



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

### Understanding [Embedded - FPGAs \(Field Programmable Gate Array\)](#)

Embedded - FPGAs, or Field Programmable Gate Arrays, are advanced integrated circuits that offer unparalleled flexibility and performance for digital systems. Unlike traditional fixed-function logic devices, FPGAs can be programmed and reprogrammed to execute a wide array of logical operations, enabling customized functionality tailored to specific applications. This reprogrammability allows developers to iterate designs quickly and implement complex functions without the need for custom hardware.

### Applications of Embedded - FPGAs

The versatility of Embedded - FPGAs makes them indispensable in numerous fields. In telecommunications.

#### Details

Product Status	Active
Number of LABs/CLBs	216000
Number of Logic Elements/Cells	3780000
Total RAM Bits	514867200
Number of I/O	448
Number of Gates	-
Voltage - Supply	0.825V ~ 0.876V
Mounting Type	Surface Mount
Operating Temperature	-40°C ~ 100°C (TJ)
Package / Case	2577-BBGA, FCBGA
Supplier Device Package	2577-FCBGA (52.5x52.5)
Purchase URL	<a href="https://www.e-xfl.com/product-detail/xilinx/xcvu13p-2flga2577i">https://www.e-xfl.com/product-detail/xilinx/xcvu13p-2flga2577i</a>

# Kintex UltraScale+ FPGA Feature Summary

Table 5: Kintex UltraScale+ FPGA Feature Summary

	KU3P	KU5P	KU9P	KU11P	KU13P	KU15P
System Logic Cells	355,950	474,600	599,550	653,100	746,550	1,143,450
CLB Flip-Flops	325,440	433,920	548,160	597,120	682,560	1,045,440
CLB LUTs	162,720	216,960	274,080	298,560	341,280	522,720
Max. Distributed RAM (Mb)	4.7	6.1	8.8	9.1	11.3	9.8
Block RAM Blocks	360	480	912	600	744	984
Block RAM (Mb)	12.7	16.9	32.1	21.1	26.2	34.6
UltraRAM Blocks	48	64	0	80	112	128
UltraRAM (Mb)	13.5	18.0	0	22.5	31.5	36.0
CMTs (1 MMCM and 2 PLLs)	4	4	4	8	4	11
Max. HP I/O <sup>(1)</sup>	208	208	208	416	208	572
Max. HD I/O <sup>(2)</sup>	96	96	96	96	96	96
DSP Slices	1,368	1,824	2,520	2,928	3,528	1,968
System Monitor	1	1	1	1	1	1
GTH Transceiver 16.3Gb/s	0	0	28	32	28	44
GTY Transceivers 32.75Gb/s <sup>(3)</sup>	16	16	0	20	0	32
Transceiver Fractional PLLs	8	8	14	26	14	38
PCIe Gen3 x16 and Gen4 x8	1	1	0	4	0	5
150G Interlaken	0	0	0	1	0	4
100G Ethernet w/RS-FEC	0	1	0	2	0	4

## 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. GTY transceiver line rates are package limited: SFVB784 to 12.5Gb/s; FFVA676, FFVD900, and FFVA1156 to 16.3Gb/s. See [Table 6](#).

# Zynq UltraScale+: CG Device Feature Summary

Table 11: Zynq UltraScale+: 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 <sup>(1)</sup>	156	156	156	156	208	416	208
Max. HD I/O <sup>(2)</sup>	96	96	96	96	120	48	120
System Monitor	2	2	2	2	2	2	2
GTH Transceiver 16.3Gb/s <sup>(3)</sup>	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 12](#).

# Zynq UltraScale+: EG Device Feature Summary

Table 13: Zynq UltraScale+: EG Device Feature Summary

	ZU2EG	ZU3EG	ZU4EG	ZU5EG	ZU6EG	ZU7EG	ZU9EG	ZU11EG	ZU15EG	ZU17EG	ZU19EG
Application Processing Unit	Quad-core ARM Cortex-A53 MPCore with CoreSight; NEON & Single/Double Precision Floating Point; 32KB/32KB L1 Cache, 1MB L2 Cache										
Real-Time Processing Unit	Dual-core ARM Cortex-R5 with CoreSight; Single/Double Precision Floating Point; 32KB/32KB L1 Cache, and TCM										
Embedded and External Memory	256KB On-Chip Memory w/ECC; External DDR4; DDR3; DDR3L; LPDDR4; LPDDR3; External Quad-SPI; NAND; eMMC										
General Connectivity	214 PS I/O; UART; CAN; USB 2.0; I2C; SPI; 32b GPIO; Real Time Clock; WatchDog Timers; Triple Timer Counters										
High-Speed Connectivity	4 PS-GTR; PCIe Gen1/2; Serial ATA 3.1; DisplayPort 1.2a; USB 3.0; SGMII										
Graphic Processing Unit	ARM Mali-400 MP2; 64KB L2 Cache										
System Logic Cells	103,320	154,350	192,150	256,200	469,446	504,000	599,550	653,100	746,550	926,194	1,143,450
CLB Flip-Flops	94,464	141,120	175,680	234,240	429,208	460,800	548,160	597,120	682,560	846,806	1,045,440
CLB LUTs	47,232	70,560	87,840	117,120	214,604	230,400	274,080	298,560	341,280	423,403	522,720
Distributed RAM (Mb)	1.2	1.8	2.6	3.5	6.9	6.2	8.8	9.1	11.3	8.0	9.8
Block RAM Blocks	150	216	128	144	714	312	912	600	744	796	984
Block RAM (Mb)	5.3	7.6	4.5	5.1	25.1	11.0	32.1	21.1	26.2	28.0	34.6
UltraRAM Blocks	0	0	48	64	0	96	0	80	112	102	128
UltraRAM (Mb)	0	0	14.0	18.0	0	27.0	0	22.5	31.5	28.7	36.0
DSP Slices	240	360	728	1,248	1,973	1,728	2,520	2,928	3,528	1,590	1,968
CMTs	3	3	4	4	4	8	4	8	4	11	11
Max. HP I/O <sup>(1)</sup>	156	156	156	156	208	416	208	416	208	572	572
Max. HD I/O <sup>(2)</sup>	96	96	96	96	120	48	120	96	120	96	96
System Monitor	2	2	2	2	2	2	2	2	2	2	2
GTH Transceiver 16.3Gb/s <sup>(3)</sup>	0	0	16	16	24	24	24	32	24	44	44
GTY Transceivers 32.75Gb/s	0	0	0	0	0	0	0	16	0	28	28
Transceiver Fractional PLLs	0	0	8	8	12	12	12	24	12	36	36
PCIe Gen3 x16 and Gen4 x8	0	0	2	2	0	2	0	4	0	4	5
150G Interlaken	0	0	0	0	0	0	0	1	0	2	4
100G Ethernet w/ RS-FEC	0	0	0	0	0	0	0	2	0	2	4

## 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 14](#).

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 an FPGA divided into regions.




Figure 2: Column-Based FPGA Divided into Clock Regions

## Processing System (PS)

Zynq UltraScale+ MPSoCs consist of a PS coupled with programmable logic. The contents of the PS varies between the different Zynq UltraScale+ devices. All devices contain an APU, an RPU, and many peripherals for connecting the multiple processing engines to external components. The EG and EV devices contain a GPU and the EV devices contain a video codec unit (VCU). The components of the PS are connected together and to the PL through a multi-layered ARM AMBA AXI non-blocking interconnect that supports multiple simultaneous master-slave transactions. Traffic through the interconnect can be regulated by the quality of service (QoS) block in the interconnect. Twelve dedicated AXI 32-bit, 64-bit, or 128-bit ports connect the PL to high-speed interconnect and DDR in the PS via a FIFO interface.

There are four independently controllable power domains: the PL plus three within the PS (full power, lower power, and battery power domains). Additionally, many peripherals support clock gating and power gating to further reduce dynamic and static power consumption.

## Application Processing Unit (APU)

The APU has a feature-rich dual-core or quad-core ARM Cortex-A53 processor. Cortex-A53 cores are 32-bit/64-bit application processors based on ARM-v8A architecture, offering the best performance-to-power ratio. The ARMv8 architecture supports hardware virtualization. Each of the Cortex-A53 cores has: 32KB of instruction and data L1 caches, with parity and ECC protection respectively; a NEON SIMD engine; and a single and double precision floating point unit. In addition to these blocks, the APU consists of a snoop control unit and a 1MB L2 cache with ECC protection to enhance system-level performance. The snoop control unit keeps the L1 caches coherent thus eliminating the need of spending software bandwidth for coherency. The APU also has a built-in interrupt controller supporting virtual interrupts. The APU communicates to the rest of the PS through 128-bit AXI coherent extension (ACE) port via Cache Coherent Interconnect (CCI) block, using the System Memory Management Unit (SMMU). The APU is also connected to the Programmable Logic (PL), through the 128-bit accelerator coherency port

(ACP), providing a low latency coherent port for accelerators in the PL. To support real-time debug and trace, each core also has an Embedded Trace Macrocell (ETM) that communicates with the ARM CoreSight™ Debug System.

## Real-Time Processing Unit (RPU)

The RPU in the PS contains a dual-core ARM Cortex-R5 PS. Cortex-R5 cores are 32-bit real-time processor cores based on ARM-v7R architecture. Each of the Cortex-R5 cores has 32KB of level-1 (L1) instruction and data cache with ECC protection. In addition to the L1 caches, each of the Cortex-R5 cores also has a 128KB tightly coupled memory (TCM) interface for real-time single cycle access. The RPU also has a dedicated interrupt controller. The RPU can operate in either split or lock-step mode. In split mode, both processors run independently of each other. In lock-step mode, they run in parallel with each other, with integrated comparator logic, and the TCMs are used as 256KB unified memory. The RPU communicates with the rest of the PS via the 128-bit AXI-4 ports connected to the low power domain switch. It also communicates directly with the PL through 128-bit low latency AXI-4 ports. To support real-time debug and trace each core also has an embedded trace macrocell (ETM) that communicates with the ARM CoreSight Debug System.

## External Memory

The PS can interface to many types of external memories through dedicated memory controllers. The dynamic memory controller supports DDR3, DDR3L, DDR4, LPDDR3, and LPDDR4 memories. The multi-protocol DDR memory controller can be configured to access a 2GB address space in 32-bit addressing mode and up to 32GB in 64-bit addressing mode using a single or dual rank configuration of 8-bit, 16-bit, or 32-bit DRAM memories. Both 32-bit and 64-bit bus access modes are protected by ECC using extra bits.

The SD/eMMC controller supports 1 and 4 bit data interfaces at low, default, high-speed, and ultra-high-speed (UHS) clock rates. This controller also supports 1-, 4-, or 8-bit-wide eMMC interfaces that are compliant to the eMMC 4.51 specification. eMMC is one of the primary boot and configuration modes for Zynq UltraScale+ MPSoCs and supports boot from managed NAND devices. The controller has a built-in DMA for enhanced performance.

The Quad-SPI controller is one of the primary boot and configuration devices. It supports 4-byte and 3-byte addressing modes. In both addressing modes, single, dual-stacked, and dual-parallel configurations are supported. Single mode supports a quad serial NOR flash memory, while in double stacked and double parallel modes, it supports two quad serial NOR flash memories.

The NAND controller is based on ONFI3.1 specification. It has an 8-pin interface and provides 200Mb/s of bandwidth in synchronous mode. It supports 24 bits of ECC thus enabling support for SLC NAND memories. It has two chip-selects to support deeper memory and a built-in DMA for enhanced performance.

## General Connectivity

There are many peripherals in the PS for connecting to external devices over industry standard protocols, including CAN2.0B, USB, Ethernet, I2C, and UART. Many of the peripherals support clock gating and power gating modes to reduce dynamic and static power consumption.

### **USB 3.0/2.0**

The pair of USB controllers can be configured as host, device, or On-The-Go (OTG). The core is compliant to USB 3.0 specification and supports super, high, full, and low speed modes in all configurations. In host mode, the USB controller is compliant with the Intel XHCI specification. In device mode, it supports up to 12 end points. While operating in USB 3.0 mode, the controller uses the serial transceiver and operates up to 5.0Gb/s. In USB 2.0 mode, the Universal Low Peripheral Interface (ULPI) is used to connect the controller to an external PHY operating up to 480Mb/s. The ULPI is also connected in USB 3.0 mode to support high-speed operations.

### **Ethernet MAC**

The four tri-speed ethernet MACs support 10Mb/s, 100Mb/s, and 1Gb/s operations. The MACs support jumbo frames and time stamping through the interfaces based on IEEE Std 1588v2. The ethernet MACs can be connected through the serial transceivers (SGMII), the MIO (RGMII), or through EMIO (GMII). The GMII interface can be converted to a different interface within the PL.

## High-Speed Connectivity

The PS includes four PS-GTR transceivers (transmit and receive), supporting data rates up to 6.0Gb/s and can interface to the peripherals for communication over PCIe, SATA, USB 3.0, SGMII, and DisplayPort.

### **PCIe**

The integrated block for PCIe is compliant with PCI Express base specification 2.1 and supports x1, x2, and x4 configurations as root complex or end point, compliant to transaction ordering rules in both configurations. It has built-in DMA, supports one virtual channel and provides fully configurable base address registers.

### **SATA**

Users can connect up to two external devices using the two SATA host port interfaces compliant to the SATA 3.1 specification. The SATA interfaces can operate at 1.5Gb/s, 3.0Gb/s, or 6.0Gb/s data rates and are compliant with advanced host controller interface (AHCI) version 1.3 supporting partial and slumber power modes.

### **DisplayPort**

The DisplayPort controller supports up to two lanes of source-only DisplayPort compliant with VESA DisplayPort v1.2a specification (source only) at 1.62Gb/s, 2.7Gb/s, and 5.4Gb/s data rates. The controller supports single stream transport (SST); video resolution up to 4Kx2K at a 30Hz frame rate; video formats Y-only, YCbCr444, YCbCr422, YCbCr420, RGB, YUV444, YUV422, xvYCC, and pixel color depth of 6, 8, 10, and 12 bits per color component.

## Graphics Processing Unit (GPU)

The dedicated ARM Mali-400 MP2 GPU in the PS supports 2D and 3D graphics acceleration up to 1080p resolution. The Mali-400 supports OpenGL ES 1.1 and 2.0 for 3D graphics and Open VG 1.1 standards for 2D vector graphics. It has a geometry processor (GP) and 2 pixel processors to perform tile rendering operations in parallel. It has dedicated Memory management units for GP and pixel processors, which supports 4 KB page size. The GPU also has 64KB level-2 (L2) read-only cache. It supports 4X and 16X Full scene Anti-Aliasing (FSAA). It is fully autonomous, enabling maximum parallelization between APU and GPU. It has built-in hardware texture decompression, allowing the texture to remain compressed (in ETC format) in graphics hardware and decompress the required samples on the fly. It also supports efficient alpha blending of multiple layers in hardware without additional bandwidth consumption. It has a pixel fill rate of 2Mpixel/sec/MHz and a triangle rate of 0.1Mvertex/sec/MHz. The GPU supports extensive texture format for RGBA 8888, 565, and 1556 in Mono 8, 16, and YUV formats. For power sensitive applications, the GPU supports clock and power gating for each GP, pixel processors, and L2 cache. During power gating, GPU does not consume any static or dynamic power; during clock gating, it only consumes static power.

## Video Codec Unit (VCU)

The video codec unit (VCU) provides multi-standard video encoding and decoding capabilities, including: High Efficiency Video Coding (HEVC), i.e., H.265; and Advanced Video Coding (AVC), i.e., H.264 standards. The VCU is capable of simultaneous encode and decode at rates up to 4Kx2K at 60 frames per second (fps) (approx. 600Mpixel/sec) or 8Kx4K at a reduced frame rate (~15fps).

---

## Input/Output

All UltraScale devices, whether FPGA or MPSoC, 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 UltraScale FPGAs and in the programmable logic of 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-range (HR), high-performance (HP), or high-density (HD). The HR I/Os offer the widest range of voltage support, from 1.2V to 3.3V. 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 or HR pins per bank or 24 HD pins per bank. Each bank has one common  $V_{CCO}$  output buffer power supply, which also powers certain input buffers. In addition, HR banks can be split into two half-banks, each with their own  $V_{CCO}$  supply. 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.



# High-Speed Serial Transceivers

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.

Three types of transceivers are used in the UltraScale architecture: GTH and GTY in FPGAs and MPSoC PL, and PS-GTR in the MPSoC PS. All transceivers are arranged in groups of four, known as a transceiver Quad. Each serial transceiver is a combined transmitter and receiver. [Table 17](#) compares the available transceivers.

**Table 17: Transceiver Information**

	Kintex UltraScale		Kintex UltraScale+		Virtex UltraScale		Virtex UltraScale+	Zynq UltraScale+		
Type	GTH	GTY	GTH	GTY	GTH	GTY	GTY	PS-GTR	GTH	GTY
Qty	16–64	0–32	20–60	0–60	20–60	0–60	40–128	4	0–44	0–28
Max. Data Rate	16.3Gb/s	16.3Gb/s	16.3Gb/s	32.75Gb/s	16.3Gb/s	30.5Gb/s	32.75Gb/s	6.0Gb/s	16.3Gb/s	32.75Gb/s
Min. Data Rate	0.5Gb/s	0.5Gb/s	0.5Gb/s	0.5Gb/s	0.5Gb/s	0.5Gb/s	0.5Gb/s	1.25Gb/s	0.5Gb/s	0.5Gb/s
Key Apps	<ul style="list-style-type: none"> <li>Backplane</li> <li>PCIe Gen4</li> <li>HMC</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>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>	<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>	<ul style="list-style-type: none"> <li>100G+ Optics</li> <li>Chip-to-Chip</li> <li>25G+ Backplane</li> <li>HMC</li> </ul>	<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 10G+ and 25G+ backplanes.

## 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 UltraScale architecture includes integrated blocks for PCIe technology that can be configured as an Endpoint or Root Port. UltraScale devices are compliant to the PCI Express Base Specification Revision 3.0. UltraScale+ devices are 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 FPGA or MPSoC.

This block is highly configurable to system design requirements and can operate up to the maximum lane widths and data rates listed in [Table 18](#).

*Table 18: PCIe Maximum Configurations*

	Kintex UltraScale	Kintex UltraScale+	Virtex UltraScale	Virtex UltraScale+	Zynq UltraScale+
Gen1 (2.5Gb/s)	x8	x16	x8	x16	x16
Gen2 (5Gb/s)	x8	x16	x8	x16	x16
Gen3 (8Gb/s)	x8	x16	x8	x16	x16
Gen4 (16Gb/s) <sup>(1)</sup>		x8		x8	x8

**Notes:**

1. Transceivers in Kintex UltraScale and Virtex UltraScale devices are capable of operating at Gen4 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, FPGA or MPSoC logic interface speeds, reference clock frequency, and base address register decoding and filtering.

---

## Cache Coherent Interconnect for Accelerators (CCIX)

CCIX is a chip-to-chip interconnect operating at data rates up to 25Gb/s that allows two or more devices to share memory in a cache coherent manner. Using PCIe for the transport layer, CCIX can operate at several standard data rates (2.5, 5, 8, and 16Gb/s) with an additional high-speed 25Gb/s option. The specification employs a subset of full coherency protocols and ensures that FPGAs used as accelerators can coherently share data with processors using different instruction set architectures.

Virtex UltraScale+ HBM devices support CCIX data rates up to 16Gb/s and contain four CCIX ports and at least four integrated blocks for PCIe. Each CCIX port requires the use of one integrated block for PCIe. If not used with a CCIX port, the integrated blocks for PCIe can still be used for PCIe communication.

---

## 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 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 operated without using the MAC.

## Stacked Silicon Interconnect (SSI) Technology

Many challenges associated with creating high-capacity devices are addressed by Xilinx with the second generation of the pioneering 3D SSI technology. SSI technology enables multiple super-logic regions (SLRs) to be combined on a passive interposer layer, using proven manufacturing and assembly techniques from industry leaders, to create a single device with more than 20,000 low-power inter-SLR connections. Dedicated interface tiles within the SLRs provide ultra-high bandwidth, low latency connectivity to other SLRs. Table 19 shows the number of SLRs in devices that use SSI technology and their dimensions.

Table 19: UltraScale and UltraScale+ 3D IC SLR Count and Dimensions

	Kintex UltraScale		Virtex UltraScale				Virtex UltraScale+								
Device	KU085	KU115	VU125	VU160	VU190	VU440	VU5P	VU7P	VU9P	VU11P	VU13P	VU31P	VU33P	VU35P	VU37P
# SLRs	2	2	2	3	3	3	2	2	3	3	4	1	1	2	3
SLR Width (in regions)	6	6	6	6	6	9	6	6	6	8	8	8	8	8	8
SLR Height (in regions)	5	5	5	5	5	5	5	5	5	4	4	4	4	4	4

## Clock Management

The clock generation and distribution components in UltraScale devices are located adjacent to the columns that contain the memory interface 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).

There are three sets of programmable frequency dividers (D, M, and O) that 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.

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 UltraScale devices 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 UltraScale device 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, UltraScale FPGAs and MPSoCs 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 devices support the highest bandwidth HMC configuration of 64 lanes with a single FPGA.

## Block RAM

Every UltraScale architecture-based device contains a number of 36 Kb 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  $512 \times 72$  (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  $512 \times 36$ . 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 18bits (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 available in UltraScale+ devices. 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. 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. If additional density is required, all the UltraRAM columns in an SLR can be connected together with a few fabric resources to create single instances of RAM approximately 100Mb in size. This makes UltraRAM an ideal solution for replacing external memories such as SRAM. Cascadable anywhere from 288Kb to 100Mb, 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.

---

## High Bandwidth Memory (HBM)

Virtex UltraScale+ HBM devices incorporate 4GB HBM stacks adjacent to the FPGA die. Using stacked silicon interconnect technology, the FPGA communicates to the HBM stacks through memory controllers that connect to dedicated low-inductance interconnect in the silicon interposer. Each Virtex UltraScale+ HBM FPGA contains one or two HBM stacks, resulting in up to 8GB of HBM per FPGA.

The FPGA has 32 HBM AXI interfaces used to communicate with the HBM. Through a built-in switch mechanism, any of the 32 HBM AXI interfaces can access any memory address on either one or both of the HBM stacks due to the flexible addressing feature. This flexible connection between the FPGA and the HBM stacks results in easy floorplanning and timing closure. The memory controllers perform read and write reordering to improve bus efficiency. Data integrity is ensured through error checking and correction (ECC) circuitry.

---

## 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.

---

## Digital Signal Processing

DSP applications use many binary multipliers and accumulators, best implemented in dedicated DSP slices. All UltraScale 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 and external channels to the ADC.

All UltraScale architecture-based devices contain at least one System Monitor. The System Monitor in UltraScale+ FPGAs and the PL of Zynq UltraScale+ MPSoCs is similar to the Kintex UltraScale and Virtex UltraScale devices but with additional features including a PMBus interface.

Zynq UltraScale+ MPSoCs contain an additional System Monitor block in the PS. See [Table 20](#).

**Table 20: Key System Monitor Features**

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

In FPGAs and the MPSoC PL, sensor outputs and up to 17 user-allocated external analog inputs are digitized using a 10-bit 200 kilo-sample-per-second (kSPS) ADC, and the measurements are stored in registers that can be accessed via internal FPGA (DRP), JTAG, PMBus, or I2C interfaces. The I2C interface and PMBus allow the on-chip monitoring to be easily accessed by the System Manager/Host before and after device configuration.

The System Monitor in the MPSoC PS uses a 10-bit, 1 mega-sample-per-second (MSPS) ADC to digitize the sensor outputs. The measurements are stored in registers and are accessed via the Advanced Peripheral Bus (APB) interface by the processors and the platform management unit (PMU) in the PS.

## Configuration

The UltraScale architecture-based devices store their customized configuration in SRAM-type internal latches. The configuration storage is volatile and must be reloaded whenever the device is powered up. This storage can also be reloaded at any time. Several methods and data formats for loading configuration are available, determined by the mode pins, with more dedicated configuration datapath pins to simplify the configuration process.

UltraScale architecture-based devices support secure and non-secure boot with optional Advanced Encryption Standard - Galois/Counter Mode (AES-GCM) decryption and authentication logic. If only authentication is required, the UltraScale architecture provides an alternative form of authentication in the form of RSA algorithms. For RSA authentication support in the Kintex UltraScale and Virtex UltraScale families, go to [UG570](#), *UltraScale Architecture Configuration User Guide*.

UltraScale architecture-based devices also have the ability to select between multiple configurations, and support robust field-update methodologies. This is especially useful for updates to a design after the end product has been shipped. Designers can release their product with an early version of the design, thus getting their product to market faster. This feature allows designers to keep their customers current with the most up-to-date design while the product is already deployed in the field.

## Booting MPSoCs

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 decryption/authentication, and 4096-bit RSA blocks decrypt and authenticate 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. One of the CPUs, Cortex-A53 or Cortex-R5, executes code out of on-chip ROM and copies the first stage boot loader (FSBL) from the boot device to the on-chip memory (OCM).

After copying the FSBL to OCM, the processor 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 executable after boot.

## Configuring FPGAs

The SPI (serial NOR) interface (x1, x2, x4, and dual x4 modes) and the BPI (parallel NOR) interface (x8 and x16 modes) are two common methods used for configuring the FPGA. Users can directly connect an SPI or BPI flash to the FPGA, and the FPGA's internal configuration logic reads the bitstream out of the flash and configures itself, eliminating the need for an external controller. The FPGA automatically detects the bus width on the fly, eliminating the need for any external controls or switches. Bus widths supported are x1, x2, x4, and dual x4 for SPI, and x8 and x16 for BPI. The larger bus widths increase configuration speed and reduce the amount of time it takes for the FPGA to start up after power-on.

In master mode, the FPGA can drive the configuration clock from an internally generated clock, or for higher speed configuration, the FPGA can use an external configuration clock source. This allows high-speed configuration with the ease of use characteristic of master mode. Slave modes up to 32 bits wide that are especially useful for processor-driven configuration are also supported by the FPGA. In addition, the new media configuration access port (MCAP) provides a direct connection between the integrated block for PCIe and the configuration logic to simplify configuration over PCIe.

SEU detection and mitigation (SEM) IP, RSA authentication, post-configuration CRC, and Security Monitor (SecMon) IP are not supported in the KU025 FPGA.

---

## Packaging

The UltraScale devices are available in a variety of organic flip-chip and lidless flip-chip packages supporting different quantities of I/Os and transceivers. Maximum supported performance can depend on the style of package and its material. Always refer to the specific device data sheet for performance specifications by package type.

In flip-chip packages, the silicon device is attached to the package substrate using a high-performance flip-chip process. Decoupling capacitors are mounted on the package substrate to optimize signal integrity under simultaneous switching of outputs (SSO) conditions.

Table 21: Speed Grade and Temperature Grade (Cont'd)

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 <sup>(2)(3)</sup> (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
					-1LI <sup>(3)</sup> (0.85V or 0.72V)
	ZU2EG ZU3EG		-2E (0.85V)		-2I (0.85V)
				-2LE <sup>(2)(3)</sup> (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
					-1LI <sup>(3)</sup> (0.85V or 0.72V)
	ZU4EG ZU5EG ZU6EG ZU7EG ZU9EG ZU11EG ZU15EG ZU17EG ZU19EG		-3E (0.90V)		
			-2E (0.85V)		-2I (0.85V)
				-2LE <sup>(2)(3)</sup> (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
					-1LI <sup>(3)</sup> (0.85V or 0.72V)
	EV Devices		-3E (0.90V)		
			-2E (0.85V)		-2I (0.85V)
				-2LE <sup>(2)(3)</sup> (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
					-1LI <sup>(3)</sup> (0.85V or 0.72V)

**Notes:**

- KU025 and KU095 are not available in -3E or -1LI speed/temperature grades.
- 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.
- In Zynq UltraScale+ MPSoCs, when operating the PL at low voltage (0.72V), the PS operates at nominal voltage (0.85V).

---

## 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.