MIL-EMBEDDED.COM

COTS and technology for the entire military life cycle

Military EMBEDDED SYSTEMS

VOLUME 4 NUMBER 1
JAN/FEB 2008

IN THIS ISSUE:

Top 10 technologies for the war fighter

Space logic

ALSO:

Open source off-the-shelf s/w

OpenSystems Publishing

Discover Chassis COTS Standards with CM COMPUTER

and enter into another world with our new CM-ATR-25/35/45 Series. Up to 15 enclosure models offering outstanding Performance, oversized Power Capacity, optimum Heat Management and a wide range of Accessories & custom Options.

Decidedly well suited for critical Defense & Aerospace systems requiring premium quality that can be provided only by leading-edge technology.

Brief Chassis Specification Guide

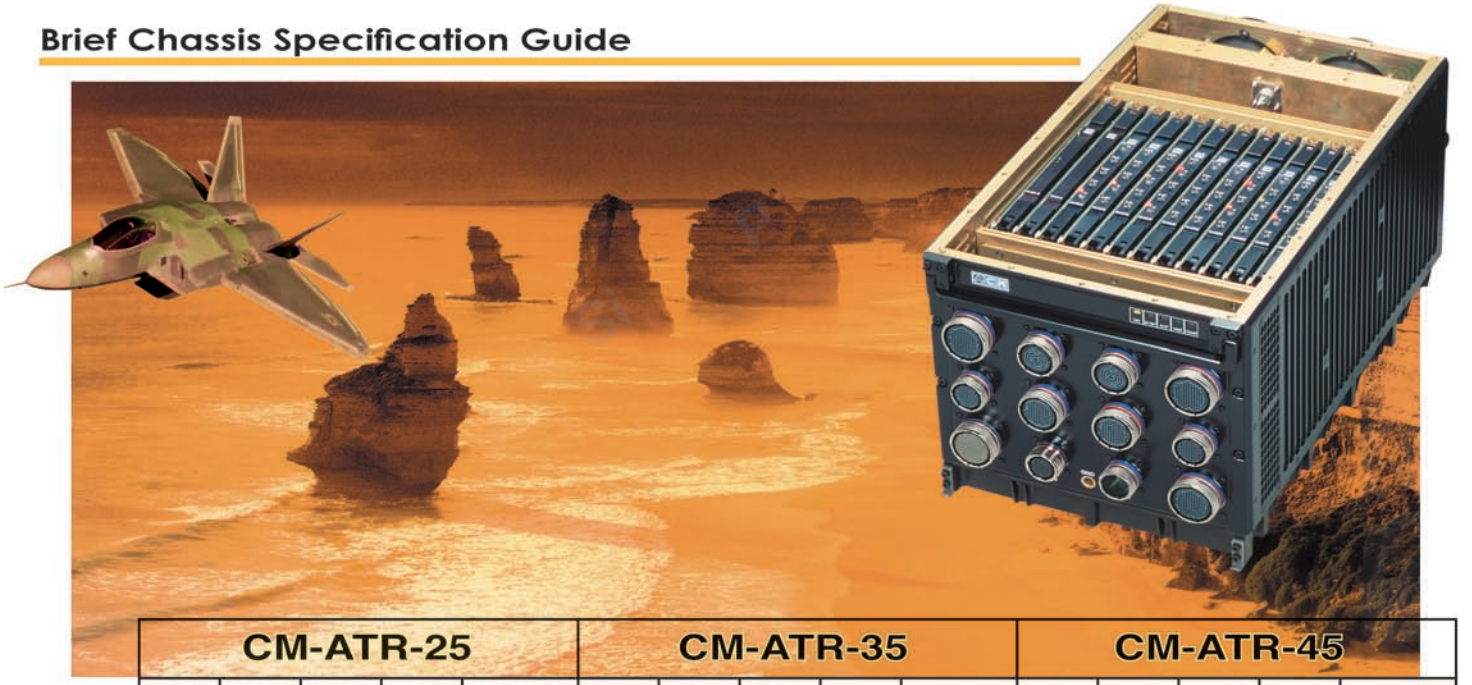

	CM-ATR-25					CM-ATR-35					CM-ATR-45				
(*)	FAC	S	SEF	HES	SIXHEX	FAC	S	SEF	HES	SIXHEX	FAC	S	SEF	HES	SIXHEX
Size / Slots	½ ATR - Long / 5 Slot 6U VME64x or cPCI					¾ ATR - Long / 7 Slot 6U VME64x or cPCI					1 ATR - Long / 12 Slot 6U VME64x or cPCI				
Power Supply	400 Watts				575 Watts	500 Watts				775 Watts	1050 Watts				1450 Watts
+5 / +3.3 VDC	20 Amp / 23 Amp				40 A / 23 A	40 Amp / 23 Amp				80 A / 23 A	80 Amp / 45 Amp				160 A / 45 A
±12 VDC	12 Amp each					12 Amp each					20 Amp each				
Chassis Vin	All PSU models accept: 28 VDC; 48 VDC; 270 VDC; Autorange 90-264 VAC @ 47-880 Hz and 200 VAC 3-Phase @ 47-880 Hz														
Board Format	Slot-by-slot user configured card-cage allows intermixing conduction-cooled IEEE-1101.2/ANSI-VITA 30.1 and air-cooled IEC-297/IEEE-1101.1 boards														

(*) Cooling Options: FAC - Flowthrough Air Cooled; **S** - Sealed; **SEF** - Sealed, Extended Fins; **HES** - Sealed, Heat Exchanger Sidewalls; **SIXHES** - Sealed, Six Heat Exchangers

CM Computer
True Military COTS Products

All our ATR products are delivered Tested and Certified by independent authorized Labs per MIL-STD-461E & MIL-STD-810F for immediate deployment in US Navy & US Air Force military Fighters and Helicopters.

For more information visit **www.cmcomputer.com** or contact us at **info@cmcomputer.com** to request our Chassis Catalog

Elcard™ Wireless LAN Modules

Designed for Industrial and Professional Applications

PC/104+ WIB4xx dual WLAN module

USB-attached WLAN

- PC/104+, USB, and PCI versions
- IEEE 802.11b/g/a/h WLAN standards
- At 2.4GHz up to 11 and 54/108Mbps bandwidths
- At 5 GHz up to 54/108Mbps bandwidths
- Dual antenna diversity
- Extended temp versions available (-40°C to 85°C and -20°C to 70°C operating)
- Rugged and shock resistant, high altitude operation
- Long term supply
- O/S support for Linux, Microsoft™ Windows™ XP/2000/NT/98SE/ME

Rugged Access Points Available

- Dual WLAN versions available (WIB400 series)
- Evaluation kits for easy start-up
- Ranges of 1 mile+ can be reached even at 100mW Tx power with our directional antennas
- Ranges of several miles can be reached with our power amps and special antennas
- WIB250 WLAN module provides dual band 802.11g/ 2.4GHz & 802.11a/5GHz with two antenna connectors

-40°C to +85°C Operating Temperature Range Versions Available

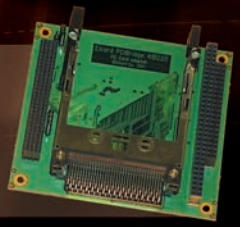

AIB220 PC/104+ Cardbus

- Cardbus/PCMCIA Adapter
- Dual Type I/II or single Type III
- Linux and Win9x/2K/XP support
- 3.3V and 5V card support
- TI PCI1420 chipset

Elcard

Elcard USA
10849 Kinghurst, Suite 105
Houston, Texas 77099
Toll Free: 800-688-4405
Phone: 281-568-4744
Fax: 281-568-4604
Email: sales@elcard-usa.com
Web: www.elcard-usa.com

Military EMBEDDED SYSTEMS

January/February Volume 4 Number 1

COLUMNS

DEPARTMENTS

ON THE COVER:
On location at Camp Liberty, Iraq. Airman Gevoyd Little of the 447th EOD (Explosive Ordnance Disposal) operates this remote explosive detection robot looking for IEDs. A COTS device itself, the robot relies heavily on some of the Top Ten Technologies we're discussing in this issue (page 34).
(U.S. Army photo taken by SSG Kevin Moses, Sr.)

WEB RESOURCES

Subscribe to the magazine or E-letter:
www.opensystems-publishing.com/subscriptions

Live industry news:
www.mil-embedded.com/news
www.opensystems-publishing.com/news/submit

Submit new products:
www.opensystems-publishing.com/vendors/submissions/np

E-LETTER

www.mil-embedded.com/eletter

CORRECTION: In the Nov/Dec *Crosshairs Editorial* column "MILCOM exposé," Roger Krone of Boeing was identified as President, Boeing Integrated Defense Systems (IDS); however, his correct title is President of Network & Space Systems (NSS), an IDS subsidiary reporting $11.9 billion in earnings for 2006. (2007 figures unavailable at press time.)

Published by: OpenSystems Publishing™

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

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

Made in the USA

Ultimate Modularity

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

Direct Seamless Connections with no Data Reduction

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

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

OpenSystems Publishing

Advertising/Business Office

30233 Jefferson Avenue
St. Clair Shores, MI 48082
Tel: 586-415-6500 ■ Fax: 586-415-4882

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Business Manager
Karen Layman

Sales Group

Dennis Doyle
Senior Account Manager
ddoyle@opensystems-publishing.com

Tom Varcie
Senior Account Manager
tvarcie@opensystems-publishing.com

Doug Cordier
Account Manager
dcordier@opensystems-publishing.com

Andrea Stabile
Advertising/Marketing Coordinator
astabile@opensystems-publishing.com

Christine Long
E-marketing Manager
clong@opensystems-publishing.com

Regional Sales Managers

Jerry Bleich
New England
jbleich@opensystems-publishing.com

Ernest Godsey
Central and Mountain States
egodsey@opensystems-publishing.com

Barbara Quinlan
Midwest/Southwest
bquinlan@opensystems-publishing.com

Denis Seger
Southern California
dseger@opensystems-publishing.com

Sydele Starr
Northern California
sydele@pacbell.net

Ron Taylor
East Coast/Mid Atlantic
rtaylor@opensystems-publishing.com

International Sales

Dan Aronovic
Account Manager – Israel
daronovic@opensystems-publishing.com

Sam Fan
Account Manager – Asia
sfan@opensystems-publishing.com

Reprints and PDFs

Nan Lamade: 800-259-0470 • mesreprints@opensystems-publishing.com

ADVERTISER INFORMATION

Page	Advertiser/Product description
11	ACCES I/O Products, Inc. – The source for all your I/O needs
31	ACT/Technico – Embedded sub-systems
21	Acumen Instruments – DataBridge SDR-CF
24	Advantech Corporation – Total rugged solutions for Mil
20	Alphi Technology Corporation – Modular solutions
5	Annapolis Micro Systems, Inc. – Wildstar 5 for IBM blade
42	Avionics – Five reasons to register now
26	BMC Communications – Protocols converter
2	CM Computer – Discover chassis COTS standard
19	Concurrent Technologies, Inc. – Dual-core processors
41	Connect Tech Inc. – Solid-state storage
52	CWCEC: Data Communications – Your data is critical
15	Excalibur Systems, Inc. – Dense
51	GE Fanuc Intelligent Platforms, Inc. – AXISLite
7	Hypertronics – Unite
3	Microbus Inc – Elcard wireless LAN modules
43	Military and Aerospace Electronics – Under the COTS Umbrella
16	MPL AG – You need
37	Performance Technologies – What's on your radar
27	Phoenix International – Mission critical
33	Schroff, a brand of Pentair Electronic Packaging – When your mission is critical
49	Super Talent Technologies – Reinventing storage
21	TEWS Technologies LLC – COTS I/O solutions
23	Tri-M Systems Inc. – Tri-M Engineering
25	Tri-M Systems Inc. – Tri-M Systems
38	VersaLogic Corp. – The only thing missing is the lead
9	VMETRO – The sky is the limit
35	White Electronic Designs – Advanced Microelectronics solutions
28	Winchester Electronics – Power connector solutions

Military EMBEDDED SYSTEMS

AN OPENSYSTEMS PUBLICATION

Military and Aerospace Group

- DSP-FPGA.com Resource Guide
- DSP-FPGA.com
- DSP-FPGA.com E-letter
- Military Embedded Systems
- Military Embedded Systems E-letter
- PC/104 and Small Form Factors
- PC/104 and Small Form Factors E-letter
- PC/104 and Small Form Factors Resource Guide
- VME and Critical Systems
- VME and Critical Systems E-letter

Group Editorial Director Chris A. Ciufo, cciufo@opensystems-publishing.com
Senior Editor (columns) Terri Thorson, tthorson@opensystems-publishing.com
Associate Editor Sharon Schnakenburg, sschnakenburg@opensystems-publishing.com
Assistant Editor Robin DiPerna
European Representative Hermann Strass, hstrass@opensystems-publishing.com
Senior Web Developer Konrad Witte
Web Content Specialist Matt Avella
C eative Director Steph Sweet
Art Director David Diomede
Graphic Coordinator Sandy Dionisio
Circulation/Office Manager Phyllis Thompson, subscriptions@opensystems-publishing.com

OpenSystems Publishing

Editorial/Production office:
16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystems-publishing.com

Publishers John Black, Michael Hopper, Wayne Kristoff
Vice President Editorial Rosemary Kristoff

Communications Group

Editorial Director Joe Pavlat
Assistant Managing Editor Anne Fisher
Senior Editor (columns) Terri Thorson
Technology Editor Curt Schwaderer
European Representative Hermann Strass
Senior Designer Joann Toth

Embedded and Test & Analysis Group

Editorial Director Jerry Gipper
Editorial Director Don Dingee
Senior Associate Editor Jennifer Hesse
Special Projects Editor Bob Stasonis
European Representative Hermann Strass

ISSN: Print 1557-3222

Military Embedded Systems (USPS 019-288) is published eight times a year (January/February, March/April, May, June, July/August, September, October, November/December) by OpenSystems Publishing LLC, 30233 Jefferson Avenue, St. Clair Shores, MI 48082.

Subscriptions are free to persons interested in the design or promotion of *Military Embedded Systems*. For others inside the US and Canada, subscriptions are $28/year. For 1st class delivery outside the US and Canada, subscriptions are $50/year (advance payment in US funds required).

Canada: Publication agreement number 40048627
Return address WDS, Station A PO Box 54, Windsor, ON N9A 615

POSTMASTER: Send address changes to *Military Embedded Systems*
16872 E. Ave of the Fountains, Ste 203, Fountain Hills, AZ 85268

HYPERTAC
HYPERTRONICS

UNITE

HIGH DENSITY AND RELIABILITY

In designing high density military applications, it is critical that the hardware components provide environmental immunity united with connection reliability to meet the demanding requirements for signal integrity, consistent resistance and operational performance in extreme conditions. Hypertronics HyperGrid™ compression mount connectors adhere to any footprint down to 0.4mm pitch and are suitable for high speed microwave MMIC, flat panel displays, flexible cables, printed circuit and parallel boards, and RoHS compliant solderless, multi-termination connections.

HyperGrid Features	**Benefits**
Micro Miniature Contacts	Reduced Weight and Space
Low Contact Self Inductance	Superior Signal Integrity
Up to 37 GHz Speeds	Exceptional AC Performance
Modular Design/Custom Footprints	Fully Adaptable Z-Height and Form Factors
100,000+ Mating Cycles	Robust Enough for Dynamic Environments

Hypertronics interconnect solutions are choice of engineers worldwide for mission critical systems.

www.hypertronics.com/unite

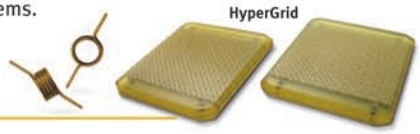

HYPERTRONICS: WHEN FAILURE IS NOT AN OPTION

smiths
bringing technology to life

Off-the-shelf software enables embedded training

Training has always been a top priority for the armed forces, and much of their time is spent honing their preparedness to deal with conflicts and operate the complex and sophisticated weapons platforms at their disposal. Training takes many forms – including classroom, procedural, simulated, and embedded through to the live exercise. Each type is designed to prepare for every eventuality from the instinctive reaction to the intuitive reasoning required to deal with a real-life scenario. This need is reflected in the design of many weapons platforms where systems are now being architected so that their human interfaces (screens, keyboards, and pointing devices) are accessible both to live sensors and weapons and to the synthetically generated scenarios of embedded training.

Being able to effectively operate each of the various subsystems making up a weapons platform is arguably the most essential training that is required. The skills to operate a complex weapon or Naval tactical console cannot be taught in the classroom environment alone and need to be reinforced regularly, which can only be achieved using a live platform. Typically such a training system will have four components:

- Scene simulator
- Scene builder
- Scenarios
- Hardware and integration

The *scene simulator* determines how the training system will display the scene to the trainee. It will offer a view of the scene as if it were being played back through a particular sensor type, for example, as perceived by a radar or an electro-optical or infrared sensor. The scene simulator also plays the scenario, moving objects and background as the training scenario plays out, reacting to inputs from the trainee and, optionally an instructor. Most systems will offer the ability to record operator inputs and to rerun any part of the exercise at the trainee's or instructor's discretion.

A key part of any training system is achieving a level of realism of the background and objects in order to fully engage the trainee in the exercise. A *scene builder* is the package that is used to create the objects that are to be displayed. Objects may be anything from people, weapons platforms, airplanes, trees, hills, houses, and so on. Objects can be defined in three dimensions and rules created for their movements and interaction with other objects. Many packages will include a library of objects that can be reused or modified as required.

Figure 1

A training *scenario* is the sequence of events that the training system plays out to the trainee. Scenarios are built offline for each training session to meet whatever the objectives may be. A platform will usually have a number of *canned* scenarios available for regular training exercises.

Because embedded training takes place on the same equipment as is normally deployed in live operations, the training system must have independent access to any display monitors and data entry devices such as keyboards, pointers, push buttons, and so on.

More modern systems have been designed to accommodate this, but older legacy systems will require extensive modification to support the addition of an embedded training system. Often such a system is PC-based with specialized interfaces for the data entry devices plus a high-performance graphics accelerator for display. Of course, training exercises could be embedded within a vehicle's overall operational software package, obviating the need for any additional hardware dedicated to training. However, this brings with it concerns for safety and security unless the training is very robustly partitioned from the operational software; firing live missiles or making compromising radio emissions must be avoided at all costs. It is also difficult and more costly to just update the training package as the entire operational software package may need to be requalified if any part of it is modified.

An example of an embeddable training package that can be used to introduce realistic training exercises into existing platforms is the SIVET-SE produced by GE Fanuc. It can be hosted on most high-performance, embeddable PCs from commercial grade through to full military specification and is targeted primarily toward the simulation of electro-optical sensor types. Figure 1 illustrates a typical workstation in an armored vehicle that could be supported by the embedded training concepts of SIVET-SE.

Training is an essential ingredient of every type of system used by the armed forces. While the eventual solution will be to equip every platform with training embedded from within the operational software, there are many legacy platforms in need of life extension where the availability of proven, off-the-shelf tools and hardware makes the provision of embedded training an economic reality. As these products become more sophisticated through continual product investment, they will converge with wholly embedded software, extending and strengthening this important COTS software market segment.

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

The sky is the limit

with FPGA-based DSP solutions for Electronic Warfare & ISR systems

FPGA-based Analog I/O:
- *3 GSPS ADC (AD3000)*
- *Dual 1.5 GSPS ADC (AD1500)*
- *Dual 105 MSPS ADC (ADC-MOD1)*
- *Dual 210 MSPS DAC (DAC-MOD1)*

FPGA-based Digital I/O:
- *Camera Link (CAML-MOD3)*
- *LVDS (LVDS-MOD3 & LVDS-MOD4)*
- *sFPDP, FPDP & FPDP-II (SFM & DPIO2)*
- *RS485/422 (RS485-MOD2)*
- *10 Gb Ethernet (V1020)*
- *Fiber Optic (PMC-FPGA03F)*

High-Performance Real-Time Processing Solutions:
- *VXS Processor with Freescale MPC8641D and Dual Xilinx® Virtex™-5 FPGAs (Phoenix VPF2)*
- *VXS Carrier with Dual XMC sites (Phoenix M6000)*
- *VXS Processor with Dual Virtex™-4 FPGAs (MM-1500 & MM-1550)*
- *VXS Recording Engine (Vortex VXS)*
- *VXS Switch Card (Phoenix CSW1)*
- *VPX Processor with Dual Virtex™-4 FPGAs (MM-1600 & MM-1650)*

Buffer Memory Nodes

Xilinx Virtex-5 based XMC/PMCs

VIRTEX™ 5

Embedded Computing - Data Recorders & Rugged Storage - Bus & Protocol Analyzers

For more information, please visit
http://www.vmetro.com or
call (281) 584-0728

VMETRO
innovation deployed

Brushing up on motor control

By John Wemekamp

Embedded technologies in ground-based fighting vehicles are often considered to be sensors and weapons plus command, control, and communications subsystems that directly assist the commander and crew to accomplish their mission. However, in addition to these high-profile, compute-intensive technologies, a modern fighting vehicle is a remarkably complex machine with large numbers of motor-driven components such as a turret, gun, hatches, fans, pumps, ammunition handling, and many others. The recent advances in high-power electric motion control and the introduction of hybrid electric drive to fighting vehicles arguably deserves a high placing in the top 10 list of technologies for the war fighter for speed of reaction, reduction of workload, reliability, and remarkable power density.

The majority of today's fighting vehicles use hydraulics for actuating their mechanical systems. Typically operating at 3,000 to 5,000 psi, hydraulics are prone to leakage and difficult to repair and maintain. In addition, hydraulic pumps, piping, and actuators are bulky and heavy compared to the latest generation of motor-driven subsystems that is set to replace them in the future. Whether the function to be performed is a simple, single-axis operation such as closing a hatch or the much more complex sequence of moves of an autoloader, all closed-loop servo systems share common control requirements to vary torque, acceleration, velocity, and position in order to accomplish their specific functions. Such servo systems use a controller/amplifier that modulates high-voltage dc (up to 610 V) to one or a number of motors to create the motion required. The controller will have a number of inputs such as analog, resolver, or encoder to sense position or velocity to close the loop.

Modern fighting vehicles have an integrated vetronics architecture supporting all the functions necessary to drive and control them. Based on redundant Ethernet or CANbus architectures, bused vetronics systems save on the mass of discrete point-to-point cabling that would otherwise be needed. All functions from external lighting to maneuvering the vehicle or rotating the turret will be transmitted over the vetronics bus as a series of commands to remote controllers located adjacent to the function to be performed. Motor controllers, like other vehicle functions, receive their commands via the vetronics bus. They are most often implemented as Line Replaceable Units (LRUs), complete with considerable intelligence to perform the servo function on command, in addition to handling background diagnostics and local error detection.

Controllers come in a range of types, controlling from one to many axes of movement from one LRU. For a complex piece of equipment such as an autoloader, commands received via the vetronics bus will choreograph a sequence of controlled motions from many motors to select, for example, the required type of ammunition from a dispenser, arm it, move it into position, and load it ready for firing. To accomplish this level of complexity, some controllers are programmable, their specific set of functions being created during development using modeling and simulation tools such as MATLAB and Simulink.

In ground-based fighting vehicles these mechanical components are often large and heavy, requiring a great deal of power to move and position them accurately and safely. Since vehicle weight and size are such critical parameters, demands for power to move heavy components must be traded against weight and physical size of the servo system. The controller is a key element and continuous development is targeted to reduce weight, power, and size of the controller and motors. The success of these efforts can be judged from Table 1, which compares the physical size and power delivery of a typical motor controller with a conduction-cooled VMEbus power supply module.

> "The recent advances in high-power electric motion control and the introduction of hybrid electric drive to fighting vehicles arguably deserves a high placing in the top 10 list of technologies for the war fighter for speed of reaction, reduction of workload, reliability, and remarkable power density."

This level of power density exceeds the capacity of conventional air- or conduction-cooling techniques, requiring both exacting thermal design of the LRU and its components and the employment of liquid cooling. This form of cooling is used in a range of motor controllers produced by Curtiss-Wright Controls Embedded Computing (CWCEC). Designed for a

	Dimensions (inches)	Volume (cubic inches)	Continuous power delivery (watts)	Power density (W/cubic inch)
VMEbus power module	6.2 x 9.3 x 1.6	92.3	750	8.13
Controller/amplifier	5.8 x 8.8 x 8	408	6,000	14.7

Table 1

wide range of high-power, vehicle-based motion control applications, the internal construction and cooling of an example controller are shown in Figure 1.

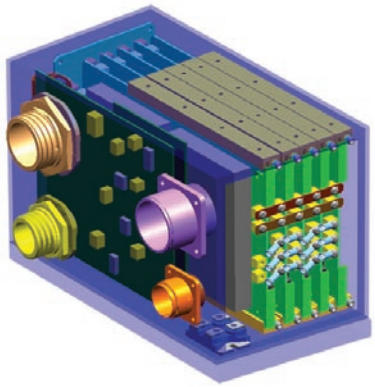

Figure 1

The precise control of the machinery's motion might not be perceived to be as glamorous as a cooperative, networked command and control system or the multiple compute nodes of a synthetic aperture radar processor, but motor control truly stretches its designers and technology to the limits. As a result of this evolution of compact motor controllers, size and weight of modern fighting vehicles will continue to be reduced. This will enable them to be more effective platforms, carrying greater payloads of troops, command and control systems, sensors, weapons, and self-protection systems.

For more information, e-mail John at john.wemekamp@curtisswright.com.

Military EMBEDDED SYSTEMS ★★★★★ Editor's Choice

Editor's Choice Product

Portable, rugged 3U RAID

Disk drives are so inexpensive these days – under $1 per GB – that we're saving everything instead of deleting it. But in deployed platforms and systems, where collecting and *storing* the data is the whole point, it's not enough just to have *beaucoup* storage. One needs assurances of reliable data storage. That's the point behind Phoenix International's RPC12 SAS/SATA II 3U RAID system: ruggedness, reliability, portability. Housing up to 12 HDDs in carriers, the 3U (5.25") system has redundancy and ruggedness built in everywhere you look.

Capable of withstanding 60 *g* shock at 2 ms, operating over 5 °C to 60 °C (even more with sealed drives), and up to 40,000 feet with sealed HDDs, this baby's in it for the long ride. It accommodates RAID 0, 1, 3, 5, 10, 50, NRAID, and JBOD as well as dual hot-swap PSUs and cooling fans. Host interfaces include 2 or 4 Gbps Fibre Channel, serial attached SCSI, or iSCSI. Input voltage ranges from 40 to 440 Hz, and from 90 to 240 VAC. As for software, the unit interfaces to Windows, Linux, and UNIX and includes a management GUI and failover software.

Phoenix International • www.phenxint.com • RSC# 35575

Daily Briefing: News Snippets

By Sharon Schnakenburg, Associate Editor

www.mil-embedded.com/dailybriefing

Australian and U.S. governments give each other space

… Or rather, the Australian and U.S governments are about to get into the same space: The U.S. Air Force has authorized Boeing to build a sixth Wideband Global SATCOM (WGS), and the Commonwealth of Australia is funding the endeavor as part of an agreement between the two nations. In exchange, the Australian Defence Force gains worldwide WGS service access. The new WGS satellite will be a Block II version of the 702 model spacecraft, supporting Radio Frequency (RF) bypass data rates of up to 311 Mbps (200x faster than typical DSL or cable connections).

Patents: Yours, mine, or ours?

In the wake of a continuing barrage of technology patent disputes, some contentions are finally getting resolved: Artesyn Technologies – recently acquired by Emerson Network Power, as was Motorola's Embedded Communications Computing group – offered a successful defense on most of the 82 patent infringement claims lodged by Power-One, Inc. surrounding 4 of Artesyn's patents. A Texas jury found only one patent infringement on a product never sold by Artesyn, and the jury determined that infringement was unintentional. A fee of $100 in damages may be awarded to Power-One, Emerson reports.

Green Hills and its Padded Cell

Is security the problem or the remedy? "Many people believe that hypervisors are the solution to security problems. But existing hypervisors actually make security problems worse by providing another avenue for attack," says Green Hills Software founder and CEO Dan O'Dowd. Accordingly, Green Hills has announced what it calls "the world's first secure hypervisor" or the *Padded Cell Secure Hypervisor,* which operates on top of the company's EAL6+ compliant INTEGRITY separation kernel and compartmentalizes and virtualizes guest operating systems. Additionally, it implements separation among virtual machines: Its virtualization software operates as an application and therefore cannot circumvent the separation kernel's security policies.

Increased optimization in a reduced size

It's all in a day's work – or at least in a day's announcement – for Intel, which unveiled 16 new lead- and halogen-free products at the International 2008 Consumer Electronics Show held in Las Vegas, Nevada. Among the offerings announced are Intel's first 45nm processors for tucking inside Intel Centrino Processor Technology based laptops. The chips leverage Intel's new transistor formula and increase PC speed, prolong battery life, reduce power consumption, and are available in smaller packages.

VITA's VPX scores with ANSI

Nearly four years in the making, VITA's *VPX base standard* (ANSI/VITA 46.0-2007) and its "dot standard" *VMEbus Signal Mapping for VPX* (ANSI/VITA 46.1-2007) have received American National Standards Institute (ANSI) ratification. VITA 46.0 describes the full suite of VPX dot standards and VPX's physical features. Meanwhile, VITA 46.1 enables legacy and latest-generation VMEbus technology to converge under VPX. The VPX backplane standard accommodates several switch fabric protocols including PCI Express, Serial RapidIO, and GbE, among others.

Double duty for new CWCEC Senior VP

Multitasking is required at most jobs these days … particularly if you're Curtiss-Wright Controls' new Senior VP Tom Quinly, who is also continuing in his role as president of the company's Embedded Computing Division. Quinly received his recent promotion to SVP, according to Curtiss-Wright Controls president David Adams, because of his strong performance within the company. Quinly has 25 years of experience in the defense industry and became part of Curtiss-Wright Controls with its acquisition of his employer, Dy4, in 2004.

Savi enters 'green' scene

Not only is "green" a growing trend in the U.S. and Europe, it's also *hot* in Kuwait, as Lockheed Martin subsidiary Savi Technology is proving. Savi's engineers recently installed Radio Frequency Identification (RFID) signposts and solar panels at a key U.S. Army supply depot in Kuwait. The new technology reportedly reduces costs, is environmentally friendly, and removes any need for electrical infrastructure installation in remote areas. Plans are also in the works with the U.S. DoD to bring solar-powered technology to additional Middle Eastern regions, Savi reports.

New iRobots to provide a theater of operations

More war fighters can stay out of harm's way, thanks to the U.S. Army's new $286 million xBot contract with iRobot Corp. The Indefinite-Delivery, Indefinite-Quantity (IDIQ) contract stipulates that the Army can acquire up to 3,000 military robots, along with training, spare parts, and repair services within the next five years. As the new robots will largely be deployed for mission-critical general infantry force purposes, the contract indicates a departure from the Army's earlier strategy to only deploy robots to Explosive Ordnance Device (EOD) specialists.

RTI and Northrop Grumman: *CLIP*ping out military data issues

Northrop Grumman and Real-Time Innovations (RTI) recently joined forces to "CLIP" a new DDS-compliant interface into the U.S. Navy and Air Force Common Link Integration Processing (CLIP) system. Aimed at solving inherent incompatibility between the two branches' Tactical Data Links (TDLs), CLIP enables legacy platforms connection to IP-based systems. It also enables TDL processing for platforms without their own data link via bridging a new TDL radio/terminal to the legacy mission computer's software. CLIP's new DDS interface is Net-Centric Enterprise Solutions for Interoperability (NESI) compliant and purported to ease TDL integration.

Chassis Plans partners up

Some say "less is more," but Chassis Plans may beg to differ. The mission-critical COTS and custom computer systems manufacturer recently earned the title of Microsoft Certified Partner, which enables its access to additional training, software licensing, and technical support channels. Steve Travis, Director of Program Management at Chassis Plans, says, "Our customers can't afford downtime with these systems installed in Iraq and other mission-critical arenas. Gaining Microsoft Certified partner status fits within our ISO:9001 plan and shows our customers our commitment to quality in the products we ship them."

Fast path to DO-178B certification

Companies laboring for DO-178B certification might find a speedier route if they use DDC-I's mixed-language development environment OpenArbor, which now supports the FAA-certified LynuxWorks LynxOS-178 RTOS. Any developers utilizing OpenArbor in C, Ada, and Embedded C++ (EC++) or any combination thereof can reach DO-178B Level A certification by using the RTOS to deploy their mission-critical applications.

Cell phone-like device protects soldiers from chemical threats

General Dynamics has chosen biological subsystems and sensor chips producer Sionex Corporation to supply Differential Mobility Spectrometry (DMS) technology for use in its JUNO "cell phone-style interface" chemical detector. JUNO is part of a contract between General Dynamics and the U.S. DoD under the Joint Chemical Agent Detector (JCAD) program's Increment 2. JUNO can simultaneously monitor levels of exposure to chemical warfare agents, nontraditional agents, and toxic industrial chemicals. It also lets users know when decontamination efforts are successful.

LETTER TO THE EDITOR

Dear Editor,

Thank you for your article on the continued health of the Ada language and its market for safety-critical systems. In the Sept/Oct 2007 issue of *Military Embedded Systems* and your editorial, *The "A's" have it,* you mentioned new support for Wind River's VxWorks 653 Platform. AdaCore has supported this platform since 2002 when GNAT Pro was selected for both the KC-767 and C-130 AMP projects. This is the GNAT Pro High-Integrity Edition for DO-178B, which has been used to successfully develop certified avionics systems on multiple aircraft, and which continues to be used in this safety-critical domain. The High-Integrity Edition for DO-178B contains a certifiable runtime system, a full ARINC-653 API, along with various support tools to aid in safety-critical certification efforts. We are delighted to see the market for Ada and safety-critical systems continuing to expand with new vendors entering the market.

Sincerely,
Greg Gicca
AdaCore
Director of Safety and Security Product Marketing

Do-it-yourself Linux – a money pit:

Commercial Linux saves time and money

By Jim Ready

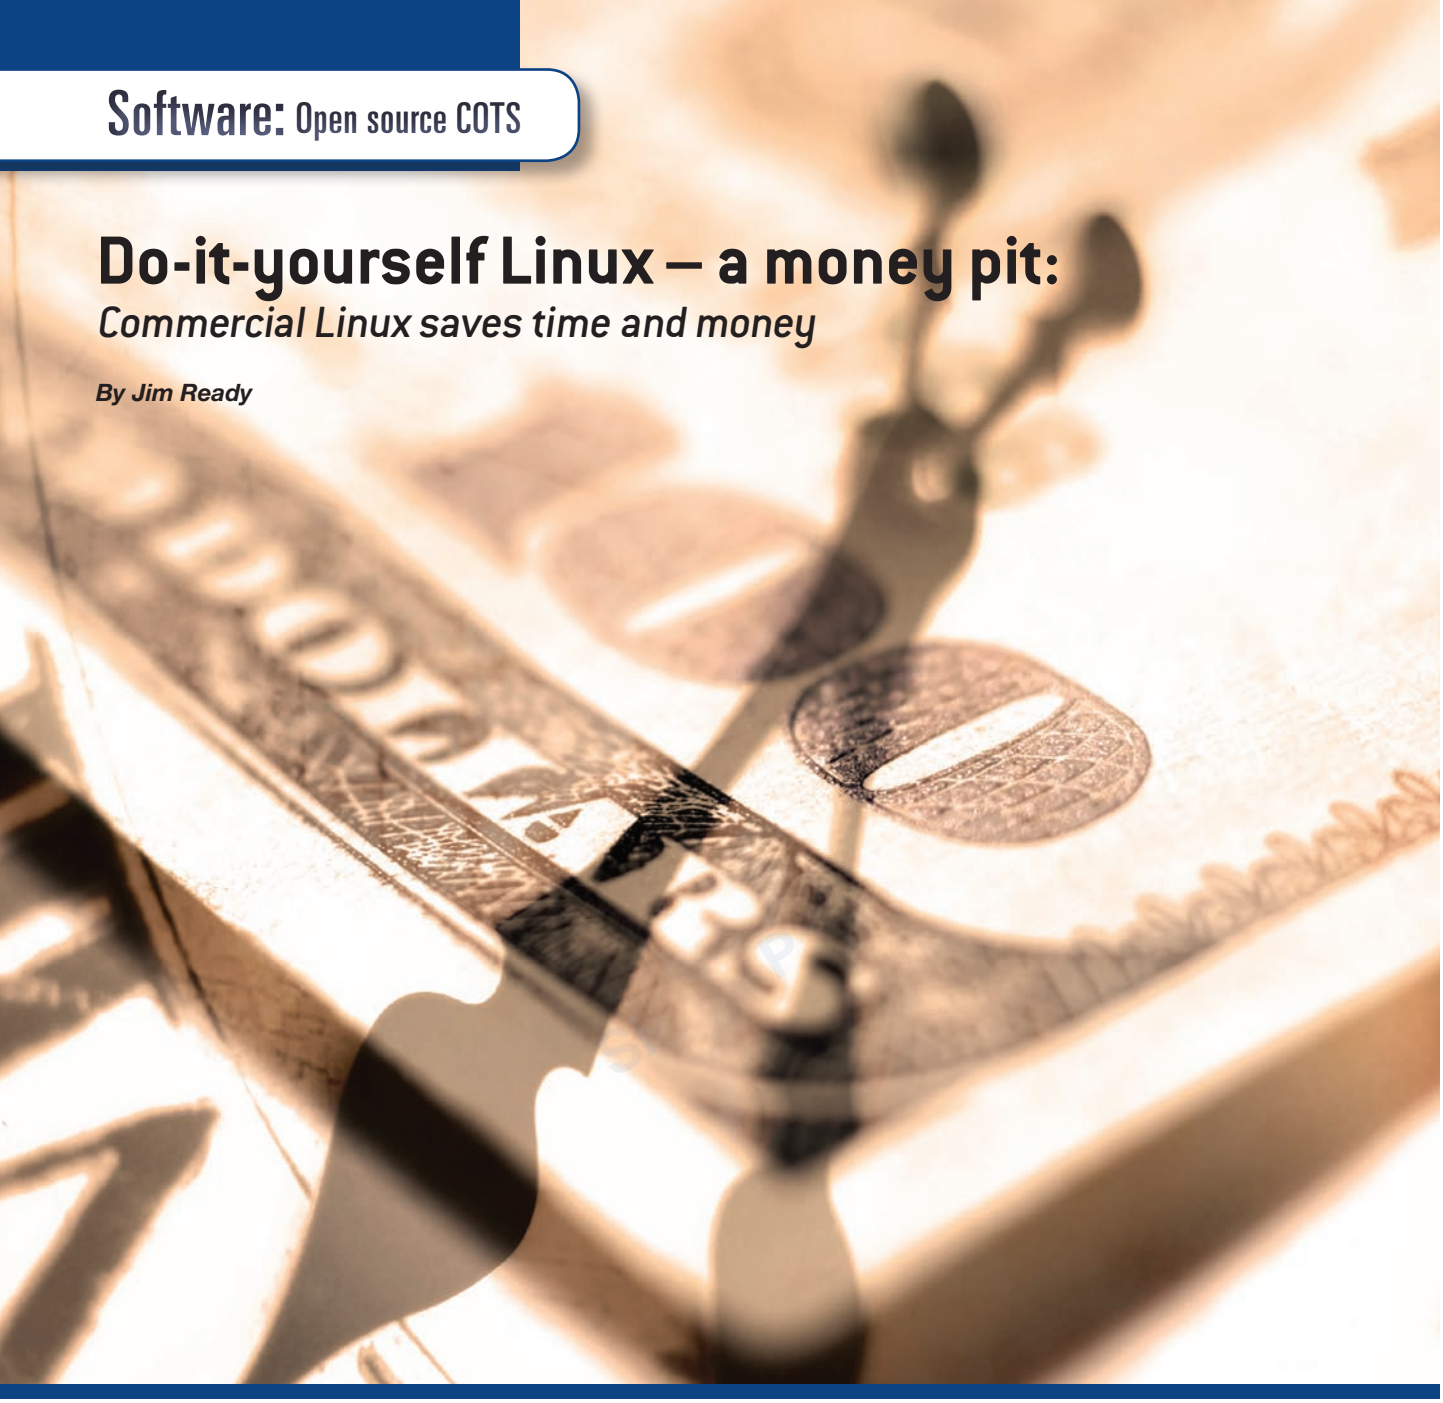

Companies that haven't clearly understood the high life-cycle costs of in-house RTOS development typically find themselves victims of the "software money pit." But in the buy-versus-make RTOS decision, commercial Linux clearly provides cost and complexity advantages versus the do-it-yourself Linux alternative.

For the past 25 years, engineering organizations developing software for embedded devices have become embroiled in the buy-versus-build decision for Real-Time Operating Systems (RTOSs). At the core of the debate is the emergence of commercial RTOSs, providing a complete system off-the-shelf and eliminating the need to write an RTOS system in-house. However, these COTS RTOSs carried a tidy price, and some entrants into the embedded market view them as more expensive than writing their own RTOS in-house.

However, most companies who choose to develop their own RTOS in-house have learned that there are many cost considerations that lie beneath the surface. Many organizations haven't fully understood the life-cycle costs encountered in developing an RTOS in-house, and most have learned the hard way about the development process costs involved. Thus, the in-house alternative often results in a "software money pit" where engineers start down the in-house development road, then pick up cost after cost until they realize it has become very expensive. By then, though, it is difficult to stop since so much has been invested.

Linux, however, is one RTOS that has established itself as the preferred Operating System (OS) for a wide variety of embedded devices. It too comes in two forms: Commercial and in-house. As time has proven, commercial embedded Linux has many advantages that developers should consider over building an in-house version of embedded Linux. An understanding of such considerations removes most of the debate as to whether or not to build in-house or purchase a commercial Linux package. Important factors developers need to consider include "obvious" and "hidden"

costs, along with the rising open-source complexity issues that affect the in-house versus COTS Linux equation.

The explosion of software content in devices

There are two dimensions of rapid expansion in the embedded system world that underlie this move to Linux over traditional RTOSs:

- The growth of new and complex product categories
- The growth of requirements for the system software in the embedded device itself

A case in point: A military manufacturer that previously dealt with small amounts of firmware needs to develop a whole new range of products based on a complete OS environment with multi-programming, TCP/IP connectivity, built-in security, and more. Consequently, this manufacturer suddenly finds its product teams forced to make decisions about the development and adoption of millions of lines of OS software into their product lines. With such an explosion in internal costs and in the complexity of product design and development, there is the risk of a software crisis throughout the technology product industries (See Figure 1). This is especially true in the embedded military market, where there is an increasing amount of scrutiny over cost, complexity, and time-to-market issues. Other considerations include:

1. The need for such system software cuts across a large number of different product organizations and teams. The external competitive situation doesn't tolerate a high cost structure for system software. Current budget constraints will not support the in-house teams formerly maintained; a company's engineering resources must focus on value-added technology to remain competitive.
2. The royalty component of current off-the-shelf system software in the cost of goods is large and impacts corporate margins.
3. The selection of a common, strategic system software platform will help a company avoid getting mired in a multiplicity of alternatives, each with a high cost structure and no overall leverage. Industry experience shows that disparate platforms inflate costs and increase product development cycles.

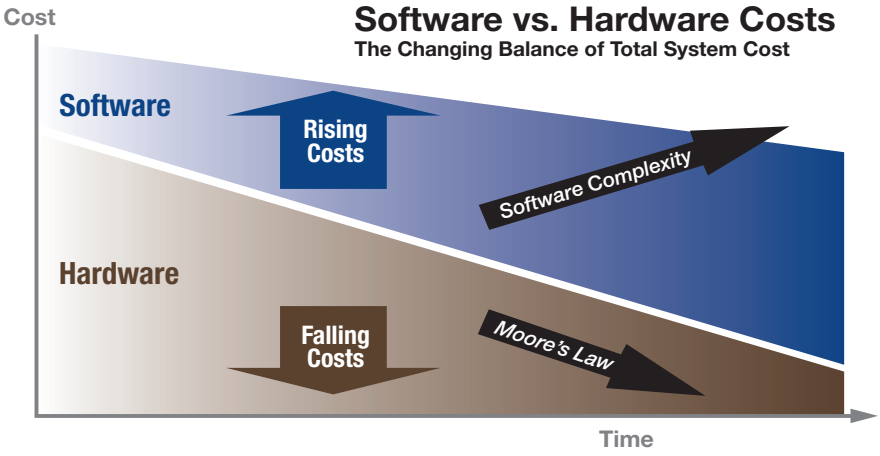

Figure 1

The growing popularity of embedded Linux as an alternative to traditional RTOSs is evident from the tens of millions of electronic devices that have already shipped with Linux as their OS. Ironically, however, even within organizations facing increasing pressure to produce software-intensive products with high functionality and quality requirements, as well as strong schedule pressure, the temptation to implement Do-It-Yourself (DIY) Linux remains alive and well, as opposed to off-the-shelf Linux, and development costs are a key factor – both obvious and hidden.

The "obvious" development costs

There are many processes and costs to understand in embedded Linux that are native to both in-house and commercial implementations; however, the difference is that these factors are typically included in the cost of commercial packages:

- It is not unusual for an embedded Linux or Linux distribution in general to comprise more than 30 million lines of source code.
- The code base can comprise 19 or more different, unsynchronized and non-integrated code repositories.
- The code base often changes daily.
- Many large companies use a wide range of embedded processor architectures requiring support for as many as 24 microprocessor architectures and their variants, and more than 100 hardware platforms.
- Multiple host computing environments and various revisions must be supported (such as Windows and Solaris, among others).
- The time to construct, test, and deliver an initial distribution can easily reach more than 30 development years (not including the ongoing maintenance, minor enhancements, or support time), which easily amounts to a cost in the seven-figure range. The ongoing maintenance, addition of minor enhancements, support, and the building of a development infrastructure also add significant ongoing cost.

Most commercial Linux distributions not only take these factors into account, but also have tools, resources, and expertise to help address them, making it an even more viable option for developers than downloading from kernel.org and building their own flavor. Considering that the typical cost of outfitting a mid-sized development team (~10 engineers) with a commercial distribution can be well under $100,000, the cost savings over a seven-figure in-house development effort reflect a real bargain.

The "hidden" development costs

In addition to the more "obvious" development activities, there are numerous "hidden" or often overlooked efforts that are essential to a successful development effort. Efforts involved in in-house development include:

- The development of a comprehensive Testing/QA capability. Test suites need to be developed for the OS itself and the wide range of I/O devices present, especially in SoC devices. (A typical SoC for a mobile device has 20-30 complex I/O devices built in – all these need to be tested.) The physical infrastructure to host all the reference boards for each architecture needs to be developed and deployed.

- The creation of a build computing environment that is efficient and capable of turning the build around as quickly as possible (hours not days); otherwise, the build process itself becomes a bottleneck in the overall project development flow.
- Tool development: cross debuggers, memory leak detection tools, performance-tuning tools, kernel-aware debuggers, and the list continues. Rarely are developmental dollars allocated to tool development, since most of the money had to be invested in developing the kernel itself.
- Training course material and course delivery. The Linux system is quite large and functional. Developers need training on its programming model, device driver structure, and development tools in order to be productive quickly.

These additional "hidden" costs can also affect delays, which can be quite costly. In fact, a December 2007 Embedded Market Forecasters report states that the average cost for each month of design delay is $386,250 assuming that the average cost per engineer is $150,000 per year and an industry average of 30.9 developers per project (www.embeddedforecast.com).

Process development, the "new complexity," and costs

The process development costs that Linux DIY projects incur are another cost consideration in the make-versus-build equation. It's clear that without the proper "Linux/open source-aware" processes in place, the development schedule and product quality will both suffer the potential for skyrocketing costs and/or outright failure.

Unlike a traditional in-house development process where the overall procedure could be strictly controlled from start to finish, in-house embedded Linux development begins with a largely uncontrolled process: open source. While an individual company might actually contribute to a particular open source project, by no means can that company totally control the process – it can only hope to influence and participate. For example, the rate of change of Linux from release to release can vary, from minor bug fixes and modest new feature enhancements to wholesale changes in underlying subsystems, introducing new capabilities as well as new instabilities and bugs.

The sheer volume of these changes can be daunting. Consider tracking the activities around an individual architecture for Linux, MIPS for example. To keep abreast of the changes occurring on a daily basis, a developer needs to monitor the email traffic of 11 different and unsynchronized open source projects: kernel.org, the core home of the Linux kernel; the gcc and glibc projects (the core tool chain and libraries from FSF at fsf.org); and at least nine other components that would typically comprise a useable Linux development environment.

Kernel.org itself may have up to 5,000 messages a day with 1,000 of these being patches that need to be evaluated and possibly applied to the source base. Simply ignoring the traffic, figuring that the system in use seems to be working well enough, can lead to disastrous consequences later. For example, a recent security patch that took all of 13 lines of code to implement against an embedded Linux system would have taken more than 800k lines of source patches to implement if the previous trail of patches had been ignored. It's a classic case of pay now or really pay later.

Adding it all up

All of the aforementioned factors combine to make developing a DIY embedded Linux distribution a significant investment (read "big bucks") in time and money. If there ever were a situation where the "software money pit" could really take hold, it's in owning 30 million lines of constantly changing source code. Even in the simplest case, the development costs are typically in the millions of dollars, with the cost for the more comprehensive implementations substantially higher (Figure 2).

Commercial Linux is the answer

Where do we stand today on this debate? On the one hand, one could argue that the build-versus-buy Linux debate will take some time to decide because on one hand, engineers can just download Linux and go. On the other hand, the industry has matured overall, and there is certainly a better recognition of the real costs of developing software.

Accordingly, complexity and development costs have grown over the past several years and are significant factors when selecting an embedded operating system. We can safely conclude that if it made economic sense years ago to buy instead of create, which has played out in the RTOS environment, it overwhelmingly makes sense now when the move is being made to a far more capable but far more complex embedded Linux-based environment.

Jim Ready *is a recognized authority in the embedded systems and real-time software industry. He founded MontaVista in 1999 to provide the Linux operating system to the embedded systems market and to offer embedded-systems expertise to the open source Linux community. He received his BA in Physiological Psychology from the University of Illinois at Urbana-Champaign and his MA in Physiological Psychology from the University of California, Berkeley. For more information, email info@mvista.com.*

MontaVista Software
408-572-8000
www.mvista.com

Develop-It-Yourself Linux
Large-Scale Deployment

Here's what it takes:

Corporate Wide Linux Deployment
- Dozens of projects
- Numerous hardware platforms
- Comprehensive feature set
- High quality requirements

Initial Investment - year 1	$7,278,750
Recurring Cost - year 2	$6,172,500
Recurring Cost - year 3	$6,172,500
Total	$19,623,750

	Initial Investment	Recurring Costs
Basic Distribution	47,600 hours	17,600 hours
Feature Enhancements	9,700 hours	9,700 hours
HW Enablement	25,000 hours	25,000 hours
Tools	6,000 hours	4,000 hours
Support	0 hours	8,000 hours
Maintenance	0 hours	18,000 hours
Documentation	1,750 hours	
Training	7,000 hours	
Total hours	**97,050 hours**	**82,300 hours**

Based on a three-year project – at $75/hr

Figure 2

Eclipse Europa maximizes embedded developers' productivity

By Doug Schaefer

IDEs such as Eclipse Europa make life easier for embedded developers. Europa introduces several productivity enhancements, including a new version of Eclipse CDT that speeds up the creation of C/C++ projects, make code structures easier to understand, and simplifies setup of hardware debugging sessions.

Before the advent of Integrated Development Environments (IDEs), most embedded developers debugged software by writing information to the console or by using command-line, source-level debuggers such as gdb. Either way, it was difficult and time consuming to debug complex data structures. With a well-designed IDE, on the other hand, a developer can view data structures easily, using a point-and-click user interface. Moreover, the IDE can present a single interface for most development tasks, eliminating the need to constantly switch between editors, compilers, debuggers, and other tools. The net result: greater productivity.

A well-designed IDE may look simple on the outside, but underneath, it's a complex application. Building one requires considerable financial resources, a suitable architecture, a skilled workforce, and a lot of time. In fact, most software vendors don't have the resources to support an IDE on multiple hardware platforms, and even those who concentrate on a single platform struggle to keep pace with evolving customer requirements.

The Eclipse open source project addresses these problems. Software vendors and code contributors from a variety of sectors work together to create the Eclipse frameworks and IDEs, which vendors then use as the basis for their commercial development environments. As a result of this approach, vendors can focus on creating innovative, value-added tools, rather than on maintaining their own homegrown IDEs. The Eclipse C/C++ Development Tools (CDT) subproject is a prime example of the Eclipse ecosystem in action. Created by several vendors working in concert, the CDT has quickly become the *de facto* standard for C and C++ development in the embedded industry. In fact, most RTOS and embedded tools vendors now base their IDEs on the CDT.

Until 2006, however, project teams at Eclipse generally released new versions of their components on independent schedules. This approach proved difficult for vendors that wanted to distribute products based on components from multiple projects. As a result, the Eclipse Foundation staff worked with the leaders of the most popular projects to coordinate project releases. The first simultaneous release of Eclipse projects, called *Callisto*, became available in June 2006; the second simultaneous release, called *Europa*, debuted in June 2007.

The Europa release includes the latest version of the CDT, version 4.0. (This is a milestone release for the CDT project, which celebrated its fifth anniversary this year.) When developing CDT 4.0, the CDT project team focused on streamlining the project creation experience for new users, making code structures more understandable, and simplifying the debugging process.

C/C++ project wizard simplifies creation

When users launch CDT 4.0, the first thing they see is an updated version of the New Project wizard (Figure 1). In previous releases, users had to select whether they were using C or C++; they also had to choose which CDT build system would build their projects. With CDT 4.0, the New Project wizard hides the details of the underlying build systems, simplifying the process of setting up and maintaining new projects.

In this wizard, the user selects the type of project (executable, shared library, static library, and so on) they are building as well as the toolchain that they will use to build it. Some project types also provide the option of using a project template to generate

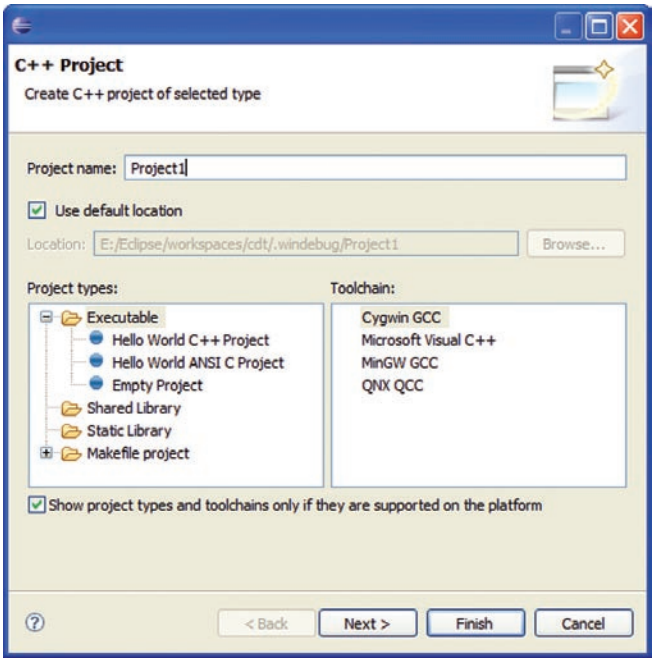

Figure 1

build settings and skeletal source code. Choosing this option automates the preliminary work otherwise required to set up a new project for a given problem domain. In the past, users would have to set up a makefile, understand and decide on compiler command-line options, decide on the project's structure, figure out include paths, and then create all of this before writing the first C file.

If developers want to use their own makefile, they can choose Makefile projects. This new option is especially useful for embedded developers who need full control of their build steps; for example, to control linker scripts or build downloadable images. Makefile projects place almost no limitation on what can be done within the makefile (arbitrary commands, for example). IDE-built projects, in comparison, allow for simpler build scenarios: change a file, recompile, or update binaries. Anyone familiar with CDT will find the Makefile projects option similar to the old Standard Make projects. However, the Makefiles option allows the developer to specify which toolchain they are using for the build to assist the CDT when creating its source code index.

The other project types, for creating executables and shared and static libraries, automate control of the build process by using CDT's internal build system. By automating the build process, the New Project wizard frees users from having to create their own makefiles and allows them to focus on writing code.

Toolchain vendors are free to extend the New Project wizard as well. They can provide integrations for their toolchains, thereby allowing their users to take advantage of CDT's build automation features. Vendors can also add project types and templates that customize the CDT for their specific environments. For instance, a vendor could create a project type that generates boilerplate code for device driver projects.

Navigation, editing features increase viewing capabilities

One of the most useful features of the CDT is its internal parsers and source code indexing component. This functionality drives a number of the CDT's source navigation and editor features. The

CDT editor now provides enhanced text coloring that highlights identifiers based on their role in the program. For example, it can use one color scheme to identify functions defined by the user and another color scheme to identify calls to functions from header files on the system include path.

A new call hierarchy view provides a tree view that shows function call relationships (Figure 2). Using this view, the user can see which functions call which other functions, without having to run and debug the application. The CDT also provides a similar view for viewing file inclusion relationships. It also provides an enhanced Type Hierarchy view to show C++ class hierarchies and the member variables and functions of those classes.

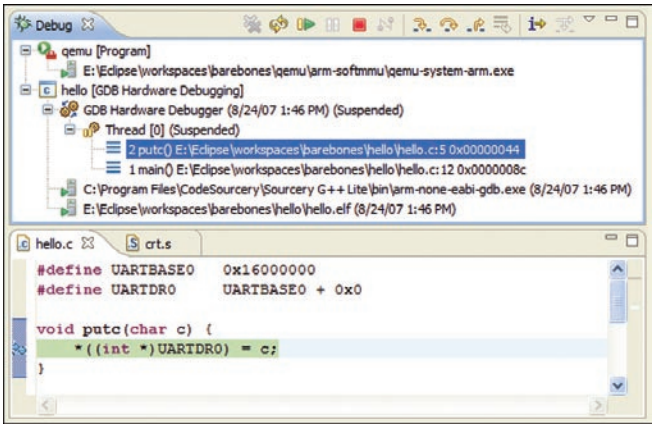

Figure 2

GDB hardware debugging

The CDT debugging component has always had an extensive integration with the GNU debugger, gdb. This debugger is very popular with embedded developers due to its free availability and support for a large variety of target architectures. JTAG/BDM vendors are also catching on to gdb's popularity. Thanks to gdb's built-in simple remote debugging protocol, it is possible to use gdb as a front-end for hardware debugging sessions. However, with previous releases of the CDT, it was very difficult to set up the CDT to use gdb with hardware debuggers.

To simplify the setup of hardware debugging sessions, CDT 4.0 introduces a debug launch configuration that lets the user enter a number of properties specific to using gdb for hardware debugging. The launch configuration includes the remote address of the hardware debugger interface, the location of the image file, if necessary, as well as commands to customize the startup of the gdb session to prepare the embedded target for execution. In addition to CDT, the Eclipse Device Software Development Platform (DSDP) provides even more functionality to aid embedded developers (see sidebar).

A complete IDE for embedded development

Europa and the CDT/DSDP projects are providing complete support for developing C and C++ code for embedded systems

Device software components

The CDT addresses the general needs of C/C++ development and provides support for the edit, build, debug cycle of embedded development. The Eclipse Device Software Development Platform (DSDP) provides additional features that focus on the specific needs of embedded developers. For Europa, the DSDP project has also released a number of components that plug into the Eclipse environment. These include the Remote System Explorer, which provides a framework and tools for accessing remote targets. This framework allows the developer to transfer files to the target and to perform remote execution and debugging of applications. DSDP also provides a new terminal view for command-line access to remote targets, using telnet or ssh over serial or network connections.

Europa also provides a preview of the DSDP's device debugging component. This component is designed to handle the specific requirements of some embedded debugging environments and will be used by a number of vendors in future Eclipse releases.

Meanwhile, Eclipse continues to provide a strong foundation for building IDEs. A large number of vendors have caught onto this already and are building their commercial IDEs based on this platform. The QNX Momentics IDE is an example of this. It adds the necessary compilers and debuggers to build applications for the QNX Neutrino operating system. QNX also created their own proprietary Eclipse components to provide support for system and application profiling.

and additional components for embedded development. Europa is also making system developers' lives easier via simplified project creation, code structuring, and hardware debugging.

Many vendors are sharing in the work of the open source Eclipse platform as they build their own customizations to create successful commercial products. It is this sharing and commercialization that creates a healthy Eclipse ecosystem and ensures the continued evolution of Eclipse technologies to further meet the productivity needs of embedded developers.

Doug Schaefer *is the project leader for the Eclipse CDT project and a tools architect for QNX Software Systems, where he is involved in all aspects of the QNX Momentics integrated development environment. He has been a tools developer for most of his career, having worked for ObjecTime, Rational, and IBM building modeling tools for embedded systems. He is a strong proponent of open source software development and is heavily involved in Eclipse, not only leading the CDT project, but also as a member of the project management committee that oversees all tools projects at Eclipse. He can be reached at DSchaefer@qnx.com.*

QNX Software Systems
613-591-0931
www.qnx.com

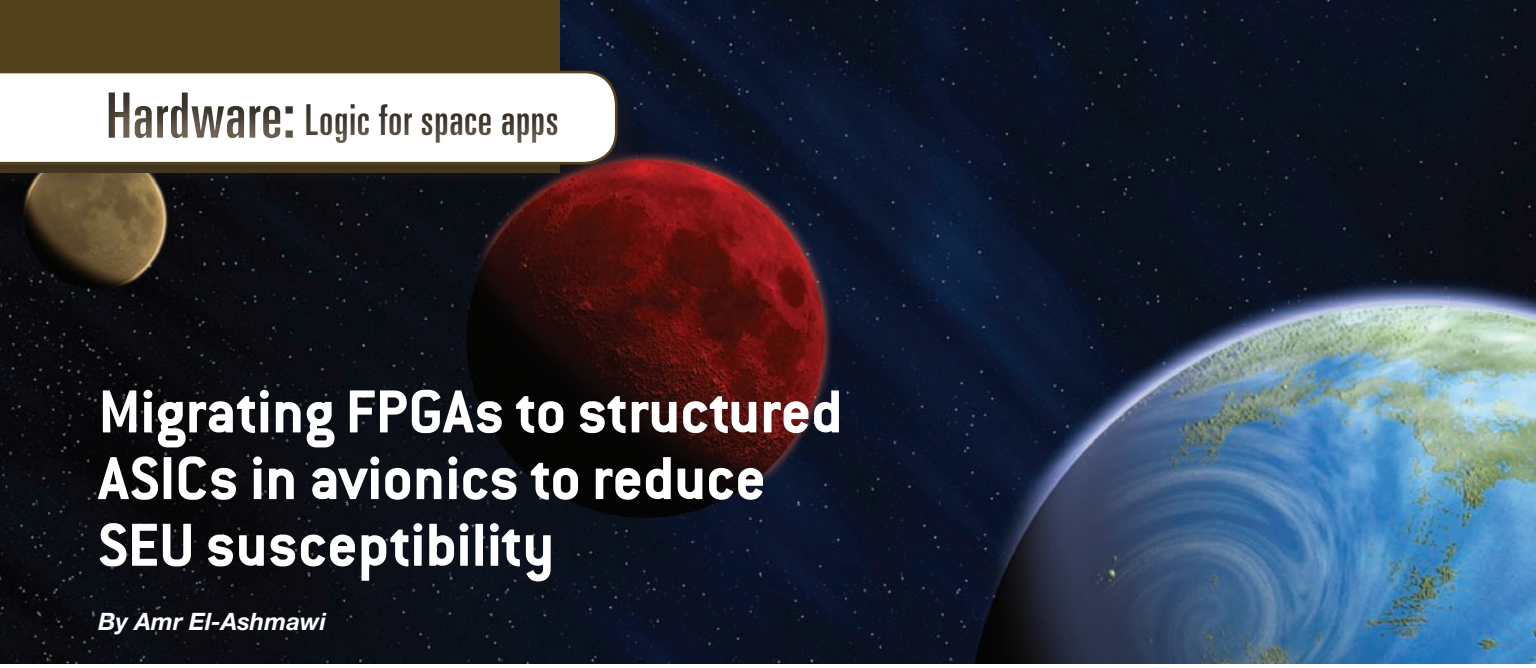

Migrating FPGAs to structured ASICs in avionics to reduce SEU susceptibility

By Amr El-Ashmawi

Today, FPGAs are being used more often in a broad spectrum of applications. With their inherent attributes, they offer avionics design engineers advantages not available in ASIC-based platforms. FPGAs are susceptible to particle-induced Single-Event Upsets (SEUs). Combining the use of FPGAs and a structured ASIC provides a clear migration path for designers to utilize.

The use of FPGAs in many applications is rapidly increasing. With the attributes of reconfigurability and design-to-working part times much faster than those of ASICs, FPGAs offer avionics engineers advantages not available in mask-programmable silicon platforms. However, SRAM-based FPGAs are susceptible to particle-induced SEUs, which make their deployment in avionics problematic. The solution to this dilemma lies in the use of both FPGAs and a structured ASIC for complete system development, with a clear migration path using a single tool flow and suite. Using an FPGA to validate an initial design and then migrating from the FPGA to a structured ASIC reduces cost, power, and system susceptibility to SEUs.

Atmospheric-induced SEUs

The atmosphere is the source of several types of ionizing subatomic particles that can interfere with the normal operation of electronic components. These particles are the result of solar rays and galactic cosmic rays that collide with the oxygen and nitrogen atoms in the Earth's atmosphere and produce high-energy protons and neutrons. Any of these particles have the potential to cause an SEU in an integrated circuit. NASA defines SEUs as "radiation-induced errors in microelectronic circuits caused when charged particles (usually from the radiation belts or from cosmic rays) lose energy by ionizing the medium through which they pass, leaving behind a wake of electron-hole pairs."[1] Figure 1 illustrates high-energy particles from space passing through silicon. They produce hole-electron pairs that can change the state of a memory cell or the configuration bit of an FPGA.

The density or flux of neutron particles peaks at several tens of thousands of feet above the ground, often at levels where avionic equipment operates. Other particles such as secondary protons are also present, but for SEU effects, neutrons have the highest probability of interfering with electronic system operation. The greatest density of these neutrons is at around 50,000 to 60,000 feet. Below that level, they are attenuated by the atmosphere, resulting in a much lower density at ground level. At 30,000 feet, a chip will have a Failure-In-Time (FIT) rate far higher than that at ground level.[2]

When a high-energy neutron impacts a transistor on a chip, potentially, the magnitude of the pulse may be sufficient to change the state of a logic node. This is called a *soft error*, since only data is affected, not the functionality of the circuitry. These errors are typically transient and are nondestructive when they occur in logic circuits. However, in a volatile

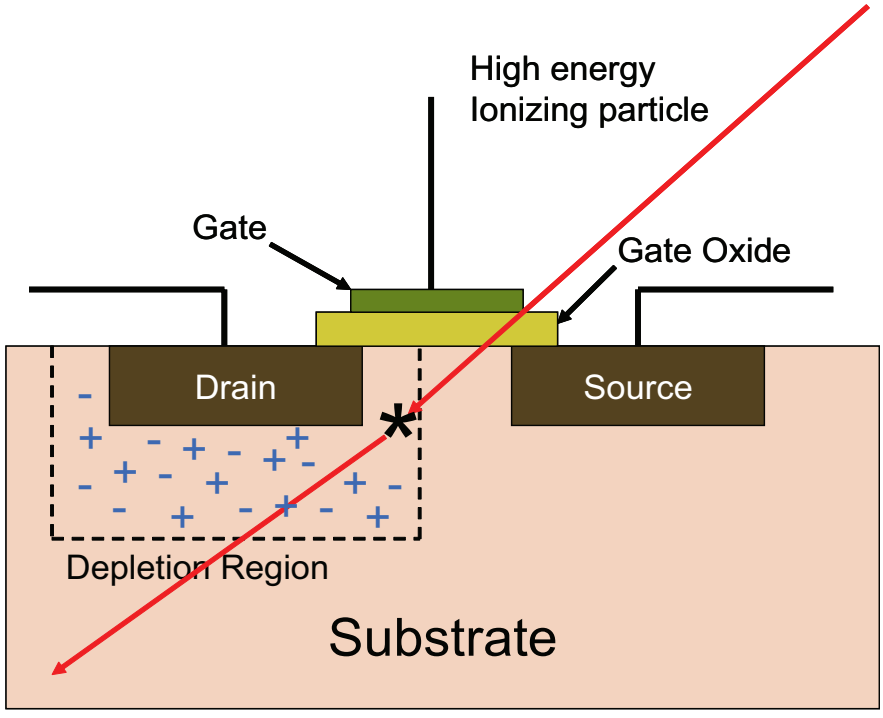

Figure 1

memory circuit or latch, such as an SRAM cell, the neutron impact may change the state of the memory, which can have consequences on the operation of the chip containing the affected memory cell.

For military and commercial mission-critical avionics applications in jet aircraft and other high-altitude vehicles, there are concerns about using SRAM-based FPGAs, because of the possibility of neutron-induced SEUs. In many cases, designers are forced to use ASICs or One-Time Programmable (OTP) devices that have density or performance limitations.

SRAM-based FPGA circuit

SRAM-based FPGAs have their functionality set by the bit pattern in an SRAM configuration memory. If an SEU upsets a configuration memory bit, the device's functionality can change, which may result in data corruption. The error, if not detected and corrected, can become a "hard" error. When an SEU-induced error is detected, the FPGA must be reconfigured, in some cases with the device's power turned off to clear the fault and then reapplied.

An FPGA's SRAM-based configuration memory, which controls the FPGA's logic and routing, is the part of the FPGA sensitive to SEUs. SEU-based configuration memory problems include modifying the chip's logic, disrupting interconnect paths, changing the global clocking, and modifying operation of the user I/O pins.

Configuration memory upsets can be detected and corrected using configuration feedback. This has led to different solutions at various cost levels, including making the SRAM cells more robust and introducing additional circuitry to support Error Correction Codes (ECCs), Triple Module Redundancy (TMR), and/or Cyclic Redundancy Code (CRC). Each of these has various pros and cons at the system level.

ECC support is provided by additional circuitry that stores the FPGA's configuration in additional bits in the memory word and calculates a checksum for the contents. This checksum is periodically compared to the checksum for the FPGA's configuration memory contents and, if a discrepancy is found, the FPGA is reconfigured. In a similar manner, extra circuitry can be used for CRC checking, to see if the configuration memory contents have changed. If they have, the FPGA is reconfigured, either partially or fully, to reproduce the correct configuration memory pattern. The additional circuitry for ECC and CRC costs chip complexity, area, and power and also produces latency between the Soft Error Rate (SER) detection and correction.

> "Using an FPGA to validate an initial design and then migrating from the FPGA to a structured ASIC reduces cost, power, and system susceptibility to SEUs."

TMR provides embedded "majority voting" circuits in flip-flops that result in high levels of tolerance to SEUs. If an error is detected the device must be reconfigured. TMR is hardwired into the design and results in significant additional silicon area, extra weight, additional power dissipation, and potentially reduced speed.

Migrating to a structured ASIC

Designers can customize SRAM-based FPGAs to application-specific algorithms for many high-altitude applications including flight control systems, displays, high-altitude electronic warfare systems, Unmanned Aerial Vehicles (UAVs), global positioning, seeker modules on missiles, and communications modules. FPGAs provide performance capabilities approaching those of ASICs and far beyond those offered by OTP devices such as UV EPROM, EEPROM, or antifuse-based devices. SRAM-based FPGA vendors are also hardening devices to be more SEU tolerant.

However, designers can further reduce their development risk and improve time to market by initiating their design-planning efforts using an SRAM-based FPGA, validating the design in-system and at-speed, and then migrating from the FPGA to a structured ASIC with lower power, SEU immunity, and a single chip.

In a structured ASIC, metal interconnects replace all the configuration RAM elements, dramatically reducing the

probability of SEU-induced errors and the need for error mitigation for the configuration SRAM. Military applications need an ITAR tool flow to assure IP security during overseas manufacturing. Altera is in the process of supporting such an effort for overseas manufacturing. All metal lines are predefined and then programmed for a particular logic configuration using vias between the lines. A structured ASIC offers faster development and lower NRE costs than a standard-cell-based ASIC, as well as significantly lower unit cost and power, and often higher performance, compared to a high-end FPGA, as illustrated in Figure 2. A structured ASIC also eliminates the programming circuitry associated with an FPGA.

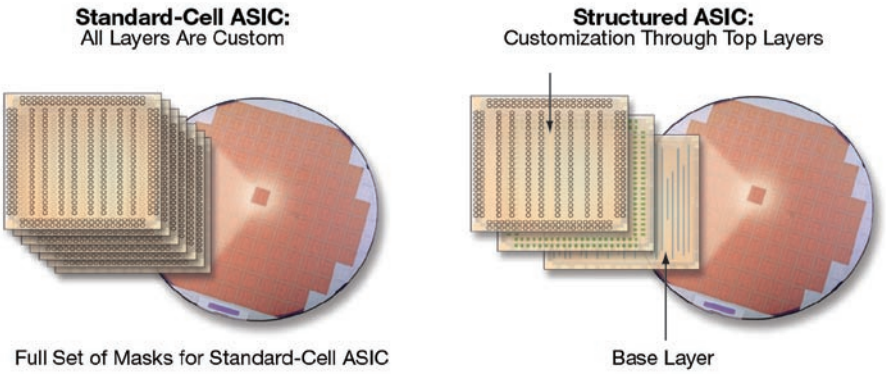

Figure 2

Registers are needed in any logic design and, like any nonvolatile element, they are susceptible to SEUs. Some older FPGAs have similar registers or latch SERs to what is found in an ordinary ASIC. Structured ASICs have architectural features to eliminate this problem and to improve overall soft error rates[3] over an FPGA:

- Increasing feedback loop gate strength
- Isolating the master and slave stages with an inverter
- Improving node capacitance via programming
- Not creating an asynchronous load MUX if it is not needed

The schematic in Figure 3 shows Altera's HardCopy II Structured ASICs implementation of logic features that are not available on an FPGA and how SEUs are reduced by adding logic to improve feedback loop gate strength.

These architectural features provide the highest level of DAL, Level A. Development Assurance Levels (DALs) are used within Aerospace Recommended Practices (ARPs) to classify system components based on their most severe failure condition associated with an aircraft-level function. DALs range from Level A – failure is catastrophic – to Level E, where failure has no safety effect. At a 130 nm process node and 1.5 V, the structured ASIC has a neutron SER cross-section (> 10 MeV, sea-level) of 7E-15. For 90 nm and 1.2 V, using hardened register design, the neutron SER cross-section (> 10 MeV, sea-level) drops below 1E-15.

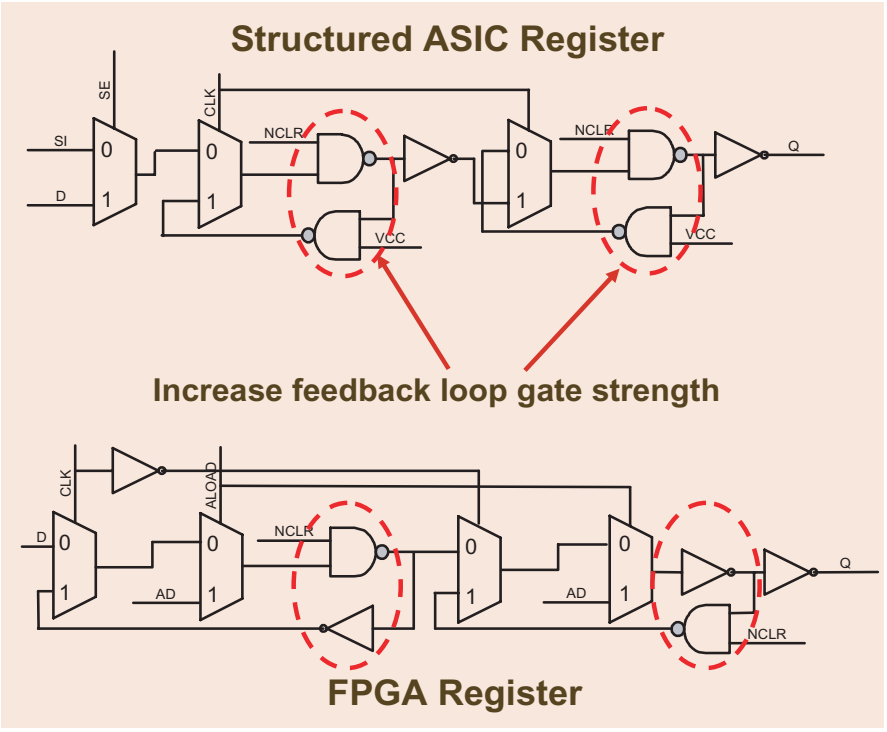

Figure 3

A well-designed structured ASIC also significantly reduces dynamic and static core power over an equivalent FPGA. For example, there are several ways to reduce static power:

- Removing the programming circuit and configuration logic
- Powering on only the logic and memory that the actual circuit uses
- Powering down the unused logic and memory with via connections

Unlike cell-based ASICs, structured ASICs can be quickly and easily manufactured with short fabrication cycles and low NRE charges for each design iteration.

Realizing design benefits using structured ASICs and FPGAs

Developing your system with a structured ASIC and an FPGA can provide identical functionality and performance during the finalization of the design algorithms, as observed in the lab with the FPGA-based prototype. Designers can be certain that they'll get a lower power, single-chip SEU-immune device without the need for an external storage device, the latter feature improving system reliability and reducing board space by eliminating a chip. A further benefit is that designers can switch back to an FPGA if they need to modify the design to accommodate a design change or target the chip for a different application.

References

1. NASA Aerospace Science and Technology Dictionary, http://www.hq nasa.gov/office/hqlibrary/aerospacedicti nary/aerodictall/s.html.

2. Graham, Paul e al, "Consequences and Categories of SRAM FPGA Configuration SEUs," Los Alamos National Laboratory, Paper LA-UR-03-6122.

3 Euzent, Bruce et al, "Structured ASIC for SEU Mitigation of SRAM-based FPGAs," presented at the 2006 MAPLD International Conference, Washington DC, Sept. 26-28, 2006.

Amr El-Ashmawi *is the senior marketing manager responsible for the military and aerospace business unit at Altera. He has more than 16 years of experience in semiconductor design, programmable logic, and working within the environment of the defense market. Prior to joining Altera, Amr was the director of marketing at SiliconExpert Technologies, where he was a founder with responsibilities in providing component management tools to the commercial and defense industries. He has also held marketing and engineering positions at Actel Corporation. He holds a BSEE and MSEM from Santa Clara University. To learn more, contact Amr at aelashma@altera.com.*

Altera Corporation
408-544-7000
www.altera.com

Deployable mass storage systems:

Bringing military applications to the front lines

By Laura Cooper

Today's military and intelligence applications increasingly require deployability. This includes everything from mobile communications to data surveillance. The system requirements for many of these applications are demanding: high-performance processing, multiple I/O options, and massive storage. Yet many systems on the market today often fall short of fulfilling all of these requirements, and may only be somewhat portable, limiting their application uses. Today's laptops max out at dual core processing technology, limited I/O options, and about 4 GB of storage. Deployability of higher-performing systems, those capable of collecting, storing, and transmitting large amounts of sensitive data, is a request that is on the rise in the military and intelligence communities.

Massive storage alone is not a useable component in a system that does not also incorporate those features an all-in-one mobile system should provide. Powerful processors and multiple I/O are needed to fully utilize the application capabilities of a system capable of storing seismic amounts of data.

A *deployable mass storage system* is one that is capable of housing loads of data, yet small enough that it is portable. For this feat to be possible, the system also must utilize advanced processing technology and I/O options such as the several PCI or PCIe slots. In effect, the system is a portable computer system that performs various tasks, ideally suited for gathering, storing, and disseminating large amounts of critical data out in the field, in a laboratory, or virtually anywhere. As technology advances, the development of systems offering massive storage in a deployable form factor, coupled with the performance and I/O needed to utilize it, mobilizes mission-critical applications that once could not be taken into the field.

Sensitive information collection, storage, and dissemination

Military intelligence is information needed to plan for our national defense. Knowledge of the number, location, and tactics of enemy forces and potential battle areas is needed to develop military plans. To acquire information, services and agencies rely on aerial photographs, electronic monitoring, data gathering, and human observation. Intelligence specialists collect and study the information required to design defense plans, strategy, and tactics.

The methods utilized to gather tactical information are at the forefront of technology, enabling data streaming at speedy rates to ensure information is as close to real time as possible. A mobile system that provides enough processing power, I/O options, and storage for the large amounts of data being gathered would speed up the collection process and foster truer analytical results. Most, if not all, data gathering applications of this type require more than 4 GB of storage, more than most laptops can provide. A portable system with high storage capacity, perhaps even up to 2 TB, would be ideal.

System requirements for information collection

There are three elements required in a system used for information gathering and transmittance, a function technically referred to as high-throughput data streaming.

The first is the availability of numerous high-speed hard drives. The number of drives is relative to the speed of data transfer, and the more drives there are, the faster the transfer rate is. Hard drive options in smaller form factors have emerged in the form of 2.5" drives, proven to be as robust and reliable as 3.5" drives. These smaller drives allow for the chassis to be smaller, and hence, more portable. Figure 1 illustrates potential hard drive configurations in a portable form factor.

Secondly, powering many high-speed drives requires the performance of top-of-the-line processor cores. AMD and Intel workstation class processors, including each company's respective implementations of quad core processing, would provide the boost needed to maximize throughput rates.

Finally, several high-performance PCI or PCIe slots enable application-specific cards to perform data transfer, sniffing, and collection. Some examples relevant to deployed, military/intelligence applications include MIL-STD-1553, NTDS, DSP cards, Fibre Channel, and other FPGA powered I/O cards. Ideally, a system would utilize open standards, easily capable of accommodating all of these types of cards and many others.

The ability to house large amounts of storage in a deployable computer system would enable a variety of uses for smaller systems. A system of this type could be used as a portable sniffer, enabling on-site troubleshooting for many intelligence applications where a data center or laboratory is not accessible. A portable data recorder, with lots of storage, would enable the

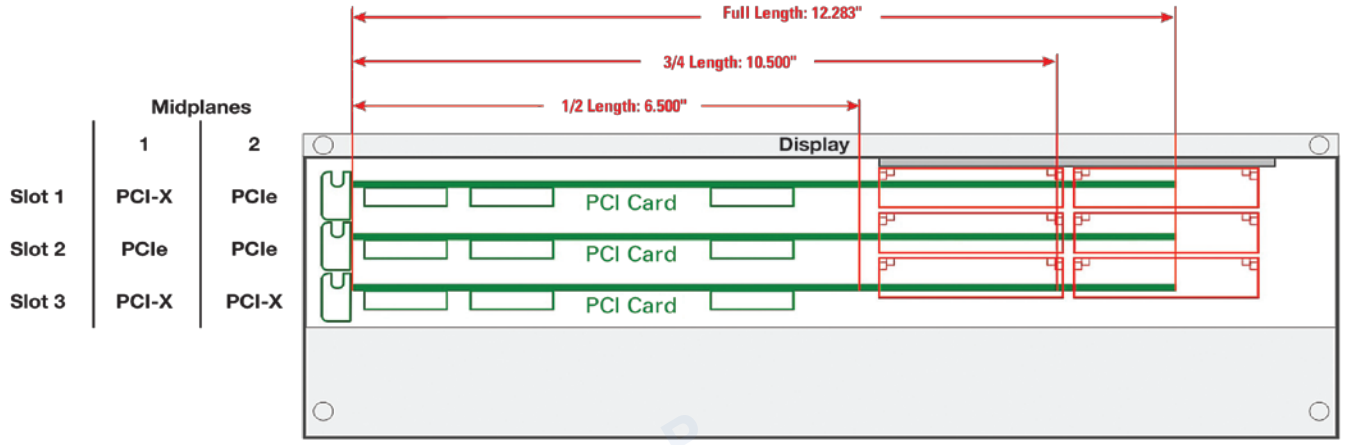

Figure 1

collection of real-time data of any type (video, audio, communication signals, or encrypted code) for on-site or later analysis.

Ideally, a system utilized for these applications would incorporate the latest in quad core processing, up to four PCI or PCIe slots, and allow for as many as 15-18 HDDs to be utilized. If capacities for each drive reach the current standard for a 2.5" drive, each would house up to 200 GB. In addition, a user interface with integrated screen and keyboard/mouse would foster ease of use. Some computer designers have begun to implement the technology in smaller chassis to allow maximum storage in a deployable form factor.

Mobile information collection and military apps

But which military intelligence agencies could utilize a deployable system with loads of storage, powerful processors, and multiple I/O? And what would they use it for? The choices are so vast that only a few are listed here.

SIGINT

SIGnals INTelligence (SIGINT) is used in missions to intercept and analyze foreign adversaries' communications signals, many of which are protected by codes and other complex countermeasures. They collect, process, and disseminate reports on foreign intelligence targets in response to requirements set at the highest levels of government. A deployable computer capable of high-speed data and signal collection and analysis would allow closer to real-time results from the outset. Sensitive information, such as that relayed by SIGINT, is more accurate the more quickly it is received. Collection of the data out in the field utilizing a deployable collection system would speed up the entire dissemination process.

C4ISR

Command, Control, Communication, Computer, Information, Surveillance, and Reconnaissance (C4ISR) missions have become increasingly deployable, utilizing portable technology on humvees and in tents, for example. The transmission of sensitive information in these applications would also benefit from a truly deployable mass storage system, rendering truer analysis results from communications brought as close to real time as possible.

TAMPS

The Tactical Aircraft Mission Planning System (TAMPS) is a computer-based method for weapons planning and optimizing mission routes against hostile targets. TAMPS is designed to provide a common automated system for rapidly processing large quantities of digitized terrain, threat, and environmental data, along with aircraft, avionics, and weapon systems parameters that assist in the precision engagement of enemy forces. TAMPS would benefit greatly from a deployable mass storage device to process and transmit various sorts of data. The system would ideally take up minimal space in the aircraft and be transportable from one spot to another in a scenario like the one shown in Figure 2.

Figure 2

Mobile TOCs

Portable Tactical Operation Centers (TOCs) are another area that would stand to benefit from a deployable system with massive storage of 4 GB or more, high performance, and multiple I/O options. TOCs are mobile, rapidly deployable, forward command facilities where commanders and their staffs can assess battle conditions and developments and execute a wide range of strategic and tactical options. These battlefield TOCs must be able to acquire and present an extensive array of real-time data and audio/visual information to commanders. Deployable technology, allowing them to disseminate information easily from center to center, would facilitate more accurate and up-to-date communications; thus, strategic planning would be optimized.

Available mass storage system options

Companies like NextComputing and Seagate Technology recognize the need for deployable computer systems incorporating massive storage, high-performance processing, and multiple I/O options in the military intelligence community and have developed the flextop computer, a briefcase-sized portable workstation first introduced in November 2005. It has increasingly become more capable, incorporating top-of-the-line processors from both AMD and Intel, multiple hard drive options, and maximum I/O options.

The NextDimension family of flextops is part of a unique breed of portable computers that achieves the combination of high-performance processing, mass storage, and multiple I/O functionality. With a form factor and durability making the line ideal for military applications, many military intelligence agencies have taken advantage of this newer technology. An optional integrated screen, external keyboard, and mouse make for ease of use wherever the system is set up.

Figure 3

Seagate Technology is responsible for the implementation of 2.5" hard drives, enabling up to (18) 200 GB drives to be incorporated into one NextDimension Evo model (Figure 3), with storage options reaching up to 3.75 TB. Many of the drives are removable, allowing quick and easy access for added flexibility, security, and ease of service.

The future of deployable mass storage systems

Expect to see an emergence of more deployable mass storage systems. As the military and intelligence communities continue to look for increasingly mobile ways to collect data, communicate, and plan, innovative technology companies are bound to take their cue and produce. Mass storage only recently showed up on the mobile technology forefront. More computer companies are bound to jump onboard in the coming months and years.

With only a couple of companies currently combining massive storage with the processing power and I/O needed to utilize it in a portable form factor, options are limited, though the products that have been brought to market were developed as a direct response to the requirements requested by many of today's mobile operations. The aforementioned options allow for a number of mission-critical applications to be mobile; more reliable data collection, faster communications, and enhanced analytic accuracy all serve to show consideration for the sensitive nature of our deployed military and intelligence projects.

***Laura Cooper** has been the marketing communications manager at NextComputing for more than three years. Previously, she gained marketing experience in the financial services industry at Merrill Lynch and SalesLink. She holds a BA in Sociology from Drew University. She can be reached at lcooper@nextcomputing.com.*

NextComputing
603-886-3874
www.nextcomputing.com

Advances in NV-SRAM technology provide robust memory in military apps

By Tim O'Connor and Grant Hulse

Improvements in NV-SRAM and the CMOS process have resulted in faster access times and higher densities for high-reliability military applications that previously struggled with the challenge of nonvolatile memory.

The fast SRAM-plus-nonvolatile memory known as *NV-SRAM* was first introduced to the military market nearly 20 years ago. Since that time, the technology has grown in popularity and has been used in a wide variety of military and avionics applications. It provides 25 ns SRAM read/write access times and immediate nonvolatile backup on any power disruption in a small package footprint, while meeting all military qualification and reliability requirements. Because each memory cell contains a full six-transistor SRAM bit, a two-transistor nonvolatile bit, and four transistors for the STORE/RECALL steering functions, the technology has only been offered at relatively low densities of 1 Mb, 256 Kb, 64 Kb, and 16 Kb (Figure 1).

The past several years, however, have seen solid NV-SRAM and CMOS process improvements, providing higher densities, faster access times, and lower price points – all without impacting the reliability, retention, and endurance specifications that the technology provides. These recent improvements breathe new life into a variety of military and high-temperature applications and help solve the military NV challenge.

The military NV dilemma

Virtually all electronic military and avionics systems require nonvolatile data storage. Much of this data must be captured, processed, and used repeatedly with new command and control information. During any power interruption or power loss, this critical information must not be lost. Restoration of power must also mean instantaneous recovery of all critical mission details at the moment before disruption occurred.

Nonvolatility has traditionally been incompatible with the fast read/write

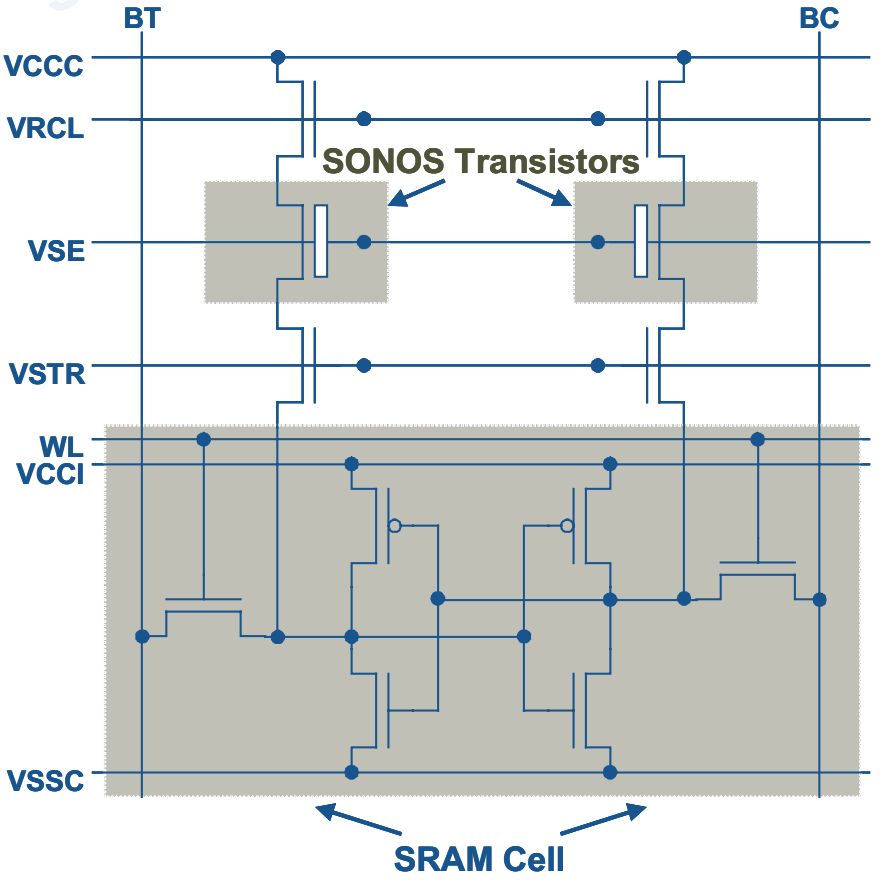

Figure 1

memory, rugged reliability, and functional integration requirements of modern military systems, requiring a trade-off that balances speed, power, density, performance, board space, reliability, and the price of available devices. The NV-SRAM was invented to provide precisely this solution, behaving as a fast SRAM until a power disruption is detected, then becoming a nonvolatile memory until power is restored. The automatic transfer of data from the SRAM to the nonvolatile memory takes place in only 15 ms, and the small amount of power is supplied by a small external capacitor.

In past applications, the NV-SRAM's smaller memory size has sometimes required military system designers to use slower, less reliable, and more complex devices like flash, EEPROM, SRAM with flash backup, battery-backed SRAM schemes, and so on. This has resulted in a growing demand from the design community to improve NV-SRAM's densities and applications reach.

Recent NV-SRAM CMOS process improvements

There are two competing CMOS structures used to provide a nonvolatile memory cell: Silicon-Oxide-Nitride-Oxide-Silicon (SONOS) and floating gate. NV-SRAM employs SONOS, the simpler of these two to manufacture, and also the most reliable with respect to endurance and retention for military applications. SONOS consists of a nonconductive nitride layer sandwiched between the substrate and a control gate.

When a high voltage is placed on the control gate, it excites electrons (or holes) from the substrate into (or out of) the nitride trap, depending on the voltage polarity. This trapped charge is the source of the stored nonvolatile information. In a loss-of-power crisis, this structure can retain the stored information reliably, to full military temperature specs, for more than 20 years. Furthermore, NV-SRAM uses the Fowler-Nordheim (F-N) charge injection technique, which causes minimal oxide damage during each STORE/ERASE event. Typically, one million STORE/ERASE events are possible using F-N injection without damaging the IC.

The floating gate memory structure is also a sandwich technology that uses a floating poly-gate layer and hot electron injection to rapidly push charges from the substrate into the charge trap. While this approach is used in most NAND and NOR flash technologies today, SONOS is now seeing growing interest in those markets as well.

Recent improvements in the SONOS technology have allowed NV-SRAMs to be manufactured in 130 nm CMOS. While this is several technology generations behind the state of the art, it nevertheless allows significant performance improvements over current competing products without compromising the military's high-reliability requirements. This process reduction has allowed memory densities to increase dramatically.

Furthermore, the Oxide-Nitride-Oxide (ONO) sandwich can now be added as a module to standard CMOS process flows without adversely impacting the performance of the CMOS gates used to construct other critical components on the same IC. This enables broad Intellectual Property (IP) block reuse and the ability to add SONOS-based nonvolatile memory blocks to virtually any CMOS design; floating gate technology cannot make such a claim. Program and erase voltages can be scaled down with much greater ease and precision using SONOS than with the floating gate approach, because both the nitride and the tunneling

oxide can be made progressively thinner, allowing the physical height of the structure to be reduced (Figure 2).

The ruggedness of SONOS has also been well documented. Rad-hard studies have rated SONOS technology in excess of 300 Krads, resulting in 10- to 100-year memory retention, with negligible change to the memory cell structure after 10,000 write cycles.

NV-SRAM advances in density and speed

Today's high-speed military microprocessor and DSP-based systems need very fast SRAMs. Current NV-SRAMs deliver access times of 25 ns across the military temperature range. The new 130 nm designs are expected to increase this to 15 ns in a 4 Mb format, which is a 40 percent speed improvement. On the planning table are densities of 8 Mb and 16 Mb. Faster access speeds are also planned in a hi-rel format at the smaller 1 Meg, 256 K, and 64 K densities.

The 130 nm NV-SRAM architecture will use the same proven one-to-one pairing of an NV bit and a fast SRAM bit in each memory cell (Figures 1 and 3). Faster speed comes from improvement in the SRAM portion of the cell, as well as in the address and data bus interface circuitry, and density improvements come largely from geometry shrink.

NV-SRAM application examples

NV-SRAM is finding its way into many high-rel military applications, including bomb, extreme-environment robotics, and avionics and networked ground vehicle "black boxes."

NV-SRAM provides bomb data

NV-SRAM is used both in bomb development programs to datalog a test bomb's behavior from launch to detonation, as well as to provide information in a live smart bomb. A good example of this is the bunker buster bomb (Figure 4). These bombs penetrate through hardened physical structures when dropped from an elevation of 40,000 feet. Falling from that height ensures that the 4,650-pound bomb can penetrate up to 100 feet of earth or 20 feet of reinforced concrete.

The nose of the bomb includes a highly complex detonation system that contains a microcontroller and a G-force sensor. The G-force sensor is used to detect how many floors the bomb has penetrated, or how deep into the earth it has traveled in order to determine when to explode the ordnance. Bunker buster bombs also have a guidance system directed by GPS. Here, mission-critical information received in real time from satellites must be saved in the NV-SRAM memory in the event of any disruption to the power supply.

NV-SRAMs help robots survive

Military and extreme-environment robotics is another NV-SRAM application area. Autonomous and remote control

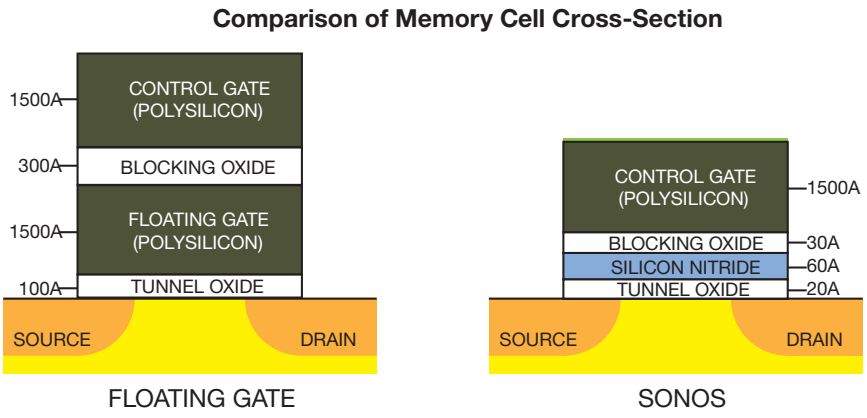

1. The stack height of floating gate can be twice the stacking of SONOS, primarily due to the difference in the thickness of the storage layer.

Figure 2

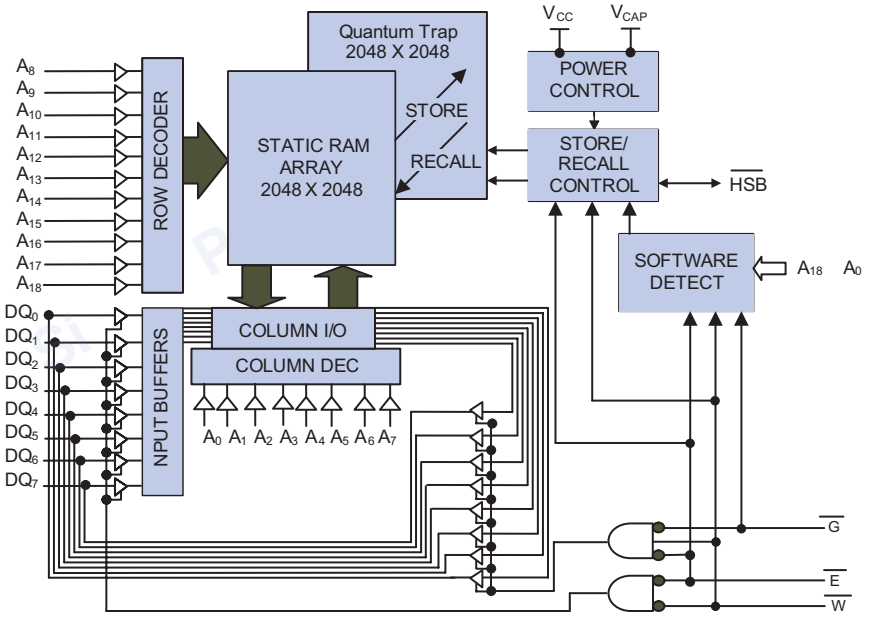

Figure 3

Figure 4

systems use NV-SRAM to add redundant backup of critical data to battery systems, including position of each axis, sensor readings, past sequence of completed activity algorithms, and so forth.

One recent example of an extreme-environment application is an oil field drill head. This includes a diamond drill bit backed by electronics that use an NV-SRAM to store physical readings such as temperature, depth, speed of the drill bit, and resistance. While friction created by drilling at great depths into extremely hard solids creates temperatures of up to 200 ˚C, NV-SRAM's tolerance to extreme heat makes it a viable solution.

Black boxes retain memory with NV-SRAM

Avionics modules in jet fighters are much like the "black boxes" used by commercial airlines. Storing both voice and flight data is an extremely important aspect of understanding aircraft malfunctions, as well as performance aspects of flight training. In cases of emergencies, the likelihood of power outages within the aircraft is extremely high. Saving data that accurately describes the flight characteristics of the aircraft is required to pinpoint the source of the problem. In a worst-case scenario of a fatal crash, the NV-SRAM helps identify which avionics component and/or systems caused the accident.

Networked ground vehicle and networked soldier applications will have similar black box applications for NV-SRAM. Additionally, any manual battery replacement sequence or momentary disconnection of the battery due to vibration, sudden impact, short circuit, or power disruption resulting from equipment damage can be better managed by using NV-SRAMs.

NV-SRAM continues to provide system gains

The increasing complexity of today's military applications will push system memory capacity requirements ever higher. Current densities of NV-SRAM will be a key part of this memory hierarchy and, as densities grow, will provide an increasing share of the total memory solution over time, while continuing to improve system reliability and performance.

Simtek will be among the first vendors to supply new 130 nm based products for high-rel military applications in 2008, including commercial-grade 4 Mb densities (512 K x 8 and 256 K x 16) in plastic packages and extended temperature versions of 256 K, 1 Mb, and 4 Mb under its nvSRAM brand name.

***Grant Hulse** is vice president worldwide marketing at Simtek Corp. Prior to Simtek, he worked as an RF product line manager at Qualcomm, director of product marketing at RF Magic, held executive roles at AMI Semiconductor, and worked in marketing at NEC Electronics. He holds a BSEE from Brigham Young University and an MSEE from Stanford University. He can be contacted at hulse@simtek.com.*

***Timothy O'Connor** is an associate applications engineer at Simtek Corp. and is currently completing his BSEE at University of California, San Diego. He can be reached at oconnor@simtek.com.*

Simtek Corporation
719-531-9444
www.simtek.com

Top 10 technologies sure to affect our war fighters

A random list of our favorite top 10

By Chris A. Ciufo

Courtesy of Pericles of Athens

For millennia, superior technology has been an integral part of battlefield victory. From iron and black powder for the ancient Chinese (see figure) to unmanned aerial vehicles first deployed by the Israelis decades ago, technology can be a game-changer in warfare. In no particular order, here's a look at 10 technologies we believe are – or will be – essential on tomorrow's battlefields.

1. Solid-State Disks (SSDs)

Storing data without moving parts evolved from core memory in the 1940s at Harvard and MIT to bubble memory from companies like Texas Instruments (TI) to UV erasable EPROMs, Electrically Erasable PROMs (EEPROMs), and finally to today's contemporary flash memory. As vendors migrated from 1 bit/cell Single Level Cell (SLC) NAND flash to today's super high-volume 2 bits/cell Multilevel Cell (MLC) flash, the price per megabyte has dropped precipitously to a mere 10x more expensive than rotating hard drives. This is a sweet deal in military systems because SSDs have no moving parts, are mostly impervious to shock and vibration, and have superior Size, Weight, and Power (SWaP) advantages at densities below 20-30 GB

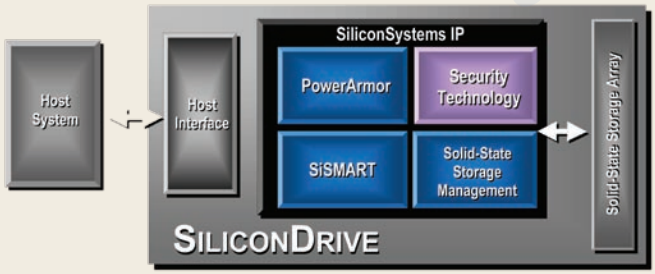

Figure 1

Wear-leveling algorithms and microcontroller-enabled SSDs such as those from SiliconSystems (Figure 1) can provide near-instantaneous secure data scrubbing, achieve 100,000 endurance cycles, and provide 10-year data retention. Although consumer SSDs such as MicroSD cards, CompactFlash, and even Microsoft's Windows ReadyBoost-equipped USB flash drives can't achieve this kind of reliability, the onslaught of storage in consumer gadgets keeps spurring flash technology forward. Embedded defense systems from soldier radios to SSD-equipped laptops are the obvious beneficiaries.

2. Multicore CPUs

Are two or four monolithic cores better than one? Darned right they are. Even if low power is the primary concern, spreading complex tasks across multiple microprocessor or peripheral cores may get the job done in less time or fewer cycles such as in the Intel Core 2 Duo shown in Figure 2. Since total energy consumed is directly proportional to a processor's power, number of clock cycles to accomplish a task, and the clock frequency, the fewer cycles required, the better. Multicore processor architectures also eliminate the redundant peripheral controllers and DC-DC converters that would otherwise feed separate discrete CPUs, thus saving additional power.

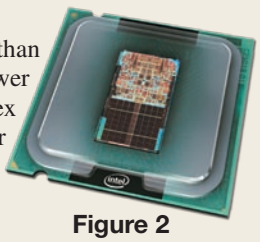

Figure 2

But power is often a secondary concern over performance in military systems such as Synthetic Aperture Radars (SARs), sonar, or Software-Defined Radios (SDRs) dominated by complex algorithms and data manipulation. Performing more on-chip calculations per unit time – as opposed to moving data between CPUs or memory – can pay big dividends. Intel's Quad-Core Xeon and Core 2 Quad CPUs or AMD's Quad-Core Opteron or Phenom follow on the heels of last year's successful dual-core CPUs. Based upon these vendors' road maps, don't expect the party to last much beyond eight cores, where physics and memory bottlenecks start to negate the multicore benefits.

3. CPU and system virtualization

On the heels of multicore architectures come virtual machines. Third-party vendors such as VMware or TenAsys have long provided multiple "desktop environments" in Windows or Real-Time Operating Systems (RTOSs). Intel processors and peripherals as well as BIOS and driver vendors are now exploiting hardware virtualization features to optimize hypervisors and system architectures by eliminating data redundancies between virtual environments. For instance, if data from memory on core 1 is destined for an Ethernet controller shared between core 1 and core 2, pointers enabled by drivers can allow the controller direct access to the core 1 data. This saves cycles and power and is highly efficient.

But in military systems, where ARINC 653 provides a guideline for saving LRU real estate by partitioning the software environment instead of splitting tasks onto separate hardware, these new virtual machines may be either a Godsend or a curse. On the one hand, hardware-based *Intel VT* (Virtualization Technology) definitely improves software-only virtualization commands. But some security experts have expressed concern that rogue processes may "leak" from one partition to another due to processor VT enhancements. It remains to be seen if virtualization done in hardware is as secure as, say, a MILS-certified OS separation kernel with partitioned memory under an MMU's control.

3. Google – the great enabler

Love it or hate it, monopolist or market catalyst, the effect that Google has had on the world is hugely profound. Google easily spiders into poorly secured DoD, program office, military contractor,

LOCATING...

Advanced microelectronics solutions partner *with*:

- MIL-PRF-38534/38535 **Class H & K certified facility**
- Semiconductor products for specific **circuit board density, weight** and **environmental** requirements
- **Enhanced** and **ruggedized** active matrix liquid crystal displays (AMLCDs)
- **RAD tolerant & RAD hard** capabilities

TARGET LOCATED

White Electronic Designs has expertise and experience in constructing extended environment, high-reliability, high-density semiconductor products and enhanced AMLCDs that deliver on time – over time – every time.

Proceed directly to:
www.whiteedc.com/ClassK

Alternate routes include **info@whiteedc.com** or **602.437.1520.**

WHITE ELECTRONIC DESIGNS
WWW.WHITEEDC.COM

or government websites and intranets, sometimes revealing documents questionably intended for public display. Try searching for Joint Publication 3-06 *Doctrine for Joint Urban Operations*, 16 September 2002[1]. One can argue that in a democracy, unfettered access is a good thing. On the other hand, Google probably also keeps DoD security managers on their toes – also a good thing.

Google is also driving Microsoft and cell phone providers insane. Google's Internet-accessible online documents can replace desktop Word or Excel, possibly saving the DoD millions in licensing fees. Google's plans to ostensibly build their own wireless Internet – bypassing ISPs such as Comcast or Verizon – are driving Wi-Fi and cellular data access technology rapidly down the price curve. Since the military uses radios galore, who's to say that new encryption technology applied to ubiquitous civilian LANs, MANs, and WANs isn't a better deal for the taxpayer than some fancy one-off DARPA mesh network? There's even some chatter that the DoD is considering an appliqué waveform to run on top of wideband CDMA.

4. Software-Defined Radios (SDRs) – Still Don't Receive

Don't get us wrong – SDR technology works, just ask Qualcomm or Nokia, or even Boeing or Harris. These companies all build software-reconfigurable radios (the former build cellular handsets, the latter build JTRS or military radios). But what we want are *cognitive radios* and *Multiple Input/Multiple Output (MIMO)* systems. Merely plugging the radio into a docking stand and re-flashing ROM to convert from, say, LINK 16 to WNW is so Year 2000.

Instead, we need radios that can "sense" the signal and network environment and adapt accordingly. Says Lee Pucker, CEO of the Software Defined Radio Forum and president of ForwardLink Consulting, radios "must efficiently utilize the available spectrum, manage their interaction with the spectrum, including adjusting their data rates and frequencies, and immediately interact with civilian and coalition networks if so programmed." For this, we need cognitive radios. As well, nifty MIMO algorithms like those from DARPA contractor Silvus Communications will optimize for ultra-high bandwidth data rates, low probability of detection, low power, or small physical size by balancing time, frequency, *and space*.

5, 6. FPGA devices and development tools

Let's get this straight. In the past 12 months, Actel, Altera, Lattice, QuickLogic, and Xilinx have all publicly (and almost forcefully) recommitted themselves to the defense market. In each case, the company is offering extended temperature parts, longer life cycles, or a commitment to obsolescence notification. You know, sort of like MIL-STD-883 and MIL-M-38510 all over again. Hello, 1985. Actel and Xilinx are both touting rad-hard tolerance; Lattice is the most likely to move up to the big table with Altera and Xilinx in the next few years; and Altera has this nifty U.S. ITAR facility that might be important to some programs of a sensitive nature.

As well, we're convinced that FPGAs will wipe out the general-purpose DSP market, relegating Analog Devices, Freescale, and TI GP DSPs to the IC graveyard. Instead, those DSP vendors are concentrating on special-purpose DSP-enabled SoCs. But beyond the wealth of signal processing IP available for FPGAs, the tool sets are finally improving. We're close to a C-to-HDL translator, and vendors are desperately trying to abstract the silicon so designers can code algorithms and functions, not individual transistors. As well, soft cores like Xilinx's MicroBlaze finally have an MMU, paving the way for a secure RTOS to run natively on-chip.

As we went to press, Xilinx had recently championed the Accelerated Computing Platform (ACP) M1 that bolts up to Intel processors' 1066 MHz front side bus. Nallatech is demonstrating their *Slipstream Module* that utilizes the Intel/Xilinx IP.

7. Simulation – cheaper training

So cruising around in a $22 million F/A-22 Raptor is an inefficient way to do basic flight and weapons systems training. While nothing is a substitute for the real aircraft, T-38s and T-6s can be equipped with synthetic systems to mimic those on other airplanes. For aircraft, beyond the multi-axis gimbaled full mock-up simulators housed in air-conditioned training facilities, in-airframe embedded systems increasingly provide more realistic exercises. The hugely popular video game America's Army proved a real catalyst for these types of realistic training systems, now spread across all of the Services.

Since our war fighters tend to be youthful and Internet-savvy, they expect the same capabilities and realism in their training simulators. Multidimensional training and simulation merges live, virtual, and constructive training vectors to provide the best of all worlds. For instance, the company Presagis, a wholly owned subsidiary of CAE, was recently stood up by merging Terrex, MultiGen-Paradigm, and Engenuity Technologies.

With their products, software and hardware systems can create a *virtual* world with objects moving, create an animated *constructive* world, and marry them in a *visualization* environment with simulated instruments and sensors (Figure 3). One can even create deployable code. We expect to see more battlefield assets interoperable in a training scenario (not unlike a LAN party or

Figure 3

[1] We first found it not long after 9/11/2001, and are shocked to still find it available at www.dtic.mil/doctrine/jpoperationsseriespubs.htm.

What's on your RADAR?

Thomson-CSF

ASTERIX

TPS-43

Marconi 10-Bit

TADIL-B

Raduga-2

NEXRAD

HDLC

X.25

Synchronous BitStream Interface

CD-2

TPS-75

Performance Technologies' NexusWare® WAN protocol software provides a seamless connection to a broad range of telecommunications, aerospace and defense, and commercial applications for both emerging and legacy networks.

With over 25 years experience developing highly reliable and scalable solutions, we can provide you with integrated hardware and software that can meet your product needs today and allow for future migration tomorrow.

Managed WAN Gateways, Radar Extraction, Recording Servers, and Lawful Intercept Platforms are a few of the many applications that utilize our broad software offering. Available on a suite of hardware platforms in PCI, PCI Express, CompactPCI®, PICMG® 2.16, and stand-alone form factors, our flexible WAN protocol solutions can easily accommodate a wide range of data and communication network requirements.

Explore the limitless possibilities for your programs by calling Performance Technologies today.

Platforms. Blades. Software. Support.
One company – many solutions.

PERFORMANCE TECHNOLOGIES

Phone: 585.256.0200 ◆ www.pt.com/rad ◆ E-mail: info-request@pt.com

Internet gaming via an Xbox 360), as well as the wider use of embedded training where actual military hardware contains a "training mode" that simulates threats or creates on-the-fly changing mission profiles.

8. Secure software – it's everywhere

We've talked about this technology extensively in the pages of *Military Embedded Systems*. From ARINC 653 partitions to DO-178B certifiable FAA code used in the military to MILS and Common Criteria, defense systems and associated software must be safe and secure. The biggest trend we see includes RTOS vendors revamping their code into a *separation kernel* upon which the entire OS is built. Products from Green Hills, LynuxWorks, Wind River, and others are all moving in this direction with secure partitions above the kernel into which *any* OS (including Windows or Linux) can be safely run.

The wild card here is what's happening in the commercial world as companies like Symbian worry about hackers cracking European cellular handsets, or as Cicso/Linksys hardens their consumer routers that increasingly carry sensitive personal or banking data. It's a fair bet that the DoD approach to securing software environments is way overkill, albeit repeatable and reliable. We expect the COTS world to advance further and faster because the market opportunities are just *huge*. From there, maybe the DoD will again adopt whatever the COTS market develops.

9, 10. Open standards, open source

Rounding out our list are the latest in open standard form factors. From the very beginning of the Non-Developmental Item (NDI) military movement 20 years ago, modules such as VME and STDbus cracked sole-source defense systems wide open. VME's latest mainstream technology uses switched serial fabrics in a crossbar architecture called VITA 41 VXS. Already there are dozens of vendor choices and military programs giving VXS the nod. From rival organization PICMG comes MicroTCA – or rather, rugged MicroTCA – a 3U-sized module and quasi-rugged chassis based upon the AMC standard. MicroTCA looks extremely compelling because it may meet the needs of all but the harshest military systems at lower price points.

Linux, the poster boy for the open source community, continues to gain ground in military systems. Sometimes Linux migrates into a program from the desktop-based advanced development phase, sometimes it's already embedded in a piece of COTS hardware, and just as often it's specified by the chief engineer so a prime can maintain revision control over the application code. Other open source technologies popular in defense include Eclipse for development environments, OpenEmbedded (based upon the OpenZaurus project) for targeting Linux on consumer doodads, and even an open source Joint Tactical Radio System Software Communications Architecture (JTRS SCA) from Virginia Tech called *OSSIE*.

Considerations in designing multicore systems

By Arun Subbarao

Multicore processing is having a tremendous impact on the development of embedded systems in the military and elsewhere. While previous multiprocessing solutions involved two or more physical chips, which doubled or more the amount of board space consumed, the introduction of multiple processing cores in a single chip allows operating systems and applications to leverage increased computing power. It has also provided access to additional computing resources without noticeably increasing the size or weight of the system.

But as developers delve deeper into the nuts and bolts of designing multicore systems and the applications that run on them, they're finding that these advantages come with certain obstacles and additional considerations that must be taken into account during the design process. The impact on software is most immediately felt in designing the OS, which must support either Asymmetric Multiprocessing (AMP) or, ideally, the more advanced Symmetric Multiprocessing (SMP) architecture. Concurrency and processor affinity issues also need to be addressed. For applications, there are considerations such as parallel execution and resource allocation.

Finding an OS to balance workloads

In multicore computing, the functions performed by the operating system become layered and more complex. The operating system design must be capable of handling the complex concurrency issues that arise with multicore architectures. Some of the generic areas of OS design affected by the presence of multiple cores are initialization, interrupt handling, scheduling, and locking. In the context of a Real-Time Operating System (RTOS), however, other key factors – such as priority scheduling, determinism, and interrupt latency – should be preserved in multicore architectures.

One such design optimization, known as *processor affinity*, may allow applications to request an affinity to a processor core. In this case, the operating system schedules the applications on the preferred processor core, as long as it does not affect overall system scheduling. A more rigid form of processor affinity is processor binding, where the task is always scheduled on the same processor core. This approach, however, may lead to priority inversions in RTOS environments. Operating system design should accommodate considerations such as processor affinity without degrading real-time determinism and responsiveness.

An RTOS such as LynuxWorks' LynxOS supports SMP and can schedule tasks dynamically and transparently between processors to efficiently balance workloads using available processors. It optimizes load-balancing support on multiple cores while preserving the key elements of real-time latency and determinism.

Optimizing applications for multicore architecture

Applications written for uniprocessor execution are not necessarily optimized for multicore architectures. Any inherent contention for resources that prevents execution parallelism could result in performance bottlenecks in multicore environments. Applications that are CPU-bound can exploit the full power of multicore architectures, while memory-bound or I/O-bound applications might need to be optimized to avoid the bottlenecks that arise due to bus contention in SMP architectures. Multithreaded applications can implicitly invoke resource contention as they request services from the operating system.

When looking at applications on multicore solutions, embedded developers must determine the allocation of functions within applications and redesign their applications to exploit parallelism. Developers must consider the design trade-offs of using multithreading versus nonmultithreading to harness the power of multiple processor cores. In some instances, applications may perform better on a single-core system.

Multicore computing continues to advance

New strides in multicore architecture make this an exciting time to be a developer. Every day, new approaches and strategies come to light that show us new ways to approach efficient multicore systems that are forming the heart of next-generation embedded computing systems. Choosing the most suitable OS for multicore architectures is a vital consideration, and overcoming issues like processor affinity, resource allocation, and parallel execution is also key to optimizing multicore systems.

Arun Subbarao *is the vice president of engineering at LynuxWorks, responsible for LynxOS technology and product development. He has more than 11 years of software experience in the embedded industry working on UNIX, Linux, and RTOS kernels, networking protocols, and high availability. He received his bachelor's degree in computer science while studying in India, his master's degree from SUNY Albany, and an MBA from Santa Clara University. He can be reached at asubbarao@lnxw.com.*

Spin-off COTS

An interview with CCAT, the DoD's very own technology incubator that operates in a very Silicon Valley-like way

EDITOR'S FOREWORD

CCAT is the Center for Commercialization of Advanced Technology, a government-funded program that exists to help commercialize technologies for the military. CCAT is a consortium of many organizations – some federal and some state – all rolled into one while sharing a common goal (see figure).

I sat down with CCAT's program directors and marketing director, who are based at two southern California universities: UC San Diego and Cal State University, San Bernardino. For convenience, all responses are aggregated into a single voice. Edited excerpts follow. – Chris Ciufo

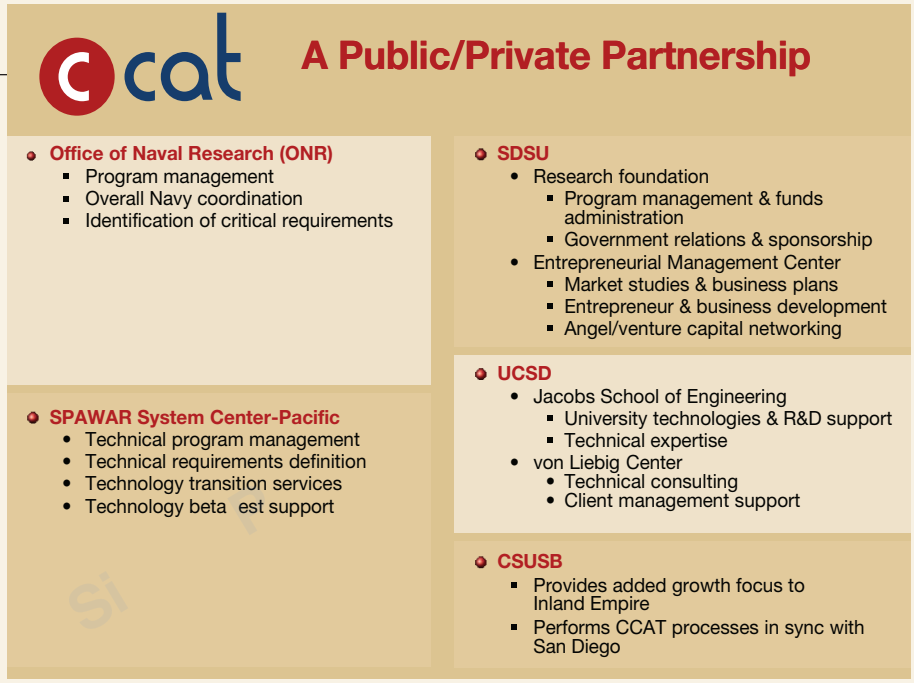

For an example of one of CCAT's successful commercialization efforts with PixonImaging, refer to www.mil-embedded.com/articles/id/?2840

MIL EMBEDDED: ***What does CCAT do?***

CCAT: CCAT's primary focus is formalization. In other words, doing what it takes to get [technology] out of a prototype stage and into the hands of first responders, into the hands of the military.

We do that in several different ways. We can do it through a product grant, which goes up to $75,000. This money is primarily used to miniaturize, to put on finishing touches, to do testing … or to do anything that it takes to get a prototype to a stage where it can be commercialized. We also offer business development services including help assistance plans, market study, and PR help. There are a lot of things that CCAT does to get a prototype into the marketplace.

MIL EMBEDDED: ***So what's the key one-liner description?***

CCAT: "Commercialization" is the operative word, as opposed to public relations or communications services. Commercialization implies funding for development, testing, and other services. The business part of this is every bit as important because we are dealing with technology to come out of research laboratories, whether they're collaborative or federal labs. A "bare" technology with a patent attached to it isn't worth very much. Its value only becomes evident as it begins to move along a process whereby it can be deployed as a product. And the farther down that path you go, the more business is actually involved than technology developed.

MIL EMBEDDED: ***Why's that?***

CCAT: Because you're not going to be trying to sell something when you haven't solved the technical problems with it first. So business is a very important part of it. Market analysis and indepth understanding of markets are extremely important because that's going to guide how you're going to shape the product in order to sell it in the right markets.

MIL EMBEDDED: ***What is the government's interest in this? Are there some stated goals?***

CCAT: The reason CCAT came into existence in the first place is that it's a new model, a very new *innovative* model on how to transition new technology to defense for defense purposes. We like to think that we fast-track the commercialization process. Many people have noted over the years that "traditional" technology transfer programs (government sponsored, that is) take a lot of time and a lot of money.

CCAT came along basically to try to improve that process. So our fast-track process basically has a goal of getting a technology into the marketplace – from prototype to marketplace – between a period of 12 to 18 months. Sometimes we make that, sometimes we don't. But we think that everybody who's looked at our program has come away fairly impressed with how limber we are, how fast we can change, and how fast we can move technology. One thing that's very important to understand is that we compete our solicitations. In other words, we get our technologies through a competitive process by working with our government sponsors who will tell us what they need at one time or another.

> **"The reason CCAT came into existence in the first place is that it's a new model, a very new *innovative* model on how to transition new technology to defense for defense purposes."**

MIL EMBEDDED: ***What are the mechanics of the solicitation announcement?***

CCAT: We do not go through CBD [Commerce Business Daily] or BAA [Broad Agency Announcement]; we market through our website, where we have a list of hundreds and hundreds of labs and companies. We also do press releases and other announcements to the press, so there are a lot of different ways to get the word out. We leave the solicitation open usually for 30 days, during which vendors and labs can submit a proposal to CCAT. We also receive applications from companies who are specifically looking for ways to grow into the government space.

MIL EMBEDDED: ***Where does CCAT's funding come from?***

CCAT: The funding comes through the Department of the Navy right now. Congress appropriates the funding that is used for both centers – San Diego and San Bernardino. The funding flows to our government sponsor, which is the Office of Naval Research (ONR). And then ONR flows the money to SPAWAR in San Diego. We have a contract between the San Diego State University Research Foundation and SPAWAR Center San Diego to manage the CCAT program on behalf of both centers. You could track our funding in the DoD budget under ONR appropriations.

MIL EMBEDDED: ***What is the alternative to the CCAT methodology?***

CCAT: The old method, for example, is that a technology might have been developed in a government lab, and the idea would have been to get it out either to a commercial or government customer via *technology transfer*. Typically government technology transfer people – while they may be good technologists or engineers – haven't the foggiest notion about how to get a technology into the marketplace. They don't know marketing, they often don't know business, and it's very difficult for them to interact with businesspeople on the other end who seek to license [their technology].

MIL EMBEDDED: ***Is CCAT a "threat" to the process of ATDs, government labs, and the existing way of doing business?***

> **"[Venture capitalists] like working with CCAT because we're able to give them a portfolio of technologies that's already been vetted."**

CCAT: No, not at all. We actually rely on them even when we're working with a small entrepreneur with a single technology. We need to work with the federal labs to be able to get those technologies tested and evaluated because the capabilities often don't exist except in those labs. CCAT's charter is not at odds with the charters of the government labs.

MIL EMBEDDED: ***How big is CCAT's annual budget?***

CCAT: It varies from year to year. Since CCAT became operational in 2001, we've received about $28 million in funding. The funding range is anywhere from $4.5 to $7 million per year. Obviously, the more the better and the more technology we can support. But even with $4.5 million, we are able to run about three solicitations per year and award or pick up anywhere from 8 to 10 new technologies at $75,000 each. In some very exceptional cases, we have gone beyond the $75,000.

Remember, one of our key benefits to commercialization goes beyond the monetary award. It's the business development, the business planning, the market study, the administering, and the "bringing together" of the venture capitalists and angel investors. They like working with CCAT because we're able to give them a portfolio of technologies that's already been vetted. Those that make it through the CCAT process and end up getting awarded are very attractive candidates for them to take a look at.

MIL EMBEDDED: ***Why are there two centers?***

CCAT: San Diego was the original center when we went operational in 2001. We had [university] informal collaboration between the two campuses, and San Bernardino's congressman became very interested in sponsoring an initiative that lead to a formal agreement between the two campuses. We agreed to split the CCAT [budget] earmarks and appropriations. The two faculties and collective group of entrepreneurs bring a rich collection of resources. CCAT also helps to bring additional investment to the Inland Empire [where CSUSB is located.]

MIL EMBEDDED: ***Are there plans to expand CCAT and institutionalize it across the DoD?***

CCAT: We have gotten lots of inquiries from across the U.S., and we've actually worked with several organizations in standing up similar programs. There's one in Hawaii called *High Tech Development Venture* with a sponsoring senator, and they've been up and running for a couple of years. We've worked with some folks in the Midwest also, so we see this as a model that's definitely portable to other regions. Every region has its own assets and capabilities, and each could customize its program accordingly.

With CCAT, we address a well-known problem in technology transfer, and that is the transition from a technology into an innovative product. It's really difficult to get money to fund those marketing studies, to fund that small additional product development that you need to do – and that's because a university or lab says "That's not our role." CCAT really feeds into that process in a way that very few programs – whether federal or university – actually do. CCAT is the only organization that we've found that provides both funding *and* assistance services to entrepreneurs or labs or the folks that actually come up with the technology.

Interviewees:

Tim Gerrity – Program Director, CCAT San Bernardino Center
Cal State San Bernardino, 909-809-9359
tgerrity@csusb.edu

Tom Sheffer – Program Director, CCAT San Diego Center
San Diego State University, 619-594-4135
tsheffer@foundation.sdsu.edu

Suzanne Finch – Director, Public Relations and Marketing
619-594-0206, sfinch@foundation.sdsu.edu

CCAT Program
www.ccatsocal.org

REGISTER TODAY! Updated Speakers and Conference Sessions

Register Online at www.milaeroforum.com

Under the COTS Umbrella

Military & Aerospace Electronics Forum™

The Conference and Expo for Innovation in Defense Applications

March 11 – 12, 2008

Pre-Conference Workshops: March 10

San Diego Convention Center • San Diego, CA

www.milaeroforum.com

Integrating COTS electronics and innovation for superiority on battlefields today and tomorrow

Participate in one of industry's premier military conferences and exhibitions!

- Explore the role of the GIG and network architectures in military systems
- Hear from leaders in DOD and the defense industry on managing COTS procurement in urgent programs
- Receive the latest updates on software defined radio applications and the effect of IPv6 on military systems
- Understand the latest challenges in ITAR compliance and export issues
- Explore thermal and power management solutions for warfighter systems
- Learn about emerging standards in information assurance and software security for defense applications
- Network with your peers from the defense electronics community

New for 2008! Pre-Conference Workshops on Monday, March 10.

New COTS Zone on the Exhibit Floor.

Visit us on our new location in San Diego, CA!

Go to www.milaeroforum.com to register online and get up-to-date information on the conference, exhibition, and other event activities.

Featured Speakers Include:

Dr. Ron Jost
Deputy Assistant Secretary of Defense for C3, Space & Spectrum
Office of the Assistant Defense Secretary for Networks and Information Integration

Dr. Stephen M. Jarrett
DBA/ITM Chief Technologist
SPAWAR Systems Center Charleston

Dr. Robin Keesee
Deputy Director for the Joint Improvised Explosive Device Defeat Organization (JIEDDO)

Col. Timothy A. Kokinda (ret)
Former Assistant Chief of Staff, G6, XVIII Airborne Corps, Fort Bragg, NC and Systems Engineer
Lockheed Martin Integrated Systems and Solutions

Howard Pace
Deputy Joint Program Executive Officer for Joint Tactical Radio System (DJPEO JTRS)

James Robles
Senior Technical Fellow
The Boeing Company

FLAGSHIP MEDIA SPONSOR: Military & Aerospace Electronics

OWNED & PRODUCED BY: PennWell

Editor's Choice Products

Multicore small form factor for mil apps

With Intel driving its processors well down into deep, deep submicron (45 nm) territory, multicore CPUs now run faster while sipping less power. Add to this the huge economic benefits of consumer volumes, and then small form factor boards like the MI935 from IBASE make good sense in many quasi-rugged military applications. The Mini-ITX form factor SBC supports Intel's Core 2 Quad/Core 2 Duo/Celeron 400 Sequence processors (Conroe-L) and Intel's latest Q35 Express chipset with a full 1,333 MHz FSB.

With notebook computer-sized dimensions of 6.7" x 6.7" (170 x 170 mm), the SBC has two DDR2 DIMM slots and can support up to 4 GB of DRAM. The board also includes a 10/100BASE-T Ethernet port, integrated VGA, 8x USB 2.0, 2x SATA II, 1x eSATA, and 2x COM ports. Optional I/O includes a 1x PCIe GbE, a second VGA via PCIe (x16), and four more COM ports. With a company focus on industrial PC products, the MI935 might be a good choice for full-featured but small-sized PC-based military systems.

IBASE Technology • www.ibase.com.tw • RSC# 35572

PC/104 does MIL-STD-1553

Recently selected by rugged shoebox supplier Parvus in a U.S. Navy aircraft upgrade program, the Data Device Corporation BU-65578Cx PC/104-*Plus* 1553 card is ideal for fitting 1553 into tight spaces. Based upon DDC's own Extended Enhanced Mini-ACE (E2MA) 1553 chipset – the offspring of an industry standard – the card can bring up to four dual-redundant 1553 channels to bear in legacy systems. With 2 MB of parity memory per channel and BIT, robust transmissions and data transfer are assured.

But interestingly, MIL-STD-1553 is rarely found alone. It's often accompanied by myriad other legacy and proprietary connectivity schemes, especially in system upgrades. For this reason, the PC/104-*Plus* card also includes five user-programmable digital I/O channels, an IRIG-B time code input, an external time-tag clock input, and a 48-bit/1 microsecond time stamp capability. As expected, a -40 °C to +85 °C temperature version is available.

Data Device Corporation • www.ddc-web.com • RSC# 35574

15 W in a mere 1.0 in^2

Looked closely at a circuit board lately? Designers have to contend with more voltages than a drawer filled with old cell phone chargers. And myriad voltages mean DC/DC converters galore, consuming power and precious real estate. RECOM Power Laboratory claims a breakthrough in high power and diminutive size. The company's RP-15xxxxSA series boasts 15 W in an incredible 1" x 1" (25.4 x 25.4 mm) package at only 0.39" (10 mm) height. This compares with competing devices measuring 1" x 2" – saving a cool square inch of board space.

The converter is galvanically isolated, operates at up to 87 percent efficiency, and can withstand an ambient temp of 60 °C without needing special cooling. That's an impressive power density of 2.3 W/cm^3. Outputs are set at 3.3 V, 5 V, 12 V, and 15 V, while inputs include: 9-18 V, 9-36 V, 18-36 V, 18-75 V, and 36-75 V. The case is a six-sided, nickel-coated copper material, and vacuum potting on the bottom assures high shock and vibration tolerance.

RECOM Power Laboratory • www.recom-development.at • RSC# 35577

Rugged LCD controllers

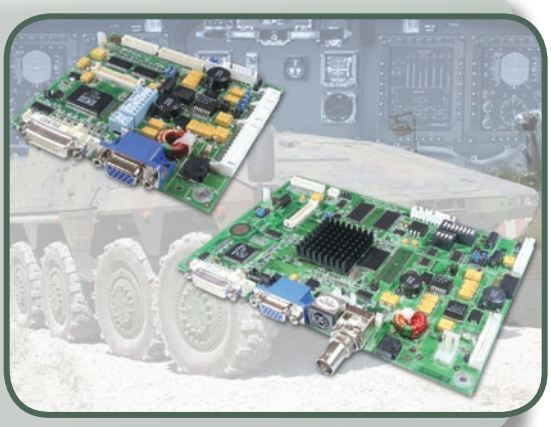

So you've hardened the chassis, ruggedized the boards, chosen solid-state storage for reliability – even shock-mounted the LCD. But what about the LCD controller board? That's the realm of Digital View's HE-1400 and HE-1600 Series rugged LCD controllers. These COTS LRUs specifically target industrial and military installations with a -40 °C to +80 °C operation, and "dirty" 12 VDC power that can vary by as much as 25 percent. Shock and vibration are handled from the board design up, including locking connectors and such attention to detail as low-mass tantalum capacitors that don't act like cantilevered masses. They are even available in conformal-coated versions.

The HE-1400 is 4.2" x 3.6" and supports LVDS and TTL LCD panels at SXGA (4:3) and WXGA (16:9 at 1,366 x 768 pixels) feeding DVI and ARGB. The HE-1600 is a fully buffered, multisync interface controller feeding both analog and digital up to UXGA resolution (that's 1,600 x 1,200) over DVI, dual VGA, composite video, S-Video, and Component Video. It too handles standard 4:3 and widescreen 16:9. Both controllers feature remote management via RS-232.

Digital View • www.digitalview.com • RSC# 34063

Multi hard disk sanitizer/duplicator

Military programs are acutely concerned with data and software security, especially when it comes to sanitizing disks used in actual missions. And the same thing applies to desktop computers from the Pentagon or to any forwardly deployed operations center. But on the flip side, sometimes it's necessary to *replicate* a disk's entire content for back-up purposes, or for use in parallel missions or platforms. While not designed with the DoD/MoD/DND in mind, Aleratec's HDD Cruiser might be just the product you're looking for.

The desktop unit can simultaneously sanitize up to four 3.5" HDDs without requiring any connection to a PC or network (2.5" drives require an adapter). Disks can be wiped using NSA-approved algorithms such as the seven-pass "DoD Wipe," or a simple one- or three-pass erase can also be performed. Once temporarily installed in a tray, the drives are mechanically locked during operation to prevent tampering or interruption. The aluminum trays are also designed to conduct heat away from drives during extended operations, and the built-in 10 ultra-quiet fans also keep temperatures down. Finally, the Cruiser can replicate one drive's contents to the other three simultaneously.

Aleratec • www.aleratec.com • RSC# 35573

Battlefield LAN? Will travel.

Ethernet is as ubiquitous on the modern battlefield as disposable batteries. All manner of equipment, be it ground-, ship-, or air-based relies on 10, 100, or 1000 Mbps Ethernet ports. So Curtiss-Wright Controls Embedded Computing decided to bring 1 and 10GbE to the up-and-coming VME VPX form factor. The company's VPX6-685 FireBlade II resides on a 6U VPX module and is specifically designed for networking in extremely harsh environments. With switching and routing capabilities, plus VITA 48 REDI two-level maintenance options, this board is ready to travel.

The heart of the board includes 12, 20, or 24 1GbE interfaces capable of autonegotiating 10/100/1000 speeds. The board can also support 4x 10GbE ports as part of a blazing-fast backbone configuration. Current versions of the board offer front-panel optical ports (1000BASE-SX), while future versions will route fiber over one of the VPX connectors. There's IPv4/v6 support, wirespeed routing, enhanced security, BIT, and a whole host (no pun) of management interfaces, protocols, and software from CLI and Telnet to SNMP and NAT.

Curtiss-Wright Controls Embedded Computing • www.cwcembedded.com • RSC# 35576

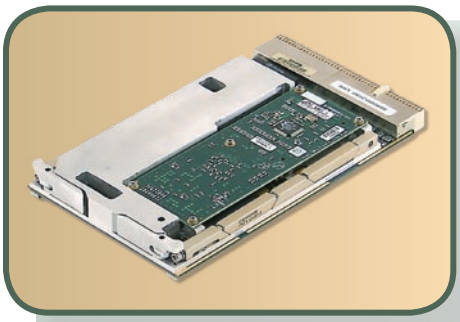

3U conduction-cooled SBC with legacy compatibility

The company SBS Technologies popularized the conduction-cooled flavor of 3U CompactPCI single board computers because the size offers such nifty space savings over 6U. The 3U module can easily fit into legacy ATR boxes, making it a nice upgrade path. Similarly, GE Fanuc Embedded Systems – which purchased SBS a couple of years ago – hasn't forgotten about legacy migration with its latest 3U conduction-cooled CompactPCI SBC, the CR5. That's because it's FFF backward compatible with the previous CR4 version.

Sporting an Intel Core Duo LV running at 1.66 GHz, the SBC includes up to 2 GB of DDR2 SDRAM with ECC, a 400 MHz memory bus via Intel's 3100 ICH memory and I/O controller, and a 667 MHz FSB. There are two GbE ports, two SATA ports, two USB 2.0 ports, and eight of the ever-popular-in-defense GPIO ports. An optional mezzanine card adds CompactFlash and SVGA graphics.

GE Fanuc Embedded Systems • www.gefanuc.com • RSC# 33448

Filter connector can withstand 1,000 cycles

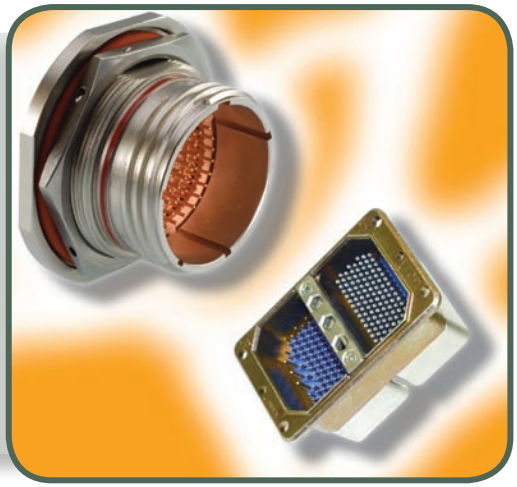

How often do you think about the screws, washers, or standoffs in an embedded system? Yeah, me neither. But you might put a little more thought into the electrical connectors – but just barely. The thing is, it's those connectors that take much of the abuse in deployed electronics, especially if the system must withstand MIL-STD-810 or MIL-STD-1344. ITT Cannon's Chip on Flex filter connector is designed to meet MIL-STD-1344 testing, boasting 1,000 thermal shock cycles from -55 °C to +125 °C.

Compared to ceramic planar array block capacitors used in conventional filter connectors, the MIL-DTL-38999 style Chip on Flex technology mounts individual chip caps on a pad adjacent to each feed through contact. This eliminates stress points, increasing connector life from a mere 50 cycles to 20 times that. The connector handles voltages from 200 Vdc to 120 VAC rms at 400 Hz. Current ratings range from 5 A, 7.5 A, and 15 A, depending upon wire gauge.

ITT Electronic Components • www.ittcannon.com • RSC# 35578

Editor's Choice Products are drawn from OSP's product database and press releases. Vendors may add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/ and submit press releases at www.opensystems-publishing.com/news/submit. OSP reserves the right to publish products based on editors' discretion alone and does not guarantee publication of any product entries.

By Sharon Schnakenburg

x86

Arcom Control Systems, Inc.

Website: www.arcom.com
Model: GEMINI
RSC No. 33648

A 5.25" EBX form factor single board computer for use in embedded and industrial applications • Intel Core 2 Duo • Ideal for graphics-intensive applications • 4x Gigabit LAN, up to 3 GB of DDR2 RAM • Dual video output • 6x USB 2.0 • 8-24 Vdc power input requirements

Avalue Technology Inc.

Website: www.avalue.com.tw
Model: ECM-GM965
RSC No: 33834

3.5" Intel GME965 Core 2 Duo single board computer • Supports 65 nm Intel FC-PGA 478/FC-BGA 479 Core 2 Duo CPU • Intel GME965 chipset • One SODIMM up to 2 GB DDR2 SDRAM • 2x 24-bit LVDS, Dual View, 5.1 CH/7.1 CH audio with S/PDIF, Dual Realtek RTL 8111B GigaLAN • One EIDE, one Type I/II CF, one Mini PCI, two SATA, four USB 2.0, 8-bit GPIO, jumperless serial port configuration • 12 V power input

Concurrent Technologies, Inc.

Website: www.gocct.com
Model: 6U CompactPCI Core 2 Duo SBC
RSC No: 34763

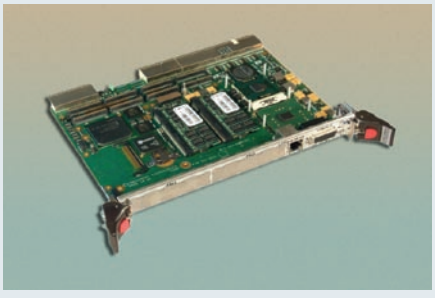

Model PP 412/03x • 6U CompactPCI Core 2 Duo SBC supporting two PMC sites, including one XMC site • Up to 4 GB DDR2 ECC SDRAM • All in a single system or peripheral slot • Two PMC sites supporting up to 100 MHz PCI-X • Front and rear I/O • One XMC x8 PCI Express • Supports 2.16 GHz or 1.5 GHz Intel Core 2 Duo processors • Uses Intel E7520 server chipset • Up to 4 GB dual-channel DDR2-400 ECC SDRAM in a single slot • 33/66 MHz CompactPCI operation in system or peripheral slot • Three GbE ports with option for PICMG 2.16 • 2x SATA-150 and EIDE interfaces via J3/J5 • Graphics front and rear • Three USB 2.0 and three RS-232 interfaces • PICMG 2.9 IPMI interface and PICMG 2.1 hot-swap support • Support for Linux, Windows XP, Windows XPE, Windows Server 2003, Windows 2000, QNX, LynxOS, VxWorks, and Solaris • Option for extended temperature operation, -25 ºC to +70 ºC

x86

Curtiss-Wright Controls Embedded Computing

Website: www.cwcembedded.com
Model: SVME-1900
RSC No: 35579

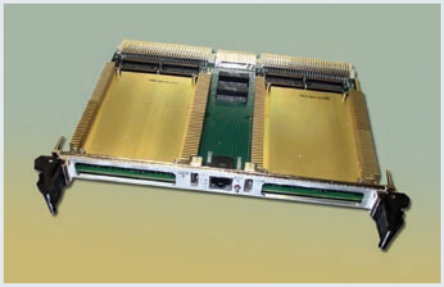

A single 6U VME64x SBC featuring two independent Ultra Low Voltage Intel Core Duo processor-based compute platforms, each running at speeds of 1.67 GHz (15 W clock rates) • Can be configured with either two dual-core Intel Core Duo processors or two single core Intel Core Solo processors • Ideal for aerospace and defense platforms • Available in air-cooled, ruggedized configurations • 2 MB L2 advanced transfer cache per CPU • 2 GB DDR2 SDRAM with ECC • 2 GB USB user flash, formatted as two individual USB drives • Intel E7520 Memory Controller Hub (MCH) • Intel 6300ESB I/O Controller Hub (ICH) • I/O includes: two PMC (64-bit) mezzanine expansion sites, two (10/100/1000) GbE ports, four USB 2.0 ports, 12 COM (four RS-232, eight RS-422) ports, two SATA ports, 12 DIO lines • Offered in a wide range of ruggedization levels • Supports Windows XPe, Solaris 10, Linux, and VxWorks

ICP America, Inc.

Website: www.icpamerica.com
Model: NANO-9453 "EPIC"
RSC No: 35070

Intel Core 2 Duo EPIC SBC with dual independent video onboard (no pin headers) • Dual Broadcom PCIe GbE, SATA/3G, PCIe mini card socket, and CompactFlash • Intel Core 2 Duo, Core Duo/Solo, 667 MHz and Celeron M 533 MHz Yonah • Memory up to 1 GB DDR2 SODIMM • Intel 945GME GMCH + ICH7M • Dual Broadcom BCM5787 GbE • Intel GMA950 PCIe graphics with dual independent VGA ports onboard and LVDS flat panel support • 6x USB 2.0, 3x RS-232, 1x RS-422/485, 2x SATA/3G, 1x IDE, 1x CFII, 1x LPT1, 1x PS/2, 7.1 HD audio, 1x IrDA • PCI-104 and PCI Express mini card expansion slots • Supports Windows Vista, XP, XP Embedded, Windows CE, and Linux

ITOX Applied Computing

Website: www.itox.com
Model: CT930-B
RSC No: 35108

An ECX SBC • Small form factor SBC is an Intel open standard with board dimensions of 105 mm x 146 mm (4.13" x 5.75") • Intel Core Duo LV and ULV processors • Intel 945GM Express chipset • Up to 2 GB of DDR2 667 MHz memory • GbE controller • Interfaces: two SATA and one IDE • LVDS, DVI, and VGA graphics ports

x86

BOSER Technology Co., Ltd.

Website: www.boser.com.tw
Model: HS-2620 **RSC No:** 34719

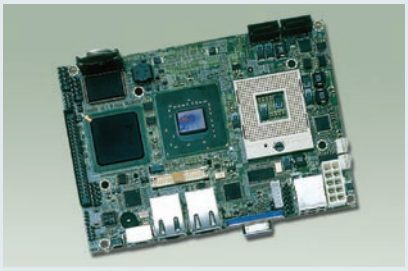

Intel Core 2 Duo Mobile procesor SBC 1.8~2.4 GHz with CF, CRT/LVDS, dual GB LAN, audio, SATA II, USB 2.0 • One SO-DDRII socket with a max. capacity of 2 GB • Intel GM965 GMCH/ICH8-M chipset • Winbond W83627DHG super I/O chipset • Intel GM965 graphics controller • 24-bit/48-bit LVDS panel display interface • Dual Marvell GbE controller • ALC883 7.1CH 3D audio controller • Intel ICH8-M Serial ATA controller • GPIO, two COM, four USB 2.0 • Hardware monitor function

Core Systems

Website: www.coresystemsusa.com
Model: CECD **RSC No:** 33470

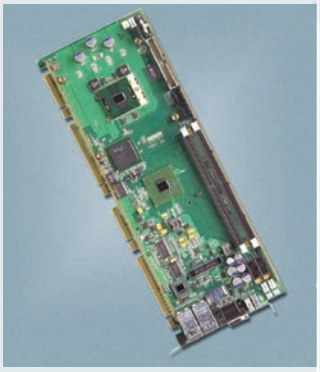

A single board computer long-life Intel Core Duo Processor (T2500) with the Intel 945G chipset • Dual-core and single-core processor options supported • PCI Express graphics-class SHB supports x16 video and graphics cards or ADD2 cards • Direct connect video option via the chipset's Intel Graphics Media Accelerator 950 • Compatible with the SHB Express (PICMG 1.3) specification • Supports dual-channel DDR2-667 memory, 4 GB maximum • Dual GbE ports plus one 10/100BASE-T backplane interface • Integrated RAID 0, 1, 5, and 10 implementation support via four SATA/300 ports

ICOP Technology, Inc.

Website: www.icop.com.tw
Model: VSX-6154 **RSC No:** 33517

A PC/104 SBC based on Vortex86SX System-on-Chip (SoC), an x86-based SoC designed for embedded applications • At 300 MHz, this low-power SBC draws 250 mA at 5 Vdc • Soldered-on RAM, fanless design uses low-profile heat sink • Integrated video/LCD, Ethernet, USB 2.0, 4 serial ports, parallel, IDE, RTC, and more • Supports Linux and Windows CE, and support is available for legacy applications running on older versions of Linux and Windows

x86

One Stop Systems Inc.

Website: www.onestopsystems.com
Model: OSS-6U-CPCIe-CPU-08 **RSC No:** 35376

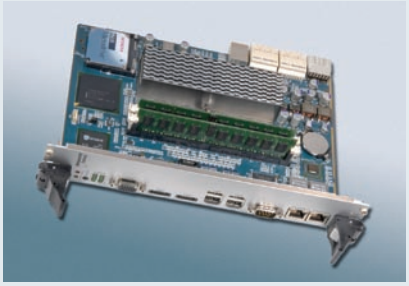

The 6U CompactPCI Express Core Duo-based CPU board is available for industrial, power, military/aerospace, medical, telecom, and scientific applications • Intel Core Duo 1.66 GHz CPU • Features front panel I/O with standard interfaces including two GbE ports, two Serial ATA ports, four USB 2.0 ports, one RS-232 port, and activity LEDs • Includes an SXGA-compatible display interface and onboard CompactFlash interface • Supports up to 4 GB of DRAM

WIN Enterprises Inc.

Website: www.win-ent.com
Model: MB-06051 **RSC No:** 35462

Fully featured SBC with Intel Pentium M processor, Intel 855 chipset, and low-power requirements • Intel Pentium M processor • CompactFlash socket • FDD interface • Four USB ports (1.1 or 2.0) • 10/100 or 10/100/1000 LAN

WinSystems, Inc.

Website: www.winsystems.com
Model: EBC-855-G-1.8-1 RSC No: 35580

1.8 GHz Pentium M EBX-compatible SBC • Operates at -40 °C to +70 °C for rugged, industrial, and test and measurement applications • Based on Intel's 855GME chipset with the ICH 4 communications controller and integrated Extreme Graphics 2 video 3D controller • RoHS-compliant, processor- and I/O-intensive board • Supports up to 1 GB SDRAM and up to 8 GB CompactFlash • 10/100BASE-T Ethernet port (with remote boot capability) • VGA and dual-channel LVDS flat panel video • Mini PCI connector for an 802.11 wireless networking module • Four USB 2.0 ports, four serial COM ports, AC97 audio (5.1 codec), LPT, and PS/2 • Operates Windows XP embedded and Linux • Operates QNX and VxWorks RTOSs • Measures 146 mm x 203 mm (5.75" x 8.0") • Requires only +5 V and typically draws 2.1 A (typical) with 1 GB of DDR SDRAM installed

PowerPC

MEN Micro Elektronik GmbH

Website: www.menmicro.com
Model: A17 6U VME SBC **RSC No:** 34777

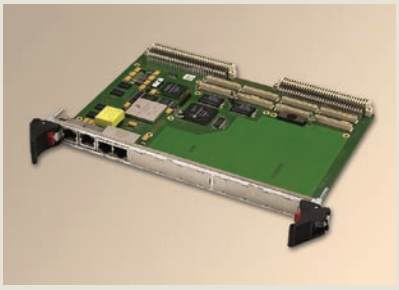

Advanced PPC-based SBC that acts as master or slave in a legacy VMEbus environment • Provides 2eSST performance levels while maintaining backwards compatibility with older standards • PowerPC MPC8548/1.33 GHz • One-slot 2eSST VMEbus master and slave using the new Tundra TSI148 bridge controller • Up to 2 GB (ECC) DDR2 RAM • Flash disk, FRAM • 2 GbE, 2 COMs at front • 2 GbE via P0 • Two PMC slots (one slot also XMC) • FPGA for individual I/O functions • MENMON BIOS for PowerPC cards

Mercury Computer Systems, Inc.

Website: www.mc.com
Model: Momentum VPA-200-SBC **RSC No:** 34950

A dual 7448 PowerPC VXS SBC • Full-featured 6U VME board offering a higher level of computing power and bandwidth for users with high-d nsity needs • VITA 41 performance and flexibility with high-speed serial fabric • Two standard PMC-X sites for off-the-shelf PCI mezzanine cards • Tundra Tsi148 VME bridge chip provides VME 2eSST support • Access to wide ecosystem with Linux or VxWorks operating systems • RapidIO on VSX P0 option

Performance Technologies

Website: www.pt.com
Model: AMC141 **RSC No:** 35527

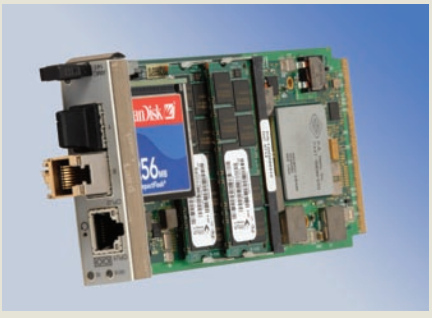

64-bit AdvancedMC compute module with the PA6T-1682 PowerPC processor designed for high-performance embedded applications • Single-width, mid-size compute module with PA6T-1682 PowerPC 2.0 GHz with 2 MB shared L2 cache • CompactFlash site for onboard program and OS storage • Up to 4 GB of DDR2 SDRAM with ECC w/PC2-5300 interface • Quad 1 G or 2.5 G Ethernet • Dual 10 G Ethernet • Supports both 32- and 64-bit versions of the NexusWare CGL OS and development environment

PowerPC

Aitech Defense Systems

Website: www.rugged.com
Model: C108 **RSC No:** 35044

6U VME high-performance addition to Aitech's C10x SBC series, for use in both new and legacy military embedded applications • 1.4 GHz of processing power via the G4+ PowerPC (PPC) MPC7448 processor • GbE port, 2 Fast Ethernet ports, 2 dual-redundant MIL-STD-1533B interfaces, 2 USB ports, 8 serial ports, and 16 discrete I/O channels • Two PMC expansion slots • Onboard, optically isolated CANbus 2.0B to interface with modern vehicle electronics and major functional subsystems • 128 MB of flash for user application storage, 32 MB of boot flash memory, and up to 4 GB of high-density flash for mass data storage • Two watchdog timers allow the user to set a programmable time-out interval

Extreme Engineering Solutions

Website: www.xes-inc.com
Model: XPedite8070 **RSC No:** 33674

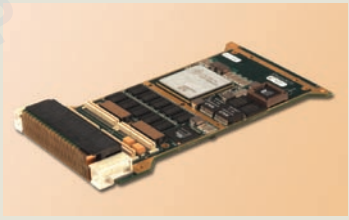

A 3U VPX-REDI (VITA 46/48) form factor SBC • Dual P.A. Semi PA6T-1682 2.0 GHz PowerPCs • SBC consumes only 34 W • 2 GB of DDR II (1 GB/core), 32 Mb NOR, and 1 GB NAND flash • I/O includes PCI Express, 10 GbE, and dual isolated 1 GbE ports • Supports Linux, VxWorks, or PNE 1.4

GE Fanuc Intelligent Platforms, Inc.

Website: www.gefanucembedded.com
Model: SBC310 **RSC No:** 35411

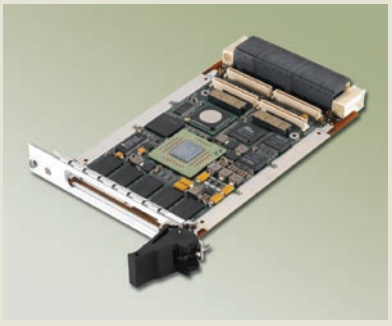

3U VPX rugged SBC that uses a dual-core MPC 8641D PowerPC at 1.0 GHz • Up to 1 GB SDRAM • One PMC or XMC site for mezzanine card • I/O includes two 10/100/1000 BASE-T Ethernet ports, two USB 2.0, and a SATA disk interface • Two RS-232 and four GPIO • Backplane connections via two 4-lane PCIe ports, with several user-defined I/O pins available

For more information

Enter the product's RSC# at
www.mil-embedded.com/rsc

SUPER ★ TALENT®
THE BEST MEMORY

Reinventing Storage

Rugged, Reliable, Unparalleled

SATA/ IDE SSD

• Performance

Operating Temp: -40 to +85C
Sequential Read Rate: 60 MB/sec (min)
Sequential Write Rate: 45 MB/sec (min)
Sequential Read IOPS: 58,675 IOPS
Access Time: 0.1ms
Power Consumption: 85% less than Average HD

• Reliability

Shock: 1500G (operating)
Vibration: 16G (operating)
MTBF: > 1,000,000

• Endurance

Read: Unlimited
Write: >140 years @ 50GB write-erase/day

"AWARD WINNING PRODUCT"

CeBIT
"Innovation of the year" Award

DISKCON USA
Best of Show Award

- 100% compatible with standard hard disk drives
- IDE form factors in 1.3" / 1.8" / 2.5"
- Designed and manufactured in USA
- Capacities range from 4GB to 256GB
- SATA form factors in 1.8" / 2.5" / 3.5"

Super Talent USA Headquarters
www.supertalent.com
Tel: 1.408.934.2560 | Email: sales@supertalent.com
2077 North Capitol Avenue San Jose, CA 95132 U.S.A

Crosshairs Editorial

Small form factors: A new SIG in town

Stood up at last fall's ESC Boston, the SFF-SIG is ready to roll in ships, tanks, trucks, and wide-body aircraft

By Chris A. Ciufo, Editor

First question: What's the most popular open standard board type used in military systems? Answer: the personal computer, both in desktop and laptop sizes. But if we narrow the question to "rugged, deployed systems," you'd primarily find 6U VME boards, followed by 3U CompactPCI, SEM-E, and then a bunch of other types too numerous to mention. For the past 15 years of the COTS era, VME's anything-but-small size has dominated rugged, deployed systems.

But VME has limitations in power (too much), interconnect (parallel bus), physical size (approximately 6" x 9"), and *cost*, forcing mil designers to increasingly look at other options in all but the harshest systems. With an eye toward PC-based, low-cost, modular, industrial and military systems, the brand-new Small Form Factor Special Interest Group (SFF-SIG) will be ready to go public with their plans in a few months[1]. We got a special sneak preview that I thought would interest our readers. If successful, the results of the SIG's efforts may wind up in your next military system.

The companies behind SFF-SIG are still in stealth mode because many participate in other standards bodies, some of which may perceive the SFF-SIG to be competition. Of these, the PC/104 Consortium has probably the most to lose since 100 percent of the SFF-SIG companies do or have participated in the PC/104 Consortium in its 15-plus year history[2,3]. The companies kicking off the SFF-SIG include Octagon Systems, Samtec, Tri-M Systems, WinSystems, and VIA. According to the SIG's president, Colin McCracken, there are three board members (VIA, Octagon, and WinSystems), seven members, and more in discussion.

SFF-SIG has three working groups – SBCs, Stackables (like PC/104 or PMC mezzanines on VME), and Computer-On-Module (such as PICMG's COM where the CPU resides on the mezzanine card plugged onto an I/O-laden carrier board). This year, the SIG plans to work on all of these three areas with a special emphasis on a new stackable spec it won't reveal until April. In the meantime, the SFF-SIG is following in the PC/104 Consortium's footsteps with its "Adopt a Spec" policy whereby existing SFF board types of value become standardized for multi-vendor adoption and market competition.

As memory serves, this is how EBX, PC/104-*Plus* (which added PCI to the ISA bus), and EPIC became PC/104 Consortium standards some years ago. I maintain a list of more than 60 different small form factors – some open, most proprietary – which should provide fertile ground for open standards adoption by the SFF-SIG. According to McCracken, some of those SFFs are already under discussion by the SFF-SIG. This is particularly exciting news for low-power and low-cost defense systems, as many of these SFFs were designed for semi-rugged consumer applications such as long-battery-life handsets or "droppable" portable mobile devices.

Key to achieving the SIG's mundane-sounding goals will be a combination of focused marketing efforts and a VME VITA-like rulebook that crisply brings topics up for a vote so that they don't become "over-engineered, academic corner cases," says McCracken. We like the sound of that. One of VME's strengths has long been its VITA Standards Organization (VSO), which is aggressively run both technically and administratively. Additionally, the SFF-SIG recognizes that it's market success that matters over technical success. Hence, the charter focuses on "*rapid* ecosystem growth" with a "divide and conquer" plan to combine the strengths of worldwide leaders.

On the technical side, right now the group is focused on creating a slightly larger-than-PC/104 SFF that is I/O-centric. Part of the problem with many of those 60 SFFs I mentioned earlier is that they were designed around the CPU (AMD, ARM, Intel, VIA … pick one), and the CPU's bus also becomes the board's bus. If the CPU changes, then the standard is in trouble.

Instead, the SFF-SIG wants to make its first new standard I/O-centric, expressly looking at forward and backward migration between processors and even other SFFs. This means that the proposed new standard would accommodate many of today's (and tomorrow's) CPUs and peripherals, while paying attention to system aspects of tech refresh and spiral insertion strategies. For the military, this is translated to code portability and Pre-planned Product Improvements (P3I). This is a page right out of VME's rule book, and one that resonates very well in defense because 10-year upgrades are just business as usual.

But small form factors come and go, as do their sponsoring companies. The same might happen to the SFF-SIG. Will it be around next year or the year after? We don't know, but we hope so. Personally, I like what I'm hearing. I can't wait to tell our readers about their first SFF when it's unveiled on April 8, 2008.

Chris A. Ciufo
cciufo@opensystems-publishing.com

[1]Full disclosure: I personally helped to catalyze the formation of the SFF-SIG, although I have no direct interest in it. OpenSystems Publishing, *Military Embedded Systems'* publisher, is an unofficial sponsor of the SFF-SIG.
[2]For more on the PC/104 Consortium, check out OpenSystems Publishing's *PC/104 and Small Form Factors* magazine at www.smallformfactors.com.
[3]By the way, VITA and PICMG memberships overlap by at least 40 percent.

GE Fanuc
Intelligent Platforms

AXISLite. For those who think that seeing is believing.

Power, productivity, portability. And now, free to try.

Since it launched, AXIS – the powerful yet easy to use multiprocessor software development environment for VME, VXS and VPX from GE Fanuc Intelligent Platforms – has been at the heart of leading edge applications that are being deployed much faster.

But we know that there are plenty more organizations out there who could benefit from its sophistication, its flexibility and its ease of use – but who we haven't yet convinced. That's why you can now download AXISLite for free from our web site – so you can see for yourself how it lives up to our claims.

Being free, AXISLite doesn't do everything AXIS does – it's a scaled down version. Even so, we believe there will be enough there for you to see what you've been missing.

If state of the art multiprocessing and DSP systems are what you're developing, and performance and portability are important to you, you'll see that nothing else comes close to what AXIS can do.

www.gefanucembedded.com/axis

AXIS™ Lite

GE FANUC

YOUR DATA IS CRITICAL. YOU CAN DEPEND ON US.

Curtiss-Wright understands that your data is critical. Whether on the battlefield or back at the base, the ability to record, store, retrieve, and playback your data is a must. Curtiss-Wright has solutions for all of your secure recording and storage needs. For in-the-field applications, our 6U VPX* NAS (Network Attached Storage) device provides unmatched rugged data storage in any environment. And back at the base, you can depend on our SDRxR Streaming Data Recorder for recording, retrieval and playback. When your data is critical, depend on Curtiss-Wright. (*VITA 46/48)

SDRxR Recorder

The SDRxR Streaming Data Recorder and VPX NAS (Networked Attached Storage) device are two of Curtiss-Wright's newest next-generation data recorder solutions designed to enable recording, archiving and retrieval of your critical data whether under harsh battlefield conditions or back at the base.

CURTISS WRIGHT Controls
Embedded Computing

www.cwcembedded.com

SECURE DATA... ABOVE & BEYOND