VME and Critical Systems

www.vmecritical.com

VOLUME 28 NUMBER 1

SPRING 2010

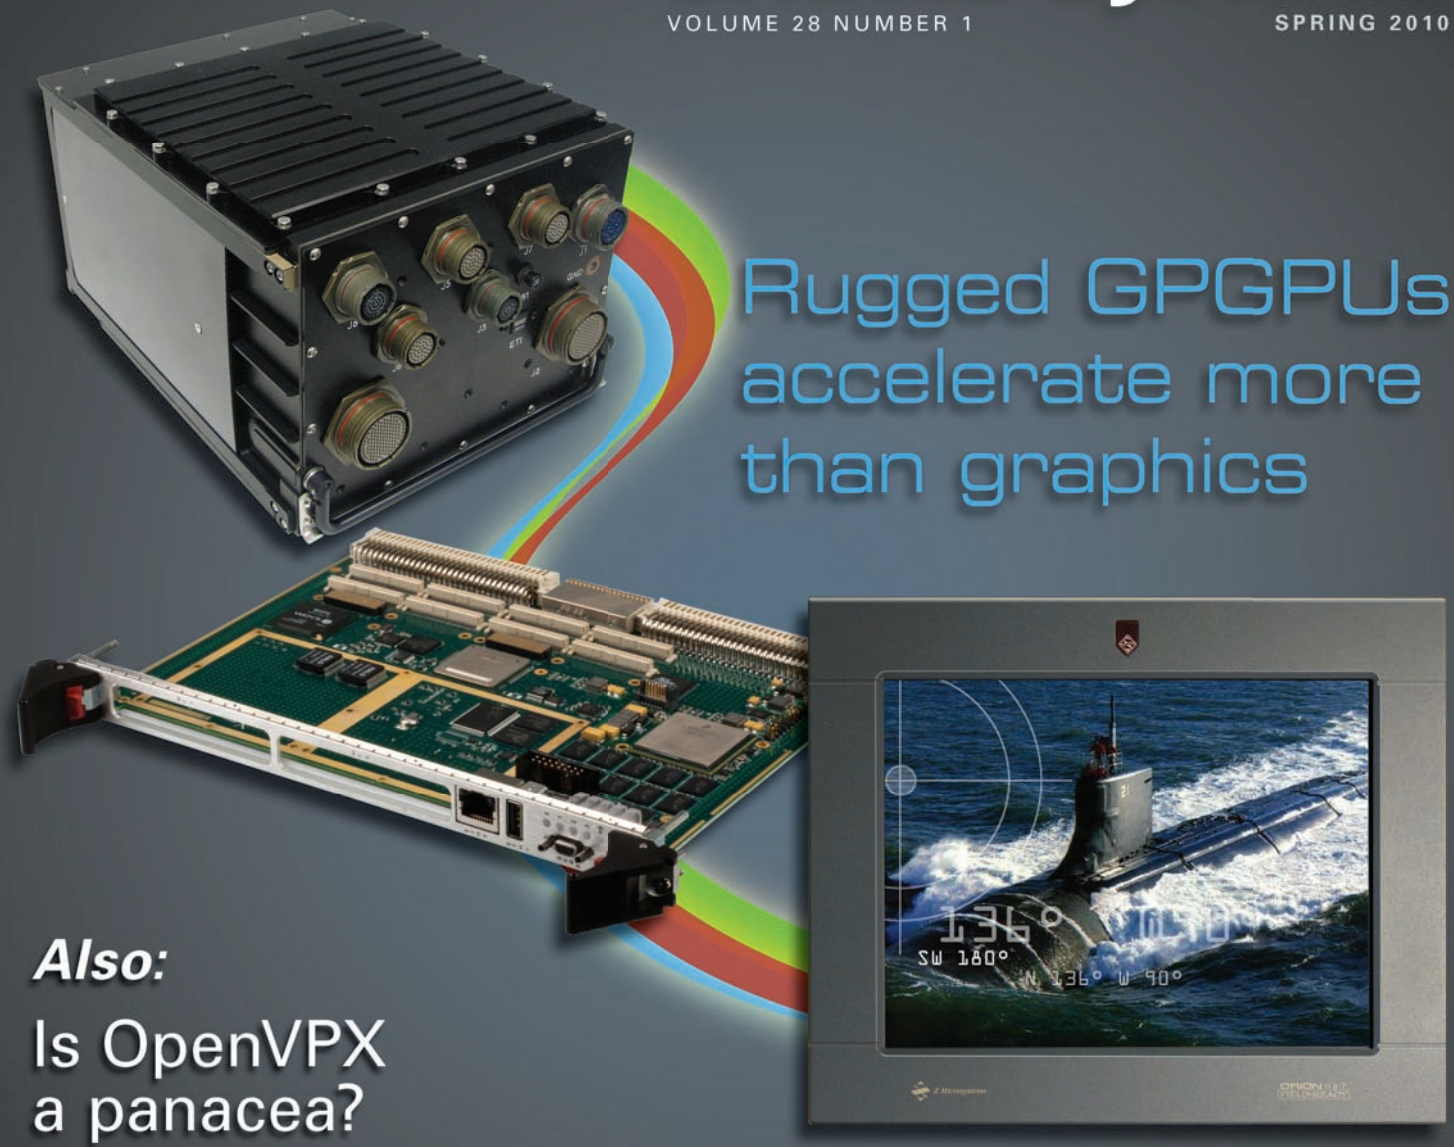

Rugged GPGPUs accelerate more than graphics

Also:

Is OpenVPX a panacea?

Migrating to OpenVPX

OpenVPX NOW

Ascend to higher levels of performance and interoperability

Migrate your performance away from a proprietary, closed architecture to flexible open solutions. Ensemble™ Series OpenVPX™ 3U/6U solutions are scalable from simple switchless to complex configurations using board types such as SBCs, switch and I/O modules, and HCD boards.

Why OpenVPX

- Interoperability ensures scalable, flexible solutions
- Migration with Mercury maintains high performance
- Mercury leadership kick-starts open standards
- Extensive product line deployed and available today

MERCURY COMPUTER SYSTEMS™

Visit mc.com/openvpxsolutions

NEW LEVELS OF INTERCONNECT RELIABILITY

For lengthy space missions tasked with discovering the effects of atmospheric phenomena on vital communication transmissions and data gathering, it's critical that equipment be long lasting and impervious to damaging effects. Hypertronics interconnect solutions featuring the Hypertac® hyperboloid contact technology offer the reliability and dependability that are required in making space discovery possible. Hypertronics is the choice of engineers worldwide for mission critical systems.

Hypertac Features	Benefits
Low Insertion/Extraction Forces	High Density Interconnect Systems
Long Contact Life	Low Cost of Ownership
Lower Contact Resistance	Low Power Consumption
Higher Current Ratings	Maximum Contact Performance
Shock and Vibration Immunity	Reliability Under Harsh Conditions

Maintain your system integrity with Hypertronics reliable connectors and dependable solutions.

www.hypertronics.com/discover

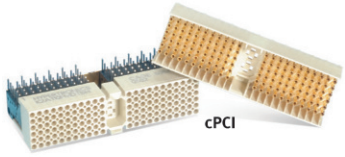

HYPERTRONICS: WHEN FAILURE IS NOT AN OPTION

smiths
bringing technology to life

www.vmecritical.com

VME and Critical Systems

Spring 2010 Volume 28 Number 1

Columns

Departments

ON THE COVER:
GPGPUs from vendors like NVIDIA do more than just crank out the polygons. Increasingly, their vector performance is being harnessed in rugged VME systems to replace CPUs and FPGAs in number-crunching applications like signal processing. See articles starting on page 12. Images courtesy of Curtiss-Wright Controls Embedded Computing (ATR), Extreme Engineering Systems (board), and Z Microsystems (rugged display).

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

Hardware: Pedal to the "mettle" – GPUs accelerate performance, video processing

Special: VPX/OpenVPX – Panacea for all that ails rugged systems?

Technology: HALT/HASS defy fear of failure

Event

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

Web resources

Subscribe to the magazine or E-letter
Live industry news
Submit new products
http://submit.opensystemsmedia.com

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

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

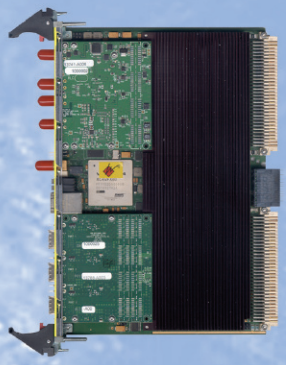

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

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

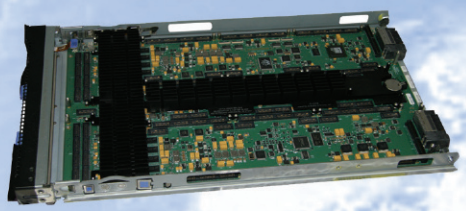

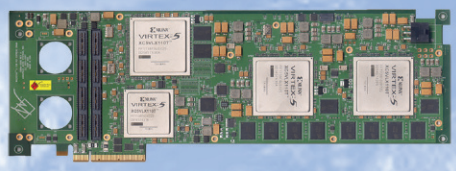

Annapolis Micro Systems

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

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

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

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

ADVERTISER INFORMATION

Page	Advertiser/Ad title
5	**Annapolis Micro Systems, Inc.** – The FPGA systems performance leader
19	**CM Computer** – Expand your military system horizons
21	**Concurrent Technologies, Inc.** – x86 processor boards
30	**Elma Electronic** – The industry's choice in VPX handles and panels
22	**Elma Electronic – Systems** – Introducing System Paks
17	**Excalibur Systems, Inc.** – Ready for the unexpected?
17	**Hartmann Elektronik** – VME
3	**Hypertronics** – Discover
32	**Kontron** – We do not build military aircraft
2	**Mercury Computer Systems** – Ascend to higher levels of performance and interoperability
29	**Nallatech** – High performance FPGA solutions
23	**Pentek, Inc.** – Stroke of genius
11	**Phoenix International** – RPC12 ruggedized 3U Fibre Channel RAID system
31	**Technobox, Inc.** – Create your own PMCs
11	**TEWS Technologies LLC** – COTS I/O solutions
28	**VEROTEC Electronics Packaging** – Good news!

E-LETTERS

Summer: www.vmecritical.com/eletter

- Getting 10 GbE onto the XMC platform: VITA 42.6 defines high-speed Ethernet standard
  By Rob Kraft, AdvancedIO Systems Inc.
- Time-triggered Ethernet – A powerful cure for what ails real-time networked applications
  By Ben Daniel, GE Intelligent Platforms
- Case study: Using VME to build an FPGA-based data acquisition system
  By Fred Joerger, Joerger Enterprises, Inc.

VME and Critical Systems

OpenSystems media

Military EMBEDDED SYSTEMS | DSP-FPGA.com | PC/104 and small form factors | INDUSTRIAL EMBEDDED SYSTEMS | CompactPCI Advanced & MicroTCA | Embedded COMPUTING DESIGN

Editorial/Business Office

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

Publishers: John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman

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

VME and Critical Systems is published five times a year (April, June, Aug, Oct, Dec) by OpenSystems Media, 16626 E. Ave of the Fountains, Ste 203, 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 203, Fountain Hills, AZ 85268

Editor's Foreword

By Chris A. Ciufo, Editor

Mentor automates requirements tracking for the military market

The Aerospace and Defense (A&D) market isn't an easy one to play in. Long design cycles, huge up-front costs with no guarantee of success, endless details buried in specifications, and disparate and sometimes warring teams spread across prime- and sub-contractors.

Sounds just like an FPGA or ASIC design, doesn't it? No wonder Mentor Graphics is leveraging their astoundingly huge product portfolio into the military market. I recently met with Embedded Systems Division company executives Glenn Perry (General Manager) and Shay Benchorin (Director) at the Embedded Systems Conference and boy, are these guys serious about A&D. Moreover, they have the resources, the tools, and the tenacity to target some great Commercial Off-the-Shelf (COTS) products at military problems.

Mentor's not new to A&D. The company's DO-254 tool suite has been around for a while and targets FAA and EASA aviation requirements compliance for safety-critical hardware, now mandated by contract for FPGAs and ASICs. As we went to press, the company announced an update to its HDL Designer tool, adding HDL coding checks for DO-254 compliance. In effect, FAA policy changes require designers to document their HDL coding standards prior to hardware design, and then review and verify compliance. The Mentor tool automates and simplifies compliance without requiring expensive third-party reviews. So it's not surprising that Mentor would expand the concept of requirements creation and compliance assurance to the entire military design process, not just the hardware or IC portion.

Sensing an opportunity, Mentor's newly released[1] ReqTracer product is a requirements tracking tool that eliminates the painful process of documenting requirements and making incessant changes in Excel spreadsheets. Though initially oriented at hardware designs – IC, board- and system-level – the tool is capable of conducting software requirements tracking a la FAA/RTCA DO-178B/C. Figure 1 shows that modern A&D platforms such as the JSF certainly have a wide variety of design and verification data to track.

But what is "requirements tracking"? In a military platform, it consists of capturing and tagging (keywords) all requirements in the SOW, applying typical DoD verbiage such as, "shall, will, may" (and avoiding definition ambiguity like "fast" or "good"), managing the gazillions of ECOs, and tracing all of this over a typical 10-year program papered with DIDs and CDRLs[2]. It's easy to see how just about any halfhearted tool could replace Excel for this n-dimensional matrix, but Mentor made sure to integrate with other widely used tools such as IBM's DOORS, Word and PDF files, advanced verification suites for HDL (Mentor is an EDA company, after all), as well as Design and Testbench Data in C/C++, Verilog, and even MATLAB. And since we're discussing military and safety-critical requirements, Mentor included a traceability matrix that generates artifacts for hardware (and soon software) certifiability.

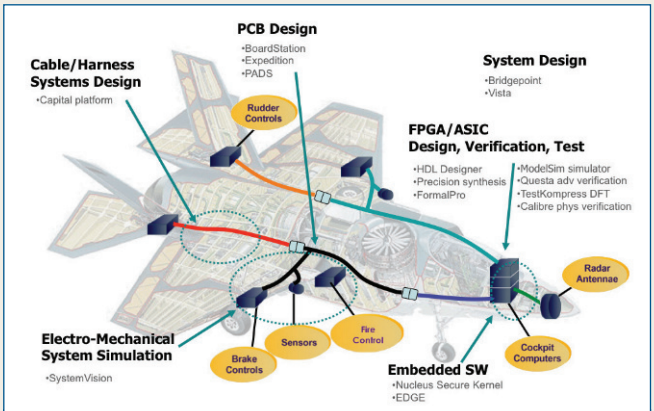

Figure 1 | Mentor Graphics is targeting aerospace and defense applications with design tools for hardware, software, and systems – including the new ReqTracer requirements management tool (not shown).

It's unclear to me how the tool integrates with other Application Life-cycle Management (ALM) tools such as offerings from Tail-f Systems and Coverity, or LDRA's newly repackaged Embed-X with TBreq. Announced in March at Embedded World, LDRA's Embed-X is an embedded ALM suite of products that "merges product requirements, business objectives, and metrics" into an actionable perspective. Embed-X focuses on certification support for DO-178B/C (the companion spec to DO-254) by automating traceability throughout the whole process. The company has partnered with Visure Solutions' "requirements solutions" to create a fully certified application and "agile development" that supposedly cuts down the iterative process from months to weeks.

While LDRA has focused on the DoD from day one, Mentor is clearly picking from its arsenal and locking on military programs. Between the two of these companies, 10-year military programs will surely benefit from requirements tracking and ALM. Time to toss that Excel spreadsheet.

Chris A. Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

[1] ReqTracer has actually been available for about 12 months, but Mentor wasn't talking much about it until recently.
[2] Data Item Descriptions; Contract Data Requirements Lists

VITA
Open Standards, Open Markets

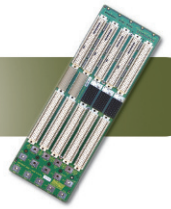

VITA News

By Ray Alderman

Out with 2009, enter 2010

The year 2009 was a tough year in some board and system market segments, but it was a good year for other segments. The telecom segment suffered the most, with declines in sales of up to 70 percent for some vendors. The industrial market bowed to the financial and economic slowdown and some vendors saw 50 percent declines in sales, especially in Europe. The medical market went on hiatus as the new health-care bill meandered through Congress and caused diagnostic equipment purchases to slow. The major risk of those investments – including the amount hospitals can charge for MRI, CAT scan, and PET scan procedures under the new plan – motivated buyers to sit and wait.

The U.S. DoD budget contained some surprises, too. Everyone expected the F-22's cancellation, in addition to some of the other program terminations (the new presidential helicopter) and reductions (slowing the replacement purchases of new aircraft carriers). However, the surprise for military vendors was how effectively the new *procurement reform* initiatives would slow spending on funded programs. Congress did not reduce the military budgets. But the new procurement rules and paperwork have definitely slowed down the issuance of contracts, having the same effect as a budget reduction.

At least for a while, we are done with most of the major reactions to the economic, financial, and technology turmoil of 2009. It is now time to continue to plan for 2010 and beyond.

Telecom: Bankruptcy or growth

The telecom service providers must upgrade their networks, and Wall Street says that both AT&T and Verizon will face serious financial challenges as we move forward. Recent reports claim that the AT&T network is crashing in places because of the data-hungry iPhones it feeds. Another report indicates AT&T needs to spend up to $18 billion upgrading their network for 3G services, but the revenues will fall far short of paying for those upgrades. This is the same for Verizon, but to a lesser degree. As a result, the telecom board segment looks questionable for 2010.

> "I expect to see the increased number of technology refreshes and upgrades to existing platforms push VME board sales to a new high."

Medical equipment markets: The socialized medicine effect

The medical market will not settle down until hospitals and doctors know how much they will be paid for their services and the use of their equipment. Municipal budgets are in shambles, so there is not much money available for public hospitals to buy new equipment. That makes this market segment, mostly comprising motherboards for image reconstruction in diagnostic procedures, also questionable until we fully understand the new health-care scheme.

Industrial segment: Pent-up demand for energy efficiency

There is pent-up demand in the industrial segment for more energy-efficient equipment to be installed in all facets of manufacturing.

Recent announcements state that the ongoing Air Traffic Control System upgrades in the United States will consume up to 250 new systems in 2010. The Obama administration has awarded $3.4 billion in grants for the *smart grid* upgrades and revamping the nation's electrical distribution and generation system. So this segment (motherboards and small form factor products) looks promising for 2010.

MIL/AERO/COTS markets: Pushing the performance curve

However, I believe the MIL/COTS market will again be the winner in 2010. I expect to see the increased number of technology refreshes and upgrades to existing platforms push VME board sales to a new high.

VPX design-ins promise very high growth rates and shipments in 2010. I see approximately 80 percent of military sales in VME, with about 20 percent in VPX in 2010. And for 2011 and 2012, I see that ratio rising to 70 to 30 percent and 60 to 40 percent, respectively.

When we get to prolific 10G silicon this year, we have some severe technical problems to resolve: Signal-to-Noise Ratios (SNRs) will approach zero, crosstalk will become severe at 10G, and signal integrity will become a nightmare. So in 2010, I think we will see vast engineering resources dedicated to new optical connectors, fibers, and silicon.

VITA sent out an industry-wide invitation in February 2010 to join a special optical feasibility study group to define the state of the art for optical backplanes and the problems that we need to solve. Sometime around 2013, we will begin to field pure optical backplane computers in many traditional applications, and we need to start the standards work now. So 2010 will usher in the transition to optical connections, and that will be an exciting time for our industry as we continue to push the bandwidth envelope.

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

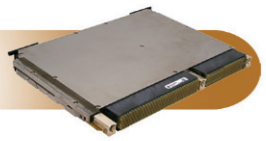

By Stewart Dewar

OPENVPX INTERCONNECTS...

OpenVPX eases harsh-environment systems integration

With OpenVPX (VITA 65), integrators are choosing interoperable module architectures for next-generation systems. Thus, high-performance embedded processing systems engineers have more reasons than ever to choose a VPX (VITA 46)-based system design. VPX was initiated to address the challenge of enabling high-speed serial interfaces using the multi-gigabit capable RT2 connector family. (And that was the easy part.) And, like any new innovation, VPX has endured its share of growing pains.

Today, however, is a different ballgame. OpenVPX – which specifies VPX system-level interoperability via defined profiles and planes – will soon be ratified, representing a great leap forward in the technical maturity of VPX. Stewart illustrates how OpenVPX simplifies system interconnection via a case study, then explains 1) how OpenVPX furthers VPX's ecosystem and 2) how VITA initiatives are broadening OpenVPX.

System interconnection made simple: Case study

OpenVPX defines a generic system interconnection architecture for VPX consisting of three primary communications planes – Data, Expansion, and Control. Currently the defined option for OpenVPX Data and Expansion planes are PCI Express and Serial RapidIO, whereas the defined options for the Control plane are various implementations of Ethernet.

To illustrate the concept of the different communication planes, Figure 1 depicts a high-performance video processing subsystem. Such a system could be deployed in a military vehicle to provide immersive visual situational awareness to the vehicle commander. The system diagram illustrates the following uses of the communication planes:

- The Image Capture module moves data to the Control Processor over the Expansion Plane.
- The Control Processor moves data to and from the Video Compression and Encryption module over the Expansion Plane.

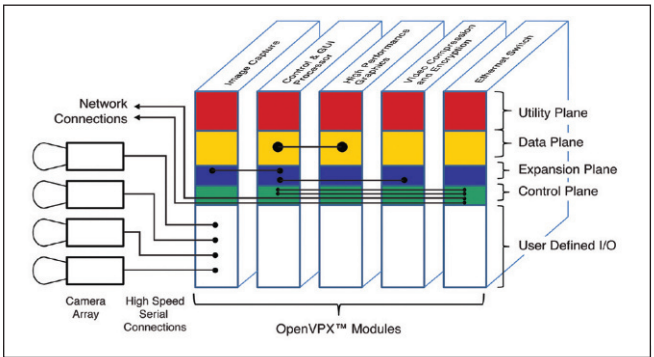

Figure 1 | A conceptual system diagram of a high-performance video processing subsystem, illustrating various usages of the OpenVPX communication planes.

- The Control Processor communicates to the High Performance Graphics module over the Data Plane; typically this would be PCI Express to provide the highest performance and most direct connection between the Control Processor CPU and the graphics processor.
- The Control Processor communicates with other system elements and to external equipment via the Ethernet-based Control Plane.

OpenVPX furthers the VPX ecosystem

No longer considered a niche technology, VPX is now mainstream in the military and aerospace industries. The VITA online directory includes nearly 100 OpenVPX-compliant products and indicates/illustrates the introduction of second- and third-generation products in core areas such as SBCs and DSPs. Perhaps the strongest indicator of VPX's product ecosystem's maturity is the appearance of specialized supporting products like test backplanes, power load modules, and a cable-terminating connector that enables cabling directly to the backside backplane connectors instead of using a rear transition module.

Ongoing VITA standards broaden OpenVPX

Looking forward, VITA continues the work needed to broaden the scope of its VPX and OpenVPX standards. Numerous initiatives are underway, including:

VITA 62 – The development of a standard power supply module for OpenVPX systems. This will allow system integrators to design a standard power supply slot in their system and choose a power supply from a competitive pool of suppliers. As a comparative measure of maturity, even the venerable VMEbus never included a standard for a power supply module.

VITA 66 – Defines an optical interconnect for use on OpenVPX modules.

VITA 67 – Standardizes RF connectors for use on OpenVPX modules and allowable combinations and positioning of RF connectors and standard connectors. This provides a standards-based approach to route RF signals such as radar or communications signals to the receiver payload module, avoiding front-panel I/O or customized solutions.

OpenVPX readies mil/aerospace market

Equipped with an ongoing arsenal of product announcements, the VPX community, now armed with the new OpenVPX standard for improved interoperability, has never been more ready to build harsh environment, computer-intensive military/aerospace systems.

***Stewart Dewar**, Product Marketing Manager at General Dynamics Canada, can be reached at Stewart.Dewar@gdcanada.com.*

VITA Standards Activity Chart

MARCH MEETING

Standard *Reaffirmed	Title	Status	*VME and CS* edition
ANSI/VITA 1.0 *2002	VME64 Standards	Released	
ANSI/VITA 1.1 *2003	VME64 Extensions	Released	Aug. 2004
ANSI/VITA 1.3 *2003	9U x 400 mm Format	Released	
ANSI/VITA 1.5	2eSST	Released	Feb. 2004
ANSI/VITA 1.6 *2005	Keying for Conduction-cooled VME	Released	
ANSI/VITA 1.7	Increased Connector Current Level	Released	
ANSI/VITA 3 *2002	Board Level Live Insertion	Released	
ANSI/VITA 4.0 *2002	IP Modules	Released	
ANSI/VITA 4.1 *2003	IP/I/O Mapping to VME64x	Released	
ANSI/VITA 5.1 *2004	RACEway Interlink	Released	
VITA 5.2	RACEway++	Withdrawn	Aug. 2004
ANSI/VITA 6.0 *2002	SCSA	Released	
ANSI/VITA 6.1 *2003	SCSA Extensions	Released	
ANSI/VITA 10 *2002	SKYchannel Packet Bus	Released	
ANSI/VITA 12 *2002	M-Modules	Released	
ANSI/VITA 13	Pin Assignments for HIC on VME	Withdrawn	
ANSI/VITA 17.0 *2004	Front Panel Data Port	Released	
ANSI/VITA 17.1	Serial Front Panel Data Port	Released	Feb. 2004
VITA 17.2	Serial Front Panel Data Port (SFPDP) Channel	Working Group	Dec. 2009
VITA 19.0	BusNet Overview	Withdrawn	
ANSI/VITA 19.1	BusNet MAC	Withdrawn	
ANSI/VITA 19.2	BusNet LLC	Withdrawn	
ANSI/VITA 20 *2005	Conduction-cooled PMC	Released	Apr. 2005
ANSI/VITA 23 *2004	VME64x Extensions for Physics	Released	
ANSI/VITA 25	VISION	Withdrawn	
ANSI/VITA 26 *2003	Myrinet-on-VME	Released	
ANSI/VITA 29	PC•MIP	Released	
ANSI/VITA 30.0 *2005	2 mm Connector Practice on Euroboard	Released	
ANSI/VITA 30.1	2 mm Conduction-cooled Euroboard	Released	
VITA 30.2	Power Connector Equipment Practice	Released	Apr. 2007
ANSI/VITA 31.1	GbE on VME64x Backplanes	Released	Feb. 2004
ANSI/VITA 32	Processor PMC	Released	Feb. 2004
VITA 34	A Scalable Electromechanical Architecture	Working Group	Apr. 2004
ANSI/VITA 35 *2005	Pin Assignments for PMC to VME	Released	
VITA 36	PMC I/O Modules	Withdrawn	Apr. 2004
ANSI/VITA 38	System Management on VME	Released	
ANSI/VITA 39	PCI-X Aux. Std. for PMCs and PrPMCs	Released	Feb. 2004
ANSI/VITA 40	Status Indicator	Released	Dec. 2009
ANSI/VITA 41.0	VXS: VME Switched Serial	Released	Oct. 2006
ANSI/VITA 41.1	VXS: InfiniBand Protocol Layer	Released	Oct. 2006
ANSI/VITA 41.2	VXS: RapidIO Protocol Layer	Released	Oct. 2006
VITA 41.3	VXS: GbE	Working Group	Apr. 2006
VITA 41.4	VXS: PCI Express	Working Group	Apr. 2006
ANSI/VITA 41.6	VXS: 1x GbE Control Channel Layer	Released	Sept. 2009
VITA 41.7	VXS: Processor Mesh Topology	Working Group	
VITA 41.8	VXS: 10 GbE Protocol Layer	Working Group	June 2009
VITA 41.10	VXS: Live Insertion Requirements for VITA 41 Boards	Working Group	Apr. 2006
VITA 41.11	VXS: Rear Transition Modules	Working Group	Apr. 2006
VITA 42.0	XMC	Released	Feb. 2009
ANSI/VITA 42.1	XMC: Parallel RapidIO	Released	Oct. 2006
ANSI/VITA 42.2	XMC: Serial RapidIO	Released	Oct. 2006
ANSI/VITA 42.3	XMC: PCI Express	Released	Oct. 2006
VITA 42.4	HyperTransport	Working Group	Apr. 2005
ANSI/VITA 42.6	XMC: 10 GbE 4-Lane Protocol Layer	Released	June 2009
VITA 42.10	XMC: General Purpose I/O	Working Group	
VITA 42.20	XMC: Dual Fabric I/O	Working Group	
VITA 43S	Hot Swap NextGen Mezzanine	Inactive	Feb. 2004
VITA 45S	Serial VME	Canceled	Apr. 2004
ANSI/VITA 46.0	VPX: Base Specification	Working Group	Feb. 2009
ANSI/VITA 46.1	VPX: VMEbus Signal Mapping	Working Group	Feb. 2008
VITA 46.3	VPX: Serial RapidIO on VPX Fabric Connector	Working Group	Spring 2010
VITA 46.4	VPX: PCIe Mapping and Advanced Switch Signal Mapping	Working Group	Spring 2010
VITA 46.5	VPX: HyperTransport	Inactive	
VITA 46.6	VPX: GbE	Working Group	Spring 2010
VITA 46.7	VPX: 10 GbE	Working Group	Spring 2010
VITA 46.9	PMC/XMC/Ethernet Signal Mapping to 3U/6U on VPX User I/O	Working Group	Spring 2010
ANSI/VITA 46.10	Rear Transition Module for VPX	Released	Dec. 2009
VITA 46.11	System Management on VPX	Working Group	Spring 2010
VITA 46.12	Fiber Optic Interconnect	See VITA 66	Dec. 2009
VITA 46.14	Mixed Signal VPX	See VITA 67	Dec. 2009
VITA 46.20	VPX Switch Slot Definition	See VITA 65	June 2009
VITA 46.21	Distributed Switching on VPX	See VITA 65	June 2009

Defining Standards

VITA 68 allows a common backplane design for multiple fabric protocols

By Bob Sullivan

VITA 68 (VPX Compliance Channel) defines a VITA 46 (VPX) compliance channel based on the channel performance methodology and metrics used in IEEE 802.3-2008 for 1000BASE-KX, 10GBASE-KX4, and 10GBASE-KR for up to 10.3 Gbaud per pair.

This methodology is comprehensive and defines a number of frequency domain channel parameters that can be adapted to the PCI Express and Serial RapidIO Gen 1 and Gen 2 baud rates. Key channel parameters are:

- Insertion Loss (IL)
- Fitted Attenuation (FA)
- Insertion Loss Deviation (ILD)
- Return Loss (RL)
- Insertion Loss to Crosstalk Ratio (ICR)
- Intra-pair skew
- Lane-lane skew

VITA 68 defines an end-to-end channel including multi-line coupled plug-in module Tx and Rx pads, vias, and traces as well as multi-line coupled backplane traces, mated connectors, and connector via footprints. It also allocates a budget for each parameter to the backplane portion only. Based on preliminary signal integrity simulation work performed by Hybricon, an initial signal integrity budget has been established for the backplane portion. This is documented in the VITA 68 draft standard.

The idea is to precisely define the backplane performance requirements for each baud rate such that any fabric protocol that runs at that rate could be supported. For example, a 3.125 Gbaud backplane would need to support 3.125 Gbaud Serial RapidIO, Ethernet XAUI, or 10 GBASE-KX4, plus lower baud rates such as 2.5 Gbaud PCI Express. Each fabric protocol needs to operate at a specified Bit Error Rate (BER); at higher rates this requires modeling the receiver equalization, which varies with different protocols. These different fabric protocols would all use the same backplane channel.

Signal integrity simulations

Since the backplane is the least common denominator for VPX interoperability, the VITA 46 and VITA 68 working groups

have agreed that the VITA 68 working group will develop a signal integrity Statement of Work (SOW). The working group will oversee the signal integrity simulation work to finalize the signal integrity budget for the backplane and to define the signal integrity requirements for the VITA 46.3 Serial RapidIO, VITA 46.4 PCI Express, and VITA 46.6/VITA 46.7 Ethernet VPX dot-specs. A draft SOW has been developed based on previous efforts in the VITA 46.3 and VITA 46.4 working groups.

Latest VITA 68 developments

VITA 68 became a working group within the VITA Standards Organization (VSO) at the September 2009 meeting. The draft specification and the draft signal integrity SOW are just now undergoing their first VITA working group ballot.

Bob Sullivan *is vice president of technology at Hybricon Corporation and is chairperson of VITA 68. He can be contacted at bobs@hybricon.com.*

Standard *Reaffirmed	Title	Status	*VME and CS* edition
ANSI/VITA 47	Env., Design and Const., Safety, and Qual. for Plug-in Units	Released	June 2006
VITA 47r1	Revisions to ANSI/VITA 47	Released	Feb. 2008
VITA 47r2	Revisions to ANSI/VITA 47	Working Group	Dec. 2009
VITA 48.0	REDI: Ruggedized Enhanced Design Implementation	Working Group	Dec. 2009
VITA 48.1	Mechanical Specs for Microcomputers Using Air Cooling	Working Group	Dec. 2009
VITA 48.2	Mechanical Specs for Microcomputers Using Conduction Cooling	Working Group	Dec. 2009
VITA 48.3	Mechanical Specs for Microcomputers Using Liquid Cooling	Working Group	
ANSI/VITA 49.0	VITA Radio Transport (VRT)	Released	May 2009
ANSI/VITA 49.1	VITA Radio Link Layer (VRL)	Released	May 2009
VITA 50	Best Practices for Electronic Module Cooling	Inactive	Dec. 2007
ANSI/VITA 51.0 *2008	Reliability Prediction	Released	Aug. 2008
ANSI/VITA 51.1 *2008	Reliability Prediction: MIL-HDBK-217 Daughter	Released	
VITA 51.2	Physics of Reliability Failure	Working Group	Feb. 2009
ANSI/VITA 51.3	Qualification and Environmental Stress Screening	Released	Spring 2010
VITA 52	Lead-free Practices	Working Group	Oct. 2006
VITA 53	Commercial Technology Market Surveillance	Working Group	Feb. 2009
VITA 54	Embedded Platform Management Architecture (EPMA)	Inactive	Aug. 2005
VITA 55	Virtual Streaming Protocol	Inactive	Feb. 2009
VITA 56	Express Mezzanine Card (EMC)	Inactive	Oct. 2007
ANSI/VITA 57 *2008	FMC: FPGA Mezzanine Card	Released	Feb. 2009
VITA 57.1	FPGA I/O Mezzanine Pin Assignments	Working Group	June 2009
ANSI/VITA 58.0	Line Replaceable Integrated Electronics Chassis	Released	May 2009
VITA 59	RSE: Rugged System-on-Module Express	Working Group	Dec. 2008
VITA 60	Alternative Connector on VPX	Working Group	Dec. 2009
VITA 61	Alternative Connector for XMC	Working Group	Dec. 2009
VITA 62	Power Supply Modules	Working Group	Dec. 2009
VITA 63	KVPX	Working Group	Feb. 2009
VITA 64	Optimized Footprint for VITA 60	Working Group	Feb. 2009
VITA 65	OpenVPX	Working Group	
VITA 66	Fiber Optic Interconnect (Formerly 46.12)	Working Group	Spring 2010
VITA 67	Analog/RF Interconnect (Formerly 46.14)	Working Group	Spring 2010
VITA 68	VPX Compliance Channel	Working Group	
VITA 69	Common Glossary	Working Group	
VITA 70	Common Standard Template	Working Group	
VITA 71	Rugged Mezzanine	Working Group	
VITA 72	Connector Comparision Testing	Working Group	

General-purpose GPUs breathe new life into high-performance embedded computing

By Anne Mascarin and Scott Thieret

The GFLOPS/Watt metric is now seen as an essential measure of embedded image processing applications in defense programs, where GFLOPS/Watt metrics are crucial. However, General-Purpose GPUs (GPGPUs) are currently being used instead of CPUs in high-performance embedded computing applications where GFLOPS/Watt metrics are paramount. Prior to deciding on the CPU or GPGPU path, though, it's important to understand the differences between GPUs and GPGPUs – and to comprehend how these GPGPUs (as opposed to CPUs) are an ideal fit for high-performance embedded computing applications.

Many advanced applications for high-performance embedded computing demand excessive amounts of computing power. Real-time imaging systems in applications such as persistent surveillance and electronic warfare applications, among others, require the highest possible GFLOPS/Watt to meet performance requirements without exceeding the power budget. Traditional CPU-based boards simply don't meet these power budget constraints.

General-Purpose GPUs (GPGPUs) are currently being used in high-performance embedded computing applications where GFLOPS/Watt metrics are paramount. Before deciding whether to embark on the CPU or GPGPU route, however, it's important to explore the differences between GPUs and GPGPUs – and to understand how these GPGPUs (as opposed to CPUs) are a natural fit for high-performance embedded computing applications.

The rise of the GPU versus CPU equation

Government programs are putting the squeeze on prime contractors to develop more warfighting capability faster. At the same time, the needs of embedded defense computing platforms are accelerating: to acquire more data and arrange and process it more quickly, with the goal of extracting actionable information immediately and making it available in real time to the warfighter. The need for creative and innovative solutions to the "actionable information" problem has never been stronger.

Government agency mandates and the requirement for actionable information aren't the only pressures that affect prime contractors. Consider Size, Weight, and Power (SWaP) and historical constraints that greatly impact the adoption and performance of deployed platforms. Together, these issues force prime contractors to turn to innovative solutions in order to squeeze every ounce of performance out of their subsystems.

GFLOPS/Watt matters

Real-time imaging systems in deployed environments such as persistent surveillance, onboard exploitation, and electronic warfare applications, among others, require the highest possible GFLOPS/Watt to meet deployed performance requirements. Frequently, these subsystems are the last to be added onto the airframe and are subsequently allotted the platform power budget's smallest portion.

Many CPU-based boards can't keep up with the stringent GFLOPS/Watt requirements. For example, peak theoretical GFLOPS/Watt for the IBM Cell processor is 1 GFLOP/Watt, while the peak theoretical GFLOPS/Watt for AMD's ATI RV770 GPGPU is 9.23 GFLOPS/Watt. Graphical Processing Units (GPUs), first introduced by NVIDIA in 1999, have always had very high GFLOPS/Watt metrics. Although the GFLOPS/Watt value increases with every new chip revision, GPUs traditionally performed best in the application for which they were designed – graphics processing in desktop systems.

Serial or parallel processing?

GPUs have been architected to maximize arithmetic and logic performance on one type of very self-similar data. CPUs, on the other hand, offer more control support and disposition of data – such as flow control and caching. In general, CPUs operate on data in a serial fashion; for example, even when a matrix operation is performed, it is necessary for the CPU to perform the overhead task of loading each input element sequentially. CPUs were architected to support flow control for comparisons and decision making, in addition to calculations. Conversely, GPUs process data in a parallel arrangement because they contain a matrix of multiple cores of simple Arithmetic-Logic Units (ALUs) to rapidly perform simple calculations in parallel. This high degree of parallelism is what makes GPUs efficient and fast image processing engines for high-performance military applications.

> "So, given their parallel performance potential and low power consumption, why haven't GPUs been utilized much for high-performance embedded computing?"

Programmability versus upgradeability

So, given their parallel performance potential and low power consumption, why haven't GPUs been utilized much for high-performance embedded computing? Programmability is one reason, and upgradeability is another.

The software environment for GPUs is notoriously non-intuitive even to proficient

embedded programmers. The environment is based on graphics primitives – not high-level language constructs or even CPU assembly variants. And the basic structure of programming tools for GPUs does not offer the optimizations that programming languages for CPUs do. GPGPUs are a relatively recent concept of GPU computing, offering a developer-friendly software environment. Software developers can now program GPGPUs with familiar constructs such as well-defined APIs and indexed matrix operations.

Historically, GPUs have not been easily upgradeable; they have been discrete components soldered directly onto the printed circuit boards. Upgrading the chip as new versions become available would require a complete board respin. Many of today's GPGPUs (from ATI, NVIDIA, and others), however, are available in a mobile PCI Express module (MXM): an easy-to-insert format that facilitates upgrades when new, faster GPGPUs are available. Adherence to the MXM specification, developed by NVIDIA and now a stand-alone specification, ensures easy upgradeability for technology updates.

Figure 1 | The Sensor Stream Computing Platform

GPGPUs: A natural fit for high-performance embedded applications

A high GFLOPS/Watt ratio, parallel processing capabilities, and a programmable software environment and upgradeability are all now available with today's GPGPUs. The application space in high-performance embedded computing for defense is clearly defined.

As mentioned, several applications in the high-performance embedded computing space could benefit from the use of GPGPUs. Persistent surveillance – an unmanned aerial vehicle application characterized by long mission duration and onboard sensor data exploitation – is a particularly good example. The long mission duration aspect of persistent surveillance demands minimal power consumption. Meanwhile, the intense computational aspect of onboard exploitation, including image stabilization and geo-registration, requires parallel processing – such as that provided by GPGPUs – to provide real-time, actionable information to the warfighter.

The missing link is a platform or environment that can support experimentation and algorithm tradeoffs. One such link is Mercury's Sensor Stream Computing Platform (SSCP, see Figure 1), a 6U VXS-development chassis that is the size of a piece of carry-on luggage, weighs 32 pounds, draws less than 600 W from a standard wall outlet, and achieves 3.84 TFLOPS (see Figure 2). The SSCP tunable power/performance operation allows the user dial-down GPU clock speed to minimize power consumption during periods of inactivity, as is required for persistent surveillance and similar applications. **CS**

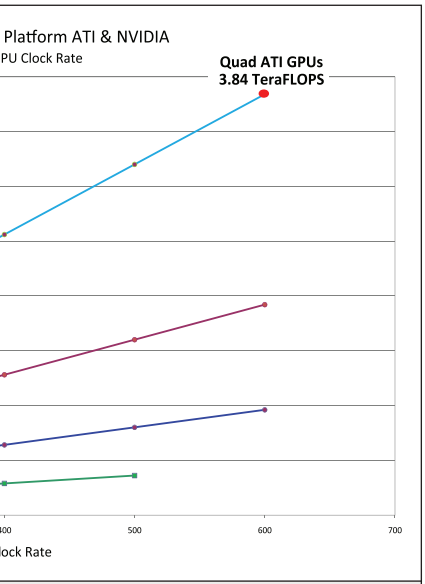

Figure 2 | Sensor Stream Computing Platform: Peak FLOPS versus GPU clock rate

Anne Mascarin *is a Product Marketing Manager at Mercury Computer Systems, where she has been employed for five years. Previously, she worked at The MathWorks and Analog Devices, Inc. Anne holds a Master of Science in Electrical Engineering from Northeastern University and a Bachelor of Arts in Economics from Boston University. She can be contacted amascari@mc.com.*

Scott Thieret *is the Technical Director for GPU Computing at Mercury Computer Systems, where he has been employed for 10 years in various positions dedicated to GPU development. Prior to Mercury, he worked at Avid, MITRE, and IBM. Scott holds a Bachelor of Science in Computer Engineering from the University of Vermont. He can be contacted at sthieret@mc.com.*

Making real-time, multichannel video processing a reality: There's an easier way

By George Schreck

Processing power … system power dissipation … data bandwidth … data connectivity … video latency. These all become critically important components of system development when designing a system to withstand the rigors of a mobile embedded system – and manage multiple video channels in a real-time environment. But the key might prove surprising: It's all in the processor(s).

To accommodate remote surveillance, vehicle autonomy, and situational awareness, there is a growing requirement for real-time, multichannel, remote video connections in the military sector. It is now relatively easy to connect desktop computers or even laptops with a video stream. However, once one attempts to connect multiple remote video streams in a low-power, non-equipment-friendly environment, it becomes painfully obvious that what appears simple in the office accelerates into a significantly larger challenge in an embedded environment with myriad interrelated problems, each affecting the others' performance or cost.

As the channel count increases, some of the more prominent challenges include: having enough total processing throughput without dissipating too much power for the environment; having enough bandwidth and therefore performance on an individual channel to always maintain acceptable video quality; and managing to move all the data on a network connection while still guaranteeing minimal latency on all channels with possible loss of video recovery. However, a dedicated processing scheme is serving to solve the issues of power, data bandwidth, and data connectivity and latency – and make possible the video capabilities so essential to successful field operations.

Processing power and power dissipation

To understand the system requirements and possible performance/power tradeoffs, it is necessary to develop a metric based on a single-processor system. The most obvious issue to examine is that of processor bandwidth. To get a scale of the total processing power required to handle up to five video streams, a simple test was run that included a commercially available computer containing a dual-core, 3.16 GHz Pentium equivalent processor, 4 GB of DDR2 SDRAM at 666 MHz, and a GbE connection. The system was running Fedora 9 Linux along with GStreamer 0.10.20 video management software. Table 1 shows the data collected using multiple remote computers to stream the video to and from the target machine.

Number of channels	Data direction	Total processor load (%)
1	Encode	34.9
2	Encode	74.8
3	Encode	113
4	Encode	148
5	Encode	190
1	Decode	11.6
2	Decode	25.3
3	Decode	40
4	Decode	55
5	Decode	70

Table 1 | Measured PC processor load with 3,000 kb H.264

The total system power was estimated at 44 W without disks, fans, or power supply inefficiencies. Since vehicle autonomy applications such as supply transport and tactical ground video often require at least four video streams, one for each side of the vehicle, a full dual-core PC of about 57 W of power would be necessary. For optimal Mean Time Between Failure (MTBF) and therefore field survivability, complexity and power dissipation must be minimized. Therefore, a simpler solution with less power dissipation is needed.

Data bandwidth

Inadequate memory bandwidth often results in non-real-time video or missing video streams. Either scenario could be devastating for remote ground support or tactical awareness. In an application where a real-time tactical view is required from a vehicle such as a Humvee or a tank, the typical video processing flow can split into two logical sections as depicted by Figure 1.

To provide video processing as detailed in Figure 1, it is critical to have adequate aggregate internal memory bandwidth to manage all the data. If a single 640 x 480 video stream is being acquired and the image were 2 bytes per pixel and 30 frames per second, the aggregate required bandwidth to move the image once would be 18.4 MBps. Assuming the video encoding algorithm is running mostly from internal processor cache and the software is structured so the data copies are minimized, the data throughput requirement for a single channel would exceed 249 MBps. This number includes:

- Buffering the incoming frames for de-interlacing – 18.4 MBps
- Copying the data for a simple "weave" de-interlace – 18.4 MBps
- Reading and writing the data for compression – 209 MBps for a typical algorithm implementation
- Writing the data for data transmission – 1 MBps (assumes less than 8 Mb result)
- Ethernet protocol buffering, packetizing, etc. – in most systems this could be two or more copies – 2 MBps

This data throughput requirement does not include any bandwidth allocation for program accesses, data encryption, issues with cache page size or line size, cache flushing, physical memory access inefficiencies, stream scaling, or encryption support.

Since many high-performance graphics chips contain hardware assist for

compressing and decompressing video, one could drop 209 MBps from the aggregate for four channels. However, even after making that adjustment, it is not possible to get the total processing system power, given in the previous example, down to a reasonable embedded target of 10 to 15 W.

However, if one were to have an independent video hardware video compression/decompression engine for each video channel and a separate memory interface, the processing requirements would be greatly reduced. And the available memory bandwidth would scale with the number of processors. This leads one quickly to the conclusion that using a low-power embedded processor, such as the Freescale i.MX27, for each channel would be a more efficient solution. In fact, when building a system with four processors running simultaneously, the processor load is less than 12 percent for each processor for the same configuration settings that were used for the PC.

Data connectivity and video latency

Assuming the encoding/decoding of channels is done on individual processors, it is still necessary to make sure that the data streams are combined in a logical fashion, that the aggregate network bandwidth is below the requirements of the remote transmission media, and that the system maintains a low latency.

Managing the source and destination of the data for the channels is relatively easy since the streams are from independent processors. Each stream can be assigned to an Ethernet address and the data streams connected to each other by a UDP connection. This can be managed through a Web page interface and would be easily reconfigured to meet the changing demands of the theater of engagement. Since the bandwidth requirement for a single channel is relatively small, 2 to 4 Mbps, the combined bandwidth can be kept to a manageable number (under 10 Mb for four channels, for example).

When independent processors are being used, the bandwidth requirements scale linearly with the number of channels being encoded. If the bandwidth requirement approaches the actual maximum throughput of the Ethernet connection, the traffic from any single connection could adversely affect the latency of the other channels. In more extreme cases, it will cause frame loss on other channels. The threshold of this effect is difficult to precisely predict since it is related to aggregate bandwidth for the wire, data buffer sizes, compression ratios, and a large number of other factors. Radio communications, such as those between a remote vehicle and a command center, are often the bottleneck for moving data. Therefore, careful examination of the application latency, bandwidth and connectivity requirements will drive the selection of the video encoding scheme.

Multiprocessor remedy

With the processing requirements being identical for all video streams, it has been shown that it is advantageous to build fully symmetric processing systems. Video bandwidth is optimized, as each processor has its own dedicated memory and is handling only a single video stream. Connectivity is simplified because each stream can be treated as a separate Ethernet address. Finally, latency is minimized because each processor is handling only one major task with the assistance of dedicated hardware. Also, each processor can be identical and, therefore, run the same software. This further reduces design cost and the complexity of managing the deployed system.

> "To accommodate remote surveillance, vehicle autonomy, and situational awareness, there is a growing requirement for real-time, multichannel, remote video connections in the military sector."

To validate the multiprocessor technology, the power consumption of the Beyond Electronics CPU-iMX27-VME video/audio processing board was measured. The resultant approximate 11 W with four processors streaming video and the onboard Ethernet switch managing the data traffic through a single Ethernet connection exhibits a sharp contrast to the commercial system described earlier at 57 W. When packaged in a conduction-cooled VMEbus form factor, this multiprocessor system is clearly suited for operation within the constraints of most tactical mobile environments. **CS**

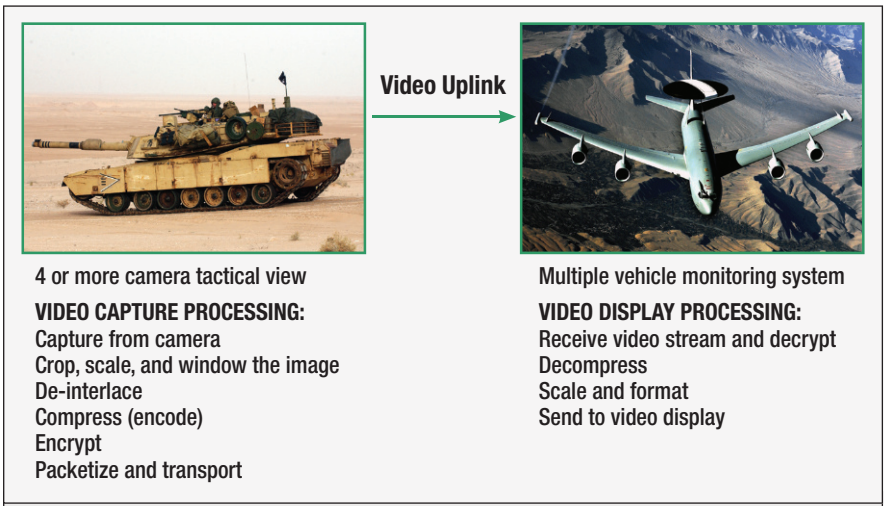

Figure 1 | Tactical video acquisition showing the processing steps for both camera and display

George Schreck is Chief Technical Officer with Beyond Electronics Corporation, where his current responsibilities include product development, product marketing, and corporate management. His experience includes 24 years of product design and marketing of high-reliability, embedded systems. He holds a Bachelor of Science from Lock Haven State University with additional studies in Physics and Computer Science at Lycoming College and Centenary College. He can be reached at gschreck@beyondelectronics.us.

Beyond Electronics Corporation
919-231-8000
www.beyondelectronics.us

Migrating to VPX – Without leaving VME behind

By Paul Mesibov

Adopting new technologies like VPX (VITA 46) doesn't have to mean abandoning previous solutions like VME. Smart product development strategies and use of XMC (VITA 42) modules enable the option to use both VME and VPX.

Many systems integrators and application providers serving the military/aerospace market have traditionally used VME system-level implementations. As COTS technology continues to make amazing performance strides, even ruggedized personal computer architectures have been deployed successfully. Board-level VME vendors wanting to leverage VPX (VITA 46) might wonder how to address the needs of this increasingly diverse "solution space" without multiplying development expenses by three or four times. Adding to this quandary is the fact that while VPX is on the verge of becoming another extremely popular element of this space, it might not yet represent a wide enough market for developers to feel completely comfortable investing in it.

So how can an investment in VME technology be preserved while upgrading to VPX at the same time? To figure this out let's review VPX and its spinoff, OpenVPX (VITA 65), and then some embedded system basics surrounding VME- and VPX-based systems. Next, a real-time beamforming system application example shows how XMC (VITA 42) modules are key in exploiting VPX without abandoning VME.

VPX: What it's all about

VPX is a 3U or 6U electromechanical arrangement that allows cards to communicate using gigabit signaling methods such as those used in PCI Express (PCIe). The VPX base specification describes the mechanical configuration, connector types, and pinouts but leaves a lot of system-related questions unanswered. This made multi-vendor interoperability difficult amongst VPX vendors. However, VPX's spinoff standard, OpenVPX, aims to solve VPX's system interoperability issues to perpetuate wide acceptance of VPX technology. OpenVPX contains standardized nomenclatures for describing card and backplane configurations plus much detail about signal usage and interconnection. And although implementing the new VPX standard into VME-based systems might seem daunting at first, it is entirely possible.

VME and VPX meet embedded system requirements

Most real-time embedded systems must perform at least two functions. They must provide command and control data paths so that system hardware can be initialized and controlled during runtime. They must also provide high-speed data paths such that the required real-time data rates can be met.

Traditionally VME met both of these demands. VMEbus transactions allow data transfer rates of 40 to 320 MBps. But VME's shared bus design made it difficult to guarantee that all bus transactions would enjoy the maximum achievable rates. Consequently, high-speed data paths have been implemented in VME systems using alternative methods: auxiliary backplanes, ribbon cable, fiber optic connection, mezzanine boards, or some combination of these.

However, the introduction of gigabit serial links and the attendant development of a rich set of communication protocols allow the same electromechanical configuration to implement the command and control functions as well as high-speed data transfer. The point-to-point nature of these serial links allows much better determinism when it comes to guaranteeing data transfer rates. This concept certainly favors VPX, which can support several different protocols simultaneously. For example, PCIe can be used for command and control while Aurora allows FPGAs to communicate directly. Alternatively, GbE can be used for command and control, and Serial RapidIO can be used for high-speed data transfer. Of course, the concept of communication protocol does not really apply to VME, which relies only on "raw" data transfer.

Folding in VPX while preserving VME

Many customers are interested in new technology developments while maximizing the performance of their typical long-term investment in VME-based systems. Thus, a popular strategy for preserving the VME investment is to implement crucial functions on a mezzanine board such as an XMC. It is then possible to implement a carrier for the various platforms such as VME and VPX to allow development on inexpensive desktop PCs. This suggests the need for a product development strategy that associates the core competencies of the board vendors with a mezzanine strategy that can be located easily in various backplane/signaling environments.

> "... While VPX is on the verge of becoming another extremely popular element of this space, it might not yet represent a wide enough market for developers to feel completely comfortable investing in it."

Many of today's military customers using VME equipment would very much like to leverage ubiquitous PC technology while moving rapidly toward the VPX backplane infrastructure. A beamforming application illustrates how this can be done.

An application example: Digital beamforming

The technical challenge in our digital beamforming application is shown in Figure 1. The basic approach is to use an arrangement of simple elements to build up a more complex system. If properly conceived and designed, we should be able to use either VME or VPX to implement our beamformer.

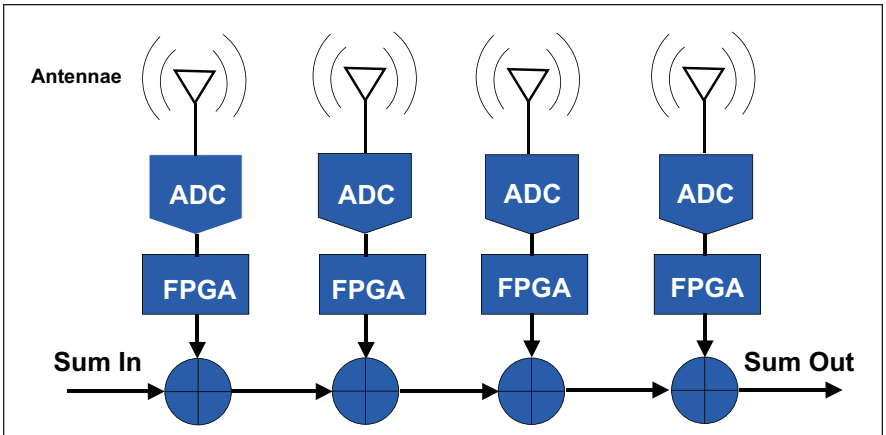

Figure 1 | It should be a straightforward effort to design a VME or VPX card to contain an ADC, FPGA, control interface, and sum-in and sum-out ports.

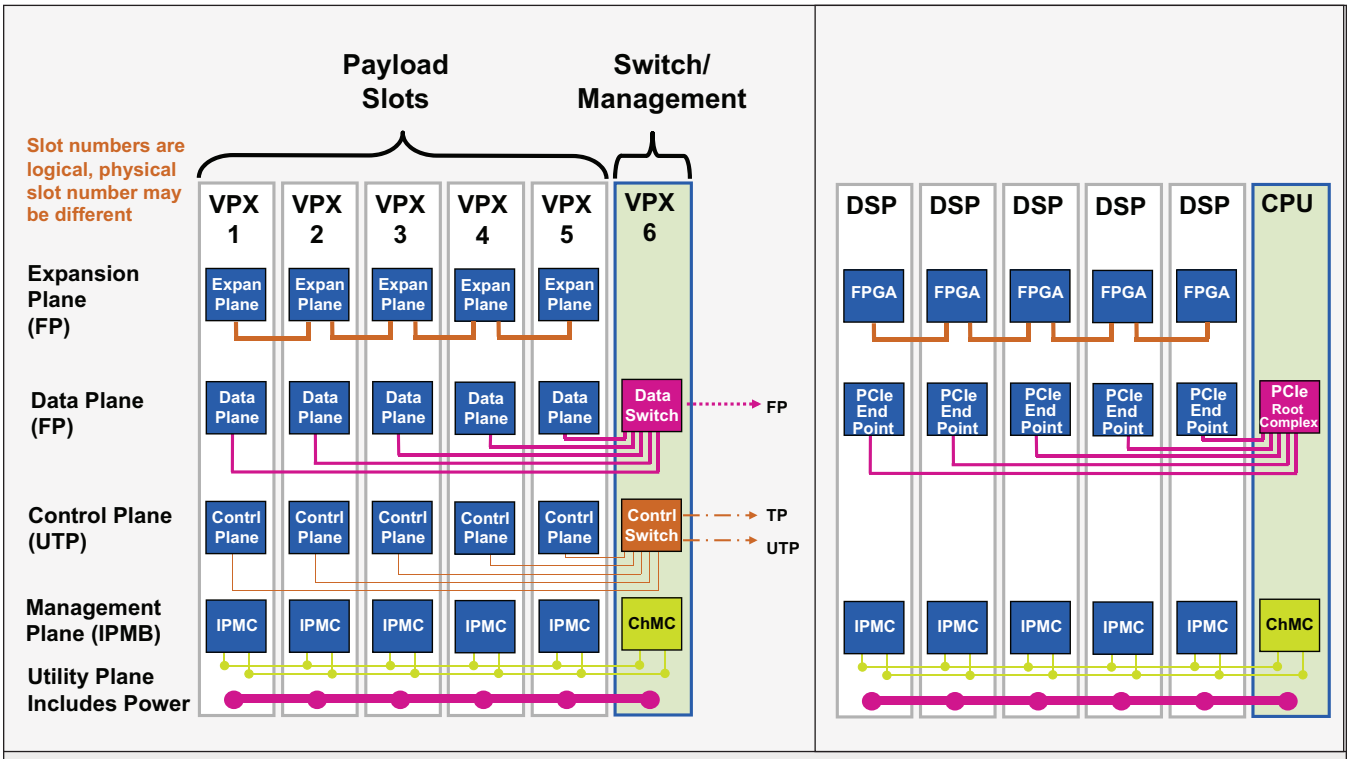

Figure 2 | Left side: A typical 3U VPX slot profile with multiple data paths. Right side: An illustration of how highly suitable VPX is to the beamforming application.

A digital beamforming system can be implemented using a cascade topology. In this case, an array of antenna elements is connected such that each antenna output is applied to some RF circuitry that then amplifies and converts to an intermediate frequency band suitable for digitization. The downconverted signal is then applied to an analog-to-digital converter that passes its output to an FPGA where the beamforming function is performed. This processing typically comprises digital frequency downconversion and a weighting and summation process.

Beamforming implementation

Since the basic elements of this system are ADCs and FPGAs, it should be possible to design a module that can be used in either VME or VPX systems. An XMC module is ideal in this case.

A particularly effective method to address this challenge is to use an XMC module that has four ADC channels and an FPGA to implement the necessary beamforming algorithm. The XMC module might also have a PCI Express interface to connect to a host CPU for command and control and a dedicated high-speed interconnect to implement the summation with the other FPGAs in the system. Xilinx FPGAs allow use of the Aurora protocol, a low-overhead, point-to-point gigabit link suited perfectly to this task.

The XMC form factor provides up to 16 gigabit serial links. Let's decide to use 8 for PCIe and 8 for Aurora. XMC carrier modules, such as Pentek's Model 4207, can assist in preserving the VME investment while updating the system to VPX. The carrier contains a flexible crosspoint switch allowing the two XMC modules to be cascaded with one another.

A 3U VPX XMC carrier, such as the Model 5300 from Pentek, can then be used to implement the beamformer in a VPX system. A typical 3U VPX slot profile with multiple data paths is shown in the left side of Figure 2. (Some of these paths are referred to in the OpenVPX nomenclature as Control Plane, Data Plane, or Expansion Plane.)

To see how this relates to the beamforming application, consider that the setting of the weights and initialization of the system are considered Control Plane functions; the receiving of processed data by the host CPU is a Data Plane function; and the Aurora FPGA interconnection is an Expansion Plane function. Meanwhile, the right side of Figure 2 shows how nicely suited VPX is to the beamforming application.

Investment preserved

In our beamforming example, the bulk of the development of precision analog-to-digital conversion circuitry and the FPGA-based signal processing is in the XMC module. Using XMC modules comprising the sophisticated signal processing necessary for digital beamforming and then employing them on VME or VPX carriers is a smart way to keep a foot in both camps. CS

Paul Mesibov *is cofounder of Pentek and serves as director of engineering. He has more than 26 years of electronics industry experience and is an active member of PICMG, VITA, PCI-SIG, SDR Forum, and IEEE. He holds Bachelors and Masters degrees in Electrical Engineering from the City University of New York. He can be contacted at paul@pentek.com.*

Pentek, Inc.
201-818-5900
www.pentek.com

Expand your military system horizons with CM Computers ATR

CM-ATR-25/35/45/S/SEF/HES/FAC

VPX, VME & cPCI True Military ATR Enclosures

SIXHEX™ SIX HEAT EXCHANGE

Up to 15 ATR COTS models available
PSUs from 300-1500 Watts
½, ¾ & 1 ATR size

CM-ATR-35/SIXHEX
¾ ATR, 7 Slot, 800W PSU

CM-ATR-35/SIXHEX

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 & US Air Force military 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 Chassis Catalog

CM Computer
True Military COTS Products

OpenVPX closes the integration gap between off-the-shelf VPX modules

By Jeff Porter

Until recently, widespread acceptance of VPX (VITA 46) suffered from the lack of a system-level interoperability specification. However, VPX systems integration just got a lot easier, thanks to OpenVPX (VITA 65).

With a module or board-level approach, the VPX (VITA 46) base specification and associated "dot specs" stop short of addressing system-level issues. Without any further definition, the VPX community was in jeopardy of fragmenting into a group of vendors with incompatible, standard-in-name-only products.

Fortunately, the VPX community united and created the OpenVPX (VITA 65) specification, paving the way toward cutting-edge VPX specifications within the rugged embedded marketplace. OpenVPX builds on the module-centric VPX specifications by providing a nomenclature for system integrators, module designers, and backplane providers to use when describing and defining aspects and characteristics of a system. OpenVPX accomplishes this via planes and profiles pertinent to the 3U and 6U form factors, providing a blend of interoperable flexibility for modern military VPX systems.

Telecom lends its planes to OpenVPX

The term "plane" is common within telecom architectures such as AdvancedTCA and MicroTCA and serves the same purpose within OpenVPX: to help categorize requirements while not changing the physical design structure of VPX. Accordingly, five planes have been defined within OpenVPX: Utility, Management, Control, Data, and Expansion (Figure 1). A closer look at each plane will provide understanding.

Utility plane

The Utility Plane encompasses the power, clock, and reset connections within the system. OpenVPX has aided interoperability within the Utility Plane by clarifying reset requirements throughout the system as well as defining power profiles for modules and development chassis. Clarification of reset requirements and the definition of power profiles were fundamental roadblocks in system interoperability prior to OpenVPX.

Management Plane

The Management Plane defines a hardware and software framework for platform-level tasks such as health monitoring, inventory management, event logging, fault detection, and fault isolation. OpenVPX supports the ongoing efforts of VITA 46.11 in defining optional system management components at the payload module level (IPMCs) and the chassis level (ChMCs). For the first time within VPX, an industry-standard IPMI-based infrastructure is being defined for the collection and distribution of module sensor data. VITA 46.11, like OpenVPX, leverages from the AdvancedTCA and MicroTCA specifications while removing focus from telecom-specific features, such as hot swap, not required within the rugged embedded market.

Control Plane

The Control Plane defines a system-level communication path for control traffic. OpenVPX has focused on using the 1000BASE-BX SERDES GbE protocol for the Control Plane. This is especially beneficial with the real-estate-constrained and pin-limited 3U profiles, as the SERDES-based GbE protocol reduces the pin count by half and removes the requirement for large on-card magnetics. For many systems, the GbE Control Plane can also take the place of the longstanding VMEbus.

Data Plane

The Data Plane provides a high-throughput mechanism based on switched fabrics for transferring data between system peers. This can be accomplished by using a distributed architecture or utilizing a switch to provide a centralized location through which modules can route communications. With the standardization of the high-speed serial switched fabric Data Plane, OpenVPX provides the bandwidth and forward compatibility to support the data throughput requirements of both current and future applications.

Expansion Plane

The Expansion Plane presents a high-throughput mechanism for transferring

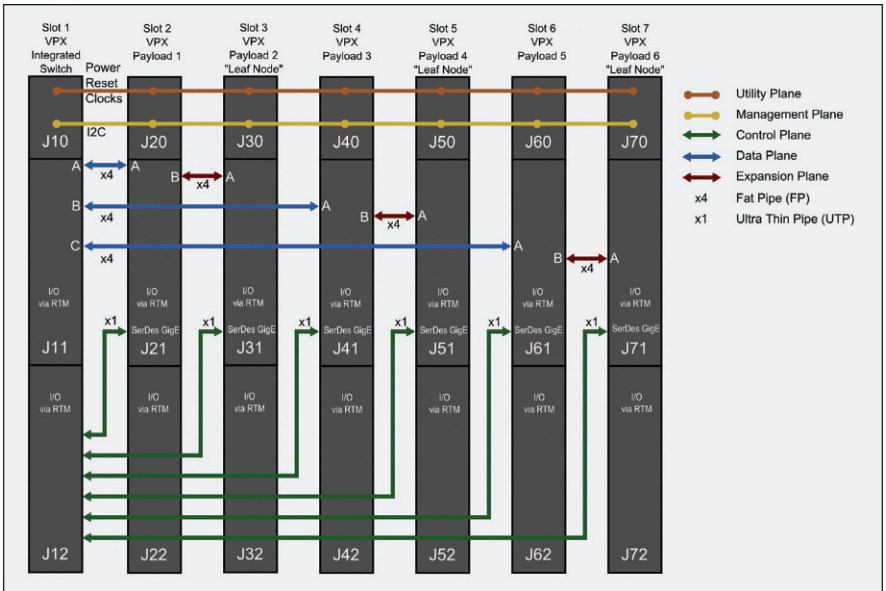

Figure 1 | A representation of the Utility, Management, Control, Data, and Expansion planes in a representative 3U OpenVPX backplane.

data, intended for sharing data between two specific entities within a system. The Expansion Plane provides a standard mechanism for implementing functions such as moving data to and from a coprocessor or off-load engine as well as moving data from a sensor interface module to a processor module.

Profiles smooth multi-vendor module integration

Before OpenVPX's clearly defined standard for backplane fabric pinouts, integrating a system utilizing VPX modules from multiple vendors was often a costly and time-consuming effort. The newly defined OpenVPX profiles for the slot, module, switch, backplane, and development chassis ease system integration. OpenVPX defines *slot profiles* as the physical location and logical definition of the planes from the backplane's perspective, while *module profiles* define protocol requirements.

To accommodate flexibility of development systems, encourage innovation, and assist the future evolution of VPX, OpenVPX has identified a number of "User Defined I/O" regions within the module profiles. While the location of "User Defined I/O" on the connectors has been set by the module profiles, the specific allocation of these interfaces is characterized by the module vendor. This approach focuses the standard on inter-module fabric communication. Since access to these "out-of-the-box" I/O interfaces within development systems is provided with the use of a cost-effective Rear Transition Module (RTM) or through front-panel connectors on the module itself, there is little advantage in specifying required I/O pinouts for modules within the development environment. Since the majority of deployed environments require the customization of the backplane to accommodate rugged box-level connectors, there is also no practical benefit to locking down user I/O pinouts for deployed systems.

OpenVPX also defines power and thermal profiles to provide chassis and module designers standard guidelines upon which to base designs. Realizing the system constraints and distribution (I2R) losses associated with the lower-voltage 5 V-based legacy systems, OpenVPX has taken the forward-looking stance of defining a higher-voltage 12 V-based power profile for 3U as well as 6U. Let's take a look at more OpenVPX profiles' advantages related to the 3U and 6U form factors specifically.

3U OpenVPX Profiles: SWaP focused

Amidst the military's clamor for SWaP-constrained wares, OpenVPX defines the multipurpose 3U Payload Slot Profiles "SLT3-PAY-2F2U-14.2.3" and "SLT3-PAY-1F1F2U-14.2.4." These particular profiles maximize the configurability of the pin-constrained 3U module by leaving about one-third of P/J1 pins open as "user defined." This provides an excellent location to route mezzanine I/O or external I/O from the VPX module. While defining a number of separate profiles increases variability within the ecosystem, it also serves to provide flexible solutions for the system integrator. Additionally, the similarities between profiles allow module vendors to design a single module that can comply with multiple module profiles.

Another important achievement for OpenVPX is that for the first time within the VPX domain, standard switch-based profiles have been fully defined: a key requirement for the proliferation of high-speed switched fabrics such as PCI Express, GbE, and 10 GbE within VPX. Additionally, the integrated 3U switch slot "SLT3-SWH-6F6U-14.4.1" profile focuses on reducing system slot count and overall system SWaP by defining a single switch slot that performs the duties of both Data Plane and Control Plane switches. The integrated switch includes the capability of providing an external access point to the system via GbE or 10 GbE through the user-defined pins at the bottom of P/J2. If the integrated

switch is also designed to accommodate a mezzanine, it provides a centralized location for the Data Plane Root Complex and Control Plane management. The Extreme Engineering Solutions (X-ES) XChange3012 (Figure 2) is an example of an OpenVPX 3U integrated switch.

In addition to the profiles discussed, yet another new concept for VPX brought forward by OpenVPX is a standardized storage module profile. The storage module profile "SLT3-STO-2U-14.5.1" provides VPX with a standardized 3U storage module pinout to enable the production of commodity VPX storage modules. This, in turn, helps reduce lead time and costs of this critical component within many systems.

6U OpenVPX profiles surpass VME

For 6U VPX, OpenVPX has further advanced the performance capabilities over legacy 6U form factors such as VME and even CompactPCI. These advancements include defining P/J2 on 6U modules for something other than legacy VME support: Previous VPX specifications reserved P/J2 on 6U modules for VME, which limited the performance impact of VPX by tying up precious pins that could otherwise be used by high-speed differential I/O, but that is not the case with 6U OpenVPX. SERDES GbE is additionally defined as a Control Plane interface with 6U OpenVPX. By focusing on SERDES GbE as opposed to 1000BASE-T for the control plane interface, 6U systems that include a Control Plane switch will benefit from increased pin availability and more available board real estate through the removal of on-card magnetics.

Figure 2 | The Extreme Engineering Solutions (X-ES) XChange3012 OpenVPX 3U integrated switch

The 6U payload profile "SLT6-PAY-4F1Q2U2T-10.2.1" is an example of a multipurpose 6U profile that utilizes the advantages of SERDES Control Plane ports, provides a large amount of user-defined I/O, and allocates P/J2 for high-throughput serial I/O. Another advantage of this particular profile is that a module designed in this profile can also be designed to comply with several other 6U OpenVPX payload and peripheral profiles.

The right stuff: Interoperability and flexibility

OpenVPX does not take the "one-profile-fits-all" approach that has limited innovation within other form factors. Instead, it addresses major interoperability issues while allowing for flexibility within the system, as enabled by its planes, backplane profiles, and flexible module profiles featuring user-defined I/O. **CS**

***Jeff Porter** is a Senior Systems Engineer at Extreme Engineering Solutions (X-ES). Jeff received his BS in Electrical and Computer Engineering from the University of Wisconsin, and has been an engineer at X-ES since its inception. He can be contacted at jporter@xes-inc.com.*

Extreme Engineering Solutions
608-833-1155
www.xes-inc.com

Stroke of Genius

Introducing the NEW Cobalt Series of Virtex-6 Modules!

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

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

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

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

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

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

OpenVPX: The future is now

By James M. Goldenberg

VPX was the natural successor to the venerable VMEbus architecture, providing significant advances in performance while providing straightforward migration. The VPX specification did not, however, address interoperability between boards – and the OpenVPX specification addresses this. It is, however, in essence and by design, a work in progress.

VME has long been the *de facto* open standard for the embedded computing industry. And VPX/VITA 46 (in addition to its more VME-like alternative, VXS/VITA 41) was one of the first VITA open standards for embedded computing boards to capitalize on emerging serial link protocols. VPX was not, however, and nor was it intended to be, an open architecture for complete systems solutions. The flexibility of VPX was both a strength and a weakness. On one hand, it allowed manufacturers significant leeway to develop innovative technologies. On the other hand, its flexibility in the use of pins for serial link communications, its mapping of those pins to various protocols, and the lack of system-level thinking meant that VPX boards from different manufacturers did not easily interoperate.

As VPX attracted more and more attention, with myriad defense vendors introducing broad ranges of products, interoperability became an increasing issue. A number of vendors, including GE Intelligent Platforms, responded in the form of the OpenVPX (VITA 65) initiative, which was designed to leverage the work done on each of the individual "dot" specifications within the VPX standard to create a comprehensive set of definitions for system-level interoperability (Figure 1). Where individual manufacturers had used VPX's plethora of pins as each saw fit, OpenVPX looked to harmonize the assignment of those pins – in effect, defining pinouts. VPX was a bottom-up approach to providing a successor to VMEbus: OpenVPX is a top-down approach that still leaves room for customization.

How is OpenVPX organized?

The OpenVPX specification has been constructed in sections that break down the standard, using a top-down approach; this makes it a different type of document from the typical specification. It is also clearly a "living" document, allowing the opportunity for additional profiles – be they slot, module, or backplane – to be added without major modification of the document. For example, slot profiles – currently in sections 6, 10, and 14 of the document – are written such that, if a new 3U slot profile is needed, it can be added to the end of a subsection within section 14. This can be done because much time and effort were put into understanding the various slot types that might be used and how they can be summarized: as payload, peripheral, switch, and miscellaneous (Figure 2).

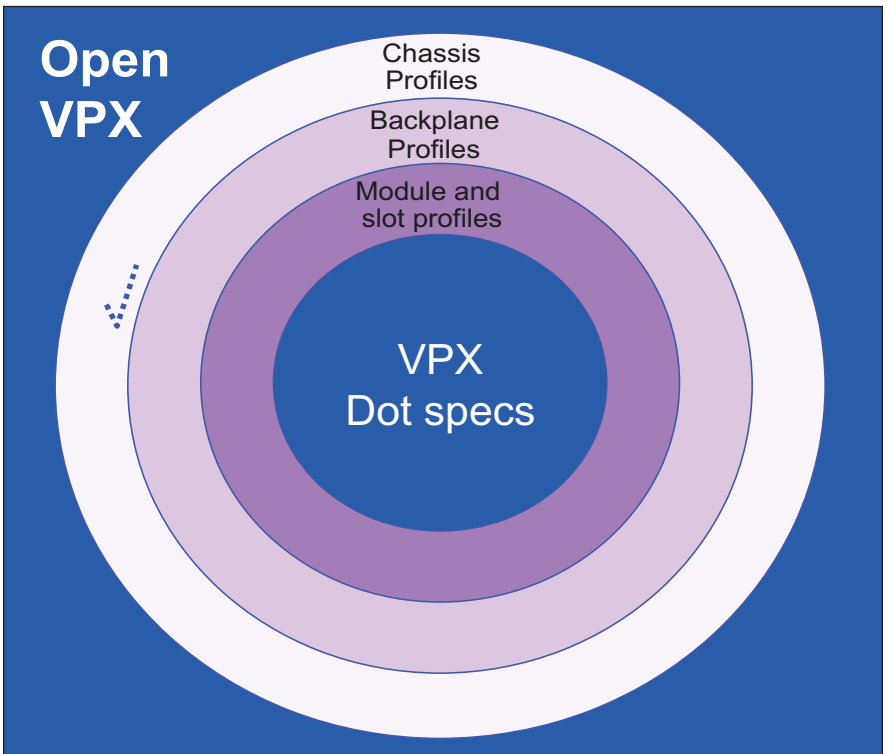

Figure 1 | "Dot" specifications are at the heart of a comprehensive set of definitions for system-level interoperability.

The "miscellaneous" category is significant because it is used to standardize pinouts that would otherwise normally not be set by the OpenVPX specification. For example, the storage slot profile defines two ultra-thin pipes that shall be used by vendors developing compliant storage modules of whichever type or capacity.

Whether of the payload, peripheral, switch, or miscellaneous type – which apply equally to 3U and 6U implementations – slot profiles use the concept of pipes and planes. Pipes are, in effect, communications pathways through the system. Pipes can be "fat" (four links of 4x Tx pairs and 4x Rx pairs), "thin" (two links of 2x Tx pairs and 2x Rx pairs), and "ultra-thin" (single link comprising 1x Tx pair and 1x Rx pair), among others called out in the specification. Note that pipes do not define the protocol used on them; this is specified by the module profile.

Meanwhile, OpenVPX system architectures can also include multiple planes, where a plane defines a type of traffic depending on its characteristics and requirements. There are five planes: management, utility, control, data, and expansion. Within each of the slot profiles, all serial link communications paths between slots are defined as either on the Data Plane or Expansion Plane: Those that would, in the traditional VME environment,

have been on the VMEbus might be on either the Data or Expansion Plane. High-speed interfaces that in VME used special backplane overlays on user I/O pins such as RACE++ or SkyChannel are now defined on the Data Plane. It is important to understand that, at the slot level, no protocols are specified – only the number and size of links. For example, one such 3U slot profile is SLT3-PAY-2F-14.2.7. This describes a Slot Profile (SLT) for a 3U (3) Payload (PAY) with two Fat Pipes (2F) as described in section 14.2.7 of the document. Module profiles and backplane profiles are similarly defined.

A module conforming to a particular definition will always be compatible with a chassis slot defined as being designed to accept that type of module. Every OpenVPX-compliant VPX module has a defined set of backplane signal assignments, mapping protocols to its ports – while OpenVPX backplanes have slot profiles designed to support specific types of modules. The slot profile provides mapping of ports to a slot's backplane connectors. A backplane profile specifies the type and number of slots along with the topology of the channels and buses that connect the slots. As of this writing, profiles had been defined for 22 3U VPX modules, 13 6U modules, 11 3U backplanes, and 13 6U backplanes.

Still a need for customization

There is, however, one area of the OpenVPX standard intentionally left to the discretion of each vendor: user-defined pins, recognizing that every program's I/O requirement will be different. For the VMEbus community, these have long been a fundamental and advantageous element of the VME architecture, and this flexibility is retained by OpenVPX. The implication, of course, is that this will often mean that 100 percent plug-and-play operation is not possible in the way that it is possible with, for example, the PC.

If the OpenVPX standard sounds prescriptive, that is because it is intentionally so: Without a strict set of definitions, the goal of interoperability cannot be obtained. It is also important to understand that the primary goal of OpenVPX is to support more rapid, more cost-effective development of applications: It implicitly recognizes that deployment is a very different consideration than development. Only development chassis and backplanes are standardized, with Rear Transition Modules (RTMs) providing user I/O. The expectation is that a custom backplane will be used for deployment, especially where the target system will be deployed in a harsh environment.

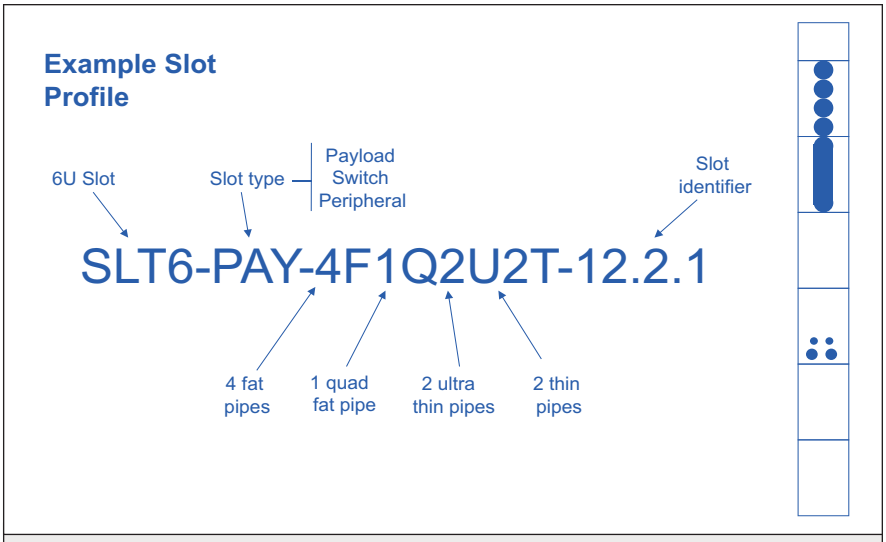

Figure 2 | An example slot profile

Typical rugged deployed systems are usually size- and weight-constrained, with the emphasis resting on obtaining the maximum performance from the minimum space envelope. This is especially true in land-based or airborne fighting vehicles. It is also often true for naval applications as more and more equipment is added to existing platforms. These systems will rarely, if ever, use RTMs as these take up far too much space and are susceptible to shock and vibration; however, such systems have complex backplanes that contain the routing for all functionality required by the application, including all user-defined I/O. As such, OpenVPX is not – yet – typically applicable to the deployed environment. However, by specifying development backplanes and chassis, OpenVPX means that the long and complex effort of developing application software can begin much earlier in the development process while the custom chassis and backplane are developed. This will shorten the time to deployment.

Currently, a limited number of backplanes have been defined for use in a limited number of chassis. That might sound like a restriction, but, in fact, it is not: Today, the OpenVPX standard is designed to apply almost exclusively to development systems in which user I/O goes to an RTM; this is where most of the custom application-specific requirements come from. As such, the number of payload modules interacting with switch and peripheral modules can be relatively easily standardized to a smaller set if the complexity of all the various user I/O types is not addressed on the backplane, but rather in an RTM ignored.

Going forward

In the near term, we will continue to see VPX specified in requests for quotation. Just as there are custom backplanes for VME in deployed systems, so there will be custom backplanes for VITA 46/VPX and VITA 65/OpenVPX. VPX and OpenVPX are joined at the hip and will live together – but as the OpenVPX specification lives and grows, it will assume increasing significance.

OpenVPX is an initiative that should be welcomed by the embedded computing community. It is, by design, a work in progress. An inherent strength of the OpenVPX initiative is that it explicitly recognizes the need for, and encourages submission of, new module and slot definitions that can be incorporated into the standard, such that it will progressively embrace a broadening range of possible scenarios. Ratification by ANSI later this year is the beginning of the journey, not the destination. **CS**

James M. Goldenberg *is Chief Engineer at GE Intelligent Platforms and has been a contributing editor to the OpenVPX initiative since it began early last year. He is also the current editor for VITA 46.9 and contributed to the VITA 41 SERDES specification. He has more than 25 years of experience in the computer industry. Prior to GE, he was VP of Engineering at SBS Technologies and Sky Computer. Jim has a Masters in Computer Science from Indiana University. He can be contacted at jim.goldenberg@ge.com.*

GE Intelligent Platforms
505-796-1763
www.ge-ip.com

HALT/HASS testing drives product reliability

By Jon Semancik

The critical nature of military electronic and other critical systems demands that fielded products are as reliable as possible. Highly Accelerated Life Testing and Stress Screening (HALT/HASS) can quickly uncover problems associated with design and production capabilities. This early problem discovery results in more reliable products, reduced life-cycle support costs, and shorter time to market. Understanding this test environment, the application challenges associated with test equipment selection and implementation, and the intricacies of measurement challenges will ensure that critical performance data is available when needed.

HALT and HASS rely on techniques that shorten the time required to identify potential causes of failure. These techniques are performed by applying much higher stresses than exist in actual product use, forcing failures to occur in significantly less time than under normal conditions.

HALT is performed during the product development cycle with the aim of shortening time to market. Identifying and correcting potential design issues early in the cycle can significantly reduce both overall development costs and schedule delays. HASS focuses on uncovering issues that might occur during the manufacturing process such as machine, process, or workmanship problems, which can result in expensive rework and scrap expense.

Since hundreds of products can be aged and monitored simultaneously, the challenge for test engineers is the large number of test channels needed, environmental conditions such as electrical noise that can be present in a production environment, as well as the cost and reusability of test systems. Therefore, designing an effective test system requires a clear understanding of the environment, the performance characteristics of the test equipment, and the implementation and measurement challenges.

HALT/HASS fundamentals

In HALT, temperature and vibration stress conditions are used during product development to find weak spots, both in the design and in the product's planned fabrication processes. Test stimuli can include humidity, thermal cycling, burn-in, overvoltage, voltage cycling, and anything else that could logically expose defects. This phase of testing requires only a few units and a short testing period to identify the fundamental limits of the technology being used. Generally, every weak point should be identified and redesigned to meet the product specifications.

In production, HASS uses the highest possible stresses, frequently well beyond the qualification level. Also proof-of-screen techniques must be used to protect good product (that is, to identify defects without damaging good product). Keep in mind that HASS usually is not as effective unless HALT is done earlier in the process, because fundamental design limitations will tend to restrict HASS stress levels. Typically, a population of products will exhibit reliability characterized by a bathtub-shaped curve (see Figure 1) with three distinct failure-rate regions.

The first region of the curve is infant mortality, which has a decreasing failure rate and is associated with built-in (not designed-in) defects. These are the types of defects that can often be identified by HALT testing. The amount of time required for a given HASS test is determined by the infant mortality region of the reliability curve. Generally speaking, the higher the stress, the sooner the failures, the narrower the infant mortality region, and the shorter the HASS test period.

As mentioned, test profiles can include temperature and vibration as well as other stimuli designed to reveal a wide range of problems. Some common profiling examples and test objectives are identified below:

Temperature:

- Determine the minimum and maximum product operating temperature limits
- Apply extreme temperatures to accelerate the aging process of the Device Under Test (DUT)

Humidity:

- Determine the effects of high/low humidity on the DUT (corrosion, electrostatic discharge, ESD damage)
- Indentify latent shorts on Printed Circuit Boards (PCBs) that are related to humidity

Vibration:

- Determine the maximum product operating vibration levels
- Find mechanical defects early in the cycle

The test parameters during any of these HALT/HASS profiles can vary but are usually limited to the signals that indicate whether or not the product is functioning properly.

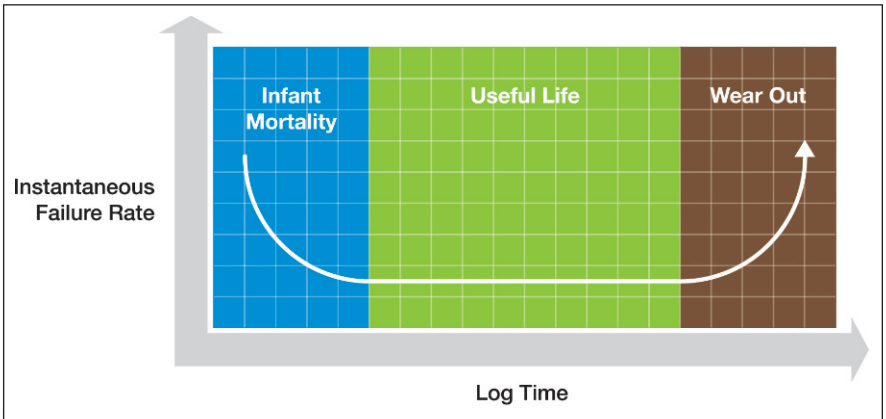

Figure 1 | Reliability curve: Typically, a population of products will exhibit reliability characterized by a bathtub-shaped curve with three distinct failure-rate regions.

Test equipment considerations

HALT/HASS are conducted in environmental chambers where high or low temperature, vibration, and other stress conditions can be applied to the DUTs. In many cases, additional stresses not associated with the chamber can also be applied, such as full power (burn-in) and overvoltage conditions.

Many considerations are associated with the selection of test hardware including channel count, signal input levels, and measurement speed and bandwidth, in addition to cabling and connectorization. Multiple DUTs are often connected through a multiplexer to maximize testing efficiencies. With appropriate characteristics, a multiplexer can increase throughput, reduce test equipment costs, maximize test equipment utilization, and improve consistency and reliability.

When a DC stimulus must be applied to the DUTs and their response measured, synchronization and triggering are issues associated with speed and accuracy. For example, parasitic capacitance in cables and connectors makes it necessary to allow for signal settling time before measurements are made; otherwise, inaccuracies can occur. This is often seen when long cable lengths are used to connect devices in a test chamber. Typical mechanical relay settling times are at a rate of <5 mS; however, additional time will be needed to compensate for the cabling effects.

Test system implementation

HASS burn-in testing is common for switching power supplies, and it is not unusual to test hundreds of these products simultaneously within a chamber. The challenge is the high levels of electrical noise produced by such a large number of switching supplies. Power supplies for computers typically have four to six voltage outputs that are usually measured while the temperature in the environmental chamber is cycled between upper and lower limits, and the power supply output is repeatedly cycled on and off.

A Digital Multimeter (DMM) is connected in parallel with the load resistance to measure the power-supply voltage. The load resistance is chosen to simulate that found in the final application, or can be chosen to force the power supply to its maximum output capacity for testing purposes. Power supply output cycling, which increases the level of stress, is a preliminary test before burn-in; this helps weed out power supplies with major problems. To capture this failure data, the output voltage can be measured each time the supply is turned on.

Precision switch/measure and I/O technology is ideal for both the voltage and temperature measurements required in this example. VTI's EX1200-3048, a 48-channel multiplexer card, is a member of VTI's EX1200 Series (Figure 2) and has two connections (HI and LO) for each power supply. Additionally, the same terminal configuration can be used for thermocouple measurements because of built-in automatic Cold-Junction Compensation (CJC). Once the field terminations are complete, the instrument is configured using the Ethernet-based LXI communications interface; this platform is scalable and can expand with application requirements.

Figure 2 | Precision switch/measure and I/O technology, such as VTI's EX1200 Series, is ideal for both the voltage and temperature measurements of HALT/HASS testing.

Measurement challenges

Environmental noise is a prime error source during HALT/HASS testing, and placing hundreds of switching power supplies in an environmental chamber creates very challenging measurement conditions. Switching power supplies radiate high-frequency noise, and if the ground connection is noisy, conventional data-acquisition systems will be unable to make satisfactory measurements. Noise not only degrades the overall quality of data, it can also result in good product being identified as unacceptable; this will result in lower throughput and higher costs.

A commonly used approach to reduce the effects of noise spikes involves varying the integration time of the instrument. Integration time is typically determined by the number of Power Line Cycles (PLC) used, based on 60 Hz, for example:

0.1 PLC = 1.667 mS

1 PLC = 16.67 mS = 1/60 Hz

10 PLC = 166.7 mS

A longer integration time provides the best common-mode and normal-mode noise rejection. This technique will result in a measurement that is less susceptible to noise but can significantly increase the measurement time; for example, scanning 288 channels at 10 PLC would take 48 seconds to complete (288 x 166.7e-3 = 48).

Relay wear is another consideration because mechanical relays have a finite life; typical mechanical relays are rated at $1x10^7$ operations, closure time <5 mS, with a path resistance of less than 500 mΩ. Relay contacts will eventually degrade and can affect the quality of the measurement due to increased contact resistance and thermal EMF; features such as built-in counters can monitor contact closures to identify units needing maintenance. Another alternative involves using solid-state relays. Switching times are fast, less than 500 microseconds, but path resistance can be as high as 10 Ω; this might not be suitable for all applications.

HALT/HASS paves road to reliability

HALT/HASS implementation does involve an initial investment and planning, but this will ultimately result in products that are more reliable and less costly to produce. Implementation costs can be minimized by selecting test equipment that is easily connected, reconfigured, and scalable, and measurement challenges can be thereby reduced. And ultimately, customer benefits such as improved uptime and performance cannot be underestimated. **CS**

Jon Semancik *is Marketing Director at VTI Instruments Corporation. He is a U.S. Navy veteran and has held engineering and project management positions on numerous programs in both the military and commercial sectors. Jon joined VTI in 2002, where he held positions as Product Support Engineer and Region Sales Manager. Jon is also on the LXI Consortium's board of directors. He holds a BSEE from Fenn College of Engineering and an MBA from the Weatherhead School of Management, Case Western Reserve University. He can be contacted at jsemancik@vtiinstruments.com.*

VTI Instruments Corporation
949-955-1894
www.vtiinstruments.com

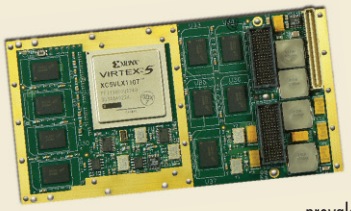

XMC card doubles the span of memory lane

As the years march on, memory shortcomings become increasingly prevalent in human beings. But thank goodness that the technologies spawned by the embedded computing industry are just the opposite: As time forges on, technology steps up and provides increasingly more advanced memory capabilities – even up to twice as much. Case in point: Curtiss-Wright Controls Embedded Computing's (CWCEC's) MM-6171 buffer memory XMC card, which the company reports as having stepped up its game from providing 4G to 8G memory for rugged air- and conduction-cooled mil apps such as image processing, SIGINT, and radar, for example.

Not only that, MM-6171 is designed to render volatile, deep storage capacities, and its two primary advantages and enablers are: 1) High-speed and faster bidirectional bandwidths enabled by serial fabric interfaces such as x4 Serial RapidIO or x4 PCI Express; and 2) Memory connection to high-speed signals via a Virtex-5 LX110T FPGA sitting on the memory card and interfacing to DDR2 of 1 to 2 GB, with 720 bit-wide memory arrays, a data path of 64 bits, and ECC. [Note that the FPGA is used as a memory controller only and is not intended for hosting User Programmable Logic (UPL).] Beyond the FPGA, the MM-6171 additionally sports a full-featured DMA engine, and a VxWorks 6.x device driver is available.

Curtiss-Wright Controls Embedded Computing
RSC# 36725 • www.cwcembedded.com

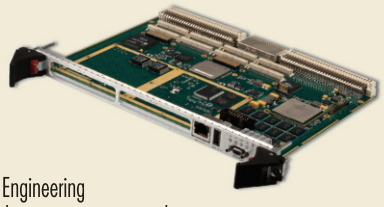

6U VME SBC rides in full-throttle style

While some industry pundits still predict the eventual demise of VME, wares like Extreme Engineering Solutions' XCalibur1531 6U VME single board computer just go to show that old faithful is still alive and kicking. And even thriving: This SBC is souped up for today's high-performance mil apps, kicked into high gear with a Freescale MPC8572E PowerQUICC III processor fully loaded with two e500 Power Architecture cores speeding along at rates up to 1.5 GHz. Also contributing to XCalibur1531's high throttle is its memory: Up to 256 MB NOR flash (with redundancy) and 16 GB NAND, in addition to up to 4 GB (two channels) of DDR2-800 SDRAM including ECC.

Meanwhile, I/O detailing includes two XMC/PrPMC interfaces, three USB 2.0 ports, and a SATA 3.0 Gbps port. The SBC also lends support for XMC I/O, PMC I/O, four GbE ports, plus P2/P0 backplane connectors and/or RS-232/422/485 serial ports on the front panel. And the behind-the-scenes driver, of course, is the Board Support Package (BSP). This one's got four BSP courses to race on: Wind River VxWorks, Green Hills INTEGRITY, Linux, or QNX Neutrino. And XCalibur1531 bodes well against strong headwinds – or other aspects of harsh environments – with 5 ruggedization levels, including rugged air-cooled and conduction-cooled, the latter of which operates at -40 °C to +85 °C. *[If you're wondering why this product looks a little familiar, it also made a cameo appearance on this edition's cover.]*

Extreme Engineering Solutions, Inc. (X-ES)
RSC# 44113 • www.xes-inc.com

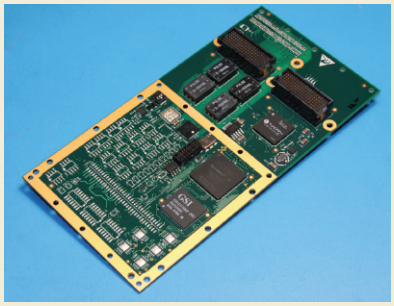

Avionics XMC module: "Wright" for days gone by?

Avionics technology's depth and breadth has risen to an all-time high, leaving Orville and Wilbur Wright's comparatively basic flying machine technology in the proverbial dust. Though the brothers' 1903 flight made history and serves as the basis of modern flight, that historic day would have undeniably occurred sooner had the Wright brothers had access to myriad embedded wares used today by Boeing, NASA, and others. One such technology is the XMC form factor, which was recently incarnated once again by GE Intelligent Platforms in the form of its RXMC-1553 high-density, rugged MIL-STD-1553 XMC module.

Ideal for avionics testing and simulation applications, the module provides 100 percent monitoring for fully loaded buses. Additionally, the conduction-cooled RXMC-1553 has either one or two channels. Each channel features 1 MB memory and is designed to flexibly accommodate MIL-STD-1553 A/B Notice II applications. RXMC-1553 also features multifunction interfaces that can work via a simultaneous bus controller, along with bus monitor functionality and up to 31 remote terminals. On the other hand, single-function interfaces do everything that the multifunction interfaces do; however, each major operational mode must operate one at a time. The RoHS compliant module also renders 45-bit microsecond message time tagging, error injection/detection, automatic/manual RT status bit and mode code responses, and an IRIG-B signal receiver/generator with GPS synchronization.

GE Intelligent Platforms • RSC# 43932 • www.ge-ip.com

VPX backplane thwarts testing challenges

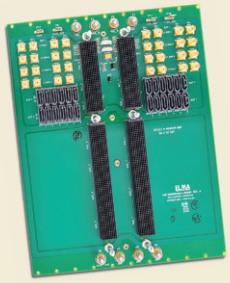

VPX has inspired the VME industry with its small form factor and extreme ruggedness, even serving as the catalyst for the OpenVPX (VITA 65) movement. But one thing that has been less than inspiring with VPX – or at least proven itself quite a challenge – is that of testing VPX designs. However, Elma Bustronic Corporation's 2-slot Test Backplane aims to thwart these VPX testing snafus. And, of course, the method depends on the challenge at hand.

One VPX testing challenge occurs because VPX cards are typically developed with an RTM that cannot access J1 fabric signals, meaning that engineers can't test J1 fabric signals at the same time the RTM module is accessing I/O signals. However, the 2-slot Test Backplane enables access to and connection with primary J1 signals without having to sacrifice simultaneous RTM use. Secondly, it has typically been necessary to design custom backplanes to connect fabric signals amongst multiple VPX blades heretofore; however, the 2-slot Test Backplane enables users to connect at least two blades prior to investing in a custom backplane. Larger chassis designs can also benefit, as it is feasible to connect all the J1 primary fabric into any desired serial topology. And a bonus capability: It can be used to test both 3U and 6U VPX designs.

Elma Bustronic • RSC# 45094 • www.bustronic.com

6U VME board melds old and new

On one hand, there are the classic technologies like VME that have been around for years. On the other, there are new processors and technologies such as Intel's Core i7. Though the latest technologies quickly gain more notoriety than the tried-and-true technologies, it doesn't have to be an "either/or" proposition. Concurrent Technologies apparently found the philosophy appealing, as manifested in its VP 717/08x family of 6U VME boards, which feature ... a Core i7 processor. Since VME is known to be hale and hearty, the VP 717/08x has followed suit, with extended-temp versions available now and ruggedized air- or conduction-cooled versions to soon follow.

Designed to handle CPU-intensive processing applications for the homeland security, defense, transportation, and industrial control industries, VP 717/08x is based on an Intel Platform Controller Hub (PCH) in addition to the Mobile Intel QM57 Express chipset. Supporting the classic VITA 31.1 GbE on a VME64x backplane standard, other highlights offered by the VP 717/08x include two XMC x8 PCI Express or 100 MHz PCI-X PMC sites, optional expansion for two additional PMC sites, 2048 x 1536 analog graphics, 2x RS-232/422/485, SATA300 and EIDE rear I/O interfaces plus an optional SATA300 2.5-inch disk drive … not to mention 2x GbE ports and 64 MB application flash memory.

Concurrent Technologies Inc.
RSC# 45092
www.gocct.com

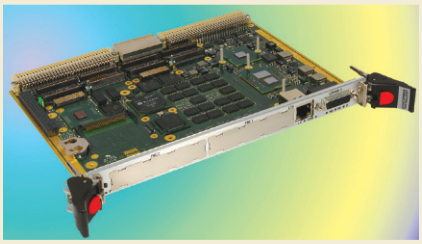

Editor's Choice Products are drawn from OSM's product database and press releases. Vendors may add their new products to our website at http://submit.opensystemsmedia.com and submit press releases at http://submit.opensystemsmedia.com. OSM reserves the right to publish products based on editors' discretion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

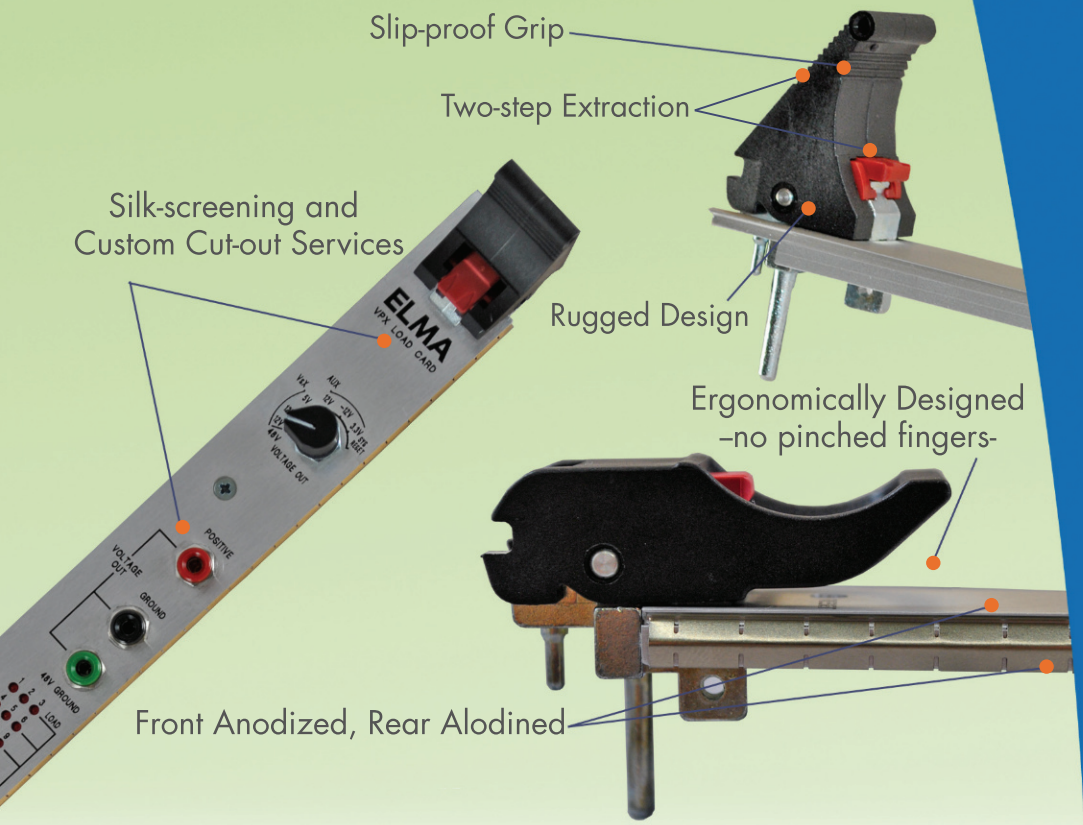

The Industry's Choice in VPX Handles and Panels

Elma's handles are the industry favorite – and for good reason. The rugged design is ergonomic, reliable, and allows simple assembly with just one screw. Our two-step latching process with push-button release makes hot swap extraction a breeze. Choose from low profile versions that don't protrude out, long styles for extra leverage, or our "famous" IEEE ergonomic design. Elma's VPX panels are specially designed for VPX spacing and board/panel assembly. Call Elma for a sample today!

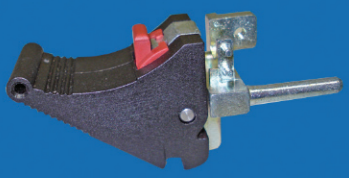

IEEE Handles

- Used for VPX, VXS, VME64x, cPCI, and custom
- Hot swap and non hot swap versions
- Two-step latching
- Easy assembly

Low Profile Handles

- Used for VPX, VXS, VME64x, cPCI, and custom
- Low profile, handle close to panel
- Increased leverage, ergonomics
- Hot swap and non hot swap versions

Panels

- VPX sizes of 1.0" and 0.80" (1.2" available upon request)
- Offset spacing for PCBs
- Custom cutout and coating options
- Silk screening and digital printing options

ELMA
Your Solution Partner

USA Elma Electronic Inc.
Phone: 510.656.3400 Email: sales@elma.com
Web: www.elma.com

Create Your Own PMCs

Custom I/O as Easy as 1, 2, 3!

The Technobox ***Micro Mezzanine System***™ is based on a simple idea – provide embedded systems designers with a foundation for innovation and flexibility. Provide a highly-granular, modular architecture featuring a range of configurable FPGA-based carrier boards and an extensive variety of micro mezzanine Electrical Conversion Modules (ECMs) that can be assembled in thousands of combinations. Provide an environment in which a designer can create an array of unique, *future-proofed*, board-level solutions. But without the costs normally associated with custom board development and manufacture, while speeding development and reducing time to market. It's the logical next step in mezzanine systems.

1 Select an FPGA-based Carrier

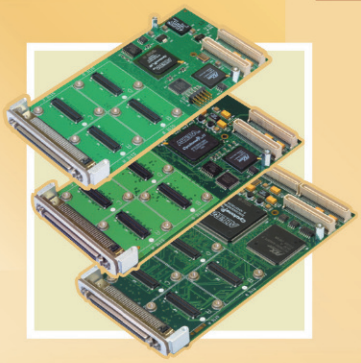

2 Choose I/O Conversion Modules

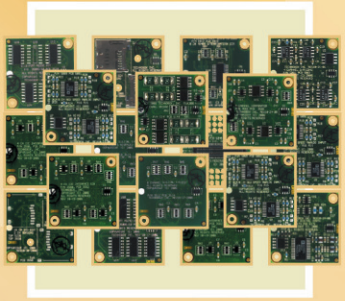

3 Assemble with IP Core and ECM Code

Patent Pending

- ***Build Your Own Board by Mixing and Matching Modular Components***
- ***Thousands of Possible Combinations***
- ***Flexible, FPGA-based, Patent-pending Architecture***
- ***Incorporate Multiple Functions on a Single Board***
- ***Design and Build Application-specific, Future-proofed Solutions***
- ***Accelerate System Development, Reduce Time to Market***

To learn more about the Technobox Micro Mezzanine System, visit:
www.technobox.com/mmsintro-fp.htm

Technobox, inc.®

Micro Mezzanine System

kontron

» We Do Not Build Military Aircraft «

Our Customers Do.

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

Radar

Situational Awareness

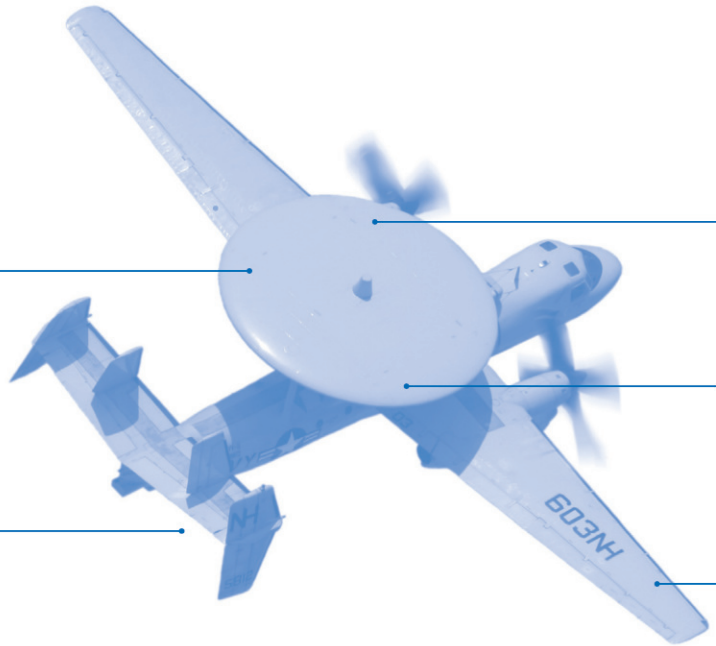

Command and Control

Satellite Communication

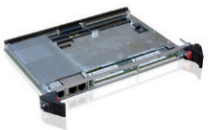

GPS

MILITARY RUGGED COTS

Call, Email or Visit today.

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

If it's embedded, it's Kontron.