

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, 154K+ Logic Cells
Operating Temperature	-40°C ~ 100°C (Tj)
Package / Case	784-BFBGA, FCBGA
Supplier Device Package	784-FCBGA (23x23)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xczu3eg-1sfvc784i

Programmable Logic (PL)

Configurable Logic Blocks (CLB)

- Look-up tables (LUT)
- Flip-flops
- Cascadable adders

36Kb Block RAM

- True dual-port
- Up to 72 bits wide
- Configurable as dual 18Kb

UltraRAM

- 288Kb dual-port
- 72 bits wide
- Error checking and correction

DSP Blocks

- 27 x 18 signed multiply
- 48-bit adder/accumulator
- 27-bit pre-adder

Programmable I/O Blocks

- Supports LVCMOS, LVDS, and SSTL
- 1.0V to 3.3V I/O
- Programmable I/O delay and SerDes

JTAG Boundary-Scan

- IEEE Std 1149.1 Compatible Test Interface

PCI Express

- Supports Root complex and End Point configurations
- Supports up to Gen4 speeds
- Up to five integrated blocks in select devices

100G Ethernet MAC/PCS

- IEEE Std 802.3 compliant
- CAUI-10 (10x 10.3125Gb/s) or CAUI-4 (4x 25.78125Gb/s)
- RSFEC (IEEE Std 802.3bj) in CAUI-4 configuration
- Up to four integrated blocks in select devices

Interlaken

- Interlaken spec 1.2 compliant
- 64/67 encoding
- 12 x 12.5Gb/s or 6 x 25Gb/s
- Up to four integrated blocks in select devices

Video Encoder/Decoder (VCU)

- Available in EV devices
- Accessible from either PS or PL
- Simultaneous encode and decode
- H.264 and H.265 support

System Monitor in PL

- On-chip voltage and temperature sensing
- 10-bit 200KSPS ADC with up to 17 external inputs

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 2: Zynq UltraScale+ MPSoC: CG Device-Package Combinations and Maximum I/Os

Package (1)(2)(3)(4)(5)	Package Dimensions (mm)	ZU2CG	ZU3CG	ZU4CG	ZU5CG	ZU6CG	ZU7CG	ZU9CG
		HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY
SBVA484 ⁽⁶⁾	19x19	24, 58 0, 0	24, 58 0, 0					
SFVA625	21x21	24, 156 0, 0	24, 156 0, 0					
SFVC784 ⁽⁷⁾	23x23	96, 156 0, 0	96, 156 0, 0	96, 156 4, 0	96, 156 4, 0			
FBVB900	31x31			48, 156 16, 0	48, 156 16, 0		48, 156 16, 0	
FFVC900	31x31					48, 156 16, 0		48, 156 16, 0
FFVB1156	35x35					120, 208 24, 0		120, 208 24, 0
FFVC1156	35x35						48, 312 20, 0	
FFVF1517	40x40						48, 416 24, 0	

Notes:

1. Go to [Ordering Information](#) for package designation details.
2. FB/FF packages have 1.0mm ball pitch. SB/SF packages have 0.8mm ball pitch.
3. All device package combinations bond out 4 PS-GTR transceivers.
4. All device package combinations bond out 214 PS I/O except ZU2CG and ZU3CG in the SBVA484 and SFVA625 packages, which bond out 170 PS I/Os.
5. Packages with the same last letter and number sequence, e.g., A484, are footprint compatible with all other UltraScale devices with the same sequence. The footprint compatible devices within this family are outlined.
6. All 58 HP I/O pins are powered by the same V_{CCO} supply.
7. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s.

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.

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

- 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

Configuration Security Unit (CSU)

- Triple redundant Secure Processor Block (SPB) with built-in ECC
- Crypto Interface Block consisting of
 - 256-bit AES-GCM
 - SHA-3/384
 - 4096-bit RSA
- Key Management Unit
- Built-in DMA
- PCAP interface
- Supports ROM validation during pre-configuration stage
- Loads First Stage Boot Loader (FSBL) into OCM in either secure or non-secure boot modes
- Supports voltage, temperature, and frequency monitoring after configuration

Xilinx Peripheral Protection Unit (XPPU)

- Provides peripheral protection support
- Up to 20 masters simultaneously
- Multiple aperture sizes
- Access control for a specified set of address apertures on a per master basis
- 64KB peripheral apertures and controls access on per peripheral basis

I/O Peripherals

The IOP unit contains the data communication peripherals. Key features of the IOP include:

Triple-Speed Gigabit Ethernet

- Compatible with IEEE Std 802.3 and supports 10/100/1000Mb/s transfer rates (Full and Half duplex)
- Supports jumbo frames
- Built-in Scatter-Gather DMA capability
- Statistics counter registers for RMON/MIB
- Multiple I/O types (1.8, 2.5, 3.3V) on RGMII interface with external PHY
- GMII interface to PL to support interfaces as: TBI, SGMII, and RGMII v2.0 support
- Automatic pad and cyclic redundancy check (CRC) generation on transmitted frames
- Transmitter and Receive IP, TCP, and UDP checksum offload
- MDIO interface for physical layer management

SPI

- Full-duplex operation offers simultaneous receive and transmit
- 128B deep read and write FIFO
- Master or slave SPI mode
- Up to 3 chip select lines
- Multi-master environment
- Identifies an error condition if more than one master detected
- Selectable master clock reference
- Software can poll for status or be interrupt driven

I2C

- 128-bit buffer size
- Both normal (100kHz) and fast bus data rates (400kHz)
- Master or slave mode
- Normal or extended addressing
- I2C bus hold for slow host service

GPIO

- Up to 128 GPIO bits
 - Up to 78-bits from MIO and 96-bits from EMIO
- Each GPIO bit can be dynamically programmed as input or output
- Independent reset values for each bit of all registers
- Interrupt request generation for each GPIO signals
- Single Channel (Bit) write capability for all control registers include data output register, direction control register, and interrupt clear register
- Read back in output mode

CAN

- Conforms to the ISO 11898 -1, CAN2.0A, and CAN 2.0B standards
- Both standard (11-bit identifier) and extended (29-bit identifier) frames
- Bit rates up to 1Mb/s
- Transmit and Receive message FIFO with a depth of 64 messages
- Watermark interrupts for TXFIFO and RXFIFO
- Automatic re-transmission on errors or arbitration loss in normal mode
- Acceptance filtering of 4 acceptance filters

HS-MIO

The function of the HS-MIO is to multiplex access from the high-speed PS peripheral to the differential pair on the PS-GTR transceiver as defined in the configuration registers. Up to 4 channels of the transceiver are available for use by the high-speed interfaces in the PS.

Table 9: HS-MIO Peripheral Interface Mapping

Peripheral Interface	Lane0	Lane1	Lane2	Lane3
PCIe (x1, x2 or x4)	PCIe0	PCIe1	PCIe2	PCIe3
SATA (1 or 2 channels)	SATA0	SATA1	SATA0	SATA1
DisplayPort (TX only)	DP1	DPO	DP1	DPO
USB0	USB0	USB0	USB0	–
USB1	–	–	–	USB1
SGMII0	SGMII0	–	–	–
SGMII1	–	SGMII1	–	–
SGMII2	–	–	SGMII2	–
SGMII3	–	–	–	SGMII3

PS-PL Interface

The PS-PL interface includes:

- AMBA AXI4 interfaces for primary data communication
 - Six 128-bit/64-bit/32-bit High Performance (HP) Slave AXI interfaces from PL to PS.
 - Four 128-bit/64-bit/32-bit HP AXI interfaces from PL to PS DDR.
 - Two 128-bit/64-bit/32-bit high-performance coherent (HPC) ports from PL to cache coherent interconnect (CCI).
 - Two 128-bit/64-bit/32-bit HP Master AXI interfaces from PS to PL.
 - One 128-bit/64-bit/32-bit interface from PL to RPU in PS (PL_LPD) for low latency access to OCM.
 - One 128-bit/64-bit/32-bit AXI interface from RPU in PS to PL (LPD_PL) for low latency access to PL.
 - One 128-bit AXI interface (ACP port) for I/O coherent access from PL to Cortex-A53 cache memory. This interface provides coherency in hardware for Cortex-A53 cache memory.
 - One 128-bit AXI interface (ACE Port) for Fully coherent access from PL to Cortex-A53. This interface provides coherency in hardware for Cortex-A53 cache memory and the PL.
- Clocks and resets
 - Four PS clock outputs to the PL with start/stop control.
 - Four PS reset outputs to the PL.

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.

Programmable Logic

This section covers the information about blocks in the Programmable Logic (PL).

Device Layout

UltraScale architecture-based devices are arranged in a column-and-grid layout. Columns of resources are combined in different ratios to provide the optimum capability for the device density, target market or application, and device cost. At the core of UltraScale+ MPSoCs is the processing system that displaces some of the full or partial columns of programmable logic resources. [Figure 1](#) shows a device-level view with resources grouped together. For simplicity, certain resources such as the processing system, integrated blocks for PCIe, configuration logic, and System Monitor are not shown.



DS891_01_012915

Figure 1: Device with Columnar Resources

Resources within the device are divided into segmented clock regions. The height of a clock region is 60 CLBs. A bank of 52 I/Os, 24 DSP slices, 12 block RAMs, or 4 transceiver channels also matches the height of a clock region. The width of a clock region is essentially the same in all cases, regardless of device size or the mix of resources in the region, enabling repeatable timing results. Each segmented clock region contains vertical and horizontal clock routing that span its full height and width. These horizontal and vertical clock routes can be segmented at the clock region boundary to provide a flexible, high-performance, low-power clock distribution architecture. [Figure 2](#) is a representation of a device divided into regions.

3-State Digitally Controlled Impedance and Low Power I/O Features

The 3-state Digitally Controlled Impedance (T_DCI) can control the output drive impedance (series termination) or can provide parallel termination of an input signal to V_{CCO} or split (Thevenin) termination to $V_{CCO}/2$. This allows users to eliminate off-chip termination for signals using T_DCI. In addition to board space savings, the termination automatically turns off when in output mode or when 3-stated, saving considerable power compared to off-chip termination. The I/Os also have low power modes for IBUF and IDELAY to provide further power savings, especially when used to implement memory interfaces.

I/O Logic

Input and Output Delay

All inputs and outputs can be configured as either combinatorial or registered. Double data rate (DDR) is supported by all inputs and outputs. Any input or output can be individually delayed by up to 1,250ps of delay with a resolution of 5–15ps. Such delays are implemented as IDELAY and ODELAY. The number of delay steps can be set by configuration and can also be incremented or decremented while in use. The IDELAY and ODELAY can be cascaded together to double the amount of delay in a single direction.

ISERDES and OSERDES

Many applications combine high-speed, bit-serial I/O with slower parallel operation inside the device. This requires a serializer and deserializer (SerDes) inside the I/O logic. Each I/O pin possesses an IOSERDES (ISERDES and OSERDES) capable of performing serial-to-parallel or parallel-to-serial conversions with programmable widths of 2, 4, or 8 bits. These I/O logic features enable high-performance interfaces, such as Gigabit Ethernet/1000BaseX/SGMII, to be moved from the transceivers to the SelectIO interface.

High-Speed Serial Transceivers

Ultra-fast serial data transmission between devices on the same PCB, over backplanes, and across even longer distances is becoming increasingly important for scaling to 100 Gb/s and 400 Gb/s line cards. Specialized dedicated on-chip circuitry and differential I/O capable of coping with the signal integrity issues are required at these high data rates.

Three types of transceivers are used in Zynq UltraScale+ MPSoCs: GTH, GTY, and PS-GTR. All transceivers are arranged in groups of four, known as a transceiver Quad. Each serial transceiver is a combined transmitter and receiver. [Table 10](#) compares the available transceivers.

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.

PLL

With fewer features than the MMCM, the two PLLs in a clock management tile are primarily present to provide the necessary clocks to the dedicated memory interface circuitry. The circuit at the center of the PLLs is similar to the MMCM, with PFD feeding a VCO and programmable M, D, and O counters. There are two divided outputs to the device fabric per PLL as well as one clock plus one enable signal to the memory interface circuitry.

Zynq UltraScale+ MPSoCs are equipped with five additional PLLs in the PS for independently configuring the four primary clock domains with the PS: the APU, the RPU, the DDR controller, and the I/O peripherals.

Clock Distribution

Clocks are distributed throughout Zynq UltraScale+ MPSoCs via buffers that drive a number of vertical and horizontal tracks. There are 24 horizontal clock routes per clock region and 24 vertical clock routes per clock region with 24 additional vertical clock routes adjacent to the MMCM and PLL. Within a clock region, clock signals are routed to the device logic (CLBs, etc.) via 16 gateable leaf clocks.

Several types of clock buffers are available. The BUFGCE and BUFCE_LEAF buffers provide clock gating at the global and leaf levels, respectively. BUFGCTRL provides glitchless clock muxing and gating capability. BUFGCE_DIV has clock gating capability and can divide a clock by 1 to 8. BUFG_GT performs clock division from 1 to 8 for the transceiver clocks. In MPSoCs, clocks can be transferred from the PS to the PL using dedicated buffers.

Memory Interfaces

Memory interface data rates continue to increase, driving the need for dedicated circuitry that enables high performance, reliable interfacing to current and next-generation memory technologies. Every Zynq UltraScale+ MPSoC includes dedicated physical interfaces (PHY) blocks located between the CMT and I/O columns that support implementation of high-performance PHY blocks to external memories such as DDR4, DDR3, QDRII+, and RLDRAM3. The PHY blocks in each I/O bank generate the address/control and data bus signaling protocols as well as the precision clock/data alignment required to reliably communicate with a variety of high-performance memory standards. Multiple I/O banks can be used to create wider memory interfaces.

As well as external parallel memory interfaces, Zynq UltraScale+ MPSoC can communicate to external serial memories, such as Hybrid Memory Cube (HMC), via the high-speed serial transceivers. All transceivers in the UltraScale architecture support the HMC protocol, up to 15Gb/s line rates. UltraScale architecture-based devices support the highest bandwidth HMC configuration of 64 lanes with a single device.

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

Ordering Information

Table 12 shows the speed and temperature grades available in the different device families.

Table 12: Speed Grade and Temperature Grade

Device Family	Devices	Speed Grade and Temperature Grade				
		Commercial (C)	Extended (E)		Industrial (I)	
		0°C to +85°C	0°C to +100°C	0°C to +110°C	-40°C to +100°C	
Zynq UltraScale+	CG Devices		-2E (0.85V)		-2I (0.85V)	
				-2LE ⁽¹⁾⁽²⁾ (0.85V or 0.72V)		
			-1E (0.85V)		-1I (0.85V)	
					-1LI ⁽²⁾ (0.85V or 0.72V)	
	ZU2EG ZU3EG		-2E (0.85V)		-2I (0.85V)	
				-2LE ⁽¹⁾⁽²⁾ (0.85V or 0.72V)		
			-1E (0.85V)		-1I (0.85V)	
					-1LI ⁽²⁾ (0.85V or 0.72V)	
	ZU4EG ZU5EG ZU6EG ZU7EG ZU9EG ZU11EG ZU15EG ZU17EG ZU19EG		-3E (0.90V)			
			-2E (0.85V)		-2I (0.85V)	
				-2LE ⁽¹⁾⁽²⁾ (0.85V or 0.72V)		
			-1E (0.85V)		-1I (0.85V)	
					-1LI ⁽²⁾ (0.85V or 0.72V)	
		EV Devices		-3E (0.90V)		
				-2E (0.85V)		-2I (0.85V)
					-2LE ⁽¹⁾⁽²⁾ (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)	
					-1LI ⁽²⁾ (0.85V or 0.72V)	

Notes:

1. In -2LE speed/temperature grade, devices can operate for a limited time with junction temperature of 110°C. Timing parameters adhere to the same speed file at 110°C as they do below 110°C, regardless of operating voltage (nominal at 0.85V or low voltage at 0.72V). Operation at 110°C Tj is limited to 1% of the device lifetime and can occur sequentially or at regular intervals as long as the total time does not exceed 1% of device lifetime.
2. In Zynq UltraScale+ MPSoCs, when operating the PL at low voltage (0.72V), the PS operates at nominal voltage (0.85V)

The ordering information shown in Figure 3 applies to all packages in the Zynq UltraScale+ MPSoCs.

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.