INCLUDING DEFENSE TECH WIRE

Military EMBEDDED SYSTEMS

March 2013
Volume 9 | Number 2

MIL-EMBEDDED.COM

Leveraging commercial avionics for military platforms

DO-332, the Liskov Substitution Principle, and local type consistency ramp up DO-178 certification
By Dr. Benjamin M. Brosgol, AdaCore

FACE, ARINC and other avionics standards enable software reuse

www.opensystemsmedia.com

A surprising blend of personalities

That's the nature of the AB2000 Avionics BusBox®

With so many built-in capabilities, the highly-flexible AB2000 family from Ballard Technology is the ideal fit for a wide range of applications. These rugged, conduction-cooled, COTS devices combine a powerful computer processor, multi-protocol databus interfaces, Ethernet, USB, other I/O, and PMC expansion in a small, lightweight package.

Discover what the AB2000 can do for you.

Visit our website or call 425.339.0281 for more information.

AB2000 Personalities...

- Embedded Computer
- Aircraft Interface Device
- Data Recorder
- Protocol Converter
- Ethernet Switch
- Power Management Unit
- Radio/Satcom Interface
- and many more

www.ballardtech.com/AB2000

Ballard TECHNOLOGY

AN ASTRONICS COMPANY

AS9100 / ISO 9001 Registered

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

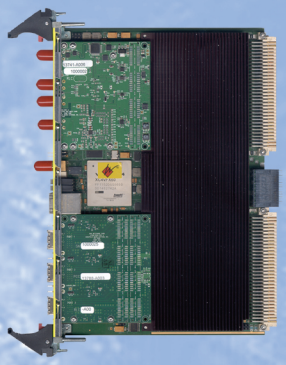

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

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

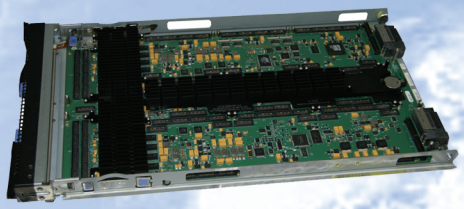

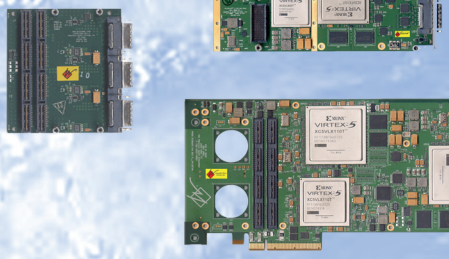

Annapolis
Micro Systems

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

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

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

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

www.mil-embedded.com

Military EMBEDDED SYSTEMS

ON THE COVER:

Top photo: A Pro Line Fusion Synthetic Vision system developed by Rockwell Collins for the Bombardier Global 5000 business jet's flight deck has been adapted for the MH-60 Black Hawk helicopter. (Photo courtesy of Rockwell Collins)

Bottom photo: Military avionics programs are taking advantage of common standards such as FACE, ARINC, and DO-178C to enable portability and affordability of avionics software across multiple platforms. Pictured: C130-J being delivered by the Lockheed Martin Aeronautics Company, Marietta, GA (Lockheed Martin Photography by Thinh D. Nguyen)

March 2013 | Volume 9 | Number 2

COLUMNS

DEPARTMENTS

E-CAST

http://ecast.opensystemsmedia.com

Managing safety and security through open architectures

March 26th • 11am MST

Presented by: Esterel, GE Intelligent Platforms, RTI, 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

Published by: OpenSystems media®

ISSN: Print 1557-3222

@military_cots

www.linkedin.com/groups/Military-Embedded-Systems-1864255

VPXpert

WE STILL HOLD THE TITLE.

VPX Solutions: X-ES Is The VPXpert

Extreme Engineering Solutions (X-ES) has the broadest range of proven VPX solutions available for rugged embedded computing applications. From single board computers, switches, storage, FPGA processing, and I/O to fully qualified integrated systems, look to X-ES for industry leading VPX solutions. VPX boards and systems from X-ES are 100% designed, developed, manufactured, and supported in the U.S. to deliver extreme performance in extreme environments.

Come discover why X-ES is *the* VPXpert. Call or visit our website today.

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

ADVERTISER INFORMATION

Page	Advertiser/Ad Title
39	**ACCES I/O Products, Inc.** – USB embedded I/O solutions
21	**Aitech Defense Systems** – Our technology investments protect yours
43	**AMP Inc. Accelerated Memory Production** – High durability standards for the DDR3 ruggedized SODIMM
3	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
2	**Ballard Technology** – A surprising blend of personalities
9	**D-TA Systems** – From DC to daylight ... scan, detect, process, record and playback RF signals at will
34	**Elma Electronic – Systems** – Small is the new big
40	**EMAC, Inc.** – SoM 9X25
24	**Engineering Design Team (EDT) Inc., a HEICO Company** – PCIe dual IF receiver
19	**Excalibur Systems, Inc.** – Convert – your avionic data of course!
47	**GE Intelligent Platforms, Inc.** – Leadership and experience in avionics bus protocols
18	**Innovative Integration** – Faster than a bat out of ...
23	**Interface Concept** – Switches and IP routers
46	**Kontron** – Highest data throughput for your application
27	**North Atlantic Industries** – Connect with us
48	**Pentek, Inc.** – Critical recording in any arena
24	**Phoenix International** – RPC24 – rugged, deployable, mission oriented data storage
31	**Proto Labs** – Real parts. Really fast.
40	**Radisys Corporation** – Radisys ruggedized CEQM77HDE COM Express module
13	**SIE Computing Solutions, Inc.** – Rugged and ready when you are
7	**SynQor** – Uninterruptible power supply
35	**Themis Computer** – Take your virtual environment to the edge
44	**Themis Computer** – High-density computing for demanding environments
41	**WinSystems, Inc.** – A 1.66 GHz Intel Atom Mil/COTS SBC of EPIC proportions
5	**X-ES** – VPXpert – we still hold the title
20	**Z Microsystems, Inc.** – Experience unprecedented clarity

Military EMBEDDED SYSTEMS

OpenSystems media®

SynQor

Uninterruptible Power Supply

Sealed

Rugged

Smaller

Lighter

More Powerful

- **Power 1U: 1250W/1500VA**
- **Weather-proof, shock-proof construction**
- **True on-line double conversion**
- **1U high rack mount (17" x 21.6" x 1.73")**
- **Low weight — 32 pounds**
- **Dual input (AC and DC)**
- **Battery run time >10 minutes @ full load**
- **Removable battery pack**
- **Full power operation: -20°C to +55°C**

Made in the United States of America.

1-978-849-0600 www.SynQor.com

SynQor®

Military not deploying to trade shows

By John McHale, Editorial Director

Do you like crowded convention centers buzzing with tales of new technology and the smell of money being made? If so, you should avoid military conferences/trade shows in the United States. Long lines and packed exhibition aisles don't exist at these events.

The past six months have been particularly hard on military conference and event planners. AFCEA's MILCOM 2012 was canceled due to Hurricane Sandy shutting down many East Coast airports, preventing a large number of attendees from traveling to Orlando to attend the event. Other event experiences such as AUSA meetings and AFCEA's West conference in San Diego experienced last-minute cancellations as the U.S. military services cut back on travel due to budget cuts, anticipation of sequestration, and some say ethical perceptions.

The Association of the U.S. Army (AUSA) Winter Symposium's last visit to Fort Lauderdale went out with a whimper rather than a roar last month. The annual gathering near the Florida beaches had a much smaller footprint this year. Show management said this year the outside exhibits, those typically featuring Army land vehicles and helicopters, were eliminated due to the decrease in business. (The Army only announced support for the show a couple weeks prior to its start.) While the symposium only attracted 1,400 attendees this year, event officials said they were pleased and said the symposium will happen again next year – just not in Fort Lauderdale. A 2014 location had not yet been selected at the writing of this column, but AUSA officials said they are looking at Raleigh, NC – to be near Ft. Bragg – or military facilities and contractors' locales in Huntsville, AL. After 2014 they would most likely start a rotation including those two locations and one or two more, according to show officials.

During my previous trips to the Winter Symposium, I found it buzzing with new Army technology with large exhibits and displays from prime contractors such as Lockheed Martin and General Dynamics. You really can't blame the primes for not going since their customer – the Army – only committed to the show a couple weeks out. Those who did attend apparently were flag officers and staff. A similar atmosphere surrounded the AUSA Annual Meeting, held every October in Washington – typically one of the largest military trade shows/conferences in the U.S. The 2012 event actually shrunk for the first time I can remember.

> *"The U.S. events that still create a 'have to be there' buzz are those related to Special Forces and unmanned systems such as SOFIC and AUVSI's Unmanned Systems respectively."*

Many in government and industry credit the travel decline to the DoD's budget constraints and sequestration fears. However, there also is a perception that events held in places like Sin City, Fort Lauderdale, and other swanky locales are boondoggles. Much of this perception can be traced to the General Services Administration (GSA) scandal in Las Vegas where agency officials reportedly spent nearly a million dollars on an event. However, I fear it will only get worse. Just this morning I read FoxNews.com coverage on a House of Representatives Oversight and Government Reform Committee report that said the DoD spent nearly $89 million on 300 conferences in 2012. No way that gets a positive reception from the public.

The only U.S. events that still seem to create a "have to be there" buzz are those related to Special Forces and unmanned systems such as the Special Operations Forces Industry Conference (SOFIC) and AUVSI's Unmanned Systems located in Washington. This makes sense as both those application areas will get the bulk of whatever funding is left after budget cuts. I hope this downturn in attendance doesn't cripple the military trade-show circuit long-term. Many of the events provide education and learning opportunities to military/government attendees as well as chances to meet with industry. Only now the events may have to take on a more modest posture to enable military personnel to get approval to attend. Fewer events will be aligned with warm beaches, Disney World, and the Vegas strip. They will be more strategically located near key military bases.

One group that is generating a lot of buzz at their events – which are located at military bases – is the Future Airborne Capability Environment (FACE) Consortium that is working on a standard to enable software to be reused across multiple avionics platforms. The next event is April 2 at Wright-Patterson Air Force Base (www.opengroup.org/face). For more on FACE and avionics trends, see the Special Report section beginning on page 16 and the Mil Tech Trends section beginning on page 22. Also in this issue, Managing Editor Sharon Hess has redesigned her Defense Tech Wire column to better reflect content from our website – www.mil-embedded.com (see pages 14-15).

John McHale
jmchale@opensystemsmedia.com

From DC to Daylight....
Scan, Detect, Process, Record & Playback RF Signals at Will

RFvision-1 Wideband (40MHz) Scanning Transceiver

Shock-protected packaging

DTA-3290 (1U)

DTA-1000-R (1U)

- 0.02 – 6 GHz frequency coverage
- 40 MHz instantaneous BW
- 16-Bit ADC & DAC
- FPGA based programmable DDC for complex baseband conversion
- Real-Time recording (over 11 hrs for 40 MHz BW)
- Optional RF playback of recorded & simulated signal
- Convenient & failsafe GUI for control and display

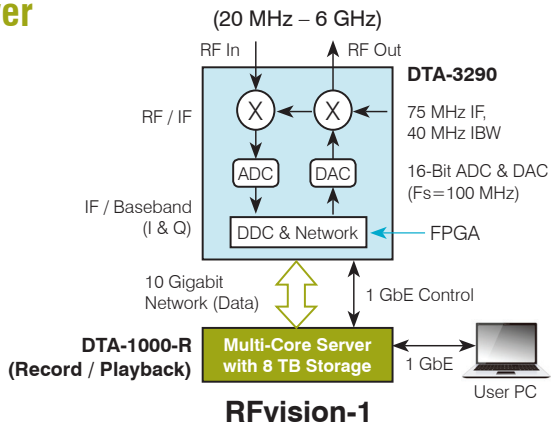

Click "RFvision-1 Wideband Scanning Transceiver" on our home page.

RFvision-2 Ultra-Wideband (500 MHz) Scanning Transceiver

Shock-protected packaging

DTA-5900 (1U)

DTA-9500 (1U)

DTA-5000 (3U)

- 0.5 – 26.5 GHz frequency coverage
- 500 MHz instantaneous BW
- 12-Bit ADC & DAC
- FPGA based DDC for complex baseband conversion
- Real-Time recording for hours
- Optional IF playback of recorded & simulated signal
- Convenient & failsafe GUI for control and display

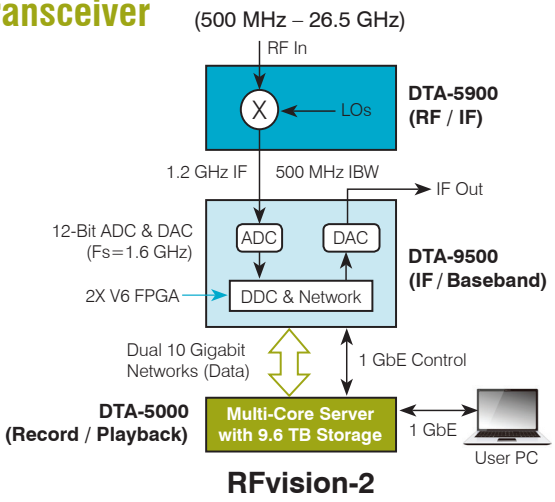

Click "RFvision-2 Ultra-Wideband Scanning Transceiver" on our home page.

SigInspector™ GUI

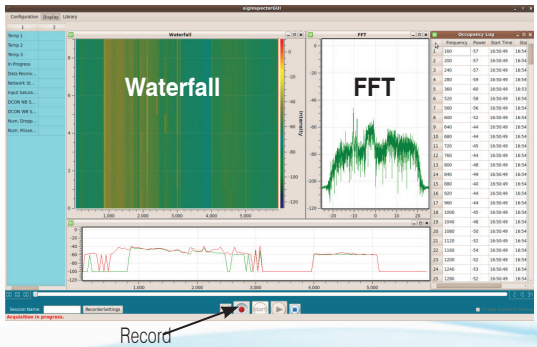

- 'Plug & Play' solutions with D-TA's SigInspector™ GUI
- Rapid user development using D-TA's Software Development Kit (SDK) and support

D-TA SYSTEMS INC.

A Sensor Interface and Processing Company

www.d-ta.com

sales@d-ta.com 1-877-382-3222

SPECTRUM INTELLIGENCE FROM D-TA
NOTHING ELSE COMES CLOSE!

Video processing on the fly

By Charlotte Adams

A GE Intelligent Platforms perspective on embedded military electronics trends

Video – in common with other sensor-derived data – plays an increasing role in today's military. Its proliferation at all levels of command reflects a deep hunger on the part of decision makers for remotely collected imagery that helps them see threats and, if necessary, deploy soldiers and weapons. This data is all the more important in maximizing the effectiveness and minimizing the vulnerability of forces as troop levels decrease in theaters of war and, in the future, as the overall footprint changes. As such, the demand for image processing solutions is widely expected to increase.

This market development is all the more likely as potential uses for video multiply in number and expand in scope. The National Aeronautics and Space Administration (NASA), for example, already uses captured video to monitor wildfires and hurricanes. The Department of Homeland Security (DHS) uses it for border surveillance. And local jurisdictions want to use vehicle-captured video in law enforcement and public safety roles.

The challenge

Increasingly, military vehicles are being equipped with high-definition video cameras and sometimes with multiple sensors and sensor types per vehicle. Imagers on legacy platforms, meanwhile, are being upgraded to provide higher-resolution data. But how can users handle this torrent of image data pouring into control centers hundreds or thousands of miles away via satellite or other links that may be much more limited in bandwidth than the data source?

In addition to the volume of data to be uploaded, compressed, manipulated, stored, or transmitted, the remoteness of video capture platforms puts a premium on timeliness. To address threats and other intelligence within tactical timelines, information must be presented with minimal latencies. This means the image processor needs to be flexible enough to accommodate multiple data streams and data formats simultaneously and adjust output bandwidth on the fly.

Technology to the rescue

Fortunately, image processing technology has evolved to keep pace with data volume. Small-format boards can handle multiple data streams, bandwidths and storage, or transmission tasks. They also can format the data according to multiple standards, convert formats on the fly, and vary compression ratios according to the limits of particular data links. Reconfiguration can be performed in real time via radio frequency link or remote commands to the host processing boards. This type of flexibility is key to military customers. Further, while multiple vendors supply image processing hardware, the most desirable solutions optimize flexibility in all areas of performance, from packaging to remote configurability and data compression.

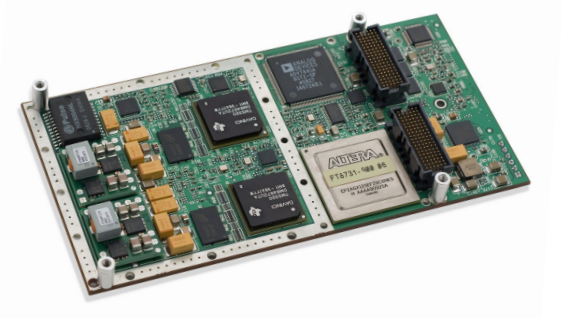

Figure 1 | Video streaming, compression, and capture boards such as GE Intelligent Platforms' ICS-8580 are satisfying today's increasing demands for optimized, manifold video processing.

An image processor's ability to modify compression parameters on the fly enables users to maximize video quality within the limits of available bandwidth. In critical phases of a mission, higher-capacity, line-of-sight data links can be used for vehicle safety. But when the vehicle is in theater within an area of interest, operators can limit its output bandwidth to suit the chosen data link. The image processor would then maximize video quality within those constraints by intelligently managing the data compression ratio.

One means of satisfying the hunger for optimized, manifold video processing is to utilize video streaming, compression, and capture boards built to provide a diverse range of video processing capability. One example is the GE Intelligent Platforms ICS-8580, a video processing board in an XMC form factor with high-speed PCI Express connectors that can be plugged into a single board computer in a larger system or function as a ruggedized stand-alone product (Figure 1). Such a board's architecture can accommodate multiple data streams simultaneously. One card can process: two channels of high-definition video (1,920 x 1,080 pixels); one channel of high-definition and two channels of standard-definition (720 x 480 pixels) video; or four standard-definition video channels. Data compression ratios are configurable and adaptable but range as high as 150:1 with 75 ms latency from input to output for full high-definition video. By comparison, JPEG 2000 compresses data at a 30:1 ratio, which could overload a narrowband data link.

Video processing: Flexible technology is key

Whether it's a question of the packaging, multidata stream capability, compression rate adaptability, or support for multiple camera input and output types, the video processing platform's flexibility is the key to image processing performance. This ensures not only that the disparate needs of military users can be met, but that new video technologies can be adopted as they become mainstream.

defense.ge-ip.com

FACE avionics standard will transform systems design

By Curtis Reichenfeld
An industry perspective from Curtiss-Wright Controls Defense Solutions

In today's avionics system world, applications tend to be costly and unique and have long lead times. Because they are proprietary they create barriers to competition, which keeps costs high in an era of constrained budgets. The Future Airborne Capability Environment (FACE) promises to usher in a brave new world, where, for example, code developed for a helicopter radio software application can be transparently reused in a fighter jet. Reuse is dissuaded today because any modifications made to software created for a unique platform will require costly and time-consuming code recertification and testing. FACE middleware will make the reuse process simple and transparent, essentially doing away with the need to modify and recertify code.

The FACE initiative is backed by the government, principle and second-tier prime contractors, and the COTS vendor community, and is getting closer to producing a set of standards that will define a middleware approach for application software. This FACE approach promises to dramatically lower system development costs through application code reuse and portability. But be advised, no COTS vendor can claim that their products are FACE compliant today. COTS vendors are only able to identify appropriate products as "FACE Candidate Hardware" or "FACE Candidate Software," indicating that they have been designed to be suitable for certification and testing once the FACE standard is finalized and those protocols have been defined.

FACE examined

The objective of the FACE Consortium is to define a standard common operating environment to support portable capability-based applications across Department of Defense (DoD) avionics systems. The consortium is currently working toward establishing a strict set of open standards that builds on Open Architecture (OA), Integrated Modular Avionics (IMA), and Modular Open Systems Approach (MOSA) concepts, all of which fit perfectly with the defense and aerospace COTS market's goals and principles.

Edition 2.0 of the FACE Technical Standard was completed on January 25, 2013 and is currently undergoing review by member companies. But FACE is coming soon; the government is already issuing RFPs and RFIs for which FACE is a requirement. FACE's mission is to define open middleware that will enable an application to talk to any hardware platform on any operating system over any I/O. Simply put, FACE has the potential to change the military avionics system landscape by making it possible to reuse standard avionics components across platforms and promote rapid technology insertion for avionics upgrades while cutting costs and reducing maintenance requirements and increasing competition among defense contractors.

Additionally, FACE will be a boon for designers of next-generation platforms who will gain access to less costly, faster-to-market, best-of-breed proven solutions. Likewise, FACE will improve technology on legacy aircraft, helping to extend useful life and cost. Leveraging FACE middleware, systems on older platforms will be upgraded more easily, quickly, and cost effectively and gain additional capabilities via new FACE-compliant applications developed first for use on next-generation platforms.

And FACE eliminates software and hardware interoperability headaches. The FACE middleware layer, which will reside on a particular COTS board's BSP, will feature a FACE-compatible API. This means that the application will not be aware of or even care what operating system is being used or whether the underlying hardware is Intel or Power Architecture. As the platform becomes transparent, COTS vendors can focus on their differentiators such as integration tools, performance, cost, and faster time to market.

Figure 1 | Curtiss-Wright Controls Defense Solutions' MPMC-9335 3U VPX Mission Computer and its VPX6-187 6U VPX SBC are examples of FACE Candidate Hardware that can host FACE Candidate Software.

As an active member of FACE, participating in the Technical Working Group and Business Working Group, Curtiss-Wright Controls Defense Solutions is helping to provide the COTS vendor perspective to the initiative and offers FACE Candidate Hardware; this includes rugged embedded open standards-based modules and systems that are able to host FACE Candidate Software applications and operating systems (Figure 1).

FACE progresses up the ranks

FACE is gaining traction and the consortium is on track, making progress toward a releasable standard. COTS vendors and their customers need to be aware of this coming standard so that they are prepared to take full advantage of the benefits it can deliver. For more information, contact the FACE Consortium at www.opengroup.org/getinvolved/becomeamember, or the Open Group at www.opengroup.org.

Curtis Reichenfeld
Chief Technical Officer
creichenfeld@curtisswright.com
www.cwcdefense.com

"Buy" over "build": The military systems migration mantra

By Steve Graves

In the shift toward greater use of Commercial Off-the-Shelf (COTS) system components, software lags hardware, and the prevalence of non-COTS technology is still relatively high in military/aerospace systems because of protracted life cycles. But contractors increasingly view system migration in terms of acquiring and integrating new prebuilt software rather than adding to their old "homegrown" code base.

The trend is particularly marked in middleware and other complex software such as protocol stacks, messaging, and data management. Consider protocol stacks, for example. High-performance embedded networking requires fast, reliable, and reentrant TCP/IP stacks, extensively tested in multiple environments. Today's COTS protocol stacks deliver this, and are also configurable, enabling applications to choose between allocating buffers in advance or at runtime.

Data management for military/aerospace systems has similarly outgrown its self-developed roots. With the advent of the Internet, high-speed communications and inexpensive memory, and relatively faster processors, embedded software now manages greater volumes of more complex data.

"... Contractors increasingly view system migration in terms of acquiring and integrating new prebuilt software rather than adding to their old 'homegrown' code base."

To meet this challenge, a new breed of COTS DataBase Management System (DBMS) has emerged – poised to replace both homegrown database code and enterprise DBMSs used in some military embedded systems – thereby offering a new migration path. The following examines some of their characteristics.

Many short, fast transactions

The transaction – a grouping of operations that must succeed or fail as a unit – is a database system's basic unit of work, whether that DBMS is part of a corporate system or buried deep within defense gear. A key difference is the required transaction speed. Embedded applications typically have high transaction rates, with each transaction of short duration. Tellingly, enterprise DBMSs are benchmarked in transactions per minute, while for today's fast embedded databases, the common metric is transactions per second.

One key change is that emergent embedded databases replace the client/server architecture of enterprise DBMSs with a design in which database operations execute within the application process. This eliminates the latency of Inter-Process Communication (IPC) between client and server, helping to deliver the needed submillisecond responsiveness.

Shared data and event processing

Most defense applications are event-driven, responding to interrupts from external sources. COTS databases for defense systems support event processing by propagating events to other "interested" software components. For example, in an Airborne Warning And Control System (AWACS), this "event" could be acquisition of data objects – airborne, seaborne, and terrestrial – from sensors, which triggers an operator console to show an alert.

Another example: In battlefield simulation, a task representing one combatant group takes some action, which is recorded in the database and in turn triggers another task (simulating another combatant) to reevaluate the battlefield scenario. Without COTS database systems to provide event notifications, developers of legacy systems would likely spend time cobbling together a solution based on homegrown message-passing code.

Complex data and design flexibility

Defense systems manage complex data, at a speed that is not amenable to using the SQL Data Definition Language (DDL) and Application Programming Interface (API) of enterprise DBMSs. For example, flight systems work with standard airport data that includes the objects DateTime, AirportStatus, WeatherSource, AirportWeather, and Airport. These objects are complex (that is, each contains records consisting of different data types) and they are nested: DateTime is a record type within WeatherSource, and WeatherSource is a type within AirportWeather.

SQL is notoriously bad at handling nested data, because processing it requires decomposing/recomposing the data to/from a normalized form (that is, third normal form, 3NF, or higher). However, COTS-based DBMSs gaining ground for mil/aero systems provide APIs and DDLs that map directly to the C programming language used widely in the industry. With support for fixed- and variable-length arrays of atomic data types as well as nested structures, opaque, or untyped data, they handle complex data with speed and fluidity.

Fault tolerance

Military technology must be resilient; therefore, a DBMS needs the capability to maintain one or more standby copies of the database on separate hardware, with failover. Today's COTS database solutions go beyond simplistic mirroring or replication. They offer sophisticated mechanisms like "2-safe" replication, in which transactions must complete on replica node(s) before finishing on the master node; clustering architectures also dramatically increase net processing power and reduce system expansion costs, while delivering a more scalable and reliable database solution.

In contrast, some legacy systems provided basic homegrown replication, but without COTS DBMS bells and whistles (such as hot synchronization or binary schema evolution), much less full-blown clustering architectures.

Preserving time and talent

It is true that the capabilities described could be developed in-house by defense contractors, as in military days of yore where proprietary systems dominated. Additionally, some engineers still relish such a challenge, and this, as much as anything, ensures that homegrown data management, protocol stacks, and messaging software won't vanish from our midst.

But assuming that COTS vendors price their technology reasonably, deploying the time, expertise, and expense to build these features in-house is becoming less and less justifiable economically – in the case of modern DBMSs and other technologies – and is an ineffective use of the organization's time and talent these days.

Steve Graves *is cofounder and CEO of McObject and can be contacted at steve.graves@mcobject.com.*

DEFENSE TECH WIRE

NEWS | TRENDS | DOD SPENDS | CONTRACTS | TECHNOLOGY UPDATES

By Sharon Hess, Managing Editor

USS Milwaukee about to cool its new, more powerful waterjets

The USS Milwaukee Littoral Combat Ship (LCS), the Navy's number five, is soon to be fitted with new, first-of-their-kind, high-power-density waterjets (Figure 1). The waterjets' mission is to reduce the wear and tear on marine propellers and other rotating machinery powering LCS. Though water is smooth, an occurrence called *cavitation* isn't, resulting from air bubbles spawned by pressure changes. These air bubbles, in turn, after several instances can result in entire metal chunks eroding away. However, the new waterjets aim to thwart this phenomenon on high-speed ships, reducing maintenance and replacement costs and increasing time between repairs. Procured by collaboration among the Naval Surface Warfare Center, the Office of Naval Research (ONR), and developer Rolls-Royce Naval Marine, one of the new Axial-Flow Waterjet Mk-1 waterjets propels almost a half-million gallons of seawater per minute. Additionally, speeds in excess of 40 knots can be achieved when four of the Axial-Flow Waterjet Mk-1s are used in tandem. Full-scale sea trials are next on the docket for USS Milwaukee, anticipated within the coming 24 months.

Figure 1 | The USS Milwaukee is soon to be fitted with new, first-of-their-kind, high-power-density waterjets dubbed "Axial-Flow Waterjet Mk-1," designed to thwart propeller wear and tear and resultant maintenance and replacement costs. Official U.S. Navy file photo

Kickback on U.S. Armed Forces virtual trainer

Step aside, gamers. Now the U.S. Armed Services have their own way to shoot targets without harming anyone or anything, too – the Engagement Skills Trainer 2000, a weapons simulator that projects big-screen images in a dark room, engendering an oversized arcade arena vibe. But the Engagement Skills Trainer 2000, located at Joint Base McGuire-Dix-Lakehurst, NJ – has a compressed air feature designed to mimic a rifle's post-firing kickback. And another bonus not available to gamers: The trainer spits out a target-shot results card – reminiscent of live-fire ranges' paper targets – after three shots. A primary goal of the trainer is to save on ammunition expenses, as it enables service members to hone their basic markmanship skills prior to actual weapons range target practice. And the trainer additionally renders realistic combat scenario simulations to provide skill building for previously deployed and yet-to-be-deployed service members.

Raytheon/U.S. Navy doubleheader

NEWS

Raytheon and the U.S. Navy put pen to paper twice in one day, spawning a contract and a modification to a previously issued contract. First up: Raytheon Missile Systems in Tucson, AZ received a contract for $14 million for Standard Missile intermediate- and depot-level maintenance, special maintenance tasks, all-up-round recertifications, and maintenance and repairs for Foreign Military Sales, all for fiscal 2013 through fiscal 2017. Work is anticipated for completion by this September. Next up: Raytheon Co. also located in Tucson, AZ was awarded a cost-plus-fixed-fee modification totaling nearly $10 million for the Cobra Judy Replacement (CJR) program's alternative architecture study. The CJR program's umbrella covers design, development, and acquisition of a mission equipment suite and functional replacement ship to supplant the present-day USNS Observation Island and Cobra Judy. The CJR is destined for an identical mission as its predecessor: foreign ballistic missile data collection pertaining to international treaty verification. Work under the modification is also slated for completion this September. The contracting activity for both the contract and the modification is the Naval Sea Systems Command in Washington, D.C.

Special Ops' TACLAN to interconnect DoD-deployed elements

Net-centricity is imperative for the DoD these days (Figure 2). Thus, U.S. Special Operations Command recently awarded iGov Technologies Inc. in Reston, VA a $500 million IDIQ contract to render its net-centricity-enabling Tactical Local Area Network (TACLAN) Family of Systems design, integration, acquisition, production, and delivery. According to the iGov website, TACLAN interconnects DoD tactical force-deployed elements by utilizing the company's mobile IT infrastructure plus its corresponding applications. These elements might comprise SIPR, NIPR, and JWICS and/or HQ-level Joint Task Force and/or teams and soldiers, who all need the military IT functional support. Such support includes mission analysis/planning/execution, logistics tracking and planning, and intelligence reporting and analysis, among other tasks. Work under the contract is slated for Tampa, FL.

Figure 2 | iGov Technologies received a $500 million IDIQ contract from U.S. Special Ops to render the Tactical Local Area Network (TACLAN) Family of Systems. Photo courtesy of U.S. Army by Col. A.T. Ball

VIDEO

Video unveils ARGUS sensor for the first time

A new PBS-produced "Nova" unveils the world's highest-resolution camera: ARGUS. "For the first time, we actually have permission from the government to show the basic capabilities," says BAE Systems Engineer Yiannis Antoniades on the video. While the sensor itself is covered on-camera by what appears to be plastic bags, because of classified concerns, the footage shown and detailed capability coverage are, nevertheless, compelling. Most notable: The wide area persistent stare ARGUS sensor renders 1.8 billion pixels and yields the same view as 100 Predators providing simultaneous ISR on a mid-sized city. Additionally, the sensor renders live-to-ground streaming and storage of a whopping 1 million TB of video daily, equaling about 5,000 hours of high-def video. And a directive from ARGUS' budget-conscious funder, the Defense Advanced Research Projects Agency (DARPA), inspired Antoniades' usage of an everyday item as ARGUS' backbone: a cell phone imaging chip.

Watch the ARGUS video at video.opensystemsmedia.com.

TECHCHANNEL

The Unmanned Systems TechChannel, designed and managed by OpenSystems Media, brings the industry the want-to-knows and breaking news on the unmanned technology forefront, pertaining to military defense technology. Included are news and facts on unmanned vehicles such as Unmanned Aerial Vehicles (UAVs), Unmanned Aerial Systems (UASs), Unmanned Ground Vehicles (UGVs), and autonomous vehicles. We also bring you the latest on other unmanned technologies like UAV ground control stations, robotics, drones, and much more. We'll also relay the latest on U.S. DoD contracts pertaining to unmanned systems, in addition to news from unmanned systems industry consortia and companies.

http://tech.opensystemsmedia.com/unmanned-systems/

INNOVATION

Gartner recognizes General Dynamics' cybersecurity product as 'Visionary'

The Fidelis XPS network security solution from General Dynamics Fidelis Cybersecurity Solutions was recognized as a "Visionary" in Gartner Inc.,'s 2012 Magic Quadrant for Content-Aware Data Loss Prevention (DLP). Fidelis XPS gives organizations network visibility, control, and analysis capabilities covering all channels and ports in real time, to stop, analyze, and/or recognize data breaches that might occur on multi-gigabit-speed networks, whether the attacks are malware based or not. Fidelis XPS also provides organizations the capability to "consume third-party threat intelligence" or to create custom rules, the company reports. General Dynamics Fidelis Cybersecurity Solutions falls under the General Dynamics Advanced Information Systems umbrella, which is a General Dynamics business unit.

Other General Dynamics news: http://opsy.st/Vy6gLE and http://opsy.st/RVWwOz

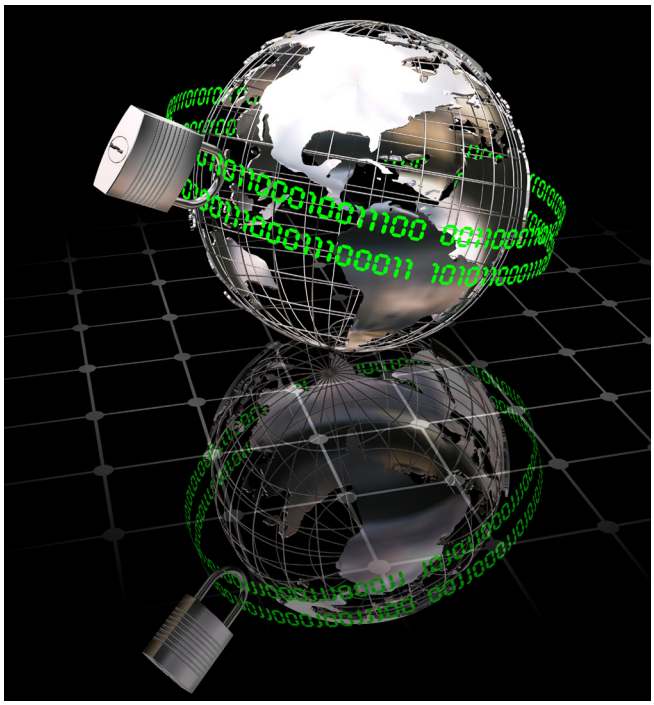

BLOG

Dissolvable military electronics investigated by DARPA program

By Brandon Lewis

The Defense Advanced Research Projects Agency (DARPA) has begun a program for the development of electronics that are capable of dissolving into the environment after use. The Vanishing Programmable Resources Program (VAPR) is aimed at developing transient electronics that are comparable in ruggedness in functionality to Commercial Off-the-Shelf (COTS) systems currently used, but that can be triggered to self-destruct in a controlled manner. The electronics themselves would be composed magnesium, silicon, and silk in ultra-thin sheets that are water soluble, and depending on the thickness of the silk, could decompose in times ranging from days to minutes. As low levels of magnesium and silicon are present within humans, DARPA Program Manager Alicia Jackson expressed DARPA's ambitions for the devices to eventually be able to "reabsorb into the body." Read more: http://opsy.st/14tvOib

Military avionics designs embrace common standards and TRLs

By John McHale, Editorial Director

The task of improving situational awareness for military pilots in a tough budget climate with little development funding available requires designers to use open architectures and common standards to keep costs down. This trend also has fueled the enthusiasm behind the FACE Consortium, which promises long-term potential savings of billions of dollars by enabling software reuse across multiple avionics platforms.

A Pro Line Fusion Synthetic Vision system developed by Rockwell Collins for the Bombardier Global 5000 business jet's flight deck has been adapted for the MH-60 Black Hawk helicopter. (Photo courtesy of Rockwell Collins)

Department of Defense (DoD) leaders want their pilots to have the best technology possible to do their jobs – whether it is a new touch-screen primary flight display, night vision goggles, 3D flight simulation systems, or even a brand-new aircraft to replace a decades-old platform. However, today's economic climate will not allow for new fighter jets or helicopters, so military program managers need to compromise and find ways to keep older aircraft platforms flying while still enhancing capability for the pilot.

"Military avionics customers in fixed- and rotary-wing platforms want improved situational awareness and to increase their mission envelope and effectiveness," says Karl Shepherd, Marketing Director for Airborne Solutions at Rockwell Collins in Cedar Rapids, IA. "For the DoD it is all about designing avionics that can reduce the workload of a pilot so they can enhance their decision making during missions. It's about helping pilots make decisions instead of taking decision making away from them."

While these enhancements do not require an entire new aircraft, the necessary upgrades can still be quite expensive – not just the up-front costs of new components and systems, but the behind-the-scenes costs such as training and maintenance. Keeping avionics retrofit costs down requires a move toward more commonality and open architecture designs.

Open architectures

"Open architectures provide the military customer with more value for his dollar," says Robert Waage, Director of Business Development at Elbit Systems of America in Fort Worth, TX. Multiple vendors can be used within one program, which helps drive costs down and enables platforms to fly for decades, he adds. Elbit's redesign and upgrade of the Apache Block III AH-64D Mission Processor will use an open architecture to accommodate new capabilities as they become available, Waage continues. The Army Apache team wants to sustain and maintain these aircraft until 2040, so they and the Apache prime contractor – Boeing – are creating road maps for future refreshes that leverage open architectures to achieve that goal, he adds. "It is about having market agility and flexibility."

"There is an increasing requirement that solutions should be based on open, interoperable, industry standards – even though the underlying technology is often less important to military integrators," says Simon Collins, Product Manager, GE Intelligent Platforms in Huntsville, AL. "Customers are more focused on how the solution meets their functional needs and how they will bring that solution to deployment."

Leveraging commercial avionics systems for military platforms

Looking to add avionics functionality and capability in an affordable way, military program managers often will look to adapt designs and applications that gained traction first in commercial aircraft cockpits.

"Designing with open architectures when upgrading military avionics platforms is crucial if we want to be able to leverage commercial processors and other components without major obsolescence issues," Shepherd says. "At Rockwell Collins we have been doing this for more than 10 years using technology developed for commercial avionics and then adapting it for military platforms. A big example of this is our Pro Line Fusion Synthetic Vision (SV) capability – developed for the business jet community and now being adapted for use in military rotary wing applications under a DARPA contract. The work essentially has a synthetic avionics backbone that hosts the SV terrain engine and fusion algorithms. The fusion algorithms process the data coming from the multifunction radar sensor along with terrain and obstacle data to provide an integrated 3D view of the operational environment on Heads-Down Displays and Heads-Up Displays in the cockpit. Real-time sensor imagery can be overlaid as well. The sensors provide real-time imagery, fused with the SV database imagery to give the pilot increased awareness when he can't see out the window.

"A primary difference between military and commercial SV applications occurs with the military platform's mission parameters," Shepherd continues. "We started with a civil-certified version of SV and made the necessary changes for high-resolution terrain data in a 3D environment out-the-window type of display. The initial rotary wing helicopters making use of SV technology will be those experiencing brownout conditions in Iraq and Afghanistan such as U.S. Army MH-60 Black Hawk and MH-47 Chinook helicopters. We see the Marine Corps and Army likely integrating the SV systems into more platforms further down the road."

FPGAs enable security in avionics systems

One area that needs more attention in avionics architecture designs is security. Each time something gets connected to the Internet, it automatically becomes vulnerable to some form of cyber attack, and avionics systems are no exception. "Security is the big white elephant in the room in electronics design. While our military avionics customers still want reduced size, weight, power, and cost, more and more they are concerned about security," says Tim Morin, Director of Product Marketing at Microsemi in San Jose, CA. "Protection from side channel attacks and enabling a secure boot process are two areas where security needs to improve." FPGAs are helping solve these challenges, he adds.

"Side channel attacks target the physical layer of a device for stealing cryptographic keys, which allow intruders to copy data and clone systems," Morin says. Different types of side channel attacks include: timing analysis, which observes differences in timing due to program branches, cache hits, and so on; Simple Power Analysis (SPA), which is where power supply currents may leak secret information directly; Differential Power Analysis (DPA), which uses statistical techniques to correlate power supply information leakage to recover critical information such as cryptographic keys; and Electromagnetic Analyses (EMA), which looks at the radiation emissions analog in power analyses, he continues. "FPGAs can prevent these by offering mitigation techniques such as DPA-resistant protocols by updating keys before leakage accumulation is significant. Microsemi's SmartFusion2 FPGA has a DPA patent portfolio license from Cryptography Research, Inc. (CRI) and is licensed to use DPA countermeasures for improved FPGA design security, such as bitstream decryption, to protect against these attacks."

Secure boot is another area where SoC FPGAs enable protection. "There are hundreds of processors onboard a typical aircraft today, and most often the first thing they do when booting is access something from external memory," Morin explains. "However, most processors do not verify if that external memory is a trusted source before they move onto the next step in the boot process. During the boot process, every step needs to be verified and validated by the previous step. None of the current processors offer that type of security while SoC FPGAs do by acting as a hardware root-of-trust."

Microsemi's SmartFusion2 acts in this way by "using the Intrinsic-ID SRAM PUF called *Quiddikey-Flex,* (either in soft or hard form) along with immutable on-chip embedded nonvolatile memory," he continues. "This along with other security features provides a root-of-trust for configuration and secure boot of the SoC FPGA itself. The SoC can then extend that trust to securely boot an external processor chip – even if the processor chip has limited or no intrinsic secure boot capability." The devices also leverage AES 256 encryption, Morin adds. For more information, visit www.microsemi.com.

High TRLs wanted

A tight budget environment may be perceived as one fostering increased use of more affordable COTS equipment since there is little funding available for development, but *COTS* has different meanings for different people. Some view it as too commercial and therefore lacking in reliability, while others feel that view is outdated and COTS has a proven performance track record. However, there is no proven metric for measuring the reliability of avionics products tagged as *COTS*, since it is more of a procurement term than a technological one. Instead, military program managers and integrators are turning toward Technology Readiness Levels (TRLs) to determine an avionics product's pedigree for military systems – whether it is COTS or not. "The

generals are demanding a high TRL level; therefore, I want and demand a high TRL level – something that has a high proven operational capability such as TRL level 7 or higher," Elbit's Waage says.

The FACE revolution

Open architectures, common industry standards, and TRL requirements have all been elements of an avionics procurement evolution driven by a need to better manage development costs without sacrificing capability and innovation. Avionics upgrades have embraced these concepts and seen major cost savings that enabled older platforms to live on for decades. Now branches within the DoD are taking cost-effectiveness to another level entirely by requiring future avionics systems to ensure portability of software across multiple platforms. This initiative has taken form as the Future Airborne Capability Environment (FACE) Consortium, and support for it is growing rapidly as DoD officials look for ways to make the funding cuts their civilian leadership is demanding.

"We see FACE – which is an open systems approach for avionics – as a natural evolution of leveraging commercial technology and common standards," Shepherd says. "It focuses on the reuse of software applications from one aircraft to another – even from one military service to another. Instead of paying each OEM or contractor every time you develop a software component, you can develop once and then redeploy the software."

"There is quite a bit of activity happening within the FACE Consortium today," says Chip Downing, Senior Director, Business Development, Aerospace & Defense at Wind River Systems in Alameda, CA. "The membership now has more than 50 organizations – including essentially every major U.S. defense prime contractor. The completion of the Technical Standard for Future Airborne Capability Environment (FACE), Edition 2.0 is expected [this month] in March 2013. The FACE Conformance Test Policy, Conformance Certification Guide, Conformance Tests, and the Conformance Verification Matrix Guides should be completed [later this month] with the official release targeted for early Q2 2013. Two major programs have also been awarded requiring FACE compliance to date."

One of these programs is the modernization of the C130-T, awarded to Lockheed Martin Mission Systems & Training in Owego, NY, by the U.S. Naval Air Systems Command. Under the $30 million contract, Lockheed Martin engineers "will deliver a suite of GFE [Government-Furnished Equipment] and CFE [Contractor-Furnished Equipment] to provide new navigation, communication, flight management, and controls and displays capability," says John Aebli, Director of Avionics Products for Lockheed Martin's Mission Systems & Training business.

"The new avionics package is built around the latest generation of FACE-conformant Lockheed Martin avionics,"

he continues. "The FACE framework enables cost savings by isolating change and variability. More specifically, individual software applications are isolated from the aircraft-specific configurations for controls and displays, other avionics, aircraft systems, and the hardware platform. This isolation enables software migration from platform to platform, as well as third-party application development, which leverages the nonproprietary software interfaces." Lockheed Martin engineers also will develop and deliver nine initial cockpit kits for the C130-T upgrade.

Figure 1 | The Navy's modernization of the C130-T's avionics by Lockheed Martin requires conformance to the FACE standard for software.

Enthusiasm for a new standard

"Whenever you talk to military leaders about FACE, they initially react as if the last thing the U.S. military needs is a new standard, and their defense shields go up," Downing says. "But once they learn how much money can be saved over the long-term, and how much [more] efficient and capable their platforms will be with FACE adoption, they can't hide their enthusiasm. FACE will potentially save the government billions of dollars in the long run. FACE has created both technical and business guidelines, and combined, both will drive more capability to the warfighter faster. Because the FACE team is getting so many work products completed, everyone can see the results, and this has energized the entire team."

"It's really vital for us to be members of FACE," GE IP's Collins says. "We see it as an incredibly important organization with objectives that are hugely significant for the industry in terms of application development, portability, reuse, and compliance, together with the inherent time- and cost-advantages of COTS solutions. When it comes to FACE, the customers we've been talking to typically fall into two camps: ones who are totally committed to going down the FACE-compliant route in their development and customers who see the benefits that FACE will bring but for now are watching and waiting."

Interoperability

"The movement behind FACE basically came from the government and DoD being fed up with being charged every time they used software programs in different platforms," says Robert Day,

Vice President of Marketing for LynuxWorks in San Jose, CA. "FACE solves this problem by creating interoperability among software platforms. Applications that have the common FACE API can move among multiple platforms, making it easier to keep legacy products and have them work with new applications and designs."

"Interoperability is an underlying theme with FACE. Software applications running on a mission computer in a helicopter platform could be reused on an Unmanned Aerial Vehicle (UAV), or the UAV system could work on a manned aircraft – all because of the common foundation," Downing says. "In this way, FACE has a similar utilization model to that of Android, where applications can be used across multiple platforms due to its underpinning architecture.

"The FACE team was smart – they optimized existing standards, ARINC 653, and POSIX to create the FACE technical standard," Downing continues. "Realizing that there is a broad range of applications in military avionics, the team created four separate profiles – Security, Minimum Safety, Extended Safety, and General Purpose. The Security and Safety profiles are designed to complement Common Criteria security and RTCA DO-178C safety certification environments. I expect that all of these profiles will be able to leverage the existing COTS certification and quality attributes of existing ARINC 653 and POSIX products from a wide range of vendors." Wind River's VxWorks 653 Real-Time-Operating-System (RTOS) for ARINC 653 time-and-space partitioned environments is FACE compliant, he adds.

"Most of the avionics funding for the next few years will go toward upgrading and modernizing current platforms," Day says. "If you don't have a FACE API, you likely will be left out of new refresh programs coming up. Our safety-critical RTOS 178 product – which has a native POSIX API – is FACE compliant."

"We have been reflecting FACE activities in our product development road map and are working with a number of software vendors to ensure their offering is compatible with the FACEREF1 Software Reference Platform that we introduced last year," Collins says. "It's a platform that features GE IP's SBC312 3U VPX single board computer and PMCCG1 graphics PMC" and enables organizations developing FACE-compliant applications to reduce risk by using preconfigured, prevalidated, and pretested COTS solutions, he adds.

Green Hills Software in Santa Barbara, CA, which makes the POSIX-compliant INTEGRTY RTOS, also is a member of FACE. The INTEGRITY-178B tuMP multicore operating system was selected for use in their Gen II Mission Computer for upgrades of the U.S. Marine Corps UH-1Y and AH-1Z helicopters. **MES**

Our technology investments protect yours.

Whatever direction you choose, Aitech has the map!

Established platform parallel bus protocols like VMEbus and CompactPCI still have their place in today's and tomorrow's harsh environment, real-time/hard-deadline embedded sub-system applications...especially when these products are upgraded and maintained to keep pace with the newest, fastest processor and memory technologies.

While there are some applications where high speed serial fabrics like VPX are ideal, there are others where VMEbus or CompactPCI still rule the roost.

One company continues to actively invest in maintaining – and **not** obsolescing – their military and space embedded computing products with a proactive 12-year minimum COTS Lifecycle+™ Program.

And one company continues to also invest in delivering the very best of the newest embedded COTS computing platforms with the new, serial fabric protocols.

And one company actively invests in technology insertion at the board level, creating backplane, pin-compatible products with the latest, next generation memory and processor technologies "on-board".

And that same company still delivers their legacy bus products at full speed **and** full capability and full mil temp range (-55 to +85°C) with those latest technologies.

The one company to do all that? Aitech. Check our website to learn more about our technology roadmaps and how they protect your investments.

Aitech Defense Systems, Inc.
19756 Prairie Street
Chatsworth, CA 91311
email: sales@rugged.com
Toll Free: 888-Aitech8 - (888) 248-3248
Fax: (818) 407-1502
www.rugged.com

Celebrating 28+ Years of Success

TRENDS IN
SMALL FORM FACTORS

Small form factor designs balance performance, power

By Brandon Lewis, Associate Editor and Monique DeVoe, Asst. Managing Editor

Whether next to the sensor on a UAV payload, in a vetronics mission computer, or in a helicopter's avionics bay, computer architectures are shrinking, requiring smaller form factors with the same performance punch as larger designs. The embedded COTS community is offering various Small Form Factor (SFF) flavors – off-the-shelf or customized – while industry groups such as VITA are developing standards to provide some order to the growing SFF demand.

Vetronics systems on vehicles such as the M-ATV benefit from smaller, high-performance SFF designs. Photo courtesy of Parvus Corporation

"Smaller, faster, cheaper" was NASA's space exploration motto a few years back, as they looked to send more affordable unmanned missions to Mars and other celestial targets. This phrase equally applies to the U.S. military's current approach to mission-critical electronics. The Department of Defense's (DoD's) program managers want to reduce Size, Weight, and Power (SWaP) in vetronics, avionics, and unmanned systems to take advantage of commercial processing performance and open architectures. In response, embedded Commercial Off-the-Shelf (COTS) suppliers are developing Small Form Factor (SFF) standards and products that leverage these processors while managing the power and thermal challenges that come with them.

Tight spaces in military platforms such as tanks drive much of this demand, as they have little room for ever-expanding electronics bays. Legacy systems already take up more than enough space, leaving little room for warfighters and their gear. To take advantage of high-performance computing architectures, the electronics systems will have to be reduced in size and power. Modern systems and SFFs can provide as much as 10 times the performance in a quarter of the space compared to systems from 5 to 10 years ago.

"In last year's tactical vehicles, there is equipment bolted to any available surface, with more equipment than available real estate. It is so bad in some vehicles that a reservist cannot fit into the crew seat wearing body armor," says Bill Ripley, Director of Business Development, Tactical Systems at Themis Computer in Fremont, CA. In some Army vehicles it is possible to easily achieve a reduction in components of around 50 percent by simply switching to modular, standards-based architectures, such as those outlined in the VITA SFF specs currently under development, he adds (see Sidebar 1).

Packing performance into SFFs

Defense customers want solutions that reduce the electronics footprint for mobile platforms such as Unmanned Aerial Vehicles (UAVs) and wearable computers. These systems require considerable processing power for compute-intensive applications such as data acquisition and analysis and graphical imaging.

"The biggest military market request is reliable equipment that solves size, cooling, and other design problems," says Ray Alderman, Executive Director of the VITA Standards Organization (VSO). "The military would love to use small form factors everywhere, but it's a challenge to get performance, enough memory, and enough processing power."

"Military customers are requiring PCI Express channels for high-speed subsystems such as video processing," explains Gary Schulz, Director of Marketing at VersaLogic in Eugene, OR. Mini PCIe cards also are being used

for system I/O expansion and memory functions, he adds.

"We're really talking about putting better processing at the edge, where you need more in-field intelligence, so you need localized processing capability," says Jack London, Embedded Products Business Unit Product Manager at Kontron in Poway, CA. Improved graphics performance is needed for better mapping and better field information, he adds. The technology is at a point where much of this processing can be done locally rather than having it sent back upstream, London continues.

"The military market is looking for higher-performance products from a computing architecture and networking architecture standpoint," says Michael Smith, Lead Engineer at Parvus Corporation in Salt Lake City. "They're asking for multicore architectures, more RAM, high-capacity SSDs, and they want it all in the smallest footprint possible." Equipment manufacturers are integrating GPUs into their CPU chipsets, an architecture that generates considerably more power consumption and, in turn, constrains the thermal limitations of systems using only passive cooling, he adds.

Small and cool

Higher performance in a smaller footprint is often referred to as the *functionality-to-size ratio*, and designers want this ratio to be as high as possible (Figure 1). However, the higher-performing processors generate an excessive amount of heat that is difficult to dissipate using only conduction and convection cooling methods. "Military embedded systems are facing the same kinds of problems that there are in cell phones – trying to get the heat out of individual modules," Ripley says.

Although advanced solutions for cooling embedded systems exist, the methods can be very expensive, exotic, and even inefficient for some airborne applications, Alderman says. Methods such as "liquid cooling cannot be used in avionics due to G forces and the uncertainty of the orientation of the aircraft," he continues. "Heat pipes that use a wick also can't be used due to G forces. Even nuclear submarines can pull Gs like an airplane, so liquid cooling can't be used there either. The optimal method for cooling is a hybrid system that depends on the application, but generally it involves a chassis that is conduction cooled, with liquid running

(continued on page 25)

Figure 1 | Themis Computer's RES-mini SFF high-density servers aim for a high functionality-to-size ratio by managing the thermal challenges that come with Intel Xeon E5 1600 and 2600 devices.

Three VITA standards aim to meet military SFF challenges

By Monique DeVoe, Assistant Managing Editor

Members of the VITA Standards Organization (VSO) are developing three small form factor standards influenced by the Department of Defense's (DoD's) Vehicle Integration for C4ISR/EW Interoperability (VICTORY) initiative and the increasing demand for SWaP optimization. Though other SFF standards also fit SWaP needs, VITA 73, 74, and 75 are being designed to address the updated VICTORY requirements and ruggedization needs of manned and unmanned vehicles (Figure 1).

	3U VPX	VITA 75	VITA 73 Double	VITA 73 Single	VITA 74
Board Size	160 cm^2	100 cm^2	144.84 cm^2	72.42 cm^2	69.40 cm^2
Total available component area^2	216.46 cm^2	128 cm^2	219.60 cm^2	109.80 cm^2	56.25 cm^2
Cubic envelope size^3	384 cm^3	240 cm^3	160 cm^3	80 cm^3	70 cm^3

Figure 1 | The developing VITA 73, 74, and 75 SFF standards target applications that require form factors smaller than 3U VPX.

VITA 73

The first of these VICTORY-inspired standards is VITA 73 – headed by PCI Systems in Sunnyvale, CA. The small form factor product it aims to develop is about the size of a 2.5" hard drive. It is driven by applications such as small UAVs, which often require a form factor smaller than 3U VPX, but with the same performance and bandwidth. VITA 73 also tries to fit the highest density of functions in the smallest possible size, says Ben Gross, VP of Sales and Marketing at PCI Systems. "We can directly compete with the bigger form factors like 3U, surpassing it in function-to-density ratio," he adds.

Their VITA 73 product does away with OpenVPX-style profiles and connects all slots with PCIe Gen 3 x16, and places an active high-speed switch to the active backplane. The active backplane deviates from the current method, but provides more options that help create a highly responsive system, Gross says. PCI Systems' VITA 73 device was used in small drones during a Navy project that simulated airborne naval threats on small UAVs. The customer preferred the VITA 73-based system over a 3U VPX solution because of its smaller size, lower weight, and easier profile options, he notes.

VITA 74

Themis Computer heads the VITA 74 initiative, which is based in part on COM Express mini (formerly known as *nanoETXexpress*). VITA 74 is about the size of a credit card at 54 mm x 85 mm and remaps the signals from the VPX standard onto a lower-cost, smaller form factor connector. VITA 74 has some performance advantages over other SFFs. "The connector used to interface between the module and the backplane can support data rates as fast as 20 Gbps, more than enough for PCIe Gen 3, 10 Gig-KR, and USB 3.0," says Dennis Smith, VP of Engineering at Themis Computer in Fremont, CA. "There is enough headroom to allow at least two more generations of fabric to be included." The fat pipes used in VITA 74 can support as much as 40 Gbps per lane, full duplex, between two card slots because of inherent connector speed and the very small spacing in the backplane. Across 16 signal lanes, 640 Gbps or 80 GBps is available.

VITA 74 also provides the thermal management needed with higher performance. According to Themis, 19 mm modules can handle power levels as high as 30 W without using exotic materials or designs. The resistivity of the thermal path from the chip to the die to the module thermal interface can be significantly less than a corresponding large form factor module, says Bill Ripley, Director of Business Development, Tactical Systems at Themis. Testing has shown that it has thermal properties that are better than required for a system of its size and power level, he adds.

VITA 75

The third VITA SFF standard effort, which is led by Curtiss-Wright Controls Defense Solutions in Ashburn, VA and Mercury Systems in Chelmsford, MA, targets the box level – enclosures and external interfaces. In the VITA 75 process so far, working group members from Curtiss-Wright, Mercury, Lockheed Martin, and other companies have gathered data on user problems with rugged mobile platform technology so they can create a standard that solves them. Work is being done to define the external dimensions of the box, the front panel profile, and cooling and mounting methods, as well as the types of modules to be used inside. Subprofiles have been released for manufacturers to design, build, and test products to gain feedback based on actual use.

A growing set of standards

On the supplier side there is a lot of interest, with several vendors offering solutions for image processing, video frame grabbing, commercial and SAASM GPS, inertial measurement, and analog and discrete I/O, Ripley says.

Though companies are already shipping VITA 73, 74, and 75 products, the standards are still in the "discovery phase," says VITA Standards Organization (VSO) Executive Director Ray Alderman. "VITA is figuring out what works, what makes sense, how to cool it, how to package it, and how to make it cost effective."

While 73, 74, and 75 development continues, not everyone is enamored with the three proposed standards. Engineers at Extreme Engineering Solutions (X-ES) in Madison, WI see problems with VITA 73 and 74 that make it a less-than-ideal choice. "For VITA 73 and 74, the XMC/PMC market is going to be nonexistent," says Jeff Porter, Senior Systems Engineer at X-ES. "There is a whole new infrastructure with 73 and 74, and we don't like that because you have to leverage a custom infrastructure."

Sidebar 1 | The VITA Standards Organization (VSO) continues to develop its triad of Vehicle Integration for C4ISR/EW Interoperability (VICTORY)-compliant small form factor standards: VITA 73, 74, and 75.

(continued from page 23)

through the end plates. SFFs mostly use conduction cooling or air cooling (which isn't viable for military applications due to the low MTBF of fans and their low reliability in rugged environments) that keep the cost low. More advanced cooling options raise costs."

Therefore, manufacturers have turned to more conventional ways of removing heat from SFF systems, most of which are targeted at modification to the board itself. "It all starts at the Integrated Circuit (IC) and works its way out to the external environment," says Doug Patterson, Vice President of the Military and Aerospace Business Sector at Aitech Defense Solutions in Chatsworth, CA. "Internal thermal impedances need to be minimized – from the active devices out to the real world – in the most cost-effective and least complicated ways."

These cost-effective and uncomplicated cooling solutions have largely presented themselves through convection and conduction, the former of which is typically accomplished through removal of active components on the SFF board. "Our military customers will say that they are trying to cut down on power, and we'll work with them very closely to contain Thermal Design Power (TDP) with depopulation of the module to reduce power consumption," London says. "We use System-on-Chip (SoC) solutions that minimize the amount of ICs on the board."

For suppliers that provide higher-end products, conduction-cooling strategies are often employed via modifications in board layout. "Since the thermally limiting component in many SFF systems is the processor, we have attached the die of the processor directly to the wall of the chassis [in the XPand6000]," says Jeff Porter, Senior Systems Architect at Extreme Engineering Solutions (X-ES) in Madison, WI (Figure 2). "Simply put, there is no better way to get the heat off the processor."

Figure 2 | The XPand6000 Series from Extreme Engineering Solutions (X-ES) attaches the die of the processor directly to the chassis wall to maximize natural conduction cooling.

Thermal management considerations can also differ based on application requirements – airborne platforms might have different thermal restrictions than a vetronics application. Complex thermal load definitions – created by processors and FPGAs that generate more or less

heat depending upon utilization and clock rates – also vary greatly depending on where the system is deployed, says Dave Caserza, Embedded Computing Architect at Elma Electronic in Atlanta. "For example, at very high altitudes, the temperature is quite cold and conduction cooling to an airframe works very well. When that same vehicle is on the ground in the desert, the conductive cooling is not as effective and some of the application features may need to be turned off."

Customizing COTS around requirements

Despite the work on SFF standards and open architectures, there will never be a one-size-fits-all approach for military electronics systems, as space constraints in aging helicopter cockpits, ground vehicles, small UAVs, and other platforms often require a somewhat custom approach. "Everything in the military market is customized to some degree, but based off of standards," Alderman says.

Attempting to mold the industry to standardized SFF platforms is a move contrary to the needs of the defense space, Porter says. "We feel that trying to get everybody's system requirements and capabilities to fit into the same-sized packages with the same connectors may be too difficult of a task to actually accomplish without limiting performance, capabilities, and innovation."

"Customers generally have a wide and diverse mix of sometimes-conflicting requirements," Patterson says. "Vendors periodically canvas the customer base to garner insight and direction for their next-gen products. Once we collect these requirements, we then develop products that can meet the majority of our customers' needs and demands with more off-the-shelf products and less-costly customization. Let's face it, money is tight right now. In just about every industry, customers want to maximize every dollar spent to position themselves with the latest technology at the highest Technology Readiness Level (TRL) to provide market leadership and value to their end customers." **MES**

Small Form Factor players

4DSP
www.4dsp.com

AAEON Electronics
www.aaeonusa.com

ACCES I/O Products
www.accesio.com

Acromag
www.acromag.com

ADL Embedded Solutions
www.adl-usa.com

ADLINK Technology
www.adlinktech.com

Advanced Micro Peripherals
www.ampltd.com

Advantech Corp.
www.advantech.com

AIM-USA
www.aim-online.com

Aitech Defense Systems
www.rugged.com

Alphi Technology
www.alphi.com

American Portwell Technology
www.american-portwell.com

BittWare
www.bittware.com

Colorado Engineering
www.coloradoengineering.com

Concurrent Technologies, Plc.
www.gocct.com

Connect Tech
www.connecttech.com

Creative Electronic Systems (CES)
www.ces.ch

Critical I/O
www.criticalio.com

Curtiss-Wright Controls Defense Solutions
www.cwcdefense.com

Data Device Corporation (DDC)
www.ddc-web.com

Diamond Systems Corporation
www.diamondsystems.com

Elma Electronic
www.elma.com

Emerson Network Power
www.emersonnetworkpower.com

Eurotech
www.eurotech.com

Excalibur Systems
www.mil-1553.com

Extreme Engineering Solutions (X-ES)
www.xes-inc.com

Fastwel
www.fastwel.com

GE Intelligent Platforms
defense.ge-ip.com

General Standards Corporation
www.generalstandards.com

Kontron
www.kontron.com

Mercury Systems
www.mrcy.com

Milcots
www.milcots.com

Molex
www.molex.com

Neuro Logic Systems
www.nlsdisplays.com

North Atlantic Industries
www.naii.com

Parvus Corp.
www.parvus.com

PCI Systems
www.pcisystems.com

Pentek
www.pentek.com

RTD Embedded Technologies
www.rtd.com

Samtec
www.samtec.com

Sealevel Systems
www.sealevel.com

SIE Computing Solutions, Inc.
www.sie-cs.com

Sundance DSP Inc.
www.sundancedsp.com

Sundance Multiprocessor Technology
www.sundance.com

TE Connectivity
www.te.com

TechSource
www.techsource.com

TEK Microsystems
www.tekmicro.com

Themis Computer Inc.
www.themis.com

Tri-M Technologies
www.tri-m.com

VectorNav
www.vectornav.com

VersaLogic
www.versalogic.com

WinSystems
www.winsystems.com

Connect with Us...

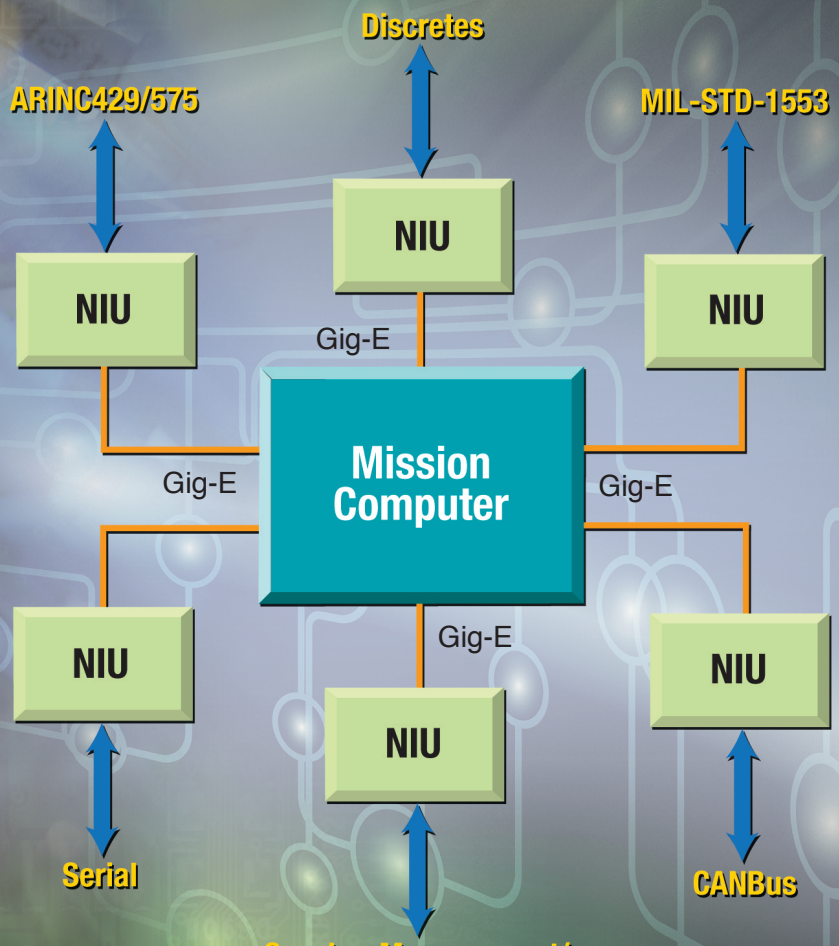

Proudly made in the USA.

Maintain Mission Computer configuration while adding I/O & Comms!

With the NIU1 Gigabit Ethernet-to-I/O interface unit, add I/O and communications functions to existing systems — **fast** with no redesign. Connect to existing Ethernet Networks for full access to all your I/O function needs.

The Nano Interface Unit (NIU1) measures only 6.5"W x 1.5"H x 1.6"D.

I/O support includes:

- High Density A/D, D/A, Discretes, Differential & TTL
- Synchro/Resolver/LVDT/RVDT Measurement, Simulation and Excitation
- SATA Flash Memory Storage, RTD, Encoder, Strain Gage
- Power over Ethernet (PoE) Support

Contact NAI today and let us show you how the NIU1 can adapt to your ever changing I/O needs!

The Single Source for Intelligent COTS Solutions

Visit www.naii.com/niu1 or call us at 631-567-1100 today.

Embedded Computing | *Power Supplies* | *Instruments*

631-567-1100 • Email: NIU1@naii.com • www.naii.com

North Atlantic Industries

Excellence in ALL we do

Cache partitioning increases CPU utilization for safety-critical multicore applications

By Tim King

Cache partitioning reduces worst-case execution time for critical tasks, thereby enhancing CPU utilization, especially for multicore applications.

U.S. Air Force photo by Tech. Sgt. Dana Rosso

One of the biggest challenges facing developers of certifiable, safety-critical software applications for MultiCore Processors (MCPs) is managing access to shared resources such as cache. MCPs significantly increase cache contention, causing Worst-Case Execution Times (WCETs) to exceed Average-Case Execution Times (ACETs) by 100 percent or more. Because safety-critical developers must budget for WCETs, tasks on average (critical and noncritical) are allocated more time than they need, resulting in significantly degraded CPU utilization. One way to address this problem is to utilize an RTOS that supports cache partitioning, which enables developers to bound and control interference patterns in a way that alleviates contention and reduces WCETs, thereby maximizing available CPU bandwidth without compromising safety criticality.

Cache contention

In a simple dual-core processor configuration (Figure 1), each core has its own CPU and L1 cache. Both cores share an L2 cache. (Note that shared memory and optional L3 are not shown.)

In this configuration, applications executing on Core 0 compete for the entire L2 cache with applications executing on Core 1. (Note that applications on the same core also compete with one another for L2; cache partitioning applies in this case as well.) If application A on Core 0 uses data that maps to the same cache line(s) as application B on Core 1, then a collision occurs.

For example, suppose A's data resides in L2; any accesses to that data will take very few processor cycles. But suppose B accesses data that happens to map to the same L2 cache line as A's data. At that point, A's data must be evicted from L2 (including a potential "write-back" to RAM), and B's data must be brought into cache from RAM. The time required to handle this collision is typically charged to B. Then, suppose A accesses its data again. Since that data is no longer in L2 (B's data is in its place), B's data must be evicted from L2 (including a potential "write-back" to RAM), and A's data must be brought back into cache from RAM. The time required to handle this collision is typically charged to A.

Most times, A and B will encounter such collisions infrequently. In those cases, their respective execution times can be considered as "average case" (ACET). However, on occasion, their data accesses will collide at a high frequency. In these cases, their respective execution times must be considered as "worst case" (WCET).

When developing certifiable, safety-critical software, one must budget an application's execution time for worst-case behavior. Such software must have

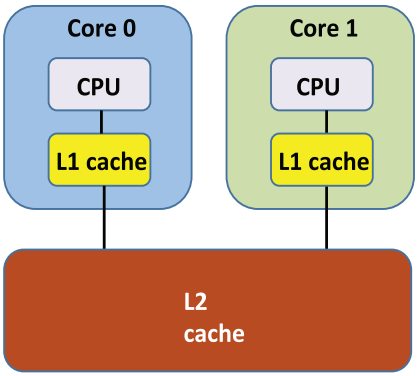

Figure 1 | Dual-core configuration without cache partitioning

adequate time budget to complete its intended function every time it executes, lest it cause an unsafe failure condition. The safety-critical RTOS must enforce time partitioning, such that each application has a fixed amount of CPU time budget to execute.

With the potential for multiple applications on multiple cores to generate contention for L2 cache, WCETs on MCPs often are considerably higher than ACETs. And since certifiable, safety-critical applications must have time budgets to accommodate their WCETs, this situation leads to a great deal of budgeted but unused time, resulting in significantly degraded CPU utilization.

Cache partitioning

Cache partitioning increases CPU utilization by reducing WCETs, thereby reducing the amount of time that must be budgeted to accommodate WCETs. Again, in a simple dual-core processor configuration (Figure 2), each core has its own CPU and L1 cache and both cores share an L2 cache.

In this configuration, the RTOS partitions the L2 cache such that each core has its own segment of L2, meaning that data used by applications on Core 0 will only be cached in Core 0's L2 partition. Similarly, data used by applications on Core 1 will only be cached in Core 1's L2 partition. This partitioning eliminates the possibility of applications on different cores interfering with one another via L2 collisions. Without such interference, the deltas between application WCETs and ACETs often are often considerably lower than is the case without cache partitioning. By bounding and controlling these interference patterns, cache partitioning makes application execution times more deterministic and enables developers to budget execution time more tightly, thereby keeping processor utilization high.

Test environment and applications

To demonstrate the benefits of cache partitioning, DDC-I used Deos (its certifiable, safety-critical, time- and space-partitioned RTOS) to run a suite of four memory-intensive test applications, all with a range of data/code sizes, sequential and random access strategies, and various working set sizes:

- Read only
- Write only
- Copy
- Code execution

The tests were performed on a 1.6 GHz Atom processor (x86) with 32 KB of L1 data cache, 24 KB of L1 instruction cache, and a 512 KB unified L2 cache. Note that while a single-core x86 processor was used for these tests, Deos' cache partitioning capability applies equally well to applications executing on the same core (which also compete for L2). Further, it does not depend on any features that are special or unique to x86 processors and applies equally well to other processor types (such as ARM or PowerPC).

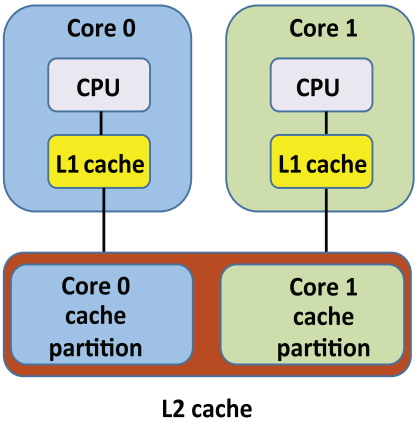

Figure 2 | Dual-core configuration with cache partitioning

The tests were run with and without a "cache trasher" application, which evicts test application data/code from L2 and "dirties" L2 with its own data/code. In effect, the cache trasher puts L2 into a worst-case state from a test application's perspective. That is, the cache trasher mimics real-world scenarios, where different applications run concurrently and compete for the shared L2 cache.

Each test application was executed under three scenarios. In scenario 1, which is conducted without cache partitioning or cache trashing, the test application competes for the entire 512 KB L2 cache along with the RTOS kernel and a variety of debug tools. This test establishes baseline average performance, wherein each test executes with an "average" amount of L2 contention.

In scenario 2, which uses no cache partitioning, the test application competes for the entire 512 KB L2 cache along with the RTOS kernel, the same set of debug tools used in scenario 1, and a rogue cache trasher application. This test establishes baseline worst-case performance, wherein each test executes with a worst-case amount of L2 interference from other applications, primarily the cache trasher.

In scenario 3, which uses cache partitioning and cache trashing, three L2 partitions are created:

- A 256 KB partition allocated to the test application
- A 64 KB partition allocated to the RTOS kernel and the same set of debug tools used in scenarios 1 and 2
- A 192 KB partition allocated to the rogue cache trasher application.

This scenario establishes optimized worst-case performance, wherein each test executes within its own L2 partition with no interference from other applications, including the cache trasher.

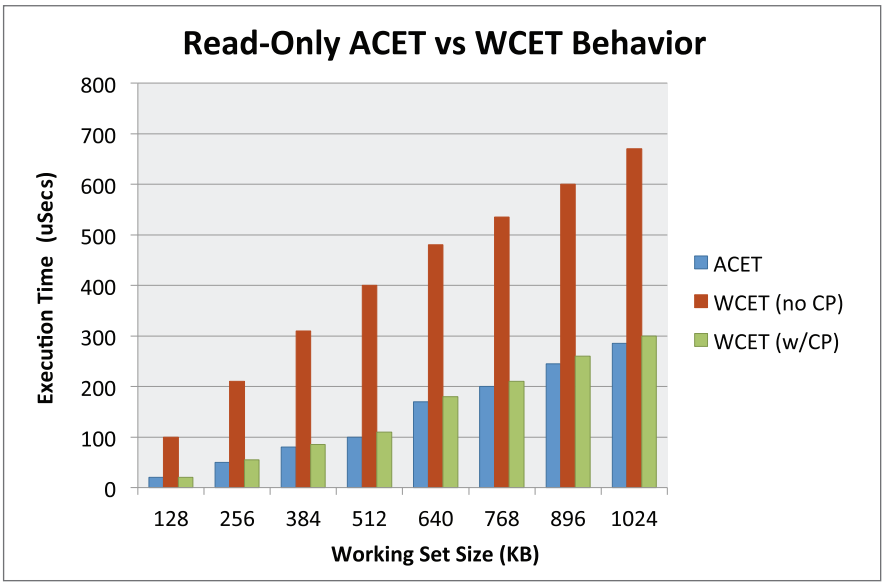

Figure 3 | Cache partitioning impact on read-only tests

Cache partitioning results, benefits

Figure 3 shows the results for the read-only test application.

For example, with no cache partitioning and no cache trashing (scenario 1, ACET), the read-only test averaged 105 microseconds per execution given a working set size of 512 KB. In scenario 2 (WCET with no partitioning, cache trashing added), the test averaged 400 microseconds per execution for the same 512 KB working set, a 280 percent increase. When cache partitioning is added (scenario 3, WCET with cache trashing), the average execution time drops to 117 microseconds, or just 11 percent higher than the ACET.

These results demonstrate the efficacy of cache partitioning for an application that performs a large number of reads per period. Though difficult to discern here due to differences in magnitude, the impact on bounding WCETs is more pronounced when the application's working set size fits within the cache partition that it's using (in this case, 256 KB). This result is expected because of the nature of cache. That said, embedded, real-time applications tend to have relatively small working set sizes, so we expect that cache partitioning will benefit most applications.

Results for the write-only test were similar to the read-only test, though more pronounced for smaller working sets. For larger working sets, results showed relatively small differences between WCETs with and without cache partitioning.

Results for the copy test were similar to the read-only test, though more pronounced for smaller working sets. For larger working sets, results were less dramatic, but still showed significant improvement (roughly 2x) in WCETs with cache partitioning.

Results for the code execution tests were similar to the read-only test, though somewhat less dramatic.

Note that it is possible for applications executing in the same cache partition to interfere with each other. However, such interference typically is much easier to analyze and bound than the unpredictable interference patterns that may occur between applications executing on different cores with shared cache. In those situations, if interference is unpredictable, then applications could be mapped to separate cache partitions.

The benchmark results clearly demonstrate that cache partitioning provides an effective means of bounding and controlling interference patterns in shared cache on an MCP. In particular, WCETs can be bounded and controlled much more tightly when the cache is partitioned. This allows application developers to set relatively tight, yet safe, execution time budgets, thereby maximizing MCP utilization.

Of course, results will vary for different applications and hardware configurations, and additional RTOS capabilities will be required to successfully certify safety-critical MCP-based systems. Regardless, these results represent a significant advancement toward the goal of using MCPs to host certifiable, safety-critical applications. **MES**

Tim King *is the Technical Marketing Manager at DDC-I. He has more than 20 years of experience developing, certifying, and marketing commercial avionics software and RTOSs. Tim is a graduate of the University of Iowa and Arizona State University, where he earned master's degrees in Computer Science and Business Administration, respectively. He can be contacted at tking@ddci.com.*

DDC-I
602-275-7172
www.ddci.com

Real parts. Really fast.

A product development team needs parts to meet its rapidly approaching deadline.

Upload 3D CAD file.

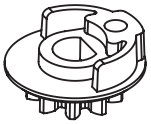

firstcut®
CNC Machining Service

Receive FirstQuote® interactive quote.

Finalize options, order with credit card or P.O. Computerized automated toolpaths sent to CNC Machine.

Machining begins.

Part ready for shipping.

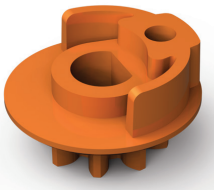

1–3 Days

1–10 parts
How Many Parts?
10–10,000+ parts

Receive a ProtoQuote® interactive quote.

Finalize quote and submit P.O.

Mold design and milling.

Receive order confirmation with gate and ejector layout. Approve.

Part production.

Parts ready for shipping.

protomold®
Injection Molding Service

1–15 Days

Check out our vitual tour!

It's easy to work with Proto Labs. Just upload your 3D CAD model and choose the best process for your project: CNC machining in 1–3 days or injection molding in 1–15 days. Real parts in real materials, in days—not weeks. And that's the real story. Call 877.479.3680 or visit www.protolabs.com

proto labs®
Real Parts. Really Fast.™

ITAR Registered

Puzzled by resin choices? Visit **www.protolabs.com/parts** today and request your FREE Resin Puzzle — nine of the most common thermoplastics used in injection molding. Enter code MY13A.

AVIONICS SAFETY CERTIFICATION

DO-332, the Liskov Substitution Principle, and local type consistency ramp up DO-178 certification

By Dr. Benjamin M. Brosgol

DO-332, the DO-178C standard's supplement on Object-Oriented Technology (OOT) and related techniques, analyzes the issues raised by object orientation in safety-critical software and supplies new guidance to deal with OOT's vulnerabilities. An important new objective of DO-332 is "Local Type Consistency Verification," which exploits a type theory result known as "the Liskov Substitution Principle" and helps address some of the key certification challenges raised by OOT's dynamic flexibility.

Object-Oriented Technology (OOT) is widely used and is supported by a range of programming languages including C++, Java, and Ada, but for various reasons its popularity has not spread to airborne and other safety-critical software. The underlying problem is the complexity of verifying software that makes use of three of OOT's basic elements: inheritance, polymorphism, and dynamic binding. (Figure 1 explains the object orientation basics.) A simple example illustrates the issues:

Suppose that the `Sensor` class is the root of an inheritance hierarchy, `ref` is a polymorphic reference to an object from any class in this hierarchy, and `Reset` is an operation defined differently for different `Sensor` classes. The statement `ref.Reset(...)` dynamically binds to the appropriate version, based on the class of the object denoted by `ref`. How does one verify that this invocation meets the requirements for the `Reset` operation?

One problem arises if inheritance is used to define a subclass that is not a specialization of `Sensor`. This subclass's `Reset` may have some effect unrelated to resetting a `Sensor`, or it may generate an exception. It would be an error for `ref` to reference an object from such a subclass, and analysis would be needed to show that the error could not occur. This complicates the verification process.

Another issue concerns structural coverage analysis. For systems at any of the three highest levels (A, B, or C) of the DO-178 standards, complete statement coverage must be demonstrated using requirements-based tests. But there are several implementation strategies that a compiler might choose for handling dynamic binding, with different implications on what "statement coverage" means. The scope of the structural coverage should not depend on the implementation strategy that the compiler uses.

DO-332[1], the OOT supplement to DO-178C[2], has addressed these issues through the new concept of *local type consistency*, which exploits the principle that inheritance should only be used for class specialization.

> “Object-Oriented Technology (OOT) is widely used and is supported by a range of programming languages including C++, Java, and Ada, but for various reasons its popularity has not spread to airborne and other safety-critical software.”

Inheritance and the Liskov Substitution Principle

In an object-oriented design, the system's architecture reflects the classes and their relationships. A particularly important relationship is specialization ("is a"), but there are many others. Implementing the design involves choosing language mechanisms for capturing the relationships.

In an object-oriented language, inheritance may be used to implement a variety of relationships between two classes. However, anomalies can arise when inheritance is used for anything besides specialization since operations defined for the superclass might not make sense for the subclass. Restricting inheritance to specialization relationships has been studied in the context of type theory, where it is known as the *Liskov Substitution Principle (LSP)*[3]. Informally, LSP means that wherever an instance of a superclass can be used, substituting an instance of any subclass should be permissible.

Using inheritance for specialization has an important interaction with an operation's preconditions and postconditions (its "contract"). A *precondition* is an assumption that the operation is making with respect to the program state when the operation is invoked. A *postcondition* is a guarantee that the operation is making on the program state when the operation is completed. Pre- and postconditions may be specified explicitly – either in the source text, as in Ada 2012[4] or SPARK[5], or separately – or they may be implicit in the logic of the operation.

If inheritance is to comply with LSP, a subclass's version of an operation should not impose a stronger (more restrictive) precondition than the superclass's version. Otherwise an invocation might succeed in some cases (on a superclass instance) but fail in others (on a subclass instance). Analogously, a subclass's version of an operation should not specify a weaker (more general) postcondition than the superclass's version.

Complying with LSP thus means satisfying two properties:

- **Contract consistency:** No subclass operation strengthens a precondition or weakens a postcondition of the superclass operation that it is overriding.
- **Behavioral consistency:** Each subclass operation meets its superclasses' requirements.

DO-332 captures these concepts in a new objective, Local Type Consistency Verification. This objective does not require demonstrating that the class hierarchy complies with LSP, which would be overly restrictive. Instead, it reflects that the verification effort is simpler for class architectures that do comply, and that the analysis only needs to consider local context.

Object orientation basics

Object-Oriented Technology (OOT) is characterized by several fundamental concepts:

Class: A template for objects. A class has members that are either operations (also known as *methods or functions*) or data fields.

Encapsulation: The ability to control access to a class's members based on a "need to know."

Object: An instance of a class.

Inheritance: A relationship between two classes, a *subclass* and a *superclass*. A *subclass* extends the superclass, inheriting its members and possibly overriding operations and/or adding new members. The classes related by inheritance, directly or indirectly, form an *inheritance hierarchy*.

Polymorphism: The ability of a variable to reference objects from different classes in the same class hierarchy at different times.

Dynamic binding: The interpretation of an operation applied to a polymorphic reference based on the class of the object currently denoted by the reference.

Figure 1 | The fundamental concepts of Object-Oriented Technology (OOT)

Local type consistency

Figure 2 shows the activity associated with verifying local type consistency, which DO-332 requires for software at levels A, B, or C. The wording "for each subtype where substitution is used" refers to contexts where dynamic binding occurs, such as `ref.Reset(...)`, and the "subtype" in question is the class of an object that `ref` could reference at that point. The potential classes will not necessarily be the full hierarchy, and different sets of classes may be applicable at different invocations of the same operation.

Local type consistency verification activity (DO-332, §OO.6.7.2)

For each subtype where substitution is used, perform one of the following activities:

- Verify substitutability using formal methods.
- Ensure that each class passes all the tests of all its parent types that the class can replace.
- For each call point, test every method that can be invoked at that call point (pessimistic testing).

Figure 2 | The activity associated with verifying local type consistency, which DO-332 requires for software at levels A, B, or C

Consider a particular occurrence of `ref.Reset(...)`, and let `HeatSensor` be one of the possible subclasses for the objects that `ref` could reference there. Local type consistency of `ref.Reset(...)` for `HeatSensor` may be demonstrated either "optimistically" or "pessimistically." The optimistic approach works if `HeatSensor` satisfies LSP, and may be carried out in two ways:

- Through formal methods, by demonstrating that `HeatSensor`'s version of `Reset` meets the requirements for `Sensor`'s version and does not strengthen the preconditions or weaken the postconditions of `Sensor`'s `Reset`.
- Through testing, by running the requirements-based tests for `Sensor`'s version of `Reset`, using an instance of `HeatSensor`.

Formal methods may be facilitated by appropriate support from the programming language and its toolset, for example, Ada 2012 or SPARK.

The optimistic approach will demonstrate contract and behavioral consistency between the superclass's and subclass's versions of the operation. Additional verification is obtained through requirements-based tests for the subclass and possibly also through formal methods.

If the classes do not comply with LSP, or if there are few dynamically bound

calls or the hierarchy is shallow and/or narrow, then it may be simplest just to test each possible case that can arise. This is the pessimistic testing specified in the third bullet item in Figure 2. Requirements-based tests are needed to exercise the operation for each subclass that could arise.

DO-332, local type consistency, and LSP guide certification

Local type consistency verification is just one aspect of using OOT safely; DO-332 contains guidance on other OOT elements as well as related techniques such as generic templates. DO-332 is "language agnostic"; more specifics are available on how to apply OOT in safety-critical or high-security systems, using Ada 2012 as the programming language[6].

DO-332's local type consistency guidance is consistent with DO-178C's general approach to verification, ensuring that all tests are based on requirements. It adapts the verification activities in a novel way to reflect the semantics of object orientation and the class structure's degree of compliance with LSP. The new guidance should help to promote the safe use of Object-Oriented Programming (OOP) in avionics and other critical domains. **MES**

References

[1] RTCA/EUROCAE DO-178C/ED-12C. Software Considerations in Airborne Systems and Equipment Certification, December 2011.

[2] RTCA/EUROCAE DO-332/ED-217. Object-Oriented Technology and Related Techniques Supplement to DO-178C and DO-278A, December 2011.

[3] B. Liskov and J. Wing. "A behavioral notion of subtyping," ACM Transactions on Programming Languages and Systems (TOPLAS), Vol. 16, Issue 6 (November 1994), pg. 1811-1841.

[4] Ada 2012 Language Reference Manual (December 2012), www.ada-auth.org/standards/ada12.html

[5] J. Barnes with Altran Praxis. SPARK – The Proven Approach to High Integrity Software. Altran Praxis, 2012.

[6] AdaCore, High-Integrity Object-Oriented Programming in Ada, July 2011. extranet.eu.adacore.com/articles/HighIntegrityAda.pdf

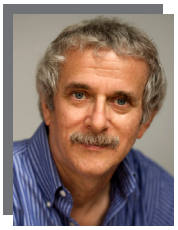

Dr. Benjamin M. Brosgol *is a senior member of the technical staff at AdaCore. He has more than 30 years of experience in the software industry, concentrating on languages and technologies for high-integrity systems. He has presented papers and tutorials on safety and security certification at numerous conferences and has published articles on this subject in a variety of technical journals. He holds a Ph.D. in Applied Mathematics from Harvard University. He can be contacted at brosgol@adacore.com.*

AdaCore
212-620-7300 • www.adacore.com
www.linkedin.com/company/adacore
www.twitter.com/AdaCoreCompany

FACE, ARINC, DO-178C avionics standards help U.S. DoD's vision of reusable technology to take off

By Bernard Dion

While historically single-vendor procured, today's airborne systems utilized by the U.S. DoD need to instead comply with the DoD's vision of more cost-efficient, reusable, modular, standards-based applications ready for flight in both manned and unmanned airspace. Standards such as the Future Airborne Capability Environment (FACE), ARINC 653 and 661, and DO-178C are helping the military avionics industry achieve this paradigm. Tools providing model based-design and verification capabilities can also assist.

Photo courtesy of Lockheed Martin

Historically, U.S. Department of Defense (DoD) airborne systems have typically been developed for a unique set of requirements by a single vendor. This form of development has served the DoD well until recently. However, this development process has major drawbacks including long development time and lack of hardware and software reuse between various aircraft platforms, which results in a platform-unique design. The advent of significantly more complex mission equipment and electronics systems has caused an increase in the cost and schedule to integrate new hardware and software into aircraft systems. This combined with the extensive testing and airworthiness qualification requirements have begun to affect the ability of the DoD aviation community to deploy these new capabilities across the DoD projects. The good news, though, is that avionics standards such as the Future Airborne Capability Environment (FACE), ARINC 653 and 661, and DO-178C can enable the aviation community to provide such technological advantages to the DoD, saving time and money. Model based-design and verification tools can also help achieve the goal.

FACE provides reusability

The FACE Consortium's FACE standard was designed as a response to the U.S. DoD aviation community's challenges. The approach used by FACE is to develop a standard for a software computing environment designed to promote software product lines that are reusable across different air platforms[1].

Several components comprise the FACE approach to software reuse. This approach allows software-based "capabilities" to be developed as components that are compatible with other software components through defined open standards interfaces. It also provides for the reuse of software across different hardware computing environments that contain differing platform devices. This includes the use of avionics functional standards such as ARINC 653 for Integrated Modular Avionics and ARINC 661 for the design of Cockpit Display Systems (CDSs) and User Applications (UAs). FACE will also support airworthiness qualification of airborne systems such as DO-178C Levels A through E and, in addition, the Common Criteria Evaluation Assurance Level 4 (EAL 4) through Level 7 (EAL 7). Ultimately, the goal of FACE is to reduce development and integration costs and reduce time to field new avionics capabilities. The FACE Consortium is hosted by the Open Group and it includes U.S. military industry suppliers, customers, and users. It meets regularly to define a reference architecture, a set of guidelines, and a business model.

ARINC 653: Multiple application hosting

The ARINC 653 standard[2] is a specification for executive software that allows hosting several avionics applications on a

single Integrated Modular Avionics (IMA) hardware platform while guaranteeing space- and time-partitioning for these critical applications. Assuming each application is certified for DO-178B/C at a given level (A through E), safety integrity of the overall avionics systems can be achieved since data integrity can be guaranteed and a sufficient time budget can be allocated to each application so that it can run in due time. ARINC 653 is a part of the ARINC 600-Series standards for Digital Aircraft and Flight Simulators.

An IMA-based methodology has several key benefits for airframers and avionics developers: It aims to minimize the number of on-board computers, resulting in lower size and weight, while enabling the use of legacy systems because components are interoperable and "plug and play." This flexible design can be changed and optimized at a late stage during development. The standard supports incremental certification, allowing airframers to better manage the DO-178B/C certification process. The ability to achieve incremental certification is a direct benefit of the ARINC 653 partitioning mechanisms since certification can be achieved first for each application and then for the overall systems. Moreover, changing an application does not require recertification of the other applications.

Figure 1 depicts the ARINC 653/IMA use model together with a model-based methodology that relies on a SysML-based description of the applications residing on the IMA and the automatic generation of the IMA configuration tables. The top-right system view exhibits three applications that reside on the IMA platform and their communications. On this basis and on the basis of a description of the IMA platform configuration (Usage Domain), a set of IMA configuration tables can be automatically generated to configure the platform. As shown on the left side of Figure 1, each application can itself be developed by using a traditional manual coding approach or, better, by using a model-based development approach.

ARINC 661 standardizes CDSs, UA communications

ARINC 661[3] is a standard that defines the Cockpit Display System (CDS) interfaces, including communication between the CDS and the avionics UAs and communication between the CDS and the pilots. It is also a part of the ARINC 600-Series Standards for Digital Aircraft and Flight Simulators.

The ARINC 661 standard normalizes the design of an interactive CDS and the manner in which these systems communicate with UAs such as flight management systems, flight control systems, flight warning systems, and so on. It does this by using a set of 67 predefined and standardized graphical widgets, some of it changeable through pilot interaction (trackball, keyboard, tactile screens, and so on), and by standardizing the communication protocol at runtime between a UA and the CDS. Thus, ARINC 661 ensures that the full CDS interactively behaves with the avionics systems in the same manner regardless of UA developer and CDS manufacturer.

ARINC 661-compliant applications provide key benefits for airframers, CDS developers, and UA developers. With model-based development tools that support the ARINC 661 standard, aircraft manufacturers, CDS developers, and avionics UA developers can ensure ARINC 661 compliance and dramatically increase productivity while achieving the highest level of graphical quality and the safety considerations that must be obeyed when certifying these applications under DO-178B/C up to Level A.

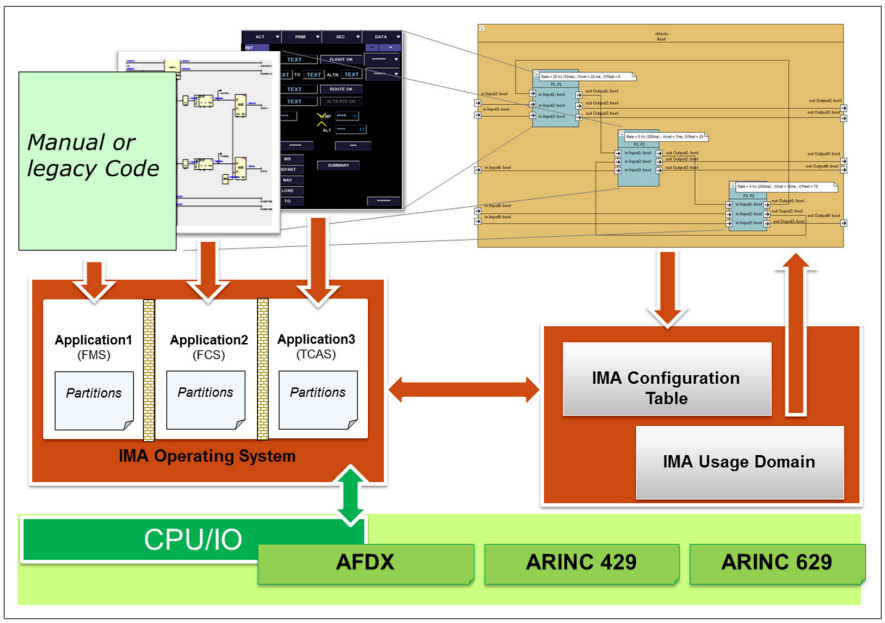

Figure 1 | The ARINC 653/IMA use model and model-based technology

Figure 2 depicts the ARINC 661 use model together with a model-based methodology that relies on specific tools to describe the graphical and behavioral properties of the widgets and the automatic generation of the ARINC 661 runtime server.

Graphics and logic modeling tools can be used to describe each widget that will reside in the CDS ARINC 661 runtime server. Moreover, for a given UA, the same graphical modeling tool can be used to describe how widgets are assembled on the pages pertaining to this application. From this description of the pages, the modeling tool is automatically generating Definition Files that are loaded into the CDS when the aircraft starts.

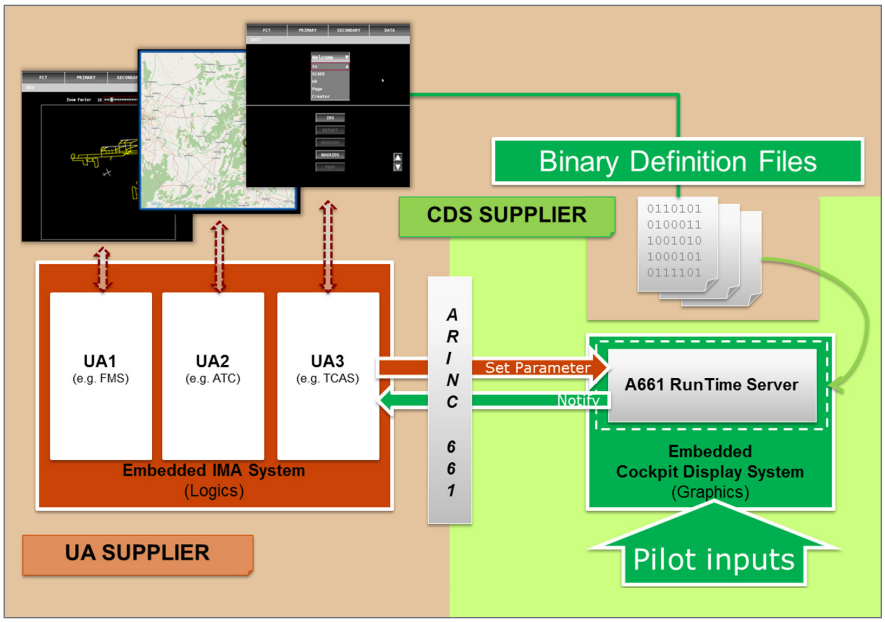

Figure 2 | ARINC 661 use model and model-based technology

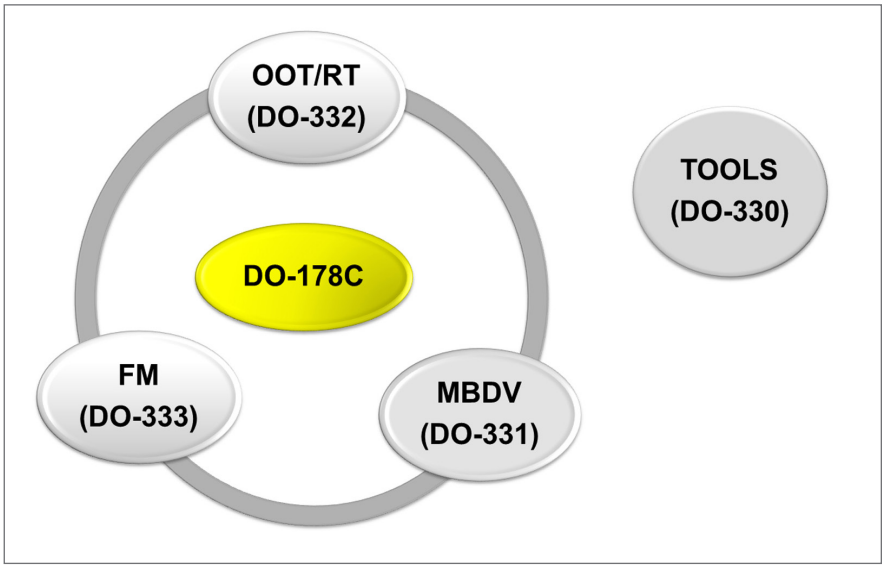

Figure 3 | The DO-178C documents structure

Understanding DO-178C

DO-178C should be approved by the U.S. Congress this year. Therefore, understanding the key differences in the standard as compared to DO-178B and how to best meet the certification requirements in the most cost- and time-effective manner possible will affect all avionics developers in the immediate future. The standard defines the objectives to be achieved for software certification of airborne systems. The overall structure of the DO-178C standard and related documents is depicted in Figure 3.

The core DO-178C standard is an evolution of the existing DO-178B standard. It defines a number of objectives regarding the processes that have to be followed to certify software at a given level (A through E). It comes with technology-specific documents pertaining to Object-Oriented Technology and Related Techniques (OO/RT), Model-Based Development and Verification (MBDV), and Formal Methods (FM) that complement the objectives to be achieved when a specific technology is used by the applicant (the aircraft manufacturer working in relation with its suppliers). In this discussion, we should put the emphasis on the Model-Based Development and Verification Supplement, DO-331[5], and the Tools Qualification document (DO-330)[6]. The rules and objectives described in these two documents allow efficient use of the methods described in the previous two sections. These methods are model-based (DO-331), and they can bring the most benefits to their users when they can rely on qualified tools (DO-330).

Standards, model-based design tools meet avionics challenge

Creating avionics applications that comply with FACE, ARINC 653 and 661, and DO-178C presents unique challenges for systems and software engineers. These applications must be designed in an effective manner to meet the requirements of these emerging standards, while taking into consideration development and certification costs.

Thus, to meet emerging avionics standards, avionics developers will have to rethink their systems and software development strategies. The time to do that is now. These standards are dictating all future development strategies, and

by utilizing tools that encourage reuse and reduce development and certification costs, the U.S. DoD, the FACE standard, and other emerging standards are leading the way into the future of aircraft development strategies.

Traditional tools relying mostly on manual coding methods will not allow meeting the challenges since they lead to high development and verification costs and they result in software code that is highly dependent on the initial technology and thus cannot be reused.

Meanwhile, model-based design and verification tools facilitate reusability and lower development and certification and therefore have unique advantages over non-model-based tools. They improve communication among system/software teams, customer, suppliers, and certification authorities because model-based design is used as a common graphical language in the development process. Model-based design tools also improve long-term maintainability, reuse, and retrofit of applications, as proposed by FACE, because they can support hardware, software, and platform independence. Also, developers can more easily maintain variations, making changes when needed in a graphical fashion, rather than having to maintain code. The SCADE product family is a good example of this type of model-based development tool. It supports both the ARINC 653 and ARINC 661 standards and automatic and DO-178B/C-certified code generation from software models. **MES**

References

[1] Future Airborne Capability Environment (FACE) Reference Architecture, the Open Group, 2012.

[2] ARINC 653, Avionics Application Software Standard Interface Software, ARINC, 2010.

[3] ARINC 661 Specification 661-4, Cockpit Display System Interfaces to User Systems, ARINC, 2010.

[4] DO-178C, Software Considerations in Airborne Systems and Equipment Certification, RTCA, 2011.

[5] DO-331, Model-Based Development and Verification, Supplement to DO-178C and DO-278A, RTCA, 2011.

[6] DO-330, Software Tool Qualification Considerations, RTCA, 2011.

Bernard Dion *is Chief Technical Officer at Esterel Technologies, which he cofounded in 1999. He received an E. Engr. degree from Ecole Centrale de Lyon and a Ph.D. in Computer Sciences from the University of Wisconsin. He then started at Honeywell Bull in the Ada language team. As a professor, he created a graduate school in systems and software engineering. He has been the European secretary of the tools qualification subgroup within the committee that created DO-178C.*

Esterel Technologies,
a subsidiary of ANSYS, Inc.
+33 1 30 68 61 60
www.esterel-technologies.com

Small Form Factors Product Spotlights

SoM 9X25

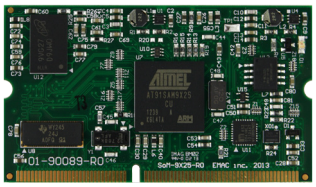

- Wide Temperature SoM 9X25 based on the Atmel AT91SAM9X25 400MHz Processor
- 10/100 BaseT Ethernet with on-board PHY (2nd Ethernet Optional)
- 6 Serial ports, 3 with handshake. Made in the USA.
- 2 USB 2.0 (Full Speed) Host ports & 1 USB 2.0 (Full Speed) Device port
- SD/MMC Flash Card Interface. Qty 1 price begins at $180
- 2 SPI & 2 I2C ports, 2 CAN 2.0B ports, 1 I2S Audio port

EMAC, inc.
EQUIPMENT MONITOR AND CONTROL

EMAC, Inc.
618-529-4525
www.emacinc.com

Radisys Ruggedized CEQM77HDE COM Express Module

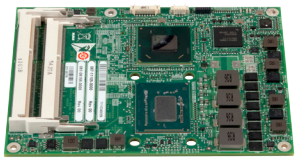

- Intel 3rd generation quad and dual core i7 processors
- Extended -40C to +85C temperature range
- Up to 16GB dual channel DDR3 memory with ECC
- Type 6 COM Express revision 2.0, Basic 125mm x 85mm form factor
- 7 x1 PCI Express ports and one x16 PEG port (Gen3), 3 DDI display ports
- Ideal for high performance communications and aerospace and defense equipment

radisys

Radisys
503-615-1100
www.radisys.com

E-cast

REGISTER NOW

for these and other free, on-demand webinars by accessing ecast.opensystemsmedia.com, then clicking on "See Archived E-casts and Events in-progress" at the bottom of the page.

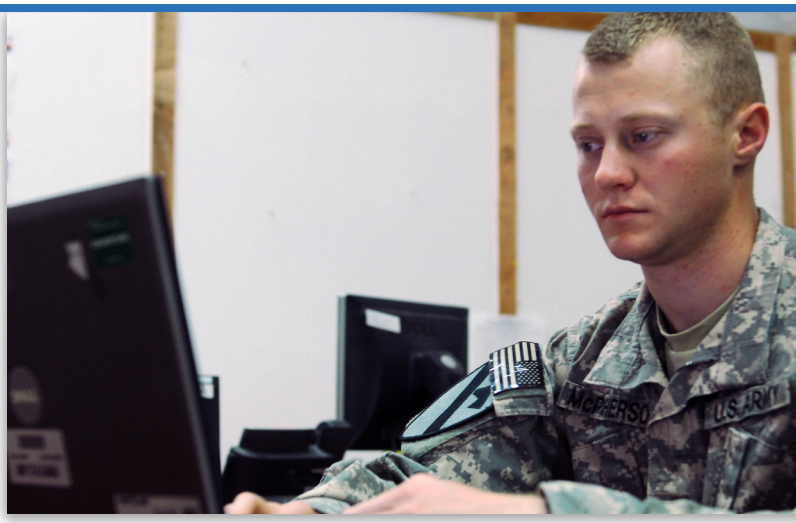

Securing Android Against Cyber Attacks

Presented by: Wind River

Cyber warfare has swiftly migrated from hacking into enterprise networks and the Internet to targeting, and being triggered from, mobile devices. With the recent explosion in the number of connected devices, security attacks can now be launched from anywhere and target any place. Securing mobile devices is crucial because the same capabilities that increase our social interactions and business effectiveness can be subverted and used for unintended and unapproved activities. As a widely deployed and robust open source platform, Android provides an ideal foundation for integrating security controls for advanced protection for high-assurance and battlefield use cases. Learn about Security Enhanced (SE) Android and about its vulnerabilities and how to address them.

Making it cool: Thermal management challenges in military electronics systems

Presented by: Mercury Systems, Pentek, and GE Intelligent Platforms

Designers of electronic systems for unmanned vehicles, radar systems, fighter aircraft, ground vehicles, helicopter avionics, and so on, all share a common challenge – how to keep electronic systems cool when modern processors and boards generate so much excessive heat and power. Today's chips from Intel, Freescale, and others provide unprecedented performance in military systems. Yet their excessive heat coupled with increased requests from DoD program managers for rugged small form factor designs and packaging at lower costs forces military embedded designers to be innovative with thermal management. This webcast of industry experts covers these Size, Weight, and Power (SWaP) challenges and more.

A 1.66GHz Intel® Atom™ Mil/COTS SBC of *EPIC* Proportions

The EPX-C380 is a rugged single board computer that provides an open powerful platform for harsh Mil/COTS applications. Powered with either a single or dual core processor, it has a wealth of onboard I/O plus expansion options. Also it supports Linux, Windows® XP embedded, and other x86 real-time operating systems.

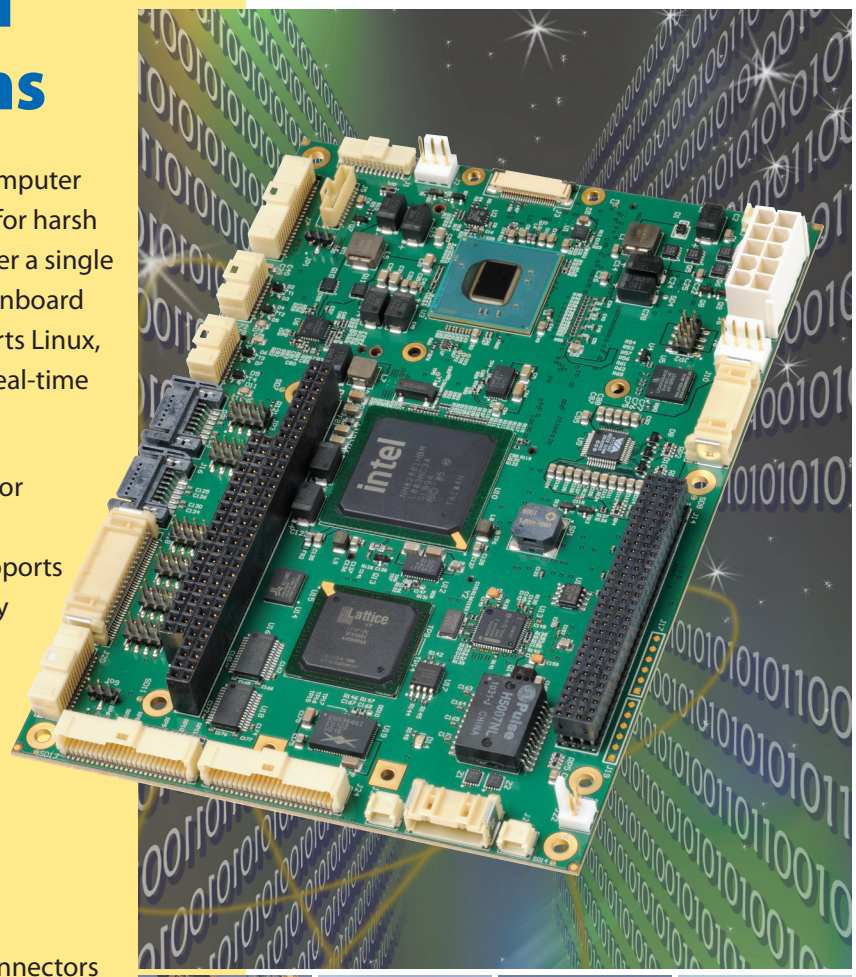

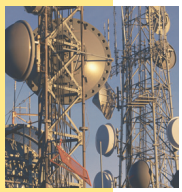

- 1.66GHz Intel® Atom™ N450 single core or D510 dual core processor available
- Embedded Gen 3.5+ GFX video core supports CRT and LVDS flat panels simultaneously
- Custom splash screen on start up
- Optional 1MB of battery backed SRAM
- Two Gigabit Ethernet ports
- Two SATA channels
- Eight USB 2.0 ports
- Four serial RS-232/422/485 channels
- 48 bi-directional digital I/O lines
- CompactFlash (CF) card supported
- MiniPCIe and PC/104-*Plus* expansion connectors
- High Definition audio supported (7.1)
- WDT, RTC, LPT, status LEDs, and beeper
- +5 volt only operation
- EPIC sized: 4.5" x 6.5" and RoHS compliant
- EBX form factor SBC also available
- Extended temperature operation
- Responsive and knowledgeable technical support
- Long-term product availability
- Quick Start Kits for software development

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

Call 817-274-7553 or Visit www.winsystems.com/EPX-C380MR

Ask about our 30-day product evaluation

WinSystems®
An Employee-Owned Company

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

ISR video enabled by rugged encoder

Engineers at Haivision in Lake Forest, IL have leveraged their commercial video encoding technology into a rugged chassis called the Makito Air for airborne and mobile Intelligence, Surveillance, and Reconnaissance (ISR) applications. Makito Air enables encoding of video and metadata transmitted over microwave (line of sight) or satellite (beyond line of sight) networks from manned and unmanned airborne platforms to operations and exploitation centers. It also has wireless transmitters to enable operation in mobile ground platforms. Makito Air is based on the same chip-level components, circuit board, and firmware as the company's Makito HD H.264 encoder.

The 3 lb. Makito Air supports constrained data links found in tough operating environments by reducing the spatial resolution of the video without affecting image aspect ratios. The encoder is housed in a ruggedized DO-160 compliant form factor with KLV/CoT metadata support. The low-latency system delivers H.264 video up to 1080p60 with metadata, encoding latency as low as 55 milliseconds, and end-to-end latency as low as 70 milliseconds when coupled with their Makito decoder. Haivision's encoders were the first to achieve JITC Certification for H.264 1080p60 HD performance for ground and airborne applications.

Haivision | www.haivision.com | www.mil-embedded.com/p371025

Three rugged displays in one chassis for extreme environments

Looking to meet needs for larger screen space in a reduced rack-mount form factor, Chassis Plans designers in San Diego created the TFX-9 trifold rack-mount LCD monitor that provides enhanced situational awareness via its three screens. The rugged monitor comes in a 2U form factor with wider environmental tolerances than previous models. The three 19-inch TFX LCD displays have a per-panel resolution of 1,280 x 1,024. The displays provide a large working space and have gas strut assisted opening and lockdown. The left/right monitor positions are adjustable.

Two models of the TFX are available – the standard TFX-19B, and the TFX-19G. Both models meet MIL-STD-810 specifications for altitude, high temperature, low temperature, humidity, transport vibration, and bench handling shock. The TFX-19G, which is a specialty model, is designed for use in extreme outdoor environments via its brightness – 1,500 cd/m^2 – and its extreme temperature operating range: 20 °C to 70 °C. It also has optional Electromagnetic Interference (EMI) protection and safety glass.

Chassis Plans | www.chassisplans.com | www.mil-embedded.com/p371026

Digital table tracks sensor data for battlefield commanders

The touch-screen and easy-to-use interfaces of digital tablets enhance situational awareness, improve communication, and enable more efficient decision making for those on the move. However, for warfighters on station, having that same functionality in a larger form factor such as a table can make more sense. Engineers at 901D provide such a system that also is rugged and can display and disseminate sensor data in harsh environments. Dubbed the Rugged Navigation and Mapping Table by 901D officials, it can be used for Computer Aided Dead Reckoning Tracing (CADRT), navigation, sand table applications, mapping, real-time intelligence, and mission planning.

The 230 lb., 46-inch TFT LCD has a 1,920 x 1,080 resolution and a 16:9 aspect ratio and can be operated in vertical and horizontal positions. 901D's digital table is 53.36 inches wide, has a depth of 43.07 inches, and measures 34.47 inches in height. The digital table runs with an embedded CPU with an Intel i7 dual core, has 2 GB RAM, 500 GB HDD storage (SSD optional), and has customizable I/O interfaces and connectors. The standard system has two 100/1000 Ethernet ports, three USB ports, and two serial ports. As many as four video feeds with embedded frame grabber and storage capability can be supported, and a video conferencing add-on module is available with amplified audio and embedded web cam.

901D | www.901d.com | www.mil-embedded.com/p371028

Secure military communications in 4 GHz spectrum

Now that the U.S. military, civilian agencies, and NATO countries have exclusive rights to the 4 GHz spectrum for mobile and fixed communications, they need solutions in this range that can operate securely in harsh environments. Engineers at Cambium Networks in Rolling Meadows, IL are meeting these requirements with two point-to-point PTP wireless broadband solutions – the PTP 45600, which operates in the 4.4 to 4.6 GHz licensed radio frequencies, and the PTP 48600 that works in the 4.7 to 5.0 frequencies. Since the 48600 uses the 4.9 band, government agencies may use the same radios for military, civilian, and NATO operations as well as collaborate with state and local agencies.

The Cambium solutions enable connectivity across long-distance line-of-sight paths and over water and through extreme weather conditions. They also support telemetry, land mobile radio backhaul, persistent awareness, and hub and other operations. Both solutions enable security for transporting data, voice, and video and for tracking users who attempt to remove, alter, or update existing hardware and software. Each comes with a FIPS 140-2 level 2 mode and has an optional module available with 1,280-bit or 256-bit AES encryption.

Cambium Networks
www.cambiumnetworks.com
www.mil-embedded.com/p371027

High-density computing for demanding environments

MIL EMBEDDED: ***What caused the demand for new high-density server solutions in the defense industry?***

SMITH: Function consolidation, virtualization, and big data analytics have driven the need for more compute capability in a smaller footprint. The DoD requires feature-rich systems that interoperate in multiple applications and allow information sharing between applications. Demand is also driven by "Common Operating Environment" requirements, the use of common components, and "right sizing" systems to deploy solutions in as many places as possible.

To support big data analytics, the DoD utilizes the Map/Reduce function initially developed by Google for search purposes and provided by Apache in Hadoop clusters. The U.S. Army utilizes Hadoop for mining sensor data in the DCGS-A program. The Army is constrained by size, weight, power consumption, and heat. Themis HD/HDS systems provide robust thermal management and double compute density with a weight savings of nearly 50 percent when compared to a 1U server stack.

MIL EMBEDDED: ***How does HD/HDS enable enterprise RAS features for embedded mission-critical systems?***

SMITH: Themis HD servers utilize the latest RAS features provided by Intel, including data and address path protection through parity and ECC for CPU and memory. These units incorporate built-in out-of-band management features for accessing system health. Through the KVM function, any network-connected client can access the console for BIOS setup, system boot, or software installation.

The fans in each server module are managed locally and are over-provisioned. In the event of fan failure, remaining fans manage the required cooling load until the failed fan is replaced. Server, storage, and power modules are hot pluggable. Front-mounted air filters protect electronic components and can be easily cleaned or replaced. When the server management controller is combined with the Themis Resource Management Appliance, full system management is at the fingertips of the IT user. When a server is attached to the appliance, the appliance automatically recognizes it and adds it to the database. The appliance acts as the system management hub to obtain health information for any network-connected client.

MIL EMBEDDED: ***How do Themis HD/HDS solutions address DoD requirements?***

SMITH: Themis HD servers deliver increased capability while allowing systems to be built up of standard, modular, lightweight, rack-mount components. Combined with a network switch and a transit case, a complete server solution can be deployed to any service region required.

The Themis HD system incorporates up to four server modules in a 2U chassis. Server modules include dual 8-core Xeon sockets, up to 256 GB of memory, dual GbE ports, and a single PCIe slot. These modules are powered by a dual redundant 1,200 W power supply and are hot pluggable.

Like the HD system, the HDS system is modular. A 2U chassis hosts up to two server-storage modules. Server-storage modules utilize the same motherboards as server modules with identical processor, memory, and network features. Each server-storage module incorporates up to four SATA 3.5-inch HDD/SSD components using 4 GB HDDs. Each HDS module provides a total storage capacity of 16 TB or 32 TB in the 2U system.

MIL EMBEDDED: ***What are the Themis server's primary Size, Weight, and Power-Cost (SWaP-C) characteristics?***

SMITH: Themis HD systems occupy 2U rack height and are 20 inches deep from front flange to the rear. The HD system weighs 43 lbs, and the HDS system weighs 50 lbs when four 3.5-inch HDDs are installed. When fully populated, the HD system power consumption is ~1,300 W and the HDS power consumption is ~750 W.

MIL EMBEDDED: ***Do HD/HDS designs enable regular technology refresh or technology insertion?***

SMITH: Yes. In all of our x86-based systems, Themis follows the Intel road map. HD systems are refreshed at the same interval. These systems enable individual module upgrades with the main chassis in place, in the rack.

MIL EMBEDDED: ***How does a stack of HD servers compete with proprietary blade servers or AdvancedTCA?***

SMITH: HD/HDS systems compare very favorably with proprietary bladed systems, as well as AdvancedTCA. The HDS system has the added advantage of including high-capacity storage. From a cost standpoint, RES-HD systems are priced closer to the 1U systems stack, and cost is significantly reduced when external storage is added. Themis RES-HD systems also incorporate thermal and kinetic management design capabilities for shock, vibration, and extended temperature.

MIL EMBEDDED: ***Where can Themis HD/HDS servers be used in applications outside of defense?***

SMITH: HD/HDS systems can be used in any application where high compute density and large, local storage are needed. Add in the robust environmental capability and they can easily be deployed in industrial or energy-exploration applications.

Dennis Smith *is the VP of Engineering at Themis. Contact him at dennis.smith@themis.com.*

OpenSystems Media

works with industry leaders to develop and publish content that educates our readers.

Check out our white papers.

http://whitepapers.opensystemsmedia.com/

Most popular topics:

AdvancedTCA
Android
Avionics Certification
Automotive
Deep Packet Inspection
GUI Linux in Medical Devices
Internet of Things
M2M
Multicore
PCI Express
Radar
SDR
Static Analysis
Switched Fabrics
Test & Measurement
UAVs

OpenSystems media®

kontron

» Highest data throughput for your application «

Latest serial interconnects via backplane on 3U CompactPCI®

» PCI Express® 3.0, Ethernet, USB 3.0, SATA 6Gb/s interconnects
» Multi-CPU implementations for multiprocessing
» PCIe over backplane for high-end graphics and visualization
» Individual services, OS integration and license management included

CompactPCI® Serial

CPS3003-SA
3rd Gen Intel® Core™ i7 CPU board

CPS3101
HDD / SSD Carrier

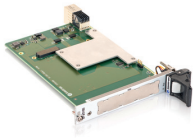

CPS310
XMC Carrier

CPS3402
Dual 10GBASE-T

CPS3410
Quad GbE Controller

CPS-ASM4 family
CompactPCI Serial Platforms

LEARN MORE AT
highspeedcompactpci.com »

Tel: 1-888-294-4558 • info@us.kontron.com • www.kontron.com

The pulse of innovation

GE
Intelligent Platforms

Leadership and experience in avionics bus protocols

Discover why more engineers continue to rely on our two decades of avionics expertise and products

For more than 20 years, GE Intelligent Platforms has been a major supporter of game-changing avionics protocols such as MIL-STD-1553 and ARINC 429 as well as newer technologies such as 10Mbit 1553 and Avionics Full Duplex Ethernet (AFDX). We continue to invest heavily in the avionics products needed for where you are today and where you'll be tomorrow. We also offer best-in-class avionics support through online knowledge bases and on-call or even on-site technical experts, as well as our Product Lifecycle Management program.

To launch your next avionics design project, download our Avionics products brochure at:
defense.ge-ip.com/av-brochure

imagination at work

Critical Recording in Any Arena

When You Can't Afford to Miss a Beat!

Introducing Pentek's expanded line of Talon™ COTS, rugged, portable and lab-based recorders. Built to capture wideband SIGINT, radar and communication signals right out-of-the-box:

- Analog RF/IF, 10 GbE, LVDS, sFPDP solutions
- Real-time sustained recording to 2.2 GB/sec
- Recording and playback operation
- Analog signal bandwidths to 1 GHz
- Shock and vibration resistant Solid State Drives
- GPS time and position stamping
- Hot-swappable storage to Windows® NTFS RAIDs
- Remote operation & multi-system synchronization
- SystemFlow® API & GUI with Signal Analyzer
- Complete documentation & lifetime support

Pentek's rugged turn-key recorders are built and tested for fast, reliable and secure operation in your environment.

Call 201-818-5900 or go to www.pentek.com/go/mestalon for your FREE online *High-Speed Recording Systems Handbook* and *Talon Recording Systems Catalog*.

TALON

PENTEK
Setting the Standard for Digital Signal Processing

25 YEARS

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