Military EMBEDDED SYSTEMS

VOLUME 6 NUMBER 3
MAY 2010

INCLUDING:

MIL-EMBEDDED.COM

GIG: Our modern Red Ball Express

Moving data, video, and ISR

Also:

Battlefield security, from the ground up

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

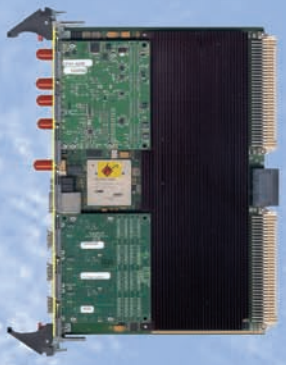

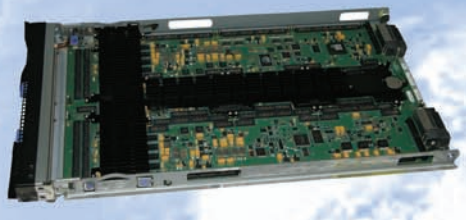

Annapolis
Micro Systems

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

D-TA SYSTEMS INC.

SENSOR INTERFACE & PROCESSING COMPANY

10 GIGABIT SENSOR PROCESSING

THE FUTURE IS NOW

DELIVERING 10 GIGABIT NETWORK ATTACHED COMPLETE SOFTWARE RADIO SOLUTIONS

Combining High-Precision RF, High-Performance Data Conversion, High-Speed FPGAs & 10 Gigabit Network Technologies, D-TA Systems is enabling you to connect sensors to your computer. D-TA products come fully integrated. Now there is no need to struggle to integrate COTS boards from multiple vendors. **D-TA PRODUCTS WORK RIGHT OUT OF THE BOX.**

DTA-2300 Synchronized Multi-Antenna Software Radio

A 16-Bit Digital IF Transceiver for up to 16 Antennas. DTA-2300 is supported by multi-channel DTA-3200 RF Front-end and DTA-5000 Record & Playback.

DTA-2210 Single Antenna Software Radio

A 16-Bit Digital IF Transceiver that costs less than typical SDR COTS boards. DTA-2210 is supported by DTA-3200L RF Front-end and DTA-5000L Record & Playback.

CALL 1-877-382-3222 FOR A SPECIAL PACKAGE OFFER FOR DTA-2210 WITH DTA-3200L TUNABLE RF AND DTA-5000L RECORD & PLAYBACK

DTA-3200L Tunable RF Transceiver

20 MHz to 1 GHz with 40 MHz BW & 70 MHz IF

DTA-5000L 10 Gigabit Record & Playback System

with 2.4 Terabyte Storage

SENSOR PROCESSORS THAT DRASTICALLY REDUCE DEPLOYMENT TIME AND COSTS

WWW.D-TA.COM

Military EMBEDDED SYSTEMS

May 2010 Volume 6 Number 3

COLUMNS

DEPARTMENTS

ON THE COVER:
During the World War II invasion of France, the Allies pressed south and eastward only by keeping the materiel flowing towards Germany. The 6x6 "deuce" trucks did the lion's share of hauling ammo, rations, and gasoline (often in the same load!). Dubbed the "Red Ball Express," these vehicles formed the backbone of Montgomery's and Patton's marches into Nazi territory. Today's military relies on the Internet as the modern backbone. Net-centric COTS standards form the battlefield's Global Information Grid (GIG). Refer to articles starting on page 24.

EVENT

www.opensystemsmedia.com/events

Sensors Expo & Conference
June 7-9, 2010 • Rosemont, IL
www.sensorsexpo.com

E-CAST

http://ecast.opensystemsmedia.com

Streamline Safety-Critical Software Development to Stop Budget and Schedule Overruns
June 8th • 1pm EDT
Presented by: LDRA, Visure Solutions

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

Stroke of Genius

Introducing the NEW Cobalt Series of Virtex-6 Modules!

Unleash your inner genius with Pentek's new 71620 Cobalt™ Virtex-6 XMC module. Built to optimize high-performance, wideband, communications, SIGINT, radar, beamforming and telemetry applications, this high-speed module includes:

- Three 200 MHz 16-bit A/Ds
- Dual 800 MHz 16-bit D/As
- Intelligent chaining DMA engines
- High dynamic range
- PCI Express x8
- Dedicated memory for each I/O stream
- Secondary x8 XMC for gigabit serial I/O
- Advanced gating and triggering modes optimized for radar applications
- Multi-channel, multi-board synchronization
- GateFlow® FPGA Design Kit
- VPX, XMC, PCI, PCIe, VME/VXS, cPCI, conduction-cooled XMC

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

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

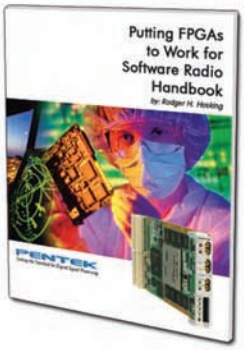

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

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

ADVERTISER INFORMATION

Page	Advertiser/Ad title
21	**Alphi Technology Corporation** – The expert in military technologies
2	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
48	**Curtiss-Wright Controls Embedded Computing** – Speed your time-to-market
3	**D-TA Systems** – 10 Gigabit sensor processing
44	**Elma Electronic** – "The industry's choice in VPX handles and panels"
33	**Elma Electronic** – Systems – 10 Gigabit Ethernet switches blade level VPX
12	**Engineering Design Team Inc.** – PCIe x8 FPGA board
13	**Excalibur Systems, Inc.** – Dragon
47	**GE Intelligent Platforms, Inc.** – Welcome to the knowledge bank
37	**InnoDisk** – Top reliability, finest solution
7	**Kontron** – We do not build military aircraft
23	**LCR Electronics** – Engineered for your application
41	**LynuxWorks, Inc.** – We won't leave you high and dry
45	**Magma** – Bridge telecom to legacy platforms
26	**Microhard Systems, Inc.** – Wireless OEM solutions
18	**Nallatech** – High performance FPGA solutions
42	**Parvus Corporation** – Qualified to perform
5	**Pentek, Inc.** – Stroke of genius
26	**Phoenix International** – Ruggedized 3U fibre channel RAID system
31	**Schroff a Brand of Pentair Electronic Packaging** – Schroff understands the demands
12	**TEWS Technologies LLC** – COTS I/O solutions
28	**Themis Computer** – Themis servers: speed to burn
17	**Trident Space & Defense** – Protect your data with Trident solid state drives
25	**Tri-M Systems Inc.** – PC/104 Can-Tainer
29	**Tri-M Systems Inc.** – 100Mhz PC/104 module
10	**White Electronic Designs** – We create space
19	**Z Microsystems, Inc.** – Rugged rack mounted servers

E-LETTER

www.mil-embedded.com/eletter

- Cloud computing offloads legacy hardware, balances loads, and puts the focus *back* on the application
  Interview with CIA veteran Bob Flores and Appistry founder Bob Lozano
- Data indexes boost embedded software's performance and efficiency
  By Steve Graves and Konstantin Knizhnik, McObject
- Static analysis improves efficiency, reduces downstream integration costs
  By Andy Chou, Coverity

OpenSystems media

Military EMBEDDED SYSTEMS

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

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

kontron

» We Do Not Build Military Aircraft «

Our Customers Do.

At the heart of many avionics and airborne computing applications are Kontron Military Rugged COTS boards and systems. Keeping us safe with smart applications, military contractors look to Kontron for superior technology, performance and life cycle management expertise.

Radar

Situational Awareness

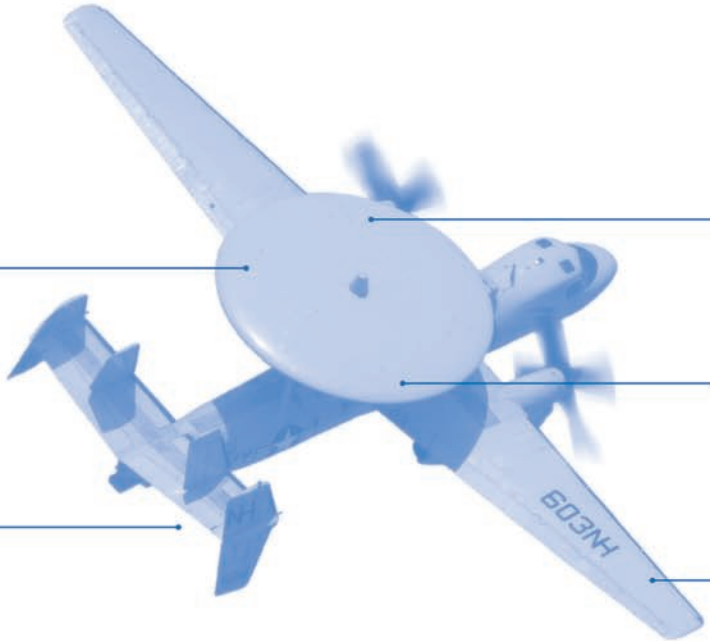

Command and Control

Satellite Communication

GPS

MILITARY RUGGED COTS

Call, Email or Visit today.

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

If it's embedded, it's Kontron.

Field Intelligence

Clustered GPUs closing in on supercomputer performance

By Duncan Young

The potential of high-end Graphical Processing Units (GPUs) to efficiently process complex, real-time sensor image and video data has been recognized for some time. The introduction of GPU-appropriate C-based computing languages such as OpenCL (Khronos Group) and CUDA (NVIDIA) gives the programmer access to arrays of processing cores within each GPU, offering typically 500 GFLOPS of performance to apply to highly parallel, data-intensive algorithms. Combining multiple GPUs into computing clusters – using industry standard, open architecture formats – will provide TFLOP performance levels for rugged, embedded sensor processing applications such as multisensor, multiplatform tactical battlefield surveillance.

The tried-and-trusted solution for complex sensor processing is a heterogeneous processing configuration based on an FPGA front end, plus a multicomputing array of vector processors, such as Freescale's 8640 Power Architecture. Using an FPGA front end is the ideal solution with its flexible high-speed I/O signals to interface to the sensor, plus it can be relatively easily applied to repetitive, parallel algorithms such as filtering and data reduction. For ultimate real-estate efficiency, multiple ASICs or FPGAs have been used to replace vector processors. However, this approach has proven to be time consuming to develop, test, and verify. Additionally, the promise of FPGA reconfigurability has never truly been realized, making overall life-cycle costs and upgradeability prohibitively expensive for any but the most generously funded projects. Much of this difficulty with portability and maintainability can be attributed to FPGA code being a specific hardware description of the problem to be solved. This code often incorporates third-party cores and logic optimizations to achieve reliable operation over temperature extremes.

Software holds key to portability

CUDA and OpenCL provide abstracted, multithreaded computing models, coded in C with extensions to suit a GPU-specific hardware architecture. As a result, the major manufacturers of GPUs, such as NVIDIA and ATI (AMD), do not need to release detailed hardware descriptions of the processor-core arrays in each device as these will change as devices evolve. Currently, each processor core is a very fast, single-precision math engine much simpler than a general-purpose processor. It has no register set as such, no predictive logic, and no pre-fetch capability, but is designed to be constantly fed with data and instructions. A typical GPU architecture uses multiples of these cores. (In the case of NVIDIA, these are referred to as *CUDA cores.*)

Eight cores form a group with common register sets and memory, known collectively as a *streaming multiprocessor*. This streaming multiprocessor runs a simple kernel to dispatch operands and organize results from the cores. Each GPU has many such multiprocessor groups. Multithreading is inherent within this type of architecture as individual cores are only required to perform a math operation on whatever data is presented to them, which can be from a different thread from cycle to cycle. Multithreading is transparent to the programmer, the compiler, and development tools dividing tasks among streaming multiprocessors to make the most efficient use of available resources.

General-purpose host CPU

OpenCL and CUDA require a general-purpose CPU to host the application and dispatch tasks to a GPU. This is typically based on the desktop PC architecture, using the GPU as a math accelerator in addition to its primary role as a graphics device. This basic CPU/GPU configuration is mirrored by many off-the-shelf, embeddable PC products in formats such as CompactPCI or VPX (VITA 46), some of which are available in ruggedized form for deployment in harsh military applications. However, by breaking the one-to-one relationship between CPU and GPU, even greater performance gains can be made by creating clusters of two or three GPUs accessible to just one CPU via a high-speed PCI Express switch. The VPX module format, in particular, provides the capability to route PCI Express offboard, through the backplane, allowing a host CPU access to a cluster of GPUs located on VPX modules within the same chassis. Such a configuration can yield TFLOPS of performance and is embeddable within any deployable environment. Depicted in Figure 1, the NPN240 from GE Intelligent Platforms adopts this architecture to pack dual NVIDIA GT240 GPUs and DDR3 local memory onto a single rugged 6U VPX module.

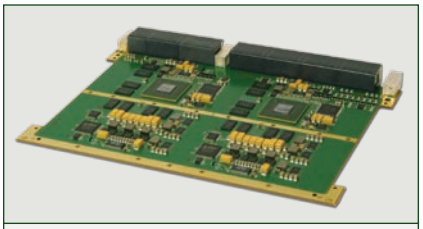

Figure 1 | The NPN240 multiprocessor from GE Intelligent Platforms

OpenCL and CUDA offer unprecedented levels of performance for the repetitive, parallel processing tasks found in advanced sensor systems. Importantly, these are software-based technologies, making them easier to develop, maintain, and port from one GPU to future generations. Looking ahead for rapid evolution, the next generations of GPU have already been announced, such as Fermi from NVIDIA. Fermi will have many more streaming multiprocessing groups, on-chip caches, 32 CUDA cores per group, more logic within each core, and full support for double precision floating-point math, suited to new high-performance, high-resolution sensor types.

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

Mil Tech Insider

Networked Attached Storage expands data storage and security on military platforms

By Steve Edwards

The pervasive use of standard network fabrics such as Ethernet on military platforms has joined with recent cost and density improvements in solid-state non-volatile memory, fueling significant growth in demand for Networked Attached Storage (NAS) devices the past couple of years. On today's military platforms, on the ground and in the air, the use of Ethernet to connect various parts of avionics architecture or crew stations, smart displays, and controls to the platform's local network has become nearly ubiquitous. System designers have found that Ethernet enables them to easily attach devices together over the network. Once networked, the platform's subsystems can take advantage of the high-speed, high-volume data storage; Space, Weight, and Power (SWaP) mitigation; and enhanced security options that NAS makes possible compared to direct attached storage alternatives.

NAS: Flexible storage in net-centric environments

Today's devices are incorporating network connectivity as system designers embrace net-centric architectures. Legacy subsystems built with older technologies – such as RS-232, 1553, and CANbus – are being bridged to Ethernet so they can communicate with other systems on the network. Now, using NAS, designers can more fully exploit the Ethernet-based LAN to centralize data storage that was formerly available only to individual devices through direct attached devices. While direct attached storage devices enabled system designers to dedicate private storage of the speed and size required by each particular subsystem, NAS provides the significant advantage of supporting easy and dynamic reallocation of available system memory. With NAS, if a new system requires storage capacity, the needed memory can be easily allocated from existing memory without incurring the additional cost, weight, and maintenance of a new storage device.

NAS benefits: Security, version control, and SWaP

Another significant advantage of NAS over direct attached storage is that NAS enables easy removal of stored data, especially classified data, so that it can be protected remotely when the platform is not in operation. Direct attached stored data is located throughout a vehicle, making it difficult to remove for security purposes. With NAS, the user can remove individual cards or the entire NAS box. To further ensure data security, today's NAS devices also incorporate encryption technology, such as 256-bit Advanced Encryption Standard (AES) encryption. Vendors such as Curtiss-Wright Controls Electronic Systems (CWCEL) are also incorporating secure-erase technology to zeroize the encryption key on the NAS device that causes the data to become declassified as it can no longer be decrypted. This approach allows terabytes of data to be sanitized within milliseconds, versus the many hours it could take to zeroize terabytes of solid-state memory data.

NAS devices ease software configuration control because the entire platform's processing elements boot from the NAS. The use of standard protocols enables all of the platform's processing systems to boot from a common storage device. The boot software and related binary programs can be loaded onto a single box, ensuring any system software updates occur simultaneously. This eliminates the risk of any single device getting updated while others are overlooked.

The need for SWaP mitigation also makes NAS attractive to military system designers. By packaging the solid-state storage into rugged, high-bandwidth 3U VPX modules within the storage device, lightweight compact NAS devices can be built. CWCEL offers the Compact Network Storage (CNS), a very compact, half-ATR-wide NAS device. It deploys one-half TB of memory on 3U VPX cards such as the VPX3-FSM (Flash Storage Module). The use of solid-state memory eliminates rotating disks' susceptibility to shock and vibration.

Solid-state memory densities rise as costs drop

It is only recently, though, that the cost of solid-state memories has made this approach economically practical. The cost of 500 GB of solid-state memory has dropped from $1 million (with 80 MBps throughput) to $35,000 (with 800 MBps throughput) in just seven years. Furthermore, that same 500 GB once required a full ATR-sized box, but today can be deployed on two 3U VPX-based cards. As memory density increases, it will be possible to deploy much more storage in the same footprint. For example, the CNS 0.5 TB is expected to increase to 1 TB by the end of the third quarter of 2010. With the higher-density networked storage made possible with NAS, system designers can capture and store more information – such as streaming data – than was possible before.

To the user, the NAS appears as a local drive on the system. Because NAS operates as a standard file server, it runs standard operating systems and major network protocols. It facilitates network architectures in military platforms in a way not possible with direct attached storage. It provides access to many network standards and communications mechanisms, bringing state-of-the-art networking into vehicles. The NAS/file server, a server designed to provide file services to clients on an Internet Protocol (IP) network, is depicted in Figure 1.

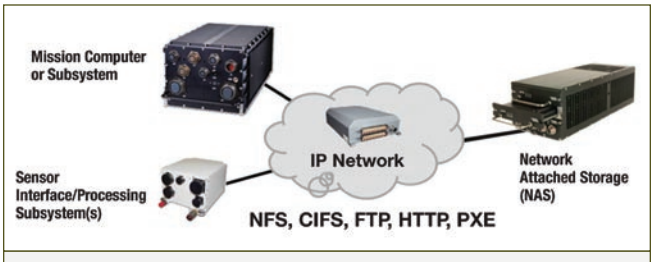

Figure 1 | The NAS/file server is designed to provide file services to clients on an IP network.

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

We create space…

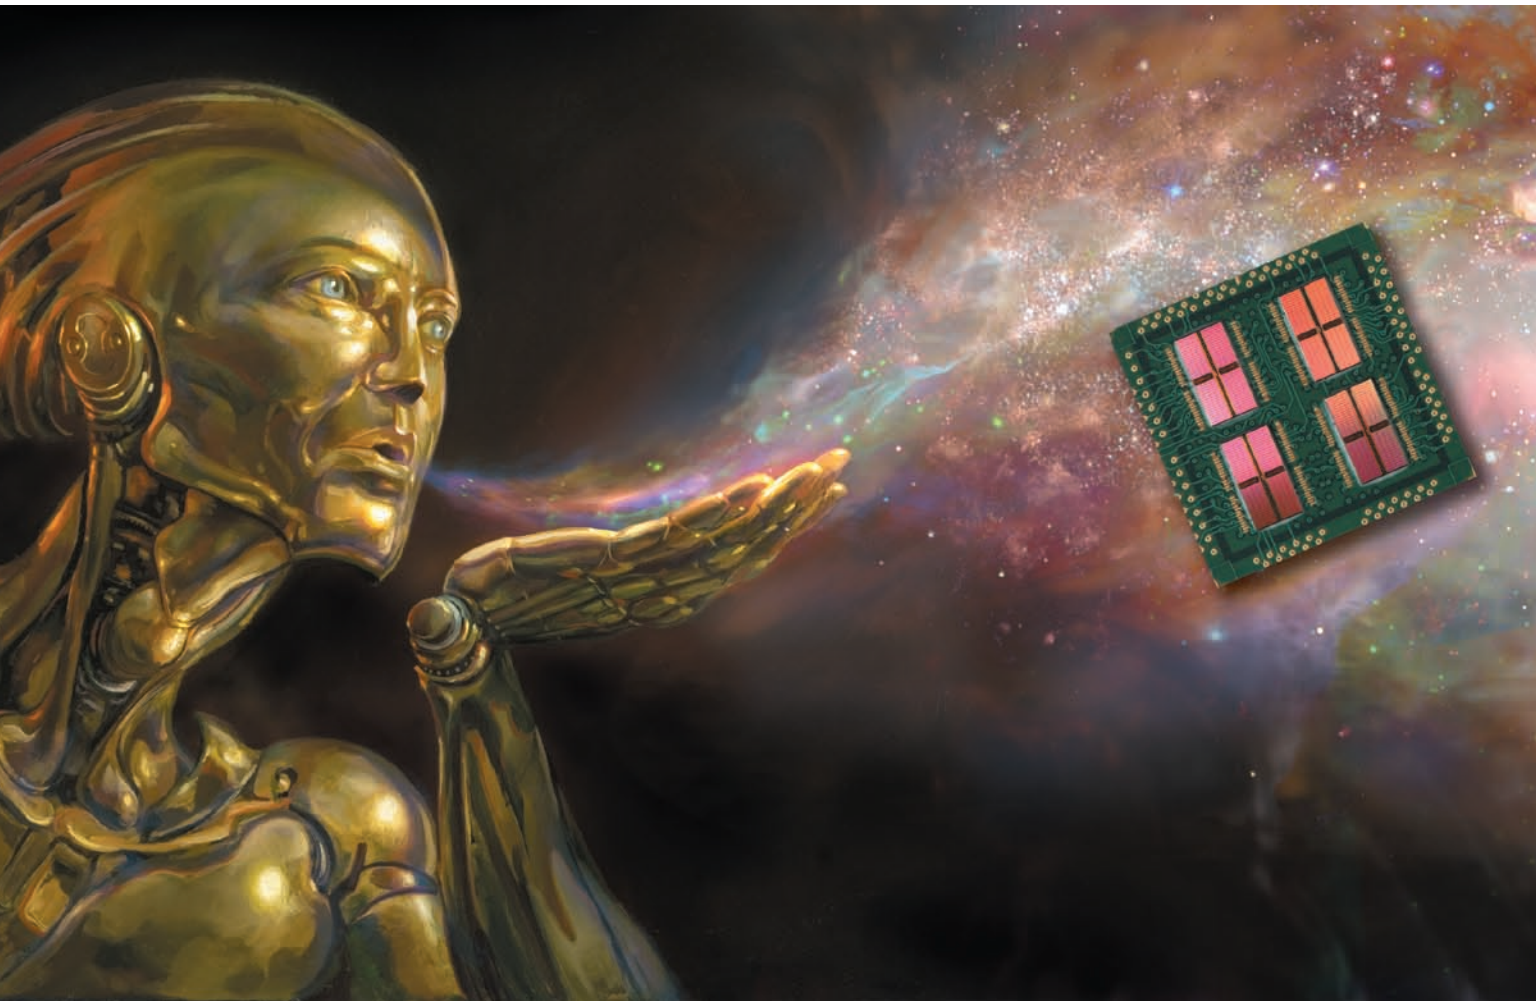

White Electronic Designs' lean microelectronic solutions increase board density, solve component incompatibilities, reduce design complexity, and extend product life and environmental performance in military applications. We offer turnkey design, assembly and test of custom multi-chip solutions, and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

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

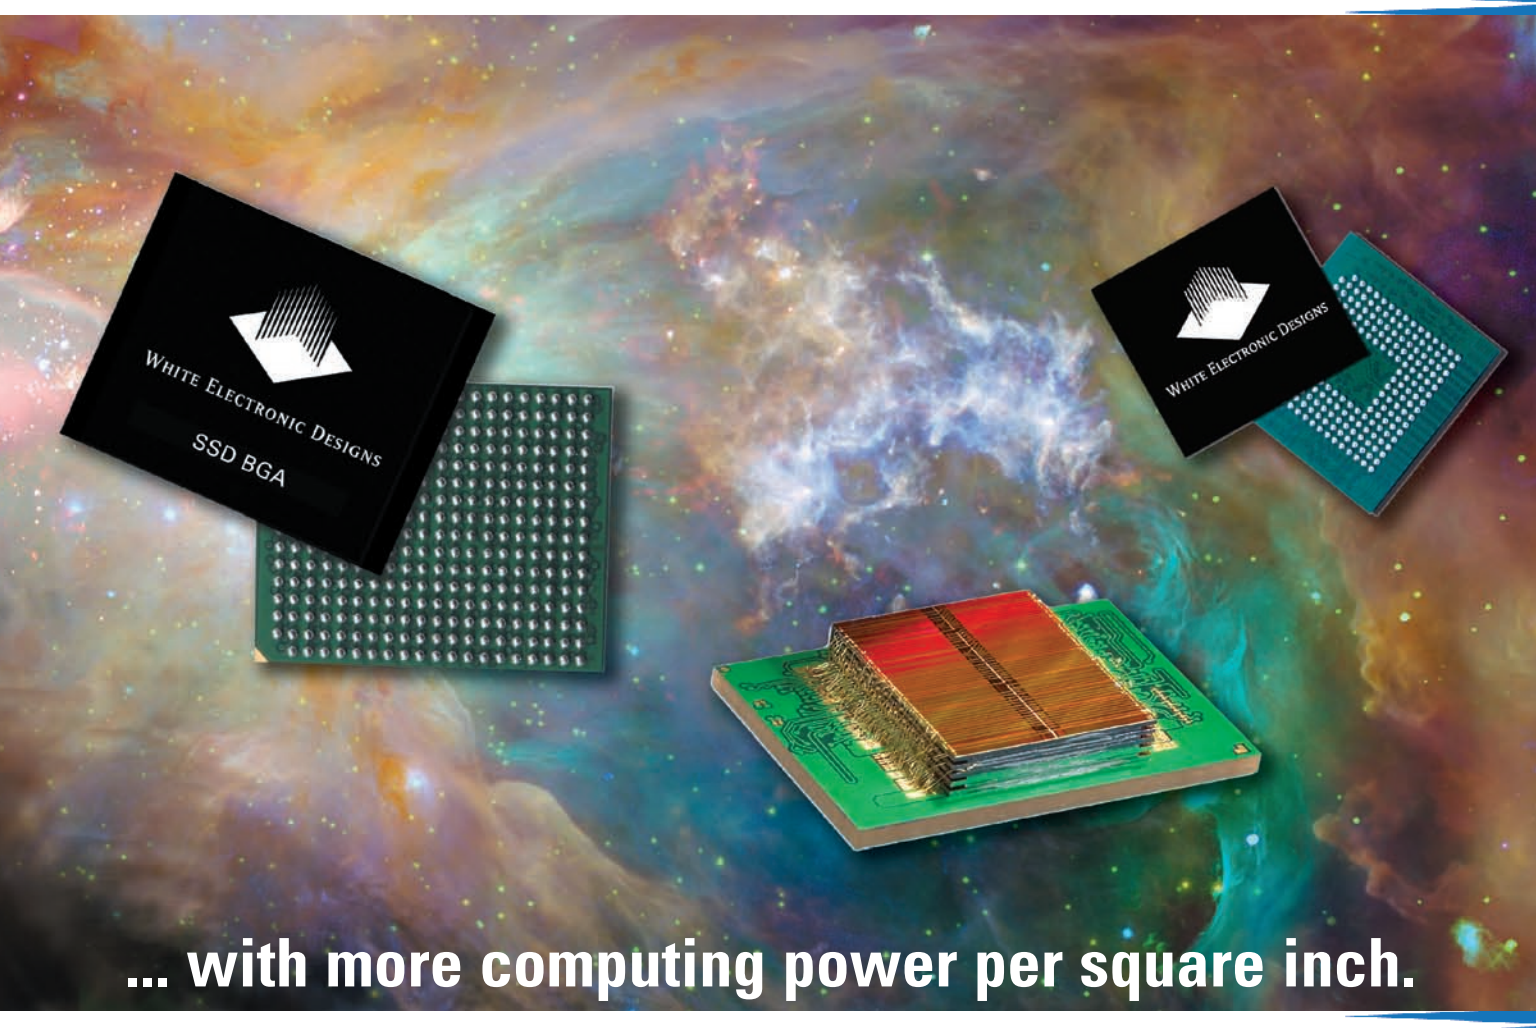

... with more computing power per square inch.

Our 4GByte SLC NAND flash SSD BGA provides big features in a miniaturized package, freeing up precious board space. It is engineered specifically for the defense market, designed to ensure data reliability in mission critical rugged and mobile systems. Constructed using a 32-bit RISC processor as its core storage controller, this SSD is also available with a hardware- and software-triggered security erase that meets military protocols.

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

602.437.1520 TEL | 602.437.9120 FAX

WHITE ELECTRONIC DESIGNS

Securing the Future Through the Power of Design

WWW.WHITEEDC.COM

Legacy Software Migration

Model Driven Architecture offers promise for migrating legacy software

By D.K. McKean

Current tool technologies based on OMG Model Driven Architecture (MDA) reduce software maintenance costs by improving software quality. MDA reduces the effort required to migrate legacy software by abstracting underlying platform technologies from software applications.

Several significant obstacles exist in migrating legacy software applications. The root cause in many cases can be traced to the overall software architecture and its constituent artifacts. Two primary culprits are software design documentation and underlying platform dependencies such as operating system, target processor, middleware, and programming language, just to name a few. OMG Model Driven Architecture (MDA) provides improvements in both of these problem areas.

Problem 1: Lacking or inaccurate software design documentation

Most legacy software developments initially produce three data repositories: requirements, design, and source code. However, at the end of most developments, the design repository loses synchronization with the source code and requirements repositories. This is because most software developments finish in a similar state: behind schedule, out of money, and with myriad software defects (a.k.a. "features"). Design repository accuracy is not considered critical to development effort completion, so it is left for the scrap heap. This occurs because many software design architectures only depict high-level views and relationships for element structures and behaviors. This problem significantly affects migration of legacy application business logic.

Problem 2: Application software underlying platform dependencies

All software applications have dependencies on the underlying platform technology. The most basic dependency is on the target computer architecture, processor(s), memory architecture, interconnection bus, peripherals, and so on. In most cases a high-level programming language, libraries, and operating system provide an abstraction for the software developer from these hardware

dependencies, creating a new set of software dependencies. Depending on the application, additional software dependencies are introduced such as a network stack, distribution middleware, database interface, user interface, security services, and so on. A replacement or significant change in any of these areas greatly affects an application's business logic even if the software is well structured.

Solution: Model Driven Architecture (MDA)

One emerging solution for both previously stated problems is the concept of the Object Management Group (OMG) Model Driven Architecture (MDA). MDA defines three model levels: Computation Independent Model (CIM), Platform Independent Model (PIM), and multiple Platform Specific Models (PSMs). The PIM models software application structure, behavior, and functionality using the Unified Modeling Language (UML), independent of the underlying platform technology. The PIM is transformed into a series of PSMs that targets specific areas of the underlying platform technology.

One example of a PIM-to-PSM transformation uses a technology known as *Executable UML,* currently supported by a number of tool vendors. Executable UML translates a PIM into source code (considered a PSM by MDA). Executable UML utilizes a subset of UML 2.x diagrams such as state charts, activity diagrams, class diagrams, and sequence diagrams used to model the PIM. A code generator analyzes the diagrams and automatically generates a source code set based on the compiler that is attached to the tool. Some of the more advanced tools allow selection of the compiler (C, C++, Java, or Ada) and compiler vendor. The more advanced tools also provide capabilities of visual debug and automated test and are integrated with a configuration management environment. It has been shown that this technology can automatically generate 75 to 80 percent of an application's business logic from UML diagrams. The set of UML diagrams that captures the software design is now used to generate the source code set.

Other PIM-to-PSM translations include UML 2.x ports, interfaces, and active objects. Ports are used to encapsulate interfaces among application business logic and areas of the underlying platform technology (middleware, operating system, and so on). Port stereotypes are used to select PIM-to-PSM translation. For example, a port could be designated as a CORBA Publisher/Subscriber for data distribution. Active objects are used to transform PIM objects to PSM threads (or tasks) in the underlying operating system.

MDA: Grist for software legacy architecture

New software technologies based on OMG MDA with currently available vendor tools offer tremendous potential to extend legacy software life expectancy. These technologies hold the promise of significant improvements in ease of maintenance, rapid integration of new capabilities, and the ability to migrate software to new platform technologies. Sounds too good to be true, doesn't it? Well, honestly, there is a price to be paid to achieve these architectural improvements. Though far from its idyllic end state, today's MDA capabilities still offer significant improvements to the current state of legacy software migration.

D.K. McKean *is CTO at Advanced Fusion Technologies. He has more than 32 years of experience as a systems engineer, software developer, and software architect, developing real-time embedded and safety-critical software. He can be contacted at david.mckean@aft-worldwide.com.*

Daily Briefing: News Snippets

By Sharon Schnakenburg-Hess, Assistant Managing Editor

www.mil-embedded.com/dailybriefing

USJFCOM releases JOE

The United States Joint Forces Command (USJFCOM) recently released its Joint Operating Environment (JOE) 2010 report. (See www.jfcom.mil/newslink/storyarchive/2010/pa031510.html to download the report.) U.S. Marines Corps General J.N. Mattis, USJFCOM commander, says JOE "in no way constitutes U.S. government policy and must necessarily be speculative in nature. It seeks to provide the Joint Force an intellectual foundation upon which we will construct the concepts to guide our future force development" in creating upcoming U.S. policies for worldwide operations (Figure 1). Regarding technology, the report predicts that key advances will continue at an explosive pace through 2030, including Electro-Magnetic Pulse (EMP) weapons, directed energy systems, laser systems, and High Powered Microwave (HPM) weapons. HPM weapons attack electrical systems, electronics, and ICs for ISR and command and control, yet remain nonlethal and nonexplosive in urban environments. Moreover, robotic systems, nanotechnology, and Nanoenergetics (NE) – touted to "dramatically increase the power and efficiency of explosives and propellants" – are also vital to the 2030 battlefield, JOE says.

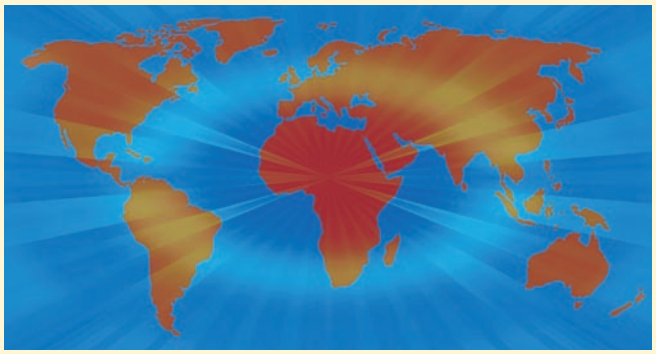

Figure 1 | The United States Joint Forces Command (USJFCOM) recently released its Joint Operating Environment (JOE) 2010 report, which is "speculative in nature" yet provides "an intellectual foundation" for future Joint Force battlefield (and other) strategies.

Partnership serves the Italian Navy

The FREMM European frigate program, a joint effort of Italy and France, will soon benefit from a new partnership between Z Microsystems and Eurolink Systems. The agreement specifies that Z Microsystems will supply the Italy-based Eurolink Systems with Z Microsystems' ZX rugged computer servers series for the FREMM program, which ultimately includes 10 frigates for the Italian Navy and 17 for the French Navy. Eurolink will then provide Z Microsystems' ZX1, ZX2, and ZX3 servers to an unidentified "global communication supplier," who will integrate the MIL-STD-810G compliant servers into FREMM's onboard computing systems. Z Microsystems' logical nomenclature designates its three "extended ATX form factor" wares as: ZX1 is 1U high, ZX2 is 2U high, and ZX3 is 3U high.

Friend or foe: The USAF will soon know

Whether in one's social life or even in business ventures, it's important to know whether one's associates are friends or foe. And it's even more important in military endeavors, where national security is at stake. Accordingly, Raytheon Company recently delivered to the USAF the first incarnation of the Identification Friend or Foe (IFF) equipment-compatible KIV-77 Mode 4/5 crypto appliqué computer. Specifically, the KIV-77 computer provides combat-identification capability to warfighters engaging in surface, land, or air combat. KIV-77 is National Security Agency (NSA) Type 1 certified, meaning that it has authority to handle classified information. Meanwhile, Mode 4 signifies legacy applications, while Mode 5 designates next-gen data links, encrypted between transponders and interrogators, to decipher whether an approaching aircraft is friend or foe. The first KIV-77's delivery was at least 60 days ahead of the contract deadline, Raytheon reports.

VME flies the (un)friendly skies

Some say VME is dead. However, the recent fulfillment of a USAF $6.5 million B-1B bomber (Figure 2) upgrade order by GE Intelligent Platforms indicates VME is still flying high. The upgrade consisted of a Vertical Situation Display Upgrade (VSDU) based on GE's 6U VME Octegra3 rugged graphics/video processor and the designed-for-Octegra3 VIM2 rugged video input mezzanine. Since the "goods" were delivered early, prime contractor Boeing has a large window of time before VSDU's flight testing, slated for early 2011. VSDU aims to enforce B-1B aircrew flight safety by providing: 1) a DO-178B certified Board Support Package (BSP), ensuring flight worthiness, thanks to GE subcontractor Ultra Electronics Controls; and 2) "improved protection against hostile action" for the aircrew.

Figure 2 | B-1B Lancer, U.S. Air Force photo by Airman 1st Class Corey Hook

Java-enabled combat system passes the test

Modernization is the name of the game for many military programs, and the USS Bunker Hill (CG 52) is a good case in point (Figure 3). Now suited with a Lockheed Martin-developed Aegis Weapon System, CG 52 recently finished its full combat systems' operational trial. Key to Aegis' operation is Atego's Aonix Java-based PERC Ultra virtual machine technology, which renders real-time, deterministic performance and provides virtual machine management and instrumentation tools to fulfill mission-critical needs. Meanwhile, the Aegis Weapon System resides on 92 ships presently in military service worldwide – and is sported by maritime vessels owned by Japan, South Korea, Australia, Norway, and Spain.

*[**Editor's note:** As* Military Embedded Systems *went to press, Atego announced the availability of PERC Ultra SMP, a variation of PERC Ultra used in the Aegis system but designed for multicore processor systems. Concurrent Java "garbage collection" doesn't require stopping the entire system for routine Java housekeeping.]*

Figure 3 | USS Bunker Hill (CG 52), now sporting an Aegis Weapon System, recently finished its full combat systems' operational trial.

A flurry of contracts appears

Curtiss-Wright (CW) has been busy lately, and is about to get even busier, thanks to three recently announced defense contracts: 1) CW will provide an "approximately $25 million" Advanced Mission Management System (AMMS) to prime contractor Northrop Grumman as part of the U.S. Navy's Broad Area Maritime Surveillance Unmanned Aircraft System (BAMS UAS). BAMS UAS proffers "persistent" ISR for identification, classification, detection, and tracking of littoral and maritime targets. CW will design and manufacture BAMS UAS AMMS units at its Santa Clarita, CA Motion Control facility with deliveries starting at the end of this year and continuing through next year. 2) A second contract between Northrop Grumman and CW for $10.5 million stipulates that CW produces an upgraded Radar Signal Processing (RSP) ware for the Joint Surveillance and Target Attack Radar System (Joint STARS) program. Work will be performed at CW's Motion Control divison in San Diego, CA and is one segment of an upgrade to the USAF's E-8 Joint STARS aircraft's Radar Airborne Signal Processor (RASP) radar signal processing system. 3) And finally, a $1.17 million contract stipulates that CW provides its rugged, high-altitude capable SANbric Storage Area Network (SAN) to prime MDA in Richmond, BC, Canada. The SAN will then take flight on the CP-140, a Canadian Forces Air Command's maritime patrol aircraft. CW's Littleton, MA Electronic Systems facility will produce the MDA SANbric units.

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

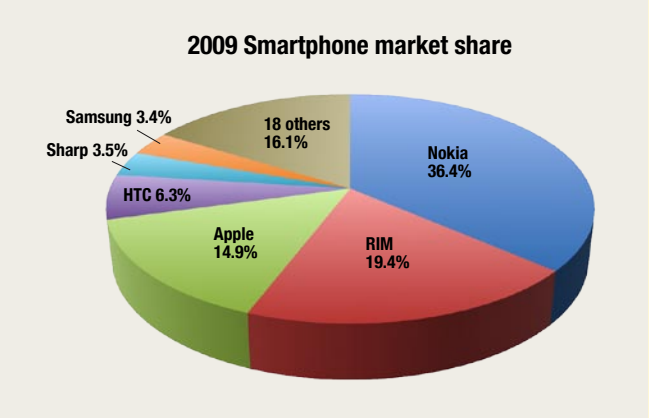

Figure 4 | Top 2009 smartphone vendors

More and more people are getting "smart"

Smartphones are of wide-ranging interest, both to the gadgetized consumer and the military commander, who thinks "if only I could check a sit rep from my iPhone." Smartphones are also of interest to Forward Concepts, who recently published a 470-page study of related trends, entitled "Smartphone Device & Chip Market Opportunities '10". The report indicates that 2009 boasted an 18% growth in worldwide smartphone shipments, rising to $67 billion and 171 million units. Smartphone mobile Internet consumption also rose by 29% in 2009 as compared to 2008. Forward Concepts forecasts that 2010 market share will feature North America leading in smartphone consumption at 22%, followed by Western Europe at 21.6% and China at 17%. Meanwhile, the top 2009 contender for smartphone market share was Nokia, followed by RIM, Apple, and others (Figure 4).

Predator climbs to 1 million

Military tech trend reports indicate that Unmanned Aircraft Systems (UASs) are increasingly being used on the modern battlefield, with flexibility in configuration and the ability to spare human lives as the most critical factors. This technology shift has been validated by a recent milestone: The Predator series of UASs recently reached 1 million flight hours. Trend statistics leading up to the milestone include annual Predator flight hours increasing from 80,000 hours in 2006 to 295,000 last year. The Predator MQ-1B (Figure 5) surpassed 700,000 flight hours this past March, with the remainder of the Predator series marking the other 300,000. Meanwhile, the Predator series includes the B/MQ-9 Reaper, C Avenger, Sky Warrior, and others on active duty with the U.S. Department of Homeland Security, U.S. Army, USAF, U.S. Navy, NASA, the U.K.'s Royal Air Force, and others.

Figure 5 | The Predator MQ-1B, U.S. Air Force photo by Tech. Sgt. Sabrina Johnson

U.S. Army photo by Pfc. Sarah De Boise

POSIX RTOS attacks SWaP and time-to-deployment issues

By Kim Rowe

Infantry are limited by the Size, Weight, and Power (SWaP) capabilities of their equipment. The next generation of handheld and portable kits can exploit COTS Microcontroller Unit (MCU) technology to vastly reduce the cost of producing small, lightweight, ultra-low power gear. Adding an ultra-tiny POSIX RTOS enables the immediate introduction of scores of standard applications and rapid adaptation of equipment to dynamic battlefield conditions, maximizing mission success.

Choosing a standard set of SoC Microcontroller Unit (MCU)-based hardware and a standards-based RTOS opens the door to fast and flexible development of tools and sensors for use in the battlefield. Millions of lines of Linux code are available, and now they can be applied to help warfighters quickly and easily – and without significant SWaP requirements. This is a lean product development approach that will minimize Total Cost of Ownership (TCO).

The key to reducing SWaP for portable electronics is:

- Minimizing total parts count
- Building with die instead of packaged components
- Utilizing multichip modules
- Exploiting consumer and automotive parts
- Exploiting IP
- Utilizing a POSIX-based RTOS

Using this approach, existing applications can be quickly shrunk onto COTS SoC MCUs, saving money, time, power, and weight.

How can a POSIX RTOS make such a difference to the hardware design? By allowing the use of SoC MCUs to replace power-hungry discrete implementations, electronics' operating life can be increased significantly while reducing size, weight, power, and cost. At the same time, the POSIX compatibility allows rapid injection of standard applications into devices, ensuring that software capabilities are not compromised. This discussion explores how modern MCUs and power electronics increase SWaP, and where software and standards-based RTOSs fit in. An ultra-tiny POSIX RTOS example is also presented.

Reducing power and weight increases survival

A few years ago, MCUs were tiny eight-bit machines with very limited memory that were clocked at a few MHz. Today, that has all changed. SoC MCUs run up to 200 MHz and have more than 2 MB of flash, 128 KB of RAM, and all peripherals onboard. High-performance floating-point hardware is available on SoC MCUs, speeding DSP applications, video processing, and automatic detection algorithms. The vastly increased computing power of these devices in tiny packages allows battlefield computers to shrink from notebook sizes to handheld devices with corresponding savings in weight.

The power consumption of MCUs has also dropped drastically. As little as one year ago, common thinking would have an 8-bit processor consuming less power than a 32-bit CPU. Today, using a 32-bit processor to replace the 8-bit processor and powering the processor down when it is idle is more effective. Today's 32-bit MCUs use a stunning 85 microwatts/MHz. On the battlefield, this means increased operational time, greater mobility, and less weight.

Power electronics that enable MCUs have also shown significant improvement. Today Field Effect Transistors (FETs) run at 900 V, providing better switching characteristics at higher voltages and reducing the power lost during switching compared to Insulated Gate Bipolar Transistors (IGBTs). Using soft switching techniques to minimize IGBT losses (1/3 factor) makes this trade-off

very application dependent. High-voltage FETs offer fewer components while IGBTs offer smaller packages. Either approach means less weight, improved reliability, and greater mobility – increasing both effectiveness and survival.

Software and standards-based RTOSs are key

Previously, MCU software was 20 percent of the problem, and the main issues were getting the hardware to work effectively. Now, 80 percent of the problem in developing an MCU system is software. The SoC MCU represents a complete system on a chip with analog and digital I/O, file systems and databases, network connectivity, and user interfaces. Further complexity comes as these MCUs are networked together to solve larger problems using less power, with less weight and greater design reliability. A portable solar converter, for example, would have 1/6 of the power consumption, 1/10 of the weight, and greater than six times the reliability for the control electronics.

The approach used to implement the software for these systems is also changing. The four approaches that users typically select are:

- Single loop of control or timer-based scheduler
- Small scheduling kernel
- RTOS (kernel plus I/O)
- Standards-based RTOS (POSIX)

The confusing part for designers is that each of these approaches can be effective depending on how complex the problem is. By minimizing the TCO for the software development and maintenance over the product line's lifetime, substantial savings are generated that are not readily apparent to traditional engineers. Figure 1 shows these four approaches and when they are most effectively used.

> "Previously, MCU software was 20 percent of the problem, and the main issues were getting the hardware to work effectively. Now, 80 percent of the problem in developing an MCU system is software."

If systems are trivial, such as those that monitor a single analog value and flag an out-of-range condition, the old-school single loop of control is best. As complexity grows, though, the advantages of having a separate stack for each thread, thread priorities, standard APIs, and a standard I/O model grows. For complex applications, a standards-based RTOS is the only way to effectively solve the problem. Proprietary

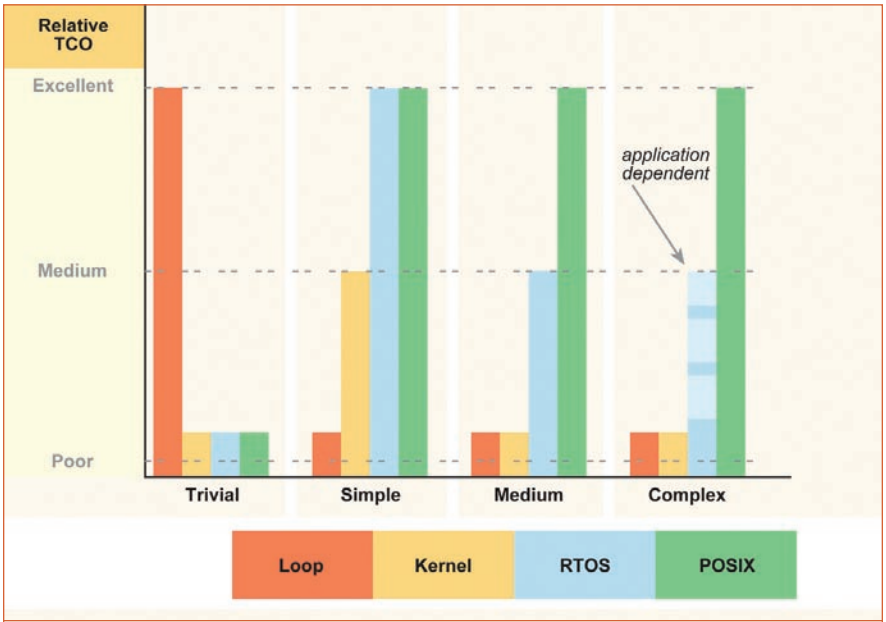

Figure 1 | A single loop of control or a simple scheduler was acceptable in the past; however, with increased complexity, MCU systems require standards-based RTOS products with commercial support.

kernels require invention of I/O models and driver development while complete proprietary RTOS solutions eliminate software reuse, cause training problems, slow field deployment, and can limit real-time response.

An alternative approach to a proprietary RTOS is to utilize a POSIX-based RTOS. However, many of the benefits of a POSIX RTOS are hidden. Application portability, elimination of training, real-time response, and reduced time to deployment in theater are required to achieve the optimum savings over the life cycle of a product line. POSIX delivers these variables, while proprietary RTOSs and other standards do not.

As developers of military systems, we must be cognizant of the immediate demands of warfighters in the field by choosing a POSIX platform application that can be immediately and easily adapted to fielded systems. On the surface, adding a podcast feature to a military radio using off-the-shelf software is a nonessential feature; however, if it is used to listen to messages from families far away, its morale boost could be substantial. Being able to quickly add features to systems using off-the-shelf software is a key feature to meet new battlefield demands.

POSIX and Linux-compatible RTOSs are optimal

As mentioned, use of a standards-based RTOS is optimal for other than nontrivial systems. By making the selected standard POSIX, substantially more benefits are accrued to the user by virtue of available applications, number of knowledgeable programmers, and training savings. The POSIX-based, Linux-compatible, 32-bit ultra tiny Unison POSIX RTOS (Figure 2), for example, utilizes the same architectural design as its nano-kernel 8/16-bit version: DSPnano.

By moving to a nano-kernel architecture, the size of the system is kept very small. Unlike other Linux variants, the tiny POSIX RTOS and its nano-kernel version can run in 2 KB RAM and 10 KB flash for an entire application including I/O. As more features are used, the kernel grows. The modularity of the nano-kernel architecture contributes substantially to the reduction in size of the system. In addition, elimination of memory management support, lightning fast interrupt handling, and embedded safety features make the POSIX RTOS highly effective on SoC MCU architectures. With its tiny memory footprint and POSIX APIs, the POSIX RTOS allows many applications to run on tiny, low-power, and low-weight SoC MCU systems where it was not possible before.

Thus, the complexity of the system is significantly reduced using the nano-kernel with a layered POSIX and Linux-compatible I/O layer in comparison to any Linux alternatives. Any engineer can understand the basic kernel and I/O model in a few minutes due to its modular architecture and isolated interfaces. Servers (or drivers plus software glue) can be added for device specific uses quickly and easily. Additionally, if a specific algorithm has a limitation in theater, adaptation can be fast and easy.

Flexible devices save lives

Using lean product development for military hardware based on SoC MCU technologies, new power technologies, and an ultra-tiny, POSIX-based RTOS or

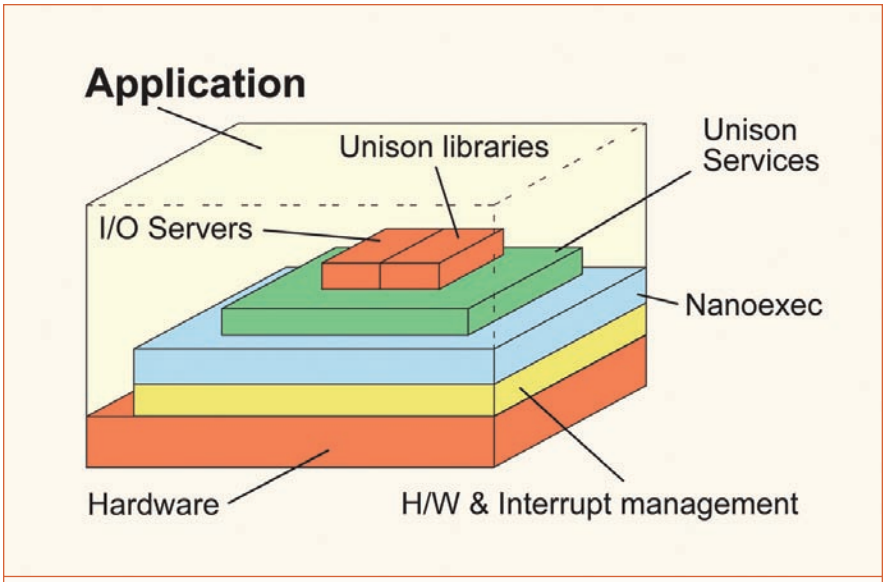

Figure 2 | The ultra-tiny Unison and nano-kernel DSPnano offer a nano-kernel architecture with an ultra-tiny memory footprint, POSIX APIs, and a modular structure.

nano-kernel RTOS increases flexibility and reduces total cost of ownership. This benefits the manufacturer and the military. Additionally, the ability to adapt an RTOS quickly in theater – coupled with lowered power consumption, smaller sizes, and reduced weight – maximizes survivability for the troops in the field.

Kim Rowe *is CTO and founder of RoweBots Research Inc. and has more than 30 years of experience in business management and systems engineering. He has been instrumental in the startup of several companies and several business units in the computer systems and services areas. He has extensive international experience, having taken a broad set of software and hardware products to market in more than 20 countries. He has published approximately 35 papers and articles in various journals and magazines. Kim holds both an MBA from the University of Ottawa and an MEng from Carleton University. He can be contacted at pkr@rowebots.com*

RoweBots Research Inc.
519-208-0189
www.rowebots.com

FPGAs balance architecture, IP, power in electro-optical/infrared systems

By Dr. Tibor Kozek, Juju Joyce, and Suhel Dhanani

The MQ-1 Predator carries the Multispectral Targeting System with AGM-114 Hellfire missile targeting capability and integrates electro-optical, infrared, laser designator and laser illuminator into a single sensor package. U.S. Air Force photo/Airman 1st Class Jonathan Steffen

Modern Electro-Optical/Infrared (EO/IR) systems have become increasingly and exceedingly complex, and therefore demand processing capabilities best offered by FPGAs. The most critical design challenge in these systems is that of combining high-performance sensor/video processing with low power consumption. To help solve this dilemma, key sensor-processing and video-processing algorithms – and how they can be implemented on FPGAs – are presented.

Portability and versatility, coupled with leading-edge COTS technology, characterize many modern Aerospace and Defense (A&D) sensor platforms. Whether mounted on Unmanned Aerial Systems (UASs), manpacks, or left behind as autonomous sensors, Electro-Optical/Infrared (EO/IR) systems have become exceedingly complex and demand processing capabilities best offered by FPGAs. Power constraints have also tightened as a function of available energy and heat dissipation in low-footprint platforms.

These military imaging systems have become increasingly more sophisticated and incorporate multiple advanced sensors – ranging from thermal infrared, to visible spectrum, to even UV focal planes. Not only do these sensor outputs have to all be corrected (defective pixel correction and color correction) and interpolated, but images from multiple sensors must be fused, overlaid, and further processed for local display and/or for transmission on the battlefield. The key design challenge in these systems is to combine high-performance sensor/video processing with low power consumption. The following discussion focuses on some of the key sensor-processing and video-processing algorithms and how these can be implemented on FPGAs. And because an FPGA-based design reduces component count while adding flexibility, system power is reduced.

Typical sensor processing system

FPGAs are the platform of choice for almost all state-of-the-art EO/IR systems – since they meet the need for requirements for programmability, high-performance sensor/video processing, and low power consumption.

In fact, each new generation of low-power FPGAs features significantly lowers static and dynamic power consumption by utilizing a combination of architectural enhancements and lower core voltages, coupled with geometric advantages resulting from shrinking silicon feature sizes.

Sensor processing

Output from image sensors used in EO/IR systems needs correction by using algorithms such as those for non-uniformity correction and pixel replacement. While these algorithms typically require only a few mathematical operations per pixel, calculations need to be done at the pixel rate and with data that might be different for every pixel. In this case, an FPGA is an ideal platform because of the inherent parallelism in the architecture, as well as the ready availability of algorithm IP to realize video-processing functions.

For non-uniformity correction, pixel-specific coefficients need to be streamed into the logic block implementing the correction formulas. For smaller sensors, these coefficients can be stored internally in FPGA memory. Depending upon the resolution of the sensor and precision of the coefficients, the memory requirement will vary. For larger sensors, this data needs to be buffered in external memory and read out in sync with the pixel stream for every video frame. In either case, there is typically a need to change the corrective data set based on some selected parameter such as the Focal Point Array (FPA) temperature that varies greatly from ground to at-altitude.

A typical data flow for such an algorithm would double-buffer the correction coefficients to allow the relatively slow data stream from flash memory to complete before a new set of data is applied. One of the key FPGA advantages is the ability to create a data pathway that fits the algorithm, rather than change the algorithm to fit a predefined architecture. This is critical for achieving low power consumption. The massive I/O and large number of logic elements in the FPGA allow easy implementation of parallelism, and off-the-shelf IP algorithms facilitate a straightforward design. Some of the more typical sensor-processing algorithms that are available for FPGAs such as those in Altera's Cyclone family are shown in Table 1.

Sensor-processing algorithms for FPGAs
Digital zoom / binning
Noise filtering
Non-uniformity correction
Wide dynamic range processing
Local-area adaptive contrast enhancement
Pixel-level adaptive image fusion
Electronic image stabilization
Super-resolution
Motion detection and multi-target tracking

Table 1 | Typical sensor processing algorithms available for FPGA devices such as Altera's Cyclone family.

Another way that FPGAs reduce power in EO/IR sensor systems is via a dramatically lower footprint. For example, Figure 1 shows an Altera Cyclone FPGA-based system that implements a thermal sensor with integrated processing. The FPGA performs real-time image enhancement, image stabilization, and digitally enhanced resolution – and can also drive an integrated microdisplay. The FPGA power consumption in this case is ~500 mW.

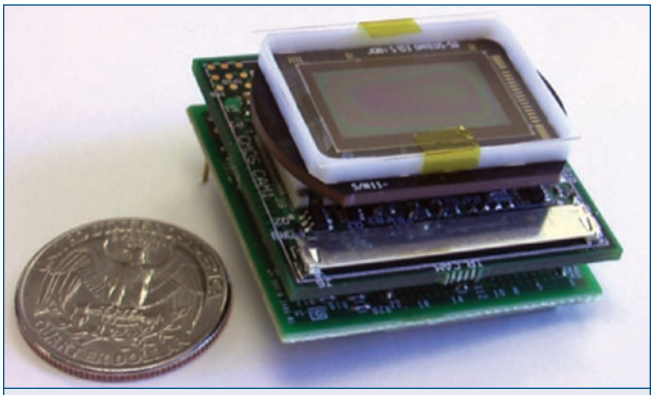

Figure 1 | A thermal sensor system with integrated image processing.

Video processing: Mixing and scaling

Military EO/IR systems often include multiple image sensors with outputs that must be fused together and displayed on a custom display with non-standard resolution. ("Non-standard" implies resolutions different from the typical desktop or laptop LCD.) A video processing system can be used to generate a composite image from two video sources on a custom display. And a video data path inside an FPGA can generate a composite image from multiple sources.

The input video is first formatted into the desired color space and is then subsequently scaled (resized) and mixed (alpha blended) with multiple other video streams. Scaling and mixing are among the most commonly used video functions, and they can be realized using off-the-shelf IP algorithms available for FPGAs.

Scaling can be as simple as copying the previous pixel (or dropping it), or it can be implemented with complex interpolation filtering techniques to generate a new pixel. Figure 2 shows the difference between the different algorithms that can be used for scaling.

The graphic illustrates all the generated pixels (shown in solid black) versus all the original pixels (shown in white). There are many ways of generating the solid black pixels; for instance, the nearest neighbor algorithm copies the preceding pixel. A more complex way would be to take an average of the two neighboring pixels in both vertical and horizontal dimension. Sometimes this is known as *bilinear scaling* – bilinear because it uses a pixel array of size 2x2 to compute the value of a single pixel.

Taking this concept further, one can compute the new pixel by using "m" pixels in the horizontal dimension and "n" pixels in the vertical dimension. Figure 2 additionally shows how a pixel is generated using four pixels in each dimension – also called a *four-tap scaling engine.*

Of course, the trick deals with the weights assigned to each pixel – also called the *coefficients* when realized algorithmically. The coefficients will determine the quality of the scaled image.

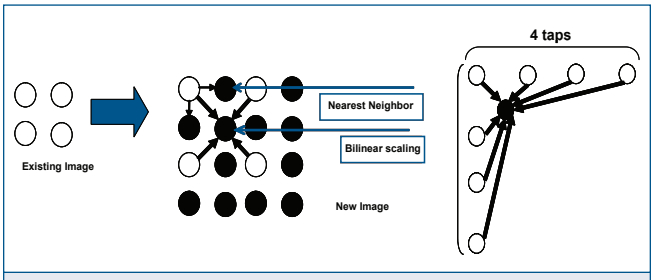

Figure 2 | Different video scaling algorithms and their implementation with a video scaling IP

Figure 2 depicts one such scaling IP core available for an Altera Cyclone FPGA. This function comes prebuilt with various "Lanczos" filter functions. The Lanczos multivariate interpolation method is used to compute new values for any digitally sampled data. When used for scaling digital images, the Lanczos function indicates which pixels and in which proportion in the original image make up each pixel of the final image.

Selecting from a range of Lanczos algorithms to scale the image or bypassing them completely in favor of custom coefficients is also possible. In either case, the function automates the tedious job of generating HDL code for what is essentially a two-dimensional filter. It also maps it to the various FPGA structures such as the DSP blocks and the embedded memory blocks, thus improving productivity and reducing overall design time.

Video overlays

The other commonly utilized function is mixing and overlay of two or more video streams. This is generally done by an alpha blending function. This is a way of generating a composite pixel from two or more pixels. One pixel is assigned an opacity value called the *alpha.* When alpha is zero, that pixel is completely transparent (that is: not displayed). When that same alpha value is 1, the pixel is completely opaque – only that pixel is seen and the other pixel is not displayed.

In mathematical terms, the value of the composite pixel is calculated as:

$$C = \alpha P1 + (1-\alpha)P2$$

Where

- α is the alpha value
- P1 is the pixel 1 from the video layer 1
- P2 is the pixel 2 from the video layer 2
- C is the composite pixel

The same technique can be used to create translucent images because the alpha value can be set anywhere between 0 and 1.

A more sophisticated way of combining information from two (or more) images is to utilize image-fusion algorithms. Imagine a thermal infrared sensor and a visible image sensor depicting the same scene, but each contains information in different parts of the image.

If alpha blending could be applied to select how dominant one or the other image should be in the combined output, there would be no single value for alpha that would extract all the information available from the sensors. Figure 3 shows such a scenario.

While the visible image on the top left contains information about the surroundings, the thermal view on the right only shows discernible features where there is a temperature difference against the background. Conversely, the performance of the thermal sensor is unaffected by the strong light source in the scene, while the visible camera provides no information in the same area as a result of saturation.

In the fused view, a seamless combination of information from both input modalities is achieved on a pixel-by-pixel basis.

Figure 3 | Sensor fusion allows simultaneous utilization of multiple EO sensors

Details missing from one modality are "filled in" from the other and vice versa. One of the simplest forms of fusion would be to apply alpha blending with a different alpha value for every pixel calculated from local image statistics. State-of-the-art fusion algorithms, however, typically go beyond that and perform a decomposition of the input images that extracts relevant features around every pixel. These features are then combined to form a fused image.

Dr. Tibor Kozek *is cofounder/Chief Technology Officer of Imagize LLC and has nearly 20 years of experience in signal and image processing. Prior to Imagize, he was chief scientist with Teraops, Inc., as well as a visiting scholar at the University of California, Berkeley.*

Juju Joyce *is Senior Strategic Marketing Engineer, Military & Aerospace Business Unit, at Altera Corporation and has more than 10 years of semiconductor industry experience. He holds a Bachelor's degree in Electrical Engineering from the University of Texas at Austin.*

Suhel Dhanani *is Senior Manager, DSP, for Altera Corporation's software, embedded, and DSP marketing group. He has more than 15 years of industry experience with Xilinx, VLSI Technology, Anadigm, and Tabula. He holds M.S.E.E. and M.B.A. degrees from Arizona State University.*

Altera Corporation
408-544-7000
newsroom@altera.com • www.altera.com

Telecom equipment nets big benefits on the battlefield

Q&A with Anthony Ambrose, VP and GM of RadiSys

EDITOR'S NOTE

Editor Chris Ciufo jumped at the chance to learn how RadiSys, one of the larger COTS telecom companies serving the embedded space, is finding homes for its technology in Aerospace and Defense (A&D). RadiSys is being pulled into the market by DoD platforms demanding network-centric equipment in semi-rugged deployments. Edited excerpts follow.

MIL EMBEDDED: ***So can you start by familiarizing our readers with RadiSys?***

AMBROSE: Sure. RadiSys has been in business for 22 years. For the first 12 to 15 years of our life, the customer would sketch what they wanted on a napkin, and we would design and deliver it to great results. And, then really after the telecom bubble of 2001, the whole dynamics of the industry changed, and not just in telecom. So it was pretty clear that we needed to go to more of a standard embedded product portfolio, targeted to customers who cannot have the products fail.

MIL EMBEDDED: ***Where does the military focus come in?***

AMBROSE: About a year or 18 months ago, we said, "We're not really targeting the military space, so why are military customers choosing our products?" We talked to some experts in the industry, and they came back and said, "Look, RadiSys, you've got some great products here. You've got about 90 percent of what the military and aerospace market already needs. You just have to make a few simple changes in some of your basic processes and basic marketing, highlight some of the things that you're already doing, and you'll have a very good mil/aero solution including ATCA and COM Express."

And it turns out that for applications like ground stations or command and control – or anything needing extremely high compute performance – ATCA is a superior solution, far in excess of what VME or CompactPCI or even VPX would deliver. And ATCA has been on the market for six or seven years so we have learned a lot and earned a high level of trust from our customers. Our ATCA products are very much field-proven.

MIL EMBEDDED: ***Skeptics might say, "Those comm guys don't speak our language, and we don't understand those telecom standards like MicroTCA or ATCA."***

> "... For applications like ground stations or command and control – or anything needing extremely high compute performance – ATCA is a superior solution, far in excess of what VME or CompactPCI or even VPX would deliver."

AMBROSE: You know, I'd be surprised if many were saying that. The reason is that ATCA has adopted almost the exact opposite approach of MicroTCA. And aside from having three letters in common, there's really not much commonality between the two. ATCA is field proven and deployed, and it's been discovered by customers as very useful for the military and aerospace markets. ATCA was "stealth" technology, meaning it had no marketing energy behind it in the military and aerospace market. So it's been selected with no hype, with no buzz, with no top-down standard because it's the best product, period.

MIL EMBEDDED: ***So will we be seeing more of RadiSys targeting the military market?***

AMBROSE: Yes, we're absolutely targeting more for military and aerospace markets. It's a pretty straightforward strategy. We've selected products that have traction in these segments already, such as ATCA for C4ISR and information assurance and COM Express for ruggedized computers, robots, UAVs, and UGVs.

MIL EMBEDDED: ***Which technologies are you bringing to the military arena that perhaps weren't there before, given your telecom background?***

AMBROSE: We've had an opportunity to work with several of our technology partners to develop a ruggedized chassis, and we've come up with a conformal coating capability for ATCA. So now we can provide a ruggedized ATCA chassis, which expands the market even further.

MIL EMBEDDED: ***I've not heard of a ruggedized ATCA chassis. So let's talk boards. A typical benign commercial temp range would be 0 to +50 °C, or maybe 0 to +70 °C. Are your boards rated for a wider temperature range than that?***

AMBROSE: Yes, we deliver a wider temperature range than that, specifically on the COM Express. ***[Editor's note:*** *RadiSys'*

Intel Atom-based Procelerant CEZ5XT and CEGSXT COM Express modules both boast an operating temp range of -25 °C to +70 °C.] Another point to bring up is that a lot of people use a simple screen when testing; that is, they test a product one time in the factory then if that happens to work at that lower temperature, they certify it as having some extended temperature properties. However, we utilize both HALT and HASS test methodologies. This ensures that a product works to the extreme of its temperature range; unlike other products certified based on a single test, it doesn't leave the factory until we see those results.

MIL EMBEDDED: ***You've described your company's ATCA products several times as "field proven," which might be analogous to "mission critical." But isn't an ATCA board just a commodity?***

AMBROSE: Well, I would characterize ATCA as an open standard, not a commodity. But even working with an open standard doesn't mean that everyone has chosen to invest at the same level, or will actually have the same experience level, or will build the same types of products at the same quality level. So, again, RadiSys focused on more of a premium position in the market, where our products go into applications that just can't fail. By "field-proven," we mean our products have been successfully deployed in these applications.

In fact, we've had our systems NEBS certified, which is very important in telecom. A lot of this is very similar to what you need in military and aerospace: Shock and vibe tests, for example in NEBS, electrostatic discharge, flame tests, things like that … the ability to work in a very hot environment. When you certify to NEBS, your product has to work at up to +55 °C.

MIL EMBEDDED: ***Why do you think your military customers are turning to telecom-based technologies?***

AMBROSE: We're seeing a new set of applications emerge, outside of the standard avionics/vetronics space. It doesn't mean customers don't want to have rugged, reliable systems, but it means that they're in a network-centric warfare mode. So it's very logical that they would want to have solutions that look more communications-centric.

MIL EMBEDDED: ***OK, so they're willing to use a different set of metrics to describe and procure those products then?***

AMBROSE: Right, they still want rugged, they still want field proven, but in a number of cases, we're selling a COTS technology. Because we're selling a commercial technology, they know it's open. When they choose ATCA, they want to know it's not something where RadiSys or anyone else can suddenly say, "Wait a minute, I own this, you can never bid it from anyone else in the future."

MIL EMBEDDED: ***What about long-life support for your ATCA products? Telecom isn't known for lengthy EOL management practices.***

AMBROSE: Actually, telecom equipment *does* need to survive long-term. Because of our history in embedded segments that demand long life, we already know how to provide long-life solutions for typical military and aerospace requirements. And we've done that. We can point to customers that say, "Look, here's a product we designed in 1995 and we supported it all the way to 2007. Here's a product we designed in 2001 and it's still going."

MIL EMBEDDED: ***Whom would you consider your biggest competitors in the market?***

AMBROSE: The biggest competitor overall across all of our businesses continues to be in-house design: The concept is a lot of people "roll their own boards," and we understand why. They used to do that in commercial, they used to do that in communications. Those markets now are rapidly adopting COTS technology that allows them to deploy their R&D resources to areas of their own differentiation.

Then you have the traditional CompactPCI and VME companies. And we're not targeting direct replacement of those form factors; but I certainly think that they're going to want to take their products and try to extend them as far as they can. Our competitors are also the PC/104 [small

form factor] companies that could potentially do something with PCs. But realistically, the COM Express way of building robots, UAVs, and ruggedized personal computers is a very robust method.

MIL EMBEDDED: ***So far we've focused on boards; let's shift gears and talk about systems. Does RadiSys do systems integration?***

AMBROSE: Yes, we've been doing that for a while. We're going to see more and more customers say at the end of the day, "Just please manage the whole thing for me."

MIL EMBEDDED: ***Which technologies do you foresee into the future?***

AMBROSE: ATCA is a very rich ground for technology. We've just announced ATCA 4.0 – our first 40 G products and strategies. You can fast-forward two or three years and this will become very important across a broad range of applications that need such high performance. And that means if you have backplane performance, you can get a lot more I/O, even if your base technology's around 10 G.

I also think the whole concept of integrating packet processing and applications processing for secure networks information assurance is gaining traction. That means you can do deep packet inspection on a lot of your communications, and you can improve security within your networking infrastructure. It also gives you the ability to have a lot of performance so that even if you're securing and encrypting information within the network, you still get a pretty reasonable user experience on either end.

MIL EMBEDDED: ***What would you say to those who remain skeptical about using telecom-based technologies such as ATCA in the military?***

AMBROSE: Unlike some other technologies that have been out there, as I mentioned, the marketing has lagged the reality in this case. We know [ATCA] products work, we know the products fit a need in the mil/aero industry, and we're going forward on that basis.

Anthony Ambrose *joined RadiSys in 2007 and is the Vice President and General Manager of Communications Networks. Previously, Anthony was Intel's General Manager of the Modular Communications Platform Division in the Communications Infrastructure Group. In this role, he led Intel's ATCA effort while his organization had the responsibility for servers, blades, boards, and software for the telecommunications industry. Anthony holds a B.S. in Engineering from Princeton University. He can be contacted at anthony.ambrose@radisys.com.*

RadiSys
503-615-1100
www.radisys.com

Designers are demanding end-to-end sensor processing, right out of the box

An interview with Dipak Roy, Chairman of D-TA Systems

EDITOR'S NOTE

Sure, there are lots of signal processing, I/O, and FPGA co-processor boards on the market. And many of them are amazingly fast, sporting leading-edge performance and software. The trouble is, designers still need to design a system around these boards, which often come from disparate suppliers. I caught up with Dipak Roy, an old friend from my COTS sonar days, who has resurfaced with his new company, D-TA Systems. Dipak is a bona fide sonar expert, and his former company, ICS (aquired by Radstone, now a part of GE Intelligent Platforms), invented the FPDP interface that became a VITA standard. The fact that Dipak is now relying on 10 GbE for data transfer came as no surprise. Edited excerpts follow.

MIL EMBEDDED: ***According to your literature, D-TA's value proposition is box-level, preconfigured systems. What does that mean?***

ROY: The conventional norm for building embedded systems for military applications is to start from COTS boards from various vendors. One has to select a computer platform, OS, enclosure, and power supply, then procure software device drivers and application routines. Next, the processes of system integration and debugging commence, followed by laboratory tests and field trials. The whole exercise is time consuming, expensive, and often leads to significant cost and time overruns.

However, D-TA Systems' approach changes the whole paradigm for demanding sensor-processing applications by offering box-level, user-reconfigurable systems based on 10 Gigabit network technology. Since our products come fully tested and require minimum customer reconfiguration, the deployment cost and time savings are dramatic.

MIL EMBEDDED: ***Describe the types of integration challenges customers face with board-level point solutions.***

ROY: Starting from board-level products (COTS boards) and building a complex multichannel system is a time consuming and expensive proposition. For example, a user trying to implement a 16-channel HF radar system with four-channel COTS PMC ADC cards (the maximum ADC count available on current PMC cards) needs to integrate four such cards, acquire a clock source, and develop a clock distribution strategy to maintain synchronization before any application development can commence.

> "... I believe that the current trend in multithreaded software processing in platforms with multiple GPP cores will revolutionize sensor processing; this will allow more real-time signal processing functions to be housed in software and GPPs, rather than in FPGAs."

To help solve these issues, a fully integrated and tested functional module that can be quickly configured for a particular application is a real time-saver. Following a successful trial, customers often require customization services for repackaging, ruggedization, functional modifications, and so on, for volume deployment.

MIL EMBEDDED: ***What are your target markets and applications?***

ROY: Defense, aerospace, wireless, and test and measurement. Our specialty is sensor-array processing, with "plug-and-play" products that cover an application space from sonar/acoustic to radio/radar.

We are targeting a very wide range, from sonar to wireless. In the sonar/acoustic area, we are looking at high-frequency and high-channel-count applications such as mine sonar, shallow-water sonar, obstacle avoidance, vibration analysis, ultrasound, and real-time acoustic simulation. In the radio and radar area, the products are particularly suited to multiantenna applications such as phased array radar, RF direction finding, smart antenna base station, signal intelligence and target location, sonobuoy, RF test, and channel simulation/emulation. Our specialty is to offer precise synchronization, from RF to baseband, across a large number of input and output channels.

Interestingly, we have also recently established an advanced sensor processing laboratory at a local university, and several faculty members and graduate students are now engaged in developing new and exciting applications.

MIL EMBEDDED: ***Where does 10 GbE fit into the need for deployment speed?***

ROY: 10 Gigabit network technology offers data throughput rates that are faster

than any computer bus. And, more importantly, it allows better system partitioning that is particularly advantageous in mixed-signal applications. Analog functions can be isolated from the "noisy" computer bus for better analog performance. Moreover, it is part of the IEEE standards process and is continually upgraded for speed and bandwidth.

MIL EMBEDDED: ***You are building systems using the same COTS hardware and software as everyone else. Why would a customer choose your solution?***

ROY: All our products are designed from the ground up. Unlike the typical COTS board vendor, we are not constrained by board size, predefined connector pin-outs, power supply ratings, enclosure size, and so on. So we can offer an uncompromising solution that is designed to solve the problems encountered in demanding sensor-processing applications: dynamic range, analog performance, processing speed, and instantaneous bandwidth, to name a few.

Also, as mentioned, virtually all of our products are 10 Gigabit network attached, and most of them are housed in a 1U high, standard 19" rack-mountable enclosure. Our mandate is to offer end-to-end solutions for demanding sensor-processing applications. For radio and radar applications, examples include tunable RF, multichannel IF (software radio), and 10 Gigabit record and playback systems. For high-frequency sonar and acoustic applications, there are high-precision signal conditioning, 24-bit digitization, and a high channel count. These products can be seamlessly connected to build any complex sensor processing systems in a matter of days, not years. Other factors in making a designer's life easier include libraries of pretested FPGA cores and real-time multithreaded software application routines. So our value proposition is very compelling: The customer can focus on their applications and not on data acquisition system development.

MIL EMBEDDED: ***What are the software considerations?***

ROY: The 10 Gigabit network connection is OS agnostic, so we typically use the 10 Gigabit links for data transfer and 1 Gigabit network for control. The base SDK supplied with the products includes a Control API and Data API. The Control API allows the control of our boxes, while the Data API allows the user to build applications using example codes provided. This structure shields users from socket calls and simplifies data access. We have also developed many DSP functions for multithreaded, multicore platforms that seamlessly attach to the base SDK.

Looking toward the future, I believe that the current trend in multithreaded software processing in platforms with multiple GPP cores will revolutionize sensor processing; this will allow more real-time signal processing functions to be housed in software and GPPs, rather than in FPGAs. It will also save development time and cost and enable rapid reconfiguration.

MIL EMBEDDED: ***Will the sensor-processing world adopt emerging 40 or even 100 GbE technology?***

ROY: D-TA Systems is the first company to introduce 10 Gigabit sensor processing. It is new and just getting started. As opposed

to datacom applications, sensor-processing applications require sustained processing at a high data rate. Most applications cannot tolerate missing samples or data packets. There is no "resend" in sensor processing applications. The processor has to keep up with the data rate. That is why the sensor-processing world trails behind the datacom world in terms of sheer network speed.

As the processor gets faster and faster, we will see introduction of 40 and 100 GbE networks for sensor processing applications for the simple reason that more signal channels with higher bandwidths can be handled by each network. Moore's law dictates that the processing power will increase to allow signal processing at these rates.

MIL EMBEDDED: ***How has sensor-processing technology evolved in the past 20 years?***

ROY: One has to remember that not long ago, all sensor processing was done in the analog domain. Over the years, it has evolved significantly and in many ways. Today, the processing has become mostly digital, primarily because of increased data converter speed and accuracy. This has not only allowed detection and processing of low-level signals, it has also allowed data converters to get closer and closer to the actual sensor, thereby simplifying front-end analog design. For example, the advent of Sigma-Delta technology has revolutionized acoustic signal processing. The advent of FPGAs has allowed many DSP functions to be implemented in real time. The availability of high-speed networks and computer buses has allowed raw or preprocessed sensor data to be accessed by a general-purpose computer for further processing or display.

MIL EMBEDDED: ***What are the biggest challenges facing your customers?***

ROY: The biggest challenge facing our customers is time to market. In the current tight-budget scenario, cost overruns cannot be tolerated. Further, most customers these days want to see a working demonstration before funding any development. D-TA products are designed to precisely address these problems: Virtually any sensor-processing systems can be built quickly by selecting products from our catalog.

MIL EMBEDDED: ***Which new technology do you think we'll see in the next two years that's not evident today?***

ROY: Going forward, we will see more evolution in data converter technology in terms of higher speed and accuracy. The other focus will be on lower-power devices that will also allow increased channel counts. During the next several years, it might be possible to connect the data converters (ADCs and DACs) directly to the antenna. This will require not only very high-speed converters but also faster FPGAs and processors to handle the very high data rate. In the short-term, we will also witness significant evolution in RFIC [Radio Frequency Integrated Circuit] technology, which will significantly reduce the size and cost of RF subsystems.

Dr. Dipak Roy *is Chairman of the Board at D-TA Systems Inc. His embedded industry experience began briefly in the form of Research Engineer. Then he took the entrepreneur route: His first company, ICS (now part of GE Intelligent Platforms), won defense contracts including the U.S. Navy's SQQ-89 sonar upgrade program. Additionally, Dipak has authored a patent and published articles in 30 technical publications. He has received many awards, notable among them is the American National Standards Institute (ANSI) award for the invention and standardization of FPDP, a high-speed data flow concept that simplifies system integration. In 2007, he was appointed by the government of Canada to the board of Sustainable Development Technology Canada (SDTC), a $1 billion foundation developing "Clean Technology" companies in Canada. He received his PhD in Electronics Engineering from Carleton University, Ottawa. Dipak can be contacted at dipak@d-ta.com.*

D-TA Systems Inc.
613-745-8713
www.d-ta.com

Elements of a deployed, modern net-centric system

By Haritha Treadway

U.S. Marine Corps photo by Cpl. Albert F. Hunt

In striving for the U.S. DoD's vision of net-centricity, military embedded designs must optimize size and power consumption, provide fast and effective graphics and visualization, keep commanders and soldiers connected, and provide hardware supported by the Global Information Grid's (GIG's) Service Oriented Architecture (SOA) framework. But don't forget the key ingredient: the processor, which heightens performance and brings all these elements together.

For the past 10 years, the United States Department of Defense has pioneered the military doctrine of net-centricity to achieve effective information sharing in a complex environment. Net-centric warfare aims to achieve a robustly networked force for shared situational awareness, which in turn dramatically increases military mission effectiveness.

Net-centric devices include soldier navigation tools, precision optics capabilities, and radio technologies. These devices rely on robust embedded computers and high-performance, low-power processors. Developing embedded systems for net-centric devices includes the ability to reduce soldier-borne weight, increase mission length, and improve soldier situational awareness and effectiveness.

To achieve these goals, embedded designs have to optimize size and power consumption without making sacrifices in equally important graphics, I/O, and communications capabilities. In addition, software plays a key role in allowing embedded designs to seamlessly and cost-effectively integrate the hardware capabilities with existing infrastructures developed to support net-centricity. Meanwhile, the underlying processor accelerates performance and ties together all these net-centric needs.

Size and power

Size and power considerations are central to the effectiveness of net-centric devices. Net-centric warfare leverages the power of real-time information transmitted over a complex grid of data to make military operations more effective over land, air, and sea. With instant access to data comes the need for soldiers to be mobile and responsive, meaning the embedded devices they operate are flexible, compact, and light, allowing them to move freely and quickly.

In addition to being physically worn or carried on a soldier, embedded devices also operate in vehicles such as UAVs, specialized planes, and helicopters that are space-constrained while still requiring the ability to interface with extensive machinery and other peripherals. As a result, size is one of the top priorities in designing embedded systems targeted for net-centric warfare.

By choosing processors with low Thermal Design Power (TDP), designers can

> “3D visualization and simulations with interactivity are critical in understanding local environments in great detail – not only for the effectiveness of the mission, but also for the safety of soldiers.”

remove the need for bulky fans or heat sinks, thereby minimizing the size of the hardware. In addition, because net-centric devices are so mobile, power consumption is equally important, with a focus on battery operation easily lasting more than 10 hours. To meet these needs, processors with features such as Dynamic Voltage and Frequency Scaling (DVFS), deep sleep, and idle states bring an obvious advantage by managing heat and power dissipation. The challenge is to meet these size and power consumption needs without sacrificing processing and performance capabilities.

Graphics and real-time visualization

As military operations become quicker and more reactive, net-centric warfare relies heavily on sophisticated graphics to provide realistic imagery of situations. 3D visualization and simulations with interactivity are critical in understanding local environments in great detail – not only for the effectiveness of the mission, but also for the safety of soldiers.

Display support is also varying, as shown in Table 1. On one hand, real-time simulation displays must be large with very high resolutions, in the range of XGA (1,024 x 768) to High Definition (1,280 x 720). For devices carried by soldiers in the field, the displays are often smaller with VGA resolution (640 x 480). Dual-display support is also important. Therefore, the key is to pick an embedded processor that has the right mix of graphics capabilities for the specific end application. An ideal choice given the space constraints of embedded devices is a CPU with an integrated graphics engine to allow for small form factor designs.

Connectivity and I/O

While size, power, and graphics are critical, another key to net-centric warfare is the communications and I/O capability of the device. Net-centric warfare relies heavily on sensors, satellite communications, and GPS. These functions will be linked together via the IP-enabled Global Information Grid (GIG), a communications project of the U.S. DoD that connects military entities and operations across the globe to a common networked

Application	Soldier in field	Control center	Simulations
Graphical requirements	Small screens VGA resolution Lightweight	Large displays XGA/HD resolution Fast response time	High definition XGA/HD resolution 2D/3D graphics and fast response time

Table 1 | Graphical needs for net-centric warfare depend on the end application.

infrastructure. Figure 1 highlights how the GIG offers an overarching structure for global, cross-platform networking.

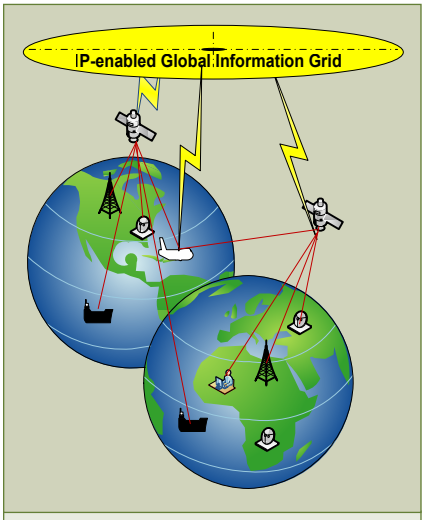

Figure 1 | IP is used to connect various communication devices (Ethernet, GPS, satellite, and radio) to a common infrastructure via the Global Information Grid.

As demonstrated by the large demand for ruggedized, military-grade routers, Fast Ethernet is also an important way to access IP. However, traditional data buses will still play an important role in supporting traditional military devices and peripheral machinery. An embedded device will have to support the latest in Wi-Fi, 1000BASE-T Ethernet, or GPS capabilities while still preserving high-speed serial ports, USB, PCIe, CAN, and more, making it rich in communications and I/O.

Hardware supported by software

With such an extensive list of hardware requirements, software is critical to bring the net-centric system together seamlessly and effectively. Service Oriented Architecture (SOA) is a key element. In fact, one of the tenets of the GIG is to use SOA capabilities to increase the flexibility and portability of applications used in net-centric operations.

The SOA strategy is data-focused rather than hardware-focused. This means more parallelism as information is readily available to several players with access to the grid regardless of the underlying hardware used by each specific player. The SOA approach is very different from the traditional approach, which tended to be more serial, using a sensor or piece of hardware to communicate serially to a data-gathering station that then dissipated the information to another entity.

An SOA framework separates the hardware from the end service or application logic. SOA bundles are more portable and accessible, easily transported from one platform to another with minimal development work. Portability is critical in net-centric warfare, especially as projects like the GIG require devices to be flexible across terrains and operations. Figure 2 shows an example framework based on Java. Here, the advantage is to separate application-specific bundles, referred to as *vertical market bundles,* from foundation bundles and hardware platforms. This makes code reusable and flexible, allowing for parallelism across platforms.

Processor accelerates performance, brings it all together

As mentioned in previous sections, a low-TDP processor is a key element within each component of net-centricity, and also heightens performance. Accordingly, the Intel Atom family offers a highly suitable solution.

Performance

The U.S. DoD and its military forces know the need for speed can make or break a mission – and ensure soldier survival. Therefore, a fast processor is integral. The Atom can meet this need by offering accelerated processing speed: CPU speeds up to 1.6 GHz allow the Atom to support the intensive processing requirements of modern net-centric military applications.

Size and power

Embedded boards targeted for net-centric warfare need to be small and low power. Thus, the Atom's small-form-factor savvy and low TDP facilitate net-centric military applications. Table 2 shows the three variations currently available. Each processor has unique capabilities and is targeted for specific applications. As the table shows, the Z5xxP stands out with 2.2 W of TDP. One of this chip's features that keeps thermal power low is its thermal monitor circuit, which brings

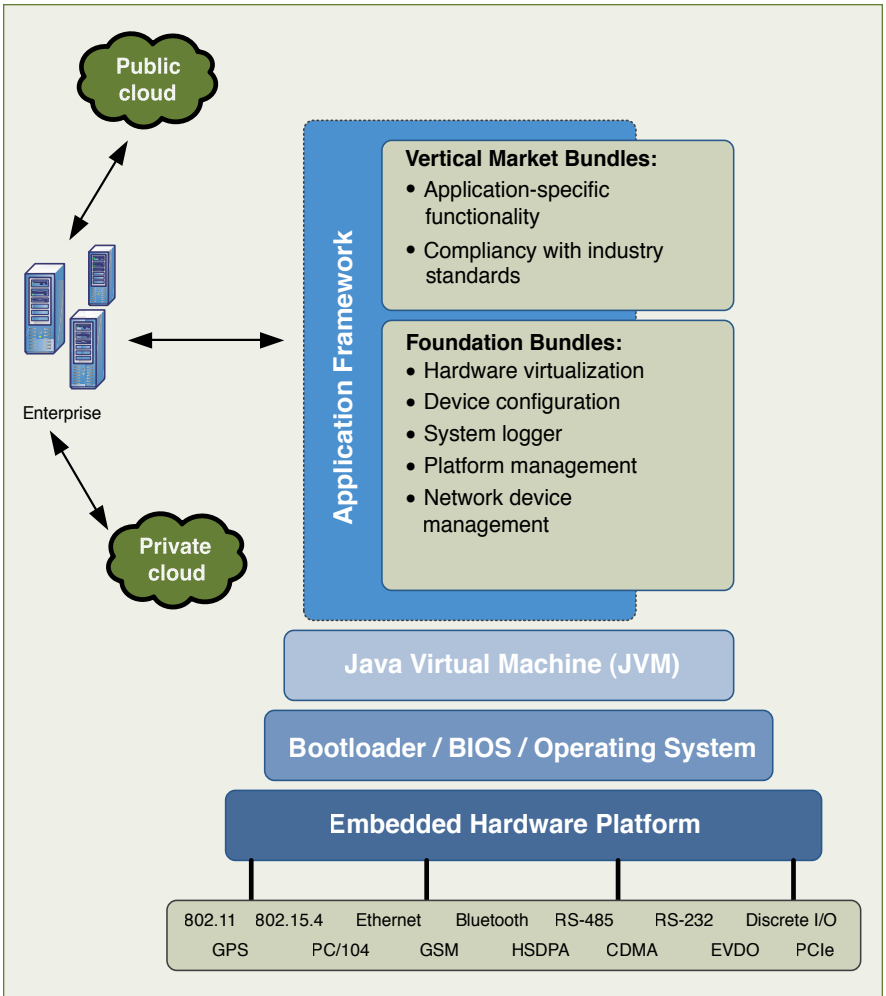

Figure 2 | Middleware allows code to be reusable and portable across hardware platforms by separating foundation bundles from application-specific bundles.

down the temperature when the processor reaches its +85 °C limit. The cooling process occurs with minimal effects on application performance.

To improve efficiency and minimize waste, Z5xxP also includes a new C6 deep sleep mode that boasts a very efficient power spec of 0.1 W. Table 2 additionally highlights the sleep/idle power consumption of other processors; these features are ideal for net-centric military devices because they eliminate the need for bulky cooling components. The end result is an extended lifetime for the embedded device without concern for unnecessary electrical or mechanical failures.

Graphics and visualization

Despite small size and low power requirements, military net-centric applications still require high-speed graphics and visualization capabilities to get mission-critical information to the warfighter on a moment's notice. The Atom D510 meets this need, with an ultra-fast 400 MHz graphics engine integrated into the CPU.

I/O and communications

In addition, with extensive I/O support, the military can interface with several devices either wirelessly or through standard means. The D510 is a good example of this, offering speeds up to 1.8 GHz with dual-core capability. This high-speed I/O support meets the processing crunch required by communications, handheld, or radio devices deployed in net-centric warfare.

Software considerations

Atom processors are compatible with an SOA framework and also support hyper-threading to improve performance in multithreading or multitasking applications, which are often required of net-centric military devices.

Embedded systems drive net-centric warfare

Embedded hardware providers can help the United States military migrate to net-centric warfare techniques by making sure their products meet performance needs, provide optimal size and power consumption, enable fast graphics and I/O, and provide SOA framework compatibility. Atom-based products – such as Eurotech's Catalyst family of embedded modules – meet these needs and are an ideal fit for modern net-centric military applications.

Processor	Processing speed	Thermal Design Power (TDP)	Deep sleep/ Idle power
Intel Atom Z5xxP	1.1 GHz – 1.6 GHz	2.2 W	0.5 W – 0.1 W
Intel Atom N270	1.6 GHz	2.5 W	0.5 W
Intel Atom D510	1.6 GHz (dual core)	16 W	4.5 W

Table 2 | Comparison of processors' speed, TDP, and deep sleep/idle power metrics

Haritha Treadway *is a product manager at Eurotech Inc., responsible for the company's ARM- and x86-based boards portfolio. She had 10 years of experience as an engineer in the semiconductor industry before joining Eurotech. Haritha received a BS in Electrical Engineering from Cornell University and an MBA from Boston College. She can be contacted at haritha.treadway@eurotech.com.*

Eurotech, Inc. • 301-490-4007 • www.eurotech.com

470th Military Intelligence Brigade Public Affairs photo by Gregory Ripps

Secure virtualization combines traditional desktop OSs and embedded RTOSs in military embedded systems

By Robert Day

Advances in software and hardware technologies now make it feasible to use both embedded and desktop operating systems in a secure military system. Robert examines enablers such as a secure separation kernel and an embedded software hypervisor, then explains uses of desktop OSs in secure military systems.

As Intel continues to bring its processor technologies into the embedded world, an interesting convergence of embedded applications with more traditional desktop applications is taking place. For military applications, desktop systems and embedded systems have traditionally been separate systems, connected over a secure network (see Figure 1). However, there is now a desire to consolidate multiple hardware platforms to reduce Size, Weight, and Power (SWaP) while maintaining the security that discrete systems traditionally offered.

By combining new software and hardware technologies, this consolidation is now a reality, without having to sacrifice either performance or security. The software technology is a secure separation kernel and embedded hypervisor, utilizing the Intel multicore virtualized hardware technology. This software platform becomes a true enabler of modern hardware functionality; however, before examining the application of the technology, it is beneficial to examine the two component parts of the software.

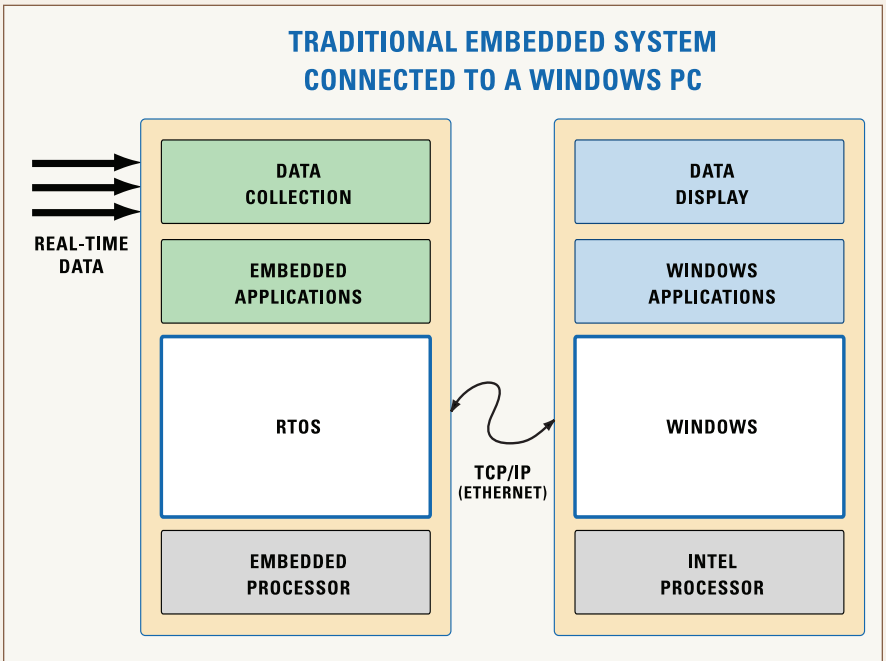

Figure 1 | Traditional systems have physically separate hardware to maintain security and performance.

Software component 1: Secure separation kernel

A separation kernel is a small, lightweight operating system that is the lowest-level connection to the processor. The separation it provides is not dissimilar to a traditional time- and space-partitioned OS (see Sidebar 1), but it also adds a secure function by enforcing predefined security policies in areas such as device management and interpartition communication. Also, the separation kernel itself does not offer traditional OS features such as disk or network access, but it does manage scheduling and memory functions. The advantage of removing many of the high-level OS features is that the separation kernel can be kept small and efficient, offering real-time application performance and secure, high-speed interpartition communication using memory rather than physical networking connections.

Safety-critical OSs vs. security separation kernels

For safety-critical systems, a new class of partitioned RTOSs has been developed in recent years, based around an ARINC-653 model. These partitioned RTOSs provide the ironclad guarantees of time and space (for example, memory) partitioning, which prevent one software function from interfering with another. They also enforce strict quotas/budgets on a function's access to system resources (RAM, CPU time) and restrictions on which a complement of software functions may be active on a given module.

As with a safety-partitioned RTOS, a separation kernel's primary job is to let developers establish partitions, then to enforce those partitions during runtime. In this sense, a separation kernel and a partitioned RTOS are very similar. However, there are key differences:

- A separation kernel is primarily concerned with partitioning a system's data and resources and controlling information flow between partitions.
- A separation kernel must enforce strict adherence to data isolation, damage limitation, and information flow policies.
- A separation kernel operates in a much more dynamic environment than a partitioned RTOS.
- A separation kernel must protect against malicious attacks, from without and within.
- When combined with hypervisor technology, the separation kernel can run both software applications and other OSs in their own secure partitions.

For high-assurance systems, the Common Criteria evaluation is more complex and hence more costly than a DO-178B certification is for safety. Meeting the highest EAL criteria requires the use of formal methods, where the separation kernel developer proves security properties of the separation kernel. While formal methods have improved significantly from their early days, it quickly becomes intractable to use them on software of even modest size/complexity (a few thousand lines, at most). Consequently, separation kernels intended for the most security-critical applications have fundamentally different architectures from those found in safety-partitioned RTOSs.

The functionality of a separation kernel is therefore quite different from that of a safety-critical OS. And the addition of the hypervisor allows "guest" OSs to sit on top of the separation kernel and give additional functionality to applications, including Linux/Windows functionality that could not be found in a safety-critical OS. Unlike the monolithic safety-critical OS, the separation kernel and hypervisor can be separated from the guest OSs when it comes to evaluation, thus reducing the amount of code that has to be evaluated to the highest levels and helping to implement a system with multiple levels of security evaluation.

Sidebar 1 | There are differences between safety-critical OSs and security separation kernels.

In the security world, this small separation kernel is the cornerstone of high-assurance systems, offering security policy enforcement and strict partitioning, using a Multiple Independent Levels of Security (MILS) architecture. This allows security engineers to build systems that need to be taken to the highest level of Common Criteria (currently EAL 7) and run applications requiring different security levels on the same physical hardware. Many separation kernels are derived from partitioned OSs by removing OS functionality and adding security features. However, to achieve the highest levels of evaluation, the software must also be proven secure by using formal-methods analysis. The separation kernel is the fundamental enabler to the secure coexistence of multiple applications on the same hardware platform. And, when united with an embedded hypervisor, the combination of desktop and embedded systems can be achieved.

Software component 2: Embedded hypervisor

A software hypervisor is a software layer that allows different guest OSs to reside on a single hardware platform. This technology is commonly used in the enterprise or data center realm to allow the IT departments to run all their required applications across multiple versions of server-based OSs. In the embedded world, the use of hypervisors is not as common. The requirement to run multiple different versions of an OS on a dedicated embedded system is not as crucial. And there have been questions over the performance of running extra layers of software in systems where real-time performance is key. When a hypervisor and a separation kernel are combined, the ability to bring desktop and embedded systems together becomes a reality (see Figure 2).

Hardware: Desktop OSs in secure military systems

With the use of Intel processors, traditional desktop OSs are also being used in many military systems. However, when multiple levels of security are required, this can stop the use of nonsecure desktop OSs. With the introduction of a secure separation kernel and hypervisor, traditional desktop OSs and applications can be run in their own unclassified partition, thus allowing for the functionality of a known user interface and applications, without compromising the security of the

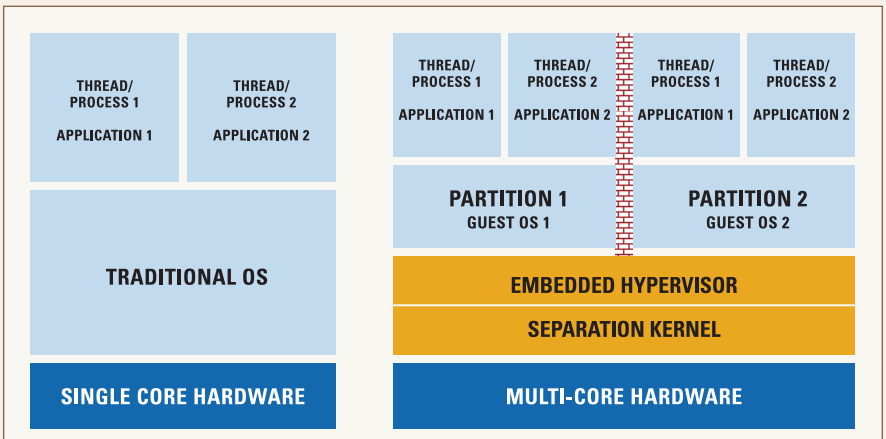

Figure 2 | The combination of separation kernel and hypervisor allows multiple OSs to be run securely on the same physical hardware.

rest of the system. Anything that enters into the desktop partition cannot breach the secure separation kernel and hence will be contained in the unclassified part of the system.

This software partitioning and virtualization also aid in the consolidation of hardware and the reduction of SWaP, which is of particular interest in many military scenarios. By running separate systems in their own partitions, and allowing for different OSs and applications to be run in those partitions, there can be a true consolidation of physically separate systems to a single physical piece of hardware.

The use of Intel multicore, virtualized processors allows the merging of a Windows or Linux desktop system with a more traditional Real-Time Operating System (RTOS), and allows the same performance and functionality of applications as if they were still running on their own dedicated hardware platforms.

An additional feature that is very compelling in regard to this approach is that of virtual networking. Here, the guest OSs and applications can communicate "virtually" with other guest OSs and applications, even though they are residing in separate partitions. The virtual network looks to the applications as a real network port, and so these applications can communicate as if they were two physically separate networked devices, even though the communication is internal. A secure separation kernel can also enforce security policies to this virtual networking and dictate which partitions can communicate with each other and in which direction (see Figure 3).

This gives a secure partitioned environment with the ability to run multiple guest OSs and applications separated from one another on the same hardware. To allow near-native performance while maintaining real-time determinism and security, hardware virtualization support for both execution and memory can be utilized by the separation kernel and hypervisor. Independent studies performed on the LynxSecure separation kernel and hypervisor have shown that running benchmark applications on a virtualized Linux OS yields less than a 5 percent performance degradation as compared to the same applications running on a native

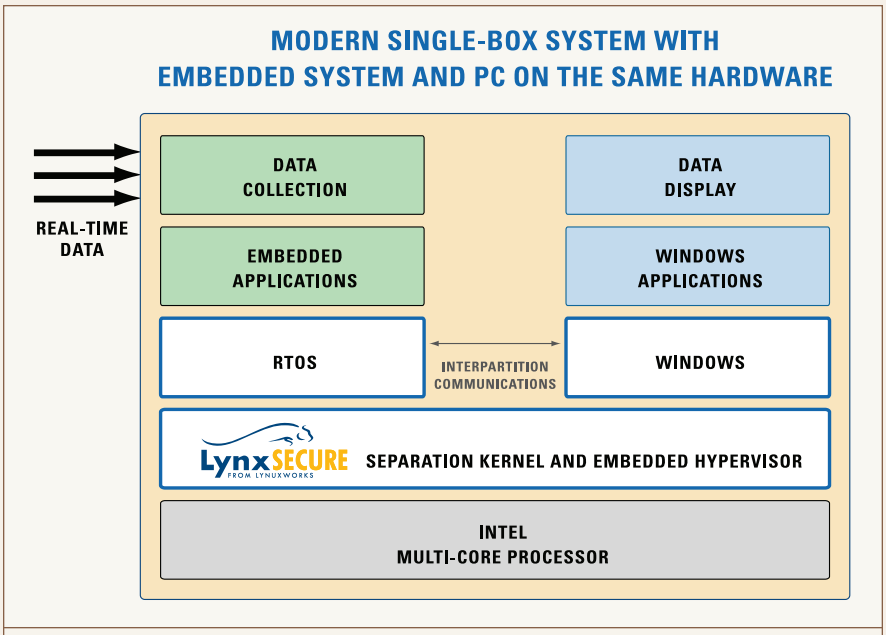

Figure 3 | The use of a separation kernel and hypervisor allows desktop OSs and RTOS to reside on the same hardware platform.

implementation of the same Linux on the same hardware.

Another benefit of the migration of desktop systems is afforded if the embedded hypervisor uses Intel's Virtualization Technology. This allows Microsoft Windows to be run in fully virtualized mode, which requires no changes to Windows to run on the hypervisor, and a combination of the software separation kernel and the hardware virtualization gives Windows the impression it has the whole system, while running in its own secure partition. If no changes are required to either Windows or its applications, this speeds the development or porting activity from a stand-alone system to a secure virtualized system.

An example of a MILS solution running on Intel virtualized hardware is LynxSecure from LynuxWorks. It is a secure separation kernel and embedded hypervisor that uniquely offers both para- and full-virtualization of guest OSs, and maintains real-time performance and MILS security that can be evaluated to the highest Common Criteria levels. It takes advantage of multicore Intel components to enable high performance even when running multiple guest OSs. Microsoft Windows can run on the same system as Linux and RTOS, with each having its own secure partition and running applications at different security classifications. For the next generation of military embedded systems, the combination of LynxSecure and Intel hardware allows the ultimate in flexibility of system and applications while maintaining the highest level of security.

Robert Day *is Vice President of Marketing for LynuxWorks, where he is responsible for all global external and internal marketing functions. With more than 20 years in the embedded industry, his most recent position prior to joining LynuxWorks was heading marketing for Mentor Graphics' embedded software division. Prior to the marketing role, he held a variety of management, sales, and engineering positions for Mentor Graphics and Microtec Research, spanning more than 18 years in total. Based in San José, California, Robert is a graduate of The University of Brighton, England, where he earned a Bachelor of Science degree in Computer Science. He can be contacted at rday@lnxw.com.*

LynuxWorks
408-979-3900
www.lynuxworks.com

INNODISK®
Beyond your imagination

innoRobust®

Top Reliability, Finest Solution

InnoRobust brings you to experience the most advanced flash storage series. It provides SSD with extreme reliability and outstanding performance. We introduce data security features targeted for military, mission critical, and aerospace applications.

- Sensitive Data Security / Protection
- Operation in Environmental Conditions
- Conformal Coating Support

-40°C~+85°C Wide Temp. Available

InnoRobust Ultra SATA SSD
Capacities: 16GB~256GB

The total solution for Memory and Flash Storage

SSD Industrial Solid State Disk -Capacity to 512GB

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

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

2.5" SATA 10000-RS
Capacities: 2GB~32GB
Read: 100MB/s(max.)
Write: 50MB/s(max.)

-40°C~+85°C Wide Temp. Available

2.5" ATA 8000-J
Capacities: 8GB~128GB
Read: 85MB/s(max.)
Write: 75MB/s(max.)

SATADOM

i-Series
Capacities: 2GB~16GB
Read: 100MB/s(max.)
Write: 50MB/s(max.)
Dimension: 25x39.4x6.5mm(WxLxH)

SATADOM i-series

8000 Series-EDC/iCF

-40°C~+85°C Wide Temp. Available

Vertical type 40pin
Capacities: 2GB~16GB

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

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

i-DIMM Series

-40°C~+85°C Wide Temp. Available

Low Profile Available

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

-40°C~+85°C Wide Temp. Available

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

www.innodisk.com

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

U.S. Army photo by Staff Sgt. Alex Licea

Developing high-performance embedded network security applications: A heterogeneous multicore processing approach

By Daniel Proch

Today's network and computing infrastructure is of critical interest to our national security. The vital communication systems supporting our multinational forces are valuable resources that share a couple of important characteristics: They require massive amounts of bandwidth to support our insatiable appetite for IP-based services and they require mechanisms offering visibility into all data protocol and application layers to ensure network security. Unfortunately, the network appliances commonly hosting security applications have failed to keep pace with improvements in network performance. However, a new heterogeneous multicore processing architecture can scale to support tomorrow's throughput needs while providing the ability to see deeply into network traffic.

The amount of network traffic in today's wired and wireless infrastructure continues to rise at dramatic rates to keep up with voice/video/data services and real-time military applications. In both classified and unclassified military networks, line rates of 10 Gbps are commonplace and are expected to quickly grow to 100 Gbps in the next several years. These throughputs are largely the result of new IP-based network-centric warfare applications like real-time battlefield monitoring, video surveillance, and fully networked forces. The scalability of the network as a whole is a major risk to the effectiveness of any network-centric warfare program.

As network throughputs explode, we also need to be able to intelligently monitor our networks for exploits and to protect confidential data sources from breaches. An increasing threat to our homeland security is the growing number of high-profile cyber attacks on major military installations, our communications systems, government agencies, and financial markets and the resultant leakage of classified information and personal data. Even Google has been attacked recently in what could be an instance of state-sponsored corporate espionage. Compounding the problem, there are countless other attacks never publicized, but rather hidden by a web of obscurity for obvious confidentiality and national security reasons. The security applications responsible for protecting these critical resources need to keep pace with these increasing network throughputs with even greater network intelligence. Thus, communications equipment must provide complete visibility into network traffic at extremely high bandwidths by using content inspection to ascertain the nature of traffic, not just its destination.

Military networks already deploy an array of security applications to protect their classified and unclassified resources. These applications include virus scanning, firewalls, Intrusion Detection and Prevention Systems (IDS/IPS), Distributed Denial of Service (DDoS) mitigation, Data Loss Prevention (DLP), test and measurement, and network forensics solutions. These applications work almost entirely by providing Deep Packet Inspection (DPI) and flow analysis, looking for known patterns in network flows and blocking or recording them. With the need for application awareness, security processing, and DPI, the amount of processing power required for these computationally intense applications grows exponentially at these increasing line rates. However, these needs for increased visibility, throughput, and

network processing power can be met by a heterogenous multicore processing architecture.

Heterogeneous multicore processing paradigm solves paradox

Network and security applications can generally be viewed in several distinct network architectures. Compute-intensive applications like intrusion prevention systems are deployed as active elements sitting directly on the network wire (inline) processing every bit of data that traverses the application in real time. These active security appliances need to operate at network line rates with very low latency. Computation that adds just microseconds of delay to network traffic can ruin the effectiveness of real-time end-user applications like sensitive military telemetry systems or Voice over IP (VoIP). Specifically, developers typically view 250 microseconds of delay that any inline network element can add (before end-user application performance degrades) as a high watermark in 1 Gbps networks.

Alternately, passive computing elements like network and computer forensics systems, intrusion detection systems, honeypots, and vulnerability scanners are not in the direct network path, but rather are deployed off a span port, network tap, or mirrored switch interface. These systems are responsible for collecting and analyzing terabytes of data from distributed sensors as would be typical in a battlefield node scenario. These passive monitoring devices can offer a thorough understanding of a network's topology and which services are available, and scan to assess which vulnerabilities might be exposed on the network.

Network appliances deployed in either a passive or active network architecture share a common trait in that they must guarantee 100 percent traffic capture across all packet sizes to be effective. Missing any portion of data in a communications stream poses a large threat, making the overriding security application ineffective.

Meeting these performance challenges warrants a new approach to the development of the high-performance systems required by the intelligent network. Such systems need to be capable of analyzing traffic at all layers of the OSI model, from the data link layer (Layer 2) all the way into the application space (Layer 7) while performing this intelligent processing on all traffic at sustained throughputs of 10 Gbps and higher. Achievement of these goals requires specialized and varied processing elements, each custom designed for a specific type of workload computation.

A heterogeneous multicore architecture sets a new performance benchmark for embedded application development though separate and discrete processing elements for packet classification, stateful flow management, and application and control plane processing, each with increasingly fine granularity. This architecture tightly couples off-the-shelf Ethernet switch processors and network flow processor cores with general-purpose multicore x86 systems over a high-speed 40 Gbps, virtualized PCIe datapath. This architecture can be scaled from very low-end systems up to appliances offering hundreds of Gigabits per second of packet analysis, stateful flow monitoring, DPI, and application throughput, all with a common software architecture. Accelerated designs based on a heterogeneous multicore architecture can enable equipment providers to deliver high-performance, flexible systems that are up to four times more efficient than systems based on x86 general purpose processors alone with standard Network Interface Cards (NICs) as shown in Figure 1.

Specialized packet, flow, and application workload processing

As shown in Figure 2, a distributed network acceleration architecture uses a multi-chip system to achieve maximum performance and application effectiveness. The three distinct processing stages function as shown.

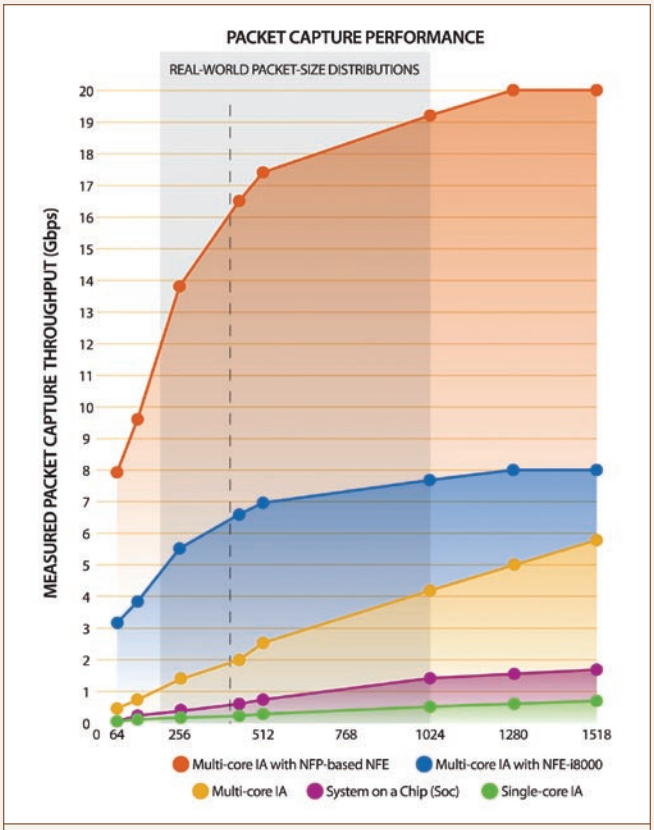

Figure 1 | Designs based on a heterogeneous multicore architecture enable high-performance, flexible systems up to four times more efficient than traditional x86 systems and standard NICs.

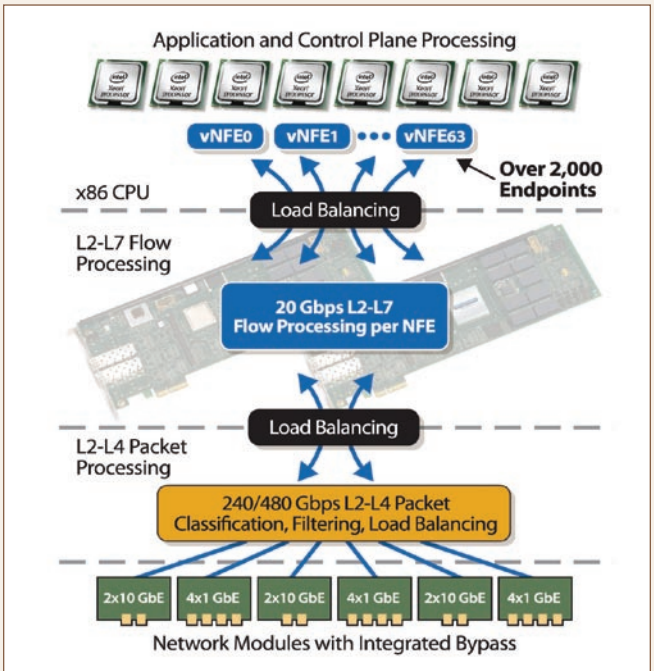

Figure 2 | A three-layered heterogeneous processing paradigm uses varied specialized processors to achieve maximum performance while keeping overall system costs low.

> "The mission-critical nature of our military networks driven by information-centric warfare and homeland security creates an opposing set of forces."

Ethernet packet processing

To heighten performance levels, off-the-shelf Ethernet switch processors are commonplace, offering up to hundreds of Gbps of configurable packet processing spanning the datalink, IP, and TCP/UDP packet layers. Traffic is classified on ingress and optionally filtered, cut-through to another network interface, or load-balanced across the Network Flow Processors (NFPs) that sit logically behind the Ethernet switch processors.

NFPs to accelerate higher-layer flow processing

NFPs containing a powerful array of microengine RISC processors are specialized, multicore devices optimized to offload burdensome workloads from general-purpose multicore CPUs. NFPs can handle lower-layer packet processing and accelerate higher-layer flow and application level processing. This accelerated architecture utilizes the network-optimized NFP cores for switching and routing, packet classification, stateful flow analysis, DPI, and dynamic flow-based load balancing. Other network processing functions such as TCP termination and SSL offload can also be performed on the NFPs and offloaded from the general-purpose CPUs. Traffic can be cleanly structured for transmission from the NFP to the general-purpose cores for application processing, thereby increasing host performance. Additionally, network flow processors provide hardware acceleration engines for PKI and bulk cryptography to assure line-rate throughput.

PCIe communications path to x86 cores

General-purpose multicore x86 CPU(s) in a system are optimized for application and control plane processing. From the network flow processors, packets are passed to the x86 cores across a high-performance, virtualization-aware PCIe communications path. An efficient zero-copy technique allows the transfers of packets directly into user-space application memory bypassing the operating system kernel, further accelerating application performance. Flows can be pinned directly to specific applications or load-balanced across parallel application instances to scale application performance.

As shown in Figure 3, through a cooperative set of software APIs between the x86 CPUs and NFP cores, the treatment of flows can also be updated in real-time, offering the ability to change the treatment of a flow after x86 analysis. This type of flexibility is essential in situations where a specific portion of a flow is of interest for inspection. After inspection is complete, all subsequent packets belonging to the flow can be filtered or cut through at the NFP layer, which conserves valuable PCIe bandwidth and reduces x86 CPU cycles. Through this heterogeneous architecture, the general-purpose multicore processors can focus on the compute workloads they are best suited for such as behavioral heuristics, Perl Compatible Regular Expression (PCRE) processing, content inspection, and analysis or other similar applications.

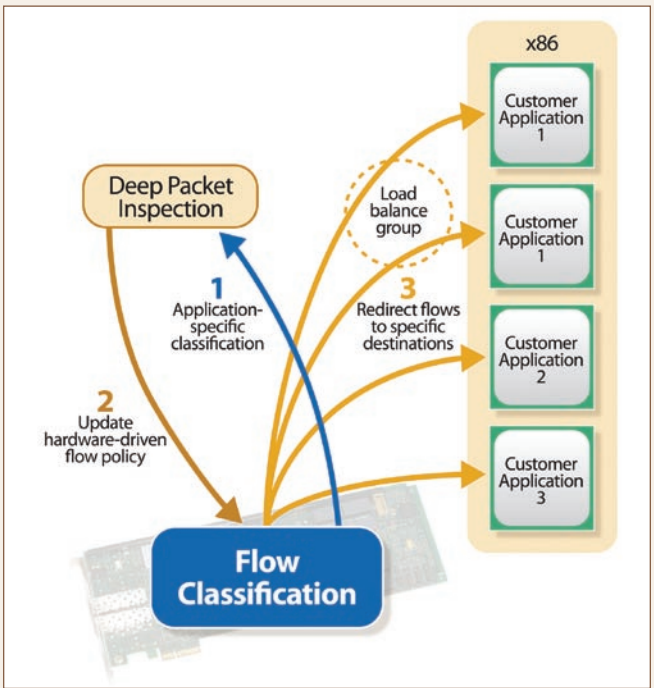

Figure 3 | Via a set of software APIs between the x86 CPUs and NFP cores, the treatment of flows can be updated in real-time. This flexibility is essential when a specific portion of a flow is of interest for inspection.

Intelligent Networks at 40 Gbps

The mission-critical nature of our military networks driven by information-centric warfare and homeland security creates an opposing set of forces. Networks need to continue to scale to meet exponentially growing bandwidth demands, and enterprises need the ability to effectively monitor these networks at all packet and content layers with stateful network intelligence. To meet these needs, a distributed, multi-chip, heterogeneous multicore architecture is required, providing specialized workload processing to effectively scale applications to 40 Gbps and beyond.

Daniel Proch *is director of product management at Netronome responsible for their line of network flow engine acceleration cards and flow management software. He has 14 years of experience in networking and telecommunications spanning product management, CTO's office, strategic planning, engineering, and technical support. Previously, Daniel was with FORE Systems and remained with the organization through acquisitions by Marconi and Ericsson. Daniel has a BS in Mechanical Engineering from Carnegie Mellon and an MS in Information Science and Telecommunications from the University of Pittsburgh. He can be contacted at daniel.proch@netronome.com.*

Netronome
724-778-3290
www.netronome.com

We Won't Leave You High and Dry.

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

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

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

LYNUXWORKS™

Open.Reliable.Safe.Secure.

*limited quantities available

Editor's Choice Products

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

Adapters preserve the legacy display investment

Sometimes the introduction and then proliferation of new technology wares yields angst for the budget-conscious consumer or DoD department heads looking to stretch defense dollars as far as they can. While Mini DisplayPort and DisplayPort rapidly become the *de facto* display-to-computer interfaces in new mobile laptop or desktop computers, many already-owned-and-paid-for displays and monitors only support older video connections including VGA, HDMI, or DVI. However, StarTech.com's Mini DisplayPort adapters link old and new – and eliminate the need to purchase a new monitor or display. Sounds perfect for a military command and control center looking to next year's budget and tightening the monetary belt in the meantime.

And the investment-preserving adapters provide many notables. Among them: Mini DisplayPort adapters offer display resolutions up to 1920 x 1200, and are software-free for easy implementation. In addition, HDTV support renders resolutions up to 1080p. And the plug-and-play Mac- and PC-compatible adapters are available in three variants: Mini DisplayPort to DVI, Mini DisplayPort to HDMI, and Mini DisplayPort to VGA.

StarTech.com • RSC# 45104 • www.startech.com

Storage system satisfies mission-critical needs

Storage is everywhere, and with good reason. The world would be chaotic without ways to store data or even personal collectibles, and work would have to be repeated incessantly to replace fleeing data or belongings. On the mission-critical side, Themis Computer's RES-XStore storage system can help satisfy the thirst for storage by providing harsh-environment storage capacity of 12 TB. Tucked into a 1U (17" or 432 mm) chassis, RES-XStore utilizes four hot-pluggable canisters with a total capacity for twelve 2.5" SATA or SAS drives. A RAID controller supporting either JBOD or single disk at levels ranging from 0 to 60 is also provided, as is RAID volume management and RAID out-of-band or in-band management. (Out-of-band management is additionally afforded via Ethernet.)

Storage system to host server communication is facilitated by PCI Express x8 and an add-in host adapter. And since this communication is paramount in mission-critical scenarios, RES-XStore can withstand operating shock at 25 *g*, 20 ms and vibration at 3.0 Grms, 8 Hz to 2,000 Hz. It also complies with MIL-STD-810G, and, sporting a SATA HDD, RES-XStore's operating temp is 0 °C to 60 °C. And another plus: Its modular design simplifies service and upgrades.

Themis Computer
RSC# 45086
www.themis.com

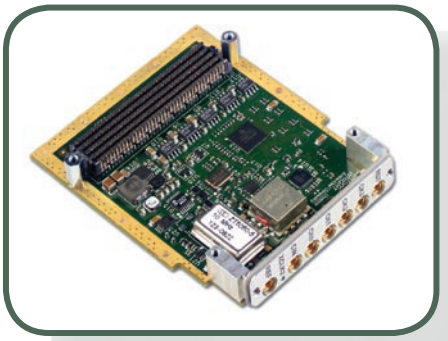

FMC clock generator boosts high-frequency data sampling

A high-performance, reliable clock generator is an integral part of high-frequency data sampling applications, and Curtiss-Wright Controls Embedded Computing's FMC-XCLK2 quad-channel clock generator FMC card provides that … and more. Designed for sampling apps such as radar, SIGINT, spectral analysis, SDR, and Electronic Counter Measures (ECM), the FMC-XCLK2 also, importantly, eases FPGAs' integration into embedded systems. The simplifying factor is its FPGA Mezzanine Card or FMC (VITA 57) form factor, designed to ease FPGA I/O at half a PMC's size.

The FMC-XCLK2 supports 50 MHz to >2 GHz RF output frequencies and provides a source for synchronizing and clocking I/O. In fact, the FMC card offers up to 4x phase matched outputs (two differential) – ideal for synching multiple I/Os – along with 10 MHz master reference output, ultra-low jitter, and selectable external/internal 10 MHz reference including RF output frequency multipliers. Input levels for the external 10 MHz reference are 50 Ohm AC coupled, with -5 to +5 dBm recommended. Meanwhile, the external RF clock input supports the same range. Effective in high-performance I/O wares like ADCs and DACs, FMC-XCLK2 is manufactured in either rugged conduction-cooled or air-cooled variants.

Curtiss-Wright Controls Embedded Computing • RSC# 45097 • www.cwcembedded.com

PCI2PMC adapter joins two worlds

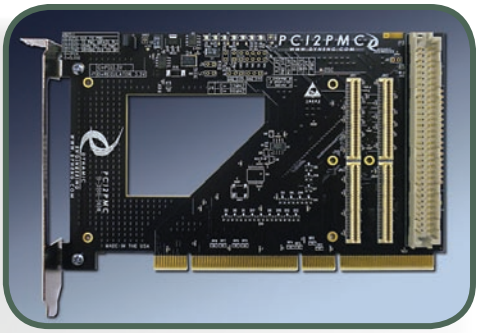

It almost looks like a texting acronym, but it's not one a teen would use. Rather, "PCI2PMC" is the model number of Dynamic Engineering's PCI-to-PMC adapter, enabling PMC card installation into either a half-length or standard PCI slot. The adapter actually measures 6.600" per the PCI specification, meaning it's an adaptable half-length ware with universal voltage. Fully loaded, PCI2PMC provides a passive design.

With 33/66 MHz bus operation, the adapter renders 32-/64-bit data transfers. PCI2PMC also supports 3.3 and 5 V PCI bus signaling, and +3.3, +5, +12, and -12 V can be supplied from the PCI backplane to the PMC with an optional jumper for PCI or a regulator at 3.3 V. Convenience in connection is provided via front- and rear-accessible connectors including a rear DIN64/SCSI connector. Compatible with third-party or Dynamic Engineering PMCs, the adapter's cut-out design enables increased airflow, and the unit can operate in the industrial temp range of -40 °C to + 85 °C. Meanwhile, build options for the 4.6 million-hour MTBF adapter include ditching the standard DIN for a SCSI connector, adding PrPMC clocking, conformal coating, and JTAG for convenient debugging, or an RoHS version. Software is low-maintenance, as the passive adapter merely enables usage of the original PMC software.

Dynamic Engineering • RSC# 45099 • www.dyneng.com

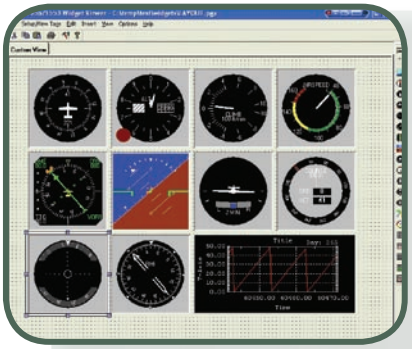

Bus analyzer GUI's got the look

Face it. Looks do matter: not just on the runway or new car lot, but certainly also when it comes to pitting character-based bus analyzers against modern Graphical User Interfaces (GUIs). Accordingly, one of the knockout features on GE Intelligent Platforms' version 7 of its BusTools-1553 bus analyzer software is its new "intuitive" GUI. The GUI simulates, tests, and analyzes 1553 data bus traffic on CompactPCI, PCI, VME, VXI, PC/104, and PC/104-*Plus* form factors, among others. The GUI interface hastens bus traffic analysis and monitoring and facilitates fast message modification/creation. It also enables simultaneous multiple-bus control in addition to error detection/injection and speedy filtering for either recorded or live displayed data.

And BusTools-1553 version 7 has even more new tricks up its sleeve. One is the Dynamic Bus Monitor stop/start feature, enabling users to achieve efficient on-the-fly 1553 bus traffic routing. Another is a one-page bus list editor, replacing the multi-page editor and permitting users to view every bus traffic message within a single window, then organize them quickly onto a highly readable one-page list. And finally, BusTools-1553 version 7's Selective Data Watch feature lets readers choose different data words from any bus message to identify elusive system issues, thanks to integrated high/low limit checking, automatic limit event logging and corresponding snapshot feature, and DDE output.

GE Intelligent Platforms • RSC# 45098 • www.ge-ip.com

ALM system offers one-stop-shop for safety-critical apps

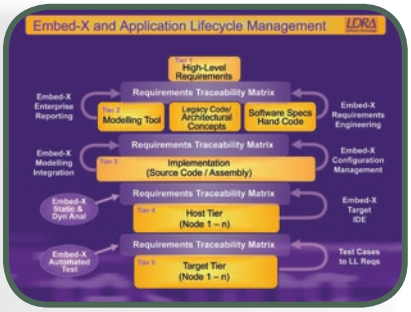

Life-cycle management is paramount in the defense industry, and not just at the component level. The application level, particularly for safety-critical apps like avionics, also holds its challenges in dovetailing all project phases. Enter Embed-X, an "end-to-end" Application Lifecycle Management (ALM) system from LDRA and Visure Solutions. Embed-X integrates the entire software engineering process, from project management, to requirements management, to architecture, code creation, software configuration, and finally testing. Now requirements, for example those for DO-178B, can be found and used dynamically within the system alongside the tangibles from all other project phases, rather than being isolated and located elsewhere.

To maintain full phase integration, Embed-X reports on the progression of development and enforces inter-phase development processes. LDRA and Visure Solutions estimate that defense and avionics vendors often experience 2x budget overruns, but suggest that Embed-X's one-stop-shop efficiency could yield as much as 50 percent in cost savings. Where does its efficiency come from? Safety-critical requirements tracing is executed via dynamic and static analysis all the way through to testing and verification. Additionally, Embed-X is compliant with MISRA, security standards including the Homeland Security Agency's Common Weakness Enumerations (CWE) and Cert C, plus safety-critical standards including DO-178B.

LDRA • RSC# 45100 • www.ldra.com / Visure Solutions • RSC# 45101 • www.visuresolutions.com

FPGA starter kit provides the whole package

FPGAs are prevalent in all sorts of military embedded applications and systems these days, and form factors such as AdvancedTCA, MicroTCA (collectively referred to as "xTCA"), and AdvancedMC (AMC) are increasingly following suit. Thing is, how are FPGAs integrated with these historically telecom form factors? Pigeon Point Systems' Module Management Controller (MMC) Board Management Reference (BMR) Starter Kit provides the key by rendering all firmware needed for FPGA development, with the aim of quicker deployment of AMC products. The kit also includes a benchtop management controller development board in AMC form factor style and a full SmartFusion FPGA design tucked inside a Libero Integrated Design Environment (IDE). A production license and "comprehensive documentation" round out the entire picture. And xTCA developers will be happy: Pigeon Point makes two benchtop variants of the kit: One for MicroTCA, one for AdvancedTCA.

Meanwhile, the kit's foundational product is the SmartFusion intelligent mixed-signal FPGA product including built-in flash. What makes SmartFusion unique is its composition: an FPGA, programmable analog, plus a 32-bit ARM Cortex-M3-based Microcontroller Subsystem (MSS) at 40 MHz. With the ARM processor experiencing zero load while SmartFusion renders advanced analog processing, the processor is left to execute xTCA analog sensor monitoring. Now that's what we call efficient.

Pigeon Point Systems • RSC# 45102 • www.pigeonpoint.com

Find errors your way – ahead of time

In safety-critical systems, failure is simply not an option. Such thinking was surely the impetus behind aicas GmbH developing their Java-based VeriFlux automatic static code analysis tool. (The tool additionally works well on non-safety-critical Java code.) While testing always mandates executing yet another run for each code path, VeriFlux can analyze all code paths simultaneously. Not only that, VeriFlux works to prevent runtime exceptions by pinpointing everything from null pointer errors to arithmetic issues. Deadlocks, especially in multicore programs, can be flushed out with the static analysis tool, and race conditions can be prevented as well. And if the system carries dead weight (aka "dead code"), it won't get past the tool's highlighting.

One thing that makes VeriFlux particularly intriguing, however, is its configurability: It takes orders from the user, meaning that it can be "told" to focus on certain aspects of the analyzed code and to provide a briefer summary for the remaining code. It can also be used anytime in the development cycle, including analyzing incomplete programs or just a single method. Meanwhile, other notables include the tool's RTSJ memory that validates all scoped memory assignments, in addition to VeriFlux's automatic simple reflection or user-programmed complex reflection. Library modules are also included in the tool to provide ease of use.

aicas GmbH
RSC# 44362 • www.aicas.com

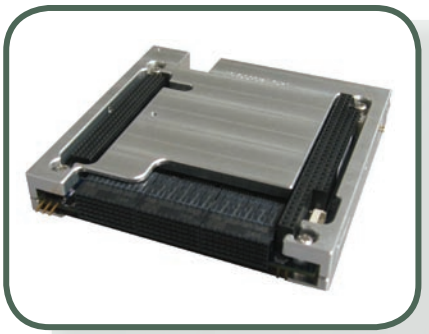

GbE switch card aids soldiers and their networks

Out in the harsh environmental terrain of Afghanistan, situational awareness is paramount to mission success and survival. One key enabler is network-centricity, allowing soldiers to receive strategic or protective orders from commanders and other personnel in the timeliest and most straightforward way. We're going to guess that developers at Parvus had just such a scenario on their minds when they developed the new COM-1268 rugged PC/104-*Plus* GbE switch card. Designed to heighten IP network-centricity and situational awareness, the 10-port board provides robust network performance even under the most demanding thermal and shock/vibe conditions.

Supporting both IPv4 and the later-and-greater IPv6 Quality of Service (QoS) traffic prioritization, the Layer 2 GbE switch survives – and thrives – in a wide temperature range of -40 °C to +85 °C. It also supports Simple Network Management Protocol (SNMP), Virtual Local-Area Network (VLAN) trunking, and Rapid Spanning Tree (RSP) redundancy. And it is vended as a stand-alone board … or not: It can also be integrated into a mission computer platform, mobile router system, or ruggedized COTS switch subsystem, for example. And there's even more good news: The COM-1268 does not need an additional processor for board operation and can also be used in non-PC/104 systems.

Parvus Corporation • www.parvus.com • RSC# 44031

Show me the data ... in SWaP-savvy style

Data gathering, formatting, and transmission are imperative to fostering mission success and soldier safety, particularly for manned military aircraft, for example. And the Aitech NightHawk RCU [Rugged Controller Unit and Data Concentrator Unit (DCU)] is a technology designed to that end – and more. Based on the compact Intel Atom N270 processor running at 1.6 GHz with 2 GB DDR2 SDRAM, the ultra-rugged 4.5-lb unit is designed for SWaP-conscious airborne or ground vehicles in the defense and industrial arenas. With a slim profile sporting natural convection/radiation cooling, the company reports that NightHawk RCU "dissipates over 22 W at +55 °C in free air, or at up to +71 °C with an optional low-pressure fan." A triad of rugged MIL-DTL-38999 connectors additionally boosts robustness.

Specifically suited for harsh-environment chemical, climatic, mechanical, and electrical apps, NightHawk RCU's data storage is accomplished via a standard offering of 4 to 8 GB SSD or an option to increase to 250 GB SSD. Accoutrements include standard I/O interfaces: dual GbE, stereo audio in/out, RS-232/422/485, PS/2 ports, video graphics, and many more. Meanwhile, NightHawk RCU's options include WAN wireless radio or Wi-Fi, to enable the RCU to execute periodic data monitoring or logging to a home base. A second option is a preformed cable set, which makes the unit prototype-ready by merely plugging it into Ethernet and a standard keyboard, mouse, and monitor.

Aitech Defense Systems, Inc. • RSC# 44263 • www.rugged.com

WIN-T router hits the streets

Many technologies are developed for a particular military program, but then what? Sell it as-is to defense and civilian markets? Yes ... if you're Juniper Networks, which recently released its LN1000 mobile secure router – used in the U.S. Army's Warfighter Information Network-Tactical (WIN-T) program – for public consumption. (Well, at least for industries such as defense, public safety, utilities/energy, smart grid, and others.)

Available in the VPX form factor (4" x 6" x .85"), the 1.5-lb LN1000 provides transmission of data, video, and audio traffic for data aggregation, surveillance, or comms apps. But the notable here is the mobile security, which sounds like an oxymoron but isn't. Indeed, LN1000 securely interconnects platforms such as remote monitoring or sensor stations, UAVs, and so forth to their operations centers or central command. Also afforded are resiliency in networking and high performance with low power consumption: a mere 35 W. Designed for tough environments, this mobile secure router withstands temps from -40 °C to +85 °C. And deployment onto existing platforms is no problem either: LN1000's conduction-cooled design eliminates the need for external power.

Juniper Networks, Inc.
RSC# 45096
www.juniper.net

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

Crosshairs Editorial

Embedded Systems Conference portends COTS' future

By Chris A. Ciufo, Editor

Every year CMP's Embedded Systems Conference offers a glimpse into technologies likely to find their way onto tomorrow's battlefield. At the close of this week's ESC (ending 29 April 2010), there were four key markets that countless vendors emphasized: Android, industrial, medical, and military. Yeah, that last one surprised me, too. Rarely do software, IC, and network infrastructure vendors talk about the military market. But unlikely companies from Mentor Graphics to Dell talked about their defense market plans. After meeting with nearly 50 vendors, here are some COTS trends to watch out for during the next 18 to 24 months.

Smartphones; video

Warfighters want the same rich media and Internet connectivity on the battlefield as they have in their off time. Hence, UAS video feeds, Blue Force Tracking, and moving GPS maps make sense on portable handhelds. And DoD brass wants to give it to them. I discovered at least two contracts – one from the NSA and one from DARPA – aiming to create Android-based and military-secure iPhone-like handsets. Already, companies like Thales offer secure BlackBerries with https and IPsec security. And consumer devices like Apple's 3GS iPhone make the video experience so compelling that it's likely all dismounted soldiers will eventually tap into some data stream, somewhere.

Cloud computing and "point-of-the spear" processing

Not all battlefield computers need heavy horsepower. Thin client devices relying on ARM Cortex M3 or QorIQ CPUs connected to fat pipes and rack-mount servers such as AdvancedTCA or 2U rigs from Emerson, RadiSys, or Z Microsystems would be cheaper, lighter, and use less power; they would also still provide access to information from elsewhere in the GIG. This is the cloud computing model first theorized in X-Windows a decade ago. Intel's MIDs vision maps perfectly to many battlefield sensors and embedded platforms to realize cloud computing. Then again, some devices such as UAS payload sensors or comms jammers need all the performance they can get in the local box. New DSP offerings in Altera Stratix V devices or Xilinx Virtex-6 FPGAs bolted to Core i5 Intel CPUs with Quick Assist bring massive IOPS to low-power embedded platforms.

Fatter pipes

The digitized battlefield needs node connectivity, whether via satellite, mesh networks such as SNAP from Synapse, or an aloft UAS relay station. But routing video everywhere and to everyone requires more bandwidth. Trouble is, current technologies are maxed out, and forklift upgrades are cost-prohibitive. Instead, COTS technologies like JPEG2000 compression, H.264, and deep packet inspection will reduce the amount of data being piped around the battlefield. New CPUs from ARM, Freescale, Intel, and VIA are all optimized for one or more of these data reduction schemes.

Application Life-cycle Management (ALM)

Not a new technology, ALM is moving into defense system designs through offerings such as LDRA's TBReq and Mentor Graphics' new ReqTracer. Tools like these collect and track requirements management snippets, making sure that complex, multiyear programs meet evolving specs. Previously, myriad SOW CDRLs were (at best) tracked via complex, iterative Excel spreadsheets.

Security

We mention them often in *Military Embedded Systems,* yet high-assurance hardware and software are still rare in all but the most classified of comm systems. But with every embedded system now a node on someone's network, information assurance to protect assets from cyber attack is becoming a baseline requirement. Secure operating systems meeting (or soon to be meeting) Common Criteria, MLS, or MILS are available from Green Hills, Wind River Systems, and LynuxWorks. For data on-the-fly, Mocana provides security assurance.

Safety-critical and MILS

These two terms are often used together but are not synonymous. Safety-critical specs such as DO-178B (soon, C) and DO-254 provide certification that code and hardware behave predictably. The same RTOS vendors listed above offer safety-critical offerings, as does AdaCore with their Praxis partnership. MILS, on the other hand, will soon extend beyond just security. How? RTOS vendors have revamped their hypervisors to provide protected partitions, which now offer the added benefit of running a guest OS in a protected partition. Vendors at ESC including Enea demonstrated how even a retail Windows XP installation could run in its own partition and if crashed or compromised, wouldn't affect the rest of an embedded system. Partitioned environments also allow legacy operating systems to be moved forward into modern, multicore hardware. Lastly, LynuxWorks demonstrated secure Windows machines where the Windows OS is run atop a MILS hypervisor.

Small-form-factor shoeboxes

I've talked about this for years. It's where a vendor such as WinSystems, ADLINK, Parvus, or others build a custom small-form-factor box with the "right" amount of I/O and processing power for the system. Despite the best efforts of consortia like SFF-SIG, PICMG, or the PC/104 Consortium, DoD designers care less about the card *inside* the box than about the box itself. And since shoeboxes are inexpensive and available in myriad ruggedization levels, they can often be tossed out during the next tech refresh.

And lastly, more with less

Right now, the industry is girding for reduced DoD budgets as O&M sucks money from RDT&E, despite President Obama's $708B request (Feb10) versus $690B in FY10. Worse, as SECDEF Gates puts the brakes on underperforming programs, primes expect to have to rejustify many line items. That means funding delays that ripple down to the COTS industry. The good news is: This is *deja vu* all over again. COTS has always been the way to get the best tech to the front, fast.

Chris A. Ciufo, cciufo@opensystemsmedia.com

GE
Intelligent Platforms

Welcome to the Knowledge Bank.

A wealth of free COTS white papers and tutorials.

Information may be the ultimate competitive advantage, and that's why we offer you free access to our extensive COTS white paper library. With dozens of papers currently available for download, you'll find information ranging from ARINC 429 and MIL-STD-1553 tutorials, to an examination of the latest Intel® Core™ i7 processors. Who knows what great ideas you may find in one of our papers? It could be the idea that gives you an important competitive edge in developing and delivering the solutions your customers are demanding.

www.ge-ip.com/cots-papers

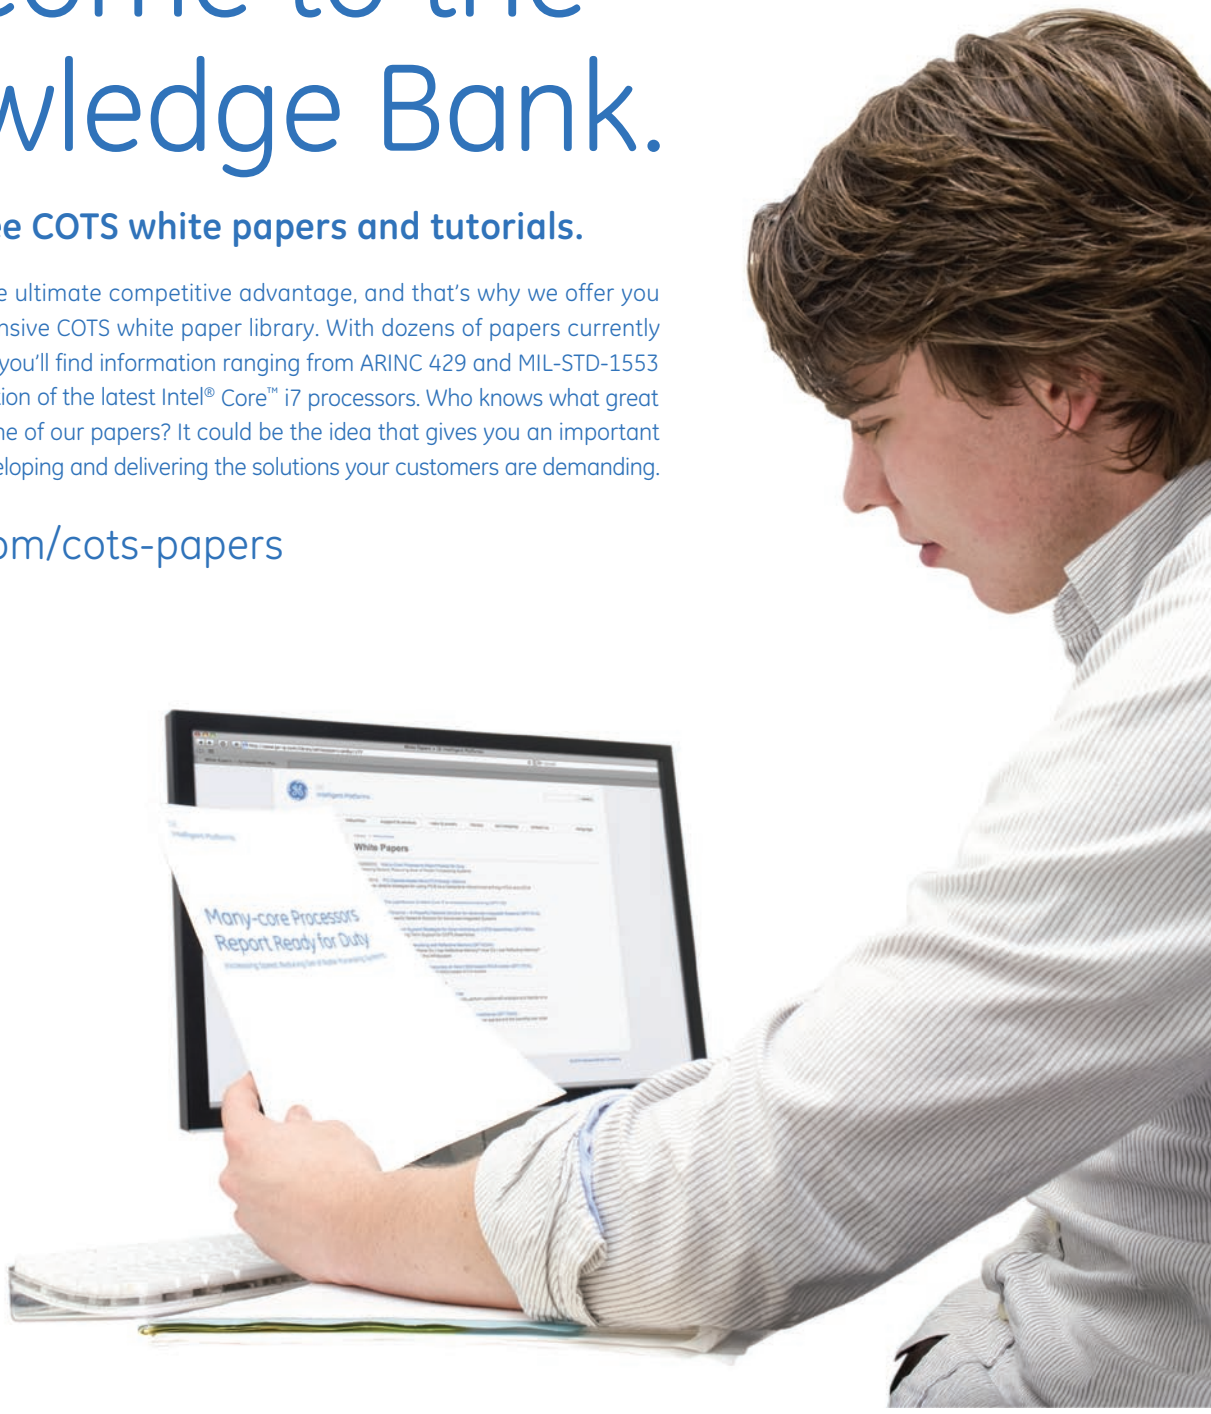

imagination at work

SPEED YOUR TIME-TO-MARKET

YOUR TOTAL SOLUTIONS PARTNER

1 Signal Acquisition
Analog & Digital I/O
ADC510 (FMC) & FMC-XCLK2

2 Radar Processing
Digital Signal Processors
CHAMP-AV5

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

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

4 Graphics Display
Video Input & Output
XMC-710

5 System Connectivity
Ethernet Switches
VPX6-684 FireBlade

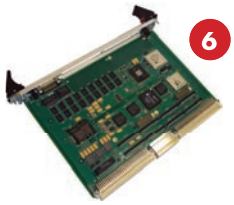

6 Command & Control
Single Board Computers
VPX6-187 & VME-1905

CURTISS WRIGHT Controls
Embedded Computing

sales@cwcembedded.com cwcembedded.com

BOARDS & SYSTEMS ABOVE & BEYOND