COTS JOURNAL®

Special Supplement:
FPGAs: The New Matrix for Design

PROCESSING CHOICES EXPAND FOR MILITARY APPS

PLUS:

Configurable Computing Solutions Ready to Take Off

Tech Focus: FPDP/FPDP II Boards

Volume 8 Number 5 May 2006 www.cotsjournalonline.com

An RTC Group Publication

Need reliability?
–40° to +85°

EPIC™ XE–900
1.0 GHz CPU

Features	XE–900	XE–800	XE–700
CPU	Via Eden	AMD Geode GXI	STPC
Clock speed	400 MHz; 733 MHz; 1.0 GHz	300 MHz	133 MHz
BIOS	General Software	Phoenix	Phoneix
DRAM support	to 256 MB	to 256 MB	32/64 MB
Compact/Flash	Type I or II	Type I or II	Type I or II
COM 1	RS–232	RS–232/422/485	RS–232
COM 2	RS–232	RS–232/422/485	RS–232/422/485
COM 3	RS–232	NA	RS–422/485
COM 4	RS–232	NA	RS–232
COM 5	RS–232/422/485	NA	NA
COM 6	RS–422/485/TTL	NA	NA
LPT1	0	0	1
EIDE	2	2	1
USB	2	6	2
CRT	1600 x 1200	1280 x 1024	1280 x 1024
Flat panel	LVDS	yes	yes
Digital I/O	24–bit prog.	48–bit prog.	24–bit prog.
Ethernet	10/100 Base–T	Dual 10/100 Base–T	10/100 Base–T
Expansion	PC/104 & *Plus*	PC/104 & *Plus*	PC/104
Power	3.6A operating	1.6A max.	1.6A max
Temp. range	–40° to 70/85° C	–40° to 80° C	–40° to 80/85° C
Shock/vibration	40/5g	40/5g	40/5g

Need Linux, QNX, Windows®?
Try our OS EMBEDDER™ KITS

Our kits are the shortest path to a successful OS on an Octagon embedded computer.

- Pick your Octagon SBC
- Pick the OS you prefer: Linux, Windows, QNX

Octagon delivers a high performance, total solution.

Typical Linux kit includes:

- *Target CPU card*
- *Preloaded OS image on 256 MB industrial CompactFlash*
- *256 MB SO–DIMM module*
- *Interface cables*
- *Hard copy of manual*
- *Mouse*
- *CPU OS bootable CD*
- *Optimized OS version*
- *Full driver support for on–board hardware*
- *X–Windows support*
- *Example applications and source code*
- *Extra documentation*

Need PC/104 expansion?
Try our XBLOKs®

XBLOKs offer the best compromise in cost and function for both PC/104 and PC/104-*Plus*. Only 44% the size of a standard PC/104 card, you can add two functions to your system but increase the stack height by only one level. –40° to 85° C. Heat diagram shows enhanced cooling.

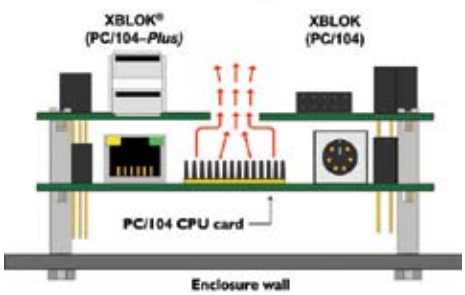

X–SRAM–2 MB
- 2 MB high speed, SRAM
- Read and write at full bus speed
- Pointers to memory saved if CPU resets or loses power

X–DIO–48 bit programmable digital I/O
- 48 digital I/O, 5V compatible
- Source and sink 16 mA per output
- Direct connection to opto-module racks

X–COM–2 dual UART
- Up to 230.4 kBaud data rate
- Supports RS–232/422/485
- RS–485 fault protected to ±60V

X–LAN–1 Ethernet LAN
- 10/100 Base–T, Intel 82551ER
- Fully plug–n–play
- High performance, PCI bus interface

X–USB–4 quad USB 2.0
- Speeds up to 480 mbps
- Mix and match USB 1.1 and 2.0
- Current–limited ports can supply 500 mA to external devices

Need a fanless system?
NEW CONDUCTION COOLING SYSTEM

Designed for the **XE-900**, our conduction cooling system eliminates a fan even at 1.0 GHz.

For a full listing of Octagon Systems products, visit us at
www.octagonsystems.com

OCTAGON SYSTEMS

RC-1 SINGLE SLOT CHASSIS

6U VMEbus or cPCI rugged enclosure, includes a built-in PMC site for I/O flexibility and safely dissipates over 80 W to free air.
www.cwcembedded.com/cha1

RC-2 AND RC-5 TWO AND FIVE SLOT CHASSIS

6U VMEbus or cPCI rugged system enclosures with a 28VDC power supply module and customizable front and rear I/O panels.
www.cwcembedded.com/cha1

OUR CHASSIS **PROTECT** MILITARY ELECTRONICS TO HELP **PROTECT** OUR TROOPS.

Modern military programs present a wide range of environmental and performance issues for system electronics. At Curtiss-Wright Controls Embedded Computing, we specialize in producing rugged enclosures which deliver steadfast electronics protection despite intense shock, vibration, corrosive atmospheres and severe temperatures. Curtiss-Wright battle-tested chassis help ensure the safety of our troops by providing reliable operation in the harshest environments.

CURTISS WRIGHT Controls Embedded Computing

Innovation In Motion

WWW.CWCEMBEDDED.COM

NEW

LIQUID-COOLED CHASSIS

Cutting-edge technology, dissipates up to 450 W per slot!
www.cwcembedded.com/cha1

1 ATR LONG SEVEN TO TEN SLOT CHASSIS

6U VMEbus or cPCI rugged system enclosure available in forced-air cooled cooling options.
www.cwcembedded.com/cha1

3/4 ATR FIVE SLOT CHASSIS

6U VEMbus or cPCI rugged system enclosure with 28VDC power supply module and forced-air cooling options.
www.cwcembedded.com/cha1

RcPCI 3U FIVE SLOT cPCI CHASSIS

Compact and powerful rugged system enclosure with forced-air cooling options, safely dissipates 130 W.
www.cwcembedded.com/cha1

COTS JOURNAL

Features

May 2006 Volume 8 Number 5

Main Feature

Tailoring Processors for Military Applications

Special Section

FPGAs: The New Matrix for Design

Technology Focus

FPDP and FPDP II Boards

COTS (kots), *n.* **1.** Commercial off-the-shelf. Terminology popularized in 1994 within U.S. DoD by SECDEF Wm. Perry's "Perry Memo" that changed military industry purchasing and design guidelines, making Mil-Specs acceptable only by waiver. COTS is generally defined for technology, goods and services as: a) using commercial business practices and specifications, b) not developed under government funding, c) offered for sale to the general market, d) still must meet the program ORD. **2.** Commercial business practices include the accepted practice of customer-paid minor modification to standard COTS products to meet the customer's unique requirements.

——**Ant.** When applied to the procurement of electronics for the U.S. Military, COTS is a procurement philosophy and does not imply commercial, office environment or any other durability grade. *E.g., rad-hard components designed and offered for sale to the general market are COTS if they were developed by the company and not under government funding.*

Departments

Coming in June...
...See Page 85

Cover Story ..... Page 16

Using a modified Boeing 707-300 series commercial aircraft, the E-8C Joint Surveillance Target Attack Radar System (Joint STARS) is an airborne battle management and command and control (C2) platform. The system conducts ground surveillance enabling commanders to develop a detailed understanding of the enemy situation, to monitor activities including treaty compliance, and to support attack operations and targeting. The program is now undergoing a major upgrade to its onboard electronics, computing, networking and display systems.

Publisher's

Notebook

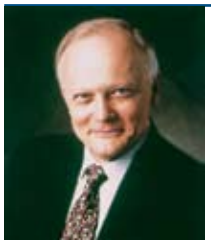

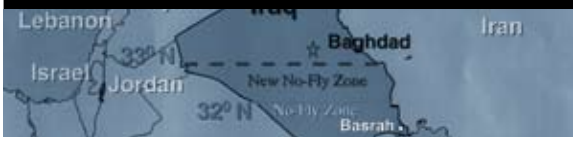

VME Skyrockets

I always get into trouble when I try to write about technology, even if it's only from a marketing viewpoint. Perhaps my age and gender dictate that although I know I shouldn't, I can't keep myself from doing it anyway. Twenty-five years ago I was probably much more naive or ignorant of what was going on than today. Back then there were only four open bus architectures—although "open" is a matter of interpretation—Multibus I, Multibus II, VERSAbus and VMEbus. I won't rehash old news, but fast-forward to today and VMEbus is still with us while the rest are in museums.

Today, I can't count the number of buses: some based on the PCI architecture, a slew of serial/switched fabrics, VME and its variants and so on. Note that I'm using a very generic interpretation of the term bus: a vehicle over which more than one electronic element exchanges data with another. There must be hundreds of buses if you count all the variants. I remember when Futurebus was talked about as being the "upgrade" path for all VME and Multibus users. It never came out of the incubator because there were too many variants in the Futurebus spec, and none of the variants developed a "critical mass" in the marketplace. How is it then that we have so many different bus architectures in existence today and so many companies that aren't going bankrupt?

Probably the primary answer to this question is that the embedded computer marketplace is much larger than it was fifteen years ago. And it continues to grow at an enormous rate. As a result, each of these technologies can grow large enough to sustain a small number of suppliers. Meanwhile, each of these buses provides sufficient diversity from the other to attract and keep a critical mass of users. These users are then nurtured by the suppliers through the addition of upgrade paths—upgrade paths that include some way to retain the user's investment. In my mind's eye, I picture the life cycle of a bus as a Fourth of July fireworks skyrocket. It shoots out, breaks into different paths, then each of those produces paths. Then, eventually, it fades and is gone.

Depending on what you've used and what interests you, every one of us has a different opinion of the bus architectures on the scene today. For shear tenacity the VMEbus has to be the most durable with a life line that as of this year stretches back twenty-five years. To what does VME—as it's now called—attribute its success? Well, like most things, luck and timing are the biggest factors. If it wasn't for its success in penetrating the military market, VME would only be an afterthought in today's market. Why did VME become the darling of the military? Several reasons: form-factor, durability, a history, large supplier base, wide array of available products and backward compatibility. But timing and luck came into play in that VME had all those things just when the military was moving to accept open standards and systems. It wasn't leading-edge performance or technology that interested the military. In fact, as the military was getting interested in VME, the rest of the world was getting itchy to move away from it to design products based on other architectures.

So, just what is VME's future? We have a whole host of VITA 40 series numbers in the VSO cranking out variations of VME. Some offer a lot of backward compatibility, some very little. Each variant is designed to solve or complicate a market need, while trying to preserve some portion of a user's prior investment in VME.

Each time a new bus architecture is being presented to the market, I get the feeling like Moses is coming off the mountain carrying tablets. I've seen a lot of these trips, and most recently it's been VITA 58. I don't fully understand it, but I'm supposed to get a complete sermon at MEECC. The initial information is that there will no longer be connector issues, nor bus architecture issues; everything will be inside a metal can and you will pull one out and replace it with the same type or a different type.

The reason we have boards and buses is to have the ability to "configure" what we need. Most of us are old enough to realize that every ten years we go through a cycle where pundits proclaim that boards/modules will become obsolete because everything will be put on silicon. Yes, we do continue to put more on silicon, then we take those denser chunks of silicon and put them on boards or modules and the modules into systems, and the beat goes on. Until we get to a level of integration where we can put every conceivable electronic function necessary to run something like a tank in one module named "tank," and only insert or remove this one module, then we'll need to have electronic subsystems. How those subsystems are divided and what they will contain will evolve and change over time. Maybe VITA 58 solves all that and we will finally see the light go out on the VME skyrockets. ▮▮

Pete Yeatman, Publisher
COTS Journal

GE Fanuc Automation

World-Class Solutions for World-Class Challenges

When the time comes to put your trust in a technology partner for the world's most demanding military applications, put your trust in GE Fanuc Embedded Systems. With GE Fanuc as your partner, you'll be teaming with an industry leader with a long and successful history of providing the military market with world-class solutions for world-class challenges.

GE Fanuc can quickly and effectively deliver a customized application for your specific requirements. We can also provide an off the shelf portfolio of SBCs, switches, controllers, rugged flat panel monitors and PCs, networking devices and much more, to meet your wide range of embedded computing needs. When the choice means this much, choose GE Fanuc.

Learn more at gefanuc.com/embedded

VME-7807
Intel® Pentium M VME
Single Slot Single Board Computer

- VITA 31.1
- Processor speeds up to 1.8 GHz
- Up to 1.5 GB DDR SDRAM
- 64-bit, 66 MHz PCI-X PMC site
- Dual 10/100/1000 Ethernet
- 10/100 Ethernet
- Four 16550-compatible serial ports
- Four USB 2.0 ports
- SVGA support on the front panel
- Serial ATA interface
- Up to 1 GB CompactFlash

RM921
VME Managed Gigabit
Ethernet Switch with IPv6

- IPv6 support
- Layer 2/3/4 switching
- Twelve and twenty-four 10/100/1000 Ethernet port (front panel) options
- Copper and fiber connection options
- Full management support
- 802.1p, 802.1Q VLAN, deep packet filtering, link aggregation, Rapid Spanning Tree (802.1w, 802.1d), broadcast storm Control, port mirroring

VME-5565
Ultrahigh Speed Fiber Optic Reflective Memory with Interrupts

- Available in PMC, PCI, and VME form factors
- Easy-to-use shared memory network-operating at 2.12 Gbaud
- Low latency for real-time applications
- Network transfer rates up to 174 Mbyte/s
- Data connections up to 10 km (single-mode fiber)
- DMA controller
- Operating system and architecture independent
- Ultra low CPU overhead, onboard hardware manages communication protocol

Sentinel
Ruggedized Flat
Panel Monitor

- 10.4, 15, 19, and 21 inch AMTFT display sizes
- Optional touchscreen, picture-in-picture video (PAL/NTSC)
- Video signal support includes: VGA, SOG, NTSC, RS170, PAL, S-video and DVI
- Compliance for shock (MIL-S-901D, Grade A), vibration (MIL-STD-167-1), EMI/EMC (MIL-STD-461E) and environmental (MIL-STD-810E)
- Anti-reflective and anit-glare coatings
- Field replaceable backlights and power supply

GE FANUC

Embedded Systems

COTS JOURNAL

Publisher

PRESIDENT
John Reardon, johnr@rtcgroup.com

PUBLISHER
Pete Yeatman, mail@yeatmangroup.com

EDITORIAL DIRECTOR / Associate Publisher
Warren Andrews, warrena@rtcgroup.com

Editorial

EDITOR-IN-CHIEF
Jeff Child, jeffc@rtcgroup.com

SENIOR EDITOR
Ann Thryft, annt@rtcgroup.com

CONTRIBUTING EDITOR
David Cotton, davidc@rtcgroup.com

MANAGING EDITOR
Marina Tringali, marinat@rtcgroup.com

STAFF EDITOR
Craig Choisser, craigc@rtcgroup.com

COPY EDITOR
Rochelle Cohn

Art/Production

CREATIVE DIRECTOR
Jenna Hazlett, jennah@rtcgroup.com

DESIGNER
Melissa Gaeta, melissag@rtcgroup.com

PRODUCTION DESIGNER
Kirsten Wyatt, kirstenw@rtcgroup.com

DIRECTOR OF WEB DEVELOPMENT
Jan Garcia, jang@rtcgroup.com

WEB DEVELOPER
Marke Hallowell, markeh@rtcgroup.com

Advertising

CALIFORNIA COASTAL ADVERTISING MANAGER
Diana Duke, dianad@rtcgroup.com
(949) 226-2011

WESTERN REGIONAL ADVERTISING MANAGER
Lea Ramirez, lear@rtcgroup.com
(949) 226-2026

EASTERN REGIONAL ADVERTISING MANAGER
Nancy Vanderslice, nancyv@rtcgroup.com
(978) 443-2402

EMEA SALES MANAGER
Marina Tringali, marinat@rtcgroup.com
(949) 226-2020

BUSINESS DEVELOPMENT MANAGER
John Koon, johnk@rtcgroup.com
(858) 755-4999

BILLING
Maggie McAuley, maggiem@rtcgroup.com
(949) 226-2024

COTS Journal

HOME OFFICE

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

EDITORIAL OFFICES

Warren Andrews, Editorial Director/Associate Publisher
39 Southport Cove, Bonita, FL 34134
Phone: (239) 992-4537 Fax: (239) 992-2396

Jeff Child, Editor-in-Chief
20A Northwest Blvd., PMB#137, Nashua, NH 03063
Phone: (603) 429-8301 Fax: (603) 424-8122

Ann Thryft, Senior Editor
15520 Big Basin Way, Boulder Creek, CA 95006
Phone: (831) 338-8228

Published by THE RTC GROUP

We think inside the box ...

because outside, it can get pretty rough!

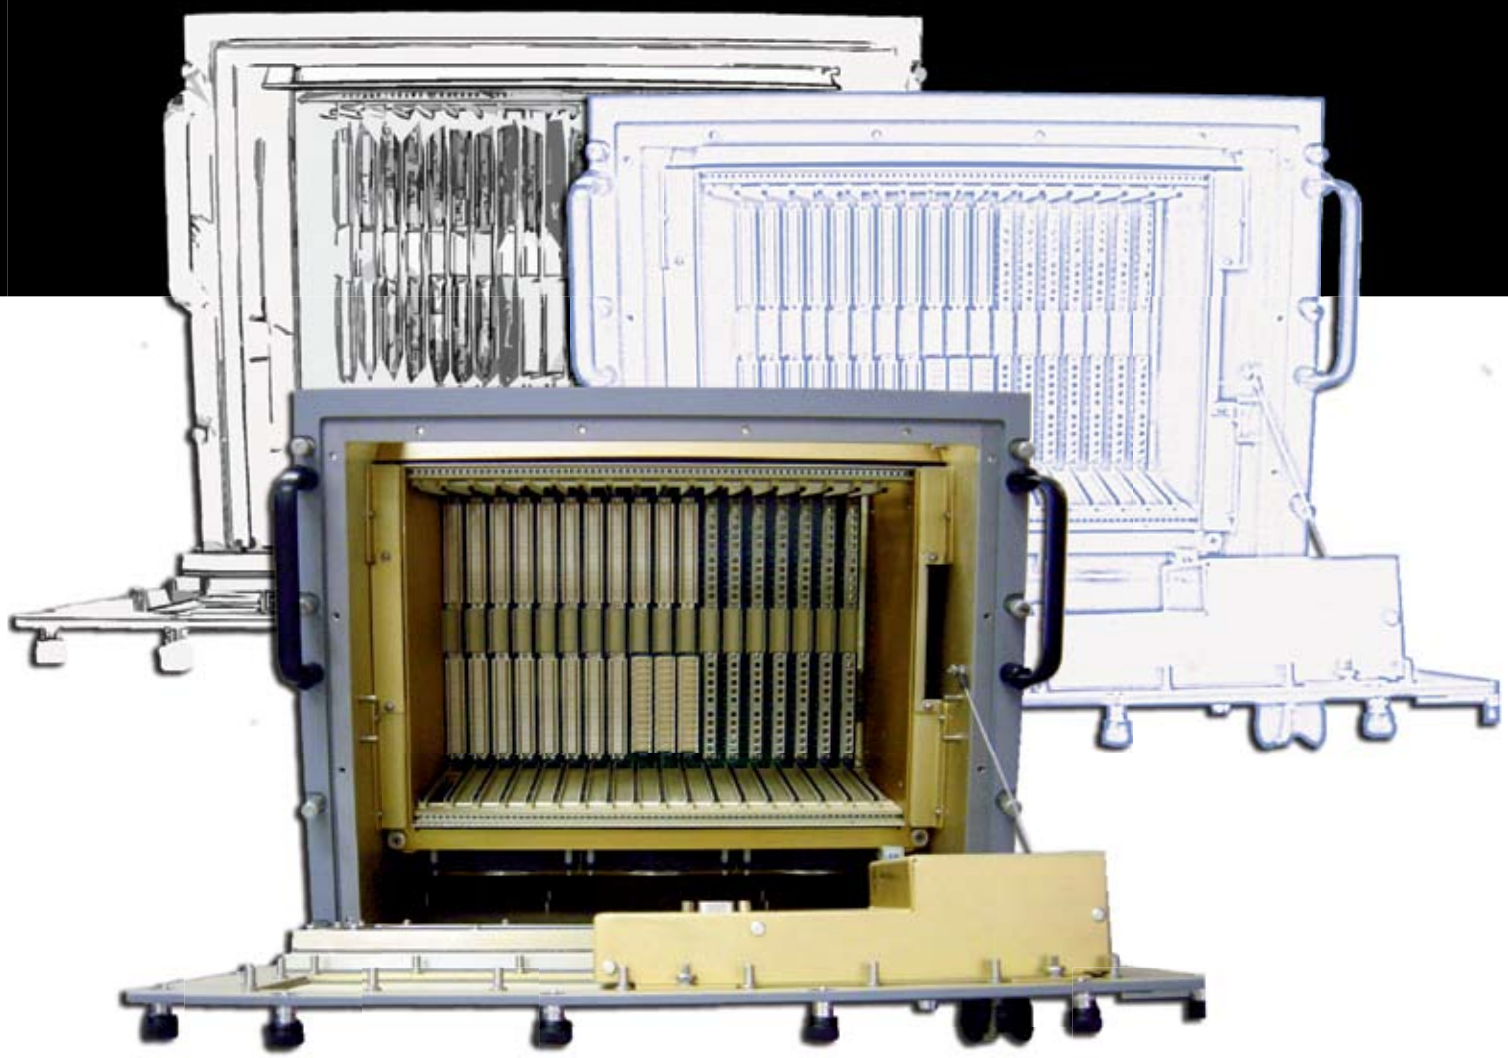

For more than 20 years, AP Labs has successfully leveraged design and manufacturing experience into system solutions that survive the harshest of environments.

Custom designs for backplanes, power systems and enclosures are possible with state-of-the-art tools, best-in-class components and quality manufacturing.

Be assured that our products specifications, warranty statements and service will exceed your expectations. Let our knowledgeable and responsive team of design and application engineers bring your project to a successful completion.

- ☑ Create Concept
- ☑ Leverage Experience
- ☑ Custom Design
- ☑ Quality Manufacture
- ☑ Compliance Test
- ☑ Rugged Solution

www.aplabs.com rugged@aplabs.com
10864 Thornmint Road, San Diego, CA 92127
Phone: 800-822-7522

AP Labs

Defining Rugged

Real teamwork means not getting in each other's way.

So why aren't you working with Radstone?

Any group of individuals can claim to be a team. But a real team is motivated by a shared vision, a united purpose. A real team doesn't compete within itself: all of its members do what they do best at the right time and in the right place, their complementary skills harnessed in support of a common goal.

At Radstone, we don't just talk the talk. Real teamwork is our way of life – and for us, our customers are part of the team. When we win, our customers win. It's a game plan based on trust.

When it comes to mission-critical embedded computing solutions, Radstone teamwork delivers with maximum performance and minimum risk. That's why we're the team-mates of choice for some of the industry's biggest names.

Shouldn't you be working with Radstone?

RADSTONE
EMBEDDED COMPUTING

Radstone. One team. One aim. One winning result.

FOR MORE INFORMATION VISIT
www.radstone.com

OR CALL TOLL FREE
800 368-2738

The Inside Track

AFRL Selects Mercury ARIES System for UAV Reconnaissance Program

The Air Force Research Lab (AFRL) at Wright-Patterson Air Force Base has selected Mercury Computer Systems to provide computer hardware and services for the Continuous Look Attack Management for Predator (CLAMP) Program. The goal of CLAMP is to develop and transition advanced sensor exploitation capabilities to the Long-Endurance Predator UAV (Figure 1).

CLAMP is the first U.S. Air Force program to undertake Mercury's ARIES (Airborne Reconnaissance Image Exploitation System) concept, in which airborne multi-sensor platforms will use stored sensor data for comparative purposes. ARIES is an adjunct processor that is designed to facilitate the migration of ground-based algorithms to the platform, adjacent to the sensor and with direct access to original sensor data, so that image exploitation can occur in real time.

As the Predator UAV flies and collects data over a designated area for hours at a time, the multi-look, multi-sensor nature of its mission can be fully exploited. By enabling the UAV to store, retrieve and process sensor data over a long period of time, warfighters will be able to detect changes in tactical conditions and allow them to pinpoint attacks or avoid dangerous situations.

In addition, CLAMP will combine other selected technologies in a systematic approach to enable the Predator UAV to fulfill its Hunter-Killer role with improved reliability and reduced collateral damage. Mercury plans to upgrade two MP-510 multiprocessor systems previously purchased by the AFRL, and to assist in the performance optimization of key applications to be tested by the CLAMP program.

Figure 1

The goal of the Continuous Look Attack Management for Predator (CLAMP) Program is to develop and transition advanced sensor exploitation capabilities to the Long-Endurance Predator UAV. The CLAMP system will enable the UAV to store, retrieve and process sensor data over a long period of time.

Mercury Computer Systems
Chelmsford, MA.
(978) 256-0052.
[www.mc.com].

EADS and Saft Team to Form U.S.-Based Thermal Battery Company

EADS and Saft America have inked a deal to form a new thermal battery subsidiary called Advanced Thermal Batteries Inc., located in Cockeysville, Maryland. This builds on the two firms' existing joint venture called ASB. ASB ranks as number one in Europe and number two worldwide in the thermal battery market, supplying thermal batteries for missiles, launchers, aircraft, torpedoes and submarines. Thermal batteries are inert batteries designed for absolute reliability and shelf storage of at least 15 years. ASB, Saft America and EADS are merging their assets to offer both development and production capacities in the United States. Advanced Thermal Batteries Inc. supplies customized thermal batteries for missiles and smart weaponry.

The new subsidiary will provide to its American customers all the technologies that have been proven by its parent companies. The new subsidiary, through ASB, already has sales in the American market, and hopes to become a leader in the thermal battery field in the United States.

EADS North America
Arlington, VA.
(703) 236-3300.
[www.northamerica.eads.net].

Saft America
Cockeysville, MD.
(410) 771-3200.
[www.saftbatteries.com].

NCSA and SRC Computers Collaborate on Reconfigurable Computing Tools

SRC Computers, a manufacturer of reconfigurable computing systems, has entered into a joint development agreement with NCSA (National Center for Supercomputing Applications). Under the agreement NCSA will develop software and processes that will allow programmers in the embedded systems and signal processing communities to write programs using MATLAB/Simulink for use on SRC's reconfigurable computing systems. The result of the collaboration will describe in detail how to use the Mathwork's MATLAB/Simulink, Xilinx DSP System Generator and the SRC Carte programming environment macro capability to implement Simulink fixed-point designs on SRC's reconfigurable MAP processors.

Get Connected with companies mentioned in this article.
www.cotsjournalonline.com/getconnected

NCSA will be using SRC's newest portable MAPstation product (Figure 2) in its development work. Originally developed for the Air Force for use in mid-size UAVs, SRC's portable MAPstation is a compact general-purpose processor that typically performs operations such as radar processing, image processing and spectrum analysis 100 times faster than a microprocessor. NCSA will develop several digital signal processing applications for SRC's portable MAPstation to showcase the MATLAB/Simulink-to-MAP programming ability. One set of these applications will build on NCSA's earlier work in software-defined radio as part of the National Center for Advanced Secure Systems Research (www.ncassr.org) funded by the Office of Naval Research (ONR). Other applications for the portable MAPstation will include speech analysis and image processing. These applications will be ready for demonstration as early as Q4 2006.

The MathWorks
Natick, MA.
(508) 647-7000.
[www.mathworks.com].

SRC Computers
Colorado Springs, CO.
(719) 262-0213.
[www.srccomputers.com].

CMC Selects Greenhills for Cockpit Avionics

CMC Electronics has selected the Green Hills Platform for Avionics with the INTEGRITY-178B RTOS to support the development of integrated cockpit avionics solutions. The RTOS will be used in the development

Figure 2

SRC's portable MAPstation is a compact general-processor system that typically performs operations such as radar processing, image processing and spectrum analysis 100 times faster than a microprocessor. The lightweight airborne enclosure weighs less than four pounds and has a maximum power consumption of 86W, 40W typical.

of a new CMC Electronics Aircraft Management System product line. The product supports an all-glass cockpit anchored by two powerful CMC Electronics Integrated Avionics Computers that are fully qualified to commercial standards. The rugged mission system uses a PowerPC 7457 processor and will support multiple display functions.

INTEGRITY-178B is a time and memory partitioned operating system, certified to DO-178B Level A with full ARINC-653-1 compliance. Support for ARINC-653-1 with its partitioning definition allows developers to deploy multiple applications on a single processor, at potentially multiple safety certification levels. That allows developers to reduce the number of onboard computers needed to support multiple software systems. The Platform for Avionics with INTEGRITY-178B is certified for multiple languages, including: Ada, C and Embedded C++.

Green Hills Software
Santa Barbara, CA.
(805) 965-6044.
[www.ghs.com].

CPU Technology Gets Subcontract for Bradley Turent Drive PIB

BAE Systems has awarded CPU Technology a follow on production subcontract for Bradley Combat Systems vehicle Turret Drive Position Interface Box (PIB) control systems. The new PIB systems will be installed in remanufactured and upgraded Bradley A3 Vehicles at BAE Systems in York, Pennsylvania. CPU Tech began providing the Bradley Position Interface Box two years ago with a modern System on Chip (SoC)-based PIB that is smaller, lighter, more reliable and lower cost than earlier systems.

The Bradley A3 (Figure 3) features an advanced digital architecture that integrates communications equipment, digital sensors, battle management systems, embedded diagnostic and training systems. Last June, BAE Systems was awarded a series of delivery orders and contract modifications worth $1.127 billion from the U.S. Army Tank-Automotive and Armaments Command (TACOM) to remanufacture and upgrade more than 500 Bradley Combat System vehicles.

CPU Technology
Pleasanton, CA.
(703) 251-2568.
[www.cputech.com].

Figure 3

Shown here are soldiers in a Bradley armored vehicle assigned to the Oregon Army National Guard's Company C, 2nd Battalion, 162nd Infantry Regiment. The A3 upgrade version of the Bradley features an advanced digital architecture that integrates communications equipment, digital sensors, battle management systems, embedded diagnostic and training systems.

COTS Websites

www.aaai.org

Artificial Intelligence Site Provides Portal to All Things AI

The DoD's interest and investment in autonomous vehicles—land, sea and air—continues to soar. As a result, interest in the science of artificial intelligence is moving into the foreground of military system designers. One source of such information is the American Association for Artificial Intelligence (AAAI), a nonprofit scientific society devoted to advancing the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines. AAAI's goals also include increasing public understanding of artificial intelligence, improving the teaching and training of AI practitioners, and providing guidance for research planners and funders concerning the importance and potential of current AI developments and future directions.

AAAI's Web site provides a wealth of information about major AAAI activities such as conferences, symposia and workshops, and its books, proceedings and reports. The list of AI topics and subtopics on the site number in the hundreds, and the links to each topic page are listed alphabetically in an index list. Of particular interest is the topic page on Autonomous Vehicles, which in and of itself packs a wealth of information and links to related sites.

American Association for Artificial Intelligence, Menlo Park, CA.
(650) 328-3123.
[www.aaai.org].

Perform like never before.

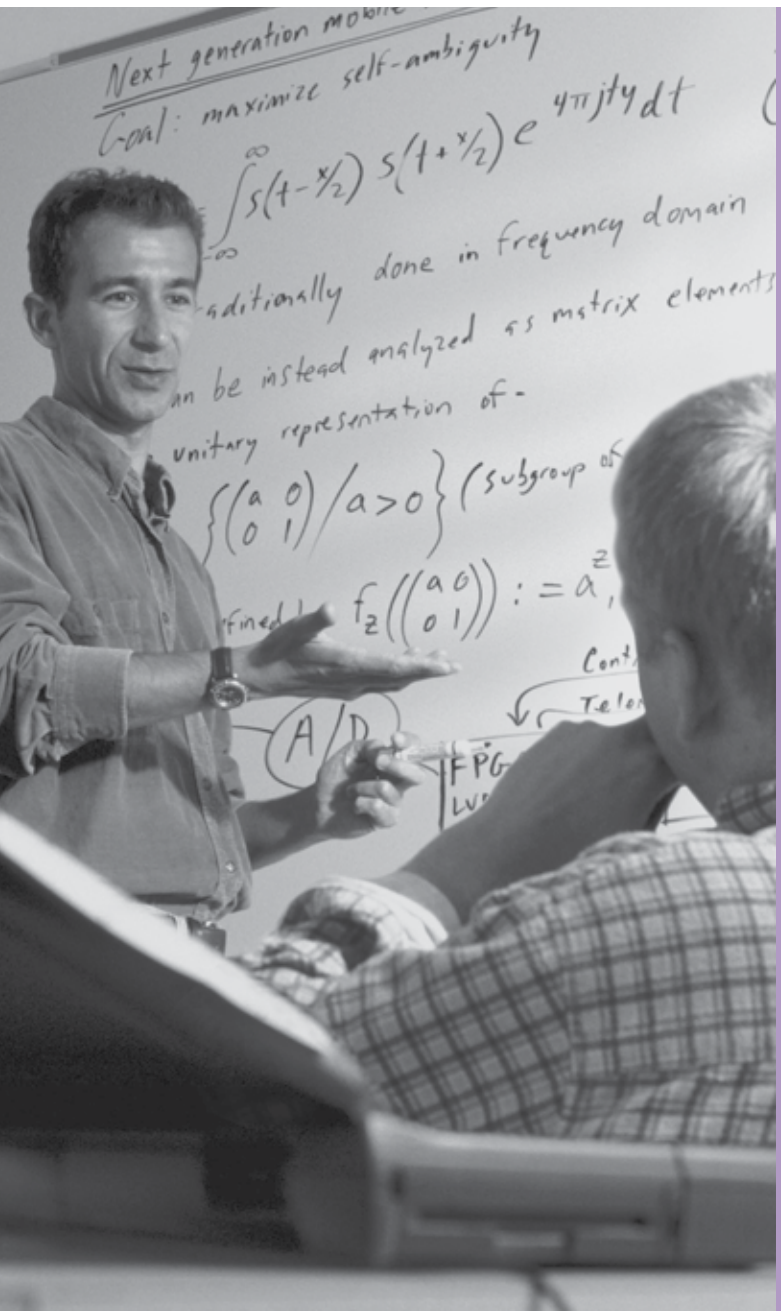

ANALOG DEVICES SHARC®

ANALOG DEVICES TigerSHARC®

From $5 to 3.6 GFLOPS, your DSP job just got easier.

Worried about rounding errors in your 16-bit arithmetic? Anxious about noise in your frequency domain processing? What about blowing your project schedule as you move your floating-point concept to fixed-point implementation? Analog Devices SHARC® and TigerSHARC® Processors solve your most frustrating numeric DSP challenges. Now you can really perform.

For more than 10 years, SHARC has dominated the floating-point DSP market, delivering exceptional core and memory performance complemented by outstanding I/O throughput. At 319 MFLOPS per dollar, the SHARC family makes floating point more accessible than ever. And for high performance multiprocessing, TigerSHARC Processors deliver GFLOPS performance with glueless, 1 Gbyte/sec multiprocessing link ports, and 24 Mbits on-chip memory.

If you're ready to speed your time to market and showcase your design's real potential, Analog Devices SHARC and TigerSHARC Processors deliver floating-point performance—and performance density—like never before.

SHARC Processor Leadership in MFLOPS/$

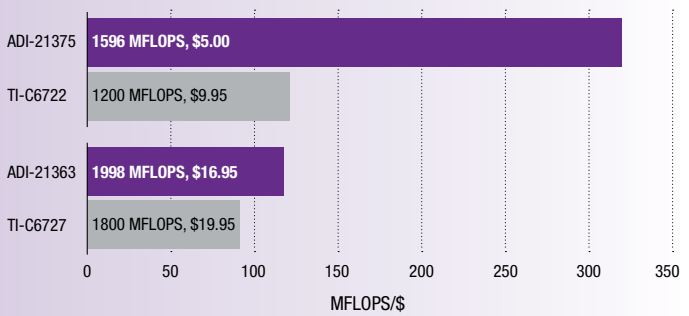

Leading Performance vs. Competition

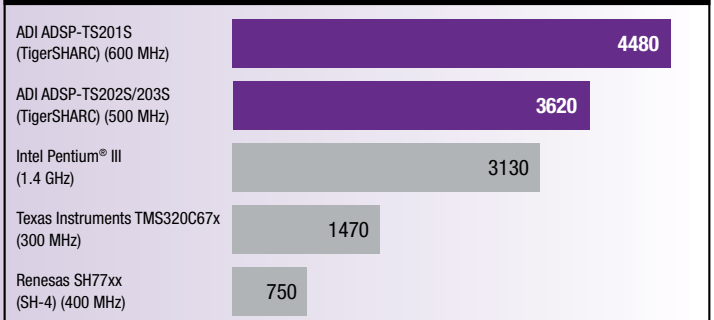

BDTi BENCHMARK PARTNER

The BDTImark2005™ provides a summary measure of DSP speed. For more information and scores see www.BDTI.com. Scores© 2005 BDTI.

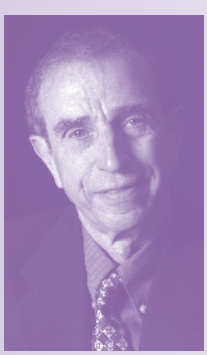

"Floating-point processing essentially eliminates concerns about dynamic range and ensures precision in the most intensive and demanding signal processing tasks—but it's been an expensive alternative in terms of both cost and power. Architectural advances are now making this capability far more widely accessible."

— Dr. Alan Oppenheim
Ford Professor of Engineering
Massachusetts Institute of Technology

SHARC Advantages

- Broad portfolio of code- and pin-compatible processors starting at $5 and providing 2.4 GFLOPS of signal processing performance
- High precision computation enabled by SHARC's 32-bit fixed- and floating-point architecture
- Up to 400 MHz core instruction rate
- Single-Instruction Multiple-Data (SIMD) computation architecture
- Simplified system design and lower system cost via a wide array of integrated peripherals

TigerSHARC Advantages

- Optimal balance of computational performance, large on-chip memory, and I/O bandwidth maximizes sustainable FFT performance
- Memory integration up to 24 Mbits accessible at core frequency enables internal execution of FFTs of 64 K points maximum
- Glueless, scalable multiprocessing capability with 5 Gbytes/sec aggregate throughput
- Interlocked pipeline for "out-of-the-box" execution of ported workstation code
- 24 Mbits of on-chip RAM—three times more than the closest competition

Mackie TT24™ Digital Live Console
Mackie uses SHARC's on-board DSP power to provide audio dynamics, equalization, and effects processing in a way that's practical, intuitive, and fast enough for live sound.

BMW M5 Series
BMW chose SHARC Processors for control of its sophisticated 500 hp V-10 motor. The V-10 engine control unit performs an amazing 200 million calculations per second to ensure optimal performance.

R&S® Radio Communication Tester CMU
With ADI floating-point processors, Rohde & Schwarz provides a versatile mobile radio test platform that conforms to multiple international standards and delivers maximum measurement speed for extensive testing of current and future mobile radio networks.

iCyt™ Cell Separation Instrument
TigerSHARC Processors are used in a high throughput, biological cell separation instrument designed and built by iCyt Visionary Bioscience, Inc. while utilizing Bittware's TigerSHARC development boards.

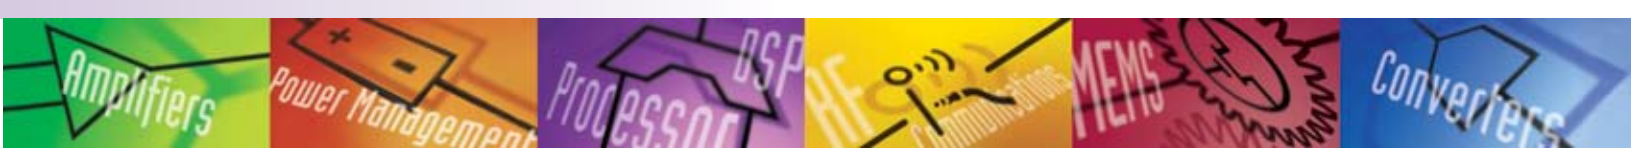

BittWare

Visit **www.analog.com/perform** for information on special bundled board offers from BittWare, Inc. and Analog Devices.

ANALOG DEVICES

Main Feature

Tailoring Processors for Military Applications

Processors Combine Forces for Military Systems

Next-generation, complex electronic defense systems often need multiple compute nodes. These are being supplied by multicore processors or by general-purpose processors in combination with FPGAs and/or DSPs.

Ann R. Thryft
Senior Editor

When it comes to the processors used in high-demand, complex compute-intensive military systems, the big news these days is the rise of FPGAs and the new generation of multicore processors.

FPGAs are now big enough and powerful enough to take over an increasing number of digital processing tasks from general-purpose DSPs, or to work with them in the same design. Military engineers are using FPGAs as coprocessors that complement high-performance general-purpose processors (GPPs) for signal processing and sensor data processing, as well as image processing. FPGAs are found in an increasing number and type of defense systems, including submarine and ordnance detection, communications intercept and direction finding, for tanks, UAVs and helicopters (Figure 1).

FPGAs, DSPs and high-performance GPPs are being combined in military designs to provide the right kind of processing punch in the right place. For repetitive signal processing, encryption/decryption or beam-forming algorithms that require fast execution, FPGAs shine. In contrast to fixed processor architectures, they are flexible and scalable at multiple levels.

Figure 1

Combinations of multiple processor types, such as general-purpose processors and FPGAs, are found in an ever increasing number of defense systems. The U.S. Navy's Airborne Laser Mine Detection System, used by pilots of H-60 helicopters, depends on this mix of processing power for processing signals and sensor data. Members of Explosive Ordnance Disposal Mobile Unit Eleven Detachment 9 fast rope from an SH-60F Seahawk helicopter assigned to Helicopter Anti-Submarine Squadron Two, during an air power demonstration in the South China Sea. (Photo by Photographer's Mate 3rd Class Jordon R. Beesley, courtesy of U.S. Navy.)

DSPs have been the processor of choice for applications requiring compute-intensive, high-speed calculations, since they can provide sustained, low-latency, high-throughput data processing. The benefits of DSPs include algorithmic capabilities optimized for fast floating-point math and programming with C. Applications that require different levels of processing will benefit from partitioning the design so it combines both FPGAs and DSPs.

Meanwhile, a new generation of multicore processors has appeared that will help improve the performance of next-generation electronic defense systems, such as radar, sonar, SIGINT and UAV control. Some of the most recent multicore architectures combine multiple cores of the same GPP on one chip so each core can run a separate program thread.

Mainstream processor vendors

Design to Deployment with One Graphical Environment

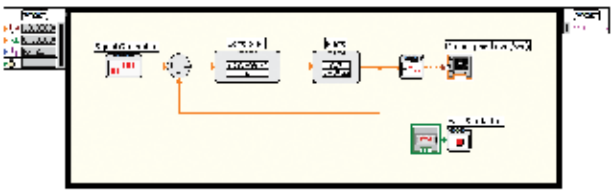

Design with LabVIEW

Interactive algorithm design tools

Native simulation capabilities

Built-in analysis and I/O

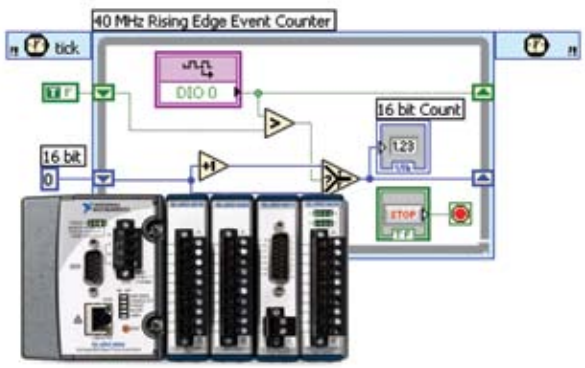

Prototype with LabVIEW

Reconfigurable FPGA I/O

COTS prototyping platform

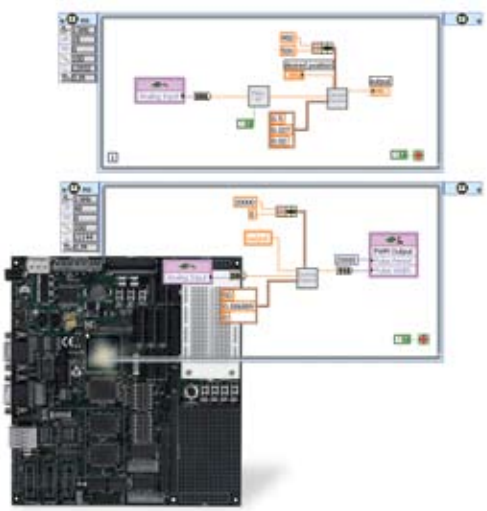

Deploy with LabVIEW

32-bit processor deployment

Same environment from algorithm design to implementation

LabVIEW

Graphical Development Platform for Design, Control, and Test

- Intuitive embedded programming representation
- Rapid application development
- Built-in I/O connectivity and ready-to-run analysis

Improve your productivity with LabVIEW graphical system design.
View technical tutorials online at **ni.com/labview/embedded**.

(800) 890 6229

2006-6658-821-101-D

NATIONAL INSTRUMENTS

FPGAs, DSPs and high-performance GPPs are being combined in military designs to provide the right kind of processing punch in the right place

such as Intel, AMD and Freescale are taking this symmetric multiprocessing approach to boost performance without burning huge amounts of power. For example, Intel's new dual-core technology—the 2 GHz Intel Core Duo T2500/1.6 GHz Core Duo L2400—boasts almost twice the performance of Pentium processors in the same space, and with only slightly higher power consumption. Other dual-core designs include the Freescale 1.5 GHz 8641 PowerPC with its dual integrated 64-bit memory controllers, and AMD's 2.2 GHz 64-bit, dual-core, 30/55/95W Opteron. Quad core designs are also in the works.

Another type of multicore processor is the tiled design. There are several different versions. The Cell Broadband Engine processor, developed by IBM, Toshiba and Sony Group, includes eight identical, synergistic processing elements in addition to a core based on IBM's Power Architecture. Originally created for the video game market, the Cell processor's peak performance is more than 200 GFLOPS, or 200 billion FLOPS.

The 64-bit, 25 GFLOPS CSX600 is a multi-threaded array processor that acts as an accelerator and coprocessor. This tiled processor is offered by Clearspeed on a PCI-X add-on board.

A tiled processor under development by the Defense Advanced Research project (DARPA) is the Morphable Networked Micro-Architected (MONARCH) chip. This will unite two radically different architectures into a device that can act either as a single chip or as a system-on-chip.

One card that beats a full house.

Intelligent Multifunction I/O Cards from North Atlantic Industries.

Function integration for reduced card count, weight, power, heat, and overall system cost.

- ~ Fill up to 6 slots per card with your choice of over 16 different function modules
- ~ **NEW!** Gigabit Ethernet Connectability
- ~ On-line Background Built-in-Test (BIT) alerts you of problems *during operation*
- ~ VME, cPCI, and PCI Platforms
- ~ Available as Conduction-Cooled/Wedge-Lock or Air-Cooled
- ~ Custom function designs
- ~ Can be configured with user definable FPGA
- ~ Extensive library of available functions
- ~ All major I/O Signal Types

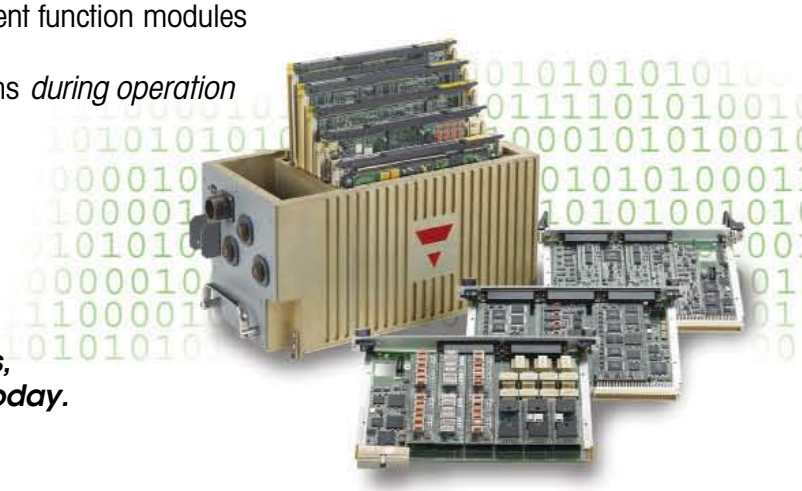

To find out more about Intelligent Multifunction Cards, visit www.naii.com/multi or call us at 631-567-1100 today.

North Atlantic Industries

Embedded Boards | Power Supplies | Instruments

631-567-1100 • Fax: 631-567-1823 • email: sales@naii.com

SBS knows MIL-STD-1553. For real time, deterministic systems, it's still the best option.

MIL-STD-1553 IS THE PREFERRED TECHNOLOGY for safety and life critical applications because it outperforms the bleeding edge. There's no other solution that offers the same reliability levels for military, avionic and space applications, and that's why SBS offers a full range of superior 1553 hardware and software.

We offer our MIL-STD-1553 products in CompactPCI®, VME, PMC, PC-104, PCMCIA and PCI as part of our commitment to provide options to system designers. Many of our offerings are specifically designed and tested to stand up to the demands of harsh environments that are common in defense and aerospace. Even now, we are designing and building rugged cards, boards and systems that will support 1553 into the next generation of land, sea and air platforms.

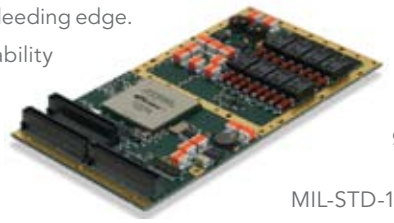

1553-PMC3
1 to 4 MIL-STD-1553 channels for PMC

MIL-STD-1553 is a proven technology that has a distinguished record of service in extremely critical situations. MIL-STD-1553 may not be the "latest and greatest," but then again, sometimes the latest isn't the greatest.

SBS Technologies®

SBS knows.
Find the 1553 solution you're looking for at **defense.sbs.com** or call **800.SBS.EMBEDDED**

Main Feature

Tailoring Processors for Military Applications

Commercial Processors in Military Applications

A wide variety of processor types is available for military applications. The use of a decision model can help narrow the choices for a given application type.

Dawn Levy, Senior Systems Application Engineer
Stephen Pearce, Technical Director
Mercury Computer Systems

The expanding array of processor types available today has significantly complicated the selection process for military imaging and signal processing applications. The various processor types range from FPGAs and DSPs to graphics processing units (GPUs) and single, vector and multicore processors.

This wide variety provides an opportunity to fine-tune hardware platform environments through an assessment of clear performance tradeoffs. Selecting the optimal processing environment to meet application requirements is critical in ensuring successful deployments, particularly in complex military environments.

Decision Model Criteria and Parameters

To simplify the selection process, a decision model can be developed with criteria such as bit resolution, software code type, field deployment requirements and applicability. This model can be used as a set of guidelines and includes parameters for, and a brief discussion of, each of these four decision criteria. It considers the following processor types: the Cell Broadband Engine (BE), GPUs, FPGAs, DSPs, PowerPCs (PPCs) and Pentiums.

The model's parameters for the first criterion, bit resolution, include single-bit, 8- or 16-bit, floating-point and double precision. Certain processor types are clearly more adept at high-precision processing than others. With their math processing engines, the Cell BE, PPC and Pentium are most qualified for double-precision operations. However, FPGAs and DSPs most skillfully perform arbitrary low-precision arithmetic. GPUs, on the other hand, are great for 8-, 16- and 32-bit floating-point operations as long as the dataflow and

Application Ranking	0	0	5	5	4	3	2	5	5	4	4	2	2	
	Resolution				Code Type			Air/Cond. cooled MIL				Applicability		
Processing Technology	Bit opns, e.g. image morph	8/16 bit opns	Float opns	Double Precn opns	Vector (SIMD)	Scalar	Condnl GPP C code	Perf./ cu. ft	Perf./ watt	Min. abs. watt/ chip	Perf ./lb	Ease of use	Range of appns	Final Score
Cell BE	1	4	6	6	6	2	2	5	5	3	3	4	3	145
GPU	2	2	5	1	6	1	1	4	4	2	2	1	1	91
FPGA	6	6	2	1	5	6	1	6	6	6	6	3	4	123
DSP	5	5	3	1	3	5	4	2	2	5	5	4	2	97
PPC	4	3	4	5	4	3	5	3	3	4	4	5	6	133
Pentium	3	1	1	4	1	4	6	1	1	1	1	6	5	84

To avoid double counting, only one rating for each pair of correlated metrics, shown in in shaded cells, is considered in the final score.

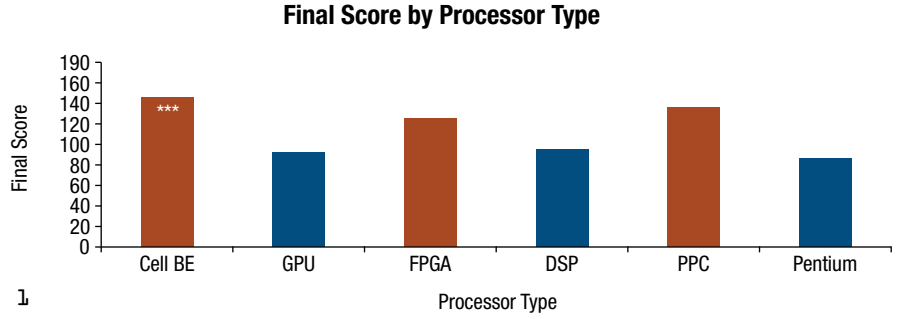

Figure 1

Processor Decision Model, Cell BE High Score example. The model's top row, highlighted in yellow, depicts application characteristics that are ranked from lowest to highest importance using a score of 0 to 5. At the bottom, the rank order from highest to lowest final scores across all processor types is shown as bar height in the graph.

Application Ranking	5	0	1	0	5	3	1	5	5	5	5	3	2	
	Resolution				Code Type			Air/Cond. cooled MIL				Applicability		
Processing Technology	Bit opns, e.g. image morph	8/16 bit opns	Float opns	Double Precn opns	Vector (SIMD)	Scalar	Condnl GPP C code	Perf./ cu. ft	Perf./ watt	Min. abs. watt / chip	Perf ./lb	Ease of use	Range of appns	Final Score

To avoid double counting, only one rating for each pair of correlated metrics, shown in in shaded cells, is considered in the final score.

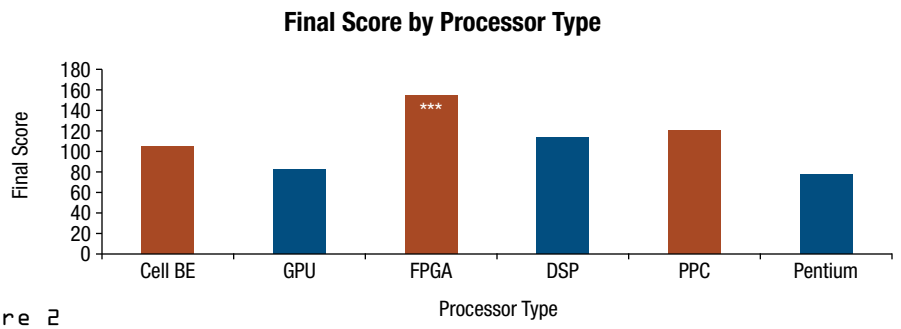

Figure 2
Processor Decision Model, FPGA High Score example.

processing types closely match their internal architecture, which is optimized for visualization.

The second criterion, software code type, characterizes application logic. Model parameters include vector single-instruction-multiple-data (SIMD) processing, scalar code and conditional, "if-then-else" general-purpose programming (GPP) C code.

Whereas the Pentium or PPC are best suited for conditional C code, the repetitive nature of vector processing lends itself best to the Cell BE, GPU or FPGA. FPGAs do very well when an algorithm must access its data in an order that is cache-unfriendly. For cache-friendly data access, a processor is often preferred. However, even in these cases an FPGA might still be the best choice if it is needed to get sensor I/O data into a system, because engineers might face the choice of adding a processor or just deploying a bigger FPGA for the I/O task.

GPUs excel when the algorithm happens to match the GPU architecture. This architecture is highly tuned to streaming matrix operations with moderately complex data structures, little conditional code and little data reuse. Data re-sampling is a GPU forte. To program a GPU, however, requires specialized training in graphics terminology.

DSPs excel in places where their instruction sets fit the application. They can handle some conditional C code. In some applications, they can provide close to the sort of performance one might expect from an FPGA. They thus occupy the middle ground between GPPs and more application-specific processing elements.

The third criterion, deployed environment, refers to requirements for ruggedization, such as air-cooled or conduction-cooled military environments. The model considers two basic parameters: performance per cubic foot and performance per pound.

Processors that perform well per cubic foot also achieve relatively strong performance per watt. Similarly, good performance per pound is typically a result of the chip's power efficiency, or minimum absolute watt per chip. To avoid double counting, however, only one rating for each pair of correlated metrics is considered in the model's final score.

In environments where weight is a critical concern, such as airborne platforms, devices that provide processing at the lowest absolute power per chip, such as DSPs and PPCs, are good candidates. In situations where the parallelism of an FPGA can be effectively exploited, the processing element's higher power

Need Military Gigabit Ethernet?

GigEXTREME™ Lightens Your Load

ET-71000 GigEXTREME™ TCP/IP Offload Engine

	GigEXTREME™	Other TOE Cards	NICs with Partial Offload
Processor Loading*	**5%**	**5%**	**50% to 70%**
RTOS Support	✓	✓	–
No Application Modification or Recompilation Needed	✓	–	✓
No Use of Upscreened Commercial Components	✓	–	–
Military Rugged	✓	**Some**	**Some**
IPv6 Capable	✓	–	–

* Processor loading measured for the TCP/IP receive network task on a 1 GHz Power PC® host.

Toll Free: 1-800-DDC-5757 **Web: www.ddc-web.com** **E-mail: info@ddc-web.com**

DDC®
©Data Device Corporation
ISO 9001:2000 REGISTERED

EMRISE
power group of companies

empowering applications

military
aerospace
transportation
telecommunication

CUSTOM AC/DC POWER SUPPLIES COTS MIL STD-461 HIGH TEMPERATURE SHOCK & VIBRATION MIL STD-704 HARSH ENVIRONMENT MIL STD-810 HIGH VOLTAGE EMI / RF

DC/DC CONVERSION PRODUCTS MIL STD-810 SINGLE & THREE PHASE INPUTS DO-160 MIL STD-461 COTS RUGGEDIZED CONSTRUCTION MIL STD-704 CUSTOM DESIGNS VME

Pascall ELECTRONICS LIMTED

Civil Aerospace DO-160
Custom Military Power Supplies
Mil –Std 704, 461, 810
VME Power Supplies

www.pascall.co.uk

XCEL Power Systems

Custom Military Power Supplies
Mil –Std 704, 461, 810
Harsh Environments
Configured Module Systems

www.xcelpower.com

RO Associates

50-1000W COTS Modules
Standard Products
DC-DC Converters
PFC front end Modules

www.roassoc.com

Power to the Power of 3

Companies are subsidiaries of **EMRISE** Corporation, Rancho Cucamonga, CA 91730

roassoc.com
North American Sales & Engineering Support

RO Associates
246 Caspian Drive
Sunnyvale, CA 94088

Tel: (800) 443-1450
Fax: (408) 744-1521
E-Mail: sales@roassoc.com

consumption is offset by its higher performance.

Interestingly, where performance per cubic foot is more important, it is the processing-per-watt metric of the processing element that most effectively helps solve the problem. Therefore, FPGAs, Cell BEs and GPUs are the most suitable in space-constrained environments such as ground-based vehicles.

Application Ranking	0	0	5	0	0	3	4	2	2	3	3	4	3	
	Resolution				Code Type			Air/Cond. cooled MIL				Applicability		
Processing Technology	Bit opns, e.g. image morph	8/16 bit opns	Float opns	Double Precn opns	Vector (SIMD)	Scalar	Condnl GPP C code	Perf./ cu. ft	Perf./ watt	Min. abs. watt/ chip	Perf ./lb	Ease of use	Range of appns	Final Score
To avoid double counting, only one rating for each pair of correlated metrics, shown in in shaded cells, is considered in the final score.														

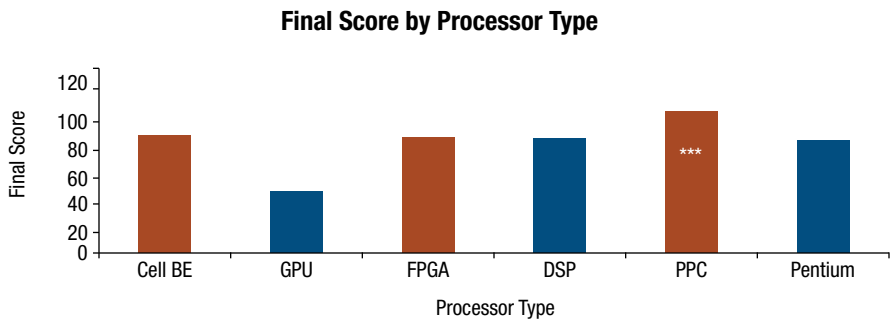

Figure 3
Processor Decision Model, PPC High Score example.

The fourth criterion, applicability, encompasses both the range of applications that can be supported on the processor and the processor's ease of use. It compares these parameters against the needs of those applications for which each processing element is most suitable. The ease-of-use parameter does not consider any specialized skillsets an organization may possess.

In this category, PPCs and Pentiums rank highest, since design engineers in nearly all organizations use at least one of these processors. In contrast, there are significant challenges in leveraging GPUs for military applications. The most significant are device drivers for non-x86 hosts and operating systems that are not Windows or mainstream Linux implementations on the x86. FPGAs are becoming more widely applicable here, while DSPs remain chiefly specific to application type. The Cell BE holds promise for a reasonably wide range of processing requirements where a lot of DSP and some general-purpose processing are required.

Processor Tradeoffs

Using the model, subjective capability ratings are assigned to each processor type for each decision parameter. In any given category, these ratings are scored from a low of 1 to a high of 6. Application characteristics are ranked from lowest to highest importance using a score of 0 to 5. Although individual model values are debatable, the aggregate result, or final score, proves useful in determining the optimal processor choice for the application, or part thereof, being assessed.

The final score is computed as the sum

of the individual parameters weighted by the application rankings. For example, the final score of 145 for the Cell BE High Score example (Figure 1) was derived as shown: {(0*1) + (0*4) + (5*6) + (5*6) + (4*6) + (3*2) + (2*2) + (5*5) + (4*3) + (2*4) + (2*3) = 145}.

Although final scores can be presented graphically, it is the ranking order from highest to lowest across all of the processor types that proves most useful to the difficult problem of prioritizing processor choices. Relative scores are less meaningful because they are not used in the assessment of each processor type against the specific criteria. When implemented as an Excel spreadsheet, the rank ordering of processors adjusts automatically to a score that indicates the order needed to consider each different processor type as a system solution to the application problem.

In the Cell BE High Score example (Figure 1), the specific application description exemplified in its high processing score contains significant floating-point and double-precision operations, indicative of radar or beam-forming applications. Moreover, the application contains a combination of vector, scalar and conditional "C" programming components. This particular application also requires strong performance per cubic foot as well as per pound, typically required for military airborne deployed systems. The applicability criterion is of limited importance in this case, which is to be expected of mission-specific systems.

Use of the model illustrates that the Cell BE, closely followed by PPCs and FPGAs, is most suitable for computationally intensive, high-precision, space-constrained, deployed applications.

In the second FPGA High Score example (Figure 2), a bit operation application with vector arithmetic favors an FPGA implementation, closely followed by the PPC. Here, size, power and weight are of critical importance. This application could be an image formation processing system that is co-located at the sensor in a manned or unmanned airborne platform. As sensor data continues to overwhelm satellite data links, the ability to process in real time close to the sensor is critical.

A very different set of issues is posed by a floating-point application that incorporates both scalar and conditional "C" code. In this PPC High Score example (Figure 3), size, weight and power management are of moderate concern, but applicability is important. In this case, the model suggests that the PPC be considered first.

This could be a target-recognition, tracking, or other rule-based system that requires some flexibility. Ideally, such post-processing applications would co-exist with the sensor and image processor, as hosting algorithms on deployable platforms in lieu of in-ground stations would further empower warfighters with real-time intelligence. The integration of new algorithms with embedded high-performance computers is the key enabler for this real-time capability. ▮▮

Mercury Computer Systems, Inc.
Chelmsford, MA.
(978) 256-1300.
[www.mc.com].

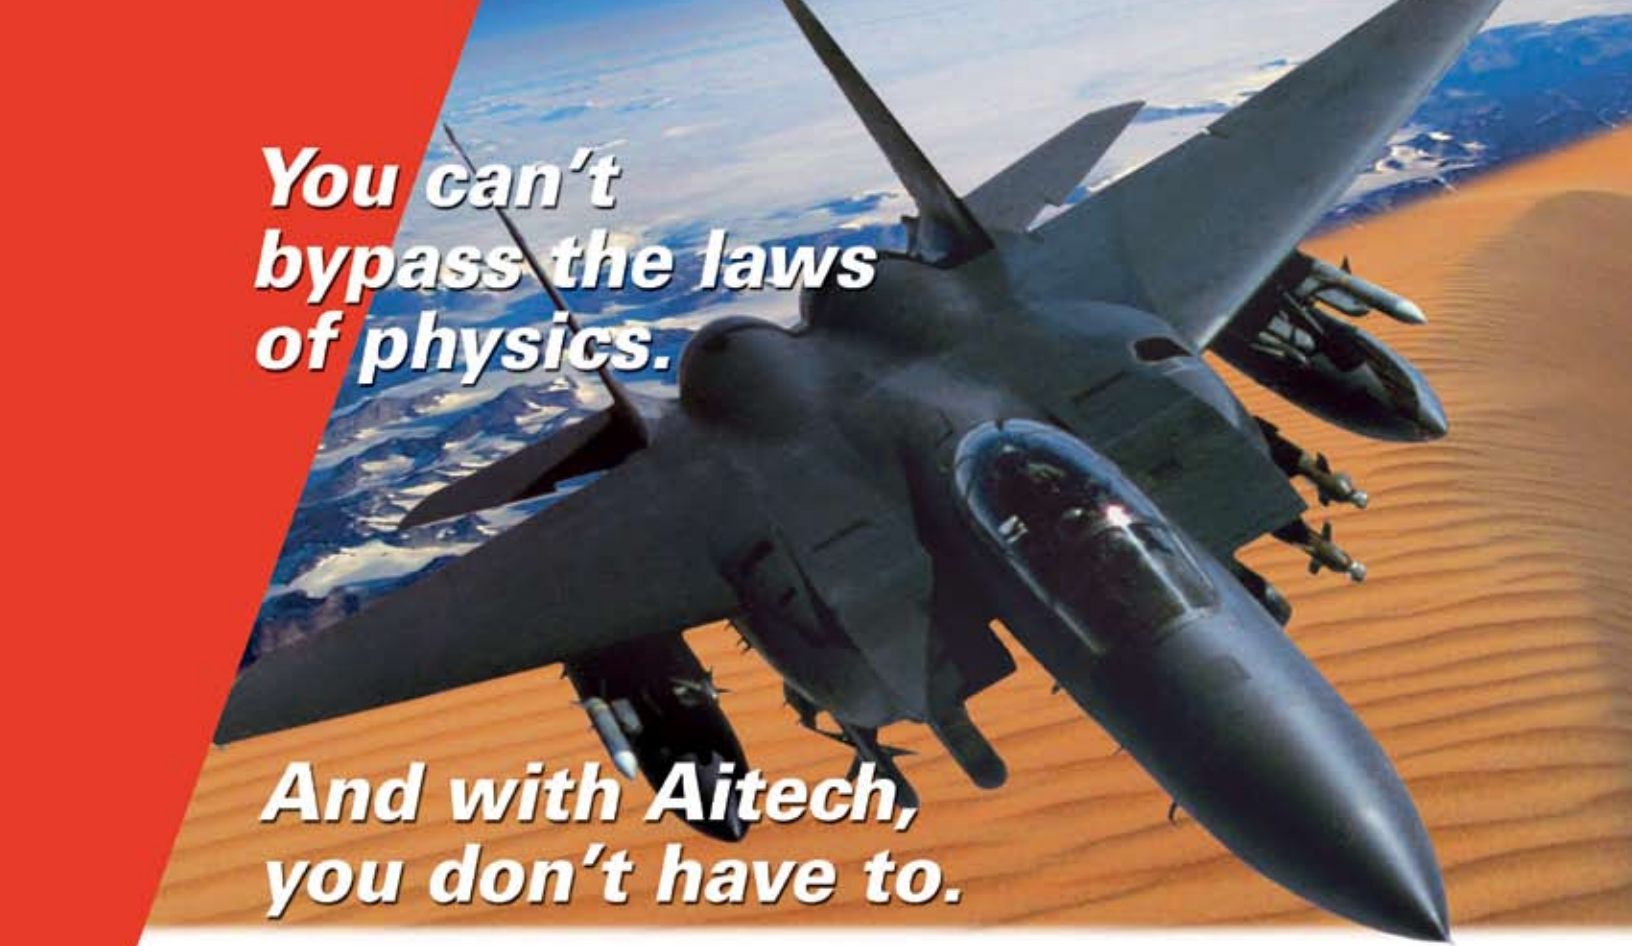

True -55°C to +85°C operation... Standard!

Aitech continues to provide industry standard open systems architectures such as VMEbus and CompactPCI products, designed and tested to -55°C to +85°C as standard, because you can't bypass the rules of engagement either.

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 there... Aitech subsystems have been proven in the world's most demanding mission-critical mil/aero applications – from complex ground, air and sea platforms to rad-tolerant solutions for the Space Shuttle, International Space Station and now earth orbiting satellites!

We've done that... 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.

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

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

Ready!

Design with Quartus® II development tools.

QUARTUS® II

Set!

Prototype with Stratix® II FPGAs.

Go to Market!

Ramp to production with Stratix II FPGAs or migrate seamlessly to HardCopy® II structured ASICs.

The quickest & easiest path to production.

ALTERA®

The Programmable Solutions Company®

www.altera.com

FPGAs: The New Matrix for Design

PCI-to-PMC

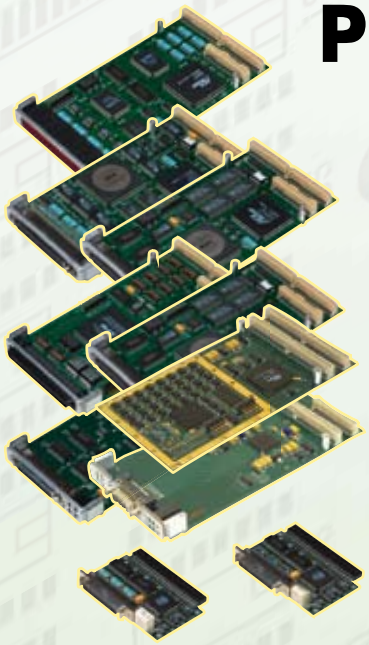

Serial

Ultra320

FPGA

Quad Ethernet

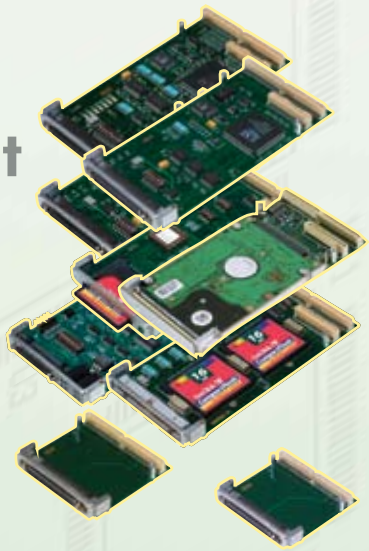

Adapters

SCSI

Async

SATA

Reconfigurable

RoHS

CardBus

IDE

Digital I/O

CompactFlash

Fast/Wide

PrPMC

SAS

PMC-to-PCI

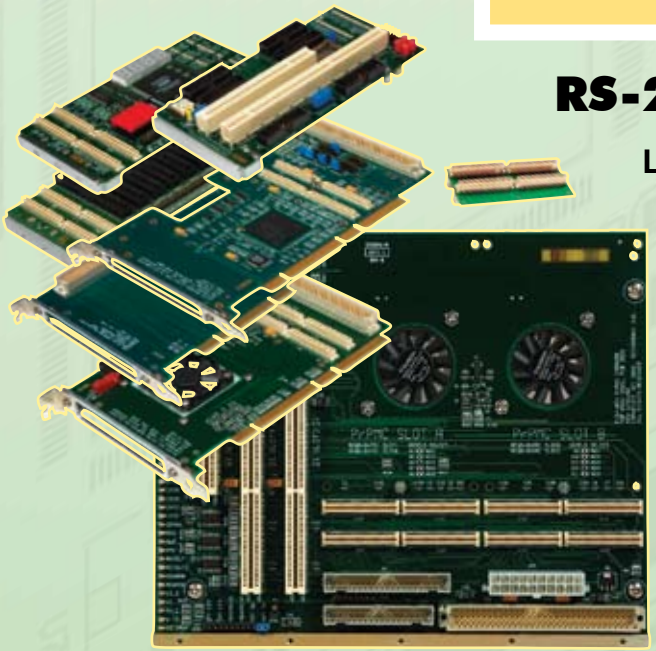

RS-232/422/485

LVD/SE

PIM

PMC-X

Test Access & Bus Analysis

Carriers

Extenders

Technobox produces **custom** and **off-the-shelf PMCs**. Our extensive range of solutions includes versatile FPGA-based digital I/O, async I/O, SATA, Ethernet, and IDE controllers. Plus a variety of **adapters, extenders, and tools** for PMC development or integration.

Make us your **first stop** for PMC Solutions.

www.technobox.com

Your Source for PMCs, Adapters, and Tools.

Technobox, inc.®

140 Mount Holly Bypass – Unit 1
Lumberton, NJ 08048
Tel. 609-267-8988
Fax 609-261-1011

Changing Horses in Midstream: Partial Reconfiguration for FPGA Designs

The ability to leverage partial reconfiguration for programmable logic opens new doors to a whole host of applications such as software defined radio, dynamic instruction set computing and automatic target recognition.

by Mark Goosman
Xilinx

Ongoing trends in logic design such as shorter product lifecycles, greater design complexity and increased scaling (e.g., "Moore's Law") are adding increasing pressure for faster design speed and lower power in a smaller physical space. Although today's advanced FPGAs are rapidly evolving to address these issues of speed, power and size, new technologies in the area of partial reconfiguration offer the promise of even greater advances.

Partial reconfiguration is a design process that allows a limited, predefined portion of an FPGA to be reconfigured while the remainder of the device continues to operate. This is especially valuable where devices operate in a mission-critical environment and cannot be disrupted while subsystems are redefined. The ability to partially reconfigure a device takes the already powerful benefits of reprogrammability to a much higher level.

The obvious benefit of reconfigurable devices, such as FPGAs, is that the functionality with which a device is configured can be changed and updated at some time in the future. As additional functionality is available or design improvements are made available, the FPGA can be shut down, completely reprogrammed with new logic and operations can be resumed. Partial reconfigurability addresses the environment where logic needs to be changed or updated within a part of an FPGA without disrupting the entire system. This may be a design comprised of several blocks of logic and, without disrupting the system and stopping the flow of data, requires an update of the functionality within one block.

Using partial reconfiguration, designers can dramatically increase the functionality of a single FPGA, allowing for fewer, smaller devices than would otherwise be needed. This allows for additional functionality, lower power, reduced cost and less physical space on the board.

Partial reconfiguration is useful for systems with multiple functions that can time-share the same FPGA device resources. In such systems, one section of the FPGA continues to operate while other sections of the FPGA are disabled and reconfigured to provide new functionality. This allows concurrent support for multiple independent applications in a single FPGA. This is somewhat analogous to dynamic task switching or multitasking of a general-purpose processor. Without this capability, it would be necessary to reconfigure the entire FPGA to support a different application, which would result in the loss of all previous applications.

Partial reconfiguration provides an advantage over multiple full bit streams in applications that require continuous operation, which is not otherwise accessible during full reconfiguration. One example is a graphics display that utilizes horizontal and vertical synchronization. Because of the environment in which this application operates, signals from radio and video links need to be preserved—but the format and data processing format may require updates and changes during operation. With partial reconfiguration, the system can maintain these real-time links while other modules within the FPGA are changed on-the-fly.

In order to implement partial reconfiguration on an FPGA, it first requires an FPGA that inherently supports the dynamic

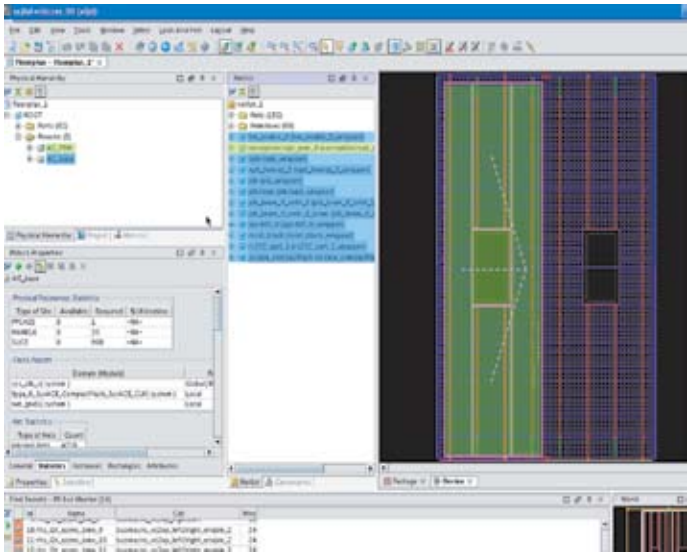

Figure 1 All static logic is grouped in a single Pblock AG_base.

reconfiguration of only portions of the device, while leaving the other portions unaffected. Then a set of software development tools are needed that support the development of applications restricted to boundaries that comply with the hardware architecture of the FPGA. Finally, some form of basic controller must be available to dynamically manage the reconfiguration of the FPGA. This could be an embedded general-purpose processor (GPP), a soft core GPP, or an external GPP connected to the FPGA. In this shared resources model, the same embedded GPP that is running the design infrastructure and operating environment is also managing the partial reconfiguration of the FPGAs.

In an FPGA, all user-programmable features are controlled by memory cells that are volatile and must be configured on power-up. These memory cells are known as the configuration memory, and define the look-up table (LUT) equations, signal routing, input/output block (IOB) voltage standards and all other aspects of the design.

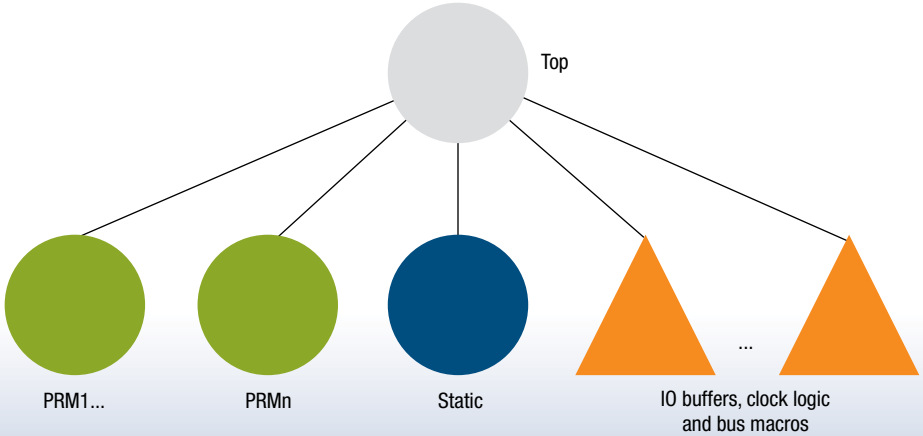

Figure 2 The hierarchy for a partial reconfiguration design.

To program configuration memory, instructions for the configuration control logic and data for the configuration memory are provided in the form of a bitstream, which is delivered to the device through the JTAG, SelectMAP, serial, or ICAP configuration interface.

Typically, a user performs the initial programming by downloading an entire bitstream to an inactive target device. Using partial reconfiguration, a subset of the FPGA can be reprogrammed using a partial bitstream. You can use partial bitstream to change the structure of one part of an FPGA design as the rest of the active device continues to operate.

A Methodology for Partial Reconfiguration

Successful implementation of a design using a partially reconfigurable flow requires following a strict design methodology. A reconfigurable design will consist of partially reconfigurable modules (PRMs) that will be swapped in and out of the FPGA and the static logic, which will remain in place. The general picture of the design flow involves the need to insert bus macros between the PRMs and the rest of the design, the static or fixed logic that remains in place. Bus macros are the channels or ports through which modules communicate and pass data. This allows a fixed communication channel for the static logic regardless of the reconfigurable logic on the other side.

Successful design requires following the guidelines of the synthesis tools to generate a partially reconfigurable net list. The synthesis tool must be configured so that no optimizations occur across hierarchical boundaries. This is generally done with a KEEP_HIERARCHY or similar directive.

The next step is to use the tool to floorplan the PRMs and cluster all static modules together and then place the bus macros between the PRMs and the static logic following PRM-specific design rules. Finally, run the partial reconfiguration implementation flow

PlanAhead 8.1 from Xilinx is an example of a single environment (or platform) used to manage the preceding guidelines, which can be broken down into the following steps:

1. Net list import
2. Floorplanning the design for partial reconfiguration
3. Design rule checks
4. Net list export
5. Implementation flow management
6. Bitstream size estimation

Although these steps are straightforward, the methodology requires meticulous implementation in order to ensure success. Changing out a portion of a complex, high-speed design does not allow much margin for error.

Use a tool like PlanAhead that works with imported net lists, such as those

from XST or Synplify, to import any hierarchical net list (single edf/ngc or multiple edf/ngc files). Then follow the regular guidelines to import the design into the tool and create a floorplan as you would with any non-partially reconfigurable design.

Floorplaning for partial reconfiguration is an important step in the partial reconfiguration flow. Floorplanning is based on design partitions referred to as physical blocks, or Pblocks. A Pblock can have an area (such as a rectangle) defined on the FPGA device to constrain the logic. The designer can define Pblocks without rectangles and the implementation software will attempt to group the logic during placement. Net list logic placed inside of Pblocks will receive AREA_GROUP constraints.

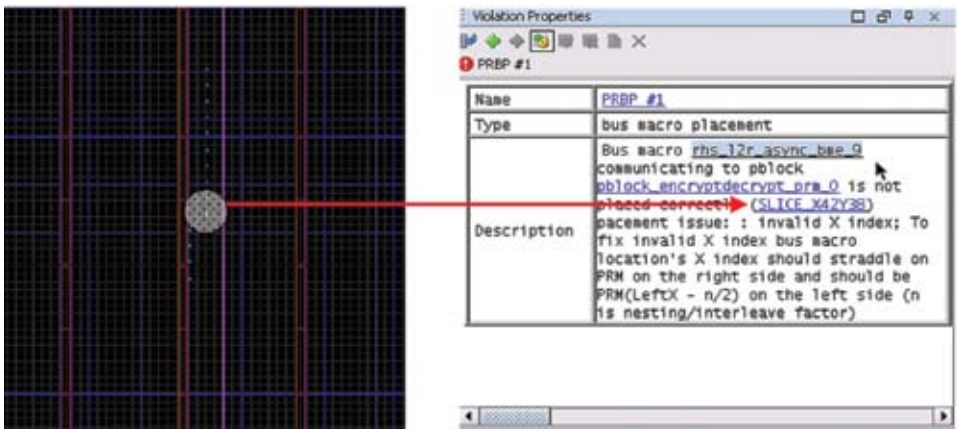

Figure 3 The PRBP DRC verifies all rules that should be followed for bus macro placement.

Floorplanning for partial reconfiguration entails several key subtasks. The first subtask is to assign an area for the PRM by creating a Pblock with an area defined within the fabric. This includes assigning the values for RANGES for the Pblock. The MODE constraint must be defined for all reconfigurable regions (MODE=RECONFIG). This constraint prevents the implementation tools from failing with unexpanded block errors during implementation of the static and reconfigurable modules.

Every top-level module, other than PRMs, should be grouped together in a single Pblock. This is called a static logic block. This block should not have a RANGE defined; this will cluster the static logic together in a single Pblock. Select all top-level modules (except the PRMs) and assign them to a Pblock. Figure 1 shows the static logic grouped in a Pblock named AG_base. When the floorplanning is completed in the design tools, the resulting physical hierarchy will be organized as shown in Figure 2.

The next step is to place the bus macros. Bus macros are physical ports that connect a PRM to static logic. Any connection from a PRM to static logic should always go through a bus macro. Bus macros are instantiated as black boxes in RTL and are filled with a predefined routing macro in the form of an .nmc file. Bus macros are placed on the PRM boundary. Static logic connected to PRMs will migrate toward the bus macro during placement.

Given the complexity of the flow, it is very common for mistakes to be introduced in the original RTL and during the floorplanning process. Any tool worth its salt will check for design violations. Also integrated into this feature, in the case of PlanAhead is the PR-Advisor, which provides feedback on how to improve your design. There are a number of design rule checks that are specific to Partial Reconfiguration.

The bus macro DRC provides verification for all design rules related to bus macro connectivity and placement. One example

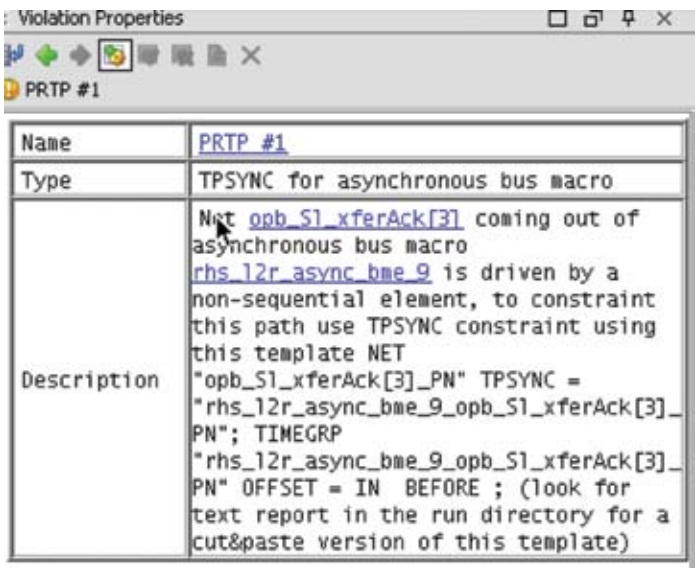

Figure 4 Special consideration should be given to timing-critical paths, which include an asynchronous path.

of a bus macro DRC is the PRBP check. This DRC checks for all rules that should be followed for bus macro placement. Figure 3 shows an example of a design that failed the PRBP DRC. In this case, the tool tells us that the interleaved/nested macro should be placed at SLICE_X41Y.

The Floorplanning DRC covers floorplanning rules. Clock objects (global clock buffers, DCM) and I/Os should be placed and static logic clustered. The glitching logic DRC verifies glitching logic elements (SRL and distributed RAM) above and below PRM regions.

Another DRC is the timing advisor/DRC. This provides a check for timing-related issues. One example of timing DRC is the PRTB check. With the PRTB check, the static module is implemented before the PRM during the implementation phase. Regular timing constraints do not cover the paths that cross between the static and a PRM module. This does not

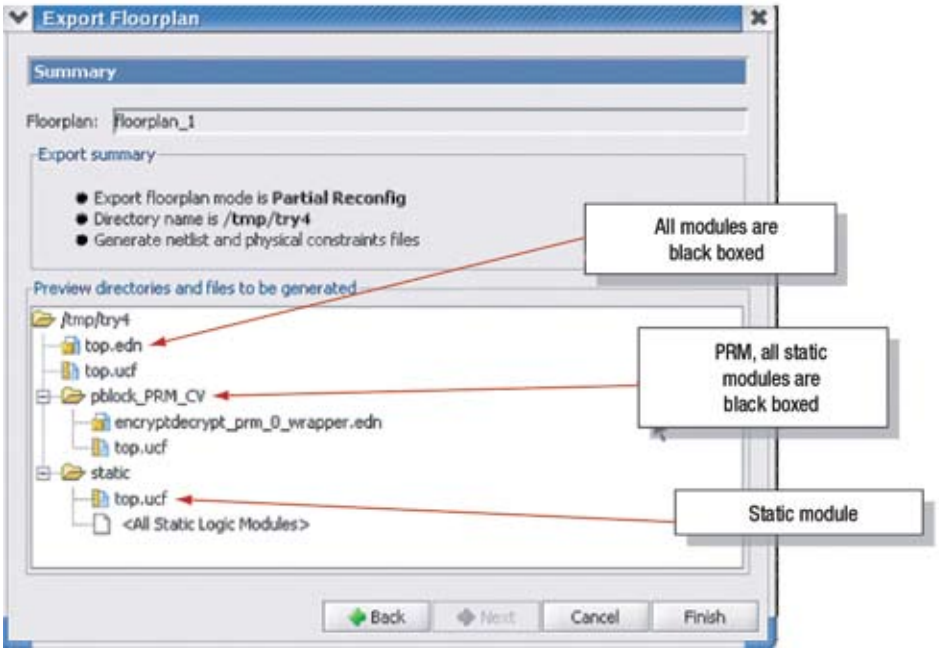

Figure 5 Upon successful completion of the floorplanning DRC processes, all modules will be displayed in black text within the Export Floorplan dialog box.

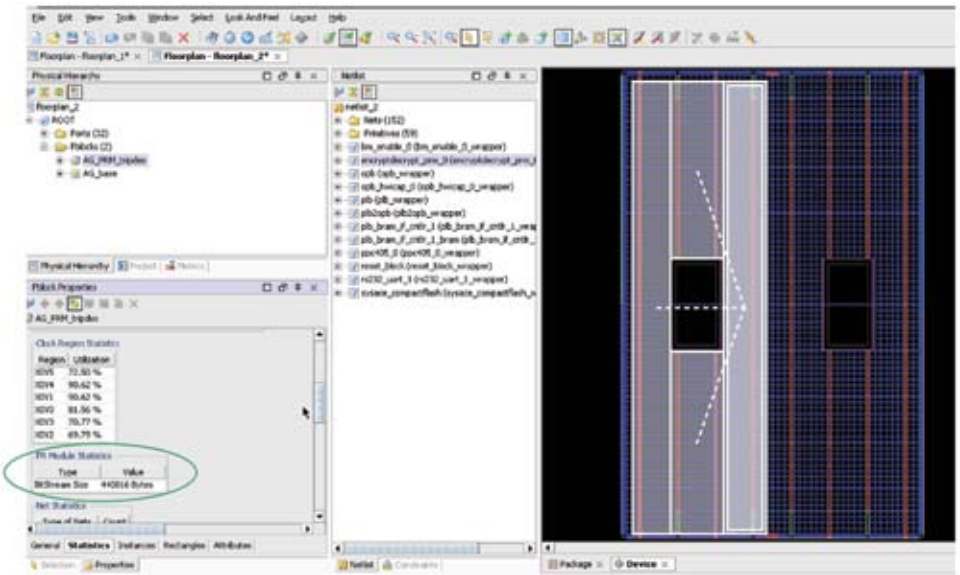

Figure 6 The display of Pblock properties includes the estimated size of the bitstream.

present a problem, provided that the bus macro is synchronous. However, if it is an asynchronous bus macro, the static module does not know about the propagating of asynchronous paths, as shown in the example in Figure 4. This could be important if these paths are timing-critical. One way to pass this information to the static module is to specify a TPSYNC constraint on the bus macro output net. PlanAhead software will recommend a TPSYNC constraint that can be added to the .UCF file.

Once the design is floorplanned and passes the DRC checker, it is ready to be exported. The design tools should take care of exporting the original hierarchical net list into a PR-style net list that has a specific format (static and PRM in separate directories). The export directory will appear as shown in Figure 5. Next, a partial reconfiguration flow wizard, shown in Figure 6, runs the partial reconfiguration implementation on the exported design. It will produce a full bitstream for the complete design and a partial bitstream for each of the PRMs. The implementation steps are:

- Initial budgeting
- Static module implementation
- PRM module implementation (one implementation for each version of every PRM)
- Assembly and bitstream generation (results are stored in the merge directory)

The Pblock statistics report includes a section that reports PRM bitstream size (Figure 6). This information can be used for estimating the size of configuration memory storage such as external flash and DDR. This information can also be used to calculate how long it will take to swap the module based on your bitstream memory interface.

Partial reconfiguration offers a tremendous opportunity for designers looking for a way to increase the functionality of their design, achieve lower power and reduce the number of devices on their board. Using new design tools and techniques becoming available for partial reconfiguration applications can greatly simplify the complexities of juggling the dynamic operating environment of these cutting-edge applications, allowing a single device to operate in applications that previously required multiple FPGAs along with the required power, board space and design overhead. ■

Xilinx
San Jose, CA.
(408) 559-7778.
[www.xilinx.com].

Acromag I/O. Customize It.

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 FPGA with 500K, 2M or 3M system gates
- 1Mb on-chip RAM, 9Mb on-board SRAM
- Fast PCI with 32-bit, 66MHz dual DMA

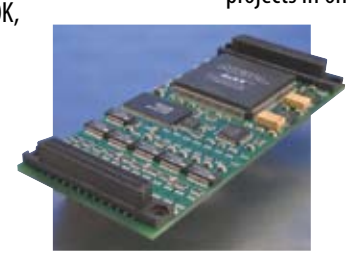

Industry Pack FPGA I/O also available

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

- Digital I/O: TTL, RS422, or LVDS I/O
- Analog I/O: four 20 or 65MHz A/D and two 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, PCI, PMC, and Industry Pack solutions.

Acromag®
THE LEADER IN INDUSTRIAL I/O

www.acromag.com

800-881-0268 or 248-624-1541

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

Manufactured in Wixom, Michigan, USA

Reduce Cost, Risk and Time-to-Market with an FPGA-to-Structured ASIC Strategy

The costs, development times and risks of ASIC designs have become all but prohibitive. Finalizing a design on an FPGA and moving it to a structured ASIC can cut cost, time and risk as well as result in smaller size and lower power consumption.

by Danny Biran
Altera

Few would argue that ASICs and ASSPs have become viable vehicles for only a few silicon vendors. As process nodes continue to shrink—to 90 nm now and to 65 nm in the near future—not many vendors are willing to take on the risk associated with an ASIC design unless they have a very high level of confidence that their chip will be sold in huge volumes to justify an investment that is often tens of millions of dollars (Figure 1). ASICs are under pressure. It is increasingly difficult for companies to justify the development of a new chip with the very high NRE price tags. If chip development cost is $30M, R&D costs are 20% of revenue, and one expects 10% market share, only a $1.5 billion market opportunity can justify the expense. There aren't many markets of this size that can be serviced with a single product.

One of the risks associated with building an ASIC is that its functionality is fixed during fabrication, resulting in the need to "get it right the first time," which can often be hard to achieve. The cost of redesign for even a small portion of an ASIC that doesn't work may be prohibitive in terms of both time and money. To mitigate this risk, FPGAs are often used as prototyping vehicles. With their high levels of logic density and performance, high-density FPGAs help developers implement complex designs that can be tested in the target system. FPGA prototypes offer many distinct advantages, allowing a design to have its hardware, application software and firmware fully developed and tested at full speed without a multi-million-dollar price tag.

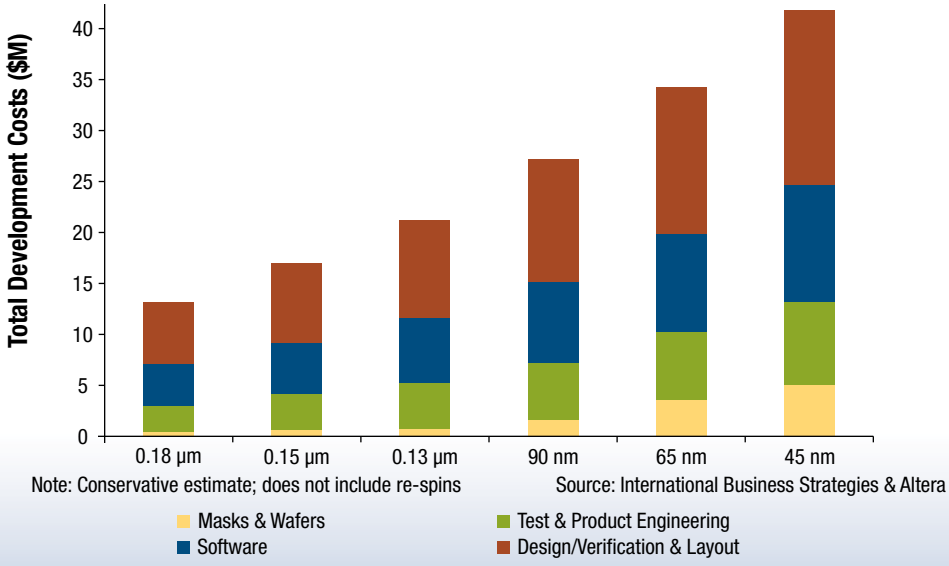

Figure 1 The cost of chip development goes up dramatically as technology nodes shrink.

Moving to a Structured ASIC

Once the design is prototyped in the FPGA, the next step is to move to an effective silicon platform. Migrating from an FPGA to an ASIC is tricky and time-consuming, since there is generally a good

deal of work required to make sure the original design will still function correctly when implemented as an ASIC.

A better solution is to move the verified FPGA design onto a structured ASIC, which has the characteristics of an FPGA and most attributes of an ASIC, and is significantly smaller than an FPGA. A structured ASIC has near-ASIC performance and power consumption, along with a shorter development cycle and lower NRE than an ASIC. However, it has a much lower unit cost than a comparable FPGA—as much as one-tenth the cost.

Structured ASICs are comprised of prefabricated base arrays with predefined and verified logic, memory, clock networks and I/O resources. The arrays are processed through manufacturing up to a certain point and then "banked" for future customization. Application-specific designs are then configured onto the base arrays using a few top metal layers, thus creating the structured ASIC. This gives a structured ASIC a post-fabrication configurability dimension that is not available with an ASIC. According to Gartner Group, structured ASICs are rapidly gaining in popularity, with revenues projected to grow from $99 million in 2004 to $848 million by the end of 2007.

Designers should be aware that different structured ASIC vendors implement their architectures in very different ways, even with different families from the same vendor. Parameters such as number of user-definable layers, number of equivalent ASIC gates, memory capacity, maximum clock rate and target processes are very different from vendor to vendor, making the migration from FPGA of one vendor to structured ASIC of another, while easier than to an ASIC, still difficult and fraught with potential errors that can add time and cost to reaching production silicon.

The Key for Structured ASIC Success

If the same silicon vendor produces both silicon platforms, moving a design from a system-verified FPGA to a structured ASIC, while maintaining functionality and meeting timing constraints, is greatly simplified and provides the customer with a low-cost and safe "path to production" strategy. A one-to-one mapping between the FPGA and structured ASIC can eliminate redesigning the board, and redeveloping and revalidating the system, resulting in significant development cost savings and time-to-market benefits. Proven IP cores, common to both platforms, a single design flow, common EDA tools and pin-to-pin compatibility between the FPGA and the structured ASIC assure that the smaller, less expensive structured ASIC will work correctly in the application.

The costs associated with a structured ASIC are much lower than those for an ASIC. For a 90-nm design with, say, 2.2 million ASIC gates, 9 Mbits of SRAM and 1.5 million gates for DSP and multipliers, NREs are often as low as $225,000 to $300,000, which is significantly lower than just the cost of a complete mask set for an ASIC. Designs prototyped and verified in an FPGA and then migrated to a structured ASIC can be used for all but the very highest performance and/or cost-sensitive applications that would justify the cost optimization realized by a full ASIC.

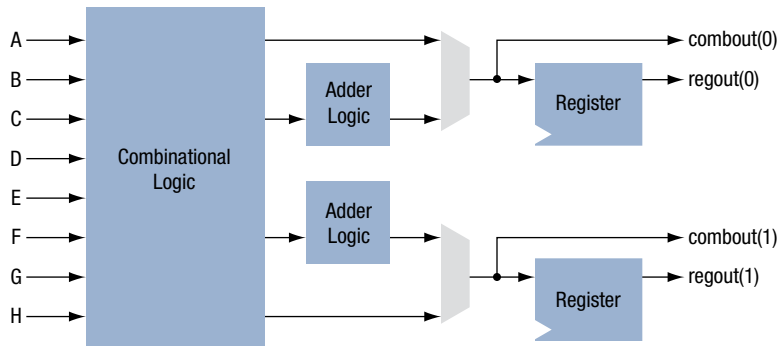

Figure 2 This Adaptive Logic Module (ALM) is the basic logic block of the Stratix II FPGA. Using ALMs instead of traditional 4-input look-up tables (LUTs) increases logic-utilization efficiency and performance.

When migrating from an FPGA to a structured ASIC, the vendor can eliminate a lot of the circuitry from the FPGA that is not required for normal chip operation. The removed circuitry includes FPGA configuration logic, programmable routing and logic and memory programmability. What needs to be added to the structured ASIC is embedded testability, since the circuitry needed to test the structured ASIC is much different than that for the FPGA. Eliminating all the extra transistors results in a much smaller chip with a corresponding unit cost reduction of as much as 90%. The structured ASIC also sees a significant power reduction. The shorter interconnect paths result in lower dynamic power dissipation and the structured ASIC has lower static power consumption with the elimination of the many transistors used for configuration and programming on the FPGA.

A key to successful FPGA and structured ASIC migration, even when both silicon platforms are from the same vendor, is how well the basic logic building blocks of the structured ASIC implement the logical functions of the FPGA. This is not a simple task, and requires a lot of work on the part of the vendor who has developed both platforms.

For example, the logic structure of the Stratix II FPGA architecture comprises basic logic units known as adaptive logic modules (ALMs). As shown in Figure 2, each ALM contains a variety of look-up table (LUT)-based resources, two full adders, carry-chain segments, two flip-flops and many additional logic enhancements that can be divided into two adaptive LUTs (ALUTs). One ALM can implement logic functions with up to seven inputs and complex logic-arithmetic functions, increasing logic efficiency and reducing routing resources.

The HardCopy II structured ASIC family comprises an array of fine-grained structured cells called HCells that are grouped

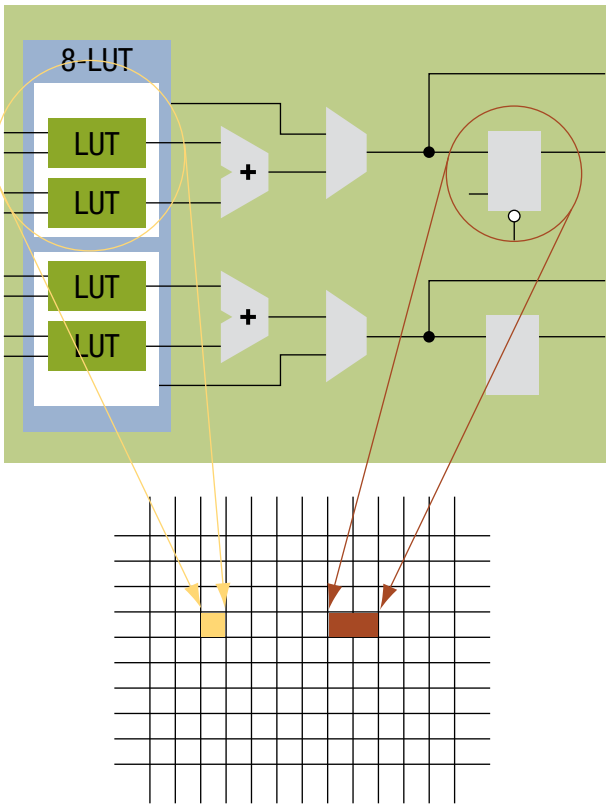

Figure 3 If sections of an ALM are not used in the FPGA design, then they are not mapped to the structured ASIC. This increases the efficiency of the FPGA-to-structured ASIC operation.

into HCell macros to implement a portion of an ALM or a section of a DSP block. The design software maintains a library containing a pre-verified, pre-characterized HCell macro for every ALM configuration, which then maps the ALMs into a structured ASIC design. The tool only maps the utilized portion of each ALM to HCell macros (Figure 3); if parts of an ALM are not used in the FPGA design, then they are not mapped to the HardCopy II device, yielding a more efficient mapping of the prototyped design.

When compared to a corresponding Stratix II device, a HardCopy II structured ASIC is 50% faster, dissipates up to 70% lower core power and has a 60-85% smaller chip size.

A HardCopy structured ASIC also offers time-to-market advantages. The typical time it takes an ASIC/ASSP to go from design initiation to completion is around two years. During that time, market opportunity for the chip can be severely hindered or it can even disappear. Taking a design to a structured ASIC typically takes half that time or less, allowing the product to reach the market significantly sooner than if it were implemented as an ASIC or ASSP.

The secret to making FPGA design an integral part of product development beyond using the FPGA only as a prototype device is to provide a clear and low-risk path from the FPGA to a production-viable silicon platform, such as a structured ASIC. With near-ASIC performance and cost, structured ASICs provide a production silicon solution for a broad range of applications that are currently filled by expensive and unreliable (in terms of cost and development schedule) ASICs and ASSPs. However, with so many different structured ASIC architectures and business models available, an important factor in successful FPGA-to-Structured ASIC migration is *re-mapping* the logic functions of one platform to the other as opposed to doing an expensive and risky architectural *conversion*. This is accomplished by working with a vendor who has developed both the FPGA and structured ASIC architectures, along with the EDA tools a designer uses to develop and verify the target design on the FPGA and then migrate the design to a structured ASIC. ■

Altera
San Jose, CA.
(408) 544-7000.
[www.altera.com].

Annapolis Micro Systems

The FPGA Systems Performance Leader!

SIGINT ELINT Digital Receivers FLIR
FOPEN Radar Systems Software Defined Radio

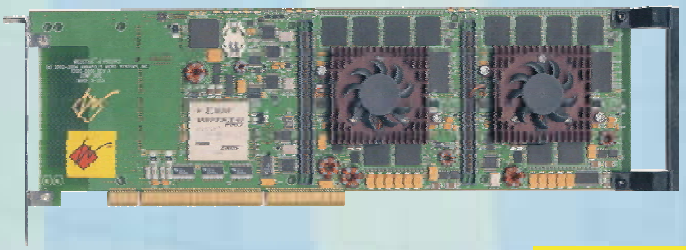

CoreFire™

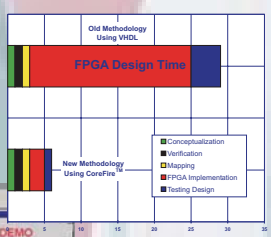

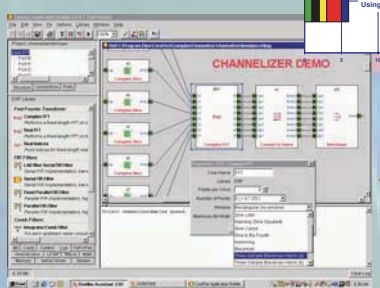

FPGA Application Builder

Annapolis Micro Systems

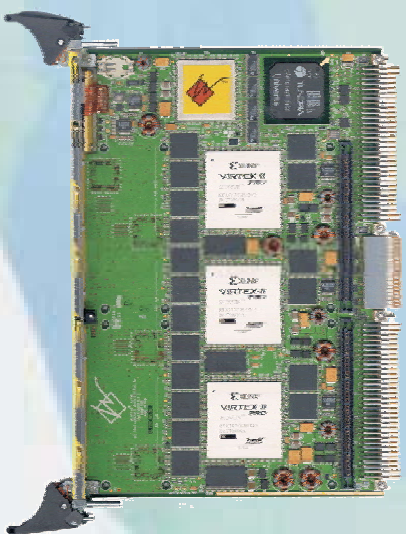

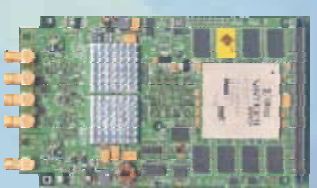

High Performance Signal Processing in Scalable FPGA Computing Fabric

Real Time, Real World.....REAL FAST!

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

Synplify® DSP

A Breakthrough in DSP Synthesis

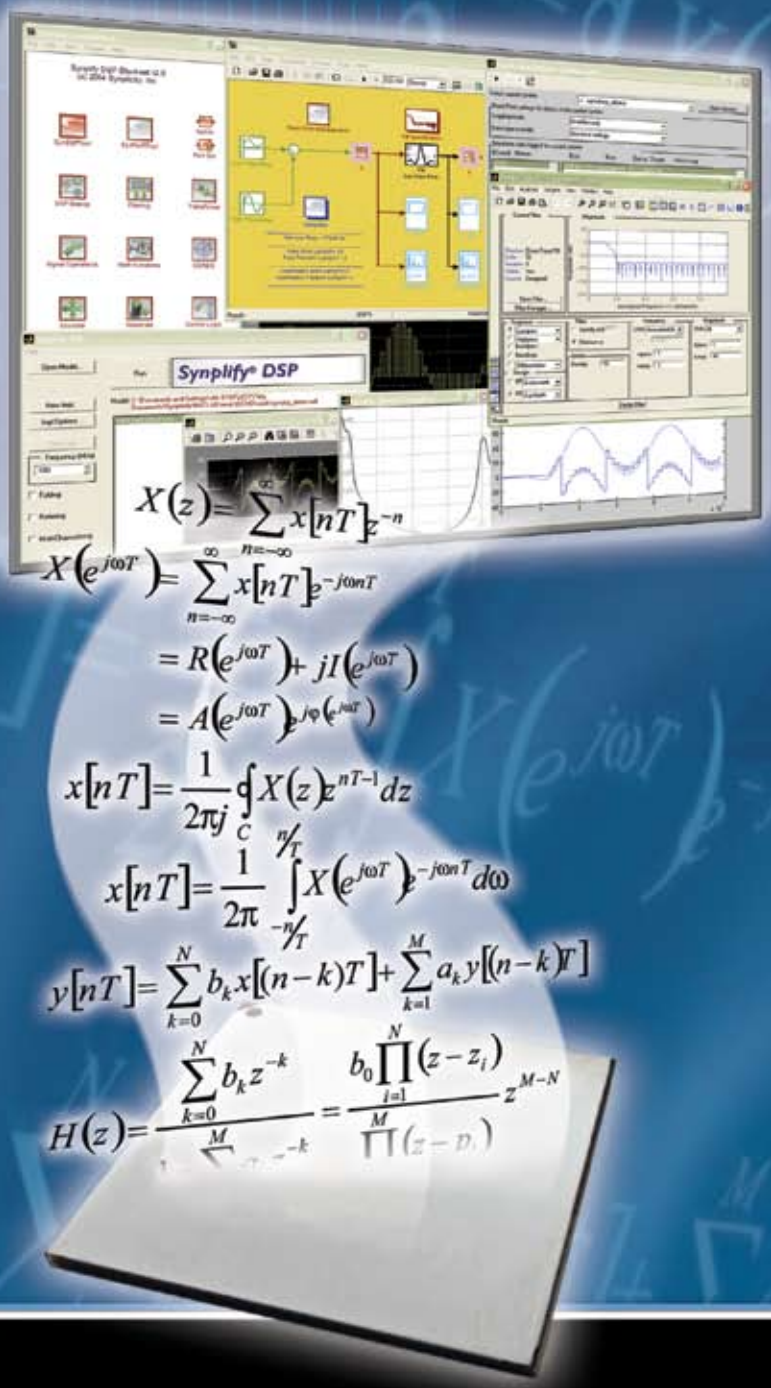

Synplicity's Synplify DSP synthesis solution offers DSP designers an efficient automated path from Simulink® to optimized RTL code. And, unlike other solutions, Synplify DSP software provides technology independent device targeting and advanced synthesis optimizations for performance and area (cost).

Synplify DSP Software Uniquely Offers:

- A User-extensible DSP library: quickly create and re-use custom functions
- Sample-rate-based synthesis engine: capture and explore algorithm behavior without worrying about target hardware
- Retiming: automatic pipelining and register insertion for higher speeds
- Folding: automatically apply resource sharing logic for lower area (device cost)
- Multichannelization: automatically apply resource sharing across multiple channels for lower area (device cost)

For more information on Synplicity's Synplify DSP solution and all of Synplicity's offerings, please visit our website at www.synplicity.com or contact info@synplicity.com

Synplicity®

Simply Better Results

Moving Real-Time Data around FPGA-Centric Systems

Support for real-time FPGA networks is emerging from board level vendors. As larger FPGAs and solutions that are more sophisticated are required, this will become increasingly important.

by Jeremy Banks
VMetro

Computationally intensive applications often involve real-time high-bandwidth data streams coupled with fast signal processing. This increasingly involves using multiple FPGAs because they are faster than CPUs or DSPs. However, coupling FPGAs (or IP cores) can be difficult and can consume unnecessarily large amounts of FPGA resources leaving fewer resources for the actual signal processing. Since the data movements can occur in parallel, the network communications can also be complex (Figure 1). The ability to establish links between FPGA IP cores, no matter where they are physically located, is important.

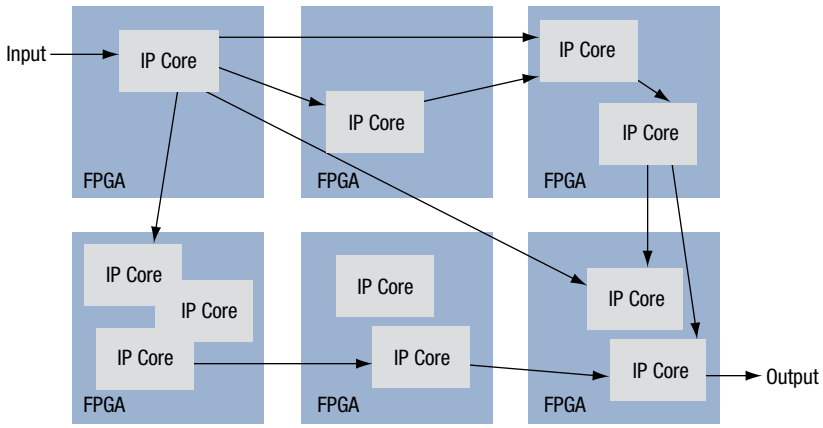

Figure 1 Data flow around networks of FPGAs can be complex.

The trend with CPU data communications is to move away from parallel bus structures in favor of multiple serial point-to-point data links. This makes systems easier to build and improves system performance because CPUs do not have to share a bus—instead there are dedicated communication paths that optimize the system communication. This also improves determinism by reducing the data traffic and the number of data sources on a link, an essential requisite for real-time solutions.

The most popular serial fabrics used for point-to-point communications include PCI Express, Serial RapidIO, Gigabit Ethernet and InfiniBand. High-speed serial communication (HSSC) using these fabrics is becoming the backbone of modular high-performance processing solutions through the adoption of standards such as VXS (VITA 41) and VPX (VITA 46) (Figure 2a and 2b). Boards and systems built using HSSC are providing high-density, tightly coupled FPGA and CPU solutions with high-bandwidth I/O. Such solutions are well suited to real-time applications.

Protocol-rich HSSC fabrics work well for CPU-centric systems, but for FPGAs, these fabrics are a luxury that sacrifices resources to implement complex communications. For systems that include both FPGAs and CPUs, using CPU-centric fabrics at the CPU-FPGA boundary makes sense. However, for direct FPGA-to-FPGA communications, an alternative is required in both the type of fabric and how it is used—after all, an FPGA is not a CPU, and treating it as such will dilute the true benefits offered by FPGAs.

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Data Flow

A common way to design a system in its early conceptual stages is to draw a block diagram. In doing this, the designer is describing the major processing blocks and how data moves through them. Memory mapping is a concept that CPUs use to arrange and process their data. Memory-mapped fabrics, such as PCI Express, fit well with CPU-centric processing models. Implementing the conceptual design onto a CPU requires "converting" the block diagram into the CPU's memory-mapped model.

By contrast, an FPGA can be used to lay out the block diagram "as is," parallel flows of data can be the same as the block diagram, as can the processing blocks. Along with increased performance, this difference of one-to-one mapping is a key advantage of FPGAs. For communication, a flow of data, or stream, is a better model to use with FPGAs than memory maps. How-

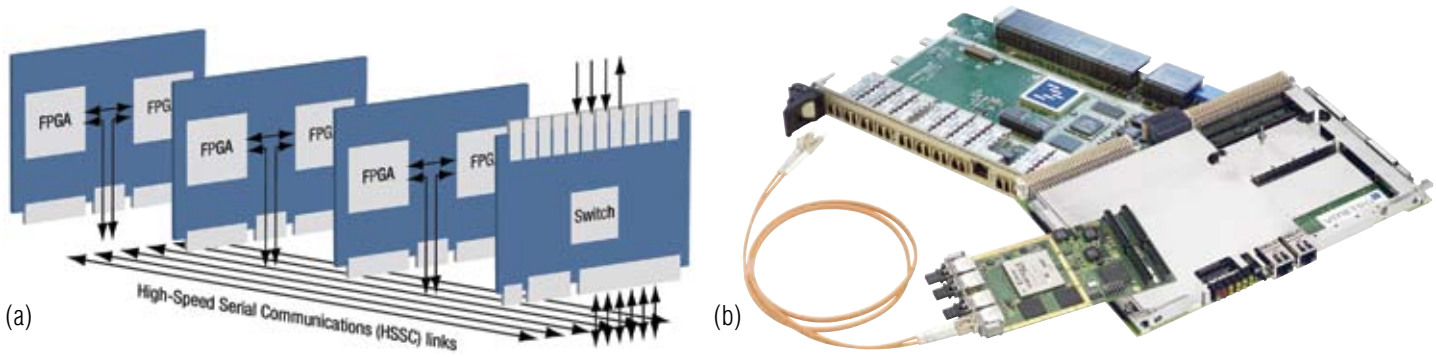

Figure 2 Example VXS board-based solution linking multiple FPGAs together over HSSC links (a), and examples of board-level components, a VXS switch, VXS FPGA and FPGA PMC with HSSC links (b).

ever, implementing the block diagram and its data flows on an FPGA presents some practical limitations. Even assuming all the processing blocks (or IP) are available, there is still the issue of whether all of this functionality can be fitted onto a single FPGA device. If not, can the parallel and independent data flows between processing blocks on different FPGAs be maintained and operate as if they are processing blocks on the same device?

Imposing CPU-style fabrics to manage the data flow would be a step backward because it would mean handling the data in different ways than those needed by the FPGA and would no longer represent a "flow." However, simple real-time communications networks based on data flows are becoming available. In their simplest form, these real-time networks need reliable data links and a tag to describe which processing block the data is to be delivered and, ideally, what the data represents. Dedicated real-time communications network IP cores can handle all of this work without having to implement memory-mapped fabrics.

Real-Time Communications Networks

Real-time systems are deterministic; for a given event, any event, a real-time system must respond within a guaranteed period. This is a common statement dictating what a real-time, deterministic system has to be able to achieve. When FPGAs are used for processing because of their speed, the ability of the network to deliver data within guaranteed constraints is even more important. Why use a device for performance if you cannot get data to it in the first place? For a system to be deterministic, the communications network must also be deterministic. This can be further complicated if there is a need to provide support for multiple data streams across a few links. This is a potential bottleneck and threat to real-time performance.

What makes a network deterministic? If the network can reliably deliver its complete payload within a guaranteed time period, no matter what is happening elsewhere on the network, it can be considered deterministic. For dedicated point-to-point connections, this is straightforward. However, what about packet switch connections where the data hops through many devices before arriving at its destination (Figure 3)? Or, what about situations where multiple data streams must share physical links?

Non-blocking switches are part of the answer, but there is no substitute for determinism by design, a higher-level system concept. One way to do this is to ensure that all data paths through the system are known (ideally fixed) with dedicated, allocated and guaranteed bandwidth—all the way through the system. If switches are used, they should be non-blocking on all data channels.

For properly designed real-time networks, the protocol can remain relatively simple and low level—ideal for FPGAs because fewer resources are used. Since data paths between nodes are known, features such as out of order data packet handling is unnecessary (as used by TCP/IP) as is flow control (other than moving data between clock domains using simple FIFOs). To a certain extent, even error correction (error detection is still important) can be simplified for real-time systems, otherwise, the question of what to do with the error is raised.

By definition, if the data has errors in the first place, then resending it is no guarantee that it will be correct the next time. If the data connection failed (errors occurred or no data at all was received), why shouldn't it fail repeatedly. This indeterminate situation is not good for a real-time system. The solution is usually a system design parameter of being able to live with the errors. For an imaging application, the processors may make the decision to throw away the data if it is useless, carry on and resynchronize. For FPGA-based designs, this simplification of the network protocol saves a large amount of FPGA resources.

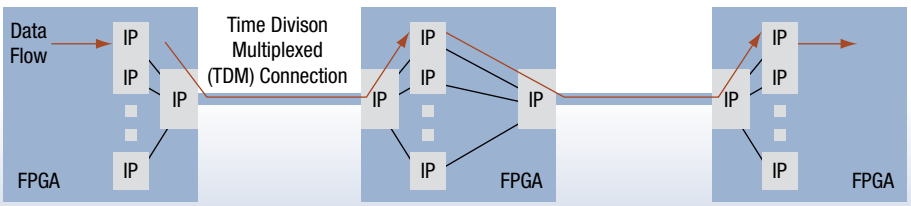

Figure 3 Transparent data flow needs to be maintained in real time–even if there is not a direct physical data link.

Efficient FPGA-Based Fabric Communications

Clusters of FPGAs that link streams of data between devices in real time need efficient point-to-point data links and protocols. Using protocols such as PCI Express consumes large amounts of FPGA resources. A PCI Express x4 core could account for as much as 30-40% of a Xilinx XC2VP50's resources. By contrast, a much simpler protocol, such as Serial FPDP uses as little

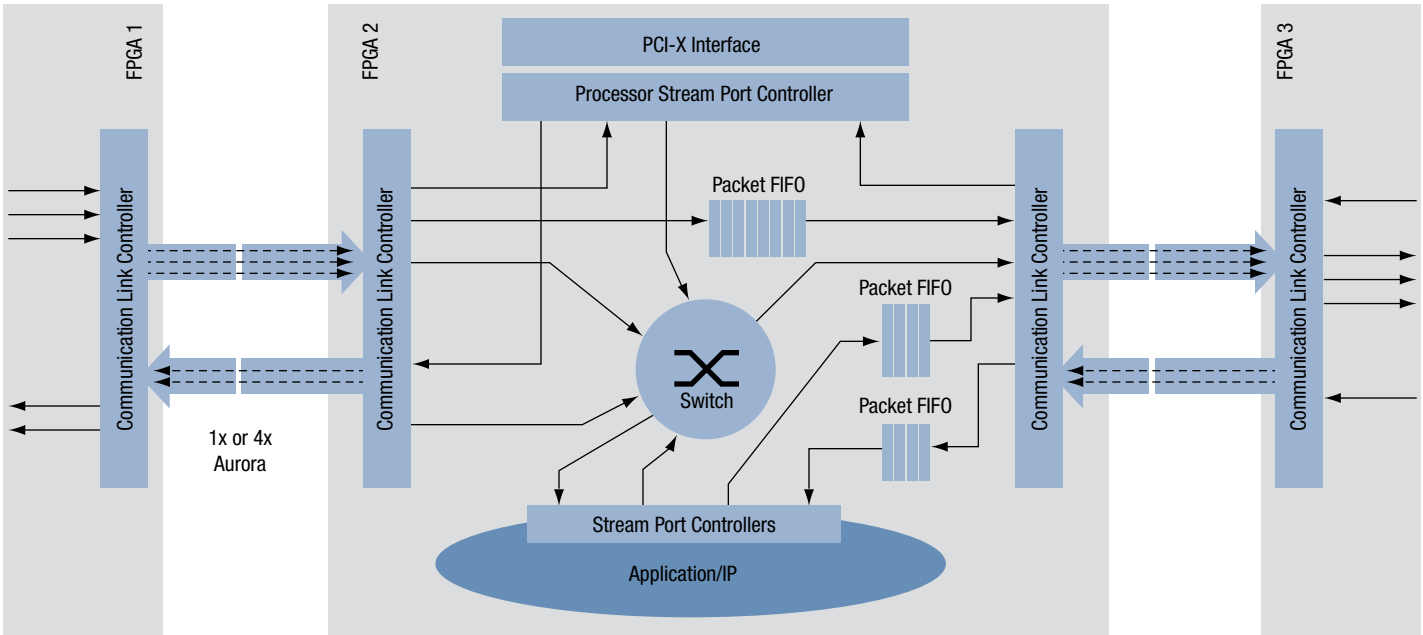

Figure 4 Example of FPGA-to-FPGA components as provided by VMetro's TransComm.

as 1% of the resources (for a x1 channel) of the same size FPGA. But Serial FPDP was developed as a simple sensor interface rather than a network protocol (Table 1). However, the Aurora protocol, developed by Xilinx, is optimized for FPGA-to-FPGA communications. By comparison, Aurora uses around 4% of an XC2VP50 FPGA; this is for a x4 channel (less for a 1x channel) and includes lane alignment, error detection and flow control—both user-defined and native.

Serial protocols using multi-Gbit/s data links are an efficient use of I/O pins for an FPGA. However, even the largest FPGA devices have no more than 20-30 full duplex serial links, or if used as x4 links for higher bandwidth, there are only four to six channels. If that is the case, what about large clusters of FPGAs with multiple data types and high-connectivity requirements? How are the limitations in the number of connections resolved? For these scenarios, the data links have to be shared with the different data flows between or through devices using a real-time network. In effect, the network must support bridging through an FPGA so that logical data paths can be established anywhere around the system.

The problem with the simple point-to-point protocols developed for FPGAs is that they do not handle bridging via an intermediary FPGA. Higher-level protocols are required that understand what data is targeted at them and if it is not, where it should go. If the data flow is self-identifying, then it should be straightforward for the real-time network to handle this. However, it is not something that the developer wants to deal with; there is an increased expectation that this is something the system vendor should be providing.

A Developer's Perspective

Handling FPGA communications in the development of an FPGA-centric design can be complex. While the performance advantages of FPGAs are well understood, implementing the network communications efficiently is critical for a successful real-time system. Ideally, a developer would like to have a toolkit of firmware IP and software components, just as they have for processing blocks.

Communications IP Core	Approximate resource usage of XC2VP50 FPGA
PCI Express x4	30-40%
Serial FPDP (x1)	1%
Aurora (x4)	4%

Table 1 FPGA Core usage for examples of communications protocols.

These components must be provided in such a way that only the components that are needed are included at compile time for optimal solutions, rather than generic code blocks that cater to all situations. This saves valuable FPGA resources by removing unnecessary IP cores for data channels, DMA controllers, etc.

The ideal components to do this include communication link controllers, which maintain the physical interfaces such as a high-speed serial communications link, or parallel LVDS ports with the ability to support virtual data streams, non-blocking switches and simple interfaces for IP to link into the fabric, etc. Such toolkits are now becoming available from companies such as VMetro with its TransComm firmware and software tools (Figure 4.). With such toolkits, creating a real-time communications fabric for FPGAs, perhaps included on analog input designs, becomes much easier with reduced risk.

Developers want to harness the power of FPGAs and focus on their own expertise, the application IP, not the network communications. Support for real-time FPGA networks is emerging from board-level vendors. As larger FPGAs and solutions that are more sophisticated are required, this will become increasingly important. □

VMetro
Houston, TX.
(281)584-0728.
[www.vmetro.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®				VIA Eden™				AMD Geode			486DX	
		CMX58886PX1400HR	CMX58886PX1100HR	CMX58886CX1000HR	CML47786CX650HR	CML47786HX650HR	CMX47786CX650HR	CMX47786HX650HR	CML36786HX1000HR	CML37786HX1000HR	CMX36786HX1000HR	CMX37786HX1000HR	CME26686HX333HR	CME27686HX333HR	CME27686CX333HR	CMV6486DX100HR	CMi6486DX100HR
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									
CPU and BIOS	CPU Max Clock Rate (MHz)	1400	1100	1000	650	650	650	650	1000	1000	1000	1000	333	333	333	100	100
	L2 Cache	2MB	2MB	512k	256k	256k	256k	256k	64k	64k	64k	64k	16K	16k	16k	16k	16k
	Intel SpeedStep Technology	✓	✓														
	ACPI Power Mgmt	2.0	2.0	2.0	1.0	1.0	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	512	256	256	256	32	32
	RTD Enhanced Flash BIOS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Nonvolatile Configuration	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Quick Boot Option Installed	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Fail Safe Boot ROM				✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	USB Boot	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
Peripherals	Watchdog Timer & RTC	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	IDE and Floppy Controllers	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓
	SSD Socket, 32 DIP					1		1	1	1	1	1	1	1		2	1
	ATA/IDE Disk Socket, 32 DIP	1	1	1	1		1								1		
	Audio	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
	Digital Video	LVDS	LVDS	LVDS	TTL	TTL	LVDS	LVDS	TTL	TTL	LVDS	LVDS	TTL	TTL	TTL		
	Analog Video	SVGA	SVGA	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	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
	USB 2.0 Ports	2	2	2													
	USB Ports				2	2	2	2	2	2	2	2	2	2	2		
	10/100Base-T Ethernet	1	1	1	1	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	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
- User defined FPGA

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

Video Controllers
- Analog VG A
- TTL and DVI panels

Communication Modules
- Copper or fiber Ethernet
- USB 2.0 and Firewire
- CAN Bus & CAN Spider
- 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
- 50/75/83/100 Watts
- Wide input range
- 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

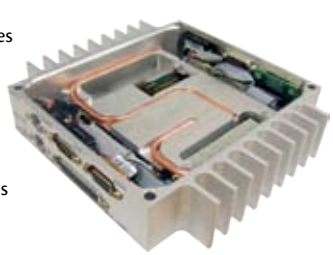

EMBEDDED-PC PC/104 MODULES

Full Product Line and Pricing Online

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

www.rtd.com

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	FPGA7800HR
Bus	AT Expansion Bus	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Bus	PCI Expansion Bus Master	✓	✓				✓							✓	✓
Bus	McBSP Serial Ports	✓	✓				✓								
Analog Input	Single-Ended Inputs	16	16	16	16	16	16								
Analog Input	Differential Inputs	8	8		8	8	8								
Analog Input	Max Throughput (kHz)	1250	1250	40	500	100	1250								
Analog Input	Max Resolution (bits)	12	12	12	12	16	12								
Analog Input	Input Ranges/Gains	3/7	3/7	3/1	3/4	1/4	3/6								
Analog Input	*Autonomous* SmartCal	✓	✓												
Analog Input	Data Marker Inputs	3	3		3		3								
Conversions	Channel-Gain Table	8k	8k		8k	8k	8k								
Conversions	Scan/Burst/Multi-Burst	✓	✓		✓	✓	✓								
Conversions	A/D FIFO Buffer	8k	8k		8k	8k	8k								
Conversions	Sample Counter	✓	✓		✓	✓	✓								
Conversions	DMA or PCI Bus Master	✓	✓		✓	✓	✓	✓						✓	
Conversions	SyncBus	✓	✓				✓								
Digital I/O	Total Digital I/O	16	16	16	16	16	16	16	48	18/9	32	64	32	48	48
Digital I/O	Bit Programmable I/O	8	8		8	8	8	8	24	6/0				48	✓‡
Digital I/O	Advanced Interrupts	2	2		2	2	2	2	2					2	
Digital I/O	Input FIFO Buffer	8k	8k		8k	8k	8k							4M	8M
Digital I/O	Opto-Isolated Inputs										16	48	16		
Digital I/O	Opto-Isolated Outputs										16	16			
Digital I/O	User Timer/Counters	3	3	3	2	3	3	3	3	3				10	6
Digital I/O	External Trigger	✓	✓		✓	✓	✓	✓	✓					✓	
Digital I/O	Incr. Encoder/PWM									3/9					✓‡
Digital I/O	Relay Outputs												16		
Analog Out	Analog Outputs	2	2		2	2	2	4							
Analog Out	Max Throughput (kHz)	200	200		200	100	200	200							
Analog Out	Resolution (bits)	12	12		12	16	12	12							
Analog Out	Output Ranges	4	4		3	1	4	4							
Analog Out	D/A FIFO Buffer	8k	8k				8k	8k							

‡ User-defined, realizable in FPGA

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**
- MIL I/O connectors
- Optional MIL-SPEC paint
- Shock-mounts 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"®

Design Methodologies Help Leverage IP for FPGA-Based Development

For FPGA-based designs today, the spectrum of IP available on the market, common implementations in industry and fine tuning with technologies under the hood make leveraging IP a strong consideration.

by Jeff Harriman, Xilinx
Jeff Meisel, National Instruments

Code reuse in FPGA designs offers advantages such as a reduction in overall software development costs and a shorter time-to-market, both of which are well documented. A feature sometimes overlooked, however, is the tremendous design flexibility gained when using intellectual property (or IP blocks) through recent technological advances in compiler optimization, place and routing, and verification tools.

As FPGAs have become faster and more powerful, the amount of work required of them has also grown. Years ago, an FPGA was expected to provide some essential glue logic to tie a board together. But today, the trend is to pull tasks historically handled by dedicated ASICs into the heart of the FPGA design. The increase in work handled by FPGAs can be attributed to the giant leaps in technology that have pushed speed and size to new levels. However, the challenge is effectively using an FPGA's fabric resources, which requires a high-level understanding of the architecture and routing process. Luckily, FPGA vendors and third-party experts offer building blocks in the form of IP blocks (also called IP cores) to simplify the design process.

IP cores cover an unlimited spectrum ranging from basic functions to extremely complex design blocks. Vendors and third parties offer cores including networking interfaces, system I/O interfaces, communications blocks, digital signal processing (DSP) functions and external memory interface controllers, as well as a suite of IP for embedded systems. These general categories feature several predefined functions such as those listed in Table 1. Cores, which can typically be parameterized to meet the specific needs of a design, are optimized to take advantage of the features of the FPGAs for which they were designed. Most cores from FPGA vendors come with documentation similar to stand-alone ASIC data sheets. You can see the growth in the popularity of code reuse by the growing online community of developers sharing open-source IP through Web sites such as OpenCores.org.

System requirements can change over time. You should expect and plan for this at the outset of your design. Figure 1 shows how a hardware developer of a DSP board could implement an FPGA-based design built almost entirely on IP blocks. The block highlighted in red, which represents the "in-house IP" that the company creates to differentiate itself from the competition, is implemented as a processor core or as state machine logic. The company also must implement a communication protocol among blocks, so data can be bused between the different components.

For instance, you can recycle a DSP application involving a specific chain of filters with different interfaces. Once you have developed an algorithm to perform the desired signal processing, you can alter the means by which data is presented to the filter relatively quickly, from an Ethernet interface to a PCI Express interface, without having to redesign the signal processing chain or the bulk of the data transfer interface.

In the current technology environment, cutting-edge designs can quickly become outdated. With IP, you can save development time and avoid chasing after technology trends. For example, National Instruments uses the PCI Express LogiCORE from Xilinx to migrate its PCI-based data acquisition systems to the newer PCI Express standard.

The design flexibility that IP brings to the table also provides you with the ability to use higher-level tools for rapid system development by simply dropping in an IP block. The perceived trade-off for going to a higher level of abstraction in programming languages is that an easier development experience comes at the cost of code optimization. However, because IP is often pre-compiled and optimized for a partic-

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

ular FPGA, the cost of going to higher-level tools is negligible. LabVIEW FPGA and Xilinx System Generator for DSP are examples of two such tools.

Figure 2 illustrates the simplicity of a LabVIEW FPGA application that you can use to implement a LogiCORE FIR filter from Xilinx for targeting a Virtex-II FPGA found on National Instruments commercial off-the-shelf (COTS) hardware. The parallel nature of graphical programming maps intuitively to an FPGA.

Audio, Video, and Image Processing	JPEG Encoder	MPEG Encoder	2-D Discrete Cosine Transform
Basic Logic	Comparators	Multiplexers	Counters
Bus Interface and I/O	SPI, I²C	PC/104, USB	CAN
Communication and Networking	Interleaver	Gigabit Ethernet Controller	Viterbi Decoder
Digital Signal Processing (DSP)	FFT, FIR, IIR, Reed-Solomon	Radar Pulse Compression	Digital Downconverter Digital Upconverter
Math	Multiplier Accumulator	Discrete Wavelet Transform	Sine Cosine Look-Up Table
Memory Interface and Storage Element	Dual-Port Memory	SDRAM Controller	Asynchronous FIFO

Table 1 Example IP blocks available from FPGA vendors and third parties.

Under the Hood

Developing intellectual property source code in an organized, modular and hierarchical fashion results in faster development, better performance and better code reuse in the long run. The individual function blocks, which serve a well-defined purpose, have clearly defined inputs, outputs and parameters. Because the block is well defined, you can more easily avoid negative feature creep—adding unrelated features that are specific to a single use case. Instead, you can add features over time that make the function block more general-purpose to suit a wider array of applications. This is "positive" feature creep. Additionally, you can test, validate and optimize function blocks on an individual basis. Knowing these key IP development principles can help you develop your own in-house IP or use commercially available blocks.

No matter the tool flow, the same design principles apply when optimizing your IP for a particular device. The key here is a solid understanding of the tools at your disposal, including the most effective ways to use them. This translates to deep architecture knowledge of your target FPGA. The standard features of programmable logic devices now include modular building blocks that improve size and performance for the majority of designs. Various device families include dedicated 18x18 multiplier blocks, 18 Kbit-lock memory, digital clock managers, FIFOs and a DSP48 slice to complement the rest of the configurable FPGA fabric. Understanding the capabilities of these blocks and the routing resources available helps you maximize your design performance.

While you can often infer these blocks using synthesis tools with generic code, you need to ensure their optimal use. A good example of this is a custom multiply-accumulate (MAC) FIR filter. A MAC FIR is comprised of storage elements, control logic, a multiplier and an adder. The specs for such a design should, at a minimum, include the data sample rate and the desired frequency response. At this point, take stock of the requested specifications and observe how they map to your hardware with regard to your target architecture. What is the ratio between the sample rate and the system clock? How many taps do you need and what bit width should you use taking quantization into account?

When you answer these questions, do not just consider the algorithm. Also examine the building blocks already available. In a Virtex-4, the DSP48 slice is ideal for performing MACs with a system clock running up to 500 MHz. Once you determine how many coefficients you need, compare that to the different memory resources available. The four input look-up tables map well to memories with depth increments of 16. However, as distributed memories become large, the speed they can run at decreases, and memory can become a bottleneck. At what point should you switch over to a dedicated block RAM? The answer to that question depends on your performance and resource requirement balance. In a predefined IP block, you never have 100 percent flexibility to control all of these details.

Tweaking the Tools

Once you have written the code and verified it behaviorally, you still have several steps left that influence the final hardware results. If you have been careful to map your design well to the device primitives, you should not need to iterate through the source

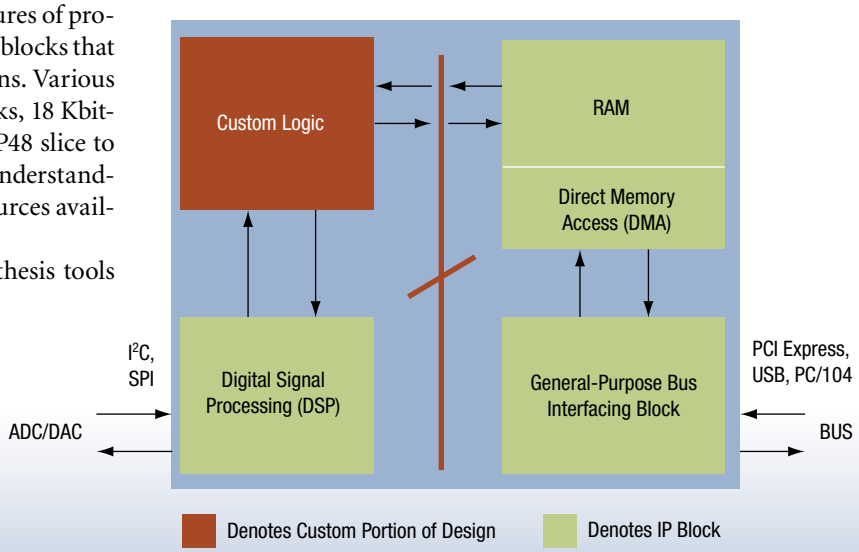

Figure 1 System-level diagram of an IP-centric FPGA design.

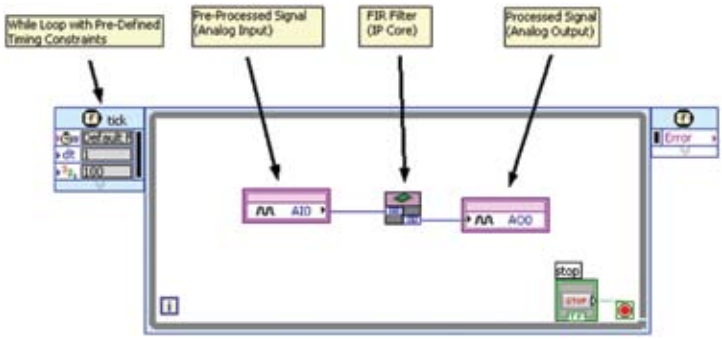

Figure 2 Example of using an IP block in a high-level programming environment (LabVIEW FPGA).

code. To maximize the speed performance and minimize the area at this stage, you must know which buttons to push in synthesis and be familiar with the back-end tools. In a large design, hierarchy is important for organization. You can keep the boundaries between modules intact or you can allow the tools to blur them.

Depending on the blocks that make up your design, maintaining the hierarchy for some parts of the design but not others may be wise. When a core has area or placement constraints that are necessary to meet a specified performance, retaining hierarchy allows the constraints to implement their intended functions. In other parts of the design, optimization across hierarchical boundaries helps the tools peer inside the netlist making up the core and look for ways to optimize signals that transmit between it and other parts of the design.

You should approach timing closure in a similar way for designs involving custom functions or designs heavy in third-party IP. Register balancing and hierarchy optimization in the early stages help the place and route tools achieve your desired performance. Occasionally, you need additional placement constraints to meet all timing constraints. Technological advances make it easier to push the maximum performance of a design to higher levels. This means you can focus on the functionality and let the back-end tools find the optimal placement for registers within a combinatorial path. These tools also use intelligent algorithms to determine when register duplication provides an additional speed boost.

Understanding the tools available to you as well as your hardware capabilities helps you produce efficient and robust designs. While creating custom functions and interfaces is part of any design, keeping reuse in mind helps you protect the investment on work you have already done. High-level tools can assist you in implementing this concept when using hierarchical and modular design flows. Finally, with IP cores, you can quickly create complex designs when your resources and time are limited. ■

Xilinx
San Jose, CA.
(408) 559-7778.
[www.xilinx.com].

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

More Features. More Rugged. More reasons to choose ICS.

The one and only ICS-8550.

Designed for high-speed data acquisition applications such as Software Defined Radio, SIGINT, tactical communications and radar, the ICS-8550 XMC module — which is available for both benign and rugged environments — can simultaneously sample two RF/IF inputs at frequencies up to 210 MHz at a resolution of 12 bits.

With the industry-leading Xilinx Virtex-4 FPGA at its heart to deliver unprecedented power and user programmability — and enabling IF/UHF signals to be processed directly on the board itself, freeing the host board for other tasks — the ICS-8550 is truly an ADC module that sets new standards.

And with up to eight lanes of high-speed serial I/O, the ICS-8550 provides the throughput to match its performance and flexibility. Configure it with Radstone's remarkable V4DSPFPGA/PowerPC processor, and the partnership is unbeatable.

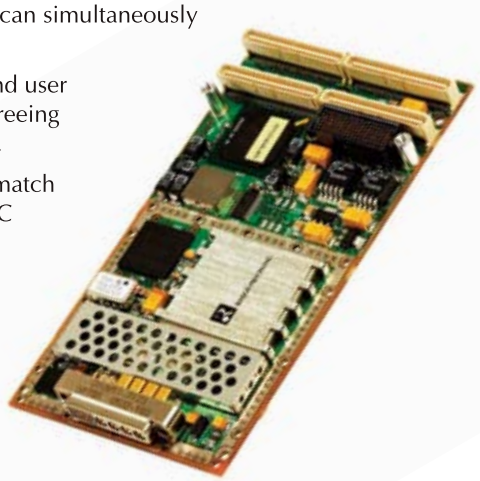

ICS. For when you need more.

ICS
SENSOR PROCESSING

CHECK THE FULL SPEC NOW AT
www.ics-ltd.com/8550

OR CALL TOLL FREE (US ONLY)
800-267-9794 or 613-749-9241

Embedded FPGA Soft Core Processor Enables Universal CompactPCI Applications

by Pat Mead, Altera
Barbara Schmitz, MEN Mikro Elektronik

Programmable logic has reached such a state of advancement in terms of speed and density that it has become a truly attractive alternative to RISC and CISC processors. It can form a "matrix" within which processing, peripherals, data paths and algorithms can be placed to create powerful, flexible and upgradable systems. Programmable logic is now available in forms and sizes that range from the traditional use as glue logic up to structured ASIC replacements and even further. To date, to fully use the advances of this key technology—re-programmability, reusability and upgradability—you need to be a FPGA expert, but these benefits need to be opened to a much wider market.

A Nios-II-CompactPCI development package designed as an open FPGA platform includes a sample design with a PCI system unit, integrating the standardized Wishbone bus and the Altera Avalon switch fabric. The PCI system unit forms the interface to the PCI bus, where the CPU board can then be addressed as a PCI slave. It connects to the Wishbone bus where a SDRAM and a flash controller are already implemented.

The 3U CompactPCI card with a Cyclone FPGA and the integrated Nios II microcontroller soft core is designed for final use in volume in production and it acts at the same time as the standard FPGA development platform for this application (Figure 1). As a universal FPGA platform, the board has a multitude of directly accessible I/O pins. The Nios II CPU in the Cyclone FPGA provides performance similar to an ARM processor. It allows the use of the CPU board, for example, as an intelligent slave on the CompactPCI bus. The FPGA and the integrated processor core support a 32-bit bus with 33 MHz, control 32 Mbyte SDRAM and support read and write to and from the 2 Mbyte flash memory. The special flash structure provides initial programming using a boundary scan interface. Once configured, the FPGA may be reconfigured at any time during operation with data from the CompactPCI bus.

The FPGA also controls four status LEDs and up to 83 user-defined I/O pins. The final functionality of the board depends entirely on the application and can be anything from a simple UART solution up to a complex analog front end with DSP-like data pre-processing. In any case, the CompactPCI card supports a nearly endless range of applications. The designer can use IP cores to configure the function. This includes different serial interfaces from RS-232 to intelligent HDLC protocols up to Fast Ethernet. Other functions include graphics, fieldbus connections or digital I/O. On the other hand, a development package included allows the user to create custom cores or to integrate third-party cores from opencores.org or from Altera.

Open-Platform Development Concept

The user can now add any kind and number of IP cores to the Wishbone bus. To do this, a "Wishbone Bus Maker Tool" has been developed, which can be used to generate the Wishbone bus and which is part of the development package. The Wishbone Bus Maker can generate multi-master and multi-slave bus systems. A Wishbone-to-Avalon-bridge, and vice versa, an Avalon-to-Wishbone-bridge (Figure 2) allow the additional integration of Avalon-based IP cores and especially of the Nios II soft processor core. Nios II connects to the Avalon switch fabric, where a GPIO module for the user LED control is already implemented as well. The user can now also add any kind and number of IP cores to the Avalon switch fabric by using the SOPC Builder tool from Altera, which is part of the Quartus II development package.

The open-platform concept for Nios-based development and integration of IP cores for Wishbone and the Avalon switch fabric can also be expanded on the hardware side. Currently a new generation of ANSI standard PMC and M-Modules is being designed that uses the Cyclone II family. The entire board logic including the individually configurable FPGA is located on the base PMC or M-Module, while the physical interfaces are implemented on an adapter card, which can be plugged either on the PMC or the M-Module. Since both PMC and M-Modules can be used on all kinds of system platforms using dedicated carrier boards, the user can now concentrate completely on the FPGA application.

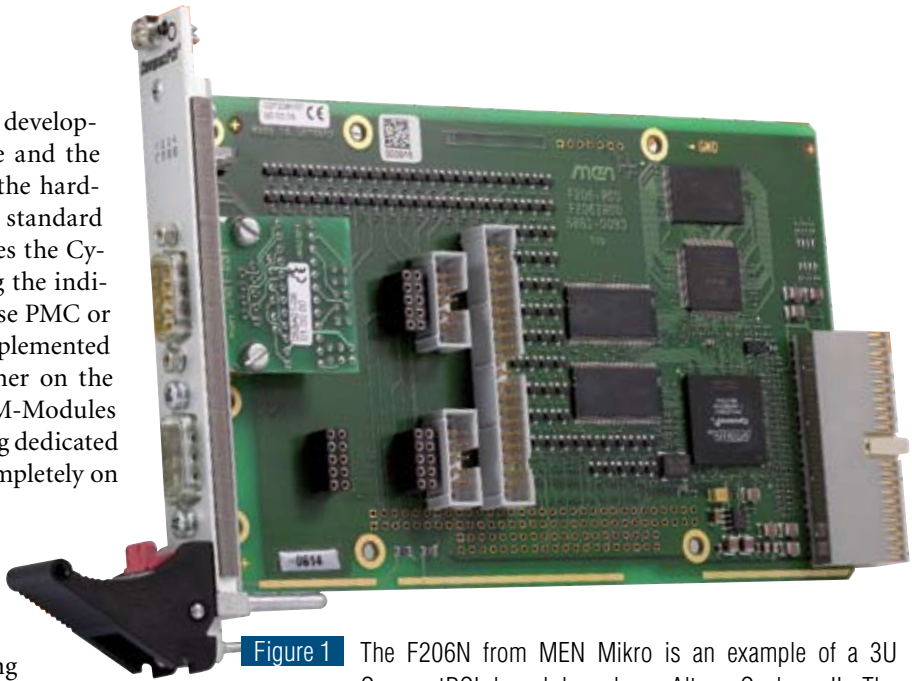

Figure 1 The F206N from MEN Mikro is an example of a 3U CompactPCI board based on Altera Cyclone II. The functionality of the board is entirely dependent on the IP programmed into the FPGA.

Flexibility—Base for Comprehensive Applications

An example application is now in use in the automated driverless underground. A leading manufacturer is using standard 19" systems. These redundant built-up 3U-CompactPCI systems feature a Pentium III CPU, analog and digital I/O, sensors for position encoders and an optional MVB link. A feedback channel allows sending back data from the vehicle permanently to the central control station.

FPGAs are used in three different boards. The FPGA of the CPU board contains a watchdog and different UARTs; the NAND flash is controlled by the Nios II microcontroller. A second board features eight UARTs completely integrated in the FPGA used for asynchronous RS-422 operation and optional synchronous HDLC (without Nios). The third board features digital I/Os, analog outputs, counter pulses, radar sensor and interrupt inputs—all implemented in the FPGA hardware without using the Nios II. This standard CompactPCI system can be configured for the requirements of the operation of the underground trains in different cities and countries by only changing the content of the FPGA—the hardware remains the same.

Modern low-cost FPGA components have a usable size. The Cyclone II family from Altera contains nearly 70,000 logic elements and their pricing is acceptable starting from just a few dollars for the smaller devices. This makes them effective factors for cost savings and time-to-market when making individual configurations of standard products. A time-consuming and expensive redesign of a board can often be avoided through application-specific integration of IP cores in the FPGA. Furthermore, FPGA technology is indispensable wherever long-term availability or harsh industrial environments are involved. IP cores per se are not threatened by discontinuation, even if an FPGA component may be replaced by a newer one after 10 years, for instance.

The Nios II family of 32-bit RISC embedded processors delivers more than 100 DMIPS of performance when implemented in the Cyclone II family. Because the processors are soft core and flexible, it is possible to choose from a nearly

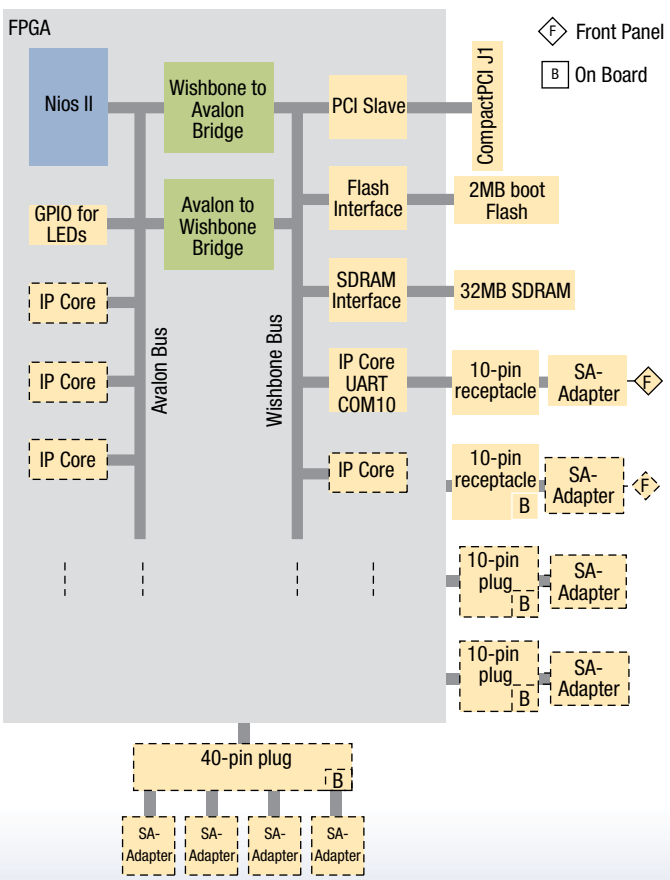

Figure 2 Bridges implemented inside the FPGA allow integration of standard IP cores along with Altera IP, including the NIOS II soft core processor connected to Altera's Avalon switch fabric.

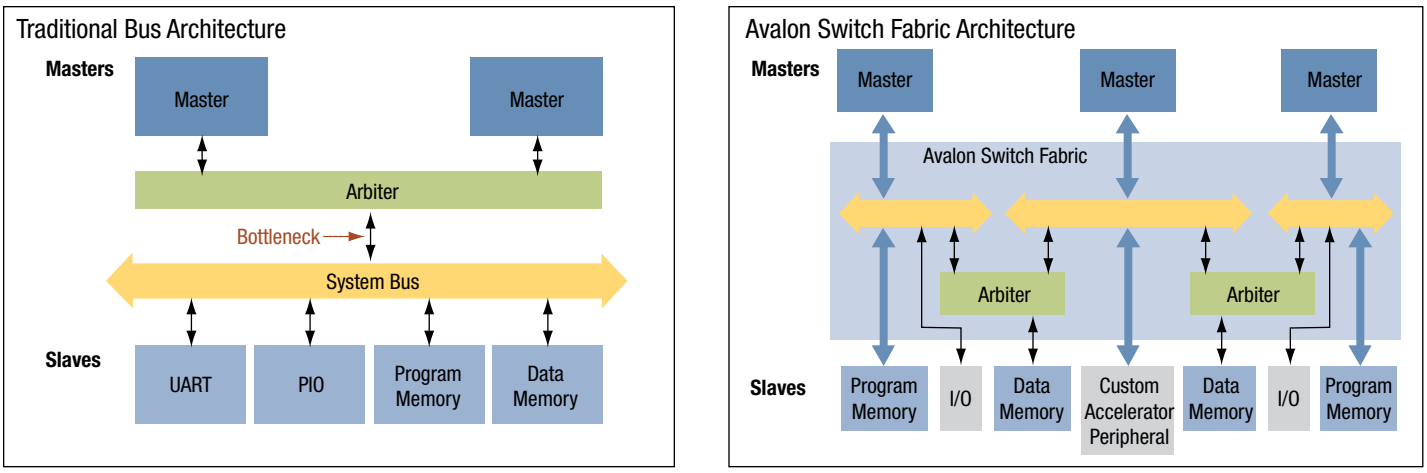

Figure 3 Avalon switched fabric architecture compared to a traditional bus architecture.

unlimited combination of system configurations to meet the required performance, features and cost. The Nios II processor family consists of three cores—fast (Nios II/f), standard (Nios II/s) and economy (NIOS II/e)—each optimized for a specific price and performance range. All three cores share a common 32-bit instruction set architecture and are 100 percent binary code compatible. A library of commonly used peripherals and interfaces is included in the Nios II development kit. A complete list of SOPC builder-ready intellectual Property (IP) and peripherals can be found at the Altera Web page. Using the interface-to-user-logic wizard in the SOPC Builder software, enables the creation of custom peripherals and their integration into Nios II processor systems.

Optimized Use of Bus Resources

The Avalon switch fabric enables multiple, simultaneous data transactions for maximum system throughput. SOPC Builder automatically generates an Avalon switch fabric optimized to the specific interconnect requirements of the final system processors and peripherals. In traditional bus architectures, a single arbiter controls the communication between the bus master and slaves. Each bus master requests control of the bus, and the arbiter then grants bus access to a single master. If multiple masters attempt to access the bus at once, the arbiter allocates bus resources to a master based on a fixed set of arbitration rules. This can lead to a bandwidth bottleneck as only one master can access the system bus and its resources at a time.

The Avalon switch fabric's simultaneous multi-master architecture increases the system's bandwidth by eliminating this bottleneck (Figure 3). Using the Avalon switch fabric, each bus master gets its own dedicated interconnect, meaning that bus masters only contend for shared slaves, not for the bus itself. Each time a component is added or the peripheral access priorities change, SOPC Builder generates a newly optimized Avalon switch fabric with a minimum of FPGA resource use. The Avalon switch fabric supports a wide range of system architectures, including single- and multiple-master systems, and allows seamless data transfers between peripherals with performance-optimized data paths. Off-chip processors and peripherals are equally well supported.

Custom instructions allow developers using Nios II processors to increase system performance by extending the CPU instruction set to accelerate time-critical software. Using custom instructions enables the optimization of system performance in a way not possible with traditional off-the-shelf processors. The Nios II family of processors supports up to 256 custom instructions to accelerate logic or mathematically complex algorithms normally handled in software. For example, a block of logic that performs a cyclic redundancy code calculation on a 64 Kbyte buffer operates 27 times faster as a custom instruction than when performed by software.

Nios II processors support fixed and variable cycle operations, include a wizard for importing user logic as a custom instruction, and automatically create software macros for use in developers´ code. Large blocks of data can be processed concurrently with CPU operation by adding application-specific hardware accelerators that act as custom co-processors within the FPGA. Using the cyclic redundancy code example, processing a 64 Kbyte buffer runs 530 times faster with hardware accelerators than software. SOPC Builder includes a wizard that allows developers to add their acceleration logic and DMA channel to the system.

A complete set of tools is available for the hardware design, including the SOPC Builder system development tool, Quartus II design software, ModelSim-Altera software and SignalTap II embedded logic analyzer. Hardware design for creating Nios II processor-based systems uses the SOPC Builder system development tool to specify, configure and generate systems. Launching from within the Quartus II design software, SOPC Builder provides an intuitive wizard-driven graphical user interface for creating, configuring and generating system-on-a-programmable-chip (SOPC) designs.

To make the software design flow as easy as possible, it is possible to accomplish all software development tasks within the Nios II IDE, including editing, building, debugging programs and flash programming. As part of the Nios II IDE, Altera partners with operating system and middleware providers for additional software development tools. A PC, an Altera FPGA device and a JTAG download cable is everything you need to develop and debug Nios II processor-based systems. The Nios II architecture supports a JTAG debug module that provides on-chip emulation features to control the processor remotely from a host PC. The IDE can communicate with the JTAG module on one or more processors. This allows downloading programs to memory, starting and stopping program execution, setting breakpoints and watch points, analysing registers and memory and collecting real-time execution data.

The instruction set simulator (ISS) makes it possible to begin developing programs before the target hardware platform is ready. Many designs also incorporate flash memory on the board. Therefore any CFI-compliant flash device connected to the FPGA can be programmed using the IDE flash programmer. The flash programmer is pre-configured to work with all of the boards available with the Nios II development kits, and can be easily ported to any custom hardware. In addition to a project set-up wizard, the IDE provides software code examples, in the form of project templates, to help bring up working systems as quickly as possible.

The IDE enables quick system customization using system software. The hardware abstraction layer (HAL) library is a lightweight runtime environment that provides a simple device driver interface for programs to communicate with underlying hardware. MicroC/OS-II from Micrium is a complete, portable ROM-able, pre-emptive real-time kernel, shipped with all development kits and includes full source code, reference manual and free developers' licence. Included in the development kit is also an open-source IwIP TCP/IP stack that is built to work with the MicroC/OS-II applications and implements the standard UNIX socket API as well as a full-featured LINUX operating system. □

MEN Mikro Elektronik
Lago Vista, TX.
(512) 267-8883.
[www.menmicro.com].

FGPA-Based Development for Defense and Aerospace Applications

The right FPGA design toolkits not only speed development, but can simplify the addition of custom IP so designs can be tailored to specific applications.

by Steve Edwards
Curtiss-Wright Controls Embedded Computing

The parallelism, speed and I/O flexibility provided by today's FPGAs make it possible for system engineers to replace multiple processor boards with a single FPGA COTS board. For the defense and aerospace market, where high performance frequently must be traded off with power and size/weight restrictions, these FPGA boards offer the best of both worlds: high performance in a single slot.

Formerly, FPGA solutions had the reputation of being costly, due to long development cycles and high development costs compared to traditional software-based solutions. But today's FPGA device families, combined with appropriate FPGA design kits, help engineers get designs to market rapidly. In addition, they offer lower costs and greater flexibility that simplifies adding the system developer's intellectual property (IP).

One of the prime advantages of today's FPGAs is the balance they provide between processing and I/O. This balanced approach makes FPGAs very efficient at simultaneously processing several high-speed, parallel data streams. Such I/O versatility means that multiple banks and types of high-speed memory, including DDR SDRAM and DDR SRAM, can be connected to an FPGA.

Many newer families of FPGAs, such as Xilinx's Virtex-II Pro and Virtex 4, feature high-speed serial transceivers, each capable of throughput of 3.125 Gbits/s or greater. These transceivers support many of the high-speed serial interfaces used in emerging defense/aerospace applications, such as Serial RapidIO (SRIO), PCI Express (PCIe), 10 Gigabit Attachment Unit Interface (XAUI) and Gigabit Ethernet.

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

The combination of parallel processing with fast, synchronous memories and high-speed serial transceivers lets system designers replace multiple processing cards with a single card containing two or more FPGAs. This results in systems with lower power requirements, weight and cost.

FPGA-Based COTS Boards

FPGA-based COTS boards targeted to the defense/aerospace market often share certain common elements. For example, they all provide one or more high-density, high-performance FPGAs, high-speed serial I/O and high-performance memory.

But hardware architecture is only part of the solution. To realize the full value that FPGAs can deliver, the success of an FPGA-based project is dependent on having the right FPGA design tools. Some of these can shorten time-to-market and lower development costs, while also providing an open architecture that makes it possible to tailor the design to specific applications.

Systems engineers evaluating FPGA boards must address certain challenges. These include how to add their algorithms to the FPGA and how to simulate the design at the system level to ensure that these algorithms work. Because FPGA development can be costly, engineers also must evaluate whether their approach will save time and money and if the resulting solution will be robust and reliable.

To understand the importance of design tools, it is useful to consider how data is handled by an FPGA board. For example, one board (Figure 1) uses two Virtex-II PRO VP-70 or VP-100 FPGAs, several high-speed serial interfaces and a large amount of high-speed memory to achieve a high-performance, reconfigurable computing engine. The combination of DDR SDRAM for bulk storage and DDR SRAM for fast, non-sequential stor-

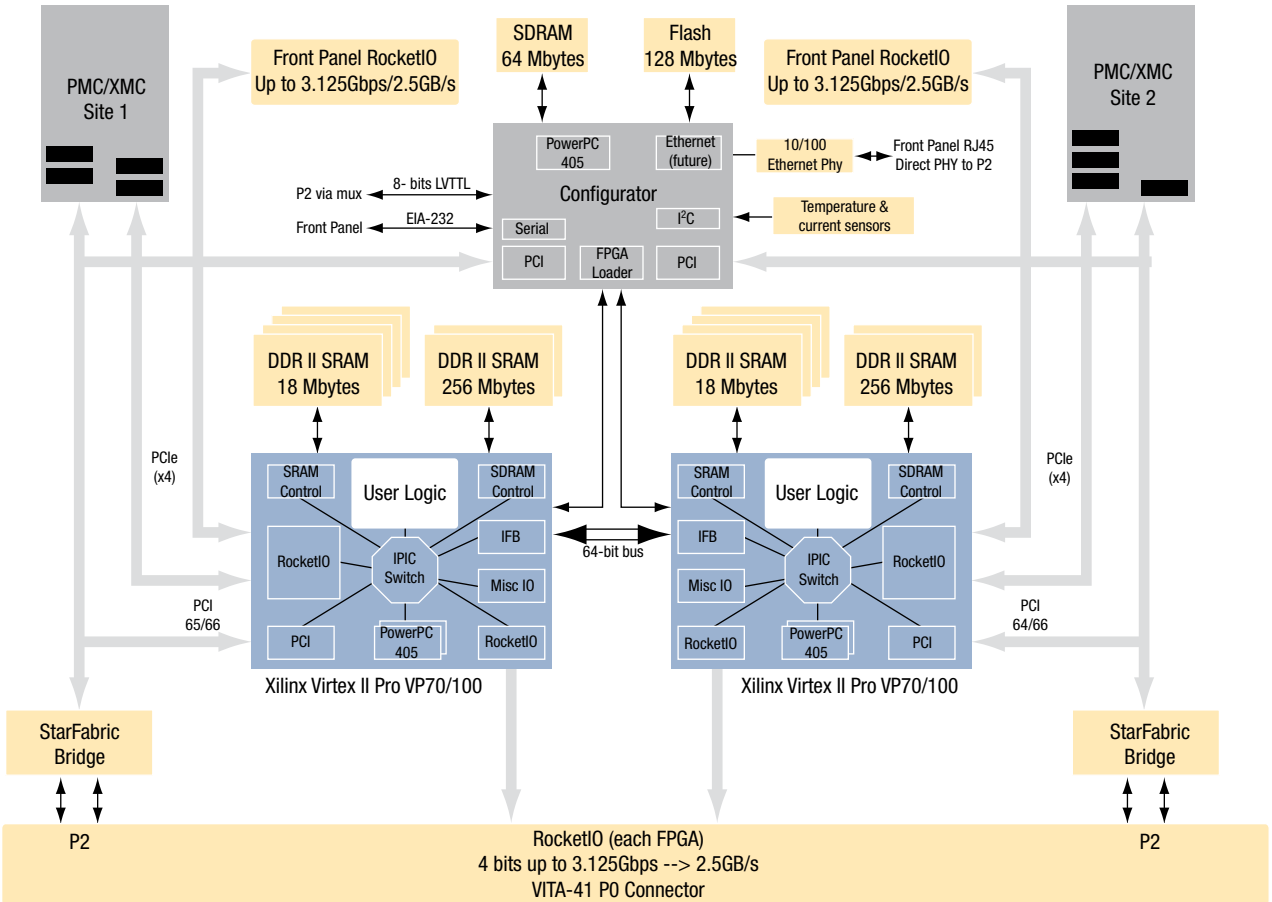

Figure 1 Curtiss-Wright's CHAMP-FX architecture balances high-speed processing, memory and I/O.

age of algorithm data allows flexibility in mapping algorithms to this architecture.

In a typical application, data comes in via a high-speed serial interface and is stored in SDRAM. This data is then pulled out of SDRAM by the end application and processed via the customer's algorithm. Intermediate results are stored in internal or external SRAM and the final result is decimated to a lower data rate.

The FPGA Developers Kit

For DDR SDRAMs, FPGA designers must confront the challenges of aligning data and data strobes, tight timing constraints, signal integrity issues and simultaneously switching output (SSO) noise. In addition, certain design issues can prolong design cycles or force them to accept reduced performance.

To make matters worse, all of these hurdles become more pronounced at high frequencies. On a read of data from SDRAM, the data is valid for only two to three nanoseconds. Significant effort is required to latch it reliably inside the FPGA and then synchronize it with the rest of the logic there. The FPGA designer is thus faced with significant challenges that may take several man-months to complete, but that can be solved by using the IP in some FPGA developers kits.

In the ideal FPGA developers kit, all of the high-speed IP provided is fixed to certain regions within the FPGA (Figure 2). This is done to ensure that all critical paths meet timing, as well as to confine the overall IP design to a small region of the chip to minimize logic resources.

FPGA IP Designed Specifically for the Hardware

Designers must also make sure that the IP works with the hardware. One potential challenge is the issue of SSO noise. In a Xilinx FPGA, when too many outputs in a particular bank switch at the same time, simultaneously switching output introduces noise into the system and causes one or more bits in a particular bank to flip to the wrong value.

This can be a tricky problem to track down if the designer is not familiar with the SSO phenomenon and with the particular board hardware involved. The FPGA board supplier should test the board with SSO in mind, and its memory interface pinouts should be selected to avoid this issue.

A DDR SDRAM read cycle provides a good example of how the FPGA IP and the board hardware must work together. In order to read data, the FPGA sends out a clock signal to the SDRAM and waits two clock cycles for a return four-word burst. The challenge is how to clock the data back into the device. Because of trace delays in the PCB, there is skew between the

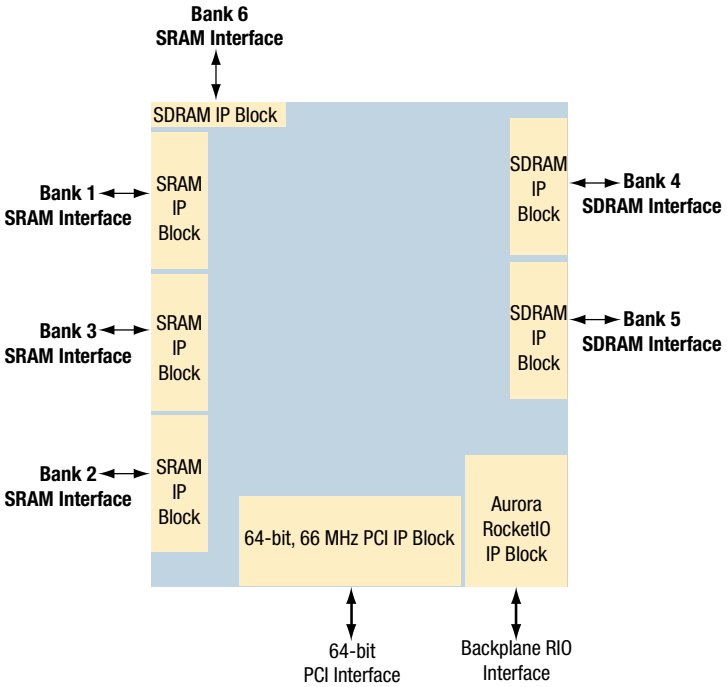

Figure 2 To minimize routing delays and improve timing, I/O IP blocks are located along the edge of the FPGA near I/O pins.

transmitted clock signal and the incoming data. This is not an optimal design approach because skew minimizes the window during which valid data can be clocked into the FPGA.

A better approach is to compensate for trace delays in the FPGA by phase shifting the internal clock relative to the incoming data. This allows the FPGA to clock in read data with the appropriate timing margins. However, to do this the FPGA IP designer must be intimately familiar with the PCB characteristics of the FPGA-SDRAM interface.

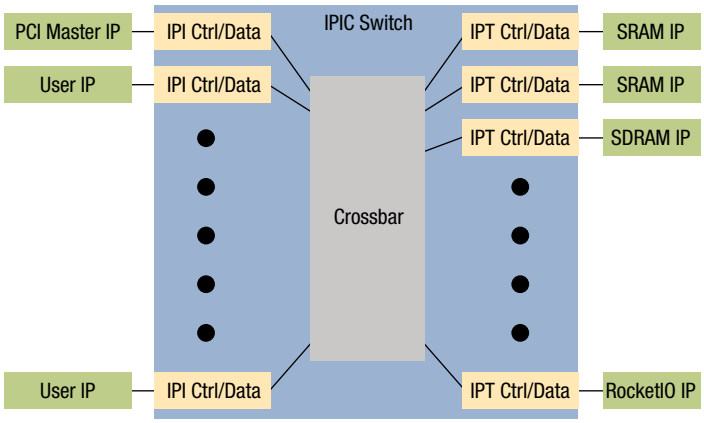

Figure 3 An internal crossbar switch is used to connect user IP to IP supplied in the CHAMPtools-FX FPGA Developers Kit.

Some argue that, to simplify the design process and reduce design cycle time, designers should use memory-controller IP cores provided by FPGA vendors or third-party suppliers. This is a valid assertion if understood in the right context. FPGA IP that has been designed for a particular hardware platform with the end application in mind, and that has been tested and qualified for use in that application, will certainly save the designer time and money.

Interfaces and IP Integration

The IP in the FPGA developers kit is extremely important. For example, IP in the CHAMPtools-FX FPGA Developers Kit, for the Xilinx-based board described above, is built around one of two standard Xilinx interfaces: IP Interconnect (IPIC) or Local Link. IP Interconnect is a standard memory-mapped interface that allows the IP block to send data to a particular memory address. SDRAM, SRAM and PCI IP are provided with IPIC interfaces.

Local Link is a streaming data interface in which the data's destination address is predetermined. Local Link interfaces are provided with DMA controllers so that engineers can set up a destination buffer to a memory-mapped location. As data comes into the Local Link interface, the DMA automatically sends it to the correct memory address.

Local Link interfaces usually include RocketIO IP. Both of these interfaces must be well documented in the FPGA developers kit. This documentation should describe the signals for each interface, as well as timing diagrams and anything else needed to develop custom IP using one of these interfaces.

Another important part of the developers kit should be an IP block called the IPIC switch.

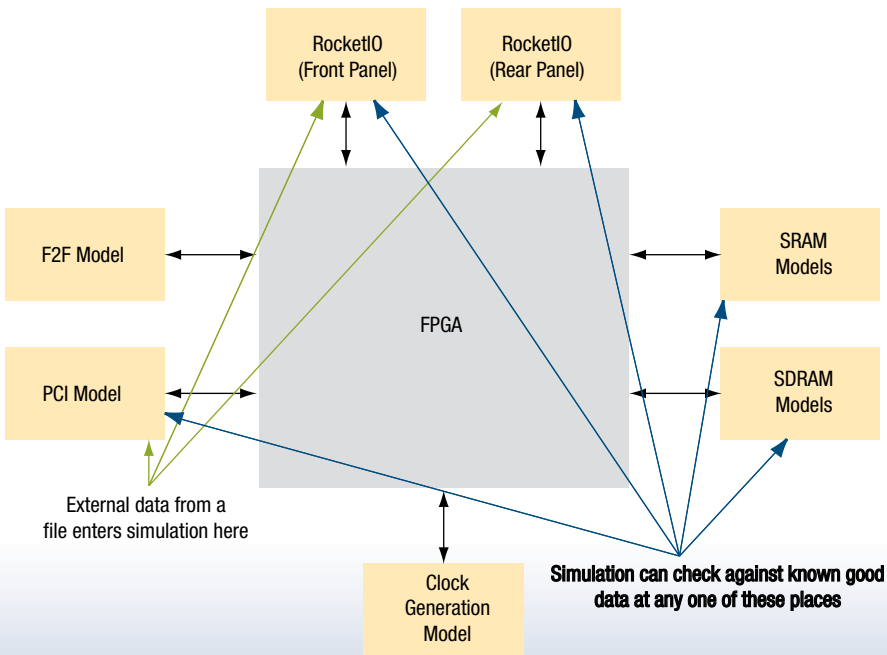

Figure 4 The ability to input known data and check results is a key feature of an FPGA developers kit simulation environment.

DSP-Based RF Solutions!

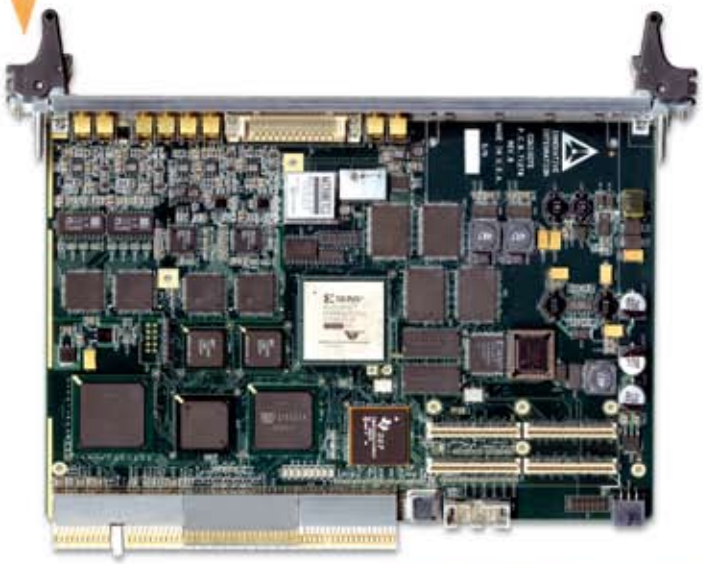

Quixote

Quixote

- 1 GHz TMS320C6416 DSP
- 2 to 6 MGATE Virtex-II FPGA
- 32 MB SDRAM, 8 MB ZBT SBSRAM
- A/D: 105 MSPS 14-bit
- 64/32 bit CompactPCI, 66 MHz, 5V/3.3V
- Complex trigger modes w/HW event log
- PMC Site with Jn4 to FPGA DIO
- PICMG2.17 StarFabric Compliant
- MatLab & Framework Logic Development

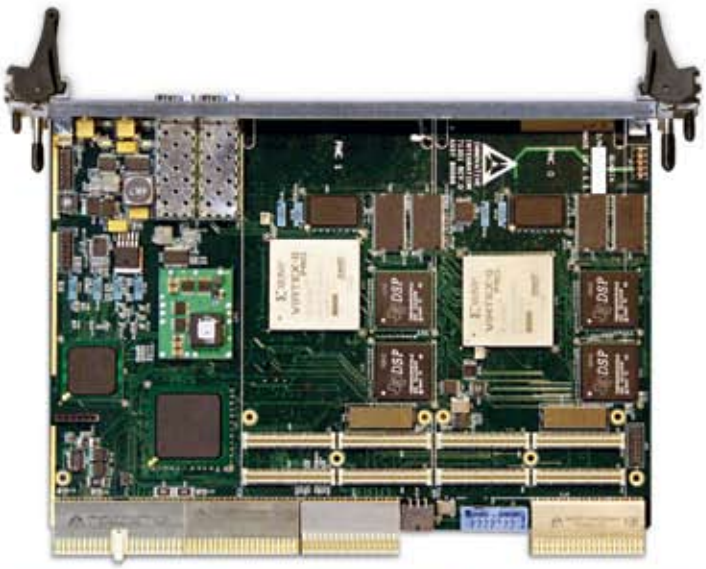

Quadia

Quadia

- 1 GHz TMS320C6416 DSP (x4)
- 64MB SDRAM per Processor
- Flexible communication mesh
- 64-bit/66MHz CompactPCI
- 2 100MHz PMC/XMC Sites - 800MB/s data rate
- External Data Port, up to 2Gbit/s
- Dual Xilinx VP40 FPGAs
- Two 2MB private SBSRAM per FPGA
- Up to 2 128MB private DDR SDRAM per FPGA
- Up to 512MB Global DDR SDRAM
- MatLab & Framework Logic Development

Complete DSP-Based PMC/XMC RF Solutions

SIO
High-Speed Tx/Rx Serial I/O
Dual 1Gbit/sec Full Duplex

UWB
Ultra-Wide Digital Receiver
Dual A/D 210MHz

DR
16 Channel Digital Receiver
Four A/D 125MHz

TX
Digital Transmitter
Four 1 GSPS DAC

Free Instant On-Line Pricing!

805.578.4260 phone
www.innovative-dsp.com

Innovative Integration
... real time solutions!

This block allows other IP blocks with IPIC or Local Link interfaces to be connected together so that data can be transferred among them (Figure 3). Multiple IPIC switches can be instantiated in the design allowing engineers flexibility in choosing the best solution.

Finally, integration is easy when all of the IP can be integrated within the Xilinx Embedded Developers Kit framework. This allows designers to easily instantiate or remove IP blocks and synthesize the design, using a GUI-based software framework.

Simulation and Hardware Test

Design verification, including functional and timing verification, typically consumes the biggest part of the FPGA development cycle. As design complexity increases, this causes a dramatic increase in simulation time. Many FPGA designers recommend budgeting 50% of IP development to simulation. Some of this time is spent simulating at the IP block level, but a good portion is also spent in full simulation of the entire design. This is the last hurdle for the FPGA designer before testing the design on the target hardware.

After all of the bugs have been removed from individual blocks, the design needs to be put together and simulated. This must be done before testing in silicon. Simulation requires a testbench environment that includes models of all of the external interfaces. It also requires the ability to initialize memory interfaces, pass data from the PCI and RocketIO into the FPGA and check memory, as well as the ability to test PCI or RocketIO data against expected data previously stored in files. The FPGA developers kit must provide all the features needed for a full functional simulation of the FPGA.

Today, many FPGA algorithms are developed using a tool such as Matlab, which enables designers to input real data into their algorithms and capture the results in a file. These input and output results can then be reused during simulation.

The simulation environment provided by the Xilinx-based board's developers kit lets engineers input data from a file into the simulation from the board's external PCI or RocketIO interfaces. Simulation data can be stored to a file at any of the memories or the PCI or RocketIO interfaces. These results can then be compared to result data captured during the Matlab simulation, enabling developers to verify the correctness of their VHDL IP model versus the mathematical model originally created in Matlab. The ability to check data at any of the memory interfaces is particularly useful for allowing developers to check integrity at multiple stages within an algorithm.

The final step is testing the design on the hardware. At this point the design is almost complete, but there is still the possibility that the test environment did not test for all possible conditions, or that there is some variation between simulation and the actual hardware. The Xilinx-based board's developers kit makes it possible to automatically insert Xilinx Chipscope IP onto either the IPIC or Local Link interface and specify which ports to monitor. This is an easy way to add a circuit logic analyzer function to the FPGA design, especially at boundary points between customer IP and IP provided by the development tools vendor.

System engineers need the right tools in their toolkit to complete the job well, on time and under budget. The ideal FPGA developers kit provides the right tools for the system engineer to design a robust, reliable product on time and under budget. It is an absolute necessity for getting the job done. ☐

Curtiss-Wright Controls Embedded Computing
Leesburg, VA.
(703) 779-7800.
[www.cwcembedded.com].

The FPGA Eye Opener!

Transmit eye diagram at 6.375 Gbps. For characterization details, see www.altera.com/stratix2gx/characterization.

The complete transceiver solution with best-in-class signal integrity.

Stratix® II GX FPGAs offer the industry's most comprehensive solution for high-speed serial applications. With significantly lower power consumption than competing devices, Stratix II GX FPGAs have up to 20 transceivers performing with exceptional noise immunity from 622 Mbps to 6.375 Gbps. Designers have immediate access to Quartus® II software support and an array of intellectual property (IP) cores, system models, development boards, and collateral—everything needed to complete a design in the shortest amount of time.

Discover Stratix II GX FPGAs today at www.altera.com.

- **Up to 20 transceivers with dynamic pre-emphasis, equalization, and output voltage control**
- **Best-in-class signal integrity from 622 Mbps to 6.375 Gbps with exceptional channel-to-channel noise immunity**
- **Low power consumption (140 mW/channel at 3.125 Gbps and 240 mW/channel at 6.375 Gbps) while retaining best-in-class jitter performance**
- **Dedicated circuitry supporting wide range of protocol standards across multiple data rates**

ALTERA®

The Programmable Solutions Company®

www.altera.com

Cool. Calm. Collected.

The Themis Slice™ architecture can handle whatever you or the outside world throw at it.

Cool.

Innovative and cost-effective cooling system can tame the hottest of processors in four 100 Watt sockets that support eight processor cores. Air and liquid heat exchanger options.

Calm.

Standardized mechanical footprint designed to survive in high shock and vibration environments - up to 40G shock and 1 grms vibration.

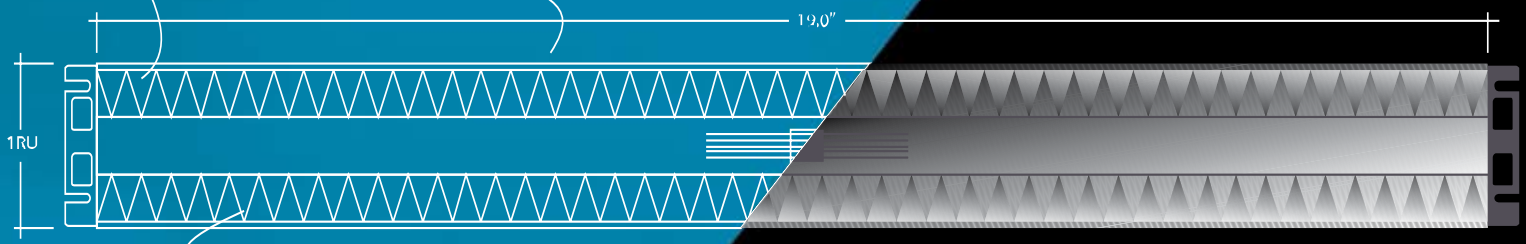

Collected.

Scalable - up to 16 processor cores. High-density, processor-independent architecture implemented with SPARC®, IBM®, PowerPC and x86 technology. Solaris™, Windows®, and Linux® operating systems. All networked in a single system. All managed in real-time.

Themis' Slice technology is a breakthrough in lifecycle management for mission-critical computing.

By functionally disaggregating commercial computing resources and housing them in a standardized footprint, purpose-built enclosure, Slice provides resilience with superior thermal and kinetic management. This open and modular design allows for spiral technology refresh, extending computing infrastructure investments for complete lifecycle management.

Themis' real-time resource manager integrates all computing resources, under dynamic policy control, to optimize application quality of service.

Themis Slice is transformational technology offering the most robust and efficient platform available. No other platform will enable the 21st century war fighter as effectively.

www.themis.com (510) 252-0870

THEMIS

Transformational.

Technology Focus

FPDP and FPDP II Boards

Outside-the-Box Features
Keep FPDP in the Game

FPDP has proven itself as an effective solution for high-throughput, point-to-point data movement. Meanwhile, serial FPDP offers a way to overcome the distance limits of its parallel FPDP predecessor.

Jeff Child
Editor-in-Chief

Sometimes simple ideas are the big winners. Such is the case with the Front Panel Data Port (FPDP) interconnect standard. Using an inexpensive ribbon cable, FPDP links boards without eating up more than a tiny amount of board space. It's particularly useful in military applications like radar and sonar where FPDP is used as the interface to sensor networks. Because FPDP operates independently of the backplane bus, it provides a deterministic sustained bandwidth free from contention.

Because FPDP does not require a host computer at the sensor end, processing overhead is reduced to the bare minimum, allowing for maximum data throughput. But more importantly, the elimination of the host computer means that no operating system or file system is required. With the OS eliminated, the complexities of the OS are also eliminated. That's key for ensuring that none of the high-speed data coming from the sensor will be dropped or lost.

Introduced to the industry by Interactive Circuits and Systems (ICS) in 1994, FPDP was one of the first successful methods for transporting large quantities of data quickly and continuously. It is an ANSI/VITA standard (ANSI/VITA 17-1998) and is widely accepted in the market as a means of moving data between sensors and DSPs effectively and efficiently. The specification defines 32 bits of data moving in one direction, over ribbon cable. The clock can operate at up to 40 MHz, resulting in a maximum data transfer rate of 160 Mbytes/s. This data rate is sufficient for many radar, sonar and other high-performance applications. The latest version, FPDP II, supports data rates up to 400 Mbytes/s and is finding acceptance despite the fact that it is not a formal standard.

Get Connected with companies mentioned in this article.
www.cotsjournalonline.com/getconnected

Figure 1

The Navy used FPDP for the NSSN sonar transmit system and NSSN Simulation/Stimulation system (Sim/Stim) aboard their Virginia Class Attack Submarines, like the USS Texas. Shown here, the Texas is being moved on land as part of its roll-out from the Modular Outfitting Facility in preparation for its christening ceremony in July of 2004.

Navy Embraces FPDP

Some of FPDP's most notable successes have been in major Navy programs. Driven by a need for rapid deployment and a desire to keep costs down, the Navy used FPDP for the NSSN sonar transmit system and NSSN Simulation/Stimulation system (Sim/Stim) aboard their Virginia Class Attack Submarines, like the USS Texas (Figure 1). FPDP, coupled with VMEbus building block modules reportedly saved the Navy millions of dollars in software development costs and cut years off the development time in those programs.

The most common use of FPDP interfaces has been on DSP systems via VME cards, PMC boards and custom designed interfaces. FPDP often shows up on PMC host bus adapters to interface to DSPs via the PCI bus. That flexibility has allowed FPDP to be used in a variety of systems. While FPDP is a very efficient data transport methodology, the ribbon cable-based, parallel physical architecture includes one fundamental limitation: distance. The ribbon cable is limited to a distance of less than 5 meters (about 15 feet). This limitation rules out using FPDP in many radar and sonar applications, or any system where the distance between the sensor and the processing system must be greater than 5 meters.

Going Serial

Overcoming those distance limitations is a more recent flavor of FPDP, called Serial Front Panel Data Port (Serial FPDP), which is an alternative that provides the high bandwidth and distance requirements found in switched serial fabrics, yet also eliminates the need for a processor and operating system—and their inherent impact on determinism and performance—on the sensor side.

Serial FPDP (ANSI/VITA 17.1-2003) was developed by Systran (now part of Curtiss-Wright) in 1997 as an extension of FPDP. In early 2003, the Serial FPDP specification was finalized and gained ANSI approval. The ANSI/VITA 17.1-2003 specification for Serial FPDP supports 1 Gbit/s, 2 Gbit/s and 2.5 Gbit/s link speeds.

With the data serialized, the Serial FPDP protocol can transmit over extended distances. Using multi-mode fiber optic cable, the maximum distance is 300 meters (approximately 1,000 feet) with 850 nm transceivers, and at least 2 kilometers with 1300 nm transceivers using single-mode cable. Even longer distances are possible using 1550 nm transceivers (up to 50 kilometers). Using copper cable, Serial FPDP can still transmit as far as 10 to 30 meters, depending on system configuration. Supporting point-to-point, broadcast chaining and single or multiple master ring topologies, Serial FPDP is rapidly becoming the interconnect of choice for streaming sensor signal processing systems.

TO GREAT LENGTHS FOR OUR CUSTOMERS.

From the toughest terrains in the universe to the antiseptic environment of the operating room, Hypertronics delivers the highest reliability electronic connectors for our customers' most demanding applications. But our commitment doesn't end there. We provide complete, custom design solutions for customers, including cabling, mechanical design, integrated electronics and electrical modeling capabilities. Hypertronics custom solutions can save you valuable engineering and manufacturing time and ensure the overall reliability of your final product. Our complete solutions are as dependable as our connectors, and our service and support is unmatched. For more information on Hypertronics reliable connectors and dependable solutions, call 1-800-225-9228, email info@hypertronics.com or visit our web site at **www.hypertronics.com/wego.html**



HYPERTRONICS: RELIABLE CONNECTORS. DEPENDABLE SOLUTIONS.

HYPERTRONICS CORPORATION Hudson, MA 1-800-225-9228 www.hypertronics.com/wego.html
ISO 9001, 13485 & 14001 certified

smiths

One World. One Analyzer.

FibreChannel, Gigabit Ethernet, Serial FPDP

Protocol Analysis & Traffic Generation at its finest

Fibre Channel ■ Gigabit Ethernet ■ Serial FPDP ■ Custom Protocols

- All NEW! 4 port card; stack to 32 ports
- Switch between protocols with Protocol-on-Demand
- Traffic generation, error insertion while analyzing
- Modify custom protocols with Editor-on-Demand
- Compatible with existing Absolute Analysis hardware

ABSOLUTE ANALYSIS

www.absoluteanalysis.com
(805) 376-6048

Technology Focus:

FPDP/FPDP II Boards Roundup

FPDP PCI Card Delivers 200 Mbytes/s of Sustained Digital I/O

Not all military uses of FPDP occur in rugged environments. For example, there are applications for FPDP in back-end processing for radar and SIGINT facilities where PCI cards make economic sense. Conduant's FPDP Digital I/O Board for the PCI bus provides 200 Mbytes/s of sustained digital I/O and numerous features, such as

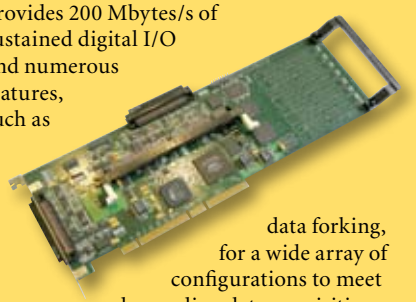

data forking, for a wide array of configurations to meet demanding data acquisition requirements. Because FPDP does not use the PCI backplane bus, there are no issues of contention and limited bandwidth. This makes the sustained transfer rate and data bandwidth known for certain during the design stage, so costly bandwidth problems do not crop up during the system integration stage.

The Conduant FPDP Digital I/O Board features dual FPDP interfaces and a universal 64-bit PCI interface. The board can be used to capture one or two FPDP input streams into a standard PC and capture 32 bits/4 bytes of parallel data at up to 50 MHz via standard FPDP Protocol. The Conduant FPDP Digital I/O Board features 132 Mbyte/s sustained throughput in data forking mode. It supports standard FPDP protocol and connectors (ANSI/VITA 17-1998). Users may choose either a single or dual FPDP input function. The Conduant FPDP I/O Interface Board operates in Windows NT/2000/XP and Linux environments. The Conduant FPDP Digital I/O Board is priced at $4,980.

Conduant
Longmont, CO.
(303) 485-2721.
[www.conduant.com].

Serial FPDP Cards Do Point-to-Point and Broadcast Links

The FPDP protocol was specifically invented to address the high-speed connection between the A/D converter of a sensor subsystem and the DSPs used in advanced image processing systems. The serial version of FPDP does that more effectively by extending the FPDP connection from 1 m to 10 km while retaining its simplicity, bandwidth and reliability. Systran, now part of Curtiss-Wright, invented Serial FPDP, and their offering is the FibreXtreme SL100/SL240 Serial FPDP data link system that connects distributed devices through the VITA 17.1-2003 FPDP communications protocol. FibreXtreme is a dedicated, Point-to-Point or Broadcast Data Link consisting of a transmitting card (source) and one or more receiving cards (destination). Cards are bi-directional, capable of performing both functions simultaneously.

The SL240 Series operates at 2.5 Gbits/s, resulting in sustained throughputs up to 247 Mbytes/s The SL100 Series, operating at a 1.062 Gbit/s baud rate, offers sustained

data throughput of 105 Mbytes/s. Other FibreXtreme features include various levels of error detection and status reporting and 32 Kbyte source and destination FIFOs to absorb bursts. The product families include PCI, PMC and CPCI source and destination cards with dual DMA controllers and PIO for data transfer over PCI. FibreXtreme data link products are available in configurations designed to handle extremely harsh environments that subject system components to additional shock, vibration and humidity and extended temperature ranges.

Curtiss-Wright Controls
Embedded Computing
Leesburg, VA.
(703) 779-7800.
[www.cwcembedded.com].

Dual Serial FPDP XMC Targets Sensor I/O

Mercury Computer Systems' Sensor I/O XMC daughtercard provides a direct interface into the RapidIO switch fabric for sensor input, enabling low-latency processing of data streaming directly from sensors. The daughtercard implements the Serial Front Panel

Data Port (sFPDP) protocol over fiber on two 2.5 Gbaud full-duplex channels.

Full system performance is enhanced, because each channel can be programmed for data distribution without processor intervention. The interface can sense signals in the data stream that indicate sensor mode changes, and route data appropriately to different processors or endpoints on the RapidIO switch fabric. The board supports connections up to 150m and a real-time latency as low as 4 microseconds. All four FPDP data modes are supported, and it provides four DMA engines with chaining and branching. Support for sFPDP is as specified by VITA 17.1-2003 and is compatible with all products supporting any subset of the VITA 17.1-2003 protocol.

The Sensor I/O XMC is software compatible with RACE++ Series RINOJ-F products, easing migration from the legacy I/O daughtercards while offering significant improvements in speed as well as configuration flexibility. The card draws approximately 6.5W of power (typical) and operates over temps of 0° to 40°C and at altitudes of 10,000 ft. A rugged version of the product is also available.

Mercury Computer Systems
Chelmsford, MA.
(978) 256-0052.
[www.mc.com].

VME Card Sports 2 Gbytes of FPDP-Accessible Memory

The point-to-point nature of FPDP makes it a popular solution for linking to large memory arrays. Exemplifying that trend, Micro Memory's MM-6496D product is a high-speed, nine-port, single slot VME memory buffer that provides up to 2 Gbytes of DRAM for FPDP, VME and RACEway systems. Data from two independent FPDP ports and two independent RACEway ports can be directly transmitted or received from any one of four "non-busy" RACEway ports at 160 Mbytes/s with a combined data rate of 640 Mbytes/s.

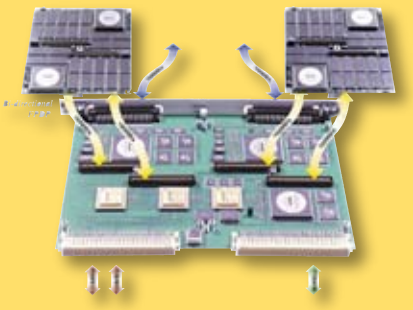

In addition to having two FPDP ports and a VME port, the card has two RACEway ports on P2 and four RACEway ports on the face of the board. Each of these four ports has an individual memory bank with up to 512 Mbytes of DRAM. Each memory bank can be accessed simultaneously, but only through its dedicated RACEway port. The FPDP ports are controlled by a full chain DMA in the RACEway interface and can be initialized by a host processor on either the VME bus or over the RACEway. Reliability is ensured by burn-in and running memory diagnostics that check operations for 48 hours while temperature-cycling boards from 0° to 60°C.

Micro Memory
Chatsworth, CA.
(818) 998-0070.
[www.micromemory.com].

VME Board Blends VXS I/O and FPDP II

Even as next-gen VITA fabric specs like VXS roll into products, the need to interface to popular interfaces like FPDP II lives on. Pentek's first VME card to blend VXS I/O with fast A/D conversion sports up to four Front Panel Data Port (FPDP II) connectors to deliver output data at up to 320 Mbytes/s each. A VXS interface provides two 1.25 Gbyte/s switched serial fabric ports to the VME backplane.

Targeting extremely high-performance DSP tasks, the Pentek Model 6822 couples its high-speed A/Ds to two Xilinx Virtex-II Pro FPGAs. A key feature of the Model 6822 is the VXS port on the P0 backplane connector, which made it one of the first boards in the industry to offer VXS. The port follows the VITA 41 specification for VME backplane switched serial fabrics and supports several different protocols through installation of one of the various fabric IP cores now available. The front panel accepts two analog inputs and delivers digital output samples over two FPDP or FPDP II connectors. An optional two-slot version of the card offers two more FPDP or LVDS connectors for digital output data. The board is available in commercial (L0), air-cooled (L1, L2) and conduction-cooled (L3, L4) versions. The Model 6822 Dual-Channel A/D VME Board with Virtex-II Pro FPGAs and VXS I/O is priced starting at $8,995.

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

Get Connected
with companies mentioned in this article.
www.cotsjournalonline.com/getconnected

Discover Condor's Avionics
Embedded Solutions

It's easy to see why we're the leader.

1553 | 429 | AFDX | MMSI

- Rugged and reliable
- High performance
- Advanced FPGA technology
- VxWorks, Integrity, LynxOS & other RTOS supported
- Test and integration tools
- Intellectual property cores
- ISO 9001:2000 certified
- Outstanding technical support

Go with the leader.

805.965.8000 • **www.condoreng.com**

CONDOR ENGINEERING

Protocol-Agnostic PMC Supports Serial FDDP

There's no doubt that FPGAs open up a treasure chest of flexibility for designers of I/O subsystems. Using FPGAs, designers can facilitate dropping in blocks of pre-integrated signal processing cores, thereby leveraging single hardware architecture across multiple applications. Following exactly that road, TEK Microsystems came out with the JazzFiber PMC, the first protocol-agnostic fiber optic PMC I/O module optimized for both streaming I/O and signal processing applications. The JazzFiber Quad Serial FPDP PMC module was the first member of Tekmicro's new family of JazzFiber FPGA-based multiprotocol fiber optic I/O modules. The JazzFiber PMC combines the advantages of the ANSI/VITA 17.1 Serial FPDP interconnect with highly integrated FPGA technology.

The use of a common FPGA architecture and software API allows applications to easily migrate between different JazzFiber solutions. The first JazzFiber protocol core supports ANSI/VITA 17.1 Serial FPDP. The JazzFiber PMC provides four fiber optic transceivers operating at up to 3.125 Gbits/s each, which can be configured as four independent interfaces or combined into a single 4x link. The JazzFiber PMC supports both PCI and PCI-X protocols at up to 133 MHz. A full gigabyte of onboard DDR SDRAM allows deep buffering of streaming data at the full 1 Gbyte/s data rate. Pricing starts at $7,995 for two-channel models and $9,995 for four-channel models in single unit quantities.

TEK Microsystems
Chelmsford, MA.
(978) 244-9200.
[www.tekmicro.com].

Data Acq Board Offers Rugged FPDP II Features

The "Front Panel" part of Front Panel Data Port is by no means the most important part of the technology. In fact, for rugged designs, rear I/O is preferred. With that in mind, ICS—part of Radstone Embedded Computing—crafted a 16-channel data acquisition card that combines rear I/O with onboard signal conditioning circuitry in a rugged, single slot package. Called the ICS-8145, the board is designed for harsh environment, confined-space high-frequency sonar, wet end of towed arrays, torpedoes, autonomous underwater vehicle (AUV) and unmanned underwater vehicle (UUV) applications.

The ICS-8145 is designed around the rugged 6U VME form-factor and features 2.5 Msamples/s oversampling ADC, gain, anti-alias filtering and clock generation resources on the card, eliminating the usual need for auxiliary signal conditioning equipment. Because the card is expected to be deployed in space-constrained areas, all I/O is routed through the backplane. The differential analog inputs are routed though VME P0, while configuration and control is carried out via the VME 64 interface. An FPDP II interface (on the VME P2 connector) allows data from multi-card systems to be easily organized and collected. The inclusion of up to 8 Mbytes of onboard storage means that fast transient captures can be read back over the VME without data loss. Availability of the ICS-8145 is immediate.

Radstone-ICS
Ottawa, Canada.
(613) 749-9241.
[www.ics-ltd.com].

PMC Board Supports the P4DP Variant of FPDP

The Front Panel Data Port (FPDP) interface is a popular data transfer solution for many military applications. But some applications need a bit more performance and can't support the complication of supporting FPDP cabling. Along those lines, Thales Computers unveiled its PMC-HTLK advanced high-speed serial data acquisition I/O mezzanine card.

Designed to meet the needs of the compact, real-time systems running on defense and

aerospace applications, the PMC-HTLK offers high-speed data acquisition and point-to-point data transmission. PMC-HTLK is a 64-bit PCI mezzanine card (PMC) that is capable of moving high-speed data to and from PCI-accessible memory using a HotLink II transmitter channel, a HotLink II receiver channel or a 32-bit wide parallel data port P4DP. P4DP is a 32-bit rear-panel variant of the front panel data port (FPDP) that runs at a software-programmed clock rate of up to 40 MHz. PMC-HTLK supports three speeds—high speed (1 Gbaud) and low speed (350 Mbaud and 320 Mbaud)—and is available in commercial and rugged air-cooled and conduction-cooled versions. Pricing for the PMC-HTLK starts at $2,690 subject to specifications.

Thales Computers
Raleigh, NC.
(919) 231-8000.
[www.thalescomputers.com].

XMX/PMC Board Offers Quad Serial FPDP

FPDP offers many advantages as a point-to-point data link, and Serial FPDP does the same only faster. VMETRO expanded its range of high-performance PMC/XMC modules with the SFM Quad Serial FPDP module. The SFM supports up to four simultaneous serial FPDP (VITA 17.1-2003) channels. Until now, Serial FPDP cards have typically had just one channel. This new four-channel interface card provides a higher level of functional density without creating a bottleneck getting the data to and from the baseboard. The functional density and high performance is especially important for high-performance data recorders, high-channel density sensor arrays and high-end DSP systems. The simplicity and wide support for Serial FPDP makes it ideal for a wide range of real-time embedded computer solutions.

In order to achieve optimal performance, VMETRO implemented the SFM with separate DMA controllers for each channel. The SFM PMC module supports PCI-X data transfers at speeds up

to 133 MHz. The SFM XMC module supports PCI Express via the XMC connectors and provides the full 2.5 Gbit/s data rate per channel. Using PCI-X and PCI Express this way enables more than one sFPDP transfer to happen simultaneously.

VMETRO
Houston, TX.
(281) 584-0728.
[www.vmetro.com].

AFDX Solutions

Hardware and software tools for Aircraft Data Networks

MIL-STD-1553
ARINC 429
ARINC 708
ARINC 717
ARINC 664/AFDX
CSDB
Discretes
Serial
Others

PCI
cPCI/PXI
PMC
PCMCIA
VME
PC/104
OmniBusBox
USB/Ethernet
Handheld

AFDX/ARINC 664 products from Ballard meet the demanding requirements of the new Ethernet based Aircraft Data Networks. They provide the advanced features you need for developing, testing, and simulating your AFDX products and systems. Intelligent hardware interface cards minimize host processing and are available for PMC, PCI, and cPCI. Run your own applications or use with CoPilot ®AFDX as a full featured databus analyzer/simulator.

Call (800) 829-1553 or email sales@ballardtech.com

Ballard Technology

Powerful solutions

Avionics databus tools & interfaces **www.ballardtech.com**

Can Your Computer Take This?

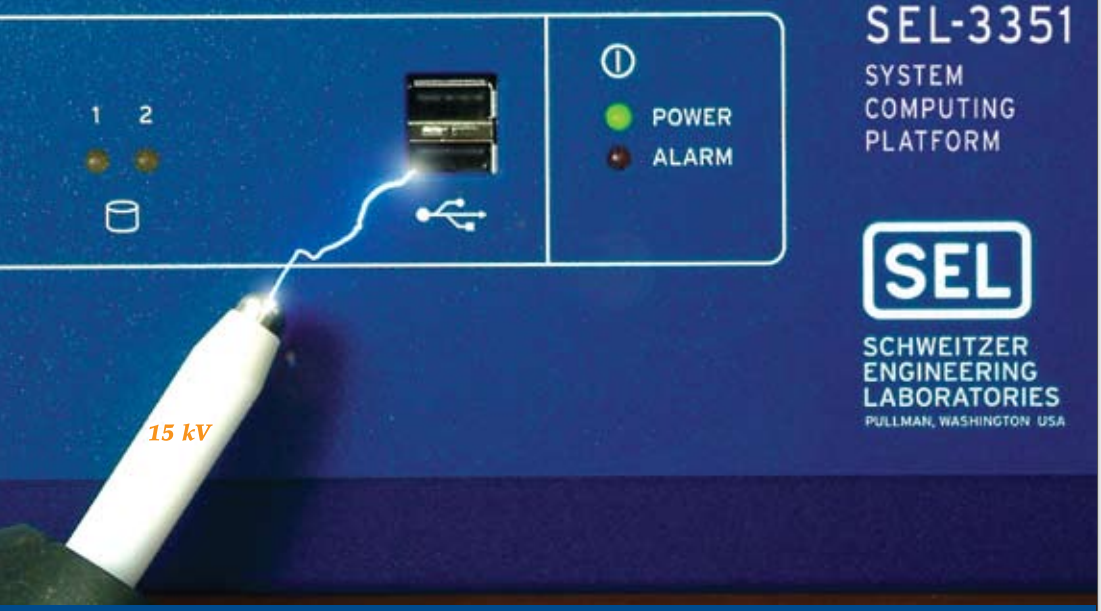

Keep the Plant Running

- Reliable for mission-critical applications, with no moving parts and no fans
- Innovative best-in-class watchdog processor
- Hardened to meet IEEE and IEC environmental, electrical surge, and fast-transient immunity tests
- Flexible hardware and software choices include Ethernet ports, 16 serial ports, ac/dc power, operating systems, and more


Heat (+75°C)

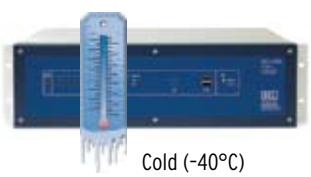
Cold (-40°C)


Vibration

SEL rugged computers are designed to perform in extreme conditions.

Quality. Reliability. Customer Service.

- Worldwide, ten-year product warranty—no questions asked!
- Worldwide technical support with products in over 100 countries
- Commitment to making electric power safer, more reliable, and more economical

WORLDWIDE 10 years WARRANTY

SEL SCHWEITZER ENGINEERING LABORATORIES

Serial and Wireless Encrypting Transceivers starting ***under $700***

Fiber-Optic Transceivers starting ***under $100***

Accurate Satellite-Synchronized IRIG-B Clocks starting ***under $500***

Large-Display Digital Clocks starting ***under $650***

GSA Schedule Contract: GS-07F-0123N

Many SEL products are GSA approved. Contact your sales representative for pricing and delivery options.

Visit **www.selinc.com/cots5** to learn more about SEL off-the-shelf solutions.

SCHWEITZER ENGINEERING LABORATORIES, INC.
www.selinc.com
info@selinc.com
509.332.1890

SEL

Effective Tools Smooth Serial FPDP Bus Analysis Task

At first blush, Serial FPDP seems simple and easy to test. But there are a lot of complexities, which only effective analyzer tools can tame.

Nigel Brownlow, Director of Product Marketing
Roger Paje, Product Marketing Manager
Absolute Analysis

Serial FPDP continues to gain popularity as a general-purpose link for applications where low latency and high throughput are desirable. The protocol has won acceptance in numerous military radar, sonar and imaging applications. Serial FPDP (ANSI/VITA 17.1-2003) addresses the distance limitations of FPDP, replacing the parallel connection with a serial interface based on the Fibre Channel physical layer. Serial FPDP retains the frame format of the original standard thus simplifying the exchange of data between parallel and serial implementations. That permits the easy exchange of data from local chassis and legacy systems using parallel interfaces to remote chassis through a Serial FPDP connection. Table 1 lists both the serial data rates and data transfer rates of Serial FPDP.

The addition of an optical connection to the physical layer provides better noise immunity and extended range with distances of up to 10 km being supported by the standard. Future versions of the protocol will support data rates of up to 10 Gbits/s and will be standardized as ANSI VITA 17.2.

Using Serial FPDP, it is essential to have the correct tools to complete and maintain the project. At first glance, the Serial FPDP protocol may appear to be relatively simple and the need for test tools minimal. Yet, like most modern standards, Serial FPDP provides a number of options and variants that complicate matters. Many remember that even for something as "simple" as the RS-232 standard, a breakout box was often essential in trying to get two systems to work together.

Point-to-Point Links

In its simplest form, Serial FPDP consists of a point-to-point connection between two devices with a single transmitter connected to a single receiver. In this instance traffic flows in one direction only from the transmitter to the receiver (Figure 1). Typical applications would include a sensor connected to a recording device. While this sounds simple enough there are a number of important issues.

An optical power meter is a useful addition to the toolkit, but it will only diagnose excess power loss. In contrast, a good

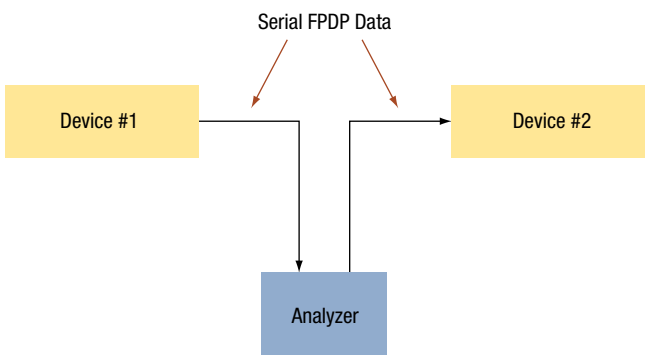

Figure 1

Using the same physical layer scheme as Fibre Channel, a Serial FPFP datalink is tested by generating worst-case noise, power and jitter test patterns and measuring the error rate on the returned data. Analyzers that implement this facility can be used to test both cables and devices for their ability to support link traffic at the desired speed. A pass using just the cable would point to the cable plant being acceptable and the problem being at the transmitter or receiver of the device at either end of the link.

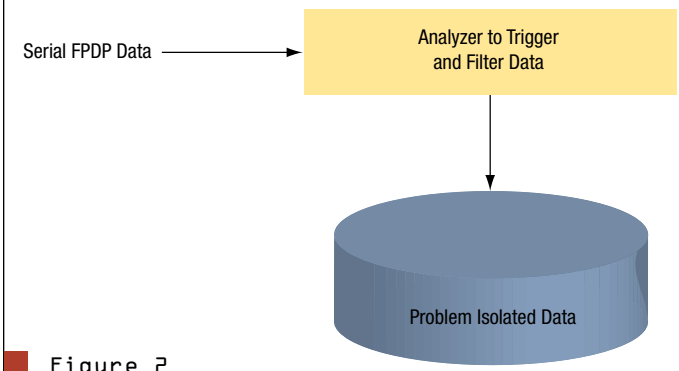

Figure 2

Triggering and filtering capabilities are key capabilities in Serial FPDP analyzers. Filtering is used to maximize capture time by removing any unwanted information. And triggering ensures that the problem is captured within the buffer.

protocol analyzer will provide indicators of link status and error conditions that are being detected. Using the analyzer, it is possible to determine that light is present, both devices are working at the same speed, the link is transmitting data and that the quality of the data is good (no encoding errors). If errors are present then noise, power or jitter may be the cause and the analyzer can be used to trace these issues to devices or cable plant.

Testing the Link

The T11 standards body, creators of the physical layer standards shared by Serial FPDP, defined standards for testing the physical layer connection. That's done by generating worst-case noise, power and jitter test patterns and measuring the error rate on the returned data. Analyzers that implement this facility can be used to test both cables and devices for their ability to support link traffic at the desired speed. A pass using just the cable would point to the cable plant being acceptable and the problem being at the transmitter or receiver of the device at either end of the link.

Serial FPDP also allows chained and loop topologies to be supported. Data is transferred from the input to output of each device using an option called Copy Mode. If Copy Mode is supported by these devices, then the same analyzer function can be used to test the individual devices and even the entire system—the cable plus connected devices. Serial FPDP limits the number of devices that can be connected in these topologies to three or six depending on the implementation due to accumulated jitter issues. The test will expose any problems along those lines.

Serial Data Rates	Corresponding Data Transfer Rates
1.0652 Gbits/s	105 Mbytes/s
2.125 Gbits/s	210 Mbytes/s
2.5 Gbits/s	247 Mbytes/s

Table 1

Listed here are the serial data rates and data transfer rates supported by Serial FPDP.

As well as testing Serial FPDP devices for their ability to accurately receive and transmit a signal, some analyzers can emulate a Serial FPDP data flow. This allows the receiving devices to be tested for their ability to process the system data. Some of the issues that may be encountered involve CRC, flow control and buffer size problems.

Detecting CRC Issues

Using CRC in the frame to protect the data is optional in Serial FPDP. Since the CRC will be indistinguishable from the data, an error in CRC handling will result in data corruption. An analyzer can be used to detect such an error. Perhaps the easiest way to do this is to view the empty data frames, as the standard requires CRC to be present even on these empty frames. If present, a single word will appear in every one of these empty frames, while implementations without CRC will be completely empty.

The choice of buffer size to implement on the receiver can have a dramatic effect on the performance of the entire system. If the transmitter generates data faster than the receiver can process it then the receiver buffer will begin to fill. Without flow control it will overflow resulting in data loss. With flow control it will stop the transmitter from sending data and either limit the performance of the system, or if the transmitter has no buffering capability, result in data loss. Analyzers with traffic generation can be used to simulate the traffic flows expected from the system, including bursty traffic and see if the receiving device is capable of meeting the performance criteria.

Flow Control Issues

Flow control is also optional in Serial FPDP and, if implemented, works through a STOP/GO word at the end of every frame. If a STOP is requested and the transmitter has flow control turned off, the transmitter can ignore the flow control character and continue transmitting, possibly resulting in a buffer overflow at the receiver. Analyzers that can simultaneously generate traffic and capture the transmitted and received data can be used to emulate the STOP condition and see both the transmitter's response and response time.

In some projects multiple Serial FPDP links are used to increase data rates or provide some degree of resilience. Serial FPDP can also be used in a switched topology. To analyze timing issues across such implementations we need a tool that time synchronizes across multiple capture channels. As the number of ports required in such implementations can often exceed the number of channels on a single analyzer card, cross-card synchronization is an important requirement in these environments. The combination of this feature with a multi-user capability can allow a multichannel device to be used in a smaller configuration by a number of users and then used in its larger configuration when required.

Capturing and Analyzing Serial FPDP Data

Capture memory size has long been a point of competition between analyzer vendors. In reality, Gbit speed links to a huge data

The leader in **Rugged Computing Technology**

Kontron

www.kontron.com

MIL-STD-810F
Shock & Vibration

IP65 Available Option

Rugged Performance

Custom Solutions

MilPAC Plus
- Fully Sealed Case-Dust, Drip and Rain Proof
- AC and DC Operation Option
- 14.1 Sunlight Viewable Display Option
- 3 Full Length PCI Slots

ECM
- Intel® Pentium® M w/802.11b/g
- Compact Embedded Module Design
- Shock Mounted Removable Hard Drive
- Extended Operating Temperature Option

TransPAC
- Rugged Battery Operated
- Customer Military Connections
- Expandable, Portable & Detachable
- MIL-STD 461E

Kontron is leading the way for customer-driven solutions that offers competitive, innovate and rugged design alternatives. Backed by a global network of engineering expertise and resources, Kontron is providing exceptional products without compromising quality. Our solutions provide designed benefits that deliver powerful performance and the most reliable state-of-the-art ruggedized portables systems. Kontron—your solutions partner for extraordinary challenges.

To see our complete rugged mobile product line, please visit www.kontron.com/mrd

Intel® Communications Alliance
Premier Member

GSA Contract # GS-35F-4082D

Computer On Modules	Blades & Mezzanines	CPU Boards	Systems	**Mobile Rugged**	Custom Solutions

kontron

1-888-294-4558 - sales@us.kontron.com — EMEA: +49 0 800 7253758 - ASIA: +886 2 2910 3532

Windows Embedded Windows Server System

OEM Technical Seminars

presented by Microsoft

Everyday developers are choosing Microsoft® Windows® Embedded for their next great device.

See for yourself how Microsoft Windows Embedded can help you optimize device design, enhance reliability and integration, and improve time to market and development costs.

These sessions include a presentation and technical demonstration of the following Microsoft Windows Embedded Products, plus time to network at the exhibition.

- Windows CE 5.0
- Windows XP Embedded with Service Pack 2
- Windows Embedded for Point of Service
- Windows Embedded Server Products
- Special Guest Presentation by intel

Win a FREE Portable Media Center!

register today www.microsoftembedded.com

Attend this event and receive a FREE evaluation copy of Windows XP Embedded or Windows CE.

2006 Strategic Locations

Mountain View, CA January 24	Reston, VA April 11
Redmond, WA February 24	Boston, MA May 23
Chicago, IL March 21	Atlanta June 20

Exhibitors

Windows Embedded Partners

ADVANTECH Applieddata.net Ardence ICOP LOGIC

Microsoft Authorized Embedded Distributors

Arrow Electronics, Inc. AVNET Bell Micro bSQUARE

*Content, topics and exhibitors may vary by city. Please visit www.microsoftembedded.com for complete information.

an RTC Group production

buffer just gives the problem a bigger place to hide. Far more important is the triggering and filtering capability of the unit. Filtering is used to maximize that capture time by removing any unwanted information (Figure 2). For example in Serial FPDP, empty data frames are constantly transmitted across the link. These frames contain no data and can quickly fill a buffer of any size.

Triggering ensures that the problem is captured within the buffer. For example, if we experienced buffer overruns and suspect a device is not honoring the STOP flow control request, it would be essential to be able to trigger on the STOP. However this is insufficient. If the analyzer stopped immediately and the STOP was detected, we would not see the effects on the transmitter. Therefore, the ability to position the trigger point within the trace and the ability to specify a time or number of events that will be captured after the trigger are essential features. With these capabilities most problems can be dealt with in 32 Mbytes or less, avoiding the need to sift through millions of data frames.

Occasionally a project will require storing several hours of data. In these cases even the largest capture buffers on the market will not suffice. In these cases a stream to disk function allows analyzers to copy the captured data to a disk array rather than the capture buffer, providing several hundred gigabytes of capture capability. Note that it is essential to have software assistance in sorting through this data to find the error condition.

Upper Layer Protocol on Serial FPDP

Because of the general-purpose nature of the Serial FPDP link, proprietary application layer protocols are typically used on top of the standard link layer. It's often desirable to decode those protocols for the purposes of detecting errors and improving performance. Sometimes security or issues of intellectual property make it not possible to divulge these protocols to a vendor for implementation on their platform. A protocol editor will allow the user to specify their own decodes to be used by the test tool. The most effective solutions also allow the use of these definitions as filters and triggers as well as decodes.

Several projects using the Serial FPDP protocol also involve storage to disk using Fibre Channel protocol. While Serial FPDP and Fibre Channel share common roots they are completely different protocols. To avoid buying a second test tool look for an analyzer that supports multiple protocols, the more the better. Each individual port pair should be capable of being set to use a different protocol to avoid having to purchase a second analyzer interface to cope with multiple protocols.

Careful choices and efficient use of test tools can be significant factors in reducing project support costs. In Serial FPDP, as well as other protocols, multi-function tools can reduce the need for Power Meters, Bit Error Rate Testers, Protocol Analyzers and Traffic Generators. That tool integration leads to faster problem resolution and improved project completion time. ▮▮

Absolute Analysis
Newbury Park, CA.
(805) 376-6048.
[www.absoluteanalysis.com].

Products

Get Connected with companies and products featured in this section.
www.cotsjournalonline.com/getconnected

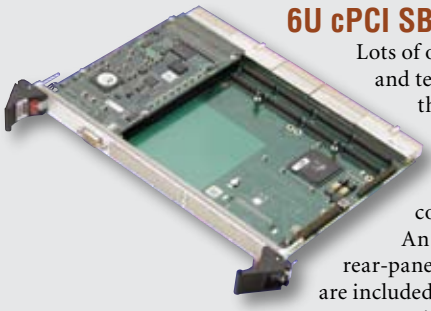

6U cPCI SBC Targets Avionics Apps

Lots of onboard compute resources along with flexible, comm-specific I/O make an SBC well suited to avionics and telematics applications. The new 6U single-slot CompactPCI D5 PowerPC SBC from Men Micro has all of this, and it operates in harsh environments from 0° to 60°C.

Based on the PowerQUICC III, the D5 features either the MPC8540 or MPC8560 CPU, dissipating only 7 to 8W. The MPC8560 features several interfaces specific to telecomm systems, such as ATM, E3/T3 and HDLC. I/O includes two slots for PMC mezzanine cards that support both front- and rear-panel connectors, and a 32/64-bit, 33/66 MHz PCI-to-cPCI bridge.

An onboard Altera Cyclone FPGA is included for implementing several functional cores that are available as rear-panel I/O. Up to 4 Gbytes of ECC DDR SDRAM main memory and 1 Gbyte or more of NAND flash memory are included for program or data storage. Two Gigabit Ethernet ports and one Fast Ethernet port are implemented as rear connections. Comprehensive board support packages are available for Linux, VxWorks and QNX as well as Men Micro's own BIOS for PowerPC processors, MENMON. Pricing starts at $2,154 for single units.

Men Micro, Lago Vista, TX. (512) 267-8883. [www.menmicro.com].

MIL-STD Processor Platform Targets Ground Vehicle Apps

Systems that must function onboard military vehicles or aircraft must meet strict requirements for shock, vibration, humidity and extreme temperatures. A rugged processor platform series from Parvus meets MIL-STD-810F requirements for shock/vibration, temperature, humidity and impact compliance.

The DuraCor 810 integrates a fanless, low-power x86-compatible processor with a vehicle-class power supply in a shock/vibration-isolated card cage with a hardened finish. Memory includes 128 Mbytes of soldered-on SDRAM and a 512 Mbyte flash disk. PC/104, 100 Mbit/s Ethernet, USB, serial, video, audio and IDE interfaces are provided. Locking MIL-grade circular connectors include hardened RJ-45 Ethernet and USB. Power is supplied from a 50W, 8-40 VDC input DC/DC converter with input protection and automotive transient voltage suppressor. Operating temperature range is -40° to +70°C.

Options include AMD Geode or Intel Celeron or Pentium III processors, integrated cards such as MIL-STD-1553, GPS, FireWire and PCMCIA, an isolated 100W DC/DC power supply and alternate IDE or CompactFlash mass storage. The unit ships with Linux and is hardware-compatible with Windows XPe, Windows CE, QNX and VxWorks. Pricing for the DuraCor 810 starts at $3,499.

Parvus, Salt Lake City, UT. (801) 483-1533. [www.parvus.com].

PMC Module Works with Multiple Networking Protocols

High-performance communications and networking military systems must do a lot in harsh environments. A new conduction-cooled communications controller from Extreme Engineering Solutions, based on the PowerQUICC II, supports multiple protocols and is software-configurable.

Based on the MPC8270 PowerQUICC II running at up to 450 MHz, the Xport2001 supports fully channelized HDLC and transparent protocols over four software-configurable T1/E1/J1 interfaces. Military system designers can change the type of port or protocol support on-the-fly, eliminating the need for manual dip switches or jumpers. CSU/DSU support plus optional Signaling System 7 (SS7) software is included. Memory provided is 32 to 256 Mbytes of DDR SDRAM and 16 to 64 Mbytes of flash.

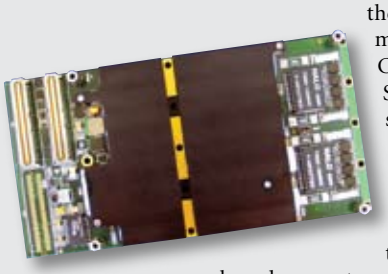

Power consumption is less than 4W. Linux and VxWorks board support packages are available. Single quantity pricing starts at $3,995, with OEM pricing below $2,000.

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

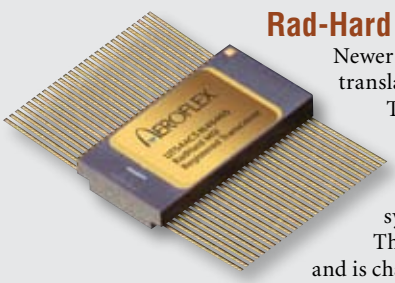

Rad-Hard MSI Logic Transceiver Has Voltage Translation

Newer subsystems used in aerospace applications need transceivers that are not only rad-hard, but also offer voltage translation. A new MSI Logic transceiver from Aeroflex withstands 100 Krads and is both 5V and 3V compatible.

The 16-bit UT54ACS164646S MultiPurpose Registered Transceiver is a high-speed, voltage translating, latchable bus transceiver that offers cold- and warm-sparing with low power consumption. It provides synchronous two-way communication and signal buffering. Applications include memory and microcontroller interface logic, backplane interface logic and voltage translation in mixed-power supply systems.

The UT54ACS164646S withstands ionizing doses of 100 Krads (Si), is immune to single event induced latch-up and is characterized over a temperature range of -55° to +125°C, making it suitable for satellite applications. It is QML Q and V compliant. Prototypes will be available in Q3 2006 and production units in Q4. Pricing for the QML Q version in lots of 100 is $400 each.

Aeroflex, Colorado Springs, CO. (800) 645-8862. [www.aeroflex.com].

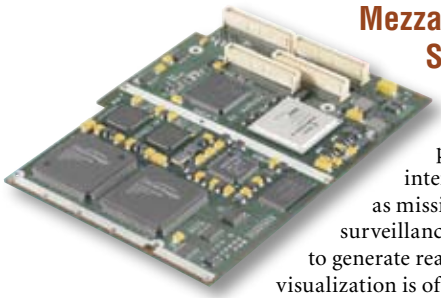

Mezzanine Supports 5 Simultaneous Video Inputs

For high-performance, video-intensive applications such as mission computers and surveillance systems, the ability to generate real-time 3D terrain visualization is often accompanied by the need for multiple video channels. A video input mezzanine card from Radstone Embedded Computing supports five simultaneous video inputs—two RGB, two TV and one DVI—with video input resolutions of up to SXGA (1280 x 1024) and digital video input resolutions up to UXGA (1600 x 1200).

Designed to complement the Octegra3 video and graphics processor, the VIM2 Video Input Mezzanine card features 539 Mbyte/s aggregate digitized video bandwidth, up/downscaling of video inputs with programmable filter and fully independent X and Y scaling. Latency can be as low as 15 milliseconds, less than one video frame, making it well suited for applications where state-of-the-art situational awareness is a requirement.

Four fully independent scalers are implemented in an FPGA. This allows substantial flexibility in configuring solutions. Scaling algorithms, appropriate to diverse input formats, are user-selectable. Support for motion-adaptive deinterlacing as well as weave deinterlacing provides flexibility and the highest possible image quality. The VIM2 is available in five ruggedization levels. Price for a single, level-one unit is $3,871 in OEM quantities.

Radstone Embedded Computing, Towcester, UK.
+44 (0) 1327 359444. [www.radstone.com].

VITA 46 Chassis Has 12-Slot Mesh Backplane

The upcoming VITA 46 and VITA 48 standards, being developed and defined by the VITA Standards Organization (VSO), will define the next major evolutionary advance in open architecture, standards-based embedded computing. A new VITA 46 chassis has been developed by Elma Electronic with a 12-slot mesh backplane.

The Elma Electronic VITA 46 ATR chassis features integrated front-to-rear airflow. It uses 120 VAC or 28 VDC power supplies with bottom patch panel access to the P2 section of the backplane. The Elma Bustronic 12-slot VITA 46 hybrid fabric backplane is a general-purpose VITA 46 implementation that supports three legacy VME64x slots and nine VITA 46 slots. The VITA 46 bus segment provides two 4-slot meshed clusters for high-performance blade computing or pipeline graphics processing. The VMEbus originates in the 3-slot legacy VMEbus segment and is continued across the P2 connectors of the VITA 46 bus segment.

The high-density MultiGig connectors in each slot have been interconnected so that each of the two 4-slot clusters provides a choice of star, mesh or ring topologies. Pricing for the VITA 46 ATR starts at under $4,500, depending on volume and features. Lead time is 6 weeks ARO.

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

Controller System Protects from AC Current Overloads

Along with the U.S. Military's growing dependence on embedded computing and networks, comes an increased importance in reliable protection from power failures and overloads. Feeding that need, American Aerospace Controls now offers an AC current controller that provides exceptional protection from system failure and current overload damage. The non-intrusive, isolated current-sensing 882 Series incorporates an output relay switch with an LED to indicate on/off status, which enables easy and accurate detection of AC current overload or underload to prevent damage to the system.

A 2A relay in the system enables the controller to turn on an alarm, shut down an electrical system and provide automatic warning alerts, making the new controllers ideal for multiple current level detection applications. The new controllers provide exceptional accuracy of ±0.5% over an extended frequency range of 50 Hz to 400 Hz and feature a field-adjustable current trip point from 20% to 100% of the full scale level range as well as high stability over the operating temperature range of -20° to +70°C (-4° to +158°F). Storage temperature is -40° to +85°C (-40° to +185°F). Pricing for an 882 Series controller starts at $147 per unit.

American Aerospace Controls, Farmingdale, NY. (631) 694-5100 [www.a-a-c.com].

PC/104 SBC Blends Fast Boot-up, Extended Temps

Many mission-critical applications can't tolerate the usual slow boot-up cycles of typical computing systems. Tackling that problem, the SBC1491ET from Micro/sys is a full-featured PC/104 computer that has a sub-five second boot-up time using an industrial BIOS or DOS operating system. The board can operate in a -40° to +85°C environment and includes standard PC features such as SVGA, dual serial ports and 10BASE-T Ethernet support. Onboard memory includes 64 Mbytes of RAM and up to 576 Mbytes solid state flash.

The SBC1491ET uses a 486/586-compatible STPC Atlas processor, with operating speeds from 120 MHz to 133 MHz, on-chip cache, 64-bit DRAM access, hardware floating-point support and AT-compatible interrupt, timer and DMA controllers. The SBC1491ET has two serial ports, COM1, COM2, SuperVGA, keyboard and mouse. The SuperVGA includes hardware acceleration and drives CRT monitors with resolutions up to 1280 x 1024. The 10 MHz Ethernet provides fast, reliable communications. I/O expansion is available by using the popular onboard PC/104 connector. The basic SBC1491ET starts at $512 in single quantity. A commercial temperature version that operates from 0° to +70°C is also available starting at $455.

Micro/sys, Montrose, CA. (818) 244-4600. [www.embeddedsys.com].

Pentium M Rides Conduction-Cooled cPCI SBC

At one time, a military embedded computer requirement for high performance, low power consumption and operation in extreme temperatures could only be met with a costly solution. Breaking that mold, General Micro Systems has introduced the Rock (CC61x), a new, rugged conduction-cooled cPCI SBC. The board boasts it is the only rugged cPCI SBC to provide full system health monitoring and reporting, meeting all PICMG 2.9 specifications.

The highest-performance version is driven by the 1.4G Hz Pentium M-738 processor with 2 Mbytes of L2 cache, or with the Celeron M-373 with 512 Kbytes of L2 cache. The health monitoring/reporting function is crafted using a microcontroller, and an FPGA performs the Baseboard Management Controller functions per PICMG 2.9, reports the results of the Built-in Test and Extended Built-in Test, and monitors and controls the baseboard temperature. The CC61x comes with up to 2 Gbytes of memory and up to 16 Gbytes of flash. The board sports dual Gbit Ethernet with a TCP/IP offloading engine on a PCI-X (66 MHz/64-bit) bus. Other key I/O on the card are two SATA ports, four USB 2.0, 512 Kbytes of user/boot flash and two serial ports. The Rock is also available in a convection-cooled version (0° to +50°C). Pricing for the conduction-cooled version starts at $4,370 (100s).

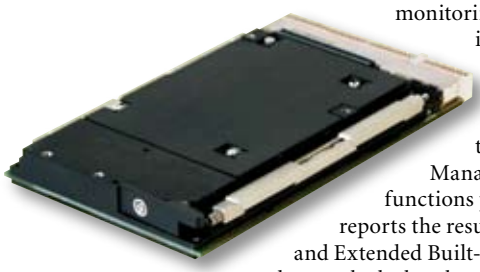

General Micro Systems, Rancho Cucamonga, CA. (909) 980-4863. [www.gms4sbc.com].

PMC Does Dual-Channel Radar Acquisition and Processing

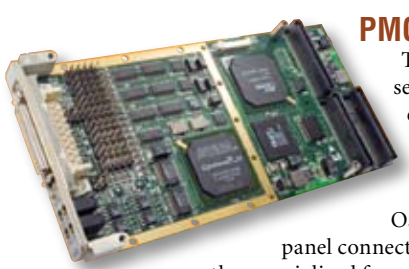

The forward march of semiconductor integration has enabled many functions that once required many separate boards to now reside in one mezzanine module. Curtiss-Wright Controls Embedded Computing exemplifies that trend with the Osiris, a new high-performance, dual channel radar interface board. Osiris, available in PMC mezzanine card or half-length PCI card formats, is a cost-effective, open standard-based radar signal processing solution ideal for use in VME and CompactPCI embedded systems (PMC version) or in PC-based workstations (PCI version).

Osiris supports two independent radar channels inputs at rates up to 16k samples per return from its front panel connector. Inputs supported include radar video, triggers and azimuth turning data for RADDS, ACP/ARP and other specialized formats. The card supports sampling at frequencies up to 50 MHz on each of two analog inputs, which can be mixed with up to 8 digital radar inputs. Additional features on Osiris include dual trigger input, gain control as a function of range, range and azimuth correlation, digital filtering, a built-in test generator, dual azimuth input and interference suppression. Synchro input is supported on the PCI variant of card with the addition of an optional synchro-adaption module. The board is priced starting at $4,500.

Curtiss-Wright Controls Embedded Computing, Leesburg, VA. (703) 779-7800. [www.cwcembedded.com].

Conduction-Cooled PICMG 2.16 SBC Draws a Mere 20W

Packing a lot of computing muscle into a small space is always a challenge, particularly so for rugged military applications where fans are frowned upon and power budgets are strict. Serving up a solution for all that, Dynatem offers the Intel Pentium M-based CRM1. The CRM1 was designed in compliance with VITA 30.1-2002, so it comes with top and bottom cooling plates that are bonded to the major components through thermal conduction, and to the heat conducting printed circuit board mechanically.

Onboard CompactFlash permits single-slot booting. I/O routed to the backplane includes an EIDE port, two Serial ATA ports, two Gbit Ethernet ports (PICMG 2.16-compatible), DVO/VGA, four USB 2.0 ports and two COM ports. Two PMC expansion sites permit system tailoring to users' application requirements. The CRM1 comes populated with 1 Gbyte of DDR-266 SDRAM with ECC and a memory bandwidth of 2.1 Gbytes/s. The 855GME offers integrated, high-performance graphics that can support resolutions up to 1600 x 1200 at 85 MHz. Pricing for the CRM1 starts at $5,300 in single quantity.

Dynatem, Mission Viejo, CA. (949) 855-3235. [www.dynatem.com].

3U cPCI Card Sports 1.4 GHz PowerPC e600 CPU

CompactPCI has become firmly entrenched as an option for defense applications, and the 3U version of cPCI provides a "here now" solution for small form-factor backplane situations. Interface Concept's latest offering along those lines is a 3U CompactPCI SBC, called IC-e6-cPCIa. This new SBC comes in addition to the existing PrPMC model (IC-e6-PMCa) and is ideal for use in defense, automation, network and imaging applications.

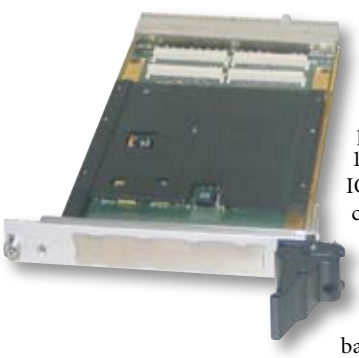

This new SBC is designed around Freescale's PowerPC e600 processors—the MPC7447A-1GHz or MPC7448-1.4GHz). The IC-e6-cPCIa board integrates many communication functions, such as two Gbit Ethernet channels, two high/full speed USB 2.0 ports and two multi-purpose serial controllers. The memory banks are made of 256/512 Mbyte DDR-ECC SDRAM, 128 Kbyte ultra-fast SRAM, 32 Kbyte FRAM, 64 Mbyte flash EPROM and up to 1 Gbyte of Nand Mass Flash. A PMC expansion slot permits the addition of a 64-bit PMC card. The IC-e6-cPCIa basic software is based on UBOOT together with a comprehensive Power-on Built-in Test and a set of maintenance tools. This board is designed to meet the most severe environments and is available in standard, extended and conduction-cooled grades. Prices start at $3,250.

Interface Concept, Briec de l'Odet, France. +33 (0)2 98 57 30 30. [www.interfaceconcept.com].

SBC for Mobile Apps Has 5 Low-Power Modes

Designers of mobile high-performance embedded systems, such as handheld and wearable computers or small unmanned vehicles, need to be able to fine-tune power consumption. The new BitsyXb SBC from Applied Data Systems features five low-power modes to help make the SBC power-stingy, as well as dynamic variable speed and voltage regulation.

The compact, 3 in. x 5 in. BitsyXb is based on Intel's 32-bit, 520 MHz XScale PXA270 CPU, with a video interface up to XGA resolution. Up to 128 Mbytes of SDRAM program memory and up to 64 Mbytes of flash memory are provided. 128 Kbytes of EPROM is included as a boot device. For expansion and connectivity, the board has a PCMCIA Type II interface, three serial ports, a USB port, an Intel QuickCapture camera sensor input bus, 10 digital I/Os, an SPI port, an I^2C bus and ADSmartIO with nine configurable inputs/outputs.

An onboard power supply has input voltage of 5V or 6-16V. The board consumes less then 1W during operation, and is ruggedized at -45° to +85°C. Windows CE .NET and Linux are supported. The BitsyXb SBC is priced in the $300s.

Applied Data Systems, Columbia, MD. (301) 490-4007. [www.applieddata.net].

3U cPCI Board Musters More Speed for the Same Power

Intel's new dual-core technology boasts almost twice the performance of Pentium processors in the same space, and with only slightly higher power consumption. A new CPU board from Kontron leverages Intel Core Duo or Core Solo processors to provide extremely high performance per watt in a 3U CompactPCI form-factor.

The CP307 is equipped with the 2 GHz Intel Core Duo T2500/1.6 GHz Core Duo L2400 or 1.6 GHz Core Solo T1300, the 945GM Express chipset and the ICH7-R Southbridge. Designed with soldered-on processor and memory for the toughest environmental conditions, the CP307 offers a 667 MHz front-side bus, up to 4 Gbytes of 533/667 MHz DDR2-SDRAM and 10.6 Gbits/s data throughput. Interfaces include two Gigabit Ethernet, up to six USB, two SATA and a CompactFlash slot.

The board comes in single-slot or dual-slot versions. The dual-slot version offers support for LPC, COM, DVI and 2.5-in. SATA. An extended temperature range version is also available. Board support packages are available for Windows XP, XP Embedded, Linux and VxWorks. Pricing starts at $2,450 in single quantities.

Kontron America, Poway, CA. (888) 294-4558. [www.kontron.com].

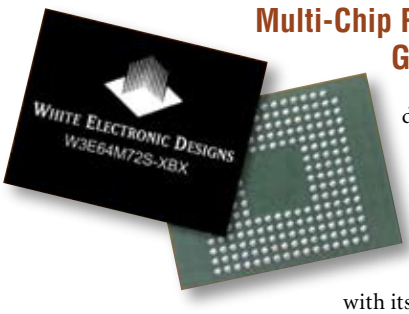

Multi-Chip Package Packs in 4 Gbits of DDR SDRAM

By packing several DRAM die onto a mini module, military system designers can pack extremely high memory densities into a reliable compact package. Doing exactly that is White Electronic Designs with its 4 Gbit (512 Mbyte) DDR SDRAM PBGA Multi-Chip Package (MCP). Organized as 64M x 72, the SDRAM is packaged in a 25 x 32 mm, 800 mm2 , 219 plastic ball grid array (PBGA). This package is designed for high-reliability applications and is available in commercial, industrial and military temperature ranges. Benefits include a 66% space savings versus a comparable density using thin small-outline packages (TSOPs), 55% I/O reduction versus TSOP, reduced trace lengths for lower parasitic capacitance and reduced part count.

WEDC's new DDR SDRAM MCP features internal pipelined double-data-rate (DDR) architecture enabling two data accesses per clock cycle. The SDRAM provides programmable Read or Write burst lengths of 2, 4 or 8, four internal banks allow concurrent operation and it features bi-directional data strobe (DQS, DQS#) per byte for transmitting and receiving data. White's 64M x 72 DDR SDRAM, designated as part number W3E64M72S-XBX, is priced at $450 (1,000s).

White Electronic Designs, Phoenix, AZ. (602) 437-1520. [www.wedc.com].

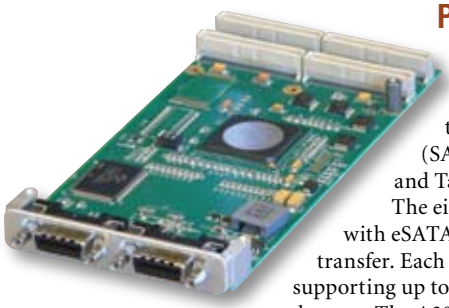

PMC Adapter Serves Up 3 Gbit/s Serial Attached Storage

The embedded computing world is smack in the middle of a transition away from parallel interconnects toward serial solutions, and the military market is riding that wave. In the storage area, Serial Attached SCSI is a favorite as a low-risk technology to ease the transition to serial storage. With that in mind, Astek has announced the A3803-PMC-01, an eight channel, 3Gbit/s Serial Attached SCSI (SAS) adapter, Host Bus Adapter (HBA) in the PMC form-factor. The A3803-PMC-01 supports Initiator and Target modes, while the A3803-PMC-01R supports Initiator and Integrated RAID modes.

The eight-channel A3803-PMC-01 offers front panel I/O with two InfiniBand connectors, plug-compatible with eSATA InfiniBand multilane x4 cabling, and PCI-X 64-bit/133 MHz support for up to 1 Gbit/s of data transfer. Each channel supports direct attach target devices (SATA or SAS), or connection to a SAS expander device, supporting up to 132 target disks per HBA. The A3803-PMC-01 supports target mode operation for builders of storage enclosures. The A3803-PMC-01R card supports on-chip RAID 0/1/1e, with support for automatic hot-spare and hot-swap of target disks for high-availability and high-reliability requirements.

Astek, Colorado Springs, CO. (719)-260-1625. [www.astekcorp.com].

CompactPCI Express

Development Systems

Available Now!

8U CPCI/CPCIe Hybrid System

Cable from a Host PC to One Stop Systems' Hybrid CPCI/CPCIe System at 10Gb per second

Ideal Platform for CPCI and CPCIe Board and Software Development at Highest Performance Available.

Interface Board to Host PC

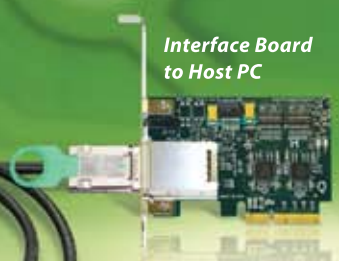

Interface Board to Host PC

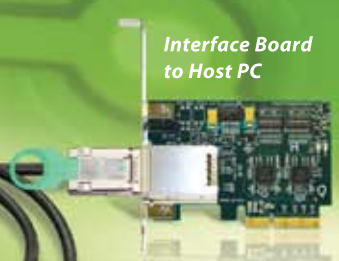

5U CPCI/CPCIe Hybrid System

MAX EXPRESS

Now available at

ONE STOP SYSTEMS

Call today.

(877) GET-CPCI

www.onestopsystems.com

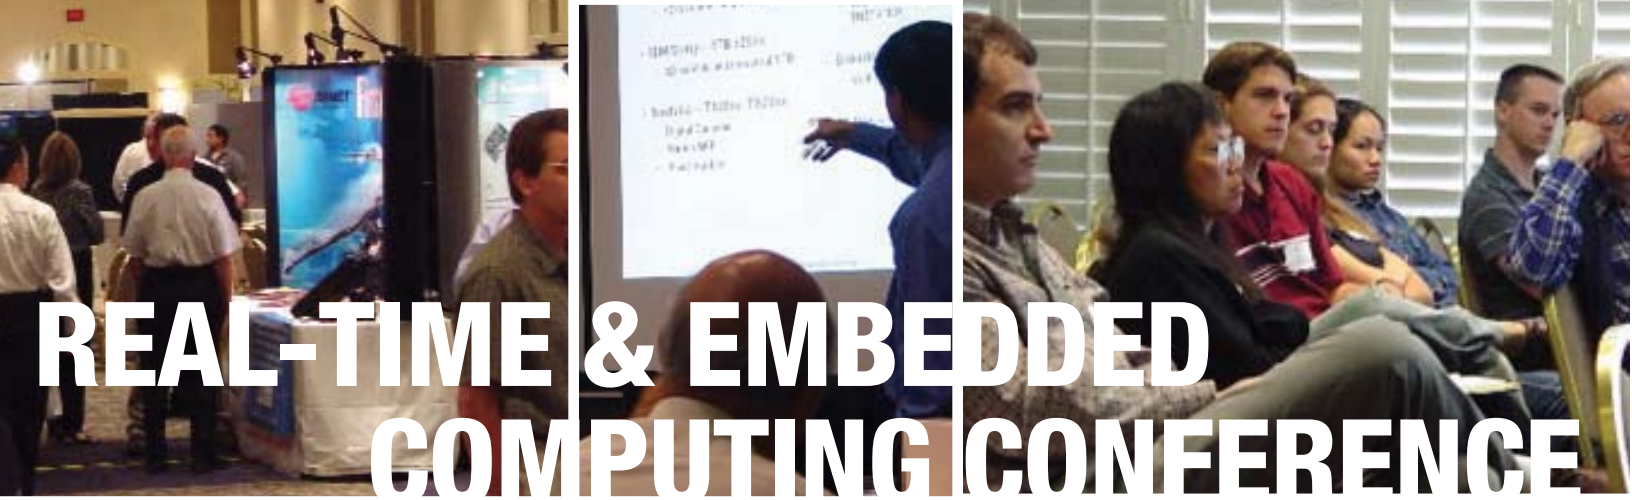

REAL-TIME & EMBEDDED COMPUTING CONFERENCE

Real-Time & Embedded computing conference

Evaluate Your Options

Register now at www.rtecc.com

Enter a World of Embedded Computing Solutions

Attend open-door technical sessions especially designed for those developing computer systems and time-critical applications. Get ahead with sessions on Embedded Linux, VME, PCI Express, ATCA, DSP, FPGA, Java, RTOS, SwitchFabric Interconnects, Windows, Wireless Connectivity, and much more.

Your Resource Opportunity

Exhibits arranged in a unique setting to talk face-to-face with technical experts. Table-top exhibits make it easy to compare technologies, ask probing questions and discover insights that will make a big difference in your embedded computing world. Join us for this complimentary event!

Sponsors

AAEON · Accelerated technology · coverity · ENEA · Green Hills Software, Inc. · LYNUXWORKS™ · QNX · SBS Technologies. · WDL Systems · WIND RIVER

2006 Locations

Dallas	Houston	Boulder	Salt Lake City
June 6	**June 8**	**June 20**	**June 22**
Detroit	Los Angeles	Taipei	Dublin
Toronto	Lucerne	Washington DC	Montreal
Calgary	Eindhoven	Shenzhen	Ottawa
Vancouver BC	Bristol	Shanghai	Seattle
San Diego	Patuxent River	Beijing	Portland

www.rtecc.com

Advertisers Index

Get Connected with technology and companies providing solutions now

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

www.cotsjournalonline.com/getconnected

Company	Page#	Website
Abbott Technologies	8	www.abbottpower.com
Absolute Analysis	64	www.absoluteanalysis.com
ACT/Technico	26	www.acttechnico.com
Aitech Defense Systems, Inc.	27	www.rugged.com
Analog Devices, Inc.	14,15	www.analog.com/perform
AP Labs	9	www.aplabs.com
Ballard Technology	71	www.ballardtech.com
Behlman Electronics	79	www.behlmancots.com
BitMicro Networks, Inc.	8	www.bitmicro.com
Condor Engineering	68	www.condoreng.com
Curtiss-Wright Controls Embedded Computing	4	www.cwcembedded.com
Data Device Corp.	23	www.ddc-web.com
Equipto Electronic Corp	25	www.equiptoelec.com
GE Fanuc Embedded Systems	7	www.gefanuc.com/embedded
Hybricon Corporation	69	www.hybricon.com
Hypertronics Corporation	63	www.hypertronics.com
Kontron America	62,75	www.kontron.com
Lippert Embedded Computers	70	www.lippert-at.com
Micro Memory LLC	88	www.micromemory.com
Microsoft Tour	76	www.microsoftembedded.com
National Instruments	17	www.ni.com
North Atlantic Industries	19	www.naii.com
Octagon Systems	2,3	www.octagonsystems.com
One Stop Systems	83	www.onestopsystems.com
Pelican Products, Inc.	18	www.pelicanoem.com
Phoenix International	81	www.phenxint.com
Pioneer Magnetics	67	www.PioneerMagnetics.com
Radstone Technology Corporation	10	www.radstone.com
Rantec Power Systems, Inc.	66	www.rantec.com
Rave Computer	13	www.rave.com
Real-Time & Embedded Comp. Conf.	84	www.rtecc.com
Red Rock Technologies, Inc.	77	www.redrocktech.com
RO Associates	24	www.roassoc.com
RTD Embedded Technologies	44,45	www.rtdusa.com/www.rtdfinland.fi
SBS Technologies	20	www.sbs.com
Schweitzer Engineering Laboratories	72	www.selinc.com
Targa Systems, L3 Communications	22	www.targasystems.com
Themis Computer	60	www.themis.com
Tri-M Systems	77,79,81	www.tri-m.com
VICOR	87	www.vicorpower.com

FPGAs: The New Matrix for Design Advertisers Index

Company	Page#	Website
Acromag	35	www.acromag.com
Aeroflex Microelectronic Solutions	38	www.aeroflex.com/radhardfpga
Altera Corporation	28, 59	www.altera.com
Annapolis Micro Systems	39	aaa.annapmicro.com
BittWare,	48	www.bittware.com
Innovative Integration	57	www.innovative-dsp.com
Interactive Circuits and Systems	49	www.ics-ltd.com
Synplicity, Inc.	40	www.synplicity.com
Technobox, Inc.	30	www.technobox.com
Themis Computer	53	www.themis.com
VPT, Inc,	58	www.vpt-inc.com

COTS Journal (ISSN#1526-4653) is published monthly at 905 Calle Amanecer, Suite 250, San Clemente, CA 92673. Periodicals Class postage paid at San Clemente and additional mailing offices. POSTMASTER: Send address changes to COTS Journal, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

Coming Next Month

June is almost on us, and that means the year's about half over. Whether that means the year is half full or half empty depends on your point of view. All we know is that we never run at half speed at *COTS Journal*. Here's how we're bringing that energy to bear for our June issue:

- *Military Batteries*. For the battlefield soldier, batteries are critical for both primary and back-up power for a wide range of communications equipment, imaging devices and weapons systems. Those demands are driving military battery manufacturers to extend the life and reduce the weight of their rechargeable and non-rechargeable battery products. Meanwhile, they're also addressing important issues such as safety and life cycle costs. Articles in this section examine the latest trends along those lines.
- *FCS Tech Update*. The Army's Future Combat Systems program ranks as one of the most complex ever taken on. Comprised of a wireless data network using advanced communications technologies, FCS links soldiers with 18 new, lightweight manned and unmanned ground vehicles, unmanned aircraft, sensors and weapons. This feature updates readers on the status of the FCS program, with a look at the central role that embedded computing architectures will play.
- *Designing for Rad-Hard Systems*. Space-based systems must be built using electronics capable of withstanding everything from intense radiation due to high-energy atoms to bombardments from neutrons and other particles. Right-sizing the appropriate level of radiation-hardening is somewhat of an art. Articles in this section explore the radiation concerns facing space designers and update readers on radiation hardened boards, subsystems as well ASICs, FPGAs and power components designed for those applications.
- *Naval I/O Update*. As the U.S. Navy works to upgrade its shipboard command and control systems and develop brand new next-generation vessels, it's making a shift from legacy systems to an open systems environment. Articles in this section update readers on technology and product trends that are critical to naval networking and I/O, such as NTDS, ATDS, TADIL A and NATO STANAG 4146 standards.

Editorial

Jeff Child, Editor-in-Chief

One thing that comes with the territory when you're a magazine editor is that a lot of people want to get your opinion about things. Because we talk to so many experts and authorities each week, even slow brains like mine pick up quite a bit of information as we go about our research and the routine dialogue we maintain with the industry our magazine covers. Case in point, I found I had a surprising amount to say when I got a call the other day from a novelist who had read an editorial I wrote last fall about the 2005 DARPA Grand Challenge.

Judging by the questions he asked me, it was clear that the novelist was interested in facts and speculation about just what is possible today with unmanned air and ground vehicles. Could a robotic vehicle be controlled by satellite? Could it execute an attack without remote control? Could it be subverted by terrorist hacking into the wireless network the vehicle is linked to? I gave him an earful, and was quite amazed at just how much there was to say on the topic.

Calling All Robot Makers

Unmanned vehicles are far from science fiction—and are no stranger on today's battlefields. Currently, there are approximately 2,000 robots in use by the U.S. military in Afghanistan and Iraq for duties such as the detection and disablement of improvised explosive devices (IEDs), surveillance and many other tasks that are deemed unsafe for human intervention. The number of robots in conflict zones is expected to double by the end of 2006 according to DoD sources. Congress set a goal in 2001 that one-third of this country's operational ground combat vehicles be unmanned by 2015. Imagine, for example, the benefits of large unmanned supply convoys trekking over vast distances through dangerous territory.

For those of you who haven't heard of it, the DARPA Grand Challenge was a $2 million prize competition sponsored by the U.S. Department of Defense to build a fully autonomous vehicle that could drive on- and off-road through the desert without human drivers. Many familiar names from the embedded computing industry donated boards and software to various competing teams in the Grand Challenge, including Hybricon, ACT Technico, Concurrent Technologies, SBS Technologies and Real-Time Innovations.

In the 2005 Grand Challenge, four autonomous vehicles successfully completed a 132-mile desert route under the required 10-hour limit, and DARPA awarded a $2 million prize to the team who developed the vehicle "Stanley" from Stanford University. The primary goal of this DoD-sponsored event was to accelerate research and development in robotic ground vehicles. Last month the PBS television show NOVA did an interesting show about the Grand Challenge. If you missed it, you can watch the hour-long program online at www.pbs.org/wgbh/nova/darpa.

In that same theme of looking outside the box for robotics expertise, the Center for Commercialization of Advanced Technology (CCAT) (www.ccatsandiego.org) in San Diego, California recently announced that their program is accepting applications from entrepreneurs, government laboratories and academic researchers who have designed robotics technologies that can be used in military and defense applications. The DoD-supported CCAT program awards funding and commercialization assistance to technologies dedicated to military, defense, homeland security and first responder operations. Award applications will be accepted through the end of this month.

Submitted technologies will be reviewed and evaluated by the CCAT program, in conjunction with the DoD Joint Robotics Program at the Space and Naval Warfare Systems Center, San Diego (SSC San Diego). The upcoming evaluation, entitled Integration of Technologies on an Unmanned Ground Vehicle for Enhanced Situational Awareness, seeks reasonably proven technologies and systems that can be integrated onto specific existing military robotic platforms. The program's primary areas of interest include miniature sensors—chemical, seismic, explosives—communications systems; autonomous operations technology; and detection and classification.

One of the questions the novelist asked me was the following: Could an autonomous robotic vehicle drive through a populated urban area? In response, I explained what a significant real-time computing challenge that would be. The robot vehicles in last year's DARPA Challenge had a hard enough time navigating a fixed route that had no moving obstacles—the vehicles raced against the clock and not each other, so they didn't even have each other as obstacles. The novelist was grateful for my insights. But only a few days after talking to him, ironically, DARPA announced its next Grand Challenge competition. And, sure enough, its next event is to be called the DARPA Urban Challenge. Scheduled for November 3, 2007, the race will feature autonomous ground vehicles executing simulated military supply missions safely and effectively in a mock urban area.

Safe operation in traffic is considered essential to U.S. military plans to use autonomous ground vehicles to conduct important missions. DARPA will award prizes for the top three autonomous ground vehicles that compete in a final event where they must safely complete a 60-mile urban area course in fewer than six hours. First prize is $2 million, second prize is $500,000 and third prize is $250,000. To succeed, vehicles must autonomously obey traffic laws while merging into moving traffic, navigating traffic circles, negotiating busy intersections and avoiding obstacles. We'll be paying close attention to the DARPA Urban Challenge. I guess fact and science fiction will be riding together. ▮▮

We're your mission support!

TALK TO US
800-496-5593

VICOR offers complete COTS power solutions. It starts with the widest range of high-density DC-DC converter modules available. Military COTS filter modules are available to complement the converter modules, offering EMI and transient compliance for both DC and AC systems. You can choose from user-configurable DC-DC power supplies based on converter modules — as well as module-based custom solutions. What's more, experienced applications engineers give you world-class technical support that stands behind you. And, all Vicor products are made in the USA.

TALK TO US, and you'll find out how you can reap the benefits of designing with Vicor high-density DC-DC converters. Faster time to market, greater power density and performance, and higher reliability. Call 800-496-5593 or go to **vicorpower.com/cj** for more information.

Always COTS

vicorpower VICOR com

CoSine™ SoC

-on-

Othello®

VME CARRIERS FOR SWITCH FABRICS

MM-15x0 & MM-16x0

- Two CoSine 2VP70™ System-on-Chips
- Two XILINX® V-4 SX55 or LX160 CoSine Companion Devices
- Two XMC/PMC sites
- 14 independent memory arrays with total bandwidth of over 40 GB/s
- Four embedded PowerPC™ computers
- Complete on-board Serial RapidIO switch fabric connectivity
- VITA 41 or VPX VITA 46 formats
- Rugged Air Cooled and Conduction Cooled versions

MM-15x0

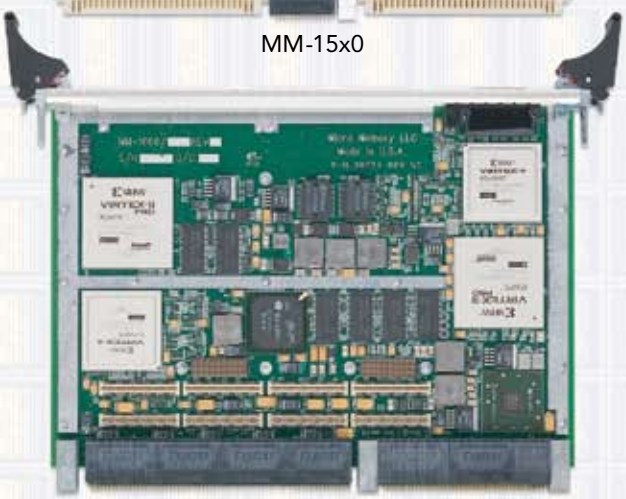

MM-16x0

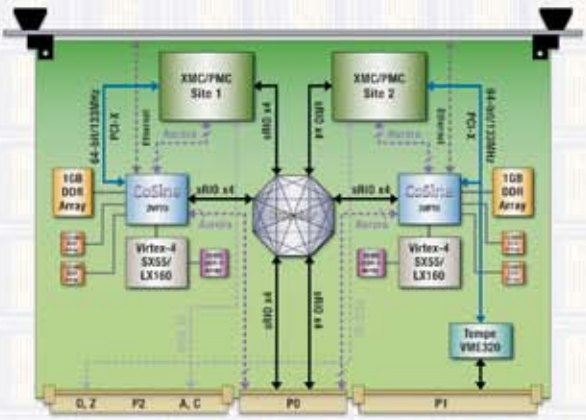

Easily the most impressive board ever built for FPGA processing.

Quite possibly the most impressive ever built for VME.

RapidIO®

micro Memory

MICRO MEMORY
9540 Vassar Avenue, Chatsworth, California 91311 U.S.A.
(US) Tel 818 998 0070 • (US) Fax 818 998 4459
www.micromemory.com • sales@micromemory.com