Daily Briefing: News Snippets

Military EMBEDDED SYSTEMS

VOLUME 6 NUMBER 8
NOV/DEC 2010

INCLUDING:

Chris A. Ciufo
A quiet revolution against RTL

Field Intelligence
COTS = SWaP + high performance

Mil Tech Insider
Beware of counterfeits

Legacy Software Migration
Micro Focus North America:
Modernization/migration are Administration's focus

MIL-EMBEDDED.COM

COTS accessories for A&D systems

Open source *is* sufficient for critical systems

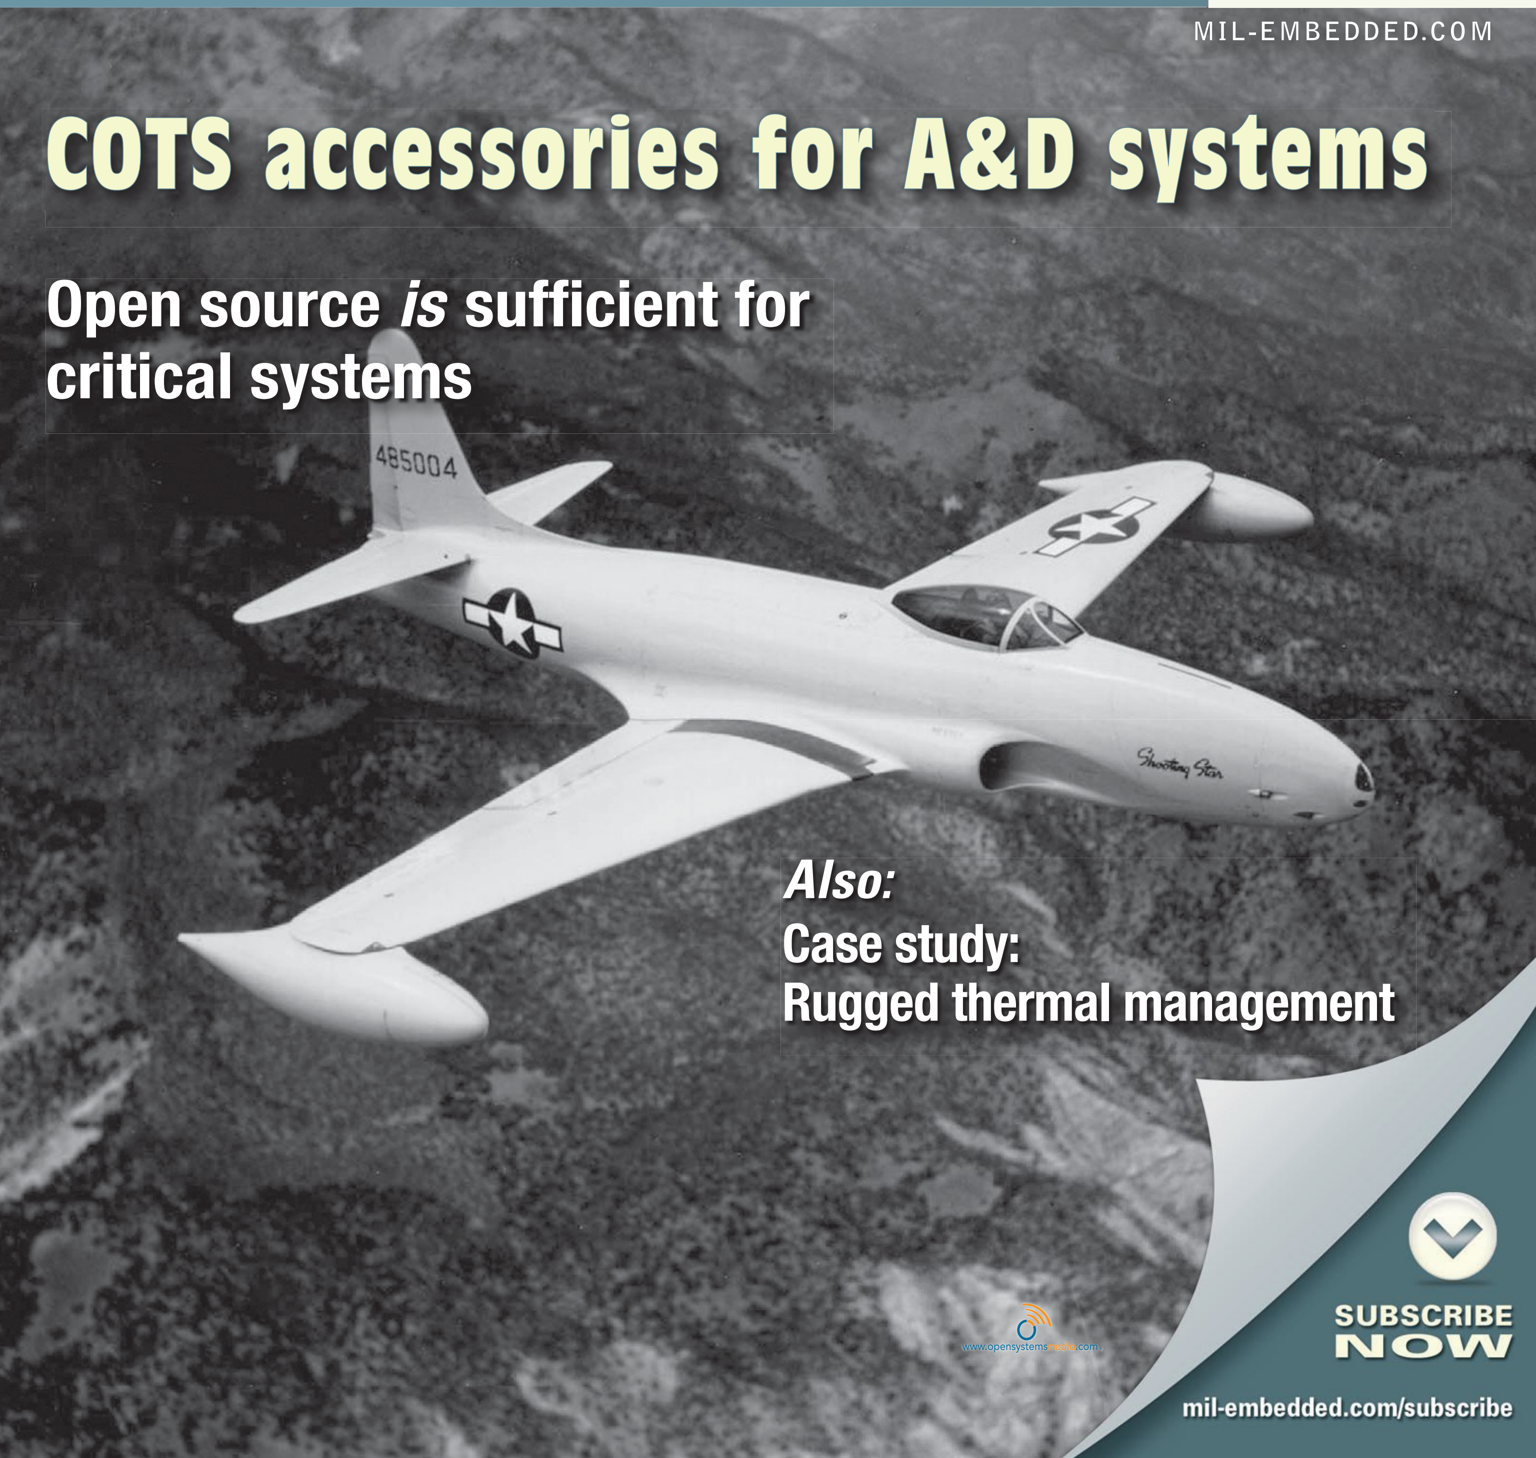

Also:

Case study:
Rugged thermal management

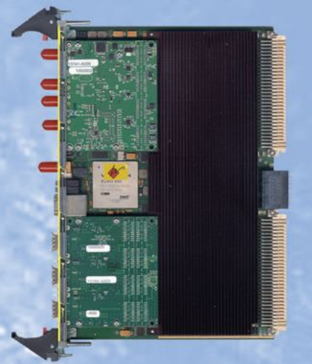

Stroke of Genius

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

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

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

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

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

Now Available in VPX!

COTS & Rugged Versions

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

Military EMBEDDED SYSTEMS

November/December 2010 Volume 6 Number 8

COLUMNS

DEPARTMENTS

EVENT

www.opensystemsmedia.com/events

CMSE – Components for Military and Space Electronics Conference and Exhibition
February 7-10, 2011 • Los Angeles, CA
www.cti-us.com

E-CASTS

http://ecast.opensystemsmedia.com

Designing Standards-Based Open Platforms in a Multi-Core Environment Using COTS Hardware Solutions
January 12, 2011 • 12pm MST
Presented by: Curtiss-Wright, Wind River, and Intel

Developing Security-Critical Applications for Multi-Level Secure (MLS) Platforms
January 18, 2011 • 12pm MST
Presented by: AdaCore and 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

ON THE COVER:
The Lockheed P-80 Shooting Star first flew in 1945 near the end of World War II. Though the Germans perfected (for the time) jet engines in the V-1 and V-2 "buzz" bombs, the ME-109 jet saw actual service during the closing months of the war, terrifying the lower and slower Allied propeller fighters. Today, of course, all fighters are jets and props are primarily used for C4ISR, scouting and forward air control, and cargo/lift. Yet all avionics platforms are mission- and safety-critical because they either carry humans or fly in controlled airspace. Safety-critical hardware and software is a requirement in modern avionics systems. Read how one company – AdaCore – believes that open source software is ready for service in avionics. Article on page 16. (U.S. Air Force photo)

Published by: OpenSystems media.

ISSN: Print 1557-3222

ADVERTISER INFORMATION

WWW.MIL-EMBEDDED.COM

OpenSystems media.

Military EMBEDDED SYSTEMS

@military_cots

DSP-FPGA.com · VME and Critical Systems · PC/104 and Small Form Factors · Industrial Embedded Systems · CompactPCI, AdvancedTCA & MicroTCA Systems · Embedded Computing Design

Military & Aerospace Group

Chris Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

Sharon Hess, Assistant Managing Editor
sharon_hess@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

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

Monique DeVoe, Web Content Editor
mdevoe@opensystemsmedia.com

Christine Capuano, Web/Editorial Assistant

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Steph Sweet, Creative Director

Joann Toth, Senior Designer

David Diomede, Art Director

Matt Jones, Web Developer

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Sales Group

Dennis Doyle, Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Bridget Hildebrand
Strategic Account Manager
bhildebrand@opensystemsmedia.com

Christine Long
Digital Content Manager
clong@opensystemsmedia.com

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

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

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

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

Ad Coordinator
Steph Sweet
ssweet@opensystemsmedia.com

Reprints and PDFs

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

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

10G SERIES RECORD & PLAYBACK SYSTEMS
ULTRA-FAST, SCALABLE & 10 GIGABIT NETWORK POWERED

Synchronized (Phase-Coherent) Recording & Playback of Multiple Wideband Radio Signals

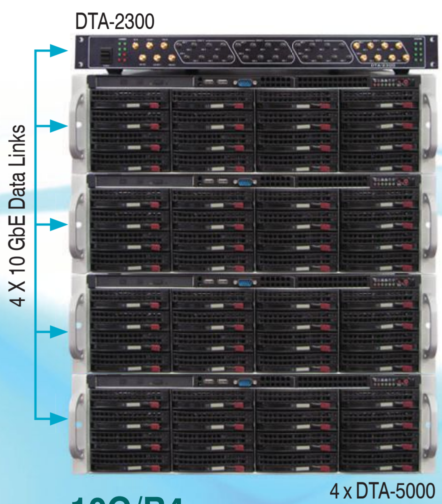

10G/R4
- 8 or 16 IF Inputs & Outputs
- 38.4 TBytes Storage Capacity
- 2.4 GBytes/s Sustained Record/Playback Rate
- Optional DTA-3200 (20MHz to 6GHz) tunable RF Front End

10G/R2
- 4 or 8 IF Inputs & Outputs
- 19.2 TBytes Storage Capacity
- 1.2 GBytes/s Sustained Record/Playback Rate
- Optional DTA-3200 (20MHz to 6GHz) tunable RF Front End

10G/R1
- 2 or 4 IF Inputs & Outputs
- 9.6 TBytes Storage Capacity
- 0.6 GBytes/s Sustained Record/Playback Rate
- Optional DTA-3200 (20MHz to 6GHz) tunable RF Front End

(User PC)

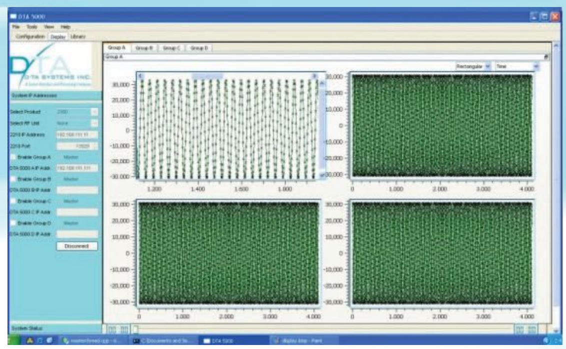

Comprehensive GUI for Failsafe Operation & Control of All System Components. Time, FFT and Waterfall Displays for Monitoring Data Quality During Record or Playback.

With 16-bit 130MHz ADCs & 500MHz DACs and programmable DDCs and DUCs, the 10G Series Systems offer high-precision and large IF Bandwidths. They are fully-integrated for immediate deployment.

1-877-382-3222
SALES@D-TA.COM

Download 10G Series Technical Presentation
WWW.D-TA.COM

NOTHING ELSE COMES CLOSE

D-TA SYSTEMS INC.
A Sensor Interface and Processing Company

SENSOR PROCESSORS THAT DRASTICALLY REDUCE DEPLOYMENT TIME AND COSTS

Field Intelligence

COTS delivers high performance and low power

By Duncan Young

Minimizing Size, Weight, and Power (SWaP) is one of the most critical requirements for deployable military equipment. Power dissipation is the key factor that drives both size and weight, as heat requires either space or mass for it to be dissipated. A typical 3U or 6U off-the-shelf, embeddable SBC dissipates between 35 and 75 W that, when combined as a subsystem with its I/O and power supplies, might easily dissipate more than 100 W. The introduction of low-power processors such as the Intel Atom can potentially reduce this power dissipation by a factor of three, while still offering performance equivalent to a Celeron or Pentium M.

Intel's Atom

Introduced in 2008, Intel's low-power Atom, with its x86 code compatibility and integrated graphics controller, fueled a generation of netbooks and portable devices. To save power, Atom's processing core utilizes a simplified architecture and extensive power management of unused resources. Complex multi-operand x86 instructions are converted to multiple micro operations for a "RISC-like" execution unit. The resultant simplification does not support preemptive or out-of-order processing, but does support multithreading, having two logical processors.

Codenamed *Menlow,* the original Atom platform can be used with a variety of existing I/O controllers connected via its Front Side Bus (FSB) to achieve the specific functionality needed. This flexibility allows the user to trade off overall power for cost or functionality. The new *Queens Bay* platform, recently announced at the Intel Developer Forum (IDF), extends the Atom processor even further as a System-on-Chip (SoC) using a 45 nm process, to include higher-performance graphics, a DDR2 memory controller, flash, and four PCI Express x1 interfaces. This device, the E6xx, dispenses with the FSB, instead using one of the PCI Express interfaces for direct connection of peripheral devices through an I/O controller hub providing CANbus, GbE, USB, and SATA. The Queens Bay offers greater processor performance and functionality than Menlow plus lower power, typically from 2.7 to 3.9 W when clocked at up to 1.6 GHz, to reduce the overall dissipation of an embeddable 3U Atom-based SBC to less than 10 W.

3U form factor

While many open architecture small form factors such as EBX, Mini-ITX, PC/104-*Plus*, or PC/104-Express are readily available for commercial applications, none rival the established position of 3U for rugged applications. CompactPCI and VPX are ideal choices with COTS vendors offering both. The ACR301 from GE Intelligent Platforms (Figure 1) uses the latest E6xx processor platform to provide a very low-power, entry-level SBC in the 3U CompactPCI form factor. The ACR301's low power and performance potential extend the range of computing technologies for critical 3U-based applications. It also offers many existing users a power saving, plug-in replacement.

Figure 1 | The ACR301 Atom-based SBC from GE Intelligent Platforms

Critical applications for low-power computing

Meanwhile, a very large class of man-wearable and portable military applications exists, ranging from personal sensors and radios to netbooks and smartphones, set to benefit from the advances in low-power processors. But these are unlikely to adopt an open architecture form factor such as 3U. These will be wholly custom designed. However, many other application spaces are available. Mission avionics applications on platforms such as helicopters or combat aircraft where weight is always a critical parameter are possibly the most SWaP-sensitive military applications utilizing off-the-shelf, embedded computing modules. Similarly, for an Unmanned Aerial Vehicle (UAV) the avionics, mission system, and sensor payload weights are critical to achieve the desired performance and endurance. Robotic Unmanned Ground Vehicles (UGVs) and scout/surveillance vehicles are complementary low-power applications on the ground. UGVs are typically battery powered and operated wirelessly or remotely via an umbilical. Scout/surveillance vehicles, which can be manned or unmanned, are used for covert surveillance operations requiring long endurance, a low infrared signature, and a passive multisensor suite such as radar, signals, organic material, audio, or LF sensing. These are all required to operate from limited battery power for extended periods.

Alternate I/O hubs

The Atom is a prime example of military embedded computing leveraging expensively developed technology from mass commercial markets. The Queens Bay platform offers CANbus, intended for future automotive "infotainment" systems but equally applicable to a military vetronics system's driver or commander's interface. The new E6xx Atom is the basis for a family of processing platforms teamed with application-specific I/O hubs. For example, Intel's new MP20 platform controller hub includes network security, encryption, and extensive audio processing capability intended for 3G handhelds and smartphones. Such hub combinations offer many new possibilities for further integration of embedded vetronic and avionic capabilities in the future network-enabled battlefield.

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

Strategies for mitigating the threat of counterfeit devices

By Steve Edwards

A strategy for dealing with the increasing problem of counterfeit devices is key to any successful life-cycle management program; it is also a key part of a security-focused COTS design philosophy. In the military COTS market, the issue of counterfeit devices is significant because of the dependence on commercial semiconductor parts, which are typically end-of-lifed at a rate driven by Moore's Law. Commercial component suppliers can expect to see certain volumes of sales over time on their devices. Once consumption of these devices starts to ramp down, suppliers may determine that sales are dwindling to an amount too low to justify production. When this occurs, a COTS board vendor – knowing their customer's long-term requirements – might be able to provide a business case to the supplier that ensures a customer base. Hence, communication is vital in combating obsolescence.

In the case that a designed-in device becomes obsolete, the COTS vendor has the opportunity to plan a technology insertion. The COTS vendor can respond with a lifetime buy to minimize the impact of the imminent nonavailability of the component. Another option is to work with the customer to determine if a specific functionality is being used. If it is not, then it might be possible to eliminate the device from the card design.

Turning to the aftermarket

In the worst-case scenario when a device is no longer being manufactured, COTS vendors can turn to the aftermarket, searching globally for a specific part. Unfortunately, the process of searching aftermarket suppliers also notifies "would be" counterfeiters as to which devices are currently in demand, and thus provides the greatest profit potential for faking components. Counterfeit components most commonly enter the supply chain through independent distributors, also known as *brokers.* They can source materials from authorized channels, acquire surplus inventory from equipment manufacturers, or buy from other brokers. The further from the original source the supply chain extends, the greater the risk of introducing counterfeit components.

There are several ways that criminals exploit aftermarket demands. Slower-speed devices can be relabeled to appear as faster versions. Parts can be desoldered from PCBs and claimed to be new, or rejected parts that have been scrapped for noncompliance to the manufacturer's specifications might be fraudulently salvaged and offered as good parts through the broker market. There have been cases where a part's packaging looks great from the outside but proves to be empty inside.

Common counterfeit methods

Counterfeit parts can be produced in a variety of ways:

- **Relabeling –** Original labeling can be physically removed and devices relabeled. These parts may or may not have similar functionality to the authentic ones.
- **Reclaiming –** Authentic parts might be salvaged from discarded circuit card assemblies. Verification of functionality within specifications will not exist.
- **Unauthorized manufacture –** Counterfeit components are manufactured and marked to duplicate authentic parts.
- **Production escapes –** Test fallout from authorized manufacturing facilities is packaged/sold as authentic parts.

The responsibility to mitigate the risks posed by counterfeit materials lies throughout the entire supply chain. The board manufacturer is in a unique position to act as a gatekeeper for its customers further along the chain. Being the last to handle individual components prior to their becoming part of a board level product, a COTS hardware manufacturer must have established processes to reduce the risk of counterfeit parts integration.

The importance of supply-chain management

A COTS board design company that operates its own manufacturing facilities has an obvious advantage regarding visibility and control over the materials used in the fabrication of its products. Avoidance of counterfeit parts is best accomplished by sourcing materials directly from the component manufacturer, from authorized aftermarket support, or from franchised distributors. Components from these sources will have supporting documentation that provides an auditable pedigree.

COTS vendors should implement enhanced controls to mitigate the risk of counterfeit devices. For example, the Curtiss-Wright Controls Embedded Computing (CWCEC) "Trusted COTS" initiative manages its supply chain by maintaining an Approved Vendors List. To qualify as an approved vendor, all suppliers are subject to audits and must be able to comply with quality clauses assigned to every purchase order.

When sourcing materials from independent distributors is considered necessary, an authorization process should be initiated. Procurement, engineering, and quality representatives should be asked to confirm first that all other viable options have been explored. Risk mitigation processes for the needed device should include comparison to known-acceptable materials and verification with the original manufacturer. Testing and analysis should include third-party destructive physical analysis, third-party electrical/functional testing, and in-house electrical/functional testing within the application. At issue is the safety and success of the warfighter. Without appropriate measures in place, products with counterfeit components can be deployed in critical military applications. These can fail at critical times, causing catastrophic results. We owe it to those who put their lives on the line for us to have appropriate measures in place to ensure that this does not happen.

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

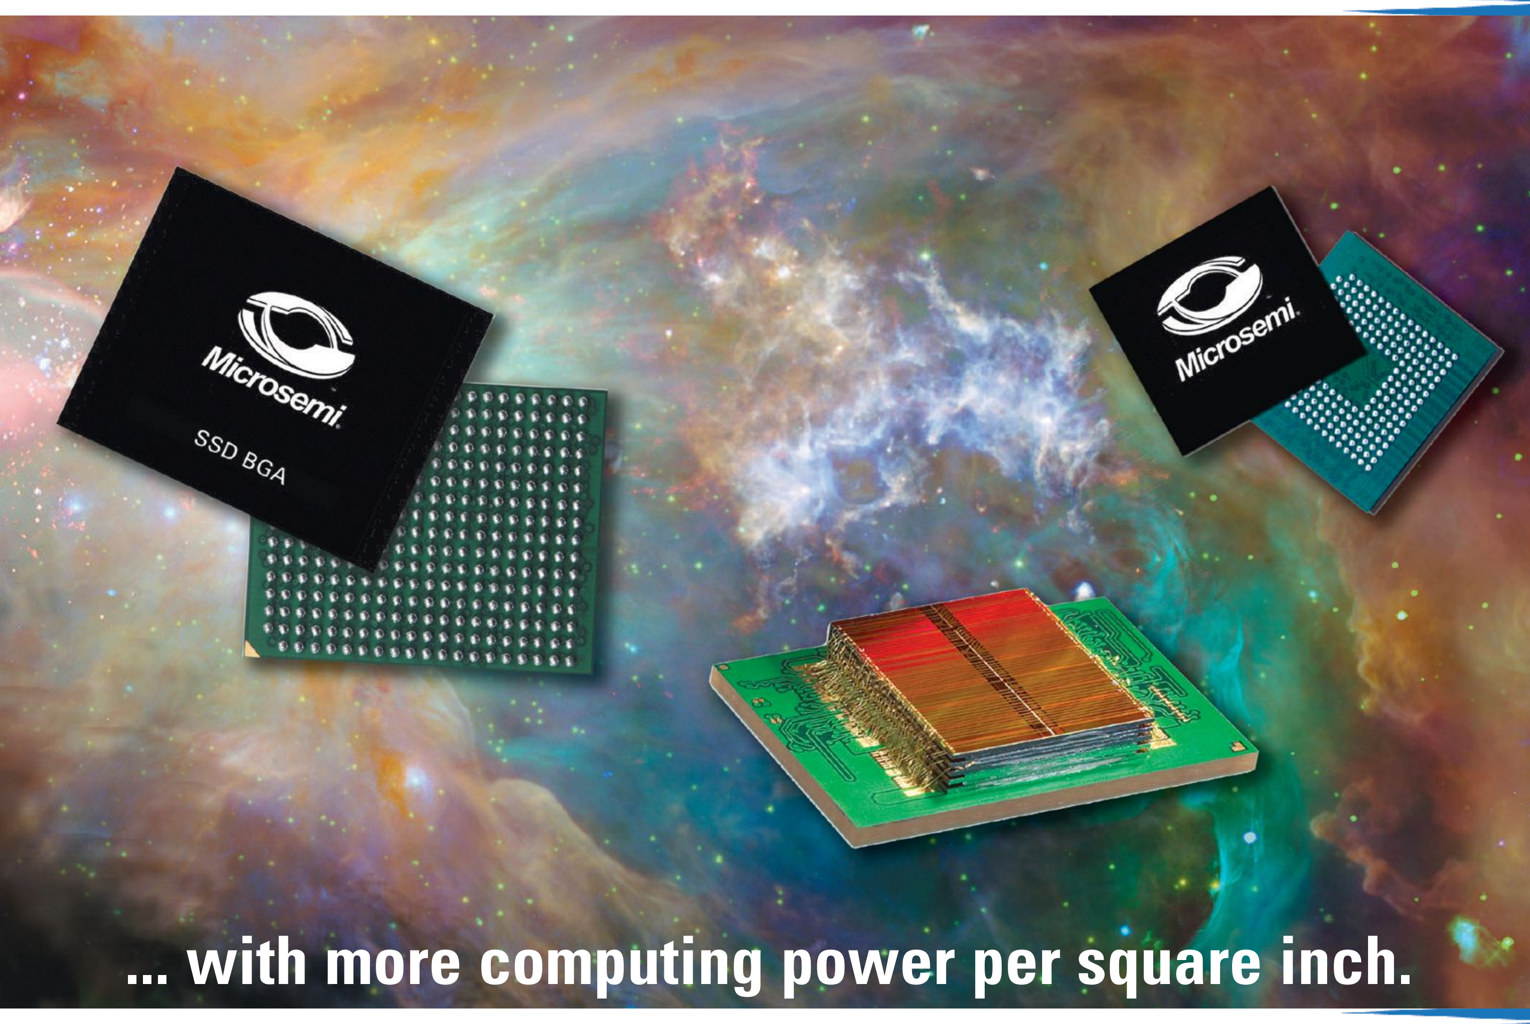

Our SLC NAND flash SSD BGA provides big features in a miniaturized package, freeing up precious board space. It is engineered specifically for the defense market, designed to ensure data reliability in mission critical rugged and mobile systems. Constructed using a 32-bit RISC processor as its core storage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

Migrating and modernizing legacy applications brings material cost savings

More than ever, the Administration is laser focused on making government agencies accountable for their IT spending, and it has publicly called on agencies to deliver more value from existing resources. One way to achieve this is through application modernization, which proffers an incremental path in migrating COBOL applications onto lower-cost platforms such as Linux – minus the code rewrites.

Under an increasingly scrutinizing eye, defense agencies are undoubtedly feeling the pressure to cut costs and do more with less, all while delivering top-tier end-user service levels. One of the most effective ways to align with this agenda is to migrate expensive-to-maintain legacy applications off proprietary mainframes and onto less expensive, nonproprietary modern platforms. However, the path to achieving this desired outcome, at least historically, has been anything but simple.

Migrating legacy applications can be a tremendous undertaking for any defense organization. One only needs to observe the Office of Management and Budget's recent modernization spending freeze on underperforming defense projects to appreciate the challenges involved. Historically, when assessing a migration/modernization project, defense agencies have considered one of these options:

- Rewriting applications to operate on a more modern or lower-cost platform
- Replacing applications entirely
- Deciding to maintain the status quo due to the costs and risks involved

However, these options are all less than ideal. Using rewrite or replacement strategies to migrate applications, especially when conducted on a large scale, invites failure on many levels. More often than not, these risks result in significant delays. Consequently, these projects frequently run over budget and exceed the needed timeframe. Maintaining the status quo is no longer acceptable either, as the costs of maintaining legacy applications on propriety platforms continue to escalate. However, there is another option. Today, many defense agencies are pursuing application modernization, a lower-risk, lower-cost strategy to migrate key applications for cost savings. In some cases, this strategy can provide up to an 80 percent reduction in costs, all without changing a single line of code.

Living in a COBOL world

COBOL code, which recently celebrated its 50th birthday, remains one of the most prevalent mission-critical application programming languages today. There is understood to be more than 200 billion lines of COBOL code in production, with hundreds more being created every day. It is equally ubiquitous within the defense department. Many mainframe applications written in COBOL continue to run the DoD's mission-critical business processes, including logistics, accounting, and personnel.

It is clear that mainframe applications written in COBOL still deliver outstanding value for defense agencies, and many contain business logic that has evolved over decades in support of the defense mission. Rather than putting decades of business logic at risk – by rewriting these COBOL-based applications in new programming languages or by replacing them entirely – agencies need to consider application modernization. Such modernization reuses proven code and delivers a much less invasive and more prudent migration approach.

Think Linux for cost savings: Application modernization provides the path

One path to application modernization, for example, is a Linux-compatible compiler, such as the one offered by Micro Focus. These compilers take COBOL source code and recompile it into an executable that runs on Linux. By encapsulating the underlying COBOL code, the application remains virtually unchanged, and it is enabled to function on the lower-cost Linux platform. This specific migration approach is gaining popularity among defense agencies that have an Integrated Facility for Linux (IFL) built into their mainframe environments to consolidate workloads. By using this migration tactic, made possible by application modernization, several large defense agencies have been successful in migrating key applications to their IFL to reduce the overall Million Instructions Per Second (MIPS) rate on the proprietary side of the mainframe. The lower the MIPS rate on the propriety platform, the lower the operating costs.

This migration approach is considered less invasive because it leaves existing code and business logic untouched, unlike rewrite or replacement strategies. It also allows for the incremental migration of IT resources, an approach gaining clout among key government leaders, including White House CIO Vivek Kundra. Kundra has advocated a more modular, service-based approach to modernizing IT assets. Using application modernization, defense agencies can literally pick and choose which applications they wish to migrate for cost savings, offering agencies much greater flexibility in defining the scope of a migration effort. Also, because the underlying code is not affected, the desire to migrate one application will not commit an agency to making system-wide changes to preserve interoperability; this further bolsters the ability to define a project's scope.

Deliver results quickly, avoid costs and risks

For agencies struggling with the high costs of maintaining applications on the proprietary mainframe, the good news is that there is a solution. Application modernization offers defense organizations a less invasive, incremental approach, allowing IT staff to migrate expensive-to-maintain COBOL applications onto distributed systems and lower-cost platforms such as Linux, all without changing a single line of code. Combined with a thorough analysis of the organization's application portfolio and accurately defining requirements, application modernization can provide a much faster and less risky path for defense agencies to realize long-term cost savings in the IT department.

Ken Powell *is President of Micro Focus North America. He can be contacted at Ken.Powell@microfocus.com.*

Daily Briefing: News Snippets

By Sharon Hess, Assistant Managing Editor

www.mil-embedded.com/dailybriefing

The end of one era, continuation of another

After nearly three decades, it's "prime" time for Boeing to wish the Space Shuttle Discovery a fond farewell (Figure 1). However, Boeing did not leave its NASA-sponsored progeny without assistance prior to Discovery's final mission launch from Kennedy Space Center to the International Space Station (ISS) on November 3; rather, the company performed shuttle payload processing, comprising the Leonardo and Express Logistics Carrier 4 (ELC4) and Permanent Multipurpose Module (PMM). More precisely, Boeing's Checkout, Assembly, and Payload Processing Services (CAPPS) team, in conjunction with PMM designer Thales Alenia Space, transformed Leonardo into a permanent fixture for extra storage and spare parts; CAPPS also integrated ELC4, designed to give the orbiting complex more storage. Boeing also serves as prime for NASA's ISS.

Figure 1 | The Space Shuttle Discovery recently set out on its final mission, from Kennedy Space Center to the International Space Station (ISS). Pictured: Discovery docked to the ISS's Pressurized Mating Adapter (PMA-2), photo courtesy of NASA Marshall Space Flight Center (NASA-MSFC)

Egyptian military and U.S. Navy go "mod"

A recent contract between Raytheon Co., Naval Sea Systems Command, and Egypt will likely produce its own version of "The Mod Squad" … figuratively speaking anyway. The $17 million contract calls for a pair of upgraded/refurbished Rolling Airframe Missile (RAM) MK 49 Mod 3 Guided Missile Launch Systems (GMLSs) and hardware accoutrements, in addition to a brand-new Mod 3 GMLS suited to Egyptian fast missile craft. The MK 49 GMLS, along with the MK 44 guided missile round pack, is part of the NATO-inspired and -managed RAM MK 31 Guided Missile Weapon System (GMWS) developed by the U.S. and Germany to fulfill Anti-Ship Missile (ASM) needs in defense systems. If all contract options are exercised, the contract's value could climb to $32 million. The contract delineates Egypt at 62.9 percent and the U.S. Navy at 37.1 percent of the purchase. Work will commence at Raytheon's Tucson, Arizona locale, and a completion date of January 2013 is expected.

OWL to spot weapons threats

Well-known for introducing the world to last year's ultra-secure BlackBerry cell phone, perhaps SRC, Inc. will become nearly as well-known for a recent entry into the annals of U.S. Army contracts: a $500,000 contract for preliminary design of the Army's radar system dubbed the "Omni-Directional Weapon Location (OWL)." The multimission-savvy OWL system proffers surveillance, not just of the local area but "over a hemispherical coverage area – a capability that is not currently achievable with existing weapon location radars," according to the company. OWL is designed to be "affordable" yet state-of-the-art, locating, tracking, and detecting rocket, cannon, and mortar firing locales within a wide threat-trajectory spectrum.

Army's network management … think "positive"

Being able to optimize and manage at-the-halt, fixed, or on-the-move networks' performance is critical to mission success on the modern battlefield, and the U.S. Army Research, Development, and Engineering Command (RDECOM) and an unnamed Coalition country's military are marching closer to the goal. Specifically, Hughes Network Systems, LLC's HX ExpertNMS network management system recently garnered "positive results from trials." Providing the aforementioned stationary or on-the-go network optimization and management capabilities, HX ExpertNMS is a "full-featured capability" iteration of Hughes' HX System commercial network management system (Figure 2). The military-flavor network management system is touted to render quick setup via an installation wizard, enabling the military to rapidly fulfill assignments so they can move on to the next locale. Additionally, any network issues are easily identified, thanks to HX ExpertNMS's quick-view dashboard offering fast diagnostics of bandwidth snafus or network malfunctions.

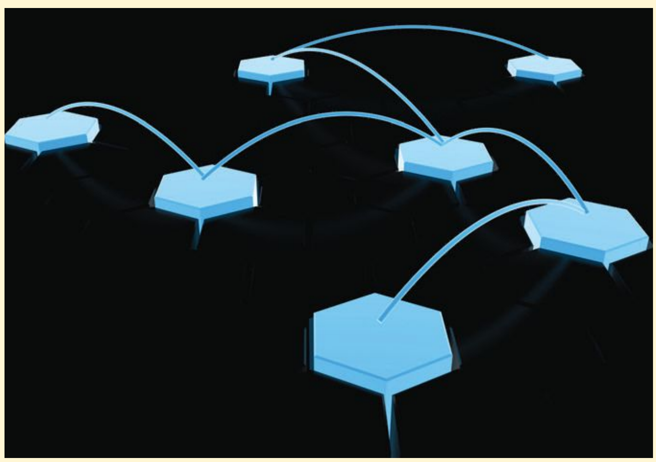

Figure 2 | RDECOM and an unnamed Coalition country's military will benefit from the recently garnered "positive results from trials" achieved by Hughes Network Systems, LLC's HX ExpertNMS network management system. Stock photo

No human photographic memory? No problem

In the absence of *homo sapiens* with 100 percent accurate photographic memories, a recent contract between EADS North America's American Eurocopter Corporation unit and Curtiss-Wright Controls, Inc. (CW) will provide the next best thing: a video management system, aerial style. Specifically, CW will proffer the Skyquest Video Management System (VMS) to American Eurocopter, who will then suit up Security and Support (S&S)-configured versions of the U.S. Army's UH-72A Lakota Light Utility Helicopter (LUH) (Figure 3) with Skyquest, as part of the LUH Batallion Mission Equipment Package (MEP). Each S&S MEP-equipped Lakota is slated to feature the VMS processor, a double-deck DVR, and a triad of multifunction mission displays. Skyquest VMS's job duties include enabling simultaneous viewing of several video sources. Meanwhile, work on the contract is anticipated through 2015, at which time the Army may have ordered the contract's ceiling of 345 helicopters. The Army currently has 187 of the 345 LUHs on order.

Figure 3 | Security and Support (S&S)-configured versions of the U.S. Army's UH-72A Lakota Light Utility Helicopter (LUH) will soon be equipped with the Skyquest aerial video management system. Photo by Sgt. 1st Class Tyrone Walker

Can the U.S. Army's radar take the heat?

Unmanned Aerial Systems (UASs) are once again making headlines, this time for hosting a Lockheed Martin-built Tactical Reconnaissance and Counter-Concealment-Enabled Radar (TRACER) in-flight. Accordingly, the MQ-9 UAS's recent voyage with the TRACER Synthetic Aperture Radar (SAR) on its back signified the first time a fixed-wing UAS flew a penetrating radar. The impetus for the test was to ascertain the dual-band TRACER's viability for remote operation on a high-endurance platform. Next up: TRACER is slated for a ride onboard the Predator B (Ikhana) NASA-operated unmanned aircraft. The testing mission: to gather high-res images and to find out how successfully TRACER's external unpressurized pod can encase the radar system's RF segment in a harsh environment.

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

Figure 4 | The U.S. Navy's CVN 78's MK-7 linear hydraulic system is about to be supplanted by a modular next-gen aircraft recovery system dubbed "Advanced Arresting Gear" (AAG). Northrop Grumman photo

Navy's unmanned system-within-a-system progresses

Of course unmanned vehicles are making a huge splash on the modern mil-embedded scene, but what about unmanned systems within the system? Perhaps those will be next to go to low- or no-maintenance, per a recent contract between U.S. Navy CVN 78 aircraft carrier vendor General Atomics and "PT" (formerly "Performance Technologies"). The contract stipulates that PT provides its COTS IPnexus Application-Ready Systems to General Atomics' Electromagnetic Systems (EMS) Division to facilitate sensor-processing capabilities on CVN 78's modular next-gen aircraft recovery system dubbed "Advanced Arresting Gear" (AAG) (Figure 4). IPnexus renders precision information reporting and data monitoring in addition to self-diagnosis accompanied by maintenance alerts. AAG aims to supplant the CVN 78's MK-7 linear hydraulic system, which requires manual adjustments of tension each time the aircraft lands. In contrast, AAG provides a digital control system proffering built-in diagnostics and testing, along with prognostics technology and health monitoring assessments to increase safety margins and reliability – while allowing military personnel to focus their energies elsewhere. Meanwhile, Northrop Grumman serves as CVN 78 prime.

VPX REDI gets thumbs-up from ANSI

Now VPX developers have something new to celebrate: the American National Standards Institute (ANSI) ratification of the VPX REDI (48.0) base standard (Ruggedized Enhanced Design Implementation) and supporting dot specs: 48.1 – Mechanical Specification for Microcomputers Using Air Cooling Applied to VPX; 48.2 – Mechanical Specification for Microcomputers Using Conduction Cooling Applied to VPX; and 48.5 – Mechanical Specification Using Air Flow-through Cooling Applied to VPX. The VPX REDI embedded computing standard (Figure 5) was developed to specify 2-level maintenance and advanced cooling for VPX (VITA 46).

VPX REDI

Figure 5 | VITA's VPX REDI (48.0) base standard and 48.1, 48.2, and 48.5 dot specs recently achieved ANSI ratification. Logo courtesy of VITA

U.S. Army photo by Drew Hamilton

COTS software in critical systems: The case for Freely Licensed Open Source Software

By Dr. Robert B.K. Dewar

COTS software helps reduce development costs for large, long-lived systems, but "COTS" does not mean "proprietary." Freely Licensed Open Source Software (FLOSS) brings COTS benefits but without the restrictions and vendor tie-in typical of proprietary products.

The International Space Station (ISS) (Figure 1) is the antithesis of the notion of Commercial Off-the-Shelf or "COTS." There is and will only be one of these amazing items built, and you certainly won't find it "on the shelf." Yet some of the critical software that runs this station was created using COTS tools including, for example, the GNAT Pro Ada compiler, used to build the software running the Canadian Space Arm.

Why the choice of COTS software for this purpose? Procuring or creating software for large, complex systems is a difficult and expensive process. Addressing this need by increasing use of COTS is a natural response to this difficulty, and it seems to offer many advantages at first glance: reduced cost, economies of scale, widespread use leading to greater reliability, and availability of people who know the system well. These considerations apply especially strongly to one-off projects like the space station, where it is otherwise expensive (and more risky) to build specialized tools.

However, in practice, the promise of COTS is not so easily achieved. One major problem is vendor tie in: If you obtain a proprietary COTS system from a vendor, you are tied into that vendor for support, modifications (in the common case where the COTS software does almost but not quite what you want and must be modified), and availability of updates and improvements. There is also the problem that the general COTS market thrives on frequent updates and rapid obsolescence of old versions. The example of XP/Vista is instructive here, where Microsoft pushed to abandon XP long before users were ready to move to Vista. Even if Microsoft is persuaded (as looks likely) to continue XP support for perhaps another decade, even that is not nearly long enough for some projects.

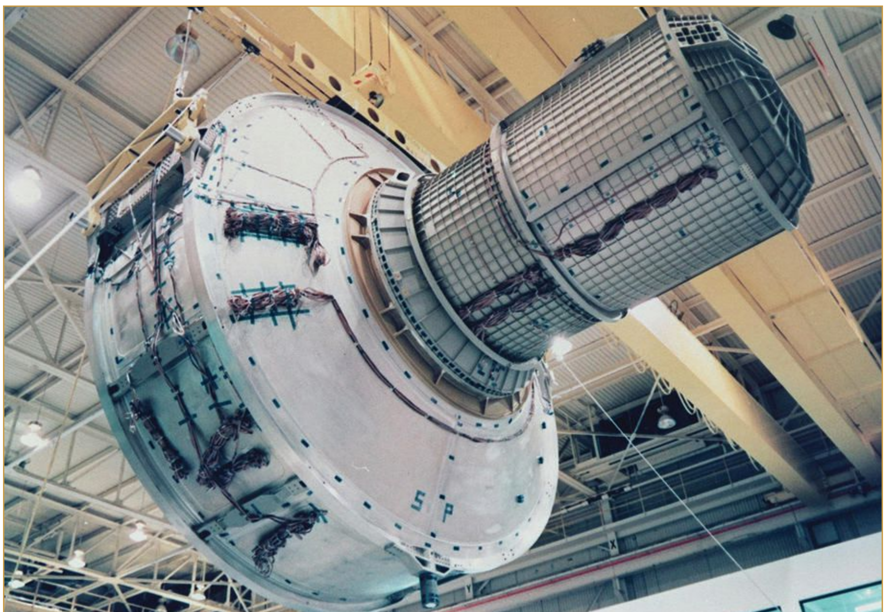

Figure 1 | The International Space Station is the antithesis of the notion of Commercial Off-the-Shelf or "COTS." NASA photo

Modern aircraft such as the Boeing 777 and Boeing 787 (Figure 2) have a very long life, measured in decades. These aircraft have very complex software systems and require the utmost reliability. Software onboard these planes is not only built using COTS tools in many cases, but there are actual COTS software components aboard – where not only the software, but also the certification materials are available "off the shelf." Given the very long expected life of these planes, concerns about continued availability of COTS tools and components become a critical issue.

One of the things that distinguishes software from hardware (both inside and outside the COTS arena) is that it is notionally easy to modify software. It probably takes no more than a few minutes to completely replace the avionics software of a plane with a new version (not counting, of course, the time and effort to prepare and certify the new version). In the case of one well-known military plane, I have been told that every tail number corresponds to a slightly different set of software components. Hence, the notion of build-once-and-forget certainly does not apply in the case of avionics software.

This ability to change software means that refreshes and improvements to software systems can be implemented in a manner that would not be practical for corresponding hardware systems. Going back to the aforementioned Canadian Space Arm, completely replacing the hardware would be unthinkable; however, there has indeed been at least one complete refresh of the software. But given the possibility of executing such updates, the long-lived availability of software tools and components becomes critical.

Additional considerations in the choice of COTS in software construction include the issues of reliability and warranties. The very fact that a given software component may be used on millions of PCs where reliability is not crucial (or at least is not seen as crucial by manufacturers) means that its use in high-reliability environments is not welcomed by the manufacturers: It is, in fact, actively discouraged with legal disclaimers. In this context, trying to get comprehensive warranties for such products may be completely infeasible.

On the other hand, extensive use can be a significant contributor to reliability, so this widespread use can also work in a positive direction. Software developers can't guarantee correctness and reliability by testing alone, but nevertheless, a product that has been very widely used tends to become more reliable over time, as problems are smoked out by actual use. In some cases, this method must be relied upon to achieve reliability. As an example, consider ground-based air traffic control systems. Very often, these have to be built on top of general-purpose operating systems such as IBM's AIX. It is out of the question to certify such systems using standards similar to DO-178. They are far too complex. However, if such a system has been in use for decades, as in this case, that experience instills vital confidence in the system.

Thus, COTS software can provide many advantages. But the fundamental difficulties of inability to make custom modifications and guaranteeing long-term support can often stand in the way of these COTS promises. Unless these problems can be solved, the military industry will fall back into the mode of procuring expensive custom technologies.

Using FLOSS COTS yields best of all worlds

One interesting response to these general problems of acquiring and using COTS software products can be found in the use of commercial Freely Licensed Open Source Software (FLOSS) products. The use of FLOSS COTS can potentially address these fundamental concerns and

Figure 2 | Given the very long expected life cycle of modern aircraft such as the Boeing 787, concerns about continued availability of COTS tools and components becomes a critical issue. Boeing photo

provide COTS advantages without the attendant disadvantages. These FLOSS advantages include:

Accessible source code
First, the fact that the sources are always available and can be freely modified means that it is at least possible to make modifications. Furthermore, these modifications can be made either by the customer, or by the manufacturer, or by a third party. In the case of proprietary products, the manufacturer can by technical and legal means maintain a firm control over such modification, but the free licensing associated with COTS FLOSS means that such control is not possible.

Perpetual availability
Second, the free licensing means that anything that is out there stays available forever. In the case of proprietary products, manufacturers can force upgrades by licensing conditions, or simply by making products stop working. (For example, some functions of Quicken stop working after a few years if you don't keep upgrading to the latest version.) In the case of FLOSS products, the licensing guarantees perpetual availability, and long-term support can be provided by the vendor, the customer, or a third party.

Emphasis on support
Finally, the commercial FLOSS market is support-oriented. It doesn't work to just dump a FLOSS product on the market with no support and expect people to buy it. The precise reason people will pay for something that might otherwise be downloaded free from the Internet is to get a high level of support. If this support means making minor modifications for specific customers or keeping support alive forever on old versions, then it will be provided at a reasonable price. If the manufacturer won't provide this at a reasonable price, a third party can step in to provide that support. In the case of the GNAT Pro Ada compiler we mentioned in the context of the space station, AdaCore makes new versions available annually, but we still have customers using versions well over 10 years old, which we expect to continue to support indefinitely. Making minor custom improvements is an important part of our commercial offering.

FLOSS for success
In summary, COTS software tools and components are here to stay in the context of critical projects, and the increasing reliance on FLOSS-based COTS tools can ensure that the promises of COTS can flourish without the common disadvantages of this approach.

***Dr. Robert B.K. Dewar** is cofounder, President, and CEO of AdaCore. He has been involved with the Ada programming language since its inception in the early 1980s. He has coauthored compilers for SPITBOL (SNOBOL), Realia COBOL for the PC (now marketed by Computer Associates), and Alsys Ada, and he is a principal architect of AdaCore's GNAT Ada technology. Robert has also written several real-time operating systems for Honeywell, Inc. He is frequently invited to conferences to make presentations on a range of topics including open-source software, programming language issues, and safety certification.*

AdaCore
1-877-787-4628
www.adacore.com

U.S. Navy photo by Mass Communication Specialist 2nd Class Ja'lon A. Rhinehart

AdvancedTCA delivers military COTS promise

By John Long

Changes in the military equipment market are driving new system designs to leverage commercial products and technologies, and AdvancedTCA is stepping up to the front lines to meet this challenge.

Changes in the military equipment market are driving new system designs to leverage commercial products and technologies. The ideal technology would offer design features readily adaptable to military applications, a modular architecture, and a diverse vendor base to keep costs down and innovation high. AdvancedTCA is an open-standard COTS form factor that – with the support of industry organizations like the Communications Platforms Trade Association (CP-TA) ensuring multivendor interoperability – is proving its ability to deliver on all these promises in military applications.

Additionally, the military market spending paradigm has shifted significantly. For instance, traditional military acquisition called for purpose-built equipment that a single prime contractor designed for a specific application. The result: a growing accumulation of proprietary system designs that could neither interoperate nor readily exchange data.

Another aspect of the military spending paradigm shift is a lag in military systems technology behind commercial products. The time required to design systems from scratch, prove their reliability, and enter production resulted in military systems that were years or even decades behind commercial systems. This technology lag, as well as the high cost of traditional acquisition, has prompted the U.S. military to increasingly require designs based on COTS products.

And finally, yet another shift in the military market is the growing need for network-centric systems. In addition to full-scale war, military actions are increasingly involving unconventional conflict, such as insurgencies, which demand that both the warfighter and upper command have rapid access to actionable intelligence. Addressing these objectives requires a greater emphasis on Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) – not only on weapons superiority. AdvancedTCA addresses these network-centric requirements by delivering built-in reliability and ruggedization as well as high performance. Meanwhile, CP-TA helps drive interoperability for AdvancedTCA components.

Interoperability is key

These C4ISR systems must exhibit a high degree of interoperability. One reason is that they must be able to handle information from a wide array of information sources, including satellites, Unmanned Aerial Vehicles (UAVs), ground troops, radar, and sonar. Systems must also be able to reliably exchange information and coordinate command decisions among differing service branches as the military moves to a network-centric doctrine. An increasing emphasis on cooperation with and reliance on allies, as directed by President Obama and Defense Secretary Gates, also demands interoperability.

Achieving such interoperability through traditional proprietary design approaches, even if COTS-based, is extremely difficult. However, one proven solution to interoperability issues is to design systems around a framework based on open standards. AdvancedTCA is one such open standard, originally created for commercial networking equipment designs, that has a broad multivendor ecosystem. VPX is another standard that is a contender for serving the military segment for next-generation high-performance computing systems. However, VPX has a limited number of vendors and is still challenged by interoperability issues, though the new OpenVPX standard is making headway on solving VPX system interoperability dilemmas. The AdvancedTCA architecture is also ideal for addressing the military's emerging network-centric design needs. AdvancedTCA is a high-reliability, modular design approach that offers the robustness and performance that the military requires. By comparison, AdvancedTCA offers nearly double the computing power of comparable VPX systems.

AdvancedTCA: Upon closer examination

AdvancedTCA's foundation is a card cage with a high-speed, protocol-agnostic, switched serial backplane that accepts blades containing computing, routing, I/O, and other networking system functions in any desired combination. Each blade is also modular. Advanced Mezzanine Card (AMC) modules on the blades carry a blade's unique computing and I/O functions (Figure 1). This level of modularity allows highly complex and optimized systems to arise from a handful of standard module types.

Figure 1 | AdvancedTCA is a highly modular architecture using a switched-serial backplane connecting large carrier cards, configurable using mezzanine cards of varying sizes.

The blades are large, supporting four AMC modules apiece, which allows a high degree of functionality on a single blade. The standards also allow for the creation of systems using AMC modules plugged directly into a backplane. Systems built to this MicroTCA standard have full architectural and software compatibility with AdvancedTCA systems and use the same AMC modules but offer a much smaller form factor.

AdvancedTCA's built-in reliability and ruggedness

Because service providers consider their telecommunications systems to be "mission critical," both high reliability and ruggedness are vital to the mil/aero industry.

High reliability

AdvancedTCA includes many high-reliability architectural features that can translate directly to the mil/aero industry. The AdvancedTCA specification calls for built-in system management and error detection, hooks for fault isolation and redundant system configuration, and the ability to perform live blade replacement (hot swap) with electronic keying to avoid errors (Figure 2). These features allow system designs that achieve greater than five-nines (99.999 percent) reliability and short Mean Time To Repair (MTTR). This level of availability enables mission-critical computing and is one reason that AdvancedTCA has gained traction into applications such as UAV ground control stations.

PICMG has recently increased the permitted power dissipation for AdvancedTCA blades and is increasing backplane speeds to support 40 Gbps links between blades.

Ruggedness

There is also effort underway to standardize methods for enhancing the ruggedness of AdvancedTCA to meet the ANSI/VITA 47-2005 (R2007) specification, also known as the Environmental, Design and Construction, Safety, and Quality for Plug-In Units standard, which in turn draws from MIL-STD-810F. The target specifications are environmental classes EAC1 and EAC4 for convection-cooled systems, which include operation with ambient temperatures from 0 ºC to +55 ºC, 2 *g* (5-100 Hz) or 8 *g* (100-1,000 Hz) vibration, and 20 *g* shock.

This enhanced ruggedness is not a first step toward addressing military needs; it's merely an additional one. The original AdvancedTCA environmental specifications follow Network Equipment Building Standard (NEBS) requirements, which go beyond the laboratory and office environment into something much more demanding – the central office.

The NEBS Central Office Relay Equipment (CORE) GR-63, Level 3 specification defines those environments and is similar to MIL-STD-167 for equipment in shock-isolated cabinets. The NEBS specifications include a requirement to withstand the mechanical shock and vibration equivalent of a magnitude 8.1 earthquake and operation in temperatures of -5 ºC to 50 ºC for up to 96 hours (continuous operating environment range is 5 ºC to 40 ºC). In addition, the equipment must operate reliably after coming out of 72 or more hours of storage at -40 ºC to +70 ºC and being brought to 25 ºC within five minutes. HALT/HASS testing can be used to ensure that the equipment will operate reliably within this range.

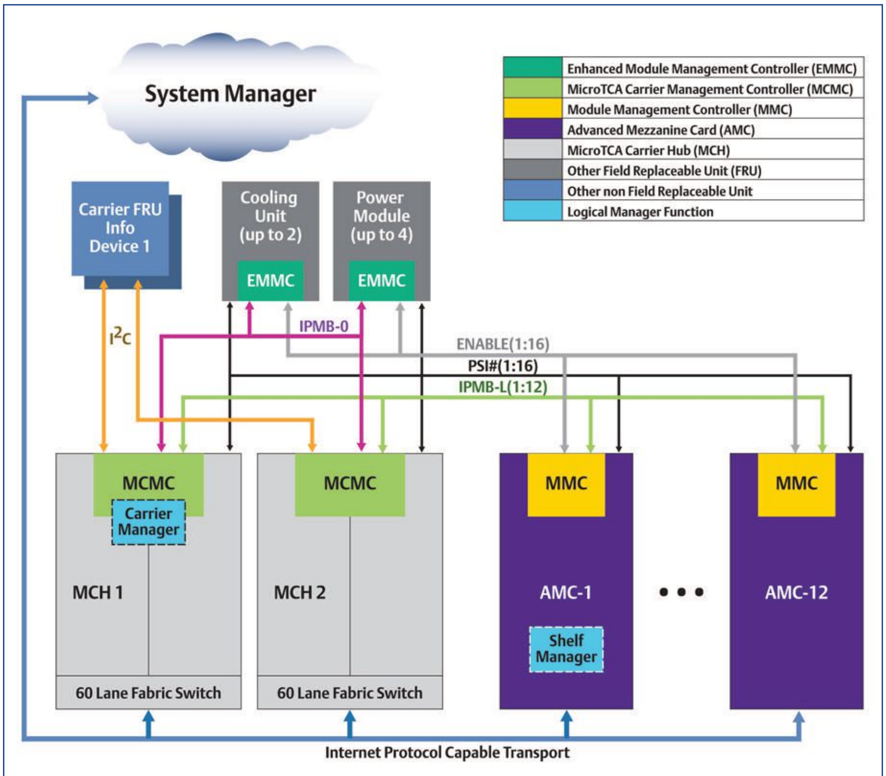

Figure 2 | AdvancedTCA includes complete system management functions for fault detection and module-level control of high-reliability features.

Proven military program performance

While AdvancedTCA's specifications are not as stringent as the military's most extreme environmental requirements, they are suitable for a majority of installations such as C4ISR and command and control centers. Compared to the battlefield, for instance, shipboard and command center environments are relatively benign and not as demanding on equipment. AdvancedTCA-based systems have already proven themselves fully able to meet the installation needs of such "benign shelter" environments. Lockheed Martin[1], SAIC[2], and Hughes Networks[3] have all publicly declared the adoption of AdvancedTCA in some of their deployed military systems for such applications. An AdvancedTCA-based design is currently providing support in the Multimission Maritime Aircraft (MMA)[4] system, for instance, as well as aboard the P8A Poseidon airframe.

AdvancedTCA is also under evaluation for programs such as the U.S. Navy's Consolidated Afloat Networks and Enterprise Services (CANES)[5] program, which is intended to consolidate the shipboard network infrastructure into a standard off-the-shelf solution. One goal of the CANES project is to support a two-year refresh cycle on system software and a four-year refresh cycle on hardware for the next 20 years. The modular nature of AdvancedTCA makes that refresh cycle easy to achieve.

AdvancedTCA: Ready for military deployment

AdvancedTCA thus fulfills the requirements of a COTS solution for many of today's military market demands. It is modular, flexible, and widely supported; meets the ruggedness requirements of most applications; and simplifies both maintenance and system evolution. It is an established and technically robust architecture with industry support to ensure multivendor interoperability and continued technical evolution. Further, it has been field-proven both in its original commercial application and in several military systems. AdvancedTCA has shown that it delivers on all the promises of COTS in military system designs.

References:

[1] "Rugged MicroTCA opens up," by Joe Pavlat, Contributing Editor, *Military Embedded Systems* magazine, www.mil-embedded.com/articles/id/?4051

[2] "AdvancedTCA joins the Navy," by John Walrod, SAIC, *CompactPCI, AdvancedTCA & MicroTCA Systems* magazine, www.compactpci-systems.com/articles/id/?2139

[3] "Iridium Selects Hughes to Develop Access Network Controller as Major Step Toward 'Iridium NEXT,'" http://investor.iridium.com/releasedetail.cfm?releaseid=443411

[4] "Why AdvancedTCA merits serious consideration for defense systems," by Scott Burgan, Consultant, The Corner Group, www.intel.com/design/embedded/military/atca.pdf

[5] "AdvancedTCA: Open standard moves beyond the Central Office," by Brian Carr, Emerson Network Power Embedded Computing, *CompactPCI, AdvancedTCA & MicroTCA Systems* magazine, www.compactpci-systems.com/articles/id/?4188

John Long, *a product line manager at RadiSys, also serves as a CP-TA member. He has more than 10 years of experience in the communications industry holding various marketing, sales, and operational positions with Intel, Dialogic, and AT&T. John has an MBA from Carnegie Mellon University. He can be reached at John.Long@radisys.com.*

CP-TA
503-619-0680 • www.cp-ta.org

VITA 65 in record time

Interview with Dr. Ian Dunn, Chief Technology Officer, Advanced Computing Solutions, Mercury Computer Systems, Inc.

EDITOR'S NOTE

Gone are the days when a VITA spec takes years to become a reality. The following interview is a behind-the-scenes look at why OpenVPX was born and how it achieved rapid ANSI approval. Editor Chris Ciufo interviewed Ian Dunn shortly after ANSI blessed the spec. Edited excerpts follow.

MIL EMBEDDED: ***Last year something called* OpenVPX *was born, amidst controversy. Where has the standard ended up?***

DUNN: We made it through standardization with VITA and ANSI. It took a few months longer than we anticipated, but that's not bad for outside standards work. The more important deadline we wanted to keep was the working group's commitment to the original date of handing it off to VITA, which occurred on time [at MILCOM in October 2009]. The original hope was to get the ANSI process through in one round with VITA.

MIL EMBEDDED: ***How many rounds did it go with VITA and ANSI?***

DUNN: It went two rounds with VITA and two rounds with ANSI, too. The official publication date of VITA 65 as an ANSI standard was June 18. I think it was May 24 that OpenVPX was first submitted to ANSI, so the entire OpenVPX process was about a year and a half.

MIL EMBEDDED: ***What were the biggest achievements – and disappointments – in the OpenVPX process?***

DUNN: An important achievement is that VITA 65 is a "living, breathing" specification. We figured out questions such as: How could it be augmented, and how could things be deleted? The idea was, for example, 1) What if someone wants to add a profile, or what if someone wants to add a definition of the architecture? The normal VITA cadence was to take years to make a major revision. We approached VITA about this, and their statement was that each working group could set its own guidelines within its own specification. For ANSI, my understanding is that if a committee ratifies an update process in the spec and ANSI signs off on it, that allows the spec to be modified without waiting for those major revision cycles.

MIL EMBEDDED: ***OK, so faster revisions … any other goals with OpenVPX?***

DUNN: There were really a couple of primary goals. First, we were very concerned with the status of the embedded industry in general. We were seeing a lot of pressure from the primes to just replace embedded form factor architectures with bladed architectures and then to start ruggedizing those. And secondly, a lot of it had to do with interoperability mismatched with mainstream market-driven technology. And then the final goal was to create an output with a minimal skill, create an environment where the amount of variability was minimized. We were convinced that we could live with a handful of profiles. Later on as we engaged everybody, we figured it was more than that. And I think in reality we've ended up with an environment now where the core usable profiles are probably a dozen if not more.

MIL EMBEDDED: ***Can you explain what a "profile" is?***

DUNN: The OpenVPX standard defines the allowable combinations of interfaces between the module, backplane, and chassis to facilitate interoperability and reuse. These interfaces are called *profiles*. For example, slot profiles define the connector type and how each pin, or pair of pins, is allocated. The module profile defines what communication protocol may be used on each module interface as defined in a corresponding slot profile, connector type, module height (6U/3U), and cooling method (forced air/conduction). The backplane profile defines which pins or set of pins is routed in the backplane, and which pins are routed to the rear transition module. The backplane profile also defines allowed slot-to-slot pitch.

MIL EMBEDDED: ***OK so any other achievements before we move on?***

DUNN: The enterprise goal – That was a very important outcome in our perspective, too: OpenVPX now can be used to build enterprise-class architectures in an embedded backplane. And so customers can now build a development environment with PCs, or blades, or servers, and they can transition that to a rugged design. This means OpenVPX vendors are in a position to sell rugged, enterprise-class products at the end of the day. *[**Editor's note:** As we went to press, Mercury announced what they call the industry's first dual quad-core Xeon embedded server on a VPX form factor. Clearly the company is executing a long-term strategic plan.]*

A very important aspect of this enterprise ability is rugged LRUs. The second is the basic ability to do giga-clustering. And then finally, the last element is the expansion plane capability, meaning that an enterprise-class architecture has a PCI-based expansion plane. We really wanted that concept to find its way into a rugged architecture.

MIL EMBEDDED: ***Mercury was very clear about allowing nothing to derail OpenVPX's schedule.***

DUNN: Yes. I've got to credit the community because all the OpenVPX players understood that. At first, I thought, "Wow, you know, we're really talking about a business decision here. So how do I get access to the business leaders of these corporations?" Well unfortunately, I had to create some controversy, and I had to put people's businesses on the line. *[**Editor's note:** See Chris Ciufo's column entitled "OpenVPX Industry Working Group: Open for business, or just controversy?" at www.mil-embedded.com/articles/id/?3818.]*

MIL EMBEDDED: ***What about the primes' viewpoint?***

DUNN: Regarding the primes who are not on the committee with us but who are just primes winning deals and winning programs: One of the outward-facing accomplishments we wanted to achieve was to reset the image of the embedded community to not be perceived as a closed architecture community. Of course this was terribly important for Mercury because we had the reputation of being the quintessential closed company with RACEway; RACE++, and so on.

After OpenVPX we've made some progress. Particularly after it was handed over to VITA, we talked to programs and offices within the business management at the primes and, even to a certain extent, the government. We talked up the fact that OpenVPX is in the spirit of what the government is trying to accomplish from an open architecture perspective. Now realistically, the government is many, many layers above the hardware, usually. But it created a good conversation because it meant the primes could use OpenVPX as a way to showcase how they were being open, and we found some very good reception.

MIL EMBEDDED: ***With all these goals already achieved, where will OpenVPX head now?***

DUNN: While there are more OpenVPX profiles than perhaps we wanted with our original goal, one of the architecture's successes is that it locked down certain categories of signals. This is the planes concept in the specification: the management plane; the control plane, which is the gigi-clustering plane; the data plane; the expansion plane; and then the user I/O [plane].

I think the next frontier is the user I/O dimension, which is the sensor plane or the "user plane." This is the area of a subsystem that is very useful to the primes because it's the area where, for instance in ground combat operations, computers now need to include user interfaces while serving as net-centric resources for the entire mission vehicle. The original VPX space had nothing locked down, with all kinds of planes everywhere. OpenVPX took six or seven of the dot specs that everybody really could agree upon and locked them down into a plane. Now what's left are the other six or seven dot specs that weren't making a lot of progress. And so it's allowed everybody to put some real core focus into the things that are still left.

MIL EMBEDDED: ***Let's switch gears and talk about Mercury's Cell BE effort of a couple years ago. What happened with that?***

DUNN: When Cell hit the street, one of the things Mercury was excited about was that it was the first processor in a while that could simultaneously execute the front and back ends of a sensor computing problem. So it could be used to do signal processing and signal exploitation, data exploitation, and data mining. And as we went around, the primes were all equally excited about this.

We developed a strategic relationship with IBM, then jointly developed some software and a value proposition. We then started helping prime contractors – both in commercial and defense – to use the Cell to create larger solutions, whether that be in industrial imaging, ISR, communications, or whatever. We also classified the Cell internally as a game chip, and lumped it in investments the rest of the industry was making in using gaming chips to execute high-performance computing.

And the Cell being the first and the most programmable variant of that, we jumped on and went after it with a vengeance. Now, when it hit the street, I was astounded at the number of customers who tried it. The main message that came back was that it was a great chip but that it was a productivity challenge: A lot of customers had a hard time getting performance out of Cell BE because of programmability issues. So Mercury offered to do customers' work for them, and we ended up doing quite a bit of that. Of course, we had a scalability issue there; we could only help so many.

MIL EMBEDDED: ***You had some commercial design wins with Cell BE, such as Massachusetts General Hospital.***

DUNN: Yes, that's right, and we also had some industrial imaging wins; however, in the meantime, on the defense side we were not as successful because in many cases we couldn't do the programming for defense customers – because of classification or security issues. And so we only ended up with one or two adoptions there.

And with the exception of the big OEMs that IBM was servicing – Sony and Toshiba – I think a lot of the opportunities that IBM saw for it as well were challenging because, as I mentioned, customers really could just not figure out how to efficiently get performance out of Cell BE even though it was a phenomenal architecture. That led to IBM not investing into a follow-on roadmap. So we've turned over all those internal investments and personnel expertise to the more generalized category of GPGPU. That way, we use our Cell BE learning and expertise to create a larger processing category that we're now building into products.

MIL EMBEDDED: ***Makes perfect sense. Last question: Curtiss-Wright recently acquired Hybricon, and Kontron acquired AP Labs. Do you have any comments on these acquisitions or why the industry might be turning toward systems integration houses?***

DUNN: Well, the way that we look at the acquisitions of some of our industry peers is that first and foremost, they're filling out their catalog, really – whether it's with boxes, backplanes, or whatever. That makes them more responsive to industry integrators, whether primes or otherwise. There's a class of companies that can do that well, but that's not really the way we think the industry's going, or where the value proposition is.

The government's moving to a model where they really want to procure capabilities, and they're looking to the primes for that. They're looking to commercial companies for that, and that's the strategy and marketplace we're after. We obviously must have our own products and third-party sourcing and relationships, but we're really focused at the solutions level, at the integrated capability level. Or maybe our goal is to be just below that – to be what we're calling the "application-ready level" to help the prime contractors build solutions faster, get them to the warfighter faster.

MIL EMBEDDED: ***Mercury is certainly among the top five systems integrators in the industry. You guys have never really been a board-level supplier. You've always been a solve-the-problem company.***

DUNN: Yeah, and it's a core capability and core asset of our company.

Dr. Ian Dunn is Mercury Computer Systems, Inc.'s CTO and senior architect responsible for technology strategy and R&D projects. He has 20 years' experience designing and programming parallel computers for real-time signal processing applications. Before joining Mercury in 2000, he consulted to Disney Imagineering and Northrop Grumman on automation and computing projects. Ian received his doctoral degree in Electrical Engineering from Johns Hopkins University and his undergraduate degree in Electrical Engineering from Oregon State University. He has authored numerous papers and a book on designing signal processing applications for high-performance computer architectures. He can be contacted at idunn@mc.com.

Mercury Computer Systems
866-627-6951
www.mc.com

U.S. Army photo by Travis McNiel

Advanced primary portable power requirements in next-generation military applications

By Eric Lind

A new generation of carbon fluoride batteries is stepping onto the front lines to power soldier-carried and unmanned aerial applications, thanks to much-increased volumetric and gravimetric energy densities.

The ideal battery would deliver ample power for the application, provide long shelf and service lives, withstand environmental extremes, and be safe to handle and use. The problem is the trade-offs required to achieve these different desirable operating characteristics. Delivering ample power in next-generation military applications requires higher energy and power densities than currently available in the batteries typically used in consumer electronics. Some military batteries, for example, employ toxic or hazardous chemicals to achieve a higher energy density, making them unsuitable for certain applications. Others may be able to deliver safe operation over a long service life, but only for low-amperage applications.

The following discussion examines the issue of energy density, both gravimetric and volumetric, and how existing primary lithium batteries stack up in next-generation military applications for soldiers and unmanned aircraft systems. The exclusive focus here on *primary* batteries – including sulfur dioxide, manganese dioxide, polycarbon monofluoride, and carbon fluoride – is not meant to diminish the growing importance of *rechargeable* batteries; rather, it is to limit the scope of the topic to one that can be covered in sufficient depth.

Lightening the load

The burden on today's soldier to carry an increasing amount of high-tech equipment, such as advanced soldier systems, next-generation radios, and imaging and sensing systems, is great and growing. Depending on the mission, the total weight of a soldier's gear ranges from 28 to more than 70 kilograms (kg), or 60 to more than 130 pounds. The batteries needed to power the electronic equipment can often constitute approximately 15 to 30 percent of the total load. And because the batteries must last an entire mission, soldiers often need to carry spares (or a charging system when rechargeable batteries are used).

A similar situation exists for a small or micro Unmanned Aerial System (UAS) where the battery needs to power the motor, controls, radio, and imaging equipment. As with the soldier, the battery must last the full duration of the mission. The allowable weight for the battery varies by the type of vehicle, of course, but is normally set to enable mission durations of between 30 minutes and two hours with existing types of batteries.

The solution to these weight and battery duration limitations is to use a lighter battery. By doubling the energy density, the weight of the battery pack needed for a mission of given duration can be cut in half. Alternatively, the same size and weight in battery pack(s) with double the energy density could double any mission's duration. This would be particularly valuable in remote reconnaissance and surveillance, in addition to target acquisition missions, where being aloft longer could make the difference between success and failure.

Doubling the energy density is, of course, far easier said than done, which is why this goal has remained elusive for the two most popular types of Commercial Off-the-Shelf batteries used in military applications today: Lithium/Sulfur Dioxide (Li/SO_2) and

Lithium/Manganese Dioxide (Li/MnO_2). While these batteries have similar energy densities of 200-250 Watt-hours/kilogram (Wh/kg), their volumetric energy densities are different at 350-450 Watt-hours/liter (Wh/l) and 500-650 Wh/l, respectively. Note that because energy densities vary with the different form factors used for different applications, the ranges used here depict typical values for cells in applications requiring moderate to high rates of amperage.

In applications where the weight is a significant design consideration, the similar gravimetric energy densities shared by sulfur dioxide and manganese dioxide batteries gives neither an advantage. But manganese dioxide batteries are increasingly preferred owing to other reasons, including their enhanced safety over pressurized sulfur dioxide batteries. And in applications where the space available for the battery is limited, the manganese dioxide battery has an even greater advantage with its 40 percent improvement in volumetric energy density.

In a BA-5X90/U battery pack, for example, manganese dioxide's increased volumetric energy density delivers about 11.5 Amp-hours (Ah) of service compared to about 7.5 Ah with sulfur dioxide. But the problem remains: the 11.5 Ah battery in this example is significantly heavier than the 7.5 Ah battery (at about 1.35 kg or 48 ounces compared to 1.0 kg or 36 ounces) because the gravimetric energy densities of both are similar. Table 1 compares and contrasts sulfur dioxide and manganese dioxide batteries, in addition to polycarbon monofluoride and carbon fluoride batteries, which will be discussed in the following section. (The table additionally shows Thionyl Chloride (Li/$SOCl_2$) batteries, also used in some military applications.)

Advanced carbon fluoride batteries

Lithium/Polycarbon Monofluoride or Li/$(CF)_n$ batteries have been around since the 1970s and were traditionally used for high-energy, low-power applications, such as memory backup systems. In the never-ending quest for a better battery, however, a new technology has emerged as an off-shoot of traditional Li/$(CF)_n$

Primary Lithium Battery Type	Manganese Dioxide Li/MnO_2	Sulfur Dioxide Li/SO_2	Thionyl Chloride Li/$SOCl_2$	Polycarbon Monofluoride Li/$(CF)_n$	Contour Fluorinetic Technology Li/CF_x
Gravimetric Energy Density (Wh/kg)	250-280	240-280	250-400	300-400	>600
Volumetric Energy Density (Wh/l)	500-650	350-450	600-900	600-800	700-1000
Operating Temperature Range (°C)	-20 to 60	-55 to 70	-55 to 150	-20 to 60	-40 to 160+
Typical Shelf Life (Years)	5-10	10	15-20	15	15
Safe (High-rate Discharge)	Yes	No	No	Yes	Yes
Environmental Impact	Moderate	High	High	Moderate	Moderate
Relative Price/Performance	Fair	Good	Fair	Poor	Good

Table 1 I A comparison of the four types of batteries covered in this discussion – sulfur dioxide, manganese dioxide, polycarbon monofluoride, and carbon fluoride – in addition to Lithium/Thionyl Chloride (Li/$SOCl_2$), which is also used in some military applications.

batteries: the Lithium/Carbon Fluoride (Li/CF_x) battery. Carbon fluoride batteries maintain the benefits of high energy and power densities, wide operating temperature range and long shelf life found in sulfur dioxide batteries, while employing a solid cathode (with no heavy metals or other toxic materials) to eliminate the safety and environmental concerns (see again Table 1). In addition, the carbon fluoride battery possesses none of the operational problems exhibited by some other batteries, such as passivation.

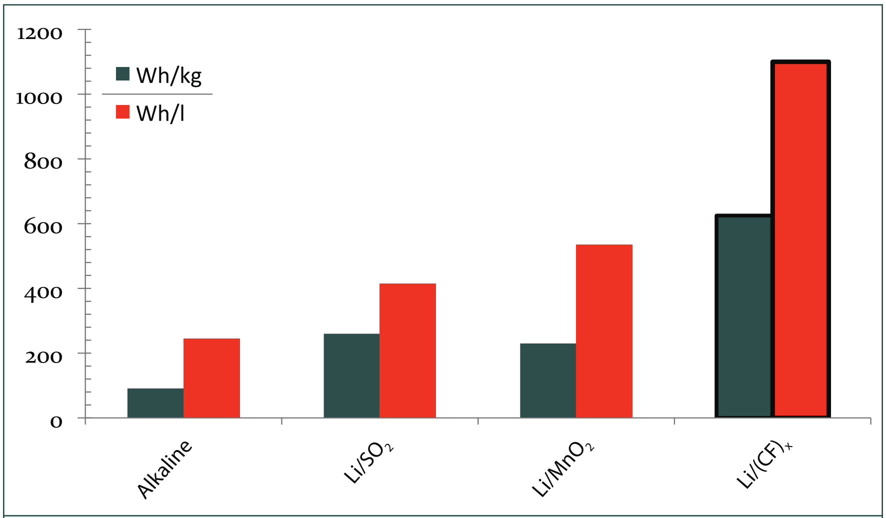

Figure 1 | This comparison of gravimetric and volumetric energy densities for different types of batteries (in addition to alkaline) demonstrates the advantage that carbon fluoride provides.

Most important is the carbon fluoride battery's higher gravimetric and volumetric energy densities of >600 Wh/kg and 700-1000 Wh/l, respectively, as compared to the other batteries discussed herein. As shown in Figure 1, the gravimetric and volumetric energy density improvements may be even greater in some configurations than the conservative estimates provided earlier.

In BA-5X90/U battery packs commonly used in military radios and other systems, a carbon fluoride version of this popular battery should be able to more than double the operating time of this battery – while weighing about the same as a traditional sulfur dioxide version of this product.

An additional major advantage of the carbon fluoride battery is its ability to exceed manganese dioxide and polycarbon monofluoride batteries, among others, in both power density and maximum safe current draw. The laboratory test results in Figure 2 demonstrated up to an 8x improvement in high-current

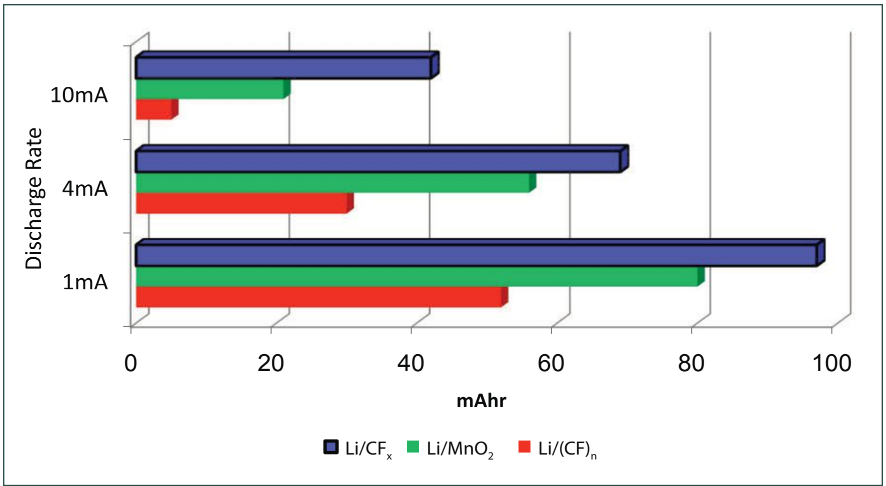

Figure 2 | These test results of available capacity at three different rates of discharge (to 2.0 V) for three different 2016 size coin cells quantify the carbon fluoride battery's improvements in power density at low, moderate, and high discharge rates.

applications and a nearly 2x improvement in low-current applications. This makes the carbon fluoride battery particularly well suited for applications that require high sustained or pulse currents.

Like the other lithium-based primary batteries, carbon fluoride batteries can be packaged in a variety of form factors, including coin, cell, film, or prismatic. This enables carbon fluoride batteries to accommodate both standard sizes and customized packs (which combine cells in series and/or in parallel to satisfy specific needs for operation in the typical military range of 6 to 30 V).

Additional "weighty" design considerations

How does the carbon fluoride battery stack up against the other battery types in other respects? As mentioned, Table 1 lists the details important in most applications. The use of only solid materials and a nontoxic electrolyte makes the carbon fluoride batteries safer than sulfur dioxide batteries, especially in those applications drawing a high, sustained current where sulfur dioxide batteries might overheat and fail. Solid materials eliminate the need for pressurized cans that can vent or leak corrosive or noxious gases, making carbon fluoride batteries safe even when mishandled or damaged or when subjected to a short-circuit condition. This is obviously of particular concern for the soldier, but even the batteries used in weapons and unmanned vehicles must still be handled during transport and replacement.

Operating temperature range is not a factor for the soldier, but can be for weapon and surveillance systems. And here, too, the carbon fluoride battery technology has made improvements over both manganese dioxide and sulfur dioxide. Indeed, the operating temperature range of carbon fluoride batteries such as Contour Energy Systems' advanced Fluorinetic batteries far exceeds the requirements of today's military applications.

Not only that, with their higher gravimetric and volumetric energy densities and other attributes described in this discussion, Contour Energy Systems' carbon fluoride Fluorinetic batteries provide a longer service life than both manganese dioxide and sulfur dioxide batteries. Just as significantly, these battery systems also afford a longer shelf life – up to 50 percent longer than either manganese dioxide or sulfur dioxide batteries.

Eric Lind is VP of Business Development at Contour Energy Systems. Before joining Contour Energy Systems, he was VP/General Manager of Ultralife Corporation's Commercial Business. Prior to that, Eric was the Director of Business Development in ConocoPhillips' Emerging Technology group. He started his career at Duracell, where he spent eight years in various engineering and management positions. Eric is a graduate of Duke University with a double major in Biomedical Engineering and Electrical Engineering, and holds an MBA from the University of Connecticut with a concentration in Marketing. He can be contacted at eric.lind@contourenergy.com.

Contour Energy Systems
626-620-0660
www.contourenergy.com

Power Product Showcase

NEW -70° to +150° C DC/DC POL Converters

Model #: MFP0507S

- Maximum Flexibility through a rich feature set; no external components required
- Output voltage adjustable from 0.64 to 3.5 VDC
- 4 pin-selectable standard output voltages
- Operates from a 3V to 6V input
- Capable of operating in parallel; Indefinite Short Circuit Protection
- Rad Hard options available

CRANE AEROSPACE & ELECTRONICS

Crane Aerospace & Electronics
425-882-3100
www.craneae.com

MilQor Mil-COTS DC/DC Power Converters & Filters

- High efficiency: >95% at full load
- Full power from -55 to +100°C
- Fixed switching frequency
- Input voltage ranges from 9V to 425V continuous
- Output power up to 600W
- No minimum load requirement; Made in USA

SynQor®

SynQor, Inc
978-849-0600
www.SynQor.com/ESM

MilQor Hi-Rel DC/DC Power Converters & Filters

- High efficiency: >90% at full load, Fixed switching frequency
- Full power from -55 to +125°C
- Parallel operation with current share
- Primary & Secondary referenced enable
- 28V and 270V nominal input voltages
- No minimum load requirement; Made in USA

SynQor®

SynQor, Inc
978-849-0600
www.SynQor.com/ESM

PFC & High Voltage DC-DC Modules

- Operating temperature range from -20 to 100 °C max baseplate for harsh environmental conditions
- Output power range from 75 - 1600 W (PFC); 65 - 600 W (DC-DC) with various output voltages
- Wide input range from 85 - 264 Vac (PFC); 250 - 420 Vdc PFC ready (DC-DC)
- DMTBF of up to 4.5 million hours for high reliability
- Vibration/shock acceleration tests (non-operating) 2.5 G/30 G

EMERSON. Network Power

Emerson Network Power
Embedded Power
760-930-4600
Emerson.com/EmbeddedPower

iMP Series Configurable Power Supplies

- Intelligent I²C control monitoring of voltage, current & temperature (PMBus/CAN/IEEE compliant)
- Six types of power modules – including single, dual & triple output units
- 25 standard output voltages from 2 - 60 Vdc
- Up to 1500 Watts & 21 separate outputs from one power supply
- Ruggedized for high shock and vibration; optional conformal coating
- Optional extended hold-up time modules

EMERSON. Network Power

Emerson Network Power
Embedded Power
760-930-4600
Emerson.com/EmbeddedPower

200W Hi-Rel DC/DC Converter

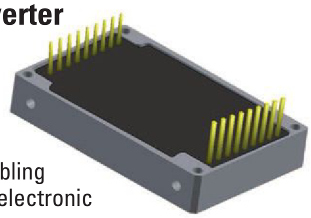

- Wide Input Voltage Range
- Redefining power converter heat dissipation
- Heat dissipation on the side enabling cooling to take place alongside electronic circuit boards
- The GAIA Converter 200W series has been designed to address power requirements in the range of 200W to 1KW in high reliability applications by offering high efficiency, up to 91%, and simple power sharing without extra components

GAÏA CONVERTER

Gaia Converter
1-888-601-3169 x202
www.gaia-converter.com

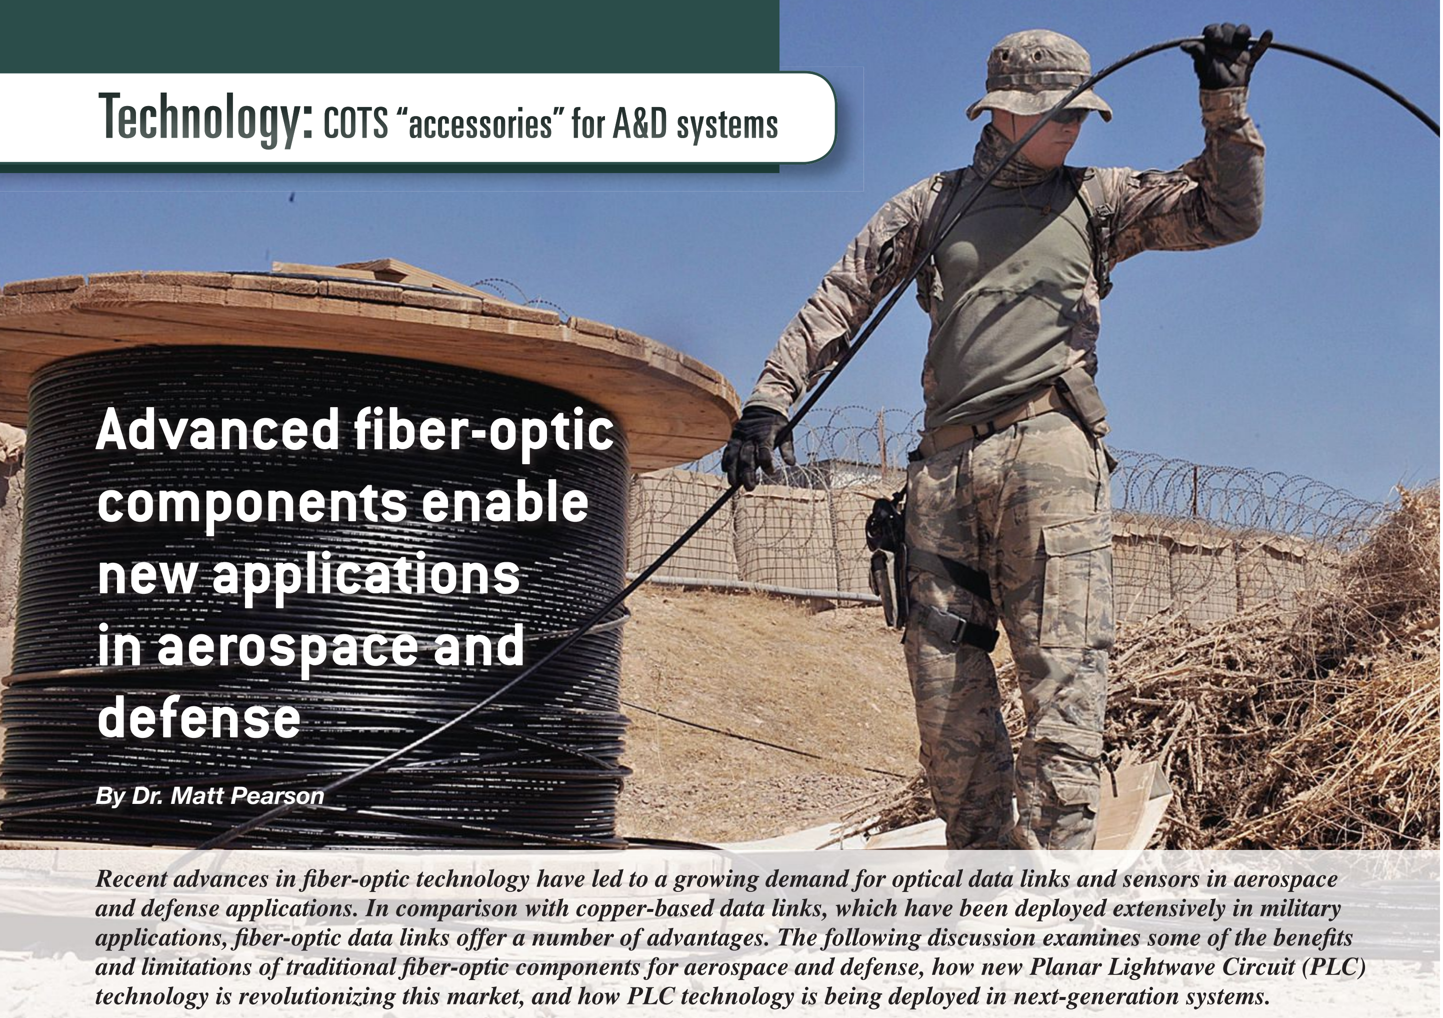

U.S. Air Force photo by Staff Sgt. Quinton Russ

Advanced fiber-optic components enable new applications in aerospace and defense

By Dr. Matt Pearson

Recent advances in fiber-optic technology have led to a growing demand for optical data links and sensors in aerospace and defense applications. In comparison with copper-based data links, which have been deployed extensively in military applications, fiber-optic data links offer a number of advantages. The following discussion examines some of the benefits and limitations of traditional fiber-optic components for aerospace and defense, how new Planar Lightwave Circuit (PLC) technology is revolutionizing this market, and how PLC technology is being deployed in next-generation systems.

The design of new military and aerospace vehicles imposes many restrictions on size and weight of components, while also requiring many more embedded sensors for monitoring the condition of the airframe and critical systems. The routing of these systems is significantly complicated by the available space and weight allowance inside the vehicle structure.

New materials have also affected the way modern vehicles are manufactured. Advanced composite materials are becoming commonplace in modern vehicles. While these materials can save a lot of weight, they are effectively transparent to electromagnetic energy like radio waves. These vehicles can therefore be more susceptible to electromagnetic interference than the more traditional metal structures that acted in part like a Faraday cage to partially protect the electronics inside.

Shipboard communication networks have also evolved dramatically in recent years. The communication infrastructure found on many ships is now more comparable to what previously might have serviced an entire city. While legacy systems might have run at 155 Mbps, new shipboard technologies can pack 10 or more channels on a single optical fiber, each channel supporting 10 Gbps. Fiber-based optical networks offer a number of significant advantages for defense and aerospace applications over conventional copper-based networks. It has been demonstrated that optical links have the capability to transport mixed digital and analog signals at speeds in excess of 10 Gbps, compared to the 100 Mbps that is more typical of copper connections. Along with higher data rates, optical fibers can also transmit light with losses as low as 0.2 dB/km, which is effectively lossless for most ship and aircraft applications, allowing power budgets to be kept low. Traditional fiber-optic components have some limitations in aerospace and defense applications, and the following discussion examines how Planar Lightwave Circuits (PLCs) are poised to displace incumbent technologies in many present and future applications.

Strengths and limitations of fiber-optics

Fiber-optics allow the use of Wavelength Division Multiplexing (WDM) technology, where many different channels can be transmitted down the same fiber by sending them on different colors (or wavelengths) of light. This can significantly reduce the size, weight, and complexity of cable routing.

Since only light signals are traveling along these optical fibers, they are effectively immune to electromagnetic interference. Moreover, optical fiber presents no spark or fire hazard, and can even be routed alongside or through fuel bays if required.

One of the most significant drawbacks of fiber-optics has been the reliance on traditional micro-optic assemblies used in most applications to date. Virtually all optical components to date have been built using miniature lenses, filters, and other optical components that are aligned by hand and glued in place. There are several potential issues with this approach.

First, the labor involved with aligning these micro-optic components is a very significant portion of the cost, which has resulted in nearly all this assembly being done overseas. Also, many piece parts epoxied together presents a risk of a subcomponent being knocked out of alignment in harsh environments, high vibration, or mechanical shock. This is particularly important in aerospace applications where vibration or shock up to 20 *g* acceleration is not uncommon. Finally, the amount of functionality that can be included in a single optical module is significantly limited by the sheer number of subcomponents that can be reliably aligned in a single package.

Planar Lightwave Circuit or *PLC* optical-chip technology addresses all of these shortcomings. This PLC platform collapses many of the same optical functions onto a silicon chip. Custom optical circuits can be designed and developed for virtually any application. These optical chips look visually similar to their electronic chip counterparts, as shown in Figure 1, and are in fact fabricated using very similar processes. However, along with the typical electrical connections found on any silicon chip, these PLC chips also have optical connections that receive or transmit signals optically into and out of the PLC chip.

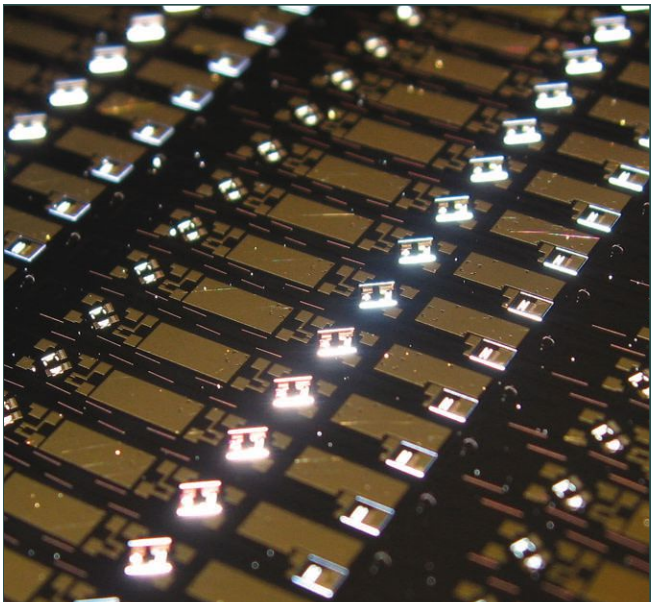

Figure 1 | Similar to their electronic counterparts, optical chips are fabricated on silicon wafers, making them small, lightweight, and rugged.

This optical chip is housed in a small package designed for soldering onto a common printed circuit board. The result is a very compact, lightweight optical module that is effectively a single chip when assembled, minimizing any chance of parts becoming misaligned during high vibration or shock. A complete bidirectional transceiver capable of transmitting and receiving data at 2.5 Gbps can be integrated on a chip that is smaller than 4 x 10 mm, weighing less than 5 grams. These chip-based transceivers can withstand harsh g-forces, including sinusoidal vibrations with amplitude of 20 *g* acceleration, followed by mechanical shock of 500 *g* acceleration. The components are capable of temperature cycling through a range of more than 125 °C.

PLCs deliver functionality not possible in other platforms

While virtually any optical function can be integrated on the silicon chip, in some cases a hybrid approach is advantageous, where small chips are also integrated onto the PLC platform. For example, very high-performance lasers can be fabricated in indium phosphide and integrated onto the PLC as part of a hybrid approach. This allows the optical chip to provide state-of-the-art performance, leveraging different material systems as needed.

All of this hybrid integration can be accomplished using robotic pick-and-place machines typically used for electronic assembly. These customized machines can bond a new chip every 20 seconds, with extremely good repeatability, making them ideally suited for volume production. Equally important, this approach requires very minimal hands-on labor, allowing these products to be manufactured domestically.

The photonic integration technology allows the development of ultra compact and lightweight communication components such as optical transceivers. Optical transceivers are essential components in fiber-optic communication systems. They contain both a transmitter and a receiver in a single housing, and can carry upstream and downstream mixed digital and analog signals through a single fiber.

At present, the optical transceivers designed and produced for aerospace and defense applications are built using traditional micro-optic technology. The complexity of assembling a large number of critically aligned micro-optic components in a single hermetic package makes these transceivers expensive and susceptible to failure in harsh military applications.

Planar Lightwave Circuit technologies have the capability to deliver outstanding improvements in the reliability and cost of optical transceivers, and represent a major step forward in aerospace and defense communication systems.

Early applications in aerospace and defense

As an example, this rugged PLC platform is ideally suited for the development of multichannel redundant transceivers. Photonic technology originally developed for the telecommunication industry has been leveraged to combine multiple signals at different wavelengths onto a common fiber, using proprietary filtering technology that is integrated into the PLC chip, requiring no external filters or lenses. If one transceiver fails, the others still provide support for critical avionics systems. The transceiver shown in Figure 2 is capable of 2.5 Gbps transmission, while simultaneously receiving downstream traffic at 2.5 Gbps, while also receiving analog video signals on a third channel.

This transceiver is packaged in a very compact housing, and designed for easy mounting and soldering onto a standard printed circuit board. This transceiver module was originally developed for use in Unmanned Aerial Vehicles (UAVs), although it has since been customized for a number of other applications.

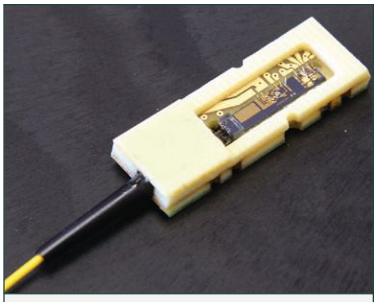

Figure 2 | This compact, lightweight, bidirectional optical transceiver can transmit and receive simultaneously at 2.5 Gbps, while also receiving analog video signals at up to 1 GHz on a third channel.

While the transceiver shown in Figure 2 supports three channels each running at up to 2.5 Gbps, some applications require even more capacity. The very flexible PLC platform is fully scalable in both channel count and data rate. Enablence Technologies manufactures PLC-based multichannel receiver modules that include an 8-channel wavelength division multiplexer, in addition to 8 receiver channels, each running over 10 Gbps, all integrated on a silicon platform. This receiver module, along with a matching transmitter module, is designed to form the basis of a shipboard communications network. Even more of the manufacturing efforts at Enablence Technologies are focused on custom modules unique to specific applications.

The future of optical components

Some basic examples of optical modules that can be developed based on Planar Lightwave Circuits have been examined herein. It is not uncommon for these modules to be significantly more advanced, often containing dozens or even hundreds of optical elements on a single chip. This level of functionality is typically not even feasible using more traditional optical platforms, and is opening up new applications for fiber-optics in aerospace and defense applications, including communications, sensors, and monitoring. When reviewing some of the strengths and limitations of traditional fiber-optics, it becomes clear how PLCs are poised to displace these technologies in present and future defense applications.

Dr. Matt Pearson presently serves as the Vice President of Technology at Enablence Technologies Inc. He has extensive experience in the design, characterization, and manufacturing of highly integrated photonic circuits. Matt holds a B.Sc. from the Royal Military College of Canada and a Ph.D. in Optoelectronics from McMaster University. He can be contacted at matt.pearson@enablence.com.

Enablence Technologies Inc.
613-270-7851
www.enablence.com

» How can I reduce my time-to-deployment? «

Kontron provides a unique blend of COTS and specialized systems, custom-tailored to meet the most demanding requirements of our customers.

» Kontron's expertise goes beyond compliance, providing systems that fully address the complex needs of customers with application specific hardware and software elements.

» With a long track-record of engineering excellence, Kontron is able to mitigate risk and reduce time-to-deployment.

» Experienced program management allows Kontron engineers to work seamlessly within complex programs and multi-partner development scenarios.

CRITICAL QUESTIONS ... ANSWERED

FS-5985

1/2 ATR Conduction Cooled Rugged Chassis - cPCI, VPX

» Suitable for MIL-STD-5400 Class I environments

» 5 Conduction Cooled 3U slots to IEEE 1101.2, 0.8" pitch

Cobalt™

Reduced SWAP - High Performance Embedded Computer

» Small Form Factor (6.5" x 9.725" 2.95")

» Intel® Core™2 Duo or Atom™ processor options

» Custom I/O option

VX3020

High Performance 3U VPX Rugged PC

» Implements VITA 46.0, 46.4, 46.7, 46.9

» 1.5 GHz Dual Core® Core™2 Duo processor

CONTACT US

Call, Email or Visit today.

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

If it's embedded, it's Kontron.

Case study: Designing for rugged thermal management to meet military needs

By Dave Turner

A Special Operations Craft Riverine (SOC-R) required a rugged shipboard computer capable of operating in an ambient temperature of up to +71 °C per the requirements of MIL-STD-810G.

As military programs continue to push the limits of computing system requirements, rugged computer systems must continue to evolve to endure the toughest conditions: extreme temperatures, shock, vibration, and humidity amidst altitude, fungus, salt fog, explosive decompression, immersion, and sand/dust exposure. To meet these stringent requirements, rugged subsystems suppliers must produce innovative thermal management technologies that can withstand situations that may threaten a computer's durability. A case study of a U.S. Navy Special Ops river craft is presented to illustrate the inner-workings of a thermal management scheme.

Thermal management remains a major component of any rugged electronic design, as heat issues are often the largest contributors to failures. Consequently, advances in thermal management will continue to rank as one of the most important trends in rugged computing design because processors continue to change along with the demand for higher processing capabilities, including multicore computing at speeds of 1.5 GHz or greater.

While many applications can still be served by energy-efficient single-core processors, new technology developments are demanding higher processing power with low power consumption. Ultra-low power Intel Atom processors cannot satisfy performance requirements in all cases. Consequently, multicore processing technology is seeing a significant boost in deployment within stand-alone rugged boxes. For demanding applications, mobile Core 2 Duo or Core i7 processors, as examples, provide attractive solutions. While presenting a challenge to manage the 10-55 W of Thermal Design Power (TDP) of these processors, they offer a level of performance that pushes the boundaries between traditional rugged computing and new tactical applications. Meanwhile, the following case study examines the thermal management of a multicore mission computer utilizing an Intel Core 2 Duo CPU.

Navy application prompts new thermal management design

Parvus engineers were tasked with engineering a passively cooled solution that met the increased thermal demands of a multicore processor, upon receipt of a contract to design a specialized version of the DuraCOR 810-Duo mission computer for a U.S. Navy special operations river craft (Figure 1). This specific military customer required a shipboard mission computer as part of its command, control, and communications system designed to operate in an ambient operating temperature of up to +71 °C per the requirements of MIL-STD-810G.

Unlike Parvus' legacy-generation mission computer based on a Pentium M processor, this new mission computer integrated an Intel Core 2 Duo CPU – which consumed about 60 percent more power and generated significantly more heat. The challenge for engineers was to design a rugged mission computer with integrated application-specific payload cards capable of cooling a multicore processor in an environment that can only support cooling by natural convection and possibly operate in a stagnant air environment.

Heat pipe not enough for increased thermal demands

To manage the thermal increase in the new mission computer, engineers initially designed the internal subsystem with an integrated heat pipe: a thermal management device that transfers heat by the evaporation and condensation of an internal fluid. Heat pipes didn't rely on any moving parts to dissipate heat from the system, plus they had been used successfully in the past for other single-core CPU system designs. Heat pipes are also well known for how quickly they can wick heat away from system hot spots. The heat pipe contained a copper outer layer with a hollow center that contained a wicking mechanism. The center was filled with a liquid capable of vaporizing, such as

Figure 1 | The DuraCOR 810-Duo is a rugged multicore mission-processor subsystem designed for high-reliability applications requiring MIL-STD-810G environmental compliance.

water, alcohol, ammonia, or methanol. When heat hit the pipe, the pipe converted the liquid to a vapor. The vapor escaped down the pipe and came in contact with the cool side of the heat sink. The vapor then cooled and was condensed back to a liquid where it was absorbed by the wick. This transferred the liquid back down the heat pipe to start the exchange over again (see Figure 2).

Engineers soon learned that a heat pipe was not going to suffice for this system, as preliminary qualification tests concluded that the computer server was experiencing difficulty when performing system stress testing beyond +65 °C ambient air temperature before the CPU's die would hit Intel's rated thresholds. The internal temperature of the heat pipe was exceeding its design limits and the core temperature of the computer system enclosure was approaching +90°C. The heat pipe's increased temperature would not allow the vapor to cool enough to convert back to a liquid – an event referred to as "saturation." This heat pipe saturation effectively eliminated continued vapor-to-liquid conversion, halting convection to the exterior of the enclosure. The only effective method of heat transfer was the conduction of the copper pipe from the heat-producing processor to the exterior heat sink interface. This condition reduced the possibilities of the system operating past +65 °C ambient air for extended periods of time.

Structural heat-spreader plate offers improvement

Through initial testing, the suggested engineering answer to this challenge was to introduce a conduction-cooling mechanism through a structural heat-spreader plate. Rather than channeling the heat from the single board computer through a copper pipe and converting the liquid to a vapor, engineers designed an aluminum heat spreader to transfer heat to the exterior chassis heat sink. When the aluminum is exposed to a heat source, the heat is conducted through the metal. Unlike the heat-pipe approach where vapors move very quickly through the pipe, the aluminum takes a slightly longer time to conduct the heat. While the design may suffer slightly in the initial heating cycle, the solid metal "spreader" offers the advantage of not losing its effectiveness by "saturating."

When the heat reaches the exterior heat-spreader plate, it interfaces with a thermally conductive material, which then transfers the heat to the chassis heat sink. With a high-finned protruding design, these extruded heat sinks increase the surface area of an enclosure housing, transferring potentially damaging heat away from a system and preventing performance degradation. By designing a direct conduction method from the internal heat-producing components to the natural convecting exterior of the enclosure, the heat can continue to move to the ambient air without the risk of saturation, allowing the system to operate in higher temperatures as required by contract: to +71˚C.

Though the DuraCOR 810-Duo described was designed for a specific application, other military customers are interested

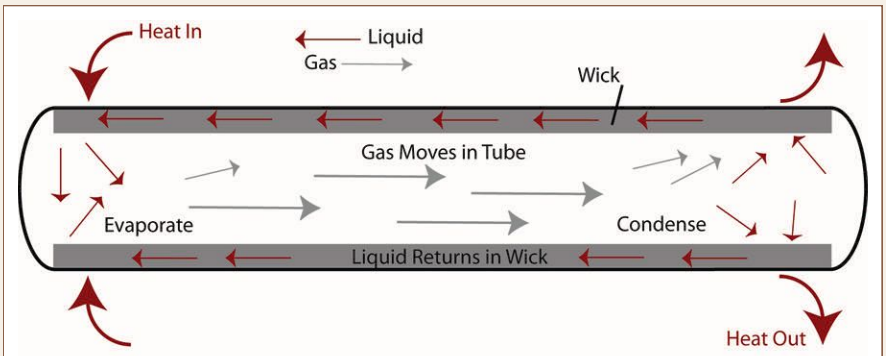

Figure 2 | The heat pipe thermal management solution transferred heat by the evaporation and condensation of an internal fluid.

in deploying similar solutions in ground, naval, and airborne vehicles. Accordingly, Parvus intends to apply its lessons learned and similar thermal management analysis to address varying program requirements. Because at 40,000 feet (12,192 meters), roughly one-fifth of the atmospheric pressure (2.7 PSI) is available for convective cooling compared to sea level (14.7 PSI), the COTS dual-core subsystem will be fitted with an optimized heat-spreader plate as the default configuration to maximize application robustness.

Continuous innovation needed for rugged design

Engineers always need to be conscious of which issues are preventing a computing system from reaching its ability to meet a customer's needs. Gaining a better understanding of which combination of thermal products and techniques help transfer heat while maintaining cost, weight, and system integrity will prove to be one of the most important elements in rugged computing design. This influences chassis and internal package design to accommodate hardware and system design features to manage the heat while still maintaining the unit's physical integrity and size and weight constraints. Consequently, the ability to analyze requirements quickly and identify potential solutions remains a key skill for the design engineer. As illustrated in the aforementioned case study, thermal engineers should be particularly cognizant of saturation of a heat-pipe design at high temperatures. Not all solutions are the answers to every design problem. Heat pipes are an excellent choice for cooling electronics and will always be one of the top choices for these high-powered applications. Diligent study of every possible alternative is always part of a design engineer's plan to succeed. Being aware of a component's limitations adds to that success.

Additionally, engineers should remain current on the best rugged design practices by continuing to expose themselves to several areas of design. By gaining a better understanding of mechanical engineering practices, electronic engineers in particular can better learn how to solve electronics design challenges, especially when dealing with rugged design requirements. The same can be said for mechanical engineers as they are tasked with the rugged packaging of electronics. As customers continue to push the limits of computing system requirements, rugged computer subsystems will continue to evolve to endure the toughest conditions, as long as good engineering practices are implemented and maintained.

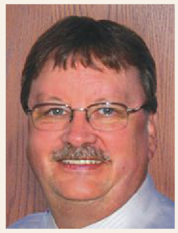

Dave Turner *is a Senior Mechanical Engineer at Parvus Corporation, based in Salt Lake City, Utah. He has more than 25 years of engineering experience, including development programs of the armed forces in areas such as B-1, B-2, F-117, F-14, A-7, F-22, space environments, missiles, ground vehicles, and shipboard equipment. Dave can be reached at dturner@parvus.com.*

Parvus Corporation
801-483-1533 • www.parvus.com

Editor's Choice Products

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

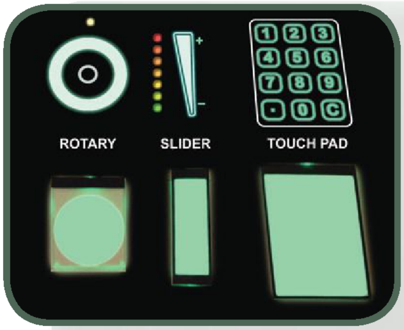

More light ... without the extra light

Like sensor/weapon controls operators don't already have enough on their mind – *and* they also have to remember to bring that extra flashlight so they can use their systems' consoles once nighttime arrives or a room darkens. That could be one reason why Global Lighting Technologies (GLT) developed its ultra-thin light guides, which backlight touch-enabled or membrane switches for myriad applications. Leveraging LED edge-lighting technology, the ultra-thin light guides render illumination of on/off buttons, directional systems, touch pads, rotary switches, logos, sliders, and other large or small graphic icons.

Working to help relieve operator errors and frustrations, the light guides' "extreme thinness" is key to their effectiveness. They are slipped in between the graphic overlay and circuit sensor or PC board and *voila*! And since precision is key in most apps – and particularly the mission-critical – the light guides' insertion will not lower application touch sensitivity. So maybe it's time to kick that extra flashlight to the curb.

Global Lighting Technologies (GLT) • www.glthome.com • www.mil-embedded.com/p46824

Catch some rays ... but don't lose your memory

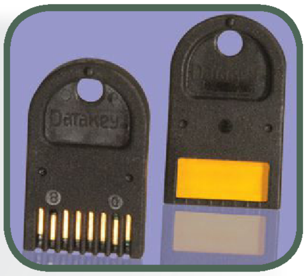

When planning for application and system security, military embedded systems designers have to think "out of the box" about what will go *in* the box. Hence, with many defense systems subject to Nuclear, Biological, and Chemical (NBC) requirements, even if humans are put at risk, it's often imperative that the equipment survive intact. And Datakey Electronics' GammaSafe portable, reprogrammable, nonvolatile memory token enables specialized military or medical technology developers to fit devices with limit-use and anti-counterfeit capabilities to either limited-use disposable or single-use attachments exposed to gamma radiation.

Accordingly, gamma radiation used in sterilizing medical devices, for example, often causes device failure and data loss in conventional technologies such as solid-state memory wares; however, the GammaSafe memory token can withstand up to 45 kGy (4.5 Mrad) gamma radiation – with all data remaining intact. The SGT4Kb GammaSafe flavor comprises 4 Kb memory accessed utilizing an SPI bus, and its functionality resembles an SPI EEPROM. And to make things more convenient, the memory token features a triad of receptacle choices: one board-mount and two panel-mount variants.

Datakey Electronics • www.datakeyelectronics.com • www.mil-embedded.com/p46825

This board's not so hot

When dealing with challenging thermal management issues, many technicians would undoubtedly like to change the axiom "If you can't stand the heat, get out of the kitchen" to "If you can't stand the heat, get out of my system!" (as if the board would understand such instructions ...). The good news is that Laird Technologies, Inc. developed its Tflex XS400 Series thermal pad to aid in eliminating such thermal management frustrations. The pad provides conductivity between a board's hot components and heat sinks at a moderately rugged level, effective in relatively benign venues such as a command and control center, for example.

Providing 2.0 W/mK thermal conductivity, the naturally tacky Tflex XS400 Series is easy to apply, voiding the need for adhesive coating. It is vended in 0.010 thickness increments ranging from 0.020" (0.50 mm) all the way to 0.200" (5.0 mm). Stable at -40 °C to + 160 °C and RoHS compliant, the gap filler provides electrical isolation and is certified to UL 94V0 for flammability. And that's not all: The thermal pad's soft texture requires only minimal pressure when conforming to the board's design, thus thermal resistance and stress on mating parts are reduced or even eliminated.

Laird Technologies, Inc. • www.lairdtech.com • www.mil-embedded.com/p46826

PC/104 saves the day, melds a trio

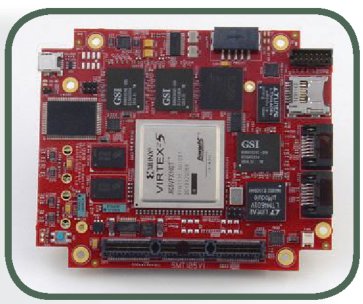

Sundance is well-known for its high-performance signal processing wares and beefy algorithms, and its tools are an industry standard. Now the company aims to make the PC/104 open standard even more of a household name, thanks to its Sundance Series 1 boards melding FPGAs, PowerPC440, and PCI Express on the small form factor. Highly suited for the image processing, high-speed data acquisition, and communication arenas, the series employs Xilinx Virtex-5 FPGAs for its two versions: The SMT100 featuring 512 MB DDR2 SDRAM (handy when it comes to data processing and buffering), in addition to the SMT105, offering a mixture of memory banks in the DDR2 SDRAM and QDR2 SRAM camps to facilitate signal processing in real time.

Then there's the plethora of I/O that's packed onto this small form factor board: MicroSD/Transflash, USB 2.0 FPGA interface, fiber-optic modules, and RS-232 interfaces. And it doesn't end there. Also included: 32-bit PCI, LVTTL I/O, SATA ports, and 1- and 4-lane PCI Express links in addition to stack-down or stack-up flavors that afford automatic detection. Many different IP cores are also vended to speed time to market and foster Series 1 design integration.

Sundance • www.sundance.com • www.mil-embedded.com/p46827

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

Crosshairs Editorial

A quiet revolution against RTL

Xilinx DSP kits develop into more than reference designs

By Chris A. Ciufo, Editor

FPGAs have gotten so large and complex that even veteran silicon designers tremble at the prospect of figuring out efficient routing to utilize all the Look Up Tables (LUTs), logic elements, routing interconnects, DSP slices, and gigabit transceivers. This is despite having the best tools money can buy from Cadence, Synopsys, or Mentor. And with Xilinx's recent announcement of stacked-chip interconnects leading up to an *unheard-of 2 million gate* Virtex-7 FPGA in 2011, things are gonna get tougher. Only hardcore FPGA designers will use the damn things, and they'll be the guys who absolutely need million-point FFTs, wacky OFDM algorithms in LTE cellular base stations, or AESA radar processors.

It doesn't have to be this way. It'll only be a tough design flow if you stick with the traditional semiconductor design route using HDL and RTL. But other options exist, including model-based design a la MathWorks, and coding in high-level languages such as C/C++. These are not so popular today, though they're certainly viable, and Xilinx is about to give non-RTL alternatives a swift kick to the front of the line. I believe the company's new *DSP Targeted Design Platform* development kits are going to be the catalysts to usher in a move to non-HDL FPGA designs (www.xilinx.com/technology/dsp.htm). After all, who the heck wants to be an IC designer when what they really want is to design *systems*?

For a couple of years, Xilinx has partnered with companies like Avnet to make available hardware reference boards called *Targeted Design Platforms* for things like image tracking and signal processing. They're doing it again, with three new ones just announced in December:

- Virtex-6 FPGA Design Kit with Integrated AD/DA
- Spartan-6 FPGA DSP Kit
- Spartan-6 FPGA/OMAP Co-processing Kit

Each is a stand-alone board complete with demos, drivers, appropriate software, I/O modules (where applicable), extensive documentation, and even royalty-free reference design docs (gerbers and schematics) so customers can replicate the design – or modify it for their own system. I could go on and on about the coolness of each of these boards and cite their technical specs, but that's not what is revolutionary about these kits. In fact, they're really quite *evolutionary* by looking at the specs. Two of them use FMC for I/O – one with an impressive card by partner 4DSP boasting a 250 MSps A/D into a DDC, routed back through the Virtex-6 logic to an 800 MSps DUC into a D/A. This particular kit is intended for gangbuster performance in systems needing up to 1,000 GMACs, implemented by seasoned FPGA designers doing "bare metal" RTL coding with optimized and probably existing legacy DSP algorithms.

"It doesn't have to be this way. It'll only be a tough design flow if you stick with the traditional semiconductor design route using HDL and RTL."

"So what," you say, "it's evolutionary" … except for the fact that Xilinx also includes the ability to code this board in C/C++ using software from vendors AutoESL or Synopsys. And they've also included interfaces to a SysGen flow that bolts onto MathWorks' Simulink model-based environment. Using Xilinx- and BDTi-generated benchmarks with mildly optimized code, the performance of this particular Virtex-6 kit is almost as good with the high-level tools as it is with RTL, and the logic utilization is also nearly as good. And to boot, changes to the FPGA design can be done in modular code blocks (in C/C++ or Simulink) and realized within *hours* as opposed to *days* when done in RTL. Simplicity and time savings are huge advantages that will save designers and their employers some real coin.

But there's more. The Virtex-6 kit is meant for FPGA designers already doing DSP. The other two kits are geared for either a non-FPGA DSP designer thinking about moving to FPGAs (the Spartan-6 DSP kit) or for processor designers (using an ARM-based TI OMAP L-138) out of horsepower who must add an FPGA to meet performance requirements. In both of these cases, Xilinx demonstrates how easy and painless it is to add an FPGA to achieve 3x, 10x, and even more performance with a mere $10 Spartan-6 device. And since these two kits are optimized for C/C++ and Simulink flows, the barrier to entry for RTL-averse system designers is practically nil. Way cool, and *revolutionary*.

Xilinx's quiet introduction of the kits is preparing the world to stick FPGAs in all designs that need more performance or headroom with minimal cost. In effect, they're set to dramatically expand their market to non-FPGA users by making it easy to choose and program an FPGA. With high-level language or model-based tools, who needs to worry about RTL? Let the IC designers fret about that. The rest of us just want to design systems. Now that's a revolution in the making.

Chris A. Ciufo, Editor
cciufo@opensystemsmedia.com

GE
Intelligent Platforms

You can only target what you can track.

As the deployment of unmanned vehicles has increased, so has the demand for high-performance video tracking capabilities. In response, GE Intelligent Platforms has developed a full range of cutting-edge, real-time COTS video tracking, compression and image processing technology that is ruggedized and low power.

We support simultaneous compression of multiple channels of high-resolution video in rugged environments. Simply connect your video inputs and begin streaming HD at full frame rates for superb tracking performance and reduced operator workload. No additional integration or software development is required.

For more information, visit:
www.ge-ip.com/image-processing

imagination at work