OpenSystems Publishing

NEW! DAILY BRIEFING: *NEWS SNIPPETS*

The COTS Technology Authority

Military EMBEDDED SYSTEMS

VOLUME 3 NUMBER 6
NOV/DEC 2007

IN THIS ISSUE:

Chris A. Ciufo
MILCOM exposé

Joe Pavlat
The U.S. Navy and AdvancedTCA

Duncan Young
All-around awareness

WWW.MIL-EMBEDDED.COM

COTS over Mars

ALSO:

Harsh environments, cool systems

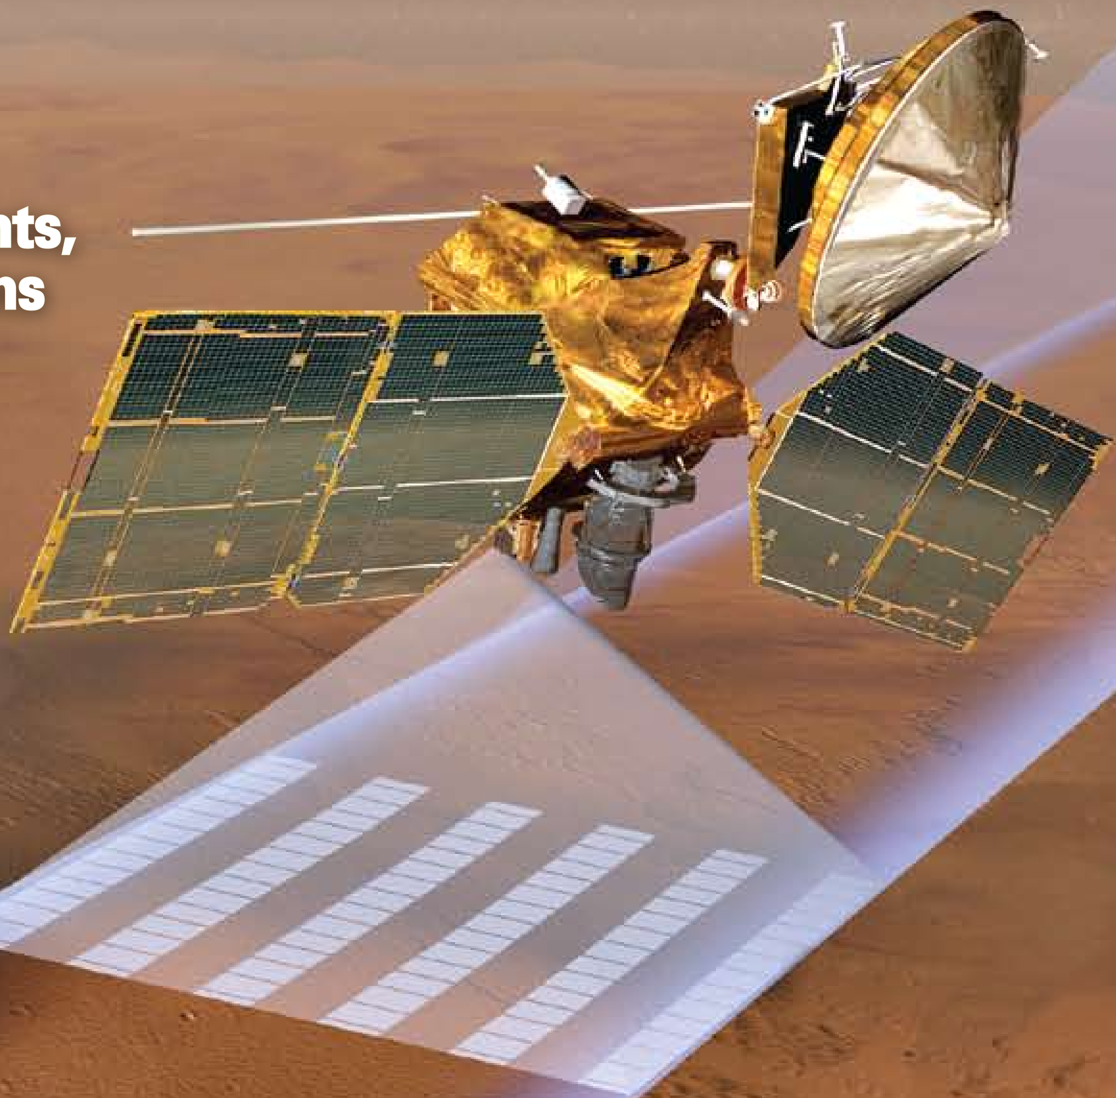

Military EMBEDDED SYSTEMS

www.mil-embedded.com

November/December Volume 3 Number 6

COLUMNS

Industry Analysis

Crosshairs Editorial

DEPARTMENTS

WEB RESOURCES

Subscribe to the magazine or E-letter:
www.opensystems-publishing.com/subscriptions

Live industry news:
www.mil-embedded.com/news
www.opensystems-publishing.com/news/submit

Submit new products:
www.opensystems-publishing.com/vendors/submissions/np

COVER
An artist's rendering of the Mars Reconnaissance Orbiter (MRO) using its Mars Climate Sounder instrument to profile temperature, pressure, dust opacity, and other planetary metrics. In our case study on page 42, COTS software by Express Logic and integrated into the MRO by Ball Aerospace controls the electronics of the High Resolution Imaging Science Experiment (HiRISE). Image courtesy of NASA/JPL-Caltech.

Hardware: ABCs of Cooling/Chassis

Special feature: COTS over Mars

Product Guide: Radio

E-LETTER

www.mil-embedded.com/eletter

FPGA coprocessing architectures for military applications
By J. Ryan Kenny and Bryce Mackin, Altera Corporation

FPGA/DSP hybrid architectures: Satisfying the reconfigurability requirements of the military
By Ron Huizen, BittWare, Inc.

New design methodologies offer easier path to custom DSP hardware
By Shawn McCloud, Mentor Graphics

Implementation trade-offs of digital FIR filters
By Ed Rocha, Quickfilter Technologies

Published by: OpenSystems Publishing™

Answer: Both.

Which path would you choose? Would you stick to the safer-looking, well-traveled route? Or would you choose the more rugged path — the one that may very well take you on an adventure?

What if you could do both? With new high-performance, low-power Intel embedded technology you get all of the familiar goodness of the most widely used processing architecture in the world. And now we've added advanced platform technologies that will help you blaze new trails and explore new ground. Like hardware-enabled virtualization and active management technology that allows for remote control of embedded platforms. All without adding more chips or increasing the complexity of your board designs.

With so much efficiency built right in, it's almost like going down two roads at the same time. And that's an adventure on a whole new level. Are you ready for it?

To learn more go to: http://developer.intel.com/design/info/899.htm

Intel Embedded Technology. Igniting Innovation.

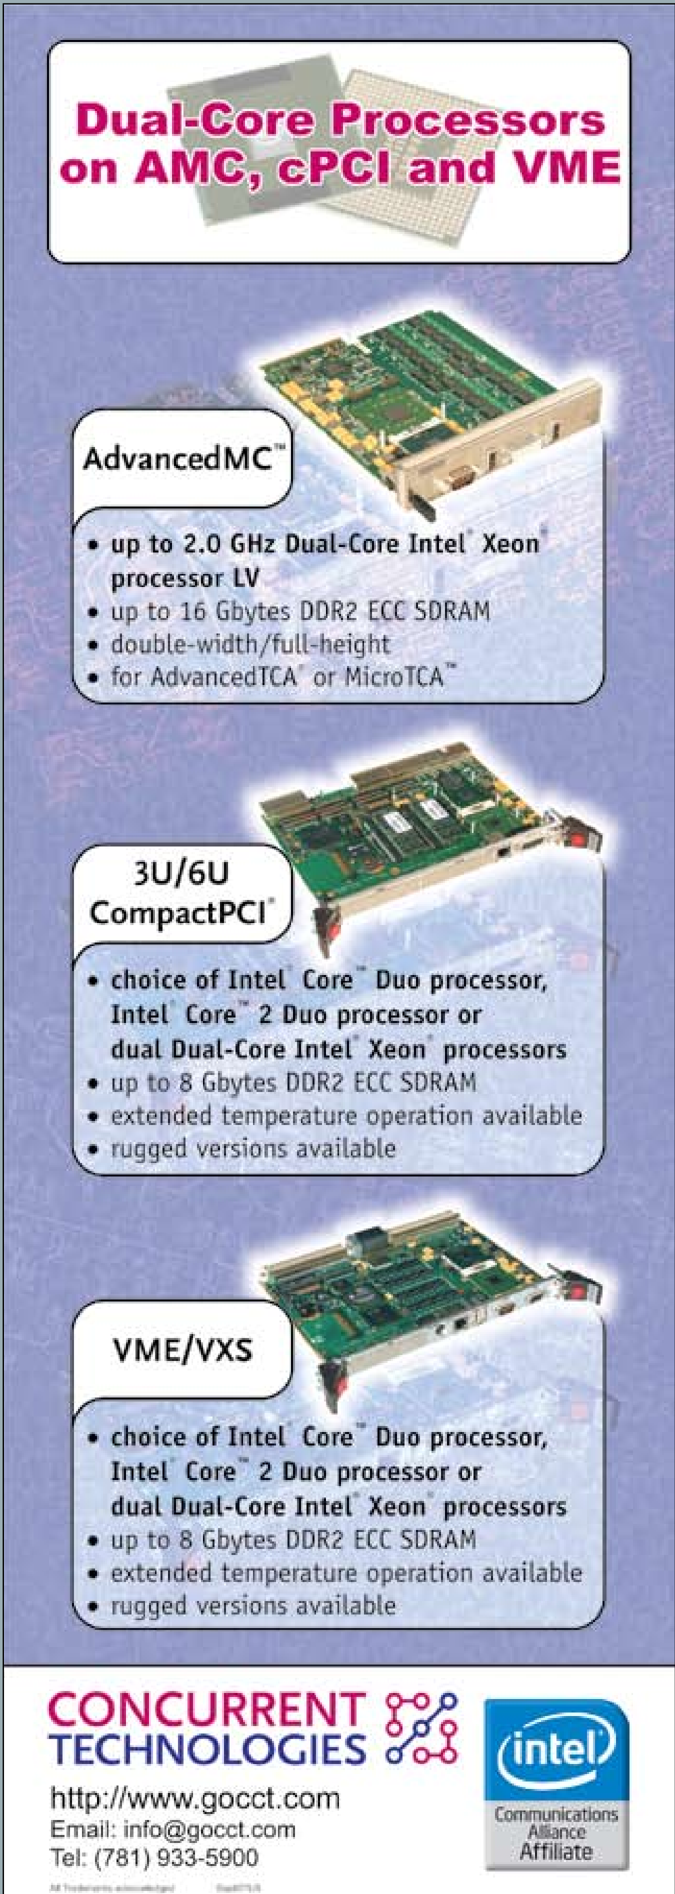

Military EMBEDDED SYSTEMS

AN OPENSYSTEMS PUBLICATION

Military and Aerospace Group

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

Group Editorial Director Chris A. Ciufo
cciufo@opensystems-publishing.com
Senior Editor (columns) Terri Thorson
tthorson@opensystems-publishing.com
Assistant Editor Sharon Schnakenburg
sschnakenburg@opensystems-publishing.com
Copy Editor Robin DiPerna
European Representative Hermann Strass
hstrass@opensystems-publishing.com
Creative Director Steph Sweet
Senior Web Developer Konrad Witte
Art Director David Diomede
Graphic Coordinator Sandy Dionisio
Circulation/Office Manager Phyllis Thompson
subscriptions@opensystems-publishing.com

OpenSystems Publishing

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

Publishers John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial Rosemary Kristoff

Communications Group

Editorial Director Joe Pavlat
Assistant Managing Editor Anne Fisher
Senior Editor (columns) Terri Thorson
Technology Editor Curt Schwaderer
European Representative Hermann Strass

Embedded and Test & Analysis Group

Editorial Director Jerry Gipper
Editorial Director Don Dingee
Associate Editor Jennifer Hesse
Special Projects Editor Bob Stasonis
European Representative Hermann Strass

ISSN: Print 1557-3222

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

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

Canada: Publication agreement number 40048627

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

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

The U.S. Navy sails with AdvancedTCA

By Joe Pavlat

With a workforce of almost 650,000, including 344,000 in active duty, the U.S. Navy maintains, trains, and equips combat-ready Naval forces capable of winning wars, deterring aggression, and maintaining freedom of the seas. Among the Navy's assets are nearly 300 deployable battle force ships and 4,000 aircraft. The U.S. Navy Program Executive Office, Ships (PEO Ships) manages acquisition and total life-cycle support for all U.S. Navy non-nuclear surface ships. This includes onboard computing systems of all flavors.

We tend to think of military electronics applications as requiring extremely rugged components, and that is often the case, but there are many places where only moderately rugged platforms are used. Large Naval vessels, including aircraft carriers, destroyers, and support ships, currently have onboard data centers not unlike those in any modern business. Interestingly, traditional 1U server racks are starting to be replaced with AdvancedTCA, and for interesting reasons.

Power, cooling, and efficiency issues are as important to these floating data centers as they are on land. Everyone is struggling to get more performance in a smaller space with existing power availability. Processor performance has outpaced memory performance for some time, and the need to get more processing power without increased power and heat has moved the industry to multithreaded, multicore processors. Sun Microsystems' UltraSPARC T2 is a fine example of this technology, containing eight cores supporting eight threads per core. This level of performance, combined with the robustness of the Solaris operating system, makes AdvancedTCA ideal for many Navy AdvancedTCA initiatives. AdvancedTCA provides the U.S. Navy several important advantages, among them better power efficiency and better volumetric efficiency than traditional 1U servers.

AdvancedTCA provides improved power efficiency because it is a DC-based system, not an AC-based one. Traditional data centers, military or otherwise, bring in AC power, usually at 480 VAC. This must be converted to DC in an uninterruptible power supply to charge batteries, then be converted back to 120 VAC for delivery to the rack. That AC is then converted back to a fairly high-voltage DC in the server's power supply, which subsequently converts it down to the required logic voltages for the semiconductors, fans, and disk drives. This all means that a lot of power, about half, is lost and plenty of heat is produced before any computing is completed. AdvancedTCA's redundant 48 VDC power structure reduces the amount of conversions that occur, and energy and thermal efficiency improvements of 20 percent or more are possible. Reliability improves too. Verizon is currently building DC-only data centers for this exact reason. Cooling is better with AdvancedTCA as well, as hot air likes to go up, not sideways. Better cable management also improves airflow.

AdvancedTCA provides better volumetric efficiency than traditional rack-mount servers. Recently, I spoke with Dave Berry, an industry veteran and senior product manager for Sun Microsystems' Netra Systems. He said Sun Microsystems has developed a useful metric that they call *SWaP* (Space, Watts, and Performance). SWaP is calculated by dividing performance (in whichever benchmark units you like) by the product of space used times power consumed. Without going through the details, SWaP demonstrates that AdvancedTCA servers can provide a SWaP metric almost three times greater than a bunch of 1U servers. While AdvancedTCA performance is a function of the CPU used and is no different from traditional servers, the power consumed and the space required are dramatically improved. The Sun Microsystems CT900 AdvancedTCA system is shown in Figure 1.

Figure 1

Dave says that there are other aspects of AdvancedTCA that the U.S. Navy likes. One is the common management interface. Another is the easy integration of Service Availability Forum compliant, high-availability middleware. They also like the NEBS standards for shock, vibration, and EMI that are widely used in the telecom environment. Then add the IP foundation of AdvancedTCA, a 10 GbE data fabric, and additional fabric headroom, and AdvancedTCA is a winner for the critical work of the U.S. Navy.

To learn more, e-mail Joe at jpavlat@opensystems-publishing.com.

All-around awareness

By Duncan Young

Traditional armored vehicles normally require both the driver and commander to be *head-out* in order to maneuver their vehicle, make tactical decisions, deliver troops, or engage targets. However, as crew survivability is a paramount consideration, the next generation of vehicles will be able to perform their entire mission with all the hatches locked down tight. Although the vehicle is likely to be connected to the "digital battlefield" – receiving and maintaining a regular tactical view of the battlespace showing maps, asset dispositions, threats, and targets – there is still a need for the crew to have a high-fidelity view of the area immediately surrounding their vehicle, as though they were head-out. This local awareness capability complements the sophisticated sensors, target acquisition, fire control, and self-defense systems that are fitted to, for example, main battle tanks, extending across the fleet to lightly armed vehicles and troop carriers.

A typical local awareness system will use from 8 to 16 cameras mounted in armored glass portholes located around the periphery of the vehicle, giving full 360-degree coverage of the immediate environment. Alternatively, the cameras could be mounted on a mast, which is harder to protect although much lower in initial cost. The concept of full coverage by 16 cameras is depicted in Figure 1. The cameras would normally operate in the visible wavelengths but could also operate in near or far infrared, if required by the role and mission of the vehicle. Each camera could simply be connected to a monitor for display, but the external environment would then be difficult to visualize from the resulting 16 monitors, whatever their physical arrangement within the vehicle.

Ideally the images from the cameras should be displayed together on a single screen to give the visual impression that the

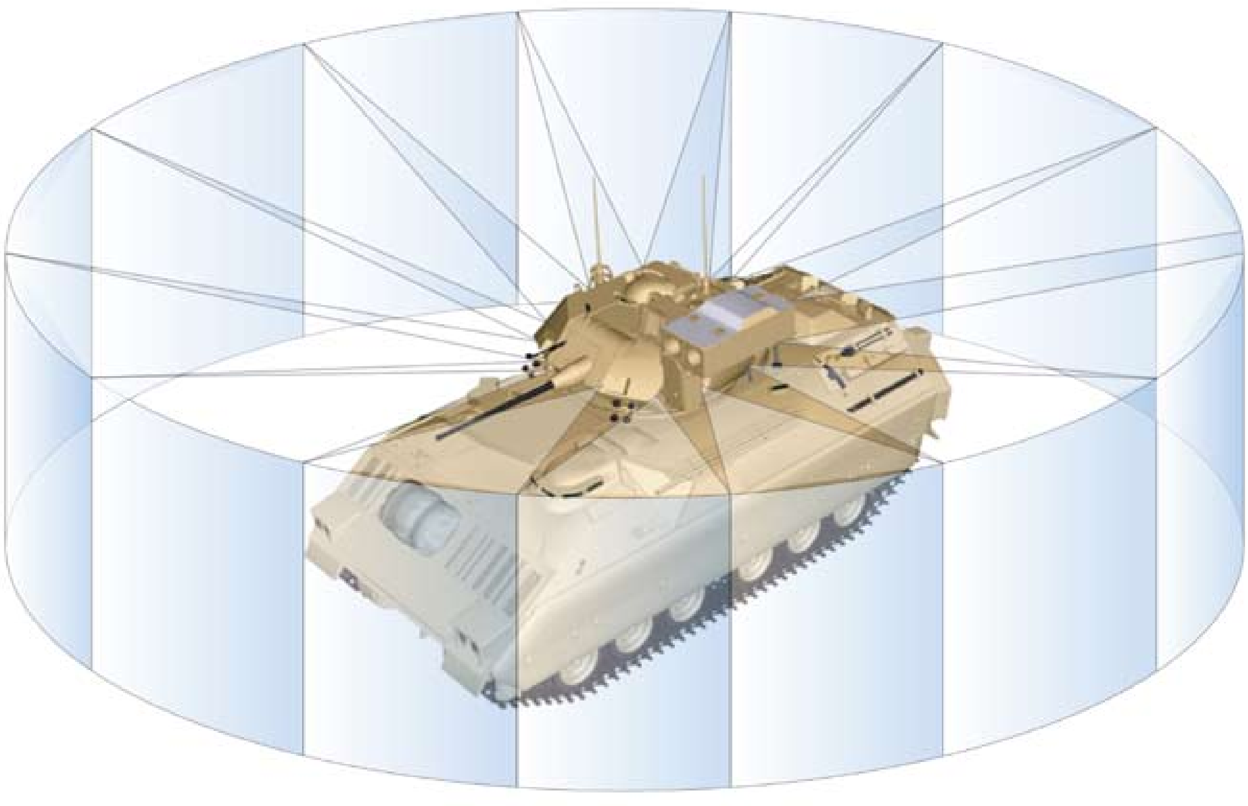

Figure 1

surrounding scene is being viewed head-out. Trying to recreate this perfectly is challenging because although the normal field of view is quite narrow, the eyes and head can move rapidly to pan a very broad area. The commander could be surrounded by a 360-degree bank of screens, but this wouldn't be a very practical solution to the problem. More practical would be to use the power of today's advanced video and graphics processors to seamlessly stitch together the incoming video streams and display segments of the resulting image through a virtual camera that can be panned and zoomed by the commander to areas of interest as required.

Newer types of military ground vehicles already have rudimentary local awareness systems allowing the commander and crew some measure of confidence in their outside environment while on the move or during an engagement. The next generation – being specified now for new vehicle fit into real, funded programs – will be much more sophisticated, offering rapid all-around visibility, visualization options, image processing capability, training, and recording of a complete mission's local environment.

To learn more, e-mail Duncan at young.duncan1@btinternet.com.

In addition to creating a usable visual display, the next logical step would be to perform further image processing on the composite panoramic video: motion detection, feature detection, image capture and enhancement, image stabilization, target classification, tracking, and so on. However, some of these applications are very complex, particularly while the vehicle is on the move, making it likely that for the most economical deployment, a system's image processing capabilities will be closely matched to its initially intended battlefield role. The use of scalable and upgradeable, off-the-shelf processing solutions will allow for these future enhancements as sensor performance increases and vehicles' roles evolve over extended periods of time.

An example of a typical local awareness/surveillance system is the ADEPT 5000 from GE Fanuc Embedded Systems. Supporting up to 16 TV-standard cameras, it has extensive and tailorable visualization capabilities. Based on industrial-packaged PC technology, the system is ideal for security monitoring of fixed installations such as airfields and Naval bases. In this case, the cameras are mounted to give a panoramic view of the installation, and image processing detects any changes or motion indicating potential threats or intrusions. It can also be configured around the latest generation of military specification, embeddable shoebox PCs such as GE Fanuc's MAGIC-1 system, with its modular architecture giving flexibility in configuration and scalability to suit a wide variety of rugged, deployable local awareness applications.

Daily Briefing: News Snippets

By Sharon Schnakenburg, Assistant Editor

www.mil-embedded.com/dailybriefing

BAE and DARPA "SWaP" in Tadiran batteries

It's "out with the old, in with the new" for DARPA's Optically Directed Attack Munitions (ODAM) program, which trades in its 60 mm mortars for a new guided munitions system featuring CR2-sized batteries. Accordingly, prime contractor BAE Systems selected Tadiran Batteries' TLM-1530HP lithium batteries for ODAM because they meet program SWaP requirements, beating out the limited-shelf-life 3.0 V CR2 cells typically found in consumer cameras. Tadiran's TLM-1530 has a 20-year shelf life, handles up to 6.5 A pulses, features a maximum continuous load of 2.5 A, and operates from -40 ˚C to +85 ˚C.

Now … *how* do we do this again?

Apparently the Joint Program Executive Office (JPEO) recognized a need for "JTRS" (Joint Tactical Radio System) to stand for something closer to "Just The Right Size." Accordingly, PrismTech and other SDR industry experts recently converged to help define guidelines for the DoD's waveform portability program, giving the industry a clarified procedure and strategy. The goal is to increase JTRS waveform portability and enhance the program overall, and the revised guidelines are scheduled for release in Q4 07.

GE Fanuc goes ballistic

Our "men in blue" around the globe are about to get more support from GE Fanuc Embedded Systems … or at least their engineering staffs will, according to GE's recent announcement of its newly created Naval Systems Engineering Group. In addition to the obvious, the group also provides LEAP ballistic prediction and systems interfacing, along with ballistics assessments using tools such as the Dynamic Gun Alignment and Analysis System (DGAAS) and the VTMS and LCASS NGS systems.

SDR Forum gets a fresh start

There's nothing like a new beginning … especially for 20-year industry veteran Lee Pucker, who was recently named the SDRF's "consulting CEO." In his two-year contract position, Pucker helps drive strategic plan development, guides Forum officers in their new strategic direction, and oversees Forum financial matters. Pucker's most recent position was as CTO and VP for corporate development at SDRF member Spectrum Signal Processing, and he is presently forming his own ForwardLink Consulting firm. Previous engineering experience includes positions at Computer Sciences/ Nichols Research and the U.S. DoD, among others.

Then two became one

Some say two (companies) are better than one, but now VMETRO and Micro Memory are refuting that notion with the former's recent acquisition of the latter. The move is touted to blend VMETRO's presence worldwide – and its two design centers in Europe – with Micro Memory's U.S.-based design center and customers. The result: VMETRO gains more opportunities to participate in U.S. ITAR-regulated programs. The acquisition also serves as VMETRO's gateway into a new market sector through Micro Memory's SSD Umem NVRAM cards, primarily used by Tier I/II OEMS in the enterprise network storage segment.

BiTMICRO: Show me the money

While SSD's internal architecture won't be turning round and round anytime soon, the wheel of fortune is … at least for SSD manufacturer BiTMICRO. The company recently completed Series F funding of $9.3 million via Woodside Investments LP's investor team. The money is slated to "ramp production, sales, and marketing operations" and to support R&D initiatives to design next-gen, flash-based E-Disk SSDs (pictured) and semiconductor storage technology. BiTMICRO's total private funding now totals $31.4 million.

Making a "CASE" for safety-critical software

Software tools vendor AdaCore isn't saying "c'est la vie" to the military's challenge of finding an open-source development environment to produce safety-critical embedded software with a 10- to 30-year lifespan. Instead, the company has just signed on to the Toolkit in Open Source for Critical Applications & System Development (TOPCASED) project, an open-source consortium comprising scientific, academic, and commercial partners hoping to remedy the problem. AdaCore's primary effort is to provide an Eclipsed-based development environment to TOPCASED, which is funded in part by the French General Business Directorate.


TOOLS OF INNOVATION
The BU-65590U — Small, lightweight, and portable.
• MIL-STD-1553 and ARINC 429
• Graphical User Interface
• High Channel Count
• USB Powered
DDC®
©Data Device Corporation
www.ddc-web.com
Toll Free: 1-800-DDC-5757

Heat pipe superconductors provide high-speed cooling

By Bahman Tavassoli, PhD

Heat pipes are an integral part of thermal management systems. Here's what you need to know.

High-performance, highly condensed electronics are running hotter than ever – and are susceptible to performance and lifespan shortcomings if temperatures exceed their safety range. Heat sinks have long been used as basic thermal management devices. Now, when combined with heat pipes they provide faster cooling, which can make all the difference if temperatures rise quickly.

Heat pipes are transport mechanisms that can carry heat fluxes ranging from 10 W/cm^2 to 20 KW/cm^2 at nearly the speed of sound (340 m/sec). Basically, they can be considered heat superconductors. Heat pipes can be used as a means to transport heat from one location to another. They can also be used as a means to isothermalize temperature distribution (equalize heat transfer across the spreading surface).

Vaporize, condense, repeat

A heat pipe has three sections: the evaporator, the adiabatic, and the condenser[1], as shown in Figure 1. The interior of the pipe is covered with a wick, and the pipe is partially filled with a liquid such as water. When the evaporator section (L_e) is exposed to a heat source, the liquid inside vaporizes and the pressure in that section increases. A level of control over the total pressure in the heat pipe can be obtained by controlling the amount of working fluid. Water, for instance, expands 1,600 times when it vaporizes at 1 atmosphere. If 1/1,600 of the volume of a heat pipe is filled with water, when all the fluid is just vaporized, the pressure will be 1 atmosphere.

The increased pressure causes the vapor to flow at near-sonic speed toward the condenser section of the heat pipe (L_c). The vapor in the condenser section loses heat to an integral heat sink and is converted back to liquid by the transfer of the latent heat of vaporization to the condenser. The liquid is then pumped back to the evaporator through the wick capillary action. Throughout this activity, the middle section of the heat pipe (L_a), the adiabatic portion, has a temperature variance of just a few degrees.

For the capillary force to drive the vapor, the wick's capillary pressure should exceed the pressure difference between the vapor and the liquid at the evaporator. If the heat pipe is operated against the force of gravity, the liquid undergoes a larger pressure drop. Depending on the heat load, this can be from 30 to 100 Pa. The result is less wick pumping with reduced heat transfer. The amount of heat transfer decrease depends on the particular heat pipe.

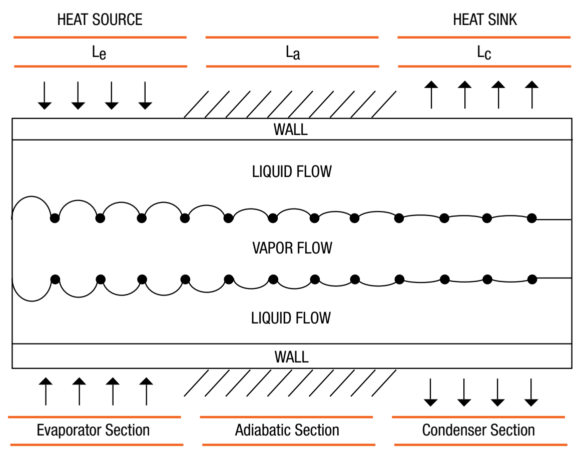

a) Liquid-vapor interface

Figure 1

Working Fluid	Operating Temperature, (ºC)	Wick Design	Wall Material	Axial Heat Transport, (W)
Methane	-140	Circumferential mesh	Stainless steel	12
Water	100	Axial grooves	Copper with rectangular cross-section	70
Sodium	430-790	Circumferential stainless steel screen	Stainless steel	1309

Table 1

Table 1 shows experimental data for the operating temperature and heat transfer for three types of heat pipes[1].

What's in a heat pipe?

A typical heat pipe is made of the following:

1. ***Metallic pipe*** **–** The metal can be aluminum, copper, or stainless steel. It must be compatible with the working fluid to prevent chemical reactions such as oxidation.
2. ***Working fluid*** **–** Several types of fluids have been used to date. These include methane, water, ammonia, and sodium. Choice of fluid also depends on the operating temperature range.
3. ***Wick*** **–** The wick structure comes in several different shapes and materials. Figure 2 shows the profiles of common wick types: axial groove, fine fiber, screen mesh, and sintering. Each wick has its own characteristics. For example, the axial groove has good conductivity, poor flow against gravity, and low thermal resistance. Conversely, a sintering wick has excellent flow in the opposite direction of gravity but has high thermal resistance.

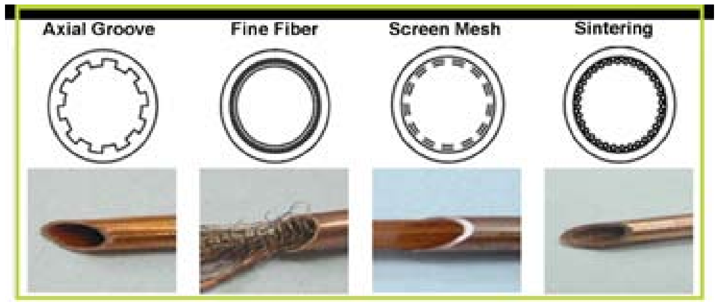

Figure 2

Certain factors can limit the maximum heat transfer rate from a heat pipe. These are classified as follows:

1. ***Capillary limit* –** Heat transfer is limited by the pumping action of the wick.
2. ***Sonic limit* –** When the vapor reaches the speed of sound, further increase in the heat transfer rate can only be achieved when the evaporator temperature increases.
3. ***Boiling limit* –** High heat fluxes can cause dry out.
4. ***Entrainment limit* –** High-speed vapor can impede the return of the liquid to the condenser.

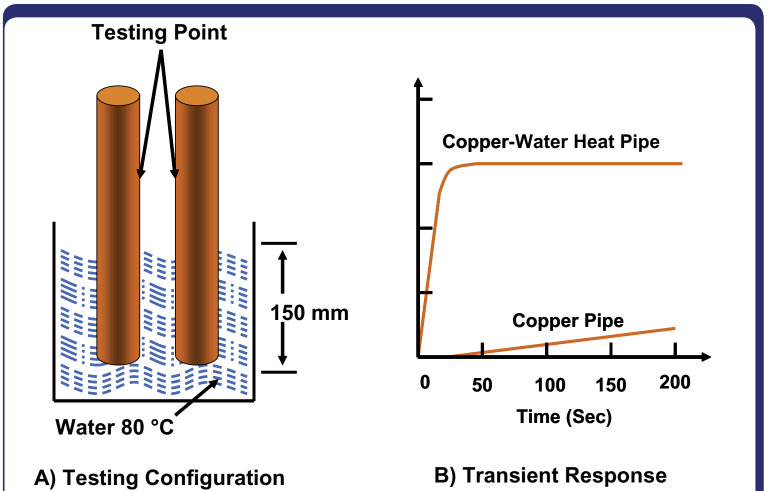

Figure 3

Piping hot conductivity

A quality heat pipe has an effective thermal conductivity much larger than that of a very good metal conductor such as copper. Figure 3 shows a copper-water heat pipe and a copper pipe dipped into an 80 °C water bath[1]. Both pipes were initially at 20 °C. The heat pipe temperature reaches the water temperature in about 25 seconds, while the copper rod reaches just 30 °C after 200 seconds. However, in an actual application when a heat pipe is soldered or affixed to the base of a heat sink with epoxy, the effective thermal conductivity of the heat pipe can be reduced fivefold or more because of the extra thermal resistance added by the bonding. A rule of thumb for effective heat pipe thermal conductivity is 4000 W/mK. Additionally, there are many heat pipe shapes in the market, but the most common are either round or flat.

Round (or tubular) heat pipes are used for transferring heat from one point to another. They can be applied in tightly spaced electronic components such as a laptop, where space for a heat sink is remote from the hot components. Heat is transferred to a different location that provides enough space to use a proper heat sink or other cooling solution.

Flat heat pipes, also called *vapor chambers* or *heat spreaders,* conceptually work the same as round heat pipes. When the heat source is much smaller than the heat sink base, flat heat pipe can be embedded in the heat sink base, or it can be attached to the base to spread the heat more uniformly on the base of the heat sink. Flat heat pipes are typically used for cooling surface-mount circuit boards and are ideal for use in military avionics, laptops, and other DoD electronics. For direct contact with a component, flat heat pipes must be used because they transport heat in two dimensions, as opposed to a single dimension with round heat pipes.

Although a vapor chamber (flat heat pipe) might be helpful in minimizing spreading resistance, it may not perform as well as a base plate made from a very high conductor such as diamond. A determining factor is the thickness of the plate. Figure 4 shows the spreading resistance for an 80 x 80 x 5 mm base plate of different materials with a 10 x 10 mm heat source[2].

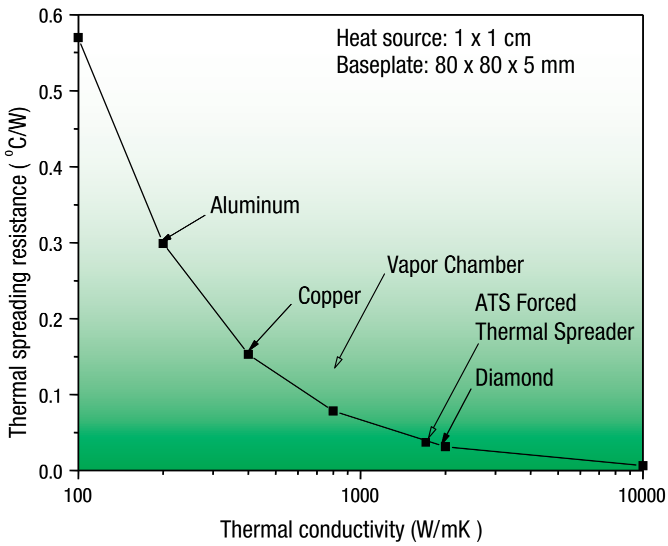

Figure 4

The vapor chamber has a spreading resistance of under 0.15 °C/W, which is better than copper (0.15 °C/W), but worse than diamond (0.05 °C/W). However, the price of the diamond heat spreader might not justify its application. Figure 4 also includes the spreading resistance from the ATS Forced Thermal Spreader (FTS), which is equal to that of diamond at a much lower cost. The FTS uses a combination of mini and micro channels to minimize the spreading resistance by circulating the liquid inside the spreader.

When cooling methods might not be so hot

Heat pipes have a very important role in the thermal management arena. With projected life spans of 129,000-260,000 hours (as claimed by their manufacturers), they will continue to be an integral part of some new thermal systems. However, with such problems as dry out, acceleration, leakage, vapor lock, reliability, and performance in ETSI- or NEBS-type environments, heat pipes should be tested prior to use and after other cooling methods have been explored and ruled as unsatisfactory.

References:

1. Faghri, A., Heat Pipe Science and Technology, Taylor & Francis, 1995.
2. Xiong, D., Azar, K., Tavossoli, B., Experimental Study on a Hybrid Liquid/Air Cooling System, Advanced Thermal Solutions, Inc., 2003.

Bahman Tavassoli, *PhD, is chief technologist at Advanced Thermal Solutions, Inc., responsible for the development of innovative electronics cooling solutions, provision of thermal design and consultation services, and ongoing research and development in electronics cooling. He is an internationally recognized expert on thermal management and a frequent speaker at industry conferences including IMAPS and ITHERM. He has gained extensive experience in computational fluid dynamics, authored several technical papers on electronics cooling and thermal engineering, and holds a patent in the area of electronics cooling. Bahman earned his undergraduate and master's degrees in Mechanical Engineering from Clarkson University and a PhD in Mechanical Engineering from the University of Connecticut. He can be reached at ats@qats.com.*

Advanced Thermal Solutions, Inc.
781-769-2800
www.qats.com

Military & Aerospace Applications:

- rocket & missile guidance
- smart ammunition
- torpedoes & mines
- sonobuoys
- ground sensors
- artillery fuses
- active decoy systems
- trajectory correction add-on kits
- bomb proximity fuses
- dispersed munitions sensors

Commercial Applications:

- medical devices
- back-up systems
- automotive
- marine/oceanographic

Powerful design ideas start with TLM batteries from Tadiran. Capable of delivering high cell voltage, high energy density, instant activation, and exceptionally long operational life even in extreme temperatures, TLM Series batteries utilize an innovative COTS technology to deliver an open circuit voltage of 4.0 volts with high current pulses of up to 15A and 5A continuous current at 3.2V. For small, space-saving TLM cells, contact the biggest name in lithium, Tadiran.

TADIRAN BATTERIES

Tadiran Batteries
2 Seaview Blvd.
Port Washington,
NY 11050
1-800-537-1368
516-621-4980

www.tadiranbat.com

Quick, easy air cooling for VME and CompactPCI systems

By Charles Linquist

VME and CompactPCI designs are still very popular for military and commercial applications due to the wide variety of boards available. These "standards-based" designs allow companies to focus on developing solutions without devoting scarce resources to hardware design. Even though the "pure electronics" has been handled for them, it is still up to the designers to specify or design the packaging and cooling system that best suits the end product's needs. And in most cases, the hand calculation method is more efficient and cost-effective than computer modeling.

Electronics are getting smaller and more powerful every day, and even though die shrinkage continues to reduce the watts per MIP, the demand for more performance is outstripping the ability to shrink device geometry. As a result, the power consumed by processors and chipsets is still increasing, and power produces heat. While commercial equipment producers have considerable flexibility in packaging their units, those forced to adhere to the VME and CompactPCI standards are severely limited in their choices of how to package and cool their systems. And, although conduction-cooled VME and CompactPCI chassis are fairly common in the military, air-cooled units still predominate because of the wide selection and low cost of the boards available that utilize this cooling method.

A VME or CompactPCI system designer, therefore, needs to quickly specify or design a chassis that is virtually guaranteed to cool today's hottest boards. And while computer modeling is a valuable tool, careful "massaging" of the input is necessary to insure an accurate result. Often, simple hand calculations can provide a more than satisfactory answer to the air-cooling equation.

Use LFM, not CFM

When calculating the cooling requirements for electronics, the formula most people refer to is the "air temperature rise formula" **CFM = (1.830*W)/dT** where *CFM* stands for airflow in Cubic Feet per Minute, W is in watts dissipated, and "dT" is the temperature rise in degrees Celsius. This formula gives the temperature rise of air going through any space when a given number of watts are being dissipated in that space.

This formula works perfectly for designing HVAC systems for buildings, but is not very useful when trying to cool small heat-producing components, since the focus of this formula is to maintain proper *air* temperature; however, in electronics cooling, we are concerned with maintaining proper *component* temperatures.

When your mission is critical...

...DEPEND on Birtcher Card Guides & Conduction Cooled Assemblies

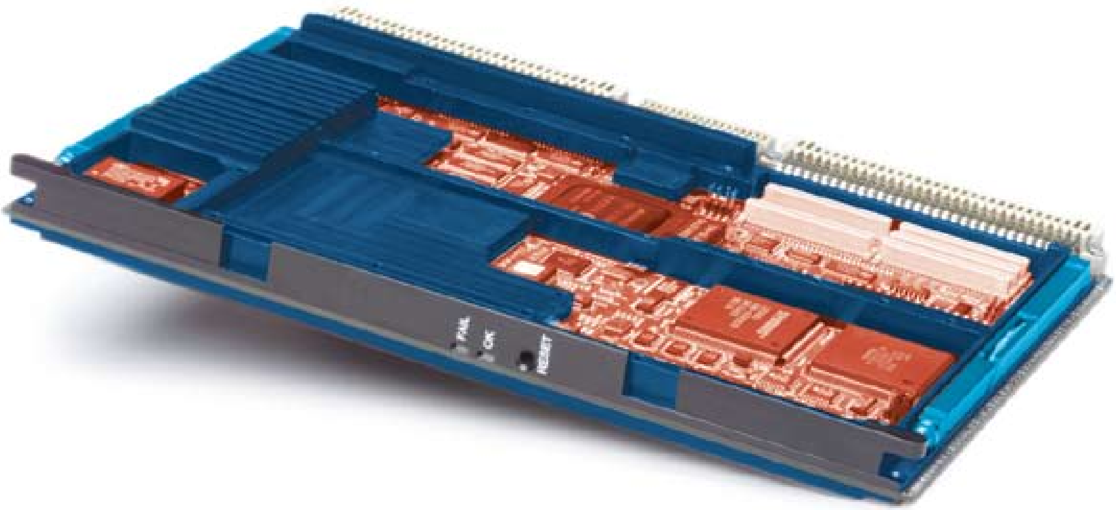

Wedge-Tainer™

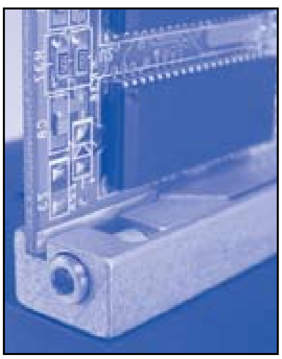

- Designed to comply with MIL-STD-810D and MIL-STD-901D in addition to MIL-E-5400

Wedge-Lok®

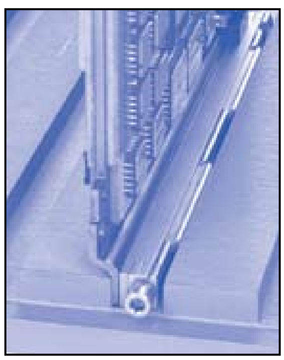

- Meets the requirements of drawing 84103, Defense Supply Center, Columbus

For 46 years, Birtcher's rugged card guides, Wedge-Loks®, and conduction cooled products have been the keystone of electronic systems. Birtcher products secure electronics with robust clamping and a protective cooling path regardless of vibration or G-loads. Whether deployed in spacecraft, aircraft, watercraft, or all-terrain vehicles, when your application is mission-critical, look no further than Birtcher.

Birtcher

HOLDING IT ALL TOGETHER

www.birtcherproducts.com • 888-957-8291

Every component on a board is a potential heat source and may generate substantial amounts of heat. If the board was designed properly, the designer will have added heat sinks to those components that have an unacceptable temperature rise. In any case, all the components must be kept below their design temperature limits if reliability is to be assured.

Whether or not a component has a heat sink, the same rules apply: As long as the air moving past a device is cooler than the device, heat will be transferred from it to the passing air. Just how much heat this airflow removes is dependent on the efficiency of the component/air interface (usually a heat sink), the temperature difference between the component and the passing air, and the velocity of that air.

Figure 1, courtesy of Aavid Thermalloy, LLC shows an extrusion. This graph is not presented to help designers pick a heat sink, rather it is provided to show which parameters are important in cooling design. Note that CFM is never mentioned on the graph. The reason is simple: Volume (CFM) does not cool individual electronic devices. Only air velocity over the component surface or through the heat sink as expressed in Linear Feet per Minute (LFM) is important. Note that the temperature rise of this heat sink (or any component connected to this heat sink) is proportional to the square root of the airflow velocity. This requirement for velocity is exactly why the fan on a desktop computer is generally mounted to blow air directly on the heat sink. Even a small fan can produce 600-800 LFM at very short distances. When looking at selection guides provided by manufacturers, note that while the actual amount of heat dissipated at a given air velocity varies according to the size of the heat sink, the shape of the temperature rise versus velocity curve remains the same for all of them.

So, what is the use of the CFM-based air temperature rise formula presented at the beginning of this section? It can tell designers how hot the air will get going *through* a chassis. If 1,000 W is dissipated, only 83 CFM is needed through the chassis to yield the commonly accepted value of 10 °C rise. This volume of air through a 21-slot wide CompactPCI or VME enclosure would produce only about 110 LFM, which is far below the 300-400 LFM required to cool most modern boards. Now it is easy to understand why the formula is inadequate for electronics cooling.

Real applications: Cooling VME or CompactPCI system boards

As a starting point, the system designer should look at the requirements given in the "power" or "environmental" sections of the datasheets for the boards the system will contain. If cooling requirements are mentioned at all, a LFM requirement will usually be given. Experience shows that most processor and DSP boards require anywhere from 250 to 400 LFM past their heat sinks to keep them operating within specifications. Of course, if the boards must run in elevated temperatures, such as those found in military environments, or are very high performance, the velocity may need to be even higher to prevent overheating.

The size and efficiency of the heat sinks on the heat-generating components have much to do with the required velocity. A small heat sink on a chip producing 65 W may require 600 LFM to cool it, while the same chip with a heat sink with four or five times the surface area will require only about 200 LFM to keep its temperature below the recommended levels.

AC versus DC fans

The chassis and cooling design process will almost always be iterative, since size determines the number and placement of fans. When faced with the problem of selecting the proper fan, most individuals will pick up one of the commonly available fan catalogs and look first for ones that fit. In that catalog, both AC and DC powered fans of the proper dimensions are usually available, but which type is best? In the past, AC fans had by far the greatest reliability because of the lack of brushes or commutation components. Now, however, brushless DC fans offer equivalent MTBFs and have several advantages over their AC counterparts, especially in smaller sizes.

DC units can be speed-controlled over a wide range simply by changing the applied voltage and can usually be obtained with tachometers, allowing for cooling performance monitoring, while AC fans with tachometers are quite rare. DC fans are often the only choice where very high flow rates are needed, since their motor speed isn't limited by the AC line frequency.

DC fans are also the obvious choice when "universal input" (110 V/220 V) is needed, since the power supply can deal with whatever voltage is available from the mains and convert it to a stable DC voltage for the fans. Of course, there is a downside in that the system power supply may need to be larger to power the fans and system electronics.

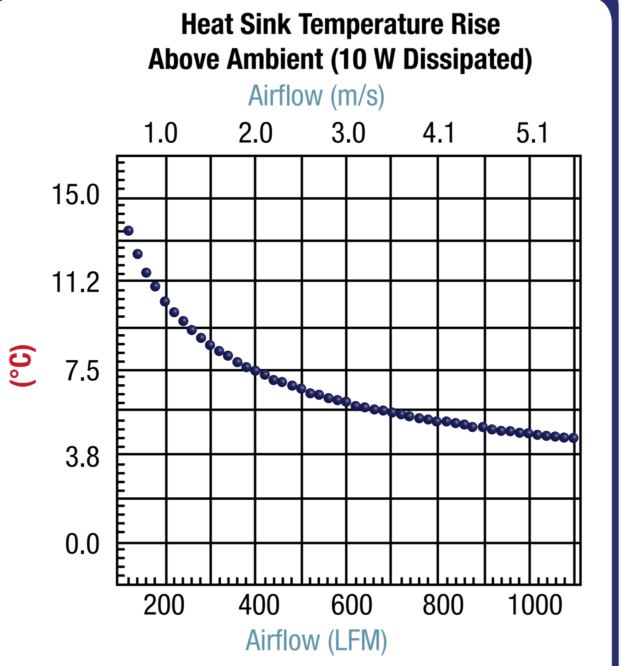

Figure 1

After a suitable size and operating voltage have been chosen, most people look in their fan catalog for high CFM numbers because it is generally thought that a fan with a higher CFM rating will move more air and therefore do a better job of cooling.

In reality, it turns out that the CFM rating is often a poor indicator of a fan's actual cooling performance. "Free air CFM" (the standard quoted figure) is only a *rough* indicator of a fan's ability to move air through a real chassis because a chassis will always have some restriction to airflow. This restriction causes *back pressure* (also called *static pressure*), and small amounts of static pressure can reduce the flow to far below the "free air" volume. To get a better picture of just what a fan can do in an actual application, designers need to look at the "fan curves" information provided by the manufacturer. They show how much air a fan can move when faced with different amounts of the inevitable back pressure.

Figure 2, courtesy of Mechatronics, shows representative fan curves of a 120 mm unit. Each letter (E, H, M, L, S) represents a fan from the same family, but each is designed to run at different Revolutions Per Minute (RPMs). The "E" fan is the fastest, while the "S" fan is the slowest. Note two items: The rated CFM of the fans is produced only at "0" static pressure, and the actual delivered CFM is heavily dependent on static pressure and fan RPM.

Since it is velocity that cools, these CFM figures must be translated into velocity. To accomplish this, the designer can simply divide the CFM through the box by the cross-sectional area in square feet. A 21-slot CompactPCI or VME enclosure has a cross-sectional area of about .767 square feet (17" x 6.5"/144). To get an average of 300 LFM through the card area, we need to move 300 * .767 = 230 CFM. If all slots are filled, it is likely that approximately half the width of each slot is occupied by the PCB and components. In that case, the area is halved, and the volume of air needed to achieve an average of 300 LFM is only 115 CFM.

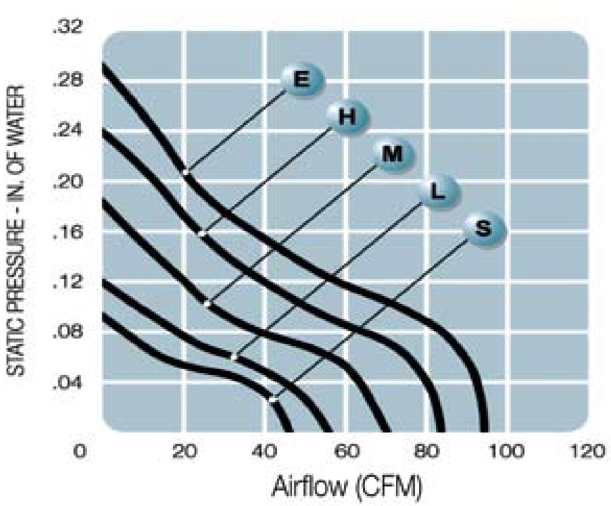

Figure 2

Since every enclosure will have some static pressure that the fan must work against, this must be considered in the calculations. The amount of static pressure the fan experiences will depend not only on the relative restrictions in the system but also on velocity. In fact, static pressure builds in proportion to the square of the velocity through an enclosure. A good first approximation to use for static pressure in VME/CompactPCI enclosures (assuming 250-300 LFM velocity) is .15" of water. According to the fan curve shown in Figure 2, the volume produced at 0.15" by the fastest fan is only about half of the fan's "free air" value. Also note that the air delivered by the slower fans is much less when presented with this static pressure. This also illustrates how difficult it is to achieve velocities above 400 LFM, since to move air twice as fast, the fan has to work against four times the static pressure.

In the CompactPCI or VME enclosure mentioned earlier, it would take three of the fastest fans (the "E" curve) to achieve 400 LFM through the card cage area. Often if high velocities are needed, a push-pull arrangement is best. This arrangement utilizes a fan on each side of the card area, one pushing air into it, and another evacuating air from the other side. By effectively putting the fans in a series, each one sees only half the static pressure.

Computer modeling: Is it really better?

Recently, computer modeling of cooling systems has become popular. Modeling can be very useful, especially when the exact configuration of the system is known and when every component is being pushed to the limit and there is no margin for error. It should be noted, however, that these programs often require careful "massaging" of input data to produce meaningful results. As an example, most software doesn't realize that small gaps and screw holes do not leak significant amounts of air; therefore, they can take hours calculating these insignificant losses. Preparing a model for analysis usually involves removing features that have no bearing on the final outcome. Also, in many cases, reworking the software models, running the analysis, and analyzing the results takes more time than actually building a test unit.

LFM formula usually wins out

In the majority of cases, the "back of the napkin" techniques outlined earlier can yield more than acceptable results in a short time and at very low total cost. This quick analysis is also useful to alert the designer when cooling margins are very slim, at which time it may be useful to perform a more extensive computer analysis.

Charles Linquist *has been director of engineering at Dawn VME Products in Fremont, California for the past five years. He also has experience in the data storage, telecommunications, commercial product, and military electronics arenas. Charles holds a BSEE from Iowa State University. He can be contacted at clinquist@dawnvme.com*

Dawn VME Products
510-657-4444, Ext. 108
www.dawnvme.com

Liquid cooling facilitates tomorrow's embedded systems

By Rex Harvey and Andy Odar

Modern embedded systems are extremely dense electronic assemblies that produce self-damaging amounts of waste heat. Traditional air, convection, and conduction cooling methods cannot accommodate these high heat densities. In contrast, liquid cooling not only has the capability of cooling these heat densities, but also has considerable margin for several spiral and/or major upgrades without replacing the system cooling infrastructure.

The power density and resulting waste heat of military embedded systems is growing exponentially with no end in sight. Even improvements in electronic component efficiencies and software optimization are quickly lost in the insatiable desire for more compute density. This means that current high-performance, high-connectivity embedded computing applications such as the emerging VPX (VITA 46) and VPX-REDI (VITA 48) board standards can generate more than 850 W of heat per board slot. State-of-the-art electronics such as these far surpass the practical cooling capability of the traditional coolant: forced air.

The reality is that air is a much better insulator than it is a heat transfer medium. (Table 1 shows a comparison of the heat transfer properties of various common fluids.) For example, a liter of the common dielectric fluid Polyalphaolefin (PAO) oil has more than 1,400 times the heat capacity of a liter of air and conducts heat more than five times more efficiently than air, which makes heat transfer into and out of PAO much easier. And water as a coolant more than doubles these figures. Liquids allow the size of the passages and coolant volumes to be very small, avoiding the large space claims of air-cooled systems that require high volumetric flow rates and large ducts and vents. The superior performance of liquid cooling allows substantial electronic upgrades with minimal cooling hardware modifications or changes in cooling system space and weight claim.

For a given cooling capacity, the prime movers, liquid system pumps, can be much smaller and require less power than air fans. Moreover, liquid is more easily routed to areas in the system that require a high level of heat removal. It is then routed to areas where the heat can be rejected remotely through liquid-air heat exchangers such as vehicle radiators, or liquid-liquid heat exchangers such as another existing liquid system or to sea water. These can be far from the heat source, if necessary, as only small liquid tubing is required. The cooled liquid then returns to the system in a very clean and efficient closed loop.

These concepts are demonstrated in the case study we present, using a chassis and board modules built to the draft VPX standard. This thermal management implementation shows how Liquid Flow-Through (LFT) cooling appears to be the intelligent approach to employ for current and future military embedded systems such as the one depicted in Figure 1.

	Thermal Properties			
Coolant	Dielectric	Density (kg/m^3)	Specific Heat (J/kg-K)	Conductivity (W/mK)
Air	Yes	1.205	1006	0.026
Fluorocarbon (FC-84)	Yes	1730	1100	0.061
Synthetic Oil (PAO)	Yes	848.8	2048	0.138
Antifreeze (PGW 50-50)	No	1041	3551	0.370
Water	No	998.2	4180	0.610

Table 1

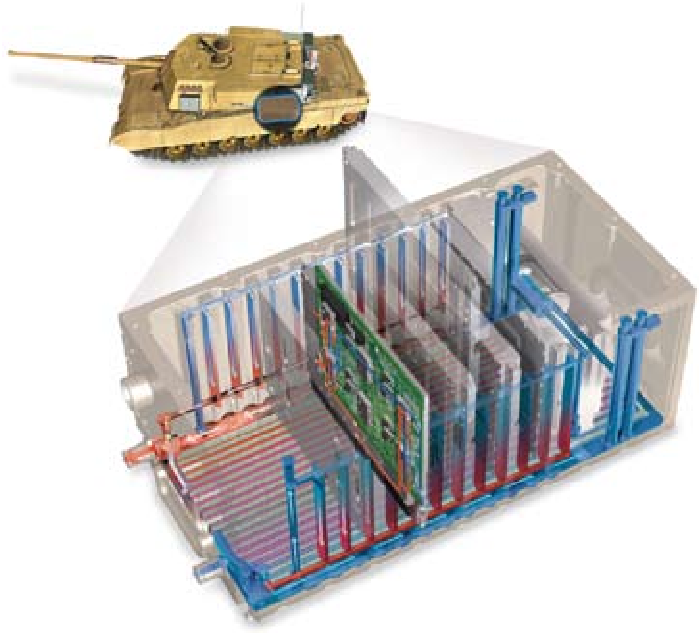

Figure 1

Cooling methods: How do they compare?

Consider a typical military embedded environment with ambient air temperature as high as 55 °C and maximum allowable device junction temperatures of 105 °C. Various cooling methods can be applied to these conditions, including natural convection cooling, forced air cooling, and conduction cooling. Figure 2 shows the comparative performance of these methods.

Natural convection cooling

Traditional natural convection cooling has a maximum cooling capacity of less than 0.05 W/cm^2 of board surface area. This means that a natural convection-cooled board of 6U x 160 mm VME size (surface area of 370 cm^2) can tolerate less than 20 W of total component waste heat.

Forced air cooling

Forced air can increase cooling to about 0.25 W/cm^2 of board surface area or just over 90 W for a typical 6U VME-sized board. Achieving this level of cooling, however, comes at the cost of providing adequate heat spreaders and heat sinks to effectively use the available airflow. Additionally, forced-air systems require fans and associated power and control functions and provisions for air duct volume and venting. Forced-air cooling is also highly dependent on altitude, often requiring as much as a tenfold increase in airflow rate (that is, going from about 10,000 feet to 30,000 feet) to compensate for the decreased air density. Increases in airflow rate can also lead to noise issues, which are detrimental in some applications.

Conduction cooling

Conduction cooling can provide substantial improvements in cooling performance to nearly 0.35 W/cm^2 of board area. It usually relies on metal frames to conduct waste heat from the high thermal output board components to the enclosure side walls. The enclosure side walls can be finned to have a relatively large thermal area, which can dissipate heat loads as high as 125 W per 6U board module. The module clamping system, typically an expanding-wedge retainer pressing the module against the slot's heat transfer area is the limiting factor for heat transfer.

Conduction cooling with liquid-cooled sidewalls

Cooling can be improved by using conduction modules with liquid-cooled side walls. This boosts cooling capacity to nearly 0.50 W/cm^2 of board area, allowing heat loads in excess of 180 W for a 6U VME-sized board. This approach, however, is still limited by the relatively high thermal resistance of the module clamping method.

Liquid cooling systems to the rescue

Meeting the thermal challenges presented by high-performance military embedded systems demands a new approach that removes the thermal roadblock at the junction between the board module and the enclosure walls. This is accomplished through the use of Liquid Flow-Through (LFT) cooling.

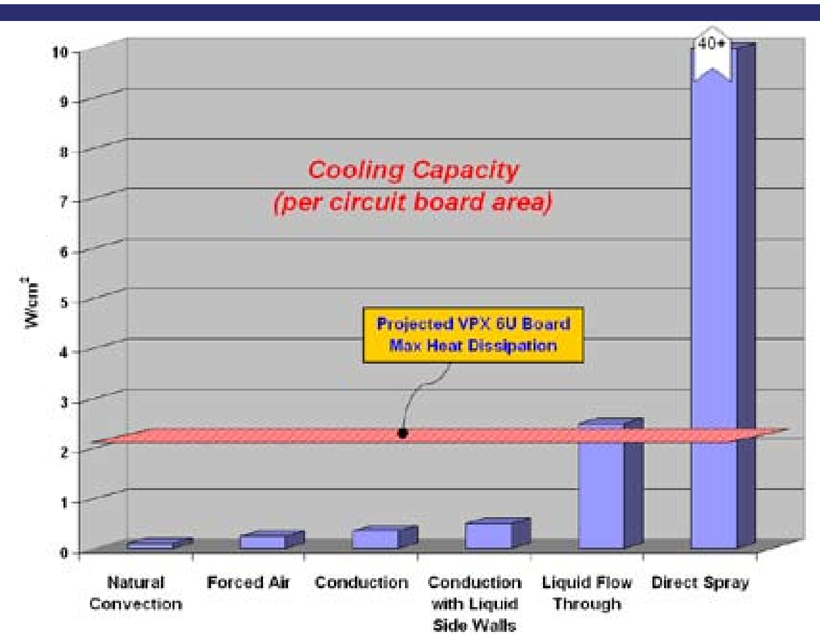

Figure 2

Liquid flow-through cooling

LFT coldplates eliminate the thermal discontinuity at the interface between board module and enclosure wall, realizing a minimum of a fivefold improvement in heat load capacity compared to conduction-cooling methods. The LFT cooling method essentially replaces the conduction frame of a board module with a liquid-cooling frame containing fluid passages and integral blind-mate fluid connectors for interfacing with the liquid supply and return manifolds in the enclosure. The liquid connectors mate at the same time as the electrical connectors when the modules are plugged into the backplane. Cooling capacities in excess of 2.50 W/cm^2 of board area are attained, making 6U VME boards with heat loads approaching 1,000 W a realistic design option. VME 6U-sized cards are currently below 200 W, so 1,000 W provides ample "thermal headroom" for future implementation of electronic upgrades without the need for comprehensive changes in the cooling system.

LFT cooling is a "dry" technique: The liquid is contained within the coldplate, never exposed to the environment nor wetting the devices being cooled or the enclosure. The cooling mode is single-phase; no liquid is vaporized, relying solely on the coolant's higher specific heat and conduction rate compared to other fluids (air) for its superior cooling ability. When optimally implemented, the LFT technique routes coolant at the lowest temperature directly to the hot devices on the board utilizing discrete, parallel paths eliminating the additive heating of coolant routed in a serial fashion.

Liquid two-phase cooling

Two-phase liquid cooling systems, such as spray evaporative cooling, can deliver even higher levels of cooling than other methods by taking advantage of the high heat of vaporization of liquid coolants compared to the sensible heat capacity of the liquid. Heat fluxes in excess of 40 W/cm^2 have been demonstrated (Figure 2). These systems, however, are extremely challenging to implement in a rugged environment because of sensitivities to heat flux levels, physical orientation, *g*-force, shock, vibration, as well as other factors. One application that has been successfully implemented is the spray evaporative cooling of the Cray model X1 Supercomputer, which is in a stationary and well-controlled environment. Parker spray caps that cover the Multi-Chip-Modules (MCMs) each contain 20 nozzles that spray dielectric fluorocarbon coolant directly onto the hot chips. This closed loop liquid cooling system is designed for 10-year operation without maintenance and, to date, has never caused a system failure.

Case study: Liquid flow-through cooling in today's military systems

The Advanced Cooling Systems team of Parker Hannifin Corporation has demonstrated the application of liquid cooling to modern military embedded systems through the hardware implementation shown in Figure 3. It utilizes 6U VPX-REDI board modules housed in an eleven 1.0-inch-pitch slot Air Transport Rack (1-ATR Long) enclosure measuring approximately 20 cm (8") (H) x 25 cm (10") (W) x 50 cm (20") (L). The modules and enclosure chassis support Two-Level Maintenance methodologies and are ruggedized and specifically designed to operate in harsh environmental conditions, for example, the vibration, shock, and acceleration requirements of MIL-STD-810F.

The liquid-cooled demonstrator is a self-contained, flexible system incorporating a 5 cm thick, plug-and-play smart pump and controller and PCB modules as shown in the left photo of Figure 3. Fluid inlets and outlets are optionally on the front or rear of the chassis. There is also an optional liquid-to-liquid heat exchanger at the bottom of the enclosure. The configuration is modular to provide versatility and adaptability as particular applications evolve and requirements become more demanding. The PCB modules interface with a dual-use chassis slot that allows use of conduction-cooled or LFT-cooled boards in the same slot. This facilitates the upgrade of individual slot locations as the need for additional cooling capacity arises. Slots configured for LFT cooling utilize dripless, blind-mate, quick-disconnect fluid couplings that also serve double duty as alignment features to guide module insertion into the backplane connectors.

Cooling performance testing has shown that device junction temperatures can be maintained below the desirable 105 °C point with heat loads of more than 650 W per VME 6U-sized module. And this cooling was done using a dielectric fluid (PAO) flow of 0.92 lpm, a fluid inlet temperature of 55 °C, and a total pressure drop of 12.5 psid across the module, all of which are within the performance and environmental requirements of the VITA 47 standard for VPX plug-in modules.

Figure 3


Annapolis Micro Systems
The FPGA Systems
Performance Leader!
FOPEN Radar Systems Software Defined Radio FLIR
SIGINT ELINT Digital Receivers Recording Systems
Annapolis
Micro Systems
Annapolis Micro
Systems, Inc.
High Performance
Signal Processing in
Scalable FPGA Computing
Above and Beyond -----
FPGA Acceleration
190 Admiral Cochrane Drive, Suite 130, Annapolis, Maryland 21401
wfinfo@annapmicro.com (410) 841-2514 www.annapmicro.com
Photo courtesy of USN

Liquid flow-through provides many benefits

The benefits of liquid flow-through cooling are many. They are derived from the superior heat-absorbing capacity of the cooling media.

LFT implementation is facilitated by the fact that liquids are easily routed to hot spots in the system and then to the most advantageous, possibly remote, location to efficiently reject heat to the ambient environment. This allows significant reductions in cooling system space claim as compared to fan-driven air-cooling approaches. Liquid-cooling systems are readily adaptable for severe environmental conditions and, depending on system design, are not impacted by altitude. Moreover, the complexities and risks inherent in "wet" direct-spray cooling are avoided.

Liquid flow-through cooling solves not only today's thermal packaging problems, but also eliminates the thermal restrictions that limit power density, giving a wide margin of headroom for future tech refresh and system upgrades without requiring major cooling system overhauls. Liquid cooling is the sensible successor to forced-air, convection, and conduction methods.

Rex Harvey *serves as a principal design engineer on Parker Hannifin Corporation's Advanced Cooling Systems team. He has spent the past 20 years of his 38-year engineering career in aerospace-related liquid spray atomization, most recently in the liquid-cooling application field. He has headed many Parker programs and holds more than a dozen patents. His technical background includes membership in AIAA, ASME, SAE, SME, and VSO professional societies and standards organizations. Rex holds a BS in Mechanical Engineering from Iowa State University and an MS in Systems Management from the University of Southern California. He can be reached at rharvey@parker.com.*

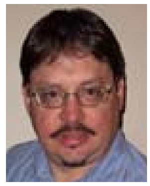

Andy Odar *is a principal systems engineer on Parker Hannifin Corporation's Advanced Cooling Systems team. He has been involved in liquid-cooling projects for the past five years, including the Cray X1 Supercomputer spray-cooling system and several cooling implementations for the Army's Future Combat Systems program. Previously, he designed fuel injectors for gas turbine engines for military, commercial, and power generation platforms, holding two patents in airblast and macrolaminate spray atomizer designs. He received his BS in Mechanical Engineering from Case Western Reserve University. He can be reached at aodar@parker.com.*

Parker Hannifin Corporation
440-954-8100
www.parker.com/advancedcooling

Conduction cooling reaches new power heights

By Ivan Straznicky

Power dissipation from rugged circuit cards continues to rise, calling into question how long incumbent card cooling approaches, like conduction, will last. Recent advances show that 160+ W can be conduction cooled at 85 °C card edge temperature. The setting of this new bar, and ongoing progress, attest to the staying power of conduction cooling, as it continues to meet the demanding requirements of advanced military electronic systems.

The thermal packaging approach of choice for circuit cards on many rugged military embedded systems is conduction cooling. This method provides inherent advantages including shock and vibration protection, the ability to house circuit cards in sealed chassis, and passive cooling. However, the conduction approach's cooling limits were considered a drawback several years ago because they were too low for advanced electronics and did not provide enough cooling headroom.

Limits of 50-60 W were once claimed to be ceilings for conduction cooling, with 100 W as a more recently assumed limit. In addition, the thermal contact resistance between conduction card edge and chassis rail was viewed as conduction's "Achilles' heel" because of the large relative temperature delta it created along the heat removal path.

During the past several years, these shortcomings have been the focus of significant research, analysis, design, and testing activity.

The results of this work have steadily advanced conduction-cooling capability, with the outcome that conduction cooling is now much more attractive for current *and* future electronics. New designs are now capable of cooling over 160 W (at 85 °C card edge). Some of the advances that have made this possible are: embedded heat pipes, new thermal interface materials, emerging standards such as VPX-REDI, thermal contact resistance reduction, and chassis-level liquid cooling.

Embedded heat pipes

A recent advance in conduction cooling is the use of embedded heat pipes in conduction frames. Heat pipes are thermally interesting because they can move large amounts of heat with very little temperature difference, with no input power requirement, and no moving parts. Heat pipes have been used for cooling since the 1960s, with early use focused on space applications because of both the aforementioned benefits and their ability

Figure 1

to work in microgravity environments. Since their introduction, the use of heat pipes has greatly increased, particularly for terrestrial commercial electronics cooling where the gravity-dependence of heat pipes is not a drawback. Another attraction of heat pipes is their relatively low cost compared to other high-power cooling solutions.

Heat pipe use in rugged military environments poses several challenges, such as orientation-dependence, shock and vibration, acceleration limitations, and cold start-up. Curtiss-Wright Controls Embedded Computing has been researching and testing heat pipe use in harsh environments. This work has led to CWCEC's patented design of an embedded heat pipe frame that successfully meets these challenges. The highly engineered design will be used for advanced processors that cannot be cooled with typical conduction approaches, such as the VPX-185 single board computer shown in Figure 1.

Thermal Interface Materials

Conduction-cooling design is also benefiting from a greater understanding of Thermal Interface Materials (TIMs). TIMs are used between mating conduction surfaces along the heat removal path and have been an essential part of conduction cooling for many years now. With the increased power dissipation and densities of recent processors, TIMs have become even more important. Not coincidentally, TIM manufacturers have introduced a steady stream of improved products that attempt to keep up with the power increases. These products include: gap pads, thermal putties, thermal adhesives, Phase Change Materials (PCMs), filled gels, and even solders. Figure 2 shows typical thermal resistances associated with use of these materials[1].

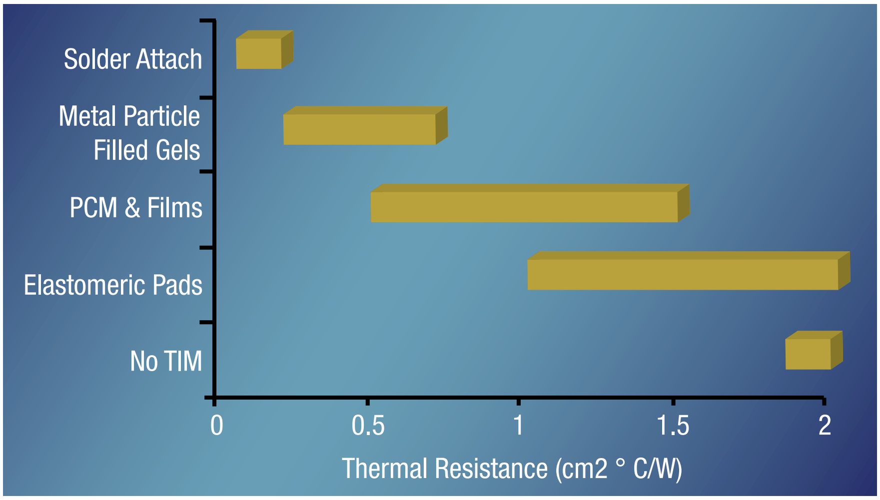

Figure 2

A key characteristic of any TIM is its thermal conductivity, or the application-dependent inverse, thermal resistance. Thermal conductivity values are published by vendors for use in thermal analyses like Finite Element Analysis (FEA). Years of testing, however, have shown that these values are consistently higher than independently measured figures. Large discrepancies are not uncommon, for example, in one case the measured value was less than 10 percent of the published value. Using the published value in a thermal analysis would have caused a substantial under-prediction in junction temperature, greatly increasing the risk of thermal failure. As more and more "new and improved" TIMs come to market, testing for thermal conductivity (and other material properties) is essential to separate the winners from the losers.

VPX-REDI (aka VITA 48)

New open standards are also helping to foster conduction-cooling improvements. VITA 48.2 is standardizing mechanical formats that are compatible with IEEE 1101.2 card slots. It also allows for increased functional density and thermal management with a 1" pitch (the IEEE 1101.2 pitch is 0.8"). The increased pitch, along with other changes, offers cooling improvements that extend conduction cooling to higher allowable powers.

Card edge thermal contact resistance

Thermal contact resistance is another field where increased research has helped improve conduction-cooling design. When two surfaces in contact have heat flowing across their junction, a measurable temperature difference arises caused by contact resistance. This thermal resistance is because surfaces are not perfectly smooth at the microscopic level (see Figure 3 for a schematic view of the interface); thus, only a few points of actual contact are made with the rest of the area consisting of air gaps. In the absence of a TIM (that is, a conduction-card edge and chassis rail), the air's low thermal conductivity creates a substantial thermal resistance between mated surfaces. The value of this resistance is a complex combination of factors including surface finish, hardness, flatness, and apparent and actual contact area and pressure.

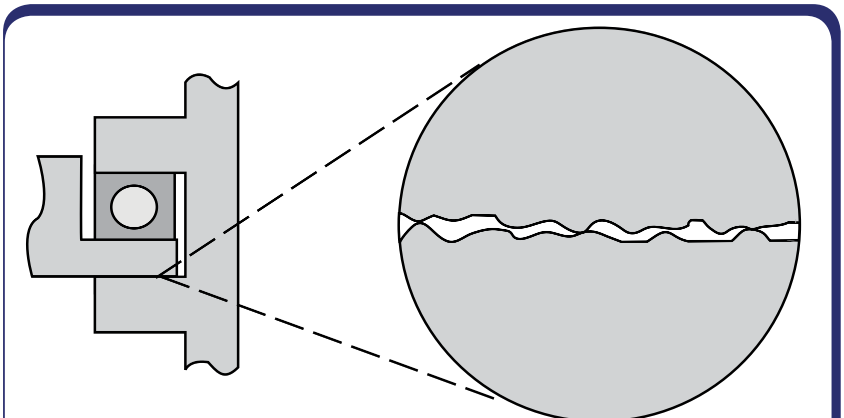

Figure 3

In the case of a typical IEEE 1101.2 conduction-card edge mated to a chassis rail, the contact resistance value used by many thermal engineers is in the range of 0.3 to 0.5 °C/W. As a result of several initiatives and based on test data, significantly lower contact resistance values have been achieved, as low as 0.1 °C/W. This reduction in thermal resistance enables the chassis rail, and more importantly its cooling medium, to be at a higher temperature for the same amount of card power dissipation and results in more efficient cooling.

Chassis cooling

While heat pipes, TIMs, VPX-REDI, and contact resistance have all improved conduction-cooling design at the card level, significant advances such as liquid cooling are also being made at the chassis level. With it now possible to cool 160+ W at the card level per slot, chassis cooling presents itself as a potential limitation. The typical chassis-level approaches of natural convection, forced air, or conduction will not be able to cool heat loads with even just a few of these powerful cards.

Fortunately, new chassis designs are being introduced that can cool significantly more heat than previous generations. Many of these designs use liquid in the sidewalls (see Figure 4, photo courtesy of Hybricon) because of inherently superior thermal properties compared to air, for example, thermal conductivity, density, and

Figure 4

specific heat. The liquid can be supplied by the platform's Environmental Control System (ECS), or some designs have optional integrated heat exchangers to cool the warmed liquid. These and other innovative chassis designs are advertising the ability to cool conduction cards of 100-200 W, making them compatible with the new generation of high-functional-density cards.

The conduction cooling outlook

The aforementioned advances continue to push the limits of conduction-cooling capability. However, diminishing returns are being seen even with increased effort in thermal research, analysis, design, and testing. Fortunately, processor power dissipations will reach a plateau in the near term, with more efficient architectures and transistor material advances. This will allow the state of the art in conduction cooling to advance further and continue to meet many rugged military system requirements.

Ivan Straznicky *is a senior staff mechanical engineer for Curtiss-Wright Controls Embedded Computing. He has 18 years of manufacturing and mechanical engineering and management experience in the military and aerospace industry. His responsibilities include advanced thermal and packaging technologies for the company's products. Ivan has a degree in Mechanical Engineering from McGill University in Montreal, Canada. He can be contacted at ivan.straznicky@curtisswright.com.*

Curtiss-Wright Controls Embedded Computing
613-599-9199
www.cwcembedded.com

References:

1. *ElectronicsCooling,* Vol. 11, No. 4, November 2005.

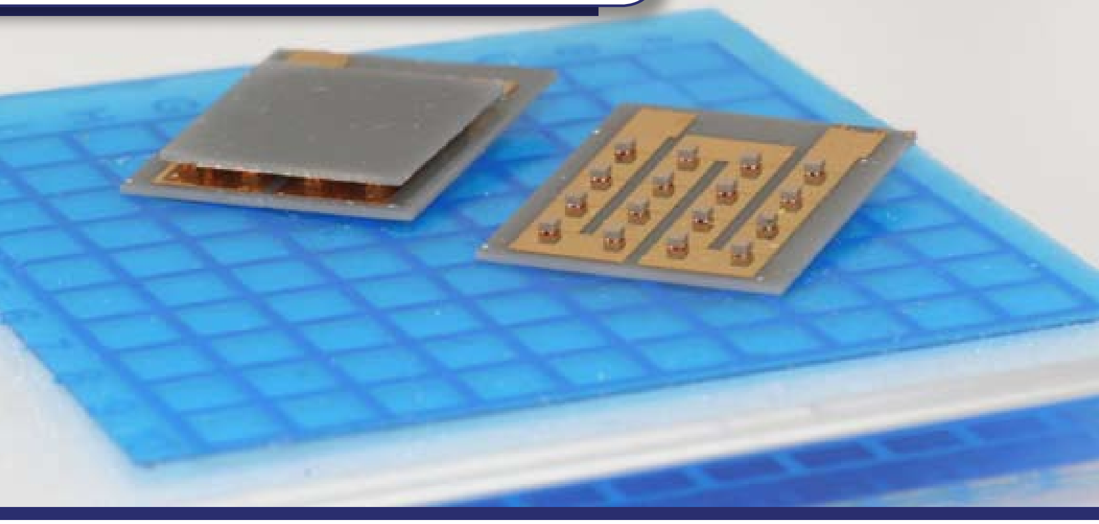

Embedded thermoelectrics: Cooling ICs in harsh environments

By Seri Lee, PhD

The demand for small, compact electronic devices with increased capabilities and functionality is a driving force in today's technology markets. Whether the device is manufactured for the consumer, military, aerospace, medical, communications, or automotive markets, its evolutionary path is the same – smaller and faster with more functionality. Design engineers face many challenges in meeting this continuously raised bar. A particularly vexing problem is removing heat from integrated circuits.

Thermoelectric modules based on thin-film thermoelectric materials have higher design flexibility, lower cost structure, and better performance than conventional thermoelectric devices. Thin-film thermoelectric modules are currently being demonstrated as a unique alternative to conventional technologies for cooling ICs in rugged and harsh environments.

Avionic and military application electronic devices and equipment used in harsh and extreme environments have many significant thermal challenges above and beyond those encountered in consumer-based products. Electronic systems in these applications consist of rack-mounted printed circuit boards stacked one on top of the other in a sealed enclosure. These systems allow no airflow for convection cooling space within the chassis to accommodate conventional thermal management techniques.

A typical option is to use a conduction plate laid over the board, in contact with the top surface of the components. This conduction plate, often made of machined aluminum, is fastened to the board and provides both structural support and thermal transport.

The thermal management challenges of ICs in harsh environment systems include:

- **» *Very small size, low weight, and very high reliability –*** These prohibit the use of large heat sinks, fans, liquid cooling, jet spray impingement, and so on.
- **» *Wide operating temperature ranges (-40 °C to + 120 °C)* –** High temperature ambient limits the "conduction thermal budget" for dissipating heat.
- **» *Use of COTS components, typically rated (-20 °C to +80 °C) in systems conforming to military specifications –*** Components need to be heated in cold ambient (that is, from -40 °C to -20 °C) and cooled in hot ambient (that is, from +120 °C to +80 °C).
- **» *Hot spot thermal management on ICs arising from die-level non-uniform power dissipation –*** Close packing of high-performance circuits results in localized high heat fluxes leading in turn to hot spots that limit performance and reliability.
- **» *Hot spot thermal management on boards arising from higher packing density –*** Requirements toward smaller form factor systems drive high-heat-generating components to be placed on a compact board layout.

Accordingly, *embedded Thermoelectric Coolers* (eTECs) utilizing thin-film technology are providing a remedy for these IC cooling issues.

Embedding a thermoelectric cooler

Thin-film thermoelectric devices are designed to address site-specific thermal management problems. Thin-film thermoelectric

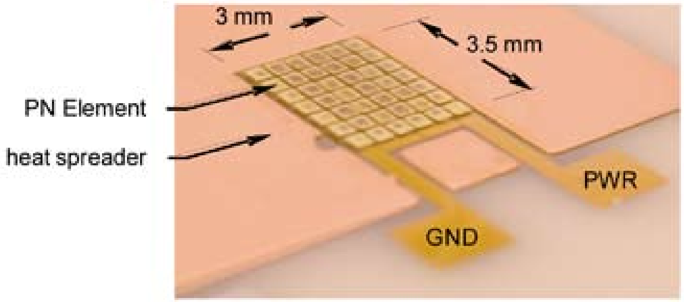

Figure 1

Another unique area of need for devices in military applications is the need to find a cost-effective method for closing the thermal gap created by the difference between the rated range of typical military operating temperatures (-40 °C to +120 °C) and the narrower range often specified for COTS components (typically -20 °C to + 80 °C). Furthermore, when implemented in conjunction with a feedback control system, a thin film thermoelectric

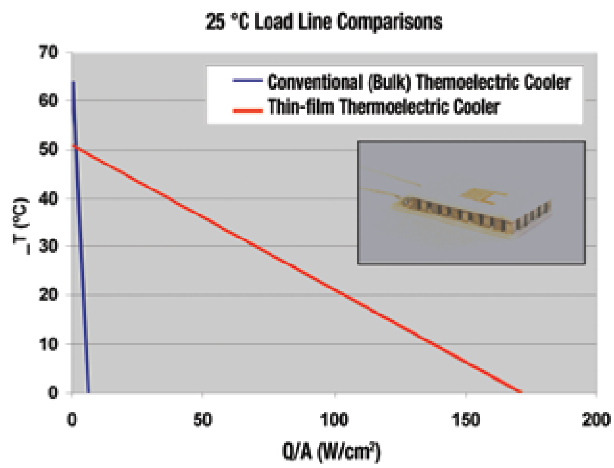

Figure 2

materials can be grown by conventional semiconductor deposition methods. The resulting device, the thin-film thermoelectric discrete module shown in Figure 1, is smaller and thinner than conventional products as it measures 3.5 mm x 3.0 mm x 0.1 mm in size. Larger devices can be fabricated, and devices consisting of a single PN-couple as small as 0.3 mm x 0.3 mm are feasible.

Thermoelectric modules are most commonly compared on the basis of their load lines. A load line is defined as the maximum temperature difference, ΔT that can be sustained across the TEC thickness as a function of heat-pumping power density, Q/A. Figure 2 shows the load lines of the eTEC pictured in Figure 1 and a typical bulk device, both measured at room temperature. For reference, the thin-film thermoelectric cooler is shown relative to a conventional bulk TEC in the inset.

The comparison in Figure 2 illustrates a fundamentally new operating regime offered by thin-film thermoelectric modules, producing a ΔT_{max} of more than 60 °C across the thickness of a piece of paper while pumping a maximum power density (Q_{max}/Area) of more than 150 W/cm^2. The eTEC's benefits over conventional bulk TECs include:

- The device is typically 10x thinner, enabling unobtrusive integration to the heat source for site-specific cooling.
- Heat-pumping density, which is inversely proportional to the thickness of the thermoelectric material, is as much as 30x greater.
- Operation in a high Coefficient of Performance (COP) regime while still pumping at a high power density.
- The input power can be dynamically controlled to provide active cooling and/or heating.
- Provides very fast thermal response time constants that are orders of magnitude smaller, allowing rapid cooling and heating to maintain a precise temperature.
- Solid-state design with no moving parts or fluids provides high reliability.

component can be used to rapidly respond and maintain the in-situ operating temperatures.

Implementing eTECs

Thin-film thermoelectric cooling module technology provides many new design “knobs” for an application. The basic PN-couple building block can be implemented with varying packing density and height in final assembly to achieve enhanced performance characteristics in a thermoelectric module (see Figure 3). A single PN-couple thin-film thermoelectric device (top-left) is as small as 0.3 mm x 0.3 mm x 0.1 mm high; these devices can be distributed with varying packing fractions to cool a larger surface area with a lower power density (top-right) or localized spots with a very high power density (bottom). Smaller yet, thermal bumps no greater than 0.06 mm high can be placed directly onto a hot spot for maximum cooling potential.

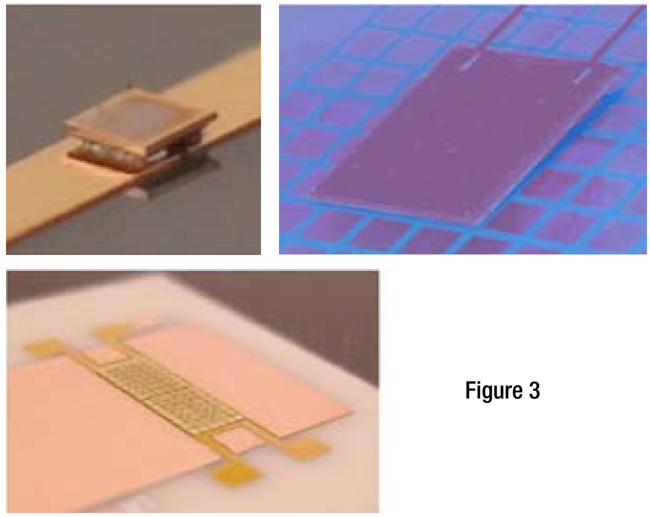

Figure 3

Thermoelectric cooling devices may be used in a variety of implementations including the following:

- Embedded in a component’s thermal solution (that is, heat sink, vapor chamber, heat pipe, and so forth).
- Embedded in a component’s package.
- Embedded in a system to locally reduce the surface temperature of a hot spot or hot zone on a board or case.

By tiling multiple thermoelectric modules in array forms, a larger cooling area is also possible. This is particularly feasible for cooling high-heat-generating packages on a board or on the surface of a conduction plate where the density of heat that needs to be removed is high, and the form factor thickness in the system does not allow other types of thermal solution options.

Thin-film technology meets thermal challenges

The need for higher-performance electronics in avionics and military applications is resulting in greater heat dissipation from the system and higher heat density from its components. Electronic devices and equipment in such harsh and extreme environments have a number of added thermal challenges. By leveraging thin-film technology advancements in the miniaturization and efficiency of solid-state cooling devices, it is possible to address these site-specific, IC thermal management problems.

__Seri Lee__ is chief technology officer at Nextreme Thermal Solutions, Inc., where he is responsible for directing advanced thermal and power generation technology projects. Before joining Nextreme, he worked at Intel as senior thermal scientist. He is a member of the ASME Heat Transfer Division K-16 Committee on Heat Transfer in Electronic Equipment and the IEEE/SemiTherm Executive Committee. He holds 16 patents covering a wide range of thermal issues in electronics and received his PhD from the University of Waterloo. Seri can be contacted at slee@nextreme.com.

Nextreme Thermal Solutions, Inc.
919-990-8300
www.nextreme.com

Accelerating algorithms in FPGAs, one process at a time

By David Pellerin

FPGAs have become commonplace in embedded systems and are now beginning to appear in high-performance, server-class computing applications as well. This shift toward hardware-accelerated computing has been driven by two factors: increased FPGA device capabilities and improved software-to-hardware tool flows. Emerging tools can help demystify FPGAs, making them more accessible to software application programmers.

Programmable logic has come a long way since its humble beginnings in the mid to late 1970s. FPGAs first appeared in 1984 and now serve as highly capable platforms for high-performance embedded computing. Developers use FPGAs to integrate embedded processors with custom hardware peripherals, forming single-chip embedded systems. FPGAs also can serve as coprocessors for enterprise and scientific computing (see Figure 1).

FPGAs have proven themselves in a broad range of hardware applications, including telecommunications, aerospace and defense, medical electronics, and consumer electronics. FPGAs can be found inside set-top boxes, big-screen televisions, and car dashboards as well as on Mars in the robot brains of NASA's Mars Exploration Rovers *Spirit* and *Opportunity*.

Because they are memory technologies, FPGAs have increased in size and density annually in accordance with Moore's Law. FPGAs have grown to the point where they can now be considered for entirely new types of computing applications – applications that until recently would have been in the domain of traditional CPU architectures. As a result, the high-performance computing community is beginning to experiment with and deploy these devices to accelerate algorithms in areas as diverse as financial modeling, scientific computing, and bioinformatics.

On paper, an FPGA sounds like a perfect fit for these types of applications. From a gigaflops-per-watt perspective, no other fundamental computing technology can match a modern FPGA. But there is a catch: FPGAs are notoriously difficult for software programmers to use. The two most common programming

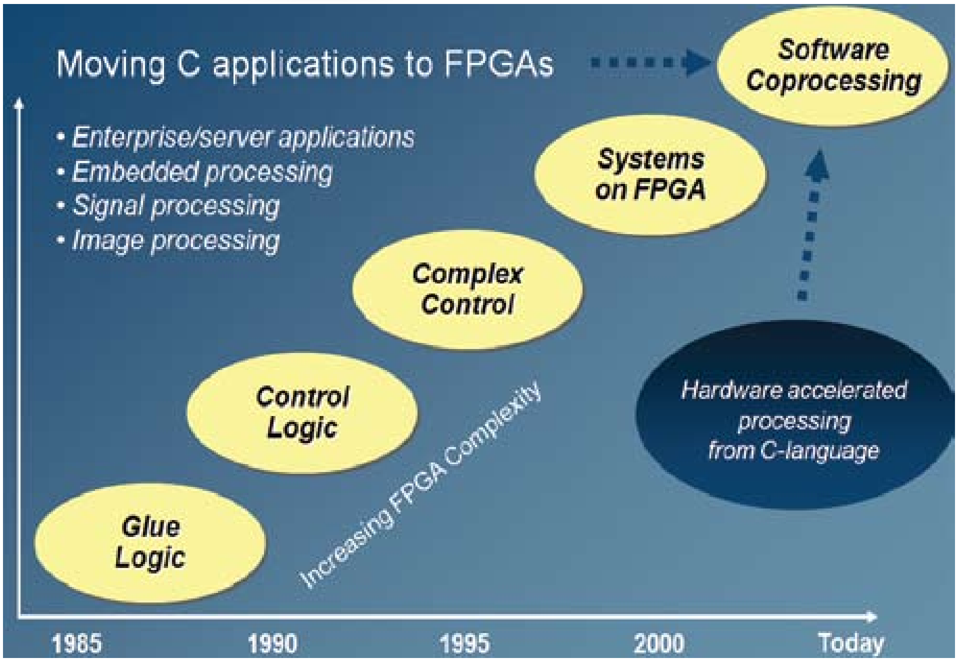

Figure 1

languages for FPGA programming, VHDL and Verilog, are inscrutable to the average software programmer, and the tool flow from those languages to working hardware can be difficult to comprehend and slow to complete.

However, software-to-FPGA compiler tools are now available and maturing every year. These tools offer significant benefits, not the least of which is design productivity. C-to-hardware tools in particular (see Figure 2) are gaining converts from both software and hardware design backgrounds. When properly applied, these tools can give software programmers access to FPGAs as computing devices.

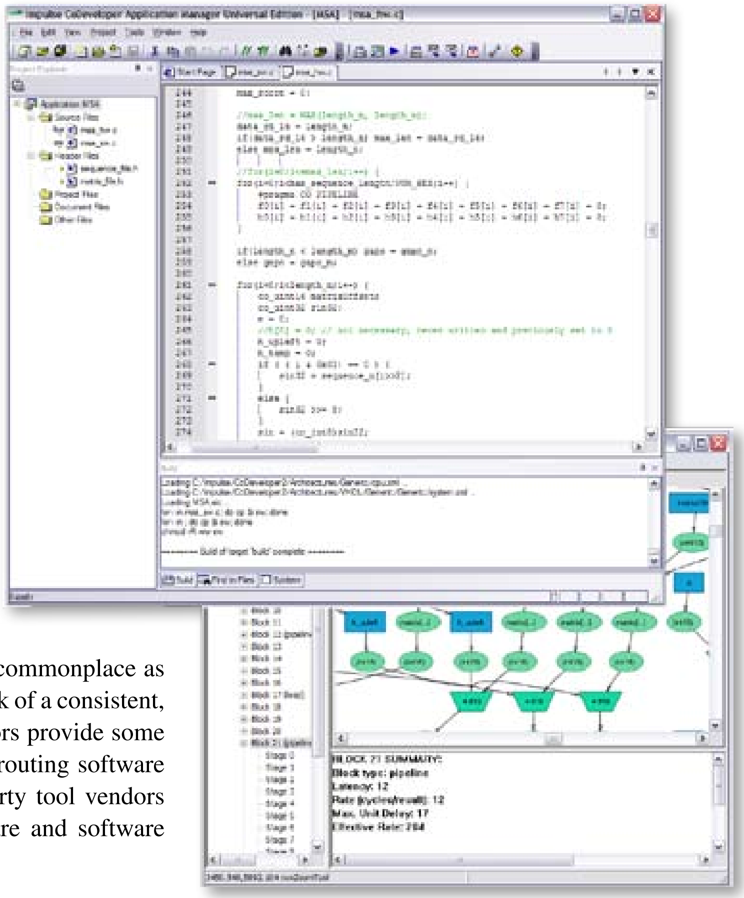

Figure 2

Hurdles to overcome

At this point in the evolution of FPGAs, two near-term hurdles must be overcome before FPGAs become truly commonplace as software acceleration platforms. The first hurdle is the lack of a consistent, integrated, and portable tool flow. FPGA platform vendors provide some parts of this tool flow, most notably the placement and routing software required to implement logic within an FPGA. Third-party tool vendors provide other parts of the flow, including both hardware and software front-end design tools.

FPGA vendors recognize the value of device-independent software tools, and have historically worked to foster and promote partner tools ecosystems for higher-level FPGA design. These tool ecosystems encourage competition and innovation, but also present a bewildering set of choices to new users, with very few standards for interoperability. For FPGAs to flourish in software applications, the tools environment for FPGAs must become more consistent and software-like, while at the same time providing ways to take direct advantage of what an FPGA has to offer in terms of performance. Organizations such as OpenFPGA can play a role in improving interoperability.

The second critical hurdle to overcome is software designers' knowledge gap concerning what an FPGA is and what it can do. What does the term *field programmable* mean to a software programmer familiar with traditional processors? What is a *gate array* and how does it relate to running software algorithms? What does the term *electronic system-level design* mean to a financial or scientific application developer?

Demystifying the FPGA as a computing platform is therefore critical. Just as a

> **"The challenge moving forward is to provide tools that best utilize on-chip processing capabilities without forcing FPGA developers to become hardware design experts."**

good development suite for traditional processors encourages efficient programming methods for processors, a tool suite for FPGAs must encourage programming methods that allow software programmers to utilize available FPGA resources and make intelligent decisions regarding hardware/software partitioning and source code optimization.

Software-to-hardware tools break down barriers

Because FPGAs have become commonplace in the domain of embedded systems, it makes sense for enterprise and scientific application developers to examine what has and has not worked for embedded systems designers.

Embedded systems design teams are often composed of hardware and software specialists. In some cases, the embedded systems designer will have the complete skill set needed to program a complex application *and* design hardware components. However, more often than not, the disciplines of hardware and software design have remained distinct. Today, software-to-hardware compiler tools are beginning to change this differentiation.

Most embedded systems design teams freeze the hardware first, then continue to work on software until later in the project cycle. Some hardware/software codevelopment may occur to ensure that the overall system achieves its objectives; nonetheless, a schedule mismatch exists between the two groups. One key appeal of programmable logic is that it allows hardware freezing to occur later in development. FPGAs encourage rapid prototyping and design flexibility and allow for future modifications without the need to completely respin the board or system. Software-to-hardware tools build on the benefits of FPGAs by reducing the time to prototype and encouraging iterative, explorative design methods.

The advantages of FPGAs and software-to-hardware tools are exactly what higher-level, enterprise-type applications need. Algorithm performance and throughput are of paramount importance to software developers. Even more so than in embedded systems, software drives the project forward, leaving the hardware platform as nothing more than a platform, replaceable at any time.

Getting from C to FPGA hardware

For a software developer – someone not experienced with FPGAs or hardware design in general – the tool flow from software to hardware is challenging. Software teams have intellectual property and expertise developed using GCC, Visual Studio, Eclipse, or other software integrated development environments.

When moving into FPGAs, these software programmers must have a software-oriented design environment in which to work. If it is a simple application that can be represented as a single FPGA process, they may be able to wrap their legacy algorithm using FPGA-friendly API functions, recompile/cross-compile the application for an FPGA target, and go. But to achieve efficient results for larger, more complex applications, software programmers need to understand more about the tool flow and methodology. For a C-to-FPGA tool, the tool flow includes the following fundamental steps:

1. Design and debug the application using standard C tools
2. Partition the C language algorithm between hardware and software
3. Optimize the C code for an FPGA, introducing parallelism
4. Synthesize and generate a bit file via FPGA place and route tools
5. Download to an FPGA development board or system for testing

While these steps sound simple, they can represent an iteration time measured in hours or even days. Most successful FPGA projects are therefore built in steps, using partitioning techniques to reduce compile times.

The methodology for a single FPGA accelerator process is straightforward. For applications requiring multiple hardware processes, however, it takes a bit more effort. Developers must focus on one critical function at a time and use iterative partitioning and optimization methods to move incrementally more of the application into the FPGA.

There is no magic tool that will, without fail, take arbitrary C code in and generate efficient hardware. C-to-FPGA tools can help make FPGAs more accessible for software/hardware design teams but require some specialized knowledge. Modern FPGAs offer impressive sets of features, including embedded microprocessor cores. The challenge moving forward is to provide tools that best utilize on-chip processing capabilities without forcing FPGA developers to become hardware design experts.

FPGA reconfiguration gaining a greater role

FPGAs' future for high-performance computing appears positive. Software developers are learning about these devices, and higher-level tools including Impulse C can help them start their first projects. Looking forward, developers can expect that dynamic FPGA reconfiguration will become increasingly important. This reconfiguration will require some level of operating system support to be efficiently performed and the ability to partially reconfigure the FPGA without resetting it. Partial reconfiguration is expected to be an enabler for shorter place and route times and correspondingly quicker design iterations. Standards must emerge to allow each tool in the chain to be evaluated on its merits.

***David Pellerin** is cofounder and CTO of Impulse Accelerated Technologies in Kirkland, Washington. He is a programmable logic industry veteran and the author of five books, including* Practical FPGA Programming in C*,* VHDL Made Easy*, and* Practical Design Using Programmable Logic. *David was formerly the president and founder of Accolade Design Automation, an FPGA tools and training provider. Prior to that, he spent more than a decade as staff member of Data I/O Corporation, where he was a key contributor to the development of programmable logic device and FPGA tools. David can be contacted at david.pellerin@ImpulseC.com.*

Impulse Accelerated Technologies, Inc.
425-605-9543
www.ImpulseC.com

Mars Reconnaissance Orbiter (MRO) uses embedded technology for martian images

By John Carbone

Image courtesy of NASA/JPL/Arizona State University

The High Resolution Imaging Science Experiment (HiRISE), now flying on the MRO, provides the highest-resolution images taken of deep space. From 190 miles away, HiRISE can identify objects only a meter across, providing a better understanding of the geologic and climate processes on Mars and making it possible to scout out possible landing sites. It will play a key role in the MRO's mission of characterizing the surface, subsurface, and atmosphere of Mars and identifying landing sites for future missions.

When they wrote the embedded software that controls HiRISE, a team of Ball Aerospace engineers led by Steve Tarr knew they only had one chance to get it right. If there was a serious flaw in the program, the $720 million spacecraft would have no more value than a digital camera dropped in a bathtub. Tarr and his team wrote 20,000 lines of code running on an Atmel SPARC TSC695F processor, using Express Logic's ThreadX RTOS. The code sets up and operates the camera and sends the resulting pictures to the spacecraft software. The spacecraft software transmits the pictures to the ground. The software has worked flawlessly, resulting in history-making photographs such as one that shows the Opportunity rover moving across the surface of Mars (Figure 1, courtesy of NASA).

The MRO has four main goals: determine whether life ever existed on Mars, characterize the climate of Mars, characterize the geology of Mars, and identify promising locations for scientific study. The HiRISE camera plays a key role in each of the objectives. HiRISE can zero in on water-related surface features such as outflow channels from ancient floods. HiRISE will look particularly for geologic settings that indicate the presence of liquid water on the surface at some point in the planet's history. Hundreds of locales will be examined in unprecedented detail to reveal water-related mineralogy and water's role in shaping the terrain. The camera can even hone in on rocks as small as 3 or 4 feet to help evaluate the safety of potential landing sites.

HiRISE provides images of the surface of Mars to a much finer resolution and higher level of contrast than ever before. Light enters the front of the camera, is gathered by a 50-centimeter diameter primary mirror, and is then sent by a series of other mirrors to be focused on the detectors. HiRISE doesn't take a single image of a scene all at once but rather grabs one row of about 20,000 pixels at a time, imaging a 6 km swath, while the spacecraft sweeps over the surface of Mars.

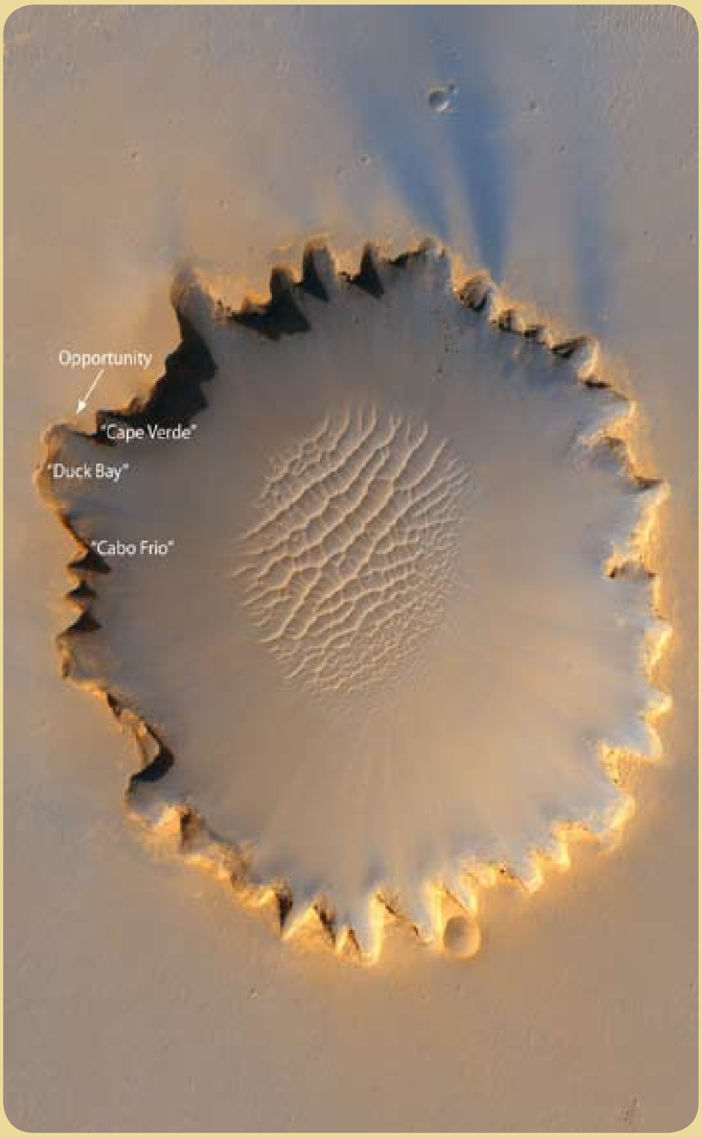

Figure 1

The Time Delay and Integration (TDI) method improves image quality as the spacecraft covers the ground at a speed of approximately 3,200 meters per second. The TDI continuously collects and reads out the accumulated signal from the Charge Coupled Device (CCD) one row at a time to match the ground. As a row is read out, the charges in the remaining rows are shifted down by one row. This has the same effect as increasing the exposure time in a conventional camera.

While HiRISE delivers a high-resolution image of Mars surfaces, the Context Camera (CTX) imager provides a wider-angle view of the same region and the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) multispectral imager obtains near-infrared spectral images to identify surface composition.

Functionality of embedded software

The HiRISE embedded software written by Tarr and his group performs a number of different functions and is integrated into an overall system of programs. For example, the camera is aimed by the spacecraft software, a separate application written by Jet Propulsion Laboratory engineers, which controls thrusters that position and orient the spacecraft so that the camera is pointing at the area that is to be photographed. When positioning is correct, the spacecraft software directs the HiRISE software to take the picture.

The HiRISE software configures the camera by setting the TDI pixel line time to match the ground velocity and maintain alignment with the spacecraft's motion over Mars. The software then turns on the camera and begins processing the images generated by the CCDs by adding header information and sending them to the spacecraft image storage system, which in turn sends them to a ground station. In the background, the HiRISE software captures engineering information from sensors located throughout the camera in order to help identify and diagnose problems. The sensor readings provide information such as the temperature of the camera's key components. HiRISE evaluates these readings and issues a warning if any parameter moves outside the normal range.

Tarr and Ball Aerospace addressed a similar challenge when they developed software for NASA's Deep Impact mission in 2003. This mission involved crashing a spacecraft into a comet while a second spacecraft photographed the collision and recorded data about the dust and gas released by the explosion. Tarr's team at Ball wrote software to control the three imaging instruments that generated images of celestial bodies with known positions so that the spacecraft could determine its position and compute its course to its destination. The instruments also photographed the exploding comet, as shown in Figure 2 (courtesy of NASA).

Figure 2

Just as with the Deep Impact mission, Tarr and his team selected the ThreadX Real-Time Operating System (RTOS) from Express Logic, San Diego, California to manage the scheduling of the application threads and service interrupts, and to pass the messages needed for the cameras to operate correctly. On July 4, 2005 the Deep Impact spacecraft kept its rendezvous with the Tempel 1 comet, plunged its impactor into the comet's surface, and transmitted more than 4,500 photographs of the event back to Earth.

Reusing code from earlier project

The success of the Deep Impact application made it natural to reuse as much of its technology and code on the HiRISE software as possible. Yet there were substantial differences in the HiRISE application that required a substantial development project.

First, the Deep Impact camera used a shutter to take a single image, while HiRISE uses a more complex push-broom approach to sweep a continuous image. This, along with TDI use, increased the complexity of the CCD interfacing task. The application captures the bits from the CCDs, then, when required by the TDI configuration, sends the data to charge the next CDI. Another difference is that HiRISE uses a heater to maintain the camera at a constant temperature to improve image quality, and the new software application must control the heater.

When the image is fully formed, the embedded code adds header information and sends the data to the navigation software, which communicates with the ground tracking system.

Multitasking critical to application

Multitasking is a critical HiRISE software requirement, such as checking sensor readings at the same time it processes the images from the CCDs. Tarr and his team provided this capability by taking advantage of the RTOS's Interrupt Service Routines (ISRs). The ISRs that poll sensors, for example, are triggered in response to periodic timers.

Since application code execution is preempted during the execution of an ISR, the Ball Aerospace team's application minimizes the amount of code in the ISR and relies instead on a different application thread to complete the processing. This approach allows the highest-priority application code to be executed as quickly as possible.

Camera has worked flawlessly

The new MRO software and the HiRISE camera have worked flawlessly since the camera took the first high-resolution pictures of Mars on September 29, 2006. HiRISE reveals unprecedented detail in its images taken of Mars, showing seemingly endless fields of sand dunes, some carved by gullies that possibly formed when sunlight heated carbon dioxide or water frost in the dunes, triggering avalanches of flowing sand.

Other HiRISE images show layered arid terrains that resemble landscapes protected as national parks on our own planet, along with a fossil delta inside a crater that once held a lake. HiRISE images resolve meter-sized blocks within the delta channel that may be blocks of sand and gravel carried along as the channels eroded.

HiRISE images also capture numerous impact craters, including the Endurance crater that NASA's Opportunity rover explored for 10 months of its now nearly four-year mission. Details visible in the HiRISE image of Opportunity's landing site show the parachute lying on the martian surface, Opportunity's heat shield at a different location, and the lander itself on the floor of the small impact crater where the airbag came to a stop.

Figure 3, courtesy of NASA, shows gullies in an unnamed Terra Sirenum crater on Mars. This region receives very little sunlight in the southern winter, and some areas consist of frost. At the latitude of this image, frost is most likely composed of water because the temperature is not low enough for carbon dioxide condensation.

Other images show layered polar terrains that likely record martian climate changes, and also polygon-patterned northern plains regions that are among candidate landing sites for the Phoenix Lander in 2008.

HiRISE: Higher resolution

The HiRISE camera is generating images of Mars at a higher resolution than ever before on a daily basis. Key to the success of this camera has been the embedded software that sets up, operates, and monitors it. This software flawlessly and efficiently handles the large number of simultaneous tasks required to keep the camera working. The result is a dramatic increase in our knowledge of the red planet.

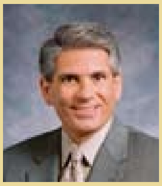

John Carbone, *vice president of marketing for Express Logic, has 35 years of experience in real-time computer systems and software, ranging from embedded system developer and FAE to vice president of sales and marketing. Prior to joining Express Logic, he was vice president of marketing for Green Hills Software. John has a BS in mathematics from Boston College. He can be contacted at sjcarbone@expresslogic.com.*

Express Logic
858-613-6640
www.expresslogic.com

See www.mil-embedded.com/articles/authors/lamie for more information on the Deep Impact mission.

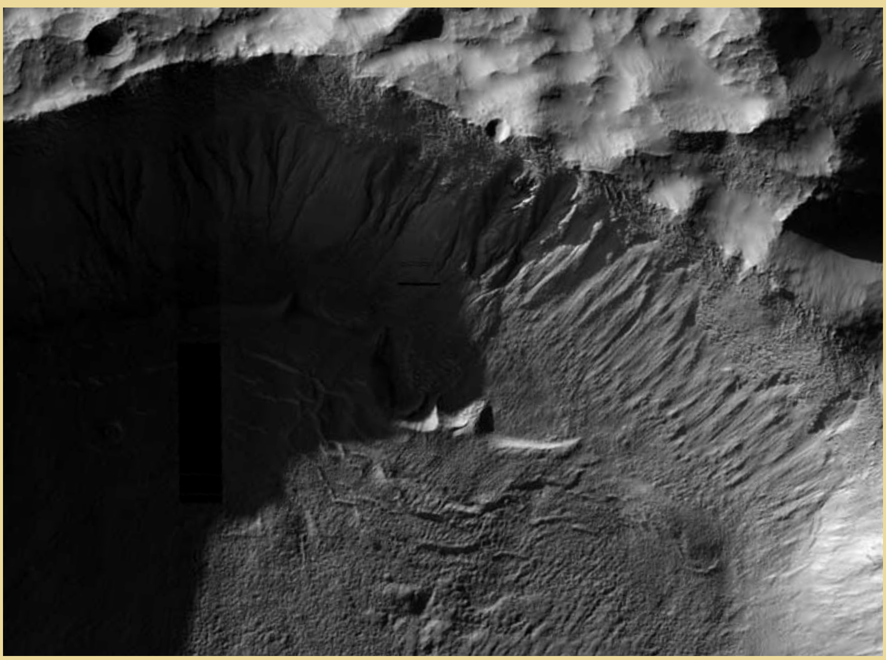

Figure 3

In space, on Earth – Relying on 3D technology to quickly deliver critical information

By Maureen Campbell

Ensuring the safety of NASA's space shuttles, the technology used to develop the Laser Camera System (LCS) is now also being used to streamline information in military and defense applications right here on Earth.

On February 1, 2003, NASA's *Columbia* shuttle suffered a catastrophic failure. The crew was tragically lost upon the shuttle's re-entry into the Earth's atmosphere. Investigations into the accident revealed that during launch, falling foam from Columbia's external tank struck the Reinforced Carbon-Carbon (RCC) panels on the leading edge of the left wing. This breach allowed superheated air to penetrate through insulation and melt the aluminum structure of the wing, thereby weakening the structure until increasing aerodynamic forces caused loss of control, failure of the wing, and ultimately the disintegration of the shuttle.

To ensure the safety of future missions, NASA required a means to determine the amount of damage – if any – sustained by a shuttle during the launch phase. In response, Neptec Design Group developed the Laser Camera System (LCS). Used to scan critical areas of the shuttle's exterior to assess areas of damage, the LCS now provides a view that was once invisible to the astronauts. Part of a larger sensor system installed on a 50-foot boom extension to the shuttles' robotic arm (Figure 1), the LCS made its first mission in July 2005 on NASA's STS-114 Return to Flight shuttle mission and is now a mandatory system for every shuttle mission.

The LCS utilizes 3D data once thought too cumbersome to process in real time. With unique data acquisition and processing methods, the root technology in the LCS is now being developed for a variety of uses in the defense industry, particularly in the area of Intelligence, Surveillance, and Reconnaissance (ISR) where the technology has been traditionally limited by low-bandwidth network environments. Such technologies include automated target recognition, pose estimation, and change detection. When sensors are capable of working autonomously (allowing the sensor to perform high-level filtering and queueing so that critical information is reported first), the bandwidth needed is dramatically reduced, providing results in near real time and reducing the use of bandwidth by orders of magnitude. The technology used to develop the LCS has proven effective in space and is now being used in a variety of terrestrial

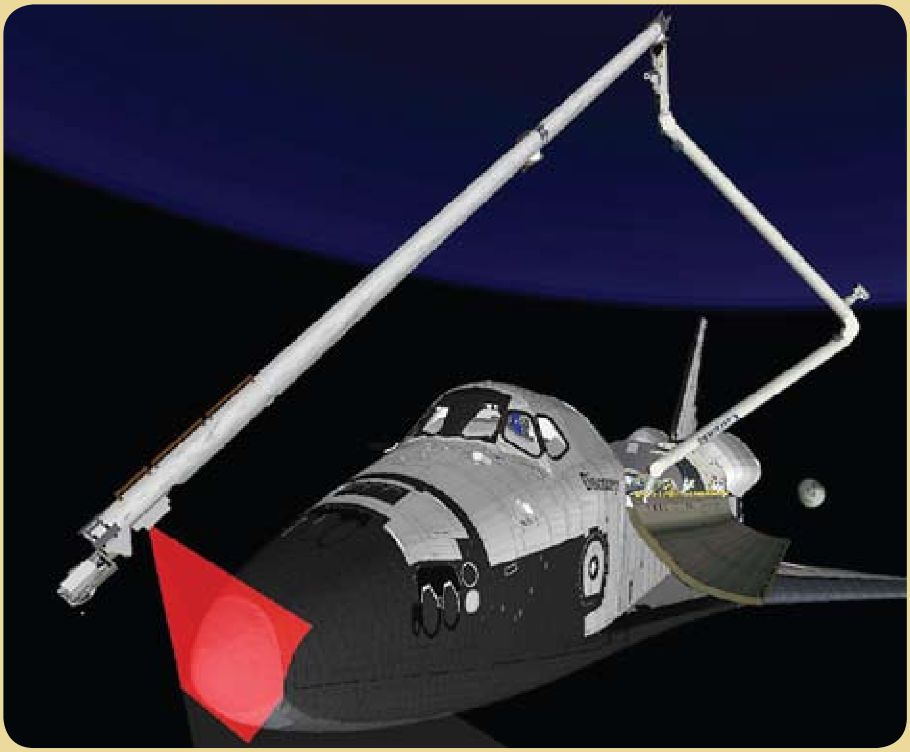

Figure 1

applications. In the defense and military industry, this technology is now being used to provide a variety of critical capabilities including object identification, inspection and tracking, and three-dimensional dynamic change detection.

The founding technology

Using a proprietary triangulation method, the LCS captures direct 3D measurements to create accurate three-dimensional models and images. Because 2D images are flat, one must rely on visual cues from the camera's perspective of the image to infer the required information. This can be ambiguous and at times misleading – especially in environments where light, gravity, and reference points vary greatly or simply do not exist. For example, imagine a picture of two buildings that appear to be the same height against a white background. Without a reference point, the height of these buildings being equal is a logical assumption. Now, consider placing one building in front of the other. Without additional visual cues, it would be concluded that one building is taller than the other, resulting in error. By also capturing direct measurements – meaning the distances from the laser source to every object within the field of view – the LCS can create an accurate three-dimensional scene from the measurements it has collected.

Fundamental features of the LCS are its data processing and acquisition methods. The data is stored so that it is easily accessible, which allows for rapid development and reconstruction of 3D models. As in the example above, if the important information is the distance between the two buildings and the square footage of each building, this can be accurately determined from the model that the LCS builds.

The system is based on an autosynchronous triangulation scanning principle originally developed by Canada's National Research Council. Using a continuous wave laser source, the LCS accurately measures both intensity and range information over a wide (30° x 30°) field of view from distances of 1 to 10 meters. This allows scanning rates that exceed other high-precision 3D metrology instruments by an order of magnitude or more. These combined measurements of angles and range can then be transformed through a previously determined calibration into "voxels" (three-dimensional pixels) for every sampled location. The two-axes steerable mirrors also provide a random access capability that allows the sensor to scan in arbitrary patterns at arbitrary resolutions to suit a wide variety of tasks, such as inspection of the shuttle's exterior and docking the shuttle at the International Space Station.

Figure 2 depicts an autosynchronous scanner using a triangulation principle to derive range information. An emitted laser beam reflects from the front surface of a double-sided mirror (1), bounces off a fixed mirror (2), reflects off the object (3), and is collected

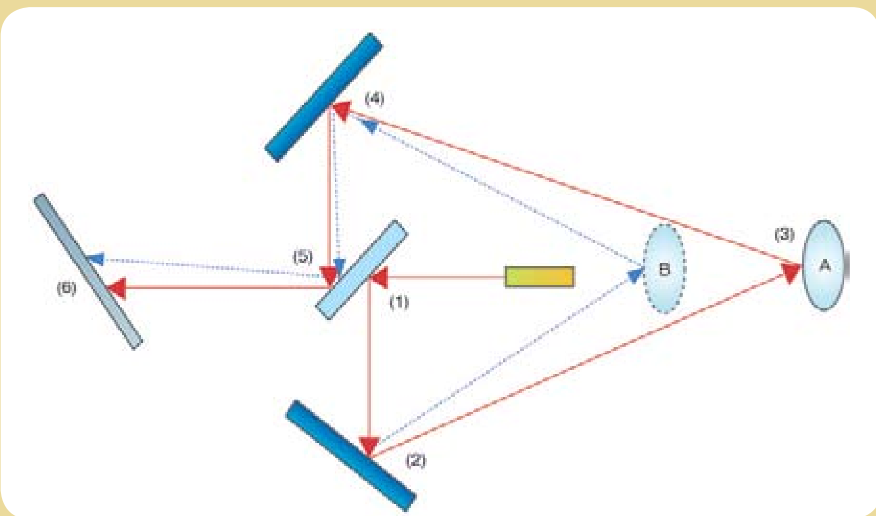

Figure 2

at another side mirror (4). From there, the beam is reflected from the back of a double-sided mirror (5) and directed onto a light-sensitive detector array (6). As the range of the object changes (from A to B), the angle of the collected ray changes, causing the position of the beam to change on the detector. Thus, the position of the laser beam on the detector is directly related to the range to the object. Scanning in one dimension is achieved by rotating the double-sided mirror with a scanning galvanometer. To sweep out a pattern in two axes, a second scanning mirror is placed after the side mirror (2).

NASA uses the LCS to inspect a shuttle's Thermal Protection System (TPS) for tiny fractures and cracks. Its primary function is to perform nondestructive, high-resolution scans. Faster than most 3D sensors, the LCS was designed for efficient acquisition and processing. The raw data is processed quickly, making the models easy to reconstruct. During missions, the LCS scans the TPS and collects 3D information in the form of binary data. This binary data is downlinked from the shuttle to Mission Control Center at Johnson Space Center in Houston, Texas. The data is then used to produce accurate three-dimensional models of the shuttle's surface (Figure 3) and inspect them for the presence of anomalies. With a variable of less than a millimeter, these simulated 3D images can be magnified and tilted to cast even the tiniest fractures and imperfections into stark relief.

Developing accurate information quickly, the speed and quality of the LCS information is dependent on a number of variables such as resolution, distance from surface, size of feature, and reflectivity of surface. For example, a full 30° x 30° field of view at 1,024 x 1,024 resolution will take approximately two to three minutes to scan. A full analysis of the data collected will take an average of 30 minutes to complete. The LCS demonstrates that 3D data acquisition and processing can work quickly and effectively.

Working in space

Space is an unforgiving environment, and space hardware must be very robust. Extreme temperature fluctuations, rapidly changing lighting conditions, lack of gravity, and vibrations due to launch are just a few of the challenges to overcome when designing and building space applications. Systems like the LCS – which is housed on the boom extension of the robotic arm – must be tough enough to survive the vibrations, as well as work and survive within the harsh elements.

In space, temperatures vary between -110 °C and +25 °C between sunrise and sunset during a 90-minute orbit around the Earth. The LCS required an insulator for the cold temperatures as well as a cooling system for the warm temperatures. Installed into the system enclosure's external walls is a highly reflective material that allows the LCS to operate in these extreme temperature fluctuations. Acting like a mirror, this reflecting material is silver-coated Teflon. With a solar absorptivity value of α=0.10, it reflects 90 percent of the energy coming from the sun, while the thermal emmissivity value is ε=0.84 and emits 84 percent of the heat

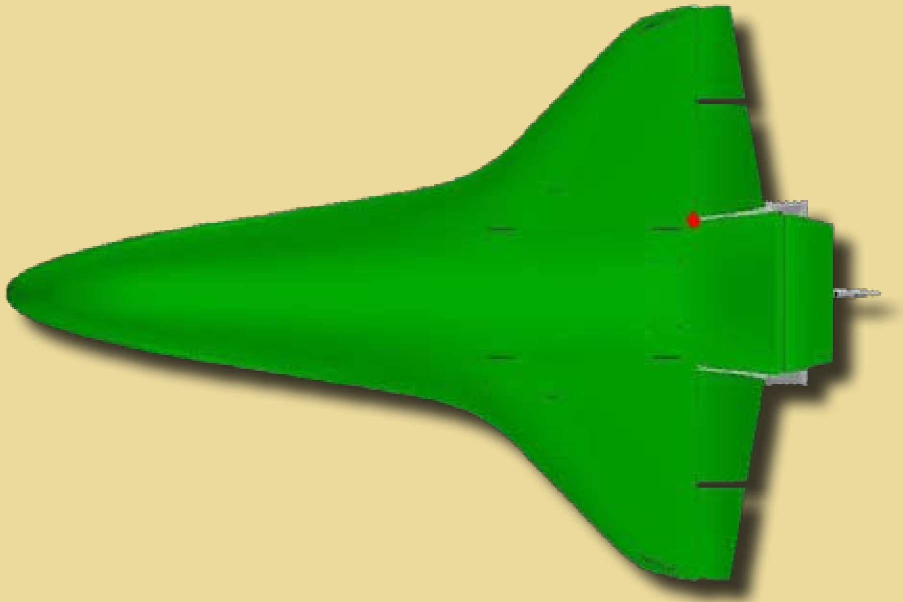

Figure 3

in the infrared range. Additionally, radiators were installed to allow the heat produced by electronics to radiate out of the enclosure.

The LCS optical path requires precise alignment that must be preserved during temperature changes and during the vibrations caused by the million pounds of thrust during liftoff and turbulence during re-entry. Optical subassemblies mounted on a rigid carbon laminate plate with a near-zero coefficient of thermal expansion counteract these vibrations.

Dynamic lighting conditions create temperature fluctuations and present challenges to the functionality of video or cameras. The wavelength of the laser source was chosen to be 1,500 nm for several reasons. First, it enhances the sensor's lighting immunity. On-orbit lighting is particularly challenging, and a laser frequency of 1,500 nm is 10 times more solar immune than a 532 nm green laser, which emits near the maximum of the solar spectrum. An optical band-pass filter located in front of the focusing lens ensures that only light of the desired wavelength reaches the detector. This allows the LCS to function over a wide range of standoff distances.

During scanning, oscillations in the shuttle arm cause the LCS position to change slightly, relative to the shuttle surface. These oscillations, which result from start-up transients at the beginning of arm motion or from thruster firings from the shuttle itself, can pose a challenge to scanning. The time to scan an individual line is short with respect to the period of oscillation, which means the point spacing along a line is undisturbed. The impact of oscillations perpendicular to the slow-scanning axis can be minimized by analyzing the resulting range image line-by-line.

Overcoming these obstacles made the LCS and its role on STS-114 very successful. It was also a pivotal point – as continued work with 3D data allowed engineers to streamline data processes and ultimately utilize this technology for a variety of purposes on Earth.

LCS technology for military applications

3D data typically had limited applications as it focused on visualization. The inclination has been to treat 3D data like 2D: developing the data into a model and then transforming it into 2D format for final delivery. The primary challenge for using 3D data in this way – for any real-time application – has been the time it took to both process and view the data, as well as the limited information the end result would provide. 3D data, however, contains proportionately more information than 2D data. Intelligent processing – that is processing the 3D data to extract relevant information on the sensor end of the system – allows sensors to work more autonomously. This addresses two critical challenges when working with 3D data: it greatly reduces the required communication bandwidth and the time to make critical decisions.

In ISR applications, new systems developed with the LCS technology provide the ability to recognize, inspect, and track objects, rather than just perceive a threat or change in an environment. The ability to collect this information from mobile platforms is essential; even more crucial is the ability to extract and deliver the best information while using as little data as possible. Because of their size, mobile platforms such as Unmanned Aerial Vehicles (UAVs), typically have limited bandwidth. Using autonomous sensors capable of intelligent processing offers an effective solution. Embedding image processing and analysis tools on the sensor itself allow the sensor to sift through the data and extract only the relevant information. The sensor will then send only this information, rather than an abundance of 3D data that has yet to be processed.

The 3D tracking capability of the LCS was the pivotal feature for new scanners designed for Earth applications, particularly in the defense and military industries (Figure 4). Designed to increase the effectiveness of existing sensor platforms, these autonomous 3D sensors are capable of supporting a variety of applications such as automated target recognition, pose estimation, change detection, and real-time object tracking.

The 3D Automatic Target Recognition (ATR) and Pose Estimation System was originally developed with Canada's Department of National Defence. Designed using two algorithms in a multi-step process, the recognition approach is unique. These algorithms provide uncorrelated metrics and, consequently, different characteristics of the target. This allows the potential target dataset to be reduced prior to final selection. In the preprocessing phase, the data from the scanned object is segmented from its surroundings and is reprojected onto an orthogonal grid, making the shape independent of range. A fast recognition algorithm accesses a 3D target model database and reduces the list of potential targets by removing unlikely cases. The second algorithm or the ATR algorithm's primary purpose is to ascertain what the object in question is. After passing a number of consistency checks that include a series of scans from simulated angles, the metrics from both algorithms are combined to provide relative probabilities for each database object. In a field test, the system recognized 9 out of 10 possible targets within a 400-meter range.

As stated, the ATR algorithm will perform a series of simulated scans on the object in question. This not only increases the recognition factor, but it also provides 3D object tracking and pose estimation, providing vital information about orientation, position, and even speed. Sensitive to occlusions, the

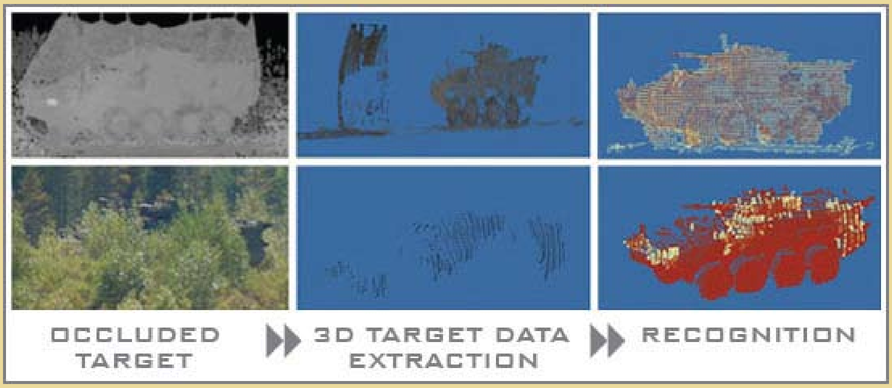

Figure 4

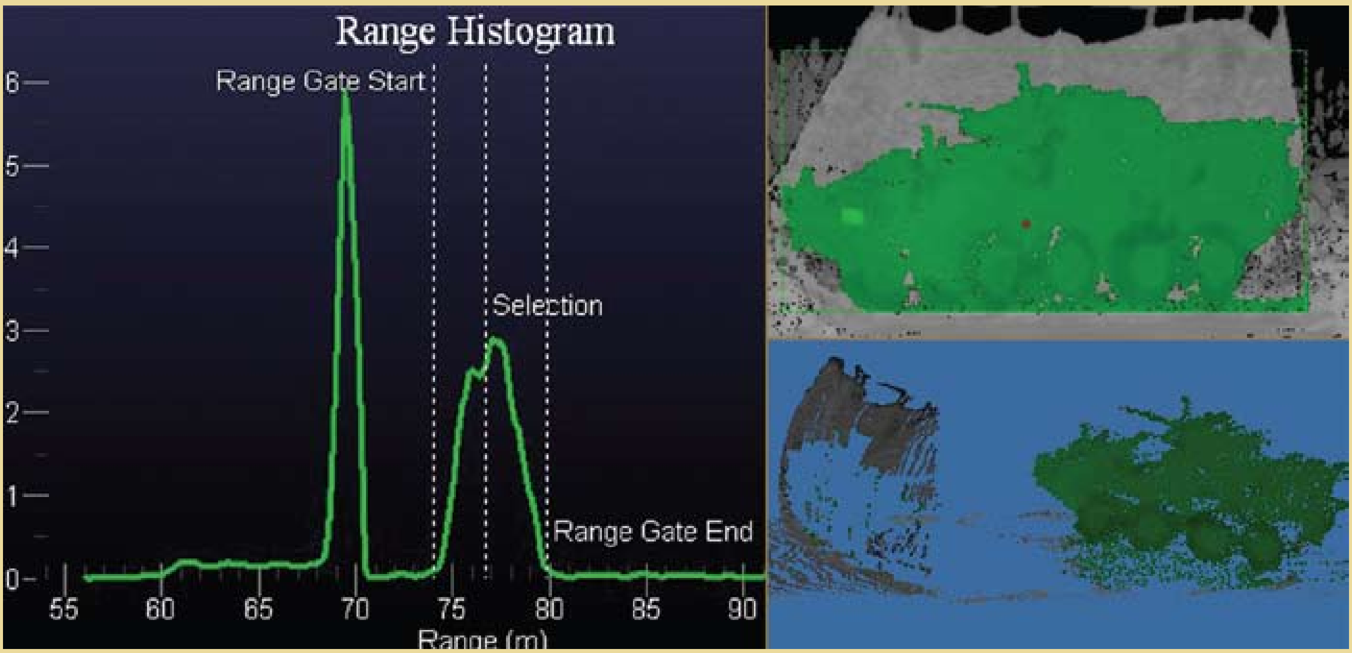

Figure 5

application needs few points to determine objects. For example, a vehicle occluded by leaves and shrubs can be identified, as well as its direction and speed when only a small portion of it is visible to the sensor (Figure 5). A further-refined ICP algorithm allowed for a full 3D pattern, matching rates in excess of 5 Hz on a standard Pentium-based laptop computer.

A change detection system designed for use on a mobile platform for ground and perimeter security applications has also been developed based on the LCS and ATR technology. Processing 3D imagery data to identify anomalies, this technology is fast and is only limited by sensor resolution. This system can sense changes as small as a few cubic centimeters at a range of a couple of meters to vehicle-sized objects at a range of a kilometer. Unlike many pixel-based change detection systems, this technology is robust to dynamic lighting conditions and changing perspectives. Maintaining its ability to provide a full 3D model, the sensor remains autonomous, delivering relevant information in real time.

Fundamental technology

The ability to visually comprehend and recognize objects, rather than to just perceive a threat or change in an environment, is fundamental. Collecting this information from mobile or stationary platforms is essential. Laser technology, when used intuitively, can be leveraged in extreme environments over low-bandwidth networks. The LCS, operating in extreme conditions, collects critical data for NASA to ensure that its astronauts return home safely.

Applying LCS core principles and technologies to terrestrial applications means that the defense and military industry can also take advantage of the many benefits that accurate 3D information provides. Now those making critical decisions don't just have information. They have the means to collect the right information.

Maureen Campbell *is a technical marketing specialist at Neptec Design Group. With more than 10 years of experience in the technology industry, she works with the research and development team at Neptec. Maureen has been with the company for two years. She can be contacted at mcampbell@neptec.com*

Neptec
613-599-7602
www.neptec.com

Product Guide: Radio front-ends, tuners, digitizers, and tool sets

Company	Model	Description	Form factor	Number of channels	A/D bit resolution	Sample rate	FPGA onboard	IP/Libraries
Aitech Defense Systems www.rugged.com	S930 I/O Card	A rad-tolerant analog I/O card	3U CompactPCI	64	12-bit			
GE Fanuc Intelligent Platforms Embedded Systems www.gefanucembedded.com	IMP1A-571 Blade	A single-slot software radio blade for space- and weight-constrained applications	3U CompactPCI				Yes	
North Atlantic Industries www.naii.com	75C2	Improved single-slot, multifunction card with synchro measurement	3U CompactPCI					
RT Logic www.rtlogic.com	DG4000 Converter	A high-rate converter	3U CompactPCI			ADC @ 213 MSps, DAC @ 500 MSps	Yes	
General Standards www.generalstandards.com	cPCI6U-24DSI32R	32-channel 24-bit delta-sigma analog input board with rear I/O	6U CompactPCI	32	24-bit	2 KSps to 200 KSps		
Pentek www.pentek.com	7340-420	Multiband transceiver	6U CompactPCI		14-bit		Yes	
CAEN Technologies www.caentechnologies.com	V1724	A module housing a flash ADC waveform digitizer with threshold auto-trigger capabilities	6U VME	8	14-bit	100 MSps		Yes
TYCO Electronics www.tycoelectronics.com	M/A-COM TU-6401	A lightweight low-power SIGINT microwave tuner designed for signal search and collection applications	6U VME					
Picor www.picorpower.com	QPI-8LZ	System-in-a-Package (SiP) designed to integrate hot-swap function with an active EMI filter	AdvancedTCA					
Red Rapids www.redrapids.com	Channel Accelerator Plus 8/1500	FPGA configurable, dual-channel receiver	CompactPCI, PMC, PCI	2	8-bit	1500 MSps	Yes	Yes
Red Rapids www.redrapids.com	PMC/PCI Dual Channel Digitizer	A dual-channel digitizer	CompactPCI, PMC, PCI	2	16-bit	130 MSps	Yes	Yes
Agilent Technologies www.agilent.com	10-bit Digitizers	Digitizers designed for optimizing high-speed ADC performance	CompactPCI, PXI					
KineticSystems www.kscorp.com	CP213	A module with 32 or 64 differential analog input channels that can be configured as 64 or 128 single-ended analog input channels	CompactPCI, PXI	64/ 128	16-bit			
ACCES I/O Products www.accesio.com	LVDT-8A	An LVDT signal conditioner board	Custom module	8				
Annapolis Micro Systems www.annapmicro.com	Quad 16 Bit I/O	An A/D I/O card	Custom module		16-bit		Yes	Yes
CAEN Technologies www.caentechnologies.com	V1720	A module housing a flash ADC waveform digitizer with threshold auto-trigger capabilities	Custom module	8	12-bit	250 MSps		Yes
Crystek Corporation www.crystek.com	CVCO55CC-2580-2860	A high-frequency VCO for use in digital radio equipment, satellite communications systems, and base stations	Custom module					
EZ-EMBEDDED www.ezembedded.com	MCK2812 DSP EVM	A DSP board (TMS320F2812)	Custom module	16	12-bit			
Lyrtech www.lyrtech.com	Quad Dual-Band RF	An RF transceiver	Custom module	4			Yes	
Millennial Net www.millennialnet.com	MeshScape EN5424	Compact, low-power embedded device for military use and to connect sensors to a wireless sensor network	Custom module	16				
Altera Corporation www.altera.com	Arria GX FPGA Family	A transceiver-based FPGA ranging from 21,580 to 90,220 logic elements	Integrated Circuit					
IntegrIT www.integrit.ru	NatureDSP IT2400	1200/2400 bps software modem for data and digital voice transmission over HF channels using commercial radios	N/A					Yes
IntegrIT www.integrit.ru	NatureDSP Squelch+	A software package providing subaudio signaling functions for FRS, PMR446, and analog radios	N/A					Yes

Company	Model	Description	Form factor	Number of channels	A/D bit resolution	Sample rate	FPGA onboard	IP/Libraries
PrismTech www.prismtechnologies.com	Spectra	Tools for SDR application development	N/A					Yes
QuickFilter Technologies www.quickfiltertech.com	QF4A512	A mixed signal FIR filter designed to function between an analog input and an MCU or FPGA	N/A		16-bit			Yes
Agilent-Acqiris www.acqiris.com	DP1400	High-speed digitizer	PCI	2	8-bit			Yes
BittWare www.bittware.com	Tetra-PMC+	High-speed A/D I/O card TigerSHARC link ports	PCI	4	14-bit		Yes	
Spectrum GmbH www.spectrum-instrumentation.com	M2i.46xx Series	A/D converter PCI board with up to 8 synchronous channels with separate A/D	PCI	8	16-bit	3 MSps		
ADLINK Technology www.adlinktech.com	DAQe-2200	Data acquisition cards	PCI Express	64	12- or 16-bit	500 KSps-3 MSps		
Strategic Test www.strategic-test.com	UF2e-2020	Two-channel 8-bit 50 MSps digitizer/oscilloscope PCI-Express x1 board	PCI Express	2	8-bit	50 MSps		Yes
Strategic Test www.strategic-test.com	UF2e-4640	Waveform digitizer/oscilloscope boards	PCI Express	4	16-bit	3 MSps		Yes
Acromag www.acromag.com	PMC-VLX/VSX	PMC module with user-configurable Virtex-5 FPGA and plug-in analog and digital I/O cards	PMC				Yes	
GE Fanuc Intelligent Platforms Embedded Systems www.gefanucembedded.com	TI C64 DSP Blade	Software radio TI C64 DSP blade	PMC				Yes	
Micro Memory www.micromemory.com	MM-7105, 7110, 7115	The MM-71xx series of compute nodes providing FPGA processing resources	PMC				Yes	Yes
VMETRO www.vmetro.com	AD3000 3GSPS ADC	A single-channel 3 GSps ADC utilizing National Semiconductor ADC083000 8-bit converters	PMC	1		3 GSps	Yes	
4DSP www.4dsp.com	AD450	A digitizer and waveform generator for MIMO systems	PMC, XMC	2	14- or 16-bit		Yes	
4DSP www.4dsp.com	AD491	A digitizer featuring two Virtex-4 devices	PMC, XMC	2	8-bit		Yes	Yes
Pentek www.pentek.com	Model 7142-428 Digital Transceiver	Digital transceiver with multiband DDC core and interpolation filter	PMC, XMC		14-bit ADC, 16-bit DAC/DUC	125 MHz, 320 MHz		Yes
Spectrum Signal Processing www.spectrumsignal.com	901-00001	An RF-to-Ethernet platform that integrates a cross band RF transceiver with a high-performance SDR-based signal processing subsystem	System					
Trimax Wireless www.trimaxwireless.com	Trimax 6000 Series	A six-radio platform built to military standards and supporting WiMAX, Wi-Fi, DECT, and 3G technologies	System					
Mercury Computer Systems www.mc.com	Echotek Series RF 1800GT	An ultra high-performance wideband microwave receiver	VME					
Annapolis Micro Systems www.annapmicro.com	WILDSTAR 4 VXS	An FPGA-based processing board	VME, VXS		8- to 16-bit	600 MSps-2.3 GSps	Yes	Yes
TEK Microsystems www.tekmicro.com	Tarvos VXS	Part of the Quixilica VXS family of FPGA-based processing hardware	VME,VXS	6 ADC, 1 DAC	16-bit	160 MSps, 500 MSps	Yes	
Innovative Integration www.innovative-dsp.com	X5-210M	An I/O module	XMC	4	14-bit	210 MSps	Yes	

Data was extracted from OSP's product database on Nov. 14, 2007. Keywords searched include Radio, Digitizer, Software Radio, Transceiver, Digital Downconverter, Tuner, and A/D products entered 1/23/07 through search date within *Military Embedded Systems* and *CompactPCI and AdvancedTCA Systems* magazines. Entries were chosen at editor's discretion due to space limitations and edited for relevance to product guide's theme. OpenSystems Publishing is not responsible for errors or omissions, and vendors are encouraged to add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/.

New Products

By Sharon Schnakenburg

10 GbE CONNECTIVITY MODULE

AdvancedIO Systems
www.advancedio.com
Model: V1020 XMC / PMC **RSC No: 33361**

10 GbE connectivity module that functions as a high-speed, flexible, open-standards-based data pipe • Using a sockets-based API, the module improves system performance by offloading communication protocol stacks from embedded processors and facilitating direct data streaming between 10 GbE and other input/output devices • Based on reprogrammable FPGA technology to accommodate added packet, data, and signal processing functions such as deep packet inspection, packet filtering, and up/down conversion, as well as protocol customizations • Supports both PCI-X and PCI Express (PCIe) • Suitable for radar, avionics, and a variety of other high-performance applications

BOUNDARY SCAN DEVELOPMENT SYSTEM

XJTAG
www.xjtag.com
Model: XJTAG version 2.0 **RSC No: 34990**

A boundary scan development system used by board developers and manufacturers to debug, test, and program complex BGA-populated printed circuit boards and systems • Automates JTAG chain set-up and enables quick and easy categorization of non-JTAG/cluster devices • New built-in netlist explorer provides a simple interface to view connectivity between devices • Includes improved memory test, real-time DFT test coverage analysis, an extended library of device-centric test scripts, improved integration with LabVIEW, and support for Xilinx's Virtex-5 FPGA System Monitor

COMPARATORS

National Semiconductor
www.national.com
Model: LMP7300 **RSC No: 34987**

A precision comparator with 2.048 V reference that provides adjustable hysteresis signal detection in power supply and battery monitors, sensor interface, and threshold detectors • Low-offset voltage of 300 µV and low supply current of 10 µA • Wide supply voltage range of 2.7 V to 12 V for use in 3.3 V, 5 V, and ± 5 V applications • A 2.048 V reference voltage with 0.25 percent maximum error provides a precise input monitoring voltage reference level • Propagation delay is less than 5 microseconds to enable fast, precise, low-power signal detection

CONNECTORS: AdvancedTCA

Harting Technology Group
www.harting.com
Model: ATCA PowerConnector **RSC No: 34949**

Advanced power connectors that offer robust and reliable performance with press-fit technology to acheive a reliable connection • Designed according to the PICMG 3.0 specification • 8 power contacts (AWG size 16) and 22 signal contacts (AWG size 22) • Power contacts rated up to 16 A • Signal contacts rated up to 1 A signal contacts • Guiding offered for blind mate capability • Assembly with standard flat rock die

CONNECTORS: VME 64x

Hypertronics
www.hypertronics.com
Model: VME64x connectors **RSC No: 35134**

VME64x connectors including a shield that prevents EMI/RFI • Connectors are mechanically compliant with IEEE-1101.2 -1992, ruggedization level 5 • Stackable design of high-speed modules features round pins to mate with Hypertac contacts • Optimized lead traces within modules provide superior performance in high-speed applications • Aluminum frames for ruggedness and conduction cooling • Keying feature assures proper mating

DUAL-CHANNEL RADAR INTERFACE

Curtiss-Wright
www.cwcembedded.com
Model: Osiris Radar Input **RSC No: 34877**

A high-performance, dual-channel radar interface board that accepts and processes analog and digital radar signals • Provides a PCI interface to applications • Dual-channel radar interface card • PMC or half-length PCI form factors • Two analog inputs • Eight digital input bits • Dual trigger inputs • Dual azimuth inputs (ACP/ARP, RADDS) • Sampling up to 50 MHz

DUAL PowerPC VXS SBC

Mercury Computer Systems, Inc.
www.mc.com
Model: MomentumVPA-200-SBC **RSC No: 34950**

A dual 7448 PowerPC VXS SBC • Full-featured 6U VME board offering a higher level of computing power and bandwidth for users with high-density needs • VITA 41 performance and flexibility with high-speed serial fabric • Two standard PMC-X sites for off-the-shelf PCI mezzanine cards • Tundra Tsi148 VME bridge chip provides VME 2eSST support • Access to wide ecosystem with Linux or VxWorks operating systems • RapidIO on VSX PO option

HIGH-SPEED TRANSIENT RECORDER

AMO GmbH
www.amo.de/dcc.html
Model: Saturn System **RSC No: 34750**

High-speed transient recorder with >1,000 channels and fiber optic isolation up to 1 Mill. Volt • Portable mini mainframe for service • Up to ±100 V input range • Multiple electrical trigger options • Powerful real-time signal analysis for complex trigger setups • Enclosures for all purposes • UPS integrated • Large acquisition memory up to 1 GB per channel • Multiple import and export formats • Multi-monitor software environment to keep overview with even largest setups • Fully automatic acquisition with analysis and reporting • Complete API to remote control Saturn System via Ethernet

LINE CONDITIONING MODULE

VPT Inc.
www.vpt-inc.com
Model: VPTPCM-1 **RSC No: 34709**

A line conditioning module that protects power systems from transient voltages to ensure smooth, consistent power through harsh conditions • Designed for rugged military and avionics applications • Continuous operation with input down to 6 V per MIL-STD-1275 • A wide input range provides dual nominal input voltages (12 V and 28 V) to satisfy both MIL-STD-1275 and MIL-STD-704 requirements • Transient suppression up to 600 VDC • Up to 120 W output power from a single unit • Operation over a wide temperature range of -55 °C to +100 °C with no power derating • Six-sided metal case intended for rugged environments • High efficiency, up to 99% • Space-saving size – just 2.350" x 1.550" x 0.465", 90 g

PORTABLE DATA RECORDER

DSPCon, Inc.
www.dspcon.com
Model: DataFlex-500 **RSC No: 34716**

Portable, high-speed data acquisition recorder with standard real-time analysis software • Records data while performing extensive displays and real-time data tests at bandwidths up to 115 KHz • Sustained, continuous recording speed of up to 256,000 samples/second • High-capacity recording – standard 750 GB SATA drive • Remote monitoring and control via dual-port Ethernet • Supports a wide variety of data extraction options, including Ethernet, DVD, or USB removable media (NTFS format) • Supports a wide variety of signal conditioning

PowerPC EMBEDDED PROCESSOR

AMCC
www.amcc.com
Model: PowerPC 460EX **RSC No: 34948**

A PowerPC-based embedded processor • Power Architecture core delivers a frequency ranging from 667 MHz to 1.2 GHz with 2.0 Dhrystone MIPS/MHz • Memory support includes 32 KB instruction and data caches, an on-chip 256 KB L2 cache, and an additional 64 KB of on-chip memory mapped SRAM • Supports up to four banks of DDR2 SDRAM memory and a maximum capacity of 16 GB • Two independent PCI Express interfaces and a 32-bit PCI v2.3 interface • Two integrated 10/100/1000 Ethernet ports with TCP/IP acceleration, QoS, and Jumbo Frame support • Features a best-in-class Turbo Security Engine that supports standard security protocols such as IPSec, SSL, and DTLS • Compliant with FIPS-140-2 and ANSI X9.17 Annex C • Typical power dissipation is an estimated 5 W @ 1 GHz for the entire chip

RACK-MOUNT CompactPCI SYSTEMS

Comtel Electronics GmbH
www.comtel-online.de
Model: 1U CPCI Systems **RSC No: 34901**

1U, 19" rack-mount system with hot removable CompactPCI 3U power supply, four fans, and CompactPCI backplane • Supports 6U cards front and rear • Horizontally mounted 6U 160 mm boards and rear I/O 80 mm • CompactPCI backplane, two slots, 64 bit, 66 MHz • EMC and ESD optimized steel construction • IEEE front and rear card guides • Hot swap CompactPCI power supplies AC wide range or 48 VDC • Active heat dissipation of the boards via convection from side to side • Variable cooling solution by using up to eight fans • Four fans 40 x 40 mm with fan fail LED on front • Optional additional three fans for push-pull cooling • One optional additional fan for push-pull cooling for RTM • Customer-specific assembly version on request

RUGGED COTS CHASSIS

Schroff, a brand of Pentair Electronic Packaging
www.schroff.us
Model: COTS RGD-4 Chassis **RSC No: 34812**

Chassis designed for rugged applications with weight restrictions and less stringent shock and vibration requirements • Provides a very stable platform for rack-mount VME applications • Complete EMC shielding and filtering • Operating temp. 0 °C to +65 °C (non-operating -20 °C to +85 °C) • Operating altitude: 18,000 ft • 20-slot VME backplane housed in a 6U x 160 mm deep card cage • 750 W power supply • Efficient, negative-pressure, filtered, convection cooling system • Switch guard for power cycle and system reset • Shock: 15 *gs* • Vibration: 2.5 *gs* RMS 15 to 2,000 Hz • Acceleration: 3 *gs*

RUGGED VMEBUS SBC

Dynatem, Inc.
www.dynatem.com
Model: DPD **RSC No: 34979**

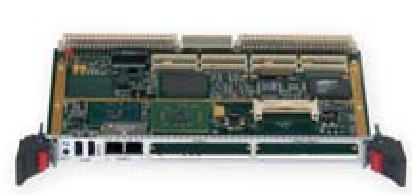

A rugged, low-power VMEbus SBC • Intel low-power Core-Duo (Yonah) processor @ 1.66 GHz • E7520 chipset for PCIe support and high memory bandwidth • Onboard SVGA controller • Two Gb LAN front panel ports plus two more routed to the backplane in compliance with VITA 31.1 • Supports two PMC sites, one of which optionally supports XMC modules • Available in conduction-cooled versions for rugged applications

For more information, enter the product's RSC# at www.mil-embedded.com/rsc

Pack four (or eight!) per suitcase

The race for more cores is on. AMD and Intel are leap-*flogging* each other with more cores, then more cores per watt, then . . . well, wait and see. Rugged suitcase computer provider NextComputing has jumped into this foray with AMD's native quad-core Opteron. Operating at 55 W Average Computing Power (ACP), AMD's 2347 HE will be the CPU of choice for all of NextComputing's flextop suitcases, including: NextDimension Pro and Pro HD, Evo and Evo HD (think it'll smoke the tires?), and the ruggedized Vigor (is it just me, or do these all sound like models from *The Fast and the Furious?*). Since too much is never enough, each suitcase can be equipped with up to two quad-core CPUs – wonder if eight is enough?

The modular suitcases are designed for performance-hungry tasks in military/intelligence, simulation, or other C4ISR applications where portability, modularity, and multiple I/O options are paramount. The AMD Opteron has a floating point unit for each core and the ability to power-down major functional blocks to save juice. On the suitcase side, models such as the Evo HD can include up to seventeen 2.5" SATA drives, up to 16 GB DDRAM, internal DVD R/RW, up to four PCIe/PCIx64 expansion slots, boatloads of I/O, and even up to 1 TB of internal RAID. Multiple 10/100/1 GbE/10 GbE and monitors can also be accommodated.

NextComputing • www.nextcomputing.com
RSC# 35180

Deployable "mini mainframe"

When it comes to field-deployable computers, you've usually got two choices: sealed, single-purpose boxes or huge rack-mounted assemblies that require a trailer wrapped around them. Or, there's the Multi Processor Unit (MPU) COMx from Z Microsystems. This open architecture modular unit allows up to seven hot-swappable processor modules with either AMD or Intel CPUs – and you can mix and match as desired. Each module can accommodate two sealed TP2 disk drives and up to 4 GB of RAM.

Based on PICMG's COM Express standard of interoperable modules, the box can house dual hot-swappable AC or DC power supplies. One can even add optional PCI Express modules with enhanced graphic capabilities to take advantage of Z Micro's C4ISR GUI software. The 65-pound box can operate over 0 to 50 °C, boasts a 10-year operating life, and even allows for field replaceable fans – an ideal feature in fine-sand desert environments. Networking and remote management are also possible.

Z Microsystems • www.zmicro.com
RSC# 35181

Image processor reduces operator overload

Human factors studies show that military operators can only handle so much information before becoming overwhelmed. The situation is exacerbated when under stress, despite best-in-class training. Instead, what's needed is to use more of the intelligence embedded in intelligent machines such as image processors. That's the intention behind GE Fanuc Intelligent Platforms' ADEPT5000 image processing subsystem. Designed to operate with a wide range of physical sensors, the unit clearly is ideal for rugged and mobile deployed applications such as in-vehicle or front-line TOCs.

Through algorithmic *data extraction processes,* the unit aggregates sensor inputs to enhance operator decision making. For instance, ADEPT5000 can use scene change detection, moving target detection, or contrast detection – all of which can alert an operator to take action instead of expecting the operator to blindly stare at an unchanging screen for hours, often missing subtle action due to sheer console fatigue. Flashing graphic symbols are used to alert operators, and picture-in-picture enhancement, recorded snapshot action, and audible alarms all seek to *enhance* the operator's decision making. And since the unit aggregates multiple sensors, it is truly a force multiplier beyond what a single sensor might normally provide.

GE Fanuc Intelligent Platforms (formerly GE Fanuc Embedded Systems) • www.gefanucembedded.com
RSC# 35182

PCI Express 2.0 switches do double time

Many new COTS ICs – processors, peripherals, and interfaces – include native PCI Express (PCIe) pinouts. That's all well and good until your server or embedded system needs to bridge between onboard PCIe or route lanes between boards. Then the reality hits you: A switch is needed. IDT's PCIe Gen2 switches, announced earlier this year, simplified the job dramatically. But in high-density server applications or in sensor-critical embedded installations, the company realized that "burst mode" operation might be handy.

Offering the ability to *double* PCIe's rate from 2.5 Gbps to 5.0 Gbps for cutting-edge designs will move data around a system at "double time" (to use the military cadence vernacular), which also has the nifty benefit of allowing a 50 percent reduction in the number of PCIe lanes and board traces necessary in cost-sensitive designs. For military systems, this translates to lower PWB costs and higher reliability, especially in conduction-cooled boards. IDT also claims industry-leading performance-per-watt and "lowest total power consumption" – we'll have to take their word for that for now. Four versions are available: a 24-lane, 6-port; a 24-lane, 3-port; a 6-lane, 6-port; and a 4-lane, 4-port.

Integrated Device Technology • www.idt.com
RSC# 35183

Quad SFPDP XMC module

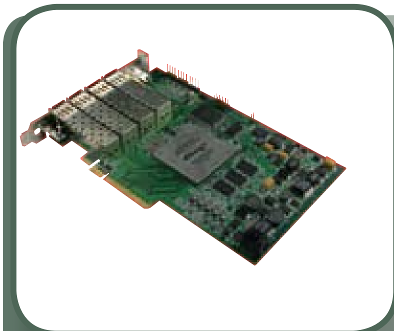

Many defense systems rely on FPDP – or the serial version called *SFPDP* – to transmit low-latency data from sensor to processor. Curtiss-Wright Controls Embedded Computing's (CWCEC's) FibreXtreme SL100/SL240 mezzanine modules boast four SFPDP channels and are available in PMC (PCI) or XMC (PCIe) flavors. There are two versions of the boards: One supports 1 Gbps data rates and the other 2.5 Gbps. With the tap wide open, each channel can sustain 247 MBps to SFPDP distances up to a staggering 50 km.

Using an Altera Stratix II GX FPGA for line processing, the FPGA's transceivers can scale up to 6 Gbps for lots of future headroom. In addition, embedded DMA controllers offload the host CPU and can blast up to 6 MB of data before toggling a CPU interrupt. Onboard data processing also includes register byte/word swapping (Big Endian, anyone?) plus myriad other preprocessing functions that designers might pour into the FPGA. The VITA 17.1-2003 link communications protocol is included, and extended temperature and conduction-cooled versions are also available.

Curtiss-Wright Controls Embedded Computing • www.curtisswright.com
RSC# 35184

Front-line dual Xeon embedded server

Known for its PICMG-compliant, Intel-powered rack-mount servers, RadiSys is gearing up to bring its servers a bit closer to the front lines. The Procelerant RMS420-5000XI (rolls right off the tongue) packs twin Intel Xeon quad-core CPUs (E5335 @ 2.0 GHz or E5345 @ 2.33 GHz), a 1.3 GHz FSB, and eight memory slots (up to 32 GB) in a svelte air-cooled chassis measuring 7" x 17" x 19.6" (H x W x D). Weighing only 44 pounds, the box is ideal for environmentally controlled installations such as trucks, command shelters, or quick-deployed reach-back installations.

The unit runs off of 100 to 240 VAC (50/60 Hz), and includes five USB 2.0, VGA/DVI, PS/2, AC 97 audio, and two GbE ports. Of course, this is a server, not a PC so there are four drive bays, three PCIe slots (1 x16, 1 x8, and 1 x4), plus two PCI-X 100 slots. RadiSys offers suggestions for dual graphics cards to turn the box into an image processing unit. The server is pretty much a commercial temp device that operates over +5 °C to +45 °C, but the company provides altitude derating criteria – suggesting that operation in wide-bodied aircraft such as AWACS, E2C, JSTARS, or Rivet Joint was a design criteria. Finally, the Procelerant can withstand nonoperating temperatures of -30 °C to +70 °C, with a 5 °C per minute maximum excursion gradient. We think this is just RadiSys's prelude of a family of ruggedized, embedded servers.

RadiSys • www.radisys.com
RSC# 35185

Can VME vendors build a better shoebox?

It's irrefutable that Motorola – ostensibly the inventor of VME – and backplane/chassis expert Hybricon know how to build boards and boxes used by the military. So it is entirely noteworthy that for the past six months, the companies have been trotting around a *Rugged MicroTCA chassis* that's designed to cocoon AdvancedMCs used for "military edge applications." Since an AMC is a tad longer than a PMC, sticking a bunch of them in a rugged enclosure qualifies for our definition of a "shoebox." Previously, the military relied on shoeboxes stuffed with proprietary boards, PC/104 boards, or 3U CompactPCI cards. The significance of this proof-of-concept chassis is that, in principle at least, the dozens of AMC cards designed to be interoperable under PICMG standards can now be *cheaply* used in rugged defense systems.

Specs on the proof-of-concept are sketchy, but we wholeheartedly endorse the concept and have awarded it an Editor's Choice mention. Here's what we can piece together for now: Air-cooled AMCs can be installed as-is, while a conduction-cooled MicroTCA chassis would require some mechanical appendages (for example, wedgelocks) affixed to unmodified AMC PWBs. Ruggedization levels are administered by PICMG's Rugged MicroTCA subcommittee and follow ANSI/VITA-47 levels (the irony!). The air-cooled proof-of-concept holds five AMC cards, blows over 32 cfm of air per slot, has room for various MIL-SPEC-style PSUs, and uses 38999-style front panel connectors mated to existing AMC front-panel connections. The chassis can be cocooned for enhanced shock and vibration, and both companies are working on a conduction-cooled frame + AMC + cold plate-to-air chassis for the ultimate in ruggedization. More info and drawings can be found at www.compactpci-systems.com/articles/id/?2231.

Hybricon • www.hybricon.com
Motorola Embedded Communications Computing (now "Emerson")
www.motorola.com

Editor's Choice Products are drawn from OSP's product database and press releases. Vendors may add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/ and submit press releases at www.opensystems-publishing.com/news/submit. OSP reserves the right to publish products based on editors' discretion alone and does not guarantee publication of any product entries.

OpenSystems Publishing

Advertising/Business Office

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

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Business Manager
Karen Layman

Sales Group

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

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

Doug Cordier
Account Manager
dcordier@opensystems-publishing.com

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

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

Regional Sales Managers

Jerry Bleich
New England
jbleich@opensystems-publishing.com

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

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

Denis Seger
Southern California
dseger@opensystems-publishing.com

Sydele Starr
Northern California
sydele@pacbell.net

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

International Sales

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

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

Reprints and PDFs

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

Advertiser Information

MILCOM exposé

Sun, Disney, Orlando play host to military communications conference demanding connectivity and interoperability

By Chris A. Ciufo, Editor

The "free world's premier military communications conference" was held Oct. 29-31 in Orlando, Florida this year. The 26th annual MILCOM event was cosponsored by AFCEA International and the IEEE Communications Society – the same partners who have successfully held this event since its inception in 1981, a few years before my own first attendance. With approximately 6,000 attendees[1], 297 sold-out exhibitor spots, more than 600 technical sessions, and parking so tight that the Gaylord Palms Convention Center had to bus attendees from neighboring hotel parking lots, it's fair to say the event was a rip-roaring success. As for me, except for Orlando's all-week rain reminiscent of my Pacific NW home (so much for Florida "pool time"), this was the most informative conference I've attended all year, military or civilian.

Looking past the aforementioned parking issue, the venue was superb. Held in a private resort complex, the Las Vegas-like hotel and adjoining convention center featured a huge 10-plus story atrium mimicking Minnesota's Mall of America – complete with shops, cafés, and private press meeting rooms where I kicked off the week at 0700 with an exclusive interview with Roger Krone of Boeing and eight of his top executives.

Krone, President of Boeing Integrated Defense Systems (IDS)[2], was very gracious to this tech-only journalist, and surprisingly knowledgeable about technology. Also, he consented to discuss in great detail some of the nuances of his top programs from JTRS, to Future Combat Systems, and GPS III. I'll publish excerpts of the interview in a future issue of *Military Embedded Systems.*

You might be asking: What's the one thing he said that surprised me the most? Our nation's urgent need for more space-based communications platforms, networks, and Global Information Grid (GIG) *ad hoc* mesh peer-to-peer platforms. My obsolete notion that rad-hard semiconductors are at least one generation out-of-date with the best-of-breed commercial is inaccurate. Boeing has been "secretly" partnering with IBM Microelectronics for years and deploying deep submicron, leading-edge processors, fabrics, and radio technology for space and LEO platforms. The on-board processing, routing, and interoperability capabilities in Boeing's birds are staggering.

It's therefore fitting that *Interoperability: Policy to Performance* was MILCOM's central theme. From the opening keynote by M.G. Jeff Sorenson (USA), Special Assistant to the Secretary of the Army, to the NSA's closing speech entitled "Shift Happens,"[3] the need for connectivity and *more bandwidth* was oft repeated. For instance, pre-9/11 the U.S. CENTOM reachback bandwidth was fixed at roughly 46 Mbps (keep in mind that today's SINCGARS radios are limited to paltry kilobits per second data rates). Today, we've widened the backbone pipe to more than 7.7 Gbps with less reliance on civilian assets that can be compromised, and a network that's almost exclusively based upon IP packets.

Interestingly – and privately – I couldn't get anyone to salute IPv6, mostly because of the lack of COTS support for it. This underscores the tenuous balance between what the DoD wants and can self-fund, versus how much it must rely on the civilian COTS marketplace. When the world ultimately runs out of IP addresses or demands better packet-in-flight security, the Internet will switch to IPv6, finally realizing the DoD's "mandate."

Down on the exhibit floor, there was the usual mix of humongous 40 x 40 booths by Tier 1 primes Boeing, Lockheed Martin, Northrop Grumman, and others; but increasingly the Tier 2 guys like DRS and Harris are trotting out their COTS-based advanced technology demonstrators. First seen at AUSA earlier this year, the ATDs are looking increasingly compelling where a cauldron of DoD-funded MIL-SPEC hardware such as Harris' Falcon radios are repurposed and integrated with JTRS clusters, trailer-mounted Cisco routers, and Intel 1U servers.

Together, this military/civilian mix is realizing new non-military capabilities. Hurricane Katrina, southern California's recent wildfires, and the Department of Homeland Security's Border Patrol and Customs initiatives may, for example, use UAVs to surveil Arizona's border or hot spots at a wildfire's edge, or deploy military perimeter detection sensors to alert first-responders of trapped victims after a natural disaster. I was extremely impressed that Harris' Falcon Watch RF-5400 Remote Surveillance System does just that. Besides meeting a need, I suspect that primes are consciously lessening their DoD reliance as budgets get tighter.

And finally, down to our own little corner of the embedded world, many familiar companies were present and rolled out some pertinent news of their own. Intel had the *Intel Chopper* in their booth (fresh from a press conference days earlier in Arizona where our very own Assistant Editor Sharon Schnakenburg got to poke and prod its embedded technology); Altera and Xilinx unveiled their increasingly mil-centric FPGA product lines; Nallatech showed off their support for Intel's FPGA coprocessor API initiative; and Motorola/Hybricon showcased a rugged-looking MicroTCA chassis that BAE Systems has allegedly chosen for a WIN-T JC4ISR demonstrator.

The week ended with a press conference by Curtiss-Wright Controls Embedded Computing, GE Fanuc Intelligent Platforms, and TEK Microsystems – announcing that ANSI has ratified VITA's VITA 46 VPX standard. This clears the way for these companies and many others to begin shipping compliant modules to harsh-environment programs waiting for the "certification green light." Boeing's own FCS program earlier this year was rumored to have committed to VPX.

---

[1]Actual figures not available at press time.
[2]$14.3 billion FY06; one of six divisions in the $61 billion company.
[3]The briefing was modeled on – and included – the popular YouTube video of the same name (www.youtube.com/watch?v=ljbI-363A2Q).

ONE OF THE MOST RUGGED PIECES
OF EQUIPMENT IN ITS CLASS.

SO IS THE TANK.

Not all rugged boards are the same. Neither are the companies that make them. At Curtiss-Wright, we're fully focused on the unique and demanding requirements of the rugged defense and aerospace market. That's why our design philosophy encompasses every critical area of rugged product design. Whether it's unique thermal demands, extreme shock and vibration, or even the latest requirements for line-replaceable modules and two-level maintenance, our approach to ruggedization goes above and beyond.

The VPX6-185 single board computer is just one example in our broad range of ruggedized board-level and subsystems products. It's VPX format is expressly designed to bring advanced serial fabric I/O performance to rugged defense and aerospace computing platforms.

CURTISS WRIGHT Controls
Embedded Computing

www.cwcembedded.com

RUGGEDIZATION... ABOVE & BEYOND