

Xcell journal

ISSUE 94, FIRST QUARTER 2016

SOLUTIONS FOR A PROGRAMMABLE WORLD

Xilinx Extends Ecosystem to Reshape the Future of Embedded Vision, IIoT System Design

Intelligent Gateways Make a Factory Smarter

How to Extend the Operating Temperature of FPGAs

FMC+ Standard Propels Embedded Design to New Levels

The latest Xilinx tool updates and patches, as of March 2016




Drone Platform Soars with the Zynq SoC

34

 **XILINX**
ALL PROGRAMMABLE™

www.xilinx.com/xcell



Lifecycle



Technology

Software-Defined Radio from Concept to Production


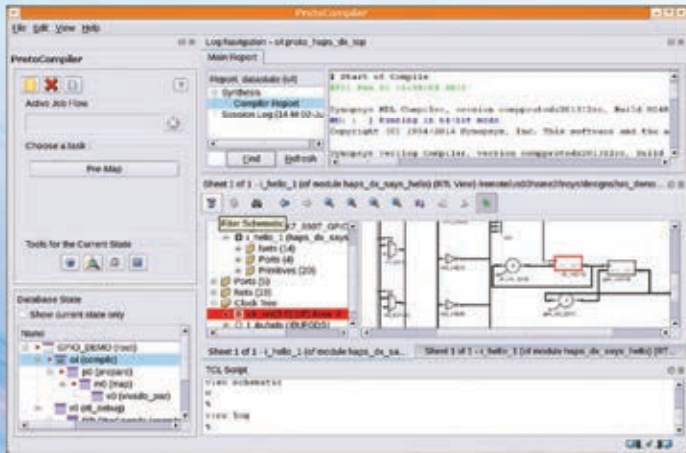
Product Benefits/Features:

- Combines the Analog Devices AD9361 integrated RF Agile Transceiver™ with the Xilinx Z7035 Zynq®-7000 All Programmable SoC
- Running Linux on the dual core ARM A9, the PicoZed SDR provides a full software environment which can be used from prototyping to production
- Userspace I/O (UIO), Industrial I/O (IIO) subsystems and drivers
- Supports Avnet camera modules with V4L2 drivers for realtime video capture
- Frequency-agile 2x2 receive and transmit paths with independent LO from 70 MHz to 6.0 GHz
- Handheld form-factor, ideal for a broad range of fixed and mobile SDR applications
- Full support in MATLAB and Simulink



For more information, visit picozed.org/sdr

HAPS-80 FPGA-Based Prototyping Solution Delivers Up to 100 MHz System Performance



- ▶ New automated high-speed pin-multiplexing delivers the highest system performance
- ▶ Integrated ProtoCompiler software automates partitioning to reduce time to first prototype to less than two weeks
- ▶ Built-in debug enables the capture of thousands of RTL signals
- ▶ Scalable architecture supports up to 1.6 billion ASIC gates based on the Xilinx Virtex UltraScale FPGA

To learn more visit: www.synopsys.com/haps-80

Xcell journal

PUBLISHER Mike Santarini
mike.santarini@xilinx.com
408-626-5981

ACTING EDITOR IN CHIEF Steve Leibson
steve.leibson@xilinx.com
408-879-2716

EDITOR Jacqueline Damian

ART DIRECTOR Scott Blair

DESIGN/PRODUCTION Teie, Gelwicks & Associates
1-800-493-5551

ADVERTISING SALES Judy Gelwicks
1-800-493-5551
xcelladsales@aol.com

INTERNATIONAL Melissa Zhang, Asia Pacific
melissa.zhang@xilinx.com

Christelle Moraga, Europe/
Middle East/Africa
christelle.moraga@xilinx.com

Tomoko Suto, Japan
tomoko@xilinx.com

REPRINT ORDERS 1-800-493-5551



www.xilinx.com/xcell/

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3400
Phone: 408-559-7778
FAX: 408-879-4780
www.xilinx.com/xcell/

© 2016 Xilinx, Inc. All rights reserved. XILINX, the Xilinx Logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

The articles, information, and other materials included in this issue are provided solely for the convenience of our readers. Xilinx makes no warranties, express, implied, statutory, or otherwise, and accepts no liability with respect to any such articles, information, or other materials or their use, and any use thereof is solely at the risk of the user. Any person or entity using such information in any way releases and waives any claim it might have against Xilinx for any loss, damage, or expense caused thereby.

The Expanding Xilinx Ecosystem

This issue of *Xcell Journal* brings you news of a radical expansion in the ecosystem of companies supporting systems development based on Xilinx® Zynq®-7000 SoCs and Zynq UltraScale+™ MPSoCs. Xilinx announced this expansion at the end of February at the Embedded World conference, held in Nuremberg, Germany. The Xilinx booth at the show offered demos of many applications created by Xilinx as well as by Xilinx ecosystem member companies and Xilinx customers, including:

- 4K Video Processing with Zynq UltraScale+ MPSoC;
- ADAS Development Assistant, presented by Xylon;
- Eyescan Test with Real-time Analytics, presented by EyeTech Digital Systems;
- Intelligent Gateway for IIoT;
- Machine Vision and Control, presented by National Instruments;
- OpenCV Hardware Acceleration for Machine Vision;
- Image Classification Training Software, presented by Xylon and eVS embedded Visions Systems;
- Real-Time Multi-Object Detection, presented by Xylon and eVS;
- Real-Time 3D Multi-Object Recognition for Smart Cameras, presented by iVeia;
- Zynq UltraScale+ MPSoC Powered Robotics; and
- Software Acceleration via Programmable Logic.


Chances are good that your next design will resemble one or more of these demonstrations, which means that the Xilinx and ecosystem IP used in the demos could easily apply to your design too. This issue's cover story, "Xilinx Extends Ecosystem to Reshape the Future of Embedded Vision, IIoT System Design," by Aaron Behman and Dan Isaacs, starting on page 8, gives you a lot more detail about this important ecosystem expansion.



Kudos and Goodbye to Mike Santarini

Xcell Journal and Xilinx owe a great debt to Mike Santarini, who served as the publisher of this magazine for the past eight years. Mike has moved on to other things, but the contribution he made was significant. Mike, all of us here at Xilinx say "Thank you!" and wish you good luck in your future endeavors.

Accelerate Synthesis Runtime by 3X with Synplify Premier



Achieve fast runtimes with:

- ▶ Incremental and fast synthesis flows
- ▶ Distributed synthesis technology
- ▶ Fast compile with multiprocessing using single or multiple machines

To learn more about how Synopsys FPGA design tools accelerate synthesis runtimes, visit: www.synopsys.com/fpga

VIEWPOINTS

Letter from the Publisher

The Expanding Xilinx Ecosystem... **4**



XCELLENCE BY DESIGN APPLICATION FEATURES

Xcellence in Industrial

Intelligent Gateways
Make a Factory Smarter... **14**

Xcellence in Communications

Evaluating an IQ Compression
Algorithm Using Vivado
High-Level Synthesis... **22**

Xcellence in Wireless Communications

Xilinx FPGA Enables Scalable
MIMO Precoding Core... **28**

Xcellence in Aeronautics

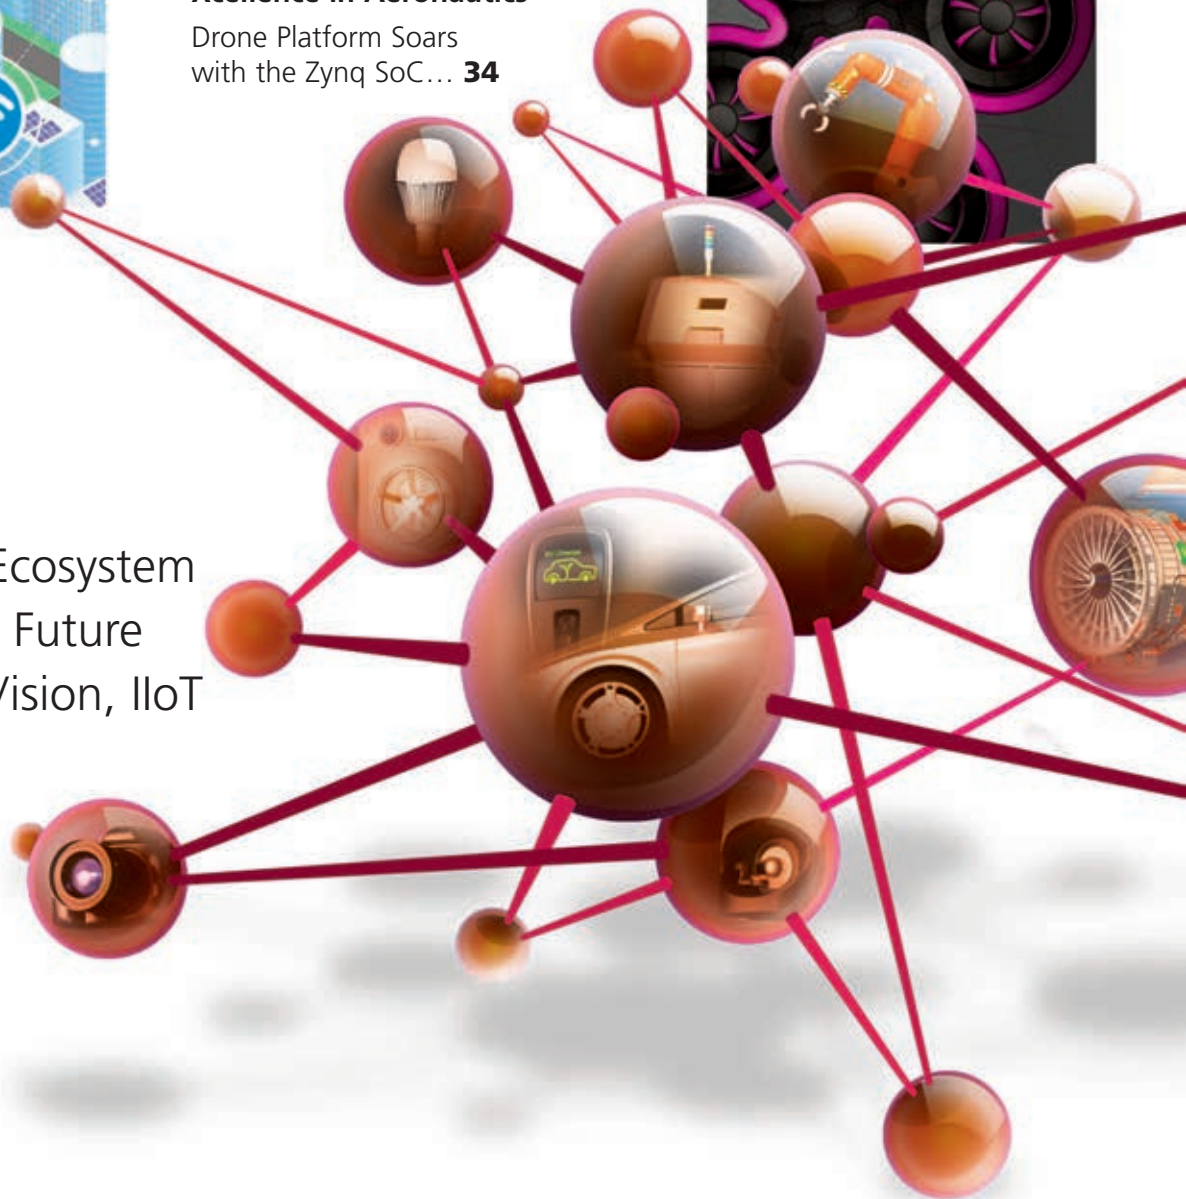
Drone Platform Soars
with the Zynq SoC... **34**



Cover Story

Xilinx Extends Ecosystem
to Reshape the Future
of Embedded Vision, IIoT
System Design

8



THE XILINX XPERIENCE FEATURES

Xperts Corner



How to Extend the Operating Temperature of FPGAs... **40**

Xplanation: FPGA 101

How to Digitize Hundreds of Signals with a Single Xilinx FPGA... **46**

Tools of Xcellence

FMC+ Standard Propels Embedded Design to New Levels... **52**



XTRA READING

Xtra, Xtra The latest Xilinx tool updates and patches, as of March 2016... **58**

Xpedite The latest and greatest from the Xilinx Alliance Program Partners... **60**

Xclamations! Share your wit and wisdom by supplying a caption for our wild and wacky artwork... **62**

Xilinx Extends Ecosystem to Reshape the Future of Embedded Vision, IIoT System Design

by Aaron Behman

Director, Video & Vision,
Corporate Strategy & Marketing
Xilinx, Inc.
aaronb@xilinx.com

Dan Isaacs

Director, Industrial IoT,
Corporate Strategy & Marketing
Xilinx, Inc.
dani@xilinx.com

Broad IP, software, hardware and design services offerings enable smarter, connected, highly differentiated systems based on Xilinx All Programmable devices.

S

Systems with unprecedented levels of software-based intelligence, optimized hardware and any-to-any connectivity are shaping the future of embedded vision and the Industrial Internet of Things (IIoT). At Embedded World in Nuremberg, Germany, in February, Xilinx® announced it had strengthened and expanded the ecosystem that supports the development of IIoT and embedded vision systems based on Xilinx All Programmable devices. Xilinx and its ecosystem partners demonstrated offerings at the show that make it easier for users to develop all manner of smarter, connected and highly differentiated systems.

The ecosystem announcement at Embedded World completes three milestones Xilinx has met in the past year. On March 9, 2015, Xilinx announced its SDSoc™ development environment, which permits the community of designers writing C/C++ programs to work with Xilinx Zynq®-7000 SoCs. The environment targets algorithm developers who are not accustomed to getting under the hood and changing hardware using Verilog or VHDL, but who can benefit from the superior performance and performance/watt of Xilinx devices. On Sept. 30, 2015, Xilinx announced that it had begun shipping the Zynq UltraScale+™ MPSoC. And on Feb. 16, 2016, Xilinx announced the strengthening and expansion of the ecosystem supporting Zynq-based designs in the embedded vision and IIoT arenas.

The exciting applications emerging in industrial/embedded vision and IIoT cut across the industrial, scientific, medical, pro A/V, consumer, aerospace and defense, and automotive market segments.

The key barrier to using the superior performance and performance/watt characteristics of Xilinx All Programmable devices has been the programming model.

With its ecosystem expansion, Xilinx is making its All Programmable devices just as easy to use as CPUs and GPUs, but with superior performance/watt.

C/C++ users are more accustomed to writing code for CPUs and, more recently, GPUs. With Xilinx's Vivado® High-Level Synthesis (HLS) for software-defined hardware and SDx™ environment for software-defined systems development, many more system developers can make use of the software-defined All Programmable devices in the Xilinx Zynq-7000 SoC and Zynq UltraScale+ MPSoC families. With its ecosystem expansion, Xilinx is making its offerings just as easy to use as CPUs and GPUs, but with superior performance and performance/watt.

The pipelines for embedded vision and IIoT systems have much in common. Both start with sensing and data acquisition. For embedded vision systems, that data takes the form of a series of images or a video stream. Sensed data for IIoT systems includes

video but also encompasses a long list of additional sensed parameters, including acceleration and vibration; acoustic/ultrasonic; chemical and gas; electric/magnetic; flow; force, load, torque and strain; humidity and moisture; leak and level; machine vision; optical; motion, velocity and displacement; position, presence and proximity; pressure; and temperature.

RISING NEED FOR SENSOR FUSION

Several embedded vision and IIoT systems require sensor fusion, or the processing and merging of data from multiple and different types of sensors into actionable intelligence. For embedded vision systems, multiple video streams may be combined to produce more-usable or more-useful video streams. For example, vehicle-based vision

systems often combine video streams from four, five, six or more video cameras to produce a single bird's-eye view that gives the driver 360° 2D planar or 3D spherical vision. Vision systems drive local displays but also send locally processed video to the cloud for further processing, for combination with other video streams and for storage.

IIoT systems may combine video with additional sensed data to define needed actions. For example, the new CPPS-Gate40 Smart Gateway from System-on-Chip engineering (SoC-e; see article, [page 14](#)) incorporates a variety of I/O ports commonly used in industrial control systems, combined with local, high-speed data processing, and places the resulting data on a dual-redundant optical Ethernet ring using High-Availability Seamless Redundancy/Parallel Redundancy Protocol (HSR/

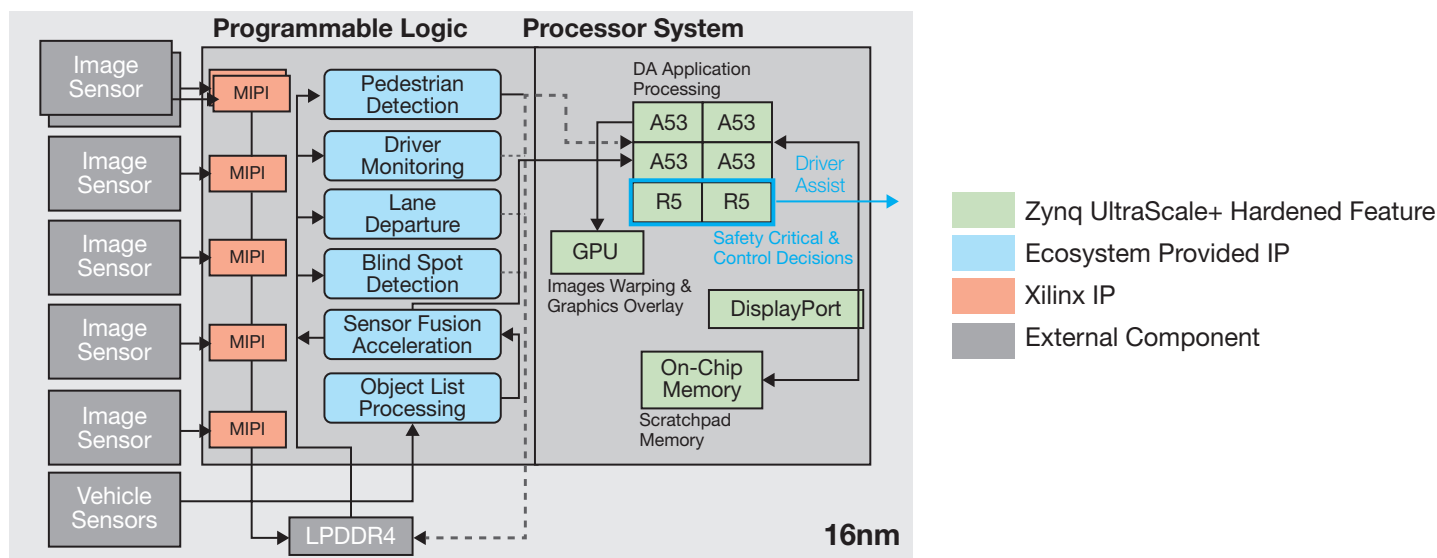


Figure 1 — This ADAS design leverages the heterogeneous processing capabilities of the ARM Cortex cores in the Zynq UltraScale+ MPSoC.

PRP). A defining characteristic of IIoT systems is the ability to use sensed data for high-speed, real-time control not possible if relying on cloud-based processing and decision-making.

Of course, there are many alternative ways to design such systems using a CPU or GPU, but Xilinx Zynq-7000 SoCs and Zynq UltraScale+ MPSoCs offer several significant advantages and benefits when you're designing differentiated systems:

1. Highest performance/watt. Xilinx All Programmable devices combine hardware, software and I/O programmability, letting you collapse a two-, three- or four-chip design into one chip to maximize system performance while lowering power consumption.

2. Sensor fusion. Xilinx All Programmable devices offer unique abilities to ingest and process multiple types of information, ranging from low-bit-rate data, such as temperature and pressure, to high-bit-rate data, including multiple, simultaneous high-definition or super-high-definition video streams.

3. Any-to-any connectivity. The programmable I/O capabilities of Xilinx Zynq-7000 SoCs and Zynq UltraScale+ MPSoCs offer an unmatched ability to adapt to nearly any conceivable sensor I/O requirement, from multiple video interface standards (such as MIPI and HDMI) to intelligent sensor interfaces (such as I²C and SPI) and high-speed A/D converters (including JESD204B and LVDS).

4. Multilevel security and multilayer safety. The Zynq UltraScale+ MPSoC's quad-core ARM® Cortex™-A53 application processor and dual-core ARM Cortex-R5 real-time processor with hardware security features offer unique abilities to implement security and functional-safety protocols.

5. “Chameleon” All Programmable platforms. The hardware and software processing and I/O flexibility of Zynq-7000 SoCs and Zynq Ultra-

Scale+ MPSoCs allow you to create reusable, software-defined hardware platforms with configurable and extensible range both up and down the end-product family cost curve, from low-cost to high-performance systems, and to extend a brand into new markets across multifunction product lines. This is not a hypothetical advantage; many Xilinx customers are already putting it into practice.

Here are four examples of Chameleon All Programmable platforms, all leveraging the Xilinx Zynq UltraScale+ MPSoC to target distinct markets.

EXAMPLE 1: ADVANCED DRIVER ASSISTANCE SYSTEM

An advanced driver assistance system (ADAS) combines video data from several video cameras and additional vehicle sensor data, including inertial navigation and even GPS map data, to make decisions about braking, steering and driver alerts. The block diagram in Figure 1 shows a typical ADAS design based on a Zynq UltraScale+ MPSoC.

As Figure 1 shows, this design leverages the heterogeneous processing capabilities of the quad-core ARM Cortex-A53 application processor and dual-core ARM Cortex-R5 real-time processor in the Xilinx Zynq UltraScale+ MPSoC. The five red boxes in the diagram depict MIPI video-interface IP available directly from Xilinx. The six blue boxes show high-speed IP processing blocks provided by other companies in the Xilinx ecosystem that implement high-level functions, including pedestrian detection, driver monitoring, lane-departure monitoring, blind-spot detection and sensor fusion.

The depicted ADAS system takes full advantage of the Zynq UltraScale+ MPSoC's any-to-any connectivity to communicate with any sensor interface, including MIPI for the video cameras. Nonprogrammable devices from competing vendors cannot easily adapt to new sensor interfaces without adding I/O chips to handle the additional I/O

interfaces and protocols. In addition, the Zynq UltraScale+ MPSoC's superior hardware-based video-processing performance allows it to handle more video channels than competing standard devices. Unlike such devices, the Zynq UltraScale+ MPSoC handles a programmable number of video streams.

Because of the Zynq UltraScale+ MPSoC's I/O flexibility and processing power, very little hardware is needed outside of the MPSoC itself, except for the sensors and external memory. The performance/watt metric for this system is approximately 3x better than for a comparable system using CPU-based silicon from a leading competitor.

EXAMPLE 2: 4K VIDEO SURVEILLANCE

A Zynq UltraScale+ MPSoC connects to multiple sensors, including different types of video cameras, in the 4K multichannel, multisensor video surveillance system shown in Figure 2. The red boxes in the block diagram again depict Xilinx interface IP for MIPI-interfaced video cameras and displays and for different I/O interfaces that connect other sensor types. The six all-blue boxes depict processing IP available from Xilinx ecosystem companies. The two red/blue boxes indicate IP blocks available both from Xilinx and from companies in its expanded ecosystem.

The performance/watt metric for this Chameleon All Programmable system is approximately 5x better than for a comparable system designed with CPU/DSP/GPU-based silicon from a leading competitor. The safety and security features of the Zynq UltraScale+ MPSoC, including ARM TrustZone® capabilities and the device's hardware-based AES encryption, are especially useful in security applications like this one.

EXAMPLE 3: SMART-GRID SUBSTATION AUTOMATION

Our third example, a substation automation system targeting smart-grid design,

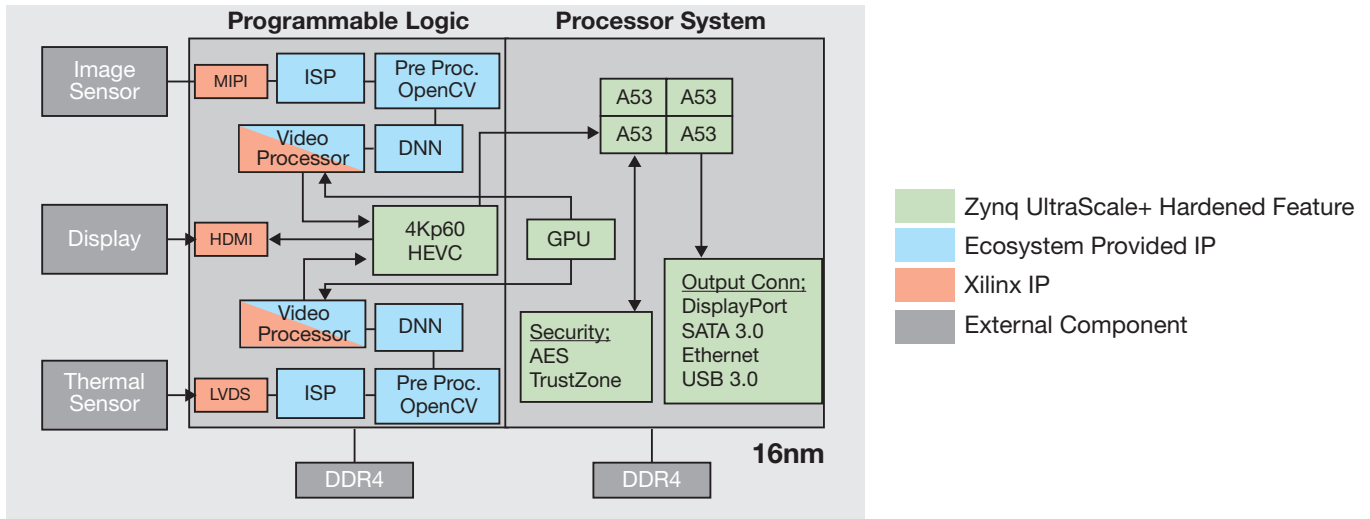


Figure 2 — This 4K multichannel/multisensor video surveillance system taps the safety and security capabilities of the Zynq UltraScale+ MPSoC.

is an IIoT application that deals with multiple Ethernet streams from various sensing components that monitor substation parametrics. A system block diagram for this Chameleon All Programmable system example appears in Figure 3.

A key feature of this example IIoT system is its ability to connect to a large number of interface units throughout the substation over standard industrial Ethernet systems using standardized IEEE-1588 Precision Timing Protocol (PTP) and IEC

62439 HSR/PRP. It does so through a compatible industrial Ethernet switch instantiated in the Zynq UltraScale+ MPSoC's programmable logic using IP sourced from SoC-e, a company in the Xilinx ecosystem. This Ethernet switch appears as the large blue box in the diagram. Data from the various sensor sources can be processed in high-speed IP blocks (represented by the red/blue box in the diagram) from Xilinx and from companies in the Xilinx ecosystem, or the processing algorithms can run on one or more

of the Zynq UltraScale+ MPSoC's six ARM processor cores, depending on performance requirements.

The performance/watt metric for this system is approximately 1.2x better than for a comparable system based on CPU/DSP silicon from competitors, and there's a 2:1 reduction in the number of chips needed in this design, thanks to the Zynq UltraScale+ MPSoC's massive programmability, processing capacity and superior I/O flexibility. Clearly, a security application must protect the power grid from malicious attack,

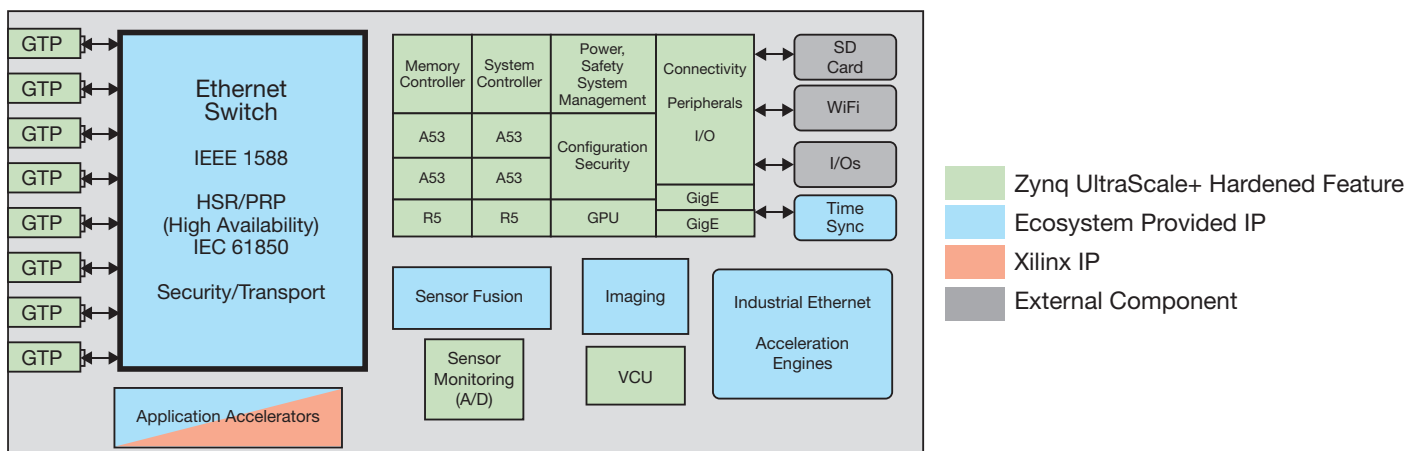


Figure 3 — In this smart-grid substation automation system, an industrial Ethernet switch instantiated in the Zynq UltraScale+ MPSoC's programmable logic sources IP from SoC-e, a company in the Xilinx IIoT ecosystem.

With this expansion of its ecosystem, Xilinx has made it easier for product design teams to achieve aggressive project goals within equally aggressive project schedules.

so the built-in functional-safety and security features of the Zynq UltraScale+ MPSoC are especially useful in this application.

EXAMPLE 4: INDUSTRIAL AUTOMATION

The final Chameleon All Programmable system example is for industrial control and might take the form of a motion controller, programmable logic controller (PLC) or human-machine interface (HMI) system. This IIoT example uses the Zynq UltraScale+ MPSoC to integrate an entire system that might otherwise require four chips (a CPU, a functional-safety processor, a shaft encoder and an FPGA for high-speed power modulation and motor control) into one device, resulting in a 30 percent improvement in performance/watt and a substantial reduction in system pc board real estate. A system block diagram appears in Figure 4.

As in the three other examples, this industrial control system benefits from the Zynq UltraScale+ MPSoC's any-to-

any connectivity and from the functional-safety features embodied in the lockstep capabilities of the integrated dual-core ARM Cortex-R5 processor.

THE ECOSYSTEM LOWDOWN

All four of these examples make ample use of hardware and software IP from Xilinx and its ecosystem member companies. This IP is essential to easing your job of creating advanced, intelligent systems, especially Chameleon platforms that pick and choose which IP to use within each product built with one hardware platform.

Xilinx ecosystem members provide hardware and software IP in four major categories:

1. Domain-specific hardware and software IP for embedded vision and IIoT applications, plus a variety of real-time operating systems;
2. Design enablement, including several high-level design tools;

3. Modules, evaluation boards and production-ready systems-on-modules (SOMs) based on the Zynq-7000 SoC or the Zynq UltraScale+ MPSoC for rapid hardware development and proliferation; and

4. Design services.

Every design team is pressed for time, even as project requirements entail ever-increasing performance and increasingly complex product features. No design team can do it all, quickly. With this newly announced expansion of its ecosystem, Xilinx has made it easier for product design teams to achieve aggressive project goals within equally aggressive project schedules.

To learn more about Xilinx's expanded and strengthened ecosystem, please visit <http://www.xilinx.com/alliance/featured-solution-partners/solutions-by-megatrend.html>. 🌟

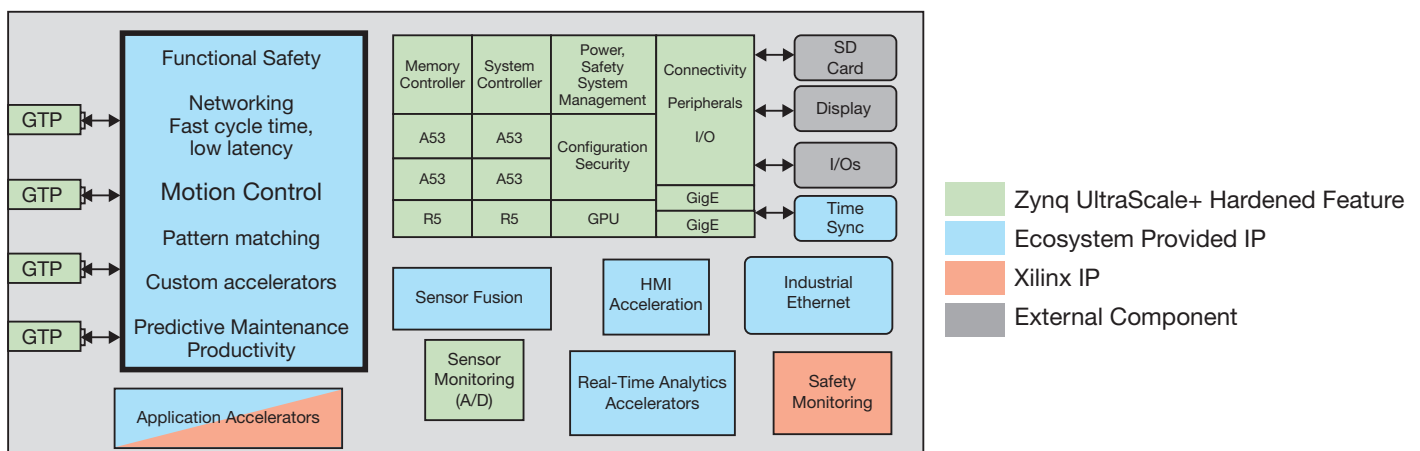


Figure 4 — This industrial automation design for IIoT uses the Xilinx Zynq UltraScale+ MPSoC to integrate an entire system that might otherwise require four chips.

Intelligent Gateways Make a Factory Smarter



by Armando Astarloa
 Founding Partner and CEO
 System-on-Chip engineering S.L.
 armando.astarloa@soc-e.com

An intelligent gateway powered by the Zynq SoC enhances productivity in a state-of-the-art manufacturing plant.

The Industrial Internet of Things—the idea that all systems should be connected on a global scale in order to share information—is quickly becoming a reality. Today, a growing number of companies, especially in the industrial equipment markets, are taking IIoT one step further by creating complex systems that integrate sensors, processing and communications to form intelligent factories, smart energy grids and even smart cities. These developments increase productivity and profitability, as well as enrich lives.

New technology implemented on a Xilinx® Zynq®-7000 All Programmable SoC is helping to bring intelligent systems into the manufacturing sector of the IIoT. The smart gateway, designed by System-on-Chip engineering S.L. (SoC-e), streamlines productivity and helps companies like Microdeco become more reliably connected and secure.

To maximize profitability, factories seek more flexibility in their layouts, more information about the process and manufactured products, more intelligence in the processing of this data and an effective integration of the human experience/interaction. However, as new technology is introduced into the factory sector, those creating it need to respect some rules. The first and most important is that production cannot stop. New technologies must be compatible with old systems and interoperability among vendors should be facilitated. Furthermore, the solutions should provide a means of taking the next step in automation, leading to more autonomous or decentralized analytics.

Factory equipment must communicate with a company's IT network. Intelligent gateways will play a vital role in offering transparent operations between these two worlds: machine and IT.

In order to achieve what many are calling the “fourth industrial revolution,” factories need infrastructure and systems to use the IT and electronics for automated production. Although many factories automated in the third industrialization wave, in many scenarios it is necessary to implement both steps simultaneously: the third and fourth evolutions of automation. This situation offers a good opportunity to integrate IT infrastructures that will fit with new requirements for smart factories but are compatible with the third-

era production-scheduling and automation systems. Figure 1 depicts the typical production system widely used in industry that helps adapt and optimize production to demand. The enterprise resource-planning (ERP) software consists of tools that support the commercial database. It defines what to fabricate. Meanwhile, the manufacturing enterprise system (MES) focuses on the scheduling of production. It uses the ERP outputs, communicates with the production plant equipment and tells the equipment what to do.

NETWORKING, PROCESSING AND SENSING IN THE SMART FACTORY

With many companies offering different types of factory equipment and many generations of that equipment coexisting, connecting equipment from different vendors and different time periods that conforms to different standards can be quite challenging. It's further complicated by the fact that this factory equipment must also communicate with a company's IT network (enterprise and/or Internet); combinations of PC-based systems; gateways, black boxes

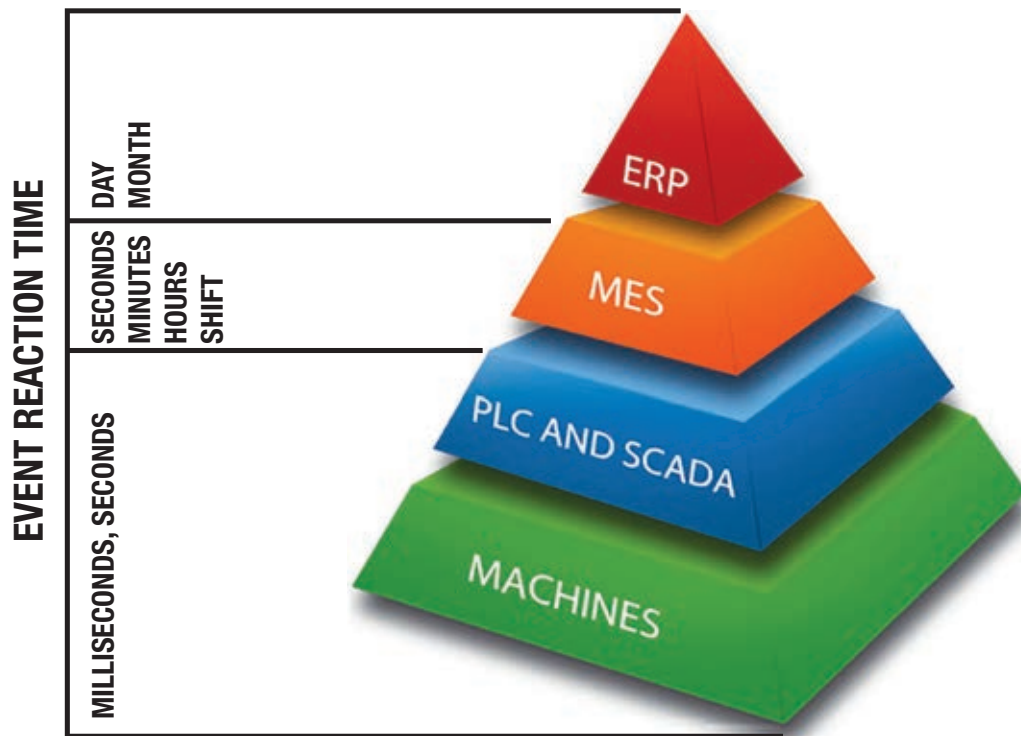


Figure 1 – Scheduling the production via ERP/MES



Figure 2 – The CPPS-Gate40 smart gateway from SoC-e

and industrial switches built around multiple protocols. As such, a factory can quickly turn into a heterogeneous nightmare, lacking the simplicity and flexibility that a “plug-and-work” operation demands. Intelligent gateways like the CPPS-Gate40 from SoC-e (Figure 2) will play a vital role in offering secure and transparent operation between both worlds (machine and IT).

Microdeco is a company that manufactures small metal parts for the automotive sector. The company is always looking for ways to enhance productivity and is at the forefront of using intelligent systems. In the company’s pilot plant, located in Ermua, Spain, Microdeco has built a networking infrastruc-

ture around the concept of smart gateways that combine in the same system networking, processing and sensing.

One of the top challenges in creating a smart factory lies in connecting the various systems. The factory includes high-speed optical links that interconnect the various cyber physical production system (CPPS) areas—that is, each production group of machines, sensors and actuators. The intelligent gateway is in charge of all the communication infrastructure. This includes, the high-speed switching for the fiber links and flexible, trispeed Ethernet ports to implement regular Ethernet or Industrial Ethernet protocols in each cell, along with serial ports to implement widely

used industrial protocols such as Modbus and Profibus.

Figure 3 shows how each smart gateway installed in each machine (CPPS area) is tied to the next one using a single fiber-optic link. The infrastructure is completed by connecting all the devices in a single ring that implements the High-Availability Seamless Redundancy (HSR) protocol. This nonproprietary (IEC 62439-3 Clause 5) Ethernet “zero-delay recovery time” solution allows operators to disconnect any equipment from the ring without adversely affecting other nodes or equipment in the factory. This real plug-and-work operation facilitates plant layout modifications. Furthermore, HSR supports the redundant IEEE 1588v2

submicrosecond synchronization protocol, which simplifies the synchronization of the system to perform precise reconstruction of the sampled sensor data or the implementation of control tasks.

In order to provide seamless redundancy, each HSR node sends the Ethernet frames through both directions of the ring. This approach allows “hot” cable or equipment plugging and unplugging. Each node is in charge of forwarding both frames, and the IEEE 1588v2 support corrects the residence and link delay times to ensure timing accuracy in the entire network. Thus, frame hardware processing is mandatory to ensure low and constant latency times in every node. Indeed, the IEC standard recom-

mends a “cut-through” approach for forwarding the frames in the ring.

To avoid circulating frames, for unicast communications the node that receives the frames is in charge of removing them from the ring. For multicast and broadcast traffic, the sender removes the frames when it sees them again in the redundant port. Additional rules regarding circulating frames (such as corrupted frames) are applied to ensure network stability.

HSR, combined in many cases with the Parallel Redundancy Protocol (PRP), is the recommended High-Availability Ethernet protocol in the standard for the automation of one of the most critical sectors worldwide: power sub-

stations. Other sectors, such as military and aerospace, are also adopting these Layer 2 solutions.

Smart gateways provide hardware switching from the Ethernet and serial ports to the HSR infrastructure ring. There are two smart gateways, represented in the left and in the right of Figure 3, that connect the HSR ring with the Ethernet-based enterprise network working as a redundancy box (Red-Box). Functionally, the access point represented on the right is optional, as it can be used to avoid the single point of failure that would appear in the case of a network using only one RedBox. We recommend implementing the dual-box setup in cases where

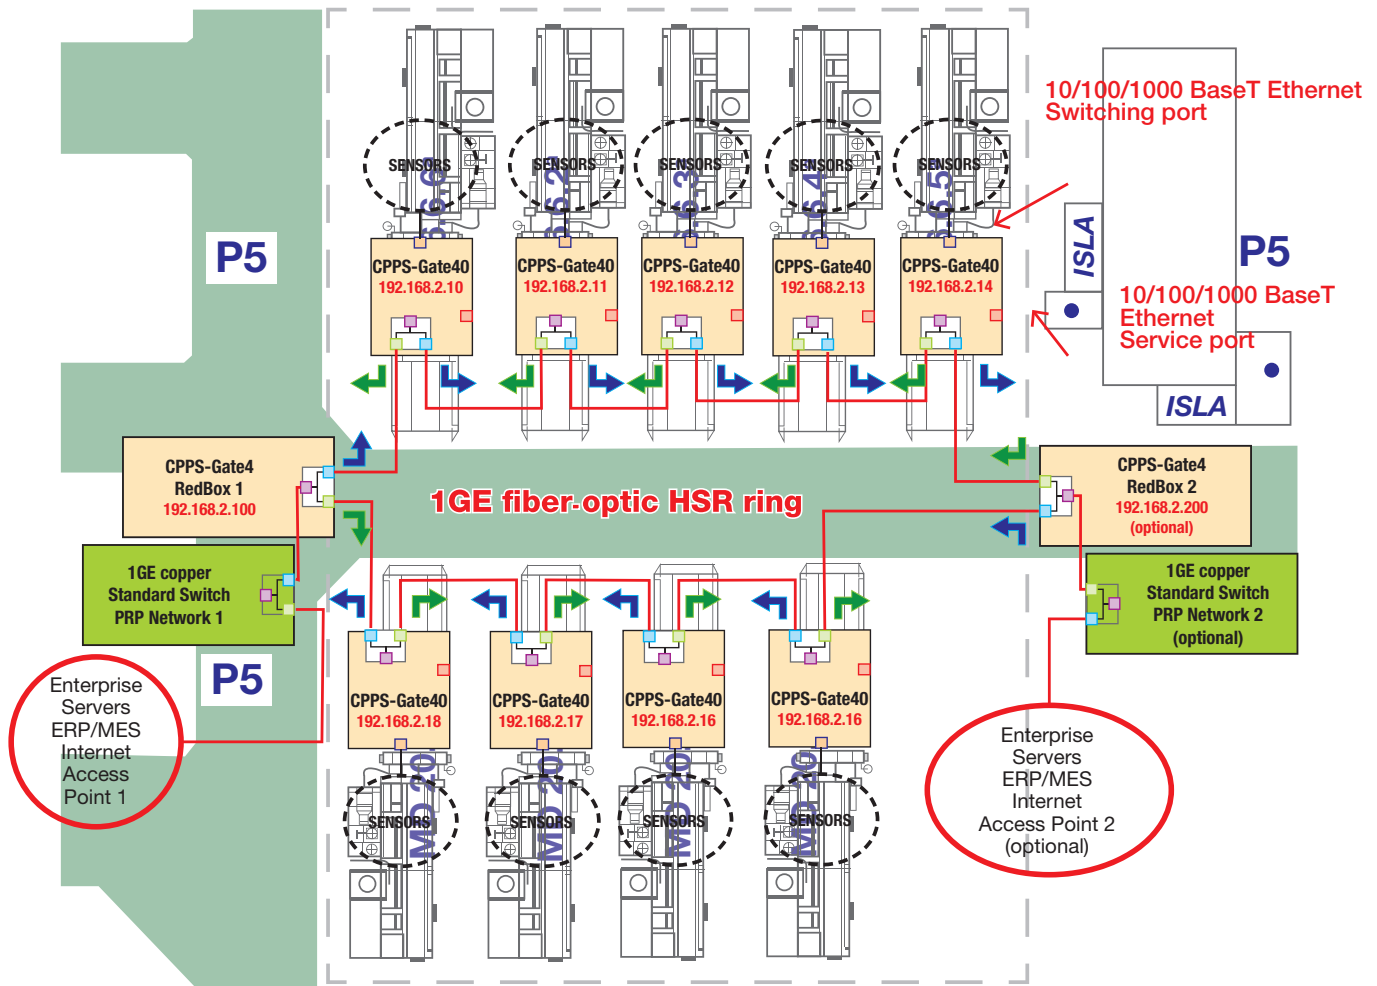


Figure 3 – Lathes section in the Microdeco factory

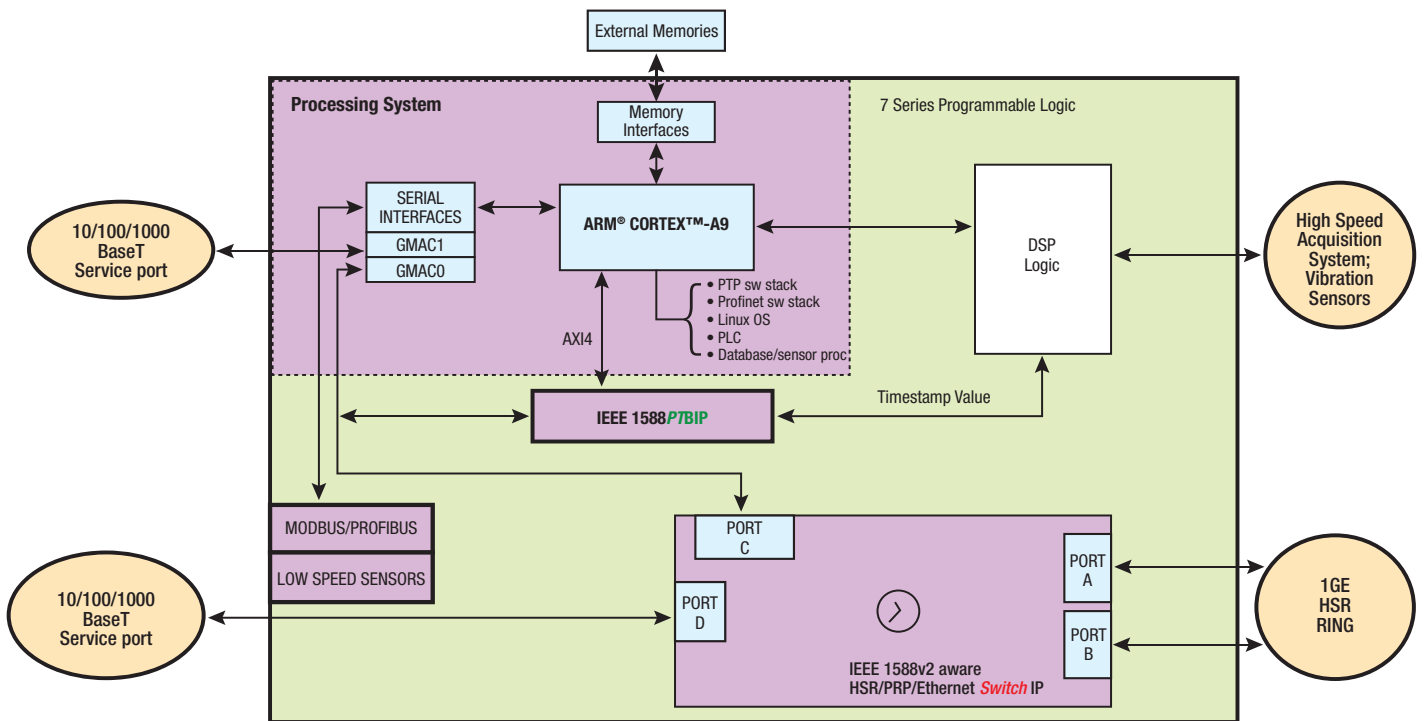


Figure 4 – Block diagram of Zynq SoC implementation

high availability is needed, or when it is necessary to manage PRP frames (IEC 62439-3 Clause 5) in the critical nodes in the enterprise network.

Additionally, there are internal networking ports in the gateway to the processing elements of the SoC device. In most cases, a “dumb” switching approach is useless to join plant and IT worlds. The heterogeneity in the data and network formats makes straightforward connections difficult. What’s needed is a powerful integrated processing system able to talk with local, enterprise or cloud databases. In addition, such a system would be in charge of translating protocols, managing HMI systems, supporting MES systems and even running soft PLCs for real-time control. But that is not all. The customer also expects such a system to perform complex sensor data preprocessing and filtering in the equipment, and of course, advanced cybersecurity operations.

The cybersecurity requirements in these kinds of advanced manufacturing facilities vary widely. Advanced security is necessary to protect the status of the production itself, avoiding any malicious or accidental interruption generated by any cyber infrastructure (device, network, software or hardware). It is also necessary to authenticate users and devices that are accessing information or any critical operation. Furthermore, this information and the control protocols need to be protected in terms of authentication and privacy, because factory networks are connected to larger IT networks in an enterprise and outside of it.

These challenges can only be addressed with a layered cybersecurity approach that takes into account each plant implementation. A common element in all the projects is the need to support secure boot and storage with encryption and authentication. This feature will make credible the implementation of secure software and secure net-

works. The trusted embedded system is more and more difficult to protect due to the increasing number of devices and their heterogeneity.

For authentication and for networking security, these systems can directly use many of the solutions present in the IT world today. Well-known authentication mechanisms like IEEE 802.1X combined with RADIUS are a good example. Many embedded systems with high-level operating systems can run cryptographic libraries (such as OpenSSL) to support all the secure Layer 3 protocols and applications useful for secure data interchange. However, a big challenge arises when it is necessary to secure Layer 2 industrial protocols with strict real-time requirements. The analysis of these scenarios shows that the software approach of protecting these frames by applying cryptographic algorithms, even using crypto accelerators, is not straightforward, and in many cases custom hardware processing is required.

In the presented topology, from the network and user point of view, it is necessary to secure three network links—the redundant HSR/PRP, the 10/100/1G switching port and the service ports—with authentication mechanisms. Furthermore, due to all the plant traffic passing through the intelligent gateway, the three links will play a vital role in monitoring traffic for potential threats.

A final concept is the integration of a sensor interface suite. As discussed, the advances in the technology should help us to simplify the installations, not make them more complex. To fulfill this demand, we integrated all the standard digital and analog interfaces in the gateway. Additionally, we also included high-end interfaces for advanced vibration sensors and high-speed data acquisition interfaces with direct access to the Zynq SoC device.

HOW SOC PROGRAMMABLE PLATFORMS DRIVE THE CHANGE

The “magic” of merging high-end networking, powerful processing and sensing capabilities has been obtained thanks to SoC programmable platforms. Our product, named CPPS-Gate40, embeds a Xilinx Zynq-7000 All Programmable SoC device implemented on the SoC-e SMARTzynq OEM module. The dual-core ARM® Cortex™-A9 MPCore™ on the device is complemented with different memory resources (DDR3, flash, massive storage units, etc.) and hardware to support multiple high-speed networking links. This infrastructure offers a huge level of freedom to partition hardware and software processing in order to face the challenges these applications present.

From the hardware perspective, the Zynq SoC’s programmable logic is the perfect candidate to implement the

low-latency networking tasks combined with the IEEE 1588v2 hardware support units. Figure 4 is a block diagram of the SoC implementation for the CPPS-Gate40 in the Microdeco implementation. The network’s switching infrastructure is coordinated by the SoC-e HSR/PRP/Ethernet switch (HPS) IP core, which ensures a constant forwarding time of 550 nanoseconds in each node of the ring and integrates internal and external trisped Ethernet ports.

The internal port is sniffed and time-stamped by the Precise Time Basic (PTB) IP core, providing support for the PTP stack. This IEEE 1588v2 infrastructure allows the smart gateway to work as master, slave, transparent clock and boundary clock. Thus, at the end, in each piece of equipment a synchronized 64-bit timer can be used for time-stamping, synchronization, control and as a common time reference

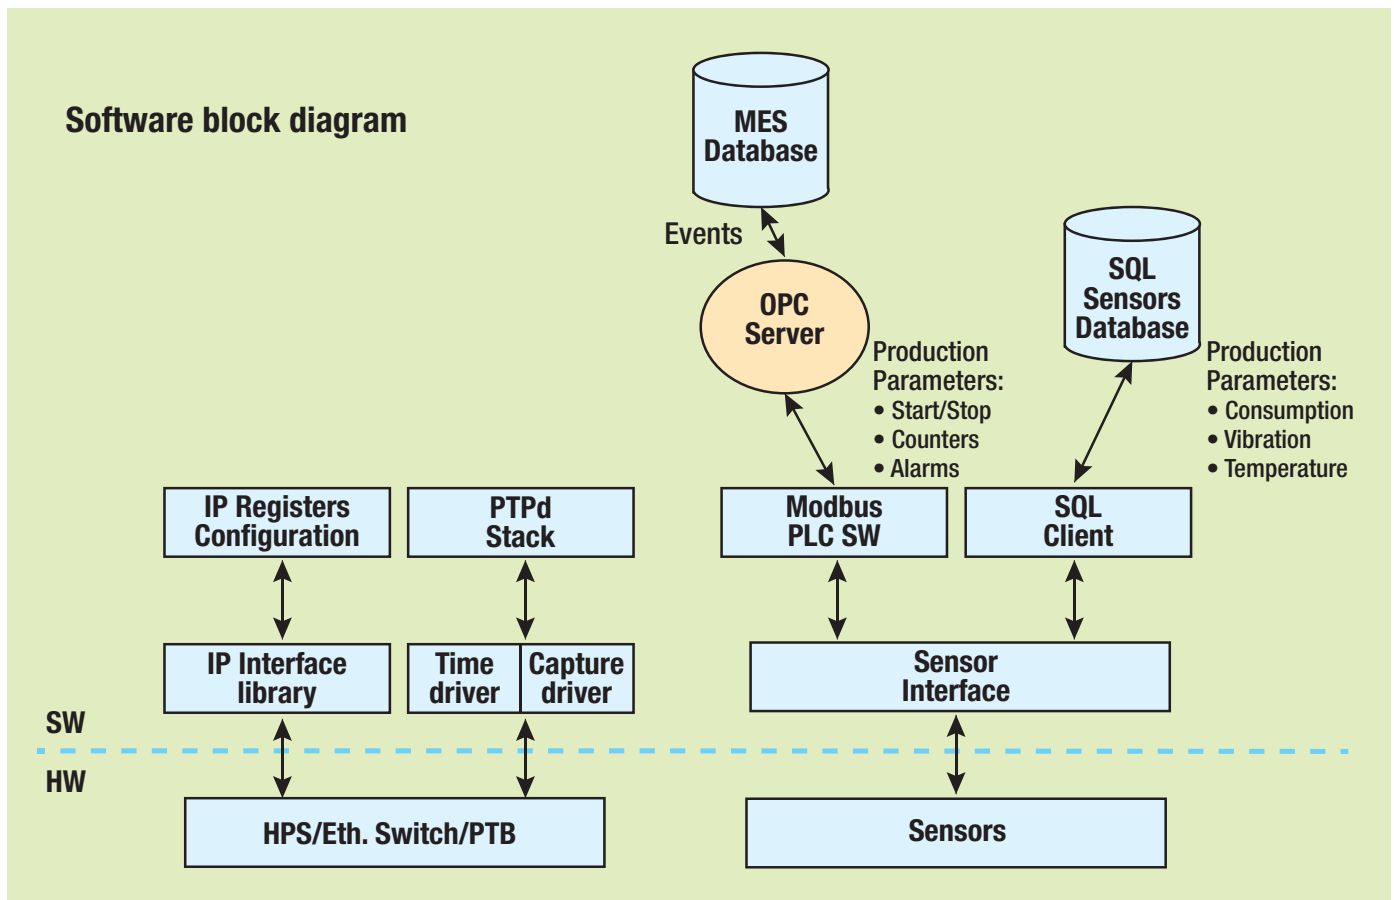


Figure 5 – Software infrastructure for the smart-factory network

to implement Time-Sensitive Networking (TSN) networks.

These networking cores implemented on the FPGA section of the Zynq SoC are also ready to support cybersecurity features such as IEEE 802.1X authentication. This mechanism, combined with an external authentication server, protects non-authorized connections to the network ports. The Zynq SoC's programmable logic also plays a vital role in securing Layer 2 control-related frames on the fly, like the authentication needed in the IEEE 1588v2 transparent clock operation.

The cybersecurity is further enhanced by the Zynq SoC's secure boot. All the external software and bitstreams external from the device, even the bootloader and OS, are stored, AES-256 encrypted and HMAC authenticated. This feature, combined with other hardware security protections included in the device, ensures that data throughout the cyber infrastructure comes from trusted origins.

Additionally, a SIEM agent installed in each CPPS-Gate40 runs (among others) the following security-related tasks: surveillance of new connections, authentication attempts, SSH connections and access to analytics tools; virus/malware detection; network attacks identification; and ARP traffic analysis.

The sensor interfaces are also implemented on the programmable logic section (high-speed data acquisition, digital filtering and FFT) and via some of the standard communication channels present on the Zynq SoC's processing system (UART, I2C, SPI).

The software infrastructure implemented on this equipment benefits from the seamless integration of Linux OS Ubuntu's distribution on the device. The list of features that Linux supports is extensive. For Microdeco's specific implementation, Figure 5 summarizes the most relevant software services implemented on top of the Linux OS.

A Python-based PLC emulator has been developed as the key piece to map sensor interfaces in a well-known Modbus TCP scheme. This approach simplifies the communication with the third-party MES

software. In parallel, a SQL client transfers raw and preprocessed sensor data packets to a remote SQL server. Specific alarms and selected data are directly published in a cloud-based couchDB database. The data analysis can be performed remotely in the enterprise or cloud server and even locally on the smart gateway. For this last purpose, the product includes a temporal database that can predict failures or other defined behaviors in the production and act locally. Big-data analysis software provided by Juxt.io is in charge of performing the predictive analytics related to machine behavior.

Network management is supported via SNMP thanks to SoC-e's Portable Tools API. The cybersecurity infrastructure is built around the hardware support of SoC-e IP and the integrated SIEM agent for network and user activity surveillance.

INCREASING PROFITS THROUGH TECHNOLOGY

Germany's Fraunhofer Institute for Industrial Engineering and Automation forecasts that Industry 4.0 may lead to a leap in productivity of 20 to 30 percent by 2025. However, the industrial sector needs progressive changes and friendly technologies and solutions. The Microdeco plant, for example, benefits from high-end technologies to integrate flexible and computationally powerful networking and processing infrastructures in its production lines.

The drivers of this approach are the adoption of open standards for networking and for the data formats; the use of extensible and repartitionable SoC reconfigurable devices; and the selection of software frameworks that offer a high level of productivity (like Python over embedded Linux). Furthermore, manufacturers can drastically reduce their time-to-market in addressing this new market by means of the ready-to-use, value-added hardware IP now available. And of course, the system must also come with the highest levels of cybersecurity at the device, software and networking levels.

For more information on SoC-e's IIoT IP portfolio, visit www.soc-e.com.

All Programmable FPGA and SoC modules



difference by design



Available SoMs

KINTEX⁷
UltraSCALE

ZYNQ⁷
UltraSCALE+

ZYNQ⁷

KINTEX⁷

ARTIX⁷

SPARTAN⁶

- customizable
- long term supply guarantee
- broadest range of products
- integration support
- automotive, industrial and commercial temperature ranges
- very small form factor
- ruggedized
- inhouse production
- rapid prototyping

Design Services

- Module customization
- Carrier board customization
- Custom project development

XILINX[®]

ALLIANCE PROGRAM
CERTIFIED MEMBER — BASE

.de
.de
.de

Evaluating an IQ Compression Algorithm Using Vivado High-Level Synthesis

by Stefan Petko

Design Engineer
Xilinx, Inc.
stefan.petko@xilinx.com

Duncan Cockburn

Design Engineer
Xilinx, Inc.
duncan.cockburn@xilinx.com

Xilinx's Vivado HLS tool helps tame the escalating cost of wireless fronthaul network infrastructure.

Wireless network operators face a major challenge in maintaining the bottom line while increasing the capacity and density of their networks. A compression scheme for wireless interfaces can help by reducing the required fronthaul network infrastructure investment.

We used the Vivado® Design Suite's high-level synthesis (HLS) tool to evaluate an Open Radio equipment Interface (ORI) standard compression scheme for E-UTRA I/Q data to estimate its impact on signal fidelity, the introduced latency and its implementation cost. We found that Xilinx®'s Vivado HLS offered an efficient platform for evaluating and implementing the selected compression algorithm.

WIRELESS BANDWIDTH PRESSURE

The ever-increasing demand for wireless bandwidth drives the need for new network capabilities such as higher-order MIMO (multiple-input, multiple-output) configurations and carrier aggregation, for example. The resulting increase in network complexity leads the operators to architectural changes such as the centralization of baseband processing to optimize their network resource utilization. While reducing baseband processing costs, the sharing of baseband processing resources increases the complexity of the fronthaul network.

These fronthaul networks, transporting the modulated antenna carrier signals between the baseband units (BBU) and remote radio heads (RRH),

are most frequently implemented using the Common Public Radio Interface (CPRI) protocol over optical fiber. The CPRI protocol requires a constant bit rate and its specification has over the years increased the maximum data rate to match the increasing bandwidth demands. Network operators are now looking at technologies that will allow them to achieve a significant hike in data rate without increasing the number of optical fibers in use, thus maintaining current capex and opex overheads associated with a cell site.

In an effort to provide a long-term solution, the network operators are looking at alternative network arrangements including rearchitecting the interface between the baseband-processing and radio units in order to reduce the fronthaul bandwidth. However, functional rearrangements can make it more difficult to meet the stringent performance requirements for some wireless interface specifications.

An alternative way to reduce bandwidth is to implement a compression/decompression (codec) scheme for wireless interfaces that are nearing or exceeding the available throughput. The achievable compression ratios depend on the specific wireless-signal characteristics such as noise levels, dynamic range and oversampling rates.

Let's take a closer look at an ORI standard compression scheme for E-UTRA IQ data—the real and imaginary components of the transmitted modulation symbols. A simplified application, shown in Figure 1, illustrates the placement of a compression-and-decompression module at the CPRI IQ input-and-output interface.

The filter design process exploits certain channel characteristics to minimize signal loss due to downsampling and upsampling.

IQ COMPRESSION ALGORITHM

The ORI standard is a refinement of the CPRI specification aiming to enable an open BBU/RRH interface. In its latest release, ORI specifies a lossy time-domain E-UTRA data compression technique for channel bandwidths of 10, 15 or 20 MHz. The combination of a fixed 3/4 rate resampling and non-linear quantization of 15-bit IQ samples achieves a 50 percent reduction in bandwidth requirements, facilitating an 8 x 8 MIMO configuration covering two sectors over a single 9.8-Gbps CPRI link, for example.

The resampling stage involves interpolating the input I and Q streams, passing the interpolated data through a low-pass filter and decimating the output data stream. The filter design process exploits specific channel characteristics in order to minimize signal loss due to downsampling and upsampling stages. For example, the effective

bandwidth of an OFDMA output for a 20-MHz E-UTRA downlink channel sampled at 30.72 MHz is 18.015 MHz, implying a lossless ideal low-pass filter response at 3/4 resampling rate.

The nonlinear quantization (NLQ) process transforms 15-bit baseband IQ samples with normal distribution into 10-bit quantized values. NLQ minimizes the quantization error by using a cumulative distribution function (CDF) with a specified standard deviation to represent amplitudes occurring more frequently with a finer granularity than amplitudes occurring less frequently. As illustrated by our results in Figure 2b, the quantized constellation fills a significantly higher proportion of its reduced number range than the input constellation shown in Figure 2a, minimizing the quantization error in comparison with alternative linear quantization schemes. Typically implemented in a

lookup table, I and Q samples can be quantized independently using their corresponding distribution functions.

We compared the ORI IQ compression performance to a Mu-Law compression algorithm implementation as specified by the ITU-T Recommendation G.711. Also a nonlinear quantization technique, Mu-Law uses a logarithmic function to redistribute the quantized values across the available number range. Unlike CDF-based quantization, which considers the statistical distribution of the input samples, Mu-Law quantized output is a function of the corresponding input sample value and the specified companding value.

To make the comparison for an equivalent compression ratio of 50 percent, we considered a 16- to 8-bit Mu-Law encoder. Because it does not require resampling, Mu-Law compression is a cheaper solution in terms of latency




Figure 1 – Simplified wireless system utilizing CPRI IQ compression

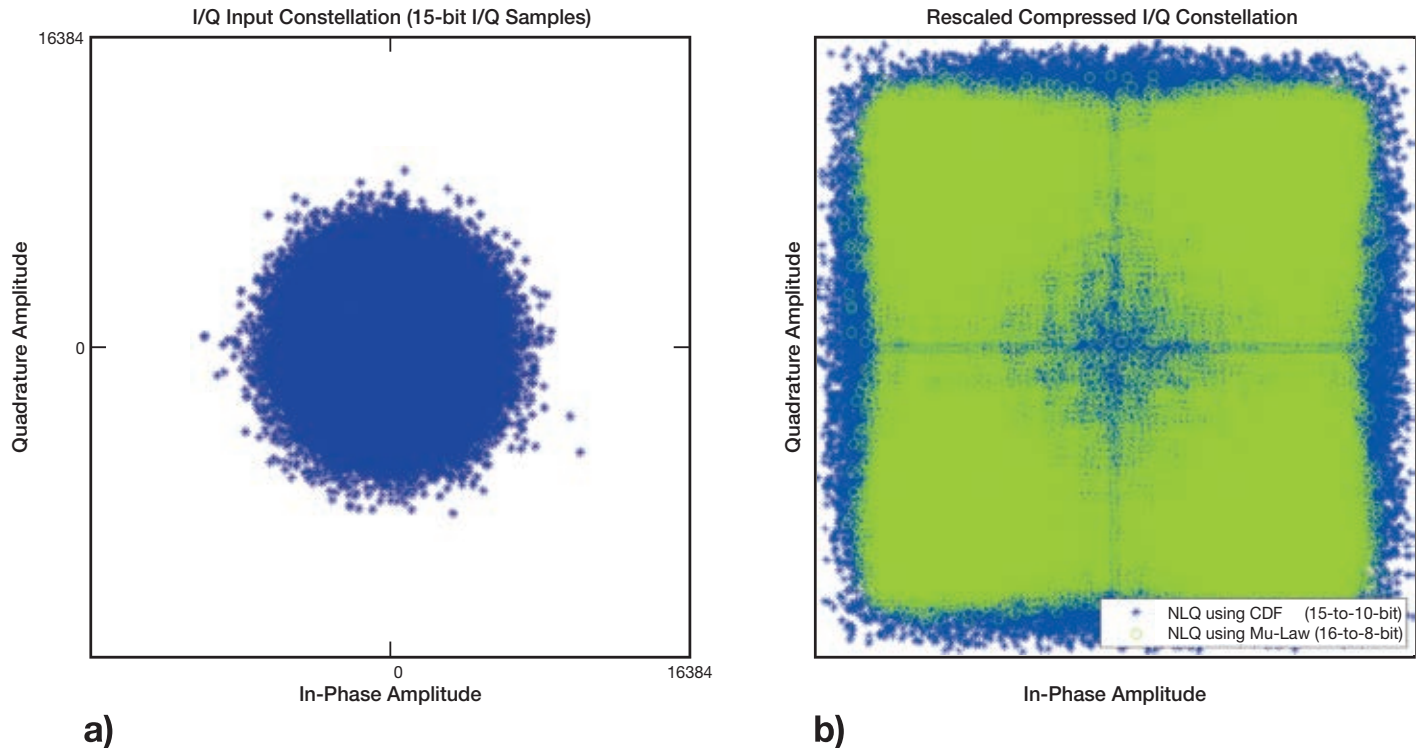


Figure 2 – IQ constellation of the reference 20-MHz E-UTRA DL channel input frame (a) and the compressed IQ data (b) scaled to illustrate the effective number range utilization for each constellation

and implementation resource cost, offering a trade-off between the design complexity and the achievable reconstructed-signal fidelity.

SCALING UP THE CODEC ARCHITECTURE

For our prototype configuration, we aimed to scale up the compression algorithm to fully utilize a 9.8304-Gbps CPRI link (line bit rate option 7). The ORI-compressed E-UTRA sample specification allows us to transport 16 compressed IQ channels (32 I and Q channels compressed independently) over a single 9.8G CPRI link. A target throughput of three compressed samples per CPRI clock is sufficient to fully pack the 32-bit Xilinx LogiCORE™ IP CPRI IQ interface, giving us the required 737.28-Msample/second compression IP output.

Targeting a single clock domain, we needed to architect the resampling filter to meet the output rate of three samples per clock cycle. Interpolating the input sample stream with 0's al-

lowed us to ignore the noncontributing input samples. The output stream becomes a function of the interpolation rate number of subfilters, each using a subset of the FIR coefficients (total number of coefficients/interpolation rate). Each of the four parallel filters operates on a subset of channels, making the overall throughput equal to the required three compressed samples per clock cycle. Along with the high throughput, the proposed architecture reduces the resampling latency, since only a fraction of coefficients is used in each subfilter.

On the compression path, we computed the NLQ quantization table using the cumulative distribution function (CDF). Assuming a symmetric IQ distribution, we reduced the size of the NLQ lookup table to 2^{14} entries of 9-bit quantized values. Because our design required three parallel lookups per clock cycle, we implemented three parallel lookup tables utilizing the same quantization values. The quantization

levels can be computed independently for I and Q samples using the expected or observed standard-deviation values. Alternatively, subsets of channels can be quantized independently based on the actual signal-level measurements or higher-level network parameters. For decompression, we used the quantile function (inverse CDF) to compute the Inverse NLQ table. Its size is limited to 2^9 entries of 14-bit values.

We tested the implemented codec algorithm using a 20-MHz LTE E-UTRA FDD channel stimulus generated with MATLAB®'s LTE System Toolbox. We then used Keysight VSA to demodulate the captured IQ data and quantify the signal distortion due to compression and decompression stages by measuring the output waveform error vector magnitude (EVM). We compared the reported output EVM measurements—which represent the difference between the ideal and the measured signal—to the reference input signal EVM.

HIGH-LEVEL MODELING AND IMPLEMENTATION FLOW

We started the implementation process by developing single-channel compression and decompression models using

the GNU Octave language, utilizing its signal-processing and statistics packages. In addition to providing useful verification reference data points, the model output also generated a set of FIR filter

coefficients and quantization tables.

From a high-level mathematical model, the Vivado HLS tool offers an obvious transition path to evaluate the proposed architecture in terms of the

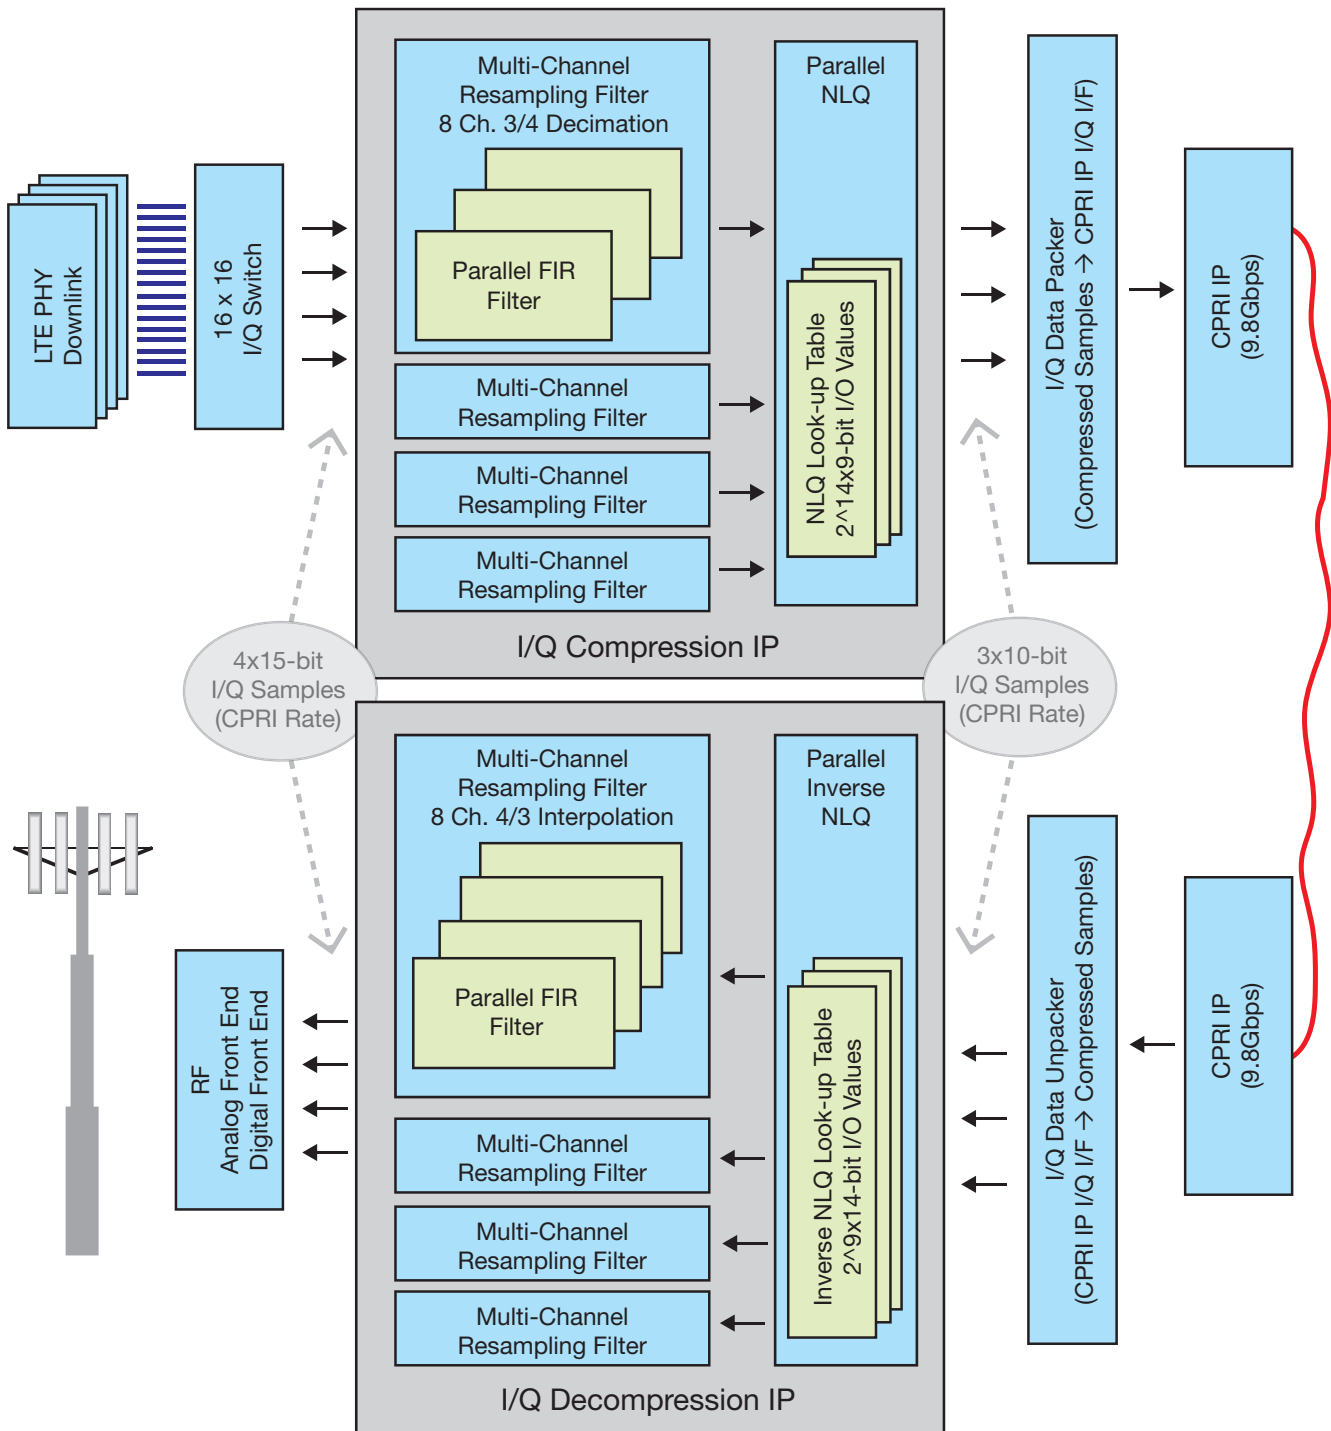


Figure 3 – IQ codec architecture indicating sample processing rates at the codec’s IP interfaces (downlink path only)

potential hardware performance and cost. We set up our C++ testbench to operate on the input data streams using the compress and decompress functions. We synthesized these functions independently, since they would be physically placed at the opposite ends of a CPRI link. We implemented all of the external and internal function interfaces using HLS streams with the interleaved channel dataflow managed with simple C++ loops.

We utilized the Vivado HLS FIR IP to prototype the resampling filter. To meet the high-throughput requirements of our design, we implemented parallel single-rate FIR filters and used a loop-based filter output decimation.

It's possible to obtain a more resource-efficient resampling filter by implementing a polyphase resampling filter. One such ready-to-use alternative supporting the ORI resampling rates is the Multi-Channel Fractional Sample Rate Conversion Filter described in Xilinx application note [XAPP1236](#), "Multi-Channel Fractional Sample Rate Conversion Filter Design Using Vivado-High-Level Synthesis."

The benefits of a fast C-level simulation become even more evident when the verification data sets are large. This is very much the case when evaluating an IQ compression algorithm since at minimum, a full radio frame of data (307,200 IQ samples per channel) is required to make use of the VSA tools for EVM measurements. We observed a simulation speed-up of two orders of magnitude for C simulation compared with C/RTL co-simulation, translating to a nine-hour co-simulation run vs. a five-minute C simulation for our compression IP test run.

Another significant HLS testbench advantage was the ease of input data use and output data capture utilizing files in conjunction with the HLS streams. The result was to provide an interface for data analysis with VSA tools or direct comparison against the Octave model output in the C++ testbench.

PERFORMANCE MEASUREMENTS

The Keysight VSA measurements reported an averaged EVM of 0.29 percent for a codec configuration with 144 FIR coefficients. Compared with the original input data with EVM RMS of 0.18 percent, the additional EVM attributable to the compression-decompression processing chain is 0.23 percent. By comparison, the Mu-Law compression algorithm operating on the equivalent input data set results in an average EVM of 1.07 percent.

The reduced latency and resource utilization cost of the Mu-Law compression would make it a better choice over the ORI-proposed IQ compression scheme in situations where you can allocate as much as 1 percent of the entire LTE downlink signal-processing chain's 8 percent EVM budget to IQ compression. However, any additional signal distortion implies tighter performance targets for the remaining system components. The potential IQ compression cost benefits may therefore be offset by cost increases of the digital front-end components and power amplifiers, for example.

Vivado high-level synthesis confirmed the required throughput reported in terms of the initiation interval—the number of clock cycles before the top-level task is ready to accept new input data. We also verified that the exported Vivado IP Integrator cores met the timing requirements for the target Kintex® UltraScale™ platform.

We limited the scope of our investigation to a small set of configurations and input data vectors. However, once the system model and the corresponding C model are in place, it will be possible to customize, implement and evaluate an alternative configuration in minutes.

DESIGN ALTERNATIVES

From the design tool perspective, Vivado HLS provides a viable hardware prototyping path. A high-level testbench fits well within a design framework that requires flow of data among a number

of design and verification tools. The main advantage of such a testbench is the ability to perform fast C-level simulations of a hardware system model. For IQ compression and similar applications, simulation runs involve frequent higher-level parameter or input data set changes, making a fast feedback essential.

Measurements show that the proposed ORI compression scheme can deliver signal distortion below 0.25 percent for a 20-MHz E-UTRA downlink channel. While the compression performance does depend on the input channel characteristics, the ORI compression scheme provides a scope for performance tuning by choosing the optimal set of filter design and quantizer parameters.

Our prototype utilizes a common static set of design parameters for all 16 antenna-carrier data streams. In a real system, the signal characteristics would be either known in advance or could be measured and used to tune the design. Alternatively, the compression performance could be dynamically adjusted by reconfiguring the quantization tables to maintain the minimal required signal fidelity.

The cost in terms of implementation resources and additional latency required to perform compression and decompression would also be considered in conjunction with the required compression performance. The resampling filter size and latency dominate the overall codec costs, and a greater EVM tolerance would allow for a design with fewer filter coefficients.

With time-to-market considerations in mind, Xilinx has created an ORI-based IQ codec proof of concept. You can read more about this scheme and request access to the design files on the Xilinx website. Visit <http://www.xilinx.com/applications/wireless-communications/wireless-connectivity.html> or e-mail Perminder Tumber, wireless business marketing manager for connectivity, at Permind@xilinx.com. 🌟

Xilinx FPGA Enables Scalable MIMO Precoding Core

by **Lei Guan**

Member of Technical Staff
Bell Laboratories, Nokia
lei.guan@ieee.org

Researchers at Bell Labs Ireland have built a frequency-dependent precoding core with high-performance FPGAs for generalized MIMO-based wireless communications systems.

Massive-MIMO wireless systems have risen to the forefront as the preferred foundation architecture for 5G wireless networks. A low-latency precoding implementation scheme is critical for enjoying the benefits of the multi-transmission architecture inherent in the multiple-input, multiple-output (MIMO) approach. Our team built a high-speed, low-latency precoding core with Xilinx® System Generator and the Vivado® Design Suite that is simple and scalable.

Due to their intrinsic multiuser spatial-multiplexing transmission capability, massive-MIMO systems significantly increase the signal-to-interference-and-noise ratio at both the legacy single-antenna user equipment and the evolved multi-antenna user terminals. The result is more network capacity, higher data throughput and more efficient spectral utilization.

But massive-MIMO technology does have its challenges. To use it, telecom engineers need to build multiple RF transceivers and multiple antennas based on a radiating phased array. They also have to utilize digital horsepower to perform the so-called precoding function.

Our solution was to build a low-latency and scalable frequency-dependent precoding piece of intellectual property (IP), which can be used in Lego fashion for both centralized and distributed massive-MIMO architectures. Key to this DSP R&D project were high-performance Xilinx 7 series FPGAs, along with Xilinx's Vivado Design Suite 2015.1 with System Generator and MATLAB®/Simulink®.

PRECODING IN GENERALIZED MIMO SYSTEMS

In a cellular network, user data streams that radiate from generalized MIMO transmitters will be “shaped” in the air by the so-called channel response between each transmitter and receiver at a particular frequency. In other words, different data streams will go through different paths, reaching the receiver at the other end of the airspace. Even the same data stream will behave differently at certain times because of a different “experience” in the frequency domain.

This inherent wireless-transmission phenomenon is equivalent to applying a finite impulse response (FIR) filter with particular frequency response on each data stream, resulting in poor system performance due to the introduced frequency “distortion” by the wireless channels. If we treat the wireless channel as a big black box, only the inputs (transmitter outputs) and outputs (receiver inputs) are apparent at the system level. We can actually add a pre-equalization black box at the MIMO transmitter side with inversed channel response to precompensate the channel black-box effects, and then the cascade system will provide reasonable “corrected” data streams at the receiver equipment.

We call this pre-equalization approach precoding, which basically means applying a group of “reshaping” coefficients at the transmitter chain. For example, if we are going to transmit N_{RX} independent data streams with N_{TX} (number of transmitters) antennas, we will need to perform a pre-equalization precoding at a cost of $N_{RX} \times N_{TX}$

temporary complex linear convolution operations and corresponding combining operations before radiating N_{TX} RF signals to the air.

A straightforward low-latency implementation of complex linear convolution is a FIR-type complex discrete digital filter in the time domain.

SYSTEM FUNCTIONAL REQUIREMENTS

Under the mission to create a low-latency precoding piece of IP, my team faced a number of essential requirements.

1. We had to precode one data stream into multiple-branch parallel data streams with different sets of coefficients.
2. We needed to place a 100-plus tap-length complex asymmetric FIR

function at each branch to provide reasonable precoding performance.

3. The precoding coefficients needed to be updated frequently.
4. The designed core must be easily updated and expanded to support different scalable system architectures.
5. Precoding latency should be as low as possible with given resource constraints.

Moreover, besides attending to the functional requirements for a particular design, we had to be mindful of hardware resource constraints as well. In other words, creating a resource-friendly algorithm implementation would be beneficial in terms of key-limited hardware resources such as DSP48s, a dedicated hardware multiplier on Xilinx FPGAs.

HIGH-SPEED, LOW-LATENCY PRECODING (HLP) CORE DESIGN

Essentially, scalability is a key feature that must be addressed before you begin a design of this nature. A scalable design will enable a sustainable infrastructure evolution in the long term and lead to an optimal, cost-effective deployment strategy in the short term. Scalability comes from modularity. Following this philosophy, we created a modularized generic complex FIR filter evaluation platform in Simulink with Xilinx System Generator.

Figure 1 illustrates the top-level system architecture. **Simulink_HLP_core** describes multibranch complex FIR filters with discrete digital filter blocks in Simulink, while **FPGA_HLP_core** realizes multibranch complex FIR filters with Xilinx resource blocks in System Generator, as shown in Figure 2.

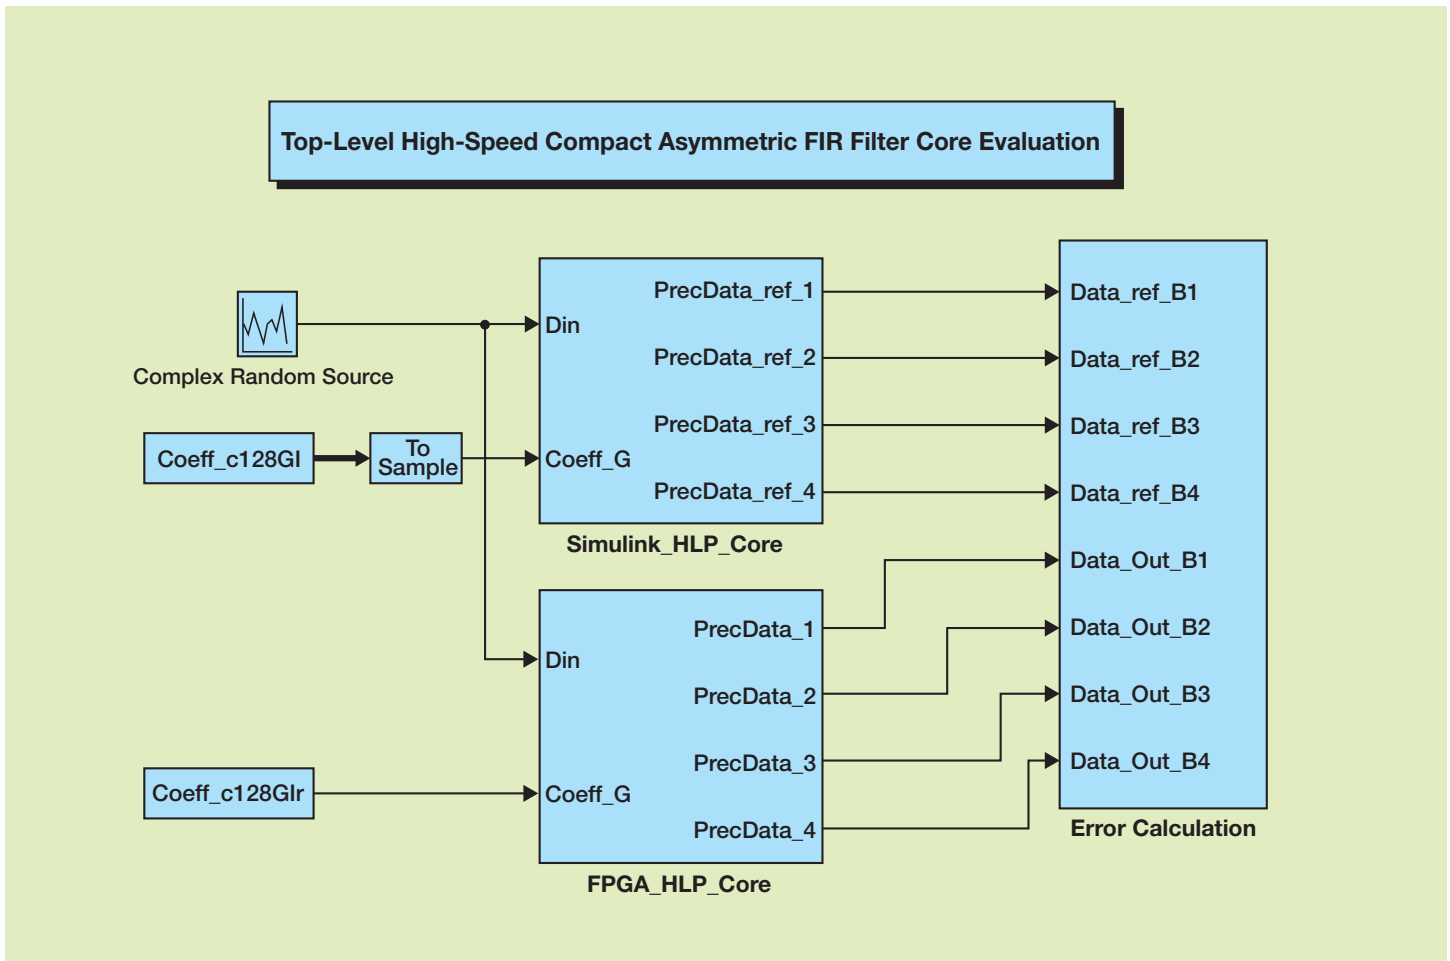


Figure 1 – Simulink evaluation of the top-level precoding core

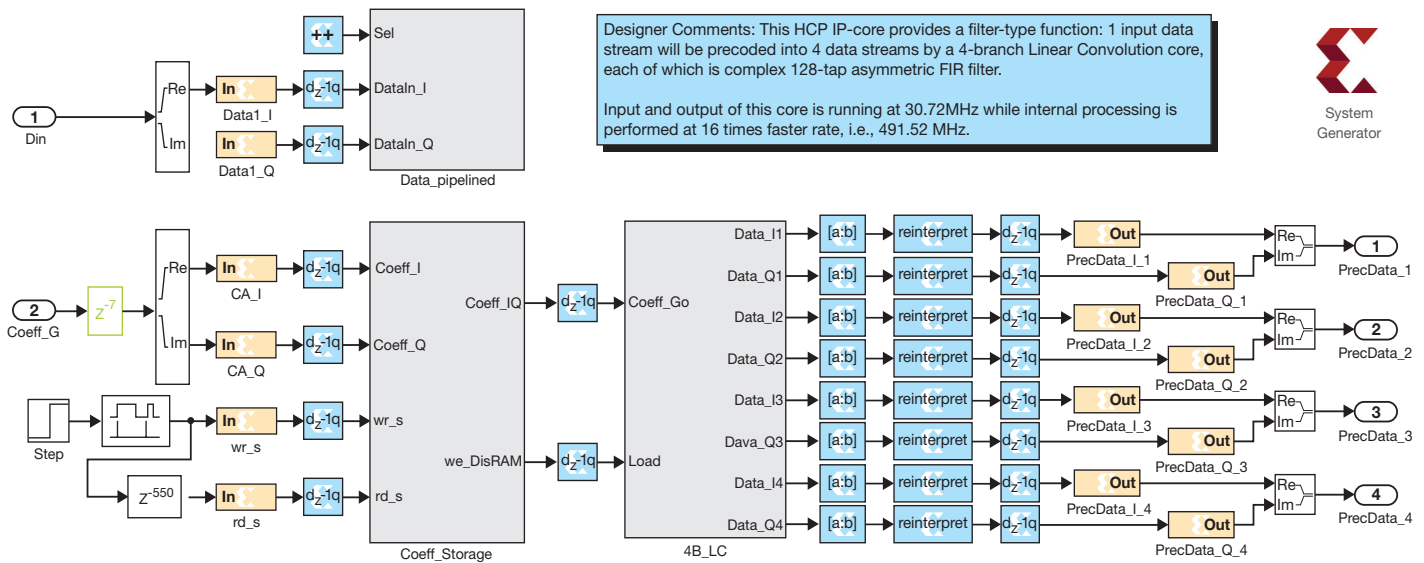


Figure 2 – Top level of the FPGA_HLP_Core

	Full Parallel	Full Serial	Partial Parallel-A*	Partial Parallel-B
Required F_{CLK} (MHz)	30.72 x 1 = 30.72	30.72 x 128 = 3932.16	30.72 X 16 = 491.52	30.72 X 8 = 245.76
1-branch LC core	128 CM	1 CM	8 CM	16 CM
4-branch LC core	512 CM	4 CM	32 CM	64 CM

* Suggested high-speed, low-latency precoding (HLP) core architecture

Table 1 – Complex multiplier (CM) utilization comparison for a 128-tap complex asymmetric FIR filter

Different FIR implementation architectures lead to different FPGA resource utilizations. Table 1 compares the complex multipliers (CM) used in a 128-tap complex asymmetric FIR filter in different implementation architectures. We assume the IQ data rate is 30.72 Msamples/second (20-MHz bandwidth LTE-Advanced signal).

The full parallel implementation architecture is quite straightforward according to its simple mapping to the direct-I FIR architecture, but it uses a lot of CM resources. A full serial implementation architecture uses the fewest CM resources

by sharing the same CM unit with 128 operations in a time-division multiplexing (TDM) manner, but runs at an impossible clock rate for the state-of-the-art FPGA.

A practical solution is to choose a partially parallel implementation architecture, which splits the sequential long filter chain into several segmental parallel stages. Two examples are shown in Table 1. We went for plan A due to its minimal CM utilization and reasonable clock rate. We can actually determine the final architecture by manipulating the data rate, clock rate and number of sequential stages thus:

$$F_{CLK} = F_{DATA} \times N_{TAP} \div N_{SS}$$

where N_{TAP} and N_{SS} represent the length of the filters and number of sequential stages.

Then we created three main modules:

- 1. Coefficients storage module:** We utilized high-performance dual-port Block RAMs to store IQ coefficients that need to be loaded to the FIR coefficient RAMs. Users can choose when to upload the coefficients to this storage and when to update the coefficients of the FIR filters by **wr** and **rd** control signals.

2. Data TDM pipelining module:

We multiplexed the incoming IQ data at a 30.72-MHz sampling rate to create eight pipelined ($N_{SS} = 8$) data streams at a higher sampling rate of $30.72 \times 128 \div 8 = 491.52$ MHz. We then fed those data streams to a four-branch linear convolution (4B-LC) module.

3. 4B-LC module: This module contains four independent complex FIR filter chains, each implemented with the same partially parallel architecture. For example, branch 1 is illustrated in Figure 3.

Branch 1 includes four subprocessing stages isolated by registers for better timing: a FIR coefficients RAM (cRAM) sequential-write and parallel-read stage; a complex multiplication stage; a complex addition stage; a complex addition

stage; and a segmental accumulation-and-downsample stage.

In order to minimize the I/O numbers for the core, our first stage involved creating a sequential write operation to load the coefficients from storage to the FIR cRAM in a TDM manner (each cRAM contains $16 = 128/8$ IQ coefficients). We designed a parallel read operation to feed the FIR coefficients to the CM core simultaneously.

In the complex multiplication stage, in order to minimize the DSP48 utilization, we chose the efficient, fully pipelined three-multiplier architecture to perform complex multiplication at a cost of six time cycles of latency.

Next, the complex addition stage aggregates the outputs of the CMs into a single stream. Finally, the segmental accumulation-and-downsample stage accumulates the temporary substreams for 16

time cycles to derive the corresponding linear convolution results of a 128-tap FIR filter, and to downsample the high-speed streams back to match the data-sampling rate of the system—here, 30.72 MHz.

DESIGN VERIFICATION

We performed the IP verification in two steps. First, we compared the outputs of the FPGA_HLP_core with the referenced double-precision multibranch FIR core in Simulink. We found we had achieved a relative amplitude error of less than 0.04 percent for a 16-bit-resolution version. A wider data width will provide better performance at the cost of more resources.

After verifying the function, it was time to validate the silicon performance. So our second step was to synthesize and implement the created IP in the Vivado Design Suite 2015.1 targeting the FPGA fabric of the Zynq®-7000

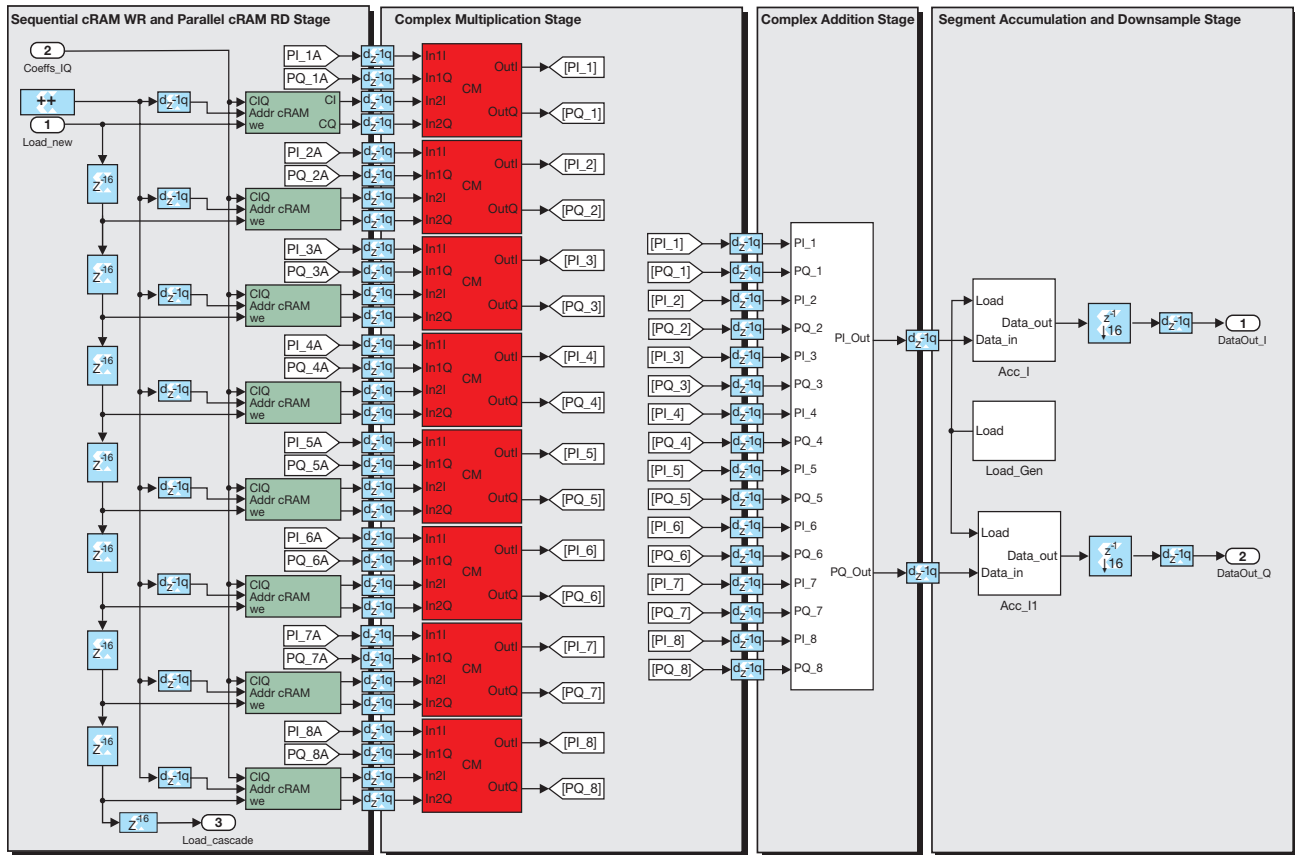


Figure 3 – Partial parallel complex FIR module

All Programmable SoC (equivalent to a Kintex® xc7k325tffg900-2). With full hierarchy in the tools' synthesize and default implementation settings, it was easy to achieve the required timing at a 491.52-MHz internal processing clock rate, since we created a fully pipelined design with clear registered hierarchies.

SCALABILITY ILLUSTRATION

The HLP IP we designed can be easily used to create a larger massive-MIMO precoding core. Table 2 presents selected application scenarios, with key resource utilizations. You will need an extra aggregation stage to deliver the final precoding results.

For example, as shown in Figure 4, it's easy to build a 4 x 4 precoding core by plugging in four HLP cores and one extra pipelined data aggregation stage.

EFFICIENT AND SCALABLE

We have illustrated how to quickly build an efficient and scalable DSP linear convolution application in the form of a massive-MIMO precoding core with Xilinx System Generator and Vivado design tools. You could expand this core to support longer-tap FIR applications by either using more sequential stages in the partially parallel architecture, or by reasonably increasing the processing clock rate to do a faster job. For the latter case, it would be helpful to identify the bottleneck and critical path of the target devices regarding the actual implementation architecture. Then, co-optimization of hardware and algorithms would be a good approach to tune the system performance, such as development of a more compact precoding algorithm regarding hardware utilization.

Initially, we focused on a precoding solution with the lowest latency. For our next step, we are going to explore an alternative solution for better resource utilization and power consumption. For more information, please contact the author by e-mail: lei.guan@ieee.org.

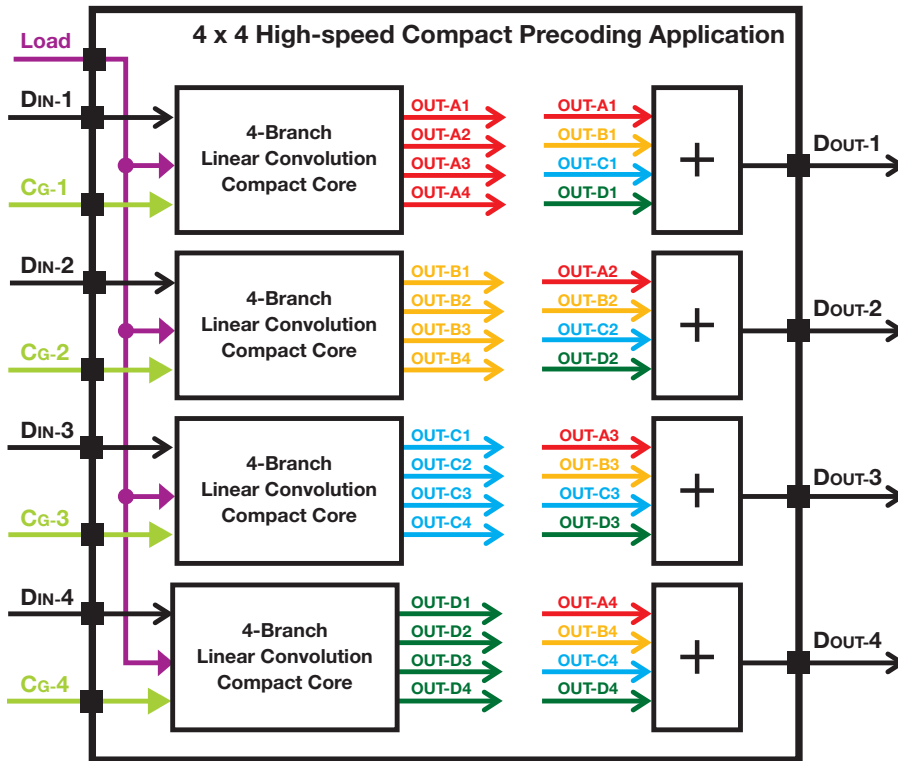


Figure 4 – Example of scaled precoding application based on the 4B-LC core

Massive MIMO Configuration (N _{RX} x N _{TX})	Number of HLP Cores	Complex Multiplier	Distributed Memory (Kb)	Block Memory (RAM18E)
4 x 4	4	128	16	4
4 x 8	8	256	32	8
4 x 16	16	512	64	16
4 x 32	32	1024	128	32

Table 2 – Example resource utilization in different application scenarios based on suggested HLP core

Drone Platform Soars with the Zynq SoC

Aerotenna achieved its first ArduPilot-compatible UAV flight by leveraging the processing power and I/O capabilities of Xilinx's Zynq SoC device.

by Zongbo Wang

Founder
Aerotenna
zongbo@aerotenna.com

Bruno Camps

Founder
Aerotenna

Yan Li

Embedded Design Engineer
Aerotenna

Joshua Fraser


Embedded Design Engineer
Aerotenna

Lianying Ji

CTO
Muniu Technology

Jie Zhang

Software Engineer
Muniu Technology



The unmanned aerial vehicle (UAV) and drone industry is quickly growing and reaching new commercial and consumer markets. The horizon of what is possible with UAVs continues to push forward into creative new applications like 3D modeling, military aid and delivery services.

The problem is that the applications are becoming increasingly more complex, requiring more and more processing power and I/O (input/output) interfaces, while the UAV platforms available are not improving at the same pace. The limits on the capabilities of most UAV platforms are being reached as the software and hardware needed for flight continue to advance.

Our team at Aerotenna has successfully flown a UAV with a board we built based on the Xilinx® Zynq®-7000 All Programmable SoC. This flight marks the beginning of our plans to release microwave sensing products that are computationally heavy. We turned to the Zynq SoC because the needed processing power was unavailable with other solutions of its class. With this new platform (Figure 1), we plan to improve the unmanned flying experience by deploying our microwave-based collision avoidance systems.

LIMITS OF TODAY'S UAV TECHNOLOGY

The main push of the UAV industry has been to make flying as affordable as possible, simplifying and stripping all unnecessary capabilities. This is a good thing if you are just looking to buy a product that does only what you want in the simplest way. But for developers like us, who seek to explore new, complex applications, it was necessary to branch out and build our own UAV platform capable of providing the processing speeds to power our ideas.

Another big limitation of today's standard UAV platforms is the lack of input-and-output connections to the processor. Thus, the flight control system easily reaches the maximum capacity of the processor and the I/O capabilities, not leaving much room for new sensors, and new applications.

Most of the I/Os included in the standard processor boards are already used up by the various components needed for flight. These functions include the inertial measurement sensors for quantifying aircraft orientation, the barometer and altimeter for determining the altitude and an RC receiver for decoding the user's input. Any I/O that's left over for adding extra features does not offer much in the way of options, generally confined to the

most popular demands such as a camera or GPS for navigation. A single platform that is compatible with a very wide range of sensors and external interfaces currently does not exist on the market.

Here at Aerotenna, we believed the way to overcome these limitations was to create a new board design from scratch. We have been working to perfect a new UAV platform that will excel in the areas where the other platforms fail. We used the Zynq SoC device provided by Xilinx to achieve this goal. Its superior design will offer the greatly increased processor speed and I/O capabilities needed for the next generation of UAVs.

WHY THE ZYNQ SOC?

We chose the All Programmable Zynq SoC as the foundation on which to build

our powerful platform. The dual-core ARM® Cortex™-A9 APU within the Zynq SoC chip allows for unparalleled processor speed. Nothing on the market for affordable UAV platform solutions compares with the Zynq SoC in chip structure, multiprocessor capabilities and I/O access speeds. Thus, the Zynq SoC is the perfect candidate for the next-generation platform.

Most of the flight control platforms currently in the market are based on a microcontroller unit (MCU). That architecture limits the potential for sensor fusion due to the limited processing power and I/O extension capabilities.

The advantage of the Zynq SoC is clear in both processing power and I/O capability: the combination of dual ARM cores plus FPGA logic enables a hardware/software co-design approach

that places some of the timing-critical processing tasks in the programmable logic. The I/O peripherals and memory interfaces are more versatile than the ones provided by MCU-based platforms.

Another reason we chose the Zynq SoC is because it is able to easily handle the complexities of flight control programs, which can be enormous and require very fast CPUs. And the device has plenty of power left over, which leaves a lot of room for expansion in the flight control programs.

There are many types of flight software programs, and they all differ in behavior and complexity. The flight control software we have decided to use is called ArduPilot, provided by APM (autopilot machine) on Dronecode. It is more complex than most, but provides a lot of functionality not included



Figure 1 – Conceptual drone that Aerotenna has been developing has advanced sensing technology and a powerful processing platform powered by the Zynq SoC.

The ArduPilot flight control system continues to grow in complexity as the open-source community adds more and more to the APM project.

in simpler programs, such as waypoint navigation and multiple flight modes to cater to the user's specific application.

WHAT IS ARDUPILOT?

ArduPilot is an open-source autopilot software program built for UAVs. It is kept up to date and improved upon by a large community of developers and enthusiasts. The APM project, which started out as much simpler software built for the open-source Arduino micro-processor, has grown much larger and more complex and is now compatible with many UAV platforms. Currently, the program contains more than 700,000 lines of code, and is a beautifully intricate flight control system.

The code is divided into two main parts: the high-level layer and the hardware abstraction layer (HAL). The high-level layer is responsible for scheduling tasks and making decisions based on incoming data. The HAL is the low-level code that accesses the hardware's memory. This separation of code structure allows the whole system to be ported to other platforms by changing only the HAL for the platform-specific memory access. And the upper-level code simply retrieves the data from the HAL the same way across all platforms.

The ArduPilot flight control system continues to grow in complexity as the open-source community adds more and more to the APM project. As a result, the industry is reaching the hardware's limit, waiting for the next platform on which to continue growing.

DRASTIC IMPROVEMENT

The initial efforts of porting ArduPilot to the Zynq SoC (led by John Williams in the drones-discuss Google group) in 2014 paved the way for porting the APM to the same Xilinx platform. Dr. Williams noticed the Zynq SoC's potential for offering custom I/O and real-time image processing as the beginning of an amazing new world for UAVs. In an interesting twist, Williams was the founder of PetaLogix, which created the original PetaLinux tools. Xilinx acquired the company in 2012.

The Aerotenna team continued these design efforts in both hardware and firmware, and accomplished the first Zynq SoC-powered ArduPilot flight in October of 2015. Our custom board runs the ArduPilot flight control software on the PetaLinux operating system. This impressive feat marks a drastic improvement in UAV technology and capability.

The dual ARM core within the Zynq SoC puts our flight control solution far ahead of any other UAV solution of its class in processing power and I/O capabilities. This leap forward will open the door to many new UAV applications that require greater computing power. We want to make sure to provide something with plenty of hardware interfaces built for the enthusiast as well as for the developer. Atop a Linux operating system, the UAV platform has much more flexibility to accommodate a very wide range of applications because of Linux's programmability and versatility. As one of the most powerful user-programmable operating systems, Linux al-

lows our team to customize the system exactly to our needs.

We accomplished the flight test on the commercial off-the-shelf DJI F550 airframe and plan to test our Zynq SoC-based flight controller on more airframes. We will soon release this platform as part of the Octagonal Pilot On Chip (OcPoC) platform.

AMBITIOUS ENDEAVOR

To start completely from scratch to make a customized flight control platform is an ambitious endeavor that takes the perfect team of engineers, and a lot of learning, to accomplish. Starting from nothing, there are a lot of decisions that need to be made about the system. In order to run a flight control program, an operating system must be used. A real-time operating system (RTOS) processes data right as it comes in, resulting in a negligible buffering delay. As a result, an RTOS is great for running time-sensitive tasks like flight control. The disadvantage is that it is difficult to interface this kind of system with ArduPilot because some of the data processing tasks would need to be reimplemented in the operating system itself.

That's why we opted instead for a Linux operating system, which is not real-time, but is much easier to implement in a hardware/software co-design, maximizing the versatility of the system. Xilinx provides a powerful embedded Linux operating system called PetaLinux that is compatible with the Zynq SoC and other Xilinx devices.

Since multirotor vehicles are naturally unstable, measuring the frame's inertia and altitude is critical for stability.

The road map to getting this system up and running seemed complicated, and we had to overcome many difficult challenges. The process began by developing the system design with FPGA development software, and writing and creating new intellectual property (IP) for driver interfaces. The cores are embedded processes at the hardware level that process data at very fast speeds. Then, a PetaLinux operating system must be deployed using the custom system design. Finally, we compiled and modified the ArduPilot system to be able to run on PetaLinux and the new platform.

We tackled the problem in stages to achieve a proof of concept. Our team

began working to first receive and detect an RC signal, and then to power a single motor. After finally demonstrating this proof of concept, we proceeded to expand the interface between the OcPoC and sensors that ArduPilot relies on to function. Writing our own device drivers from scratch was a major challenge. The most critical sensors for achieving a successful flight include accelerometers, gyroscopes and a barometer. Since multirotor vehicles are naturally unstable, measuring the frame's inertia and altitude is critical for stability. All these had to be configured in our FPGA hardware design with the correct communication protocol, and eventually included into our PetaLinux operating system.

The ArduPilot code base consists of more than 700,000 lines of code, so one big task was to get the system operational on a brand-new platform. With no easy interface to calibrate the inertial sensors, motors and RC controller (normally done inside a nice graphical user interface for other platforms), we had to manually calibrate the whole system by tweaking the hundreds of stored parameter values. Calibration is necessary, since each hardware component is slightly different and will produce slightly different outputs. So you must define the maximum and minimum values produced by each component. The process finally ended in a smooth and sustained flight for the Aerotenna team.



Figure 2 – The OcPoC system will be the first commercially available flight control platform powered by the Zynq SoC chip.




Figure 3 – OcPoC is designed as a ready-to-fly box, integrated with IMU sensors and GPS, with versatile I/Os to interface with external devices.

INTRODUCING THE OCPOC

The OcPoC project (Figure 2) is Aerotenna’s UAV flight control platform. With it we plan to meet the needs of the drone community with greatly enhanced processing capability, I/O expansion and much more flexibility in programming than other solutions. Though using the Zynq SoC to power our system may seem to be overkill for running the current ArduPilot release, we foresee this industry continuing to expand and want to provide the potential that will soon be utilized.

This architecture paves the way for developers to create and design with all the processing power they need. With this new platform, we plan to introduce new applications of microwave technology in imaging, mapping and proximity detection compatible with the OcPoC. Our system will also be able to perform onboard data capture and analysis through the processing capabilities of the Zynq SoC chip.

Our platform will provide integrated IMU data acquisition, which will create a

ready-to-fly “box” without any additional sensor setup (Figure 3). It will also provide integrated navigation interfaces for any type of wireless navigation control. What takes our platform a step ahead is the ability for any external sensor data to be directed through the Zynq SoC to perform high-speed data processing simultaneous to the ArduPilot program. This is not possible with MCU-based platforms.

The Zynq SoC’s extra processing power can also handle more complicated flight control systems to more finely tune the performance of the UAV. This includes expanding the I/O capabilities to a much wider range of external interfaces and sensor options, like real-time video streaming, microwave proximity sensors and Bluetooth.

Our hope is that by making our platform easy to test and develop new ideas on, many other companies and individuals will contribute novel creations to the drone industry, unhindered by the processing limitations of today’s available hardware. ●●



XEM7350 KINTEX⁷

- ✓ Available K70T, K160T, and K410T
- ✓ FMC HPC Carrier
- ✓ USB 3.0 transfers over 340 MB/s
- ✓ 8 Gigabit Transceiver Lanes
- ✓ 512 MiB DDR3
- ✓ **Excellent JESD204B Host**

FrontPanel SDK



Software API and Driver



Device Firmware



FPGA IP

www.opalkelly.com

How to Extend the Operating Temperature of FPGAs

by **Arnaud Darmont**

Founder, CEO and CTO

Aphesa

arnaud.darmont@aphesa.com



Some applications require the electronics to run hotter than the maximum operating junction temperature specified for the device. Case in point: an oil well camera design.

The lifetime of any electronic device depends on its operating temperature. At elevated temperatures, devices age faster and their lifetime is reduced. Yet some applications require the electronics to operate beyond the specified maximum operating junction temperature of the device. An example from the oil-and-gas industry illustrates the problem and ways to solve it.

A customer asked our team at Aphesa to design a high-temperature camera that will operate inside oil wells (Figure 1). The device required a rather large FPGA and had temperature requirements up to at least 125 degrees Celsius—the system's operating temperature. As a consultancy that develops custom cameras and custom electronics including FPGA code and embedded software, we have experience with high-temperature operation. But for this project, we had to go the extra mile.

The product is a down-hole, dual, color camera designed for use in oil well inspection (Figure 2). It performs embedded image processing, color reconstruction and communication. The system has memory, LED drivers and a high-dynamic-range (HDR) imaging capability. For this project we chose to use the XA6SLX45 device from Xilinx® (Spartan®-6 LX45 automotive) because of its wide temperature range, robustness, small package, large embedded memory and large cell count.

This project was quite challenging also lots of fun to do. Let's examine how we put it all together, starting with a re-

view of some concepts of temperature, including junction temperature, thermal resistance and other phenomena. We will look into the causes of temperature increase in a device and will list our solutions to prevent it. We will also address the possible hot-spot issues and propose solutions to them.

In this particular project, the use of thermoelectric cooling was limited and we had to find other solutions.

TEMPERATURE VARIATION

Electronic devices are usually specified with their maximum junction temperature. Unfortunately, what the system designers are concerned with is the ambient temperature. The difference between the ambient and the junction temperatures will depend on the capability of the package to transfer heat and the cooling system to dissipate this heat outside of the system's enclosure.

Thermal resistance is a heat property and a measurement of how much a given material resists heat flow. Because of thermal resistance, temperatures inside and outside of a component through which heat flows will differ, in just the same way that the voltage on each side of a resistor is not the same because of current flow. For a temperature difference of 20 degrees between the inside and the outside of a body, a device with 125 degrees of maximum junction temperature can operate in an environment of up to 105 degrees. The thermal resistance is expressed in degrees per watt; for 1 W of heat to dissipate, the temperature difference between the inside and the outside will be equal to the

Thermal resistance is a heat property and a measurement of how much a given material resists heat flow.

thermal resistance. This relationship is formalized in the Figure 3 equation.

The energy dissipated as heat depends on the device, the circuit, the clock frequency and the code running in the device. The temperature difference between the inside of the device (the junction temperature) and its environment (the ambient temperature) therefore depends on the device, the code and the schematic.

USUAL COOLING SOLUTIONS

In most designs where cooling is required, designers use either passive cooling (a heat sink that helps dissipate heat into the air by increasing the surface area in contact with air) or active cooling. Active-cooling solutions typically force an airflow in order to help renew the cold air that absorbs heat above the

device. The capability of air to absorb heat depends on the temperature difference between the air and the device, as well as on the pressure of the air. Other solutions include liquid cooling, where the liquid, usually water, replaces air for more efficiency. The capability of a mass of air or fluid to absorb heat is given by the heat absorption equation in Figure 4.

The final approach that designers often use is thermoelectric cooling, where the Peltier effect—a temperature difference created by applying a voltage between two electrodes connected to a sample of semiconductor material—is used to cool one side of a cooling plate while heating the other. Although this phenomenon helps to move heat away from the device to be cooled, Peltier cooling has one big disadvantage: It requires significant external power.

In our case, airflow was not a solution because the quantity of air in the enclosure is limited and the air quickly equalizes in temperature. Water cooling was not possible either, because of the long distance between a water source and the tool. So for us, the Peltier effect was the only cooling option. Since the ambient temperature is fixed (we can't heat up the large fluid quantity as in Figure 3's formula for a very large value of mass), the thermoelectric cooler will actually reduce the temperature of the electronics. Unfortunately, as a high current is required for the cooling devices and a very long conductor is used to connect from the surface to the tool, only limited current is actually available for cooling and only a small temperature drop is reached.




Figure 1 – The design of a high-temperature camera to operate inside oil wells (as shown at left) required extending the operating temperature beyond the qualified maximum temperature of the device. A close-up of the camera is at right.




Figure 2 – The high-temperature camera and high-temperature processing board are equipped with Xilinx Spartan-6 FPGAs.

Moreover, our device is a camera and image quality decreases exponentially with temperature. Therefore, we had to optimize our cooling strategies to cool down as much as possible the image sensors and not the other devices such as FPGAs, memories, LED drivers or power supply circuits.

Since cooling of the FPGA was almost impossible due to the limited choice of Peltier cooling on the image sensors only, our only option was to reduce the peak temperature inside the FPGA.

CAUSE OF HOT SPOTS AND RISING TEMPERATURES

There are three sources of power consumption in a digital device: dynamic, static and Joule effect. Dynamic power consumption is the power required to

charge and discharge route capacitance when a gate toggles; it is proportional to the clock rate and the total capacitance. Static power consumption is a function of the device’s type, core voltage and technology. The power can be consumed by the core or by the I/Os.

When heat is generated at a point in space, it will spread around and cause the surrounding area to heat up. If the surrounding area is not a heat source, then heat will diffuse and the temperature rise will be limited. The temperature will eventually become uniform in the whole device if one waits long enough. If the surrounding area is made of other heat sources, then there will be a net increase in temperature as each source will bring heat to the other.

If many heat sources are grouped in a small area, then the temperature of

that area will heat up more than in the other places, causing a hot spot.

As the device’s junction temperature is limited, it actually means that the temperature of the hottest spot should not exceed the maximum junction temperature. Knowing the power consumption of a device and the temperature of the package, all we can estimate is the average temperature of the junction.

The last source of heat is related to the current flow in conductors resulting in Joule effect, $P = \frac{0L}{S} I^2$.

WHAT HAPPENS IF THE MAXIMUM TEMPERATURE IS EXCEEDED?

As operating temperatures rise, the lifetime of a device decreases and the part will age faster. Some aging processes, such as electromigration and corrosion, only take place at higher temperatures.

Electromigration occurs in the presence of humidity and an electric field when atoms of a conductor move in ionic form from their original location to reposition themselves somewhere else, leaving a gap behind. The gap will reduce the effective width of the conductor at that location and will cause an increase of the local electric field, hence inducing more electromigration. This chain reaction will result

$$T_j = T_a + (R_{th, package} + R_{th, ambient}) \cdot P$$

Figure 3 – Relationship between ambient and junction temperatures where T_j is the junction temperature, T_a is the ambient temperature, $R_{th, package}$ is the thermal resistance between the junction and the outside surface of the package, $R_{th, ambient}$ is the thermal resistance between the outside surface of the package and the ambient air (it is zero if there is no heat sink or air flow) and P is the power dissipated by the device.

There are three sources of power consumption in a digital device: dynamic, static and Joule effect.

$$Q = mc\Delta T$$

Figure 4 – Heat absorption equation, where Q is the maximum heat that can be absorbed, m is the mass of the heat-absorbing material, c is a constant representative of the material and ΔT is the difference between the ambient temperature of the heat-absorbing material at the beginning and the final temperature. This formula is valid for a nonrenewing heat-absorbing material and a net amount of heat to be absorbed, an unrealistic situation but which already shows that pressure (mass), type of material (c) and outside temperature play a role in the efficiency of cooling.

in a crack (open circuit) where the atoms are removed or a short where the atoms relocate (dendrites). A few layers of water molecules may be sufficient to initiate the ionization process of the metal and trigger electromigration. The phenomenon drastically increases with temperature.

Corrosion, as one can see in the rusting of iron, involves moisture and a hostile gas. Semiconductor materials are encapsulated in their protective packages. Such packages are usually highly absorbent to moisture but are made of materials that do not easily produce corrosive ionic solutions. This corrosion mostly affects the lead frame and the bond wires. The most important hostile materials are the phosphorus present in the silicon's passivation layer and some contaminants left from the semiconductor-manufacturing process or the packaging process. Contact with human skin and other chemicals during transport, soldering and assembly are other potential sources of contamination by hostile atoms.

When dissimilar materials are connected together, the less-noble material will corrode relative to the more-noble one (galvanic corrosion). This type of corrosion is another cause of degradation over time.

When the junction temperature is ex-

ceeded, the lifetime of the device is no longer guaranteed and can be drastically reduced. If the temperature continues to rise, the device may fail immediately.

The performance of a device also depends on speed. Devices tend to be slower at elevated temperatures and therefore their maximum clock rate is reduced.

The reason why the Spartan-6 XA (automotive) FPGA is limited to 125 degrees is because of minimum lifetime requirements (reliability concerns) and guaranteed clock frequency capability (performance requirements). Other reasons include the leakage of RAM cells and as a consequence, bit errors.

SEVERAL SOLUTIONS

To overcome these varied challenges in our design for an oil well camera, we implemented several solutions.

One of the most important decisions was choosing the right size device. Larger devices have more static power consumption but allow for more spread of heat into the device, avoiding hot spots.

Devices qualified for automotive use have a long lifetime at elevated temperature and it is acceptable to have a shorter lifetime in industrial applications. We have evaluated the code in LX25 and LX45 devices of the XA (au-

tomotive) series and measured the total power consumption and the temperature of the device's body. Sometimes it is acceptable to have a higher average device temperature if the peak temperature is less. We also evaluated the lifetime in accelerated aging tests.

Our next design choice was to place a limit on the device usage. In order to reduce the heat dissipated by the device, we avoided using all possible logic cells and memory. The unused parts of the device consume static power but not dynamic power.

We also performed clock gating. As dynamic power depends on the clock rate, we can use clock gating to cancel the dynamic power consumption of the blocks that are not in use. If the clock tree does not toggle, the power consumption in that part of the device is reduced.

We also kept the number of I/Os we were using to a minimum. This in turn lowered the power consumption of the I/O banks.

Then, by using some I/Os as virtual grounds, we reduced the distance traveled by the current inside the device and therefore reduced the Joule effect in power routing. The virtual grounds also help with conducting heat into the ground plane.

Since we did not want to use all I/Os and all logic cells, we chose to spread the design over two FPGAs (Figure 5). This means that the heat is dissipated at two separate locations.

We also used multiple ground planes. This technique helps conduct the heat from the warmer areas to the cooler areas and also provides additional heat capacitance. It's important to design the thermal planes for board reliability against delamination during temperature cycles.




Figure 5 – To avoid using all the I/Os and logic cells (top), the design uses two Spartan-6 FPGAs rather than one. This means the heat is dissipated at two separate locations.

Another important step was to optimize our code to reduce the clock rate. Reducing the clock rate reduces the power consumption but also allows the device to run at a higher temperature. As an example, we evaluated the trade-off between slow parallel design and fast pipelined design.

To improve the design, we ensured the components were dried before their final assembly and applied a protective coating to repel moisture.

At elevated temperatures, devices will age faster. A product qualification can be used to measure the actual lifetime of the device vs. temperature.

We also employed a burn-in process in production to pre-age the device and reject the parts that seemed to age faster than others (early fails), therefore only keeping the best.

Also important to our design process was the use of cyclic redundancy checking (CRC) and other types of error detection and correction. We used these techniques in various plac-

es in the design in order to recover, or at least detect, bit errors in memory cells or communication. State machines also recover if they end up in an unused state.

We found that using the Xilinx Power Estimator (XPE) was a good first step in executing our design. The Vivado® Design Suite has power estimation tools for designs on newer devices. But the measurement of power consumption on the real devices and comparisons between versions of the code proved to be the best and most accurate practice.

NO THERMOELECTRIC COOLING

Combining all of the above techniques, we came away with a camera that is able to work at 125 degrees of ambient temperature with SDRAM management, communication buses and image processing, even though it is by specification limited to 125 degrees of junction temperature. Moreover, we managed to reach 125 degrees without thermoelectric cooling.

Everything FPGA.

1. MARS ZX2
Zynq-7020 SoC Module

- Xilinx Zynq-7010/7020 SoC FPGA
- Up to 1 GB DDR3L SDRAM
- 64 MB quad SPI flash
- USB 2.0
- Gigabit Ethernet
- Up to 85,120 LUT4-eq
- 108 user I/Os
- 3.3 V single supply
- 67.6 x 30 mm SO-DIMM

from \$127

2. MERCURY ZX5
Zynq™-7015/30 SoC Module

- Xilinx® Zynq-7015/30 SoC
- 1 GB DDR3L SDRAM
- 64 MB quad SPI flash
- PCIe® 2.0 x4 endpoint
- 4 x 6.25/6.6 Gbps MGT
- USB 2.0 Device
- Gigabit Ethernet
- Up to 125,000 LUT4-eq
- 178 user I/Os
- 5-15 V single supply
- 56 x 54 mm

3. MERCURY ZX1
Zynq-7030/35/45 SoC Module

- Xilinx Zynq-7030/35/45 SoC
- 1 GB DDR3L SDRAM
- 64 MB quad SPI flash
- PCIe 2.0 x8 endpoint¹
- 8 x 6.6/10.3125 Gbps MGT²
- USB 2.0 Device
- Gigabit & Dual Fast Ethernet
- Up to 350,000 LUT4-eq
- 174 user I/Os
- 5-15 V single supply
- 64 x 54 mm

1, 2: Zynq-7030 has 4 MGTs/PCIe lanes.

4. FPGA MANAGER
IP Solution

→ USB 3.0

→ PCIe® Gen2

→ Gigabit Ethernet

Streaming,
made simple.

One tool for all FPGA communications. Stream data from FPGA to host over USB 3.0, PCIe, or Gigabit Ethernet – all with one simple API.

Design Center • FPGA Modules
Base Boards • IP Cores

ENCLUSTRA
FPGA SOLUTIONS

How to Digitize Hundreds of Signals with a Single Xilinx FPGA

by William D. Richard
Associate Professor
Washington University in St. Louis
wdr@wustl.edu

Mitchell Manar
Student
Washington University in St. Louis

Jeremy Tang
Student
Washington University in St. Louis

Using just one resistor and one capacitor per input channel, it's possible to digitize high-frequency analog input signals.



Using the low-voltage differential signaling (LVDS) inputs on a modern Xilinx® FPGA, it is possible to digitize an analog input signal with nothing but one resistor and one capacitor. Since hundreds of LVDS inputs reside on a current-generation Xilinx device, it is theoretically possible to digitize hundreds of analog signals with a single FPGA.

Our team recently explored one corner of the possible design space by digitizing a band-limited input signal with a 3.75-MHz center frequency with 5 bits of resolution while investigating options for digitizing the signals from a 128-element linear ultrasound array transducer. Let's take a look at the details of that demonstration project.

In 2009, Xilinx introduced a LogiCORE™ soft IP core that, along with an external comparator, one resistor and one capacitor, implements an analog-to-digital converter (ADC) capable of digitizing inputs with frequencies up to 1.205 kHz [1].

Using an FPGA's LVDS inputs instead of an external comparator, in conjunction with a delta modulator ADC architecture, it is possible to digitize much higher-frequency analog input signals with just one resistor and one capacitor.

ADC TOPOLOGY AND EXPERIMENTAL PLATFORM

The block diagram of a one-channel delta modulator ADC [2] implemented using the LVDS inputs on a Xilinx FPGA is shown in Figure 1. Here, the analog input drives the noninverting LVDS_33 buffer input, and the input signal range is essentially 0 to 3.3 volts. The output of the LDVS_33 buffer is sampled at a clock frequency much higher than the input analog signal frequency and fed back through an LVCMOS33 output

Simplicity and low component count make this approach attractive. And since the LVDS_33 input buffer has a relatively high input impedance, in many applications the sensor output can be directly connected to the FPGA input with no preamplifier or buffer.

buffer and an external, first-order RC filter to the inverting LVDS_33 buffer input. With just this circuitry, the feedback signal, given an appropriate selection of clock frequency (F), resistance (R) and capacitance (C), will track the input analog signal.

As an example, Figure 2 shows an input signal in yellow (channel 1) and the feedback signal in blue (channel 2) for $F = 240$ MHz, $R = 2$ K and $C = 47$ pF. The input signal shown was produced by an Agilent 33250A function generator using its 200-MHz, 12-bit, arbitrary output function capability. The Fourier transform of the input signal as

computed by the Tektronix DPO 3054 oscilloscope we used is shown in red (channel M). At these frequencies, the input capacitance of the oscilloscope probe (as well as grounding issues) did degrade the integrity of the feedback signal shown in the oscilloscope trace, but Figure 2 does illustrate operation of the circuit.

We defined the band-limited input signal shown in Figure 2 by applying a Blackman-Nuttall window to a $1-V_{pp}$ 3.75-MHz sine wave. While the noise floor associated with the theoretical windowed signal is almost 100 dB below the amplitude associated with the

center frequency, the 200-MHz sample frequency and 12-bit resolution of the Agilent 33250A function generator result in a far-less-ideal demonstration signal. The output signals produced by many ultrasound transducers with center frequencies near 3.75 MHz are naturally band-limited, due to the mechanical properties of the transducers, and are therefore ideal signal sources for use with this approach.

We obtained the plot shown in Figure 2 using a Digilent Cmod S6 development module [3] with a Xilinx Spartan®-6 XC6SLX4 FPGA mounted on a small, custom printed-circuit board with eight

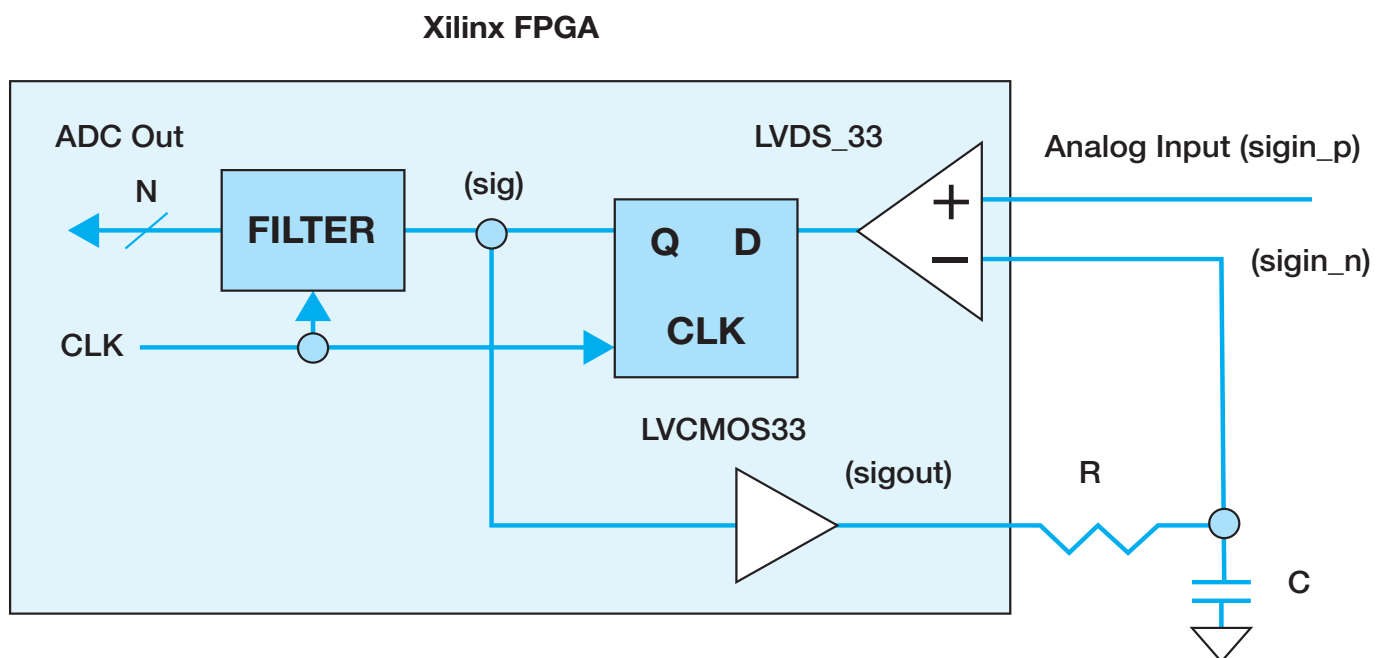


Figure 1 – A one-channel delta modulator ADC uses one external resistor and one external capacitor.

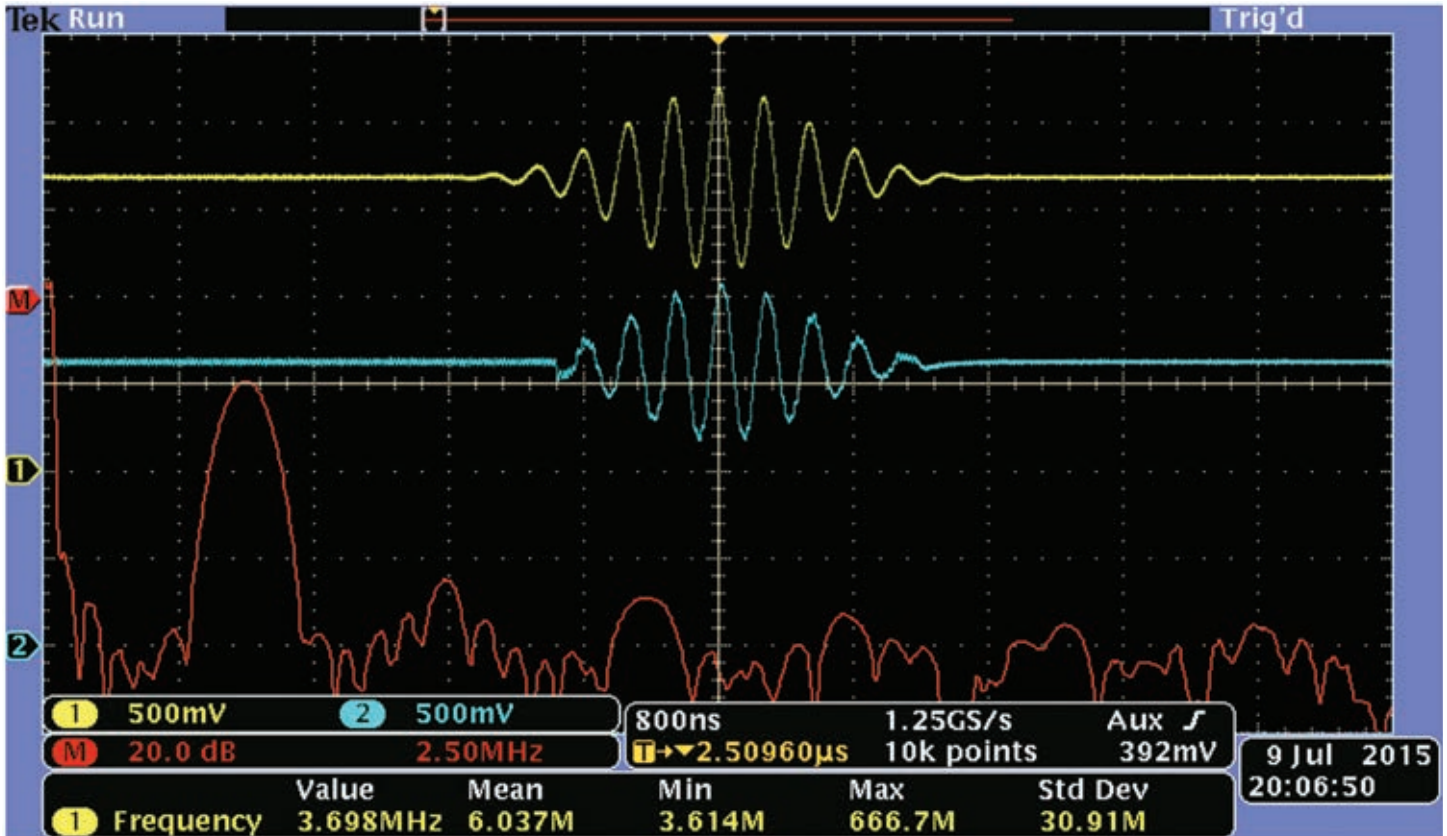


Figure 2 – This oscilloscope plot shows the 3.75-MHz input signal produced by the Agilent 33250A function generator in yellow (channel 1) and the feedback signal in blue (channel 2) for the case where $F = 240$ MHz, $R = 2$ K and $C = 47$ pF. The Fourier transform of the input signal as computed by the Tektronix DPO 3054 oscilloscope is shown in red (channel M).

R/C networks and input connectors, allowing the prototype system to digitize up to eight signals simultaneously. Each channel was parallel-terminated with 50 ohms to ground to properly terminate the coaxial cable from the signal generator. It is important to note that to achieve this performance, we set the drive strength of the LVCMOS33 buffers to 24 mA and the slew rate to FAST, as documented in the example VHDL source in Figure 5.

The custom prototype board also supported the use of an FTDI FT2232H USB 2.0 Mini-Module [4] that we used to transfer packetized serial bitstreams to a host PC for analysis. Figure 3 shows the magnitude of the Fourier transform of the bitstream the prototype board produced when fed the analog signal of Figure 2. Peaks associated with subharmonics of the

240-MHz sampling frequency are clearly visible, along with a peak at 3.75 MHz associated with the input signal.

LARGE NUMBER OF TAPS

By applying a bandpass finite impulse response (FIR) filter to the bitstream, it is possible to produce an N-bit binary representation of the analog input signal: the ADC output. Since the digital bitstream is at a much higher frequency than the analog input signal, however, you need to use FIR filters with a large number of taps. The data being filtered, however, only has values of zero (0) and one (1), so multipliers are not needed (only adders to add the FIR filter coefficients).

The ADC output shown in Figure 4 was produced on the host PC using an 801-tap bandpass filter centered at 3.75 MHz that we designed using

the free, online TFilter FIR filter design tool [5]. This filter had 36 dB or more of attenuation outside the 2.5- to 5-MHz passband and 0.58 dB of ripple between 3 and 4.5 MHz.

The ADC output signal shown in Figure 4 has a resolution of approximately 5 bits. This is ultimately a function of the oversampling rate, and you can achieve higher resolution with designs optimized for lower input frequencies.

The ADC output signal shown in Figure 4 is also severely oversampled at 240 MHz and can be decimated to reduce the ADC output bandwidth. In a hardware implementation of the bandpass filter and decimation blocks, it would be possible to only compute every 16th filter output value when decimating by a factor of 16 down to an effective sample rate of 15 MHz (three times faster than the highest frequency

| DFT | in dB vs. F in MHz




Figure 3 – This graph shows the Fourier transform of the bitstream produced by the configuration associated with Figure 2.

Bandpass Filter Results (Fc = 3.75 MHz)




Figure 4 – The ADC output was produced using an 801-tap bandpass FIR filter centered at 3.75 MHz.

in the band-limited input signal), reducing the hardware requirements.

Figure 5 shows the VHDL source used with the Digilent Cmod S6 development module to produce the feedback signal shown in Figure 2, along with the bitstream data associated with the Fourier transform of Figure 3. An LVDS_33 input buffer is instan-

tiated directly and connected to the analog input and feedback signals, **sigin_p** and **sigin_n**, respectively. The internal signal **sig** is driven by the output of the LVDS_33 buffer and sampled by the implied flip-flop to produce **sigout**. The signal **sigout** is the serial bitstream that is filtered to produce the N-bit ADC output. We used the free

Xilinx ISE® Webpack tools to implement the project [6].

Figure 5 shows the VHDL code and the portion of the UCF file associated with the circuitry of Figure 1.

LOW COMPONENT COUNT

The ADC architecture we have described has been inaccurately referred

VHDL SOURCE

```

LIBRARY IEEE ;
USE IEEE.STD_LOGIC_1164.ALL ;
LIBRARY UNISIM ;
USE UNISIM.VCOMPONENTS.ALL ;

ENTITY deltasigma IS
    PORT (clk          : IN  STD_LOGIC ;
          signin_p     : IN  STD_LOGIC ;
          signin_n     : IN  STD_LOGIC ;
          sigout        : OUT STD_LOGIC) ;
END deltasigma ;

ARCHITECTURE XCellExample OF deltasigma IS

    SIGNAL sig : STD_LOGIC ;

BEGIN

    myibufds:IBUFDS
    GENERIC MAP (DIFF_TERM    => FALSE,
                IBUF_LOW_PWR => FALSE,
                IOSTANDARD    => "DEFAULT")
    PORT MAP (O => sig,
             I => signin_p,
             IB => signin_n);

    mydeltasigma:PROCESS(clk)
    BEGIN
        IF (clk = '1' AND clk'EVENT) THEN
            sigout <= sig ;
        END IF ;
    END PROCESS mydeltasigma ;

END XCellExample ;

```

UCF FILE

```

NET "clk"      LOC = J1 | IOSTANDARD = LVCMOS33;
NET "signin_p" LOC = N12 | IOSTANDARD = LVDS_33;
NET "signin_n" LOC = P12 | IOSTANDARD = LVDS_33;

NET "sigout"   LOC = P7 | IOSTANDARD = LVCMOS33 |
    SLEW = FAST | DRIVE = 24;

```

Figure 5 – The VHDL source code and UCF file contents

to in several recent articles as a delta-sigma architecture [7]. But while true delta-sigma ADCs have advantages, the simplicity of this approach and low component count make it attractive for some applications. And since the LVDS_33 input buffer has a relatively high input impedance, in many applications the sensor output can be directly connected to the FPGA input without the need for a preamplifier or buffer. This can be very advantageous in many systems.

Another advantage of our approach is that superposition makes it possible to “mix” several serial bitstreams and apply a single filter to recover the output signal. In array-based ultrasound systems, for example, the serial bitstreams can be time-delayed to implement a focus algorithm, and then added in vector fashion, and a single filter used to recover the digitized, focused ultrasound vector.

Using an FIR filter to produce the ADC output is a straightforward, brute-force approach used here primarily for illustrative purposes. In most implementations, the ADC output will be produced using the traditional integrator/lowpass filter demodulator topology [2].

References

1. XPS Delta-Sigma Analog to Digital Converter (ADC) V1.01A, DS587, Dec. 2, 2009
2. R. Steele, *Delta Modulation Systems*, Pentech Press (London), 1975
3. Digilent Cmod S6 Reference Manual, Digilent Inc., Sept. 4, 2014
4. FT2232H Mini-Module Datasheet, V1.7, Future Technology Devices International Ltd., 2012
5. TFilter, The Free Online FIR Filter Design Tool, <http://t-filter.engineerjs.com/>
6. USE In-Depth Tutorial UG695 (V13.1), Xilinx, Inc., 2011
7. M. Bolatkale and L.J. Breems, *High-Speed and Wide-Bandwidth Delta-Sigma ADCs*, Springer, May 2014 Edition

FMC+ Standard Propels Embedded Design to New Levels

Updated FPGA Mezzanine Card specification promises unparalleled I/O density, backward compatibility.

by Jeremy Banks
Product Manager
Curtiss-Wright
Jeremy.banks@curtisswright.com

Jim Everett
Product Manager
Xilinx, Inc.
jeverett@xilinx.com

A

A new mezzanine card standard called FMC+, an important development for embedded computing designs using FPGAs and high-speed I/O, will extend the total number of gigabit transceivers (GTs) in a card from 10 to 32 and increase the maximum data rate from 10 to 28 Gbits per second while maintaining backward compatibility with the current FMC standard.

These capabilities mesh nicely with new devices such as those using the JESD204B serial interface standard, as well as 10G and 40G fiber optics and high-speed serial memory. FMC+ addresses the most challenging I/O requirements, offering developers the best of two worlds: the flexibility of a mezzanine card with the I/O density of a monolithic design.

The FMC+ specification has been developed and refined over the last year. The VITA 57.4 working group has approved the spec and will present it for ANSI balloting in early 2016. Let's take a closer look at this important new standard to see its implications for advanced embedded design.

THE MEZZANINE CARD ADVANTAGE

Mezzanine cards are an effective and widely used way to add specialized functions to an embedded system. Because they attach to a base or carrier card, rather than plugging directly into

a backplane, mezzanine cards can be easily changed. For system designers, this means both configuration flexibility and an easier path to technology upgrades. However, this flexibility usually comes at the cost of functionality due to either connectivity issues or the extra real estate needed to fit on the board.

For FPGAs, the primary open standard is ANSI/VITA 57.1, otherwise known as the FPGA Mezzanine Card (FMC) specification. A new version dubbed FMC+ (or, more formally, VITA 57.4) extends the capabilities of the current FMC standard with a major enhancement to gigabit serial interface functionality.

FMC+ addresses many of the drawbacks of mezzanine-based I/O, compared with monolithic solutions, simultaneously delivering both flexibility and performance. At the same time, the FMC+ standard stays true to the FMC history and its installed base by supporting backward compatibility.

The FMC standard defines a small-format mezzanine card, similar in width and height to the long-established XMCs or PMCs, but about half the length. This means FMCs have less component real estate than open-standard formats. However, FMCs do not need bus interfaces, such as PCI-X, which often take a considerable amount of board real estate. Instead, FMCs have direct I/O to the host FPGA, with simplified power supply requirements. This means that despite their size, FMCs could actually

have more I/O capacity than their XMC counterparts. As with the PMC and XMC specification, FMC and FMC+ define options for both air and conduction cooling, thereby serving both benign and rugged applications in commercial and defense markets.

The anatomy of the FMC specification is simple. The standard allows for up to 160 single-ended or 80 differential parallel I/O signals for high-pin-count (HPC) designs or half that number for low-pin-count (LPC) variants. Up to 10 full-duplex GT connections are specified. The GTs are useful for fiber optics or other serial interfaces. In addition, the FMC specification defines key clock signals. All of this I/O is optional, though most hosts now support the full connectivity.

The FMC specification also defines a mix of power inputs, though the primary power supply, defined by the mezzanine, is supplied by the host. This approach works by partially powering up the mezzanine such that the host can interrogate the FMC, which responds by defining a voltage range for the V_{ADJ} . Assuming the host can provide this range, then all should be well. Not having the primary regulation on the mezzanine saves space and reduces mezzanine power dissipation.

FMCs FOR ANALOG I/O

Designers can use FMCs for any function that you might want to connect to an FPGA, such as digital I/O, fiber optics,

control interfaces, memory or additional processing. But analog I/O is the most common use for FMC technology. The FMC specification affords a great deal of scope for fast, high-resolution I/O, but there are still trade-offs—especially with high-speed parts using parallel interfaces.

For example, Texas Instruments' ADC12D2000RF dual-channel, 2-Gsample/second (Gsp/s) 12-bit ADCs use a 1:4 multiplexed bus interface, so the bus speed is not too fast for the host FPGA. The digital data interface alone requires 96 signals (48 LVDS pairs). For a device of this class, FMC can support only one of these ADCs, even if there is sufficient space for more, because it is limited to 160 signals. Lower-resolution devices, even at higher speeds, such as those with 8-bit data paths, may allow more channels even with the increased requirements of the front-end analog coupling of the baluns or amplifiers, clocking and the like.

The FMC specification starts to run out of steam with analog interfaces delivering more than 8 bits of resolution at around 5 or 6 Gsp/s (throughputs of > 50 Gbps) using parallel interfaces. From a market perspective, leading FMCs based on channel density, speed and resolution are in the 25- to 50-Gbps throughput range. This functionality results from a trade-off between physical package sizes and available connectivity to the host FPGA.

In addition to the parallel connections, the FMC specification supports up to 10 full-duplex high-speed serial (GT) links.




Figure 1 – The effect of package shrink on FMC through JESD204B

Function	FMC	FMC+	FMC+ with HSPCe extension
Maximum # parallel I/Os	80 diff/160 single ended	80 diff/160 SE	80 diff/160 SE
Clocks	4	4	4
Maximum # GTs	10	24	32
GT clocks	2	6	8
Miscellaneous	JTAG, SYNC, power good, geographic address	JTAG, SYNC, power good, geographic address	JTAG, SYNC, power good, geographic address
Power supplies	V _{ADJ} * (4 pins), 3V3 (4 pins), 12V (2 pins), 3V3 Aux (1 pin)	V _{ADJ} * (4 pins), 3V3 (8 pins), 12V (4 pins), 3V3 Aux (1 pin)	V _{ADJ} (4 pins), 3V3 (8 pins), 12V (4 pins), 3V3 Aux (1 pin)

* V_{ADJ}: mezzanine defined for voltage level, but provided by host

Table 1 – Summary of FMC and FMC+ connectivity

These interfaces are useful for such functionality as fiber-optic I/O, Ethernet, emerging technologies like Hybrid Memory Cube (HMC) and the Bandwidth Engine, and newer-generation analog I/O devices that use the JESD204B interface.

ENTER JESD204B

Although the JESD204 serial-interface standard, currently at revision “B,” has been around for a while, only recently has it gained wider market penetration and become the serial interface of choice for newer generations of high-sampling data converters. This wide adoption has been stoked by the telecommunications industry’s thirst for ever-smaller, lower-power and lower-cost devices.

As mentioned earlier, a dual-channel 2-Gsps, 12-bit ADC with a parallel interface requires a large number of I/O signals. This requirement directly impacts the package size, in this case mandating a 292-pin package measuring roughly 27 x 27 mm (though newer-generation pin geometry could shrink the package size to something less than 20 x 20 mm).

A JESD204B-connected equivalent device can be provided in a 68-pin, 10 x 10-mm package—with reduced pow-

er. This dramatic reduction in package size marries well with evolving FPGAs, which are providing ever more GT links at higher and higher speeds. Figure 1 illustrates an example of package size and FMC/FMC+ board size.

Typical high-speed ADCs and DACs using the JESD204B interface have between one and eight GT links operating at 3 to 12 Gbps each, depending on the data throughput required based on sample rate, resolution and number of analog I/O channels.

The FMC specification defines a relatively small mezzanine card, but with the emergence of JESD204B devices there is room to fit more parts onto the available real estate. The maximum of 10 GT links defined by the FMC specification is a useful quantity; even this limited number of GT links provide 80 Gbps or more of throughput while using a fraction of the pins otherwise required for parallel I/O.

The emergence of serially connected I/O devices, not just those using JESD204B, does have drawbacks for some application segments in electronic warfare, such as digital radio frequency memory (DRFM). Serial interfaces invariably introduce additional latency

due to longer data pipelines. For DRFM applications, latency for data-in to data-out is a fundamental performance parameter. Although latency is likely to vary widely between serially connected devices, new generations of devices will push data through the pipelines faster and faster, with some promising the ability to tune the depth of the pipeline. It remains to be seen how much of an improvement is to be realized.

Some standard ADC devices sampling at >1 Gsps today have latency below 100 nanoseconds. Other applications can tolerate this latency, or do not care about it, including software-defined radio (SDR), radar warning receivers and other SIGINT segments. These applications gain large advantages by using a new generation of RF ADCs and DACs, a technology driven by the mass-market telecommunications infrastructure.

Outside of the FPGA community, newer DSP devices are also starting to adopt JESD204B. However, FPGAs are likely to remain the stronghold in taking full advantage of the capabilities of wideband analog I/O devices. That’s because FPGAs can deal with vast data volumes with better parallelization.




Figure 2 – FMC vs. FMC+ digitizer throughput capability

THE EVOLUTION OF FMC+

To move FMC to the next level, the VITA 57.4 working group has created a specification with an increased number of GT links operating at increased speed. FMC+ maintains full FMC backward compatibility by adding to the FMC connector’s outer columns for the additional signals and not changing any of the board profiles or mechanics.

The additional rows will be part of an enhanced connector that will minimize any impact on available real estate. The FMC+ specification increases the maximum number of available GT links from 10 to 24, with the option of adding another eight links, for a total of 32 full duplex. The additional links use a separate connector, referred to as an HSPCe (HSPC being the main connector). Table 1 summarizes FMC and FMC+ connectivity.

Multiple independent signal integrity teams characterized and validated the higher 28-Gbps data rate. The maximum full-duplex throughput can now exceed 900 Gbps in each direction, when the parallel interface is included. See Figure 2 for an outline of the net throughputs that can be expected for digitizer

solutions supporting the different capabilities of FMC and FMC+.

Designers can use the increased throughput enabled by FMC+ to take advantage of new devices that offer huge I/O bandwidth. There will still be trade-offs, such as how many devices can fit on the mezzanine’s real estate budget, but for a moderate number of channels the realizable throughput is a huge leap over today’s FMC specification.

NEXT-GENERATION ADCS AND DACS

In the next few years, it is reasonable to expect high-resolution ADCs and DACs to break through the 10-Gsps barrier to support very wideband communications with direct RF samplings for L-, S- and even C-band frequencies. Below 10 Gsps, converters are emerging with 12-, 14- and even 16-bit resolutions, with some supporting multiple channels. The majority of these devices will be using JESD204B (or a newer revision) signaling with 12-Gbps channels until newer generations inevitably boost this speed even further. These fast-moving advances are fueled by the telecommunications industry, but the

military community can take advantage of them to meet SWAP-C requirements.

OTHER ADVANTAGES AND USES OF FMC+

Although FMC+, like FMC, is likely to be dominated by ADC, DAC and transceiver products, the increased GT density provided by FPGAs makes it useful for other functions. Two functions of note are fiber optics and new serial memories.

As with JESD204B, there are requirements for faster, denser fiber optics. Those based on fiber-optic ribbon cables offer the smallest parts. Because the FMC+ footprint readily supports 24 full-duplex fiber-optic links, this application is likely where the higher speeds supported by FMC+ will first be realized. Bandwidths of 28 Gbps per fiber will take the throughputs quickly past 100G and 400G speeds on a single mezzanine. Optical throughput of 100G is emerging today on the current FMC format.

Another emerging area suitable for FMC+ is serial memory such as Hybrid Memory Cube and MoSys’ Bandwidth Engine. These novel devices represent an entirely new category of




Figure 3 – Xilinx's KCU114 Demonstration Board featuring FMC+

high-performance memory, delivering unprecedented system performance and bandwidth by utilizing GT connectivity. (*Xcell Journal issue 88* examines these new memory types.)

ALIVE AND KICKING


A new generation of the FMC specification has been introduced and is adapting to new technology driven by serial connected devices. Key players in the FMC industry have already begun adopting this specification. Figure 3 shows the first Xilinx® demonstration board featuring FMC+, the KCU114. The FMC standard, through its new incarnation FMC+, is alive and kicking and is prepared for the next generation of high-performance, FPGA-driven applications. 🌟

TRACE32®

Support for Xilinx Zynq® Ultrascale+™ and Zynq-7000 families

- ▶ Concurrent debugging and tracing of ARM Cortex-A53/-R5, -A9 and MicroBlaze™ soft processor core
- ▶ RTOS support, including Linux kernel and process debugging
- ▶ Run time analysis of functions and tasks, code coverage, system trace

LAUTERBACH
DEVELOPMENT TOOLS



www.lauterbach.com

New Vivado HLx Editions Deliver High Productivity

Xilinx is continually improving its products, IP and design tools as it strives to help designers work more effectively. Here, we report on the most current updates to Xilinx design tools including the Vivado® Design Suite, a revolutionary system- and IP-centric design environment built from the ground up to accelerate the design of Xilinx® All Programmable devices. For more information about the Vivado Design Suite, please visit www.xilinx.com/vivado.

Product updates offer significant enhancements and new features to the Xilinx design tools. Keeping your installation up to date is an easy way to ensure the best results for your design.

VIVADO DESIGN SUITE HLx EDITIONS

In order to enable ultrahigh productivity for designing All Programmable SoCs and FPGAs, Xilinx has launched the Vivado Design Suite HLx Editions with the Vivado Design Suite 2015.4 release. The HLx Editions include Vivado High-Level Synthesis (HLS), complete with C/C++ libraries; Vivado IP Integrator (IPI); LogicCORE™ IP subsystems; and the full Vivado implementation tool suite. When coupled with the new UltraFast™ High-Level Productivity Design Methodology Guide, the tools enable users to realize a 10x to 15x productivity gain over traditional approaches. This methodology enables:

- Rapid creation of connected platforms using IPI and connectivity IP subsystems
- Rapid development of differentiated logic using HLS and libraries based on C/C++/OpenCL

The Vivado Design Suite HLx Editions are available from the Xilinx Download Center at www.xilinx.com/download.

Public Access Support for 16nm UltraScale+ Family

After announcing the first customer shipment of the Zynq® UltraScale+™ MPSoC in September 2015, Xilinx has announced the industry's first publicly available tools for FinFET-based programmable devices, the UltraScale+ family, in the Vivado Design Suite 2015.4 release.

These devices are supported in the Vivado Design Suite HLx Editions, in embedded software development tools and in the Xilinx Power Estimator. Designers can now validate the UltraScale+ portfolio's 2x to 5x performance/watt improvement over 28nm offerings for their specific designs.

The announcement marks the industry's first publicly available tools for 16nm devices, enabling broad market adoption.

LEARN MORE

See the Vivado Design Suite 2015.4 Release Notes for more information.

QuickTake Video Tutorials

Vivado Design Suite QuickTake video tutorials are how-to videos that take a look inside the features of the Vivado Design Suite and UltraFast Design Methodology.

See all Quick Take Videos here www.xilinx.com/training/vivado.

Training

For instructor-led training on the Vivado Design Suite, UltraFast Design Methodology and more, visit www.xilinx.com/training.


Download Vivado Design Suite HLx Editions today at <http://www.xilinx.com/download>.

Techway
The way of innovation

TigerFMC

200 Gbps bandwidth
VITA 57 compliant

The highest
optical
bandwidth
for FPGA
carrier



www.techway.eu

This year's best release.

The definitive resource for software developers speeding C/C++ & OpenCL code with Xilinx SDx IDEs & devices

The Award-winning Xilinx Publication Group is rolling out a brand new trade journal specifically for the programmable FPGA software industry, focusing on users of Xilinx SDx™ development environments and high-level entry methods for programming Xilinx All Programmable devices.



Xcell Publications

Solutions
for a
Programmable
World



This is where you come in.

Xcell Software Journal is now accepting reservations for advertising opportunities in this new, beautifully designed and written resource. Don't miss this great opportunity to get your product or service into the minds of those who matter most. Call or write today for your free advertising packet!

For advertising inquiries contact
xcelladsales@aol.com
or call: 408-842-2627.

 **XILINX**
ALL PROGRAMMABLE™

Latest and Greatest from the Xilinx Alliance Program Partners

Xpedite highlights the latest technology updates from the Xilinx Alliance Program ecosystem.

The Xilinx® Alliance Program is a worldwide ecosystem of qualified companies that collaborate with Xilinx to further the development of All Programmable technologies. Xilinx has built this ecosystem, leveraging open platforms and standards, to meet customer needs and is committed to its long-term success. Alliance members—including IP providers, EDA vendors, embedded software providers, system integrators and hardware suppliers—help accelerate your design productivity while minimizing risk. Here are some highlights.

PRO-AV BOARD PROVIDES 4K/UHD HDMI TRANSPORT OVER IP

The Barco Silex Viper-HV-4K OEM board for 4K HDMI transport over Internet Protocol, which is built around a Zynq®-7000 All Programmable SoC, compresses 4K/UHD video using hardware-based, real-time SMPTE 2042 VC-2 LD High Quality (visually lossless) video compression and sends it over 1-Gbps Ethernet. Using it, developers of professional AV systems can create point-to-point unicast or multicast distribution networks with low-cost, widely available IP networking equipment and standard, low-cost Ethernet cables.

For more information, please visit <http://www.barco-silex.com/>.

FPGA PROTOTYPING BOARD HAS RATED CAPACITY OF 30M ASIC GATES

As the name suggests, the UltraScale™ XCVU440 FPGA Module from proFPGA is based on a Xilinx® Virtex®

UltraScale VU440 and the prototyping board has a rated capacity of 30 million ASIC gates. This FPGA Prototyping board offers 10 extension sites with as many as 1,327 user I/Os for daughterboards (e.g. memory boards, interface boards), interconnecting cables or customer-specific application boards. The product works in combination with proFPGA's uno, duo or quad motherboards, which respectively carry one, two or four members of the growing family of proFPGA's FPGA Prototyping boards. That's 120M ASIC gates worth of prototyping capacity when you fully load a quad motherboard with four UltraScale XCVU440 FPGA Modules.

For more information, please visit <http://www.prodesign-europe.com/profpga/>.

VPX MODULE MERGES VIRTEX-7 FPGA, FREESCALE POWER PROCESSOR

The X-ES (Extreme Engineering Solutions) XPedite2470 merges a Freescale P1010 QorIQ processor based on an 800-MHz Power Architecture e500-2

microprocessor core with the considerable processing power and programmable-I/O flexibility of a Xilinx Virtex-7 FPGA on one ruggedized, air-cooled, 3U VPX module measuring a mere 100 x 160 mm. The module accommodates as much as 512 Mbytes of DDR3L-800 SDRAM, 1 Gbyte of NAND flash memory and 16 Mbytes of NOR flash memory connected to the P1010 processor, and as much as 4 Gbytes of DDR3L SDRAM and 128 Mbytes of NOR flash memory connected to the Virtex-7 FPGA.

For more information, please visit <http://www.xes-inc.com/>.

CARDSHARP SBC PUTS ZYNQ SOC INTO RUGGED XMC FORM FACTOR

The Cardsharp single-board computer (SBC) from Innovative Integration packs a Xilinx Zynq-7000 All Programmable SoC (the Z-7045 device) along with 3 Gbytes of DDR3 SDRAM, 16 Mbytes of QSPI flash memory and 32 Gbytes of eMMC flash memory into

a 149 x 74-mm XMC form factor for rugged application. The card, which has built-in 1-Gbps Ethernet and USB 2.0 ports, runs many of the Zynq SoC's programmable I/O pins including all eight of the 12.5-Gbps GTX transceiver ports to an HPC FMC I/O site for maximum interfacing flexibility.

For more information, please visit <http://www.innovative-dsp.com/>.

DSMART IP CORE LETS XILINX DEVICES READ, WRITE SMART CARDS

Ubiquitous in Europe and now becoming a staple in the United States as well, smart cards based on chip-and-PIN cryptosecurity represent a leading world standard for secure financial and related transactions. Digital Core Design's (DCD) DSMART IP Core—compatible with the ISO/IEC 7816 - 3: 2006 and EMV 4.1 security standards—makes it possible to implement a contact smart-card reader using the programmable logic in an FPGA. The core has been tested with and implemented by a Digital Core Design licensee using a Xilinx Artix®-7 FPGA.

For more information, please visit <http://www.digitalcoredesign.com/>.

LOTS OF DSP PROTO CAPABILITY IN S2C'S KU115 LOGIC MODULE

S2C has announced the availability of its Single KU115 Prodigy Logic Module, which incorporates one Xilinx Kintex UltraScale XCKU115 FPGA with—among other on-chip resources—a whopping 5,520 enhanced UltraScale DSP48E2 slices. The Prodigy KU Logic Module is well suited for calculation-intensive applications and, according to S2C, offers more DSP resources than any other prototyping board on the market. Along with the many thousands of DSP slices, the Single KU115 Prodigy Logic

Module offers forty-eight 16.3-Gbps GTH transceivers—also supplied by the Xilinx Kintex UltraScale KU115 FPGA—for high-speed communication throughput.

For more information, please visit <http://www.s2cinc.com/>.

ANALOG DISCOVERY 2 PACKS DSO, LOGIC ANALYZER, POWER SUPPLY

Just over a year ago, Digilent rolled out the extremely versatile Analog Discovery multifunction instrument, which is based on a Xilinx Spartan®-6 FPGA. Now Digilent has come out with an encore product, the Analog Discovery 2, that implements additional features. Among them are a two-channel, 100-Msample/second USB digital oscilloscope with 14 bits of resolution and differential inputs; a two-channel, 100-Msps, 10-MHz arbitrary function generator with AM/FM modulation; and a stereo audio amplifier to drive external headphones or speakers. Also onboard are the following:

- A 16-channel, 100-Msps digital logic analyzer and bus decoder
- A 16-channel, 100-Msps pattern generator
- A 16-channel virtual digital I/O port for including buttons, switches and LEDs
- A two-input/output digital trigger for linking multiple instruments
- A single-channel DVM (AC, DC, ±25V)
- A 1-Hz to 10-MHz network analyzer (Bode, Nyquist, Nichols transfer diagrams)
- A spectrum analyzer that performs power spectrum and spectral measurements (noise floor, SFDR, SNR, THD, etc.)

Multiple digital bus analyzers (SPI, I²C, UART, parallel) are also included, along with two programmable power supplies (0 to +5V, 0 to -5V at 50 to 700 mA, 500 mW max).

For more information, please visit <http://store.digilentinc.com/>.

\$975 KINTEX ULTRASCALE FPGA DEVELOPMENT KIT IS A REAL BARGAIN

Do you want immediate access to the most advanced Xilinx FPGA architecture so that you can start your next system design—like right now—but you're looking for a solution that's easy on your development budget? Avnet has a new development kit for you. It's the \$975 Xilinx Kintex UltraScale Development Kit featuring a 20-nanometer Kintex UltraScale XCKU040-1FB-VA676 device. That's a 27 x 27-mm device, the smallest Kintex UltraScale device, with 1-mm ball pitch.

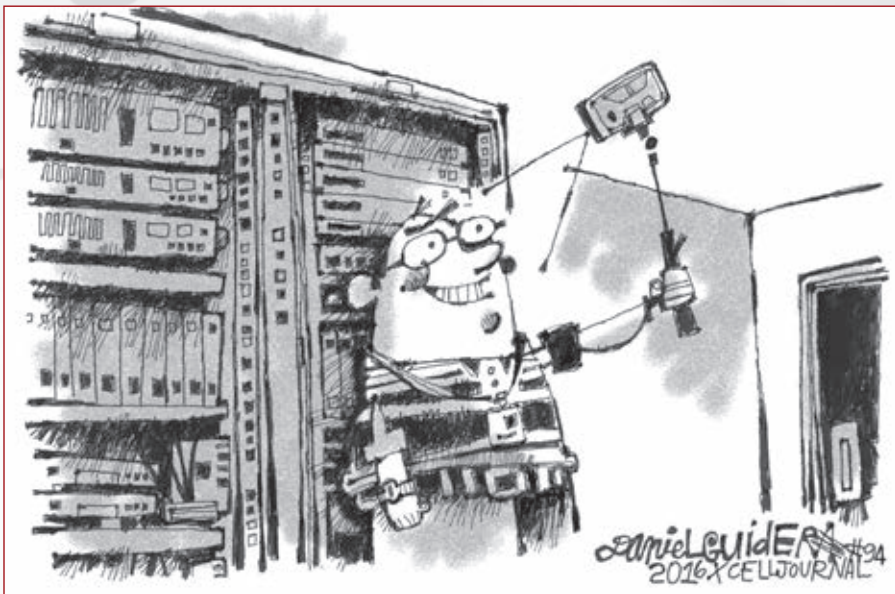
For more information, please visit <http://www.avnet.com/>.

€99 DEV BOARD PUTS ZYNQ SOC INTO RASPBERRY PI FORM FACTOR

The €99 Trenz Electronic TE0726 ZynqBerry development board puts a Xilinx Zynq-7010 SoC into a Raspberry Pi-compatible form factor with 64 Mbytes of LPDDR2 SDRAM, four USB ports (in a hub configuration), a 100-Mbps Ethernet port, an HDMI port, MIPI DSI and CSI-2 connectors, a PWM digital audio jack and 128 Mbits of flash memory for configuration and operation.

For more information, please visit <http://www.trenz-electronic.de/>.

Xpress Yourself in Our Caption Contest



DANIEL GUIDERA

There's a time and a place for selfies, and the workplace may or may not be one of them. You decide as you take a gander at our cartoon engineer setting himself up for a glamour shot against the backdrop of a server farm. We invite you to Xercise your funny bone and submit an engineering- or technology-related caption for this self-centered cartoon. The image might inspire a caption like "En route to finalizing his self-correcting network design, Sam decided a selfie was in order."

Send your entries to xcell@xilinx.com. Include your name, job title, company affiliation and location, and indicate that you have read the contest rules at www.xilinx.com/xcellcontest. After due deliberation, we will print the submissions we like the best in the next issue of *Xcell Journal*. The winner will receive a Digilent Zynq Zybo board, featuring the Xilinx® Zynq®-7000 All Programmable SoC (<http://www.xilinx.com/products/boards-and-kits/1-4AZFTE.htm>). Two runners-up will gain notoriety and fame when we print their captions, names and affiliations in the next issue.

The contest begins at 12:01 a.m. Pacific Time on March 1, 2016. All entries must be received by the sponsor by 5 p.m. PT on May 1, 2016.

ADAM HALE, a computer engineering student at Texas A&M University (College Station, Texas), won a shiny new Digilent Zynq Zybo board with this caption for the buggy cartoon in Issue 93 of *Xcell Journal*:



"Frank now realized why they asked him if he could troubleshoot problematic chips on the fly in his recent interview."

Congratulations as well to our two runners-up:


"Carl's old-school manual-debugging skills made quite an impression on the team."

— *Luis Benites, staff hardware FPGA engineer, Spirent (Agoura Hills, Calif.)*

"Joe finally understands all the buzz around IoT."

— *Daniel C. Kline, CEO, ParaTechnica (West Bend, Wis.)*

GET ON TARGET



LET XCELL PUBLICATIONS HELP YOU GET YOUR MESSAGE OUT TO THOUSANDS OF PROGRAMMABLE LOGIC USERS.


Hit your marketing target by advertising your product or service in the Xilinx *Xcell Journal*, you'll reach thousands of engineers, designers, and engineering managers worldwide!

The Xilinx *Xcell Journal* is an award-winning publication, dedicated specifically to helping programmable logic users – and it works.

We offer affordable advertising rates and a variety of advertisement sizes to meet any budget!

Call today :
(800) 493-5551
or e-mail us at
xcelladsales@aol.com


Join the other leaders in our industry
and advertise in the *Xcell Journal*!



www.xilinx.com/xcell/



Xcell Journal Adds New Daily Blog



Xilinx has extended the Award Winning Journal and added an exciting new *Xcell Daily Blog*. The new site provides dedicated readers with a frequent flow of content to help engineers leverage the flexibility and extensive capabilities of Xilinx products, ecosystem, and customers to create All Programmable and Smarter Systems.

Recent

- [4 Minutes to Error-Free 100G Ethernet Operation using Xilinx UltraScale+ Integrated 100G IP](#)
- [Avnet introduces \\$699 Zynq-based, Multiprotocol, Industry 4.0/IoT MicroZed Kit](#)
- [3 Eyes are Better than One for 56Gbps PAM4 Communications: Xilinx silicon goes 56Gbps for future Ethernet](#)
- [Tiny Zynq Board "does absolutely nothing, and that's a good thing" explains Servaes Joordens](#)
- [How I got lost and beer-buzzed in Prague just so I could see Sphericam's Zynq-based, 360-degree, 4K video VR camera](#)

Visit Blog: www.forums.xilinx.com/t5/Xcell-Daily/bg-p/Xcell