OpenSystems Publishing

PowerPC: STILL TOPS FOR DEFENSE

Military EMBEDDED SYSTEMS

The COTS Technology Authority

VOLUME 3 NUMBER 5
SEPT/OCT 2007

INCLUDING:

MIL/COTS DIGEST SUPPLEMENT

Chris A. Ciufo
The "A's" have it

Jerry Gipper
Business end of standards

Duncan Young
Small form factors

WWW.MIL-EMBEDDED.COM

FPGA vs. DSP? Together is better

Industry execs weigh in:

Aitech, Intel, LynuxWorks, Themis

SPECIAL

INCLUDED WITH THIS ISSUE:

Advantech
Ruggedized Solutions Extend Application Range

Ethernet

ARINC 429

Mil-Std-1553

USB

Digital I/O

RS232/422

Small Box, many Interfaces.

Ballard's Avionics BusBox® *is a small, lightweight, embedded computer with built-in avionics databus interfaces. Use it for protocol conversion, radio tuning, or as an inexpensive LRU.*

Powerful, flexible, small and cost-effective.
Isn't it time you got MORE for your money?

Conduction cooled box coming soon!

Ballard Technology

Powerful Solutions

(425) 339-0281 **www.ballardtech.com**

Military EMBEDDED SYSTEMS

www.mil-embedded.com

September/October Volume 3 Number 5

COLUMNS

DEPARTMENTS

E-LETTER

www.mil-embedded.com/eletter

COVER
A USN F/A-18 Hornet pilot checks his AGM-65F Maverick missile during pre-flight as part of OIF. Raytheon's Maverick is a "smart" missile, employing COTS electronics throughout its Imaging Infrared (IIR) seeker, Harris digital data link, and GPS/INU subsystems. More importantly, FPGAs play a critical role in the electronics – including the MTS-206 Maverick Missile field test sets provided by Geotest-Marvin Test Systems. To learn more about the FPGA's increasing role in myriad defense applications, refer to the articles starting on page 12. (Image courtesy U.S. Navy, photographed by Phan (AW) Tommy Gilligan, USN)

FEATURES

WEB RESOURCES

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

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

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

Published by: OpenSystems Publishing™

Military EMBEDDED SYSTEMS

AN OPENSYSTEMS PUBLICATION

ISSN: Print 1557-3222

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

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

Canada: Publication agreement number 40048627

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

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

The business end of standards

By Jerry Gipper

New computer technology specifications emerge every day that become potential candidates for standardization through the efforts of existing standards organizations. The primary reason these specifications strive for standard status is to make developers and users confident that the specification will not change without due process and that it will be widely used, thus driving down costs. If this is not a goal, then staying at specification status is just fine.

Organizations such as the Institute of Electrical and Electronics Engineers, Inc. (IEEE) and the American National Standards Institute (ANSI) have rigorous processes in place to ensure that the approved specification has been through a thorough vetting process. The resulting documents have been revised and improved upon through the work of many concerned parties. Anyone who has been involved in the standards process knows the immense amount of time and energy required to make it to the approved stage.

What appears to be missing is adequate justification to expend all the necessary effort to convert a specification to a standard. My observation is that launching a standards development effort appears too easy. The business side needs more representation to justify the tremendous amount of effort it takes to develop a ratified standard.

From the IEEE Standards Association (IEEE-SA) operations manual, "Approval of a standard by the IEEE-SA signifies that the IEEE believes the document to be consistent with good engineering practice and that it represents a consensus of representatives from materially affected industries, governments, or public interests." This is consistent with most other standards bodies' operations procedures, but it only addresses the results of the technical standards development effort. It does not address the potential business and market impact of the standard. No basic requirements are established for meeting good business goals.

Many consortia and standards bodies have both a technical and business structure. There is usually a team of engineers working on technical challenges and issues; meanwhile, the marketing and public relations types sit in a separate room, trying to understand what the technical side is working on or how to position and message the technical side's output. Usually the technical side is a strong team that is excellent at coming up with solutions, creating a draft specification, and pushing it through the standards process.

However, in many cases, the business side is weak or even nonexistent. All too often, no one is there to catch and run with the ball as a standard pops out. Even worse, most consortia do not have a strong group that can gather and analyze market problems before the technical teams start their work. Sometimes an influential individual or company provides some guidance, but as the individual moves on or the companies shift strategies, business leadership falters.

I propose that a better job could be done to create standards if more marketing was done at the beginning. A product lifecycle-like process should be in place to force the steps needed to build and justify a business case. A much more rigorous data gathering and analysis process should be in place. A business case should be developed, reviewed, and approved by the governing board of a consortia or standards body before any standards work is expended. The typical process used by many groups requires that three or more companies sponsor a proposal for a working group to be established. This process is very inadequate. The sponsors need to be called to the mat to demonstrate a business case. Processes similar to what an investor or venture capitalist follows before making an investment demonstrate the type of rigor that should be followed for a major endeavor.

> **... In many cases, the business side is weak or even nonexistent. All too often, no one is there to catch and run with the ball as a standard pops out.**

The effort in launching most standards should be as consistent as the effort to launch a major new product line. The market problems being addressed by the specification should be researched and understood. The market potential and adoption plan should be presented. A plan for evangelizing and marketing the standard should also be included, and a budget and sources for the budget should be considered.

The industry and consumers would then benefit from better choices, and developers would be able to focus on efforts that stand a much greater chance of sticking and reaching levels of success. The success level of any individual standard should increase and the confusion factor with conflicting standards should decrease if the homework is done before launching a significant standards effort.

I would love to hear your opinion on this topic.

For more information, contact Jerry at jgipper@opensystems-publishing.com.

Small form factor embedded subsystems

By Duncan Young

The drive toward battlefield digitization plus the effects of ongoing conflicts around the world are driving the military to upgrade existing vehicle fleets as a higher priority than introducing wholly new fleets, as envisaged by the Future Combat Systems (FCS) program. This is ramping up demand for quick-turn applications using new communications bandwidth. These new applications are often not critical to a fleet's war-fighting capability but are targeted toward greater information exchange, data presentation, training, and general purpose tools for the crew. Many of these new capabilities are being added on top of existing systems, rather than being fully integrated, creating a new phenomenon: the shoebox-sized, rugged embedded PC, which is easy and familiar to use and small enough to be mounted almost anywhere that panel space is available.

Systems for new platforms will continue to be implemented using well-developed open architecture embedded computing standards such as VMEbus, VPX, and CompactPCI. These will allow the configuration of very powerful, fully network-enabled, integrated systems that will be maintainable and upgradeable in the field over the life of each vehicle fleet. However, as these standards are generally 6U in size (233.4 x 160 mm), a subsystem configured from them will be quite bulky for some of the new rapid upgrade programs. For example, it is much harder to pack even a small number of these bulky boxes into light helicopters, Unmanned Aerial Vehicles (UAVs), and wheeled vehicles than it would be if using larger numbers of much smaller box sizes. With electronics miniaturization taking such strides, the larger board sizes may also contain too much functionality and performance for convenient modularization for these types of relatively simple add-on capabilities.

Embeddable PCs are widely used in industrial control systems, displacing many earlier VMEbus systems. They can be custom built for very specialized systems or, more frequently, configured from one of the many industry-standard form factors. Table 1 lists the characteristics of some of the most commonly used varieties, all of which either have been or could be used as the basis of a rugged embeddable PC.

There is room in the market for many different solution sets based around the concept of a self-contained, fully functional PC running Windows or Linux. As these are add-on units, the vehicle to which they will be fitted is unlikely to have any provision for cooling. So there will be no conditioned air or pumped liquids available. Baseplate or natural convection are the most feasible methods of cooling, thus steering shoebox designs to be fully enclosed and use only the outside surfaces of the box for cooling and connector space. The selection of the form factor and box size most suited to an application can be complex, but there are two simple relationships that help:

performance = watts = size
ruggedness = weight

For example, an embedded PC based on the PC/104-*Plus* form factor will offer relatively low-end performance, limited by its small size and hence its ability to dissipate a lot of heat. PC/104-*Plus* is not inherently rugged but can be made so by adding weight in the form of supporting structures and enclosures. Where top-end performance is required, larger size and/or more efficient cooling such as conduction will be required. Conduction cooling also provides much additional stiffening to a module, making it the ideal choice for prolonged operational use in the harshest military environments.

GE Fanuc Embedded Systems' MAGIC-1 is an example of a fully ruggedized, preconfigured embeddable PC. Based on the conduction-cooled 3U VPX (VITA 46)

Board type	Form factor mm	Real estate cm²	Conduction cooling	PCI	Network support (Ethernet)	PCI Express and fabric support
PC/104-*Plus*	96 x 90	86	x	✓	✓	x
EPIC	115 x 165	190	x	✓	✓	x
EBX	203 x 146	297	x	✓	✓	x
3U VPX	160 x 100	160	✓	x	✓	✓
MicroTCA single	74 x 182	134	proposed	x	✓	✓
MicroTCA double	149 x 182	270	proposed	x	✓	✓
ETX Express	95 x 125	119	x	x	✓	✓

Table 1

Figure 1

form factor, MAGIC-1 (Figure 1) is optimized for graphical applications. It uses a 2 GHz Intel Core Duo processor with a 16-lane PCI Express interface to an NVIDIA G73 graphics processing unit to achieve its high-performance goals.

Embedded training systems located in active fighting vehicles are an excellent example of an application area that such an embeddable PC with high-resolution, truly realistic graphics performance is targeted toward. Training exercises – particularly in multisensor, cooperative environments – are designed to develop more intuitive and interpretive skills than the traditional *action-reaction* type of training. This does not require an exact real-time mimic of actual equipment operation, hence hard-deadline, real-time performance parameters are less critical. But because the training is performed in the deployed environment, embedded training systems must be fully ruggedized. Another typical application for this kind of preconfigured subsystem could be to drive a large-screen display for planning or briefing troops inside expeditionary armored vehicles.

As lessons learned from ongoing conflicts percolate down into urgent operational requirements, the military COTS market is developing new directions and creating new standards with innovative technology solutions. Flagship applications for the more mature 6U form factors still abound, particularly in new build platforms where the real benefits of fully integrated applications and reduced levels of maintenance can be achieved. However, existing platforms – many of which were introduced into active service prior to the COTS era – are, of necessity, rapidly evolving in an incremental manner to counter new and more demanding threats. To satisfy this demand, cycle times for new subsystems from design to deployment must be reduced dramatically. And it is in these applications that the new breed of small, rugged, preconfigured PCs will be so successful.

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

Daily Briefing: News Snippets

By Sharon Schnakenburg, Assistant Editor

www.mil-embedded.com/dailybriefing

A (life)time to achieve

It's that time of year, and Frost & Sullivan is at it again. This time, the organization, which hands out beaucoup awards annually, has announced that Quantum3D, Inc.'s cofounder and president Ross Q. Smith is one of its 2007 Lifetime Achievement Award winners. Frost & Sullivan's director of production events Caryn Brown says Smith was nominated for several reasons: His role as a "pioneer" in PC-based real-time visual computing, his service as a "driving force" in the movement to turn the embedded visual computing and visual simulation/training markets into COTS/open architectures, and his "strategic vision" of watershed real-time visual computing products. Smith – along with honorees from other industries including healthcare, automotive and transportation, and financial services, among others – is honored at the 3rd Annual Growth Innovation and Leadership 2007: A Frost & Sullivan Executive Congress on Corporate Growth in Monterey, Sept. 16-18.

FPGAs: Up in space

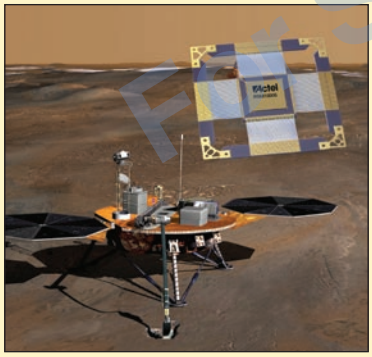

How well do FPGAs and "mission-critical" really blend together in space? Very well, according to representatives of the Actel Corporation, whose RTAX1000S-CQ352 FPGA device flew out of the stratosphere at NASA's Phoenix Mars mission launch at Cape Canaveral Aug. 4. Actel's RTAX-S FPGA is part of the spacecraft's Meteorological Station (MET) instrument, designed to gather, process, and send pressure and temperature data back to Earth-stationed researchers. Actel's high-reliability, high-density RTAX-S family includes 250,000 to 4 million equivalent system gates, along with usable error-corrected onboard memory, Single-Event Upset (SEU) hardened flip-flops, and large amounts of user-defined I/O.

Headache-eliminating partnership

Engineering departments can now eliminate the hassle of sending their test and measurement instruments back to several different manufacturers for calibration: Sypris Test & Measurement and Tektronix recently formed a one-stop-shop partnership where customers can receive onsite, single-source calibration services for about 80,000 different instruments (nearly 90 percent of those typically used, according to Sypris), no matter who the manufacturer is. The major motivations behind the union reportedly include cutting customers' costs, increasing customer convenience, and eliminating the inconsistency that can arise when dealing with many different vendors.

U.S. Army gains one more

Boeing photo taken by Fred Trato

The U.S. Army has a new addition (or is that a new *edition*?): Boeing's new CH-47F Chinook helicopter was recently pronounced combat ready by Army officials and has joined the ranks for service to its country. The helicopter – assigned to the 7th Battalion, 101st Aviation Regiment (Air Assault), based at Ft. Campbell, Ky. – derives its power from two 4,868-horsepower Honeywell engines and travels at speeds in excess of 175 mph. The CH-47F comprises a BAE Digital Advanced Flight Control System (DAFCS), a Rockwell Collins Common Avionics Architecture System (CAAS), and a "modernized" airframe. It also includes the Robertson Aviation Extended Range Fuel System – which provides a 400+ nautical mile mission radius – and "improved survivability features" such as Improved Countermeasure Dispenser Systems and Common Missile Warning.

Semiconductor sales: Moving up in the world

In the common vernacular, "It's all good." Well, at least most of the recent semiconductor sales news is, according to a report by the Semiconductor Industry Association (SIA). Tallies indicate that worldwide semiconductor sales reached $20.3 billion in May 2007, representing a 2.4 percent growth hike from $19.8 billion in May 2006. NAND flash and microprocessors experienced the largest sales volume gains, likely contributing to the nearly 10 percent increase in PC and cell phone market unit sales; meanwhile, DRAM experienced strong price pressures, resulting in total DRAM sales revenue losses of 8 percent even though unit shipments were 7 percent higher. While admitting to a "sequential decline" this past April, SIA president George Scalise remains optimistic about this year's semiconductor sales overall, citing a 3.1 percent increase in the first five months of 2007 as compared to the same time frame in 2006.

DDC®

©Data Device Corporation

LEADERSHIP BUILT ON OVER 40 YEARS OF INNOVATION

Four decades of quality manufacturing, design innovation, and process control has earned us the time-honored trust and confidence of a global network of customers.

DDC IS YOUR DATA NETWORKING SOLUTION

- MIL-STD-1553 COTS Boards and Components
- Fibre Channel
- Gigabit Ethernet
- High Speed 1760
- Enhanced Bit Rate 1553 (EBR)
- ARINC 429

Toll Free: 1-800-DDC-5757 **www.ddc-web.com**

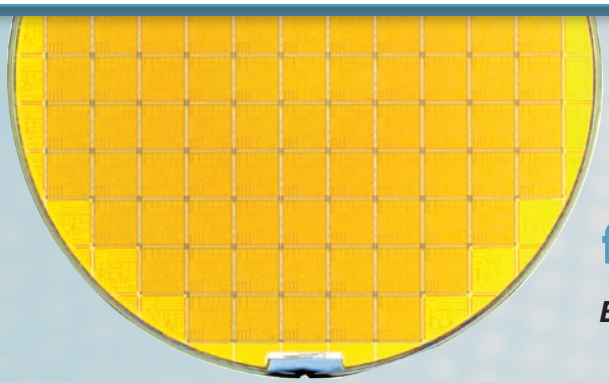

FPGA coprocessing architectures for military applications

By J. Ryan Kenny and Bryce Mackin

Military electronic equipment manufacturers at one time were categorized into two fairly distinct design types: hardware based and software based. Each offered significant advantages and disadvantages in speed and flexibility. Newer, more flexible war-fighter support systems today take advantage of both hardware and software using hardware acceleration with FPGAs.

Hardware acceleration with FPGAs is important in military sensor applications for three reasons. First is to reduce system latency so that defensive systems can react faster to enemy threats, such as jamming or blinding. Second, as military threats have become more elusive and blend in with urban settings, systems are constantly under pressure to increase sensor resolution. And lastly, as military threats become more tactical, the Size, Weight, and Power (SWaP) of a system must be reduced.

Given these technology drivers, the military market for high-performance processing is expanding rapidly, though there are currently few application-specific (tailored application) coprocessing FPGA solutions available.

With the proliferation of distributed and multicore processing computers, complex signal processing applications can be accelerated with distributed functions. The processing cores, however, are still limited by cache size and coherency, memory bandwidth, and in some cases, power and cooling.

One approach to improving processing power and efficiency in single-core and multicore processors is to use FPGAs as application-specific coprocessors. These devices can be designed to fit into Intel Xeon and Advanced Micro Devices (AMD) Opteron sockets to replace one of the processors with pipelined FPGA logic. Alternately, FPGA "drop-in" modules can be added to processor memory subsystems to offload specific hardware pipeline functions. The advantages of these approaches have been explored by Altera Corporation and applied in some cases by FPGA users. Many significant advantages to sensor signal processing problems can be examined with three sample layouts: two example architecture sets for multicore processing and one architecture for single-core processing.

Potential algorithmic returns

Hardware acceleration using an FPGA as a coprocessor is intended to offer 10x to 100x speed improvement for tailored algorithms, and anywhere from 3x to 50x speed improvement for the user application. These target numbers are based on commercial applications such as financial analysis, data warehousing, and biosciences.

Application	Processor Only	FPGA Coprocessing	Speed Up
(1) Hough and inverse Hough processing	12 minutes processing time Pentium 4-3 GHz	2 seconds of processing time @ 20 MHz	370x faster
(2) Spatial Statistics (Two Point Angular Correlation Cosmology)	3,397 CPU Hours with 2.8 GHz Pentium (Approximate Solution)	36 hours (exact solution)	96x faster
(2) Black-Scholes (Financial Application (single precision floating point 2M points))	2.3M experiments/sec with a 2.8 GHz processor	299M experiments/sec	130x faster
(1) Smith Waterman ssearch34 from FASTA	6,461 sec processing time (Opteron)	100 sec FPGA processing	64x faster
(3) Prewitt Edge Detection (compute-intensive video and image processing)	327M Clks (1 GHz processing power)	131K Clks @ .33 MHz	83x faster
(1) Monte Carlo Radiative Heat Transfer	60 ns processing time (3 GHz processor)	6.12 ns of processing time	10x faster
(1) BJM Financial Analysis (5M paths)	6,300 sec processing time (Pentium 4-1.5 GHz)	242 sec of processing @ 61MHz FPGA	26x faster

Table 1

Several selected algorithms taken from these commercial applications have been implemented and tested on FPGA coprocessors for evaluation. These are shown in Table 1.

In addition to the performance advantages of hardware acceleration, replacement of either an Intel or AMD processor significantly reduces power consumption and heating in a system. This is a result of the streamlining of software and hardware functions where they are best suited. Initial benchmarks on the expected power savings are algorithm dependent.

Single-core coprocessing architecture

The majority of military systems utilize single-core CPU architectures. Migrating to higher-performance processing systems, while still meeting government software code reuse objectives, is a problem that can be solved by offloading software functions into FPGA hardware.

Some signal processing problems are fairly simple, involving a single sensor with a single data processor. In order to optimize the data flow and computational bandwidth in this system, the coprocessing problem is very straightforward. Large functions that can be mathematically isolated and pipelined are performed outside a processor using a systems analysis process.

One such process utilizes a hardware and software architecture called IMPLICIT+EXPLICIT. It can be used with single or multiple microprocessor boards. This architecture combines an FPGA-based reconfigurable processor, MAP, in a peer relationship with a microprocessor. This peer relationship is achieved by connecting programmable logic to the microprocessor using the microprocessor's memory DIMM slots. FPGAs are used to implement pipelined direct execution logic to perform functions such as DSP logic. This can accelerate sensor processing, allowing for less power consumption and greater sensor resolution.

Another processing challenge in sensor systems is moderating the information between sensor and microprocessor with minimal latency. This is performed with user logic chips diagrammed (Figure 1) in generic blocks. The controller is an FPGA that performs several functions including virtual to physical address conversion and DMA packet generation. The SDRAMs shown are part of a 64-bit global shared address space that includes the microprocessor memory. The direct execution logic that implements the user's program resides in the two user logic FPGAs. These can perform up to sixteen 64-bit references per clock cycle to the 8 SRAM banks. In addition to the connection to the DIMM bus, data may also be received or sent using the GPIO port. This port is very useful for delivering sensor data directly to the user logic chips.

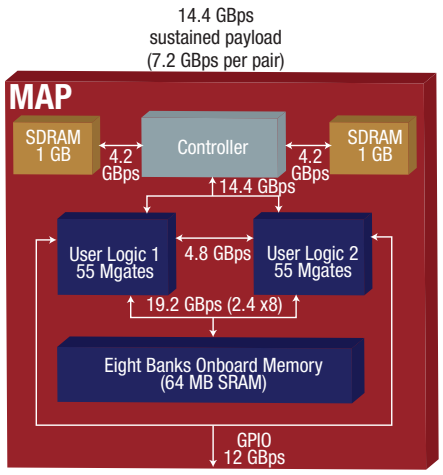

Figure 1

The division of the user's program between the microprocessor and programmable logic is accomplished using the Carte high-level language programming environment. This allows programmers to use ANSI C and FORTRAN to generate a single executable program that will control both FPGAs and the microprocessor. This allows the division of logic between hardware and software to be performed by a single architect.

Performance gains of this coprocessing approach have been demonstrated to several government customers, resulting in orders of magnitude improvement over a stand-alone 2.8 GHz Xeon processor. These applications include military imaging, radar signal processing, and image target recognition. Benchmark results for the complete application including data movement time are shown in Figure 2.

Some military applications may be more advanced, utilizing multiple sensors, sensors in multiple modes, or rapid system reconfiguration. For these systems where the user is willing to pursue a more difficult segmentation of logic in pursuit of higher performance, multicore processing is one of the latest trends in high-performance computing.

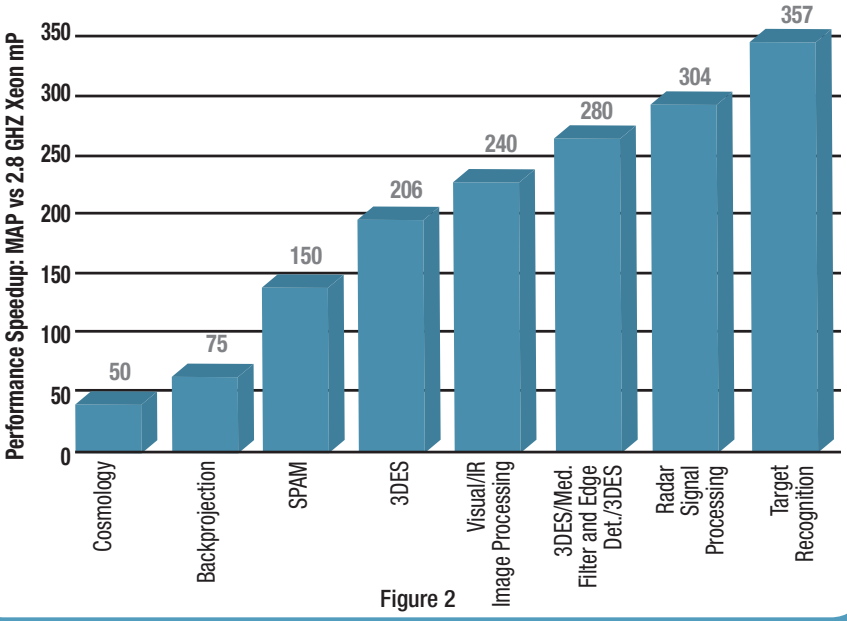

Figure 2

Multicore coprocessing architecture

Multicore processing is utilizing multiple CPUs to execute code in parallel on the same algorithm, but separate process threads. Multicore processing requirements are being significantly pushed by military applications in imaging and data processing where performance gains in single-core processors are becoming marginal. Providers of both Intel and AMD based architectures are being asked for their solutions.

Accordingly, the goal of multicore coprocessing is the same as single core: to identify and isolate math functions that can be efficiently offloaded from software into hardware in order to accelerate sensor system performance. A set of notional coprocessing architectures for the Intel Xeon Quad Core (Figure 3) and AMD Opteron (Figure 4) layouts are shown. Figure 5 shows a representative system developed by XtremeData that utilizes an FPGA on an XDI module. In all three layouts, FPGAs are placed in CPU sockets and require application hardware to best utilize the FPGA logic.

In the Xeon architecture, a processor and a coprocessor are connected using the Intel Front Side Bus (FSB) architecture. A Northbridge (information available from Intel) is used to connect each FSB to one another. These are the normal pathways used by multicore CPU instructions. Access to front side bus interface standards is becoming more widely available to developers.

The AMD Opteron architecture uses direct HyperTransport interconnects between each processor/coprocessor socket. This 32-bit packetized data connection is similar to 2.5 V LVDS and can easily be integrated using available HyperTransport FPGA cores.

The Xtreme Data architecture is an implementation currently in use in the several nonmilitary markets. It utilizes a multiprocessor motherboard with an AMD and Intel socket-compatible module interface for the FPGA and other components. This multisocket architecture allows the application to take advantage of the processing power available in the FPGA. This coprocessor solution offloads high-resolution DSP algorithms in military imaging systems and performs Black-Scholes simulations for financial modeling.

The motherboard plugs into a standard commercial enterprise rack or blade server. It utilizes a low-latency, high-bandwidth HyperTransport interface to the other processor. The FPGA uses all resources intended for a CPU – power supply, heat sinks, HyperTransport links – and is programmed with on-chip memory controllers.

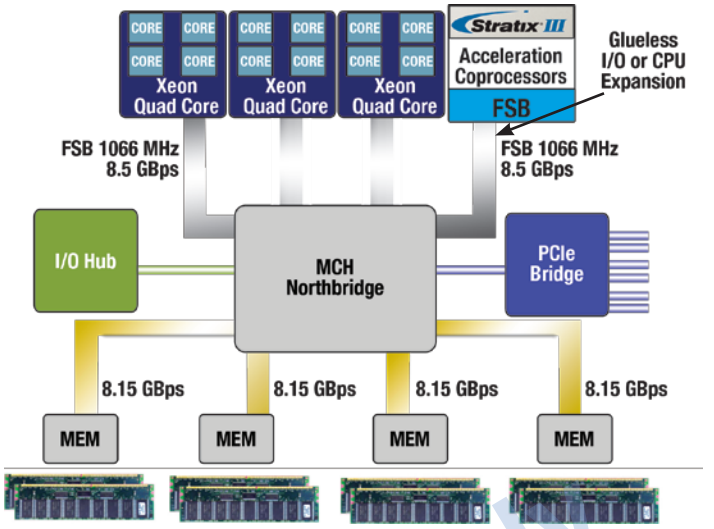

Figure 3

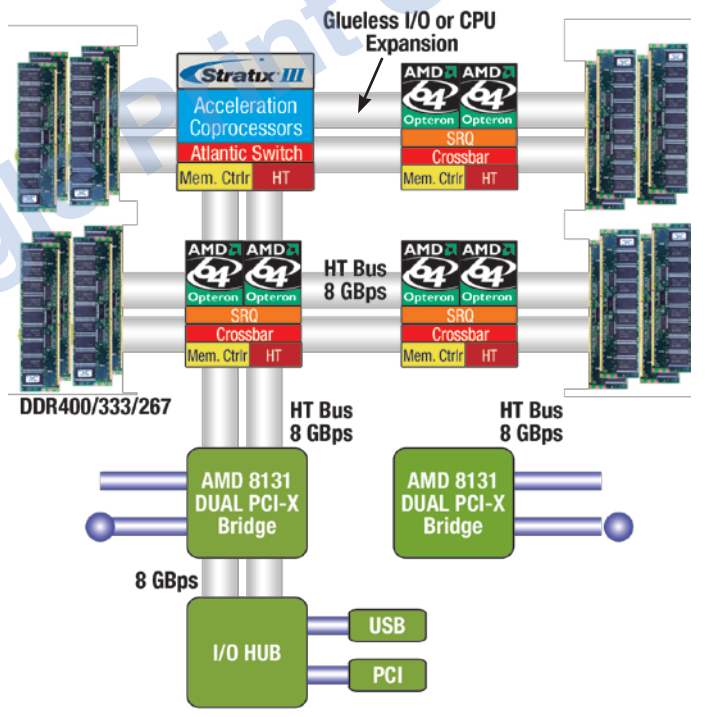

Figure 4

Implementation

In addition to the FPGA design and drop-in module, an application engineer has the task of modifying the application software to exercise the FPGA coprocessor for tailored hardware operations. There are several approaches to this process, which can be logically broken out into two steps.

First, a systems engineer needs to determine which functions in the system will be offloaded into the FPGA for accelerated processing. This can be accomplished by using traditional software profiling tools.

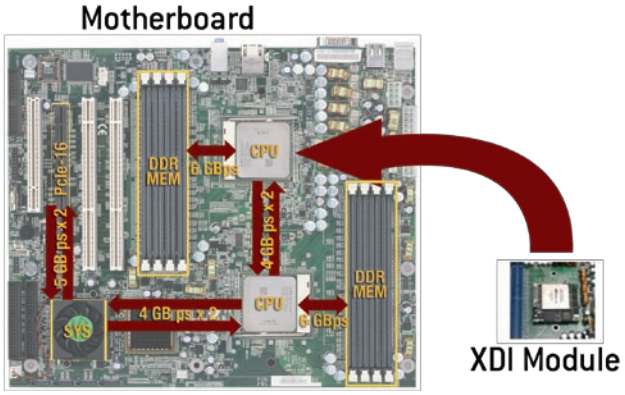

Figure 5

Second, the application engineer must identify the function calls and interfaces for coprocessing, as well as the task distribution among the remaining processors. In the case of the Carte Programming Environment, users have equated this effort as comparable to the effort required to convert a uniprocessor application into an MPI application.

As multicore processing becomes more prevalent, software tools will become available to help designers parallelize and compile their systems for N hardware nodes from other vendors. Function calls can be translated from C to efficient HDL using proprietary tools offered by FPGA providers. This application tailoring requires only basic FPGA knowledge, allowing the system designer to build algorithm architectures confidently without a full FPGA development staff.

Simplifying coprocessing

Using design software developed by FPGA providers and hardware partners, hardware acceleration is evolving from a specialized engineering effort into a powerful systems design tool. It is accessible to software designers today, and will only become more accessible in the future.

The next generation of embedded computation for military applications is likely to involve strong design partnerships between providers of software and pipelined hardware. More of these will be seen in the future as holistic development tools allowing simple partitioning of applications into the processor domain that better suits their execution.

As computing architectures become optimized for military applications, code portability and supportability may suffer. Engineers will then have a new challenge: balancing tailored performance versus code reuse.

J. Ryan Kenny *is technical marketing manager in Altera's military and aerospace business unit. He is responsible for creating FPGA-based technical solutions for the military data and signal processing market. He joined Altera in March 2007 and has more than 10 years of experience in space and defense electronics in the U.S. Air Force and at Lockheed Martin. He graduated from the U.S. Air Force Academy and completed an MSEE and MBA from California State University Northridge and Santa Clara University respectively. He can be contacted at rkenny@altera.com.*

Bryce Mackin *is the strategic marketing manager in Altera's computer and storage business unit. His focus is on FPGA coprocessors for the high-performance computing market. Bryce has been in the computer and storage market for more than 10 years. Prior to joining Altera, he served as a strategic marketing manager for Xilinx, worked at Adaptec in several product marketing roles, and held the position of marketing chairman for the Storage Networking Industry Association (SNIA) IP Storage Forum. Bryce has a degree in Industrial Engineering from California State University in San Jose. He can be contacted at bmackin@altera.com.*

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

FPGA/DSP hybrid architectures: Satisfying the reconfigurability requirements of the military

By Ron Huizen

Many signal processing applications in the military require reconfigurability coupled with the speed and dynamic range of a floating-point DSP. The solution is a hybrid architecture that employs an onboard FPGA framework to create a seamless transition between I/O, FPGAs, and DSPs.

Embedded signal processing applications in the evolving modern-day military increasingly require flexibility, adaptability, and reprogrammability. The ability to modify a military system in real time based on what is happening in the real world – adapting to weather conditions, new threats, mission changes – is crucial. When this is coupled with the requirements of floating-point signal processing – low latency, high throughput, extended precision, and dynamic range – the ability to properly satisfy all requirements using one processing technology diminishes.

Many times, a design that includes both FPGAs and DSPs – a hybrid architecture – provides the best solution. The challenge then becomes one of integrating two quite different technologies efficiently and effectively so that the end system functions as one entity. A software control framework that can successfully merge the two technologies, allowing the system to play on the strengths of each, is a necessity.

Leveraging strengths, mitigating weaknesses

Embedded signal processing applications that require sustained, low latency, high-throughput data processing traditionally rely on DSPs. When additional dynamic range is needed, a floating-point DSP has been the processor of choice, also ensuring that the system will support future, more advanced requirements. On the other hand, applications that require enough flexibility to be easily updated and modified usually rely on FPGAs, which, in recent years, have also gained the ability to provide embedded signal processing. When an application requires both the flexibility of an FPGA and the advanced signal processing capabilities of a floating-point DSP, the problem arises of which technology to use.

FPGAs are extremely effective for well-defined, straightforward, high-speed, repetitive problems requiring data flexibility and parallelism, the type of processing that is very often needed at the front end of the signal processing system. They can provide reconfigurable interfacing, with many of the newly released FPGAs able to implement switch fabrics such as Serial RapidIO and PCI Express, making them ideal for on- and off-board data transfer. They have been touted as do-everything silicon, able to perform high-end processing and provide the dynamic range required by many military signal processing applications.

In practice, though, difficulty in programming these devices for floating point can greatly increase time-to-market, making them unlikely candidates for high-end signal processing at any point in the near future. Another drawback of these devices is power inefficiency, a common trade-off with flexibility. A similar function implemented in an FPGA might take two to three times as much power versus a DSP. Note, however, this inefficiency may be an acceptable trade-off for the sake of flexibility.

Floating-point DSPs, on the other hand, are much better at implementing a broad range of highly complex algorithms that require floating-point math and have the tendency to change frequently. Any application that requires some sort of decision making or adaptive processing, where the next processing step is dependent upon the current result, requires a floating-point DSP. They are better in terms of power consumption and have the benefit of programming in C, decreasing development time. What they cannot do efficiently is handle the repetitive, straightforward processing that is many times required on the front end of the signal processing application; they can actually waste large amounts of processing efforts on these types of tasks, translating to a less efficient and more costly system. It's not uncommon for an FPGA implementation of a straightforward, parallel, repetitive task to require two to five times as many DSPs to complete the same task.

Given the strengths and weaknesses of each, it makes sense to use them as complementary technologies, mitigating the risks and weaknesses of each technology when used on its own. A hybrid architecture can intelligently combine the two, providing efficient performance, high throughput, and data control, all within a reasonable power budget. This type of system also has a higher amount of inherent flexibility, providing reprogrammability for bug fixes and giving system designers the ability to support a different application in a very short amount of time.

The FPGA framework solution

While the benefits of hybrid design are many, along with these benefits come some significant design issues. A successful hybrid design needs to be able to properly allocate data bandwidth (including memory and I/O) among the FPGA(s) and DSP(s), easily connect the onboard data options while retaining the ability to modify them, and last but not least, successfully integrate the FPGA processing with the DSP so that each compute element is handling the part of the problem it is most proficient at. A software framework that can handle each of these is essential. The framework then functions as two separate entities: an I/O interfacing and routing device and a configurable FPGA pre-, post-, or co-processing engine. One example implementation is BittWare's integrated system framework ATLANTiS, shown in Figure 1.

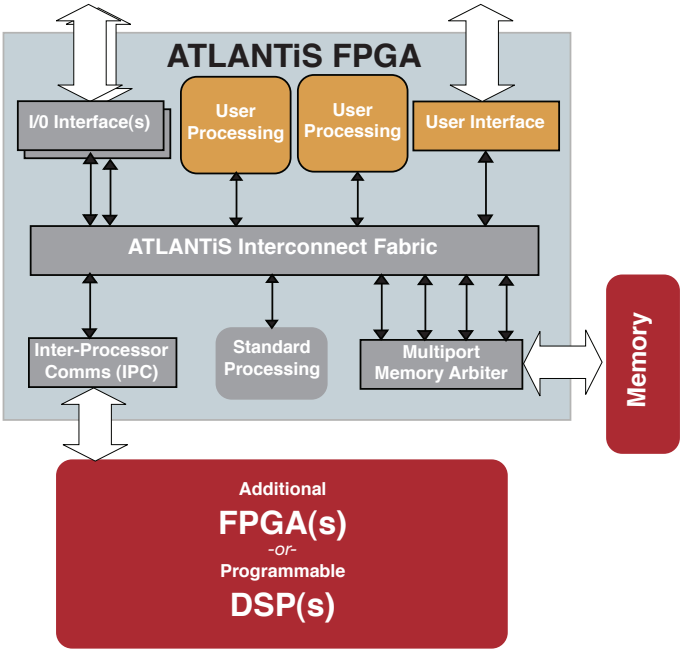

Figure 1

Data interfacing and routing

Given that FPGAs are better at the repetitive, well-defined, straightforward processing – as well as being known for their flexibility – it makes sense that the board's onboard entry and exit points reside on the FPGA, and that the entire software framework handling this data also resides on the FPGA. This FPGA framework can then handle all data transfer on and off the board and between itself and the DSP(s), acting as a software programmable cable and enabling the dynamic connection of every I/O to any other I/O, where connections can be created and broken without the need to recompile or change cables. The options for data interfacing can include: digital I/O (LVDS or single ended), general purpose I/O, flags, interrupts, link ports, high-speed serial links, rear panel connectors, off-board connectors, board-to-board links, cluster-to-cluster links, and multiport memories.

The ability to easily connect (and disconnect) the external data sources to each other, to IP processing modules within the FPGA, and to the onboard DSPs creates an incredibly flexible system that can be instantly modified for an immediate need, while also retaining the ability to be updated for future modification and expansion. This is achieved via the main framework switch, shown in Figure 2, which controls all connectivity and bandwidth allocation.

The switch is configured via registers by the host or any of the onboard DSPs and comprises two smaller switches, each with eight inputs and eight outputs. The switch connectivity is defined by the specific FPGA load and is controlled by the configuration registers, which allow source and destination time slices. This is all controlled by the system designer via a Windows navigator, command scripts, or C functions. Each switch accommodates single point, multipoint, or broadcast switching amongst any of its eight inputs at up to 125 MHz clock rate at 128 bits per I/O. The bandwidth can be allocated to any or all data inputs by 32 configuration registers that evenly divide the bandwidth slots into 62.5 MBps and can be changed "on the fly."

> "The ability to easily connect (and disconnect) the external data sources to each other, to IP processing modules within the FPGA, and to the onboard DSPs creates an incredibly flexible system that can be instantly modified for an immediate need…"

Configurable FPGA processing engine

Should the signal processing application in question require the FPGA to provide some of the signal processing, the FPGA framework also needs to enable the user to add FPGA processing blocks into the data flow at any time while the application is running. This can also be achieved through register-controlled data routing between all modules within the FPGA, enabling the insertion of FPGA processing modules at any point. As discussed in the previous section, the framework's memory space is accessible by the onboard DSPs or the host via either a peripheral bus or a cluster bus and is accessible to the designer via a GUI that provides a graphical representation of all possible source and destination I/Os. This enables DSP or host control and status, and R/W access to all of the frameworks resources; the most important resource is the main framework switch. Access to the memory

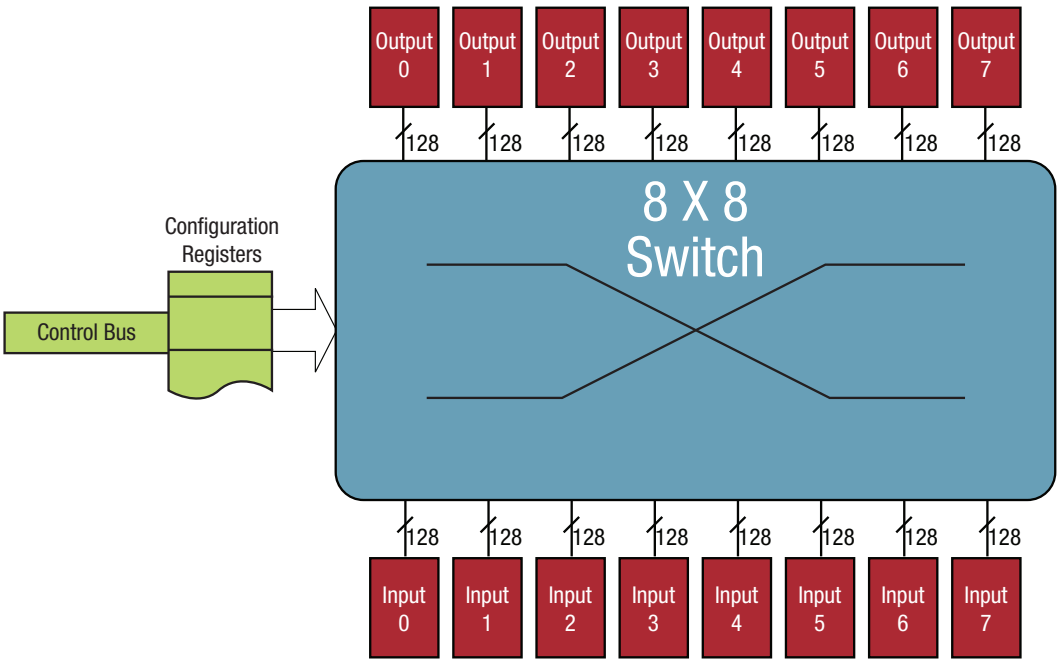

Figure 2

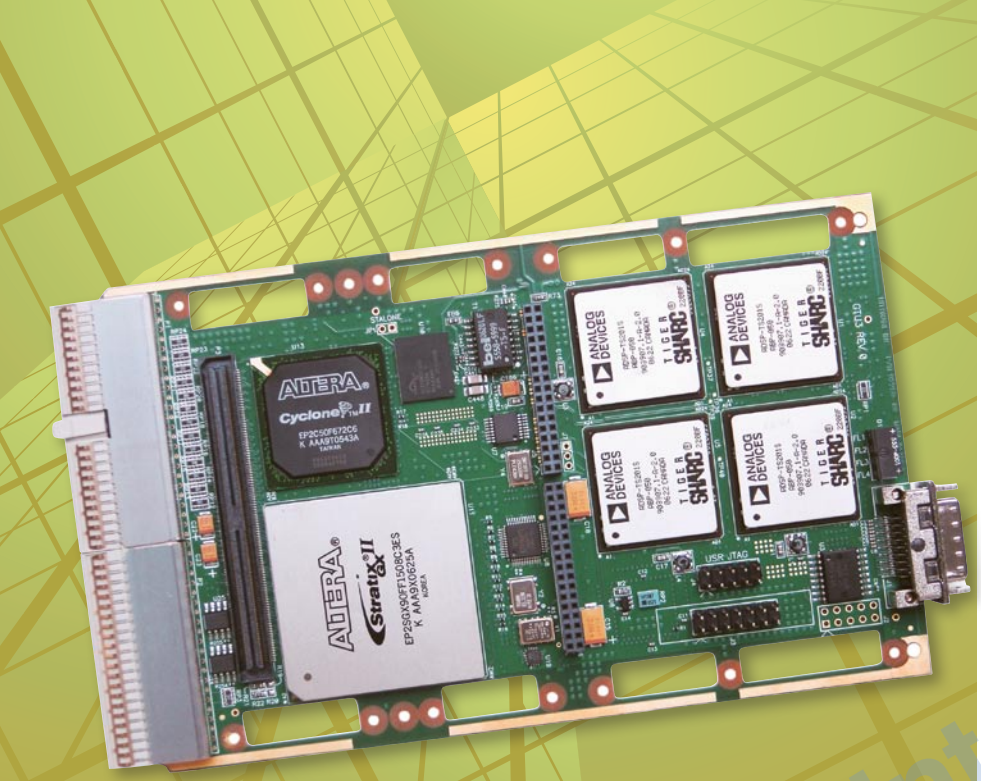

Figure 3

space enables instant I/O routing changes and the ability to insert standard and/or custom FPGA processing blocks into the queue at any point during data transfer.

When the user configures the framework for a specific load, the master switch is programmed, defining source and destinations for each I/O from all available connections; the 32 configuration registers are memory mapped by a DSP or the host. Should a different load be required, the framework is reset and then configured with the next load. FPGA IP blocks can be added into the switch to provide additional processing at any point during the data flow. The result is the efficient integration of two different technologies. With the FPGA framework handling all data routing, as well as providing any pre- or post-processing, the onboard DSPs are freed to handle the highly complex signal processing for which they have been designed.

A flexible, high-end signal processing system

Many of today's military signal processing applications require the flexibility and reprogrammability of an FPGA combined with the ease-of-use and floating-point processing provided by DSPs. The solution is a hybrid architecture that efficiently and effectively combines both by way of an FPGA framework. BittWare's GT-3U-cPCI, shown in Figure 3, is one example of this architecture. The board combines the Altera Stratix II GX FPGA with Analog Devices TigerSHARC TS201 DSPs, creating an intelligent and flexible system with an easy path for future enhancements and upgrades.

Ron Huizen *is VP of technology at BittWare, where he focuses on technology direction and new product concept definition. In a previous role at BittWare, Ron oversaw all aspects of product development. Before joining BittWare, he held various roles in electronic product development at Amirix Systems. During his tenure at Amirix, Ron started Cabotel Systems, a spin-off company focused on developing semi-custom electronic products. Prior to Amirix, he worked at Nortel Networks for several years on SS7 switching systems. He holds a Bachelors of Computer Science from Acadia University and a Masters of Computer Science from Carleton University. He can be contacted at rhuizen@bittware.com.*

BittWare, Inc.
603-226-0404
www.bittware.com

New design methodologies offer easier path to custom DSP hardware

By Shawn McCloud

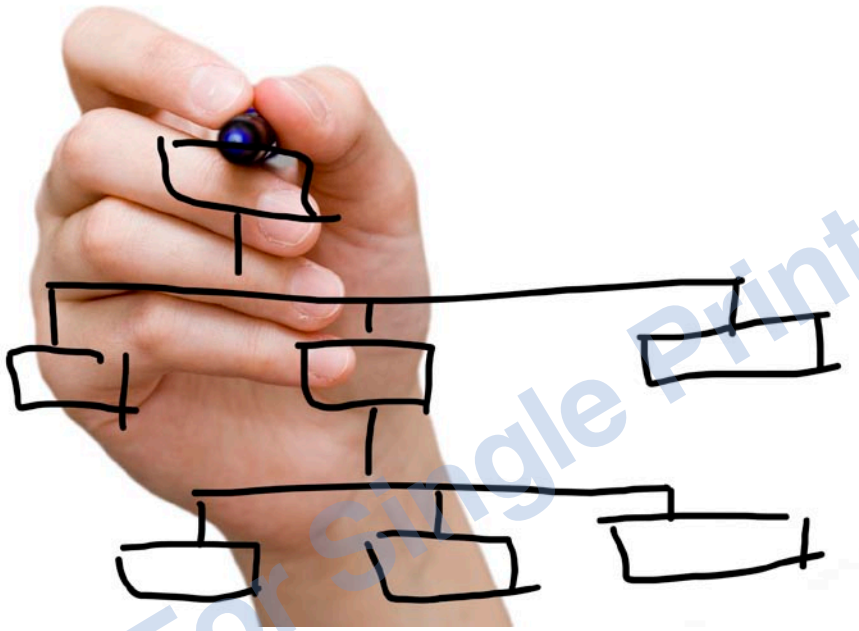

For years, system designers have weighed the trade-offs between the programming flexibility of discrete DSPs against the performance, area, and power benefits of custom hardware. The effort involved in creating DSP hardware was a major deterrent for all but the most performance-hungry applications. However, a new design methodology called algorithmic synthesis *makes it easier to create high-performance dedicated DSP hardware. With this methodology, designers can automatically create Register Transfer Level (RTL) implementations in seconds, compare multiple microarchitectural options, and quickly achieve designs that are optimized for the performance, area, and power consumption needs of systems as diverse as a land-based surveillance system, an airborne reconnaissance system, and a personal night-vision system. With these capabilities, designers find themselves seriously rethinking their overall design flow for implementing DSP algorithms.*

The processing performance required for next-generation compute-intensive applications, including wireless communication and image processing, has created a gap between off-the-shelf DSP performance and market needs. More and more, discrete DSP devices fall short of performance requirements for leading-edge communications and multimedia subsystems designed into modern mil/aero applications. In recent years, system designers have increasingly looked beyond programmable DSPs toward dedicated hardware solutions, such as FPGAs and ASICs that deliver increased levels of performance.

However, manually implementing DSP algorithms in hardware can be an expensive, time-consuming process. Hand-coding hardware descriptions in RTL can take a design team weeks or months, with verification and optimization doubling or even tripling the total time required to implement complex DSP algorithms (see sidebar). This effort and expense meant that ASICs and FPGAs were only used in demanding niche applications. Now, a new class of *algorithmic synthesis* design tools makes it faster and easier to implement DSP algorithms in hardware and put within easy reach hardware implementations that are optimized for performance, area, or power consumption.

RTL: DSP algorithms the hard way

Using traditional Register Transfer Level (RTL) methodologies to create hardware implementations for complex DSP algorithms, design teams had to iterate through several steps, including microarchitecture definition, hand-coding the RTL, and area/speed/power optimization through iterative RTL synthesis. This manual process was, and still is, slow and error-prone, introducing up to 60 percent of the bugs found in RTL through design misinterpretation from original specification. In the final result, both the microarchitecture and technology characteristics become hard-coded into the RTL description. This effect severely limits the notion of RTL reuse or retargeting for real applications and leads to inefficient, overbuilt designs and wasted silicon.

Algorithmic synthesis moves hardware design decisions to a higher level of abstraction. As a point of reference, one could compare the automated C-to-RTL design flow (Figure 1) used in algorithmic synthesis to the traditional DSP software programming flow. In both flows, algorithm designers develop a floating-point model of an algorithm, then convert that to a fixed-point model, typically in C++.

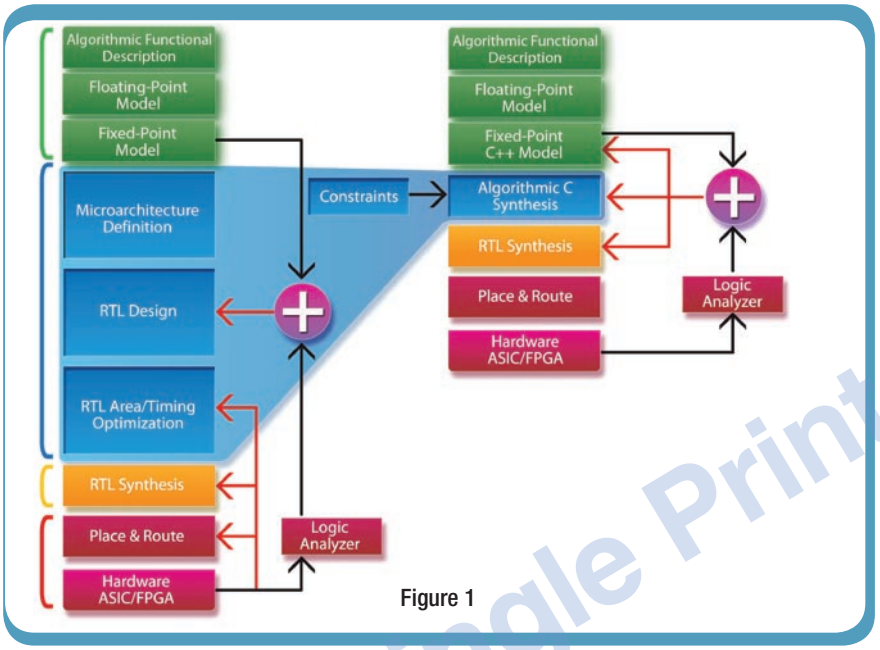

Figure 1

At this point in the traditional flow, software developers use a compiler to automatically compile the C code for an off-the-shelf DSP. With algorithmic synthesis, a hardware designer uses an algorithmic synthesis tool to automatically create an RTL description of DSP hardware that is tuned for a specific ASIC or FPGA technology platform, dramatically shortening the design hardware design flow. In fact, the technology-independent ANSI C++ description enables the algorithmic synthesis tool to target either ASIC or FPGA implementations.

Algorithmic synthesis enables designers to tune the design to exactly match the performance required for a specific application, including latency, throughput, power consumption, and frequency, helping designers avoid the common problem of overbuilt hardware, for example, hardware that exceeds performance requirements at the cost of increased silicon size and/or power consumption. Since the C representation is completely abstracted

from the final implementation, designers can later use the constraints in the algorithmic synthesis tool to easily retarget the same representation for different microarchitectures and ASIC/FPGA implementations.

Abstraction: The key to greater exploration and optimization

Algorithmic synthesis based on pure ANSI C++ speeds the hardware design process 10-100x. By nature, an ANSI C++ source is a purely functional description of the algorithm and therefore independent of the target architecture and technology. Lower levels of abstraction necessarily include more information on hardware structure, which limits the freedom a designer has when exploring alternative microarchitectures. In addition, automating the painstaking RTL creation process means designers have more time to consider architectural decisions that can have a tremendous impact on design performance, area, and power consumption. The user can apply synthesis constraints to specify the target technology (ASIC or FPGA), the amount of parallelism, and desired performance. These constraints, combined with increased productivity, give the designer both the ability and time to explore different trade-offs, resulting in a wider optimization scope for their design (Figure 2).

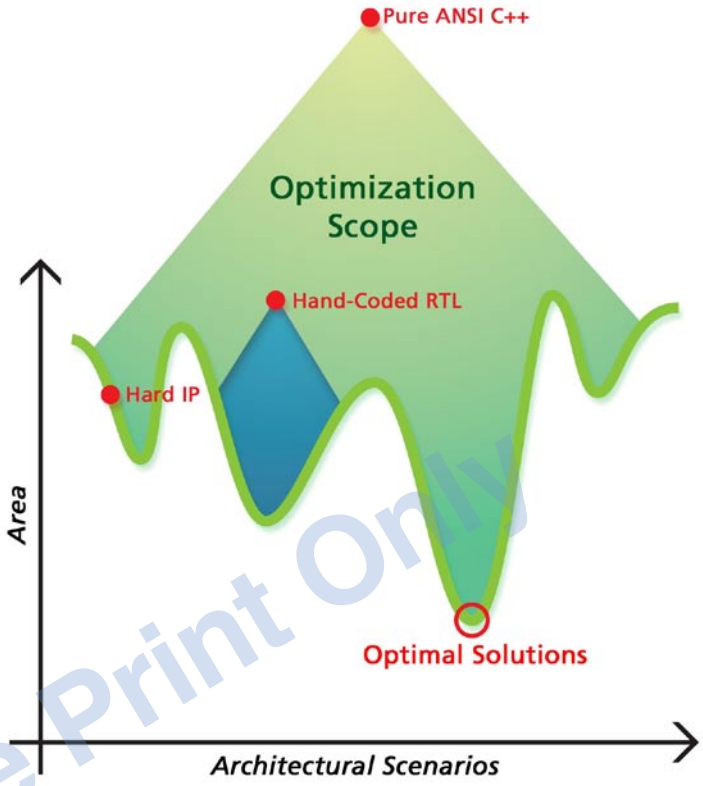

Figure 2

These factors enable designers to create hardware designs more tightly tuned to a specific application than traditional RTL methods. For example, the same pixel-pipe algorithm for video might be used in a land-based surveillance system, an airborne reconnaissance system, and a personal night-vision system, but each implementation would look very different. Trade-offs to achieve these different implementations can be made using a variety of constraints within an algorithmic synthesis tool. These constraints can include loop unrolling or pipelining, loop merging, RAM, ROM, and FIFO array mapping, memory resource merging, and memory bit-width resizing. Using this methodology, hardware designers can easily perform "what if" trade-offs evaluating area, latency, power consumption, throughput, and clock frequency for each microarchitecture, all the while leaving the original pure ANSI C/C++ source unchanged.

Example: Power optimization

The increasing number of battery-powered military applications that rely on power-efficient algorithms – such as wireless communications, portable data devices, and video systems – has made power optimization a higher priority. Well-known tactics like clock gating, optimizing memory accesses, controlling clock rates, and changing state machine encoding, for example, are typical RTL methods for power-efficient design. Most of these design techniques are available through algorithmic synthesis tools.

Typically, the ability to influence factors like power consumption, performance, and area is much greater at higher levels of abstraction. Unfortunately, accuracy of power, performance, and area estimation is inversely proportional to a design's abstraction level (Figure 3). However, the ability to compare power consumption estimates for the various algorithms or microarchitectures within an algorithm gives designers an invaluable advantage and allows them to have a much greater impact on power-related decisions earlier in the design process.

Case study: Finite Impulse Response (FIR)

Using a FIR as an example (Figure 4), let's look at how this simple algorithm can be implemented in multiple ways using algorithmic synthesis. Used to restore signal clarity in transmission systems, the FIR filter is one of the most common components used in signal processing applications. This example will evaluate trade-offs for an eight-tap FIR filter with a performance requirement of 400 MHz, targeted at a 90 nm ASIC technology.

One of the most common structures of a FIR filter is the literal implementation or the direct form implementation, where the data is moved through a shift-register based delay line and each register's output is multiplied with corresponding coefficients. The resulting outputs of all the multipliers are summed up to create the filter's output. Typically, this implementation delivers the highest throughput. As most FIR filter coefficients are symmetrical, this tradi-

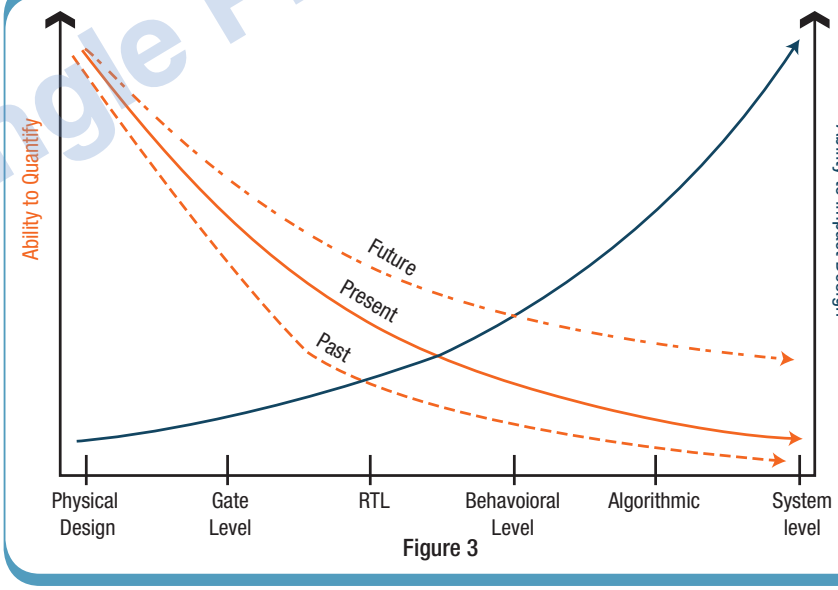

Figure 3

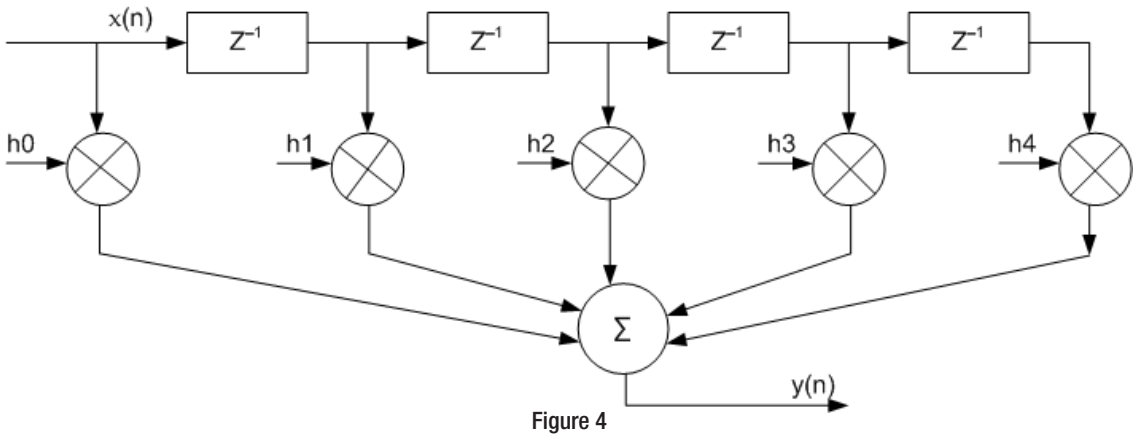

Figure 4

tional architecture could also be optimized by *folding* the structure, thus reducing the number of multipliers required. Folding is a common optimization strategy that helps reduce area utilization by reusing the same area hardware components (such as multipliers) for multiple streams of data. Folding involves a trade-off between the area utilization of the hardware and the higher clock rate required to maintain similar data rates. Figure 5 shows an RTL schematic view of a pipelined implementation of a direct form FIR filter.

Another FIR filter implementation typically used when the filter has a low tap number is the structure wherein the taps are rotated through a shift register with only the end tap being indexed. This implementation typically results in a lower area structure. Figure 6 shows a schematic view of a register-based rotate implementation of a four-tap FIR filter. Of course, there are many other logically equivalent popular FIR implementations such as a transpose format or circular buffer using memory (for a larger number of taps); it is up to the designer to choose one that best fits performance needs. In this article, we will experiment with the direct form and register-based rotate implementations of the FIR filter and examine them with respect to power consumption.

Using an advanced algorithmic synthesis tool, such as Catapult C Synthesis from Mentor Graphics, one can rapidly create various microarchitectures for any given algorithm. For example, a traditional or direct form implementation of the FIR filter can be designed with minimal resources or as a parallel fully pipelined system. Though similar in functionality, the effect on performance – especially with respect to power – for each of these implementations is quite different and can be clearly seen (Figure 7). The fully pipelined solution runs with the highest throughput rate, but also has larger area and higher estimated power usage.

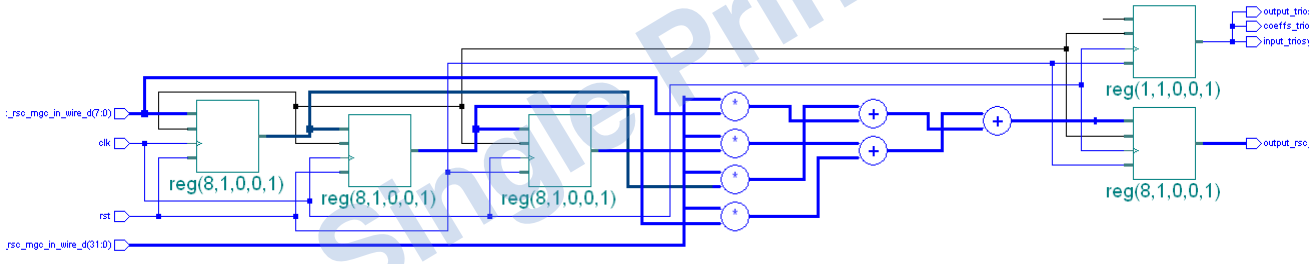

Figure 5

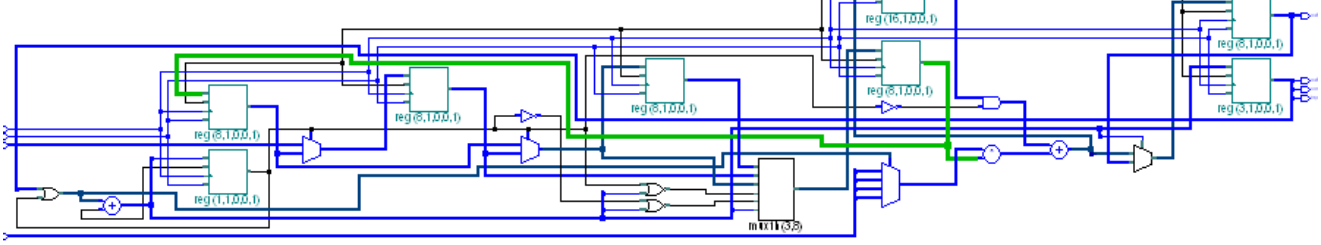

Figure 6

Project Files | XY Plot | Bar Chart | Table

Solution	Status	Total Area Score	Latency Cycles	Latency Time	Throughput Cycles	Throughput Time	Power Consumption
DIRECT_FORM	Passed extract	5140.63	18	45.00	18	45.00	2.66 mW
DIRECT_FORM_THRUPUT_1	Passed extract	30212.40	1	2.50	1	2.50	6.66 mW
DIRECT_FORM_AREA_OPT	Passed extract	4040.64	18	45.00	16	40.00	2.68 mW
REG_ROTATE	Passed extract	6685.45	63	157.50	63	157.50	2.56 mW
REG_ROTATE_LP	Passed extract	4822.77	56	140.00	56	140.00	2.03 mW
REG_ROTATE_AREA_OPT	Passed extract	3368.06	8	20.00	8	20.00	3.03 mW
REG_ROTATE_PARTIAL	Passed extract	6201.16	4	10.00	4	10.00	3.71 mW

Figure 7

Similar experimental implementations can be created for the register-based rotate version of the FIR filter algorithm. As expected, this algorithm's implementation uses less area compared to the shift-register based version and also consumes less power.

Verification benefits

Another important benefit derived from algorithmic synthesis is the greatly enhanced quality of the RTL code. Since the lower-level RTL code is automatically generated from the system specification, there are fewer bugs introduced into the design – up to 60 percent less. By eliminating errors that invariably crop up during manual RTL generation, algorithmic synthesis shortens the verification effort, thereby moving a design to completion faster.

For those bugs that stem from design-related decisions, the same algorithmic source description used to generate the design can be used to automatically create a consistent verification environment including high-speed system models. Advanced algorithmic synthesis tools automatically create SystemC wrappers, allowing designers to rapidly verify their designs 20-100x faster than traditional RTL. A test bench can also be generated that automatically compares the ANSI C/C++ input to the RTL output, providing debug information for specific synchronization points in the case of a simulation mismatch.

Rapidly implementing algorithms in hardware

Even with an automated approach such as algorithmic synthesis, the hardware design process is a departure from the traditional DSP software programming flow. For companies that simply must achieve higher performance, lower power consumption, or more efficient implementations, algorithmic synthesis is a cornerstone design technology. Algorithmic synthesis based on pure ANSI C++ eliminates the manual effort involved in hardware creation, helping designers explore more design options in less time, plus reduce overall design time by 10-100x.

Shawn McCloud *has been the product line manager and director for Mentor Graphics' high-level synthesis technology for the past four years. Prior to that, his positions at Mentor Graphics included technical and product marketing focused on RTL and high-level synthesis. From 1986 to 1994, Shawn worked for Motorola as a senior systems engineer responsible for RISC- and CISC-based microprocessor design. He holds a BS in Electrical and Computer Engineering from Case Western Reserve University in 1986. E-mail him at shawn_mccloud@mentor.com.*

Mentor Graphics
503-685-7000
www.mentor.com

Implementation trade-offs of digital FIR filters

By Ed Rocha

High-precision FIR filters are used in myriad medical, military, and high-volume consumer applications. However, given the choice of stand-alone Microcontroller Units (MCUs), DSPs, FPGAs, or dedicated Finite Impulse Response (FIR) ICs, only the latter balances cost, power, size, and precision performance.

FIR filters form the basis of wireless systems in medical devices, industrial control, consumer electronics, and cellular infrastructure. In fact, they are one of the most common types of digital filters. These linear time-invariant style filters rely solely on current and past input samples and not on past outputs, making the resultant signal directly proportional to the number of taps (summation series) and a coefficient set used to calculate the output.

FIR filters can be calculated relatively easily using several different IC implementations, including microcontrollers, general purpose DSP chips, FPGAs, and purpose-built dedicated FIR devices. While each type has advantages and disadvantages, in systems requiring the combination of high precision, low cost, reasonable size, and power, purpose-built FIRs plus low-cost MCUs represent the best of all worlds. Let's examine each option in detail.

Implementation comparison

Several implementation choices are available to designers. Each has its benefits and risks (Table 1). The first implementation is that of a microcontroller with a built-in hardware multiplier and enough data RAM to store data samples to be used in the computation. The major benefits are that this is a small, low-power implementation. It is also highly integrated. The microcontroller not only implements the FIR filter, but can also function as the overall system controller. The disadvantage is that a microcontroller, even after adding a hardware multiplier on-chip, still has very limited computational abilities and is only a viable solution for relatively simple FIR filters at low sample rates. Also, adding resources to a microcontroller like a hardware multiplier and memory significantly increases the cost of the chip.

The second implementation is that of a DSP. The major benefit of using a DSP is that it provides a large amount of computation horsepower and a virtually infinite amount of flexibility in design and algorithm choices. The DSP can also handle medium to large FIR filters up to very high sample frequencies. The downside of using DSPs, though, is that they are large devices that burn a lot of power. They are also relatively expensive and possibly overkill in systems that run at lower sample frequencies.

The third choice is that of an FPGA. The major benefit of an FPGA is the ability to implement algorithms in hardware without losing versatility. FPGAs can support an almost infinite number of algorithmic choices and provide the highest performance of all. The downside of using FPGAs is that they are also large

FIR basis	Pro	Con
Microcontroller (MCU)	SWaP, integration, CPU decision making	Implements only simple FIR, limited computation, cost increases with complexity
Dedicated DSP	Flexibility, nearly limitless horsepower, implements large FIRs and high sample rates	Power hungry, expensive, often overkill in some systems
FPGA	H/W implementation yet still versatile, very flexible algorithms	Large and power hungry, requires skilled IC designers, and cost per chip is high
Purpose-built FIR chip (such as the SavFIRe)	Fast time-to-market, programming ease, low cost	Fixed functionality, requires host controller

Table 1

devices that burn more power than any of the other alternatives. They also require skilled engineers and considerable design time to implement. The cost per chip is the highest of all competing solutions for a single filter such as the one described earlier.

The final implementation choice is a Simple and Versatile FIR Engine or *SavFIRe*. The dedicated FIR chip is the easiest to program and represents the quickest time-to-market of all solutions. It is the least expensive of all the solutions and is by far the smallest and most power-efficient method discussed. The downside of using this part is that it is a dedicated FIR filter and cannot be reused or reconfigured to be anything else. Using SavFIRe requires some sort of host controller to run the system. However, an MCU with filter data storage can cost as low as $1. The combination of an inexpensive, low-power microcontroller along with the SavFIRe chip represents the best total system solution in terms of power, size, time-to-market, and cost.

Example FIR filter for comparison

To make a reasonable comparison between the different platforms for implementing the FIR filter, it helps to have an example. This way we can not only compute the ability of each platform to perform the task, but also estimate comparisons between the platforms' size, power, and cost.

The FIR filter we will design is a notched low-pass filter and has the specifications shown in Figure 1.

- System sample frequency = 1 KHz
- Notch center frequency = 60 Hz
- Top notch BW = 50 Hz
- Bottom notch BW = 38 Hz
- Notch attenuation = 60 dB
- Pass-band upper frequency = 250 Hz
- Pass-band lower frequency = 260 Hz
- Pass-band ripple = 0.1 dB
- Stop-band attenuation = 60 dB

The approximate number of taps for this filter is 467.

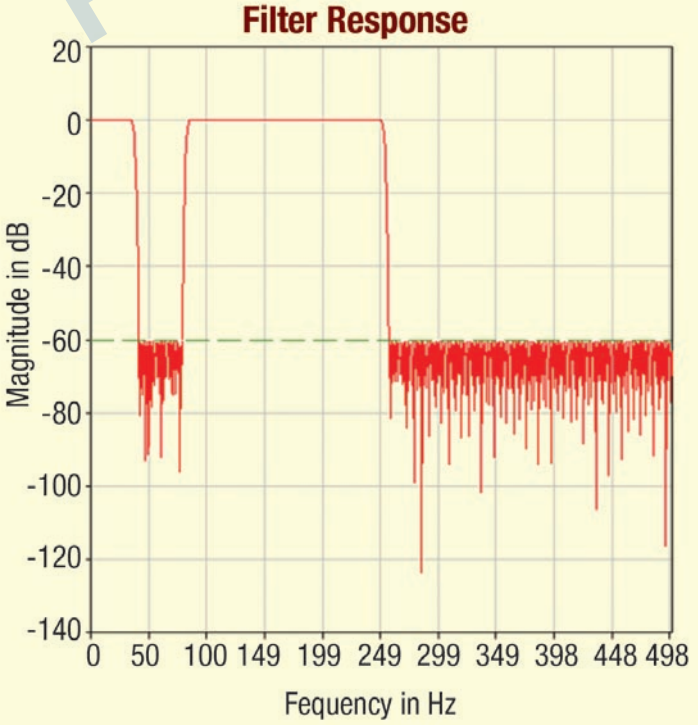

Figure 1

Microcontroller implementation

A good example of a microcontroller that can handle a considerable FIR load is Texas Instruments' MSP430F169. It has both a built-in MAC unit and multiple DMA channels to assist in moving data and coefficients to and from data memory. Another significant specification is its 2 KB of memory. This is important because the microcontroller must have enough memory to buffer data samples equal to the number of taps.

Figure 2 shows the data flow of the FIR filter. The steps needed to calculate each output of the FIR filter are as follows:

- Get input sample from ADC
- Loop N-1 times: Move two input samples from memory to MAC; move coefficient from memory to MAC; perform MAC operation
- Retrieve output from MAC and store to memory

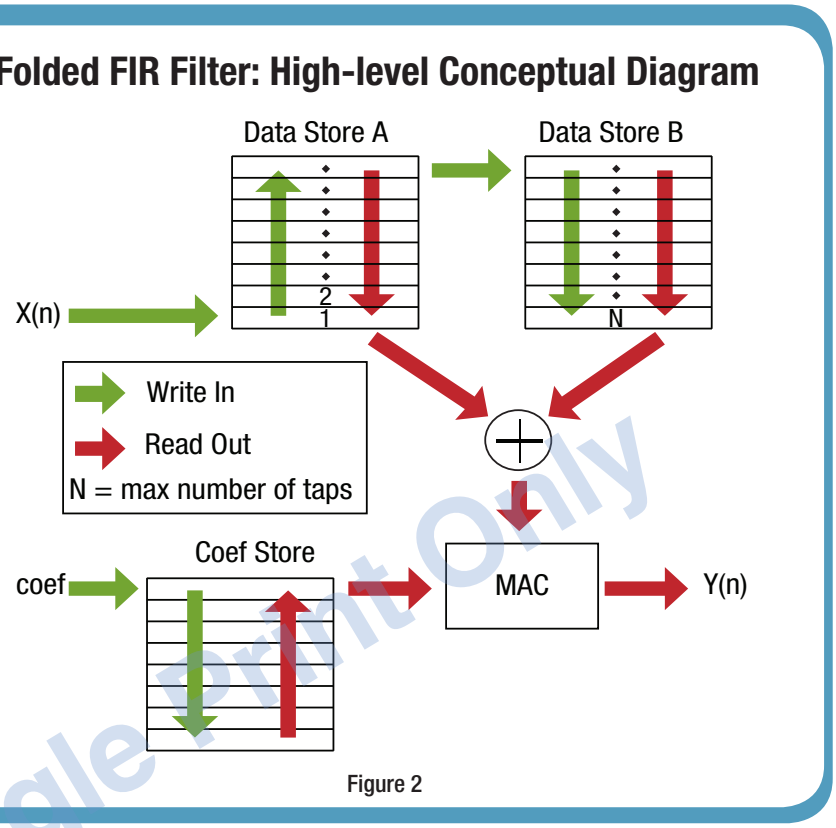

Figure 2

Table 2 represents some benchmarks for the MSP430F169. Using the data in the table, we can estimate the number of clock cycles needed to calculate a 467-tap filter. After some interpolation, we come up with about 3,200 clock cycles. This indicates a maximum sample rate of 2,500 Hz. We are only sampling at 1,000 Hz. If we add a bit of overhead to retrieve samples and deliver results to memory or perform any type of analysis on the samples, it is fair to assume that our filter will result in approximately a 50 percent loading on the CPU. Table 2 specifies the benchmarks of the MSP430F169.

Number of Taps, N	Number of CPU Clocks	Max Sampling Rate (KHz) (100% CPU loading)
32	310	25
50	410	19
100	730	11

Table 2

The MSP430 family only has a 16 x 16 multiplier in its MAC unit. This means that designers are limited to 16-bit data and 16-bit coefficients if they are to achieve the performance just mentioned. In many systems, both data and coefficients can be larger than 16 bits, which could easily result in a performance degradation of 10x or more.

DSP implementation

The steps involved in FIR calculations are very similar for a DSP compared to a microcontroller. Differences occur in the memory usage where the DSP treats the memory as a circular buffer with pointers that automatically update. Also, because of the high amount of parallelism and pipelining in a DSP, most DSPs can execute the entire MAC operation including fetching and storing data in one clock cycle per MAC operation. There is a slight bit of overhead associated with the beginning and end of the computation loops, but this is negligible for a large number of taps.

Using a Texas Instruments lower-end, fixed-point DSP, the TMS320C55x, let's look at our previous filter design. This DSP is capable of 160 MIPS with one 32 x 32 multiply each clock cycle. This equates to 160M/467 taps per sample = 342 KSps at full speed; this is more than adequate for our example application.

FPGA implementation

The available hardware resources in any given FPGA vary widely depending on the FPGA vendor, part family, and the size of the device chosen.

For the 467-tap FIR in our example, let's use the smallest FPGA in the Cyclone II family from Altera. The EP2C5 has about 4.6 K logic elements, 26 M4K blocks (4 Kb RAMs), and 13 embedded multipliers that are each 18 x 18 bits wide. To implement the 467-tap filter in our example, there are many different architectural choices available using the resources in the FPGA.

The simplest architecture is to use a single multiplier and run it 467/2=234 times for each input sample. The sample rate in this example is only 1 KHz, which implies that our multiplier would have to run at:

1 KHz * 467/2 = 234 KHz

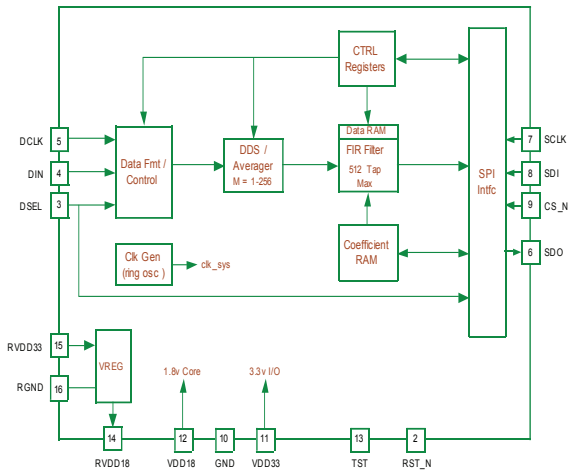

Figure 3

This is easily achievable in today's FPGAs and represents a minimal hardware resource usage.

The multipliers in most low-cost FPGAs are 18 x 18 as mentioned earlier. That means to process larger than 18-bit data or 18-bit coefficients or both, designers will need to use four of these multipliers to create a 36 x 36 multiplier.

Let's assume a typical industry worst case of 24-bit data words and 32-bit coefficients. For a 467-tap filter, there must be enough memory to buffer up 467 data samples and to store 467/2 coefficients. Ignoring some of the implementation details of the FPGA RAM architecture, this implies that a data RAM of 12 Kb and a coefficient RAM of 8 Kb are needed. Translating this into the M4K blocks of the Cyclone II FPGAs would mean a total of five M4K blocks.

SavFIRe implementation

Quickfilter Technologies has developed a SavFIRe chip that is a dedicated FIR filter, supporting up to 512 taps with 32-bit coefficients and anywhere from 12- to 24-bit data words. The QF1D512 supports sample rates from 1 Hz up to 500 KHz. A block diagram of the chip is provided in Figure 3.

The main advantages that come from a dedicated IC all stem from the ability to design a chip specifically for the purpose in mind and optimize all of the associated characteristics of the device. For example, this device is 3 x 3 mm in size; 16 of these will fit in the same footprint

Elcard™ Wireless LAN Modules

Designed for Industrial and Professional Applications

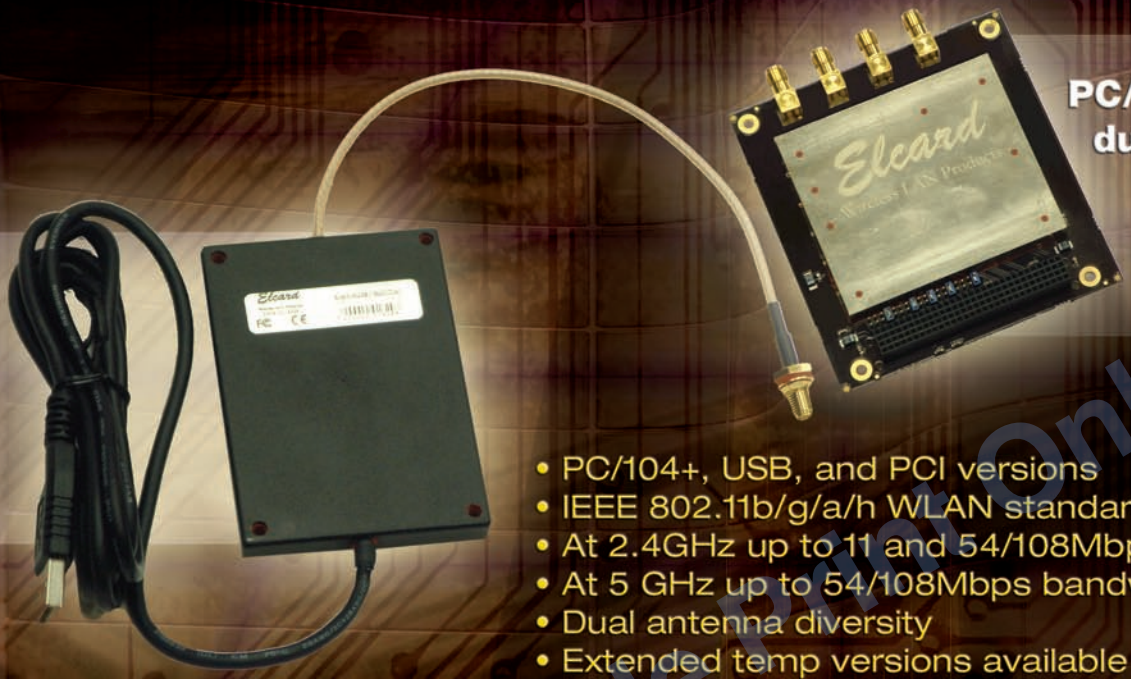

PC/104+ WIB4xx dual WLAN module

USB-attached WLAN

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

Rugged Access Points Available

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

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

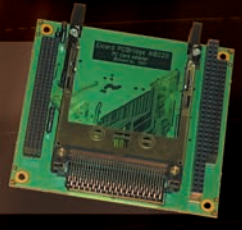

AIB220 PC/104+ Cardbus

- Cardbus/PCMCIA Adapter
- Dual Type I/II or single Type III
- Linux and Win9x/2K/XP support
- 3.3V and 5V card support
- TI PCI1420 chipset

Elcard

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

RTD Embedded Technologies, Inc.

"MIL Value for COTS prices"™

Geode cpuModules™

Pentium® M cpuModules™

8000 MIPS dspModules™

cpuModules™ –40 to +85°C

		Pentium® M				Intel® Celeron®						AMD Geode		
		CMX158886PX1400HR	CMD158886PX1400HR	CMX158886PX1400HR-BRG	CMD158886PX1400HR-BRG	CME147786CX400HR	CME147786CX650HR	CML147786CX400HR	CML147786CX650HR	CMX147786CX400HR	CMX147786CX650HR	CME26686CX333HR	CME27686CX333HR	CMC26686CX333HR
Bus	AT Expansion Bus			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PCI Universal Expansion Bus	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	
	PCI Bus Masters	4	4	4	4	4	4	4	4	4	4		4	
	APIC (add'l PCI interrupts)	9	9	9	9	9	9	9	9	9	9			
CPU and BIOS	CPU Max Clock Rate (MHz)	1400	1400	1400	1400	400	650	400	650	400	650	333	333	333
	L2 Cache	2MB	2MB	2MB	2MB	256k	256k	256k	256k	256k	256k	16K	16k	16k
	Intel SpeedStep Technology	✓	✓	✓	✓									
	ACPI Power Mgmt	2.0	2.0	2.0	2.0	1.0	1.0	1.0	1.0	1.0	1.0			
	Max Onboard DRAM (MB)	512	512	512	512	512	512	512	512	512	512	256	256	256
	RTD Enhanced Flash BIOS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Nonvolatile Configuration	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Quick Boot Option Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	USB Boot	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			
Peripherals	Watchdog Timer & RTC	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	IDE and Floppy Controllers	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	ATA/IDE Disk Socket, 32 DIP	1	1	1	1	1	1	1	1	1	1	1	1	1
	Audio	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			
	Digital Video	LVDS	LVDS	LVDS	LVDS			TTL	TTL	LVDS	LVDS	TTL	TTL	
	Analog Video	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	SVGA	
	AT Keyboard/Utility Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PS/2 Mouse	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	USB Mouse/Keyboard	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
I/O	RS-232/422/485 Ports	2	1	2	1	2	2	2	2	2	2	2	2	2
	USB 2.0 Ports	2	4	2	4									
	USB Ports					2	2	2	2	2	2	2	2	2
	10/100Base-T Ethernet	1		1		1	1	1	1	1	1	1	1	
	ECP Parallel Port	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	aDIO(Advanced Digital I/O)	18	18	18	18	18	18	18	18	18	18	18	18	18
	multiPort(aDIO, ECP, FDC)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SW	ROM-DOS Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	DOS, Windows, Linux	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

utilityModules™ –40 to +85°C

dspModules™
- Coprocessors
- Accelerators

Specialty I/O
- Pulse width modulator
- Incremental encoder
- Opto-isolated MOSFET

Frame Grabbers
- Single or multi-channel
- MPEG-2 compression

Video Controllers
- Analog VGA
- TTL and DVI panels

Communication Modules
- Copper or fiber Ethernet
- USB 2.0 and Firewire
- CAN Bus & CAN Spider
- Dual Synchronous Serial
- Quad Serial w/ Ethernet
- Octal PCI Serial

Wireless Telematics
- GSM, GSM-R, CDMA
- EDGE, GPRS, SMS
- GPS, Wi-Fi, Bluetooth

Motion Controllers
- DC motor controllers
- Synchro, resolver, LVDT

Power Supplies
- 50/75/83/88/100 Watts
- Wide input range
- ATX Power Supply
- UPS backup
- MIL-STD-704/461

Mass Storage
- 1.8/2.5" IDE & PCMCIA
- CompactFlash

IDAN™— Intelligent Data Acquisition Node
- Easily build your PC/104 system
- Rugged PC/104 stackable framed modules
- Quick interchangeability and expansion
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- Standard PC connectors
- Optional MIL-SPEC paint & shock mounts
- –40 to +85 °C

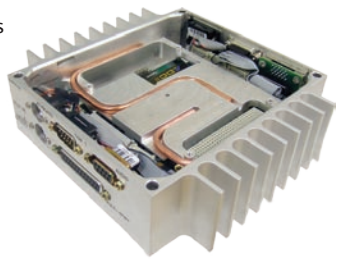

Full Product Line and Pricing Online

A Founder of the PC/104 Consortium • ISO9001:2001 Certified

HighRel PC/PCI-104 Modules and Systems

–40 to +85°C

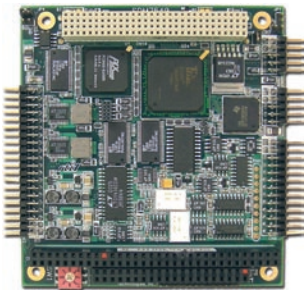

Autonomous SmartCal™

Wireless Telematics

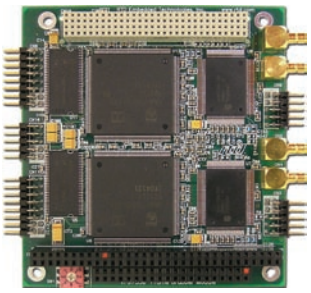

Frame Grabbers

dataModules® –40 to +85°C		Smart A/D		Analog I/O					Digital I/O						
		SDM7540HR	SDM8540HR	DM6210HR	DM6420HR	DM6430HR	DM7520HR	DM6620HR	DM6812HR	DM6814/16HR	DM6856HR	DM6888HR	DM6956HR	DM7820HR
Bus	AT Expansion Bus	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	PCI Expansion Bus Master	✓	✓				✓							✓
	McBSP Serial Ports	✓	✓				✓							
Analog Input	Single-Ended Inputs	16	16	16	16	16	16							
	Differential Inputs	8	8		8	8	8							
	Max Throughput (kHz)	1250	1250	40	500	100	1250							
	Max Resolution (bits)	12	12	12	12	16	12							
	Input Ranges/Gains	3/7	3/7	3/1	3/4	1/4	3/6							
	Autonomous SmartCal	✓	✓											
	Data Marker Inputs	3	3		3		3							
Conversions	Channel-Gain Table	8k	8k		8k	8k	8k							
	Scan/Burst/Multi-Burst	✓	✓		✓	✓	✓							
	A/D FIFO Buffer	8k	8k		8k	8k	8k							
	Sample Counter	✓	✓		✓	✓	✓							
	DMA or PCI Bus Master	✓	✓		✓	✓	✓	✓						✓
	SyncBus	✓	✓				✓							
Digital I/O	Total Digital I/O	16	16	16	16	16	16	16	48	18/9	32	64	32	48
	Bit Programmable I/O	8	8		8	8	8	8	24	6/0				48
	Advanced Interrupts	2	2		2	2	2	2	2					2
	Input FIFO Buffer	8k	8k		8k	8k	8k							4M
	Opto-Isolated Inputs										16	48	16	
	Opto-Isolated Outputs										16	16		
	User Timer/Counters	3	3	3	2	3	3	3	3	3				10
	External Trigger	✓	✓		✓	✓	✓	✓	✓					✓
	Incr. Encoder/PWM									3/9				
	Relay Outputs												16	
Analog Out	Analog Outputs	2	2		2	2	2	4						
	Max Throughput (kHz)	200	200		200	100	200	200						
	Resolution (bits)	12	12		12	16	12	12						
	Output Ranges	4	4		3	1	4	4						
	D/A FIFO Buffer	8k	8k				8k	8k						

RTD FieldPads™

- Ruggedized, embedded computer systems
- User-specified CPU and PC/PCI-104 expansion
- Weathertight components
- Integrated 6.5-inch video panel, keyboard
- Heat pipes for high performance CPUs
- User-defined MIL connectors
- Internal and external battery packs
- Expand with any RTD PC/PCI-104 product

Industrial FieldPad™

Ideal for control and monitoring of processes on factory floors or industrial installations. Mounting flanges allow the unit to be installed on machinery or walls, enabling standard PC access in a rugged enclosure resistant to industrial environments.

Tactical FieldPad™

Designed for mobile and portable applications where the angled panel and ergonomic design allow for optimal viewing with flexible positioning. Data collection/downloading and information access accomplished through wired or wireless connections.

HiDAN™ and HiDAN*plus*™— HighRel Intelligent Data Acquisition Node

- HiDAN is a rugged, watertight enclosure for a stack of PC/104 modules
- HiDAN*plus* combines the modularity of IDAN with the environmental ruggedness of HiDAN
- Integrated tongue and groove O-ring for environmental sealing and EMI suppression
- Structural heat sinks and heat pipes
- Optional cooling fins
- Milled aluminum frames
- **Stackable signal raceway**
- Optional MIL-SPEC paint
- MIL I/O connectors
- Shock-mount optional
- **–40** to **+85** °C

www.rtd.com
Specifications, manuals, drivers, and plant tour

RTD Embedded Technologies, Inc.
103 Innovation Blvd • State College, PA 16803
T: 814-234-8087 • F: 814-234-5218

"Accessing the Analog World"®

> **It literally takes minutes for an inexperienced designer to design a high-precision filter.**

area of an FPGA configured for the same function. Power and cost can also be fully optimized for the function at hand.

The disadvantage of using an IC tailored for a specific function like FIR filters is that it is not very flexible for use in applications outside of FIR filtering. Another requirement is that SavFIRe cannot be used autonomously in a system. There must be some sort of host controller to configure the chip and run sample data through it.

For comparison to the other implementation methods, let's use the same 467-tap FIR example from above. This filter can be created with 32-bit coefficients and use from 12- to 24-bit data samples with the current architecture in SavFIRe. Also, 1 KHz is no problem since SavFIRe works from 1 Hz to 500 KHz.

So what about design effort? There is also a development package, QF1D512-DK, comprising a development board and software development tool. The software tool enables a designer to design a filter by simply specifying the filter parameters. Once the parameters are entered, the software generates all the filter coefficients and necessary register values to get the chip up and running, making design time absolutely minimal. It literally takes minutes for an inexperienced designer to design a high-precision filter.

Another advantage to using SavFIRe is power consumption. The average power consumption for SavFIRe running a 467-tap filter at 1 KHz is about 50 µW. This is orders of magnitude lower than any competing solution. The power scales directly with the sample rate, giving an equivalent power of 0.5 mW at 10 KHz and 5 mW at 100 KHz.

Implementation comparison

We have outlined four popular choices for implementing a high-precision digital FIR filter. All of these implementations have their bright spots as well as their downfalls, which are summarized in Table 3. Again, the SavFIRe choice appears to come out ahead.

***Ed Rocha** is a consultant and the owner of Ed Rocha Consulting. He consulted on the architecture and development of Quickfilter's software development tools for both the QF1D512 and QF4A512. He can be contacted at erocha@quickfilter.net.*

Quickfilter Technologies
214-547-0460
www.quickfiltertech.com

	Compute Power	Flexibility	Size	Power	Cost	Time-to-Market
Micro-controller		✓		✓		
DSP	✓	✓				
FPGA	✓	✓				
SavFIRe	✓		✓	✓	✓	✓

Table 3

Annapolis Micro Systems

The FPGA Systems Performance Leader!

FOPEN **Radar Systems** **Software Defined Radio** **FLIR**
SIGINT **ELINT** **Digital Receivers** **Recording Systems**

Annapolis
Micro Systems

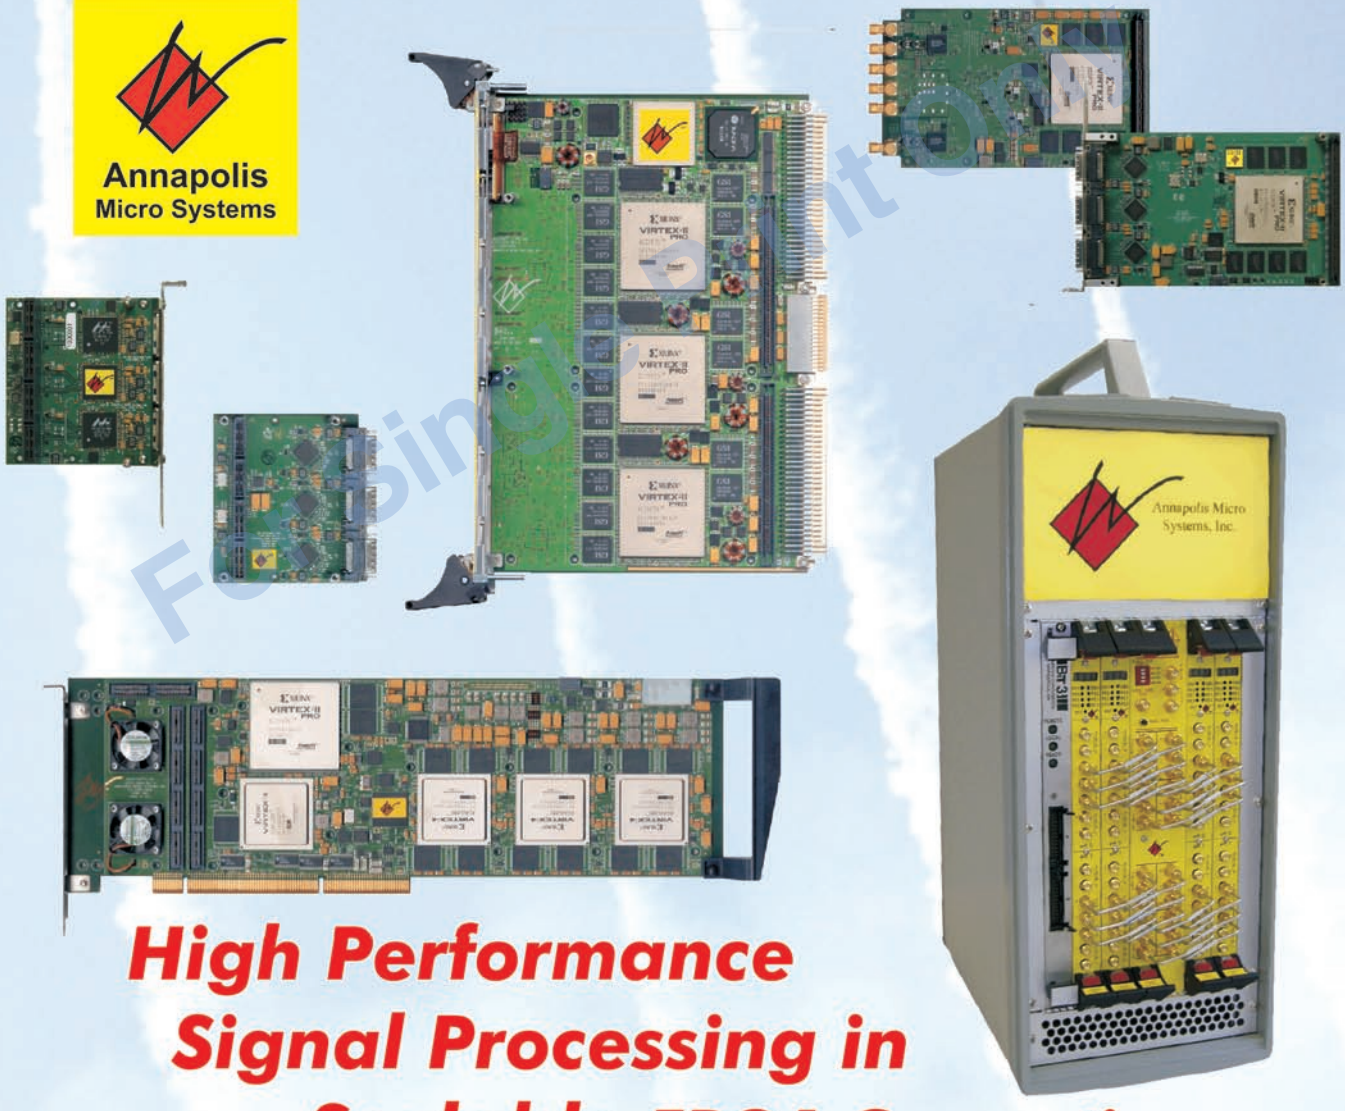

High Performance Signal Processing in Scalable FPGA Computing

Above and Beyond -----
FPGA Acceleration

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

Photo courtesy of USN

Focus on SWaP, energy management, and QoS

An exclusive interview with Bill Kehret, founder and CEO of Themis

EDITOR'S FOREWORD

Regretfully, my time with Bill Kehret ended all too soon. This edited and excerpted Q&A is so chock full o' technical goodness that I hated to leave anything out. But sadly, we had a space budget. However, this one is a must-read for anyone dealing with boards, software, power, multicore, or fabrics. Uh, pretty much all embedded applications. – Chris Ciufo

MIL EMBEDDED: ***I think of Themis mostly as a SPARC server provider. How's this processor working for you?***

KEHRET: First, let me say that we also field innovative, multicore Power-64 architecture products, as well as Intel and AMD multi-socket, multicore products; we're not all about SPARC.

Specifically, regarding your question about SPARC servers, I'd like to point out that SPARC had the first SIMD instruction set enhancements, the first dual core processors, and now the first eight-core SMP chips. Sun's latest chip multithreading effort has industry-leading floating point as well as hardware support for up to 64 threads.

Themis has been a pioneer in bringing these innovations to the embedded marketplace, having deployed many thousands of SPARC and Solaris "seats." We continue to innovate with all of these technologies for ecosystems as diverse as VME, ATCA [AdvancedTCA, a PICMG specification] and their variants, as well as rack-mounting blade servers. We've even deployed blade solutions integrated with storage and InfiniBand switching, complete with flow-through liquid cooling.

> **"... SPARC had the first SIMD instruction set enhancements, the first dual core processors, and now the first eight-core SMP chips."**

MIL EMBEDDED: ***How are multicore processors [AMD, Intel, IBM] affecting your market?***

KEHRET: Multicore is the right answer for SWaP-constrained applications. Sun probably got that right first, though it's rarely credited with those innovations. Multicore provides an energy-efficient alternative to scaling core speed and can be a cost-effective way of fitting SMP onto a small platform. The solutions provided by AMD, Intel, and IBM vary greatly. For embedded applications, the adoption of multicore is lagging behind the commercial enterprise because embedded operating systems and applications evolved in a high clock rate, uniprocessor environment. This is especially true of real-time OSs that rely on a simple hardware model.

For our traditional SPARC market, customers using Solaris are accustomed to running SMP, whether it is implemented over multiple sockets or as a single socket with multiple cores. This class of customers will make the move to multicore and multithreading over a shorter timeframe than those whose experience is limited to uniprocessor environments. It is a natural consequence that all Themis technology platforms have multicore variants.

MIL EMBEDDED: ***Which switched fabrics are having the greatest impact on your customers? Why?***

KEHRET: Switched Ethernet, InfiniBand, and PCIe are the standards being deployed by our customers. Our OEM customers are coming up to speed on the use of these fabrics, as they move from "stovepipe" solutions, to the end customer's computing needs, to more open and scalable architectures. Customers and users alike perceive Ethernet as the standard solution to their system interconnect requirements. Sensor I/O is making a change to support Ethernet as opposed to the more traditional serial or 1553 connections. Since there are many legacy interfaces in service, the move to fabrics for I/O will come slowly.

Also, we're increasingly seeing intersystem communication moving to 10 GbE along with the emergence of 10 GbE MACs, with TOE [TCP/IP Offload Engine] and HW encryption support. For TOE and encryption, SoC [System-on-Chip] processing platforms will surely accelerate this trend. For the future, the choice will depend on whether the application can tolerate the protocol stack latency as well as the nonpredictable delivery of data over the network.

MODULAR SOLUTIONS FOR MIL-STD-1553

MAXIMIZE YOUR SPACE CUSTOMIZE YOUR 1553

PMC-1553 Dual SUMMIT
Dual UTMC 1553 SUMMIT Controller
Bus Controller, Remote Terminal,
Bus Monitor modes
- 1 or 2 UTMC SUMMIT
- UT69151DX-GPC
- BC / RT / BM
- On-chip Transceivers

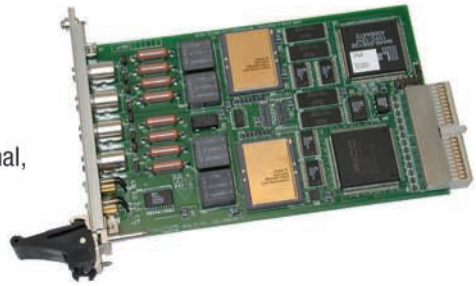

CPCI-1553 SUMMIT
Single or Dual UTMC 1553
SUMMIT Controller
Bus Controller, Remote Terminal,
Bus Monitor modes
- UTMC SUMMIT
- UT69151DX-GPC
- 128Kbytes SRAM
- On-chip Transceivers

PMC-1553 DDC
Single or Dual DDC ACE 1553 Controller
Bus Controller, Remote Terminal, Bus
Monitor modes
- DDC ACE BU61580S3
- 128Kbytes SRAM
- On-chip Transceivers

Whether for PCI, CPCI, IP or PMC, **ALPHI** provides the MIL-STD-1553 solution that you need for your Mission Critical Systems.

From 1 to 4 channels, Aeroflex SUMMIT, or DDC ACE, MiniACE and Micro ACE.

ALPHI also offers an intelligent choice of IndustryPack Carriers allowing you to create your own dedicated MIL-STD-1553 solution.

Contact us at: 480-838-2428
or visit our Web site at
www.alphitech.com

ALPHI
TECHNOLOGY CORPORATION

MIL EMBEDDED: ***Can you contrast the market for air-cooled, rack-mount equipment versus embedded, rugged, deployed equipment? What are the trends?***

KEHRET: If the distinction is a choice between rack-mounting COTS chassis, including "blade servers" and more robust bus-based board ecosystems such as VME and ATCA, then I'd have to say that the market for generic COTS boxes is growing at the expense of all of the board-based ecosystems. But there is a SWaP and cost penalty to be paid for "cocooning" lightly ruggedized COTS boxes, which can typically handle less than half the *"g"* loads of average VMEbus boards, with vibration modal frequencies significantly lower than for a VME ecosystem.

MIL EMBEDDED: ***"Mission critical" describes many applications today. What are the technical issues to consider in a mission-critical environment?***

KEHRET: Regardless of how robust we make the hardware and software, redundancy is necessary to achieve truly high levels of availability. The only question is how to size and control the overhead. Fully redundant systems are a luxury few applications can afford, so the trick is turning *over-provisioned* resources into *productive* resources. Real-time computing resource management is critical to achieving contracted Application Quality of Service [AQoS] for true mission-critical computing. Flexible policies determine the priority of applications loaded on the system. The system's mission drives those policies. Real-time automated control is the only way to wring out all of the available performance for a given amount of computing resources and application load.

MIL EMBEDDED: ***What are the technology trends you're seeing, looking out three to five years?***

KEHRET: Connectivity, both in terms of hardware and software infrastructure, will scale vertically and horizontally by two powers of 2. By this I mean that endpoint-pair bandwidths will double twice and the number of switch endpoints will double and possible quadruple. Entities under management will scale equivalently. Applications will continue to be virtualized, to support the scaled number of "clients."

Security will be pushed further into the IT fabric and processor chips so that multilevel security will be supported on operating systems as well as data links, making networks and IT centers more impervious to malicious attack. Service Oriented Architectures [SOA] will fuel the demand for scaling processor cores and threads and the virtualization of OS and application instantiations on those processing resources. Automatic computing resource management and policy languages will become as commonplace as high-level hardware definition languages and "silicon compilers," as required to cope with the avalanche of computing complexity.

MIL EMBEDDED: ***As processors and attendant systems have become more powerful, there is a need to develop clever cooling techniques. What are today's challenges, and how are Themis and other companies addressing the problems?***

KEHRET: We think there's lots more life left for air-cooled systems, and conduction cooling is a key to extending the life of air-cooled systems. The question is where to put the heat sink. We like to use as much of the board surface area as possible for the heat exchanger (heat sink) on our high power density boards. We pioneered four-way SMP servers in the VME form factor, using discretionary pins and 12 V power to get the power into the slot and wall-to-wall heat sinks that slid into the VME board card guides. Interestingly, VITA 46/48 takes a page out of this book and adds liquid flow-through cooling in the bargain.

The trouble with bus and board open standards is the freedom that vendors have for interpretation. ATCA especially suffers from all the choices of how to use cooling air channels on the board. Our industry likes to go crazy with modularity. The downside of this is that cooling air channels get messed up by a proliferation of AMCs [Advanced Mezzanine Cards].

At the level of the shelf or chassis, there is very little effort made to regulate the available slot-to-slot inlet air pressure, so much of the touted power density headroom is fiction. These problems put a special burden on the system integrator and can, in part, be mitigated by the ecosystem vendor. However, that means the vendor that builds the high power density boards also needs to design the shelf/chassis and its cooling system.

Bill Kehret *has been active in the embedded computing industry for more than 30 years and CEO of Themis since it was founded in 1989.*

Themis
510-252-0870
www.Themis.com

Coordinated Performance

with Phoenix Systems

Real-Time Multi-Processors with High-Speed Serial Communications

Powerful

- Customer programmable Xilinx FPGA compute engines
- Multiple PowerPC processors
- Multi-channel Gbit/sec backplane communications
- Fiber Optic Front Panel I/O, including Serial FPDP
- High-speed Analog and Digital I/O

Flexible

- Combine FPGA and PowerPC Processors to match the application requirements
- Boards or complete sub-systems
- Open standards including VITA 41/VXS and VITA 42/XMC
- Air-Cooled to rugged conduction-cooled versions

Innovative

- Zero Latency Switch technology
- TransComm™ inter-processor communications suite for simplified application development

Processing and FPGA - Input/Output - Data Recording - Bus Analyzers

For more information, please visit
http://processor.vmetro.com/phoenix
or call (281) 584-0728

VMETRO
innovation deployed

Guest editorial: Intel commits to embedded, communications

By Anthony Neal-Graves

EDITOR'S FOREWORD

I'm often asked if pure commercial telecom and communications equipment applies to the defense and aerospace markets. It most certainly does. This question is usually posed by someone with a "hidden agenda" hoping to cast aspersions on vendors and supporters of telecom markets. Pure commercial telecom equipment – such as that found at the edge of the cloud in special-purpose systems such as packet shapers, voice-to-text-to-voice systems, and other specialty applications – has direct applicability to the military. This equipment forms the backbone of the non-classified sections of the Global Information Grid, reachback installations, and provides interoperability with the rest of the civilian "POTS" world. And no company has been more involved with these telecom markets than Intel.

To address those naysayers who say COTS suppliers don't have a long-term commitment to telecom (and hence to A&D), I present this letter from Intel showing the company's strong commitment to these applications, and to telecom in general. Intel prepared this statement exclusively for our readers at my request. – Chris Ciufo

Intel is completely committed to the embedded and communications markets, as demonstrated by its 30 year history supporting products for long lifecycles of at least 5 to 7 years, and that it continues to align its product roadmap with real customer needs for performance and power. This could not be possible without the rich array of hardware and software providers in Intel's ecosystem. Intel first launched an ecosystem program in 1998, named the Intel Communications Alliance (ICA) in 2002. Intel and its ICA currently support over 3,000 customers in more than 30 market segments – a testament to Intel's commitment to work with industry developers to bring new solutions and standards to wide adoption.

For example, real world deployments and adoption by leading Telecom Equipment Manufacturers (TEMs) and service providers worldwide have made AdvancedTCA (ATCA) the de facto platform for the fastest growing IP services and wireless segments, such as IPTV and IP services. In 2006, Intel saw a doubling in volumes of ATCA products shipped to the communications market. Intel's recently released fifth generation ATCA Single Board Computer (SBC) brings to market all the performance and performance per watt benefits of the Dual-Core Intel® Xeon® Processor 5100 Series to these customers. With a wide variety of ATCA products currently deployed on Intel Architecture, IA is clearly the leading architecture of choice for ATCA.

In addition, in the embedded industry, Intel is starting to see momentum pick up in market segments such as Military and Government. These customers have the same requirements for rugged, reliable and long life systems as the telecommunications market – and ATCA and MicroTCA are a natural fit. Intel continues to work closely with its ecosystem and standards bodies such as PICMG to ensure that Intel Architecture based solutions are available on ATCA and MicroTCA for customers in all these market segments.

Anthony Neal-Graves, General Manager
Modular Communications Platform Division, Intel

When your mission is critical...

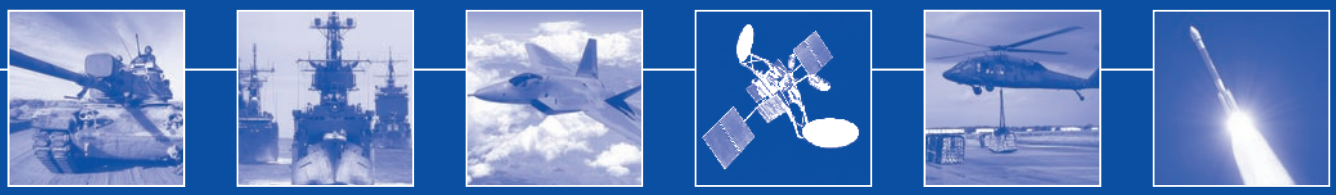

...DEPEND on Birtcher Card Guides & Conduction Cooled Assemblies

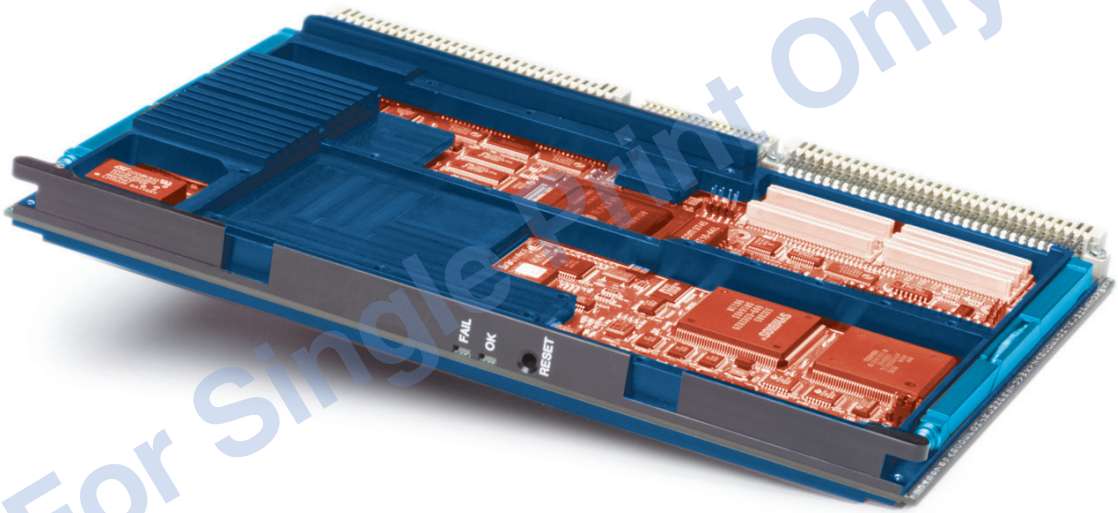

Wedge-Tainer™

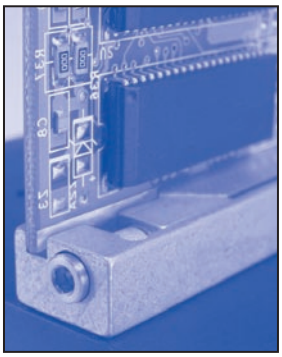

- Designed to comply with MIL-STD-810D and MIL-STD-901D in addition to MIL-E-5400

Wedge-Lok®

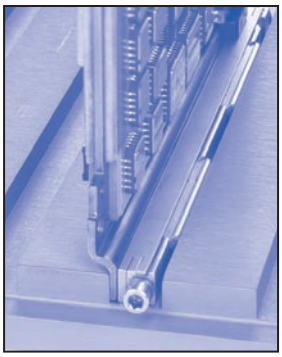

- Meets the requirements of drawing 84103, Defense Supply Center, Columbus

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

Birtcher

HOLDING IT ALL TOGETHER

www.birtcherproducts.com • 888-957-8291

Linux meets the needs of security-critical, network-centric computing

By Dr. Inder M. Singh

The flexibility and availability of Linux applications provide the best way to manage costs, leverage commercial technologies, and provide flexibility in the military's emerging network-centric computing systems. Separation kernel technology provides the framework for meeting the critical security requirements of these systems.

The information age dramatically changed the modern battlefield. The U.S. military has been deploying a new network-centric framework characterized by information sharing – shared situational awareness and knowledge of the commanders' intent throughout the force, regardless of the individual's or unit's precise location. Under the direction of the U.S. Department of Defense's (DoD's) Office of Force Transformation (www.oft.osd.gov), all advanced weapons platforms, sensor systems, and command centers are to be linked via the Global Information Grid (GIG). Still in its burgeoning phase, Network Centric Operations (NCO), also known as *network-centric warfare*, has already proven that networked forces can be more effective than those using traditional means of communication.

Implementation of the DoD's massive NCO systems will be the most ambitious application of the network-centric computing concept ever attempted, and it is likely to become a model for building many large security-critical, network-centric systems in the commercial market as well.

Designing an extended network on this scale is obviously a massive undertaking. The military recognizes that designing its "system of systems" to leverage the flexibility and availability of Linux applications is the best way to manage costs, leverage commercial technologies, and provide flexibility in the network.

Benefits of Linux

Linux has played a major role in the development of networking, and it is widely used across the Internet infrastructure. The wide availability of Linux applications for defense systems, combined with other benefits of using open source technology, can help to reduce the cost and time for building large, complex applications. Many of the subsystems involved in NCO (see Table 1) use embedded processors and software, and some are used in other parts of these network-centric systems. Linux is unique in its scalability across this large range of available hardware.

Moving toward security

The difficulty in transferring information between networks with different security classifications makes designing the NCO tricky. While conventional wisdom once stated that open source software ran the

Subsystems in Network Centric Operations
Weapons systems
Sensors
Controls
Servers
Workstations
Personal computers
Handheld terminals

Table 1

risk of malicious developers inserting code that would make networks vulnerable to attack or misuse, today the tables have turned. In fact, open source Linux is a better match for mission-critical applications than other proprietary COTS operating systems, as its power and flexibility lie in the robust code that anyone can view and improve upon.

However, the security requirements of mission-critical military systems that handle multilevel classified data cannot be met by any existing operating system, including Linux. Security has often been addressed by placing application guards and security functions at appropriate places in the network, thereby continuing to leverage Linux's flexibility while addressing security requirements. But ad hoc solutions of this kind are not easily scalable to the large, ambitious systems envisioned in NCO. A more powerful solution is provided by separation kernel technology. Secure separation kernels offer the ability to sequester the Linux kernel in a separate partition, separate from highly trusted software that provides the mission-critical security functionality. These secure separation kernels are designed to meet the highest levels of embedded software security.

Separation kernels enable one operating system (the "guest" operating system) to run within the same environment as another, supporting multiple operating system environments on the same physical hardware. With a separation kernel, security-related functions can be run in separate, highly trusted partitions, isolated from the Linux partitions, so that the security of the system is not dependent on the operating system. At the same time, legacy Linux applications software can be run unmodified, saving dollars and staff years in programming time.

Combining separation kernels and hardware virtualization technologies such as Intel's VT provides high-performance operation of "guest" operating systems running on separation kernels, comparable to running the operating system directly on the hardware. Unlike traditional security kernels, which perform all trusted functions for a secure operating system, secure separation kernels partition data and system resources, providing dependable data isolation between partitions, controlling information flow among the partitions, and enforcing policy for how partitions can communicate.

The future of Linux networks

Many companies and organizations – such as the nonprofit Network Centric Operations Industry Consortium (NCOIC) at www.ncoic.org – are working toward bringing Linux to the forefront in secure, network-centric operations, integrating open standards into a global framework based on a common set of principles and processes to facilitate the global deployment of network-centric applications.

With its international roster of defense contractors such as General Dynamics, Boeing, and Lockheed Martin, as well as others in security-sensitive industries, the NCOIC is working to develop a secure information management architectural framework for collaboration and projects such as the GIG and others.

As part of these efforts, LynuxWorks has developed LynxSecure, an embedded secure separation kernel that will be certifiable to Common Criteria EAL 6+ and DO-178B Level A. Its approach to certification, correct by construction, provides the most cost-effective approach to certifying a secure kernel and minimizing the costs of sustaining engineering throughout a program's life.

Dr. Inder M. Singh, *LynuxWorks chairman, served as CEO until 2006. He has founded and led numerous companies including Exelan and Kalpana. He was formerly board chairman and president for the Embedded Linux Consortium. He holds PhD and M. Phil. degrees in Computer Science from Yale University, and an MSEE from Polytechnic Institute of New York. He can be contacted at inside@lnxw.com.*

LynuxWorks
800-255-3969
www.lynuxworks.com

The costs of doing business in the new RoHS world

By Doug Patterson

In addition to the design flaws and increased system failure rate it has produced, the inane RoHS compliance mandate has cost multiple industries millions of dollars in new development and remediation costs. Many questions still remain as to the usefulness of this directive in view of the performance sacrifices manufacturers and end users have had to endure.

Adapting to change and assuming risk are generally considered costs of doing business. As industries and technologies shift and new challenges present themselves, businesses must be ready for these costs. However, when an irrational and politically motivated change within an industry forces companies to assume potentially detrimental risks without proof of a truly meaningful benefit, it becomes something more than just the cost of doing business. It becomes *the costly way* of doing business.

The new way of doing business

Within a very short time frame and with seemingly utter disregard for the lasting impact it will have on the reliability of critical embedded systems, the RoHS directive has forced our industry to adopt what is touted to be a more environmentally safe alternative – the elimination of lead (Pb) in electronic components.

But, time and again, we continue to see far more negative consequences than perceived environmental benefits. Research dating back several decades has documented cases of tin (Sn) whiskers causing failures in several mission- and life-critical applications. Missile and satellite programs, medical devices like pacemakers, consumer electronics, and the nuclear power generation industry have all documented and suffered from the detrimental effects of tin whiskers. Wasn't this problem one of the reasons tin-lead solder alloys came into play in the first place?

The price to be lead-free

The RoHS directive has already incurred astronomical costs. The University of Maryland's Center for Advanced Life Cycle Engineering (CALCE) has documented several ongoing failures within electronics systems specifically related to tin whisker growth. A 2002 report on these failures states that one particular recall within a missile program cost taxpayers upwards of $5 million, and that was more than five years ago.

In a presentation at the IEEE's 2005 Microelectronics Reliability & Qualification Workshop in Manhattan Beach, California, Northrop Grumman noted several satellite programs from 1998 to 2002 that completely or partially failed, including the Galaxy 3, Solidaridad 1, and HS 601. With this known problem of tin whisker growth, the reason why the RoHS directive is being so adamantly upheld must be questioned. Without any plan in place to identify an alternative solution, RoHS eliminates an essential element (lead) in compounds that helps lessen the formation of tin whiskers.

The same 2005 IEEE presentation identified some potential mitigation techniques (Table 1), but nothing proved reliable enough to entirely eliminate tin whisker formation. In fact, industry research shows that conformal coating hasn't stopped whisker growth. In a 2000 paper entitled "Effects of Conformal Coat on Tin Whisker Growth," Unisys and NASA documented that tin whiskers continued to develop even with a Uralane 5750 conformal coating in place. The study tracked the whisker development over a six-month period and a one-year period, with both instances indicating that growth was continuing and coating penetration was inevitable.

Lead-free tin whisker mitigation techniques*

Technique	Result/Status
Matte tin (dull finish)	Proved more resistant to whiskering than bright tin
Annealing tin	Reduces some stresses in plating that contribute to growth
Soldering or solder dipping with lead-tin solder	Proved useful for most components, but 100 percent coverage is mandatory. Only robotic solutions shown reliable to date.
Stripping finishes; replating with lead-tin solder	May be possible for some components
Conformal coatings	Can be applied with some degree of success

*No technique has been proven effective in a long-term high-stress environment.
Presented at the December 2005 Microelectronics Reliability & Qualification Workshop

Table 1

Effects on manufacturing and beyond

And what about the costs to develop and manufacture *two component product lines* – one lead-free version and another that includes a tin-lead alloy? Next, consider the costs to companies that purchase these components while providing the next higher-level subassemblies that now must maintain segregated stock and run separate processes – one RoHS compliant, the other compliant with the waivers and exemptions allowing tin-lead alloys.

Lead-free alloys tested

Alloy	Soldering method	Advantages[a]	Disadvantages[a]
Sn-3.9Ag-0.6Cu [b] (SAC)	Wave Reflow Manual	• M.pt < SnAg • Good mechanical properties • Better wetting than SnAg	• Cost and supply of Ag for wave soldering • Component thermal damage
Sn-0.7Cu-0.05Ni [c]	Wave Manual	• Supply of material • Lower cost for use in wave soldering	• M.pt > SnAg • Component thermal damage • Lower creep resistance
Sn-3.4Ag-1.0Cu-3.3Bi [d] (SACB)	Reflow Manual	• M.pt < SnAg and SnAgCu • Promising long-term reliability	• Potential for weak joints with Pb • Commercial availability

Control alloy: Sn37Pb

a Adapted from "Lead-free Electronics" (2004) for general alloy family; pros and cons may not apply for exact alloy shown
b NEMI-tested alloy
c SN100C from NIHON Superior
d NCMS-tested alloy
Source: BAE Systems 2005 testing of lead-free solder or high-reliability applications

Table 2

And how much lead are we saving with RoHS? Are the costs really justifiable? Texas Instruments has estimated that its conversion to lead-free will save the equivalent of just 10 automobile batteries annually on a worldwide scale – a minuscule amount when you consider the millions of dollars and hours spent scrambling for compliance and trying to counter the effects of using lead-free components.

RoHS seems to be here to stay, which means tin whiskers aren't going away either, until new, cost-effective soldering alloys and compounds are found. Even though tin whiskers were discovered several decades ago, no one clearly understands what makes them grow. CALCE proposes several fundamental questions aimed at discovering the root cause of whiskers:

- How much mechanical or environmental stress is needed to initiate growth?
- Does temperature cycling or steady-state temperature aging worsen growth?
- Could passivation (or surface oxide) layers reduce tin whiskers?

To date, more than 11 months after the RoHS directive was put into full effect, the industry has not yet produced an effective, standard replacement for the tin-lead alloy. However, some companies have tried to find a solution. BAE Systems researched three tested alternatives, but none proved to match the caliber of tin-lead (Table 2).

Feeling the full effect

This lead-free issue affects many other businesses besides those in the medical, communications, and defense and aerospace industries. For example, in the power-generation industry, the U.S. Nuclear Regulatory Commission has documentation dating back to 1999 from The Foxboro Company of Foxboro, Massachusetts, indicating that nuclear reactors had to perform a SCRAM (an emergency nuclear reactor shutdown) because of tin whisker growth embedded deep within the control electronics. A tin whisker was also the reported cause of a 2005 plant shutdown at the Millstone Power Station in Millstone, Connecticut.

LOCATING...

Advanced microelectronics solutions partner *with*:

- **MIL-PRF-38534/38535** facility
- Semiconductor products for specific **circuit board density, weight** and **environmental** requirements
- **Enhanced** and **ruggedized** active matrix liquid crystal displays (AMLCDs)
- **RAD tolerant & RAD hard** capabilities

TARGET LOCATED

White Electronic Designs has expertise and experience in constructing extended environment, high-reliability, high-density semiconductor products and enhanced AMLCDs that deliver on time – over time – every time.

Proceed directly to:
www.whiteedc.com/Delivers

Alternate routes include **info@whiteedc.com** or **602.437.1520.**

WHITE ELECTRONIC DESIGNS
WWW.WHITEEDC.COM

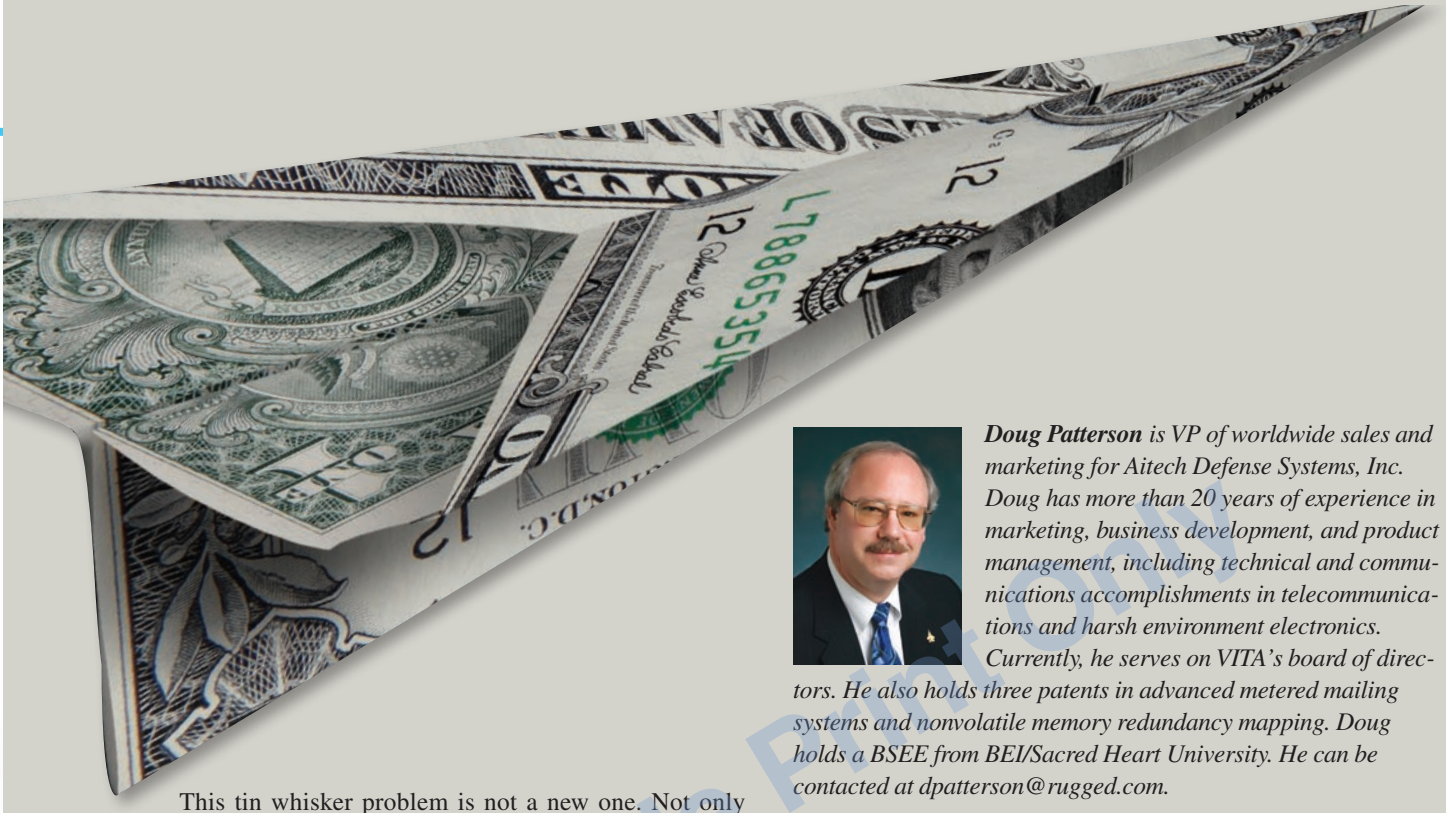

Doug Patterson *is VP of worldwide sales and marketing for Aitech Defense Systems, Inc. Doug has more than 20 years of experience in marketing, business development, and product management, including technical and communications accomplishments in telecommunications and harsh environment electronics. Currently, he serves on VITA's board of directors. He also holds three patents in advanced metered mailing systems and nonvolatile memory redundancy mapping. Doug holds a BSEE from BEI/Sacred Heart University. He can be contacted at dpatterson@rugged.com.*

Aitech Defense Systems, Inc.
888-248-3248
www.rugged.com

This tin whisker problem is not a new one. Not only has it reared its ugly head before, but it also has been exacerbated by electronics' decreasing geometries and smaller device size. Components are now physically closer together on densely packed circuit boards, and the distance a tin whisker must cover before reaching another component and causing a short, or in some instances, a catastrophic system failure, is decreasing.

The negative effects of RoHS are not limited to tin whiskers. Additional revelations continue to surface. Last January, researchers at Ford Automotive's Lincoln Design Center identified the cause of malfunctioning computers – lead-free solder that corroded when it came in contact with modeling clay that contained high levels of sulfur. The older computers with the traditional 60/40 tin-lead solder were unaffected. What other risks inadvertently created by this transition await us down the road, and what will the real costs be?

To mandate a directive with known and unknown consequences within such a short time period seems an extremely expensive and haphazard way of doing business. Change can be good, and taking risks helps businesses and industries grow. But, it is also necessary to find a balance so that the cost of doing business does not exceed the expenses needed to maintain compliance with changes of dubious benefit.

Advantages of Power Architecture technology for military applications

By Katie Butler

The PowerPC and the subsequent development of the complete Power Architecture technology family have become synonymous with military embedded computing applications, powering many high-profile programs from the M1-A2 Abrams main battle tank to the F/A-18E/F. We look at its evolution and how it has maintained a technological advantage over competitive processing architectures.

Power Architecture technology was first introduced as the PowerPC core by Motorola and IBM in the early 1990s. The PowerPC core was a new breed of superscalar RISC microprocessor that offered very high levels of performance with relatively modest clock rates. Its performance was achieved through more efficient forms of pipelining and branch prediction than previous processor architectures, resulting in improved performance per watt of power dissipation. Since then Power Architecture technology vendors have continually improved and innovated new microprocessor solutions, contributing to the architecture's success in embedded computing markets such as automotive, telecommunications, aerospace, and defense.

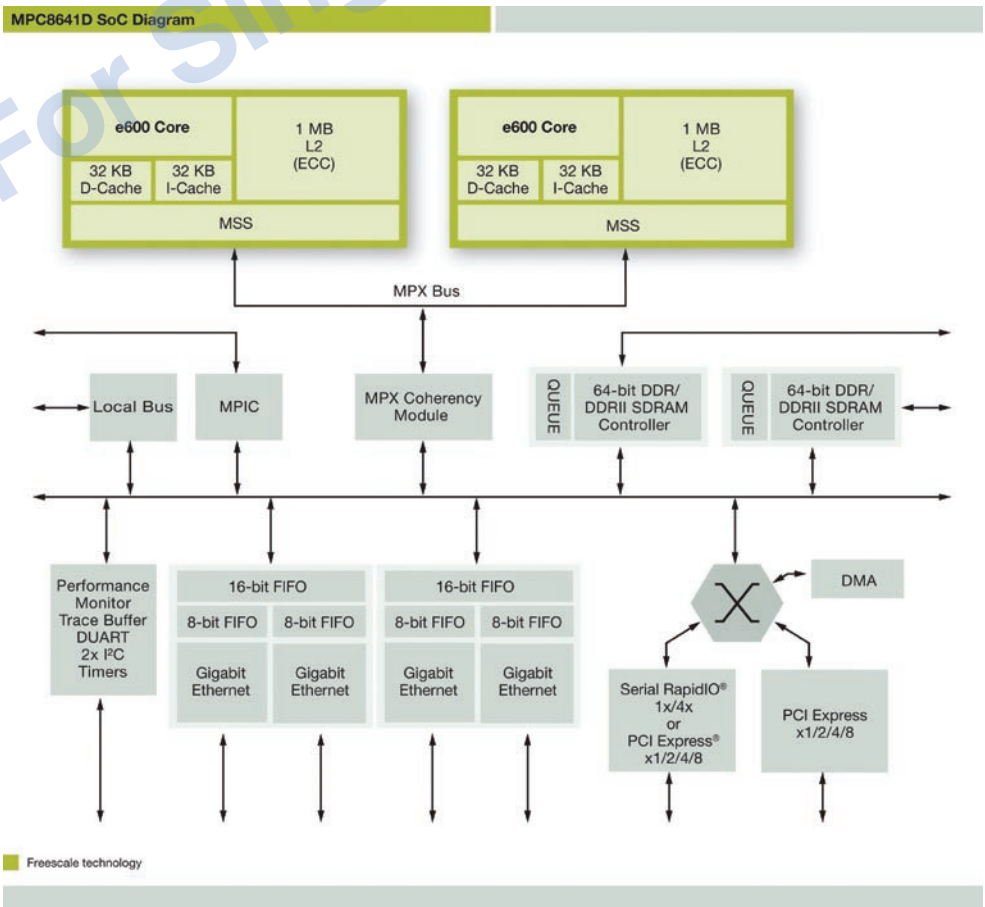

Figure 1

The stringent selection criteria required by embedded systems deployed on military front lines are reflected in the choice of processor that can be used for each application. Key selection characteristics include the following:

- High processor performance for compute-intensive and embedded applications
- High levels of integration to reduce supporting device real-estate and power dissipation
- Vector processing unit for high-performance digital signal processing applications
- Power dissipation as low as possible. Power translates into space and weight, which are premium commodities in military platforms.
- Software compatibility between generations for spiral development and long-term sustainability

Power Architecture technology addresses these criteria by having a broad product portfolio from a variety of vendors to balance the performance-versus-power trade-off. Furthermore, many of the products are targeted for long embedded application life cycles. For example, the MPC603e PowerPC device from Freescale, which was introduced in 1996, is still shipping in the industry today. By meeting the selection criteria of high-performance processors with a wide temperature range and an embedded power envelope, Power Architecture technology has become one of the *de facto* processor architectures in the aerospace and defense market today.

Power Architecture technology evolution

The first generation of PowerPC core based processors comprised general purpose execution and arithmetic units only, requiring external support for caches, memory controllers, bus interfaces, and I/O interfaces to create a complete embeddable general purpose processing subsystem. As semiconductor process technology improved and transistors became smaller, many of the functions were absorbed into the processor itself. One example was Freescale's PowerQUICC family of processors, which integrated a communications offload engine able to perform stack processing and thereby offload the CPU. This was a pointer to the future of processor architecture development, as today many functions have been integrated in hardware.

Power Architecture technology based devices such as Freescale's single and dualcore version of the MPC8641 (Figure 1) have comprehensive processing subsystems all in the same device, and yet still maintain an overall power dissipation of no more than .30 W. This compares favorably to earlier Power Architecture based devices such as the MPC7447/7448 which, when configured with external devices of equivalent functionality, would typically dissipate similar power for only one processor device.

High performance and integration

MIPS used to be the primary indicator of a processor's performance, as it is a measurement of the instruction rate of a particular device. However, the MPC8641's high degree of integration shows that overall system performance, instead of just core instruction rate, is becoming the industry's key metric.

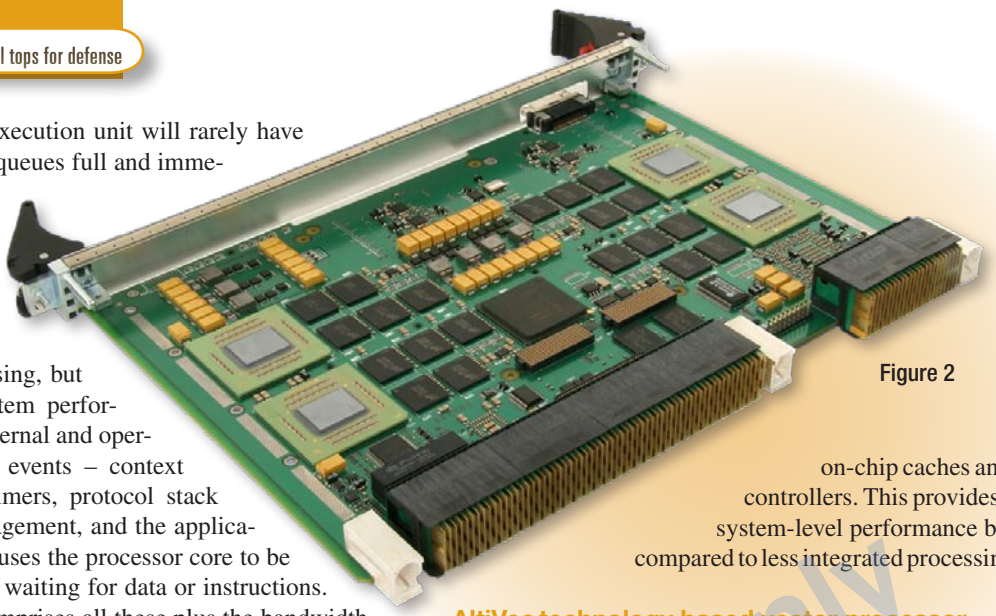

Figure 2

In today's systems an execution unit will rarely have its instruction and data queues full and immediately available for processing at the maximum instruction rate. There are exceptions to this, such as some classes of digital signal processing, but these are unusual. System performance is affected by external and operating system-generated events – context switching, interrupts, timers, protocol stack processing, buffer management, and the application itself. All of this causes the processor core to be idle for part of the time waiting for data or instructions. System performance comprises all these plus the bandwidth and management of the data paths and memory such that optimal processing performance, I/O processing, network performance, and data flows can be achieved without compromise.

Incorporating all this functionality onto a single device has been made possible by improved process technologies yielding more transistors per die. The current process allows the inclusion of hardware implementations of many I/O and communications subsystems such as multiple Ethernet controllers, Serial RapidIO interconnect, and PCI Express expansion cards in addition to on-chip caches and memory controllers. This provides a massive system-level performance boost when compared to less integrated processing devices.

AltiVec technology based vector processor

Integration has become one of the defining technology innovations of Power Architecture technology devices. One such example is the AltiVec technology, which revolutionized military digital signal processing and floating point applications, first integrated into Freescale's family of processors containing PowerPC cores in 1999. Operating in parallel with the PowerPC's execution unit, AltiVec technology is a SIMD vector processor with thirty-two 128-bit registers supporting many data types and complex instructions. This makes it highly suitable for applications such as image, sonar, and radar processing. AltiVec technology was recently enhanced to include out-of-order execution to increase its efficiency when working with multiple data streams.

Digital signal processing applications require specialized architectures that support scalability and rapid data movement from the sensors and between many processing nodes. The MPC8641D, with its dual cores and support for Serial RapidIO interconnect, makes an ideal basis for such a signal processing system. Serial RapidIO provides a high-speed, peer-to-peer switched fabric with very low latency, which can be configured in many different topologies to provide the bandwidths required between multiple processing nodes.

An example of how the Power Architecture technology's integrated features and AltiVec technology can be harnessed is the Curtiss-Wright Controls Embedded Computing CHAMP-AV 6 advanced multicomputer for digital signal processing applications, shown in Figure 2. This platform, based on the new VPX (VITA 46) standard, integrates quad MPC8641D devices, each with its own 1 GB DDR2 memory subsystem, for a total of eight processor cores and a high-speed Serial RapidIO fabric to distribute data between processor nodes.

Power dissipation

Another key issue in processor architecture today is power dissipation. Much in the way that overall performance measurements have changed from just core clock frequency, the ways of describing processor power dissipation have changed, especially with complex power management schemes. Many devices in consumer markets benefit from power management schemes where parts of the device can be powered down when not in use, or clocks can be slowed to reduce power drain.

However, many embedded applications cannot handle reduced performance to counteract an overheated processor. A primary flight control computer for a single-seat aircraft cannot operate more slowly just because it gets warm. Therefore, it is important that maximum power dissipation is always stated at maximum performance and maximum temperature. This allows designers to properly characterize their embedded systems within their operational limits.

Software compatibility

Beyond evaluating hardware platform performance, system designers must also evaluate the available software solutions. Aerospace and defense vendors make a significant investment in software. Because military platforms such as warships, aircraft, and ground vehicles have a very long service life, companies must invest in architectures that have software compatibility across platforms. Power Architecture technology vendors such as Freescale Semiconductor offer software compatibility across their Power Architecture cores and products, making it very easy to scale platforms. Preservation of the integrator's software investment is key to the success of these programs, along with a commitment to long-term supply.

The future of Power Architecture technology

The future of Power Architecture technology lies in Power.Org (www.power.org), an open forum focused on promoting the architecture and bringing together the current Power Architecture technology specifications to the open community. Both Freescale Semiconductor and IBM participate in Power.Org, and recently the organization announced the Power ISA 2.03, which merges previous versions of the Instruction Set Architecture. This merger results in more consistency and compatibility across platforms based on Power Architecture technology, a significant benefit for software and system designers. Power.Org works to ensure the future compatibility of new Power Architecture technology products and the continued growth of the supporting ecosystem of developers and suppliers.

Power Architecture technology has become one of the prominent processor architectures in the military and aerospace industry today. Its historic reputation for low power and high performance, as well as its strength in software compatibility and embedded life cycles, has resulted in highly integrated products that will continue to serve the market.

Katie Butler *is the aerospace and defense product marketer for the networking computer systems group within Freescale. She supports aerospace and defense customers that integrate PowerPC and PowerQUICC products into their solutions. Katie received her Bachelors and Masters degrees in Electrical Engineering from the Massachusetts Institute of Technology, where her graduate research work was focused on power dissipation in optical networking. She can be contacted at katie.butler@freescale.com.*

Freescale Semiconductor
512-996-5370
www.freescale.com

Product Guide: PowerPC products

Company name	Form factor	Model number	Processor/Frequency	Rugged Version(s) available	Description
Aitech Defense Systems www.rugged.com	3U CompactPCI	C903	IBM 750L, 600 MHz	✓	A rugged 3U CompactPCI single-slot SBC in conduction- and air-cooled versions
GE Fanuc Embedded Systems www.gefanucembedded.com	3U CompactPCI	IMP1A-571 Blade	PowerPC 7410	✓	A single-slot software radio blade
GE Fanuc Embedded Systems www.gefanucembedded.com	3U VPX	SBC330	PowerPC 8641D, 1.5 GHz	✓	A single board computer featuring x8 PCI Express fabric, 2 GB memory, and two GbE ports
General Micro Systems www.gms4vme.com	6U CompactPCI	C512 Freedom	AMCC PPC, 800 MHz	✓	Ultra low-cost single board computer with power consumption as low as 7 W
Interface Concept www.interfaceconcept.com	6U CompactPCI	IC-e6-cPCIb	MPC7448 PowerPC, 1.4 GHz (2 optional)	✓	A high-performance board in standard, extended, and rugged grades
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	6U CompactPCI	CPCI-6190 Processor Board	IBM 750GX, 1 GHz		A universal-mode SBC with IPMI system management support (PICMG 2.9)
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	6U CompactPCI	CPCI-6106 Universal Processor Board	IBM 750GX, 1 GHz		A CompactPCI high-performance universal processor board
Orion Technologies www.otisolutions.com	6U CompactPCI	CPC7520	IBM 750GX, 1 GHz or IBM 750FX, 600/800 MHz	✓	A PowerPC 6U single board computer
Orion Technologies www.otisolutions.com	6U CompactPCI	CPC7510	PowerPC IBM 750FX/GX, 1 GHz	✓	A hot-swappable, high-availability PowerPC 6U SBC
Performance Technologies www.pt.com	6U CompactPCI	CPC5900	AMCC High Performance PowerPC		A high-availability NAS or SAN storage blade supporting 800 GB of storage
Actis www.actis-computer.com	6U VME	ACTIS-VSBC-6848	PowerQUICC II MPC8248E, 400 MHz		A PowerQUICC II SBC based on the MPC8248 processor
CES www.ces.ch	6U VME	RIO4 Family	PowerPC G4	✓	A family of reconfigurable PowerPC-based 6U VME 2eSST computers for harsh environments
GE Fanuc Embedded Systems www.gefanucembedded.com	6U VME	PPC8A-572 Blade	PowerPC 7447 G4+, 1.2 GHz		A single-slot blade system
Aitech Defense Systems www.rugged.com	6U VME64x	Dual-Process SBC	7448 RISC	✓	A combined C102 SBC processor and M222 flash memory PMC to help solve data throughput problems and provide lower power dissipation and large mission data mass storage
Extreme Engineering www.xes-inc.com	AdvancedTCA	XCalibur1210	MPC7448 PowerPC, 1.7 GHz (2)		A versatile AdvancedTCA mezzanine card carrier that can be factory configured to support any combination of PrPMC, XMC, and AdvancedMC modules
Interphase www.interphase.com	AdvancedTCA	iNAV 74K	Freescale 8641D (4)		A blade targeted at high-volume transaction and user-plane traffic processing required for next-generation wireless infrastructure and IMS control planes
Kontron www.kontron.com	AdvancedTCA, MicroTCA	AM4100	Dual Core Freescale PowerPC MPC8641D, 1.5 GHz	✓	An AdvancedMC module
Eridon www.eridon.com	Development Tool	UnifiedLogic			A development framework comprising hardware, development tools, and software
Jungo www.jungo.com	Development Tool	Embedded USB			A development tool used to incorporate complete USB software connectivity

Company name	Form factor	Model number	Processor/Frequency	Rugged Version(s) available	Description
Macraigor www.macraigor.com	Development Tool	Eclipse+GNU Tools			An implementation that packages Eclipse 3.2.1, several of the open-source GNU tools/utilities, and a program called OCDRemote
MICETEK www.micetek.com	Development Tool	USB TAP For PowerPC			A high-speed in-circuit emulator where code is loaded by JTAG/COP/BDM interface at speeds up to 1.6 Mbps
MICETEK www.micetek.com	Development Tool	JediView for PowerPC			An Integrated Development Environment (IDE) that provides a simple, versatile graphical user interface
MICETEK www.micetek.com	Development Tool	Tools for PowerPC			PowerPC development tools, including in-circuit emulator, evaluation board, and IDE, along with EV board models such as MPC8313E, 8323E, 8347E, 8349E, and others
Embedded Planet www.embeddedplanet.com	EPIC	EP440xS	AMCC PowerPC 440EPx, 667 MHz		A credit card sized SBC (3.5" x 2.2")
Embedded Planet www.embeddedplanet.com	EPIC	EP440K	AMCC PowerPC 440GR, 667 MHz		A feature-packed computer delivering high levels of performance and reliability in a small, power-efficient package
MPL www.mpl.ch	PC/104	MIP470	MPC8347 PowerPC, 400 MHz	✓	A controller board with 400 MHz PowerPC and FPU
MEN Micro www.menmicro.com	PCI-104	EM9	MPC8548, 1.33 GHz or MPC8543, 800 MHz	✓	An embedded system module
GE Fanuc Embedded Systems www.gefanucembedded.com	PMC	ICS-1580	IBM 400 MHz PowerPC		An FPGA DSP board with high-speed inputs in PMC format
Micro Memory www.micromemory.com	PMC	MM-71xx Series	Virtex-4 FPGAs	✓	A series of compute nodes providing FPGA processing resources on a PMC form factor, targeting applications such as synthetic aperture and phased array radar, Software-Defined Radio, and signal intelligence
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	PMC	PrPMC6001 Processor PMC	MPC7448 PowerPC G4, 1.4 GHz		A PMC designed for embedded applications requiring low power and high performance
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	PMC	PrPMC2750 Processor PMC	PowerPC 750 GX, 1.0 GHz (2 optional)		A high-performance PowerPC processor PMC module
Orion Technologies www.otisolutions.com	PMC	PMC7525	PowerPC IBM 750FX/GX	✓	An extended temperature, conduction-cooled PMC7525 PowerPC PMC SBC
4DSP www.4dsp.com	PMC/XMC	AD491	PowerPC RISC		A dual-channel, 8-bit, 1 GHz PMC/XMC digitizer
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	PrPMC	PrPMC8005E Processor PMC	MPC7410 or MPC750 PowerPC	✓	A high-performance Processor PMC that operates in extended temperatures
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	PrPMC	PrPMC8150 Processor PMC	PowerPC 750FX, 800 MHz		A high-performance Processor PMC with optional encryption
Orion Technologies www.otisolutions.com	PrPMC	PMC7520	IBM 750GX, 1 GHz or IBM 750FX, 600/800 MHz		A PowerPC PrPMC single board computer
MICETEK www.micetek.com	System	MPC8313E-RDB	PowerQUICC II Pro	✓	A low-cost, high-performance system that includes a built-in security accelerator
GE Fanuc Embedded Systems www.gefanucembedded.com	VME	PowerPC DSP Blade	7410 PowerPC DSP (4)		PCI software radio PowerPC DSP blade

Company name	Form factor	Model number	Processor/Frequency	Rugged Version(s) available	Description
GE Fanuc Embedded Systems www.gefanucembedded.com	VME	ICS-8550	G4 PowerPC (4)		A DSP board with two PMC sites integrated with GE Fanuc radio PMC modules to offer a complete receive and/or transmit solution in only one VMEbus slot
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	VME	MVME6100	MPC7457 PowerPC, 1.267 GHz		Motorola's original SBC designed with the 2eSST VMEbus protocol
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	VME	MVME5500	MPC7457, 1 GHz		Motorola's latest VME single board computer
Motorola Inc., Embedded Communications Computing www.motorola.com/computing	VME	MVME3100	MPC8540, 800 MHz		Motorola's second VMEbus SBC designed with the 2eSST VMEbus protocol
CSPI www.cspi.com	VXS	3000 SERIES	Freescale 8641D		Hybrid computing platforms that integrate the Freescale 8641D multicore General Purpose Processor (GPP), Xilinx Virtex-5 FPGA, and Myricom's 10 GbE Myri-10G clustering technology on a VXS platform
VMETRO www.vmetro.com	VXS	Phoenix VPF1	PowerPC 7447A (2)	✓	A digital signal processor supporting VITA 41 backplane switch fabric communications via VXS over P0

Data was extracted from OSP's product database on July 31, 2007. Search keyword was "PowerPC" on products entered January 1, 2007 through June 30, 2007 within all OSP magazines. Products were also selected according to relevance to the product guide's theme. Entries have been edited for publication, and OpenSystems Publishing is not responsible for errors or omissions. Vendors are encouraged to add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np.

Acromag I/O. Customize It.

Acromag 50 YEARS
Excellence in I/O since 1957

Acromag introduces affordable FPGA I/O. For ALL your projects.

As an engineer, your projects are unique, ever-changing, and budget-bound. That's why our new PMC modules give you an affordable solution to create custom I/O boards.

But if you thought FPGA computing was only for top-end applications, think again. Our PMCs are ideal for protocol conversion, simulation, in-circuit-testing, and much more.

So, why settle for generic I/O when you can design exactly what you need while staying in budget and reducing your time to market?

- Virtex®-II, Virtex-4, Acex®, and Cyclone®-II FPGAs
- Large DRAM buffers and dual-ported SRAM
- Conduction-cooled models available

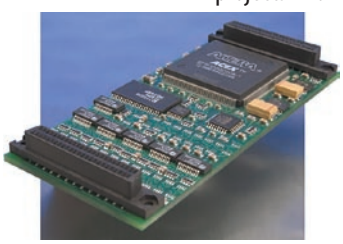

Industry Pack FPGA I/O also available

Cost-effective custom I/O
Choose from a variety of I/O configurations:

- Digital I/O: TTL, CMOS, RS422, or LVDS I/O
- Analog I/O: 16-bit 100MHz A/D, 900KHz D/A

Faster time to market
Why waste precious time building a board from scratch? Our new FPGA modules let you process your I/O signals any way you want. Quickly.

Flexibility to meet unexpected challenges
Acromag FPGA I/O will help you bring your projects in on time and under budget. And with FPGAs, you'll be ready to adapt to all the inevitable changes.

Thinking about FPGA I/O?
Think flexible. Think affordable.
Think Acromag.

Call or visit our website today – for VME, CompactPCI, PMC, and Industry Pack solutions.

Acromag®
THE LEADER IN INDUSTRIAL I/O

www.acromagembedded.com
877-295-7088 or 248-295-0310

ANALOG I/O | DIGITAL I/O | SERIAL I/O | FPGA | COUNTER/TIMER | QUADRATURE

Manufactured in Wixom, Michigan, USA

Unless you're in the NSA, read this …

NSA-encrypted red/black systems are without a doubt the most secure systems on the planet – they have to be. Trouble is, all the rest of the DoD's equipment might literally include computers and peripherals purchased off the GSA from Best Buy, Wal-Mart, Radio Shack, and other civilian COTS suppliers. That's the benefit of COTS, folks. But there's still the problem of lost laptops, unencrypted files, malware, forgotten passwords and cookies, and all the other crud that affects our personal computers every day. So Lavasoft, a recognizable name in the anti-spyware business, has created a trio of digital protection tools that apply to DoD users just as well as consumers.

Sure, you can do some of these same functions with open source tools like Eraser or TrueCrypt, but Lavasoft's are commercial-grade COTS packages that are designed for utility and ease of use. The Digital Lock – File Encryptor product includes multiple encryption algorithms such as AES standard 256 bits, multilayered encryption, shredding after encryption, and an e-mail security channel. File Shredder's drag-and-drop features work with Word, Excel, PowerPoint, videos, and other multimedia files, and has some tricks to work with Windows Vista. And the Privacy Toolbox combines both products into one. A 30-day trial is available on Lavasoft's website.

Lavasoft • www.lavasoft.com
RSC# 34057

Serial RapidIO AdvancedMC for radar, sonar, telecom

The AdvancedMC card – used on AdvancedTCA and MicroTCA platforms – is turning into quite the mezzanine for communications connectivity. Originally designed for telecommunications applications, AdvancedMCs are appealing to the military primarily because of their COTS status and myriad networking options. If rugged MicroTCA gains traction in defense, AdvancedMCs such as Mercury Computer's Ensemble MPC-102 might be coming to a TOC or AWACS near you.

Based on the Power Architecture dual-core Freescale 8641D, the MPC-102 offers Serial RapidIO or PCI Express connections to the carrier card, along with four GbE connections (two to the panel and two to the AdvancedMC connector), a SATA port, and dual RS-232 connections. With up to 10 Gbps raw I/O bandwidth on Serial RapidIO and the on-the-fly processing afforded by the 1.3 GHz PowerPC and 2 GB DDR2 SDRAM combo, the module is ideal for networking applications that process radar or sonar data.

Mercury Computer Systems • www.mc.com
RSC# 34056

Brightness beyond "nits"

How many times have you tried to use that fast Wi-Fi pipe out on your porch, patio, or deck – only to find that the sunlight washed out your laptop's display so badly that it was illegible? Me, too. Now imagine how this affects soldiers and Marines in the Middle East – whether they're dismounted, under sun shades, or in HWMMVs. That's why General Dynamics Itronix invented the DynaVue all-light viewable touchscreen display. A manufacturer of rugged laptops, notebooks, and Tablet PCs, Itronix deploys its portables into all kinds of battlefield conditions – including bright sunlight.

Focusing on contrast ratio and polarization techniques, the patent-pending light filtering technology increases the display's contrast ratio to provide rich color saturation and improved fine detail visibility. With a single cold cathode fluorescent backlight lamp, the computer saves battery power because the display requires less energy to maintain sunlight readability. General Dynamics Itronix says DynaVue will debut on the company's next-gen computers.

General Dynamics Itronix • www.gd-itronix.com
RSC# 34053

32 instruments on one PC

Think of it as a data acquisition Swiss Army knife – one of those ridiculously large ones that even has a fishing pole built in, tucked next to the shortwave transceiver in the handle. Well, you get the idea. The Version 6.5 MultiScope software from DSPCon runs on a PC and provides display power for 32 test and lab instruments on a single PC. "Feature rich" is an understatement for this product. The package replaces oscilloscopes, spectrum analyzers, data loggers, recorders, pianos (kidding about that last one), and several other instruments.

The GUI desktop works with DSPCon's Piranha III and DataFlex-1000 products for acquiring, analyzing, and displaying data from up to 32 test channels. MultiScope can display data in Time Series, FFT, Strip Chart, Lissajous, Campbell, Spectrum, and myriad other modes – with up to 32 scopes per user-defined page. And, users can create multiple pages just like "sheets" inside of Microsoft's Excel environment. The convenience of having one multichannel tool that does it all – on a single PC – with the ability to display, analyze, process, record, and playback is extremely compelling for design and test engineers. Amazing.

DSPCon • www.dspcon.com
RSC# 33488

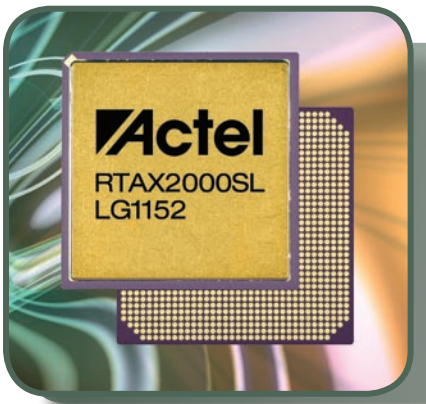

FPGAs for Mars

Picture this: Ground control, Major Tom, little green men. What image comes to mind? That's right: Mars, space, the Final Frontier. A truly hostile environment, no matter how glamorous Hollywood makes it out to be. Surviving in space means constant radiation bombardment, temperature extremes, and imperative energy vigilance as batteries charge in sunlight, then go flat in the inky darkness.

Actel's RTAX-SL FPGAs were designed with just this environment in mind. The new "L" variation of the company's RTAX line includes the RTAX250SL (30,000 ASIC gates), RTAX1000SL (125,000 ASIC gates), and RTAX2000SL (250,000 ASIC gates) and promises 50 percent lower standby current "than the leading space FPGA at 125 °C." (We're assuming they mean a product from "A" or "X.") The devices use SEU-hardened registers, eliminate the need for TMR, and promise an SEU rate of less than 10E-10 errors/bit-day in worst-case geosynchronous orbit. The devices are also all available in QML Class V per MIL-PRF-38535.

Actel Corporation • www.actel.com
RSC# 34051

Itty bitty 15 A fuse

In electric shop during my childhood, I built my first project power supply out of an old MIL-SPEC power supply. I remember vividly the mombo fuses that thing had because every time I was careless with the leads: Pop! I fried one of them. Those suckers musta been 0.5" diameter and more than 1.0" long. My, how times have changed. Today you can surface mount a 125 V/15 A subminiature fuse that's about the size of Lincoln's head (see photo).

The 465 Series Littelfuse devices are RoHS compliant and offer 20 A and 30 A ratings in 0.397" (10.10 mm) x 0.123" (3.12 mm). Even more impressive are their interrupting ratings of 100 A at 125 Vac and 300 A at 65 Vac, respectively. Designed for overcurrent protection in space-constrained applications, they're ideal for vehicle-mounted military comms and telemetry equipment that relies on "shore power" but must fit into suitcase-size cubbyholes.

Littelfuse • www.littelfuse.com
RSC# 34052

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

A Menu for All Tastes

AMPRO Embedded Solutions

Hungry for Grade A hardware with a side of embedded software support? Ampro serves up a wide range of embedded system boards that support all major software platforms. Whether you are running Linux® or Vx Works® on a system platform, or Windows XP® or QNX® on a Mini-ITX board for low cost applications, we're ready to take your order. Virtually every combination of software platform and hardware form factor is available through Ampro. For over 2 decades Ampro has been delivering industry specific modular embedded computing solutions that set the bar for quality, performance and value. Ampro's solutions include quality SBC's pre-configured for almost every purpose to a family of turnkey, industrial computer systems that meet different application requirements. All systems are delivered running the software platform that you request. With Ampro you'll get no bad aftertaste. Just the right products at the right price.

Contact us today and find out how Ampro can help satisfy your appetite for embedded systems.

Call today: (408) 360-0200 or toll free: (800) 966-5200 or visit: www.ampro.com

AMPRO®

Dual-Core Xeon security server

Besides courage and guts, the military relies on communications equipment – from extremely rugged to civilian COTS gear. The 1U rack-mount PL-01039 from WIN Enterprises is a security server/switch with eight GbE ports. The 2.0 GHz Dual-Core Xeon LV/ULV processors run two Ethernet modules, one powered by the Intel 82571EB (four ports) and the other by the Intel 82573L (four ports). There's also a 32-bit PCI connector, Mini PCI socket, a USB 2.0 port, and an RS-232 port.

But the heart of this system is the security element. The PL-01039 can run multiple security apps, including IDS/IPS, firewall, VPN gateway, Unified Threat Management (UTM), anti-spam, anti-virus, and several other security profiles. Available with copper, fiber, or mixed media, the unit also has its own onboard battery and can run three hours without external power.

WIN Enterprises • www.win-ent.com
RSC# 34054

More mobile graphics realism

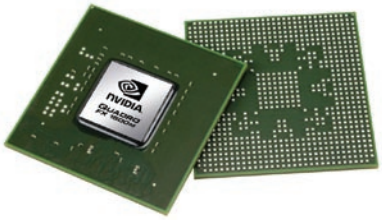

Though designed for notebook computers, NVIDIA's Quadro professional GPU architecture is ideal for power-constrained military boards and systems that create intense graphics workloads. Borrowing from the company's recent Quadro FX 5600 and FX 4600 workstation GPUs, three new mobile chips carry over the Shader Model 4.0 and uniform feature set. All the devices are aimed at engineers, scientists, artists, and other users of large datasets, models, and high-res images.

The unified architecture is capable of dynamically allocating compute, geometry, shading, and pixel processing power. Shader Model 4.0 adds next-gen vertex and pixel programmability with ultra-realistic OpenGL and DirectX 10 effects. The NVIDIA CUDA technology aids in solving complex visualization problems, and the PowerMizer 7.0 technology is the company's seventh-generation hardware power management feature set that brings lower power consumption and longer battery life to portable systems. It all adds up to kick-butt realistic images and displays in defense applications.

NVIDIA Corporation • www.nvidia.com
RSC# 34055

OpenSystems Publishing

Advertising/Business Office

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

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

Business Manager
Karen Layman

Sales Group

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

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

Doug Cordier
Account Manager
dcordier@opensystems-publishing.com

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

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

Regional Sales Managers

Jerry Bleich
New England
jbleich@opensystems-publishing.com

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

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

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

International Sales

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

Reprints and PDFs

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

Advertiser Information

The "A's" have it

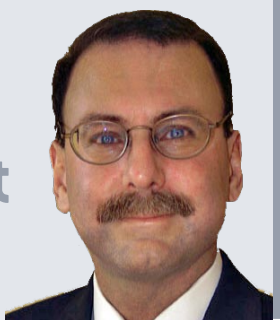

By Chris A. Ciufo

Last month, a coincidence of media events and new product releases aligned (or maybe "conspired") on my desk in such a way that they all came from companies starting with the letter "A": Altera, Aonix, AMD. I didn't plan it this way, folks. But I might as well take advantage of this woo-woo *Twilight Zone* moment. Here's what's new, and why you should care.

Altera steps up mil tempo

In the world of big FPGAs, there's Altera and there's Xilinx. Using public information, I estimate that about 25 percent of Xilinx's $1.8 billion revenue is due to the military market, while only about 10 percent of Altera's $1.29 billion revenue comes from defense. But Altera is looking to change all that in a big way by focusing on SWaP, SDR, VPX, mil temp, and AQEC.

In a recent interview with company bigwigs, I got the G2 on this alphabet soup. Size, Weight, and Power (SWaP) are important in many applications these days, but even more so in Software-Defined Radios (SDRs) that are often either deployed as handhelds, or rely on limited power sources in avionics bays or wheeled vehicles. The company's new Programmable Power Technology in Stratix III devices only activates essential transistors, leaving unused gates in low power standby. Altera says this saves up to 90 percent of active power in SDR waveforms such as Soldier Radio Waveform (SRW). Also, the company surprisingly has VITA's new 3U VPX form factor on its charts as a key platform for military SWaP applications.

Altera also maintains a secure ITAR design facility, and was one of the *very first participants* in the DoD's AQEC IC design flow. Now sponsored by GEIA, AQEC (GEIA-STD-0002-1A) is a voluntary IC spec, and companies promise to provide design data to aid mil contractors in making uprating and life-cycle decisions. Finally, and perhaps most importantly, the company designs all of its devices for industrial temp [-40 °C to +100 °C (Tj)], and some Stratix and HardCopy devices extend to the MIL-STD-883 range of -55 °C to +125 °C (Tj). Altera seems very keen on growing military beyond that 10 percent in FY07 (ending December 2007).

Aonix offers Ada for safety critical

You're as surprised as I am that Ada is still out there kicking around. Or, maybe you're not surprised if you're working on legacy defense systems. Ada's strong typing and rigid syntax make it ideal for military applications where code ambiguity – things like priority inversions or oddly defined variables – can't be tolerated lest lives be lost.

During September's Embedded Systems Conference, Aonix announced a version of their ObjectAda RAVEN that works with Wind River's VxWorks ARINC 653 for PowerPC processors. Although the previous sentence is a mouthful, what this means is important for several reasons. First is that users now have a way to run legacy Ada code on top of VxWorks, marrying the old with the new. Secondly, by using the ARINC-653 API, VxWorks 653 creates partitioned environments above the kernel into which multiple distinct and segregated environments can run.

This partitioned model is the core of ARINC-653 safety-critical systems. The idea is that a failure in one module, such as an Ada DO-178B application, would have no effect on the other partitions and hence the overall system's integrity is maintained. Additionally, the Aonix ObjectAda RAVEN has an ACATS 2.5 Ada 95 compiler and tools that help create or recompile Ada code in a partition; it also relies on the APEX communications API for easy communication with the VxWorks 653 executive API. According to Aonix, this Ada + VxWorks + ARINC-653 combo is unique in the industry. Well, except for something from AdaCore that doesn't have the APEX API. (Notice the preponderance of "A's" again. Coincidence? I think not.)

AMD adds multicore "triple threat"

Just as we went to press, AMD one-upped Intel by introducing the triple-core desktop Phenom x86 processor. Expected to ship in Q1 2008, the device sports three cores on the same die – a feat not to be confused with multicore CPUs having multiple dies in one package. Although it's rumored that the triple-core Phenom might be a quad-core Phenom with one core "disabled," this doesn't detract from AMD's achievement. (In fact, the original 8031 MCU was just a mask-programmed 8051 with the ROM disabled, and today's flash memories contain mask-enabled memory arrays.)

AMD's Balanced Smart Cache speeds access to memory and the shared L3 cache aids in multithreaded applications such as HD video, games, and as-yet-unannounced software applications. Bill Mitchell, corporate VP of the Windows Ecosystem at Microsoft, said "Microsoft is excited to see AMD creating new technologies." Well, sure. The device sports HyperTransport 3.0 with up to 16 GBps I/O transfer bandwidth, and each core can run at a separate frequency via Cool 'n' Quiet 2.0 technology. SYSmark 2007 and 3DMark06 benchmarks are due out soon, as is the quad-core version of the Phenom. We'll keep an eye on this multicore race between titans.

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

GE Fanuc
Embedded Systems

We walk the talk.

Real world, deliverable VPX from GE Fanuc Embedded Systems.

While others were talking about their commitment to VPX, we were walking the talk. And as a result, we can now offer you an impressive selection of VPX products. Not words, products. Like our Intel- and PowerPC-based single board computers, graphics processors, switches and full VPX systems. With more in the pipeline.

Take, for example, the MAGIC1 Rugged Display Computer. Because it's VPX-based, it is able to connect a 2.0 GHz Intel® Core™ Duo to an NVIDIA® G73 with 16-lane PCI Express®. The result is an integrated graphics subsystem that leads the market in performance, flexibility and ruggedness. And because it is so small and lightweight, it's perfect for embedded training, simulation and mission rehearsal.

To us, VPX isn't just cool technology. It's technology that solves problems for our customers. For more information about VPX and our real world products, give us a call or visit our site. No one knows more about how VPX can solve problems for you.

www.gefanucembedded.com/vpx

MAGIC1
3U VPX system with leading edge graphics performance

VPX

FANUC

RADSTONE
EMBEDDED COMPUTING

Now a part of GE Fanuc Embedded Systems

TOUGH HARDWARE EASY SOFTWARE.

Software shouldn't be hard. That's why we back our rugged boards and subsystems with all the tools you need to easily and rapidly design, integrate and support your deployed electronics systems. You can depend on our board support packages, built-in test routines, diagnostic utilities, extensive software libraries and broad operating system support to reduce design costs and get your product to market faster. That's hard evidence that makes Curtiss-Wright the easy choice.

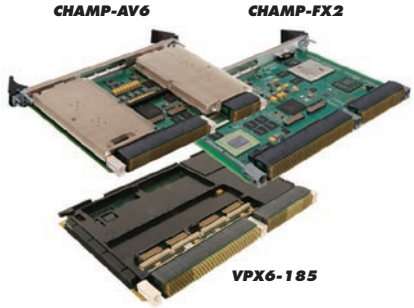

The CHAMP-AV6, CHAMP-FX2 and VPX6-185 are three of our latest generation of VPX/VPX-REDI boards and include a rich set of COTS Continuum system and support software designed to ease integration and reduce system development time.

CURTISS WRIGHT Controls
Embedded Computing

www.cwcembedded.com

EMPOWERED SOFTWARE... *ABOVE & BEYOND*

MIL/COTS DIGEST

The Defense Electronic Product Source

Sept/Oct 2007

In This Issue

All the products you *need* for Aerospace and Defense

You're holding in your hands the second installment of *MIL/COTS DIGEST (MCD)* now part of *Military Embedded Systems* magazine. My vision for MCD is to dramatically expand our coverage of new COTS products for the Aerospace and Defense (A&D) markets. Why's that? Because we get about 100 new product announcements *per day!* How else can we tell you about all this great new stuff? MCD gives you a quickie "thumbnail" view of some of my favorite ones – in an easy-to-scan print format.

Even in an era of way-cool Web pages, old-fashioned paper is still handy for perusing new product listings, and easier to read while leaning back in a chair or on an airplane. In this eight-page supplement, you'll find 28 products, ranging from boards and systems to blades, vision systems, cameras, rugged box-level computers, LCDs, components, and ... well, you get the idea. A *wide variety* of new products awaits you.

Go ahead and start thumbing through. Something's bound to catch your interest.

(By the way ... one of the most intriguing products is the "hygienic design" industrial PC from noax Technologies. Think you can just hose it off?)

Chris Ciufo, *Editor*
cciufo@opensystems-publishing.com

PS: In 2008, *MCD* will increase to quarterly as *Military Embedded Systems* increases to eight issues. So, look for even more products in the future.

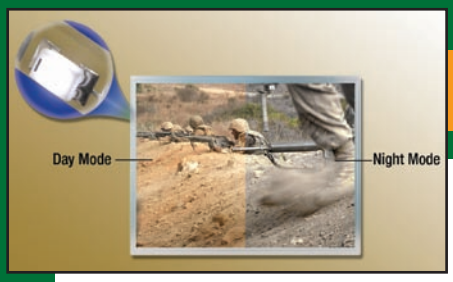

Day-to-night mode LCDs

The NVIS A and NVIS B are VGA and XGA resolution TFT LCDs backlit by LED rails that can be switched back and forth between day and night mode via the company's standard PRISMA II industrial controller board. Compatible with Night Vision Imaging Systems (NVIS) technology, NVIS A is for upper echelon applications, while NVIS B is for applications with less stringent requirements. The units feature updated technology that eliminates the cumbersome and expensive filters – large glass overlays bonded to the outside of the display – previously used. NVIS LCDs can be used in military aircraft cockpits, tanks, trucks, ground mobile applications, and communications systems. The LCDs are 8.4" diagonal, scalable up to 15" diagonal, and fully RoHS compliant.

www.apollodisplays.com **APOLLO DISPLAY TECHNOLOGIES, LLC**

3U VPX SBC and carrier board

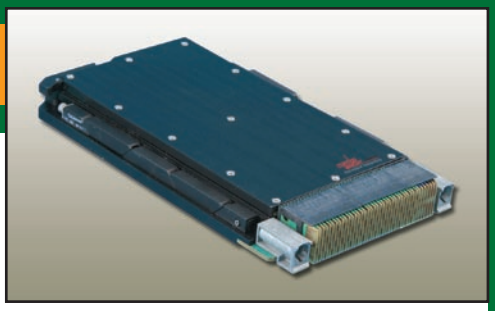

The VPX3-125 is a 3U VPX SBC featuring a single or dual core P.A. Semi PWRficient PA6T-1682M processor at 1.5 GHz. Highly suited to harsh environment A&D applications, the VPX3-125 provides 512 MB/1 GB DDR2 memory @ 400 MHz, 128 MB NOR flash, 1 GB NAND flash, and 512 KB NVRAM. I/O includes one XMC/PMC site, two 10/100/1000 Ethernet ports, RS-232 and RS-422 serial channels, and two x4 lane PCI Express egress ports off-board. Its optional 3U VPX I/O expansion companion carrier board, the ExpressReach, provides an XMC/PMC site supporting x8 lane PCI Express port (VITA 42.4) with differential pair I/O routing, along with two x4 PCI Express switch ports to simplify expansion.

www.cwcembedded.com **CURTISS-WRIGHT**

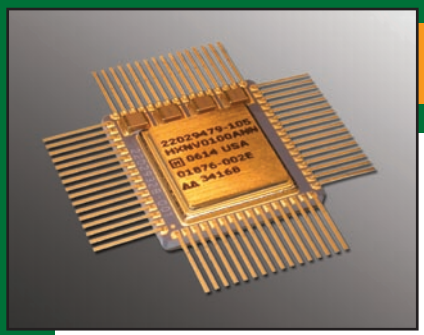

One-million bit nonvolatile memory

The HXNV0100 is a one-million bit nonvolatile static memory component for strategic space electronics applications. The memory array and control electronics are both radiation hardened, providing high reliability for low-voltage systems operating in radiation environments. The magnetic RAM runs from a 3.3 V power supply and offers >1e15 read/write cycles. Fabricated with Honeywell's Silicon-On-Insulator (SOI) Complementary Metal Oxide Semiconductor (CMOS) technology and combined with magnetic thin films, the new memory component can either replace plated wire memory or be integrated into new systems designs.

www.honeywell.com **HONEYWELL INTERNATIONAL**

All the Right Stuff for Avionics and UAV Applications:

Designing and building board-level products and integrating sub-systems for today's advanced avionics and UAV applications is tough enough. Doing it in a true COTS environment is even tougher. But Aitech is more than equal to the challenge. Aitech delivers more than two decades of harsh environment, open systems architecture expertise and proven solutions in UAV and avionics applications such as Predator, Global Hawk, C-130, F-18, F-16, UCAV/J-UCAS, and many more...

And Aitech continues to provide advanced VMEbus and CompactPCI products designed, built and tested to -55°C to +85°C as standard. Because when it comes to manned or unmanned air vehicles, you can't bypass the laws of physics! And when it comes to light weight, low power, extreme reliability, ease of maintenance, minimized development costs, ease of technology insertion and upgrades, Aitech delivers every time.

We can't change the physics... but we can ensure your COTS sub-systems are designed, built, and tested to perform reliably at the temperature extremes of your specification – without custom development, "work-arounds", or compromises.

We take the extra steps... including pre-screened parts qualification, HALT, and 100% HASS/ESS testing to ensure that every standard Aitech product meets all your temperature and rugged performance specifications...standard.

We've been doing it... meeting full temperature-range specifications with standard products is just part of our 20+ year heritage and commitment to COTS advancements – from the first conduction-cooled Mil-Spec VME board in 1984, to today's highest functionality MIPS/Watt boards, multi-Gigabyte mass Flash mass memory cards, and high-speed mezzanines...for manned and unmanned applications.

We have the proof... visit our web site or call for more information and our catalog of proven solutions.

COTS Compliant Solutions from Aitech.

Aitech

Aitech Defense Systems, Inc.
9301 Oakdale Avenue
Chatsworth, CA 91311
email: sales@rugged.com
Toll Free: 888-Aitech8 - (888) 248-3248
Fax: (818) 718-9787
www.rugged.com

3U CompactPCI codec board

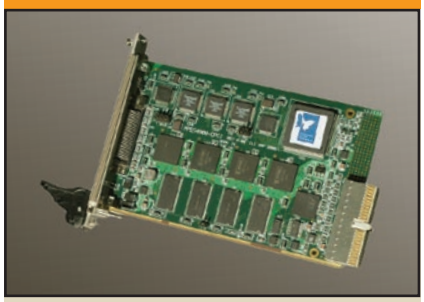

The MPEG4CPCI is a single 3U form factor CompactPCI MPEG4 codec board that encodes up to four concurrent full-size real-time analog inputs at a full frame rate of 25/30 fps from PAL or NTSC video and audio sources. The board decodes and plays back video and audio recordings from storage to display, and a preview feature allows incoming video to be viewed on the host screen in parallel with the recording process. The MPEG4CPCI features a 32-bit PCI architecture and supports Windows NT/2000/XP, Linux, and QNX. Text and graphics annotation can be superimposed on any channel, and audio/video synchronization is provided on each channel. The device operates from 0 °C to +60 °C, and an extended temperature version (-40 °C to +85 °C) is available.

www.ampltd.com

ADVANCED MICRO PERIPHERALS

XScale touch panel computer

The TPC-120H is an Intel XScale PXA 12.1" SVGA TFT touch panel computer featuring two 10/100BASE-T Ethernet ports. The computer's resolution is 800 x 600, and it comes complete with a Windows CE operating system to support thin-client solutions. I/O includes two RS-232 and one RS-232/485 serial port, two USB 1.1 ports, and one VGA port. The fanless TPC-120H provides one CompactFlash expansion slot and 64 MB SDRAM in a super slim design.

www.eAutomationPro.com

ADVANTECH eAUTOMATION GROUP

Dual-channel PMC/XMC digitizer

The AD491 – an 8-bit, 1 GHz PMC/XMC digitizer featuring two Virtex-4 FPGAs – includes two ADC channels that provide a 30 MHz to 1,000 MHz (1 GHz) sampling range. Useful for connecting to a remote server or storage system, the AD491 provides 8-bit data resolution and onboard clock generation in steps of 1 MHz. Memory comprises 2x 32 M x 16 DDR2 SDRAM (128 MB), 4x 2 M x 32 QDR2 SRAM devices (32 MB), and a 128 Mb flash device. Several interfaces – PCI; PCI-X 64-bit 133 MHz, 3.3 V; PCI-X/PCI 64/32-bit 66 MHz, 3.3 V; PCI 64/32-bit 33 MHz, 3.3 V; and four-lane PCI Express – are available. Includes 4x 2.5 Gbps optical transceivers for SFPDP, Fibre Channel, GbE, and InfiniBand applications.

www.4dsp.com

4DSP

Conduction-cooled FPGA PMC modules

The PMC-CX Series comprises user-configurable Virtex-II FPGA PMC modules with conduction cooling and differential digital I/O. Accordingly, I/O includes 16 bidirectional CMOS I/O lines, 24 bidirectional RS-422/485 differential I/O lines, and rear I/O connection. The series features customizable FPGAs with 11,500 or 24,192 logic cells (Xilinx Virtex-II XC2V1000 or XC2V2000), and FPGA code loads from PCI bus or flash memory. PMC-CX modules feature 256 K x 36-bit dual ported SRAM memory, along with support for dual DMA channel data transfer to the CPU and both 5 V and 3.3 V signaling. The series operates at extended temperatures of -40 °C to +85 °C.

www.acromag.com

ACROMAG

High-availability MicroTCA server

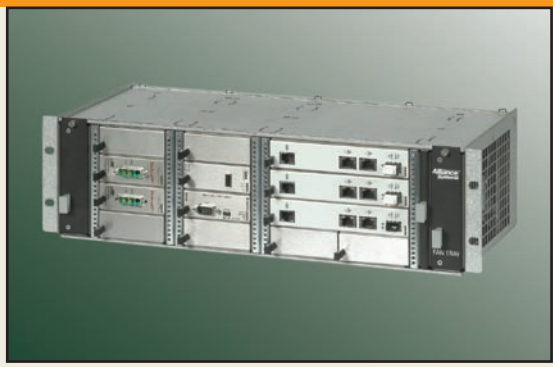

The U-3000 – a NEBS Level 3/ETSI compliant, high-availability MicroTCA communications server – provides fully redundant power, cooling, and front-accessible MCHs. The 3U rack-mount MicroTCA system accommodates two-post 19" to 24" racks and features dual redundant, hot-swap 400 W AC or -48 Vdc power supplies with power and cooling for up to 80 W per slot. The server features a dual-star topology with a fully passive MTCA.0 R1.0 backplane for 10 Gbps throughput. The U-3000's horizontal, flexible design enables installation of double-wide and/or single-wide AdvancedMCs: 4 double-wide and 2 single-wide AdvancedMCs or 10 single-wide AdvancedMCs.

www.alliancesystems.com

ALLIANCE SYSTEMS

Quad-core Xeon appliance

The FW-8890 is a 2U, two-way, quad-core Xeon rack-mount appliance for enterprise-grade network security and management applications. Its Intel core micro-architecture increases performance headroom for single- and multi-threaded applications. An onboard Cavium Nitrox CN1010 VPN Accelerator increases total system throughput by processing high-level IPsec and IKE, IPv6, SSL, and wireless LAN security protocol macro commands, reducing host I/O traffic and offloading the system processor. The FW-8890 provides up to 12 enterprise-grade Intel Gigabit NICs, and Intel Virtualization technology enables migration of more environments including 64-bit applications and operating systems to virtual environments.

www.lannerinc.com

LANNER ELECTRONICS

Software radio FPGA blade

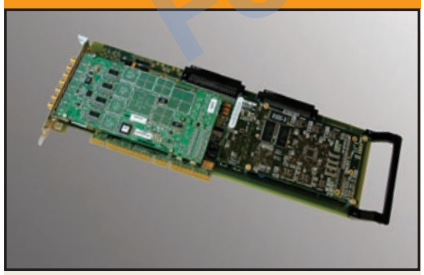

The Xilinx FPGA Blade is a PCI software radio FPGA blade consisting of GE Fanuc radio PMC modules and an FPGA PMC module mounted on a PCI carrier card. The circuit card assembly occupies one PCI bus slot, and the Xilinx FPGA module includes a Virtex-II FPGA with up to 8 million gates. The module also includes 128 MB (64-bit wide) SDRAM, six QDR SRAM banks, and a PCI 2.2 compliant high-performance bus interface. The device provides direct data transfer to processor via P4 user I/O and has a 64-bit, 66 MHz PCI host interface on the PMC modules. The Xilinx FPGA Blade supports Windows, Linux, and Solaris. Baseband processing custom coding in the Xilinx FPGA is also supported.

www.gefanucembedded.com

GE FANUC EMBEDDED SYSTEMS

Multi-threat security appliances

The FortiGate-3810A and FortiGate-3016B are multi-threat security appliances for large enterprise and MSSP applications. The appliances are highly scalable and deliver up to 26 Gbps of firewall performance. Ideally suited for network core and data center deployments, FortiGate-3810A and the FortiGate-3016B combine Fortinet's existing FortiASIC-CP6 content processor with its new FortiASIC-NP2 network processor, which enables firewall throughput of up to 26 Gbps from a single device. Additionally, hardware-accelerated GbE interfaces are available for all FortiGate-3000 Series platforms to enable wire speed firewall and near wire speed VPN performance; this ensures that time-sensitive applications like VoIP or IPTV do not suffer from network latency or jitter. The FortiGate-3810A and FortiGate-3016B also accommodate optional AdvancedMC expansion, providing hardware-accelerated 10-gigabit XFP interfaces, gigabit SFP interfaces, and hard drive storage options.

www.fortinet.com

FORTINET

DSP motion control kit

The MCK2812 is a DSP motion control kit that contains three parts: MCK2812 DSP board, PM50 power module, and IB-2812 interface board. The MCK2812 can be used for a variety of applications, including digital motor control (DC brush/brushless servo motor, stepping servo motor, AC servo motor), as a Variable Frequency Control (VFC) packing mechanism, as a digital control machine tool, and for electrical control. The DSP board features a DSP controller (TMS320F2812) operating at 150 MHz with a single DC power supply of +5 V or +24 V, along with an RS-232 serial port and opto-isolated CAN communication interface. Memory includes 128 K word on-chip flash program memory, 18 K word on-chip data/program RAM memory, and 128 K word on-board data/program RAM memory.

www.ezembedded.com

EZ-EMBEDDED

Intel-based processor blade

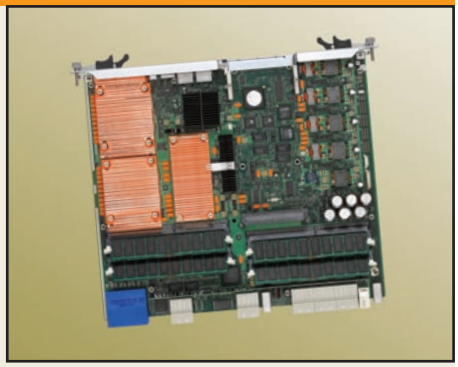

The ATCA-7140 blade is an Intel-based processor blade featuring two Dual-Core Intel Xeon (2.13 GHz) LV processors. The blade suits any 32- or 64-bit application requiring high-performance processing and provides SMP support. The blade also supports the PICMG 3.0 GbE base interface and PICMG 3.1, Option 1 and 2 fabric interfaces. The ATCA-7140 has an AdvancedMC site for I/O, coprocessing, or SAS hard disk drive flexibility, and is designed for NEBS and ETSI compliance. The ATCA-7140, which also provides multiple software packages and an operating system, is RoHS (6 of 6) compliant.

www.motorola.com/computing

MOTOROLA INC.

Xilinx DSP platform family

The Spartan-3A DSP is a digital signal processing platform family featuring the Xilinx XtremeDSP slice, which can be interconnected in many different ways on-chip and provides an 18-bit x 18-bit multiplier, 18-bit pre-adder, 48-bit post-adder/accumulator, and cascade capabilities for various DSP applications. The Spartan-3A DSP provides up to 2,200 Gbps memory bandwidth, and the chip's DSP48A slices can implement wide math functions, DSP filters, and complex arithmetic, all at low power consumption. The device provides up to 53,712 logic cells, 2,268 Kb block RAM, and 373 Kb distributed RAM. Xilinx development tools including System Generator for DSP and AccelDSP synthesis have been updated for the Spartan-3A DSP.

www.xilinx.com

XILINX

Ultra-compact video-over-Ethernet camera

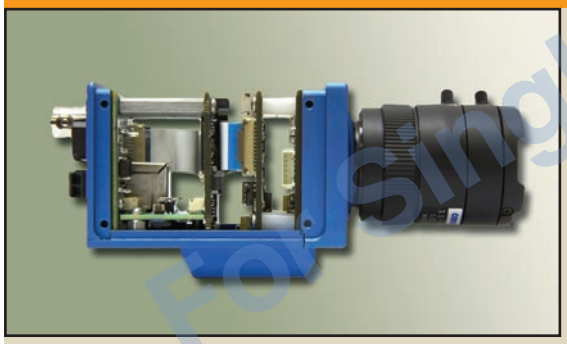

The Nuvation IP Camera is a production-ready camera for streaming real-time, full-color, compressed 720 x 480 (D1 resolution) video over Ethernet at 30 frames per second. Progressive image capture is -540 HTVL equivalent, and the camera has a wide dynamic range of 102 dB typical/120 dB max. (Video encoding options include H.264, MPEG-4/H.263, and JPEG.) The embedded Linux-enabled camera features a TI DaVinci DM6446 Dual-Core DSP with ARM926, along with Pixim's ultra Wide Dynamic Range (WDR) imaging technology. The Nuvation IP Camera comes in an ultra low-power, compact form factor measuring (L) 3.25" (79 mm), (H) 1.7" (43 mm), and (W) 1.8" (47 mm), excluding mounting bracket and lens. Interfaces include USB 2.0, analog video (NTSC/PAL), and Ethernet (RJ-45). The device is also RoHS compliant.

www.nuvation.com

NUVATION

Multifunction, single-slot VME card

The 64CS4 is a GbE-capable, multifunction, single-slot VME card designed to eliminate the complexity of using multiple, independent, single-function cards. The 64CS4 – ideal for avionics, ground mobile, and C3I applications – accommodates up to five independent function modules that may be selected from the included library. The card's GbE interface transfers data to and from the board without a VME backplane bus, enabling the board to be used as a stand-alone remote sensor interface without a separate computer board. The 64CS4 operates at -40 °C to +85 °C and 0 °C to +70 °C, and conduction cooling is available.

www.naii.com

NORTH ATLANTIC INDUSTRIES

Rugged USB interface

The SiliconDrive USB CF is a USB interface in the rugged and industry-standard CF (Type 1) form factor. It provides enhanced protection from host system voltage and power anomalies, preventing data and drive corruption. It also incorporates a more precise useable-life monitoring system to compensate for lower-endurance storage components. The SiliconDrive USB CF incorporates a robust storage management architecture to ensure long product life; therefore, it is well-suited to space-constrained embedded applications such as SBCs, edge routers, wearable computers, and medical devices that demand high-performance, high-reliability storage solutions. The SiliconDrive USB CF reduces host system design complexity, board space, and susceptibility to electrical noise by requiring that only 4 signals be routed on the host system motherboard versus the 50 required for standard CF interfaces.

www.siliconsystems.com

SILICONSYSTEMS

VXS switch fabric backplane

The VITA 41 VXS is a switch fabric backplane for military and aerospace applications, enabling users to integrate a cost-effective, high-performance upgrade to the latest VITA standards. The 12-slot backplane provides 10 VME64x payload slots and two fabric switch slots. Other features include interswitch links, SMT passive termination, and decoupling capacitors. A passive ABG option is also available. The VITA 41 VXS backplane comes in a dual star configuration and can also be customized with multiple configurations.

www.gavazzi-computing.com

CARLO GAVAZZI CS

COTS power management

The COTS Continuum Power Management features address power dissipation issues to benefit customers who want to use new technologies in legacy systems: Power Management Software API – A standardized API across product lines ensures board configuration control and system power optimization; Power Disconnect – Users power down a module via an external hardware mechanism, supporting the fail-over concept in systems with redundant chassis-level CPUs; CPU Low Power Mode – The CPU operates at lower power modes under software control, providing a single variant configuration to support both low- or high-performance mode; Peripherals Component Low Power Mode – Reduces power dissipation for component functions not being used by a given application, via software control; and Power Surge Prevention – Enables control of system module power-up sequencing.

www.cwcembedded.com

CURTISS-WRIGHT

High-power DC-DC converter

The DC500 converter is a high-power DC-DC unit that converts inputs of 18 V to 32 V to outputs of 28 Vdc with 90 percent efficiency. Customized output voltages are also available, and single outputs of up to 560 W are provided as standard. Input transients are 50 V, and input inrush current is limited to 50 A. DC500 converters are unaffected by time, temperature, and radiation, and they provide output short-circuit protection with autorecovery.

www.centuryele.com

CENTURY ELECTRONICS

Vertical PCB interconnects

The HPC-Z interconnects solve thick board drilling and wire density problems. They also meet the need for functional isolation with the capability of isolating high-speed channels on one subassembly and low-speed channels on another. HPC-Z provides interconnections within the board, and eliminates unnecessary Plated-Thru-Holes (PTH) drilling. The interconnections provide signal connections only where desired, with functional isolation for single- or double-side surface-mount partitions. HPC-Z eliminates PTHs that are blocking other wiring channels, which increases wireability, eliminates PTH stubs, and results in a thinner, faster board. HPC-Z provides a "PTH-like" connection at lamination with a pad-to-pad connection made with conductive epoxy. Since there is vertical connection with a lamination in place of a PTH, aspect ratios for drill and plate are limited to subassembly dimensions.

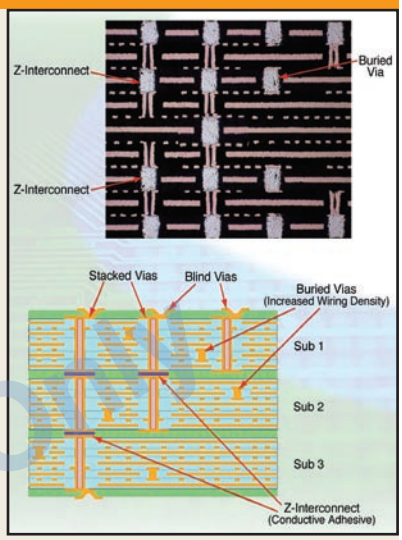

www.eitny.com

ENDICOTT INTERCONNECT TECHNOLOGIES, INC.

Rugged, ultra-slim laptop

The M230 is a rugged laptop designed for field-based, harsh environment applications. The device packs high performance into an ultra-portable, slim design that meets MIL-STD 810F and IP54 standards for durability. It includes a large 14" XGA or 15" SXGA TFT LCD high-resolution display, an optional ATI M54 chip for additional graphics processing power, and storage up to 120 GB. The M230 accommodates a wide range of plug-ins and offers an Intel Core Duo L2400 1.66 GHz processor, 2 MB L2 cache, and 512 MB DDR2 (expandable to 4 GB). Options include a sunlight readable display, touch screen, waterproof membrane keyboard, TPM 1.2 and Kensington lock security features, and a smart card reader. It also accommodates an optional GPS receiver.

www.getac.com

GETAC INC.

RISC-based, ready-to-run computer

The ThinkCore IA240-241 is a RISC-based, industrial ready-to-run embedded Linux computer featuring a MOXA ART 32-bit ARM9 industrial communication processor. Memory comprises 64 MB onboard RAM and a 16 MB flash disk. Also provided are four RS-232/422/485 serial ports, four-channel digital input, and four-channel digital output, along with dual 10/100M Ethernet for redundant networking. ThinkCore IA240-241 also offers wireless LAN expansion (802.11b/802.11g), a robust fanless design, and an IP30 protection mechanism. SD socket storage expansion is supported, and the unit can be utilized via DIN-rail or wall-mount installation. ThinkCore IA240-241 operates at -40 °C to +75 °C.

www.moxaUSA.com

MOXA TECHNOLOGIES

Industrial-grade box computer

The Matrix 520 is an industrial-grade, Linux-ready box computer. The unit includes an ATMEL AT91RM9200 (ARM9-core) CPU and a Linux 2.6.x prebuilt operating system. Memory comprises 32 MB SDRAM and 16 MB flash, and storage is provided via an internal SD memory card slot. I/O includes two host and one client USB 2.0 ports; two 10/100 Mbps Ethernet ports; 21 pins, TTL-level GPIO; and eight 921.6 Kbps baud TTY (serial) ports. A GNU C/C++ tool chain is included, and power input is 9~40 Vdc, 300 mA @ 12 V.

www.artila.com

ARTILA ELECTRONICS

Embedded system module

The EM6 SBC is an Embedded System Module (ESM) ideal for embedded industrial and communications applications where high computing power is necessary. The single board computer is powered by an Intel CoreT Duo or single core Celeron M processor, along with the Intel 3100 chipset. The EM6's Altera Cyclone FPGA provides flexibility, and required system I/O can be tailored for specific applications using IP cores such as IDE, graphics, additional fieldbus and legacy interfaces, and binary I/O. Memory includes 1 GB flash and DDR2 SDRAM with ECC for shock and vibration resistance. Front I/O comprises two GbE controllers via PCI Express and two COM interfaces via RJ-45 connectors. The EM6 can function as a stand-alone system, with an application-specific carrier card, or be connected to additional PCI-104 modules.

www.menmicro.com

MEN MICRO

"Hygienic design" industrial PC

The S19 is a 19" stand-alone industrial PC in a V2A (304) stainless steel enclosure. The enclosure is completely sealed in a "hygienic design" with no external fan. The NEMA12- and NEMA4-compliant computer features an Intel Celeron M @ 1.0 GHz or an Intel Pentium M @ 1.4 GHz. The S19's TFT XGA (1,280 x 1,024) 300 cd/m^2 display has a resistive analog touch screen with protective foil. Memory ranges from 512 MB to 2 GB, and I/O includes COM 1 RS-232, COM 2 RS-232, LPT, and 4x USB 2.0. The unit measures 19.6" (W) x 17.1" (H) x 5.4" (D) (498 x 435 x 136 mm) and weighs just under 50 lbs. The S19 operates at temperatures of 32 °F to +104 °F (0 °C to +40 °C).

www.noax.com

NOAX TECHNOLOGIES AG

Rugged SSD card

Advanced Media introduces its harsh environment RiDATA Solid State Disk (SSD) card, geared toward the aerospace and aviation markets. The nonvolatile flash SSD card features 16 GB and 32 GB capacities. RiDATA uses NAND flash memory technology and offers 1.8" and 2.5" module-type flash SSD for an ATA/IDE interface and 2.5" for a SATA interface. It also supports up to PIO Mode-4 and up to Multiword DMA Mode-2. The device has a flash media interface of 8- or 16-bit access and supports up to eight flash-media devices directly and 32 devices through external decoding logic. Its ECC corrects up to three random 12-bit symbols per 512-byte sector. RiDATA operates at 0 °C to +70 °C for commercial operation and -40 °C to +85 °C for industrial operation.

www.ritekusa.com

ADVANCED MEDIA (RITEK USA)

Nano-structured thermoelectric cooler

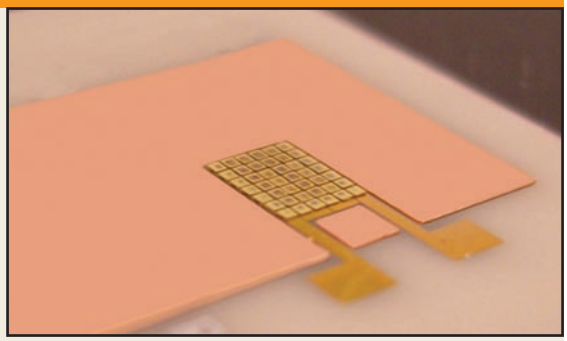

Nextreme's thin film Embedded Thermoelectric Cooler (eTEC) addresses thermal management needs of the electronics, photonics, biotech, and defense/aerospace industries. Functioning as a miniature, solid-state heat pump, eTEC is ideal for cooling hot spots that result from localized areas of high heat flux on an IC. The nanostructured devices are designed to add only 100 microns of height to a heat spreader, enabling unobtrusive integration close to the heat source. The eTEC has an ultra-fast, millisecond response time for rapid cooling and heating to maintain an application's precise temperature. The device pumps a maximum heat flux of 150 W/cm^2 with some designs delivering as much as 400 W/cm^2.

www.nextreme.com

NEXTREME

IN THE HEAT OF BATTLE WE'LL HELP YOU KEEP YOUR COOL.

Temperature extremes on the battlefield can be brutal. They play havoc with the thermal limits of today's high-performance electronics and make aerospace and defense system integrators fight a two-front challenge – maximizing performance while beating the heat. We can help. We're experts in power management and thermal design. From intelligent component selection process to innovative, patented board and system cooling technologies, we knock-out heat so you can take full advantage of today's cutting-edge processing power.

Our new 3U VPX and CompactPCI Single Board Computers utilize ultra-low power processors from P.A. Semi™ and Intel® to maximize performance-per-watt in weight and space-constrained environments.

CURTISS WRIGHT Controls
Embedded Computing

www.cwcembedded.com

POWER MANAGEMENT... ABOVE & BEYOND