MIL-EMBEDDED.COM

VPX rising ... past VME?

Military EMBEDDED SYSTEMS

VOLUME 5 NUMBER 7
OCTOBER 2009

INCLUDING:

Mil apps blend real-time with realism

Also:

Does EAL 6+ make the grade?
An interview with Wind River's Marc Brown

Supplement inside:
MIL/COTS DIGEST

DDC®
©Data Device Corporation

Total-ACE™

1 Complete MIL-STD-1553 Single Package Solution

WORLD'S 1ST Fully Integrated,
MIL-STD-1553 Terminal & Transformer Solution!

LESS … SPACE!

LESS … COMPLEXITY!

LESS … COST!

DDC®
©Data Device Corporation

www.ddc-web.com Toll Free: 1-800-DDC-5757

Innovation sets us apart

Avionics Databus Test, Simulation & Embedded Solutions

Lightweight, small, rugged, reliable, feature rich and cost-effective. These are only a few of the benefits of the new AB2000 series of Avionics BusBoxes® from Ballard Technology. They deliver outstanding avionics performance with multi-protocol and flexible expansion capabilities. Conductively cooled, the AB2000 is ideal for rugged avionics applications such as serving and recording data, converting protocols and providing federated control. Find out today how AB2000 innovation can work for you.

The AB2000 offers multi-protocols in a single, lightweight rugged format.

ARINC 429
MIL-STD-1553
Ethernet
RS232/422
Digital I/O
USB

Innovation for your success

For more information call 425.339.0281 or visit us at www.ballardtech.com

Ballard TECHNOLOGY

Military EMBEDDED SYSTEMS

October 2009 Volume 5 Number 7

COLUMNS

DEPARTMENTS

ON THE COVER:

Simulated but realistic graphics environments are an essential part of military operations. They can be used for training purposes, honing operator skills without the risk of life or equipment loss. They're also how pilots operate Unmanned Aerial Systems (UASs) remotely, as drones replace manned aircraft. In this issue of *Military Embedded Systems*, Yannick Lefebvre of Presagis (page 23) describes ways to create simulated environments for UAS operator training. And our resident "gray beard" Duncan Young follows upon last month's Quantum3D article on mobile graphics technology with his own opinion on the subject of Graphics Processor Units or GPUs (page 8). [Cover photo: A B-1B Lancer flies a combat patrol over Afghanistan in support of Operation Enduring Freedom, while a civilian aircraft approach is mapped out virtually. B-1B photo courtesy of U.S. Air Force/Staff Sgt. Aaron Allmon; inset courtesy of Quantum3D]

EVENTS

www.opensystemsmedia.com/events

SAE AeroTech Congress & Exhibition
November 10-12, 2009 • Seattle, WA
www.sae.org/events/conferences/aerospace/#2009

SDR'09 Technical Conference and Product Exposition
December 1-4, 2009 • Washington, DC
www.sdrforum.org/SDR09

WEB RESOURCES

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

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

Published by: OpenSystems media

ISSN: Print 1557-3222

We Won't Leave You High and Dry.

Recent consolidations in the embedded industry have left many PowerPC® users wondering where their next real-time or embedded Linux® operating system will come from.

At LynuxWorks, we are here to help. Our new PowerPC rescue package offers an open-standards based RTOS or embedded Linux, and a migration kit to help PowerPC users stay with their processor of choice.

Visit www.lynuxworks.com/rescue-me to request a migration whitepaper and an emergency hand-crank flashlight-radio*.

LYNUXWORKS™

Open.Reliable.Safe.Secure.

*limited quantities available

ADVERTISER INFORMATION

Page	Advertiser/Ad title
34 MCD	**ACT/Technico** – Solid state storage solutions
51	**Alphi Technology Corporation** – The expert in Military Technologies
7	**Annapolis Micro Systems, Inc.** – WILDSTAR 5
30 MCD	**Annapolis Micro Systems, Inc.** – WILDSTAR 5
3	**Ballard Technology** – Innovation sets us apart
17	**Conec Corporation** – CONEC Embedded Series
56	**Curtiss-Wright Controls Embedded Computing** – Rugged deployed electronic systems
2	**Data Device Corp. (DDC)** – Total-ACE
38 MCD	**ELMA Electronic** – Cutting edge backplanes
43 MCD	**ELMA Electronic** – Rethink MicroTCA
12	**ELMA Electronic** – VPX
35 MCD	**Excalibur Systems, Inc.** – Ready for the unexpected?
13	**Excalibur Systems, Inc.** – Dense?
55	**GE Fanuc Intelligent Platforms, Inc.** – Which of these platforms use GE Fanuc
21	**Great River Technology, Inc.** – ARINC 818
33 MCD	**Hartmann Elektronik** – VME
19	**Kontron** – We do not build Navy ships
44 MCD	**Kontron** – We do not build Navy ships
42 MCD	**LCR Electronics** – Three rugged AdvancedTCA chassis
5	**LynuxWorks, Inc.** – We won't leave you high and dry.
42 MCD	**MEN Micro Elektronik GmbH** – PICMG 2.30 CompactPCI PlusIO SBC
18	**Nallatech** – High performance FPGA solutions
25	**Parvus Corporation** – Qualified to perform
22	**Phoenix International** – Data storage
31 MCD	**Positronic Industries** – PICMG Power Connectors
37 MCD	**RadiSys Corporation** – Next Generation COTS-Based Solutions
42 MCD	**RadiSys Corporation** – Leveraging a ruggedized AdvancedTCA platform
40 MCD	**Schroff a Brand of Pentair Electronic Packaging** – Performance
33 MCD	**Schroff a Brand of Pentair Electronic Packaging** – XTCA Cube
39 MCD	**Schroff a Brand of Pentair Electronic Packaging** – Planning your next product?
53	**SDR Forum** – SDR '09
22	**TEWS Technologies LLC** – COTS I/O Solutions
27	**Themis Computer** – Themis has ... Speed to Burn
28	**Trident Space and Defense** – Protect your data
24	**Tri-M Systems Inc.** – 100Mhz PC/104 Module
47	**Tri-M Systems Inc.** – PC/104 Can-Tainer
50	**VersaLogic Corp.** – When you're at 30,000 feet, quality matters
49	**VPT Inc.** – Efficent, reliable power
10	**White Electronic Designs** – Trapped in your design?
41 MCD	**XTech** – The XTech Files

E-LETTER

Nov/Dec: www.mil-embedded.com/eletter

- Cell phones: The next great *software* driver
  By Chris A. Ciufo
- Pulling out secrets: New technologies help military find intelligence in degraded documents
  By Dr. Steven Schlosser, NovoDynamics

OpenSystems media

Military EMBEDDED SYSTEMS

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

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

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

Made in the USA

Ultimate Modularity

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

Direct Seamless Connections with no Data Reduction

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

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

Releasing the processing potential of GPUs

By Duncan Young

The internal computing architecture of high-performance Graphics Processing Units (GPUs) has evolved from fixed-function graphics execution units to arrays of fully programmable Single Instruction Multiple Data (SIMD) processors. This evolution has been driven by the demand from the video gaming community to perform generic physics calculations in parallel to give greater realism to the behavior of, for example, smoke, debris, fire, and fluids. The ability to offload and accelerate these same types of repetitive parallel calculations onto a GPU offers great potential to military technologies such as radar, sonar, and image processing. The key to efficient implementation is a new generation of tools such as OpenCL and CUDA, which integrate code development across heterogeneous CPU/GPU architectures along with the memory and I/O bandwidth to support them.

SIMD processing rays

At its heart, a high-performance GPU device will typically have up to 128, 32-bit single precision processor cores clocked at 1 GHz or more. These are organized as parallel SIMD arrays so that groups of processors can execute the same instructions on different data sets in parallel. When operating as a GPU, the primary requirement is to utilize animated 3D graphics functions such as shaders. However, GPUs are evolving away from being specific shader processors and are becoming more generic math processors, now referred to as "stream processors." With the right tools, GPUs can be applied much more broadly to accelerate many kinds of PC-based applications such as genetic research, seismic processing, meteorological processing, and DSP, at much lower cost than other more specific forms of hardware acceleration.

One major GPU manufacturer, NVIDIA, has developed a software environment known as CUDA to release the potential of the GPU into these other application areas. CUDA supports the combination of CPU and GPU by allowing inline C code development through an abstracted library of functions that hides the specifics of the GPU's stream processors and their interface with the CPU. This provides a very flexible programming interface and permits growth or even radical change of the stream processors in the future, without impacting existing code. To reduce the scope for errors, CUDA adopts a straightforward programming model that manages multiple threads internally to optimize processor utilization so that there is no need to write explicitly threaded code.

GPU without graphics

Paradoxically there will be a class of embedded applications that will not generate any local graphical output at all. Typically, this class could comprise image processing in UAVs or underwater Remotely Operated Vehicles (ROVs) or many other types of unmanned sensors. An embedded PC with GPU becomes an ideal platform for image enhancement, stabilization, pattern recognition, target tracking, video encoding, or encryption/decryption. These are all applications that can be written in regular C code to run on a high-performance PC but could be accelerated to run orders of magnitude faster by a GPU stream processor. The GPU provides generic parallel processing already integrated into many PC configurations and requires less specialized skills than, for example, an FPGA development, by using off-the-shelf tools such as CUDA, the MathWorks' MATLAB, and ported VSIPL DSP libraries to construct, test, and verify the application.

PCI Express key to performance

An embedded sensor processing application requires high data bandwidth to receive and process the continuous stream of incoming raw image data. The ability of CUDA to handle multithreading and consequently maximize the processing load of the GPU's SIMD arrays is dependent on the performance of both the external interface and its local memory interface. High-end GPU devices will use 16-lane PCI Express 2.0, doubling the earlier PCI Express 1.0 data rates, to give a theoretical 500 MBps per lane. For rugged embedded applications, this fits well with both the popular 3U and 6U formats of the VPX (ANSI/VITA 46) packaging standards with their extended high-speed connectivity. The MAGIC1 rugged embedded PC from GE Fanuc Intelligent Platforms, pictured in Figure 1, is based on the 3U VPX form factor and has been reengineered and enhanced to support CUDA-capable GPUs from NVIDIA. While such an embedded PC fits comfortably within the 3U format, a 6U profile also has the real estate and greater connectivity to potentially make a new class of powerful multicomputing engine based on a number of multicore processors and GPUs using PCI Express 2.0 as the interconnect.

The GPU is evolving rapidly, creating a processing capability with broad application across many diverse markets. CUDA and similar development environments provide accessibility to this untapped performance reserve. Consequently, the rugged military and aerospace domain appears set to change the way complex, time-consuming sensor applications can be developed, tested, verified, and successfully deployed.

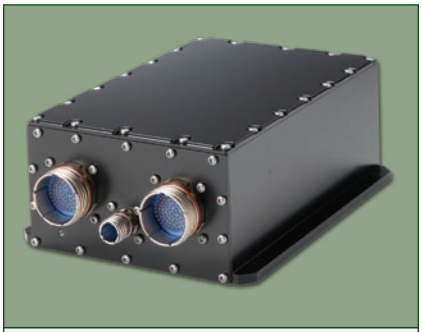

Figure 1 | MAGIC1 is a rugged, embedded PC from GE Fanuc Intelligent Platforms, based on the 3U VPX form factor and reengineered and enhanced to support CUDA-capable GPUs from NVIDIA.

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

Preparing for lead-free electronics

By John Wemekamp

Commercial and domestic electronics products of all types have been heavily influenced by the European directives for Restriction of the use of certain Hazardous Substances (RoHS) and Waste Electrical and Electronic Equipment (WEEE). Across industry, the most significant change has been the rapid adoption of lead-free solder. The worldwide component supply chain is now fully geared up to provide lead-free terminations for all types of components as first choice, while some manufacturers are still offering traditional tin/lead as an alternative but at an increasingly higher cost. Military and aerospace equipment vendors, integrators, and end users are still apprehensive about using lead-free materials. This is because of lingering doubts about the critical long-term reliability of lead-free electronics in the more highly stressed temperature and mechanical environments of many military applications.

GEIA-STD-0005-1

As a result of the worldwide lead-free initiatives, future military contracts will require lead-free project plans. For the U.S. DoD, these will need to be in accordance with GEIA-STD-0005-1 "Performance Standard for Aerospace and High Performance Electronic Systems Containing Lead-free Solder." Such a plan will include reliability, configuration control, risks and limitations of use, and the effects of tin whiskers, plus the repair, rework, maintenance, and support of the equipment. However, the preparation of a standard or plan does not necessarily impose the use of lead-free electronics. Instead, it demonstrates an understanding of the impact of the component and materials supply situation in order to mitigate risks. Although the U.S. is likely to adopt similar legislation pertaining to RoHS in the near future – and despite possible exemptions for military and aerospace – some newer programs such as the Future Combat System (FCS) are already being selectively proactive in moving to lead free.

Lead-free solder

Using lead-free solder is the most obvious change, and it has a number of secondary effects on reliability. A leading contender for lead-free solder is SAC305 (Ag – silver 3.0 percent, Cu – copper 0.5 percent, the remainder being Sn – tin). This type of solder has some undesirable mechanical characteristics compared to tin/lead:

- Its higher melting point exacerbates manufacturing and rework stresses in components and the Printing Wiring Board (PWB).
- Its microstructure is less homogeneous, causing its properties to be more orientation dependent.
- The absence of lead results in stiffer solder joints. During extended thermal cycling, this increased stiffness can become a critical reliability issue for larger Ball Grid Arrays (BGAs) due to stresses imposed by differential rates of expansion between the device and PWB.

The long-term solution does not lie with expedients such as interposers or replacing solder balls with tin/lead but instead lies in additional material research. New, more stable PWB materials have been developed to reduce some of these effects but in turn have introduced the risk of pads tearing from the PWB, known as pad cratering, illustrated in Figure 1.

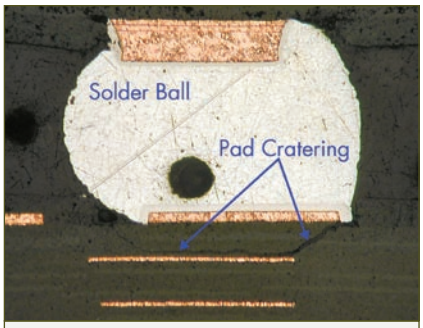

Figure 1 | The cross-section photo shows pad cratering on a lead-free solder joint. Stiffer lead-free solders and more brittle PWB materials are leading to an increased risk of pad cratering on lead-free boards.

Tin whiskers

Tin whiskers is a phenomenon whereby thin conductive filaments grow from the surface of a tin-plated termination that can potentially short-circuit adjacent pins of devices. The cause of whiskers is not entirely understood but is believed to be related to compressive stresses in the plating. It can happen at any time during the equipment's life and is therefore difficult to predict or mitigate. However, the most effective mitigations for tin whiskers include design rules that reduce the risk of shorts between solder pads, in addition to some types of conformal coatings that counter whisker growth.

Impact on COTS vendors

Suppliers of COTS embedded computing equipment are planning to introduce lead-free assemblies for military applications. But, as lead-free technology is introduced, it will require much additional test and qualification data to support its use in deployable projects. There will be a lengthy transition period during which COTS vendors will continue to manufacture and support both tin/lead and lead-free alternatives of the same product lines and provide the levels of configuration and materials management required to maintain each type of product over the life cycle. Curtiss-Wright Controls Embedded Computing (CWCEC) has adopted a proactive approach, having already amassed much test data, introduced new materials and test regimes such as Interconnect Stress Testing (IST), and developed the full manufacturing and life-cycle capability to introduce lead-free assemblies.

Lead-free products will require additional testing and tailoring during development to optimize long-term reliability, and COTS vendors will need to incorporate this data into their customers' projects lead-free plans. The rate of adoption of lead-free electronics for use in harsh environments will be determined partly by legislation but primarily by both vendors' and integrators' accumulated knowledge base and continuing research and development efforts.

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

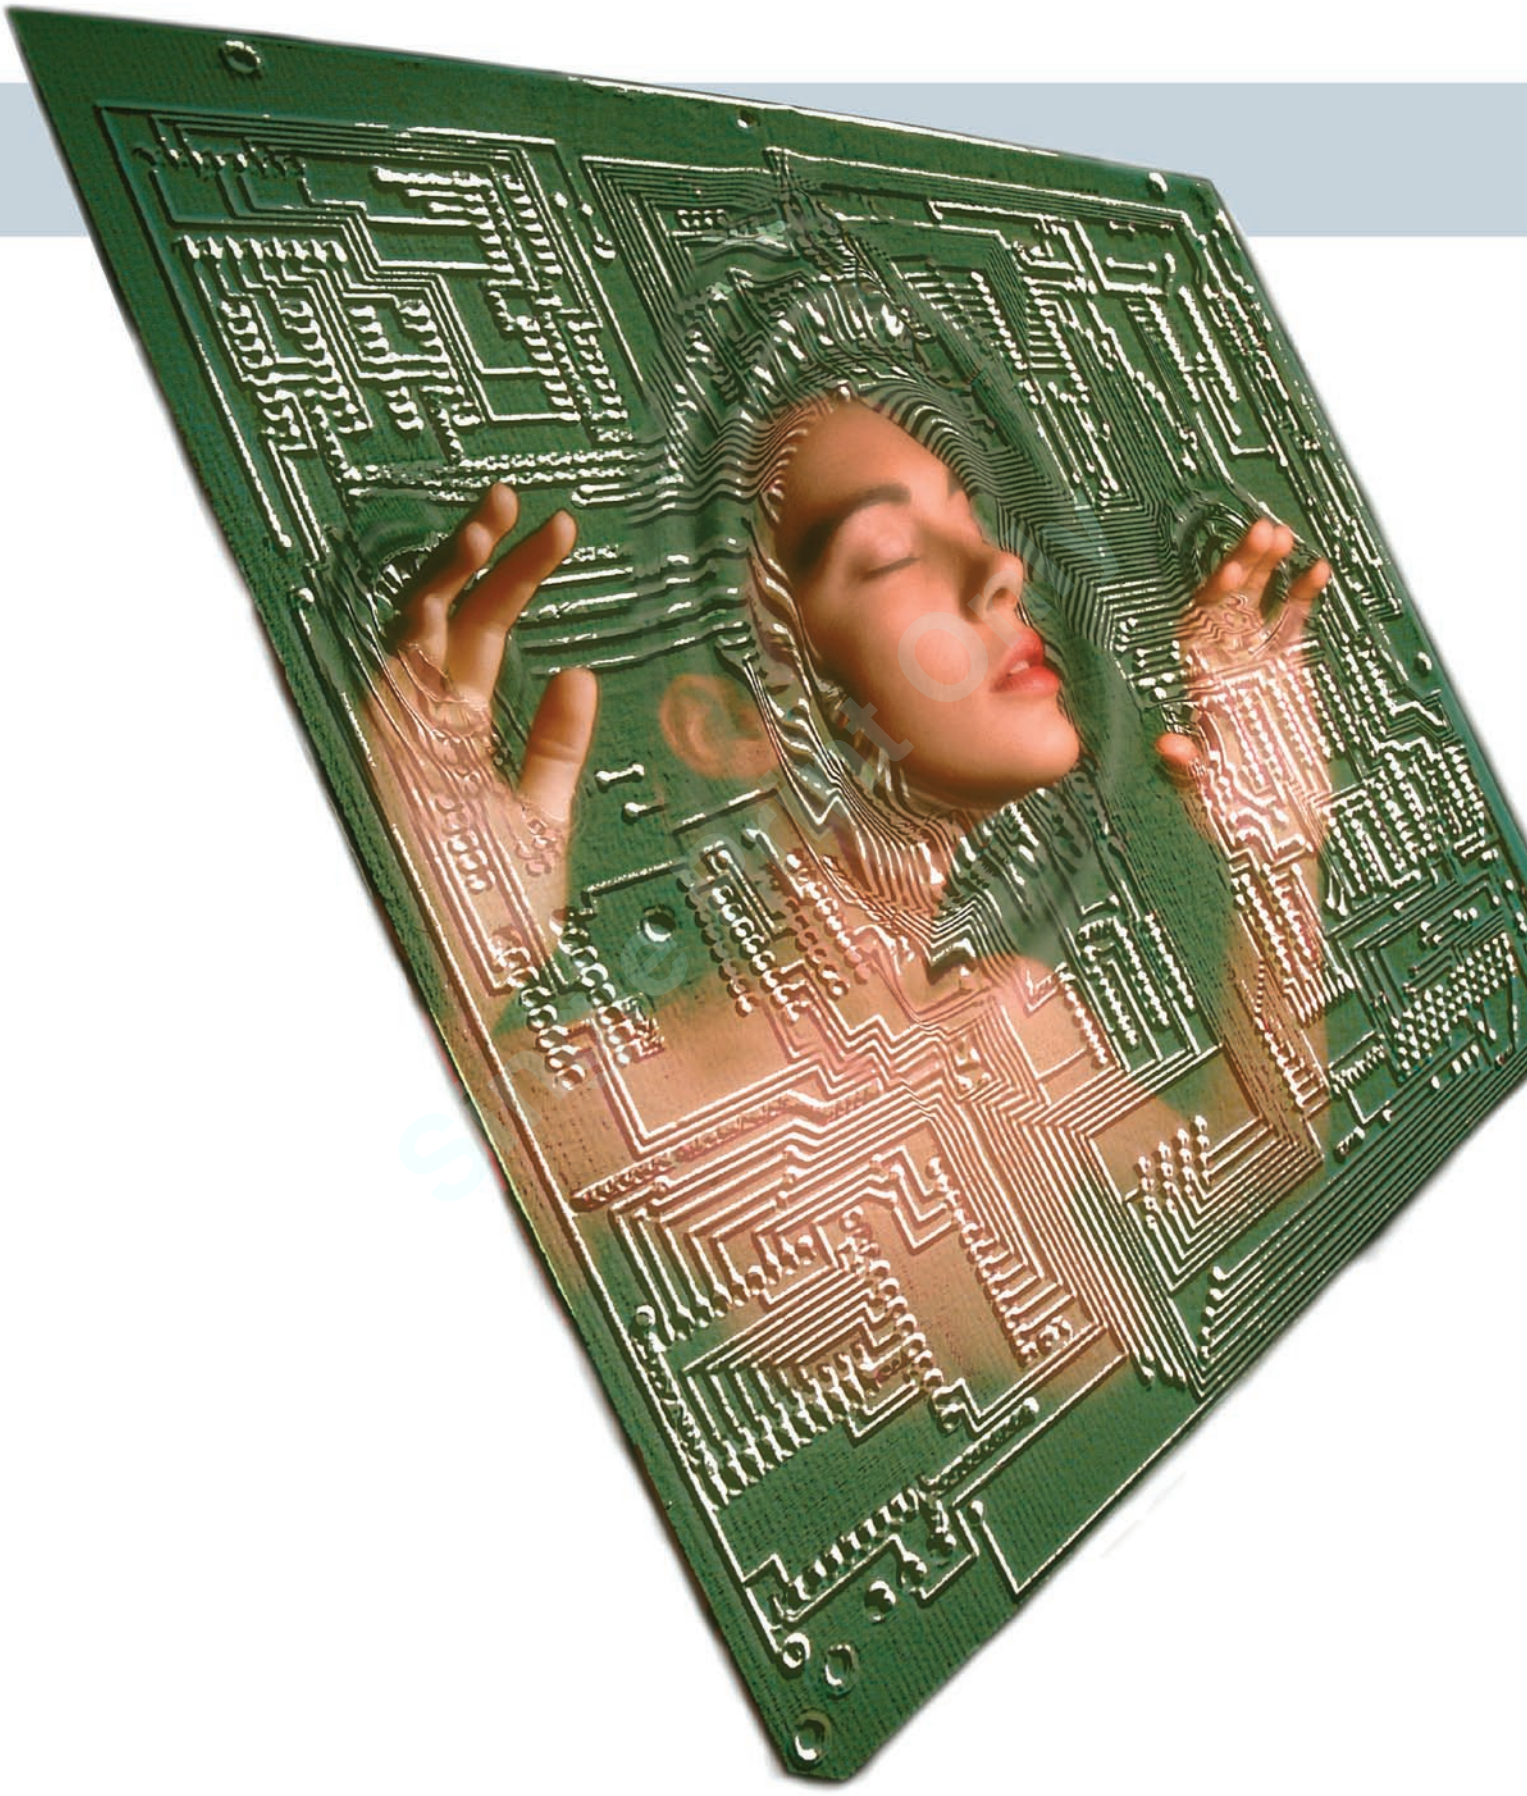

Trapped in Your Design?

Break free with our space-saving solutions.

Escape the limitations of your board design with advanced microelectronics solutions from White Electronic Designs. Use our compact, high density memory modules to free up board space and improve performance.

Try our new DDR2 SDRAM that packs 1 GByte into a SnPb BGA package less than 1/2 in^2 … or our new embedded SLC NAND flash BGA that has the performance of a solid state disk drive in a form factor that is one-fourth the volume of a comparable compact flash-based product.

WEDC lean microelectronic solutions are designed to solve component incompatibilities, reduce design complexity, and extend product life and environmental performance in military applications. We offer turnkey design, assembly and test of custom multi-chip solutions, and a wide range of standard military off-the-shelf space-saving solutions.

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

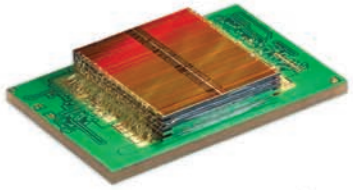

602.437.1520 TEL | 602.437.9120 FAX

WHITE ELECTRONIC DESIGNS

www.whiteedc.com

Customer First. Quality Always.

Legacy Software Migration

By Tom Schultz

Modernizing legacy military systems with top-grade static analysis

Outdated legacy code bases in warfare and defense systems create vulnerabilities to safety and prohibit the introduction of new, network-centric technology. With static analysis, these systems can be retrofitted to be high integrity, making devices more technologically competitive.

The need to modernize

The world of warfare and defense has changed. In the past 10 years alone, the main factor in ranking the best military has changed from merely possessing the largest and most powerful fighting force to having the best possible communication between personnel and military machines on the battlefield. These military machines are no longer as hardware-centric as they were decades ago; most employ embedded software systems to perform essential functions. The Abrams M1 tank currently deployed in Iraq and Afghanistan, for example, has a computer-aided targeting system that allows precise aiming in conjunction with a thermal imaging system.

These types of embedded systems are designed to be as robust and reliable as possible since lives depend on their effectiveness. These systems are highly specialized and – once built and deployed to the field – are generally meant to be used without further retrofitting for the lifetime of the unit. Unfortunately, such a design philosophy does not take into account the increasing roles these systems are required to perform eventually, as technology progresses, and safety can become an issue.

In the same way, increasingly, networking has taken on an expanded role in military devices over the past few decades. Modernizing legacy systems to take advantage of new networking functionalities, which rely on software, has become mandatory. Today, tanks and other vehicles are now nodes in a larger interconnected system that continually shares information. Each node in the system can work with other nodes in its network (or even nodes in other networks) to solve problems that could not be handled by any one system.

Software-driven systems like Mine-Resistant Ambush Protected (MRAP) vehicles containing Communication-On-The-Move (COTM) systems depend on software architecture and must be modernized to become network-centric,

for example. And modernization is a good strategy, even when optional, because doing so can reduce software maintenance costs in the long run. Once built, legacy applications should theoretically last forever, but they seldom do because routine maintenance ultimately causes the code base to deteriorate. Maintaining old code is very expensive (as high as 10 percent of the initial development budget per year) because it involves patching vulnerabilities and improving reliability. This is usually done by expensive developers possessing a rare right skill set. In contrast, new code is much cheaper. Although modernizing code has a higher price up front, it is often better to do so because regular long-term maintenance of legacy code costs more in the long run.

Legacy systems are generally modernized (but not fully replaced with new units due to cost) because many currently deployed military systems have at least another decade of planned functional lifetime left. Since safety and reliability are paramount, it is essential to modernize software to include the desired new capabilities without compromising previous quality or going over budget. Static analysis makes this task easier than it would be otherwise.

Static analysis for modernizing legacy code

Retrofitting legacy code for new functionality or to bring it up to new standards should not be done lightly. Since legacy systems are often mission critical and already deployed in the field, revisions should introduce minimal new coding defects into existing code. As a code base becomes more complicated, it becomes more difficult to isolate and eliminate defects since there are more variables that might interact with and stack on top of other defects in unforeseen ways.

While older tools use syntax and style-based testing, modern static analysis uses path simulation through dataflow analysis in conjunction with Boolean satisfiability. Path simulation works by emulating what a program will do once it has been compiled and executes each function in its code base. This allows for the added detection of resource allocation issues, pointer mismanagement, improper buffer and string usage, and tainted data handling. Boolean satisfiability translates every operation in a program into Boolean variables and uses this data to test for potential runtime problems like integer overflows. These methods generate a false-positive rate as low as 15 percent.

To the rescue, however, is static analysis, which helps developers refactor and restructure the code bases of legacy programs to bring them up to modern standards. Static analysis reveals the steps performed at the creation of a legacy project so developers can understand the current architecture and original intention of the program. This feature innovation alleviates much of the guesswork of detecting resource allocation issues, pointer mismanagement, improper buffer, string usage, and tainted data handling when trying to upgrade code. Static analysis tools, such as those offered by Coverity, are also able to display information about any defect in question and show how it adversely affects the overall program.

Defense, similar to other software dependent industries, cannot afford to allow its technology to lag behind the times. The modernization of legacy code is inevitable, but advanced static analysis provides a sure remedy.

Tom Schultz *brings 30 years' experience to his role as Director of Products at Coverity. Previously, Tom was founder and CTO of Codefast, which Coverity acquired in 2007. Tom also spent more than 10 years at Rational Software and later IBM Rational. He holds a B.S. in Biology-Geology from the University of Rochester. He can be contacted at tschultz@coverity.com.*

Daily Briefing: News Snippets

By Sharon Schnakenburg, Assistant Managing Editor

www.mil-embedded.com/dailybriefing

'Background' technology moves forward in development

While the Orion Crew Exploration Vehicle (CEV) gets all the limelight for its part of NASA's year 2020 moon-and-beyond Constellation Program, the Orion CEV wouldn't get off the ground, literally, without its more "background" technology: the perhaps less-glamorous (but just as essential) 25-ton payload capacity Ares 1 vehicle booster (Figure 1). Consequently, Ares 1 development is progressing with a recent contract between Ball Aerospace & Technologies Corp. and Aitech Defense Systems. The contract specifies that Aitech provides its S950 single board computer, in addition to its S750 PMC, to be utilized for Ares 1's Command Telemetry Computer (CTC) and Instrument Unit Avionics (IUA) Flight Computer (FC) systems. CTC and IUA FC systems proffer the navigation, control, and guidance hardware for Ares 1 and act as a vehicle control subsystem when the rocket ascends into orbit. The S950 will lend its 1 GHz processing capability, while S750 will render high-speed data imaging to Orion CEV's ground support system and solid-state recorder.

Figure 1 | Ares 1 photo courtesy of Aitech

Turkish Armed Forces build SCA repertoire

Prototyping … porting … testing … it's all in a day's work when developing SDR waveforms. Thus, Spectrum Signal Processing by Vecima's ("Spectrum's") flexComm SDR-4000 multipurpose, SCA-compliant transceiver will be used by The Scientific and Technological Research Council of Turkey (TUBITAK) and the National Research Institute of Electronics and Cryptology (UEKAE) for the aforestated triad of purposes. Specifically, "TUBITAK UAKAE" will use flexComm SDR-4000 within the final phase of its twofold SDR project: Building a test bed for testing and developing SCA-compliant waveforms per Turkish Armed Forces (TAF) specifications, and SCA verification and testing on SDRs to be utilized by the TAF.

Helmet jettison device to enlighten Air Force

Sometimes the full ramifications of events are not understood until much later, if at all. However, the U.S. Air Force is aiming to change that with a recent Small Business Innovation Research (SBIR) Phase 1 grant to Diversified Technical Systems, Inc. (DTS) for a helmet jettison device. The project's threefold objective is to design a helmet-mounted device that can execute real-time analysis of sensor data; to create simulated eject conditions under which the device is tested; and to perform analysis on existing testing data on ejection events and Air Combat Maneuver (ACM) events. Based on a customized version of DTS's e-SENSING Embedded Recorder (ER), the smart helmet-mounted sensor system will specifically be used in the measurement and detection of blunt impact and blast events and will send a release signal to the Helmet-Mounted Display System (HMDS) during eject events.

Iraqi Air Force's vision widens, courtesy of USAF

Situational awareness makes or breaks a mission – and the safety of military personnel and nearby civilians. Consequently, the USAF and Lockheed Martin recently penned a $28.1 million Phase 1 contract to enhance the situational awareness of the Iraqi Air Force in the civilian and military airspace. The Iraq Air Command and Control System contract, awarded by Hanscom Air Force Base's Electronic Systems Center in Massachusetts, is designed to afford the Iraqi Air Force with modernized C2 technology interoperable with Coalition forces. Specifically, the agreement stipulates that Lockheed Martin renders the Iraqi Air Force an integrated technology network comprising: a Ground-to-Air Transmitter and Receiver (GATR) site, a networked comms infrastructure consisting of a training suite and Sector Operation Center (SOC), and a long-range AN/TPS-77 radar system (Figure 2) for air surveillance. Delivery is slated to occur inside the next 27 months.

Figure 2 | AN/TPS-77 radar system, photo courtesy of Lockheed Martin

3U VPX drafted for M1A2

The 3U VPX (VITA 46) form factor will soon be tucked inside the U.S. Army's rough-and-ready Abrams M1A2 tank (Figure 3), according to a recent contract between prime General Dynamics Land Systems (GDLS) and GE Fanuc Intelligent Platforms. As part of the Army's M1A2 Abrams Evolutionary Design (AED) program, GE will supply 3U VPX wares including: graphics processors, single board computers, switches, and disk subsystems. The bonus capability: The wares are all additionally incarnated in accordance with the REDI (VITA 48) specification, yielding resiliency in the harshest environments and easing Two Level Maintenance via their Line Replaceable Module (LRM) format. Specified under the contract are GE's SBC340, an SBC based on the Intel Core 2 Duo processor; SBC310, a Freescale 8641D-based SBC; SDD910 solid-state SATA disk module; GBX410 GbE switch featuring 16 ports; PEX430 PMC/XMC carrier and PCI Express switch; and the GRA110 NVIDIA G73 graphics processor. AED will gradually supersede the current GE Fanuc-supplied 6U VME-based systems under the Abrams Continuous Electronics Enhancement Program (CEEP) and System Enhancement Program (SEP V2).

Figure 3 | M1A2 Abrams tank, U.S. Army photo

PHM extends 'advanced aircraft' life cycles

Though it has, of course, been targeting matters of national security for decades, the U.S. Air Force has a relatively new target: the identification and analysis of jet engine equipment degradation. Accordingly, the USAF recently awarded DSPCon, Inc. a Phase II Small Business Innovation Research (SBIR) contract for ongoing optimized-processing-platform development, in accordance with the USAF's goal of heightening Propulsion Health Management (PHM) systems' signal processing capabilities. As the platform in development provides real-time data capture and enhanced diagnostics, the USAF can detect jet engine equipment degradation faster and more thoroughly, before equipment fails. The net result: Extended and improved equipment life cycles. PHM technology will be utilized in "advanced aircraft" and commercial aircraft applications.

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

Figure 4 | AH-64 Apache, U.S. Army photo by Sgt. 1st Class Brent Hunt

Boeing demo embodies Future Force

The U.S. Army knows that networking technology such as its Wideband Networking Waveform (WNW) can bring to life its vision of Future Force networking, enabling video, voice, and data communications between military personnel in the sky and on Earth. Accordingly, Boeing recently held its first WNW integrated capabilities demonstration. Held at the U.S. Army's C4ISR On-The-Move (OTM) Event 09, the demonstration comprised an airborne AH-64 Apache helicopter (Figure 4), a tactical operations center, several High Mobility Multipurpose Wheeled Vehicles, and an airship, all equipped with surrogate JTRS radios sporting WNW version 3.1a. The WNW network was monitored by a Boeing-designed network management software that proved compatible with the network operation manager system of Warfighter Information Network-Tactical (WIN-T). Demonstration information transmitted by WNW consisted of imagery, streaming video, spot reports on hypothetical enemy forces, and more.

USMC radios: Making the grade?

Whether in legacy or new digital style, U.S. Marine Corps' radios need testing to keep them at "grade A" level. Thus, the USMC recently granted Aeroflex a $40.5 million, five-year contract for Ground Radio Maintenance Automatic Test Systems (GRMATS): Aeroflex's 7200 Configurable Automated Test Set (CATS). The 7200 CATS (Figure 5) is designed to test SDRs such as military tactical radios, with a modular design to easily accommodate future technology. The 7200 CATS is also crafted for compatibility with the Joint Tactical Radio System's (JTRS's) Software Communications Architecture (SCA) and includes multi-gbps data buses featuring plug-and-play software architecture. Additional highlights include MIL-PRF-28800F Class 3 packaging, 90 MHz signal generation and instantaneous digitization bandwidth, 2.6 GHz frequency coverage, and phase noise capability of more than -108 dBc/Hz.

Figure 5 | 7200 CATS test system, photo courtesy of Aeroflex

‘EAL 6+’ says it all – or does it?

Q&A with Marc Brown, Vice President, VxWorks Product Strategy and Marketing for Wind River Systems

EDITOR’S NOTE

To be or not to be … EAL 6+ certified: That is the question. Apologies to William Shakespeare, but you get the point. With EAL 6 (or 6+) certification becoming the latest trend among embedded software vendors, the criteria for achieving such should be fairly straightforward … or is it? Marc Brown, VP over VxWorks Product Strategy and Marketing at Wind River Systems, reveals what he claims are misassumptions about EAL 6+, as the company’s VxWorks MILS Platform 2.0 undergoes NSA/NIAP evaluation for EAL 6+ compliance. He also has a thing or two to say about the company’s recent acquision by Intel, multicore, and a clarified focus for the future. Edited excerpts follow.

MIL EMBEDDED: ***Wind River recently announced its VxWorks MILS Platform 2.0, which is now listed on the NSA-operated National Information Assurance Partnership’s (NIAP’s) website[1] as “In Evaluation” for compliance to High Robustness and Common Criteria Evaluation Assurance Level (EAL) 6+. What was the impetus?***

BROWN: We’ve got a pretty long history with products that have been safety-certified. And there are quite a few people using VxWorks 653 as well as our certified VxWorks products for safety-critical applications such as avionics. VxWorks MILS Platform 2.0, though, is really our first widely available platform for highly secure environments. It’s designed for cases where there are different levels of data, whether it’s data at different classification levels – maybe unclassified, classified, secret, or top secret – or data from coalition partners, for example, and these need to coexist securely on the same hardware system. There are many more applications these days that have security requirements on how to handle that data. Unfortunately, for many organizations today, the only way they know how to deal with different levels of security is to use redundant hardware or systems. Instead of having one computer on your desk, there might be two or three computers, each based on different security levels and the actual hardware line that they plug into.

VxWorks MILS provides a development environment that allows systems with multiple levels of security to be collapsed into a single system, and the operating system itself provides the data separation required to deal with these multiple levels of security or data classification. This reduces the costs for our military: For example, soldiers on the battlefield might need to carry two or three mobile handheld radios, each handling communications at different security levels, but now they can use one single radio that provides the same secure communication services. Separation is becoming a mandatory requirement these days.

> “… There are two or three requirements that actually need to be added into the “high robustness” definition to fulfill the requirements of EAL 6.”

MIL EMBEDDED: ***How have NSA and NIAP played into VxWorks MILS 2.0’s development?***

BROWN: We’re working very closely with NSA and NIAP, certainly, to do our evaluation. We’ve really tried to take a different approach from some other companies in that we’ve developed the MILS separation kernel from the ground up to be compliant with the U.S. government protection profile for separation kernels[2]. We’re not tweaking or trying to force any new requirements into an existing product.

We’re working closely with NSA and the University of Idaho on the formal methods evaluation. The end goal is that once someone buys this particular environment, if they need to change the hardware, a re-evaluation will be as cost-effective as possible in that they won’t have to recertify the entire system. They’ll be able to reuse certification artifacts for the majority of the system, and only need to re-evaluate the lowest layers of software that have been affected by the change in hardware.

MIL EMBEDDED: ***Tell me a little more about VxWorks MILS 2.0 itself.***

BROWN: VxWorks MILS 2.0 provides a Separation Kernel (SK) as well as a number of other key capabilities, but I won’t go through all of them now. VxWorks MILS 2.0 is really focused on providing key SK capabilities, interpartition communication capabilities, and networking

[1] NIAP’s listing of validated products can be found at: www.niap-ccevs.org/cc-scheme/vpl
[2] NIAP’s certification documentation can be found (and subsequently downloaded) from: www.niap-ccevs.org/cc-scheme/st/vid10119

capabilities. This release supports the Power Architecture. Other architectures will be supported in the near future.

MIL EMBEDDED: ***What's the significance of the SK technology in VxWorks MILS 2.0?***

BROWN: I think one thing of significance to note with VxWorks MILS 2.0, at least from a technical perspective, is that VxWorks MILS 2.0 was built using a Type 1 hypervisor technology. There are two approaches: You can either go with a Type 2 hypervisor, which basically utilizes an operating system to act as a hypervisor. Or you can use a Type 1 hypervisor that's a separate technology that's very lightweight, very small, and provides the fundamental services. In this case, it provides the services necessary to be compliant with the Separation Kernel Protection Profile (SKPP), suitable for EAL 6+ certification levels, but provides for better system performance than a Type 2 hypervisor-based system. Our VxWorks 653 platform uses Type 1 hypervisor-based virtualization to achieve very high performance, even with dozens of user-mode partitions.

In addition, we have the VxWorks guest operating system that can be utilized in the partitions, really giving VxWorks users great API compatibility so they can take their applications and plug them into a MILS environment. And at that point, they've got these applications that they can now reuse in a fully secure environment.

MIL EMBEDDED: ***When you say a Type 1 hypervisor is "very small," how many lines of code are we talking about?***

BROWN: I don't think we've released the number of lines of code, but it is definitely in the lower thousands of lines of code.

MIL EMBEDDED: ***What role does the SKPP play in EAL 6+ certification?***

BROWN: The SKPP provides guidelines for developing a separation kernel, which is what actually goes through certification. We have a developed a separation kernel based on the NSA-approved SKPP that will go through the EAL 6+/high robustness certification.

MIL EMBEDDED: ***How is this different or similar to how Green Hills got its INTEGRITY-178 RTOS certified?***

BROWN: It's different. Green Hills took a Type 2 hypervisor approach. And they only certified under "high robustness"; they did not certify under EAL 6+, as they did not actually add in the necessary requirements to comply with EAL 6. We're actually taking a very similar Type 1 hypervisor-based approach as LynuxWorks, but we think it's important for us to be certified under both the common criteria EAL 6 and also under "high robustness."

MIL EMBEDDED: ***You're telling me certification to EAL 6+ is different from being certified and categorized as "high robustness." Those two terms are not analogous[3]?***

BROWN: That's correct. They're very similar, but there are two or three requirements that actually need to be added into the "high robustness" definition to fulfill the requirements of EAL 6. One of the reasons Green Hills is not listed on NIAP's website under systems evaluated at EAL 6+ is that there are certain requirements that have to be satisfied in addition to developing in compliance to the SKPP.

MIL EMBEDDED: ***When will VxWorks MILS 2.0 be certified?***

BROWN: Our current plan is for certification to be accomplished by the end of 2011.

MIL EMBEDDED: ***Wind River's ARINC 653 partitioned OS is a completely different kernel from the one we're talking about now, correct?***

BROWN: Yes, exactly. They're both based on virtualization technologies, but VxWorks MILS 2.0 was designed to comply with the guidelines in the SKPP to achieve certification to EAL 6+ and "high robustness."

[3] An interview with Green Hills CTO Dave Kleidermacher can be found at: www.vmecritical.com/articles/id/?3747

MIL EMBEDDED: ***By now everyone knows about Intel's recent acquisition of Wind River. Can you make any comments in respect to what we've talked about?***

BROWN: There's intent for Wind River to maintain its independence and continue driving forward with its defined plans. I think we've got some really strong roadmaps and strategies laid out – for various architectures such as Intel, PowerPC, ARM, and MIPS – that span the next 18 months to 3 years that we're going to continue to drive forward.

MIL EMBEDDED: ***About five years ago, we were hearing lots of awful things about the platform strategy that Wind River had launched, which was confusing a lot of people. There was a transition out of Tornado. There was also the "Hell no, we'll never do Linux," then "Guess what – now we're doing Linux." But I think Wind River is well sorted at this point in time. Where is your focus now?***

BROWN: We've clarified our focus much more. It's great because we've got broad coverage across many vertical markets. And then some of the new technology disruptions, especially multicore, have been really good for Wind River in that they have allowed us to apply a lot of the technological history that our company's built up. It's amazing to me when I hear about processors in the works with hundreds of cores versus the dual- or quad-core processors that people are already struggling with. I think customers will need operating systems that can somewhat shield them from the complexity of the processor while also being optimized. That's a big area that we're highly focused on.

MIL EMBEDDED: ***With FPGA companies focusing on interfacing with lots of different cores, might it seem natural to Wind River that operating systems like VxWorks are going to have to find their way a lot more formally into FPGA-based and multicore designs?***

BROWN: Yes, definitely. It's been amazing for us to watch some of the new processors coming to market with dedicated acceleration engines – built into a multicore design – that are very vertical market-centric. They're certainly useful in the networking and wireless markets, and we have heard that the industry is going to see more and more multicore chips with FPGAs. If there's an FPGA on the silicon, you can basically tailor that chip for any particular vertical market. So it's only natural that an operating system such as VxWorks or Linux actually will be able to support that.

Marc Brown *is Vice President, VxWorks Product Strategy and Marketing for Wind River Systems. He has more than 18 years of experience in the development and deployment of highly technical mission-critical systems. He also served as VP of Product Marketing and Strategy for Borland Software and held senior management positions at IBM and Rational and technical positions at Motorola and Corning. He can be contacted at marc.brown@windriver.com.*

Wind River
510-749-2238
www.windriver.com

kontron

» We Do Not Build Navy Ships «

Our Customers Do.

At the heart of the most advanced command and control systems deployed in Navy ships are Kontron Military Rugged COTS boards and systems. Keeping us safe with smart applications, military contractors look to Kontron for superior technology, performance and life cycle management expertise.

MILITARY RUGGED COTS

Call, Email or Visit today.

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

If it's embedded, it's Kontron.

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

10 GbE in net-centric warfare: Why commercial network cards can't drive the application

By Rob Kraft

10 GbE provides a standards-based "fat pipe" to move data. However, real-time applications present unique challenges that must be addressed at the outset.

The world of technology has flipped 180 degrees: No longer is the best technology developed for the military and then later commercialized. Now, the defense department is playing catch-up to off-the-shelf technology. These days, civilians are accustomed to a profoundly high level of networking connectivity, enabled largely by Ethernet and Internet Protocol (IP). The military wants the same.

The appeal of Ethernet-based networking in the digitized battlefield is clear. And with the advent of 10 GbE, Ethernet is no longer limited to the command and control fabric in systems. It can also function as a "fat pipe" that reaches all the way down to high-bandwidth sensors. So, it's not surprising that 10 GbE is being designed into more and more defense applications, both in the main network and as the pipe to and from sensors or effectors.

But using 10 GbE is not without problems.

The fire hose crisis

Ever try to drink from a fire hose? It's not a matter of digesting or processing the water. You can't even ingest it. A processor faced with a 10 GbE pipe has the same problem.

It takes about 1 GHz of CPU processing per Gb of Ethernet traffic. CPUs fall quite short of bandwidth trying to run Ethernet protocol at 10 Gbps.

This problem also applies to commercial servers, but adopting commercial-world solutions for net-centric defense applications doesn't work. Let's look at why commercial network cards don't fit the bill.

The problem of handling bursts

High-performance real-time sensor applications, such as ELINT analysis systems found on platforms like Boeing's P-8A Poseidon (Figure 1), often need to sample in high-speed (hundreds of MSps) bursts lasting at least a few milliseconds. For a single channel, this leads to an incoming large data blast of at least 10 MB to 30 MB. On a 10 GbE link, this corresponds to receiving 1,000 to 20,000 back-to-back packets at line rate (depending on the MTU size employed).

Figure 1 | High-performance real-time sensor applications, such as ELINT analysis systems found on platforms like Boeing's P-8A Poseidon (pictured), often need to sample in high-speed (hundreds of MSps) bursts lasting a few milliseconds. Boeing photo by Ed Turner

If the receiving interface cannot absorb all this data into system memory, data is dropped. In a commercial server environment, dropped packets are resent. But in sensor applications, there is no time and/or no hardware facility to resend the data. Typical commercial 10 GbE interface cards, designed and cost-optimized for an environment where re-transmission of data is permitted and line rate bursts are rare, simply cannot address this problem.

Beyond protocol offload

In many real-time applications, offloading the transport protocol provides only part of the solution. For instance, COMINT or ELINT direction finding, network surveillance, and intercept are examples of applications that collect sensor data to synthesize multidimensional models of the environment. Such applications rely on the fusion of the sensor data. Bringing the data together requires accurate and precise time stamping of the individual data streams. Outgoing data in Electronic Countermeasures (ECM) or simulation systems must similarly be precisely time gated and synchronized to other events.

Precise time stamping and gating can only be performed through a deterministic interface. The CPU, vulnerable to software interrupt latencies and inconsistent access to buses, cannot provide the necessary time-stamp precision. Similarly, commercial Ethernet cards, with or without offload, get traffic from place to place. Time stamping is typically not a part of their functionality.

Time stamping is an example of additional special functionality required by these real-time applications. Another kind of problem occurs when offloading the protocol still results in data rates that are too high to be managed after they leave the interface. Take a multicamera GigE Vision-based situational awareness system that runs on top of standard UDP protocol or an intrusion prevention system guarding a network from viruses, Trojans, and other cyber threats coming from the WAN where the contents of packets are inspected at line rate. Both applications execute on payload data arriving at a 10 Gbps rate. So, offloading only the transport protocol, as a commercial offload solution would do, still leaves the CPU struggling to process data at 10 Gbps.

Making a mountain out of a molehill

Real-time digitized signals are destined for signal processing such as low or band pass filtering or error correction. These algorithms correct or compensate for small runs of errors, such as the error in a few consecutive data points. When faced with a large run of consecutive errors (such as when a large group of samples is missing), the algorithms break down.

This is particularly relevant for Ethernet data that transmits in packets with checksums and Cyclic Redundancy Check (CRC) fields that check for errors in the packets. Per the protocol standard, when errors are detected, even if they were caused by an error in a single data point, the entire packet is discarded. This means that as many as 9,000 consecutive bytes go missing, subsequently choking the signal-processing algorithm (Figure 2). If those few errors had instead been delivered to the signal-processing algorithm, they would not have caused a problem. So, the standard protocol stack behavior can make a mountain out of a molehill in a real-time signal processing system.

The real-time 10 GbE alternative

These problems can be circumvented with an intelligent 10 GbE interface. FPGA-based solutions are capable of the line-rate performance and flexible enough to run algorithms optimized for application requirements. Unlike CPUs with their sequential manipulation paradigms, data flowing into an FPGA cascades like

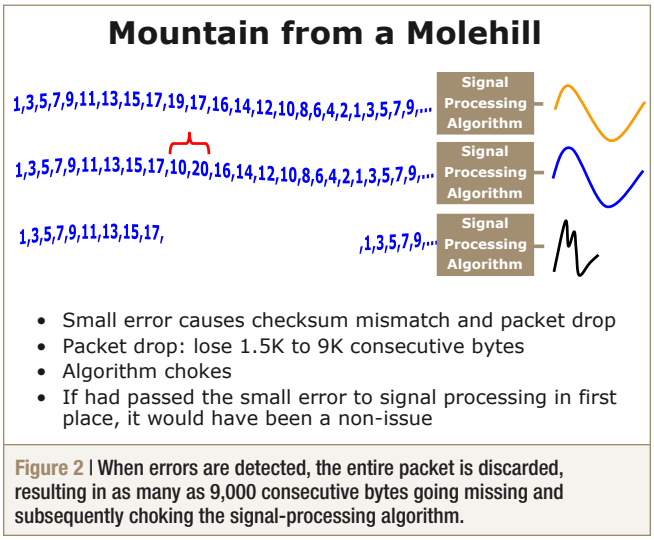

Figure 2 | When errors are detected, the entire packet is discarded, resulting in as many as 9,000 consecutive bytes going missing and subsequently choking the signal-processing algorithm.

a waterfall over massive amounts of parallel configured logic. In today's large FPGAs, the data flows over pipelines of hundreds of thousands of logic cells that can effectively be running at a few hundred megahertz. For the I/O, modern FPGAs are equipped with high-speed transceivers running at up to 11 Gbps. This allows them to interface to high-speed serializers or directly into the high-speed optical or copper communications interfaces. Using clever designs, even modern medium-sized FPGAs – containing, say, a hundred thousand cells – can fit multiple channels of real-time 10 GbE interfaces along with higher-layer application functionality.

A real-time 10 GbE system should include:

- Sufficient memory to accommodate the extended duration, line-rate bursts
- A time stamp and synchronization interface to deterministically and precisely stamp packets entering or exiting

These systems should additionally comprise algorithms optimized for real-time requirements that might include:

- A transport layer protocol offload tailored to the application, bus interface, and processor
- A process that alleviates CPU burden by offloading intensive application processing operations or inspecting and dropping packets that are uninteresting before they ever get to the CPU
- Modifying the standard transport protocol behavior to tag but not drop the packets received with checksum errors

No question, civilian COTS technologies – from the ubiquitous Internet to the gaming console in your living room – are finding their way into the digitized battlefield and changing both expectations and the nature and capabilities of net-centric warfare.

When it comes to using 10 GbE in these high-performance real-time applications, it is not enough to simply grab the mass-market technologies off the shelf. Instead, while understanding and respecting the unique requirements of net-centric warfare applications, we need to tailor the implementation of the standardized 10 GbE interface to meet these real-time needs.

Rob Kraft *is VP of Marketing at AdvancedIO Systems Inc. He has more than 13 years of experience in systems engineering and business roles in the embedded real-time computing industry. Prior to joining AdvancedIO, he worked at Spectrum Signal Processing and AlliedSignal Aerospace. Rob has an MASc in Electrical Engineering from the University of Toronto. He can be contacted at RKraft@advancedio.com.*

AdvancedIO Systems Inc.
604-331-1600 • www.advancedio.com

U.S. Air Force photo by Rich McFadden

Creating a simulated environment for UAS operator training

By Yannick Lefebvre

The military's increased worldwide use of Unmanned Aerial Systems (UASs) is driving the demand for highly realistic training environments. Using a mix of integrated 2D and 3D COTS graphics tools can greatly accelerate the development and realism of simulators to fulfill this growing need.

Unmanned Aerial Systems (UASs) are increasingly being used by military forces around the world. Whether for reconnaissance work or tactical engagements, unmanned vehicles are remotely controlled from a ground station and provide extraordinary capabilities in range and agility while eliminating risk to human life. To execute tasks, operators need to be trained in a variety of skills, from the interpretation of visual information that the UAS provides to reacting correctly based on the situation at hand.

Simulating a UAS ground station is complex and requires many different technologies to work together to produce the final result for trainees. When creating a simulated environment for operator training, critical elements such as the UAS's sensor output and graphical overlays that display data used to pilot or control the vehicle need to be delivered in an intuitive way to ensure an effective training environment is provided.

From a system design perspective, it may be easier to simulate sensor outputs and graphical overlays on several different screens using disparate technologies. From a trainee's perspective, this approach can create unnecessary workload and the significance of an overlay's exact positioning over the UAS's sensor feed could be lost. The optimal solution is to merge both the UASs sensor feed and its visual overlays into a single screen using a Commercial Off-the-Shelf (COTS) 3D image generator with a COTS 2D human-machine interface modeling tool to ease the process of merging these applications and provide a more effective and intuitive UAS training system (Figure 1).

Matching requirements with technology

COTS 3D image generation systems are available to simulate UAS sensor feeds. This type of application can render synthetic animated environments on-screen such as terrain, buildings, moving vehicles, atmospheric effects, and any other details that make a virtual world come to life. Beyond receiving continual product updates over the lifespan of a project and benefiting from a proven framework,

Figure 1 | A simulated 3D camera view with 2D static and interactive graphical overlays.

opting for a COTS image generator brings additional functionality that elevates the realism of a UAS ground station simulator such as simulating high-fidelity sensors (night vision goggles and infrared, for example). In contrast, developing such advanced visual features in-house is cost-prohibitive, time-consuming, and quite challenging as it requires a very specific expertise.

For the 2D overlays that are presented to the operators on top of the sensor feed, such as a Heads-Up Display (HUD) or a targeting reticule, a COTS Human-Machine-Interface (HMI) design software is best suited as it enables developers to visually define the elements of the overlay along with their respective functionality, data ranges, and information sources without needing to write code manually. Once a design is finalized, an automatic code generator packages the display contents as a self-contained executable program. Using a COTS HMI design software for this part of the development introduces advanced authoring concepts such as graphical logic creation and accelerates the development cycle by using a graphical user interface instead of manual coding.

Integrating 2D and 3D graphics

Once the 2D overlays and 3D graphics have been developed, they need to be integrated into one cohesive method to provide the trainee an integrated view that accurately represents the real-life equipment.

For this to be possible, both the 2D HMI design software and the 3D image generator need to use the same graphical language to draw graphics on-screen. If they don't, it will be difficult – if not impossible – to integrate them within the same environment. In today's desktop and embedded computing world, the graphical language of choice is usually OpenGL.

Another concern that needs addressing is that the graphics code produced by the HMI design software needs to adapt itself to being rendered in a larger environment and not clear the screen buffer when it renders its graphics so that the underlying 3D graphics don't get erased.

One last consideration before starting the integration of these two technologies is

> "... Opting for a COTS image generator brings additional functionality that elevates the realism of a UAS ground station simulator such as simulating high-fidelity sensors ..."

the authoring of the 2D overlay graphics. Since they will be rendered on top of the 3D environment with transparent elements, graphical techniques such as masking (using visual elements that are the same color as the background as masks) should be avoided as they would produce undesirable visual artifacts.

There are two common approaches that can be followed to perform this integration. The first and more manual method is to take the generated code from the HMI design tool and call it from a post-draw function within the 3D image generator. A post-draw function is a programming hook that allows users to render their own graphics once the 3D virtual world has been completely displayed. This integration technique requires skilled programmers who are familiar with both technologies at hand.

An easier approach that is provided through the use of COTS HMI design software is to encapsulate the 2D graphical overlay as a Dynamic Link Library (DLL) plug-in with a defined communications interface and to load the resulting component in a COTS 3D environment that supports external plug-ins and offers a visual configuration utility. In addition to simplifying the integration of the visual overlays with the UAS sensor feed simulation, this technique makes it easy to iterate through the development of an application by only having to generate a new plug-in file after making changes to the overlay graphics.

Once the display code is loaded in the 3D display, it needs to be animated. Most often, this is done by sending data to the 2D overlay using a communications structure that is defined in the HMI design. Once data is assigned to the interface, the display elements that are linked to incoming variables automatically update

themselves to reflect new values. Figure 2 depicts the 3D rendering pipeline with 2D overlay integration.

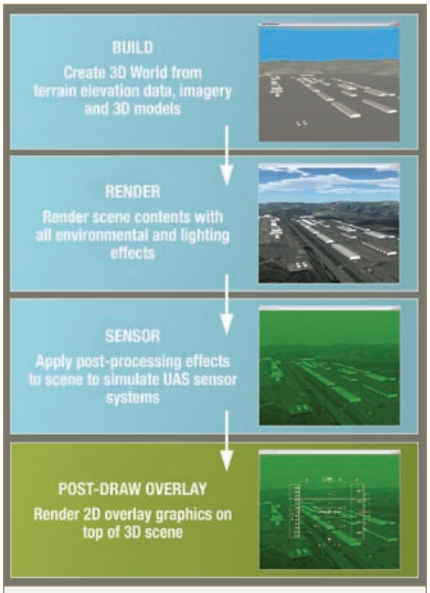

Figure 2 | 3D rendering pipeline with 2D overlay integration.

In terms of data transport mechanisms, information can be sent by using API functions defined in the DLL plug-in or through direct communication between the overlay code and external data sources. The first method would be used if data displayed by the overlay is shared with the 3D scene (for example, pitch, roll, and altitude). The second technique is used in situations where the overlay data is hosted outside the visual system (for example, targeting reticule or engine data) and needs to be transported over protocols such as UDP, TCP, or shared memory.

The benefits of an integrated solution

The use and integration of COTS 2D and 3D graphics is the best way to create an environment that is closest to the real equipment that trainees need to become familiar with in a UAS ground control station. The need for dynamic, intuitive and realistic training environments continues to increase in tandem with the growth of the UAS market. Developers play a key role in the creation of accurate and immersive training programs and must understand how COTS technologies such as HMI design software and a 3D image generating system help them deliver high-quality simulators efficiently. By eliminating the need for hand-coding and using tools that offer advanced logic authoring and the ability to easily integrate 2D and 3D, developers will stay one step ahead of the market curve.

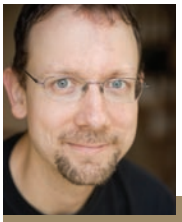

Yannick Lefebvre *is a senior application developer at Presagis. With a background in computer sciences and 12 years of experience in modeling and simulation, Yannick has provided counsel on hundreds of simulation and embedded display programs globally and is considered an expert in the industry. He can be reached at yannick.lefebvre@presagis.com.*

Presagis
514-341-3874
www.presagis.com

Data-centric architecture for military embedded systems

By Dr. Stan A. Schneider

Publish-subscribe middleware is the key technology that enables data-centric architecture. The DDS standard precisely defines how the middleware controls and transmits information. Successful applications include many shipboard and UAV systems.

Modern networked applications must connect and coordinate many resources. A typical distributed application integrates CPU-intensive servers, data stores, user interfaces, and real-time sensors and actuators. Each of these applications demands a different set of information, at varying rates, with varying urgency, and with varying reliability.

Getting the right data to the right place at the right time in this complex environment is perhaps the greatest integration challenge of the new networked reality. The "traditional" way to connect these disparate systems is to first examine the data requirements of each device or processing endpoint, and then design point-to-point interfaces between those devices. Once the relationship has been established, the data can be passed between those two devices, likely using direct messaging. If many devices need the same information, a client-server design allows distributed access.

However, this design assumes the network is relatively static, servers are always present and accessible, the server/client relationships are clear, clients know where and – most importantly – when to request data, and all nodes have similar delivery requirements. This "server-centric" design quickly breaks down as complexity grows.

Data-centric design offers an alternative. With data-centric design, developers specify only the data requirements, inputs, and outputs of each subsystem. Integration middleware discovers the producers and consumers of information and provides the data immediately when needed. This design greatly simplifies integration of systems with complex data requirements. Driven by the rapid adoption of the Object Management Group (OMG) Data Distribution Service (DDS) standard[1], many fielded systems – including shipboard systems like Aegis, UAVs like Scan Eagle, and base stations like the Advanced Ground Control Station for Predator – are adopting this DDS and data-centric design, empowered by publish-subscribe middleware.

Data-centric design

Instead of focusing on endpoint applications (devices, processing nodes, console applications) and how they individually interact, the data-centric approach begins the design process from an information perspective. What data does this application produce? What does it need? When? These questions decouple the implementation of the application from the other parts of the system. Data-centric design can greatly simplify system concepts.

Thus, a data-centric designer first defines an information model that captures the essential state and events in the system, then creates data input and output specifications, and then develops components that can produce and process that information. Rather than deriving specific data-interface requirements between components, the designer determines how to represent the state of the system and the external or internal events that can affect it. This "data model" captures the essential elements of the physical system as well as the processing logic. The model decouples applications; data can be provided by any (authorized) process and used by any other process. Applications must specify when and how they can supply information, but they do not need to know when or where that data might be used.

The enabling technology for data-centric design is publish-subscribe messaging, often shortened to "pub-sub." In this model, data sources, or producers, publish data into and subscribe to data from an "information bus" or "cloud" (Figure 1). Note that although there is a concept of data "in the cloud," it is a virtual concept. The actual data exists only in the publisher and subscriber endpoints. The pub-sub system connects endpoints by sending messages from the publishers to subscribers over a variety of transports, including direct-memory transfers, switched fabrics, or multicast or unicast over Ethernet. Transports, operating systems, and other location details do not need to be known, decoupling the design and allowing adaptation to performance, scalability, and fault-tolerance requirements.[2]

On a ship, for example, a GPS receiver can publish position data. Navigation computers and targeting systems can all

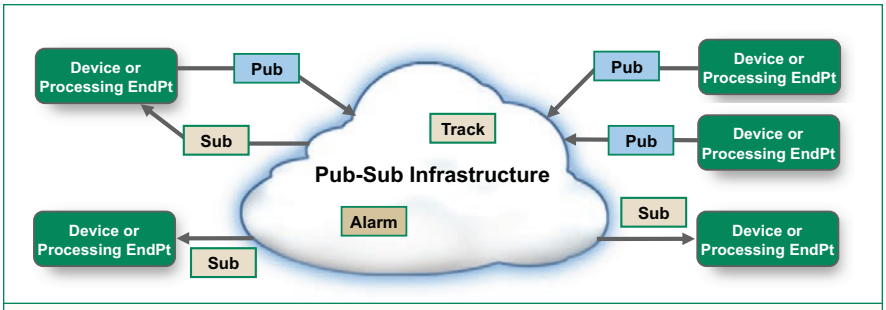

Figure 1 | Data-centric design revolves around the information itself. The information model captures the essential state and events in the system. Components are then built to interact with the information model "cloud," rather than with each other directly. The pub-sub infrastructure connects all the pieces.

subscribe to the GPS location data. The publications go directly from the GPS to the targeting system, even though both conceptually simply publish or subscribe to the "cloud."

By converting data within the middleware, publish-subscribe models can also connect systems with unmatched data formats. By enforcing quality of service parameters such as timing specifications and buffering, pub-sub models can connect systems with disparate delivery requirements, even trading off delivery reliability with timing constraints.

But the most important advantage of the publish-subscribe approach is decoupling. Since only the data interactions are specified, devices can be upgraded or added without the need to change code and exhaustively retest every configuration. If new data is available on the network, other devices might require additional code to make use of that data, but in practice this is significantly simpler than modifying and testing a large number of specific point-to-point connections.

Decoupling also makes distributed applications highly scalable. Because there are no fragile point-to-point data connections and devices can be added with little or no change to underlying code, expanding the application to include a larger network with more endpoints is simplified.

Designing with data-centric principles

Data-centricity provides a guide for how to design distributed applications in general. Many system architects of distributed applications today use procedural or object-oriented principles in creating the fundamental design, often using UML sequence, class, and state diagrams.

These design methodologies tend to treat transporting and consuming data as second-class citizens, focusing instead on the step-by-step processes by which devices make computations and produce actionable results. A data-oriented methodology, on the other hand, focuses on the flow of data through the application.

In general, the tenets of data-oriented programming include the following principles:

- Expose the data. Ensure that the data is visible throughout the entire system. Hiding the data makes it difficult for new processing endpoints to identify data needs and gain access to that data.
- Hide the code. Conversely, there is no reason for any of the computational endpoints to be cognizant of one another's code. By abstracting the code, data is free to be used by any process, no matter where it was generated. This allows for data sharing across the distributed application

and for the application to be modified and enhanced during its life cycle.

- Separate data and code into data-handling and data-processing components. Data handling is required because of differing data formats, persistence, and timeliness, and is likely to change during the application life cycle. Conversely, data processing requirements are likely to remain much more stable. By separating the two, the application becomes easier to maintain and modify over time.
- Clarify data interfaces. Interfaces define the data inputs and outputs of a given process. Having well-defined inputs and outputs makes it possible to understand and even automate the interfaces to data processing code. By specifying exactly how and when components produce or consume information, new applications can join and interact without impacting the system function.
- Loosely couple all modules. With well-defined interfaces and abstracted computational processes, devices and their computation can be changed with little or no impact on the distributed application as a whole.

These and other principles are summarized in Table 1, along with a comparison with object-oriented development tenets.

A data-centric approach simplifies design

Data-centric architecture decouples designs. It simplifies communication while increasing capability and easing system evolution. It especially simplifies developing distributed applications with complex components on remote network nodes, integrating new functionality into those applications and maintaining and changing the components independently.

Data-centric design requires a change in perspective. Rather than conceiving the system primarily as a collection of interacting programs, the system must be first regarded as an information model that supports applications that contribute and use information. Networking middleware enables this view by providing information anonymously and consistently with data delivery properties and (importantly) timing. The OMG Data Distribution Service networking standard is serving as the catalyst for adoption of this approach by many fielded systems including shipboard systems, aircraft, and UAV base stations.

References

[1] Gerardo Pardo-Castellote. "OMG Data-Distribution Service: Architectural Overview," Proceedings of the 23rd International Conference on Distributed Computing Systems, May 2003.

[2] Darby Mitchell. "Applying Publish-Subscribe to Communications-on-the-Move Node Control," MIT Journal, Volume 16, Number 2, 2007.

***Dr. Stan A. Schneider** is the CEO of RTI. His expertise is in architectures and tools for real-time systems. Before RTI, Stan managed a Stanford laboratory's intelligent mechanical systems, developed communication and computer systems, and researched automotive safety. He holds a PhD in EE/CS from Stanford and an MSEE and BS from the University of Michigan. He can be reached at stan@rti.com.*

Real-Time Innovations, Inc.
408-990-7415
www.rti.com

Object-oriented programming	Data-oriented programming
Hide the data (encapsulation)	Expose the data
Expose methods – code	Hide the code
Intermix data and code	Separate data and code
Define object interfaces	Agree on data labels and schemas
Invoke/Implement operations on objects	Send/Receive messages
Combined processing, no restrictions	Strict separation of parser, validator, transformer, and logic
Changes: Read and change code	Changes: Change declarative data definition
Tightly coupled	Loosely coupled

Table 1 | A comparison of data-oriented programming with object-oriented programming. The data-oriented approach enforces attention on the data rather than on the processes that manipulate the data.

MIL/COTS DIGEST

The Defense Electronic Product Source

October 2009

In This Issue

The modern military relies on comms gear

JDAMS, Hellfires, and Tomahawks are exciting to write about because they blow things up and rely on electronic guidance and/or fire control systems to propagate "smart" warfare. But communications networks, office environments, and COTS-based telecom and datacom equipment form the backbone of high-tech weaponry. These LANs and WANs connect sensors to shooters and their upper-echelon decision makers; without this gear, weapons and warfare would still be line-of-sight.

In this special *Supplement on PICMG and communications-related products,* we highlight some of modern warfare's unsung technology heroes. Databases need network-attached storage appliances; we've got one here for you. EMI shielding is essential to interconnectivity – we've got that too, along with a rugged deployable ATR chassis said to increase payload MBTF and solutions for those seeking IPMI management. The tasks performed by the products described herein range from pushing for near real-time analog-to-digital translation to shielding against radiation to cryptography to standing up to 15 *g*/11 ms shock. The PICMG products included here are AdvancedTCA (1), MicroTCA (2), Advanced Mezzanine Card (1), and 3U CompactPCI (2). In addition to the AMC mentioned, there are three other mezzanines, in PMC and newer XMC flavors.

Many of the new products in this *Supplement* were not originally specified for the battlefield. These products are the embodiment of COTS: designed for communications and data applications in central offices, cellular base stations, or edge equipment, yet they've been successfully adapted to fit aerospace and defense systems.

Chris A. Ciufo, Group Editorial Director

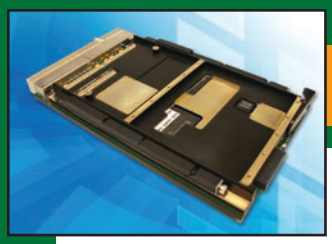

3U CompactPCI board at home in space

The S950-02 is an enhanced, 1 GHz version of Aitech's space-flown S950 3U CompactPCI radiation-tolerant SBC. It outpaces the S950 when it comes to increased processing power and data throughput. Its high-performance PowerPC 750GX runs at 1 GHz, and the board's power appetite can be satiated at less than 10 W thanks to a diet of Silicon on Insulator (SOI) PowerPC technology and conduction cooling. Designed for numerous space applications, the S950-02 can shield against Total Ionization Dose (TID) greater than 15 krad (Si). With a low SEU rate of less than one upset per 900 days of operation in LEO with considerations for the worst case solar flare and the South American Anomaly (SAA), the SB950-02 syncs with the needs of harsh, mission-critical systems. Redundant mission computers, flight guidance and navigation computers, command and data handling computers as well as solid state recorders, video controllers and manipulation controllers fall squarely in its bailiwick.

www.rugged.com

AITECH DEFENSE SYSTEMS

Analog-to-digital closes in on real-time

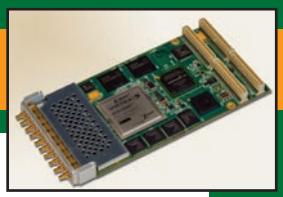

Suitable for CompactPCI carriers, this 400 MHz bandwidth, 14-bit ADC-PMC-module, dubbed the ICS-1556B, addresses software defined radio, signals intelligence (SIGINT), digital receivers, and tactical communications. If you are working with apps where you need to translate analog data into digital information in as close to real-time as possible, you will want to note that the ICS-1556B, with a Xilinx Virtex-5 SX95T FPGA onboard, lets you develop your own front end signal processing functionality and execute it at high speed to maximize overall system-level performance. Its four 14-bit ADCs sample synchronously at frequencies up to 400 MHz, and a fast (64-bit/133 MHz) PCI-X interface is included. The FPGA's signal processing capability makes possible standard functions such as wideband Digital Down Conversion (DDC), Fast Fourier Transform (FFT), and time stamping – or to implement any required functionality. GE Fanuc calls the ICS-1556B a solution for a broad range of signal processing applications that can be used with any type of carrier card that accepts a PMC module.

www.gefanuc.com/embedded

GE FANUC INTELLIGENT PLATFORMS, INC.

Tailoring MCHs to match mil and more

Working from the premise that MicroTCA has a home in the defense and aerospace, telecom, medical, and industrial markets – each with different packaging requirements – Emerson has introduced four MicroTCA Carrier Hub (MCH) variants. A base module includes management, telecom clock switching, and Layer 2 unmanaged Gigabit Ethernet switching to each of 12 payload slots, and the company has also created a lower-cost base model without telecom clocking. One variant has an additional PCI Express fabric switch complete with spread spectrum clocking functionality for those applications that need switched PCI Express connections and centralized clock distribution, also available in a non-telecom-clocking version that costs less. The U.S. Military has increased its emphasis on high availability, which is where the modules' SpiderWareM3 and IPMI management functions come in. Two Gigabit Ethernet fabric expansion/uplink RJ-45 faceplate connections are included. The full-size modules are designed to fit into the MCH slot of MicroTCA platforms.

www.emersonnetworkpower.com/embeddedcomputing

EMERSON NETWORK POWER EMBEDDED COMPUTING

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

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

Made in the USA

Ultimate Modularity

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

Direct Seamless Connections with no Data Reduction

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

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

PICMG® Power Connectors

MTCA.1 12 volt

New!

- QBH5W5 and QBH15W4 connectors are compliant to MTCA.1 R1.0 specifications
- Power contact carries 50 amps minimum at a 30°C temperature rise prior to derating.

MTCA.0 24 volt

- QBH9W4 series is compliant to MTCA.0 R1.0 specifications
- Power contacts carries 49 amps minimum at a 30°C temperature rise prior to derating.

MTCA.0 48 volt

- QB7W2 series is compliant to MTCA.0 R1.0 specifications
- Power contacts carries 40 amps minimum at a 30°C temperature rise prior to derating.

Features:

- low contact resistance
- minimized height above PCB
- uni port and dual port configurations
- cable connector options available

Positronic offers all µTCA® power input connector variants!

Widest Variety of Power Distribution Interconnects

Compact Power Connectors

P47 variant is compliant to CompactPCI® requirements. Five package sizes available. High current through a small package. A variety of termination styles are available.

High Density Combo-D Connectors

New!

High density connectors use contact sizes 22, 16, and 8, providing previously unavailable density. Professional, industrial and military performance levels. Offering signal, power coax, high voltage and air coupling contacts.

VP Series Connectors

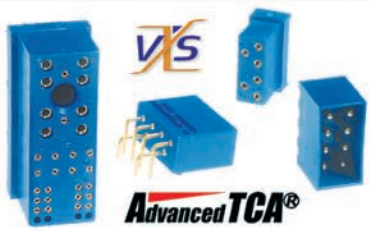

Dedicated power interface connectors for plug-in cards. Applicable connectors are compliant to ATCA Zone 1 and Vita 41 requirements. A variety of other contact variants and package sizes.

POSITRONIC®
GLOBAL Connector SOLUTIONS

Quality Solutions
Reliable Service

Springfield, Missouri USA • 800.641.4054 • info@connectpositronic.com

Visit www.connectpositronic.com for details!

U.S. Air Force photo/Senior Airman Julianne Showalter

CompactPCI solution targets rugged apps

By Jeff Munch

Jeff walks us through the features of a processor board design that was the result of efforts to stay within a desired thermal envelope yet boost performance.

One of the challenges in the architecture and design of CompactPCI processor board is staying within a "reasonable" thermal envelope. In general, "reasonable" is something under 50 W. Developers know, of course, what is available in their local computer stores, and have an expectation that solutions in the CompactPCI form factor deliver the same level of performance. This has led to somewhat of an arms race in CompactPCI processor designs as vendors strive to take the latest and greatest processor and chipset combinations and cram them onto CompactPCI designs. The results can be processor blades that approach 100 W. Fortunately, 45 nm process technology from Intel has resulted in a processor and chipset combination that will help CompactPCI stay within acceptable power levels while providing a significant increase in performance and functionality over earlier products. This article describes one example of a blade using 45 nm process technology to meet performance objectives without jumping outside the thermal envelope (Figure 1). The blade consumes less than 50 W when configured with a 2.53 GHz Core 2 Duo processor, 4 GB of memory, and a 100 GB SATA hard disk drive.

The Intel Penryn processor we used contains several advancements compared to the previous generation, most notably the 45 nm process technology, which enables a smaller, higher-performance die. The smaller die provides more room for a larger 6 MB cache. (In the previous generation cache was 4 MB.) The front side bus increases from 800 MHz to 1066 MHz. Intel Trusted Execution Technology (TXT) has been added. TXT is a hardware-based mechanism to protect against software-based attacks. Enhanced virtualization technology (VT-d1) improves the reliability and performance of virtualized environments.

Figure 1 | The cPCI-6880 blade from ADLINK uses 45 nm process technology.

Processor	Frequency	L2 Cache Size	Front Side Bus Speed	Number of Cores	Power Dissipated
Core 2 Duo T9400	2.53 GHz	6 MB	1066 MHz	2	35 W
Core 2 Duo P8400	2.26 GHz	3 MB	1066 MHz	2	25 W
Core 2 Duo T7500	2.2 GHz	4 MB	800 MHz	2	35 W
Celeron 575	2.0 GHz	1 MB	667 MHz	1	35 W

Table 1 | Intel Penryn processor selections available with the ADLINK cPCI-6880 blade

The Penryn continues Intel's Turbo Mode support, which allows the processor to run at speeds above the rated clock frequency as long as the processor stays within its thermal envelope. Table 1 details the Penryn processor selections available with the blade.

In addition to finding enhancements in the CPU, designers will also discover that the chipset used boasts an improved video subsystem. The graphics engine is fifth generation and runs at 533 MHz with up to 384 MB of shared memory. Hardware accelerated decode for MPEG2,

VC-1, and H.264 is provided in the chipset. Supported display interfaces include single- and dual-channel 24 bpp LVDS, and DVI-I. Display resolutions are up to 1600 x 1200 for dual independent displays. The graphics performance is 30-40 percent above previous generation chipsets. Improved 3D capabilities include support for DX10 and OpenGL 2.0.

Harsh environment ready

ADLINK has made the new blade available with an extended operating temperature range of -20 ºC to +70 ºC. To withstand high-vibration environments, the blade can be configured with just the 4 GB of solder-down memory, and it contains an on-card 4 GB USB based NAND Flash. The blade also supports a conduction-cooled option.

With up to 8 GB of system memory, the cPCI-6880 has a GM45 memory controller that supports two channels of memory. The blade routes one channel to 4 GB of solder-down memory and the second channel to an SO-DIMM connector that can support an additional 4 GB of memory. The ability to handle 8 GB of memory as well as its VT-d1 support matches up well with virtualization environments. Dual BIOS PROMs provide redundancy in the event of a BIOS corruption or a need to roll back the BIOS to a previous known good revision. Other peripherals include four PCI Express based 10/100/1000BASE-T Ethernet ports, one DVI-I port, one analog CRT port, three USB 2.0 ports, a serial port, an onboard hard drive, 4 GB of NAND Flash, and a PMC site. Additional peripherals are available on the rear transition module.

The blade can do its job residing in either a CompactPCI system slot or a peripheral slot and supports 5.0 V and 3.3 V backplane V (I/O). A PCI Express to PCI bridge makes it possible for a 66 MHz 64-bit PCI interface to be used for the PMC and bridged for use on the CompactPCI bus. It is fully compliant with PICMG specifications 2.0 R3.0, 2.1 R2.0, 2.9 R1.0, and 2.16 R1.0.

Looking ahead

The continuing need to avoid thermal excess across a wide range of environments, including military and aerospace arenas, means that blades that support a 2.53 GHz processor, yet don't cross that 50 W line, will find themselves in a ballooning market. If these same blades support 4 GB of solder-down memory, on-card NAND Flash disk, and conduction-cooled and extended operating temperature options, so much the better for an installed base that is always on the lookout for viable technology refresh.

Jeff Munch is CTO of ADLINK and heads all R&D operations in North America and Asia, and is responsible for building ADLINK's presence throughout the world. Jeff has more than 20 years of experience in hardware design, software development, and engineering resource management. Before joining the company, he spent five years at Motorola Computer Group as Director of Engineering. Jeff is also Chair of the AdvancedTCA Subcommittee and Interim Chairman of the PICMG COM Express Plug-and-Play Subcommittee. He brings a wealth of technical knowledge and experience to ADLINK's management team. He may be reached at info@adlinktech.com.

Ampro ADLINK Technology, Inc.
800-966-5200
www.adlinktech.com

3U CompactPCI switch/mil-aero

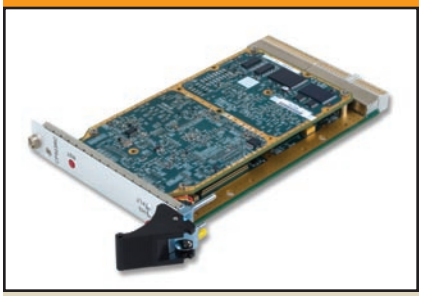

Available in rugged air-cooled and conduction-cooled versions for non-benign environments, the NETernity CP923RC 3U CompactPCI fully managed Layer 2/3+ Gigabit Ethernet switch is said by GE Fanuc to deliver optimum density. Ten 10/100/1000BASE-T ports route to rear I/O. It features OpenWare switch management software for ease of use. Network traffic switching can take place externally to (as well as within) a chassis. It is designed to enhance support for mobile devices and a significantly larger number of Internet addresses.

www. gefanuc.com/embedded

GE FANUC INTELLLIGENT PLATFORMS

CompactPCI blade for extremes

Soldered memory and an extended temperature range of -20 °C to +70 °C help peg the 6U cPCI-6880 as aimed at defense and aerospace apps. ADLINK has built the cPCI-6880 series around the 45 nm Intel Core 2 Duo processor with a 2.53 GHz core speed. This approach is key, says ADLINK, to achieving an optimum performance and total power consumption balance. The cPCI-6880 also features the latest Mobile Intel GM45 Graphics Memory Controller Hub.

www.adlinktech.com

ADLINK TECHNOLGY INC.

Ad hoc battlefield networks

Portable, secure in-the-field wireless connectivity is the point of the XMC 660, which can add trusted wireless communications to CompactPCI embedded systems. The company anticipates system designers will use the designed-for-rugged-use XMC 660 to connect portable PCs and wearable systems to access points within ground vehicles or base camps (or to other nodes) for ad hoc wireless networks and thus an infrastructure for conveying voice, video, and data among authenticated users.

www.cwcembedded.com

CURTISS-WRIGHT CONTROLS EMBEDDED COMPUTING

EMI hood for aerospace

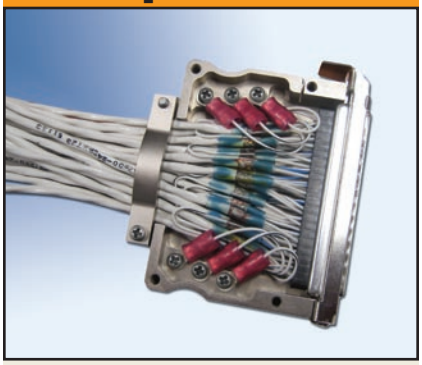

According to Positronic, the wall thickness of this higher-reliability lightweight D-sub EMI hood (backshell) for aerospace application is thick enough to be robust, but not so thick that it adds unnecessary weight. The cable opening is large enough to accommodate high-density wire bundles using twisted, shielded pairs. Up to six ground points reside inside the hood – four for shell sizes 1 through 3 and six for shell sizes 4 through 6.

www.connectpositronic.com

POSITRONIC INDUSTRIES

Guarding linked devices

Mocana's 5.1 Device Security Framework, according to the company, safeguards the integrity of connected devices. Both source and binary versions include full support for NSA's Suite B algorithms, so government agencies and contractors can obtain secure communications when classified "Suite A" algorithms are inappropriate or not available. The release comprises a series of components. One of is these is NanoCrypto, a miniaturized, device-optimized cryptographic engine that supports NSA Suite B cryptography, RSA, AES, and Elliptic Curve encryption.

www.mocana.com

MOCANA

AdvancedMC for an application-ready system

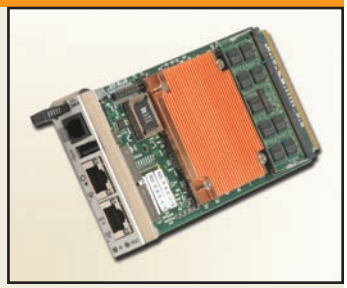

Targeting COTS-based defense, aerospace, telecom, and wireless applications, Performance Technologies' AdvancedMC compute module (AMC123) opted for the Intel EP80579 Integrated Processor, described by Performance Technologies as an innovative and highly integrated system-on-a-chip (SoC) processor that supports up to 4 GB DDR2 memory with ECC. This compute module option for the company's IPnexus Application-Ready Systems, which integrates a MicroTCA-based platform, other configurable AdvancedMC modules, and the NexusWare family of Linux software products, is said by Performance Technologies to address needs for cost-effective and modular processing in a low-profile, appliance-style form factor. This AdvancedMC also includes: 3 Gigabit Ethernet Ports 0 and 1; Dual SATA II Interface Ports 2 and 3 and onboard SATA Flash module up to 16 GB; USB 2.0; PCI Express x1, x4, or x8, Port 4-7, 8-11.

www.pt.com

PERFORMANCE TECHNOLOGIES

ATR chassis keeps its cool in more ways than one

The seven universal slots in the CM-ATR-35 high-performance ¾ ATR chassis handle CompactPCI, VME, and VPX conduction- or air-cooled 6U boards at up to 150 W per slot. Its sealed contaminant-free enclosure incorporates six air-to-air Heat Exchangers (SixHex series). According to the company the chassis significantly boosts payload MTBF and is also notable for weighing in at a fighting trim 13 kg. Other features of note include: its Integrated Temperature Supervisory Unit (TSU); an 800 W PSU that accepts all military standard input voltages; available mounting tray with quick release system; easy customizable front panel and flexible top and bottom I/O wiring. The CM-ATR-35 is tested per MIL-STD-810F and MIL-STD-461E and certified for immediate deployment in defense and aerospace systems.

www.cmcomputer.com

CM COMPUTER

Rugged family welcomes new member

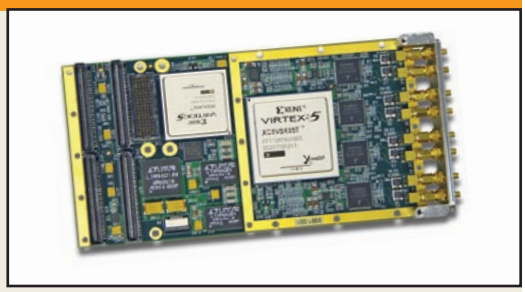

An addition to the company's REVX series for rugged embedded computing, this PMC/XMC mezzanine, the XMC-240, has four channels of analog input that can sample analog signals at up to 250 MSPS with 14 bits resolution. Developers will find it suitable for SIGINT, Software-Defined Radio, and radar requirements. The channels tie directly to a Xilinx Virtex-5 SX95T FPGA. High-performance DSP applications process the data captured from the A/D converters. Comprehensive clock management and either onboard clocking or external clock input are included. Linux and Wind River VxWorks development kits comprise a Software Development Kit (SDK) with all necessary software drivers and libraries plus an FPGA Development Kit (FDK), with FPGA cores supplied to interface to FPGA peripherals including memories, data interfaces, and analog interfaces.

www.nallatech.com

NALLATECH

NEXT GENERATION COTS-BASED SOLUTIONS

RUGGEDIZED QUALIFIED FIELDED

RadiSys is the industry leader in ATCA and COM Express. Our deep partnerships with Intel and other industry leaders will provide you with the latest COTS technologies that will put you ahead of the competition and win your next contract. Our products are ruggedized above and beyond industry standards, tested for extreme temperatures and integrated with ruggedized platforms to provide you a reliable solution that you can count on.

RadiSys has years of proven deployments and success stories in demanding telecom and medical markets. We look forward to doing the same for you so you can **BE RIGHT. NOW.**

RadiSys.
THE POWER OF WE
radisys.com/milaero

intel
Embedded and Communications Alliance
Premier

Cutting-edge backplanes, old-fashioned passion

At Bustronic, backplane design is more than just offering the latest technology. It's about the passion and perseverance that goes into each design. We continuously strive for the cleanest signals, superior quality, and more cost-effective solutions. When it comes to new architectures, we pride ourselves in being able to answer your toughest questions. Focused on backplanes, our team has the knowledge, ingenuity, and drive you need to make your design a success. Learn how we can put our experience to work for you. Contact Bustronic today.

bustronic
Elma Bustronic

Tel: 510.490.7388 www.elmabustronic.com info@elmabustronic.com

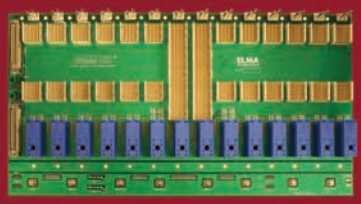

AdvancedTCA – Dual Star

AdvancedTCA – Mesh

MicroTCA – Star, Dual Star

MicroTCA – Cube, Pico

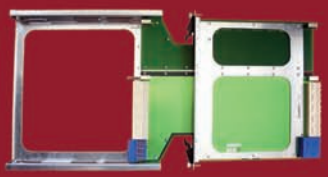

ATCA & MTCA load boards, extenders, and accessories

ELMA
Your Solution Partner

Rugged PICMG 3.0 chassis

Not easily shaken (or stirred for that matter) the 13U Gemini, 13U KBWD, and 12U KBWD-2 are rugged chassis that conform to MIL-STD-810 Method 514.5 and MIL-STD-167 Type 1, para. 5.1 for vibration as well as MIL-STD-810 Method 513.5 for acceleration. The Gemini is built to PICMG 3.0 Rev. 2, while the KBWD adheres to PICMG 3.0 Rev. 3, and both come in dual-star 14-slot versions or in mesh 3-, 5- and 7-slot configurations. KBWD-2 (PICMG 3.0 Rev. 3) has a dual 6-slot design, with each electrically isolated, full-mesh backplane incorporating four node slots and two hub slots. Dual redundant fan trays on the KBWD models shoo away up to 300 W per slot, aided by RTM cooling. The Gemini dissipates 200 W per slot using a single fan tray with RTM cooling available.

www.lcr-inc.com
LCR ELECTRONICS

High-speed storage

The CP-ASM3-RAID fanless CompactPCI RAID server combines CompactPCI technology with high-speed SATA connections via the backplane. This modular high-speed RAID array comprises up to 8 SATA-II hard disks (SATA 3.0 Gbps) or Solid State Drives and according to Kontron excels with high availability and high shock and vibration resistance. Secure protection against humidity and dust is afforded by the protective coating on the server's electronic components. The system is fanless. Kontron notes a long life cycle in operating temperatures of 0 °C to 55 °C is possible for the CP-ASM3 even when using standard hard disks.

www.kontron.com **KONTRON**

Extended temperature PMC withstands shock and vibration

The P507 has 4 UART interfaces providing a 60-byte transmit/receive buffer and can be used on all PMC-compliant host carrier boards in any type of bus system independent of the application, or on any standalone SBC. The PMC's temperature range is -40 °C to +85 °C, and it is prepared for coating against dust and humidity. It can withstand 15 *g*/11 ms for shock, 10 *g*/16 ms for bump, and 1 *g*/10 Hz to 150 Hz for vibration (sinusoidal). Independent 500 V isolation per channel makes data rates of up to 921,600 bps possible.

www.menmicro.com **MEN MICRO GMBH**

No data loss as MIL-S-901D Grade A and MIL-S-167 are met

The RAIDStor NAS Appliance is a compact, rugged platform from ACT/Technico, which recently became part of the Elma Electronic Inc. Systems division. Able to withstand high shock and vibration without losing data or sustaining hard drive damage, this 2U convection-cooled platform gathers data with high reliability and redundancy in harsh military environments. Elma describes the platform as a match for any application requiring a mission-critical data repository where data preservation is essential, including distributed communications processing.

www.elma.com

ELMA ELECTRONIC (ACT/TECHNICO)

Compute blade at home in mil

PDSi describes the CP86-N1 processor blade as the answer for mil-aero and telecom OEMs seeking configurable, scalable, high-reliability CompactPCI solutions. It's built around the Intel 45 nm Penryn Core 2 Duo processor and server-grade "Eagle Lake" chipset (5100 MCH/ICH9R) supporting ECC memory. The blade includes a standard PMC/XMC site for I/O expansion and is offered in two alternate models, one featuring an onboard SATA hard drive plus high-resolution graphics, the other providing a second PMC expansion site.

www.pdsi.com

PINNACLE DATA SYSTEMS INC.

MicroTCA rack saves space

Designed for high availability, the EDGE2000_RACK is a 10-slot MicroTCA compliant platform designed to rack mount up to five EDGE2000 MicroTCA systems. It fits space-constrained environments such as those found in avionics and other applications, including enterprise applications. The EDGE2000_RACK takes up only 4U. Mounting is in either direction and airflow is front to rear. The platform is fully compliant with MicroTCA.0, AMC.0, AMC.1, AMC.2, and AMC.3, and the company describes pricing for the EDGE2000_RACK as "price aggressive."

www.armvida.com

ARMVIDA, LLC

THE XTECH FILES

Case #4 2hp AMC Extender

Case #2 Hybrid AMC

Case #3 Adjustable

Case #1 FMC Solution

Case #5 Reinforced MicroTCA Front Panels

CASE SOLVED

Challenge:

XTech salutes military call for robust, reliable front panels

Leveraging its AMC expertise, XTech engineers have developed a family of robust, reliable MicroTCA front panels that can be used in military and aerospace environments—where shock and vibration are always a factor.

The XTech front panels feature front-locking attachment flanges with M3 captive screws, top and bottom, which are inserted into a lightweight extruded aluminum profile. The reinforced fasteners offer increased security while mobile vehicles encounter rough terrain or turbulent air.

XTech provides tools to enable designers to easily add modifications or cutouts and labeling to the front panel— speeding prototyping and delivery time. The robust front panels are fully assembled and are available in the full range of AMC sizes.

XTECH

www.XTech-Outside.com

781-963-7200

XTech – *The problem solver*

XTech is a leading global supplier of front panels and mechanical solutions for circuit board technology and other embedded systems.

XTech offers you in-house design assistance, engineering support, finishing and labeling options, component kitting, and management for your supply chain.

Discover XTech's fast turnaround capabilities, variable manufacturing efficiencies for prototype through production, and precision manufacturing in extruded aluminum, steel and plastic.

Executive Speakouts

Leveraging a ruggedized AdvancedTCA platform for command and control

By John Long

Ruggedized AdvancedTCA standards-based solutions provide both reliability and high performance under extreme conditions. RadiSys, the industry leader in COTS-based systems, has a history of proven deployment in exacting industries. Leveraging its pioneering engineering and manufacturing capabilities, RadiSys offers the Military, Aerospace, and Government (MAG) market innovative solutions featuring the latest next-generation technology. One such solution is the Promentum® ATCA-4500. This Single Board Computer (SBC) combines high performance with large memory capacity and expansion flexibility, enabling MAG customers to leverage a ruggedized AdvancedTCA platform for command and control as well as network appliance solutions that require high bandwidth. The ATCA-4500 also features an innovative virtualization solution that streamlines the complexities of network element configurations and eases the porting of legacy application software and operating systems. The ATCA-4500 has passed the official certification of VMware ESX 4.0. In addition, RadiSys' leading portfolio of Promentum® ATCA products are already being deployed in a multitude of communications networks under stringent reliability environments such as NEBS.

RADISYS CORPORATION **www.radisys.com**

Three rugged AdvancedTCA chassis bolster design flexibility

By Tom Malek

LCR Electronics' new AdvancedTCA systems rugged chassis offer increased design flexibility in a wide array of industrial, commercial, and rugged military environments. All of the new models conform to MIL-STD-810 Method 514.5 and MIL-STD-167 Type 1, para. 5.1 for vibration as well as MIL-STD-810 Method 513.5 for acceleration, making LCR's AdvancedTCA systems ideal for airborne, shipboard, and mobile applications. The Gemini, built to PICMG 3.0 Rev. 2, and the KBWD, built to PICMG 3.0 Rev. 3, are both 13U chassis available in dual-star 14-slot versions or in mesh 3-, 5-, and 7-slot configurations. Also built to PICMG 3.0 Rev. 3, the 12U KBWD-2 features a dual 6-slot design, with each electrically isolated, full-mesh backplane featuring 4 node slots and 2 hub slots. Each uses an Intelligent Platform Management Bus (IPMB). The chassis can be used as stand-alone units or mounted vertically in a standard rack system.

LCR ELECTRONICS **www.lcr-inc.com**

Introducing a new PICMG 2.30 CompactPCI PlusIO standard SBC *and* a new family of rugged CompactPCI SBCs

By Stephen Cunha

MEN Micro's F19P is the first 3U Intel-based Single Board Computer (SBC) that conforms to the new PICMG 2.30 CompactPCI PlusIO standard. Available with Intel processors ranging from 1.20 to 2.26 GHz, and dissipating between 5.5 W and 25 W, the 32-bit/33 MHz SBC was developed for low-power embedded applications with high performance and reliability requirements. An integrated graphics controller makes the F19P especially suited for not only applications needing high graphics performance, but also for monitoring or control as well as test and measurement applications. The F19P's specially developed heat sink and soldered components meet shock and vibration requirements to enable the board's use in mobile applications. The F19P includes serial functions on the J2 rear I/O connector in CompactPCI systems for added flexibility and versatility as well as easier system upgrades. It can be used in CompactPCI-only systems and in hybrid systems using both CompactPCI and CompactPCI Plus.

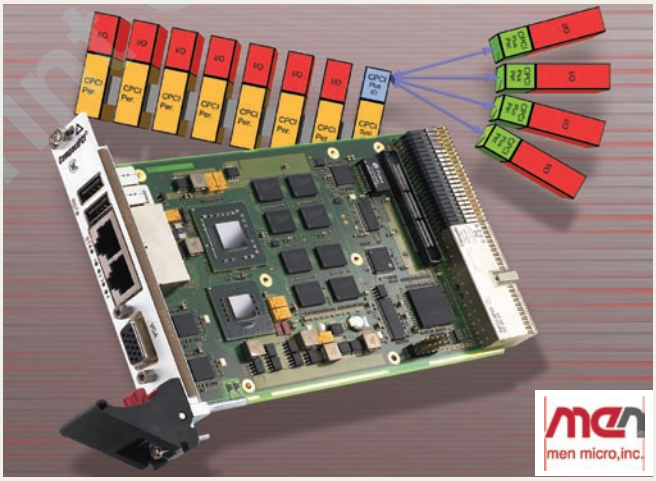

MEN Micro's new rugged 3U CompactPCI family of two SBCs and a rack provide reliable performance in temperatures as extreme as -40 °C and +85 °C. In addition to the convection-cooled F50P SBC, the new line includes the conduction-cooled F50C that easily fits into the new conduction-cooled rack. The new rack accommodates readily available, less costly 3U cards designed for ventilated systems when mounted within conduction-cooled adapter frames. Both SBCs offer the choice of either MPC8548 and MPC8543 PowerPC processors running at speeds up to 1.5 GHz or up to 2 MB of SDRAM, 2 MB non-volatile SRAM, 128 K non-volatile FRAM, and up to 16 GB of Flash solid state disk storage. Each board provides versatile front- and rear-connection I/O options including a minimum of four USB ports, up to three 10/100/1000BASE-T Ethernet channels, and up to two SATA ports.

MEN MICRO GMBH
www.menmicro.com

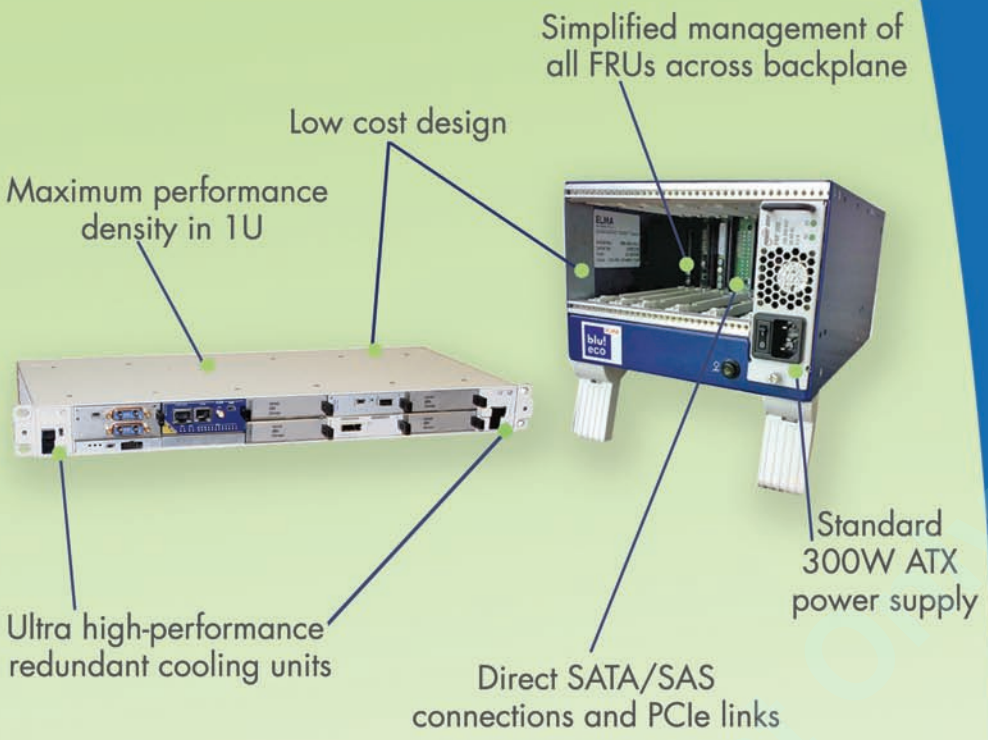

MicroSlim Chassis Line

Ruggedized MicroTCA

MicroTCA Load Boards and Extenders

Rethink MicroTCA.

Thought MicroTCA was out of reach for your design? Think again. Elma's ultra-compact MicroSlim™ and cost effective EcoBox™ innovations are making MicroTCA more competitive in a wide range of applications. From Components to Backplanes to System Platforms to Integration, Elma is Your Solution Partner for MicroTCA. Call us today!

ELMA

Your Solution Partner

USA Elma Electronic Inc.

Phone: 510.656.3400 Email: sales@elma.com
Web: www.elma.com

Elma Product Divisions

System Platforms & Integrated Products

Enclosures & Components

Backplanes

Cabinets

Rotary Switches

kontron

» We Do Not Build Navy Ships «

Our Customers Do.

At the heart of the most advanced command and control systems deployed in Navy ships are Kontron Military Rugged COTS boards and systems. Keeping us safe with smart applications, military contractors look to Kontron for superior technology, performance and life cycle management expertise.

MILITARY RUGGED COTS

Call, Email or Visit today.

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

If it's embedded, it's Kontron.

Protecting today's military electronics systems with real-time hardware/software protection countermeasures

By Paul Bradley

Today's military electronic systems have undergone a rapid evolution, and as a result, these systems and the data stored and transferred within are increasingly vulnerable to harmful manipulation caused by both unintentional and malicious actions. While systems designers are becoming more security-conscious and security requirements are increasingly being included in the design and development process, there are many challenges to address when attempting to secure electronics systems. The optimal approach is to incorporate both hardware-based and software-based security measures to protect against piracy, reverse engineering, and unauthorized use.

U.S. Navy photo by Mass Communication Specialist 2nd Class Barry Hirayama

Through many advances, including semiconductor integration and the use of third-party software and hardware IP, electronic systems have undergone a rapid evolution in both complexity and construction. Systems that were once composed of a few million transistors and a few hundred thousand lines of code are being replaced with systems constructed with hundreds of millions of transistors and millions of lines of code.

As a result, modern military electronic systems and the data stored and transferred within are increasingly vulnerable to harmful manipulation, caused by both unintentional and malicious actions. In terms of unintentional vulnerabilities, the majority are introduced during the system design phase, when security concerns are inadequately defined or insufficiently considered, and simple design flaws occur. Unfortunately, these unintentional vulnerabilities are increasingly being exploited by ever-more-sophisticated and malicious threats.

Recent news attests to these vulnerabilities. Over the past year, a number of high-profile cyber assaults were launched against defense, government, utility, and financial sector infrastructures, including the infiltration of the Pentagon's highly secret fighter jet project and a breach of the Air Force's air traffic control system – both of which represent shocking blows to the nation's defense capabilities. Evidence increasingly indicates that cyber attacks that were once perpetrated by a few poorly funded individuals are now being executed by a larger number of well-funded and well-organized institutions with scientific resources, aided by an incredibly efficient mass communication system – the Internet.

While systems designers are becoming more security-conscious and security requirements are increasingly being included in the design and development process, there are many challenges to address when attempting to secure electronics systems. The optimal approach is to incorporate both hardware-based and software-based security measures to protect against piracy, reverse engineering, and unauthorized use. These software and hardware systems must be bound tightly together and operate in concert to ensure maximum protection. Our discussion will demonstrate the feasibility of a hardware-assisted technology for adaptive software protection. It will address the elements of a software/hardware approach, and explore how a programmable design can provide real-time security monitoring to detect unexpected or illegal behavior.

The programmable hardware-assisted technology

A programmable hardware-assisted technology provides many benefits over an exclusively software-based alternative, including the ability to monitor low-level hardware functions in real time and the ability to monitor multiple points in a design simultaneously without impacting CPU performance. The most notable benefit, however, is that the hardware-assisted technology is programmable, allowing a variety of monitoring and countermeasure functions to be executed during runtime and providing a means for new security monitors and countermeasures to be created on-the-fly, or uploaded post-deployment to address any unexpected threats or latent design flaws.

This "defense logic" consists of distributed programmable instruments that can be configured repeatedly to dynamically implement different security checks that constantly monitor the system operation to detect unexpected or illegal behavior. Programmability also allows a large number of checks to be implemented by time-sharing the same hardware.

Technology that couples an existing software protection product, such as the ARM TrustZone (TZ) or Green Hills Software's

INTEGRITY RTOS, with security analysis software functions and a new programmable hardware mechanism will defend against a broad spectrum of attacks including Class I (outsiders or hackers), Class II (insiders), and Class III (well-funded organizations or nation-states). Reported attacks can be quickly analyzed and classified so that an appropriate-level response is issued based on the severity of the detected attack.

System architecture

The programmable defense logic is distributed through much of the hardware subsystem to provide ubiquitous coverage, with the specific location determined at design time. The location can be constructed differently for each system, providing a unique security scheme for each hardware design. Even systems composed of the same hardware design can have unique security schemes via the configurable programmable defense logic methodology, and this logic can be controlled through a secure JTAG port and/or by an embedded processor via a secure internal interface.

The security monitor is a programmable transaction engine configured to implement finite state machines to check user-specified behavior properties such as memory access privileges, bus performance levels, boot sequences, and operation signatures, all using the signals brought to its inputs to be analyzed. Signal probe networks are then configured to select a subset of the monitored signals and transport them to the security monitors. To implement different security checks, supervisor security software (or a specialized hardware controller) configures the signal probe networks to select the groups of signals to be checked by security monitors and configures them to perform the required checks. All security instrument configuration programs are encrypted and stored in one or more locations, including hardware-controlled secure flash, secure OTP memory, or software-controlled flash memory.

A dynamic wrapper is used in conjunction with the security monitor to provide a variety of real-time countermeasures. Any wrapped signal or set of signals can be controlled in real time in the event a security monitor (or the supervisor security software) detects misbehavior. Dynamic wrappers can be used to isolate logic, protect memory, reset peripherals, create decoy transactions, and erase keys or other sensitive data within memory.

Figure 1 provides an illustration of the main components of the programmable defense logic mechanism inserted in an SoC with a processor running a secure RTOS. (The SoC may be contained in an ASIC, FPGA, or a collection of off-the-shelf electronics mounted on a printed circuit board assembly.) Additionally, Figure 2 illustrates the hardware and software architecture where the programmable defense logic is configured and controlled by supervisor security software that runs in a protected software area and performs a variety of functions, such as managing the time-sharing security monitoring functions, activating countermeasures, and more.

System security checks

System security checks are application-dependent and circuit-dependent. The programmable defense logic element checks both the operation of the hardware and the integrity of the software. The checks target the entire SoC without distinguishing between the secure and non-secure environments. A first

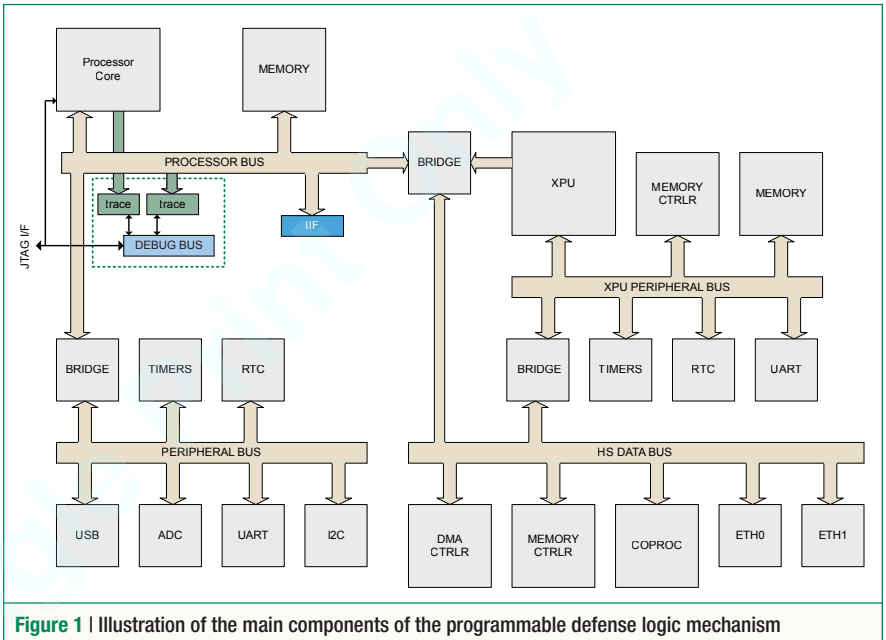

Figure 1 | Illustration of the main components of the programmable defense logic mechanism (shown in blue) inserted in an SoC with a processor running a secure RTOS.

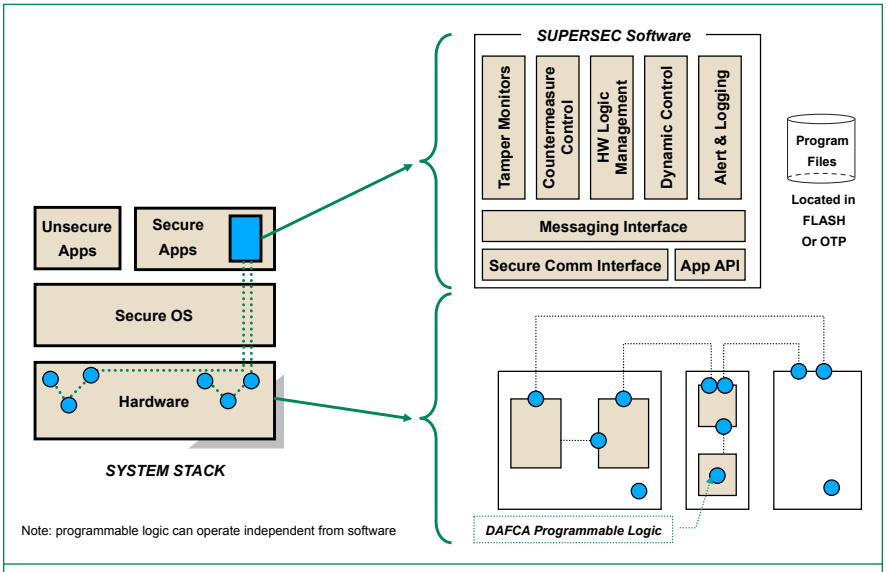

Figure 2 | Illustration of the software and hardware architecture.

category of checks is looking for a set of user-specified security violations such as:

- An attempt to access a restricted address space
- Denial of service
- A change in static characteristics (for example, checksum) of loaded software
- An unexpected output change on a disabled core (for example, an internal logic block)
- An illegal mode of operation on a core
- A change in a dynamic characteristic (for example, signature) of code execution
- Excessive resource utilization – beyond some defined limit
- An attempt to tamper with the boot or BIOS code
- Clock glitching and clock modification attacks
- Tampering attacks that change environmental conditions such as temperature or voltage – requires analog sensors with A-to-D converters coupled with a security monitor

The second category of checks examines general correctness properties (that is, assertions) of system behavior. The rationale for assertion checking is that tampering attacks often cause the system to operate in an incorrect way. Some checks will be based on assertions used in predeployment simulation to verify the correct implementation of the standard communication protocols used within the SoC (AMBA, PCI, and so on) or the behavior of a specific block.

All checks are prepared and verified predeployment in a secure environment, and their corresponding configurations are preloaded into one or more secure SoC memory locations. Not only are these configurations encrypted, but they are unique to each design – and difficult to understand without access to the functional design database. Moreover, in a powered-off state, the programmable defense logic is "blank" (unprogrammed); thus, its function is concealed from unauthorized persons and attackers trying to reverse engineer the system. The programmable defense logic is invisible to both the mission logic and the application software running in the non-secure environment. As such, these security functions can be hidden from "untrusted" system and chip manufacturers and/or others with access to the system in the supply chain.

Implementing countermeasures

When the defense logic detects an attack, it reports it to the supervisor security software via a high-priority or privileged "security" interrupt, along with information detailing the nature and location of the attack. The supervisor security software analyzes the received information, determines the severity of the attack, and deploys the appropriate countermeasures.

The defense logic implements countermeasures by controlling specified signals. For example, if a core exhibits illegal behavior, various countermeasures may isolate that core by disabling its clock, powering it off, holding it in a reset state, or forcing safe values on its outputs. In the software domain, a corrupted routine may be reloaded from the memory or disk. Urgent countermeasures that need to be deployed in real time can be implemented directly by security monitors without requiring supervisor security software.

The programmable defense logic can also be deployed within fail-safe and recovery security applications. System-level countermeasures and recovery may combine techniques such as provision of fail-safe states, spare logic to replace misbehaving logic, and check pointing to return the system to a known-good state.

Taking action

Securing our military electronic systems is critical to the nation's future. The first step in achieving this is to begin incorporating hardware- and software-based security measures early in the design stage to ensure the most comprehensive protection against the serious threats posed by piracy, reverse engineering, and unauthorized use.

Paul Bradley *is chief technical officer of DAFCA, Inc. He has more than 20 years' experience in electronics and systems design, and specializes in product development and engineering leadership in emerging technology markets. He has held numerous engineering and technical leadership positions at Motorola, Nortel, CrossComm, Sonoma Systems, and Internet Photonics prior to joining DAFCA. He can be reached at paul.bradley@dafca.com.*

DAFCA, Inc.
774-204-0220
www.dafca.com

U.S. Navy photo by Mass Communication Specialist 2nd Class Nardelito Gervacio

Taking the VPX ecosystem to the next level

By Justin Moll

The VPX industry needs to add more test/prototype and development tools to its ecosystem, and careful consideration of modules, cooling methods, and subsystem integration is also necessary in furthering the high-speed architecture's ecosystem.

There are several elements to an architecture's ecosystem, including development products, deployable products, test tools, and related services. Imagine how much time can be shaved off project development times with specific devices for a specified form factor that aid prototyping/testing and development. Without devices such as load boards, RTMs, and development systems, the systems engineer is forced to use less accurate alternatives. When development is quicker and easier, a larger amount of products can be brought to market more efficiently. This in turn creates a cyclical churning that encourages higher adoption rates (with cheaper and more accurate products), more development, and so on.

During the past few years, the VPX architecture has seen an accumulation of boards, chassis, and backplanes to provide more choices for the customer. However, an architecture needs more to thrive. As mentioned, without adequate development systems and testing/prototyping tools, development can be hindered. Thus, the VPX ecosystem also needs these important elements. Cooling methods and subsystem integration are also vital considerations linked to the VPX ecosystem.

Development systems

Development chassis are a key part of the ecosystem, as they provide a rugged enclosure for prototyping and testing. They will often have top carrying handles for portability and rugged or scratch-proof covers to survive demos, shows and events, or lab use. However, as VPX card wattages continue to increase, the development chassis needs to evolve with it. This includes having higher-power options for larger VPX backplanes. Some of these development chassis will need power supplies exceeding 1,300 W, and the enclosure will need to adequately cool the higher-wattage cards.

More fans with higher Cubic Feet per Minute (CFM) may be required in these development chassis. Take, for example, a five-slot VPX backplane in a small portable chassis with protective covers, and 2 x 190 CFM fans. (If this were a VME64x architecture, it would have just a 1 x 90 CFM fan.) Then consider an open-frame chassis designed specifically for VPX. To handle the more intensive power and cooling requirements, this unit has 3 x 170 CFM fans and a 1,300 W power supply often used for larger VPX backplanes. With a small backplane, the five-slot chassis was modified with the more robust fans to cool a VPX design. For larger backplanes, a design (like the open frame version) allowing more fans with high CFMs is likely a necessity for adequate cooling.

To make a development chassis simpler and more cost-effective, a two-slot 6U VPX backplane implementation is needed. The backplane would allow point-to-point signals between the slots and its flexible design would ensure compatible use with a large amount of VPX boards. A small backplane like this with fewer slots and layers could also help provide a useful way of prototyping/testing boards without a large expense. The small size would make it easier to provide power and cooling, thus reducing costs as well. A two-slot VPX backplane is in development.

Going hand-in-hand with the development chassis is the extender board. VPX extender boards are coming out in both 3U and 6U versions. These allow the signals to be extended outside of the card cage for easier test and debugging. This is particularly important for the immediate adjacent cards in the system. As the mating right-angle connectors are not available in the market for plugging the VPX cards into the extender, a rigid-flex-rigid design can overcome the problem.

> During the past few years, the VPX architecture has seen an accumulation of boards, chassis, and backplanes to provide more choices for the customer. However, an architecture needs more to thrive.

Test tools: Proof of performance

Test and debug is an important part of the design process and the more testing/prototyping tools the industry can provide for VPX, the better. With the performance of VPX, the design engineer will witness lots of power, and plenty of heat. Thus, VPX systems must be carefully simulated and tested. The VPX ecosystem has been short of these types of signal performance testing tools.

SERDES test modules: "The health monitor"

However, one tool that will help facilitate the goal of building the VPX ecosystem is the VPX SERDES test module: a pluggable 6U VPX card that acts like a health monitor, checking the "vital stats" of the VPX system. In minutes, a designer or tester can check the Bit Error Rate (BER), jitter, skew, and so on – and execute pattern generation such as eye diagrams for the signals. It can be used to test a VPX board, a VPX backplane, or the full interconnect path between boards across the backplane.

Load boards: Power 'em up

Another beneficial tool for VPX is a load board. It can aid the system designer in assuring adequate chassis cooling and verifying that the VPX chassis is capable of meeting the power requirements of the system (or VITA specs). The load board functions to test a system's cooling capabilities by first applying the load to the power supply for verification and creating the necessary heat to confirm chassis cooling. By helping the test engineer locate hot spots in the chassis, he can verify where to optimally redirect the airflow to prevent overheating (see Figure 1). The load boards can be designed in either 3U or 6U versions. Without a load board, the test engineer simply could not test the system at various wattage load levels and make necessary adjustments to the system.

Important VPX considerations: RTMs and cooling

Development and testing products are very important for the VPX ecosystem. But the standard chassis and backplanes need some creative solutions too, like RTMs.

RTMs ease I/O

Also a viable consideration affecting the VPX ecosystem is the Rear Transition Module (RTM), which brings I/O off a backplane in a pluggable format. Otherwise, designers need to use ribbon cable or wire wraps that are less reliable, not designed for rugged use, and so on. However, an RTM for VPX can be designed for various configurations of the J0-J6 connectors, and a version designed as a Universal VPX RTM break-out board can allow a test engineer to access I/O signals on custom-built VPX boards. The board would not be intended for high-speed

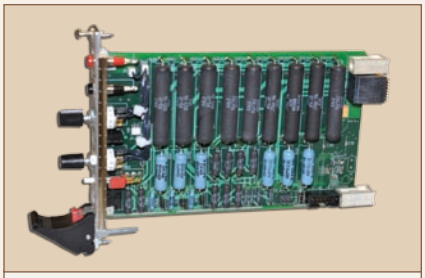

Figure 1 | A VPX load board allows the test engineer to step up the voltages and power levels as desired to apply various loads for testing.

When You're At 30,000 Feet Quality Matters

Our products and service are second to none

Photo Credit: U.S. Air Force photo by Staff Sgt. Samuel Morse (Released)

At VersaLogic we know about building dependable, high reliability embedded computers for mission critical applications. Our products have been designed into some of the most sophisticated and high profile military and aerospace equipment used today. Every board we make is run through exhaustive inspections and quality tests, ensuring that we deliver only the best. And with our world class service and five year availability guarantee, we'll always be there when you need us. Whether you need one of our standard products or a specially customized version, our skilled technical staff will work with you to meet your exact specifications. So before your program launches, choose the company with the quality and service to take you where you need to go.

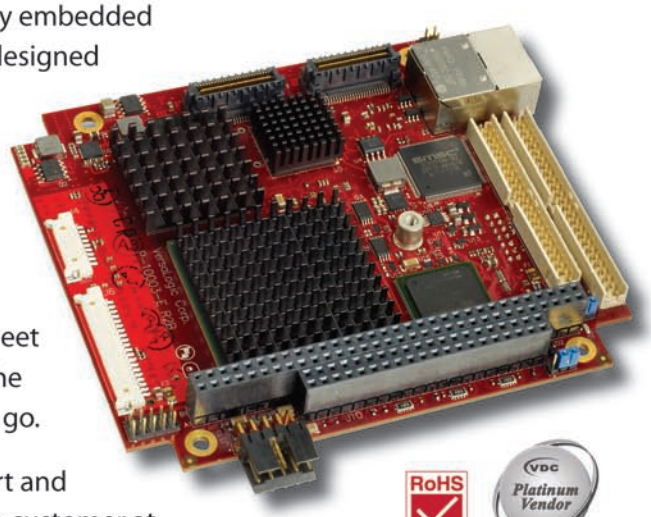

With more than 30 years experience delivering extraordinary support and on-time delivery, VersaLogic has perfected the fine art of service, one customer at a time. Experience it for yourself; call 800-824-3163 for more information!

Recipient of the Platinum Vendor Award for 5 years running!

1-800-824-3163
1-541-485-8575
www.VersaLogic.com/COTS

VERSALOGIC CORPORATION

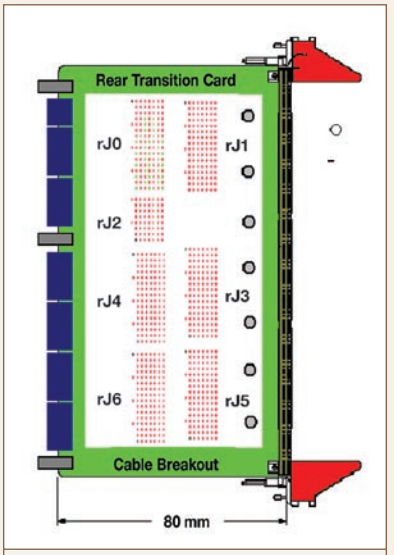

Figure 2 | A Universal VPX RTM board can have open connections out the back for directly wiring any pin.

multigig differential signals but would typically be used to bring out single-ended TTL signals that might be part of a customer's custom I/O board. (Figure 2 illustrates this type of interface.)

Development with conduction/ liquid cooling

In addition to the aforementioned metrics and considerations for the VPX ecosystem, cooling is imperative to VPX's market proliferation. Whether using 3U or 6U high cards, typical VPX wattage per slot today is 100-120 W and rising. Therefore, enclosures need to cool these chassis in creative ways, especially in densely packaged designs. With applications such as UAVs, forced-air cooling may not be an option due to lack of available air in the enclosed design or in high altitudes. Conduction cooling can be effective, but conduction alone typically cannot dissipate the 100-120+ W levels. VITA 48 will provide a pathway to liquid cooling through the individual modules.

Although progress is being made in the VITA 48 specification, there is more work to be done. The VPX community needs a cost-effective way to develop higher-wattage conduction-cooled boards. One solution in the meantime is to have the liquid go through the chassis sidewalls. This can provide up to 150 W per slot of cooling, without the tricky issue of having the liquid go through each card. This alternative is more than adequate for many applications – and is a highly simplified and cost-effective design consideration.

Summing it all up: Fostering the VPX ecosystem

From the already-present VPX boards, backplanes, and chassis to the still-needed development and test tools to the RTM, cooling, and subsystem integration considerations, the VPX ecosystem needs to continue to grow. This in turn will make a design and test engineer's life easier and provide superior VPX solutions in the industry.

Justin Moll *is director of marketing at Elma Bustronic, where he has worked since 2000. He is active in VITA and PICMG and has been the VP of marketing for the StarFabric Trade Association. He received his Bachelor of Science degree in Business Administration from the University of California, Riverside. He can be reached at Justin.Moll@elmabustronic.com.*

Elma Bustronic
510-490-7388
www.bustronic.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.*

Rugged 3U VPX SBC designed for SWaP

GE's on quite a roll these days with their A&D program design wins. One common theme? 3U VPX, which shoehorns nicely into legacy spaces while providing ultra-rugged, conduction-cooled, high-density computing. Take their rugged VPXcel3 SBC341, for instance. Based upon a 2.26 GHz Intel Core 2 Duo (or 17 W 1.86 GHz version), this board runs 13 percent faster than its predecessor, but at 20 percent lower Thermal Design Power (TDP). The board's ideal for SWaP designs and can be equipped with up to 8 GB of DDR3 memory.

With VPX, there's ample ability to get I/O to and from the board. This one has GbE, x16 PCIe, x4 PCIe, and x1 PCIe. There are also four USB ports, COM1 and COM2 ports, and two SATA interfaces. As for software, this one's designed to provide desktop-like performance in an embedded system. There's Windows XP, XP Embedded, Windows Vista, and Linux. A BSP for VxWorks 6.6 is also available.

GE Fanuc Intelligent Platforms • www.gefanucembedded.com • RSC# 42970

Reballed ICs solve RoHS woes

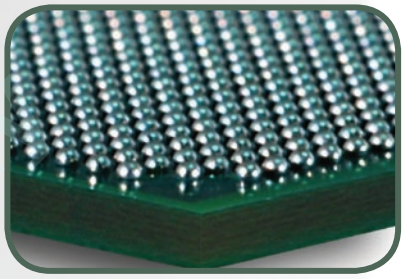

Back in the day, there were cans, TO packages, DIPs, and sidebrazed packages. When IC vendors increased pin count and migrated to Leadless Chip Carriers (LCCs), the defense industry had a collective heart attack because thermal mismatch and vibration profiles proved the unreliability of LCCs. The solution then was to solder on individual "J" leads to each package position. Today the semiconductor industry has evolved to even higher-density Ball Grid Array (BGA) packages, which have better shear strength to solve the vibration problems but aren't without downsides such as visual inspection or thermal mismatch.

But of *primary concern* these days is lead versus lead-free solder. Interconnect Systems, Inc. (ISI) has found a way to remove lead-free (RoHS) BGA solder balls and reattach good old lead-based solder balls with Sn63Pb37 eutectic solder spheres. These "old school" bumps ease the concern of tin whiskers and long-term solder reliability by reverting to tried-and-true materials used in DoD systems for decades. Utilizing an IPC/JEDEC J-STD MSL Classification process, ISI deballs and adds lead spheres – never subjecting the package or die to a temperature cycle above 50 percent of the supplier's recommended cycle. This avoids inadvertent excursions beyond glass transition temperatures, which can overstress a package and induce hermeticity problems. Spheres are reattached in a nitrogen environment and ISI performs 100 percent optical inspection and then certifies the parts to original OEM specs for flatness and coplanarity. And, while it ain't glamorous or the latest whiz-bang microprocessor, ISI's value-add is absolutely crucial to *using* the latest COTS IC devices.

Interconnect Systems, Inc. • www.isipkg.com • RSC# 42832

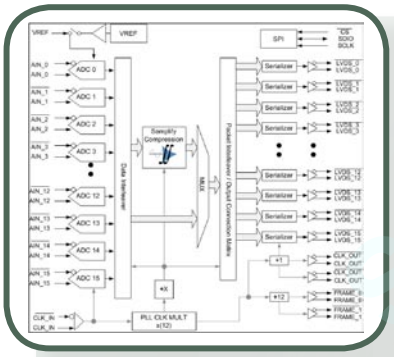

Programmable compression frees up A/D bandwidth

This is one of our favorite new products because of cleverness and the potential wide-ranging applicability to defense and civilian markets. Samplify is a fabless mixed-signal semiconductor company whose Prism 3.0 real-time signal compression algorithm boasts a 200 percent increase in decompression performance while reducing data latency to a mere 1 microsecond. Samplify's mission is to create high-performance A/D converters using existing *digital* semiconductor foundry processes, thus avoiding the need for proprietary analog fabs. Their "secret recipe" combines an analog front-end with a high-speed compressor that mimics the best-in-breed high-speed analog converter. Digital data is then shipped over system pipelines to a downstream decompressor. This architecture claims the best of both worlds: low-cost but high-speed A/D devices in high-volume digital fabs.

The company's Prism 3.0 compression algorithm improves on this process and can be instantiated in FPGAs to solve transmission issues in high-data-bandwidth applications including radar, sonar, wireless base stations, and CT medical systems. The new decompressor uses 50 percent fewer FPGA resources or can be *software decoded* on an IA processor with a claimed 200 percent speed improvement. Flow-through latency (compression-to-decompression) is under 1 microsecond and independent of packet size. User-selectable compression modes allow algorithm tuning, and the SignalTrak mode maintains signal quality for bursting signals with wide dynamic ranges. We're going to keep our eyes on this technology because we expect to see it show up in other devices.

Samplify Systems • www.samplify.com • RSC# 42833

It's an SD card ... isn't it?

So what … it's an SD card, right? How different could yet another SD card be? That's what we thought – until we read the specs. Swissbit, a European company specializing in packaging for more than 18 years, has designed a rugged SD card that's so different from the competition that it should be in its own unique category. Sure, it conforms to the SD card spec for removable flash storage. But the S-200 is designed specifically for the needs of industrial and extended temperature applications. The attention to design detail is absolutely startling.

Available up to 8 GB using robust, single-level cell flash chips (unlike the cheaper MLC commercial SD cards), there's an onboard 32-bit RISC controller to provide ECC, bad block management, and wear-leveling algorithms and to control power-loss protection and power-saving modes. The CPU also runs Swissbit software to determine wear level classes and bad block counts, data that can factor into maintenance plans. The SD card's ABS/PC housing is designed for 15 KV ESD protection and contains provisions to resist bending torque and achieve 10,000 insertions via gold-plated connectors. Internal write protection and lock/unlock features adhere to SD Specification 1.01, 1.10, and 2.0. All of these specs help reinforce our original assertion that "No, this isn't just another SD card."

Swissbit AG • www.swissbitna.com • RSC# 42834

Talk without towers

For Army tactical radios, there's SINCGARS, Tri-TAC, WIN-T, and eventually JTRS. Then there are Family Radio Service, satellite phones, and cell phones. They're all used – depending upon the scenario – but nearly all of them require some sort of infrastructure, behemoth antennae, or are just plain impractical in many situations. But the WZRDnet *ad hoc* mesh network phone from Telegrid aims to simplify tactical voice and text communications. Using a ruggedized mobile phone that resembles an FRS handset, the unit can connect hundreds of users in a mesh network that "hops" data packets between available radios from source to destination. The mesh doesn't require anything more than the radios themselves.

Utilizing the 2.4 GHz spectrum (the same as WiFi), IEEE 802.15.4 ZigBee, plus a proprietary codec and compression algorithm, the radios can move up to 250 Kbps over this self-forming/self-healing network. Battlefield relay stations can extend the range beyond 32 hops or the typical 400 meter handset-to-handset distance, while gateway devices bridge the mesh into cellular, PSTN (POTS), and other networks. Handsets rely on USB rechargeable (or discardable) lithium ion batteries that last up to 55 hours (5 percent talk, 5 percent receive, 90 percent standby), or up to 17 hours of pure talk time. The units also support intra-handset text messaging and have built-in GPS receivers that can provide GPS coordinates in peer-to-peer or broadcast arrangements.

Telegrid Technologies
www.telegrid.com
RSC# 42969

Eight-channel USB data acquisition

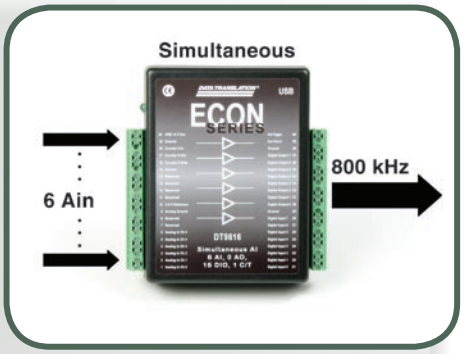

Military systems use lots of I/O, especially during laboratory and field trial testing. And while MIL-SPEC test equipment is appropriate for depot support, low-cost COTS testers are perfectly legit during a program's early phases. That's where Data Translation's low-cost DT9816-S USB-powered eight-channel data acquisition module comes in. Sampling each channel up to 800 KHz, the shielded and ruggedized enclosure is ideal for easy setup and medium-rate sampling. Each channel boasts 16-bit resolution, and the whole system can support 4.8 MHz of total throughput across six total channels.

Signal sampling ranges from ± 10 V and ± 5 V, and the unit offers monitor and control via eight digital inputs and eight digital outputs. A single 16-bit counter/timer is included for reference and real-time event logging. Powered from a standard USB connector, the DT9816-S also ships with a reasonable complement of test and measurement software, ready-to-measure application software, and the company's Measure Foundry – a "powerful drag and drop test and measurement application." And as much as we hate to say this: The unit sells for a mere $595. That's COTS for you.

Data Translation
www.datatranslation.com
RSC# 42831

Fast erase, secure SSD in 3U VPX

Seems like Solid State Disks (SSDs) are available everywhere. But there's a big difference between high-density disks with commercial-level reliability and SSDs designed for A&D applications. The Curtiss-Wright Controls Embedded Computing (CWCEC) VPX3-FSM flash storage module is designed for rugged, high-security platforms. Available in conduction-cooled 3U VPX (VITA 46) and VPX-REDI (VITA 48.2) versions (at 1.0" pitch), the unit boasts up to 256 GB of reliable SLC NAND flash. (SLC is the more expensive kind with lower density, but it offers greatly improved longevity.) CWCEC says their FSM has an MTBF of 2,000,000 hours, with each memory cell good for 100,000 writes.

Organized as 64 GB banks, the SSD can operate as four SATA drives or a single RAIDØ SATA. In RAIDØ, 160 MBps R/W is possible; 75 MBps per port is possible in a JBOD configuration. Besides the expected wear leveling and bad block management, there's also an MCU-based temperature monitor, BIT, and RS-232 or I2C controller interface. But what really sets this SSD apart from the wannabees is the five-level, 256-bit AES encryption. Internal or external key storage is possible, with onboard, battery-backed SRAM or EEPROM. A fast erase wipes the drive for declassification purposes, initiated remotely or via a front-panel push button. Wait 500 ns for the AES key to disappear and the drive is rendered unreadable.

Curtiss-Wright Controls Embedded Computing
www.cwcontrols.com
RSC# 42615

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.

OpenVPX goes from conflict to collaboration

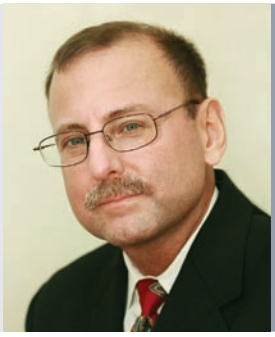

By Chris A. Ciufo, Editor

The back story

Just about the most exciting thing that's happened in the VME ecosystem this year was when a gang of rebels egged on by Mercury Computer Systems started the OpenVPX Working Group back in January. Why? They were appalled by the lack of interoperability built into the 3-years-young VPX/VITA 46 specifications[1]. "Shockingly," Mercury and friends wanted to follow the PICMG model of control-data-management planes to assure high-availability military systems that would accommodate LRUs from multiple vendors with minimal backplane changes.

How VITA 46 evolved – with Curtiss-Wright Controls Embedded Computing (CWCEC), Mercury, GE Fanuc, Tek Microsystems, and many others driving the specs – allowed for too much vendor-specific pinout customization. Mercury (and their customers such as Boeing) essentially argued that once a VITA 46 design was won, it was damn near impossible to design out the winning vendor from a quasi-sole-source position[2].

But besides breaking ranks with the VITA 46 team in the VSO, what made OpenVPX such an evolving soap opera was that VITA's executive director Ray Alderman openly endorsed the OpenVPX Working Group's efforts[3] to work outside of VITA to realize Mercury's vision of interoperable VPX boards – now just called "OpenVPX." Alderman also seemed to tacitly endorse PICMG's architecture as the basis for OpenVPX. And if all this weren't enough, the OpenVPX team refused to let CWCEC play in their sandbox. Since CWCEC appeared to have more VITA 46/VPX design wins than anyone, their customer base fired back at the OpenVPX Working Group with hate mail aplenty[4].

Under extreme pressure (and after signing a few secret NDAs and agreements), CWCEC was admitted to the OpenVPX Working Group and even agreed to chair the VITA 65 committee that would eventually reintegrate OpenVPX spec results into VITA. By Fall 2009, the plan was to present the specs back to the VITA 65 committee and dissolve the OpenVPX Working Group[5].

The present: MILCOM announcements

I'm happy to report that everything went as planned, and on October 19, 2009 at the AFCEA/IEEE MILCOM show in Boston, MA, the OpenVPX Working Group presented to VITA the results of nine months of public collaboration resulting in a set of interoperable OpenVPX specifications. The VITA 65 committee will now formalize for ANSI approval a set of documents that builds upon the previous VITA 46 (VPX) and VITA 48 (REDI) specifications, but adds air-cooled LRU plug-and-play vendor interoperability via Rear Transition Modules (RTMs).

CWCEC and Hybricon's OpenVPX mission computer

CWCEC and Hybricon demonstrated at MILCOM that real working VITA 65/OpenVPX systems are available and ready to solve military problems. Comprised of a four-slot vetronics-style cold-plate chassis, the Hybricon SFF-4 Small Form Factor box swallows 3U OpenVPX-style cards plus a plug-in PSU. Of interest is the lack of any cabling in the box; Hybricon relies on a plug-in front panel PWB for mounting the 38999 connectors.

Based upon the VITA 48 REDI 1-inch pitch backplane, the chassis is designed for MIL-STD-704F aircraft and MIL-STD-1275B vehicles. In the MILCOM demo configuration, it's populated with two 3U OpenVPX PowerPC SBCs (VPX3-127), an FPGA module (VPX3-450), a carrier card for PMCs or XMCs, and two mezzanine modules: one with an eight-port GbE switch (PMC-650) and the other with a dual output graphics controller (XMC-710). All modules are rugged and conduction cooled. Collectively, the cards talk on several OpenVPX backplane profiles that incorporate data, control, and management planes based upon Mercury's original vision of OpenVPX and implemented as payload, peripheral, and switch architectures.

Hybricon's chassis, originally scheduled for production in mid-2010, was accelerated by more than six months to meet market demand. And CWCEC's modules, though redesigned from the previous VITA 46 "VPX" to the new VITA 65 "OpenVPX" flavors, are representative of all the company's modules transitioning to OpenVPX. What's most interesting is how quickly CWCEC went from being the underdog at OpenVPX to being the first company to announce working modules. (CWCEC says 6U-sized cards will be available soon.) They beat Mercury by a scant 24 hours.

Mercury at MILCOM: "Broad range of products"

You'd expect Mercury Computer Systems to announce a bunch of OpenVPX cards at MILCOM, since it was their idea in the first place. Compliant with the "new V1.0 OpenVPX and draft VITA 65 Specifications," Mercury's Ensemble 3000 Series 3UOpenVPX and Ensemble 600 Series 6UOpenVPX product lines comprise SBCs, switch and I/O modules, and high-compute density boards. While no specific cards were announced at MILCOM, a trip to the company's website lists six 3U OpenVPX choices, from complete systems to Virtex-5 FPGA nodes and OpenVPX Ethernet switch hubs, to RapidIO data plane switches. In the 6U category, Mercury has nearly 10 products listed, from 8640D dual-core PowerPC SBCs to serial FPDP I/O modules.

What does it all mean?

So forget about VPX as Gen 1. Gen 2 is OpenVPX. The controversy's over. The conflict is resolved. Collaboration among the 28 OpenVPX Working Group members has shown that VME's future is in OpenVPX under VITA 65.

For the full-length version of this article, go to www.vmecritical.com/articles/id/?4264.

[1] *OpenVPX Industry Working Group: Open for business, or just controversy?,* by Chris A. Ciufo, www.mil-embedded.com/articles/id/?3818

[2] *Mercury's "OpenVPX Industry Working Group" colors outside VSO's lines: An efficient technology fast-track or Pandora's box? Q&A with Mercury Computer Systems' Greg Tiedemann,* www.vmecritical.com/articles/id/?3884

[3] *Those VPX surprises keep a'coming*, by Sharon Schnakenburg, www.vmecritical.com/articles/id/?3771; and *OpenVPX: It's all about the backplane,* by Ray Alderman, www.vmecritical.com/articles/id/?3880

[4] *VITA 65 versus OpenVPX: Why can't we all just get along?,* by Chris A. Ciufo, www.vmecritical.com/articles/id/?3885

[5] *The best way to run a railroad,* by Chris A. Ciufo, www.mil-embedded.com//articles/id/?3925

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

Which of these platforms use GE Fanuc hardware?

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

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

GE Fanuc
Intelligent Platforms

RUGGED DEPLOYED ELECTRONIC SYSTEMS

YOUR TOTAL SOLUTIONS PARTNER

Photo courtesy of General Dynamics Land Systems Inc.

Cooling...Size...Weight...Power...

These are the system design hurdles you face. Curtiss-Wright Controls Electronic Systems helps by designing next generation high performance interoperable systems with more computer power and bandwidth. With fully qualified, turnkey rugged subsystems proven to work together, you can replace cumbersome, power-hungry vehicle systems with compact, power-managed packaged solutions. We are your integration partner. We'll handle the design details so that you can focus on taking your system capabilities to the next level. We have system solutions for servo motion and fire control, turret and platform stabilization sensors, and design subsystems that use COTS and custom electronics. Our success stories include HMMWV, Armored Security Vehicle, Stryker, Abrams and Bradley Fighting Vehicle.

High Voltage Servo Motion Control
610VDC Servo Motor Controller

Fire Control Solutions
Fire Control & Communications Processor

Ammunition Autoloader Controls
Large Caliber (105mm & 120mm) Ammunition Handling Systems

Sight Head, Gun & Turret Stabilization
Fiber-Optic Gyro Rate Sensor Assembly

CURTISS WRIGHT Controls
Electronic Systems

Tel: +1 (661) 257-4430
cwcelectronicsystems.com
systeminfo@curtisswright.com

ABOVE & BEYOND