



WP493 (v1.0) September 6, 2017

Key Attributes of an Intelligent IIoT Edge Platform

By: Chetan Khona

Xilinx® All Programmable SoCs and 7 series FPGAs provide the widest breadth of capabilities for Industrial Internet of Things (IIoT) platforms today and offer maximum flexibility for the future—enabling the highest return on investment and lowest total cost of ownership over the life cycle of industrial systems.

ABSTRACT

Using the examples of industrial communications, cybersecurity, and edge compute (from simple data optimization to machine learning), this white paper focuses on the fit and benefit of Zynq®-7000 SoC and Zynq UltraScale+™ MPSoC families for IIoT embedded systems. These devices combine ARM® application processors with FPGA logic (programmable hardware), peripherals, and other embedded blocks to enable users to strike the ideal balance between software intelligence and hardware optimization for their systems. The white paper explores industrial equipment life cycles and how the combination of software and programmable hardware facilitates a fundamentally more capable system in the short term, but how it also extends the usable life of the system within the rapidly shifting market trends of IIoT. The white paper introduces software tools that facilitate the distribution of a function between software and programmable hardware and highlights the business risks and costs of not choosing an All Programmable solution.

Approaches to IT-OT Convergence

Industrial IoT (IIoT) refers to a multidimensional and tightly coupled chain of systems involving edge devices, cloud applications, sensors, algorithms, safety, security, vast protocol libraries, human-machine interface (HMI), and other elements that must interoperate. Some describe the vision of IIoT as the convergence of operational technology (OT) and information technology (IT), but the goal is actually more significant. The time-sensitive nature of OT applications and the data-intensive nature of IT applications require all these elements to come together and perform critical tasks *reliably and on schedule*. When coupled with another fundamental requirement—longevity—conflict can arise. Longevity ensures return on investment of these IIoT systems for both the systems supplier and its customers. Significant advances are being made in a number of underlying areas that are foundational to IIoT systems, such as analytics, machine learning, cybersecurity, and others. However, when modifications or upgrades are made over an extended life cycle, the tight integration requirements can create unwanted ripple effects to the time-critical nature of these systems.

The most common preemptive response to this challenge is to seek embedded electronics, the heart of IIoT edge systems, which offer the best specifications available to create margin for the unknown. Edge systems are the deterministic embedded communication and real-time control engines that reside at the edge of the network and closest to the physical world of factories and other industrial environments, e.g., motion controllers, protection relays, programmable logic controllers, and similar systems. Clock frequencies in gigahertz, larger memory sizes, higher numbers of input/output ports, and the latest encryption engines might seem to offer solutions for future requirements that are as yet unknown. However, when dealing with the timescale of industrial equipment, which has critical subsystems that operate on a scale of hundreds of microseconds (or less) but need to operate in factories and remote locations for decades, relying solely on a cutting-edge multicore embedded processor to scale in the IIoT space is, at best, unimaginative. At worst, it is a short-sighted catalyst leading to a series of difficult and costly marketing and engineering trade-offs focused on managing functional timing issues stemming from performance bottlenecks. A much higher degree of freedom in scaling is desperately needed at the IIoT edge due to the timescales involved. Such scaling freedom can be unlocked by using programmable hardware that augments the software running on the embedded processor cores. This is a more consistent approach that allows determinism, latency, and performance to be easily managed and eliminates interference between the IT and OT domains and within subsystems in the OT domain.


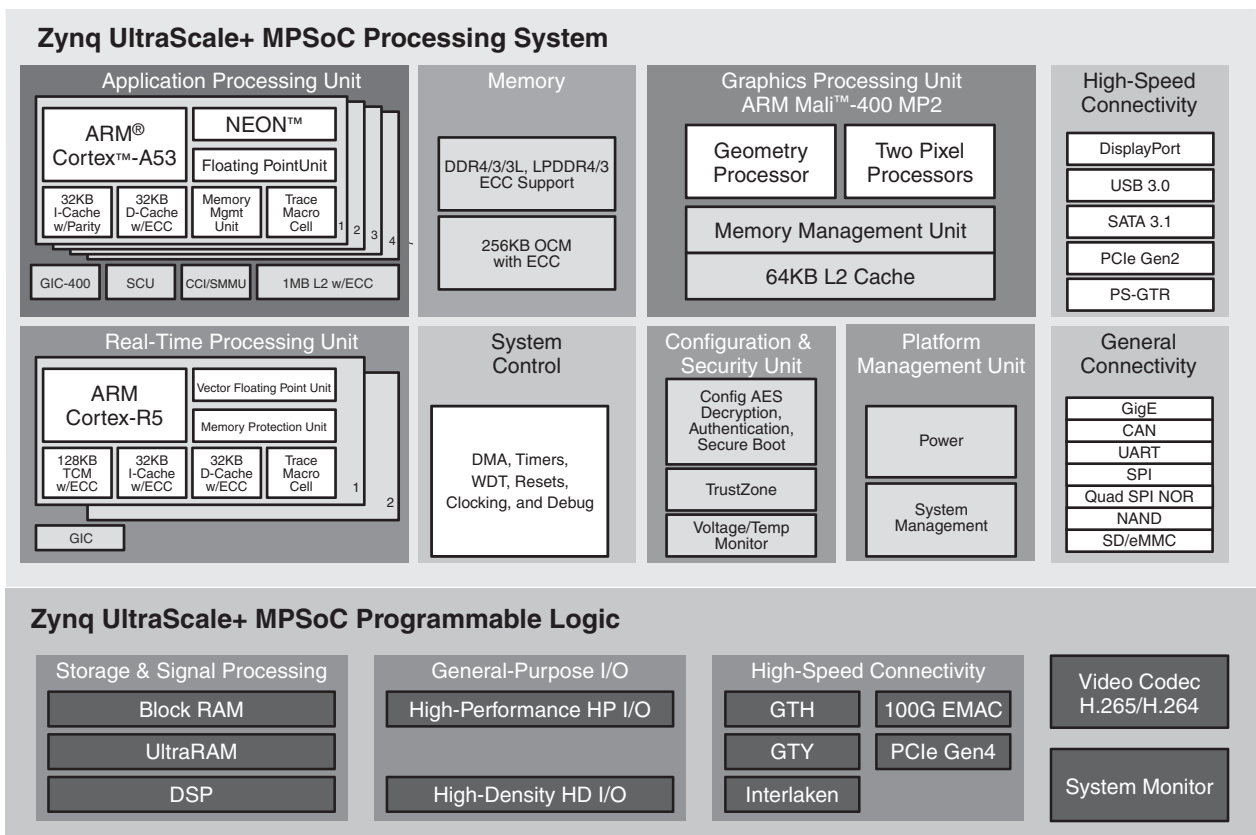


Figure 1: Timescale, Attributes, and Life Cycle of Typical IIoT Systems

In parallel, sustainable value is available with processors that provide features like hardware virtualization, which allows architects to incorporate new guest operating systems and provides levels of autonomy and isolation where needed. Consistently useful features are also available, such as memory protections (parity, or—preferably—error-correction code [ECC]) that likely will

never outlive their usefulness. Augmenting static processor architectures with specialized hardware to create a division of labor that is balanced and ideally suited for the pending tasks is not a new paradigm for the embedded electronics world. Garnering more attention is the need to adapt both the tasks and the division of labor itself over time. An example is a new predictive maintenance algorithm that requires more sensor inputs than previous inputs. Offloading the incremental calculations to hardware maintains the overall loading and—most importantly—the cycle time of the processor subsystem. This flexibility can pay significant dividends for both the customers who purchase and install the system and the systems suppliers who can extract multiple value-added software-based service revenue streams from that equipment for decades in the future.

This white paper examines three key applications areas that comprise the foundation of IIoT—connectivity, cybersecurity, and edge compute—within the context of selecting an IIoT edge platform that can adapt to the impact of market trends over time. It is vital to have an IIoT platform that is extraordinarily flexible, scalable, and equally capable of dealing with both OT and IT technologies. An All Programmable SoC, i.e., a system-on-chip that is hardware and software programmable, is the ideal solution. This white paper also covers two technology topics that are relevant to All Programmable SoCs: software-defined hardware and All Programmable SoCs vs. FPGAs companions to discrete embedded processors. Xilinx offers the Zynq-7000 SoC and Zynq UltraScale+ MPSoC families, which uniquely offer comprehensive coverage for IT and OT tasks. See Figure 2.



WP493_02_042517

Figure 2: Zynq UltraScale+ Block Diagram

Connectivity: From Legacy Standards to Future Protocols

Connectivity in the age of IIoT is moving toward a streamlined approach, but this transition introduces new complexities. Edge and system-wide protocols like the OPC Foundation Open Platform Communications-Unified Architecture (OPC-UA) and Data Distribution Service for Real-Time Systems (DDS) are gaining significant momentum in their respective application areas. Both benefit from the emergence of time-sensitive networking (TSN), a deterministic Ethernet-based transport that can manage mixed criticality streams. TSN significantly enables the vision of a unified network protocol across the edge and throughout the majority of the IIoT solution chain, because it supports varying degrees of scheduled traffic alongside best-effort traffic. TSN is an evolving standard, and dedicated chipsets (e.g., ASIC or ASSP) advertising standards compliance before all aspects of the standard and the endmarket-specific profiles are finalized is fraught with risk. In a similar way, attempting to add support for TSN to an existing controller that manages real-time data via a purely software-based approach might result in unpredictable timing behavior—at best. The likely result is the degradation of interrupt responsiveness, memory access timing, etc. Ultimately, this is not a reasonable solution, because TSN requires a form of time-awareness not in controllers today. Even if an external TSN switch is added to the system, without *integrated* TSN in the same device (to manage control functions, e.g., the Endpoint), the switch connecting the various Endpoints will likely produce backward Ethernet compatibility support for a non-TSN enabled controller. The goal is to get TSN integrated into the Endpoint to enable scheduled traffic versus best-effort traffic with a minimum impact to control function timing. See [Figure 3](#).

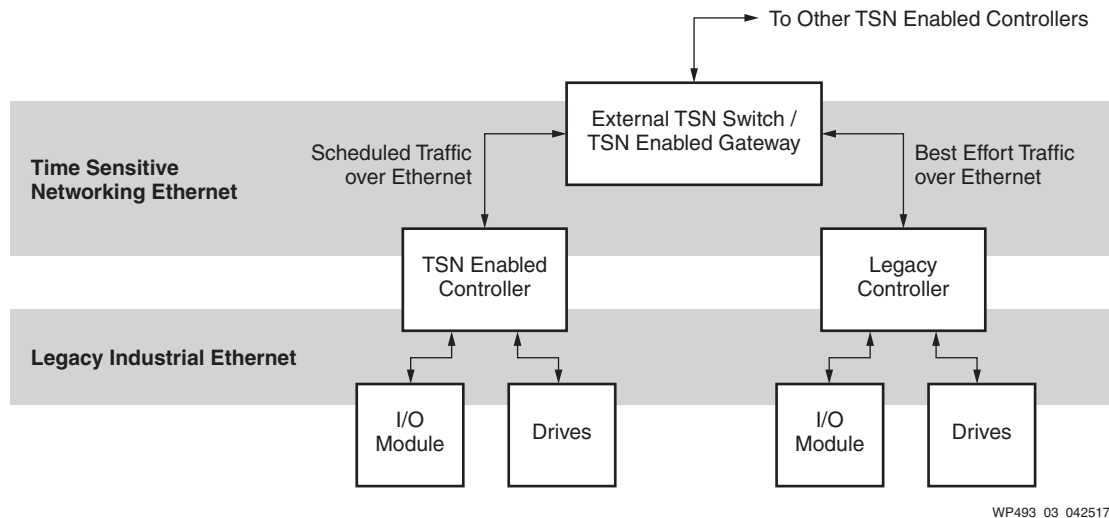


Figure 3: TSN Topologies and Benefits

Integrating an All Programmable implementation of TSN in the controller minimizes the effects of the change by enabling the implementation of bandwidth-intensive, time-critical functions in hardware—without significant impact to the software timing. The designer can implement pure Endpoints or bridged Endpoints by using Xilinx’s internally developed, fully standards-compatible, and optimized implementation of TSN. Whether upgrading a controller that is designed with an All Programmable SoC from standard Ethernet to TSN or designing a new controller with the evolving TSN standard, Xilinx’s All Programmable approach enables the designer to make changes with the least impact to critical timing and is future proofed (vs. ASICs and ASSPs).

An alternate but equally common use case is also worth considering. Because IIoT is not a new industry, it still needs to support the lengthy list of legacy industrial protocols that have been in use throughout the industry's fragmented past and present. Most modern SoCs do not offer support for even a sizable fraction of these protocols. Also, the number of network interfaces can exceed the I/O capabilities of most of these fixed SoCs. In contrast, Xilinx's All Programmable SoCs enable the creation of a system that can withstand customer-specific customization, such as support for legacy protocols and their associated I/O connectivity. Whether the protocol requires a 250 μ s or 64 μ s cycle time, the fully encapsulated and hardware-offloaded implementation of these industrial communication controllers eliminates the cost of additional devices, without causing the side effects to mainstream software and firmware that a software-based approach might cause. Whether with TSN, with legacy industrial protocols, or the most likely scenario—a mix of both bridging the past and future, Xilinx offers any-to-any connectivity that is deterministic by design.

Cybersecurity: Hardened and Adaptable to Future Threats

IIoT thought leaders employ a “defense-in-depth” approach to the broad topic of cybersecurity. Defense in depth is a form of multilayered security that starts at the supply chain of suppliers and reaches the end customers' enterprise and cloud application software, even extending to the things to which that software might connect. In this section, the scope is the chain of trust for deployed embedded electronics at the IIoT Edge. With the network extending to the analog-digital boundary, data needs to be secured as soon as it enters the digital domain. Defense-in-depth security requires a strong hardware root of trust that enables secure and measured boot operations, run-time security through isolation of hardware, operating systems, and software, and secure communications. Independent validation of credentials through trusted remote attestation servers, certificate authorities, and so forth should be employed throughout the chain.

With cybersecurity attacks expected to become more frequent, security is not a static proposition but an ever-evolving one. For example, since 1995, five notable revisions were made to the transport layer security (TLS) secure messaging protocol, with more to come. IIoT system suppliers and their customers need to know how to mitigate security risks that evolve over time while maximizing the life and utility of costly assets. The cryptographic algorithms that underscore protocols like TLS can often be implemented in software, but with the move toward IT-OT convergence, these changes on the IT side can create adverse effects on time-critical OT performance. To reduce this impact, some software architectural tools such as hypervisors and other isolation methods are available. It is possible to pair these software concepts with the ability to offload and support new, currently undefined cryptographic functionality that uses programmable hardware years after product field deployment. This approach provides a stronger risk mitigation plan and might avoid costly recalls, patches, and the potential threat of litigation.

Software-Defined Hardware

Hardware offload, as mentioned in the [Cybersecurity: Hardened and Adaptable to Future Threats](#) section, is not just supported in the programmable hardware of an All Programmable SoC. Achieving the full vision requires software automation that streamlines the technology. A tool like the SDSoC™ development environment enables users to write C/C++/OpenCL, among a growing list of languages, and to partition all or part of the function in programmable hardware or software. The SDSoC development environment also generates the data movement engines and infrastructure between the processor and the programmable hardware. In 2015, the SDSoC tool

was used with an advanced encryption standard (AES)-256 algorithm to demonstrate 4X improvement in performance when partially moving the algorithm to the programmable hardware. See the *Accelerate AES Encryption with SDSoC* article in the [Xcell Software Journal](#).

That benchmark focused on exploring the optimal balance of software intelligence and programmable hardware optimization. But the tool can also completely offload the function to programmable hardware as well. Similarly, motor control loop closure times via hardware acceleration engines are shown to offer 30 to 40 times the performance of a software-only implementation. See [Figure 4](#).

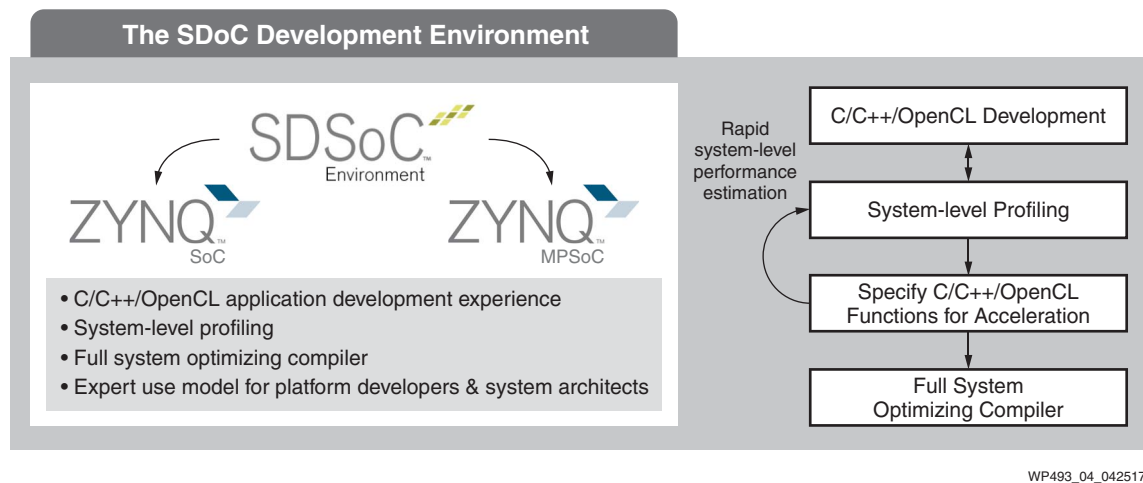


Figure 4: SDSoC Design Environment Design Flow

WP493_04_042517

Edge Compute: Scalable, Cost-Efficient, and Real-time

Just as with communications and security, edge compute is evolving in multiple directions. The computational power of the cloud, which operates on streams of unlocked data from previously inaccessible systems, provides users with actionable insights previously unseen or not understood. This creates a set of expectations, or table stakes, which serve as a new baseline. Just as relying on GPS-based navigation systems is making most highway maps obsolete, purchasers and users of industrial equipment have different expectations of feedback from their IIoT systems. Currently, the trend is to push the generation of these insights from the cloud to the edge, as driven by three primary factors:

- The desire to apply insights more quickly than possible in a round-trip from the edge to the cloud
- The cost of sending (in many cases) exorbitant amounts of data to the cloud
- Security, reliability, and privacy concerns of sending data to the cloud

These are industry trends and should not be viewed in terms of absolutes. Even merely pre-processing data locally and sending optimized, obfuscated data to the cloud can be hugely beneficial in addressing some of these security and privacy concerns. An extremely simple example is to apply a low-pass, or averaging filter, to time series data at the controller that is responsible for a machine. The result is to simultaneously reduce the number of data points that are sent to the cloud and also to suppress outlier data. Programmable hardware enables you to apply these

optimization functions to the data as it is streaming off the machine, which enables the most efficient processing of that data versus using complex memory transactions that compromise the response time of any potential decisions based on the data. This example can be stated in terms of a single data stream from a single sensor, but in actuality, most industrial systems consist of hundreds or even thousands of simultaneous data streams. The number of connections amplifies the problem and the value of the solution that is provided by programmable hardware through various sensor fusion techniques and on-chip analysis.

In the example described here, intelligence is embedded at the controller to make local adjustments for time-sensitive feedback items and push the less time-critical data to the cloud in a condensed format. This is a good example of the edge and cloud complementing each other. This description of embedded intelligence and edge-cloud cooperation can also apply to machine learning at the edge, a topic that is rapidly becoming more relevant in IIoT. Machine learning—which includes neural network-based machine learning inference and deployment, as well as classical techniques such as regression and others—is extremely well suited for the power-efficient, customizable, and massively parallel compute architecture of programmable hardware. Because of this, programmable hardware-based acceleration cards are used extensively in the cloud. The same All Programmable technology is available for use at the edge, offering the lowest latency, power, and cost for multi-sensor machine learning applications. The ability to efficiently support all the foundational aspects of the IT-OT convergence while simultaneously offering superior capability in burgeoning areas enables All Programmable technology to claim the widest IIoT application coverage in a single device. For example, combining applications like motor control, machine vision, network communications, functional safety, cybersecurity, etc., with edge analytics and machine learning is the expected use case for All Programmable technology in IIoT. By using tools like the SDSoC development environment with its supporting libraries, users can implement substantial algorithms in a fraction of the smallest All Programmable devices. See [Figure 5](#).

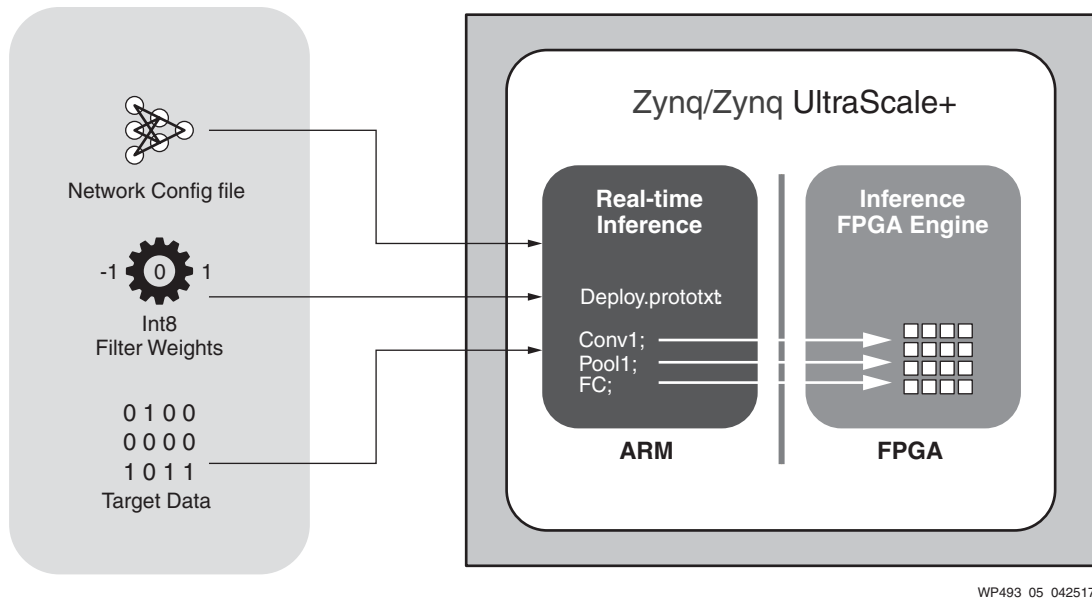


Figure 5: Machine Learning Inference Flow for Zynq-7000 and Zynq UltraScale+ SoCs

FPGA Companion for Legacy Processors

To implement the wide range of IT-OT functions, the most common choice for IIoT Edge platforms are All Programmable SoCs. These devices offer the integration benefits that were described previously, as well as power and cost savings. Another option exists for real-world situations in which a pre-existing architecture is in place, perhaps with legacy code that is tied to a legacy embedded processor. In these cases, some of the benefits described are still valid through the use of programmable hardware-only devices that are called FPGAs. FPGAs operating as a companion device are easily interfaced to the main embedded processor. These FPGAs can act as co-processors to the main embedded processor and offer the option to implement a compact microcontroller or microprocessor (such as the Xilinx MicroBlaze™ processor). These soft (built out of programmable hardware) processors support a wide range of operating systems and real-time operating systems. Offloading evolving or time-critical functions in the context of legacy systems can still be achieved by utilizing these options. Xilinx's All Programmable portfolio, which includes FPGAs and SoCs, enable long life cycles, high-availability silicon in extended temperature ranges, and the ability to reconfigure the entire device or partial areas, even during operation. Multiple FPGA options exist that share footprint compatibility, to allow for a measure of platforming. The two-chip approach lacks the high bandwidth between the processor and the FPGA, compared to an All Programmable SoC. This bandwidth and the number of connections within the monolithic SoC facilitate the dynamic hardware-software division of labor (that previous examples rely on), which a two-chip solution cannot replicate. Even with these limitations, the value of programmable hardware is large enough that more and more embedded processors are advertising dedicated FPGA interfaces (typically built out of standards such as: PCIe, SPI, QSPI) in their data sheets.

Hardware and Software Programmability for Longevity in a New Industrial Era

Industrial control systems that are driven by electric means have been available for over a century. With some referring to IIoT as the fourth Industrial Revolution, not only have the available technology and required tasks changed, but the pace of the industry in aggregate has increased in the rate of disruption. Today, new technologies are available as building blocks of IIoT Edge platforms that fundamentally offer better coverage for the expansive breadth and depth of IT-OT tasks over time. All Programmable SoCs, such as Zynq-7000 and Zynq UltraScale+ devices, use software and hardware programmability to keep assets useful for longer periods of time as compared with the traditional embedded architectural building blocks of the previous twenty or thirty years. The use of a different embedded controller for each end product without regard to their connection to the same cloud infrastructure is a failing approach, since approximately 75% of the costs in IIoT systems development lie in cloud and embedded software development. What is most important to system suppliers is a common platform that enables them to invest their research and development time and funds creating value through software services, rather than re-inventing communication interfaces, security infrastructures, control loop timing, data analytic algorithms, and so on. An FPGA-based approach offers many of these benefits to suppliers who must contend with a legacy processor system. An All-Programmable SoC approach helps maximize available options and is key to increased return on investment for both industrial system suppliers and their customers.

Conclusion

In summary, this white paper has highlighted how Xilinx All Programmable SoCs and FPGAs maximize the return on investment (ROI) for the systems supplier and its customers through:

- Long-term availability of Xilinx All Programmable SoCs and FPGAs combined with their inherent HW/SW reprogrammability to enable field updates and eliminate risk to follow emerging IIoT standards and trends
- Scalability across multiple Xilinx device families for system supplier platforms enabling lower Total Cost of Ownership across a comprehensive product line-up
- Integration of multiple IIoT functions from the IT and OT domains into a single flexible, low latency, and power-efficient device

Getting Started

Xilinx is the leading solutions provider for scalable and comprehensive IIoT edge platforms and continues to offer a growing number of recorded webinars, application notes, reference platforms, and evaluation kits, in conjunction with other leading companies in the IIoT industry. Examples include (but are not limited to) Industrial Ethernet connectivity, motor control, secure and measured boot for HW root of trust, machine learning, and so forth. The Avnet Industrial IoT Starter Kit and the Xilinx All Programmable Industrial Control System (APICS), in particular, bring together a variety of technologies from edge to cloud, to illustrate the benefits of integration and the capability of All Programmable SoCs.

For more information on All Programmable SoCs and FPGAs in IIoT, go to:
<https://www.xilinx.com/applications/megatrends/industrial-iot.html>

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
09/06/2017	1.0	Initial Xilinx release.

Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx’s limited warranty, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

AUTOMOTIVE PRODUCTS (IDENTIFIED AS “XA” IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE (“SAFETY APPLICATION”) UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD (“SAFETY DESIGN”). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.