Backgrounder

Xilinx Delivers Zynq UltraScale+ RFSoCs Enabling the RF Signal Chain for 5G Wireless, Cable Remote-PHY, and Radar

Silicon Shipped to Multiple Customers, Early Access Program Available Now

Introduction

In February of this year, Xilinx unveiled the monolithic integration of high performance RF data converters onto its SoC platform with its "RF-Analog" technology for commercial deployment of 5G radio and wireless backhaul. Xilinx is now disclosing details of its entire Zynq® UltraScale+™ RFSoC product line and shipping devices to tier-1 vendors developing multiple 5G end-applications, cable access remote-PHY nodes, and Electronic-Warfare / Radar applications. Xilinx RFSoC customers are now equipped with design and implementation tools, evaluation platforms, and silicon samples for end-application development. Beyond wireless, cable, and EW / Radar, RFSoCs are applicable to a breadth of RF-based applications, including test & measurement, SatCom, and military radio, among others.

As part of the RF signal chain, Zynq UltraScale+ RFSoC devices feature integrated Soft-Decision Forward Error Correction (SD-FEC) cores, addressing signal impairments in the RF signal chain and meeting 5G and DOCSIS 3.1 specifications. With these high-density, high throughput error correction cores, Xilinx RFSoCs serve 5G baseband as well as encapsulate modem functionality for 5G mobile backhaul and DOCSIS 3.1 Remote-PHY.

Monolithic RF-Analog Integration for 50-75% Power and Footprint Reduction

The All-Programmable RFSoC integrates up to 16x16 carrier-grade RF sampling ADCs and DACs tightly coupled to programmable logic and an ARM multi-processing subsystem.

By eliminating discrete ADC and DAC components, systems can achieve up to 50-75% reduction in system power and footprint, as shown in Figure 1.

Figure 1: Power and Footprint Impact from Integration of Discrete ADC and DAC Components

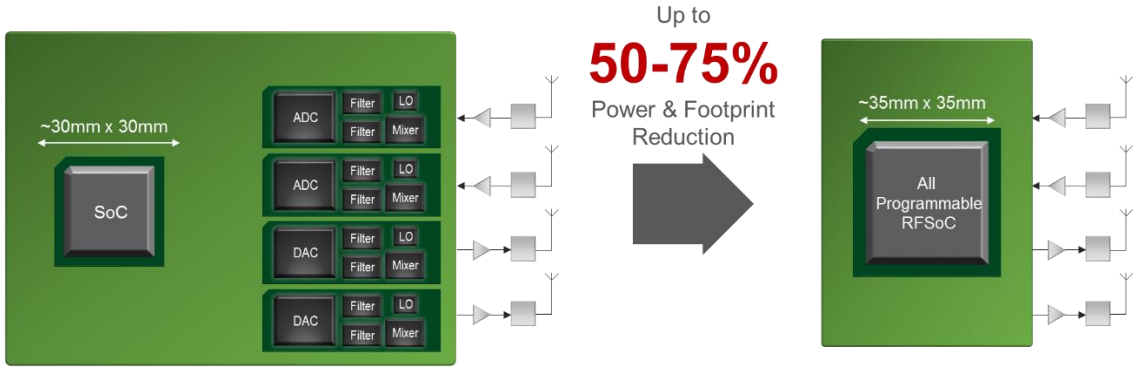

More than just RF data converters, the integrated block includes a power-efficient DSP subsystem for flexible configuration and RF signal conditioning. Specifically, the subsystem includes:

- Eight 4 GSPS or sixteen 2 GSPS 12-bit ADCs, with digital down-conversion (DUC)
- Eight to sixteen 6.4 GSPS 14-bit DACs, with digital up-conversion (DDC)
- Direct RF sampling for flexible analog design, greater accuracy, and lower power

Direct RF sampling, or the ability to sample incoming signals directly without initial down conversion to an intermediate frequency (IF), provides RF designers greater flexibility. Digitizing the signal directly and then applying modern DSP techniques for signal conditioning yields better performance and programmability in the digital domain, particularly on an advanced 16nm FinFET process. To date, direct RF techniques have been incremental in adoption due to the economics and power inefficiency. Monolithically integrating this technology into the SoC itself increases this RF technique's viability to the broader market.

Soft-Decision LDPC Forward Error Correction for 5G Access/Backhaul and DOCSIS 3.1

Critical to the RF signal chain and all communications is forward error correction—a DSP technique to correct signal impairments (e.g., degradation, distortion) in data transmission across various mediums, e.g., copper cable, optical fiber, or air interface. Xilinx has delivered FEC IP for over 20 years, and in 2015 introduced a hardened Reed-Solomon FEC (RS-FEC) in 16nm UltraScale+ FPGAs and MPSoCs primarily for wired communication applications.

With the ambitious throughput requirements of next generation wireless and cable broadband, both the 5G access/backhaul and DOCSIS 3.1 standards require a more compute-intensive FEC coding scheme—Low Density Parity Check (LDPC) codes— to maximize spectral efficiency in RF transmission.

While LDPC implementations can range from soft IP in FPGAs to fixed and hardened cores in ASSPs or ASICs—the Zynq UltraScale+ RFSoC balances flexibility, throughput, and power efficiency with the world's first hardened and fully programmable LDPC encoding/decoding cores in a programmable device.

Figure 2: Integrated and Fully Programmable LDPC Codec for 10-20X Throughput vs. Soft Core

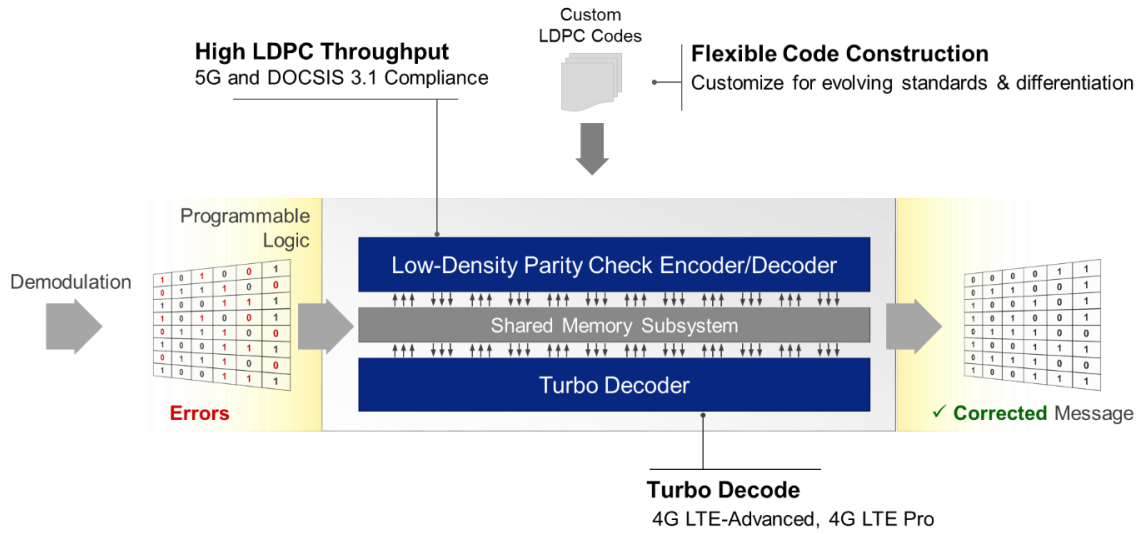

Capabilities include:

- Up to 42 Gb/s LDPC Encode and 10 Gb/s Decode System Throughput
- Turbo Decode for LTE backward compatibility with 4G LTE-Pro and LTE-Advanced
- 80% less dynamic power than a soft IP implementation
- Flexible, customizable LDPC codes for evolving standards and differentiation
- Soft-decision decoding for greater reliability

With compute-intensive matrix multiplication and continuous read and write to memory, hardening the SD-FEC can meet next generation standards for high throughput systems, such as 5G baseband.

A Comprehensive RF Signal Chain

From direct-RF conversion and signal conditioning, to acceleration and differentiation in FPGA logic, to correction of signal impairments, Zynq UltraScale+ integrates a comprehensive RF signal chain as shown in Figure *3*. Capabilities across the signal chain include:

- RF signal conditioning in the integrated data converters
- Linearization, correction, differentiation, and acceleration with soft IP and custom FPGA logic
- Baseband processing and acceleration, and error correction with integrated DSP and FEC cores
- Multiple 100G Ethernet MAC cores
- 33 Gb/s transceivers for evolving connectivity standards

Tightly coupled to an ARM® processing subsystem, the RFSoC provides a unified platform for analog, digital, and embedded design—simplifying calibration and synchronization along the signal chain.

Figure 3: Integrating the RF Signal Chain from Radio to Modem

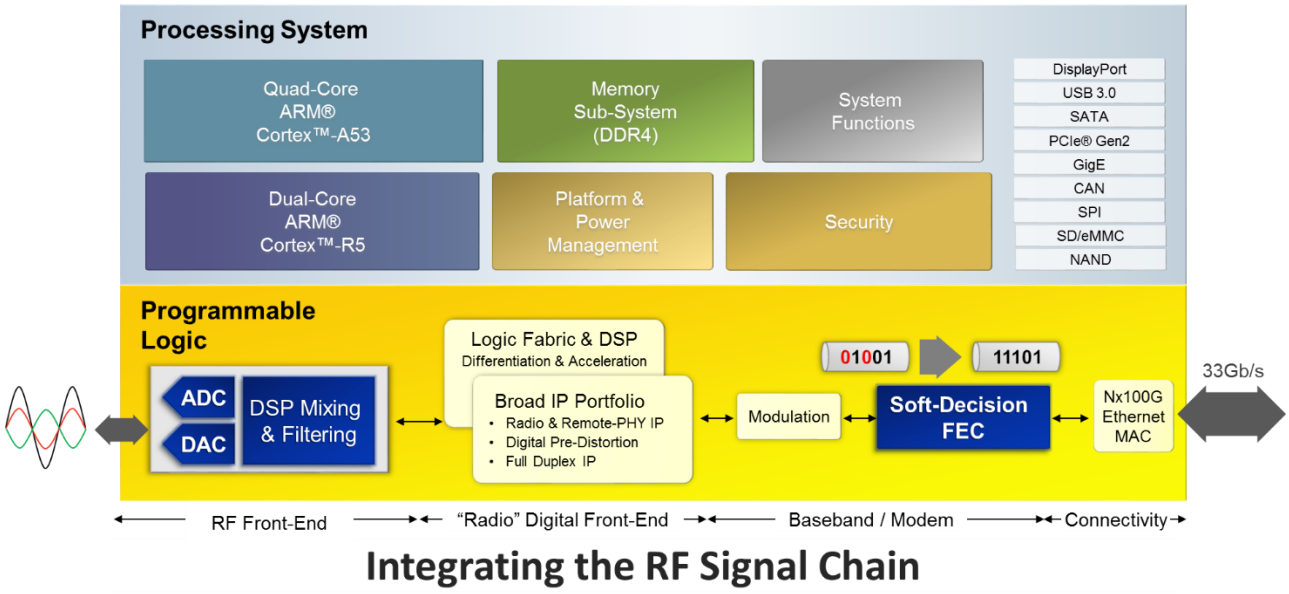

5G Wireless

There is no single solution to realize the broad range of use cases envisioned for fifth generation wireless systems. Massive MIMO, new beamforming techniques, millimeter wave transmission, and other enablers all contribute to next generation imperatives such as spectral efficiency, power efficiency, and network densification. Targeting remote radio, baseband, and wireless backhaul, Zynq UltraScale+ RFSoCs play a central role in meeting many of these imperatives.

Figure 4: All Programmable RFSoCs Enabling 5G Architectures

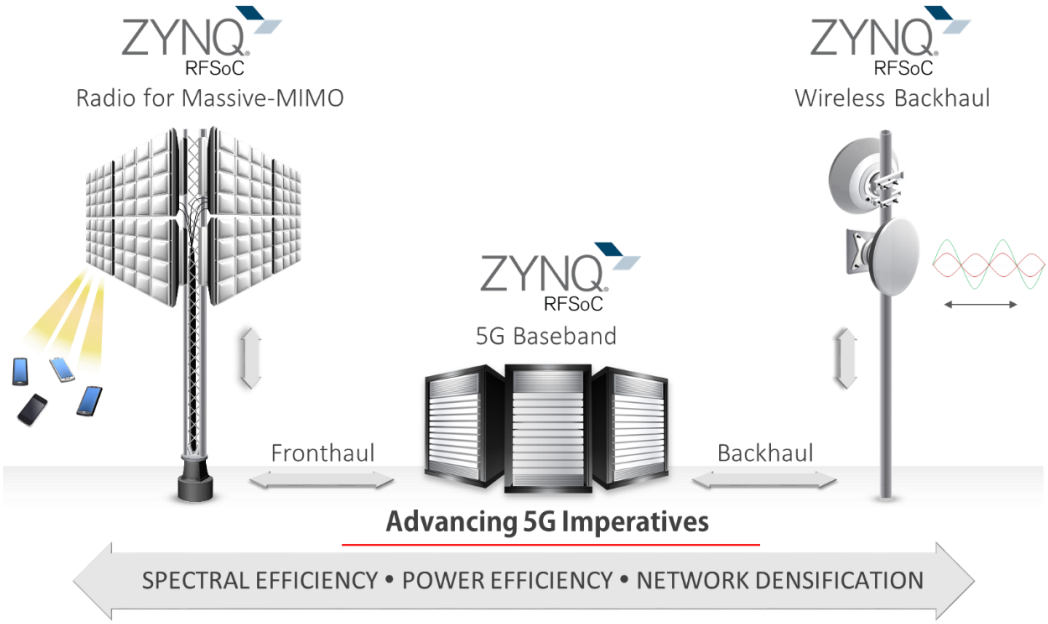

5G New Radio (NR) and Massive-MIMO

Remote radio head design and implementation of 5G New Radio is one of the most well suited applications for Xilinx RFSoCs. Since Massive MIMO transceivers can typically incorporate a large number of antennas (e.g., 32, 64, and over 1,024+) in the form of two-dimensional arrays—the number of discrete, sizable components eliminated using Zynq UltraScale+ is groundbreaking. A common SoC platform simplifies design and integration overhead associated with moving data between the RF front-ends (RFFE) and the digital front-end (DFE).

An RFSoC implementation, however, entails more than just ADC and DAC integration for 5G radio. Both the programmable logic and processing subsystem in Zynq UltraScale+ encompass the radio digital front-end. Specifically, the RFSoC provides:

- Push-button 491MHz fabric performance on UltraScale+ 16nm FinFET silicon
- Embedded DSP blocks for energy efficient signal processing and beamforming
- IP for digital pre-distortion (DPD), DUC, DDC, and crest factor reduction (CFR)
- ARM Cortex-A53 and/or Cortex-R5 for Operation & Maintenance (O&M)
- ARM Cortex-A53 cores for RF calibration, DPD software, and beamforming control
- Flexible L1-Split IP, reducing fronthaul throughput requirements
- 33 Gb/s transceivers CPRI interface for high performance fronthaul network interface

Figure 5: Implementing RF Front-End and Digital Front-End Radio with Zynq UltraScale+ RFSoCs

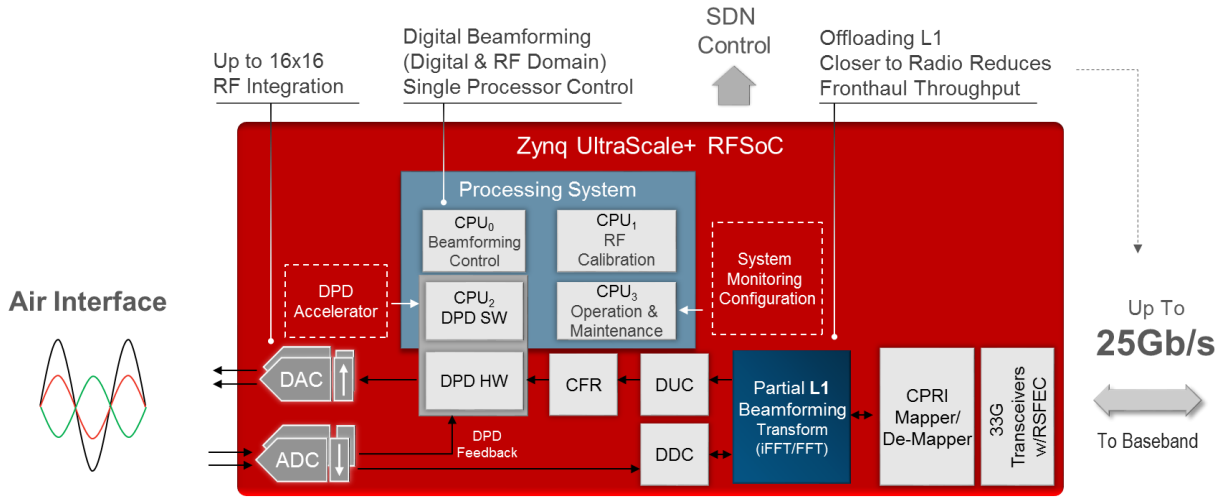

Baseband

Baseband units carry one of the most computationally intensive workloads in a Radio Access Network (RAN). For greater cost efficiency and the ability to centrally manage different radio sites, baseband units in 4G/LTE networks are commonly pooled in centralized offices, while radios are co-located with the antenna arrays. Within a baseband card, the workload is typically shared across FPGAs, DSPs, ASICs, and GPPUs. But with wider carrier bandwidths, more carrier aggregation, and the matrix computations needed for beamforming, 5G baseband demands hardware parallelism beyond what is available now in 4G/LTE.

For massive throughput in deployable baseband cards, manufacturers can delegate the most compute-intensive tasks—Layer-1 PHY acceleration and offloading—to Zynq UltraScale+ RFSoCs, a workload traditionally done by the DSP or ASIC. L1 acceleration in 5G is even more computationally intensive than in 4G-LTE in light of bandwidth requirements, the number of antenna paths, flexible numerology (waveform parameterization), and LDPC throughput for spectral efficiency.

As shown in Delivering more than just throughput, SD-FEC flexibility supports the LDPC coding scheme specified in the latest 3GPP release for 5G—with the ability to use custom codes for vendor differentiation. Turbo decode support provides 4G LTE-Advanced and 4G LTE-Advanced Pro for scenarios where 5G deployment is gradual. In short, for wireless, the SD-FEC combines the flexibility of a soft core with the performance of an ASIC. Manufacturers can evolve with the 3GPP standard and explore proprietary LDPC codes for differentiation with an RFSoC solution.

Figure *6*, the SoC becomes a major compute engine in baseband with

- Up to 42 Gb/s encode and 10Gb/s decode LDPC SD-FEC throughput
- Over 6,000 GMACs of signal processing bandwidth enabling L1 Offloading
- Power-optimized 33 Gb/s transceivers for 12.2G CPRI /eCPRI and expansion to 16G & 25G CPRI

Delivering more than just throughput, SD-FEC flexibility supports the LDPC coding scheme specified in the latest 3GPP release for 5G—with the ability to use custom codes for vendor differentiation. Turbo decode support provides 4G LTE-Advanced and 4G LTE-Advanced Pro for scenarios where 5G deployment is

gradual. In short, for wireless, the SD-FEC combines the flexibility of a soft core with the performance of an ASIC. Manufacturers can evolve with the 3GPP standard and explore proprietary LDPC codes for differentiation with an RFSoC solution.

Figure 6: Zynq UltraScale+ RFSoC in 5G Baseband for L1 Acceleration

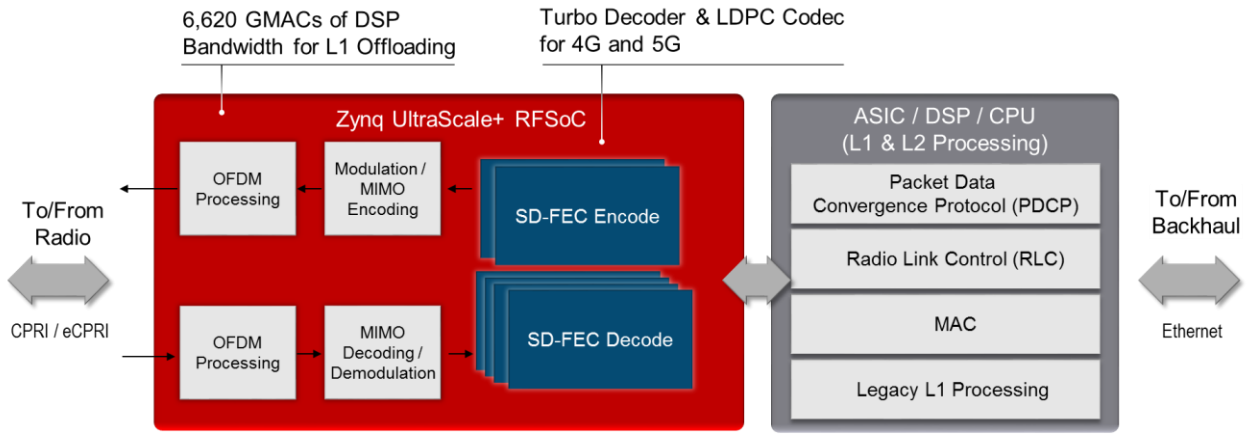

Power efficiency is critical in baseband processing, and a single Zynq UltraScale+ SD-FEC core in an LDPC configuration would take 300K logic cells in programmable logic while consuming 5X the dynamic power. RFSoCs realize 10-20X throughput vs. soft cores while meeting the stringent power and thermal constraints for centralized baseband.

Wireless Backhaul

With denser cellular deployments, backhaul links connecting distributed/central units to the core networks demand greater throughput and capacity per cell site. Wireless connectivity continues to be the most flexible method of backhaul, and 5G takes advantage of the millimeter wave spectrum (ranging from 30GHz to 300GHz) for "fiber-like" connection speeds.

As full outdoor units mounted on cell towers or other urban infrastructure, wireless backhaul systems implemented with Zynq UltraScale+ RFSoCs can encapsulate the RF front-end, digital radio front-end, modem, and perform packet processing functions. Unlike radio or baseband, backhaul takes advantage of both RF-Analog and SD-FEC technologies of the RFSoC.

With leading backhaul architectures employing 4x4 channel systems, the power and footprint reduction can range from 30-50% through integration of discrete ADC and DAC components. Though the number of channels is fewer than in remote radio, the throughput for backhaul is significant, with uplink rates of 30 gigabits-per-second. As shown in Figure 7, a Zynq UltraScale+ RFSoC implementation of wireless backhaul includes

- ARM processing system for O&M functionality
- Up to 4x4 TX/RX channels for RF integration (device capable of 8x8)
- LDPC SD-FEC with optional by-pass mode (for transmissions with lower correction rates)

Figure 7: RF-Analog and SD-FEC Integration in Wireless Backhaul

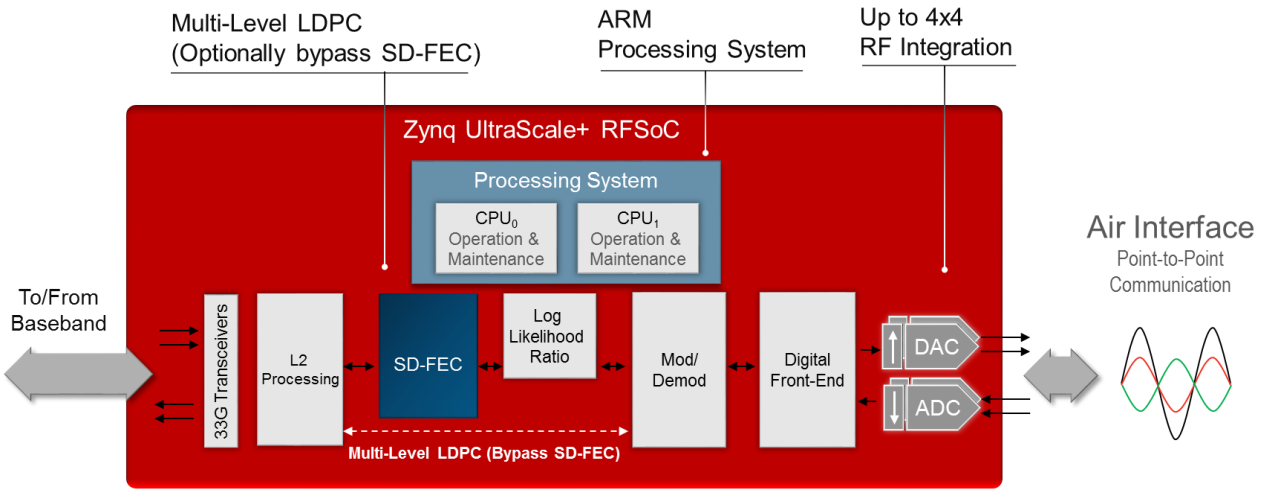

Cable Remote-PHY

Experiencing similar pressures as mobile network providers to evolve infrastructure for greater network capacity, cable multi-service operators (MSOs) have followed (and continue to follow) an evolutionary path similar to 4G-LTE and 5G wireless in terms of decentralized architectures and network virtualization. To realize 10X data rates (10Gb/s downstream) over previous generation cable (DOCSIS 3.0), the DOCSIS 3.1 standard seeks to achieve more efficiency in the RF spectrum. In turn, MSOs naturally seek to achieve this with the lowest overhead.

Deep Fiber and Remote-PHY for Distributed Access Architectures

As part of the MSO strategy, cable providers are gradually migrating to distributed access architectures (DAA) where various functions previously done in the headend (central office) are moved closer to the subscriber premises to distribute the workload and fanout while minimizing the cost of power, maintenance, and upgradability. Analogous to distributed base stations in cellular networks, DAA involves moving analog conversion and front-end processing closer to the subscriber premises. 'Fiber deep' and Remote-PHY node deployment are part of this re-architecture.

In 'fiber deep', digital fiber replaces analog fiber from the headend office and strung even closer to the subscriber premises—i.e., closer to the home. Replacement of analog with digital fiber to the node (box or cabinet) eliminates the need for RF amplification through coax.

At the node itself is the Remote-PHY system, capable of fanning out to more subscribers through shorter, passive coaxial cable. The combination of deep fiber and remote RF conversion/processing dramatically improves power- and spectral-efficiency for multi-gigabit subscriber bandwidth.

RFSoCs for Remote-PHY: RF Integration and LDPC FEC

Zynq UltraScale+ RFSoCs deliver the form factor, power efficiency, and DOCSIS 3.1 compliance to make Remote-PHY deployment not only viable but upgradable and future proof. Remote-PHY integrates the RF- and digital front-end, as well as an LDPC FEC-enabled modem.

Figure 8: Deployment of Remote-PHY Nodes with All Programmable RFSoCs

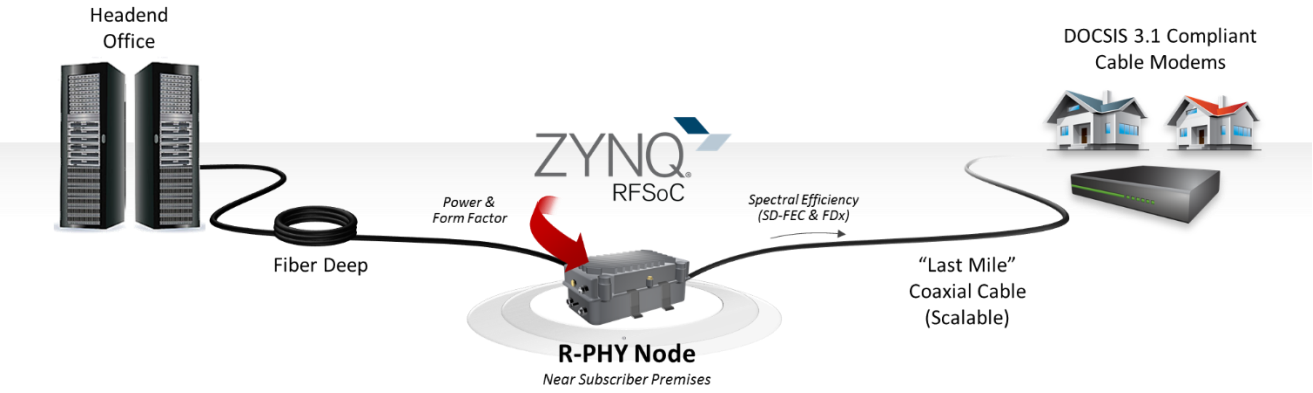

As a multi-channel system of varied configurations, R-PHY demands power efficiency given the greater number of nodes vs. headend equipment. A basic 1 downstream x 2 upstream channel system with Zynq UltraScale+ RFSoC can reduce power and footprint of the Remote-PHY Device by 30-50% vs. a traditional FPGA/ASSP implementation through integration of data converters, all while offering scalability of channel count for flexible provisioning of varied services and number of subscribers.

Power efficiency is even greater through the hardening of the SD-FEC—whose LDPC support is a key requirement of DOCSIS 3.1 and critical to the spectral efficiency requirements across coax. A soft LDPC core would demand a multi-device implementation and waste programmable logic that can be used for differentiation.

Flexibility with Programmable Logic and a Complete R-PHY IP Portfolio

In fact the programmable logic enables flexibility for the rest of the system, allowing for upgradability and evolution with standards. A comprehensive IP portfolio includes all the necessary components for a complete R-PHY device, allowing for rapid bring up and a foundation for further differentiation. With FPGA fabric MSOs can field-upgrade hardware without overhauling or uprooting their infrastructure.

Part of the portfolio is Full Duplex DOCSIS (FDx) Modulation IP—providing the ability for up-stream and downstream services across the same frequency. The industry is looking to FDx for even greater spectral efficiency across coax, but whose use model is dependent on Remote-PHY node deployment. FDx algorithms, however, are not yet standardized. Using programmable logic in an RFSoC and soft Xilinx IP as a foundation, manufacturers and MSOs can evolve future-proof their FDx solution as the specification matures.

As shown in Figure 9, the Zynq UltraScale+ RFSoC solution for R-PHY includes

- Complete R-PHY IP Portfolio, e.g., R-DEPI (downstream), R-UEPI (upstream), OFDM/A, FDx
- ARM Cortex-A53 for traffic management and DPD control
- Up to 8x8 TX/RX data converter integration for scalability
- LDPC Codec for DOCSIS 3.1 compliance
- 100G Ethernet MAC cores for packet processing, with MACSEC IP

Figure 9: Zynq UltraScale+ RFSoC in Remote-PHY Node with Comprehensive IP Portfolio

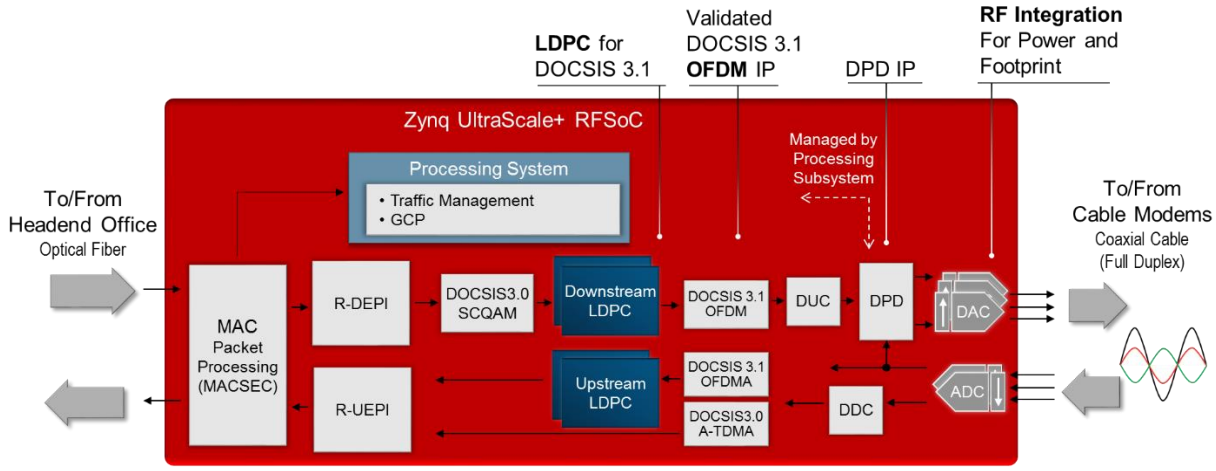

Radar and Electronic Warfare

The paradigm shift from ad-hoc discrete systems to integrated multi-channel phased array platforms presents a multi-faceted set of challenges for today's system engineering teams developing weather and defense radar platforms. Shrinking the form factor, lowering power requirements, improving thermal management, and reducing system latencies are all simultaneously on top of most program's lists. The Xilinx Zynq UltraScale+ RFSoCs provide weather and defense radar programs with the technology to address these challenges with the added flexibility of an all programmable cost-effective device.

For Electronic Warfare (EW)—and particularly multi-phased array radar—Zynq UltraScale+ RFSoCs deliver the low latency digital signal processing capability needed for rapid turnaround from detection to countermeasure together with a clear power, footprint and cost advantages to prior-art solutions. To implement the most versatile, flexible, and scalable EW processing systems, government programs can now benefit from DARPA's investments into commercial technology, innovation, and the improved performance-per-watt capabilities of 16nm FinFET+ technology. Previously cost prohibitive solutions are now possible and within budget.

T/R Module: Fundamental Building Block of Phased Array Radar

For modularity, phased array radar systems are built to handle *thousands* of Transmit/Receive (T/R) modules, which are the tightly coupled building blocks managing the front-end EW processing. Modular, stackable, and sized to fit within the lattice of a phased array, these modules integrate both digital and analog functionality.

Because the cost of radar programs are often hampered by an inability to adapt, the flexibility of the universal T/R module is of paramount importance. Most modules today are based on DSP-rich FPGAs or SoCs with integrated FPGA fabric. Reprogrammable hardware keeps the RADAR system "relevant" and "future-proof" with the ability to update waveforms and algorithms. Unfortunately, discrete RF data converters—considered the 'last ASICs' left on the board—are a 'bottleneck' to flexibility.

Figure 10: 2x2 Transmit / Receive Module & Building Block for Phased Array Radar

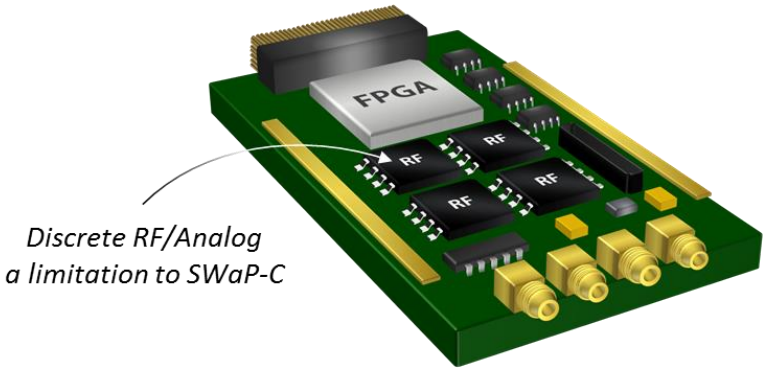

A Joint Program for RFSoC Adoption

Multiple government programs need a monolithic integration of RF and FPGA capabilities, such as the RFSoCs versus alternate Multi-Chip Modules (MCMs) or System-in-Package (SiP) solutions, which only address size challenges. In short, multiple governing bodies see All Programmable RFSoCs as an urgently needed breakthrough.

Specifically, the U.S. government's Multi-function Phased Array Radar (MPAR) initiative has combined the functions of several national radar networks into a single radar system for aircraft and weather surveillance. The joint-program sees adoption of Zynq UltraScale+ RFSoCs as a way to push past traditional barriers for delivering radar arrays.

The Spectrum Efficient National Surveillance Radar's (SENSR) is now collaborating with MPAR to move this effort forward, combining interests of the National Oceanic and Atmospheric Administration (NOAA) as well as the Federal Aviation Administration (FAA) with the rest of the defense community.

Low Power, Reduced Footprint, and Low Latency in EW Processing

For large-scale phased arrays, the SWaP-C advantages of RFSoCs are staggering when the number of modules are taken into consideration. The ideal choice for high channel density radar applications is the Zynq UltraScale+ RFSoC (ZU29DR) providing 16 ADCs at 2 GSPS with the power and footprint reductions needed.

In EW applications, systems designers need high channel density, and reduced power and footprint *with lower latency*. They must meet the low latency imperative demands of signal acquisition (ADC), followed by digital signal processing (DSP), and then the appropriate response (DAC). Whether it be to attack, impede, or control spectrum through methods such as jamming and spoofing, the data path delay through data converters and signal processing engine needs to be on the order of nanoseconds. Modern ADC and DAC devices with JESD204B chip-to-chip serial interfaces provide increased bandwidth, but also introduce unacceptable increases in system latencies for typical EW systems.

Figure 11: Reduced Power and Footprint, and Low Latency with RFSoCs

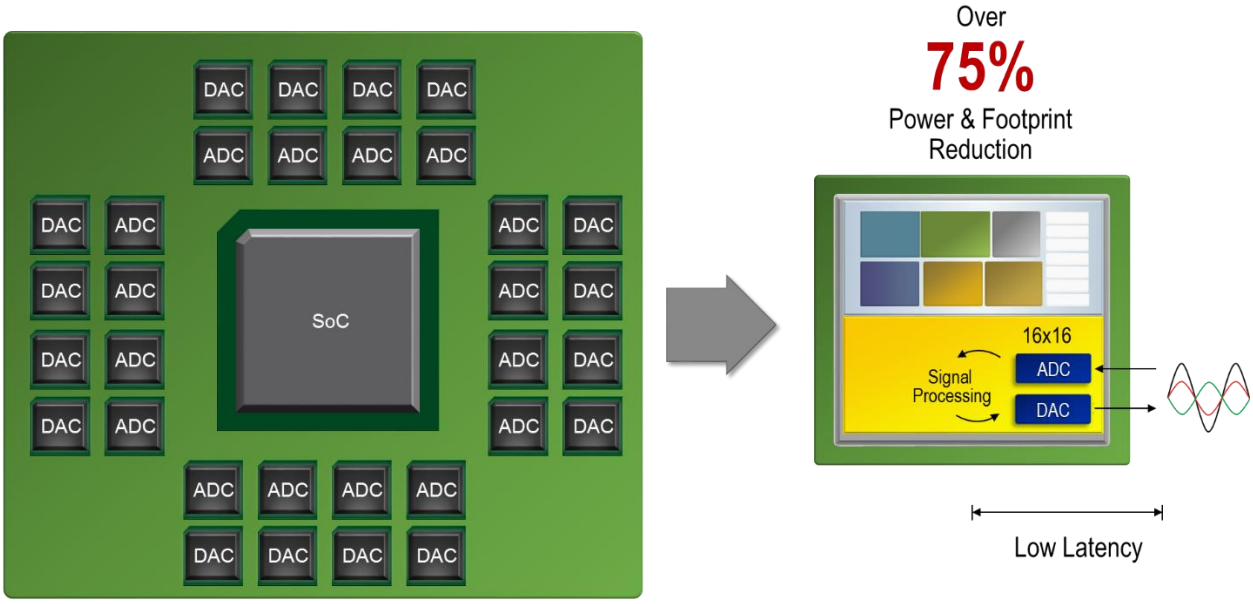

ADC and DAC integration, coupled with integrated digital signal processing engines on 16nm FinFET+ address the low latency requirements for rapid response in A2AD (Anti-Access and Area Denial) events. High channel count systems can achieve over 75% reduction in power and footprint. The single chip, modular 16x16 T/R module can then be arrayed into a 2D array for a fully programmable digital radar platform.

Figure 12: Single Chip T/R Modules for Programmable Digital Radar

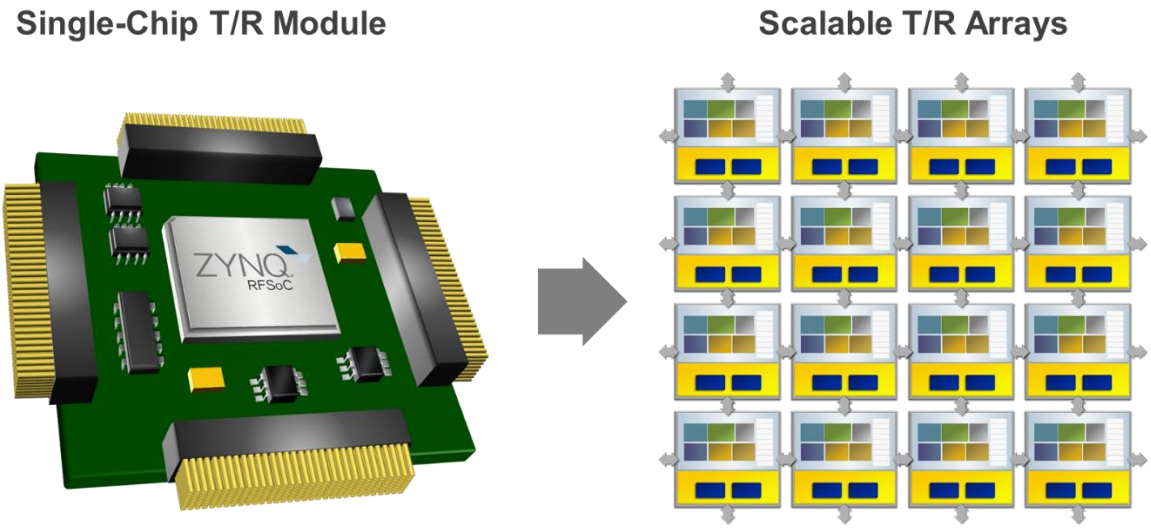

Approaching RF Convergence

Bringing RF, digital, and software design into a monolithic device brings the defense industry one step closer to meeting the vision of RF hardware/software convergence—where a single platform can combine functions for different end-systems. Given that nearly all defense electronic systems have RF elements, a

modular piece of hardware that can be reprogrammed and reused not only reduces SWaP-C across end-applications, but cuts development costs while enabling a longer system lifecycle.

An RFSoC-based platform that enables EW Radar can be reprogrammed for use in an airborne radio or a SIGINT receiver, ensuring compatibility, interoperability, and re-use. The key to fulfilling this vision of modularity and scalability is not just the hardware, but a common framework across applications and design domains. Open architecture and open standards initiatives together with high level abstraction tools for RF bring hardware and software designers closer together in the system design process.

Tools, IP, and Evaluation Platforms

Delivered with the All Programmable RFSoCs comes a complete analog-to-digital design solution. The RFFE and radio DFE can be developed and co-verified more seamlessly versus two discrete subsystems. Xilinx provides all the elements of a complete solution for system-wide design exploration, bring-up, validation, and final implementation. A comprehensive IP portfolio, tool suite, and base evaluation platform make this possible.

The IP portfolio brings both RF and digital designers up and running in short order. Validated RF sampling IP comes complete with software drivers. Xilinx's extensive IP portfolio includes wireless radio digital front end IP, such as DPD, DUC, DDC, and CFR cores—enabling off-the-shelf usage and the ability to customize for differentiation. Similarly, a complete IP portfolio is now available for Remote-PHY systems, including R-DEPI, R-UEPI, and OFDM cores.

The integrated tool flow covers all design domains, encompassing RF design and simulation, digital logic design, and embedded SW development. Vivado® IP Integrator simplifies reuse of IP and subsystems across multiple designs. DSP and system architects designing for the RF or digital domain can leverage Vivado High Level Synthesis for design in C, C++, or OpenCL, or they can take advantage of System Generator for designs in MATLAB® or Simulink®.

Zynq UltraScale+ evaluation kits—complete with reference designs—are also available for technology evaluation and application development.

A Scalable SoC Platform for Analog & Digital Design

With All Programmable RFSoCs, diverse end-markets can leverage the same analog and digital hardware not just for power and footprint reduction, but for scalability and design variants. RF, digital, and embedded designers may always have to reconcile conflicting requirements, but a single platform simplifies integration of subsystems, synchronization and calibration across domains, and the ability to adapt to evolving standards and requirements. No single radio variant will address the evolving requirements of 5G, nor any one Remote-PHY configuration for cable access, nor any single EW processing platform for phased array radar. For these and other end-markets, integration is the most seamless approach to realize—and adapt—the system.