



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	MPU, FPGA
Core Processor	Quad ARM® Cortex®-A53 MPCore™ with CoreSight™, Dual ARM®Cortex™-R5 with CoreSight™, ARM Mali™ -400 MP2
Flash Size	-
RAM Size	1.2MB
Peripherals	DMA, WDT
Connectivity	CANbus, I ² C, SPI, UART/USART, USB
Speed	500MHz, 1.2GHz
Primary Attributes	Zynq@UltraScale+™ FPGA, 103K+ Logic Cells
Operating Temperature	-40°C ~ 125°C (Tj)
Package / Case	784-BFBGA, FCBGA
Supplier Device Package	784-FCBGA (23x23)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xazu2eg-1sfvc784q

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 500MHz
 - 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)
- 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 600MHz
- Single Geometry Processor and two Pixel processor
- Vertex processing: 66 M Triangles/s
- Pixel processing: 1.2 G Pixels/s
- 64KB Level 2 Cache (read-only)
- 4X and 16X Anti-aliasing Support
- ETC 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 with performance scalable up to 1080p 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 2,400Mb/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

- 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

- 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
 - 4,096-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 10/100/1,000Mb/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

- Full duplex flow control with recognition of incoming pause frames and hardware generation of transmitted pause frames
- 802.1Q VLAN tagging with recognition of incoming VLAN and priority tagged frames
- Supports IEEE Std 1588 v2

SD/SDIO 3.0 Controller

In addition to secure digital (SD) devices, this controller also supports eMMC 4.51.

- Host mode support only
- Built-in DMA
- 1/4-Bit SD Specification, version 3.0
- 1/4/8-Bit eMMC Specification, version 4.51
- Supports primary boot from SD Card and eMMC (Managed NAND)
- High speed, default speed, and low-speed support
- 1 and 4-bit data interface support
 - Low-speed clock 0–400KHz
 - Default speed 0–25MHz
 - High speed clock 0–50MHz
- High-speed Interface
 - SD UHS-1: 208MHz
 - eMMC HS200: 200MHz
- Memory, I/O, and SD cards
- Power control modes
- Data FIFO interface up to 512B

UART

- Programmable baud rate generator
- 6, 7, or 8 data bits
- 1, 1.5, or 2 stop bits
- Odd, even, space, mark, or no parity
- Parity, framing, and overrun error detection
- Line break generation and detection
- Automatic echo, local loopback, and remote loopback channel modes
- Modem control signals: CTS, RTS, DSR, DTR, RI, and DCD (from EMIO only)

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

- 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

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.

Table 6: MIO Peripheral Interface Mapping

Peripheral Interface	MIO	EMIO
Quad-SPI NAND	Yes	No
USB2.0: 0,1	Yes: External PHY	No
SDIO 0,1	Yes	Yes
SPI: 0,1 I2C: 0,1 CAN: 0,1 GPIO	Yes CAN: External PHY GPIO: Up to 78 bits	Yes CAN: External PHY GPIO: Up to 96 bits
GigE: 0,1,2,3	RGMII v2.0: External PHY	Supports GMII, RGMII v2.0 (HSTL), RGMII v1.3, MII, SGMII, and 1000BASE-X in Programmable Logic
UART: 0,1	Simple UART: Only two pins (TX and RX)	Full UART (TX, RX, DTR, DCD, DSR, RI, RTS, and CTS) requires either: <ul style="list-style-type: none"> • Two Processing System (PS) pins (RX and TX) through MIO and six additional Programmable Logic (PL) pins, <i>or</i> • Eight Programmable Logic (PL) pins
Debug Trace Ports	Yes: Up to 16 trace bits	Yes: Up to 32 trace bits
Processor JTAG	Yes	Yes

Transceiver (PS-GTR)

The four PS-GTR transceivers, which reside in the full power domain (FPD), support data rates of up to 6.0Gb/s. All the protocols cannot be pinned out at the same time. At any given time, four differential pairs can be pinned out using the transceivers. This is user programmable via the high-speed I/O multiplexer (HS-MIO).


- A Quad transceiver PS-GTR (TX/RX pair) able to support following standards simultaneously
 - x1, x2, or x4 lane of PCIe at Gen1 (2.5Gb/s) or Gen2 (5.0Gb/s) rates
 - 1 or 2 lanes of DisplayPort (TX only) at 1.62Gb/s, 2.7Gb/s, or 5.4Gb/s
 - 1 or 2 SATA channels at 1.5Gb/s, 3.0Gb/s, or 6.0Gb/s
 - 1 or 2 USB3.0 channels at 5.0Gb/s
 - 1-4 Ethernet SGMII channels at 1.25Gb/s
- Provides flexible host-programmable multiplexing function for connecting the transceiver resources to the PS masters (DisplayPort, PCIe, Serial-ATA, USB3.0, and GigE).

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 100Gb/s and 400Gb/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.

Two types of transceivers are used in the XA Zynq UltraScale+ MPSoC: GTH and PS-GTR. Both transceivers are arranged in groups of four, known as a transceiver Quad. Each serial transceiver is a combined transmitter and receiver. [Table 8](#) compares the available transceivers.

Table 8: Transceiver Information

	XA Zynq UltraScale+ MPSoC	
Type	PS-GTR	GTH
Qty	4	0–4
Max. Data Rate	6.0Gb/s	12.5Gb/s
Min. Data Rate	1.25Gb/s	0.5Gb/s
Applications	<ul style="list-style-type: none"> • PCIe Gen2 • USB • Ethernet 	<ul style="list-style-type: none"> • Backplane • PCIe Gen3 • HMC

The following information in this section pertains to the GTH 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. 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. 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.

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.

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 x 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-Level Features

Several functions span both the PS and PL and include:

- Reset Management
- Clock Management
- Power Domains
- PS Boot and Device Configuration
- Hardware and Software Debug Support

Reset Management

The reset management function provides the ability to reset the entire device or individual units within it. The PS supports these reset functions and signals:

- External and internal power-on reset signal
- Warm reset
- Watchdog timer reset
- User resets to PL
- Software, watchdog timer, or JTAG provided resets
- Security violation reset (locked down reset)

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

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 4,096-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.

Ordering Information

Table 10 shows the speed and temperature grades available in the different device families.

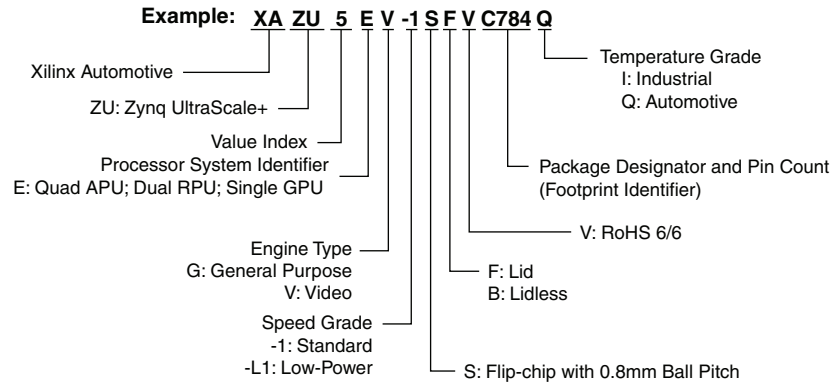
Table 10: Speed Grades, Temperature Grades, and Operating Voltages

Device Family	Devices	Speed Grade and Temperature Grade	
		Industrial (I)	Automotive (Q)
		-40°C to +100°C	-40°C to +125°C
XA Zynq UltraScale+	EG and EV Devices	-1I (0.85V)	-1Q (0.85V)
		-1L (0.72V)	

Notes:

- 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 XA Zynq UltraScale+ MPSoC.



-L1 is the ordering code for the -1L speed grade.

DS894_03_022317

Figure 3: XA Zynq UltraScale+ MPSoC Ordering Information