now included and can deliver over 2x higher performance in Virtex-5 (and soon) Virtex-6 silicon. And interestingly, an Enhanced Fixed Point report helps designers decide the trade-offs in moving from floating to fixed point math. There's lots more to ISE 11.1 than we can cover here.

Across the board, ISE 11.1 is supposed to deliver 10 percent lower dynamic power by reducing the number of switching elements and gated clocks. It promises 2x faster runtimes, while delivering 28 to 30 percent better host memory utilization (the difference between a low-cost or big bucks machine). As well, Xilinx has finally (!) implemented a FLEXnet license system with either floating or locked (cheaper) options. As well, the company is aligning the 25 development reference boards developed for the platform-oriented Virtex-5 devices into Design Platform versions, complete with an industry-standard FMC mezzanine card interface.

Altera: 40 nm and still shrinking

In February, Altera followed up on last year's 40 nm news by making it official: Stratix IV GT and Arria II GX have screaming transceivers. Capable of pumping 155 Mbps up to a whopping 11.3 Gbps, the high-end hardware in Stratix IV GT brings to bear 24 high-speed transceivers, and an additional 24 6.6 Gbps transceivers, all in a density with up to 530K LEs, 20.3 Mbits of RAM, and 1,288 18 x 18 multipliers for DSP implementations. So that's the hardware; what about the tools?

Altera has made numerous changes and improvements to the company's (mostly) free Quartus II software, now at version 9.0. Many of the changes are specifically geared to optimize high-speed transceiver designs in the Stratix, Arria, and HardCopy FPGA families. In fact, this is a major advantage of Altera's own tool versus a third-party broadline EDA suite: Quartus II version 9.0 is unified across the company's entire product line. Designers use one familiar software suite for any IC within Altera's portfolio, which probably saves learning curve time while allowing hardware optimization for low-cost, speed-sensitive, or eventual high-volume ASIC-like HardCopy designs.

New features in version 9.0 include a Simultaneous Switching Noise (SSN) tool that warns a designer of possible pin assignment cross-talk with the high-speed transceivers. While it seems obvious

that high-frequency buses shouldn't all switch at the same time (hello?), routing complex signals while optimizing board real estate could easily result in switching violations. We think the SSN tool is a nice feature. New metastability analysis capability tools will auto-recognize potential metastability issues, and the TimeQuest static timing analysis tool is now automated to report MTBF values. Finally, the pin planner's clock network view helps to better manage resources. The company also juggled around some of the previous built-in feature packages. ModelSim Web Edition is now included in Quartus II 9.0 as "Starter Edition," which Altera says increases simulation speeds by 50 percent. The ModelSim Altera Edition is also available stand-alone, with another 33 percent speed improvement on top of Starter Edition.

On the DSP side of the house, the company has partnered with The MathWorks in a series of "Military, Aerospace and Defense DSP Productivity Seminars," showcasing model-based design. Both companies will delve into their DSP design flows in a Stratix IV GX realizing a sustained 175 GMAC sensor fusion application. The customer-location seminar schedule can be found at www.altera.com/pr/military-dsp_resources.

Update

As we went to press, Altera announced the Transceiver Signal Integrity Development kit for the Stratix IV GX. With 8.5 Gbps transceivers, the board includes an FPGA and 8 channels routed to SMA connectors plus USB, flash memory, and user software via a GUI that lets designers try out applications such as video over IP, 4G wireless, and digital TV.

DSP-FPGA.com Editor's Choice

Cost-effective prototyping for space-flight FPGAs

Though Altera and Xilinx FPGAs are used in space-based applications such as satellites and spacecraft, Actel has the market cornered on Class-S style rad-hard FPGAs. Trouble is, they aren't cheap. And RTAX anti-fuse devices are one-time programmable, making them expensive to burn during prototyping. The solution? Use an Aldec development adapter board that's populated with lower-price, reprogrammable Actel ProASIC FPGAs. The RTAX4000S prototyping adapter board includes a unique dual-stacked configuration with a ProASIC3 A3PE3000 on each board, which collectively "replace" the more expensive RTAX4000S device.

The bottom part of the board is footprint compatible with the space-type RTAX CQ352 devices. This allows soldering the adapter board to a PCB just as if it was a regular device. As well, it can be soldered and unsoldered several times should the need arise during prototyping. And the ProASICs are flash-based, so they can be reprogrammed during development. According to Actel, the JTAG programming connector allows on-the-fly device reprogramming without detaching the adapter, and an EDIF netlist converter lets designers migrate from RTAX-S/SL and RTSX-SU space devices to ProASIC3E devices.

Aldec • www.aldec.com • RSC# 41500

Multiprocessor debugging challenges

By ROB HOYECKI

Recognizing the significant impact on time to market that the task of debugging complex multiprocessing systems presents, vendors are introducing solutions that capitalize on new processor devices with multiple cores and heterogeneous systems based on FPGAs and multiple processors.

Two key hurdles facing developers of large DSP systems are the management and debugging of these complex configurations. A look at a typical configuration reveals many processing nodes interconnected by a switched fabric such as Serial RapidIO. Developers have approached the challenge by wielding source-level debugging tools originally designed for single processor applications only to find that these don't scale well when applied to a complex DSP system with real-time dataflow. One of the most important reasons is that each instance of a debugger requires its own connection.

These large multiprocessor DSP systems can't quit their day (or night) jobs. Performing such tasks as searching large parts of the electromagnetic spectrum for emissions, plus monitoring and processing many hundreds of active channels, demands that all the processors in a system be operating and interacting optimally with each other. For defense and aerospace applications including synthetic aperture multimode radar or signals intelligence, if errors occur or performance is degraded, it is very difficult to identify where this might be taking place and yet more difficult to name the individual processor at the root of the problem.

Resource management

With systems comprising tens or even hundreds of processors the typical edit/compile/load/debug cycle becomes tedious and error prone, even if scripted. The system may contain processing nodes of different revision states, plus processors may each have their own image, or groups of processors may need to be loaded with the same image. Whenever a hardware or configuration change is made, revision states, as well as the required images, may change. Tools are becoming available to identify processors, their resources, groups of processors, and overall system configuration in order to automate the downloading and verification of current or latest versions of images before and during debugging sessions. This practice ensures that each processing node is of a known configuration and contains the appropriate image before debugging starts. These tools not only reduce the number of initialization errors, but save considerable time. For example, a typical time saving of five to ten minutes per reset can add up to over an hour of active debugging time saved each day.

Grouping for an answer?

Traditional source-level debugging tools were created for single-processor environments. However, a typical complex DSP has many processors, often organized as groups sharing data or executing similar code and communicating within the group via a switched fabric or network.

Performance degradation or erroneous results may be caused by many potential issues such as:

- Race conditions
- Memory leaks
- Buffer starvation
- Buffer overflow
- Loss of synchronization

All these issues are difficult to pinpoint to a particular time period or processor. Developers have come up with various methods to monitor and debug groups of processors. One method uses instrumentation to capture and time-stamp the state of each processor at predetermined events within the user's code, or at the operating system call level, using a tool such as Wind River Systems' System Viewer. The instrumentation approach generally allows the system's performance to be monitored visually at the coarse level online, but so much data can be captured that off-line analysis is often the only method to obtain the necessary granularity. Instrumentation is most useful to gain insights into performance degradation before source-level debugging an individual processor.

Another method is to include in-line logging within the code, sending results to a console port. However, this is not really a practical approach for more than a small number of nodes and, being intrusive, it affects the performance and determinism of each processor. Similarly, using multiple instances of a single processor debugger is clearly unwieldy beyond a small number. Neither does it address the issue of time-coherence, as breakpoints on each processor may occur at different times or be triggered by different events.

Breakpoint breakthrough

A better approach, and that relies on a single breakpoint on a single processor, addresses many of the downsides that the instrumentation or in-line logging methods present.

Setting a single breakpoint on one processor and using that breakpoint to halt selected other processors at the same instant enables single stepping through the problem, using one or many processors, until the error point is located. Curtiss-Wright Controls Embedded Computing (CWCEC) has developed this unique capability to breakpoint multiple processors as part of its Continuum Insights tools for DSP development, which also offers extensive complementary resource and Flash management tools for use with its dual and quad multicomputing engines.

Until recently many of the management and debugging tools for multiprocessing environments were limited in their ability to resolve the complex, interactive, and time-critical problems of such systems. The additional time burden and effort faced by system engineers who must debug these complex multiprocessing systems can detrimentally affect time to market for systems critically needed by today's warfighter. Vendors have now recognized these needs, and solutions to meet multiprocessor debugging challenges are set for further rapid expansion that takes advantage of new processor devices with multiple cores and heterogeneous systems based on FPGAs and multiple processors. The ability to efficiently debug multiprocessor systems will become increasingly important as these systems become the norm for the next generations of multispectral, reconfigurable sensors.

Robert Hoyecki *is Director of Advanced Multi-Computing at Curtiss-Wright Controls Embedded Computing. Rob has 15 years of experience in embedded computing with a focus on signal process products. He has held numerous leadership positions such as Application Engineering Manager and Product Marketing Manager. Rob earned a Bachelor of Science degree in Electrical Engineering Technology from Rochester Institute of Technology.*

Rob can be reached at info@cwcembedded.com.

We've Hatched the Next Generation of Software Radio Solutions

Pentek delivers board and system-level SDR products that include digital downconverters, upconverters and transceivers. Since all of these products are FPGA-based, Pentek offers powerful factory-installed IP cores plus the GateFlow FPGA Design Kit for custom development. These software radio solutions are perfect for applications in communications, SIGINT, radar, direction finding and many more.

Let Pentek help you get your software radio systems out of the nest on time!

Call 201-818-5900 or go to www.pentek.com/go/dspsdregg for your FREE Software Defined Radio Handbook, technical datasheets and to request pricing.

- XMC/PMC, PCI, PCIe, VME/VXS, cPCI
- Over 150 software radio modules
- Up to 512 channels of DDC per slot
- New 200 MHz, 16-bit A/Ds
- FPGA-installed SDR IP cores
- Recording systems with rates up to 800 MB/sec

PENTEK
Setting the Standard for Digital Signal Processing

ISO 9001:2000 CERTIFIED QUALITY MANAGEMENT SYSTEM

Pentek, Inc., One Park Way, Upper Saddle River, NJ 07458 Phone: 201.818.5900 Fax: 201.818.5904 e-mail:info@pentek.com www.pentek.com

The drive to lower power

By Christian Plante

Christian explains why the demand for portable, lower-power electronics is growing and describes applications where FPGAs have a role today.

As process technology continues to improve, device sizes shrink, and designers pack more onto silicon real estate, FPGAs increasingly find themselves the center of attention. FPGA differentiation as measured by flexibility and time to market is finding new applications and responsibilities.

The shift toward portability and power-conscious electronics has led to the demand for low-power components of every variety, and now that includes FPGAs. Today power is a more important design consideration than performance. For some vendors this has meant offering incrementally lower-power devices, with next-generation devices being 60 percent lower power than previous generations.

The low-power trend is taking root in a variety of applications:

- In medical devices, where there is a clamoring for portable medical equipment for quick and easy diagnosis.
- In displays, which are being incorporated into innumerable systems.
- In industrial equipment, where motor control design is seen as a key way to improve overall energy efficiency of products, and ultimately, of the world.

Low-power applications represent a new role for FPGAs. Designers have traditionally relied on ASICs, not FPGAs, to meet their low-power constraints. But hardwired ASICs – with longer time-to-market, rising Non-Recurring Engineering charges (NREs), and a lack of flexibility to address changing standards and late-stage design modifications – are riskier and often impractical for applications with short product life cycles or evolving standards.

Similarly, Complex Programmable Logic Devices (CPLDs), used in some low-power applications, are losing their effectiveness due to relatively high costs and the increased demand for high-end features and extra logic. CPLDs do not offer the level of integration, flexibility, or sophistication required for most of today's applications.

Designers are finding that a low-power, reprogrammable solution is required to adapt to evolving standards, speed time to market, offer products with multiple personalities, and deliver the footprint and power consumption that cutting-edge electronics designs require.

The power-miserly FPGA

Not all programmable logic is well suited to address low-power needs. In fact, some of today's "low-power" FPGAs draw upwards of 30 mA, which is often an order of magnitude or two higher than typical power-sensitive, battery-operated applications can tolerate. SRAM-based devices experience well-documented inrush and boot-up configuration power spikes during system initialization that can drain a battery quickly.

But single-chip, flash-based devices do not require an external configuration device (for example, a boot prom or microcontroller) to support device programming at every power-up cycle. And the live-at-power-up feature eliminates the need for an external device to assist in system boot up. Removing the additional parts required by SRAM-based FPGAs not only reduces board space and system power consumption, but also increases reliability, simplifies inventory management, and lowers total system costs by as much as 70 percent compared with similar SRAM-based FPGA solutions.

Once the FPGA is on and configured, power consumption takes two basic forms – static and dynamic. Static power consumption is the current drawn by an FPGA when it is powered up, configured, and doing nothing. Dynamic power is consumed when devices are actively working. Until recently (see Figure 1), dynamic power was the dominant source of power consumption. Device supply voltages (Vcc) once helped manage dynamic power but scaled downward with process shrinks and subsequent lower system voltages. However, the days of continued scaling are coming to a close.

Compounding the issue, each process node shrink means additional static power consumption for transistor-heavy SRAM-based FPGAs. This is due to worsening problems like quantum tunneling and sub-threshold leakage, which create real challenges for devices targeted to power-conscious applications. And, with leakage worsening, static power has begun to dominate the power consumption equation as the biggest concern.

The SRAM cell structure incurs substantial leakage and requires power-consuming configuration memory. In sharp contrast, flash-based cells have no leakage path and thus have 1,000 times lower leakage per cell than SRAM.

Alternative approaches

To address some of these power concerns, several suppliers of SRAM-based FPGAs claim to offer "single-chip, flash-based" solutions. These "hybrid" solutions are merely combinations of flash memory components with the underlying SRAM FPGA technology – integrated with the FPGA die into a single package, or stacked, or placed side-by-side. Unfortunately, the FPGA array is still volatile and is subject to the power drawbacks associated with these types of devices. With these solutions, the embedded flash memory blocks control only the initial configuration of the devices during power-up.

Certainly, both the Silicon-in-Package (SIP) and the Multichip Package (MCP) hybrid approaches overcome some of the limitations of traditional SRAM-based solutions with a smaller footprint, a minor reduction in power consumption, and small advances in power-up time and security. But these are only incremental improvements over their pure SRAM-based peers.

True nonvolatility

True nonvolatile FPGAs are those that contain a nonvolatile FPGA array, reducing power consumption, improving response times, and delivering unparalleled reliability and security. Because true nonvolatile flash-based FPGAs don't use millions of power-hungry SRAM configuration bit cells, they have significantly lower static power than SRAM-based solutions, making them ideal for power-sensitive applications. In fact, the many flavors of flash-based, low-cost FPGAs include devices that have been optimized for power, speed, and I/O, some of the fundamental design requirements for power- and cost-sensitive design.

Low-power FPGAs at work

Portable medical equipment made big news at the 2008 Beijing Olympics when companies, such as GE, were test-driving portable MRIs and other imaging and diagnostic tools at the Games. FPGAs are playing a big part, allowing developers to program various features, design for differing geographical standards, and imbue the devices with multiple personalities, while at the same time keeping power as well as design costs low.

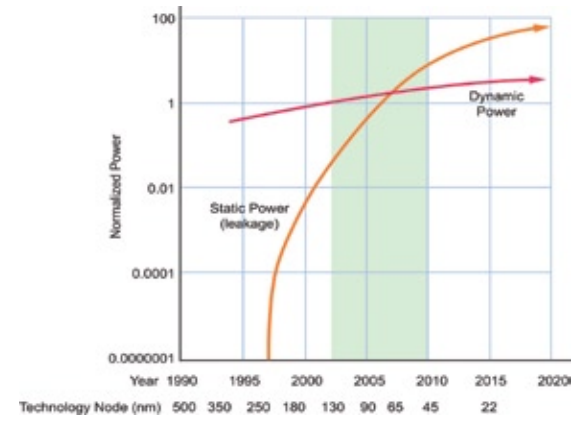

Figure 1 | Traditionally, dynamic power has been the dominant source of power consumption. Device supply voltages (Vcc) have scaled downward with process shrinks and the resulting lower system voltages, but continued scaling downward cannot be ongoing.

This trend towards miniaturization and portability for home, clinical, and imaging medical devices presents a significant opportunity for medical equipment designers to use FPGAs in developing efficient and flexible designs.

Medical devices have high-reliability requirements, demand multifunctionality (integrated capabilities), and require data logging and transmission capabilities. Yet these devices must consume the lowest amount of power. In the most basic form, portable medical devices are all battery-operated, microcontrolled handheld devices that take and analyze measurements using the various bio-sensors a patient's treatment plan needs.

Home-based and consumer medical devices – digital blood pressure meters, blood gas meters, and blood glucose meters – have been traditionally used for testing and monitoring. Today, medical devices are expected to do much more than just test and monitor. Some now log and analyze data and communicate accurate results to the health provider. Blood pressure meters benefit from a more extensive data-logging feature as well as communication ports for real-time information sharing with the health provider. Insulin meters are now equipped with communication ports (IR/wireless) to transfer real-time measurement to the PC or to the insulin pump to effectively treat the disease.

In Figure 2 the shaded functional blocks represent some possible functions that can be implemented in FPGA devices. These functions can be either individually addressed as needed by smaller low-power reprogrammable FPGA devices or can be integrated into larger FPGA devices. These ultra-low-power FPGA families offer gate capacity from 15 k gates up to 3 million gates.

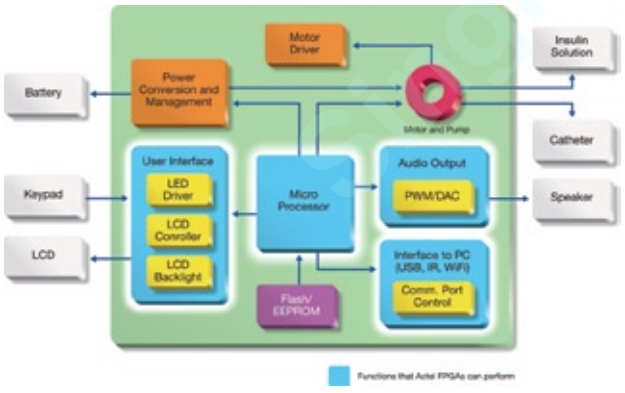

Figure 2 | Gate capacity ranging from 15 k up to 3 million is available from ultra-low-power FPGA families.

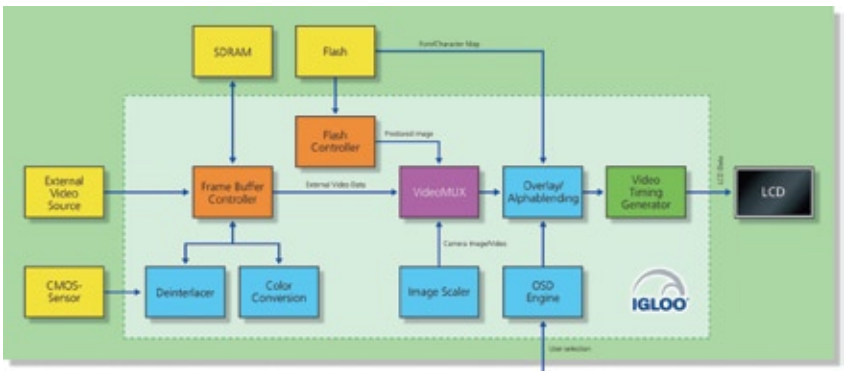

Figure 3 | Taking on a complex LCD controller role, FPGAs can be a power-efficient option.

Pressure to reduce health care costs is butting into the demand for more portable devices, which, when deployed in the home, can save the system money. But these devices need to be low power, flexible, and cost-effective for wide deployment.

Display dynamic

Let's look at another application that is increasingly leaning on programmable logic to ensure flexibility and time to market speed: Displays.

Lower costs and ease of mass manufacturing have increased LCD panel demand in medical markets. When creating these devices to meet consumer demands, designers select LCD panels based on critical factors, such as size, resolution, reliability, power consumption, and product life cycle. As newer displays with enhanced capabilities and features are continuously launched, designers face challenges in keeping up with technology by redesigning the display controller and need solutions enabling them to incorporate the latest technology with minimal cost and effort.

In portable devices, LCDs can consume up to 50 percent of the application's power budget, escalating the need for a power-efficient solution. Power-aware attention to the design does not rule out the use of FPGAs, which often are seen as power inefficient choices. Careful design considerations in managing power using FPGAs as a complex LCD controller, for example, can reduce power (see Figure 3).

Some FPGAs offer low-power advantages, down to 5 mW, while retaining the contents of the system memory and data registers. As a result, the FPGA approach can enable both the LCD panel and the controller to function in a power-saving mode with the LCD data and backlight disabled, achieving significant battery savings.

Motor control

New designs for AC and DC motor control must be highly efficient and consume little power, offering longer operation without affecting performance quality. The need to implement smaller, more cost-effective motors in traditional motor applications is also influencing electronic motor control techniques for the industrial sectors.

Expensive computer and power electronics have been significant obstacles to overcome for motor control applications. Integration that combines analog, flash memory, and FPGA fabric in a monolithic device is helping designers face these challenges. For the first time, engineers can combine the motor control analog front end, high-speed flash lookup tables, and deterministic algorithm processing capabilities of programmable logic into a single-chip solution.

Trade-offs? What trade-offs?

Low-power solutions are the biggest growth

(CHAIS)ing the dream: 100 GSps for analog-to-digital converters

Reality doesn't byte. The real world is analog, so converting signals into the digital realm is really important. Fujitsu Microelectronics Europe thinks so, too. The company's new CHArge-mode Interleaved Sampler technology (CHAIS) achieves sampling rates up to 100 GSps for Analog-to-Digital Converters (ADCs). This takes place in plain old CMOS instead of high-power SiGe. Moreover, multiple ADCs with tens of millions of gates of signal processing logic and memory can now be combined on a single chip, avoiding the mixed-signal multichip modules of other devices.

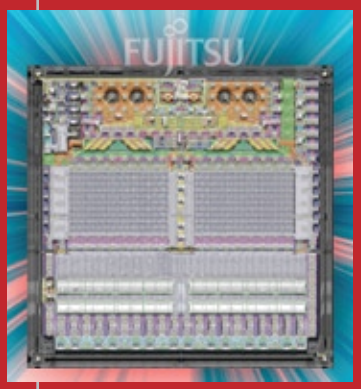

Targeting the company's 65 nm CMOS process, power consumption for a typical ADC is 2 W per channel at 56 GSps, which the company says is "unheard of." A half-speed mode boasts 28 GSps at 1 W per channel. Internal sampling clocks at under 100 fs total root mean square jitter and under 500 fs I/Q ADC skew. When combined on-chip with other logic, an ASSP can exceed 50 million gates, so super high-end communications rates such as fiber-optic 100 G Ethernet and OTU-4 can now be realized. The company's first production device is the DP-DQPSK coherent receiver for 100 G optical networks with 4 channels of 56 GSps 8-bit ADCs, logic, and memory all on a single-chip PHY device. Of course, with all this on-chip capability, it's no longer necessary to move terabits of data between an ADC front end and an IF DSP processor.

Fujitsu • emea.fujitsu.com/microelectronics • RSC# 41856

Mentor's Precision Synthesis melds with Altera Stratix IV GT

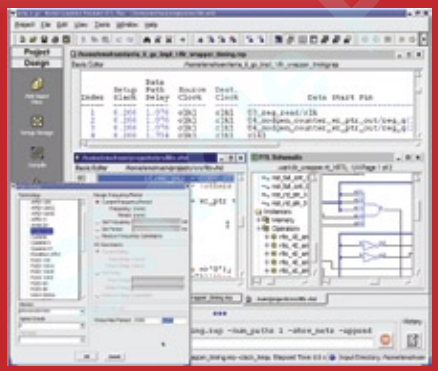

The thing about the competitive world of high-end FPGAs is that they get bigger (gate density-wise), faster, and more complex. Xilinx now has Virtex-6; Altera's got Stratix IV GT. If you stick with these vendors' own tools, you're golden. But you might be missing some advanced features, too. That's why Mentor Graphics revised the company's Precision Synthesis tool to dovetail with the latest from Altera's 40 nm family of Stratix and Arria devices: Stratix IV GT and Arria II GX. These devices boast a bit more density and lower power than their equivalent predecessors. Vastly improved LVDS SERDES I/O, suitable for 40 GbE and 100 GbE systems, is also part of the mix. In the area of transceiver technology Mentor's tool really shines.

The vendor-independent tool set allows a common environment for developing with multiple FPGA types and vendors. Precision Synthesis is called "the centerpiece" of Mentor's FPGA flow, offering push-button synthesis, HDL language support (including SystemVerilog), and automatic incremental synthesis to save time and facilitate modular design practices. Designers can cross-probe between multiple views or conduct "what if?" static timing runs. Mentor claims that the tool reduces design iterations while allowing designs to be completed faster. If you're already a Mentor user, your suite can be upgraded to work with Altera's latest FPGAs.

Mentor Graphics • www.mentor.com • RSC# 41857

segment for electronics design, but until recently that design imperative often came at the cost of designing out performance or functionality or missing out on integration opportunities.

Increasingly, however, those traditional trade-offs are not so onerous. FPGA vendors have integrated microcontroller and microprocessor cores into their devices. Some of these approaches feature very low operating current and static power, consuming only 24 µA in static mode and 3 µA in sleep mode. The Flash Freeze mode that Actel employs, which enables easy entry and exit from ultra-low-power modes while retaining SRAM and register data, reduces quiescent current to 20 µA. The feature also allows instant on/off cycling of the processor core for maximum performance and minimum power consumption.

This is approximately 200 times less static power than competitive FPGA offerings and delivers more than 10 times the battery life of the leading programmable logic devices in portable applications.

There's a similar advantage where analog meets digital. Mixed-signal FPGAs integrate programmable logic, RAM, flash, and analog onto a single chip, while lowering overall system power. A flash-based approach to mixed-signal FPGAs does not require additional configuration nonvolatile memory in order to load the device configuration data at every system power-up, which reduces cost and increases security and system reliability. Increased functionality can remove several additional components from the board, such as flash memory, discrete analog ICs, clock sources, EEPROM, and real-time clocks, thereby reducing system cost and board space requirements.

Conclusion

Today, FPGA technology is increasingly used in low-power applications. FPGAs have been adopted widely in recent years due to advanced technology that lowered the unit price, but the price reductions have come with higher power consumption due to higher transistor leakage.

In addition, smaller cores, better IP and development environments, and the double-time march of technology have created a situation in which processor and controller cores and mixed-signal functionality can now be part and parcel of FPGAs. This integration provides a single-chip solution that lowers overall power, cuts costs, and reduces design complexity. It also broadens a designer's application reach and time to market.

FPGAs are bringing to market low-power features at exactly the moment when power conservation and design flexibility are most required to meet evolving customer demand. PJ

Christian Plante *is Director of Marketing low-power and mixed-signal FPGAs. Christian joined Actel in 2008, bringing with him more than 14 years of experience in the semiconductor and computer hardware industries. Prior to joining Actel, he was the Director of Customer Marketing at Cswitch Corporation and also held several senior marketing positions at Altera Corporation. He holds a Bachelor's of Science degree in electrical engineering from Laval, Quebec City, Canada and a Master's degree in business administration from Queen's University, Kingston, Canada.*

Actel
newsroom@actel.com • www.actel.com

Create Your Own PMCs

Custom I/O as Easy as 1, 2, 3!

The Technobox ***Micro Mezzanine System***™ is based on a simple idea – provide embedded systems designers with a foundation for innovation and flexibility. Provide a highly-granular, modular architecture featuring a range of configurable FPGA-based carrier boards and an extensive variety of micro mezzanine Electrical Conversion Modules (ECMs) that can be assembled in thousands of combinations. Provide an environment in which a designer can create an array of unique, *future-proofed*, board-level solutions. But without the costs normally associated with custom board development and manufacture, while speeding development and reducing time to market. It's the logical next step in mezzanine systems.

1 **Select an FPGA-based Carrier**

2 **Choose I/O Conversion Modules**

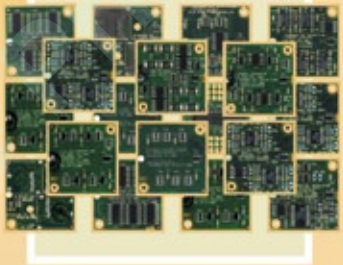

3 **Assemble with IP Core and ECM Code**

Patent Pending

- ***Build Your Own Board by Mixing and Matching Modular Components***
- ***Thousands of Possible Combinations***
- ***Flexible, FPGA-based, Patent-pending Architecture***
- ***Incorporate Multiple Functions on a Single Board***
- ***Design and Build Application-specific, Future-proofed Solutions***
- ***Accelerate System Development, Reduce Time to Market***

To learn more about the Technobox Micro Mezzanine System, visit:
www.technobox.com/mmsintro-fp.htm

Technobox, inc.®

Micro Mezzanine System

Trends in FPGA testing and validation

By Paul **Bradley**

Aiming for an at-speed hardware/software co-validation environment, developers want eminent domain visibility – for software and hardware.

FPGA vendors are feeling the effects of the economic downturn, but they appear to be in a position to weather the storm more easily than others in the semiconductor industry due to their broad appeal and compelling economic advantages. Despite adverse conditions, FPGAs continue to make inroads into new markets, and to expand in capability and capacity.

The infusion of FPGAs into the embedded systems community and the general convergence of hardware platforms is a positive development. Industry trends are pushing FPGAs, structured ASICs, DSPs, and conventional processors toward some virtual point of convergence. FPGAs started as almost purely programmable fabric but have since added embedded processor cores, memory, sophisticated I/O capabilities, hardwired multipliers, and dedicated peripherals to increase their versatility as embedded platforms.

While the advancement of these devices is good news, it has brought some associated headaches to the systems designers who work with them. The phenomenal growth in design size and complexity has made FPGA debug and design verification a challenge. Although the capabilities and performance of FPGAs have advanced, the debug and design validation techniques have not kept pace. Even the simplest of RTL designs need to be debugged and validated. This requires an intensive amount of engineering effort and man-hours, the results of which are not predictable and do not always prove the design is functionally flawless.

A new approach

Today's more sophisticated FPGAs require more sophisticated solutions – in particular, tools that can provide advanced controls and views into the embedded system.

FPGA vendors have a history of using their own on-chip analysis tools, but these solutions are no longer enough. A new generation of tools is emerging that delivers a more sophisticated on-chip instrumentation scheme and off-chip development and test environment. These new tools feature enhanced capabilities beyond what conventional FPGA tools and hardware debug/logic analyzers offer, including on-chip stimulus, on-chip analysis (such as performance monitoring), hardware-software correlation, assertions, transaction analysis, and multi-FPGA visibility, with the added ability to use the same technology in simulation and emulation environments or even for ASICs, if required.

While traditional tools primarily target hardware designers performing on-chip FPGA debug, the new tools serve a wider spectrum of system, software, and hardware engineers. For example, extending visibility beyond the software domain into the hardware interactions is a constant challenge for those who work with FPGAs. New toolsets can traverse the gap between the two domains to create an at-speed hardware/software co-validation environment, as illustrated in Figure 1. These next-generation tools are also more flexible than earlier offerings, because they don't target only certain FPGA families, but can perform validation across all FPGA devices.

Pre-silicon instrumentation and at-speed post-silicon validation tools can enhance observability and control of internal signals. They employ on-chip reconfigurable instruments that developers can insert and customize at design time for at-speed data acquisition, performance monitoring, stimulus, and fault injection functions. Programming post-silicon while the system is in operation offers a significant advantage, because the instruments can be repurposed to serve a variety of functions without requiring a resynthesis of the design. With in-depth visibility and stimulus-injection capabilities, these next-generation tools offer validation engineers not only the ability to observe deeply embedded signals, but also the capability to control those signals.

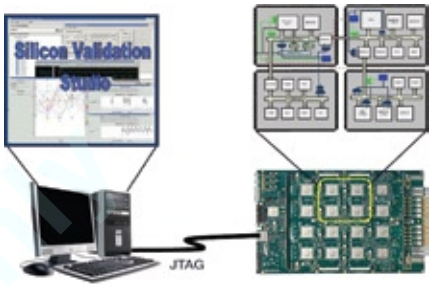

Figure 1 | New FPGA debug and validation tools' on-chip/off-chip architecture creates an at-speed hardware/software co-validation environment for multi-FPGA partitioning.

Key requirements

System, software, and hardware engineers who seek to develop more complex FPGA designs without increasing their costs or time to market must begin to use debug and validation tools that meet the following requirements:

- **Multi-FPGA designs** – to ensure observability and control across FPGA boundaries
- **Multi-FPGA partitioning** – to mitigate, rather than exacerbate, design partitioning challenges
- **Multicore/multichip visibility and control** – to ensure maximum insight and access into increasingly complex SoC designs
- **FPGA-agnostic** – usable on any FPGA, regardless of manufacturer
- **Flexible** – work on multiple platforms, including FPGAs, simulators, emulators, and ASICs
- **Comprehensive instrumentation** – with the fabric of multiplexors and transaction engines embedded in the user design, anywhere in the hierarchy
- **Configurable** – users can customize instrumentation at design time
- **In-system programmable** – the FPGA design does not have to be resynthesized to change instrument function
- **Offer on-chip stimulus, including transaction stimulus, fault insertion, and stress testing** – to achieve significant acceleration and benefits in delivering robust, high-quality embedded software
- **Comprehensive capabilities** – go beyond a logic analyzer to include performance monitoring, assertions, and transaction stimulus

It is clear that a new approach to FPGA debugging and validation is needed to reduce costs and time to market. As the market for FPGAs continues to evolve, the tools that support their cost-effective development must evolve as well.

***Paul Bradley** is Chief Technical Officer of DAFCA, Inc. Paul has over 20 years' experience in electronics and systems design, and specializes in product development and engineering leadership in emerging technology markets. He has held numerous engineering and technical leadership positions at Motorola, Nortel, CrossComm, Sonoma Systems, and Internet Photonics prior to joining DAFCA.*

DAFCA • 774-204-0020 • paul.bradley@dafca.com • www.dafca.com

At Rittal, We Handle CompactPCI System Designs With Experience.

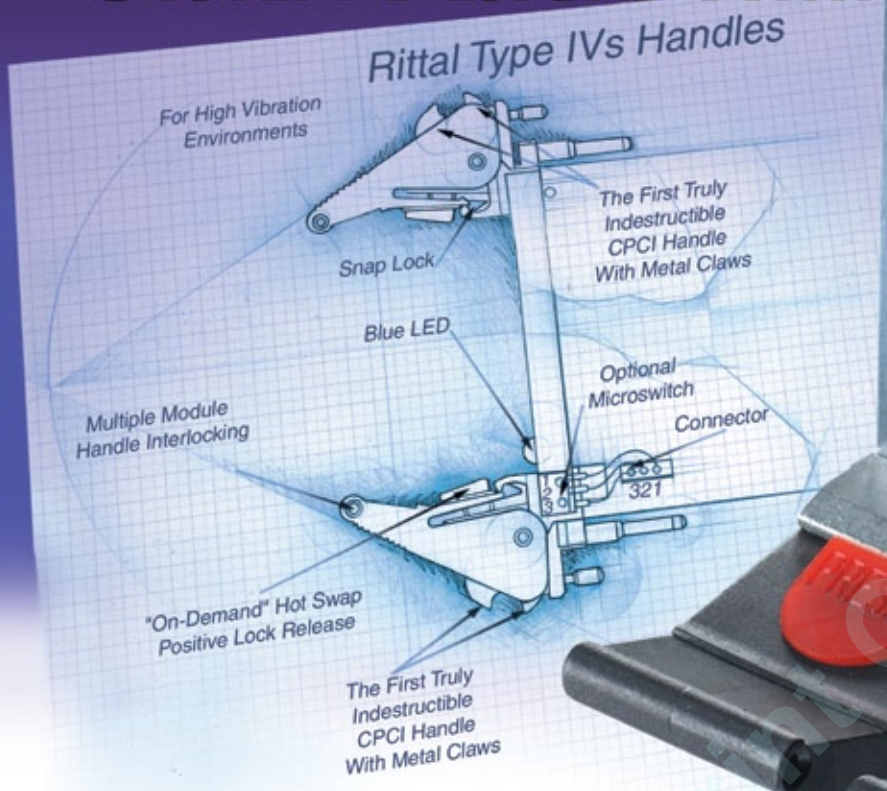

Type IVs Metal Claw Handle Features

- *Indestructible* ***METAL CLAW*** *design for long life*
- *Easy no-nonsense injection/ extraction process for 3U, 6U, even 9U CPCI/VIPA modules*
- *Shock/vibration proof locking (even without screw)*
- *Safe microswitch actuation "on-demand" permits electrical loading effect without activating the inserted board*
- *Type IVs Version I – the classic CPCI hot swap and VME 64 Ext./ VIPA handle*
- *Type IVs Version II – features a .10" offset for increased component side 2 space*

Experience produces knowledge. And no one surpasses Rittal's expertise in designing smart packaging solutions for CompactPCI applications. Like our new generation of indestructible Type IVs CompactPCI metal handles that are capable of injecting/extracting a 6U fully connector loaded board at 118 lbs./525 N. Combine these or any other of our high-performance handles with our front panel customization services and you have a complete CompactPCI front panel solution. So to make your next CompactPCI design a Rittal design, contact Rittal at **(817) 447-9420** or email to **kdubois@rittal-corp.com**.

Rittal's low-profile Telco Type VII metal handles will handle 9U boards up to 180 lbs./815 N and are also available in plastic.

RITTAL®

WWW.RITTAL.COM

DSP-FPGA.com Editor's Choice

Digitizer boards have what it takes for high channel count signal processing

Ultra-wideband signal processing requires getting the real-world signals into and out of the front-end DSP processors with minimal latency and maximum signal integrity. The Triton V-5 from TEK Microsystems combines ultra-fast ADC and DAC channels with three Xilinx Virtex-5 FPGAs, while the company's Orion V-5 boasts twin ADCs with three Virtex-5s. Each board is designed around the ANSI/VITA 41 VXS 6U form factor, which marries the best of industry-standard serial fabrics with interoperable board designs. VITA is the trade consortium that created the FPGA Mezzanine Card (FMC) standard, which is now becoming popular on Xilinx reference designs and mezzanine modules.

The Triton V-5 features a 10-bit, 2.2 GSps ADC and one 12-bit, 3.0 GSps DAC feeding into the company's QuiXilica-V5 architecture, delivering 2,336 DSP slices with 1.3 TeraMAC/s of signal processing. The board is ideal for signal jamming, radar range rate calculations, or data mining. Based on the Euvis MD653D noninterpolating DAC and the E2V AT84AS008 ADC, the board has a common clock but allows multiboard clock synchronization. Using two Euvis MD653D DACs with a 4 GHz bandwidth, the Orion-V5 has similar DSP performance and includes an AC-coupled output that can be single-ended or differential.

TEK Microsystems
www.tekmicro.com • RSC# 41847

COVER PHOTO: Xilinx, Inc. (NASDAQ: XLNX) is the worldwide leader in complete programmable logic solutions, with more than 50 percent market share and US$1.91B in revenues for calendar year 2008. Xilinx award-winning programmable solutions provide design teams with enhanced system-level performance, expanded flexibility and increased design environment productivity for a broad range of applications.

ADVERTISER INFORMATION

Eye of the beholder:
FPGAs convert CMOS imager output to human-viewable RGB

By PAUL **NICKELSBERG**

CMOS imager Bayer-pattern pixel data conversion to RGB pixel data can be readily accomplished using today's low-cost FPGA technology devices. Pixels within modern CMOS image sensors are typically arranged in a Bayer color pattern. Figure 1 presents an example of that pattern.

Image data is scanned from left to right and from top to bottom. Each scanned raw pixel is a monochrome color element, either Red or Green or Blue. Inherent characteristics of the human eye make it advantageous to have twice as many green pixels as red or blue ones.

In order to render a CMOS imager's output human-viewable, a conversion from the Bayer image data format to an RGB image data format is necessary. A common method of converting Bayer image data to RGB data is the bilinear-interpolation method in which a 3 x 3 Bayer-pixel matrix is mathematically combined into a single pixel with separate Red, Green, and Blue data component values. Thus a 3x3 matrix of monochrome pixels is transformed into a single pixel with three component values. The math for bilinear interpolation (as well as algorithms for other image conversion methods) may be readily found in industry textbooks.

At Orchid Technologies Engineering and Consulting, Inc. we implemented the bi-linear interpolation based image conversion using an Altera Cyclone III FPGA. We selected the Cyclone III FPGA device for its low-cost, resource-rich architecture.

We began our FPGA implementation by dividing the complex function up into seven individual building blocks. These blocks were:

- Input Data Formatter
- Image Line Data Memory
- Image Line Data Selector
- Bayer Matrix Calculator
- Output Data Formatter
- PLL Clocking and Control
- Built-in Test Pattern Generator

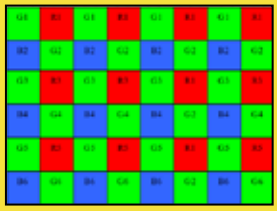

Figure 1 | Modern CMOS image sensors arrange pixels in a Bayer color pattern.

Working with Altera Corporation's Quartus II development platform we implemented and tested each subsystem. Our goal was the development of a 10-bit wide conversion system that preserved data width throughout the process. Preservation of color depth and resolution was specifically necessary for our application. Many low-cost commercial imagers will provide 16-bit wide RGB 565 data, YUV 4:2:2 data, or some other truncated color depth output. Unique to our conversion approach was the 30-bit wide color depth with 10 bits per pixel per color. Our approach was designed for a maximum line width of 1024 pixels at a maximum input pixel rate of 27 MHz.

Our Cyclone III implementation required 1 PLL Block, 462 logic cells, 219 logic registers, 81920 internal memory bits, and 40 I/O pins. Implemented in an EP3C5E144 device, our resource usage was under 10 percent for logic element utilization and under 20 percent for memory bit utilization. Clearly, plenty of resources remain for the implementation of other complex system features. Our image processor was an easy fit for the Cyclone III device. PJ

***Paul Nickelsberg**, President and CTO of Orchid Technologies Engineering and Consulting, Inc., has more than 20 years of experience in electronic products design.*

Orchid Technologies Engineering and Consulting, Inc.
978-461-2000 • paul@orchid-tech.com • www.orchid-tech.com

For Systems
that Demand
the Most

XG-5000K

Our most advanced FPGA Based PC/104 Plus Circuit Board

Features

- Xilinx Spartan 3, 5 Million Gate FPGA
- 264 User Configurable I/O
- 256 MByte of SRAM
- 32 MByte of Flash
- 6 MByte of Serial Page DataFlash
- CompactFlash Type 1 Connector
- Secondary Xilinx Spartan 3, 400,000 Gate FPGA
- 10/100Base-T Ethernet
- 2 RS232 Ports
- 25MHz FPGA Master Clock Source
- Secondary 0 - 25MHz user programmable DDS FPGA Master Clock Source
- Incorporates Xilinx's design revisioning technology and can retain onboard as many as 16 partial or 4 complete FPGA design BIT files
- Can be powered from the PC/104 connector or an external DC power supply
- PC/104 Plus Form Factor; ISA and PCI 33MHz Bus Interfaces
- Extended temperature -40C to +85C

Jacyl Technology is a Team Member under Lear Siegler Services, Inc for the US Governments CECOM Rapid Response Contract

JACYL TECHNOLOGY

Advancing Today's Technology into Tomorrow

www.jacyltechnology.com

Open source, high availability middleware for the military environment

By John Fryer

U.S. Marine Corps photo by Lance Cpl. Kelsey J. Green

High availability is a key requirement for many military embedded systems. To meet this need, OpenSAF, an open source middleware implementation of the Service Availability Forum's high-availability specifications, provides new options for this critical software layer. This open source middleware offers several benefits, comprises the HPI and AIS open specifications, and is highly suitable for use in security-conscious environments.

The availability of systems in many harsh environments is an essential requirement of continuous operation. Because military systems must increasingly leverage COTS technology to provide superior performance in all situations while optimizing life-cycle costs, the availability of these systems is critical. Consequently, implementing high-availability middleware – such as that developed by the OpenSAF (Open Service Availability Framework – www.opensaf.org) – is an essential aspect of meeting this challenge.

Early versions of the Service Availability Forum (SA Forum) specifications are already deployed in some updated aspects of the Aegis Combat System, with the potential for broader deployment of additional features in future technology refreshes. Other examples of applicability include radar control systems, unmanned system ground control stations, and control systems on ships, aircraft, missiles, and space shuttles. Meanwhile, there are several important aspects to understand for this high-availability middleware approach, including the benefits of using open source, the open specifications available, and how these specifications can be used effectively in a security-conscious environment.

Benefits of open source

High availability is not a new concept in many industries and is often COTS-driven these days. Defense systems have incorporated redundancy in many areas as a fundamental aspect to continuous operation in battlefield situations. With increasing use of and reliance on technology, it is essential that systems can survive catastrophic failure and continue to operate seamlessly.

The traditional approach to high availability has been to design to specific system requirements, particularly when proprietary or specialized hardware has been used. While this might yield an optimal solution for a particular system, it generally results in software systems that are tailored to the specific hardware and architectures where the software applications are tightly integrated with high-availability middleware. Extension of these systems requires the expertise of the design team who understands the philosophy and implementation behind the system. As a result, open source works best in environments where functionality is necessary but where multiple proprietary or commercial implementations hinder the overall objective.

Open specifications are key

With the increasing use of COTS and open specification technologies such as AdvancedTCA and MicroTCA, the perspective is radically different, and optimization can only be achieved if the high-availability environment is common to all applications' interfaces and underlying implementation. Thus, the SA Forum has developed two sets of

OpenSAF: From proprietary to COTS

The OpenSAF project used an implementation from Emerson Network Power Embedded Computing as its starting point. This code base had already been deployed in telecom networks in proprietary form, but has undergone additional extensive testing and modifications by multiple contributing companies to become available as OpenSAF. The result is a rapidly maturing code base implemented on platforms including AdvancedTCA, MicroTCA, enterprise servers, and proprietary hardware. The AIS and HPI specifications are key and make OpenSAF appropriate for use in security-conscious military environments.

high-availability middleware specifications that provide the layering of high-availability services: the Hardware Platform Interface (HPI) and the Application Interface Specification (AIS).

HPI: Low-level abstracting

HPI is used to abstract low-level information from hardware so that it can be accessed and programmed through common interfaces. This enables applications directly accessing hardware functions and receiving hardware events to run on multiple platforms with minimal modification. Indeed, HPI is now implemented in many commercial and proprietary platforms and is viewed as a market success. HPI exposes a set of platform-defined management instruments, examples of which are shown in Figure 1. Through the HPI interface, the various instruments can be read and configured. Common application triggers, such as voltage drops or watchdog timer expirations, constitute failure "events," which serve as inputs to AIS high-availability middleware. The specifications also allow for instrument grouping to create resource records that can then be further grouped into domains with a common set of capabilities.

AIS serves up high availability

AIS is significantly more sophisticated as it provides the set of services necessary to support highly available software applications, as shown in Figure 2. All high-availability middleware implements most or all of these services, as they are fundamentally necessary for an "always on" system. What is different is the layered approach and open forum collaboration to create application- and platform-agnostic architectural models with a rich set of APIs. Remember that AIS is driven and configured by its application environment, and it is the common approach to the middleware that enables rapid portability across multiple systems and between multiple applications.

The OpenSAF project

Open source and commercial implementations of the HPI and AIS specifications are now available in the marketplace. OpenSAF is actively backed by the OpenSAF Foundation, which provides financial, legal, and marketing support. The open source project itself is open to anyone who wishes to participate and is organized by industry experts who form a technical leadership council. Actively participating companies include Ericsson, Huawei Technologies, Nokia Siemens Networks, HP, Emerson Network Power, Enea, and Wind River Systems.

OpenSAF Release 2 provides capabilities beyond the scope of the SA Forum, such as porting and messaging infrastructure, and it also provides additional functionality such as a message-based checkpointing service. Release 2 is stable and field deployed with the communications industry.

Release 3 of OpenSAF is currently undergoing testing within the OpenSAF community. It is targeted for release in the second quarter of 2009. The flexibility of OpenSAF will be extended to include more ports and hardware platforms beyond the current base of AdvancedTCA, enterprise and proprietary hardware, and a broad array of Linux variants.

Security-conscious environments meet OpenSAF

OpenSAF is available under the Lesser GNU Public License v2.1 (LGPL v2.1), a

critical consideration for defense industry adoption. Under LGPL v2.1, any applications that link to the code base and do not modify the code base itself are not subject to the license terms. LGPL v2.1 states that modifications made to the code base must be made available for free, if requested, although there is no obligation to offer these as contributions to the OpenSAF project. The license also enables multiple companies from an industry segment to contribute to the project without fear that a competitor can derive direct financial benefit from the software, as no charge can be made for direct licensing of the code base.

The OpenSAF mailing lists are mechanisms to contribute fixes and enter into discussion with experts on specific applications of OpenSAF. For most organizations, the preferred method is to work with companies who offer commercial distributions of the OpenSAF code base, which include training support and services.

OpenSAF gains momentum

The modern military environment calls for highly available systems in an increasingly technological environment, and it might be impractical for trained technicians to be present in all situations. The transition to COTS and accelerated technology refresh rates means that standardization of high-availability middleware is an increasing requirement. The SA Forum HPI and AIS specifications provide the groundwork for this transition, and they have already been deployed in early versions in the military world.

The OpenSAF open source project provides an active and robust ecosystem, using a proven stable implementation of the AIS services as a starting point to accelerate development of a common high-availability infrastructure in a collaborative environment. This enables sharing of resources and an increased focus on additional functionality. It also provides a choice of how to adopt OpenSAF, either directly by downloading the code or through a commercial distribution, similar to the various Linux models. In this environment, OpenSAF is an increasingly interesting option for an SA Forum based middleware implementation.

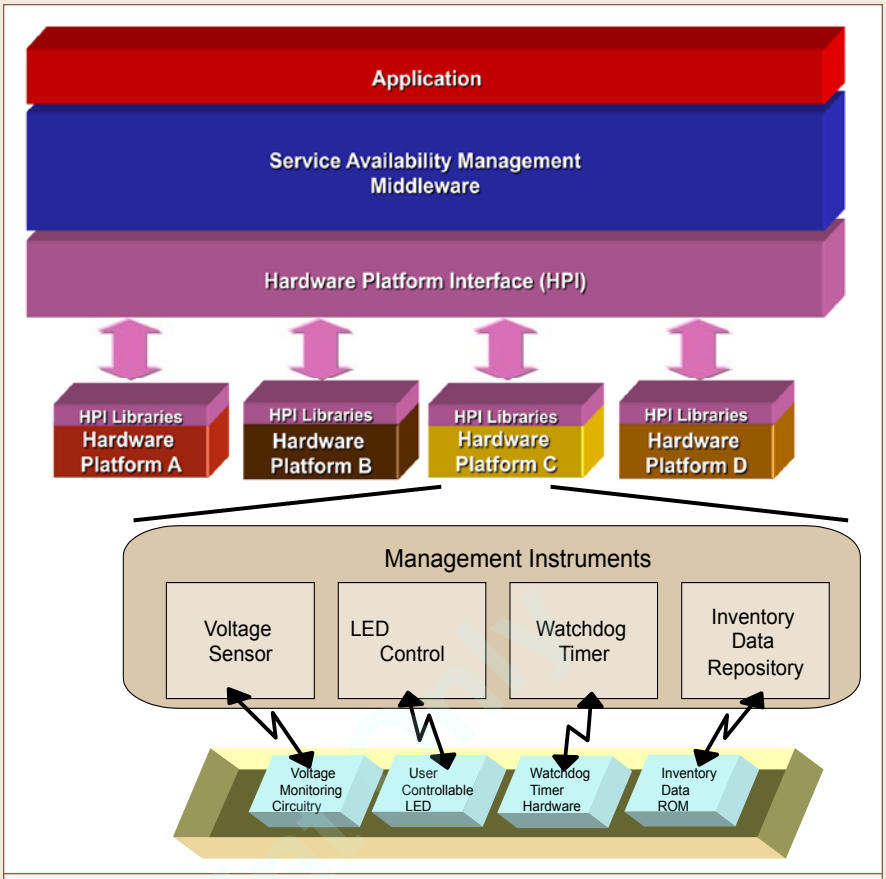

FIGURE 1: HPI exposes a set of platform-defined management instruments.

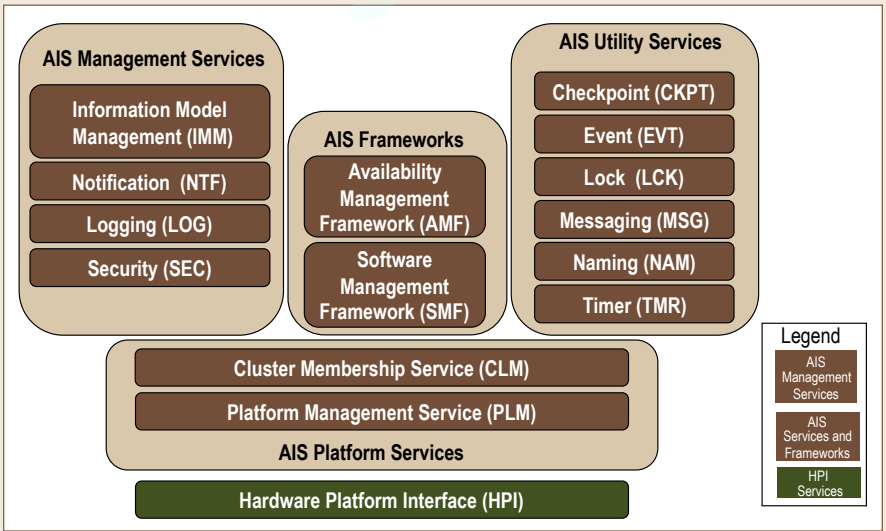

FIGURE 2: AIS is significantly more sophisticated than HPI and provides the set of services necessary to support highly available software applications.

John Fryer is president of the OpenSAF Foundation and serves on the Service Availability Forum's board of directors representing Emerson Network Power, where he is director of technology marketing for Emerson's Embedded Computing business. Previously, he was responsible for the worldwide product marketing of AdvancedTCA platforms at Motorola. Prior to joining Motorola, he was vice president of marketing for control plane and data plane software applications at NetPlane Systems. John has more than 25 years of experience in the communications industry in a variety of marketing and engineering positions. He holds a B.Sc. with honors in mathematics from the University of Nottingham, England. He can be contacted at John.Fryer@Emerson.com.

OpenSAF Foundation
925-964-9298
www.opensaf.org

Military EMBEDDED SYSTEMS ★★★★★ Editor's Choice

Editor's Choice Products

Editor's note: Military Embedded Systems *is "hip" to the whole Web 2.0 social networking revolution. While we don't know which of today's buzzy trends will last, we're going to start including links to vendors' social networks, when provided. You can also reach us on Twitter, Facebook, and LinkedIn ... and that's just for this week. Next week there'll undoubtedly be more new sites.*

When your grip on power fluxuates: Use 54 KW AC sources

So *embedded* usually means DC-powered electronics sometimes run by batteries. But military systems are different: They can require high DC voltages (28 VDC for vetronics) and even high-current AC in many airborne platforms. Alternating current is more efficient in these instances and is eventually converted into DC for the electronic suites. Chroma Systems Solutions offers a way to rig three 18 KW 1- and 3-phase units in parallel to source up to 54 KW. The lab testing-oriented 61500 and 61600 series are programmable and use advanced DSP algorithms to generate very clean sine waves with a total distortion that's less than 0.5 percent at 50/60 Hz.

Designed to test aerospace systems in accordance with MIL-STD-704F, RTCA DO-160D, and ABD100, the units offer power line disturbance simulation, programmable output impedance, wave-shape synthesis, and the ability to simulate harmonic components in the waveforms to simulate dirty power often found in deployed applications. Front-panel programmability includes an LCD and keypad, along with remote control via GPIB, RS-232, USB, or Ethernet. LabVIEW drivers are even available. Despite reviewing the company's website, we were unable to determine the differences between the 61500 and the 61600, though perhaps the 61500 allows precompliance test software configuration.

Chroma Systems Solutions • www.chromausa.com
Twitter: http://twitter.com/chroma_usa • RSC# 42075

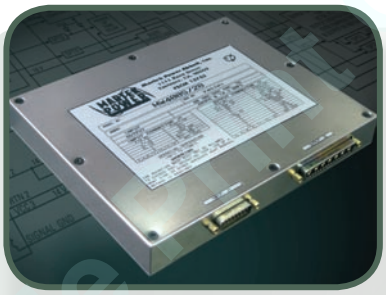

All-in-one AC-DC brick builds rugged military systems

Have an AC-powered system that needs to supply DC voltages from 2 VDC up to 28 VDC? Martek's 400 W MW400S power module with active Power Factor Correction (PFC) is designed specifically for A&D as well as industrial apps. Feed the unit 90-265 VAC, single phase, at 47-440 Hz, and it'll deliver 2, 3.3, 5, 12, 15, and 28 VDC outputs at .99 PFC with 82 percent efficiency at full load. Pretty impressive. And it's clean, too. Line harmonics meet MIL-STD-1399 for 60 Hz or 400 Hz, and EMI meets CE101 and CE102 of MIL-STD-461.

Designed to operate from -40 °C to +100 °C (from zero to full load), the brick can be remotely turned on and off via TTL. MW400S also has features that microelectronics designers love. They include: output voltage trim, overvoltage/overcurrent protection, AC and DC "good TTL signal," and a fixed-frequency (500 to 550 KHz) conversion with synchronization input. Even better, the module is reasonably small considering all these features and power: 6.5" x 8.0" x 1.0" (165 mm x 203 mm x 25.4 mm).

Martek Power • www.martekpower.com • RSC# 42093

High-functionality, rugged Core 2 Duo CompactPCI SBC

With all the hubbub surrounding VITA's large I/O VPX moving to OpenVPX and then to VITA 65, A&D designers are understandably wary. Instead, they might reconsider 6U CompactPCI's high-pin-count backplane. Supporting PICMG's 2.16 Dual GbE Packet Switching backplanes – the "original fabric" implementation – Concurrent Technologies' PP 452/03x uses one of two Intel Core 2 Duo CPUs to move data around. The two PMC/XMC sites add additional I/O functionality, and the board is available in ruggedized and conduction-cooled variations. Soldered CPU options include Intel's L7400 (1.5 GHz) or ULV U7500 (1.06 GHz), each with 4 MB of L2 cache for 64-bit operating systems.

Chipsets are Intel's E7520 and 6300ESB ICH – both designed with lower-power applications in mind. The CPU can talk with up to 4 GB of DDR2-400 ECC at up to 6.4 GBps. I/O from the PMC/XMCs routes down via 66 MHz PCI-X and x4 PCIe, and XMC site number 1 also supports x8 PCIe. Front- and rear-I/O options abound. PP 452/03x features two SATA150, one RS-232, three USB 2.0, and dual GbE ports. We already mentioned PICMG 2.16 (Ethernet fabric), but there's also PICMG 2.9 (IPMI) and PICMG 2.1 (hot swap). An onboard 4 GB flash disk is included, as are options for EIDE CompactFlash. If your system needs this kind of server-class I/O and processing while operating on a power budget of less than 35 W, consider this one.

Concurrent Technologies • www.gocct.com • RSC# 42092

Flight-qualified rugged SFF PC

For years Parvus has been bringing rugged Small Form Factor (SFF) products to the civilian and military markets. The company's latest effort, the DuraCOR 820 shown here, is a svelte 3" high and weighs about 3 lbs. The "little cube that could" is designed for command and control on the move, unmanned vehicle operator control, or C4ISR applications. Stuffed with a PC/104-*Plus* Pentium M-based SBC, the rugged cube is fully qualified to MIL-STD-810F for altitude, thermal profile, shock, vibration, and humidity. It can operate over -40 °C to +71 °C, withstand 15 *g* operating shock, 95 percent humidity, 60,000 feet altitude, and survive many aircraft and helicopter vibration profiles.

Equipped with an Intel 1.4 GHz Pentium M CPU and SSD preloaded with WinXP or Linux, the watertight cube has multiple I/O options on the front panel, including dual 10/100 Ethernet, MIL-STD-704E power input, analog video, three USB ports, two RS-232 ports, digital I/O, keyboard, and mouse. Since the inside modules are PC/104-*Plus*, additional I/O is easily added, such as a 12-channel GPS, additional serial RS-232/422/485, multiple 1553 channels, or even customer-specific I/O.

Parvus • www.parvus.com • RSC# 36887

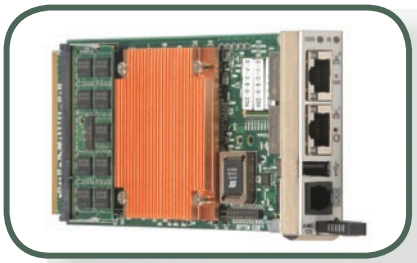

Single board cluster, LAN, WAN computer

MicroTCA continues to buzz around military applications because it brings leading-edge communications capabilities from civilian land and cellular networks into Aerospace and Defense (A&D). When bolted atop a 1U MicroTCA carrier, Performance Technologies' AMC123 provides a ready-to-go, out-of-the-box network designed for LTE, WiMAX, media and signaling gateways ... or military comms systems. The x86-based module uses Intel's EP80579 integrated processor, one of the first of Intel's reentry into the microcontroller market a la their former 80196 series. This Pentium M-based MCU includes a 32 KB L1 (split for I and D), 256 KB two-way L2, integrated I/O controller hub, four-channel EDMA, SpeedStep power and clock management, and a 400/533 MHz FSB.

Up to 16 GB onboard flash stores the OS, program, and data sets, while up to 4 GB of PC2-5300 DRAM with ECC provides scratchpad storage. Flash can also be configured to run via USB 2.0. Three GbE ports can be routed as: one to the front panel and two to the AMC card edge, or two at the front panel and one to the card edge. AMC123 is also equipped with a front-panel USB 2.0 port and status LEDs, plus an eight-lane PCIe and two SATA ports routed to the AMC connector. But hardware's just hardware. What makes the AMC123 special is Performance Technologies' system features, including NexusWare, which is the company's family of Linux software products. There's the Carrier Grade Linux package, NexusWare Studio for software development, and a remote systems management tool called NexusWare Portal. After all, Performance Technologies is the company that brought PICMG 2.16 Ethernet to the masses, along with myriad other innovations now standard across the communications and VME industries.

Performance Technologies • www.pt.com • RSC# 42065

RF transceiver that records, plays back

No question about it: One can build a very capable RF front-end transceiver out of VME and PMC modules. In fact, Pentek makes some very good boards. But if you want to record, analyze, and play back received data ... well, now you've got to add an SBC, an operating system, some disk controllers, a chassis, cables, yadda, yadda. For laboratory and "gently rugged" deployments, Pentek offers a better solution: the RTS 2701 Real-Time Data Recording and Playback System with Multiband Transceiver. This Windows XP workstation is designed to be bolted into a 19" rack or used on a desk. The ready-to-use PC instrument uses a multiband transceiver that's amazingly flexible. Voice at 8 KHz can be input, all the way up to 60 MHz for wideband radar. Narrow or wideband inputs are allowed, and programmable IF frequencies provide for great RF variety.

Twin-channel inputs feed 125 MHz A/D outputting 250 MBps, which eventually gets streamed to up to 5 TB of NTFS disks (RAID 0, 1, 45, 6, 10, and 50) at an aggregate sustained 500 MBps. The system includes a bevy of out-of-the-box software to accomplish the mission, including Pentek's SystemFlow package with client/server GUI, digital oscilloscope, and spectrum analyzer. Data can be time stamped, and other real-time data can be appended. Finally, the unit can be remotely mounted and connected via Ethernet to allow Internet access via a handy built-in API. What impresses us most about this system is that Pentek appears to have thought of everything. Who would want to try to build one of these, when you can just buy it?

Pentek • www.pentek.com • RSC# 38752

Editor's Choice Products are drawn from OSM's product database and press releases. Vendors may add their new products to our website at http://submit.opensystemsmedia.com and submit press releases at http://submit.opensystemsmedia.com. OSM reserves the right to publish products based on editors' discretion alone and does not guarantee publication of any product entries.

Crosshairs Editorial

MIDs course correction
Intel buys Wind River

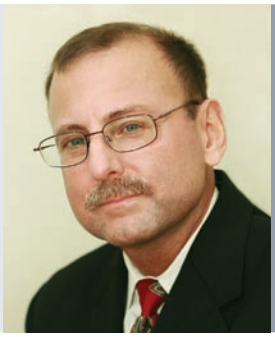

By Chris A. Ciufo, Editor

The announcement yesterday that Intel was acquiring Wind River Systems (ticker: WIND) for approximately $884 million caught the industry by surprise. For us at *Military Embedded Systems*, it was literally a 'Stop the presses!' moment as we pulled back one story in order to cover this significant breaking news.

First, a disclaimer. I'm a huge fan of Intel, and I frequently write glowing things about the company because I'm convinced they're the primary technology driver for the embedded industries we follow here. As a result, they invite me to their annual "Intel Developer's Forum" and give me front-row seating[1]. We're buds, so to speak. As for Wind River Systems (WRS), I tell the story of getting a frantic call from founder Jerry Fiddler asking for a replacement copy of his own PowerPoint presentation because of a laptop crash. I sent it to him immediately. So yeah, I like Wind River Systems, too. Sounds like one big happy relationship, doesn't it?

Except for this nagging fact: I can't figure out the reason for this acquisition. Intel offered $11.50 a share for WRS, which was a whopping 43 percent increase over the company's previous day's price. Today, JP Morgan downgraded WRS, causing speculation that there may be other buyers willing to pay more, or – as other analysts speculate – maybe Intel is paying too much. But I'm not an investment journalist, and this isn't my area of expertise. Go call your broker for financial advice.

However, I *am* a technology expert. I know Intel's dying to get into the 15 *billion* unit Mobile Internet Devices (MIDs) forecast, I just don't see how paying nearly a billion dollars for Wind River is good for either company. What's clear, though, is that as the desktop and server markets saturate, Intel needs other sources of revenue. The good news is: The company has market opportunities aplenty, and the technology *cojones* to succeed. When Transmeta made low power (instead of clock speed) the key CPU differentiator, Intel wiped the floor with them with the Centrino and Pentium M. Ditto with multicore dabbled in by MIPS, Freescale, and ARM, giving us Core 2 Duo and Quad, Core i7, and now SoC devices for digital TV and other high-volume consumer markets. Well, all but *one* really important market: cell phones, and increasingly *smart phones*. Intel's foray into smart phones such as Samsung's i600 with the Xscale ARM derivative ended when Marvell bought the line for $600 million in 2006 (according to Jordan Robertson, AP technology writer). Intel's awesome new Atom CPU, the most frugal of which sips a mere 0.65 W, is a definite ARM contender in the 164 million unit smart phone market (per Forward Concepts; Will Strauss). But Intel's been married to Microsoft for years – though Windows CE/Windows Smart Phone/Mobile isn't as sexy as RIM's BlackBerry or Apple's iPhone (running OS X, likely on a Samsung-based ARM ASSP).

So what does Wind River bring to the smart phone market? I suppose Wind River's Linux distro is a possibility, although I'd speculate why not buy MontaVista's or TimeSys' offerings for much less? Or, if they want a light OS for small devices, why not Express Logic's ThreadX, QNX, or even Mentor's Nucleus? Perhaps Intel wants to try Apple's H/W + S/W approach – but Apple makes the whole *phone*, not the ICs and software combined.

For non-desktop/server markets, that still leaves telecom (Intel sold off their Communications Alliance/PICMG boards to RadiSys, and their Dialogic and Ziatech acquisitions), medical, automotive, "other consumer embedded," plus Aerospace and Defense (A&D). I doubt Intel's interested in Wind River's A&D dominance, though darn near all new Intel embedded devices go on their seven-year life-cycle roadmap, which appeals to the military. Let's put automotive aside for the moment, although custom, long-annuity MCUs and CPUs with bolt-on software *could* be Intel's plan in automotive. No, I really think that Intel sees the "other consumer embedded" as the target. Those mobile Internet devices basically are everything that's attached to what Intel calls "The 4th Stage of the Internet": appliances, kiosks, handheld doodads yet to be designed, telemedicine, games, and so on. Let ARM have the cell phone market (for now), position and posture for the global auto industry's resurgence, keep A&D happy, but provide custom software for the huge-and-growing "other" category. At last year's IDF, Intel forecasted "15B Devices by 2015" (see figure). By adding system value through software IP, Intel can command higher prices for its ICs, or still win designs when the IC choice is a price-sensitive commodity. This latter strategy is the only thing that makes sense to me at this moment.

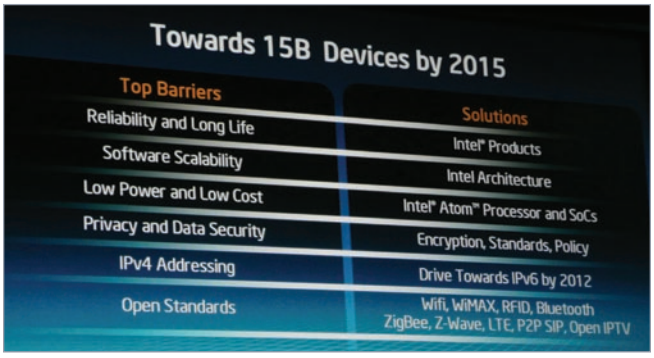

Still, Intel's smarter than I am. You can bet that for $884 million they did lots of due diligence discovery. I recall an industry social event where I asked Intel's then-CEO (now chairman) Craig Barrett about the 3GIO fabric replacing the PCI bus. He smiled and politely said: "No, you want to watch PCI Express." At the time, PCIe was a clearly inferior underdog. But once Intel threw its weight behind it, the rest is history. With Wind River, we'll have to wait and see if the company puts the wind in Intel's mobile Internet device sails, or blows them off course.

Chris A. Ciufo, cciufo@opensystemsmedia.com

[1] To the extent that I received a cameo appearance in the TV show "Orange County Choppers" when Intel was presented the *Intel Chopper* in 2007.

MLRS Multiple Launch Rocket System AH-64 Helicopter T-6A Peloton
CV90 Armored Vehicle DDG-1000 Multi-Mission Destroyer Roland Air Defense System
HIMARS Artillery Rocket System B-2 Stealth Bomber F-35 Lightning II
LHD Class Amphibious Assault Ship Expeditionary Fighting Vehicle Gripen Fighter
Littoral Combat Ship A400M Transport NSSN Virginia Class Submarine
F-117 Stealth Fighter MEADS Air Defense System EMB-145 Challenger 2 Tank
M2/M3 Bradley C-5 Transport NH-90 ASW Helicopter AH-1W Helicopter
Predator RQ-1 E-2C/D Early Warning Aircraft Neuron UCAV V-22 Osprey

Which of these platforms use GE Fanuc hardware?

Fire Scout UAV F-15 Fighter SSN Astute Class Submarine Tornado Fighter
Nimrod Aircraft RQ-4B Global Hawk AMX Fighter P-3C Maritime Patrol Aircraft
C-17 Globemaster K1A1 Tank LPD 17 Landing Platform B-1B Bomber
Taranis UCAV Avenger Air Defense System F-22 Raptor C-130 Transport
Patriot Missile System 737 Wedgetail Arleigh Burke Class Destroyer
E-3 AWACS M1A2 Abrams Tank AV-8B Harrier II Plus Eurofighter Typhoon
F-16 Fighter Merlin ASW Helicopter Ticonderoga Class Cruiser T-6B Trainer
EA-6B B-52H Long Range Multi-Role Bomber Barracuda UAV Demonstrator

ANSWER | **GE Fanuc offers more COTS products than any other embedded company and we are part of ALL these platforms. To find out more, and join this growing list, visit www.gefanucdefense.com**

GE Fanuc
Intelligent Platforms

YOUR DATA IS CRITICAL. DEPEND ON US.

Whether on the battlefield or at the base, the ability to record, store, retrieve, playback and analyze your data is a must. Curtiss-Wright provides solutions for all of your secure recording and storage needs.

Vortex Data Recorders

- Data recording, playback and analysis at sustained rates up to 800MB/s in VXS (VITA 41), VME, CompactPCI and PC form factor
- Customer programmable or application specific solutions with a wide range of inputs supported
- Virtually limitless storage capacity with solid state drive (SSD) and hard disk drive (HDD) options

SANbric High Altitude Rugged Storage

- Rugged 2.7TB Storage System & Shock Isolation Units
- Removable drive module with easy insertion/extraction handle
- Dual Fibre Channel Interfaces

Find out more about data recording solutions by contacting us at +1(937) 610.5457 or visiting www.cwcembedded.com/recorders.

CURTISS WRIGHT Controls
Embedded Computing

cwcembedded.com/recorders
recorders@cwcembedded.com

Vortex
Data Recorders

Vortex VXS
Data Recording & Playback Engine & GUI

SANbric
High Altitude, 2.7TB Rugged Storage

DATA RECORDING & STORAGE... *ABOVE & BEYOND*