Military EMBEDDED SYSTEMS

VOLUME 7 NUMBER 2
MAR/APR 2011
MIL-EMBEDDED.COM

INCLUDING:

Chris A. Ciufo
Civilian COTS spurs competition

Field Intelligence
SBCs gain 256-bit vector processing

Mil Tech Insider
A case for RADHARD COTS

Legacy Software Migration
LDRA:
When code is born of old

Civilian tech serves up battlefield innovation

CPUs and servers used creatively

Also:
Rapid deployment of UAV ISR payloads

LTC William Geesey, PM MC4
"Telemedicine brings e-Health to the battlefield"

Military EMBEDDED SYSTEMS

March/April 2011 Volume 7 Number 2

COLUMNS

Field Intelligence

8 SBCs gain 256-bit vector processing
By Duncan Young

Mil Tech Insider

12 Off-the-shelf radiation resilience: A case for RADHARD COTS
By Steve Edwards

Legacy Software Migration

14 When code is born of old: Architectural obsolescence within the development life cycle
By Mark Pitchford, LDRA

Crosshairs Editorial

46 COTS means Competition, and that's good
By Chris A. Ciufo

DEPARTMENTS

16-17 Daily Briefing: News Snippets
By Sharon Hess

43-45 Editor's Choice Products

Hardware: Breakthroughs for the battlefield

18 MC4 telemedicine brings e-Health to the battlefield:
Interview with Army Lieutenant Colonel William Geesey, product manager, MC4

22 "I'm excited about the processor road maps."
Q&A with Matt Tracewell, Executive Vice President at Tracewell Systems

Software: Complex system? Model-Based Design's got it covered

24 Increasingly complex system designs need simulation, code generation tools
Interview with Dr. Jon Friedman, MathWorks

Mil Tech Trends: Controlling the UAV overhead (Part 1 of 2)

28 Case study: Secure FPGA technology enables UAV communications and control
By Jim Anderson, Xilinx

34 Rapid deployment of scalable, dense ISR payloads for UAVs
By David Pointer, SRC Computers, LLC

40 Autonomous ISR software reduces operator workload
An interview with executives from Kitware

E-CASTS

http://ecast.opensystemsmedia.com

Level-up with the 2nd Generation Intel Core Processor Family
March 29, 2011 • 10 a.m. MST
Presented by: Kontron, LynuxWorks, Intel Embedded Alliance

OpenVPX: From Specs to Solutions
March 30, 2011 • 11 a.m. MST
Presented by: Pentek, Curtiss-Wright Controls Electronic Systems, GE Intelligent Platforms, Curtiss-Wright Embedded Controls

Mitigating the risks of adopting multi-core processors for 40G telecom equipment
April 13, 2011 • 11 a.m. MST
Presented by: GE Intelligent Platforms, 6WIND

EVENT

http://events.opensystemsmedia.com

Embedded Systems Conference - Silicon Valley
May 2-5, 2011 • San Jose, CA
http://esc-sv09.techinsightsevents.com

WEB RESOURCES

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

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

ON THE COVER:
USMC Machine gunners Cpl. Charles Trask (aka "Gnome Warrior") and Cpl. Jimmy Miller (aka "Home-Grown," background) were awarded Purple Heart medals for battlefield injuries. They are painfully aware of the harshness of battle, and how hard the body and electronics have to work to survive. In this issue of *Military Embedded Systems*, we profile numerous new technologies designed to survive on the battlefield – and to make more successful *and lethal* the job of Marines and other warfighters. (Image courtesy of the U.S. Marine Corps)

Published by: OpenSystems media

ISSN: Print 1557-3222

kontron

» Where can I find a single source for my embedded computing project? «

Kontron is your single source for COTS products and specialized systems designs - built to meet the most demanding requirements.

COTS Boards & Enclosures
- Standard form factor
- Semi & full custom

Development Tools
- Software tools and development platforms for rapid application modeling and testing

Software Services
- BSPs: VxWorks, Lynxworks, Linux, Windows, ...
- Virtualisation
- IPMI

Systems Integration
- Full range of environmental compliance
- Choice of form factor or architecture
- Backplane or module based

Total Custom Solution
- Application ready
- Concept to production

CRITICAL QUESTIONS ... ANSWERED

VX6060
6U VPX Dual Intel® Core™ i7 Single Board Computer

- Available in air-cooled & conduction-cooled versions
- Compliant to VITA46 (VPX), VITA65 (OpenVPX) and VITA48 (VPX REDI)

Kontron EAPI
Kontron Embedded Application Programming Interface

- Unified middleware for access & control of hardware
- Includes standard libraries & APIs

Cobalt™
Reduced SWAP - High Performance Embedded Computer

- Small Form Factor
- Intel® Core™2 Duo or Atom™ processor options
- Custom I/O option

CONTACT US

AP Labs is now a part of kontron

☎ 1-888-294-4558

✉ info@us.kontron.com

www.kontron.com/military

intel Embedded Alliance Premier

If it's embedded, it's Kontron.

ADVERTISER INFORMATION

Page	Advertiser/Ad title
20	**ACCES I/O Products, Inc.** – USB embedded I/O solutions
31	**Aitech Defense Systems** – Rad-tolerant, rad-hard solutions
3	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
45	**Artisan Scientific** – Solutions for legacy systems
2	**Ballard Technology** – A surprising blend of personalities
26	**BittWare, Inc.** – Size, weight, and power issues?
44	**Concurrent Technologies, Inc.** – 2nd generation Intel Core processors
30	**CoreSolid Storage Solutions** – Mini PCI Express DiskOnModule
36	**Creative Electronic Systems** – Safe technical solutions
48	**Curtiss-Wright Controls Embedded Computing** – Defy size, weight and power
23	**D-TA Systems** – N X 10 Gigabit sensor processing
19	**Elma Electronic** – Your guide to high-speed backplane architectures
33	**Elma Electronic – Systems** – 3U CompactPCI system platforms ready to go
35	**Excalibur Systems, Inc.** – Dragon
21	**Extreme Engineering Solutions** – Designed with your application in mind
41	**Galleon Embedded Computing** – Setting new standards for rugged storage
47	**GE Intelligent Platforms, Inc.** – Why haven't you unleashed the power of embedded GPGPU?
9	**Hypertronics** – Hypertac contact technology
27	**InnoDisk** – i-DIMM industrial DRAM module
5	**Kontron** – Where can I find a single source for my embedded computing project?
33	**Kontron** – VX3035 – 2nd generation Intel Core i7 single board computer
25	**LCR Electronics** – Engineered for your application
38	**Memoright Memoritech (Shenzhen) Co., Ltd.** – The SSD expert
29	**Microhard Systems, Inc.** – Wireless innovation
10, 11	**Microsemi** – We create space
15	**Nallatech** – One stop SWaP
32	**Octagon Systems** – Redefining rugged
37	**Pentek, Inc.** – Stroke of genius
29	**Phoenix International** – VC1-250-SSD conduction cooled SATA solid state disk
39	**Sealevel Systems, Inc.** – USB to synchronous adapter
14	**VersaLogic Corp.** – Extreme embedded performance
7	**Wind River Aerospace & Defense Division** – Innovation
13	**WinSystems, Inc.** – 1 GHz PC/104 SBC supports networking and communications

OpenSystems media

Military EMBEDDED SYSTEMS

@military_cots

DSP-FPGA.com · VME and Critical Systems · PC/104 and Small Form Factors – The Journal of Modular Embedded Design · Industrial Embedded Systems · CompactPCI Advanced TCA & MicroTCA Systems · Embedded Computing Design

Military & Aerospace Group

Chris Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

Sharon Hess, Assistant Managing Editor
sharon_hess@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

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

Monique DeVoe, Web Content Editor
mdevoe@opensystemsmedia.com

Christine Capuano, Web/Editorial Assistant

Brandon Lewis, Copy Editor

Hermann Strass
European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Steph Sweet, Creative Director

Joann Toth, Senior Designer

David Diomede, Art Director

Matt Jones, Web Developer

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Sales Group

Dennis Doyle, Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Christine Long
Digital Content Manager
clong@opensystemsmedia.com

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

Regional Sales Managers
Barbara Quinlan, Midwest/Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

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

Ad Coordinator
Steph Sweet
ssweet@opensystemsmedia.com

Reprints and PDFs

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

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

INNOVATION IS THE FIRST UNMANNED AIRCRAFT THAT LANDS AT SEA.

Northrop Grumman X-47B
First tailless unmanned aircraft designed for autonomous, carrier-based capability.

Initial Flight:
Edwards AFB: 29 Minutes
4 Februrary 2011

How does the latest breakthrough in Unmanned Aircraft Systems go from concept to carrier deck? With help from Wind River. Our VxWorks RTOS was chosen by the innovators at Northrop Grumman Corporation as the software platform for their Unmanned Combat Air System-Demonstration (UCAS-D) program, and by GE Aviation as the foundation for the Common Core System, the backbone of the UCAS-D computers, networks and interfacing electronics. Building upon VxWorks proven reliability and unmatched performance, project engineers were able to rapidly create, deploy and maintain safety-critical control systems. All while reducing development costs and maintaining schedule integrity. Proof that when innovators work together, the sky is hardly the limit.

Please visit **www.windriver.com/customers** to see how Wind River customers have reached new heights.

WIND RIVER

INNOVATORS START HERE.

Field Intelligence

SBCs gain 256-bit vector processing

By Duncan Young

SBCs are key elements of military embedded computing subsystems and systems. They represent the lifeblood of every COTS vendor's offerings, from the smallest computer on a module to the extensive, sophisticated offerings in 6U VME, CompactPCI, and VPX (VITA 46). Intel's Core i7 processor family has become widely adopted by these COTS vendors for its ease of use, high performance, and good thermal characteristics, specifically in its mobile and embedded forms. The latest Core i7, based on Intel's Sandy Bridge microarchitecture, heralds a new level of capability for embedded military computing applications. In addition to improvements in processor performance, it incorporates a high-end, on-chip graphics processor plus the first mainstream implementation of a 256-bit Single Instruction Multiple Data (SIMD) vector processor, known as *Advanced Vector Extensions* (AVX).

Microarchitecture

Sandy Bridge is the second generation of Intel's Core i7 processor. It uses a different microarchitecture from its Arrandale predecessor, with Sandy Bridge typically offering a 20 percent increase in processing performance, improved power/system management, plus a big boost to the on-chip graphics capability. Sandy Bridge uses a highly parallel internal ring bus architecture to interconnect the processor cores, L3 cache, and the new graphics processor. Also on this ring is the system agent (effectively replacing the Northbridge) that provides the external DDR3 memory interface and 16x PCI Express 2.0 for all off-chip I/O connections.

Unlike NVIDIA's or ATI's separate, high-performance Graphics Processor Units (GPUs) using massive arrays of programmable SIMD processor cores, Sandy Bridge's graphics engine implements the majority of its complex graphics functions in dedicated hardware. However, it also includes an array of SIMD shader processors plus larger register files to enable the dispatch of more instruction threads without starvation by the register file. Key to improved performance is the graphics processor's equal L3 cache access rights, equivalent to that of any processor cores on the internal ring bus.

The 256-bit vector processor

By processing two vectors in a single execution cycle, each Sandy Bridge core's 256-bit vector processing engine offers twice the data processing capability of the earlier Core i7's AVX processor. This performance boost for floating-point applications will renew pressure on vendors such as Freescale (which recently announced reintroduction of the AltiVec to their QorIQ line) and others. These vendors will need to keep up, as vector processing finds more applications in wireless base stations, communications, security, and signals intelligence. In addition to these high-volume applications in relatively benign environments, there are many rugged, real-time military applications – such as ground mobile or airborne radar and video sensor suites – that will benefit from the performance of AVX-based Digital Signal Processing (DSP).

Despite the performance potential of AVX, Core i7 was not designed solely for vector processing applications. Its primary role is clearly for general-purpose computing, server, or embedded applications across the entire spectrum of potential uses. In rugged applications these roles might range from equipment control and monitoring, fire control, and communications, to complex mission computing and sensor processing. SBCs for rugged, embedded applications are equally general purpose in nature, typically offering maximum flexibility and customization through the addition of PMC/XMC mezzanine modules, networking, and many onboard I/O interfaces.

Core i7's system agent (Northbridge) provides PCI Express 2.0 to simplify the connection of external interfaces such as Serial ATA (SATA), Universal Serial Bus (USB), serial ports, and network ports. PCI Express also offers SBC vendors the opportunity to differentiate their product offerings by providing, for example, avionics interfaces or multiple network or serial fabric choices on a single board. GE Intelligent Platforms' SBC624 6U VPX SBC (Figure 1) uses the dual- or quad-core, embedded 26xx or 27xx Core i7 device running at 2.1 GHz. The SBC also adds InfiniBand to the onboard I/O options as a low-latency, high-bandwidth, switched-fabric alternative to Ethernet for multicomputing and server applications.

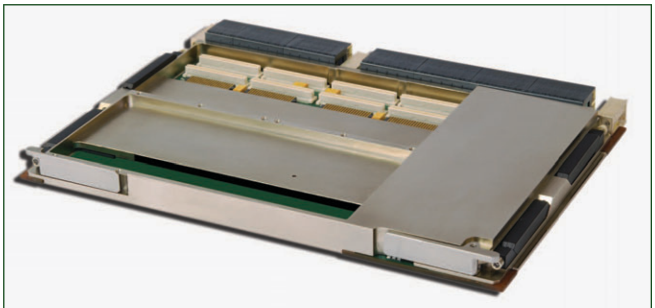

Figure 1 | The SBC624 6U VPX Core i7-based SBC from GE Intelligent Platforms

Range of DSP capability

The introduction of Sandy Bridge with its 256-bit vector processing capability to a broad range of SBC types, from the physically very compact 3U VPX profile upwards, will further encourage the assimilation of many DSP applications directly onto the SBC itself. Similarly, at the other extreme, multicompute DSP engines will benefit from quad-core devices, either reducing board count or increasing performance, to create a seamless range of DSP capability using the same AVX technology and software tools. Currently based on 32 nm process technology, a further boost in performance and reduction in power requirements is assured by the planned introduction of 22 nm process technology for the next generation of the Core i7 family in early 2012.

To learn more, e-mail Duncan at duncan_young1@sky.com.


We create space...
Microsemi

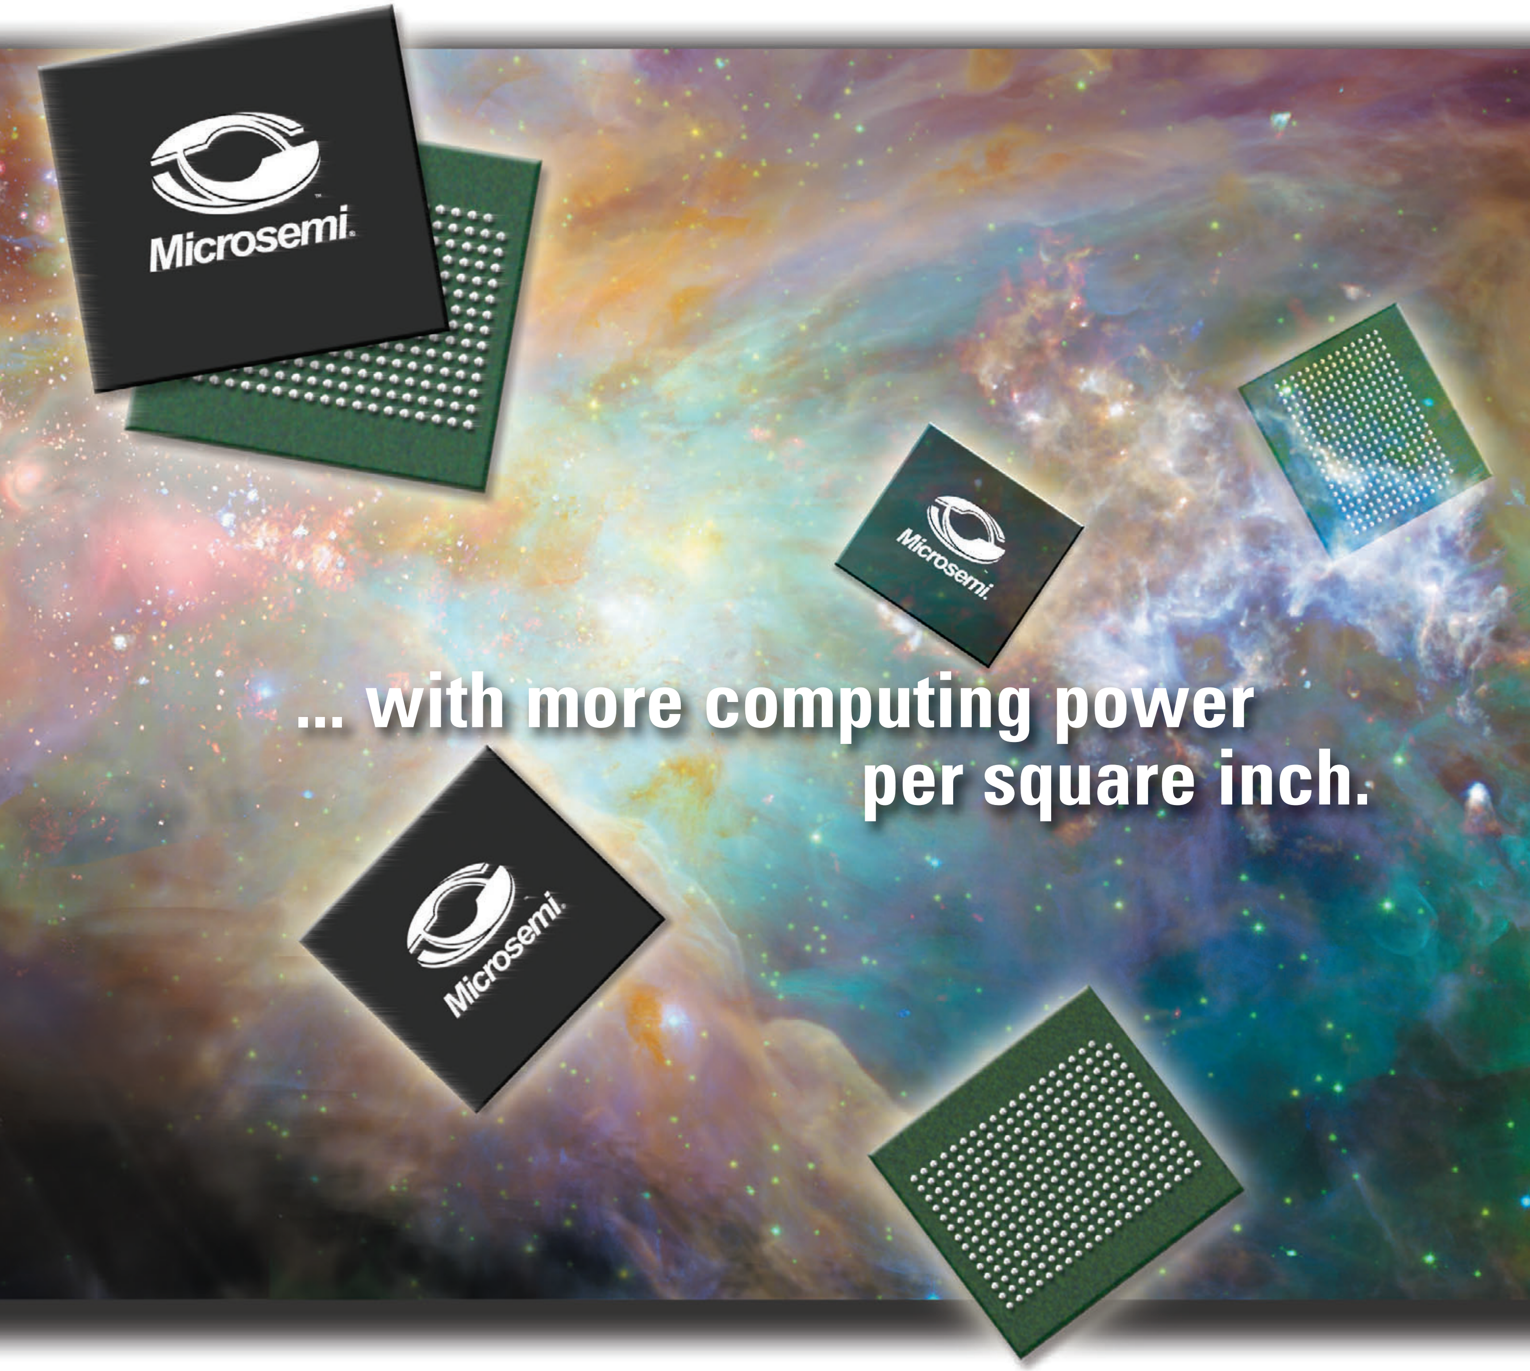
Microsemi
Microsemi
... with more computing power
per square inch.
Microsemi

Off-the-shelf radiation resilience: A case for RADHARD COTS

By Steve Edwards

Today's vehicle system designer must plan for the battlefield survivability of critical electronics during nuclear events. Electronic components deployed in space flight for flight-critical systems must be designed and hardened to continue to function while being exposed to high levels of gamma and neutron radiation for repeated or extended periods of time. In comparison, the radiation protection needed for electronic systems on ground vehicles – such as fire control systems and mission control computers – need not ensure continued functionality through nuclear events. Instead, there are many ground vehicle electronic systems that need only survive a nuclear event by detecting damaging levels of radiation, then shutting down and restarting after a specified period of time to then operate within higher levels of residual radiation. COTS components can mitigate these levels of battlefield radiation, achieving a satisfactory level of radiation resilience at a cost and complexity of 2 to 5 times less than that needed for space radiation hardening.

Radiation-resilient COTS

Radiation-resilient COTS products (or "RADHARD COTS") can address a large percentage of today's ground-based vehicle applications through a process that includes proper component selection, component and card-level testing, and evaluation. The harm to electronics depends on the density of the radiation created by a nuclear event in the form of gamma and neutron radiation. While there are always ambient (or background) levels of radiation present, they have no adverse effect on electronic circuitry. Today, most circuitry in COTS products is based on Complementary Metal Oxide Semiconductor (CMOS) technology.

Gamma radiation can cause an exposed powered CMOS circuit to upset or latchup. *Upset* is a condition where the device loses memory, changes state, or does not operate properly after the gamma dose. *Latchup,* the more damaging of the two, occurs when the parasitic Silicon-Controlled Rectifiers (SCRs) in the CMOS cells start conducting when exposed to a high gamma dose, resulting in a high current to the substrate and burnout. The typical latchup circumvention is to detect dangerous levels of gamma rays and turn off power to the circuitry before the latchup causes burnout. This can be done at the card or system (box) level. The traditional approach is to perform detection at the box level, protecting all cards in the system. If detection is done at the card level, all cards used in the system must have detection circuitry requiring the redesign of a portion of the cards used in the system.

Neutron radiation must also be mitigated. The CMOS devices are inherently resistant to most tactical neutron dose levels but can still experience upset. Neutron radiation can flip bits or flip the state of latches used in the circuitry. Compared to gamma radiation, which is gone quickly after a blast, neutron radiation lingers longer in the background radiation. Typical neutron-radiation circumvention includes providing Error Correction Code (ECC) on memories [especially Double Data Rate (DDR)], and parity bits or checksums to detect whether the accessed data is correct. If a transmission is bad, the system is allowed to request a retransmit or identify a failure and make the appropriate correction.

Getting to radiation resilience

Radiation resilience can be achieved with the proper selection of COTS components. The first step is for the system designer to review the system and hardware requirements to determine whether a RADHARD COTS product is suitable, rather than a more costly and complex radiation hardened product. For example, long-duration space applications are not generally well suited to RADHARD COTS.

BOM analysis and component selection

After the hardware requirements are determined for vehicle systems that might be exposed to nuclear radiation, the next step is to analyze the components used in the design, followed by proper component selection. Selection should be done to identify components that will not degrade when exposed to higher levels of radiation. Functionality such as ECC, parity, or checksums should be added to the system design to enable error detection and correction. Materials should also be analyzed for their suitability. For example, different types of optical fibers respond differently to nuclear events. Some fibers fail to function after exposure to radiation because the event detrimentally changes the structure of the fiber.

Lastly, radiation-resilient COTS components must be tested to ensure they will survive nuclear events. Increased customer requirements for ground vehicle programs have led Curtiss-Wright Controls Embedded Computing (CWCEC) to launch its RAD-HARD READY initiative and test radiation-resilient variants of standard VME and OpenVPX SBCs such as the SVME/DMV-183 at the White Sands Missile Range in New Mexico (Figure 1).

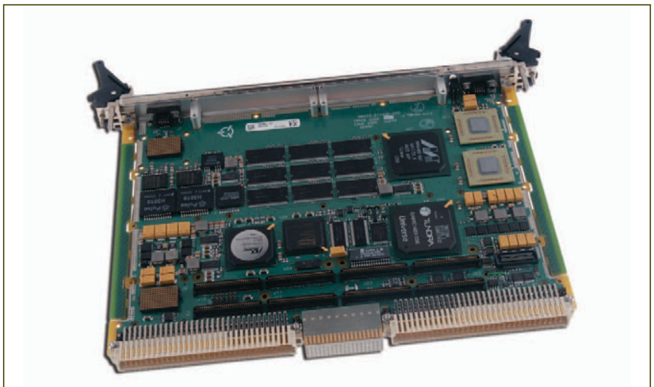

Figure 1 | The SVME/DMV-183 SBC from Curtiss-Wright Controls Embedded Computing

To learn more, e-mail Steve at Steve.Edwards@curtisswright.com.

1 GHz PC/104 SBC
Supports Networking and Communications

Hardened for harsh, rugged environments, the low power PCM-VDX-2 PC/104 single board computer is ideal for Mil/COTS applications. It is packed with serial, USB, and Ethernet ports, plus GPIO. With its diverse range of functions, most applications will not require additional I/O cards; however, it has Mini PCI and PC/104 connectors for specialty I/O expansion.

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

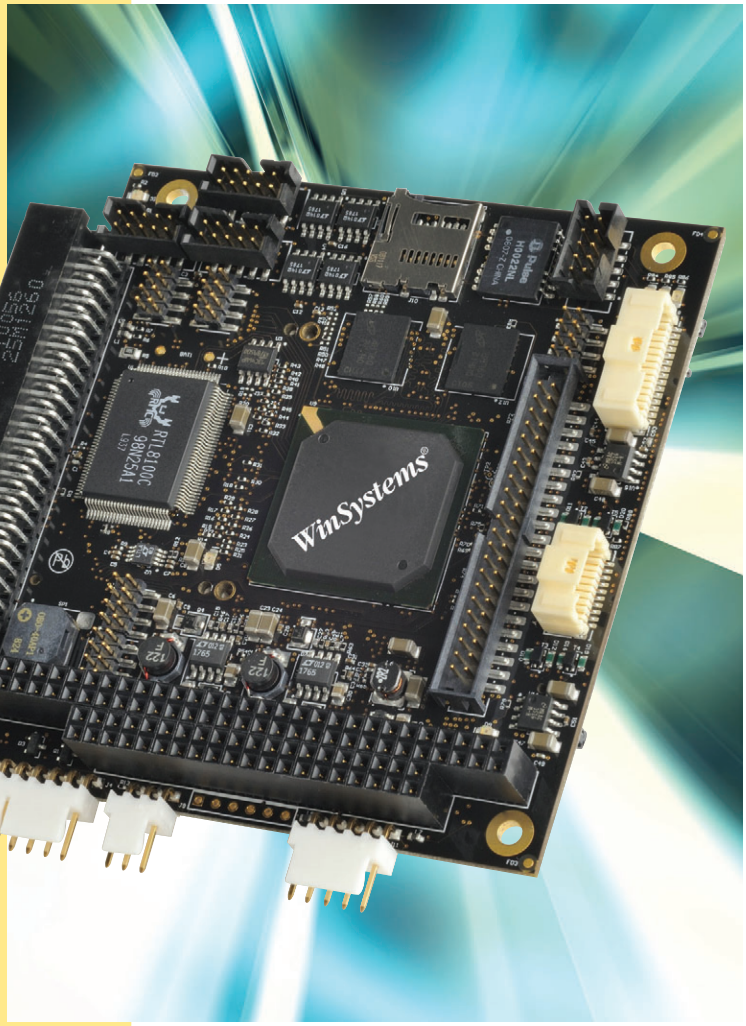


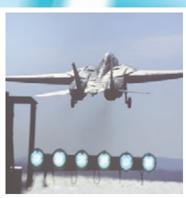

Understanding long-term product availability is a critical issue for Mil/COTS customers, the PCM-VDX-2 is offered beyond 2017.

Contact us for additional product information and pricing. Our factory application engineers look forward to working with you.

Ask about our 30-day product evaluation.

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

WinSystems®
An Employee-Owned Company

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

M2M 100 COMPANIES 2010 · sfig · Linux · EMBEDDED-PC PC/104 MODULES · Extended Temperature · RoHS compliant

Legacy Software Migration

By Mark Pitchford

When code is born old:

Architectural obsolescence within the development life cycle

Portability is not just relevant to software migration during the lifetime of a product. In complex military projects where the development cycle can easily take up to 10 years, software might have to be ported to new architectures two or three times. Unless specific steps are taken to ensure platform portability, such migration accelerates project overage, extending work schedules right through testing and verification. The ideal would be to anticipate such an occurrence and plan to minimize its impact, but the right software test tools can smooth the way even when porting legacy code written without such foresight.

Anticipating platform obsolescence

Choosing appropriate software test tools is always an essential step for safety-critical applications, and that is particularly true when portability is vital. It is well worth a little additional lateral thinking about how those tools might be applied to this specific circumstance.

The purpose of DO-178B is to provide guidelines for the production of software for airborne systems and equipment that performs its intended function while providing safety criticality. Although DO-178B demands the use of appropriate coding rules, it does not dictate what those rules should be.

Other standards, such as the Motor Industry Software Reliability Association C coding standard (MISRA C), MISRA C++, and Joint Strike Fighter Air Vehicle C++ Coding Standard (JSF++AV), specify detailed rules about how such safety-critical code should be written, and are found in an increasingly broad spectrum of applications. For instance, MISRA C was initially developed for use in automotive applications but has since been adopted and used across a wide variety of industries, including the medical, rail, aerospace, and military sectors. Sure, standards such as DO-178B will insist on the use of a rule set such as the MISRA and JSF++AV standards. The problem is that although these all acknowledge that portability is desirable and include rules relating to it, it is rarely the primary concern.

Adopting an appropriate rule set

Choosing a static analysis tool that encompasses many such standards will provide access to the portability rules of all of them. Considering additional portability rules in tandem with those from a chosen standard will ensure more focus on that element of the development work, and hence, ensure the minimum possible impact in the event of such a change.

There is no consequential compromise in the adherence to the standard of choice. And by choosing a rule superset of the nominated standard, it is possible to ensure that a change in architecture will have the minimum possible impact. With the correct tools, it is also perfectly possible to cross-reference Code Review Reports to a nominated standard and still include the additional portability rules.

Adapting to new target hardware

When the time comes to port code to a new target, it is useful to focus on only those parts of the rule set that have implications for portability, regardless of whether target obsolescence results from lengthy development times.

If the portability-focused rule set has been in place from the start, this exercise should largely be one of confirming that there are no problems. If not, the ability to filter static code review reporting to focus on portability becomes invaluable.

In this case, there are likely to be changes spread through the code base. So once the changes have been made, how can it be proven that the code functionality has not been compromised?

Ensuring uncompromised functionality

Unit test tools often encourage the collation of a suite of regression tests as development progresses. If such tests exist, rerunning them for a new target environment is a trivial exercise. If there is no such existing portfolio of unit test cases, an alternative approach needs to be considered.

For example, it is possible to automatically generate a set of unit test cases based on code analysis. These test cases can then be used to prove that functionality is identical before and after modification. While unit test cases not based on requirements are generally inferior because of a lack of independence, this approach bears scrutiny here because the primary requirement is to confirm that functionality has not changed.

Another approach involves comparing "baselines" from test tools, which can show how code has been changed. These can be used in conjunction with the tools' graphical representations of code architecture to ensure that even unchanged code is not affected by changes made elsewhere.

Such possibilities exist because modern test tool suites include a collection of tools that, although frequently used in a prescribed sequence, is often much more adaptable than such a rigid definition implies. By considering them not only as part of the development process but also as a tool kit, they have the flexibility to help in a plethora of situations … including this one.

Mark Pitchford *has more than 25 years of experience in software development for engineering applications. Since 2001, he has specialized in software testing and works as a Field Applications Engineer with LDRA. He can be contacted at Mark.Pitchford@ldra.com.*

Daily Briefing: News Snippets

By Sharon Hess, Assistant Managing Editor

www.mil-embedded.com/dailybriefing

Navy experiences open seas and open architectures

A recent $26 million IDIQ contract between the U.S. Navy and General Dynamics Advanced Information Systems falls into lockstep with the Navy's open-architecture COTS initiative. Specifically, General Dynamics will ramp up production for the Common Display System (CDS), a Grade-A shock-qualified display system designed for myriad Navy platforms. CDS is based on commercial software and hardware and slated to ride on Navy aircraft, subs, and ships, with the possibility of future use by the USMC and U.S. allies. CDS is expected to adhere to the same display requirements as other Navy programs specify, including the Aegis modernization program (Figure 1). The contract will be fulfilled at Tallman, NY; Fairfax, VA; Smithfield, PA; and Freemont, CA by June 2013.

Figure 1 | General Dynamics will soon ramp up production of the CDS, which meets the display requirements of other Navy programs, including Aegis modernization. Aegis includes vessels such as this USS Arleigh Burke (DDG-51), for example. U.S. Navy photo

IT: Federal agencies get some support

The IT support realm is a real hodgepodge, and no two entities ever have identical IT needs. However, a recent $400 million Information Technology Services – Small Business (ITS-SB) contract between the U.S. Army and MicroTech aims to solve federal agencies' IT dilemmas, anyway. The IDIQ contract states that MicroTech will render IT services such as IPv6 engineering and information assurance, along with support of the Electronic Product Environment Assessment Tool (EPEAT) Green technology program and Independent Verification & Validation (IV&V). MicroTech will also handle integration, migration, maintenance, and warranty work. Entities to benefit from the contract include the U.S. DoD and U.S. Army, in addition to other Federal agencies.

SN and GMD programs go the distance

Just like the Energizer Bunny, military programs keep going and going and going … and the Missile Defense Agency (MDA) is following the trend. Case in point: MDA recently awarded a nearly $110 million contract modification to The Boeing Co. to keep MDA's Sensors (SN) and Ground Missile Defense (GMD) programs operable. Notably, Boeing will keep up its sustainment and operation services for SN and GMD to the tune of $37,900,000 for SN and $72,000,000 for GMD. Contract completion will occur at Colorado Springs, CO and Fort Greely, AK, with SN work scheduled for March through December 2011 and GMD commencing March until August 2011.

U.S. Army software products find common ground

It would be so much easier if all U.S. Army technology platforms were compatible and could "speak" to each other and operate cohesively. Such a paradigm might soon be reality, thanks to a recent up-to-$33 million contract between the U.S. Army and prime Future Skies, Inc. Under the contract, Future Skies will render evolution and maintenance of "Common Software products" (Figure 2). The contract includes work on the Command and Control Registry (C2R); Data Dissemination Service (DDS) Client Interface (DCI); Battle Command Staff Trainer (BCST); Common Messaging Processor (CMP); and Joint Consultation, Command, and Control Information Exchange Data Model (JC3IEDM) Validator, among others. The goal of the contract is to further Common Software's role as encompassed within the Project Manager Battle Command: to streamline systems into a common architecture to provide commanders heightened capabilities to collaborate across the spectrum of military operations, relative to mission execution and planning. The contract consists of two base years plus two option years.

Figure 2 | Future Skies, Inc. will provide evolution and maintenance of the U.S. Army's "Common Software products."

FLIR and EOSS get a new lease on life

After the glory of initial development and deployment, the more behind-the-scenes requirement for life-cycle support steps into the limelight. Case in point: U.S. Special Operations Command and Raytheon recently incarnated a life-cycle contractor support contract with a $48 million ceiling. The contract specifies that Raytheon affords Electro-Optical Sensor System (EOSS) and Forward Looking Infrared Radar System (FLIR) system and ancillary-equipment subsystem life-cycle support for the 160th Special Operations Aviation Regiment (Airborne) and Marine, Air Force, Navy, and Army components (Figure 3). Contract fulfillment is slated for El Segundo, CA and McKinney, TX through Feb. 2015.

Figure 3 | A new up-to-$48 million contract calls for Raytheon to render life-cycle contractor support for EOSS and FLIR systems. FLIR photo courtesy of the U.S. Air Force, by Staff Sgt. Quinton Russ

A330 MRTT flies to DO-178B Level A

With passenger safety on the line, there shouldn't be any guesswork when it comes to software coding and resultant functioning. Thus, certifying airborne software is imperative … and not easy. However, Airbus Military's A330 Multi Role Tanker Transport (MRTT) recently achieved DO-178B level A software certification, with the aid of AdaCore's GNATcheck tool. GNATcheck verifies software-coding compliance to DO-178B by enabling engineers to define a set of rules derived from a coding standard such as DO-178B. Meanwhile, the A330 MRTT – designed to transport aircraft and administer air-to-air refueling – is a military incarnation of Airbus's A330-200 airliner. Other notables of A330 MRTT include its full fly-by-wire, all-electrical flight control system, along with autonomous disconnect for both tanker and receiver.

For consideration in Daily Briefings, submit your press releases at http://submit.opensystemsmedia.com. Submission does not guarantee inclusion.

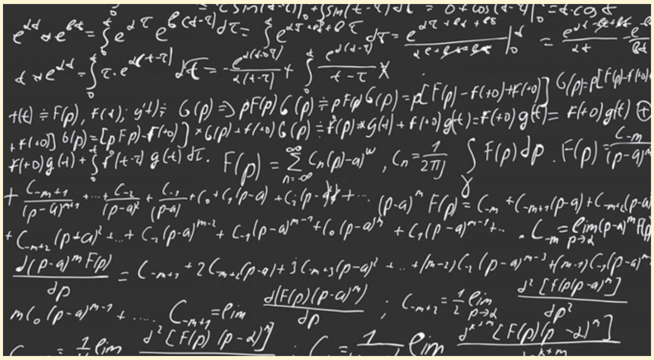

Figure 4 | A recent U.S. Navy contract stipulates that Centurum Information Technology, Inc. renders cryptologic systems support to both the U.S. Navy and the U.K. Navy. Stock photo

Navy systems get more secure

The ability to securely relay military intelligence and commands is paramount to mission success – and to national or international security. Thus, the U.S. Navy and Centurum Information Technology, Inc. recently penned a $15 million IDIQ contract for tactical cryptologic systems support (Figure 4). Specifically, Centurum Information Technology is to provide test and evaluation services for the U.K. Navy and U.S. Navy shore and shipboard cryptologic systems (i.e., systems engaging in scientific study of cryptanalysis and cryptography) globally. Work completion is anticipated to occur in Charleston, SC, by March 2012, or March 2016 if all contract options are activated. All options, if exercised, could escalate the contract's total to about $82 million.

Northrop Grumman sets the stage

When the project is as large an undertaking as the USS Abraham Lincoln's (CVN 72's) overhaul, the recent exercise of a $206 million option to an existing Northrop Grumman/U.S. Navy contract is likely not a surprise. Accordingly, under the option, Northrop Grumman Shipbuilding, Inc. will keep up its ongoing advance planning efforts to prepare for the overhaul. Said efforts include engineering, procurement, design, fabrication, and shipyard work as routine ship work, modernization, and refueling continue for CVN 72 and its reactor plants (Figure 5). Contract option activity takes place at Newport News, VA, slated for completion in February 2012. The contracting activity is the Naval Sea Systems Command in Washington, DC.

Figure 5 | A recently exercised $206 million contract option has Northrop Grumman prepping for CVN-72's overhaul. U.S. Navy photo by Capt. Lee Apsley

MC4 telemedicine brings e-Health to the battlefield

Interview with Army Lieutenant Colonel William Geesey, product manager, MC4

EDITOR'S NOTE

Battlefield medicine has come a long, long way since the Civil War, and advanced many clicks beyond what was commonplace only 10 years ago in OIF. The U.S. Army's Medical Communications for Combat Casualty Care (MC4) is integrating and deploying the best COTS technology available to further soldiers' golden hour, aid in mental health, and even facilitate "remote" surgeries directed by off-site surgeons. Increasingly, MC4's technological breakthroughs are rivaling those available to the general civilian population in America, as lumped under the latest "e-Health" buzzword. I had the privilege of speaking with Lieutenant Colonel William Geesey, product manager and "chief visionary" for MC4. Edited excerpts follow. – Chris A. Ciufo, Editor

MIL EMBEDDED: ***Let's take it from the top: Where are you based, and what is MC4?***

GEESEY: I'm based in Fort Detrick, Maryland, and MC4's higher headquarters is the Program Executive Office Enterprise Information Systems (PEO EIS). MC4 fields, trains, and sustains a comprehensive medical IM/IT [Information Management/Information Technology] system that has three primary capabilities: 1) an electronic medical record capability; 2) a medical logistics suite of applications; and 3) medical surveillance and command and control. Additionally, there's a suite of joint software applications developed by the military health system and employed by all of the services. We also have several other innovations in the works.

MIL EMBEDDED: ***Let's dig a bit deeper. First, what does MC4 do, more specifically?***

GEESEY: In addition to what I've mentioned, we also have a responsibility to implement Army-unique requirements. One of those is the Joint Theater Trauma Registry [JTTR], a software application in which medics in-theater can capture all data available from point of injury through the evacuation chain to facilitate research on better ways to create material solutions for saving soldiers' lives.

Another thing specifically mentioned in our ORD [Operational Requirements Document] is telemedicine/telehealth. In the past 18 months, we had a successful demonstration of telesurgery capability in Iraq in which a general surgeon performed a very difficult, complicated surgical procedure and saved a patient's life (Figure 1). This particular surgical procedure was not one the surgeon was familiar with, but was able to perform via telemedicine consultation with a CONUS [Continental United States] expert. We've also done some experimenting with TeleBurn, where a burn surgeon in Iraq wanted to provide consultation to other physicians with burn patients.

MIL EMBEDDED: ***You've described Army-unique requirements like the Joint Theater Trauma Registry, and also telemedicine – anything else for MC4?***

GEESEY: Yes, we have something that's really current: a pilot telebehavioral health project in Afghanistan and Iraq. MC4, PM Army Knowledge Online, the Office of the Surgeon General for the Army, and TATRC (the Telemedicine Advanced Technology Research Center) have collaborated on this project. The purpose is to develop new capabilities to improve the delivery of behavioral health services to deployed forces in Operation Enduring Freedom and Operation New Dawn.

There might be a situation, for example, in Afghanistan where soldiers are spread throughout many different bases. A divisional unit in Afghanistan might have one psychiatrist assigned, who would review the types of medications soldiers might take for depression or related issues. We've been able to set up a teleconsultive environment where a soldier goes into a clinic at one site with a non-mental-health-care provider, say a family practice physician, and engages in a teleconsultation with a psychiatrist located elsewhere. Of course, that allows the psychiatrist to perform a greater number of consults because they don't have to travel everywhere.

Figure 1 I Lt. Col. T. Sloane Guy IV received real-time consultation from providers at Brooke Army Medical Center (BAMC) in Fort Sam Houston, Texas, during a complex and rare surgical procedure while deployed with the 47th Combat Support Hospital in Mosul, Iraq. Photo courtesy of U.S. Army/MC4

MIL EMBEDDED: ***What's involved in these teleconsultations, technically?***

GEESEY: In Iraq, they're doing physician consults using the TANBERG system, which provides a VTC [Video Teleconferencing] capability using a camera and a computer. There are a lot of bandwidth restrictions and challenges in-theater, and this particular solution can regulate the flow of voice and video based on available bandwidth.

MIL EMBEDDED: ***Does the computer use a satellite uplink/downlink, or is it just plugged onto SIPR/NIPR?***

GEESEY: It can do both SIPR and NIPR in Iraq, and they're currently operating it on the NIPR. In Afghanistan, we're currently operating it on the SIPR because there's greater SIPR connectivity than NIPR connectivity there. So right now, they're using the TANBERG in Iraq, but in Afghanistan they're using MC4-supplied hardware, software, webcams, and headsets for provider-to-provider video teleconferencing.

MIL EMBEDDED: ***What's the next step for MC4 technology?***

GEESEY: The next step is to adapt a text-based secure chat capability using Army Knowledge Online [AKO]. A solider would be able to log in using their common access card that identifies who the individual is. And then there would be a link on the AKO website that says there's a mental health provider available; then if soldier wants to chat with someone, they could go ahead and click on that. This would be very useful in the field, as soldiers might receive the Dear Jane letter from home or something like that, or be bothered by a death of a family member or some other stressful event, combat event, or whatever. The part we're working on now is what we call a *chat queuing feature.* So if today I go in and speak with Dr. Smith, and the next day I want to follow up with Dr. Smith, I'm able to choose Dr. Smith versus just the next available person. And with the right security features, the soldier could sit in the privacy of their own room and talk to a mental health provider without anybody knowing.

MIL EMBEDDED: ***Makes sense. You talked about telemedicine, remote surgery teleconsultations, and mental health professional conferencing and secure chats. Are these all to be implemented by DoD personnel?***

GEESEY: Currently, all of the MC4 telemedicine activities are with either military providers or DoD-employed civilians. For CONUS, though, there's a network and there are consultations provided through the contracted support, contractor network, and Tri-Care support network. But right now in-theater, it's just military-to-military providers.

MIL EMBEDDED: ***Are there active plans within MC4 to describe to the civilian population the things implemented in MC4?***

GEESEY: There are several media outlet methods for that. As far as peer-reviewed journals, there has been some documentation and there's actually an article that I'm collaborating on that we're looking to get into a peer-reviewed journal at some point, talking about some of these telemedicine accomplishments and uses in the combat zone.

MIL EMBEDDED: ***Going back to the technology … what about other future technologies for MC4 – anything else in the works?***

GEESEY: Yes, one thing is hands-free electronic medical records. The point of injury where a medic or a buddy is rendering aid is called *care under fire,* which is very difficult to document. A number of different handheld devices require a stylus, but it's hard for anyone to document care at the same time they're applying a tourniquet. And, of course, that handheld doesn't necessarily have connectivity because the envisioned network doesn't really exist yet to allow data to flow ahead of the patient.

So what we envision in the future is PEO Soldier putting radios on every deployed soldier. So we're working with DARPA, TATRC, and a couple other organizations on a digital voice recorder that can suppress the background noise of helicopters, gunfire, and those sorts of things. And then we're also working on automatic speech recognition. To process spoken words, technology such as Dragon NaturallySpeaking version 10 is actually on our hardware configurations, available for providers to use in-theater. So the medic at the point of injury would speak into that microphone, which would create a .wav file that would then be transferred back to the combat support hospital. That spoken text would then be processed into recognizable text and inserted into the medical record.

MIL EMBEDDED: ***Any plans for smart devices?***

GEESEY: We've been experimenting with putting our applications on the iPad, iPhone, and other Droid-type devices, and we've been able to get our applications to work on those devices. The Army also has the *Army Training Network to Go* in which a soldier can download manuals onto their iPad or iPhone. MC4 training products are being piloted on the next version of ATN2GO. As far as training, we're leveraging Defense Connect Online to make our applications and systems training materials more readily available to soldiers using a 'Go mobile device.

MIL EMBEDDED: ***Wow, you folks really have a lot going on. But if you wanted to close with a mission statement for MC4, what would it be?***

GEESEY: We are working to improve the availability of medical information – whether it's patient health records, medical supply information, or situational awareness in command and control. Our goal is to seamlessly make that data available throughout the chain of care from the far-forward point of injury all the way back to our CONUS-based hospitals like Walter Reed Army Medical Center. By having lifelong electronic medical records available to all providers along that route, the results include better patient outcomes and availability of information to facilitate medical R&D. Also, once the soldier transitions, we want to make those records available to the Veterans' Administration so those soldiers receive the benefits to which they're entitled.

__Lieutenant Colonel William E. Geesey__ is the product manager for the Medical Communications for Combat Casualty Care (MC4) Product Management Office. With more than 26 years of active, reserve, and National Guard service, he has held a variety of management positions in medical logistics and healthcare administration.

"I'm excited about the processor road maps."

Q&A with Matt Tracewell, Executive Vice President of Tracewell Systems

Recent reviews by leading computer magazines tout Intel's Core processor family as spewing the most performance ever seen out of desktop and laptop computers. With next-gen Sandy Bridge versions of Core devices already in fab, 2011 promises to ratchet up performance even more. As well, Freescale's QorIQ devices are targeting cellular base stations, and NVIDIA and VIA are pouncing on GPU and low-power niches, respectively. It's no wonder systems integration expert Matt Tracewell has something to be thrilled about. Interviewed by Editor Chris A. Ciufo. Edited excerpts follow.

> Which technology are you most excited about today?

TRACEWELL: Really, I'm excited about all the [micro]processor road maps. So whether we're talking the IBM Power product line or the Intel Sandy Bridge-type technology, it's phenomenal what those mean to capability. They have so many applications and are so capable of transforming the market at a price point and a Size, Weight, and Power [SWaP] level that wasn't available previously; you just couldn't solve the problems before now. I'm all about enabling technology and lowering barriers to entry, which are key to numerous [military] programs.

GPUs are also revolutionary. So because of the processing horsepower, there are moves to things like OpenCL where you can common-out the programming language and have portability on different processing platforms. And so with that, you can benchmark against a number of best-of-breed types of either processors or techniques. And still today you get variations in performance: CPU/GPU combinations, or what's projected on Sandy Bridge, or what IBM is doing on the Power line. But as you move further and further into the future, that all kind of melds together.

> So which technologies are you excited about in the short-term future?

TRACEWELL: The coming-of-age of nano materials. Look at all the DARPA announcements and awards over the past four or five years and you will find maturing technologies that are now being used in small scale for very pinpoint-specific problems, whether it's a thermal issue, conductivity, or trying to make different types of composites with thermal, mechanical, and EMI shielding properties. And within that, what makes all of them work are what I'll call *nano projects* – pieces of technologies that have been matured enough that they're starting to be usable. And what I see happening is that as those first-out-of-the-gate usable products start to join together to create the next level of integration, the result is all kinds of functionality the industry hasn't seen before.

> Can you point to any particular DARPA program?

TRACEWELL: The ones right now that I have my hands on best are the materials ones. So you've got materials, and that ties in to sensors as well.

> What's missing from our technology future?

TRACEWELL: Energy. We're moving toward energy as "energy." And it really doesn't make any difference what domain it sits in. And you sense it and process it, and that ties back to this high-end processing capability we're discussing.

> Can you be more specific? That's a little bit too ephemeral.

TRACEWELL: Well, OK. So I've got all this capability and I'm making it mobile. The equipment is portable; it gets sent all over the place. We've been talking about fuel cells forever. We've been pushing battery technology. You've got some nano capability showing up in batteries.

You have processor power consumption coming down, but the issue still is: "In the future, I've got this great capability and I'd like to put it someplace and forget about it. I'd like it to run for two years on a charge, or on its stored energy, in a reasonable manner." How does that happen? We're still missing those pieces.

So there's work being done, but there's no shining hope out there for dramatic energy savings or harnessing the energy that's available all around us. There's work being done on RF energy-harvesting chips. So there are little pieces like that, but it's thousands of dollars still to do that. We're still waiting.

Matt Tracewell *is Executive Vice President, Tracewell Systems. He pioneered Tracewell's concept of Advanced Form Factor Engineering, a process where COTS technology and products are modified to meet the needs of the DoD and companies in aerospace, healthcare, telecom, and other industries. Matt holds four registered patents and is a graduate of Ohio State University. He can be reached at mtracewell@tracewell.com.*

Tracewell Systems
614-846-6175
www.tracewellsystems.com

Increasingly complex system designs need simulation, code generation tools

Interview with Dr. Jon Friedman, MathWorks

EDITOR'S NOTE

"I think automatic code generation, test vector generation, or property proving aren't about taking the human out of that process. They're about adding tools to the engineer's quiver." So said our interviewee, Jon Friedman from MathWorks. In this exclusive interview, editor Chris Ciufo gets a firsthand education in Model-Based Design and why it's becoming essential in air- and spacecraft systems. Edited excerpts follow.

What is Model-Based Design, and why is it important?

FRIEDMAN: Model-Based Design is the way the industry has been moving for a while, because sometimes (such as in the case of spacecraft), you can't really stick a prototype up in space and start tweaking it. And with aircraft, nobody wants to fly in an aircraft that you're changing the wings on as it flies. What's taken hold probably in the last decade or so is the notion that modeling is an important part of all the system aspects including the physical, mechanical, electrical, or hydraulics components as well as the embedded software components.

At the center of Model-Based Design is an executable model that simulates the overall system-level design. So engineers start with an executable specification, linked back to the base-level requirements. And a lot of times those requirements are captured in the natural language environment such as Word, Excel, DOORS, or some requirements tool.

Why use Model-Based Design and what are the perhaps less-obvious benefits?

FRIEDMAN: There are two main things: One is that testing can start at the model level. Another is [handling] functional requirements, both in terms of how the system should perform and things the system should *not* do. Returning to our spacecraft example, the solar panels should not deploy until the vehicle is on orbit. That type of logic can start to be tested in a simulation environment, and worst-case, Monte Carlo, robustness, and many other types of simulations can be used. So [basically it's] designing with simulation.

Where does automatic code generation fit in – or does it?

FRIEDMAN: Yes, it fits in. Since everything is in the model including the algorithm, the next step is to use automatic code generation technology. Automatic code generation technology has been around for probably about 10 or 15 years. And the idea there is that code can be generated from the physical models to do HIL [Hardware in the Loop] testing, or the algorithm code can be generated and used to do real-time testing. Then the same models can also be used for production code generation of the flight control or on a commercial aircraft for the avionics system. Because the models have links to the original requirements, the code that gets generated also contains those links. This capability is particularly important when engineers have to conform to DO-178B, where every line of code has to be traced back to a requirement. All that traceability is established automatically.

What about legacy code – does it work well in Model-Based Design environments?

FRIEDMAN: Yes. Many applications are built from both existing legacy code and automatically generated code. What engineers tend to do a lot is reuse existing designs and add or update parts of the design. For this work, they pull the legacy code into the modeling environment and create models of the new parts. They can then complete their design work in the modeling environment and code generated from the models. Next they can integrate the legacy code and generated code together and pull the application back into the modeling environment to make sure the integrated code still provides the same functionality that they designed and verified by performing model-in-the-loop testing. From the model-level testing, they will often move to hardware-in-the-loop testing in a real-time simulator to perform further verification.

What about testing with Model-Based Design?

FRIEDMAN: The testing and verification activity with Model-Based Design is paramount. Traditionally using a V or waterfall method, testing gets done at the end. Within Model-Based Design, when the executable specifications are created, requirements tests can be developed inside the modeling environment to ensure that requirements have been met. So engineers are testing continuously throughout the design process in the modeling environment. Then they can use those same tests at the code level by either linking to an IDE to enable PIL (Processor-in-the-Loop) testing – or engineers can pull code into the modeling environment and do software-in-the-loop testing.

Which MathWorks offerings provide Model-Based Design?

FRIEDMAN: Our platform product for Model-Based Design is called *Simulink*. It's a graphical environment, which uses block diagrams to allow engineers to build up models. The models can be dynamic models built from predefined blocks in Simulink. In addition to a time-based engine, Simulink has an acausal solver that allows engineers to create mechanical, electrical, or hydraulic linkages without having to first solve the closed form differential equations. This means the solver figures out where the equilibrium point is and then starts time moving forward. Simulink has a finite state engine in it so engineers can model logic. For example, a logic state might include how a spacecraft needs to behave differently when in ascent mode or in station-keeping mode. The algorithm focuses on mode change, not how long the spacecraft is in a particular mode.

There's also a discrete event engine inside of Simulink that allows queuing models such as communication protocols or air traffic control strategies to be constructed. The platform has these different engines all communicating with each other so the engine can sort of move time forward and backward in the simulation environment to let the logic engine, discrete event simulation engine, or acausal engine work in parallel.

Briefly remind our readers what MATLAB is, please.

FRIEDMAN: MATLAB is our technical computing environment. The best way to describe MATLAB is that it is a combination of a programming language and a powerful data analysis and visualization environment. Engineers will both build programs in MATLAB and deploy them. There's a MATLAB complier for deploying outside of MATLAB, and there's a lot of analysis that gets done inside MATLAB. And so engineers can use MATLAB to essentially pull data in, fit a model to it, test and analyze the model, and then sometimes deploy that model into Simulink as part of the modeling process.

Also, for verification tests, engineers can take all the data into MATLAB, write a script in MATLAB to analyze the data, and identify any anomalies or passing of the test. Engineers also write signal-processing algorithms sometimes inside of MATLAB because of its array processing capabilities.

Let's switch gears and zero in on your recent announcement, the real purpose for today's discussion.

FRIEDMAN: Sure. The announcement is about what the Swedish Space Corporation was able to achieve with the Prisma project, a civil mission. The important thing I would say about the Prisma project is that it's actually the second project. Their first project, another civil mission, was the SMART-1 satellite of the European Space Agency, which they had also developed using Model-Based Design. When they went about developing SMART-1, they saved a tremendous amount of time and they were very happy with the results. And so they reused a lot of the models from SMART-1 in the Prisma project because both systems are general-purpose geostationary platforms.

How did the company use Model-Based Design for the satellites, technically speaking?

FRIEDMAN: The Swedish Space Corporation reused about 70 percent of the attitude control models that they developed for the SMART system. The other thing that they did [was to] generate code not just for production flight control, but [for] for real-time simulation. And on this particular project, the engineers generated code from the plant models, not just from the controller model, and deployed that to the MathWorks real-time environment called xPC Target. They were then able to run real-time simulations of the spacecraft control and to verify that it met the overall requirements and real-time constraints.

Where did the compiler fit in with the satellite project?

FRIEDMAN: In this case, the engineers built a MATLAB application to display and analyze some of their satellite flight data, and they used the MATLAB compiler to deploy that application. So that data coming in off the satellite was sent through a compiled MATLAB application, and displayed the analysis results to the engineers. This particular project also had a degree of difficulty because it was autonomous formation flying, where there were multiple satellites that needed to fly in a particular formation. The controller allowed them to verify that capability.

How would you contrast this methodology and this customer specifically with how, say, JPL would have done it with one of the Rovers?

FRIEDMAN: JPL engineers did use [automatic code generation] for the Mars Rovers Model-Based Design. And also for one of their earlier satellites, they used automatic code generation. I would describe JPL and those engineers as the leading edge.

This all sounds great, but are there any drawbacks, perceived or otherwise, in automating so much of the design process?

FRIEDMAN: Somebody once asked me, "These tools seem to automate a lot of things. Are you trying to automate out the engineer?" Not at all. That's the furthest thing from our minds at MathWorks. It's about trying to automate the things that can be automated so engineers can be freed up to be creative and solve problems and not have to redo everything over and over again. Traditionally, an engineer would build a simulation model of an algorithm, then hand it off to another engineer who would write the code for that algorithm. In this scenario, there's potential for misunderstanding what was specified and there's potential that the code is wrong. I think automatic code generation, test vector generation, and property proving aren't about taking the human out of that process: They're about adding tools to the engineer's quiver.

***Dr. Jon Friedman** is the aerospace & defense and automotive industry marketing manager at MathWorks. Prior to joining MathWorks, he worked at Ford Motor Company and as an independent consultant on projects for Delphi, General Motors, Chrysler, and the U.S. Army Tank-automotive and Armaments Command (TACOM). He holds a B.S.E., M.S.E., and Ph.D. in Aerospace Engineering and a Masters in Business Administration, all from the University of Michigan. Contact him at jon.friedman@mathworks.com.*

MathWorks
508-647-7000
www.mathworks.com

InnoDisk®
Beyond your imagination

i-DIMM
INDUSTRIAL DRAM MODULE

*Fully support 0˚C~-40˚C, and 85˚C~95˚C operation temperature.
*Better signal quality via 30um gold connector
*ESD and steam proof protection by PCB surface protection.
*Support 7.8 us Average periodic refresh interval (Auto-Refresh Rate)
*Reduce 50% power consumption on 85˚C ~95˚C

-40˚C~+85˚C Wide Temp. Available

Low Profile Available

The total solution for Memory and Flash Storage

SSD Industrial Solid State Disk

2.5" SATA 20000
Capacities: 16GB~256GB
Read:230MB/s(max.)
Write:160MB/s(max.)

1.8" SATA 10000
Capacities: 8GB~64GB
Read:160MB/s(max.)
Write:140MB/s(max.)

EverGreen
Industrial Flash Storage

Capacity: 8GB-256GB
Read: 200MB/sec(max)
Write: 80MB/sec(max)

SATA Slim

SATA Slim J200
Capacities:2GB~32GB
Read: 230MB/sec. (max.)
Write: 180MB/sec. (max.)

CFast

Capacities:2GB~ 16GB
Read:100MB/sec.(max.)
Write: 50MB/sec.(max.)

8000 Series-EDC/iCF

-40˚C~+85˚C Wide Temp. Available

Vertical type 40pin
Capacities: 2GB~16GB

iCF 8000
Read: 80MB/s(max.)

mSATA

Capacities:2GB~16GB
Read: 120MB/sec. (max.)
Write: 100MB/sec. (max.)

SATADOM

-40˚C~+85˚C Wide Temp. Available

Low Profile / Vertical Type
Capacities: 128MB~8GB
Read: 24MB/s(max.)
Write: 14MB/s(max.)
Dimension: 0.81x1.54x0.31(inch)
Dimension: 0.90x1.28x0.31(inch)

www.innodisk.com

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

Case study: Secure FPGA technology enables UAV communications and control

By Jim Anderson

Single-chip cryptography enables a cost-effective implementation of a UAV command and control system in a single FPGA. Partial reconfiguration capabilities in the programmable IC add SWaP-C savings because a less-dense, lower-power FPGA can host the design.

Editor's note: Due to popular demand (and about twice as many articles and interviews as we could print in this edition alone), we share Part 1 of our special coverage on UAVs in this edition. Look for Part 2 in our June edition.

Over the past few years, the U.S. military and its allies have come to increasingly rely on Unmanned Aerial Vehicle (UAV) systems to carry out surveillance and combat missions around the world. Secure communication links are vital for UAV operation, both to control the UAV based on mission objectives and to reliably deliver actionable data to mission controllers on the ground. Encryption and decryption are inherent requirements, adding complexity and cost in the UAV electronics package. But with a single FPGA capable of meeting Type 1 cryptographic requirements, design teams can leverage reprogrammability and realize Size, Weight, Power, and Cost savings – referred to as *SWaP-C* savings. Xilinx and Advanced Communications Concepts, Inc. (ACCI) have demonstrated one such UAV communication and control system based on an FPGA.

The UAV application relies on a Single-Chip Crypto (SCC) design implemented in an FPGA to protect communications between the ground control stations and the UAV. The implementation completely safeguards the telemetry, video, and control data. The example system relies on the power of FPGA partial reconfiguration to provide algorithm swapping within a field-upgradable solution, all in a small product footprint.

Xilinx worked with the leading defense solution developers and key government agencies to develop an FPGA design flow and verification process that enables a single FPGA to meet Type 1 cryptographic requirements. The older method for meeting Type 1 cryptographic requirements employed two FPGAs – one to securely partition the unencrypted portions of the design. In the single-chip implementation, unused logic elements serve to implement the partitions.

The design flow isolates regions of the FPGA that handle red and black data and the encryption/decryption function (Figure 1). The red portions of the design deal with unencrypted data and must be isolated from the portions that deal with encrypted data. The SCC sits functionally between the red and black sides. The UAV example described here is based on a Virtex-5 FPGA using the SCC technology.

The UAV demo

At conferences such as MILCOM, Xilinx and ACCI have demonstrated an FPGA-equipped UAV providing a real-time encrypted flow of control, telemetry, and video data between the UAV and the ruggedized, laptop computer-based ground-control stations (Figure 2). The live-fly version has flown at events such as Air Force Joint Forcible Entry Exercise (JFEX) and the SOCOM/NPS Tactical Network Topology (TNT) exercises. They are being evaluated for use in various planes and systems, including UAVs.

The UAV command and control system uses a Virtex-5 FPGA with an integrated PowerPC processor. The system requires little more than the FPGA, MEMS accelerometers, and a physical layer for the wireless communications link. In developing the system, ACCI started with the SCC design flow and Xilinx's

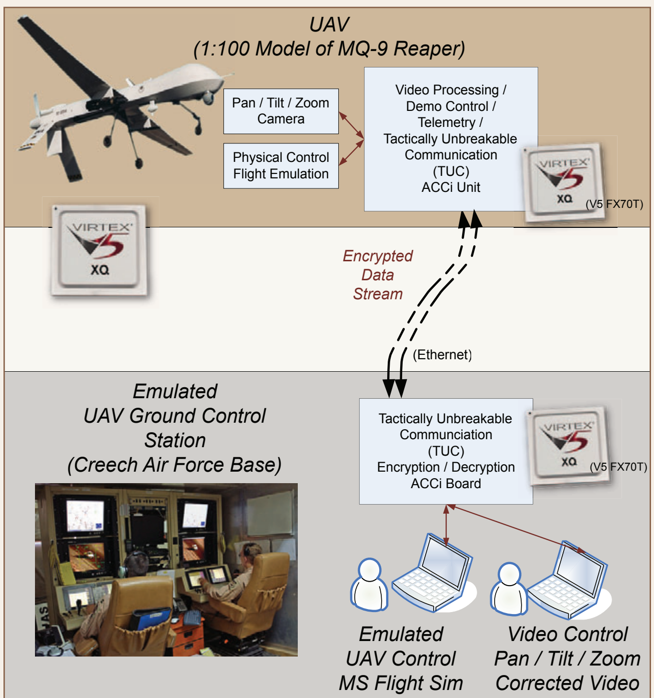

Figure 1 | Type 1 cryptography demands isolation between encrypted (black) and unencrypted (red) regions of the UAV system.

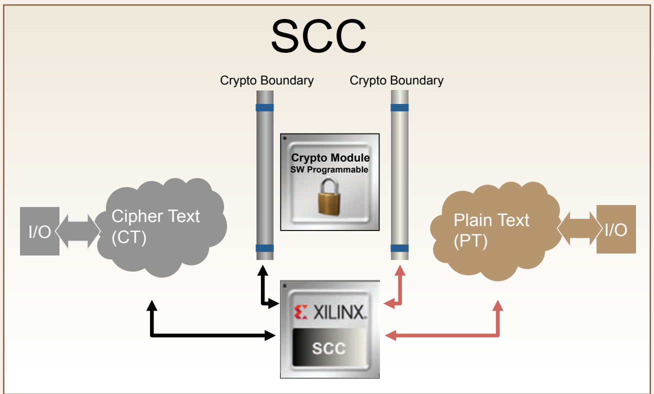

Figure 2 | The ACCI UAV demonstrates the power of single-chip cryptography and partial reconfiguration.

information assurance techniques, and added a secure communications layer called *Tactically Unbreakable Security Communications* or *TUCNet*. TUCNet can encrypt any digital data stream. For example, it can handle video, telemetry, control, or even voice data packets.

ACCI doesn't reveal TUCNet technical details for both security and competitive reasons. But more broadly, the company relied on features such as protocol hopping and encryption-scheme hopping to deliver a network layer that is secure over any type of wired or wireless network.

To meet Type 1 requirements ACCI had to isolate each of the regions of the FPGA according to defense agency specifications. Using Xilinx's SCC methodology along with the Isolation Verification Tool (IVT), ACCI was able to implement this solution and provide necessary documentation validating the isolation.

Most notably, ACCI implemented the Type 1 requirements in a single FPGA. Prior to Xilinx's work with government agencies and the validation of Type 1 cryptographic capabilities, a design would have used multiple ICs or subsystems to isolate the red and black data and the algorithms that operate on each. The SCC technology simplified the system implementation, resulting in SWaP-C savings. At a minimum, the SCC technology eliminated one FPGA from the implementation, halving the PC-board real estate needed to host the design. Power and cost aren't halved, because the two-chip implementation might have used slightly less dense FPGAs, but the savings are significant and even the weight is reduced by a small amount.

Compounding the SWaP-C advantages

While meeting all security requirements for Type 1 cryptographic certification, ACCI's FPGA algorithms and processing implementation compounded the SWaP-C advantages in the UAV application by using *dynamic partial reconfiguration*. The work Xilinx did on Type 1 cryptographic systems proved the ability to maintain proper isolation of red and black data even when reconfiguring a portion of the FPGA on the fly. With dynamic partial reconfiguration, the FPGA does not have to be big enough to hold all of the processing algorithms. It only needs to be big enough to simultaneously hold the single largest data-processing algorithm, the main control algorithm, and the SCC implementation.

ACCI utilized unique dynamic partial reconfiguration to add to the capabilities of the UAV control and communication system, and to minimize the SWaP-C burden of doing so. The system has a proprietary Hardware Operating System (HardwareOS) that is static in the FPGA. HardwareOS provides system resource allocation and system services functions that an OS would provide in a traditional software based system architecture.

The UAV system relies on a library of application or algorithm accelerators developed by ACCI. The TUC-enabled algorithmic accelerators enable, besides security functions, on-UAV, real-time manipulation of telemetry and video data streams and data transcoding functions. For example, if the UAV is in a banked turn, the video frame is horizontally distorted by both the pitch and roll angles of both the UAV frame and the camera pan and tilt settings. This problem was solved by dynamically loading and running an algorithm to "counter-rotate" video frames in real time into the proper orientation.

The TUC system also transcodes the digital video from RS-170 format to MPEG-2 and H.264 formats, among others. The system then combines the transcoded video with the telemetry from the autopilot, and other onboard sensors, into an MPEG transport stream that correctly emulates a Predator data download format. This allows the UAV data to be utilized by any system that currently handles Predator formatted data streams. And all of the data streams are encrypted for ground transmission.

The system can load every data packet transmitted to the UAV or every packet of captured telemetry or video data into static Block RAM (BRAM) on the FPGA, and then dynamically apply any desired sequence of algorithms to each packet as required. With TUC hardware acceleration, the entire frame processing of video stabilization, horizontal correction, Predator format transcoding, transport stream packaging, and encryption is done in less than 12 milliseconds. At 30 frames per second from the camera, 33 milliseconds are available between frames, thus allowing ample processing resources for future planned enhancements, such as automated target tracking and direct autopilot control.

A closer look: Dynamic partial reconfiguration

While using SCC flow to help maintain Type 1 requirements, the real advantage of using dynamic partial reconfiguration is apparent: The system can reconfigure the FPGA more than 100,000 times per second. Moreover, the data flow and parallel processing inherent in the FPGA fabric minimize latency and


VISIT US AT
ESC SILICON VALLEY
BOOTH #1346
Proven in Over
38 Million Field Hours
Designed, Manufactured
and Certified in the US
Wide Temperature
Range -40° C to 75° C
ISO 9001 Certified
Modular Design Ensures
Rapid Customization
RE-DEFINING
RUGGED
GPS
LAN
Option
Power
ANT1
LNK / ACT
PWR
Display
Multi I/O
COM/CAN
ANT2
OCTAGON SYSTEMS
RMB-C2
RUGGED COMPUTER
SOLUTIONS FOR
EXTREME APPLICATIONS
OCTAGON SYSTEMS
www.octagonsystems.com 303-430-1500 sales@octagonsystems.com

enable real-time manipulation to optimize gathered data for transmission to the ground station. The ACCI system encrypts the preprocessed data and transmits the secured data to the ground control station. The laptops used in the demo decrypt the data and present it to the user.

The partial reconfiguration capability enabled additional savings by allowing ACCI to utilize the smallest, ruggedized, defense-grade Virtex-5 family member that integrates a PowerPC. The XQ5VFX70T device chosen includes 11,200 Configurable Logic Blocks (CLBs) and a single PowerPC core. Without partial reconfiguration, the design requires a larger FPGA that would cost more and use more power. As an example, this can mean a 5x savings in just static quiescent power consumption between the smaller product and the next larger product in the Virtex-5Q family.

ACCI and Xilinx are working on a new version of the UAV demonstration system that will leverage the defense-grade Virtex-6 family and further compound the SWaP-C benefits. Virtex-6 FPGAs consume 50 percent less power than Virtex-5 FPGAs with a similar number of CFBs. Moreover, the Virtex-6 family is manufactured in a 45 nm process technology, versus a 65 nm process for the Virtex-5 family. Rather than requiring an FPGA with an integrated PowerPC hard core, the new version of the UAV system will yield further savings through the use of a soft core MicroBlaze processor.

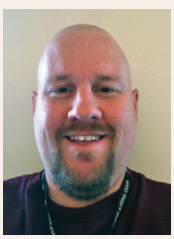

Jim Anderson *is a Senior Staff Defense Application and Architecture Engineer at Xilinx. Prior to joining Xilinx, Jim was an engineer at Philips Semiconductors and Sandia National Laboratories. He holds a BSEE from the University of New Mexico. Contact him at Jim.Anderson@xilinx.com.* *The author thanks Jonathan Ellis, CEO of Advanced Communications Concepts, Inc. of Austin, TX for his assistance with this article.*

Xilinx
408-559-7778
www.xilinx.com

U.S. Navy photo by Mass Communication Specialist 1st Class Chris Fahey

Rapid deployment of scalable, dense ISR payloads for UAVs

By David Pointer

In light of increasing demands for advanced rapid prototyping and deployment of ISR applications on UAV platforms, systems engineers would be wise to examine reconfigurable architectures and flexible software tools that match the job's qualifications.

Reconfigurable computing systems are an excellent and increasingly popular choice to provide heterogeneous Digital Signal Processing (DSP) compute solutions for Intelligence, Surveillance, and Reconnaissance (ISR) applications on Unmanned Aerial Vehicles (UAVs). Reconfigurable systems can be created with lower Size, Weight, and Power (SWaP) and higher computational density for an ISR unmanned airframe application compared to other types of systems.

In addition, programming environments are now available using ANSI standard C or FORTRAN, and enable a programmer to extract all possible compute performance from the hardware. The use of standard programming languages greatly reduces the learning curve, enabling much quicker application deployment when compared to programming environments that utilize a proprietary non-standard language or purely hardware languages. Thus, UAV ISR systems developers should consider reconfigurable computing architectures and flexible software tools imperative in their designs.

Reconfigurable computing for UAV payloads

When considering rapid deployment of UAV payloads, reconfigurable hardware is key. At the core of a reconfigurable system is the FPGA integrated circuit. This device may be explicitly programmed to execute application-specific algorithms, and yields very high computational efficiencies relative to a general-purpose device such as a CPU or a General Purpose Graphics Processor Unit (GPGPU). This computational efficiency, in turn, generates high performance per watt for applications executing on reconfigurable systems, which then enables the creation of computationally dense/low SWaP designs for UAV-based ISR applications.

Table 1 compares the execution of a single precision, matrix multiplication benchmark on a CPU-only system, a CPU paired with a GPGPU, and a CPU paired with

	2 CPUs	CPU+GPGPU	CPU+FPGA
Matrix multiplier	2 quad-core CPUs	CPU+NVIDIA Tesla S870	CPU+XD2000i+ EP3S260
Floating-point precision	Single	Single	Single
Sustained GFLOPS	140	140	157
Power (watts)	~150	~245	~110
GFLOPS per watt	0.9	0.6	1.5

Table 1 I Matrix multiplication performance/power ratio. Source: Altera Corporation, "FPGA Coprocessing Evolution: Sustained Performance Approaches Peak Performance," WP-01031-1.1, June 2009 version 1.1.

an FPGA. The reconfigurable system's performance per watt is 1.7 times better than the CPU-only system, and 2.5 times better than the CPU/GPGPU system. Note the total power consumption for the CPU/FPGA combination is the lowest power consumption.

A major factor in rapid UAV payload deployment is the availability of standard programming tools that are closely coupled with the reconfigurable computer hardware. The piecemeal method of integrating compilers, software tools, FPGA boards, and CPU boards from different vendors slows down deployment. If the tools and hardware have different suppliers, a system runtime environment must be created to unify the system before meaningful application work may begin.

Once real application development finally begins, one vendor's compiler macro libraries (assuming they do have libraries) will not be optimized for another vendor's FPGA boards, and so the application development must either experience a reduction in the performance specification or a schedule extension. These and other inherent difficulties in the piecemeal method of UAV payload development conspire to decouple the word "rapid" from "deployment," usually in an unanticipated manner, and always during application development. Thus, a complete, well-integrated software and hardware package from one supplier is typically the best route for rapidly deploying a UAV payload.

Scalable systems and parallel programming

A modular, scalable system lends itself well to code reuse, which also accelerates deployment for UAV payloads. Modularity in software design allows proven code to be reused across several ISR applications, while hardware modularity supports easy scaling of an ISR application according to mission parameters and a UAV airframe's SWaP requirements. Most heterogeneous systems today use one form or another of PCIe to give coprocessors access to system memory through the CPU. However, the effective scalability of PCIe is limited by its blocking "many-to-one" architecture (Figure 1). PCIe coprocessors access data and communicate with each other only through system memory on the other side of the CPU. While PCIe does provide point-to-point connectivity

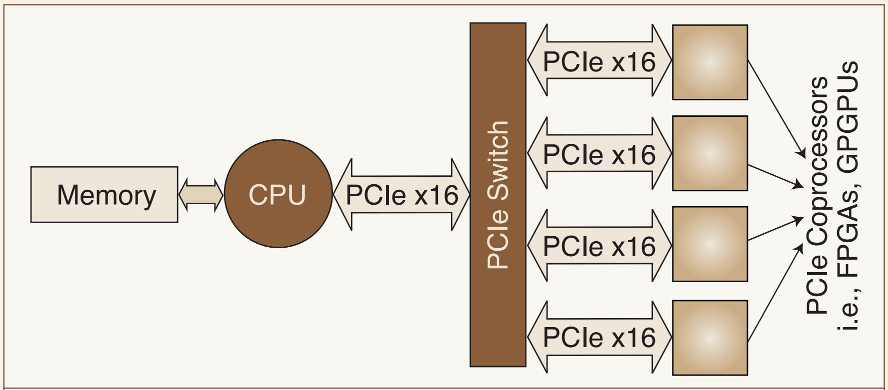

Figure 1 | Typical PCIe-based coprocessor architecture with limited scalability because of its blocking "many-to-one" point-to-point architecture

through the PCIe switch, one point is always CPU memory and the other point is always a PCIe device. Truly usable scaling requires distributed memory and a switched network with nonblocking, "many-to-many" connectivity to or from any module on the switch.

A modular hardware system that scales well is only a good start. Software tools must provide the programmer with intuitive or automatic access to the dense computational efficiency in the reconfigurable system. UAV payload development and deployment are impacted if a programmer has to drill down into a system's architecture to find and develop solutions to meet an application's performance requirements. There are standard computer languages available for programming the CPU and FPGA in a reconfigurable system, but simply using C or FORTRAN alone will not achieve ISR application performance requirements. C and FORTRAN are serial programming languages traditionally used for CPUs, where instructions are executed serially, one instruction at a time. Performance in a reconfigurable system is achieved by parallel programming: multiple streams of program instructions acting upon multiple streams of data at the same time.

Fortunately, the scientific supercomputing community has already developed parallelization techniques for C and FORTRAN, many of which have been adopted by some reconfigurable system compilers. One method of programmatically specifying parallelism is the OpenMP parallel section pragma statement. On a traditional large cluster of microprocessors, the code blocks enclosed by the parallel section pragma may be executed in parallel on the CPUs. On reconfigurable systems, the code blocks specified by OpenMP-style pragmas are instantiated in the FPGAs so that these "hardware code blocks" execute in parallel. Another method is data streaming, where a series of calculations is overlapped in time; for example, a computational block may start executing when the first results of the previous block are received instead of waiting for all results from the previous computational block to be produced before starting.

"There are standard computer languages available for programming the CPU and FPGA in a reconfigurable system, but simply using C or FORTRAN alone will not achieve ISR application performance requirements."

Along with parallelization techniques borrowed from scientific supercomputing, most reconfigurable system compilers perform automatic loop pipelining for execution performance. In addition, reconfigurable system compilers automatically create in FPGA hardware all

Stroke of Genius

Now Shipping the NEW Cobalt Series of Virtex-6 Modules!

Unleash your inner genius with Pentek's new Cobalt™ Virtex-6 XMC modules. Built to optimize high-performance, wideband, communications, SIGINT, radar, beamforming and telemetry applications, the high-speed modules include:

- A/Ds with sampling ranges from 10 MHz to 1 GHz
- 800 MHz D/As, DUCs & Multiband DDCs
- Dedicated memory for each I/O stream
- Intelligent chaining DMA engines
- Secondary serial gigabit interface
- Multichannel, multiboard synchronization
- ReadyFlow® Board Support Libraries
- GateFlow® FPGA Design Kit & Installed IP
- VPX, XMC, PCI, PCIe, cPCI, VME/VXS, rugged XMC

With extensive resources and new enhancements in DSP engines, block RAM, logic and clocking technology, the Virtex-6 family represents the industry's most advanced FPGAs.

Now Available in VPX!

COTS & Rugged Versions

Call 201-818-5900 or go to www.pentek.com/go/mescobalt for your FREE *Putting FPGAs to Work in Software Radio Handbook*, technical datasheets and price quotations.

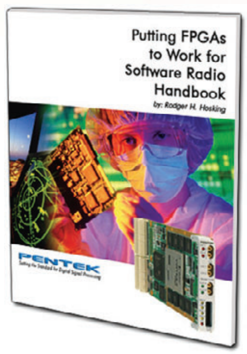

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

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

arithmetic operations in a program loop, all of which execute in parallel. Contrast this with a microprocessor compiler, which is limited by the number of arithmetic compute elements available in the CPU's design. The real issue here is the effectiveness of a given set of software tools for a given system. But software tools do exist to provide programmers of compute-intensive ISR applications access to the potential performance in reconfigurable hardware.

Reconfigurability facilitates rapid deployment, SWaP

Reconfigurable systems and tools are available today for rapid development and deployment of ISR applications. The dense computational nature of reconfigurable systems makes them an ideal choice for solutions where Size, Weight, and Power consumption matters. Software tools closely coupled with the hardware allow software programmers to quickly achieve high performance in a low-SWaP processor payload.

Accordingly, SRC Computers provides modular, scalable, reconfigurable low-SWaP systems with software tools and libraries, using ANSI standard languages for rapid deployment of compute-intensive ISR applications for UAVs. SRC Computers has developed a high-bandwidth, low-latency network switch that provides the necessary non-blocking "many-to-many" module, in addition to data stream constructs for intermodule communication, OpenMP-style pragmas for code block parallelization, and automatic loop pipelining for instruction-level parallelization.

David Pointer is Director of System Applications for SRC Computers, LLC (Colorado Springs, CO). His career spans 30 years in industry and academia, including Hewlett-Packard Laboratories and the National Center for Supercomputing Applications. He holds an M.S. in Electrical Engineering and an M.S. in Computer Science. He may be contacted at dpointer@srccomputers.com.

SRC Computers, LLC
719-262-0213
www.srccomputers.com

Autonomous ISR software reduces operator workload

An interview with executives from Kitware

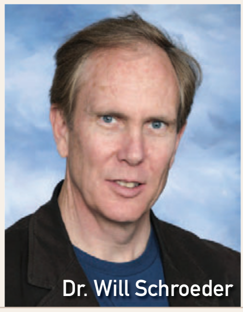

Dr. Will Schroeder

Dr. Anthony Hoogs

EDITOR'S NOTE

European citizens are familiar with the ubiquity of cameras everywhere, monitoring high-traffic public places and roadways. In most cases, video feeds back to operations centers where operators look for specific events. On today's digital battlefield, ISR sensor data and video from UAS platforms, satellites, aircraft, and ground forces feed into command centers and TOCs where military analysts perform similar functions to their civilian counterparts, but with deadly seriousness. One missed event could mean death by IED to a warfighter. Yet how can a human watch and correctly identify threats within hundreds of hours of reconnaissance data? Open source software company Kitware, with some help from DARPA, has the answer. I sat down with executives from Kitware to learn about leading-edge autonomous software. Edited excerpts follow.
– Chris A. Ciufo, Editor

Let's start by talking about Kitware.

SCHROEDER: We've been around since 1998. We're an open source company, and we develop technologies in the scientific computing arena. In most cases, we give that technology away in the form of open source code available from our website.

How do you license it, and how do you make any money?

SCHROEDER: It's a very permissive BSD licensing agreement, which allows people to use it without necessarily giving anything back to us. We make money as a service company: By giving this stuff away, people start using it and they need help using it, integrating it into their systems.

Does your company maintain any intellectual property?

SCHROEDER: It's minimal. The majority of our "intellectual property" is customer relationships and customer contact lists. We have a small number of patents and some patents pending. There are proprietary software systems we've developed for our customers, but we can't necessarily sell those. We have a couple proprietary products, too. The value in our company is the employees and the knowledge that travels the open source world.

Which software technology areas do you focus on?

SCHROEDER: We have five key areas. We're a software company, so we focus a lot on software quality and software process. Around that software core, we have four technology areas. One is computer vision, which Anthony [Hoogs] can talk more about.

HOOGS: Sure. Most of our funding comes from DARPA, which is mostly interested in aerial video. So primarily what we do is have a computer examine the video and determine whether the video contains anything of interest to a video analyst. The overall goal is really to make this enormous amount of video being collected by the military indexable and accessible so that you can know what's in it without having humans look at it – until there's an event that requires operator intervention.

Let's set the stage. Are these super high-res cameras? Are there ever image quality issues?

HOOGS: The military uses consumer cameras that are high resolution and create lots of data. Sometimes image quality issues can arise, but they are not because of poor cameras; rather, they are because the scene is miles away from the UAV, such that atmospheric disturbances and imaging conditions have an effect. Additionally, if the camera's on a moving sensor or platform, … you have issues of parallax: 3D objects in the scene like a building rising above the ground plane with apparent motion.

So the intention is to make this poor-quality video into high-quality video while discerning people and events.

HOOGS: Right. [A computer-flagged event] might be someone implanting IEDs; the computer flags things like that. For example, there might be a Predator sweeping along a road because the analyst is following a vehicle of interest. The camera might sweep right over somebody digging a hole off the side of the road. If the analyst is watching that vehicle, he may not even notice that person digging a hole. However, the computer doesn't get distracted by watching the vehicle versus something on the side of the road.

Is a second or two of video enough to save that portion of the image?

HOOGS: The amount of time required to recognize a certain type of event or action depends on the event. It does vary quite a bit, but typically it's just a few seconds. We don't have any kind of control feedback in the systems now, but that's certainly conceivable.

How does the software report the flagged events?

HOOGS: The core of the system is the ability to catalog everything happening in the video. What you do with that information can vary quite a bit. If you're getting a video stream in real time, then the analyst can set up a standing alert that says, "Tell me whenever you see anything like this." Maybe it's digging, maybe it's somebody loading a vehicle. The analyst can program their own set of alerts that can be at a given location; they can be contextual, in this type of data or at this time of day. Then whenever any of the computed information matches the alert criteria, the analyst is notified.

OK back to you, Will, now that we've delved into Kitware's 1) software infrastructure and 2) computer vision, what were the other three Kitware core areas?

SCHROEDER: One of the other three areas is 3) scientific visualization. When I say *visualization,* most people think about 2D graphics. But scientific visualization involves large-scale 3D, 4D [time] graphics, and visualization where we take the output of an MRI scan, or an engineering simulation, or some large nuclear physics or oil and gas simulation or analysis.

And then our last two core areas are 4) medical imaging and 5) data management (Figure 1). Relative to scientific data management, we found that a lot of our customers have large data sets such as computer vision video streams or large biomedical MRI, CT, or confocal microscopy, and so on. And they're not doing a very good job of managing these very large and complicated things. Traditional database methods don't really manage scientific data very well, so we help them with data management, viewing, processing, and so on.

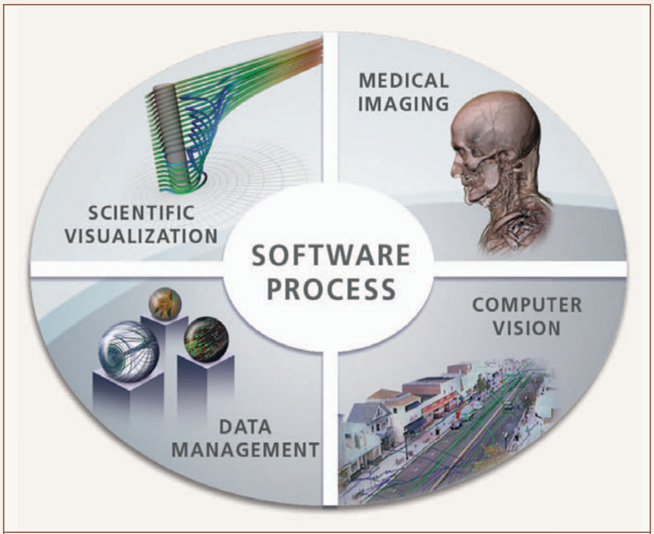

Figure 1 | Kitware's key technology focal points

Scientific data management doesn't necessarily have to be imagery data though, correct?

SCHROEDER: That's right. When we say *scientific visualization,* that actually is producing video or images, but the input to that process is large data sets that might come from a supercomputing simulation. So the data actually can be a range of things, from input decks for simulation codes all the way to the output, which might be graphics. The sensors are becoming more accurate and resolution is increasing very quickly. So now the data sizes are no longer 640 x 480 video streams. We're talking HDTV multiplied by 100. And confocal microscopy or electron microscopy are down to 5 nm in electron microscopy, which means you end up with images that are like 100,000^2 in depth or it might be 40,000 of these 100,000^2 images.

Do your algorithms and software work with sensor images such as radar?

HOOGS: Often it's less rich in information, which is part of the issue. So some of the algorithms do work because we develop them to mostly work off trajectory-level information. As long as you have a track you can reason about that track. Is it starting or stopping? Are two vehicles coming together, and so on. And that is agnostic to how the tracks were created. So if you want to create tracks from radar, GMTI, or even radar imagery, then, in principle, it would work.

What about your software itself? Tell me about it.

HOOGS: It's mostly C++, and our core framework is C++. But we work with a lot of universities. They work in MATLAB and sometimes Java, but the ability to integrate across languages is increasingly mature. So on our [Video and Imagery Retrieval and Analysis Toolkit] VIRAT program, which is a technology we're talking about here, we define an API that allows MATLAB modules to be dropped into the C++ system. And most of it runs on Linux and Windows, too.

VIRAT is source code. Are the MATLAB models also part of the source code?

HOOGS: No, that's just the research software version. A deployed version wouldn't have MATLAB.

VIRAT is a series of algorithmic modules. It's generally pipeline software that's typically at least six or eight processing stages, starting with video pixels and going through a bunch of different things, ending up in a database. So we take all this content in the video that we're computing and we put it in a database and index it. We had to develop special customized indexing software because these descriptors and the approximate matching functions needed to look them up on are not really suitable for typical relational database systems.

How would you characterize your database?

HOOGS: It's really a set of databases. We have different algorithms to describe the content in the video, and each has a set of what you might think of as tables in the database. (It doesn't actually map that way, but that's a good approximation.) A typical relational database has these structured fields. A given table might have 5 fields or 10 fields. And a field you can think of as a dimension. A field might contain an integer, or a string or something. And we have some descriptors, which are really mathematical ways of representing a piece of the videos, like "keep track of an object and how its appearance changes over time." Some of these descriptors end up needing thousands of fields, thousands of floating-point numbers, to represent them on every frame of the video, or for every track.

Is VIRAT open source?

HOOGS: At this point, none of the code for the VIRAT system is open source. Many subcontractors and universities have contributed to this, but none of the companies are open source companies like Kitware. DARPA might not necessarily want it to be open source for security reasons. However, some components are good candidates to become open source, and we hope to get approval for them at some point.

What about your company's products is open source then?

HOOGS: Well, in general, Kitware is built on an open source foundation. There's the Visualization Toolkit, the Insight Toolkit, these big C++ systems for visualization, medical image analysis, and scientific data visualization and management. Some of those are used on our computer vision programs, and through those we are slurping in a lot of open source. We're also developing elements for those toolkits, which will go back into the open source parts as we get approval. But the core vision stuff right now is not out in those open source toolkits.

Tell me about the processing requirements to execute the code you're talking about.

HOOGS: It's all on desktop PCs. Typically we like x86-based quad cores, or at least two or three cores. Our system can use as many cores as are thrown in its pipeline.

What are Kitware's future technologies?

SCHROEDER: One thing we're starting to put more work into is computational chemistry, which is becoming extremely important in the DoD.

Another area we've been growing is *informatics,* also known as *information visualization,* which is an extremely important field for homeland security issues. And so instead of looking at data that's spatial-temporal – like a CT scan or a particle physics simulation – you're looking at data that's not related to space and time. And we're working on automatic analysis of wide-area video, mostly on a DARPA program called *PerSEAS.* On PerSEAS we're developing algorithms to automatically detect threats and insurgent activities in city-wide video.

Dr. Will Schroeder is President, CEO, and cofounder of Kitware, Inc. He has an M.S. in Applied Mathematics, and a Ph.D. in Mathematics from Rensselaer Polytechnic Institute. He can be contacted at Will.schroeder@kitware.com.

Dr. Anthony Hoogs is the Director of Computer Vision at Kitware. He holds a Ph.D. in Computer and Information Science from the University of Pennsylvania and an M.S. from the University of Illinois at Urbana-Champaign. He can be contacted at Anthony.hoogs@kitware.com.

Kitware
518-371-3971 • www.kitware.com

Editor's Choice Products

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

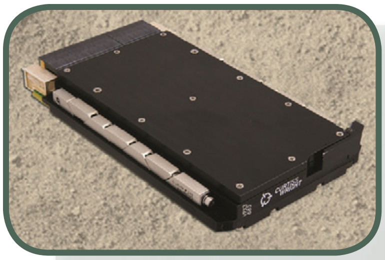

VPX router gets dolled up for FIPS crypto cert

This kind of certification is the stuff dreams are made of … well, if you're a marketing guy for a VPX company, that is. Curtiss-Wright Controls Embedded Computing's VPX3-685 secure router is undergoing FIPS 140-2 Level 2 cryptographic validation at NIST under the Cryptographic Module Validation Program (CMVP). This is a whole lot of alphabet soup jargon to say that once certified, the single-slot VPX module will ready for deployment in the most information sensitive data communications systems. The certification is the equivalent of winning an Oscar, if you're a crypto kinda gal. But this certification is accepted by the U.S., Canada, and many other NATO countries.

"Authentication," "integrity," and "confidentiality" are the watchwords for data passing through this module. Corsec Security, Inc., is helping CWCEC attain the certification. The module itself can be configured with up to 20 GbE interfaces and up to two 10 GbE ports for switch-to-switch comms and backbones. Like any router, it thwarts IP spoofing, Denial of Service (DoS), and Trojan horses. There's VLAN and VPN (IPsec/PPTP/L2TP) support, a firewall, NAT routing for IPv4, access control lists, and IPv6 with IPsec tunneling. The list of additional advanced router features is too long to mention here, but you'd think this thing was a full-blown Cisco rack … yet it's "only" a 3U VPX module. If you're running a red/black network on the battlefield, this router's worth a close look.

Curtiss-Wright Controls Embedded Computing • www.cwcembedded.com • www.mil-embedded.com/p47648

VMware-certified AdvancedTCA node board

AdvancedTCA is making inroads into moderate environment deployed military racks. The promise of exceptional horsepower and scads of I/O in a small 1U-sized space makes the interoperable PICMG 3.0 (AdvancedTCA, Option 9) standard quite appealing. But the true promise of AdvancedTCA servers and processor nodes lies in their virtualization potential, as exemplified by Diversified Technology's ATC6239. This AMD Opteron Socket F board is VMware compatible and certified for VMware's ESX 4.1 and ESXi 4.1 for virtualization. This means the DoD user is assured that the industry's most popular virtualization software's gonna run on this baby – no questions asked.

Virtualization takes advantage of excess CPU and I/O resources and makes this AdvancedTCA board appear as multiple processing resources for applications including: LTE/4G wireless and other wireline networks, WiMAX, IMS applications, IPTV, radio network controlling, security and traffic handling, computational clustering, and video/audio transcoding and encoding. Four different kinds of Opterons are supported with a HyperTransport subsystem. Additionally, there's up to 32 GB of ECC DDRII plus some onboard CompactFlash, and a Broadcom HT2100/HT1000 acts as traffic cop. I/O includes 1 GbE (front) and dual 1 GbE with dual 10 GbE (back), plus 2 USB, SAS/SATA (both to the AMC and RTM), plus video and three more USB to the RTM. And DTI has even tested this card with Windows Server 2008 R2 Hyper-V and Xen virtualization.

Diversified Technology • www.dtims.com • www.mil-embedded.com/p47649

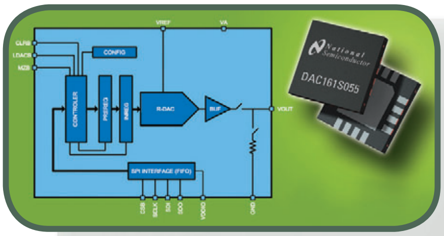

National's new DAC: Bit on bits, low on power

A DAC's a DAC, right? Well, maybe not. National Semiconductor's latest 16-bit digital-to-analog converter, the DAC161S055, has some interesting features up its sleeve, er, SPI. This precision DAC is available over a wide -40 °C to +105 °C temperature range, so it's ideal for high-rel applications. More importantly, the four-wire double-buffer SPI interface operates at 20 MHz and can be daisy-chained so multiple DACs in series can be connected via a single SPI I/F. This is but one aspect of the device's flexibility that might save the designer some board space.

Designed for industrial I/O, automated test equipment, and other data acquisition, the fast, buffered output will settle in only 5 microseconds. As well, DC specs show off the device's precision: +1LSB INL (max) and +1LSB DNL (max). The DAC161S055 has selectable power up to either 0 V or midscale, and the rail-to-rail buffered output has a very low noise of 120 nv/√Hz. There's read-back on the SPI, an asynchronous load capability, reset pins, and a wide reference voltage of +2.5 V to VA. The analog supply is 2.7 to 5.25 V and supports digital interfaces down to 1.7 V. Oh, and the best part? The little 16-pinner sips only 5.5 mW at 5.25 V (max).

National Semiconductor • www.national.com • www.mil-embedded.com/p46948

Rugged digital touch screens, from 10 to 19 inches

We doubt that VIA, the number three x86 processor vendor, set out to supply rugged LCDs to the military. But these ruggedized and open-frame LCD assemblies are solidly designed for kiosks, POS terminals, and other bright sunlight applications so they ought to do well in military systems … an Abrams here, an AC-130 there. Available to operate over a -10 °C to +70 °C "harsh industrial environment," the Vision Intelligence Display assemblies are also water resistant to NEMA IP65.

The SXGA resolution screens (1,280 x 1,024) are available in 7-, 10.4-, 12.1-, 15-, and 19-inch TFT sizes, with bright 600 cd/m2 LED backlights that consume up to 30 percent less power compared to traditional CFL backlights. As well, they boast 178-degree wide viewing angles and a five-wire USSB-based resistive touch interface, which VIA calls "excellently responsive." There's an available onscreen display panel plus pluggable options for VIA's own EITX-3001 and AMOS 3001 boxed PCs to feed data to and from the panel. We think there's a growing market for pre-packed, self-contained panel PCs, and this collection from VIA is a great place to start your search.

VIA Technologies • www.via.com.tw • www.mil-embedded.com/p47651

Continued on page 44

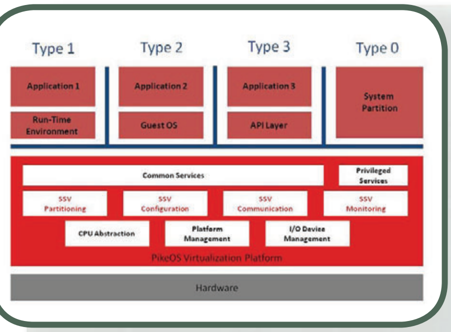

Android developers benefit from safe RTOS PikeOS

So, have you noticed the preponderance of Android on *everything* and *everywhere*? Yep – even in rugged applications, my friends. So those zany developers at SYSGO have decided to toss Android developers a lifeline by letting it run as a guest OS in version 3.2 of PikeOS. SYSGO will develop a new *Personality* executive environment for Android within the virtualized safety and/or security constraints on the same hardware host device.

The Android Personality will work on separate sets of resources within the same machine, with resource separation enforced by PikeOS's microkernel. This microkernel is lightweight and designed to be used in cost-sensitive, resource-constrained devices in both single- and multi-core environments. Typical processor architectures include: x86, PowerPC, QorIQ, MIPS, ARM, SPARC, and SH. Both AMP and SMP architectures are supported, and the RTOS is certifiable to safety standards including DO-178B, IEC 61508, and EN 50128.

SYSGO • www.sysgo.com • www.mil-embedded.com/p47652

Double-wide AMC plus double cores completes MicroTCA equation

Crunching math and computing algorithms forms the basis of many C4ISR applications. The sensors reel in the data, while the processors slice and dice the bytes into a reconnaissance picture. Adding maximum horsepower in the minimum available space is the norm, and Kontron's AM5020 Advanced Mezzanine Card (AMC) is one way of leveraging telecom processing in military applications. Based on an Intel dual-core Core i7 CPU, the AMC bolts onto a base card (not unlike the COM Express concept for which Kontron is famous) to form a MicroTCA system. MicroTCA, you might recall, has gained some favor in military applications because it can be ruggedized, leverages COTS modules, and eschews the backplane concept found in VME and VPX LRUs.

With memory controller, PCIe Gen2 I/O, and Intel HD graphics, the processor couples tightly with an Intel QM 57 platform controller hub. This module, with a MicroTCA.1 faceplate and PICMG AMC.1/.2/.3 sub-spec compatibility, makes a complete system. There are 8 PCIe lanes, 4 GbE interfaces, and 6 SATA II channels. The AM5020 can be equipped with an onboard 2.5-inch HDD or up to 32 GB of securely fastened SATA flash. Two USB interfaces, a DVI-D port, two more GbE, and a serial port route to the front panel. And since this is an AMC designed to work with AdvancedTCA, that means IPMI, shelf management, Module Management Controller (MMC), and other high-availability PICMG features are baked in. This is one equation that all adds up.

Kontron • www.kontron.com • www.mil-embedded.com/p47435

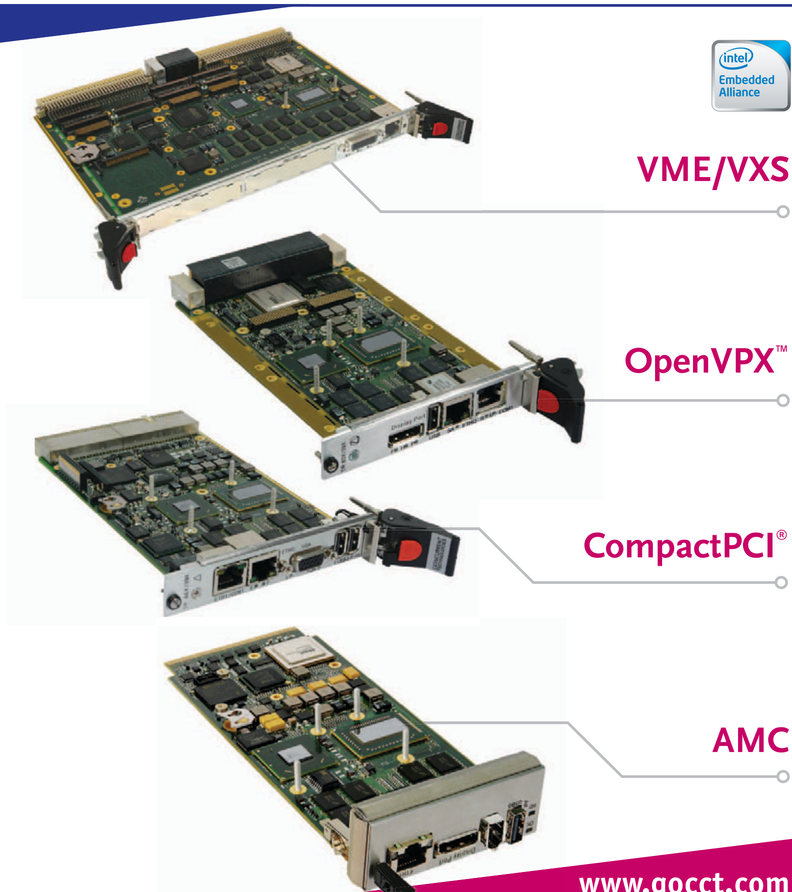

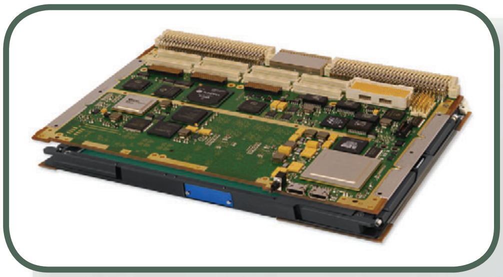

Increased performance without all that nasty power

Remember when Motorola/Freescale rocked our military worlds with the PowerPC and AltiVec? Yeah, we kinda remember that, too. But it's been a while and we're anxiously awaiting the revival of AltiVec *back* into the QorIQ family of Freescale processors. So here's our prediction: GE's 6U VME PPC10A, which was announced late last year and profiled herein, will soon be equipped with whichever Freescale CPU includes an AltiVec. Meanwhile, GE's PPC10A with QorIQ P4080 has a rockin' 8-core CPU yet still drinks no more power than the predecessor board called the *PPC9A* (which used Freescale's 2-core 8641D). So that's 6 extra cores in the same power budget … and soon an AltiVec. (We hope.) Too cool!

The PPC10A's cores zip along at 1.5 GHz and chat with up to 8 GB of dual-channel DDR3. There are 2 GbE ports (with 2 more as options), 2 SATA, 2 USB (with 3 more as options), and 21 GPIO ports. Two XMC/PMC sites promise even more add-on I/O or processing, and GE's AFIX expansion site allows customization with SCSI, VGA, 1553, a flash drive, and an AltiVec. Just kidding about that last one, at least until Freescale reintroduces a QorIQ with AltiVec. And since this is GE Intelligent Platforms, there are six ruggedization levels, including conduction-cooled flavors. Even though VPX is all the rage, lots of DoD programs have VME backplanes, and the PPC10A is the top shelf for VME-based Freescale SBCs.

GE Intelligent Platforms • www.ge-ip.com • www.mil-embedded.com/p47653

COM Express SBC tucks neatly into rugged shoeboxes

The trend in many military systems is a purpose-built box that fits the required space and I/O ports. This means that what's inside the box is insignificant – and increasingly, COM Express is what's so "insignificant." One perfect example is congatec AG's conga-CA6 COM Express Type 2 module, based upon Intel's E6xx Atom and EG20 platform controller hub. Consuming only 5 W at 95 x 95 mm, the COM boasts low battery life and all-rugged components with long-life longevity.

There are four speed grades available for the Atom: 600 MHz, 1.0 GHz, 1.3 GHz, and 1.6 GHz. Each accesses up to 2 GB of directly soldered DDR2 SDRAM. The onboard graphics controller has a 256 MB frame buffer, and Intel has made it 50 percent faster on tasks like DirectX 9.0E, OpenGL 2.0, and MPEG-2/-4. Graphics output is via 24-bit LVDS or SDVO. Other onboard I/O includes 3 PCIe x1, 2x SATA, PCI, EIDE, GbE, High-Def audio, and 6 USB. congatec can equip the conga-CA6 with up to 32 GB of flash memory, accessed via the SATA interface.

congatec AG • www.congatec.us • www.mil-embedded.com/p47546

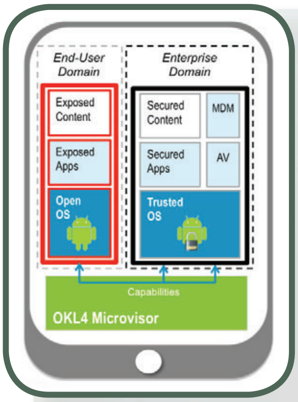

Companies show initiative when it comes to Android

You can't go anywhere these days without seeing someone whip out their smartphone to look up a restaurant, or to Google the answer to some question that arose during dinner with friends. And the military might not be far behind in adopting smartphones or tablets featuring Android and other open mobile OSs – thanks to OK Labs and Fixmo Inc.'s new "joint initiative to create and commercialize a highly secure COTS mobile platform." The platform will target civilian, federal government, and military needs, and will culminate in a "reference solution" – for tablet and smartphone OEMS – that melds system software, firmware, and mobile virtualization for information assurance. The primary goal is to isolate the organization's apps and data from the operator's personal user stuff … thus reinforcing the mantra: It's all about the security.

In engineering speak, this means: 1) The new initiative's progeny platform will include Fixmo's Sentinel security monitoring software running on an OKL4 Microvisor-based independent virtual machine to, once again, keep secure things separate on a single mobile device; and 2) The joint platform solution offers image profiling and tracking by checksums and size, in addition to handy-dandy activity auditing. And OK Labs and Fixmo aren't planning on resting on their laurels: They are also targeting the financial services and healthcare industries. But if we've whet your appetite, you'll still have to wait a few months: The platform isn't slated for market delivery until Q3 2011.

OK Labs • www.ok-labs.com • www.mil-embedded.com/p47654
Fixmo Inc. • www.fixmo.com • www.mil-embedded.com/p47655

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

Crosshairs Editorial

COTS means Competition, and that's good

By Chris A. Ciufo, Editor

On February 15, I received a text message from AT&T on my iPhone:

> *"AT&T is increasing the amount of data included in your DataPro Tethering plan from 2 GB to 4 GB. … The monthly charge … will not change …"*

Wow! I won the data jackpot and just got 2x more for my money. But why? It's because Verizon was soon to launch its iPhone with (then) unlimited data plans, and AT&T didn't want to see me defect. In essence, the Verizon competition forced AT&T to up the ante to keep its customers happy and want to stay put. Competition has a funny way of leveling the playing field, keeping prices down[1], and forcing competitors to find new ways to stay in the game. This happens all the time in the COTS industry, and military programs are the beneficiaries.

Let's deviate from smartphones for a moment, with a counter example. Before Secretary of Defense William Perry's famous "Perry Memo," MIL-SPEC ICs were still available and controlled by DESC drawings. Manufacturers took a hard look at sales forecasts before committing to a new device in 883 or 38510 flavor since although the price was 10x higher, so were the costs. Class S devices for radiation were even higher: typically 100x. This wasn't sustainable as the PC market exploded, since more money could be made on new ICs like modems, DSPs, and microcontrollers – which, incidentally, were designed in and out of systems roughly every 18 months on the "COMDEX cycle."[2]

Perry's Memo told DoD designers to use whatever IC met the ORD, essentially creating competition for the stoic MIL-SPEC ICs. One result: The nonmilitary PowerPC (PPC) got deployed and blew the doors off of 68000 CPUs, and the PPC's AltiVec brought mainstream DSP processing to any high-rel system. Slot count in ATR boxes declined as horsepower/box increased with other processors and FPGAs. Image, radar, and SIGINT processing systems like Rivet Joint blossomed with unheard of capabilities. As well, heavy processing boxes like BAE's ATIRCM with Common Missile Warning System got the chance to prove themselves in lifesaving airborne countermeasures. Civilian ICs helped out a lot.

It's a good thing that competition even happens within the established military industry. VME – the military's favorite rugged board type – started to get stagnant as the limitations of 80 MBps parallel backplane buses mounted. System engineers watched the personal computer switch from PCI to LVDS-based serial "buses" like PCI Express. Heck, why couldn't VME switch to a serial backplane? VPX, and later interoperable OpenVPX, became available. There were more than 20 OpenVPX vendors listed in our online database (www.vmecritical.com/products) in 2010, up from only 10 in 2009. In 24 months, we've logged 261 VPX or OpenVPX products in our online database, too. It took years to get this many VME products. Competition is still a good thing.

Yet VPX itself is under pressure from Small Form Factor (SFF) rugged "shoeboxes," designed for myriad rugged non-mil applications but easily adapted for military use. Parvus (now part of Eurotech) pioneered this concept in the 1990s with their rubber-bumpered boxes, and now VME vendors Curtiss-Wright, Kontron, Themis, and others build application-specific SFF shoeboxes. GE's latest, the oddly named CRS-C2P-3CC1, is stuffed with CompactPCI while similar products from Kontron/AP Labs use COM Express. Who cares what's inside the box?

VITA, the VME standards organization, recognizes the competing trend toward SFFs and shoeboxes and has launched three new standards activities. VITA Executive Director Ray Alderman cites SFFs as a growth area in his "2011 State of the VITA Technology Industry" white paper just published as we went to press. Competition in this segment will come from PC/104 stackables and PICMG COM Express vendors such as the Taiwanese heavyweight Advantech. And if an LCD is involved in the system, look for industrial and kiosk rugged panel PC competitors such as IEE and VIA. An SFF or shoebox is increasingly a better choice in some systems than VME or VPX.

Lastly, whether software for military platforms was done in-house or purchased from established RTOS vendors like Wind River, today's COTS choices are numerous. Besides flavors of Windows running on extremely capable Intel Atom or Core processors, there's Google's Chrome, Linux, and Android. Since its introduction in 2009 in smartphones, Android-based devices now control 29 percent of the U.S. market, said analyst firm Nielsen in their report published in January 2011.

Android's shake-up in civilian smartphones has spilled over into defense as well. According to Jaime Rubscha, Practice Lead, Defense Programs at DoD contractor Product Development Technologies, the market is booming in Android accessory platforms for the warfighter. Primes like Boeing, Harris, and Thales are looking for wearable Android-based computers, biometric sensors, and GPS units to bolt onto JTRS-like, high-bandwidth tactical radios. Five years ago, it would've been ludicrous to consider using a cell phone OS in a soldier's computer. Yet now, because of competition in the red-hot smartphone market, a warfighter can even download a mil-only RECON Android app on a scout mission.

---

[1] This even works with airlines, though oil north of $100 a barrel has ended airline price wars as they all raise fares. The International Air Transport Association estimates 29 percent of operating costs are in fuel. No amount of competition can erase this away.

[2] Now dead, COMDEX was the biggest industry tradeshow. Today, that honor goes to the Consumer Electronics Show (CES), where this year 20+ new tablets were born and Internet-enabled TVs are a consumer reality.