



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	34260
Number of Logic Elements/Cells	599550
Total RAM Bits	41881600
Number of I/O	304
Number of Gates	-
Voltage - Supply	0.825V ~ 0.876V
Mounting Type	Surface Mount
Operating Temperature	0°C ~ 100°C (TJ)
Package / Case	900-BBGA, FCBGA
Supplier Device Package	900-FCBGA (31x31)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xcku9p-2ffve900e

Migrating Devices

UltraScale and UltraScale+ families provide footprint compatibility to enable users to migrate designs from one device or family to another. Any two packages with the same footprint identifier code are footprint compatible. For example, Kintex UltraScale devices in the A1156 packages are footprint compatible with Kintex UltraScale+ devices in the A1156 packages. Likewise, Virtex UltraScale devices in the B2104 packages are compatible with Virtex UltraScale+ devices and Kintex UltraScale devices in the B2104 packages. All valid device/package combinations are provided in the Device-Package Combinations and Maximum I/Os tables in this document. Refer to [UG583](#), *UltraScale Architecture PCB Design User Guide* for more detail on migrating between UltraScale and UltraScale+ devices and packages.

Kintex UltraScale Device-Package Combinations and Maximum I/Os

Table 4: Kintex UltraScale Device-Package Combinations and Maximum I/Os

Package (1)(2)(3)	Package Dimensions (mm)	KU025	KU035	KU040	KU060	KU085	KU095	KU115
		HR, HP GTH	HR, HP GTH	HR, HP GTH	HR, HP GTH	HR, HP GTH	HR, HP GTH, GTY ⁽⁴⁾	HR, HP GTH
SFVA784 ⁽⁵⁾	23x23		104, 364 8	104, 364 8				
FBVA676 ⁽⁵⁾	27x27		104, 208 16	104, 208 16				
FBVA900 ⁽⁵⁾	31x31		104, 364 16	104, 364 16				
FFVA1156	35x35	104, 208 12	104, 416 16	104, 416 20	104, 416 28		52, 468 20, 8	
FFVA1517	40x40				104, 520 32			
FLVA1517	40x40					104, 520 48		104, 520 48
FFVC1517	40x40						52, 468 20, 20	
FLVD1517	40x40							104, 234 64
FFVB1760	42.5x42.5						52, 650 32, 16	
FLVB1760	42.5x42.5					104, 572 44		104, 598 52
FLVD1924	45x45							156, 676 52
FLVF1924	45x45					104, 520 56		104, 624 64
FLVA2104	47.5x47.5							156, 676 52
FFVB2104	47.5x47.5						52, 650 32, 32	
FLVB2104	47.5x47.5							104, 598 64

Notes:

- Go to [Ordering Information](#) for package designation details.
- FB/FF/FL packages have 1.0mm ball pitch. SF packages have 0.8mm ball pitch.
- Packages with the same last letter and number sequence, e.g., A2104, are footprint compatible with all other UltraScale architecture-based devices with the same sequence. The footprint compatible devices within this family are outlined. See the [UltraScale Architecture Product Selection Guide](#) for details on inter-family migration.
- GTY transceivers in Kintex UltraScale devices support data rates up to 16.3Gb/s.
- GTH transceivers in SF/FB packages support data rates up to 12.5Gb/s.

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 ⁽¹⁾	208	208	208	416	208	572
Max. HD I/O ⁽²⁾	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 ⁽³⁾	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](#).

Kintex UltraScale+ Device-Package Combinations and Maximum I/Os

Table 6: Kintex UltraScale+ Device-Package Combinations and Maximum I/Os

Package (1)(2)(4)	Package Dimensions (mm)	KU3P	KU5P	KU9P	KU11P	KU13P	KU15P
		HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY
SFVB784(3)	23x23	96, 208 0, 16	96, 208 0, 16				
FFVA676(3)	27x27	48, 208 0, 16	48, 208 0, 16				
FFVB676	27x27	72, 208 0, 16	72, 208 0, 16				
FFVD900(3)	31x31	96, 208 0, 16	96, 208 0, 16		96, 312 16, 0		
FFVE900	31x31			96, 208 28, 0		96, 208 28, 0	
FFVA1156(3)	35x35				48, 416 20, 8		48, 468 20, 8
FFVE1517	40x40				96, 416 32, 20		96, 416 32, 24
FFVA1760	42.5x42.5						96, 416 44, 32
FFVE1760	42.5x42.5						96, 572 32, 24

Notes:

1. Go to [Ordering Information](#) for package designation details.
2. FF packages have 1.0mm ball pitch. SF packages have 0.8mm ball pitch.
3. GTY transceiver line rates are package limited: SFVB784 to 12.5Gb/s; FFVA676, FFVD900, and FFVA1156 to 16.3Gb/s.
4. Packages with the same last letter and number sequence, e.g., A676, are footprint compatible with all other UltraScale architecture-based devices with the same sequence. The footprint compatible devices within this family are outlined. See the [UltraScale Architecture Product Selection Guide](#) for details on inter-family migration.

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

Table 14: Zynq UltraScale+: EG Device-Package Combinations and Maximum I/Os

Package (1)(2)(3)(4)(5)	Package Dimensions (mm)	ZU2EG	ZU3EG	ZU4EG	ZU5EG	ZU6EG	ZU7EG	ZU9EG	ZU11EG	ZU15EG	ZU17EG	ZU19EG
		HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY
SBVA484(6)	19x19	24, 58 0, 0	24, 58 0, 0									
SFVA625	21x21	24, 156 0, 0	24, 156 0, 0									
SFVC784(7)	23x23	96, 156 0, 0	96, 156 0, 0	96, 156 4, 0	96, 156 4, 0							
FBVB900	31x31			48, 156 16, 0	48, 156 16, 0		48, 156 16, 0					
FFVC900	31x31					48, 156 16, 0		48, 156 16, 0		48, 156 16, 0		
FFVB1156	35x35					120, 208 24, 0		120, 208 24, 0		120, 208 24, 0		
FFVC1156	35x35						48, 312 20, 0		48, 312 20, 0			
FFVB1517	40x40								72, 416 16, 0		72, 572 16, 0	72, 572 16, 0
FFVF1517	40x40						48, 416 24, 0		48, 416 32, 0			
FFVC1760	42.5x42.5								96, 416 32, 16		96, 416 32, 16	96, 416 32, 16
FFVD1760	42.5x42.5										48, 260 44, 28	48, 260 44, 28
FFVE1924	45x45										96, 572 44, 0	96, 572 44, 0

Notes:

- Go to [Ordering Information](#) for package designation details.
- FB/FF packages have 1.0mm ball pitch. SB/SF packages have 0.8mm ball pitch.
- All device package combinations bond out 4 PS-GTR transceivers.
- All device package combinations bond out 214 PS I/O except ZU2EG and ZU3EG in the SBVA484 and SFVA625 packages, which bond out 170 PS I/Os.
- Packages with the same last letter and number sequence, e.g., A484, are footprint compatible with all other UltraScale architecture-based devices with the same sequence. The footprint compatible devices within this family are outlined.
- All 58 HP I/O pins are powered by the same V_{CCO} supply.
- GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s.

Zynq UltraScale+: EG Device Feature Summary

Table 15: Zynq UltraScale+: EV Device Feature Summary

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

Notes:

1. HP = High-performance I/O with support for I/O voltage from 1.0V to 1.8V.
2. HD = High-density I/O with support for I/O voltage from 1.2V to 3.3V.
3. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s. See [Table 16](#).

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

Table 16: Zynq UltraScale+: EV Device-Package Combinations and Maximum I/Os


Package (1)(2)(3)(4)	Package Dimensions (mm)	ZU4EV	ZU5EV	ZU7EV
		HD, HP GTH, GTY	HD, HP GTH, GTY	HD, HP GTH, GTY
SFVC784(5)	23x23	96, 156 4, 0	96, 156 4, 0	
FBVB900	31x31	48, 156 16, 0	48, 156 16, 0	48, 156 16, 0
FFVC1156	35x35			48, 312 20, 0
FFVF1517	40x40			48, 416 24, 0

Notes:

1. Go to [Ordering Information](#) for package designation details.
2. FB/FF packages have 1.0mm ball pitch. SF packages have 0.8mm ball pitch.
3. All device package combinations bond out 4 PS-GTR transceivers.
4. GTH transceivers in the SFVC784 package support data rates up to 12.5Gb/s.
5. Packages with the same last letter and number sequence, e.g., B900, are footprint compatible with all other UltraScale architecture-based devices with the same sequence. The footprint compatible devices within this family are outlined.

Device Layout

UltraScale 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.



DS890_01_101712

Figure 1: FPGA 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 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.

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"> Backplane PCIe Gen4 HMC 	<ul style="list-style-type: none"> Backplane PCIe Gen4 HMC 	<ul style="list-style-type: none"> Backplane PCIe Gen4 HMC 	<ul style="list-style-type: none"> 100G+ Optics Chip-to-Chip 25G+ Backplane HMC 	<ul style="list-style-type: none"> Backplane PCIe Gen4 HMC 	<ul style="list-style-type: none"> 100G+ Optics Chip-to-Chip 25G+ Backplane HMC 	<ul style="list-style-type: none"> 100G+ Optics Chip-to-Chip 25G+ Backplane HMC 	<ul style="list-style-type: none"> PCIe Gen2 USB Ethernet 	<ul style="list-style-type: none"> Backplane PCIe Gen4 HMC 	<ul style="list-style-type: none"> 100G+ Optics Chip-to-Chip 25G+ Backplane HMC

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.

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.

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 27×18 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.

Ordering Information

Table 21 shows the speed and temperature grades available in the different device families. V_{CCINT} supply voltage is listed in parentheses.

Table 21: Speed Grade and Temperature Grade

Device Family	Devices	Speed Grade and Temperature Grade			
		Commercial (C)	Extended (E)		Industrial (I)
		0°C to +85°C	0°C to +100°C	0°C to +110°C	–40°C to +100°C
Kintex UltraScale	All		-3E ⁽¹⁾ (1.0V)		
			-2E (0.95V)		-2I (0.95V)
		-1C (0.95V)			-1I (0.95V)
					-1LI ⁽¹⁾ (0.95V or 0.90V)
Kintex UltraScale+	All		-3E (0.90V)		
			-2E (0.85V)		-2I (0.85V)
				-2LE ⁽²⁾ (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
					-1LI (0.85V or 0.72V)
Virtex UltraScale	VU065 VU080 VU095 VU125 VU160 VU190		-3E (1.0V)		
			-2E (0.95V)		-2I (0.95V)
			-1HE (0.95V or 1.0V)		-1I (0.95V)
	VU440		-3E (1.0V)		
			-2E (0.95V)		-2I (0.95V)
Virtex UltraScale+	VU3P VU5P VU7P VU9P VU11P VU13P		-3E (0.90V)		
			-2E (0.85V)		-2I (0.85V)
				-2LE ⁽²⁾ (0.85V or 0.72V)	
			-1E (0.85V)		-1I (0.85V)
	VU31P VU33P VU35P VU37P		-3E (0.90V)		
			-2E (0.85V)		
				-2LE ⁽²⁾ (0.85V or 0.72V)	
			-1E (0.85V)		

Date	Version	Description of Revisions
02/06/2014	1.1	Updated PCIe information in Table 1 and Table 3 . Added FFVJ1924 package to Table 8 .
12/10/2013	1.0	Initial Xilinx release.