



Welcome to [E-XFL.COM](https://www.e-xfl.com)

Embedded - System On Chip (SoC): The Heart of Modern Embedded Systems

Embedded - System On Chip (SoC) refers to an integrated circuit that consolidates all the essential components of a computer system into a single chip. This includes a microprocessor, memory, and other peripherals, all packed into one compact and efficient package. SoCs are designed to provide a complete computing solution, optimizing both space and power consumption, making them ideal for a wide range of embedded applications.

What are Embedded - System On Chip (SoC)?

System On Chip (SoC) integrates multiple functions of a computer or electronic system onto a single chip. Unlike traditional multi-chip solutions, SoCs combine a central

Details

Product Status	Active
Architecture	MCU, FPGA
Core Processor	Quad ARM® Cortex®-A53 MPCore™ with CoreSight™, Dual ARM®Cortex™-R5 with CoreSight™, ARM Mali™ -400 MP2
Flash Size	-
RAM Size	256KB
Peripherals	DMA, WDT
Connectivity	CANbus, EBI/EMI, Ethernet, I ² C, MMC/SD/SDIO, SPI, UART/USART, USB OTG
Speed	500MHz, 600MHz, 1.2GHz
Primary Attributes	Zynq@UltraScale+™ FPGA, 256K+ Logic Cells
Operating Temperature	-40°C ~ 100°C (Tj)
Package / Case	900-BBGA, FCBGA
Supplier Device Package	900-FCBGA (31x31)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xczu5ev-1fbvb900i

ARM Mali-400 Based GPU

- Supports OpenGL ES 1.1 and 2.0
- Supports OpenVG 1.1
- GPU frequency: Up to 667MHz
- Single Geometry Processor, Two Pixel Processors
- Pixel Fill Rate: 2 Mpixels/sec/MHz
- Triangle Rate: 0.11 Mtriangles/sec/MHz
- 64KB L2 Cache
- Power island gating

External Memory Interfaces

- Multi-protocol dynamic memory controller
- 32-bit or 64-bit interfaces to DDR4, DDR3, DDR3L, or LPDDR3 memories, and 32-bit interface to LPDDR4 memory
- ECC support in 64-bit and 32-bit modes
- Up to 32GB of address space using single or dual rank of 8-, 16-, or 32-bit-wide memories
- Static memory interfaces
 - eMMC4.51 Managed NAND flash support
 - ONFI3.1 NAND flash with 24-bit ECC
 - 1-bit SPI, 2-bit SPI, 4-bit SPI (Quad-SPI), or two Quad-SPI (8-bit) serial NOR flash

8-Channel DMA Controller

- Two DMA controllers of 8-channels each
- Memory-to-memory, memory-to-peripheral, peripheral-to-memory, and scatter-gather transaction support

Serial Transceivers

- Four dedicated PS-GTR receivers and transmitters supports up to 6.0Gb/s data rates
 - Supports SGMII tri-speed Ethernet, PCI Express® Gen2, Serial-ATA (SATA), USB3.0, and DisplayPort

Dedicated I/O Peripherals and Interfaces

- PCI Express — Compliant with PCIe® 2.1 base specification
 - Root complex and End Point configurations
 - x1, x2, and x4 at Gen1 or Gen2 rates
- SATA Host
 - 1.5, 3.0, and 6.0Gb/s data rates as defined by SATA Specification, revision 3.1
 - Supports up to two channels
- DisplayPort Controller
 - Up to 5.4Gb/s rate
 - Up to two TX lanes (no RX support)

- Four 10/100/1000 tri-speed Ethernet MAC peripherals with IEEE Std 802.3 and IEEE Std 1588 revision 2.0 support
 - Scatter-gather DMA capability
 - Recognition of IEEE Std 1588 rev.2 PTP frames
 - GMII, RGMII, and SGMII interfaces
 - Jumbo frames
- Two USB 3.0/2.0 Device, Host, or OTG peripherals, each supporting up to 12 endpoints
 - USB 3.0/2.0 compliant device IP core
 - Super-speed, high-speed, full-speed, and low-speed modes
 - Intel XHCI-compliant USB host
- Two full CAN 2.0B-compliant CAN bus interfaces
 - CAN 2.0-A and CAN 2.0-B and ISO 118981-1 standard compliant
- Two SD/SDIO 2.0/eMMC4.51 compliant controllers
- Two full-duplex SPI ports with three peripheral chip selects
- Two high-speed UARTs (up to 1Mb/s)
- Two master and slave I2C interfaces
- Up to 78 flexible multiplexed I/O (MIO) (up to three banks of 26 I/Os) for peripheral pin assignment
- Up to 96 EMIOs (up to three banks of 32 I/Os) connected to the PL

Interconnect

- High-bandwidth connectivity within PS and between PS and PL
- ARM AMBA® AXI4-based
- QoS support for latency and bandwidth control
- Cache Coherent Interconnect (CCI)

System Memory Management

- System Memory Management Unit (SMMU)
- Xilinx Memory Protection Unit (XMPU)

Platform Management Unit

- Power gates PS peripherals, power islands, and power domains
- Clock gates PS peripheral user firmware option

Configuration and Security Unit

- Boots PS and configures PL
- Supports secure and non-secure boot modes

System Monitor in PS

- On-chip voltage and temperature sensing

Feature Summary

Table 1: Zynq UltraScale+ MPSoC: CG Device Feature Summary

	ZU2CG	ZU3CG	ZU4CG	ZU5CG	ZU6CG	ZU7CG	ZU9CG
Application Processing Unit	Dual-core ARM Cortex-A53 MPCore with CoreSight; NEON & Single/Double Precision Floating Point; 32KB/32KB L1 Cache, 1MB L2 Cache						
Real-Time Processing Unit	Dual-core ARM Cortex-R5 with CoreSight; Single/Double Precision Floating Point; 32KB/32KB L1 Cache, and TCM						
Embedded and External Memory	256KB On-Chip Memory w/ECC; External DDR4; DDR3; DDR3L; LPDDR4; LPDDR3; External Quad-SPI; NAND; eMMC						
General Connectivity	214 PS I/O; UART; CAN; USB 2.0; I2C; SPI; 32b GPIO; Real Time Clock; WatchDog Timers; Triple Timer Counters						
High-Speed Connectivity	4 PS-GTR; PCIe Gen1/2; Serial ATA 3.1; DisplayPort 1.2a; USB 3.0; SGMII						
System Logic Cells	103,320	154,350	192,150	256,200	469,446	504,000	599,550
CLB Flip-Flops	94,464	141,120	175,680	234,240	429,208	460,800	548,160
CLB LUTs	47,232	70,560	87,840	117,120	214,604	230,400	274,080
Distributed RAM (Mb)	1.2	1.8	2.6	3.5	6.9	6.2	8.8
Block RAM Blocks	150	216	128	144	714	312	912
Block RAM (Mb)	5.3	7.6	4.5	5.1	25.1	11.0	32.1
UltraRAM Blocks	0	0	48	64	0	96	0
UltraRAM (Mb)	0	0	14.0	18.0	0	27.0	0
DSP Slices	240	360	728	1,248	1,973	1,728	2,520
CMTs	3	3	4	4	4	8	4
Max. HP I/O ⁽¹⁾	156	156	156	156	208	416	208
Max. HD I/O ⁽²⁾	96	96	96	96	120	48	120
System Monitor	2	2	2	2	2	2	2
GTH Transceiver 16.3Gb/s ⁽³⁾	0	0	16	16	24	24	24
GTY Transceivers 32.75Gb/s	0	0	0	0	0	0	0
Transceiver Fractional PLLs	0	0	8	8	12	12	12
PCIe Gen3 x16 and Gen4 x8	0	0	2	2	0	2	0
150G Interlaken	0	0	0	0	0	0	0
100G Ethernet w/ RS-FEC	0	0	0	0	0	0	0

Notes:

1. HP = High-performance I/O with support for I/O voltage from 1.0V to 1.8V.
2. HD = High-density I/O with support for I/O voltage from 1.2V to 3.3V.
3. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s. See [Table 2](#).

Table 5: Zynq UltraScale+ MPSoC: EV Device Feature Summary

	ZU4EV	ZU5EV	ZU7EV
Application Processing Unit	Quad-core ARM Cortex-A53 MPCore with CoreSight; NEON & Single/Double Precision Floating Point; 32KB/32KB L1 Cache, 1MB L2 Cache		
Real-Time Processing Unit	Dual-core ARM Cortex-R5 with CoreSight; Single/Double Precision Floating Point; 32KB/32KB L1 Cache, and TCM		
Embedded and External Memory	256KB On-Chip Memory w/ECC; External DDR4; DDR3; DDR3L; LPDDR4; LPDDR3; External Quad-SPI; NAND; eMMC		
General Connectivity	214 PS I/O; UART; CAN; USB 2.0; I2C; SPI; 32b GPIO; Real Time Clock; WatchDog Timers; Triple Timer Counters		
High-Speed Connectivity	4 PS-GTR; PCIe Gen1/2; Serial ATA 3.1; DisplayPort 1.2a; USB 3.0; SGMII		
Graphic Processing Unit	ARM Mali™-400 MP2; 64KB L2 Cache		
Video Codec	1	1	1
System Logic Cells	192,150	256,200	504,000
CLB Flip-Flops	175,680	234,240	460,800
CLB LUTs	87,840	117,120	230,400
Distributed RAM (Mb)	2.6	3.5	6.2
Block RAM Blocks	128	144	312
Block RAM (Mb)	4.5	5.1	11.0
UltraRAM Blocks	48	64	96
UltraRAM (Mb)	14.0	18.0	27.0
DSP Slices	728	1,248	1,728
CMTs	4	4	8
Max. HP I/O ⁽¹⁾	156	156	416
Max. HD I/O ⁽²⁾	96	96	48
System Monitor	2	2	2
GTH Transceiver 16.3Gb/s ⁽³⁾	16	16	24
GTY Transceivers 32.75Gb/s	0	0	0
Transceiver Fractional PLLs	8	8	12
PCIe Gen3 x16 and Gen4 x8	2	2	2
150G Interlaken	0	0	0
100G Ethernet w/ RS-FEC	0	0	0

Notes:

1. HP = High-performance I/O with support for I/O voltage from 1.0V to 1.8V.
2. HD = High-density I/O with support for I/O voltage from 1.2V to 3.3V.
3. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s. See [Table 6](#).

ASIC-class capabilities afforded by the UltraScale MPSoC architecture while supporting rapid system development.

The inclusion of an application processor enables high-level operating system support, e.g., Linux. Other standard operating systems used with the Cortex-A53 processor are also available for the Zynq UltraScale+ MPSoC family. The PS and the PL are on separate power domains, enabling users to power down the PL for power management if required. The processors in the PS always boot first, allowing a software centric approach for PL configuration. PL configuration is managed by software running on the CPU, so it boots similar to an ASSP.

Processing System

Application Processing Unit (APU)

The key features of the APU include:

- 64-bit quad-core ARM Cortex-A53 MPCores. Features associated with each core include:
 - ARM v8-A Architecture
 - Operating target frequency: up to 1.5GHz
 - Single and double precision floating point: 4 SP / 2 DP FLOPs
 - NEON Advanced SIMD support with single and double precision floating point instructions
 - A64 instruction set in 64-bit operating mode, A32/T32 instruction set in 32-bit operating mode
 - Level 1 cache (separate instruction and data, 32KB each for each Cortex-A53 CPU)
 - 2-way set-associative Instruction Cache with parity support
 - 4-way set-associative Data Cache with ECC support
 - Integrated memory management unit (MMU) per processor core
 - TrustZone for secure mode operation
 - Virtualization support
- Ability to operate in single processor, symmetric quad processor, and asymmetric quad-processor modes
- Integrated 16-way set-associative 1MB Unified Level 2 cache with ECC support
- Interrupts and Timers
 - Generic interrupt controller (GIC-400)
 - ARM generic timers (4 timers per CPU)
 - One watchdog timer (WDT)
 - One global timer
 - Two triple timers/counters (TTC)
- Little and big endian support
 - Big endian support in BE8 mode
- CoreSight debug and trace support
 - Embedded Trace Macrocell (ETM) for instruction trace
 - Cross trigger interface (CTI) enabling hardware breakpoints and triggers
- ACP interface to PL for I/O coherency and Level 2 cache allocation
- ACE interface to PL for full coherency
- Power island gating on each processor core
- Optional eFUSE disable per core

Real-Time Processing Unit (RPU)

- Dual-core ARM Cortex-R5 MPCores. Features associated with each core include:
 - ARM v7-R Architecture (32-bit)
 - Operating target frequency: Up to 600MHz
 - A32/T32 instruction set support
 - 4-way set-associative Level 1 caches (separate instruction and data, 32KB each) with ECC support
 - Integrated Memory Protection Unit (MPU) per processor
 - 128KB Tightly Coupled Memory (TCM) with ECC support
 - TCMs can be combined to become 256KB in lockstep mode
- Ability to operate in single-processor or dual-processor modes (split and lock-step)
- Little and big endian support
- Dedicated SWDT and two Triple Timer Counters (TTC)
- CoreSight debug and trace support
 - Embedded Trace Macrocell (ETM) for instruction and trace
 - Cross trigger interface (CTI) enabling hardware breakpoints and triggers
- Optional eFUSE disable

Full-Power Domain DMA (FPD-DMA) and Low-Power Domain DMA (LPD-DMA)

- Two general-purpose DMA controllers one in the full-power domain (FPD-DMA) and one in the low-power domain (LPD-DMA)
- Eight independent channels per DMA
- Multiple transfer types:
 - Memory-to-memory
 - Memory-to-peripheral
 - Peripheral-to-memory and
 - Scatter-gather
- 8 peripheral interfaces per DMA
- TrustZone per DMA for optional secure operation

Xilinx Memory Protection Unit (XMPU)

- Region based memory protection unit
- Up to 16 regions
- Each region supports address alignment of 1MB or 4KB
- Regions can overlap; the higher region number has priority
- Each region can be independently enabled or disabled
- Each region has a start and end address

Graphics Processing Unit (GPU)

- Supports OpenGL ES 1.1 & 2.0
- Supports OpenVG 1.1
- Operating target frequency: up to 667MHz
- Single Geometry Processor and two Pixel processor
- Pixel Fill Rate: 2 Mpixel/sec/MHz
- Triangle Rate: 0.11 Mtriangles/sec/MHz
- 64KB Level 2 Cache (read-only)
- 4X and 16X Anti-aliasing Support
- ETC1 texture compression to reduce external memory bandwidth
- Extensive texture format support
 - RGBA 8888, 565, 1556
 - Mono 8, 16
 - YUV format support
- Automatic load balancing across different graphics shader engines
- 2D and 3D graphic acceleration
- Up to 4K texture input and 4K render output resolutions
- Each geometry processor and pixel processor supports 4KB page MMU
- Power island gating on each GPU engine and shared cache
- Optional eFUSE disable

Dynamic Memory Controller (DDRC)

- DDR3, DDR3L, DDR4, LPDDR3, LPDDR4
- Target data rate: Up to 2400Mb/s DDR4 operation in -1 speed grade
- 32-bit and 64-bit bus width support for DDR4, DDR3, DDR3L, or LPDDR3 memories, and 32-bit bus width support for LPDDR4 memory
- ECC support (using extra bits)
- Up to a total DRAM capacity of 32GB

SATA

- Compliant with SATA 3.1 Specification
- SATA host port supports up to 2 external devices
- Compliant with Advanced Host Controller Interface ('AHCI') ver. 1.3
- 1.5Gb/s, 3.0Gb/s, and 6.0Gb/s data rates
- Power management features: supports partial and slumber modes

USB 3.0

- Two USB controllers (configurable as USB 2.0 or USB 3.0)
- Up to 5.0Gb/s data rate
- Host and Device modes
 - Super Speed, High Speed, Full Speed, and Low Speed
 - Up to 12 endpoints
 - The USB host controller registers and data structures are compliant to Intel xHCI specifications
 - 64-bit AXI master port with built-in DMA
 - Power management features: Hibernation mode

DisplayPort Controller

- 4K Display Processing with DisplayPort output
 - Maximum resolution of 4K x 2K-30 (30Hz pixel rate)
 - DisplayPort AUX channel, and Hot Plug Detect (HPD) on the output
 - RGB YCbCr, 4:2:0; 4:2:2, 4:4:4 with 6, 8, 10, and 12b/c
 - Y-only, xvYCC, RGB 4:4:4, YCbCr 4:4:4, YCbCr 4:2:2, and YCbCr 4:2:0 video format with 6,8,10 and 12-bits per color component
 - 256-color palette
 - Multiple frame buffer formats
 - 1, 2, 4, 8 bits per pixel (bpp) via a palette
 - 16, 24, 32bpp
 - Graphics formats such as RGBA8888, RGB555, etc.
- Accepts streaming video from the PL or dedicated DMA controller
- Enables Alpha blending of graphics and Chroma keying

- Audio support
 - A single stream carries up to 8 LPCM channels at 192kHz with 24-bit resolution
 - Supports compressed formats including DRA, Dolby MAT, and DTS HD
 - Multi-Stream Transport can extend the number of audio channels
 - Audio copy protection
 - 2-channel streaming or input from the PL
 - Multi-channel non-streaming audio from a memory audio frame buffer
- Includes a System Time Clock (STC) compliant with ISO/IEC 13818-1
- Boot-time display using minimum resources

Platform Management Unit (PMU)

- Performs system initialization during boot
- Acts as a delegate to the application and real-time processors during sleep state
- Initiates power-up and restart after the wake-up request
- Maintains the system power state at all time
- Manages the sequence of low-level events required for power-up, power-down, reset, clock gating, and power gating of islands and domains
- Provides error management (error handling and reporting)
- Provides safety check functions (e.g., memory scrubbing)

The PMU includes the following blocks:

- Platform management processor
- Fixed ROM for boot-up of the device
- 128KB RAM with ECC for optional user/firmware code
- Local and global registers to manage power-down, power-up, reset, clock gating, and power gating requests
- Interrupt controller with 16 interrupts from other modules and the inter-processor communication interface (IPI)
- GPI and GPO interfaces to and from PS I/O and PL
- JTAG interface for PMU debug
- Optional User-Defined Firmware

- Sleep Mode with automatic wake-up
- Snoop Mode
- 16-bit timestamping for receive messages
- Both internal generated reference clock and external reference clock input from MIO
- Guarantee clock sampling edge between 80 to 83% at 24MHz reference clock input
- Optional eFUSE disable per port

USB 2.0

- Two USB controllers (configurable as USB 2.0 or USB 3.0)
- Host, device and On-The-Go (OTG) modes
- High Speed, Full Speed, and Low Speed
- Up to 12 endpoints
- 8-bit ULPI External PHY Interface
- The USB host controller registers and data structures are compliant to Intel xHCI specifications.
- 64-bit AXI master port with built-in DMA
- Power management features: hibernation mode

Static Memory Interfaces

The static memory interfaces support external static memories.

- ONFI 3.1 NAND flash support with up to 24-bit ECC
- 1-bit SPI, 2-bit SPI, 4-bit SPI (Quad-SPI), or two Quad-SPI (8-bit) serial NOR flash
- 8-bit eMMC interface supporting managed NAND flash

NAND ONFI 3.1 Flash Controller

- ONFI 3.1 compliant
- Supports chip select reduction per ONFI 3.1 spec
- SLC NAND for boot/configuration and data storage
- ECC options based on SLC NAND
 - 1, 4, or 8 bits per 512+spare bytes
 - 24 bits per 1024+spare bytes
- Maximum throughput as follows
 - Asynchronous mode (SDR) 24.3MB/s
 - Synchronous mode (NV-DDR) 112MB/s (for 100MHz flash clock)
- 8-bit SDR NAND interface

- 2 chip selects
- Programmable access timing
- 1.8V and 3.3V I/O
- Built-in DMA for improved performance

Quad-SPI Controller

- 4 bytes (32-bit) and 3 bytes (24-bit) address width
- Maximum SPI Clock at Master Mode at 150MHz
- Single, Dual-Parallel, and Dual-Stacked mode
- 32-bit AXI Linear Address Mapping Interface for read operation
- Up to 2 chip select signals
- Write Protection Signal
- Hold signals
- 4-bit bidirectional I/O signals
- x1/x2/x4 Read speed required
- x1 write speed required only
- 64 byte Entry FIFO depth to improve QSPI read efficiency
- Built-in DMA for improved performance

Video Encoder/Decoder (VCU)

Zynq UltraScale+ MPSoCs include a Video codec (encoder/decoder) available in the devices designated with the EV suffix. The VCU is located in the PL and can be accessed from either the PL or PS.

- Simultaneous Encode and Decode through separate cores
- H.264 high profile level 5.2 (4Kx2K-60)
- H.265 (HEVC) main, main10 profile, level 5.1, high Tier, up to 4Kx2K-60 rate
- 8 and 10 bit encoding
- 4:2:0 and 4:2:2 chroma sampling
- 8Kx4K-15 rate
- Multi-stream up to total of 4Kx2K-60 rate
- Low Latency mode
- Can share the PS DRAM or use dedicated DRAM in the PL
- Clock/power management
- OpenMax Linux drivers

Interconnect

All the blocks are connected to each other and to the PL through a multi-layered ARM Advanced Microprocessor Bus Architecture (AMBA) AXI interconnect. The interconnect is non-blocking and supports multiple simultaneous master-slave transactions.

The interconnect is designed with latency sensitive masters, such as the ARM CPU, having the shortest paths to memory, and bandwidth critical masters, such as the potential PL masters, having high throughput connections to the slaves with which they need to communicate.

Traffic through the interconnect can be regulated through the Quality of Service (QoS) block in the interconnect. The QoS feature is used to regulate traffic generated by the CPU, DMA controller, and a combined entity representing the masters in the IOP.

PS Interfaces

PS interfaces include external interfaces going off-chip or signals going from PS to PL.

PS External Interfaces

The Zynq UltraScale+ MPSoC's external interfaces use dedicated pins that cannot be assigned as PL pins. These include:

- Clock, reset, boot mode, and voltage reference
- Up to 78 dedicated multiplexed I/O (MIO) pins, software-configurable to connect to any of the internal I/O peripherals and static memory controllers
- 32-bit or 64-bit DDR4/DDR3/DDR3L/LPDDR3 memories with optional ECC
- 32-bit LPDDR4 memory with optional ECC
- 4 channels (TX and RX pair) for transceivers

MIO Overview

The IOP peripherals communicate to external devices through a shared pool of up to 78 dedicated multiplexed I/O (MIO) pins. Each peripheral can be assigned one of several pre-defined groups of pins, enabling a flexible assignment of multiple devices simultaneously. Although 78 pins are not enough for simultaneous use of all the I/O peripherals, most IOP interface signals are available to the PL, allowing use of standard PL I/O pins when powered up and properly configured. Extended multiplexed I/O (EMIO) allows unmapped PS peripherals to access PL I/O.

Port mappings can appear in multiple locations. For example, there are up to 12 possible port mappings for CAN pins. The PS Configuration Wizard (PCW) tool aids in peripheral and static memory pin mapping.

High-Performance AXI Ports

The high-performance AXI4 ports provide access from the PL to DDR and high-speed interconnect in the PS. The six dedicated AXI memory ports from the PL to the PS are configurable as either 128-bit, 64-bit, or 32-bit interfaces. These interfaces connect the PL to the memory interconnect via a FIFO interface. Two of the AXI interfaces support I/O coherent access to the APU caches.

Each high-performance AXI port has these characteristics:

- Reduced latency between PL and processing system memory
- 1KB deep FIFO
- Configurable either as 128-bit, 64-bit, or 32-bit AXI interfaces
- Multiple AXI command issuing to DDR

Accelerator Coherency Port (ACP)

The Zynq UltraScale+ MPSoC accelerator coherency port (ACP) is a 64-bit AXI slave interface that provides connectivity between the APU and a potential accelerator function in the PL. The ACP directly connects the PL to the snoop control unit (SCU) of the ARM Cortex-A53 processors, enabling cache-coherent access to CPU data in the L2 cache. The ACP provides a low latency path between the PS and a PL-based accelerator when compared with a legacy cache flushing and loading scheme. The ACP only snoops access in the CPU L2 cache, providing coherency in hardware. It does not support coherency on the PL side. So this interface is ideal for a DMA or an accelerator in the PL that only requires coherency on the CPU cache memories. For example, if a MicroBlaze™ processor in the PL is attached to the ACP interface, the cache of MicroBlaze processor will not be coherent with Cortex-A53 caches.

AXI Coherency Extension (ACE)

The Zynq UltraScale+ MPSoC AXI coherency extension (ACE) is a 64-bit AXI4 slave interface that provides connectivity between the APU and a potential accelerator function in the PL. The ACE directly connects the PL to the snoop control unit (SCU) of the ARM Cortex-A53 processors, enabling cache-coherent access to Cache Coherent Interconnect (CCI). The ACE provides a low-latency path between the PS and a PL-based accelerator when compared with a legacy cache flushing and loading scheme. The ACE snoops accesses to the CCI and the PL side, thus, providing full coherency in hardware. This interface can be used to hook up a cached interface in the PL to the PS as caches on both the Cortex-A53 memories and the PL master are snooped thus providing full coherency. For example, if a MicroBlaze processor in the PL is hooked up using an ACE interface, then Cortex-A53 and MicroBlaze processor caches will be coherent with each other.

Out-of-Band Signaling

The transceivers provide out-of-band (OOB) signaling, often used to send low-speed signals from the transmitter to the receiver while high-speed serial data transmission is not active. This is typically done when the link is in a powered-down state or has not yet been initialized. This benefits PCIe and SATA/SAS and QPI applications.

Integrated Interface Blocks for PCI Express Designs

The MPSoC PL includes integrated blocks for PCIe technology that can be configured as an Endpoint or Root Port, compliant to the PCI Express Base Specification Revision 3.1 for Gen3 and lower data rates and compatible with the PCI Express Base Specification Revision 4.0 (rev 0.5) for Gen4 data rates. The Root Port can be used to build the basis for a compatible Root Complex, to allow custom chip-to-chip communication via the PCI Express protocol, and to attach ASSP Endpoint devices, such as Ethernet Controllers or Fibre Channel HBAs, to the MPSoC.

This block is highly configurable to system design requirements and can operate 1, 2, 4, 8, or 16 lanes at up to 2.5Gb/s, 5.0Gb/s, 8.0Gb/s, or 16Gb/s data rates. For high-performance applications, advanced buffering techniques of the block offer a flexible maximum payload size of up to 1,024 bytes. The integrated block interfaces to the integrated high-speed transceivers for serial connectivity and to block RAMs for data buffering. Combined, these elements implement the Physical Layer, Data Link Layer, and Transaction Layer of the PCI Express protocol.

Xilinx provides a light-weight, configurable, easy-to-use LogiCORE™ IP wrapper that ties the various building blocks (the integrated block for PCIe, the transceivers, block RAM, and clocking resources) into an Endpoint or Root Port solution. The system designer has control over many configurable parameters: link width and speed, maximum payload size, MPSoC logic interface speeds, reference clock frequency, and base address register decoding and filtering.

Integrated Block for Interlaken

Some UltraScale architecture-based devices include integrated blocks for Interlaken. Interlaken is a scalable chip-to-chip interconnect protocol designed to enable transmission speeds from 10Gb/s to 150Gb/s. The Interlaken integrated block in the UltraScale architecture is compliant to revision 1.2 of the Interlaken specification with data striping and de-striping across 1 to 12 lanes. Permitted configurations are: 1 to 12 lanes at up to 12.5Gb/s and 1 to 6 lanes at up to 25.78125Gb/s, enabling flexible support for up to 150Gb/s per integrated block. With multiple Interlaken blocks, certain UltraScale architecture-based devices enable easy, reliable Interlaken switches and bridges.

Integrated Block for 100G Ethernet

Compliant to the IEEE Std 802.3ba, the 100G Ethernet integrated blocks in the UltraScale architecture provide low latency 100Gb/s Ethernet ports with a wide range of user customization and statistics gathering. With support for 10 x 10.3125Gb/s (CAUI) and 4 x 25.78125Gb/s (CAUI-4) configurations, the integrated block includes both the 100G MAC and PCS logic with support for IEEE Std 1588v2 1-step and 2-step hardware timestamping.

In UltraScale+ devices, the 100G Ethernet blocks contain a Reed Solomon Forward Error Correction (RS-FEC) block, compliant to IEEE Std 802.3bj, that can be used with the Ethernet block or stand alone in user applications. These families also support OTN mapping mode in which the PCS can be operate without using the MAC.

Clock Management

The clock generation and distribution components in UltraScale architecture-based devices are located adjacent to the columns that contain the memory interfacing and input and output circuitry. This tight coupling of clocking and I/O provides low-latency clocking to the I/O for memory interfaces and other I/O protocols. Within every clock management tile (CMT) resides one mixed-mode clock manager (MMCM), two PLLs, clock distribution buffers and routing, and dedicated circuitry for implementing external memory interfaces.

Mixed-Mode Clock Manager

The mixed-mode clock manager (MMCM) can serve as a frequency synthesizer for a wide range of frequencies and as a jitter filter for incoming clocks. At the center of the MMCM is a voltage-controlled oscillator (VCO), which speeds up and slows down depending on the input voltage it receives from the phase frequency detector (PFD).

Three sets of programmable frequency dividers (D, M, and O) are programmable by configuration and during normal operation via the Dynamic Reconfiguration Port (DRP). The pre-divider D reduces the input frequency and feeds one input of the phase/frequency comparator. The feedback divider M acts as a multiplier because it divides the VCO output frequency before feeding the other input of the phase comparator. D and M must be chosen appropriately to keep the VCO within its specified frequency range. The VCO has eight equally-spaced output phases (0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°). Each phase can be selected to drive one of the output dividers, and each divider is programmable by configuration to divide by any integer from 1 to 128.

The MMCM has three input-jitter filter options: low bandwidth, high bandwidth, or optimized mode. Low-Bandwidth mode has the best jitter attenuation. High-Bandwidth mode has the best phase offset. Optimized mode allows the tools to find the best setting.

The MMCM can have a fractional counter in either the feedback path (acting as a multiplier) or in one output path. Fractional counters allow non-integer increments of 1/8 and can thus increase frequency synthesis capabilities by a factor of 8. The MMCM can also provide fixed or dynamic phase shift in small increments that depend on the VCO frequency. At 1,600MHz, the phase-shift timing increment is 11.2ps.

Configurable Logic Block

Every Configurable Logic Block (CLB) in the UltraScale architecture contains 8 LUTs and 16 flip-flops. The LUTs can be configured as either one 6-input LUT with one output, or as two 5-input LUTs with separate outputs but common inputs. Each LUT can optionally be registered in a flip-flop. In addition to the LUTs and flip-flops, the CLB contains arithmetic carry logic and multiplexers to create wider logic functions.

Each CLB contains one slice. There are two types of slices: SLICEL and SLICEM. LUTs in the SLICEM can be configured as 64-bit RAM, as 32-bit shift registers (SRL32), or as two SRL16s. CLBs in the UltraScale architecture have increased routing and connectivity compared to CLBs in previous-generation Xilinx devices. They also have additional control signals to enable superior register packing, resulting in overall higher device utilization.

Interconnect

Various length vertical and horizontal routing resources in the UltraScale architecture that span 1, 2, 4, 5, 12, or 16 CLBs ensure that all signals can be transported from source to destination with ease, providing support for the next generation of wide data buses to be routed across even the highest capacity devices while simultaneously improving quality of results and software run time.

Block RAM

Every UltraScale architecture-based device contains a number of 36Kb block RAMs, each with two completely independent ports that share only the stored data. Each block RAM can be configured as one 36Kb RAM or two independent 18Kb RAMs. Each memory access, read or write, is controlled by the clock. Connections in every block RAM column enable signals to be cascaded between vertically adjacent block RAMs, providing an easy method to create large, fast memory arrays, and FIFOs with greatly reduced power consumption.

All inputs, data, address, clock enables, and write enables are registered. The input address is always clocked (unless address latching is turned off), retaining data until the next operation. An optional output data pipeline register allows higher clock rates at the cost of an extra cycle of latency. During a write operation, the data output can reflect either the previously stored data or the newly written data, or it can remain unchanged. Block RAM sites that remain unused in the user design are automatically powered down to reduce total power consumption. There is an additional pin on every block RAM to control the dynamic power gating feature.

Programmable Data Width

Each port can be configured as $32K \times 1$; $16K \times 2$; $8K \times 4$; $4K \times 9$ (or 8); $2K \times 18$ (or 16); $1K \times 36$ (or 32); or 512×72 (or 64). Whether configured as block RAM or FIFO, the two ports can have different aspect ratios without any constraints. Each block RAM can be divided into two completely independent 18Kb block RAMs that can each be configured to any aspect ratio from $16K \times 1$ to 512×36 . Everything described previously for the full 36Kb block RAM also applies to each of the smaller 18Kb block RAMs. Only in simple dual-port (SDP) mode can data widths of greater than 18 bits (18Kb RAM) or 36 bits (36Kb RAM) be accessed. In this mode, one port is dedicated to read operation, the other to write operation. In SDP mode, one side (read or write) can be variable, while the other is fixed to 32/36 or 64/72. Both sides of the dual-port 36Kb RAM can be of variable width.

Error Detection and Correction

Each 64-bit-wide block RAM can generate, store, and utilize eight additional Hamming code bits and perform single-bit error correction and double-bit error detection (ECC) during the read process. The ECC logic can also be used when writing to or reading from external 64- to 72-bit-wide memories.

FIFO Controller

Each block RAM can be configured as a 36Kb FIFO or an 18Kb FIFO. The built-in FIFO controller for single-clock (synchronous) or dual-clock (asynchronous or multirate) operation increments the internal addresses and provides four handshaking flags: full, empty, programmable full, and programmable empty. The programmable flags allow the user to specify the FIFO counter values that make these flags go active. The FIFO width and depth are programmable with support for different read port and write port widths on a single FIFO. A dedicated cascade path allows for easy creation of deeper FIFOs.

UltraRAM

UltraRAM is a high-density, dual-port, synchronous memory block used in some UltraScale+ families. Both of the ports share the same clock and can address all of the $4K \times 72$ bits. Each port can independently read from or write to the memory array. UltraRAM supports two types of write enable schemes. The first mode is consistent with the block RAM byte write enable mode. The second mode allows gating the data and parity byte writes separately. Multiple UltraRAM blocks can be cascaded together to create larger memory arrays. UltraRAM blocks can be connected together to create larger memory arrays. Dedicated routing in the UltraRAM column enables the entire column height to be connected together. This makes UltraRAM an ideal solution for replacing external memories such as SRAM. Cascadable anywhere from 288Kb to 36Mb, UltraRAM provides the flexibility to fulfill many different memory requirements.

Error Detection and Correction

Each 64-bit-wide UltraRAM can generate, store and utilize eight additional Hamming code bits and perform single-bit error correction and double-bit error detection (ECC) during the read process.

Digital Signal Processing

DSP applications use many binary multipliers and accumulators, best implemented in dedicated DSP slices. All UltraScale architecture-based devices have many dedicated, low-power DSP slices, combining high speed with small size while retaining system design flexibility.

Each DSP slice fundamentally consists of a dedicated 27×18 bit twos complement multiplier and a 48-bit accumulator. The multiplier can be dynamically bypassed, and two 48-bit inputs can feed a single-instruction-multiple-data (SIMD) arithmetic unit (dual 24-bit add/subtract/accumulate or quad 12-bit add/subtract/accumulate), or a logic unit that can generate any one of ten different logic functions of the two operands.

The DSP includes an additional pre-adder, typically used in symmetrical filters. This pre-adder improves performance in densely packed designs and reduces the DSP slice count by up to 50%. The 96-bit-wide XOR function, programmable to 12, 24, 48, or 96-bit widths, enables performance improvements when implementing forward error correction and cyclic redundancy checking algorithms.

The DSP also includes a 48-bit-wide pattern detector that can be used for convergent or symmetric rounding. The pattern detector is also capable of implementing 96-bit-wide logic functions when used in conjunction with the logic unit.

The DSP slice provides extensive pipelining and extension capabilities that enhance the speed and efficiency of many applications beyond digital signal processing, such as wide dynamic bus shifters, memory address generators, wide bus multiplexers, and memory-mapped I/O register files. The accumulator can also be used as a synchronous up/down counter.

System Monitor

The System Monitor blocks in the UltraScale architecture are used to enhance the overall safety, security, and reliability of the system by monitoring the physical environment via on-chip power supply and temperature sensors.

All UltraScale architecture-based devices contain at least one System Monitor. The System Monitor in UltraScale+ devices is similar to the Kintex UltraScale and Virtex UltraScale devices but with the addition of a PMBus interface.

Zynq UltraScale+ MPSoCs contain one System Monitor in the PL and an additional block in the PS. The System Monitor in the PL has the same features as the block in UltraScale+ FPGAs. See [Table 11](#).

Table 11: Key System Monitor Features

	Zynq UltraScale+ MPSoC PL	Zynq UltraScale+ MPSoC PS
ADC	10-bit 200kSPS	10-bit 1MSPS
Interfaces	JTAG, I2C, DRP, PMBus	APB

PS Boot and Device Configuration

Zynq UltraScale+ MPSoCs use a multi-stage boot process that supports both a non-secure and a secure boot. The PS is the master of the boot and configuration process. For a secure boot, the AES-GCM, SHA-3/384 decrypts and authenticates the images while the 4096-bit RSA block authenticates the image.

Upon reset, the device mode pins are read to determine the primary boot device to be used: NAND, Quad-SPI, SD, eMMC, or JTAG. JTAG can only be used as a non-secure boot source and is intended for debugging purposes. The CSU executes code out of on-chip ROM and copies the first stage boot loader (FSBL) from the boot device to the OCM.

After copying the FSBL to OCM, one of the processors, either the Cortex-A53 or Cortex-R5, executes the FSBL. Xilinx supplies example FSBLs or users can create their own. The FSBL initiates the boot of the PS and can load and configure the PL, or configuration of the PL can be deferred to a later stage. The FSBL typically loads either a user application or an optional second stage boot loader (SSBL), such as U-Boot. Users obtain example SSBL from Xilinx or a third party, or they can create their own SSBL. The SSBL continues the boot process by loading code from any of the primary boot devices or from other sources such as USB, Ethernet, etc. If the FSBL did not configure the PL, the SSBL can do so, or again, the configuration can be deferred to a later stage.

The static memory interface controller (NAND, eMMC, or Quad-SPI) is configured using default settings. To improve device configuration speed, these settings can be modified by information provided in the boot image header. The ROM boot image is not user readable or callable after boot.

Hardware and Software Debug Support

The debug system used in Zynq UltraScale+ MPSoCs is based on the ARM CoreSight architecture. It uses ARM CoreSight components including an embedded trace controller (ETC), an embedded trace Macrocell (ETM) for each Cortex-A53 and Cortex-R5 processor, and a system trace Macrocell (STM). This enables advanced debug features like event trace, debug breakpoints and triggers, cross-trigger, and debug bus dump to memory. The programmable logic can be debugged with the Xilinx Vivado Logic Analyzer.

Debug Ports

Three JTAG ports are available and can be chained together or used separately. When chained together, a single port is used for chip-level JTAG functions, ARM processor code downloads and run-time control operations, PL configuration, and PL debug with the Vivado Logic Analyzer. This enables tools such as the Xilinx Software Development Kit (SDK) and Vivado Logic Analyzer to share a single download cable from Xilinx.

When the JTAG chain is split, one port is used to directly access the ARM DAP interface. This CoreSight interface enables the use of ARM-compliant debug and software development tools such as Development Studio 5 (DS-5™). The other JTAG port can then be used by the Xilinx FPGA tools for access to the PL, including configuration bitstream downloads and PL debug with the Vivado Logic Analyzer. In this mode, users can download to and debug the PL in the same manner as a stand-alone FPGA.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
02/15/2017	1.4	Updated DSP count in Table 1 , Table 3 , and Table 5 . Updated I/O Electrical Characteristics . Updated Table 12 with -2E speed grade.
09/23/2016	1.3	Updated Table 2 ; Table 3 ; Table 4 ; Table 6 ; Graphics Processing Unit (GPU) ; and NAND ONFI 3.1 Flash Controller .
06/03/2016	1.2	Added CG devices: Updated Table 1 ; Table 2 ; Table 3 ; Table 4 ; Table 5 ; Table 6 ; and Table 12 . Added Video Encoder/Decoder (VCU) ; Table 7 ; and Power Examples (removed XPE Computed Range table). Updated: General Description ; ARM Cortex-A53 Based Application Processing Unit (APU) ; Zynq UltraScale+ MPSoCs ; Dynamic Memory Controller (DDRC) ; and Figure 3 .
01/28/2016	1.1	Updated Table 1 and Table 2 .
11/24/2015	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>.

This document contains preliminary information and is subject to change without notice. Information provided herein relates to products and/or services not yet available for sale, and provided solely for information purposes and are not intended, or to be construed, as an offer for sale or an attempted commercialization of the products and/or services referred to herein.

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.