Daily Briefing: *News Snippets*

Military EMBEDDED SYSTEMS

VOLUME 6 NUMBER 5
JULY/AUG 2010

INCLUDING:

Chris A. Ciufo
Thinking outside the box

Field Intelligence
1553's well-deserved longevity

Mil Tech Insider
I/O mezzanines in FPGA computing

Legacy Software Migration
ALT Software:
Managing legacy graphics software

MIL-EMBEDDED.COM

FPGAS ARE THE MILITARY'S NEW SECRET WEAPON

Also:
UAVs step up to the front line

www.opensystemsmedia.com

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

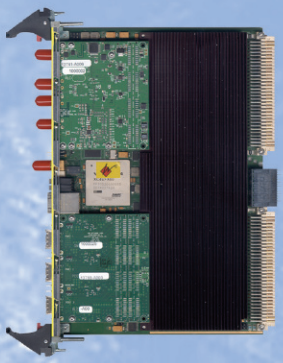

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

Direct Seamless Connections with no Data Reduction
Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE
Between FPGAs and Standard Output Modules
Between FPGAs and Storage Arrays

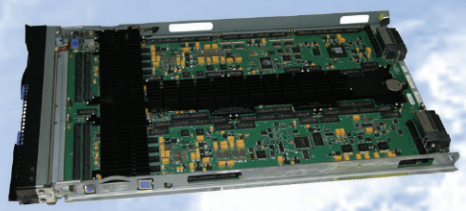

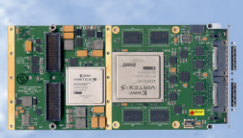

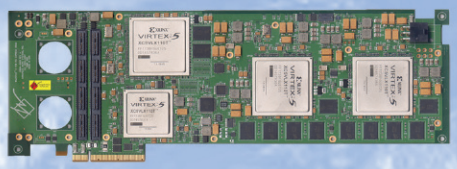

Annapolis Micro Systems

Ultimate Modularity
From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
Input/Output Modules Include:
Quad 130 MSPS thru Quad 550 MSPS A/D
1.5 GSps thru 5.0 GSps A/D, Quad 600 MSps D/A,
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10G, 40G or 100G Ethernet or SFPDP

VME/VXS/VPX, IBM Blade, PCI-X/PCI Express, PMC/XMC, MicroTCA

No Other FPGA Board Vendor Streams This Volume of Data Real Time Straight Into the Heart of the Processing Elements and Then Straight Back Out Again

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

INNODISK®
Beyond your imagination

Experience
Fast and Furious!

PCIe Solid State Disk

Matador Series

Custom FW for special needs: IOPS,wear Leveling, Security
Fast access on small file size (DRAM cache buffer)
Wide range of capacity support
Small form factor to fit into small chassis
Onboard RAID/SSD controller for better compatibility
Simple integration RAID system
Power Guard – To ensure data integrity

SSD Industrial Solid State Disk -Capacity to 512GB

2.5" SATA 20000
Read:230MB/s(max.)
Write:160MB/s(max.)

1.8" SATA 10000
Read:160MB/s(max.)
Write:140MB/s(max.)

2.5" ATA 8000-J
Read: 85MB/s(max.)
Write: 75MB/s(max.)

i-DIMM Series

Embedded Solution:
Low Profile, Wide Temp.
Server/Workstation Solution:
w/ECC, w/ECC w/PLL, w/ECC w/R
Printer Solution

Mini PCIe SSD

Mini PCIe DOM

MiniPCIe Series

8000 Series-EDC/iCF

Vertical type 40pin

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

iCF Write Protect Series
Write: 75MB/s(max.)

CFast 100

Capacity	2GB~16GB
Data Transfer Rate	Read-100MB/sec. (max.)
	Write-50MB/sec. (max.)

Low Profile / Vertical Type
Read: 24MB/s(max.)
Write: 14MB/s(max.)

www.innodisk.com

InnoDisk USA Corporation
43154 Christy Street Fremont, CA 94538 Tel: 510-770-9421 Email: usasales@innodisk.com

Military EMBEDDED SYSTEMS

July/August 2010 Volume 6 Number 5

COLUMNS

DEPARTMENTS

ON THE COVER:
In July 1945, the Army's successful Trinity test was the culmination of America's Manhattan Project, and proved that a nuclear weapon was not only feasible, but frighteningly lethal. During the war, the two brute force nuclear bombs required specially modified B-29s to deliver ordnance on target. Today, precision strike weapons carried by smaller UAS platforms can eliminate hard and soft targets with minimal collateral damage. The secret to these weapons is low-cost, lightweight COTS hardware and sophisticated software. See stories on pages 22, 26, 32, and 36 (Image courtesy of http://listsoplenty.com)

Published by: OpenSystems media

ISSN: Print 1557-3222

E-CASTS

http://ecast.opensystemsmedia.com

Seven strategies to optimize quality with frictionless delivery
August 31st • 11am MST
Presented by: IBM, Wind River

DO-178B Level A Programming Environment
September 9th • 10am MST
Presented by: Esterel, Wind River

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

Mission Critical

Control the sea

Cutting-edge reliability

Royal Navy Astute Class nuclear-powered attack submarine.

Thales' periscope provides a 360° scan of the surface above with minimal risk of detection.

Wind River embedded solutions deliver the breakthrough dependability and performance essential to innovation.

To control the sea, a submarine depends on remaining invisible. But when designing and building a sub, visibility is critical. That's why Thales partnered with Wind River to create a breakthrough in periscope design for the Royal Navy's new Astute-Class submarine.

Relying upon the proven innovation, reliability and performance of our VxWorks RTOS platform, Thales developed a state-of-the-art optronic imaging system that provides stable, high-resolution views in the world's most demanding conditions.

It's the kind of teamwork and support that's made Wind River a trusted leading provider of advanced embedded solutions for aerospace and defense.

To see how Wind River can help you innovate with confidence, download our Mission Critical Toolkit at www.windriver.com/missioncritical/security.

WIND RIVER

ADVERTISER INFORMATION

Page	Advertiser/Ad title
2	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
35	**BittWare, Inc.** – Essential building blocks
17	**CMP Techinsights** – A new era, a new ESC
44	**Curtiss-Wright Controls Embedded Computing** – Reduce your development time and costs
37	**Elma Electronic** – The industry's choice in VPX handles and panels
21	**Engineering Design Team Inc.** – PCIe x8 LX + PCIe 1553b
13	**Excalibur Systems, Inc.** – Dragon
43	**GE Intelligent Platforms, Inc.** – It's an open specification
3	**InnoDisk** – Experience fast and furious
30	**Interface Concept** – Trust a worldwide expert
19	**Microhard Systems Inc.** – Wireless innovation
24	**Nallatech** – High performance FPGA solutions
28	**North Atlantic Industries** – Multifunction I/O boards
31	**Parvus Corporation** – Qualified to perform
12	**Phoenix International** – RPC12 ruggedized 3U Fibre Channel RAID system
39	**Sanmina-SCI** – Embedded SSD solutions
27	**Schroff a brand of Pentair Electronic Packaging** – Schroff understands the demands
25	**Sealevel Systems, Inc.** – The right connection creates incredible power
21	**TEWS Technologies LLC** – COTS I/O solutions
31	**Trident Space & Defense** – Proteus Series SSDs
29	**Tri-M Systems Inc.** – PC/104 Can-Tainer
34	**Tri-M Systems Inc.** – 100Mhz PC/104 module
7	**Vector Electronics & Technology, Inc.** – VME/cPCI ruggedized chassis
10	**White Electronic Designs/Microsemi** – We create space
5	**Wind River Aerospace & Defense Division** – Mission critical
23	**Z Microsystems, Inc.** – Rugged rack mounted servers

E-LETTER

www.mil-embedded.com/eletter

- Securing wireless Local Area Network interconnections with Layer 2 encryption
  By Juan Asenjo, Thales e-Security
- Raising the bar for security needs: What does "secure boot" *really* mean?
  By J. Ryan Kenny, CPU Tech
- Requirements tracing tools can manage legacy systems, too
  By Pete Decher, Mentor Graphics

OpenSystems media

Military EMBEDDED SYSTEMS

DSP-FPGA.com | VME and Critical Systems | PC/104 and small form factors THE JOURNAL of MODULAR EMBEDDED DESIGN | INDUSTRIAL EMBEDDED SYSTEMS | CompactPCI Advanced & Micro TCA SYSTEMS | Embedded COMPUTING DESIGN

ITAR REGISTERED
NO. M26764

VECTOR

ELECTRONICS & TECHNOLOGY, INC.
A FINE TECHNOLOGY GROUP

VME / cPCI® Ruggedized Chassis

- Well suited for Military / Aeronautics
- VME64x, VME64 or cPCI® compatible
- Mil 461-D, E for shock/vibration, EMI/RFI
- CE and FCC class B, Part 15
- Many power options

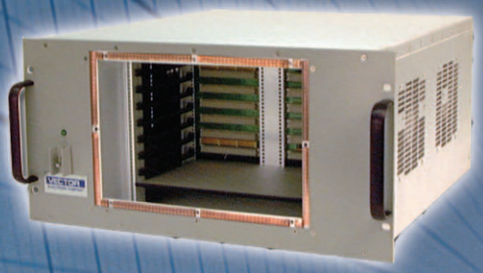

790 Chassis

792 Chassis

Series 445, Modified
Full EMI/RFI

VECTOR SPECIALIZES IN:

UNPARALLED FACTORY SUPPORT
CUSTOM CONFIGURATIONS
SHORTER LEAD TIMES

Made in U.S.A

(800)423-5659

www.vectorelect.com

MIL-STD-1553's longevity is well deserved

By Duncan Young

After almost 40 years of deployment, the MIL-STD-1553 databus protocol still holds a strong position on military platforms. How can a technology that old still have a place in mission-critical systems when so many newer databus protocols have appeared on the scene? Real-time deterministic capabilities and reliability, coupled with a vast installed base, have ensured 1553's ubiquity. Current and future military platforms will engage newer databus protocols, but 1553 will still be around for quite some time.

The evolution of 1553

In 1968, the Society of Automotive Engineers (SAE), a technical body of military and industrial members, established a subcommittee to define a serial data bus to meet the needs of the military avionics community. Several years of military and government reviews and changes led to the August 1973 release of MIL-STD-1553, with the U.S. Air Force's F-16 as the initial standard's primary user.

After some "real world" experience, program managers soon realized that further definitions and additional capabilities to the standard were needed. The SAE group spent three years of concentrated effort to produce the follow-on standard, MIL-STD-1553B, which was released in 1978. At that point, the government decided to *freeze* the standard at the "B" level to allow component manufacturers to develop products that could be immediately deployed. And today, the venerable 1553B standard is still in use.

A look inside MIL-STD-1553

The MIL-STD-1553 time division multiplexing data bus is the most ubiquitous military data bus today, used in platforms where data integrity and system reliability are critical. MIL-STD-1553 defines a method that permits digital signal communications along a common databus network with the real-time characteristics of a direct cable connection. It saves the weight and expense of miles of individual cables used to connect subsystems such as aircraft control surfaces and navigation controls.

By defining the standards for the functional, mechanical, and electrical characteristics of the serial data bus commonly used in the avionics of military aircraft, MIL-STD-1553 has been utilized to integrate computers into aircraft platforms and provide greatly enhanced reliability and performance while reducing costs.

MIL-STD-1553 applications

Since its inception, MIL-STD-1553 has found numerous applications. Notice 2 to the standard has removed all references to *aircraft* or *airborne* so as to not limit its use to aircraft platforms.

Space and military venues

Although the standard has been applied to satellites and payloads within the space shuttle and International Space Station realms, its military applications are the most numerous and far ranging. It is employed on large transports, submarines, aircraft refuelers, bombers, tactical fighters, and helicopters. It is even contained within missiles and serves in some cases as the primary interface between the aircraft and missile. The U.S. Navy has applied the data bus to accommodate both surface and subsurface ships. And the U.S. Army utilizes 1553 in its helicopters, tanks, and howitzers.

Additionally, the central role that the 1553 standard plays in weapons delivery will remain strong, even with the latest fighters. For example, the Joint Strike Fighter F-35's avionics data bus is not limited to MIL-STD-1553B, but also includes IEEE1394B Firewire and Fibre Channel for use in flight control and data display.

Because of its longevity, MIL-STD-1553 products, reference guides, and tutorials are readily available. Rugged COTS MIL-STD-1553 interface hardware, such as the RXMC1553 XMC.0 Mezzanine Card from GE Intelligent Platforms, are also building the 1553 ecosystem. Depicted in Figure 1, the XMC combines high-speed encoding and decoding and intelligent protocol processing with advanced Application Programming Interface (API) software that reduces 1553 application development time.

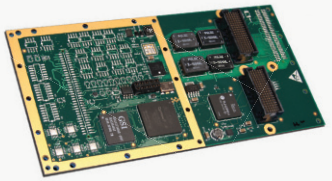

Figure 1 | The RXMC1553 XMC.0 Mezzanine Card from GE Intelligent Platforms

Transportation and government

MIL-STD-1553 products have even found their way into transportation applications such as San Francisco's Bay Area Rapid Transit (BART) system. Other non-military applications include nuclear reactors and manufacturing production lines.

Meanwhile, government adoption of MIL-STD-1553B includes acceptance and implementation by the North Atlantic Treaty Organization (NATO) and many foreign governments. The UK has issued Def Stan 00-18 (Part 2), and NATO has published STANAG 3838 AVS, both of which are versions of MIL-STD-1553B.

Looking ahead: Is MIL-STD-1553 a panacea?

Higher-performance data bus technologies such as IEEE1394B Firewire, Fibre Channel, and GbE are moving into areas once solely occupied by MIL-STD-1553. Today's modern aircraft use a mix of high-performance data buses and 1553. These newer technologies bring higher network speeds than 1553's 1 Mbps and are well suited for high-bandwidth applications such as video transmission and sensor displays, coexisting with MIL-STD-1553 but not supplanting it completely.

The 1553 protocol's extensive use in platforms and applications with service lives spanning decades will ensure the standard will be supported and employed by the military, government, transportation, and commercial sectors for years to come.

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

Mil Tech Insider

I/O mezzanines for FPGA-based reconfigurable computing

By Steve Edwards

Which mezzanine format – the PCI Mezzanine Card (PMC), Switched Mezzanine Card (XMC), or FPGA Mezzanine Card (FMC) – is best for reconfigurable FPGA-based computing applications? The answer depends on design issues such as application details, perception of risk, development timeline, and personal preference. FPGAs are ideal for high-bandwidth I/O applications because they connect directly to I/O devices to ensure low latency and support high-speed front-end DSP.

Any of these three mezzanine approaches should be compared with monolithic board technology, for example, a single Printing Wiring Board (PWB) with all FPGA functionality onboard. While representing the least flexible and possibly most costly approach, the monolithic card usually provides the best technical option because it does not have the restrictions imposed by segmenting the design, such as the number of connector I/O pins to the mezzanine. However, depending on application requirements, it is probable that rugged FPGA I/O will be served by the PMC, XMC, or FMC open standards.

PMC and XMC with FPGAs?

PCI-X (PMC) or PCI Express (XMC) latency is typically in the order of 1 to 2 microseconds and delivers bandwidths of just a few GBps, which is sufficient for some of the more demanding applications, but not for all. For some FPGA applications, PCI, PCI-X, PCI Express, and Serial RapidIO can actually dilute the advantages of using FPGAs, which excel with parallel streaming dataflows. The FMC module, similar in height and width to a PMC but approximately half the length, is designed to carry only the I/O devices and connect them directly to the FPGA device on the host board. This approach enables interface optimization among the I/Os, and the FPGA and also delivers a reduction in real estate, cost, latency, and power while boosting bandwidth.

FPGAs are ideal for high-bandwidth, high-resolution I/O because of the many I/O connections provided. This is especially true if high-speed memory is required to buffer the input or output data streams. These applications usually demand the largest FPGA packages with the most available I/O pins. Because these devices tend to be large, measuring 35 x 35 mm or more, they can violate the PMC or XMC specifications' "no go" area across the module's middle where components are banned. This restricted area forms part of the card's primary thermal interface for conduction-cooled cards and serves as a mechanical fixing area to marry with stiffening bars on the host. The result: Using large FPGA packages on PMC or XMC cards can encroach on the real estate where the designer would ideally want I/O devices placed.

Relative sizes of PMC/XMC and FMC modules

The FMC has a simplified power requirement that frees up valuable module real estate for additional I/O. FMC measures only about a half the PWB area of a PMC or XMC, but an FMC can often provide greater I/O functionality compared to PMCs and XMCs because of the PMC and XMC "no go" restriction mentioned earlier. The useful space for I/O devices on PMC/XMC can sometimes end up only a quarter of the overall real estate of the XMC. Figure 1 shows the relative sizes of PMC, XMC, and FMC modules.

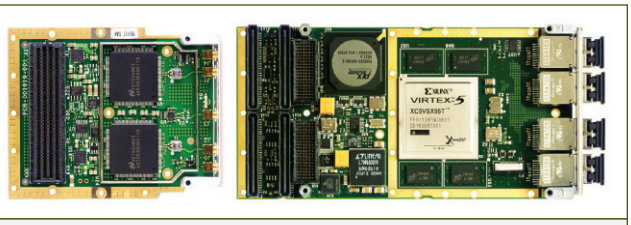

Figure 1 | The relative sizes of PMC, XMC, and FMC modules

Rugged PMC and XMC versus FMC cooling

FMCs also simplify cooling challenges, especially on small form factor cards. When plugged onto a 3U host card such as 3U VPX, a PMC or XMC covers the majority of the host's real estate. Any hot devices on the host will be located beneath the XMC, seriously affecting cooling. Furthermore, an XMC mezzanine's devices are located facing the host, which places the heat generating devices opposite those on the host and increases the cooling problem. In 6U designs, the situation is not much better. While some of the 6U host's real estate is not covered by the mezzanine card(s), the thermal paths to either the cooling air inlet or cold wall interface are longer.

Because of FMC's smaller size, less of the host board is covered by the mezzanine. Appropriate FMC host design allows for suitable heat sinks to be implemented in the areas not restricted by mezzanine placement. And since only the I/O devices, not the FPGA, are on the module, the FMC is easier to cool. The FMC specification limits the power dissipation of a single-width module to 10 W.

FMC complements PMC and XMC

While FMCs were designed for front-panel I/O, Curtiss-Wright Controls Embedded Computing (CWCEC) addresses this limitation with right-angle connectors and a specially designed strain-relief bracket to secure the connectors and minimize vibration damage to delicate connectors. This has been implemented on the 3U FPE320 board and is designed into or planned for all of the company's future FMC-based products.

FMC does not compete with PMC or XMC, but rather complements it, specifically for high-bandwidth, low-latency applications. Depending on the application, it is likely that rugged FPGA I/O will be serviced by one of these open standards. Where they cannot meet the requirements, the monolithic approach provides the safety net. Determining which approach works best will depend on application specifics.

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

We create space…

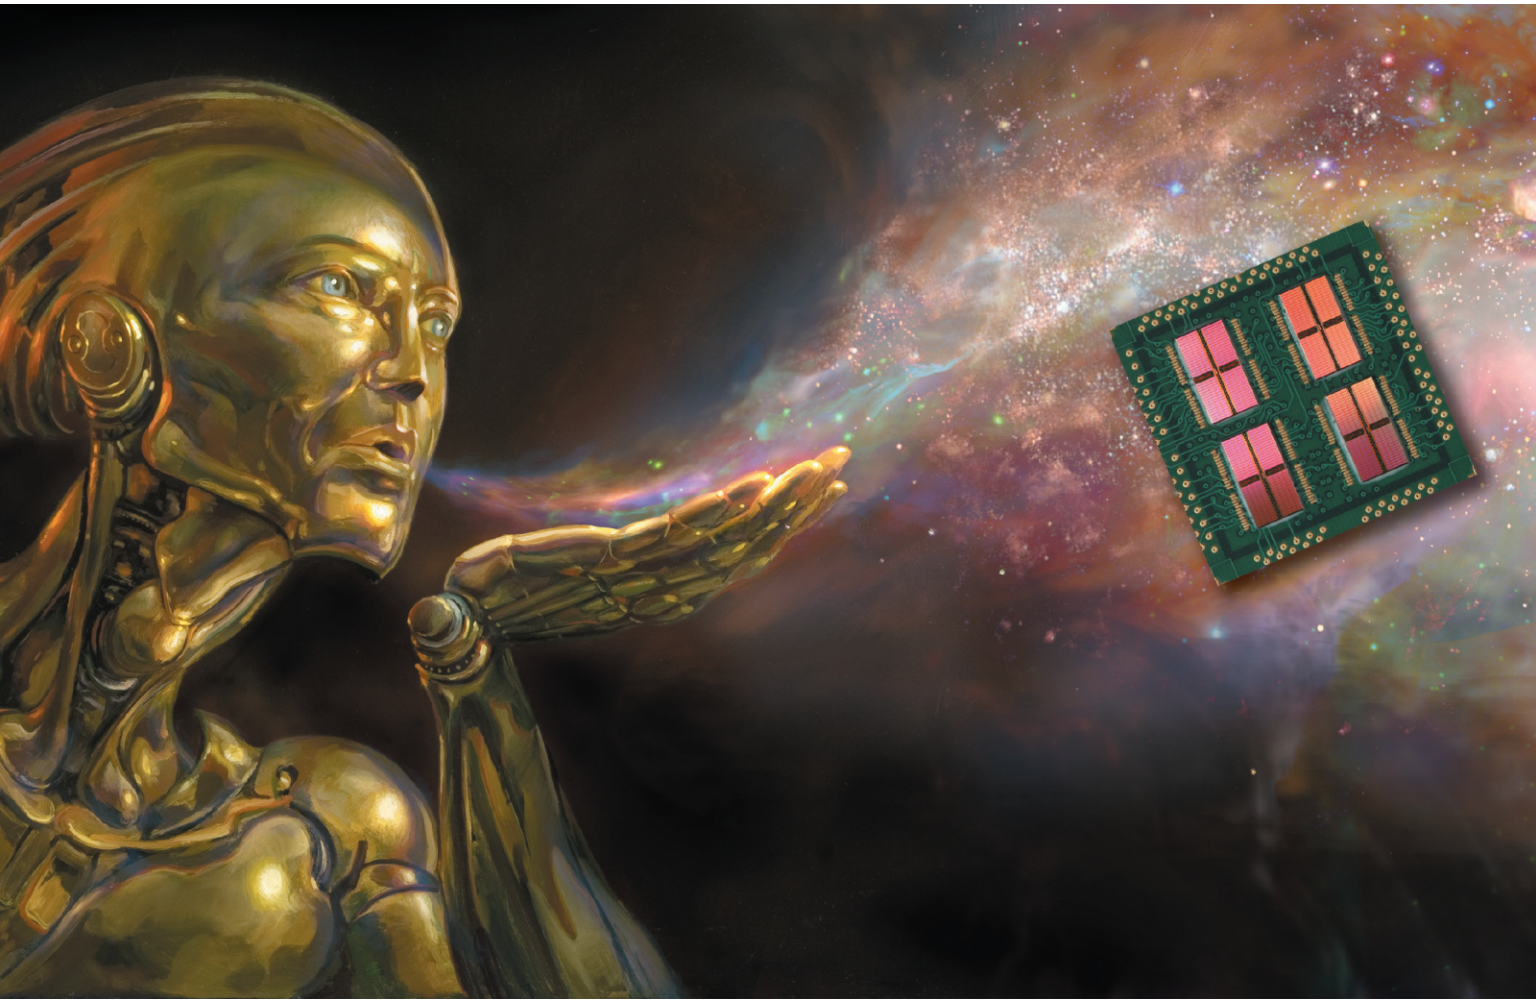

Microsemi's lean microelectronic solutions increase board density, solve component incompatibilities, reduce design complexity, and extend product life and environmental performance in military applications. We offer turnkey design, assembly and test of custom multi-chip solutions, and a wide range of standard military off-the-shelf space-saving solutions.

Expand the possibilities; visit **www.whiteedc.com/spacesaving.**

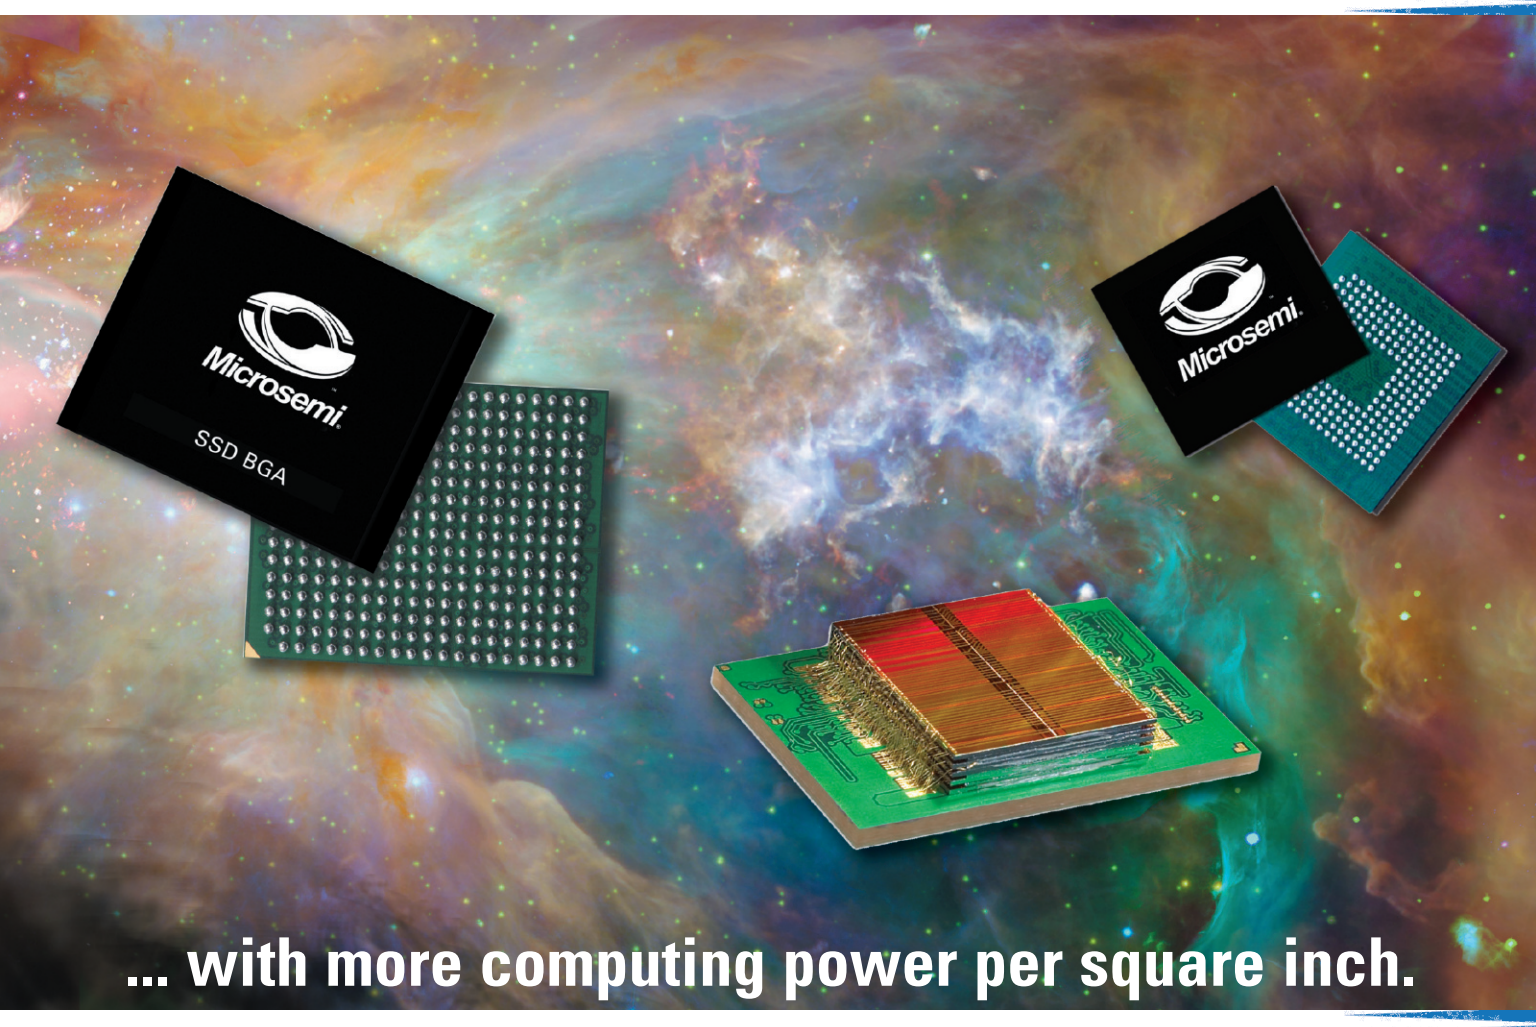

... with more computing power per square inch.

Our SLC NAND flash SSD BGA provides big features in a miniaturized package, freeing up precious board space. It is engineered specifically for the defense market, designed to ensure data reliability in mission critical rugged and mobile systems. Constructed using a 32-bit RISC processor as its core storage controller, this SSD is available in 4, 8 and 16 GByte densities.

Learn more at **www.whiteedc.com/ssd.**

602.437.1520 TEL | 602.437.9120 FAX

Microsemi

Securing the Future Through the Power of Design

WWW.WHITEEDC.COM

Legacy Software Migration

By Mark Snyder

Managing legacy graphics and user interface software

Graphics represent a challenging area for managing software obsolescence. Most applications that render to the screen or manage a user interface make use of low-level driver calls, such as OpenGL, and extensive User Interface (UI) libraries such as X-Windows or other UI libraries. Deploying these software applications on new platforms when technology insertion is needed because of performance requirements changes or obsolescence is challenging. However, a methodology whereby graphics and UI interfaces can be abstracted without sacrificing performance or flexibility is vastly extending the useful life of legacy software when technology insertion is required.

Why are graphics and UI a legacy problem?

User interface and graphics software represents a significant proportion of any software system that requires user interactions. In military systems, User Interfaces (UIs) range from real-time embedded displays through extensive digital map and geospatial C4ISR applications. User interface software is unique in that it is highly visible to the computer user while at the same time interfacing directly to graphics hardware through driver calls. Such driver calls include those provided by standards such as OpenGL or by APIs such as rendering libraries provided for a radar processor to draw reticles or text.

In addition, user interface software is highly dependent on operating system constructs such as a windows manager (for example, X-Windows), often also used to provide 2D graphics capabilities. Many times the chief impediment to software portability is in the UI and graphics area. This is even more true when trying to manage legacy software migration to systems that require technology insertion, such as replacement of obsolete graphics chips.

Many tools and methods have been developed during the years to aid in making UI and graphics applications more portable and resistant to obsolescence. For instance, many popular software tools and libraries that abstract the interface to the graphics hardware are available for developers' use. These tools might employ code generation (example: Presagis VAPS) or pass graphics rendering calls through an abstraction library (example: Nokia QT).

These development tools are very beneficial, but they do not help much with a legacy code base that has extensive dependence on graphics calls. In many cases, technology insertion requires investment in new graphics drivers or window management software to match legacy software's API needs. The API calls required by the legacy application are implemented in the hardware device driver. In practice, this means that hardware evolution is slowed, and that hardware capabilities do not evolve as they should. For military systems, sadly, it also often means that developers cannot take advantage of new technologies that would greatly benefit the system, such as low-power, low-cost computing platforms that are revolutionizing mobile computing. Industry standard groups such as Khronos (www.khronos.org) have paved the way for new, low-power graphics and UI systems through defining and supporting new standards such as OpenGL ES for embedded graphics. However, these standards are designed to enable UIs and rich graphics experience with drastically reduced APIs and embedded driver sizes. Few would argue that they deliver rich, useful user interfaces, but legacy software cannot take advantage because it was not written to support these new, streamlined APIs.

Simply put, graphics software obsolescence is leaving military systems stuck in the past. This means that soldiers, sailors, and airmen are using computers that cost more and consume more power than they need to, and often deliver substandard operation. In addition, new technologies that make smartphones, TabletPCs, and automobile computers ever more user-friendly and useful cannot find their way to military embedded systems.

New software approaches

Several approaches exist to bring legacy graphics software applications onto these new hardware platforms. One approach that offers much promise is to leverage the concept of virtualization and apply this to graphics and UI systems. Using this approach, graphics-rendering commands – whether legacy APIs, 2D window-management commands, or 3D rendering – can be implemented through a virtualization mechanism similar to those used for other OS and device interfaces. The main challenge in implementing such a scheme lies in preserving the performance and real-time interface to the hardware and graphics resources while still supporting the needed APIs. For this reason, a good virtualization solution must be built from the device-driver level to ensure success. A "hybrid" between a traditional driver and a virtualized one is necessary.

Such a "hybrid" approach can offer high-performance, divergent back-end implementations for new devices in addition to front-end APIs that legacy applications expect. Graphics resources and commands must be marshaled by this layer to deliver real-time performance. The virtualization layer must be prepared to manage graphics resource memory and virtualized device interfaces such as windows. It must also be ready to perform seamless low-overhead API conversions as required.

Embracing the future

The best approach for managing obsolescence in graphics and UI applications is to design and develop for it up front, by choosing a development strategy and tools that attempt to "future proof" an application and by adhering to relevant standards such as those offered by Khronos. When this is not possible, a virtualization strategy, or "hybrid" device driver, can offer a promising alternative. Work remains to be done, but graphics virtualization is a new and promising area. It offers a way forward to bring legacy military graphics systems into the new computing world vigorously charted by smartphones, TabletPCs, and the devices that are changing everyday lives.

> "Simply put, graphics software obsolescence is leaving military systems stuck in the past. This means that soldiers, sailors, and airmen are using computers that cost more and consume more power than they need to, and often deliver substandard operation."

Mark Snyder *is Vice President of Product Marketing for ALT Software. He has been responsible for many innovations in user interface and real-time computer graphics software architectures and tool chains during his 25-year career working at Quantum3D, Honeywell International, and the U.S. Air Force. He can be contacted at msnyder@altsoftware.com.*

Daily Briefing: News Snippets

By Sharon Hess, Assistant Managing Editor

www.mil-embedded.com/dailybriefing

NASA's TDRS gets more modern

NASA's Tracking Data and Relay Satellite (TDRS) constellation is about to get a facelift, per a recent contract awarded by NASA to General Dynamics C4 Systems. The IDIQ contract, with an option-inclusive potential value of $642.2 million, stipulates that General Dynamics modernizes TDRS's network and ground system. Specifically, a new ground-system architecture will be implemented to perpetuate tracking coverage and space-to-ground telecommunications for near-Earth and low-Earth missions. Because the TDRS network supports the entire gamut of NASA human flight and scientific space missions, including the International Space Station and Hubble Space Telescope (Figure 1), integration will be a strong focus to avoid interruption of present space network operations. Contract work is slated through June 2017. The TDRS modernization falls under the NASA Space Network Ground Segment Sustainment (SGSS) project umbrella.

Figure 1 | Hubble Space Telescope image courtesy of NASA/hubblesite.org

"GMTI radar aboard the WASP UAV?"

… That is the question on DARPA officials' minds these days, thus the agency recently commissioned OEwaves to find out just whether GMTI onboard the WASP UAV was feasible – or not. OEwaves conducted the study utilizing both conventional RF technologies and the company's proprietary microwave photonics technologies. OEwaves demonstrated, via Whispering Gallery Mode (WGM) optical resonator technology, that a GMTI-parameter-resembling transmit/receive capability was possible – and compatible with the WASP's power envelope, weight, and size. The study's findings are anticipated for use in an unnamed future DARPA project. Meanwhile, OEwaves' crystalline WGM optical resonator technology, in addition to the photodetector and miniature resonator's small bandwidth, is touted to deliver "one-step photonic mixing," as opposed to traditional super heterodyne receivers that require several steps in filtering and mixing.

DARPA project aims to reduce surveillance woes

Sorting through copious amounts of surveillance video data can prove a daunting task – and one that typically results in many false positives – but one that can be managed more effectively. Case in point: DARPA and Kitware recently signed a nearly $14 million contract stipulating that Kitware develops software that can interactively and automatically glean actionable intelligence derived from Wide Area Motion Imagery (WAMI) video of rural, suburban, and urban areas. The software aims to thwart the stops and occlusions that can fragment tracks, particularly in urban locales. The problem solver: algorithms that can integrate local events, identify threats, then conduct forensic analysis. The algorithms will then spawn the Persistent Stare Exploitation and Analysis System (PerSEAS) software prototype, which will utilize context, normalcy, and activity models to provide operator ease in identifying new intelligence. To reduce the frequency of false alarms, the models are relationship-based to make suspicious anomalies more evident versus benign events. The contract's Phase I research is slated for completion in 2012.

Raytheon APY-10 goes international

A recent contract between Boeing and Raytheon Company serves as impetus for development of the first international version of Raytheon's APY-10 surveillance radar. The international APY-10 iteration is slated for installation on Boeing's in-development P-8I maritime surveillance aircraft, the P-8A's first international variant procured by the Indian Navy. APY-10's mission is to provide actionable, accurate intelligence for ISR and anti-surface and anti-submarine warfare, come rain or shine. Raytheon is additionally working to fulfill its U.S. Navy/Boeing contract to render six APY-10s for the P-8A (Figure 2), four of which were already delivered.

Figure 2 | P-8A Poseidon, photo courtesy of Boeing

Simulation eases training at Whidbey Island

Tactical training just got a whole lot easier, not to mention safer, for EA-18G Growler and EA-6B Prowler (Figure 3) operators stationed at Whidbey Island, the locale the entire host of the aforementioned aircraft attack squadrons call "home." The facilitator: a recent contract between the U.S. Navy and Tactical Communications Group (TCG), specifying that TCG's Ground Tactical Data Link System (GTS) tactical training and simulation system be used for rendering realistic tactical training for ground operational personnel and pilots. GTS, geared for terrestrial TDL networks, enables military air operators to establish live tactical data link networks and serves to automate and simplify Link 16 network entry processes. Combat scenarios are also realistically rendered via dual-screen display of both the cockpit and the full tactical situation at hand, to increase situational awareness capabilities.

Figure 3 | Training just got a whole lot easier for EA-6B Prowler (pictured) and EA-18G Growler pilots and operators at Whidbey Island. U.S. Navy photo by Mass Communication Specialist 2nd Class John Philip Wagner Jr.

Harris Corp.'s Falcon radios travel to Asia

Harris Corporation's Falcon radios are about to embark on a very long voyage: The company has reported multiple orders, amounting to $99 million, from an unnamed military customer in Asia. Specifically, the RF-7800S, RF-7800V, RF-7800M, RF-5800H, and RF-7800W Falcon II and Falcon III software-defined tactical radio systems will form the backbone of a multilevel C4I system's next phase, to be completed by Harris. Falcon radios are suited for vehicular, handheld, or manpack radios and additionally meet the U.S. military's JTRS specifications for global netcentricity. The RF-7800S wearable secure radio affords soldier communications over distances exceeding 2 km; the RF-7800V renders data connectivity up to 192 Kbps; the RF-7800M wideband tactical radio offers networked comms over the 30 MHz to 2 GHz frequency range; RF-5800H is a High Frequency (HF) radio providing Automatic Link Establishment (3G-ALE); and RF-7800W is a wireless IP networking radio providing secure broadband connectivity.

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

Figure 4 | The HULC exoskeleton will soon undergo $1 million worth of testing. Lockheed Martin photo

Lockheed Martin's HULC flexes its muscles

Perhaps similar to its sound-alike namesake, Lockheed Martin's HULC will soon get an opportunity for muscle flexing in the form of testing – per a recent $1.1 million contract between the company and the U.S. Army Natick Soldier Center. The contract affords testing of the upgraded HULC system, a hydraulic-actuated anthropomorphic exoskeleton featuring robotic legs powered by an onboard microcomputer and optimized control software to facilitate wearer-synchronized exoskeleton movement. HULC is designed to ease dismounted soldiers' injuries and fatigue resulting from carrying heavy loads of combat equipment. HULC is battery powered and can provide crawls, deep squats, and upper-body lifting, with virtually no operator exertion. Testing will comprise HULC's effects on operator performance, soldier energy expenditures during HULC use, and operator learning speeds when moving around or carrying different loads.

USMC "System of Systems" garners upgrades

The Marine Corps Systems Command in Quantico, Va., and L-3 Communications, Nova Engineering, Inc. recently put pen to paper for upgrade/repairs, equipment systems, and program management pertaining to the Tactical Remote Sensor System (TRSS) System of Systems (SoS). The IDIQ contract is dollar-based, as opposed to quantity-based, with a ceiling of $52,781,000. TRSS SoS is designed to provide remote activity monitoring close to or in a specified locale, and in all weather conditions (Figure 5). The first delivery order on contract is anticipated at $10,854,132, and work will commence in Cincinnati, Ohio. Meanwhile, contract completion is scheduled for July 2015.

Figure 5 | The Tactical Remote Sensor System (TRSS) System of Systems (SoS) will soon be getting upgrades, repairs, equipment systems, and program management per a recent contract. U.S. Air Force photo by Senior Airman Alexandra Sandoval

RFID and asset authentication: Enabling true security measures

By Neil Mitchell

Tracking military wares is integral to efficient logistics, cost savings, and most importantly, security. RFID tags are often used for security purposes, but the results can be disastrous when RFID tampering occurs. However, there is a remedy.

Streamlining logistics, compliance demands, and tuning processes to increase cost savings, efficiencies, and revenues is only part of the overall benefits RFID can deliver. In areas where security is of great importance, authentication is a piece of the RFID system that is often overlooked. Without the proper technologies in place, certainty that an RFID tag is tracking the item it is intended to track just cannot be guaranteed.

For instance, what if an RFID tag separates from its asset or is applied to another object or container? Or, worse yet, what if the RFID tag remains on a container but the contents of that container are removed or replaced? In either case, the accuracy of the RFID system is compromised.

These are just a few examples of why it is critical to maintain a one-to-one relationship between the tag and the asset. Without this, the benefits of RFID will not be fully realized. Physical tamper protection of the RFID tag itself is the solution to this issue (Figure 1). If a tag is tampered with, the tag is permanently altered or destroyed.

If implemented correctly, the RFID tag can have the option to secure assets by serving as a seal. With special consideration given to the tag and reader location, as well as the need for inclusion of other electronics, the ability to read the tag confirms the authenticity and wholesomeness of the asset.

Figure 1 | A tamper-evident RFID tag in the process of being destroyed during a tamper event

An RFID system without these types of physical tamper-evident indicating tags is an incomplete system that invites fraud and compromises security.

Neil Mitchell, *VP Marketing for MIKOH Corporation Limited, delivers expertise in security solutions for automatic vehicle identification, secure asset tracking, digital marking, and secure network online identification. Over the past 20+ years, he has worked for STMicroelectronics, NXP Semiconductors, and Xceive. Mitchell can be contacted at nmitchell@mikoh.com.*

MIKOH Corporation Limited
www.MIKOH.com

A new era, a new ESC

New tracks. New courses. New focus.

Entering into its 21st year in Boston, ESC brings together systems architects, design engineers, suppliers, analysts, and media from across the globe. With cutting edge product demonstrations, visionary keynotes, and hundreds of essential training classes, ESC is the ideal conference for the embedded design community to learn, collaborate, and recognize excellence.

ESC Boston 2010 Tracks Include:

- Aerospace and Military
- Designing with Open Source Software, including Linux and Android
- Developing for Windows Embedded
- Graphics, Display and Lighting
- Industrial, Automation, and Robotics
- Medical
- Multicore and Virtualization
- Multimedia and Signal Processing
- Networking, Connectivity, and Embedded Internet
- Programmable Logic
- Project Management
- Real Time System Development
- Safety and Security
- Software Debugging Techniques
- Systems Integration and Test

Start your own personal development at ESC 2010. You can't afford to miss it.
Expo Floor Access is FREE!

Boston, MA
Hynes Convention Center
Conference: September 20-23, 2010
Expo: September 21-22, 2010

Register Today.
www.embedded.com/boston

Learn today. Design tomorrow.
ESC

Rapid prototyping helps U.S. military achieve optimal DSP systems for ISR missions

By Dr. James A. DeBardelaben

Confronted by rapidly evolving threats, the DoD can look to COTS-based rapid prototyping of cost-effective, high-performance Digital Signal Processing systems to meet stringent Size, Weight, and Power (SWaP) constraints and mission objectives for Intelligence, Surveillance, and Reconnaissance (ISR) operations.

U.S. Air Force photo by Staff Sgt. Joshua Strang

Modern warfare is forcing the U.S. military to quickly adapt to rapidly evolving asymmetric threats in an effort to maintain U.S. military tactical situational awareness. However, the traditional dependence on custom, stove-piped Digital Signal Processing (DSP) system implementations consisting of Application-Specific Integrated Circuits (ASICs) and Application-Specific Standard Products (ASSPs) is restricting the agility of U.S. military Intelligence, Surveillance, and Reconnaissance (ISR) operations.

To address these limitations and volatile threats, the U.S. DoD is increasingly funding ISR Quick Reaction Capability (QRC). However, this new acquisition model challenges tactical ISR system developers to rapidly prototype cost-effective, high-performance DSP systems while meeting stringent Size, Weight, and Power (SWaP) program requirements. To satisfy these demands, a novel rapid prototyping methodology leverages COTS hardware/software signal processing technology and cost modeling to close the data collection/analysis gap, thwart the limits of traditional ISR design methodologies, and reduce development time and expense while meeting challenging ISR mission requirements.

Closing the collection-analysis gap in ISR mission scenarios

The increased demand for ISR capabilities has led to an exponential increase in data collection capacity over the past few years and will continue to do so for the foreseeable future. However, ISR data Processing, Exploitation, and Dissemination (PED) processes have only improved linearly over the same period, leaving a critical gap between collection and analysis capabilities. To close the collection-analysis gap, the DoD needs high-performance DSP-based ISR systems that enable automated, real-time processing of massive amounts of data and dissemination of actionable intelligence directly to the warfighter in the field.

High-end ISR applications such as real-time, automated PED push the limits of state-of-the-art COTS DSP technology. Throughput requirements may exceed tens of tera-ops/s. New many-core Graphics Processing Unit (GPU) and General Purpose Processor (GPP) architectures appear to be theoretically capable of satisfying such high-end performance requirements. It is, however, an extremely difficult task to develop parallel software algorithms to fully exploit more than only a fraction of the peak performance of many-core architectures.

The rapid prototyping of cost-effective system implementations that meet such extreme performance requirements under severe SWaP constraints is a monumental task. Detailed trade-off analysis and extensive architectural exploration during the front-end design process are critical to accomplishing this goal. Table 1 shows a comparison of many-core GPU, multi/many-core GPP, multi-core DSP, and Field Programmable Gate Array (FPGA) technologies with respect to processor-specific software tool availability, processor peak throughput, and power efficiency. The maturity of processor-specific software development tools has a significant impact on system development effort, cost, and schedule. System developers must optimally trade off the maturity of software tool support with hardware performance and power efficiency to meet high-end ISR application requirements.

Limitations of traditional ISR design methodologies

Most of today's ISR systems follow the "waterfall" development method, which dictates a sequential process. Current waterfall-type design processes for high-performance ISR systems impose a number of limitations, including:

- Limited architectural exploration
- Lengthy prototyping times
- High cost of design
- Lack of systematic hardware/software reuse
- In-cycle hardware fabrication and testing

	Processor-Specific SW Development Tool Availability			Performance and Power Efficiency Parameters	
Processor	**Software Tool Usage Description**	**SW Tool Usage Rating**	**Effort Multiplier**	**Peak Throughput per chip**	**Power Efficiency**
Tilera TILE64 Many-core GPP	ANSI standard C / C++ compiler; profiling and debugging tools for multi-core programming; SMP Linux; multi-core component libraries for interprocessor core communication	**Nominal** (Lacks DSP software library support)	1.0	Fixed Point 443 GOPS	22 W (all cores active at 700 MHz)
NVIDIA Tesla C2050 Many-core GPU	C / C++ / Fortran compiler; CUDA-GDB debugger; CUDA Visual Profiler; OpenCL Visual Profiler; GPU-accelerated BLAS library; GPU-accelerated FFT library	**High** (Mature tool and library support for DSP parallel programming)	0.86	Single Precision Float 1.03 TFLOPS Double Precision Float 515 GFLOPS	247 W TDP
Intel Core i7 -920 Multi-core GPP	C / C++ / Fortran compiler; Math Kernel Library; integrated performance primitives for multimedia, data processing and communications; threading building blocks; performance analyzers	**Very High** (Very mature and well integrated tools)	0.72	Single Precision Float 85 GFLOPS	130 W TDP
TMS320C6472 Multi-core DSP	Code Composer Studio IDE; data visualization kit; math library; signal processing library; image library; floating point emulation	**Nominal** (Good library support; basic multi-core programming support)	1.0	Fixed Point 256 GMACS Single Precision Float 128 GFLOPS	5.4 W TDP
Altera Stratix V FPGA	Quartus II design software including support for planning, simulation, timing analysis, debugging, verification, and synthesis	**Very Low** (Time-intensive verification and validation process)	1.24	Fixed Point 1840 GMACS Single Precision Float 1TFLOPS	110 W TDP

Table 1 | System developers must optimally trade off the maturity of software tool support with hardware performance and power efficiency to meet high-end ISR application requirements.

Most design automation activities have focused on leveraging tool support for detailed system behavioral design, as opposed to early architecture design where much of the system cost is committed. Current industrial practice predominately relies upon designer experience to select system architectures and allocate algorithm functionality. Furthermore, for fully customized ISR systems, hardware and software subsystems are not integrated until after hardware is fabricated, making design errors very costly.

Grounded innovation: COTS-based rapid prototyping

The increased frequency of ISR system development cost overruns and schedule delays has compelled the DoD acquisition community to launch numerous initiatives that encourage the contractor community to better leverage COTS DSP hardware boards and system components.

In COTS hardware-based systems, the time and cost of software development can, however, dominate the schedule and budget. Parametric studies based on historical project data show that designing and testing embedded software is particularly difficult if margins of slack for processor and memory resources are too restrictive. Severe resource constraints may prohibit embedded software developers from leveraging high-level programming tools, thereby requiring direct interaction with the hardware and/or operating system to optimize code to meet system requirements.

One innovative solution garnering attention in the industry is a rapid prototyping methodology that exploits the use of COTS hardware/software signal processing technologies and

cost modeling to achieve significant reductions in total ISR system cost and development time. This approach leverages a COTS library-based optimization framework that includes trade-offs in throughput, hardware/software development costs, and schedule, procurement costs, and SWaP. This rapid prototyping methodology maximizes system architectural exploration in the front-end design process. The resulting solutions are cost-effective DSP embedded systems that exploit the flexibility of many-core GPU, multi/many-core GPP, multi-core DSP, and/or FPGA technologies, while satisfying stringent SWaP constraints dictated by mission objectives.

Figure 1 illustrates the COTS-based rapid DSP system prototyping methodology. The process starts by translating written system requirements into executable requirements and specifications using signal processing libraries and integrated Graphical User Interface (GUI) toolkits such as those available in MATLAB. The executable requirements and specifications provide an early prototype to the customer to validate original requirements and to remove ambiguities. This feedback allows for receiving any requirement alterations as early as possible, which is critical to minimizing the high cost of requirements creep, one of the most common risks in software projects.

After the validation of system requirements, system-level cost parameters, application requirements, and performance statistics, these components feed the architecture selection and partitioning optimization process. System developers can use parametric cost models, such as COCOMO II, to drive the architecture trade-off analysis, producing hardware/software architectural candidates that minimize total system cost and development time. The cost parameters include: software cost driver attributes (size, product, platform, personnel, and project), COTS hardware procurement costs, product deployment deadlines, schedule constraints, and labor costs and constraints. Application requirements include SWaP, environmental, precedence, and real-time constraints, as well as functional, memory, and communication requirements. Performance statistics consist of benchmark time measurements of DSP primitives (for example, fast Fourier transform) executing on the DSP processor boards (for example, many-core GPU, multi/many-core GPP, multi-core DSP, FPGA) contained in the reuse library.

System developers can then simulate the resulting architectural candidates using dynamic performance modeling tools, such as Simulink, to verify that an architecture meets system-level requirements. After performance modeling, the system

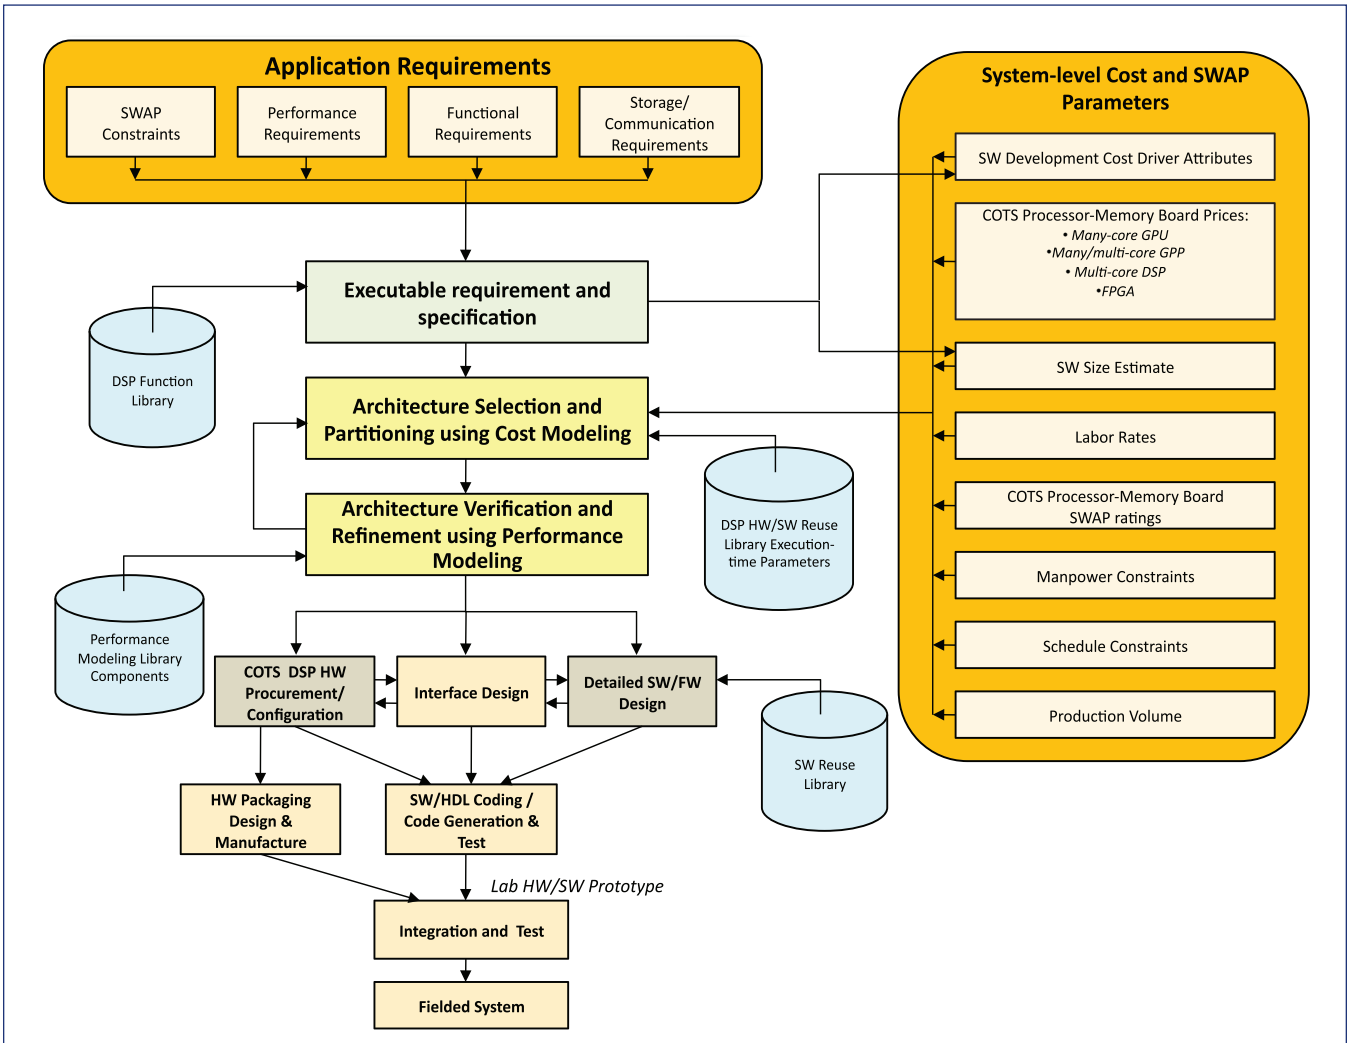

Figure 1 | The COTS-based rapid prototyping methodology prioritizes both application requirements and cost modeling, critical to minimizing the high cost of requirements creep.

architect feeds communication overhead parameters such as communication queuing delays and bottlenecks back to the architecture selection stage for refinement. The methodology produces new architecture candidates with the updated model parameters and repeats the process until the architecture meets performance requirements and no longer changes between successive iterations.

The refined hardware/software architectural candidate moves on to the detailed architecture design stage for detailed software and/or firmware design, hardware/software interface design, and COTS procurement. Depending on the COTS DSP hardware platform and architecture selected, the DSP software and/or firmware design process heavily leverages reusable libraries developed in previous projects. The cost of assessing, selecting, assimilating, and modifying the reusable component must also be minimized to significantly reduce software development expense and time. System designers can refine the candidate architecture's performance model and executable signal processing algorithm specifications to permit automatic code generation into the C programming language or a hardware description language. To enable the use of automated code generation tools, sufficient hardware resource slack margins and high-level software tool support must exist for the target DSP board architecture.

The high-level virtual prototypes of the system allow the system designer to catch hardware/software integration errors early in the design process. This approach allows for low-level performance limitations to be identified and corrected before costly hardware packaging assembly and field testing.

Rapid prototyping helps ISR systems keep pace

Traditional DSP implementations for ISR systems can no longer keep pace with modern warfare. System developers need a rapid prototyping methodology that exploits the use of COTS hardware/software signal processing technologies and cost modeling to achieve significant reductions in total ISR system cost and development time. IvySys Technologies' Real-Time Intelligence Analysis methodology fully leverages this COTS-based rapid prototyping approach. The front-end design process is automated by incorporating software cost and development time models. The design optimization process reduces development time and cost by as much as a factor of four. The IvySys COTS-based rapid prototyping methodology enables the DoD and intelligence community to quickly adapt to rapidly evolving asymmetric threats and to thereby maintain U.S. tactical situational awareness.

***Dr. James A. DeBardelaben** is the President and CEO of IvySys Technologies, LLC. Prior to founding IvySys, he worked as an expert consultant in Special Operations ISR, real-time embedded software, and tactical distributed systems and networking. James received a Ph.D. in Electrical and Computer Engineering from the Georgia Institute of Technology, an M.S.E in Computer Engineering from Princeton University, and a B.S. in Electrical Engineering with honors from Brown University. He may be reached at jd@ivysys.com.*

IvySys Technologies, LLC
703-414-5665 • www.ivysys.com

UAV system architectures: The next evolution

By David W. Lee

U.S. Army photo by Sgt. Travis Zielinski

Network-centricity and recent advances in high-speed serial fabric technologies, along with legacy device bridging methodologies, enable highly flexible UAV system architectures to meet the rapidly changing demands of the 21st century battlefield, while shortening time-to-market and mitigating risks.

The roles of Unmanned Aerial Vehicles (UAVs), like the Global Hawk UAV shown in Figure 1, are expanding. As UAVs' strategical and tactical mission capabilities continue to expand, there is an ever-increasing need for highly reliable, high-performance electronic systems that can support high-bandwidth network connectivity within the platform – and from the platform to the Command and Control infrastructure. UAVs are increasingly becoming networked systems; they are one element of the entire Unmanned Aircraft System (UAS) that communicates and shares mission-critical payload and sensor information with other elements such as ground control stations and ground/remote terminals to provide a common operating picture.

This migration from platform-centric to a network-centric systems approach poses new challenges that are met by the use of open architectures. Formerly, UAVs adopted platform-centric architectures that relied on proprietary interfaces. Many of these design concerns can now be overcome by ruggedizing commercial electronics, enabling subsystem technologies containing COTS components and allowing faster time-to-market. The result is rapid battlefield deployment of the newest technologies.

Advances in open standards performance, such as VPX high-speed serial fabrics, along with the ability to bridge legacy devices, make it possible to develop scalable and interoperable network-centric subsystems at significantly lower development and logistical risks.

Figure 1 | The demand for UAV technology is ever-increasing, and the Global Hawk is one example of such technology.

Network-centricity eases interfacing

Net-centric architecture enables the UAV system designer to standardize the interface between the computing elements and the interface elements. For example, a COTS SBC might not have the needed ADC/DAC onboard, but a network-attached Remote Interface Unit (RIU) with ADC/DAC capabilities can be provided at or near the vehicle's sensors and actuators. (See Figure 2 for a network-centric reference design.) The RIU in this case would digitize the inputs, communicate with the SBC through a standard network interface such as GbE, and generate analog outputs, therefore creating a standardized interface between the processor and the I/Os. The same approach is also applicable to Network Attached Storage (NAS). NAS improves data accessibility via Network File System (NFS) protocol while making it seamless to grow the storage capacity on the network at a later time. As long as the new device conforms to the same GbE and the file sharing protocols, the rest of the system is not impacted.

Traditionally, the inputs from platform-specific sensors are either hardwired directly to the processor or transmitted over MIL-STD-1553 or ARINC-429 buses, and the processor board must be designed with dedicated signals such as discrete, analog I/O MIL-STD-1553/ARINC-429 bus interfaces to handle that communication. This introduces an obsolescence risk as the support for

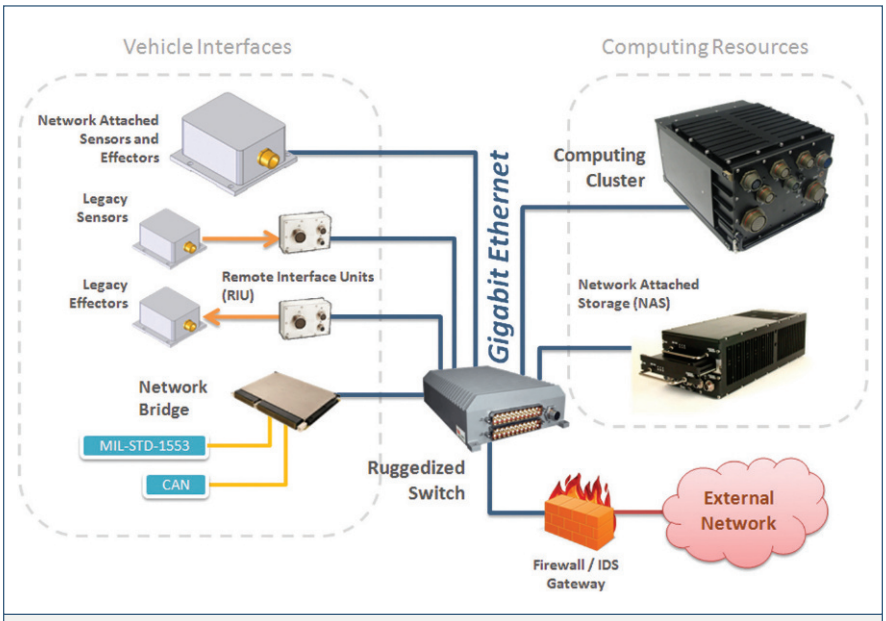

Figure 2 | A network-centric reference design

these special signals is not standardized by the industry. In other words, a custom product has to be made, which comes at steep nonrecurring and recurring prices. To eliminate those issues, taking the RIU approach, network-attached sensors and effectors can use GbE to communicate through network switches, freeing the processor card from requiring any special onboard I/O. This permits the sensor to stay in place while the processor card or electronic system can be replaced or upgraded separately, with the added benefit of reduced wire counts in the main harness assemblies. Combined with deterministic Ethernet, like ARINC-664/AFDX for avionics, flight control needs can be met.

VPX and high-speed serial fabrics

Another advantage of embracing new open architectures on UAVs is their support for distributed computing. In today's systems, it is typical for each of the multiple SBCs to be assigned specific functions. Density limits on SBCs, currently at around 18 GFLOPS for a conduction-cooled 3U card, mean that the requirements frequently can't be addressed in a single slot. A distributed or *cluster* environment, on the other hand, enables tasks to be shared amongst the SBCs, effectively multiplying CPU power to the 100s of GFLOPS required by video processing and intelligent algorithms. Working in unison with RIUs, the processors can now be physically relocated to practically anywhere within the airframe where CPU power, electronics packaging, power supplies, and cooling capacities can be managed more effectively and where space is available. Figure 3 shows a ruggedized electronic enclosure that supports high-speed interfaces and can be cooled in different ways in accordance with VITA 48.x.

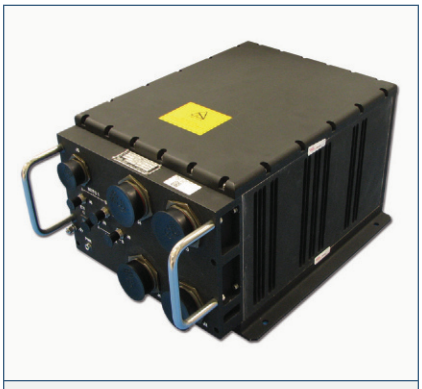

Figure 3 | A ruggedized electronic enclosure that can support high-speed interfaces and be cooled in different ways in accordance with VITA 48.x.

This type of processor-to-processor collaboration essentially takes the Symmetric Multi-Processing (SMP) and multicore model beyond the slot level to the chassis level and beyond. This collaboration requires high-speed interconnects, often at multiple Gigabytes-per-second speeds, and is now possible via open standards such as VPX (VITA 46) and its support for serial fabrics such as Serial RapidIO and PCIe. The newer OpenVPX (VITA 65) further refines pinout definition for interoperability. External communications outside the chassis can be achieved with GbE and 10 GbE. One way to satisfy the growing hunger for more processing is to standardize on communications protocols that support distributed computing.

Bridging with legacy devices

As UAV platforms migrate to a network-centric approach, legacy buses such as 1553, RS-232, and CANbus will need to communicate with newer GbE-based interfaces through bridge devices. This approach allows for an incremental upgrade that matches today's existing

budget constraints, because it is unlikely that all older LRUs can be replaced at the same time. There will be a period of time during which new electronic systems and older ones based on the legacy interfaces must coexist.

Curtiss-Wright Controls Electronic Systems' network bridge technology is used in sensor management and payload management computers to integrate legacy, proprietary, and emerging interfaces in addition to performing data fusion by translating the communications protocols from one side to another. For example, when it is time to upgrade a Remote Terminal (RT) from the 1553 network and put it onto the GbE side, the network bridge can transparently emulate the RT's I/O as though it were still on the 1553 side. The other LRUs on the UAV won't even detect that the RT has been transitioned to the GbE side. The same can be achieved for Bus Monitors (BMs) and Bus Controllers (BCs).

Thinking ahead

The demand for electronic systems with advanced multiprocessor compute technologies on tactical and strategic UAVs creates many challenges. Open interface standards such as VPX (VITA 46), OpenVPX (VITA 65), and VPX-REDI (VITA 48) – along with adoption of standardized high-speed interconnects like GbE and PCIe – enable designers to overcome these challenges. Also critical to the equation is the ability to bridge legacy devices. Hence, now is the time for designers to optimize their electronic systems with a network-centric, open architecture approach to achieve new levels of processing efficiencies and performance.

David W. Lee *is a Project Engineer for Curtiss-Wright Controls Electronic Systems. He has more than 12 years of experience in avionics and embedded systems. He holds a degree in Computer Science from the University of Buffalo and an MS in Computer Science from Binghamton University. He can be contacted at dlee@curtisswright.com.*

Curtiss-Wright Controls Electronic Systems
661-257-4430
www.cwcelectronicsystems.com

Open architecture COTS form factors: The optimal approach for UAV/UAS applications

By David French

A broad range of embedded computing form factor options matches diverse airframe size and mission requirements for today's UAVs and UASs.

"We have just won a war with a lot of heroes flying around in planes. The next war may be fought by airplanes with no men in them at all. It certainly will be fought with planes so far superior to those we have now that there will be no basis for comparison. Take everything you've learned about aviation in war, throw it out of the window, and let's go to work on tomorrow's aviation. It will be different from anything the world has ever seen," said Gen Hap Arnold, U.S. Army Air Forces (USAAF), on VJ Day, 1945[1].

Unmanned Aerial Vehicles (UAVs) and Unmanned Aerial Systems (UASs) are proving General Arnold's visionary thinking, as they emerge as two of today's most adaptable, in-demand military technologies. These programs offer endurance, efficiency, flexible mission management, attack capability, information collection, and connectivity: valuable attributes that enable military command to effectively and safely multiply forces. The vision for these unmanned systems is to harness their intrinsic benefits for automated, modular, globally connected and sustainable multi-mission purposes that maximize force utilization. The U.S. Air Force's UAS Flight Plan has predicted a leaner, more adaptable and efficient air force encompassing a broad family of unmanned aircraft that reaches well past the earliest strictly surveillance-based predecessors.

As new UAS applications go beyond mere surveillance to encompass small man-portable vehicles (even micro- and nano-sized), medium "fighter sized" vehicles, large "tanker sized" vehicles, and special vehicles with unique capabilities, there is a significant requirement to expand the embedded computing technology that controls them. This increased diversity in UAS applications and functionality supports the U.S. DoD's long-term commitment to sophisticated, technology-supported warfare. Making

this commitment a reality will challenge future UAS designs to combine multi-mission, ultra-rugged, net-centric, modular, open architecture features capable of carrying any standard payload within their given performance envelope.

Embedded technology advancements and a broad range of COTS form factors have become key enablers of new modular UAS/UAV platforms. The good news is that UAS designers have a wide arsenal of open architecture COTS form factors that will satisfy evolving SWaP, performance, space constraint, extended temperature, and sophisticated signal-processing requirements. Determining the optimal form factor for a particular UAS/UAV design requires a deep level of expertise and an understanding of the advantages and design trade-offs provided by the range of embedded computing alternatives.

Open architecture COTS form factor options

In this maze of embedded design challenges, sticking to a single computing form factor for all UAV/UAS designs is not really a relevant approach. Instead, designers must leverage and evaluate a varied slate of standards-based options. The optimal embedded computing choice meets the desired program and UAS/UAV platform goals by matching the airframe requirements and creating interoperable, COTS-based systems that integrate with the Joint Forces' worldwide information network. This includes adding capabilities such as payload, networking, and effective data processing, analysis, and dissemination. The demand is rising for net-centric implementations that enable soldiers, military systems, and equipment to share real-time information with higher bandwidth and increased computing performance. So is the need to upgrade embedded computers in existing deployed units.

From the ultra-rugged VPX to small Computers-on-Module (COMs), high-bandwidth MicroTCA, and proven CompactPCI, designers now have a richer and broader variety of proven COTS embedded computing form factor options. Designers can also leverage continued advancements to standards-based technologies to achieve their varied and evolving UAS/UAV platform objectives. Since the landscape of available open architecture form factors covers a wide set of technologies, it is wise to evaluate each on its strengths for a particular application.

VPX for high-speed signal processing

UAS payloads (short-, medium-, or long-range airframes) may include communications, signals intelligence, high-resolution Synthetic Aperture Radar (SAR), imaging systems, or even weaponry. Improved situational awareness is the primary goal, downloading compressed, live video or other information to a portable device on the ground. As a result, image compression and bandwidth demand much of the system designer's attention. Largely based on its ability to provide high-frequency processing as well as a reliable fabric solution, rugged VPX is emerging as an ideal platform for these data-intensive UAV applications (Figure 1).

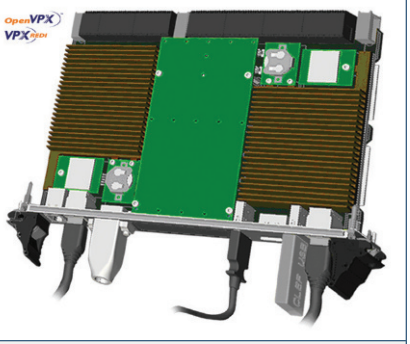

Figure 1 | Rugged VPX technologies are emerging as an ideal platform for data-intensive UAV applications. Pictured: The Kontron VPX Blade VX6060, with two independently implemented Intel Core i7 processing nodes linked to an Ethernet and PCIe infrastructure.

VPX enables higher performance processing per slot but also higher-speed interconnects between processing and I/O elements using PCIe, 10 GbE, or Serial RapidIO. These interconnects provide 10 Gbps between elements or several hundred GBps in aggregate, depending on the system implementation. Handling the increased signal rates that characterize high-bandwidth, high-performance UAS applications, VPX systems can achieve greater than 5 GBps using a number of different serial fabric technologies. Full-mesh architectures, not available in VMEbus systems, enable system aggregate bandwidth greater than 100 GBps.

The VPX platform builds on VME's processing capabilities, combining robustness and excellent EMC – fundamental strengths of the VMEbus architecture – with new high-bandwidth serial interconnects for high-speed differential signaling over the backplane. VPX leverages more I/O per slot – for example, taking data in a fast 10 GbE data rate and dispatching it to several processors that manage the workload in parallel – and maximizes higher computing density from available processors and chipsets. This is in contrast to VMEbus platforms primarily using parallel bus technology that provided successive bandwidth improvement over the years, starting with 40 MBps and evolving to today's 320 MBps.

As an example of VPX's suitability for applications that require video-streaming capabilities, VPX can be integrated with codecs including ITU-T H.263, H.264 (MPEG-4 part 10), and JPEG2000 to provide very efficient image compression. The H.264 codec is particularly optimized for streaming and offers extremely efficient compression by providing the capability to trade off image quality or compression as the available bandwidth changes. This flexibility makes using VPX with this codec a viable option for supporting UAS video payload applications that can be required to operate over a number of data link options and operational scenarios.

MicroTCA meets rugged, high-bandwidth needs

Range and altitude may be deciding factors in choosing an embedded computing form factor. UAS platforms that need to perform in longer-range missions will typically need to operate at higher altitudes as well (Figure 2). Then there are range and altitude requirements for extended loitering capabilities, too, which all translates into the need for a rugged and reliable form factor in terms of power, weight, and thermal management.

A case in point is that a larger airframe may have an avionics bay containing forced-air cooling, where many smaller airframes do not offer that capacity. Rugged air-cooled MicroTCA (MTCA.1), hardened MicroTCA (MTCA.2), and conduction-cooled MicroTCA (MTCA.3) leverage the ANSI/VITA 47 specification to define these types of environmental requirements.

MTCA.1 extends MicroTCA into more rugged military environments as defined by ANSI/VITA 47's EAC6 environmental class and V2 vibration class. The new MTCA.3 spec is now underway with PICMG and is expected to define a conduction-cooled interface that meets the most extreme thermal, shock, and vibration profiles defined in ANSI/VITA 47 (that is, performing in conduction-cooled systems with no airflow within sealed environments, common to space-constrained UAS applications).

Figure 2 | Vertical Takeoff and Landing Tactical UAVs (VTUAVs) can take off and land in rugged or unimproved terrain in close proximity to troop and tactical operations centers. Acting as a communications node within the Joint Tactical Radio System (JTRS), they extend both network effectiveness and flexibility.

These options highlight MicroTCA as a cost-effective alternative to the ultra-rugged VPX; MicroTCA is characterized by high processing capacity, extremely high communication bandwidth, and high availability designed into a small form factor. A 3U or 4U system, tapping up to 12 compute blades on a single backplane all potentially using a multicore processor, could be integrated with as many as 24 cores in a very small footprint. Designers using MicroTCA can also leverage as many as 21 high-speed serial connections on the backplane, each delivering bandwidth of 2.5 Gbps. Depending on the airframe or its ground control system, as well as how each system is implemented, an extensive range of MicroTCA-based communications bandwidth capacities is available, ranging from 40 Gbps to >1 Tbps. When contrasted to VPX's 80 GBps aggregate bandwidth, designers need to additionally understand that VPX also defines a larger quantity of user I/O options not defined in the MicroTCA spec.

CompactPCI: Proven rugged, long-term reliability

CompactPCI also provides notable bandwidth, performance, and cooling options. Its status as a mainstay UAS platform is based on advantages such as its inherent ruggedness, rear I/O, support for the full range of high-speed interfaces, and an extensive range of PCI-compatible software. With improvements in space and energy savings enabled by advancements such as Intel's 32 nm, Atom-based low-power processors and multicore architectures, many designers consider CompactPCI ideal for a self-contained network implementation managing multiple blades communicating over GbE in a single backplane. Even though CompactPCI is largely limited to 1 GB speeds and parallel PCI – dramatically reduced bandwidth as compared to a VPX backplane – it uses the same 3U/6U mechanics and so can readily accommodate the same rugged environments.

Integrated 32 nm processor technology along with multicore performance benefits mean CompactPCI systems can enable new and more compute-intensive applications such as UAVs and UASs, characterized by extreme conditions, round-the-clock performance, and high-speed processing. Multicore processor technologies and power-aware designs are enabling higher overall performance within the same or lower thermal and power envelope, essentially extending the life of existing form factor standards such as CompactPCI and delivering options for upgrading and extending the life of currently deployed systems. Existing systems could trim 10 CompactPCI 2.16 single-core boards down to just two dual- or quad-core boards. Multiple quad-core boards can be implemented

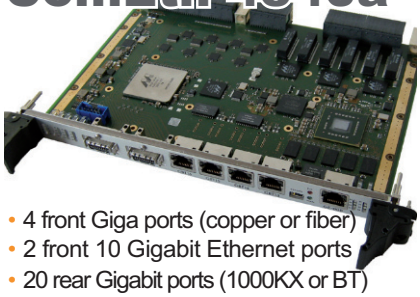

when ultra-high performance is required for extremely data hungry applications common to UASs. And low-power CompactPCI boards achieve suitable performance-to-power ratios based on passive cooling for convection-cooled and forced-airflow applications (Table 1).

Computer-On-Modules: Power and performance for constrained spaces

The small UAS class of airframes represents profound technological advancements in air warfare, with proven utility and capability during the earliest phases of Operation Iraqi Freedom. Providing not only the commander but individual service members with life-saving situational awareness, full-motion video is perhaps the most important mission of these small but imperative UAV devices.

Based on the compact Intel Atom solution, these COMs balance performance with the ability to satisfy SWaP requirements critical to these small and extremely energy-efficient UAS devices. For instance, the 45nm Intel Atom processor architecture achieves fast performance (with clock speeds between 1.1 GHz and 1.6 GHz) in a sub W thermal power envelope. It features a power-optimized front side bus of up to 533 MHz for faster data transfer. This in turn enables the development of energy-saving, high-end graphics devices such as those in UAVs or UASs without leaving the safe and proven development path of COMs as an established and future-proof industry standard. Available now, COMs that integrate Intel Core i7 advancements deliver even greater design flexibility in terms of performance and onboard features. Performance-per-watt is as good as it gets, with virtually no performance trade-offs via the enhanced I/O capabilities.

Overall, Core i7-based platforms incorporate a more efficient two-chip solution for better signal integrity and minimized board space, in turn enabling higher performance in smaller, power-constrained portable designs. For visually demanding UAV/UAS applications such as compute- and graphics-intensive imaging or persistent surveillance applications, this technology also delivers significantly enhanced integrated graphics capabilities and data flow performance.

> "The U.S. Air Force's Flight Plan states that 'standards and interoperability are keys to the Joint Forces gaining Information Superiority in today's network enabled environment.'"

Planning ahead

The U.S. Air Force's Flight Plan states that "standards and interoperability are keys to the Joint Forces gaining Information Superiority in today's network enabled environment."[2] These goals mandate a common set of airframes within a family of systems, all based on standard interfaces and interoperable plug-and-play payloads that are tailored to support one or more of the Joint Forces' priorities. This strategy represents a major shift that is significant for military embedded designers. No longer a platform-centric model, achieving success means designers must address not only how the system works but also how it integrates with the Joint Forces' worldwide information network.

With assorted and extensive application requirements that continue to evolve, designers of UASs/UAVs are wise to take an open architecture COTS form factor approach such as VPX, MicroTCA, CompactPCI, or COMs. These can deliver advancements in greater compute density, low power consumption, small footprint, enhanced thermal management, increased I/O bandwidth, and optimized processing architectures. The result: Designers have a wealth of options for proven modular systems that meet their design and platform objectives.

Processor	Thermal Design Power (TDP)	Performance in GFLOPS
Intel Pentium M (~2 GHz): 5-year-old technology	22 W	1.1
Intel Core i7 (~2 GHz): Today's technology	25 W	24

Table 1 | Performance-to-power ratios have improved, even within the past five years. In the same Thermal Design Power (TDP) envelope, designers can now get a 24x performance increase using processors capable of operating in extended-temperature environments.

***David French** is Director of Military and Aerospace Business Development at Kontron. He has more than 25 years of experience in the embedded computing industry across multiple market verticals with significant experience in defense and aerospace electronics. David has worked in both engineering and sales/business development capacities. He holds a BSEE from San Diego State University and an MSEE from Rensselaer Polytechnic Institute. He can be contacted at David.French@us.kontron.com.*

Kontron
858-677-0877 • www.kontron.com

References:

[1] *Words on War: Military Quotes from Ancient Times to the Present*, by Jay Shafritz, Prentice Hall, New York, 1990, pg. 104.
[2] USAF Flight Plan, http://www.globalsecurity.org/military/library/policy/usaf/usaf-uas-flight-plan_2009-2047.htm, page 56.

HD streaming video in wide-area surveillance systems: A reconfigurable hardware processor approach

By Adam Smith

The demand for higher-quality video imagery in surveillance systems requires an advanced link between the sensors and the usable video data. An image processing system consisting exclusively of software processing isn't capable of meeting these demands, and an ASIC-based system doesn't provide the needed flexibility. Because of the compute-intensive nature of pixel processing, an FPGA-based imaging processing design using JPEG2000 encoding provides a powerful and flexible methodology.

Air Force photo by Tech. Sgt. Erik Gudmundson

High-definition digital video has become commonplace; most of us have HD equipment right in our own homes. However, moving the technology to real-time military surveillance applications poses major challenges in capturing the HD images and distributing the enormous amounts of data. Image sensors are now being produced in the gigapixel realm, implemented using arrays of smaller megapixel detectors. Image data acquisition must be flexible to facilitate interfacing with a variety of sensors and must be closely coupled to data processing and compression in order to distribute data. Utilizing an FPGA coprocessor has many benefits over the traditional software or ASIC imaging processing system, and packs the complex algorithms into a small, powerful device.

Accordingly, FPGA-based processors can provide the critical link between HD image sensors and low-bandwidth JPEG2000 compressed data. This trend in reconfigurable computing makes it possible to satisfy the needs of powerful imaging systems in airborne and ground vehicle platforms such as Predator UAVs, Reaper UAVs, Hummingbird UAV helicopters, MULE autonomous ground robots, aerostats, and many more. High-performance, FPGA-based computing allows military surveillance systems to receive and process HD sensor data from a variety of sources and compress the data for transmission using JPEG2000.

Pixel processing is painless for FPGAs

The highly computational nature of image pixel processing can be efficiently implemented in modern FPGA devices. A Virtex-6 SX475T FPGA has more than 475 K logic cells and 2,016 embedded DSP blocks, flexible connections to the rest of the system including integrated 6.5 Gb transceivers, and many other powerful elements. These features allow many of the image processing functions to be pipelined together in a single FPGA device for multiple image sensors. The example image processing design detailed in Table 1 shows that less than 7 percent of FPGA resources is needed for a full-function image processing and compression channel.

Image processing system needs: Tough enough?

It's desired for a military surveillance system to have the capability to view a wide area, for example greater than 16 km^2, since not much information can be gained by "looking through a soda straw." Furthermore, the surveillance system must typically operate at distances many kilometers away from the target to preserve safety while generating high-resolution images. All this requires a colossal amount of image data to be captured and processed. For example, the DARPA ARGUS-IS system uses 368 5-megapixel detectors (a total of 1.8 gigapixels) capable of generating more than 260 Gbps of raw data.

This data must be captured and processed in a real-time system, often with extreme SWaP limitations. Additionally, the imaging system might need many other desired features for use after capturing raw pixel data:

- Defective pixel correction
- Color filter array interpolation
- Color correction
- Gamma correction
- Color space conversion
- Image noise reduction
- Edge enhancement
- Video scaling
- Video windowing – Object tracking
- Image stabilization
- Video compression (JPEG2000, H.264, MPEG-4)
- Data encryption
- Video output for transmission and/or storage

Feature		Logic Cells (2)	DSP Blocks	Memory
Image Processing	Defective Pixel Correction(1)	665	0	0
	Color Filter Array Interpolation(1)	8444	8	81 KB
	Color Correction(1)	421	9	0
	Gamma Correction(1)	65	0	40 KB
	Color Space Conversion(1)	812	4	0
	Image Noise Reduction(1)	3980	3	18 KB
	Edge Enhancement(1)	3546	1	54 KB
JPEG2000 Compression (ADV212 Pair Interface)(3)		14886	0	54 KB
Example Image Processing System (Resource usage in Virtex-6 SX475T device)		*32,819 (6.9%)*	*25 (1.2%)*	*247 KB (5.2%)*

(1) Example data based on Xilinx® image processing pipeline v1.0 12-bit data width, 2047x2047 resolution
(2) Estimates based on Xilinx® Virtex-5 / Virtex-6 logic cell rating equivalent
(3) Estimates based on Alpha-Data CameraLink SDK example design

Table 1 | Example image processing system – FPGA resource requirements

Reconfigurable computing technologies based on FPGA processors are excellent for meeting these demanding needs. This methodology is smaller, faster, and more efficient than an exclusively software processing system – and much more flexible and cost-effective than ASIC approaches. The essential link between the image sensors and software processing is a perfect fit for FPGA technology.

Adapting to various image-sensor interfaces

Surveillance systems and other imaging platforms are quickly evolving and expanding to utilize a variety of sensors for optical, infrared, and synthetic aperture radar technologies, and the image sizes are always increasing. The physical interface to the image sensors varies, and therefore the processing system needs a flexible and modular interface in order to scale and expand with these systems. Reconfigurable computing systems meet this need by supporting incoming data in many formats and physical interfaces, for example, CameraLink, SDI, raw pixels, LVDS, multi-gigabit serial interfaces, fiber-optics, and so on. An FPGA-based system is well suited for acquiring the image data from a variety of sources and protocols because of the programmable I/O blocks, integrated transceivers, and SERDES blocks.

Additionally, a surveillance system may have many different mission objectives that require flexibility in processing the video data. For example, the system may need to provide full-resolution images of the entire field of view every couple of seconds, or to track a moving target with a smaller image window that updates at 10 Hz. The system may need to control the data bandwidth that the image system has to downlink, or to store the data for later transmission. The FPGA-based image processor can be reconfigured in the system, allowing for upgrades and additional features such as pattern matching or motion detection to be added without modifying the hardware. This is a great advantage over any predefined frame grabber or ASIC-based hardware. Figure 1 shows the flexibility of a reconfigurable image processing system capable of receiving image data from a variety of sources and interfaces, implementing a selection of computationally complex image processing functions and outputting a compressed video stream via a choice of interfaces.

Upfront data compression using JPEG2000

The data received from color image sensors in any system must first be interpolated to get the RGB components for each pixel, and then typically converted to YCrCb color space to reduce the amount of data using chroma subsampling. However, by utilizing an FPGA processing system, the video

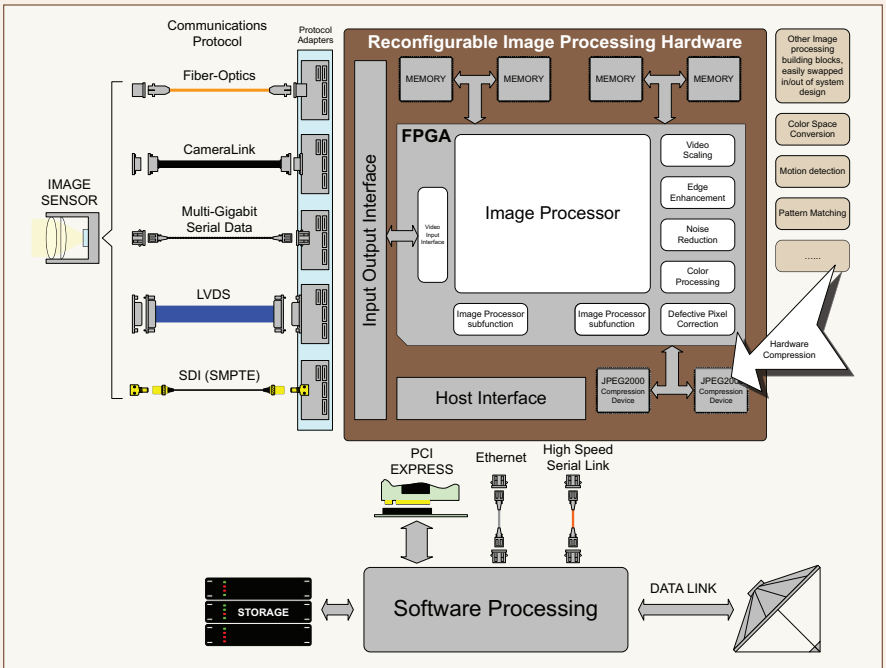

Figure 1 | Reconfigurable image processing hardware's flexibility

data can be further compressed with JPEG2000 and the bandwidth reduced by an additional factor of 2:1 in a lossless compression, 10:1 in a visually lossless compression, and up to 100:1 in a lossy compression. The compressed data provided by the FPGA image processor allows the system to easily store, further process by software, and transmit the data to users. Image data compression using JPEG2000 provides superior results for a system without upfront data compression and reduces data for the system to handle. Figure 2 shows an example of an image compressed at different JPEG2000 ratios, along with the example data sizes this provides.

The JPEG2000 codec is especially popular for military surveillance applications, as it offers very high performance, low latency, the ability to perform lossless compression, and other very useful features. This compression standard works by breaking down individual image frames into increasingly lower resolution sub-bands using wavelet transformations. The data is then arranged into packets that contain information about groups of pixels in a succession. The resulting image stream proffers useful features such as selectable areas, selectable resolution, error resilience, and bandwidth limiting, in addition to allowing a system to work with it in the image stream's compressed format. For instance, the JPEG2000 data can be used in remote surveillance systems to downlink flexible combinations of small HD image windows or wide area snapshots from the same image sensors, all without having to decompress the data first.

It can also be used at a specific bandwidth by simply extracting a portion of the JPEG2000 data and truncating the higher-resolution portion of the stream. This is very useful for transmitting the image data, as the most important information in an image can be delivered using less than 20 percent of the JPEG2000 data stream. This error resilience is furthered because each video frame is individually compressed, which allows the data link to operate with less error correction. Therefore, there are fewer concerns over data loss with JPEG2000 as compared with other compression standards, which require the information from multiple frames for decoding. Finally, JPEG2000 data can be viewed at different resolutions and quality from the same original stream.

Although a JPEG2000 encoder is quite complex to implement, an FPGA provides plenty of resources to include this feature in an image processing design. There are IP cores readily available to integrate into an FPGA design, or an FPGA can be used to interface to ADV212 ASICs (JPEG2000 System-on-Chip). Both of these options have their own benefits and trade-offs (Table 2), and either of these codec options can be used on COTS hardware such as the Alpha Data ADM-XRC-5T2-ADV mezzanine FPGA board.

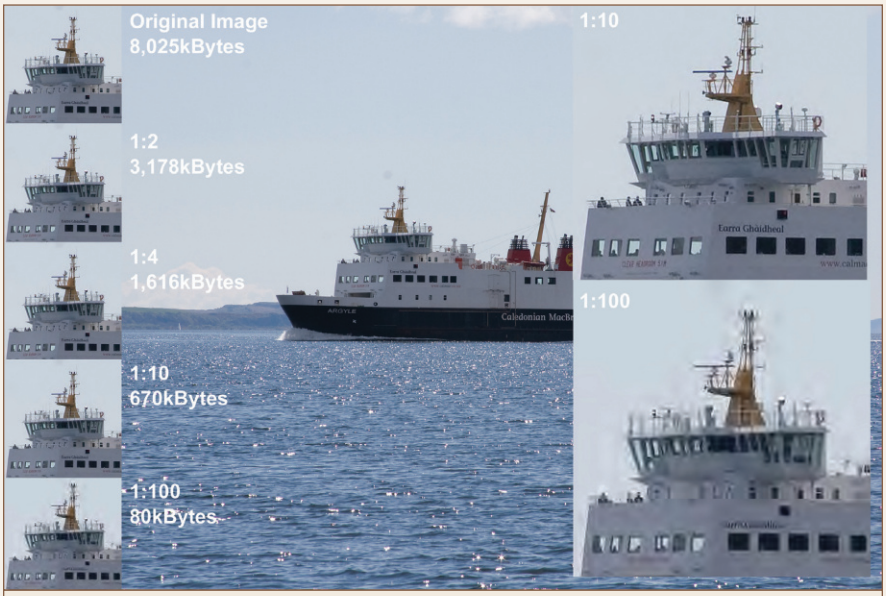

Figure 2 | JPEG2000 image compression example

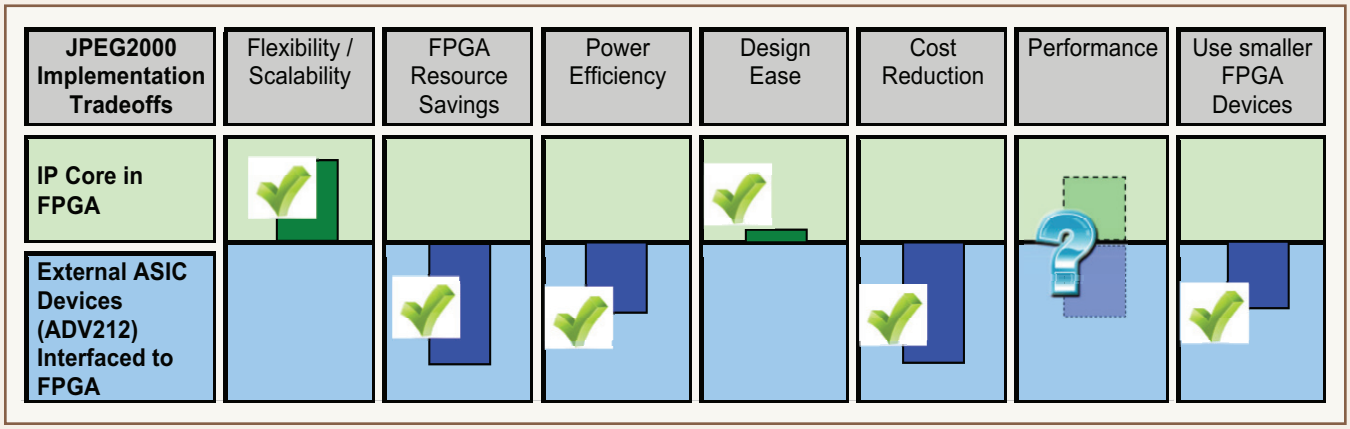

Table 2 | Design trade-offs when implementing JPEG2000 image compression using an IP core in an FPGA versus using external ASIC devices interfaced to an FPGA.

Utilizing COTS FPGA hardware for image processors

To facilitate advanced image data processing, a modular FPGA image processor board – such as those that are PMC/XMC-based – can be used in industry-standard embedded computing systems including those in the CompactPCI, VME, VXS, VPX, PCI, PCI Express, and other form factors. These small (74 mm x 144 mm) modules are available with the largest FPGA devices offered and can be used in workstation development systems, or moved directly into ruggedized deployed surveillance systems. For JPEG2000 compression, COTS boards offer large amounts of onboard memory that IP cores require and/or dedicated JPEG2000 ASICs. The host computing system can easily control the hardware image processor cards to apply user settings, reconfigure the FPGA for various video operations, control image windowing, command target tracking, and receive the output video streams. Such COTS hardware is typically provided with software drivers and API interfaces that make the control and data transfer seamless in the system.

***Adam Smith** is Vice President of Sales and Support at Alpha Data Inc. His background is in designing FPGA firmware, digital, and analog electronics. Adam has enjoyed a career with Lockheed Martin Space Systems and Alpha Data after earning a BSEE from Colorado School of Mines and an MBA from the University of Colorado. He can be contacted at adam.smith@alpha-data.com.*

Alpha Data Inc.
408-467-5076 • www.alpha-data.com

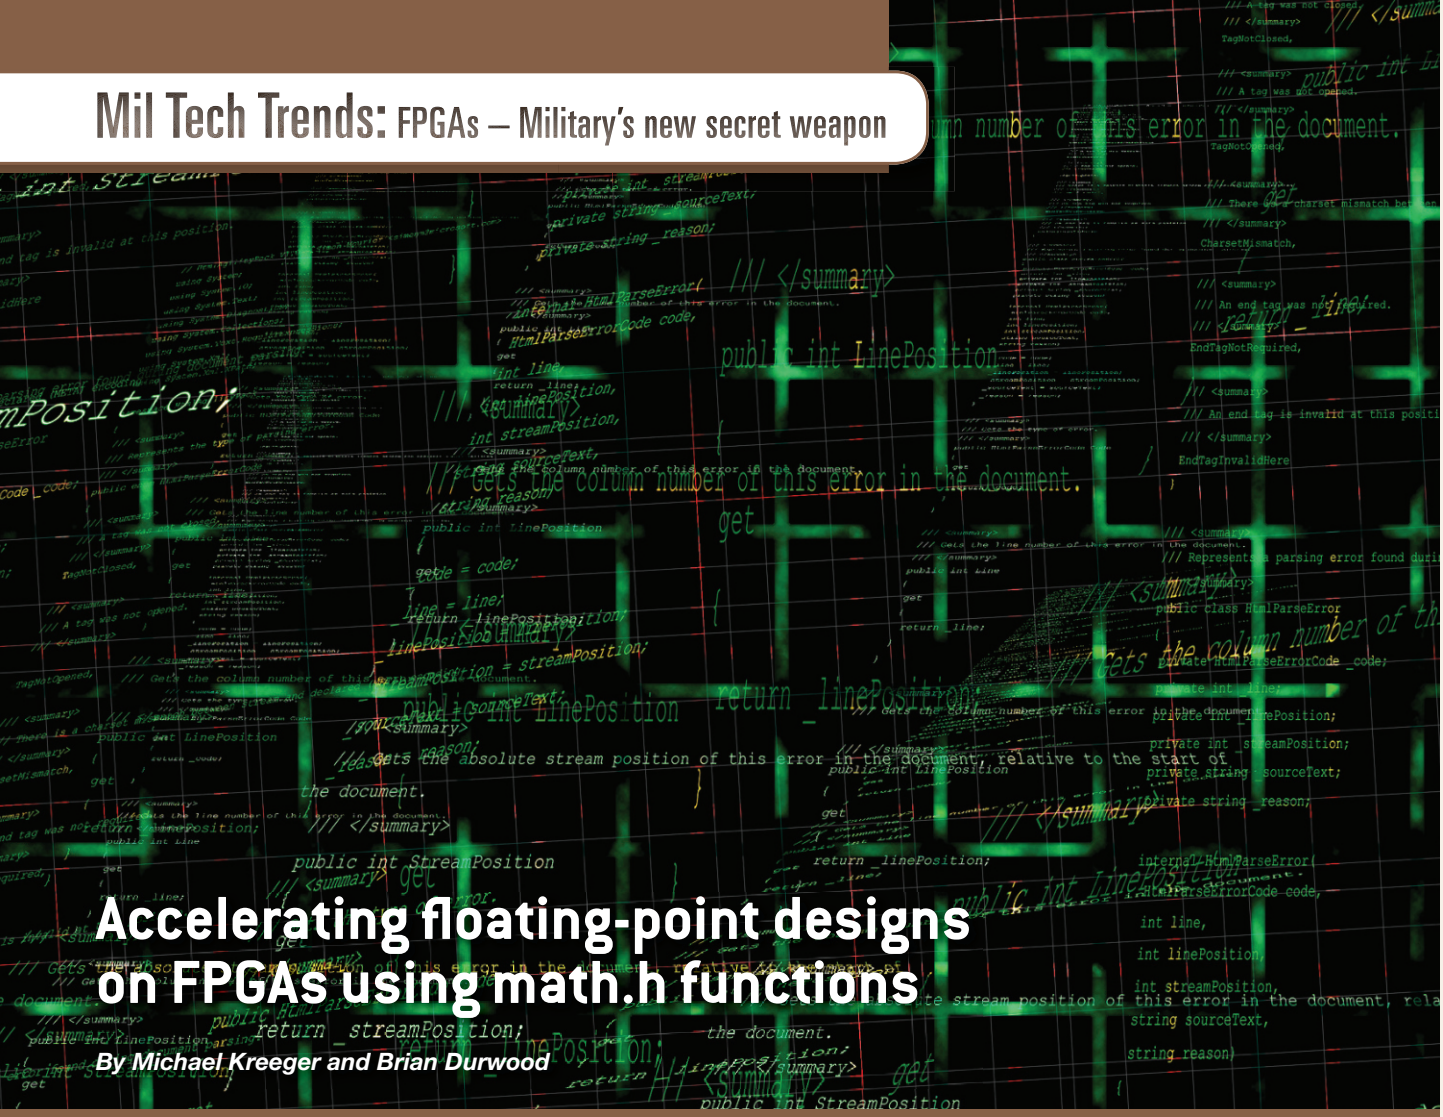

Accelerating floating-point designs on FPGAs using math.h functions

By Michael Kreeger and Brian Durwood

Floating-point math is increasingly necessary for 10/10ths accuracy in high-performance computing. Because hardware resources are not infinite for most of us, an understanding of practical approaches such as using math.h library functions for floating-point designs on FPGAs becomes integral.

Because of their flexibility advantages, FPGAs hold great promise for in-line hardware acceleration of processes that benefit from floating point. As such, FPGAs are finding their way into computational finance, scientific computing, and government and military applications.

FPGA adoption in floating-point design has been impeded by 1) issues with ease of use; 2) the need to write VHDL interfaces for C functions; and 3) inability or limitations in using known-good code. However, a new floating-point FPGA library is accelerating FPGA-based algorithm development. This library and tool flow is particularly beneficial for software developers experimenting with FPGAs for hardware acceleration of their code. The enabling elements are the emergence of robust tools for moving C to the FPGA without having to be a hardware guru, and most recently, flexible math.h libraries for common integer and floating-point computations.

Collectively the newer approach of using a math.h library to floating point on FPGAs has the advantages of:

1. Recycling known-good code to reduce user coding errors.
2. Automating pipelining and synchronization issues.
3. Shortening algorithm iteration by up to 8x over HDL development – a case study is presented.

Reusing known-good code

Higher-level math is used to calculate functions enabling machine vision to track moving objects, such as threat identification systems being developed for fighter jets. Engineers need trigonometry to figure out distances to objects, angles of view, perspectives, and relative positions. In algorithms like these, in which moving objects are tracked, vector math is heavily used. These problems have been often solved in microprocessor or DSP software code, resulting in significant power and space demands. To execute algorithms reliably in hardware, using tested hardware equivalents of floating-point functions is much safer than creating new ones.

Hence, using known-good math.h library functions reduces errors that can, by the geometric nature of the moving targets, amplify small mistakes beyond tolerance. Ironically, this is also true in long-range financial options calculations. There was an instance in European options a few years ago, where a small math imprecision created hundreds of thousands of dollars of errors when extended out 18 months, for which the bank was liable.

Until fairly recently, FPGAs were less common as devices for floating-point implementation. A leading engineer at an FPGA manufacturer pointed out that it can be difficult to predict the dynamic range of a digital signal processing system and implement

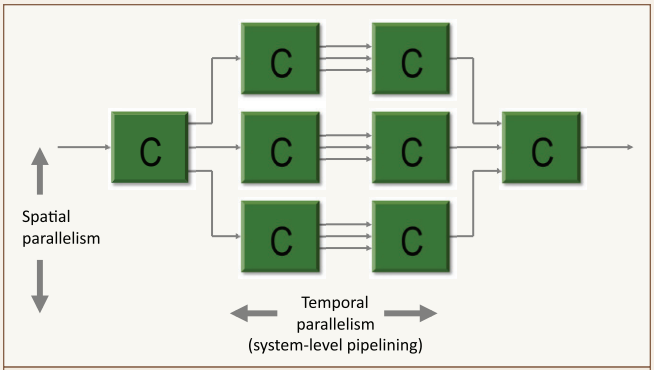

Figure 1 | Parallelism occurs in either a spatial or temporal methodology.

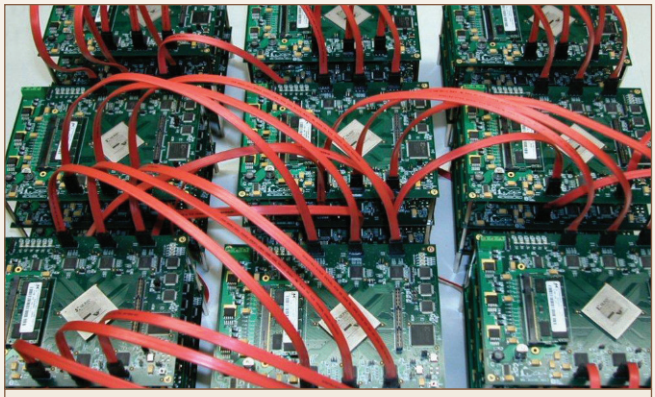

Figure 2 | Early "desktop grid computer" developed by Assistant Professor Ross Snider and his students at Montana State University, part of an experiment with parallelism as part of a shift to "coarse-grained" streaming. Deployment of math.h functions fits within this coarse-grained architecture methodology.

a fixed point system to match. Given that FPGA sizes continue to increase, while logic costs continue to decrease, it is reasonable to assume that signal processing functions will migrate away from fixed point arithmetic and move towards floating-point arithmetic. Basically, that's a nice way to explain that cheap gates will make FPGA-based floating point less of a luxury.

Parallelism eases pipelining, synchronization

Supercomputer architectures are taking an interesting shift towards slower clock speed, massively parallel arrays. Logic optimized for these architectures relies on a shift to individual processes that are synchronized via methods such as shared memories. This parallelism can be implemented in terms of spatial or temporal shifting (Figure 1). Logic challenges with significant non-sequential processes are good candidates for this type of FPGA-enabled implementation. Within this category, floating-point designs can be challenging to implement and optimize in an FPGA.

First, the typical floating-point design is optimized for a single- or dual-streaming microprocessor or DSP. FPGAs achieve performance through slower clock speeds, lower power, and massive parallelism, such as those featured in the early parallel processor array in Figure 2. (The figure depicts an early "desktop grid computer" developed by Assistant Professor Ross Snider and his students at Montana State University, who experimented with parallelism for neurological emulation in a 27-FPGA array using low-cost COTS devices[1].) Intrinsically, designs need to be refactored to run in parallel. This shift to "coarse-grained" streaming parallelism is not rocket science but is needed for the range of optimizing C compilers that work with FPGAs to unroll and parallelize the files to run in FPGAs. Deployment of math.h functions fits within this coarse-grained architecture methodology, which enables the algorithms to be machine optimized for parallel operation.

Case study: Algorithm iteration 8x faster than HDL

In a step towards making FPGA algorithms even more predictable, the preconfigured math.h library can be used to accelerate FPGA-based algorithm development and extend math operations. ANSI C-compatible tool sets can work in front of synthesis tools from Xilinx, Altera, Synopsys, Mentor Graphics, and others.

The math.h library components are provided with standard C-language function prototypes, allowing them to be invoked from C, using common function-calling methods. For instance, for functions like y = atan(x), in VHDL the developer would have to write interfaces, entities, and more versus describing the algorithm in C and having the math.h library's optimizing C compiler generate the VHDL. The VHDL is then either used just as VHDL, or is synthesizable into multiple brands of FPGA. This technique can reduce initial design time by about 1/3 and iteration time by 7/8ths.

The math.h library's composition

The math.h library is provided as a set of HDL files and related configuration files implementing common mathematical operations. The currently available functions (more are being added) include: sin, cos, tan, exp, log, log10, pow, asin, acos, atan, sqrt, and fabs. Most of these functions involve the use of floating-point numbers, either single- or double-precision.

Unlike math.h functions that run "native" in embedded processors, this library is implemented directly in FPGA hardware, and supports refactoring into multiple, pipelined parallel processes. When used in this way, the math.h functions operate 2 to 10x faster than on embedded processors. The math.h library adds more scientific, algorithmic, and engineering functions to the existing C to FPGA floating-point support. These C-callable functions represent optimized math elements that are instantiated, through the use of synthesis and place-and-route tools, in the target FPGA. Running floating-point functions directly in hardware helps accelerate the performance of embedded applications.

In the "why bother" category, let's look at the "before and after" of a common function. As implemented without invoking a library function (that is, the "before"):

```
library ieee;
use ieee.std_logic_1164.all;

library impulse;
use impulse.components.all;

use impulse.xilinx_float_math_fast.all;

use impulse.xilinx_float_fast.all;

entity dut is
  port (signal reset : in std_ulogic;
    signal sclk : in std_ulogic;
    signal clk : in std_ulogic;
    signal y_out_rdy : in std_ulogic;
    signal y_out_data : out std_ulogic_vector
     (31 downto 0));
end dut;

architecture rtl of dut is
       signal val_x : std_ulogic_vector (31 downto 0);
       signal val_y : std_ulogic_vector (31 downto 0);
begin

  val_x <= X"0x4039999a" -- represents a float value of 2.9

  logf_fast_0: logf_fast
    port map (
      clk => clk,
      a => val_x,
      go => '1',
      result => y_out_rdy,
      pipeEn => y_out_rdy);

end rtl;
```

As implemented with a known good math.h library, it's a bit shorter to say the least (aka, the "after"):

```
#include <math.h>

void main(void) {
  float x,y;
  x = 2.9;
  y = logf(x);
  printf("result=%f\n",y);

}
```

The typical design flow is vastly streamlined:

1. Combine math.h library function calls C-language to create complex systems.
2. Integrate pre-optimized library blocks from FPGA manufacturer libraries.
3. Analyze, refactor, compile, and iterate to optimize FPGA performance.
4. Verify C code functionally in a desktop environment such as Visual Studio, Eclipse, and GCC-based tools.
5. Export synthesizable VHDL or Verilog to FPGA synthesis and platform tools.

Within the math.h library toolset, support for module generation allows hardware IP blocks to be generated from the C language, using named ports and streaming API functions to integrate these blocks with the overall design. IP blocks can be mixed with Verilog or VHDL or with IP created using FPGA manufacturers' tools. For video applications, the C API functions can be used to combine multiple streaming C-language processes to create highly pipelined, high-throughput systems (see Figure 3). Math.h operators are easily called within this tool flow, providing parallelized, optimized versions of math.h elements that integrate seamlessly into Impulse CoDeveloper Version 3.

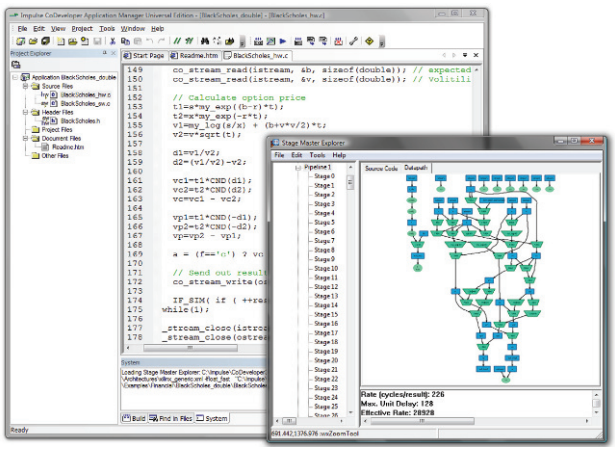

Figure 3 | A depiction of the compact C-code design description, rendered in Microsoft Visual Studio, refactored into machine synthesizable VHDL. (top screen). The pop-up flow window is a graphic interface that can be used to profile code by showing where bottlenecks result.

Figure 3 (top screen), shows the VisualStudio rendition of the original C code with some refactoring into multiple streams. The pop-up flow window is a graphic interface that can be used to profile code by showing where bottlenecks result: where there is potential for acceleration via refactoring that code block into more streams. For more detailed refinement, the same code feeds into an integrated stage-delay analysis tool (the tree-type diagram), which shows the propagation in maximum detail.

Math.h library eases FPGA development pain

For software developers exploring hardware acceleration, automating floating-point library acceleration lowers a hurdle for software-to-hardware compilation. Within the realm of Impulse Accelerated Technologies' Impulse tools, software developers are able to refactor C code for massive parallelizing to exploit available FPGA resources. It is our experience that FPGAs are moving more into high-performance computing and, as such, will become the pre- or co-processing element in a design responsible for the heavy lifting on non-standard logic constructs.

Michael Kreeger *is Principal at Kreeger Research, where he designs high-performance, communications, and math-oriented FPGA-based libraries for government, industrial, and wireless development teams. He graduated with a BSEE Cum Laude from the University of Massachusetts. He can be contacted at mike@kreegerresearch.com.*

Brian Durwood *founded Impulse Accelerated Technologies with David Pellerin in 2002 to provide C-to-FPGA based tools, training, and IP. He was a VP at Tektronix's high-frequency MCM division in the '90s, an original member of the ABEL team in the '80s, and is a Business graduate of Brown University. He also received an MBA from University of Pennsylvania's Wharton School of Finance. He can be contacted at brian.durwood@ImpulseC.com.*

Impulse Accelerated Technologies
425-605-9543
www.impulseaccelerated.com

References:

[1] *Developing a Data Driven System for Computational Neuroscience,* by Ross Snider and Yongming Zhu, Springer Berlin/Heidelberg, http://www.springerlink.com/content/ejc270741d7cmn6n/

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.*

Scalable, rack-mount MicroTCA box

One European vendor recently revealed that their robust 1,500 unit MicroTCA backlog demonstrates the market's keen interest in using Advanced Mezzanine Card (AMC) COTS boards, but in a smaller and complete rack-mount chassis. Mercury Computer's Ensemble 2000 MicroTCA platform is just such a chassis, incorporating six slots and a built-in system manager function bolted to the backplane to eliminate the need for a separate MCH module. Designed for stand-alone entry-level or multi-chassis systems, the standards-based unit is designed to maximize use of AMC, Ethernet, IPMI, RapidIO, and MicroTCA industry standards.

Backplane switches for GbE are built in, and other available communications fabrics include RapidIO, 10 GbE, and PCI Express. Not requiring an MCH module saves a slot, so all six slots can accept AMCs. The chassis can be stacked in a 19" rack, and Ethernet, communications fabrics, and system clock can be daisy-chained across multiple boxes. Mercury's own AMCs include FPGA and processor cards based upon Xilinx FPGAs, TI DSPs, Freescale PowerQUICC or MPC 8641D, and Intel Penryn processors.

Mercury Computer Systems • www.mc.com • www.mil-embedded.com/p45264

Riding the rails with this DC/DC converter

Defense integrators utilize COTS products from whichever industry meets their needs, including automotive and railway transportation. We think the Martek MBRH Series of railway DC/DC converters might have some applicability in mil apps, especially in "big iron" vetronics chassis. After all, a big old locomotive or a big old M1A2 Abrams are, well, big, hot, and loaded with "dirty" power.

The fully encapsulated converters come in single- or dual-output flavors, sourcing from 5 VDC to 48 VDC, with inputs including 24, 36, and 110 VDC. And a 230 VAC input model is soon available. Based on the company's 15 W models, the new series puts out 25 W in the same bulkhead-mounted package – a footprint we think might be handy in DoD retrofits where in-chassis space is at a premium. Input ripple is per RIA 13 and EN50155, and output transients are in conformance to RIA 12 – go and dig up the railway specs for details. The MBRH Series is manufactured in the UK – handy if a program needs some FMS set-aside or offset forgiveness.

Martek Power • www.martekpower.com • www.mil-embedded.com/p45532

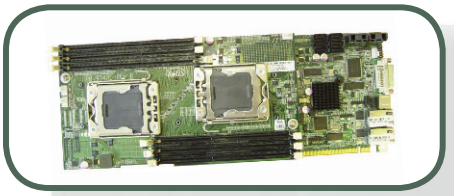

In the fast lane: PICMG style

When WIN Enterprises incarnated its MB-80100 SHB, it undoubtedly had high performance on its mind. Evidentiary support includes the SHB's PICMG 1.3 style design, facilitating twin Intel Westmere and Nehalem multicore processors. Case in point: MB-80100 supports quad- and six-core Westmere, and dual- or quad-core Nelahem – and the company even promises that its SHB will keep pace once the Westmere reaches its anticipated 8- and 12-core iterations. Not only that, the MB-80100 provides a whopping 20 lanes of 32- or 64-bit PCI Express Gen 2.0 – perfect for military, seismic, and scientific applications, just to name a few.

But there's more to MB-80100 than meets the performance-critical eye. It also provides support for an MXM-II interface or a custom MXM-II pass-through card that can bridge x16 lanes. And those pesky coding errors are no problem, thanks to MB-80100's six ECC registered DDR3 240-pin slots, which equip each processor with triple channel memory; meanwhile, those who are perhaps more daring and prefer to not go the ECC route can also opt for 240 non-ECC pin slots. In addition, individual CPU and memory power planes are also rendered, as are FireWire support via a custom card and a custom HD Audio output card featuring an HD Audio bus header. And finally, dual 1 GbE LANs are provided as standard. But MB-80100 just couldn't function without all the support it gets from its OSs: Windows Vista 32/64 bit, Windows Server 2008, Windows 7 (32/64 bit), and Windows XP 32/64 bit.

WIN Enterprises, Inc. • www.win-ent.com • www.mil-embedded.com/p45533

GUI for every device

Graphical user interfaces have become expected on most devices these days, and we've all had the displeasure of trying to interpret complicated commands using a few buttons and a measly three LEDs. A GUI proves a rich experience, as long as the CPU/chipset can drive an LCD (or equivalent) display. Blue Water Embedded's Prism is a suite of GUI development tools optimized for low-cost, resource-constrained embedded devices. The suite consists of the Prism Runtime Framework as a GUI toolkit; Prism Micro for monochrome and 8-bit color depth targets; and Prism Insight, which is the GUI design and resource editing tool.

The Framework component automates design and deployment with a widget set, drawing engine, event manager, and screen manager. Collectively, the tools create rich animations, screen transitions, anti-aliasing (a *must* on low-res screens), as well as blending and canvas transformations. Prism Micro is a version of the Runtime Framework that's designed specifically for cost-constrained and low-color-depth targets. Prism Insight is a drag-and-drop WYSIWYG with TrueType fonts and widgets to create buttons and other screen elements. Output files can be C/C++, XML, or binary.

Blue Water Embedded • www.bwembedded.com • www.mil-embedded.com/p45534

iPhone-sized SBC defines portable power

Weighing a mere 1/3 pound and sized similarly to an iPhone, the Atom XPC40x by General Micro Systems is the company's latest rugged mil/aero all-in-one SBC. Sporting a 1.6 GHz Intel Atom with 512 KB of cache, the 3.5" x 2.5" x 0.5" module comes in air- ("XP") and conduction-cooled ("XPC") flavors ready to operate over -40 °C to +85 °C. The SBC portion is pretty standard stuff: expandable 533 MHz DDR-2 SDRAM with 3D accelerated graphics, but there are five USB ports.

Not only that, the packaged version (shown) has support for two Express Mini Cards, so a user (or GMS) can add Wi-Fi, GPS, CANBus, or myriad other I/O functions. The packaged version also offers additional security for storage, since the *entire computer* contains the storage and can be moved from location to location. Finally, there's the portability factor. The Atom CPU runs on 3 W (typ) and 10 W (peak). The whole XPC40x, depending upon I/O and peripherals, is ideal for battery-operated applications such as UAVs, remote sensing, or handheld/manpack equipment.

General Micro Systems • www.gms4sbc.com • www.mil-embedded.com/p45386

Mini, wide DC-DC converters

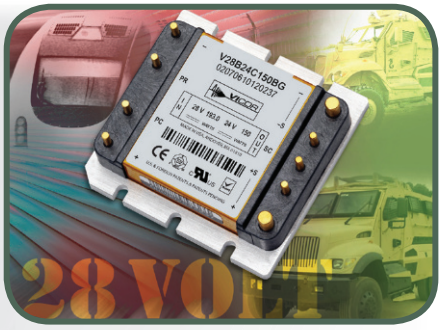

It sounds like a line from a country western song, but the "28V Wide Input Mini Family of DC-DC Converters" high-efficiency DC-DC converters from Vicor are both "mini" and "wide." How's that? Their small "mini" size fits between the vendor's micro and maxi sizes at 2.28" x 2.2" x 0.5" (59.9 x 55.9 x 12.7 mm), making them ideal in size-constrained, high-rel applications. Four temperature grades are available, the most robust capable of -55 °C operation at 90 percent energy efficiency.

More important is the "wide" part. Designed for vehicles with 12 V or 24 V battery power, the converters can accept voltage inputs from 9 to 36 V (for the 12 VDC unit) or 18 to 36 V (for the 24 VDC unit). Vehicle charging systems, high battery loads, and discharged cells can easily swing available voltages within the limits of these converters. All the while, each is capable of outputting a stable voltage in eight model increments from 3.3 V to 48 V. Finally, mounting options include slotted, threaded, or through-hole, and the device's encapsulation is assured void-free to maximize energy and heat efficiency, along with reliability.

Vicor • www.vicr.com • www.mil-embedded.com/p45535

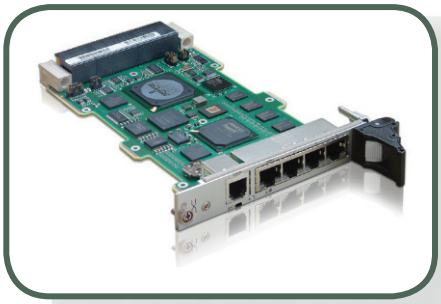

3U OpenVPX GbE switch

The verdict's in, folks: Ethernet is the winner. So it's no wonder that more defense systems than ever use Ethernet as the high-speed conduit on- and off-board. Kontron's Gigabit Ethernet Switch VX3910 performs enterprise-class switching functionality on 28 ports. It also has advanced management features, like you'd expect to find in carrier-grade switches. Even better: It's available in VITA VPX (VITA 46.x) and OpenVPX (ANSI/VITA 65) versions, in 3U size and in air- and conduction-cooled formats.

There are 20 Gigabit ports available to the backplane, four 2.5 Gigabit ports to the backplane for redundancy in multi-switch systems, and an additional four 1000BASE-T uplinks on the front panel in air-cooled versions (shown). Conduction-cooled boards operate at -40 °C to +85 °C. Kontron's Embedded Network Technology design keeps the same feature set and interfaces in these modules as in other Kontron boards, simplifying programming and interoperability. Ethernet ports are provided by Broadcom silicon, and the system controller is an AMCC PowerPC switch device with customizable software and comprehensive firmware. L2 (and optional L3) switching is supported, as is VLAN tagging (IEEE 802.3ac), dynamic VLAN registration (IEEE 802.1Q), and QoS (IEEE 802.1p).

Kontron • www.kontron.com • www.mil-embedded.com/p45497

OFDM FPGA core

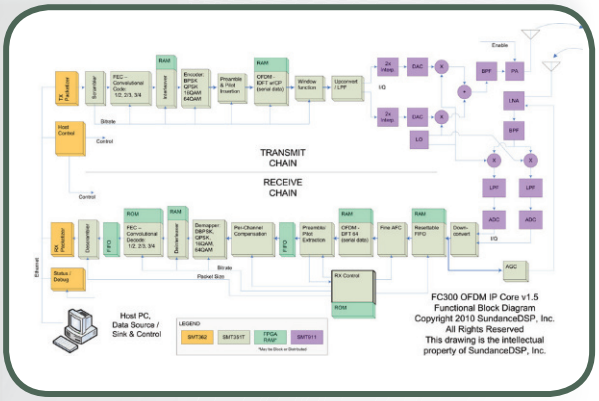

Orthogonal Frequency Division Multiplexing (OFDM) is the core of most broadband wireless waveforms, including Wi-Fi, cellular, and other spread-spectrum implementations. Available as an FPGA IP core targeting a Xilinx XC5VSX95T, the FC300 core took Sundance DSP considerable time to develop, test, and implement. Intended for use on the company's SMT351T FPGA module and bolted to an SMT911 module, the core itself doesn't use any off-FPGA resources and can be retargeted to any FPGA of sufficient size, including Spartan to Virtex-7, or for Altera's Cyclone or Spartan devices. The core only assumes an external I/Q baseband DAC/ADC interface, making retargeting hardware also straightforward.

The core contains Tx and Rx chains, and for Wi-Fi, includes 802.11a/g data scrambling, forward error correction, subchannel coding (BPSK, QPSK, 16QAM, or 64QAM), pilot insertion, and IFFT/FFT algorithms. On the Rx side, there are estimators for center frequency offset and fractional sample signal timing. The core was developed for a Chinese Sundance customer and is now available to the general market.

Sundance DSP • www.sundancedsp.com • www.mil-embedded.com/p45483

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.

Thinking outside the box:

Why did Curtiss-Wright and Kontron both just buy systems integration houses?

By Chris A. Ciufo, Editor

On May 21, 2010 Kontron announced that they were buying AP Labs for an undisclosed sum. Six days later, Curtiss-Wright said they were buying Hybricon. Coincidence? I doubt it. I first conceived the "COTS funnel" concept, which seems to be at play here, in 1997 at Dy4 Systems: As more military programs use open-standard increasingly commoditized COTS embedded hardware and software and thereby fall down the funnel, vendors are striving to add more value into the funnel through systems expertise at (and beyond) the box level.

Advanced Processing (AP) Labs has a strong COTS history as a $30 million systems house, *par extraordinaire.* Famous for their advertisements showing a monstrous cargo plane propeller motor "umbilicaled" to an AP Labs test rig, the 26-year-old company has, in the past seven years, moved into conduction- and liquid-cooled COTS chassis and systems. In the past three years, they've begun servicing the commercial aerospace markets with program wins in avionics and air-to-ground cellular and wireless communications. Actually composed of three business units – AP Labs, AP Avionix, and AP Parpro[1] – Kontron is buying program backlog, systems integration and chassis design expertise, and a Mexican manufacturing facility.

Located just 5 miles from Kontron's U.S. headquarters near San Diego, California, a spokesperson told me that "100 percent of AP Labs' employees" will move into Kontron's facility. Having toured Kontron on several occasions, they have plenty of engineering and prototyping space in the area formerly occupied by what was then the Dolch rugged computer group (sold off to Crane in 2007). AP Labs will complement Kontron's 2008 "lucky buy" of Thales Computers, which added PowerPC VME and 2eSST rugged boards sold primarily into U.S. and French military programs.

Kontron needed Thales to be a legitimate player in DoD programs and to jumpstart the company's anemic VME SBC business; you're not credible in hardcore Aerospace and Defense (A&D) with just Intel-based COM and PC/104 boards. But Thales had let their SBCs languish too long with minimal investment from mother France, so Kontron also needed table stakes in systems integration to compete with GE Intelligent Platforms, Mercury Computer Systems, and Curtiss-Wright Controls Embedded Computing (CWCEC).

AP Labs is a "strategic acquisition" for Kontron, says Western Regional Sales Manager Michael Humphrey, giving Kontron a "force multiplier" that allows DoD primes to focus instead on their core competencies: system application software and platforms. Kontron's military and aerospace group ("MAG," following Intel's nomenclature) is a double-digit and growing portion of the company's worldwide revenue as Corporate in Germany focuses on "high-margin complete solutions." Perhaps gone are the days of Kontron offering every conceivable module form-factor under the sun; I always wondered how the company could support so many sizes and processor flavors.

AP Labs' services will be sold both by Kontron and AP Labs sales forces with continued "board-level agnostic" technology solutions. But I'm sure Kontron's boards will be just a bit more equal than other vendors' – especially with the company's Premier Level partnership with Intel's CPU road map: a truly powerful combination that CWCEC can't match yet.

For $19 million, Curtiss-Wright Controls picked up packaging expert Hybricon, located in Massachussetts not far from Controls' own "electronic manufacturing services" facility in Littleton (the former Lau Defense Systems in which I played a bit role in helping CW to acquire). At $17 million in revenue, Hybricon is primarily a military COTS packaging house that CW was using for VPX and OpenVPX chassis and storage architectures for the Dayton facility (formerly Systran).

Though CW also needs systems integration capabilities just like Kontron, the company's Electronic Systems group (CWCEL, which is different from CWCEC) in Valencia, California is a "full service subsystems supplier," says VP and GM David Dietz, and I know from firsthand experience that they're already highly credible with design wins on Global Hawk, BAMS, and Bradley A3 (via Littleton)[2]. Instead, they needed Hybricon's gigabit signal integrity and thermal design intellectual property – both essential with VPX-based systems.

Like Kontron and AP Labs, CWCEL (plus Hybricon) and CWCEC each have their own sales forces, possibly resulting in some customer confusion for now. And although AP Labs adds systems integration to Kontron's board-focused businesses and Hybricon adds mechanical and electrical design to CW's primarily board and systems businesses – each company emerges about at-par against the other. Kontron has a broader range of modules and that Intel connection; CW has more military stripes than Kontron could ever hope for. It'll be interesting to watch these two titans compete … out of the box.

Chris A. Ciufo, Editor
cciufo@opensystemsmedia.com

[1] Fun fact: Back when AP Labs announced their Parpro organization, the industry watched in horror as the company's financials tanked with a market downturn. Many of us believed that it was a bridge too far, excessively straining the company's capabilities at a time when cheap manufacturing was instead moving to southeast Asia, not Mexico. Despite NAFTA, ITAR rules still don't allow full military production in Mexico, so I wonder if Parpro will eventually go on the block.

[2] Curtiss-Wright's acquisitions are hard to follow, but I worked at VISTA Controls, which was acquired by Lau Defense. Lau acquired Dy4 Systems, where I also once worked. When CW bought Lau, they ended up with VISTA in Valencia, Lau in Littleton, and Dy4 in Ottawa, Canada.

GE
Intelligent Platforms

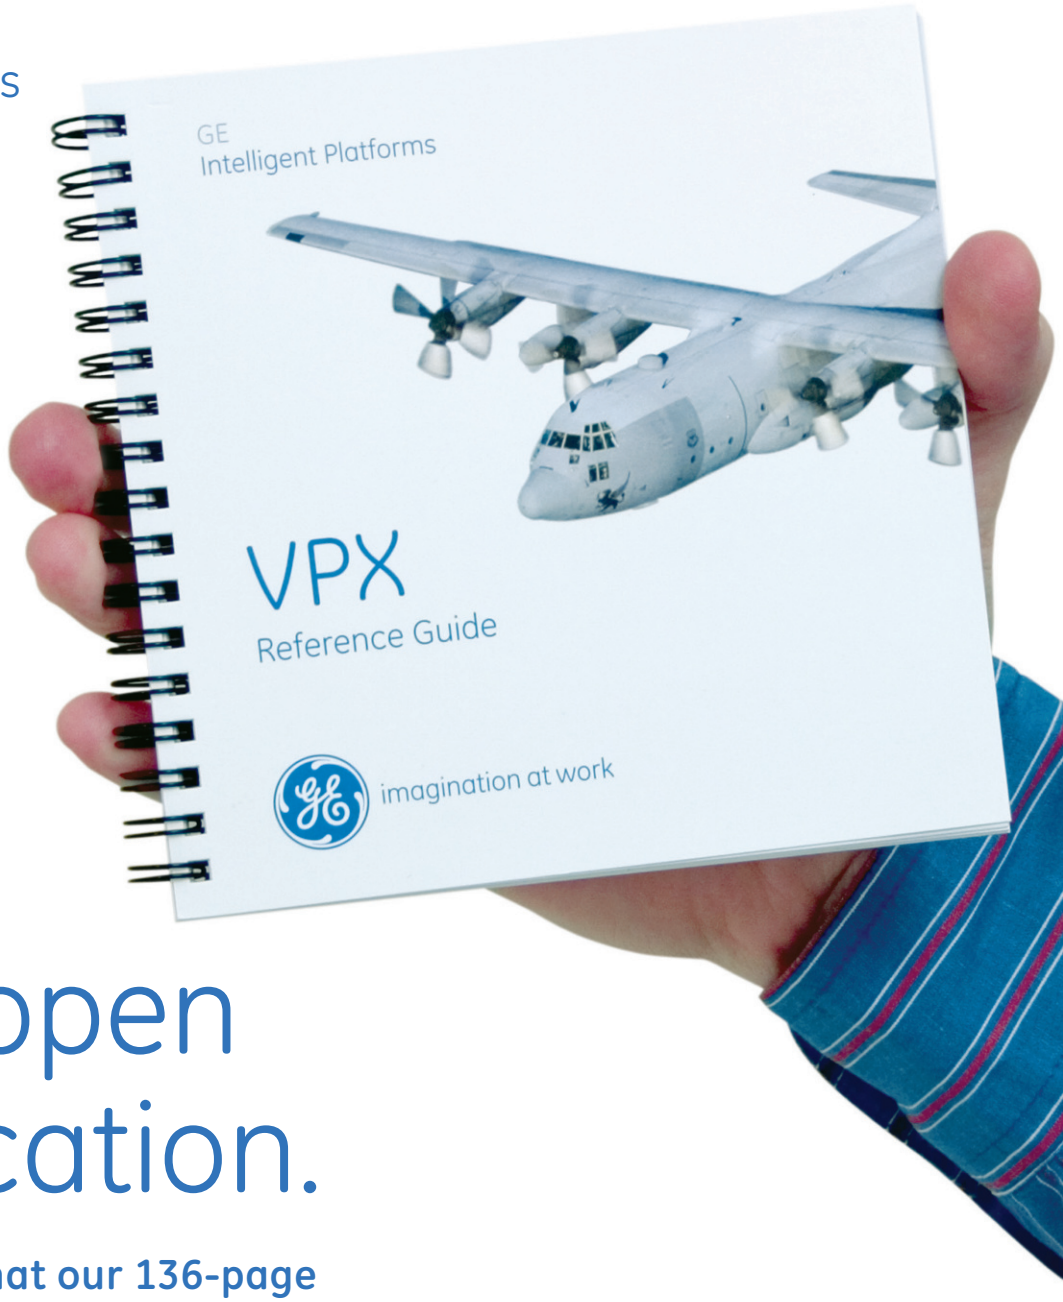

It's an open specification.

So it's only fitting that our 136-page reference guide is free.

Think of this reference book as Cliffs Notes for VPX and OpenVPX. It's that comprehensive and easy to read. It covers all the key topics, including the VITA 46 (VPX) specification, VITA 48 (VPX REDI), and VITA 65 (OpenVPX). It covers the history and rationale for VPX, its underlying principles and features, describes in detail important technical information, and illustrates the content with over 50 diagrams, charts and photos. Ordering one is a no-brainer.

Order your free copy at www.ge-ip.com/vpx

imagination at work

OpenVPX™

REDUCE YOUR DEVELOPMENT TIME & COSTS

YOUR TOTAL SOLUTIONS PARTNER

Photo courtesy of General Dynamics Land Systems Inc.

1 Image Processing
FPGA Processors
VPX3-450

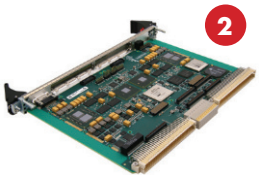

2 Radar Processing
Digital Signal Processors
CHAMP-AV5

3 Flight Control
Multi-Platform Mission Computers
MPMC-9350 & MPMC-9310 (VPX, VME or CompactPCI)

Do you need innovative solutions or complete integration services for high-density data processing now? From board-level products to fully integrated systems, we provide advanced technology solutions that perform under the most rugged operating conditions. Speed your time-to-market and lower your overall program development costs utilizing our leading edge COTS products, or modified COTS (MCOTS) and engineering services. **Ask your representative about the new VPX system specification, OpenVPX.**

4 Graphics Display
Video Input & Output
XMC-710

5 System Connectivity
SRIO/Ethernet Switches
VPX6-6902

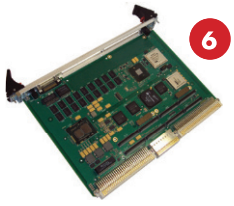

6 Command & Control
Single Board Computers
VPX6-187 & VME-1905

CURTISS WRIGHT Controls
Embedded Computing

cwcembedded.com

BOARDS & SYSTEMS ABOVE & BEYOND