Military EMBEDDED SYSTEMS

Jan/Feb 2013
Volume 9 | Number 1

MIL-EMBEDDED.COM

Next-gen radars condense sensor-to-shooter cycle

10 GbE enables real-time remote desktops for C4ISR
By Jeff Malacarne, Barco

FPGAs, commercial processors drive radar signal processing

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

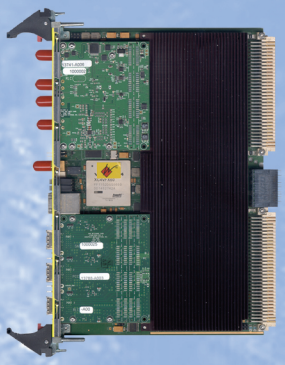

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

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

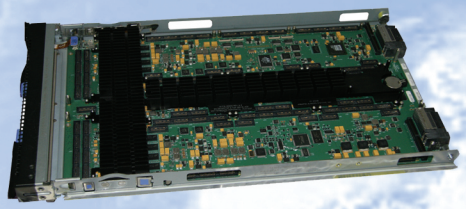

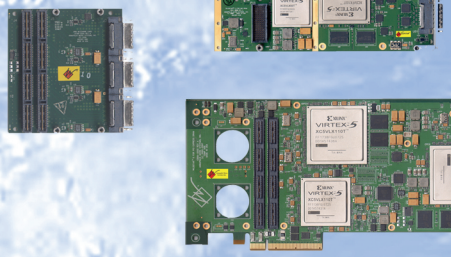

Annapolis
Micro Systems

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

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

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

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

A 1.66GHz Intel® Atom™ Mil/COTS SBC of *EPIC* Proportions

The EPX-C380 is a rugged single board computer that provides an open powerful platform for harsh Mil/COTS applications. Powered with either a single or dual core processor, it has a wealth of onboard I/O plus expansion options. Also it supports Linux, Windows® XP embedded, and other x86 real-time operating systems.

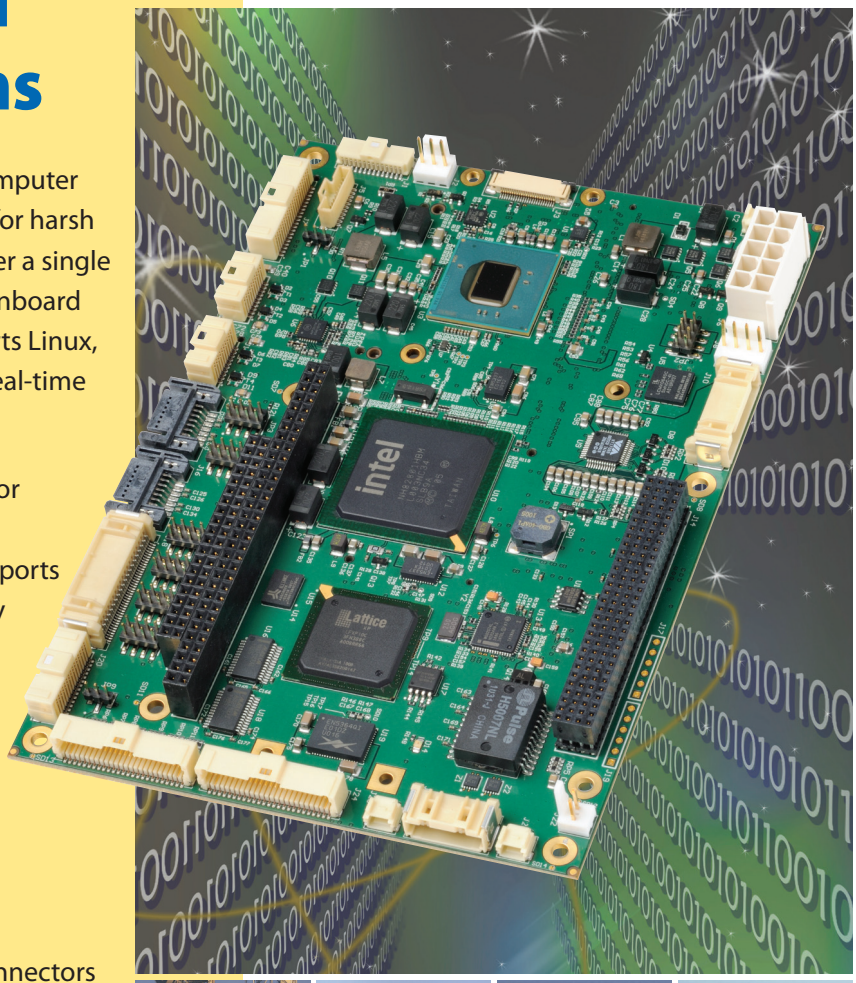

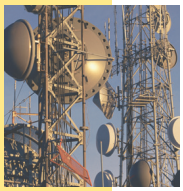

- 1.66GHz Intel® Atom™ N450 single core or D510 dual core processor available
- Embedded Gen 3.5+ GFX video core supports CRT and LVDS flat panels simultaneously
- Custom splash screen on start up
- Optional 1MB of battery backed SRAM
- Two Gigabit Ethernet ports
- Two SATA channels
- Eight USB 2.0 ports
- Four serial RS-232/422/485 channels
- 48 bi-directional digital I/O lines
- CompactFlash (CF) card supported
- MiniPCIe and PC/104-*Plus* expansion connectors
- High Definition audio supported (7.1)
- WDT, RTC, LPT, status LEDs, and beeper
- +5 volt only operation
- EPIC sized: 4.5" x 6.5" and RoHS compliant
- EBX form factor SBC also available
- Extended temperature operation
- Responsive and knowledgeable technical support
- Long-term product availability
- Quick Start Kits for software development

Contact us for additional information, pricing, and custom configurations. Our factory application engineers look forward to working with you.

Call 817-274-7553 or Visit www.winsystems.com/EPX-C380MR

Ask about our 30-day product evaluation

WinSystems®
An Employee-Owned Company

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

www.mil-embedded.com

Military EMBEDDED SYSTEMS

ON THE COVER:

Top photo: The AN/TPQ-53 (Q-53) counterfire target acquisition radar from Lockheed Martin provides 360-degree protection from mortar, rocket, and artillery fire – tracking targets in their upward path and enabling a counterattack before the projectile makes impact. Photo courtesy of Lockheed Martin

Bottom photo: Radar onboard U.S. Navy Aegis-class ships uses PowerStream technology from Mercury Systems for radar signal processing. Photo courtesy of Lockheed Martin

January/February 2013 | Volume 9 | Number 1

COLUMNS

E-CASTS

http://ecast.opensystemsmedia.com

Radar and electronic warfare drive signal processing innovation
Feb. 27, 12pm MST
Presented by: Altera, Pentek, Inc., Mercury, Rohde & Schwarz

Managing safety and security through open architectures
March 26, 11am MST
Presented by: Esterel, GE Intelligent Platforms, RTI, Wind River

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

Published by: OpenSystems media®

ISSN: Print 1557-3222

@military_cots

www.linkedin.com/groups/Military-Embedded-Systems-1864255

SEARCHING FOR Intel® Core™ i7 PROCESSOR SOLUTIONS?

Your search is over!

CompactPCI • COM Express • VME • PrPMC/XMC • VPX • Custom

X-ES delivers the latest 3rd generation Intel Core i7 processor solutions on the widest range of standard and custom form factors in the industry. With our proven record of meeting aggressive schedules, you can count on X-ES to deliver Intel Core i7 processor solutions on time, and with unparalleled customer support. **Call us today to learn more.**

You need it, we have it! That's Extreme.

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

ADVERTISER INFORMATION

Page	Advertiser/Ad Title
29	**Aitech Defense Systems** – Our technology investments protect yours
41	**Alphi Technology Corporation** – Looking for I/O for your mission computer?
37	**Altera Corporation** – Altera Stratix V GS FPGAs
42	**Altera Corporation** – Enabling advanced designs for radar and electronic warfare systems with FPGAs
2	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
37	**Annapolis Micro Systems, Inc.** – WILD storage solution
37	**Annapolis Micro Systems, Inc.** – WILD OpenVPX switch
13	**Blue Pearl Software** – Is your FPGA design signoff quality?
43	**Creative Electronic Systems** – Prospects of COTS products for future radar systems – challenges and opportunities
40	**DiSTI** – Interface with reality
17	**Elma Electronic – Systems** – Small is the new big!
46	**Embedded World** – It's a smarter world
11	**Enginasion Inc.** – Engineering imagination
15	**Excalibur Systems, Inc.** – MACCsimize your data conversion ability
47	**GE Intelligent Platforms, Inc.** – High-performance radar with out-of-box readiness
37	**GSI Technology** – Ultra-high performance memories from GSI Technology
14	**Innovative Integration** – Radar wireless solutions
19	**Interface Concept** – Switches and IP routers
24-25	**Mercury Systems** – Today we become Mercury Systems
48	**Pentek, Inc.** – Critical recording in any arena
40	**Phoenix International** – Introducing RPC24
37	**Rakon** – Digital Pulse Compression Sub-Systems (DPCSS)
44	**Rakon** – Radar retrofits, subsystems are key in today's market
37	**Spectracom Corporation** – Fast and affordable microwave frequency analysis
3	**WinSystems, Inc.** – A 1.66 GHz Intel Atom Mil/COTS SBC of EPIC proportions
5	**X-ES** – Searching for Intel Core i7 processor solutions?

Military EMBEDDED SYSTEMS

OpenSystems media®

Radar, electronic warfare steadier parts of military market

By John McHale, Editorial Director

"Uncertain" was the keyword most of our readers and advertisers attached to the outlook for electronics funding and defense spending worldwide last year. Threats of sequestration and a stagnant global economy fueled the doubt, which unfortunately continues today. The sequestration threat was only postponed, not removed. However, if Department of Defense (DoD) spending were to be dramatically slashed in such a scenario, many feel sensor procurement and development – specifically radar and electronic warfare technology – would remain strong.

As we reduce our military footprint worldwide, we will rely more than ever on sensor data from Intelligence, Surveillance, and Reconnaissance (ISR) systems onboard unmanned aircraft, in satellites, and on the ground. Accurate, real-time information from radar and other sensors is crucial to improving situational awareness. Improving bandwidth in military systems is of equal importance because it will help reduce the sensor to shooter cycle – the time it takes for sensor data to reach a shooter so it can eliminate a threat.

High-performance radar systems will need to provide much of the key intelligence to the warfighter. Radar requirements coming out of the DoD are calling for detection of small, slow-moving objects such as dismounted individuals on the ground and obscure targets in littoral waters. Market numbers show how much activity occurred in military radar contracts just in 2012 alone – indicating radar technology might be the best bet for military electronics investment over the next few years. According to Brad Curran, an analyst with Frost & Sullivan in San Antonio, TX, U.S. radar contracts for 2012 totaled $4.26 billion, an increase of almost a $1 billion over 2011 and about $1.5 billion over 2010 numbers. The U.S. radar market has shown significant increases the past three years, driven mostly by missile defense and the U.S. Army's counter-battery systems, he adds. "There were 457 companies total with U.S. contracts in 2012 radar technology and/or services. Raytheon is by far the largest radar producer in the U.S., winning contracts totaling $1.39 billion in 2012, while Lockheed Martin and Northrop Grumman were second and third respectively. Raytheon has been number one the last three years."

... U.S. radar contracts for 2012 totaled $4.26 billion, an increase of almost a $1 billion over 2011 and about $1.5 billion over 2010 numbers.

The robust radar market and the focus military embedded system suppliers have on radar products led us to gear our content this issue toward radar technology. Our Special Report section's lead article covers general radar trends such as dismount detection, the Lockheed AN/TPQ-53 (Q-53) counter-fire target acquisition radar (pictured on our cover this edition), and also looks at how Lockheed Martin engineers are mitigating the challenges wind farms create for long-range radar systems. Also included in the Special Report section is an article from Dan Kinney and Andy Johnston of Parker Aerospace on how two-phase cooling meets thermal management challenges in modern radar applications.

The lead article in our Mil Tech Trends section looks at how radar and electronic warfare designers are getting around poor-bandwidth technology by placing the processing power of a ground station next to the sensor on the platform. This paradigm enables data analysis and exploitation to be performed before data is sent to the ground on slow RF links, enabling real-time actionable intelligence to be sent to the warfighter. The article covers how FPGAs and commercial processors are driving radar and electronic warfare sensor processing designs.

Mil Tech Trends also includes two articles on Modular Open Systems Architecture, or MOSA. In the first article, Michael Stern of GE Intelligent Platforms talks about how to exploit MOSA software for high-performance ISR and signal applications. The second article, written by Chris Lewis and Alton Graves of Mercury Systems, covers how MOSA enables flexibility and cost savings for electronic warfare digital receivers. Also in this section, Jeff Malacarne of Barco contributed a piece on how "10 GbE enables real-time remote desktops for C4ISR." Meanwhile, in our Industry Spotlight section, Rodger Hosking of Pentek discusses how new Virtex-7 FPGA technology boosts radar performance.

In future radar designs, keep an eye out for InfiniBand. The switched fabric created a buzz at the Embedded Tech Trends (ETT) conference in Long Beach, CA, in January, where CSPI, GE Intelligent Platforms, Mercury Systems, and others in attendance launched new products based on the standard. InfiniBand along with 40 GbE will be implemented for radar and other intensive signal processing applications as both have long-term road maps, speed, and low latency, noted Marc Couture of Mercury Systems during ETT.

John McHale
jmchale@opensystemsmedia.com

Radar's brave new world: OSA

By Charlotte Adams

A GE Intelligent Platforms perspective on embedded military electronics trends

High-Performance Embedded Computing (HPEC) is paramount within the tight real estate confines of defense systems today. HPEC's high-scale parallel processing technology, based on Open Systems Architectures (OSAs), also increases the power of sensor applications such as radar.

But high performance has not always gone hand in hand with openness. Until recently, the best results in radar were achieved through proprietary software and hardware, which customers preferred even though it might lock them into particular vendors. Programming was difficult because the parallel execution so necessary for radar processing had to be built in at the application level. If six processors were required, six programs had to be written, each of which made calls to the math and communication libraries. These programs were difficult, time consuming, and expensive to write as well as challenging to upgrade. But these days OSA- and open standards-based software tools can change this paradigm.

Radar: Shift to OSA

As technology evolved – spurred by the consumer market – Commercial Off-the-Shelf (COTS) and OSA hardware and middleware products (as well as the toolkits supporting them) reached the volume and acceptance level required to support shifts in the design of compute-intensive applications such as radar. Open architecture means that necessary programming interfaces are available and published, giving rise to a greater wealth of competing solutions, easier interoperability, and lower life-cycle costs as compared to proprietary solutions. High-performance middleware also means that some of the parallelization required by sensor applications – and previously done manually – can now be executed automatically through such aids such as signal processing and math libraries.

At the programming level, OSA approaches have caught up with proprietary systems through middleware such as VSIPL++, the latest version of the Vector Signal Image Processing Library. Application development toolkits have incorporated such middleware and other advances to make it easier for customers to use these OSA technologies.

Pentagon buy-in

The Pentagon acknowledged these trends in 2010, when Under Secretary of Defense Ashton Carter called for the use of OSA as part of acquisition reform. This high-level support has flowed down to new procurements and upgrades. After all, since new hardware will make proprietary applications obsolete, customers might as well migrate to open systems middleware, too. Additionally, radar systems already had begun to reflect the technology shift before the Carter memorandum. MIT Lincoln Lab, for example, has employed the Radar Open System Architecture (ROSA), involving COTS hardware and open systems. ROSA has been applied in the prototype Cobra Gemini radar and the modernization of radars at the Kwajalein Missile Range.

Figure 1 | GE's AXIS software development environment enables the rapid development of OSA radar applications on hardware platforms like the HPEC Starter System.

The emphasis on OSAs in radar development continues. Requests For Information (RFIs) for new radar systems and upgrades are asking for open standards middleware technologies such as VSIPL++ and the open source Linux operating system.

Software development toolkits with OSA

One obvious way to lessen the difficulties of moving from a proprietary to an OSA-based radar system is to use software development suites with built-in OSA standards. These frameworks help developers meet deadlines by using math and signal processing functions such as the Fastest Fourier Transform in the West FFTW), VSIPL and VSIPL++, and data movement libraries such as Message Passing Interface (MPI) and Data Distribution Service (DDS).

Toolkits for radar system migration or design are available from multiple vendors. GE's Advanced Multiprocessor Integrated Software (AXIS) environment, for example, includes math libraries, visualization tools, and interprocessor communications optimizers that run on multiprocessor boards – even hybrids combining traditional and highly parallel chip architectures (Figure 1). The software suite also gives programmers access to standard middleware such as MPI, OpenFabrics Enterprise Distribution – or OFED, and VSIPL++. (MPI integration into the AXIS toolkit is expected in the near term.)

Toolkits also can accelerate radar application development by helping to partition code across potentially hundreds of processors. Then programmers can assess software performance visually, viewing a system as it is running to check the efficiency of the code and whether it's meeting its real-time constraints.

Low pain, high gain

The transition from proprietary systems to open systems architectures, assisted by the latest standards-based software tools, should be relatively painless for radar systems engineers and should lead to greater system performance, programmer productivity, and cost-effectiveness.

SiGe-based ADCs/DACs double COTS Electronic Warfare processing performance

By David Jedynak

An industry perspective from Curtiss-Wright Controls Defense Solutions

New ADC and DAC technology based on Silicon-Germanium (SiGe) promises unprecedented levels of functionality and capability for demanding signal processing applications. These new devices, which bring the advantages of SiGe to rugged deployed military systems for the first time, can deliver 2x the performance of currently available ADC/DAC devices, establishing a new class of processing performance for defense and aerospace applications. When deployed on open architecture platforms utilizing OpenVPX COTS boards with latest-generation FPGAs, the new SiGe-based technology will provide an ideal platform for radar, SIGINT, and EW applications.

New levels of performance

These types of EW applications require a balance between speed and resolution. Compared to earlier designs, these new SiGe-based ADCs/DACs, supplied by vendors such as Tektronix, feature higher sample rate performance. They leverage high-performance data conversion techniques to optimize device performance characteristics such as calibration, power, and signal/noise ratios. These ADCs, for example, deliver the best speed and Effective Number of Bits (ENOB) currently available from a commercial device. An additional advantage of SiGe-based devices is their low latency, an important feature for bandwidth-sensitive EW applications.

Until recently, COTS ADC devices on the market have topped out at 3 GSps at 8 bits of resolution. In the past few years though, we've begun to see devices that can perform at up to 6 GSps at 8 bits. The new SiGe-based generation of ADCs is delivering the next big performance leap, doubling bandwidth speeds up to 12 GSps. For EW applications, the benefit is straightforward: The higher sample rates and associated bandwidth ensure better spectrum coverage and improved Probability Of Intercept (POI) for signals of interest. In addition, the performance of these 8-bit parts surpasses off-the-shelf 10-bit ADCs in terms of Spurious Free Dynamic Range (SFDR). SFDR is a measure of the performance of the ADC, and higher SFDR ensures improved identification of signals of interest in a crowded spectral environment. Typically, a first pass of the spectrum segment is done at high bandwidth to pull in as much data as possible to obtain areas of interest to analyze, after which a higher-resolution, lower-bandwidth solution is leveraged to focus on specific targets. As warfighters see a far greater range of the spectrum, more lives are saved and mission success probability is increased because of faster, more accurate identification of threats and improved response options.

With this new generation of ADCs and DACs, EW system designers get increased bandwidth with sufficient resolution. It's a win-win with high-quality signal identification and improved immunity from noise that supports real-time analysis of larger amounts of data. While it's possible to obtain ADCs that operate at 14- to 16-bit resolution rates, these devices typically sample in the hundreds of Msample range, far below the 12 GSps rates at 8 bits now reachable.

SiGe-based ADCs/DACs also deliver lower power performance (as measured in Watt/GHz). In addition to their higher speed and lower power, these ADCs/DACs also offer reduced leakage current and less sensitivity to temperature fluctuation, which becomes more critical in EW electronics as process architectures shrink.

Faster I/O devices meld with OpenVPX and FPGAs

These faster ADCs and DACs can be readily built into rugged open architecture OpenVPX-based EW systems using FPGAs to perform high-speed algorithm processing in the digital domain; this paradigm minimizes the need for performing downconversions or other filtering stages that would typically be handled in an external analog tuner logic, slowing down performance and requiring additional on-board components that use valuable board real estate and add unwanted heat.

Figure 1 | Tektronix Component Solutions' SiGe-based ADC provides 12.5 GSps digitization performance to the TADF-4300 module.

The new ADC/DAC components can be deployed on OpenVPX hosts so that the system designer has the flexibility to swap out different front-end configurations as required while maintaining a common back-end and software interface to the FPGA to address different types of applications. An example of an OpenVPX board that delivers the latest generation of devices is Curtiss-Wright's rugged CHAMP-WB-DRFM 6U card set, combining a Tektronix TADF-4300 module featuring a SiGe-based 12.5 GSps 8-bit ADC and a 12 GSps 10-bit DAC (Figure 1), on a Xilinx Virtex-7 FPGA-based 6U VPX card, the CHAMP-WB. This modular design approach actually provides designers with two levels of modularity or reconfigurability: The first level is the ability to swap out different mezzanines as needed, and the second level is the inherent reconfigurability of the FPGA itself. This serves to benefit today's cutting-edge EW applications.

David Jednyak
CTO/COTS Technologies
Curtiss-Wright Controls
Defense Solutions
www.cwcdefense.com

Static analysis exposes latent defects in legacy software

By Paul Anderson

When migrating from a legacy software-based system to new technology, it is important to be able to reuse as much code as possible. Even if such code has been tested thoroughly and has proved reliable in practice in the old system, it might still contain latent bugs. Those bugs might have never been triggered in the legacy system because of very specific properties of that system such as the toolchain used to compile the code, the processor architecture, or the host operating system. When ported to a new system where those properties are different, the latent defects might be manifest as harmful bugs. But the good news is that advanced static analysis tools can flush out these latent defects to help combat the challenge.

Updating the system, revealing coding flaws

One of the most important motivations for migrating legacy systems is to take advantage of advances in hardware technology since the original system was first deployed. Probably the most common benefit is increased performance because of the adoption of a newer and faster processor. This is also the single most significant change from the perspective of the code. The new processor can have a different bit width, or endianness, and the number of available cores can be different. During code porting from the old platform to the next, much of the recoding effort will go into adapting the code to those differences.

Compilers, toolchains, and latent bugs

There are many other less obvious differences than implementing a new processor, and these subtle nuances can be easy to overlook. Take, for example, the toolchain used to compile the code. On the face of it, this should not make much of a difference. After all, if the code was written to be ANSI C compliant, and if the compilers claim to support ANSI C, then surely the code will have the same semantics when compiled by either compiler? Unfortunately not. The C and C++ standards are riddled with clauses that are "compiler dependent," meaning that the standard does not dictate exactly how to compile certain constructs, and the choice is up to the compiler writer. Many of these are obvious and well-known to programmers, such as the order in which operands are evaluated, but others are very subtle. A latent bug might be harmless on the legacy system because the compiler chooses to compile it a particular way, but hazardous on the new system because the new compiler makes a different choice.

> *"Compilers are programs too, of course, and are themselves not free of defects. A recent study of C compilers found code generation defects in every compiler they tested."*

Compilers are programs too, of course, and are themselves not free of defects. A recent study of C compilers found code generation defects in every compiler they tested[1]. The treatment of the volatile keyword, which is of vital importance in embedded safety-critical software because it is frequently used to read sensor data, is especially prone to compiler errors that cause changes to sensor values to be silently ignored. A program's correct functioning might even come to rely on such flaws.

Another danger zone: Standard libraries

Another subtle software migration difference that can cause latent defects to become hazardous involves the standard libraries that interface to the operating system. One might hope that such libraries would be consistent across platforms, but this is rarely so. The most significant difference is with respect to error handling. A new platform can have completely different failure modes than the legacy platform, and the code might need to be changed to be able to handle those differences. Worse still, error cases appear to be very badly documented, according to a recent study[2].

Static analysis wins out, complements traditional testing

Clearly, any legacy migration project must include a significant amount of time for testing the software in its new incarnation. However, test results are only as good as the test inputs. If a test case fails to exercise a path on which a bug occurs, that defect might go undetected. Generating new test cases is expensive too. Hence, a sensible tactic to flush out these latent defects is to use advanced static analysis tools as part of the legacy transformation effort. Such tools are capable of finding defects such as those described herein, including those that depend on platform subtleties. They are especially good at finding concurrency defects such as data races that are exceedingly difficult to find using traditional testing methodologies. They are also good at finding instances of code that, while not definitively wrong, is highly correlated with errors or is particularly risky when ported to a different environment.

References

[1] Finding and Understanding Bugs in C Compilers, Xuejun Yang, Yang Chen, Eric Eide, John Regehr, 2011. www.cs.utah.edu/~regehr/papers/pldi11-preprint.pdf

[2] Expect the Unexpected: Error Code Mismatches Between Documentation and the Real World, Cindy Rubio-González, Ben Liblit, 2010. www.eecs.berkeley.edu/~rubio/includes/paste10.pdf

***Paul Anderson** is VP of Engineering at GrammaTech and can be contacted at paul@grammatech.com.*

ENGINASION
Engineering Imagination

Company Fact Sheet
2013

Enginasion™ solves some of the most difficult technical problems involving process control, manufacturing automation, and product development, providing clients with electronic, mechanical and software engineering services from concept to production. For two decades, Enginasion (formerly Industrial Automation Systems) has been an invaluable partner to companies in the Industrial, Medical Technology, Military and Pharmaceutical sectors that need a sophisticated yet affordable engineering resource to help them overcome critical hurdles related to automation, control electronics, embedded software, and/or integrated product development. Since 1991, the Company has been providing key elements of successful R&D projects, manufacturing processes, and new high-tech products at innovative companies in the Northeast U.S. Enginasion provides the following resources:

Test calibration and equipment designed with the operator in mind. Virtual instrument software using touch screens makes operator control fast and inexpensive. We integrate off-the-shelf and custom electronic and physical components with personal computers to create sophisticated test equipment at minimal cost.

Product development, from tiny watch-sized electronics to highly specialized powder dispensing devices to large networked systems, Enginasion can handle the electronic design and PCB layout, real-time software, packaging, production prototypes, low-volume production test equipment and documentation for large-scale production.

Process control systems for test or manufacture—measuring and controlling parameters such as machine movement, particle density, electrical charge, temperature, humidity, pressure, flow rate, Ph, viscosity, and droplet size, using embedded controllers and PLCs communicating via CAN, I2C Modbus, Bitbus, Ethernet, Firewire, Intellution or any other bus or network protocol.

Quality control systems that turn quality measurements into real-time or manual process feedback, using custom data-gathering electronics, bar-coding, touch memories, RFID, statistics, reports and real-time process controls.

Clients of Enginasion include "big" and "little" companies—some of whom have in-house engineering resources, and some of whom do not. Enginasion has helped clients in the Industrial, Medical Technology, Military and Pharmaceutical sectors to design and build hardware and software to monitor, measure and control mechanical and electrical systems. Enginasion specializes in solving the most difficult automation problems, applying ingenuity, in-depth understanding of physical phenomena and materials, and wide-ranging expertise in digital and analog electronics, software, user interfaces, networks, fluid controls, motion control systems and measurement techniques.

CONTACTS & KEY INFORMATION

Headquarters	32 Almanac Way West Boylston, MA 01583 **+1 508 835 6353**
Website	www.enginasion.com
Sectors Served	Industrial, Medical Technology, Military, Pharmaceutical
Engineering R&D Resources	Controls...Sensors...Mechanics... Networks...Software...Patents... Electronics Engineering
Some of the Select Clients Served	GE; The L.S. Starrett Company; Subliminary Artworks; Renaissance Electronics; IMA-Nova® Packaging; MannKind
Sales Contact	**David Bonneau, CEO** **+1 508 835-6353** **sales@enginasion.com**
Media Contact	Ronald C. Trahan, APR Ronald Trahan Associates Inc. +1 508 359 4005, ext. 108

ENGINEERS WITH IMAGINATION

David Bonneau, BSEE
Chief Executive Officer
Skillsets: embedded development; metrology; sensors; process controls; systems arcitecture

Roy Greenwood, BSEE
VP, Electronic Software & Hardware Design
Skillsets: embedded systems, networking, electronics, systems software and architecture

EXAMPLES OF CLIENT SOLUTIONS BY ENGINEERS WITH IMAGINATION

Test Calibration and Equipment

Machine vision gauging • Multi-axis measuring machines • Spark gap analyzer, 32-channel, 1KV • Touch memory read/write probe • Watt-hour meter calibration/hi-pot/continuity test • Corrugation dimensional profiling • Alcohol sensor calibration • Digital caliper functional test, LCD, and calibration.

Product Development

PC board text fixtures and software • Four-button mouse controller • Millimeter wave camera • Piezoelectric fluid nozzle control • Portable pepper spray delivery system • Wristwatch alcohol sensor/data logger • electronic nose • Multi-axis helmet shock data logger • Wireless hand tools • Portable, residential electrolyzer.

Process Control Systems

Yarn tension velocity profiler • Fluid bed particle dispersion controls • Humidity chamber control • Low pressure fluid controls • Metering pump final test systems • High-power electrolyzer controls • PWM to Analog to PWM controller • Embedded software validator • Microgram powder dispenser array.

Quality Control Systems

Electrolysis control systems • Hi-speed fiber analysis equipment • 0.001" grinding wheel gage • Sub-micron particle size analyzer • Vision-based LCD quality tester • Hi-capacity tablet imaging counter • Cosmetic tube label registration gauge • Military electrolyzer training simulator • Hi-pot test enclosures • Lab data-logging systems.

Next-gen radars: Seeing through the clutter

By John McHale, Editorial Director

Next-generation radar systems are improving accuracy in high-clutter environments such as littoral waters, wind farm locations, and slow-moving ground objects – and at faster than ever speeds, reducing the sensor to shooter cycle.

The TPS-77 portable long-range radar from Lockheed Martin Mission Systems & Training is mitigating, via sophisticated signal processing techniques, many of the false alarms caused by wind farms.

Military radar systems remain ever vigilant in their mission to detect superfast, incoming missiles and supersonic enemy fighter jets. High-speed threats are not going away, but military officials today also want radar systems to detect slower-moving objects on the ground or in shallow, littoral waters, requiring complex algorithms and signal processing techniques.

"Radar capability demands are trending toward dismount detection technology to detect people moving on the ground," says John Fanelle, Program Director, Radar Systems, Reconnaissance Systems Group at General Atomics Aeronautical Systems Inc. (GA-ASI) in San Diego. "We have incorporated this capability into our Lynx Block 20A Multi-mode Radar and also added a maritime capability to cross-cue the Electro-Optic/Infrared (EO/IR) ball on the Predator B/MQ-9 Reaper Remotely Piloted Aircraft (RPA) fleet to detect shorter-range maritime targets. A maritime mode for the radar is currently being developed to detect specific small items in the water. For the maritime mode we can leverage different algorithms that are available now. It's just a matter of getting them to work in real time."

Accurate, real-time information from radar and other sensors brought together in one picture helps save lives by reducing the sensor to shooter cycle – the time it takes for sensor data to reach a shooter, such as a UAS, fighter jet, tank, sniper, and so on, so it can eliminate a threat. "The trends are to more automation, smaller crews, radars that do multiple things, and radars that can talk to other sensors on the battlefield," says Lee Flake, Program Director for counterfire target acquisition radar programs at Lockheed Martin Mission Systems and Training. The AN/TPQ-53 (Q-53) counterfire target acquisition radar from Lockheed Martin is that type of radar. The system – which has 360- or 90-degree modes – reduces sensor to shooter time by acquiring the projectiles such as mortars, rockets, and artillery while they are in the air, and sends information on their point of origin immediately, enabling it to destroy the object before it hits its intended target, Flake adds. "It all takes place in seconds."

The system, already operating in Iraq and Afghanistan, can be rapidly deployed, automatically leveled, and remotely operated as far as a kilometer away with a laptop computer or from a fully equipped climate-controlled command vehicle. "The radar uses a custom signal processing system designed by our partner, Syracuse Research Corp. (SRC) in Syracuse, NY," Flake says. "We designed the wireless component for the system. This approach is cutting edge and what the Army wants their sensors to do. When designing it, we started thinking 'What else can we add to this radar – what technology can we insert?' We are only beginning to attack the capability of this radar, which will be around for decades."

Adding capability without overhauling the system enables cost-effective and quick enhancements for meeting urgent requirements. "The beauty of our Lynx Block 20A is that it has enough capability to add these modes without redesigning the radar," Fanelle says (Figure 1). The multifunction radar functions in high-resolution Synthetic Aperture Radar (SAR) and Ground Moving Target Indicator (GMTI) modes. Lynx provides photographic-quality images through rain, clouds, rain, smoke, dust, and fog, day or night. It flies on manned and remotely piloted platforms and has a range as long as 100 kilometers. In radar systems, situational awareness also is enhanced through 3D images. "Lynx has the ability to do 3-Dimensional (3D) targeting using radar imagery," Fanelle says. "With 3D targeting, you need longitude, latitude, and altitude. If you look at a target such as a ship at sea, longitude and latitude would suffice, but if the target is also high up in the mountains, on a plane, or a plateau, then you need to factor in altitude as well to properly fix the crosshairs on the target.

Very high precision is associated with 3D targeting, which enables the use of small weapons that minimalize collateral damage."

Military radar users want quickly deployable systems that have more capability in smaller packages, that can run on less power, and that have more agility in how they are deployed – airborne, ground vehicle, man portable, and so on, says Pierre Poitevin, General Manager at FLIR Radars in Montreal. Time to market can be a challenge today because radar products are more complex and to deploy a system as quickly as possible requires a delicate balance between functionality in the product versus how soon a customer might want it. FLIR engineers developed a dual-mode ground surveillance radar – the Ranger R3D – that has a fast camera mode and a Doppler mode in a small form factor – about 15 inches in height by 18 inches in diameter and less than 30 lbs. It consumes only 30 W and can be used in ground vehicles or man-portable applications for persistent surveillance.

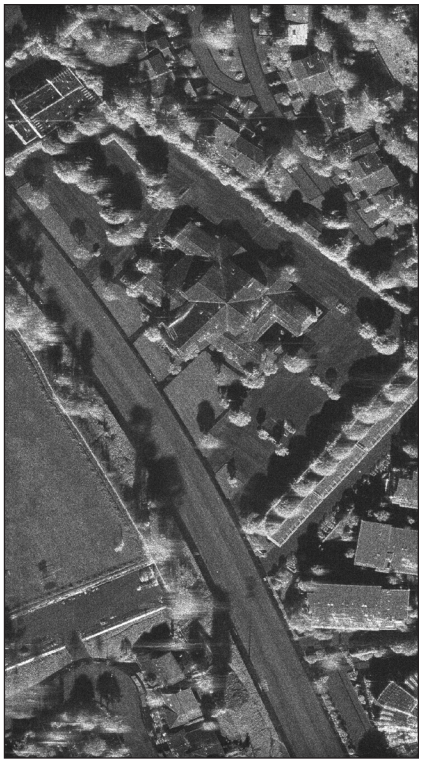

Figure 1 | Photographic-quality imagery is generated by the multimode Lynx Block 20A radar system from General Atomics Aeronautical Systems, Inc. (GA-ASI).

Open architectures

Quickly deployable systems with an architecture that can accommodate future capability upgrades need to leverage Commercial Off-the-Shelf (COTS) hardware and software as well as common standards. The use of common standards with products with high Technical Readiness Levels (TRLs) also reduces risk by encouraging interoperability.

"We leverage open architectures, and our radar compute engines use COTS components," Fanelle says. "Graphics Processing Units (GPUs) are creating a compute engine in the gaming industry that can do amazing stuff, and we are evaluating these devices to harness some of that power. The one thing the gaming industry doesn't have is environmental restrictions on its devices. We need to take a gaming GPU and make sure it can function within the environmental constraints in military applications."

"Implementing open architectures and common standards as part of our design process has allowed us to greatly reduce the number of parts for the same mission we were doing 30 years ago," says Rick Herodes, Business Manager for FPS-117/TPS-77 Radar at Lockheed Martin Mission Systems & Training in Syracuse, NY. "All three of our surveillance radar products – TPS-77, the TPS-59 Marine Expeditionary Radar, and the FPS-117 long-range surveillance radar for the Air Force all use the same signal processing architecture. Treating all three as a product line enables us to leverage investment and support costs across multiple customers. For example, we use exactly the same parts in the TPS-77 as in the other radars such as the WILDFIRE FPGA boards from Annapolis Microsystems." (For more on COTS signal processing trends, see this edition's Mil Tech Trends coverage, beginning on page 20.)

Mitigating wind farm challenges

High-clutter littoral waters are not the only tough environment vexing radar designers. The areas around wind farms, cause quite a few headaches for radar operators. Engineers at Lockheed Martin Mission Systems & Training say their TPS-77 portable long-range radar is mitigating many of the false alarms caused by wind farms through its architecture and signal processing techniques that leverage COTS computing systems.

"The basic problem with wind turbines is they are large and keep getting larger, are semi-metallic to protect the turbines from getting struck by lightning, and are high velocity. With the new generation of turbines, blade tip velocity can approach hundreds of miles an hour," says Chris Atherton, Technical Director for Long-Range Radar at Lockheed Martin Mission Systems & Training in Syracuse, NY. "These are all characteristics radar systems are designed to see, so they end up generating radar signatures that are orders of magnitude larger than a 747. Wind farms create complex clutter returns – zero Doppler (non-moving) from the tower and nacelle and non-zero Doppler (moving) from blade movement. The clutter returns

can increase false alarms and report aircraft where there are none. This clutter can also desensitize the radar such that it cannot detect aircraft in the vicinity of wind farms. The turbines also are typically located in the same areas as radars – high up and visible to get as much wind as possible. One mitigation technique used has been to prevent wind farms from being built in 'line of sight' with radars – causing conflict with wind farms development."

"To solve the wind farm challenges, we deploy a pencil beam radar system architecture in the TPS-77 transportable radar as opposed to a stack beam system," Herodes says. "The TPS-77 has the capability to produce a usable air picture within two hours of touching down at any given radar site to a range of 250 nautical miles. Many other long-range radars use multiple stacked beams to detect a target and estimate its height. Since stacked beam transmission is simultaneous, wind farms interfere with all the beams, confusing the radar's signal processing. Stack beam systems illuminate a much larger area – as large as a horizon to 40,000 feet. Comparing that to a pencil beam illumination is like comparing a flood light to a laser beam."

"A pencil beam system like the TPS-77 focuses energy very narrowly while illuminating the targeted vicinity," Atherton says. "It localizes the wind farm and enables proper detection of an aircraft immediately around the wind farm. Because it is using focused energy, it is less likely to pick up the false alarms associated with the wind farm signatures. Transmission and processing are independent, preventing clutter in one beam from affecting another. An important result of the pencil beam architecture is the ability to detect uncooperative aircraft such as those that turn their transponders off so they cannot be seen. The fundamental architecture of pencil beam separates the returns and differentiates aircraft from the clutter returns.

"The TPS-77 3D Scanning Pencil Beam Phased Array system uses advanced clutter reduction techniques such as full monopulse processing, range adaptive MTI/Doppler processing, and advanced clutter mapping," Atherton continues. "Advanced signal processing techniques such as monopulse processing are necessary for detecting objects in littoral clutter and water/land transitions, which can change clutter characteristics fairly dramatically. Monopulse processing breaks a single beam into three beams and compares the results to determine target location. From the same transmission you can get many different looks, in many different pieces, and then integrate those pieces to get a better view of the returns. We only need a single return to get full stated accuracy of a radar system with monopulse processing. Other radars use multiple returns and employ averaging techniques, whereas monopulse is one hit, which is important for high-clutter environments around wind farms. Adaptive MTI and Doppler are excellent at suppressing targets of different speeds and velocities – like wind turbine blades." **MES**

Two-phase cooling meets the challenges of modern radar applications

By Dan Kinney and Andy Johnston

As radar systems' heat flux and thermal loads continue to increase alongside the pace of technology, two-phase liquid cooling is winning the race, beating more traditional methodologies (such as air-based and single-phase liquid cooling) to the finish line.

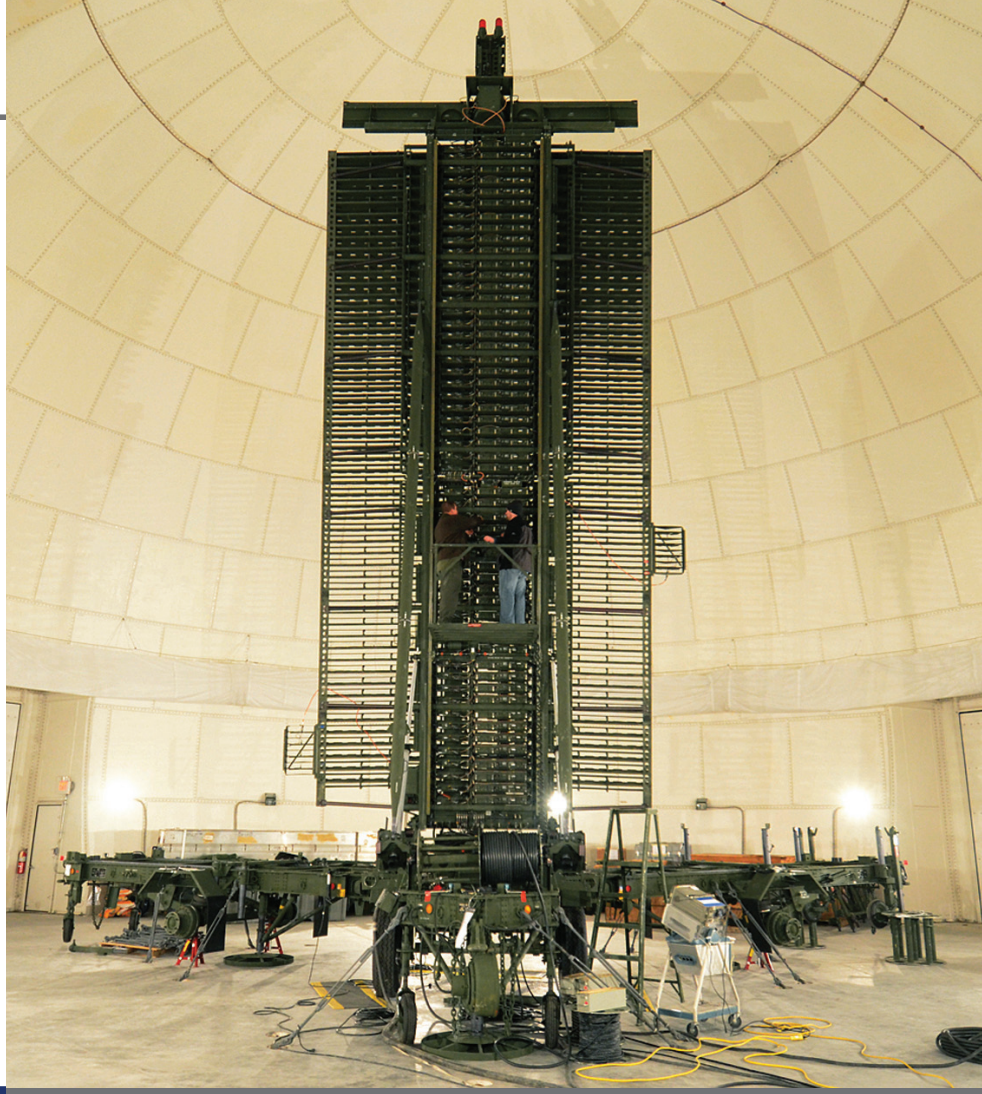

U.S. Army photo by Steve Grzezdzinski, CECOM

Customers today are seeking increased capability and performance from advanced military systems such as radar and directed energy weapons. The ever increasing thermal load and heat flux (heat rate per unit area) of such systems are challenging the practical limits of conventional air-based and single-phase liquid cooling solutions. Thermal performance as well as Size, Weight, and Power (SWaP) constraints are also causing engineers to turn to more advanced cooling methods. Active Electronically Scanned Array (AESA) radar using high-density Transmit/Receive (T/R) modules require special attention. In fact, single-phase liquid cooling solutions are projected to be insufficient for future programs such as the Navy's Next Generation Jammer (NGJ) and Integrated Topside (InTop), both of which require AESA technology. For these reasons, two-phase liquid cooling technologies are gaining increased focus and attention, and will allow engineers to overcome the design challenges of advanced radar systems.

Radar and its thermal management evolve

In the '70s, radar systems had low heat flux and could generally be managed by blowing air over the heated surfaces. By the '90s, heat fluxes had increased enough to where air cooling was oftentimes no longer sufficient, prompting engineers to turn to closed-loop single-phase cooling systems using fluids such as Coolanol and mixtures of Ethylene Glycol and Water (EGW) or Propylene Glycol and Water (PGW). Although more complex and expensive from a system perspective, single-phase cooling systems have orders of magnitude greater capacity to remove waste heat as compared to air cooling. In aerospace applications, the advantage of liquid cooling over air cooling is amplified because of the decreasing density of air with altitude. For instance, a constant speed cooling fan at a 35,000 foot altitude delivers an air mass flow rate that is approximately 40 percent of the flow rate it can provide at sea level, greatly limiting the cooling capacity of the system. Today, single-phase liquid cooling is well understood and is already in use on many military platforms including the F-22 and F-35. These cooling solutions have met the thermal demands of the presently fielded radar and tactical environment and have proven to be reliable and robust.

However, next-generation radar and jamming systems have thermal load and heat flux requirements that are starting to exceed the practical cooling capacity of single-phase liquid cooling solutions, necessitating the deployment of two-phase cooling systems. Two phase-cooling systems outperform single-phase cooling systems for two key reasons. First, two-phase cooling takes advantage of not only the higher heat absorption capacity of the liquid (that is, sensible heat), but also the heat absorbed when the liquid vaporizes (that is, the heat of vaporization). Second, two-phase flow in cooling

channels can support orders of magnitude higher heat flux than single-phase liquid flow at comparable coolant flow rates and thermal conditions.

Two-phase cooling advantages, considerations

For some applications, a two-phase cooling system can operate with less than 50 percent of the flow rate required in single-phase systems. This directly translates into smaller pumps, fluid lines, and quick disconnects, and less pumping power, resulting in a much smaller thermal management system. The resulting higher energy dissipation capacity of a two-phase system is especially important for aircraft and pod applications where SWaP is critical.

Additional benefits of two-phase cooling for radar systems are performance characteristics such as improved range and target detection as well as reliability gains from surface temperature uniformity of cooled components. Surface temperature uniformity can be achieved by ensuring that the phase change from liquid to vapor takes place at near constant pressure (relative to system pressure), as the two-phase mixture flows over the heated surface or through cooling channels. As a result, the bulk temperature of the two-phase mixture stays nearly constant even as the mixture absorbs substantially more amounts of heat from adjacent electronic components. In contrast, in single-phase liquid cooling systems, the coolant temperature rises as it absorbs sensible heat, easily rising between 10 °C and 20 °C in practical applications, which can have significant impact on radar performance.

To understand better the advantages of two-phase cooling systems for radar applications, it is useful to consider Newton's law of cooling. Newton's law of cooling states that heat removed from a surface is proportional to the area of the surface and the temperature difference between the surface and the bulk fluid that cools the surface, or $Q=h*A*(T_{wall}-T_{fluid})$. The proportionality constant, h, is the well-known heat transfer coefficient, which has units of W/m^2C. In radar cooling applications, an order of magnitude higher heat transfer coefficient can be achieved by employing two-phase flow than can be achieved with any practical single-phase cooling solution. For a desired heat removal rate through a fixed surface area, the higher heat transfer coefficient allows for a smaller temperature difference (ΔT) between the fluid and the surface being cooled. A smaller ΔT, in turn, means that the coolant flowing through a cold plate can have a higher temperature while maintaining the same surface temperature. Radar system requirements have a worst case ambient temperature and maximum radar temperature (fixed system temperature difference); delivering hotter fluid to the cold plates allows thermal engineers significant system-level SWaP benefits such as designing smaller heat exchangers or eliminating vapor

compression cycles to condition the cold plate supply fluid. (For a detailed case study, see Sidebar 1.)

Ultimately, the decision on whether to adopt a two-phase system or stay with a single-phase solution will depend on a complete evaluation of system requirements along with its capabilities and operational limitations. On one hand, a well-designed and efficient two-phase system will usually require more time to design and implement, and much of the design time comes from the increased complexity of thermal analyses required because of the nature of two-phase flow. In addition, two-phase systems, specifically systems using refrigerants, will have slightly different operating and servicing requirements compared to single-phase systems. These differences can, however, be managed and are similar to servicing an air conditioner for a home or automobile.

On the other hand, the benefits of a smaller thermal management system have far reaching impact on the overall radar system. Consider, for example, the NGJ Broad Area Announcement (BAA), which states that 60 kW to 90 kW of power must be dissipated by the cooling system[1]. Allowing a

Liquid cooling case study

Liquid cooling case study

The advantage of two-phase liquid cooling over single-phase cooling can be demonstrated in the following cold plate case study. Assume a 0.5 m x 0.5 m radar panel dissipates 30 kW of heat evenly distributed. This panel is to be cooled by a cold plate while maintaining the surface of the cold plate at 70 °C in a typical ground or shipboard 49 °C hot day environment. (Or this ambient temperature can be matched in stagnation temperature by flying around 30 kft at nearly mach 1.) The total system temperature difference (ΔT) is 70 °C - 49 °C = 21 °C. Let's also assume that half of the temperature difference is allotted to heat acquisition and half to the heat rejection. This sets the cold plate inlet temperature to 49 °C + 10.5 °C = 59.5 °C. Since much of the design challenge of liquid cooling is at the radar, the case study will focus on the cold plate only. Table 1 shows the results of a single-phase and two-phase cold plate design exercise.

Design cases 1-3 demonstrate the flow rates need to meet the surface temperature requirement in an easily machined cold plate with 150 channels of a 2.5 mm x 2.5 mm cross-section. Both the PAO and 60 percent Propylene Glycol and Water (PGW) cold plates require very high to unreasonable flow rates to achieve the necessary heat transfer coefficients to adequately cool the cold plate surface. Case 3 is the two-phase flow R134a cold plate, which requires only 7 Gallons Per Minute (GPM) at 1 psi pressure drop, an order of magnitude lower flow rate and pressure drop compared to Cases 1 and 2. Cases 1-3 help illustrate the system-level pumping improvement of two-phase liquid cooling that results from higher heat transfer coefficients at lower coolant flow rates and pressure drops.

To overcome the lower heat transfer coefficients, single-phase cold plates are typically designed to have the maximum heat transfer surface area with the penalty of higher pressure drop and a heavier cold plate. Cases 4 and 5 show the results of a more typical single-phase cold plate with folded straight fin with a 100-fold increase in number of fins and 3-fold increase in fin height over cases 1-3, constraining the design to a reasonable pressure drop and flow rate. Unfortunately, even at the highest fin density cases, neither of the single-phase cases is able to meet the surface temperature requirement and would require fluid conditioning through a vapor compression cycle, relief on the requirements, or dedicating more of the total system temperature drop to the cold plates at the cost of a larger heat rejection heat exchanger.

Case	Fluid	H_fin	W_fin	Tin	Tsurface	dP	vdot	Convection coefficient	Fluid dT
(-)	(-)	in	in	C	C	PSID	GPM	W/m^{2}K	C
Machined: 2.5mm x 2.5mm (0.1in x 0.1in) channels									
1	60% PGW	0.1	0.1	59.5	70	11	66	7300	2
2	PAO	0.1	0.1	59.5	69	49	179	7500	1.5
3	R134a	0.1	0.1	59.5	66	1	7	11300	0.2
Folded fin: 4331 Fins per meter (110 fins per inch)									
4	PAO	0.3	0.005	59.5	77	43	22	2400	12
5	60% PGW	0.3	0.005	59.5	75	20	11	5400	12

Table 1 | Liquid-cooled radar case study results

Sidebar 1 | A cold plate case study illustrates the advantages two-phase liquid cooling offers versus single-phase cooling.

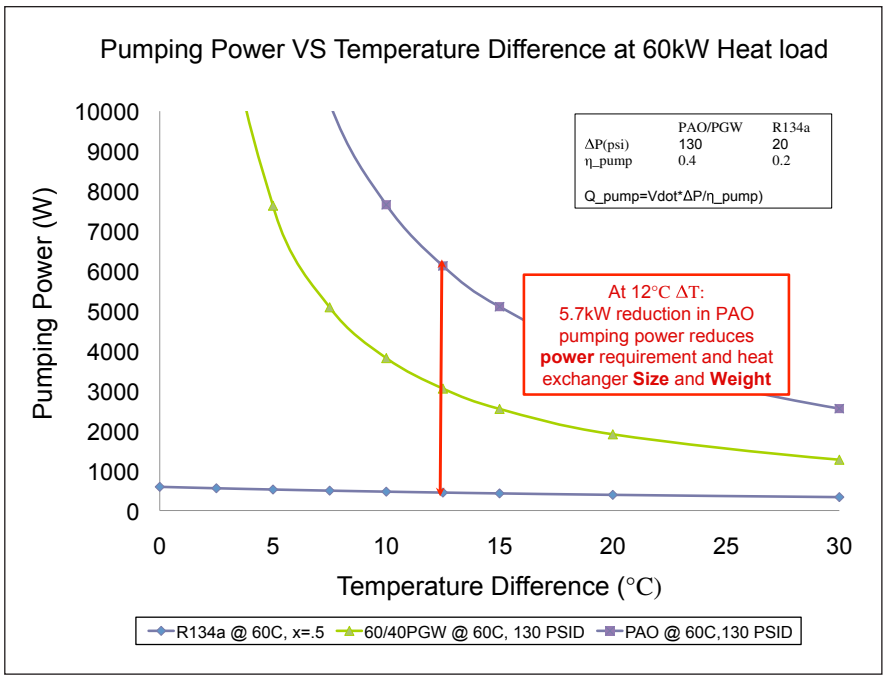

Figure 1 | System pumping power versus coolant temperature difference for single- and two-phase systems

single-phase system to operate with a (typical) 12 °C ΔT at the cold plate, a Polyalphaolefin (PAO) system would require more than 3 kW of pumping power (Figure 1) compared to about 0.5 kW for an R134a two-phase system. Figure 1 also shows that a 60 percent EGW single-phase system requires more than 6 kW of pumping power under the same system requirements. The two-phase system, operating at a lower flow rate and pressure drop, can therefore use a smaller pump and heat rejection heat exchanger because less pump waste heat must be rejected.

Matching the thermal management system

As previously discussed, the demand for two-phase cooling systems, especially in advanced radar systems, continues to increase. Matching the right thermal management system to a particular application will yield the best solution for the warfighter. Understanding differences as well as advantages and disadvantages of the various types of thermal management solutions is essential. Fortunately, Parker Aerospace has been designing, building, and fielding robust single- and two-phase liquid-cooled systems for more than 20 years. Some of these systems have been developed for fixed wing aircraft including manned and unmanned vehicles, rotorcraft, ground mobile vehicles, shipboard applications, and even missiles. Many of the next-generation systems, where heat loads and heat flux are increasing so dramatically, stand to benefit the most from efficiencies gained by two-phase liquid cooling. **MES**

Reference:

[1] Navy SBIR FY2008.1, Topic N08-035, Pod Mechanical Power Production, SBIR/STTR interactive topic information system, www.navysbir.com/08_1/100.htm

Dan Kinney *is the Business Development Manager for Parker Hannifin Corporation's Aerospace Group Thermal Management Systems (TMS) team. Contact him at dan.kinney@parker.com.*

Andy Johnston *is an engineer for Parker Hannifin Corporation's Aerospace Group Advanced Development Team. Contact him at andrew.johnston@parker.com.*

Parker Hannifin Corp.
www.parker.com

SIGNAL PROCESSING
TRENDS IN RADAR, SONAR,
AND ELECTRONIC WARFARE

Fast processors, FPGAs fuel radar/EW signal processing performance

By John McHale, Editorial Director

Radar and Electronic Warfare (EW) designers' thirst for more and more data is driving innovation at the signal processing level as embedded computing suppliers work magic with FPGAs and processors to create intelligent, fast sensor networks.

Radar onboard U.S. Navy Aegis-class ships uses PowerStream technology from Mercury Systems for radar signal processing.

The need to know where and when the enemy will strike is no less important today than when the allies cracked the allegedly unbreakable German Enigma code during World War II. The only difference today is the sophistication of the technology on either side. Modern Electronic Warfare (EW) systems can detect, collect, and catalog just about every signal, while radar systems being developed today will be able to track dismounted personnel, small objects at sea, as well as enemy fighters. To make these capabilities possible, embedded computing wizards are packing as much processing power as possible at the payload level on the platform, right next to the sensor. They mix together components such as GPGPUs, multicore processors, FPGAs, and complex software algorithms in different configurations designed to meet the low Size, Weight, and Power (SWaP) requirements demanded by their military customers.

Sensor technology and the need for greater bandwidth are driving military electronics designs, especially in radar and electronic warfare applications, says Ray Alderman, Executive Director of VITA. "What the processor and sensor technology enables us to do is absolutely incredible. Every aircraft has a signature. When we hit an enemy aircraft with radar we can figure out where it came from, destroy that location, then figure out where it is going and eliminate that destination as well. The only problem with increasing bandwidth on the sensors is the ability to stream data over the RF links is still poor. Therefore, the focus going forward is to move the processor next to the sensor."

"Department of Defense (DoD)-oriented customers want to be able to perform data exploitation onboard platforms to provide immediate actionable intelligence to the warfighter – avoiding delays and bottlenecks encountered when sending large amounts of data to a ground station," says Paul Monticciolo, Chief Technology Officer at Mercury Systems in Chelmsford, MA.

"In the past Intelligence, Surveillance, and Reconnaissance (ISR) systems were aimed at finding that particular truck or plane in an area of interest; now they are required to have the capability to pick individuals out of a crowd," says Vincent Chuffart, Embedded Computing Specialist at Kontron in Poway, CA. "Today's sensor systems can detect everything, so integrators want to process that sensor data and extract all they can out of it at the sensor level before they transmit it to the warfighter." Kontron puts a lot into small form factor designs such as 3U VPX to handle the high-speed processing and meet SWaP requirements, he continues. There is a trend toward smaller and cooler systems, especially with the vast majority of upgrades in front of us being just a refresh of existing equipment, Chuffart adds.

"Users are trying to detect and decode complex radar or communication signals so they need more channels and these channels need to be accommodated simultaneously," says Rodger Hosking, Vice President at Pentek in Upper Saddle River, NJ. "An example would be a communication system that listens to multiple radios or communicates with multiple radios at the same time or a system that requires wider signal bandwidths for radar. Once these signals are digitized, higher data rates are required

to move the data, which puts a bigger load on the DSP engine to accommodate the data so it can keep up in real time." A higher level of system performance is required through more sophisticated signal processing techniques, higher-speed A/D converters to capture wideband signals, and faster FPGAs to do signal processing and process algorithms at higher rates, he adds.

Harnessing commercial processor technology

Embedded computing companies are leveraging commercial processor technology – whether from Intel, NVIDIA, or other companies – to drive the performance of radar and electronic warfare systems. "FPGAs and GPGPUs are excellent for front-end sensor processing, but the SWaP characteristics of multicore processors are enabling us to provide the type of performance and analysis of data that is typically done in a ground station," Monticciolo says. "Some companies are working with mobile-class processors, but when you start doing massive correlations and graph-type processing, you need the performance of a multicore processor in a server-class solution. This also guarantees code portability – being able to run the same code on a 1U server in the ground station and in a small embedded processing system up on the platform." Mercury has a high-performance embedded computing solution with their PowerStream technology performing radar processing onboard the Navy's Aegis-class ships.

"On Intel Architecture, the advent of AVX and the coming of AVX2 with a fused multiply-add pipeline are significantly improving the applicability of Intel devices to signal processing applications, as is the increasing performance of the integrated GPUs," says Peter Thompson, Senior Business Development Manager for high-performance embedded computing at GE Intelligent Platforms in Huntsville, AL. "Now that NVIDIA is shipping Kepler GPUs, we are seeing significantly better performance per watt than the previous generation Fermi offered. GPUDirect is helping us to reduce sensor-to-processor latency, and opening up some new applications such as EW that are time sensitive. The interconnects are keeping pace too – PCI Express Gen 3, 40 GbE, and 56 Gbps InfiniBand are starting to become available or are already out there, and are allowing us to keep the processing pipelines fed."

"For many EW and DSP applications, our customers have found that the Intel Core i7 processor provides them with the right balance of performance and SWaP," says Ben Klam, Vice President of Engineering at Extreme Engineering Solutions (X-ES) in Madison, WI. "The Intel Advanced Vector Extensions (AVX) supported by the Intel Core i7 processor provides excellent DSP performance with support for operations on 256-bit vectors."

Managing power consumption and signal integrity

Processors do enable amazing applications, but also create serious headaches for embedded system designers with the heat they generate. Cooling the systems and keeping their power consumption low can be quite complex, especially as system designs trend toward smaller form factors.

"Two big challenges faced by embedded signal processing designers are power consumption and signal integrity," says Denis Smetana, Product Marketing Manager for FPGA products at Curtiss-Wright Controls Defense Solutions in Ashburn, VA. "Devices continue to run faster, which requires more power. With so many high-speed signals, more exotic PWB material needs to be used along with special handling of high-speed traces and more detailed signal integrity and power integrity analysis. The super high-speed signaling is running the processors so fast that a lot of heat is being dissipated and as geometries keep getting smaller, leakage current gets bigger as a percentage of total power. And leakage current is very sensitive to temperature. This results in a significant power increase as temperature rises. Software/firmware also is needed to be able to utilize the larger processing performance."

"There are some applications with DSP performance requirements that exceed the capability of the Intel Core i7 processor where customers add GPGPUs or FPGAs into the mix, but the cost is high," Klam explains. "GPGPUs typically consume more power than embedded General Purpose Processors (GPPs), which in turn creates more heat that has to be dissipated. GPGPUs are driven by the consumer market, so product obsolescence can be a big problem for long-life embedded systems. And development is much more complex – software development in the case of GPGPUs and VHDL and software development in the case of FPGAs. We are also seeing a lot of demand for systems that require something smaller than 3U VPX can support," Klam says. "We have developed a small form factor system, the XPand6000 Series, that utilizes COTS components – COM Express, PMC/XMC, and SSDs – to enable customers to rapidly prototype and deploy small form factor solutions."

FPGAs and the front end

For the front end of signal processing solutions – where the signals are received by the embedded computing system – designers far and wide sing the praises of today's FPGAs, whether Xilinx or Altera, for enabling the capabilities of modern radar and electronic warfare platforms.

"We are seeing people wanting to put more components of their radar systems in an FPGA," says Jeff Milrod, CEO of

Bittware in Concord, NH. "Altera's Stratix technology is supporting this by placing a lot of floating-point capability in the front end. That way the system can aggressively integrate full imaging and even identify areas of interest before it sends data to the ground. FPGAs are needed because the rates are so high, now everybody seems to want direct RF conversion – so they get gigahertz sample rates flying and then handle them in the FPGAs."

Compared with a GPP, the FPGAs are much better suited for real-time embedded systems. The GPP, while very fast, is not well connected to I/O and does not do real-time data processing as well as an FPGA. Pentek's Onyx 71720 software radio module is used in radar, UAV, and communication signal processing applications and is based on the Xilinx Virtex-7 FPGA, used in radar and communication signal processing applications.

"FPGAs will continue to dominate digital signal processing for a myriad of reasons: mainly processor technology isolation and control of one's own IP – the critical element in any real-time, signal processing platform," says Doug Patterson, Vice President of the Military & Aerospace Sector at Aitech in Chatsworth, CA. Aitech's military-grade 3U CompactPCI, 3U VPX, 6U VPX, and 6U VMEbus systems leverage FPGAs for radar and fire-control applications.

FPGAs process everything very fast and enable radar/EW integrators to have more control, to adapt and change their applications based on mission results, says Jane Donaldson, President of Annapolis Microsystems, in Annapolis, MD. "A major drawback with FPGAs for many is the expense of programming in VHDL, which adds labor costs and lengthens the design cycle," she continues. To reduce time to market and development costs on their WILDFIRE FPGA boards, Annapolis engineers use their CoreFire solution to program FPGAs. It enables software programmers working on GPPs to also program FPGAs, Donaldson adds.

"The combination of FPGAs with SBCs, GPGPUs, and multicore processors in one system is how requirements are trending," Smetana says. "For example, on the front end, FPGAs are well suited for parallel processing of sensor data, but then may feed the data to GPGPUs for additional parallel processing or multicore processors for sequential processing." Curtiss-Wright is working with Tektronix to improve performance for wideband, low-latency processing for Digital Radio Frequency Memory (DRFM), electronic warfare, signal intelligence, and electronic countermeasure applications. Under the collaboration, Curtiss-Wright's CHAMP-WB ("Wideband") board will work with the Tektronix ADC/DAC FMC module, the TADF-4300, to become the CHAMP-WB-DRFM utilizing a Xilinx Virtex-7 FPGA.

FPGAs embracing ARM and OpenCL

"In the future we see embedded systems generally moving toward embracing ARM processor technology," Milrod says. "There is a huge user community and infrastructure, it is quite efficient, and the performance continues to improve dramatically with 64-bit now emerging. Although currently there is no ARM COTS community in the military, interest in these designs is exploding in the Linux world and could very well catch on in military designs. Both Altera and Xilinx are integrating ARM into all their FPGAs from now on. At BittWare, we've started to integrate them onto our high-end VPX boards. The ARM can handle housekeeping functions while the massive bulk processing onboard is handled by the FPGAs and/or floating point coprocessors such as our Anemone many-core processor that is based on Adapteva's Epiphany architecture.

"One standard that is getting interest in the defense community is the OpenCL

Radar signal capture and collection speed enhanced by new software tool

Engineers at Rohde & Schwarz in Beaverton, OR, designed a new software tool – Technical Pulse Analysis (TPA) – to help Electronic Intelligence (ELINT) operators and analysts cut down on uncollected radar signals and speed up analysis.

"Modern Radar Electronic Support Measures (RESM) or Radar Warning Receivers (RWRs) must rely on the collected and analyzed data," says Darren McCarthy, Marketing Manager for Rohde & Schwarz (R&S). "As new radars increase in frequency and agility as well as in implementing new types of operating modes, the efficiency of these support measures must be able to adapt to evolving modes and capabilities. While some of these systems can provide some level of signal collection, it is only by capturing the IQDW signal and subsequent analysis that can reliably identify the capabilities of the new radar signal. Radar signals moved down to classical communication bands and vice versa." The TPA tool is an offline software analysis tool for extracting and displaying radar signal Pulse Description Word (PDW) and intrapulse description word (IQDW) parameters. The ease of use and analysis capability enable the ELINT signal analyst to rapidly perform tasks such as deinterleaving multiple different signals from a single collection, determining scan patterns, determining intrapulse modulation techniques, and performing radar mode profiling, he adds.

"Typical ELINT collection systems have tightly integrated hardware and software capabilities and are sold as a system," McCarthy continues. "As radar signals increase beyond the hardware limitations of the collection system, some signals go uncollected. Hardware collection systems that focus on signal collection (digital I/Q), especially wideband collection systems, do not typically focus on efficient signal extraction and pulse analysis for the ELINT collector." Via full IQ processing capabilities, TPA provides a view on radar signals including FMCW types and is not limited to pulsed signals only, he adds.

Radar/EW Signal Processing companies

AccessIO
www.accesio.com

ADLINK Technology
www.adlinktech.com

Aitech
www.rugged.com

Altera
www.altera.com

AMD
www.amd.com

Annapolis Micro Systems
www.annapmicro.com

BittWare
www.bittware.com

Cambridge Pixel
www.cambridgepixel.com

Coherent Logix
www.coherentlogix.com

Concurrent Technologies, Plc.
www.gocct.com

Creative Electronic Systems (CES)
www.ces.ch

Critical I/O
www.criticalio.com

CSPI
www.cspi.com

Curtiss-Wright Controls Defense Solutions
www.cwcdefense.com

Data Device Corp.
www.ddc-web.com

Elma Electronic
www.elma.com

Emerson Network Power
www.emersonnetworkpower.com

Enginasion
www.enginasion.com

Extreme Engineering Solutions (X-ES)
www.xes-inc.com

FLIR Systems
www.flir.com

Freescale
www.freescale.com

Galleon Embedded Computing
www.galleonembedded.com

General Atomics Aeronautical Systems
www.ga-asi.com

General Dynamics C4 Systems
www.gdc4s.com

General Electric Intelligent Platforms
www.ge-ip.com

General Micro Systems
www.gms4sbc.com

Hartmann Elektronik
www.hartmannelectronic.com

Innovative Integration
www.innovative-dsp.com

Intel
www.intel.com

Kontron
www.kontron.com

Mercury Systems
www.mrcy.com

Microsemi
www.microsemi.com

Nallatech
www.nallatech.com

North Atlantic Industries
www.naii.com

NVIDIA
www.nvidia.com

Radisys
www.radisys.com

Rakon
www.rakon.com

Parker Aerospace
www.parker.com

Parvus Corp.
www.parvus.com

Pentek
www.pentek.com

Raytheon-Applied Signal Technology
www.appsig.com

Rohde & Schwarz
www.rohde-schwarz.com/us

Spectracom
www.spectracom.com

Spectrum Signal Processing
www.spectrumsignal.com

SRC, Inc.
www.srcinc.com

TechSource
www.techsource.com

TEK Microsystems
www.tekmicro.com

Telephonics
www.telephonics.com

Themis Computer Inc.
www.themis.com

Trenton Systems
www.trentonsystems.com

TTTech
www.tttech.com

Versalogic
www.versalogic.com

Xilinx
www.xilinx.com

Z Microsystems
www.zmicro.com

framework," Milrod says. "Altera has made a big push with it and they've really driven OpenCL on their FPGA designs. Within the defense community we've seen attention for OpenCL because some designers are dissatisfied with inefficiencies of running on GPUs, and then the difficulties they have when porting GPGPUs to FPGAs."

Software challenges

Managing processors and the heat and power they generate is only part of the difficulty in creating complex signal processing solutions. "Integrating new, untested technologies while developing new, critical application software drains budgets and saps our customers' resources and energy," Patterson says. "Serial buses and parallel buses operate very differently, especially when time-critical messages must occur in order and in phase. Multiple serial buses can easily become out of sync and message passing becomes tricky with RTOS overhead to sort out the time sync issues. Command/response buses with tight, time-aware and synchronized higher-level data protocols are critical in real-time process control systems."

"Radar and EW systems are so complicated with so many levels that one of our biggest challenges lies in offering software board support packages that work with various types of middleware," Pentek's Hosking says. "Often the code and middleware don't interact well and don't do what people expect them to do. This is the biggest challenge we face – providing tools that enable our customers to efficiently develop their unique and highly complex applications under popular operating systems like Windows, Linux, and VxWorks."

Software solutions also help radar display processing engineers to combine "graphical data with real-time radar data and to create pictures using standard commercial technology," says David Johnson, Managing Director for Cambridge Pixel in Royston, England. "The amount of data being processed and displayed is more complex than ever, so to create complex pictures of graphical data, we use high-end graphics chips from companies such as AMD and NVIDIA. We present a representation of the data generated by the front-end processors that makes sense – taking multiple layers of radar data, map data, chart data, and other sensor data fused together in one composite display. The system integrator decides how he wants the ISR information displayed." Cambridge Pixel's standard product – the SPx Radar Development Library – is a collection of software modules that can be used for radar displays as well as recording or tracking systems, he adds. **MES**

TODAY

Today we become Mercury Systems.

Today our capabilities are far broader and deeper than ever.

Today Mercury Systems is a best-of-breed provider of commercially developed, open sensor and Big Data processing systems, software and services for critical commercial, defense and intelligence applications.

Today our innovative capabilities span three strategic areas:

Mercury Commercial Electronics – We deliver innovative, commercially developed, open sensor and Big Data processing systems for critical commercial, defense and intelligence applications. We deliver solutions that are secure and based upon open architectures and widely adopted industry standards. We deliver rapid time-to-value and world-class service and support to our commercial and prime contractor customers.

Mercury Defense Systems – We deliver innovative, open sensor processing solutions to key prime contractors leveraging commercially available technologies and solutions from our Commercial Electronics business. We leverage this technology to develop integrated sensor processing subsystems, often including classified application-specific software and IP for the C4ISR and EW markets.

Mercury Intelligence Systems – We deliver Big Data processing systems, software and services to solve complex Intelligence Community challenges with innovative technology positioned to support key intelligence missions. Our unique approach and solutions facilitate the transformation of raw data into information and then to actionable intelligence.

Check out the new Mercury . . . Today


Mercury Commercial Electronics
Microwave solutions
Mixed analog/digital solutions
Embedded signal and image processing
Mercury Defense Systems
Electronic warfare
Signals intelligence
Electro-optical/infrared
Test and simulation
Mercury Intelligence Systems
Big Data analytic processing
Predictive analytics
Multi-intelligence analysis
MERCURY
SYSTEMS™
Innovation That Matters™
Today Mercury Systems is a new name, but not a new company.
And 'Innovation That Matters' is not just a slogan. It describes
how we approach the problems you face.
Visit mrcy.com/newmercury or call 866.627.6951

10 GbE enables real-time remote desktops for C4ISR

By Jeff Malacarne

Virtualization trends in commercial computing offer benefits for cost, reliability, and security, but pose a challenge for military operators who need to visualize lossless imagery in real time. 10 GbE technology enables a standard zero client solution for viewing pixel-perfect C4ISR sensor and graphics information with near zero interactive latency.

U.S. Marines photo by Pfc. Sarah Anderson

For C4ISR systems, ready access to and sharing of visual information at any operator position can increase situational awareness and mission effectiveness. Operators utilize multiple information sources including computers and camera feeds, as well as high-fidelity radar and sonar imagery. Deterministic real-time interaction with remote computers and sensors is required to shorten decision loops and enable rapid actions.

A zero client represents the smallest hardware footprint available for manned positions in a distributed computing environment. Zero clients provide user access to remote computers through a networked remote desktop connection or virtual desktop infrastructure. Utilizing a 10 GbE media network for interconnecting multiple computers, sensors, and clients provides the real-time performance and image quality required for critical visualization operations. The cost of deploying a 10 GbE infrastructure is falling rapidly and 10G/40G has become the baseline for data center server interconnect. Additionally, deploying common multifunction crew-station equipment at all operator positions brings system-level cost and logistics benefits. The following discussion examines the evolution to thinner clients and the path to a real-time service-oriented architecture, in addition to looking at zero client benefits and applications.

Evolution to thinner clients

For military C4ISR, capabilities provided by legacy stovepipe implementations are being consolidated into networked multifunction systems of systems. To accomplish this, open standards and rapidly advancing technologies for service-oriented architectures are being leveraged (Figure 1). For crew-station equipment, this drives an evolution from dedicated high-power workstations toward thinner client equipment at user locations. Computing equipment is being consolidated away from the operators into one or more data centers. This leaves the crew station with a remote connection to system resources, but does not ease the requirement for high-performance access to visual information. 10 GbE provides the client/server connection performance necessary for real-time remote communication.

Workstations at operator positions normally run software applications locally and provide dedicated resources for data and graphics processing. Server-based data processing and networked sensor distribution systems have moved much of the application processing away from the operator. This can simplify the job of system administration and maintenance and enables multiple users to access the same capabilities. However, much of the processing for presenting images to operators can be unique to the individual needs for varying roles at each position.

Thin clients can be utilized to provide dedicated graphics and video processing horsepower for user-specific visualization operations such as windowing, rendering, and mixing multiple data and sensor sources. Dedicated local graphics processing power can be important for critical real-time operations or for interfacing to servers without high-performance graphics capabilities. This makes a thin "networked visualization client" a flexible option for multifunction crew stations that must interface with both legacy and newer service-oriented systems.

For commercial computing systems, a major push is underway to move high-performance graphics capability into the data center servers. This can be implemented via dedicated workstations for each crew station, virtualized compute engines with dedicated graphics for each crew station, or completely virtualized environments with networked image distribution. Virtualization provides a means to share CPU and GPU compute cycles between multiple users, gaining efficiency from higher utilization of system hardware resources. However, for mission-critical C4ISR systems, a deterministic Quality-of-Service level for performance, reliability, and security must be maintained.

For systems with both computing and graphics processing located away from the operator, zero clients provide network-attached displays with audio and user input devices (keyboard, mouse, and touch screen). Minimizing size, weight, and power at the operator position brings many benefits, but performance depends on the remote visualization processing capabilities and the communication channel. To match workstation performance, a consistent human-computer interaction latency of less than 50 ms must be provided.

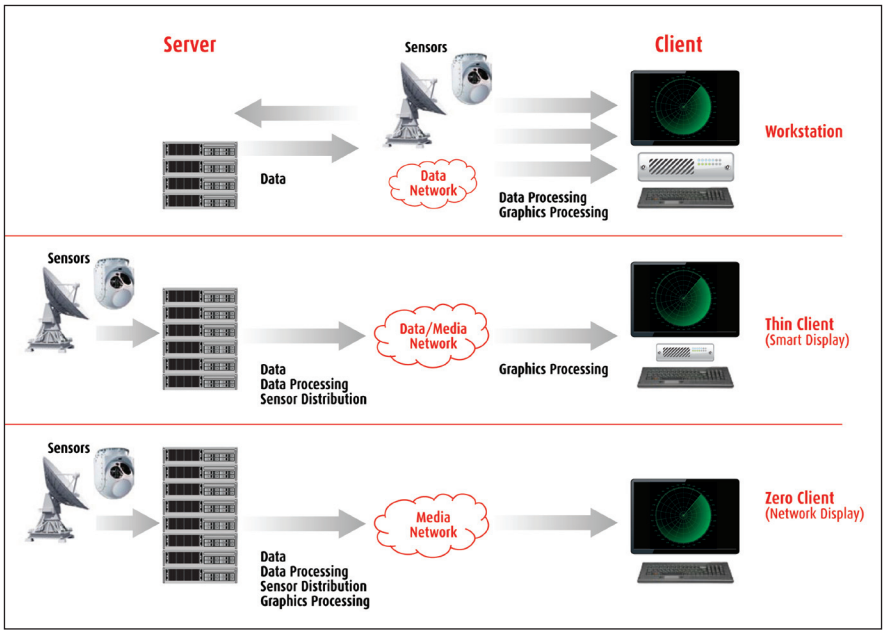

Figure 1 | Client/server evolution: Increasing communications bandwidth enables more service-oriented computing and "thinner" clients.

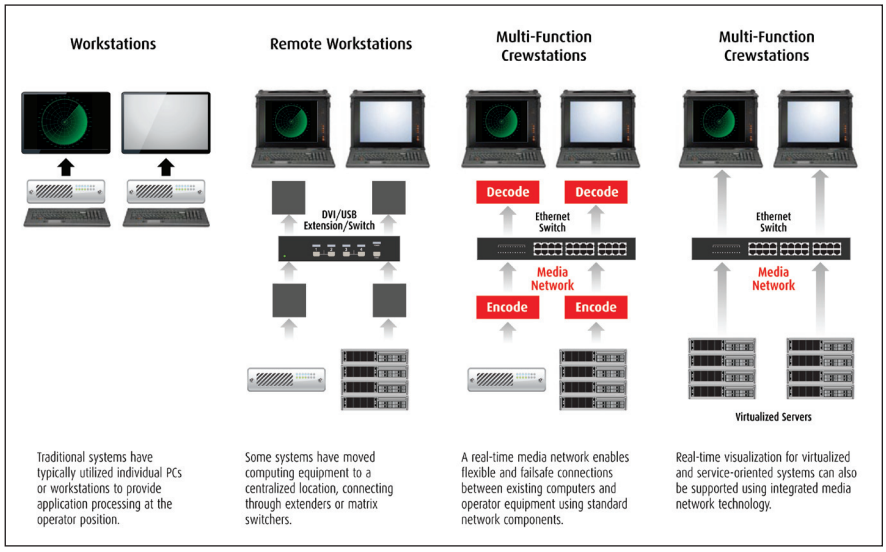

Figure 2 | Crew-station evolution to a service-oriented architecture

Path to a real-time service-oriented architecture

System architects need a graceful technology insertion path that leverages the benefits of thinner clients (Figure 2). One approach for centralizing computing equipment while maintaining performance is to simply move the workstations to the data center and extend the interfaces to the display and input devices. This maintains the dedicated computing resources for critical operations. Video and device interface extension can be accomplished via extenders or switch matrices to provide connections between operators and computers.

A more flexible approach is to utilize a standard network to support highly configurable access to all workstation resources from any operator position. With this approach, any user can connect to any image source and user screens can be shared with collaborative remote displays or other users. This also enables growth to a service-oriented "cloud" architecture that follows the trend for general-purpose IT and data processing systems. However, commercial IT products do not always meet the performance, reliability, security, or logistics requirements for mission-critical C4ISR systems.

To leverage this computing trend for real-time applications, a standard 10 GbE media network can be utilized to connect multiple zero clients to multiple remote graphics and sensor sources. Lossless distribution is supported for high-quality text, dynamic 2D/3D graphics, HD video, radar, and sonar imagery. Compositing multiple sources onto a single screen can be performed at the zero client or by networked video processing services. Near-zero latency interaction and video distribution are

now possible and support deterministic performance and real-time dynamic visualization at any operator position.

One full-resolution (1,920 x 1,200) lossless channel at 60 Hz with 24-bit color requires 3.3 Gbps of bandwidth. Therefore, one 10 GbE connection can support a dual-head crew station at full frame rate with audio and USB support. However, many visual applications require no more than a 30 Hz update rate (including 1,080p/30 HD full motion video), which reduces the bandwidth to 1.7 Gbps per channel. This enables triple-head crew stations with audio and USB support over a single 10 GbE connection. Dual Ethernet ports at the zero client can also be provided to support more video channels, higher frame rates, and/or redundant connections.

Zero client benefits

Compared to workstations, zero clients provide several benefits, including lower TCO, reduced SWaP, higher system availability, and more system security and agility.

Reduced total cost of ownership

Zero clients provide the smallest, simplest, and most maintainable equipment available for the operator position. This means lower initial investment costs as well as lower operating and maintenance costs throughout the system life cycle. System modularity and standard interfaces support seamless technology refresh as new computing and display equipment becomes available. 10 GbE has been widely adopted for data centers and standard component costs are declining rapidly. When compared to legacy stovepipe systems, networked systems also greatly reduce the amount of dedicated cabling required.

Reduced size, weight, and power

Only video, audio, and USB encoding/decoding functions are required with a zero client. These are packaged as small dongles or integrated into the display. Small packaging enables new options for lightweight operator consoles with increased ergonomics, as well as reducing noise and the burden on cooling systems for manned areas.

High system availability

System uptime and reliability benefit from consolidating all computing elements into managed data centers. Common equipment at multiple operator positions and redundant network connections support rapid recovery from computer, client, or network equipment failures.

High system security

Security risks are reduced through centralized administration and access authentication at the data center. Additionally, stateless zero client equipment outside the data center and encrypted communications between all components assure system confidentiality and integrity.

System agility

Systems using common crew-station equipment can be reconfigured by software for different mission roles and objectives. Additional clients can be added quickly to extend the system. Also, as computing systems evolve with new virtual desktop infrastructures, today's investment in zero client equipment is preserved through standard interfaces for video, audio, and user input devices including DVI, PC audio, and USB.

Applications of a zero client

In addition to the benefits of a zero client, the technology's agility also enables a range of applications using common equipment. For example, remote crew stations can now be smaller, lighter, and more versatile, and operator equipment can be located at remote locations not previously possible. Noisy, heat-generating computing equipment can be moved away from operator positions.

Another application highly suited to zero client utilization is the multifunction crew station. Common crew-station equipment can be used to access multiple computers and sensor sources under secure software control. This supports the capability for dynamic access to multiple systems from a single location. Systems can be rapidly reconfigured for different mission objectives, operating roles, or failure recovery.

Collaborative and remote displays also benefit from zero client usage. Unmanned displays can be attached to the network for sharing real-time visual information for dissemination and collaboration. Large area displays for several viewers can receive multiple feeds with full performance. Additionally, selected sources can be compressed and transmitted through secure routers for wider area distribution.

Using zero client technology for networked multifunction crew stations enables the integration of legacy capabilities into a consolidated operating environment as well as the development of new concepts of operation. One example of this is Barco's zero client technology, which brings the benefits of state-of-the-art computing architectures into mission-critical C4ISR systems involving advanced visualization.

Mission-critical solution

Leveraging commercial computing trends and standards provides significant cost and capability benefits. However, the level of real-time performance, mission assurance, and information assurance required for mission-critical C4ISR systems must be achieved. Zero client technology enabled by 10 GbE provides the necessary pixel-perfect viewing of graphics and sensor information for these demanding applications. **MES**

Jeff Malacarne *is a Technical Director for Barco Defense, focusing on networked visualization product strategy. He has more than 30 years of experience developing graphics and video products for air traffic control, naval, unmanned, and ground vehicle systems. Jeff holds a BSEE from SUNY at Buffalo and an MS in Computer Engineering from the University of Southern California. He can be contacted at Jeff.Malacarne@barco.com.*

Barco
www.barco.com

Our technology investments protect yours.

Whatever direction you choose, Aitech has the map!

Established platform parallel bus protocols like VMEbus and CompactPCI still have their place in today's and tomorrow's harsh environment, real-time/hard-deadline embedded sub-system applications...especially when these products are upgraded and maintained to keep pace with the newest, fastest processor and memory technologies.

While there are some applications where high speed serial fabrics like VPX are ideal, there are others where VMEbus or CompactPCI still rule the roost.

One company continues to actively invest in maintaining – and **not** obsolescing – their military and space embedded computing products with a proactive 12-year minimum COTS Lifecycle+™ Program.

And one company continues to also invest in delivering the very best of the newest embedded COTS computing platforms with the new, serial fabric protocols.

And one company actively invests in technology insertion at the board level, creating backplane, pin-compatible products with the latest, next generation memory and processor technologies "on-board".

And that same company still delivers their legacy bus products at full speed **and** full capability and full mil temp range (-55 to +85°C) with those latest technologies.

The one company to do all that? Aitech. Check our website to learn more about our technology roadmaps and how they protect your investments.

Aitech

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

Celebrating 28+ Years of Success

Exploiting MOSA software for high-performance ISR and signal processing applications

By Michael Stern

U.S. Marines photo by Cpl. Jeremy Ross

High Performance Embedded Computing (HPEC) is rapidly becoming mainstream in the mil/aero embedded computing world, bringing with it the potential to solve increasingly challenging problems such as that of processing ISR and other signal processing data. The key for developers, however, is to make the right architectural and technology choices as mandated by the DoD's Modular Open Systems Architecture (MOSA) initiative – and then to exploit those architectures and technologies, rather than be constrained by them. To help accomplish this, software tools utilizing advanced modeling and optimization techniques are available to enable designers to analyze their MOSA-based implementations for maximum efficacy.

Today, we live in a connected world in which the on-demand information we take for granted is made available through densely populated data processing centers. These centers process what is sometimes referred to as *Big Data,* and they rely on High Performance Computing (HPC) technology built out on a large scale measured in acres of rack-mounted servers and megawatts of power consumption.

The mil/aero world is seeing a similar revolution in terms of connectivity – except that what is being connected isn't cellphones, tablets, and netbooks. It's sensors – video, infrared, radar, sonar, and so on. And those sensors are becoming capable of capturing data at higher and higher resolutions. It is also becoming common to see dedicated processors give way to systems that process and fuse data from multiple modalities. For instance, a UAV may carry EO, IR, and SAR sensor suites, but might not have the Size, Weight, and Power (SWaP) budget for a dedicated processor for each as can be the case on larger airframes. It is becoming common practice to feed all these data streams across a standard interface (such as Ethernet) into one processing system that employs multiple multicore devices to provide the compute power required to form the individual products and to then merge and exploit them.

That need to process more data from more sources at higher resolutions at faster speeds gave rise to the High Performance Embedded Computing (HPEC) phenomenon. It is the direct descendant of HPC, but instead of occupying acres of air-conditioned data centers consuming megawatts of power, it's being deployed in 6U – and even 3U – chassis in environments that are invariably SWaP constrained.

The key for developers of radar, signal processing, ISR systems, and other mil/aero technologies is to identify the optimum architecture and technology choices. The right choices – industry standards such as OpenVPX and commercial technologies such as InfiniBand and 10 GbE switch fabrics – lead to reduced development time, reduced risk, reduced cost, and reduced time to revenue. The wrong choices lead to avoidable risk, cost, and effort – and, no less significantly, dead ends that can substantially impair the long-term viability and supportability of a program. Being based on industry standards, these technologies are also compliant with the demands of the Modular Open Systems Architecture (MOSA). And, enabled by modern software tools featuring advanced modeling

and optimization capabilities to ensure maximum MOSA efficacy, the sky's the limit when it comes to performance.

Technologies: Look to the commercial world

MOSA is the DoD's approach to these choices. It is not prescriptive in terms of architectures and technologies; rather, it requires developers to leverage widely supported, consensus-based standards, whatever those might be and wherever they might be found. The MOSA approach looks to enhanced combat capability for the warfighter through enhanced interoperability, reduced lifecycle costs, and a shortened cycle time.

The Commercial Off-the-Shelf (COTS) philosophy has now long been ingrained in the minds of mil/aero systems developers, such as those designing signal processing and ISR systems, and can be considered a precursor to MOSA with its emphasis on access to the latest commercial technologies that are supported by not just a broad array of competitively priced software and hardware products, but also by an ecosystem of skills and expertise. HPEC applications such as signal and image processing can derive similar benefits. The HPC world provides a model, and developers wanting to create sophisticated HPEC applications should turn toward the commercial world.

Today's Big Data centers and communications networks run server-grade Linux operating systems and open standard middleware upon which enterprise-level applications are built. These software layers are optimized to provide very high system availability as well as the highest levels of interprocess communication performance across multiple compute nodes, compute clusters, and data center networks.

Typical hardware platforms comprise thousands of x86 multicore CPUs and can also harness thousands of many-core, parallel processor nodes such as Graphics Processor Units (GPUs) to accelerate compute-intensive tasks such as radar and other sensor data processing. These CPU/GPU clusters are interconnected over high-speed serial networks based on protocols such as PCI Express, Ethernet, and InfiniBand from companies such as Intel, NVIDIA, AMD, Mellanox, and other silicon vendors.

Mil/aero embedded computing developers can also benefit by sourcing these same open architecture platforms in form factors such as OpenVPX that are designed for rugged, extended temperature, shock, and vibration environments typical of deployed defense and aerospace platforms. Such open architecture platforms enable the same applications developed on commercial server clusters to be deployed on rugged tactical systems in the theater of operations. The benefits? Faster development, lower cost, reduced risk, earlier deployment. In addition, by being based on open architectures at the silicon, board, system and middleware layers, such platforms are inherently aligned with the guiding principles of MOSA.

ISR: Sensor-to-user platforms are on the rise

Deployed Intelligence, Surveillance, and Reconnaissance (ISR) platforms comprise sensors that acquire and process real-world signals at very high resolutions. Such sensors include antennae for radar and signals intelligence as well as a variety of cameras and other interfaces. The analog signals captured at the front end of the platform are digitized and then processed in various stages to extract meaningful information for users within various timeframes depending on the needs of the user communities. For example, the platform itself might require output from the data processing pipeline to help control the vehicle and execute the mission, especially when the platform is unmanned. Event-driven control loops in HPC systems can usually meet the needs of their user communities within varying response times measured in hundreds of milliseconds or even seconds rather than microseconds.

ARGUS-IS, for instance, has 368 five-megapixel cameras that combine for an equivalent of 1.8 gigapixels total, and can supply 65 separate VGA resolution video streams to warfighters. A 15 Hz frame rate yields an aggregated data rate of some 27 gigapixels per second. Given that the ground link is a common data link with a bit rate of 274 Mbps, it is easy to deduce that a huge data reduction must take place on the platform, and that will require at least 10 TeraOPS of processing power.

Another typical ISR processing platform fuses data from a synthetic aperture radar with a step-stare optical camera and other sensors in a package designed for deployment on a SWaP-constrained UAV platform, with a processing load that demands close to a TFLOP of compute power in a compact footprint. The original incarnation of the system used nonrugged commercial servers based on Intel processors and Ethernet interconnects, with software layered over standard math libraries, MPI, and Linux. The migration from this to a fully rugged system was simple because of support of the same MOSA-compliant software elements.

A perfect example of the architecture/technology that underpins MOSA is Linux. Linux-based platforms used in Big Data centers support stringent quality of service and availability imperatives by harnessing open standards developed across a very wide user community and support infrastructure – and can certainly provide the raw compute power required to process large streams of digitized data from high-resolution, real-world sensors. However, such sensor processing subsystems are not typically relied on to deliver time-critical flight control commands aimed at ensuring the flightworthiness of an Unmanned Aerial Vehicle (UAV), for example. It is therefore critical that system architects carefully consider how best to deploy such Linux-based open platform capabilities as subsystem components within deployed ISR platforms.

Fortunately, system integrators can now characterize the performance of ISR processing subsystems using a variety of application development environments. Figures 1 and 2 on page 32 show a task-level event analyzer to assist the application developer in identifying timing anomalies that can affect application-level response times and system variability or jitter. GE's AXIS Advanced Multiprocessor Integrated Software environment lets application developers quickly create worst-case test harnesses running on the deployable HPEC platform before finalizing the system architecture.

Using this environment on a multinode/multiboard system running Linux and open standards-compliant middleware such as the OpenFabrics Enterprise Distribution (OFED) Remote Direct Memory Access (RDMA) driver over 10 GbE, it is possible to demonstrate near wire speed interprocess communication with zero CPU load and typical application-level, memory-to-memory latencies of less than 5 microseconds with worst-case jitter not exceeding 12 microseconds. This performance was achieved running standard Red Hat Enterprise Linux (RHEL) over a multiboard system with 10 GbE data plane. Further performance improvements are possible by moving to RHEL-compatible, real-time versions of Linux.

Such task-level performance analysis enables rapid prototyping of multiprocessor ISR and signal processing systems. Developers can produce evidence such as histograms to demonstrate response times, interprocess communication performance, and jitter over extended operational timeframes to validate system architecture decisions and algorithm design early in the platform development cycle to de-risk the delivery of SWaP-critical tactical systems to their end user customers. These parameters are of particular interest to engineers developing radar and Electronic Intelligence (ELINT) systems. ELINT is particularly sensitive to end-to-end latency of a system, as a characterization of and response to a threat must happen within a tightly defined time window to be effective. Being able to characterize latency and the variation in it is key to validating a system.

The promise of open architecture platforms

Defense and aerospace system integrators can meet the needs of deployed ISR and other signal processing systems by adopting open architecture platforms such as those envisioned by the MOSA initiative. Such a strategy holds out the promise of expanded mission capabilities now and increased performance in the future with the insertion of next-generation hardware from Big Data commerce.

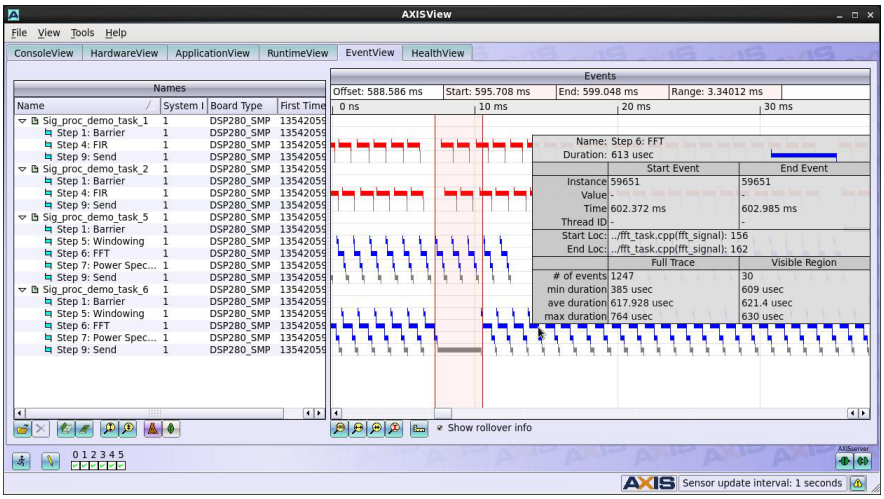

Figure 1 | AXISView – EventView task-level event analysis tool running on an HPEC platform. Screen shot taken before application tuning.

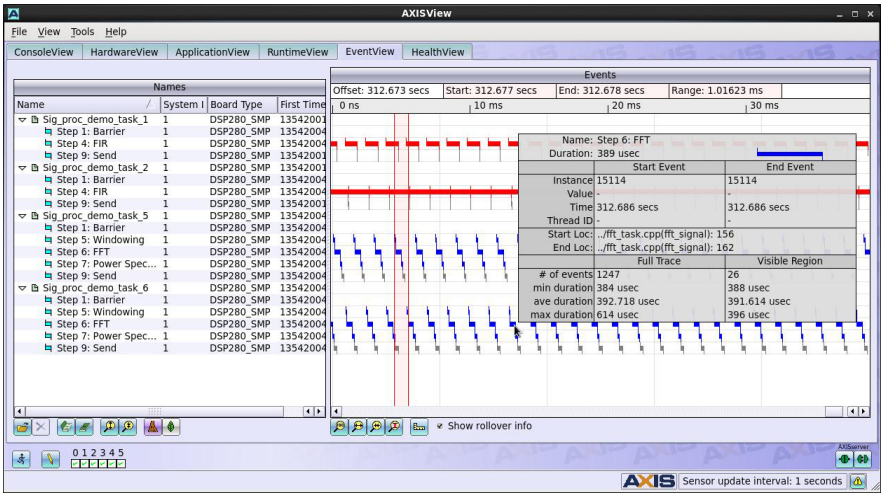

Figure 2 | AXISView – EventView task-level event analysis tool running on an HPEC platform. Screen shot taken after application tuning.

Importantly, where questions exist about the viability of those open standards for mission-critical applications, the availability of sophisticated software tools allows their implementation to be thoroughly evaluated using advanced modeling and optimization techniques. System architects and application developers can accurately characterize system performance while benefiting from wide community support for open operating systems, middleware, and tools in support of multiyear programs. This approach can also enable technology reuse across multiple platforms when applications developed on commercial HPC clusters are rehosted on a range of deployable platforms.

MOSA software is not a destination – it is a journey. **MES**

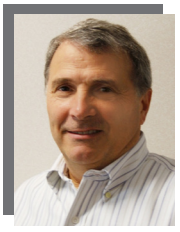

Michael Stern *is a Senior Product Manager at GE Intelligent Platforms, where he is responsible for the company's High-Performance Embedded Computing (HPEC) product line. He has spent 20 years within the defense and aerospace market. He can be reached at michael.stern@ge.com.*

GE Intelligent Platforms • defense.ge-ip.com

OpenSystems Media

works with industry leaders to develop and publish content that educates our readers.

Check out our white papers.

http://whitepapers.opensystemsmedia.com/

Most popular topics:

AdvancedTCA
Android
Avionics Certification
Automotive
Deep Packet Inspection
GUI Linux in Medical Devices
Internet of Things
M2M
Multicore
PCI Express
Radar
SDR
Static Analysis
Switched Fabrics
Test & Measurement
UAVs

MOSA provides cost savings, performance, flexibility for EW digital receivers

By Chris Lewis and Alton Graves

The use of a Modular Open System Architecture (MOSA) can provide cost savings, and performance and upgrade flexibility gains, in digital receiver designs for use in Electronic Warfare (EW). Some approaches to employing MOSA in digital receivers are illustrated, as well as the benefits achieved.

U.S. Army photo by Sgt. Austan R. Owen

The current environment for Electronic Warfare (EW) programs is a challenging one, with new and emerging threats, asymmetrical problem sets, and a rapidly evolving opponent that makes great use of off-the-shelf hardware to rapidly deploy capabilities (for example, garage door openers as IED triggers), to name just a few issues. Traditional EW solutions targeted radar systems and conventional communications networks. They were not designed to counter the unconventional use of modern convenience appliances, nor were they designed to support incorporating the degree of rapid innovation and threat adaptation demonstrated by adversaries. Not only that, historically, digital receiver implementations – including Digital RF Memories (DRFMs) – have been application specific to support very specific limitations or constraints in the upgrade to fielded systems.

Thus, the challenge is integrating the benefits of an application-specific digital receiver with upgrade flexibility and cost savings – while still achieving performance objectives. One approach to reduce EW and signal processing lifecycle costs while shortening the time to deliver solutions to the warfighter has been the adaptation of a Modular Open System Architecture (MOSA) to permit easier technology refresh implementation paths and use of nonproprietary hardware and software solutions (see Sidebar 1). And a MOSA-based digital receiver that relies on open standards *can* still achieve performance objectives, as illustrated in the following discussion.

Open digital receiver challenge examples

As mentioned, digital receivers and DRFMs have typically been application specific; thus, the challenge in implementing a digital receiver into an EW system or application lies in achieving application-specific performance levels while utilizing a cost-efficient, flexible, open standards-based device. For operational effectiveness, a digital receiver needs to have very high spurious performance, high selectivity, and wide dynamic range performance, all of which can be challenging to achieve while fully supporting an open standard interface. A flexible but generally high sample rate and wide instantaneous bandwidth are also required key performance parameters if a digital downconversion is to be employed.

In the examples presented here, two sets of custom, application-specific hardware are replaced with a MOSA solution set with a high degree of commonality between two similar but unique functions. A varying channel count and varying sample rate and resolution on

MOSA explained

MOSA explained

A Modular Open System Architecture is achieved through the use of standard interfaces and operating systems to yield elements to build a system. In the simplest form, MOSA is "plug and play," producing complex interchangeable assemblies. This permits easy scalability and provides a technology upgrade path through the replacement of assemblies that maintain the common interface.

The use of MOSA-compliant designs can help significantly reduce a program's life-cycle costs. At the same time, the designs provide higher performance compared to older systems as well as providing rapid innovation paths to the warfighter. Cost reductions include: 1) The ability to leverage commercial operations to benefit from the larger quantity volumes produced; 2) a reduction in unique designs across multiple targets through design reuse and rehosting; and 3) a reduction in development time through a spiral development process that enhances concurrent development by identifying and reducing development risks.

The use of MOSA designs provides a more rapid development path for new EW and other technologies by enabling replaceable mezzanine/daughtercards. The newest technology can then be folded into a system by removing/replacing the modules and utilizing the existing common interface for I/O and data processing without a wholesale system replacement. This can eliminate the need to dispose of the old system; it also keeps the newest technology available where it is needed most.

MOSA might not be an optimal solution for every situation. Standard interfaces often require more up-front analysis to meet a broad range of applications; as a result, sometimes they can be "over-engineered" and try to address too many applications. The MOSA standard (VME, VPX, VXS, and so on) might not cover all features required. Furthermore, it might be at cross-purposes for range of system performance parameters, or might have capability in excess of that required. Functionality or performance in excess of the requirement in a Size, Weight, and Power (SWaP)-limited environment can be a challenge to value appropriately. Alternatively, the performance might be more readily solvable through a unique technology solution that does not fit within the boundaries of the MOSA standard. In spite of these challenges, MOSA is still an attractive approach to achieving cost savings in electronic warfare digital receivers.

Sidebar 1 | A closer look at the Modular Open System Architecture (MOSA) standard

the Analog-to-Digital Converters (ADCs) are what differentiate the receivers.

SIGINT/ELINT digital receiver example

The first MOSA implementation example comprises two digital receivers: a Signals Intelligence (SIGINT) direction finding receiver and an Electronic Intelligence (ELINT) wideband detection receiver. In each case, the digital receiver base card is identical between the two implementations, with the FMC mezzanines acting as "personality modules" to define the functionality of the open-based architecture (see Table 1 and Figure 1). A common FPGA interface to the FMC mezzanine is easily adaptable for each. This provides the advantage of the FPGA data plane (data movement) infrastructure being common between different applications. Moreover, the control software interface is reusable for both applications, thus permitting the end user to focus on algorithmic development and employ existing data and control plane functions.

SIGINT Direction Finding Receiver	ELINT Wideband Detection System Receiver
8 Channel, 16-bit, 250 MHz IBW	4 Channel, 12-bit, 1,500 MHz S/R
(4) Dual Channel RF Receiver – COTS	(2) Dual Channel RF Receiver – COTS
(1) DDS Synthesizer – COTS	(1) DDS Synthesizer – COTS
(1) VXS Digital Receiver Card – COTS + IP	(1) VXS Digital Receiver Card – COTS + IP

Table 1 | Comparison of two MOSA-based receivers: a SIGINT direction finding receiver and an ELINT wideband detection receiver

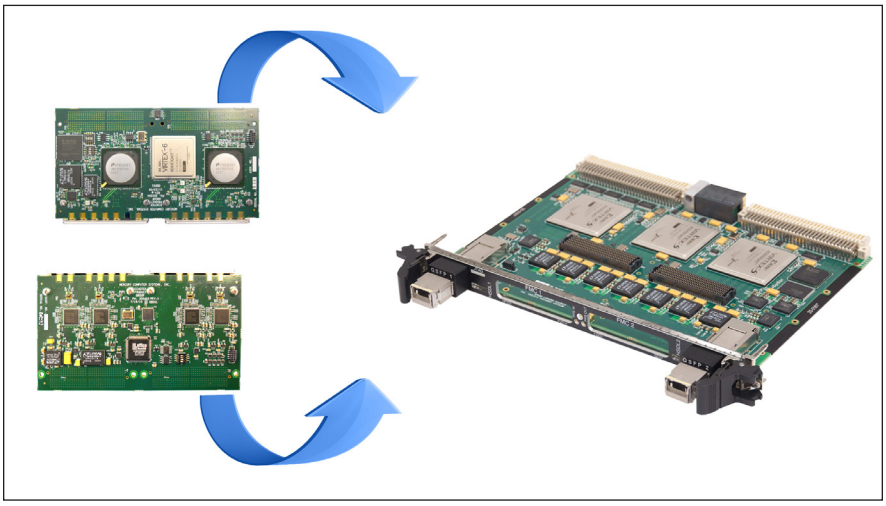

Figure 1 | MOSA SIGINT/ELINT digital receiver

A DRFM/digital receiver example

This example of a MOSA implementation is representative of an application-specific design for an airborne pod/UAV installation. In this system, a digital receiver is to provide signal environment situational awareness when coupled with a DRFM. As a cost-reduction effort over the fully custom configuration, the system is partitioned into: 1) The wideband situational awareness OpenVPX receiver; 2) the OpenVPX DRFM; and 3) the fast tuning RF OpenVPX Down/Upconverter RF control. For this implementation, the digital receiver and DRFM base cards are identical, while the FMC mezzanines are reprogrammed to define the functionality of the open-based architecture for each function. The FPGA interface to the mezzanine is readily adaptable for each FMC.

The open DRFM function needs to have a low latency/insertion delay, while fully supporting an open standard interface. A flexible, but high sample rate and wide instantaneous bandwidth are required, while incorporating a very fine delay resolution, high spectral purity/low spurious content, and wide instantaneous dynamic range.

The digital receiver needs to have very high spurious performance, high selectivity, and wide dynamic range performance, while fully supporting an open standard interface. A flexible, but generally high sample rate and wide instantaneous bandwidth are also required key performance parameters if a digital downconversion is to be employed. In this example, we see a single base card for both the wideband digital receiver and the DRFM. The base card uses a common FMC card for each with unique IP within the FPGA to achieve the performance standards needed.

Future savings with MOSA

The examples provided show that the using a Modular Open System Architecture can provide cost savings in digital receiver designs for use in electronic warfare through: shortening the time to deliver solutions; providing greater flexibility for incorporating feature growth as the systems mature; and serving as nonproprietary hardware and software solutions that provide performance comparable with or exceeding that of application-specific EW digital receivers. **MES**

Open DRFM	ELINT Wideband Detection System Receiver
Single Channel, 12 bit/14 bit, 1,600 MHz S/R (1) FMC Digital Receiver Card – COTS + IP	2 Channel, 12 bit, 1,600 MHz S/R (1) FMC Digital Receiver Card – COTS + IP

Table 2 | Open DRFM function and ELINT wideband detection system can both utilize MOSA

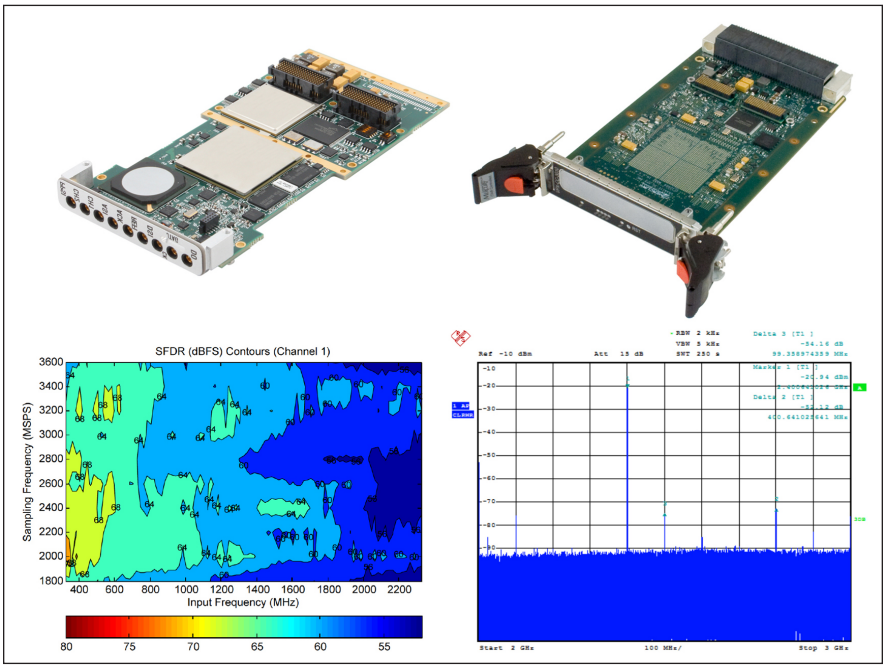

Figure 2 | MOSA 3U VPX digital receiver and DRFM

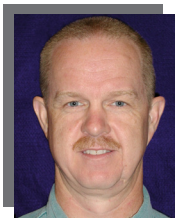

__Chris Lewis__ is Chief Technical Officer at Mercury Defense Systems. He helped cofound Mercury Defense Systems (formerly KOR Electronic) in 1986. His background in the design of high-speed multi-bit A/D and D/A modules contributed to improved performance of flight-qualified DRFM-based EA/EP equipment and radar environment simulator systems. He has prior electronic warfare experience at Design Engineering Laboratories, Hughes Aircraft Company, and Boeing Aerospace Company. He can be contacted at CLewis@Defense.MRCY.com.

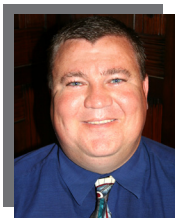

*__Alton Graves__ is Technical Director for Microwave and Digital Systems at Mercury Systems and serves as the Lead Hardware Architect for IF & Digital products within Mercury Systems' Microwave and Digital Solutions Group. Alton was a key contributor in the architecture for Mercury's (*VME and Critical Systems *magazine Editor's Choice) award-winning SCFE product. He has 15 years of experience designing mixed-signal circuitry and FPGAs. Before joining Mercury in 2003, he was a senior engineer for a major networking company. Prior to that, he served as both a Xilinx and Altera field applications engineer. He can be contacted at AGraves@MRCY.com.*

Mercury Defense Systems is an operating subsidiary of Mercury Systems, Inc.
www.mrcy.com

Radar Product Spotlights

Fast and Affordable Microwave Frequency Analysis

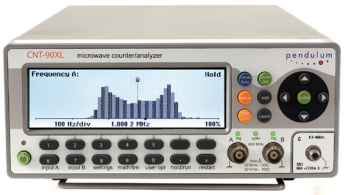

- Any pulsed or CW radar, up to 60 GHz
- Pulse width, PRI/PRF in burst down to 30 ns
- Frequency in burst down to 100 ns, power in burst
- High resolution analysis, track fast frequency changes
- Intra-pulse frequency profiling
- Battery option

spectracom

Spectracom
+1 585-321-5800
sales@spectracomcorp.com
www.spectracomcorp.com

Ultra-High Performance Memories from GSI Technology

- High Speed SRAM & Low Latency DRAM
- Full Military Temp Catalog with Long Production Lifespans
- Leaded and Lead-free BGA Packages
- On-Die ECC Options
- Up to 675 MHz Operation
- For More Info: http://www.gsitechnology.com/Patton/Military_SRAM_GSI.htm

GSI TECHNOLOGY

GSI Technology
www.gsitechnology.com
408-331-8800
MilTemp@gsitechnology.com

Altera® Stratix® V GS FPGAs

Enabling next-generation DSP radar and EW systems

- IEEE 754 compliant floating point operator
- Largest portfolio of optimized floating point IP cores
- Variable precision floating point DSP Block
- 28Gbps integrated transceiver
- Up to 50% lower power transceivers
- Advanced comprehensive radar and EW reference design

ALTERA® MEASURABLE ADVANTAGE™

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

Rakon's Digital Pulse Compression Sub-Systems (DPCSS)

- Digital technology for radar maintenance and life extension.
- Military and civil SAW based pulse compression radars.
- Low cost off-the-shelf expander and compressor unit.
- Customized unit, form, fit and function compatible with existing units.
- Improves overall system performance. High instantaneous dynamic range.
- High flexibility (programmable parameters), reproducibility and stability.

rakon

Rakon
www.rakon.com
sales@rakon.com

WILD Storage Solution

- 2, 4 or 8 Terabytes per OpenVPX Slot
- 4 GBytes per Second Write and 8 GBytes per Second Read
- Up to 40 Gbit Ethernet or QDR InfiniBand on each of 4 Fat Pipes on P1 for a total of 20 GBytes per Second.
- PCI Express 8x Gen 1, Gen 2 or Gen 3 on P2 and P5 of the OpenVPX Backplane
- Hot Swappable Canister, Up to 16 Pluggable 1.8" SSD SATA 3.x
- API for Command and Control of the Storage Process

Annapolis Micro Systems

Annapolis Micro Systems
410-841-2514
www.annapmicro.com

WILD OpenVPX Switch

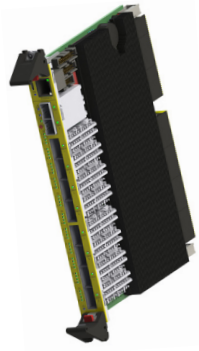

- Multiprotocol Switch – SDR/DDR/QDR/FDR InfiniBand and 1/10/20/40 Gb Ethernet
- Up to 4 Terabytes per Second Non-Blocking Switching Capacity with up to 8 Switch Partitions
- 6U OpenVPX Board – Support OpenVPX Switch Profiles
- InfiniBand and IP Routing, Ethernet Gateways
- Front Panel: Up to 8 QSFP+, Up to 2 SFP+, RJ45 Management Port, USB USART, LED Status

Annapolis Micro Systems

Annapolis Micro Systems
410-841-2514
www.annapmicro.com

Virtex-7 FPGA technology boosts radar performance

By Rodger Hosking

The latest Virtex-7 FPGAs from Xilinx deliver significant benefits for radar systems compared to previous-generation devices. Higher-density silicon with lower power consumption, more resources, faster interfaces, and more and faster memory not only enhances performance, it also opens new application spaces.

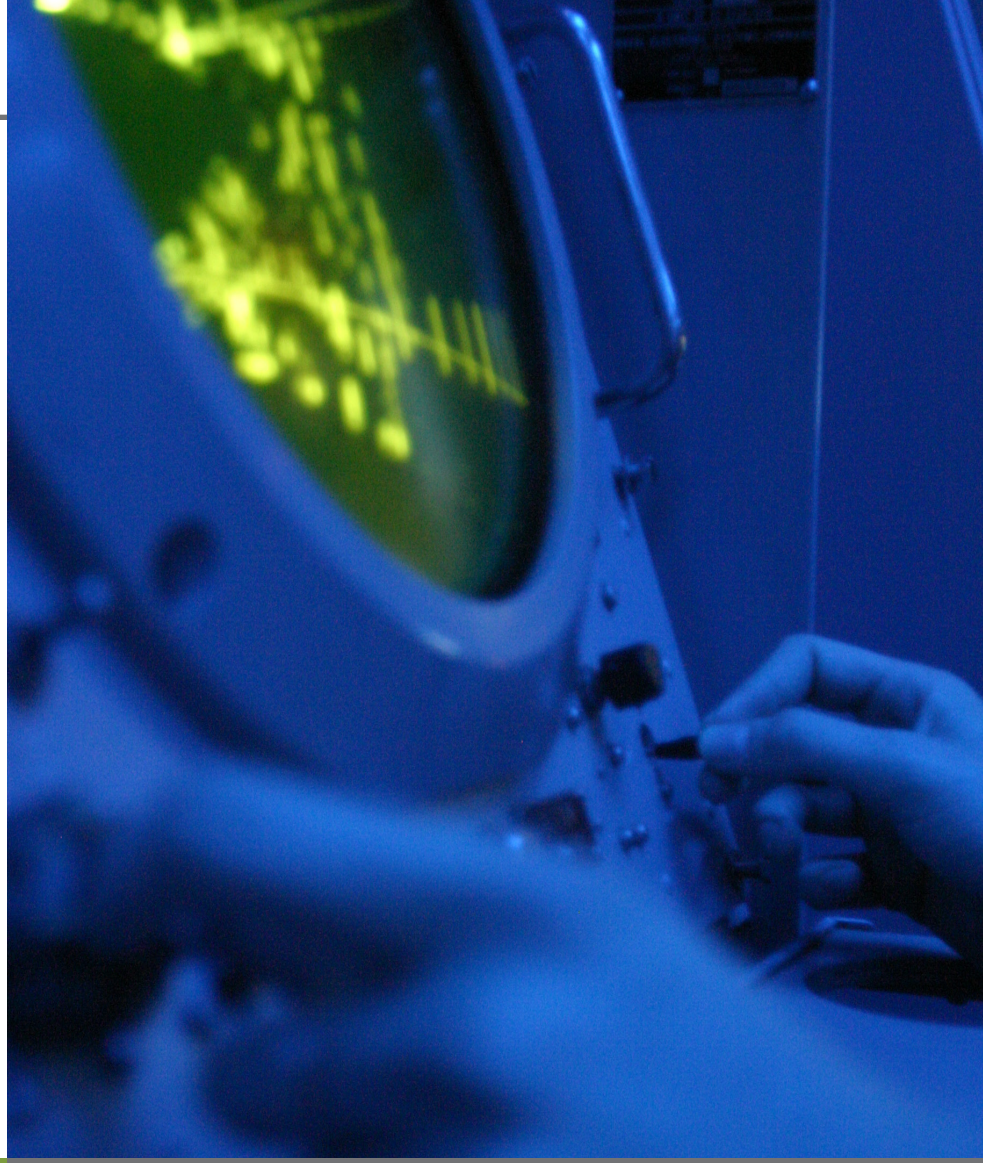

U.S. Marines photo by Lance Cpl. W. Zach Griffith

Since the advent of the first military radar systems more than 80 years ago, engineers have continuously enhanced capabilities and performance levels by harnessing new technology and components. The perpetual leapfrog race between improved detection and better countermeasures offers a steady stream of recurring opportunities for system designers to deploy advanced techniques and algorithms with increasingly higher levels of sophistication.

Because of their inherent reconfigurability and real-time signal processing resources, FPGAs are almost always the critical link for advancing radar technology to the next level. Xilinx's latest generation of FPGAs, the Series 7, consists of three families addressing a range of price and performance markets, the most powerful being the Virtex-7. By using a new 28 nm High Performance Low power (HPL) process technology, the Virtex-7 boasts twice the performance and half the power consumption of its predecessor, the Virtex-6. Figure 1 summarizes relative resource comparisons between the two families, clearly showing dramatic improvements.

Faster peripheral interfaces, more DSP engines and block RAM, enhanced memory interfaces, and faster PCIe interfaces in the Virtex-7 all directly impact radar performance. At the same time, power dissipation for a given function has dropped by half compared to previous-generation devices, opening up new applications for smaller unmanned vehicles.

Faster A/D and D/A interfaces

Figure 2 shows a generic block diagram illustrating the basic functions of a typical FPGA-based radar receiver. An analog RF tuner translates the antenna frequency band of interest down to a lower IF frequency that can be sampled by an A/D converter. Thereafter, all signal processing, control, storage, and system interface functions are handled by the FPGA.

Signal bandwidths in modern radar systems can exceed 500 MHz, pushing the sampling rates for A/D and D/A converters well into the GHz range in some cases. These increasingly higher rates not only challenge interface speeds of FPGAs, but various voltages levels and clocking schemes used in the A/Ds also present special data packing and formatting requirements.

Very high-speed A/D converters demultiplex the output samples onto multiple buses to gain a more manageable bus clock rate. To support them, the Virtex-7 FPGAs offer a direct connection with LVDS DDR I/O transfer rates reaching 1,600 MHz, up from 1,400 MHz

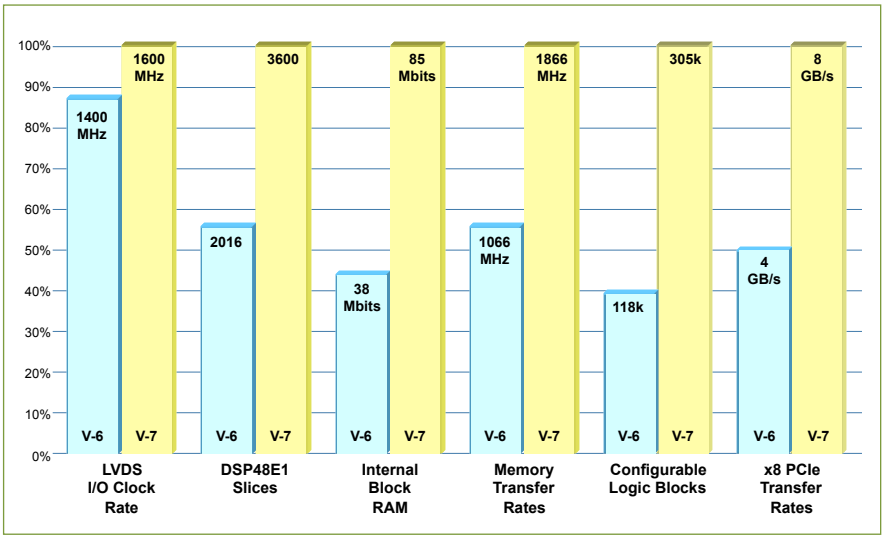

Figure 1 | Radar resources comparison between Virtex-6 and Virtex-7

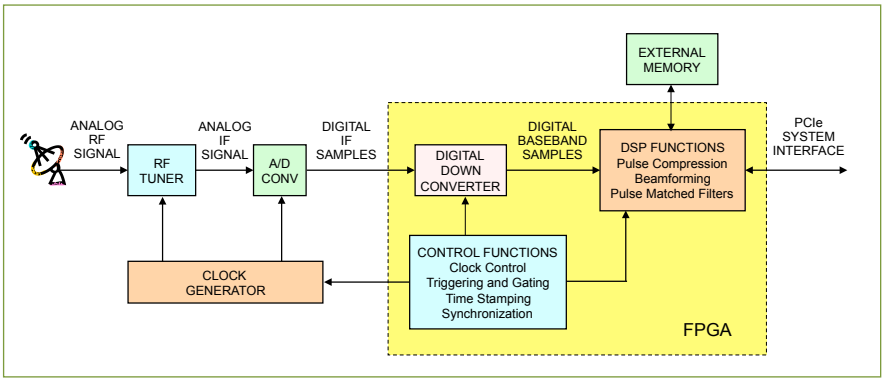

Figure 2 | Functional blocks of a typical FPGA-based radar receiver

in the Virtex-6. Now, for example, a four-way demultiplexed A/D can operate at up to 6.4 MSamples/sec, so it can digitize a wideband radar signal with an instantaneous RF bandwidth approaching 3 GHz.

These new highly configurable Virtex-7 interfaces include per-bit skew adjustments to help align bits in a data word to ease stringent trace length matching in printed circuit board design. Also, digitally controlled termination networks eliminate the need for external discrete impedance-matching resistors to achieve reliable data connections.

The architecture for FPGA-based radar transmitters is analogous, but with inverse signal flow structures including D/A converters, RF upconverters, and power amplifiers. Similar requirements for extremely fast interfaces to wideband D/A converters benefit from the same Virtex-7 speed and configurability improvements discussed.

Enhanced DSP resources for radar

Transmitted radar signals are becoming increasingly sophisticated to improve performance in target range, position and classification, resistance to clutter for low-altitude targets and weather, and resistance to electronic countermeasures and detection. Outgoing radar pulses must be precisely crafted for amplitude, frequency, and phase, and often these characteristics must change dynamically to adapt to different targets, trajectories, or conditions.

Complex outgoing radar pulses previously required the use of signal synthesis tools to precompute the waveforms, which were stored in memory for D/A playback. FPGA DSP engines can now be harnessed to compute the required transmit waveforms in real time, a major benefit for intelligent, adaptive radar systems.

Different signal processing is required for received radar signals. Some of the more important algorithms are digital downconversion, pulse matched filters and pulse compression, forward and inverse FFTs, windowing, and matrix operations. Before FPGAs took over the job, massive DSP subsystems were required to handle these compute-intensive operations because they had to be computed in real time.

Additionally, one of the most critical DSP operations for radar is beamforming, which governs the operation of phased array antennas, both for transmitting and receiving. By carefully adjusting the relative phase of signals for each antenna element, the antenna array can be electronically steered for maximum signal strength in any direction. The phase shifts are introduced electronically in the transmit and receive paths of each element through DSP operations easily implemented in the FPGA, eliminating the need for cumbersome mechanical structures.

Accordingly, Virtex-7 FPGAs use the same DSP48E1 engines first introduced in the Virtex-6 family, each containing a pre-adder, a 25 x 18 multiplier, an adder, and an accumulator. However, to support all of these increasingly demanding signal processing requirements for complex, wideband radar signals, the quantity of DSP48E1 engines in the Virtex-7 has been boosted to 3,600. This represents an increase of nearly 80 percent more than the largest Virtex-6 device, directly enhancing each of the critical transmit and receive capabilities needed for the latest radars.

Faster and deeper memory

Fast local memory provides an enormous benefit to receive-side signal processing because most radars use range gating to capture the reflected signals during a specific time window relative to the

outgoing pulse. The ratio of the range gate interval to the pulse repetition period determines the acquisition duty cycle. By storing the incoming data samples from the A/D converter in real time during the range gate, DSP blocks have more time to complete the signal processing tasks, since no new data will be acquired until the next range gate. This can dramatically improve the efficiency of the DSP hardware by several times, depending on the duty cycle.

The largest Virtex-7 devices now offer more than 85 Mb of internal block RAM for intensive on-chip DSP operations, more than twice as much as Virtex-6. But transient capture and buffering of large blocks of data require external memory that can keep up with the A/D sample rates.

Synchronous DRAMs deliver extremely fast read/write rates by transferring data on both edges of the clock. They also offer the densest and most economical solution for large memory arrays. The latest Virtex-7 devices can support DDR3 devices running a bit transfer rates up to 1.866 Gbps, far above the 1.066 Gbps for Virtex-6.

The Virtex-7 achieves these speeds by boosting the maximum 1:2 ratio between the fabric logic clock and the memory transfer rate on the Virtex-6 to 1:4 on the Virtex-7. The Virtex-7 also introduces the Phaser clock generator to maintain real-time clock-to-data timing to within 7 psec for direct, glueless connections to these fast memories.

Timing is everything

All radar systems require precise timing between the transmit pulses and the acquisition of the return signals at the receiver. FPGAs abound in configurable logic resources ideal for state machines and counters to create extremely sophisticated timing scenarios for range gating and adaptive algorithms to maximize target tracking and identification.

> "The quantity of DSP48E1 engines in the Virtex-7 has been boosted to 3,600. This represents an increase of nearly 80 percent more than the largest Virtex-6 device, directly enhancing each of the critical transmit and receive capabilities needed for the latest radars."

Coupled with agile beamforming capabilities, FPGAs can implement multiple target tracking algorithms using a single fixed array.

The number of available configurable logic blocks in the largest Virtex-7 device is 305,400, compared to the Virtex-6 maximum quantity of 118,500. This 2.5-fold increase can also be taken advantage of to implement new features such as GPS receiver time stamping of received radar pulses to capture not only the precise acquisition time, but also geographic coordinates.

Zero bottleneck system interface

A fast system interface ensures no data bottlenecks for wideband radar signals on the path to or from a downstage processor or storage facility. Because PCIe has become a de facto industry standard interface solution for embedded system boards, Xilinx includes silicon PCIe interface blocks incorporating the key layers of the PCIe protocol stack and fast gigabit serial transceivers.

The Virtex-6 PCIe interface accommodates PCIe Gen 2 x8, delivering a peak rate of 4 GBps using a 5 GHz bit clock and 8B10B encoding on each of the eight bidirectional lanes. With faster transceivers and enhanced circuitry, the Virtex-7 uses an 8 MHz bit clock and 64/66 channel coding to support PCIe Gen 3 x8, delivering a peak transfer rate of 8 GBps, doubling throughput to the system.

Virtex-7 scores high marks for radar

Virtex-7 delivers key improvements in peripheral speed, DSP resources, internal and external memory size and speed, logic resources, and PCIe system interfaces, while at the same time dropping power consumption. Again, the perpetual leapfrog race in radar between improved detection and better countermeasures offers significant new opportunities for system engineers to take advantage of each of these new features. An example is Pentek's Virtex-7 XMC module, targeting wideband radar applications with three 200 MHz 16-bit A/Ds, two 800 MHz 16-bit D/As, and 4 GB of SDRAM. It takes advantage of all of the resources discussed and represents a quantum step in performance over earlier Virtex-6 products. **MES**

Rodger H. Hosking *is Vice President and cofounder of Pentek, Inc., where he is responsible for new product definition, technology development, and strategic alliances. With more than 30 years in the electronics industry, he has authored hundreds of articles about digital signal processing. Prior to his current position, he served as Engineering Manager at Wavetek/Rockland, and he holds patents in frequency synthesis and spectrum analysis techniques. He can be contacted at rodger@pentek.com.*

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

Enabling advanced designs for radar and Electronic Warfare systems with FPGAs

MIL EMBEDDED: ***What are your customers looking for from signal processing technology in radar and electronic warfare applications?***

HU: In radar and Electronic Warfare (EW) applications, the key signal-processing-related parameters are low latency, large data volume, high data rate, rate conversion, signal estimation, and target tracking capabilities. For this discussion, let's look at the dynamic range, which greatly impacts system performance and capabilities.

The dynamic range of the signal path in radar and EW systems can easily exceed the standard 16-bit or even 32-bit twos-complement range that is generally available in Digital Signal Processing (DSP) blocks for fixed-point math. As a result, many DSP blocks now also provide support for floating-point data types.

In an FPGA, the trade-off is more flexible than a DSP block because the bit width is not limited to standard 16- or 32-bit fixed-point data types. In the past, an algorithm that was developed using floating-point math was first translated into fixed-point math for implementation in an FPGA. This required careful scaling the dynamic range in the signal path, typically allowing the bit width to grow as necessary. However, excessive bit widths impact FPGA resources and speed. This was a time-consuming part of the overall FPGA design process. In some applications, floating-point math is needed to support the high dynamic range, as we will discuss later.

Until recently, floating-point math was not practical in FPGA devices. It is well known that the efficiency of floating-point calculations can be improved dramatically by grouping operations, known as *fused operators*. The efficiency is due to eliminating redundant normalization and denormalization logic. Published results suggest a 40 percent relative improvement in both resource utilization and power dissipation. The Datapath Compiler with IEEE 754 floating-point compliance provides the same benefits of fused operators to Altera® FPGA designs.

Doppler processing requires the radar to be coherent. It measures the relative changes in phase on returned signals. Pulsed radars compromise between Doppler ambiguity and range ambiguity by varying the Pulse Repetition Frequency (PRF). The PRF must be twice the largest Doppler frequency (fastest target) to meet the Nyquist rate, so the higher the PRF, the less aliasing/ambiguity in the Doppler domain and the higher the range ambiguity. The range ambiguity is optimal when the Pulse Repetition Interval (PRI) length matches the round-trip time of the pulse to the target.

The pulse compression output feeds data into a large storage memory that holds the number of PRIs required and implements coherent integration on each range gate. The memory size is driven by the largest image size possible. Once the integration dwell is completed, with or without overlap, a corner-turn operation is performed and Fast Fourier Transforms (FFTs) are implemented on each range cell to create the range-Doppler image. As the dwell gets longer (FFT increases), the coherent integration gain goes up. The downside of this increase is that it also increases the latency. Most of this latency occurs due to the collection time of more PRIs. Larger FFTs can be implemented on smaller data sets with zero padding to increase the resolution of the Doppler offset. Most importantly, the larger the FFT, the finer the frequency resolution; thus, the bit width growth quickly exceeds the fixed-point arithmetic limit in larger FFT operations. Floating-point arithmetic is an ideal solution to address this challenge.

Most coherent radars implement overlapping of the input image to catch transients at the image boundaries. This overlapping must be traded off with the increased processing time and increased data rate. An FFT cannot be processed for conversion until the first point of the last PRI is collected. The FFT processing must be fast to minimize its latency contribution. In the worst case, an FFT cannot take more than the collection time for the next image or fraction of image in an overlapped case. Furthermore, the FFT produces a sampled version of a continuous frequency variable that must be interpolated to determine finer frequency components. The output of this processing yields the range and frequency of the targets in the radar beams.

A look at how Doppler processing works quickly reveals how the need for high-performance, low-latency, large-memory, floating-point, and DSP capabilities can greatly enhance system performance. Integration of these capabilities into one component assures more system-level innovation and advancements in future generations of radar and EW systems.

Ching Hu *is Senior Strategic Marketing Manager, Military Business Unit, at Altera Corporation. Contact him at ching.hu@altera.com.*

Prospects of COTS products for future radar systems – challenges and opportunities

MIL EMBEDDED: ***What are your customers looking for from signal processing technology in radar and electronic warfare applications?***

SCHAUFLER: Radar systems are built out of three principal elements: 1) an ADC that acquires the data, 2) high-speed interconnects that deliver that data to 3) the processing elements. For system providers, this poses two essential challenges: massive computing power and very high-speed interconnects. In real-world applications, the processing aspect is typically implemented by CPUs, FPGAs, or DSPs. Some systems focus the processing element on one of these, while others achieve their goals through a combination of processing technologies.

Straightforward system concepts link a large number of CPUs and/or FPGAs through a high-speed protocol, for example, PCIe or SFPDP, making them an ideal candidate for using VXS or VPX COTS hardware. However, an element easily underestimated in these concepts is the software component. Drivers and easy-to-use APIs are key elements for success on these systems.

In more FPGA-centered radar systems, the CPU's function is typically limited to controlling functions. Compared to CPU boards, FPGA boards pose less design efforts. Yet, combining FPGA code from two suppliers tends to be a much harder challenge than integrating software from more than one vendor. For CPU boards, COTS hardware is a straightforward option. However, COTS FPGA boards require a close relationship between the hardware supplier and the system integrator.

Legacy aspects are also an important factor. For example, VXS radar designs are still being upgraded for higher computing power. For smaller systems, the limited number of high-speed connections often does not pose a problem. The limiting factor here is the power envelope, which blocks the use of modern high-performance CPUs. However, as long as there is an upgrade path, there are good reasons to extend the life of proven designs. Low risk and low cost are the key factors here. When the market requires higher and higher performances at a stable or even declining system cost, COTS solutions have good chances to deliver.

With 10G Ethernet evolving rapidly, it might be a very interesting option for the transport protocol in radar systems. Ethernet, as a proven design for this level of performance, could be an answer to cost pressure, due to wide availability. The availability of 10G Ethernet in the PC world allows a development process where a large part of the software implementation can be done on relatively cheap PC hardware. Building blocks can then easily be moved to the embedded hardware, one at a time. Together with the wide availability of debugging tools, this could come close to the ideal platform for radar system developers and integrators. However, a question that remains open is connectivity between FPGAs and CPUs over 10G Ethernet.

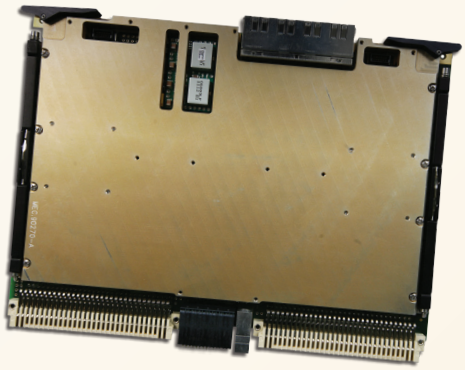

Figure 1 | The RIO5-8088 from CES – Creative Electronic Systems

CES has been providing leading-edge solutions based on PowerPC® processors and a proprietary FPGA-based high-speed interconnect protocol that allows direct connections between two FPGAs, as well as between an FPGA and a CPU. Recent real-world deployed radar systems have seen the introduction of multi-core technology, PCIe and SFPDP in the RIO5 generation of CES products (Figure 1). Among others, a key element has been CES' combined know-how on hardware, FPGAs, operating systems (Linux®, VxWorks®, and others) and software driver architecture to create the optimal mix of performance, efficiency and end-user experience. Following the announcement of Freescale's next generation of QorIQ™ processors, which provides a large number of cores with AltiVec support, an upgrade path based on the Power Architecture is now open for current VXS designs, as well as for new high-performance systems in VPX format. These CPUs also include a 10G Ethernet-enabled offload engine, which makes them ideal candidates for the next-generation interconnect systems. CES is also developing VPX FPGA boards using Xilinx Kintex-7 technology as a successor to previous PMC-based FPGA boards. The introduction of the FMC standard will aid to save engineering costs by allowing the adaptation of FPGA I/Os through a simple "piggy board."

Together with high-performance CPU boards based on Freescale and Intel technology, as well as optimized software, the building blocks are complete to build future CPU-centric, FPGA-centric or mixed radar systems based on the latest available technology.

Andreas Schaufler *is Project Manager at CES, where he focuses on COTS SBC product development including HW, SW and FPGA aspects. Contact him at ces@ces.ch.*

CES • +41.22.884.51.00 • www.ces.ch

Radar retrofits, subsystems are key in today's market

MIL EMBEDDED: ***Will there be more opportunities for embedded COTS suppliers in radar upgrades and new platforms?***

TATOPOULOS: Radar manufacturers commonly use components or subsystems from Commercial Off-the-Shelf (COTS) suppliers. This principle allows customers to access the latest technologies while optimizing development cost, time, and mitigation of component obsolescence. For embedded COTS suppliers, dealing with such projects means to assure delivery and maintenance of products for many years.

From the user side, due to the huge investment that represents new platform purchasing, the key factor is to increase radar equipment lifetime and performance. This requirement is even more important today due to the tough economic period, which has drastically decreased budgets. Consequently, radar upgrade and life extension has become a compelling trend in the radar market: It is always better replacing one subsystem rack to extend equipment life for 10 or 20 years rather than replacing the complete system.

Rakon, a global high-technology company that designs and manufactures world-leading frequency control solutions, has been a key supplier to the radar retrofit market for many years. Rakon's expertise in this technology domain has identified a critical part in which obsolescence creates trouble in radar maintenance: the Surface Acoustic Wave (SAW) pulse compression subsystem. In pulse compression radar, the expander (which generates chirp into the transmitter path) and the compressor (for pulse compression into the receiver path) are usually done using SAW dispersive delay lines. However, these very specific components are hardly replaceable for radar developed 20 or 30 years ago! To answer this problem, Rakon has developed a Digital Pulse Compression Subsystem (DPCSS) to be the ideal solution when retrofitting or upgrading existing SAW-based pulse compression radars (Figure 1). Not only does it enable a remarkable improvement in the overall system performance and extend the life of the system, but it also can be form, fit, and function compatible with existing SAW or digital units – making it easy and convenient to use in all military and civil radars.

Thanks to a powerful and scalable FPGA-based digital technology, the performance achieved is close to the theoretical limits, enabling improved aircraft detection. This means a higher instantaneous dynamic range (side lobe level 5 dB better than SAW-based systems), higher accuracy angular measurements (main lobe without base widening), higher flexibility (programmable parameters enabling customization in waveforms and software upgrades), and repeatable performance from one device to the other as well as along the operating temperature range.

Rakon's DPCSS has been successfully implemented in several military and ATC radar upgrade projects since 2006 and is expected to become the leading solution to answer the growing demand in radar maintenance and life extension.

However, radar upgrades are only a starting point within the global radar market as further requirements have been identified. The increasing demand for high-value-add subsystems is accelerating the transformation of radar manufacturers towards a systems integrator role and as a way to access innovative technologies that otherwise would have been out of investment range.

As a key technology partner with strong experience acquired in the field of radar retrofit, Rakon is able to provide additional embedded functions, including signal generation in transmitter part, IF amplification, pulse compression, and detection on receiver parts. Combined with Rakon's expertise in the frequency control domain – where ultra-low noise OCXO (BAW technology for 10-100 MHz band) and OCSO (SAW technology for 300-2,000 MHz band) can be provided as very high-performing local oscillators – Rakon has a portfolio of leading solutions for radar modernization programs and new platform developments.

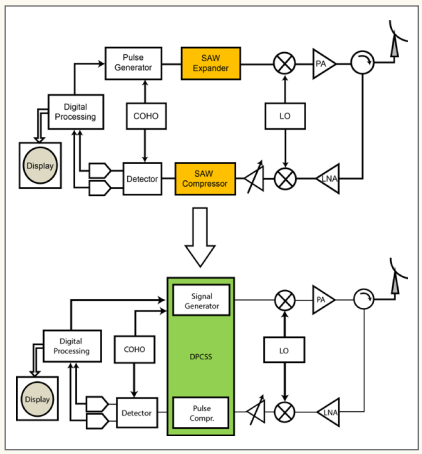

Figure 1 | Rakon's Digital Pulse Compression Subsystem (DPCSS) is an upgrading solution for SAW-based radars.

This new level of integration into a so-called Digital Subsystem (DSS) simplifies the global radar architecture, thus reducing mechanical size (less racks), cost, and maintenance.

In the near future, Rakon's capability to integrate even more analog functions into a highly integrated digital subsystem will offer the radar market even more sophisticated options.

For the radar manufacturer – beyond the economic interest, flexibility, and strong performance improvement – the significant advantage resides in obsolescence management. Based on well standardized tools and platforms, Rakon's digital solutions drastically lower obsolescence risk or at least keeps it easily manageable. For COTS suppliers, the main challenge will be to keep delivering technological advances combined with market-competitive commercial solutions to meet the demanding requirements of the radar manufacturers.

Xavier Tatopoulos *is R&D Manager at Rakon France.*

rakon

Rakon • info@rakon.fr
www.rakon.com/products/families/pcss

Coprocessor solution improves FPGA efficiency

Dubbed "Anemone accereration technology," a coprocessor floating-point signal processing solution from BittWare in Concord, NH, mitigates FPGA programming challenges by offloading C-language processing tasks from an FPGA. The FPGA handles all the I/O interfacing, memory, protocol processing, and special functions, as well as any computational tasks it might need to perform, leaving Anemone free to handle more complex processing functions. Anemone is an implementation of the Epiphany architecture from Adapteva in Lexington, MA. BittWare's Anemone family leverages Altera Stratix FPGAs in a PCIe product line, OpenVPX-based products, and a VITA 57 FMC device.

The S5-6U-VPX (S56X) is a rugged 6U VPX card using the Altera Stratix V GX/GS FPGA. The FPGA can be paired with BittWare's Anemone FPGA coprocessor, the ARM Cortex-A8 control processor, and the ATLANTiS FrameWork FPGA development kit, for high-performance signal processing and data acquisition applications such as radar and signal intelligence. The board includes a configurable 48-port multigigabit transceiver interface supporting protocols such as Serial RapidIO, PCI Express, and 10 GbE. Two optional Anemone104 processors are available for the S56X as coprocessors for the Stratix V FPGA. BittWare also offers an S5-6U-VPX device with FMC I/O and has a 3U OpenVPX product coming soon.

BittWare | www.bittware.com | www.mil-embedded.com /p370840

VICTORY-compliant rugged switch enables vetronics networks

Combining a 16-port GbE network switch with a high-performance vehicle management computer, the rugged Digital Beachhead vehicle networking system from Curtiss-Wright Controls Defense Solutions enables users to quickly introduce VICTORY-compliant network hardware into military ground vehicles – and all for less than $5,000 and burning less than 20 W of power. It is designed specifically to meet the guidelines of the U.S. Army's new Vehicular Integration for C4ISR/EW Interoperability (VICTORY) initiative for interoperable digital network services. Driven by a tight budget environment, Digital Beachhead leverages commercially available technology from other industries – network switches for cell towers and ARM processors – in an open systems approach. Its ARM processor-based system computer manages the health management (HUMS/CBM+) software on the vehicle and monitors the vehicle's primary systems.

Housed in a compact, rugged, lightweight enclosure, Digital Beachhead comes preinstalled with Curtiss-Wright's Vehicle Management Framework (VMF) software and has 16 ports of tri-speed 10/100/1000BASE-T Ethernet and network switching features. Curtiss-Wright also uses the Selective Availability/Anti-Spoofing Module (SAASM)-based MPE-S GPS receiver module from Rockwell Collins in Digital Beachhead.

Curtiss-Wright Controls Defense Solutions | www.cwcdefense.com/digital-beachhead | www.mil-embedded.com /p369574

New SFF device proves multivendor OpenVPX interoperability

OpenVPX has promised that the "Open" part meant that VPX systems from different suppliers could be interoperable and work in one system – preventing integrators from being tied to one vendor. The new SigPro 1 Small Form Factor (SFF) acquisition system that combines OpenVPX devices from Elma Electronic, Pentek, and Concurrent Technologies has proven that a multiparty VPX box not only works, but also can be developed in a short period of time. The multivendor platform is targeted at radar, signal processing, high-speed data acquisition, and beamforming applications.

SigPro 1 is based on the Pentek Cobalt FPGA XMC products with a high-performance analog-to-digital front end and the 803 Series 3U VPX-REDI Core i7 Single Board Computer (SBC) from Concurrent Technologies for processing functions and a storage subsystem that is based on a solid-state disk from Elma. The new solution uses a three-slot VPX backplane in a conduction-cooled chassis and uses sealed MIL-STD DTL 38999 connectors for I/O. The system can record at more than 200 MBps and has two GbE ports, two USB ports, and an external VGA port. Any Pentek FPGA card can be implemented in the system by using Pentek's Talon data recording software.

Elma – www.elma.com | Pentek – www.pentek.com | Concurrent Technologies – www.gocct.com

www.mil-embedded.com/p370841

Nürnberg, Germany

we are the future

February 26 – 28, 2013

embedded world 2013

Exhibition & Conference

... it´s a smarter world

Register now for your free entrance ticket:

embedded-world.de

Programmed for success.

With some 900 exhibitors, embedded world is the top pioneering exhibition for embedded technologies worldwide. Make a note of the date now!

Exhibition organizer
NürnbergMesse GmbH
Tel +49 (0) 9 11. 86 06-49 12
visitorservice@nuernbergmesse.de

Conference organizer
WEKA FACHMEDIEN GmbH
Tel +49 (0) 89. 255 56-13 49
info@embedded-world.eu

Media partners

MEDIZIN elektronik
Fachmedium für Elektronik in der Medizintechnik
elektroniknet.de
computer-automation.de
energie-und-technik.de
MEDIZIN-UND-elektronik.DE

DESIGN & ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER
Markt&Technik
Die unabhängige Wochenzeitung für Elektronik
Elektronik
Fachmedium für industrielle Anwender und Entwickler

Elektronik automotive
Fachmedium für professionelle Automobilelektronik
ENERGIE & TECHNIK
Fachmedium für Energieeffizienz
Computer & AUTOMATION
Fachmedium der Automatisierungstechnik

NÜRNBERG MESSE

GE
Intelligent Platforms

High-performance RADAR with out-of-box readiness.

In today's rapid prototyping environment, RADAR system developers cannot afford to let performance and scheduling risks impact their projects. GE's rich subsystem development heritage coupled with our proven COTS components and state-of-the-art AXIS integrated software development tools virtually eliminate many avenues of risk. Our RADAR subsystems are shipped with high manufacturing and technology readiness levels that speed system development and reduce program costs.

Let our COTS-based computing and signal processing products help you boost the readiness of your next RADAR project.

defense.ge-ip.com

imagination at work

Critical Recording in Any Arena

When You Can't Afford to Miss a Beat!

Introducing Pentek's expanded line of Talon™ COTS, rugged, portable and lab-based recorders. Built to capture wideband SIGINT, radar and communication signals right out-of-the-box:

- Analog RF/IF, 10 GbE, LVDS, sFPDP solutions
- Real-time sustained recording to 2.2 GB/sec
- Recording and playback operation
- Analog signal bandwidths to 1 GHz
- Shock and vibration resistant Solid State Drives
- GPS time and position stamping
- Hot-swappable storage to Windows® NTFS RAIDs
- Remote operation & multi-system synchronization
- SystemFlow® API & GUI with Signal Analyzer
- Complete documentation & lifetime support

Pentek's rugged turn-key recorders are built and tested for fast, reliable and secure operation in your environment.

Call 201-818-5900 or go to www.pentek.com/go/mestalon for your FREE online *High-Speed Recording Systems Handbook* and *Talon Recording Systems Catalog*.

TALON

PENTEK
Setting the Standard for Digital Signal Processing

25 YEARS

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