

VOLUME 4 NUMBER 3  
MAY 2008

## John Wemekamp

10 GbE bridges  
technology gap



## Refurbish or perish

INCLUDED WITH THIS ISSUE:  
*Programmable Logic:*  
**FPGAs &  
Reconfigurable Computing**



For Systems  
that Demand  
the Most



## XG-5000K

Our most advanced  
FPGA Based  
PC/104 Plus  
Circuit Board

### Features

- Xilinx Spartan 3, 5 Million Gate FPGA
- 264 User Configurable I/O
- 256 MByte of SRAM
- 32 MByte of Flash
- 6 MByte of Serial Page DataFlash
- CompactFlash Type 1 Connector
- Secondary Xilinx Spartan 3, 400,000 Gate FPGA
- 10/100Base-T Ethernet
- 2 RS232 Ports
- 25MHz FPGA Master Clock Source
- Secondary 0 - 25MHz user programmable DDS FPGA Master Clock Source
- Incorporates Xilinx's design revisioning technology and can retain onboard as many as 16 partial or 4 complete FPGA design BIT files
- Can be powered from the PC/104 connector or an external DC power supply
- PC/104 Plus Form Factor;
- ISA and PCI 33MHz Bus Interfaces
- Extended temperature -40C to +85C



Jacyl Technology is a Team Member  
under Lear Siegler Services, Inc for the  
US Governments CECOM Rapid  
Response Contract



Advancing Today's Technology into Tomorrow  
[www.jacyltechnology.com](http://www.jacyltechnology.com)



# Simply Powerful. Powerfully Simple.



## High-performance FPGA I/O designed for faster and easier development.

Only Acromag gives you a reconfigurable FPGA wrapped with just the features you need for fast and easy implementation. These PMC modules provide a high-speed I/O interface with plenty of memory for efficient data handling.

- Virtex-5 FPGAs (VLX155T/110T/85T or VSX95T) are optimized for logic or DSP
- Two banks of high speed 32M x 16-bit DDR2 SDRAM handle high volumes of data
- Two banks of 256K x 32-bit dual-port SRAM optimize DMA transfer to the bus
- PCI-X 133MHz 64-bit interface ensures fast data throughput
- Engineering Design Kit speeds development of custom FPGA applications.

### Plug-in I/O cards extend the I/O capabilities

AXM I/O extension modules expand interface support for a wide variety of analog and digital I/O signals.



### Affordable FPGA-based I/O for all your projects

Acromag's full line of industry pack and PMC modules feature Xilinx FPGAs at a range of price and performance points. Why settle for generic I/O when you can cost-effectively design just what you need?

For details on the fastest and easiest way to develop FPGA I/O visit

**AcromagEmbedded.com,**  
or call 1-877-295-7088

**Acromag**   
THE LEADER IN INDUSTRIAL I/O

# Military

## EMBEDDED SYSTEMS

May 2008 Volume 4 Number 3

### COLUMNS

#### Field Intelligence

- 8 Built-in-Test plays a key role in system integrity**  
*By Duncan Young*

#### Mil Tech Insider

- 10 Rugged 10 GbE closes technology gap for military**  
*By John Wemekamp*

#### Legacy Software Migration

- 12 Ada matters!**  
*By Dr. Robert Dewar, AdaCore*

#### Crosshairs Editorial

- 58 Calling all COTS**  
*By Chris A. Ciuffo*

### DEPARTMENTS

- 53,55 Editor's Choice Products**

- 14 Daily Briefing: News Snippets**  
*By Sharon Schnakenburg*

- 54 New Products**  
*By Terri Thorson*

#### ON THE COVER:

In World War II, soldiers, sailors, and Marines living in barracks kept their personal items in footlockers like this one. Today, information is power, and battlefield data can be as essential as water, bullets, and MREs. Today's "footlockers" are mass storage devices – from solid-state disks, to nvSRAMs, to hard disk drives. But there's more than just storing data – keeping it safe, secure, and routinely backed up is mandatory. Check out the articles starting on page 26.

### WEB RESOURCES

#### Subscribe to the magazine or E-letter:

[www.opensystems-publishing.com/subscriptions](http://www.opensystems-publishing.com/subscriptions)

#### Live industry news:

[www.mil-embedded.com/news](http://www.mil-embedded.com/news)  
[www.opensystems-publishing.com/news/submit](http://www.opensystems-publishing.com/news/submit)

#### Submit new products:

[www.opensystems-publishing.com/vendors/submissions/np](http://www.opensystems-publishing.com/vendors/submissions/np)

### Hardware: Systems – big and small

- 16 The Wishbone II transaction bus: Another grade of speed**  
*By Uros Platise, andEuros*
- 20 Thinking inside the box: Boosting the effectiveness of air cooling**  
*By David Lippincott, Chassis Plans*

### Systems: Mass storage

- 26 More than just cryptography: The performance advantages of Suite B compliance**  
*By William Lattin, Certicom Corp.*
- 32 Flash FPGAs: Often imitated, never duplicated**  
*By Hezi Saar, Actel Corporation*
- 34 Remote backup: A lifeline to businesses' greatest asset**  
*By Rob Cosgrove, Remote Backup Systems*

### Special: Interview with Boeing's Roger Krone

- 36 JTRS, FCS, WGS, and DHS: Chances are Boeing Network & Space Systems builds it, talks to it, or helps the mission succeed**  
*By Chris A. Ciuffo*

### Mil Tech Trends: Obsolescence and DMSMS

- 40 DoD continues refurbishing trend**  
*By Bob Smith, Dynamics Research Corporation*
- 44 Remanufacturing thwarts obsolescence in mission-critical systems**  
*By Arlin Niernberger, GD California, Inc.*

### Technology: Considering COTS costs

- 48 Dancing with wolves: The COTS illusion**  
*By Ian Colville, Aculab*

### EVENTS

#### Intel Developer Forum (IDF)

August 19-21, 2008

Moscone Center • San Francisco, California

[www.intel.com/idf](http://www.intel.com/idf)

**CORRECTION:** In the Mar/April issue, the Sealevel Systems Editor's Choice write-up implied that the PDA-184 software was created by Sealevel for their innovative ACC-188. The software was actually developed by the Defense Information Systems Agency (DISA) but is shipped by Sealevel as part of the USB Synchronous Interface Adapter package. Additionally, the product works with tactical data networks that are not likely to interface with the public Internet. *MES* regrets this error.



Published by:

All registered brands and trademarks within *Military Embedded Systems* magazine are the property of their respective owners.  
© 2008 OpenSystems Publishing © 2008 Military Embedded Systems



# **Annapolis Micro Systems**

*The FPGA Systems Performance Leader*

## **WILDSTAR 5 for IBM Blade**

### **The Perfect Blend of Processors and FPGAs**

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



**Made in the USA**

### **Ultimate Modularity**

**From 2 to 8 Virtex 5 FPGA/Memory Modules**

**Input / Output Modules Include:**

**Quad 130 MSps thru Quad 500 MSps A/D**  
**1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D**  
**Dual 1.5 GSps thru 4.0 GSps D/A**  
**Infiniband, 10 G Ethernet, FC4, SFPDP**

***Direct Seamless Connections with no Data Reduction***

***Between External Sensors and FPGAs***

***Between FPGAs and Processors over IB or 10GE Backplane***

***Between FPGAs and Standard Output Modules***

**190 Admiral Cochrane Drive, Suite 130, Annapolis, Maryland USA 21401**  
**[wfinfo@annapmicro.com](mailto:wfinfo@annapmicro.com) (410) 841-2514 [www.annapmicro.com](http://www.annapmicro.com)**





## Advertising/Business Office

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

## Vice President Marketing & Sales

Patrick Hopper  
phopper@opensystems-publishing.com

## Business Manager

Karen Layman

## Sales Group

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

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

Doug Cordier  
Account Manager  
dcordier@opensystems-publishing.com

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

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

## Regional Sales Managers

Jerry Bleich  
New England  
jbleich@opensystems-publishing.com

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

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

Denis Seger  
Southern California  
dseger@opensystems-publishing.com

Sydele Starr  
Northern California  
sydele@pacbell.net

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

## International Sales

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

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

## Reprints and PDFs

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

## ADVERTISER INFORMATION

Page	Advertiser/Ad title
3	Acromag, Inc. – Simply powerful
55	Aculab – Military commanders must
5	Annapolis Micro Systems, Inc. – WILDSTAR
22	BMC Communications – ARINC and MIL-STD-1553
28	Connect Tech Inc. – Solid-state storage for rugged applications
60	Curtiss-Wright – Tough DSP solutions
9	Data Device Corp. (DDC) – Tools of innovation
31	DIGITAL-LOGIC AG – PCI/104-Express with Core 2 Duo
25	Elma – Tailored to your exact specifications
17	Excalibur Systems, Inc. – Dense?
59	GE Fanuc Intelligent Platforms, Inc. – Here comes the future
38	Harting Technology Group – DIN 41 612
47	Hypertronics – Protect against shock and vibration
2	Jacyl Technology Inc. – XG-5000K
11	Microbus Inc. – Elcard Wireless LAN
19	MPL AG – Fanless and customizable
49	Octagon Systems – Core systems
50	Pentek, Inc. – Introducing the one board
35	Performance Technologies – Looking for a first-class CompactPCI solution
21	Phoenix International – Data storage
7	PQI Corporation – Storage leader
45	Technobox, Inc. – XMC and PCIe
22	TEWS Technologies LLC – COTS I/O solution
42	Themis Computer – New rugged, reliable servers
27	Tri-M Systems Inc. – PC/104 solutions
29	Tri-M Systems Inc. – TRI-M Engineering
13	VersaLogic Corp. – When You Design
39	VMETRO – Rugged high-speed data recording and storage
43	White Electronic Designs – Solutions are in sight
24	Winchester Electronics – Power connector solutions

# Military EMBEDDED SYSTEMS

AN OPENSYS TEMS PUBLICATION

## Military and Aerospace Group

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

**Group Editorial Director** Chris A. Ciufu  
cciufo@opensystems-publishing.com

**Senior Editor (columns)** Terri Thorson  
tthorson@opensystems-publishing.com

**Associate Editor** Sharon Schnakenburg  
sschnakenburg@opensystems-publishing.com

**Assistant Editor** Robin DiPerna

**European Representative** Hermann Strass  
hstrass@opensystems-publishing.com

**Senior Web Developer** Konrad Witte

**Web Content Specialist** Matt Avella

**Creative Director** Steph Sweet

**Art Director** David Diomede

**Graphic Coordinator** Sandy Dionisio

**Circulation/Office Manager** Phyllis Thompson  
subscriptions@opensystems-publishing.com



## OpenSystems Publishing

Editorial/Production office:

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

**Publishers** John Black, Michael Hopper, Wayne Kristoff

**Vice President Editorial** Rosemary Kristoff

## Communications Group

**Editorial Director** Joe Pavlat

**Assistant Managing Editor** Anne Fisher

**Senior Editor (columns)** Terri Thorson

**Technology Editor** Curt Schwaderer

**European Representative** Hermann Strass

**Senior Designer** Joann Toth

## Embedded and Test & Analysis Group

**Editorial Director** Jerry Gipper

**Editorial Director** Don Dingee

**Senior Associate Editor** Jennifer Hesse

**Special Projects Editor** Bob Stasonis

**European Representative** Hermann Strass

ISSN: Print 1557-3222

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

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

Canada: Publication agreement number 40048627

Return address WDS, Station A PO Box 54, Windsor, ON N9A 615

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



# Storage Leader !

.Integration  
.Reliability  
.Performance

New!



PQI Corporation  
Tel:(510)651-7281  
Fax:(510)651-7240  
Learn more at: [www.pqimemory.com](http://www.pqimemory.com)







# Built-in-Test plays a key role in system integrity

By Duncan Young



Often perceived as just a tick in the box during the selection process, Built-in-Test (BIT) is an invaluable component of modular, embedded systems that are used for critical applications such as avionics mission systems, sensors, and weapons. BIT provides a level of confidence in the correct operation of each module at both power-up and during normal operation. As more of these critical embedded systems are assembled from off-the-shelf hardware and software components, it is increasingly important to evaluate BIT's performance and its potential for interaction with software such as RTOS.

These types of systems are typically configured from VMEbus, VXS, VPX, or CompactPCI from vendors such as GE Fanuc, Curtiss-Wright, Mercury, and Thales using OSs from Green Hills Software, LynxOS, and Wind River. The way embedded computing systems are architected is changing rapidly: more highly integrated at the platform level plus extensive communications and network connectivity for cooperative engagement in the field. BIT supplied by hardware vendors is changing to reflect these new requirements by moving away from the traditional three-layer approach toward greater BIT functionality running concurrently with the application.

### Three-layer approach

The traditional approach to BIT is representative of a much less complex era, when federated subsystems could be tested individually without impacting other subsystems. It has three layers:

- **Power-up BIT (PBIT)** – Comprehensive tests of hardware functionality extending as close to the edge of a module as possible.
- **Initiated BIT (IBIT)** – Usually invoked following a failure, IBIT initiates tests to isolate a fault within a system or subsystem. In general, IBIT will halt the current application, run its tests, and either resume or restart the application.
- **Continuous BIT (CBIT)** – These are any tests performed while the application is active. They may be part of the application itself or might be supplied by the hardware module vendors to run periodically.

### IBIT becoming impractical

PBIT is still a necessary requirement to provide confidence in the hardware integrity. It is usually the SBC vendor that provides the basic PBIT package enhanced with PBIT modules for each of the additional hardware components. The concept of IBIT makes good sense where there are many discrete subsystems that do not interact with each other, for example, a Naval radar receiver or Electronic Counter-Measures (ECM) system. However, on other weapons platforms such as Unmanned Aerial Vehicles (UAVs), single seat combat aircraft, helicopters, or ground vehicles, the operators will not have the time nor the skills to employ intrusive test methods to diagnose and repair faulty equipment during a

mission. Modern avionics systems are becoming more integrated as their mission computers perform many different functions using the same hardware, making it more difficult for an individual function to be diagnosed and rectified using IBIT.

### Online BIT holds key for the future

As the value of IBIT diminishes with more system integration, the burden of hardware confidence must pass to either the application software, the operating system, or the hardware vendor's CBIT. Deployable systems make use of all of these; nevertheless, as the OS and CBIT are off-the-shelf products, their vendors have a responsibility to ensure tight integration without degrading the OS's key performance parameters or determinism. To be nonintrusive, CBIT must have very low latency and always return control to the application with the hardware environment unchanged. This can be achieved by running CBIT in a secure partition and ensuring that tests run with the minimum impact on the operational hardware environment (for example, by running memory tests in dedicated segments per bank). These techniques are used by GE Fanuc Intelligent Platform's Deployed Test Solutions product, which offers extensive CBIT testing in the form of Background Condition Screening (BCS).

It has always been a goal of the military to operate their platforms with some subsystems inoperative or with reduced functionality in the event of insufficient maintenance, spare parts, or battle damage. In the era of federated systems, this was easily achieved by disconnecting or ignoring nonserviceable subsystems. However, this approach is not always possible with an integrated system or with functions of a system that rely on extended connectivity through a network for their correct operation. CBIT has an important role to play by monitoring and logging the state of the hardware environment which, together with data gathered by the application, will allow a user to make an informed decision on which operations can be trusted and which should not.

Taking a glimpse into the future, one of the higher integration costs of embedded COTS subsystems is integrating BIT across dissimilar vendors' hardware. This might be perceived as one of the few remaining areas for market differentiation between vendors but could, usefully, be targeted for standardization. Additionally, in the future it is likely that as transistor count of processor cores continues to multiply, there will be a surplus of processing capacity that could be used exclusively for BIT, decreasing its latency and increasing its coverage. BIT is not a static piece of firmware in flash EPROM. It is continuing to evolve from a diagnostic maintenance aid to meeting the user's needs for high system availability by providing practical feedback on the ability of the hardware to meet its critical-mission requirements.

*To learn more, e-mail Duncan Young at [young.duncan1@btinternet.com](mailto:young.duncan1@btinternet.com).*





## TOOLS OF INNOVATION



Ideal for  
lab applications

**DDC's Family of USB Interfaces...  
Small, Lightweight, and Portable.**

- MIL-STD-1553 and ARINC 429
- Graphical User Interface
- High Channel Count
- USB Powered



Ideal for  
rugged/embedded  
applications

[www.ddc-web.com](http://www.ddc-web.com) Toll Free: 1-800-DDC-5757



## Rugged 10 GbE closes technology gap for military



By John Wemekamp



The military was once perceived as a very conservative adopter of new technology, even though it was the dominant force behind investment in microelectronics. This was because the military cycle from development to deployment was so long that commercial markets could exploit the technology more rapidly. When the funding finally ended, development was already driven by the commercial, consumer, and telecommunications markets, where the massive development costs get recovered in volume sales in time to feed the cycle again. But the success of COTS rugged embedded computer vendors and the development of new platform standards have placed the military back in the forefront of influence and deployment of commercial technologies, such as Ethernet and IPv6, in parallel with their other non-rugged market users.

10 GbE is still in its infancy in any market segment, with standards being developed as it is introduced into new application areas. The most common physical standard is known as XAUI (Roman *ten* Attached Unit Interface), which uses a 16-pin connector for interconnection over a few meters. For longer distances XAUI can be converted to fiber, and a number of standards exist for this. But these optical standards are more difficult to implement and have not received wide acceptance for critical embedded computing applications such as blade servers, routers, switches, and multiple computing clusters. This prompted the development of a serialized four-lane 10 GbE backplane standard, signaling at 3.125 GHz, referred to as 10GBASE-KX4.

### Evolving backplane requirements

Parallel backplane buses are being displaced from embedded computing applications, particularly where multiple computing elements are involved. High-performance processor devices and hence, SBCs or server blades, now incorporate two or more integrated GbE

ports. In a multicompute environment, these ports are connected to each other via switches to create a control/management plane which, in some cases, also serves as a data plane. Since Ethernet is a preferred medium for subsystem-to-subsystem connection, as well as a platform backbone, using common IP standards within and between subsystems simplifies development and integration of an overall platform's application. No other type of network technology offers this all-around capability: board to board, subsystem to subsystem, access to common applications and databases, the use of Network Attached Storage (NAS), plus connection to external networks such as the evolving Global Information Grid (GIG).

The availability of embeddable Ethernet switches has made it easier to integrate subsystems in racks for rugged applications, such as new-generation ground vehicles, that require the most efficient use of weight, space, and cooling. New embedded switches are being introduced with both GbE and 10 GbE ports. SBCs or blades are normally connected to the switch via its GbE ports. Switch-to-switch and external paths are connected using the extra performance of 10 GbE to avoid potential bottlenecks at switches. However, where additional bandwidth is required, individual 10 GbE links can be added to SBCs and blades by means of PMC/XMC mezzanines giving additional high-performance point-to-point or networked ports. Until recently, backplane technology has limited the connectivity of 10 GbE to front panels of SBCs, PMC/XMC mezzanines, and switches, making the implementation of truly rugged systems difficult to achieve with any level of standardization among vendors.

The rapid adoption by COTS vendors and systems integrators of the VPX (VITA 46) standard is creating a freedom of topology for rugged embedded computing applications, enabling the

use of both GbE (1000BASE-T) and 10 GbE (10BASE-KX4, VITA 46.7) across the backplane and the incorporation of dedicated switch slots within a rack (VITA 46.20). The 6U VPX6-684 FireBlade II Ethernet switch from Curtiss-Wright Controls Embedded Computing (CWCEC) is shown in Figure 1.




Figure 1

In addition to 10 GbE, the introduction of IPv6 is an important part of the DoD's future strategy. Although IPv6 provides enhanced security features, it cannot provide the unbreakable levels of protection needed in the information war zone. One approach to improving security is to add PMC/XMC mezzanines to critical SBCs and switches providing payload encryption for GbE ports using Advanced Encryption Standard (AES) and its forebear, Triple Data Encryption Standard (3DES), for high levels of classification.


The military and its supporting COTS vendors actively participate in technology development such as VPX, as well as proposing and supporting Request for Comments (RFCs) applicable to new military requirements through the Internet Engineering Task Force (IETF). This renewed focus by the DoD and other government agencies on Ethernet and IPv6 is key to their future enterprise-wide and operational strategy, closing the deployable technology gap with other leading markets.

To learn more, e-mail John at [john.wemekamp@curtisswright.com](mailto:john.wemekamp@curtisswright.com).



# Elcard™ Wireless LAN Modules

## Designed for Industrial and Professional Applications



**PC/104+ WIB4xx  
dual WLAN module**


### USB-attached WLAN

**Rugged  
Access  
Points  
Available**

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

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

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



### AIB220 PC/104+ Cardbus

- Cardbus/PCMCIA Adapter
- Dual Type I/II or single Type III
- Linux and Win9x/2K/XP support
- 3.3V and 5V card support
- TI PCI420 chipset

# Elcard

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



# Legacy Software Migration

By Dr. Robert Dewar



## Ada matters!

*The computing industry is rather peculiar: If a technology is not the latest and greatest, people think it has vanished entirely. You can, for example, often meet otherwise knowledgeable people who think that mainframes and COBOL have completely disappeared.*

Here at AdaCore, we often find people who think that Ada has disappeared and are surprised to find out that not only is it still around – in the domain of large, critical applications it is especially alive and well. For example, the new air-traffic control system for the UK, iFACTS ([www.praxis-his.com/news/radarMar2007.asp](http://www.praxis-his.com/news/radarMar2007.asp)), is being written from the ground up in Ada using the SPARK formal methods approach.

So why is Ada still around when many college students are learning Java, and there are even people around who think C and C++ are disappearing? The answer is remarkably simple. A recent National Academy of Sciences report, “Software for Dependable Systems: Sufficient Evidence?” ([http://books.nap.edu/catalog.php?record\\_id=11923](http://books.nap.edu/catalog.php?record_id=11923)), addresses the issue of security-critical software, and states:

“The overwhelming majority of security vulnerabilities reported in software products – and exploited to attack the users of such products – are at the implementation level. The prevalence of code-related problems, however, is a direct consequence of higher-level decisions to use programming languages, design methods, and libraries that admit these problems.”

There are no silver bullets when it comes to safety- and security-critical software. No tools can guarantee success, but as this quote indicates, it helps to use appropriate technology, including the right programming language. Last year’s High Confidence Software and Systems (HCSS) conference, sponsored by NSA to address security-critical issues, featured an interesting presentation from Microsoft addressing such issues in the context of Windows. The primary sources of problems in Microsoft’s experience are buffer overruns and integer overflow problems. Reasonable enough, so how can a programming language help?

The answer is: quite a bit. However, if we look at C or C++, these languages notoriously have absolutely nothing to help avoid overruns. Microsoft is trying to add assertions to C to help, and found about 100,000 possible buffer overflows in the Vista code. Ada, by contrast, fully checks all array references as a matter of course, to eliminate the possibility of buffer overflow.

Although Java is often cited as a safer alternative to C and C++, when it comes to integer overflow, the semantics for Java can introduce some dangerous vulnerabilities. Overflows wrap silently, so that adding 1 to a positive number can suddenly make

it negative. In Ada, every integer value can have a sensible range assigned, and automatic checks ensure that values do not go outside this range.

---

“Ada ... fully checks all array references as a matter of course, to eliminate the possibility of buffer overflow.”

---

These are just a couple of small examples of how a language can help. It’s not at all surprising that Ada should address issues like these; it was designed for the purpose of writing large critical programs, something that cannot be said of languages like C, C++, or Java. If you search language standards for features related to safety and security, you will draw a blank, unless you are looking at the Ada standard. With Ada, you will find a whole annex devoted to high-integrity programming.

Ada does not guarantee success, but as the NAS study indicates, it makes sense to use tools suitable for the purpose; if one’s purpose is writing large safety- or security-critical programs, then Ada is a natural choice. That’s why it should not be surprising that Ada continues to play a critical role and that Ada is a programmer’s frequent companion. Many might not be aware of that fact, however. An example, though, is the avionics system of the new Boeing 787 Dreamliner, which makes extensive use of Ada. Of course, travelers won’t be aware of this when they fly, and that’s the point. Software that works properly is out of sight, and out of mind. Programmers who build such systems know the advantages of Ada – and continue to choose it as one key element of critical systems.


*Dr. Robert Dewar is cofounder, president, and CEO of AdaCore; he also has had a distinguished career as a professor of Computer Science at the Courant Institute of New York University. He has been involved with the Ada programming language since its inception in the early 1980s and, as codirector of both the Ada-Ed and the GNAT projects, led the NYU team that developed the first validated Ada compiler. Robert was one of the authors of the requirements document for the Ada 95 revision, and he served as a distinguished reviewer for both Ada 83 and Ada 95. He has coauthored compilers for SPITBOL (SNOBOL), Realia COBOL for the PC (now marketed by Computer Associates), and Alslys Ada. He is also a principal architect of AdaCore’s GNAT Ada technology. He has also written several real-time operating systems for Honeywell Inc. and frequently shares his thoughts on computers and on open-source software at conferences. He can be reached at [dewar@adacore.com](mailto:dewar@adacore.com).*



# ***When You Design for the Extreme Choose Your Partner Wisely***

**Did we mention "Reliability" is our middle name?**

**I**ndustrial equipment needs to perform flawlessly, night and day, under even the most extreme conditions. Whether designing for the manufacturing floor, clean room or the field, you can depend on VersaLogic to deliver the highest quality embedded computer products, from prototyping and design-in, through years of product production. We design our boards for high reliability and long-term availability, then run them through exhaustive quality tests, ensuring that we deliver only the best. And with our world class service and five year availability guarantee, we'll always be there when you need us. Whether you need one of our standard products or a customized version, our skilled technical staff will work to meet your exact requirements.



Contact us to find out how for more than 30 years we've been perfecting the fine art of extra-ordinary support and on-time delivery: One customer at a time.

1.800.824.3163  
1.541.485.8575  
[www.VersaLogic.com/ind](http://www.VersaLogic.com/ind)



**VERSALOGIC**  
CORPORATION



# Daily Briefing:

News Snippets

By Sharon Schnakenburg, Associate Editor

[www.mil-embedded.com/dailybriefing](http://www.mil-embedded.com/dailybriefing)

Photo courtesy of Boeing



## Boeing: Joining A and B

Proving that the proverbial A plus B still equals C, the Boeing Company was recently chosen to integrate A (Joint Helmet-Mounted Cueing System or JHMCS) into B (145 F-15E aircraft), equaling the terms of C (a \$49.5 million U.S. Air Force contract to execute the integration project). Accordingly, the JHMCS integration contract comprises aircraft installation services and hardware,

along with initial pilot equipment including visors and helmets. The first F-15E installation is slated for October; meanwhile, contract completion is expected in December 2010. Boeing's is the ninth contract for JHMCS, which aims to increase pilots' situational awareness.

## Northrop Grumman's software by any other name ...

... is not the same. It's actually better, so the Marine Corps Systems Command in Quantico, Va. hopes. The \$29 million modified contract between the Marines and Northrop Grumman Mission Systems in Reston, Va. specifies that Northrop Grumman rearchitects the Command and Control Personal Computer (C2PC) software tool. C2PC was originally developed in 1994 by International Research Institute Inc., later acquired by Northrop Grumman. It serves as client software for Global Command and Control System (GCCS), along with other Common Operating Environment (COE) systems to ensure a common view of the battlespace is shared by all users. As it evolved, C2PC has also satisfied tactical requirements in the Data-Automated Computer Terminal (DACT) and other programs operating in intermittent communication, low-bandwidth environments. New project completion is expected by September 2009, and follows previous C2PC versions 6.2 and 7.0.

## Kontron picks up where Thales signs off

The high-tech world is always changing ... and changing hands. Case in point: Thales recently sold its Thales Computers SA subsidiary in Toulon to the Germany-based Kontron Group. Thales Computers SA holds revenues of more than €20 million and has about 100 employees; its specialty is manufacturing and designing standardized embedded boards. Meanwhile, Kontron Group is a global player in the embedded computing industry, with about 2,500 employees, revenues exceeding €400 million, and 30 percent "internal growth" in 2006, Thales reports.

## DoD releases report on China's military

In compliance with the National Defense Authorization Act for Fiscal Year 2000, the U.S. Office of the Secretary of Defense recently released its *Annual Report to Congress: Military Power of the People's Republic of China 2007*. While China's long-distance capabilities to sustain military power are presently limited, the 2006 *Quadrennial Defense Review Report* (referenced in the new report's exec summary) revealed China "has the greatest potential to compete militarily with the United States and field disruptive military technologies that could over time offset traditional U.S. military advantages." The 50-page report reveals electronic warfare strategies to interrupt battlefield network information systems and thereby gain dominance in conflict, among other topics (see [www.defenselink.mil/pubs/pdfs/070523-China-Military-Power-final.pdf](http://www.defenselink.mil/pubs/pdfs/070523-China-Military-Power-final.pdf)).



Photo courtesy of NASA/John Frassanito and Associates

## Lockheed Martin and LDRA team up

Prime contractor Lockheed Martin (LM) needed software tools to assist Orion Crew Exploration Vehicle (CEV) developers in achieving program software development goals. Enter safety-critical vendor LDRA, with its Testbed and TBrun tools to aid software developers in code coverage analysis, software standards checking, object code verification, and unit testing for the program. The rest, as they say, is history – or just might make history: NASA's Orion CEV aims to transfer astronauts to and from the International Space Station (ISS), Mars, the moon, and other locations safely – and replaces the space shuttle that retires in 2010. Orion CEV's manned mission debut is slated for 2014.



## U.S. Army's WIN-T "Expand"s optimization contract

The U.S. Army, along with General Dynamics C4 Systems (GDC4S), recently "Expand"ed its Warfighter Information Network – Tactical (WIN-T) Increment 2–3 TCP Performance Enhancing Proxy (PEP) contract to include Expand Networks. WAN optimization provider Expand Networks' role is to port the program's Accelerator Operating System to a WIN-T hardware platform developed by General Dynamics. Expand Networks' PEP is designed to operate in a mobile environment featuring dynamic outbound links frequently broken and created, to provide war fighters with highly efficient communications while halted or on the move. Meanwhile, WIN-T, the U.S. Army's high-speed/capacity communications network, continues to link ground-level soldiers to their commanders and the GIG.



Photo courtesy of the U.S. Army

## What's it like out there?

Space Micro is helping NASA answer that question, via its space dosimeter card recently integrated into the Living with a Star (LWS), Space Environments Testbed (SET) experiment hardware at NASA Goddard. Dr. Peter McNulty and a team of designers at Clemson University designed the hardware in 2007, and Space Micro then produced and tested the dosimeter card in accordance with space quality requirements. The technology will be used in the Dosimetry Intercomparison and Miniaturization Experiment (DIME), where several radiation detector technologies monitor the radiation dose in space. Five COTS micro-dosimeters will also be used to characterize single event effects and radiation-induced total ionizing dose, in addition to displacement damage. Within 12 months, the DIME card is slated for secondary payload launch on an Air Force booster, Space Micro reports.

## Mercury rises to prime JCREW role


Mercury Federal Systems Inc., a Mercury Computer Systems subsidiary, recently signed on as prime contractor to show the U.S. Army just how the Joint Counter Radio-Controlled Improvised Explosive Device Electronic Warfare (JCREW) program is done – literally. The \$2.5 million contract between Mercury and the Army's New Jersey-located Ft. Monmouth CECOM stipulates that Mercury develops an open-systems DSP architecture testbed and technology, to be utilized by the U.S. government for JCREW counter-IED (Improvised Explosive Device) systems research and development; Mercury has also agreed to provide systems engineering services to contractors using the testbed and to the U.S. government. Additionally, Mercury announced its selection of EW technologies supplier ITT Electronic Systems to collaborate on initial development and demonstration of the Constant Amplitude Network Waveform Adaptive Combiner (CANWAC) algorithm's portability.

## Let's figure this out

So many security tools ... so little systems analysis, reports Joe Schlesselman, RTI's director of market development for aerospace and defense, referring to content in the Air Force's original solicitation for the "Proactive Determination of Networked Node Vulnerability" project. RTI was recently chosen for the project, initiated by the U.S. Air Force Research Laboratories (AFRL), which focuses not only on automating network node vulnerability scanning, but also on analyzing the vulnerabilities' impact on the network. RTI's primary responsibilities encompass developing and researching a Data Distribution Service (DDS) to execute said security vulnerability scanning. The company plans to combine open-source technologies, mainstream commercial products, and research via security partners including Promia, Inc. and Partners ObjectSecurity LLC. Several unnamed prime contractors will also participate in evaluation and testing. The technology is scheduled for demonstration in mid 2008.

## Working to receive mixed signals

Radar ... communication ... radar ... communication ... The days of switching back and forth between communications signals and radar signals are over, thanks to the multichannel SR-5124 signals acquisition system. The SR-5124 SIGINT recorder, recently selected by both the German Ministry of Defense and the U.S. Navy for their next-gen EW systems, can detect, store, and analyze both COMINT and ELINT signals synchronously and simultaneously. The device, produced by EONIC B.V., also allows users to integrate proprietary algorithms for data analysis into the processing flow and signals acquisition. Altera, whose Stratix II FPGAs were recently chosen by EONIC B.V., is now onboard the project as well.





# The Wishbone II transaction bus: Another grade of speed

By Uros Platise

*Wishbone specifications have been released by OpenCores and Silicore with the aim to provide a standard IP core interconnection scheme to fulfill requirements of modern System-on-Chip (SoC) designs, including CPUs, DMA engines, memory interfaces, peripheral interfaces, and so on. The andEuros company has used the Wishbone specification since its inception and has developed an improved version of the Wishbone bus, called Wishbone II, to propose an advanced pipelined architecture where read and write transactions are separated and the bus acts as a transaction bus. In this way, multiple transactions can take place at the same time, removing all latencies along the path and stalling RMW cycles by incorporating a new per-cell locking concept. The ultimate benefit, of course, is that finally bus throughput has been increased to the maximum.*

Design and development of large-scale FPGA/ASIC SoC designs have forced designers to implement a modular architecture with a standardized module interface that connects various IP modules in any possible configuration. One of the most popular interconnect architectures was released by OpenCores called the *Wishbone B.3 bus* ([www.opencores.org](http://www.opencores.org)). In a similar way, Altera has introduced its own interconnect scheme called *Avalon Bus* ([www.altera.com](http://www.altera.com)) around which SOPC Builder and Nios (II) Systems are made. Xilinx has also introduced its own bus called the *On-Chip Peripheral Bus* combined with the *Processor Local Bus* ([www.xilinx.com](http://www.xilinx.com)).

These interconnect architectures are single transaction master/slave oriented, meaning that a CPU requesting a word from a given address stalls itself and a path (bus) to the destination for as long as this word is not received. Lots of bus cycles are lost in this way, giving lower actual data throughput than expected despite the relatively high system bus frequency. Even with fast burst reads and writes introduced by special signals, bus cycles are still lost until the first word is received at the additional cost of doubling the burst logic at both sides, source and destination. Bus stalling is more evident when accessing slower modules with greater latencies. In these cases, system performance degrades dramatically; for example, a 100 MHz system may see its throughput fall as low as a few MB per second.

That is why there was a desperate need to develop bus architectures employing new concepts. Some new signals have been introduced to support new transaction bus concepts based on the Wishbone B.3 architecture, overcoming latency issues while maintaining backwards compatibility.

### Wishbone II transaction bus concept

In our proposed bus, transactions are represented by a *transaction vector* containing:

- Source (module) address
- Destination (module) address
- Operator
- Data

Source and destination addresses define the path; the operator describes one or

more operations to be executed along the path and/or at the destination address; and some operations require supplemental data given to complete the transaction. Actual implementation requires additional handshaking signals.

Transaction vectors are placed onto a *transaction bus* transporting the vector from source to destination, and executing bus-oriented operations as requested by the vector. Once the transaction vector is placed (sent), the source has no further responsibility and the transaction bus takes complete control over it. The source is then ready to issue the next transaction vector. Multiple tasks or requests may be issued beforehand, one per bus cycle, which reduces the need for any prediction logic at the destination module to support burst reads or writes as prediction logic for various kinds of burst reads.

There are two kinds of transactions:

- Independent
- Dependent (when their order is important)

To support dependent transactions, the transaction bus must never change the order of already placed transactions. The transaction bus features a fully acknowledged mechanism to accept new transaction vectors, execute internal forwarding, and deliver to the destination module. The transparent architecture reflects itself as a simple input-output black box; however, the implementation is based on a multi-pipelined structure where each (FIFO) line holds one transaction vector.



The Wishbone II transaction bus proposes four basic operations only:

- Single read
- Single write
- Cell lock
- Bus lock

Single read and write are issued by modules, where cell and bus locking operations are in the transaction bus domain. Burst reads and burst writes are accomplished by issuing a stream of read or write transactions. RMW cycles are supported through the bus, or even better, they can be facilitated using the new cell locking concept, which instead of stalling the complete SoC bus locks a single or multiple memory cells only to a given owner. These cells cannot be accessed by others as long as they are not unlocked.

### Wishbone II signals

A Wishbone II transaction vector is composed from the Wishbone B.3 specifications by introducing the following new signals:

**WB\_ACW** Write Acknowledge

**WB\_ACR** Read Acknowledge

**WB\_TGA** Address Tag in both directions

**WB\_ALK** Address Lock

In the further text, prefix **WB** may be changed to **WBM** denoting a master interface, and **WBS** denotes a slave interface or can be left blank to describe any master or slave interfaces. Input signals are appended **\_I** at the end and output signals with **\_O**. The proposed bus discards the Wishbone B.3 **ACK** signal since its functionality is now split among the **ACR** and **ACW** signals. Complete basic signal descriptions for master and slave are listed in Table 1. New signals are marked in **bold**.

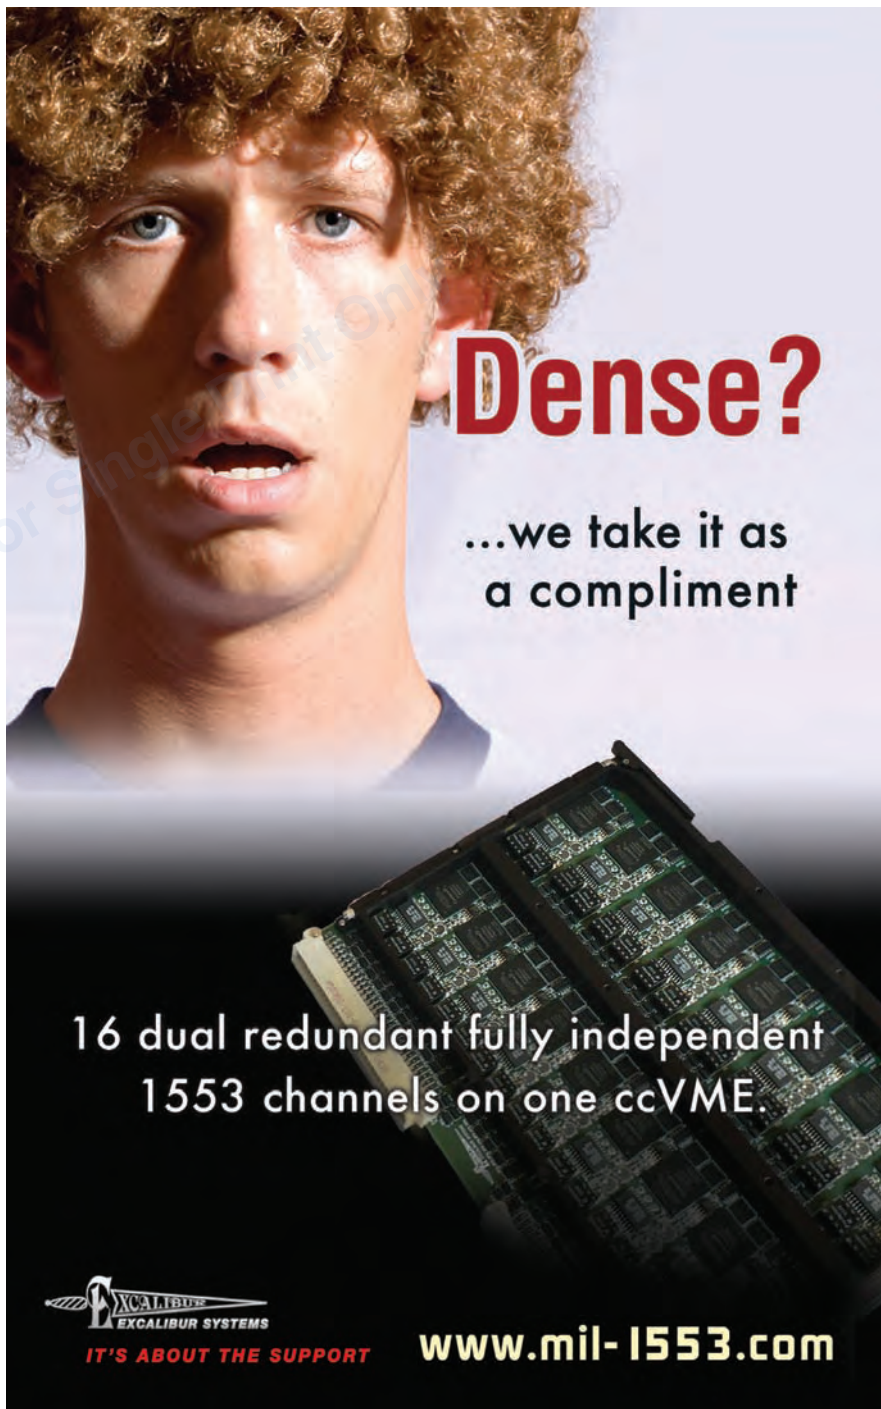
### Wishbone II bus transactions

#### Write transactions

A write transaction is almost identical to the write transaction given in the Wishbone B.3 specifications, except Wishbone II uses the **ACW** signal to acknowledge a write cycle. A read and write transaction is composed of read requests that are identical to write transactions except that the destination operation signal **WE** is set.

DESCRIPTION	MASTER	SLAVE
Data from master to slave (Data)	WBM_DAT_O	WBS_DAT_I
Data from slave to master (Data)	WBM_DAT_I	WBS_DAT_O
Slave (Destination) Address	WBM_ADR_O	WBS_ADR_I
Transaction Strobe (Handshaking)	WBM_STB_O	WBS_STB_I
Destination Operation (Operator)	WBM_WE_O	WBS_WE_I
Bus Lock (Operator)	WBM_LOCK_O	WBS_LOCK_I
Write Acknowledge (Handshaking)	WBM_ACW_I	WBS_ACW_O
Read Acknowledge (Handshaking)	WBM_ACR_I	WBS_ACR_O
Address Tag Write (Source)	WBM_TGA_O	WBS_TGA_I
Address Tag Read (Destination)	WBM_TGA_I	WBS_TGA_O
Address Lock (Operator)	WBM_ALK_O	WBS_ALK_I

Table 1



**Dense?**

...we take it as a compliment

16 dual redundant fully independent  
1553 channels on one ccVME.

**EXCALIBUR SYSTEMS**  
IT'S ABOUT THE SUPPORT

[www.mil-1553.com](http://www.mil-1553.com)



### Read transactions

A read transaction is composed of two transactions:

- Read request transaction issued by source
- Read response transaction issued by destination

A read request is sent by the master module representing a source by first issuing a write transaction with the destination operation **WE** set to read. The Master should set the Address Tag Write vector to identify read response. (If there is a single master, this is not necessary.) The read request transaction is acknowledged in the same way as the write transaction.

The destination completes the transaction by returning a separate read response transaction marked by the acknowledge signal **ACR** and providing valid data and Address Tag Read information. Address Tag Read is a copy of the Address Tag Write.

Figure 1 shows an example system with one pipeline stage on write (input) and read (output) paths between the source (master) and destination (slave) devices. The system has 1 cycle directions on both directions; therefore, a request-response loop takes at least 2 wait cycles. Slave (memory) may also perform some internal management like refresh, which adds up to the total number of wait states.

You can see that Figure 2 depicts a transaction bus data flow diagram for the given example of the three read request transactions placed by the master as **AD0**, **AD1**, and **AD2**, and the associated returned read response transactions as **DO0**, **DO1**, and **DO2**. The signal **WE** is assumed to be cleared for all three transactions to indicate read operations. Transactions **AD0** and **AD1** are burst transactions, meaning that **AD1 = AD0 + 1**, and the **AD2** is an independent transaction triggered meantime that could be a cause of an external interrupt that loads its interrupt vector, and so forth.

Each read request transaction is acknowledged by the **ACW** signal, and the returned read response transaction is marked (acknowledged) by the **ACR** signal. Note that the latency order may not be the same, due to other higher priority master(s) or memory refresh functions, and so on. In the previous example,




Figure 1

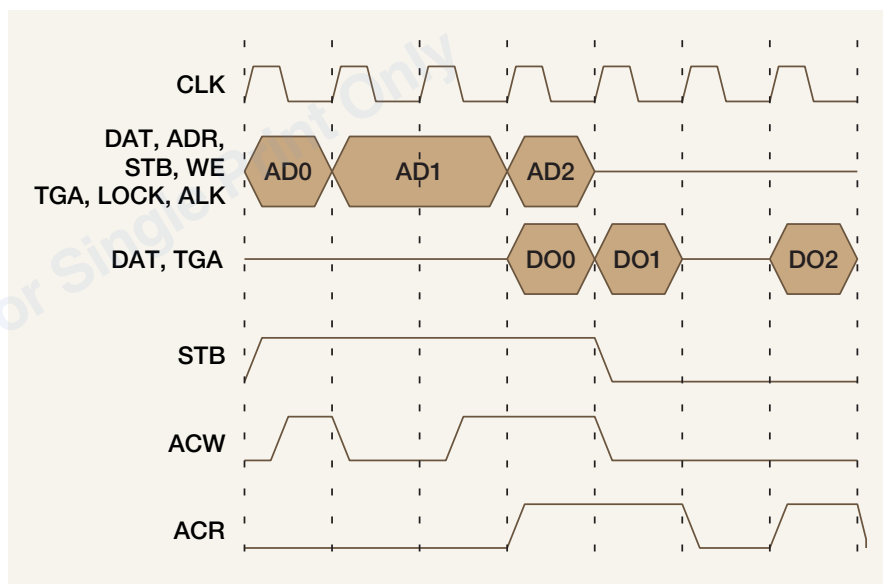


Figure 2

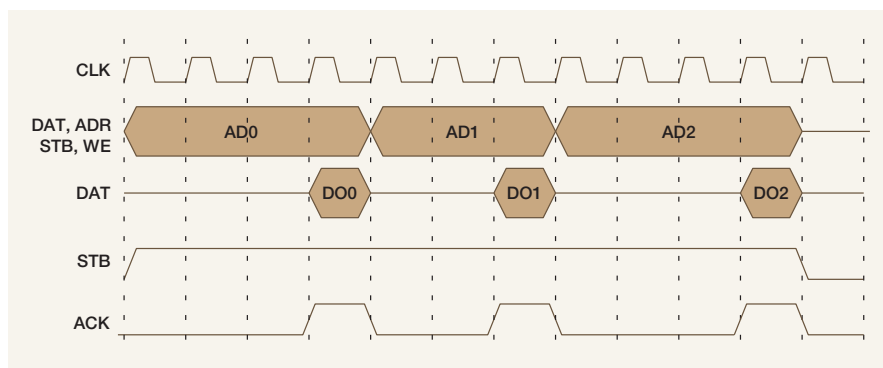


Figure 3

the **AD0** is immediately acknowledged but it takes 3 wait cycles to return the **DO0**; the **AD1** is acknowledged 1 cycle later while the **DO1** is returned in

2 wait cycles only, and the **DO2** again takes 3 wait cycles. All three transactions are completed in 9 cycles; theoretically, without adding two illustrative wait



cycles, they would complete in 7 cycles only. Using the Wishbone B.3 specifications, the same scenario is shown in Figure 3.

Where again **AD0** and **AD1** are bursts, **AD1 = AD0 + 1**, and the **AD2** is an independent request. All three transactions are completed in 12 cycles, decreasing performance for 41 percent (*at a minimum 7 cycles in Wishbone II*) even at additional silicon cost, a memory burst logic implementation on both sides: source and destination.

Imagine a continuous burst Wishbone II would perform with 0 wait cycles (latency is completely removed) and absolutely no loss (again 0 wait cycles) at the slave side when more than just one master coexists in the system for issuing the first word. To be more illustrative for a system running at 150 MHz, long bursts with fixed latency of 2 cycles would yield a Wishbone II bandwidth of 150 Mwords, and Wishbone B.3 of 50 Mwords only.

#### Read-modify-write cycles and exclusive bus/Address locking

A read-modify-write cycle can be made using the bus **LOCK** signal by issuing the read request and **LOCK** signal set, waiting for the read response, followed by a write, and finally releasing the **LOCK** afterward. To not stall the complete bus, Wishbone II introduces a per-cell memory locking feature using the **ALK** signal, which is used in almost the same way as Wishbone **LOCK** signal, just that it doesn't stall the complete bus but grants exclusive permissions to a given module distinguished by the source **TGA**.

#### Wishbone II races into the future

The Wishbone II bus proposes an advanced transaction bus-oriented architecture for SoC designs for FPGAs and ASICs in which architecture write and read operations are handled as separate write and read transactions. Each transaction is stored in a single line, and the multi-pipeline architecture acts as a FIFO buffer transporting multiple transactions from and to multiple source and destination modules. An advanced locking mechanism prevents the complete bus from stalling due to the RMW cycles using a temporary per-cell locking mechanism. In this way, overall design data throughput is increased just up to the maximum while the design successfully integrates slow- and high-speed, low- and high-latency peripherals and CPUs. ⚡



*Uros Platise has been R&D manager for more than 10 years at andEuros, specializing in electronics, robotics, and software engineering. His expertise includes FPGA architectures, price/performance optimizations, communication protocols, sensor networks, and so forth. He will receive his PhD in Isotropic Networks, from JSI in Ljubljana, Slovenia. Uros can be contacted at [uros@andEuros.org](mailto:uros@andEuros.org).*

andEuros • +385-52-777-341 • [www.andEuros.org/erd](http://www.andEuros.org/erd)



To find out more about Wishbone II, visit [www.andEuros.com/erd](http://www.andEuros.com/erd).

## Fanless & Customizable up to Pentium M 1.8 GHz soldered ECC SDRAM



typical sample  
others on request

## IP67 Military PC

For operation temperature between -40°C up to +75°C!  
Rugged, compact and reliable solutions.

WIND RIVER  
VxWorks



[info@mpl.ch](mailto:info@mpl.ch)

**MPL**  
High-Tech · Made in Switzerland

[www.mpl.ch](http://www.mpl.ch)

MPL AG, Täfernstr. 20, CH-5405 Dättwil/Switzerland  
Phone +41 56 483 34 34, Fax +41 56 493 30 20



# Thinking inside the box: Boosting the effectiveness of air cooling

By David Lippincott

*Cooling is an integral part of any system, and particularly those used in harsh environments. David discusses cooling system parameters as applied to an industrial computer installed in a military environment where the high ambient temperature and airborne dust will impact system reliability. Equations describing electronic component life versus temperature and system airflow, along with inductance, are examined. A case study is also presented, illustrating that internal system temperature considerations are critical.*

System software performance continues to be a critical factor in the embedded military space. In the end, however, environmental issues will often determine the success or failure of integrated computing platforms installed in mission-critical applications such as Command and Control, modeling and simulation, and communications. Computers intended for military applications are subjected to a broad range of punishment, including elevated ambient temperature. While shock, vibration, and power quality typically cause immediate failure, operating at elevated temperatures is more insidious, shortening the system's Mean Time Between Failure (MTBF). Studies have demonstrated that component life is typically reduced by 40 to 50 percent for every 10 °C increase in operating temperature. Thus, it is paramount to reduce the temperature of all components inside a computer system as much as possible. Filtering out the environment, calculating the airflow, and factoring in impedance are vital to conquering this challenge.

### The high-temperature challenge

In understanding the effects of increased temperatures, system life expectancy decreases exponentially with increased temperature. The extent of this degradation is generally governed by the Arrhenius equation, which describes how component age accelerates with increasing temperature.

### Arrhenius Equation

$$A_f = e^{\left( \frac{E_a}{k} \left\{ \frac{1}{T_u} - \frac{1}{T_t} \right\} \right)}$$

$A_f$  = acceleration factor

$E_a$  = activation energy in electron-volts (eV)

$k$  = Boltzmann's constant ( $k = 8.617 \times 10^{-5} \text{ eV/Tk}$ )

$Tk$  = Kelvin

$T_u$  = reference junction temperature, in degrees Kelvin

$T_t$  = junction temperature during test, in degrees Kelvin

$e$  = 2.71828 (base of the natural logarithms)

Simplified, this equation supports the rule of thumb that component life halves for every 10 °C rise in component temperature. Thus it is easy to see that MTBF goes down significantly with any increase in temperature beyond the original calculations.

Providing sufficient cooling, especially for high-power multi-processor systems in the present combat theater, requires a disciplined approach to cooling system design. The end result is a combination of several smaller contributions where any degradation in any one will impact system reliability.

### Filtering out the environment

Military grade computers expected to operate in ambient temperatures in excess of 120 °F will employ multiple fans mounted either at the front or near the center of the chassis as dictated by the installed components. These configurations provide high airflow across a wide area as well as a positive pressure within the chassis. Mounting fans at the rear of a chassis as exhaust fans induces a negative pressure inside the chassis, which will pull unfiltered dirty air in through any opening in the chassis.

A filter should be installed in the inlet path to remove dirt from the air stream. Even small amounts of dirt, in combination with moisture, can be conductive and cause intermittent operation. Dirt acts as a heat insulator and, even in small amounts, will somewhat degrade cooling effectiveness. In larger amounts, dirt can completely obscure the passages in the heat sinks on the processors and chipsets, reducing the cooling to a small fraction of that required by these devices. Any insulative effect will raise component temperatures and lead to a shorter life. Software operation may also be impacted if the processor goes into throttle mode because of overheating, automatically reducing the clock rate to reduce power consumption and heat generation.

Specification of the filter media and size is important because a filter reduces the airflow through a system. A trade-off in filter efficiency versus airflow is required to assure adequate system cooling flow. Fans with higher-pressure capability should be specified for use with filters. Filters also need regular cleaning as the trapped dirt will increase the pressure drop across the filter. A filter can become completely blocked, reducing airflow through the chassis to almost zero.

An important design criterion is completely sealing the fan bulk-head including cable passages to prevent air recirculation within the chassis. Improper sealing will allow higher-pressure hot air in the rear of the chassis to circulate around to the inlet side of the fans, significantly degrading the cooling effectiveness. Not only is the airflow through the chassis reduced, the air being circulated through the chassis will be hotter than external air. These openings do not have to be very large, on the order of a couple of square inches combined area, to reduce the airflow through the chassis to almost zero. The result is a chassis poorly cooled only through the skin and the small amount of flow through the power supply.

### Figuring the airflow

One factor to consider is the impact of using smaller fans, such as in cases where chassis height must be reduced. Smaller fans provide reduced airflow in a nonlinear fashion. Examining the Fan Laws, which describe the basic relationship between fan speed, flow, pressure, and power, shows  $Q \sim ND^3$ , where Q is airflow, N is speed, and D is diameter. Thus, airflow is governed by the blade diameter cubed. That is, a fan half the size will move one-eighth the air for the same RPM. Alternatively, half-size fans must theoretically spin eight times faster to deliver the same amount of air. However, there is a limit to fan speed with the result that fans with adequate flow rates may not be available for high-power systems. Higher-speed fans are louder, and the higher frequency makes them seem noisier. This can be a critical fac-

tor when multiple systems are operating in a confined space and there is a strict noise limit imposed. Therefore, systems should be designed with the largest fans possible given space constraints.

To conservatively calculate the airflow requirements in a configured system, a simplified formula is:

$$CFM = 3.16 \times \text{Watts Dissipated} / \text{allowed } \Delta T \text{ } ^\circ F$$

CFM = Cubic Feet per Minute actual airflow

Watts Dissipated = Actual power consumption of installed components

$\Delta T$  = Allowed temperature difference between inlet and outlet in °F

Note that ambient temperature does not enter into the equation. Higher ambient temperatures will somewhat reduce the output of the fans due to the lower density, but this is a negligible effect. Ambient temperature is important as it provides the baseline for the system temperature. If the ambient temperature rises 20 degrees above the tested temperature, the internal temperature will rise the same amount.

A fan requirement determination first requires an analysis to find the system component with the lowest maximum temperature rating. Generally this is a disk drive, but these are typically installed directly in the incoming air stream and not heated by other system components. But it might also be a plug-in card that will be mounted near and downstream of the hot processors. Or it could be a processor with a small margin to maximum operating temperature. Thus, if a component has a 125 °F limit and the

## Data Storage Technology

*Proven. Rugged. COTS.*



*We Put the State of the Art to Work<sup>SM</sup>*

At Phoenix International, we design and build Rugged COTS Data Storage Systems that plug and play in any application – from multi-terabyte RAID and Storage Area Network configurations to plug-in VME/cPCI Storage Modules. Our worldwide reputation for excellence is earned by manufacturing highest quality data storage products supported by unparalleled customer service and support.



**PHOENIX**  
INTERNATIONAL

Phoenix International Systems, Inc. An ISO 9001:2000 Certified SDVOB  
**714-283-4800 • 800-203-4800 • [www.phenxint.com](http://www.phenxint.com)**



system will be used in an ambient temperature of 115 °F, then the allowable temperature rise through the system, conservatively, would be 10 °F.

As an example, consider a 4U system using 500 W and a target of a 10 °F temperature rise:

$$\text{CFM} = 3.16 \times 500 \text{ Watts} / 10 \text{ °F} = 158 \text{ CFM}$$

This is the airflow through a system, not the free rating of the fan. Even in a well-designed system, fans will output at most 50 to 60 percent of their rated flow. Thus, in this case, a starting point of 300 CFM would be indicated for initial fan selection.

An important point to note is placement of the power supply, in regard to airflow. A rear-mounted supply will exhaust its hot air directly out the rear of the chassis. A front-mounted power supply will exhaust into the chassis. In general, a switching power supply is approximately 70 to 80 percent efficient. This inefficiency needs to be added to the total watts dissipation calculations. System components using 500 W, for example, would lead to total system power usage of  $500 / .70 = 714 \text{ W}$  for a 70 percent efficient supply. This is the number to use for a CFM calculation, but only for a front-mounted supply.

### Chassis flow impedance

Another challenge when selecting fans is to understand chassis flow impedance, which will determine the volume of air that

specific fans can push through a particular enclosure. Chassis flow impedance is defined as the resistance of airflow from the point of intake, through the enclosure, to the point of exhaust. Size and number of vent openings, filter media, internal structure, and installed components all contribute to increased impedance. Due to the infinite variations presented in chassis design and the configuration of installed components, the impedance can only be discerned by measuring the actual flow versus the required pressure. Mathematical tools such as Computational Fluid Dynamics (CFD) can be used to approximate flow. However, a measurement of the actual chassis impedance should be performed to validate system performance and the CFD model.


Additionally, pressure gradients drive airflow. There is low pressure in front of the fans with higher pressure outside the chassis so air flows into the chassis. The fans boost the pressure above ambient so air then flows into the rear and out the back of the chassis. Physics shows that airflow is proportional to the square of the pressure. That is, to double the flow through a system requires four times the pressure. All things being equal, adding a second fan will not double the flow but will increase it, at most, by the square root of 2 (1.41) or about 40 percent. To double the airflow would require four times the pressure.


Flow through a chassis can be approximately determined by laying a fan curve over a chassis impedance curve. The intersection of the two curves gives the airflow.

### COTS I/O Solutions for:


**IndustryPack®, PMC, CompactPCI, PCI**  
with Outstanding Software Support.

- CPU Carriers
- IP and PMC Carriers
- Ethernet
- Communication
- CAN Bus
- Field Bus
- Digital I/O
- Analog I/O
- PC Card/CardBus
- Motion Control
- Memory
- User-programmable FPGA





- VxWorks
- Linux
- Windows
- LynxOS
- QNX
- OS-9



[www.tews.com](http://www.tews.com)

© 2008 TEWS TECHNOLOGIES GmbH, all rights reserved. IndustryPack is a registered trademark of SBS Technologies, Inc. All other trademarks mentioned are property of their respective owners.

TEWS TECHNOLOGIES LLC: 9190 Double Diamond Parkway, Suite 127 • Reno, NV 89521/USA  
Phone: +1 (775) 850 5830 • Fax: +1 (775) 201 0347 • E-mail: [usasales@tews.com](mailto:usasales@tews.com)

TEWS TECHNOLOGIES GmbH: Am Bahnhof 7 • 25469 Halstenbek/Germany  
Phone: +49 (0)4101-4058-0 • Fax: +49 (0)4101-4058-19 • E-mail: [info@tews.com](mailto:info@tews.com)



**Your Leading Partner To Success**

**MIL-STD-1553 STANAG 3838**

**ARINC 453/708**

**ARINC 429/575/572/571/581/582..**

**RS 232/422**

**AFDX**

### Protocols Converter

## ARINC & MIL-STD-1553 SOLUTIONS

### UADI - Universal Avionic Design Interface

# BOARD



→ ISA

→ PCI

→ PC104

→ cPCI

→ PMC

→ VME

→ SBC

→ Custom

→ FOUR MIL-STD-1553 1MHz to 100 MHz

→ FOUR ARINC - 429/575/572...up to 1MHz

→ FOUR ARINC - 453/708 1MHz to 100 MHz



# IC



# IP

## www.bmccorp.com

67 Bond Street Westbury NY 11590 Tel: (516) 997-2118 Fax: (516) 997-2129

Figure 1 plots impedance curves for several different chassis. A restrictive configuration with high impedance is shown as 'A' while an open, well-ventilated chassis is shown as D. B and C show impedance curves for chassis with intermediate impedance ranges. A single fan (black line), two series fans (red line), and two parallel fans (blue line) are overlaid on the impedance curves. The amount of air that will flow through the 'D' chassis can be seen for the three fan configurations, ranging from 66 CFM for a single fan to only 71 CFM for two fans in series. A dramatic improvement to 92 CFM is shown for two fans in parallel.




Figure 1

It is evident from this graph that for a chassis with high impedance, two fans in series provides better cooling; meanwhile, for an open chassis with low impedance, two fans in parallel flow better. The B and C curves show very little difference between series and parallel fan installation, though there is improvement over a single fan.

Note that D impedance curve intersects the parallel fan curve (blue line) at the area of fan instability where minute changes in pressure have a large effect on airflow. Selection of different fans would be indicated in this situation to assure more determinate airflow.

As the airflow increases with an additional fan, the internal pressure also increases, thus reducing each fan's output. Note the fan curve for the second fan in series simply duplicates the curve for a single fan with double the flow for each pressure point. There is no additional pressure available. The same applies to fans in parallel, except flow doubles with no increase in available pressure.

To visualize the effect this would have on chassis flow, imagine a chassis with zero impedance, essentially an open box. Since there is no back pressure, two fans in series won't move more air than a single fan, but two fans in parallel would double the flow. At the other extreme, imagine a chassis with very high impedance, such as a closed box with just a small opening. Two fans in parallel will not increase the pressure, so no additional air would

flow. On the other hand, two fans in series would double the pressure, giving 40 percent more flow through the system. In the real world, chassis impedance is somewhere between these extremes, so that fan selection and configuration is determined by comparing the curves to determine best flow. Another consideration is that fans are not linear devices and have areas of instability. Careful selection of the fans based on the intersection of the curves will prevent operation in these areas of instability and maximize flow through the system.

### Case study: 4U Joint Range Extension system

As a real-world example of cooling analysis and solution for a high-power military computer, consider the problem of installing two single board computers inside one 4U enclosure. When L-3 Com ESD needed a rugged 4U enclosure to support its Joint Range Extension (JRE) program, Chassis Plans engineers designed a comprehensive solution tailored to JRE's unique mechanical and environmental system requirements. Joint Range Extension is a combined hardware and software system that receives battlefield information transmitted on a tactical data link in a particular area of operations, then forwards that information to another tactical Data Link Terminal (DLT) located at a point beyond the line of sight.

Chassis Plans' JRE-DLT (Figure 2) provides two dual processor Xeon XPT single board computers, one running Windows XP and the other Solaris. Due to the level of heat generated by hosting two computers in one enclosure, the remedy for the JRE-DLT enclosure was to create a "wall of air" by mounting four high-velocity 92 mm hot-swap fans mid-chassis and sealing the airflow path to eliminate recirculation. Because this system is used in very dirty environments, a pair of 30 ppi reticulated polyurethane foam filters were provided on the front door.



Figure 2

The system includes two independent backplanes that occupy the full width of the chassis. This dictated the use of a front-mounted power supply. System power requirements were calculated at 450 W. With a front-mounted power supply with 78 percent efficiency, system heat generation was calculated to be 576 W. With a target temperature rise of 15 °F, target system airflow was calculated to be 121 CFM. The fans were rated for 76 CFM each, for a total airflow of 305 CFM free air rating. System impedance reduced the fan output to approximately 150 CFM, meeting the target system cooling parameters with flow margin. Careful attention to chassis design details, in particular airflow paths and fan selection, provided sufficient cooling to assure adequate performance and component life from this high-power system.



### Cooling electronics down when temperatures rise

Punishing thermal environments and demand for reliable system operation require a careful analysis of airflow through a military computer system. If component life is halved for every 10 °C rise in operating temperature, then it is imperative to minimize component temperatures by maximizing cooling airflow. Optimum fan selection depends on the flow resistance imposed by chassis design, air filter media, and installed components; additionally, the best fans to use can be determined by plotting chassis impedance versus the fan flow/pressure curves. Only through complete system engineering including cooling can a system's operational potential be realized. +

## Another low-temp alternative

Blowers are an alternative to fans. The advantage of a blower is the ability to generate very high static pressure. The downside is blowers do not move a lot of air for their size. In a high-impedance system such as a 1U enclosure, a blower will generally provide more airflow than a bank of 40 mm fans due to that high pressure.



*David Lippincott founded Chassis Plans in 1997 as a design firm specializing in rugged industrial computers. In 2001, he reformed Chassis Plans with his partner, Steve Travis, in order to manufacture and integrate those same types of computers, with special emphasis on airflow, heat dissipation, and shock and vibration mitigation. Prior to founding Chassis Plans, David was cofounder and engineering vice president of Industrial Computer Source, a company that manufactured industrial computers and products. He has a degree in Aerospace Engineering Science from the University of California, San Diego, and has spent his career in industrial computing systems, industrial control, and computer design. He can be reached at davidl@chassisplans.com.*

Chassis Plans  
858-571-4330  
[www.chassisplans.com](http://www.chassisplans.com)

*Power  
Connector  
Solutions*

**WINCHESTER**  
ELECTRONICS  
CONNECTING INNOVATION TO APPLICATION™

62 Barnes Industrial Road North  
Wallingford, Connecticut 06492  
Phone: 203-741-5400  
Fax: 203-741-5500  
[www.winchesterelectronics.com](http://www.winchesterelectronics.com)



## Tailored to your *exact* specifications

Whether you need one cabinet enclosure or one thousand, Optima EPS' custom design services will fit you like a glove. Based on proven platforms, Optima's modular design allows customization to be faster, easier, and less costly. With proven strength-to-weight ratios, and MIL-tested rugged designs, Optima can provide the Seismic, MIL-COTS, or Harsh Environment model cabinet you require. Call Optima today – we've got your needs covered from head to toe.

**Optima EPS**  
Cabinets & Enclosures  
An ELMA company

Phone: 770-496-4000 Web: [www.optimaeps.com](http://www.optimaeps.com)  
Email: [sales@optimaeps.com](mailto:sales@optimaeps.com)



### MIL Rugged Cabinets

- Designed to MIL-STD: 461D, 810F, 167, 901D
- Custom engineering support
- Integrated cooling, cabling and power management



### Seismic Cabinets

- Meets Seismic Zone 4 per GR-63-CORE
- Tested to MIL-STD-810F "Deployed Transportation"
- EMC and NEMA-rated versions available



# More than just cryptography: The performance advantages of Suite B compliance

By William Lattin

*Security is the objective of the National Security Agency's (NSA's) Suite B cryptosystem specification. But what about performance? Especially in military applications, where information has to be both secure and available right now, cryptographic protection can't be allowed to interfere with functionality. The ciphers and other algorithms approved for Suite B cryptography achieve tremendous efficiency, holding the potential to enhance device performance. Accordingly, Suite B-compliant key agreement and authentication can be accelerated by capitalizing on the inherent efficiencies of Elliptic Curve Diffie-Hellman (ECDH), Elliptic Curve Menezes-Qu-Vanstone (ECMQV), and the Elliptic Curve Digital Signature Algorithm (ECDSA).*

Security is always a balancing act. For proof, you need to look only as far as today's long airport lines. On one hand, there's the requirement to screen passengers thoroughly; on the other, the need to get people onto planes and headed to their destinations. Stringency and efficiency trade off against each other.

Cryptographers have always been aware of the need for compromise in security. Their job is to encode and protect information in ways that present sufficient computational challenges to opponents but remain decipherable by friendly recipients – in a reasonable amount of time and using available resources.

That said, *compromise* is a relative concept. Today, especially for military embedded systems, what's acceptable in balancing security with functionality, performance, and cost is defined by clear and very high expectations. The compromise is to be as uncompromising as possible. Therefore, accelerating Suite B-compliant key

agreement authentication is necessary and attainable through ECDH, ECMQV, and ECDSA.

### Not for strength alone

The collection of commercial cryptographic primitives approved by the National Security Agency for protecting both Sensitive But Unclassified (SBU) and classified information – aka Suite B – sets the bar high: 128-bit security for SBU information and 256-bit security for classified. RSA, Diffie-Hellman, Triple DES, SHA-1, and other familiar security mechanisms don't make the cut. The approved primitives are ECC for asymmetric (public key) cryptography, AES for symmetric cryptography, and SHA-256/384 for hash operations.

These were chosen not only for their security strength, but also for their ability to deliver that strength efficiently. To match AES at 128 bits, RSA would require 3,072 bits – a computational burden to say the least, particularly compared with

ECC's 256 bits for the equivalent strength. This is obviously crucial for embedded systems, such as remote sensors, which are typically resource constrained.

An example of a sensor network (Figure 1) may be helpful for making this discussion concrete in a military context. Consider deploying a series of remote, wireless sensors around a field camp to monitor the perimeter. Such an arrangement requires many independent devices to interact in a mesh architecture, all of them registering and exchanging information. Giving each sensor the autonomy it needs to do its duty essentially creates multiple access points – and therefore multiple points of vulnerability. To mitigate that vulnerability, each sensor must be authorized and authenticated individually. This has to be done in an economical way both computationally and in terms of power consumption, because the sensors must be compact and must also run on batteries.




Photo courtesy of the U.S. Air Force by TSGT Rick Lisum



How these hypothetical sensors generate and exchange authenticated keys, therefore, becomes a critical consideration. The primitives and protocols approved for Suite B enable these devices to perform all the necessary tasks leanly and with robust security.

### Key agreement: The process in brief

By way of a refresher, key agreement protocols are schemes that require each communicating party to contribute material used in generating a final, secret key for encrypting and decrypting communications. The basic steps are explained in Figure 2.

The authentication piece is absolutely necessary. Participants in a key exchange who share only ephemeral keys cannot know the true identity of the person from whom a key is coming, because ephemeral keys are not bound to the parties in any way.

Suite B mandates that NSA-approved systems use either ECDH or ECMQV for key agreement and ECDSA for authentication.

### About the protocols

ECDH can be used for key agreement between parties that have had no prior




Figure 1

### KEY AGREEMENT PROTOCOLS: THE PROCESS IN BRIEF

#### 1 Key contribution

Each party randomly generates a short-term (ephemeral) public key pair and communicates the ephemeral public key to the other party (but not the private key). These are created offline through a process of fixed-point multiplication.

#### 2 Key authentication

Each party verifies the authenticity of the static key of the other party. This also happens online; the public key certificate of each participant is verified.

#### 3 Key establishment

Each party computes the shared key based on the static and ephemeral public keys received from the other party and the static and ephemeral private keys it generated itself. This is done online using variable-point multiplication.

#### 4 Key confirmation

Each party provides evidence of possession of the shared key to the other party. This step can also be used to confirm its true identity to the other party.

Figure 2

# TRI-M

SYSTEMS

proudly distributes

## TRI-M

ENGINEERING

100MHz PC/104 Module



Featuring the new edition ZF86  
FailSafe® Embedded PC-on-a-Chip  
Dual watchdog timers, Phoenix  
BIOS and FAILSAFE Boot ROM  
Extended temperature -40°C to 85°C

## TRI-M

ENGINEERING

PC/104 VersaTainer



The VT-104 VersaTainer is a rugged aluminum enclosure that can be used as either a PC/104, PC/104+ or EBX enclosure. The solid one-piece extruded body provides dual internal shock and vibration protection.

## TRI-M

ENGINEERING

75 Watt High Efficiency PC/104



75 Watt output  
+5V, +12V, -12V outputs  
6V to 40V Dc input range  
PC/104 compliant

www.tri-m.com info@tri-m.com

1.800.665.5600

HEAD OFFICE: VANCOUVER

tel: 604.945.9565 fax: 604.945.9566



“ The authentication piece is absolutely necessary. Participants in a key exchange who share only ephemeral keys cannot know the true identity of the person from whom a key is coming, because ephemeral keys are not bound to the parties in any way. ”

contact. ECMQV on the other hand, establishes shared secrets between entities that already have trusted copies of each other's public keys.

In ECMQV, both parties still generate and exchange ephemeral public keys, but on receipt each also calculates an *implicit signature* using their own long-term and ephemeral private keys and their other ephemeral public key (Figure 3). Shared secrets are then generated using the implicit signatures and the other party's

ephemeral and long-term public keys. Agreement between the shared secrets can be taken as implicit verification that each was generated by the authentic party. If either party's public keys are not used in the process, the shared secrets won't agree.

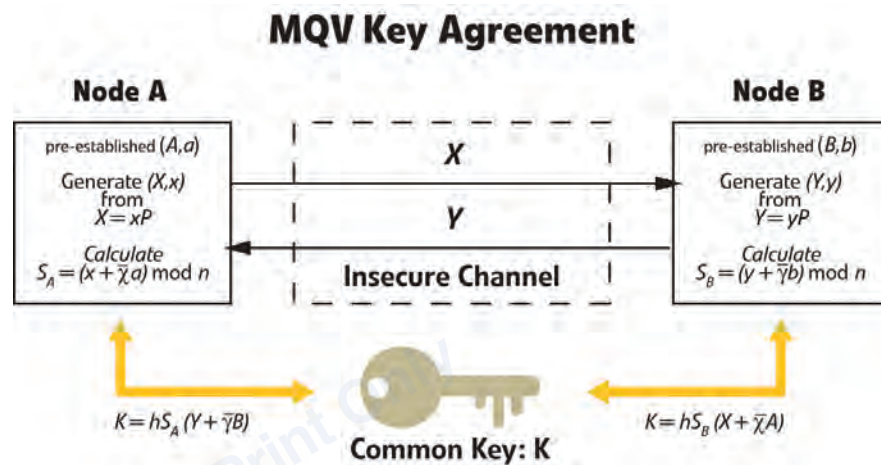


Figure 3

## SOLID STATE STORAGE for Rugged Applications



Connect Tech's FlashDrive/104 provides a durable, high performance solid state disk storage solution for harsh environments.

- 4, 8, 16 or 32 GB storage
- Immune to shock, vibration and extreme temperatures
- Industry standard ATA/IDE interfaces
- 10 MB/second read/write speeds
- Suitable for any platform
- Driverless operation



**Connect Tech Inc.**  
Industrial Strength Communications

sales@connecttech.com | 519-836-1291 | www.connecttech.com

ECDSA authenticates the long-term public keys of the ECDH/ECMQV key exchange through the use of signed certificates that identify and are bound to each participant in an exchange. To validate these certificates, each party uses the public ECDSA *verification* key of the trusted third party (called a *certificate authority*) that created the certificates.

### Streamlining key computation and authentication

Typically, key establishment and authentication are performed as follows:

#### KEY ESTABLISHMENT

Compute expression  $K := aB$

$K$  is the shared secret that will be used to generate the final shared secret key. Note that  $a$  is one party's private key (Party 1) and  $B$  is the other's public key (Party 2). The public keys in this scenario are derived from certificates issued to each user by a certificate authority within a Public Key Infrastructure (PKI). Using certificates provides greater control and management of the system as a whole.

The computational cost of this is a full-sized multiple of the unknown point  $B$ . Computational cost is the largest calculation required for a given function.

#### KEY AUTHENTICATION

Here (and later), we describe the calculations of fast ECDSA verification.

Evaluate expression  $s^{-1}(eG + rQ) - R = O$

$e$  is the hash value of certificate information (including Party 2,  $B$ );  $Q$  is the public key of the certificate authority; and  $(r, s)$  is the ECDSA signature across the certificate information.

The cost here is a linear combination of the known point  $G$  and the unknown point  $Q$ .

Significant economies can be gained, however, by combining key establishment and authentication into a single step.

#### COMBINED KEY ESTABLISHMENT AND AUTHENTICATION

$K' := aB + \lambda(s^{-1}(eG + rQ) - R)$

Here, as above,  $aB$  is the key expression.  $\lambda$  is the randomizer, and  $s^{-1}(eG + rQ)$  is the verification expression (which can also be indicated simply as  $\Delta$ ).

In fact, the entire equation can be generalized, boiled down to:

$K' := K + \lambda \Delta$

This works because Party 1 can only compute  $K'$  if the certificate is "correct" (for example,  $\Delta = O$ ); otherwise,  $K'$  will be random and therefore invalid.

The upshot of this is that a mechanism for implicitly authenticating keys is built right into the key establishment step, accelerating the entire key agreement process by eliminating half of what are known as *point doubles*: pairs of points along a curve. Computationally, the cost is the linear combination of known point  $G$  and unknown points  $B$ ,  $Q$ , and  $R$ .



PC/104 Can-Tainer



CT104

Rugged anodized aluminum PC/104 enclosure designed for harsh environments.

Isolating shock mount and an internal stack vibration mount provides maximum protection from high frequency vibrations and low frequency G-forces.

108 Watt PC/104+ Power Supply



HE104+DX

+3.3V, +5V, +12V & -12V DC output

6V to 40V DC input range

High Efficiency up to 95%

PC/104 compliant

Extended temperature: -40°C to +85°C

168 Watt Max with HPS-UPS firmware.



HPS3612

Total power: 168 Watt with ATX interface

+3.3V, +5V, 12V outputs

6V to 40V DC input range

PC/104 size and mounting holes

Built in temperature sensor

www.tri-m.com      info@tri-m.com

**1.800.665.5600**

HEAD OFFICE: VANCOUVER

tel: 604.945.9565      fax: 604.945.9566



### Working out the gains

Tables 1, 2, and 3 give a clear picture of the processing speed gained by combining key establishment and authentication. Further acceleration is achieved through the use of fast ECDSA verification. Ordinarily, EDCSA verification looks like this:

1. Compute  $e := h(m)$
2. Compute  $R' := (e s^{-1}) G + (r s^{-1}) Q$
3. Check that  $R'$  maps to  $r$

Fast verification speeds up the computation by *reconstructing*  $R$  from  $r$  in step 2 and then checking that  $R = (e s^{-1}) G + (r s^{-1}) Q$ . Tables 1 and 2 show the difference.

The overall total costs of acceleration are depicted in Table 3.

The label *static ECDH* in the charts refers to ECDH schemes that include authentication, as opposed to *anonymous* ECDH schemes, which do not.

A final comparison illustrates just how significant the efficiencies of the ECC-based protocols truly are – specifically ECDSA certificates compared to RSA certificates. As Table 4 shows, the higher the security level, the greater the benefits of ECDSA both in size and speed.

Importantly, no measure of security is lost by streamlining key establishment and authentication: The end product is every bit as strong as that of the underlying DH-based key agreement scheme or ECDSA signature scheme. At the same time, computations for all NIST prime curves (designated for use by the U.S. National Institute of Standards and Technology) are accelerated through the converged process by 45 percent for ECDH in combination with ECDSA and 40 percent for ECMQV in combination for ECDSA. ECDSA itself is sped up by as much as 2.4 times the norm compared to situations in which it is used on its own.

All of these are advantageous enough, but the combined process affords the added benefit of stronger implementation security, providing simple side-channel resistance virtually for free, as a by-product.

#### Static ECDH with ECDSA certificates

- **Curve:** Nist prime curve P-384 with 192-bit security (Suite B)
- **Integer representation:** NAF, joint sparse form (JSF)
- **Coordinate system:** Jacobian coordinates

P-384 CURVE	ECC Operations	ECDH Key	ECDSA (incremental cost)		
			SEPARATELY		Combined with ECDH
			Ordinary	Fast	
	Add	128	194	196	195
	Double	384	384	192	–
	Total <sup>1</sup>	393	459	328	195

<sup>1</sup> Normalized (add/double ratio: 0.69)

#### Speed-up incremental cost ECDSA verify

compared to separate approach: 2.4x (ordinary ECDSA verify)  
1.7x (Fast ECDSA verify)

Table 1

#### ECMQV with ECDSA certificates

- **Curve:** Nist prime curve P-384 with 192-bit security (Suite B)
- **Integer representation:** NAF, joint sparse form (JSF)
- **Coordinate system:** Jacobian coordinates

P-384 CURVE	ECC Operations	ECMQV Key	ECDSA (incremental cost)		
			SEPARATELY		Combined with ECMQV
			Ordinary	Fast	
	Add	194	194	196	196
	Double	384	384	192	–
	Total <sup>1</sup>	459	459	328	196

<sup>1</sup> Normalized (add/double ratio: 0.69)

#### Speed-up incremental cost ECDSA verify

compared to separate approach: 2.3x (ordinary ECDSA verify)  
1.7x (Fast ECDSA verify)

Table 2

#### Static ECDH and ECMQV with ECDSA certificates

P-384 CURVE	Total ECC Operations <sup>1</sup>	Key Computation	Key Computation + ECDSA (total cost)		
			ECDSA SEPARATELY		ECDSA Combined
			Ordinary	Fast	
ECDH		393	852	721	588
ECMQV		459	918	787	655

<sup>1</sup> Normalized (add/double ratio: 0.69)

#### Speed-up total cost ECDH + ECDSA

compared to separate approach: +45% (ordinary ECDSA verify)  
+23% (Fast ECDSA verify)

#### Speed-up total cost ECMQV + ECDSA

compared to separate approach: +40% (ordinary ECDSA verify)  
+20% (Fast ECDSA verify)

Table 3

#### Incremental verification cost of ECDSA certificates vs. RSA certificates

- RSA: public exponent  $e = 2^{16} + 1$
- ECDSA: NIST prime curves
- Platform: HP iPAQ 3950, Intel PXA250 processor 400 MHz

Security Level (bits)	Certificate size <sup>1</sup> (bytes)		Ratio ECC/RSA certificates	Verify – incremental cost (ms)			Ratio ECDSA verify vs. RSA verify
				RSA <sup>2</sup>	ECDSA		
	ECDSA	RSA			ordinary <sup>2</sup>	combined <sup>3</sup>	
80	72	256	4x smaller	1.4	4.0	1.7	0.8x smaller
112	84	512	6x smaller	5.2	7.7	3.2	1.6x smaller
128	96	768	8x smaller	11.0	11.8	4.9	2.2x smaller
192	144	1920	13x smaller	65.8	32.9	13.7	4.8x smaller
256	198	3840	19x smaller	285.0	73.2	30.5	9.3x smaller

<sup>1</sup> Excluding (fixed) overhead of identification data

<sup>2</sup> Certicom Security Builder

<sup>3</sup> Estimate

Table 4

Looking at the incremental cost of signature verification – the most valid metric for evaluating efficiency in this context – the formerly accepted advantage of combining RSA certificates with ECDH schemes is essentially nullified; even at 80-bit security, ECDSA takes the lead.

#### Back to the real world

How does all of this apply to our example of remote sensors monitoring a base-camp perimeter? By taking advantage of implicit verification through a streamlined key establishment and authentication process – and by capitalizing on the option of fast ECDSA verification – each sensor can perform all the processing needed to participate securely and independently in the mesh network. Due to the leanness of this key-agreement approach, system resources on each sensor device are conserved, security function processing is accelerated, and power demands are reduced. No sensor is a “weak link” in some larger security construct; each one upholds the mandatory security requirements for the whole.

Combining verification and key computation is applicable beyond the kind of key-agreement scenario described so far. It can also be used to execute key computation with ECDH schemes in ANSI X9.63, NIST SP800-56a [including Elliptic Curve Integrated Encryption Scheme (ECIES), Unified Model, STS, ECMQV, and ElGamal encryption]. Nonsecret ECC points can be calculated by the method described, provided that the correctness of the solution can be checked. The same is true for computing multiple ECC points.

The streamlined approach can also be used to verify multiple ECDSA signatures (called *certificate chains*) as well as any elliptic curve equation and even multiple elliptic curve equations in batches. Finally, it can be used to carry out operations in other algebraic structures, including hyper-elliptic curves and identity-based cryptosystems.

In the constant struggle for balancing security with functionality, the Suite B primitives and protocols give developers the freedom to load up the scales on both sides, allowing them to explore the full potential of their applications with the confidence of strong security. ✚



**William Lattin** is chief technology officer of Certicom Corp. During the past nine years, he was the managing director of SecureField, an information security

consultancy that specializes in cryptographic product design and network security. He is chair of the Standards for Efficient Cryptography Group (SECG), an industry consortium that develops commercial standards to facilitate the adoption of efficient cryptography and interoperability across a wide range of computing platforms. He holds a BSEE and an MSEE from Stanford University and UC Santa Barbara respectively. He can be reached at [blattin@certicom.com](mailto:blattin@certicom.com).

Certicom Corp.  
650-655-3950  
[www.certicom.com](http://www.certicom.com)

## PCI/104-Express with Core™2 Duo PCI Express

Applications with two displays video streaming HD-videos



### MICROSPACE® MSM945 incl. SMX945

- PCI/104-Express baseboard
- Intel® Core™ Duo LV L2400 / Intel® Celeron® M ULV423
- 1.0GHz to 2x 1.6GHz
- i945GM-PCI Express chipset
- CRT, SDVO, 224MB VRAM
- KB/MS, FD, 1x P-ATA, 2x S-ATA, 2x COM RS232, LPT1, 8x USB 2.0, 1x LAN 10/100 BASE-T, AC97-7.1 HDA
- DDR2-RAM 256 – 2048MB
- IDE
- Watchdog
- Power 5W / typ 10-20W
- Smart cooling concept
- -25°C to +60°C (Option -40°C to +70°C)

DIGITAL-LOGIC AG offers reliable Embedded Computers in PC/104, 3.5", EPIC, EBX, smartModule, COM Express and other formats.



[www.digitallogic.com](http://www.digitallogic.com)  
**DIGITAL-LOGIC**  
smart embedded computers



# Flash FPGAs: Often imitated, never duplicated

By Hezi Saar

*Not all “nonvolatile” or “flash” FPGA devices are created equal. Explored here, the benefits of true nonvolatile FPGAs – including significantly reduced power consumption, faster response times, unparalleled reliability, and uncompromising security – cannot be duplicated.*

Nonvolatile flash memory brings many benefits to FPGAs. True nonvolatile flash-based FPGAs, or those that contain a non-volatile FPGA array, enable significantly reduced power consumption, offer faster response times, and deliver unparalleled reliability and uncompromising security.

Validating the merits of these advantages over their own volatile devices, several suppliers of SRAM-based FPGAs claim to offer “single-chip, flash-based” solutions. These “hybrid” solutions are merely combinations of flash memory components with the underlying SRAM FPGA technology – either integrated with the FPGA die into a single package or, alternatively, stacked or placed side by side. Unfortunately, the FPGA array is still volatile and is subject to the power, reliability, security, and slow power-up drawbacks associated with these types of devices.

Certainly, both the Silicon-In-Package (SIP) and the MultiChip Package (MCP) “hybrid” approaches overcome some of the limitations of traditional SRAM-based solutions by providing a smaller footprint, a minor reduction in power consumption, and small advances in power-up time and security. But these are only incremental improvements over their pure SRAM-based brothers. In order to realize the full benefits of a true flash-based solution, it’s important to understand the primary differences between

these “hybrid” approaches and true flash-based FPGAs, as well as the benefits that true flash-based solutions offer.

### Power matters

Perhaps the single biggest benefit of a true nonvolatile flash-based FPGA array is the significantly reduced power consumption. As mentioned earlier, hybrid solutions simply combine flash memory and SRAM-based FPGA die together, meaning that the inherent architecture is still SRAM-based. As a result, they are subject to the well-documented leakage current issues and power spikes

during system initialization associated with SRAM-based FPGA solutions.

Figure 1 illustrates the power advantages of a true flash-based solution, the lowest power programmable device on the market today, over SRAM-based solutions. The comparison is based on equivalent 15k system gates or 128 macro-cells. Competitor A is an SRAM-based hybrid CPLD and Competitor B is an SRAM-based “low-power” CPLD. The 5 microwatt true flash-based FPGA exhibits 10x lower static power than the alternatives.




Figure 1



## True “live at power-up”

True flash-based FPGAs retain configuration memory while their flash-based SRAM counterparts must rely on flash memory for configuration when the power is turned off. As a result, they not only consume more power, but also cannot achieve “instant on” Live-At-Power-Up (LAPU) status. “Instant on” means that devices are operational as soon as the system voltage has reached its minimum level. Hybrid solutions can be configured more quickly than traditional SRAM-based FPGAs (Figure 2), though they are still up to 40x slower than their true flash-based “instant on” FPGA counterparts (Figure 3). This behavioral profile shows that hybrid device configuration takes more than 200 milliseconds, causing a delay in operation after power-up. Meanwhile, true flash-based FPGAs are live-at-power-up. Thus, true nonvolatile

FPGAs are active immediately after the voltage trigger and operational before power-up.

## Reliability

Hybrid FPGAs also present reliability concerns. Because the underlying architecture is still SRAM-based, these solutions are subject to radiation effects. High-energy neutrons are present in the atmosphere at ground level and can impact the logic modules and routing matrixes of SRAM-based FPGAs, causing firm errors and complete system failures. True nonvolatile FPGAs deliver critical firm error immunity for system-critical and mission-critical functions.

## Preventing malicious attacks and hacking

While the security threat has been lessened by moving the flash memory into the

**“Imitation is the sincerest form of flattery.”**

— Charles Caleb Colton (1820)

same package, the transfer of configuration data from flash memory to the SRAM portion of the hybrid device still makes that data stream vulnerable to hacking and malicious attacks. Comparatively, true flash-based FPGAs don’t require reconfiguration at power-up, eliminating a serious risk. With additional security, such as a 128-bit Advanced Encryption Standard (AES) decryption core, true flash FPGAs give IP providers peace of mind and help designers guard against security issues like cloning, reverse engineering, and denial-of-service attacks.

## Imitation doesn’t equal duplication

Clearly, not all “nonvolatile” or “flash” FPGA devices are created equal. Some devices labeled as “flash” simply integrate flash together with an SRAM-based FPGA to minimize the footprint. True flash-based FPGAs are those that contain a nonvolatile FPGA array, enabling significantly reduced power consumption, improving faster response times, and delivering unparalleled reliability and uncompromising security. The advantages of a true flash FPGA solution can be imitated, but cannot be duplicated. ⊕

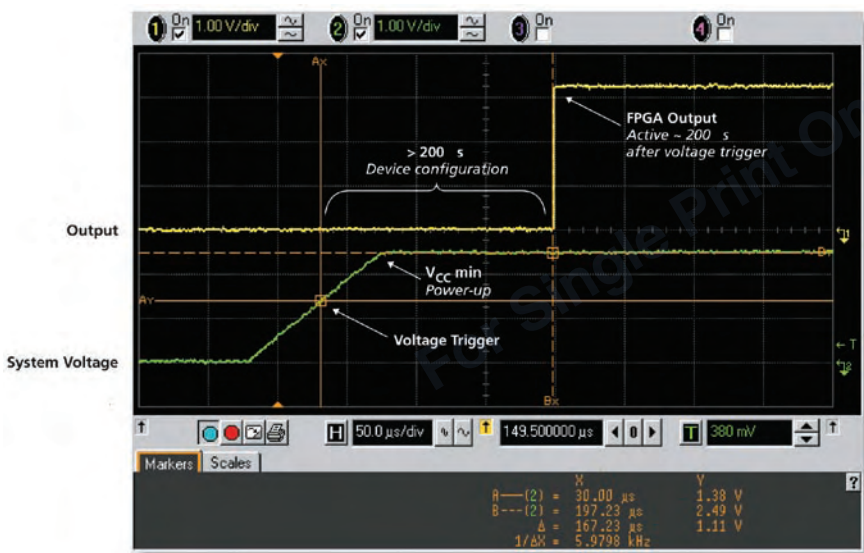


Figure 2

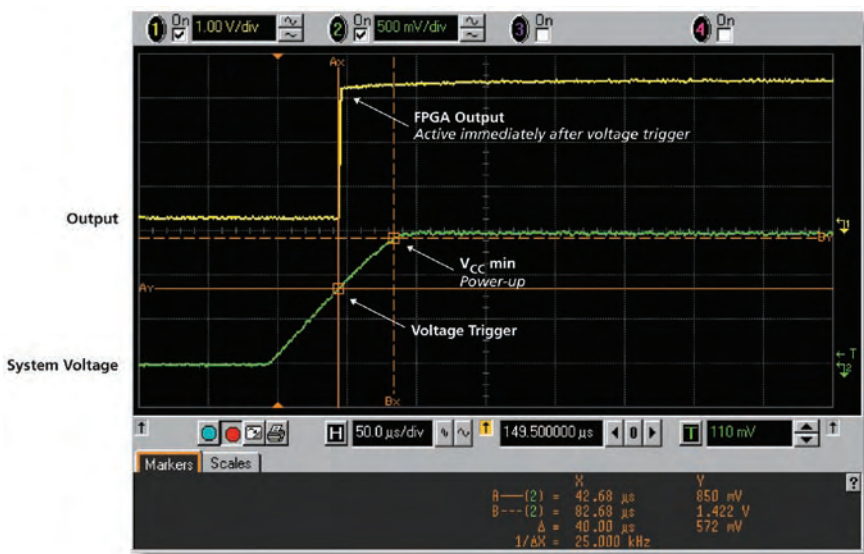


Figure 3



**Hezi Saar** is a senior product marketing manager responsible for Actel’s ultra-low power IGLOO as well as its third-generation

ProASIC3/E/L flash FPGAs. He has more than 12 years of experience in the semiconductor and electronics industries in embedded systems and marketing positions. Hezi holds a B.Sc. from Tel Aviv University in Computer Science and Economics, and an MBA from CSU. For more information, e-mail [articlefeedback@actel.com](mailto:articlefeedback@actel.com).

**Actel Corporation**  
650-318-4200  
[www.actel.com](http://www.actel.com)





# Remote backup: A lifeline to businesses' greatest asset

By Rob Cosgrove

### Save your work!

The single most valuable asset of most businesses is not the inventory of products or services that they provide, or the process by which they deliver them. It isn't the hardware and machinery that they employ in their efforts. It is not even the building the business is housed in, nor the various holdings of the company. The most valuable asset, by far, of today's businesses – is data. Stored innocuously on multiple computer hard drives scattered throughout the enterprise, this data is the lifeblood coursing through the veins of a business entity. Without it, statistically speaking, these companies are simply dead. A comprehensive remote backup plan can dramatically shorten the downtime associated with data loss, and thereby lengthen the lifespan of most companies.

Consider the integral role that computer data plays in the daily workflow of any business. Customer databases, electronic communications and distribution groups, forms, contracts, e-templates, online information stores, and other critical data are the virtual "tools of the trades" being plied by today's companies. Just as many businesses decentralize their materials warehousing operations in order to maintain a backup supply in the event of unexpected loss or destruction of their primary stockpile, thousands of businesses are now turning to automated remote backup as a means of insuring their critical computer data.

### What's the plan?

Processes for smoothing the transition and replacement of key personnel, infra-

structure, and other components of a typical businesses workflow can be invaluable in a time of crisis. Items as basic as building evacuation plans, documented sales and marketing plans, and established processes for replacing key materials and equipment can be used by new and existing employees to further the goals of the company even under extreme conditions. As important as one component in a star-performing division is, in many cases that component can be replaced rather quickly and according to established protocol. If a major piece of equipment ceases to function or a key player on a project team is suddenly out of the picture, contingency plans kick in and the transition process begins. While not always comfortable or even necessarily successful, the fact is that the plan, and at least a process outline, usually exists.

Now consider the processes in place for the safekeeping, replacement, and transition of computer data at most businesses. Many simply don't have a plan or a process in place to ensure the redundant availability or integrity of their critical business data. Especially at Small-to-Medium-sized Businesses (SMBs), the concept of critical data loss and the implications of that loss are simply swept aside in many cases, to be dealt with "later." The pressures of making the sale, completing the process, or simply staying afloat in a competitive market frequently mean that time isn't allocated to consider and plan for this critical aspect of business insurance. Business continuity planning, for many businesses, appears only as dollars in a bank account – not as a critical data backup strategy.

### The cost of inaction

Industry figures show that most data loss occurs in two distinct exposure areas – systems or hardware malfunction, and human error. In fact, some industry estimates have these two areas accounting for 89 percent of all data loss in the business sector. This indicates that the vast majority of data loss events are brought on not by external events such as natural disasters, but rather are caused by the very people and systems that create the data in the first place. Remote backup offsets some of this internal liability by storing complete, accurate data sets in a separate – perhaps less volatile – environment. By securely transmitting a copy of the quality data offsite and employing separate hardware, software, and network resources in the storage of that data, normal catastrophic hazards such as fire, flood, and windstorms are also somewhat neutralized.

### The bottom line

With the proliferation of high-speed connectivity, network resources at virtually any business are suitable for sending large amounts of data "over the wire" to an offsite server for storage. Security protocols of the best remote backup software products, including tight encryption and compression of the data, shorten data transfer times and ensure absolute security-in-transit of even the most sensitive and valuable data. Unlike some legacy backup processes, the better remote backup software and services usually provide the ability to receive verification notices that the backup sessions completed successfully, and also include the ability to restore data quickly, without IT staff involvement. When compared to the painful and lengthy manual reentry of lost data, which can take days (if not weeks or months) to complete, remote data backup is a relatively low-cost way of insuring that your client's data and business can be back online quickly after a catastrophe or other data loss event. ✚



**Rob Cosgrove** is  
CEO of Remote  
Backup Systems,  
Inc., which provides  
Internet-based  
online backups.

Remote Backup Systems  
800-945-4491  
[www.remote-backup.com](http://www.remote-backup.com)



## Looking for a First-Class CompactPCI® Solution that Really Delivers?

Whether you are seeking to reduce development time, advance communications infrastructures, increase efficiencies, or develop high availability solutions, our PICMG® 2.16 products deliver the first-class performance, speed, and reliability required to complete your mission-critical objectives.

By originating the PICMG 2.16 specification, we revolutionized CompactPCI® to meet the need for increased bandwidth in next-generation platforms. We continue to lead today by offering a wide range of 2.16-compatible platforms and blades that provide the building blocks for reliable, high-performance systems.

Your first-class solution is here. Call now, or visit us at [www.pt.com](http://www.pt.com) today for your upgrade.

Get onboard today  
with our CompactPCI® whitepaper!

**CompactPCI + Ethernet**  
The Winning Alternative to Aging VMEbus Systems

Download at: [www.pt.com/whitepaper](http://www.pt.com/whitepaper)



**PERFORMANCE  
TECHNOLOGIES**



**CPC5564**

64-Bit Optron™ Single Board Computer



**CPC6622**

24 1Gb and 2 10Gb Ethernet Switch



**CPC324**

24-Port T1/E1/J1 TDM/IP Edge Processor



**CPC5900**

RAID-ready SATA Storage Blade



### JTRS, FCS, WGS, and DHS:

### Chances are Boeing Network & Space Systems builds it, talks to it, or helps the mission succeed

**An exclusive Q&A with Roger Krone, President of Network & Space Systems (N&SS), a business unit of Boeing Integrated Defense Systems (IDS)**



#### EDITOR'S NOTE

At last fall's MILCOM conference in Florida, AFCEA and the IEEE did a fabulous job creating a convergence between government and industry luminaries. (See the "MILCOM Exposé" Crosshairs Editorial in our November 2007 issue at [www.MIL-EMBEDDED.com](http://www.MIL-EMBEDDED.com).) One of these high-ranking officials was Boeing's Roger Krone, and I jumped at the chance for an exclusive interview with him on topics ranging from TSAT and missile defense, to network-centric warfare and rad-hard technology. Edited excerpts follow. — Chris A. Ciufu

**MIL EMBEDDED:** *Boeing is a veritable household name in military embedded, but can you enlighten our readers about its organization and focus?*

**KRONE:** Boeing is one company with two distinct parts [commercial airplanes and defense]. We're in the \$60 billion range, split about 50-50, although in 2006 the defense business was actually larger than the commercial aircraft business. The Precision Engagement and Mobility Systems business unit of IDS is essentially the airplane company and includes C-17, F-18, F-15, and some weapons. Meanwhile, our N&SS business unit is the network and space side. I always like to say, "We're everything else," with a real bent toward pulling together all our net-enabled, network-centric technology into one group. We also have some things that frankly didn't fit anywhere else, including the missile defense systems division. Then we have something called *Combat Systems*, which is primarily Future Combat Systems (FCS).

**MIL EMBEDDED:** *What are a couple of your present programs?*

**KRONE:** We just won the UK version of FCS called *FRES*, the Future Rapid Effects Systems, which will do the same thing we are trying to do at FCS: take a network and over-layer it on their medium-weight capability programs. We're going to upgrade all their vehicles. And we won the system of systems contract with Thales. Our role will be to do systems integration work, to provide the

transport layer in the network, along with the network management.

We've also started up a new business division called *Intelligence and Security Systems*. It includes the SBI-net border security program, which is the application of net-enabled technology to the problem of protecting the southern U.S. border, on behalf of Customs and Border Patrol (CBP), under the Department of Homeland Security [DHS]. Under "Project 28," we build nine towers across the Arizona border around Sasabe. On the towers are personal radars and then IR and black and white or color TV cameras. We use the towers with a control center to actually detect illegal immigrants coming across the border and then the CBP can dispatch patrols to those locations.

Additionally, we do the combat search and rescue radio called *Combat Survivor Evader Locator* (CSEL) radio, which is the handheld radio that pilots carry in their flight suits. It's actually been used in combat to locate and communicate with soldiers. So it's a pretty exciting thing.

**MIL EMBEDDED:** *What about on the space side of things?*

**KRONE:** We provide technology for a wide range of programs for commercial, civil, and military government customers, including NASA, the DoD, and classified agencies. Our active, major satellite programs today are the Wideband Global SATCOM (WGS) system and the Global Positioning System satellites.

We also have another satellite on orbit, or on its way to its final destination, Spaceway F3. We launched two Spaceways and then changed the mission and used them for DirectTV services, which shows how flexible that design was. This third one is going up and will be used for its original purpose of satellite-based Internet services.

Besides that, we're a major subcontractor to NASA for the space shuttle missions. We built the Orbiters, so we've got the actual space shuttle part of the space shuttle system. We're the prime integrator on the International Space Station. We just won the contract to build the upper stage of the new launch vehicle called the *Aries Upper Stage Production* contract.

**MIL EMBEDDED:** *So, what about space-based routers? Instead of uplinking and downlinking a bunch of channels and bits all the time, what kind of data processing is now available in space?*

**KRONE:** In space, bandwidth is both the Holy Grail and the ultimate limitation. So, having a router in space allows you not to waste that bandwidth, in the sense that only when you're talking or only when you're sending new bits, you can get very sophisticated with data compression.

So routers in space can get very sophisticated with only having to send and update and actually use bandwidth when there's something new that has to be transmitted. You can also say to a router in space, "I want this to go point-to-point to 20 other people." It can also have algorithms

onboard; it can find out if it dropped a bit – and retransmit it without having to go back, in a much quicker timeframe. And it also eliminates having to hop. All this is done on the spacecraft, so you don't have it going up to the spacecraft and down on the ground, then back up, giving it a double-hop term.

**MIL EMBEDDED:** *Tell us about the choices and componentry needed to facilitate routers in space. Surely you don't have the same technology found in a CISCO router in a wiring closet.*

**KRONE:** The satellite that's on its way to orbit right now – Spaceway 3 – has a 10 Gbps capacity router. It's got around a dozen ASICS, unique designs, that are 10 million gates per ASIC. That was designed about five, six years ago, working with IBM. We've had a relationship with IBM for well over 10 years. We're using their design rules and we have a very strict process that we follow.

Besides that, we're moving on with the Transformational Communications Satellite (TSAT). Same partners, same approach and we're taking it up there. And there we're bringing in the algorithms that CISCO uses in its routers for terrestrial programs today. So yes, the same type of CISCO technology found in your wiring closet is also found in our satellites.

**MIL EMBEDDED:** *Routers are a perfect segue into interoperability of multiple assets, few of which were designed to talk to each other. Which technologies are needed to solve this problem?*

**KRONE:** Part of the technology that we're developing in the Joint Tactical Radio System [JTRS] program is for use with multiple legacy waveforms, different pieces of software. So fundamentally, we have an operating environment with a confidential service layer that allows interoperation with different legacy waveforms, so we're able to bring in SINCGARS, EPLRS, and the new Wideband Networking Waveform (WNW). The idea is that by creating this open architecture, more waveforms can be added at a later date, providing users interoperability on the battlefield. So it's Software-Defined Radio.

**MIL EMBEDDED:** *Do you really think that the SCA is the way to go to facilitate that?*

**KRONE:** The hardware architecture and a portion of the functionality between the

different sets of processors within that has a big play in the porting aspects. So if we had more of a standard architecture on the hardware side, it would result in better portability and reuse. By "hardware architecture," I mean the digital signal processor, the general purpose processor, and the gate arrays.

**MIL EMBEDDED:** *How should cross-banding be approached, for example, SINCGARS talking to EPLRS?*

**KRONE:** From an interoperability viewpoint, one of the first things you have to do is have a wide-open front end that goes from 2 MHz or so up to 2 GHz, so now you can see the signals' different waveforms. Now you've got the data from two different waveforms, two different modems, and now it's in a common format. One of the things that we did with our demos was to hook up a PRC 117 to a

“ So yes, the same type of CISCO technology found in your wiring closet is also found in our satellites. ”

JTRS Ground Mobile Radio (GMR) and send over to a SINCGARS RT-1523E that had never talked to each other before. But once we got them into the JTRS radio, the data was digitized, and we just turned around and sent it out. Not only can two operators talk together, a third person can also take part in a conference call. There's a lot of potential there.

**MIL EMBEDDED:** *So, switching gears, what about low power? FPGAs are certainly not known for their low power. What kind of technology leaps and bounds do we need to be able to do more?*

**KRONE:** We can get to lower power by taking advantage of what they've been doing on GMR. When we put together the MILS [Multiple Independent Levels of Security] system, we had a much more severe SWaP requirement. So, to meet the SWaP, we went in and said, "OK, what are the power sources that you guys have?" Then we went ahead and put together some key differences that allowed us to get the power down to what is required in a much smaller configuration.

**MIL EMBEDDED:** *Regarding the GIG [Global Information Grid] and ad hoc networking, what needs to happen in*

*terms of the ability to jump on/off and talk to whomever is available, as long as you can authenticate?*

**KRONE:** You have to have a lot of nodes that can talk to each other. The more nodes you have, the power of the network goes up by the squared factor. So once you get one or two gateways out there, you have an issue where the data is throttled and it doesn't work like cell phone towers. So, again, you must have nodes that talk to each other and waveforms that are compatible with each other.

**MIL EMBEDDED:** *What are the key challenges?*

**KRONE:** The authentication and the information assurance, so that you're protecting that network, are really key challenges in making sure that people who are trying to register are the right people – that they are who they indicate they are.

If you have a pervasive network, you'd better have pervasive security that's distributed and able to protect the network and truncate branches that are questionable or have been compromised. It can heal itself later and rejoin to get back in the network if, in fact, it was a good guy. But if it was a bad guy, you haven't let the rest of the network become infected.

**MIL EMBEDDED:** *What are the top three technologies that Boeing sees as absolutely critical to accomplishing its myriad programs?*

**KRONE:** Multilayer security and fast switching, to name a couple. Besides those, like we talked about, there are always power issues. Sometimes we design great packaging, and then we have a power amp that's twice the size of the radio.

Processor speed is key as well – to have a lot more overhead and speed in the processors to allow you to crunch those algorithms, especially with the security issues. The network protocols are also vital. We're trying to design a network in an *ad hoc* mobile environment, especially the Ground Mobile Radio environment [one of the JTRS domain radios], which turns out to be even more severe with soldiers that go behind hills and go in ditches and break the network and then rejoin. So those network protocols are very important, but not as sexy as some of the other stuff.



**MIL EMBEDDED:** *Are you inventing some of these, or can you use technology straight out of the commercial world?*

**KRONE:** In this latter case, commercial world technology definitely does not work. Everybody says, "Gee, I've got my cell phone or my Blackberry; why doesn't it work like that?" And it's a completely different system and architecture, so it makes sense that the protocols that they have developed for that wouldn't be directly applicable. So we've taken some of the pieces of those where we could, as we demonstrated on TSAT, but they mostly had to be adapted because it doesn't work the same way.

**MIL EMBEDDED:** *My perception of Boeing before today has been "thought leader domain experts, yet no COTS." I don't think your COTS intentions necessarily come out to the world very often.*

**KRONE:** If you look at what we're doing with IBM, that is COTS in a way, but to a

very different extreme. And with CISCO, that's COTS in a way. It's just not "off-the-shelf and you plug it in." But it's getting to the foundation of "how do you leverage IBM, how do you leverage a CISCO – and their investments." It's about being smart, about picking the places you do it in and knowing where you can't.

**MIL EMBEDDED:** *Your competitors on these programs have the same ORD to meet as Boeing. The perception is that your competitors will go buy components at RadioShack and clearly you folks wouldn't.*

**KRONE:** We've always been thought of as the high-end, in-house-developed, big thick piece of intellectual property type of company. That's great for the one power user who will pay to buy the best, but we want to be broader in the marketplace. There are a lot of programs that we didn't talk about today where we've partnered with COTS providers. And we built COTS. SBI-net's Project 28 is an example of that. SBI-net is almost all COTS, and we are integrating it with off-

the-shelf equipment. And we had to go contract and field it in nine months. Well, that's the paradigm for us, probably why we struggled with this. But it's helping us to create a different mindset.

**Roger A. Krone** is president of N&SS, a business of Boeing IDS. Prior to his assignment at N&SS, his most recent roles at Boeing include: vice president and general manager of Army Systems for Boeing IDS, as well as vice president of strategic programs at Boeing's corporate headquarters. He earned a bachelor's degree in Aerospace Engineering from the Georgia Institute of Technology, a master's degree in Aerospace Engineering from the University of Texas at Arlington, and a master's degree in Business Administration from the Harvard Graduate School of Business. Roger, a commercial pilot, is also a member of the American Institute of Aeronautics and Astronautics (AIAA).

Boeing Integrated Defense Systems  
703-414-6461  
[www.boeing.com](http://www.boeing.com)

## DIN 41 612

People | Power | Partnership



[www.HARTING-usa.com](http://www.HARTING-usa.com)

### HARTING DIN 41 612 Connectors

HARTING's comprehensive range of DIN 41 612 connectors includes press-in terminated versions as well as SMC types for high temperature reflow. All components are manufactured in accordance with IEC 60 603-2 specifications.

Mixed Technology versions combine signal/power/co-axial solutions, using the same board space as standard 3 row signal connectors.

The new D20 zinc die cast housing incorporates multiple design improvements, greatly improving handling and assembly processes.



**HARTING**  
Pushing Performance


HARTING North America | 1370 Bowes Road | Elgin, IL 60123 | +1 (877) 741-1500 | [more.info@HARTING.com](mailto:more.info@HARTING.com)




# Rugged High-Speed Data Recording and Storage

SerialFPDP 10GbEthernet Custom

Vortex Data Recording and Playback Systems



Fibre Channel Storage Area Network



Solid State or Rotating Media

## Powerful

- Streaming data recording platforms with up to **720 MB/s** per recording engine sustained recording performance
- Scalable Fibre Channel SAN architecture providing virtually limitless storage capacity
- Ready-to-run application examples

## Flexible

- Customer programmable or application specific recording and playback systems
- VXS (VITA 41), VME, CompactPCI and Industrial PC recording engines
- Solid state or rotating media options from rugged to commercial

## Innovative

- Access and control using web browser or XML-RPC
- Disk grouping and intelligent disk management



Embedded Computing - Data Recording/Rugged Storage - Bus/Protocol Analyzers

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

**VMETRO**   
innovation deployed



# DoD continues refurbishing trend

By Bob Smith



*The mission of the DoD logistics community is to provide worldwide, integrated logistics/supply chain and distribution management, depot-level maintenance management, and strategic prepositioning capability in support of operating forces and other supported units. The goal is to maximize their readiness and sustainability and to support enterprise and program level total life cycle management. However, the Diminishing Manufacturing Sources and Material Shortages (DMSMS) issue continues to relentlessly plague DoD planners, but technology reengineering/obsolescence management is helping to provide a remedy.*

### Prolonged conflicts and limited funding

Because of the Iraq war and its debilitating effect on available funding, new systems cannot be acquired as fast as desired, thereby delaying the deployment of new, technologically advanced replacement weapons. As a result, the DoD is extending the service life of its mature platforms by leveraging emerging technology. This provides solutions to immediate mission needs caused by DMSMS and provides the armed forces with functional weapons and communications and IT systems.

This turmoil results from a significant evolution in the world's political climate during the past 15 years. Starting with the end of the Cold War and the demise of the traditional "super power" paradigm, enemies have become less centralized, often combating any type of formal political sponsorship. The U.S. military has responded accordingly by fighting guerrilla forces or terrorist groups on a smaller, more focused scale. In turn, funds have been redistributed to support these efforts, including the war in Iraq,

occupation in Afghanistan, and other areas of U.S. military deployment around the globe. The end result is that there is little funding left for research, development, and procurement activity involving new weapons.

Therefore, the demand for new and advanced platforms, while current, remains a "want" as opposed to "must-have" capability. Instead, the military has aggressively embraced an initiative to refurbish mature platforms and to refit them to today's standards.

Government-mandated initiatives to refurbish aged weapons systems to save time and money have emerged as the "new" logical solution in the DoD community for supporting the war fighter. However, to extend the service lives of aging platforms, the industry must work diligently to resolve component obsolescence and reliability issues while also addressing the decreasing product life cycle for high-tech components.

In parallel with this premise is the need to introduce state-of-the-art repair method-

ologies that reduce cost and improve reliability. This is especially true in today's war on terrorism with its higher "up-tempo" environment and increased wear and tear on many critical components. These include Air Force, Navy, Army, or USMC systems such as avionics, flight control, and communication equipment that currently require repair or reengineering to solve DMSMS issues and extend their useful life. By making these repairs, or through the use of reengineering applications, outdated systems can be modernized via technology refresh, thereby enabling them to accommodate the military's new tactical demands.

### Technology life-extension alternatives and the DoD

At the moment, there is a small community of aftermarket suppliers that focuses on supporting these life-extension alternatives. The DoD frequently turns to these mid-tier solution providers to fulfill their needs because many of the Original Equipment Manufacturers (OEMs) have disappeared, been acquired by other companies, or are focusing their attention on other system-acquisition initiatives.

These smaller, mid-tier companies are more willing to do the retrofitting and modernization at a lower price and margin, thereby allowing the military to keep aging platforms in the field much longer. The reason they can do this is primarily because they do not have to maintain large-scale R&D and production programs like their OEM brethren, and can therefore offer more “bang for the buck.”

Many times, larger companies will partner with these smaller, more nimble integrators to take advantage of these cost efficiencies and to stay involved in the overall hardware process because the military is not buying as much new equipment.

#### Case study:

##### *Military electronics and obsolescence*

An example that clearly illustrates life-extension alternatives is the replacement of the Bulk Storage Unit Circuit Card Assembly (BSU CCA), shown in Figure 1, for one of the military services. The BSU CCA was used in a portable telephone switchboard system that was initially designed in the early 1980s. The BSU CCA itself was reengineered in the mid 1990s as a replacement for the original tape-cartridge-based memory CCA.

The operating software for the portable telephone switch is stored in a custom 2 MB solid-state flash cartridge that was hot swappable into the BSU CCA. The dynamic telephone network configuration data was also stored in this flash cartridge. The BSU CCA itself was controlled by an onboard Intel 8751 microcontroller that had 4 KB onboard assembly-language-based software. The 8751 microcontroller software was designed so that the 8751 interacted with the portable telephone switch's Z8 processor via a unique communication protocol over the switchboard's backplane. The BSU CCA communication with the processor was implemented with a memory mapped register set, 128K x8 SRAM and 1K by 8-bit




Figure 1

wide input/output data FIFOs, all on the BSU CCA. The switchboard's design predates the proliferation of industry standard circuit card specifications and standards and the widespread introduction of COTS technology into military electronics.

The mating connector for the solid-state cartridge and the cartridge itself became obsolete, and by 2003 the CCA had become a critical military supply issue with stock levels depleted. The BSU CCA suffered from an increasing attrition rate through the many years of the telephone switchboard's operation. The BSU CCA solid-state memory cartridge was handled daily by service technicians in a manual process that called for a second physical solid-state memory cartridge to be swapped into the BSU CCA to affect a backup of the telephone network's dynamic data every four hours.

In 2004, the service contracted to design a replacement: the Memory (MEM) CCA (see Figure 2). The primary design objective was to eliminate the memory cartridge while retaining the ability to have portability of the saved telephone network configuration data. The service also requested the new design provide a different means for reprogramming and testing the new CCA in the field at the point of use. Reprogramming the legacy solid-state memory cartridge installed on the BSU CCA required that the solid-state cartridge be removed from the BSU CCA and shipped to a maintenance facility. Once there, the memory cartridge would be installed in a unique programming interface box driven by a DOS PC application. The process was then reversed, and the BSU CCA with its reprogrammed memory cartridge was shipped back and reinstalled into its host equipment.

The design for the replacement for the BSU CCA combined COTS integrated circuits and USB technology to provide a cost-effective solution. The new MEM CCA replicated the communications protocol with the host processor in the telephone switch by rehosting the Intel 8751 assembly language software in a third-regeneration 8051 variant, the Atmel AT89C51RD. The replication of the hardware-based memory-mapped registers and FIFOs in the BSU CCA was accomplished in a Xilinx 512 macrocell CPLD and 512Kx8 SRAM. All of the BSU CCA's 54 series gated logic was incorporated into the Xilinx CPLD.




Figure 2

To replicate the BSU CCA's dedicated integrated circuit 1Kx8 input and output FIFOs, a FIFO controller was designed into the CPLD. To complete the replication, the 1K memory blocks were partitioned out of the Cypress 512Kx8 SRAMs on the MEM CCA.

To meet the requirement of portable storage of the telephone switchboard's dynamic configuration data, USB 2.0 technology was added to the MEM CCA with a Philips ISP1161. The ISP1161 and new C language USB drivers, written for and hosted by the AT89C51RD, were designed so that the dynamic data could be written in and out of commercial USB memory sticks. The protocol, designed for reading and writing this dynamic data out to the USB memory stick, featured data error detection and correction features and also implemented shadowing of the same data with an onboard Intel 8 MB x 8 flash memory device on the MEM CCA itself. In this way, switchboard operation would not be dependent on the highly portable USB memory stick.

The ISP1161 and new C language USB drivers written for the AT89C51RD also provided the MEM CCA a slave USB port. The MEM CCA's slave USB interface and new PC application, written in LabVIEW and run on a laptop or desktop Windows platform, provided the desired ability to test and reprogram the MEM CCA at its point of use anywhere in the field. The telephone switchboard's host processor software was stored into an Intel 8 MB x 8 flash memory device on the MEM CCA. The protocol for reprogramming the telephone switchboard's host processor software incorporated ping-pong buffers in the onboard flash memory that ensured the MEM CCA would always be programmed to support the switchboard host processor.

The MEM CCA also incorporated a new power architecture compared to the BSU CCA it replaced. The power architecture took advantage of powering the MEM CCA through the slave USB port when



hosted by a PC platform. Through this, the MEM CCA could be tested by the PC-based LabVIEW application without the need for a separate power supply. In addition, the power architecture incorporated power switching and detection such that the MEM CCA could be left in the telephone switchboard and tested while the switchboard's power was off. When the switchboard's power was reapplied to the MEM CCA, however, any testing or reprogramming in progress would be abandoned by the MEM CCA and its

primary role of uploading the switchboard host processor's operating software would be supported.

The MEM CCA is totally interchangeable with the BSU CCA. Since reprogramming can be an almost daily requirement, the new design has eliminated the root cause of the reliability problem because of wear and tear on the connectors. However, as an added benefit, the process of reprogramming has been greatly simplified by utilizing

state-of-the-art USB technology. Under the production contract for the MEM CCA, nearly 500 field modification kits and additional spare MEM CCAs have been shipped. These modification kits have been installed worldwide to resolve the DMSMS reliability issue that led to supply line problems.

The replacement of the BSU CCA with the MEM CCA is a perfect example of effective DoD "tech refresh" in action. The replacement of depleted DoD inventory through an aggressive repair and reengineering program will continue to serve as a cost-effective and efficient way to resolve the growing DMSMS dilemma.

### Reengineering and the DoD's bottom line

While not the only avenue open to the DoD, the reengineering approach has proven itself beneficial. Now system designers can design out the mechanism that led to higher failure rates in the first place, rather than simply building replacement CCAs to the original legacy design (which would have required redesign because of multiple obsolescence issues). It is expected that reengineering will continue to grow exponentially in the DoD, in direct relationship to the dwindling amount of new weapons system development dollars. ✚



**Bob Smith** is senior vice president and general manager, Acquisition Management and Engineering Group (AMEG), Dynamics Research Corporation.

Prior to joining DRC, he served as an aviation officer and UH-60 pilot and maintenance officer with the United States Army. Bob holds a Bachelor of Science degree in Engineering from the United States Military Academy and a Master of Science degree in Engineering Technology from Murray State University. He also attended the Stanford University Executive Education program and is a retired Army reservist. For more information, contact Bob at [RSmith@drc.com](mailto:RSmith@drc.com).

**Dynamics Research Corporation**  
978-289-1822  
[www.drc.com](http://www.drc.com)

## New Rugged, Reliable Servers for Mission-Critical Applications

### 1RU RES Servers

- 1 or 2 AMD® Opteron™ or Intel® Xeon® CPUs (dual- and quad-core processors)
- Up to 32GB ECC SDRAM
- Up to 2 removable HDD
- Vibration – 0.9g from 10 – 2000Hz
- Shock – up to 25G @ 40ms
- 2RU RES Servers also available

### 3RU RES Servers

- 1 or 2 AMD Opteron or Intel Xeon CPUs (dual- and quad-core processors)
- Up to 32GB ECC SDRAM
- Up to 5 removable HDD
- Vibration – 0.9g from 10 – 2000Hz
- Shock – up to 25G @ 40ms

### For Sun® Solaris™, Linux®, and Microsoft® Windows® environments

When it comes to computing in harsh environments, nobody is more at home with its surroundings than Themis Computer. For over a decade, Themis has delivered high performance, high availability computing for the most demanding military, aerospace and communications applications.

Themis' new family of 1RU, 2RU and 3RU Rugged Enterprise Servers™ (RES) includes the latest dual- and quad-core processors from AMD and Intel, offering ruggedized systems with the fastest, widest range of performance options and scalability. Themis servers provide far greater reliability, improved life cycle management and substantially lower TCO than other COTS systems solutions.

For more information on rugged Themis servers and its new family of scalable blade systems featuring innovative CoolShell Technology™ please visit us at [www.themis.com](http://www.themis.com) or call (510) 252-0870

**THEMIS**  
Transformational.

©2008. Themis Computer, Themis, the Themis logo, Rugged Enterprise Servers, and CoolShell Technology are trademarks or registered trademarks of Themis Computer. All other trademarks are the property of their respective owners.

# SOLUTIONS ARE IN SIGHT

Few industries have applications with the complexity, strict safety standards, and low error tolerance of

## DEFENSE & AEROSPACE

Work with an advanced microelectronic  
and display partner who offers:

- Semiconductor packaging for specific circuit board density, weight and environmental requirements
- Hundreds of COTS memory devices and processors for high-reliability defense applications
- Enhanced and ruggedized LCD panels for high-performance aerospace applications
- Interface and electromechanical devices built to precise performance and design specifications
- AS/9100 certified manufacturing facilities



**REVIEW OPTIONS.  
GET ANSWERS.**

Visit [www.whiteedc.com/delivers](http://www.whiteedc.com/delivers)  
or call 602.437.1520.



WHITE ELECTRONIC DESIGNS

[WWW.WHITEEDC.COM](http://WWW.WHITEEDC.COM)



# Remanufacturing thwarts obsolescence in mission-critical systems

By Arlin Niernberger



Photo courtesy of the U.S. Army

*The resolution of obsolescence for mission-critical electronics systems is of utmost importance. Obsolescence issues affect almost every electronics system that is maintained long-term. Arlin describes specific issues and suggests solutions to sustain those systems, focusing on the option to remanufacture obsolete systems. Key considerations include startup and material costs, along with hardware design and program planning.*

Organizations that provide support for obsolescent systems face many challenges that “normal” manufacturing operations do not. Diminishing material resources is perhaps the most critical factor to be considered.

Most aging electronics systems function at a totally acceptable level as designed; yet as the system ages, the Mean Time Between Failure (MTBF) becomes unacceptable, thus jeopardizing the mission itself.

A 2006 GEIA study indicated that throughout the next decade, the average age of military air fleets will exceed 20 years ([http://mae.pennnet.com/articles/article\\_display.cfm?article\\_id=276029](http://mae.pennnet.com/articles/article_display.cfm?article_id=276029)). This indicates that MTBF issues and solutions to reduce those problems will become of greater concern over the next 10 years.

Most end users face the alternatives for support listed in Table 1. No single one of these options is a solution for all situations, but all of these options are usually considered when extending the service life of aging electronics systems. Our discussion will focus on the third option listed in Table 1, *remanufacturing* obsolete systems.

What are the pros and cons of this method, and how do we best implement it? More importantly, how can we make remanufacturing a cost-effective alternative in most cases? How should we plan in advance?

### Benefits of remanufacturing

Most supply chain managers, when faced with the dilemma of whether to

repair or replace obsolete systems, can appreciate the advantages provided by remanufacturing:

- **Existing system software can be used.** No software changes are required to keep the system operating as before. No software development is required or programming time necessary.

### Alternatives for sustaining obsolescent systems

System replacement	As new hardware technology becomes available, it is common to find that the electronics on one redesigned board can replace the electronics in a completely obsolete system. This option may be practical when there are no software changes to be required. However, the cost for design resources, time to implementation, and system requalification are important considerations and may require significant resources.
Limited redesign	Replace only those portions of the system that have become obsolete: The latest developments in software-defined hardware have allowed for module (or board) replacement at an economical cost. However, as cost effective and time efficient as this alternative may seem, requalification would still be required, which may induce considerable cost.
Remanufacturing	Remanufacturing the unit using the same components as the original design: This includes remanufacturing components where necessary. In many cases, this is the best solution for the customer. Remanufacturing a product requires no requalification, software will not be affected, and resources for redesign would not be required.

Table 1

- **Hardware design resources are not required for remanufacturing.** Some components may be required to be re-specified. However, that function is usually supported by sustaining rather than design engineering.
- **There is no need to requalify the system.** In many situations where hardware or software changes are made, the system has to be requalified. Often the original qualification processes must be used. This can be costly as well as time consuming.

## Remanufacturing challenges

On the other hand, there can be disadvantages to the remanufacturing approach, such as IP issues and obsolete/counterfeit components.

### Intellectual Property

Intellectual Property (IP) is the information necessary to manufacture or test the board or system. Pieces of the IP may be missing, incomplete, or incorrect.

An example of this would be where a bill of materials may be incorrect, or "approved suppliers" are no longer in business. Fabrication drawings may be missing, incomplete, or wrong. Diagnostics and test fixtures required to verify the build and to validate final test functions may no longer be available.

### Obsolete or counterfeit components

As systems age, the pool of available components diminishes. Electronic components have on the average a four- to seven-year manufacturing life as compared to a life cycle of 20 or more years for mission-critical systems. As components become in short supply, the demand for components attracts counterfeiters such that few reliable sources remain available. Dealing with unreliable com-

ponent sources becomes a huge burden in time and money to many companies.

This is where legacy companies such as GD California (GDCA) are a critical asset in providing ongoing support of obsolescent circuit card assemblies and systems.

## Weighing the costs

A case for remanufacturing as the most cost-effective alternative is best made by presenting a brief cost analysis. Two primary cost considerations are project startup costs and material costs.

### Startup costs

The numbers generally referenced when discussing the startup costs of obsolescence support are the numbers generated by the Defense Micro Electronics Activity (DMEA). The metrics generated by that organization are used by DoD programs to report cost avoidance associated with implementing the best resolution in line with program requirements and cost constraints ([www.dmea.osd.mil/docs/cost\\_metrics\\_revision1.pdf](http://www.dmea.osd.mil/docs/cost_metrics_revision1.pdf)).

Though the DMEA uses a comprehensive description of issues related to diminishing resources, some legacy companies have expanded cost avoidance to include program preplanning at the design stage, end-of-life buys, and component redesign as important alternatives. Figure 1 lists some DMSMS legacy options available to the end user.

Costs shown on the graph depict actual average program costs relative to specific programs. They may not apply to all vendors' specific projects. However, the costs listed show an accurate cost relationship between the different alternatives and indicate relative costs in performing activities required to sustain a product.

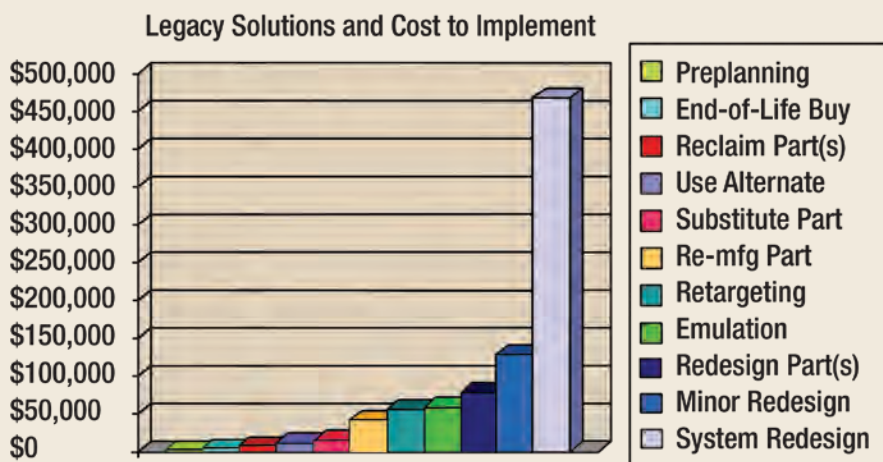


Figure 1

## XMC and PCIe Adapters and Tools for Development, Test Access, and Integration.

# Technobox, inc.®

For details, visit our web site  
[www.technobox.com](http://www.technobox.com)



Logistical options (first six listed) are much more cost effective than the engineering alternatives (last five listed). The last choice listed (major system redesign) is the most expensive option and the least appealing to apply in most circumstances.

## **Material costs**

As components become scarce, two situations affect cost: (1) If the components are obsolete now, component prices can be very high; (2) If the components have not been discontinued yet but soon will be, the prices may be quite reasonable.

## **When a component is obsolete now**

Material costs are related to market availability of the components, and that changes rather rapidly and sometimes unpredictably. Creative sourcing is required where components have been obsolete for many years, and many larger organizations have developed dedicated resources to deal with locating obsolete components.

Developing and maintaining this capability in an organization is in itself an operations cost related to materials that cuts into company profitability. A good course of action is to eliminate as much of this type of activity as possible to reduce costs.

## **When a component is going obsolete soon**

To keep your organization from having to conduct heroic buys on an ongoing basis, it is important to monitor component availability and plan program support. Most companies do this to a certain degree. Numerous software tracking packages make obsolescence tracking fairly straightforward. Once a component shows up as going obsolete, a last-time buy can be scheduled.

Because of the cost involved in last-time buys and the unpredictability of end user requirements, this option may not be practical in some cases. Other options can be pursued before material shortages create a critical support issue. The most practical option is to respecify components where a component that is in current production is used in place of the original. Requalification usually consists of lab and systems testing. The use of alternate components at this point also heads off future issues.

## **Planning ahead**

Legacy solutions can be implemented at any stage in a program. The earlier in the program the initiative is taken, the less expensive obsolescence support becomes.

With that assumption in mind, when designing new products, design criteria should include a plan that encompasses hardware and program planning considerations for obsolescence.

## **Hardware design planning**

Several design rules for obsolescence planning at the design stage exist. One is to place designs that are device independent in programmable semiconductor components. This ensures that migration to new devices is a simple process as long as the intellectual property and design environment are preserved.

Applying circuit design that can be sustained easily over a long period of time is another effective hardware remedy. This may be more difficult to implement in a design but may prove more cost effective in the long term. Selection of components in such a design may consist of going "lower tech" rather than using the latest technology.

A design strategy long taken by many system designers is to use COTS electronics in their system designs. Medical, industrial, and government organizations have made this a preferred path where possible. One great advantage in using COTS products is that many COTS products are supported indefinitely under license by GDCA.

## **Program planning**

After a product has been launched, the forward-thinking project manager may consider obsolescence planning long before a product's components become obsolete. One such process is the AAP.

The AAP is an agreement binding the end user, manufacturer, and a legacy company such as GDCA. It ensures that when the manufacturer stops producing a product, the legacy company will step in to continue manufacturing and repair support over the life of the agreement. This gives the end user uninterrupted product support, an insurance policy against obsolescence.

In cases where there are immediate obsolescence issues, legacy vendors use the PSA agreement. In this instance, they procure and hold end user inventory until required for manufacturing and repair.

In both the AAP and PSA, the end user controls the materials commitment over the remaining project life cycle. If requirements are declining at a diminishing run rate that will become zero in seven years, the end user can lay in enough stock (on the few critical parts required) to last until the projected product life will no longer require support. Should the run rate change, the end user can decide to increase or decrease stock in reserve.

## **Extinguishing the fires of obsolescence**

In general, any organization that has a vested interest in sustaining its equipment long term cannot afford to solve issues that arise in due course by "fighting fires" or waiting until a system goes down to find out repairs cannot be made. Therefore, organizations must act proactively and decisively using cost-effective tools – including preplanning at the design stage and end-of-life programs that ensure availability of components and support – that are readily available to solve obsolescence issues before they arise. ⊕



**Arlin Niernberger** is the director of engineered legacy solutions at GD California in Livermore, California. He has worked in electrical

and mechanical design management for 21 years. In his current position, he delivers legacy solutions to custom medical, military, and aerospace customers. He is also working with the Connecticut Center for Advanced Technology and the U.S. Department of Commerce to provide legacy solutions to military and aerospace organizations. He can be contacted at [arlin@gdca.com](mailto:arlin@gdca.com).

**GD California, Inc.**  
925-456-9900  
[www.gdca.com](http://www.gdca.com)



**HYPERTAC  
HYPERTRONICS**

# PROTECT

## AGAINST SHOCK AND VIBRATION

Global military forces protect their borders with the most advanced equipment available. Exposure to extreme shock and vibration can render standard COTS connectors useless. Hypertronics interconnect solutions featuring Hypertac® hyperboloid contact technology protect equipment from failing in such environments. Hypertronics is the choice of engineers worldwide for mission critical systems.

### Hypertac Features


- Low Insertion/Extraction Forces
- Long Contact Life
- Lower Contact Resistance
- Higher Current Ratings
- Shock and Vibration Immunity

### Benefits

- High Density Interconnect Systems
- Low Cost of Ownership
- Low Power Consumption
- Maximum Contact Performance
- Reliability Under Harsh Conditions

Protect your system integrity with Hypertronics ruggedized interconnect solutions.

[www.hypertronics.com/protect](http://www.hypertronics.com/protect)



VME64x



# Dancing with wolves: The COTS illusion

By Ian Colville



Photo courtesy of the U.S. Air Force by Airman 1st Class Marc J. Lane

*Enhanced Voice Processor (EVP) units are at the heart of many a military communications system. When sourcing Commercial Off-The-Shelf (COTS) enabling technology for these communications, there are always two key requirements that come to the fore: high density and low cost per channel. Although COTS is perceived as a one-size-fits-all solution, COTS products rarely meet military systems' needs all by themselves; often add-on or enabling technologies must be implemented, too. Thus, companies offering COTS equipment for integration into military standard communications systems should be aware of the contradiction in the real-world use of the term "off-the-shelf." Equally, systems suppliers looking to win defense contracts – and needing to source components to fulfill their procurement specifications – should be prepared for Non-Recurring Engineering (NRE) costs for adapting COTS products.*

Military technology has come a long way since the days of William Wallace, who advised his troops before the first battle of Falkirk (on July 22, 1298), "I haif brocht ye to the ring, hop if ye can!" This translates in part as "dance as best you can."

These days, it is no longer a case of simply doing the best you can – making do with what you've got in terms of skill, weapons, and equipment. Today's military forces, from NATO and KFOR to Iraq and ISAF, demand and get the latest in weapons and technology systems. This is also true of communications, with forces around the world being better equipped than ever before in terms of Command, Control, Communications, Computers, and Intelligence (C4I).

Key to these communications systems are Enhanced Voice Processor units (see sidebar below), where interoperability and high functionality are required. However, it is clear that COTS products are unlikely to meet all such requirements themselves, without utilizing "add-on" functionality to provide increased channel density and low cost per channel, for example. Thus, systems suppliers desiring to win defense contracts should be prepared for NRE costs when adopting COTS products.

### Enabling technologies

In light of modern military communications systems' stringent requirements for high levels of functionality, what is typically most important to a defense contractor is a high-density channel count. However, most COTS products cannot meet such stringent requirements on their own. But such requirements can be met via technology "add-ons" to the COTS hardware variant. Often the core enabling technology has to be adapted, such as adding a new codec, for example. And in providing high density, the contractor benefits are twofold: gaining reduced cost and increased value per channel.

Enabling technology encompasses a wide range of hardware and software building blocks for the development of high-performance, wired and wireless, IP- and TDM-based, enterprise and telco communications solutions.

## Enhanced voice processor units

A foundation of the communications element of C4I is the Enhanced Voice Processor (EVP) unit. This forms part of a larger gateway system between disparate networks and is an infrastructure designed to offer more than simply one-to-one gateway functionality. The EVP within this multifaceted gateway often includes voice, fax, data, and video functions.

The multifunction gateway is typically used to provide secure communications at the rear of the battlespace, linking senior commanders back to HQ and forward to the battle area. Systems allow interconnection on a backbone network across a large geographical area as well as the means to interconnect with single service and multinational systems.

Designed to link all the elements of a Joint Force, systems can be deployed in peacekeeping or offensive roles. In practice, systems are fully containerized and can be operated in mounted or dismounted mode. They are also designed to be scalable.





# We Simplify Technology

## CORE SYSTEMS

Reliable operation in severe environments

- Integrates hardware mechanically, electrically and thermally
- Fanless, total thermal solution
- Mini PCI & PC/104 expansion
- MIL-810F, EN50155 available
- Installed Windows® XPe or Linux available
- Virtually indestructible mobile power supply
- NEMA 4X
- -40° to 85° C



- Transportation
- Mining vehicles
- Outdoor kiosks
- Industrial servers
- Digital signage
- Military and police
- Traffic management
- Digital video surveillance
- Asset management

303-430-1500

**OCTAGON**  
SYSTEMS

[www.octagonsystems.com](http://www.octagonsystems.com)





## Arrivals

Origin	Board	Model	Status
FIBRE CHANNEL	PENTEK	4207	NO DELAY
SERIAL RAPID IO	PENTEK	4207	NO DELAY
PCI EXPRESS	PENTEK	4207	NO DELAY
GIGABIT ETHERNET	PENTEK	4207	NO DELAY



## Departures

Destination	Board	Model	Status
PCI-X	PENTEK	4207	NO DELAY
VXS	PENTEK	4207	NO DELAY
PMC / XMC	PENTEK	4207	NO DELAY
ROCKET IO	PENTEK	4207	NO DELAY

# Introducing the One Board with all the Right Connections



Targeting high-performance I/O and processing applications, the Model 4207 delivers an on-time solution for your real-time application using the latest high-speed serial interfaces.

On board PowerPCs, FPGA, PMC/XMC sites, and extensive memory resources are connected to the industry's fastest interfaces including VXS, Gigabit Ethernet and dual Fibre Channel ports for high-speed recording.

All board resources and interfaces are joined through a configurable zero-latency crossbar switch using protocol-transparent gigabit serial data paths.

With so many connections to choose from, you are bound to find just what you need for your next system.

- Dual PMC/XMC Sites
- Freescale MPC8641D Dual Core AltiVec PowerPC Processor to 1.5 GHz
- XC4VFX60 or XC4VFX100 FPGA
- 4 GB DDR2 SDRAM
- PCI Express, Serial RapidIO™, RocketIO™
- Dual Gigabit Ethernet Ports
- Dual 4 Gbit Fibre Channel Interface
- Fabric-transparent Gigabit Crossbar Switch
- VXS, VME64x, Gigabit Ethernet-X
- Linux and VxWorks Support
- ReadyFlow® Board Support Libraries
- GateFlow® FPGA Design Resources



**Call today to have it all!**

201-818-5900 or go to  
[www.pentek.com/go/mesconnect](http://www.pentek.com/go/mesconnect) for the  
4207 brochure, technical datasheet  
and to request pricing.

**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](mailto:info@pentek.com) [www.pentek.com](http://www.pentek.com)  
Worldwide Distribution and Support  
Copyright © 2007, Pentek, Inc. Pentek, ReadyFlow and GateFlow are registered trademarks of Pentek, Inc.  
Other trademarks are properties of their respective owners.



Accordingly, enabling technologies for EVPs embrace two main categories: media processing resources and digital network access. Classic examples of media processing resources include voicemail and Interactive Voice Response (IVR) functionality, where record and playback is used alongside Dual Tone Multi-Frequency (DTMF) detection.

For digital network access, it is important for the vendor to offer connectivity regardless of network type or geographical location. Requirements for COTS amendments could incorporate any or all of a wide range of CAS, H.323, ISDN, SIP, and SS7 protocols. In addition to support for the "vanilla" protocol specifications, the ability to offer national or country-specific variants is vital. Thus, for a contractor, finding a vendor with the willingness and ability to make changes (usually under NRE arrangements) is vital.

### NRE and NDI: One-size COTS doesn't fit all

Contractors are following a trend of incorporating a growing proportion of COTS products to take advantage of technological breakthroughs in commercial markets. These commercial products have the advantage of being more powerful, more adaptable, and less costly than militarized products.

COTS products are characterized by:

- Reduced risk, because the equipment is available and has been proven commercially
- Faster time to market
- The latest technology, because vendors operate in a competitive environment
- Open architecture
- Support of international standards

Notwithstanding this, COTS doesn't often fit the bill "out of the box"; hence, the recurrent need for NRE. COTS procurement is about making effective use of commercially available enabling technology that can be readily adapted to enable the completion of an end-user system. Since one-size COTS doesn't fit all, again, "add-ons" must be implemented.

There is no formula for the percentage of NRE costs; however, it should be clear that, besides the aforementioned benefits, choosing COTS is fundamentally less costly than scratch development.


Of the many selection criteria for COTS – functional, performance, environmental, reliability – choosing the right vendor is key. COTS products offer functionality that often only partially meets a project's specific needs; therefore, a supplier with the temperament to offer an adaptable, responsive approach to Non-Developmental Items (NDI) is critical. An illustration of this required vendor flexibility in a COTS adaption within a military communications application follows.

### Case study: Blending COTS and NRE

One example of how an established COTS product was enhanced through an NRE project to meet high-density requirements within tight timeframes follows:

The client – a major defense contractor in the UK – recently sought a COTS solution featuring an adequate channel count for G.711 A-law, G.729d, and MELPe codecs to develop a voice and

### Industrial-strength A/D/A module, passively cooled



Sure, just about any COTS module can be made to work over an extended temperature: Add fans to cool it and even heaters to warm it. But in many industrial and defense applications, moving parts mean reduced reliability. That's why Sundance Multiprocessor Technology *removed* the fan and standard heat sink on their flagship ADC and DAC module called the SMT370. Instead, a custom-made anodized aluminum heat sink and Faraday cage were added. In a zero airflow environment, the module is rated to 40 °C ambient. Pretty impressive.

Of course, this is partially accomplished with industrial-temperature components. Now dubbed the SMT370-I, there are twin 14-bit ADCs clocked at 105 MHz, along with a pair of 16-bit DACs clocked up to 400 MHz with interpolation. Preprocessing and data manipulation are facilitated via the onboard Xilinx Virtex-II FPGA, and Sundance has partnered with 3L Diamond FPGA to offer cores and other software-configured functions. Two Sundance High-Speed Bus connectors and two 20 MBps comm ports complement user-defined pins used for external connectors. As with all the company's modules, the SMT370 is supported with firmware and development tools from vendors including TI, Xilinx, and The MathWorks.

**Sundance Multiprocessor Technology**  
[www.sundance.com/smt370](http://www.sundance.com/smt370)  
**RSC# 36498**

### Thin MicroTCA chassis, perfect for Humvee

For the past several years, I've noticed a proliferation of rack-mount equipment strapped down or shock-isolated in the back of Humvees. The reason is that the go-anywhere vehicle often serves as the platform for comms gear or other special-purpose electronic systems. But 4U-sized servers and drawer-sized boxes consume way too much space. The 1U MicroTCA MTC5070 from Performance Technologies seeks to change that with an innovative design meant to house six single AdvancedMC modules horizontally, not vertically. Flow-through cooling and internal baffles, complete with two sets of push-pull fans, allow 40 W per slot, a 300 W PSU, and a slim 1U height that's perfect for managed vehicle-mount installations.

Each of the six slots supports PCIe and Ethernet switching, and there is SATA/SAS slot-to-slot connectivity. An integrated dual 10/100/1000 GbE switch provides inter-box communications, while onboard MicroTCA carrier and shelf managers provide PICMG-style management, IMPI, and remote diagnostics. Performance Technologies offers numerous AdvancedMC cards (such as x86 or PowerPC processor nodes), storage, and video cards. The company's NexusWare Linux distribution meets Carrier Grade Linux (CGL) 4.0 requirements. The MTC5070 1U MicroTCA platform can run off of AC or DC power.

**Performance Technologies • [www.pt.com](http://www.pt.com) • RSC# 36463**





data gateway with secure interoperability features to offer into the defense market. The technical challenge was to provide a guaranteed channel count performance, regardless of the combination of codecs in use, within a three-month timeframe.

Having the other codecs already available, Aculab agreed to port the Compandent delivery of MELPe to its COTS product under an NRE project. This was done for an agreed cost, enabling the client to meet its specifications and deadlines, with a high-density, small form factor, RoHS compliant, single board CompactPCI gateway (see sidebar, next page).

The result: the first such compliant device in the UK, with secure interoperability features. It was designed to provide up to 120 channels of secure, high-density voice and data to operate in either a red or black network. It allows many conformant devices to achieve end-to-end encrypted communications between secure networks as illustrated in Figure 1. Supporting both commercial and military standard voice codecs, the gateway converts speech and V.32/V.14 modem data from ISDN (Q.SIG and Q.931-based EuroISDN) to RTP/UDP/IP packet data.

Aculab's Prosody X CompactPCI board was chosen as the enabling technology platform employing Freescale's multicore

StarCore DSP technology, which functions under Aculab's proprietary DSP operating system. The available suite of media processing algorithms operates under the DSP kernel, which readily enables porting of new firmware, such as the MELPe codec.

The board's architecture offered the inherent flexibility needed to cope with the tough demands made by the client's requirements specification. The end result is a classic example of NDI being applied to an industry-proven COTS product through a short-term NRE project. The contractor's project was brought back on track, averting failure or a doubling of costs from using two boards.

### Stand easy

When sourcing COTS enabling technology for military communications, two requirements come to the fore – high density and low cost per channel. However, it is clear that functionality is the overriding consideration as the application or system must clearly deliver on the specification. It is apparent that a responsive commitment to NDI through bespoke NRE is often needed from COTS suppliers. This is likely to be true for media processing resources, the kind of functions found in commercial media servers or gateways.

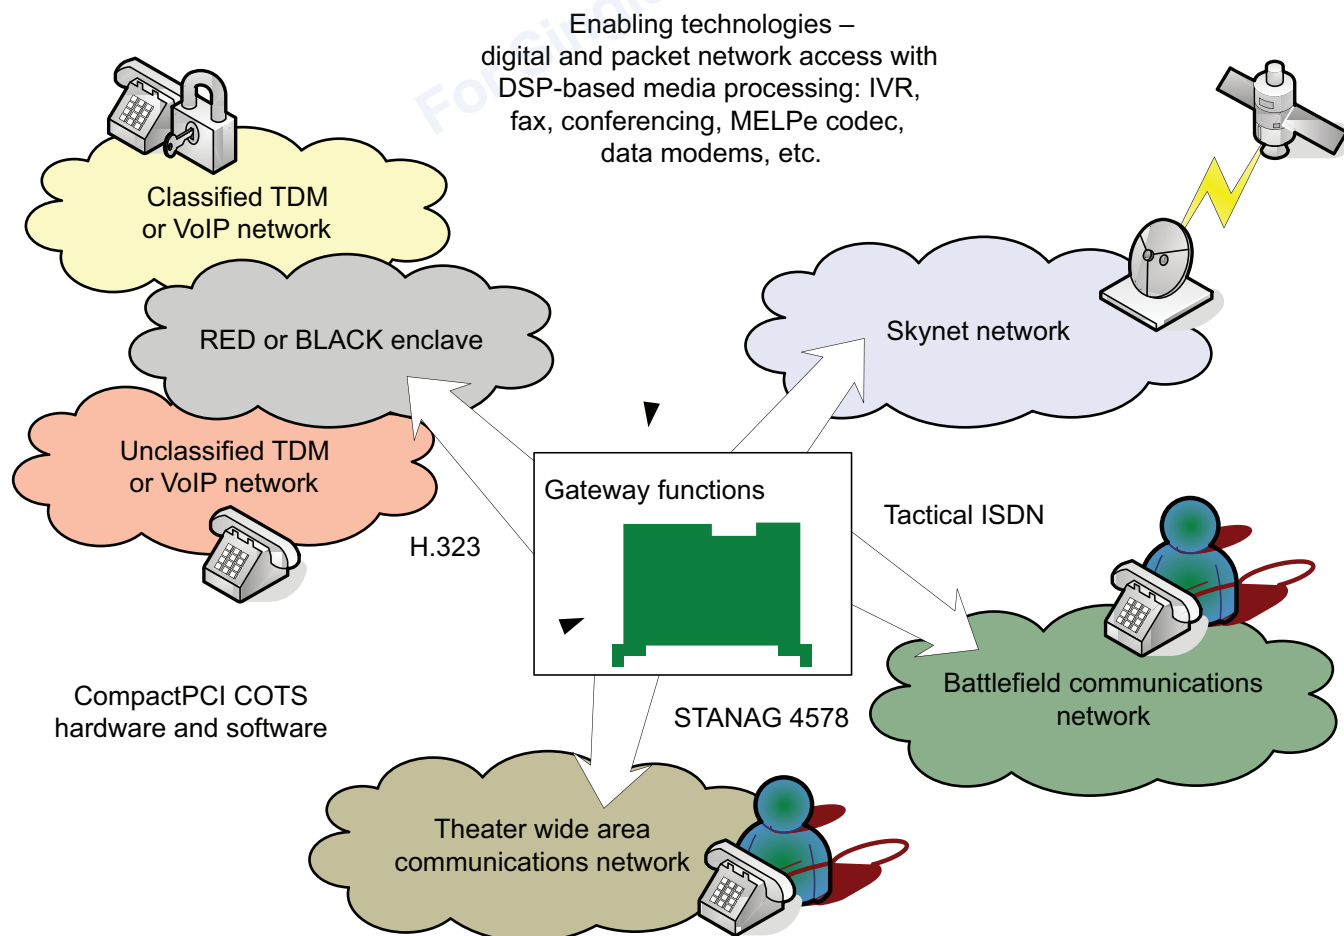


Figure 1

## Gateway function

One example of an essential CompactPCI gateway function is to receive data from an ISDN channel as G.711 A-law and to encode that data as MELPe, G.729d or G.711 A-law, then to packetize the data and encapsulate it within RTP/UDP/IP for transmission over an IP (Ethernet) interface. In the opposite direction, it recovers the raw data from the RTP/UDP/IP packets, decodes the data from MELPe, G.729d or G.711 A-law and transmits it over an ISDN channel as G.711 A-law.

The payload size and format of the RTP/UDP/IP packets generated by the board in the case study met the rigorous demands for data rates, latency (end-to-end delay), Mean Opinion Score (MOS) assessments, and throughput performance stipulated by the client.

MELPe RTP frames contain 22.5 ms of coded audio and are transmitted on 10 ms boundaries (DSPs run on a 10 ms epoch) so as to average 1 packet every 22.5 ms (1 packet at 30, 50, 70, and 90 ms). G.729d frames will contain 10 ms of coded audio and are transmitted every 10 ms. Each MELPe frame and each G.729d frame are transmitted in separate packets to minimize the end-to-end delay of the system. Delay is measured from receipt of each byte to the transmission of the same byte on the corresponding interface and is under 100 ms for both MELPe and G.729d.

The board is capable of buffering up to N frames of data from the RTP/UDP/IP packets whilst the current frame is being transmitted over the ISDN channel. Jitter buffer size is software selectable and adaptive within  $N = >0$  frames.

Dancing as best you can is no longer OK when it comes to military equipment, and companies offering COTS equipment to this sector will realize that "add-ons" must usually be implemented. Equally, systems suppliers looking to win defense contracts should be prepared for those NRE costs. ✚



**Ian Colville** is a product manager at Aculab, where his key role includes support for the company's global sales force. Ian has spoken at a variety of customer seminars on various subjects since joining the company in 2000. He has also contributed technical documentation, including product literature and several published articles. He has broad communications industry knowledge gained during a number of years employed in a variety of management roles by a major telecommunications manufacturer. Ian's industry experience spans marketing, sales, customer service, and project management. He can be contacted at [ian.colville@aculab.com](mailto:ian.colville@aculab.com).


Aculab  
+44-1908-273-800  
[www.aculab.com](http://www.aculab.com)

## Do-it-yourself spy kit

While the title above is intended to be whimsical, the Avnet Electronics Digital Video Surveillance Kit is serious business. Created in conjunction with Analog Devices and Micron Technology, the second-generation stand-alone kit features Ethernet networking, streaming video, and onboard DSP. The kit is but one of Avnet's growing cadre of value-add designs. (At April's Embedded Systems Conference in San Jose, Avnet introduced three more FPGA-based development kits intended to ease designers' headaches and shorten time-to-market.)

The surveillance kit is a simplified prototype development platform that uses ADI's Blackfin DSP for video processing, CMOS image sensors from Micron Technology, and Ethernet for connectivity. Software provides motion detection, threat analysis, and alert notification based upon preset events; images can also be sent over IP. Streaming video is supported, or images and alarms can be sent over the network. The kit is designed for multiple image sensors, from VGA to a 5 Megapixel digital still camera with a dynamic sensor range that exceeds 100 dB — great resolution with minimal noise. The kit also includes an evaluation copy of Analog Devices' VisualDSP++ IDE for the Blackfin processor. Designers can program to their hearts' content, and even deploy the design in an end system.

**Avnet Electronics Marketing**  
<http://em.avnet.com/adi-vs3>  
RSC# 35841




## 8 Mb nvSRAM never loses data

In an era of falling NAND flash ROM prices, many of us carry around over 1 GB of data on our keychains. So why is 8 Mb — a mere fraction of that amount — something to get excited about? The answer is simple: When data is written to these SRAMs and a tiny glitch such as a power spike or other anomaly occurs, the data is safe and secure. In defense systems, such a power interruption *cannot be tolerated* else lives might be lost. Previously, small amounts of serial EEPROMs or battery-backed SRAMs were used to prevent critical data loss, but they were either slow, required CPU intervention of an unforeseen event, or had that nasty battery to maintain.

Simtek's STK14EE8 and STK14EE16 nvSRAMs are available in x8 or x16 widths and use a captured charge to maintain data integrity. Read access is as fast as 25 ns with a 45 ns R/W cycle time. Unlike flash, they feature unlimited endurance, and data is automatically recalled from buried nonvolatile store when the power returns; no CPU is required. Data retention exceeds 20 years, they are powered from a standard 3.0 VDC supply, and, of course, they're available in industrial temperatures from -40 °C to +85 °C. These fast, reliable nvSRAMs are available in svelte 44-pin/54-pin (x8/x16) TSOPII or 48-pin BGA packages. Future versions of the family will be denser and might have security features such as scrubbing.

**Simtek Corporation**  
[www.simtek.com](http://www.simtek.com)  
RSC# 36499





## USB ANALOG INPUT MODULES

**ACCES I/O Products, Inc.**  
**Website:** [www.accesio.com](http://www.accesio.com)  
**Model:** USB-AI Series

**RSC No:** 36141



A series of USB analog input modules • Units can sample inputs up to 500 KHz for the board's 16 single-ended or 8 differential input channels • High-speed USB 2.0 device with up to 500 KHz sampling rate • All functions fully software configurable • 16-bit and 12-bit models with 16 single-ended or 8 differential inputs • Eight input ranges, unipolar or bipolar • Auto calibration

and real-time hardware calibration and oversampling for accurate data • Unique channel-by-channel programmable gain feature • Data buffer for A/D • Synchronous, asynchronous, and timed trigger modes • 16 high-current digital I/O lines • 16-bit counter/timer for event counting or frequency generation • USB/104 form factor for OEM embedded applications

## XMC/PMC MODULE

**AdvancedIO Systems**  
**Website:** [www.advancedio.com](http://www.advancedio.com)  
**Model:** V1021

**RSC No:** 36417

An intelligent XMC/PMC module that provides 10 GbE connectivity for demanding real-time applications • Enables use of 10 GbE as the high-performance fabric • VITA 42.3 XMC/PMC module • Xilinx Virtex-II Pro FPGA • Standard 10 GbE SFP+ optical interface • PCIe x4 and PCI-X 133 MHz interfaces • udpXG protocol offload or streamXG direct data streaming (optional)



## MULTIFUNCTION I/O MODULE

**Advantech eAutomation Group**  
**Website:** [www.eAutomationPro.com](http://www.eAutomationPro.com)  
**Model:** USB-4716

**RSC No:** 32994



A 200 KSps 16-bit multifunction USB 2.0 I/O module • Designed for industrial environments • Includes a USB connector with screw fasteners and a DIN-rail mounting kit • Draws power directly from the USB port and therefore does not

require an external power supply • Relying on USB's plug-and-play features, device requires less than five minutes of user setup before acquiring data • I/O includes 16x single-ended (8x differential), 16-bit 200 KSps analog inputs, 2x 16-bit analog outputs, 8x digit inputs, 8x digital outputs, and 1x 32-bit event counter • WaveScan trending and data logging software program • Windows APIs, Active X controls, and LabVIEW drivers • Provides programmers an easy interface for application development like data logging, test stand data acquisition, waveform analysis, and mobile measurement applications

## INDUSTRIAL EMBEDDED COMPUTER

**Artila Electronics Co., Ltd.**  
**Website:** [www.artila.com](http://www.artila.com)  
**Model:** Matrix-514

**RSC No:** 36119



An ARM9-based Linux ready industrial embedded computer • Four Ethernet ports, four RS-232/422/485 ports, USB host, and SD socket • Low-power ARM9 RISC CPU • 16 MB flash and 64 MB SDRAM • Pre-installed Linux 2.6 OS • Four 10/100 Mbps Ethernet ports • GNU C/C++ tool chain included

## VME/VXS BOARD

**BittWare, Inc.**  
**Website:** [www.bittware.com](http://www.bittware.com)  
**Model:** GT-6U-VME

**RSC No:** 36356

Ruggedizable hybrid signal processing 6U VME/VXS (VITA 41) board • Two Altera Stratix II GX FPGAs • Two clusters of two ADSP-TS201S TigerSHARC DSPs 57.5 GOPS 16-bit fixed point, 14.4 GFLOPS floating-point processing power • Two link ports routed to ATLANTIS FPGA • Two link ports routed for interprocessor communications, 24 Mb of on-chip RAM per DSP • Four link ports at up to 1 GB each routed from the onboard DSPs • 32 LVDS pairs (64 pins), 12 channels of high-speed SERDES transceivers • Up to 1 GB onboard DDR2 SDRAM • BittWare's FINE PCI bridge chip, 32-bit/66 MHz PCI, GbE • 128 MB flash memory for booting DSPs and FPGAs



## DEVELOPMENT BOARD

**Bluewater Systems**  
**Website:** [www.bluewatersys.com](http://www.bluewatersys.com)  
**Model:** Rig 200

**RSC No:** 36096



An off-the-shelf development board for embedded electronic development and rapid prototyping • Starting point for developing a vast array of electronic products ranging from GPS location units, to mobile phones, PDAs, and specialized embedded electronic devices • Provides for a wide range of ARM9 processor cores and supports a range of add-on functionality including

GPS location, GPRS data transfer, Wi-Fi, Bluetooth, and a three mega pixel camera • Using the Snapper System modules to provide the core processor function, developers can choose a Marvel PXA270, PXA255, or Cirrus Logic EP9315 CPU • Ethernet, USB, Serial

*Continued on page 56*



## Military commanders must always feel secure with their communications systems... **DO YOU?**

### Our expertise

Aculab has become a center of excellence for many agencies delivering network centric communications solutions for deployment in tactical and infrastructure networks throughout the defense technology sector. Our advanced enabling technologies and responsive commitment to bespoke engineering development provide essential interoperability and highly effective, redundant, reliable and scalable components that enable our partners to deliver secure applications for joint land, sea and air operations.

### Our difference

When you invest in a partnership with Aculab, you get more than the highest quality media processing technology that has been optimized for your sector. Whether it's for military applications, contact centers, conferencing or fax technologies, you get the reassurance of a stable and proactive partner, offering continued support and enabling you to adapt to change before it happens.

### Our technology brings you

- Reduced operating costs
- Faster time to market
- Improved margins and ROI

### For further information

[www.aculab.com/military](http://www.aculab.com/military)  
[info@aculab.com](mailto:info@aculab.com)

+1 781 433 6000

+44 (0) 1908 273802



The technology **driving** your success





# New Products

Continued from page 54


## PCI EXPRESS/PCI-X BACKPLANE

### Core Systems

Website: [www.coresystemsusa.com](http://www.coresystemsusa.com)

Model: SEBP16-6/8/2

RSC No: 33025



A 16-slot PCI Express/PCI-X graphics backplane • Six x8 PCI Express slots, eight x4 PCI Express slots, and two PCI-X slots • Backplane architecture uses four 32-lane PCI Express switches to support a nonblocking switch fabric between the 14 PCI Express slots and the host system • Nonblocking

switch fabric supports complex peer-to-peer and host-to-peer data flows • One PICMG 1.3 SBC slot

## AdvancedMC CARD


### Embedded Planet

Website: [www.embeddedplanet.com](http://www.embeddedplanet.com)

Model: EP440xS

RSC No: 33082

A full-sized AdvancedMC card • EP440xS and its companion I/O board create a compact, cost-effective, and powerful platform for developing high-performance networked control platforms • 440EPx PowerPC processor at up to 667 MHz providing 1,334 DMIPS



and including an integrated five-stage floating-point unit • Up to 256 MB of DDRII RAM with ECC support for a high-reliability and high-bandwidth memory interface • Up to 128 MB of NOR flash on the main board and up to 512 MB of NAND flash on the I/O board • One RS-232 serial port on the main board and 12 RS-232 serial ports, one USB host port, and one USB device port on the I/O board • Two GbE ports available on the I/O board for network-enabled applications • Supported by U-Boot bootloader and Linux BSP with VxWorks and INTEGRITY BSPs available upon request

## RUGGED COMPACT NOTEBOOK

### Industrial Computing

Website: [www.industcomputing.com](http://www.industcomputing.com)

Model: Convertible Notebook

RSC No: 36415



A military grade fully-rugged compact notebook that can be used as a laptop or, by rotating the display up to 180 degrees transforms into a tablet PC slate • Intel Yonah U2500 Core Duo ULV 1.2 GHz processor • 533 MHz FSB • 2 MB L2 cache • Display convertible for laptop or tablet use • 10.4" screen or 12.1" wide screen • MIL-STD 810F and IP54 compliance • Shock-mounted removable HDD • Integrated GPS and wireless access capable • Water-

proof reversible camera • Sunlight-readable LCD screen • Supports Genuine Windows XP Professional and Genuine Windows XP Tablet PC Edition 2005

## MULTI-PLATFORM CHASSIS

### ISR

Website: [www.spraycool.com](http://www.spraycool.com)

Model: MPE Enclosure

RSC No: 36347



A Multi-Platform Enclosure (MPE) chassis employing patented two-phase cooling technology (SprayCool) for use in mission-critical military applications • Scalable: 4-21 slots for 6U x 160 mm VME, VPX, VXS, CompactPCI, or CompactPCIe (EXP.0) electronics • High performance: 100-500 W slot with optional MIL-STD-704 power supply, accepts both COTS

and proprietary cards, including air-cooled cards, with minimal modifications • Tested to MIL-STD-810, MIL-STD-461 • Maximizes size, weight, and power consumption at up to 2x improvement over traditional cooled enclosures • 1,500 ft. up to 100,000 ft. including unpressurized environments, -65 °C to +71 °C

## VXI INSTRUMENT CARD

### North Atlantic Industries

Website: [www.naii.com](http://www.naii.com)

Model: 65CS4

RSC No: 36322



A high-accuracy and high-density synchro/resolver VXI instrument card • Up to 0.005 accuracy provided for measurement and stimulus channels • Up to four synchro/resolver instrument-grade measurement channels and up to four synchro/resolver instrument-grade stimulus

channels; or up to eight synchro/resolver embedded-grade stimulus channels and up to six programmable reference supplies • All functions are independent, user programmable for either synchro or resolver format, and may be formatted for single-speed or multi-speed applications • Incorporates an internal wrap-around self-test function that does not require external hardware or software

## PowerPC MULTIPROCESSING ENGINE


### VMETRO

Website: [www.vmetro.com](http://www.vmetro.com)

Model: MPE730 VPX-REDI

RSC No: 36435


6U VPX-REDI quad-core PowerPC multiprocessing engine with Serial RapidIO • Dual Freescale MPC8641D processors with support for 1.0, 1.33, and 1.5 GHz clock rate • Dual PMC/XMC mezzanine sites with support for PCIe x8, PCI/PCI-X, and Serial RapidIO x4 • Up to 1 GB of DDR2 memory per core, or 4 GB of memory per board • Full Serial RapidIO fabric connectivity • Xilinx Virtex-5 FPGA system control node for chassis management



## PMC I/O MODULE

**Acromag, Inc.**  
**Website:** [www.acromag.com](http://www.acromag.com)  
**Model:** PMC-VLX

**RSC No:** 36088



PMC I/O modules that provide a Virtex-5 FPGA for fast processing of custom logic routines • Supported by large banks of high-speed memory, a high-throughput PCI-X interface, and plug-in I/O • Customizable FPGA (Xilinx Virtex-5) with up to 110 K logic cells and 64 DSP48E slices • Supports both front and rear I/O • PCI-X bus 133 MHz 64-bit

interface • 64 I/O or 32 LVDS lines with direct connection to FPGA via rear (J4) connector • Plug-in I/O modules available for front mezzanine • FPGA code loads from PCI bus or flash memory • Two 256 Kb x 32-bit dual-ported SRAM buffers • Two 32 Mb x 16-bit DDR DRAM banks • Other memory options available (contact factory) • Supports dual DMA channel data transfer to CPU/bus • Supports 3.3 V signaling

## RUGGED ENCLOSURE – 2U

**Hybricon Corp.**  
**Website:** [www.hybricon.com](http://www.hybricon.com)  
**Model:** 2U Rugged Enclosure

**RSC No:** 36386

Rugged 2U enclosure is part of an Undersea Warfare Combat System, providing surface warships with a seamlessly integrated undersea warfare detection, localization, classification, and targeting capability • System presents an integrated picture of the tactical situation by receiving, combining, and processing active and passive sensor data from the hull-mounted array, towed array, and sonobuoys • 19" rack-mount • 6U horizontal card cage • Four-slot VME64x backplane • Front-to-rear airflow • 291 W power supply • System monitoring • Air filtering • Applications: Version 1: Surface Ship; Version 2: Undersea (MIL-STD-461)



## INDUSTRIAL COMMUNICATIONS COMPUTER

**Korenix Technology**  
**Website:** [www.korenix.com](http://www.korenix.com)  
**Model:** JetBox 8100  
**RSC No:** 36404

A 500 MHz AMD Geode LX800-based embedded industrial communications computer • System memory 256 MB SDRAM default (512 MB optional) • 64 MB VGA small volume • I/O: Ethernet, two COM, two USB, VGA, audio, PS2 (KB/mouse), 2.5" HDD IDE slot • Ready to use • Embedded Linux/Win CE/XPe OS ready • System management utility • Application development environment • MODBUS, MODBUS/TCIP



## A/D CONVERTERS

**Murata Power Solutions**  
**Website:** [www.murata-ps.com](http://www.murata-ps.com)  
**Model:** ADSD-1420

**RSC No:** 36115

Dual 14-bit A/D converters provide 20 MSps performance for commercial, industrial, and military applications • 14-bit resolution; 20 MSps sampling rate • Functionally complete;  $\pm 2.5$  V input range • No missing codes over full temperature range • Edge-triggered •  $\pm 5$  V supplies, 1.6 W • 75 dB SNR, -80 dB THD • Ideal for both time- and frequency-domain applications



## SOFTWARE RADIO TRANSCEIVER

**Pentek, Inc.**  
**Website:** [www.pentek.com](http://www.pentek.com)  
**Model:** Model 7141

**RSC No:** 35608



Dual-channel software radio transceiver with FPGA • PMC/XMC suitable for connection to HF or IF ports in a communications system • VITA 42.0 XMC compatible with switched fabric interfaces • Two 125 MHz 14-bit A/Ds and two 500 MHz 16-bit D/As • Four digital downconverters and one digital upconverter • Xilinx Virtex-II Pro


FPGA • 512 MB of DDR SDRAM • Dual timing buses for independent input and output clock rates • LVDS clock/sync bus for multiple module synchronization • Optional factory-installed IP cores available • Ruggedized and conduction-cooled versions available

## SINE/PULSE GENERATOR CARD

**Precision Analog Systems (PAS)**  
**Website:** [www.precisionanalog.com](http://www.precisionanalog.com)  
**Model:** PAS 9883/GEN

**RSC No:** 35798


An eight-channel, synchronized sine/pulse generator card • Provides eight-function generator channels on a 6U VMEbus card • VME interface A32, A24, A16; D16 slave, no interrupts • All channels can be synchronized with constant phase offset • Phase offset registers: two 12-bit registers per channel • All DDS chips are calibrated with a precision onboard voltage reference • VME ID PAS 9883/GEN A A0, programmed in FPGA • DAC power supply – onboard  $\pm 15$  V DC-to-DC converter • Operating temperature range 0 to +60 °C



## For more information

Enter the product's RSC# at  
[www.mil-embedded.com/rsc](http://www.mil-embedded.com/rsc)





By Chris A. Ciufo, Editor

## Calling all COTS

*Everyone wants COTS in the military, but with upside comes risks*



In this magazine we focus on “COTS and technology for the entire military life cycle,” and it says so right on the cover. While I prefer to emphasize the “life cycle” part of our mission statement with stories such as our new guest column series *Legacy Software Migration*, Commercial Off-The-Shelf plays a huge part of what we write about. If recent conferences and trade shows are any indication, more companies than ever want to sell their COTS wares to the military. But as always, there’s good news and bad news.

**First the bad news.** Congress and the GAO are waking up to a squad-sized list of military programs that have gone way over budget. Even at a time when the DoD has doubled planned expenditures on new weapons from \$790 billion in 2000 to a whopping \$1.6 trillion in 2007 (Source: *Aviation Week & Space Technology*, April 7, 2008), the GAO recently reported widespread skepticism at successful program outcomes. Sixty-three percent of the programs reviewed had cost increases. Moreover, parsing the data further, many of these problems can be attributed to design changes, some of which certainly can be blamed on COTS. Part of the problem is how out-of-phase the commercial market technology cycle is with the military’s *long life cycle* of intolerably slow steps from concept to fielding.

Let’s look at the WIN-T program as an example. The Warfighter Information Network-Tactical is a key component of the Army’s Future Combat Systems – arguably the United States’ biggest defense program *ever*. The ORD for WIN-T was written in June 1999 to provide a tactical telecommunications infrastructure as part of the GIG, “Joint Vision 2010” and the “Army 2010 and Beyond” vision. But as recently as March 2008, the DoD published a Selected Acquisition Report (SAR) on WIN-T, which is mandated when a program experiences at least a 15 percent cost increase or schedule delays of at least six months.

General Dynamics, Lockheed Martin, and other teammates were awarded the development contract for WIN-T in September 2004, which was estimated to cost \$7 billion through 2018. Initial Operational Capability (IOC) was slated for 2008. Since 2004, the program has been sliced into several pieces, including Increment 1 (formerly Joint Network Node) and Increment 2 (Initial Networking on the Move), both of which recently fell under the SAR’s black mark. In the meantime, also in March 2008 GD won a contract for 299 Satellite Transportable Terminals (STTs) totaling \$108 million. And what about the whole WIN-T program? Pundits now estimate the program will cost about \$17 billion when complete – more than double the original estimate.

To be fair, the military’s needs have slid to the right, too. Do some technical digging into this fairly complex program and two things become apparent: First, since WIN-T is fundamentally a telecom network, it’s loaded with COTS equipment such as CompactPCI (now AdvancedTCA?), WindowsXP (now Vista?), and Intel Pentium processors (now Core 2 Duo?). Secondly, as the Internet has changed, so have the Army’s expectations for connectivity and content – from voice and primitive data, to Blue Force Tracking, moving GPS-enabled maps, and live video feeds – causing a near constant escalation in the composition of COTS equipment. Add all of this up and the conclusion is that COTS is both a blessing and a curse.

**But everyone wants in.** Our war fighters can’t live without COTS technology, and the vendor community is anxious to oblige. Intel, for example, never warmed up to GEIA’s AQEC quasi-MIL-SPEC way of building semiconductors, but the company recently announced a life-cycle extension to *seven years* for devices on the company’s “Embedded Roadmap.” Military, medical, and industrial markets are the prime recipients of Intel’s largesse.

Additionally, at last week’s Embedded Systems Conference in San Jose, California, attendees couldn’t walk 10 feet without tripping over another vendor’s hardware or software product that was “ruggedized,” targeting “extended temperatures,” or directly marketed to “military applications.” Pictures of AH-1 helicopters, F-22 Raptors, and Aegis destroyers on vendor booth signs were astounding, considering that ESC isn’t a military show.

Germany’s Kontron Modular Computers, for example, recently completed the acquisition of France’s \$40 million Thales Computers ([www.mil-embedded.com/news/db/?10811](http://www.mil-embedded.com/news/db/?10811)). Although both companies are international, Kontron marketing executives told me that Thales was targeted *specifically to get Kontron back into the U.S. VME defense business*. Similarly, Taiwan-based ADLINK surprised the market last month by acquiring PC/104 inventor Ampro Computers for \$20 million. In a conversation with ADLINK’s U.S. general manager, I learned that the acquisition was motivated primarily to broaden ADLINK *into* the U.S. defense market. Ampro’s 70 employees give ADLINK a solid footprint on U.S. soil and the possibility to target the company’s small form factor boards and systems into hard-core DoD programs.

And there are more examples. Sun started the JSR302 safety-critical Java effort partly to offer a migration path for the military away from Ada. HP printer RTOS supplier Express Logic is finding a tidy little business putting ThreadX into space probes. Will we see an end to this COTS proliferation? I doubt it, and future defense requirements will continue to *call on COTS*.

Chris A. Ciufo  
Group Editorial Director  
[cciufo@opensystems-publishing.com](mailto:cciufo@opensystems-publishing.com)



# Here comes the future, all engines full throttle.

**VPX is revolutionizing the world of rugged COTS systems.**

VPX systems are transforming rugged mil/aero applications, as our MAGIC1 clearly illustrates. It combines desktop gaming silicon from NVIDIA® with the latest Intel® Core™2 Duo processor and x16 PCI Express® technology. The result is a complete, integrated graphics subsystem designed to deliver blazing performance in the harshest environments.

With support for industry standard graphical environments such as VAPS, GL Studio and iData, the MAGIC1 is the industry's highest performance,

ready-to-run, rugged display processor. That makes it ideal for sophisticated, real time applications that demand the utmost in 2D and 3D graphics and video processing. Think embedded training. Think simulation. Think mission rehearsal. Think digital mapping. Think vehicle displays.


For more information on VPX and the MAGIC1, please visit our web site at [www.gefanucembedded.com/vpx](http://www.gefanucembedded.com/vpx) and be prepared for a mind-changing experience.



**MAGIC1**  
3U VPX system with leading  
edge graphics performance







## TOUGH DSP SOLUTIONS

## FOR TOUGH ENVIRONMENTS

Today's critical missions demand tough solutions. Solutions built to survive the harshest environments. Today's warfighters rely on rugged deployed sensor platforms with sophisticated signal and image processing for vital real-time information. Now, FPGA-based embedded computing takes sensor platform capabilities to a whole new level, delivering dense processing on low-power platforms.

Curtiss-Wright offers a range of rugged FPGA-based processing platforms. They include our full-featured VPX-based CHAMP-FX2 heterogeneous processing module and the new compact XMC-E2201 signal acquisition mezzanine board. Each features comprehensive board support packages, software drivers, and advanced FPGA development kits to speed and ease application development and system integration.

If your system requires high performance sensor data processing, or your challenge is to reduce space, weight, and power consumption, Curtiss-Wright has the rugged, embedded FPGA platform for you. You work in a tough world. We thrive on the tough solutions.

**CURTISS  
WRIGHT** Controls  
Embedded Computing

[www.cwcembedded.com](http://www.cwcembedded.com)

**CHAMP-FX2**



**XMC-E2201**



**XMC-442**



Curtiss-Wright's FPGA processing boards for tough DSP, image processing, or signal acquisition include the CHAMP-FX2 and the XMC-442 FPGA processing modules, and the XMC-E2201 signal acquisition module.

RUGGED FPGA PROCESSING... **ABOVE & BEYOND**



Programmable Logic:

# FPGAs & Reconfigurable Computing



## Interviews

- 4 John East, President and CEO, Actel
- 8 Jeff Kodosky, Co-Founder, National Instruments
- 9 Jeff Milrod, President and CEO, BittWare
- 12 Chris Fanning and David Lee Rutledge, VPs, Lattice

## Features

- 16 Projecting images in radar and medical applications
- 19 The nuances of SIGINT in the real world
- 22 Object Array enters the fray
- 25 The key to effective interface design



# ***Annapolis Micro Systems***

## ***The FPGA Systems Performance Leader***

# **WILDSTAR 5 for IBM Blade**

## **The Perfect Blend of Processors and FPGAs**

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



**Made in the USA**

### **Ultimate Modularity**

**From 2 to 8 Virtex 5 FPGA/Memory Modules**

**Input / Output Modules Include:**

**Quad 130 MSps thru Quad 500 MSps A/D**

**1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D**

**Dual 1.5 GSps thru 4.0 GSps D/A**

**Infiniband, 10 G Ethernet, FC4, SFPDP**

### ***Direct Seamless Connections with no Data Reduction***

***Between External Sensors and FPGAs***

***Between FPGAs and Processors over IB or 10GE Backplane***

***Between FPGAs and Standard Output Modules***

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



For Systems  
that Demand  
the Most



## XG-5000K

Our most advanced  
FPGA Based  
PC/104 Plus  
Circuit Board

### Features

- Xilinx Spartan 3, 5 Million Gate FPGA
- 264 User Configurable I/O
- 256 MByte of SRAM
- 32 MByte of Flash
- 6 MByte of Serial Page DataFlash
- CompactFlash Type 1 Connector
- Secondary Xilinx Spartan 3, 400,000 Gate FPGA
- 10/100Base-T Ethernet
- 2 RS232 Ports
- 25MHz FPGA Master Clock Source
- Secondary 0 - 25MHz user programmable DDS FPGA Master Clock Source
- Incorporates Xilinx's design revisioning technology and can retain onboard as many as 16 partial or 4 complete FPGA design BIT files
- Can be powered from the PC/104 connector or an external DC power supply
- PC/104 Plus Form Factor;
- ISA and PCI 33MHz Bus Interfaces
- Extended temperature -40C to +85C



Jacyl Technology is a Team Member  
under Lear Siegler Services, Inc for the  
US Governments CECOM Rapid  
Response Contract



Advancing Today's Technology into Tomorrow  
[www.jacyltechnology.com](http://www.jacyltechnology.com)



# Power consumption as the competitive advantage



**Q&A with John East,**  
*President and CEO, Actel Corporation*



*Our exclusive interview sheds light on how Actel's strategy and technology are reshaping programmable logic for the latest power-sensitive applications.*

**PL: Actel's strategy has Design for Energy Efficiency written all over it. How did you arrive at that as a core strategy for the company, and how is it changing your business?**

**JE:** I guess the real question is why a company *wouldn't* arrive at energy efficiency as a core strategy. The industry is ripe with discussion of "green design" and industry leaders are called upon every day to act as socially responsible citizens.

Power matters. With the world's increasing use of electronic devices, Actel has the opportunity to play a major role in resolving the world's global warming problems with further technological change.

We're changing FPGAs from "power hogs" into "power savers." We've designed non-volatile, reprogrammable flash-based devices offering 200 times less static power and more than 10 times the battery life than competitive SRAM-based programmable logic solutions.

We are in the business of enabling designers to save power – from low-power FPGAs and power-efficient programmable system chips to power-optimized tools and power-smart IP.

**PL: A good example is the IGLOO. What makes the IGLOO architecture a low power solution?**

**JE:** There are three main points of interest in IGLOO.

The differentiating factor for the majority of our product line is the use of a true flash-based architecture. True nonvolatile flash-based FPGAs contain a nonvolatile FPGA array, and do not suffer from the leakage current issues associated with SRAM-based solutions.

Secondly, we felt it paramount that we select a foundry partner with a proven low-power flash process technology. We chose UMC's leading-edge 0.13-micron low-power and e-Flash processes. Combined with IGLOO's power mode options and Flash\*Freeze technology, UMC's processes enable power consumption as low as 5 microwatts.

Finally, the IGLOO family has several power modes to optimize power consumption – the Flash\*Freeze mode, a low-power active mode, and a sleep mode. While in Flash\*Freeze mode, the device conserves power while maintaining FPGA content. Device I/Os are tri-stated and SRAM and register content is maintained, but not clocked. Further, the Flash\*Freeze pin enables designers to quickly and easily enter or exit the Flash\*Freeze mode within 1 microsecond. Alternatively, the low-power active mode allows the IGLOO device to directly control when the system goes into low-power mode. While in the low-power active mode, the FPGA core, clocks, and all I/O are functional.

**PL: What have you done to aid designers in power consumption analysis?**

**JE:** The Libero 8.1 IDE has specific features to further optimize FPGA designs for low power.

The Libero IDE enables creation of realistic power profiles based on design parameters or metrics. An example is the percent of time the FPGA will be in a combination of custom or functional modes, such as Active, Standby, or Flash\*Freeze. The tool then displays a realistic and accurate report of power consumption based on the true power profile of the target FPGA.

The design analysis data is then utilized to engage in power-driven layout. Use of this tool enables users to quickly realize

# LOWEST TOTAL COST... PERIOD.



## GET UP TO 50% LOWER COST

Xilinx Spartan™-3 Generation devices are the only low-cost FPGA family to cut your total cost by half... and we can prove it! We offer the industry's widest range of devices, most comprehensive IP library (8x the nearest competitor!), a huge selection of development boards and we've minimized the need for external components.

Visit us at [www.xilinx.com/spartan](http://www.xilinx.com/spartan) to download our free ISE™ WebPACK™ design tools and start your Spartan-3 design today. You too can experience the lowest total cost. Period.

 **XILINX®**  
[www.xilinx.com/spartan](http://www.xilinx.com/spartan)

**The World's Most Widely Adopted Low-Cost FPGAs**

©2004 Xilinx, Inc. All rights reserved. XSPD, the Xilinx logo and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.



dynamic power savings through the reduction of the capacitive loading of the nets. While average IGLOO power consumption is reduced by 13 percent, some designs can realize a 30 percent reduction in power consumption.

**PL: Where do you see processor cores in FPGAs headed?**

**JE:** We see real value there and have several cores – CoreABC, Core8051, Core8051s, and CoreMP7, as well as the FPGA-optimized ARM Cortex-M1 and the Leon3 SPARC processor.

These offerings bring the flexibility, fast time to market, and low implementation cost of FPGAs to users who can't justify the expense of ASICs for their application. We'll continue to explore processor core technologies with more choices, better tools, and improved power optimization.

**PL: I see you're betting on MicroTCA specifically – why?**

**JE:** We're really interested in system management. At the high end, there are standards-based and proprietary solutions for telecommunications such as AdvancedTCA and MicroTCA. At the low end embedded applications need low cost, less complex solutions. We're doing both.

Our mixed-signal Fusion Programmable System Chip (PSC) can perform system management functions, including power and thermal management, data logging, and system diagnostics, within high-end and low-end electronic systems. And it's all in a single chip, with a cost and space savings of 50 percent or greater relative to current implementations while also improving reliability.

**PL: How can designers get an advantage in their next FPGA-based design?**

**JE:** Nearly every vendor claims low-power leadership. Ignore the claims and check the datasheets – static power, dynamic power, inrush, configuration power, and low-power modes.

Comparing one-million gate devices, for example, SRAM-based FPGAs marketed as the "lowest power FPGAs" state static power consumption between 40 milliwatts and 150 milliwatts, which is 800 to 3000 times higher than Actel's 0.05 milliwatt flash-based IGLOO device. Similarly, when comparing "zero power" CPLD solutions against IGLOO, these devices consume 10 times more static power.

Then look at tools to optimize power. How does the development environment like Libero help the designer? What power-smart IP, like FPGA-optimized 32-bit ARM processors, is available? Tools are just as important as architecture.

**PL: What's next for the FPGA industry?**

**JE:** Today, FPGAs are capable of handling so much more within a design – whether it be the interfacing and control within a smart phone or in a system-critical automotive or medical application. FPGAs are expanding into so many different applications that can leverage their flexibility, low power, and low cost.

Looking ahead, I believe a continued focus on power reduction will drive even more applications. There continues to be tremendous demand for low-power programmable solutions at lower densities to handle the control and bridging functions in portable designs. We're now hitting the \$1 price point with these solutions, particularly for price-sensitive consumer portable applications. The technology will just continue to improve.



*John East is President and CEO of Actel Corporation, with nearly 40 years experience in the semiconductor industry. Prior to joining Actel in 1988, East held positions with AMD, Raytheon Semiconductor, and Fairchild Semiconductor. He has a bachelor's of science degree in electrical engineering and a master's degree in business administration from the University of California, Berkeley.*

**Actel**

650-318-4200 • [www.actel.com](http://www.actel.com)

**DSP-FPGA.com**

**Editorial**

Don Dingee  
Editorial Director  
[ddingee@opensystems-publishing.com](mailto:ddingee@opensystems-publishing.com)

Chris Ciuffo  
Group Editorial Director  
[cciuffo@opensystems-publishing.com](mailto:cciuffo@opensystems-publishing.com)

Robin DiPerna  
Assistant Editor

Rosemary Kristoff  
Vice President Editorial

Joann Toth  
Senior Designer

David Diomedea  
Art Director

Konrad Witte  
Senior Web Developer

Sandy Dionisio  
Graphic Coordinator

**Advertising/Business Office**

OpenSystems Publishing  
30233 Jefferson Avenue  
St. Clair Shores, MI 48082  
586-415-6500

Patrick Hopper  
Vice President Marketing & Sales  
[phopper@opensystems-publishing.com](mailto:phopper@opensystems-publishing.com)

Karen Layman  
Business Manager

**Sales Group**

Dennis Doyle  
[didoyle@opensystems-publishing.com](mailto:didoyle@opensystems-publishing.com)

Tom Varcie  
[tvarcie@opensystems-publishing.com](mailto:tvarcie@opensystems-publishing.com)

Doug Cordier  
[dcordier@opensystems-publishing.com](mailto:dcordier@opensystems-publishing.com)

Andrea Stabile  
[astabile@opensystems-publishing.com](mailto:astabile@opensystems-publishing.com)

Christine Long  
[clong@opensystems-publishing.com](mailto:clong@opensystems-publishing.com)

Ernest Godsey  
[egodsey@opensystems-publishing.com](mailto:egodsey@opensystems-publishing.com)

Barbara Quinlan  
[bquinlan@opensystems-publishing.com](mailto:bquinlan@opensystems-publishing.com)

Denis Seger  
[dseger@opensystems-publishing.com](mailto:dseger@opensystems-publishing.com)

Sydele Starr  
[sydele@pacbell.net](mailto:sydele@pacbell.net)

Ron Taylor  
[rtaylor@opensystems-publishing.com](mailto:rtaylor@opensystems-publishing.com)

**International**

Dan Aronovic  
Account Manager – Israel  
[daronovic@opensystems-publishing.com](mailto:daronovic@opensystems-publishing.com)

Sam Fan  
Account Manager – Asia  
[sfan@opensystems-publishing.com](mailto:sfan@opensystems-publishing.com)

**Reprints and PDFs**

Nan Lamade: 800-259-0470  
[dspreprints@opensystems-publishing.com](mailto:dspreprints@opensystems-publishing.com)

# Reconfigurable FPGA Processing Cards



## Ultra Performance X5 Modules

Xilinx Virtex5 FPGA logic

4MB QDR SRAM

8 Rocket IO Private Links  
2.5 Gbps each

>1GB/s  
8-lane PCI Express Host

Available I/O Solutions:

### X5-400M

400 MSPS 14-bit A/D (x2)  
500 MSPS 16-bit D/A (x2)  
1GB DDR2 DRAM

### X5-210M

210 MSPS  
14-bit A/D (x4), 512MB DDR2 DRAM



## eInstruments

User-customizable, embedded instrumentation!  
COTS solution combining two XMC modules plus intelligent  
x86-based SBC in rugged enclosure

### Capabilities

- Embedded software radio and IF processors; DSP within FPGA
- FPGA-based real-time (servo) control/analog I/O at rates to 400 MHz
- Perfect for portable or data loggers or waveform generators

### Distributed Data Acquisition

- Co-locate eInstrument and unit-under-test for optimized signal quality/minimized cabling
- Limitless expansion via multiple nodes

### Single Board Computer (SBC)

- Scalable CPU up to Core 2 Duo w/ 4 GB memory
- USB2.0 x6, Gigabit Ethernet, SATA x4, VGA, AC 97 audio
- GPS timebase support
- Cabled PCI Express expansion to other eInstruments
- Windows or Linux from USB FLASH (diskless) or SATA HDD

### Rugged & Portable

- 12V DC-only operation via Battery or AC Adapter
- Conductive cooling plus optional fan
- 240 x 170 x 50 mm
- Optional aluminum instrument enclosure



## Low Cost X3 Modules

Xilinx Spartan 3A DSP 1.4M

Two 2MB SRAMs, 48-bit DIO

PCI Express with >200 MB/s data rates

Available I/O Solutions:

X3-25M: 16-bit, 130MHz, A/D (x4), D/A (x4)

X3-A4D4: 16-bit, 4 MHz A/D+ D/A (x4)

X3-DIO: 64-bit, 66 MHz LVDS

X3-SD: 216 kHz, 16-bit, A/D (x16)

X3-SDF: 20 MHz, 24-bit, A/D (x4)

X3-Servo: 16-bit, 250 KSPS, A/D (x12)  
16-bit, 1 KSPS D/A (x12)  
(ultra low-latency & glitch energy)

Using Innovative Integration's FrameWork Logic VHDL source-code or Simulink/MATLAB board support package you can readily customize FPGA functionality to include real-time processing such as independent FIR and IIR filters on each channel, real-time FFT processing, ultra-fast feedback and control loops and much more.

For more information visit [www.innovative-dsp.com/e](http://www.innovative-dsp.com/e)



FrameWork Logic



[www.innovative-dsp.com](http://www.innovative-dsp.com) • [sales@innovative-dsp.com](mailto:sales@innovative-dsp.com) • 805-578-4260

**Innovative  
Integration**  
... real time solutions!



# Stretching the boundaries for FPGAs and people



**Q&A with Jeff Kodosky,**  
*Co-Founder and Technology Fellow,*  
*National Instruments*

**J**eff K, as he's known in the LabVIEW community, has a unique vision for how people should program and who should program. He shares some exclusive thoughts with us on where he sees NI's technology today and in the future.

**PL:** We've heard from Dr. Truchard about graphical system design. How else has the FPGA capability of NI CompactRIO and LabVIEW changed the way people think about designs?

**JK:** The combination of an intuitive graphical programming language and a heterogeneous hardware platform (MPU + FPGA) has really changed not only how people think about design but also who is thinking about design.

Hardware and software design engineers see this tightly integrated, embedded platform as a solution to their time-to-market pressures. Regardless of their deployment platforms, these experts use NI LabVIEW software and CompactRIO hardware to prototype and iterate on their designs and algorithms at a fraction of the time even imagined in the past.

But even more exciting to me is how this marriage of graphical programming and the latest programmable logic technology is changing the "who" of the design world. Now domain experts – the physicists, robotics engineers, mechatronics and mechanical engineers, and scientists – can actually implement their innovative ideas in hardware themselves, instead of hiring an embedded design expert or giving up seeing their concepts come to fruition.

**PL:** Where are the boundaries today, and where would you like them to be soon, in choosing to implement on an FPGA or a DSP?

**JK:** The characteristics of an application typically suggest which alternative might be more appropriate. If the application is computation-bound and involves a great deal of complicated mathematical computations, a DSP may be the best implementation target. An FPGA works better if the application is highly parallel or pipelined, or logic-bound with many bit-level operations – for example, when implementing digital protocols. An FPGA also works better if the application uses custom I/O requiring high-speed timing and triggering logic or special "glue" logic. Finally, an FPGA may be the best choice in some safety-critical situations because the safety-critical portion can run in parallel without interference from anything else; there is no operating system, device driver, critical section, or interrupt to delay or interfere with the execution of the safety-critical actions.

In the long term, I would like to see a convergence between FPGA and DSP targets where highly optimized ALU components are generously sprinkled around in an FPGA

connection and logic fabric. Such a hybrid target would be ideal for complex embedded applications that combine lots of parallel computation with high-speed timing and triggering.

**PL:** At NIWeek 2007, OEM board versions of CompactRIO were shown – give us an example of how that is impacting fielded applications.

**JK:** The ability to move down what we call the *RIO deployment curve* is very exciting. Engineers and scientists can use powerful platforms like PCs, PXI, and plug-in FPGA devices for rapid prototyping. Then they can easily move to smaller, rugged systems like CompactRIO, and even integrated systems combining a real-time processor and reconfigurable FPGA within the same chassis.

Sanarus, a medical device start-up company, has developed plans for a potentially revolutionary product that could change the way doctors treat benign tumors. With this device, based on LabVIEW and CompactRIO, doctors can eliminate tumors by freezing and killing them in an outpatient procedure, a dramatic change from in-patient surgery or the "wait and see" approach used previously.

The Visica2 Treatment System is an instrument for use in a doctor's office or clinic. The procedure is performed with local anesthesia and uses a real-time, ultrasound-guided probe that is virtually painless. The treatment, which lasts 10 to 20 minutes, freezes and destroys targeted tissue through an incision so small that it does not require stitches.

*National Instruments, continued on page 10*

# Signal processing in FPGAs goes mainstream



**Q&A with Jeffrey Milrod,**  
*President and CEO, BittWare*

**J**eff gives us further insight in this exclusive interview as to how signal processing using FPGAs has changed, and how customers are benefitting from the improvements.

**PL:** What's different in 2008, from a couple years ago, in how people design a signal processing system with FPGAs?

**JM:** A few years ago, FPGAs were used mostly, but not exclusively, for communications, interface, and glue-type logic. Only the brave, lunatic, or desperate few implemented processing algorithms in the FPGA itself. Even then, these were mostly very simple, straight-forward, and repetitive types of algorithms.

Today, things are very different and the pendulum has swung the other way. It now seems that everyone wants to use FPGAs as a signal processing resource, but clearly not everyone understands the implications of that. Implementation is still quite challenging and not for the faint of heart. While FPGA vendors and third party tools suppliers have made great strides, there is still a gap between the high-level algorithm implementation and the HW.

**PL:** I see FPGAs, DSPs, and even GPPs in your diagrams. Where are the boundaries in partitioning a system?

**JM:** There really are no strict boundaries other than those imposed by hard performance requirements. However, partitioning can have a huge impact on development time and effort.

Traditional DSPs are still much easier for implementing complex algorithms and algorithms that are likely to be frequently modified, since development can be done in higher level languages like C. This also facilitates code reuse – many users already have a great deal of code working on DSPs, with no compelling reason to port to FPGAs.

If code reuse is not an issue, and the algorithm is fairly well-defined, standard, or straight-forward, FPGAs are very attractive since they can offer compelling performance advantages – both in terms of size and speed. Often, signal processing algorithms force the use of FPGAs because they can't be reasonably implemented in DSPs.

FPGAs provide some future-proofing – as they become bigger, faster, lower power, and easier to use, it's likely that they will dominate signal processing. FPGA development investments today will be the code reuse of tomorrow.

As for GPPs, we only use them to do command and control processing. This can greatly ease the processing and development

burden of the target DSPs and FPGAs, so that these valuable resources only do what they're best at.

**PL:** What does your FPGA design toolset look like, and why are those pieces important?

**JM:** The generic block template in the toolset consists of a function implemented with a standard data interface for sourcing and sinking, along with a standard memory mapped control interface. Many of the specific blocks provide board-level physical interfacing, but some provide data switching and routing functions, and others provide resource arbitration, DMA engines, control block, and utility functions.

These pieces are not overly exciting in and of themselves; the real value lies in the fact that they are constructed to be building blocks in a framework – BittWare's ATLANTiS – allowing users to more quickly and easily get their special algorithms and applications up and running in the FPGA on a real board.

**PL:** What makes the ATLANTiS framework special?

**JM:** In ATLANTiS, we're using standard interconnects and an orthogonal control plane, and we've done all the low-level physical interfaces and data movement structures that are standard in microprocessors or DSPs. (Figure 1.)




Figure 1

Rather than starting with a blank slate, as it were, and having to deal with all the low level interfacing, data movement, and control, our ATLANTiS modular framework enables users to focus on adding their unique value to the FPGA – much like one



would focus on writing algorithmic code and setting-up memory structures in a processor without worrying about creating a data transport layer, DRAM controller, or arbitrating for resources.

**PL: How are people succeeding at shrinking their systems with FPGAs? Please give us an example.**

**JM:** We have seen many examples of this in military, instrumentation, and communication applications where an FPGA is used to implement an algorithm that is particularly difficult to do in standard DSPs or GPPs.

One specific example of this is a high-end cytometry (cell-sorting) instrument from iCyt. Their flow cytometer uses a laser to detect and sort the cells, with optical sensors sampled at over 100 MHz. Using DSPs to detect the cells required several DSPs to perform a weighted threshold on every sample from every sensor. We consulted with them and helped implement the cell detection algorithm in a pre-processing FPGA, thereby only requiring DSPs for the cell analysis – cutting the number of DSP boards in half.

**PL: What's the next wave look like, and where can competitive advantage be gained?**

**JM:** The next generation of FPGAs that we've seen through our partnership with Altera is simply amazing. The speed, feature set, and capabilities are astounding, and the power is lower than I'd hoped. The challenge is in harnessing all that tremendous signal processing potential in a practical, timely, and cost effective way.

Competitive advantage will be gained by those who facilitate code reuse and can reduce the effort required to get real signals in and processed in the FPGA. There is a great deal of work being done to improve the algorithmic implementation process, but often the harder part is integrating the algorithms into the real world of signal processing.

**PI**

*Jeffrey Milrod is BittWare's president and CEO. He holds a bachelor's degree in Physics from the University of Maryland and an MSEE from John Hopkins University.*

**BittWare**

603-226-0404 • [www.BittWare.com](http://www.BittWare.com)

Because this device will be used in offices and clinics around the country, the machine has to be cost-effective. The lower-cost hardware options from NI, still based on the unique reconfigurable (RIO) architecture, enable Sanarus to meet its volume, cost, and technology challenges.

**PL: Also at NIWeek 2007 a vision was outlined to get LabVIEW on a chip. Where is the vision is going, and what is it going to take to get there?**

**JK:** It has always been my goal to ensure LabVIEW can reach the lowest levels of hardware and the smallest and lowest-cost programmable targets. This goal is captured nicely by the phrase "LabVIEW on a chip."

In one sense, LabVIEW FPGA already achieves this goal. If an ASIC were made from the FPGA then a LabVIEW application would really be on a chip, but we haven't seen a high enough volume application to justify making an ASIC.

A related interesting topic is whether LabVIEW on a processor could be accelerated by having some code in the run-time engine actually programmed into an FPGA sitting alongside the processor, or maybe directly in the firmware of the processor itself. This is something we continue to brainstorm, and we are beginning to talk about it informally with some processor vendors. We are also developing relationships with several universities to explore these and related ideas.

Another way to think about LabVIEW on a chip is to envision a processor chip that runs LabVIEW "natively," where software modules (virtual instruments) would be loaded and run analogously to the way a typical operating system loads and runs applications. This could be done with a typical Von Neumann architecture processor today, but it is interesting to envision what the ideal underlying hardware might be to fully realize the potential of the LabVIEW parallel language.

In my view, the ideal hardware would be a "super" FPGA containing lots of optimized embedded components sprinkled throughout the interconnect and logic fabric, such as integer and floating-point MACs and ALUs, FIFOs, blocks of memory, and so on. This super FPGA would consist of a large number of separately reconfigurable regions where the configuration memory itself would be multi-buffered so that the contents of



a region could be changed in a single clock cycle. I think LabVIEW would be an excellent tool to program such architecture, and, who knows, maybe something like this will eventually be developed.

**PL: What should designers be focused on doing differently now to get an advantage in their next FPGA design?**

**JK:** To get the most out of their next multicore or FPGA applications, designers should expose more of the parallelism by drawing more computations in parallel and pipelining computations in a loop. This is already fairly easy and safe to do in LabVIEW, and it will be increasingly important in achieving the highest performance on an FPGA when using newer, higher-speed I/O modules.

**PI**

*Jeff Kodosky cofounded National Instruments with Dr. James Truchard and William Nowlin in 1976. Known as the "father of LabVIEW," he invented the graphical programming language that defines the software. Kodosky was named an NI business and technology fellow in 2000. He received his bachelor's degree in physics from Rensselaer Polytechnic Institute in Troy, NY.*

**National Instruments**

512-683-0100 • [www.ni.com](http://www.ni.com)

## MIMO advanced development solution

### SDR development platform

## Most versatile programmable logic for telecommunications

- ✧ Complete baseband-to-RF integration
- ✧ Highly modular and scalable
- ✧ Complete board support package, including examples for fast kick-off prototyping
- ✧ Unsurpassed model-based design support package — Lyrtech products offer the best and most complete I/O and interface integration within configurable Simulink blocks.



What's more...

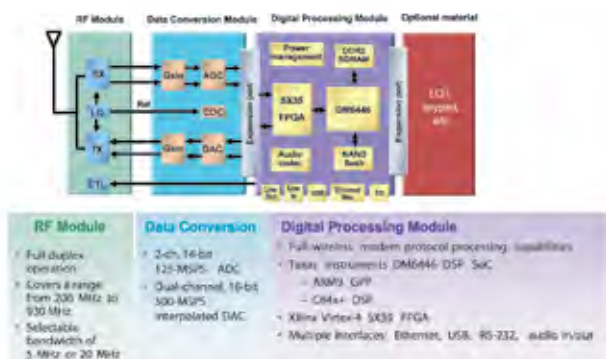
The small form factor SDR development platform is:

- Power efficient
- Capable of component-level measurement and repartitioning
- Real-time power monitoring

GSM layer 1/2 software IP on DSP-FPGA mixed architecture

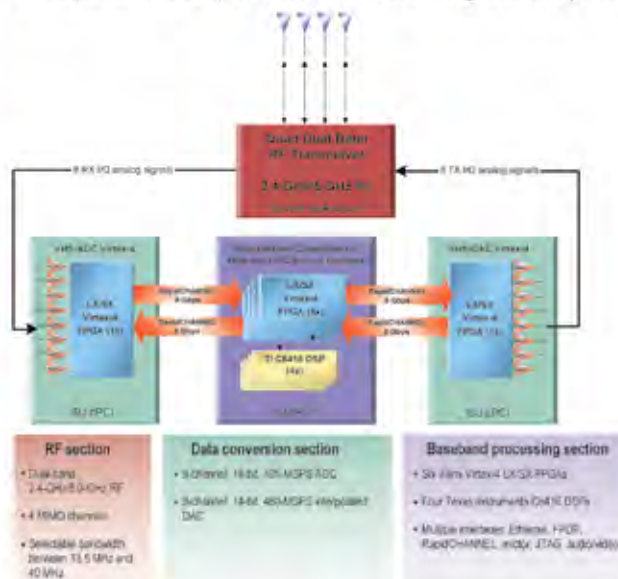
Also features an SCA version:

- All-CORBA enabled (FPGA, DSP, and GPP)
- All-SCA enabled
- Predefined device managers



The MIMO advanced development solution is:

- The **only** complete 802.11n solution on the market today
- Capable of multichannel vector signals generation and analysis
- Equipped with real-time host data exchange tools
- Capable of 1-Gbps data bandwidth throughout its system



[www.lyrtech.com](http://www.lyrtech.com)



# Driving both technology and cost/performance



**Q&A with Chris Fanning  
and David Lee Rutledge,**  
*Lattice Semiconductor Corporation*



**T**wo senior execs at Lattice give us their views on the Lattice strategy, with an in-depth look at non-volatile FPGA technology, in this exclusive interview and short article.

**PL:** Lattice is making pushes in a lot of areas. Why such diversity in the offering, and how does it tie together from a strategy standpoint?

**DLR:** When Lattice entered the FPGA market, we chose to focus on two key differentiators – Non-Volatility (NV) and Low-Cost – and set about developing technology and products in each area.

NV is the technology of choice for programmable logic, providing system-level benefits such as single-chip solutions, instant-on capabilities, and in-system programmability. We knew that we could leverage our NV expertise to deliver these benefits to FPGA users.

Lattice saw an opportunity to fill a void in the market by developing a Low-Cost, Feature-Rich FPGA and CPLD product line. Our strategy has been to develop products that work very well for most applications, allowing us to substantially reduce cost with little impact on performance.

So, while our products may seem diverse, they have been developed based on a consistent strategy to deliver value-added products, with NV and Low-Cost, Feature-Rich as our key differentiators.

**PL:** What's the biggest technology trend driving your FPGA implementations today, and how do you see that affecting what you can do in the near future?

**DLR:** Achieving higher levels of functional integration and performance while maintaining acceptable power consumption is a major challenge. Historically, the scaling of supply voltages and transistor feature sizes has provided great benefits in functional integration levels and performance, while lowering power consumption. In the future, these techniques will still provide substantially higher levels of functional density and performance, but at the expense of increased levels of static power consumption.

At 45 nm there will be a more direct trade-off between speed and power. For example, a high-density FPGA (~500 K LUT) built on 45 nm technology could easily consume over 10 W of standby power, with no clocks running, at the commercial +85 °C junction temperature limit.

Innovation in power management to optimize this trade-off is a high priority. Innovation is required concurrently across multiple disciplines, ranging from optimizing the basic process technology through the development of new “power-optimized” product architectures and also through the development of “power-aware” design tools.

Also, there is increased use of embedded SERDES channels as high-bandwidth chip-to-chip communication links. An FPGA with 40 channels of soon-to-be-available 10 G SERDES will have an incredible processing capacity of 400 Gbps. This trend will accelerate the development of radically new FPGA-based High-Performance Computing (HPC) platforms.

**PL:** On the tools side, I see some additions of synthesis and simulation capability, and I'm guessing you're expanding further. What's the latest technology you are working on, and what's the impact?

**CF:** Lattice has made a very significant investment in its design tool, ispLEVER. Two initiatives we see having great impact are physical synthesis and incremental design.

Physical synthesis should enable more optimal results in improved device performance and help accelerate the debug process. Incremental design is a collection of technologies that ultimately provides

more rapid turnaround time for achieving timing closure. Both these enable customers to meet timing requirements more easily and with fewer design iterations, even as FPGA devices become even larger and more complex.

Other Lattice-driven design tool advancements include:

- Power Calculator allows specifying parameters such as voltage, temperature, process variations, airflow, heat sink, resource utilization, activity, and frequency, and then calculates static (DC) and dynamic (AC) power consumption.
- Reveal uses a signal-centric model for embedded logic debug. Signals of interest are user-defined, and the tool adds instrumentation along with the proper connections to enable the required in-system analysis can then be performed. Users can specify complex, multi-event triggering sequences that make system-level design debug smoother and faster.

**PL: What about embedding operating systems on an FPGA core? How is this changing the way people design?**

**CF:** uClinux support expands our commitment to the open source model, and it consistently appears at the top of designer surveys we've seen as the preferred RTOS for embedded design. uClinux and the LatticeMico32 core allow designers to implement control systems in a design flow that builds on Lattice's open source, embedded solutions approach.

The adoption of embedded products has increased among FPGA designers in the last two to three years as the processors provided by FPGA vendors have dramatically improved in functionality, increasing designer productivity and lowering design risk. Embedded processors increasingly include a robust assortment of middleware such as DDR, DDR2, SDRAM memory controllers, Tri-Speed Ethernet Media Access Controller, and PCI 33 MHz Target, which automatically integrate into Lattice's Mico System Builder. Middleware has enabled designers to quickly and confidently configure a microprocessor in their design at very low cost, or no cost at all.

**PL: What should designers be doing differently now to get an advantage with their next FPGA-based design?**

**DLR:** We think in terms of two fundamental types of FPGA applications: *control-oriented* applications and *data path* applications. Control-oriented applications really have not evolved much over the years, implementing numerous small finite state machines and utilize many parallel I/Os for system monitoring and control. There are not too many new issues to deal with in this area.

Data path applications, however, have continued to grow and evolve. It is more important than ever for a system designer to consider how to best architect systems to effectively leverage the new high-speed SERDES-based capabilities of FPGAs. This increased data processing bandwidth will allow for radically new system-level architectures that can provide dramatically higher levels of cost/performance.

**CF:** And there's increasingly sophisticated functionality in ispLEVER to help the architect. Also, FPGA IP is particularly important, offering customers time to market and risk management advantages by providing proven, hardware-validated solutions

that are typically parameterizable. These pre-engineered IP cores are cost-effective, and can save countless hours of development and validation effort.



**Chris Fanning** is Corporate Vice President of Enterprise Solutions and manages Lattice's software, intellectual property, and technical support businesses. He previously worked at The Boston Consulting Group. He holds an MS in Computer Science from Worcester Polytechnic Institute, and an MBA from The University of Chicago.

**David Lee Rutledge** is Corporate Vice President of FPGA Product Development, with two decades of programmable logic experience including development of the first in-system programmable PLD. After working for Harris Semiconductor, Rutledge joined Lattice in 1983 as its 31st employee. He holds a BSEE and MSEE from Purdue University.

Lattice Semiconductor

503-268-8000

[www.latticesemi.com](http://www.latticesemi.com)

**There's a chameleon in your stack**

**FPGAs that transform to meet your needs**

**FreeForm/PCI-104**  
Virtex-5 FPGA  
Flash & DDR2-400 memory  
64 single ended or 32 LVDS GPIO  
Designed for embedded processing

**FreeForm/104**  
Spartan-3E FPGA  
Field reprogrammability  
Standard core includes: 96 digital I/O,  
6 counter/timers, Opto-22 compatibility

**Connect Tech Inc.**  
Industrial Strength Communications  
[sales@connecttech.com](mailto:sales@connecttech.com) | 519-836-1291 | [www.connecttech.com](http://www.connecttech.com)



# Non-volatility: the choice is clear

By David Lee Rutledge

Non-Volatility (NV) offers the same compelling benefits to FPGA users as it always has to CPLD users: a single-chip solution, instant-on capabilities, design security, and more. In fact, these features are arguably more valuable to FPGA applications than to CPLD applications.

So, why isn't everyone making NV FPGAs?

NV FPGAs require much more sophisticated process technology and circuit design techniques than their SRAM-based cousins. The key to NV success is providing sophistication without significantly impacting cost or performance, compared to an equivalent SRAM FPGA.

## Anatomy of a true NV FPGA

A true NV FPGA must employ high-performance, embedded flash memory technology (Figure 1). Historically, this approach had relatively low performance due to the additional wafer processing steps required to embed flash memory in a logic process. Lattice has worked closely with our foundry partner Fujitsu to develop proprietary embedded Flash technologies that enable exceptionally high performance for minimal additional cost.

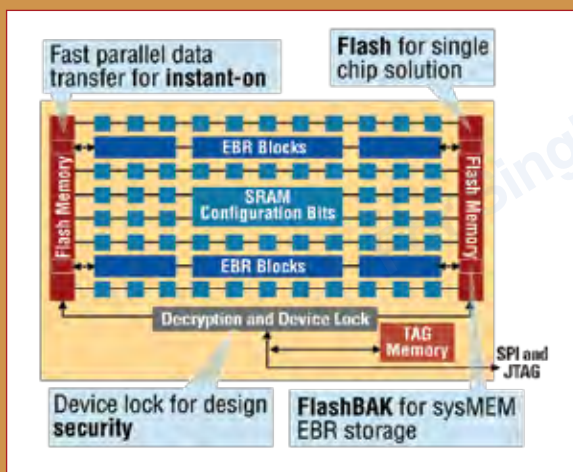


Figure 1

*Instant-on*, a feature that allows an FPGA to be used immediately upon startup, is essential in many applications. Typical SRAM FPGAs require 10s to 100s of milliseconds at startup using external boot PROMs. This prevents the use of SRAM FPGAs in certain critical control logic sections of systems that must be "live" immediately upon the application of power. Lattice has developed proprietary circuit design techniques that allow its NV FPGAs to become active virtually instantly after startup (within 1-2 ms).

Design security is also a key advantage for our NV FPGAs, in two distinct ways. First, our embedded flash memory allows the entire user design to be stored in on-chip Flash memory. The configuration bitstream is never exposed, as it is when using SRAM FPGAs with boot PROMs. It is virtually impossible for a hacker to access the configuration data of a NV FPGA.

Second, our embedded flash memory allows support of true AES encryption of the configuration bitstreams (Figure 2). This is a new feature that has been added to the latest 90 nm LatticeXP2 family. It supports the storage of a user-programmable, on-chip 128-bit encryption key. Once this key is stored on-chip, the user can encrypt the configuration bitstream for a given design and transmit it to the FPGA (over unsecured transmission lines). Once on-chip, the bitstream is decrypted using the stored AES key, and is never exposed off-chip.

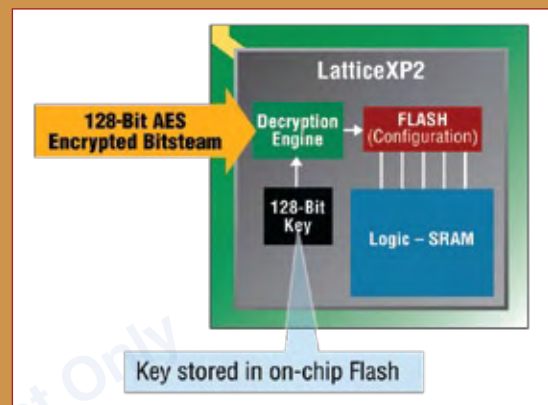


Figure 2

## Alternative NV technologies

FPGAs have been developed that are called NV, but they are not true single-chip, embedded flash solutions. Instead, they combine two independent chips (a flash memory and an SRAM-based FPGA) in a single package using a stacked die technology. Compared to true NV FPGAs, these hybrid products have severe limitations.

For example, stacked die technology cannot support instant-on because essentially it is identical to two-chip, SRAM-based technology, and so has the same lengthy boot time as an SRAM FPGA (Figure 3).




Figure 3

Stacked die technology also does not enable ultimate design security because any stacked-die implementation must transmit the unencrypted configuration data from the boot PROM to the SRAM FPGA over bond wires within the package. A hacker need only perform minor surgery to gain access to the wires that transmit the configuration data. **PI**

57,930  $\mu$ W  
35,210  $\mu$ W

# FPGA Static Power



## WHAT DO WE HAVE TO DO, DRAW YOU A PICTURE?

Only Actel gets you this close to zero. Any other claims of low power superiority are just that. According to their own data, Altera® and Xilinx® use between 10 and 1700 times the power of Actel IGLOO® FPGAs, depending on device and mode. Want specifics? Visit us to get the whole picture, including a video of actual measurements.

\* more proof and pictures at [actel.com/power](http://actel.com/power)

**Actel**  
POWER MATTERS



# "Projecting" images in radar and medical applications

By David Pointer

**E**xamples of computationally intense systems being shrunk by use of FPGAs are of keen interest to us. Filtered back-projection is finding its way into both radar and medical imaging applications, and is well served by FPGAs to handle a portion of the algorithm. The results are outstanding.

Searching for an enemy vehicle from a high-flying Unmanned Aerial Vehicle (UAV). Taking a child after a tumble from a bike to the hospital for a CT scan of their broken arm. The imaging techniques in these two seemingly unrelated scenarios rely on the same technology to produce accurate results: the Filtered Backprojection (FBP) algorithm.

The FBP algorithm reconstructs an image of an object by calculating an exact solution for each pixel from a series of planar image data sets (projections) of the object. Both radar and medical imaging require a high degree of imaging accuracy. The FBP algorithm provides excellent imaging accuracy but at a large computational cost.

Heterogeneous reconfigurable systems provide the required computational power for image reconstruction to both Synthetic Aperture Radar (SAR) Backprojection imaging and Computed Tomography (CT) imaging while delivering such significant benefits as smaller form factors and reduced power consumption.

## Partitioning the FBP

The general implementation of the FBP algorithm is quite straightforward. A two-dimensional set of floating point data representing the energy detected from illuminating an object (using radar energy for SAR, or X-Rays for CT) is filtered to remove noise, then each projection's contribution to the reconstruction is summed into each pixel in the image.

For images of a useful size, the summation operation over all pixels is the computational bottleneck. Summing each projection's contribution for a single pixel is fast; summing the effect of each projection over one million pixels for a 1024 x 1024 image – or even larger images – is intensely time-consuming.

Programmers can achieve high performance for FBP algorithms by spreading the application across an FPGA and traditional

CPU combined in a heterogeneous reconfigurable system as shown in Figure 1. An FPGA contributes to a reconfigurable system's performance by allowing a programmer to explicitly and completely dedicate a device to the solution of the regular, uninterrupted streaming aspects of a program. Along with this computational intensity, an FPGA also provides parallelism which can be exploited. The CPU, with its high clock rate, is an excellent workhorse for executing the irregular and conditional aspects of any program.

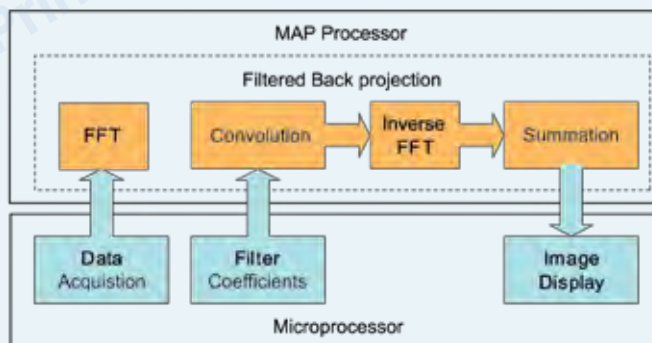


Figure 1

Equally sized sets of projection data stream into the FPGA for filtering, which is easily implemented in the frequency domain as a multiplication operation. Spatial domain projection data sets are streamed through a *fast fourier transform* (fft), multiplied by a filter coefficient matrix, and then converted back to the spatial domain using an inverse fft. These three operations are all pipelined (overlapped in time) for maximum data throughput through the FPGA.

After filtering, the backprojection (summation) step of the FBP algorithm is independent for each pixel, and so may be implemented as a set of simultaneously executing parallel summation units in the FPGA program for maximum application performance.

In a typical imaging application, the microprocessor generates appropriate filter coefficients and acquires the sensor input data. After FBP processing by the MAP processor, the final image data may be displayed or stored by the microprocessor.

## Shrinking an SAR system

Given the high computational cost of the backprojection algorithm, it can be an engineering challenge to deploy such radar imaging in portable computing systems and small UAVs while keeping overall vehicle size, weight, and power (SWaP) down to acceptable levels.

During the early specification requirement studies for various UAV programs, the United States Air Force was concerned about the feasibility of deploying SAR backprojection on mid-sized UAVs. Could a reconfigurable SAR system meet both SWaP requirements and processing requirements?

To address these concerns, engineers from the Air Force Research Laboratory (AFRL) and SRC Computers jointly benchmarked the Spotlight Synthetic Aperture Radar (SAR) algorithm in 2005. The results obtained on the SRC-6 Portable MAPstation, which contains both an Intel 1.6 MHz Pentium M CPU and SRC's FPGA-based Series F MAP processor, showed a 75x performance increase using both the CPU and MAP over using just the CPU. The absolute wall clock time for the benchmark execution exceeded the Air Force requirements.

Based upon these published results, Lockheed Martin recently selected SRC to provide the Signal Data Processor (SDP) for the United States Army's Tactical Reconnaissance and Counter-Concealment Enabled Radar (TRACER) program. This system, scheduled to fly on the Warrior UAV in 2008, contains multiple MAP processors for even greater throughput.

### Extending to medical imaging

Last year, discussions with customers in the medical imaging field indicated that CT scan image reconstruction used the same FBP algorithm as SAR imaging. The application engineers at SRC Computers found an open source CT scanner simulation program, CTsim, and ported the FBP portion of this medical imaging application to the SRC-6 MAPstation.

The results of this implementation, reconstructing a 1024 x 1024 image showed a 22x performance improvement with the combined CPU and Series E MAP over the MAPstation's 2.8 GHz Xeon 32-bit CPU. These results indicate that manufacturers of CT scan equipment could increase the resolution and clarity of the output of their equipment, draw significantly less power than traditional CPU-based systems, and still enjoy faster image reconstruction.

### SRC-7 MAPstation: Early results

When the new SRC-7 MAPstation started shipping in 2007, application engineers at SRC began work to obtain updated application performance data on several programs, including CTsim and SAR back projection.

For CTsim, a direct recompilation of the existing Series E MAP processor code without using any of the Series H MAP enhancements yields a 29x performance improvement when compared to a 3.0 GHz Xeon 64-bit CPU. Proposed modifications to the FBP implementation in CTsim in order to take advantage of the Series H MAP (shown in Figure 2) suggest that a 60x performance improvement may be realized with minimal effort.




Figure 2

*An FPGA contributes to a reconfigurable system's performance by allowing a programmer to explicitly and completely dedicate a device to the solution of the regular, uninterrupted streaming aspects of a program.*

For SAR backprojection, optimization work is underway to port this code to the Series H MAP. Analysis indicates that a 104x performance improvement (relative to a 3.0 GHz Xeon 64-bit CPU) is a reasonable expectation when this work is complete. It is worth noting that the high-performance fftw library was used to obtain the CPU-only execution measurements. The final performance results for CTsim and SAR backprojection executing on the SRC-7 will be published as soon as they are verified.

Some of the SRC-7 system improvements that enable this performance include a 50 percent faster FPGA clock rate, a five-fold increase in system interconnect payload bandwidth (from 2.8 GBps to 14.4 GBps), and the adoption of Altera FPGAs, resulting in a more than three-fold increase in the MAP processor's floating point performance.

### Reconfiguring system performance

Actual, realizable application performance on reconfigurable systems is important, but it is also good to consider other aspects of application development. FPGAs were once the domain of hardware engineers and Hardware Descriptor Languages (HDLs) like Verilog and VHDL, and programming FPGAs was notoriously difficult.

Recent advances, however, now make reconfigurable systems accessible to traditional software programmers. SRC Computers provides a unified programming environment, called Carte, which allows the programmer to use ANSI C or FORTRAN languages for combined CPU and MAP coding.

Software engineers, take note. Reconfigurable computing is spanning more industries and applications than ever before, and is no longer just a playground for hardware engineers.



**David Pointer** is Director of System Applications for SRC Computers, Inc. (Colorado Springs, CO). His engineering career spans 29 years in both industry and academia, most notably at Hewlett-Packard Laboratories in Palo Alto, CA, and the University of Illinois at Urbana-Champaign. He received an M.S. in Electrical Engineering in 1982 from Northern Illinois University and an M.S. in Computer Science in 1996 from the University of Illinois at Urbana-Champaign.



**SRC Computers**

719-262-0213 • [www.srccomputers.com](http://www.srccomputers.com)



# MATHSTAR™

## 1 GHz

# Field Programmable Object Array™



### **FPOA Advantages:**

- 1 GHz Operation
- 1 GHz Interconnect Fabric
- No Timing Closure Required
- Faster Time-to-Market

**Discover our FPOA product family at:**

**[www.mathstar.com](http://www.mathstar.com)**

# SIGINT in the real world presents nuances

By Dr. Malachy Devlin



*We've heard that FPGAs can do signal processing algorithms. But what about the rest of the problem: controlling analog-to-digital converters, compensating for effects of voltage and temperature, and more? Here, FPGAs get serious.*

The growing use of FPGAs in SIGINT applications is enabling a significant increase in the capability of defense systems. With a wider variety of threats than ever before, the desire to monitor more of the frequency spectrum necessitates a high performance front-end processing engine. A SIGINT system that needs to monitor VHF (30-300 MHz), UHF (300-1000 MHz) communications and navigational aids, through to radar and satellite communications in the L band (1-2 GHz) frequencies, requires an analog bandwidth of 2 GHz for complete and instantaneous coverage.

Since there are a wide range of frequencies and consequently a wider range of signal modulation schemes, the processing system needs to have significant I/O bandwidth to receive the data as well as the computing performance to implement real time demodulation and analysis algorithms. Typical algorithms include Digital Down conversion, ffts, filtering, and I/Q demodulation.

## DSPs short on two counts

While many of the algorithms used are familiar, it is the bandwidth of the input data that makes the implementation much more challenging. As an example, taking a simple 10 tap FIR filter and using this with a high end ADC, such as a 3 GSps part, the computing performance required is 30 GMACps. While this significantly exceeds the performance of high end DSP processors that offer around 1 GMACps, it is a fraction of the performance offered by high end FPGAs that are capable of up to 350 GMACps. This means that FPGAs have horsepower to spare to carry out additional tasks which will be required by typical applications.

Not only are the performance capabilities of DSPs under-powered for high data rate processing, but the I/O bandwidth to move these high speed data streams is unavailable. While there could be a data stream of 3 GBps from an ADC, DSPs typically provide bandwidth of 1 GBps. Clearly DSPs are not capable of handling this rate of data ingress for a single data channel. This problem is exacerbated considering antenna arrays and techniques such

as beamforming are common approaches to improve the SNR of received signals.

FPGAs can tackle the problem in a much smaller space. Figure 1 shows a 6U CompactPCI card with 4 daughtercards, providing a total of 8 channels that can each capture 3 GSps at 8-bits (aggregate of 24 Gbps) into 5 FPGAs that can deliver up to 1 TeraMAC/sec of processing performance. The equivalent implementation using DSP processors would require at least a full 42" rack.




Figure 1

## More control over digitization

In addition to the I/O and processing power that FPGAs provide, they bring further benefits to front-end processing applications by affording the user a finer level of control over the digitization process. FPGAs can directly control and manipulate the interface signals of ADCs and supporting peripherals such as clock synthesizers and clock path selection. Coupling high-speed data converters with FPGAs allows direct RF sampling, enabling digital manipulation of data and hence reducing complex analog RF circuitry for tuning and demodulation. Performing these functions in the digital domain rather than analog reduces sensitivity to environmental noise and improves the repeatability of processed results.

For example, incoming SIGINT data from antenna arrays must be captured with virtually no phase skew between channels. The challenge is ensuring two ADCs are accurately synchronized in order that the samples on each analog input are coincident in time. Channel synchronization is similarly critical when capturing I/Q signals, where the SNR will be degraded if the data from the two channels is misaligned resulting in reduced receiver performance and sensitivity.



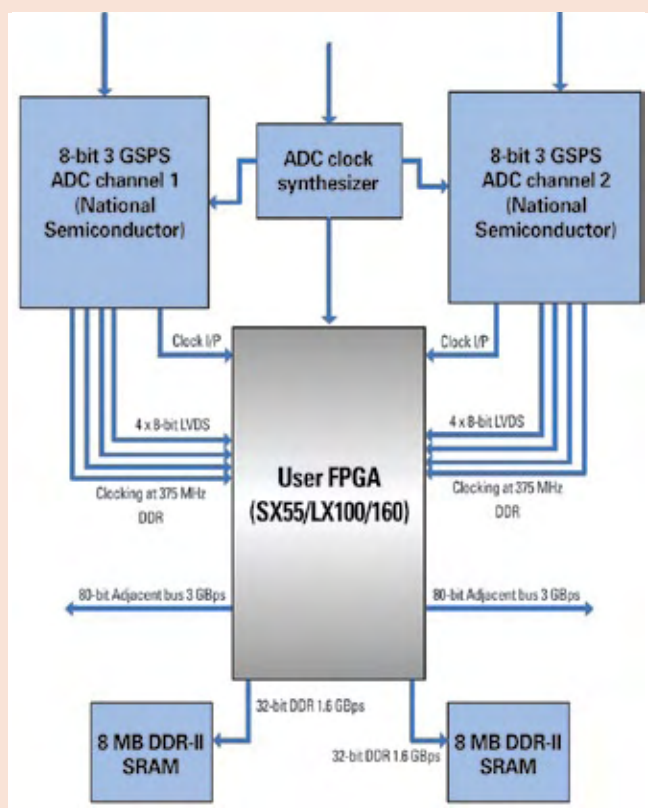


Figure 2

### A lot goes into synchronicity

Figure 2 is a functional diagram of the daughtercard module (with four shown in Figure 1) with two National Semiconductor ADC083000 3 GSps ADCs connected directly to a Xilinx Virtex-4 FPGA. The key to synchronizing the ADCs is deasserting the device reset signal.

There are many parameters that need to be considered to ensure synchronicity of the channels, particularly when having to deal with the real world differences of time of arrival of the reset signals. In this specific implementation, the key factors that had to be considered in the analysis were:

1. Skew between clock inputs of the two ADCs
2. Skew between FPGA reference clock and ADC clocks
3. Reference clock delay through FPGA
4. Skew between the two ADC reset signals include FPGA clock to out skew
5. Relative jitter of reference FPGA clock to ADC clocks
6. Relative jitter between ADC clocks
7. Setup and hold of reset signal

Although the FPGA is used for handling the timing of the reset, some of these parameters are out of the control of the application developer (namely, items 1, 5, 6, and 7) and so need to be characterized in order to be managed. Fortunately, the flexibility of the FPGA means it can be used to run a dynamically configurable measurement application that measures the effect of the *uncontrollable* parameters. Using such an application, the valid window for the reset signal was measured as 306 ps, from

*Fortunately, the flexibility of the FPGA means it can be used to run a dynamically configurable measurement application that measures the effect of the 'uncontrollable' parameters.*

a total period of 666 ps (assuming a 1.5 GHz clock that is used for 3 GSps data capture).

The critical element to maximize the reliability of channel synchronization is to deassert the reset signal in the middle of this data valid window. This shifting of the reset edge is achieved using the digital clock managers (DCMs) of the FPGA, which enable fine-tuning of the clock phase within the FPGA. The two reset signals to the ADCs are transmitted from the FPGA via flip-flops so when the phase of the FPGA clock is changed, the phase relationship between the reset signal and the ADC clocks is also changed (the ADCs are clocked directly from the clock synthesizer and are therefore independent of the FPGA). This technique enables the phase of the reset signal in relation to the ADC clocks to be changed in increments of 10.4 ps when sampling at 3 GSps.

### Compensating for environmental variables

At this stage it seems a straightforward task to run the measurement design, find the center of the data valid window, and pre-configure the phase of the DCM to align the reset signal at that location. Unfortunately, environmental parameters affect skew and jitter performance.

Testing showed that the phase of the valid window for the reset signal (relative to the ADC clocks) changed with temperature. A second order non-linear relationship was found between the phase setting and the frequency and temperature values, which also varied depending on the FPGA device in use. The resultant formula for a Virtex-4 SX55 device is:

$$\text{Phase}(\text{SX55}) = \text{round}(\text{mod}(1.5479\text{e-}006 * f * f - 0.00045165 * f * t - 0.24071 * f - 0.0033431 * t * t + 0.27627 * t + 348,256))$$

Where:  $t$  is the temperature in °C  
 $f$  is the clock frequency of the ADCs in MHz

This result greatly simplifies the system design as the phase setting can be computed based on two easily known variables. The alternative is to implement a calibration circuit in the FPGA and provide synchronized analog inputs on every startup, a more error prone technique.

### Space saved

This application would require several hundred DSPs to implement, not to mention the need for a PLD to interface between the ADCs and processors, illustrating why FPGAs have become the first choice technology for front-end sensor processing. Not only do FPGAs provide computational prowess that reduce system SWaP by orders of magnitude, they have the I/O bandwidth and flexibility needed to implement the latest and most complex sensor interfaces.



**Dr. Malachy Devlin** is CTO of Nallatech (Eldersburg, MD). He joined Nallatech in 1996, and is recognized worldwide as an industry expert on FPGA technologies. His experience includes roles with the National Engineering Laboratory (UK), Telia AB, and Hughes Microelectronics. He has a PhD from the University of Strathclyde, Glasgow, Scotland.




**Nallatech**

410-552-3352 • [www.nallatech.com](http://www.nallatech.com)

# Kiss your FPGA good-bye



eASIC's no minimum order, Nextreme zero mask-charge ASICs are ideal for replacing your expensive and power hungry FPGAs. They are simple to design and you can have production ready devices back in only 4 weeks. With Nextreme you can also design complete systems that include embedded processors such as ARM or Tensilica cores.




## Up to 10X Lower Cost Than FPGAs

- Nextreme replaces SRAM interconnect with single via layer single
- Dramatically reduces die size & cost

## Up to 15X Lower Power Than FPGAs

- Significantly reduce both standby and dynamic power
- Reduce power supply complexity and cost
- Eliminate expensive FPGA heat-sinks and fans



Learn more at  
[www.eASIC.com](http://www.eASIC.com)



# Big images, complex processing? Think objects

By Tom Diamond

**G**ot a 4K x 4K image to process? Equalizing, correcting, interpolating, adjusting, and filtering every pixel at 35 fps? The Field Programmable Object Array enters the fray and shrinks the processing hardware required.

Today's real-time image processing hardware is feeling the pinch. Gigabit Ethernet links feed multi-million pixel full color space video to frame grabbers. Multi-core processors, fed by high bandwidth PCI Express ports, consume image processed data as fast as it can be delivered. It's no wonder that traditional reprogrammable devices are straining to keep pace across a range of automated inspection, security/surveillance, and professional video applications.

The Field Programmable Object Array (FPOA) is a high performance device that operates at speeds up to 1 GHz and is programmed at the object level. FPOAs are especially well suited to meet the requirements of ultra-fast high resolution image processing applications. The key to their performance is hundreds of objects that pass data and control to each other through a patented interconnect fabric.

The MathStar Arrix family of FPOAs provides 256 Arithmetic Logic Unit (ALU), 80 Register File (RF), and 64 Multiply Accumulator (MAC) objects. The objects and the interconnect fabric run on a common core clock, operating deterministically at 1 GHz in the Arrix Family's fastest device. This determinism enables

a designer to select a core clock that meets the desired memory, I/O, and processing requirements.

## An image processing example

In a high-end image processing implementation, 4K x 4K pixel color images are fed to an FPOA from the electro-optics of a high performance color CCD or CMOS camera as a single color channel at 35 frames per second. Use of a 900 MHz Arrix device enables external memory transfers at the FPOA's 300 MHz top memory bandwidth.

The implementation has several stages:

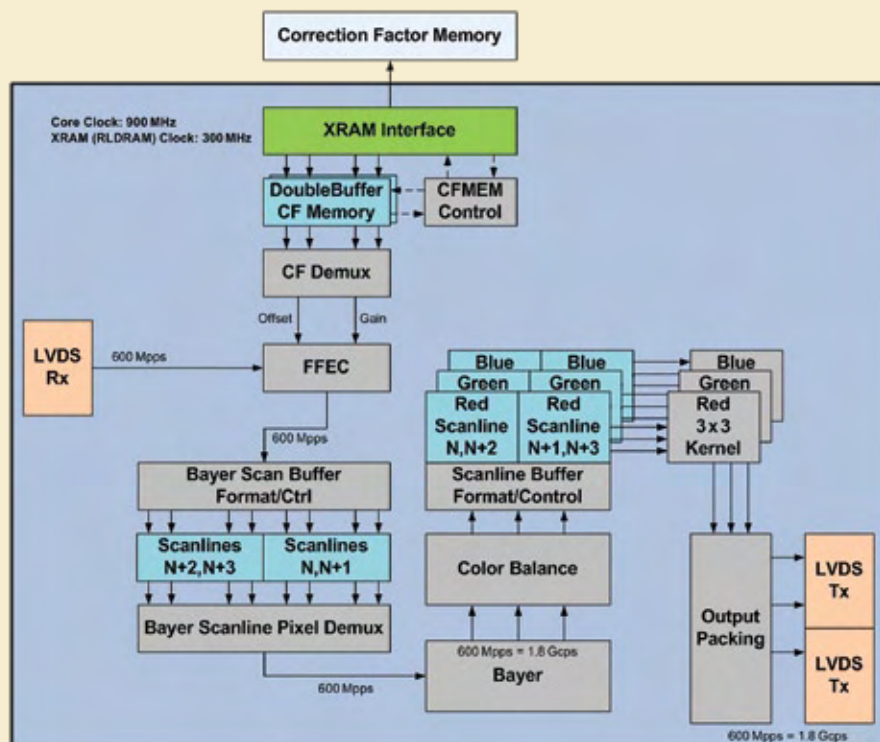
- Non-linearities from process variations in the sensor array are corrected in the flat-field processing stage, equalizing pixel gain across the array and correcting for dark current flow in the sensor array.
- The Bayer demosaicing process generates RGB pixels by interpolating the color value for each pixel based on the color filter array architecture. Demosaicing increases the number of color channels to a total of three, effectively tripling the processing load.
- RGB data is then adjusted across all color-space components for the ambient lighting qualities or targeted display qualities in the color balancing processing block.
- Lastly, the image is filtered using a 3 x 3 convolution kernel before being driven out of the FPOA at the pixel rate.

Flat-field correction, Bayer demosaicing, color balancing and spatial filtering are mapped to the FPOA, as shown in Figure 1.

## Flat-field correction

Flat-field correction is used to adjust image-sensor output data to ensure that constant intensity images generate constant pixel values at various levels of image intensity. This process addresses three types of pixel-based non-uniformities: gain, dark current offset, and defective pixels.

*The key to this performance is hundreds of objects that pass data and control to each other through a patented interconnect fabric [running at up to 1 GHz].*




### Figure 1

To rectify these non-uniformities, both a calibration and correction process must be performed. The calibration process determines the correction factors for pixel gain and offset and generates a defective pixel map. The correction process calculates an appropriate value for non-uniform pixels, and includes a dead pixel correction unit. Gain and offset correction can be performed using a single MAC object running at the core clock rate. Correction factor memory resides off-chip and can be accessed as two correction factor pairs every memory clock cycle.

### *Bayer demosaicing*

A Bayer color filter array filter is an optical filter that routes specific wavelengths of light to a particular photo-sensor in a digital camera's sensor array. Each photosensor will detect the intensity of a certain light band. After filtering, the classical Bayer filter decomposes light into red, green, and blue components.

The digital processing associated with Bayer filters determines the missing color values for a given pixel (for example, the red and green levels for a blue sensitive sensor). Reconstructing the missing pixel colors of the CFA is known as demosaicing. This processing stage interpolates the missing color values of the color filter array (CFA) sensor pattern and generates a fully populated RGB image.



### Figure 2

Pixel size at the input is a single red, green, or blue value and, at the output, consists of red, green and blue components for each pixel. Given a frame rate of 35 fps for the 4K x 4K image, this interpolation stage increases the processing throughput from

600 M to 1.8 G color components/sec. 18 ALUs support a 600 Mps input, generating three parallel outputs for each pixel (shown in Figure 2), interpolating in a 3 x 3 region about the pixel of interest.

### *Color balancing*

The color balancing operation utilizes a nine-entry color transform matrix to correct for monitor or lighting irregularities. A fully parallel color balancer implementation at the core clock rate utilizes three MAC objects for each row of the color transform. A single RF per MAC for table distribution utilizes the local memory resources, localizes data into each MAC, and provides for support of up to 64 correction tables.

### *Spatial filtering*

Image filtering for smoothing, gradient or edge enhancement, and sharpening is accomplished with a two-dimensional convolution using a 3 x 3 pixel mask. A fully parallel implementation of the filter is capable of calculating the result for one input every core clock. Since each color component needs to be processed separately with potentially separate masks per color component, a total of three filter processing blocks are required.

## Summary

Performing flat-field correction, demosaicing, balancing, and filtering stages for a 4K x 4K, 35 frame per second, full color space frame grabber is a tall order, but is easily handled by a high performance FPOA. Further, this image processing chain utilizes just 50 percent of the objects in a single FPOA, providing room for additional image processing functions.



**Tom Diamond** is marketing director for MathStar, Inc. Previously, he managed roadmap planning and product definition for Intel's LAN Access Division, and has worked in consumer electronics, test and measurement, and high volume silicon firms. Tom holds a BSME from Duke University and an MBA from the University of Washington.



**MathStar, Inc.**  
503-726-5500  
[www.mathstar.com](http://www.mathstar.com)




# LatticeXP2 & MachXO

## More of the Best

- ✓ Low Cost Mainstream FPGA Architecture
- ✓ Single Chip Integration
  - Instant-on through on-chip Flash
- ✓ On-chip RAM and DSP
- ✓ Flexible I/O
  - Up to 540 I/O
- ✓ Small Footprint
  - 8 x 8mm csBGA packaging

**Live Update Programmability**

Fast, secure design changes.



## The Best Low Cost, Non-Volatile Logic

Fewer chips, smaller footprint, instant-on start-up ... those are just a few of the benefits of Lattice's non-volatile programmable logic devices. Lattice offers the broadest range of non-volatile industry-standard architectures available. With on-chip block RAM, enhanced design security, and the flexibility to accelerate time-to-market, our latest LatticeXP2™ and MachXO™ devices deliver More of the Best.

LatticeXP2 devices combine an efficient FPGA fabric with non-volatile Flash cells in an architecture that was optimized from the outset with high performance and low cost in mind. LatticeXP2 FPGAs offer 5K to 40K LUTs, distributed and embedded memory, Phase Locked Loops (PLLs), pre-engineered source synchronous I/O and enhanced sysDSP™ blocks.



**MachXO Evaluation Board**

Complete FPGA starter kits  
starting at only \$99!  
[www.latticesemi.com/99](http://www.latticesemi.com/99)

The MachXO family combines the best attributes of CPLDs and FPGAs to serve applications that traditionally were implemented in CPLDs or low capacity FPGAs. MachXO devices include PLLs and embedded memory along with other features you wouldn't normally expect in a CPLD.

Make your next design your best with non-volatile programmable logic. Go to [www.latticesemi.com/nv](http://www.latticesemi.com/nv).

**Lattice**  
Semiconductor  
Corporation  
*More of the Best*

# Programmable logic: The key to effective interface design

By Doug Morrissey



*On the topic of shrinking designs using programmable logic, Doug walks us through how complex systems like media gateways are improved using FPGAs to incorporate faster, and sometimes missing, interfaces to DSPs.*

Now that telecommunications OEMs are seeing sustained growth in the VoIP arena, margins are being squeezed and equipment designers are being asked to do more with less. For example, IP PBX and multiservice access nodes (MSANs) need to be price competitive at many different channel densities, and yet, the engineering cost of designing many different platforms is exorbitant.

Media gateway functionalities are best offered today using DSP technology. A few select vendors offer field-proven products that include a complete hardware and software solution. But in order to properly take advantage of these solutions, hardware designers are asked to make compromises.

The main problem is that DSPs don't always offer the right combination of DSP performance and interfaces. For example, very often it is only the largest DSP in a family that supports high-end interfaces such as SGMII Ethernet or RapidIO. When it comes time to interconnect many DSPs in larger systems, an FPGA is often used to glue devices together and perform functions such as back-pressure on packet transmission or packet classification. Hardware designers have already understood the value of using programmable logic in these scenarios.

## Same DSP, but different interfaces

While different types of equipment can use the same underlying media processing technology, they each have specific requirements and constraints. Very often, the same DSP devices are targeted at all of the following applications:

- Large-scale media gateways
- IP PBX
- HMP media server off-loading
- MSAN/Access equipment

While the basic feature set may be the same, all applications require VoIP processing with a list of codecs and voice quality

features; the interfaces required for each may be quite different. For example, a PCI card designed for host-based media processing (HMP) off-loading would require a PCI Express interface, whereas a MSAN may require a lightweight interface to an Ethernet PHY and to T1/E1 framers, as well as a lot of glue logic to the various components on the board. Sometimes the ideal solution is "off the beaten track." Designers may sometimes find a more cost-effective way of doing things, but find that they are limited by what the DSPs support. In the IP PBX space, certain vendors have discovered that USB is the ideal interconnect for low-cost systems, instead of Ethernet.

USB provides many advantages. First, it is hot-pluggable by design. More importantly, it offers a cost structure that is more advantageous than Ethernet. A USB hub device is less expensive and less complex than an Ethernet switch, which means that the base cost for a system is lower. Furthermore, the incremental cost is lower, because for each processing blade added to the system, very simple, low-cost RISC processors are available with USB interfaces. Most DSPs offered into the VoIP gateway market do not support USB, whereas FPGA vendors provide IP cores for such an interface. Leveraging FPGAs allows for this kind of innovation.

## Combining DSPs with FPGAs

Ideally, OEMs should be able to get exactly the interfaces they need, without the cost and drawbacks of buying a fully-featured DSP. Typically, the DSPs that contain high-speed interfaces tend to have a large number of different interfaces, and be in a very large package with a very high pin count. These large packages are dense BGAs that are expensive and require a PCB with a high number of layers, which can all contribute to design delays.

By combining an FPGA with a DSP, the hardware designer can get the best of both worlds. In cases where the DSP interfaces are judged inadequate, or customization was required, DSPs are often front-ended with an FPGA. Adding an FPGA to an already fully-featured DSP is overkill. It adds unnecessary cost and board real estate, only to correct short comings in the DSP.

*Ideally, OEMs should be able to get exactly the interfaces they need, without the cost and drawbacks of buying a fully-featured DSP.*



A different approach can be taken when combining a DSP and FPGA. In this method, the DSP serves only as a DSP. The physical package contains a high-performance, multi-core device and uses a single memory interface to communicate with the FPGA. In this case, there is no carrying cost for the interfaces inside the DSP, keeping area and power to an absolute minimum. The DSP vendor provides turnkey FPGA designs for each major application, which can then be customized by the OEM if desired. The FPGAs used can be high-end devices or very small, low-cost devices, depending on the target application (Figure 1).




Figure 1

The main advantage to this approach is the flexibility that it offers. FPGAs can be reprogrammed in the field if necessary. Furthermore, FPGAs are always among the first devices to be released at each new process node. Therefore, they always support the latest I/O technologies with optimal power performance. As interconnect standards continue to evolve, it is possible to retarget a design very quickly to the latest fabric, whether it's PCI Express, RapidIO, or a proprietary interface.

The combined DSP and FPGA approach offers many cost savings. It is important to look at the total solution when evaluating cost. For example, there is a plurality of Ethernet standards available: MII, RMI, GMII, RGMII, SGMI, and so on. When deciding how to interconnect many VoIP DSPs together, hardware designers can choose from a large number of Ethernet switch ICs, or consider a FPGA solution. If a certain Ethernet switch offers the ideal price point, but does not support the interfaces provided by the DSP, then a more expensive switch must be used.

## More savings possible

The cost savings can be further extended beyond the electrical interface. For example, many large OEMs use proprietary headers or flags to help classify packet traffic. With an FPGA front-end, this becomes trivial and can greatly increase performance by off-loading certain tasks from the DSP. This type of header manipulation is best done in hardware by an FPGA. This offers deterministic performance, and keeps control of the design in the hands of the OEMs.

A media gateway DSP often incorporates a TDM interface that is connected to timeslot interchangers (TSIs) or T1/E1 framers. Very often, an FPGA is already present between the DSP and these TDM-based devices, or the TDM backplane. FPGAs are currently used because custom logic is required, or because the FPGAs offer more reliable I/Os that can sustain the high voltage spikes often seen on backplanes. These FPGA tasks can now be incorporated into the front-end FPGA saving more money and board area.

Another major advantage of this FPGA front-end is to make life easier for designers who again find themselves "off the beaten track." DSP manufacturers release products with feature sets and interfaces for specific market segments, usually the largest ones. It is easy to find a DSP that can serve as a generic TDM to VoIP gateway, but if the job is to design a high-performance telepresence video over IP conferencing unit, then the choice of DSPs is quite slim. The amount of bandwidth required for passing compressed or uncompressed video streams between DSPs is immense. Finding a DSP with the right interface for this or any other emerging market is quite difficult and costly.

Instead, using an FPGA front-end allows access to the latest I/O technology, as well as IP blocks for the latest standards. The designer can now get the right interfaces, without making changes to the DSP software. In organizations with large FPGA teams, the designer may choose to develop a proprietary interface that meets his or her needs perfectly and provides a competitive edge.

Media gateway designers are faced with many challenges including increased price pressure and shortened design cycles. By using a DSP that is tightly coupled to an FPGA, designers are able to get exactly the interfaces they need, without the added cost or debugging of useless I/O banks. This greatly simplifies the design and allows OEMs to add their own unique features and differentiate their system. This kind of innovation is what allows OEMs to address current and future customer requirements while maintaining a common DSP software platform.



**Doug Morrissey** is Vice President and Chief Technology Officer at Octasic and has over 10 years of experience in the definition and marketing of semiconductor devices. Joining Octasic in 1999, he strategically focuses on issues within the Voice over Packet market. Previously, he worked as Marketing Manager for ATM and DSL products for Agere (formerly Lucent Technologies, Microelectronics Group), and was Senior Systems Architect at Unisys Corporation. Morrissey holds a BSc from Rochester Institute of Technology.



## COVER PHOTOS:

Top left: Xilinx Virtex-5  
 Top center: Jacyl Technology XG-5000K  
 Top right: eASIC Nextreme  
 Middle left: Altera Stratix III  
 Middle right: Actel IGLOO  
 Bottom left: MathStar Arrix  
 Bottom middle: Lyrtech SFF SDR platform  
 Bottom right: Lattice Semiconductor LatticeXP2

## ADVERTISER INFORMATION

Page	Advertiser/Ad title
15	Actel Corporation – FPGA static power
28	Altera Corporation – Full throttle
2	Annapolis Micro Systems, Inc. – WILDSTAR
13	Connect Tech, Inc. (CTI) – There's a chameleon in your stack
21	eASIC – Kiss your FPGA good-bye
7	Innovative Integration – Reconfigurable FPGA Processing
3	Jacyl Technology Inc. – XG-5000K
24	Lattice Semiconductor Corporation – Lattice XP2 & MachXO
11	Lyrtech – Most versatile programmable logic
18	MathStar – 1 GHz Field Programmable Object Array™
27	Pentek, Inc. – Introducing the One Board with all the Right Connections
5	XILINX, Inc. – Lowers Total Cost... Period.

Octasic Inc.

514-282-8858 • [www.octasic.com](http://www.octasic.com)





## Arrivals

Origin	Board	Model	Status
FIBRE CHANNEL	PENTEK	4207	NO DELAY
SERIAL RAPID I/O	PENTEK	4207	NO DELAY
PCI EXPRESS	PENTEK	4207	NO DELAY
GIGABIT ETHERNET	PENTEK	4207	NO DELAY



## Departures

Destination	Board	Model	Status
PCI-X	PENTEK	4207	NO DELAY
VXS	PENTEK	4207	NO DELAY
PMC / XMC	PENTEK	4207	NO DELAY
ROCKET I/O	PENTEK	4207	NO DELAY

# Introducing the One Board with all the Right Connections



Targeting high-performance I/O and processing applications, the Model 4207 delivers an on-time solution for your real-time application using the latest high-speed serial interfaces.

On board PowerPCs, FPGA, PMC/XMC sites, and extensive memory resources are connected to the industry's fastest interfaces including VXS, Gigabit Ethernet and dual Fibre Channel ports for high-speed recording.

All board resources and interfaces are joined through a configurable zero-latency crossbar switch using protocol-transparent gigabit serial data paths.

With so many connections to choose from, you are bound to find just what you need for your next system.

- Dual PMC/XMC Sites
- Freescale MPC8641D Dual Core Altivec PowerPC Processor to 1.5 GHz
- XC4VFX60 or XC4VFX100 FPGA
- 4 GB DDR2 SDRAM
- PCI Express, Serial RapidIO™, RocketIO™
- Dual Gigabit Ethernet Ports
- Dual 4 Gbit Fibre Channel Interface
- Fabric-transparent Gigabit Crossbar Switch
- VXS, VME64x, Gigabit Ethernet-X
- Linux and VxWorks Support
- ReadyFlow® Board Support Libraries
- GateFlow® FPGA Design Resources



### Call today to have it all!

201-818-5900 or go to  
[www.pentek.com/go/dspconnect](http://www.pentek.com/go/dspconnect) for the  
4207 brochure, technical datasheet  
and to request pricing.

**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](mailto:info@pentek.com) [www.pentek.com](http://www.pentek.com)  
Worldwide Distribution and Support  
Copyright © 2007, Pentek, Inc. Pentek, ReadyFlow and GateFlow are registered trademarks of Pentek, Inc.  
Other trademarks are properties of their respective owners.



25% higher performance

29% lower power

1067-Mbps DDR3 support

# FULL THROTTLE



**Stratix III**

Stratix® III FPGAs, with blazing 1067-Mbps (533-MHz) DDR3 performance, are 25% faster than the competition while consuming 29% less power. With the highest density of any FPGA on the market, Stratix III devices allow you to integrate more functions in the same board space. Add Altera's Quartus® II software to the mix, and you'll complete designs in less time while meeting your power budgets. So, who wants to go faster?

**ALTERA**

[www.altera.com](http://www.altera.com)

