SIGNAL PROCESSING DESIGN

@signal_proc

From the publishers of DSP-FPGA.com

SUMMER 2015 | VOLUME 12 NUMBER 1
signal-processing.mil-embedded.com

2015 RESOURCE GUIDE

PG. 27

ANNAPOLIS MICRO
WILDSTAR 7 for OpenVPX 6U
p. 32

COTS SIGNAL PROCESSING CONTINUES FOR DESY MACHINE BEAM CONTROLS

PG. 22

CRITICAL LINK
MitySOM-5CSx Cyclone V SoC-based SOM
p. 34

SIGNAL PROCESSING

FPGA WITH HARD FLOATING-POINT — NO COMPROMISE

PG. 10

SIGNAL PROCESSING DESIGN

SUMMER 2015

@signal_proc

VOLUME 12 | NUMBER 1

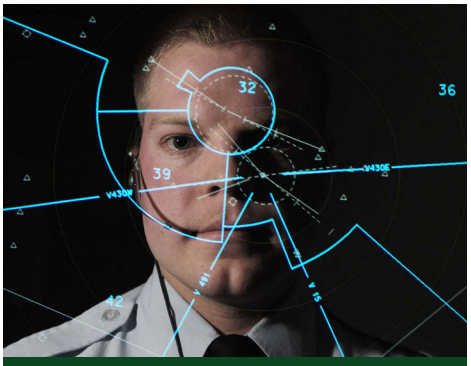

Signal processing FPGAs with hard floating-point – No compromise
By Jeff Milrod, Bittware, Inc.

» p. 10

Design of FMCW radars for active safety applications
By John Zhao, Marco Roggero, and Giorgia Zucchelli, MathWorks

» p. 18

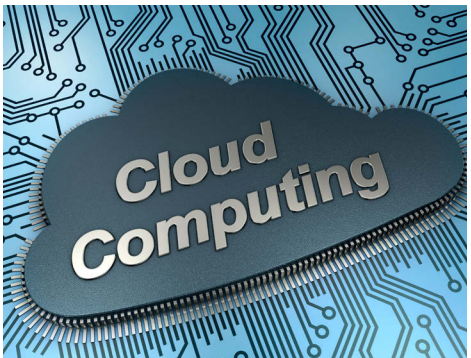

Edge-based analytics using general-purpose processors
By Nigel Forrester, Concurrent Technologies

» p. 24

DEPARTMENTS

FEATURES

On the cover

Continuing the heritage of DSP-FPGA.com, the 2015 Signal Processing Design Resource Guide features products such as the Annapolis Micro Systems WILDSTAR 7 for OpenVPX and Critical Link's MitySOM-5CSx. (Underground civil engineering work for the European XFEL laser in Schenefeld completed in 2013. Photo courtesy DESY.)

All registered brands and trademarks within *Signal Processing Design* magazine are the property of their respective owners.
© 2015 OpenSystems Media © 2015 *Signal Processing Design*

ENVIROINK
The inks used to print the body of this publication contain a minimum of 20%, by weight, renewable resources.

SIGNAL PROCESSING DESIGN

OpenSystems media.

Signal Processing Design Editorial/Production Staff

John McHale, Group Editorial Director
jmchale@opensystemsmedia.com

Brandon Lewis, Assistant Managing Editor
blewis@opensystemsmedia.com

Lisa Daigle, Assistant Managing Editor
ldaigle@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Tom Varcie, Sales Manager
tvarcie@opensystemsmedia.com
(586) 415-6500

Rebecca Barker, Strategic Account Manager
rbarker@opensystemsmedia.com
(281) 724-8021

Eric Henry, Strategic Account Manager
ehenry@opensystemsmedia.com
(541) 760-5361

Kathleen Wackowski
Strategic Account Manager
kwackowski@opensystemsmedia.com
(978) 888-7367

Asia-Pacific Sales
Elvi Lee, Account Manager
elvi@aceforum.com.tw

Regional Sales Managers
Barbara Quinlan, Southwest
bquinlan@opensystemsmedia.com
(480) 236-8818

Denis Seger, Southern California
dseger@opensystemsmedia.com
(760) 518-5222

Sydele Starr, Northern California
sstarr@opensystemsmedia.com
(775) 299-4148

Europe Sales
James Rhoades-Brown
james.rhoadesbrown@husonmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

OpenSystems Media Editorial/Production Staff

John McHale, Group Editorial Director
Military Embedded Systems
PC/104 and Small Form Factors
PICMG Systems & Technology
VITA Technologies
jmchale@opensystemsmedia.com

Rich Nass
Embedded Computing Brand Director
Embedded Computing Design
rnass@opensystemsmedia.com

Joe Pavlat, Editorial Director
PICMG Systems & Technology
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VITA Technologies
jgipper@opensystemsmedia.com

Curt Schwaderer, Editorial Director
Embedded Computing Design
cschwaderer@opensystemsmedia.com

Rory Dear, Technical Contributor
Embedded Computing Design
rdear@opensystemsmedia.com

Monique DeVoe, Managing Editor
Embedded Computing Design
mdevoe@opensystemsmedia.com

Joy Gilmore, E-cast Manager
jgilmore@opensystemsmedia.com

Brandon Lewis
Assistant Managing Editor
Embedded Computing Design
Industrial Embedded Systems
PICMG Systems & Technology
Signal Processing Design
blewis@opensystemsmedia.com

Lisa Daigle, Assistant Managing Editor
Military Embedded Systems
PC/104 and Small Form Factors
ldaigle@opensystemsmedia.com

Mariana Iriarte, Assistant Editor
Military Embedded Systems
miriarte@opensystemsmedia.com

Sally Cole, Senior Editor
Military Embedded Systems
scole@opensystemsmedia.com

Konrad Witte, Senior Web Developer

Steph Sweet, Creative Director

Dave Diomede, Creative Services Director

Joann Toth, Senior Designer

Chris Rassiccia, Graphic Designer

Corporate

www.opensystemsmedia.com

Patrick Hopper, Publisher
phopper@opensystemsmedia.com

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

John McHale, Executive Vice President
jmchale@opensystemsmedia.com

Rich Nass, Executive Vice President
rnass@opensystemsmedia.com

Wayne Kristoff, CTO

Emily Verhoeks, Financial Assistant

Headquarters – ARIZONA:
16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: (480) 967-5581

MICHIGAN:
30233 Jefferson
St. Clair Shores, MI 48082

Subscriptions

www.opensystemsmedia.com/subscriptions

Tel: 586-415-6500

subscriptions@opensystemsmedia.com

ADVERTISER INDEX

PROFILE INDEX

PROFILE INDEX

SIGNAL FILTERING

By Brandon Lewis, Assistant Managing Editor

@BrandonLewis13 blewis@opensystemsmedia.com

SpaceVPX adds fault tolerance, cuts cost for space-bound HPEC systems

In the rigors of space, fault tolerance is a prerequisite for high-performance embedded computing (HPEC) systems responsible for the transmission of mission-critical data. Faced with extreme temperatures and radiation levels, extended deployment, and the improbability of hardware maintenance, space-bound technology platforms must be designed to standards of reliability and resilience seldom matched by applications on Earth. Meanwhile, space agencies are being asked to do more with less since the effects of sequestration began to take hold in the earlier part of the decade, as are their suppliers, who are perpetually tasked with delivering components that are smaller, more powerful, and more affordable.

Though certain off-the-shelf architectures have had successful flight heritages in space, their use beyond prototyping and development has been largely spotty. On the whole, COTS technology has been viewed as incapable of offering the robustness required by an industry that has traditionally turned to custom solutions, but as the space sector becomes increasingly international and more commercialized, historical perceptions are beginning to change. This is evident through the ratification of VITA 78, or SpaceVPX, in April of this year.

SpaceVPX: Dual-redundant, fault-tolerant COTS

SpaceVPX began as part of an industry/government collaboration called the Next Generation Space Interconnect Standard (NGSIS), and is a retrofit of the VITA OpenVPX specification. that adds fault tolerance to meet the demands of space flight. Looking to define a modular open systems architecture (MOSA) for space system interconnects that removed bandwidth constraints, the NGSIS SpaceVPX working group (headed by Patrick Collier of the Air Force Research Laboratory (AFRL) and comprised of individuals from BAE Systems, Lockheed Martin, Northrop Grumman, TE Connectivity, and Elma Electronic, among others) selected OpenVPX as a physical baseline for SpaceVPX due to its broad

North Atlantic Industries
www.naii.com

Accelerate Your Time-to-Mission™

COSA™
Custom-On-Standard Architecture™

Solve challenging design issues with more intelligence, functionality and I/O density — all in a smaller footprint.

- Smarter systems
- Faster system integration
- SWaP efficient
- Less program risk
- Little or no NRE

Configure boards and systems easily for a complete COTS solution

- Multi-function I/O Boards
- Rugged Systems
- Single Board Computers
- Rugged Power Supplies

Distributed I/O & Processing • Lower Deployment Costs • Shorter Development Cycles

Ideally suited for a broad range of air, land and sea applications

Made in USA
A certified small business






- Communications
- Engine Control
- Guidance
- Machinery & Control
- Motion Control
- Sense & Response

... SPACEVPX PLATFORMS ARE NOT ONLY CURRENTLY BEING EVALUATED BY NASA, BUT ARE ALSO DRAWING "SIGNIFICANT INTEREST" FROM THE EUROPEAN SPACE AGENCY ...

ecosystem support, which helps reduce cost and promotes technology reuse.

To achieve a greater degree of resiliency than is provided in OpenVPX, however, SpaceVPX adds a Space Utility Management (SpaceUM) module that receives utility and management signals from an independent management controller for each of up to eight system slots. These signals, along with power bussed from an independent power supply, are then distributed to the various SpaceVPX modules to provide a dual-redundant architecture that eliminates single points of failure.

Further, the SpaceVPX specification limits the number of OpenVPX slot profiles to a subset that accommodates the current and future needs of the space electronics community. Of the 17 total backplane profiles, one of the payload slot profiles and one of the data switch profiles were directly mapped from OpenVPX, in addition to several new backward-compatible profiles and special slot definitions that enable the bridging of SpaceVPX modules to heritage CompactPCI (cPCI) modules.

Commercializing cost for space electronics

Rather than building custom development systems, the goal of the NGSIS and SpaceVPX specification is to provide vendors with the ability to conduct much of their design, testing, and integration work on inexpensive OpenVPX platforms before porting it to the SpaceVPX architecture. Over the life of a program using this whole lifecycle COTS approach can yield significant costs savings, and as the SpaceVPX market grows within the space community and into other verticals, these savings could possibly be augmented by a reduction in hardware costs resulting from larger scale production.

In a recent interview with Patrick Collier, he noted that SpaceVPX platforms are not only currently being evaluated by NASA, but are also drawing "significant interest" from the European Space Agency (ESA), whose suppliers include Airbus, Thales, and others. He also noted the promise for SpaceVPX in high-reliability non-space applications, such as air platforms, as possible drivers of the market.

As it stands, Collier notes BAE Systems and Honeywell as publicly developing SpaceVPX products, and points to other members of the VITA 78 working group as likely vendors of the platforms in the near future. Though there is not yet a definitive timetable, he expects the first SpaceVPX products to hit the market sometime in the first half of 2016.

Moving forward, VITA has also begun work on a VITA 78 dot spec focused on 3U SpaceVPX systems that targets low size, weight, and power (SWaP) applications. **SPD**

Digital Transceiver
V616
Turnkey Solution for
Digital Receiver • Surveillance
Arbitrary Waveform Generator/Player
Digital Transceiver V616 – a PC-based Instrument with Four 16-bit DDC channels, Two 16-bit DUC channels, Spectrum Analyzer and Two XMC I/O Module Sites
Features
• Intel i7 Quad Core
• Win 7 Pro 64-bit
• 8 GB RAM, 240 GB SSD
• Two, independent XMC module sites
DOWNLOAD DATASHEETS NOW!
Customization Available
Per-module Features
• Two 12-bit, 1000 MHz ADCs
Analog Bandwidth: 1000 MHz (AC Coupled)
• Two 16-bit, 1000 MHz DACs
Analog Bandwidth: 1000 MHz (AC Coupled)
• Xilinx Virtex-6 SX475T-2 FPGA
Digital Down-Converter (DDC)
• Four independent 16-bit DDC channels
• Programmable tuner: 1 - 1000 MHz
Resolution 0.2328 Hz
• Programmable bandwidth: 10 KHz - 100 MHz
Digital Up-Converter (DUC)
• Two independent 16-bit DUC channels
• Programmable tuner: 1 - 1000 MHz
Resolution 0.2328 Hz
Spectrum Analyzer
• Single wide-band/narrow-band 32K points FFT
• Eight Windowing available
Innovative Integration
a subsidiary of Interconnect Systems, Inc.
... real time solutions!
805.383.8994 phone • www.innovative-dsp.com

OPTICAL REVOLUTION

By Ray Alderman, Chairman of the Board, VITA

www.vita.com exec@vita.com

Optical data links could be in your (near) future

This article is not about the technical aspects of optical data links like light frequencies, single-mode versus multi-mode fibers, optical waveguides, PIN diodes, laser diodes, light beam losses over distance, connectors for optical links, optical engines, etc. There are plenty of website articles covering those topics, so let's focus on why optical could be in your future. We moved from buses to differential serial fabrics because we encountered the bandwidth limit of single-ended signals on copper. Very soon, we will hit the bandwidth limit of serial links on copper and be forced to move to optical. How close are we? Alderman's Laws of data transmission define that for backplane-based embedded computers:

Law 1: When (f+P) is increasing, then (s/n) is decreasing
On copper traces, as the frequency (f) of a link increases and the electrical signaling protocol (P) becomes more complex, the signal-to-noise ratio (s/n) declines dramatically. The higher the frequency and the more complex the signaling protocol on copper, the more the signal looks like noise.

Law 2: 2Bf=d/2
Every time the frequency of the signals going through a copper connector and along the copper backplane traces doubles (2Bf), the distance those signals can run declines by 50 percent (d/2). Otherwise, they look like noise again.

Law 3: 2Nf=B/2
Every time the network link frequency doubles (2Nf), the demand for copper-based backplanes and their associated connectors declines by 50 percent (B/2). When the network outperforms the backplane, there's no need for a backplane.

Law 4: if Nf > Bf, then 2Nf = B/2
If the network link frequency (Nf) is greater than the backplane link frequency between the boards (Bf), then backplanes lose their primary performance advantage. Today, we are running multiple links on copper traces to each board to keep up with the network link bandwidth. When we run out of connector pins on the backplane for those additional copper links, Law 3 is invoked.

Understand that backplane-based computers are centralized systems. Each board in the rack is connected to the others through copper traces on a common backplane, and all the boards share a central power supply and chassis. Backplanes have two basic benefits: high-performance data links and modularity. Modularity, in turn, has two sub-benefits: maintainability and upgradeability. When the bandwidth of the network cable or optical fiber is greater than the bandwidth of the copper traces between the boards on the backplane, the boards in the rack will be broken out into separate boxes, into a distributed system, and connected together with the network cable (Law 3). That's higher performance at a lower cost, since the rack, the big power supply, and the backplane are more expensive than the power supply and packaging for a single board (Law 4).

... UP TO 1994, ALL COMPUTERS WERE CPU-BOUND: THE DATA LINKS COULD DELIVER MORE DATA THAN THE CPU COULD PROCESS. AFTER 1994, AND THE ADVENT OF GHz MULTICORE PROCESSORS, WE HAVE BEEN I/O-BOUND: THE CPU CAN PROCESS MORE DATA THAN THE LINKS CAN DELIVER.

Also understand that up to 1994, all computers were CPU-bound: the data links could deliver more data than the CPU could process. After 1994, and the advent of GHz multicore processors, we have been I/O-bound: the CPU can process more data than the links can deliver. And, that situation is getting worse. Just look at the latest NVIDIA Tegra X-1 GPU. It has 256 GPU cores, eight 64-bit CPU cores (ARM), can handle 4K video at 60 frames per second (fps), and process data at the rate of one Teraflop (10^12 floating point operations per second (FLOPS)). You can feed data to this monster on copper connections if you only run them a few inches (Law 2). But, if you are bringing in data from remotely-mounted cameras on a Global Hawk, Predator, or a Reaper, you'll have to go optical on those links or the cores will be data starved (Law 1).

If you are making products for the industrial markets, those users can run most of their applications on an Apple Watch CPU, so they don't need the bandwidth. Industrial apps are register-oriented sequencers ("bingo" machines). If you are making products for telecom, they will stay with copper because they cannot afford optical. Telecom machines route data, they don't process it. Industry leaders in telecom or industrial board sales are just the lepers with the most fingers. If you are anticipating making products for the IoT (Internet of Trash) market, they neither need the bandwidth nor care about optical.

But, if you are making radar, sonar, medical, or signals intelligence machines (streaming data processing), you have some studying to do. **SPD**

INTEGRATION ISSUES

By Mark Honman, Sundance Multiprocessor Technology, Ltd.

www.sundance.com mark.h@sundance.com

If wishes were fishes – An accidental computer scientist goes embedded

Can't we just TALK?

There is a well-worn cliché in computing, "Standards are great – there are so many to choose from!" In computing this is perhaps less true than it used to be, thanks to the Internet and commoditization of end-user computing.

However, in the world of embedded systems there has not been the same kind of convergence in interface standards. While the DSP and FPGA vendor platforms have become evermore powerful and applicable to a broader range of applications, there seems to be little attention given to combining different types of processing elements in a single hardware module.

Before I go further, is it even still relevant to discuss the topic of integrating DSP and FPGA devices in the same hardware? After all, DSP vendors are now offering data acquisition interfaces that are supposed to make FPGAs obsolete, and FPGA vendors are offering software compilation to RTL that is supposed to make DSPs obsolete. Not to mention that everyone's space is being invaded by ARM and Linux-wielding refugees like myself from the "real computing" domain.

Perhaps the best reasons are productivity and time-to-market. After two decades of refinement, DSP architectures and surrounding tools and support libraries are wonderfully mature, and FPGA tools are similarly all-encompassing. The problem, as I see it, is that the overlap between DSP and FPGA capabilities only goes so far. There will continue to be projects that require processing for which no DSP has a suitable accelerator or interface, and others where there is a large computationally intensive code base.

Since it is easy to become so familiar with one vendor's platform that everything else appears alien, there is a trap waiting for all of us: "When all you have is a hammer, everything looks like a nail." We can generally sense when it is becoming more difficult – or risky – to implement functionality within the platform chosen for the current project. The bad news is that while adding a DSP or FPGA may save time by reducing the number of workarounds, compromises in the interface often mean that the cure is worse than the disease.

The success of the ARM-enhanced DSPs and FPGAs is with good reason – so much time can be saved by offloading non-critical functionality to ARM cores running Linux. And, at least where this code does not touch vendor-specific communication libraries, it is possible to reuse the software in a subsequent project based on another vendor's processor.

So, wouldn't our lives be simpler – and more interesting – if there was a simple, standardized, low-latency interface that made it possible to harness the best offerings of multiple vendors?

... WOULDN'T OUR LIVES BE SIMPLER – AND MORE INTERESTING – IF THERE WAS A SIMPLE, STANDARDIZED, LOW-LATENCY INTERFACE THAT MADE IT POSSIBLE TO HARNESS THE BEST OFFERINGS OF MULTIPLE VENDORS?

Now, it could be that that there is a solution that already fits the bill, but for the moment let's assume that is not the case. Other than low latency, there are two characteristics that are essential for a successful interface: simplicity and decoupling.

Without a powerful customer insisting that there is a common standard, it only makes sense to work together if using the standard saves us time and money compared to the effort we put into its development. A simple standard keeps the effort low and makes it relatively easy to adapt some new platform or bridge to a proprietary on-chip communication system.

Decoupling is a gift to the developers of the software or firmware on either side of the interface, allowing development of a communicating module that focuses on what it is supposed to do rather than being shaped by the quirks of its communication partners. When processes communicate there will always be some form of coupling, but much pain can be avoided by making these interactions predictable and consistent in their behavior. By making it easier to reason about process state, debugging is simpler and, even better, less likely to be needed.

At this point I am half hoping that someone will step forward and say "this has already been done," for surely after 25 years of embedded parallel processing this is a problem that has been solved many times over? If that is you, please write and put me out of my misery!

Next I will consider some of the options that could be used in the world of real project timescales and budgets, though from time to time you will also find me dreaming ... **SPD**

SENSOR FUSION FOCUS

By Jacques Trichet, Freescale Semiconductor

Basics of 6DOF and 9DOF sensor fusion

Inertial and magnetic sensors are becoming quite pervasive. When driving our cars or as simple pedestrians, most of us would have a hard time finding our way around without the navigation applications they support. Our smartphones rely on electronic compasses (eCompasses) built on inertial and magnetic sensor technology to determine direction, and with the help of GPS and maps, guide us blindly through the urban maze.

The eCompass is built on a very sensitive 3-axis magnetometer that strives to measure the Earth's geomagnetic field. It is associated with a 3-axis accelerometer whose main purpose is to provide a horizontality reference to the Earth's magnetic field vector, similar to how a mechanical magnetic compass must be kept horizontal for its needle to rotate freely and point toward the North Pole.

The combination of these two sensors provides six measureable data points, often referred to as six degrees of freedom (DOF), for a "tilt-compensated" eCompass. The key benefit of this two-sensor combination is that it delivers the North direction regardless of how a user holds their smartphone or tablet (in a flat or vertical position, in portrait or landscape orientation, for example). A six DOF eCompass can further provide device attitude and heading in 3D space by simply measuring the Earth's gravity and geomagnetic field, provided those two vectors are not collinear (such as at the Poles). This is how sky map applications can point at visible stars and teach you the constellations.

Whereas off-the-shelf accuracy of accelerometers is satisfactory, magnetometer measurements are susceptible to distortion and errors that can impact calibration and correction in the final environment. Indeed, even if the magnetometer has been thoroughly calibrated by the vendor, once mounted in an end product it will suffer local magnetic field perturbations such as hard-iron and soft-iron distortion, mainly due to the presence of magnetic materials in the vicinity of the sensor. Magnetic interferences can be up to

PICMG®

Open Modular Computing Standards

PICMG is the leading industry consortium for developing open standards for embedded computing. Our strength comes from our proven processes, intellectual property policies, and the diverse engineering talents provided by our hundreds of member companies.

In the last 20 years, we have released over 50 specifications leading to thousands of products, including:

- **Com Express®** - a scalable small form factor standard that powers millions of applications worldwide and is ideal for the emerging IoT revolution
- **AdvancedTCA®** - the platform of choice for high-performance, high-availability communications applications, providing NEBS and ETSI compliance and terrabit/second backplane switching and processing
- **CompactPCI® and CompactPCI® Serial** - modular architectures used in industrial, medical, military, and aerospace applications

Do you want to help set the standards for next generation architectures and systems?

Visit picmg.org and Join Today!

20x larger than the Earth's geomagnetic field of interest, so they absolutely have to be cancelled out.

One notable field magnetometer calibration procedure is the "figure 8" gesture that personal devices instruct their users to perform. Here, calibration is performed in the environment by simply leveraging the various natural orientations of a user's device, either through direct manipulation of the device or indirectly through the user's movements. Most accurate calibration algorithms are based on fitting the scatter plot of actual magnetometer measurements to an ellipsoid, after which correction parameters are used to transform the ellipsoid into a sphere.

THE OUTCOME GREATLY EXCEEDS THE SUM OF INDIVIDUAL SENSORS' CONTRIBUTIONS, WHICH IS THE ESSENCE OF SENSOR FUSION.

Once properly calibrated, the 6DOF eCompass is able to deliver accurate orientation of the device. Moreover, the sophisticated ellipsoid-to-sphere fitting calibration algorithm can run in the background and regularly update correction parameters to track changes such as ambient temperature. Nevertheless, the 6DOF eCompass still suffers a few weaknesses:

- Accelerometer measurements reflect the Earth gravity vector only when at rest. Whenever the device is moving, linear acceleration is also captured, which can corrupt the horizontality estimate for tilt-correction.
- Despite being calibrated, magnetometers are still highly sensitive to time-varying magnetic disturbances. Any changing "magnetic environment" will cause a measurement distortion that cannot be recovered.
- Output data from both sensors is usually low-pass filtered to remove measurement noise (especially for the magnetometer) and reduce perturbation impact. This results in lags in system response, as computed orientation can be heavily smoothed out.

To alleviate 6DOF shortfalls, a 3-axis gyroscope can be added, creating a 9DOF, or "gyro-stabilized," eCompass solution. The gyroscope provides the system with an independent measurement of instantaneous rotation speed, which complements the 6DOF computed orientation angles. Still, gyroscopes have an annoying "zero-drift offset" drawback, as a strictly motionless gyroscope will provide a residual rotation speed offset instead of the theoretical zero degrees per second. It is important to correct this error because the rotation speed measurement error will accumulate over time through integral calculus that is matched with the 6DOF rotation angle estimate.

Similar to how the 6DOF accelerometer/magnetometer data is low-pass filtered, the 3DOF gyroscope data could be high-pass filtered in order to remove the unwanted DC offset. As a matter of fact, those offsets are even estimated thanks to dedicated data processing using an extended Kalman filter. A comprehensive and instructive way to tackle 6DOF and 9DOF sensor fusion rudiments is experimenting with Freescale's Sensor Fusion Toolbox, which can be downloaded at www.freescale.com/sensorfusion.

Thanks to the third sensor and sensor fusion algorithms, the 9DOF tilt-compensated and gyro-stabilized eCompass is now gathering the best data from all three sensors, identifying and compensating for the flaws of some with the strengths of the others. It provides an accurate, low noise, smooth but responsive estimate of the device orientation angles, rotation speed, linear acceleration, Earth gravity, and geomagnetic fields; delivers sensor calibration error information; and is more resilient to magnetic interferences. The outcome greatly exceeds the sum of individual sensors' contributions, which is the essence of sensor fusion. **SPD**

ARIES E3800 / Bay Trail CPU
- 1.91GHz quad core & 1.46GHz dual core processors
- 2 or 4 GB memory down
- 2 Gigabit Ethernet
- 3 USB 2.0 & 1 USB 3.0
- 4 RS-232/422/485
- Industry leading data acquisition:
  - 16 16-bit A/D
  - 4 16-bit D/A
  - 22 digital I/O
- PCIe Minicard and SATA DOM
- PC/104 & PC/104-*Plus* expansion
- -40°C to +85°C
- 4.0" x 4.5" PC/104-compatible form factor

ATLAS N2800 CPU
- 1.86GHz quad core processor
- 2 or 4GB memory down
- 1 Gigabit Ethernet
- 4 USB 2.0
- 6 RS-232/422/485
- PCIe MiniCard socket
- PCI-104 & PCIe/104 expansion
- -40°C to +75°C
- 4.5" x 3.8" PCIe/104 form factor

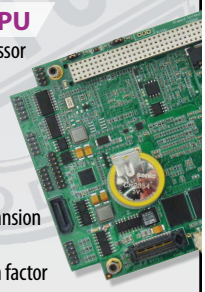

VEGA Core i7 COM Express
- COM-based SBC for long lifecycle
- Quad/Dual core COM Express CPU:
  - 2.1GHz Intel Core i7-3612QE
  - 1.7GHz Intel Core i7-3517UE
- 2 Gigabit Ethernet
- 4 USB 2.0
- 4 RS-232/422/485
- Industry leading data acquisition:
  - 16 16-bit A/D, 8 16-bit D/A
  - 30 digital I/O
- PCIe Minicard and EMX expansion
- 7 to 36V DC/DC power supply
- -40°C to +85°C
- 4.9" x 3.7" EMX form factor

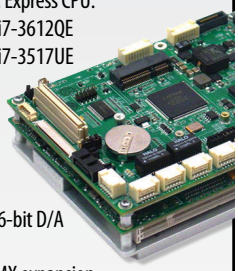

EMX EmbeddedXpress
emxbus.org

DIAMOND SYSTEMS
The Perfect Fit for Imperfect Environments

800-36-PC104 (800-367-2104) • sales@diamondsystems.com
diamondsystems.com

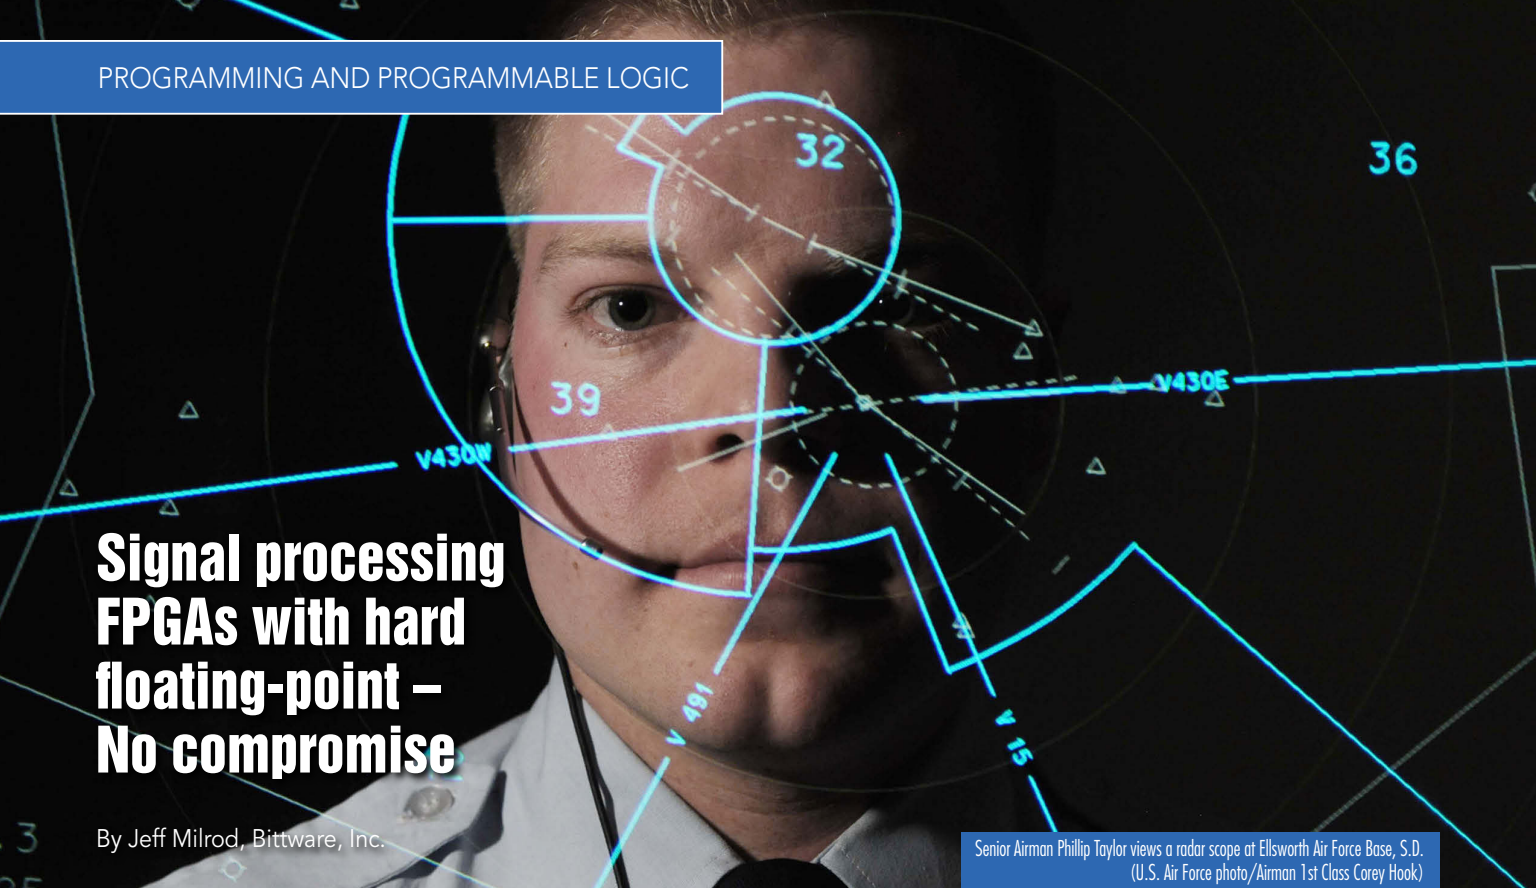

Signal processing FPGAs with hard floating-point – No compromise

By Jeff Milrod, Bittware, Inc.

Senior Airman Phillip Taylor views a radar scope at Ellsworth Air Force Base, S.D. (U.S. Air Force photo/Airman 1st Class Corey Hook)

Since the dissolution of cutting-edge digital signal processor (DSP) product lines designers have been forced to develop using either FPGAs integrated with time-consuming fixed-point DSP blocks, or floating-point general-purpose graphics processing units (GPGPUs) that leave performance on the table in high-end signal processing systems. But now, with the release of Altera's Generation 10 FPGAs that integrate hardened IEEE 754-compliant floating-point operators, why compromise?

The days of high-end, general-purpose DSPs effectively ended a decade ago with the demise of Analog Device's TigerSHARC roadmap after Texas Instruments (TI) had previously discontinued their high-end roadmap. Since then, TI has brought back some high-end parts for targeted applications, but they still have not presented a roadmap for increased performance. There are still plenty of application-specific (ASIC/ASSP) and low-end DSP processors available, but general-purpose signal processing applications requiring high performance must now rely almost exclusively on GPUs and FPGAs.

FPGAs have been used to implement high-performance DSP algorithms for a long time, but they require specialized and complex development. Many years ago FPGA vendors made a big leap by adding hard DSP blocks within the gate array to significantly improve signal processing performance and simplify algorithm implementation. Unfortunately, those hard DSP blocks were all fixed-point.

Floating-point clearly a superior format

Much can be said about numerical representation of data for signal processing. When using fixed-point there are numerous options, each with their own benefits and limitations, including integer, fractional, signed, unsigned, block exponents, and combinations thereof; however, most of the complications and tradeoffs encountered with fixed-point data representations can simply be avoided by using a floating-point format (a concise overview of these issues can be found in Michael Parker's book "Digital Signal Processing 101"). Suffice it to say that no one prefers fixed-point over floating-point, as floating-point is clearly a superior data format, and as a result virtually all algorithmic development and simulation is done in floating-point during system design. A very high-level comparison of the relevant characteristics of fixed- and floating-point implementation can be found in Table 1.

Several years ago, GPUs added floating-point capabilities to their parallel network of computer engines, creating the "GPGPU," a general-purpose GPU useful for much more than just graphics. Because they were originally architected for graphics

Characteristic:	Floating-Point	Fixed-Point
Dynamic Range	Much Larger	Smaller
Resolution	Comparable	
Algorithm Design Time	Much Less	More
Ease of Use	Better	Worse
Compiler Efficiency	Much Better	Much Worse

Table 1 | A high-level comparison of fixed- and floating-point.

processing, however, and not signal processing, GPGPUs still suffer from a number of challenges when used in general-purpose signal processing applications, including high power, single instruction/multiple data (SIMD) programming complexities, memory bottlenecks, and often disappointingly low effective processing rates relative to peak capabilities. In addition, GPGPUs are limited by a lack of flexibility, limitations of the memory types they can support, and the inability to interface directly to analog converters for signal I/O, or any I/O for that matter. Finally, the GPU can become data starved unless there is a high degree of calculation to be done on each data point, since the host GPU must provide data over a PCIe link to the GPU, which is a liability for most traditional stream-based signal processing applications. Therefore, FPGAs have continued to thrive in high-performance signal processing applications despite their lack of floating-point resources.

FPGA designers have become extremely skilled at fixed-point implementations, but there is still a significant cost associated with these fixed-point implementations. Once the algorithm simulation is completed in floating-point, there is typically a further six- to 12-month effort to analyze, convert, and verify a floating-point algorithm in a fixed-point implementation. First, the floating-point design must be converted manually to fixed-point, which requires an experienced engineer. Second, any later changes in the algorithm must be converted manually again; also, any steps taken to optimize the fixed-point algorithm in the system are now not reflected in the simulation. Third, as problems arise during system integration and testing, debug time increases inordinately as the possible causes could be an error in the conversion process, a numerical-accuracy problem, or the fact that the algorithm itself is just defective. The advantages of floating-point are so compelling that often "soft" floating-point is implemented using hard fixed-point DSPs, which causes other deleterious impact on design time while consuming significantly greater FPGA resources and slowing performance.

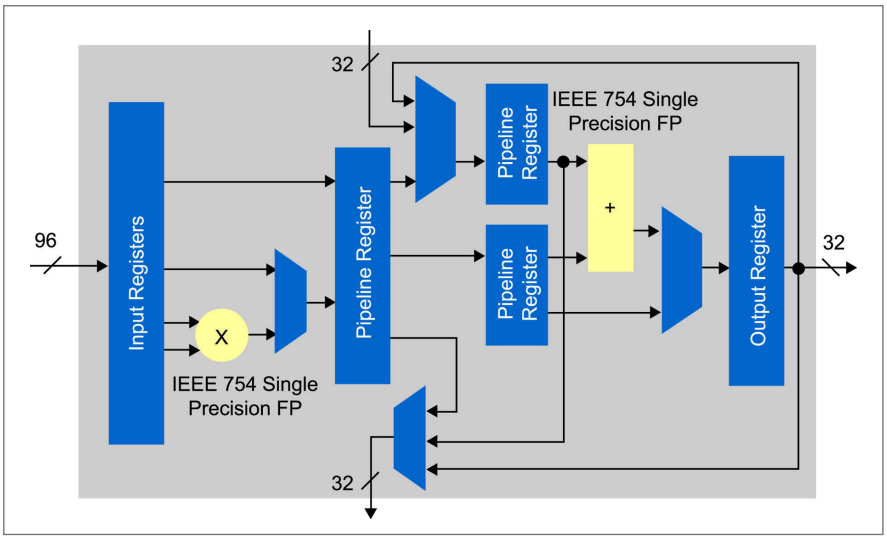

Figure 1 | The variable precision DSP block of Altera's Generation 10 FPGAs is shown here in floating-point mode.

Therefore, since the demise of dedicated high-end DSP processors, applications requiring high-end signal processing have been forced to choose between the lesser of two evils: fixed-point FPGAs or floating-point GPGPUs. While both of those options have had a great deal of effort put into them to make them less "evil" via tricks, techniques, and prebuilt intellectual property (IP), these options have generally required significant compromise – until now.

Floating-point comes to FPGAs

With the recent introduction of its Generation 10 FPGAs, Altera has become the first to integrate hardened IEEE 754-compliant floating-point operators in an FPGA, as shown in Figure 1. These hardened floating-point DSP blocks change the decade-old paradigm of having to choose between the lesser of two evils, and removes the need to compromise.

The floating-point computational units, both multiplier and adder, are seamlessly integrated with existing variable-precision fixed-point modes. This provides a 1:1 ratio of floating-point multipliers and adders, which can be used independently as a multiply-adder or multiply-accumulator. Since all the complexities of IEEE 754 floating-point are within the hard logic of the DSP blocks, no programmable logic is consumed and similar clock rates as used in fixed-point designs can be supported in floating-point, even when 100 percent of the DSP blocks are used. In addition, while designers still have access to all the fixed-point DSP processing features used in their existing designs that support backward compatibility, they can easily add or upgrade all or part of the design to single-precision floating-point as desired.

With thousands of floating-point operators built into these hardened DSP blocks, Arria 10 FPGAs are available from 140 GigaFLOPS (GFLOPS) to 1.5 TeraFLOPS (TFLOPS) across the 20 nm family. Altera's 14 nm Stratix 10 FPGA family will use the same architecture, extending the performance range right up to 10 TFLOPS, the highest ever in a single device. This situation means that FPGAs can now compete directly with GPGPUs for raw processing performance without compromising the FPGA's previous advantages of inherent flexibility, support for a variety of memory and I/O types, and the ability to directly connect to signals.

Floating-point simplifies FPGA development

The addition of native floating-point also greatly improves the ability to leverage higher-level languages, tools, and compilers for coding of FPGA applications, thus addressing any lingering ease-of-use concerns. Existing model-based flows such as

DSP Builder Advanced Blockset and MathWorks' MATLAB and Simulink tools, as well as higher-level language compilers such as OpenCL, are now able to be far better and more efficient without having to map to fixed-point numerical issues such as managing bit growth, truncation, saturation, and the like, allowing the use of integers to be restricted to more effective roles in semaphores, memory indexing, and loop counters.

While OpenCL, which is the open standard equivalent of CUDA, can be used for both FPGAs and GPGPUs, there are notable differences in how algorithms are implemented. GPGPUs use a parallel-processor architecture, with thousands of small floating-point mult-add units operating in parallel. The algorithm must be broken up into thousands of threads, which are mapped to the available computational units as the data is made available. On the other hand, FPGAs use a pipelined logic architecture where the thousands of computational units are usually arranged into a streaming data flow circuit, which operates on vectors; this setup is more typical of signal processing components such as an FFT, filters, or Cholesky decomposition.

Floating-point enables FPGAs to compete in high-performance computing

While FPGAs are relatively new to high-performance computing, they can provide some compelling advantages. First, due to the pipelined logic architecture, the latency for processing a given data stream is much lower than on a GPU. This can be a key advantage for some applications, such as financial trading algorithms. Second, FPGAs with native floating-point achieve four to eight times higher GFLOPS per watt than GPGPUs. This improved efficiency can be critical in applications such as high-performance embedded computing (HPEC), where an FPGA can perform far more computations within a limited power budget; it's also becoming a huge advantage in big-data processing and data centers due to the reduced operating costs. Third, the FPGA has an incredibly versatile and ubiquitous connectivity. The FPGA can be placed directly in the data path and process the data as it streams through. Altera has specifically added the option of data streaming to its OpenCL tools, which is in compliance with the OpenCL vendor extension rules.

FPGAS with native floating-point easy to deploy

One final and oft-overlooked aspect of deploying signal processing implementations is the availability of quality hardware. Even the greatest FPGA with the most robust and easiest to use development tools won't do much good if it takes a team of hardware engineers six to 12 months to build and debug a board to host the application. Fortunately, like GPGPUs, high-quality deployable board-level solutions for FPGAs with native floating-point are available off-the-shelf. Vendors such as BittWare offer a variety of PCIe and embedded formats (VPX and AMC). Complete with drivers, system integration software, FPGA development kits (FDKs), board monitors, and board support packages (BSPs) for OpenCL, boards such as the A10PL4

Figure 2 | BittWare's A10PL4 low-profile PCIe board is based on Altera's Arria 10 FPGA, and includes 8x 10 GbE, 2x 40 GbE, or 2x 100 GbE network interfaces and up to 32 gigabytes of DDR4.

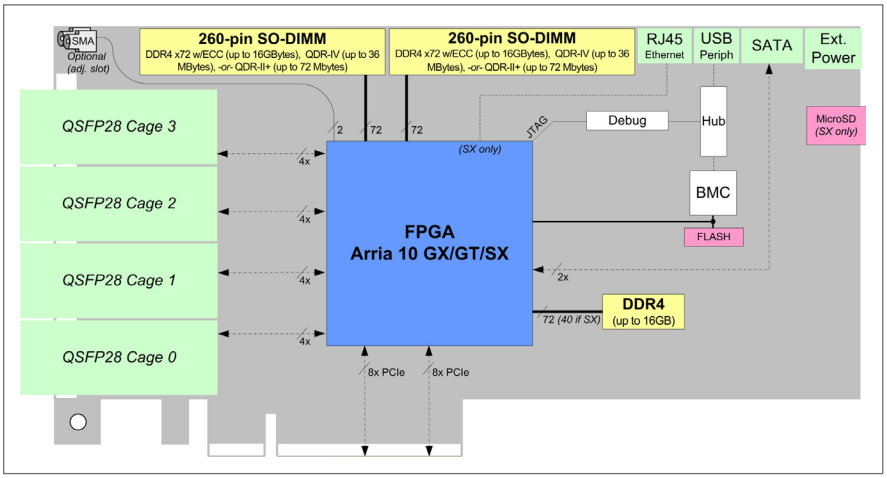

Figure 3 | This block diagram depicts BittWare's A10P3S half-length PCIe board with Arria 10 FPGA.

and A10P3S greatly lower the barrier to entry for implementing high-performance signal processing on FPGAs (Figures 2 and 3).

With integrated hard floating-point DSPs, high-level development tools, and sophisticated processing boards, FPGAs now have everything needed for efficiently and effectively implementing high-performance signal processing applications. There's no longer a need to compromise.

Jeff Milrod *received his bachelor's degree in physics from the University of Maryland, and MSEE degree from Johns Hopkins University. After gaining extensive design experience at NASA and business experience at Booz, Allen & Hamilton, Jeff started Ixthos in 1991, one of the first companies dedicated to COTS DSP. He ran Ixthos until it was acquired by DY4 Systems (now Curtiss-Wright Defense Solutions) in 1997 before taking the helm of BittWare, where he is President and CEO.*

Bittware, Inc.
www.bittware.com
j_milrod@bittware.com
www.linkedin.com/company/137501
www.facebook.com/pages/BittWare/336897291583?ref=sgm
ewww.youtube.com/user/bittwareinc

INTERFACE CONCEPT
ADVANCED ELECTRONIC SOLUTIONS

Build your own VPX system !

INTERFACE CONCEPT product range of Single Board Computers, FPGA boards, ADC/DAC FMC and Graphic boards are ideal to devise a complete VPX system for compute intensive and image processing applications (radar, electronic warfare, electro optical and IR, visualization systems)

Intel® Core™ i7 SBC

- Two **Core™ i7** Processors (Dual / Quad Core)
- One Ethernet switch, XMC slot...
- One **Kintex™ 7** FPGA & FMC site

Virtex®-7 FPGA Boards

- Two **Virtex®-7** 690T & FMC sites
- One Freescale™ **QorIQ T1042** (or T2081)

Graphic Boards

- One AMD **Radeon™ E8860**
- One **Kintex™-7** 325T FPGA
- Support for DP, HDMI, VGA, Stanag3350, Arinc8181...
- One PMC/XMC site

www.interfaceconcept.com
+33 (0)2 98 57 30 30

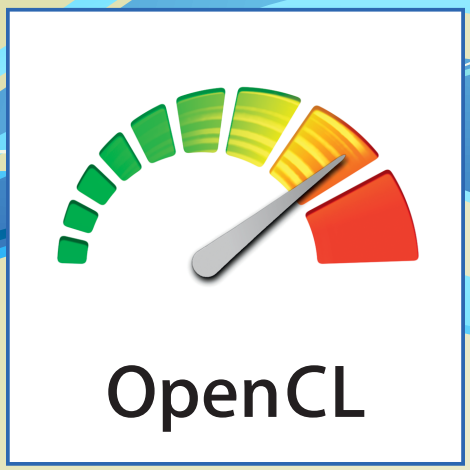

Designing a 2 million-point frequency domain filter using OpenCL for FPGA

By Dmitry Denisenko and Mykhailo Popryaga, Altera Corporation

Fast Fourier transform (FFT) is the backbone of signal processing applications. For a long time now, FPGA vendors have been providing well-tuned FFT libraries to process data sets that fit in FPGA on-chip memory. But what do you do if your data set is too large? To solve this problem, the FPGA designer must now make multiple intertwined design decisions, such as considering on-chip FFT core configuration options, how many to include, how they connect and access external memory, synchronization among multiple cores, and many others. Exploring all such design decisions to create the perfect combination for the product at hand while coding in HDL is just too time-consuming and can actually leave performance on the table. With a higher level programming language such as OpenCL, however, system design exploration can be done in days.

Consider the creation of a frequency domain filter supporting between 1 million and 16 million points on current FPGA architectures with sample rates from 120 million to 240 million samples per second. The example looks at design decision options for a 2 million-point single-precision frequency domain filter using OpenCL.

Such a filter translates its input into the frequency domain using a multimillion-point one-dimensional (1D) FFT, multiplies each frequency and phase component by a separate user-provided value, and then translates the result back into the time domain with an inverse FFT. The overall target performance requirement of the whole system is 150 million samples per second (MSPS) for a 2 million-point sample size on a current-generation FPGA with two DDR3 external memory banks. Inputs and outputs come over 10 gigabit Ethernet (GbE) directly into the FPGA.

The design uses the Altera SDK for OpenCL FPGA compiler targeting a BittWare S5-PCIe-HQ board with a Stratix V GSD8 FPGA. OpenCL, instead of a lower-level language, is used for two reasons:

1. The first reason is that designing multimillion-point filters requires building a complicated yet highly efficient external memory system. With lower-level design tools, creating individual blocks such as an on-chip FFT or a corner-turn is relatively easy (especially because every FPGA vendor already provides libraries containing such blocks). However, creating the external memory system would normally require a lot of HDL work. This situation can be especially challenging, as we will see later, because the configuration of the overall system is unknown at the very beginning.
2. The second reason for choosing OpenCL is host-level control over the FPGA logic. For this design, it is clear from the start that two full copies of multimillion point FFT cores will not fit on a single device, so a single data set will have to pass over the FPGA logic at least twice before producing the final output. Coordinating such sharing while also allowing features such as dynamically changing data set size, multiplication coefficients, and even completely changing FPGA functionality for something else is best left to a CPU.

Using the OpenCL compiler for FPGAs solves both of these challenges as it builds a customized and highly efficient external memory system while allowing fine-grained control over the FPGA logic.

On-chip FFT

For this design, it's assumed that we already have an FFT core that can handle data sizes that fully fit on an FPGA (referred to as "on-chip FFT"), as every FPGA vendor provides such cores. Such a core is parameterizable in at least the following ways:

- Data type (fixed or single-precision floating point)
- Number of points to process (N)
- Number of points to process in parallel (POINTS)
- Dynamic support for changing the number of points to process

Given such an on-chip FFT core, building the overall system requires two steps: First, building an FFT core that can handle multimillion points, and second, stitching two such cores together with complex multiplication between them to create the whole system.

Multimillion-point FFT

The classic way to implement an FFT with external storage is the six-step algorithm shown in Figure 1 that treats a single, one-dimensional data set as two-dimensional (2M = 2K x 1K)[1].

The six-step algorithm shows both the separate computation kernels and external memory buffers. The "Fetch" kernel reads data from external memory, optionally transposing it, and outputs it into a channel (also known as a "pipe" in OpenCL 2.0 nomenclature). In hardware, a channel is implemented as a FIFO with compiler-calculated depth. An "on-chip 1D FFT" is an unmodified vendor's FFT core, taking input and producing a bit-reversed output using channels. "Transpose" always transposes the data it reads from its input channel, optionally multiplying it by special twiddle factors, and writes outputs in natural order to external memory.

As shown in the diagram, the data is sent over Fetch → 1D FFT → Transpose (F1T) pipeline twice to produce final output. This gives us our first important design choice: Either have one copy of the F1T pipeline to save area, or two copies for a higher possible throughput.

Initial prototyping of this algorithm is done in an emulator to ensure that

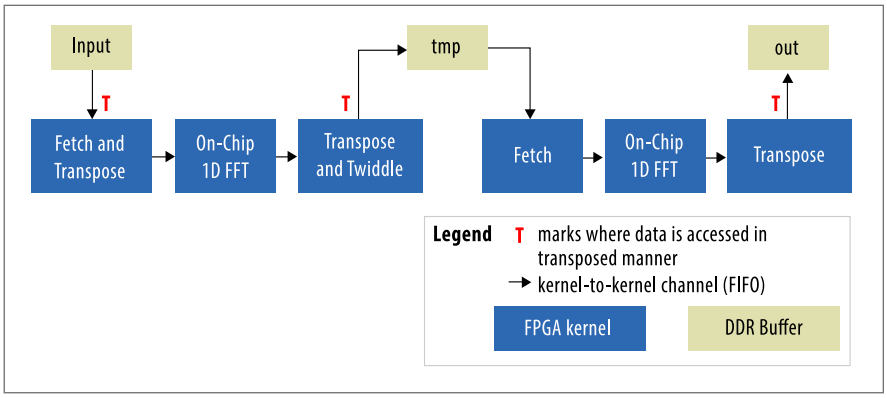

Figure 1 | A logical view of the six-step FFT algorithm.

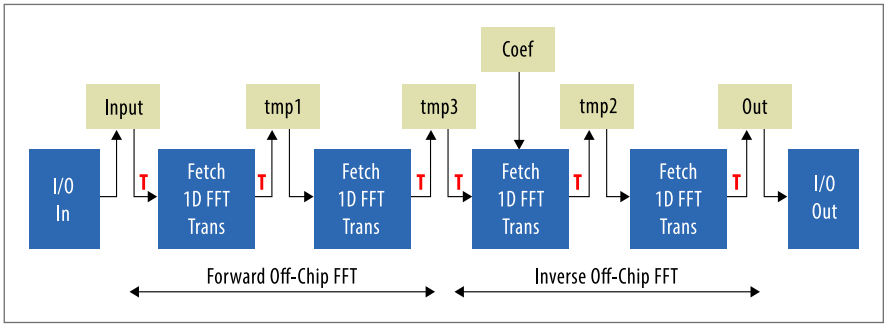

Figure 2 | This logical view of the Full Filter System shows F1T pipelines represented as single blocks for brevity.

address manipulation for transpositions and twiddle factors is correct. An emulator compiles OpenCL kernels to x86-64 binary code that can be run on a development machine without an FPGA. Going from emulator to hardware compile is a painless step, as functionally correct code in the emulator became functionally correct code in the hardware so no simulation is needed.

The only aspect that has to be modified, for performance and area reasons, is the local memory system used by the Fetch and Transpose kernels. Efficient transposition requires buffering *POINTS* columns/rows of data in local memory. The OpenCL compiler analyzes all accesses to local memory in the OpenCL code and creates a custom on-chip memory system optimized for that code. In the case of POINTS=4, the original transposition kernels had four writes and four reads. An on-chip RAM block, double-pumped, can service at most four separate requests with at most two of these being writes. To support four writes and four reads, on-chip memory needs to be both duplicated and contain request arbitration logic, which causes area bloat and performance loss. However, the write pattern can be changed to make all four writes consecutive. These four writes were grouped by the OpenCL compiler into a single, wide write, giving only five accesses to the local memory system: one write and four reads. With that change, the compiler automatically builds a much smaller five-ported memory system that could service all five requests on every clock cycle without stalling.

Once the design is compiled to hardware, it's time to measure performance. With a single copy of the F1T pipeline on the FPGA, we measure 217 MSPS with POINTS=4 and 457 MSPS with POINTS=8 for a 4 million-point FFT[2]. The POINTS=8 version used twice as many on-chip Block RAMs, and two copies in this configuration will not fit. This gives us the first design dimension to explore – the number of points to process in parallel versus area.

Full-filter design

Now that we have a multimillion point FFT, we are ready to put the whole design together. Simply stitching two off-chip FFTs gives us the logical view of the whole pipeline in Figure 2.

Besides duplicating a single off-chip FFT computation pipeline, the following parts are added to the system:

3. Complex multiplication in the frequency domain is absorbed into a third F1T block. The `coef` buffer is holding two million complex multiplication coefficients.
4. I/O in and I/O out kernels are added to realistically model the additional load of 10 GbE channels on external memory. With these kernels we can continue purely software-based development and leave Ethernet channel integration until after the core computation pipeline is fully optimized. The I/O in kernel generates a single sample every clock cycle, and I/O out consumes a single sample every clock cycle.

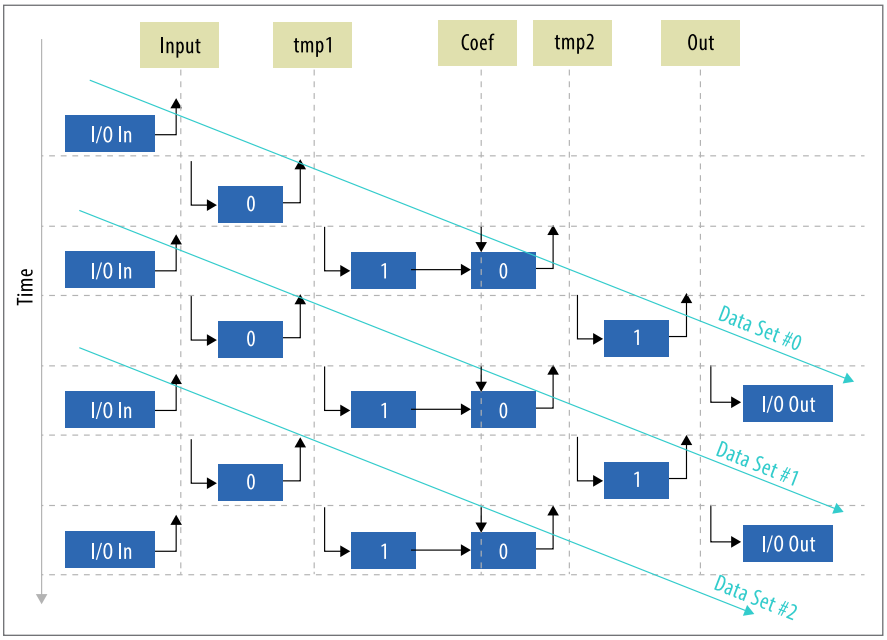

Figure 3 | In terms of kernel scheduling, "0" is the first physical copy of the F1T pipeline and "1" is the second copy. Purple arrows follow a single data set through the pipeline.

As experiments with off-chip FFT showed, we can fit only two F1T blocks, and only with POINTS=4. Therefore, the data has to pass through the hardware twice for full computation. That gives us an overall system throughput of only 120 MSPS for 2 million points, below our target of 150 MSPS. However, by reducing the data size to 1 million points, we are able to fit a POINTS=8 version and get throughput of 198 MSPS. That shows that there is still performance to be had, if only we can make a POINTS=8 version fit for 2 million points.

Picking an optimized structure of the full pipeline in Figure 2 is the next step in the overall design process. The first improvement we can make is to remove `tmp3` buffer. Both sides access it in the same way (transposed write and read), and therefore the second and third F1T blocks can be connected directly by a channel. This requires making the Transpose kernel either write its output to external memory or into a channel, and a similar change for Fetch. Such a change is dynamically controlled by the host, so a single physical instance of Fetch can be used. Note that this changes our connectivity to external memory, but this is something we don't have to worry about at all because the OpenCL compiler always generates an efficient custom external memory interconnect for our system.

A further improvement would be to move the second transpose "T" from writing to `tmp1` to reading from `tmp1` (the data in `tmp1` will be stored differently but the net effect is the same). This eliminates the need for one local memory buffer used by transpositions. Even though this change is not hard to implement, we decide to forgo it in lieu of a more radical idea.

Our original transposition implementation is done in two stages:

First all the required data is loaded into local memory and then read from local memory using transposed addresses. To efficiently utilize such a pipeline, the OpenCL compiler automatically double-buffers the local memory system. This

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

Choosing the right video interface for military vision systems

By Pleoria Technologies

This paper discusses how GigE Vision video interfaces – the technology used to transfer data from a camera or image sensor to a mission computer or display – help designers reduce the cost and complexity of military imaging systems, while also improving usability and increasing intelligence for end users.

Check out our white papers
whitepapers.opensystemsmedia.com/

way, the loading part of the pipeline can load the data into one copy while the reading part can read previous data sets from another copy.

This automatic double buffering is the right thing to do for our transposition algorithm, but it's expensive. Instead, we rewrite the transposition kernels to be in-place. Such a kernel only needs only a single buffer and supports reading and writing multiple data points at the same time (but we'll leave a detailed description of this transposition kernel for another time).

With these changes we are able to fit a 2 million-point FFT in a POINTS=8 configuration and achieve 164 MSPS throughput.

Scheduling

Only two copies of F1T could fit, but Figure 3 shows how the data flow can be scheduled to fully utilize the pipeline. Notice that in a steady state, the pipeline alternates between processing two and three data sets at a time without additional buffers. This scheduling is controlled by the host program running on a CPU and verified using the Dynamic Profiler tool.

Buffer allocation

In an OpenCL system the host program controls which DDR bank contains which buffers. Since a DDR bank is most efficient when it's either read from or written to, but not both, we can split the five buffers among two DDR banks as follows:

- DDR bank #0 gets `input` and `tmp2`
- DDR bank #1 gets `tmp1`, `coef`, and `out`

Assigning a buffer to a DDR bank is a one-line change in the OpenCL host program. The compiler and the underlying platform take care of the rest. Given such automation, we can experiment on 2-DDR and 4-DDR boards to find the best mapping of buffers to banks for each board.

Conclusion

This article describes how to design a 2 million-point frequency domain filter using the Altera OpenCL SDK for FPGAs. All functional verification was done using software-style emulation, and every single hardware compile worked correctly. We did not open a hardware simulator and never worried about timing closure.

References:

[1] Bailey, D.H. "FFTs in external of hierarchical memory." Proceedings of Supercomputing'89 (SC89), pp. 234-242.

[2] "FFT (1D) Off-Chip Design Example." Altera Corporation. www.altera.com/support/support-resources/design-examples/design-software/opencl/fft-1d-offchip.html.

Dmitry Denisenko *is a senior software engineer at Altera Corporation.*

Mykhailo Popryaga *is a software engineer at Altera Corporation.*

Altera Corporation • www.altera.com

@alteracorp

www.facebook.com/altera

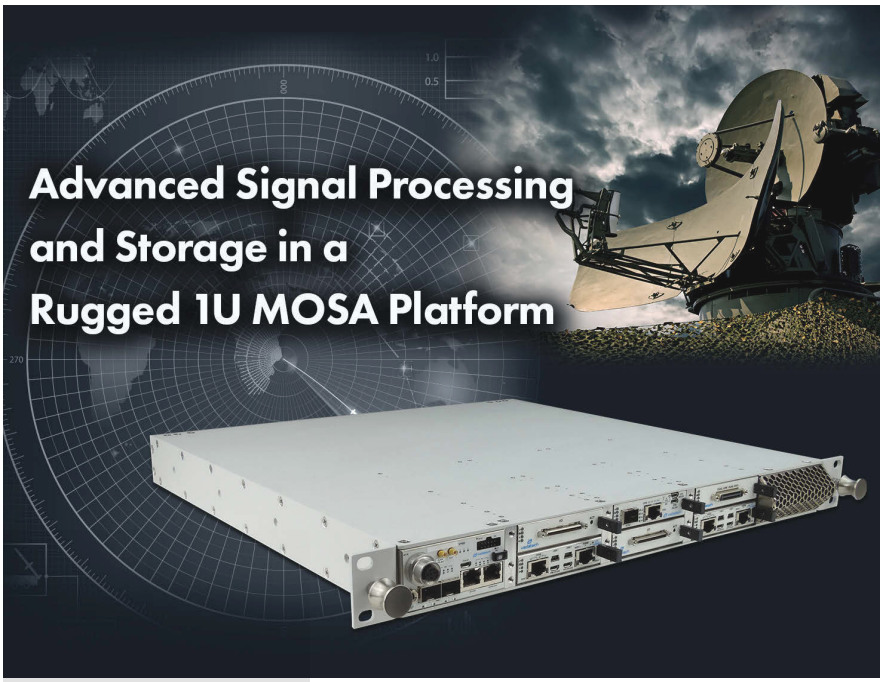

ADC & DAC Modules
- Multi GSPS, Xilinx® suite of FPGAs
- Various channels and resolution
- FMC versions

Processors
- Wide range of Intel® & Freescale™ processors

Storage Modules
- 2.5" SAS-3 or SATA III disks
- RAID options to RAID 60
- Removable or fixed options

Finally, a MIL rugged open-standard platform in 1/3 the rack height and weight of competing solutions! With built-in precision timing and time-stamping options, the modular 1U platform allows up to 6 high performance modules (digitizing, storage, etc) to be incorporated in your scalable solution.

VadaTech – Redefining performance density!

vadatech
THE POWER OF VISION

www.vadatech.com • info@vadatech.com • 702.896.3337

Design of FMCW radars for active safety applications

By John Zhao, Marco Roggero, and Giorgia Zucchelli, MathWorks

Image courtesy Bosch Mobility Solutions.

Car manufacturers, automotive electronics suppliers, and universities are working to develop new electronic systems for advanced driver assistance systems (ADAS). Frequency-modulated continuous waveform (FMCW) radars fit the requirements of automotive active safety systems because of their accurate short-range measurements, low sensitivity to clutter, and easy integration. FMCW radars are widely used as a component of ADAS systems in the automotive industry.

The following shows a unique tool chain for modeling and simulating a complete 77 GHz FMCW radar system, including waveform generation, antenna characterization, channel interference and noise, and digital signal processing (DSP) algorithms for range and speed determination. Simulation and modeling of RF impairments such as noise, nonlinearity, and frequency dependencies enable us to test the behavior of "off the shelf" components described with datasheet parameters, and provide information about the performance achievable with a specific component configuration and related costs.

Frequency-modulated continuous waveform (FMCW) radars are becoming increasingly popular, especially in automotive applications such as adaptive cruise control (ACC). The transmitter of an FMCW system sends a chirp signal with high frequency and large bandwidth. The transmitted signal hits the target and is reflected back toward the receivers with a time delay and a frequency shift that depends on the target distance and relative speed.

By mixing the transmitted and the received signal, the time delay corresponds to a frequency difference that generates a beat frequency. This allows a very accurate and reliable estimation of the target distance[1]. Often, multiple antennas are used for spatial processing and beamforming to make the detection more reliable or to have a directional system, as depicted in Figure 1.

In the design, modeling, and simulation of an FMCW radar, the designer must take into account more than just the nominal behavior. After using the radar equation to determine the fundamental design parameters, the designer must analyze the impact of imperfections introduced by the RF front-end. Nonlinearity, noise, frequency selectivity, and mismatches between components operating over ultra-large bandwidth reduce the actual dynamic range of the detectable signal.

By accurately modeling the RF front-end, designers can make complexity tradeoffs between the hardware architecture and the DSP algorithms. Moreover, they can assess whether previous implementations can be reused to retarget the radar for augmented specifications, or whether off-the-shelf components can be directly used for the front-end implementation.

Determination of FMCW waveform

The first problem we have to cope with when designing a new radar system is to determine the parameters of the triangular chirp waveform in order to achieve desired

BY ACCURATELY MODELING THE RF FRONT-END, DESIGNERS CAN MAKE COMPLEXITY TRADEOFFS BETWEEN THE HARDWARE ARCHITECTURE AND THE DIGITAL SIGNAL PROCESSING ALGORITHMS.

resolution with the specified range. We consider an automotive long-range radar used for automatic cruise control, which usually occupies the band around 77 GHz[2, 3].

As shown in Figure 2, the received signal is an attenuated and time-delayed copy of the transmitted signal where the delay Δt is related to the distance of the target. Because the signal is always sweeping through a frequency band, at any moment during the sweep, the frequency difference fb, usually called beat frequency, between the transmitted signal and the received signal is constant. Because the sweep is linear, one can derive the time delay from the beat frequency and then the distance of the target from the time delay.

Using MATLAB and Phased Array System Toolbox functionality, we can easily determine the fundamental waveform parameters for a radar working at 77 GHz, such as sweep bandwidth and slop, maximum beat frequency, and sample frequency based on user-specified range resolution and maximum speed as shown in Figure 3.

Modeling RF components, noise, and nonlinearity

Once the chirp parameters have been determined, we can proceed with modeling the transceiver of the radar system.

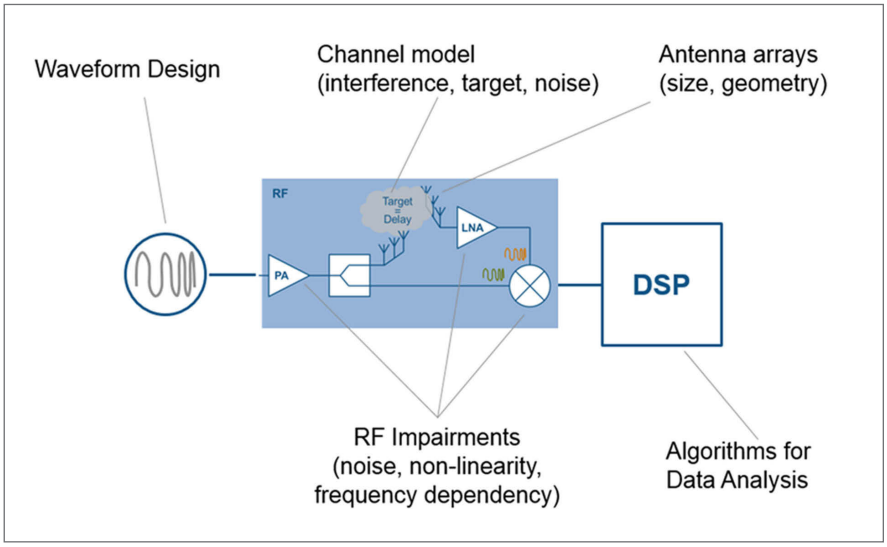

Figure 1 | Structure of an FMCW radar system.

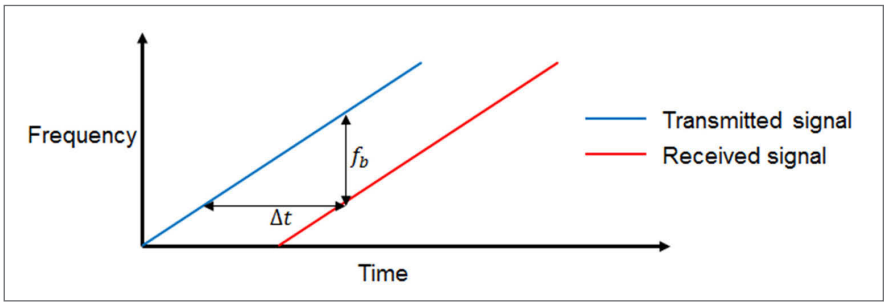

Figure 2 | Waveforms of transmitted and received signals.

```
fc = 77e9;
c = 3e8;
lambda = c/fc;
range_res = 1;                % resolution = 1m
v_max = 230*1000/3600;        % max speed = 230km/h

bw = range2bw(range_res,c); % sweep bandwidth
range_max = 200;
sweep_time = 2e-4;
sweep_slope = bw/sweep_time; % sweep slope

fr_max = range2beat(range_max,sweep_slope,c);
fd_max = speed2dop(2*v_max,lambda); % Doppler shift
fb_max = fr_max+fd_max;       % max beat frequency
fs = max(2*fb_max,bw);        % sample frequency
tstep = 1/fs;                 % sample rate
```

Figure 3 | Determining the parameters of the FMCW chirp waveform.

The front-end of the radar system includes the transmitter, the receiver, and the antenna. These models are provided in the toolbox. We parameterize these models with desired values, such as phase noise and thermal noise. Alternatively, we can model the transmitter and receiver using RF components provided in Simulink using SimRF to model the effect of component-level noise, nonlinearity, and frequency selection. Figure 4 shows how we have modeled the RF front-end using SimRF blocks. This library provides a circuit envelope solver for the rapid simulation of RF systems and components such as amplifiers, mixers, and S-parameter blocks.

We can describe in detail the architecture of the transceiver and use datasheet parameters for each front-end element. Taking as an example the direct conversion I/Q mixer, we have modeled it as illustrated in Figure 5. This element demodulates the received signal, multiplying it with the originally transmitted waveform.

Parameters of the two multipliers used in the I/Q mixer have been set directly on the blocks or using workspace variables.

With this configuration, it is easy to try different setups and explore design spaces by using different datasheet parameters for simulating off-the-shelf components.

Complete system simulation

After all components of the radar system have been properly parameterized, we can proceed with a complete desktop simulation to test whether the system will work properly under different test conditions.

When running this simulation, the model not only provides the estimated values of relative speed and object distance, but also visualizes the spectrum of transmitted and received signals, as shown in Figure 6.

A first simulation running under ideal conditions (absence of noise and distortion) shows that speed and position can be detected correctly for all targets in use. This simulation validates the test

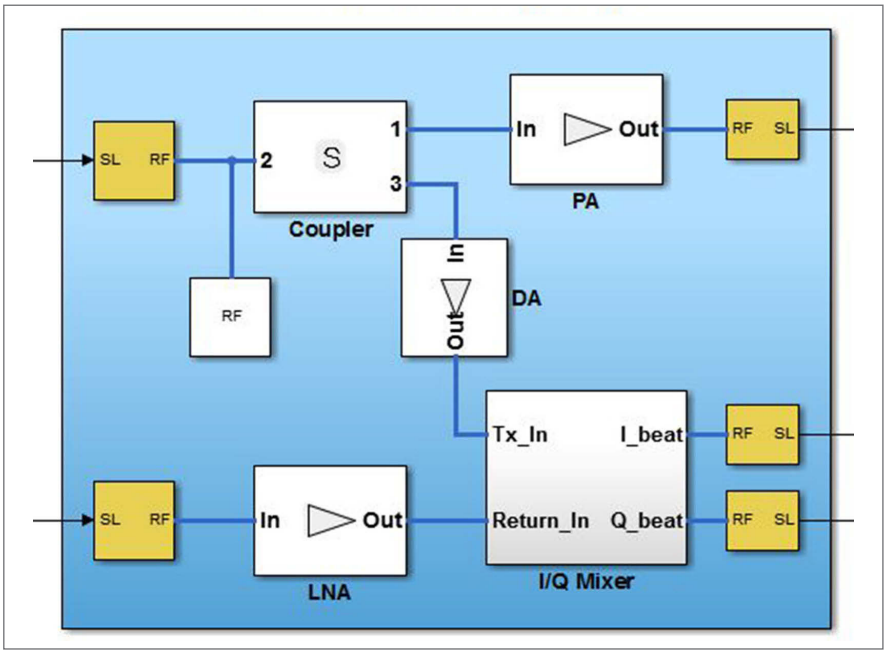

Figure 4 | RF elements modeled in Simulink using SimRF circuit envelope blocks.

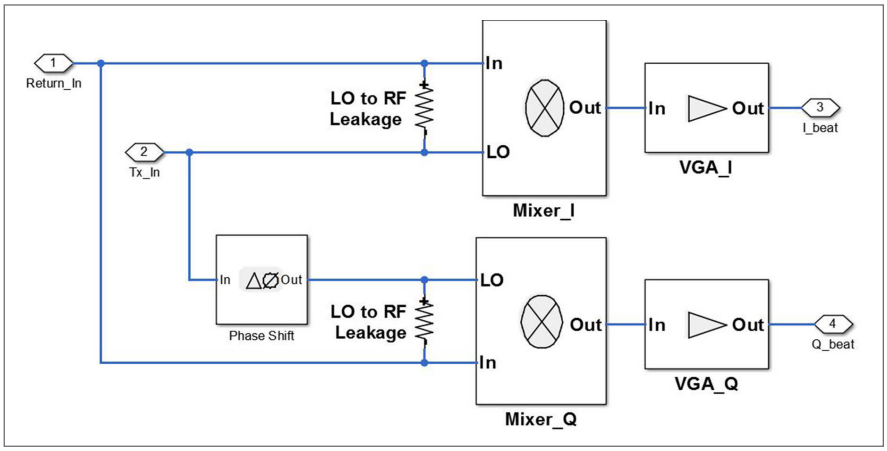

Figure 5 | Structure of the I/Q direct conversion mixer.

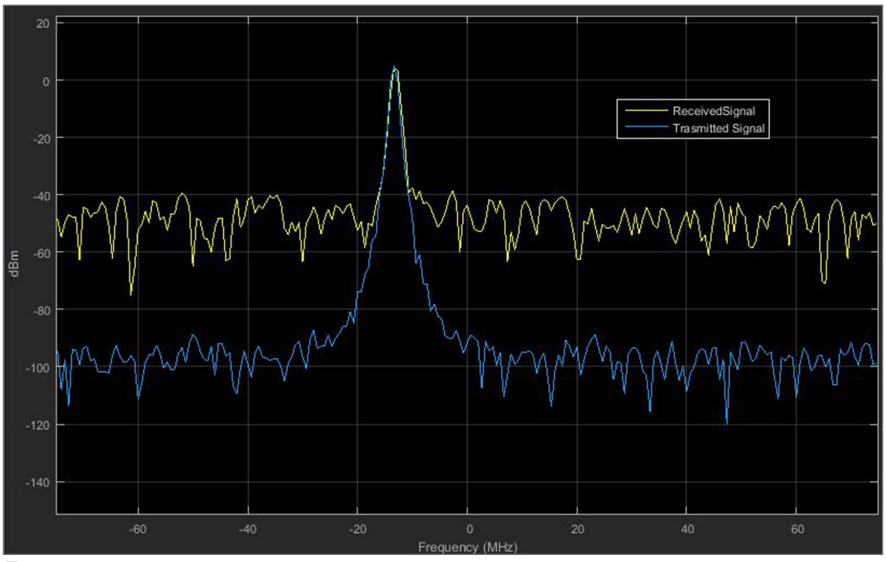

Figure 6 | Spectrum of the transmitted and received signals.

environment and the DSP algorithms. For subsequent simulations with added transceiver nonlinearity and noise, the radar deviates from the ideal behavior and cannot detect cars when they are far away.

After increasing the isolation of the mixer and the gain of the power amplifier, the radar system extends the detection range, and the simulation once again correctly estimates the target speed and range.

It is necessary to carefully trade off the gain of the different stages in order to avoid having the receiver operating in saturation. This model allowed us to simulate using a different set of parameters. It also helped us to select the suitable components for the radar implementation and to verify their impact on the radar performance.

Conclusion

This article has covered the modeling and simulation of a complete FMCW radar system for automotive active safety applications using a MATLAB-based tool chain. The proposed workflow allows us to simulate RF components within a complete system-level model, including DSP algorithms. This approach reduces both the time needed for radar development and the complexity of system tests, making the development cycle less costly.

To learn more, see Phased Array System Toolbox at se.mathworks.com/products/phased-array.

References:

[1] Design and Verify RF Transceivers for Radar Systems. Giorgia Zucchelli, MathWorks. mathworks.com/videos/design-and-verify-rf-transceivers-for-radar-systems-81990.html.

[2] Automotive Adaptive Cruise Control Using FMCW Technology. mathworks.com/help/phased/examples/automotive-adaptive-cruise-control-using-fmcw-technology.html.

[3] Karnfelt, C., et al. 77 GHz ACC Radar Simulation Platform, IEEE International Conferences on Intelligent Transport Systems Telecommunications (ITST), 2009.

John Zhao *is a product manager at MathWorks.*

Marco Roggero *is an application engineer at MathWorks GmbH.*

Giorgia Zucchelli *is a technical marketing manager at MathWorks.*

MathWorks
www.mathworks.com
john.zhao@mathworks.com • marco.roggero@mathworks.de
giorgia.zucchelli@mathworks.nl
@MATLAB
www.linkedin.com/company/the-mathworks_2
www.facebook.com/MATLAB
https://plus.google.com/+matlab/posts

MicroTCA.4 continues culture of COTS signal processing at DESY

Interview with Holger Schlarb and Michael Fenner, DESY

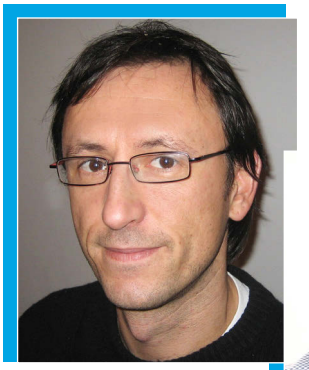

Holger Schlarb

Michael Fenner

In this interview with engineers in the Machine Beam Controls Group at DESY, Holger Schlarb and Michael Fenner discuss the data acquisition requirements of the accelerator community, as well as how a tradition of using COTS signal processing solutions continues to improve uptime and maximize performance for some of the most complex machines on Earth.

Give us a little bit of background on the Machine Beam Controls Group at DESY.

SCHLARB: We are in charge of all of the fast feedback systems that control the electron beam in the DESY accelerators. That includes orbit stabilization systems and arrival time stabilization, RF controls of the accelerating structures, timing distribution, RF synchronization, and optical synchronization down to the femtosecond level. Our part of the group mostly does the analog and digital hardware design for systems that preprocess data and then digitize it later on. We also do quite a bit of digital design because of the high data rate and the feedback system topologies that are built into the FPGAs.

For the accelerator, high data rates are required to optimize the slow controls or to debug the machine, but most of the data is just very strongly compressed at the end and then thrown away, so there's not really physics behind it, it's just to get the maximum out of the controls or to increase reliability. What we do is take rather high sampling rates of all analog data – typically 125 megasamples per second (MSPS), 200 MSPS, and sometimes 500 MSPS at 14-16-bit resolution. This is then compressed in FPGAs. So there is a first data reduction that is done, then it's shipped over to the CPU and a second data reduction is done, and what's remaining is what will go into the actual system. This data will be kept for a week or two weeks before being further compressed.

FENNER: We also have extremely high demands regarding analog performance, so we need an environment that allows extremely fast analog design – we're talking about low noise, we're talking about space for the analog electronics. At the same time on the digital side we have a need for very high data throughput, so digital links with speeds up to 10 Gbps using PCI Express Gen3. In addition to that we have high computing power requirements, so we use Intel Core i7 CPUs inside the system.

SCHLARB: What you need, specifically for analog performance, is very good clocks going into the analog-to-digital converters (ADCs), otherwise you ruin your data and you have to clean it up locally. Also, everything is triggered, so the machine is always synchronously running, so there is a trigger sent out, and at that time every station knows the electron beam is coming. The precision we need here is on the order of 10 picoseconds for many applications, and some high-end applications deal with sub-hundred femtosecond clocks. So basically we try to reach the upper limits of the ADCs, which is on the order of 50-60 femtoseconds, which is pretty demanding and not easy to achieve.

What we've observed over the past few years is that the power of FPGAs has increased so dramatically that we can now do things that we couldn't have done 5 or 10 years ago. This allows us a much higher sampling rate for preprocessing, which gives us more capability in controlling the machines to get more precise data. With the improvement in digital electronics, suddenly there are much higher data rates we have to cope with. The other thing we see that increases the data rate is that we are now able to use camera systems to monitor the beams and beam sizes, which is also contributing to these really high data rates. This was not feasible a few years ago. So the hunger for having more bandwidth, more processing power comes from these developments and allows us to be in a better position when managing the controls, which is really an advantage for rather little money I would say because when you look at the investment cost of the European XFEL that we're currently building of about 1 billion euros, machine beam controls is maybe 5 percent of that. So the development and the entire volume of investment is 20-25 million euros, which is a rather small amount but you can gain a lot if you just increase the reliability by 5 percent. You make back the beam controls investment after just a few years of operation.

Given these factors, can you provide insight into DESY's involvement in MicroTCA.4?

SCHLARB: In 2005 the decision was made that for the next collider we would move away from the VME architectures that are commonly used in our community because the bus is parallel, it's too slow, it's very noisy, and all of the reliability issues would not have been met. So this time we decided to look at the xTCA family because of reliability reasons, and in 2007 the decision was made to develop with xTCA for the European XFEL accelerator.

We had two groups, one that looked at MicroTCA.0 (mTCA.0) and one that looked at AdvancedTCA (ATCA). The conclusion was that the form factor of ATCA was too big because we have a lot of distributed systems – we have

200 racks in the XFEL distributed over 3 km, so one about every 10 meters. ATCA was not really suited for that because there are a lot of cables you have to bring to one point, and mTCA.0 wasn't suitable because the PCB area was not enough for the analog preprocessing electronics to be installed. At this point there was an effort triggered by the Stanford Linear Accelerator Center (SLAC) and DESY, along with many others, to adapt the xTCA family to the physics community. This was mTCA.4, which was driven by need to develop a specification that enabled all the benefits of a telecommunications architecture like reliability and high data throughput, but also incorporate the demands of the physics community. The outcome was a platform that is a little bit bigger, but especially that also made the rear transition modules (RTMs) larger as the RTMs were very small at that time.

FENNER: A unique selling point of mTCA.4 is that it provides as much space for the analog part as for the digital part. That is one point that is not found in, for example, ATCA, and in addition to that we paid a lot of attention to getting clean power supplies from the market. Because of our high requirements on the analog side, we have much stronger requirements for noise in such a system, which is why the power supplies that were good enough for digital mTCA.0 applications were not suitable for us. So we put a lot of effort into pushing the manufacturers to develop very clean and very low noise power supplies, and we supported the rack manufacturers in making systems that were low noise in terms of EMI.

What also differentiates xTCA from other standards is the full management of the boards, so supplying power, removing power, supervising temperature, increasing fan speeds if something gets too hot, hot plug capability – all of this is something that is new compared to old systems like VME. It's of high importance that all of this exists, because without this we would not have such high reliability in the system.

SCHLARB: As I already mentioned, we have about 200 racks and 5,000 complex electronic AdvancedMC (AMC) boards with a lot of components on them, and the other thing is that these electronics are installed in the accelerator tunnel so there is limited access. We have access once a week for a few hours, so reliability and diagnostics is very important for us because if something breaks we have to be able to fix it externally, or at least be able to diagnose what is broken so that on the next maintenance day we can exchange the particular component and immediately continue running. For the controls group this is very important because there are only one or two dozen people that have to maintain all of this. There, standardization is really important so that not everyone is using specialized electronics that have different interfaces, and so on. Standardization is what makes it possible to run such a complex system with a very limited number of personnel.

Now, we are pretty happy. At the moment we can say that this standard meets the analog performance requirements, and we also have the reliability and digital performance from the telecommunications background. We are in large-scale production now, so we are ordering all of the boards and starting to install them into the facility.

What other improvements does mTCA.4 include?

SCHLARB: What we found for RF in the mTCA.4 specification as it's written right now is that we have to deal with very complex cable management because we put a lot of RF cables into the system, and there were a lot of redundant cables. We introduced a second digital transfer AMC backplane, but behind this is an analog backplane that transports only analog signals that are very high quality. This is becoming part of the mTCA.4 standard, and it basically allows us to ease cable management because it reduces the number of cables you need. This is of general interest because the radar community or the lidar community, anyone who deals with a lot of RF signals or has a lot of sensors to integrate runs into similar problems. The RF backplane allows you to reduce this cable management issue significantly, and there are also some additional modules for that. We aren't limited in performance compared to proprietary, isolated RF modules, as we can transport RF signals with femtosecond precision over this RF backplane, which meets the requirements of the market.

A particular point of interest as well is the separation between AMCs and RTMs, because AMC boards are usually more complex digital boards that have big FPGAs and modern ADCs, while the RTM is more for analog, which are cheaper and have much longer lifecycles because analog electronics don't change quickly (it's normally about 10-20 years before you reach end of life). For the AMC boards you have much faster turnaround time because of the FPGAs, CPUs, and DSPs that are used, so it's more like two or three years until new ones reach the market. This separates the lifecycle management of your electronics, and it also separates the developer groups – so one is more of a digital developer group, and the other is analog electronics people that have no clue about how MicroTCA works, but they know analog development. That's a big plus for this standard.

Why are standards-based solutions so important for an organization like DESY?

SCHLARB: The open architecture is very important for us because we have to worry about vendor lock-in situations. We have to be very careful that we can always supply our machine with electronics, so we have to make sure we protect ourselves against companies that might make different strategic decisions. So multiple vendors is very important for us for our components, because obviously we don't want to build standard components ourselves. We are not building CPUs – there are highly specialized companies for that. There is also a pretty long tradition at DESY of supporting open standards.

***Holger Schlarb** is Group Leader of Machine Beam Controls at DESY.*

***Michael Fenner** is Chief Engineer for Digital Electronics Development in the Machine Beam Controls Group at DESY.*

DESY
www.desy.de | @desy
www.youtube.com/channel/UCcTXBDaaoJ8JLIxI9-voMRg

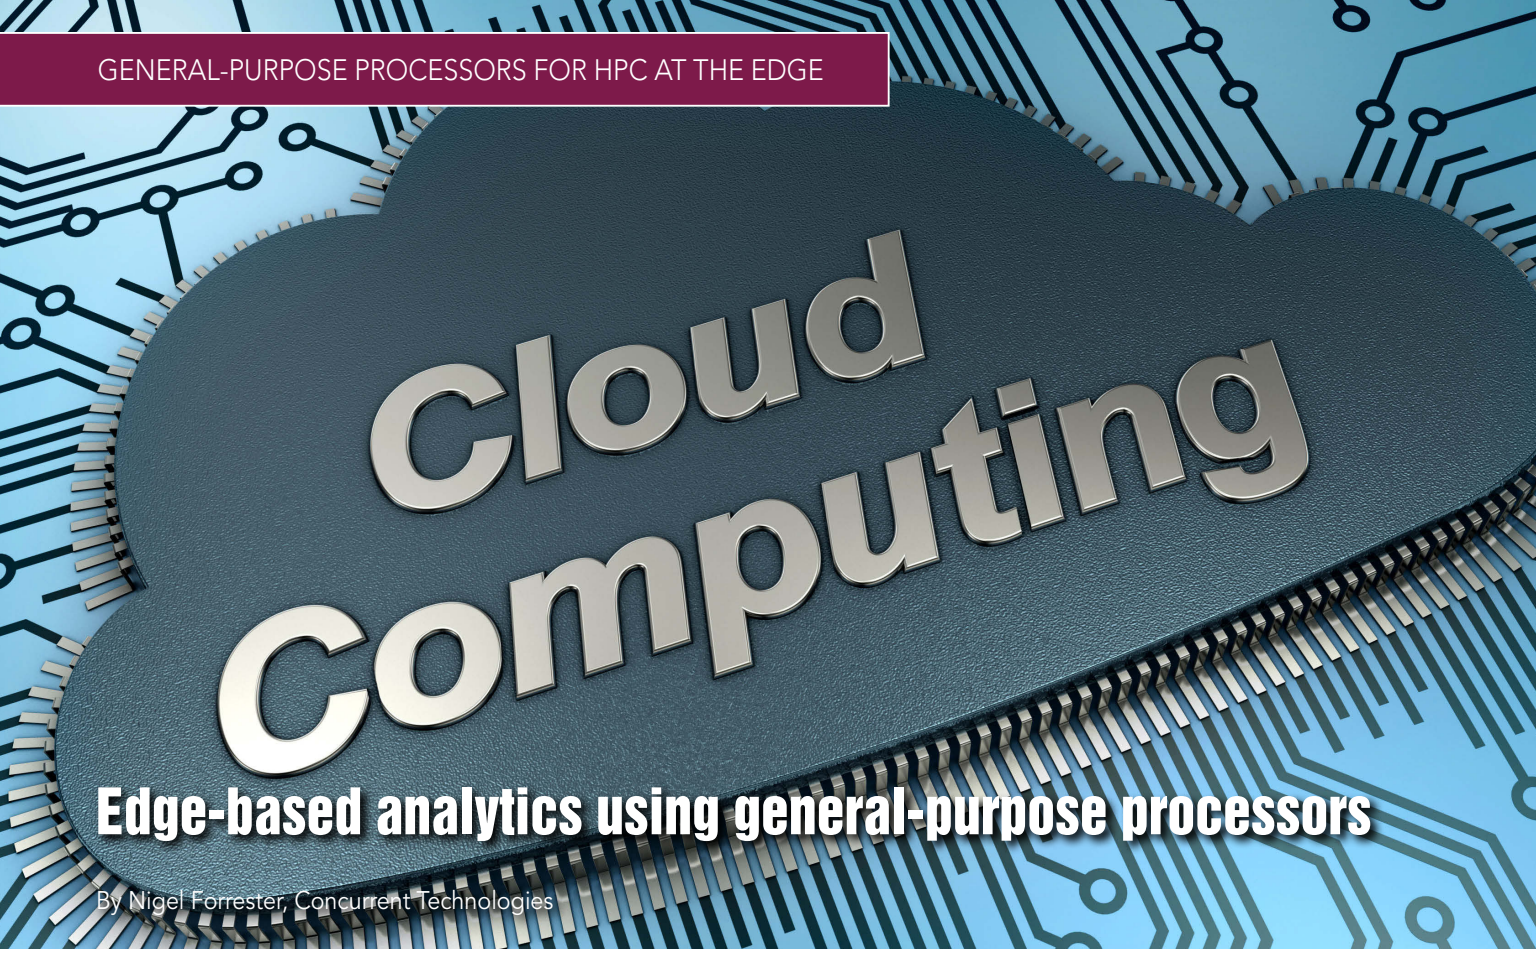

Edge-based analytics using general-purpose processors

By Nigel Forrester, Concurrent Technologies

Increasing demand for bandwidth at the network edge is driving the need for easily programmable, low-cost signal processing systems that can scale as data volumes grow. This is leading to the deployment of general-purpose processors in domains traditionally occupied by digital signal processor (DSP) and FPGA devices, with off-the-shelf module vendors innovating on low-latency interconnect architectures and middleware solutions that unlock the value of commercial processors in the mobile and machine-to-machine (M2M) era.

By definition, embedded systems are usually situated close to the application that they are controlling or measuring – irrespective of the end market – whether it be communications or industrial automation, military/aerospace or transportation, science or security. Embedded systems are often constructed from a combination of open standards-based building blocks, including general-purpose processors (typically Intel x86-based), FPGAs, DSPs, and I/O devices. A system will gather data from various input sensors, execute some heavy-weight number crunching and analysis using appropriate algorithms, and then send the output to a general-purpose processor for local storage, display, and control, as well as an uplink to a centralized data repository. General-purpose computer boards are relatively simple to program, as they are essentially long-life-cycle versions of the same x86 architecture processor boards widely used in desktop and laptop computers, which means that development tools and experience abound. On the other hand, FPGA and DSP devices are relatively specialized and can be more difficult to program, especially as there are a lot fewer engineers with the necessary experience. As the amount of data grows and the need to process it faster and more locally increases, designers need to think about potential solutions that enable general-purpose processors to do more FPGA and DSP tasks.

Maximizing throughput and reducing latency on general-purpose processors with RapidIO

One of the biggest challenges in enabling general-purpose processors to compute large, highly complex, large data sets in a deterministic manner has been coupling them tightly enough. Historically, general-purpose processors have been architected to stand alone, as in a laptop or desktop, or racked together in a datacenter-type environment by using Ethernet or InfiniBand connections via top-of-rack (TOR) switches. Standalone devices typically connect via relatively low-bandwidth connections like Wi-Fi or 4G wireless, whereas datacenter-based kits will have much higher bandwidth connections but are still unlikely to meet the very-low-latency and deterministic performance needed for a real-time application. As an example, a payment system needs to have sufficient bandwidth to service the customer base, but while the number of transactions per second may vary widely during the day, it is not critically important if the time to process a transaction varies by a few milliseconds. Indeed, it's likely that the end customer won't be significantly affected if there is a delay of a second in transaction time. On the other hand, a real-time system will fail if there are delays of milliseconds or even microseconds, depending on the type of data. A good example of a real-time system is a wireless base station that must capture and

process data on the fly so that calls are not dropped or scrambled; it does this by guaranteeing the latency between elements.

One solution for embedded signal processing applications based general-purpose computers is a module architected to have a RapidIO interconnect fabric across the backplane rather than the more generic PCI Express or Ethernet connections. RapidIO has some key advantages:

- it is designed to scale to very large systems unlike PCI Express which is based on the concept of a tree structure with a defined root RapidIO defines direct memory mapped read/write capability allowing elements to share data widely
- RapidIO was designed for low latency applications and has been successfully deployed over a number of years in real time applications

Using bridge devices it is possible to convert the PCI Express connectivity that comes standard off of most general-purpose processors to single or double x4 RapidIO links, as shown in Figure 1.

While having a bridge device to convert between PCI Express and RapidIO enables a much better degree of scalability and allows external devices to more easily share information between multiple processors, it's not obvious that this setup will perform any better than the native PCI Express interface. The situation is potentially exacerbated because it appears to add another level of complexity, especially as many applications are written with the expectation that information is shared using Ethernet socket connections. Fortunately, there is a solution in a middleware layer called FIN-S. FIN-S allows applications to communicate between processors using standard socket-based interfaces without any knowledge of the underlying RapidIO fabric, and also works over PCI Express if this is the customer's preferred backplane interface technology.

Interestingly, there are some useful side benefits when using FIN-S as well. The graph in Figure 2 plots the data

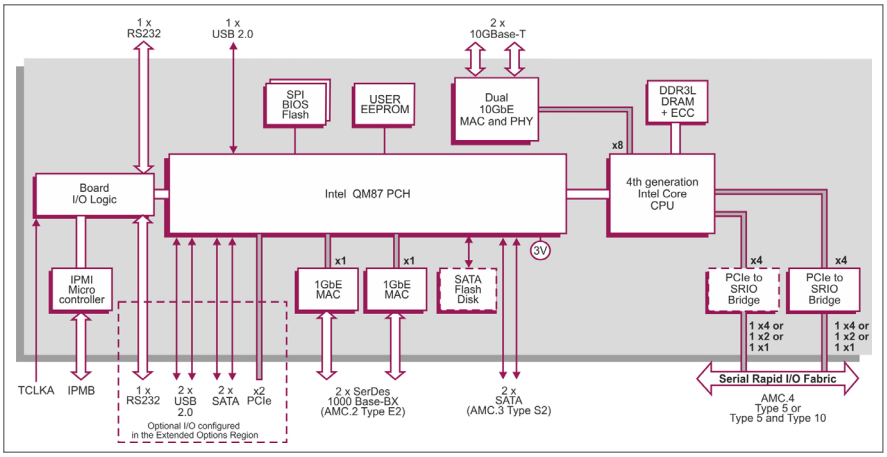

Figure 1 | The AM C1x/msd from Concurrent Technologies is based on a 4th generation Intel Core i7 and uses bridge devices to convert the processor's native PCI Express connectivity to 1 x4 or 2 x4 RapidIO lanes, as shown in the bottom right of this block diagram.

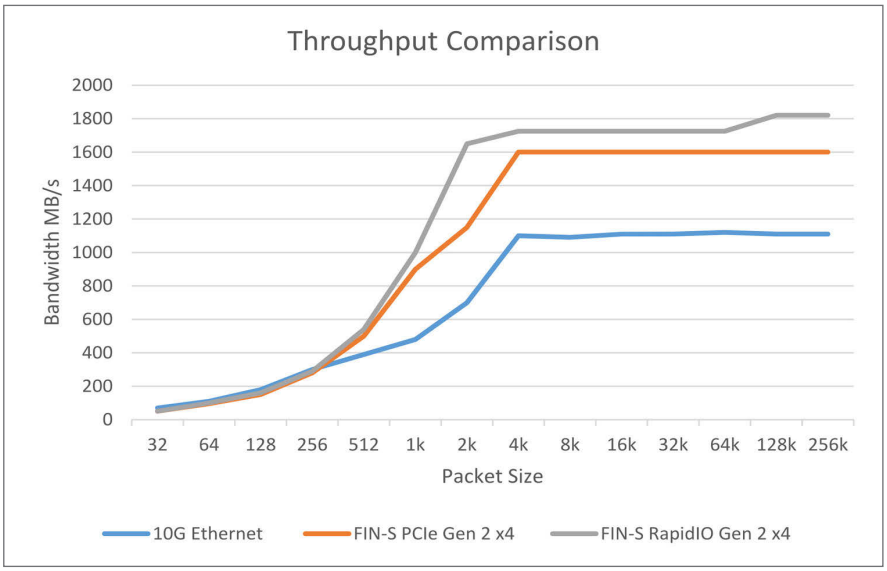

Figure 2 | With FIN-S middleware it is possible to achieve throughputs of up to 1.8 GBps over RapidIO Gen2 while using only 5 percent of a general-purpose processor's resources.

bandwidth obtained using the standard Linux "iperf" network benchmarking utility for various packet sizes using a 10 Gigabit Ethernet (GbE) adapter with a x8 Gen2 PCI Express link to the CPU, FIN-S running on PCI Express Gen2 x4 links, as well as FIN-S running on 5 Gbps RapidIO Gen2 x4 links. As can be seen, when using FIN-S a sustained throughput of more than 1.8 GBps over RapidIO Gen2 is achievable using the standard TCP/IP stack with a CPU utilization of around 5 percent. This level of CPU utilization, and particularly bandwidth, is better than using either PCI Express or 10 GbE.

Of more interest for real-time systems is the maximum latency. Again, by using FIN-S over RapidIO, good results can be achieved irrespective of the packet size (Figure 3, page 26).

Supercomputing at the network edge

Taking this a step further, there has been a lot of innovation over the last few years in using graphics processors to augment the general-purpose processing element. Highly parallel, power-efficient general-purpose graphics processing units (GPGPUs) are now available that are easy to program and can be used to accelerate edge analytics applications. For example, the Intel Core i7-based AM C1x/msd module from Concurrent Technologies contains both a central processing unit (CPU) and up to 20 GPUs. Rather than being used just for display purposes, these GPUs can be used for general-purpose computing using OpenCL.

While this additional processing capability is good to have, a more radical solution for significantly improved performance might be to use a dedicated GPGPU module. Concurrent Technologies' AG A1x/m1d is based on multiple NVIDIA Tegra K1 mobile processors that use the same NVIDIA Kepler GPU core designed into supercomputers around the world. Fitted with either two or four Tegra K1 processors, each of which provides 192 Kepler CUDA GPU cores and a 4-PLUS-1 quad-core ARM Cortex-A15 CPU, the AG A1x/m1d has a maximum performance of 1.3 teraflops (TFLOPS). Since the AG A1x/msd is based on the same open-standard AdvancedMC (AMC) form factor as the AM C1x/msd, embedded system integrators can build heterogeneous systems using combinations of CPU and GPGPU modules to suit their application – GPGPU modules offer a better power-to-performance ratio for applications that can be channelized or split up into of fragments, while CPU modules offer better command and control capability with more local storage options and a wider I/O range (Figure 4).

Combinations of these easily programmed processing modules can be used to create high-performance embedded computing (HPEC) devices that can be deployed at the edge of a network or system. Such devices enable a completely new class of applications that were not easily (or commercially) possible before, as local data had to be sent to a remote cloud-based system for storage and processing. Now, so-called "supercomputing at the edge" with off-the-shelf modules precludes the need to send data back to a datacenter for processing simply to have the results returned to the same location because the bulk of analysis can be performed locally.

For example, many wireless network operators are looking for ways to maximize available bandwidth, especially during heavy demand periods and in high-density geographies. Orange Silicon Valley, the research arm of one of the world's leading telecommunications operators, made an announcement earlier this year in conjunction with IDT and NVIDIA about a program they are developing to provide supercomputer performance at the edge of their existing wireless networks as part of 4G base station deployments. Colocating significant processing capability within a base station-type environment enables a subset of "big data" analysis at a local level without the overhead of transferring the data back and forth to cloud-based servers. This type of local computing capability opens up new opportunities, particularly the possibility to offer location-based services that are easily scalable and support low latency.

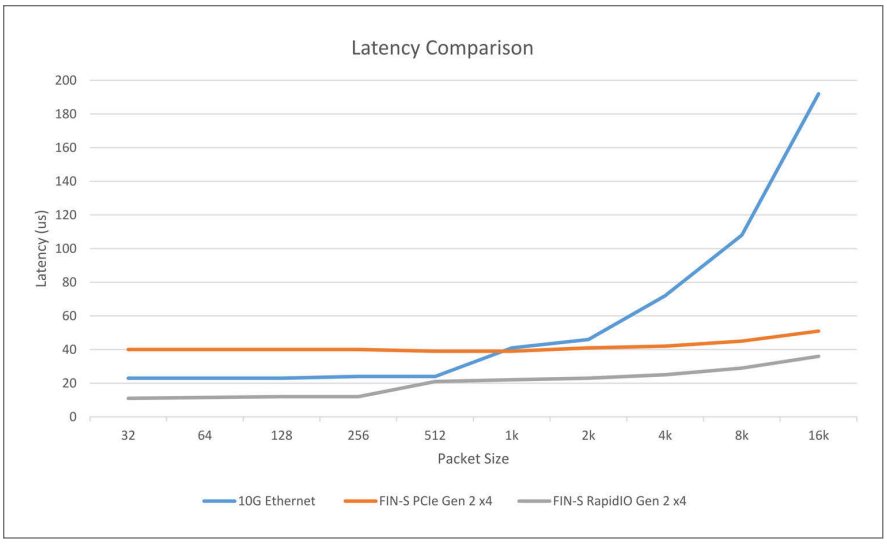

Figure 3 | The use of FIN-S middleware also helps reduce latencies over alternative fabric architectures, regardless of packet size.

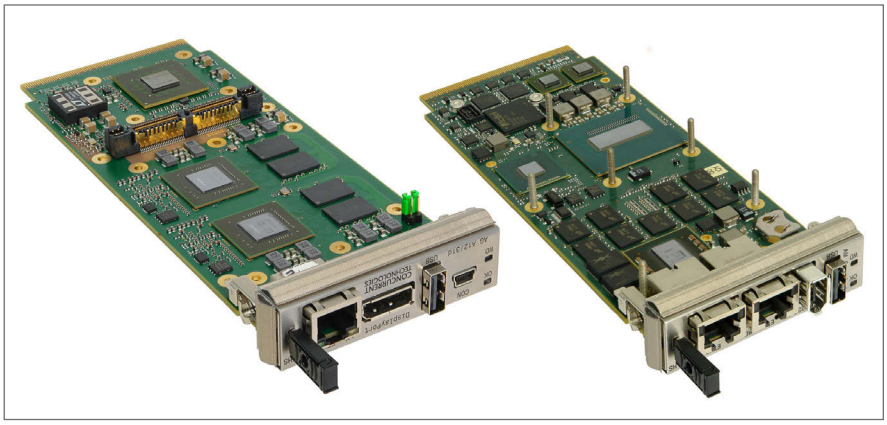

Figure 4 | Both based on PICMG's AdvancedMC (AMC) architecture, the AG A1x/m1d (left) and AM C1x/msd (right) provide system integrators with a modular approach to optimizing power and performance or command and control.

General-purpose innovation

The demand for edge-based analytics is increasing, both for human consumers and the increasing amount of M2M data capture. With the advent of more easily programmed and scalable building blocks, it is possible to create innovative solutions based on general-purpose CPU and GPGPU capabilities that meet application needs.

Nigel Forrester *is Technical Manager at Concurrent Technologies.*

Concurrent Technologies
www.gocct.com
sales@gocct.com

A.R.BAYER
DSP Systeme GmbH

Passenger²

Passenger² – Efficient OFDM Modem for Wireless and Wired Communication

PASSENGER² is a core technology for wired or wireless data links based on the industry's most advanced C-OFDM modem implementation boasting an efficiency greater than 10bits/s/Hz.

Typical use cases include data links over the air (mobile, portable, or stationary), or in existing infrastructure, e.g. 4-wire lines, coax cables, power lines, sub-sea cables, twisted pairs, and even barb wires.

Applications include intercoms (audio, video, data), ad-hoc disaster communication, industrial control, telemetry and monitoring, WAN extensions, kiosks and remote access control.

The field-proven algorithm is extremely robust and available for catalog signal processors such as SHARC and fixed-point DSP. It uses C-OFDM multi-carrier modulation (QAM-4 to QAM-4096) with up to 100 sub-carriers and 4D-Trellis coded modulation with Trellis shaping and external FEC.

Modem hardware with complete infrastructure and software is available for integration into target applications.

The modem algorithm is also available as object code for SHARC and other DSP.

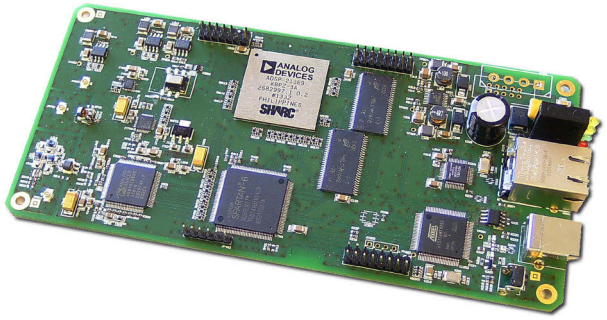

FEATURES

- **Transmission Modes:** FDD, TDD, Simplex, Half Duplex
- **Frequency Range:** 10kHz to 1000kHz in wired, 30MHz to 1GHz in wireless applications
- **Channel Bandwidth:** 12.5kHz up to 200kHz
- **Data Rates:** up to 2.048 Mbps
- **Wireless Transmission Range:** 30mls/50km max. (P2P), LOS, P_{mean}=30dBm, FDX
- **Wired Line Operation:** at up to 60dB Attenuation

signal-processing.mil-embedded.com/p372791

A.R. Bayer DSP Systeme GmbH
http://dsp-sys.de/ofdm-modem

gmbh@dsp-sys.de
de.linkedin.com/in/andreasbayer
+49-211-2714630
@AndreasBayer

E-CAST

Keeping it cool: Solving military electronics thermal management challenges

Sponsored by Kontron

Radar, electronic warfare, unmanned aircraft, combat aircraft avionics, etc. all demand high-performance signal processing to meet capability requirements. However, while faster processors and FPGAs create unprecedented performance, they also generate excessive amounts of heat. Aspects such as ambient temperature, altitude, power density, and power dissipation have to be evaluated early in the design process. This presentation focuses on the principle techniques for thermal management, their relative efficiencies, and examples of real-world solutions.

GO TO E-CAST:
ECAST.OPENSYSTEMSMEDIA.COM/550

VEROTEC INTEGRATED PACKAGING
Pluggable PSUs and fan trays

PICMG 2.11 compliant for cPCI systems
Autoranging AC - DC
200 - 500W
General purpose
20 - 240W, 1 - 4 outputs
AC - DC and DC -DC

1U 19" fan trays
Intelligent, filtered
and standard versions
3 and 6 fan versions
115/230VAC
12/24 and 48VDC

Continuing 50 years of excellence in functional and elegant enclosures for cPCI, VME, VME64x, VPX, other major bus structures and general electronics

VERO ELECTRONICS PACKAGING ➤ apw ➤ VEROTEC Electronics Packaging

Ph: 603.821.9921 • sales@verotec.us • www.verotec.us

Annapolis Micro Systems

CoreFire Next Design Suite

The CoreFire Next Design Suite (compatible with all Annapolis Virtex™ 7 and Altera Stratix® V FPGA processor and I/O boards) is a dataflow-based development system that brings new levels of ease and speed to FPGA programming on Annapolis Micro Systems, Inc. high-performance motherboards, I/O cards, and mezzanine cards.

The CoreFire Next environment supplies user-made connections between ready-made programming modules, or cores, and manages multiple domain requirements automatically. CoreFire Next eliminates the need for hardware design languages: the user simply creates dataflow diagrams by dragging and dropping cores, or building blocks, from the libraries, and connecting their ports. Cores automatically work together to handle synchronization, manage clocks and other low level hardware signals, and guarantee correct control by design. CoreFire Next allows standard data types (see Data Types and Values) and supports data type propagation, where modifying a data type will automatically propagate through the rest of the CoreFire Next design. As a result, CoreFire Next allows the user to program and debug complex FPGA designs at a high level of proficiency.

CoreFire Next's drag-and-drop method of building designs allows for ease of use, which helps make the tool easy to learn. CoreFire Next presents the user with a simple way of visualizing designs, rather than extensive and confusing code. Because of this, the user does not have to be a skilled or experienced digital hardware designer. Designers of many different disciplines can use CoreFire Next to create applications.

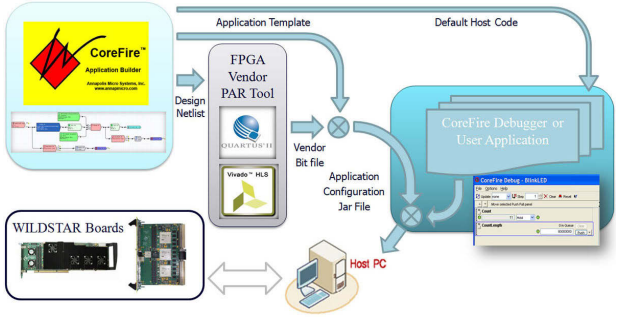

FEATURES

- Build Designs for FPGAs on WILD™ Boards
- Works from High Level, Data Flow Concept of the Application
- Combines GUI Design Entry and Debug Tools with Tested, Optimized CoreFire Next™ IP Cores
- Drag and Drop High and Low Level Modules
- CoreFire Next™ Modules Incorporate Years of Application Development Experience – Highly Optimized and Tested
- CoreFire Next™ Tools and Modules are Intelligent
- Modules Automatically Handle Synchronization
- Manage Clocks and Other Low Level Hardware Signals
- Guarantee Correct Control by Design Modules "Know How" to Interact with Each Other
- Board Support Packages Incorporate Hardware Details of the Boards – Invisible to Users
- Single Precision Floating Point, Integer and Floating Point Complex Data Types and Array Types. Provides Java File
- Supports Conversion Between Data Types – Bit, Signed and Unsigned Integers Single Precision Floating Point, Integer and Floating Point Complex Data Types and Array Types
- Integrates with Matlab™ Simulation Flow
- Works with all Annapolis Virtex™ 7 and Altera Stratix® V FPGA processor and I/O boards

Benefits

- Save Time to Market
- Save Development Dollars
- Easy to Learn, Easy to Use
- Works with Proven COTS Boards
- Concentrate on Solving Your Problem
- Reuse Your Design
- Training Classes, Application Support

signal-processing.mil-embedded.com/p372787

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

WILD Data Storage Solution

When Storage capability is needed, Annapolis offers the highest density OpenVPX storage solutions on the market with up to 9.3 TB of capacity in a single 1" slot with up to 4.5 GB/s of write bandwidth. It also features a removable hot swappable canister with a connector rated for 10,000+ mating cycles. The WILD Data Storage Solution comes with standard images to support XAUI, 40GbE and AnnapMicro Protocol (Annapolis low FPGA utilization, full flow control protocol ideal for inter-FPGA communication).

The WILD Data Storage Solution is comprised of two pieces fitting in a single 1" OpenVPX slot, the "storage canister" which holds up to 12 1.8" SATA disks, and the "Storage Carrier" that plugs into the VPX backplane and holds the disk canister.

The Storage Carrier/Canister is specifically designed to support 10,000+ insertion cycles of the disk canister for frequent drive removal. Both the canister and the entire assembly (Storage Canister + Storage Carrier) are also hot swappable for minimum system down time and highest reliability. This OpenVPX compliant payload card supports 40Gb serial I/O on the VPX Data Plane on P1 to support four channels of 40GbE (proper backplane required for faster rates).

To ensure safe and reliable processing, WILD Data Storage Solution boards come equipped with a proactive thermal management system. Sensors across the board monitor power and temperature, with automatic shutdown capability to prevent excessive heat buildup. WILD Data Storage Solution boards are built with a rugged, durable design. Sensors can be accessed with a chassis manager (ChMC).

New heatsinks have been tested with great success on WILD Data Storage Solution boards. These larger heatsinks also act as stiffeners for the boards, making them sturdier.

FEATURES

› **General Features**
- 9.3 TB of Storage Per Each 6U VITA 65 Compliant OpenVPX Slot
- Up to 4.5 GB/s Write and Up to 5 GB/s Read Bandwidth (write bandwidth determined by system environmentals)
- Scalable Depth and Bandwidth
- Hot Swappable Drive Canister with 10,000 Insertion Cycles & Hot Swappable Carrier (exclusive to WILDSTAR OpenVPX EcoSystem)

› **Backplane I/O**
- Up to 40Gb Ethernet on each of Four Fat Pipes on P1, for a total of 20GB/s on P1
- 1 Additional Fat Pipe on P4 providing QSFP+ connection via RTM
- 1Gb Ethernet Connection on P4

› **System Management**
- Client/Server Interface for WILDSTAR FPGA Boards and Linux and Windows-based CPU systems
- Extensive System and Drive Diagnostic Monitoring and Configuration over 1 Gb Ethernet via P1 and P4 Ethernet
- Standard Intelligent Platform Management Interface (IPMI) to Monitor Current, Voltage and Temperature
- Front Panel Status LEDs for all 12 SSDs and all Backplane Control and Data Plane Connections

› **Physical Features**
- 6U OpenVPX (VITA 65) Compliant, 1" VITA 48.1 spacing
- Supports OpenVPX Payload Profile: MOD6-PAY-4F1Q2U2T-12.2.1-n
- Integrated Heat Sink
- Air Cooled with Product Path to Conduction Cooling

signal-processing.mil-embedded.com/p372456

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

WILD OpenVPX 40 Gb Ethernet and FDR Infiniband Switch

The WILD OpenVPX 14Gbit Switch Card supports 1GbE, XAUI, 10GbE, 40GbE, 56GbE, and SDR/DDR/QDR/FDR Infiniband. It has dual 1/10GbE SFP+ front panel control plane uplinks and eight front panel data plane QSFPs. It supports up to 20 1GbE backplane control plane connections and 20 XAUI/40GbE/Infiniband data plane connections.

The WILD OpenVPX 14Gbit Switch Card is extremely versatile since it is capable of switching both Infiniband (SDR, DDR, QDR, FDR) and Ethernet (1Gb, 10Gb, 40Gb, 56Gb) traffic with up to 4 Tb/s of non-blocking switching capacity. The WILD OpenVPX 40 Gb Ethernet and FDR Infiniband Switch also supports chassis management and can act as a Chassis Manager (ChMC). The 1Gb Ethernet control plane supports up to 20 backplane ports and two front panel SFP+ which can run at 1GbE or 10GbE. The data and control planes are located on different virtual networks to ensure best performance on each.

Basic configuration is streamlined where all required features are selected by DIP switches. A front panel USB serial port allows configuration of management Ethernet interfaces if needed (DHCP is the default configuration). Software updates, if needed, are completed via a simple web interface which can also be disabled via USB serial console. Front panel status LEDs show the status of every switch port (link/activity) as well as overall status and health of the WILD OpenVPX 40 Gb Ethernet and FDR Infiniband Switch.

The front panel RJ45 10/100/1000 BASE-T Ethernet port is used for switch management and is connected directly to the on-board PowerPC. There is also an optional "in band" Ethernet connection from the PowerPC to the control plane. Note that not all switch configurations support the "in-band" management connection.

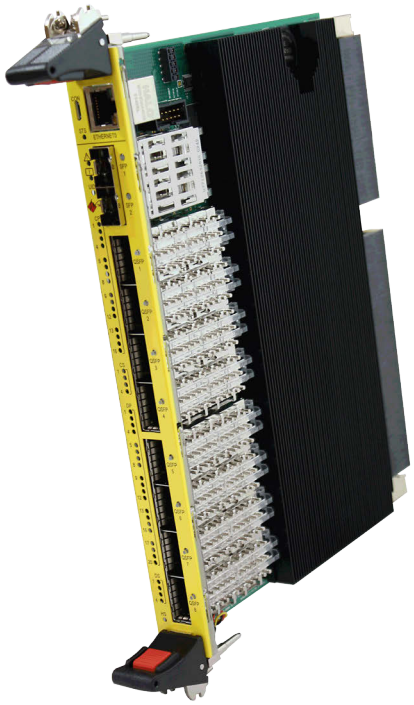

FEATURES

- **MultiProtocol Switch**
  - 1/10/40/56 Gb Ethernet and SDR/DDR/QDR/FDR Infiniband
  - Up to Four Tb/s Non-Blocking Switching Capacity with up to Eight Switch Partitions
- **Backplane & Front panel I/O**
  - Backplane Ports: Twenty High Speed Four Lane Data Plane Connections, Sixteen 1Gb Ethernet Lanes
  - Front Panel Ports: Eight QSFP+, Two SFP+, RJ45 Management Port, USB UART, Status LEDs
  - Each Backplane and Front Panel Port can be Configured for either Infiniband or Ethernet
  - Infiniband and IP Routing
  - Ethernet Gateways
- **System Management**
  - System Management using Intelligent Platform Management Interface (IPMI)
  - Diagnostic monitoring and configuration
  - Current, Voltage and Temperature Monitoring Sensors
  - Hot Swappable (exclusive to WILDSTAR OpenVPX EcoSystem)
- **Mechanical and Environmental**
  - 6U OpenVPX (VITA 65) Compliant, 1" VITA 48.1 spacing
  - Supports OpenVPX Slot Profile: SLT6-SWH-16U20F-10.4.2-n
  - Integrated Heat Sink and Board Stiffener

signal-processing.mil-embedded.com/p372674

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

WILDSTAR 7 for OpenVPX 3U

The WILDSTAR 7 for OpenVPX 3U contains one VX690T or VX980T Virtex 7 FPGA per board with up to 2 GB of DDR3 DRAM for 12.8 GB/s of DRAM bandwidth and up to 32 MB of QDRII+ SRAM for 8 GB/s of SRAM bandwidth. It has up to 1 million logic cells and 1.6 million multiplier bits per board.

These FPGA boards include a Xilinx Virtex 7 FPGA with 64 High Speed Serial connections performing up to 13.1 Gbps. There is two 36-bit QDRII+ SRAM interfaces clocked up to 500 MHz and two 32-bit DDR3 DRAM ports clocked at up to 800 MHz.

With included High Speed Serial (HSS) FPGA cores (including 40GBASE-KR), there is up 10 GB/s of bandwidth on the VPX data plane which can go directly to other VPX cards or to a switch, depending on backplane topology. In addition, there is up to 20 GB/s of bandwidth on the VPX Expansion Place. When using 40GBASE-KR, there is the added reliability of Forward Error Correction (FEC) to achieve a much lower Bit Error Rate (BER).

If IO is required, Annapolis offers extraordinary density, bandwidth and analog conversion choices. Each 3U card has 1 mezzanine IO sites which can support up to 2 WILDSTAR Mezzanine cards as well as a QSFP+ option (on WS7 and WS A5 board) that allows for 3 QSFP+ transceivers per slot. These options can be mix and matched to meet customer needs. Some configurations utilize a second slot (for example the QSFP+ option and WILDSTAR Mezzanine card used in a single IO Site).

WILDSTAR A5 and V7 FPGA boards are hot swappable allowing for more system reliability. This feature is unique to Annapolis and was developed because our experience with OpenVPX systems has shown it invaluable so a whole chassis does not need to be shutdown to remove a single board.

Annapolis OpenVPX FPGA cards include an on-board dual core 1.2 GHz PowerPC with direct FPGA 4x PCIe connection which can be used by customers for application requirements. It is also used query board health like FPGA temperature and power. It is connected to the OpenVPX control plane via 1GbE.

There are also plenty of user backplane signals available on the Annapolis 6U Rear Transition Module (RTM) such as LVDS, FPGA HSS, IRIG, Ethernet and clocking. RTM HSS is also capable of 10Gbps signalling and supports multiple channels of 40GbE.

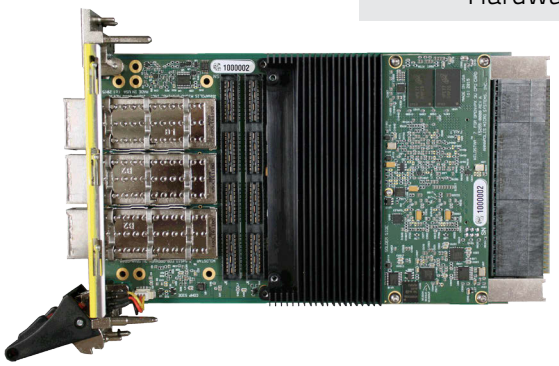

FEATURES

- **General Features**
  - One Xilinx Virtex 7 VX690T or VX980T FPGA
  - Up to 2 GB of DDR3 DRAM for 12.8 GB/s of DRAM bandwidth
  - Up to 32 MB of QDRII+ SRAM for 8 GB/s of SRAM bandwidth
- **Backplane I/O**
  - 24x High Speed Serial IO lanes to VPX Backplane (P1/P2) for 30 GB/s of Full Duplex Bandwidth
  - Two PCIe Gen3 8x Connections to VPX Backplane (P1)
  - Eight LVDS lines to P2
  - Backplane Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and user designed protocols
  - External clock and IRIG-B Support via Backplane
  - Radial Backplane Clock Support for OpenVPX backplane signals AUXCLK and REFCLK
    - Allows points-to-point, very high quality backplane connections to payload cards
    - Allows 10MHz clock and trigger from backplane to synchronize and clock compatible ADC/DAC mezzanine cards without front panel connections needed
    - Allows 1000s of analog channels across many backplanes/chassis to be synchronized via backplane
- **Front Panel I/O**
  - Accepts Standard Annapolis WILDSTAR Mezzanine Cards, including a wide variety of WILDSTAR ADC and DAC Mezzanine Cards
  - Three optional built-in Front Panel QSFP+ Transceivers running at up to 52.4 Gbps each for 39 GB/s of Full Duplex Bandwidth
  - Simultaneous QSFP and Mezzanine Card use
  - QSFP+ Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and userdesigned protocols
- **Dual Core Processor APM86290**
  - Host Software: Linux API and Device Drivers
  - Each core runs up to 1.2 GHz
  - 2 GB of DDR3 DRAM
  - 4 GB SATA SSD and 16MB NOR Boot Flash
  - 4x PCIe Gen2 connection to Virtex 7 FPGA
- **Application Development**
  - Full CoreFire Next™ Board Support Package for Fast and Easy Application Development
  - 10/40Gb Ethernet and AnnapMicro Protocol Cores Included
  - Open VHDL Model including Source Code for Hardware Interfaces
  - Open VHDL IP Package for Communication Interfaces
  - Chipscope Access through RTM
- **System Management**
  - System Management using Intelligent Platform Management Interface (IPMI)
  - Diagnostic monitoring and configuration
  - Current, Voltage and Temperature Monitoring Sensors
  - Hot Swappable (exclusive to WILDSTAR OpenVPX EcoSystem)
- **Mechanical and Environmental**
  - 3U OpenVPX (VITA 65) Compliant, 1" VITA 48.1 spacing
  - Supports OpenVPX payload profile: MOD3-PAY-2F4F2U-16.2.10-n
  - Integrated Heat Sink and Board Stiffener
  - Available in Extended Temperature Grades
  - Air Cooled with Conduction Cooled path
  - RTM available for additional I/O

signal-processing.mil-embedded.com/p372457

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

WILDSTAR 7 Conduction Cooled for OpenVPX 6U

WILDSTAR 7 Conduction Cooled for OpenVPX 6U boards provide up to two Xilinx Virtex 7 FPGAs per board with VX690T or VX980T FPGAs, up to 8 GB of DDR3 DRAM for 51.2 GB/s of DRAM bandwidth or up to 64 MB of QDRII+ SRAM for 32 GB/s of SRAM bandwidth. Up to 1.9 million logic cells and 3.3 million multiplier bits per board. Air or Conduction Cooled.

These FPGA boards include two Xilinx Virtex 7 FPGAs with 64 High Speed Serial connections performing up to 13.1 Gbps. The IO Processing Element (IOPE) FPGA has a choice of QDRII+ SRAM or DDR3 DRAM. The DRAM option has four 32-bit DDR3 DRAM ports clocked at up to 800 MHz while the SRAM option has two 72-bit QDRII+ SRAM interfaces clocked up to 500 MHz.

With included High Speed Serial (HSS) FPGA cores (including 40GBASE-KR), there is up 20 GB/s of bandwidth on the VPX data plane which can go directly to other VPX cards or to a switch, depending on backplane topology. In addition, there is 16 GB/s of PCI Express Gen 3 bandwidth on the VPX Expansion Plane with an 8x Gen3 connection to each FPGA through a non-blocking PCIe switch. When using 40GBASE-KR, there is the added reliability of Forward Error Correction (FEC) to achieve a much lower Bit Error Rate (BER).

If IO is required, Annapolis offers extraordinary density, bandwidth and analog conversion choices. Each 6U card has 2 mezzanine IO sites which can support up to four WILDSTAR Mezzanine cards as well as a QSFP+ option (on WS7 and WS A5 board) that allows for six QSFP+ transceivers per slot. These options can be mix and matched to meet customer needs. Some configurations utilize a second slot (for example the QSFP+ option and WILDSTAR Mezzanine card used in a single IO Site).

WILDSTAR A5 and V7 FPGA boards are hot swappable allowing for more system reliability. This feature is unique to Annapolis and was developed because our experience with OpenVPX systems has shown it invaluable so a whole chassis does not need to be shutdown to remove a single board.

Annapolis OpenVPX FPGA cards include an on-board dual core 1.2 GHz PowerPC. This also has a connection to PCIe infrastructure (which includes FPGAs) and can be used by customers for application requirements. It is also used query board health like FPGA temperature and power. It is connected to the OpenVPX control plane via 1GbE.

There are also plenty of user backplane signals available on the Annapolis 6U Rear Transition Module (RTM) such as LVDS, FPGA HSS, IRIG, Ethernet and clocking. RTM HSS is also capable of 10Gbps signalling and supports multiple channels of 40GbE.

FEATURES

- **One or Two XILINX VIRTEX 7 FPGAS**
  - VX690T or VX980T
  - Up to 8 GB of DDR3 DRAM for 51.2 GB/s of DRAM bandwidth or up to 64 MB of QDRII+ SRAM for 32 GB/s of SRAM bandwidth
  - PCIe Gen3 8x from each FPGA to on-board PCIe switch
- **Backplane I/O**
  - 16x High Speed Serial IO lanes to VPX Data Plane (P1) for 20 GB/s of Full Duplex Bandwidth
  - 16x High Speed Serial FPGA connections to P5
  - 8x High Speed Serial IO lanes to P4
  - Two PCIe Gen3 8x Connections to VPX Expansion Plane (P2)
  - 24 LVDS and 8 Single Ended lines to P3
  - Backplane Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and user designed protocols
- **Front Panel I/O**
  - Accepts Standard Annapolis WILDSTAR Mezzanine Cards, including a wide variety of WILDSTAR ADC and DAC Mezzanine Cards
  - Three or six optional built-in Front Panel QSFP+ Transceivers running at up to 52.4 Gbps each for 39 GB/s of Full Duplex Bandwidth
  - 1 Gb Ethernet RJ45 connector for Remote Host Access
  - External clock and IRIG-B Support via Front Panel SMA
  - QSFP+ Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and user-designed protocols
- **Dual Core Processor APM86290**
  - Host Software: Linux API and Device Drivers
  - Each core runs up to 1.2 GHz
  - 2 GB of DDR3 DRAM
  - 4 GB SATA SSD and 16MB NOR Boot Flash
  - 4x PCIe Gen2 connection to on-board PCIe Switch
- **Application Development**
  - Full CoreFire Next™ Board Support Package for Fast and Easy Application Development
  - 10/40Gb Ethernet and AnnapMicro Protocol Cores Included
  - Open VHDL Model including Source Code for Hardware Interfaces
  - Open VHDL IP Package for Communication Interfaces
  - Chipscope Access through RTM
- **System Management**
  - System Management using Intelligent Platform Management Interface (IPMI)
  - Diagnostic monitoring and configuration
  - Current, Voltage and Temperature Monitoring Sensors
  - Hot Swappable (exclusive to WILDSTAR OpenVPX EcoSystem)
- **Mechanical and Environmental**
  - 6U OpenVPX (VITA 65) Compliant, 1" VITA 48.1 spacing
  - Supports OpenVPX payload profile: MOD6-PAY-4F1Q2U2T-12.2.1-n
  - Integrated Heat Sink and Board Stiffener
  - Available in Extended Temperature Grades
  - Air Cooled with Conduction Cooled path
  - RTM available for additional I/O

signal-processing.mil-embedded.com/p372789

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

WILDSTAR A5 for OpenVPX 6U

WILDSTAR A5 for OpenVPX 6U boards provide up to three Altera Stratix® V FPGAs per board with choice of GX parts up to 5SGXAB or GS parts up to 5SGSD8, up to 96 MB of QDRII+ SRAM for 59 GB/s of SRAM bandwidth and up to 8 GB of DDR3 DRAM for 51.2 GB/s of DRAM bandwidth. Up to 2.8 million logic elements and 4.3 million multiplier bits per board. Air Cooled Only.

These FPGA boards include 3 Altera Stratix V FPGAs with 48 High Speed Serial connections performing up to 14.1 Gbps. On each Compute Processing Element (CPE) FPGA there is six 72-bit QDRII+ SRAM interfaces clocked up to 550 MHz. The IO Processing Element (IOPE) FPGA has four 32-bit DDR3 DRAM ports clocked at up to 800 MHz.

With included High Speed Serial (HSS) FPGA cores (including 40GBASE-KR), there is up 20 GB/s of bandwidth on the VPX data plane which can go directly to other VPX cards or to a switch, depending on backplane topology. In addition, there is 16 GB/s of PCI Express Gen 3 bandwidth on the VPX Expansion Plane with an 8x Gen3 connection to each FPGA through a non-blocking PCIe switch. When using 40GBASE-KR, there is the added reliability of Forward Error Correction (FEC) to achieve a much lower Bit Error Rate (BER).

If IO is required, Annapolis offers extraordinary density, bandwidth and analog conversion choices. Each 6U card has 2 mezzanine IO sites which can support up to 4 WILDSTAR Mezzanine cards as well as a QSFP+ option (on WS7 and WS A5 board) that allows for 6 QSFP+ transceivers per slot. These options can be mix and matched to meet customer needs. Some configurations utilize a second slot (for example the QSFP+ option and WILDSTAR Mezzanine card used in a single IO Site).

WILDSTAR A5 and V7 FPGA boards are hot swappable allowing for more system reliability. This feature is unique to Annapolis and was developed because our experience with OpenVPX systems has shown it invaluable so a whole chassis does not need to be shutdown to remove a single board.

Annapolis OpenVPX FPGA cards include an on-board dual core 1.2 GHz PowerPC. This also has a connection to PCIe infrastructure (which includes FPGAs) and can be used by customers for application requirements. It is also used query board health like FPGA temperature and power. It is connected to the OpenVPX control plane via 1GbE.

There are also plenty of user backplane signals available on the Annapolis 6U Rear Transition Module (RTM) such as LVDS, FPGA HSS, IRIG, Ethernet and clocking. RTM HSS is also capable of 10Gbps signalling and supports multiple channels of 40GbE.

FEATURES

- **One, Two or Three ALTERA STRATIX® V FPGAS**
  - Up to three Altera Stratix® V FPGA Processing Elements: 5SGSD6, 5SGSD8, 5SGXA7, 5SGXA9, 5SGXAB
  - Up to 8 GB of DDR3 DRAM for 51.2 GB/s of DRAM bandwidth
  - Up to 96 MB of QDRII+ SRAM for 48 GB/s of SRAM bandwidth
  - PCIe Gen3 8x from each FPGA to on-board PCIe switch
- **Backplane I/O**
  - 16x High Speed Serial IO lanes to VPX Data Plane (P1) for 20 GB/s of Full Duplex Bandwidth
  - Up to 16x High Speed Serial FPGA connections to P5
  - 8x High Speed Serial IO lanes to P4
  - Two PCIe Gen3 8x Connections to VPX Expansion Plane (P2)
  - 32 LVDS and 8 Single Ended lines to P3
  - Backplane Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and user designed protocols
- **Front Panel I/O**
  - Accepts Standard Annapolis WILDSTAR Mezzanine Cards, including a wide variety of WILDSTAR ADC and DAC Mezzanine Cards
  - Three or six optional built-in Front Panel QSFP+ Transceivers running at up to 56.4 Gbps each for 42.3 GB/s of Full Duplex Bandwidth
  - 1 Gb Ethernet RJ45 connector for Remote Host Access
  - External clock and IRIG-B Support via Front Panel SMA
  - QSFP+ Protocol Agnostic connections support 10/40Gb Ethernet, SDR/DDR/QDR Infiniband, AnnapMicro protocol and user-designed protocols
- **Dual Core Processor APM86290**
  - Host Software: Linux API and Device Drivers
  - Each core runs up to 1.2 GHz
  - 2 GB of DDR3 DRAM
  - 4 GB SATA SSD and 16MB NOR Boot Flash
  - 4x PCIe Gen2 connection to on-board PCIe Switch
- **Application Development**
  - Full CoreFire Next™ Board Support Package for Fast and Easy Application Development
  - 10/40Gb Ethernet and AnnapMicro Protocol Cores Included
  - Open VHDL Model including Source Code for Hardware Interfaces
  - Open VHDL IP Package for Communication Interfaces
  - SignalTap Access through RTM
- **System Management**
  - System Management using Intelligent Platform Management Interface (IPMI)
  - Diagnostic monitoring and configuration
  - Current, Voltage and Temperature Monitoring Sensors
  - Hot Swappable (exclusive to WILDSTAR OpenVPX EcoSystem)
- **Mechanical and Environmental**
  - 6U OpenVPX (VITA 65) Compliant, 1" VITA 48.1 spacing
  - Supports OpenVPX payload profile: MOD6-PAY-4F1Q2U2T-12.2.1-n
  - Integrated Heat Sink and Board Stiffener
  - Available in Extended Temperature Grades
  - Air Cooled with Conduction Cooled path
  - RTM available for additional I/O

signal-processing.mil-embedded.com/p368675

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Hardware

CRITICAL LINK

MitySOM-5CSx: Altera Cyclone V SoC-based SOM

The MitySOM-5CSx combines the Altera Cyclone V SoC, memory subsystems and onboard power supplies into a highly-configurable, small form-factor System on Module (SOM). All products in the MitySOM-5CSx family are pin-for-pin compatible, allowing development teams room to grow and the flexibility to quickly and cost-effectively meet customers' ever-changing needs.

The MitySOM-5CSx family offers a wide range of processing densities, speed grades, and temperature options at competitive costs. Critical Link designed the SOM family with rigorous industrial, medical, and defense applications in mind, ensuring long-term production and professional support for our customers.

Standard SOMs and development kits are available today from numerous major distributors. If a standard variant does not meet your specification, contact Critical Link to discuss developing a custom solution.

MitySOM-5CSx features a Hard Processor System (HPS) providing up to 4,000 MIPS at speeds of up to 925MHz per core and is combined with a NEON coprocessor with double-precision FPU. The MitySOM-5CSx combines a Cyclone V with up to 2GB of DDR3 CPU/FPGA RAM with ECC, 512MB of dedicated DDR3 FPGA RAM (optional) and up to 48MB of QSPI NOR Flash creating a high-bandwidth system for embedded applications. The ARM architecture supports several high level operating systems, including Embedded Linux, Micrium uC/OS, Android, QNX, and Windows Embedded Compact.

By combining six 3.125Gbps transceivers, one PCIe hard core, up to 133 user I/O, and dual Gigabit Ethernet interfaces, the system can simultaneously acquire and efficiently process large amounts of data.

signal-processing.mil-embedded.com/p372770

Critical Link
www.criticallink.com/product/mitysom-5csx/

info@criticallink.com
www.linkedin.com/company/critical-link-llc
315-425-4045
twitter.com/Critical_Link

Hardware

INTERMAS

InterShell Enclosures

Intermas develops electronic enclosure systems:

Cabinets, housings, subracks, and an extensive range of accessories for the 19" rack systems and small form factors used in the fields of PCI, VME/VME64x, cPCI, IEEE, and communication applications with state-of-the-art EMI- and RFI-shielded protection.

Intermas has an extensive product range of more than 10,000 separate components and more than 30 years' experience.

Go to
www.Intermas-US.com
for our new catalog.

FEATURES

- InterShell is a new aluminum housing enclosure composed of a top and bottom, two front panels, and four screws.
- The simple housing design offers an uncomplicated solution for small form factors for easy and quick assembly.
- Color options are unlimited and customer-specific print is possible.
- InterShell is used for the packaging of small electronic units such as Eurocard formats with 100x160 mm, universal formats, or as mITX formats for example.
- Excellent EMC compliancy.
- InterShell is available in the following standard dimensions (h/w/d) as well as customized formats:
  - 40x106.6x168.6 mm
  - 60x150x120 mm
  - 50x190x190 mm

signal-processing.mil-embedded.com/p372684

Intermas US LLC
www.Intermas-US.com

intermas@intermas-us.com
800-811-0236

PENTEK

Model 71624

The Model 71624 dual channel, 34 signal, adaptive IF (Intermediate Frequency) relay XMC with a Virtex-6 FPGA. As an IF relay, it accepts two IF analog input channels, modifies up to 34 signals, and then delivers them to two analog IF outputs. Any signal within each IF band can be independently enabled or disabled, and changed in both frequency and amplitude as it passes through the module.

The Model 71624 supports many useful functions for both commercial and military communications systems including signal drop/add/replace, frequency shifting and hopping, amplitude equalization, and bandwidth consolidation. Targeted applications include signal monitoring, signal jamming, channel security, countermeasures, beamforming and radar applications.

The Model 71624 features 34 digital down converters (DDCs), each independently tunable across a 100 MHz input IF range, handling signal bandwidths from 20 kHz up to 312 kHz. The DDCs deliver 34 baseband signals to the host computer, which determines how each signal is dropped, replaced, or changed in amplitude and frequency. The modified signals are then combined and delivered as an analog IF output.

"The flexibility of the Model 71624 allows each signal to be manipulated in its own way," said Rodger Hosking, vice-president of Pentek. "Because the host computer is in the control loop, it supports adaptive applications such as signal scanning and tracking, as well as cognitive radio functions like jamming, denial of service and other countermeasures," he added.

Adaptive Relay

The Model 71624 comes preconfigured with a suite of time-tested Pentek IP functions for DDC, DUC, DMA transfers and digital summation. Captured signals may be relayed with gain and frequency modifications directly through the module, or the host system may substitute a modified signal on any channel instead. The Model 71624 features an onboard programmable sample clock synthesizer that is phase-lockable to an external frequency reference input. Multiple units can be synchronized for high-channel count systems.

Form Factors, Development Systems & Software Support

The Cobalt Model 71624 XMC is designed for commercial, rugged or conduction cooled operating environments. It is also available in several form factors, including 3U and 6U VPX (52624/53624 & 57624/58624), 3U and 6U cPCI (72624/73624/74624), AMC (56624) and PCIe (78624).

With a Pentek SPARK® Model 8266 PC development system or Model 8267 VPX development system, designers can start working on their application immediately. Created to save engineers and system integrators the time and expense associated with building and testing a system, it ensures optimum performance of Pentek boards and working examples.

The Pentek ReadyFlow® Board Support Package is available for Windows and Linux operating systems. ReadyFlow is provided as a C-callable library, the complete suite of initialization, control and status functions, as well as a rich set of precompiled, ready-to-run-examples, accelerate application development.

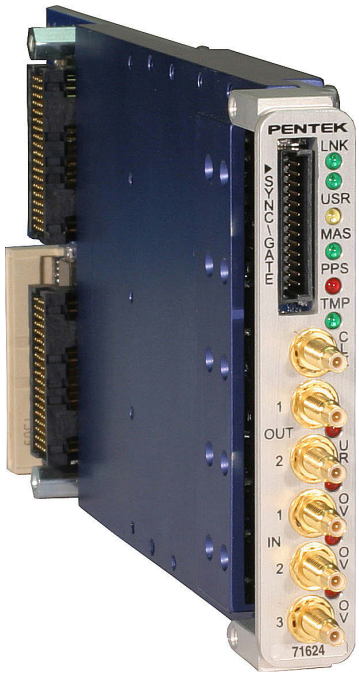

FEATURES

- Modifies 34 IF signals between input and output
- Up to 80 MHz IF bandwidth
- Two 200 MHz 16-bit A/Ds
- Two 800 MHz 16-bit D/As
- 34 DDCs and 34 DUCs (digital downconverters and digital upconverters)
- Signal drop/add/replace
- Frequency shifting and hopping
- Amplitude boost and attenutation
- PCI Express Gen 1: x4 or x8
- Synchronize up to eight modules with Model 7893 System Synchronization and Distribution Amplifier – PCIe
- Model 8266 SPARK Development System for Flexor (FMC), Onyx (Xilinx Virtex-7) and Cobalt (Xilinx Virtex-6) PCIe Boards
- Available in the following form factors:
  - Model 52624 – VPX 3U
  - Model 53624 – VPX 3U
  - Model 56624 – AMC
  - Model 57624 – VPX 6U
  - Model 58624 – VPX 6U
  - Model 72624 – cPCI 6U
  - Model 73624 – cPCI 3U
  - Model 74624 – cPCI 6U
  - Model 78624 – PCIe

signal-processing.mil-embedded.com/p372794

Pentek
http://www.pentek.com/go/sigprocrg71624

sales@pentek.com
www.linkedin.com/company/pentek
201-818-5900
www.twitter.com/pentekinc

Hardware

NH30 29-Gigabit Ethernet Switch

Designed for use in trains, mass transit or industrial environments, MEN Micro's robust NH30 29-Gigabit Ethernet switch offers carrier grade Layer 2 and Layer 3 VLAN routing capabilities.

The pre-configured CompactPCI Serial NH30, which comes in a half 19-inch format, is an additional member of the modular, built-to-order system family from MEN Micro. Redundant power supplies, PoE+ support and self-monitoring make the switch a reliable, flexible long-term solution.

The NH30 supports a total of 25 ports at the front. The system can be configured with up to six line cards as either active or passive copper connections or as SFP interfaces. In addition, the ports have PoE+ to enable the connection of different devices, such as a camera with a maximum of 60 W. Pre-installed carrier-grade, managed firmware from Vitesse supports almost all popular protocols on the market, including standard ring and protection switching/ring protocols.

The reliable NH30 Ethernet switch has self-control and self-monitoring capabilities. It manages lifetime information about critical operating components, such as the shelf fans and power supplies, and can be used for manual shutdown. For security and monitoring purposes, the shelf controller data is available via network interfaces and CLI.

Two AC or DC power supplies ensure reliability and redundancy. When used with one of the power supplies, an uninterruptible power supply (UPS) can protect the system against power fluctuations.

FEATURES

- Compact 25-port, 40 HP managed, rugged Layer 2/3 Gigabit Ethernet switch
- 29 Gbit/s carrier grade TCAM switching matrix and firmware
- Fanless operation or forced-air cooling (depending on configuration)
- Single or redundant power supplies or uninterruptible power supply (AC or DC)
- Support for up to six line cards with fiber, RJ45 or M12 connectors
- PoE+/non-PoE power sourcing Ethernet ports in mixed configuration
- One "alive" relay output port
- -40 to +85°C with qualified components
- EN 50155 class TX compliant (railways)

signal-processing.mil-embedded.com/p372792

MEN Micro Inc.
www.menmicro.com

Stephen.Cunha@menmicro.com
215-542-9575
www.linkedin.com/company/men-micro-inc-
twitter.com/MENMicro

Standards and Systems

Signal Processing Solutions

VadaTech provides Signal Processing Solutions that redefine performance density with superior SWaP-C. From RADAR data acquisition & signal processing, FPGA/data processing, video/image processing, & more, VadaTech has a solution for you.

Our standard modular building blocks include high-performance multi-core Intel, Tilera, and Freescale processors; Altera, Xilinx, and TI FPGAs; ATI & other graphics processors; & powerful ADCs and DACs.

AS9100 certified, VadaTech's state-of-the-art production facility ensures the highest product quality and consistency.

FEATURES

- Modular Open Standard Architectures (MOSA)
- Approx 1/2 the SWaP-C of competing products
- Rugged 1U Signal Processing platforms with time stamping options
- SFF & Cube modular platforms
- Virtex-7, Kintex-7, and Zynq FPGAs
- High GSPS ADCs & DACs
- Designed & manufactured in the USA
- Software enhancement services

signal-processing.mil-embedded.com/p372793

VadaTech
www.vadatech.com

sales@vadatech.com
702-896-3337

Annapolis Micro Systems

Wild40 12-Slot OpenVPX 6U Chassis

11U Rack mountable, 12-slot OpenVPX chassis with OpenVPX switched topology backplane capable of 10Gbps+ signalling compromising of 2 switch and 10 payload 1" slots. Option of additional secondary 4-slot OpenVPX power-only (Shown) or 5-slot VME/VXS backplane.

The Wild40 12-Slot OpenVPX 6U Chassis is an OpenVPX-compatible (VITA 65) chassis capable of accepting up to ten 6U tall by 160mm OpenVPX Payload Front Plug-in Modules (FPMs) and two 6U tall by 160mm OpenVPX Switch FPMs and up to twelve 6U tall by 80mm Rear Transition Modules (RTMs) in its Primary Backplane. Plug-in Module slot spacing is 1 VITA 48.1.

The Wild40 12-Slot OpenVPX 6U Chassis' Primary Backplane is a very high performance backplane which is capable of Serial I/O signaling at rates up to 10Gbps on the Data Plane and up to 8Gbps on the Expansion Plane. The Data Plane of the backplane is arranged in a dual-star configuration with two Fat Pipe connections from each Switch Slot to each Payload Slot. The Expansion Plane is a chain connecting adjacent Payload Slots.

In addition to the Primary Backplane there is also an option for a Secondary 4-Slot VPX Power-Only or 5-slot VME/VXS Backplane. The 4-slot VPX backplane supports four OpenVPX VITA65 slots with a 1 VITA 48.1 slot spacing. These slots are not connected to each other on the Data or Expansion Planes, instead all of their connections go straight through the backplane to the RTM backplane connectors. These slots are ideally suited for Clock Distribution boards, Tuners or other non-IO intensive FPMs.

The chassis includes a Chassis Monitoring system which displays DC voltages, slot temperatures and fan Revolutions Per Minute (RPMs) on the front panel of the chassis and can be used to set fan speed. The Chassis Monitor can be accessed and controlled remotely via the Serial or Ethernet interfaces.

The card cage is recessed from the front of the chassis so that cabling can be used between Plug-in Modules and be contained within the frame of the chassis.

FEATURES

- 10U High with Front Mounted OpenVPX Card Cage
- Primary 12 Slot OpenVPX High Speed Switched Backplane with RTM Support
- Optional Secondary 5 Slot VME/VXS or 4 slot VPX Backplane for Power Only Payload Cards
- Up to 3200 Watt Power Supply
- Backplane Profile: BKP6-CEN12-11.2.X
- Payload Profile: SLT6-PAY-4F1Q2U2T-10.2.1
- Switch Profile: SLT6-SWH-16U20F-10.4.2

Annapolis is famous for the high quality of our products and for our unparalleled dedication to ensuring that the customer's applications succeed.

We offer training and exceptional special application development support, as well as more conventional support.

signal-processing.mil-embedded.com/p372673

Annapolis Micro Systems, Inc.
www.annapmicro.com

wfinfo@annapmicro.com
410-841-2514

Innovative Integration
a subsidiary of Interconnect Systems, Inc.
... real time solutions!

V616 Digital Transceiver

The **V616 digital transceiver** supports one or two plug-in modules, each featuring four independent channels of DDC, two DUC, and one spectrum analyzer embedded in the Xilinx Virtex-6 FPGA. It allows users to monitor the wide-band or narrow band spectrum and record the data directly from the ADCs or down-convert the channels modulated on the IF band. The embedded PC can do contiguous recording at 2,000 MByte until running out of disk space. The transmitter can play the recorded baseband waveform of different bandwidth, up-convert and modulate it on the IF band on the DACs.

Each DDC/DUC has its own programmable tuner, programmable low-pass filtering, gain control, and decimation/interpolation setting, supporting independent channel bandwidths up-to 100 MHz. The DDC data is packetized in VITA-49 format with the accurate timestamps, synchronous to the external PPS signal. The DDC channel can be enabled and disabled on the fly to save the storage and bandwidth to the host computer. The embedded power meter monitors the power (dBFS) of the ADC inputs, allowing users to perform possibly analog gain control in the external front-end device, not included in the system.

The spectrum analyzer, including windowing, calculates the wide- band spectrum of the ADC data or the narrow-band spectrum of the DDC output data at the programmable update rate. The maximum hold helps to retain the information in the spectrum and the programmable threshold monitoring spectrum detects the spectral activities up-to 512 bins.

A development kit is available for the custom design. Users can insert custom-made cores to create even more advanced applications.

Download Pricing & Data Sheets Now!

FEATURES

- Intel i7 Quad Core, 8 GB RAM, 240 GB SSD, Win 7 Pro 64-bit
- Two, independent XMC module sites
- Sustained logging rate up-to 2,000 MByte/s

Per-Module Features

- Two 12-bit, 1000 MHz ADCs; analog bandwidth: 1000 MHz (AC Coupled)
- Two 16-bit, 1000 MHz DACs; analog bandwidth: 1000 MHz (AC Coupled)
- Xilinx Virtex-6 SX475T-2 FPGA
- Embedded power meter
- PCI Express Gen 2 (3,200 MByte/s)

Digital Down-Converter (DDC)

- Four independent 16-bit DDC channels
- Programmable tuner: 1 – 1000 MHz; resolution 0.2328 Hz
- Programmable bandwidth: 10 KHz – 100 MHz
- DDCoutputsSNR>64dB; SFDR>80dB
- Spectrum inversion for ADC under-sampling
- Support synchronous down-sampling on multiple channels and modules using external clock/trigger

Digital Up-Converter (DUC)

- Two independent 16-bit DUC channels
- Programmable tuner: 1 – 1000 MHz; resolution 0.2328 Hz
- Programmable bandwidth: 10 KHz – 100 MHz
- Spectrum inversion

Spectrum Analyzer

- Single wide-band/narrow-band 32K points FFT
- Eight Windowing available
- Programmable FFT update rate
- Programmable maximum hold mode
- Threshold limited spectrum monitoring up to 512 bins

Applications include Digital Receiver, Surveillance and Arbitrary Waveform Generator/Player.

signal-processing.mil-embedded.com/p372705

Innovative Integration
www.innovative-dsp.com/products.php?product=Digital%20Transceiver%20V616

sales@innovative-dsp.com
805-383-8994

Annapolis Micro Systems

FPGA Application Development Team

Annapolis Micro Systems employs a select team of experienced FPGA designers dedicated to helping you solve your real-world processing challenges while minimizing your Time to Market. Our patented CoreFire and Corefire Next Design Suites provide high-speed solutions that get the application developed in record time, drastically reducing program cost and risk.

Annapolis Application Engineers are here to help you in the following ways:

Architect A Custom System Using Our Powerful COTS Products

Annapolis Application Engineers will help you create customized system architectures and determine appropriate hardware configurations using our COTS products. This service is provided for free in order to accelerate your program's progress and guarantee the best and most efficient system architecture available.

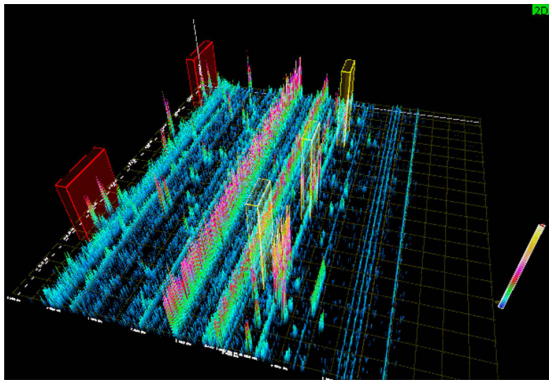

Our process begins with a discussion of feasibility with a member of our sales team, and then moves to a roundtable conversation with our lead engineers. We can collaborate with your staff to determine your project's ideal system architecture and establish its hardware requirements.

Some customers choose to task their own in-house engineers on the design phase. If that option is selected, we can provide training for both our CoreFire and VHDL Application Development paths.

Learn more about our training classes.

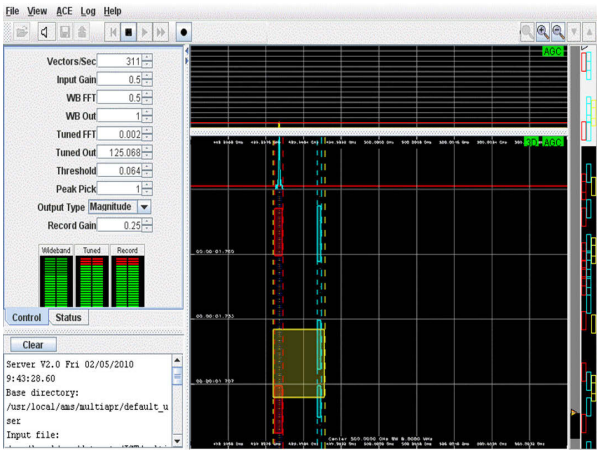

Develop FPGA Application and Integrate into a Turn-Key system

Annapolis can also help you with your FPGA Application Development. In these instances, we will work with your team to document the project's specifications, including its functional needs, schedule, and First Article Test (FAT) requirements.

Depending on the customer's desire the application may involve passing data around the key data paths and the customer will implement the processing IP later, while in other cases Annapolis will implement the whole application including the processing and GUI as a Turn-Key system.

The effort required to develop the application is then estimated and you are provided with a quote. If agreeable, we contract the project and our Applications Engineers develop the application with the CoreFire (Next) Design Suite, perform the FAT, and deliver the application. In addition we can provide integration support at the customer site when needed.

signal-processing.mil-embedded.com/p372789

Annapolis Micro Systems, Inc.
www.annapmicro.com/services-support/application-development-team/
wfinfo@annapmicro.com
410-841-2514

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR OpenVPX Ecosystem

FPGA Processing Boards
1 to 3
Altera Stratix V or
Xilinx Virtex 6 or 7
FPGAs per Slot

Input/Output Modules Include:
Quad 130 MSps thru Quad 550 MSps A/D
1.5 GSps thru 5.0 GSps A/D
Quad 600 MSps D/A
Dual 1.5 GSps thru 4.0 GSps D/A

1 to 40 Gbit Ethernet
SDR to FDR Infiniband

Open VPX Storage
Up to 8 TBytes Per Slot
4 - 8 GBytes Per Second

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

Open VPX Switch
1 to 40 Gbit Ethernet
SDR to FDR Infiniband

Chassis
4, 6 or 12 Slot
Up to 14G

Annapolis
Micro Systems

High Performance Signal and Data Processing in Scalable COTS FPGA Computing Fabric

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