



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	82920
Number of Logic Elements/Cells	1451100
Total RAM Bits	77721600
Number of I/O	624
Number of Gates	-
Voltage - Supply	0.922V ~ 0.979V
Mounting Type	Surface Mount
Operating Temperature	-40°C ~ 100°C (TJ)
Package / Case	1517-BBGA, FCBGA
Supplier Device Package	1517-FCBGA (40x40)
Purchase URL	https://www.e-xfl.com/product-detail/xilinx/xcku115-2flva1517i

I/O, Transceiver, PCIe, 100G Ethernet, and 150G Interlaken

Data is transported on and off chip through a combination of the high-performance parallel SelectIO™ interface and high-speed serial transceiver connectivity. I/O blocks provide support for cutting-edge memory interface and network protocols through flexible I/O standard and voltage support. The serial transceivers in the UltraScale architecture-based devices transfer data up to 32.75Gb/s, enabling 25G+ backplane designs with dramatically lower power per bit than previous generation transceivers. All transceivers, except the PS-GTR, support the required data rates for PCIe Gen3, and Gen4 (rev 0.5), and integrated blocks for PCIe enable UltraScale devices to support up to Gen4 x8 and Gen3 x16 Endpoint and Root Port designs. Integrated blocks for 150Gb/s Interlaken and 100Gb/s Ethernet (100G MAC/PCS) extend the capabilities of UltraScale devices, enabling simple, reliable support for Nx100G switch and bridge applications. Virtex UltraScale+ HBM devices include Cache Coherent Interconnect for Accelerators (CCIX) ports for coherently sharing data with different processors.

Clocks and Memory Interfaces

UltraScale devices contain powerful clock management circuitry, including clock synthesis, buffering, and routing components that together provide a highly capable framework to meet design requirements. The clock network allows for extremely flexible distribution of clocks to minimize the skew, power consumption, and delay associated with clock signals. The clock management technology is tightly integrated with dedicated memory interface circuitry to enable support for high-performance external memories, including DDR4. In addition to parallel memory interfaces, UltraScale devices support serial memories, such as hybrid memory cube (HMC).

Routing, SSI, Logic, Storage, and Signal Processing

Configurable Logic Blocks (CLBs) containing 6-input look-up tables (LUTs) and flip-flops, DSP slices with 27x18 multipliers, 36Kb block RAMs with built-in FIFO and ECC support, and 4Kx72 UltraRAM blocks (in UltraScale+ devices) are all connected with an abundance of high-performance, low-latency interconnect. In addition to logical functions, the CLB provides shift register, multiplexer, and carry logic functionality as well as the ability to configure the LUTs as distributed memory to complement the highly capable and configurable block RAMs. The DSP slice, with its 96-bit-wide XOR functionality, 27-bit pre-adder, and 30-bit A input, performs numerous independent functions including multiply accumulate, multiply add, and pattern detect. In addition to the device interconnect, in devices using SSI technology, signals can cross between super-logic regions (SLRs) using dedicated, low-latency interface tiles. These combined routing resources enable easy support for next-generation bus data widths. Virtex UltraScale+ HBM devices include up to 8GB of high bandwidth memory.

Configuration, Encryption, and System Monitoring

The configuration and encryption block performs numerous device-level functions critical to the successful operation of the FPGA or MPSoC. This high-performance configuration block enables device configuration from external media through various protocols, including PCIe, often with no requirement to use multi-function I/O pins during configuration. The configuration block also provides 256-bit AES-GCM decryption capability at the same performance as unencrypted configuration. Additional features include SEU detection and correction, partial reconfiguration support, and battery-backed RAM or eFUSE technology for AES key storage to provide additional security. The System Monitor enables the monitoring of the physical environment via on-chip temperature and supply sensors and can also monitor up to 17 external analog inputs. With UltraScale+ MPSoCs, the device is booted via the Configuration and Security Unit (CSU), which supports secure boot via the 256-bit AES-GCM and SHA/384 blocks. The cryptographic engines in the CSU can be used in the MPSoC after boot for user encryption.

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+ 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](#).

Virtex UltraScale FPGA Feature Summary

Table 7: Virtex UltraScale FPGA Feature Summary

	VU065	VU080	VU095	VU125	VU160	VU190	VU440
System Logic Cells	783,300	975,000	1,176,000	1,566,600	2,026,500	2,349,900	5,540,850
CLB Flip-Flops	716,160	891,424	1,075,200	1,432,320	1,852,800	2,148,480	5,065,920
CLB LUTs	358,080	445,712	537,600	716,160	926,400	1,074,240	2,532,960
Maximum Distributed RAM (Mb)	4.8	3.9	4.8	9.7	12.7	14.5	28.7
Block RAM Blocks	1,260	1,421	1,728	2,520	3,276	3,780	2,520
Block RAM (Mb)	44.3	50.0	60.8	88.6	115.2	132.9	88.6
CMT (1 MMCM, 2 PLLs)	10	16	16	20	28	30	30
I/O DLLs	40	64	64	80	120	120	120
Maximum HP I/Os ⁽¹⁾	468	780	780	780	650	650	1,404
Maximum HR I/Os ⁽²⁾	52	52	52	104	52	52	52
DSP Slices	600	672	768	1,200	1,560	1,800	2,880
System Monitor	1	1	1	2	3	3	3
PCIe Gen3 x8	2	4	4	4	4	6	6
150G Interlaken	3	6	6	6	8	9	0
100G Ethernet	3	4	4	6	9	9	3
GTH 16.3Gb/s Transceivers	20	32	32	40	52	60	48
GTY 30.5Gb/s Transceivers	20	32	32	40	52	60	0
Transceiver Fractional PLLs	10	16	16	20	26	30	0

Notes:

1. HP = High-performance I/O with support for I/O voltage from 1.0V to 1.8V.
2. HR = High-range I/O with support for I/O voltage from 1.2V to 3.3V.

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

Table 12: Zynq UltraScale+: CG Device-Package Combinations and Maximum I/Os

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

Notes:

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

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](#).

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

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.

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 512×72 (or 64). Whether configured as block RAM or FIFO, the two ports can have different aspect ratios without any constraints. Each block RAM can be divided into two completely independent 18Kb block RAMs that can each be configured to any aspect ratio from $16K \times 1$ to 512×36 . Everything described previously for the full 36Kb block RAM also applies to each of the smaller 18Kb block RAMs. Only in simple dual-port (SDP) mode can data widths of greater than 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 27×18 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.

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)
		-1C (0.95V)			-1I (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)		

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

The ordering information shown in Figure 3 applies to all packages in the Kintex UltraScale and Virtex UltraScale FPGAs. Refer to the Package Marking section of [UG575](#), *UltraScale and UltraScale+ FPGAs Packaging and Pinouts User Guide* for a more detailed explanation of the device markings.




Figure 3: Kintex UltraScale and Virtex UltraScale FPGA Ordering Information

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.