December 2016

EMBEDDED SYSTEMS ENGINEERING

powered by EECatalog

Guiding Embedded Designers on Systems and Technologies

Engineers' Guide to FPGA, PLD and SoC Solutions

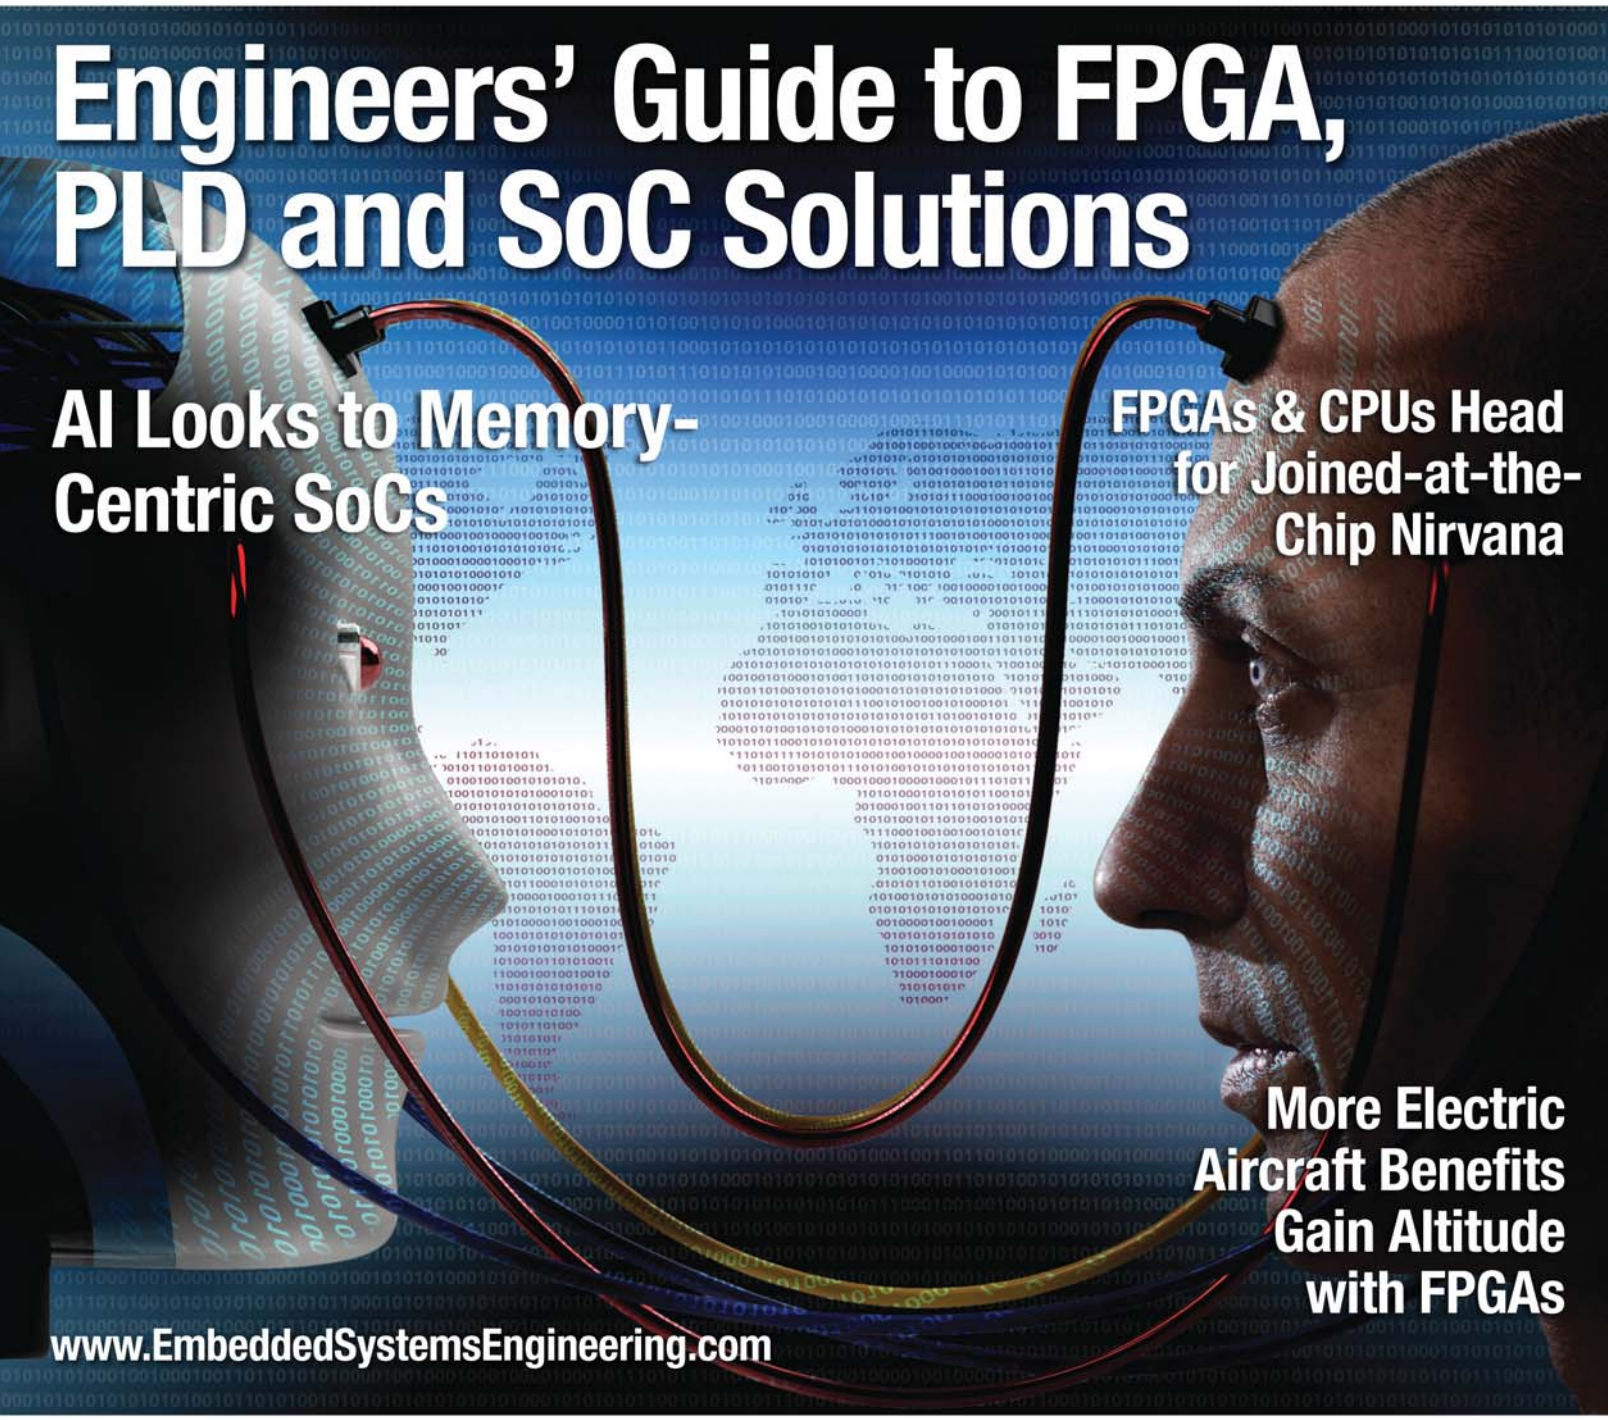

AI Looks to Memory-Centric SoCs

FPGAs & CPUs Head for Joined-at-the-Chip Nirvana

More Electric Aircraft Benefits Gain Altitude with FPGAs

www.EmbeddedSystemsEngineering.com

Xilinx Virtex-7 FPGA-Based Modules

Ideal for Customizable, High-Bandwidth, Rapid Signal-Processing Applications and Available with a Comprehensive Firmware Development Kit to Accelerate Application Development

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

MADE IN THE USA

Designed, manufactured, and supported in the USA

Nürnberg, Germany

March 14 – 16, 2017

embeddedworld

Exhibition&Conference

... it´s a smarter world

Media cooperation

elektroniknet.de

computer-automation.de

ENERGIE & TECHNIK
Solutions for a Smarter World

DESIGN & ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER

MEDIZIN-UND-elektronik.DE

Elektronik
Fachmedium für industrielle Anwender und Entwickler

Elektronik automotive
Fachmedium für professionelle Automobilelektronik

Markt&Technik

Computer & AUTOMATION
Fachmedium der Automatisierungstechnik

MEDIZIN elektronik
Fachmedium für Elektronik in der Medizintechnik

Workplace Romance

FPGAs and SoCs are destined to have a working relationship that gets ever closer, something applications that range from aerospace to virtual reality to machine learning won't mind at all.

By Anne Fisher, Managing Editor

Nishant Mittal, Systems Engineer at Cypress Semiconductor, explains how the growing demand for miniaturization is intertwined with the increasing popularity of SoCs. In this issue Mittal discusses SoCs in the context of such goals as keeping power consumption to a trickle, isolating noise, and optimizing speed.

While wearables and smartphones have been pressing for miniaturization, individuals advocating for More Electric Aircraft (MEA) are urging power electronics to become more reliable, integrated, and higher performing. Shane O'Donnell, manager for Microsemi Corporation's Aviation Center of Excellence, recently participated, along with Microsemi CTO Jim Aralis, in an EECatalog Q&A. O'Donnell wants aerospace manufacturers to examine what newer technologies, including wide bandgap (WBG) semiconductors can bring to the party. He tells us about the company's Power Core Module and about the benefits that occur when SiC MOSFETs replace Si IGBTs. Aralis discusses the greater number of roles FPGAs are taking on—coprocessors in data centers; voice recognition; server applications for baseband generation—as FPGA performance improves. For aircraft, Aralis notes, having "the system on a chip inside an FPGA" can make possible the processing power needed, for example, to optimize a jet engine "on a very quick cycle."

BRIDGE PLAYERS

The bridge game designers play today involves those legacy systems, which, Grant Jennings reminds us in this issue, can include I/O which runs the gamut from "legacy CMOS and proprietary interfaces to source synchronous standards to interfaces that feature an embedded clock for larger systems." Jennings, a Senior Strategic Planning Engineer for Lattice, goes on to explain that systems today have a great number of potential bridging applications. He makes the case for considering an alternative to a "one-size-fits-all" approach.

Also on the topic of connections, Executive Editor Lynnette Reese writes here about a means to achieve higher throughput and lower latency: Remotely connect FPGAs through PCI Express over cables or between backplane segments. Reese spoke with Dolphin Interconnect Solutions about how that company is applying technologies such as device lending and multi-cast to improve the outlook for real-time imaging, time-sensitive transactions, and more.

How we're going to cross the bridge from Artificial Intelligence (AI) and machine learning experiments to accelerated adoption and monetization of these technologies is the topic of the article by Sylvain Dubois, Crossbar, this month. And it is not just bridges, but bottlenecks, that are on Dubois's mind as he explains that an approach to taking on today's processor memory bottlenecks begins with "memory technologies that can be directly integrated on-chip with the processing logic" to realize memory-centric SoC architectures.

Integration that has been the cause of some anxiety among designers, that of Altera with Intel, is one topic our Technology Editor Dave Bursky addresses here. Bursky summarizes for us the presentations Intel gave at the 2016 IDF, where the company revealed its product development directions and both software and hardware applications of Intel Programmable Solutions Group FPGAs.

Revelations in the form of news, videos, white papers, opinion and articles is waiting for you at http://eecatalog.com/fpga/ and all our channels.

EMBEDDED SYSTEMS ENGINEERING 2016
www.embeddedsystemsengineering.com

Vice President & Publisher
Clair Bright

Editorial
Executive Editor
Lynnette Reese
lreese@extensionmedia.com
Managing Editor
Anne Fisher
afisher@extensionmedia.com
Senior Editors
Chris Ciufo
cciufo@extensionmedia.com
Caroline Hayes
Gabe Moretti
Dave Bursky

Creative/Production
Production Traffic Coordinator
LS Jerrett
Graphic Designers
Nicky Jacobson
Simone Bradley
Senior Web Developers
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales
Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15
Sales Manager
Elizabeth Thoma
(415) 255-0390 ext. 17
ethoma@extensionmedia.com

Marketing/Circulation
Jenna Johnson
jjohnson@extensionmedia.com

To Subscribe
www.eecatalog.com

Extension MEDIA

Extension Media, LLC Corporate Office
President and Publisher
Vince Ridley
vridley@extensionmedia.com
(415) 255-0390 ext. 18
Vice President & Publisher
Clair Bright
cbright@extensionmedia.com
Vice President, Business Development
Melissa Sterling
msterling@extensionmedia.com
Human Resources / Administration
Darla Rovetti

Special Thanks to Our Sponsors

Dolphin PENTEK Technologic Systems X-ES

Embedded Systems Engineering is published by Extension Media LLC, 1786 18th Street, San Francisco, CA 94107.

www.embeddedARM.com

INTERNET OF THINGS

SMART
INDUSTRIAL
PRODUCTIVE
MESHED
LOW POWER
INTELLIGENT
CUSTOMIZED
LINKED
PERCEPTIVE
CONNECTED
SCRIPTED
HELPFUL
RUGGED
RELIABLE
WEARABLE
ALL THE
OPEN SOURCE

Single Board Computer

TS-7680

Low Power Industrial Single Board Computer with WiFi and Bluetooth

Starting at
$159
Qty 100

Computer on Module

TS-4900

Industrial High Performance i.MX6 Module with Wireless Connectivity and Flash Storage

Starting at
$109
Qty 100

Touch Panel PC

TS-TPC-7990

7" High End i.MX6 Mountable Panel PC with Dev Tools Such as Debian GNU and QTCreator

Starting at
$399
Qty 100

Technologic Systems

CONTENTS

EMBEDDED SYSTEMS ENGINEERING

Departments

FPGA, PLD & SoC Solutions

Product Showcases

FPGA Boards

Boards and Kits


wh a.
Yep, the impossible just happened.
CES
Consumer Technology Association™
Jan. 5-8, 2017 | Las Vegas, NV
Register now at CES.tech
#CES2017
Celebrating 50 YEARS
CES | 1967-2017

Remotely Connecting FPGAs through PCI Express

An approach that lowers latency and increases throughput can extend benefits to such applications as medical imaging and financial trading.

By Lynnette Reese, Executive Editor

Relatively recent developments in PCI Express (PCIe), a well-known serial expansion bus standard for computers, have resulted in connectivity enhancements. FPGAs can now be remotely connected through PCI Express over cables or between backplane segments for higher throughput and lower latency. New techniques such as Device Lending from Dolphin Interconnect Solutions, combined with enhancements to existing features of PCIe Multi-cast and peer-to-peer transfers, afford FPGAs faster communication to remote resources. Today, PCIe attached FPGAs can transfer data faster with less overhead or be shared as resources in a cluster in new exciting ways.

PCIe connectivity enhancements are profound if an application's goal is high throughput or extremely low latency, which means real-time imaging, time-sensitive transactions, or other demanding functions benefit. A discernible lag in virtual reality isn't acceptable, and neither is a lag in imaging for a colonoscopy when the scope has already crept past what is identifiable on the screen as a polyp. One company, Dolphin Interconnect Solutions, has made inroads into this technology. Dolphin has enhanced the potential for streaming peer-to-peer (P2P) technology, enabled devices to be dynamically lent between systems, and enhanced the capability of PCIe multi-cast. All while achieving amazingly low latencies, just 0.54 microseconds for data transfers.

Herman Paraison of Dolphin Interconnect Solutions explains it this way, "Dolphin specializes in several areas that have opened up around PCIe technology. From device lending to PCIe multi-cast, PCIe enables you to do more with limited resources, ultimately saving money while improving latency and throughput. FPGAs, GPUs, NVMes, or another high-performance computing device can take advantage of the inherent performance feature and flexibility of PCIe."

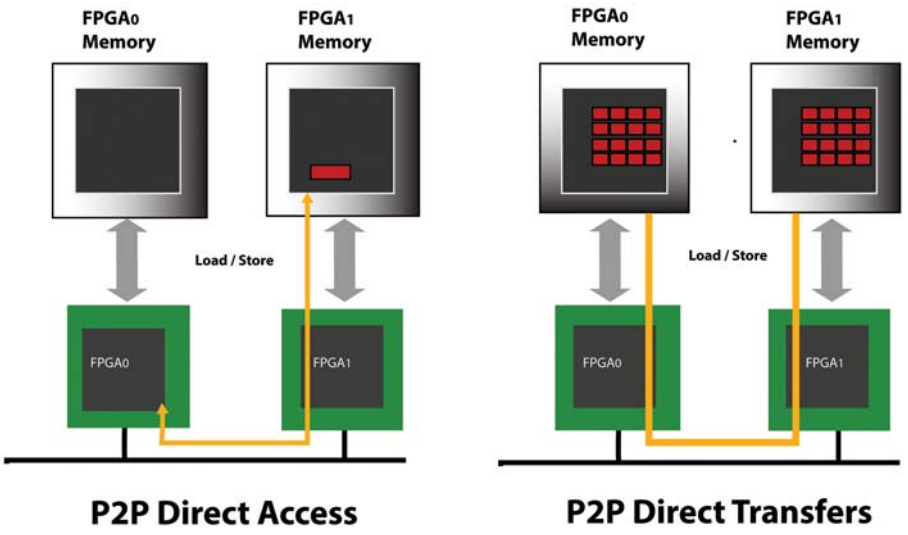

Figure 1: A standard capability of peer-to-peer (P2P) is for FPGAs to be set up to do direct access/RDMA operations between FPGAs within a system.

REMOTE FPGA PEER-TO-PEER TRANSFERS

There have been many implementations of Peer-to-peer with FPGAs communicating with GPUs and NVMe drives. Some PCIe chipset and systems support peer-to-peer communication between slots in a single system. Applications such as GPU direct enable peer-to-peer transfers across PCIe, allowing direct access to GPU memory within the same system. A standard capability of peer-to-peer

is for FPGAs to be setup to do direct access/RDMA operations between FPGAs within a system. This is illustrated in Figure 1.

Remote peer-to-peer is supported. Dolphin extends this concept to multiple systems or backplane segments.

Figure 2 illustrates a remote peer-to-peer configuration.

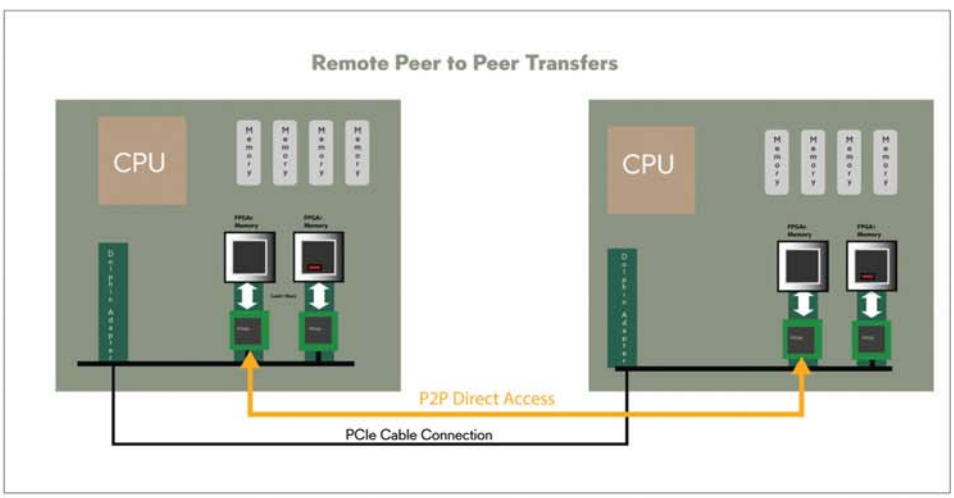

Figure 2: A remote peer-to-peer configuration.

The PCIe Network is used to connect two systems via a fiber or copper cable. These systems can then implement remote P2P to communicate over PCIe. Like standard PCIe peer-to-peer transfers, there is no requirement to utilize the CPU or memory. Data is transferred directly to a memory region in the remote FPGA. The target device can vary. It can be a CPU, GPU, or remote memory. In all instances, remote peer-to-peer can be used. Now, distributed FPGA systems can take advantage of the low latency of PCIe. This reduced overhead way of communicating will benefit real-time distributed systems such as financial trading systems seeking the lowest latency for trading transactions.

MULTI-CASTING, OR REFLECTIVE MEMORY

Multi-casting is another PCIe capability that should not be left unexplored, as it establishes a

deterministic, low-latency, high-speed interface for sharing data. Also referred to as "reflective memory," or "reflected memory," multi-casting is an advanced PCIe solution that gives instantaneous and reliable data distribution. Multi-casting enables an FPGA to deliver data to multiple nodes without incurring extra overhead for the FPGA. In plain English, a single FPGA can transmit data to multiple other devices, instantly at the same time, as if it were communicating only to one device. The heavy lifting is accomplished by implementing advanced PCIe features. Dolphin's PCIe multi-casting solution enables significantly higher performance at a lower cost than similar solutions. Connections with copper, long distance fiber-optic or mixing fiber and copper are possible.

DEVICE LENDING

The flexibility of being able to pool or move resources within a cluster is very attractive for system architects seeking to maximize usage of high valued resources such as FPGAs. Device lending is a concept targeted at maximizing device utilization. With device lending, FPGAs, GPUs, NVMe drives or other devices can be dynamically lent to remote systems from other nodes in a PCIe network. Device lending doesn't require a CPU in a data path; PCIe enables direct configuration of resources into a PCIe domain. Therefore, no changes are required to operating systems, device drivers, or applications. Devices that are "lent out" may be removed and added to an OS during run-time for devices that support Hot-Add and Hot-Removal of PCIe devices. Clusters of FPGAs, GPUs, and CPUs can all share the same device resource pool.

For example, real-time, computer-aided diagnosis-support applications benefit from being able to apply more computing resources as necessary for more accurate diagnosis. With device lending, compute resources can be added to a system in the PCIe cluster from another system in the cluster without needing to physically install the device. In the case of computer-aided diagnosis, GPUs can be added to a thin client in an operating theater running, for example, polyp detection software. The GPUs can be installed in a remote system in a server room. The thin client in the operating theater will borrow the required number of GPUs from the remote server. The polyp detection software will use the GPUs as if they were locally installed in the thin client. The advantage is the thin client doesn't need the power, size, or noise of the remote server. The software doesn't need to change, and the GPUs can be returned to the remote server for use by another client in another operating room.

Dolphin Interconnect Solutions' expertise in utilizing the newest PCIe technologies to competitive advantage have enhanced networking and benefitted many embedded applications. Financial applications have found and adopted Dolphin's solutions as a means to trade with ultra-low latency for high-speed, high-stakes trading.

The growth in complexity of FPGAs stems from increasingly demanding applications as markets seek to gain an edge. PCIe, although considered by many as a simple connectivity solution for personal computers, has the potential to be the tipping point for any application where low latency, high-speed throughput, and extremely low computing overhead are at the top of the "must-haves" list.

Lynnette Reese is Executive Editor, Embedded Systems Engineering and Embedded Intel Solutions, and has been working in various roles as an electrical engineer for over two decades. She is interested in open source software and hardware, the maker movement, and in increasing the number of women working in STEM so she has a greater chance of talking about something other than football at the water cooler.

Memory-Centric SoCs and AI

What can be realized as high-performing and low-power memory-centric SoCs pair with IoT devices and cloud-based servers?

By Sylvain Dubois, Crossbar

According to a recent Stanford report[1], artificial intelligence (AI) will not replace the need for humans any time soon, but it will have a profound impact on everyday lives, transforming industries from transportation, education, medical and entertainment. The exciting developments that deep neural network algorithms and machine learning can produce will result in new intelligent interfaces, new virtual assistants, and advancements in a wide range of applications across multiple industries—once only dreamed about in research laboratories and now becoming possible for consumer adoption.

"40% of enterprises use hyperconverged units as a standard building block in the data center, climbing rapidly over the next two years."

—451 Research—Voice of the Enterprise

The first examples of machine learning algorithms arrived to a larger audience when computers started to defeat chess players and more recently the best Go player (Figure 1). However there are many more possibilities. We are now entering an era where surveillance cameras and autonomous driving are real-life applications of advanced computer vision powered by machine learning techniques. Voice recognition and smart sensors using deep learning algorithms can increase context awareness for robotics applications. Cloud based analytics for e-commerce and advertisement recommendations, and business analytics or medical treatment recommendations are other examples....and the list goes on.

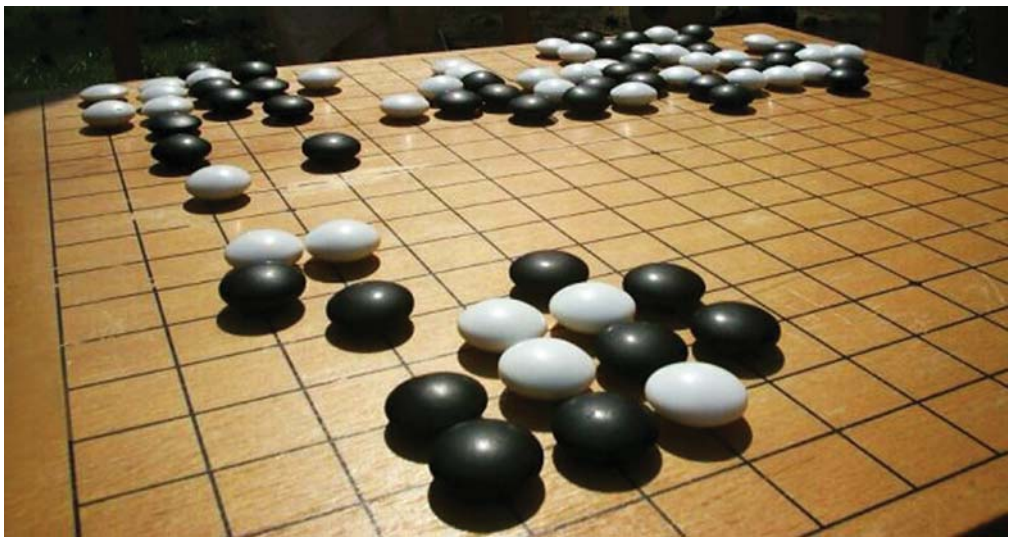

Figure 1: Courtesy https://de.wikipedia.org/wiki/Go_(Spiel)#/media/File:Go_board.jpg Donarreiskoffer - Selbst fotografiert

To transition AI beyond experiments, the user experience needs to improve in order to accelerate adoption and enable companies to make money. AI and more specifically machine learning are at the core of how nearly every business of the future will make money and have greater differentiation compared to their competition. In the past, storing data was a cost center, but today it has the opportunity to serve as a profit center as it trains algorithms and creates new applications.

1. https://ai100.stanford.edu/sites/default/files/ai_100_report_0906fnlc_single.pdf

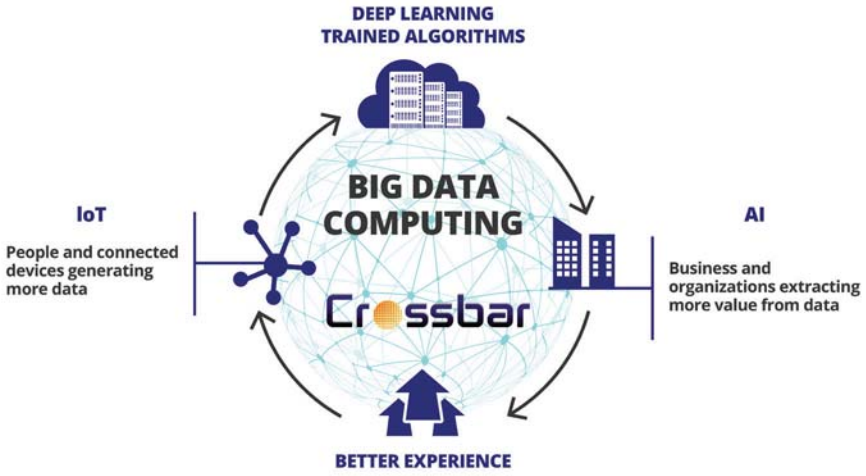

Figure 2: AI can turn data into experience.

Algorithms require large data sets to be trained. The algorithms needed to process that data are increasingly complex, and everything needs to happen in real time. For several years, many companies well understood that data gathering was the most valuable asset of their business model. There is no doubt now why social media companies provided a very handy and free platform to share pictures and post comments. This huge amount of qualified data is now a very valuable asset and justifies the high valuation of these companies. Most of the data collected by their applications or devices is directly fed into deep neural networks to train them. Zettabytes of data storage that can be quickly accessed to train machine-learning algorithms will enable even more innovative products and applications.

IOT AS STEPPING STONE TO AI

IoT is a natural stepping-stone to AI, and eventually the two will merge when all these connected devices become smarter and more predictive. While IoT is about data collection, AI is about data consumption. The more data, the smarter the AI algorithms will be.

Voice activated virtual assistants built into smart devices like Apple's Siri, Amazon's Alexa, or Google's Home are getting more and more traction. Experiences similar to the intelligent virtual assistant Samantha from the romantic science-fiction film "Her" from Spike Jonze may only be a few months away. At a more short-term scale, growth of AI is increasing in surveillance cameras, where data has to be analyzed locally. Performance and cost challenges to upload streams of 4K video to the cloud for processing and downloading will prevent the required action to be taken. Imagine driving assistance malfunctions when the wireless connection is poor. Most smart devices need a better understanding of their environment and consumers' habits to better perform according to users' expectations. Evolution is based on learning from our past experiences and improving our future behavior.

For these technologies to produce rich sets of data that can be analyzed and acted upon, innovative memory technologies are needed to deliver high performance and low energy.

Non-volatile memory technologies such as Crossbar RRAM are helping address the performance and energy challenges of embedded IoT by delivering lower power and lower voltage operation, monolithic integration with computing cores, faster read, and byte-addressable writes. RRAM is the lowest energy memory technology for IoT applications and can be integrated with processing cores on a single chip solution.

On-chip integration with dedicated logic makes RRAM adequate to accelerate deep neural network algorithms. Data coming from sensors can be stored on-chip and directly fed through deep neural networks to take direct actions. Objects become smart not by fetching lines of software but by reacting to external data coming from various sensors. Data collection and processing could be integrated onto a single chip solution with embedded RRAM. By integrating high-density memory on-chip with the processor at the same node, the inherent latencies associated with moving data from processors to off-chip memory subsystems and back again are improved. As a result, RRAM technology is an important innovation to accelerate the potential of the new big data, artificial intelligence universe, enabling a multitude of applications that can speed performance, dramatically improve energy efficiency, enable advanced security and reduce chip count and size.

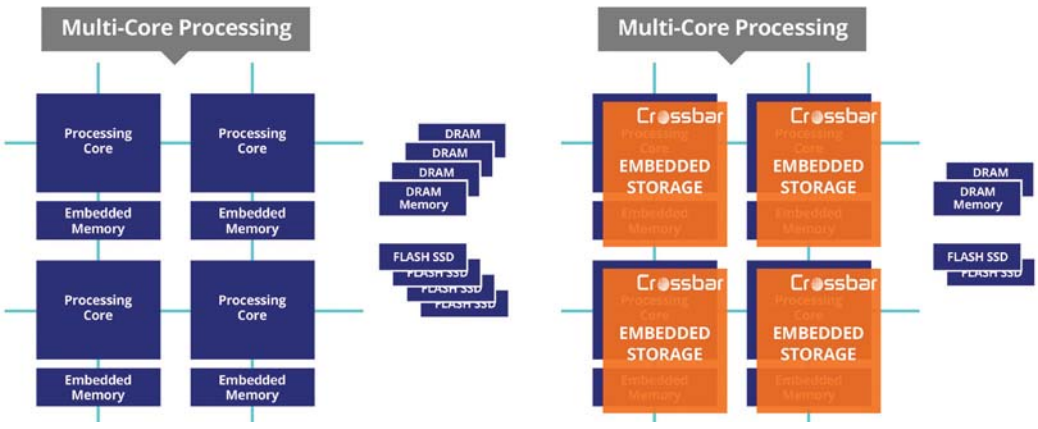

Figure 3: Non-volatile memory and computing logic on same silicon speeds data access and improves energy efficiency.

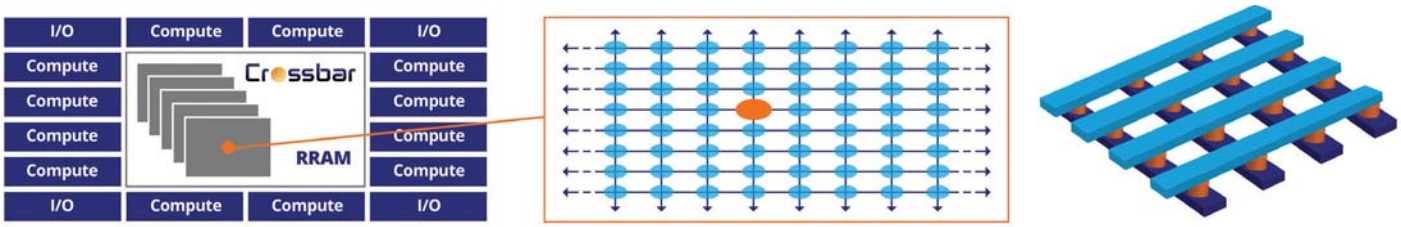

Figure 4: An RRAM-centric parallel computing platform for new AI experiences.

NEW ARCHITECTURE NEEDED FOR AI

Moore's Law is screeching to a halt and CPU refreshes are less frequent. The von Neumann memory bottleneck of current Intel-based processors could be solved with new system architectures that will be more memory-centric. The performance gap between storage technologies and computing has to be reduced. Traditional Flash-based storage solutions deliver read latencies in the 100µs range while Crossbar 3D RRAM reaches 1µs. In addition to Crossbar's RRAM, there are several other initiatives powered by emerging memory technologies like Intel Micron XPoint PCM, Everspin, and Avalanche MRAM all trying to solve this challenge.

Traditional architectures in data centers usually have three distinct and stand-alone parts:

- Computing with best-in-class processing cores and attached DRAM memories
- Data storage with lowest cost per bit SSDs or HDDs
- Networking to interconnect the computing part and data storage part

Analysts are observing a massive trend towards higher integration in data centers where these three parts: computing, storage, and networking are condensed in a compact form-factor, called the hyperconverged server.

Based on some recent reports, already 40% of the enterprise data centers use hyperconverged units, and this market is expected to grow close to 80% in the next five years.

The data center infrastructure market is evolving rapidly to handle lower latencies across the various elements of a server. Bringing all components into the same box is a way to reduce latencies, total cost of ownership, and power. In hyperconverged servers, compute, storage and networking can now be interconnected in a more efficient way. Emerging industry consortiums are arising to handle high-bandwidth low latency data accesses across processors, storage and IOs. In high performance computing applications, Intel Micron 3D XPoint PCM and Crossbar 3D RRAM are expected to provide significant improvements by reducing the performance gap between storage and computing.

Memory technologies that can be directly integrated on-chip with the processing logic will enable brand new memory-centric SoC architectures. When non-volatile memory and computing logic share the same silicon, the performance bottleneck of the external bus is removed. One particular application for embedded persistent memory technology is deep learning hardware acceleration built on memory-centric SoC. Artificial intelligence and deep learning represent the most probable evolution of computing in the decade. The deep neural network is all about data and how trained algorithms react to new sets of data. There is tremendous interest from the industry to develop new computing platforms with massive parallelism coming from multiple processing engines with dedicated embedded RRAM cores. Several companies are already working on using RRAM cells as synapses in neuromorphic processing architectures.

Bringing these high-performing and low-power memory-centric SoCs to IoT devices and cloud-based servers will make data and computing ubiquitous—available to users and applications whenever, wherever they need it. These solutions will deliver not only the "BIG" data to the cloud, in terms of capacity and scale, but also the energy efficiency, security, and blazing performance with low latency required for embedded applications to make new artificial intelligence applications possible.

Sylvain Dubois joined the Crossbar management team in 2013 as Vice President of Strategic Marketing and Business Development. Prior to joining Crossbar, Mr. Dubois led strategic product positioning and market engagement for developing new products at Spansion. He holds a Master of Science in Microelectronics from E.S.I.E.E. (Paris), University of Southampton (UK) and Universidad Pontifica Comillas (Spain).

System Solutions with Intel CPUs and FPGAs

The combined forces of Intel and Altera give designers a wide range of system options, with more to come as Intel further integrates its CPUs with FPGAs.

By Dave Bursky, Technology Editor, Extension Media

The purchase of field-programmable gate array (FPGA) supplier Altera by Intel® in 2015 surprised many companies in the electronics industry. And now that Altera is integrated into Intel, everyone is concerned about the future of the Altera programmable products and what direction Intel will take to leverage the configurable technology.

At this year's Intel Developer Forum Intel hosted an entire day of presentations outlining product development directions and both software and hardware applications of the FPGAs. During his keynote presentation Intel CEO Brian Krzanich covered two major issues on the top of all attendees' minds. First, he promised to keep the existing FPGAs with ARM cores and not replace the cores with x86 compatible cores. And, second, he gave a short look at the future roadmap for the FPGAs, which include system-on-a-chip solutions that combine the i86 cores and FPGA fabric, initially in the same package and eventually on the same chip. In these future products he indicated that x86 cores would play a significant role.

MANY APPLICATION OPPORTUNITIES

The role FPGAs will play in many applications was outlined in paper SOCTS02, in which Ian Land, the FPGA Marketing Manager for the Data Center Group at Intel, discusses how FPGAs can accelerate many functions within today's legacy data centers (Figure 1a). In a typical data center, the FPGAs supplement the host CPU and can serve multiple roles: for example, in the network processing area the FPGAs can perform in-line processing, pre-processing, pre-filtering, cryptography, compression, protocol bridging and still other functions. And in the compute section, the FPGAs can accelerate machine learning algorithms, video transcoding, custom algorithms and still other tasks. Additionally, storage-related tasks such as compression, indexing and cryptography can also be accelerated. The flexibility of the FPGA SoCs that contain embedded processor cores allow them to also handle board management, protocol bridging and security tasks.

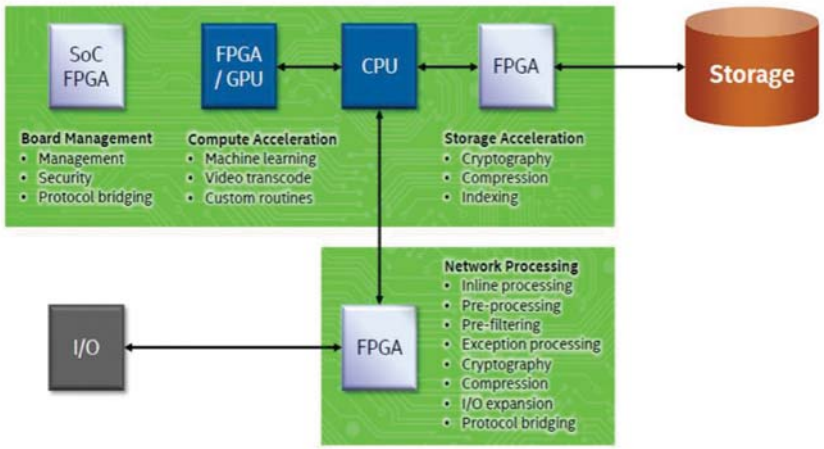

Figure 1a: In today's legacy data centers, FPGAs are typically found in different sections of system, handling tasks such as network processing, CPU acceleration, storage management and board management functions. (Image courtesy of Intel)

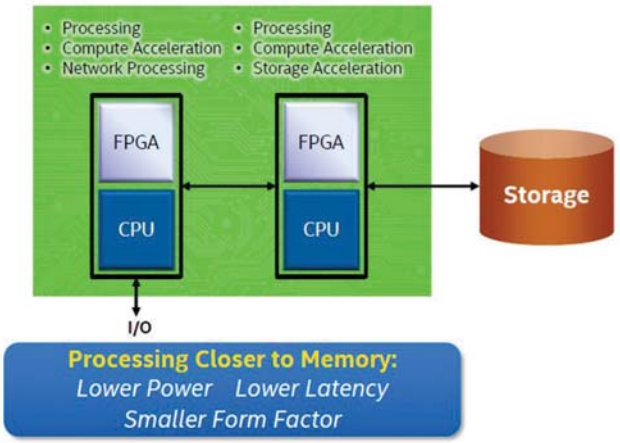

Figure 1b: The close integration of the CPU and FPGA improves system performance by reducing latency, while lowering system power and reducing the board area. (Image courtesy of Intel)

Further integration of the CPU and FPGA fabrics can reduce system complexity and bring the processing closer to memory to achieve higher performance by reducing latency (Figure 1b). Such a future data center design would also benefit since the integration would reduce power consumption and reduce the board area required. CPUs such as the Intel® Xeon® or Xeon® Phi™ would work side by side with FPGAs such as the Cyclone or Arria 10.

COMMUNICATIONS SOLUTIONS

In the same session, Mike Fitton, from Intel's Wireless and Access BD Programmable Solutions Group, examined how the FPGAs play a role in base stations and support computationally complex signal processing, new usage models, and the ability to handle emerging standards such as 5G communications. The embedded DSP functions, the programmable fabrics and other on-chip resources suit the FPGAs to handle the challenges posed by the evolving wireless infrastructure requirements (Figure 2).

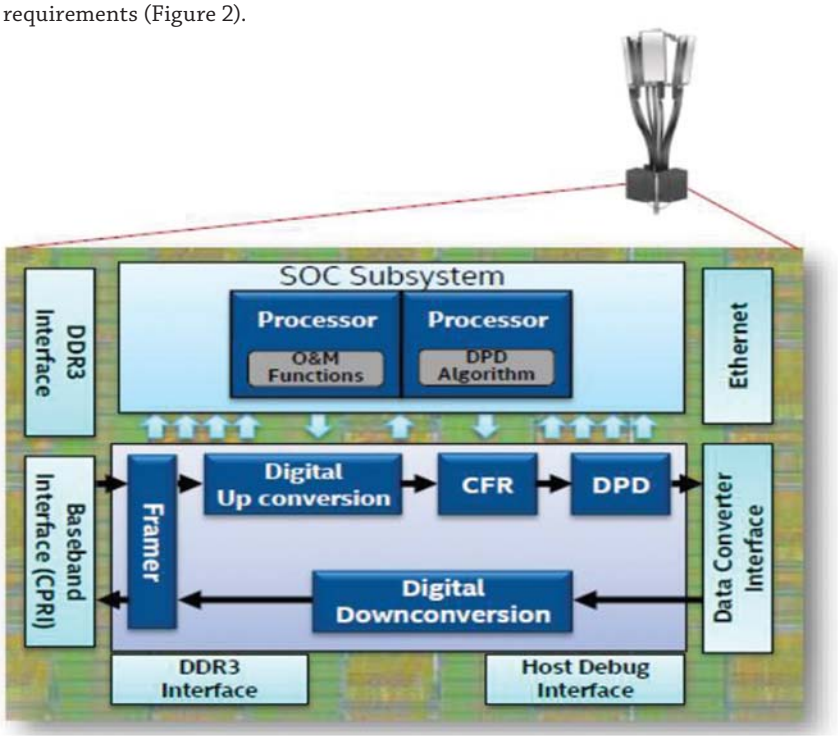

Figure 2: In a typical basestation's radio subsystem, a system-on-a-chip FPGA with embedded processors provides a flexible solution to handle both the system management and signal processing requirements. (Image courtesy of Intel)

Arria 10 FPGA provides a system-on-a-chip (SoC) solution for the radio subsystem. The programmable logic fabric and embedded DSP blocks readily handle the filtering and signal conditioning and can perform amplifier digital predistortion. Supporting the functions implemented in the FPGA fabric, the embedded processor cores also perform all the housekeeping functions and system management, while at the same time handling some of the complex algorithms.

The growth of the Internet of things (IoT) in the industrial market segment and the use of FPGAs to support that market was discussed by Joerg Bertholdt of Intel's Programmable Solutions Group. The huge amounts of data generated by various endpoints such as smart meters (35 Gbytes/day), jet planes (1 Tbyte/flight), and many other devices require significant preprocessing of the data in a smart gateway to reduce the amount of data sent to the cloud. By better analyzing the collected data, system efficiencies can be improved, failure prediction can prevent downtime by replacing parts before they fail, and that, in turn, improves the bottom line. The gateway in the example Bertholdt discussed was based on a Cyclone V SoC that includes dual ARM Cortex A9 processor cores that run a Linux kernel, implement OpenSSL encryption, and provide multiple Ethernet ports.

DATA AND SYSTEM SECURITY KEY FOR IOT AND DATA CENTERS

Keeping systems secure is a challenge for every product, ranging from servers in a data center to the endpoints in IoT applications. In these areas, the FPGAs can also play an important role as Ryan Kenny, a technical marketing manager at Intel explained in his presentation (SOCTS04). The wide range of places where the data and intellectual property can be compromised include the supply chain, data in transit, data at rest, remote access/updates, and physical and reverse engineering. Thus preventing the circuit design or data from being compromised is a huge challenge since attacks can take on many forms:

- Non-invasive such as side-channel, boot code, over-voltage, over temp, clock skew, glitching, and others
- Invasive such as decapping a package, microprobing, chemical eroding, focused ion beam probing
- Supply chain such as mask set analysis, cloning, counterfeiting, boot code changes
- Software such as the addition of extra circuits
- Intellectual property such as the addition of extra circuits, hidden failure mode, extraction of data and keys

In addition to these well-established techniques to compromise the systems, new attacks are also emerging as systems get more interconnected. New types of attacks include denial of service (DoS), injection of ARM malware into SoC type FPGAs, and inter-FPGA attacks that include unauthorized partial

reconfiguration, creation of routing violations, local thermal effects and partial reconfiguration timing attacks. To mitigate such attacks, FPGAs have included fuse scrambling to prevent embedded keys from being read out. Additionally, with the Stratix 10 series, physically unclonable functions (PUF) are used to protect the keys. In future versions of the FPGAs, Intel plans on providing features such as scripted zeroization, user-accessible PUF, secure RMA/debut/upgrade, whole device encryption/authentication, and even total domestic manufacturing.

A related presentation covering Security with Software (SOCTS07) presented by Rod Frazer, an embedded specialist field application engineer for Intel's Programmable Solutions Group, along with Michel Chabroux from Wind River® and Lisko Lappalainen from MontaVista Software examined various techniques software can use to provide security. Frazer led off by detailing secure-boot techniques for both the Cyclone V and Arria 10 FPGA families. Intel and third-party partners also offer multiple hardware IP cores that could be incorporated into the secure boot design for either family. The secure boot capability on the Arria 10 FPGA includes an authentication capability that can be used independently or combined with encryption. A public authentication key (PAK) certification authority infrastructure (256 bit ECDSA authentication) and key authorization key (KAK) storage on the FPGA give designers more options.

The presentation by Wind River examined the security issues with using an embedded real-time operating system with IoT devices. For IoT systems, security should be integrated in every aspect of the device to prevent any point from being compromised. To that end, Chabroux highlighted key features in VxWorks® that provide security support—encrypted containers, disk, secure boot, digitally signed binaries, TPM and TrouSerS (Trusted Computing Group software stack) to verify system integrity at runtime, security events handler, TEE (trusted execution environment) support (isolate credit-card application), network security (OpenSSL, Firewall, IKE, SCEP, etc.) to separate operations and business networks, and still other features. Lastly, Lappalainen rounded out the security discussion by examining how MontaVista's embedded Linux software solution can secure SoC FPGAs. The use of a simple Linux root of trust, supplied by a hardware component, provides a chain of attestation through the bootloader, kernel, and userspace. That, combined with the integrity management architecture in Linux, allows transparent checking of file and FS metadata signatures before executing or reading.

The IDF2016 conference covered many other topics—the use of FPGAs for software acceleration, performance optimization, and applications in motor control, smart drones, memory configuration, and still other topics. For access to the presentations, go to: http://www.intel.com/content/www/us/en/intel-developer-forum-idf/san-francisco/2016/idf-2016-san-francisco-technical-sessions.html and click on "Full Technical Session Catalog".

Dave Bursky, the Semiconductor Technology Editor for Chip Design magazine, is also president of PRN Engineering, a Technical Writing and Market Consulting company. In addition to joining Chip Design magazine in 2006, he was also the Technical Editorial Manager at Maxim Integrated Products, and prior to Maxim, Dave spent over 35 years working as an editor and engineer.

Time for a New Approach to Video Bridging

As the I/O landscape for automotive, industrial, and other applications continues to evolve at near breakneck speed, designers need a single programmable video bridge that supports multiple uses.

By Grant Jennings, Lattice Semiconductor

Take a walk around a recent MIPI Development Conference and it's easy to see how fast the I/O landscape is changing. The rapid proliferation of image sensors and displays and the escalating adoption of new low cost MIPI standard interfaces has spurred innovation and allowed designers to take advances from the mobile arena into a far wider array of applications. But to reap those enticing power and cost benefits, today's designers must typically interface to legacy systems that in many cases represent years of investment. They must cope with I/O that can range from legacy CMOS and proprietary interfaces to source synchronous standards to interfaces that feature an embedded clock for larger systems.

The number of potential bridge applications in today's systems is enormous. Every time an interface doesn't match between an image sensor and a processor, or between a processor and a display, a bridge is needed. And every time the number of interfaces between a processor and an image sensor or display does not match, a bridge is also needed.

Given the large number of interface permutations such as the number of PHYs and bits per pixel, not to mention the specific cost, power, and footprint requirements

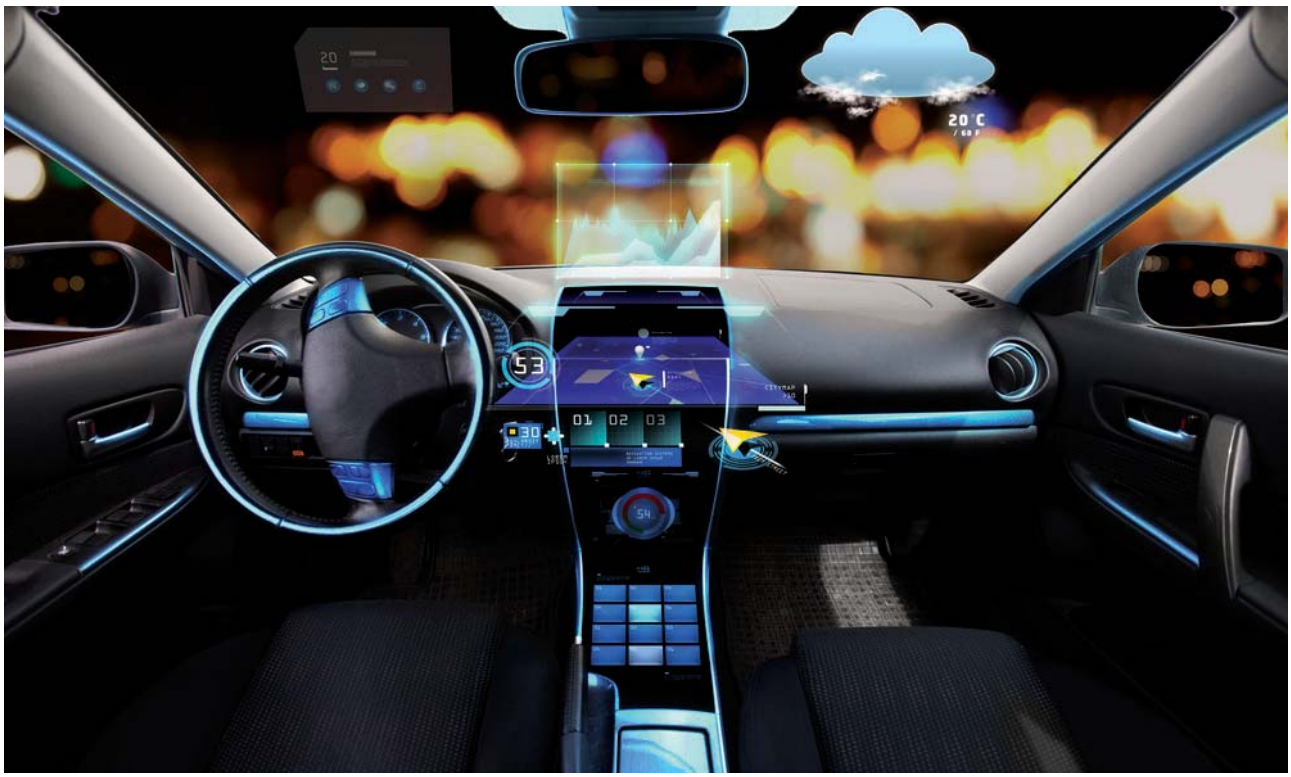

Figure 1: Encompassing support for larger pitch packages, required by industrial and automotive applications, must be a part of any successful programmable bridging solution, notes the author. [Source: Shutterstock]

unique to each application, designers will probably never have a one-size-fits-all solution. From the silicon manufacturer's view, developing a bridge IC for a singular use case is typically not cost effective. So it's difficult to find silicon for non-typical applications and to develop the IP and architectural knowledge needed both from a historical point of view and looking forward.

What designers need is a single programmable video bridge that's capable of supporting multiple uses. Imagine, for example, the same programmable device that could serve as a camera interface bridge between the image sensor and the applications processor or as a display interface bridge between the applications processor and a display.

PROGRAMMABLE BRIDGING

What form would this device take? While FPGAs have been used in non-typical bridging applications, they have never been able to meet the low cost of an Application Specific Standard Part (ASSP). Ideally designers need a device that combines the design flexibility and fast time-to-market of an FPGA with the power and functional optimization of an ASSP. Whatever form this device takes, it must provide an end-to-end solution for common use cases. To achieve that, it must supply the IP the market requires. For display applications this might include:

- DSI-to-DSI
- DSI-to-dual DSI
- Dual DSI to dual DSI
- DSI to LVDS
- DSI to DPI
- Single LVDS to DSI
- DPI to DSI

On the camera side, it should support IP for bridging applications including:

- Dual CSI-2 to CSI-2
- CSI-2 to parallel CMOS
- Parallel CMOS to CSI-2
- SubLVDS to CSI-2

This programmable bridging solution would ideally support configurable data types, configurable numbers of data lanes and PHYs, adjustable input and output frequencies, DCS command ROM and IP customization. Finally, any bridging solution of this type cannot force designers to pay a penalty for using programmable logic. It must offer those benefits in a form factor that is competitive in terms of power, size, and cost compared to other silicon solutions on the market. This includes wafer-scale packaging and larger pitch packages to support industrial and automotive applications.

WHY "JUST ENOUGH" WORKS

Lattice has recently released a device that offers a variety of interfaces, some programmable and some hard, surrounding a programmable FPGA fabric. By offering just enough FPGA fabric and RAM, it can supply all the logic and memory designers need to support muxing, merging, demuxing, arbitration, splitting, data conversion, and other key functions.

This programmable ASSP (pASSP) comes with all the required IP to cover common use cases and allows the designer to derive other capabilities from there. System developers can use different firmware loads to most efficiently leverage the programmable logic and memory on-chip. By selecting which interfaces are needed and which are not, as well as I/O speed, the designer can generate bridges for multiple use cases using the same programmable device.

Clearly bridging solutions have played a key role in innovation that extends well beyond simply resolving unexpected interface issues or enabling backward compatibility. Over the past few years bridging solutions have helped designers create entirely new use cases far beyond their original intention. Look for this trend to accelerate as designers adopt new approaches to bridging that combine the best aspects of programmable logic and ASSPs.

Grant Jennings is a Senior Strategic Planning Engineer for Lattice focused on identifying new ways to drive adoption of programmable technologies through next-generation solutions based on leading connectivity and interface standards. He has more than a decade of experience in FPGA design, ASIC prototyping and system architecture. Jennings received his bachelor's degree in electrical engineering from Iowa State University and his MBA from Texas A&M University.

More Electric Aircraft Movement Looks to Innovative Technologies: Q&A with Microsemi

Aerospace applications are benefiting from higher processing power FPGAs and the steps taken to optimize power core modules.

By Anne Fisher, Managing Editor

Shane O'Donnell

Jim Aralis

You may be a jaded traveler by now, but the early pilots, and first-time pilots today no doubt, find becoming airborne an electrifying experience. Now, a push is on to tie flight and electricity together in a more literal sense. More Electric Aircraft (MEA) will rely on advances in technology to strengthen aviation performance and reliability, among other benefits.

Figure 1: Tailoring unit construction to meet the needs of an aileron system, according to whether that system is part of a wide- or narrow-body aircraft, is one example of useful PCM customization.

Microsemi is one of the companies making the technology improvements the MEA movement demands, including with its announcement last month that it had introduced an aerospace power core module (PCM) designed to optimize power electronics. Our thanks to Jim Aralis and Shane O'Donnell of Microsemi, who responded to our questions about the new PCM and MEA goals following the announcement. Aralis is Microsemi's CTO. O'Donnell is Aerospace Product Development & Technology Manager, Space & Power Management Group, Microsemi. He is also manager for Microsemi's Aviation Center of Excellence.

EECatalog: What are three steps Microsemi took to make the new power module construction innovative?

Shane O'Donnell, Microsemi: Three of the steps involved in making our new Aerospace Power Core Module innovative while optimizing for power density, performance, reliability and weight are:

1. The product architecture has been designed for aerospace applications. The power substrate, containing the power SiC MOSFETs and diodes, is integrated seamlessly with the driver and controller printed circuit board assemblies, resulting in a low-weight solution with excellent electrical, mechanical and thermal properties. Furthermore, the architecture facilitates numerous customization options, such as adding power capability or functionality without increasing the module sizing.

2. Microsemi has selected materials and components to achieve the highest performance and smallest form factor while designing for reliability, cost and manufacturability. For example, the PCM contains an AlSiC baseplate, SiC MOSFETs and an FPGA within a 5 kVA module, which weighs less than 300g, and with an MTBF in excess of 300,000 hours in a harsh environment.

3. Using advanced electrical, thermal and system analysis, Microsemi can customize the PCM, not only for different aerospace applications, but for the mission of the aircraft. For example, Microsemi can differentiate the performance of the PCM used in an aileron system on narrow or wide body aircraft and adjust the unit construction accordingly to meet the customer's requirements.

EECatalog: One topic of the MEA conference last month was the need for aerospace manufacturers to change their thinking for electric aviation. What are

your thoughts on this and how can Microsemi's knowledge and expertise guide this change in thinking?

O'Donnell, Microsemi: Power Electronics is an enabling technology in MEA. The move toward more electric aircraft (MEA) continues to drive demand for increased levels of performance, reliability and integration in the area of power electronics. When the reliability has been proven, traditional aerospace manufacturers need to consider the adoption of new technologies such as wide bandgap (WBG) semiconductors instead of their silicon alternatives, with strong heritage, due to the significant gains in performance, size, weight and reliability which can be obtained.

Microsemi has been supplying products into high reliability application areas such as commercial aviation, defense and space for over 55 years. The company's high reliability philosophy, together with our technological expertise in semiconductors, has resulted in the development of wide bandgap devices such as silicon-carbide (SiC) MOSFETs and SiC diodes suitable for MEA.

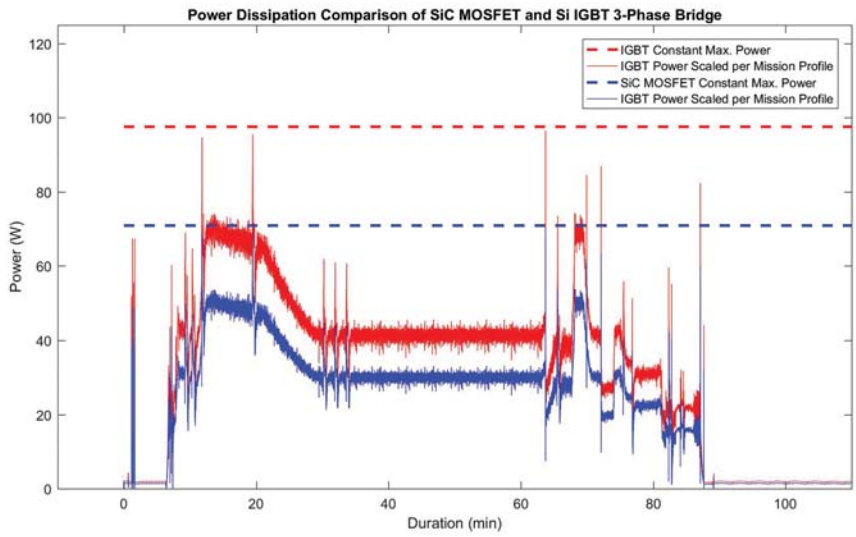

Figure 2

While aerospace manufacturers may traditionally be focused on low-cost and low-weight solutions, a more comprehensive and integrated systems perspective deserves consideration. Although SiC MOSFETs are more expensive than silicon IGBTs, their improved power density, higher switching speed capability and lower switching losses can reduce component size and system weight and improve reliability.

Microsemi's Power Core Module integrates the power switches, driver circuitry and high-speed communications interface in a single package. The replacement of Si IGBTs with SiC MOSFETs results in a power dissipation improvement of approximately 25%, with a resultant system gain in thermal performance (Figure 2).

EECatalog: What is it important to know today about FPGA development environments and the options designers have?

Jim Aralis, Microsemi: The density of course has gone dramatically up, so the things that we can do with FPGAs are much more impressive. We used to relegate them to doing things that were slow or were not too big or were specialized control functions, as opposed to being in data paths with data processing elements.

So the performance of FPGAs is getting to the point where a lot of things that could not be done by them before they can now do. FPGAs are coprocessors in data centers and servers, as well as taking on such tasks as voice recognition and server applications for doing baseband generation. Baseband virtual environments are being enabled because FPGAs cannot only process the data at high data processing rates and complexities, they can also be retargeted to new applications on an ongoing basis, which, again, is another new way of people using FPGAs, not just in embedded systems, but as well in systems where the complexity is a level above that, meaning servers, basestations, and things of that nature.

At Microsemi, we represent FPGAs that are for whole systems. We have nonvolatile memory for use inside, but also for programmability, meaning you don't have this high rate of interface to the outside world to get your programming materials—it can be a standalone part and look just like a part you buy that has already been programmed, as opposed to the system-level stuff that has to be done with other people's FPGAs.

We are seeing a high level of security. The security that is being built into the FPGAs now— both ours and our competitors—to enable network security as well as data security is unprecedented even in dedicated ASICs because these algorithms and these capabilities are evolving at a very high rate.

FPGAs have core capabilities like [physically unclonable functions] PUF, which are unique functions that identify each individual FPGA. We have very high level elliptic curve cryptography. We have a lot of very sophisticated blocks that build security capability, but then you put them together with an FPGA, where you structure an architecture security for your applications separately, and they become very difficult for somebody to go back and decode or break. Also, the economic benefit of breaking it lessens because the application is not all-encompassing.

Also pretty far along and getting even more important is the system on a chip inside an FPGA. Processors, memories, multiprocessors, and digital signal processing is now inside chips. We're also putting in specialized accelerators and data accelerators that can be used with the fabric as well. There is a fundamental change coming with our next generation of FPGAs, and probably with the next generation of our competitors as well: you almost can see it as an embedded FPGA because so many capabilities are hard wired in because of the evolution of SoCs and systems in general.

To take what I've just mentioned a little further, we build a lot of SoCs for storage, for networking that are not FPGAs, and 80 to 85 percent of those circuits are actually IP—meaning they are almost common between something even as disparate as a PCIe controller, memory controller and switch—they have very similar elements in them—the only difference is how you hook them up and some of the dedicated logics. So you end up with an FPGA in which 80 percent of its functionality is in the number of applications processors, the number of DSPs, the number of controllers, memory, and then some dedicated logic, which can be implemented in an FPGA. You will find an FPGA that will compete directly with an ASIC in those areas because not too much of it has to be programmable in a logic sense. It's programmed by micro-code and even Linux code directly.

We are going to run out of steam in process technology somewhere at four or five nanometers, which is one generation from where we are doing our designs right now at seven [nanometers]. When that happens the complexity of these chips is going to be so great that the implementation methodology of preference will be processors and FPGAs.

It costs too much to build a custom chip for every application, and nobody will be able to do it. So the people who know how to mix processors and FPGAs will be the people who will win the application space for doing economical application-specific SoCs.

EECatalog: What does greater processing power mean for the More Electric Aircraft movement?

Aralis, Microsemi: Much more analysis of the data can occur in real time, and we can make adjustments. In the past we had to take data from, let's say, a jet engine, and use that data (very good data, "big data," as everyone calls it) to optimize the jet engine for the general case of when you took the data with maybe a couple of points that say: this is takeoff configuration; this is high cruise; this is low cruise; this is landing—and so you could configure the jet engines to have performance at certain different areas. Say it was cars rather than jet engines, you would have: driving offroad; driving onroad; driving slippery.

But now we're seeing the advent of this capability of very, very high processing and that of deep learning, the idea of using learned data in order to do things. We combine those things now to monitor jet engines, for example, and learn how to dynamically change the engine's performance. We can do the equivalent of all that data analysis that in the past took place around one week after data collection. But now rather than doing the analysis a week later, it's done real time—so it is no longer low cruise/high cruise—it is wherever you are. It's optimizing the engine on a very quick cycle so that you can have it optimized no matter where you are in the cycle.

We couple that with the ability to do things, even acoustic vibration, [both] sub-acoustic and super-acoustic where we look at what is going on in the engine in order to predict its reliability. The work we are doing with some of our partners is to make aircraft more reliable, and cheaper.

An aircraft is one of the few places where if you make something last longer, you actually reduce the costs—If you can make an engine last twice as long, you cut its cost in less than half because you don't have to replace it and you can use it twice as long. It is not like a pair of shoes or even an iPhone that goes out of style. If you double in time the length that it works, you get less than half of the cost.

Our heuristics can predict closer to where it is going to fail. If you can do that you can keep the engine in the airplane much longer and even at a safer level to cut the costs. You can cut it in half. Our goal is to build these systems with twice the MTBF and with that you have actually cut the thing in less than half cost of operation.

So those are the kinds of things we apply directly to aircraft in order to increase reliability, and lower their cost as well as their weight.

As well as seeing to it that when you do have to change something, you do it at a time and place that is much cheaper. You don't want your planes to show failure in a remote location—it takes a lot of money to get the parts and the people there.

Design Advantages of Programmable, Efficient SoC Architectures

Flexible interfacing, hardware reconfigurability, and improved power efficiency are among the benefits programmable SoCs offer.

By Nishant Mittal, Cypress Semiconductor

System-on-Chip (SoC) architectures are an innovative technology designed to miniaturize devices for integrated systems, where an entire system resides on a single chip. This can halve board size and enable systems to have plug and play interfaces, further simplifying design. For example, external sensing devices can directly connect to the chip. The chip is capable of driving the pre-processing, converting and post-processing of signals. In addition to lowering system cost, an SoC can improve power efficiency and overall system speed by reducing wire delays caused by PCB tracks and parasitic capacitances.

SoC devices started becoming popular when miniature devices like smartphones came onto the market. With each new generation, SoCs have gotten thinner and smaller, while integrating more advanced technologies like gesture recognition, fingerprint sensors, touch-screen capabilities and health monitoring sensors directly into the device. As the number of interfaces increased, so did the need for the miniaturization of devices at the board level. This triggered the need of advanced SoCs to also replace active and passive components.

SYSTEM-ON-CHIP ARCHITECTURE

While SoCs may look simple on paper, they have complex architectures. A system-on-chip follows the basic principle of incorporating all the drivers, signal conditioning circuits and digital communication blocks inside the chip, so developers can implement most of the electronics of a product using a single chip. As shown in Figure 1, a basic SoC should have an integrated ADC, PWM, CAN, UART, JTAG, TWI, Ethernet, Timer/Counter and the core CPU, which controls all the operations.

Consider interfacing an LM35 temperature sensor with an 8051 microcontroller. The V_{OUT} of the temperature sensor cannot be directly connected to the microcontroller. Instead, the signal must be amplified using an op amp, and a suitable analog-to-digital converter (ADC) is needed to convert the analog signal into a digital one. This creates quite a complex system for just a temperature sensor application. With an SoC, it is possible to directly interface the temperature sensor (or any other sensor) to the SoC without requiring external circuits apart from some few biasing resistors as external circuits.

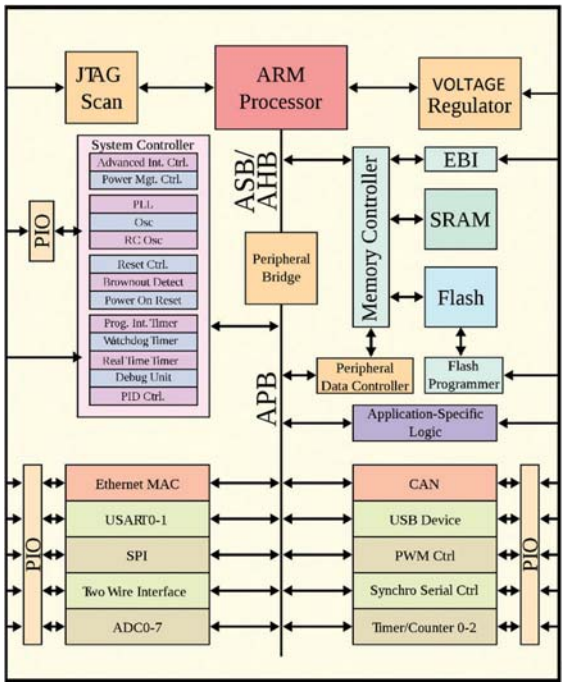

Figure 1: General SoC Architecture (Source: Wikipedia)

Some SoCs are able to interface to high current requirement components such as motors or servo/stepper motors. Rather than require an external driver circuitry to enable the motor to interface to the SoC, this driver circuitry can be integrated into the SoC, greatly simplifying design.

Internally, the integrated complexity of SoC architectures introduces challenges because of power consumption, speed optimization and noise isolation. The analog section needs to be isolated from the digital section to reduce cross talk. Special care needs to be taken while routing from high-speed lines to maintain signal integrity. A proper clocking scheme, like that of

the H clocking scheme shown in Figure 2, should be used to prevent improper clock timing in different SoC areas.

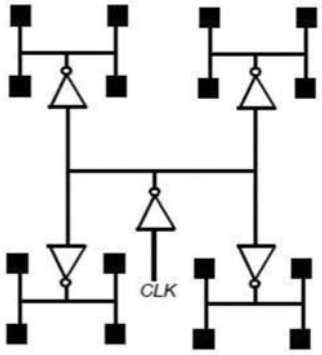

Figure 2: H Clocking Scheme

As the available surface area continues to shrink, power consumption becomes a major concern. The power being carried by the nm tracks in the SoC must be properly optimized to consume less power. Thus, designing the system for power monitoring and optimization becomes important, preventing power issues from making the system unusable.

Many SoCs are available today, from FPGA-based SoCs (e.g. switching-based designs) to RTOS SoCs and ASIC SoCs, as well as programmable SoCs. This last type of SoC offers advantages such as flexible interfacing, power performance and reconfigurability of hardware.

PROGRAMMABLE SOCS

Programmable SoCs, such as the PSoC family from Cypress Semiconductor, provide efficiency through software as well as hardware programmability. A development environment, such as PSoC Creator, manages programmability for both software and hardware.

Figure 3 shows the architecture of a programmable SoC. Integrated peripherals include a huge variety of ADCs, timers, counters, PWM, CAN, I2C, SPI, WCO, DMA, capacitive sensing, DAC, and programmable analog, among others. DACs are especially difficult to integrate, as they pose many noise and signal power optimization challenges. However, DACs can be highly efficient with a very good SNR and bandwidth when operating at a few Mbps data rate, which is more than enough for sensor and automation applications. To fill out the signal conversion components, a programmable SoC can include Delta Sigma ADC, as well as SAR ADCs.

Programmability gives developers the flexibility to trade off performance against power consumption. For

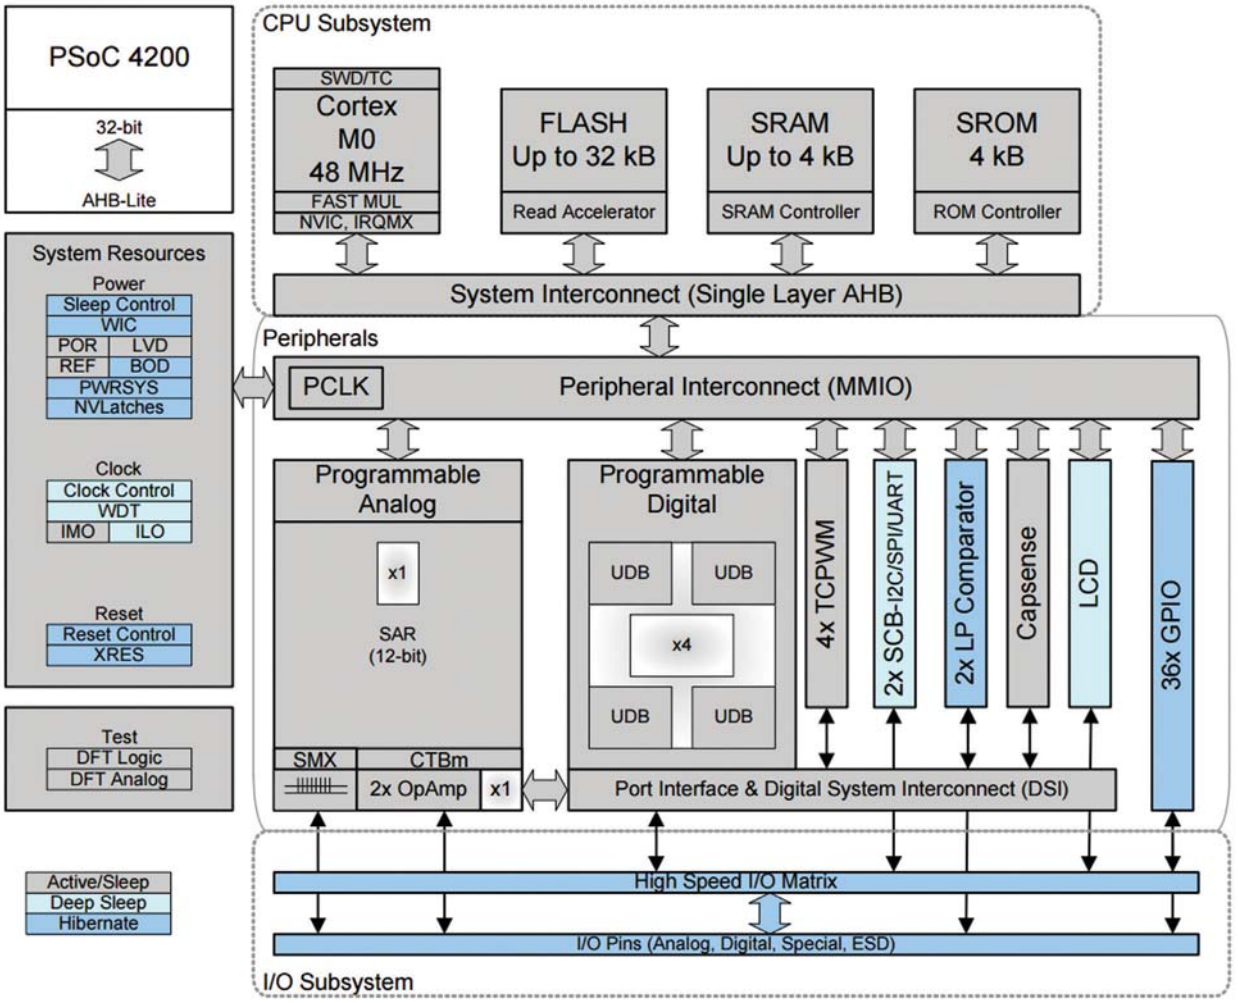

Figure 3: Programmable System-on-Chip Architecture (Cypress PSoC) [Source: Cypress Semiconductor]

example, configurable resolution, data rate and operating modes like continuous or single shot. An integrated capacitive sensing block makes the world of touch sensing available. Capacitive touch sensors, such as sliders and wheels, can be connected directly to the SoC without any external circuitry, enabling a small footprint for wearable and smartphone applications. Even a complex system like digital audio players and automotive controllers can be built around a programmable SoC.

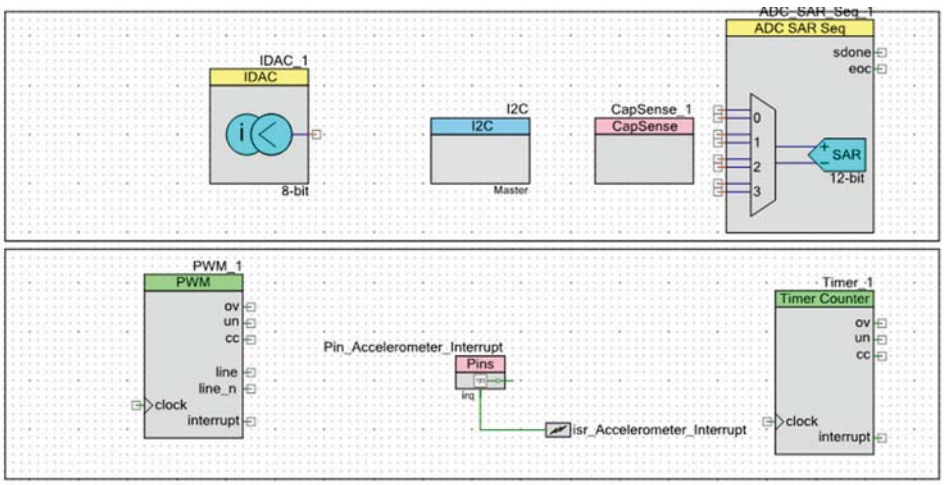

Figure 4: PSoC Creator Component Declaration User Interface

The programmable hardware of a PSoC uses universal digital blocks (UDBs), allowing developers to create various digital designs using hardware like D flip-flops and combinational components. This allows reconfigurable functionality to be implemented in hardware without writing a code. Blocks can be designed using an intuitive state machine or Verilog codes.

Today, software design environments for SoCs need to correlate program code with hardware components. With a drag and drop GUI, developers can quickly "connect" software with hardware, then the tool generates the APIs and internal connection binaries for easily interfacing a component.

Consider interfacing a temperature sensor with a programmable SoC chip compared to the 8051 example above. The lines of code required to implement the same algorithm drops by half with the internally configurable components of a programmable SoC. As shown in Figure 4, many components like ADC, DAC, Temperature sensor blocks, I2C, etc. are available in a drag and drop GUI. The developer configures the block and writes code using this block to connect to general purpose I/O (GPIO) to perform a specific application. In addition, the developer does not need to code the clock synchronization of the ADC. Instead, dropping in the component causes the tool to generate the appropriate APIs, which can then be coded to start and initiate the component.

A program can be completed in just a few lines of code. Note that developers have the flexibility to go further into the code if they desire and use advanced APIs for more complex operations like interrupt control or multiplexing.

ANALOG PROGRAMMABILITY

In addition, some programmable SoCs offer analog hardware programmability. The basic analog component is the op amp. Op amps are configurable for power, input/output routing and gain-bandwidth. They can be used as instrumentation amplifiers or as a buffer using the components. Because SoCs are designed with power as a primary constraint, low current capabilities combined with sleep and deep sleep features enable the SoC to use the least current when the CPU is not being used.

As described earlier, analog capacitive touch sensors or any other capacitive sensors can be directly interfaced to capacitive sensing capable I/O of some SoCs. For example, with PSoC, developers can select a capacitive sensing component in the development environment to easily add touch sensing to their applications.

Another important component of a programmable SoC architecture is GPIO. GPIO connects the system to the external world. A Cypress PSoC provides various categories of I/Os to interface with, including strong drive, pull up and pull down, and high impedance to address various current requirements and circuitry of the external world. In addition, a hot swap feature in the GPIO enables users to prevent input from being clamped to the I/O power supply level when the input voltage is above the I/O supply voltage.

Integrated PWMs and timer counter blocks in a programmable SoC architecture enable developers to more easily write complex timer operations through the configuration of TCPWM blocks as shown in Figure 4. By doing some simple calculations to understand the clock frequency and the compare signal, developers can adjust the duty cycle of a PWM without writing any code, making it simple to configure PWM to the requirements of a specific application.

Finally, many other communication blocks are available. For example, PSoC includes I2C, USB, CAN, SPI, and TWD with JTAG (for debugging). This makes programmable SoCs capable of addressing the latest trending products in the market at a low cost.

Nishant Mittal is a Systems Engineer at Cypress Semiconductor, Bangalore. He received his Master of Technology degree with a specialization in Electronic Systems from IIT Bombay in Mumbai, India.

Acromag, Inc.

AcroPack APA7-200

Supported Products: Xilinx Artix-7

New APA7-200 AcroPack® series provides a user-configurable Xilinx® Artix®-7 FPGA on an AcroPack mezzanine module for the best value for a variety of cost- and power-sensitive applications. This next generation I/O module offers a PCIe interface format suitable for VPX and ruggedized PCIe-based computers. The module plugs into a connector on an AcroPack carrier with no cables needed. The APA7-200 Series offers a choice of four models all offer high channel count digital interface with RS485/422, RS-485 and RS422, LVDS, and TTL interface options. Mix and match I/O combinations in a single slot for embedded applications running on Linux®, Windows®, or VxWorks® operating systems.

The APA7-200 series modules are 70mm long. This is 19.05mm longer than the full-length mini-PCIe card at 50.95mm. The boards width is the same as mPCIe board of 30mm and they use the same mPCIe standard board hold down standoff and screw keep out areas.

A down facing 100 pin Samtec connector mates with the carrier card. Fifty of these pins are available for field I/O signals.

The Engineering Design Kit provides users with basic information required to develop custom FPGA firmware for download to the Xilinx FPGA. Example FPGA design code is provided as a Vivado IP Integrator project for functions such as a one-lane PCI Express interface, DMA, digital I/O control register, and more.

Designed for COTS applications and SWaP-C optimized these FPGA-based digital I/O modules deliver user-customizable I/O, high-density, high-reliability, and high-performance at a low cost.

FEATURES & BENEFITS

- PCI Express Generation 1 interface
- Reconfigurable Xilinx® Artix-7 FPGA
- High channel count digital interface: RS485, LVDS and TTL interface options.
- 32Mb quad serial Flash memory
- 33,280 logic cells

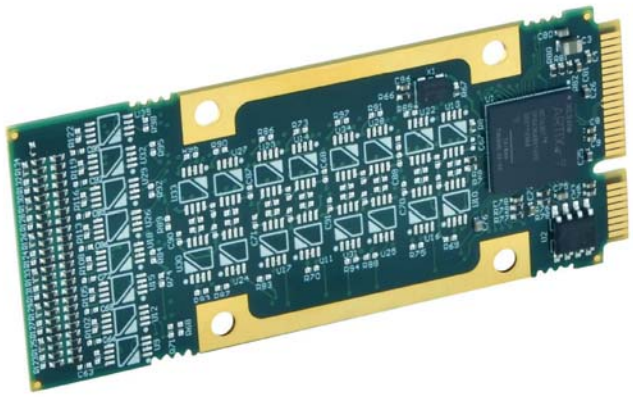

TECHNICAL SPECS

- Conduction-cooled operating temperature: -40 to 85 deg.C.
- Power: +3.3V (+/-5%).
- Example FPGA program included in EDK kit.
- 41,600 Flip flops
- 1,800 kb block RAM

AVAILABILITY

Available now for on-line ordering or call 248-295-0310.

APPLICATION AREAS

Aerospace/Defense, Data Processing and Storage, Industrial Automation, Wired Communcations, Scientific, Automation, and Development Labs

CONTACT INFORMATION

Acromag, Inc.
30765 S. Wixom Rd.
Wixom, MI 48393
USA
Tel: (248)295-0310
solution@acromag.com
www.acromag.com

Pentek

Model 71861 4-Channel 200 MHz A/D with DDC, Kintex UltraScale FPGA - XMC Module

Supported Products: Kintex UltraScale

Jade is based on the Xilinx flagship Kintex UltraScale FPGA, which significantly raises the digital signal processing (DSP) performance by over 50%. Equally impressive are reductions in cost by 39% and power dissipation by 18%.

The Navigator™ Design Suite was designed from the ground up to work with Pentek's new Jade architecture and provides an unparalleled plug-and-play solution to the complex task of IP and control software creation and compatibility. The Navigator Design Suite consists of two components: Navigator FDK (FPGA Design Kit) for integrating custom IP into Pentek sourced designs and Navigator BSP (Board Support Package) for creating host applications. Users are able to work efficiently at the API level for software development and with an intuitive graphical interface for IP design.

TECHNICAL SPECS

- Complete radar and software radio interface solution
- Supports Xilinx Kintex UltraScale FPGAs
- Four 200 MHz 16-bit A/Ds
- 3U & 6U VPX, AMC, 3U & 6U cPCI, PCIe, ruggedized & conduction-cooled versions available Model 78861 - PCIe
- Model 8266 SPARK Development System for Jade (Xilinx UltraScale), Flexor (FMC), Onyx (Xilinx Virtex-7) and Cobalt (Xilinx Virtex-6) PCIe Boards

CONTACT INFORMATION

PENTEK

Pentek
One Park Way
Upper Saddle River, NJ 07458
USA
Tel: (201)818-5900
Fax: (201)818-5904
sales@pentek.com
www.pentek.com/go/maebjade

The Premier Conferences Devoted to Implementing IoT Technology in Embedded Systems

INTERNET OF THINGS DEVELOPERS CONFERENCE

The Big Event: April 26-27, 2017 | Santa Clara Convention Center

IoT Device Security Summit | September 14, 2017 | Santa Clara, CA

Connected Devices & Gateways Summit | November 30, 2017 | Santa Clara, CA

Plan Now to Attend

www.iot-devcon.com

Xilinx Virtex-7 FPGA-Based Modules

Ideal for Customizable, High-Bandwidth, Rapid Signal-Processing Applications and Available with a Comprehensive Firmware Development Kit to Accelerate Application Development

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

MADE IN THE USA

Designed, manufactured, and supported in the USA