

General Description

Xilinx® UltraScale™ architecture comprises two high-performance FPGA families that address a vast spectrum of system requirements with a focus on lowering total power consumption through numerous innovative technological advancements.

Kintex® UltraScale FPGAs: High-performance FPGAs with a focus on price/performance, using both monolithic and next generation stacked silicon interconnect (SSI) technology. High DSP and block RAM-to-logic ratios and next generation transceivers, combined with low-cost packaging, enable an optimum blend of capability and cost.

Virtex® UltraScale FPGAs: The industry's most capable high-performance FPGAs enabled using both monolithic and next generation SSI technology to achieve the highest system capacity, bandwidth, and performance. Variants of the Virtex UltraScale family are optimized to address key market and application requirements through integration of various system-level functions, delivering unprecedented embedded memory and serial connectivity capabilities.

Virtex and Kintex UltraScale Families Comparison

Table 1: Device Resources

Range	Kintex UltraScale	Virtex UltraScale
Logic Cells (K)	355–1,160	627–4,407
Block Memory (Mb)	19.0–75.9	44.3–115.2
DSP (Slices)	1,700–5,520	600–2,880
DSP Performance (GMAC/s)	8,180	4,268
Transceivers	16–64	36–104
Peak Transceiver Speed (Gb/s)	16	33
Peak Serial Bandwidth (full duplex) (Gb/s)	2,086	5,101
PCIe® Interface	2–4	2–6
Memory Interface Performance (Mb/s)	2,400	2,400
I/O Pins	312–832	364–1,456
I/O Voltage	1.0–3.3V	1.0–3.3V

Summary of Features

I/O, Transceiver, PCIe, and MAC

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 32.75 Gb/s, enabling 25G+ backplane designs with dramatically lower power per bit than previous generation transceivers. All transceivers support the required data rates for PCI Express® Gen3 and Gen4, and integrated blocks for PCI Express enable UltraScale architecture-based FPGAs to support up to x8 Gen3 Endpoint and Root Port designs. Integrated blocks for 150 Gb/s Interlaken and 100 Gb/s Ethernet (100G MAC) extend the capabilities of UltraScale architecture-based FPGAs, enabling simple, reliable support for Nx100G switch and bridge applications.

Clocks and Memory Interface

UltraScale architecture-based FPGAs 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 within the FPGA to minimize the skew, power consumption, and delay associated with clock signals. In addition, the clock management technology is tightly integrated with dedicated memory interfacing circuitry to enable support for high-performance external memories, including DDR4.

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, and 36 Kb block RAMs with built-in FIFO and ECC support are all connected together 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 using dedicated, low-latency interface tiles. These combined routing resources enable easy support for next generation bus data widths.

Configuration, Encryption, and System Monitoring

The configuration and encryption block performs numerous device-level functions critical to the successful operation of the FPGA. This high-performance configuration block enables device configuration from external media through various protocols, including PCI Express®, often with no requirement to use multi-function I/O pins during configuration. The configuration block also provides 256-bit AES 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 (SYSMON) block enables the monitoring of the physical environment via on-chip temperature and supply sensors and can also monitor up to 17 external analog inputs.

Kintex UltraScale FPGA Feature Summary

Table 2: Kintex UltraScale FPGA Feature Summary

	XCKU035	XCKU040	XCKU060	XCKU075	XCKU100	XCKU115
CLBs	25,391	30,300	41,460	54,000	68,460	82,920
Logic Cells	355,474	424,200	580,440	756,000	958,440	1,160,880
CLB Flip-Flops	406,256	484,800	663,360	864,000	1,095,360	1,326,720
Maximum Distributed RAM (Mb)	5.9	7.0	9.1	7.2	12.8	18.3
Block RAM/FIFO w/ECC (36 Kb each)	540	600	1,080	1,188	1,680	2,160
Total Block RAM (Mb)	19.0	21.1	38.0	41.8	59.1	75.9
CMTs (1 MMCM, 2 PLLs)	8	8	12	14	24	24
I/O DLLs	40	40	48	64	64	64
Maximum HP I/Os	416	416	520	624	676	676
Maximum HR I/Os	104	104	104	104	156	156
DSP Slices	1,700	1,920	2,760	2,592	4,200	5,520
System Monitor	1	1	1	1	2	2
PCIe® Blocks	2	3	2	4	4	4
Interlaken	0	0	0	2	0	0
100G Ethernet	0	0	0	1	0	0
GTH 16 Gb/s Transceivers	16	20	32	52	64	64

Notes:

1. HR = High Range I/O with support for I/O voltage from 1.0V to 3.3V.
2. HP = High Performance I/O with support for I/O voltage from 1.0V to 1.8V.

Kintex UltraScale Device-Package Combinations and Maximum I/Os

Table 3: Kintex UltraScale Device-Package Combinations and Maximum I/Os

Package	Size (mm)	XCKU035	XCKU040	XCKU060	XCKU075	XCKU100	XCKU115
		HR, HP, GTH	HR, HP, GTH	HR, HP, GTH	HR, HP, GTH	HR, HP, GTH	HR, HP, GTH
FBVA676	27 x 27	104, 208, 16	104, 208, 16				
FBVA900	31 x 31	104, 364, 16	104, 364, 16				
FFVA1156	35 x 35	104, 416, 16	104, 416, 20	104, 416, 28	104, 416, 28		
FFVA1517	40 x 40			104, 520, 32	104, 520, 48		
FLVA1517	40 x 40					104, 520, 48	104, 520, 48
FLVB1517	40 x 40					104, 260, 64	104, 260, 64
FFVA1760	42.5 x 42.5				104, 624, 52		
FLVA1760	42.5 x 42.5					104, 624, 52	104, 624, 52
FLVD1924	45 x 45					156, 676, 52	156, 676, 52
FLVF1924	45 x 45					104, 624, 64	104, 624, 64

Notes:

1. All packages have 1.0 mm ball pitch.
2. FFV and FLV packages are footprint compatible when package code letter designator and pin count are identical.

Virtex UltraScale FPGA Feature Summary

Table 4: Virtex UltraScale FPGA Feature Summary

	XCVU065	XCVU080	XCVU095	XCVU125	XCVU145	XCVU160	XCVU440
CLB	44,760	55,714	67,200	89,520	102,500	115,800	314,820
Logic Cells	626,640	780,000	940,800	1,253,280	1,435,000	1,621,200	4,407,480
CLB Flip-Flops	716,160	891,429	1,075,200	1,432,320	1,640,000	1,852,800	5,037,120
Maximum Distributed RAM (Mb)	4.2	3.9	4.8	8.4	9.4	10.7	28.7
Block RAM/FIFO w/ECC (36 Kb each)	1,260	1,421	1,728	2,520	2,787	3,276	2,520
Total Block RAM (Mb)	44.3	50.0	60.8	88.6	98.0	115.2	88.6
CMT (1 MMCM, 2 PLLs)	10	16	16	20	24	26	30
I/O DLLs	40	64	64	80	96	104	120
Fractional PLLs	5	8	8	10	13	13	0
Maximum HP I/Os	468	780	780	936	988	988	1,404
Maximum HR I/Os	52	52	52	104	52	52	52
DSP Slices	600	672	768	1,200	1,365	1,560	2,880
System Monitor	1	1	1	2	3	3	3
PCIe Blocks	2	4	4	4	4	4	6
Interlaken	3	6	6	6	9	9	0
100G Ethernet	3	4	4	6	7	7	3
GTH 16 Gb/s Transceivers	20	32	32	40	52	52	48
GTY 33 Gb/s Transceivers	20	32	32	40	52	52	0

Notes:

1. HR = High Range I/O with support for I/O voltage from 1.0V to 3.3V.
2. HP = High Performance I/O with support for I/O voltage from 1.0V to 1.8V.

Virtex UltraScale Device-Package Combinations and Maximum I/Os

Table 5: Virtex UltraScale Device-Package Combinations and Maximum I/Os

Package	Size (mm)	XCVU065	XCVU080	XCVU095	XCVU125	XCVU145	XCVU160	XCVU440
		HR, HP GTH, GTY	HR, HP GTH, GTY	HR, HP GTH, GTY	HR, HP GTH, GTY	HR, HP GTH, GTY	HR, HP GTH, GTY	HR, HP GTH, GTY
FFVC1517	40 x 40	52, 468 20, 20	52, 468 24, 24	52, 468 24, 24				
FFVB1517	40 x 40		52, 312 32, 32	52, 312 32, 32				
FLVB1517	40 x 40				52, 312 40, 32			
FFVA1760	42.5 x 42.5		52, 676 32, 16	52, 676 32, 16				
FLVA1760	42.5 x 42.5				52, 676 36, 16			
FFVD1924	45 x 45		52, 780 28, 24	52, 780 28, 24				
FLVD1924	45 x 45				52, 780 28, 24	52, 780 28, 24	52, 780 28, 24	
FFVE1924	45 x 45		52, 624 32, 32	52, 624 32, 32				
FLVE1924	45 x 45				52, 624 36, 36	52, 624 36, 36	52, 624 36, 36	
FLVJ1924	45 x 45				52, 312 40, 40	52, 312 52, 52	52, 312 52, 52	
FLVA2377	50 x 50				104, 936 28, 24	52, 988 28, 24	52, 988 28, 24	
FLVB2377	50 x 50							52, 1248 36, 0
FLVA2892	55 x 55							52, 1404 48, 0

Notes:

1. All packages have 1.0 mm ball pitch.
2. FFV and FLV packages are footprint compatible when package code letter designator and pin count are identical.


Device Layout

UltraScale architecture-based FPGAs 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. Figure 1 shows a device-level view with resources grouped together. For simplicity, certain resources such as integrated blocks for PCI Express, Configuration logic, and System Monitor are not shown. The UltraScale architecture-based FPGAs have one or two columns of transceivers. Figure 1 shows a device with two columns of transceivers.



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.



For graphical representation only, does not represent a real device. DS890_02_032513

Figure 2: Column-Based FPGA Divided into Clock Regions

Input/Output

The number of I/O pins 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 either High Range (HR) or High Performance (HP). The HR I/Os offer the widest range of voltage support, from 1.0V to 3.3V. The HP I/Os are optimized for highest performance operation, from 1.0V to 1.8V.

All I/O pins are organized in banks, with 52 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.

I/O Electrical Characteristics

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

Most signal pin pairs can be configured as differential input pairs or output pairs. Differential input pin pairs can optionally be terminated with a 100 Ω internal resistor. All UltraScale architecture-based devices support differential standards beyond LVDS, including RSDS, BLVDS, differential SSTL, and differential HSTL. Each of the I/Os supports memory I/O standards, such as single-ended and differential HSTL as well as single-ended and differential SSTL.

3-State Digitally Controlled Impedance and Low Power I/O Features

The 3-state Digitally Controlled Impedance (T_DCI) can control the output drive impedance (series termination) or can provide parallel termination of an input signal to V_{CCO} or split (Thevenin) termination to $V_{CCO}/2$. This allows users to eliminate off-chip termination for signals using T_DCI. In addition to board space savings, the termination automatically turns off when in output mode or when 3-stated, saving considerable power compared to off-chip termination. The I/Os also have low power modes for IBUF and IDELAY to provide further power savings, especially when used to implement memory interfaces.

I/O Logic

Input and Output Delay

All inputs and outputs can be configured as either combinatorial or registered. Double data rate (DDR) is supported by all inputs and outputs. Any input or output can be individually delayed by up to 1,250 ps of delay with a resolution of 5–15 ps. Such delays are implemented as IDELAY and ODELAY. The number of delay steps can be set by configuration and can also be incremented or decremented while in use. The IDELAY and ODELAY can be cascaded together to double the amount of delay in a single direction.

ISERDES and OSERDES

Many applications combine high-speed, bit-serial I/O with slower parallel operation inside the device. This requires a serializer and deserializer (SerDes) inside the I/O logic. Each I/O pin possesses an IOSERDES (ISERDES and OSERDES) capable of performing serial-to-parallel or parallel-to-serial conversions with programmable widths of 2, 4, or 8 bits. These I/O logic features enable high-performance interfaces, such as Gigabit Ethernet/1000BaseX/SGMII, to be moved from the transceivers to the SelectIO interface.

High-Speed Serial Transceivers

Ultra-fast serial data transmission between devices on the same PCB, over backplanes, and across even longer distances is becoming increasingly important for scaling to 100 Gb/s and 400 Gb/s line cards. Specialized dedicated on-chip circuitry and differential I/O capable of coping with the signal integrity issues are required at these high data rates.

There are two styles of transceivers used in UltraScale architecture-based FPGAs: GTH and GTY, arranged in groups of four, known as a transceiver Quad. Each serial transceiver is a combined transmitter and receiver. [Table 6](#) compares the available transceivers.

Table 6: Transceiver Information

	Kintex UltraScale	Virtex UltraScale	
Style of Transceiver	GTH	GTH	GTY
Quantity of Transceivers	16–64	20–52	0–52
Maximum Data Rate	16.3 Gb/s	16.3 Gb/s	32.75 Gb/s
Minimum Data Rate	0.5 Gb/s	0.5 Gb/s	0.5 Gb/s
Example Applications	Backplane, PCI Express Gen4	Backplane, PCI Express Gen4	100G+ Optical, Chip to Chip, 25G+ Backplane

The serial transmitter and receiver are independent circuits that use an advanced 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 FPGA 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 guarantees sufficient data transitions by using the selected encoding scheme. Parallel data is then transferred into the FPGA logic using the RXUSRCLK clock. For short channels, the transceivers offer a special low-power mode (LPM) to reduce power consumption by approximately 30%. The receiver DC automatic gain control and linear and decision feedback equalizers can optionally “auto-adapt” to automatically learn and compensate for different interconnect characteristics. This enables even more margin for tough 10G+ and 25G+ backplanes.

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 PCI Express and SATA/SAS and QPI applications.

Integrated Interface Blocks for PCI Express Designs

All UltraScale architecture-based FPGAs include at least one integrated block for PCI Express technology that can be configured as an Endpoint or Root Port, compliant to the PCI Express Base Specification Revision 3.0. The Root Port can be used to build the basis for a compatible Root Complex, to allow custom FPGA-to-FPGA communication via the PCI Express protocol, and to attach ASSP Endpoint devices, such as Ethernet Controllers or Fibre Channel HBAs, to the FPGA.

This block is highly configurable to system design requirements and can operate 1, 2, 4, or 8 lanes at the 2.5 Gb/s, 5.0 Gb/s, and 8.0 Gb/s 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 PCI Express, the transceivers, block RAM, and clocking resources) into an Endpoint or Root Port solution. The system designer has control over many configurable parameters: lane width, maximum payload size, FPGA logic interface speeds, reference clock frequency, and base address register decoding and filtering.

Integrated Block for Interlaken

Some UltraScale architecture-based FPGAs include integrated blocks for Interlaken. Interlaken is a scalable chip-to-chip interconnect protocol designed to enable transmission speeds from 10 Gb/s to 150 Gb/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 10.3125 Gb/s; 1 to 12 lanes at 12.5 Gb/s and 1 to 6 lanes at 25.78125 Gb/s, enabling flexible support for up to 150 Gb/s per integrated block. With multiple Interlaken blocks, certain UltraScale architecture-based FPGAs enable easy, reliable Interlaken switches and bridges.

Integrated Block for Ethernet

Compliant to the IEEE Std 802.3ba, the 100G Ethernet integrated blocks in the UltraScale architecture provide low latency 100 Gb/s Ethernet ports with a wide range of user customization and statistics gathering. With support for 10 x 10.3125 Gb/s (CAUI) and 4 x 25.78125 Gb/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 PTP time stamping.

Clock Management

The clock generation and distribution components in UltraScale architecture-based FPGAs are located adjacent to the columns that contain the memory interfacing and input and output circuitry. This tight coupling of clocking and I/O provides low-latency clocking to the I/O for memory interfacing 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 1,600 MHz, the phase-shift timing increment is 11.2 ps.

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 FPGA fabric per PLL as well as one clock plus one enable signal to the memory interface circuitry.

Clock Distribution

Clocks are distributed throughout UltraScale architecture-based FPGAs 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.

Memory Interfacing

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 architecture-based FPGA includes dedicated PHY blocks located between the CMT and I/O columns that support implementation of high-performance physical interfaces (PHY) 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.

Configurable Logic Block

Every Configurable Logic Block (CLB) in UltraScale architecture-based FPGAs 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

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 UltraScale architecture-based FPGAs have increased routing and connectivity compared to CLBs in previous generation Xilinx® FPGAs. 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-based FPGAs 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 run time.

Stacked Silicon Interconnect (SSI) Technology

Many challenges associated with creating high capacity FPGAs 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 FPGA with more than ten thousand low-power inter-SLR connections. Dedicated interface tiles within the SLRs provide ultra-high bandwidth, low latency connectivity to other SLRs.

Block RAM

Every UltraScale architecture-based FPGA 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 36 Kb RAM or two independent 18 Kb 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 $32\text{K} \times 1$; $16\text{K} \times 2$; $8\text{K} \times 4$; $4\text{K} \times 9$ (or 8); $2\text{K} \times 18$ (or 16); $1\text{K} \times 36$ (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 18 Kb block RAMs that can each be configured to any aspect ratio from $16\text{K} \times 1$ to 512×36 . Everything described previously for the full 36 Kb block RAM also applies to each of the smaller 18 Kb block RAMs. Only in simple dual-port (SDP) mode can data widths of greater than 18 bits (18 Kb RAM) or 36 bits (36 Kb 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 36 Kb 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 36 Kb FIFO or an 18 Kb 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.

Digital Signal Processing

DSP applications use many binary multipliers and accumulators, best implemented in dedicated DSP slices. All UltraScale architecture-based FPGAs 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 two's 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 block in the UltraScale architecture is used to enhance the overall safety, security, and reliability of the system by monitoring of the physical environment via on-chip supply and temperature sensors as well as up to 17 user-allocated external analog inputs. The System Monitor supports on-chip monitoring of all the device's major supply voltages (i.e., V_{CCINT} , V_{CCAUX} , V_{CCBRAM} , and V_{CCO}).

The sensor outputs and 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, JTAG, or I2C interfaces. The I2C interface allows the on-chip monitoring to be easily accessed by the System Manager/Host before and after FPGA configuration.

Configuration

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

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. 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 control access port (MCAP) provides a direct connection between the integrated block for PCI Express and the configuration logic to simplify configuration over PCI Express.

The FPGA has the ability to reconfigure itself with a different image using SPI or BPI flash, eliminating the need for an external controller. The FPGA can reload its original design in case there are any errors in the data transmission, ensuring an operational FPGA at the end of the process. This is especially useful for updates to a design after the end product has been shipped. Customers can ship their products with an early version of the design, thus getting their products to market faster. This feature allows customers to keep their end users current with the most up-to-date designs while the product is already in the field.

Packaging

The UltraScale architecture-based FPGAs 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

This document is a pre-release document, provided ahead of silicon ordering availability. Therefore, Xilinx UltraScale architecture-based FPGA ordering information is not yet available. Please contact your Xilinx sales representative for more information on UltraScale devices' Early Access Programs.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
12/10/2013	1.0	Initial Xilinx release.

Notice of 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

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS.