



AMD Artix UltraScale+ XA FPGA and Product Data Sheet: Overview

DS886 (v1.1) September 26, 2024

Product Specification

General Description

Using the UltraScale™ architecture, the AMD Artix™ UltraScale+™ XA FPGAs offer the highest serial bandwidth and signal compute density in a cost-optimized device for critical networking applications, vision and video processing, and secured connectivity.

Summary of Features

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 58.0 Gb/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 8.0 GT/s (Gen3), and 16.0 GT/s (Gen4) for PCIe®. The integrated blocks for PCIe can be configured for Endpoint or Root Port, supporting a variety of link widths and speeds depending on the targeted device speed grade and package. Integrated blocks for 150 Gb/s Interlaken and 100 Gb/s Ethernet (100G MAC/PCS) extend the capabilities of UltraScale devices, enabling simple, reliable support for Nx100G switch and bridge applications.

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, 36 Kb 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,

AMD Adaptive Computing is creating an environment where employees, customers, and partners feel welcome and included. To that end, we're removing noninclusive language from our products and related collateral. We've launched an internal initiative to remove language that could exclude people or reinforce historical biases, including terms embedded in our software and IPs. You may still find examples of non-inclusive language in our older products as we work to make these changes and align with evolving industry standards. Follow this [link](#) for more information.

DS886 (v1.1) September 26, 2024

Product Specification

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.

Configuration, Encryption, and System Monitoring

The configuration and encryption block performs numerous device-level functions critical to the successful operation of the FPGAs. 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.

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.

Artix UltraScale+ XA FPGA Feature Summary

Table 1: Artix UltraScale+ XA FPGA Feature Summary

	XAAU7P	XAAU10P	XAAU15P
System Logic Cells	81,900	96,250	170,100
CLB Flip-Flops	74,880	88,000	155,520
CLB LUTs	37,440	44,000	77,760
Max. Distributed RAM (Mb)	1.1	1.0	2.5
Block RAM Blocks	108	100	144
Block RAM (Mb)	3.8	3.5	5.1
CMTs (1 MMCM and 2 PLLs)	2	3	3
Max. HP I/O ⁽¹⁾	104	156	156
Max. HD I/O ⁽²⁾	144	72	72
DSP Slices	216	400	576
System Monitor	1	1	1
GTH Transceiver ⁽³⁾	4	12	12
Transceiver Fractional PLLs	2	6	6
PCIE4C ⁽⁴⁾	1	1 ⁽⁵⁾	1 ⁽⁵⁾

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 transceiver line rates are package limited: SBVB484 to 12.5 Gb/s.
4. This block operates in compatibility mode for 16.0 GT/s (Gen4) operation. Go to [PG213](#), *UltraScale+ Devices Integrated Block for PCI Express Product Guide*, for details on compatibility mode.
5. PCIe Gen4 is available in XAAU10P and XAAU15P in the FFVB676 package only.

Artix UltraScale+ XA Device-Package Combinations and Maximum I/Os

Table 2: Artix UltraScale+ Device-Package Combinations and Maximum I/Os

Package ⁽¹⁾⁽²⁾⁽³⁾	Package Dimensions (mm)	XAAU7P	XAAU10P	XAAU15P
		HDIO, HPPIO, GTH, GTY		
FCVA289	9x9	72, 58, 4, 0		
SBVB484	19x19		48, 156, 12, 0	48, 156, 12, 0
SBVC484	19x19	144, 104, 4, 0		
FFVB676	27x27		72, 156, 12, 0	72, 156, 12, 0

Notes:

1. Go to [Ordering Information](#) for package designation details.
2. FF packages have 1.0mm ball pitch. SB packages have 0.8mm ball pitch.
3. Packages with the same last letter and number sequence, e.g., B676, 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.

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. [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.




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

Input/Output

All UltraScale devices have I/O pins for communicating to external components. In addition, in the 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 or RFSoc interfacing capability, referred to as extended MIO (EMIO).

The number of I/O pins in UltraScale FPGAs 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.

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 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. UltraScale+ families add support for MIPI with a dedicated D-PHY in the I/O bank.

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

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.

Artix UltraScale+ XA devices use GTH transceivers. All transceivers are arranged in groups of four, known as a transceiver Quad. Each serial transceiver is a combined transmitter and receiver. [Table 3](#) compares the available transceivers.

Table 3: Transceiver Information

	XA Artix UltraScale+
Type	GTH
Max. Data Rate	16.3 Gb/s
Min. Data Rate	0.5 Gb/s

GTH Transceivers

The serial transmitter and receiver are independent circuits that use an advanced phase-locked loop (PLL) architecture to multiply the reference frequency input by certain programmable numbers between 4 and 25 to become the bit-serial data clock. Each transceiver has a large number of user-definable features and parameters. All of these can be defined during device configuration, and many can also be modified during operation.

Transmitter

The transmitter is fundamentally a parallel-to-serial converter with a conversion ratio of 16, 20, 32, 40, 64, or 80 for the GTH. This allows the designer to trade off datapath width against timing margin in high-performance designs. These transmitter outputs drive the PC board with a single-channel differential output signal. TXOUTCLK is the appropriately divided serial data clock and can be used directly to register the parallel data coming from the internal logic. The incoming parallel data is fed through an optional FIFO and has additional hardware support for the 8B/10B, 64B/66B, or 64B/67B encoding schemes to provide a sufficient number of transitions. The bit-serial output signal drives two package pins with differential signals. This output signal pair has programmable signal swing as well as programmable pre- and post-emphasis to compensate for PC board losses and other interconnect characteristics. For shorter channels, the swing can be reduced to reduce power consumption.

Receiver

The receiver is fundamentally a serial-to-parallel converter, changing the incoming bit-serial differential signal into a parallel stream of words, each 16, 20, 32, 40, 64, or 80 bits in the GTH. This allows the designer to trade off internal datapath width against logic timing margin. The receiver takes the incoming differential data stream, feeds it through programmable DC automatic gain control, linear and decision feedback equalizers (to compensate for PC board, cable, optical and other interconnect characteristics), and uses the reference clock input to initiate clock recognition. There is no need for a separate clock line. The data pattern uses non-return-to-zero (NRZ) encoding and optionally ensures sufficient data transitions by using the selected encoding scheme. Parallel data is then transferred into the device logic

using the RXUSRCLK clock. For short channels, the transceivers offer a special low-power mode (LPM) to reduce power consumption by approximately 30%. The receiver DC automatic gain control and linear and decision feedback equalizers can optionally “auto-adapt” to automatically learn and compensate for different interconnect characteristics. This enables even more margin for 10G+ and 25G+ backplanes.

Out-of-Band Signaling

The transceivers provide out-of-band (OOB) signaling, often used to send low-speed signals from the transmitter to the receiver while high-speed serial data transmission is not active. This is typically done when the link is in a powered-down state or has not yet been initialized. This benefits PCIe and SATA/SAS and QPI applications.

Integrated Interface Blocks for PCI Express Designs

Artix UltraScale+ XA devices use PCIe4C, which are compliant to PCI Express Base Specification v3.1 and support up to Gen3 x8, and can also be configured for lower link width and speeds.

PCIe4C blocks are:

- Compliant to the PCI Express Base Specification v3.1 supporting up to 8.0 GT/s (Gen3)
- Compatible with PCI Express Base Specification v4.0 supporting up to 16.0 GT/s (Gen4)
- Compliant with CCIX Base Specification v1.0 Version 0.9, supporting speeds up to 16.0 GT/s
- Support up to 8 lanes at Gen3 or up to 8 lanes at Gen4 and can be configured for lower link widths and speeds to conserve resources and power

All integrated blocks for PCIe in the UltraScale architecture can be configured as Endpoint or Root Port. The Root Port can be used to build the basis for a compatible Root Complex, to allow custom chip-to-chip communication via the PCI Express protocol, and to attach ASSP Endpoint devices, such as Ethernet Controllers or Fibre Channel HBAs, to the FPGA, MPSoC, or RFSoc.

The maximum lane widths and data rates per family are listed in [Table 4](#).

Table 4: PCIe Maximum Configurations

	Artix UltraScale+ XA
Gen1 (2.5 GT/s)	x8
Gen2 (5 GT/s)	x8
Gen3 (8 GT/s)	x8
Gen4 (16 GT/s)	x8 ⁽¹⁾

Notes:

1. PCIe Gen4 is available in XAAU10P and XAAU15P in the FFVB676 package.

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.

AMD provides LogiCORE™ IP options to configure the integrated blocks for PCIe in all Artix UltraScale+ XA devices. This includes AXI Streaming interfaces at the PCIe packet level and more advanced IP such as AXI to PCIe Bridges and DMA engines. This IP gives the designer control over many configurable parameters such as link width and speed, maximum payload size, and reference clock frequency. For a complete list of features that can be configured for each of the IP, go to the specific Product Guide.

Cache Coherent Interconnect for Accelerators (CCIX)

CCIX is a chip-to-chip interconnect operating at data rates up to 25 Gb/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 16 Gb/s) with an additional high-speed 25 Gb/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.

PCIe4C blocks support CCIX data rates up to 16 Gb/s and contain one CCIX port. 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 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 up to 12.5 Gb/s and 1 to 6 lanes at up to 25.78125 Gb/s, enabling flexible support for up to 150 Gb/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 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 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.

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 1,600 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 device fabric per PLL as well as one clock plus one enable signal to the memory interface circuitry.

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 and RFSocS, 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 architecture-based devices 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 15 Gb/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 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 $32K \times 1$; $16K \times 2$; $8K \times 4$; $4K \times 9$ (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 18 Kb block RAMs that can each be configured to any aspect ratio from $16K \times 1$ 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.

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 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.

Table 5: Key System Monitor Features

	XA Artix UltraScale+
ADC	10-bit, 200 kSPS
Interfaces	JTAG, I2C, DRP, PMBus

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.

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.

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.

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

UltraScale devices are available in a variety of organic 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 6 shows the speed and temperature grades available in the different device families. V_{CCINT} supply voltage is listed in parentheses.

Table 6: Speed Grade and Temperature Grade


Device Family	XA Devices	Speed Grade	Temperature Grade
		Industrial (I)	Automotive (Q)
		-40°C to +100°C	-40°C to +125°C
XA Artix UltraScale+ ⁽¹⁾	All	-1I	-1Q
		-1L	

Notes:

1. In Artix UltraScale+, 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 Artix UltraScale+ FPGAs.

The -1L speed grades in the UltraScale+ families can run at one of two different V_{CCINT} operating voltages. At 0.72V, they operate at similar performance to the Kintex UltraScale and Virtex UltraScale devices with up to 30% reduction in power consumption. At 0.85V, they consume similar power to the Kintex UltraScale and Virtex UltraScale devices, but operate over 30% faster.



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

DS886_03_022123

Figure 3: Ordering Information

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
09/26/2024	1.1	Added package for XAAU7P in Table 2 .
05/15/2023	1.0	Initial release.

Please Read: Important Legal Notices

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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.

© Copyright 2023–2024 Advanced Micro Devices, Inc. AMD, the AMD Arrow logo, Artix, UltraScale+, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. PCI, PCIe, and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.