RTC
The magazine of record for the embedded computing industry
September 2010 www.rtcmagazine.com

OpenVPX Takes on Tough Industrial Tasks

Upgrade Your Industrial Net to Optical

Getting the Heat Out of Tight Spaces

Robotic Systems: Sense, Think, Act!

An RTC Group Publication

avalue Technology Inc.

www.avalue-tech.com
Dedicated to customer's success

Delivers enhanced performance, energy efficiency, manageability, security functions & smoother visual experiences

The new Intel® Core™ i7 processor, 32nm process technology

Avalue QM57 series verified Intel® AMT6.0 compliant

3.5" SBC

ECM-QM57

Intel® Core™ i7-620LE/ 620UE 3.5" Micro Module with Intel® QM57 Chipset

- Onboard Intel® Core™ i7-620LE/ 620UE Processor
- Intel® QM57 Chipset
- One 204-pin DDR3 SODIMM Up to 4GB DDR3 800/ 1066 SDRAM
- Dual View, 2-CH LVDS, VGA, HDMI
- 5.1-CH Audio, Dual Intel® Gigabit Ethernet
- 2 SATA, 2 COM, 7 USB, 16-bit GPIO
- Support iAMT 6.0

EPIC

EPI-QM57

Intel® Core™ i7-620LE/ 620UE EPIC Module with Intel® QM57 Chipset

- Onboard LV/ ULV Intel® Core™ i7 Processor
- Intel® QM57 Chipset
- One 204-pin SODIMM Up to 4GB DDR3 800/ 1066 SDRAM, Non-ECC
- Dual View, 2-CH LVDS, DVI-I
- Dual Intel® Gigabit Ethernet
- 1 CF, 4 SATA, 2 COM, 8 USB, 16-bit GPIO
- Options: Express Card/ 34mm, Built-in Touch Screen Interface
- Support iAMT 6.0
- Support RAID 0/ 1/ 5/ 10

Avalue Technology Inc.
200 Tornillo Way, Suite 210, Tinton Falls, NJ 07712
Tel: (732) 578-0200 Fax: (732) 578-0250
E-mail: sales@avalue-usa.com www.avalue-tech.com

OpenVPX
Takes on Tough Industrial Tasks

Cover Photo: New Open VPX boards: The Pentek Rugged Quad 200 MHz, 16-bit A/D with Virtex-5 FPGAs 3U VPX Module, the Elma T4340 6U OpenVPX switch and the Extreme Engineering XPedite7370 3U VPX Intel Core i7-based SBC.

CPU Cooler Improves Airflow, Water Block and Heat Dissipation Capacity

Sound and Vibration Software Supports Data Translation IEPE Modules

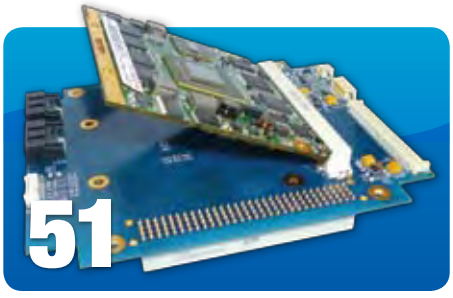

PC/104-Express Board Incorporates Qseven Modules

SEPTEMBER 2010

TABLE OF CONTENTS

RTC

DEPARTMENTS

EDITOR'S REPORT

Robotic Research

TECHNOLOGY IN CONTEXT

OpenVPX

TECHNOLOGY CONNECTED

Options for Industrial Networks

TECHNOLOGY IN SYSTEMS

Thermal Management in Tight Places

TECHNOLOGY DEPLOYED

Robotic Systems

INDUSTRY WATCH

Medical Devices

Rugged SBC Showcase

Featuring the latest in Rugged SBC technologies

ADLGS45PC - Intel Core 2 Duo / Celeron M 1.20GHz - 2.26GHz – PCI/104-Express v1.0

- Intel® Celeron® M / Core™ 2 Duo (SFF)
- Intel® GS45 / ICH9M-E Chipset / DDR3-1066MHz DRAM – Up to 4GB
- LAN Controllers 2x 1Gbit LAN, CRT/LVDS
- 4x SATA 3GB/s with RAID Support
- 8x USB2.0 Ports, 2x COM, LPT, SM-Bus TPM
- For High Performance, Extreme Rugged Applications

Advanced Digital Logic (ADL)
Phone: (858) 490-0597
Fax: (858) 490-0599
E-mail: sales@adl-usa.com
Web: www.adl-usa.com

Echotek® Series DCM-V6-XMC

- High speed transceiver powered by Virtex™-6 FPGA
- Four-channel 16-bit A/D conversion at up to 160 MSPS
- Single-channel 16-bit D/A conversion at up to 1 GHz (optional)
- Complementary to the Echotek Series RF 3000T RF Tuner and Upconverter

Mercury Computer Systems
Phone: (866) 627-6951
Fax: (978) 256-3599
E-mail: sales@mc.com
Web: www.mc.com

USB Embedded Modem Modules

- USB modems, in module or stand-alone form factor
- Linux, Windows and Mac O/S support
- -40C to +85C operating temperature (Module)
- Compact size: 1" x 1" x 0.2" (Module)
- USB 2.0 compatible
- up to 56K bps data rate, fax and voice
- AT command
- Transferable FCC68, CS03, CTR21 telecom certifications
- Global safety: IEC60950-1, IEC60601-1 (Medical) approved
- CE marking

Radicom Research, Inc.
Phone: (408) 383-9006
Fax: (408) 383-9007
E-mail: sales@radi.com
Web: www.radi.com

SEPTEMBER 2010

RTC

Publisher

PRESIDENT
John Reardon, johnr@rtcgroup.com

Editorial

EDITOR-IN-CHIEF
Tom Williams, tomw@rtcgroup.com

CONTRIBUTING EDITORS
Colin McCracken and Paul Rosenfeld

MANAGING EDITOR
Marina Tringali, marinat@rtcgroup.com

COPY EDITOR
Rochelle Cohn

Art/Production

CREATIVE DIRECTOR
Jason Van Dorn, jasonv@rtcgroup.com

ART DIRECTOR
Kirsten Wyatt, kirstenw@rtcgroup.com

GRAPHIC DESIGNER
Christopher Saucier, chriss@rtcgroup.com

GRAPHIC DESIGNER
Maream Milik, mareamm@rtcgroup.com

WEB DEVELOPER
Hari Nayar, harin@rtcgroup.com

Advertising/Web Advertising

WESTERN REGIONAL ADVERTISING MANAGER
Stacy Mannik, stacym@rtcgroup.com
(949) 226-2024

WESTERN REGIONAL ADVERTISING MANAGER
Lauren Trudeau, laurent@rtcgroup.com
(949) 226-2014

EASTERN REGIONAL ADVERTISING MANAGER
Shandi Ricciotti, shandir@rtcgroup.com
(949) 573-7660

ASIAN REGIONAL ADVERTISING MANAGER
Jessica Marinescu, jessicam@rtcgroup.com
(+852) 2548 5100

Billing

Cindy Muir, cmuir@rtcgroup.com
(949) 226-2021

To Contact RTC magazine:

HOME OFFICE
The RTC Group, 905 Calle Amanecer, Suite 250, San Clemente, CA 92673
Phone: (949) 226-2000 Fax: (949) 226-2050, www.rtcgroup.com

Editorial Office
Tom Williams, Editor-in-Chief
245-M Mt. Hermon Rd., PMB#F, Scotts Valley, CA 95066
Phone: (831) 335-1509 Fax: (408) 904-7214

Published by The RTC Group

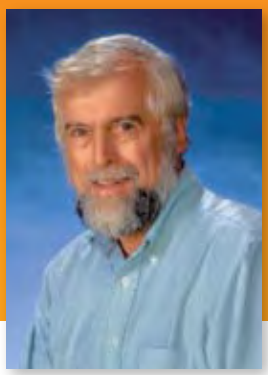

Tom Williams
Editor-in-Chief

Solid State Storage: Will the Enterprise Fuel an Upheaval in the Embedded Space?

An undisclosed number of years ago I was in the office of Alan Shugart, who was at that time the CEO of disk drive manufacturer Seagate Technology here in the throbbing metropolis of Scotts Valley, California. Shugart had previously been one of the pioneers at IBM on the team that developed the very first hard disk drive, a large portrait of which was on the wall of his office.

It was an enormous thing that looked to have been in a cabinet about eight feet long and six feet tall. On one end was a glass panel through which you could see a spindle that held possibly five or six platters. Beside the platters there was another vertical shaft that carried the arm with the read/write head, and attached to that was a rubber air hose to produce an air cushion to keep the head from physically contacting the disk surface. The arm could be withdrawn from one platter and moved up or down and then swung in again to access another platter. The whole monstrous thing had a capacity of five Megabytes—don't ask about the access speed.

Shugart said that at the time the other people in the company thought his team was crazy. "What would anybody do with five Megabytes?" Well now, of course, we know. We carry around iPods that have small rotating media containing multiple Gigabytes. That's if we're old codgers, of course. The newer devices have solid state storage—NAND flash.

Now flash memory is nothing new. It has been around—especially in embedded and mobile devices—for years. But its use has until recently been confined to relatively modest data storage tasks. Recently, however, its use for storing ever more data has been growing, and with the incorporation of things like SATA interfaces on small embedded modules, flash memory has taken on an increasing role as a solid state drive (SSD) for embedded applications. These modules realize they now have a larger capacity, low-power, small, rugged and reliable storage medium that can be used to accommodate those newer applications that suddenly have more data they need to store.

At the recent Flash Memory Summit in Santa Clara, the main concern of the vendors there appeared to be to move flash memory into the turf of enterprise storage. Attendees were flocking like packs of teenage Hannah Montana fans to sessions on performance, benchmarks, storage for enterprise and data centers and more, all lured by the promise of truly vast sales of NAND and controller silicon when the ubiquitous hard drive is pushed further to the sidelines. And truly, there were examples aplenty of SSDs with hundreds of Gigabytes of capacity and impressive performance.

Why should such a development targeted at the enterprise be of interest to the embedded community?" Just remember: "What would anybody do with five Megabytes?" What would an embedded controller or a portable medical device do with 280 Gigabytes? At the moment I'm sure I don't know, but I do know that someone will find a compelling use for such capacity if it fits the size, weight, power, performance, ruggedness and capacity needs of the application. And applications tend to evolve to overtax the capacity of the available hardware.

The conquest of the enterprise space by the flash-based SSD vendors can only be a good thing as the resulting cost reductions and technology improvements proliferate and become attractive to embedded developers. Along with the rush to the enterprise, we are already seeing numerous examples of higher-end flash storage appearing in form factors and with connectors that are clearly aimed at the needs of embedded systems.

And yet, this level of storage is but one element of some rather interesting advances that have yet to come together in actual systems. We are on the threshold of PCI Express 3.0, USB 3.0, faster multicore processors, connectivity such as Intel's LightPeak optical technology and more. And we haven't even mentioned some of the things that are waiting in the wings behind flash, such as phase change memory. Get out the popcorn. It should be quite a show. ◢

INDUSTRY INSIDER

SEPTEMBER 2010

Intel Buys McAfee for Over $7 Billion

In what turns out to be the biggest purchase in its history, Intel has acquired anti-virus software maker McAfee for $7.68 billion. The deal appears to have people scratching their heads speculating about just what it is that Intel has in mind. Of course, there is the revenue stream currently generated by McAfee as the second largest seller of security software in the PC market. But the big question is what the strategic intent may be.

That may include but most definitely is not limited to selling security for PCs based on Intel processors. A clue was given by Intel CEO Paul Otellini, who said, "Everywhere we sell a microprocessor, there's an opportunity for a security software sale to go with it. It's not just the opportunity to co-sell; it's the opportunity to deeply integrate these into the architecture of the products."

From that it would appear that Intel is interested not only in the products and services that McAfee currently offers, but also in the underlying technology that can be integrated with the existing Intel architecture as hardware and/or firmware enhancements. This would take the scope beyond the world of PCs, notebooks and netbooks and into the world of connected devices, which are projected to grow into the billions. Add to that the growing trend toward cloud computing and the fact that McAfee had recently added Cloud Secure to its offerings, and it looks even more attractive. Potentially everything with an IP address represents a potential access point for hackers and malware.

It will be interesting to see how Intel goes about crafting joint products, especially in the embedded space to which it has been devoting more attention than it has traditionally been wont to do. Since the acceptance of the Atom by vast numbers of embedded developers, Intel has used the Atom technology in the new development of the Tunnel Creek device, which it claims for the first time to have conceived specifically for the embedded market. Will we see security chips or processors with built-in hardware support for security? It should be interesting.

SCSI Trade Association Announces MultiLink SAS Connectivity

The SCSI Trade Association (STA) announced the MultiLink SAS initiative at the Flash Memory Summit on August 17. The initiative's purpose is to improve how slot-oriented Solid State Drive (SSD) devices can be configured to improve I/O performance. The externally accessible backplane slot-based drive architecture will be fully compatible for use with existing SAS/SATA storage devices as well as new devices designed to achieve higher performance.

A new form factor compatible connector will extend SAS to a 4-port configuration. When running at 12 Gbit/s, a single slot will be capable of providing up to 96 Gbit/s of bandwidth (full duplex). Additional signals will be provided for general purpose use within the same connector. The MultiLink SAS architecture is a slot-compatible implementation and will accommodate a variety of SSD form factors as well as existing Hard Disk Drives (HDDs).

Minimizing the impact to protocol changes makes MultiLink SAS primarily an enhancement to the existing connector. It was decided to maintain the existing Small Form Factor (SFF) slot dimensions for ease of refitting an existing system and for providing maximum system flexibility for storage OEMs. STA will be working with T10 and the Small Form Factor (SFF) committee to standardize this latest SAS innovation.

Fast/Rugged SDD Technology Poised for Growth

Solid state drives are poised for rapid growth in certain niche markets according to a new report just published by Objective Analysis, *Solid State Disk Market Outlook 2010*. Although SSDs have not found widespread acceptance in general-purpose PCs, those applications that benefit from this technology will drive the client SSD market to grow at a predicted rate of 60%. Objective Analysis predicts that in 2015 nearly 40 million SSDs will ship, accounting for over $7 billion in revenues.

"The PC market for SSDs has been slow to develop," said the report's author, Jim Handy. "The strongest growth has occurred in areas where HDDs simply will not operate and in systems for which users are willing to pay a significant premium for an SSD's faster speed or greater durability."

The report spells out details of the SSD market, its mechanics and anticipated growth. Analysis is based upon numerous interviews with both manufacturers and users of the technology, and explains both what will become of this market and why it will develop the way that it will.

Portwell Sponsors Winning Autonomous Underwater Vehicle Team

The Cornell University Autonomous Underwater Vehicle (CUAUV) team took top honors in the 2010 RoboSub competition using a robotic submarine powered by American Portwell Technology's WADE-8067 Mini-ITX embedded board. The Cornell team builds robotic submarines for both competition and research purposes. They approached American Portwell for sponsorship when they were considering an Intel Core processor to power their AUV because they knew their onboard computer would be subject to severe size constraints and were restricted to a Mini-ITX form factor or smaller. Jack Lam, American Portwell's senior product marketing manager, recommended the WADE-8067 Mini-ITX form factor embedded board, a combination that utilizes the Intel Core 2 processor most effectively.

The Association for Unmanned Vehicle Systems International (AUVSI) Foundation organizes the RoboSub competition, and Cornell's Tachyon AUV beat out 23 other teams from five countries to take first place at the 13th annual competition held from July 13-18, 2010 in San Diego, California. The competition required the autonomous submarine to hit a targeted buoy, send torpedoes into specific windows and drop markers in bins. According to Daryl Davidson, AUVSI Foundation's executive director, the course elements are designed so that at least one or two of the obstacles prove too challenging

for most teams. "However, we are pleased to say that Cornell has now proven us wrong twice," he explains.

Open Screen Project to Deliver Seamless Web Experience Across Connected Devices

Wind River has announced it is building on its collaboration with Adobe by participating in the Open Screen Project and becoming a worldwide scaling partner to bring the Adobe Flash Platform to Internet-enabled devices for rich and engaging Web experiences. The Open Screen Project is an industry wide initiative of more than 70 industry partners led by Adobe to provide consumers consistent Internet and rich media experiences across the broadest possible range of consumer electronics.

Wind River will license, distribute and support Adobe Flash Player 10.x, Adobe AIR 2.x and Flash Lite 4.x across its portfolio of Internet-facing software platforms as well as offer integration, certification and support for these products. Joining the Open Screen Project as a scaling partner for Adobe, Wind River is one of eight global scaling partners entrusted to offer licenses for these products directly to companies worldwide. To kick off this initiative, Flash Player 10.1 and AIR will first be incorporated with Wind River Platform for Android. As an Open Screen Project participant, Wind River will work with customers to integrate Adobe Flash and AIR with their devices, and ensure their devices are compliant with the Open Screen Project certification test suites.

With support for Flash Lite already on products such as Wind River Platform for Android, Wind River is extending its collaboration with Adobe to further pave the way for customers to create products that deliver the full experience of the Internet with Flash Player 10.1, for a variety of market segments. With active participation in the Open Screen Project, Wind River will keep Flash open and updatable in its software stack, as well as offer support customers with upgrade services. By integrating Flash into Wind River's products, device manufacturers can benefit from faster time-to-market and reduced cost and engineering effort. Additionally, application providers can be confident that their Flash-based content and applications will run smoothly

RTEC10

RTEC10 is an index made up of 10 public companies which have revenue that is derived primarily from sales in the embedded sector. The companies are made up of both software and hardware companies being traded on public exchanges. All numbers are reflected in U.S. Dollars. Learn more at **rtcmagazine.com**

	Closing Price	52 Week Low	52 Week High	Market Cap
RTEC10 Index	**46.85**	—	—	**159.65**
Company Market Performance				
Adlink Technology	1.66	1.63	1.69	199.14M
Advantech	2.65	2.55	2.68	1.00M
Elma Electronic	419.35	412.44	419.35	95.70M
Enea	6.88	6.52	6.89	124.35M
Interphase Corporation	1.70	1.64	1.70	11.61M
Kontron	8.67	8.57	8.75	482.66M
Mercury Computer Systems	12.72	12.72	12.90	302.52M
Performance Technologies	2.08	2.07	2.13	23.12M
PLX Technology	3.60	3.52	3.75	133.54M
RadiSys Corporation	9.23	9.21	9.38	222.89M

RTEC10 involves time sensitive information and currency conversions to determine the current value. All values converted to USD. Please note that these values are subject to certain delays and inaccuracies. Do not use for buying or selling of securities.

across devices that use Wind River's Internet-facing software platforms.

CANopen Profiles for Laboratory Automation

CAN in Automation (CiA), the international users' and manufacturers' group for CAN (Controller Area Network), has released additional device interfaces for laboratory automation. The CiA 434 specifications (part 2 and part 3) describe device profiles for heating, cooling and shaking units as well as dispensers, dilutors and pumps.

The profiles specify process data and configuration parameters. This makes the device in a CANopen network interoperable and even partly exchangeable, simplifying system design.

CANopen is an internationally standardized communication system (EN 50325-4), which is used in many different application fields. The CANopen profiles for laboratory automation have been developed by market-leading companies and are designed especially for pipette automation systems, but are also suitable for other laboratory equipment. In such systems, IEC 61131-3 programmable devices are increasingly used, which are originally developed for industrial automation. The CiA 434 specification has been published for CiA internally; excerpts are available for non-members.

AIT, Formerly AIM-USA, Signs Strategic Alliance with TTTech

AIT is pleased to announce the formation of a long-term strategic alliance with TTTech North America to add ARINC664/AFDX, Ethernet, Time-Triggered Ethernet and Time-Triggered Protocol to their product portfolio. AIT will provide local sales and support for North American customers and provide increased engagement in the U.S. government and aerospace market for TTTech. This tightly coupled partnership will capitalize on AIT's experience in product design, sales, production and support, as well as TTTech's leadership in time-triggered communication technologies.

AIT provides a suite of test and simulation products for a wide variety of avionics bus applications, including MIL-STD-1553, ARINC429, ARINC615A, Fibre Channel and MIL-STD-1760E. With support from TTTech, ARINC664/AFDX, Time-Triggered Ethernet (TTE) and Time-Triggered Protocol (TTP) will be added to AIT's portfolio. AIT will remain an independent company and will continue to enhance and supply the newly combined variety of avionics products to aerospace customers.

The TTP ASICs are mature DO-254/DO-178B certified for design of critical embedded flight systems. TTEthernet and TTP are open industry standards (SAE AS6802 and AS6003 respectively) that offer higher bandwidth when compared to CAN, MIL-STD-1553 and ARINC429. They provide significant advantages in terms of reliability, modularity, lower weight, certification, reduced cost and faster time-to-market for aerospace systems.

NI LabVIEW

Limited Only by Your Imagination

RF

Medical

Robotics

Multicore

LabVIEW graphical programming software and modular NI hardware, such as CompactRIO and PXI, are helping engineers develop fully autonomous robotics systems, including unmanned vehicles designed to compete in DARPA Grand Challenge events.

PRODUCT PLATFORM

NI LabVIEW graphical and textual programming

NI CompactRIO embedded control hardware

NI LabVIEW Real-Time Module

NI LabVIEW FPGA Module

>> Find out what else LabVIEW can do at **ni.com/imagine/robotics** **866 337 5041**

GSA Contract Holder
FSS Contract GS-24F-0007K

NATIONAL INSTRUMENTS

SMALL FORM FACTOR **FORUM**

Colin McCracken **&** Paul Rosenfeld

The Emperor's New Clothes

News flash: The latest SFF standard has been approved, and products are rolling off the line soon. Order your sample or dev kit today!

Not so fast. How do you know this new "thing" will stand the test of time? What if the spec was drafted by a lone technical guru, rubber-stamped by only a handful of others, and then launched under the pretense of industry-wide support from a respected trade group? Would that perspective lessen your design-in urge? Do we compliment the Emperor on his or her new clothes, or form an opinion based on our own thorough technical and market evaluation?

To analyze this consider how Corporations win or lose in competitive markets based upon the relevance of their offerings. The market decides. What would happen if your marketing folks defined products without regard for what can actually be built? At the other end of the spectrum, what would happen if your engineers designed products in a vacuum without market input? The days of "build-it-and-they-will-come" are over. Successful new ideas come from an engineering / marketing partnership where available technologies are applied to real customer needs.

Sounds simple. Yet the small form factor community is discovering déjà vu all over again. Processor and chipset vendors have appealing new products, and updated standards are needed to take advantage of the new bells and whistles. Whether board-level or SSD-level, a number of trade groups appear to have created new standards incorporating new features that embedded system OEMs don't need or want simply because a processor or chipset offers such a feature. Can you say, "type 2" or "type 6"? What is going wrong?

The value of a standard derives from the ability of customers to apply a variety of compliant products to solve design, manufacturing and lifecycle management challenges over time. Standards that do not meet the needs of a particular target market likely won't stand the test of time. However, in this market it takes many years to determine success or failure of a standard, especially if the solution looks far forward into the future. Marketing professionals are always refining their view of target customers based upon perceived current requirements (the famous moving target so detested by engineers everywhere) and upcoming challenges to reduce size, weight, power, cost and so on. Trade groups could benefit from the same level of market analysis that system OEMs routinely use.

In creating these new standards or pinout variations, it is very tempting for engineers merely to look at the latest chipsets and map the buses and I/O to off-board connectors. But this tends to disregard the installed base. The entire set of signals can be massive overkill, even for mainstream applications. High-density connectors allow many more pins in the same space. It doesn't mean we need to use them all. Smaller connectors with a well-chosen pinout provide the opportunity to shrink overall system size. Consumer flash modules offer greater data bandwidth at the cost of higher power consumption, but in some cases these modules are literally too hot to handle by embedded OEMs.

A standard specification is the product of the input, creation and review process of a trade group. It has to stand the test of the market, regardless of the marketing hype surrounding its introduction. Results can range from acceptance of a carefully tested simple migratory step to an un-validated misfire, and all shades of gray between. Diversity of thoughts and ideas is critical to the development process. The embedded community must not be afraid to engage in painful debates during the standards creation process about what features must stay and what can go. There just isn't enough space, cost and power for kitchen-sink solutions.

Users of standards-based products are faced with many choices for next-generation designs, upgrades and retrofits. To make truly informed decisions, each OEM must research potential solutions using as many independent sources as practical. Naturally, suppliers will position their products in the most favorable light possible, so one must dig deeper. Don't assume fitness for use given prior successes or reputation of a trade group. Check references and independent articles, evaluate standards against system-level requirements, and then choose wisely. Irrational exuberance doesn't guarantee winning system designs. So don't kiss up to the Emperor. Join the debate and be prepared to share your requirements rather than gush about the latest spec or pinout type. Get involved!

As usual, comments about this topic can be mailed to sf3@rtcgroup.com.

Portwell

Built Tough for Broader Embedded Applications

PEB-2738

- Intel® Atom™ processor (Z510P, Z510PT, Z520PT or Z530P) with industrial temperature range
- Intel® Embedded Compact Extended Form Factor Intel® ECX Form Factor
- Dual display (LVDS and SDVO)
- Gigabit Ethernet
- Low power, fanless & small footprint

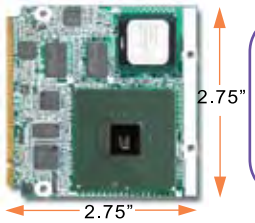

The QSeven standard defines a quadratic, 70mm x 70mm module board size that is 45 percent smaller than the already compact ETX form factor. The QSeven modular board and the carrier board are interconnected through an MXM (Mobile PCI Express Module) edge connector that is widely used in notebook PC manufacturing. Portwell's Qseven solutions will benefit you time to market, space saving, energy saving and cost effectiveness. Typical applications include Automation, DIN Rail Systems, Automotive, Ultra Mobile and Mobile.

ECX — PEB-2771VG2A

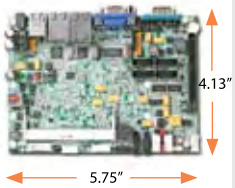

- Intel® Atom™ processor D525
- DDR3 SDRAM up to 4GB
- Dual display (VGA and LVDS)
- PATA, Compact Flash, SATA
- Dual Gigabit Ethernet
- Supports DC 12V input
- Low power, fanless and small footprint

Nano-ITX — WEBS-1012

- Intel® Atom™ processor Z510 or Z530
- DDR2 SDRAM up to 2GB
- Type II Compact Flash & 2.5" SATA
- Gigabit Ethernet
- Supports DC 12V input
- Versatile mounting - wall, panel or DIN
- Low power, fanless and compact

SBC — ROBO-8120VG2R

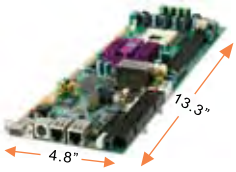

- Dual/quad-core Intel® Xeon® processor
- DDR3 SDRAM up to 32GB
- Dual Gigabit Ethernet
- XGI Z11 graphics controller
- Great I/O expansion capability
- World-class performance

COM Express — PCOM-B216VG

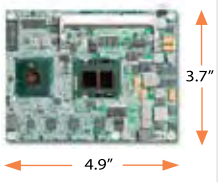

- Intel® Core™ i7/i5/i3 processor
- Mobile Intel® QM57 Express chipset
- Multiple USB ports
- Gigabit Ethernet
- EIDE and SATA
- Up to 8GB DDR3 800/1066 SDRAM
- Intel® Active Management Technology 6.0

Portwell's extensive product portfolio includes single-board computers, embedded computers, specialty computer platforms, rackmount computers, communication appliances, and human-machine interfaces.

We provides both off-the-shelf and versatile custom solutions for applications in the medical equipment, factory automation, retail automation, semiconductor equipment, financial automation, mission critical and network security markets.

American Portwell is both an ISO 9001:2000 and ISO 13485:2003 certified company.

intel Embedded Alliance

Portwell
www.portwell.com
1-877-278-8899

EDITOR'S REPORT

Robotic Research

True Robots Differ Substantially from Other Automated Systems

Yes, it can put those cans on the pallet using its arm and camera like a champ. But can it then run over to the fridge and get me a beer—avoiding my kids' toys in the way? What makes a system really a robot?

by Tom Williams, Editor-in-Chief

The topic of robotics comes with a certain number of preconceived notions. On one end, robots are ambulatory, linguistically endowed anthropomorphic intelligent machines—the stuff of science fiction. On the other end, they are synonymous with most semi-autonomous automated control systems such as those found on the factory floor. In actuality, today's robots are neither of these things. Rather, they are systems at some point in the transition from mundane machine to as far toward the science fiction image as technology and ingenuity can take them. But they are far from that goal despite some fascinating advances.

So how do we differentiate between an automated machine and a robot? According to Siddhartha Srinivasa, Senior Research Scientist for Intel, two things really distinguish robots: the ability to do numerous adaptive general-purpose tasks and the ability to operate in uncertain, unstructured environments. For example, an automated factory machine—and this includes those electromechanical arms that are often referred to as "industrial robots"—works really well in a structured environment such as a factory floor doing one defined task that is defined for it. Those tasks can, of course, be changed by switching out equipment (a welding tip for a paint sprayer) and loading a different program.

Robots, on the other hand, are distinguished by their ability to perform many general-purpose tasks and tasks that may be similar but differ in terms of objects, distances and other variables. For example, the robot that can pick up a cup from a coffee table and hand it to you should be equally capable of moving across the room, picking up a beer mug from a counter top and bringing it back to you without reprogramming. That same robot, in moving across the room, should be able to recognize and avoid obstacles even if they have been recently moved. These two little stipulations bring with them an enormous amount of added complexity, the need for large amounts of computational power and creative developments in machine intelligence. Such machines need to be automatically adaptable both at the task level and at the level of the surrounding environment.

One big issue of trying to write algorithms for robotics, according to Srinivasa, is "to try to write them as general as possible using words that have very general meaning so that at the application level they can be put together in different ways to make different paragraphs, stories and meanings." He calls these "building blocks of autonomy" so that the application developer does not, for example, have to worry about how many degrees of freedom the arm has but can specify instructions to "Pick up an object and put the object there and don't spill the coffee in the object."

Then, of course, there is the question of how one uses such a level of abstraction to instruct the robot to "Pick up the glass." That concept is translated in the human brain from its linguistic generality to very specific arm and hand motions that carry out the task for any number of specific locations and circumstances. By the same token, a robotic system must be able to take a general description of picking up the glass and apply it to many specialized instances.

In the case of the robot used by Intel Labs in Pittsburgh— HERB, the Home Exploring Robotic Butler (Figure 1)—this is done by literally taking the hand and arm of the robot and moving it to the object, wrapping the fingers of the hand around the object and lifting it. That involves, in this one teaching instance—a large series of specific movements of motor encoders and other devices within the machine. These are associated with algorithms stored and classified in a very large database. The robot tries to capture all the possible states plus what the object looks like, where it is located in its coordinate space in addition to how its arm is moving. From this example, the robot builds a model at a higher level of abstraction within its brain. This internal model is then used to search out and apply specific algorithms and values to fit a different instance of "Pick up the glass."

In addition to controlling major peripherals like its arms, HERB must also integrate and constantly update information about its surroundings. To that end, it incorporates a vision system and a laser-based coordinate system. The laser generates light pulses around the robot and measures the frequency of the returning beam to generate 40,000 points per second around the robot. The data from the laser system is used to build a 3D model of its surrounding world. In addition, a camera running vision processing algorithms is

used to recognize and manipulate objects. The robot can pick up an object, twist it around and build a 3D model that is stored in its database.

Srinivasa stresses that the recent advances in compute power have been a tremendous boon for robotics, especially in providing the ability to search large spaces to find the proper motion algorithms for a given task. The compute power in the robot is also highly distributed with motor controllers at the lowest level—right at the robot's joints. These are very fast, specific-purpose devices that talk to the motors at 1,000 Hz, and their algorithms are at the lowest level of the software hierarchy. The next level consists of behavioral loops, such as image acquisition, that run at about 10 Hz. Then at the highest level are the planning algorithms that take the data and make longer term plans to carry out a task like picking up a glass.

To do this, the robot must execute one of its general-purpose models, possibly named "pick up the glass," and adopt it to the current situation. Thus it will not be executing the exact same routines that were invoked when it learned the task. Rather, it will assess the situation given the coordinates from the laser system of objects in its surroundings and images from the vision system to invoke the proper model. Then it will plan the execution of the task by searching its database for the most appropriate algorithms for that particular instance of the task and arranging them in a sequence, setting variables for those algorithms that have been computed from the coordinate space.

It is this adaptability that sets a robotic system apart from a simpler semi-autonomous automated system. The robot selects a method that is similar to what it has learned before. It then executes it while asking if the object is still there (vision system) and if it is feeling the forces it should be feeling (tactile feedback). It is also moving its arm according to the algorithms that have been set up based on the coordinate space measured by the laser system. If it notices an error, it propagates that error back to the "brain," which is the planning level. The brain has a state machine that reacts to errors. Recognizing, interpreting and correcting for errors is one of the more advanced areas of robotic research.

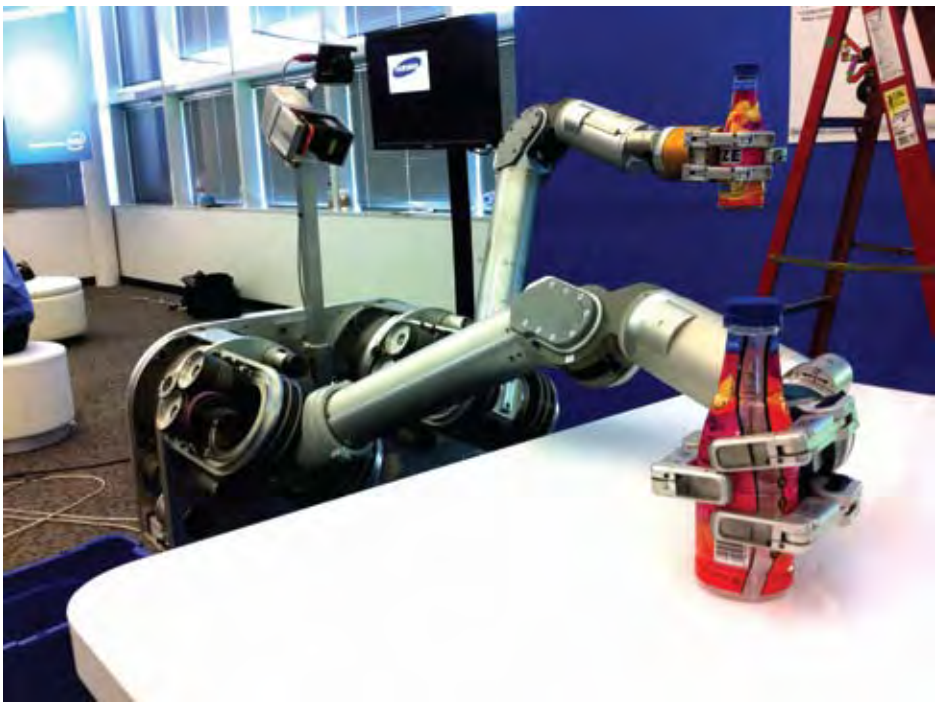

FIGURE 1

The Home Exploring Robotic Butler—HERB—is an Intel research project for proof of concept development in robotics. The system has two arms, a laser system and a visual system for navigation and recognition, and a hierarchical software architecture for adapting task models to particular situations for execution.

Of course, not all robots—even at the research level—use exactly the same mechanisms as the Intel HERB, but to be truly robots as distinguished from automated control systems, they must be able to generalize, adapt and manage an unstructured environment. One of the best moments in his research, according to Srinivasa, was "when I had never programmed the robot to pick up a given object, but it figured it out from what it had learned before."

The question then naturally arises, "Where are we going and what are we getting from robotics research?" Interestingly, much of the long-term goal seems to be directed at things like personal and home robots to take care of ordinary chores. The word "robot," after all, comes from the Czech word "robota," which means "work" or "drudgery." Obviously, the same class of machines could and is being used for work in harsh environments like space. There are aspects of robotics in unmanned aerial vehicles (UAVs), even though these are also subject to direct human control as well.

There are annual competitions involving autonomous vehicles and autonomous submersible vehicles, all of which have attractive possibilities for applications. Although we do not yet have commercially available robotic cars, we do have some advanced automobiles like the Lexus that are capable of autonomous parallel parking. This latter task must meet the more stringent criteria for a robotic system in that it must adapt a general task for parallel parking to each individual situation—especially if it involves a Hell's Angels bike. There are further more immediate applications in health care, and there are other aspects of current research that are being examined for possible spin-offs for applications. Along the way to 3-CPO we will definitely find creative and useful ways to make use of more autonomous and adaptable electromechanical systems no matter what we call them.

Intel
Santa Clara, CA.
(408) 765-8080.
[www.intel.com].

TECHNOLOGY IN
CONTEXT

OpenVPX

The Quest to Navigate the OpenVPX Standard: VITA 65

OpenVPX offers huge potential for performance, ruggedness and I/O in a wide variety of system configurations. Navigating through the specification can be intimidating at first, but once the terms have been recognized and the map laid out, the path to truly functional system topologies will open before us.

by Ken Grob, Elma Electronic

Much has been written recently about the new OpenVPX standard known as VITA 65. This article provides an introduction to the structure of the specification. In order to define a system, it's important to understand how to properly navigate through and decipher the different sections of the specification and its lexicon. Part two, a follow-on article scheduled for the November issue, will discuss how VITA 65 enables a user to build OpenVPX systems by combining slot and backplane profiles that support the establishment of an end system topology.

The OpenVPX standard has been brought to fruition through an intense effort driven first outside, then subsequently within the OpenVPX VITA Working Group. In a few months, a short list of companies and a team of dedicated industry veterans have brought us a 400+ page document that provides concepts and methods to describe system topologies using a new breed of serial fabric technologies and high-speed backplanes.

FIGURE 1

VPX wafer-based connectors fitting into backplane sockets.

As discussed in recent articles, OpenVPX is based on prior VITA standards that initially addressed VPX, including VITA 46.0 and VITA 46.1. These standards formed a good base to allow the design and implementation of new high-speed, high-power systems, but they fell short when it came to fostering interoperability among offerings from different manufacturers of VPX boards and backplanes. So off on a quest went these VPX Knights—a new standard they sought, to alleviate these plights. Contending with MultiGig-2 wafer-based connectors good for data rates up to 10 Gbit/s and new backplane materials, including FR-408 and Nelco 4000-13SI, OpenVPX is taking us to new data transfer rates where we have not gone before (Figure 1).

SerDes-based physical interfaces supporting baud rates of 3.125, 5.0 and 6.25 Gbit/s are now common within the OpenVPX Module and Backplane lexicon. OpenVPX has generated a well structured specification volume, which at first glance would send Don Quixote back to the windmill. Introducing a new set of terms for describing lane-based point-to-point interconnects, the document quickly grew large driven by the unique topology required of each backplane described, and using equation-based formulas that specify Slot, Module and Backplane profiles.

OpenVPX - the Standard

The standard was created to allow for definition of system topologies and to promote interoperability. The specification is divided into 16 sections. Section One covers structure and defines terminology. Key Words are defined. The concepts of Profile Names are introduced and are summarized in Figure 2.

Slot (SLT) in Types: Payload (PAY), Peripheral (PER), Switch (SWH) and Storage (STO).

Module (MOD): Same Attributes as Slot but specific to the module (board) and defines the protocol associated with the ports.

Backplane (BKP) in Types: CEN, DIS, HYB, BRG, where:

- Central (Star)
- Distributed (Mesh)
- Hybrid (VME & VPX)
- BRG (Bridge, e.g., parallel VME to VPX)

Naming conventions for profiles are described to allow a user to create a name

to define specific Module, Slot and Backplane profiles. Figure 3 is an example of a 3U payload slot profile with one fat pipe data plane, two fat pipe expansion planes and two ultra thin pipe control planes. All the additional attributes are found in the document in Section 14.2.2 of VITA 65. Module profiles and backplane profiles follow a similar naming convention that always includes the section number where the profile is defined in full detail.

Section two of the VITA 65 standard addresses compliance, which is an important topic since it describes how one must consider and comply with over 400 rules, permissions and recommendations. These requirements have been created to ensure interoperability between a backplane and the chosen module that one is about to plug into a defined OpenVPX Slot. Every VPX rule also has a compliance requirement that must be documented and established by one of four methods. These four compliance methods are defined within the OpenVPX standard as follows:

Inspection: The Inspection method primarily uses a static, visual means to demonstrate conformity.

Demonstration: The Demonstration method primarily uses a dynamic, visual means of showing functionality to demonstrate conformity. While test equipment may be required as part of the demonstration setup, measurements are typically not required.

Analysis: The Analysis method primarily uses theoretical means to demonstrate conformity. Analysis 1063 input parameters may be based on component datasheet or empirically derived parameters.

Test: The Test method primarily uses physical measurements and test procedures to demonstrate 1073 conformity. The Testing method is necessary when inspection, demonstration and analysis 1074 methods are inadequate, not supported by tools, or cost prohibitive.

Section three of VITA 65 discusses the Utility Plane, Power Distribution, System Control Signals, the Reference Clocks and the GPIO Signals. Pin Assignments are defined for J0/P0 and J1/P1. In addition, Section four covers the mechanical specifications. Described here are Slot Pitch, Connectors, Keying and RTM Connections.

In Section five, we are getting into the good stuff. This section covers the Fabric Protocols referenced today by VITA 65. Three major Protocols are defined, including Ethernet, Serial Rapid I/O (SRIO) and PCI Express (Table 1).

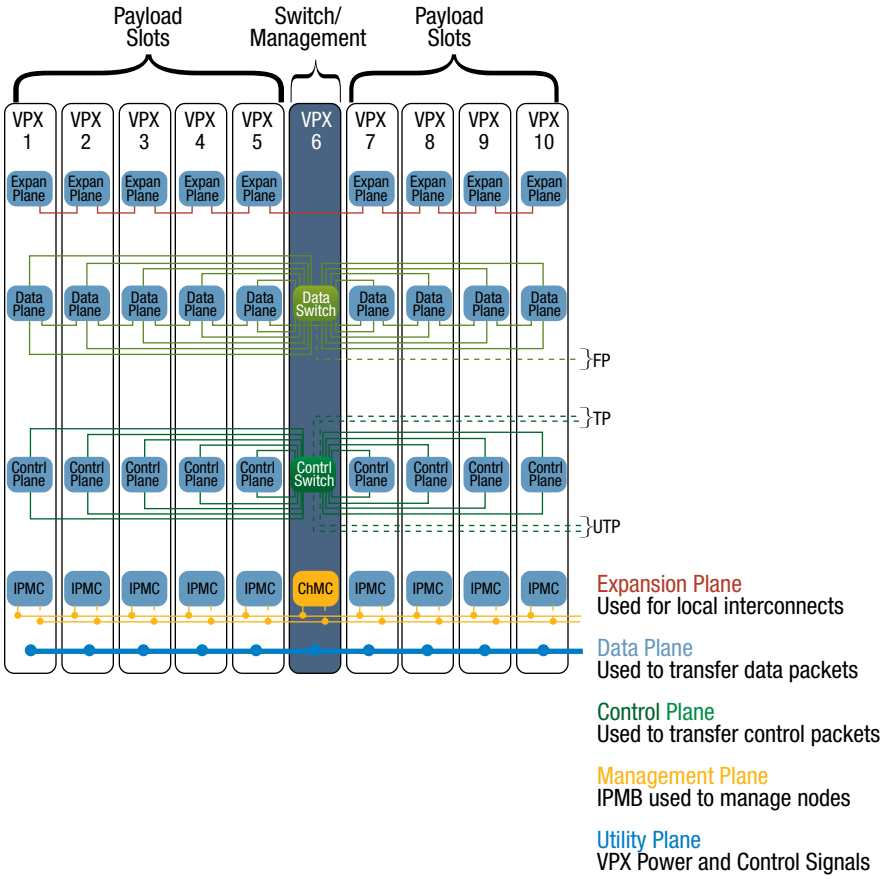

FIGURE 2

Planes in the VPX architecture are defined as Expansion, Data, Control, Management and Utility.

OpenVPX Fabric Protocols		
Ethernet	**Serial Rapid I/O**	**PCI Express**
1000Base-BX	SRIO 1.3 – 3.125 Gbps	PCI Express Gen 1 – 2.5 Gbps
1000Base-T	SRIO 2.0 – 5.0 Gbps	PCI Express Gen 2 – 5.0 Gbps
1000Base-KX	SRIO 2.0 – 6.25 Gbps	
10GBase-BX4	SRIO 2.1 – 5.0 Gbps	
10GBase-KX4	SRIO 2.1 – 6.25 Gbp	

TABLE 1

Summary of fabric protocols that can be used in OpenVPX

Sections six, seven and eight discuss Slot Profiles, Backplane Profiles and Module Profiles. Profiles are used as the central graphical representations of slots and backplanes in OpenVPX. With that said, it's time to introduce a table called Profiles at a Glance that show the relationship between Slot, Module and Backplane profiles (Table 2).

Our Knights had found themselves banished from the Castle of OpenVPX until they could decipher the cryptic terminology to find the profile they were looking for, which would in turn enable them to find their path through the document—the path that would lead to the definition of a system. So armed with new knowledge they added royal tools to the document to help find the fair profile in distress, at a glance. These regal tools included very nice hyperlinked tables, which summarize available 6U Module Profiles in Table 11.2-1, and a 6U Module to Backplane Profile reference Table 11.2.2-1. The 3U tables are Table 15.2-1 and Table 15.2.2, specifying similar information for the 3U form factor. These tables are a great Module and Backplane navigational reference. Having deciphered the clues, they can now read the tables, one example of which is shown in Table 3.

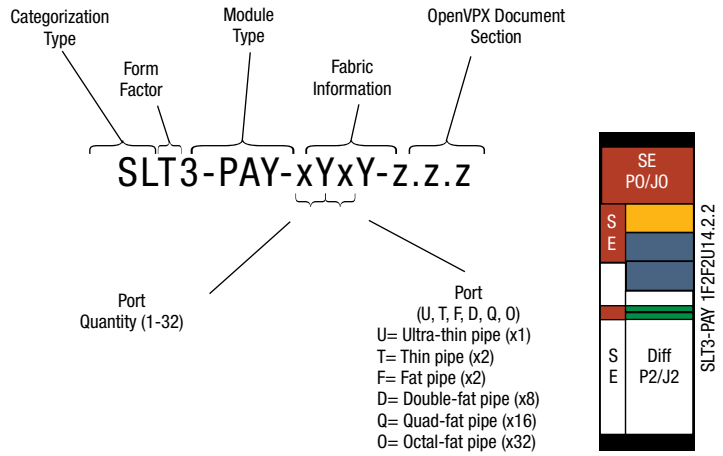

FIGURE 3

Example of the nomenclature used to describe a VPX slot profile. A similar schema applies to module and backplane profiles.

Module	Slot	Backplane
A Module Profile is: Slot Profile + Protocol Defined Defines: • Module Number is keyed to Slot Number • Described by a Table not a Graphic **Relationship to Slot Profile** • Suffix (X) Defines Protocol in a table - Module…-12,4,1-(X)	**A Slot Profile is:** Protocol Independent Mapping of Slot I/O Defines: • User I/O • SE=Single Ended • DIF=Differential • Port Mapping of Slot • Form Factor Defined **Types: PAY, PER, SWH, STO** Relationship to Module Profile • Slot Number is Keyed to Module Number • Module 12.4.1-(X) < --> Slot 10.4.1	**A Backplane Profile is:** A collection of interconnected slots Defines: • Form Factor: 3U, 6U • Pitch: .8'', 1'' • Number of Slots • Slot Profiles Used **Types CEN, DIS, HYB, BRG** • Channel Baud Rate • Parameter for speed
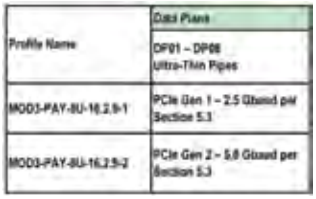	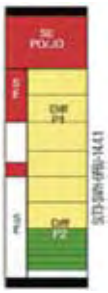	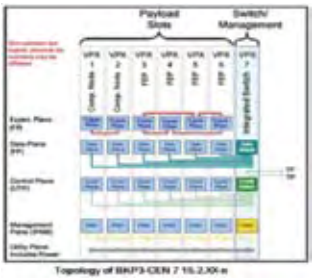
MOD-SWH-20U19F-12.4.1-5	SLT6-SWH-20U19F-10.4.1	BKP6-CEN05-11.2.5-1

TABLE 2

Examples of the graphical representation of module, slot and backplane profiles in OpenVPX.

Other new concepts include the definition of Lanes, Channels, Ports and Pipes to describe how bidirectional serial lanes are grouped into different width channels. Pipes, for example, come in different sizes:

- Ultra Thin Pipe = 1 Lane
- Thin Pipe = 2 Lanes
- Fat Pipe = Four Lanes
- Double Fat Pipe =8 Lanes
- Quad Fat Pipe = 16 Lanes

What Does OpenVPX Do for the System Architect?

The specification gives the system designer a set of terms, a common language if you will, to describe a system uniquely. The syntax defined uses graphical icons that allow the visualization and description of the system topology. The definition of slot and module profiles establishes rules for mapping pins to slots.

Systems are described by the backplane profile that is comprised of a set of slot profiles. A backplane is simply a set of interconnected slot profiles where each pipe in one slot is mapped to a pipe in a second slot. This point-to-point mapping results in a unique topology for each backplane profile. It is unique because the lane pairs in a serial fabric will be connected point-to-point as required. The backplane in Figure 4 is full mesh, made up of five instances of Payload slot profile SLT-PAY-4F 10.3.1, connected in slots one through five. Each slot has a fat pipe connection to

each of the other four slots creating a full mesh as shown in Figure 4.

In closing, one other navigational tip would be useful to note about profiles. Below is a Payload Slot profile. Note that the planes are described from top to bottom of the profiles and are associated with colors. However, the color key for the slot profiles is never explicitedly defined within the VITA 65 document nor has the convention of working down from the top of the connector to assign planes ever been explicitedly explained. Rather, these conventions are only implied and left for the reader to figure out on his or her own. Take for example the following 3U slot profile:

SLT3-PAY-2F1F2U-14.2.1
Color Code Key: Yellow - Data Plane,
Blue - Expansion Plane,
Green - Control Plane

The first two yellow data plane ports are described by the first field as 2F, shown in yellow representing two fat pipes. The second field, 1F, describes the expansion plane in blue with one fat pipe. Finally, the green section represents the 3rd field, and the 3rd position down in the connector, representing the Control Plane shows 2U; two ultra thin pipes used for Ethernet connections. These can be followed in the slot profile diagram in Figure 5 giving the adventurous reader a more convenient means of following the formula and being able to compare it quickly with other profile diagrams.

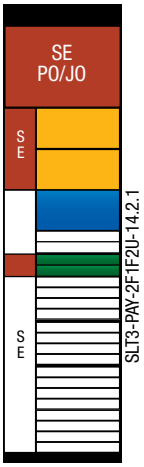

FIGURE 5

The unofficial color scheme and top-down reading convention in OpenVPX can be helpful in navigating profile descriptions.

Stay tuned for November when our brave Knights will set off on a quest to follow the profiles, charts and documents to establish an end user topology and define a real-world system. ◢

Elma Electronic
Fremont, CA.
(510) 490-7388.
[www.elma.com].

BKP6-DIS05-11.2.16.n

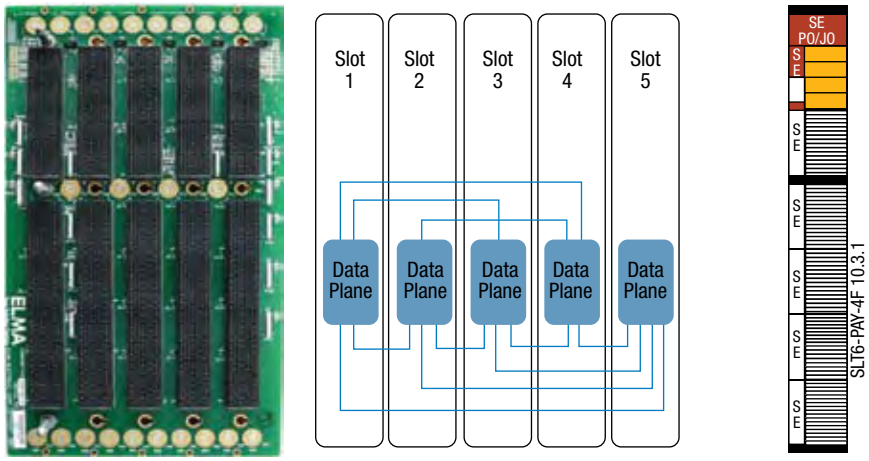

FIGURE 4

"The DIS05" in the backplane profile here indicates a distributed, or mesh, connection of five instances of the payload slot profile (right), which has four connected fat pipes.

Backplane Profiles Prefix for all names in this column is: BKP6-	Payload Modules Prefix for all names in this column is: MOD6-PAY-	Switch Modules Prefix for all names in this column is: MOD6-SWH-	Peripheral Modules Prefix for all names in this column is: MOD6-PER-	Miscellaneous Modules Prefix for all names in this column is: MOD6-
CEN16-11.2.2-n	4F1Q2U2T-12.2.1-n	20U19F-12.4.1-n		
CEN20-11.2.3-n	4F1Q2U2T-12.2.1-n	20U19F-12.4.1-n		
CEN10-11.2.4-n	4F1Q2U2T-12.2.1-n	20U19F-12.4.1-n		
CEN10-11.2.5-n	4F1Q2U2T-12.2.1-n	16U20F-12.4.2-n		
CEN10-11.2.6-n	4F1Q2U2T-12.2.1-n	16U20F-12.4.2-n		
CEN10-11.2.7-n	4F2T-12.2.2-n	24F-12.4.3		
CEN06-11.2.8-n	4F2T-12.2.2-n	24F-12.4.3		
CEN12-11.2.9-n	4F2T-12.2.2-n	24F-12.4.3		
DIS06-11.2.10-n	4F2T-12.2.2-n	4F24T-12.4.4		
HYB17-11.2.11-n	4F2T-12.2.2-n	4F24T-12.4.4		BGR-4F1V2T-12.5.1-n
HYB08-11.2.11-n	8F-12.2.3-n		2F-12.3.2-n	BRG-4F1V-12.5.2-n
CEN09-11.2.13-n	8F-12.2.3-n		2F-12.3.2-n	
CEN06-11.2.14-n	8F-12.2.3-n		2F-12.3.2-n	
DIS06-11.2.15-n	4F2T-12.2.2-n	4F24T-12.4.4		
DIS05-11.2.16-n			4F-12.3.1-n	

TABLE 3

VITA 65 establishes a family of 3U and 6U standard backplanes for development applications. Each backplane is comprised of a number of slot profiles which in turn correspond to various module profiles. This chart summarizes all the defined 6U backplane profiles and shows the module profiles that are compatible with the different slots that comprise each of the 15 different 6U backplane profiles.

Is There Life Beyond Defense and Aerospace for VPX?

Developed primarily with military applications in mind, the VPX standard has characteristics of ruggedness, high performance and high-speed I/O that lend themselves naturally to non-military, commercial environments where harsh conditions demand top of the line performance and reliability.

by Ben Klam and Dave Barker, Extreme Engineering Solutions

What does a Navy SH-60 helicopter landing on a ship have in common with an oil rig and a coal mine? Answer, they are all extremely dangerous places with harsh environments of temperature and humidity extremes, tremendous amounts of shock and vibration, and gaseous and liquid contaminants.

In 2010, unfortunately, we have become all too familiar with just how dangerous coal mines and off-shore oil rigs can be. Earlier this year, the Upper Big Branch mine explosion in West Virginia killed 25 miners. And more recently, we experienced the devastating consequences of the explosion on the Deepwater Horizon rig that killed 11 and is still causing untold damage to the Gulf of Mexico's environment and economy.

In the aftermath of these disasters, the federal and state governments will most likely tighten the regulation of the oil, gas and mining industries. It is not a stretch to imagine that in order for operators to meet the regulations and avert future disasters there will be a need for improved real-time monitoring, analysis and reaction. Before each of these disasters occurred there were triggers that, if properly monitored and acted upon, could have avoided or minimized the impact. Improved real-time monitoring could detect increased methane levels in mines and problems at the well head 5,000 feet below the surface sooner. The technologies currently used in these industries do not provide the level of real-time monitoring necessary for operators to be able to save lives and avert disasters.

FIGURE 1

The XPedite5470 from Extreme Engineering Solutions is an example of a conduction-cooled 3U VPX Freescale QorIQ P4080-based Single Board Computer; (b) the XCalibur4341 is an example of a conduction-cooled 6U VPX Intel Core i7 processor-based Single Board Computer.

An existing technology that has proven its mettle in deployed embedded real-time military applications is VPX. It is the standard of choice for new systems going into the Navy SH-60 helicopter and many other deployed military applications. It was developed specifically with deployed military applications in mind and supports both 3U and 6U form factors (Figure 1).

Most deployed military applications fit into the C4ISR (Command, Control, Compute, Communications, Intelligence, Surveillance, Reconnaissance) classification. Many deployed C4ISR applications share several characteristics. First, very large amounts of high-speed data stream into these systems from sensors such as digital receivers, A/Ds and cameras. Second, the large amounts of high-speed streaming data have to be moved through the system and processed in real time. Third, these systems are deployed in harsh environments of extreme temperatures, shock and vibration, and exposure to dust, sea salt, chemicals, etc. Fourth, since these systems are deployed on vehicles and aircraft including UAVs, they have severe Size, Weight and Power (SWaP) constraints.

This same technology, with its ability to operate in harsh environments of military applications, is very well suited to handling oil, gas and mine real-time monitoring in support of the stricter requirements these industries will likely incur moving forward. The three primary features of VPX—its ability to operate in harsh environments, handle large amounts of high-speed I/O and process large amounts of data in real time—make it a practical choice for oil, gas and mining monitoring. We will now take a closer look at each of these features.

VPX and OpenVPX

First, a short overview of the VPX standard for readers not familiar with VPX. VITA developed VPX as an open industry standard. It defines a modular embedded computing platform based on the familiar 3U and 6U form factors used by VME and CompactPCI. VPX defines a common set of attributes including physical form factors, signal and power supply interfaces, connectors and power supplies. One of the key attributes of VPX is the choice of backplane connectors. These high-performance connectors enable high-speed switched serial fabrics, such as PCI Express and Gigabit Ethernet, to be used to move data between boards and into and out of the system.

As a module, or board-level specification, VPX does not address system-level issues. To address these issues, VITA developed OpenVPX. OpenVPX is a system-level specification that builds on the module-centric VPX specifications. It provides a nomenclature for system integrators, module designers and backplane providers to describe and define aspects and characteristics of a system. OpenVPX addresses interoperability of modules, backplanes, power supplies, enclosures and other system-level components to make it easy for system designers to integrate components from different vendors into a system.

The VPX standard was developed by VITA to address the harsh environments that many military and aerospace applications operate in. Many deployed military applications face temperature extremes, shock, vibration, humidity, dust, airborne and liquid contaminants, and electromagnetic interference (EMI). If that isn't enough, many have to contend with dirty power supplied by vehicle or aircraft electrical systems.

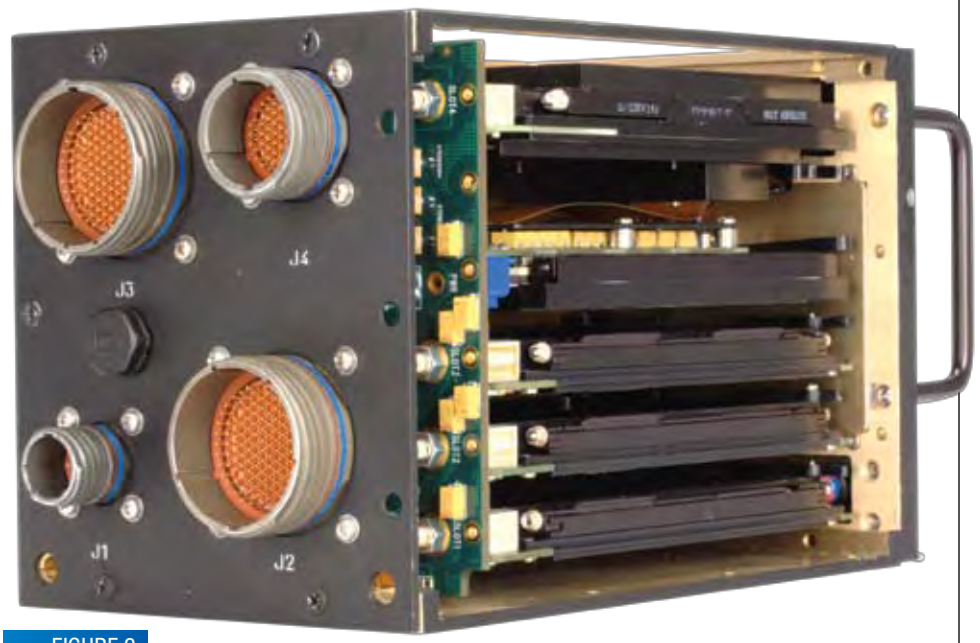

FIGURE 2

The XPand3200 (with a sidewall removed) from Extreme Engineering Solutions is an example of a 1/2 Air Transport Rack (ATR) conduction-cooled chassis that effectively isolates 3U VPX modules from the harsh environment in which it is deployed.

To address these issues, the VPX specifications define a number of standard techniques to isolate system-level components from their environment such as conduction-cooling, full product encapsulation and two-level maintenance. VPX was designed from the ground up to adhere to the strict military environmental testing methods of MIL-STD-810 and electromagnetic interface testing methods of MIL-STD-461. Systems constructed using VPX modular components can survive exposure to the worst case environments.

In addition to addressing environmental issues, VPX has also defined a modular power supply designed to handle the normal, abnormal and emergency power characteristics outlined in MIL-STD-704. Most rugged applications share common requirements such as transient, overvoltage and under voltage conditions specified in this military standard. By leveraging a modular, military ruggedized power supply approach, VPX technology allows system designers in any market to maximize design reuse and efficiency while minimizing program risk and cost.

The VPX standard was developed to support large amounts of high-speed I/O through the backplane connectors. VPX supports both front-panel I/O and rear I/O. 3U VPX supports a total of 64 differential pairs on the backplane connectors, which can be divided between data, control and I/O. 6U VPX extends this to a total of 160 differential pairs. Some of the backplane pins are dedicated for communication between modules while other pins are dedicated to external I/O. The VPX backplane connectors support signaling rates in excess of 6.25 Gbit/s. This provides enough bandwidth to support data and control fabrics consisting of the latest high-speed serial fabric protocols while also providing enough external I/O for raw sensor data.

To support the high-bandwidth, low-latency and low-overhead communication requirements of many C4ISR applications, a switched serial fabric such as PCI Express is utilized to move data through a system. Currently, VPX systems utilize three high-speed serial fabric protocols: Serial RapidIO, PCI Express and Gigabit Ethernet. As an example of a typical ap-

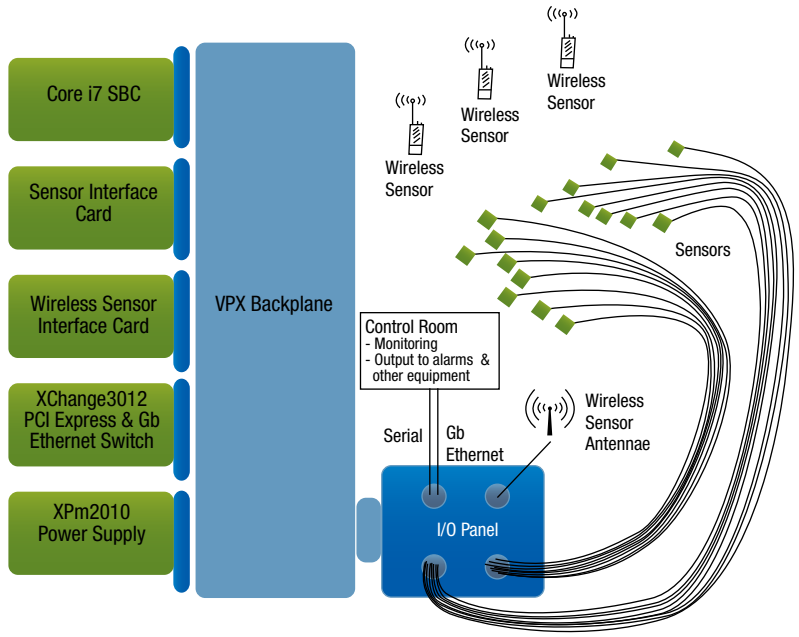

FIGURE 3

An example of real-time monitoring system that is monitoring a number of wired and wireless sensors. With an Intel Core i7 processor-based SBC, the system can process and analyze the sensor data in real time. When a problem is detected, using Gb Ethernet, the system can alert operators and interface directly to other equipment that can mitigate or resolve the problem.

plication, let us consider running a VPX x8 PCI Express link (16 differential pairs) between two adjacent modules in a VPX system. PCI Express 2.0 has a bandwidth of 500 Mbyte/s per lane, which would yield 4 Gbyte/s bandwidth for an x8 Link.

To ensure optimal system performance, the computation bandwidth much be matched with the communication bandwidth. The VPX standard was developed with this in mind. If VPX systems have the processing power to detect incoming missiles traveling at over mach 2, track them, and launch a counterstrike against them, they can handle the processing tasks required for real-time mine and oil rig monitoring. A variety of VPX single- and multi-processor boards are available utilizing today's state-of-the-art processing technologies such as the Freescale QorIQ processor, the Intel Core i7 processor and the Xilinx Virtex-6 FPGA.

With the large amount of processing power that can be put into a system, a very important consideration for high-performance embedded systems is cooling. A typical amount of power consumed by 3U VPX processing cards is in the range of 30 to 70 watts. The amount of power that can be dissipated by a module is heavily dependent on the type of cooling method employed (conduction or air-cooled) as well as the materials and techniques used to extract power. 3U VPX ATR conduction-cooled boxes, when designed properly, can address the thermal challenges of most applications. 6U VPX cards offer a larger surface area and thus improve air-cooling capacity, and therefore work very well in forced air-cooled systems. However, many applications are constrained to conduction-cooling, and it should be noted that 6U cards have the same amount of card edge rail area for conduction-cooling as 3U solutions.

Expanding the Scope of Applications

As we have seen, VPX systems are rugged, they can handle a large amount of high-speed I/O, and they have the capability to process large amounts of data in real time. These are important attributes to achieving more effective real-time mining and oil rig monitoring. Mines and drilling rigs are very harsh environments—having ruggedized systems that can adequately

protect the embedded computing hardware from the harsh environment allows for optimal placement of monitoring systems deep in a mine, on an oil rig, or even at the well head. This in turn makes it easier to optimally place sensors that are being monitored. Because of VPX's inherent I/O capabilities, VPX systems can monitor a very large number of sensors. Once data is brought in from the sensors, VPX systems have the processing bandwidth to perform real-time processing and analysis of the data to quickly and effectively deal with a situation before it turns into a disaster.

If real-time mining and oil rig monitoring systems are to be deployed in harsh environments, their internal processing elements need to be isolated from their surrounding environments. Air Transport Rack (ATR) enclosures are a proven method of achieving this. While not part of the VPX specifications, conduction-cooled ATR chassis have been used for years to house military systems deployed in harsh environments of ground vehicles, aircraft and sea vessels (Figure 2). These same enclosures can be leveraged for real-time monitoring systems deployed in the harsh environments of mines and oil rigs.

Using VPX systems for real-time mine and oil rig monitoring shows how a technology developed for one industry and its associated applications, specifically high-end deployed C4ISR systems, can be utilized within other industries. Applications that have similar requirements, namely ruggedization, high communication bandwidth and high computation bandwidth, can leverage this established standard. Leveraging VPX technology allows system designers in any market to maximize design reuse and efficiency while minimizing program risk and cost. They gain access to a thriving and competitive market of Commercial Off the Shelf (COTS) products from a number of vendors. They can develop their own in-house products designed to the VPX specifications. And, they can easily integrate COTS products and VPX products they develop in-house into systems (Figure 3). One other important aspect to VPX that system designers can leverage is software support. There is wide OS and Real-Time Operating System (RTOS) support across VPX products including Windows, Linux, Wind River Systems VxWorks, Green Hills Software Integrity and LynuxWorks LynxOS.

So, to answer the question, "Is there life beyond defense and aerospace for VPX?" Only time will tell. VPX certainly has the features and market support to make it a viable and practical choice for high-end embedded computing applications in industries other than military. With the emergence of VPX as an established and viable technology, a new generation of real-time monitoring equipment can be developed to meet more rigorous mining and oil industry regulations and to help these industries protect miners, oil rig workers, sea life, coastal environments and coastal economies.

Extreme Engineering Solutions
Middleton, WI.
(608) 833-1155.
[www.xes-inc.com].

Beamforming Systems Moving Toward New VPX and FPGA Solutions

Today's FPGAs incorporate so much functionality, including DSP capabilities plus gigabit serial interfaces, that they can become the heart of high-end signal processing systems. Combining these capabilities with the serial system architecture provided by OpenVPX can lead to truly powerful and expandable systems.

by Rodger Hosking, Pentek

Beamforming is a signal processing technique that utilizes an array of sensors to achieve directionality, increase the strength of transmitted signals and improve the quality of received signals. Beamforming applications span frequencies from sub-audio to light, and encompass a diverse range of critical applications for commerce, industry, government and defense. Systems developers are continuously exploiting new technology to boost performance for specific applications, with significant emphasis on communications and signals intelligence.

Let's first look at the general principles of beamforming as applied to the reception of radio frequency signals, followed by a basic discussion of the required signal processing algorithms. For this technology, the latest generation of field programmable gate arrays (FPGAs) is available to cover several important roles in beamforming, including both the key DSP resources, and the fast links for data transmission. Finally, a powerful new embedded system architecture called OpenVPX is now at hand, whose features are extremely well suited for implementing deployable, highly scalable beamforming systems.

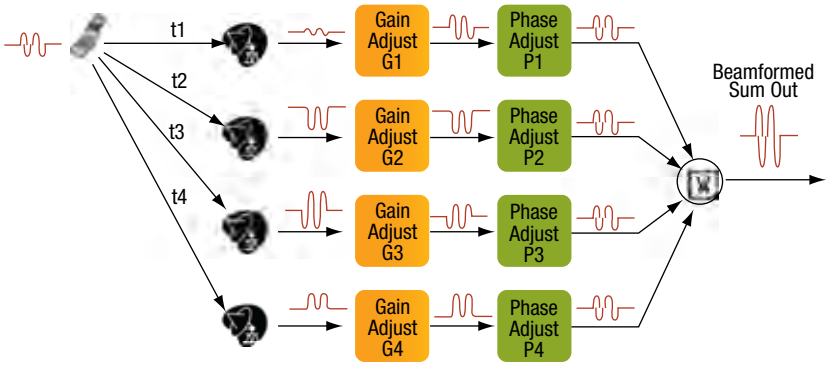

FIGURE 1

Beamforming adjusts phase and gain of signals from each antenna in an array to compensate for differential delays (tn) and attenuation, so that signals arriving from a particular angle relative to the array add constructively when combined in the summer.

For software radio systems, the beamforming sensors are transmit and receive antennas. For receiver systems, the signal arrival delay at each antenna is directly proportional to the path distance from the source. The beamforming process adjusts the gain and phase of each antenna signal to cancel the delay path differences for signals arriving from a particular direction. Aligned signals are then summed together to produce high signal to noise reception in the chosen direction (Figure 1).

By adjusting gain and phase in each path, the antenna is electronically

A Great Listener

wireless ip cores

X6 rx

Features

- Four 160 MSPS, 16-bit A/D channels
- Down-Converter ASIC supporting up 24 Narrowband or 8 Wideband Channels
- +/-1V, AC-Coupled, 50 ohm, SMA inputs
- Xilinx Virtex6 SX315T/SX475T or LX240T
- 4 Banks of 128MB DRAM
- Ultra-low jitter programmable clock
- x8 PCI Express Gen2, providing 2 GB/s sustained transfer rates
- PCI 32-bit, 66 MHz with P4 to Host card
- PMC/XMC Module (75x150 mm)
- < 15W typical
- Conduction Cooling per VITA 20
- Ruggedization Levels for Wide Temperature Operation
- Adapters for VPX, Compact PCI, desktop PCI and cabled PCI Express systems

Ideal for

- Wireless Receiver
- WLAN, WCDMA, WiMAX front end
- RADAR
- Medical Imaging
- High Speed Data Recording and Playback
- IP development

6 VIRTEX

FrameWork Logic

RoHS

805.520.4260 phone • www.innovative-dsp.com

Innovative Integration
... real time solutions!

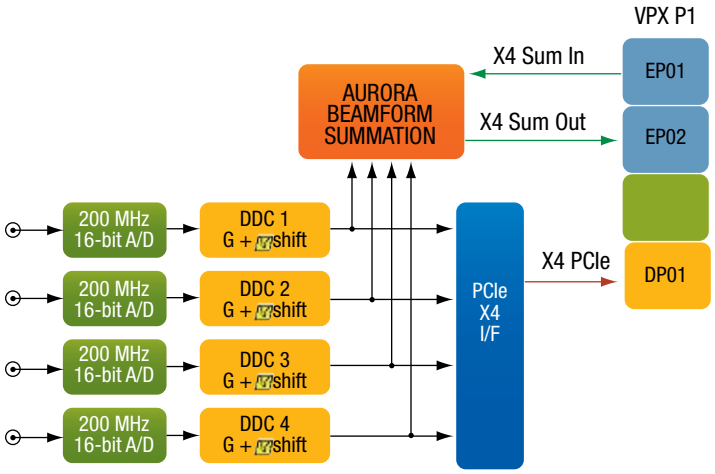

FIGURE 2

Model 5353 3U VPX Beamformer Module with four A/Ds, four DDCs, X4 PCIe interface, gain adjusts, phase shifters and summation engine for beamforming.

"steered" without the need for moving mechanical structures. An example of a software radio application that uses beamforming is direction finding, in which a beamformed antenna can be steered to locate the arrival angle of a signal source. Two or more of the arrays can be used to triangulate the exact location of the source, which is essential for many signal intelligence and counter terrorism efforts. In addition to directionality, beamforming also improves reception in so-called "diversity receivers." The combined signal from multiple antennas boosts the signal-to-noise ratio compared to a single antenna, thus extending the operational range of the receiver system.

Missile detection and countermeasure applications use beamforming to improve tracking of an object allowing for early detection and improved responsiveness. With no moving mechanisms, airborne arrays take full advantage of electronic steering to dramatically improve the range and target resolution. And lastly, beamforming allows spatial frequency sharing for commercial mobile phone carriers by dividing one cell into several beamformed pie-slice sectors that can share the same frequency.

FPGAs: Ideal Beamforming Engines

Each new generation of FPGAs delivers new features, higher levels of performance and reduced power consumption for a given function. The latest announcement from Xilinx is the new 7 Series featuring 28 nm technology, serial gigabit transceivers with rates up to 13 GHz, fast

PCI Express end points, nearly 4,000 DSP engines and over 900,000 logic cells. FPGAs have become extremely popular for embedded software radio functions, and they are especially appropriate for beamforming applications.

After amplification and analog down conversion to an IF frequency, each antenna signal must be digitized with an A/D converter. To handle this task, FPGAs offer high-speed LVDS interfaces supporting data converter peripherals operating at sample rates to 500 MHz and higher.

The next task is digital down conversion of the IF signal to complex baseband, which is performed by mixing the input signal with the signal from a numerically controlled local oscillator (NCO). The mixer employs one of the multipliers in an FPGA DSP block, and the NCO is a phase accumulator (also part of the DSP block) followed by a sine look-up table. This is followed by a low pass filter (using multipliers, registers and adders) set equal to the signal

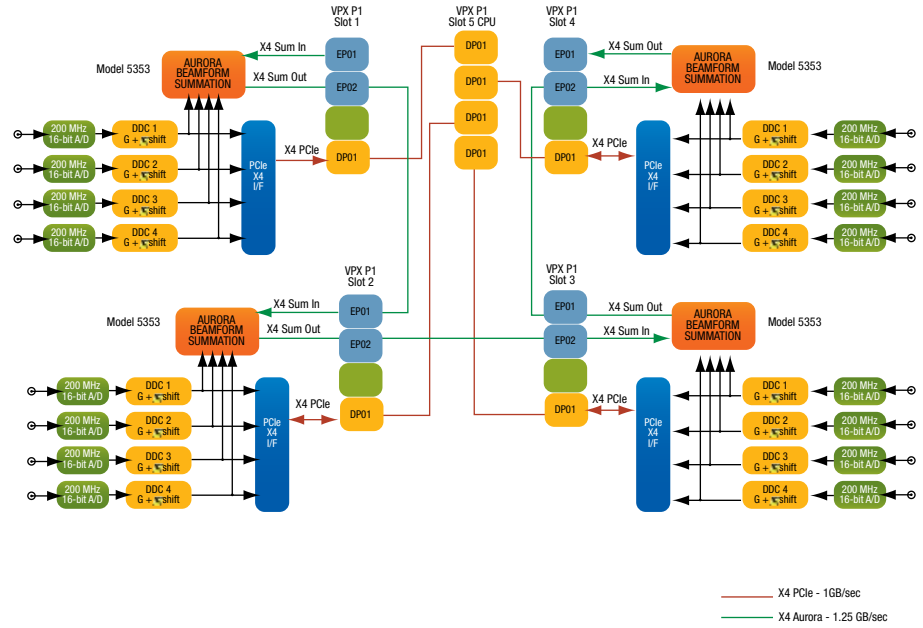

FIGURE 3

16-Channel VPX Beamforming System. Each Model 5353 creates four beamformed signals that contribute to a summation signal propagated through X4 Aurora gigabit serial backplane links on the expansion plane. A CPU module connects to all four 5353s using x4 PCIe using data plane gigabit serial backplane links on the data plane.

channel bandwidth. Together, the mixer, NCO and filter all comprise the DDC, or digital down converter. Special circuitry incorporated in the DDC allows the user to adjust the phase and gain of the down converted signals to support the special requirements of beamforming. The adjusted outputs of each DDC are summed together, again taking advantage of the adder in the DSP block.

A/D conversion plus all beamforming DSP operations for up to four channels can be handled in a contemporary FPGA-based software radio module. However, for larger systems with many antennas, the summation must accommodate multiple modules that must operate with synchronous sampling and must preserve strict alignment of DDC samples forming the final sum.

VPX Beamforming Module

OpenVPX, based on the recently adopted VITA 65 standard, provides an effective taxonomy for describing VPX components, and also defines numerous "profiles" for boards, slots and gigabit serial backplanes that detail specific configurations of channels, interconnections and fabrics. Instead of starting from scratch each time, designers can browse through these standardized profiles to find one that satisfies the objectives of each new system. By narrowing the field of configurations, these profiles boost reusability and interoperability between vendors.

Pentek's Model 5353 Software Radio Beamformer is a 3U OpenVPX module, featuring four 200 MHz 16-bit A/D converters and two Virtex-5 FPGAs. Inside the first FPGA are four digital down converters (DDCs) with programmable phase shift and gain, four power meters at each DDC output and interfaces to the four A/D converters. A simplified block diagram of the 5353 is shown in Figure 2.

The Model 5353 also includes a summation block that adds the DDC outputs to form a four-channel beamforming sum. This block also accepts a propagated "sum in" signal from another module and generates a propagated "sum out" signal to the next module. The sum in and sum out signals use two x4 Aurora gigabit serial links connected to the VPX P1 backplane connector, each capable of moving data at 1.25 Gbyte/s peak.

To support a 20 MHz IF channel bandwidth with a 25% filter margin, the DDC outputs deliver complex 16-bit I+Q samples at 25 MHz, or 100 Mbyte/s. The propagated sum in/sum out signals also operate at 100 Mbyte/s and are thus easily handled by the 1.25 Gbyte/s x4 Aurora links.

The 5353 system interface for control and data is a x4 PCIe port, also connected to P1. Bandwidth requirements for the control and data port are dominated by delivery of the final beamformed sum out to the control processor. This 100 Mbyte/s stream falls well within the 2 Gbyte/s peak rate of the x4 PCIe port when operating in Gen 2 mode.

A programmable, fabric-transparent crossbar switch allows free assignment of

the two x4 Aurora ports and the x4 PCIe port, in any combination, to the four x4 links on P1. This flexibility allows the 5353 to accommodate various VPX slot profiles and backplanes.

For a larger system with 16 antennas, a total of four 4-channel 5353 modules are required. Since the summation chain requires the same data rate as each DDC, the two sum ports must simultaneously handle 100 Mbyte/s each. This class of signal falls under the definition of "expansion plane" in the OpenVPX specification.

The x4 PCIe interface to handle the data initialization and delivery of beamforming parameters is described as the "control plane" under OpenVPX. Final delivery of the beamformed result to the system control processor is best classified as the "data plane" in OpenVPX definitions.

16-Channel VPX Beamforming System

A 16-channel VPX-based FPGA-based beamformer system using four Model 5353 3U VPX modules is shown in Figure 3. Using the dedicated x4 PCIe links, each module connects to a CPU card located in slot 5, serving as the control and status processor for the system. Each 5353 module digitizes four IF signals from four antennas in the array. Four digital down converters translate the antenna IF signals to baseband as complex digital samples and then perform beamforming signal processing, including phase shifts and gain adjustments. The OpenVPX backplane most appropriate for this system is a five-slot "full mesh" topology design, which provides one x4 link from each slot to every other slot.

Each summation engine accepts the propagated sum from the previous module, adds the four channels from the local module, and then generates a new sum signal for delivery to the sum input of the next module in the chain. The summation paths use Aurora x4 gigabit serial links for the expansion plane connections across the backplane. The final 16-channel sum is delivered from the 5353 in slot 4 to the CPU card in slot 5 across the x4 PCIe interface.

There are several major benefits to this system architecture. First, FPGAs implement the data acquisition, digital down conversion, beamforming DSP tasks and gigabit serial system interfaces. Secondly, all interboard summation paths are supported with existing links on the OpenVPX backplane. Third, the system is highly modular and scalable: additional 5353 modules can be added as required to increase the number of antenna channels. Finally, OpenVPX provides a ruggedized solution capable of operating in a wide range of deployed environments. ◢

Pentek
Upper Saddle River, NJ.
(201) 818-5900.
[www.pentek.com].

Upgrade Existing Industrial Networks with Fiber Optics

Fiber optic technology has many benefits for industrial networks including high levels of electrical insulation and isolation, easy installation, survivability in hostile environments and EMI immunity. Upgrading can bring these and more, such as greater security, robustness and signal integrity.

by Mickaël Marie, Avago Technologies

Factory automation, control and management have become crucial to the world's supply chain. High-volume consumer product and automotive manufacturing, for example, require secure, cost-effective and robust data communications in order to react to problems that, if not corrected quickly, can cause substantial profit loss and customer ill will.

Designers from many industries turn to fiber optic data links as an alternative to copper media. Fiber optic solutions result in reliable data links that are capable of communicating over distances—ranging from inches to kilometers—and are more immune to noise. Industrial network applications benefiting from fiber optic solutions span in-flight infotainment, locomotive transportation, medical equipment, casinos and wind and solar photovoltaic renewable installations.

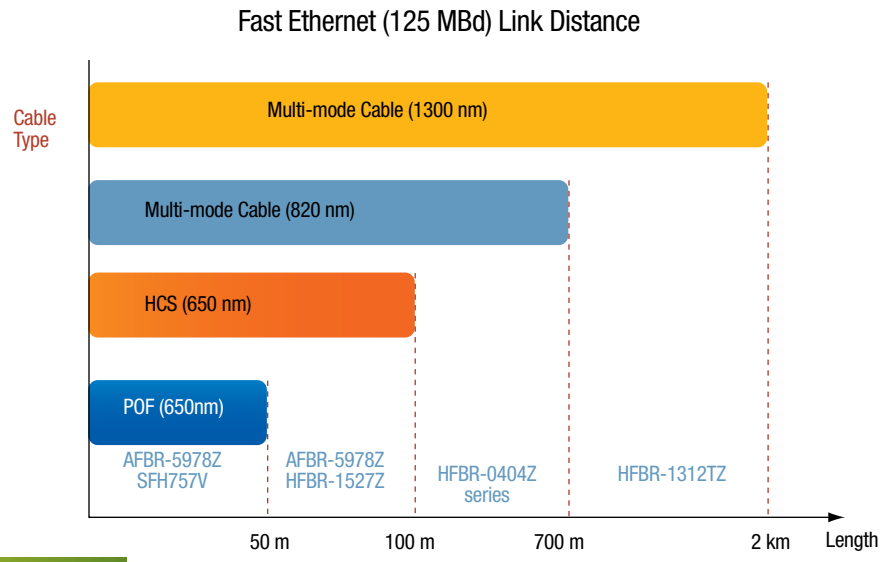

FIGURE 1

Fast Ethernet communication distances extend to kilometers with fiber technology.

Easy to Install Fiber Optics Replace Copper Data Links

Although both copper and fiber are used as a transmission medium, fiber optic solutions offer some clear benefits for the system designer. Industrial Fast Ethernet working over plastic optical fiber (POF) or hard clad silica (HCS), for longer data links, has numerous advantages over copper solutions. While copper-based communication links are susceptible to electromagnetic (EM) fields and emit EM noise, which may interfere with other instrumentation, fiber optic links are immune to EM fields and do not generate any electromagnetic interference (EMI).

Other advantages of choosing fiber over copper include: low weight, com-

Embedded Atom™ SBC with SUMIT Expansion Runs Fast, Stays Cool

Our fanless EBC-Z510-G is a powerful, next-generation EBX-size computing platform. Its outstanding complement of onboard I/O is augmented with MiniPCIe, SUMIT, and PC/104 connectors for high-bandwidth PCIe and USB 2.0 interfaces plus standard legacy bus expansion. This extended temperature SBC is compact and flexible enough to meet a broad range of application requirements.

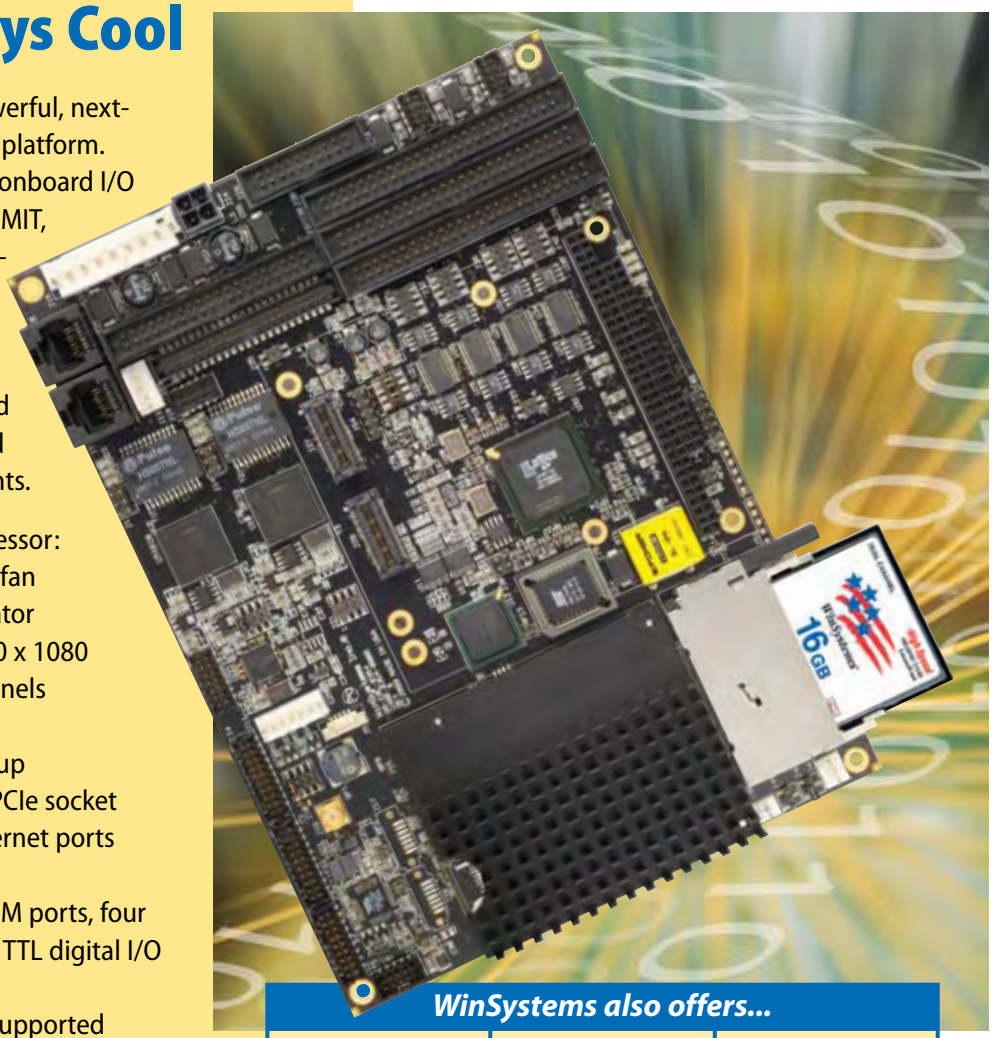

- Low power Intel® Atom™ processor: fanless 1.1GHz or 1.6GHz with fan
- Intel® Graphics Media Accelerator displays resolutions up to 1920 x 1080
- Supports CRT and LVDS flat panels simultaneously
- Custom splash screen on startup
- Wireless supported with MiniPCIe socket
- Two independent Gigabit Ethernet ports
- CompactFlash (CF) connector
- Four serial RS-232/422/485 COM ports, four USB 2.0 ports, 48 bidirectional TTL digital I/O lines, and UDMA IDE port
- HD audio, LPT, and keyboard supported
- EBX sized: 5.75" x 8.0"
- Free software drivers for Windows® XPe and Linux
- SUMIT-AB and PC/104 connectors for additional I/O expansion capability
- -40° to +70°C operational temperature range
- Responsive and knowledgeable technical support
- Quick Start Kit for software development

Contact our factory application engineers for additional product information, custom configurations, and pricing.

Ask about our 30-day product evaluation.

WinSystems also offers...

Quick Start Kits	CompactFlash	PC/104 Modules
		
Convenient and efficient	Industrial-grade Reliability	With SUMIT expansion

Call 817-274-7553 or Visit
www.winsystems.com/EBC-Z510

WinSystems®
An Employee-Owned Company

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

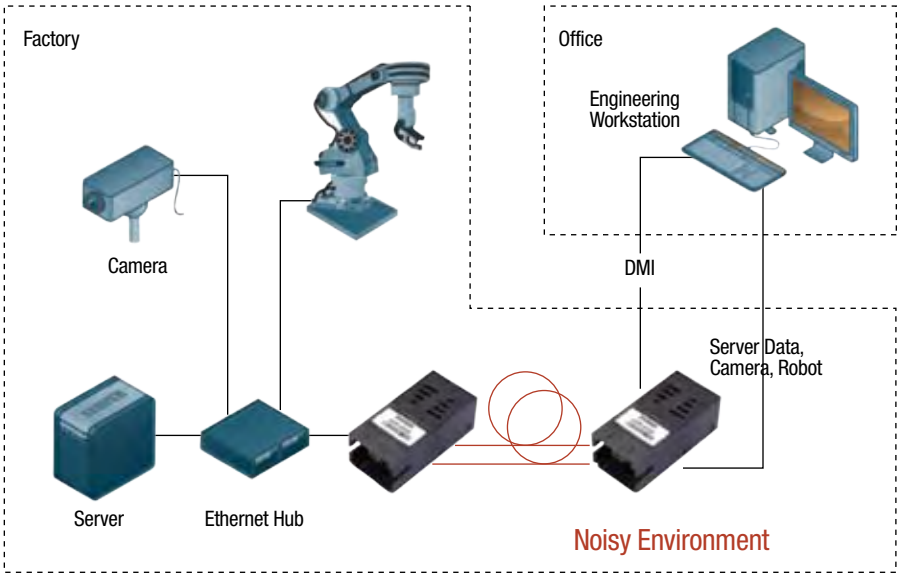

FIGURE 2

Fiber optic links are immune to the EMI generated in industrial manufacturing environments.

plete galvanic separation between link partners, easy field termination and maintenance, easier installation due to short bending radius, and less susceptibility to performance changes caused by temperature extremes and humidity. Fiber optic solutions are also well suited for noisy, industrial environments that have motors and high-voltage; fast-switching circuits, such as in power conversion; automotive manufacturing; medical systems and renewable energy applications, such as wind and solar photovoltaic farms. Fiber optic technology also offers data security since, for instance, it is almost impossible to tap light wave communications.

As shown in Figure 1, multimode (MM), hard clad silica (HCS) and POF cables provide the highest levels of signal integrity over long distances. Multimode cables are especially suited for control systems in offshore wind turbines (over 3 MW) where they are used as a link between the nacelle and the ground. In these applications link length is more than 200m.

Fiber Media for Higher Reliability and Better Safety

Given that the fiber optic link media is essentially glass or plastic, insulation and isolation characteristics are superior to common copper links in that they are immune from EMI. Fiber optic cables (e.g., 5m POF or 100m HCS as indicated by the red lines in Figure 2), offer advantages in noisy environments compared with copper media. With fiber optic solutions, there is no crosstalk between fiber cables or between fiber and copper cable, which makes data transmission more secure. Fiber cable has high immunity to lightning strikes and helps to eliminate ground loop induced errors. When hazardous conditions exist or the environment is potentially explosive, properly used fiber solutions can help increase safety by lowering the potential damage from lightning strikes and ignition caused by electrical sparking.

In addition to being a more cost-effective solution than copper media, optical fiber—such as single-mode (SM), MM, HCS and POF—can be routed in cable ducts, regardless of nearby power conductors; and it is easier to achieve compliance with electromagnetic compliance (EMC) directives and rules (Figure 3).

Additionally, in the case of plastic optical fiber, no special tools are needed and installation training is easy. POF cable installations have been used in rugged industrial environments, such as automotive assembly, for over 15 years. POF is also suitable to use for short to moderate link distances, is easy to field install, and helps simplify connections to equipment and exiting networks. Another key design advantage in industrial applications is that maintenance is low.

Fiber cable and optical transceivers/receivers/transmitters all pass rigorous quality standards. Fiber optics technology has been proven in high-capacity telecommunication links where systems may have a product life time exceeding 10 years. In addition, some fiber component suppliers are vertically integrated and manufacture their own laser diodes for transmitters and PIN diodes for receivers. This provides better control over quality and delivery.

Fiber Offers Bandwidth, Weight and EMC Advantages

In-flight infotainment systems are becoming more sophisticated and advanced. While video quality is improving, the use of larger video screens is becoming more widespread. Airplane manufacturers always look for ways to reduce plane weight. Since fiber optic solutions are much lighter and transmit more data than copper, within a single line, fiber optic solutions are a good choice for use in aviation infotainment systems.

Fiber optics technology also shows up in aircraft carriers and cruise liners. In addition, many modern trains also use fiber cable and optical transceivers, receivers and transmitters. Reliable operation of the train's power source, propulsion system and coach control systems are necessary. However, passenger convenience, information and entertainment systems can also take advantage of the benefits from using fiber optics—high bandwidth over long distance, inherent galvanic isolation, excellent EMC characteristics, and electrostatic discharge (ESD) resistance.

The IEC 61375 Train Communication Network (TCN) standard was created to define communication architecture and protocols for trains. In general, the TCN defines a Wire Train Bus (WTB) and Multifunction Vehicles Bus (MVB). WTB connects the vehicles while MVB connects equipment in a vehicle or group of vehicles. MVB operates over three media types: RS-485 for short distance, transformer-

coupled twisted wire pairs for distances up to 200 meters, and optical glass fiber for distances up to two kilometers. Optical glass fiber is often the preferred media in a locomotive MVB since it has high immunity to electrical noise. Optical fiber connects the controller to devices and subsystems, such as power electronics, motor controllers, brakes and radios. MVB also connects equipment in a coach to control lights, doors, air conditioning and passenger convenience displays for train station and arrival information. Redundancy improves reliability since MVB is backed up by a redundant fiber line and devices transmit on both lines. If one line fails, the other line is available for communication.

Train networks have become more complex, where each station is connected to a central computer for scheduling and event updates. Stations needing to send data back and forth could be a few hundred meters or kilometers away. By using fiber optic cable, more data can be transmitted over longer distance—including video—more reliably than copper cable. In addition, these applications have wires placed side-by-side, running from station to station and from one train compartment to another. When copper wire is used, these adjacent wires can cause interference. Fiber optic cables' EMI advantages, on the other hand, make them immune from this problem.

In trains powered by the electrical grid, single-phase power is taken from the 3-phase AC power grid line to supply the train's 2-phase AC power line. This creates an unbalance in the grid that must be compensated for. One of the most common methods of balancing and restoring the power quality of the grid uses Static Var Compensation (SVC) with Thyristor-Switched Capacitors (TSCs) and a Thyristor-Controlled Reactor (TCR). The TSCs and TCR operate and switch on/off at high voltage and current. This creates very high electromagnetic fields that will induce electrical noise into nearby copper lines. Fiber optic cables are the best medium for sending control signals to the devices in SVC systems because of their immunity to electromagnetic fields.

Train signaling can also take advantage of fiber optic technology's long, reliable data transmission capability in harsh physical environments and in the presence of very high EMI.

Fiber optic technology has become common in medical imaging equipment such as MRI and X-Ray machines. With all the motors and electromagnetic radiation present, these machines generate high levels of EMI. Their communication and control links must have high EMI immunity for reliable and safe operation. High EMI immunity products and electrical isolation are always necessary components for patient safety, so fiber optic products provide a well-suited solution.

and photovoltaic installations operate in remote areas and even offshore, the wide adoption of fiber technology speaks to its reliability and robustness. Downtime and unscheduled maintenance can significantly impact renewable energy costs and adoption, which must be avoided from the system design level. Fiber links are used inside the wind turbine nacelle, which can be over 100 meters above ground or sea level, as well as for data links between turbines and remote management locations.

Fiber's inherent isolation has led to the development of fiber products that need to communicate only over a few

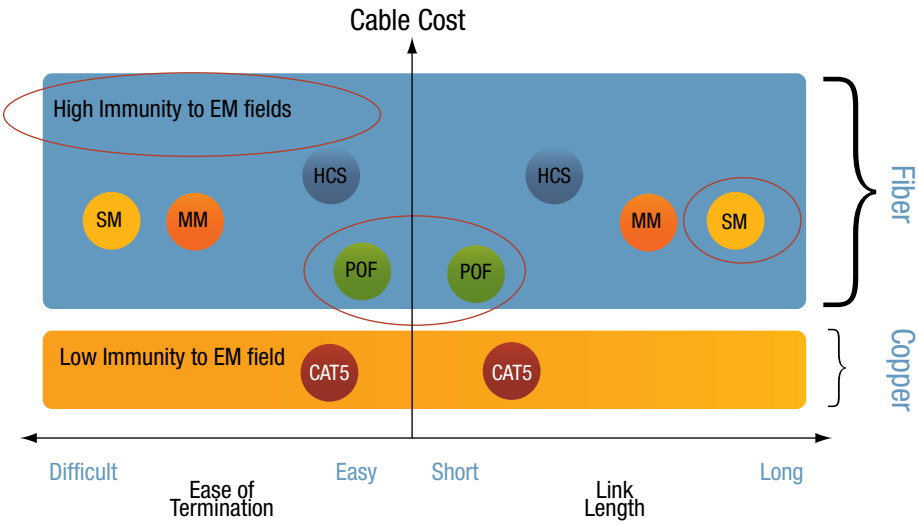

POF = Polymer [plastic] Optical Fiber, SM = Single Mode fiber, MM = Multi-mode fiber, CAT5 = Category 5 copper cable, HCS = Hard Clad Silica

FIGURE 3

Fiber cable has EMI immunity and communication link distances copper media cannot match. POF= Polymer (plastic) Optical Fiber; SM=Single-Mode fiber; MM=Multi-Mode Fiber; HCS=Hard Clad Silica and CAT5= Category 5 copper cable.

In the entertainment area, casinos have machines connected to the central computer/server for data processing, marketing programs and video surveillance. Security is the most important need for a casino's communication network. Fiber optics offer a safe and secure network for casino operators, as it's almost impossible to tap into the signal of the fiber optic cable.

Finally, the renewable energy market has adopted fiber technology. In some wind farm applications, fiber optics is the only suitable communications technology because of EMI. As wind farms

inches. A fiber optic short link device is a cost-effective transmitter and receiver that can provide up to 12 kV transient galvanic isolation on a single PCB. Such a device based on 650 nm fiber optic technology is suitable for use in applications such as inverters (i.e. for wind turbines), IGBT/MOSFET drives and medical equipment. Adding a metal shield can help provide even higher reliability.

Fiber Optics Help Prevent Down Time

In terms of speed, reliability and proactive system monitoring, the Avago

AFBR-5978Z Industrial Fast Ethernet Transceiver is an example of what can be achieved with the use of optical transceivers. In addition, the device's temperature range (-25° to 85°C) helps ensure that the transceiver can stand up to the rugged environment of industrial applications.

The AFBR-5978Z transceiver features an enhanced digital diagnostic interface, compliant to the "Digital Diagnostic Monitoring Interface (DMI) for Optical Transceivers" referenced in the multisource agreement (MSA) SFF-8472, in which fiber optics manufacturers propose similar product housing and features for easy replacement of the transceiver. Industrial Fast Ethernet fiber optic networks represent an upgrade path for Fieldbus networks, with a speed of up to 125 Mbaud compared to 2 Mbaud for Interbus, 12 Mbaud for Profibus, and 12 Mbaud and 16 Mbaud for SERCOS. It also provides the open-architecture, multi-protocol interface that permits both standard and proprietary Fieldbuses to interoperate. Upgrading to industrial Fast Ethernet allows machinery on the factory floor to be assigned IP or MAC addresses, which enables high-speed remote diagnostics and machine sequence changes via Internet access.

DMI provides real-time operational information from the transceiver module. Parameters reported include module temperature, power supply voltage level, and receiver input average optical power level. Also with DMI, the user gains the capability of performing component monitoring, fault isolation and failure prediction in their transceiver-based application. In addition, DMI fully incorporates the functionality needed to implement digital alarms and warnings.

At the higher networking level, industrial Ethernet connects engineering and management workstations to industrial Ethernet hubs for data sharing and control across the enterprise. The value proposition is significant. Fiber solutions are available with various data rates and connectors that serve industrial communications and factory automation applications.

The choice of being able to use discrete or integrated fiber optic component solutions gives the designer the ability to focus on specific design goals. Discrete components give customers the design flexibility to meet their very specific requirements, while the integrated component solution saves design effort, minimizes risk and reduces cost. Best of all, fiber's inherent technology edge in EMI immunity and isolation characteristics suits the factory environment well.

Avago Technologies
San Jose, CA.
(800) 235-0313.
[www.avagotech.com].

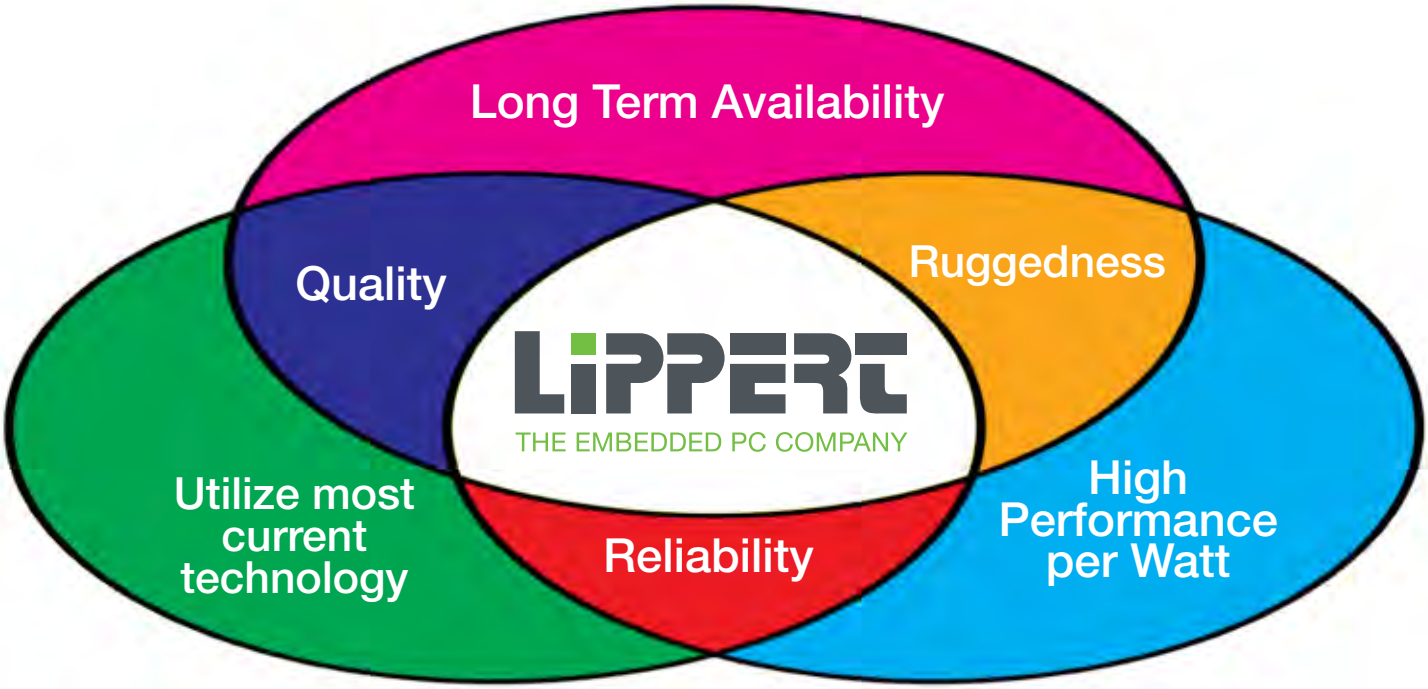

Focused on the essentials

Maximum Performance at Minimum Consumption

Utilizing highest integration levels, latest technologies, and the most powerful processors, for more than 25 years LIPPERT's embedded PCs provide the maximum possible performance while maintaining minimum power dissipation.

Integrated LEMT functions provide condition monitoring.

LiPPERT - efficient and rugged components for market segments as diverse as automotive, automation, medical, military, and aerospace.

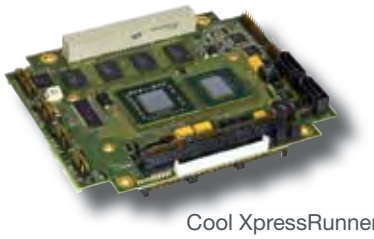

Cool XpressRunner-GS45

Model	Format	CPU, Clock	RAM (max)	Miscellaneous	LEMT
CoreExpress-ECO	CoreExpress®	Intel® Atom™ processor Z5xx, 1.1 ... 1.6 GHz	2 GB soldered	2 x PCI Express, SDVO, USB	✔
Cool LiteRunner-ECO	PC/104	Intel® Atom™ processor Z5xx, 1.1 ... 1.6 GHz	2 GB soldered	LAN, SATA, USB, Micro SD	✔
Cool LiteRunner-LX800	PC/104	Geode™ LX800 500MHz	256 MB soldered	Mini-PCI, CF slot, 2 x LAN	✔
Cool SpaceRunner-LX800	PC/104-Plus	Geode™ LX800 500 MHz	256 MB soldered	2 GB SSD	✔
Cool RoadRunner-945GSE	PC/104-Plus	Intel® Atom™ processor N270, 1.6 GHz	2 GB soldered	2 GB SSD, SATA	✔
Cool RoadRunner-LX800	PC/104-Plus	Geode™ LX800 500 MHz	1 GB	CF slot	✔
Cool XpressRunner-GS45	PCI/104-Express	Intel® Core™ 2 Duo processor, 1.2 ... 2.26 GHz	1 GB	GB LAN, SATA, USB	✔
Hurricane-LX800	EPIC	Geode™ LX800 500 MHz	1 GB	LAN, USB, uDiskOnChip	✔
Hurricane-QM57	EPIC-Express	Intel® Core™ i7 processor, 1.06 ... 2.53 GHz	4 GB	LAN, SATA, USB, Display Port	✔
Thunderbird-GM45	Mini-ITX	Intel® Core™ 2 Duo processor, 2.53 GHz	4 GB	LAN, SATA, USB, AMT, Adaptive-IO™	✔
Thunderbird-E3100	Mini-ITX	Intel® Core™ 2 Duo processor, 1.06 ... 2.2 GHz	8 GB ECC	LAN, ASF, SATA, USB, MiniCard	✔

LiPPERT Embedded Computers Inc.
5555 Glenridge Connector, Suite 200, Atlanta, GA 30342
Phone (404) 459 2870 · Fax (404) 459 2871
ussales@lippertembedded.com · www.lippertembedded.com

LIPPERT
THE EMBEDDED PC COMPANY

Thermal Management and Power Integrity in Tight Spaces

PCB thermal management has traditionally been seen as secondary to signal integrity. But due to tremendous evolution of power densities in transistors, PCB thermal management has now become a serious issue that must be considered early in the design.

by Syed W. Ali, Nexlogic

The advent of programmable logic devices (PLDs) increased circuit component miniaturization and increasing circuit densities on PCBs have introduced an exponential increase in power densities. Usually PCB thermal management is considered as secondary to signal integrity or power management. However, this task can no longer be ignored. If ample amount of work is not done at the layout level, and heat sinking techniques are not employed, heat generated by devices can result in parasitic effects to the circuit performance along with the reduced reliability and lifetime of the product.

Failure rate or mean time between failures (MTBF) measures component reliability. According to the Arrhenius equation, every 10°C rise in temperature above the component's maximum operating temperature can increase the failure rate by as much as 50%. Hence the reliability or life expectancy of a component is related to its operating temperature. Additionally, above certain temperatures, a component can be irreversibly destroyed. The task of thermal management is to make sure that electronic assemblies are maintained below the maximum rated temperatures of all the components. Maximum operating temperatures of typical components are 85°-125°C for capacitors and 125°-175°C for integrated circuits.

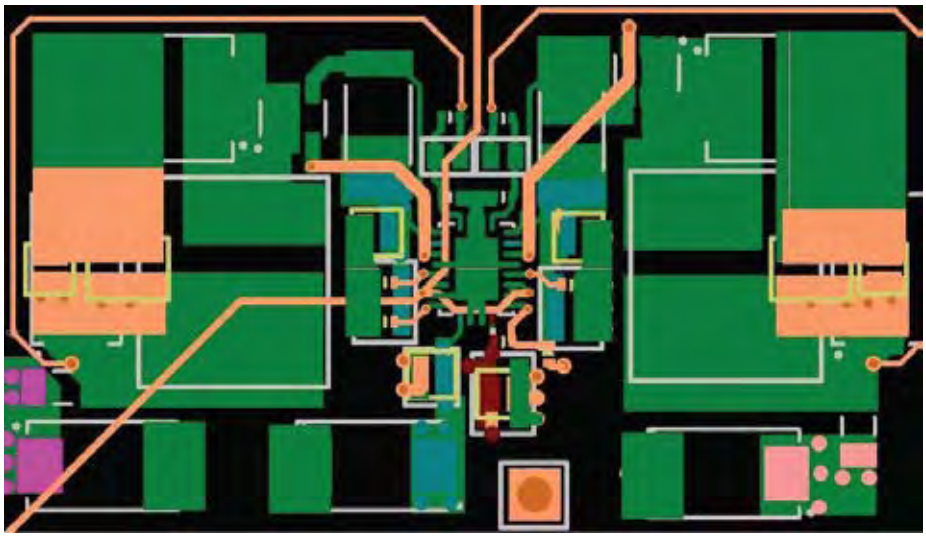

FIGURE 1

Layout of a triple switching power supply with 2 ounces of copper.

Secondly, self-heating can adversely affect device performance. The value of output voltages and the references of certain logic families can be affected by changes in junction temperatures. Such changes, if drastic, can alter the voltage levels. The logic family that experiences a significant effect is the emitter-coupled logic (ECL). It has also been noted that thermal offsets can trigger unhealthy effects on the performance of high power

microwave transistors. This is due to the fact that under high power conditions, device lattice temperature increases. This causes a drop in carrier mobility thereby negatively affecting device performance.

PCB Thermal Management

The first stage of thermal management is at the IC design level. Unfortunately, design engineers and layout designers have little or no power at this stage. Hence, it behooves the chip designer to consider any IC hotspots at design stage.

The circuit design level is the second stage of thermal management. The design engineer can take into account thermal characteristics of the circuit. For example, when designing the matching networks of power amplifiers, they can choose to implement harmonic tuning, which can reduce the amplifier's operating temperatures. The downside of this method is that it may increase the circuit complexity.

The third stage is the layout level for which both the design engineer and the PCB designers are responsible. It's in the best interest of the design engineer to identify PCB hotspots that need to be communicated to the designer, who in turn will use good layout practices with PCB thermal performance in mind.

There are three mechanisms by which heat can be transferred away from the power dissipating devices: conduction, convection and radiation. Conduction is the primary means by which heat can be transferred through a solid, which in this case is the PCB. Heat transferred by conduction is directly proportional to the cross-sectional area of the material and inversely proportional to the thickness of the material. Convection involves the transfer of heat between the solid and fluid, which in this case would be the PCB or component body and air. The rate of transfer is primarily a function of exposed surface area of the solid and the temperature gradient between the solid and the fluid. Radiation involves the transfer of heat as electromagnetic radiation, which does not require a medium

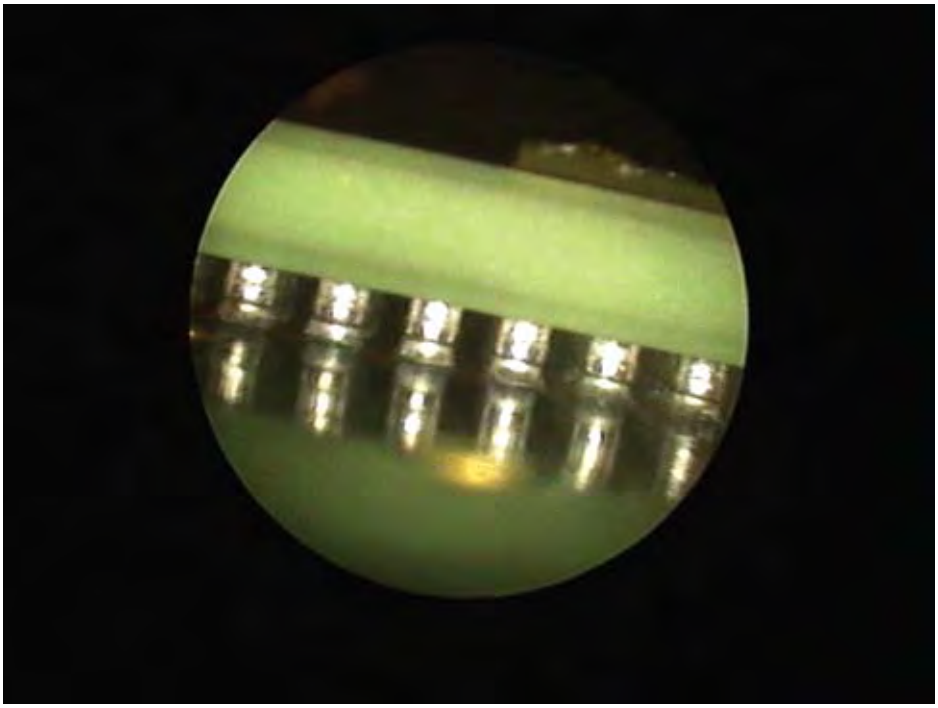

FIGURE 2

Microscopic lateral view of the BGA.

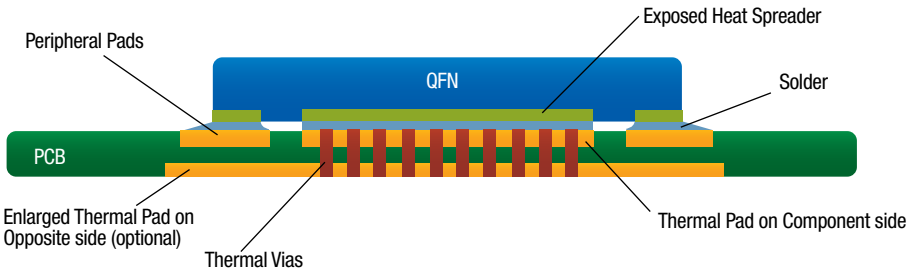

FIGURE 3

Cross-section of a typical QFN package mounted on a PCB.

to travel. The rate of this transfer depends on the surface area of the solid and the temperature. Standard PCB materials like FR-4 have low thermal conductivity, hence copper will usually dominate the heat flow on the PCB.

Layout Guidelines for Thermal Management

Certain procedures must be employed during the layout stage and in fact should be used as a standard practice while laying out boards with tight spaces with due regard for thermal performance. The designer should start with the power section of the board.The reasons are that switching power supplies are noisy and their thermal management can be challenging. That's because power dissipating components need to be close to each other for electromagnetic compatibility (EMC) purposes and space constraints. A balance is required between too close for thermal reasons and too far for EMC rea-

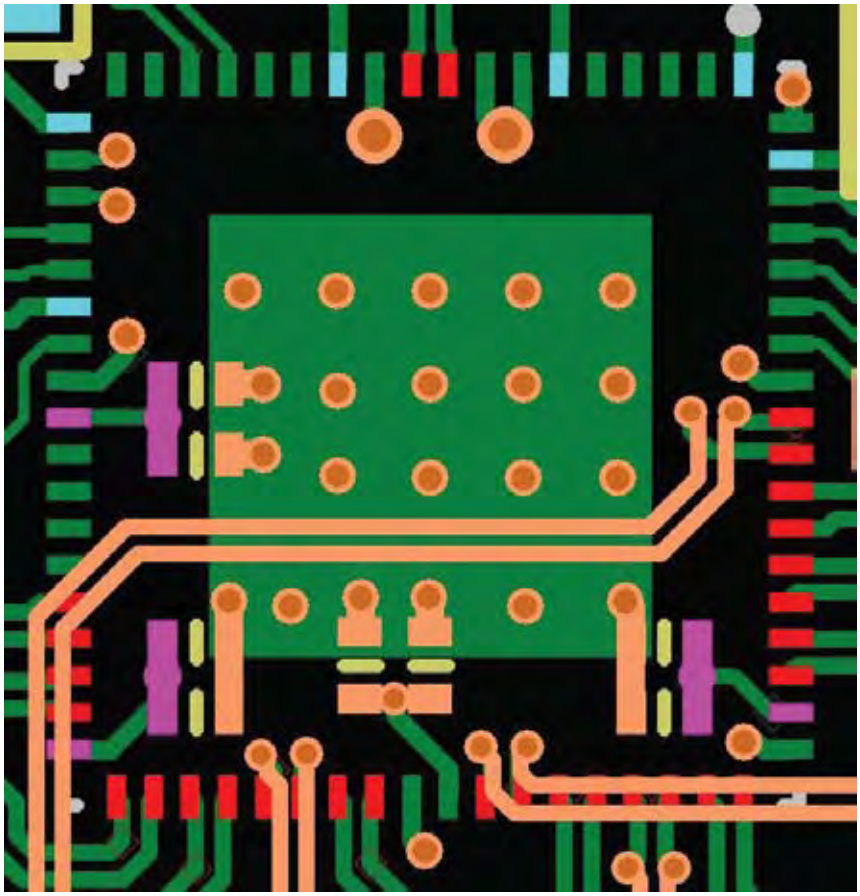

FIGURE 4

PCB Layout of a typical QFN package. The thermal vias have been connected to internal GND plane.

sons. It is always a good practice to have some distance between switching transistors and bulk electrolytic capacitors. One method that the designer would use to reduce power supply temperature is to create large areas of copper with the power supply circuitry.

This method does work some of the time, but overestimating the amount of copper required can be both ineffective as well as detrimental to circuit performance because it contributes to electromagnetic interference (EMI). Remember that perfect thermal conduction requires not only an ample surface area, but also a heat sink that is thick enough to radiate enough heat away. Unfortunately PCB copper is not thick enough for perfect conduction. Hence at some point just increasing the area results in diminished returns. Typically, for 1 ounce of copper, a copper area of more than 1 square inch will not give reasonable results. If space of about 3 square inches is available, then the ideal copper weight would be 2 ounces or more. Figure 1 shows a triple switching power supply layout. Note that the copper weight used in this case was 2 ounces.

Once the power layout is complete, and all the critical mechanical issues have been dealt with, the layout designer should start with the placement of the critical power dissipating devices. To avoid clustering the hot components on the PCB, power dissipating devices should be distributed across the PCB and positioned close to heat sinks. This step is very important for tightly packed boards.

After these critical components have been placed, heat sinking guidelines associated with all leaded and non-leaded devices like the ball grid arrays (BGAs), quad flat packs (QFPs) and quad flat no-leads (QFNs) must be followed. BGAs are relatively better in terms of thermal performance when compared to other types of devices because they allow for airflow beneath the device and can perform well when coupled with fan cooling methods. Figure 2 shows one side of the BGA as seen under a microscope.

RF power devices typically generate heat that must be conducted away from the device through its central pad. Generally housed in QFP or QFN packages, such devices have a row of perimeter pads around a larger central PAD encapsulated in a plastic body.

The pad lowers the thermal resistance of the package. The device runs cooler, which translates into higher reliability. The device center pad works both ways as the device ground as well as the primary conduction path to remove the package heat. The advantage of having the center pad at zero potential is that large copper planes will not contribute to EMI.

During the layout of such devices, the thermal pad on the PCB should ideally be created as the same size as the spreader underneath the package. This will help the device to self-center during reflow. Secondly, the device central pad requires an extensive thermal via structure that can route the heat out to the cooler regions of the PCB. The drill size of the thermal vias should ideally be below 0.3 mm. This will allow it to be completely filled during reflow. Filled vias help in two ways. The vias will efficiently conduct heat to the other side, and the vias will not starve the thermal pad of solder and pull down on the device, which may have detrimental effect on the device assembly. Figure 3 shows a cross-section of a typical QFN package mounted on a PCB.

It is a good idea to use as many vias as possible placed on a 25 mil grid on the thermal pad. These vias can be connected to internal ground planes, or better yet, to

FIGURE 5

Fin-based heat sink mounted for high-powered analog circuitry.

an exposed pad on the other side of the board, to draw the maximum amount of heat. Use of thermal vias along with the thermal pad can increase heat dissipation by as much as 70%. Such vias act like thermal shunts. A typical layout of QFN is shown in Figure 4.

If the steps mentioned above are not adequate to keep the junction temperatures of components below their rated values, then heat sinks need to be employed. In fact, most of the time the use of a heat sink is the most effective method of thermal management. The optimization criterion is to minimize the exposed heat exchanger's surface area while minimizing the weight of the heat sink and distance from the component to be cooled.

Heat sinks are made of thermally high conductive materials like copper or aluminum. Increasing their surface area by the use of fins removes the heat to the ambient. The interface between heat sink and the device is also important for efficient heat transfer. The connecting surfaces should be as even as possible and they should be attached by using a high conductive grease or elastomer (Figure 5).

Finally, the use of ground planes helps in a lot of ways including improved PCB thermal performance. Ground planes help get returns directly underneath their signals, which becomes mandatory at higher frequencies, and they provide a capacitive link to noisy signals and hence aid in crosstalk issues.

The PCB must be designed so that all the semiconductor devices on board are maintained at or below their maximum rated temperature. Thermal management can sometimes be problematic for dense boards employing fine pitch devices. But if certain layout guidelines are not followed and due considerations are not given to the PCB's thermal performance, one can end up with a product that has sub-par performance and reliability in the field. ◢

Nexlogic Technologies
San Jose, CA.
(408) 436-8150.
[www.nexlogic.com].

Prototyping Autonomous Robots with FPGAs

Producing a working prototype of a robotic design is an important step in getting to a finished product. The use of FPGAs presents the ability to rapidly try new ideas and add I/O elements in the process of completing a forward looking system to investors, managers and customers.

by Jamie Brettle, National Instruments

It may sound like science fiction, but robots are becoming more and more integral parts of our lives—from modern applications like cleaning floors and performing surgeries to future applications like autonomously operating cars. The fact is these applications that were once only in the realm of our imagination are quickly becoming a reality. However, in order for robots to gain more autonomy to perform these higher-level tasks, they require sensors to provide information about their environment such as audio, video and proximity to obstacles. As more and more sensors are required for robotic development, it becomes imperative that engineers use prototyping tools and platforms that allow them to design and iterate quickly. By developing functional robotic prototypes with Field Programmable Gate Arrays (FPGAs) and commercial off-the-shelf (COTS) hardware, engineers can test ideas, algorithms and I/O combinations efficiently to make next-generation robots a reality.

FPGA Technology

Fundamentally, FPGAs are reprogrammable silicon chips. Engineers can use FPGAs to implement application-specific custom hardware by using prebuilt logic blocks and routing resources to configure the chip. The process of developing custom FPGA hardware is a three-step process:

1. Develop digital logic using computer software to describe the desired functionality. This can be done in several ways including the use of text-based hardware description languages, preconfigured intellectual property (IP) or high-level graphical programming tools such as National Instruments LabView.
2. The description of the digital logic is then compiled down to a configuration file or bitstream that contains information on how the FPGA components should be connected to implement the digital logic. The tools for compiling and creating the bitstreams are typically provided by the FPGA hardware vendor.
3. Lastly, the FPGA bitstream is loaded onto the FPGA, which configures its operation as described in software. The FPGA is said to take on a new "personality," which can then be updated by following these three steps again, allowing it to be reconfigured as many times as necessary.

Traditionally, FPGAs have been viewed as a tool used solely by digital design engineers with an advanced understanding of hardware description languages. However, the rise of high-level and graphical programming tools for FPGAs is enabling an increasing number of domain experts to take advantage of the benefits FPGAs provide. This includes designers who are responsible for creating the electronic subsystems of robotic applications.

One of the main reasons developers choose FPGAs in their designs is because they provide the benefits of an application-specific integrated circuit (ASIC) combined with the flexibility of a software design. FPGAs feature hardware-timed speed and reliability, yet they do not require the high volumes typically needed to justify the large upfront expense of custom ASIC design. Reprogrammable silicon also has the same flexibility as software running on a processor-based system, but it is not limited by the

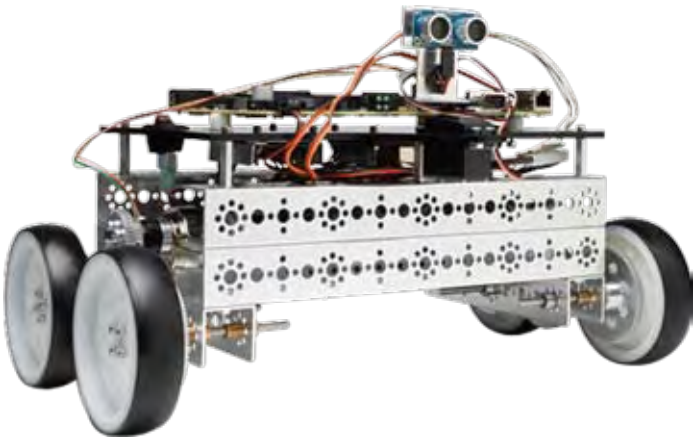

FIGURE 1

Engineers can get robotic prototypes running quickly with COTS hardware.

number of processing cores available. Unlike processors, FPGAs are truly parallel in nature, which means different processing operations do not have to compete for the same resources. Each independent processing task is assigned to a dedicated section of the chip and can function autonomously without any influence from other logic blocks.

Coupling an FPGA with COTS hardware and modular I/O drastically improves the flexibility of a design, allowing robotic developers the ability to quickly iterate on their prototypes.

As mentioned previously, the amount of sensory data required for a truly autonomous robotic system is increasing. Robotic developers face the challenging prospect of changing feature sets and unknown requirements in terms of I/O. For example, your robot may contain a single low-resolution sensor for capturing video data. However, throughout the lifecycle of your project, the price of the sensor may become cheaper, making it feasible to add additional cameras to the system. Robotic developers need to create on platforms that allow them to quickly swap out I/O to make the best possible designs. To mitigate risk, developers can choose COTS hardware that saves them board bring-up time while maintaining the flexibility to swap out different I/O configurations to meet their needs. This is especially important during time-sensitive endeavors like creating a first functional prototype (Figure 1).

FIGURE 2

FPGAs can help engineers create advanced, high-level perception and planning algorithms for autonomous vehicles.

Tips for Developing Your Prototype

Creating a prototype is a commonly overlooked step during development—often it is viewed as a cost and time impediment to creating the final product. However, by selecting a platform that allows engineers to rapidly produce and change designs, prototyping can become an invaluable phase during development. Prototypes are useful for proving the value of a particular design and, in emerging fields like robotics, can be essential to reaching the next stage of investor funding while better understanding customer needs.

One of the first tips to creating robotic system prototypes is to develop and validate FPGA IP piece by piece. Robotic applications can often be separated into "Sense," "Think" and "Act" components, which means a portion of the code is responsible for reading sensory input; an algorithmic portion decides what action to take; and lastly sends output to the actuators to drive the robot autonomously. By modularizing these three tasks, developers are able to decouple the I/O from the algorithms. This benefits the engineer by allowing them to test and validate code as independent units and make modifications to subsystems without impacting the stability of the entire prototype. The parallel nature of an FPGA design makes it an ideal candidate for this type of architecture as a variety of I/O components can be brought into and out of the FPGA independently, while the reconfigurable fabric allows continuous modifications to algorithm level decisions. Once sections of IP have been proven to work independently, the sense, think and act code segments can be integrated to form a fully functional embedded system for robotics that can be tested and validated as a whole (Figure 2).

A trap that engineers often fall into when developing a prototype is worrying about the cost of a system too early in the design process. For the hardware components of a robot, one of the great time sinks and potential failure points are early cost optimizations. Engineers can always try to trim the cost of a project by searching for cheaper components, smaller memory and fewer hardware connectors—however these benefits are most measurable once the project is in a moderately large deployment. The engineering time spent reducing the cost of the hardware bill of materials can actually result in a project failing before it even gets out the door. While cost is a factor, the goal of developing a prototype is to remain within striking distance of creating a profitable design.

An FPGA may cost more than an ASIC in some scenarios, but it has the flexibility to consolidate multiple components into a single package. For example, its reconfigurable nature allows developers to configure different types of peripheral resources on the FPGA such as serial communication interfaces. If at the onset of a project you are unsure of the required number of interfaces, an FPGA allows you to programmatically configure what you need, freeing up the additional cost of chips and potential redesign work. As a result, engineers can focus on proving the value of their design rather than the individual component costs. By developing a robot prototype using an FPGA, one can concentrate on securing the first set of customers and then work on cost optimization.

When designing the prototype it is worthwhile to develop it such that the maximum amount of IP can be reused in the final design. A couple of factors should be considered—anticipated deployment numbers and the market the robot will be deployed into.

Engineers should look at hardware platforms that can scale to the number of robots they plan on selling while still maintaining the profit margin that their company requires, such as using COTS hardware that has deployment options. In terms of market pressures, certain components need to be certified in different ways depending on the environment. For example, robots for space applications often require protection from radiation. In this case, it would be possible to prototype using an FPGA, understanding there are FPGA components that are radiation hardened such that the prototyped IP could be used in a final design.

The ability to demonstrate the behavior of a system in front of customers, investors and potential employees is one of the greatest benefits to creating a functional prototype of your robotic system. In order to make sure that your development team reaches the prototype stages, it is important to select tools that allow you to reach the end goal efficiently. By incorporating system architecture that is flexible in terms of I/O and programming capabilities, engineers can make necessary changes to the system based on client or investor feedback. FPGAs provide a rugged, stable and reconfigurable platform for integrating I/O and programming autonomy into robotics. Combined with the versatility of COTS hardware, the next generation of robotics has the ability to produce amazing results and become a more prevalent part of our everyday lives.

National Instruments
Austin, TX.
(512) 794-0100
[www.ni.com]

PCIe over Cable

I/O EXPANSION

- **PCIe x1, x4 and x8 adapters**
- **Transfer rate up to 20Gb/s**
  Gen 2 up to 80Gb/s
- **Software transparent**

HIGH-SPEED NETWORKING

with **SuperSwitch™**

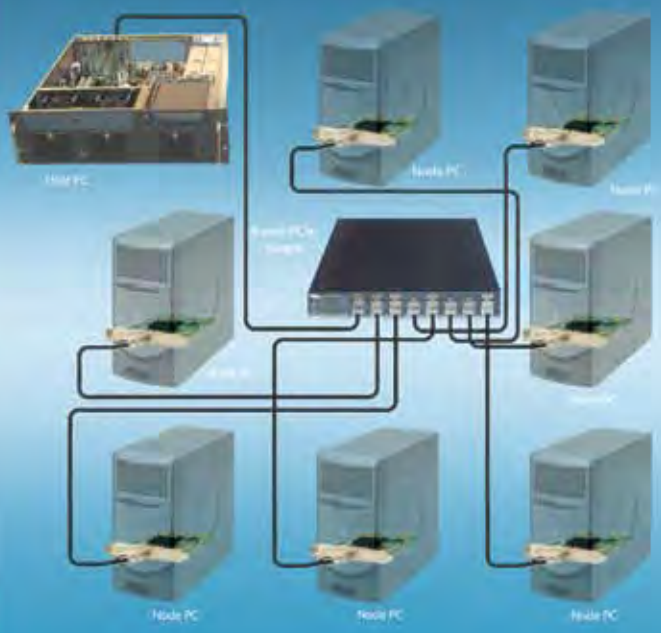

- **x4 at 10Gb/s**
  Gen 2 up to 80Gb/s
- **ExpressNet™ Software**
  Direct data transfer
  TCP/IP

MAX EXPRESS™

ONE STOP SYSTEMS

2235 Enterprise, Suite 110
Escondido, CA 92029
Tel (760) 745-9883 Fax (760) 745-9824

(877) GET-CPCI
(438-2724)

WWW.ONESTOPSYSTEMS.COM

Transitioning from Analog to Digital in Medical Designs

The move from analog to digital design in medical devices enables smaller size, lower power, greater noise immunity and lower parts count for powerful, portable solutions in the health care sector.

by Joseph Sankman, Microchip Technology

The burgeoning medical device industry stands to make significant advances with a new generation of microcontrollers that boasts high performance and low power consumption. These microcontrollers are integrated with a full complement of peripheral devices that meets the noise and accuracy requirements of medical devices. Traditionally, medical designs have relied on discrete analog circuit blocks, but digital microcontrollers are now powerful enough to assume the functions of their analog counterparts; sacrificing nothing in terms of speed and accuracy, while gaining reliability and smaller system volume. By moving to predominantly digital designs, cost minimization, design flexibility and time-to-market are improved, since software alterations are trivial in comparison to hardware redesigns. Additionally, there is a wide spectrum of available microcontrollers to meet the needs of next-generation, digital medical devices.

An Example Application

A pulse oximeter is an excellent application to demonstrate the shift to digital design. This noninvasive medical device shines red and infrared (IR) light through a patient's finger or ear, and measures the absorption at each wavelength to determine blood-oxygen saturation. In addition, the pulsation of a patient's heart is detectable, allowing the heart rate to be calculated (see "What is Pulse Oximetry?" p.xx). Portable oximeters have several critical requirements: low power dissipation for maximum battery life, small size that does not encumber the user, and high accuracy and repeatability. The last requirement is particularly important, since incorrect blood-oxygen saturation readings could endanger the health of the user.

As with any analog-based product, a gamut of factors affects the performance and design. Semiconductor products, such as operational amplifiers (op amps), are sensitive to temperature variations. Specifications, such as offset voltage and input offset current, drift with temperature and lead to measurement variations. $1/f$ and broadband noise also play a part in corrupting the accuracy of measurements. Another issue to consider is system size. The dramatic, ongoing reductions in the surface area of digital integrated circuits have not been matched by analog chips, giving digital-based devices an advantage in reducing system volume. Designers usually need many discrete components for analog designs, which could introduce problems with reliability and increased cost.

The transition to digital alleviates many of the problems that analog implementations incur. Generally speaking, digital designs are inherently smaller than analog designs. For example, analog signal processing, which commonly involves one or more op amps and a number of passive components, can be completely translated into software routines. Not only does this transition reduce the number of parts and system volume, it also frees more design time, since revisions only require code alternations. Digital circuitry has much better noise immunity compared to analog circuits and is not as susceptible to temperature drift. For wireless medical devices, the superior noise immunity of digital circuitry also improves the rejection of EMI noise.

Microchip's dsPIC33F digital signal controller (DSC) is well suited to realize a

digital pulse oximeter, because the digital signal processing (DSP) capabilities of the dsPIC33F eliminate the need for substantial analog signal-processing stages. Many resources available in books and journals are dated, and recommend the heavy use of analog circuitry. The only book available on oximeter design, *Design of Pulse Oximeters* edited by J.G. Webster, was published in 1997 and recommends the use of discrete sample-and-hold amplifiers, multiplexers and analog filters. Microcontrollers and DSCs have advanced tremendously in the 13 years since, rendering these solutions obsolete.

Figure 1 illustrates two different signal-flow paths in a pulse oximeter: the traditional, analog-centric design and a new, digital approach. In the traditional design, sample-and-holds separate and demodulate the red and IR signals, which bandpass filters (BPFs) then condition (shown in the top of Figure 1). Programmable gain amplifiers (PGAs) amplify the signals for the microcontroller A/D converter to sample. The microcontroller D/A converter also applies DC offsets to ensure that the signals are within the measurable range of the A/D converter. The signals are sampled before the bandpass filters, to determine their DC levels and adjust the red and IR LED brightness to equalize the DC levels. Measuring the DC level is important in order to set LED brightness and to simplify calculations. Since the red and IR DC levels are due to tissues that have constant absorption, if the two DC levels are equalized, they can be eliminated from the oxygen saturation calculation; only AC levels matter. Traditional analog designs rely on bulky DC pass filters with cutoff frequencies around 0.5 Hz. In digital designs, however, red and IR data can be sampled and filtered quickly with a moving average filter or other digital DC pass filter.

The Benefits of Digital Signal Processing

Utilizing a microcontroller with integrated DSP capabilities, known as a digital signal controller (DSC), reduces the entire stage into a single PGA. The DSC performs the signal separation, demodulation and filtering (bottom of Figure 1). The compression of this stage into the digital

What is Pulse Oximetry?

Pulse oximetry is a noninvasive method of determining blood-oxygen saturation and heart rate. Red and IR light (typically 660 nm and 940 nm) is pulsed through a patient's finger or ear, and the absorption is measured by a photodiode placed on the opposite side of the patient's digit or ear. The small current developed in the photodiode is amplified into a voltage read by a microcontroller. The signals then undergo processing to determine blood-oxygen saturation and heart rate.

The ratio of red and IR light is proportional to blood-oxygen saturation over a range useful for consumers: 100% saturation corresponds to R = 0.5 and 80% saturation corresponds to R = 1.0.

$$R = \frac{AC_{RED} / DC_{RED}}{AC_{IR} / DC_{IR}}$$

The DC levels are due to tissue as well as arterial and venous blood that should not change in absorption. The AC levels are due to pulsating blood flow that will change based on oxygen saturation.

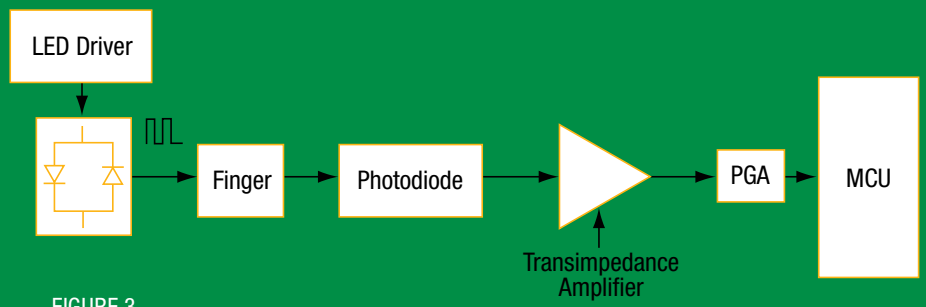

FIGURE 3

Block diagram of a modern, digital pulse oximeter. Typically, the microcontroller provides feedback to the PGA and LED driver to control gain and LED brightness, respectively.

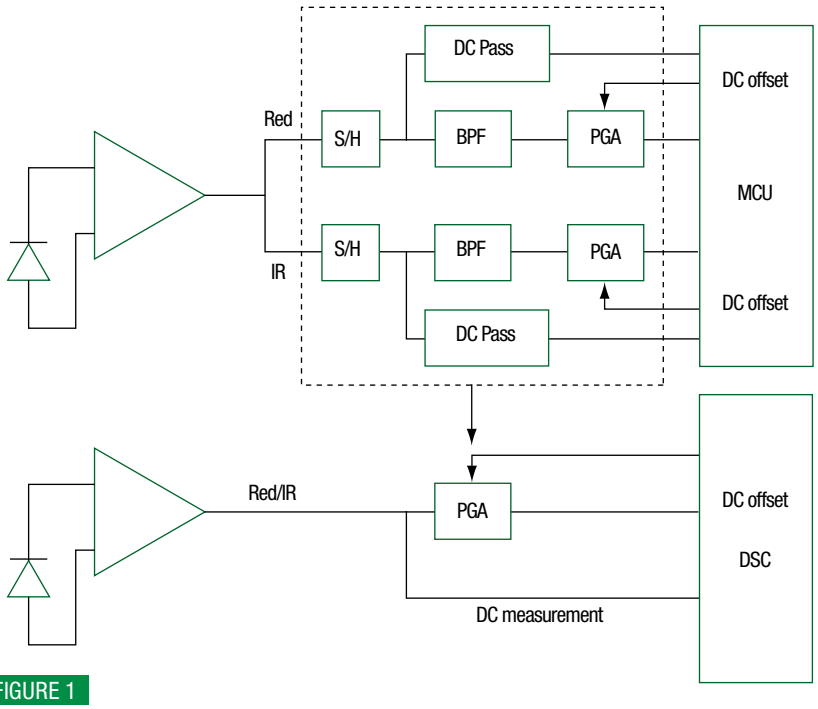

FIGURE 1

Signal flow path, from photodiode to microcontroller with analog signal processing (top) and digital signal processing (bottom). The DSC topology drastically reduces the number of discrete components by relying on software instead. Note: anti-aliasing filters not shown.

domain has the greatest impact on system volume, reliability and cost, since multiple discrete ICs, op amps and passive components are eliminated. The elimination of these sample-and-holds and analog filters helps reduce variation in the measurements that stem from temperature drift in the components, and also improves overall system reliability by reducing the number of components that could eventually fail.

Moving to digital has an impact on power consumption that is dependent on several factors. If the analog filters, sample-and-holds and PGAs use tens of miliamperes, shifting to digital will likely improve power usage for a microcontroller operating at a few MIPS. Many microcontrollers also have advanced power-saving modes that can reduce the power consumption of the microcontroller to mere nanoamps while idle.

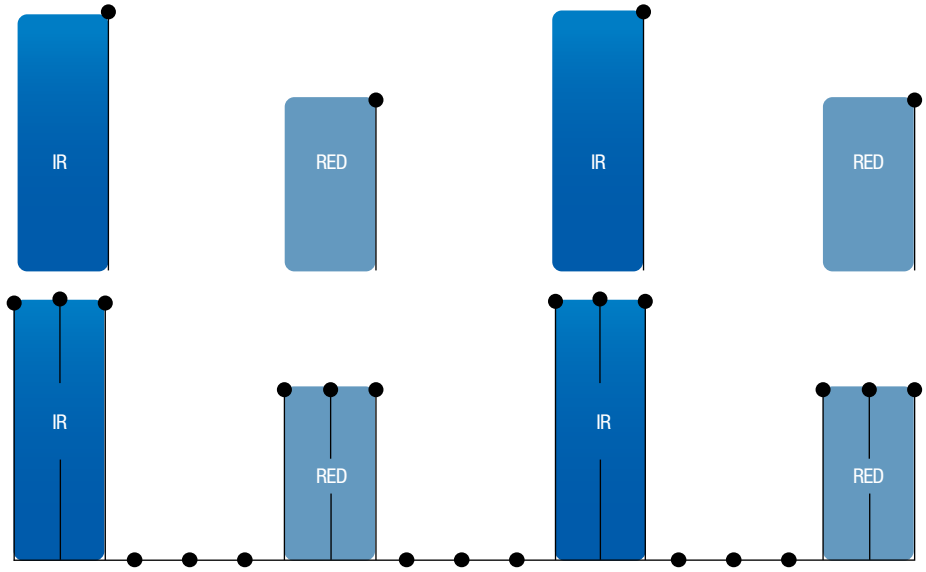

FIGURE 2

Diagrams for low sampling rate (top) and high sampling rate (bottom). Lower sampling rates require lower power dissipation, but at the cost of greater aliased noise. Note: the timing of red and IR pulses varies among designers. In some designs, red and IR pulses have no separation, and delays occur in between pairs of red and IR pulses.

Analog vs. Digital Filtering

Sample-and-hold circuits are a composite of op amps, switches and capacitors that require a nontrivial amount of system volume, if they are built discretely. Most modern microcontrollers have sample-and-hold amplifiers integrated with their A/D converters, which obviate discrete sample-and-holds. A designer might make the case that analog filtering, such as in the top of Figure 1, is worth the volume, power and reliability trade-offs, because the aliasing incurred by an A/D converter

can be avoided; with the filtering occurring before the signals are digitized. However, this is not true; the sample-and-holds still alias noise. Rather than the aliased signals appearing in the digital domain after digitization, they appear in the analog domain after the sample-and-holds.

Since oximeter bandpass filters have very low cutoff frequencies (0.5 Hz and 5 Hz), large-valued components may be required. In addition, narrow stopband width and strong stopband attenuation require cascading op-amp stages. Because analog filters require a large number of components and significant system volume, the trade-offs that digital filtering incurs are minimal. Pulse oximeter signal processing has several flavors, and if a designer needs a DC-reject, DC-pass, low-pass, or high-pass filter, Microchip's free DSP software libraries support the Finite Impulse Response and Infinite Impulse Response realizations of these filters.

Data Acquisition Options

Designers have several choices regarding data acquisition. If there is minimal high-frequency noise, sampling at the end of each pulse of red or IR light is sufficient, as in the top of Figure 2. However, if the accuracy and precision requirements are stringent, and aliased high-frequency noise is a problem, a designer can increase the sampling frequency to mitigate aliasing, as in the bottom of Figure 2. Of course, the trade-offs are higher clock frequency and higher power consumption, which are required to support the larger data set and higher order filters. Pulse oximeters commonly have a blood-oxygen saturation accuracy of ±2%, which a 10-bit A/D—often integrated with a COTS microcontroller—can support. For higher precision, designers may want to use a 12-bit A/D, which are also available in COTS microcontrollers.

As a result of a new generation of microcontrollers and DSCs, medical devices are trending toward digital signal processing and the elimination of analog signal conditioning. As the healthcare market shifts toward early diagnosis and personal health monitoring, reliable, cost-effective medical devices will be necessary. The transition to digital design from traditional analog will enable consumer medical devices to proliferate, and predominantly digital designs will allow in-the-field upgradability and low-cost redesigns. As the market for consumer medical devices expands, making the leap to digital will be necessary to deliver high-performance products at a low cost with expedient time-to-market. ◢

Microchip Technology
Chandler, AZ.
(480) 792-7200.
[www.microchip.com].

PRODUCTS & TECHNOLOGY

CPU Cooler Improves Airflow, Water Block and Heat Dissipation Capacity

A CPU cooler is designed to provide an extremely efficient cooling solution for CPUs from Intel and AMD. The Hydro Series H70CPU cooler from Corsair is an evolution of the Hydro Series H50, with several significant upgrades that enable it to deliver even greater cooling performance. These upgrades include a double-thickness (50 mm) radiator with higher heat-exchanging capacity and a pump/cold plate unit with increased efficiency. The H70 also features two 120 mm speed-switchable cooling fans in a push-pull configuration to provide increased airflow at low noise levels.

Like the H50, the Hydro Series H70 provides the benefits of water cooling in a sealed and pre-filled unit, with no maintenance required. The low profile cold plate is extremely space efficient, and is very low in mass compared to heat pipe based solutions, putting less stress on the system's motherboard. The H70 includes mounting hardware for most common AMD and Intel CPUs, and, unlike some competitive CPU coolers, includes all necessary fans for high performance operation. The Corsair Hydro Series H70 CPU Cooler is supplied with a two-year warranty, and is backed by Corsair's customer service and technical support.

Corsair, Fremont, CA. (510) 657-8747. [www.corsair.com].

Full Size PICMG 1.3 SHB Supports Intel Core i7/i5/i3 Processors and Q57 Express Chipset

A new PICMG 1.3 full-sized System Host Board (SHB) supports next-generation Intel Core i7/i5/i3 processors at clocks speeds up to 3.33 GHz, and dual-channel DDR3 1066/1333 MHz memory up to a maximum of 8 Gbytes in two DIMM slots. The NuPRO-E330 from Adlink Technology is equipped with the Intel Q57 Express Chipset to provide higher data transfer rates. The NuPRO-E330 is an appropriate solution for applications requiring high-performance computing, data security and maximum data access bandwidth such as industrial control, automation and industrial vision.

For modern industrial applications where remote management and data protection are critical, the NuPRO-E330 supports Intel Active Management Technology 6.0 (Intel AMT 6.0). Intel AMT 6.0 allows users to remotely manage assets, reduce downtime and minimize on-site visits. Utilizing out-of-band system access, Intel AMT 6.0 allows users to remotely detect, isolate and recover systems, even when they are powered off. The NuPRO-E330 provides additional data security with six SATA II interfaces supporting Intel Rapid Storage Technology for RAID 0/1/5/10 functionality.

The NuPRO-E330 features the latest Intel Core i7 processor supporting Intel Turbo Boost Technology, which allows processor cores to run faster than the rated operating frequency if the processor is operating below power, temperature and current specification limits. For more effective use of processor resources, Intel Hyper-Threading Technology provides thread-level parallelism on each processor core. The NuPRO-E330 also offers DDR3 memory support at higher data transfer rates with less power consumption.

Adlink's NuPRO-E330 provides a wide range of storage, I/O and expansion connectivity. It supports one PCI Express x16 and four PCI Express x1 (or one x4) links to the backplane. To provide the latest in operation system support, the NuPRO-E330 has been verified for both 32-bit and 64-bit versions of Windows 7.

ADLINK Technology, San Jose, CA (408) 495-5557. [www.adlinktech.com].

3U VPX Backplane Equipped with RF Connectors

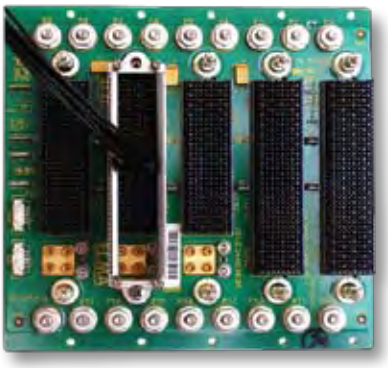

A new 3U VPX Backplane comes with connectors for RF and analog signals. The use of RF interconnects is expected to be an important issue for many military and aerospace applications where the VITA 46/65 specifications for VPX are highly targeted.

The VPX Backplane from Elma Bustronic is designed to meet OpenVPX design considerations. Work is underway on a VITA 67 specification for RF signals over VPX, and Bustronic plans to develop a version that will meet the specification when it is finalized. The backplane features a 5-slot mesh routing configuration with the RF connectors in three slots. The 4-cavity RF connector is installed in the lower half of the standard J2 connector. This corresponds to rows 9-16 of slots 3, 4 and 5. The coax interconnects provide better signaling, higher data rates, less noise and less crosstalk than discrete wire.

Bustronic also offers a wide selection of VPX and OpenVPX Backplanes in 3U and 6U heights. The company has the industry's largest offering of VPX accessories including 3U and 6U load boards (both convection- and conduction-cooled styles), test extenders, air baffles, RTMs and SerDes test devices. Pricing for the 3U VPX Backplane including the RF connectors is under $2,500 depending on volume and configuration.

Elma Bustronic, Fremont, CA. (510) 490-7388. [www.elma.com].

1U Network Appliance Based on the Quad-Core Intel Xeon 3400 Series

A 1U Communications Appliance is based on the quad-core Intel Xeon processor X3400 or L3400 (formerly codenamed Lynnfield) or Intel Core i5 or Core i3 processor. The CAR-4003 from American Portwell is a solution targeted for security applications such as enterprise Internet security, firewall, WAN optimization, unified threat management, network access control, network behavior analysis and triple play broadcasting for both the enterprise and medium-sized offices.

Powered by the Intel Xeon processor 3400 series, the CAR-4003 supports four DIMM slots with high-speed 1333 MHz dual-channel DDR3 ECC memory; dual x8 PCI Express Gen2 slots (8 Gbyte/s bi-directional) for connecting expansion modules; expansion capabilities include two PCIe x8 interfaces for a modular bay for Portwell's ABN/NIP module product family; fiber and copper port connections including dual-port 10G readiness (Intel 82598EB, Intel 92599ES with SFP+ Interface and Intel 82599EB 10Gbase-T Copper Interface) and PCIe Gen2 Quad Port GbE and SFP (Intel 82580) module.

In addition, the CAR-4003 provides Intelligent Platform Management Interface (IPMI) v2.0 functions for remote management, including IPMI over LAN, IP over KVM (iKVM), Serial Over LAN (SOL) redirection, Event logging and OS independent Hardware Health Monitoring. IPMI helps lower the overall costs of server management by enabling customers to save time, maximize IT resources and potentially manage multivendor environments in the same way.

The IPMI consists of a main controller called the baseboard management controller (BMC) and other management controllers distributed among different system modules that are referred to as "satellite" controllers. User can access IPMI functionality through the command line with the IPMItool utility either in-band or out-of-band. Additionally, user can generate an IPMI-specific trap from the Web interface, or manage the server's IPMI functions from any external management solution that is IPMI v1.5 or v2.0 compliant.

The IPMI operates independently of the operating system and allows administrators to manage a system remotely even in the absence of an operating system or the system management software, or even if the monitored system is powered off, while connected to a power source. IPMI can also function after the operating system has started, and offers enhanced features when used with system management software. IPMI prescribes only the structure and format of the interfaces as a standard.

American Portwell Technology, Fremont, CA. (510) 403-3399. [www.portwell.com].

Certified Modules Ease Adding Wi-Fi to Device Designs

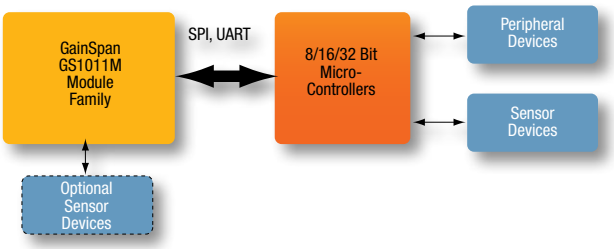

A family of certified modules offers a quick, easy and cost-effective way for device and appliances manufacturers to add Wi-Fi capabilities to their products. The GA1011M module from Gainspan provides a serial UART or SPI interface, enabling connection to any embedded design utilizing a 8/16/32-bit microcontroller via simple commands. The GS1011M is a suitable solution for organizations with limited or no Wi-Fi or RF expertise, as it not only dramatically reduces RF design time, but also removes the burden of testing and certification, allowing customers to focus on their core application, product or expertise. The module supports data rates up to 11 Mbit/s, is compliant with 802.11b and meets regulatory and Wi-Fi Alliance requirements.

Multiple software configurations are available for the stack running on the module. For applications utilizing a small 8-bit microcontroller host, the module supports a serial to Wi-Fi function and runs the full Wi-Fi and TCP/IP networking stacks, completely offloading the host. For applications utilizing more powerful microcontrollers, the networking stack and services can reside on the host while the module provides the IP to Wi-Fi functionality. In addition, it supports WEP/WPA/WPA2 security, Adhoc as well as Wi-Fi Protected Setup (WPS) for ease of provisioning.

Gainspan, San Jose, CA. (408) 454-6630. [www.gainspan.com].

Low-Power Processor Brings 64 Bits to COM Express

Via Technologies, Inc, a leading innovator and developer of embedded silicon and platform technologies, today announced the addition of the Via COME8X80 module to its COM Express portfolio.

The first COM Express type product to incorporate the 64-bit Via Nano E-Series processor is targeted at IPC and larger OEM customers in a range of embedded segments including gaming, healthcare and industrial automation. Customers can also take advantage of a specially developed mutli-I/O baseboard for evaluation purposes, as well as Via's technical assistance in developing custom baseboard designs. The Via COME8X80 from Via Technologies is available with a choice of either 1.3+ GHz or 800 MHz Via Nano E-Series processors, combining native 64-bit and virtualization support with a high-performance superscalar architecture in a low power thermal envelope.

The COME8X80 also integrates the unified Via VX800 media system processor, bringing Via Chrome9 integrated graphics, support for VGA and dual channel LVDS displays, video acceleration for MPEG-2, MPEG-4, WMV9 and VC1 video formats, plus a VMR capable HD video processor. All Via components have a guaranteed longevity of seven years. Additional highlights include three PCI and one x4 plus two x1 PCIe slots, up to 2 Gbytes of DDR2 SO-DIMM memory, 10/100 Mbit/s Ethernet, up to six USB 2.0 ports and up to two SATA devices plus one IDE.

VIA Technologies Fremont, CA. (510) 683-3300. [www.viaembedded.com].

Mass Storage Mezzanine Offers Fast Secure Erasure and Write Protection

A mass storage mezzanine for securely erasing and write-protecting sensitive data in harsh environments such as those often found in military and defense applications is built on open standards-based platforms like CompactPCI, ATCA, VME and OpenVPX. The new, high-capacity Secure PMCDisk mezzanine from Elma Electronic Systems supports 10 major defense agency secure erasure procedures via a front panel push button with an LED indicator to confirm data erasure. It also enables write protection via a front panel toggle switch.

Specific defense agency standards supported by the Secure PMCDisk include fast clear/initialize; clear; DoD NISPOM 5220.22-M and 5220.22-M-Sup 1; NSA/CSS Manuals 130-2 and 9-12; AR 380-19; NAVSO P-5239-26; AFSSI-5020; and IRIG 106-07.

The new RoHS-compliant Secure PMCDisk ships complete with a 2.5" solid state flash drive. Current capacities support up to 128 Gbytes. Software drivers are available for VxWorks, Linux and Windows. It can be used on any board with an IEEE1386.1-compliant PMC site, such as CompactPCI, VMEbus, OpenVPX, VXS and ATCA, replacing external hard drives or disk modules that require additional backplane slots or external connections.

The storage module supports SMART (self-monitoring analysis and reporting technology), which detects remaining useful drive life in order to avoid system failures that lead to catastrophic shutdowns, a key factor for applications requiring data safety. In addition, the SLC (single level cell) 2.5" SATA solid state storage device further increases data reliability and retention. The ruggedized mezzanine withstands shock to 1,500G at 0.5 ms half-sine and vibration to 16.4G rms (10 Hz to 2,000 Hz random). Pricing starts at $2,700 in low quantities.

Elma Electronic Systems, Fremont, CA. (510) 490-7388. [www.elma.com].

Family of IP Display Engines for Networked Video Connectivity

Targeted at networked video connectivity solutions for mission-critical systems, a new family of video receivers is designed specifically for viewing stations and specialized processing appliances on high-performance video networks. The vDisplay family of IP engines from Pleora is comprised of compact, purpose-built hardware that allows high-resolution video streams on GigE (Gigabit Ethernet) networks to be displayed directly on monitors, in real time, without the need for a PC. They can also be used for real-time video capture in specialized processing appliances. The engines comply fully with the open, global GigE Vision and GenICam standards.

By replacing PCs with ultra-efficient hardware, vDisplay engines shrink the size, lower the cost, reduce the power consumption and improve the reliability of viewing stations and video processing appliances. They are suited for OEMs and integrators building high-performance video products and systems for the military, medical and manufacturing sectors.

Pleora will initially offer the vDisplay HDMI-Pro IP engine. The HDMI-Pro engine converts streaming IP video to standard HDMI/DVI (High-Definition Multimedia Interface/Digital Visual Interface) formats for real-time display on off-the-shelf monitors. It auto-senses the display capabilities of the attached monitor and adjusts the image formats and resolution of the incoming video stream to match the monitor's refresh rate and resolution. The HDMI-Pro engine is available as a compact OEM board set or a small enclosed unit.

Pleora Technologies, Kanata, Ontario, Canada. (613) 270-0625. [www.pleora.com].

Sound and Vibration Software Supports Data Translation IEPE Modules

A new software application for sound and vibration measurement provides an easy-to-use interface for individual channel configuration and display for performing vibration analysis. The VIBpoint Framework Application from Data Translation now supports the company's DT9837 and DT8837 series of sound and vibration measurement products. Functions include: spectrum/spectral display, power/magnitude, RMS/peak, dB/linear and more. This hardware/software combination supports both USB and Ethernet (LXI-compliant) vibration analysis and monitoring applications including noise emission monitoring, predictive maintenance and shock analysis.

VIBpoint Framework Application supports all the features of the hardware modules including discovery and selection of available hardware, configuration, loading and saving of hardware configurations as well as individual configuration of each channel for analysis and display. It supports the hardware's per channel FFT parameter configuration including: spectrum/spectral density, power/magnitude, RMS/peak, dB/linear and more for maximum vibration analysis flexibility. It also includes the FFT averaging modes for linear, exponential and peak, and it can save, display or analyze data with Excel. A 14-day free trial of VIBpoint Framework Application is available for download, or a license key can be purchased for $995.

Data Translation, Marlboro, MA. (508) 481-8620. [www.datatranslation.com].

Big Performance, Little Footprint: Rugged System Beats ATR-Short

A fast, powerful, rugged computer system packs features and options into a small, easily mounted package that is one-quarter the size and weight of a ½ ATR-Short package. With a Core2 Duo processor running at up to 2.256 GHz and 8 Gbytes of main memory, the Golden Eye II from General Microsystems is a super lightweight, rugged system that brings high performance in harsh environments together with advanced graphics processing capabilities, making it suitable for vehicular and avionic programs. Golden Eye II is available in several standard configurations: S802-R with one removable Solid State Drive (SSD) offering up to 500 Gbytes of storage, and S802-R4, which has four removable SSDs for a total of 2 Tbytes of storage. S802-R4 units are available with a combination of drive bays for removable SSD and PCMCIA slots to accommodate Army legacy systems.

Golden Eye II in its S802-R configuration measures only 5.25"x5.25"x2", and weighs just 2.5 pounds. The unit literally can replace four 6U cards offering equivalent or better performance, and its peak drain on the system's power bus never exceeds 25 watts. GMS's patent-pending cooling design enables its systems to easily operate in environments from -40° to +85°C. Because of its small footprint, about one-quarter the size of a standard ½ Air Transport Rack (ATR)-Short, it is well suited for UAVs and vetronics applications. Golden Eye II was designed for rugged, military applications and is compliant to MIL-STD-810F, MIL-STD-704E and MIL-STD-461E. Completely validated and certified, it also uses GPS for time stamping every data packet within the stringent Army requirements of less than one micro second.

With dual pipe support for high-performance RGB video (2048x1536 at 32-bit color), DVI-D and NTSC/PAL composite video, Golden Eye II offers a broad range of video output options.

Four-channel video capture, two Dual Redundant 1553 channels, Wi-Fi and Bluetooth give the unit a huge range of options for communications and interface. Software support includes Windows XP/XPE/7, Linux and VxWorks and Network Attached Storage (NAS). It can replace four servers with one Trusted Platform Module (TPM) virtual machine running up to four different operating systems. Single unit price (S802-R) starts under $10,000.

General Microsystems, Rancho Cucamonga, CA. (909) 980-4863. [www.gms4sbc.com].

Solid State Drive in DIMM Form Factor Expands Storage Options

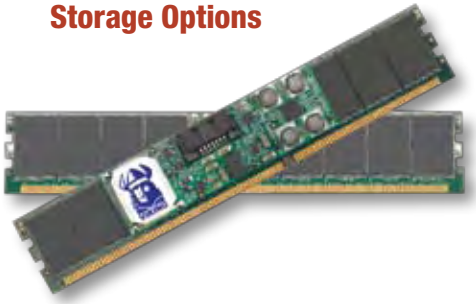

An enterprise class SATA II Solid State Drive (SSD) in an industry-standard DDR3 240-pin DIMM form factor is now available. This technology from Viking Modular Solutions allows users to significantly increase the capacity and storage performance of their existing server, storage or cloud computing solutions. It also allows system designers new options and greater flexibility when creating new server designs—and not coincidentally, it also lends itself to embedded designs.

The SATADIMM offers a high-performance and high-availability solid state drive in the DDR3 240-pin DIMM form factor (JEDEC MO-269). By adding SATA to an existing DDR3 DIMM socket, Viking Modular has delivered the performance and power savings desired from an SSD at a fraction of the space required to house a typical 2.5" SSD. This cost-effective solution provides dramatically increased capacity and performance.

Features of the SATADIMM include best-in-class sequential and random performance (30,000 IOPS), intelligent write management techniques for optimized endurance and protection against catastrophic flash failures, and power fail data protection enabled by super capacitor integration on the DIMM. SATADIMM is available in 50 Gbyte, 100 Gbyte or 200 Gbyte capacities.

For current server storage appliances, the SATADIMM has been designed to take advantage of any available 240-pin DDR3 DIMM socket. The SATADIMM SSD derives its power from the 1.5V supply to the DIMM socket, and data transfer is enabled by using a standard SATA cable. For new designs, system architects will appreciate that the SATA data signals can be routed directly to the socket, thus eliminating the need for any cables.

Viking Modular Solutions, Foothill Ranch, CA. (949) 643-7255. [www.vikingmodular.com].

2U Rackmount Network System Supports Xeon, Core i7 and up to 32 x GbE LAN

A 2U rackmounted hardware platform designed with a redundant power supply is targeted for network service applications. Built with Intel Embedded IA components with warranty of longevity, the PL-80160 from Win Enterprises supports Intel quad-core processors with Intel Hyper-Threading technology, including the Xeon, Core i7, Core i5, Core i3 and Pentium Dual-Core processors.

The platform supports four unbuffered ECC or non-ECC DDR3 1066/1333 MHz DIMM sockets with memory up to 16 Gbytes. In order to provide the best network performance and utilization, the powerful storage interfaces include one 3.5" SATA HDD and CompactFlash. The optional onboard Cavium Nitrox PX cn16xx security coprocessor supports multi-security protocol commands that can offload the CPU thus increasing overall system throughput performance.

This platform affords 8 GbE and max to 32 GbE Ethernet ports via PCI-E by8 on the front panel. To prevent network problems when the platform shuts down, PL 80160 supports two segments of LAN bypass function through WDT and GPIO pin definitions. The front panel also has dual USB 2.0 ports, one RJ-45 console port and LED indicators that monitor power and storage device activities for local system management, maintenance and diagnostics. In addition, the PL-80160 supports one PCI-E x8 slot, one PCI-E x8 Golden Fingers, one PCI slot and is RoHS, FCC and CE compliant. Software support includes Windows XP, Server 2003, Server 2008, Windows 7 and Linux. OEM quantity pricing for PL-80160 begins at $1,391.

WIN Enterprises, North Andover, MA. (978) 688-2000. [www.win-ent.com].

SAS Solid State Drive with Enterprise-Grade Multi-Level Cell Technology

A new development in flash memory is a serial-attached SCSI (SAS) solid state drive (SSD) equipped with enterprise-grade multi-level cell (E-MLC) NAND flash technology. The 2.5" XceedIOPS SAS SSD from Smart delivers superior endurance in enterprise environments due to the use of the latest 34nm E-MLC NAND flash technology. Specified endurance for E-MLC flash is 30,000 program/erase (P/E) cycles, whereas competing commercial MLC (C-MLC) technologies typically demonstrate only 1,500 to 5,000 P/E cycles. Endurance is further enhanced by an advanced wear leveling algorithm combined with advanced data management hardware, delivering the industry's lowest levels of write amplification while also significantly boosting performance. As a result, the Smart XceedIOPS SAS SSD family of products will deliver a minimum of five years of operational life in environments that write up to 10x the device's total data capacity per day.

Specifically optimized for high-performance enterprise storage and server systems, the XceedIOPS SAS SSD achieves up to 26,000/20,000 IOPS random read/write and 250/230 Mbyte/s sustained read/write. Available in 100, 200 and 400 Gbyte capacities, the new XceedIOPS SAS SSD offers high reliability and data integrity due to extensive error-correction and detection capabilities, multi-level data-path and code protection, data-fail recovery and data-integrity monitoring. Designed to fit the restricted power envelope of enterprise storage environments, the XceedIOPS SAS SSD incorporates staggered power-on support. In addition, the new XceedIOPS SAS SSD supports long data sector (LDS), which allows host transfer sizes of 512, 520 and 528 bytes.

SMART Modular Technologies, Newark, CA. (510) 623-1231. [www.smarm.com].

Battery-Free Energy Modules Back Up Critical Data During Power Loss

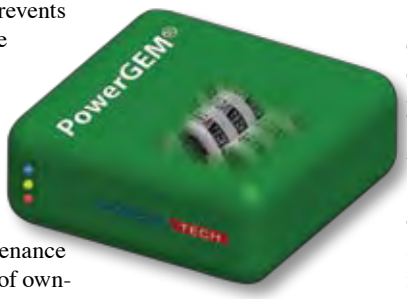

A new battery-free intelligent power subsystem prevents the loss of critical data due to a power interruption. The PowerGEM (Green Energy Module) power subsystem from Agiga Tech uses ultracapacitors (or "ultracaps") to provide temporary power in the event of a system power loss.

The PowerGEM modules allow designers to avoid problems of battery-powered energy sources, including environmental impact, increased design complexity, maintenance and conditioning, short operating life and a high total cost of ownership. PowerGEM, when paired with an AgigaRAM memory module, provides a complete secure and reliable non-volatile memory subsystem for mission-critical data back-up. In a related announcement, AgigA Tech has introduced a high-speed and high-density DDR3 solution, ranging from 1 Gbyte to 8 Gbyte.

When used as a write cache, AgigaRAM provides a performance-boosting building block while guarding against power failures and data loss. In addition to providing power, PowerGEM manages the charging/discharging, cycling, wear-monitoring and other module details to ensure long life and high reliability.

The entire AgigaRAM Non-Volatile System (NVS) implements the AgigaSafe control protocol, a simple-to-use host-controlled $I2^C$ programming interface developed by AgigA Tech that ensures safe, reliable, secure operation at the system level. This protocol allows fine control over internal system functions such as managing the NAND Flash while providing precise health monitoring and tracking. The complex system-readiness system has also been reduced to a single "Good To Go" (or GTG) signal.

AgigA Tech, Poway, CA. (858) 375-4530. [www.agigatech.com].

Auto Tuning for Easy-to-Implement Capacitive Touch Sensing Solution

Replacing the billions of mechanical buttons in mobile handsets, laptops, consumer electronics, white goods, automotive applications and virtually any system that has a mechanical button or switch, is the job of capacitive touch sensing. Now the new CapSense capacitive touch-sensing controller from Cypress Semiconductor enables designers to achieve mechanical button replacement (MBR) without having to write firmware or learn to use new software tools. The CapSense Express eliminates the requirement for system tuning. Cypress's accompanying design toolbox provides detailed resources to ensure optimal interface performance, and advanced system debug features allow taking designs directly to production for significantly shorter time-to-market. The controller delivers robust touch-sensing with ultra-low power consumption to extend battery life in a wide range of handheld products.

The hardware-configurable CY8C-MBR2044 CapSense Express Mechanical Button Replacement controller operates from 1.7 to 5.5V and offers low overall power consumption with supply current in run mode as low as 15 uA per button, and industry-best deep sleep current of 100 nA. The devices offer reliable operation in harsh sensing conditions, and Cypress's patented CSD (CapSense Sigma Delta) sensing method provides superior immunity to conducted and radiated noise. With SmartSense auto-tuning, the device dynamically optimizes the baseline and detection threshold and adjusts for the optimal capacitance sensing range at power up and during runtime as environmental conditions change

Eliminating the need to tune the UI subsystem is a significant advantage for large and small manufacturers alike, as it saves engineering time and yield loss that can occur with even slight variations in manufacturing tolerances. This savings is greatly multiplied for customers with a global factory footprint and supply chain. SmartSense auto-tuning can eliminate the need for additional test steps currently required with competing solutions to address the vendor-to-vendor variations in PCBs and overlays.

Cypress Semiconductor, San Jose, CA. (408) 943-2600. [www.cypress.com].

Embedded SATA SSD Designed to the JEDEC MO-297Spec

A case-less, small form factor and rugged SATA SSD is an embedded solid state drive solution that delivers outstanding performance in a standard form factor—and is less than half the size of a 2.5" SSD. With its 3 Gbit SATA II interface and sustained performance of up to 260 Mbyte/s the Slim SATA SSD from Viking Modular is ideal for telecommunications, gaming, embedded server & storage systems, field computing and defense & aerospace applications.

Viking Modular's Slim SATA SSD features intelligent write management techniques to optimize endurance, and offers densities of 25 Gbyte, 60 Gbyte and 120 Gbyte. Slim SATA's exclusive advanced controller features protection against normally catastrophic flash page and block failures, as well as data path CRC error detection for unparalleled data integrity. Viking Modular's Slim SATA SSD conforms to the JEDEC MO-297 industry standard and uses a standard SATA connector. It is available in a RoHS-compliant package and can be configured using MLC or SLC flash technologies.

Viking Modular Solutions, Foothill Ranch, CA. ((949) 643-7255. [www.vikingmodular.com].

Protocol-Agnostic, Multi-Fabric Interconnect for Intel Embedded Computing Line

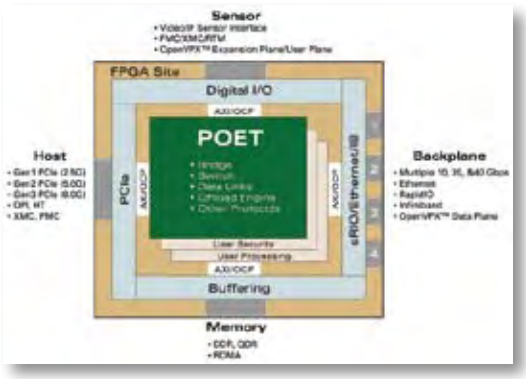

A new multi-fabric connectivity solution for Intel processors empowers companies to migrate to open solutions, future-proof their applications and guard their investment by ensuring a wide range of standards can be implemented without costly hardware changes. The protocol-agnostic, multi-standard switch fabric technology or Protocol Offload Engine Technology (POET) from Mercury Computer Systems implements, at its highest level, a collection of standard interfaces to create a bridge between processors and switched fabrics. Initially POET will support both Serial RapidIO (sRIO), the overwhelmingly preferred switch fabric for the defense industry, and 10 Gigabit Ethernet, the most common commercial fabric.

The 10 Gigabit Ethernet fabric is traditionally a best-efforts network. By offloading and accelerating standard protocols, POET simultaneously facilitates high bandwidth and low latency operation giving 10 Gigabit Ethernet the guaranteed, deterministic delivery of Serial RapidIO. POET can improve size, weight and power (SWaP) constraints in subsystems by providing local switching in the FPGA between multiple fabric ports to create distributed mesh networks, thereby eliminating the need for a centralized switched fabric card. In future releases POET will support next-generation switched fabrics such as 40 Gigabit Ethernet and InfiniBand as well as offload protocols such as RDMA over Converged Ethernet (RoCE), an InfiniBand over Ethernet standard. Through industry standard internal interfaces, customers can integrate their own downloadable firmware with POET to add value to their subsystems, or engage with Mercury's Services and Systems Integration group to obtain customization services.

Mercury Computer Systems, Chelmsford, MA. [www.mc.com].

PCI/104-Express Board Incorporates Qseven Modules

A new PCI/104-Express module incorporates industry standard Qseven modules. The Xtreme/CPU from Connect Tech was designed based on Qseven modules provided by congatec, which specializes in the development and marketing of industrial computer modules using standard form factors including Qseven, COM Express, XTX and ETX. The Connect Tech Xtreme/CPU is an embedded carrier board based on the PCI/104-Express form factor. This compact processor module enables complete integration with any industry standard Qseven module. The Xtreme/CPU conveniently provides onboard connectors allowing for instant access to Qseven features.

The Xtreme/CPU Features a PCI/104-Express form factor with 4 x1 PCIe lanes, two SATA interfaces, two each RS-232 and RS-422/485 connectors, four USB 2.0 ports, one Gigabit Ethernet connection, LVDS and VGA Video. Additionally, Xtreme/CPU gives instant access to a full range of PCI/104-Express add-on cards, and incorporating Qseven creates an I/O platform with a scalable CPU. Simply changing the Qseven module provides users with instant access to Atom-based products that are easily upgradeable to take advantage of future advancements of the Atom.

Xtreme/CPU is an attractive off-the-shelf solution for a broad range of applications. Achieve rapid proof of concept and deploy systems for field trials with off-the-shelf hardware. Connect Tech's engineering services offer customers the option of modified designs. The Xtreme/CPU PCI/104-Express processor module is compatible with all Qseven modules that are available at the time of this release.

Connect Tech, Guelph, Ontario. (519) 836-1291. [www.connecttech.com].

Highest-Density Non-Volatile DDR3 Memory Offers up to 8 Gbyte

A non-volatile system (NVS) memory family delivers densities up to 8 Gbytes providing flexibility to system architects and designers to tailor non-volatile memory for specific application requirements. AgigA Tech, a subsidiary of Cypress Semiconductor, now extends the market's portfolio of battery-free non-volatile memory solutions with the Agigaram product family. The technology merges NAND Flash, DRAM and a battery-free ultracapacitor power source into a highly reliable non-volatile memory system.

When used as a write cache in enterprise-class applications, Agigaram provides a performance-boosting building block while guarding against power failures and consequent loss of critical data. The DDR3 products deliver data transfer speeds up to 1333 megatransfers per second (MT/s). AgigA Tech offers product support to boost reliability, including extensive ultracapacitor testing with a dedicated test lab, in-system health monitoring, a system safe control protocol, and a product warranty up to the operating life of the system

Complementing the nvSRAM offerings from Cypress, the Agigaram system scales and extends non-volatile solutions to much higher densities. The next best high-density alternative, battery-backed memories, can offer high speeds but are subject to numerous problems, such as hazardous material issues, increased design complexity, long charge times, limited operating life and a high total cost of ownership.

The Agigaram system solves these problems with a novel use of a battery-free power subsystem, teamed with high-speed SDRAM, NAND Flash, intelligent power management and a proprietary system controller. During normal operation, the Agigaram system appears as a DDR3-registered DIMM to the host system, providing all the benefits and speed of a standard high-speed, high-density SDRAM. In the event of a power loss, the Agigaram system can be commanded to take control of the SDRAM and transfer its contents to flash memory using energy from its ultracapacitor power source, thereby preserving all the SDRAM data. After power is restored, the Agigaram system can be commanded to transfer the contents back into the SDRAM and returns control to the host system. This functionality can be used for power interruption/loss immunity, write caching and posting, data logging and journaling, instant-on recovery, and service and maintenance processing.

AgigA Tech, Poway, CA. (858) 375-4530. [www.agigatech.com].

COM Express Module Supports Multimedia

A fanless COM Express form factor module offers a complete, multimedia-capable platform for a variety of embedded applications. Available with a choice of Via C7 or Via Eden processors, the Via COME7N80 from Via Technologies also integrates Via CN896 North Bridge and VT8251 South Bridge chipsets providing a flexible and comprehensive computer-on-module product. Targeted at industrial PC and large OEM customers focused on dynamic application segments, including gaming, healthcare and industrial automation, customers can take advantage of a proprietary multi-I/O baseboard, or can utilize Via's extensive technical support in developing a custom baseboard.

The COM Express specification integrates core CPU, chipset and memory on the module, providing support for extensive connectivity options, including USB, audio, graphics and Ethernet, through board-to-board connectors to an I/O baseboard. The Via COME7N80 features the Type 2 COM Express standard, three PCI and three x1 lanes plus a x16 PCIe slot and up to 2 Gbyte of DDR2. In addition, the module features 10/100 Ethernet, up to eight USB 2.0 ports and up to four SATA devices.

VIA Technologies Fremont, CA. (510) 683-3300. [www.viaembedded.com].

Accelerating Your Success.™

BELL MICRO

AVNET Embedded

Embedded Solutions

Three Times the **Power**

People. Products. Services. The powerful combination of Avnet and Bell Microproducts provides the expertise you need to accelerate your success. Our combined team gives you access to world-class resources. Bringing industry leading line cards together, we now deliver the most extensive inventory of brand name systems, embedded hardware, displays, storage and software. And, with our enhanced services you have access to Avnet's leading ISO integration centers, financial solutions and supply chain strategies.

AVNET Embedded

1 800 332 8638
www.em.avnet.com/embedded

Follow us on Twitter!
www.twitter.com/avnetdesignwire

Advertiser **Index**

Get Connected with technology and companies providing solutions now

Get Connected is a new resource for further exploration into products, technologies and companies. Whether your goal is to research the latest datasheet from a company, speak directly with an Application Engineer, or jump to a company's technical page, the goal of Get Connected is to put you in touch with the right resource. Whichever level of service you require for whatever type of technology, Get Connected will help you connect with the companies and products you are searching for.

www.rtcmagazine.com/getconnected

RTC (Issn#1092-1524) magazine is published monthly at 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673. Periodical postage paid at San Clemente and at additional mailing offices.
POSTMASTER: Send address changes to RTC, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

Imagine...

Harnessing **Renewable Energy** More Efficiently

Your Imagination: Efficiently control and protect power generation systems from lightning, EMI interference, extreme temperatures and high winds, while maximizing the energy produced.

Our Innovation: Avago's fiber optics, isolated gate drivers, isolation amplifiers and digital optocouplers deliver solutions for wind turbine and solar energy designers.

Avago delivers:

- High voltage isolation
- EMI immunity
- No ground potential issues
- Increased safety and reliability
- Ultra-low power (1.3 mA) solutions

If a commitment to technical excellence and success is critical to your new designs – contact Avago at:
www.avagoresponsecenter.com/302

AVAGO TECHNOLOGIES

Your Imagination, Our Innovation

VPX, VME & cPCI True Military ATR Enclosures

16HP - Embedded Heat Pipe transfer system

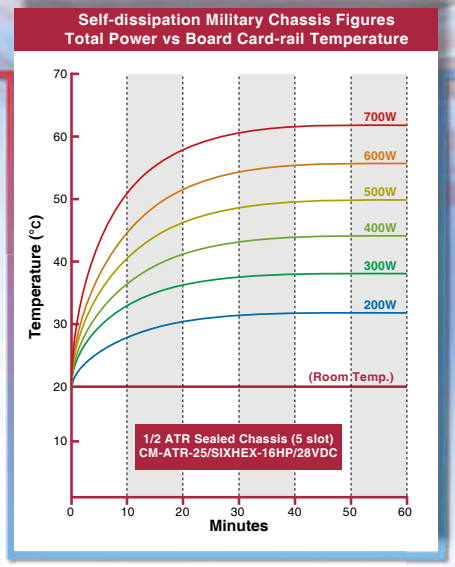

The perfect Sealed COTS solution

We have incorporated 16 Heat Pipes within our latest Six Heat Exchanger chassis series in order to improve thermal performance a step further. The new SixHex-16HP ATR range will increase your system MTBF by 40% when compared to our outstanding SixHex series as a result of the additional reduction in payload operating temperatures.

This achievement surpassed our design expectations. Customers will appreciate the superb quality and performance delivered in such compact enclosures. Today, even our smallest 9 kg ½ ATR chassis is capable of handling up to 700 Watts. This COTS industry record establishes an unbeatable sealed platform for the most demanding Military applications.

CM-ATR-25/SIXHEX-16HP : ½ ATR, 5 Slot 6U, 575W PSU, 9Kg

Product Highlights

- Contaminant-free enclosure
- Available in ½, ¾ & 1 ATR size
- VPX, VME64 & cPCI ready
- Accepts Conduction & Air-cooled 6Us
- Flexible Top & Bottom I/O wiring
- Six internal Heat Exchangers
- 16 integrated Heat Pipes
- Up to 1.8 KW total Power Dissipation
- Up to 150 W per slot
- Integrated Temperature Control Unit
- Dramatically increases payload MTBF
- 2 User defined PSU DC outputs
- 25°C less than heat exchanger ATRs
- 50°C less than conventional ATRs
- In-line EMI/EMC MIL-STD 461E Filter
- Stand alone low weight solution
- Customizable to specific requirements
- Mounting Tray with quick release system

CM-ATR-25 SixHex-16HP Power Supply Options		Vin Options		Backplane DC				Suggested BUS			Max Input Power
		28VDC	OTHER†	+5VDC	+3.3VDC	+12VDC	-12VDC	VME64	cPCI	VPX	Excluding Rear Fans
PSU Models	A-475W	✓		40A	22A	8A	8A	✓			700W
	A-575W		✓	40A	22A	12A	12A	✓			850W
	B-450W	✓		20A	45A	8A	8A		✓		700W
	B-550W		✓	20A	45A	12A	12A		✓		800W
	C-475W	✓		20A	22A	16A	8A			✓	700W
	C-575W		✓	20A	22A	21A	12A			✓	850W

† All Inputs except 28VDC : 48VDC / 270VDC / Autorange 90-264VAC @ 47-880Hz / 200VAC-3Phase @ 47-880Hz

It has been demonstrated in the field that military ATR enclosures are crucial to your end system reliability and performance. CM Computers is proud to develop superior products that guarantee your payload electronics are matched with excellence.

INSIDE US MILITARY COMPONENTS

All our chassis products are delivered Tested and Certified by independent authorized Labs per **MIL-STD-461E** & **MIL-STD-810F** for immediate deployment in US Navy and US Air Force military UAVs, Fighters and Helicopters.

CM Computer SixHex: Pure Power, Pure Dissipation, Pure Thermodynamics

Visit **www.cmcomputer.com** or contact us at **info@cmcomputer.com** to request our latest Chassis Catalog.

CM Computer
True Military COTS Products