OpenSystems Publishing

BONUS SHOW DISTRIBUTION: GSPx, MILCOM, SDR FORUM

The COTS Technology Authority

Military EMBEDDED SYSTEMS

VOLUME 2 NUMBER 2
SUMMER 2006

IN THIS ISSUE:

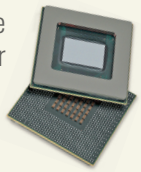

WWW.MIL-EMBEDDED.COM

Special Issue

Trends in Software-Defined Radios

EXCLUSIVE INTERVIEW

Col. Steven MacLaird

Former JTRS JPO Prog. Exec. Dir.

Can FPGAs replace stand-alone DSPs?

PM40048627

PRST STD
U.S. POSTAGE
PAID
BLAINE, WA
PERMIT NO. 226

Keep your SPARC supply lines open.

USP IIe™ VMEbus Computer

650-MHz UltraSPARC® IIi+ CPUs
4GB SDRAM
Optional TGA3D+™ and TGA-100™ graphics support
Gigabit Ethernet
PCI expansion to four slots
Four serial ports
VME Interface – VME64X via Tundra Universe II
PS/2 or USB keyboard/mouse
Solaris™ 8/9/10 Support

Stock up on Themis SPARC SBCs while supplies last.

USP IIIi™ VMEbus Computer

Two UltraSPARC IIIi 1.2 GHz CPUs
4GB 266DDR SDRAM per CPU, 8GB total
Optional TGA3D+ and TGA-100 graphics
Gigabit Ethernet port
One 64-bit/66-MHz PMC slot
Up to four PMC expansion slots
Dual FC-AL/SCSI ports
Additional VME-backplane access
Up to 30G shock
Solaris OS Support

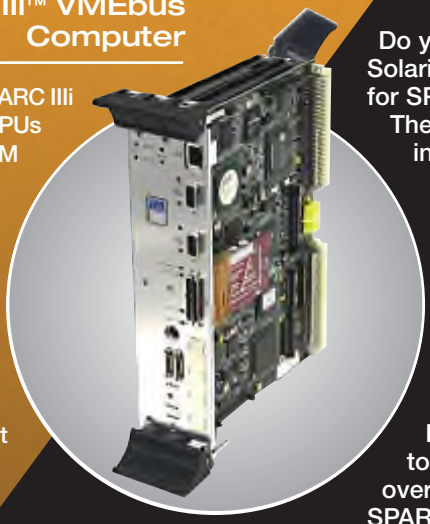

Do your mission critical applications depend on SPARC/Solaris solutions? Many suppliers are ending their support for SPARC-based single board computers. Not Themis. Themis is committed to providing its customers and industry with high performance UltraSPARC VMEbus single board computers.

Themis has solutions. We are committed to providing you with the industry's widest breadth of high performance UltraSPARC-based VME single board computers. To support thousands of Solaris platform solutions for the most demanding mission critical applications.

Don't be forced into changing your computing platform today. Extend the life of your system and gain control over your program lifecycle and budget with Themis SPARC solutions.

Themis is keeping the SPARC supply lines open so call Themis today.

www.themis.com (510) 252-0870

THEMIS

Transformational.

RSC# 2 @www.mil-embedded.com/rsc

Discover Condor's Avionics Embedded Solutions

It's easy to see why we're the leader.

1553 | 429 | AFDX | MMSI

- Rugged and reliable
- High performance
- Advanced FPGA technology
- VxWorks, Integrity, LynxOS & other RTOS supported
- Test and integration tools
- Intellectual property cores
- ISO 9001:2000 certified
- Outstanding technical support

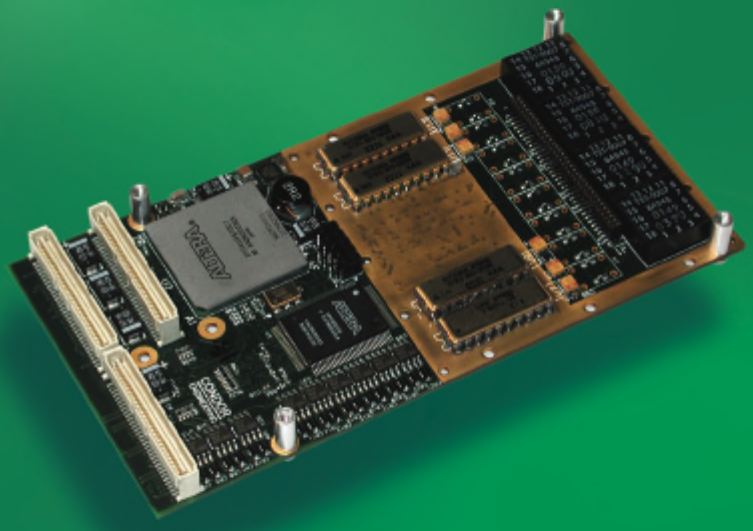

Go with the leader.

805.965.8000 • **www.condoreng.com**

RSC# 3 @www.mil-embedded.com/rsc

CONDOR ENGINEERING

Military EMBEDDED SYSTEMS

VOLUME 2
NUMBER 2

SUMMER 2006

www.mil-embedded.com

DEPARTMENTS

EVENTS

MILCOM 2006
Oct. 23-25, 2006 • Washington, DC
www.milcom.org

GSPx
Oct. 30-Nov. 2, 2006 • Santa Clara, CA
www.gspx.com

SDR Forum
2006 Software-Defined Radio Technical Conference and Product Exposition
Nov. 13-17, 2006 • Orlando, FL
www.sdrforum.org

COVER
The U.S. Army's Land Warrior program calls for network-centric assets relying on a *Soldier of One.* Lightweight reconfigurable Joint Tactical Radio System (JTRS) equipment is essential for digital information exchange – depending on critical DSP, FPGA, and COTS software. Related articles appear on pages 12 and 22. (Image courtesy of the U.S. Army)

On the cover
The IBM Cell processor implements a 64-bit PowerPC core with eight additional synergistic cores, allowing up to 10 simultaneous threads and more than 128 outstanding memory requests.

Published by: OpenSystems Publishing™

FEATURES

E-LETTER

Visit us online for the complete list of articles at www.mil-embedded.com/eletter.

WEB RESOURCES

Subscribe to the magazine or E-letter:
www.opensystems-publishing.com/subscriptions

Industry news:
Read: www.mil-embedded.com/news
Submit: www.opensystems-publishing.com/news/submit

Submit new products:
www.opensystems-publishing.com/vendors/submissions/np

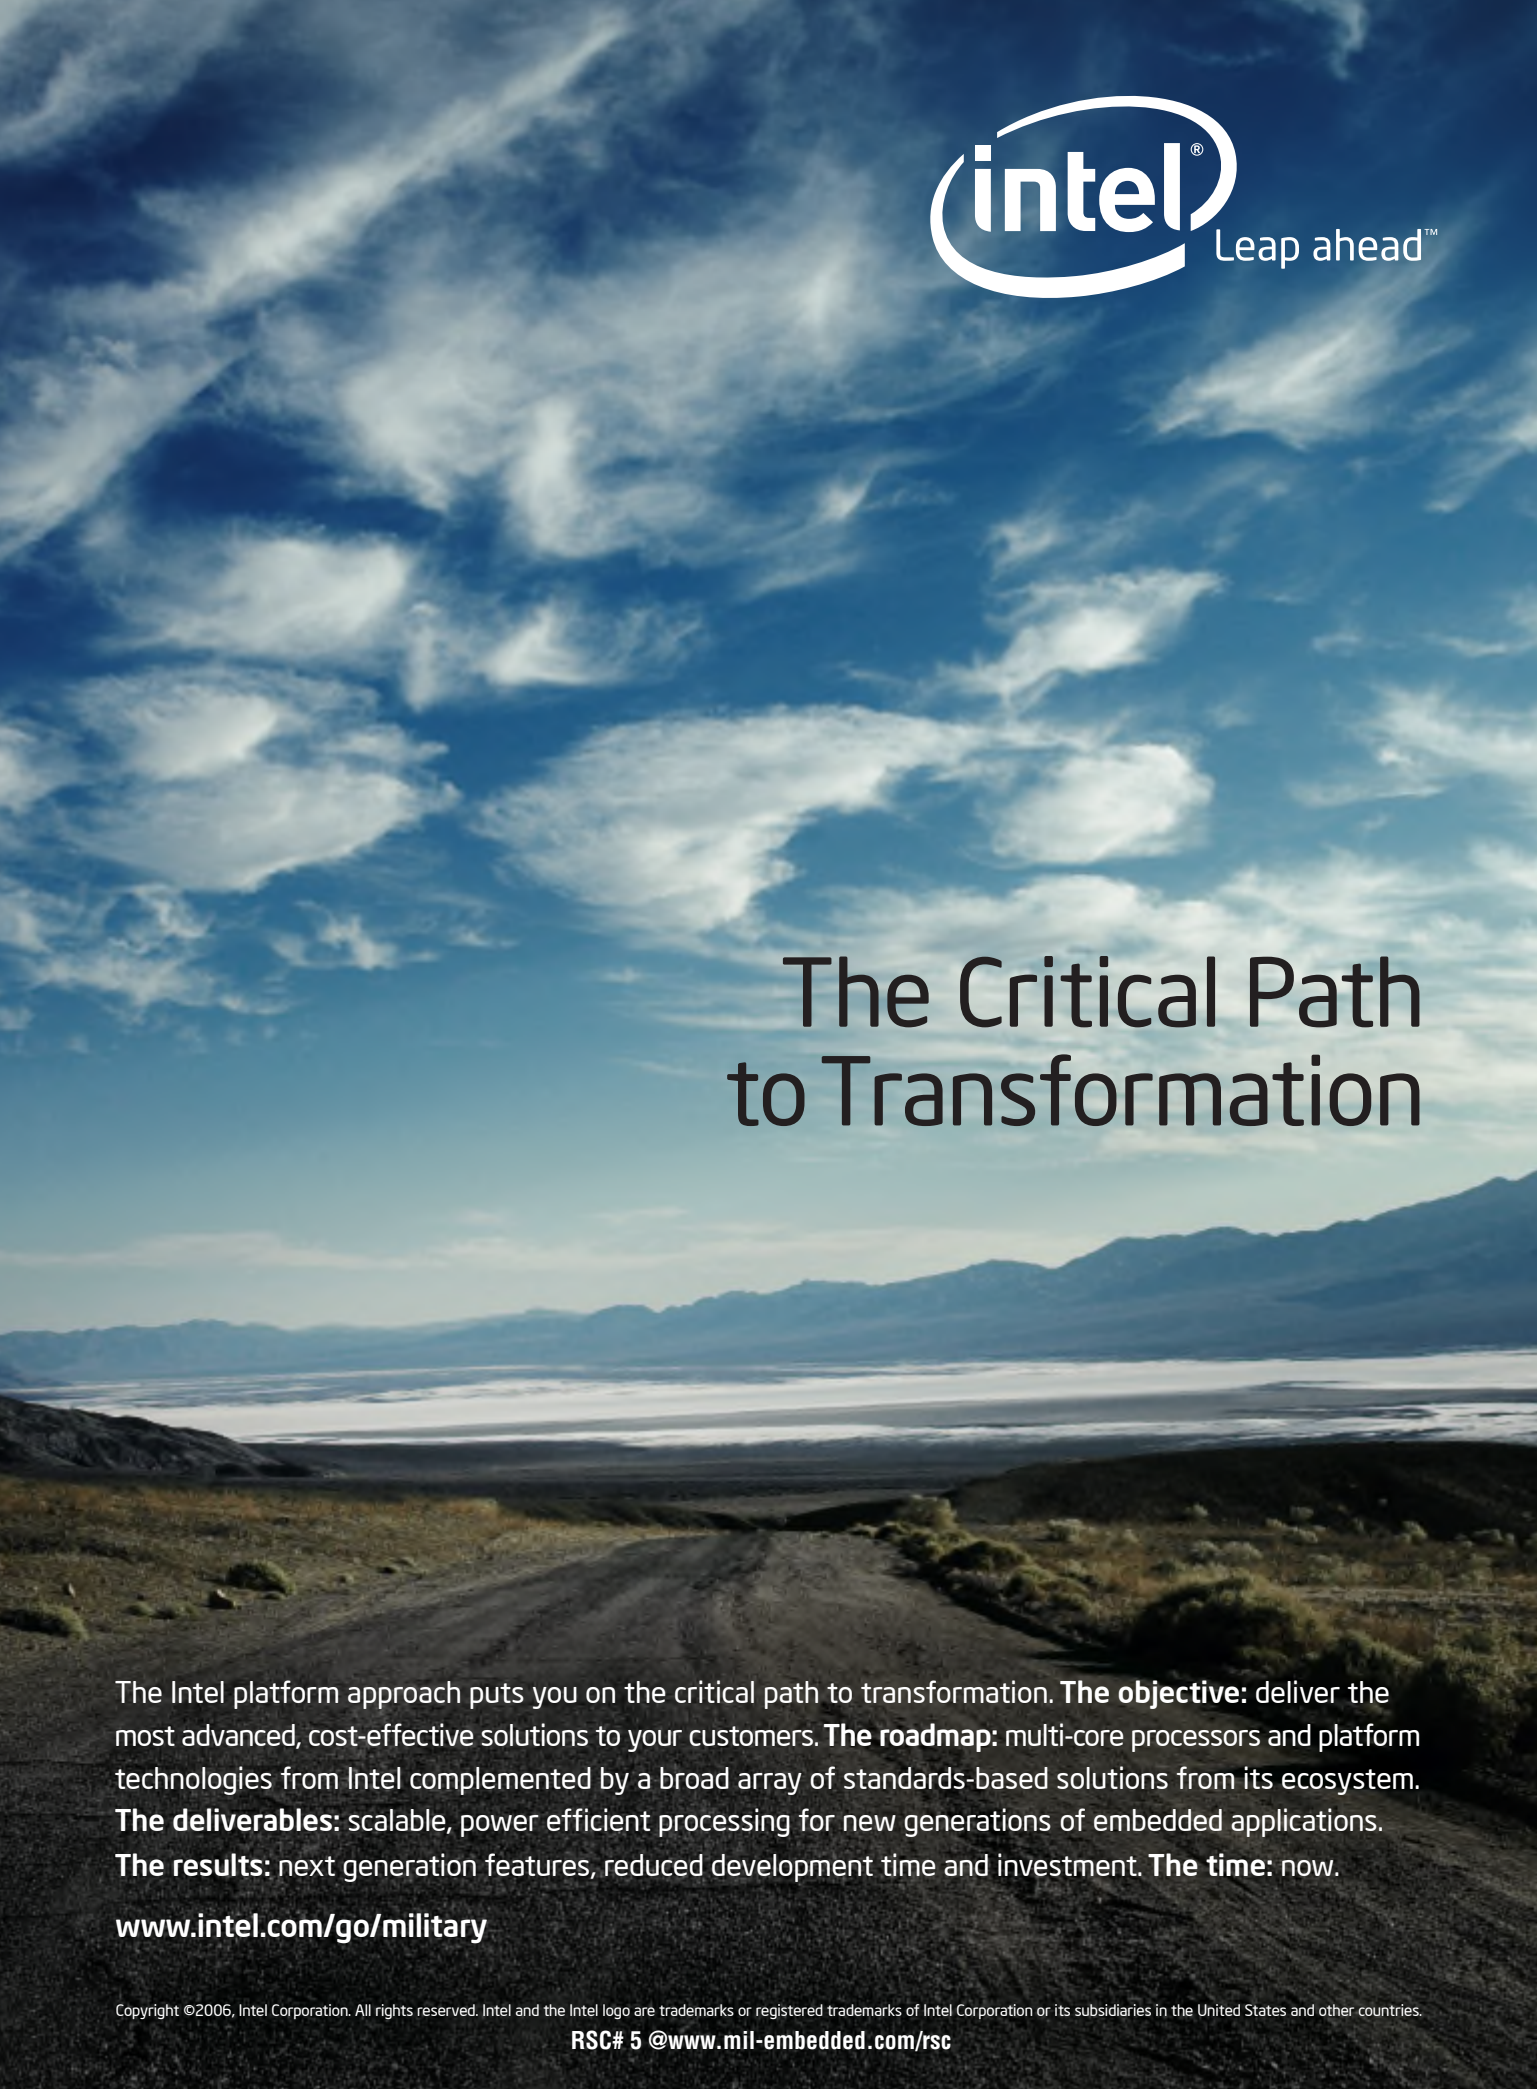

RSC# 5 @www.mil-embedded.com/rsc

Advertiser Information

Page/RSC#	Advertiser/Product description
52	ACT/Technico – PMC Storage
48	Advantech Corporation – Stackable SBCs
17	Annapolis Micro Systems – FPGA Systems
46	BittWare – 6U VME/VXS Board
3	Condor Engineering – Avionics Embedded Solutions
56	Curtiss Wright – High-performance Graphics Products
2701	Datametrics – Rugged Workstations
39	DIGITAL-LOGIC – Microspace MSM855
44	Embedded Planet – EP8548A Serial RapidIO AMC
31	Excalibur Systems – Avionics Communications
55	GE Fanuc Embedded Systems – Embedded Systems
21	General Micro Systems – V394, CC61X
43	GSPx – GSPx Silicon Valley 2006
7	Harris RF Communications Division – Encryption Modules
19	Hunt Engineering – USB-connected Programmable FPGA Systems
6	Hybricon – Liquid-cooled 100 W/Slot 1 ATR
33	Hypertronics – Ruggedized Connector Solutions
9	ICS Sensor Processing – ICS-8550
5	Intel – Military Platforms
11	Jacyl Technology – Digital FPGA and Analog FPAA
45	MILCOM 2006 – MILCOM 2006
41	Military and Aerospace Electronics – Technical Conference and Exhibition 2006
53	MPL AG – Rugged Embedded Computers
3701	North Atlantic Industries – Military Power Supplies
47	Optima EPS – Custom Cabinets/Enclosures
2702	Phoenix International – Data Storage Modules
51	Real-Time Innovations – Data-critical Networked Applications
28	RTD Embedded Technologies – HighRel PC/PCI-104 Modules and Systems
2501	TEWS Technologies – Embedded I/O Solutions
2502	Thales – PENTXM2
2	Themis Computer – SPARC SBCs
3702	Tri-M Systems – Flash Solutions
35	Tri-M Systems – CT104 Enclosure
23	VMETRO – Real-Time Multiprocessors

ISSN: Print 1557-3222

Military Embedded Systems (USPS 019-288) is published four times a year (Spring, Summer, Fall, Winter) by OpenSystems Publishing LLC, 30233 Jefferson Avenue, St. Clair Shores, MI 48082.

Subscriptions are free to persons interested in the design or promotion of *Military Embedded Systems.* For others inside the US and Canada, subscriptions are $28/year. For 1st class delivery outside the US and Canada, subscriptions are $50/year (advance payment in US funds required).

Canada: Publication agreement number 40048627

Return address WDS, Station A PO Box 54, Windsor, ON N9A 615

POSTMASTER: Send address changes to *Military Embedded Systems* 16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268

Take a closer look at your encryption.

Are you getting the protection you need?

Members of the Sierra™ family of encryption modules are easily integrated into communications devices of all kinds. Sierra encrypts classified information, processes data at a higher speed, is extremely power efficient, and allows modules to be reprogrammed as missions change.

It's benefits like these that led the U.S. Department of Defense to select Sierra II for one of the most important communications programs in recent history. Sierra II has been chosen to encrypt 100 percent of the radios under Cluster 1 of the U.S. Joint Tactical Radio System program.

Learn more at **www.rfcomm.harris.com.**

www.harris.com

Broadcast
Microwave
RF
Government

HARRIS®

assuredcommunications™

RSC# 7 @www.mil-embedded.com/rsc

sec

Has *crummy old technology* survived?

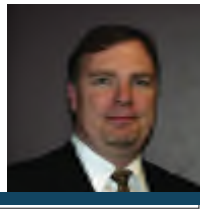

By Don Dingee

The seminal February 1994 *Mandate for Change* speech called for Commercial Off-the-Shelf (COTS) computing equipment, targeting:

- Reduced system costs
- Decreased system design time
- Improved system performance
- Reduced dependence on single-source suppliers
- Combined effects reducing total acquisition costs

Has COTS worked, or has *crummy old technology* survived? There are examples for both sides.

Direct hit

One case where COTS has gone right is the U.S. Navy's AN/BQQ-10 Acoustic Rapid COTS Insertion (A-RCI) program. A-RCI integrates COTS technology into submarine sonar systems in the Los Angeles SSN-688I class and has greatly increased anti-submarine warfare processing power at the Navy's disposal.

The A-RCI design teams at Lockheed Martin packed more than twice the sonar processing power onto a single sub than existed in the entire submarine fleet theretofore. Living the philosophy of *maximum density*, they traversed several form factors – 6U and 9U VMEbus boards and packaging from various suppliers, 2U rack-mount Pentium servers, even Apple Xserve boards repackaged to fit, looking to pack more processing into each design generation. The result: reduced system costs of about 20 percent over each generation of deployment. This program definitely hit the target intended for COTS.

Wide left and wide right

I'm aware of two recent examples of COTS selection that I believe went wrong, for different reasons:

- The SSGN program converts U.S. Navy ballistic missile submarines to cruise missile firing and special operations deployment capabilities. Its Vertical Launch Subsystem (VLS) called for a VMEbus processor board. During competitive bids in 2002, the prime contractor requested the bill of material for each board. Fifteen obsolete parts were found on the winning board at the time of selection, and the plan was for the U.S. Navy to procure a lifetime inventory for each of those obsolete parts. Buying and storing more and more *crummy old technology* as the list of obsolete components on a board grows over time doesn't make much sense.
- The Digital Airport Surveillance Radar (DASR), jointly managed by the U.S. Air Force and Federal Aviation Administration, detects aircraft position and weather conditions near civilian and military airfields. The ASR-11 – or its analog predecessor – is the large, orange, rotating *cattle gate* seen on the perimeter of most large airfields today. ASR-11 decided to borrow proven COTS technology from the Mode Select Beacon (or Mode-S) program, which completed a successful upgrade to a 68040 VMEbus board between 1998 and 2003. One small problem: By the time ASR-11 started deployment in 2003, the 68040 VMEbus board was ready to retire and finally did in 2004. The selection of *crummy old technology* for a compute engine can be a major setback, even if it feels safer at the time.

Old habits die hard

What conditions help COTS computing gear *work* for a program? Three seem to be important:

1. **Aligned thinking.** Directorate and program office teams and the contractor management teams must not only stand behind the concept of COTS but proactively break down traditional bureaucratic barriers and thinking. And vendors need to be involved in the thinking, too.
2. **Shorter life cycles.** Keeping design cycles as short as possible provides greater freedom in selecting computing technology, and planning for faster insertions of new technology keeps the total system life cycle short.
3. **Fixed scope.** Design teams should know more about their scope such as how many units they have to field, when decisions must be made, and when systems must be deployed. This helps contain risks, leverage knowledge from previous design cycles, and succeed while moving forward quickly. If the scope increases, the technology should be able to change.

Get off that *crummy old technology*

Creating any significant change is hard. Vendors, contractors, and the military need to work together to get COTS right more often. COTS board and system vendors are very hard-pressed to satisfy demands for the latest and greatest technology while simultaneously keeping older generations of product alive for what in today's technology cycles can be an epoch. Get off of the old stuff.

Is COTS working for your program or not, and why? We would love to hear from you, and you can reach me at ddingee@opensystems-publishing.com.

More Features. More Rugged. More reasons to choose ICS.

The one and only ICS-8550.

Designed for high-speed data acquisition applications such as Software Defined Radio, SIGINT, tactical communications and radar, the ICS-8550 XMC module — which is available for both benign and rugged environments — can simultaneously sample two RF/IF inputs at frequencies up to 210 MHz at a resolution of 12 bits.

With the industry-leading Xilinx Virtex-4 FPGA at its heart to deliver unprecedented power and user programmability — and enabling IF/UHF signals to be processed directly on the board itself, freeing the host board for other tasks — the ICS-8550 is truly an ADC module that sets new standards.

And with up to eight lanes of high-speed serial I/O, the ICS-8550 provides the throughput to match its performance and flexibility. Configure it with Radstone's remarkable V4DSPFPGA/PowerPC processor, and the partnership is unbeatable.

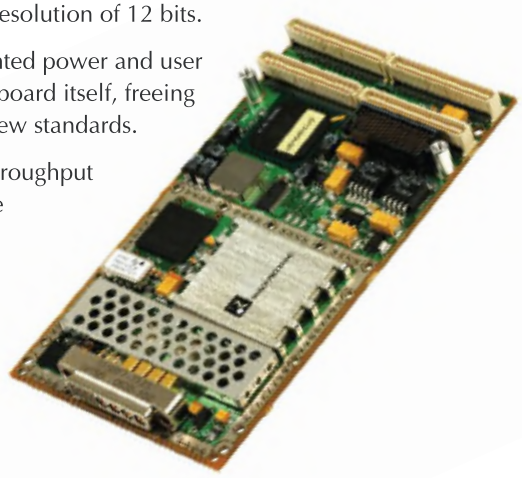

ICS. For when you need more.

ICS
SENSOR PROCESSING

CHECK THE FULL SPEC NOW AT
www.ics-ltd.com/8550

OR CALL TOLL FREE (US ONLY)
800-267-9794 or 613-749-9241

RSC# 9 @www.mil-embedded.com/rsc

Editor's Choice Products

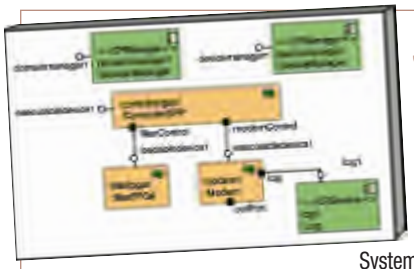

JTRS waveforms made ~~easy~~ easier

The DoD's Joint Tactical Radio System (JTRS) is a huge program that has boatloads of software at its core. Of primary importance are the SCA core framework and the portable waveforms that must work on myriad equipment while interoperating with older, hard-wired radios. Therefore, writing code for JTRS equipment can be a real nightmare. At least, that's the problem that Zeligsoft is trying to solve with its Component Enabler (CE) 2.4. This JTRS development tool helps software designers determine the usability of their software *components* in the field, and verifies compliance against the SCA.

CE 2.4 takes into account the target hardware *deployment platform* that the JTRS radio will ultimately run on. It also allows designers to write code and then iterate that code to balance the metrics of JTRS compliance, hardware choices, and waveform interoperability. Because of the iteration capability, the product speeds development, aids in the test phase, and ultimately saves costs.

Also available in CE 2.4 are modeling, validation, and runtime analysis capabilities. Multiple application views are intended for teams working on different portions of the overall code set, while keeping track of the overall combined set of software modules. A validation feature links to SCA validation and the latest version of the SCA specs, flagging rule violations. Writing and validating code for a JTRS radio will never be easy, but Zeligsoft's Component Enabler makes the task easier.

Zeligsoft
www.zeligsoft.com
RSC# 30700

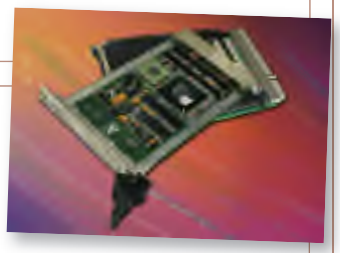

Rugged 3U CompactPCI SBC

Until the new 3U VITA 46 slim VME form factor becomes a reality later this year, 3U CompactPCI is still the format of choice for retrofitting space-constrained defense systems. With ample I/O capability and a nice size that fits well into ATR boxes and some SEM-E envelopes, 3U CompactPCI is an ideal choice. Aitech Defense Systems agrees, to which their C900 MPC7447A/7448-based PowerPC SBC can attest. Screaming along at 1.167 GHz with AltiVec signal processing support, the board includes up to 1 GB of DDR SDRAM complete with ECC. (Aitech is big into space-based apps, so you'd expect they'd include ECC or other provisions to deal with radiation effects.)

The board also includes 64 MB of user flash and a whopping 1 GB of NAND flash for program stores. There's also a thoughtful 128 kB NVRAM for application variables and 32 MB of boot flash. A Marvell MV64460 Discovery III chipset offers a PCI-to-PCI-X bridge. There's also a PMC mezzanine provision, as well as a bunch of I/O options. Two GbE ports, two USB 2.0 ports, and up to eight discrete I/O channels (or four RS-422 serial channels) round out the "goes-innas and goes-outtas." As you'd expect from Aitech, air- and conduction-cooled versions are available. Betcha they'd even make you one that's rad-hard.

Aitech Defense Systems
www.rugged.com
RSC# 23045

Ultra-low power FPGA

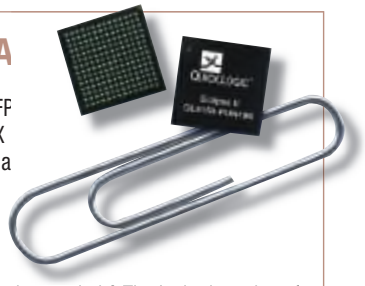

The gazillion gate highly integrated FP from companies ranging from A to X performance to spare but come with a heavy power penalty. Although VME boards have ample power budgets of up to 100 W these days, that power has to be dumped into the system somehow. If you can save power, why not do it? That's the intention of QuickLogic's QL8150 Eclipse II FPGAs. Designed for light-density logic applications such as handheld devices, they're also ideal for fixed-function interface controllers on VME basecards, PMC mezzanines, active backplanes, and chassis front panels.

With 188,946 maximum gates and 640 logic cells in a 32 x 32 logic array, this small 8 x 8 mm footprint BGA package uses only 196 fine-pitch balls. Power consumption varies, but think *battery powered* and you're in the right realm. The devices are designed to operate over a -40 °C to +100 °C temperature range, so deployment in conduction-cooled VME chassis should be no problem.

QuickLogic
www.quicklogic.com
RSC# 30701

PDA's protective skin saves soldiers' lives

Anyone with an iPod knows that a protective case called a *skin* protects the unit from scuffs and scratches and cushions it for the inevitable drop onto concrete. OtterBox cases protect iPods, tablet PCs, portable GPS units, and myriad other handheld electronic devices. But recently, Otter Products' OtterBox 1900 PDA case has been used to save lives when coupled with an HP iPAQ hx4700.

The company worked with the U.S. Army's Medical Research and Materiel Command's Medical Communication for Combat Casualty Care (MC4) and the Telemedicine and Advanced Technology Research Center (TATRC) to cocoon the iPAQ for battlefield medical needs. TATRC's Battlefield Medical Information System Tactical – Joint (BMIST-J) allows Army tactical medical forces to quickly and securely document patient information at the point of injury. The system replaces the World War II vintage paper DD Form 1380. Without the OtterBox 1900 PDA case, the iPAQ would never survive the rigors of life on the front lines, including moisture, dust, shock, and vibration.

MC4 has deployed more than 12,000 OtterBox and BMIST-J systems, with a savings of more than $1,000 per unit compared to a purpose-built rugged PDA costing upwards of $1,500 each. Otter Products also manufactures numerous other protective skins for civilian and custom handhelds.

Otter Products
www.otterbox.com
RSC# 31423

Editor's Choice Products are drawn from OSP's product database and press releases. Vendors may add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/ and submit press releases at www.opensystems-publishing.com/news/submit. OSP reserves the right to publish products based on editors' discretion alone, and does not guarantee publication of any product entries.

JACYL TECHNOLOGY

Advancing Today's Technology Into Tomorrow

We specialize in completely reconfigurable and reprogrammable Digital FPGA and Analog FPAA PC/104 Boards.

We're Your Project's Solution.

www.jacyltechnology.com

RSC# 11 @www.mil-embedded.com/rsc

Board vendor FPGA toolkits make or break your project

By Mark Littlefield

FPGAs dramatically accelerate DSP designs while bringing reconfigurability to the battlefield. But to wring out performance, ease-of-use, and overall program benefits, a design toolkit is needed. Better choose wisely.

Increasingly, the military community is recognizing that there is a class of signal processing that is now best accomplished via a reconfigurable computing implementation instead of the traditional method of software running on microprocessors. Modern Field Programmable Gate Arrays (FPGAs) have reached a level of density, speed, and cost such that system designers can now often achieve a tenfold reduction in size and power when compared with the traditional microprocessor-based approach. The power of FPGAs stems from the opportunity to *parallelize* operations that a microprocessor must do sequentially.

For many signal processing designers, the question is not whether the use of FPGAs will result in higher performance – performance/power/cost – but whether the R&D investment will pay off. The simple fact is that developing signal processing functions in FPGAs has high technical risk, which can result in cost and schedule overruns.

When selecting a COTS board level product, a system integrator should be aware that the board vendor FPGA supporting toolkit will play a large role in reducing (or not) this technical risk. Herein, we will acquaint the reader with the nature of such tools and their effect on the cost of developing an FPGA compute solution.

What is an FPGA toolkit?

In our context, an FPGA toolkit is the collection of supporting IP (VHDL designs) and other *software* components that are offered for use by the vendor of a board level product. We are not talking about the synthesis and simulation tools that are the domain of the FPGA silicon vendors and other tool specialists.

The components of an FPGA board toolkit fall into four categories:

- IP designs to control hardware features, for example, an SDRAM controller
- IP infrastructure to connect IP blocks together
- Software for system services, data movement, and general system integration
- Simulation and test

To appreciate the key attributes of these toolkit components, consider the Curtiss-Wright CHAMP-FX, an FPGA processing board designed for military signal processing applications. The CHAMP-FX, illustrated in Figure 1, integrates the Xilinx Virtex-II Pro FPGAs (VP70/100) with local memories, PCI interfaces, and high-speed serial interfaces.

The combination of DDR SDRAM for bulk storage and DDR SRAM for fast, nonsequential storage of algorithm data allows flexibility in mapping algorithms to this architecture. In a typical application, data may first flow into SDRAM with intermediate storage in SRAM for algorithm processing with results output via the PCIbus to StarFabric or an alternate interface provided by a PMC module. In such dataflow architectures, the performance and ease of use of the memory controllers can dictate the overall performance of the application.

Robust IP blocks are critical

Implementing a memory controller provides an example of the criticality of the design kit to the success of a project. While it is possible to obtain *free* designs for an SDRAM controller from FPGA vendors, be advised that these are often limited-function reference designs. Experience has shown that it takes many man-months of experienced FPGA designer time to implement a high-performance, reliable SDRAM controller for FPGA-based hardware. The challenges that emerged – not unexpected – emanate from classic FPGA design issues. Examples of some of these design issues follow:

During the read cycle of a DDR SDRAM, the FPGA sends out a clock to the SDRAM and waits two cycles for a four-word burst to return. The challenge is in clocking the data back from the SDRAM. There is skew between the

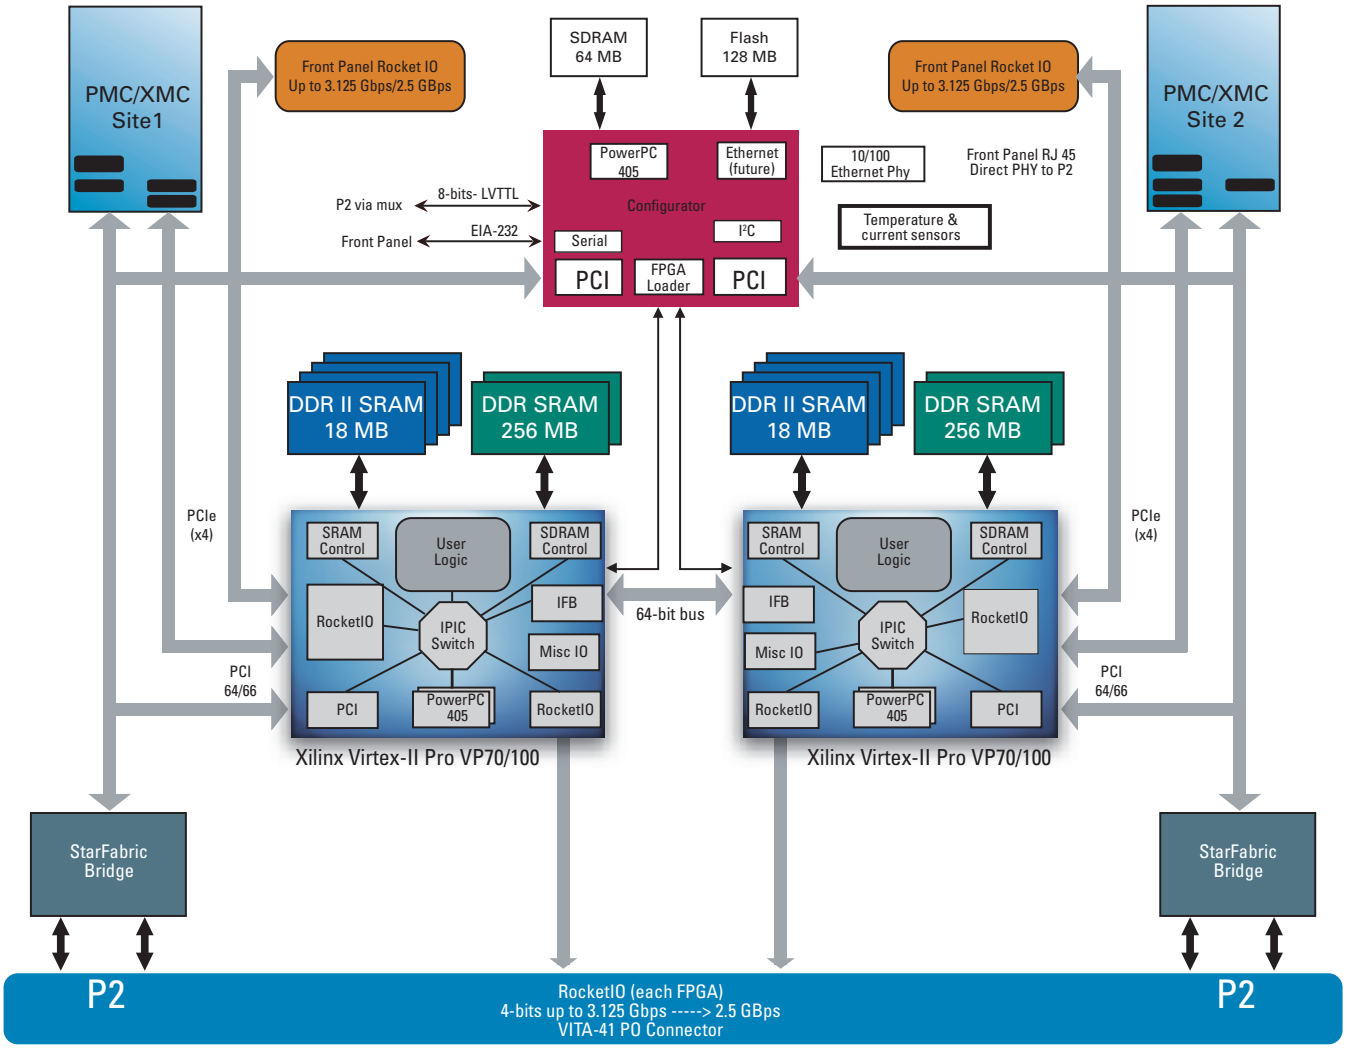

Figure 1

transmitted clock and the data (due to 180 ps/inch PWB trace delay), so this is not an optimal approach to use. A better approach is to use a phase-shifted clock within the SDRAM controller IP that can account for the PWB trace length and transmission characteristics. This allows the FPGA to clock in read data with the appropriate timing margins. Because the FPGA SDRAM controller is working with externally connected devices, the design must take into account the specific FPGA pins used for the interface, the placement of the controller within the FPGA, and the PWB track lengths. At the 132 MHz operating frequency of the SDRAM interface, there is the narrow 2-3 ns window during which the data is valid. Debugging the interface between FPGA and memories is further complicated by modern BGA packages that preclude the use of oscilloscopes and logic analyzers. If it is not working, it is very difficult to diagnose the cause of the problem.

FGPA designs must adhere to simultaneous switching output rules. When too many outputs from a particular bank switch at the same time, noise is introduced into the system, which can cause data errors to occur. This is another subtlety of FPGA IP design that is difficult to debug because the errors are somewhat random and analog in nature.

There is significant effort required to get a memory controller to function correctly. Getting controllers to work at high speeds (200 MHz SDRAM, 132 MHz SDRAM) also requires particular attention be paid to the placement of the IP blocks within the FPGA. A robust design optimized for a particular hardware implementation will include carefully selected and tested constraint information that ensures that the complete design, when integrated with the user's logic, will continue to operate within timing requirements. Figure 2, a sample floorplan, shows an example of the placement of IP blocks that have been constrained to particular placement within the FPGA. This is taken from the Xilinx ISE floor planning tool. The colors represent different IP blocks. The constraint files tell the placement tool to place all the logic for a given block within a certain area. This helps ensure repeatable timing within the block and makes it less sensitive to the logic the user is adding.

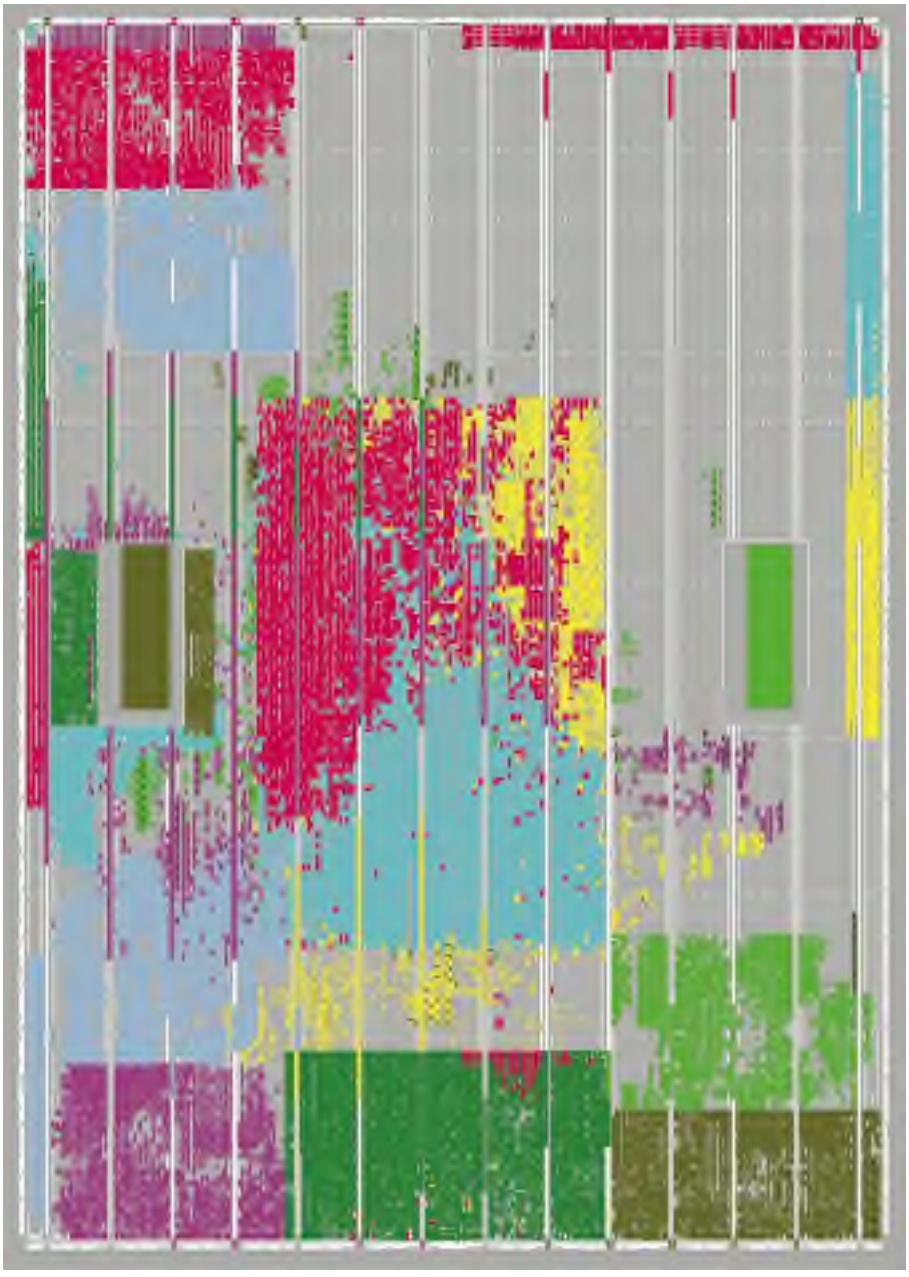

Figure 2

Last but not least is the added challenge presented by the wide temperature ranges encountered in military applications. The typical environmental requirement of -40 ºC to +85 ºC translates to an even wider range of -40 ºC to +110 ºC for the silicon. The design tools provide estimates of the timing effects over temperature, but only costly environment testing can verify the real effects, which will rarely be exactly as predicted. The design kit IP, if prequalified to this temperature, removes this as a risk item.

Interconnecting IP blocks

It is desirable for IP controllers and functional blocks provided as part of a toolkit to adhere to a consistent interface standard. Common interfaces ease IP block integration. We have adopted two interface conventions defined by Xilinx, IP Interconnect (IPIC) and Local Link. Memory mapped interfaces adhere to the IPIC standard, allowing an IP block to read or write to a memory address or register. SDRAM, SRAM, and the PCIbus are provided with IPIC interfaces.

Local Link is a packet-oriented interface for streaming data applications. In a DSP application, it is common for sensor data – digitized RF, electro-optical, and sonar – to arrive in nonaddressed packet form. The RocketIO controllers (3.125 Gbps) are supported with Local Link interfaces and DMA controllers to direct streaming data into memory-addressed locations.

To simplify the interconnection of user IP with the toolkit IP modules, the company has developed an IPIC switch, which allows modules with IPIC or Local Link interfaces to connect together seamlessly (see Figure 3). Multiple IPIC switches can be instantiated in the design, allowing users flexibility in choosing the best solution.

Simulation: Don't take it for granted

FPGA signal processor designs are complex and, unlike software, are not easily instrumented for testing purposes. (There is no `printf`()!) The rule of thumb for a complex project is that simulation is 50 percent of the effort. It is fairly obvious that the FPGA vendor simulation tools can simulate the logic within the FPGA, but unless models are provided for external devices connected to the FPGA, the simulation cannot include interaction with these devices. A good FPGA toolkit will offer a test-bench environment that includes models of all the external interfaces, the capability to initialize memory interfaces and pass data from PCI and RocketIO into the FPGA and check memory, and the capability to test PCI or RocketIO data against expected data previously stored in files. The toolkit should provide a scripting language to facilitate the simulation of the design with test data and the capture of output data into files for comparison against expected results.

Unless a good simulation environment is part of the toolkit, the customer will be faced with a great deal of unplanned effort to establish a test environment. Figure 4 depicts the simulation models that are supplied with the CHAMP-FX design kit, and the logical points where data can be input or output from the simulation.

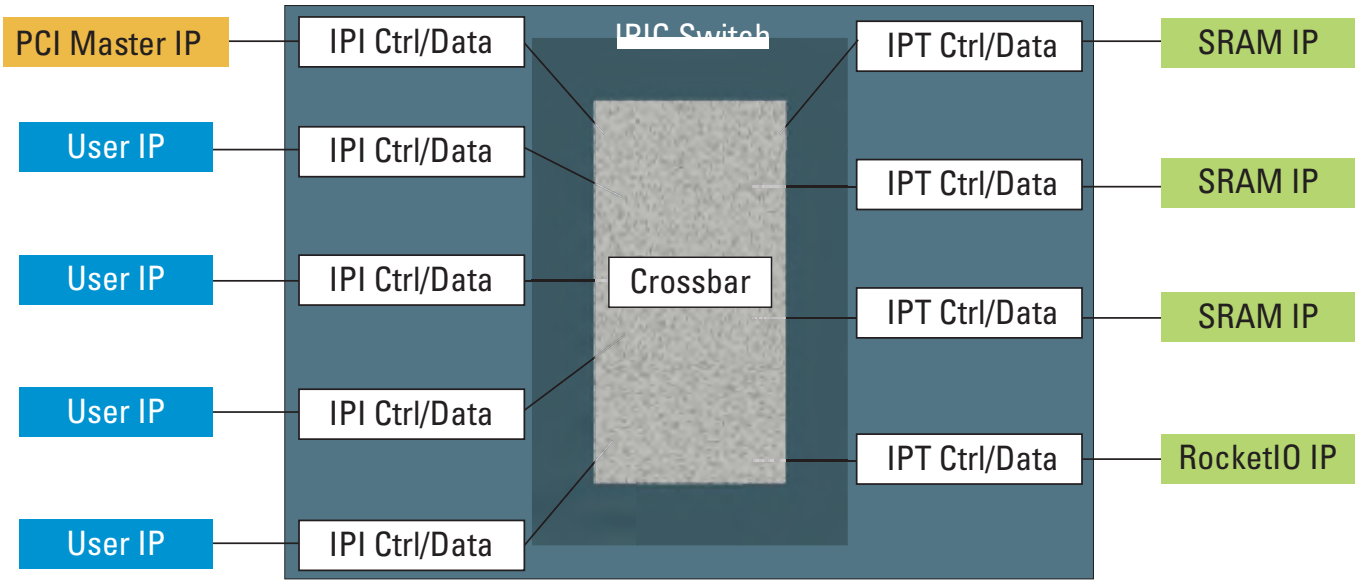

Figure 3

The final stage of testing requires loading and running the logic on the target hardware and integrating the design with the rest of the system. Despite efforts to accurately simulate a design, there remains a possibility that it will not work as intended. The Xilinx Chipscope logic analyzer is a great tool to probe the internals of a design in real time and to deduce where a problem resides. The use of Chipscope, however, requires some support within the FPGA toolkit to simplify including it in the design, so a prospective customer would be advised to check into its support. Once the FPGA bitstream design is stable, the integration effort is eased by the availability of a rich set of software APIs for device control, data movement, and interprocessor synchronization.

The quality and coverage of the FPGA toolkit that accompanies a COTS FPGA board is critical to the schedule and cost of an FPGA-based project. Since FPGA toolkits are like software in some respects, the quality, performance, and functionality can take some extra effort to ascertain. This extra time is well spent, since any one of the issues discussed herein could result in many man-months to resolve.

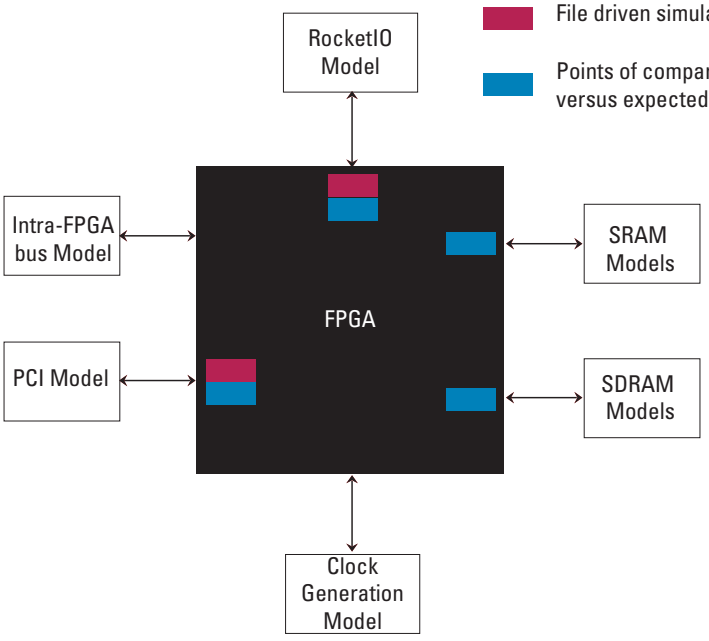

Figure 4

Mark Littlefield *is the product marketing manager for Curtiss-Wright Controls Embedded Computing's FPGA computing product line. He has more than 15 years of experience in the embedded computing industry, first as an engineer developing robot vision systems for NASA, and later as a field applications engineer, technical program manager, and product manager. Mark has a BS and an MS degree in Control Systems Engineering from the University of West Florida.*

To learn more, contact Mark at:

Curtiss-Wright Controls Embedded Computing
741-G Miller Dr. SE
Leesburg, VA 20175
Tel: 703-779-7800
Fax: 703-779-7805
E-mail:
Mark.Littlefield@curtisswright.com
Website: www.cwcembedded.com

Design strategies for an FPGA-based 256-channel digital down converter

By Rodger H. Hosking

FPGAs can replace traditional ASIC-based digital down converters in high channel count Software-Defined Radios. Their inherent parallelism allows multiple digital receiver channels per chip, and available COTS IP cores can be used to realize up to 256 independently controlled channels in a Xilinx Virtex family FPGA.

As Software-Defined Radio technology further penetrates large communication systems for battlefield military radio networks, commercial wireless systems, manned and unmanned aerial vehicles, and monitoring facilities for SIGINT and COMINT, the need to accommodate a large number of agile frequency channels for radio receivers is quite apparent. In each of these applications, the same critical metrics apply: size, weight, power, and cost for each receiver channel.

Traditional Digital Down Converter (DDC) ASIC devices feature only one to four channels per chip, and straightforward implementations of DDCs in FPGAs consume a significant percentage of available resources. A new approach to DDC design takes advantage of the parallelism of FPGAs to create a highly efficient architecture for multichannel receivers.

Basics of digital down converters

DDCs, often called *digital receivers,* perform the two essential software radio functions: frequency translation and channel filtering. In a basic DDC shown in Figure 1, a mixer and local oscillator perform the frequency translation.

The local oscillator consists of a digital phase accumulator that advances each clock by a programmable increment equal to the tuning frequency. The phase accumulator is a register whose full-scale value represents 360 degrees of a sinusoid. A sine/cosine lookup table converts the phase angle of the accumulator to the digital voltage value of the sinusoid. The higher the increment, the faster the phase accumulator steps through the sine table. It naturally overflows at the top, preserving any residue left in the register as a phase offset for the first sample of the next cycle. As a result, the output sinusoid is directly proportional to the phase increment or frequency setting. This block is a classic Numerically Controlled Oscillator (NCO), also often called a Direct Digital Synthesizer (DDS).

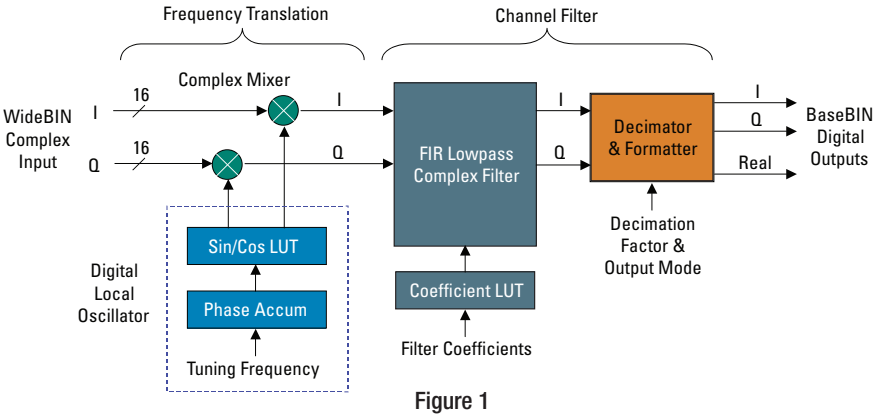

Figure 1

The mixer consists of two digital multipliers that accept complex sine/cosine outputs from the local oscillator and digital samples of the receiver input signal produced by an A/D converter. Multiplication in the time domain produces a sum and difference signal in the frequency domain. If the local oscillator is set to the frequency of the input signal of interest, the difference term will be that input signal translated down to 0 Hz. Since the mixer is complex, the upper and lower sidebands of the input signal will be translated to negative and positive frequencies centered at 0 Hz.

The filter is a complex low-pass digital filter with two parallel I and Q arms whose coefficients are programmed for a pass band equal to the channel bandwidth. Because the output of the filter is bandlimited, the output decimation stage can drop the sampling rate accordingly.

DDCs are grouped into two main categories. Wideband DDCs have output channel bandwidths typically above 1 MHz and are appropriate for wideband Code Division Multiple Access (CMDA) and radar applications. Narrowband DDCs with bandwidths below 1 MHz are widely used for Frequency Division Multiplexed (FDM) systems including voice and music channels for telecom and commercial broadcast systems. While the mixer and local oscillator sections are quite similar for all DDCs, the best filter design depends on the filter bandwidth. For wideband channels, a conventional FIR filter is best (as shown in Figure 1). For narrowband channels, a multistage Cascaded Integrator-Comb (CIC) filter followed by an FIR to correct frequency droop is more efficient.

For narrowband applications, both ASIC and FPGA Intellectual Property (IP) cores are available using CIC filter designs.

Annapolis Micro Systems

The FPGA Systems Performance Leader!

SIGINT ELINT Digital Receivers FLIR
FOPEN Radar Systems Software Defined Radio

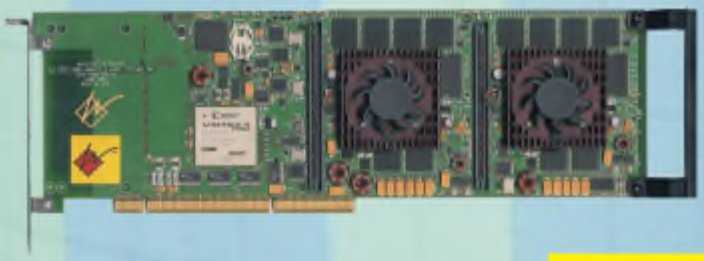

CoreFire™

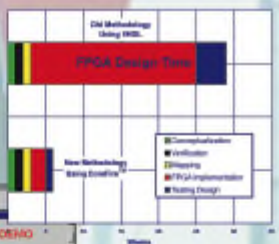

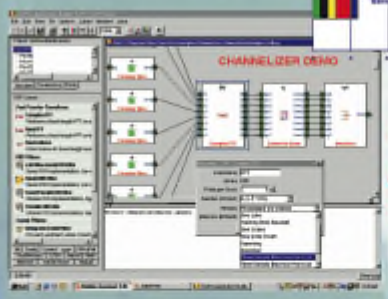

FPGA Application Builder

Annapolis Micro Systems

High Performance Signal Processing in Scalable FPGA Computing Fabric

Real Time, Real World.....REAL FAST!

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

Commercial ASICs feature as many as four channels per chip, like the popular Texas Instruments/Graychip GC4016.

IP core DDCs, like the LogiCore DDC from Xilinx for its Virtex-II Pro, can be scaled for various levels of Spurious-Free Dynamic Range (SFDR) performance to use more or less of the available resources. For example, a complex DDC with 84 dB SFDR consumes approximately 1,700 slices. In a mid-sized FPGA device with 24,000 available slices, only about 14 DDC channels can be accommodated. For applications requiring several dozen or even hundreds of channels, this approach can become impractical.

Channelizers

Because of the extremely fine resolution of its NCO tuning frequency, a true DDC can translate any input frequency component down to 0 Hz, often with 32-bit accuracy. This ability makes DDCs ideal for applications that require precise changes in tuning such as in continuous Doppler correction for satellite tracking systems.

However, in other applications, a *channelizer* approach may be sufficient. This is a bank of equally spaced, fixed frequency band pass filters whose outputs are translated to baseband (0 Hz). One crude example of a channelizer familiar to everyone is a simple FFT. It converts a block of N time samples equally spaced in time into block of N frequency samples equally spaced in frequency. For a continuous stream of input time sample blocks, samples at a given point in successive output blocks represent a translated, band pass frequency signal or *bin*.

By selecting the output of a particular bin, a channelizer can serve as a primitive DDC, but with extremely coarse tuning resolution determined by the number of points in the FFT, as shown in Figure 2.

Another serious limitation of the FFT as a DDC is the frequency response (pass band flatness) of the bin, and rejection of energy from adjacent bins (stop band rejection). Other channelizer designs use various digital filtering techniques to split the bands with better flatness and adjacent channel rejection, but they usually require significantly more hardware than an FFT for a comparable number of bins. Regardless of its design, the tuning resolution of any channelizer is simply equal to the number of bins or channel filters. As a result, channelizers may be useful for spectrum analyzers, scanners, and energy survey equipment but they are rarely used as substitutes for DDCs in software radio communication systems.

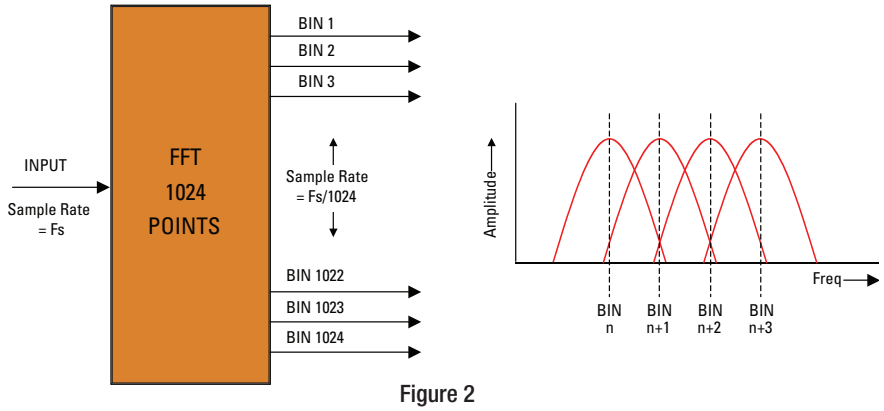

Figure 2

Rethinking the multichannel DDC

The software radio market generates a growing number of requests for DDC solutions with densities higher than the 16 or 32 channels provided per board using ASICs or standard FPGA designs. Therefore, we embarked upon a mission to develop a signal processing architecture for a narrowband DDC with 64 channels or more, with full tuning resolution, but with much more efficient use of FPGA resources than deploying a farm of conventional DDC cores.

Each conventional DDC requires its own local oscillator (phase accumulator and sine table), mixer (two multipliers), and FIR filter (multipliers and accumulators). All of this hardware must operate at the full input sample clock rate, and clock rates for A/Ds commonly used in software radios range between 100 and 200 MHz. Since this is the same clock range rating for commercial DDC IP cores, all of the hardware resources used for each channel must be dedicated to that channel.

However, imagine that the input data sample rate is reduced by a factor N. By operating the DDC hardware resources required for one channel at the full clock rate, those same resources can then be multiplexed (time shared) across N channels. Of course, provisions must be made for buffering the data for all channels while multiplexing. This is usually done in RAM or in delay memory, a common feature in FPGAs.

One way to achieve this input rate reduction is to split the input signal into a bank of N adjacent frequency bands using a channelizer. Then, the output sample rate for each band can be reduced by a factor of N. The output from the band containing the signal of interest can be selected as the input to any given DDC to fine tune within that band.

The tradeoff question becomes: Are the resources freed up by multiplexing the DDCs more than the resources required for the channelizer? The answer lies in how efficient the channelizer can be.

Realizing the design

Figure 3 shows an FPGA-based 256-channel DDC IP core that combines a channelizer stage with a multiplexed DDC stage.

The crucial part of this design is the channelizer stage. It accepts a single wideband input stream and delivers a channel bank of 1,024 output bands equally

spaced in frequency, but with significant overlap between adjacent bands.

The output sample rate of each band equals the input sample rate (Fs) divided by 256, rather than 1,024, as would be expected with a simple FFT. In fact, inside the channelizer are four high-speed 1,024-point FFTs running in parallel using a proprietary windowing and overlap processing technique. The outputs of the four separate FFTs deliver samples at a rate of Fs/1024. These outputs are combined to form a single output at a sample rate of Fs/256, supporting the wider bandwidth that will sufficiently overlap adjacent bands.

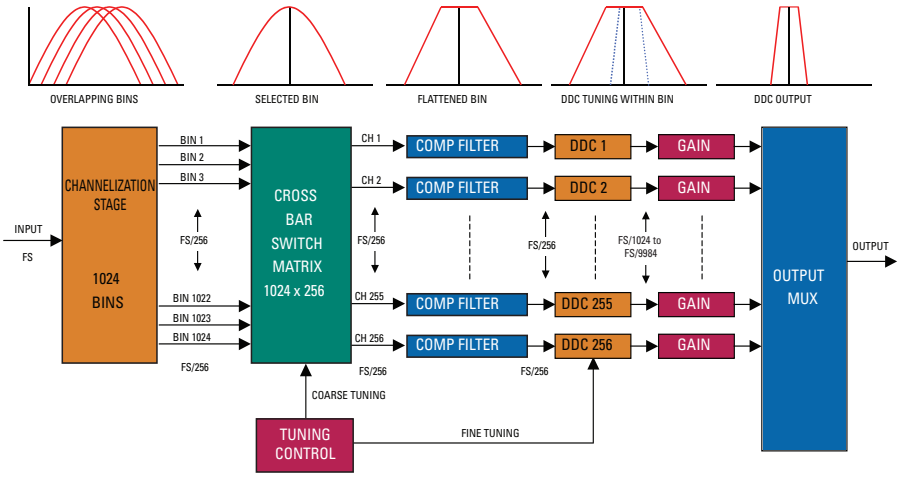

Figure 3

The next stage is a crossbar switch matrix that accepts 1,024 inputs from the channelizer and delivers 256 outputs, one to each DDC channel. The switch is nonblocking so that any of the 256 outputs can be independently sourced from any of the 1,024 channelizer bands with no restrictions.

Each of the 256 channels is tuned by a separate 32-bit frequency word, with the most significant bits sent to the switch matrix for coarse tuning. This selects the correct channel band for each channel. The least significant bits of the frequency word are used by the DDC stage for fine tuning within the selected band.

Because the channelizer outputs exhibit frequency droop at the band edges, a fixed compensation FIR filter flattens the pass band to within 1 dB across a span equal to twice the band-to-band spacing.

A bank of 256 independently tuned DDC sections, each with its own local oscillator, mixer, and FIR filter, processes the 256 compensated switch matrix outputs. Because the channelizer has dramatically reduced the input sampling rate to each DDC section by a factor of 256, the DDCs are implemented using highly multiplexed hardware resources and block RAM to preserve the data for each channel. A gain stage, output multiplexer, and data formatter complete the design.

Performance and tradeoffs

The maximum output bandwidth of this design equals the channelizer band-to-band spacing of Fs/1024. For an input sample rate of 100 MHz, this spacing is about 100 kHz. And because of the broadened response, each channelizer output has a clean pass band equal to twice the band spacing, or about 200 MHz.

This allows the DDC to perform fine tuning by sliding its local oscillator frequency ± 100 kHz across the selected 200 kHz channelizer band to precisely center the DDC output. Choosing a wider DDC output bandwidth would restrict the DDC tuning range, since the edge of that wider bandwidth would cross the edge of the flat, spurious-free region of the channelizer output.

Samples of the translated signal from the mixer arrive at the decimating FIR at the channelizer output sample rate of Fs/256. Since the maximum available DDC output bandwidth is Fs/1024, the lowest decimation factor allowed in the FIR is 4.

For narrower output bandwidths, the maximum decimation factor is determined by the complexity (number of taps) of the FIR filter, which must perform at least as well as the channelizer filter in order to maintain that dynamic range of 75 dB. Choosing a reasonable number of multiplier/accumulator stages yields an FIR filter suitable for decimation factors from 4 to 39 in steps of 1. The number of filter taps is equal to 26 times the decimation factor of the filter.

Since the channelizer decimation (256) and FIR filter decimation (4 to 39) multiply, the overall range of decimation range for the entire core is 1,024 to 9,984 in steps of 256. Each of these 36 available decimation factors requires its own set of filter coefficients, which are stored in a table within the FPGA. For an input clock of Fs = 100 MHz, the range of output bandwidths using the default 80 percent filter characteristic is approximately 8 kHz to 80 kHz. For any decimation setting, the overall DDC channel characteristics, including the channelizer response, are shown in Figure 4.

Because of the multiplexed DDC hardware, all 256 channels must have the same decimation factor setting. For high-channel count systems, this limitation is usually not an issue since it is quite common for all such channels to have the same bandwidth.

Overall performance of the complete 256-channel FPGA-based DDC IP core includes a spurious-free dynamic range of 75 dB, a pass band ripple of 0.4 dB, a pass band edge droop of 1.0 dB, and frequency tuning resolution of $Fs/2^{32}$. The maximum clock frequency depends on implementation details, but can be as high as 185 MHz in a Virtex-4 FPGA with speed grade 12.

The core consumes approximately 18,000 logic slices of a Virtex-4 device, compared to 1,700 slices for a single channel DDC LogiCore reference design. Although there are some limitations in decimation factors and dynamic range, this new core represents an improvement in the channel-per-slice ratio by a factor of more than 20.

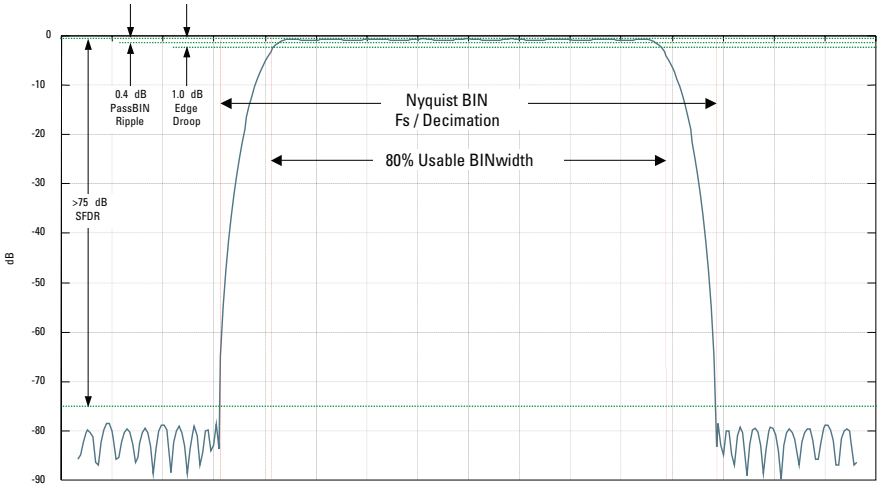

Figure 4

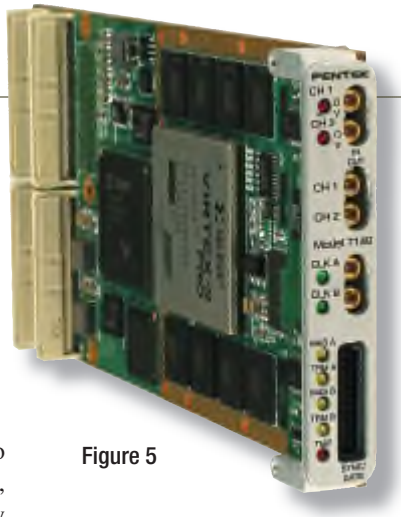

Figure 5

This 256-channel DDC core is available as a member of the GateFlow IP Core Library suitable for use with any Virtex-II, Virtex-II Pro, or Virtex-4 product.

For customers preferring to avoid FPGA development, it can be ordered as a factory installed option to the Pentek Model 7140 Dual Transceiver PMC module (shown in Figure 5) where it occupies approximately 76 percent of the Virtex-II Pro XC2VP50.

Creativity beats crunch

In order to keep pace with a steady flow of new FPGA device offerings, designers must continually evaluate, and often reinvent, real-time embedded computing strategies for critical military and commercial applications. Armed with a detailed understanding of new device resources, creative engineers can often approach a tough problem from a radically new angle to gain a major advantage. While many new FPGA design tools offer impressive features and improved efficiencies, these truly significant leaps in FPGA performance usually come from inspiration, not from automation.

Rodger H. Hosking *is Vice President and Cofounder of Pentek, Inc., where he is responsible for new product definition, technology development, and strategic alliances. With more than 30 years of experience in the electronics industry, he has authored hundreds of articles about digital signal processing. More than 10 years ago, Rodger introduced digital receiver technology to the government electronics industry through numerous presentations and his widely read publication* The Digital Receiver Handbook. *Prior to his current position, he served as Engineering Manager at Wavetek/Rockland and holds patents in frequency synthesis and spectrum analysis techniques. He received a BS degree in Physics from Allegheny College in Pennsylvania and BSEE and MSEE degrees from Columbia University in New York.*

For more information, contact Rodger at:

Pentek, Inc.
One Park Way
Upper Saddle River, NJ 07458
Tel: 201-818-5900, Ext. 228
E-mail: rodger@pentek.com
Website: www.pentek.com

ENGINEERED LIKE NO OTHER™

FROM CONCEPTS TO DEPLOYMENT

V394 "MAVERICK"

Ultra High Performance, Dual PMC, PowerPC VME SBC

- Dual 7447A/7448 PowerPC up to 1.2GHz
- Up to 1MB of L2 Cache
- Full Symmetric Multi Processing (SMP) support via Discovery®-III
- 128-bit AltiVec™ technology support
- Dual 64-bit 66/100MHz PCI-X local buses
- Up to 1GB of 333 MHz DDR SDRAM via SODIMM modules
- Two 64-bit 66/100MHz PMC expansion sites, one with rear I/O
- Ultra SCSI 160 with differential output
- Three Gigabit Ethernet ports
- Dual USB, IDE DMA-100, Serial, PS2 Mouse and Keyboard
- 32MB of application Flash
- One Multi Protocol Sync/Async Serial port
- 2MB of Bootable Flash
- Onboard 2.5 inch IDE Hard Disk Drive (Optional)
- Parallel port and 21 discreet digital I/O lines
- Optional triple 66MHz, 64-bit PMC expansion module
- Supports VME-64 via Tundra Universe-II®
- Support for 5-row or 3-row VME card cages
- Front panel LEDs for POST code and Status display
- Support for VxWorks®, Linux® and QNX® Operating Systems

CC61X "ROCK"

High Performance, 3U Conduction Cooled, Rugged CompactPCI Pentium® M SBC

- Up to 1.4GHz Pentium® M processor with up to 2MB of L2 Cache
- Ultra-low power requirements as low as 12W Max
- Up to 2GB of 266MHz DDR SDRAM
- Up to 16GB of Bootable Flash memory
- Baseboard Management Controller (BMC) to meet PICMG2.9
- Full Health reporting and monitoring with extensive BIT/EBIT
- Dual Gigabit Ethernet ports with TCP/IP Offloading Engine
- Dual Video to support DVI and RGB
- Full Power Management control:
  - ACPI 2.0 compliant (Suspend to memory/disk etc.)
  - Geyserville® III support (Speed and Voltage stepping)
  - Support for battery operation and standby power
- Four USB-2.0 ports, two Serial ATA and two Serial ports
- 512KB of BIOS/user Flash and 2Kb of Serial EEPROM
- RTC with onboard and external battery support
- High speed I/O via SAM™ bus expansion connector (PCI-X or AGP)
- Available in standard 0-55°C or extended temperature -40° to +85°C
- Onboard heaters allow extended low temperature storage/operation
- Support for Windows® XP/2000, VxWorks® and Linux®

LEADING THE EMBEDDED MARKET SINCE 1979.

PROUDLY MADE IN THE USA · USA GSA · GSA APPROVED

GMS

COMPUTING ENGINES

PERFORMANCE, RELIABILITY, LONGEVITY

UL REGISTERED FIRM

GENERAL MICRO SYSTEMS, INC.
REGISTERED TO ISO9001
FILE NUMBER A8073

GENERAL MICRO SYSTEMS, INC.

TEL(800) 307-4863 • gms4sbc.com

RSC# 21 @www.mil-embedded.com/rsc

SDR and JTRS: Lessons learned

An interview with Col. Steven MacLaird, USAF (ret.) and former Program Executive Director of the Joint Tactical Radio System JPO

EDITOR'S FOREWORD

Col. Steven MacLaird managed the Joint Tactical Radio System (JTRS) Joint Program Office (JPO) from 2000 until his retirement from the U.S. Air Force in 2005. JTRS is one of the DoD's largest programs, which I include with the "big three" of the Joint Strike Fighter F-35 and Future Combat Systems. Last year, I estimated that over the program's life, somewhere between $12-$15 billion would be spent on JTRS radios and infrastructure (download available at http://meecc.com/presentations/CIUFO.pdf).

Steven successfully navigated a challenging program, rife with technical, financial, and programmatic obstacles. Today, the JTRS clusters have been renamed and shuffled around, and the program is realigning its timetable for initial deployment over the next 12 months. He has been selected to sit on the board of directors of PrismTech, a COTS supplier of the Software Communications Architecture (SCA) that's essential to JTRS. I had the privilege of speaking with him last April. Edited excerpts from that conversation follow. – Chris Ciufo

MIL EMBEDDED: ***Please provide a brief overview of the current JTRS radios and explain how they relate to the previous clusters.***

MACLAIRD: In March of last year a hand off was initiated between the old [Joint Program Office] organization and the new, and I went in and briefed Dennis Bauman on the program. And we provided to him a way ahead on the organizational structure. He's adopted a large set of that, which you have undoubtedly seen in the press.

The ground domain consists of Ground Mobile Vehicles (GMV) radios and the Handheld/Manpack/Small form fit (HMS) radios, i.e., *Cluster 1 and Cluster 2.* Special radios used to be *Cluster 2,* also known as JTRS Enhanced MBTR (JEM). The airborne maritime fixed site domain is made up by the AMF program, which previously, probably about November 2004 was to be the airborne fixed site program, known as *Cluster 3* and *Cluster 4*, and the Air Force and Navy came together and consolidated them into the AMF. Finally, the waveform and crypto program was transitioned into the Network Enterprise Domain (NED).

MIL EMBEDDED: ***Why did you recommend changing the clusters around?***

MACLAIRD: There was a strong service equity issue within the program when I took it over in June of 2001, which led to some dysfunctional approaches to delivering capabability to the joint force. Also, when you looked at the history of the original intent of the program organization, the organization that I inherited did not reflect the intent of the program.

MIL EMBEDDED: ***Can you define the term*** **service equity?**

MACLAIRD: The way the money was laid out – actually, the way the program was laid out – was the funds for the particular programs resided in the service's top line [budget]. So if you talk about what used to be Cluster 1, that was predominately Army funded. The Army ran that program. And although the JPO director had oversight authority over the program, the actual funding went through the Army, then through the JPO and separately through the Army CECOM for Cluster 1. And the leadership of the organization reported through another organizational structure up at [U.S. Army] CECOM.

MIL EMBEDDED: ***Let's talk about some of the technologies. Can you comment on COTS technology and this program … where it's come from to where it is today?***

MACLAIRD: *Proprietary* is a pretty good summation of the program. What's really important to understand is that you took a program that started in 2000 that not only was trying to build hardware but software operating systems and waveforms, *and* build new standards *and* try to deliver new capability all at once.

What we can do today on FPGAs, GPPs, and DSPs and what people thought we could do back in 2001 when I took over the program has greatly expanded. Spectrum Signal [Processing], Harris Corporation, and General Dynamics are out there running

VMETRO 20 years! 1986 - 2006

Coordinated Performance

with Phoenix Systems

Real-Time Multi-Processors with High-Speed Serial Communications

Powerful

- Customer programmable Xilinx FPGA compute engines
- Multiple PowerPC processors
- Multi-channel Gbit/sec backplane communications
- Fiber Optic Front Panel I/O, including Serial FPDP
- High-speed Analog and Digital I/O

Flexible

- Combine FPGA and PowerPC Processors to match the application requirements
- Boards or complete sub-systems
- Open standards including VITA 41/VXS and VITA 42/XMC
- Air-Cooled to rugged conduction-cooled versions

Innovative

- Zero Latency Switch technology
- TransComm™ inter-processor communications suite for simplified application development

Processing and FPGA - Input/Output - Data Recording - Bus Analyzers

For more information, please visit
processor.vmetro.com/phoenix
or call (281) 584 0728

VMETRO
innovation deployed

RSC# 23 @www.mil-embedded.com/rsc

with the capability. So if you look at that and what technologies are out there, you know that open architectures with POSIX corporate middleware are critical.

Also, in February 2004 I spoke with NASA and they're now adopting the SCA. They actually had been looking at building their own standards and looked at ours and did a 168-page summary that said, bottom line: It's already here; we just have to adapt it for a space-based environment. They told me in 2003/2004 that they had four satellites waiting for proprietary software to show up so they could go launch the systems.

MIL EMBEDDED: ***What sort of technologies do we need going forward to actually bring JTRS into deployable fruition?***

MACLAIRD: Our biggest challenge is in the processing because of heat and the issues of quickly and efficiently dissipating that heat. In the past, we've done that with size and with fans, which is not the best way to operate radios in desert environments. Balancing environmental factors plays a big part in how you design and deliver capability to the war fighter. But that's not the only place where JTRS technology needs to be focused.

Other system needs include focusing on new battery and antenna technology, smaller, lighter, more capable, and relying on easy access to commercial markets. I've seen needs in the Special Operation Forces where they desire more batteries that are easily accessible (AA batteries, for example), and how do you do that – by going into a local Iraqi CVS or grocery store? War fighters are asking, "Why can't you build me a military radio capable of accepting commercial battery sources? And by the way, I want it to operate for 8 to 10 hours without taking out the batteries."

> **We're talking about SDR, with the emphasis on *software-defined* and yet every conversation we have in the marketplace is about the hardware.**

Also, JTRS radios have to talk to satellites, so we need transmit power and some proprietary [military] technology, but the radios still need to be reduced in size and weight. And we also need faster [red/black] encryption security chips – the ability to encode and decode and process information factors into the heat and battery issue, too.

MIL EMBEDDED: ***So what do you think the software impediment is in JTRS?***

MACLAIRD: From my perspective of watching how we do things in government, the issue has to do with the program pace. You go in, you buy a capability, and you want to build it in 36-42 months. By the time you get there, things have changed.

We're talking about SDR, with the emphasis on *software-defined* and yet every conversation we have in the marketplace is about the hardware. A lot of people don't get the fact that the value content in the future of these radios is going to be all software. Software needs to be produced very efficiently and, until now, that has been very difficult to do because open standards such as the SCA have been evolving.

MIL EMBEDDED: ***Let's go down that path for a moment. So what's the current state of the SCA?***

MACLAIRD: I understand that they recently released version 2.2.2, which is basically a debugged version of the SCA. Just prior to my departure, we had pulled a team of experts together from government and industry, people to attack current and future concerns. People like: Vanu's John Chapin and PrismTech's Dom Paniscotti and Jerry Bickle, Spacecoast Communications President John Bard, and Lee Pucker of Spectrum Signal along with some others. We had laid out a road map with what we called *SCA 3.x* that would allow fixing of some of the problems of SCA 2.2 and migrate to a capability for above 2 GHz [RF]. And it also addressed a high-order language that some call *Modem HW Abstraction Layer* (MHAL). It created a similar higher order language or higher order abstract language capability that would allow the migration and maturation of the program.

MIL EMBEDDED: ***I would argue that the idea behind the SCA core framework was a good one, but here we are four years later and we're still not actually shipping JTRS. Are we going down the right path with SCA?***

MACLAIRD: As you go through any process, the issue is getting it to be socialized with the right people and adopted. There are a lot of people who have problems with opening up the architecture because they have proprietary solutions. PrismTech's view is that there is nothing fundamentally wrong with SCA. The 2.2 version does the job for which it was designed: The JTRS program provides radios, a common standard that has been reviewed and endorsed by the 130+ member Software Defined Radio Forum (SDRF) and the 880+ member Object Management Group (OMG). Both of these are well known standards bodies.

MIL EMBEDDED: ***If the military had to do this all over again, what do you think that this program should do differently, knowing what you now know?***

MACLAIRD: Organization and financial structure were, in my mind, the biggest obstacles. Technology was there or would have evolved to get us to where we needed to be.

MIL EMBEDDED: ***Are there any technology issues that, if done differently, might have achieved more success sooner?***

MACLAIRD: I would not have put the majority of the waveforms on a single contract. We put 20 or 21 Waveforms (Wf) on the contract with Cluster 1, and it became a big output to manage

when the contract focus and difficulties seemed to be more on the hardware. I think if the organizational and financial structure and political structure had been set up right, more along the lines of what it is today, we could've been more successful than we were.

Col. Steven MacLaird *(USAF ret.) served as Program Director of the U.S. Department of Defense's JTRS Program Office. His responsibilities included: development and acquisition of a new family of SDRs for joint use throughout the armed services with the goal of replacing 750,000 radios in the 2 Mhz to 2 GHz radio spectrum; coordinating the development of the radio SCA into commercial and international standards; the development of JTRS radio families to serve domestic and international uses as well as encouraging their commercial use; overseeing five Service cluster acquisition programs valued at more than $9 billion. Steven is a 1978 distinguished graduate of Kansas State University's Reserve Officer Training Corps program.*

For more information, contact:

PrismTech
6 Lincoln Knoll Lane, Suite 100
Burlington, MA, 01803
Tel: 781-270-1177
Fax: 781-238-1700
Website: www.prismtechnologies.com

The next advancements in Software-Defined Radio

By Joseph M. Jacob

The commercial delivery of the first generation of Software-Defined Radios has shown that the technology is viable and scalable. In order to push into new markets and domains, Software-Defined Radio must meet the challenges of security, safety, and a smaller footprint.

As the first wave of Software-Defined Radios (SDRs) built for compliance with the U.S. Military's Joint Tactical Radio System(JTRS)SoftwareCommunications Architecture (SCA) becomes available, radio manufacturers continue to increase their use of COTS tools. Their goal is to reduce development and deployment costs while enabling more design options throughout the engineering process. Commercial companies are looking to leverage this substantial investment in technology and tools, in order to extend SCA-based SDRs to new domains and to new types of devices.

SDR defined

What is Software-Defined Radio?

In an effort to improve upon the flexibility, usability, and extensibility of radios, a variety of companies have been working for almost 20 years to create radios where the core functionality is implemented in software rather than hardware. A Software-Defined Radio (SDR) is a radio in which 100 percent of the modulation and demodulation is defined in software instead of being hardwired into the electronics. This means that the frequency band, performance, and functionality can be upgraded with a simple software download and update. In essence, an SDR is a radio that is substantially defined in software, with a physical layer behavior that can be significantly altered through changes to its software. SDR provides an efficient and comparatively inexpensive solution to the problem of building multimode, multiband, multifunctional wireless devices. An SDR is capable of being reconfigured to operate with different waveforms and protocols through dynamic loading of new waveforms and protocols. A waveform can contain a number of different parts: It may not be just *AM vs. FM* but also have security and safety characteristics built into the waveform itself. The figure below, courtesy of the Communications Research Centre of Canada, shows the steps in the Software-Defined Radio life cycle.

SDR Development Lifecycle

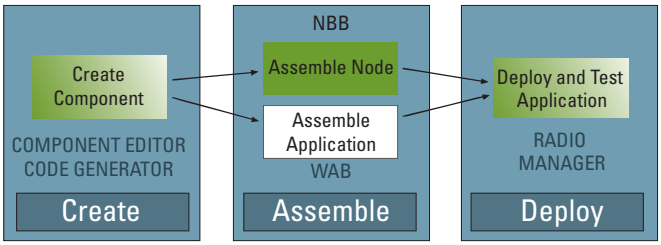

Software Communications Architecture CRC

Various groups see SDR as a solution to these needs. The military needs *smart radios* that can flexibly work in whatever country they are deployed, since they may be interacting with local forces on different networks. Cell phone makers need to consolidate the multimode radios they are building into their handsets and provide bug fixes with downloaded software. Public safety officials need a way to enable interagency communications problems during a crisis.

In order to achieve these goals, SDR technology needs to be further advanced to satisfy security, safety-critical, and footprint issues.

Security as a priority

Ensuring that radio communications are secure is one of the highest priorities for both military and commercial radio markets. In order for an SDR to be effective, the radio must implement robust security. In the military, security of communications on the battlefield is simply a basic requirement. Without it, the radio is useless, even harmful.

Radio often provides the only means of communication in high-threat military environments. Unfortunately, military personnel have not yet been able to trust that radios will be effective at keeping multiple levels of classified and unclassified transmissions separate. They need to have confidence that secret communications on one channel intended for U.S. forces only will not bleed into

unclassified channels, or be compromised by hostile third parties. It is important to note that SDRs transmit not only voice but also data. Voice transmission becomes only a small part of the overall usage of SDRs.

Both in the military and commercial markets, wireless transmission of information will continue to grow. Whether browsing on the Web, transmitting video, sending private financial information, or simply sending E-mails, the data component usage of SDR is the most significant component, and the one that often requires the highest security levels.

For public safety, security of communications during a time of crisis is essential to ensuring that public responders can get their job done. In the commercial world, handheld cell phones or radios will be used to do a variety of different tasks, including speaking with friends, sending text or video messages to colleagues, and communicating with the bank to engage in financial transactions. Consumer confidence in these devices will depend in large part on the level of security that they provide. If that device is accepting E-mail and instant messages while transmitting credit card and bank information, consumers will want assurance that subversive code contained in an E-mail or IM will not have any effect on their private financial information.

Virtually all Software-Defined Radios use a Real-Time Operating System (RTOS) as the core operating system for the radio. A remarkable amount of work has been done during the past five years in collaboration with the U.S. Air Force Research Laboratory and the National Security Agency to create highly secure versions of some of these RTOSs. The result of this effort is the Multiple Independent Levels of Security (MILS) architecture (www.mils.us). The MILS versions of the RTOSs are called *separation kernels*. Currently, three different RTOS vendors have publicly announced separation kernels: Green Hills, LynuxWorks, and Wind River. Although the initial development and deployment of these separation kernels is for defense applications, there is growing demand for these high-assurance MILS separation kernels throughout commercial domains as well.

The core security functionality of the MILS architecture is to keep data separate and to control information flows on a highly secure basis. The core foundational communications middleware for MILS, the Partitioning Communication System (PCS), ensures that only authenticated and authorized parties are allowed to exchange data, that their communications are secure and unbreakable, and that communications of data from different domains are kept separate over just one communications channel.

RTD Embedded Technologies, Inc.

"MIL Value for COTS prices"™

Geode cpuModules™

Pentium® M cpuModules™

8000 MIPS dspModules™

	cpuModules™ –40 to +85°C	Pentium® M				Intel® Celeron®						AMD Geode		
		CMX58886PX1400HR	CMV58886PX1400HR	CMX58886CX1000HR	CMV58886CX1000HR	CME47786CX650HR	CME47786HX650HR	CML47786CX650HR	CML47786HX650HR	CMX47786CX650HR	CMX47786HX650HR	CME26686HX333HR	CME27686HX333HR	CME27686CX333HR
Bus	AT Expansion Bus					✓	✓	✓	✓	✓	✓	✓	✓	✓
Bus	PCI Universal Expansion Bus	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓
Bus	PCI Bus Masters	4	4	4	4	4	4	4	4	4	4		4	4
Bus	APIC (add'l PCI interrupts)	9	9	9	9	9	9	9	9	9	9			
CPU and BIOS	CPU Max Clock Rate (MHz)	1400	1400	1000	1000	650	650	650	650	650	650	333	333	333
CPU and BIOS	L2 Cache	2MB	2MB	512k	512k	256k	256k	256k	256k	256k	256k	16K	16k	16k
CPU and BIOS	Intel SpeedStep Technology	✓	✓											
CPU and BIOS	ACPI Power Mgmt	2.0	2.0	2.0	2.0	1.0	1.0	1.0	1.0	1.0	1.0			
CPU and BIOS	Max Onboard DRAM (MB)	512	512	512	512	512	512	512	512	512	512	256	256	256
CPU and BIOS	RTD Enhanced Flash BIOS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
CPU and BIOS	Nonvolatile Configuration	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
CPU and BIOS	Quick Boot Option Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
CPU and BIOS	Fail Safe Boot ROM					✓	✓	✓	✓	✓	✓	✓	✓	✓
CPU and BIOS	USB Boot	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			
Peripherals	Watchdog Timer & RTC	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Peripherals	IDE and Floppy Controllers	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Peripherals	SSD Socket, 32 DIP						1		1		1	1	1	
Peripherals	ATA/IDE Disk Socket, 32 DIP	1	1	1	1	1		1		1				1
Peripherals	Audio	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			
Peripherals	Digital Video	LVDS	LVDS	LVDS	LVDS			TTL	TTL	LVDS	LVDS	TTL	TTL	TTL
Peripherals	Analog Video	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA
Peripherals	AT Keyboard/Utility Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Peripherals	PS/2 Mouse	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Peripherals	USB Mouse/Keyboard	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
I/O	RS-232/422/485 Ports	2	2	2	2	2	2	2	2	2	2	2	2	2
I/O	USB 2.0 Ports	2	4	2	4									
I/O	USB Ports					2	2	2	2	2	2	2	2	2
I/O	10/100Base-T Ethernet	1		1		1	1	1	1	1	1	1	1	1
I/O	ECP Parallel Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
I/O	aDIO(Advanced Digital I/O)	18	18	18	18	18	18	18	18	18	18	18	18	18
I/O	multiPort(aDIO, ECP, FDC)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SW	ROM-DOS Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SW	DOS, Windows, Linux	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

utilityModules™ –40 to +85°C

dspModules™
- Coprocessors
- Accelerators

Specialty I/O
- Pulse width modulator
- Incremental encoder
- Opto-isolated MOSFET

Frame Grabbers
- Single or multi-channel
- MPEG-2 compression

Video Controllers
- Analog VG A
- TTL and DVI panels

Communication Modules
- Copper or fiber Ethernet
- USB 2.0 and Firewire
- CAN Bus & CAN Spider
- Quad Serial w/ Ethernet
- Octal PCI Serial

Wireless Telematics
- GSM, GSM-R, CDMA
- EDGE, GPRS, SMS
- GPS, Wi-Fi, Bluetooth

Motion Controllers
- DC motor controllers
- Synchro, resolver, LVDT

Power Supplies
- 50/75/83/100 Watts
- Wide input range
- UPS backup
- MIL-STD-704/461

Mass Storage
- 1.8/2.5" IDE & PCMCIA
- CompactFlash

IDAN™— Intelligent Data Acquisition Node
- Easily build your PC/104 system
- Rugged PC/104 stackable framed modules
- Quick interchangeability and expansion
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- Standard PC connectors
- Optional MIL-SPEC paint & shock mounts
- –40 to +85 °C

Full Product Line and Pricing Online

A Founder of the PC/104 Consortium • ISO9001:2001 Certified

RSC# 28 @www.mil-embedded.com/rsc

www.rtd.com

HighRel PC/PCI-104 Modules and Systems

–40 to +85°C

Autonomous SmartCal™

Wireless Telematics

Frame Grabbers

	dataModules® –40 to +85°C	Smart A/D		Analog I/O					Digital I/O					
		SDM7540HR	SDM8540HR	DM6210HR	DM6420HR	DM6430HR	DM7520HR	DM6620HR	DM6812HR	DM6814/16HR	DM6856HR	DM6888HR	DM6956HR	DM7820HR
Bus	AT Expansion Bus	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PCI Expansion Bus Master	✓	✓				✓							✓
	McBSP Serial Ports	✓	✓				✓							
Analog Input	Single-Ended Inputs	16	16	16	16	16	16							
	Differential Inputs	8	8		8	8	8							
	Max Throughput (kHz)	1250	1250	40	500	100	1250							
	Max Resolution (bits)	12	12	12	12	16	12							
	Input Ranges/Gains	3/7	3/7	3/1	3/4	1/4	3/6							
	Autonomous SmartCal	✓	✓											
	Data Marker Inputs	3	3		3		3							
Conversions	Channel-Gain Table	8k	8k		8k	8k	8k							
	Scan/Burst/Multi-Burst	✓	✓		✓	✓	✓							
	A/D FIFO Buffer	8k	8k		8k	8k	8k							
	Sample Counter	✓	✓		✓	✓	✓							
	DMA or PCI Bus Master	✓	✓		✓	✓	✓	✓						✓
	SyncBus	✓	✓				✓							
Digital I/O	Total Digital I/O	16	16	16	16	16	16	16	48	18/9	32	64	32	48
	Bit Programmable I/O	8	8		8	8	8	8	24	6/0				48
	Advanced Interrupts	2	2		2	2	2	2	2					2
	Input FIFO Buffer	8k	8k		8k	8k	8k							4M
	Opto-Isolated Inputs										16	48	16	
	Opto-Isolated Outputs										16	16		
	User Timer/Counters	3	3	3	2	3	3	3	3	3				10
	External Trigger	✓	✓		✓	✓	✓	✓	✓					✓
	Incr. Encoder/PWM									3/9				
	Relay Outputs												16	
Analog Out	Analog Outputs	2	2		2	2	2	4						
	Max Throughput (kHz)	200	200		200	100	200	200						
	Resolution (bits)	12	12		12	16	12	12						
	Output Ranges	4	4		3	1	4	4						
	D/A FIFO Buffer	8k	8k				8k	8k						

RTD FieldPads™

- Ruggedized, embedded computer systems
- User-specified CPU and PC/PCI-104 expansion
- Weathertight components
- Integrated 6.5-inch video panel, keyboard
- Heat pipes for high performance CPUs
- User-defined MIL connectors
- Internal and external battery packs
- Expand with any RTD PC/PCI-104 product

Industrial FieldPad™
Ideal for control and monitoring of processes on factory floors or industrial installations. Mounting flanges allow the unit to be installed on machinery or walls, enabling standard PC access in a rugged enclosure resistant to industrial environments.

Tactical FieldPad™
Designed for mobile and portable applications where the angled panel and ergonomic design allow for optimal viewing with flexible positioning. Data collection/downloading and information access accomplished through wired or wireless connections.

HiDAN™ and HiDAN*plus*™— HighRel Intelligent Data Acquisition Node

- HiDAN is a rugged, watertight enclosure for a stack of PC/104 modules
- HiDAN*plus* combines the modularity of IDAN with the environmental ruggedness of HiDAN
- Integrated tongue and groove O-ring for environmental sealing and EMI suppression
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- **Stackable signal raceway**
- Optional MIL-SPEC paint
- MIL I/O connectors
- Shock-mount optional
- **–40** to **+85** °C

www.rtd.com
Specifications, manuals, drivers, and plant tour

RTD Embedded Technologies, Inc.
103 Innovation Blvd • State College, PA 16803
T: 814-234-8087 • F: 814-234-5218

"Accessing the Analog World"®

These separation kernels, as well as implementations of the PCS, are being certified to at least an EAL 6+ level under the Common Criteria, an internationally accepted security standard. This will be the first time that true COTS software is certified to such high levels. Both military and commercial radios will be able to use these validated, highly secure COTS components to build highly secure radios. Table 1 shows the different common criteria evaluation levels and the rigor with which software must be developed and analyzed at each level.

MILS is a perfect match for SDR because it ensures a high level of security while enabling modularity of new capabilities, provisioning of bandwidth and channels, and backwards compatibility with legacy radios. It also supports dynamic intra- and inter-network routing of data transparent to the radio operator. A number of different military SDR programs around the world are beginning development of radios using the MILS architecture to guarantee highly robust security.

Common criteria evaluation levels

Common criteria level	EAL description
EAL 1	**Functionally tested**
EAL 2	**Structurally tested**
EAL 3	**Methodically tested and checked**
EAL 4	**Methodically designed, tested, and reviewed**
EAL 5	**Semiformally designed and tested –** Must ensure resistance to penetration attacks with a moderate potential. Covert channel analysis and modular design are required.
EAL 6	**Semiformally verified, designed, and tested –** Must ensure resistance to penetration attacks with a high potential. The search for covert channels must be systemic. Development environment and configuration management controls are further strengthened.
EAL 7	**Formally verified, designed, and tested –** The formal model is supplemented by a formal presentation of the functional specifications and high-level design showing correspondence. Evidence of developer *white box* testing and complete independent confirmation of developer test results are required. Complexity of the design must be minimized.

Table 1

DO-178B levels

Level	Effect of anomalous behavior
A	Catastrophic failure condition for the aircraft (ex: aircraft crash)
B	Hazardous/severe failure condition for the aircraft (ex: several persons could be injured)
C	Major failure condition for the aircraft (ex: flight management system could be down and the pilot would have to complete it manually)
D	Minor failure condition for the aircraft (ex: some pilot-ground communications could have to be done manually)
E	No effect on aircraft operation or pilot workload (ex: entertainment features may be down)

Table 2

Safety-critical issues and requirements

Aircraft and avionics manufacturers are trying to achieve size, weight, and power savings by using multichannel, multiband Software-Defined Radios instead of separate radios for each waveform application. The software used in these radios will require certification by the Federal Aviation Administration in accordance with *DO-178B* at a level dictated by the impact to flight safety in the event of the radio's failure.

Although a great deal of the initial work for safety-critical SDRs is under the rubric of avionics systems, the results of this work apply to other safety-critical areas as well. These areas include medical devices, industrial process control, robotics, and so on, where communications via wireless devices are necessary and the flexibility of SDR is very useful, but where failure could lead to significant damage or even loss of life.

The FAA's Advisory Circular AC20-115B, produced by the Radio Technical Commission for Aeronautics (www.rtca.org), established DO-178B as the accepted means of certifying all new aviation software. The targeted DO-178B certification levels are either A, B, C, D, or E. Correspondingly, these DO-178B levels describe the consequences of a potential failure of the software: catastrophic, hazardous-severe, major, minor, or no effect. Table 2 shows the different DO-178B certification levels and the type of consequence if there is a failure at that level.

The generally accepted standard is that a normal aviation radio would be certified under Level C. However, if the SDR functionality will be implemented in a flight-critical component (for example, an instrument landing system), then the certification could be as high as Level A.

There are significant issues for certifying SDR for safety-critical applications.

The primary issue is that the fundamental benefit of SDR, which is its flexibility in dynamically updating and changing the system (including waveforms), rubs against traditional safety-critical principles of maintaining static configurations. For example, the ability to dynamically load and unload a waveform in flight runs counter to the traditional safety-critical principle of having a static configuration for the radio so that new possible threats to the radio's integrity are not introduced in-flight. Rather than try to shoehorn SDR into the traditional DO-178B analysis, the Avionics Special Interest Group (SIG) of the SDR Forum is taking a somewhat different approach. As discussed below, the Avionics SIG is working with accreditation agencies to demonstrate how modern software tools and techniques can ensure that certain types of dynamic behavior can meet and exceed the strictest levels of safety-critical scrutiny.

There is currently an effort underway by U.S. and European aviation certification authorities (together with RTCA) to revise DO-178B. This project, referred to as *Special Committee 205,* aims to create DO-178C, which will revise and update DO-178B. The Avionics SIG is working with SC-205 as well as with the entities responsible for the Integrated Modular Avionics certification issues. The focus of this effort is to keep the flexible and dynamic benefits of SDR, and to use modern software development tools and techniques to achieve the high level of safety-critical protection required for avionics. For example, code coverage tools, static code analyzers, and other new software testing tools can create a high level of assurance that flexible, dynamic object-oriented code meets Level A safety-critical protection. Likewise, the aforementioned high-assurance security techniques can potentially allow, for the first time, dynamic behavior previously prohibited. For example, loading and unloading waveforms in flight might now be allowed if a particular waveform has a high-assurance, FAA-digitally signed authorization certificate allowing it to be implemented dynamically into an SDR in-flight.

The immediate benefit of this work will be SDRs that can be used in military avionics where they might have a DO-178B/C requirement because of their use of civilian airspace, as well as in commercial avionics where the flexibility of SDR provides increased functionality for the aircraft. Work in this area will continue throughout 2006 and 2007 to create a set of standards that will allow SDRs to be certified to high levels of safety-critical standards.

Reducing the footprint and improving the performance of SCA radios

In order to ensure that SDRs can fit into smaller and smaller devices to

extend their usability into new domains, there are a number of different efforts underway to reduce the size and improve the performance of SCA-based SDR. Smaller code size results in lower power requirements, since the device needs less processing power and less memory in order to achieve its functions. Two of the most significant ongoing efforts are SCA-Lite, and CORBA ORBs in an FPGA.

Adopting an SCA-based SDR architecture provides several benefits; among them are software re-use, well-tested COTS tools, field upgradeability, common hardware, and software platforms to reduce production cost. At the same time, some developers of very small form factor devices have raised concerns about the cost of full SCA compliance in terms of size, cost, and power. A number of companies interested in SDR have expressed concern that the current versions of the SCA may not fit into the very small form factors they plan on using for their Software-Defined Radios. At the SDR Forum (www.sdrforum.org), the leading organization for builders and users of Software-Defined Radios with more than 100 organizational members, there is work underway to rethink certain aspects of the SCA. The goal is to reduce and modify the required SCA components for systems that require a significantly smaller footprint than standard radios. The result will be to implement an *SCA-Lite* core framework for very small form factor commercial applications while preserving core functionalities.

At the same time, regardless of which version of the SCA they are using, radio builders are finding that they can never have too much processing power. As waveforms become larger and more complex, many radio builders are running into a performance wall as they try to manage multiple large waveforms at the same time. Software vendors have responded by providing technologies to increase performance. One of the most important technologies is to implement elements of an ORB directly into IP blocks on an FPGA. This can increase performance 10 to 100 times for selected functions. Radio builders can use a standard ORB for the general purpose processor and ORB IP blocks for selected functionality to significantly improve overall throughput of the radio. This will allow larger waveforms to be used in radios without overloading the processing power of a standard general purpose processor or DSP.

The next generation of SDR

As the SCA matures and radio builders acquire experience in developing and deploying SCA-based radios, they are already thinking about the next generation of SDR devices. In order to meet the performance requirements for increased waveform utilization while ensuring lower development and deployment costs, these radios will need to utilize the security, safety-critical, and footprint/power enhancements being developed and provided by COTS component vendors.

Joseph M. Jacob *is Senior Vice President, Sales and Marketing, at Objective Interface Systems, Inc., and is co-chair of the Avionics Special Interest Group at the SDR Forum. Since joining Objective Interface in 2003, Joe leads the company's sales, marketing, business development, and product management teams. Prior to joining Objective Interface, his professional experience included leading the international business development and strategy function for America Online. He holds a B.A. from the University of Illinois at Urbana-Champaign with a double major in Economics and Political Science. He also holds a J.D. degree from Harvard Law School.*

To learn more, contact Joseph at:

Objective Interface Systems, Inc.
13873 Park Center Road, Suite 360
Herndon, VA, 20171-3247
Tel: 703-295-6500
E-mail: info@ois.com

History of SDR

In the '90s, efforts began to intensify to find standards-based approaches to SDR. The U.S. military, through the JTRS program, worked with its allies to establish the internationally endorsed open Software Communications Architecture (SCA). This standard uses Common Object Request Broker Architecture (CORBA) on POSIX operating systems as a framework for the various software modules that define the radio's behavior.

As with any new, disruptive technology, the development of powerful SDRs has had its fair share of ups and downs. In the past few years, the technology has matured to the point where solid SCA-based SDR implementations in small form factors are now becoming available. For example, Thales Communications recently announced the first SCA-certified radio for the U.S. military, and several more are to follow shortly in the pipeline.

In fact, a number of turnkey development platforms are now available, providing an integrated COTS development platform – hardware, operating system, ORB, SCA core framework, and SDR development tools – that allows researchers and developers to begin building waveforms immediately. This substantially reduces the time and risk of developing a working small form factor SDR. In order for SDR to thrive and prosper, it needs to adapt to users' demands for safe, secure, smaller, and more efficient implementations.

When failure is not an option.

Extreme Environments?

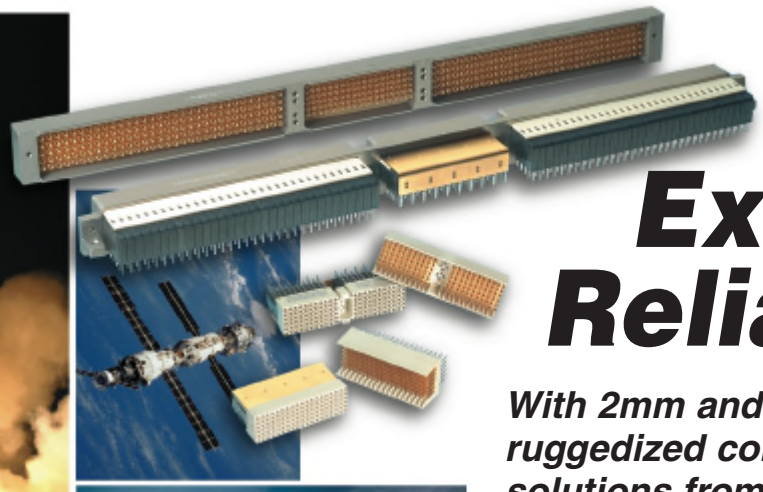

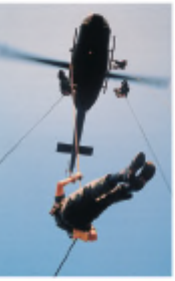

Get Extreme Reliability

With 2mm and VME64x ruggedized connector solutions from Hypertronics.

Whether your application is Land, Sea or Air, the legendary Hypertac® contact system provides:

- *Immunity to shock and vibration*
- *Up to 100,000 mating cycles*
- *Nearly half the contact resistance of conventional contact designs*
- *Extremely low contact insertion and extraction forces*

Hypertronics' 2mm and VME64x connector solutions are designed to be interchangeable with existing COTs systems.

Extreme environments need Hypertronics interconnect solutions - do it right the first time.

Visit www.hypertronics.com
for more information and to configure and download
3D connector models or call 800.225.9228 ext. 376

HYPERTAC
HYPERTRONICS
INTERCONNECT SOLUTIONS

ISO 9001, 13485, & 14001 certified *Hudson, MA* *800.225.9228* *www.hypertronics.com*

smiths

RSC# 33 @www.mil-embedded.com/rsc

Applying Model-Based Design to a Fault Detection, Isolation, and Recovery system

By Jason Ghidella, PhD,
and Pieter J. Mosterman, PhD

Model-Based Design facilitates verification and validation of an executable system specification, which prevents errors from persisting into later stages of the design process where they are more costly to eliminate. Studying a Fault Detection, Isolation, and Recovery (FDIR) system, in this case an aircraft elevator redundancy control system, demonstrates how to trace requirements to a design, create tests based on those requirements, and perform coverage analysis, which in turn reveals untested, missing, and ambiguous requirements as well as superfluous functionality in the specification.

The complexity of modern control systems makes it difficult to deliver high-quality, reliable, and timely products using traditional development processes. In such development processes, engineers first gather high-level requirements in text. These requirements form the system specification that is gradually refined into a detailed design that can be implemented. Structured test scenarios are then used during implementation to validate system behavior against the original requirements. However, errors found during implementation are costly and time consuming to fix, and can cause product delays, missed opportunities, and reduced functionality.

Companies are now adopting Model-Based Design to overcome these limitations. With Model-Based Design, engineers use a graphical model, often a block diagram, to capture requirements. This model produces an executable specification that describes the system behavior and can be gradually extended into an increasingly detailed design from which the implementation code can be automatically generated. Verification and validation can occur earlier in the system design process, which reduces costly iterations across many design steps.

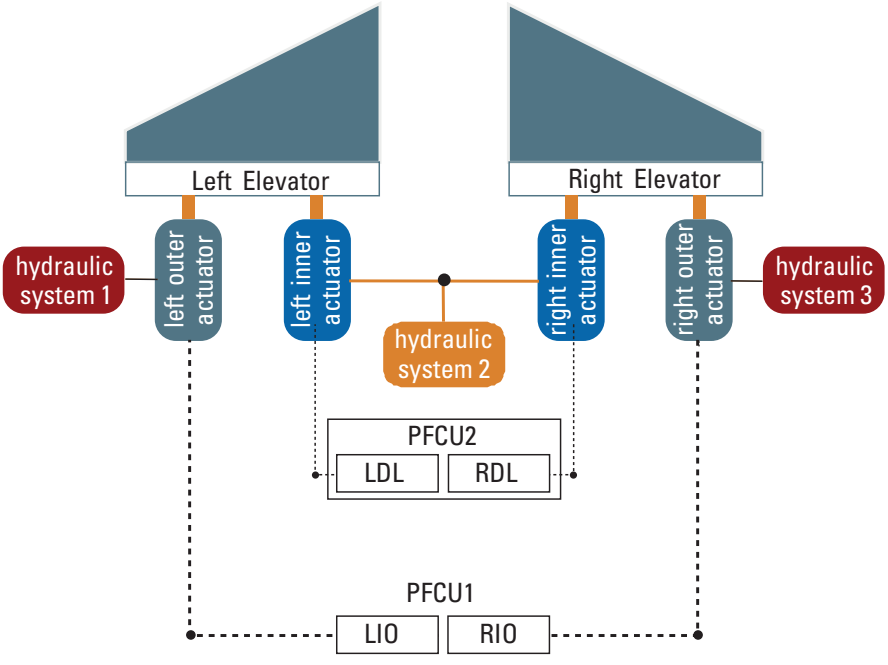

Figure 1

An FDIR application for a redundant actuator control system is designed using Model-Based Design. Associations are formed between textual requirements, items in the graphical model, and test cases used to verify and validate the design. Test completeness is determined by collecting coverage data that is derived, for example, by doing a Modified Condition/Decision Coverage (MC/DC) analysis of the design model as the tests are executed. Because this activity occurs at the model level, design errors and requirement omissions are identified and corrected early on, at less cost.

Designing an FDIR system

Figure 1 shows the typical redundancy of an aircraft elevator system with one elevator on the left and one on the right. Each elevator can be positioned by two hydraulic actuators, only one of which

should be active at any given time. Three separate hydraulic circuits drive the four actuators. The outer actuators use circuits 1 and 3, while the inner actuators share circuit 2. A Primary Flight Control Unit (PFCU), having a sophisticated Input-Output (I/O) control law, controls the left (LIO) and right (RIO) outer actuator. In case of a failure, a Direct-Link (DL) control law with reduced functionality handles the left (LDL) and right (RDL) inner actuators. Mode logic choreographs the redundancy to assure continual operation of the system.

Truth tables are used to model the combinational logic that detects sensor failures and isolates any faults in the system. Figure 2 shows the truth table (top) that isolates faults for the right-hand side actuators and a state transition diagram (bottom) that is used to handle the system recovery from the isolated fault. The state transition diagram shows the five possible actuator modes available to the four actuators LIO, LDL, RIO, and RDL: Isolated, Off, Passive, Standby, and Active. Note that even though the system has four actuators, there are symmetrical and dependence constraints placed upon them.

Developing detailed requirements

Table 1 shows a small set of high-level requirements used to guide the design of the mode-logic component. Such high-level requirements can be incomplete, inconsistent, and difficult to interpret. Design errors can be introduced simply by misinterpreting the high-level requirements into more detailed requirements. For example, the combination of requirements 2 and 3 from Table 1 creates an ambiguity:

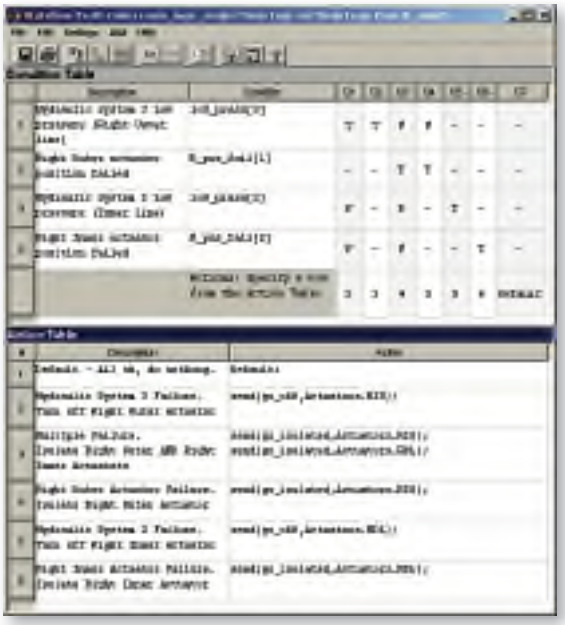

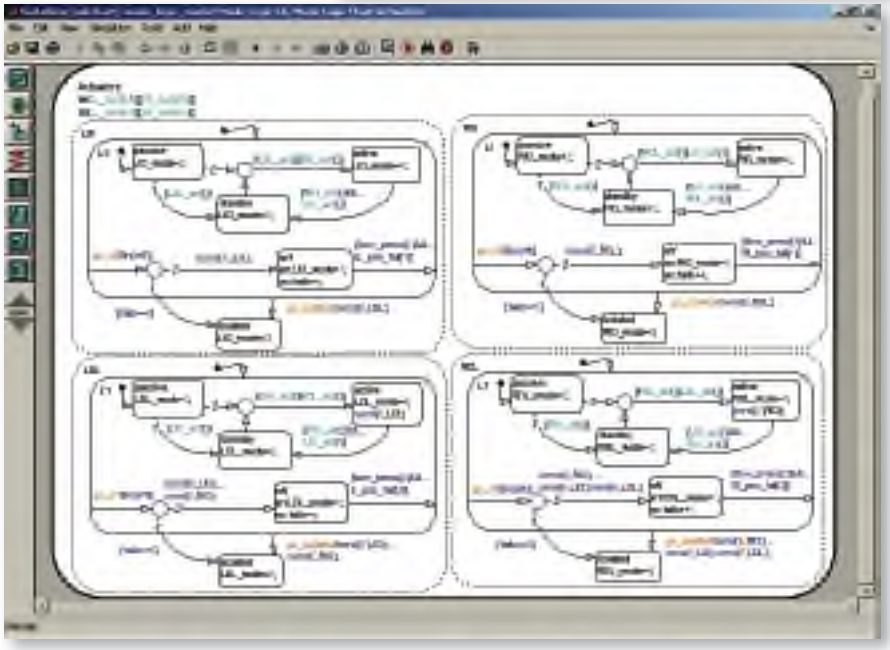

Figure 2

ID	Description
1	Each actuator will have five modes: Isolated, Off, Passive, Standby, and Active.
2	If possible, the same control law should be active for both the left and right elevators.
3	If available, the I/O control law should be active instead of the DL control law.
4	The actuator that is not active should be in standby.
5	If the pressure of the hydraulic circuit is low and the position measurement fails, the corresponding actuator should be switched to Off.
6	If the pressure of the hydraulic circuit is nominal and the position measurement fails, the corresponding actuator should be switched to Isolated.
7	Controller state changes should be made only in response to failure events.

Table 1

If one actuator can be operated only in DL, should the other still be operated in I/O? Therefore, it is necessary to test the detailed requirements early in the development process to ensure their accuracy.

Tracing requirements to design

Implementing the mode logic component as software requires translating the detailed requirements in Table 2 into an executable language such as C. Going directly to a programming language from the requirements, however, can leave too much room for interpretation. As a result, errors in the requirements or design may not be found until late in the development process, when subsystem and system integration is done, causing costly

2.1.1 Hydraulic pressure 1 failure If a failure is detected in the hydraulic pressure 1 system, while there are no other failures, isolate the fault by switching the left outer actuator to the Off mode.
2.1.2 Hydraulic pressure 1 fails and then recovers If a failure is detected in the hydraulic pressure 1 system and the system then recovers, switch the left outer actuator to the Standby mode.
2.1.3 Hydraulic pressure 2 failure If a failure is detected in the hydraulic pressure 2 system, while there are no other failures, isolate the fault by switching the left inner actuator and the right inner actuator to the Off mode.
2.1.4 Hydraulic pressure 2 fails and then recovers If a failure is detected in the hydraulic pressure 2 system and the system then recovers, switch the left inner actuator and the right inner actuator to the Standby mode.
2.1.5 Hydraulic pressure 3 failure If a failure is detected in the hydraulic pressure 3 system, while there are no other failures, isolate the fault by switching the right outer actuator to the Off mode.
2.1.6 Hydraulic pressure 3 fails and then recovers If a failure is detected in the hydraulic pressure 3 system and the system then recovers, switch the right outer actuator to the Standby mode.
2.2.1 Default start-up condition If there have been no failures detected, the outer actuators have priority over the inner actuators. Therefore, the elevator actuators should default to the following modes. The left outer and right outer actuators should transition from the Passive mode to the Active mode. The left inner and right inner actuators should transition from the Passive mode to the Standby mode.

Table 2

design iterations. Instead, modeling the design in a higher-level language, such as Simulink and Stateflow from The MathWorks, facilitates understanding of the design through simulation, so that system interactions can be studied much earlier. C code can then be automatically generated from the model, eliminating manual coding errors.

In constructing the design model, detailed requirements can be associated with elements in the model by using Simulink Verification and Validation. This association provides requirements traceability, which helps ensure that all the requirements have been incorporated into the design early in the development process.

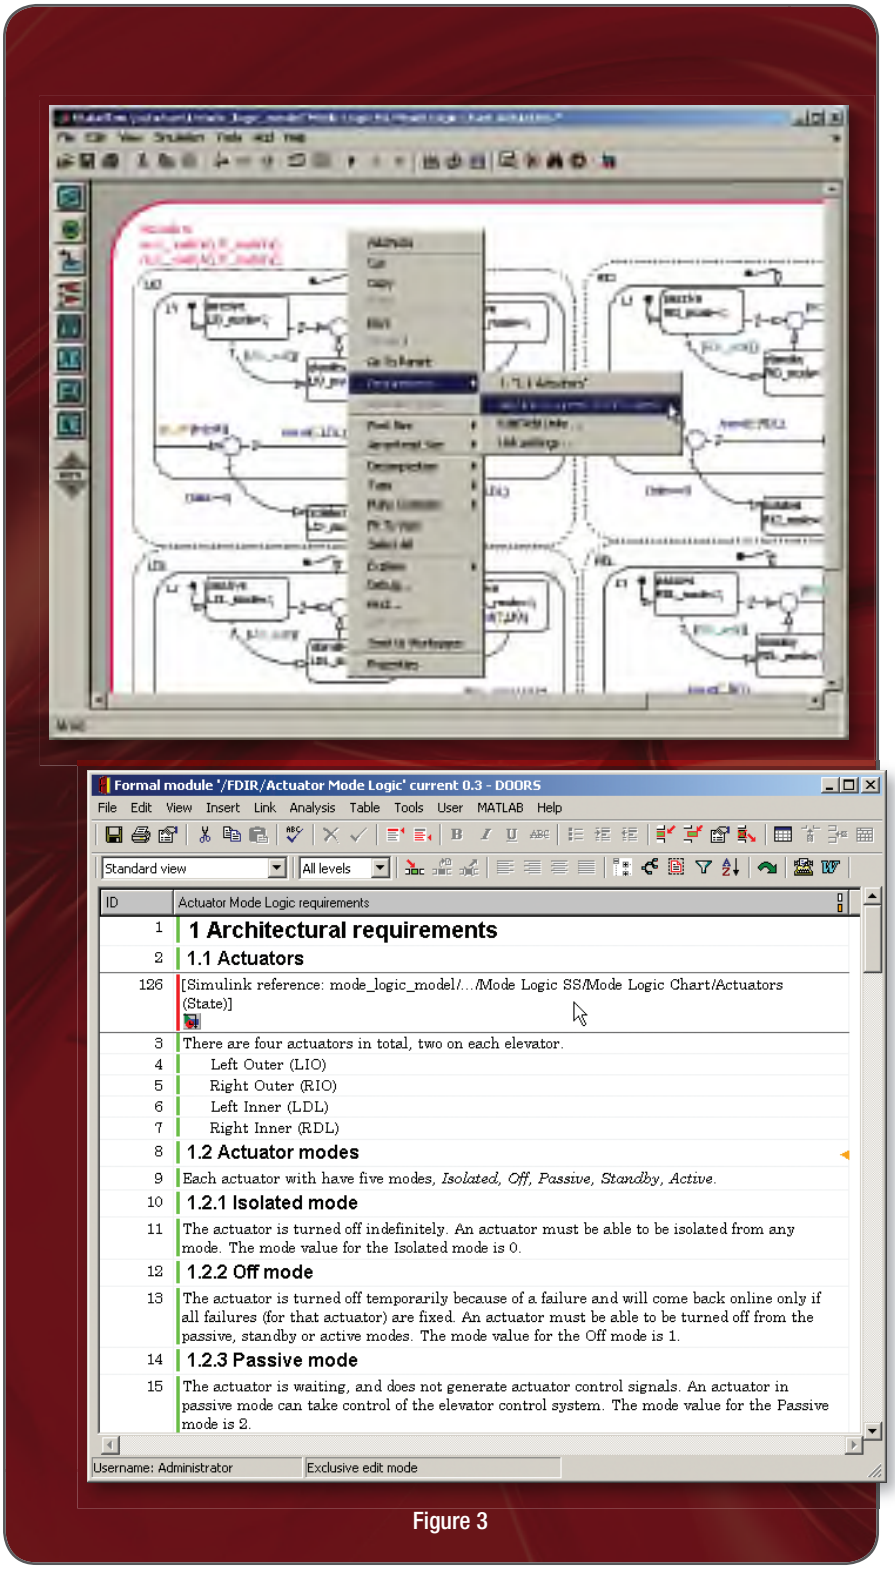

Figure 3

The requirements associations must be robust and easy to make. Context menus can be used to establish links between Simulink and Stateflow models and requirements-management systems, such as Telelogic DOORS, or documents such as Microsoft Word, Microsoft Excel, PDF, and HTML files. Figure 3 shows how an association is formed to DOORS between the design element *Actuators* state (top) and the detailed DOORS requirement 2, "1.1 Actuators" (bottom). Using an optional setting, item 126 is also inserted into the DOORS module for navigation back from DOORS to the design model, providing bidirectional connectivity.

Requirements, tests, and assertions

To test the requirements, a test harness is established in Simulink, where inputs are created in the Signal Builder block and system outputs are checked using verification blocks such as assertions. Different assertions can be used for each test case. Additionally, each test case can be associated with the requirement(s) that it tests. Figure 4 shows the test case where the hydraulic pressure of circuit 1 (H1) and the left outer actuator position sensor (LO_pos) both fail.

The right-hand panes of the test case show the verification blocks and requirements associations. For this test case, the expected output of the mode logic component is for the left outer actuator to move into the *Off* mode, the left inner and right inner actuator to move into the *Active* mode, and the right outer actuator to move into the *Standby* mode. The verification blocks check that this happens. There are two requirements associated with this test case, one that describes what happens to the left outer actuator when both the hydraulic pressure and actuator position fails, and another one to define the configuration of the remaining actuators in the system. Both associated requirements, when double-clicked, provide navigation back to the item in the requirements document.

Tests designed for all requirements can be combined into a test harness and applied to the design model for requirements-based testing. Tests can be executed individually or in batch mode. If a test runs without any verification blocks asserting, then the design passes the test. If an assertion is detected, the simulation stops and the verification block issuing the assertion is highlighted, which helps diagnose why the test failed.

Coverage

Creating and executing requirements-based tests to ensure that the design

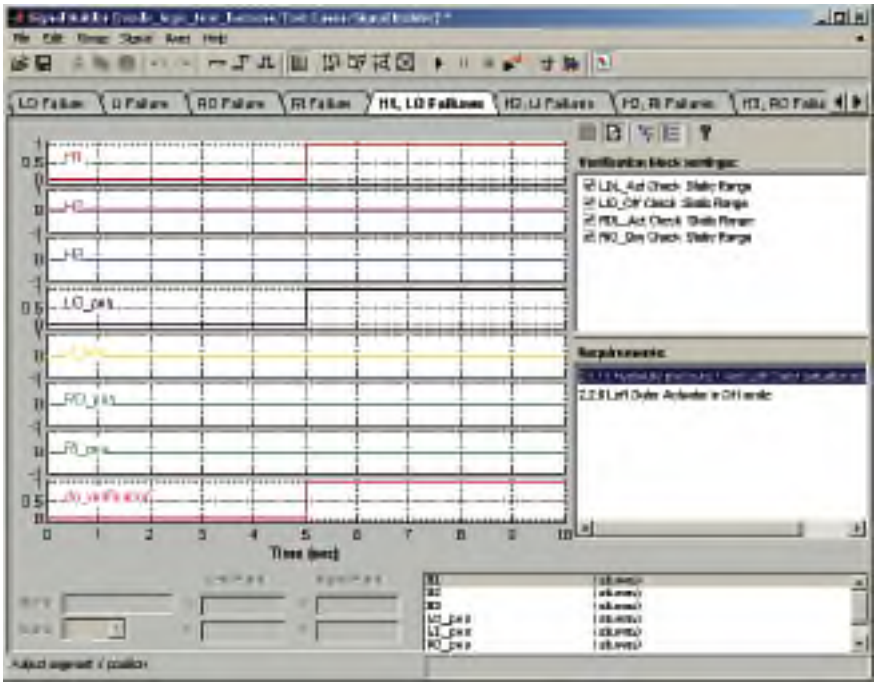

Figure 4

behaves as expected is not the same as fully testing the design. Some requirements may lack tests, the requirements themselves may be ambiguous or incomplete, and the design may contain superfluous elements.

By using Simulink Verification and Validation, coverage metrics can be collected during simulation to indicate untested design elements. The coverage metrics are displayed directly in the model using colored highlights and in a dialog box with summary information. Additionally, a detailed report of the coverage analysis is created. The coverage metrics collected include cyclomatic complexity, decision coverage, condition coverage, MC/DC, lookup table coverage, and signal range coverage.

The MC/DC metric is essential for DO-178B certification. It requires the execution of each separate input to a logical expression that can independently affect the outcome of the decision while other conditions are held constant.

Results of coverage analysis

Figure 5 shows the coverage analysis for the right outer actuator recovery logic in response to the execution of 23 requirements-based tests. Model elements, highlighted in red, indicate that full coverage of the design is not achieved. By selecting transition "[!RDL_act()|LIO_act()]" in the model, the dialog box shown in the bottom left-hand corner of the stateflow diagram summarizes the coverage analysis: *Full decision coverage. Condition "LIO_act()" was never true*. Clicking on the hyperlink in the dialog box displays the detailed coverage report in the bottom of Figure 5. Evaluating the MC/DC analysis reveals that the right outer actuator never transitioned into the *Active* mode because the left outer actuator was in the *Active* mode and the right inner actuator was also in the *Active* mode, a scenario that had not been set forth in the requirements. Because this scenario was missing, a test case had not been constructed. Additional requirements are thus necessary to fully describe the design.

In other design parts, coverage analysis helped discover and remove redundant design elements to simplify the design. Ambiguous requirements were discovered that resulted in an incorrect set of tests being executed on the design. After assessing all the design elements with incomplete coverage, two additional requirements were added, another two requirements

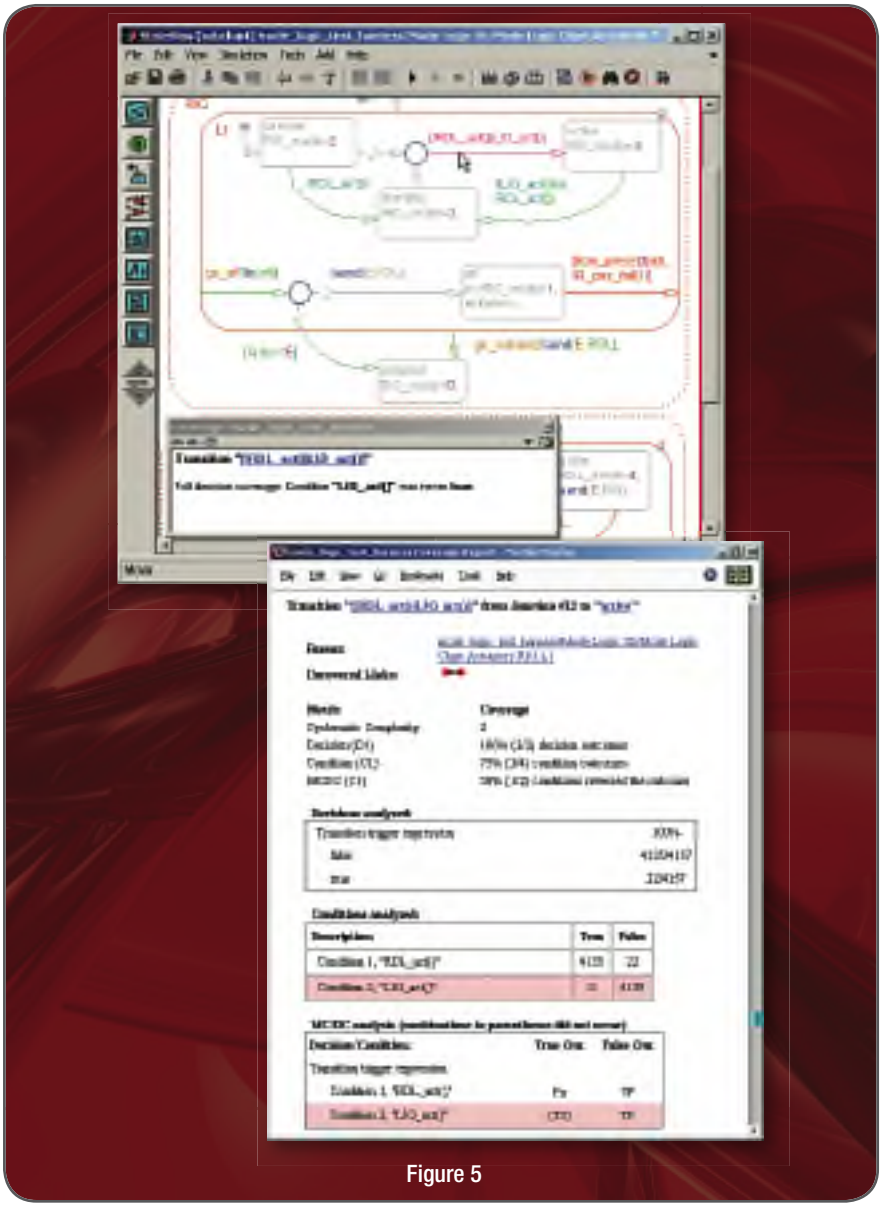

Figure 5

were rewritten to remove ambiguity, 11 tests were added, and unnecessary elements were removed. The resulting design has a cyclomatic complexity that is 8 percentage points less than the initial design and attains full MC/DC coverage.

Software implementation

Performing verification and validation tasks such as MC/DC analysis early in the development process helps demonstrate the correctness of the design and that implementation in software can be assumed. Automatic code generation provides a software implementation efficiently without the need to reinterpret the design by software engineers and eliminates potential coding errors. Traceability between requirements and generated code is achieved by including comments in code generated for each element in the design that has associated requirements.

A winning combination: Model-Based Design and MC/DC

In general, the requirements for engineered systems are ambiguous, not rigorous, and even inconsistent. It is desirable to detect such problems as early as possible in the system design process. Model-Based Design can be used to achieve this goal by facilitating executable specifications that allow testing to start at the model level instead of after system realization. MC/DC coverage analysis can help make the original requirements consistent and unambiguous, ensuring a minimal design and allowing the determination of test vectors needed for the given requirements to be established before software is implemented.

Jason R. Ghidella *works for The MathWorks as a senior technical lead responsible for Simulink platform product marketing. Prior to joining The MathWorks, Jason worked as an applications engineer at The MathWorks distributors in Australia and the Netherlands. Previously, Jason was a research scientist for two years at DSTO in Australia, working on fatigue analysis in the airframes and engines division. Jason earned a BE in Mechanical Engineering from James Cook University of North Queensland, Australia, and a PhD in Aeronautical Engineering from the University of Sydney.*

Pieter J. Mosterman *is a senior research scientist in modeling and simulation at The MathWorks, with a special interest in applying computer-automated multiparadigm modeling to model-based diagnosis and training systems. Previously, Pieter worked as a research associate at the German Aerospace Center in Oberpfaffenhofen through a grant awarded by the German Science Foundation (DFG). Pieter is editor-in-chief of* Simulation: Transactions of the Society for Modeling and Simulation International *for the methodology section and associate editor of* IEEE Transactions on Control Systems Technology *and of* Applied Intelligence. *He earned an M.Sc. in Electrical Engineering from the University of Twente in Enschede, the Netherlands, and a PhD in Electrical and Computer Engineering from Vanderbilt University.*

To learn more, contact Jason and Pieter at:

The MathWorks, Inc.
3 Apple Hill Drive
Natick, MA 01760-2098
E-mail: jason.ghidella@mathworks.com
or pieter.mosterman@mathworks.com
Website: www.mathworks.com

THE FUTURE OF ELECTRONIC SYSTEMS FOR MILITARY, DEFENCE AND AEROSPACE





Military & Aerospace Electronics
Technical Conference & Exhibition 2006

26th September 2006 • Ascot Racecourse, Ascot, Berkshire, UK

Register your place **NOW** for **MAE06** - the key technology event for all engineers, technical management, and research & development specialists involved in the design, integration, test and analysis of electronic components and embedded computing systems for the European military, defence and aerospace industries.

Conference topics will include:

- Designing with COTS
- Real-Time Operating Systems & Middleware
- Development Tools
- Microprocessors and Microcontrollers
- Power Electronics
- Optoelectronics
- Test and Measurement
- Rugged Technologies

Each of the seminars forming the three-track conference programme will be carefully selected by an independent panel of experts to enable you to learn about the latest technologies, and explore and interrogate the issues that affect your daily design, purchase and implementation decisions.

At **MAE06** you will get more than just the theory. Leading-edge solutions providers will be onsite to show you some of the most innovative, secure and reliable technologies for operation in harsh and mission-critical environments available today. As the only event in the UK to pack such a structured learning, fact-finding, and networking opportunity into a single day, **MAE06** is your must-attend event.

Platinum sponsors

FLOMERICS · Green Hills Software Ltd. · SBS Technologies · LYNUXWORKS™

Supported by

RADSTONE Embedded Computing · SAE-UK.org · IET · VITA Open Standards, Open Markets

Main media partner

OpenSystems Publishing™ · Military Embedded Systems · VMEbus Systems

Other media partners

MTe

Organised by

UMR

REGISTER ONLINE NOW AT: WWW.MAE-SHOW.COM
OR TELEPHONE OUR BOOKING HOTLINE ON: +44 (0) 1635 588860

RSC# 41 @www.mil-embedded.com/rsc

Migrating legacy applications to multicore processors

By Paul Leroux and Robert Craig, PhD

Compared to conventional processors, multicore chips offer a significant boost in processing capacity while consuming less power and less board space. But before migrating to multicore hardware, systems designers and software developers must choose a multiprocessing model that can maximize performance gains while minimizing modifications to existing software assets.

The footprint of multiprocessing systems has shrunk dramatically. A decade ago, they typically consisted of multiple interconnected boards, each running a single processor. They then evolved into multiple processors running on the same board, with each processor plugged into a separate board socket. And now, with the advent of multicore processors, they consist of multiple processing cores, all integrated on a single chip.

Much ink has been spilled over how this new era of "multiprocessing on a chip" will benefit desktop PCs, corporate servers, game consoles, and home multimedia centers. But the benefits for Defense and Aerospace (D&A) systems are, if anything, greater. Like systems in every other industry, mission computers and subsystems for radar, flight control, and sensor fusion are growing in complexity, with a voracious appetite for computational power. However, these systems must also satisfy rigorous requirements for low weight, low power consumption, and low heat dissipation, while adhering to existing form factors, backplanes, and chassis specifications.

Multicore chips satisfy these requirements by providing significantly greater processing capacity per ounce, per watt, and per square inch than their uniprocessor predecessors. By extension, boards based on multicore chips can lower the slot count and thereby reduce a system's weight, cost, power consumption, and overall chassis size.

Systems designers and software developers must learn to embrace multicore technology, not only because of its attendant benefits, but because it is quickly becoming the basis of most new processor designs. Nonetheless, multicore poses a significant software migration challenge for two reasons: a) few D&A developers have experience in multiprocessing systems, and b) the vast majority of legacy code in D&A systems was designed for uniprocessor, not multiprocessor, environments. Developers must graduate from a serial execution model, where software tasks take turns running on a single processor, to a concurrent execution model, where multiple software tasks run simultaneously. The more concurrency developers can achieve, the better their multicore systems will perform.

The first and most important decision developers must make when migrating to multicore is to select the appropriate form of multiprocessing for their application requirements. The choice will help determine how easily both new and existing code can achieve maximum concurrency. As Table 1 illustrates, developers have three basic forms: Asymmetric Multiprocessing (AMP), Symmetric Multiprocessing (SMP), and Bound Multiprocessing (BMP).

A familiar environment

AMP provides an execution environment similar to that of conventional uniprocessor systems. Consequently, it offers a relatively

Model	How it works	Key advantages
Asymmetric Multiprocessing (AMP)	A separate OS, or a separate copy of the same OS, manages each core. Typically, each software process is locked to a single core (for example, process A runs only on core 1, process B runs only on core 2, etc.). AMP is also known as *independent node architecture.*	Provides an execution environment similar to that of uniprocessor systems, allowing simple migration of legacy code. Also allows developers to manage each core independently.
Symmetric Multiprocessing (SMP)	A single OS manages all processor cores simultaneously. Processes can dynamically float to any core, enabling full utilization of all cores.	Can provide greater scalability and concurrency than AMP, along with simpler resource management.
Bound Multiprocessing (BMP)	A single OS manages all cores simultaneously. Processes can dynamically float to any core or be locked to a specific core.	Combines the scalability and transparent resource management of SMP with the developer control of AMP. The option to lock processes to specific cores allows simple migration of legacy code.

Table 1

GSPx Silicon Valley 2006

GLOBAL SIGNAL PROCESSING

Who attends GSPx:

To register visit
www.gspx.com/gspx

USE PROMOTIONAL CODE OSPGTC WHEN REGISTERING ON-LINE

Engineers, developers, project managers
The GSPx Technical Conference and Technical Workshops have been developed to quickly immerse you in the latest technological innovations.

C-Level executives
At GSPx you'll find a unique opportunity to network with your colleagues, gain insight into the trends that will affect your business and enhance your ROI.

Industry analysts
At GSPx, you'll find the content to track industry developments, identify emerging technologies and assess the state of the market

Co-located with:

TV2Mobile '06
::Business opportunities at the reach of your hand
Conference and Exhibition | Silicon Valley

GSPx Multi-core Applications
Super-charging your everyday life and security

GSPx is produced by Global Technology Conferences, Inc. | 1320 Centre St. Suite 200, Newton, MA 02459
+1 617.243.9777 | www.gspx.com | info@gspx.com

straightforward path for porting legacy code. It also allows developers to directly control how each CPU core is used and, in most cases, works with standard debugging tools and techniques.

AMP can be either homogeneous, where each core runs the same type and version of OS, or heterogeneous, where each core runs either a different OS or a different version of the same OS. In a homogeneous environment, developers can make best use of the multiple cores by choosing an OS, such as the QNX Neutrino RTOS, that supports a distributed programming model. Properly implemented, the model will allow applications running on one core to communicate transparently with applications and services (for example, device drivers, protocol stacks, and so forth) on other cores, but without the high CPU utilization imposed by traditional forms of interprocessor communication.

A heterogeneous environment has somewhat different requirements. In this case, developers must either implement a proprietary communications scheme or choose two OSs that share common protocols (likely IP-based) for interprocessor communications. To help avoid resource conflicts, the two OSs should also provide standardized mechanisms for accessing shared hardware components.

While useful for many applications, especially legacy code, AMP can result in underutilization of processor cores. For instance, if one core becomes busy, applications running on that core cannot, in most cases, migrate to a core that has more CPU cycles available. While such dynamic migration is possible, it typically involves complex checkpointing of the application's state and can result in a service interruption while the application is stopped on one core and restarted on another. This migration becomes even more difficult, if not impossible, if the cores use different OSs.

In AMP, neither OS *owns* the whole system. Consequently, the application designer, not the OS, must handle the complex task of managing shared hardware resources, including physical memory, peripheral usage, and interrupt handling. Resource contention can crop up during system initialization, during normal operations, on interrupts, and when errors occur. The application designer must design the system to accommodate all of these scenarios. The complexity of this task increases significantly as more cores are added, making AMP ill-suited to processors that integrate more than two cores.

Transparent resource management

Allocating resources in a multicore design can be difficult, especially when multiple software components are unaware of how other components are using those resources. SMP addresses the

REGISTER NOW

MILCOM 2006

NETWORKING THE WARFIGHTER

23-25 October 2006

Marriott Wardman Park and
Omni Shoreham Hotels
Ronald Reagan International
Trade Center
Washington, D.C.

The premier technical communications, networking, and information event!

Featuring classified and unclassified sessions, discussion panels, and tutorials, focused on advanced technology in military communications.

SESSIONS INCLUDE:

Ad Hoc Networking • Battlefield Network Vulnerability & Security
Black Core • Enterprise Architecture • IPv6 • Laser Communications
Mesh Networking • SATCOM • UAV Technology • Vulnerability

SPECIAL EVENTS:

Mission Possible – Battlefield of the Future!
A professionally staged mock exercise.

MILCOM Gala – Celebrating 25 Years of MILCOM with an unforgettable Silver Anniversary program of food, music, entertainment, and fun!

Oktoberfest – Raise a stein with your colleagues in the exhibit hall!

INCREDIBLE SPEAKER LINE-UP INCLUDES:

The Honorable John Grimes – Assistant Secretary of Defense for Networks and Information Integration and Department of Defense Chief Information Officer

Lt Gen Charles Croom, Jr., USAF – Director, Defense Information Systems Agency, MILCOM DoD Advisor

RDML Victor See, Jr., USN – Director, Communications Systems Acquisition and Operations Directorate, National Reconnaissance Office

Mr. Tom Brokaw

Former Anchor and Managing Editor, NBC Nightly News and Author, *A Long Way from Home* and *The Greatest Generation*

REGISTER NOW
WWW.MILCOM.ORG

NORTHROP GRUMMAN
DEFINING THE FUTURE

CISCO SYSTEMS

CORPORATE HOSTS

AFCEA

IEEE COMMUNICATIONS SOCIETY

MITRE
Solutions that make a difference

DISA

CO-SPONSORS

TECHNICAL SUPPORT

DOD ADVISOR

RSC# 45 @www.mil-embedded.com/rsc

issue by running only one copy of an OS on all of the chip's cores. Because the OS has insight into all system elements at all times, it can:

- Transparently allocate shared resources on the multiple cores with little or no input from the application designer
- Dynamically schedule any thread or application to run on any available processor core, allowing every core to be utilized as fully as possible
- Provide dynamic memory allocation, allowing all cores to draw on the full pool of available memory, without a performance penalty
- Allow applications running on different cores to communicate via simple POSIX primitives, such as semaphores; these primitives offer much higher performance and simpler synchronization than the networking protocols required in AMP systems

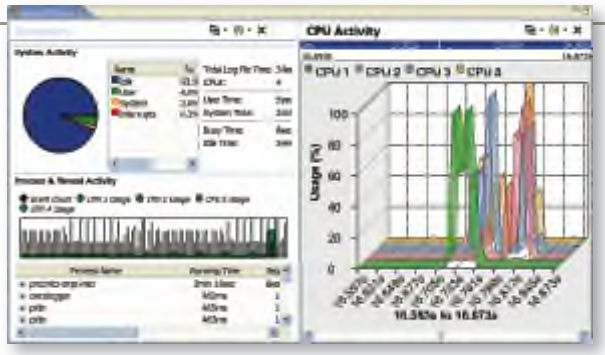

Figure 1

As an added benefit, SMP allows system tracing tools to gather operating statistics for the multicore chip as a whole, giving developers valuable insight into how to optimize and debug applications. Developers can track thread migration from one core to another, as well as OS primitive usage, scheduling events, core-to-core messaging, and other information useful for maximizing utilization of every core. In AMP, developers have to gather this information separately from each core and then somehow combine it for analysis. Figure 1 shows a system tracing tool, the QNX Momentics system profiler, being used to analyze a quad-core SMP system.

In Figure 2, a sonar system is running in SMP mode, which allows any thread in any process – data collection, signal processing, target tracking, and so forth – to run on any core. For instance, a target tracking thread can run one core while a signal processing thread performs compute-intensive calculations on another core. To implement the high-speed thread-to-thread communications needed for this scenario, developers can use either local OS primitives or synchronized protected shared memory structures.

Though it offers many advantages, SMP isn't a panacea. In particular, legacy applications with poor synchronization among threads may work incorrectly in the truly concurrent environment provided by SMP. This may not present a problem with software developed in-house but can create difficulties when a system must support software from multiple third-party suppliers.

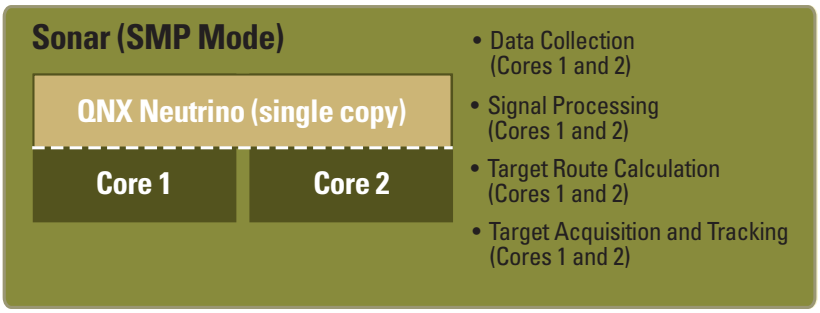

Figure 2

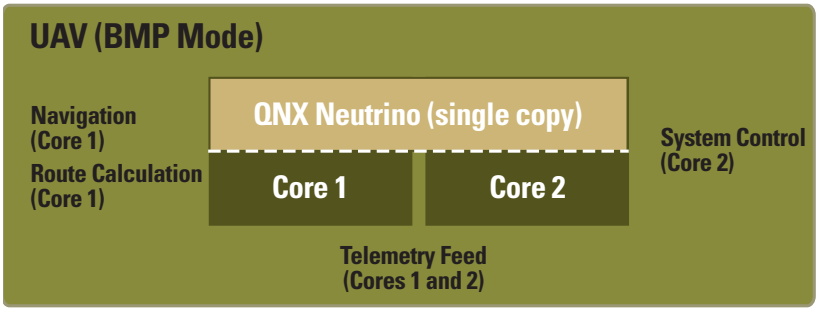

Figure 3

The best of both worlds

AMP offers greater developer control and compatibility with legacy code, whereas SMP offers greater scalability and simpler resource management. A third approach, BMP, combines benefits of both.

Like SMP, BMP uses a single copy of the OS to maintain an overall view of all system resources. Thus, the benefits of SMP remain intact. BMP goes beyond SMP, however, by giving developers the freedom to *lock* any applications to a specific core. This approach yields several benefits:

- Allows legacy applications written for uniprocessor environments to run correctly in a concurrent, multicore environment, without modifications
- Allows legacy applications to coexist with newer applications that take full advantage of the concurrent processing and dynamic load balancing enabled by multicore hardware
- Eliminates the processor-cache *thrashing* that can reduce performance in an SMP system by allowing applications that share the same data to run exclusively on the same core
- Enables simpler application debugging than traditional SMP by restricting all execution threads within an application to run on a single core

As in SMP, the OS is fully aware of what all the cores are doing, making performance information for the system as a whole readily available to system tracing tools. Using this information, developers can isolate potential concurrency issues down to the application and thread level. Resolving these issues can allow even legacy software to run with full concurrency, thereby maximizing performance gains provided by the multicore processor.

In Figure 3, a UAV subsystem is running in BMP mode on a dual-core processor, where navigation and route calculations are locked on one core and system control is locked on the other. The telemetry feed, which has less-intensive processing demands, can dynamically float to whichever core has the most available CPU cycles.

A matter of choice

The OS plays a key role in helping developers leverage the hardware parallelism offered by multicore processors. Unfortunately, the legacy RTOSs used in most D&A applications offer incomplete support for multiple processors running on the same computing platform, board, or chip. In most

cases, the kernels in these RTOSs can control only one CPU or processor core at a time. This restriction may be acceptable in an independent node (AMP) architecture but hampers flexibility of design and prohibits developers from achieving the considerable benefits offered by SMP and BMP. It's important, therefore, that the RTOS chosen for multicore designs can control multiple cores simultaneously and offer robust support for each multiprocessing model, giving developers the flexibility to choose the best form of multiprocessing for the job at hand. As Table 2 illustrates, the flexibility to choose from any of these models enables developers to strike an optimal balance between performance, scalability, and ease of migration.

	SMP	BMP	AMP
Seamless resource sharing	Yes	Yes	—
Scalable beyond dual core	Yes	Yes	Limited
Mixed OS environment (for instance, QNX Neutrino + Linux)	—	—	Yes
Dedicated processor by function	—	Yes	Yes
Inter-core messaging	Fast (OS primitives)	Fast (OS primitives)	Slower (application)
Thread synchronization between cores	Yes	Yes	—
Dynamic load balancing	Yes	Yes	—
System-wide debug and optimization	Yes	Yes	—

Table 2

Dr. Robert Craig is Senior Software Engineer for the OS Kernel Group at QNX Software Systems. With more than 12 years' experience in embedded systems design, Robert has worked extensively as both software architect and development team manager for various data communications companies. He holds a BS in Computer Science and Physics and a PhD in Physics with a focus on optical computing technologies.

Paul Leroux is a technology analyst at QNX Software Systems whose interests include high-availability design, multiprocessing systems, and OS kernel architecture.

For more information, contact Robert or Paul at:

QNX Software Systems
175 Terence Matthews Crescent
Ottawa, Ontario
Canada, K2M 1W8
Tel: 613-591-0931 or 800-676-0566
E-mail: info@qnx.com
Website: www.qnx.com

Processor type/ Max # processors*	Company/Website Model number	Description	Form factor
AMD Opteron 940 single core 2.2 GHz, or dual core 1.8 GHz processor 1*	**Performance Technologies** www.pt.com **CPC5564 SBC**	A 64-bit single slot CompactPCI single board computer designed for high-performance embedded applications	CompactPCI
Analog Devices ADSP-21062 SHARC 6	**Cornet Technology Communications** www.cornet.com/ctc **CVME-21062**	A 6U high-performance, multiprocessor OEM design solution for VME system designers	6U VME
Analog Devices ADSP-TS201 TigerSHARC 8	**BittWare** www.bittware.com **TS201 6U VME board**	A 6U VME board featuring eight ADSP-TS201 TigerSHARC DSPs	6U VME with VXS
Freescale 8641 PowerPC 2	**Curtiss-Wright Embedded** www.cwcembedded.com **VPX6-185 Dual PowerPC 8641**	A powerful general purpose single board computer with Freescale 8641 PowerPC processor	VME
Freescale dual AltiVec-based PowerPC, Dual Xilinx Virtex-II Pro XC2VP70 FPGAs 4	**VMETRO** www.vmetro.com **Phoenix VPF1**	A conduction-cooled, rugged version of the Phoenix VPF1 dual-FPGA/dual-PowerPC VME/VXS processing card	VME with VXS
Freescale dual PowerPC 2	**VMETRO** www.vmetro.com **Phoenix VXS Systems**	Systems built around high-performance processing, I/O, and multichannel Gbps serial communications with supporting software and firmware	VME with VXS
Freescale dual/single PowerPC 7457 2	**CES** www.ces.ch **RIO4 8070**	A PowerPC-based VME 2eSST reconfigurable computer	VME
Freescale dual/single PowerPC 750GX 2	**CES** www.ces.ch **RIO4 8076**	A conduction-cooled version of the RIO4 8070	VME
Freescale dual/single PowerPC G4+: MPC7448 2	**Aitech Defense Systems** www.rugged.com **C102**	A rugged dual PowerPC 7448 VME SBC	VME
Freescale MPC7448 2	**Mercury Computer Systems** www.mc.com **VPA-200**	VITA 41/VME single board computer	VME
Freescale MPC7448 or MPC7447A 4	**GE Fanuc Automation** www.gefanuc.com/embedded **Nexus Quattro**	A quad-processor 6U VME board based on four Freescale MPC7448 processors at up to 1.4 GHz or four MPC7447A processors at 1.1 GHz	6U VME
Freescale MPC8540 processor, dual C6416 DSP processor 2	**Beijing Fountain Microsystems** www.fountainsys.com **FTC-6000**	A high-performance DSP resource board designed for processing performance, network capability, and video/ voice processing in telecommunication, industrial control, and broadcast applications	CompactPCI
Freescale PowerPC 2	**Radstone Embedded Computing** www.radstone.com **PPCM2**	A high-performance 6U VME dual PowerPC processor	6U VME
Freescale PowerPC 7448 2	**Mercury Computer Systems** www.mc.com **Momentum Series CP3-102**	A 3U CompactPCI conduction-cooled SBC with dual PowerPC 7448 processors	3U CompactPCI
Freescale PowerPC 4	**GE Fanuc Automation** www.gefanuc.com/embedded **NEXUS Forte**	A PowerPC-based 6U VME board designed to meet demanding needs in communications, military/aerospace, homeland security, and life sciences applications	6U VME
Freescale quad 1 GHz PowerPC 7457 4	**Curtiss-Wright Embedded** www.cwcembedded.com **Manta QX3**	Quad PowerPC-based DSP 6U VME64x engine	6U VME
Freescale quad PowerPC 4	**Cornet Technology Communications** www.cornet.com/ctc **Celero CVME-7410**	A quad PowerPC single board computer	VME

***Note: "1" indicates a dual-core CPU**

Product Guide
MULTIPROCESSORS

Multiprocessor/multicore single board computers and DSP modules

Processor type/ Max # processors*	Company/Website Model number	Description	Form factor
Freescale single or dual MPC-7410 2	**Curtiss-Wright Embedded** www.cwcembedded.com **Falcon/51x**	A single and dual processor VMEbus SBC designed for commercial and semi-rugged military application environments	VME
Freescale single or dual MPC-7457 2	**Curtiss-Wright Embedded** www.cwcembedded.com **Manta/61x**	A single and dual processor VMEbus SBC designed for commercial and semi-rugged military application environments	VME
Freescale single or dual MPC-7457 2	**Curtiss-Wright Embedded** www.cwcembedded.com **Raptor/55x**	A single and dual processor VMEbus SBC designed for commercial and semi-rugged military application environments	VME
Freescale single or dual PowerPC 7447A 2	**Curtiss-Wright Embedded** www.cwcembedded.com **PMC-106 PowerPC PMC**	A high-performance PowerPC processor PMC	PMC
IBM PowerPC 750FX 2	**Thales** www.thalescomputers.com **PowerEngineC7**	A 3U single-slot PowerPC 750GX CompactPCI embedded computer	3U CompactPCI
IBM PowerPC PPC970 2	**Thales** www.thalescomputers.com **EasyG5**	An industrial version of the IBM Power Architecture technology in a convenient 6U format using dual PPC970 processors on a VME single board computer	6U VME
IBM single/dual PowerPC750GX 2	**Thales** www.thalescomputers.com **PowerEngine7-RC**	Conduction-cooled (RC) rugged versions of the PowerEngine7 SBC	6U VME
Intel Core Duo 1*	**BCM Advanced Research** www.BCMCOM.com **EBC5945GM, 5.25 SBC**	An SBC in a micro 5.25" EBX form factor board (5.75" x 8")	EBX
Intel Core Duo (T2500/L2400), or 1.6 GHz Core Solo (T1300) 1*	**Kontron** www.kontron.com **CP307**	A 3U Compact PCI processor board	3U CompactPCI
Intel Core Duo processor, Intel Core Solo processor 1*	**BCM Advanced Research** www.BCMCOM.com **EBC5945GM**	An Intel Core Duo/Solo processor mPGA 478 socket 5.25" SBC	EBX
Intel dual 3.4 GHz, 3.6 GHz, or LV 2.8 GHz Xeon 2	**One Stop Systems** www.onestopsystems.com **Dual Xeon SHB**	PCI Express-based system host board	PCI Express
Intel dual Pentium M 2	**General Micro Systems** www.gms4vme.com **S620 Hawk**	A Mini-ITX with dual PMC Pentium M SBC	Mini-ITX
Intel Pentium and PowerPC 2	**Thales** www.thalescomputers.com **PowerMP4**	Unique Pentium and PowerPC VME turnkey multiprocessor computer	VME
Intel Pentium D processor featuring dual-core architecture, Intel Pentium 4 processor supporting Hyper-Threading Technology, Intel Celeron D processor, and Intel Celeron processor 1*	**ITOX** www.itox.com **G7L330-B microATX**	A MicroATX motherboard utilizing the Intel 945G Express chipset	MicroATX
Intel Pentium D processor featuring dual-core architecture, Intel Pentium 4 processor supporting Hyper-Threading Technology, Intel Celeron D processor, and Intel Celeron processor 1*	**ITOX** www.itox.com **G7L630-B ATX**	A MicroATX motherboard utilizing the Intel 945G Express chipset	MicroATX
Intel Pentium D, dual-core 1*	**One Stop Systems** www.onestopsystems.com **MAX Express**	A graphics-class, PICMG 1.3 system host board	PCI Express

***Note: "1" indicates a dual-core CPU**

Product Guide
MULTIPROCESSORS

Multiprocessor/multicore single board computers and DSP modules

Processor type/ Max # processors*	Company/Website Model number	Description	Form factor
Intel Pentium M 2	**General Micro Systems** www.gms4vme.com **V469 PATRIOT**	A dual processor VXS 4.3 SBC	VME with VXS
Intel Xeon, dual-core 2	**Kontron** www.kontron.com **AT8020**	An open modular processing platform featuring two Intel Dual-Core Xeon processors and support for two AdvancedMC modules	AdvancedTCA
Xilinx dual Virtex-4 compute nodes and a PowerPC 744x 3	**Radstone Embedded Computing** www.radstone.com **V4DSP**	Dual Virtex-4 FX FPGA processor with MPC7448 GPP node	VME

Data was extracted from OSP's product database on August 3, 2006. First search included "Processors" category and "dual," "quad," and "multi" description keywords on products entered 1/1/06 through search date within *VMEbus Systems, CompactPCI and Advanced TCA Systems, Military Embedded Systems,* and *Embedded Computing Design* magazines. Second, additional search included "DSP Resource Boards" category keyword and the descriptors "dual," "quad," and "multi" on products entered 7/1/05 through search date in all OpenSystems Publishing magazines. Entries have been edited for publication, and OpenSystems Publishing is not responsible for errors or omissions. Vendors are encouraged to add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/.

***Note: "1" indicates a dual-core CPU**

New Products

By Sharon Schnakenburg

COMMUNICATIONS CONTROLLER

WIN Enterprises Inc.
Website: www.win-ent.com
Model: PL-01027 **RSC No:** 31216

A 1U scalable network appliance that provides a flexible platform for network security and other network-based applications • Fulfills several roles, including: policy gateway, wireless security gateway, content filtering, Network Attached Storage (NAS), and firewall • Architected using the VIA C7/ Eden processor with VIA CN700 and VIA VT8237R chipset • Available with a variety of VIA processors to meet a range of application requirements, including 400 MHz, 1.0 GHz, and 1.5 GHz • Four 10/100 Mpbs PCI bus Ethernet with two ports bypass function • Supports E-ATA/SATA HDD • One USB 2.0 port and one console port • One Mini-PCI expansion socket • One 50-pin CompactFlash type II socket • RoHS compliance • Compact size: 9.1"(W) x 6"(D) x 1.7"(H); 232 mm (W) x 153.3 mm (D) x 44 mm (H)

New Products continued on page 52

MIL-STD-1553

AcQ InduCom
Website: www.acq.nl
Model: USB403 RSC No: 30417

A portable interface unit providing full MIL-STD-1553 test, simulation, and bus analysis capability in a compact, self-contained unit, to connect via a USB interface to suitable host systems • MIL-STD-1553A/B, STANAG 3838 compatible • McAir compatible • 2 MB dual port RAM (all modes) • BIT and diagnostics • All setups programmable in real time • Programmable for direct or TX coupling • Programmable TX amplitude • Supports concurrent Bus Controller (BC) and up to 31 Remote Terminals (RTs) with Bus Monitor (BM) • Full error injection (BC and RT) • Full error detection (all modes)

POWER SUPPLY

Microchip Technology, Inc.
Website: www.microchip.com
Model: dsPIC DSCs RSC No: 31189

Digital signal controllers for digital control of the power-conversion feedback loop • 1 ns resolution PWM • 2 MSps, 10-bit A/D converter • Onboard analog comparators • Capable of running four simultaneous loops • 24-bit wide instructions, 16-bit data path • Up to 30 MIPS operation • 12 kB flash on-chip program space • 512 bytes on-chip data RAM • Modified Harvard architecture • C compiler-optimized instruction set architecture

PROCESSOR: PENTIUM M

VersaLogic Corp.
Website: www.versalogic.com
Model: Cheetah RSC No: 30338

A high-performance PC/104-*Plus* SBC featuring the 1.6 GHz Pentium M • Extreme Graphics 2 video: high-speed rendering and MPEG-2 support • SODIMM Memory Socket accommodates up to 1 GB of DDR RAM • Onboard sound, two USB 2.0 ports, two COM ports (one 422/485/232 configurable), IDE interface • TVS protection, which provides enhanced SD resistance • CompactFlash socket • Embedded BIOS • 400 MHz processor-side bus provides improved system throughput

Military EMBEDDED SYSTEMS

AN OPENSYSTEMS PUBLICATION

Military and Aerospace Group

- DSP&FPGA Product Resource Guide
- DSP-FPGA.com
- DSP-FPGA.com E-letter
- Military Embedded Systems
- Military Embedded Systems E-letter
- PC/104 Embedded Solutions
- PC/104 Embedded Solutions E-letter
- PC/104 & Small Form Factors Resource Guide
- VMEbus Systems
- VMEbus Systems E-letter

Group Editorial Director Chris A. Ciufo
cciufo@opensystems-publishing.com

Senior Editor (columns) Terri Thorson
tthorson@opensystems-publishing.com

Assistant Editor Sharon Schnakenburg
sschnakenburg@opensystems-publishing.com

European Representative Hermann Strass
hstrass@opensystems-publishing.com

Art Director Steph Sweet

Senior Web Developer Konrad Witte

Graphic Specialist David Diomede

Circulation/Office Manager Phyllis Thompson
subscriptions@opensystems-publishing.com

OpenSystems Publishing

Editorial/Production office:
16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystems-publishing.com

Publishers John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial Rosemary Kristoff

Advertising/Business Office
30233 Jefferson Avenue, St. Clair Shores, MI 48082
Tel: 586-415-6500 ■ Fax: 586-415-4882

Vice President Marketing & Sales Patrick Hopper
phopper@opensystems-publishing.com

Senior Account Manager Dennis Doyle
ddoyle@opensystems-publishing.com

Account Manager Doug Cordier
dcordier@opensystems-publishing.com

Account Manager Barbara Quinlan
bquinlan@opensystems-publishing.com

Account Manager Tom Varcie
tvarcie@opensystems-publishing.com

Advertising/Marketing Coordinator Andrea Stabile
astabile@opensystems-publishing.com

E-marketing Manager Christine Long
clong@opensystems-publishing.com

Regional Sales

Regional Manager – California Jane Hayward
jhayward@opensystems-publishing.com

Regional Manager – East Coast Phil Arndt
parndt@opensystems-publishing.com

Regional Manager – West Coast Richard Ayer
rayer@opensystems-publishing.com

International Sales

European Bureau Chief Stefan Baginski
sbaginski@opensystems-publishing.com

Account Manager – Israel Dan Aronovic
daronovic@opensystems-publishing.com

Reprints and PDFs

Call the sales office: 586-415-6500

Crosshairs Editorial

By Chris A. Ciufo
Group Editorial Director

What's up with Intel and AMD?

Everyone loses a little as titans collide, especially the mil market

It's pretty hard lately to miss the sparks between processor giants AMD and Intel, but isn't this really the same old rivalry? Not really. The prestige of having the highest performance mainstream microprocessor is still at stake, but both companies seem to be foolishly streamlining their product offerings to focus on winning this *epic conquest.* In the process, a lot of good technology is getting tossed out and the military market is going to suffer some pain. In this non-winnable battle, everyone loses.

I joined Advanced Micro Devices (AMD) in 1984 in the company's military group and quickly learned that there was bad blood between AMD and Intel. According to *Computer Chronicles,* it all started in 1976 when Intel agreed to cross-license the microcode to IT's iAPX family of microprocessors (8085, 8086, 80x86) and peripherals (8237, 8255, and so on). This was done solely so Intel could provide a second source to its target customers who would soon include IBM – the originator of the first mass market desktop *personal computer*. At the time, every semiconductor company had its own CPU, and developing market momentum behind one family was essential. You might recall that Fairchild, Motorola, National, Signetics, Texas Instruments, Zilog, and others all had processors in those early 8-bit days, so Intel's success with the x86 was far from assured.

AMD was little threat at the time, as the company was known simply as a *fast follower,* mostly selling versions of other companies' products with the unusual benefit of extended temperature MIL-SPEC versions as "883 For Free." Throughout my tenure, we routinely groused that Intel was slow to transfer its IP to AMD, thus allegedly keeping AMD at a disadvantage in the marketplace – giving rise to frequent lawsuits between the two companies over the years. By the time the 80386 hit the market in the late '80s, AMD began to create its own core microcode and occasionally outmaneuver Intel with higher performance versions. As I recall, AMD's '386 wasn't subject to the famous arithmetic 32-bit multiply instruction that left the high word in the result undefined[1]. By then, Intel had had enough and soon outdistanced AMD with a new revolutionary-for-the-time *Intel Inside* branding campaign, better marketing, and typically better CPUs.

Today, most experts agree that AMD has been the undisputed x86 performance leader for the past year or so, constantly turning in benchmarks such as those by EEMBC that surpass Intel. AMD was the first to introduce a 64-bit architecture and twist Microsoft's arm to create 64-bit versions of Windows XP, and was the first to endorse dual core versions. Intel's embarrassment has been widely reported – along with the company's sagging fortunes – despite Intel's overwhelming success with low-power Pentium M processors and the Centrino chipset that made Wi-Fi mainstream. With the bragging rights so high, AMD and Intel now seem predisposed to beat the other, regardless of the cost.[2]

And the cost is terribly high for the non-desktop PC market. In June, Intel sold its ARM10-based Xscale embedded processor line to peripheral chipset maker Marvell, while AMD dumped its still-fresh MIPS-based Alchemy embedded line to Raza Micro. Astoundingly, both of these product families were lightyears ahead of competing embedded offerings in terms of power/performance metrics, and one wonders if they'll ever re-emerge successfully under their new, less deep-pocketed owners. The market may suffer the pain.

In May, Intel also quietly announced the obsolescence of most of the x86 processors discussed previously. Before the end of 2006, finding '186, '386, and '486 microprocessors is going to be tough, though the company will accept orders until March 2007. Also affected are 8051 variants and i960 embedded processors – all of which are designed into countless military systems. While *Silicon Insider* analyst Jim Turley reports that the onus was really fab process capabilities, when you couple these actions with Intel's strategy shifts in telecom equipment and how it appears to be backing away from the ASI-SIG (a version of PCI Express), it's clear to me that Intel is laser-focused on retaining leading-edge processor bragging rights.

Not to be outdone, AMD just announced the huge $5.4 billion acquisition of GPU supplier ATI Technologies, effectively removing one of the two leading suppliers of graphics chips for the military market. It's clear to me that Intel's "integrated" chipset graphics solutions were a threat to AMD's processor strategy, so it had to add that capability right quick. Only NVIDIA remains as the dominant supplier in high-end graphics processors – a capability critical to military displays and simulation suppliers such as Quantum3D. But soon we'll have NVIDIA as a sole-source option.

Where's it going to end? As AMD and Intel shed product lines and business units in their struggle for processor bragging rights, the question becomes this: Is the rest of the market being moved forward by the processors' new capabilities, or dragged down by the carnage along the way?

1. This gave rise to Intel's 16-bit only 80386SX – a low-cost variant that set the stage for modern two-tier flavors such as Pentiums and Celerons.

2. As we went to press, benchmarks for Intel's very latest Core 2 Extreme processor (code named "Conroe") have become available. According to Loyd Case of Ziff Davis Media's ExtremeTech.com, the processor is "crazy fast" and blows most AMD benchmarks out of the water. Similar praise is handed out by Wil Harris at bit-tech.net, a highly respected independent benchmarking website.

GE Fanuc
Embedded Systems

Imagine what we can do for you now.

Our vision is to offer you the kind of products and services you've only dreamed of.

If you could build the perfect embedded products company, it might be something like the one we've created by combining the imagination, energy, and expertise of the people at GE Fanuc Embedded Systems, Condor Engineering and SBS Technologies.

It would be a company with the experience, resources and courage to develop truly innovative products—which is a cornerstone of the General Electric Company. It would be a company with a wide range of available products—everything from avionics and industrial controls to networking, communications, and fully integrated systems—which is exactly what Condor Engineering, SBS Technologies, and GE Fanuc bring to the table.

By creating this new company, we've taken a giant step toward our vision of a different and better kind of embedded company. So we invite you to open your mind and let your imagination run wild. Because our goal is to be right there with you helping make your inspirations into realities.

www.gefanucembedded.com

GE FANUC

CONDOR ENGINEERING

SBS Technologies

Now a part of GE Fanuc Embedded Systems

RSC# 55 @www.mil-embedded.com/rsc

PMC-706

PMC-706 is a ruggedized PMC graphics card supporting high-performance, dual headed synthetic graphics output. It utilizes the ATI MOBILITY RADEON™ 9000 (M9) processor to provide dual independent Analog and Digital video output in a wide variety of formats.

www.cwcembedded.com/grv1

PMC-724

The PMC-724 mezzanine module is a ruggedized IEEE 1386.1 PMC Frame Grabber providing high-performance image capture capabilities. It captures both analog and digital video input formats and supports high speed transfer to system memory.

www.cwcembedded.com/grv1

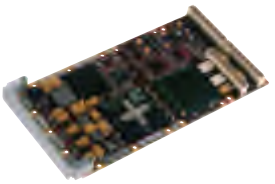

PMC-704

The PMC-704 is a ruggedized, high-performance, dual input, dual output, PMC graphics controller that is one of the first to utilize the industry leading ATI MOBILITY RADEON™ 9000 (M9) Visual Processing Unit.

www.cwcembedded.com/grv1

A HUMAN INTERFACE THAT'S DESIGNED FOR THE ELEMENTS.

Today's high-tech battlefield is a rough and challenging environment filled with manned and unmanned vehicles and data gathering devices. The human interface to these systems is more important than ever before, with sophisticated sensors feeding data to processors and high-resolution displays. That's why we've developed rugged, high-performance graphics products designed both for the elements and for the people who use them – and why you should partner with Curtiss-Wright.

Innovation In Motion

CURTISS WRIGHT Controls Embedded Computing

WWW.CWCEMBEDDED.COM

RSC# 56 @www.mil-embedded.com/rsc

SVME/DMV-183

Using single or dual Motorola MPC7447A/7448 PowerPC™ processors with AltiVec™ technology and 1 Gbyte of state-of-the-art DDR SDRAM, the SVME/DMV-183 is an ideal host card for any of our high-performance graphics PMCs.

www.cwcembedded.com/grv1

SOFTWARE

The Graphics Software Suite (GSS) can provide full X11 Server and OpenGL support. GSS provides optimal integration between graphics software and hardware and allows developers to more effectively develop graphical applications.

www.cwcembedded.com/grv1

GUARDIAN SELECT

Embedded Program Partner

GUARDIANSELECT

Curtiss-Wright offers system integrators a single end-to-end supplier of commercial and rugged grade COTS computing solutions and services. GuardianSelect services span the reach of complex program requirements – from design through to post production services.

www.cwcembedded.com/grv1

Images courtesy of www.af.mil and www.rafmarham.co.uk.