

Military

EMBEDDED SYSTEMS

VOLUME 6 NUMBER 4
JUNE 2010

INCLUDING:

Chris A. Ciuffo
Can bio-science grow in military applications?

Field Intelligence
Multicore packet processors boost performance

Mil Tech Insider
Harnessing multicore processors

Legacy Software Migration
Requirements tools for legacy systems

MIL-EMBEDDED.COM

Special report:
Can mil systems
be hacked?

Also:
Google joins the GIG
How cool: Nano fluids







Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

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

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



Ultimate Modularity

**From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
Input/Output Modules Include:**

**Quad 130 MSPS thru Quad 550 MSPS A/D
1.5 GSps thru 5.0 GSps A/D, Quad 600 MSps D/A,
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10G, 40G or 100G Ethernet or SFPDP**

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

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

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

DTA

D-TA SYSTEMS INC.

SENSOR INTERFACE & PROCESSING COMPANY

10 GIGABIT SENSOR PROCESSING

THE FUTURE IS NOW



DELIVERING 10 GIGABIT NETWORK ATTACHED COMPLETE SOFTWARE RADIO SOLUTIONS

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



DTA-2300 Synchronized Multi-Antenna Software Radio

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

DTA-2210 Single Antenna Software Radio

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



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



DTA-3200L Tunable RF Transceiver

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



DTA-5000L 10 Gigabit Record & Playback System

with 2.4 Terabyte Storage

SENSOR PROCESSORS THAT DRASTICALLY REDUCE DEPLOYMENT TIME AND COSTS

WWW.D-TA.COM

Military

EMBEDDED SYSTEMS

June 2010 Volume 6 Number 4

COLUMNS

Field Intelligence

- 8 **Multicore packet processors boost system and I/O performance**

By Duncan Young

Mil Tech Insider

- 9 **Getting the most out of multicore processors**

By Steve Edwards

Legacy Software Migration

- 12 **Requirements tracing tools can manage legacy systems, too**

By Pete Decher, Mentor Graphics

Crosshairs Editorial

- 46 **Can bio-science grow in military applications?**

By Chris A. Ciuffo

DEPARTMENTS

- 14-15 **Daily Briefing: News Snippets**

By Sharon Schnakenburg-Hess

- 44-45 **Editor's Choice Products**

ON THE COVER:

Germany's famous *Enigma* code cypher machine was "hacked" by British Intelligence during WWII, allowing the Allies unprecedented access to the Third Reich's top secret communiqués. On the modern battlefield, voice, data, and video communications have taken cryptography to new heights. Yet most military systems are vulnerable to hackers and cyber threats because they "touch" the wide-open Internet somewhere, somehow. Also, many embedded systems are designed with unintentional but inherent security vulnerabilities. Our special report begins on page 24. (Image courtesy of: www.flickr.com/photos/zoonaabar/4340886441/)

Published by:  OpenSystems media.

ISSN: Print 1557-3222

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

© 2010 OpenSystems Media © 2010 Military Embedded Systems



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

Software: Google joins the GIG

- 16 **Bringing Android to military communications devices**

By David Kleidermacher, Green Hills Software, Inc.

Hardware: FPGAs do crypto

- 20 **Accelerating cryptography with FPGA clusters**

By David Hulton and David Pellerin, Pico Computing

Technology: Can mil systems be hacked?

- 24 **Raising the bar for security needs: What does "secure boot" really mean?**

By J. Ryan Kenny, CPU Tech

- 28 **On dangerous ground: The rise and fall of military systems power Q&A with Dr. Pankaj Rohatgi, technical director at Cryptography Research, Inc.**

- 32 **Securing wireless Local Area Network interconnections with Layer 2 encryption**

By Juan Asenjo, Thales e-Security

- 36 **Guest opinion: Solid state security – A potential threat**

By Mark Downey, White Electronic Designs/Microsemi

Mil Tech Trends:

Nano fluids keep "your" cool

- 40 **New European program exhorts nano fluids to teach systems how to "be cool"**
Interview with David Mullen, NanoHex project director and mechanical engineer at Thermacore

EVENTS

www.opensystemsmedia.com/events

Autotestcon 2010

Sept. 13-16, 2010 • Orlando, Florida

www.autotestcon.com

ESC Boston

Sept. 20-23, 2010 • Boston, Massachusetts

<http://events.ubm.com/event?eid=441>

WEB RESOURCES

Subscribe to the magazine or E-letter

Live industry news • Submit new products

<http://submit.opensystemsmedia.com>

White papers:

Read: <http://whitepapers.opensystemsmedia.com>

Submit: <http://submit.opensystemsmedia.com>

Mission

Control the sea

Critical

Cutting-edge reliability



Royal Navy Astute
Class nuclear-powered
attack submarine.

Thales' periscope provides
a 360° scan of the surface
above with minimal risk
of detection.

Wind River embedded solutions deliver the breakthrough dependability and performance essential to innovation.

To control the sea, a submarine depends on remaining invisible. But when designing and building a sub, visibility is critical. That's why Thales partnered with Wind River to create a breakthrough in periscope design for the Royal Navy's new Astute-Class submarine.

Relying upon the proven innovation, reliability and performance of our VxWorks RTOS platform, Thales developed a state-of-the-art optronic imaging system that provides stable, high-resolution views in the world's most demanding conditions.

It's the kind of teamwork and support that's made Wind River a trusted leading provider of advanced embedded solutions for aerospace and defense.

To see how Wind River can help you innovate with confidence, download our Mission Critical Toolkit at www.windriver.com/missioncritical/security.

WIND RIVER

© 2010 Wind River Systems, Inc. The Wind River logo is a trademark, and Wind River is a registered trademark of Wind River Systems, Inc. Other marks are the property of their respective owners. Photograph by: Jonathan Massey; © Crown Copyright/MOD, image from www.photos.mod.uk

ADVERTISER INFORMATION

Page	Advertiser/Ad title
34	AdaCore Technologies – Safety, security, reliability
2	Annapolis Micro Systems, Inc. – High performance signal and data processing
27	Apacer – The most reliable storage for industries
12	CPU Tech – Acalis
48	Curtiss-Wright Controls Embedded Computing – Speed your time-to-market
3	D-TA Systems – 10 Gigabit sensor processing
37	Elma Electronic – The industry's choice in VPX handles and panels
30	Elma Electronic – Systems – Embedded storage for secure data
23	Engineering Design Team Inc. – Record/playback system for signals or images
13	Excalibur Systems, Inc. – Dragon
31	Extreme Engineering Solutions – Intel Core i7 processor solutions
47	GE Intelligent Platforms, Inc. – Welcome to the knowledge bank
39	Interface Concept – Switches & IP routers
7	Kontron – We do not build military aircraft
18	Nallatech – High performance FPGA solutions
35	Parvus Corporation – Qualified to perform
43	Phoenix International – RPC12 ruggedized 3U fibre channel RAID system
42	TEWS Technologies LLC – COTS I/O solutions
26	Themis Computer – Themis servers have speed to burn
17	Trident Space & Defense – Protect your data with Trident solid state drives
33	Tri-M Systems Inc. – PC/104 Can-Tainer
41	Tri-M Systems Inc. – 100Mhz PC/104 Module
10	White Electronic Designs/Microsemi – We create space
5	Wind River Aerospace & Defense Division – Mission critical
25	Wolf Industrial Systems Inc. – Aggressive? You bet!
21	Z Microsystems – Rugged rack mounted servers

WWW.MIL-EMBEDDED.COM

OpenSystems media.

Military EMBEDDED SYSTEMS

DSP-FPGA.com

VME
Critical Systems

PC/104 and
small form factors
THE JOURNAL of MODULAR EMBEDDED DESIGN

INDUSTRIAL
EMBEDDED SYSTEMS

CompactPCI
Advanced
& Micro
TCA
SYSTEMS

Embedded COMPUTING
DESIGN

Military & Aerospace Group

Chris Ciufo, Group Editorial Director
cciufo@opensystemsmedia.com

Sharon Schnakenburg-Hess
Assistant Managing Editor
sschnakenburg@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

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

Monique DeVoe, Web Content Editor

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Steph Sweet, Creative Director

Joann Toth, Senior Designer

David Diomede, Art Director

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Sales Group

Dennis Doyle, Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker, Strategic Account Manager
rbarker@opensystemsmedia.com

Christine Long, Digital Content Manager
clong@opensystemsmedia.com

International Sales

Dan Aronovic, Account Manager – Israel
daronovic@opensystemsmedia.com

Sally Hsiao, Account Manager – Asia
sally@aceforum.com.tw

Regional Sales Managers

Barbara Quinlan, Midwest/Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

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

Reprints and PDFs

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

Editorial/Business Office

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

Publishers: John Black, Michael Hopper,
Wayne Kristoff

Vice President Editorial: Rosemary Kristoff

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

Business Manager: Karen Layman



» We Do Not Build Military Aircraft «

Our Customers Do.

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




Radar



Command and Control



Situational Awareness



Satellite Communication



GPS

MILITARY RUGGED COTS

Call, Email or Visit today.

Call: 1-888-294-4558

Email: info@us.kontron.com

Visit: www.kontron.com/military

If it's embedded, it's Kontron.



By Duncan Young

Multicore packet processors boost system and I/O performance



Multicore packet processor devices have found widespread use in commercial information processing and networking applications such as file servers, secure gateways, protocol offload engines, and so on. They are equally applicable to wider DoD and international defense infrastructure projects, but also offer great scope for innovation within deeply embedded military systems. Packet processors with typically 8 to 16 industry standard, general-purpose cores based on either Power Architecture or MIPS64 (Cavium Networks) can run multiple independent or cooperative tasks on a single device. While many network-oriented applications are available off-the-shelf, the wealth of tools and environments for these well established processing cores facilitates additional task development and integration for each new project requirement.

Typical file server application

A typical use for packet processors is to improve an existing installation's service by, for example, replacing an off-the-shelf file server's NIC card. Such servers are often limited by I/O bandwidth and the computing power needed to service the many users and network connections. The processing functions for protocol stacks, payload encryption (for IPv6), and network services and security can all be offloaded to a packet processor to restore the server host processor's performance and responsiveness.

SoCs increase computing density

Embedded military applications in small platforms such as land vehicles, helicopters, combat aircraft, and Unmanned Aerial Vehicles (UAVs) are rarely as conveniently segmented as the commercial file server example. In addition to network performance and security issues, these systems are required to resolve a complex set of real-time problems such as sensor and I/O processing, sensor fusion, target tracking, weapons direction, and platform management while consuming the least possible space, weight, and power. To achieve the next levels of processing density, system

designers are exploring new territory by using one or a number of the many recent types of System-on-Chip (SoC) multicore application accelerators. For example, the latest generation of high-performance Graphics Processor Units (GPUs) has large arrays of processing cores optimized for repetitive, multithreaded video and image processing. Similar in concept but intended for quite different tasks, packet processors can be used to offload tasks such as I/O interfacing, complex decision-making algorithm execution, system security, or network management from a subsystem's host processor.

As a result, packet processing devices primarily designed for high-volume, commercial applications are receiving increasing attention from embedded designers. A device like Cavium Network's OCTEON will typically incorporate 8 to 16 independent MIPS64 processing cores, running at close to GHz clock rates with options for integrated protocol, compression, and encryption engines, plus support for file server and network security applications. The I/O is a mix of 1 GHz and 10 GHz Ethernet ports plus PCI Express ports. Both of these interface types are commonly used by sensor equipment such as radar, sonar, and Electro-Optical (EO) to stream digital sensor video. Also like other SoC-based technologies, packet processors require a host PC-based processor. A Linux SDK package, including ported OS and communications with the host, provides a platform for developing tasks for individual cores as well as the integration of off-the-shelf communication and network packages to provide the complete system alternative.

Packet processors provide network security

Weapons platforms usually comprise a number of sensor and processing subsystems that communicate intensively with each other and externally with other platforms and networks. A packet processor can also be used as a secure gateway for the platform's external communications to provide the perimeter defense

needed to protect the onboard embedded subsystems. In addition to the common IP-related routing and Denial of Service (DoS) attacks, a packet processor has the performance to inspect the header and payload of packets at wire speed to detect and prevent malicious code, viruses, or deliberate data obfuscation. This also applies to collecting statistical data on potentially abnormal network traffic patterns.

For embedded systems, the OCTEON packet processor is available from GE Intelligent Platforms in open architecture formats such as PCI Express or AdvancedMC. The AdvancedMC format offers compatibility with the increasingly popular MicroTCA equipment practice for military applications. Meanwhile, the PCI Express format, depicted in Figure 1, is intended for use in benign applications or the development laboratory.




Figure 1 | The PCI Express format OCTEON-based packet processor

Meeting military needs

The packet processor is an example of how technology can be successfully adopted from other markets and used in innovative ways for deployment in military embedded subsystems. However, because of the nature of military development projects, the technology must be sustainable, with long-term commitments to roadmaps and support before it can be considered. Packet processors meet these criteria and have the performance and capability to step beyond their communications roots and offload many other critical tasks in embedded environments.

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

Getting the most out of multicore processors



By Steve Edwards



During the past 10 years, the embedded market has undergone an SBC evolution, migrating from a single processor to two processors, then to dual-core process chips. More recently, the industry has seen the emergence of multicore System-on-Chip (SoC) devices like the P4080 from Freescale, which is available with eight processing cores. With all of these cores now at their disposal, system integrators are confronted with the challenge of how to optimize all these cores in embedded computing applications.

Saving system slots with multicore processors

One way to exploit these new micro-processor designs is to save board slots by taking applications that today require two or more SBCs and moving them onto a single SBC. For example, with a dual-core processor, a user can choose to treat the processor as two single-processor cores that share I/O and memory. Each individual core is then used to execute what would have otherwise been handled by separate SBCs. This approach can often be realized with a minimum of software code writing effort.

The benefits of SMP

Another option for taking advantage of multiple cores is to collapse the applications that would have otherwise been run on the two independent SBCs, then run them on the two Freescale 8640 cores in a Symmetric Multiprocessing (SMP) mode. This method offers the benefit of increased application efficiency, since the application can be run on either of the two cores at any time. This approach can require much more software coding than the former approach, because the system designer needs to ensure that the applications are aware that they are no longer running on a single processing core.

Interest in SMP is likely to grow as the number of cores increases. This expected growth in popularity is because SMP enables the system designer to concentrate more on the application itself, rather

than on how to split the application(s) across the multiple cores now available. SMP frees the designer from having to program each core individually. But, not surprisingly, there is no free lunch. In an SMP configuration, one of the cores must serve as the *gatekeeper*. As the number of cores (and tasks) increases, the more the gatekeeper core is required to do. The danger is that demands on the gatekeeper can result in performance decreases. In this case, "more" does not necessarily mean "better."

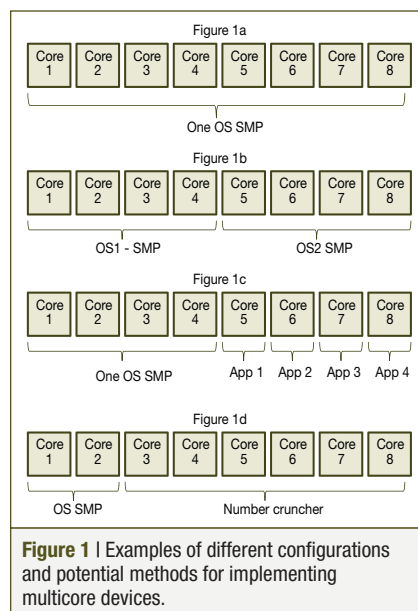


Figure 1 | Examples of different configurations and potential methods for implementing multicore devices.

Implementing multicore devices

In all likelihood, most embedded system designers will likely end up using a combination of these two methods to address their application requirements. Figure 1, depicting four segments (a through d), presents examples of different configurations and potential methods for implementing multicore devices. The examples consider an eight-processing-core device.

Figure 1a illustrates the configuration discussed in the first method described, where all the cores are running under a single OS. This SBC configuration can be very effective if the applications do not require a lot of real-time processing.

With the large number of available cores, one design approach is to split the cores and run two OSs, either the same or similar, across a designated group of the cores. Figure 1b depicts the assigned cores, four per OS. This approach is better at addressing real-time issues than the first example. It also offers the associated advantages of running two different OSs and applications. But depending on the OSs used, there might be restrictions on which data can be accessed by each of them.

Figure 1c illustrates a configuration that allocates a separate core for some specific applications. This approach is well suited for applications that are real-time sensitive. In this type of configuration, for example, the first four cores can be assigned the general-purpose applications, while Core 5 and Core 6 could be dedicated to network stacks, Core 7 could be handling a secure gateway (encryption, decryption, and security protocols), and Core 8 could be responsible for high-speed serial communications.


Figure 1d shows yet another possible configuration. In this case, the first two cores are used in SMP mode and handle all the *housekeeping* functions, such as I/O and so on. The remaining six cores are used as processing engines, perhaps running a DSP algorithm or some other parallel-processing scheme, using the multiple cores to increase processing power and reduce processing time.

Meeting the multicore challenge

System integrators, given access to more and faster cores on a single processor, are just now realizing how best to optimize and harness all of the newly available processing power. Multiple cores can have great benefits and provide impressive results when used properly. But the proper method for leveraging multiple cores in a given application might not be immediately obvious. Trial and error might be needed before the proper configuration is achieved.




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

We create space...



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

Expand the possibilities; visit [**www.whiteedc.com/spacesaving**](http://www.whiteedc.com/spacesaving).



... with more computing power per square inch.

Our SLC NAND flash SSD BGA provides big features in a miniaturized package, freeing up precious board space. It is engineered specifically for the defense market, designed to ensure data reliability in mission critical rugged and mobile systems. Constructed using a 32-bit RISC processor as its core storage controller, this SSD is available in 4, 8 and 16 GByte densities.

Learn more at www.whiteedc.com/ssd.

602.437.1520 TEL | 602.437.9120 FAX

 **Microsemi**
Securing the Future Through the Power of Design
WWW.WHITEEDC.COM

Legacy Software Migration

By Pete Decher



Requirements tracing tools can manage legacy systems, too

In military and aerospace industries, life spans for electronics, computing, and communication systems are long. Hence, these systems often far outlast the teams that develop the original design and implementation. So, the next generation of engineers tasked with maintaining and upgrading these systems is often left to sort through vast quantities of requirements documents, source code, test plans, log files, and other design artifacts. This is surely a mass of information difficult to organize and comprehend, even for the most talented organizations. However, requirements tracing tools are providing a remedy to this ongoing challenge.

When it comes to managing legacy systems, one option, of course, is to stay with the status quo. But the status quo typically amounts to many man hours spent manually slogging through the system's legacy artifacts and hoping nothing important gets overlooked. A more effective approach is to connect these legacy artifacts to an automated requirements tracing tool, however, which can help

show how legacy devices and systems are interconnected. There are four reasons why this second option of utilizing automated requirements tracing tools might make sense.

The new normal: Requirements-driven process

First, using requirements tracing tools is consistent with a subtle but clear shift in the military and aerospace industry's approach. Military contractors today are unlikely to obtain customer or regulatory sign-off on a project if the contractor has merely practiced a minimal, after-the-fact approach to compliance. The new goal is embracing standard requirements-driven processes that encourage more efficiency, reuse, and integration at all project stages.

One example of this shift is the DO-254 standard, formally implemented as policy by the FAA in 2005. DO-254 covers complex electronic hardware for avionics systems, mandating requirements capture and tracking throughout the design and verification process for devices such as FPGAs, PLDs, and ASICs.

The bottom line is that in the years ahead, military and aerospace customers will likely demand increasingly transparent and requirements-driven approaches to all technical projects. For now the focus is on new devices and systems, but it seems inevitable that soon enough, the same will be required of legacy technology. A requirements tracing tool can prepare the industry now for what is to come.

More control of component-level legacy issues

Second, requirements tracing tools can deliver results today, particularly when it comes to managing issues with legacy components. Consider the case of a company that builds electronic components for passenger aircraft. When a 15-year-old component containing a circuit board with a simple PLD controlling the deployment of the evacuation slide needed to be updated, the company had to scramble because the original PLD was no longer available. Its efforts to procure and configure a new PLD would have been easier and less risky if it could have had clear visibility of the relationships of all the legacy documentation, especially the test cases mapping to the specific device requirements. A properly configured requirements tracing tool would have made such visibility possible.

Help in handing off system-level projects

Third, such tools can manage legacy issues at the system level, as well. Consider another example, this one concerning an engineering team working on virtual prototypes for key performance aspects of a new missile system. The team has generated vast numbers of SysML or UML files, and then uses these files to build software-based simulation models to help understand if the missile will fly well in an expected set of conditions (air pressure, thrust, and so on).

- ▶ Advanced, Built-In Tamper Protection
- ▶ Highly Integrated with On-Chip DRAM
- ▶ Fabricated at TAPO / IBM Trusted Foundry

CPU TECH®
Leading Solutions You Can Trust
WWW.CPU TECH.COM

These simulation models, virtual prototypes that can be shared with customers, can become legacy artifacts of their own. This is especially true for large, long-lived platform projects that typically involve developing multiple product variants over the life of the platform. Again, a requirements tracing tool is key in alleviating any future difficulties.

Data adds transparency:

A relief to risk-averse cultures

At one firm, a technical manager on the maintenance side of a bifurcated engineering division mentioned that his group should not even accept projects from the new product development team unless all the relevant documentation, test files, and code are loaded in a requirements tracing tool.

How hard this technical manager pushes for requirements tracing tool adoption depends on the culture of his employer. In general, change is slow in the military and aerospace sectors, which are relatively top down and risk averse. So ... the fourth and possibly most important benefit to utilizing requirements tracing tools is this: Properly connected to project data, such requirements tracing tools, including Mentor Graphics' ReqTracer, for example, can be used as real-time project dashboards. These tools enable all project stakeholders to better understand dependencies among system artifacts – and the risks involved in making needed changes and upgrades.

Ultimately, manual processes mean more risk

In short, yes, it is always challenging to propose adopting a new tool. But the bigger potential time and resource sink is the effort

“ The bottom line is that in the years ahead, military and aerospace customers will likely demand increasingly transparent and requirements-driven approaches to all technical projects. ”

and associated risk of trying to manually trace documentation and design artifacts written by long-retired colleagues.

Pete Decher is the program manager for requirements tools in the Design Creation and Synthesis Division at Mentor Graphics. He has more than 30 years of experience in EDA and electronic system design and test. Pete holds a B.S. in Electrical Engineering from Georgia Tech and an M.S.E.E. from Stanford University. He can be contacted at pete_decher@mentor.com.

EXCALIBUR
EXCALIBUR SYSTEMS

DRAGON

- Ruggedized enclosure system for PC/104
- Multiple MIL-STD-1553 and/or ARINC-429
- User configurable • Environmentally tested
- Expandable • Accepts third party cards

www.mil-1553.com

CAMELOT DRAGON LANC

NSI AS9000 Registered ISO 9001 Registered UL UKAS

Daily Briefing:

By Sharon Schnakenburg-Hess, Assistant Managing Editor

News Snippets

www.mil-embedded.com/dailybriefing

New RTOS and tool chain tell Joint Strike Missile what to do

Hardware is 100 percent dormant without software giving it life and direction. That goes for all mil-embedded systems, including the Joint Strike Missile (JSM) presently in development for Lockheed Martin's F-35 Joint Strike Fighter (Figure 1). JSM is headed by Norway-based prime Kongsberg Defence Systems, which recently chose Green Hills Software's INTEGRITY RTOS, along with GH's Probe networking middleware and debugger and MULTI integrated development environment, for the program. The triad of wares is slated for utilization as JSM's planning, safety launch, and telemetric software. Specifically, INTEGRITY will power several Freescale Power Architecture processor-based multicore computers within JSM; meanwhile, JSM is to be mounted internally or externally to the F-35's bomb bay. Incarnated for the Royal Norwegian Air Force, JSM will work in naval fire support missions and anti-surface warfare over land and sea.



Figure 1 | An F-35 Lightning II Joint Strike Fighter test aircraft flies over Eglin Air Force Base in Florida: the future home of the JSF training facility. U.S. Air Force photo by Senior Airman Julianne Showalter

U.S. Special Ops makes something known

While many things are unknown about the inner workings of U.S. Special Operations Command, one thing that *is* known is its recent penning of a now-not-to-exceed \$464 million contract modification with Harris Co. The new ceiling provided for the originally \$422 million, 10-year IDIQ contract grants Harris a five-year base timeframe plus a five-year option in sustaining and procuring enhanced-capability special ops high-frequency manpack radio systems. The contract's work continues in Rochester, New York; meanwhile, the contract is slated for completion on April 30, 2012, or April 30, 2017 if the second five-year option is activated.

Navy's C2 and C4ISR to work better together

Long on nomenclature, Command, Control, Communications, Computer, Intelligence, Surveillance, and Reconnaissance (C4ISR) systems are paramount to winning any military conflict. Case in point: a recent \$73 million contract between the U.S. Navy and Eagan, McAllister Associates, Inc., wherein the latter provides the former with tactical Command and Control (C2) integration services for C4ISR systems. Apparently a high-value proposition, the IDIQ contract's options could boost the contract's cumulative total to \$533 million. Work completion is anticipated by April 2011 – or April 2013 if all options are executed, with Charleston, South Carolina as the primary contract fulfillment locale. Work will also be performed at the company's Norfolk, Virginia and Lexington Park, Maryland locations. The contracting activity is the Space and Naval Warfare Systems Center Atlantic, also situated in Charleston, South Carolina.

Raytheon to keep PATRIOT on course

The venerable PATRIOT [Phased Array Tracking Intercept of Target] missile system needs no introduction these days – but it does still need some engineering services, per a recent contract between Raytheon Company and the U.S. Army. The \$105 million cost-plus-fixed-fee contract stipulates that Raytheon provides a precise 907,043 man-hours of engineering services. The PATRIOT (see Figure 2) contract is anticipated to run through January 31, 2014, and work takes place at Tewksbury, Andover, and Burlington, Massachusetts, in addition to Huntsville, Alabama and El Paso, Texas. The contracting activity is the Aviation & Missile Command Contracting Center in Huntsville, Alabama.



Figure 2 | MIM-104 PATRIOT surface-to-air missile system, U.S. Air Force photo by Staff Sgt. Robert Barney

Linux steers U.S. Navy ships

Open architectures are making waves when compared against old-school proprietary wares. And Global Technical Systems (GTS) recently choosing SteelEye Technology Inc.'s Linux-based software to climb aboard U.S. Navy vessels exemplifies this trend (Figure 3). Specifically, SteelEye Technology will provide its SteelEye Protection Suite for Linux Multi-Site Cluster Edition (SLMSCE) software to GTS, which heads the U.S. Navy's Common Processing (CPS) program in partnership with Northrop Grumman. CPS renders a processing system as part of the Navy's anticipated Open Architecture (OA) initiative for its combat systems. CPS is slated for development centering around commercial software and hardware, including I/O interfaces, memory, data storage, and computer processing in support of software apps tucked inside combat systems powered by Red Hat Linux. The goal: to assure that the entire Navy fleet's host software applications are continuously available. Key to the equation are SLMSCE's cascading multiple node failover capabilities and multi-site cluster configurations, which will enable it to safeguard the Navy's IBM BladeCenter BCHT infrastructure amidst unplanned and planned network downtime.



Figure 3 | The USS Dwight D. Eisenhower (CVN 69) Nimitz-class aircraft carrier and the USS Farragut (DDG 99) guided-missile destroyer – both part of the Eisenhower Carrier Strike Group – pull along both sides of the USNS Supply (T-AOE 6) Military Sealift Command fast combat support ship. U.S. Navy photo by Naval Air Crewman 3rd Class Ruben N. Coss

Troops get no interference from mini actuator

Keeping a military submarine's or ship's navigation systems interference-free is paramount to mission success and troop safety. Accordingly, a recent Phase I Small Business Innovation Research (SBIR) contract between the U.S. Navy and New Scale Technologies, Inc. for a non-inductive rotary actuator system illustrates this philosophy. The contract stipulates that New Scale Technologies provides its tiny 1"-cubed closed-loop actuator system to execute flight control surface movement within mini precision-guided munitions. And because it's non-magnetic, the actuator can be positioned close to larger navigation systems utilizing the Earth's magnetic field to gather roll rate and orientation information. The small-sized actuator comprises control electronics such as a microprocessor, position sensor, and drive ICs, in addition to several piezoelectric micro motors. The piezoelectric motors serve to rotate a shaft, without mechanical linkages or gears, at 300 degrees per second with 0.2 N-m torque.

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



Figure 4 | An M1 Abrams tank, U.S. Army photo by Corporal Lee Sang-Jun

Abrams tank gets an overhaul

Though the Abrams tank is easily recognized by its external appearance, it's indeed what's inside that counts more – at least as far as a recent contract between the U.S. Army and Honeywell International, Inc. is concerned. The \$93 million firm-fixed-price contract is for the Total Integrated Engine Revitalization program's year 5, which specifies that Honeywell proffers support and parts for overhauling 1,500 engines or equivalents in addition to 1,000 automotive gas turbines for Abrams tanks and Abrams derivatives (Figure 4), as well as Army stock spares. The majority of the contract will be completed in Phoenix, Arizona, with work also transpiring in Rocky Mount, North Carolina; Greer, South Carolina; and Anniston, Alabama. The anticipated contract completion date is December 31, 2011.

USAF gets more secure – 9 times over

Clearly, security is a major concern for the DoD these days. The evidence: The recent origination (and all in one day) of nine separate Information Assurance (IA) contracts between the United States Air Force and Booz Allen & Hamilton, Inc.: 1) a \$24 million contract to arm combat-ready forces with secure cyber operations; 2) another \$24 million to render IA recommendations to the Systems Center Atlantic, with Offutt Air Force Base, Nebraska as the contracting activity; 3) a \$23 million contract to increase IA for U.S. Space Command's cyber activities; 4) \$19 million to insert IA into present-day command and control systems and networks; 5) \$19 million for IA technical analysis for future and emerging satellite comms systems and also secure comms to warfighters in the field; 6) \$15 million for secure, high-rel networks for Air Combat Command; 7) \$14 million for IA for comms systems on the ground and in the sky; 8) \$14 million for IA capabilities to enhance secure information's availability and interoperability; and, finally, 9) \$8 million to develop network defense and cyber security for Air Force information systems (Figure 5).



Figure 5 | The USAF is beefing up its cyber security per nine recent contracts with Booz Allen & Hamilton. U.S. Air Force photo by Captain Carrie Kessler

Bringing Android to military communications devices

By David Kleidermacher

With the advent of multicore and hypervisor-enabled mobile SoCs, military mobile devices such as Type-1 PDAs and Software-Defined Radio (SDR) can now keep pace with the latest commercial mobile software initiatives such as Google Android while reducing costs and meeting the most stringent security requirements. In addition to the hardware advances, this goal is made possible with software technologies and architectures that include a layered approach to virtualization using a Multiple Independent Levels of Security (MILS) separation kernel and its high-assurance partitioning policies.

Red-black separation

Communications devices that process classified data are typically architected with a standard red-black separation in which a red-side processor is responsible for cryptographic functions, and the black-side processor is responsible for communication stacks and drivers. On egress, classified information originated in the red side is encrypted and sent over some interconnect to the black side for transmission. On ingress, information received by black-side drivers is passed across the interconnect for decryption and any other red-side processing, such as guards and authentication. This general architecture is shown in Figure 1.

Note that the logical red-side crypto component may include an attached special-purpose hardware device, such as an FPGA, for cryptographic algorithm execution, key storage, and physical redundancy.

Any software running within the cryptographic boundary falls under the intense scrutiny of the National Security Agency (NSA) during Type-1 certification. This

naturally imposes a strict assurance demand and complexity limit upon the red-side software.

Software-Defined Radio

Now let's take a look at the red-black architecture in more detail, using the example of a Software-Defined Radio (SDR). More specifically, let's consider devices that employ the Software Communications Architecture (SCA), an open framework used to develop SDRs. The SCA provides standardized APIs for managing waveforms and other radio-relevant resources. The SCA lives on top of an SCA-compliant operating system. In an SDR, the SCA may be used in both the red and black sides. The black side will typically include components required to manage radio communications, including the waveforms themselves. Therefore, a Real-Time Operating System (RTOS) is often essential.

The red side may use less of the SCA's functionality. At a minimum, SCA-standardized message passing and component management functions would be included. Of critical note, however,

is the role of the red side in managing plaintext. Human-Computer Interface (HCI) components, such as keyboard drivers, voice codecs, and touch-screen management software, will be included in the red-side processing. For example, a handheld radio operator may speak voice data that must be collected by the red side for cryptographic processing before it can be sent across to the black side for transmission.

As with any Type-1 communications device, minimizing the certification-relevant software content in the red side is a strategic goal of both SDR developers as well as the certifying agency. Because of its real-time and security-critical requirements, a trusted RTOS is a natural fit. An example architecture showing the major red- and black-side SDR components is shown in Figure 2.

Enter Android

While an analog voice radio has a relatively simple HCI, smartphones and multipurpose digital voice/data radios are but two examples of devices that may incorporate a far more sophisticated HCI.

U.S. Air Force photo by Staff Sgt. Jacob N. Bailey

Android has quickly become one of the most popular HCI frameworks used in consumer electronics. Android, however, is built upon Linux, a monolithic operating system that does not meet high-assurance certification requirements and is not appropriate for low-latency, hard real-time tasks.

To incorporate Android into the red-side HCI of military communications devices without bringing its entire multimillion-line source code base into the cryptographic boundary, an obvious choice is to incorporate a second applications processor dedicated to Android. Adding a second processor, however, increases footprint, power, production cost, and system complexity. This is especially problematic in resource-constrained devices, such as battery-powered radios.

When they have been incorporated using secondary red-side processors, general-purpose operating systems are typically only used for unclassified communications. Classified communications require

custom HCI interfaces to achieve high assurance of the integrity and availability of sensitive data and commands.

The Holy Grail, of course, is to incorporate Android into the original red processor without sacrificing real-time performance, jeopardizing security, or increasing certification overhead. Furthermore, we want to use Android for classified as well as unclassified interfaces, enabling the warfighter to fully realize the rich productivity benefits of Android. This goal is made possible using a specialized form of system virtualization, called Multiple Independent Levels of Security (MILS) virtualization. MILS virtualization requires a MILS separation kernel.

MILS separation kernel

MILS is a security architecture based on the concepts of high-assurance components and the strict isolation and controlled information flow between those components. At the lowest level, MILS policies are enforced by a separation kernel, a specialized RTOS designed

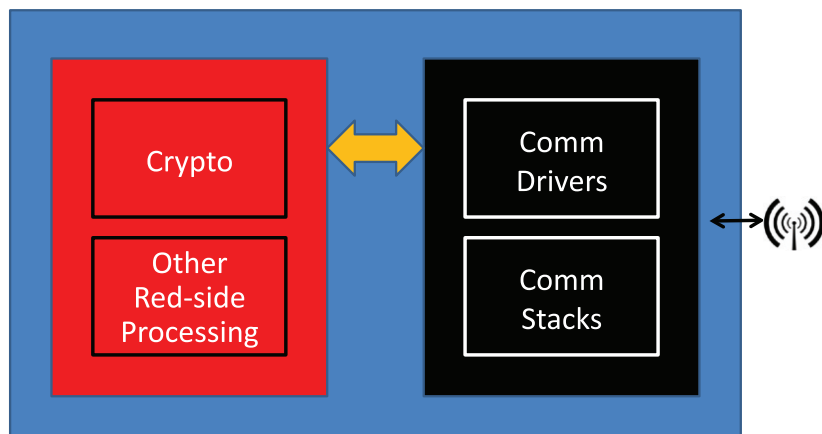


Figure 1 | Communications device red-black architecture

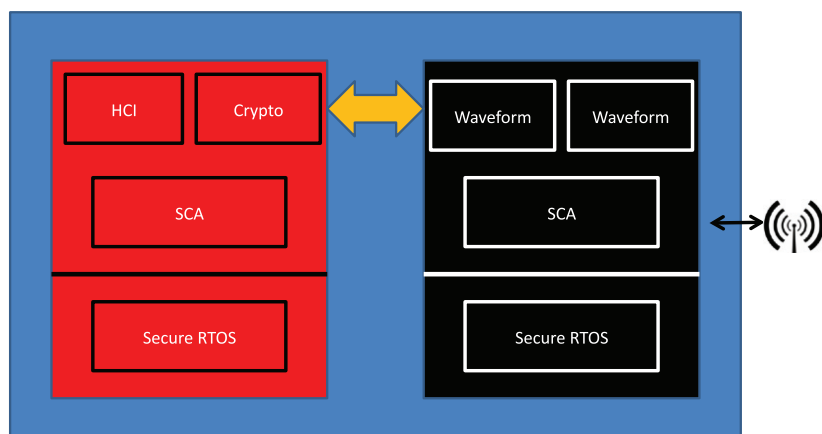


Figure 2 | Example SDR red-black architecture

**Protect your data with
Trident Solid State Drives**
Extremely Rugged for an Extreme World




**Up to 128GB Capacity with
Fast / Secure Erase**

Military Grade Solid State Drives:

- Fast / Secure Erase for complete media de-classification
- Purpose built for military applications utilizing Industrial Grade SLC NAND Flash (-40 to 85°C)
- Standard 2.5" Form Factor with anodized aluminum case
- Standard IDE or SATA interface with sustained speeds up to 100MB/s
- BGADrive available for direct placement and reflow onto PCB
- Available with custom options
- Tested to Mil-Std 810F

Trident BGADrive



Further Details at:
www.tridentsd.com



Phone: 310-214-5500
Email: sales@tridentsd.com

to meet the highest levels of assurance while providing a powerful environment for hosting both general-purpose and security-critical applications.

Assessment of software-security assurance is performed using the international standard for information technology security product evaluation: Common Criteria (ISO/IEC 15408). Common Criteria Evaluated Assurance Levels (EAL) range from 1 to 7. Most general-purpose products, such as Windows, Linux, VMware, Web servers, firewalls,

and so on are certified to level 4 or lower. Level 6+ corresponds to high assurance. In 2008, Green Hills Software's INTEGRITY separation kernel became the first software technology to achieve a high-assurance Common Criteria certification. Assurance must-haves at this level include numerous rigorous development process controls, formal mathematical proof of security policy, and NSA penetration testing with full access to source code. The same RTOS technology is widely used in NSA Type-1 certified communications devices.

EAL6+ represents the U.S. government's mapping of "high robustness" to the Common Criteria. High robustness is the level of security recommended when a communications device is managing high-value information in a high-threat environment. If either the information value or the threat environment is low, then a medium robustness (EAL 4) solution may be sufficient.

MILS virtualization

As described earlier, the power of a MILS separation kernel lies in its ability to enable coexistence of security-critical applications with general-purpose applications. In some cases, these general-purpose subsystems are full "guest" operating systems, running in a virtual machine under control of the MILS separation kernel. Unlike traditional hypervisors, the separation kernel can host native applications as well as guests. The separation kernel's strict resource scheduling and protection mechanisms ensure that the virtual machine and its constituent applications are unable to impact the execution of critical applications.

The separation kernel is the only software that runs in the processor's most privileged mode. The mechanism of system virtualization depends on the processor's specific hardware capabilities. Increasingly, modern embedded processors are incorporating hardware virtualization acceleration, which enables virtual machine management to be as simple and efficient as possible.

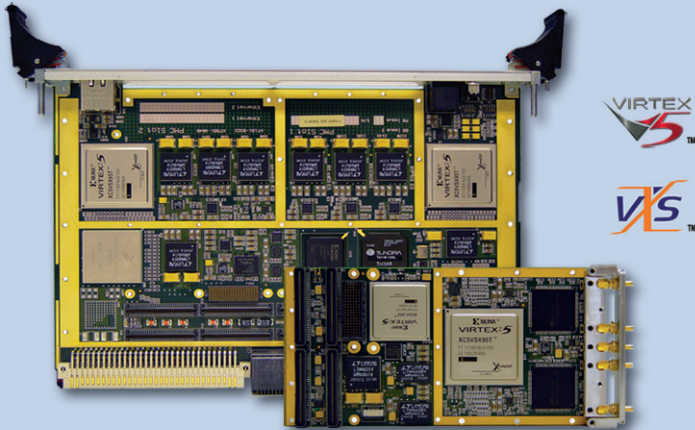
The separation kernel can provide a strictly controlled Interprocess Communication (IPC) path between the Android HCI and other red-side applications as needed. For example, Internet Protocol data originating from the Android network subsystem can be transmitted over the IPC pipe to the crypto subsystem for encryption and transmission. Application of the MILS virtualization concept to the SDR example is shown in Figure 3.

The difficulty of using Android for managing classified communications, however, remains. For example, consider the warfighter using an Android touch-screen widget to enter the sensitive command, "Zeroize all cryptographic keys now." Mission success may depend on the command information passing correctly through Android to the crypto component.


HIGH PERFORMANCE FPGA SOLUTIONS

RUGGED EMBEDDED VXS & XMC SOLUTIONS


- PowerPC & Dual FPGA VXS compute cards
- Analog & Digital I/O and FPGA processing XMC mezzanines
- Processing and I/O for SigInt Radar and SDR applications
- Rugged solutions to support multiple target applications




PCI EXPRESS AND
PCI-104 SOLUTIONS




INTEL FRONT SIDE BUS (FSB)
FPGA MODULES



FPGA MINIATURIZED
MODULES



Nallatech Inc.
Toll Free: 1-877-44-NALLA
contact@nallatech.com
www.nallatech.com



Nallatech

a subsidiary of
Interconnect Systems Inc.

© 2009 Nallatech Inc. All Rights reserved. All trademarks or registered trademarks are the property of their respective owners.

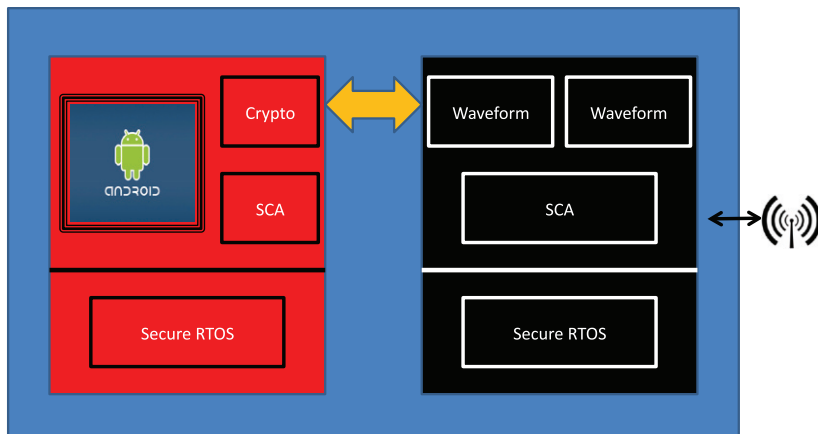


Figure 3 | Mils virtualization SDR architecture

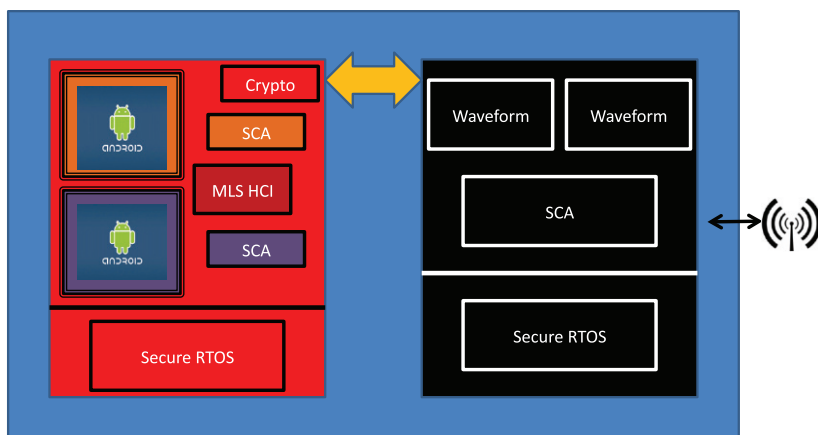


Figure 4 | Multidomain Mils virtualization

Yet the low assurance of Android makes it impossible to make the necessary integrity and availability guarantees.

Again, MILS virtualization provides a solution. In the MILS virtualization architecture, the physical graphics device is controlled exclusively by a trusted application running on the MILS separation kernel; Android has access only to virtualized devices. Thus, when a command is entered through Android and across the virtualized interface, a trusted application can check the command. The command is only committed if the warfighter is satisfied that Android has faithfully transmitted his/her intent. This verification stage provides a MILS-enforced trusted path from warfighter to high-assurance components; Android is completely out of the loop.

Multi-domain MILS virtualization

Some military communications devices, such as radios, smartphones, and network interface cards, must manage information

at varying classification levels. Default policy requires that information at different levels be kept physically isolated. Typically, this is accomplished in one of two ways. One method is to require a cold restart, in which writable hardware resources are zeroized, to safely switch the device between security levels. This approach is not only unfriendly to users, but the reboot could delay communications and impact mission effectiveness. The second approach is to incorporate a discrete red-side processor for each security level and only require a secure switch of the human interface devices (for example, a display or keyboard). Once again, the extra processing elements will increase the footprint, production cost, and system complexity of the device.

MILS virtualization solves the multi-domain problem as well. Separate instances of any required subsystems, including virtual machines and SCA frameworks, can be strictly isolated and scheduled by the MILS separation kernel. In addition, the

separation kernel's ability to host native applications can solve the shared human interface device problem: A trusted Multi-Level Secure (MLS) HCI application can run directly on the separation kernel, multiplexing multidomain I/O based on user-derived focus input captured directly by the MLS component. The multi-domain MILS virtualization architecture for an SDR is shown in Figure 4.

While the architectures described in figures 3 and 4 can be implemented using a single-core processor, the advent of multicore embedded processors can make them more practical. As shown by these examples, the sophisticated red-side processing includes numerous components, many of which can execute concurrently on multiple cores, under supervision of a multicore-aware separation kernel. This extra horsepower enables good performance for virtual machines while ensuring real-time behavior for critical applications.

Military 'Droid

The rich functionality of the latest multimedia software packages, such as Android, can now be incorporated into even the most demanding real-time, secure military communications devices without increasing hardware footprint, cost, or certification burden. The key innovation is MILS virtualization, which exploits the resource management capabilities, native applications environment, and assurance pedigree of a trusted separation kernel and modern hypervisor techniques to effectively consolidate general-purpose and critical subsystems. ✚



David Kleidermacher is Chief Technology Officer at Green Hills Software, where he is responsible for technology strategy, platform planning, solutions design, and technical evangelism. He is a leading authority in systems software and security, including secure operating systems and virtualization technology. David earned his Bachelor of Science in Computer Science from Cornell University and is an active writer and speaker on technology subjects. He has been with Green Hills Software since 1991. He can be contacted at davek@ghs.com.

Green Hills Software, Inc.
805-965-6044
www.ghs.com

Accelerating cryptography with FPGA clusters

By David Hulton and David Pellerin

There is a quiet, international battle underway, a battle that impacts every data consumer and producer. On one side of this battle are the cryptographers who work to protect our national security and the privacy of our personal information through increasingly strong methods of encryption, and through penetration testing. On the other side are the hackers, criminals, and unfriendly governments who use increasingly sophisticated code cracking methods to access seemingly secure data. The weapons in this war include increasingly sophisticated, massively parallel computers that are built using FPGA devices.

Cryptography – which includes code creation, analysis, and breaking – has been one of the most important drivers of computing technology for more than six decades. In fact, the development of modern computers owes much to cryptography research performed during and immediately after World War II, by computing pioneers such as Alan Turing and Claude Shannon.

In the modern era of ubiquitous wired and wireless communications, cryptography research is once again leading to dramatic advances in computing. As more of our personal, commercial, military, and national security data are transmitted and managed online, opportunities for mischief have become increasingly available, and stronger methods of encryption are therefore critical. Research into improved encryption requires analyzing and auditing systems in widespread use today, and systems of the future.

The following discussion focuses on how cryptographic algorithms, and code cracking in particular, can be accelerated using clusters of widely available Field Programmable Gate Arrays (FPGAs). We first explain the challenges behind cryptographic computing, then present a case study on how FPGA-accelerated methods can speed the decryption of widely used data security methods including DES, WPA, WEP, and GSM and we compare the FPGA-based approach to more traditional, software-oriented methods. And finally, we demonstrate that a cluster of commodity FPGAs, contained in a single 4U chassis and consuming less than 1,400 W, provides the computational equivalent of more than 1,000 dual-core processors.

White hats versus black hats

Code cracking is a controversial activity, with significant legal perils for those who pursue it to access otherwise off-limits

information. Likewise, there are significant risks of financial loss or sensitive data breaches for those who are the targets of such intrusions. All digital data users have cause to be worried if the systems they rely on are vulnerable to unauthorized access. For this reason, many organizations and governments make use of cryptography experts for security auditing. These professional hackers engage in legitimate attempts to break into systems, with permission of their owners, testing commercial and public data network security.

When auditing a presumably secure system, a "white hat" hacker will deploy much the same software algorithms and hardware as their "black hat" counterparts. This activity, which is referred to as *penetration testing*, is aimed at finding and reporting vulnerabilities in the data security strategy. Are inadequate or obsolete encryption methods being used? Are weak passwords commonly chosen by system users? Are there operating system or application-layer vulnerabilities?

Another legitimate reason for such cracking is to recover lost and forgotten passwords. Data locked up and unavailable because of password mistakes results in enormous amounts of lost productivity each year. Confidential documents and entire databases can be rendered useless, even to their rightful owners. Password recovery software tools based on black hat cracking methods are therefore used daily in government, industry, and academia to regain access to critical data.

The challenge of cryptographic computing

The algorithms that form the basis for these security cracking and password recovery tools require large amounts of computing power and can require very long runtimes to complete. In fact, while widely available software tools running on a single computer can crack a dictionary-based password in minutes, recovering a relatively strong password might require years of iterative computations.

In recent years, computer clusters have been augmented by the addition of GPU and FPGA accelerators. These devices, by virtue of their parallel structures, can provide substantial acceleration for many computer-intensive algorithms including password recovery.

Because of their roots in graphics processing for standard, off-the-shelf computer systems, GPUs are well known to most programmers. However, although they can provide a significant performance boost for code cracking algorithms[1], GPUs are not well optimized for this application. In particular, bitwise encryption algorithms such as the Digital Encryption Standard (DES) and SAFER (as used in Bluetooth) do not lend themselves to efficient implementations on either CPUs or GPUs. GPUs also need large amounts of power relative to the computations being performed. GPUs provide a high level of acceleration, but at a significant cost in power consumption. Instead, FPGAs offer a viable alternative.



RUGGED RACK MOUNTED SERVERS

from Z Microsystems



Built to thrive in harsh environments

- Choice of 1U, 2U, and 3U systems
- Incorporates latest Intel or AMD CPUs
- Supports up to 6 removable hard drives
- 20" deep to save valuable space



MISSION-READY

For more information, visit www.zmicro.com/zx or call 858.831.7054.

Case study: FPGAs for high-performance password cracking

To demonstrate how FPGAs can address the need for faster password recovery, we selected a representative set of password-cracking problems and deployed highly parallel recovery algorithms on FPGA clusters similar to the one shown in Figure 1.




Figure 1 | FPGA cluster populated with Xilinx Spartan FPGAs

The systems used in this demonstration included a Pico SC3 cluster containing 77 Xilinx Virtex-5 FPGA devices, an SC4 cluster containing as many as 176 Spartan-3 FPGAs, and an SC5 cluster containing 72 larger-capacity Spartan-6 FPGAs. In all cases, these clusters were housed in a single 4U chassis drawing under 1,400 W. Using FPGA hardware design methods and tools, we successfully deployed recovery algorithms for the following encryption methods:

FileVault

FileVault is based on the Advanced Encryption Standard (AES) and provides an encrypted file system for the Apple Macintosh operating system. Recovering FileVault passwords requires hashing a possible password with the SHA-1 hash function thousands of times. This derivation method, known as a *Password-Based Key Derivation Function*, or *PBKDF2*, is used iteratively in an attempt to decrypt the FileVault image. Using the 72-FPGA SC5 cluster, we were able to speed the FileVault key recovery application by a factor of 498 times when compared to the original software implementation running on an Intel Core i7 processor at 2.93 GHz. The result was a reduction in runtime from 21 hours to just 2.5 minutes (see Figure 2).

Wi-Fi Protected Access (WPA)

WPA is used to secure wireless networks and is implemented via a temporal key integrity protocol. WPA and WPA2 are intended to replace the previous-generation Wired Equivalent Privacy (WEP) protocol. WPA recovery is characterized by performing a PBKDF2 SHA-1 algorithm to the candidate passwords for the network and then comparing authentication hashes to determine if the password is correct. When run on our FPGA-based cluster, this algorithm showed acceleration of 498 times as compared to the Intel Core i7 processor software runtimes.

FileVault keys/sec




Figure 2 | FileVault cracking is accelerated 498x over the CPU, using an FPGA cluster

WEP[2] (or *Wired Equivalent Privacy*, as mentioned previously) was first introduced in 1997, and is still commonly used in homes and businesses to protect wireless access points from intrusion. WEP is significantly less secure than WPA; hence for organizations requiring high levels of security, including government and military, WEP is now being replaced by WPA and other more secure systems. Breaking WEP involves trying every possible key formed from a pseudo-random stream of bits that has been combined (using an exclusive-OR operation) with known plaintext and checksums present in each WEP packet. Statistical approaches can speed up this attack by reducing the size of the key space and the corresponding number of keys to be tested. When cracking WEP, we brought the time needed for brute-forcing the entire 40-bit keyspace from an estimated 42 days on the IntelCore i7 processor down to just 4.7 minutes on our FPGA cluster. This performance comes from generating 3.8 billion keys per second in the FPGA cluster, representing an astonishing 12,933x increase in performance as illustrated in Figure 3.

WEP keys/sec




Figure 3 | WEP cracking is accelerated 12,933x over the CPU, using an FPGA cluster

The vulnerability of digital walls

FPGA-based methods can be used to crack many data encryption schemes that once appeared to be strong. Using a single FPGA cluster equipped with 176 FPGA devices, we recently achieved the highest-known benchmark speeds for 56-bit DES decryption

using a single, FPGA-accelerated 4U server, with throughput exceeding 280 billion keys per second.

The parallel DES cracking algorithm in this demonstration used brute-force methods to analyze the entire DES 56-bit keyspace, iteratively decrypting fixed-size blocks of data to find keys that decrypted some portion of the data into sequences of ASCII numbers. This technique is often used for recovering the keys of encrypted files containing known types of data, for example encrypted documents that include financial or military information. DES cracking software running on current-generation CPU cores can process approximately 16 million DES key operations per second. A modern GPU card such as the NVIDIA Tesla can handle more than 10x that number, or approximately 250 million DES operations per second.

When using an FPGA cluster based on a single, off-the-shelf motherboard, however, each FPGA was able to perform 1.6 billion DES operations per second. This means that a key recovery that would take years to perform on a PC, even with GPU acceleration, could be accomplished in less than three days on a 4U FPGA cluster. As additional FPGAs are deployed in ever-larger clusters, the time required to decrypt data decreases in a near-linear manner.

Although DES is now considered obsolete because of its demonstrated vulnerability to attack, there are other encryption methods in common use today that are also at risk. Consider Global System for Mobile Communications (GSM), a standard that carries the overwhelming majority of cellular phone calls worldwide. At the Chaos Communication Congress held in Berlin in December of 2009, cryptographic researcher Karsten Nohl announced that his team, a group of hackers working collaboratively to create a distributed computing cluster, had cracked GSM encryption by creating an enormous, 2 TB "rainbow table" of hash values[3]. In simplistic terms, the rainbow table provides a cracking program with a reverse-lookup scheme that can quickly decrypt the wireless voice data.

During his talk, Nohl stated that a person or group wanting to eavesdrop on GSM calls would currently need to spend around \$100,000 on hardware to crack an A5/1 encrypted call in one second or less. This research illustrates the importance of security auditing to expose and correct vulnerabilities in widely used data security standards. FPGA accelerators can play an important role in such auditing.

Cracking the code faster, thanks to FPGAs

"Black hat" and "white hat" hackers alike have discovered the power of FPGA computing for cryptography. As in the earliest years of computing, researchers in other application domains including such diverse areas as genomics, physics, and finance are also realizing the benefits of FPGA-based parallel platforms.

The most computationally intensive parts of code cracking and key recovery algorithms are relatively simple operations that must be performed iteratively, and that can be parallelized for high performance. When deployed on FPGAs, these algorithms can use available FPGA resources with extreme efficiency because of the simple nature of the tasks being performed. Not all software algorithms fit so neatly into FPGA clusters, but many complex computing problems can be attacked in the same way. ✦

References:

- [1] "Distributed Password Recovery: How GPU Acceleration Works," www.elcomsoft.com/distributed_password_recovery.html
- [2] "WEP key wireless cracking made easy," John Leyden, *The Register*, April 4, 2007, www.theregister.co.uk/2007/04/04/wireless_code_cracking/
- [3] "Security fear as mobile phone code is cracked," Maija Palmer, *Financial Times*, January 2, 2010, www.ft.com/cms/s/0/ee35fcdc-f70c-11de-9fb5-00144feab49a.html



David Hulton is Director of Security Applications for Pico Computing and an acknowledged expert on FPGA-based code cracking. He is a past chairman of the ToorCon security conference and has spoken at events that include Black Hat, DEF CON, and many others. He can be contacted at dhulton@picocomputing.com.

David Pellerin, Director of Strategic Marketing for Pico Computing, has more than 25 years of experience with programmable logic. He is the author of four books on FPGA-related topics, most recently "Practical FPGA Programming in C." He can be contacted at dpellerin@picocomputing.com.



Pico Computing
206-283-2178
www.picocomputing.com

Record/playback system for signals or images



- Up to 10 Gb/sec
- 16 TB of storage
- Analog (L-band or IF)
- Video (Camera Link or ASI/SMPTE)
- Optical or electrical (up to 10 Gb/sec)
- Recorder software and utilities included



www.edt.com | 800-435-4320

Raising the bar for security needs: What does “secure boot” really mean?

By J. Ryan Kenny

Embedded boot code security is an important area of vulnerability analysis being investigated by technology providers. After adding a digital signature or authentication step, however, marketing immediately labels the solution “secure boot.” It is time to examine what “secure boot” really means, and how to grade secure boot technologies against the security risks in end systems. Key considerations include what is being protected and how, secure boot capabilities available, matching requirements to capabilities and claims, and open standards efforts.

As “Cyber Security Awareness” becomes a top White House priority in 2010, it is becoming apparent that most of today’s awareness efforts are focused on Internet systems. Very little outreach focuses on the security vulnerabilities of embedded systems, unless such a security breach makes the news.

The number one vulnerability source in embedded processors is the initial boot phase of the device. This is an assertion made by the author, but the truth of this statement can be independently assessed through a search of academic papers on embedded security, discussions at security conferences, and feedback from equipment manufacturers who can attest to the ways in which their products are tampered with and cloned in the marketplace.

The volume of these published and unpublished security incidents in the past few years has driven the software industry to focus on the security of micro-processors’ and embedded processors’ boot process. The reason for the analysis

presented herein is straightforward: By assessing how much security is really designed into the system with any secure boot solution, designers can make a more informed decision on how much actual end-system security they are getting for their R&D and procurement dollars.

What are you protecting – and how?

Two key questions should be posed by the architect responsible for system security (Figure 1). The first question in assessing the level of security in secure boot technology is whether the capability is actually preventing the exposure of the boot code

and software IP of the embedded device or system. Can the most sophisticated antagonists who have access to a boot device or process access the code, then use it for reverse-engineering purposes?

If so, the hacker will then have all the information needed to clone the system, insert malware, develop countermeasures, or utilize other methods of disabling the system based on a vulnerability or flaw. Secure boot technology that does not protect the code itself offers no real security against the most determined attackers, and might reveal critical intellectual property.

Questions to analyze secure boot requirements:

1. Do you need to protect boot code IP itself?
2. What is the threat being addressed by secure boot?

Figure 1 | Two key questions for secure boot requirements

Senior Airman Kyle Stackman (front) and Tech. Sgt. James Thoni with the Information Protection Operations office at Nellis Air Force Base, Nev., monitor the base's computer network to keep it secure. U.S. Air Force photo/Master Sgt. Jack Braden

The next question to ask about secure boot technology is: Which kinds of threats does it address? This question is a little bit more complicated. Some secure boot technologies aim to prevent the substitution of altered boot code ("unauthorized replacement"), which might introduce malware or security backdoors into a processor once it is initialized. Other capabilities attempt to limit changeable boot parameters such as multi-stage boot code sourcing in the device as it loads, so that an antagonist cannot interrupt the boot process and substitute false commands or security backdoors into the device setup.

Secure boot capabilities

A variety of methodologies are implemented today as secure boot technologies. This includes digitally signed binaries, secure and trusted boot loaders, boot file encryption, and security microprocessors.

Digitally signed boot files provide an important first step in preventing some of the most widespread boot-loading attacks tracked on the Internet. While some manufacturers call this capability "secure boot" because of its increased resistance to repetitive attacks, it is still susceptible to attacks in the verification procedure if the verifying module is not integrated into the embedded processor. In addition, this kind of secure boot does not address the protection of proprietary and sensitive information embedded in the boot file itself.

Improving the trust and security levels of boot loaders is also an important step in industry security and awareness. Proving the security or trust level of the boot loader is itself a multi-variable problem that depends a great deal on the complexity of the boot process. When loading boot code over a network, for example, the security and trust of the boot loader are likely only as secure as the network itself. Securing a large multi-party network is a far larger problem than the secure boot of an embedded processor. This is a security factor in systems that are enabled to receive remote firmware updates.

Boot file encryption/decryption and dedicated security microprocessors are relatively recent design features of secure embedded processors. They offer silicon circuit embedded capabilities to protect operating code and manage secure boot.

Matching requirements to capabilities, claims, and standards

The primary mechanism today leading to the component feature "secure boot" is a digitally signed software boot image verified by the embedded processor during the boot process. This process can coexist with virtually any software image encryption scheme.

Without question, adding a digital signature verification step adds security compared to systems with no digital verification. But to call this system "secure boot" prior to examining the verification process, and the methodology of employing

the digital signature and verifying it, is not necessarily a genuine claim. It is also a nomenclature that does not allow for comparing one "secure boot" scheme to another to determine which is more secure, or which offers more layered secure boot features.

Only one major industry group is currently focused on setting commercially available standards for the boot integrity of processors and embedded processors: the Trusted Computing Group (TCG). This international industry consortium includes many of the largest providers of operating systems and processors; the

Aggressive? You bet!

Wolf announces new PMC and XMC embedded graphics modules for VME, cPCI and VPX architectures.

Military, Aerospace, Space, Industrial and Medical OEMs may now specify Wolf plug-in replacement graphics boards that offer greatly increased performance. Based on an embedded version of AMD's new E4690 graphics chip, they offer over 10 times the 3D rendering speed of earlier solutions, with low CPU utilization and brilliant picture quality.

Wolf proudly offers an outstanding new portfolio of extended temperature video graphic boards for XMC, PMC, VPX, cPCI, PCIe, and VME-64 architectures. Select modules offer up to 28 standard combinations of dual independent display output and up to 19 combinations of dual channel input. All Wolf video graphic products conform to Mil-810 environmental: Shock, Vibration and Extended Temperature Operation and 10-plus years of availability.

Visit www.wolf.ca/products for information and other advanced video boards.

Features:

- Plug-in high performance video upgrade for OEMs
- 10x faster 2D & 3D than previous generation
- Three versions available: (1) Frame Grabber, (2) Multiple Video I/O and (3) Video Output only
- 28 combinations of dual independent video outputs
- 19 combinations of dual video inputs
- Low CPU utilization and brilliant picture quality
- Extended temp -40C to +85C operating environment
- Embedded memory version of AMD E4690 (512MB) graphics chip
- Reduced power modes and improved conductive cooling
- OpenGL drivers, DO-178B and real time operating systems support supplied from ALT Software



Wolf Industrial Systems Inc.
5 Foxfire Chase, Uxbridge, ON L9P 1R4
1.800.931.4114 Fax: 905.852.1735
Online: www.wolf.ca
Please contact sales@wolf.ca



Custom designs, from Concept to Off-The-Shelf in weeks.

TCG consortium is dedicated to creating a common non-proprietary standard for improving the "trust" level of a user computing environment. This group has significantly advanced the cause of improving the integrity of runtime code execution.

The TCG has developed a common standard called the "Trusted Platform Module" (TPM)[1], which serves to monitor the operating kernel within a system, both in boot phase and in operation. A Mobile Trust Module (MTM) is being specified as well. The TPM is designed to digitally sign loaded operating system images and determine

if there has been any tampering in the boot phase or otherwise. There has been some sporadic adoption and specification of more recent versions of the TPM standard, particularly in government systems[2].

Using a commercially available open standard is usually the most cost-effective approach to implementing new features and capabilities, though not always the most secure or tamper-proof. Military customers favor strong defense-in-depth approaches to layered security and redundancy, which often run counter to the cost-effectiveness goals of open standards.

What "secure" really means for products

When engineers at CPU Tech talk about "secure boot" in embedded systems, they define it as a fully encrypted boot file with multiple levels of encryption, implemented by a security engineer within a controlled environment. This involves configuring the device so that it boots only from an encrypted file that can be matched to the secure processor through a unique hardware ID. This process protects the boot code and any proprietary information within, manages the boot process through a dedicated internal security processor to prevent code tampering, and follows boot loader guidelines for secure operating systems to constrain unknown configuration states. These elements of secure boot can be seen in Figure 2.

End goal: More secure systems

As systems architects are well aware, secure boot is not a goal in and of itself. The purpose of a secure boot capability is to design a more secure overall system, where the initialization process and the operating code itself are protected from tampering and reverse-engineering. No system is fool-proof, and so we should be cautious and judicious in the use of the term "secure boot" in an embedded system design.

The author advocates that embedded technology providers and engineers designate the descriptors "digitally signed boot files," "encrypted boot," "boot code authentication," and "trusted boot loader" as separate features in both requirements and system specifications. Additionally, buyers of extremely high-security systems might want to specify boot features in even greater detail. This will avoid the current confusion surrounding the term "secure boot," which is now an abstract term referring to a variety of point solutions and features. ✚

References:

1. Trusted Computing Group. "Trusted Platform Module Main Specification," July 2007. www.trustedcomputinggroup.org/solutions/authentication.
2. Rolf von Roessing, *Computer Weekly*. "ISACA: Users reject trusted computing because of privacy and security concerns," 17 November, 2009. www.computerweekly.com/Articles/2009/11/17/239169/isaca-users-reject-trusted-computing-because-of-privacy-and-security.htm

Themis' New Rugged Servers Have Speed to Burn and Keep Their Cool.

New! 1RU RES Servers

- One or two Intel® Quad-Core 5500 Series Xeon® CPUs with Intel Nehalem Microarchitecture
- Up to 96GB ECC SDRAM
- Up to 3 removable and lockable 2.5" HDDs
- One PCI-E 2.0 x16 slot, optional SAS expansion
- 2RU RES Servers also available



RES-12XR3 server shown with optional filter door panels.

A New Era of Performance and Rugged Reliability

Themis' new family of XR3 Series of Rugged Enterprise Servers™ (RES) includes the latest Quad-Core Xeon processors and Nehalem Microarchitecture from Intel. These new Intel chips revolutionize server performance, and Themis' robust designs - only 20" depth - provide the reliability to keep mission critical applications running. Themis servers provide far greater reliability, improved life cycle management and substantially lower TCO than other COTS systems solutions.

Features in the RES-XR3 servers include:

- Dual redundant, hot-swappable power supplies
- Dual redundant DC power option
- Operating shock - 3 axis, 25G, 20ms
- Operating vibration - 3.0 Grms, 8Hz - 2000Hz
- Light weight, corrosion resistant, 20" depth chassis
- Optional air filter door panels

So when the environment gets tough and your data is critical, turn to the company that builds systems to perform in the harshest conditions. For Sun® Solaris™, Linux®, and Microsoft® Windows® environments. For more information on Themis' rugged new servers, please visit www.themis.com.

Themis rugged, mission-critical computers. Designed to take it.

(510) 252-0870.

THEMIS

Transformational.

©2009. Themis Computer, Themis, the Themis logo, and Rugged Enterprise Servers are trademarks or registered trademarks of Themis Computer. All other trademarks are the property of their respective owners.

Some
Elements of

Secure Boot

Digitally Signed/
Encrypted Boot Files

On- and Off-chip
Boot Authentication

Trusted Boot
Loaders and
Environments

Boot Managed by
Dedicated Security
Processor or Module



J. Ryan Kenny is a product manager at CPU Tech. He is responsible for developing security requirements and certification roadmaps for

the Acalis line of secure embedded processors. He joined CPU Tech in February 2009 and has more than 10 years of experience in space and defense electronics in the U.S. Air Force and defense systems engineering. He graduated from the U.S. Air Force Academy and completed an MSEE and MBA from California State University and Santa Clara University respectively. He can be contacted at rkenny@cputech.com.

CPU Tech
925-224-9920
www.cputech.com

Figure 2 | Separated elements of secure boot systems

Apacer
Access the best

THE MOST RELIABLE STORAGE FOR INDUSTRIES



SDM II 7P/180D LP



Unique Hook Design



SDM SATA Disk Module

- Compliant with SATA 3Gb/sec interface
- Unique locking mechanism
- Perfect solution for 1U chassis-27.8mm height
- Patented Power Cableless Solution
- Sustained Read / Write Speed: up to 27 / 27 (MB/sec)
- Capacity: 512MB-8GB
- Operating Temp.: 0°C~+70°C (Standard)
- MTBF: over 2,000,000 hours (Est.)

SAFD Serial ATA Flash Drive

- Compliant with SATA 3Gb/sec interface
- Perfect replacement of 1.8" / 2.5" HDD devices
- Sustained Read / Write Speed: up to 160 / 135 (MB/sec)
- Capacity: 4GB-128GB
- Operating Temp.: 0°C~+70°C (Standard)
- -40°C~85°C (Extended Temperature)
- *only for SAFD 254/SAFD 181
- MTBF: over 2,000,000 hours (Est.)

Industrial CF Industrial CompactFlash

- Compliant with CFA3.0 Specifications for CFCIII/CFA4.1 for CFC4
- Secured Protection Zone / Quick Erase for ATA CF
- Sustained Read / Write Speed: CFCIII and ATA CF: Up to 35/25 MB/sec; CFC4: Up to 50/21 MB/sec
- Capacity: 128MB-16GB
- Operating Temp.: 0°C~+70°C (Standard)
- -40°C~85°C (Extended Temperature)
- MTBF: over 2,000,000 hours (Est.)

Apacer Memory America, Inc.

Sales Inquiry: ssdsales@apacerus.com

<http://usa.apacer.com>

Tech Support: ssdfe@apacerus.com

On dangerous ground: The rise and fall of military systems power

**Q&A with Dr. Pankaj Rohatgi, technical director
at Cryptography Research, Inc.**



EDITOR'S NOTE

When we in the industry think about military embedded systems power, we typically think about Size, Weight, and Power (SWaP) requirements. But a new – and perilous – practice is arising pertaining to mil embedded power: Simple Power Analysis (SPA) or Differential Power Analysis (DPA) attacks, which proffer sensitive military algorithms to hackers. And these noninvasive attacks are so passive that they can't be detected by the device until it's too late. Editor Sharon Schnakenburg-Hess's recent interview with Cryptography Research Inc.'s Dr. Pankaj Rohatgi reveals more about this security threat – and how to thwart it. Edited excerpts follow.

MIL EMBEDDED: *Security for military embedded electronics is a hot topic these days. What do you think is the most serious type of security breach in that industry?*

ROHATGI: One very strong area of concern is the extraction of keys or the reverse-engineering of sensitive military algorithms using Simple Power Analysis [SPA] and Differential Power Analysis [DPA].

These attacks involve measuring and analyzing the power consumed by a device while it is performing its normal operations with secret keys and algorithms. Such passive, noninvasive attacks cannot be detected or audited by the device.

MIL EMBEDDED: *That makes me wonder two things: Which types of devices are most vulnerable, and are these attacks truly never detected?*

ROHATGI: To answer your first question, portable electronics, communications gear, and “leave-behind” equipment are the most vulnerable: They are easiest for an enemy to acquire and access. After conducting the attack, the enemy could eavesdrop on military communications and forge command-and-control messages.

The answer to your second question about detection of attacks is that commercial

systems are regularly attacked using SPA and DPA, but these attacks are discovered when fraud or crimes are committed with the extracted information. In a military setting, the enemy will be much stealthier, and successful attacks might not get discovered until it is too late.

“ ... Commercial systems are regularly attacked using SPA and DPA, but these attacks are discovered when fraud or crimes are committed with the extracted information. In a military setting, the enemy will be much stealthier, and successful attacks might not get discovered until it is too late. ”

MIL EMBEDDED: *So what are SPA and DPA attacks, technically speaking? How are they used to access secret information?*

ROHATGI: The energy consumed by a hardware device depends on the switching activity of its transistors, which in turn depends on the operations it is performing. An attacker who is passively measuring a device's power consumption or electromagnetic emissions will recover some aggregated and noisy information about the sensitive data being processed. As I mentioned, SPA and DPA attacks use the information available from power measurements to extract secret keys from a device.

MIL EMBEDDED: *OK, let's start with SPA attacks and go into more depth.*

ROHATGI: OK, so in an SPA attack, the attacker recovers the secret keys by directly observing features within individual power consumption measurements. Implementations that have very different power consumption profiles for different keys are most vulnerable to SPA. For example, implementations of modular exponentiation for public-key cryptography algorithms such as RSA [Rivest-Shamir-Adleman] and Diffie-Hellman may use a key-dependent sequence of square and multiply operations. And scalar multiplication in Elliptic Curve Cryptography [ECC] may be implemented using a

key-dependent sequence of double and add operations. Such implementations can leak the value of the key from a single operation. [Figure 1 shows the power trace from an RSA operation using the key-dependent square and multiply sequence.] The square and multiply operations have visibly different power profiles that are easy to distinguish. The secret exponent has been recovered from the sequence of squares and multiplies.

MIL EMBEDDED: *Now that we've covered SPA attacks, let's delve a bit more into what a DPA attack is.*

ROHATGI: Right. DPA attacks employ statistical techniques over multiple power

consumption measurements to extract secrets, even when the information available within any individual measurement is small and masked by other activity and noise. The basic concept behind DPA is that the overall power consumption is correlated to the bits of computational intermediates during device computation. By focusing on intermediates that depend only on a few bits of the key, it is possible to use power measurements to determine those bits of key. For every possible guess of these key bits, the attacker can predict the computational intermediate and compute the correlation between the power measurements and bits of the predicted intermediate. [As shown in Figure 2], for a correct guess of these key bits, the attacker will observe correlation spikes whenever the intermediate is being processed. For an incorrect guess, there won't be any correlation spikes or the spikes would be smaller. Once these key bits are determined, the same divide-and-conquer approach can be repeated with other intermediates to determine the other bits of the key.

MIL EMBEDDED: *So are these attacks executed by sophisticated attackers or ordinary Joes trying to get ahead?*

ROHATGI: The level of sophistication and investment required to perform these attacks is quite low. SPA and DPA are straightforward to implement and can be performed with less than \$5,000 of standard lab equipment.

MIL EMBEDDED: *What makes a system or application most vulnerable to DPA or SPA attacks, specifically?*

ROHATGI: All tamper-resistant devices and cryptographic algorithms are susceptible to these attacks if they do not contain countermeasures. In our experience, systems without countermeasures are often broken in a matter of hours. Devices where the power supply is inaccessible remain vulnerable to DPA-like attacks on the device's electromagnetic emissions.

The degree of vulnerability depends on the ratio of information leakage relative to noise. Implementations that perform different sets of operations for different values of the key are the most vulnerable to SPA attacks and may be broken from a single power measurement.




Figure 1 | SPA – Power trace of an RSA exponentiation showing the square and multiply sequence and the recovered secret exponent.

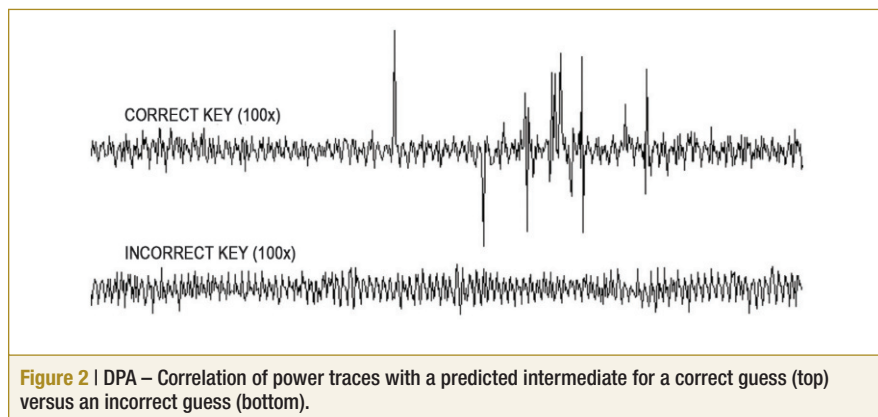


Figure 2 | DPA – Correlation of power traces with a predicted intermediate for a correct guess (top) versus an incorrect guess (bottom).

MIL EMBEDDED: *You are on record as saying that applications utilizing FPGAs are particularly vulnerable to DPA. Why is that? Are you suggesting that military technology should not include FPGAs?*

ROHATGI: Without countermeasures, all forms of silicon, including FPGAs and ASICs, are susceptible to DPA. General-purpose silicon such as FPGAs and microprocessors receive more attacker attention because they are easier to acquire. And by investing time in characterizing the specific leakages present in that device, the attacker can target multiple systems that use the same FPGA or microprocessor.

Security differences between FPGAs and ASICs are relatively minor, however, and both technologies will require countermeasures for security. Once countermeasures have been adopted, FPGAs offer valuable advantages: It is much easier to test and refine FPGA designs for DPA resistance, and there is the flexibility to upgrade a design in the field if a security vulnerability is discovered.

MIL EMBEDDED: *DPA attacks are alive and well, but how can they be prevented or circumvented – or can they?*

ROHATGI: Cryptography Research discovered SPA and DPA in the 1990s and developed the fundamental techniques for securing systems against DPA. Defending against DPA is quite feasible: Commercial products such as chip-card based payment systems routinely pass stringent requirements and tests for DPA resistance.

MIL EMBEDDED: *So countermeasures can be invoked, as you said?*

ROHATGI: That's correct. At a general level, the fundamental categories of countermeasures to DPA include things like *leakage reduction*. Leakage reduction includes techniques to make the set or sequence of operations independent of the key as well as hardware and software balancing techniques to reduce variation in the power consumption for different data. This reduces the leakage-signal to noise ratio and increases the number of power measurements needed for a successful attack.

Then there's the *noise introduction method*. This includes techniques for adding different types of noise into the power consumption measurements available to the attacker, thus reducing the leakage-signal to noise ratio.

Another method is *obfuscation*: By keeping algorithms secret, the attacker is forced to perform reverse-engineering along with power analysis. While we do not recommend this countermeasure, it is better than having no countermeasure at all.

MIL EMBEDDED: *So we've got leakage reduction, noise introduction, and obfuscation. Have we missed any methods?*

ROHATGI: Those wanting to prevent SPA or DPA attacks can also *incorporate randomness*. This includes methods for randomizing the data manipulated by the device in a way that still produces the correct result and encompasses techniques such as the masking or blinding of data and keys. These techniques force the attacker to employ more complex attacks, such as higher-order DPA and a larger number of measurements.

And finally, there are *protocol-level countermeasures* that can be applied when there is flexibility to modify cryptographic protocols used by the device. Protocols are modified so that secrets can be continually refreshed and updated during the lifetime of the device, so that an attacker is never able to get sufficient information about any particular secret.

MIL EMBEDDED: *So which of these countermeasure methods is best – or should they be used in concert?*

ROHATGI: Because DPA attacks amplify leaked information through signal processing, systems generally benefit from using multiple countermeasures. However, several commercial and government security standards require leakage testing to determine the overall effectiveness of a countermeasure implementation.

MIL EMBEDDED: *Where does Cryptography Research come into the picture?*

ROHATGI: There are several labs and testing platforms, including our DPA Workstation, which can evaluate devices for information leakage and DPA vulnerabilities. Given the magnitude of the threat and the increasing requirements for DPA resistance, we recommend that vendors use these or similar testing resources to obtain a baseline assessment of information leakage from their cryptographic device – before the perhaps unimaginable becomes an unfortunate reality. ✚

Dr. Pankaj Rohatgi is Technical Director, hardware security solutions at Cryptography Research, Inc. Prior to joining Cryptography Research, Pankaj spent 13 years at IBM, conducting research in cryptography, secure hardware, systems, and network security. Prior to IBM, he was the security architect for the OpenTV operating system at Thomson R&D labs and at a Thomson/Sun Microsystems joint venture. He received his Ph.D. in Computer Science from Cornell University. Pankaj can be reached at pankaj.rohatgi@cryptography.com.

Cryptography Research, Inc.
415-397-0123
www.cryptography.com



Network

SBCs

Storage

I/O
FPGA

Software

Systems



We have the right piece of the puzzle.
Find out now...

Elma offers way more than chassis

Embedded Storage for Secure Data



Elma Electronic's Systems division has a complete line of storage solutions to meet tough environmental requirements and secure data needs:

- Mezzanine based CF and SSD (PMC, XMC, and AMC)
- Hot swap removable shuttle based storage on PMC
- Blade level Network Attached and RAID storage (VME, VPX and cPCI)
- Hardware enabled Secure Erase and Write Protect features
- Available in convection and conduction cooled

Call us or visit our website for more details.

elma.com • acttechnico.com
(510) 656-3400 • (215) 956-1200



Embedded Computing Solutions from Elma


Intel® Core™ i7 Processor Solutions

Optimized For Embedded Computing Applications.

6U VPX • 3U VPX

6U cPCI • 3U cPCI

XMC • VME



X-ES Intel® Core™ i7 Processor Solutions: Unmatched Performance

Extreme Engineering Solutions, Inc. (X-ES) unleashes the performance of the Intel Core i7 processor for embedded computing. By utilizing a processor with integrated graphics, PCIe, and ECC DDR3 memory controllers, the X-ES solutions deliver unmatched power savings and processing performance for compute intensive commercial and military applications.

X-ES offers an extensive product portfolio that includes commercial and ruggedized single board computers, high-performance processor modules, multipurpose I/O modules, backplanes, enclosures, and fully integrated systems.

Intel Core i7 processor solutions available now in a variety of form factors. Call or visit our website today.

X-ES

Extreme Engineering Solutions

608-833-1155 • www.x-es.com

Securing wireless Local Area Network interconnections with Layer 2 encryption

By Juan Asenjo

Enabling military and civilian government operations to dynamically interconnect Local Area Networks (LANs), wireless technologies are a lifesaver in environments where wired connections are cost-prohibitive or just not practical. However, transmitting sensitive information over the airwaves presents security challenges including passive attacks and active attacks. Enter Layer 2 encryption, which can effectively thwart these security challenges.

Wireless technologies enable military and civilian government operations to dynamically interconnect Local Area Networks (LANs) quickly and reliably in environments where wired connections are impractical and cost-prohibitive. This connection of LANs over the air without the use of a fixed, wired medium is typically referred to as *wireless interconnectivity*. Under this umbrella, a number of specific connection technologies are used including radio frequency, microwave, and free-space optics.

While popular from an operational perspective, as mentioned, wireless LAN interconnections suffer significant

drawbacks when it comes to security. As with any open medium, ensuring the confidentiality and integrity of sensitive data traveling across these networks is of paramount importance, particularly in government and military applications. These security challenges incurred by transmission of sensitive information over the airwaves include both passive and active attacks. *Passive attacks* occur when perpetrators collect and read sensitive data, whereas *active attacks* occur when perpetrators inject new traffic and network integrity is breached.

To provide insight into remedying these challenges in a government arena, the

following discussion examines LAN operational advantages and associated vulnerabilities – and explores Layer 2 versus Layer 3 alternatives for enhanced security.

Wireless technologies for LANs

The proliferation of wireless LAN interconnections within government and enterprise has come as a result of LAN flexibility, ease of deployment, and cost savings. As alluded to previously, outdoor wireless interconnections over radio frequency, microwave, and free-space optic mediums allow system architects to connect LANs dynamically without having to physically lay cable




Figure 1 | Wireless LAN interconnection in a forward-deployed tactical battlefield environment

or provision a service. In military environments in particular, wireless LAN interconnections can be established and dismantled at a moment's notice in accordance with changing tactical and strategic battlefield conditions. Examples of this include forward-deployed tactical units and strategic intra-base virtual campus topologies such as military clinics and hospitals. A schematic representation of this environment is shown in Figure 1.

While providing quick setup and complete ownership of the backbone wireless LAN links, the connections offer no inherent level of security. Wireless LAN interconnections are vulnerable to interception, and therefore, must be secured to ensure the confidentiality and integrity of the data traveling across them. As a result of this vulnerability, the U.S. government has developed regulations to mitigate the threat of interception and specifies encryption as the preferred mechanism for protecting sensitive data. Within the Department of Defense (DoD), directives DoDD 8500.2 and DoDD 8100.2 mandate that Sensitive But Unclassified (SBU) data be encrypted using FIPS 140-2 approved equipment employing the Advanced Encryption Standard (AES) algorithm when employing wireless systems.

In theory, encryption across LANs can be done at any of the seven layers defined by the Open System Interconnection (OSI) model for data networking (Figure 2). The OSI architecture model defines the functions and components that establish a data connection. Depending on where encryption is employed in the layered model, the more transparent and therefore effective it can become. Higher in the model (at Layer 7), specific applications are considered, while at the bottom (Layer 1), the general physical medium

is addressed. Data encryption is generally done at the frame (Ethernet Layer 2) or packet (IP Layer 3) levels.




Figure 2 | OSI reference model for data networking

Layer 2 versus Layer 3: Advantages and vulnerabilities

While the application of encryption technologies to protect LAN interconnections can thus be made at either Layer 2 or Layer 3, with the proliferation of the Internet, most encryption devices available in the market until just recently were packet encryptors operating strictly at IP Layer 3 using the IP Security (IPsec) encryption standard. However, with increased traffic volumes and growing use of latency-sensitive applications such as voice, video, and multimedia, IPsec has shown significant limitations that impact operational performance. Given the nature of deployed battlefield communications, Layer 3 interconnections using IPsec encryption have proven impractical.

Also, unlike Virtual Private Networks (VPNs) that use IPsec with intricate security associations, Layer 2 makes the process simple and independent of complex routing tables that add unnecessary overhead to the operation. By creating "tunnels" across the LANs being interconnected, private traffic is segregated and protected on the otherwise open medium.

TRI-M ENGINEERING

PC/104 Can-Tainer



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



+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.



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

Additionally, Layer 2 establishes the physical connection between the local telecommunication devices and remote destinations, and defines the data frame as the physical transmission medium between nodes. Layer 2 connections are primarily used for high-speed/high-data throughput applications between telecommunication facilities. When this layer is used to connect telecommunications facilities on high-speed lines, encryption mechanisms encapsulate all higher-level protocols crossing the link.

Delving deeper in our discussion, Layer 2 encryption typically is performed either in bulk where the entire medium is encrypted, or in a tunneled mode where only the data payload of the Ethernet frame is secured. Wireless point-to-point applications are typically bulk encrypted as they usually only connect two discrete sites. Applications where a switched network is employed use tunneling to maintain the frame header information in the clear while encrypting the rest of the payload.

Compared to Layer 3 IPsec encryption – which significantly impacts throughput – Layer 2 encryption offers the best performance. Ethernet bulk mode yields full throughput with no frame expansion. Tunneling typically expands frames no more than 22 bytes. When comparing the

performance offered by Layer 2 encryption versus Layer 3, throughput for small frame/packet size (64 byte – typically employed in applications such as voice over IP, video, and multimedia) can be improved as much as 60 percent versus Layer 2 encryption, as illustrated in Figure 3.

Theoretical Throughput [Mbps]




Figure 3 | Throughput comparison (Mbps) between Layer 3 IPsec and Layer 2 Ethernet encryption (Source: Rochester Institute of Technology)

Safety, Security, Reliability

AdaCore, your partner for high-integrity software development.

Expert consulting and tool support
 MC/DC code coverage
 Static code analysis
 Automatic peer review
 Tools for formal verification and correctness
 Certification and tool qualification materials
 DO-178B / C
 IEC 61508



www.adacore.com

AdaCore
 The GNAT Pro Company

Latency, the amount of time needed for data to go through an encryption device, varies depending on the frame sizes utilized and generally decreases as the frame size increases. Depending on line speeds, Layer 2 tunneling will normally vary between 4 and 40 microseconds. IPsec, on the other hand, can add millisecond latencies, severely impacting performance.

In contrast to IPsec, encryption of the Ethernet frame over Layer 2 solves many of these challenges, providing line-speed encryption with minimum frame expansion and delay. Besides these advantages, Layer 2 encryption also eliminates the need for complex router and configuration management, and can be deployed as a “bump-in-the wire” security solution allowing higher-level protocols to be encapsulated within the encrypted Ethernet frame. Because of this characteristic, Layer 2 encryption can typically fit within the existing IT infrastructure, requiring minimal incremental resources and training to operate.

Enhancing LAN security

LANs are known for their ease-of-use and quick setup. However, LAN security is only as good as the weakest links that tie the wireless network together. Numerous protection challenges including strong access control mechanisms, intrusion detection and prevention systems, firewalls, malware removal, and encryption are often employed within LANs. However, if these methodologies are not connected securely, tremendous data compromise and interception vulnerabilities will result.

However, there is hope. Deployment of wireless LAN interconnectivity using Layer 2 cryptography allows government, military, and civilian organizations to implement robust (meaning: secure) encryption quickly and with minimal network disruption, while typically preserving current investments. Layer 2 also enables military organizations to protect sensitive information exchanged across wireless LANs without affecting operational performance. To meet security requirements and reduce overall network complexity, wireless LAN interconnections operating at speeds up to 10 Gbps using Layer 2 security are becoming increasingly popular. ✚



Juan Asenjo is a Senior Product Marketing Manager at Thales e- Security, where he manages the company's network security product line. Juan has worked in the information security field for 24 years, including more than 10 years in government and military environments. He has degrees in Engineering and Business, and is a Certified Information System Security Professional (CISSP). Juan can be contacted at Juan.Asenjo@thalessec.com.

Thales e-Security
954-888-6200
<http://iss.thalesgroup.com>

QUALIFIED TO PERFORM
RUGGED C4 SUBSYSTEMS

Ultra-Rugged COTS Computers, Routers, & Switches Qualified for MIL-STD Performance in Platform Modernization & Situational Awareness

Dual-Core Modular Mission Computer

10-Port Gigabit Ethernet Switch

Cisco IP Router w/ GigE Switch

Parvus
www.parvus.com | sales@parvus.com | 800.483.3152

Photo Courtesy of Department of Defense



Guest opinion: Solid state security – A potential threat

By Mark Downey

History has proven that an effective strategy for protecting the security of information is one of the most critical factors in wartime success. When it comes to solid state information storage, the biggest threat to security might reside in the technology and manufacture of the storage device itself. Reliability is also imperative in ensuring information integrity – and ensuring warfighter safety.

Defense and aerospace applications such as rugged computing and mobile data acquisition systems were among the first to utilize solid state storage technology. This adoption leadership was realized because designers in these markets understood the harsh environmental survivability benefits of solid state technology when compared to mechanical hard drives – and decided solid state was worth the price premium. Solid State Drive (SSD) developers then began creating innovative SSD technologies in response to the defense market's values, requirements, and priorities. Thus, the door for implementation of solid state storage technology was opened wide, and myriad benefits ensued.

As time marched on, development cycles continued to turn out SSD technologies that reflected the defense and aerospace market's desires. These SSD wares also positioned the U.S. military with the unsurpassed storage systems that we have come to appreciate and value. However, the recent explosion of SSD technologies into more commoditized, commercial applications has disrupted this trend, stealing the focus away from the defense and aerospace markets in exchange for the handsome prospects of larger markets. The price tag: A downturn in SSD security and reliability that could put our nation's information – and warfighters – at risk.

The way it was

Throughout the '90s, design engineers developed system technologies for defense and aerospace applications and garnered the full attention of the many SSD developers. By being the most significant adopters of the technology, defense industry needs came first and helped to mold and define development goals. Most SSD providers were relatively small companies and would readily go the extra mile to customize or design harsh-environment, ruggedized products needing to fulfill specific security and reliability requirements. However, over the past few years, SSD production has begun to explode in both the enterprise and notebook markets, in turn forever changing the SSD provider marketplace.

The way it is


Today, military systems designers are faced with a different paradigm. No longer

do they define the storage requirements they would like to see from the SSD providers, but rather they must settle for what might or might not be available. They find themselves wading through a large number of SSD technologies, most of which hold hidden risks or subtleties that might be difficult to discern:

- How is the SSD actually tested?
- Will the device meet the specification over the entire life of the product?

- Will it survive extended temperature operation and if so, for how long?
- How will obsolescence or life-cycle changes be managed?

Program requirements continue to push cost savings and, therefore, the use of Commercial Off-the-Shelf (COTS) products, further complicating the already cumbersome decision process. Today, more than ever, directives are often based on what is available rather than what is



The Industry's Choice in VPX Handles and Panels

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

ELMA
Your Solution Partner

USA Elma Electronic Inc.

Phone: 510.656.3400 Email: sales@elma.com

Web: www.elma.com




IEEE Handles

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



Low Profile Handles

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



Panels

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

possible. Designers find themselves working to solve problems associated with using commercial-environment products in a harsher and more security-critical defense and aerospace environment where reliability is key (Figure 1). Typical questions might include:

- How does one make this Serial Advanced Technology Attachment (SATA) connector work in a high-vibration and high-shock environment?
- Is this reliability information really valid at industrial temperatures?
- Can one get an SSD built with high-grade flash devices that are burnt-in to ensure a higher level of fielded reliability?

The bottom line is that it is becoming increasingly more difficult to identify an SSD manufacturing partner willing to address these real technical challenges, never mind one whose original design intention embodies the reliability and security priorities that mission-critical environments demand.

Security threat

One conflict evolving with this recent trend relates to the topic of security. As a primary requirement for successful adoption, the larger and more dominant SSD markets continue to pressure providers for lower costs. To meet these reduced costs,

trade-offs are made that can threaten overall security in military applications. Price pressures might force SSD manufacturers to utilize non-U.S. designers and labor, resulting in potential security risks for U.S. DoD and U.S. military applications that might go unrecognized.

SSDs are constructed around a controlling management device, often a processor, which also contains firmware. This construction makes them susceptible to certain types of attacks if proper precautions are not taken. These attacks might be subtle, only preventing the erasure of critical data in times of duress – or they might allow some master key to be used, permitting an enemy to extract encrypted data without the original key. Thus, maintaining control over security in tomorrow's warfighting systems becomes more challenging, yet increasingly urgent and essential. Data is moved and stored throughout networks like never before, and the storage element of the overall system is one of the most vulnerable areas. The question of securing this storage effectively must remain in the crosshairs of design priorities and cannot be ignored or traded to save a few dollars. Security remains of paramount importance and must not be diminished.

How does one accomplish world-class security with products that are sold globally? Is it even possible? Isn't

“ The question of securing this storage effectively must remain in the crosshairs of design priorities and cannot be ignored or traded to save a few dollars. ”

“standardized security” akin to “fat-free oil”? Standardization is important, as long as we find a way to maintain security; where conflicts arise, one must choose security. Pentagon spokesman Bryan Whitman in a recent press conference warned, “As you develop those [technologies], you have to be mindful of how the enemy can counteract any technology you have. That's why you always have a constant review process in place to not only improve that capability, but address any vulnerability it may have.” His statement preceded a discussion of a security breach of the unmanned Predator drones used in Afghanistan and Iraq. Shi'ite fighters in Iraq used open market software costing less than \$26 to intercept and monitor video feeds from U.S. drones.

Such incidents demand attention and require consideration as a threat to warfighter network security. How does one place a value on, or define a cost target for, security? Designers must be aware of these risks and consider that SSD choices will continue to be influenced by the requirements of the larger enterprise and/or notebook markets.

Reliability threat

Along with security, a second area of defense industry concern regarding SSD technologies is reliability. While many have had a PDA or laptop freeze up at an inopportune time, the potential consequences of that kind of malfunction can be disastrous for the warfighter. The soldier under fire in the mountains of Afghanistan would not want to know



Figure 1 | The potential consequences of an SSD malfunction in military applications can be disastrous.



Figure 2 | There are problems associated with commercially designed SSDs used in security-critical defense and aerospace environments, where reliability is key.

that the SSD on which he is relying has components designed for use in an air-conditioned computer lab – or that design trade-offs were made such that it could meet the \$150 price point required for sale into the soon-to-be explosive Netbook markets of India. The implications of using commercial-grade SSDs in military applications designed for such disparate markets could be tragic (Figure 2).

A call to action

So what is the answer to this perplexing paradigm? Government agencies and the private sector must work together to establish clear and definable safeguards against such evolving SSD security and reliability risks. Price pressures will continue to dominate the decision process until something tragic triggers a reactive response in the supply channel.

With the shift of attention away from military/aerospace requirements comes the risk of settling for mediocrity in solution sets. Both reliability and security are at risk of being watered down or lost with the disruptive, commercially driven changes in SSD development. The warfighter in the field will no longer be getting the best of what one can design for them, but rather the closest fit of what has been designed for other more dominant markets. Although this might appear as a

subtle difference today, given sufficient time and product design cycles, the technology gap between would have been and what *is* will certainly expand. †

References:

- [1] Malicious Firmware Could Sabotage Military, Security Systems, By Stew Magnuson, National Defense magazine, Feb. 2010, <http://www.nationaldefensemagazine.org/ARCHIVE/2010/FEBRUARY/Pages/MaliciousFirmwareCouldSabotageMilitarySecuritySystems.aspx>



Mark Downey is former Director of Defense Technology at White Electronic Designs. He has 18 years' experience in the electronics industry in the

areas of solid state storage technology, high-speed memory bus architecture, electrical simulation, thermal management, complex microelectronic packaging, and process development. He co-holds a U.S. patent in the area of dense memory module packaging. Mark also holds a BSEE from the University of Massachusetts at Lowell.

White Electronic Designs
(now Microsemi)
602-437-1520
www.wedc.com

Trust a world-wide expert for your embedded critical network application

SWITCHES & COMETH & IP ROUTERS

More than 30 models... VME, cPCI, VPX

ComEth 4300a

- 4 front Giga ports (copper or fiber)
- 2 front 10 Gigabit Ethernet ports
- 20 rear Gigabit Ethernet ports

10Gigabit Ethernet Routers

Picmg2.16 & VPX

ComEth 4340a

SBCs PREMIUM

Intel® & Freescale® processors

IG-De6-VMEb

amazing processing unit

- Two MPC 8640 single or dual core
- One embedded Ethernet Switch
- One Open FPGA VIRTEX 5
- PMC/XMC, USB, RS232/RS422, IOs

New OpenVPX Range

- Intel® Core 2 Duo®, MPC8640, Virtex5&6

OpenVPX 3U Single Board Computers

For more information on our products and custom design services...

www.interfaceconcept.com

+33 (0)298 577 176

New European program exhorts nano fluids to teach systems how to “be cool”

Interview with David Mullen, NanoHex project director and mechanical engineer at Thermacore



EDITOR'S NOTE

No one's really 100 percent certain why the heat capacity of nano fluids is so much higher than other cooling fluids, but the fact is: They're really cool for hot electronics. They can offer up to a 350 percent cooling improvement over conventional coolants. But manufacturing nano coolants is slow, and the European Commission (EC) is looking for ways to ramp production and bring down the price. The avenue: the NanoHex nano cooling fluid project, headed by David Mullen of Thermacore, whom editor Chris Ciufu recently interviewed to get the whys and wherefores. Edited excerpts follow.



Lockheed Martin photo by Tom Harvey

“ What the EC would like us to do is to take promising laboratory-based results, which can only produce perhaps a few kilograms of nano fluid coolant currently, and upscale that to create a pilot line to result in the eventual manufacture of a few tons of nano fluid coolants. ”

MIL EMBEDDED: *Let's start with a brief overview of what Thermacore is and what it provides.*

MULLEN: OK. Thermacore is a global company with corporate headquarters in Lancaster, Pennsylvania that specializes in advanced thermal solutions. However, I work for the Thermacore Europe subsidiary in the United Kingdom. Thermacore serves a variety of OEM applications across a diversified set of global markets that include military/aerospace, medical, computers, communication, energy conversion, power electronics, and government.

Essentially Thermacore's core technology is based on heat pipes and heat pipe assemblies, but we also offer other thermal solutions such as liquid-cooling cold-plate components and systems, enclosure heat exchangers, and complex vacuum brazing heat exchangers. Thermacore has specialized in the development and application of heat pipe technology for more than 40 years, and our new NanoHex nano fluid coolant program [www.nanohex.eu] is hopefully going to allow us to extend liquid-cooling technology as well.

MIL EMBEDDED: *Can you explain briefly how a heat pipe works, and how or when you implement them?*

MULLEN: Sure. Essentially a heat pipe is a heat superconductor. A heat pipe utilizes two-phase heat transfer to passively [no external energy] and efficiently move heat from a concentrated area and transfer it to a remote area where it can be removed to the external environment through natural convection, forced convection, or radiation. Thermacore's engineering services include design using CFD [Computational Fluid Dynamics] analysis, Finite Element Analysis (FEA), as well as development and prototyping, and product qualification. Many of our customers come to Thermacore late in the design stage after they have discovered a thermal problem within their system. Obviously, the late-design stage challenge restricts us a little, though, and we prefer to be involved in the initial concept stages of the electronic or system design.

MIL EMBEDDED: *OK so switching gears, let's talk about your recent announcement of the NanoHex project. Set the stage for that, if you would, please.*

MULLEN: Thermacore Europe has an ongoing program of recruiting PhD students, and then about 90 percent of those will be taken on as engineers by Thermacore after graduation. So we're in quite a good position to take on forward-looking technologies. And liquid cold plates is one area that we want to try and improve the thermal technology.

Within Europe, there is a program called a *framework program*, funded by the European Commission [EC] to enhance the EU's technical base by targeting technologies in European companies. The EC had identified 53 billion euros to invest in new technologies. So we learned that one EC "call" [meaning, a project bid request] was to upscale manufacturing on nano technology.

So we decided to, along with a few academic institutions with which we already had contacts, set up a project to target this particular call to develop nano fluid coolants.

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 VT104 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

MIL EMBEDDED: *When was all this, and how much was the final price tag?*

MULLEN: Let me put it this way: After three years, we won the NanoHex project from the European Commission. We believe if it's not the biggest, it's at least one of the biggest nano technology projects in the world. The EC funded it with 6.1 million euros, with additional contributions from our partners making it an 8.3 million euro project. The project actually started on September 1 of last year and will run for 3 to 3.5 years.

MIL EMBEDDED: *Can you describe what the NanoHex project entails?*

MULLEN: The original call was to upscale the production of nano technology, and in our case that's nano fluid coolants. What the EC would like us to do is to take promising laboratory-based results, which can only produce perhaps a few kilograms of nano fluid coolant currently, and upscale that to create a pilot line to result in the eventual manufacture of a few tons of nano fluid coolants.

MIL EMBEDDED: *So the challenge is manufacturing, which is a surprise to me. What are the additional end goals?*

MULLEN: Beyond manufacturing a pilot line and producing a few tons of this material, per the EC's directive, we also need to

sell and commercialize the nano fluid coolants: The EC wants to see that we're looking at even more opportunities in the industry that might be a fit for using the fluid. So, within the project we've identified two main areas: data center cooling and traction control IGBT cooling.

MIL EMBEDDED: *Any other areas for possible usage, beyond those two?*

MULLEN: One of the key technologies we're targeting for the nano fluid coolant is housed within data centers, where there is an ever-pressing need to reduce energy costs. Predominantly, energy is used to cool the servers or cool the environment in which they work. What we hope to offer the data center is nano fluid coolants with up to a 350 percent improvement when compared to a traditional coolant. We hope to utilize it in the data centers' hot spots within cabinets, and to provide some exceptional cooling that will ultimately lead to reduction in energy costs, perhaps even increasing the actual server's density so that we can even look to reduce the overall size of the data centers. There are lots of possibilities.

MIL EMBEDDED: *Can cooling racks of equipment reduce the energy consumed by the servers, or would it reduce the energy used in cooling the servers themselves?*

MULLEN: Predominantly, it's the energy used in cooling the servers.

MIL EMBEDDED: *I see. So the servers are still going to consume what they consume but you could reduce the amount of cooling required because you're putting it right where it needs to be.*

MULLEN: Yes. And our technology goes even a step beyond that because we want to take the nano liquid coolant directly to the chip. So it's taking the heat directly from the chip or any other component within the server blade, out – and then perhaps into a heat exchanger on the cabinet and then to a heat exchanger outside the building. It could even then be used, as some people have started to use the coolant, to heat their office space, for example.

MIL EMBEDDED: *With a heat exchanger, you take the hot air out of the data center and then route it into the rest of the offices of the building?*

MULLEN: Not the air. What we're doing is taking a nano fluid coolant and circulating that around a data cabinet where the servers are, taking more heat out because of the extra performance we can get from nano fluids, into a heat exchanger on the back of the cabinet. And then more coolant, not air, would take that heat energy outside of the building via a heat exchanger.


MIL EMBEDDED: *What's the practical distance of that, within the building itself?*

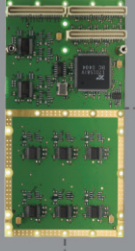
MULLEN: Data centers can vary in size from half a football pitch to even four football pitches or the size of a small office block. But a lot of data centers have been built with the infrastructure to support liquid cooling underneath the floors. Pipe works are

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

TEWS TECHNOLOGIES

www.tews.com

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
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

© 2006 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.

generally already there. So we'll be utilizing what's already there to take the heat from the processor within the server cabinet to the heat exchanger or to that pipe or that infrastructure that already exists.

MIL EMBEDDED: *Is the nano cooling fluid proprietary, patented, standard materials, or what?*

MULLEN: It's not proprietary. It's patented by a couple of the partners within the project. There are 12 project partners in total, many academic and several industrial. We have two partners – a UK company called Dispersia and also ITN Nanovation in Germany – who already produce nano fluid coolants on only a laboratory scale. They employ two different techniques of producing nano fluid. So if we find one is too risky to upscale, we've always got an alternative. But at the moment, both technologies appear feasible in achieving the pilot lines.

MIL EMBEDDED: *What are the technical challenges in producing the fluid itself?*

MULLEN: One challenge is in the manufacturing, that is, the actual particle conditioning and also coating the particles' surfaces to create a nano fluid coolant where the particles don't agglomerate. That way, they remain very stable and very dispersed within the actual base carrier fluid. Another challenge is to upscale the fluid and be able to repeat what we see in the laboratory at an industrial level.

And the final challenge is to turn the nano fluid coolant into something that can be used, so we're looking to develop very advanced cooling devices because the condition of the fluid as it flows through these heating devices is very, very critical to its performance.

Lots of people have done work on nano fluid, but nobody understands the fundamentals of how it transfers that heat, which is what has primarily hindered a lot of nano fluid development worldwide. We have a good knowledge, and a certain part of the project entails understanding that mechanism of heat transfer and developing very specific transfer devices or cold plates that can harness that. ⊕

David Mullen is a Senior Research and Development Engineer at Thermacore Europe Ltd. Since 1987, he has worked in a number of senior engineering capacities and is a chartered member of the Institution of Mechanical Engineers. Predominantly, he has specialized in engineering project management, including a number of multimillion-dollar projects. NanoHex, which David will be coordinating, represents a significant opportunity to develop an advanced thermal solution to reduce energy usage around the world. He can be contacted at d.mullen@thermacore.com.

Thermacore
+44 (0) 1670 859518
www.thermacore.com

For more information on the NanoHex project, go to www.nanohex.eu.

DATA STORAGE TECHNOLOGY
RPC12 Ruggedized 3U Fibre Channel RAID System

Phoenix International designs and builds rugged COTS Data Storage Systems that plug and play in any application -- from Multi-Terabyte Fibre Channel RAID and Storage Area Network configurations to plug-in Solid State Disk Drive VME/cPCI Storage Modules.

Low Operational Temperature -20° C

High Operational Temperature +60° C

Operational Altitude to 45,000 feet

- Operational altitude to 45,000 feet
- Operational Temperature -20° to +60° C
- Redundant, hot swap components/FRU's
- 40Hz to 440Hz, 90/240 VAC Input Operation

PHOENIX INTERNATIONAL

See us at: www.phenixint.com or contact us at: 714-283-4800 • info@phenixint.com
An AS 9100 / ISO 9001: 2000 Certified Service Disabled Veteran Owned Small Business

We Put the State of the Art to Work™

OpenVPX Channel

Visit: channels.opensystemsmedia.com/OpenVPX

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

"Hardcore" liquid-submersed workstation goes and goes ...

When those internal components' temperatures go up, system performance goes down. Such is the natural order of the embedded electronics world. However, Hardcore Computer, Inc.'s new Detonator workstation, cooled by liquid submersion methodology, is designed to crush the natural order. Praising liquid submersion for its performance levels as compared to air cooling, Hardcore states that its Core Coolant, used in the process, "is 1,350 times better than air, by volume, at heat removal" to provide what they say is improved reliability of 24/7/365.

Ideal for compute- and GPU-intensive applications such as military programs, simulation, real-time situational awareness, CAD, and digital content creation, the Detonator gets its power from up to two 5500 or 5600 Xeon processors. The workstation also includes optimized data pathways and the Intel 5520 chipset on a purpose-built motherboard. Other notables include 4x DDR3 slots for each CPU; 2x 1.0 Gbps LAN; 3x 3.0 Gbps SATA II slots on the motherboard for up to 3x 2.5" internal SSDs; 4-drive RAID 0, 1, 5, or 10; or 5-drive RAID 0, 1, or 5. The Detonator also has a Creative Labs X-Fi audio processor (20K2) enabled, in addition to Dolby 7.1 audio sound.

Hardcore Computer, Inc. • RSC# 45294 • www.hardcorecomputer.com

ISR subsystem makes things a lot easier

Wouldn't life be easier if engineers could just put an entire already-made subsystem into their design and move on, exponentially speeding development time and ease? Sounds too good to be true, but that's precisely what Mercury Computer Systems, Inc. is offering in the form of its OpenVPX (VITA 65)-based ISR subsystem. The subsystem, already deployed in an unnamed rugged platform, is crafted to execute Processing, Exploitation, and Dissemination (PED) within the ISR realm. After all, this is just the sort of "high-end" image- and signal-processing subsystem that deployed warfighters need to heighten their own situational awareness and provide the parallel data stream computing capabilities paramount to mission success.

The ISR subsystem's primary enabler is Mercury's Ensemble 6000 Series GSC6200 OpenVPX-based GPU processing module. Not only does the OpenVPX form factor make the subsystem wide open for compatibility in cutting-edge rugged defense systems, but GPU incorporation also renders some strong SWaP advantages. GSC6200 features the MxM GPU form factor to expedite ATI or NVIDIA GPU integration or upgrades. Not only that, the ISR subsystem uses open standards-based APIs, simplifying the process of mixing and matching with other wares.

Mercury Computer Systems • RSC# 45292 • www.mc.com

Bumps in the night are no problem for TabletPC

Protecting data from things that go bump in the night — or during the day — is the recently upgraded E100 fully rugged TabletPC from Getac Inc. Designed to meet both IP65 and MIL-STD-810G specifications, the 3 lb TabletPC is well-protected against water, dirt, vibration, weather extremes, and many other harmful-to-electronics variables. It is also intended as an all-in-one ware, eliminating the need for numerous handheld devices for in-field warfighters and personnel in any out-in-the-field industry. And for times when the sun is shining relentlessly on soldiers or industrial workers just trying to get the job done, E100 provides an 800 NITs sunlight-readable display to ease viewing and speed workflow.

But that's just the beginning. Things new and improved by Getac this time around include a faster 1.6 GHz processor, along with a solid state drive and expanded storage capacity. Users will also appreciate the new embedded backup battery, which facilitates spare battery hot-swapping. The real benefit to users is that they don't have to shut down the TabletPC when they replace the battery packs, and yet all data is preserved. Additional plusses of the E100 include its RS-232 serial port, Giga LAN and Bluetooth 2.0, 812.11a/g/n wireless LAN, and optional 3G Network connectivity and/or dual smart card readers. Getac can also lend its "build to order" services to render a portable dual-battery charger or vehicle dock.


Getac Inc. • RSC# 45293 • us.getac.com


XMC card doubles the span of memory lane

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

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

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





Will the real automatic target tracker please stand up?

Clutter might be a nuisance in the kitchen or conference room, but it's even more imperative to clean it up when we're talking about video imaging used for target tracking and detection. Of course, atmospheric challenges such as dust and fog will always be there, but those are no problem for GE Intelligent Platforms' spruced-up ADEPT5000 multi-target video tracker series. Now sporting improved algorithms and processing power, in addition to precision platform dynamics, better optics, and higher-resolution sensors, the rugged ADEPT5000 multi-target video tracker can distinguish decoys and spot coordinated evasive maneuvers involving multiple targets. Obscured targets are no problem either, and ADEPT5000 also aids operators in assessing target priority.

The 54 gram device, designed for military and aerospace electro-optical applications and typically consuming a miniscule 8 to 12 W, utilizes open standards form factor carrier cards to ease integration: 3U VME, 6U VME, 3U VPX, PC/104, and PCI Express. Off-the-shelf PC development tools additionally boost productivity, and graphics rendering affords clear outputs. Digital video input includes GigE Vision, CameraLink, HDSI, and DVI-D. Meanwhile, analog video input comprises NTSC, RS170, PAL, or CCIR. Sensor format auto-detection, context-based settings, and fast-start application modes also decrease development man-hours.

GE Intelligent Platforms • RSC# 45295 • www.ge-ip.com

EMI won't blow a gasket

Electromagnetic Interference or EMI (aka RFI or Radio Frequency Interference) can be annoying when it affects devices such as consumers' cell phones or televisions. But EMI's presence in military embedded electronics can prove deadly or disastrous as it degrades, obstructs, or interrupts electrical circuit performance. We're guessing these issues were what motivated Leader Tech engineers to develop the TechMESH knitted wire gasket, designed to serve as a high-attenuation EMI shield for military-grade enclosures.

Why talk gaskets though, you ask? Perhaps more "vanilla" than some of the other splashier and flashier components of military embedded systems, gaskets such as these are absolutely critical in many DoD systems' performance. And the TechMESH gasket does indeed enhance interference-free, high performance levels: It offers EMI H-field shielding up to 80 dB and E-Field shielding up to 130 dB. And TechMESH doesn't always have to be configured the same old way: The gaskets are available in elastomer or all-mesh incarnations in aluminum, monel, or tin-plated copper-clad steel materials. And you get it your way in yet another way: TechMESH can be vended by cutting to customers' exact requirements or provided on 25-foot spools.

Leader Tech • RSC# 45296 • www.leadertechinc.com



Tiny DC-DC converter goes on power trip

Able to conquer a threefold mission in a single bound, Picor's PI3101 Cool-Power DC-DC converters render output regulation, voltage transformation, and isolation all inside a Power-System-in-Package (PSiP) package measuring a mere .87 inches (l) x .65 inches (w) x .27 inches (h). The PI3101 renders point-of-load isolation and conversion, which in turn eliminates an overall electrical efficiency-degrading stage. And it doesn't end there . . . Designed for an input voltage range of 36 V to 75 VDC, the PI3101 is suited to high-speed server platforms, Power-over-Ethernet, and wireless infrastructure and networking/communications applications such as those used by Command and Control or even DoD offices.

Output includes a regulated 3.3 V at up to 18 A output current. The device also stands steadfast in the midst of input voltage transients up to 100 V for 100 ms, in addition to providing input-to-output isolation of 2,250 V. Another plus of the PI3101, which boasts up

to 87 percent efficiency, is that it can reduce real-estate usage versus discrete-based or sixteenth-bricks designs by 50 percent. Additionally, system designers can position any necessary isolated voltage "at the point-of-load virtually anywhere it is needed." This is true while 48 V is available, and is not contingent upon bus architecture topology. With switching frequencies faster than 1 MHz, PI3101 also renders under-/over-voltage lockout and dual current limit threshold plus programmable features such as soft-start capability, temperature monitoring, and ± 10 percent output voltage trimming.

Picor, a subsidiary of Vicor Corporation • RSC# 45297 • www.picorpower.com


Virtex-6 FPGAs just got a whole lot easier

The military embedded signal processing arena, like all embedded technologies, is constantly "evolving." Of course, that's code for: "Companies have to spend a lot to keep up with the times." Using reprogrammable FPGAs in DSP applications can help companies stay on the cusp of technology by saving some hardware expense, but the programming and development costs associated with FPGAs can be high. That's undoubtedly why Avnet, Inc. created its Xilinx Virtex-6 FPGA DSP Development Kit. Geared toward military, wireless, instrumentation, and many other compute-intensive applications where high-performance DSP is a must, the kit is part of the Xilinx Targeted Design Platform for DSP. The kit comprises Virtex-6 FPGAs, DSP IP, a scalable development board, cables, documentation, and DSP development tools. The kit is even touted to offer up to a 10x productivity boost.

Meanwhile, the kit's key offering is its Virtex-6 DSP Targeted Reference Design, which conceptualizes DSP design flows and methodologies for Virtex-6. The Digital Up Converter (DUC)/Digital Down Converter (DDC) Targeted Reference Design gives users the how-to's for advanced features including time division multiplexing, clock over sampling, and DSP48 slice utilization for high performance. Since it is based on MATLAB and Simulink, the Targeted Reference Design comprises a familiar environment sans RTL. But designers knowledgeable in RTL can also use the Targeted Reference Design's ISE Design Suite and LogicCore DSP IP. Other highlights include testbenches, design synthesis parameters, and Simulink and RTL design source files.

Avnet, Inc. • RSC# 45298 • www.avnet.com





By Chris A. Ciufo, Editor

Can bio-science grow in military applications?



A recent *FORTUNE Magazine* article described Hewlett-Packard's (HP's) Central Nervous System for Earth (CeNSE) network consisting of a trillion tiny Micro-electromechanical Systems (MEMS) scattered around the Earth. The tiny sensors, bundled with wireless transmitters, are envisioned to collect vibration and motion in soil, on buildings, and in any geographical location where scientists want to collect data. Petrochemical companies (such as Shell Oil) and seismologists looking for the next Big One have expressed interest in these sensors, cheaply manufactured in huge volumes originally for HP's inkjet printers. One can easily imagine use in military applications for monitoring runways, no-go interdiction zones, or in perimeter defense – a smaller and less costly alternative to radar, IR, or even CCTV sensors, which all require operator intervention or sophisticated embedded computers to interpret sensor events.

So I got to thinking: Like the HP MEMS, can COTS bio-sensors benefit the warfighter or DoD, or find utility in any homeland security application? I'm not sure, but here are some examples of the "precision plant measurement" instruments manufactured by Camas, Washington-based CID Bio-Science (www.cid-inc.com).

Leaf area meters

Look out your window and you'll see that leaves vary in size, shape, thickness, and color. And thinking back to your high school biology class, you'll recall that a plant's leaves convert sunlight into food energy while emitting healthy oxygen back to the environment. When determining a plant's overall bio-structure and health, scientists need to know the total area of a plant's leaves. But what on Earth would you measure that with?

CID's handheld laser leaf area meter and similar (guts-wise) portable laser leaf area meter use a 650 nm laser diode scanner element to easily identify and calculate leaf area (versus the non-leaf background). For the portable device, simply pulling a leaf through the handheld wand measures area, width, perimeter, shape factor¹, and aspect ratio. Resolution up to 0.025 mm² at +1 percent with a scanning rate of 127 mm/s (palette version) or 0.01mm² at 200 mm/s (wand) provides reasonably high-resolution images and area accuracy of fine-veined leaves and even evergreen needles.

How might these battery-operated instruments be used in military applications? When I asked the engineers at CID, they really had no idea. (They're biologists after all.) But while perusing the DARPA website, I saw a number of R&D sensor programs that seem related to this kind of instrument. For example, sensors that monitor a soldier's or Marine's health might measure hair thickness or skin surface area. It's possible that measuring the amount of Nuclear, Biological, Chemical (NBC) agent on the skin could give a prognosis to those exposed to hazards. A leaf area sensor could easily be used to measure body hair or skin – and feed that information in as part of an overall prognosis.

As well, these instruments can measure any kind of area "stain" – from blood or chemical residue on a briefcase to oil washed onshore and clinging to saw grass in the Gulf of Mexico. And it seems to me that cheaply measuring ammonium nitrate residue on objects might provide advanced warning of bomb threat material. Sure, fancy spectrographic analyzers and gas chromatographs easily provide chemical concentration charts, but they aren't handheld and surely aren't cheap.

Digital root imager

A derivation of the leaf imager, the *in situ* root imager slides into a clear tube in the ground and non-destructively measures plant roots and other buried objects. For botanists, this application showcases a plant's health by examining the root structure without having to yank the plant from the earth. The 360-degree scanning sensor records full-color images at 1,200 dpi, with a 21.6 x 19.6 cm image containing up to 188 million pixels.

In military applications, one could envision using a similar application to look for visible objects in soil, liquid, or other hard-to-reach places where a low-light video image would record results poorly versus a laser-scanner-based imager. As before, ground-penetrating radar or sonar also provides non-destructive (though only indirectly representative) object images. But portability and sheer low cost might make the digital root imager much more appealing and deployable to warfighters hunting for IEDs or other paraphernalia.

Digital plant canopy imager

This last device is pretty unique, though I struggle to find direct applicability in military applications. Using an upward-looking Pentax fish-eye 150-degree wide-angle lens on a simple free-wheeling gimbaled boom-based video camera at 768 x 494 pixels, sophisticated software interprets overhead forest canopy images and calculates the area of coverage and other biological metrics including Leaf Area Index (LAI) and Photosynthetically Active Radiation (PAR) levels. I think the magic here is the COTS software that digitizes and manipulates captured images based upon a simple azimuthal and zenith angle entered by the operator. Sunshine that makes it through is counted as "not canopy," and color shades can be interpreted as vertical canopy density. The gap-fractional inversion procedure algorithm is at the heart of this system, providing an amazingly accurate estimate of canopy coverage².

How might this apply to defense applications? Looking for patterns in clouds or aerial formations perhaps? Passively estimating vertical distances without signal emanation? Performing ceiling tile reconnaissance? I'm interested in your ideas about these three technologies – and I'll publish them. Let me know.

Chris A. Ciufo, Editor
cciufo@opensystemsmedia.com

¹ Shape factor = ratio of area to perimeter. Aspect ratio = ratio of length to maximum leaf width.


² Courtesy info from CID: J.M. Norman and G.S. Campbell (1989) Canopy Structure. In: *Plant Physiological Ecology: Field methods and instrumentation*. (eds. R. W. Pearcy, J. R. Ehleringer, H. Mooney, and P.W. Rundel), Chapman & Hall, London and New York, pp. 301-325.

Welcome to the Knowledge Bank.

A wealth of free COTS white papers and tutorials.

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

www.ge-ip.com/cots-papers




imagination at work

SPEED YOUR TIME-TO-MARKET

YOUR TOTAL SOLUTIONS PARTNER



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




- 2 Radar Processing**
Digital Signal Processors
CHAMP-AV5



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


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



- 4 Graphics Display**
Video Input & Output
XMC-710



- 5 System Connectivity**
Ethernet Switches
VPX6-684 FireBlade



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

**CURTISS
WRIGHT** Controls
Embedded Computing

sales@cwembedded.com

cwembedded.com

BOARDS & SYSTEMS

ABOVE & BEYOND