Software special: Programming tools come of age

VME and Critical Systems

www.vmecritical.com

VOLUME 29 NUMBER 1

SPRING 2011

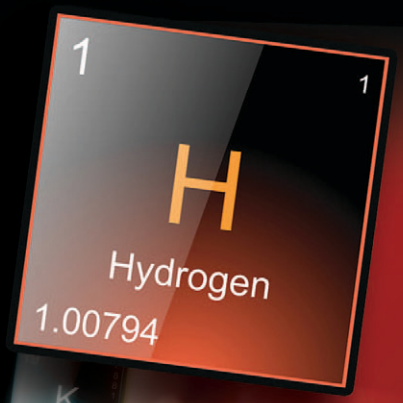

Atomic-energy batteries power memories for decades

Also:

Latest VSO updates on VITA 62, 65, and 66

www.opensystemsmedia.com

» Where can I find a single source for my embedded computing project? «

Kontron is your single source for COTS products and specialized systems designs - built to meet the most demanding requirements.

COTS Boards & Enclosures
- » Standard form factor
- » Semi & full custom

Development Tools
- » Software tools and development platforms for rapid application modeling and testing

Software Services
- » BSPs: VxWorks, Lynuxworks, Linux, Windows, ...
- » Virtualisation
- » IPMI

Systems Integration
- » Full range of environmental compliance
- » Choice of form factor or architecture
- » Backplane or module based

Total Custom Solution
- » Application ready
- » Concept to production

CRITICAL QUESTIONS ... ANSWERED

VX6060
6U VPX Dual Intel® Core™ i7 Single Board Computer

- » Available in air-cooled & conduction-cooled versions
- » Compliant to VITA46 (VPX), VITA65 (OpenVPX) and VITA48 (VPX REDI)

Kontron EAPI
Kontron Embedded Application Programming Interface

- » Unified middleware for access & control of hardware
- » Includes standard libraries & APIs

Cobalt™
Reduced SWAP - High Performance Embedded Computer

- » Small Form Factor
- » Intel® Core™2 Duo or Atom™ processor options
- » Custom I/O option

CONTACT US

AP Labs is now a part of **kontron**

☏ 1-888-294-4558 ✉ info@us.kontron.com www.kontron.com/military

intel Embedded Alliance Premier

If it's embedded, it's Kontron.

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

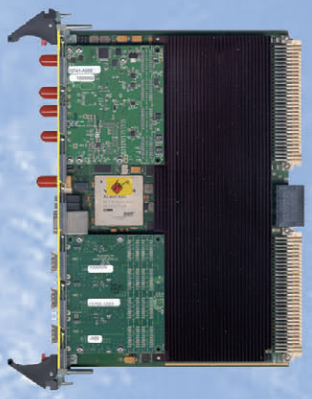

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

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

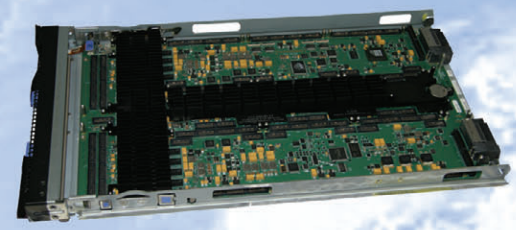

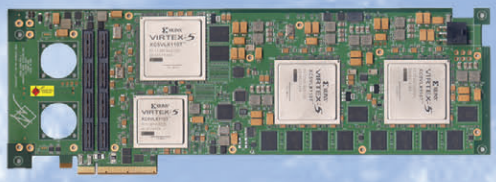

Ultimate Modularity
From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
Input/Output Modules Include:
Quad 130 MSPS thru Quad 550 MSPS A/D
1.5 GSps thru 5.0 GSps A/D, Quad 600 MSps D/A,
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10G, 40G or 100G Ethernet or SFPDP

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

Annapolis
Micro Systems

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

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

www.vmecritical.com

VME and Critical Systems

Spring 2011 Volume 29 Number 1

Columns

Departments

Events

American Telemedicine Association 2011 Annual Meeting
May 1 - 3, 2011 • Tampa, FL
www.americantelemed.org

Embedded Systems Conference - Silicon Valley
May 2- 5, 2011 • San Jose, CA
esc-sv09.techinsightsevents.com

ON THE COVER: The radioactive isotope Tritium (Hydrogen-3) illuminates watch dials and exit signs ... and also batteries. An interview with start-up City Labs reveals how nano-voltaic batteries work like solar cells, using Tritium particles. See story on page 16.

Published by: OpenSystems media.

ENVIROINK™
The inks used to print the body of this publication contain a minimum of 20%, by weight, renewable resources.

Features

Web resources

Subscribe to the magazine or E-letter
Live industry news
http://opensystemsmedia.com/subscriptions

Submit information:
http://submit.opensystemsmedia.com

White papers:
Read: http://whitepapers.opensystemsmedia.com

ISSN: Print 1941-3807, ISSN Online 1550-0403

VME and Critical Systems is published four times a year (Spring, Summer, Fall and Winter) by OpenSystems Media, 16626 E. Ave of the Fountains, Ste 201, Fountain Hills, AZ 85268. ***VME and Critical Systems*** is free to qualified engineers or management dealing with or considering open-system technologies. For others, paid subscription rates inside the US and Canada are $45/year. For first-class delivery outside the US and Canada, subscriptions are $60/year (advance payment in US funds required). Periodicals postage paid at St. Clair Shores, MI, and at additional mailing offices.

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

POSTMASTER: Send address changes to ***VME and Critical Systems***
16626 E. Ave of the Fountains, Ste 201, Fountain Hills, AZ 85268

ADVERTISER INFORMATION

E-CASTS

Level-up with the 2nd Generation Intel Core Processor Family
March 29, 2011 • 10am MST
Presented by: Kontron, LynuxWorks, Intel Embedded Alliance

OpenVPX: From Specs to Solutions
March 30, 2011 • 11am MST
Presented by: Pentek, Curtiss-Wright Controls Electronic Systems, GE Intelligent Platforms, Curtiss-Wright Embedded Controls

VME and Critical Systems

OpenSystems media.

@VME_VPX_mag

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

Military & Aerospace Group

Chris Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

Sharon Hess, Assistant Managing Editor
sharon_hess@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

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

Monique DeVoe, Web Content Editor
mdevoe@opensystemsmedia.com

Christine Capuano, Web/Editorial Assistant

Sydney Lines, Products Editor

Brandon Lewis, Copy Editor

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

David Diomede, Art Director

Steph Sweet, Creative Director

Joann Toth, Senior Designer

Matt Jones, Web Developer

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Sales Group

Dennis Doyle, Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Christine Long
Digital Content Manager
clong@opensystemsmedia.com

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

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

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

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

Ad Coordinator
Steph Sweet
ssweet@opensystemsmedia.com

Reprints and PDFs

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

Editorial/Business Office

16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystemsmedia.com

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

Editor's Foreword

By Chris A. Ciufo, Editor

Optional ethics? No freakin' way

If you make a mistake, own up to it. And always stick to your ethical principles, else they're worth as much as how you'll eventually be treated. I was reminded of these mantras this week, right before I penned this column. So this month, we're taking a break from Core i7, Thunderbolt, and OpenVPX. Consider it "Chicken Soup for the VME Soul."

Sometimes I impress myself too much and am carried away by my own "brilliance." The end result is always the same: I outsmart myself and screw up … while usually laughing at my own stupidity. Such is the explanation for why a half-page advertisement is at the bottom of this page.

Simply stated, we ran out of space in this magazine for a last-minute ad, but I incorrectly approved its addition at the last minute. Here's where my brilliance "shined" ever so brightly: I failed to ask if there was actually available space in the magazine.

In my defense, several people involved in this fiasco also failed to notice the error of adding 0.5 pages into a magazine with exactly 0.0 pages left. But I made the call and the error, so I'm taking the blame for the mistake – and penalizing myself by slicing 0.5 pages from this column.

Which (finally!) brings me to my point about ethics: Doing the right thing ... is always the right thing. It's not an optional tactic that can be employed when convenient. In the ad screw-up scenario, I made the mistake, pure and simple. Blaming someone on my staff would've been counterproductive and unethical. At 20 years old at my first job, my mentor drilled ethics into my head. But what is "ethical behavior"? It's as hard to define as "irony," but most non-psychotics recognize an ethical decision as one you feel good about after doing. Speaking of ethics, one of my favorite ethical statements is read by John Rynearson of VITA's VSO before every meeting:

> "VITA's patent policy regarding the use of patented technology in standards is posted on its website. … Working group members should read these policies. ... Members who are aware of any patents or patent applications that might be essential to the standards proposed by this working group are required to disclose them."

This isn't just a legal threat, it's a statement about doing the right thing. Past VSO disclosures have yielded second sources for connectors, and predatory patents have been avoided.

So do the right thing – be ethical in all of your personal and professional dealings. And remember my own screw-up and subsequent public admission of guilt. Oh, by the way, please also patronize the advertiser below. If you buy something, maybe my boss will also forgive my mistake.

– *Chris A. Ciufo, Editor*

VITA
Open Standards, Open Markets

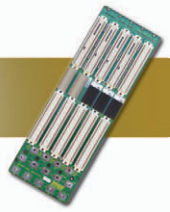

VITA News

By Ray Alderman

The big shift in 2011

Defense Secretary Robert Gates' January 2011 memo on DoD budget cuts gave a good idea of what to expect in the coming years: a big shift from weapons platform spending to intelligence platform spending. In this memo, he put the Marine Expeditionary Fighting Vehicle program on two-year probation and said the DoD will revise and refresh existing ground vehicles such as the Bradley Fighting Vehicle and the Mine Resistant Ambush Protected (MRAP) vehicle. So, there will not be any major ground vehicle programs in the near future. Also, Gates reduces the number of F-22 and F-35 fighter purchases. For now, the DoD will purchase more F-18s and refresh the older F-15s.

Spending on UAVs

Spending on Unmanned Aerial Vehicle (UAV) Intelligence, Surveillance, and Reconnaissance (ISR) platforms will definitely increase for many years to come. The Global Hawk platforms have been refreshed several times since their debut, and the new high-altitude, long-endurance, hydrogen-powered AeroVironment Global Observer UAV embarked on test flights in January. The Gorgon Stare ISR system flew missions for the first time last December in Afghanistan, attached to the underside of a Predator UAV.

Several things are clear here. We are beginning to field new ISR modules on older, first-generation, propeller-driven UAVs such as the Global Hawk, Predator, and Reaper. Also in the test phase are second-generation, jet-powered UAVs like Boeing's Phantom Ray and Northrop Grumman's X-47B UAV platform. The first-generation UAVs will be refreshed with new ISR capabilities for a while, and more of these platforms will be put into service as the testing and refinement of the next jet-powered UAVs continue. Then, focus will shift to deploying second-generation, jet-powered UAV platforms.

Refreshes with VME

If there are VME card cages in older mil/aero and ground vehicle systems, they will remain VME. They will be continually refreshed for many years to come with newer VME boards (faster processors, faster A/D boards, and so on). My best estimate indicates that in 2010 there was more than $600 million in VME boards sold to refresh programs. The Gates memo suggests that VME sales to refresh programs will probably increase in the next few years because the purchase of new ground vehicles and weapons systems is unfunded in the new DoD budgets. That means that VME will remain the top COTS technology (based on sales) for military applications for several years to come.

New platforms and VPX

The shift to additional spending on ISR platforms suggests that VPX (VITA 46) will be the primary technology used. VPX contains the trusted supply chain, processing power, ruggedness, and bandwidth to handle the massive amounts of data these new ISR systems will produce. Preliminary reports on the Gorgon Stare platform operation in Afghanistan last December showed that the amount of data produced overwhelmed the intelligence analysts. More supercomputer processing power will be needed on the ground to run algorithms against all the data coming from ISR systems in the air, and to refine and filter data that analysts must view and analyze. Those systems will also be VPX-based, possibly using General Purpose Graphics Processor Units (GPGPUs) such as the CUDA architectures from NVIDIA.

Small form factors

VITA 73/74/75 Small Form Factor (SFF) products will find their way into these new ISR systems where Size, Weight, and Power (SWaP) consumption are critical: in UAVs (both refreshes to the older platforms and the new jet-powered vehicles), upgrades to existing ground vehicles, and perhaps in upgrades to F-18s and F-15s. At some point, SFF cubes can be used to terminate the old copper wire 1553 bus and translate that into a high-speed fiber-optic network on the aircraft.

We all feared that the DoD budget cuts would dampen our opportunities, but that is not the case. As it turns out, they just shift spending from big weapons system platforms to new, more sophisticated ISR platforms. That is good news.

For more information, contact Ray at exec@vita.com.

Stroke of Genius

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

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

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

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

Now Available in VPX!

COTS & Rugged Versions

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

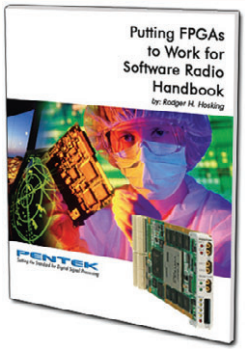

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

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

By Gregory Bolstad

OPENVPX INTERCONNECTS...

SSD technology drives OpenVPX storage blades to new levels

Data storage is an increasingly important component of many OpenVPX (VITA 65) systems. Intelligence, Surveillance, and Reconnaissance (ISR) systems, for example, generate large volumes of data that must be recorded in real time for later analysis. These systems require high capacity and performance, often combined with strict Size, Weight, and Power (SWaP) constraints.

Storage implementations used with these systems have historically fallen into two categories: 1) low-capacity, low-performance embedded storage boards, or 2) higher-capacity, higher-performance but physically larger and heavier external storage boxes or subsystems. However, current flash-based Solid-State Drive (SSD) technology, combined with optimized storage controller architectures, has fueled the development of embedded storage blades that provide very high levels of consistent performance, reliability, and capacity.

This new generation of storage blades enables fully embedded blade-level storage that yields the same (or better) performance and capacity as large external storage boxes, but combined with ultra low SWaP. A single slot 6U OpenVPX embedded storage blade can now offer greater than 600 MBps of sustained read/write performance, a capacity of more than 3 TB, and power consumption of less than 30 W. This blade-level solution can be contrasted with external RAID or NAS boxes that may consume hundreds of watts and weigh 50 or more pounds.

OpenVPX connection supports many storage models

The rich backplane I/O capabilities provided in OpenVPX systems (particularly as compared to VME or VXS) allow embedded storage blades to support a wide variety of host interfaces, protocols, and storage models. Storage interface options often used in OpenVPX include 1/10 GbE, Fibre Channel, backplane PCIe, and others. Typical storage models include: Direct Attached Storage (for block-level data access, or *RAID*); Network Attached Storage (for file sharing) using NFS, FTP, CIFS/SMB, and other protocols; and high-performance data recording modes. The underlying storage controller may provide support for both RAID0 (highest performance) as well as RAID5 (redundant data protection).

As a primary or secondary OpenVPX board-to-board interconnect, PCIe and 1/10 GbE provide an efficient interface to embedded storage blades. A four-lane "Gen 1" PCIe backplane interface provides about 2 GBps aggregate storage access bandwidth, while newer "Gen 2" interfaces provide twice that. Backplane 1/10 GbE provides an efficient (but lower performance) method for hosts to access storage blades in a file sharing or NAS mode, as well as providing RAID/DAS access using iSCSI or Fibre Channel over Ethernet (FCoE).

High-performance RAID enabled by next-gen SSDs

Flash-based SSDs are the foundation of the new generation of OpenVPX embedded storage blades. While flash SSDs have been available for years, early versions were unsuitable for high-performance embedded usage, hampered by low capacity and poor (and highly inconsistent) performance. SSDs are still not all created equal, and thus it is critical to match SSD characteristics to the end application.

Three key performance metrics for SSDs in OpenVPX and other systems are: 1) sequential write performance (measured in MBps); 2) random write performance (measured in IO/s); and 3) what is perhaps the most important for real-time systems: *consistency in performance*. The first two characteristics are easy to measure, but the third is not. Nearly all SSDs suffer from a significant "memory" effect, where previous usage (and access patterns) can dramatically affect future performance. A "fresh" SSD might write sequential data at 200 MBps (sustained). But that same SSD, after having been exposed to a write-intensive access pattern for only a short time, might only be able to write data at 50 MBps.

Advancing VITA 65 systems with SSD-based blades

Storage blade and controller architectures that leverage the strengths of new-generation SSDs (and compensate for their shortcomings) result in powerful and compelling blade-level OpenVPX storage. The flexibility and performance offered by this new generation of embedded storage blades suits them to many OpenVPX system applications, replacing large, power-hungry external RAID or file server boxes with a high-reliability, single-blade storage solution (Figure 1).

Figure 1 | Critical I/O's new OpenVPX StoreEngine storage blade provides simultaneous block (DAS/RAID), file (NAS), and high-speed recording access to up to 3 TB of onboard RAID 0/5 at sustained rates up to 600 MBps.

***Gregory Bolstad** is Chief Systems Architect at Critical I/O, LLC. Prior to Critical I/O, he was CTO of Broadband Storage and held system architecture design positions at Raytheon and Hughes Aircraft Company. Contact him at gbolstad@criticalio.com.*

Next-generation programming tools come of age

By Kumaravel Pandurangan

The landscape of software development is changing. It's no longer just creating new software from the ground up. In today's programming environment, engineers are often taking an existing software application and updating it with new features and functionality. However, there's a catch. These new features are being added along with the addition of new application requirements such as updated hardware, a new programming language, or a new operating system or RTOS. This method of updating software is beneficial in that it's allowed for applications to advance at an accelerated pace. However, the arsenal of tools needed by software engineers on today's projects is much different from what was required in the past.

The software world is at a new level. The availability of software has increased exponentially within the past 3 decades. This has spawned a need for different skill sets for software engineers for a couple of reasons.

First, as new hardware vendors introduced new hardware platforms and semiconductors, they often provided software as demo code to show off the various capabilities of the hardware. Over time, this demo code evolved and became more complex, containing startup code, advanced demos, and initialization routines. Then more middleware vendors started supplying software only, and the creation of middleware packages and library code was introduced. Of course, the Web helped to make much of this code available to almost anyone. This availability and sharing gained mainstream traction through the open source movement. It is now possible to find enough "pieces" of code to create a solid starting point for almost any project.

Second, there is a shift to software integration rather than always writing an application from the ground up. As many in the embedded defense industry know, commercial and government projects often do not "start from scratch." After money, time, and effort are spent building and debugging a software application, it is often the mandate from management and project leaders that software developers start new projects using software from existing projects already deployed. These projects aren't reconstructed by starting over, but are essentially "debugged into existence." This occurs when new parameters – including hardware, new features, a new language, new middleware, or a new RTOS – are introduced to an existing, working code base. The team is required to take the existing code and make it run using the newly defined operating parameters.

Let's look at the three key elements of this new coding integration paradigm and examine where modern programming tools fit in:

1. Debugging into existence, not debugging for errors
2. Maintaining a minimal level of performance before optimizing
3. Eliminating unproductive work

Debugging into existence

Compare "debugging into existence" to "debugging for errors." When first writing a software application, programmers spend a certain amount of time debugging the program, while checking for operational issues and functionality that comply with the *original* written specifications. When integrating software, they are debugging to make the application run in the *new* environment with new hardware or a new operating system or RTOS. This integration approach demands a different method of debugging, because it requires a familiarity with all the new components. This is a problem because all that learning adds a lot more time to the development schedule. In this situation, the use of language conversion or porting tools to handle some of the code transition can save weeks, months, or even years.

Let's look at a simple example of task creation in an application. Here is a sample call to a common RTOS application function.

```
void SampleCreateTask( void* parameters)
```

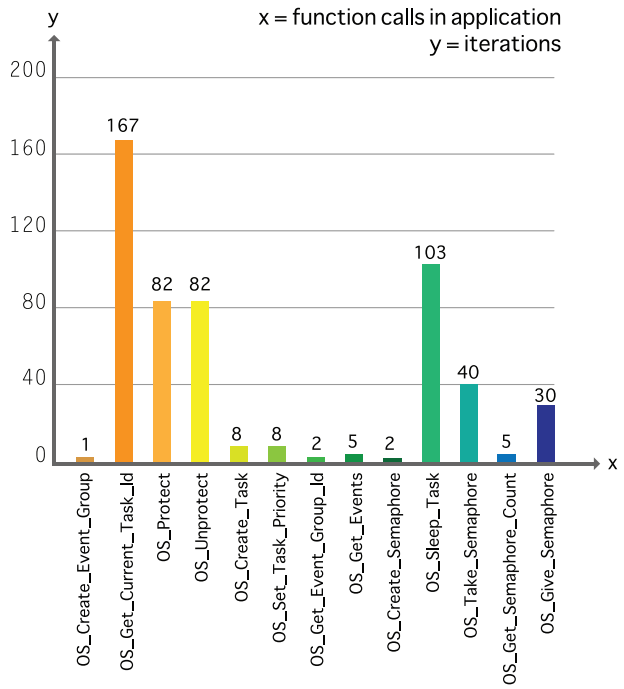

Figure 1 | A graphical view generated from a graphical profiling tool, showing the number of times functions have been called

This call is passing a pointer to a control block that holds the individual parameters for the task. When switching from one RTOS to another, the programmer must have knowledge not only of how many times this function is called, but they must also have a deep understanding of both task structures in each RTOS. When you extend this knowledge out to all of the components of an RTOS – initialization, tasks, scheduling, mailboxes, queues, and so on – the use of a language conversion tool to automate the transitions makes sense.

Maintaining minimal pre-optimization performance

During the initial development of a software application, much work is put into the optimization and performance. Often times, many hours are spent tweaking the software application to be as small as possible or to run as fast as possible. When updating or integrating projects, the initial goal isn't to make the application run as fast as possible, but to run at least as fast as the prior version of the application. How is this done? Using hooks inside the code for timing and execution speed, and then also using graphical profiling tools to help indicate the bottlenecks in the new code.

For profiling the application, a programmer can manually insert hooks into the code. This normally includes searching through the application code and inserting appropriate profiling functions on the entry and exit of the function call. When the application consists of thousands of functions and millions of lines of code, this method is overwhelming. When using an advanced profiling tool with the application and an RTOS, there is the ability to automatically insert hooks into the RTOS and application code. The end result is a graphical view of many important timing and performance metrics, including the number of function calls made, time spent inside a function, memory usage, and so on. Figure 1 shows a sample report from a graphical profiling tool used to automate the process.

Elimination of unproductive work

When one of the components in the application is changed – such as an RTOS or hardware component – some work can be automated and some can't. For instance, using Search and Replace to update function names is error prone and tedious. When changing the programming language, the error-prone work exists when programmers attempt to manually convert one language to another. The controversy comes in because many software developers consider important all the work they do. However, in reality, there are tools that exist that can help automate much of the software work when moving the software application to a new RTOS or new hardware.

For example, consider the C/C++ and Ada languages. Both are very powerful and used in many military software applications. Many existing applications today are integrating both Ada code and C/C++ code. The problems occur when programmers attempt to integrate or translate code by hand. The differences between the two languages are huge, and a simple example can illustrate some of the many differences. Table 1 shows the various operators in C/C++ and Ada. Trying to convert code by hand is an example of unproductive work because

Operator	C/C++	Ada
Assignment	=	:=
Equality	==	=
NonEquality	!=	/=
PlusEquals	+=	
SubtractEquals	-=	
MultiplyEquals	*=	
DivisionEquals	/=	
OrEquals	|=	
AndEquals	&=	
Modulus	%	mod

Table 1 | Operators for C/C++ and Ada vary immensely. Empty cells on the Ada side indicate no direct equivalent in C/C++.

"Imagine trying to debug the software application if there is a tiny mixup in hand-converted code. It seems better to rely on automated tools for this type of work. The use of an automated tool could condense weeks and months of manual code conversion down to a few hours."

there are tools like MapuSoft's Ada-C/C++ Changer, a language conversion tool that can do this automatically. A programmer's time is better spent elsewhere on the project. Further, imagine trying to debug the software application if there is a tiny mixup in hand-converted code. It seems better to rely on automated tools for this type of work. The use of an automated tool could condense weeks and months of manual code conversion down to a few hours.

The changing role of programming and tools

The role of the programming has changed dramatically over the past few years. The bulk of the work now is on integration and changing code for new situations, rather than creating new software from the ground up. Thus, engineers must be equipped with a new arsenal of tools, including those that render the ability to debug into existence, maintain minimal pre-optimization application performance levels, and eliminate unproductive manual procedures. MapuSoft Technologies' conversion tools, including its OS PAL product, help accommodate this change. CS

Kumaravel Pandurangan *was appointed Director of Engineering of MapuSoft Technologies' India office in August 2007. He brings more than 12 years technical and managerial experience in real-time and embedded systems to MapuSoft's team. He handles project management, software implementation, network security management, and team management at MapuSoft's India Development Center. He graduated from Bharathiar University in 1995 with a Bachelors Degree in Engineering, specializing in electronics and instrumentation. Contact him at kumaravel@mapusoft.com.*

MapuSoft
251-665-0280
www.mapusoft.com

NightHawk RCU™ Rugged Computer Unit

Light weight. Light years ahead. Hi-performance. Lo-SWaP. And OTS!

Harness improved processing and SWaP efficiency in the free-air-cooled Aitech NightHawk RCU™ Rugged Computing Unit – without sacrificing durability or versatility. The Remote Interface/Data Concentrator capabilities of this PC-based embedded unit are ideal for demanding applications in defense, aerospace, and rugged industrial environments, such as CBM, GPS, and RFID monitoring in manned or unmanned vehicles.

*Minimize weight...*at just 4.5 lbs., this space-saving unit delivers exceptionally efficient SWaP performance in a slim-profile 195-in.^{3} package.

*Expand I/O versatility...*with 18-36 vdc input, standard PC I/O interfaces, video-cam input, independent hi-def XGA and standard TV video outputs, plus extensive internal expansion support to add peripheral I/O interfaces for military/industrial embedded apps.

*Enhance your processing efficiency...*with a low-power Intel® Atom™ (Navy Pier) 1.6 GHz RISC microprocessor, 2 GB DDR2 SDRAM and 8 GB SSD memory, plus optional CompactFlash or 250 GB Solid State Disk (SSD) storage with room for customization.

*Take rugged to greater extremes...*with an environmentally sealed, extruded aluminum housing and innovative convection/radiation cooling design that dissipates more than 22 W at +55°C. With MIL-DTL 38999 connectors, the sealed NightHawk protects against windblown fine dust and sand particles, while numerous EMI/RFI features meet MIL-STD 461 emission and susceptibility limits, all in a shock-, vibration-, and corrosion-resistant design for harsh mechanical, climatic, chemical and electrical environments.

*Limitless applications...*maximize your options with the flexibility to configure the NightHawk RCU™ into the control unit you need it to be. Contact us today for detailed power and I/O specifications.

Aitech

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

Celebrating 25 Years of Success

Photography Courtesy of Lockheed Martin

Atomic technology powers crypto keys for the long haul

20-year betavoltaic batteries rely on tritium electrons

I remember the first time I played with a solar cell as a child experimenter. I used a flashlight, solar cell, and another flashlight bulb to create a weakly lit filament. It would take me another 10 years and a college education before I understood the principle behind using photons to create an electrical current. City Labs has performed a similar miracle, creating a nanocurrent betavoltaic battery that not only holds a charge for 20 years, it meets NRC standards so anyone can buy it. This technology is a genuine breakthrough and I'm particularly enthused at the horizons opened by the company's NanoTritium batteries. I spoke with City Labs founder Peter Cabauy to get the inside glow, so to speak. Edited excerpts follow. *Chris A. Ciufo, Editor*

VME: ***I am thoroughly unfamiliar with your company. Let's start there.***

CABAUY: We started in this industry in 2005 as a couple of scientists and engineers just saying, "Well, we want to do some cool technology." But we didn't necessarily fall in love with batteries. However, we saw the need for advanced battery technology in FPGAs. Early on, we realized there was a market here because Lockheed Martin, Raytheon, and Northrop Grumman wanted to buy these batteries, but at the time these types of batteries didn't exist. Even then, Lockheed Martin, Orlando, FL gave us a purchase order for prototype batteries.

VME: ***OK. City Labs issued a press release in December saying that the company was awarded a million-dollar Air Force Research Labs [AFRL] contract for your 20-year lifespan NanoTritium batteries for microelectronics.***

CABAUY: Yes. We are making a betavoltaic battery. Betavoltaics are very much like photovoltaics, except that instead of receiving sunlight to make a current out of a solar cell, we are receiving beta radiation, caused by beta electrons from tritium. Tritium is the most benign radioisotope, and is commonly used for medical tracers, exit signs, and luminous watch dials, for example.

VME: ***Why use betavoltaic batteries? What are the advantages?***

CABAUY: Instead of having the "betas" or electrons coming off the tritium and hitting a phosphor so that it glows, they are hitting a semiconductor PN junction and thereby creating a current just like a solar cell would. The advantage is that while none of the chemical batteries work well through all temperatures, betavoltaic batteries are very resilient and can withstand temperatures of -50 °C up to +150 °C. Also, tritium's half-life is 12.5 years. So the battery could last more than 20 years.

VME: ***So tritium emits electrons that strike the PN junction and then a current flows?***

CABAUY: Well, the electron strikes the PN junction and creates electron-hole pairs. And a PN junction is able to separate those electron-hole pairs and capture them and make them into usable electrical current. As I mentioned, it's very similar to the way a solar cell would work. In fact, it is a very tweaked-up solar cell that can receive these very low-energy electrons coming off of tritium. Tritium is such a good radioisotope for this type of application because the electrons – or as they call them, "betas" – can be stopped with just a thin sheet of paper. So, for example, someone wearing a diver's wristwatch isn't worried about radiation affecting them in any way. Tritium is a very benign radioisotope, as far as radioisotopes go.

VME: ***No current flows until there's a load put on the device?***

CABAUY: Since we are creating electron-hole pairs, there always needs to be a load, but it is very different from a chemical battery. A solar cell has an open-circuit voltage and a short-circuit current. If the solar cell is connected to an application or load, the maximum voltage and maximum current will rise and are contingent upon the connected device.

VME: ***With a typical beta hitting a typical PN junction, are we talking nano-amps of currents?***

CABAUY: Just like solar cells, you can connect as many as needed to get the desired current or desired voltage, but typically it's somewhere between 50 and 100 nano-amps per square centimeter, depending on tritium density and semiconductor efficiency. At City Labs, we've achieved 25 to 50 nanowatts per centimeter squared. The semiconductor can be made quite thin, so it's just a matter of stacking. Right now, it would be about 10 microwatts per cubic centimeter by stacking, but we expect to achieve up to 50 microwatts per cubic centimeter.

Keep in mind that I was referring to the power density, but the energy density is very large. The average energy density over 20 years is about 5x the energy density of a lithium thionyl chloride battery, which is the highest-energy density lithium battery.

VME: ***Over 20 years, what changes within the battery?***

CABAUY: Imagine that you're just pumping out little nano-amps and micro-amps through time, so it's like a little trickle. But if it's a continuous trickle – and with radioisotopes it's always continuous – even continuous trickling of water can create a Grand Canyon over a long period of time. It's the same thing with these betavoltaics – a few drops of this trickle charging current will create a lot of energy over time. Connecting these to a lithium thin-film battery or an ultacapacitor can do quite well over time for powering up a chip or RF device.

VME: ***Is that the typical system configuration: One of these betavoltaics connected to a lithium ion battery to power the circuitry?***

CABAUY: That's one configuration. The central application we've been looking at is to take the betavoltaic battery and use it to power up FPGA encryption keys held in SRAM – at any temperature. Nanoamps is typically sufficient and it can easily be provided by the NanoTritium battery. Furthermore, the NanoTritium battery can reliably withstand extreme temperature conditions.

Battery-backed SRAM power is extremely important in fielded military applications such as a UAV, when it's flying about 30-40,000 feet at -60 °C, then comes back down and lands in a very hot desert. Chemical batteries can see failure pretty quickly in a case like that. And when you're dealing with encryption keys, failure is not something you want to happen lest the system go brain dead upon power failure.

VME: ***Do you have an off-the-shelf version in existence yet?***

CABAUY: We're finalizing the design components and assemblies for the EOL200Y20 series, which is the FPGA encryption key battery and the first COTS betavoltaic battery. We're projecting a summer release.

And I'd like to point out that the betavoltaic battery is nothing new. This idea has been around since Paul Rappaport, one of the fathers of solar cells, did his groundbreaking solar cell development [in the 1950s].

VME: ***It sounds like these advancements have been a long time coming.***

CABAUY: Yes. It took decades of research and work for people to figure out the most viable candidates for radioisotopes: tritium, promethium, and Nickel 63.

Promethium is interesting; only has a half-life of 2.6 years, and it's very difficult to get these days. Promethium 147 requires a bit of metal to be able to shield out the radiation coming out of it.

On the other hand, nickel is expensive and the flux is very, very low, so the usable power that comes out of it is not as good. Of the three, tritium is the preferred choice to generate betas.

VME: ***Let's talk about the coming COTS betavoltaic battery you mentioned earlier, the EOL200Y20 Series. "EOL" can have a negative connotation in the defense industry.***

CABAUY: It actually *does* stand for "end-of-life," yes, but what we're basically saying is that at end-of-life, it will have 200 nanowatts at year 20. That's why our NanoTritium betavoltaic battery is called "EOL200Y20."

VME: ***What are the negatives of this technology? Does it outgas? Does it tend to be explosive?***

CABAUY: I think the only negative is that it's taken this long to commercialize a betavoltaic battery. As I mentioned, this idea came out in the 1950s and it seems like it's only now at the point of usability, mostly because the FPGA and microelectronics industries have lowered their power requirements to start using long-term power like this.

VME: ***Which other applications do you foresee for betavoltaic batteries?***

CABAUY: Within the defense industry, we see specifics with security supervisors, low-power microcontrollers like the MSP430 from Texas Instruments, or even ubiquitous computing. For instance, they're usable for these very low-power microprocessors that can sleep on picowatts and operate at nanowatts.

Also, these batteries could power the University of Michigan's Phoenix Processor, a very low-power device implanted

transdermally to detect whether a tumor is regrowing. Or for a glaucoma patient, they put sensors and processors inside the eye. NanoTritium battery technology, which can be made as thin and small as a cubic millimeter, can be combined with a battery chip sensor that can then be implanted into all sorts of medical applications that no one's even thought about.

VME: ***Any other possible applications before we move on?***

CABAUY: There are also applications within Homeland Security, intelligence sensors, structural integrity sensors, sensor actuators, space satellite pro-power sources, and crystal oscillators, if they're within 15 microwatts. Combine these batteries with a lithium thin-film battery, some type of lithium ion battery, or an ultracapacitor, and that can keep these batteries topped off for years until a usage need arises. So the battery's self-discharge is eliminated and then it can be used like a reserve battery. Or it can be used in combination with duty-cycle and it can operate and provide a power burst every week, every day, for RF signals and so forth. The applications are unlimited.

VME: ***As we wrap up, is there anything else we should know about City Labs or tritium-based batteries?***

CABAUY: Yes. Since the NanoTritium batteries are radioisotope powered, you can't just go buy them at the corner market. Typically, their use requires radiation training and a radioactive license on behalf of the user or the individual purchasing them.

VME: ***What does this mean? That the user needs to be a nuclear engineer?***

CABAUY: Normally, radioisotope products are classified under three general headings: specific license, general license, or exempt license. The smoke alarm you see in your home or office is an exempt license. The Exit sign you see at Home Depot or in a commercial aircraft is illuminated with tritium, also falling under a general license. And, medical tracers and the things the Department of Energy is using – those are under the specific license.

Very few products obtain the general license or exempt license status, because they have to be proven inherently safe through all sorts of testing processes. But City Labs recently achieved regulatory approval in the form of a general license issued by the Florida Department of Health - Bureau of Radiation Control; as an Agreement State, Florida's general license mirrors the policies and guidelines established by the Nuclear Regulatory Commission. So now we can sell our NanoTritium batteries anywhere in the United States. You can just put the order in and receive it without any training, without any prior radiation license. There's never been a general license classification for a battery before, so that says that we broke some ground in the industry. **CS**

Peter Cabauy *cofounded City Labs, Inc. in 2005 and serves as CEO. Peter received his Ph.D. from the University of Michigan in Applied Physics and a BS from Florida International University. Contact him at pcabauy@citylabs.net.*

City Labs
www.citylabs.net

VITA 62 brings COTS power supplies to VPX

By Patrick Shaw

Though VITA 46 (VPX) and VITA 65 (OpenVPX) systems are rapidly becoming a defense industry mainstay, a system-level interoperability standard for these systems' power requirements has been a missing ingredient. However, the VITA 62 (Power Supply Modules) working group now aims to fill the void.

VITA 65 (OpenVPX) and VITA 46 (VPX) systems are revolutionizing the defense industry, having overcome initial system-level interoperability issues – for the most part. However, a system-level interoperability void still exists when it comes to power generation for these advanced systems. This gap has led to *ad hoc* power designs that are system specific but not interoperable between VPX systems. However, VITA 62 (Power Supply Modules) allows systems designs, using OpenVPX as a guideline, to specify power requirements for VITA 62 compatible systems. This will create an ecosystem of many VITA 62 products that are market driven for the large and growing VPX community.

Accordingly, the VITA 62 working group was started in September 2008 with the support of both module vendors and system integration companies. (It currently has about 40 participating companies.) The mandate of the committee is to provide a standard set of design guidelines for both hardware and systems to allow vendors to create power supply modules to power VPX systems. The development taking place in the VPX community and in the OpenVPX initiative has primarily focused on the development of modules, slots, and backplane profiles for applications, with little emphasis on how these systems are to be powered.

The VITA 62 standard creates a set of design guidelines by which industry power supply vendors can build products that will be intermateable to VITA 62-compatible backplanes. OpenVPX currently only refines the power guidelines defined within VITA 46. It does not address any issues pertaining to powering such OpenVPX systems. When VITA 62 products are later defined in OpenVPX module, slot, and backplane profiles, they will allow system architects to specify power requirements for OpenVPX systems.

In addition to power conversion/conditioning, VITA 62 also defines secondary or utility slots for such system functions as inrush current, output control, Electro-Magnetic Interference (EMI), synchronization, and nuclear detection. VITA 62 has addressed and resolved many of these issues in the current draft, and some highlights of the accomplishments to date are described in the following sections.

Connectors and pins: Vital to VITA 62

VITA 62 leverages all the effort within the VITA 46, VITA 47 (environmentals), VITA 48 (VPX-REDI), and VITA 65 standards while providing a power-generation architecture that fully supports the needs of the critical-systems embedded community. Key to that goal is VITA 62's new multivendor connector set for both 3U and 6U modules. The connectors fit into the current VITA 46 mechanical outline, thus facilitating the retrofitting of existing backplanes. The complete connector design, which facilitates efficient power conversion/conditioning, is included in the specification to support multiple connector companies' implementations of VITA 62 connector sets. The connector reuses the same guide pins as defined in the ANSI/VITA 46.0 standard. The pins' placement, also ANSI/VITA 46.0 compliant, eases modifications of the existing PWB designs to support VITA 62.

The VITA 62 power supply connector pin allocation was the main driving force in balancing the application's needs versus availability of pins in these new connectors. The 3U posed the greatest challenge with its limited connector real estate. Figure 1 shows the current 3U connector set. These samples are being used at General Dynamics Canada to create a set of product offerings for VITA 62.

The VITA 62 pins' definitions provide a good summary as to the design decisions that have been made in the selection of pins for VPX systems:

```
-DC_IN/CAN., +DC_IN/ACL CHASSIS,UD1,UD2,
UD3,UD4,VBAT,FAIL*,INHIBIT*,ENABLE*,UD0,
+12V_AUX,NED,NED_RETURN,3.3V_AUX,

GA0*,GA1*,SM0,SM1,SM2,SM3,-12V_
AUX,SYSRESET*,VS1_SHARE,VS2_SHARE,VS3_
SHARE,SIGNAL_RETURN,VS1_SENSE,VS2_SENSE

VS3_SENSE, SENSE_RETURN, VS3, POWER_
RETURN, POWER_RETURN, VS2, VS1
```

The VITA 62 committee decided to leverage the VITA 46 P0 utility connector pins to support requirements such as geographical addressing, system management, and other control functions while supporting maximum power generation. The connector supports both DC and AC input voltages, with three-phase AC supported only on the 6U connectors. The generation of a maximum

Figure 1 | The current 3U connector set for VITA 62

output power is limited by the system architecture, and the power demands on the VS1/2/3 voltage rails. A maximum current of up to 50 A per output rail supports power outputs as high as 1,000 W for 3U and 2,000 W for 6U. When the shared pins are utilized with multiple supplies, there are no real limitations on achievable power levels.

What about the other OpenVPX issues?

Besides defining effective remedies to OpenVPX system power issues, VITA 62 also defines limits to myriad system-level issues such as the following.

Inrush current: This often occurs on initial system powering, where all modules attempt to fully power the components on the boards. This creates an initial current surge on the power supply, which, if not limited, could result in output rail voltages dropping below required levels. VITA 62 compliant systems will provide VPX module designers tolerances on the amount of current that the module can draw in excess of their steady state requirements. In so doing, the architect of the VITA 62 compliant system will be able to correctly size the power requirements needed to be supported by the VITA 62 power module(s). This will eliminate potential system issues where VPX modules cannot power up completely or are in an unknown state due to lack of current or undervoltage during initialization. This can lead to only certain board portions or certain components being properly initialized. These are challenging issues that can be difficult to detect and resolve.

Output control: The *ENABLE and *INHIBIT are input pins that control the output power generation. This allows systems to support the use of an external on/off switch and the disabling of output rails.

EMI: Pertaining to EMI, VITA 62 provides a set of guidelines for standards like MIL-STD-461F and MIL-STD-1275D, allowing board designers to specify their compliance to these specifications and how they have been met. Power supply vendors are free to provide onboard emission control, or to define requirements for external components.

Synchronization: Synchronizing power supplies is a recently discovered requirement for high-speed data acquisition systems. These systems typically need to synchronize data acquisition sampling times with power-supply switching times. VITA 62 allocates pins to allow the routing of synchronization signals to modules that require this capability.

Nuclear detection: The system also defines pins for use in nuclear detection. These systems typically need to crowbar output rails on detections of an event. VITA 62 allows for both internal and external use of detection devices to support this capability.

Putting it all together

The existing VITA 46 key pins and keying definitions are used for VITA 62. This provides a means to identify up to five unique input powers. These keys will protect against the accidental plugging in of an incompatible power supply card. The 3U connectors provide two keying pins, and the 6U has a set of three keying pins. A single key on each connector defines the input voltage setting of the power supply. The keys and connector are designed to accommodate the use of liquid-cooling connectors as replacements for standard keys. The middle key is used on a 6U connector as the key for input power. This will accommodate the use of liquid-cooling quick connects as replacements for the outer keys' pins. Figure 2 shows the final configuration of a 6U conduction-cooled module. The module is fully compliant with the VITA 48 mechanical profile.

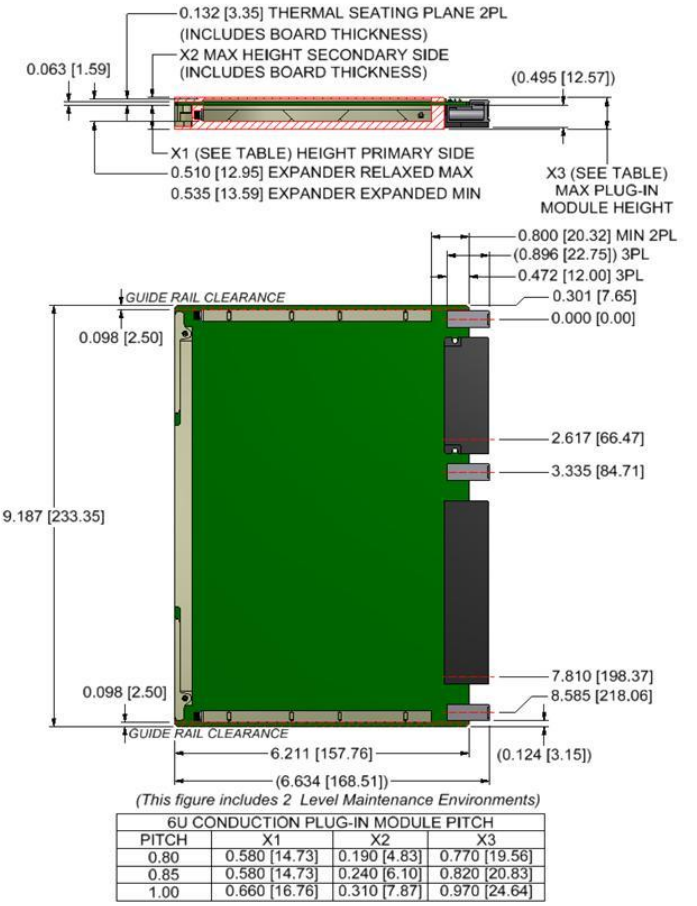

6U CONDUCTION PLUG-IN MODULE PITCH			
PITCH	X1	X2	X3
0.80	0.580 [14.73]	0.190 [4.83]	0.770 [19.56]
0.85	0.580 [14.73]	0.240 [6.10]	0.820 [20.83]
1.00	0.660 [16.76]	0.310 [7.87]	0.970 [24.64]

Figure 2 | A 6U conduction-cooled VITA 62 module, which is fully compliant with the VITA 48 mechanical profile.

Since all VITA 62 slots are connected in parallel, this allows secondary slots to add key system features without affecting primary slot functionality. A module designer could provide extra capability on a utility VITA 62 module, which, upon detection of loss of primary power, would "hold up" the inputs to the VITA 62 power supply for up to 50 ms or more. The VBAT pin can be used to maintain a low-power trickle charge to modules in the systems when main power is no longer available, thus allowing CMOS settings and "time of day" to remain active. The connection to the system management infrastructure provides the ability for a 46.11 shelf manager to reside in the VITA 62 utility slot.

What all this means

The diverse nature of the VITA 62 slot in system backplanes naturally lends itself to supporting important features that traditionally have been configured in a more ad hoc manner. And, VITA 62, once completed, will be integrated into the current OpenVPX standards by adding profiles for slots, modules, and backplanes, thus allowing system architects to fully design new VPX systems using a complete set of COTS modules. VITA 62 requires that the complete set of capabilities supported on the module be published by the vendors, so that system architect can then choose the module which best meets their needs. VITA 62 also solves the issues of managing inrush current, output control, EMI, synchronization, and nuclear detection. CS

Patrick Shaw is the chair of VITA 62 and participates in many of the other VITA committees. He is the chief BIOS engineer at General Dynamics Canada and acts as a system architect for single board computer products. He can be reached at Pat.Shaw@gdcanada.com.

General Dynamics Canada
613-596-7000
www.gdcanada.com

Outpacing VME: OpenVPX fast-tracks technologies to the front lines

By Brian Roberts

VPX is replacing VME. OpenVPX is that part of VPX that facilitates multivendor COTS system-level interoperability including modules, backplanes, and development chassis. Since embedded defense systems must often be developed and deployed rapidly, it is very important that potential conflicts are eliminated as quickly as possible during the design phase.

VME has long been used in mission-critical embedded systems. However, its VPX (VITA 46) progeny surpasses VME in bandwidth and power. But VPX designers soon found that VPX yielded its share of system-level interoperability issues: Multiple vendors' boards and components were often incompatible within the larger system design, presenting a major obstacle to VPX system development. The answer: OpenVPX (VITA 65), a leading system architecture standard in field-deployed, mission-critical embedded systems.

OpenVPX's goal is to provide a framework for system-level interoperability for COTS-based VPX Line Replaceable Units (LRUs). LRUs' main plus is that they are Two-Level Maintenance (2LM) compatible. The ultimate goal of two-level maintenance in all complex defense systems, including OpenVPX, is to enable a field technician without special maintenance tools to quickly identify a faulty board and swap it in a harsh and hostile environment – and therefore quickly return the system to operation.

The need for speed is also relevant to the OpenVPX design phase. Since embedded defense systems must often be developed and deployed rapidly, it is very important that potential conflicts between the various profiles are eliminated as quickly as possible. SWaP and thermal management, along with other requirements, must be resolved faster so that advanced capabilities can be swiftly engaged.

Accordingly, OpenVPX's proven field-deployed readiness and advanced systems design flexibility hasten the delivery of next-generation customizable COTS solutions to the front lines, today and tomorrow – and at a rapid pace. The following discussion examines VME and OpenVPX pertaining to:

1. SWaP and thermal management
2. Technical considerations for connectors and backplanes

SWaP and thermal management

Current and future embedded processors are pushing and exceeding the once prime-time VME64x 35 watts/slot at +12 V limit. Processor core densities and the number of cores on a chip are increasing the input power required by processor boards. Even as the power efficiency of processor technology increases, demand for increased operation rates and complexity will probably take up the slack in order to more rapidly decrease OpenVPX-based technologies to the warfighter. A look at SWaP and thermal management considerations will lend insight.

Currently, OpenVPX uses 3U (100 mm x 160 mm) and 6U (233 mm x 160 mm) form factors. Since more applications are moving to the lighter, smaller 3U form factor, heat management has become more important than ever, which, of

course, ties into power considerations. SWaP is not specified by VME or VPX, but OpenVPX is defined to supply 120 W at 5 V and 383 W at 12 V. The challenge is that this increased power creates heat management issues.

While OpenVPX does not consider specific temperature range or heat-flow-rate cooling strategies, the dimensions of board features such as side rail wedge mounts to facilitate conduction cooling are specified. Rugged and serviceable OpenVPX COTS is standardized in VITA 46 and VITA 48 (VPX-REDI). Dimensional specifications for cooling are given in VITA 48.1 for convection, VITA 48.2 for conduction, as well as VITA 48.3 for liquid. Covering all cooling methods in one standard enables board manufacturers to offer each board design with various cooling options.

Although all forms of heat management are supported by OpenVPX, conduction cooling is becoming the method of choice for ruggedized military and aerospace applications. Air is often contaminated or unavailable, and convection is often inadequate. Forced-air and liquid cooling require fans and pumps. Mechanical devices fail quicker than electronics and can severely reduce MTBF below 200,000 hours, which is often the minimum required for mission-critical systems. The development of heat pipes that can carry more than 150 times the heat flow per volume than aluminum is increasing the popularity of conduction-cooled OpenVPX thermal management systems.

Making the connection

OpenVPX connectors are a significant advance over VME alternatives. They provide greater bandwidth, inlet power, number of pins, and serviceability (Table 1). OpenVPX connectors also facilitate system-level development of VPX LRUs, to speed deployment and help maintain OpenVPX technologies' field readiness. The MultiGig RT family of connectors, for example, provides much higher bandwidth, pin density, and mission-critical ruggedness than VME64x's DIN 41612 connectors.

Specifically, 3U OpenVPX boards have three MultiGig RT2 connectors (P0 – P2):

- P0 has 56 pins for power, addressing, system management, and other utility signals.
- P1/P2 have 112 pins each that can be used as differential pairs or as single-ended conductors.

The 6U boards have seven connectors. P3 through P6 can be used for differential pairs or single-ended signals, with one row of P3 through P6 reserved for single-ended signals. Other signals requiring coax or fiber optics can be connected through P5 or P6. Currently, coax and fiber optics for 3U boards are not prescribed by OpenVPX. However, they can be custom connected from the front panel.

The MultiGig RT connector's ruggedness facilitates military 2LM. It permits boards to be removed and inserted by reasonable manual force from front handles. The MultiGig RT connector has been tested for thermal and mechanical shock, vibration, and contaminants. Through severe testing, pin contact resistance remains low. The MultiGig RT connector has electrically grounded blades adjacent to sensitive recessed pins, making it virtually impossi-

	VME64x	OpenVPX (VITA 65)
Bus characteristics		
Common bus used	VMEbus	PCI Express (PCIe)
Typical bus bandwidth	Supports 80 MBps for standard 64-bit VME, up to 320 MBps for the 2eSST standard	PCIe v2.x = 500 MBps per lane, 16 lanes (PCIe v2.x = 8 GBps)
Bus minimum pin count	106 (approx.)	4 = 1 lane
Minimum bus width	64 bit	1 lanes (1 TX diff. pair + 1 RX diff. pair)
Standard max bus width	64 bit	16 lanes (PCIe v2.x = 8 GBps)
Bus Size Variation	64 bit	1,2,4,8,16 lane variations
Additional commonly available bus interface	100 Mb/1 GbE	Serial RapidIO/10 GbE
Bus links configurable by FMM	No	Yes
User-defined connector features		
Common user I/O pin counts	205	6U = 360 , 3U = 108 (varies by module profile type - important to review card selected)
Pin arrangement in connector	Low-speed pin construction with minimal GND shielding	Integrated controlled impedance differential pair conductors with GND shielding between each pair for improved signal integrity
User I/O variation above standard pin count	No	Might increase depending on bus width implementation
Different module interface and user I/O "profiles"	No	Yes - Important to evaluate profile compatibility
Connector variations	No	Differential pair or single-ended pin arrangements
I/O routable by FMM	No	Yes
Other feature sets		
Thermal models	Air/conduction cooled	Air/conduction/liquid cooled
Standard form factors	6U	6U and 3U
PMC/XMC capable	Yes (XMC support varies)	Yes
Standard voltages	+3.3V, +5V,+12V,-12V,Optional (-V1,+V1,-V2,+V2)	3U = (VS2)+3.3V,(VS3) +5V,(VS1) +12V, Aux+12V,Aux-12V, Aux+3.3V
		6U = (VS2)+12V,(VS3) +5V,(VS1) +12V, Aux+12V,Aux-12V, Aux+3.3V

Table 1 | A comparison of VME64x and OpenVPX (VITA 65)

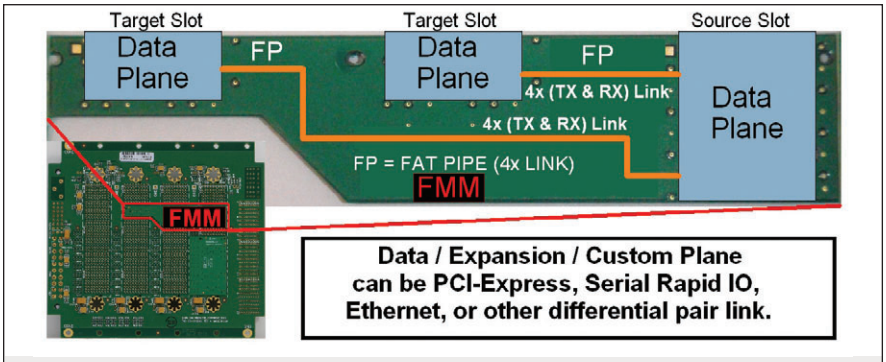

Figure 1 | Expansion diagram of FMM micro overlay with two fat pipe PCI Express inks

ble to touch a pin without being grounded. This eliminates the possibility of electro-static damage through the connector. The covers on both sides protect the board from electrostatic damage and debris.

Backplane to the future

If the backplane can be chosen quickly and well, OpenVPX systems can be deployed faster. And historically, high-speed switched serial fabric presented the most compelling demand for a new backplane architecture. The bandwidth of the OpenVPX backplane is greatly improved over the VMEbus 40-60 MBps range. The OpenVPX backplane can be configured into many network topologies such as mesh, star, dual-star, ring, or daisy chain. These networks permit multiple signals to be routed such that several cards can talk to each other simultaneously, achieving an aggregate bandwidth well over 100 GBps.

The OpenVPX backplane profile supports 10 GbE, PCI Express, Serial RapidIO, and SATA for nonvolatile memory. New revisions of these standards will push the limits of differential copper pairs; OpenVPX provides for coax and fiber-optic connections for higher-speed data and other signal formats.

OpenVPX COTS backplanes can also provide more options if problems arise later in the design or manufacturing phase. For example, a multi-Gigahertz tested Fabric Mapping Module (FMM) can connect to the backplane and modify its topology, making it more flexible. It provides quick-turn backplane customization, thus eliminating interconnect conflicts. Dawn VME Products' FMMs use BGA (Ball Grid Array) substrate technology as a "micro overlay," similar in function to overlays used on VME backplanes except that FMMs can be used on VPX backplanes and are tuned for high-speed signal transmission (Figure 1). They can:

- Directly connect PCI Express or SerialRapidIO to multiple cards or cards and switches
- Link SATA from a CPU card to a Solid State Drive (SSD) carrier
- Enable XMC cards to talk to other XMC cards

Dawn engineers have successfully used Fabric Mapping Modules to solve many OpenVPX application problems in the design phase.

Designing with OpenVPX for rapid deployment

At one time, VME was found in military systems aplenty. But modern weapon and other embedded defense systems are becoming more modular and much simpler and faster to service using LRUs. Thus VME is being outpaced. OpenVPX, on the other hand, facilitates VPX LRU development by providing system-level, multivendor interoperability.

Since these embedded defense systems and VPX LRUs must often be developed and deployed rapidly, it is very important that potential conflicts between various profiles are eliminated as quickly as possible during the OpenVPX design phase. SWaP and thermal management, in addition to connector and backplane technical considerations, must be addressed and resolved faster so that advanced capabilities can be swiftly engaged to the warfighter. CS

__Brian Roberts__ is Senior Designer at Dawn VME Products. He has 21 years of design experience developing electronic products with Commercial and Military companies based in Silicon Valley. He can be contacted at robertsb@dawnvme.com.

VITA 66 expands VPX with fiber-optic connectivity

By Gregory Powers

As rugged embedded computing systems strive for higher operating speeds and densities, adding fiber-optic connectivity to the proliferating VITA 46 (VPX) market space brings benefits in higher bandwidth, longer transmission distances, and lower weight. VITA 66 defines fiber-optic connector modules for VPX using well-established and proven optical termini. Each style of termini offers different benefits in terms of density, ruggedness, repairability, and other characteristics.

The thirst for higher bandwidth and faster processing speeds in military and aerospace applications seems unquenchable. VITA 46 (VPX), with a backplane connector system supporting 6.25 Gbps in a switched-fabric architecture, is the latest generation of VMEbus and offers new levels of performance for embedded computer systems. VPX systems are designed for the flexible application of demanding high-speed protocols – such as 10 GbE RapidIO, InfiniBand, and HyperTransport – in ground, aerospace, and marine applications. While the basic architecture of VPX is defined in VITA 46 and 65 (OpenVPX), ongoing efforts by the VITA Standards Organization (VSO) to enhance VPX capabilities include VITA 66 – "Optical Interconnect on VPX" – which adds fiber-optic interfaces to the connector platform.

The issue at hand is one of elegant, high-volume data transport. As operating speeds increase, copper cables become increasingly limited in transmission distance, besides becoming heavier and more expensive to deploy. For example, Category 6 twisted-pair cable runs 100 m at 1 Gbps. As the I/O speed increases to 10 Gbps, the recommended cabling distance reduces to 55 m in a benign electrical environment. Where high levels of alien crosstalk may exist – such as with high densities of closely packed cables – the distance is reduced to 37 m. Military and aerospace applications also can involve substantial levels of EMI, elevating the need for cable shielding, which increases the size, weight, and complexity of the interconnection system.

Fiber-based I/O does not offer the same trade-offs in bandwidth and distance. Among the well-known advantages of fiber-optic cables are: lighter weight, higher bandwidth, and longer transmission distances. For instance, with single-mode fibers multikilometer transmission distances at multigigabit speeds are possible. In addition, fiber-optic cables are inherently EMI resistant and require no electromagnetic shielding.

The installed advantages are clear, too. In cutting-edge composite vehicles, fiber-optic lines minimize shielding and bonding challenges. Problematic installation of heavy metallic components is eliminated, saving weight and time and reducing risk. Vehicle designers achieve location-independent architecture. Boxes that are meters or kilometers apart communicate as though colocated. VITA 66, by defining optical I/O capabilities, additionally provides improved density, ruggedness, and repairability via three termini styles.

Flexible options in optical connections

The VITA 66.0 base specification defines the common mounting interface requirements for the various fiber-optic interconnects within 3U and 6U VPX applications. This includes definition of the mounting provisions, permitted locations, and range of engagement. Additional "dot" specifications define requirements for the specific optic module interfaces. A fundamental goal of VITA 66 is to offer designers multiple existing and fielded mil/aero termini technologies, allowing them to quickly and confidently implement the best solution for specific applications. The three module varieties are based upon proven optical termini:

- MT ferrule (VITA 66.1)
- ARINC 801 termini (VITA 66.2)
- Expanded Beam (EB) insert (VITA 66.3)

Each style of termini offers different benefits in terms of density, ruggedness, repairability, and other characteristics. In addition, the modules are designed to meet the requirements of VITA 47 (Environments, Design and Construction, Safety, and Quality for Plug-In Units), which covers environmental and mechanical ruggedness for VPX systems, including temperature cycling, vibration, shock, altitude, and more.

MT ferrule for high fiber counts

VITA 66.1 modules use the MT ferrule, configured to enable up to 24 optical fibers per the standard. With two ferrules per module, this equates to 48 fibers in a 3U system and up to 240 fibers in a 6U system. Of all industry-standard optical connectors, the MT ferrule provides the highest-density interconnections for both multimode and single-mode fibers in a ferrule with an end-face only 6.4 mm by 2.4 mm. The MT, with its history in rugged box-level applications, is ideal for switches and concentrators.

Application considerations for the MT ferrule are the inability to perform field terminations and to replace individual fibers. In addition, the MT is a physical contact style termini, meaning the glass end-faces are in direct contact. While this style of termini initially can provide very low loss, polish degradation via end-face abrasion is possible.

The interconnect for MT ferrules is the first "dot" specification to be published, and will be VITA 66.1. The intention is to have it serve as a template, allowing the additional "dot" specifications to be created and ratified quickly. It is anticipated that VITA 66.0 and 66.1 will be submitted to ANSI in mid-2011, with VITA 66.2 and 66.3 following later in the year.

ARINC 801 termini for highest optical performance

ARINC 801 termini, as used in VITA 66.2 modules, are based on industry-standard 1.25 mm ceramic ferrules, bringing all the advantages of discrete ceramic ferrule connectivity. The ceramic ferrules offer high-performance features that include physical contact technology for very low insertion loss, angled polishes for minimal reflection loss, and keyed

orientation for optimal single-mode performance.

In contrast to the MT ferrule, ARINC 801 termini allow individual optical fibers to be installed or removed. The contacts are a standard technology within commercial aviation, but they are also finding wider use in military/aerospace, marine, and ground applications. Like MT ferrules, the ARINC 801 termini are a physical contact technology and may experience higher loss over time due to end-face abrasion.

Expanded Beam (EB) termini for extreme environments

EB connectors are widely known for being extremely robust.

The EB insert, defined in VITA 66.3 and based on ball lens technology in MIL-DTL-83526/20 and 21, is the most rugged interface and supports up to four fibers per module. The ball lens technology expands the effective diameter of the channeled light and collimates the beam, projecting it across an air gap to the receiving ball lens, which reverses process and reintroduces the light to the awaiting fiber core. Where a speck of dust might completely block a 50 micron or smaller fiber core, the expanded beam minimizes the impact of dirt and debris on insertion loss. A small dirt particle on the glass face of the expanded beam connector does not dramatically block the light as it can on the end-face of a typical ceramic-based ferrule.

The non-contacting EB interface eliminates potential abrasion induced by shock, vibration, or repeated mating/unmating. The fiber ends are basically protected behind safety glass, allowing very easy cleaning and inspection. The EB interface makes them suited to two-level maintenance or applications calling for frequent insertion/extraction, such as a secure storage device requiring removal after each flight operation.

Performance characteristics of VITA 66 optical termini

		MT ferrule	ARINC 801	Expanded beam
Standard		VITA 66.1	VITA 66.2	VITA 66.3
Fibers/Module		48	4	4
Fibers/Connector	3U	48	4	4
	6U	240	20	20
Durability (Mating cycles)		100*	100	>3,000
Ruggedness		Low	Moderate	High
Individual fiber repair/Replacement		No	Yes	Yes
Contacting Interface		Yes	Yes	No
Weight		Low	Moderate	High
Fiber type		MM and SM	MM and SM	MM and SM
Insertion Loss (dB), typ	SM	0.25	0.15	0.8
	MM	0.2	0.1	0.7
Return loss (dB)	SM	>55	>50	>34
	MM	>20	>20	14

Table 1 | Performance characteristics of VITA 66 optical termini. (*proposed)

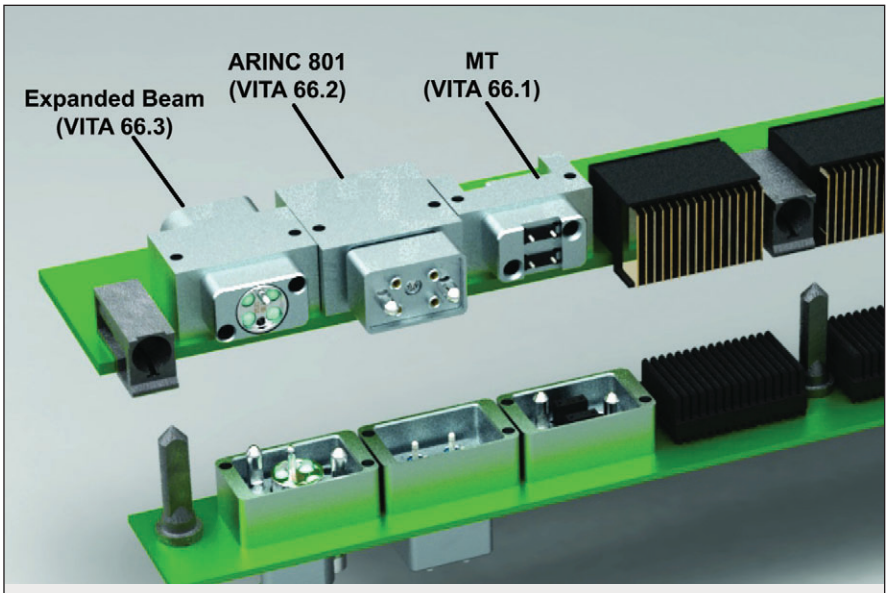

Figure 1 | VITA 66 modules applied

Insertion losses—the loss of optical power as light passes through the connectors—for EB connectors are higher initially. But unlike MT and ARINC 801 termini, the losses remain constant over the life of the connector without degradation because its non-contacting optic interface does not become damaged. The EB interface is a better solution for applications with higher levels of vibration, dirt, and mating cycles. Where multiple mating cycles are expected, it only takes one mating of an MT or 801 terminus without adequate cleaning to permanently damage the fiber endface. Table 1 summarizes the performance characteristics of the different optical styles.

While VITA 46.0 defines P5/J5 and P6/J6 as user defined locations in the 6U implementation, and P2/J2 in the 3U, the VITA 66 fiber-optic modules have been designed for installation in P2/J2 through P6/J6. The integrator could even mix the ferrule technologies along the same card edge if the situation warrants doing so. Figure 1 illustrates the three module types on a notional 6U VPX card edge.

VITA 66: Flexibility in VPX optical interconnects

VITA 66 is an important addition to the VPX family and further empowers the rugged embedded computer architect. It supports a wide range of fiber-optic interconnection needs, from intrasystem board-to-board links to intersystem networks spanning many kilometers. Because application requirements can vary greatly between airborne, shipboard, and ground defense needs, the standard offers a range of high-integrity choices. By using well-established termini, VITA 66 will also speed adoption of optical interconnections in VPX systems. Distance, weight, EMI, and bandwidth limitations are no longer insurmountable challenges for data transport in rugged embedded computing. CS

Gregory Powers serves as Market Development Manager for the Electronic Systems and Space segments within the Global Aerospace, Defense & Marine business unit of Tyco Electronics. Mr. Powers received a BMSE from Syracuse University, has completed numerous graduate-level studies, and holds two patents relative to optic datacom devices. He can be contacted at gregory.powers@te.com.

Tyco Electronics
425-413-1910
www.te.com

Latest Intel processors advance embedded DSP and SBC system design

By Ian Stalker and Alan Baldus

Intel's latest generation of Core i7 processors is a game-changer for embedded military DSP and SBC system designs. For the first time, they bring support for Serial RapidIO, the OpenVPX (VITA 65) board interconnect fabric of choice. Even better, the performance of vector math processing, critical for DSP applications, is effectively doubled with Intel's new 256-bit AVX instruction set.

The key performance metric of interest for military DSP systems is the speed of performing floating arithmetic operations, which is referred to typically as *GFLOPS* when discussing the speed of computers. In recent history, these DSP systems were commonly built using Texas Instruments 320C40 and 320C6701k and Analog Devices SHARC dedicated DSP processors, which were themselves followed by a number of generations of PowerPC/Power Architecture processors with AltiVec. All of these processors offered good floating point performance per watt and all were available from vendors with a history and track record of support for military embedded customers. Now, with the introduction of Intel's 2nd Generation Core i7-2715QE quad-core processor, the design of x86-based embedded military DSP systems and high-performance SBCs takes a significant leap forward.

Intel refers to their product introduction cadence as the "Tick-Tock" model. A "tick" is when Intel delivers new silicon process technology with increased transistor density, and enhanced performance and energy efficiency within a smaller version of an existing microarchitecture. The 2nd Generation Intel Core i7 is a "tock," which is when an entirely new microarchitecture is introduced on an existing semiconductor process technolgoy. Using the 32 nm process introduced with the Westmere generation, the 2nd Generation Core i7 (previously code-named "Sandy Bridge") features many architectural improvements (especially in the cache subsystem) that lead to improved performance per clock cycle. It is the nature of microprocessor design that revised architectures typically provide incremental performance improvements. However, the 2nd Generation Core i7 has delivered a major leap forward in the signal processing capability of the processor, thanks to the new 256-bit wide Intel Advanced Vector Extensions (AVX) floating-point instruction set, which supercedes the earlier 128-bit Streaming SIMD Extensions (SSE) instructions.

While the new Core i7 brings many advantages for DSP system designs, SBCs used in conjunction with Core i7-based DSP engines also benefit. SBCs can now take advantage of the first ever support for Serial RapidIO on Intel Architecture, as the result of an upcoming PCIe2-to-Serial RapidIO2 bridge chip from IDT that will provide a common communications path and improve interoperability in a complete system. The new Intel processor also supports 16 lanes of Gen2 PCIe for full-bandwidth communications across high-performance processor cards. Intel's hyperthreading technology provides for running two execution threads on each core, enabling greater utilization of the execution units and providing improved power efficiency. Published reports show performance increases of 7 to 34 percent due to hyperthreading alone.

The AVX 256-bit difference

Prior to the introduction of Intel's new 256-bit AVX, developers of military DSP systems typically turned to 128-bit AltiVec-enabled CPUs for vectorized signal processing functions. In the past few years, development of new AltiVec-enabled processors slowed significantly, leaving DSP system developers with limited options. In the meantime, Intel continued to invest in and enhance its own high-performance vectorized processing solution with continual enhancements to Intel Streaming SIMD Extensions, a 128-bit wide processing unit predecessor to AVX, capable of simultaneously operating on four 32-bit floating point values. Intel SSE also featured support for double-precision floating point, a feature not available in AltiVec. In Intel's earlier multicore processors, each core was provided with its own SSE unit, so raw floating-point performance scaled with the number of cores. In the new Core i7 Intel has upgraded SSE with AVX, doubling the size to 256-bits wide.

This doubled vector processing performance is a significant milestone in DSP system design. DSP algorithms used in critical military applications such as radar, SIGINT, and image processing depend on the precision achieved with floating point numbers combined with the speed of processing. The new Core i7 doubles the peak performance of SSE. When compared to SSE in actual FFT kernels, AVX has been benchmarked up to 1.8x faster than SSE (Figure 1). The AVX instruction set was designed to support future extensions, which hints at wider implementations in the future.

Serial RapidIO onboard

Serial RapidIO is the preferred fabric for the types of processor-to-processor communications required by demanding military DSP systems. This is because of Serial RapidIO's reliable packet transmission and ability to deliver low and predictable latencies. These benefits of RapidIO messaging are ideal for large peer-to-peer clusters of processors typically used in complex signal processing applications. With the Intel 2nd Generation Core i7, Serial RapidIO is supported on Intel architecture-based OpenVPX/VITA 65 embedded boards for the first time with an easy, cost-effective interconnect provided by IDT's upcoming PCI Express (PCIe) Gen2-to-Serial RapidIO protocol conversion bridging semiconductor product.

Before this newest generation of Core i7, the lack of support for Serial RapidIO for Intel platforms severely limited the viability of using Intel architecture in DSP multiprocessor system designs. Solutions for Intel have included support for fabrics such as InfiniBand and Gigabit/10 Gigabit Ethernet, which are not embraced in military applications because of their non-industrial temperature silicon and relatively high power consumption. For SBCs, where the requirement is typically a single processor communicating with I/O, these fabrics have been sufficient,

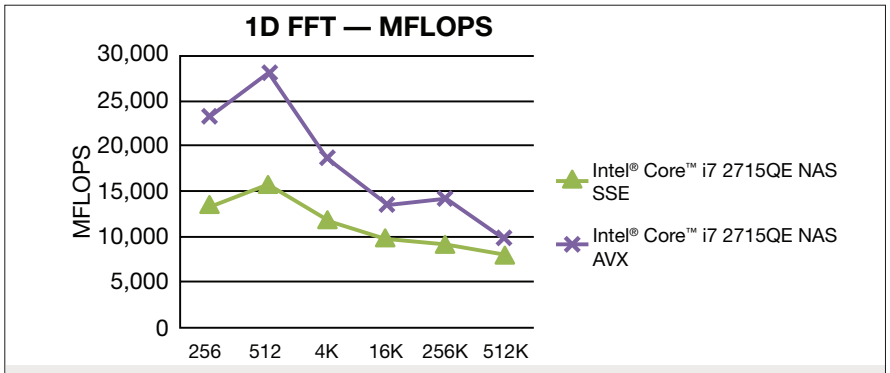

Figure 1 | A comparison of complex FFT performance using 128-bit SSE versus 256-bit AVX instructions – The code was run on the same processor and clock rate. Data was provided by Intel.

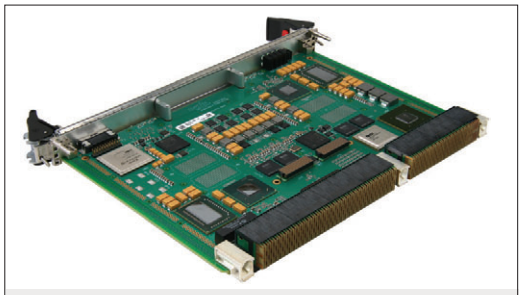

Figure 2 | The Curtiss-Wright CHAMP-AV8 VPX board with 269 GFLOPS and featuring dual Intel Core i7 2715QE CPUs with four Gen2 Serial RapidIO ports and 80 Gbps to the backplane.

but would-be Intel-based DSP military designers were deprived of the option to design systems around Serial RapidIO, the multiprocessor fabric of choice.

IDT's PCIe-to-Serial RapidIO bridge and new Gen2 Serial RapidIO switches will enable system designers to build Intel architecture-based processing engines with much more fabric bandwidth than that offered by any other currently available technology. The upcoming IDT bridge product supports 5 Gbps interfaces on both PCIe and Serial RapidIO ports. With the advantage of small size and low power consumption, system designers can add bandwidth by using multiple PCIe2-to-Serial RapidIO2 bridges connected directly to the processors or via a PCIe switch. This performance can scale at the system level with the new Gen2 Serial RapidIO: This new generation of systems will deliver double the backplane bandwidth provided by the already fast 3.125 Gbps Gen1 Serial RapidIO technology. A 19" rack, OpenVPX processing system will be able to deploy 1.2 terabits per second of fabric bandwidth. The Intel/Serial RapidIO combination is also suited for SwaP-constrained systems, as designers can maximize the power available for actual computing knowing that Serial RapidIO fabric technology provides the best bandwidth/watt.

Serial RapidIO bridges implemented in FPGAs don't support high-performance messaging, a feature which directly maps to higher-level software APIs such as MPI. IDT's new bridge product will support the two main Serial RapidIO transfer modes, Serial RapidIO messaging, and memory-mapped transfers. Another benefit of the IDT silicon is the inclusion of DMA engines that speed computation while offloading the host processor. Intel processors typically don't have DMA engines on-chip, but depend instead on the peripheral chip to move data. Without a DMA engine, moving data can require a large amount of the host processor's attention, with the result that a multicore processor might have one of its cores (and associated power) largely consumed by moving data, which is all the more burdensome because it has to be done in code.

Another advantage of Serial RapidIO for SWaP-constrained military systems is its ability to support distributed switch and centralized switch architectures. Distributed switch systems (an example is the VITA 65 BPK6-CEN05-11.2.5-n backplane profile) can make use of the local Serial RapidIO switch and thus avoid the need for a separate switch card. For example, if the system were using a ½ ATR Short enclosure (four 1" slots), this capability saves 25 percent of the space and a considerable amount of power. For large systems, centralized switch architectures are often preferred, and Serial RapidIO is equally adept at this approach.

An example of a high-performance DSP engine designed to take full advantage of the latest offering for Intel's Core i7 is the new CHAMP-AV8 from Curtiss-Wright Controls Embedded Computing (Figure 2). The CHAMP-AV8 is an Intel Core i7-2715QE-based rugged, high-performance OpenVPX DSP engine. Performance of this dual Core i7 DSP engine is rated at up to 269 GFLOPS. It also supports the IDT Gen2 PCIe-to-Serial RapidIO bridge product, effectively tripling the bandwidth of first-generation VPX products with up to 240 Gbps of fabric performance. **CS**

Ian Stalker *is the DSP product manager for Curtiss-Wright Controls Embedded Computing. He has more than 20 years of experience in the embedded industry and holds a degree in Electronics Engineering. Contact him at ian.stalker@curtisswright.com.*

Alan Baldus, *Intel SBC Product Marketing Manager at Curtiss-Wright Embedded Controls, has more than 15 years in the embedded computer industry. He holds a BS in Electrical Engineering. Contact him at abaldus@curtisswright.com.*

Curtiss-Wright Controls Embedded Computing
703-779-7800
www.cwcembedded.com

SBC with lots o' I/O

Paper or plastic? Fries or chips? The options are sometimes overwhelming, as they might also seem to be with this single-slot 64EP3 SBC from North Atlantic Industries that offers plug-in modules for the processor subsystem as well as various I/O options. But when all stacked up with four modules, this PowerPC or Blackfin SBC replaces multiple boards with just one. I/O options include typical defense and automotive analog and digital sensor types, as well as communications schemes like MIL-STD-1553 and GbE. Interestingly, the board comes with a dedicated FPGA and function libraries designed to do pre- and post-data processing, depending upon the I/O subsystem.

The CPU can be either an ADI Blackfin 500 MHz BF-533 DSP processor with Visual DSP++ 4.5 libraries and emulator, or a Freescale MPC8536 PowerQUICC III running at 1.25 GHz. Each CPU subsystem is equipped with flash, an FPGA, and direct I/O to front panel or rear panel I/O when ordered as a conduction-cooled LRU. The PowerPC has DDR2, while the DSP has SRAM. These CPU modules then order around the many I/O choices, which include ARINC 429/575, RS-232/422/485, or CANbus. There's also myriad other less sexy, though no less important, I/O: A/D, D/A, discrete/TTL/CMOS/differential; RTD, synchro/resolver, LVDT/RVDT, or an encoder. With four mix-and-match slots, the 64EP3 offers lots o' I/O in not a lot of space.

North Atlantic Industries • www.naii.com
www.vmecritical.com/p47552

Two cores ... with more

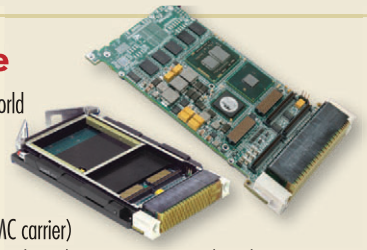

As Intel's Core family of CPUs takes the civilian world by storm, the military is closer behind than ever before. Gone are the days of waiting two or more years for the rugged, embedded version. Take Aitech's C870 (SBC) and CM870 (PMC/XMC carrier) boards, for instance. These single-slot, 3U VPX boards pack more into two 3U boards than was previously available from two 6U boards. The SBC is equipped with Intel's Nehalem-based Core i7 dual-core CPU running at 2.53 GHz for performance, all the way down to 1.33 GHz for "ultra-low power." Since Aitech knows rugged (sounds like a slogan, doesn't it?), the 4 GB of DDR3 SDRAM with ECC is soldered to the board. There's also 4 MB of flash BIOS and 8 GB of SSD storage.

I/O on the C870 CPU includes 2 1000BASE-T Ethernet and 2 1000-BASE-BX/KX (backplane Ethernet), 2 SATA II ports, 4 USB 2.0, and 8 discrete I/O. There's also an Intel graphics controller onboard, so HDMI/DVI and CRI outputs are also present. Operating systems include Windows XP and Embedded, plus Linux and VxWorks. The CM870 carrier card is designed for PMC and XMC boards and supports PCIe x4, 64-bit 66/133 PCI-X, and plain old PCI. For XMC boards, there's PCIe x8. All of this routing is compliant with VITA 46.3, 46.4, and 46.9 for various VPX configurations. And as we said earlier, since Aitech knows rugged, both boards are available in air- and conduction-cooled flavors. Amazingly, both cards weigh less than 0.7 pounds each in conduction-cooled armor.

Aitech Defense Systems • www.rugged.com (Say, maybe this really *is* their slogan!)
www.vmecritical.com/p47553

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

GE
Intelligent Platforms

We earn your trust, one platform at a time

Discover why GE Intelligent Platforms is one of the most trusted suppliers of military and aerospace COTS solutions

Program managers around the world consistently put their trust in GE Intelligent Platforms for COTS computing and communications products for ground-up developments as well as technology insertions into existing platforms. For a hundred years, GE has been supplying the defense industry with innovative products. GE Intelligent Platforms carries forward that legacy with embedded computers that reduce costs, leverage the latest commercial technologies, slash time-to-market, and reduce engineering risk.

Numerous high-visibility programs for manned and unmanned ground, air, and sea platforms have incorporated GE Intelligent Platforms products to achieve targeted hardware and software benchmarks with a minimum of program risk. Let us help you achieve that same level of success with your next program.

To learn more, visit:
defense.ge-ip.com

For more information about QRcodes, visit www.ge-ip.com/qrcodes

imagination at work

VPX, VME & cPCI True Military ATR Enclosures

16HP - Embedded Heat Pipe transfer system

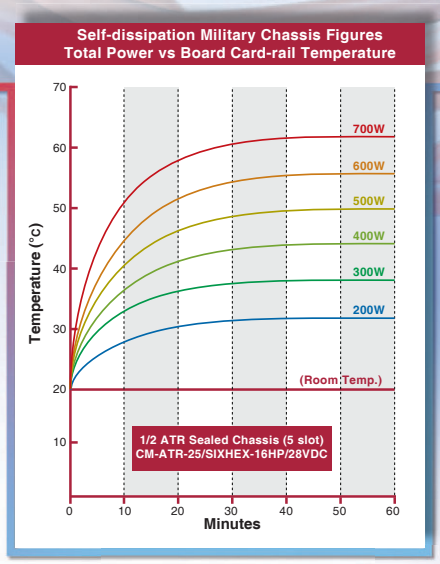

The perfect Sealed COTS solution

We have incorporated 16 Heat Pipes within our latest Six Heat Exchanger chassis series in order to improve thermal performance a step further. The new SixHex-16HP ATR range will increase your system MTBF by 40% when compared to our outstanding SixHex series as a result of the additional reduction in payload operating temperatures.

This achievement surpassed our design expectations. Customers will appreciate the superb quality and performance delivered in such compact enclosures. Today, even our smallest 9 kg ½ ATR chassis is capable of handling up to 700 Watts. This COTS industry record establishes an unbeatable sealed platform for the most demanding Military applications.

CM-ATR-25/SIXHEX-16HP : ½ ATR, 5 Slot 6U, 575W PSU, 9Kg

Product Highlights

- Contaminant-free enclosure
- Available in ½, ¾ & 1 ATR size
- VPX, VME64 & cPCI ready
- Accepts Conduction & Air-cooled 6Us
- Flexible Top & Bottom I/O wiring
- Six internal Heat Exchangers
- 16 integrated Heat Pipes
- Up to 1.8 KW total Power Dissipation
- Up to 150 W per slot
- Integrated Temperature Control Unit
- Dramatically increases payload MTBF
- 2 User defined PSU DC outputs
- 25°C less than heat exchanger ATRs
- 50°C less than conventional ATRs
- In-line EMI/EMC MIL-STD 461E Filter
- Stand alone low weight solution
- Customizable to specific requirements
- Mounting Tray with quick release system

CM-ATR-25 SixHex-16HP Power Supply Options		Vin Options		Backplane DC				Suggested BUS			Max Input Power
		28VDC	OTHER†	+5VDC	+3.3VDC	+12VDC	-12VDC	VME64	cPCI	VPX	Excluding Rear Fans
PSU Models	A-475W	✓		40A	22A	8A	8A	✓			700W
	A-575W		✓	40A	22A	12A	12A	✓			850W
	B-450W	✓		20A	45A	8A	8A		✓		700W
	B-550W		✓	20A	45A	12A	12A		✓		800W
	C-475W	✓		20A	22A	16A	8A			✓	700W
	C-575W		✓	20A	22A	21A	12A			✓	850W

† All Inputs except 28VDC : 48VDC / 270VDC / Autorange 90-264VAC @ 47-880Hz / 200VAC-3Phase @ 47-880Hz

It has been demonstrated in the field that military ATR enclosures are crucial to your end system reliability and performance. CM Computers is proud to develop superior products that guarantee your payload electronics are matched with excellence.

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

CM Computer SixHex: Pure Power, Pure Dissipation, Pure Thermodynamics

Visit **www.cmcomputer.com** or contact us at **info@cmcomputer.com** to request our latest Chassis Catalog.

CM Computer
True Military COTS Products