

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 <sup>2</sup> C, MMC/SD/SDIO, SPI, UART/USART, USB OTG
Speed	500MHz, 600MHz, 1.2GHz
Primary Attributes	Zynq@UltraScale+™ FPGA, 469K+ Logic Cells
Operating Temperature	-40°C ~ 100°C (Tj)
Package / Case	900-BBGA, FCBGA
Supplier Device Package	900-FCBGA (31x31)
Purchase URL	<a href="https://www.e-xfl.com/product-detail/xilinx/xczu6eg-l1ffvc900i">https://www.e-xfl.com/product-detail/xilinx/xczu6eg-l1ffvc900i</a>

Table 4: Zynq UltraScale+ MPSoC: EG Device-Package Combinations and Maximum I/Os

Package (1)(2)(3)(4)(5)	Package Dimensions (mm)	ZU2EG	ZU3EG	ZU4EG	ZU5EG	ZU6EG	ZU7EG	ZU9EG	ZU11EG	ZU15EG	ZU17EG	ZU19EG
		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	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		48, 156 16, 0		
FFVB1156	35x35					120, 208 24, 0		120, 208 24, 0		120, 208 24, 0		
FFVC1156	35x35						48, 312 20, 0		48, 312 20, 0			
FFVB1517	40x40								72, 416 16, 0		72, 572 16, 0	72, 572 16, 0
FFVF1517	40x40						48, 416 24, 0		48, 416 32, 0			
FFVC1760	42.5x42.5								96, 416 32, 16		96, 416 32, 16	96, 416 32, 16
FFVD1760	42.5x42.5										48, 260 44, 28	48, 260 44, 28
FFVE1924	45x45										96, 572 44, 0	96, 572 44, 0

**Notes:**

1. Go to [Ordering Information](#) for package designation details.(5)
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 ZU2EG and ZU3EG 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  $V_{CC0}$  supply.
7. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s.

## Zynq UltraScale+ MPSoCs

A comprehensive device family, Zynq UltraScale+ MPSoCs offer single-chip, all programmable, heterogeneous multiprocessors that provide designers with software, hardware, interconnect, power, security, and I/O programmability. The range of devices in the Zynq UltraScale+ MPSoC family allows designers to target cost-sensitive as well as high-performance applications from a single platform using industry-standard tools. While each Zynq UltraScale+ MPSoC contains the same PS, the PL, Video hard blocks, and I/O resources vary between the devices.

Table 7: Zynq UltraScale+ MPSoC Device Features

	CG Devices	EG Devices	EV Devices
APU	Dual-core ARM Cortex-A53	Quad-core ARM Cortex-A53	Quad-core ARM Cortex-A53
RPU	Dual-core ARM Cortex-R5	Dual-core ARM Cortex-R5	Dual-core ARM Cortex-R5
GPU	–	Mali-400MP2	Mali-400MP2
VCU	–	–	H.264/H.265

The Zynq UltraScale+ MPSoCs are able to serve a wide range of applications including:

- Automotive: Driver assistance, driver information, and infotainment
- Wireless Communications: Support for multiple spectral bands and smart antennas
- Wired Communications: Multiple wired communications standards and context-aware network services
- Data Centers: Software Defined Networks (SDN), data pre-processing, and analytics
- Smarter Vision: Evolving video-processing algorithms, object detection, and analytics
- Connected Control/M2M: Flexible/adaptable manufacturing, factory throughput, quality, and safety

The UltraScale MPSoC architecture provides processor scalability from 32 to 64 bits with support for virtualization, the combination of soft and hard engines for real-time control, graphics/video processing, waveform and packet processing, next-generation interconnect and memory, advanced power management, and technology enhancements that deliver multi-level security, safety, and reliability. Xilinx offers a large number of soft IP for the Zynq UltraScale+ MPSoC family. Stand-alone and Linux device drivers are available for the peripherals in the PS and the PL. Xilinx's Vivado® Design Suite, SDK™, and PetaLinux development environments enable rapid product development for software, hardware, and systems engineers. The ARM-based PS also brings a broad range of third-party tools and IP providers in combination with Xilinx's existing PL ecosystem.

The Zynq UltraScale+ MPSoC family delivers unprecedented processing, I/O, and memory bandwidth in the form of an optimized mix of heterogeneous processing engines embedded in a next-generation, high-performance, on-chip interconnect with appropriate on-chip memory subsystems. The heterogeneous processing and programmable engines, which are optimized for different application tasks, enable the Zynq UltraScale+ MPSoCs to deliver the extensive performance and efficiency required to address next-generation smarter systems while retaining backwards compatibility with the original Zynq-7000 All Programmable SoC family. The UltraScale MPSoC architecture also incorporates multiple levels of security, increased safety, and advanced power management, which are critical requirements of next-generation smarter systems. Xilinx's embedded UltraFast™ design methodology fully exploits the

- Low power modes
  - Active/precharge power down
  - Self-refresh, including clean exit from self-refresh after a controller power cycle
- Enhanced DDR training by allowing software to measure read/write eye and make delay adjustments dynamically
- Independent performance monitors for read path and write path
- Integration of PHY Debug Access Port (DAP) into JTAG for testing

The DDR memory controller is multi-ported and enables the PS and the PL to have shared access to a common memory. The DDR controller features six AXI slave ports for this purpose:

- Two 128-bit AXI ports from the ARM Cortex-A53 CPU(s), RPU (ARM Cortex-R5 and LPD peripherals), GPU, high speed peripherals (USB3, PCIe & SATA), and High Performance Ports (HP0 & HP1) from the PL through the Cache Coherent Interconnect (CCI)
- One 64-bit port is dedicated for the ARM Cortex-R5 CPU(s)
- One 128-bit AXI port from the DisplayPort and HP2 port from the PL
- One 128-bit AXI port from HP3 and HP4 ports from the PL
- One 128-bit AXI port from General DMA and HP5 from the PL

## High-Speed Connectivity Peripherals

### *PCIe*

- Compliant with the PCI Express Base Specification 2.1
- Fully compliant with PCI Express transaction ordering rules
- Lane width: x1, x2, or x4 at Gen1 or Gen2 rates
- 1 Virtual Channel
- Full duplex PCIe port
- End Point and single PCIe link Root Port
- Root Port supports Enhanced Configuration Access Mechanism (ECAM), Cfg Transaction generation
- Root Port support for INTx, and MSI
- Endpoint support for MSI or MSI-X
  - 1 physical function, no SR-IOV
  - No relaxed or ID ordering
  - Fully configurable BARs
  - INTx not recommended, but can be generated
  - Endpoint to support configurable target/slave apertures with address translation and Interrupt capability

## **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

## 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

- 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

## 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)**

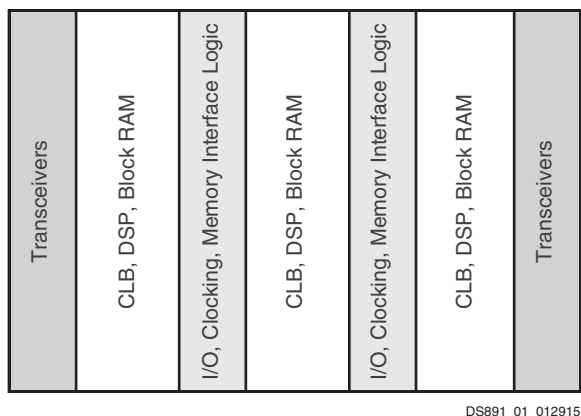
The Zynq UltraScale+ MPSoC AXI coherency extension (ACE) is a 64-bit AXI4 slave interface that provides connectivity between the APU and a potential accelerator function in the PL. The ACE directly connects the PL to the snoop control unit (SCU) of the ARM Cortex-A53 processors, enabling cache-coherent 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.




Figure 2: Column-Based Device Divided into Clock Regions

## Input/Output

All Zynq UltraScale+ MPSoCs have I/O pins for communicating to external components. In addition, in the MPSoC's PS, there are another 78 I/Os that the I/O peripherals use to communicate to external components, referred to as multiplexed I/O (MIO). If more than 78 pins are required by the I/O peripherals, the I/O pins in the PL can be used to extend the MPSoC interfacing capability, referred to as extended MIO (EMIO).

The number of I/O pins in the PL of Zynq UltraScale+ MPSoCs varies depending on device and package. Each I/O is configurable and can comply with a large number of I/O standards. The I/Os are classed as high-performance (HP), or high-density (HD). The HP I/Os are optimized for highest performance operation, from 1.0V to 1.8V. The HD I/Os are reduced-feature I/Os organized in banks of 24, providing voltage support from 1.2V to 3.3V.

All I/O pins are organized in banks, with 52 HP pins per bank or 24 HD pins per bank. Each bank has one common  $V_{CC0}$  output buffer power supply, which also powers certain input buffers. Some single-ended input buffers require an internally generated or an externally applied reference voltage ( $V_{REF}$ ).  $V_{REF}$  pins can be driven directly from the PCB or internally generated using the internal  $V_{REF}$  generator circuitry present in each bank.

### I/O Electrical Characteristics

Single-ended outputs use a conventional CMOS push/pull output structure driving High towards  $V_{CC0}$  or Low towards ground, and can be put into a high-Z state. The system designer can specify the slew rate and the output strength. The input is always active but is usually ignored while the output is active. Each pin can optionally have a weak pull-up or a weak pull-down resistor.

Most signal pin pairs can be configured as differential input pairs or output pairs. Differential input pin pairs can optionally be terminated with a  $100\Omega$  internal resistor. All UltraScale architecture-based devices support differential standards beyond LVDS, including RSDS, BLVDS, differential SSTL, and differential HSTL. Each of the I/Os supports memory I/O standards, such as single-ended and differential HSTL as well as single-ended and differential SSTL. The Zynq UltraScale+ family includes support for MIPI with a dedicated D-PHY in the I/O bank.

### ***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.

**Table 10: Transceiver Information**

	<b>Zynq UltraScale+ MPSoCs</b>		
Type	PS-GTR	GTH	GTY
Qty	4	0–44	0–28
Max. Data Rate	6.0Gb/s	16.3Gb/s	32.75Gb/s
Min. Data Rate	1.25Gb/s	0.5Gb/s	0.5Gb/s
Applications	<ul style="list-style-type: none"> <li>• PCIe Gen2</li> <li>• USB</li> <li>• Ethernet</li> </ul>	<ul style="list-style-type: none"> <li>• Backplane</li> <li>• PCIe Gen4</li> <li>• HMC</li> </ul>	<ul style="list-style-type: none"> <li>• 100G+ Optics</li> <li>• Chip-to-Chip</li> <li>• 25G+ Backplane</li> <li>• HMC</li> </ul>

The following information in this section pertains to the GTH and GTY only.

The serial transmitter and receiver are independent circuits that use an advanced phase-locked loop (PLL) architecture to multiply the reference frequency input by certain programmable numbers between 4 and 25 to become the bit-serial data clock. Each transceiver has a large number of user-definable features and parameters. All of these can be defined during device configuration, and many can also be modified during operation.

## Transmitter

The transmitter is fundamentally a parallel-to-serial converter with a conversion ratio of 16, 20, 32, 40, 64, or 80 for the GTH and 16, 20, 32, 40, 64, 80, 128, or 160 for the GTY. This allows the designer to trade off datapath width against timing margin in high-performance designs. These transmitter outputs drive the PC board with a single-channel differential output signal. TXOUTCLK is the appropriately divided serial data clock and can be used directly to register the parallel data coming from the internal logic. The incoming parallel data is fed through an optional FIFO and has additional hardware support for the 8B/10B, 64B/66B, or 64B/67B encoding schemes to provide a sufficient number of transitions. The bit-serial output signal drives two package pins with differential signals. This output signal pair has programmable signal swing as well as programmable pre- and post-emphasis to compensate for PC board losses and other interconnect characteristics. For shorter channels, the swing can be reduced to reduce power consumption.

## Receiver

The receiver is fundamentally a serial-to-parallel converter, changing the incoming bit-serial differential signal into a parallel stream of words, each 16, 20, 32, 40, 64, or 80 bits in the GTH or 16, 20, 32, 40, 64, 80, 128, or 160 for the GTY. This allows the designer to trade off internal datapath width against logic timing margin. The receiver takes the incoming differential data stream, feeds it through programmable DC automatic gain control, linear and decision feedback equalizers (to compensate for PC board, cable, optical and other interconnect characteristics), and uses the reference clock input to initiate clock recognition. There is no need for a separate clock line. The data pattern uses non-return-to-zero (NRZ) encoding and optionally ensures sufficient data transitions by using the selected encoding scheme. Parallel data is then transferred into the device logic using the RXUSRCLK clock. For short channels, the transceivers offer a special low-power mode (LPM) to reduce power consumption by approximately 30%. The receiver DC automatic gain control and linear and decision feedback equalizers can optionally “auto-adapt” to automatically learn and compensate for different interconnect characteristics. This enables even more margin for tough 10G+ and 25G+ backplanes.

---

## 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.

## 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  $27 \times 18$  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

## Clock Management

The PS in Zynq UltraScale+ MPSoCs is equipped with five phase-locked loops (PLLs), providing flexibility in configuring the clock domains within the PS. There are four primary clock domains of interest within the PS. These include the APU, the RPU, the DDR controller, and the I/O peripherals (IOP). The frequencies of all of these domains can be configured independently under software control.

## Power Domains

The Zynq UltraScale+ MPSoC contains four separate power domains. When they are connected to separate power supplies, they can be completely powered down independently of each other without consuming any dynamic or static power. The processing system includes:

- Full Power Domain (FPD)
- Low Power Domain (LPD)
- Battery Powered Domain (BPD)

In addition to these three Processing System power domains, the PL can also be completely powered down if connected to separate power supplies.

The Full Power Domain (FPD) consists of the following major blocks:

- Application Processing Unit (APU)
- DMA (FP-DMA)
- Graphics Processing Unit (GPU)
- Dynamic Memory Controller (DDRC)
- High-Speed I/O Peripherals

The Low Power Domain (LPD) consists of the following major blocks:

- Real-Time Processing Unit (RPU)
- DMA (LP-DMA)
- Platform Management Unit (PMU)
- Configuration Security Unit (CSU)
- Low-Speed I/O Peripherals
- Static Memory Interfaces

The Battery Power Domain (BPD) is the lowest power domain of the Zynq UltraScale+ MPSoC processing system. In this mode, all the PS is powered off except the Real-Time Clock (RTC) and battery-backed RAM (BBRAM).

### ***Power Examples***

Power for the Zynq UltraScale+ MPSoCs varies depending on the utilization of the PL resources, and the frequency of the PS and PL. To estimate power, use the Xilinx Power Estimator (XPE) at:

[http://www.xilinx.com/products/design\\_tools/logic\\_design/xpe.htm](http://www.xilinx.com/products/design_tools/logic_design/xpe.htm)

## 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 <a href="#">Table 1</a> , <a href="#">Table 3</a> , and <a href="#">Table 5</a> . Updated <a href="#">I/O Electrical Characteristics</a> . Updated <a href="#">Table 12</a> with -2E speed grade.
09/23/2016	1.3	Updated <a href="#">Table 2</a> ; <a href="#">Table 3</a> ; <a href="#">Table 4</a> ; <a href="#">Table 6</a> ; <a href="#">Graphics Processing Unit (GPU)</a> ; and <a href="#">NAND ONFI 3.1 Flash Controller</a> .
06/03/2016	1.2	Added CG devices: Updated <a href="#">Table 1</a> ; <a href="#">Table 2</a> ; <a href="#">Table 3</a> ; <a href="#">Table 4</a> ; <a href="#">Table 5</a> ; <a href="#">Table 6</a> ; and <a href="#">Table 12</a> . Added <a href="#">Video Encoder/Decoder (VCU)</a> ; <a href="#">Table 7</a> ; and <a href="#">Power Examples</a> (removed XPE Computed Range table). Updated: <a href="#">General Description</a> ; <a href="#">ARM Cortex-A53 Based Application Processing Unit (APU)</a> ; <a href="#">Zynq UltraScale+ MPSoCs</a> ; <a href="#">Dynamic Memory Controller (DDRC)</a> ; and <a href="#">Figure 3</a> .
01/28/2016	1.1	Updated <a href="#">Table 1</a> and <a href="#">Table 2</a> .
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.