



WP497 (v1.0) October 23, 2017

Xilinx Advanced Multimedia Solutions with Video Codec/Graphics Engines

By: Yashu Gosain and Alok Gupta

Xilinx's efficient power-optimized All-Programmable System-on-Chip multimedia solution with integrated GPU and codec, designed for Ultra High-Definition (UHD) video, provides exceptional design flexibility and enables innovative product differentiation.

ABSTRACT

This white paper explores the multimedia aspects of the Zynq® UltraScale+™ MPSoC. Based on the Zynq-7000 SoC, the MPSoC includes a processing system (PS) bridged with programmable logic (PL)—but it is the additional extensions in the Zynq UltraScale+ MPSoC that make it a perfect fit for multimedia applications. Key differentiators include:

- Integrated video codec unit that supports H.264/H.265 for UHD-4K video, suitable but not limited to surveillance, video conferencing, embedded vision, and some broadcast applications, with the ability to achieve low latency for encoding, while being more power-efficient relative to a server or programmable logic.
- The power-optimized graphics subsystem is based on the ARM® Mali-400, a proven and established architecture, ideal for the 2D and 3D graphics.
- PL programmability provides: (a) flexibility and scalability in terms of any-to-any high-speed video/audio interface, and (b) differentiation to enable customized image and video processing for multimedia pipelines.

The white paper also highlights how the processing engines, together with the unique configuration of the Zynq UltraScale+ MPSoC, target the next generation of multimedia needs.

Introduction

Media networks from glass-to-glass involve the acquisition edge, where content is captured or created, and the consumption edge, where content is watched or consumed on a growing array of devices. At the acquisition edge, the proliferation of advanced sensors capable of 4K resolution (and beyond), combined with the growing adoption of 4K-capable (3840x2160) consumer displays at the consumption edge is enabling not only higher resolution content, but also exciting innovations such as 360 degree video and VR (virtual reality). Additional technology enablers of compelling content include HDR (high-dynamic range) for HD and UHD resolutions, and WCG (wide color gamut) for deep colors.

Improved sensors, better displays, higher resolutions, HDR, WCG, and higher frame rates combine to enable much higher video quality with whiter whites, blacker blacks, wider and better color accuracy, and fewer motion artifacts.

Compression is a standard technique to reduce data bandwidth and storage size, and a number of codecs are available for different use cases, trading off key parameters such as output quality, latency, size, power consumption, and cost. Advancement in the codec standard is driven by consumers' desire to experience ever-better HD or UHD quality on any screen. These higher expectations are intimately linked with ensuring delivery to a maximum number of consumers through low-bandwidth networks, where the threats of limited access, slow playback capabilities, and unreliable connections are constantly looming.

Today, H.264 (MPEG-4, AVC (Part 10)) is the most widely accepted and adopted format in the online and broadcast domains for high-definition content compression and distribution. But consumer desire for even higher quality video has driven migration from high-definition content to UHD video content, which requires four times the storage or transmission bandwidth. It poses a huge challenge to the current deployed infrastructure to distribute the video content. Consequently, with respect to the existing codec standard, there is a need to compress more without impacting quality.

Advances in video compression technology have introduced dramatic efficiency in terms of system bandwidth and storage. High Efficiency Video Coding (HEVC), or H.265 compression, is becoming a key enabler of UHD; it promises to provide up to a 50% increase in bandwidth and storage savings compared to the currently deployed H.264 compression standard. The bandwidth efficiency associated with the adoption of HEVC promises to greatly expand delivery of high-quality, high-resolution video over bandwidth-constrained networks. For instance, imagine a multiview application that takes multiple videos from different sources and stitches them together using real-time video processing (like scaling, combining, or overlaying) for display on a single 4K output monitor. (See [Figure 1.](#))

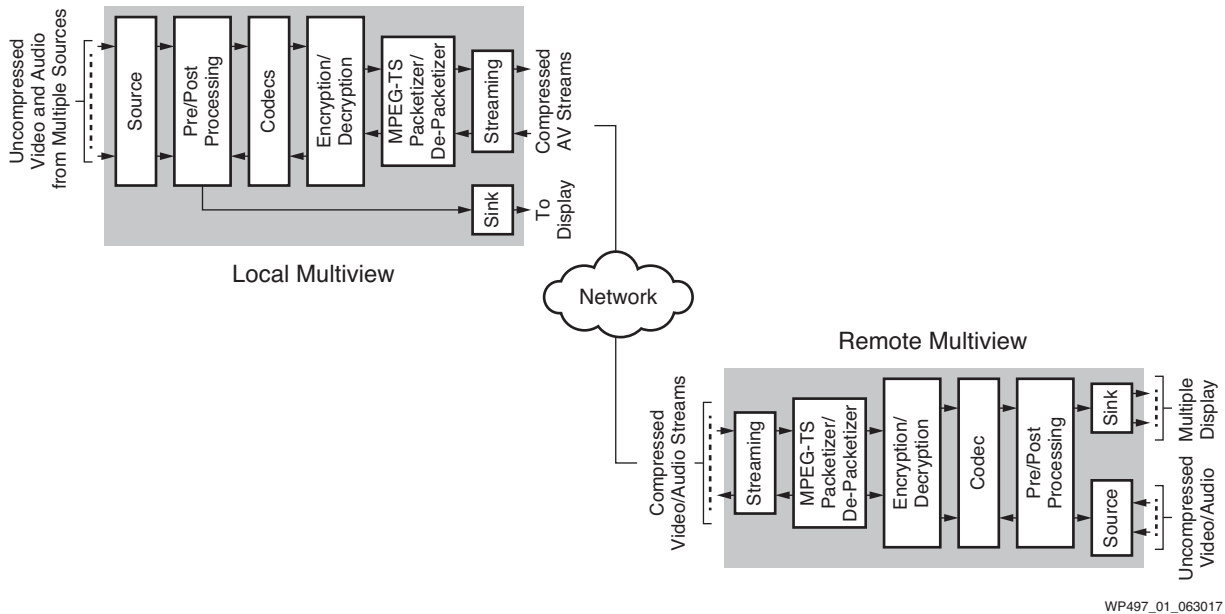


Figure 1: Multiview Application with Real-Time Processing

The video source interface and video processing can be run in real time to provide an uncompressed output. The multiview output can be compressed significantly using a codec, allowing the output to be transmitted over Ethernet. Transmission of video over IP brings serious attention to security and privacy issues today.

To protect digital content, codec output gets encrypted and packetized in transport streams so the content can be streamed to a remote monitoring station or digital signage anywhere in the world.

At remote monitoring stations, multiple streams of compressed video received in real time is depacketized and decrypted to extract the video content. The decrypted video output is decompressed using a given codec. The video stream (decompressed output from the codec) is then given to a video post-processing block to perform overlays, noise cancellation, composition, etc., for single or multiple 4k display. A remote monitoring application can also have a return channel to capture multiple video sources, and then can compress, encrypt, and packetize the stream to transmit over the network.

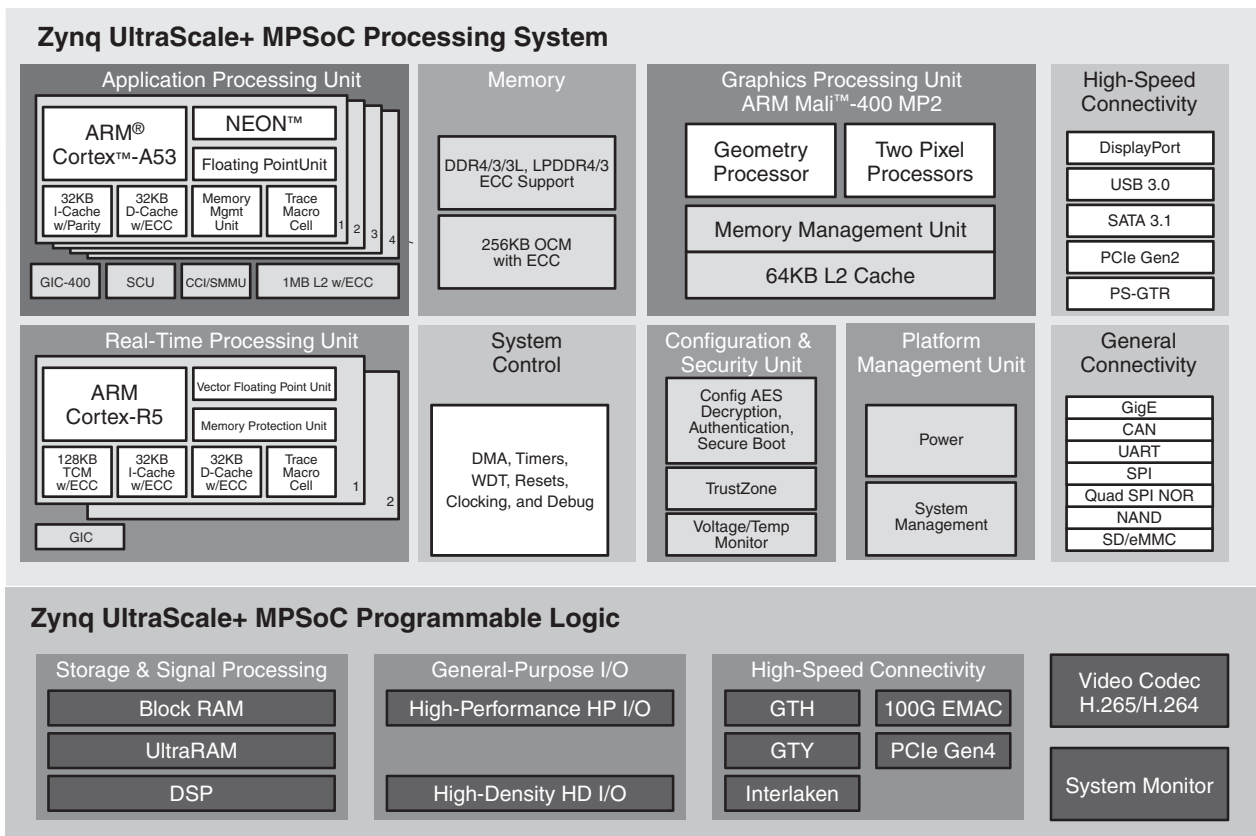
To achieve this kind of advanced topology, a system must support various codecs and graphics engines, as well as offer tight integration with multimedia connectivity interfaces, real-time video processing pipelines, and IP networking capabilities.

Such complex topology is simplified by adopting the Xilinx® Zynq® UltraScale+™ MPSoC as a platform, which integrates a multicore processor subsystem, graphics processing unit, HEVC/AVC-compliant video codec unit, and a wide range of on-chip peripherals tightly coupled to the programmable logic (PL).

This white paper explores the multimedia aspects of the Zynq UltraScale+ MPSoC EV devices and highlights how the processing engines, together with the unique device configuration, target the next generation of multimedia needs.

Zynq UltraScale+ MPSoC Multimedia Solution

The Zynq UltraScale+ MPSoC is a heterogeneous SoC comprising many processing engines, a list of high-speed peripherals, advanced I/O capabilities, and PL. (See [Figure 2](#).) The processing engines include the quad-core ARM® Cortex™ A53-based APU, dual-Core ARM Cortex R5-based RPU, Mali graphics processing unit, platform management unit, and video codec unit (VCU). It can offload critical applications like graphics and video pipelining to dedicated processing blocks, as well as turn blocks on and off through efficient power domains and gated power islands. With a wide range of interconnect options, DSP blocks, and PL choices, the Zynq UltraScale+ MPSoC has the overall flexibility to fit many user application needs.



WP470_01_051016

Figure 2: Block Diagram

The scalability of the product family provides designers with the perfect fit for both cost-sensitive and high-performance applications using a single platform and industry-standard tools.

To achieve high-performance, advanced multimedia systems must have the right processing engines and the capability to add custom logic for differentiation. They must also support any-to-any connectivity required by the wide range of multimedia devices available today. Traditionally, these requirements suggest a multi-chip solution. While a multi-chip solution might provide the required multimedia and connectivity functions, it can also lead to high power consumption. A single-chip solution like the Zynq UltraScale+ MPSoC solves this problem. The existence of custom logic for hardware acceleration or any-to-any connectivity on the same device can lead to significant power savings. In addition to the mix of processing engines, hardware

codecs (with their own power management features), and support for custom logic, the Zynq UltraScale+ MPSoC places these components in different power domains with independent power rails. This configuration can be used to design optimized power management schemes for the entire system. The Zynq UltraScale+ MPSoC is built upon the 16nm FinFET process node from TSMC, resulting in greater performance and lower power consumption, enabling the design of power-efficient next-generation multimedia systems.

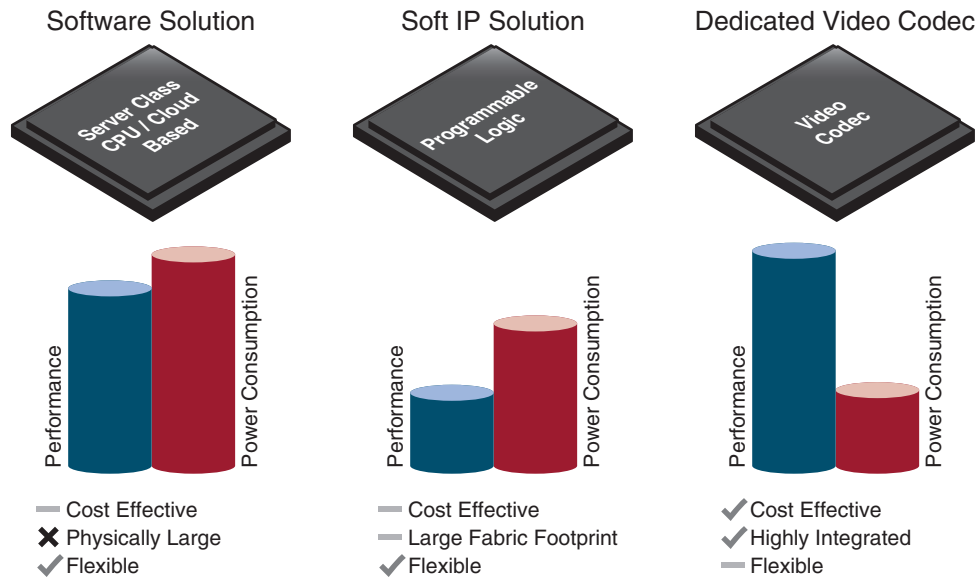
Integrated Video Codec Unit (VCU)

There are several software video codec implementations, but all software, in fact, runs on hardware. With software-based codec solutions, achieving high compression ratios with excellent video quality (VQ) at reasonable frame rates requires several power-hungry server-class CPUs, which are often located in the cloud.

Being necessarily remote from the device, real-time compression cannot typically be implemented. However, being implemented in software, and thus extremely flexible, the system is highly subject to optimization. An alternative solution might be a soft-IP block implemented in the PL. Such an implementation might be lower in performance relative to an offline software model, but it can achieve low-latency compression. Power consumption is improved over a server-class CPU, but not as much as that attained by a fully integrated solution. Usually an advanced codec requires a relatively large programmable logic device; however, post-deployment programmability provides a high degree of flexibility for this solution.

The Zynq UltraScale+ MPSoCs follow an alternative strategy of having a hardened dedicated video codec block. Due to the monolithic die, which includes the hardened VCU, the Zynq UltraScale+ MPSoC provides real time compression with a low latency path and less power consumption than other strategies. Since it is a dedicated codec engine, the application processing unit (APU) and/or the PL can be used for the other computation.

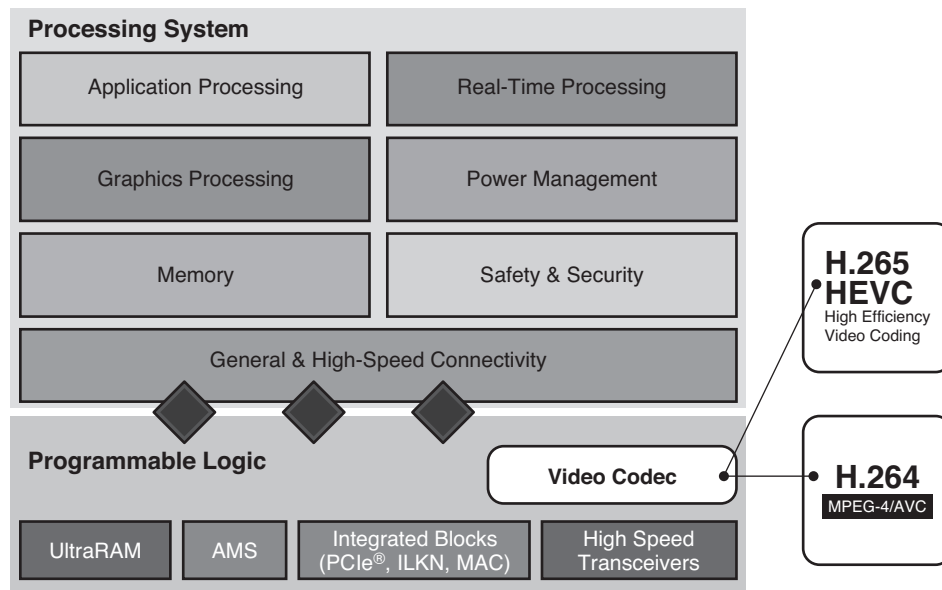
The VCU provides multi-standard encoding and decoding, including the AVC/H.264 and HEVC/H.265 standards. It contains both encode (compress) and decode (decompress) functions, and is capable of simultaneous encode and decode at rates up to 4K UHD at 60Hz. Larger resolutions such as DCI at 4K or higher can be supported at a lower frame rate. The VCU simultaneously supports up to eight different video streams. The total bandwidth requirement of handling different streams simultaneously should not exceed the maximum VCU processing throughput of 4K UHD at 60fps. [Figure 3](#) shows the different strategies for video codec implementation.



WP497_03_101317

Figure 3: Video Codec Implementation Strategies

The VCU is an integrated block in the PL that supports H.265 (HEVC) and H.264 (AVC) standards. There are no direct (hardwired) connections to the processing system (PS). Using the Xilinx Vivado® IP integrator, the designer connects the VCU through the application processing unit (APU) and the memory subsystem in the PS and/or PL. See Figure 4.



WP497_04_063017

Figure 4: Video Codec Block

VCU features:

- Multi-standard encoding/decoding support, including:
 - ISO MPEG-4 Part 10: Advanced Video Coding (AVC)/ITU H.264
 - ISO MPEG-H Part 2: High Efficiency Video Coding (HEVC)/ITU H.265
 - HEVC: Main, Main Intra, Main10, Main10 Intra, Main 4:2:2 10, Main 4:2:2 10 Intra up to 5.1 High Tier
 - AVC: Baseline, Main, High, High10, High 4:2:2 up to 5.2 level
- Supports up to eight streams of encode and eight streams of decode simultaneously
- On-the-fly change of multiple encoding parameters
- Flexible rate control: CBR, VBR, and Constant QP
- Supports simultaneous encoding and decoding up to 4K UHD resolution at 60Hz
- Supports 8K UHD at reduced frame rate (~15Hz)
- Progressive only support for H.264/H.265 standard
- Video input
 - YCbCr 4:2:2, YCbCr 4:2:0, and Y-only
 - 8-bit and 10-bit / component
- Flexible group of picture (GOP) configuration
- Low latency mode
- Power management
 - Active clock gating enabled when idle
- Performance monitoring:
 - Measures task execution time
 - Measures bandwidth and AXI transaction count
 - Measures min, max, and average latency

Unlike software codecs, the VCU in Zynq UltraScale+ MPSoC EV devices provides low-power, high-performance compression and decompression of H.264/H.265 video data. This also makes it an ideal candidate for real-time streaming of UHD videos on the network, where it can save a considerable amount of storage and network bandwidth. VCU support for both H.264 and H.265 standards provides a compelling capability to develop solutions in line with current market needs (H.264) as well as advance-generation requirements (H.265). The ability to encode and decode simultaneously with low latency makes it a perfect fit for video conferencing and transcoding from H.264 to H.265 or vice-versa. Multi-stream multi-codec encoding and decoding suits the requirement for video surveillance applications such as DVRs, video servers, and multistream IP camera head-ends. Supporting video formats up to 4:2:2 10-bit UHD-4K makes the VCU suitable for the professional and prosumer production and post-production solutions. The All Programmable Zynq UltraScale+ MPSoC EV device together with integrated VCU specifications makes it a very compelling cost-effective, single-chip solution for numerous multimedia markets.

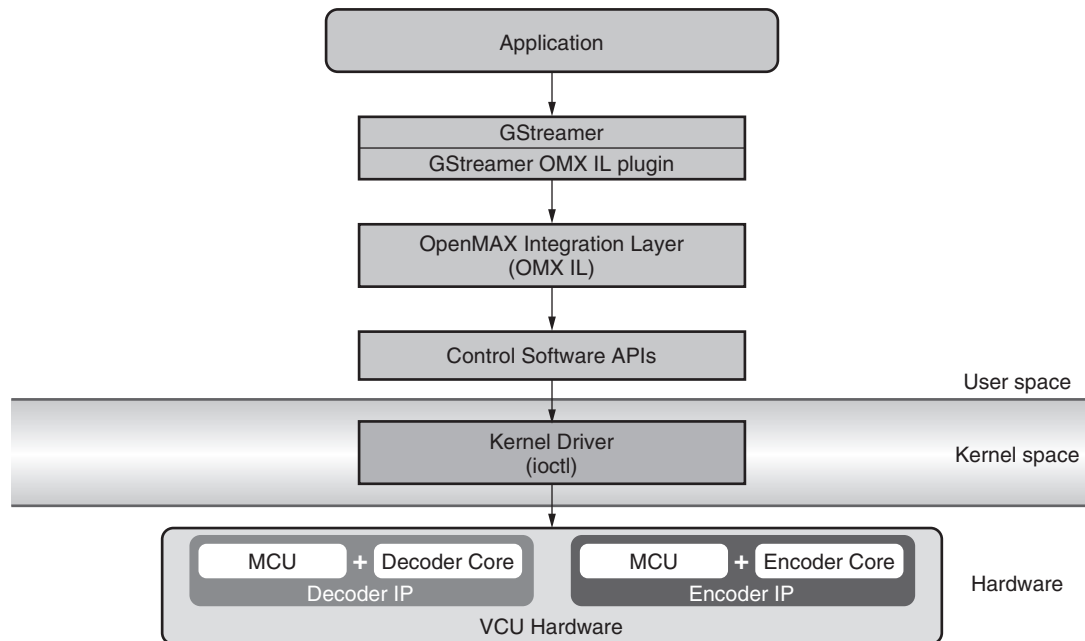
VCU Software Stack

Hardware-accelerated multimedia applications can be developed on the Zynq UltraScale+ MPSoC with a commonly used multimedia framework called GStreamer. GStreamer is based on a plug-in model that is categorized into three functions: source, filter, and sink plug-ins. By chaining together several such plug-ins/elements, a pipeline can be created to do a specific task, e.g, media playback or capture. GStreamer provides a *gst-omx* plug-in, which can perform hardware-accelerated video encoding and decoding on a Zynq UltraScale+ MPSoC. The GStreamer application interacts with an OpenMAX integration layer via the *gst-omx* plug-in. OMX IL defines a standardized media component interface to communicate with a VCU implemented in hardware. This layer interacts with control software (user-space driver) APIs, which in turn call kernel-space drivers.

Kernel drivers communicate with an embedded microcontroller unit (MCU), which is integrated in both the encoder and decoder engines in the VCU. The low-level firmware runs on the MCU to control the encoder/decoder engine. The kernel driver sends the frame-level command to the MCU and waits for the MCU's response. Upon receiving a command from a kernel driver, the MCU schedules the task to an encoder/decoder engine and returns status back to the kernel driver.

The VCU software stack provides the flexibility to be used at any level, depending on the requirements of the user application. The developer can interact with VCU hardware using control software APIs, OMX IL, or GStreamer framework to develop the multimedia pipeline.

The flow for the VCU Linux software stack is depicted in Figure 5.



WP497_05_092117

Figure 5: VCU Software Architecture

Developers can opt to use a specific VCU framework layer based on the complexity of the multimedia application to be developed. GStreamer is a popular choice because of its cross-platform filter, codec support, and ease of use in creating applications such as video editors, transcoders, streaming media broadcasters, and media players, where the system reads files in one format, processes them, and exports them in another format, using a plug-and-play model. It provides the flexibility to integrate pluggable components that can be mixed and matched into arbitrary pipelines so that it is possible to write a full-fledged multimedia application. Alternatively, to develop customized multimedia applications/frameworks with simple, light-weight design, an application can opt to use control software APIs.

Integrated Graphics Processing Unit (GPU)

The GPU in the Zynq UltraScale+ MPSoC processing system (PS) is the ARM Mali-400 MP2, which is tied directly to the APU and can optionally accelerate the rendering of video images in a frame buffer for display output. (Refer to [Figure 6](#).)

The GPU can perform pixel rendering through its dedicated parallel engines much faster than competing ASSPs, which rely on the CPU to handle graphics processing, and cheaper and with less power consumption than solutions that rely on the designer to add an off-chip GPU engine.

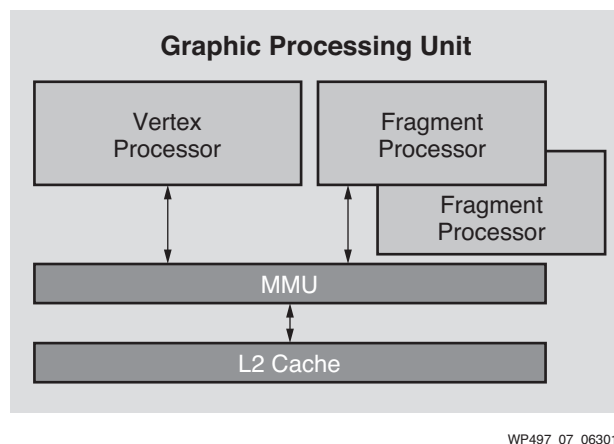


Figure 6: Mali-400 MP2 Architecture Design

The GPU accelerates both 2D and 3D graphics with a fully programmable architecture that provides support for both shader-based and fixed-function graphics APIs. It includes anti-aliasing for optimal image quality, with virtually no additional performance overhead. A full set of proven and tested drivers for Linux is included; these can handle automatic offloading of graphics commands from the APU to the CPU.

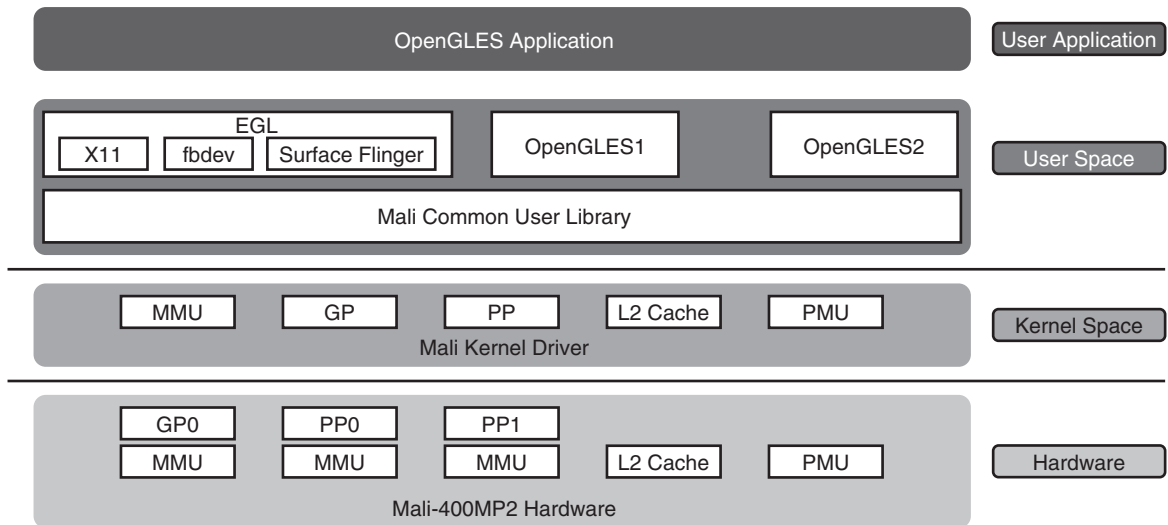
Highlights of the Zynq UltraScale+ MPSoC GPU:

- ARM Mali-400 MP2
- Up to 667MHz performance in the fastest speed grade
- One geometry processor, two pixel processors
- Dedicated 64KB shared L2 cache
- Dedicated memory management unit
- OpenGL ES 2.0 and OpenGL ES 1.1 support
- OpenVG 1.1 API support
- Individual power gating on each of the three engines
- 1,334 Mpixels/sec pixel fill rate
- 72.6 Mtriangles/sec
- 12Gflops floating-point shading

These performance numbers are based on the maximum GPU clock rate of 667MHz.

GPU Software Stack

The Zynq Ultra Scale+ MPSoC graphics software stack is divided in three major layers: kernel space, user space, and application space. See [Figure 7](#)



WP497_08_092717

Figure 7: Graphics Software Stack

Kernel Space

The kernel-space driver interacts with GPU hardware to provide the information necessary to offload tasks to integrated hardware shaders. It also performs low-level interrupt handling and memory management.

User Space

The user-space driver features the following components:

- *Mali Common User Library*: This layer is responsible for GPU job creation as well as scheduling the task for the graphics rendering pipeline.
- *Extended Graphics Library (EGL)*: This layer is used to interact with underneath windowing layer used to interact with the display driver. It handles resource management, such as window creation, handles graphics context management, surface /buffer binding, rendering synchronization, etc.
- *OpenGL ES*: This layer provides standardized abstraction APIs, which can be used by the application.

Application Space

The software packages available for the GPU include OpenGL ES 1.1 and OpenGL ES 2.0; low-level graphics libraries; and tools (Streamliner and Graphics Debugger) provided by ARM. Open GL ES is OS-neutral, providing high portability for applications. Other software components are expected to be available on demand.

Display Controller

The Zynq UltraScale+ MPSoC includes an integrated DisplayPort interface module as part of its available high-speed connectivity peripherals. The DisplayPort interface is located in the PS and can be multiplexed to two of the four dedicated high-speed serial transceivers operating at up to 6Gb/s. This eliminates the need for additional display chips, further reducing system BOM cost.

The DisplayPort interface is based on the VESA DisplayPort Standard Version 1, Revision 2a and provides multiple interfaces that process live audio/video feeds from either the PS or the PL, or store audio/video from memory frame buffers. It simultaneously supports two audio/video pipelines, providing on-the-fly rendering features like alpha blending, chroma resampling, color-space conversion, and audio mixing. The DisplayPort can use one of PS PLLs or the clock from PL to generate the pixel clock.

Integrated Programmable Logic (PL)

In addition to the video codec and graphics processing, there are other important component requirements in any multimedia application, e.g., input and output management for video data as well as the capability to process high-speed video data. The ARM Cortex-A53 cores, along with the memory unit and the many peripherals on the Zynq UltraScale+ MPSoC, play a strong role in managing and capturing the data from many different sources before making it available to the VCU. PS peripherals, e.g., USB and Ethernet, can be used for streaming video devices like a camcorder, Network camera, webcams, etc. Custom logic can be designed in the PL to capture video from live sources. For example, SDI RX, HDMI RX, MIPI CSI IPs (and so forth) can be used to capture raw video from different sources. The VCU can then encode the raw video into AVC or HEVC compressed bitstream format. Similarly, decoded raw video can be transmitted to an external display unit. It can then be displayed using the DisplayPort controller (DP controller) in the PS or by creating a relevant IP using other protocols like HDMI TX, SDI TX MIPI DSI, etc.

Xilinx provides a collection of display interface IPs, available in the Xilinx Vivado IP catalog, that the user can license to implement a display interface in the PL:

- HDMI 1.4/2.0 transmitter subsystem and HDMI 1.4/2.0 receiver subsystem
- MIPI CSI-2 receiver subsystem
- MIPI DSI transmitter subsystem
- UHD-SDI transmitter subsystem and UHD-SDI receiver subsystem

Vision algorithms can be used to gather important information from raw data—e.g., identification of street signs and detection of motion for driver assistance, identification of human faces in surveillance, object identification and motion detection in advanced photography, etc. In addition to collecting data, the algorithms can also be used to process and manipulate raw data in audio/video broadcast and video conferencing use cases. Given the inevitability of even higher resolution video in the coming years, these algorithms need to perform extremely fast. The PL provides the required hardware acceleration for such algorithms, resulting in numerous improvements in performance suitable to next-generation technology.

Any-to-Any Connectivity Using Vivado and SDK Tools

For hardware differentiation, many platform developers use the Programmable Logic for their “any-to-any” connectivity. The Vivado Design Suite helps system designers develop their systems with ease of use by providing the following design suites:

- IP integrator (IPI) tool
 - Uses a block-diagram approach, integrating IPs together to create an overall system
- PS configuration wizard (PCW)
 - Users can configure, enable, or disable PS peripherals, and implement clock and memory configuration, etc.
- IP catalog
 - Using IPI, IPs can be instantiated and integrated into a larger system. In addition, users can integrate their own custom IP repositories to the IP catalog.

The Vivado Design Suite also provides IP simulation, synthesis, place-and-route, configuration of the target device, along with an export-to-hardware feature that provides hardware configuration information to the software development tool. This is used by the Xilinx Software Development Kit (SDK) to generate a boot loader and to pass hardware-specific information to the software during execution phase.

SDK is used for C/C++ embedded software development and debugging. It is built on the Eclipse open-source framework and provides system performance monitoring so that the user can monitor system performance at run time.

For the Linux OS user, Xilinx provides PetaLinux, an embedded Linux SDK. It provides a multi-faceted Linux OS tool flow, enabling a complete configuration, build, and deployment environment.

The Zynq UltraScale+ MPSoC is a perfect solution for teams with limited or no hardware resources who have been challenged in the past due to the RTL (VHDL or Verilog) development expertise needed to realize the full benefits of a programmable device. Recently, Xilinx released the SDSoC™ tool, providing a familiar embedded C/C++/OpenCL application development experience, including an easy-to-use Eclipse IDE and a comprehensive design environment for heterogeneous Zynq UltraScale+ MPSoC deployment. To shorten programming time, the SDSoC tool provides system-level profiling, automated software acceleration, automated system connectivity generation, and libraries. It also enables end-user developers to rapidly define, integrate, and verify system-level solutions and provide their end customers with a customized programming environment. Complex algorithms for video data analytics and processing, written in high-level software, can be accelerated in PL using the SDSoC tool. It shortens the development cycle as well as boosts performance.

Application Examples

The scalable power, high performance, and dedicated engines within the Zynq UltraScale+ MPSoC make it ideal for many applications.

Video Conferencing Application

The Zynq UltraScale+ MPSoC supports high-end video conferencing endpoints. Video conferencing endpoints are terminals used to make point-to-point video-enabled calls. A complete two-way video conferencing system consists of a video camera and a base unit that is interfaced to a video display at each location. Communication between the two base units occurs over an IP network. The video base unit consists of DSPs with integrated on-chip video ports to process audio and video encoding/decoding.

The logic resources and transceivers can be used to connect UHD-4K video cameras to capture and feed the raw video into connected memories. The DSP block resources are used to perform image processing on the captured data, and the processed data is forwarded to the integrated VCU to do compression. The compressed data is packetized using application software running on the APU subsystem and streaming it out via Ethernet to a remote endpoint.

The terminal endpoint can also receive compressed data from the remote endpoint, decode it using VCU, and store it in connected memory. A soft video processing block IP running in the PL can scale down the original raw data captured from a camera and mix it with decoded data, then give it to the display controller, implemented either in the PS or in the PL's soft IP.

The GPU can be used to create an on-screen display (OSD); its output can be blended with video output from the video processing unit (VPU), sending it to the display controller, which displays it on the monitor. See [Figure 8](#) and [Figure 9](#).

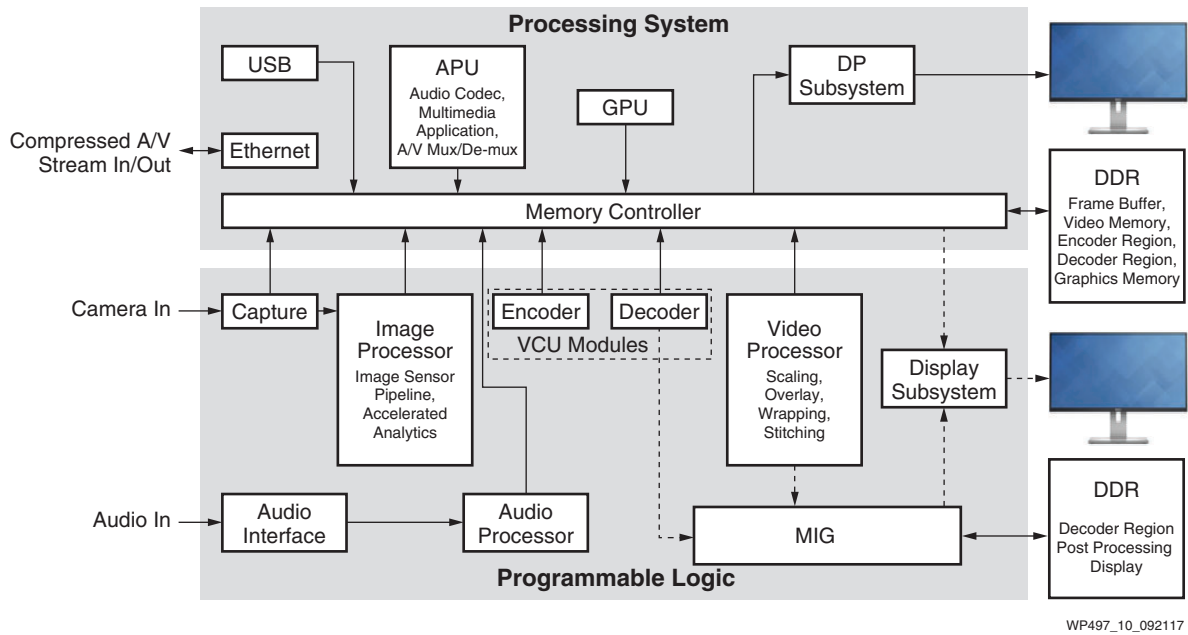


Figure 8: Video Conferencing Application

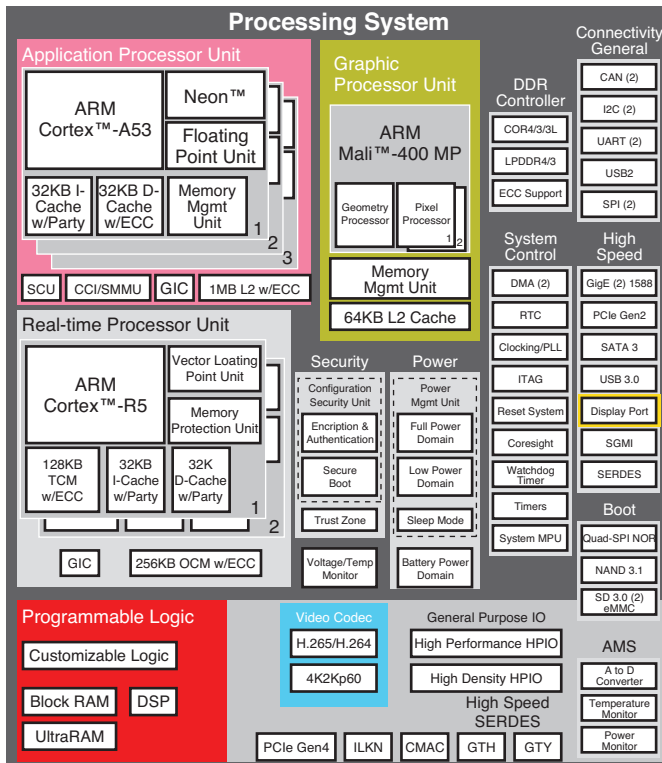


Figure 9: Leveraging the Zynq UltraScale+ MPSoC



Video Transcoding in Data Center and Cloud Computing

Video transcoding is the process that converts a media asset from one format to another, to make videos viewable across different platforms and devices. Most of the time, transcoding a video is performed due to one or more of the circumstances listed below:

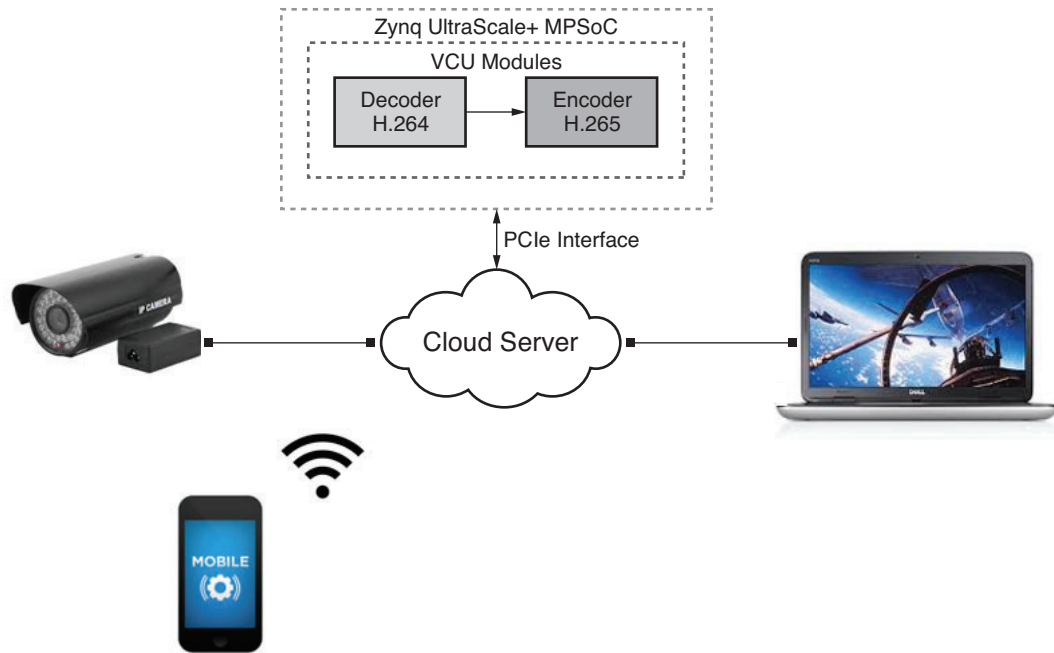
- Target device does not support the format that the original data is in.
- Target device has a reduced capacity, requiring the size of the original file to be reduced.
- Incompatible and obsolete file types must be converted into a modern format that is better supported by the new device.

The process of video transcoding is normally a two-step process. The first part of the process is decoding, whereby the original data is transferred to an uncompressed format. The second part of the process is re-encoding data that can now be transferred to the new device in the desired format.

With the increasing growth of video streaming on the Internet over popular websites such as Netflix and YouTube, and with UHD 4K cameras gaining prominence in the market, a considerable amount of storage and bandwidth is required. The increasing diversity of services, the growing popularity of HD video, and the emergence of beyond-HD formats (e.g., UHD) are creating strong need for coding efficiency superior to the capabilities of existing codecs such as H.264/AVC. The need is even stronger when higher resolution is accompanied by stereo or multi-view capture and display. The H.265/HEVC codec has been designed to support UHD4K and UHD8K, and has tools that take advantage of parallel processing architectures. The HEVC standard brings promise of huge bandwidth savings of approximately 50% over H.264 encoded content with similar quality.

The majority of data centers store videos in different compressed formats and might need to stream the videos based on the receiver's supported format. To support this topology, data centers are required to transcode one video format to other formats based on the target-supported format.

In this kind of topology, the Zynq UltraScale+ MPSoC is a suitable device for high-performance transcoding, as it has support for currently deployed technology (H.264/AVC) and the future generation codec standard (HEVC). It also has the flexibility to program different hardware codecs in the Programmable Logic to support a wide variety of codec standards. The VCU on the Zynq UltraScale+ MPSoC can accelerate the transcoding process if deployed as an accelerator in a data center or cloud computing system. For example, a mobile user records the live video in H.264 format and uploads the video to the cloud server for future playback. If the user wants to play back the same video stored in the cloud on a workstation that supports only HEVC format, the workstation can request that the server transcode the content and send the video in HEVC format. Upon receiving a request from the workstation, the server uses its integrated VCU as a codec accelerator to do the transcoding from H.264 to HEVC and stream the video to the workstation. See [Figure 10](#).



WP497_12_063017

Figure 10: Video Transcoding

Another application of live transcoding is the surveillance camera, which supports H.264 compression but the playback device supports the HEVC codec. In this scenario, the Zynq UltraScale+ MPSoC VCU perfectly matches the requirement because it can receive the H.264 compressed data from a surveillance network IP camera and then transcode it to H.265 using the VCU block, streaming it to the target playback device for display.

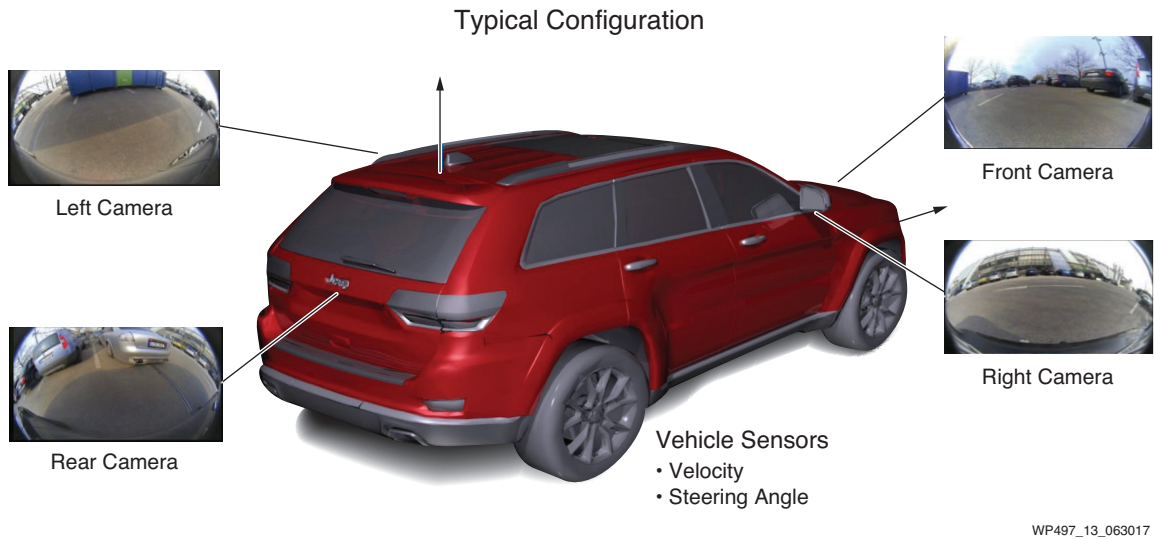
Automobile Surround Viewing System

Surround view in cars has changed the driving experience, and the GPU is the major factor in cultivating that technology. The surround view camera system is an emerging automotive Advanced Driver Assistance System (ADAS) technology that assists the driver in safely parking the vehicle by allowing driver to see a top-down 360 degree view of the vehicle’s surroundings.

Real-time rendering for the surround vision image-based solution uses the GPU to reconstruct full visual odometry using virtual cameras, wherein the virtual camera parameters are cloned from real camera and are implemented on GPUs to generate an image of the virtual scene with parameters changing in real-time. The Zynq UltraScale+ MPSoC is well suited for this domain because of extremely efficient fixed function units in GPU.

GPU horsepower is used to do the synthesis for illustrating a composite view, generating an output pixel. The generated pixel can be either a combination of two pixels (if the output is in an overlapping region) or a single pixel (if the output is in a non-overlapping region), which is fetched from the input frames using a geometric look-up table (LUT). Each entry in the geometric LUT specifies the camera ID and coordinates of the input frame to generate the output pixel at the current location. A mesh table procedure can be followed to generate the output frame. Given a specific output resolution, each output location can be back-mapped to input images via projective and lens transformation to a location in the input image. The mesh table consists of 3D world

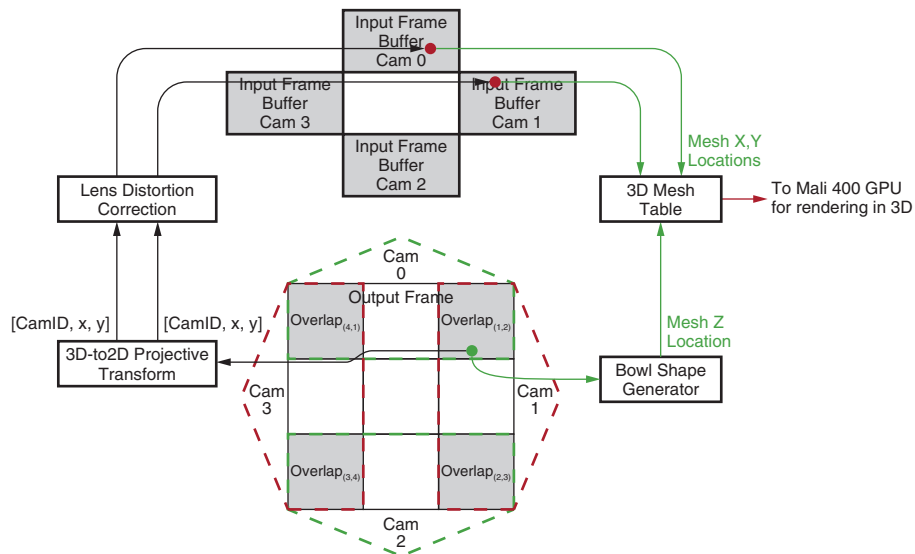
coordinates for locations in the vehicle's surroundings, and the associated input locations for texture mapping from adjacent cameras viewing the scene for a given location. The output is represented as a bowl, whose height varies as a function of the distance from the center of the vehicle. The collection of the mesh table, which includes output mesh and the associated texture mapping, is passed on to the graphics processor for further rendering. Along with the mesh table, the application can also generate a blending LUT, which encodes the weights for linear combination of the image intensity information received at each location. Using GL_OES_EGL_image_external extension, camera YUV images are passed on to the GPU as textures for rendering. See [Figure 11](#).



WP497_13_063017

Figure 11: Surround View

Surround vision systems receive input video streams from four fish-eye cameras and create a composite surround view. This system creates the stitched output image using the mapping encoded in the geometric LUT. See [Figure 12](#).



WP497_14_101317

Figure 12: Synthesizing a Composite View

The Zynq UltraScale+ MPSoC Advantage

The flexibility of Zynq UltraScale+ MPSoC allows accelerating compute intensive applications to share the workloads between the GPU, CPU, and PL, where the complex arithmetic calculations can be offloaded for hardware acceleration in the PL and OpenGL Shading Language (GLSL) uniforms can be pre-calculated on the APU. Calculations on the GPU shader core make sense only for values that vary between vertices and fragments. Everything that is constant among a whole batch of vertices is most efficiently dealt with on the CPU.

There are certain tasks where a CPU can easily outperform a GPU, even for very large data sets. Alternatively, there is programmable logic on the platform that can provide the hardware acceleration for the OpenGL matrix and lightning calculations for complex 3D models. The graphics system generates images through a pipelined sequence of operations, where the slowest stage is often called the pipeline bottleneck. A single graphics primitive, e.g., a triangle, has a single graphic pipeline bottleneck. However, the bottleneck might change when rendering a graphics frame that contains multiple primitives. For example, if the application first renders a group of lines and then a group of lit and shaded triangles, the bottleneck will likely change. Because some of the pipeline stages are executed on the CPU, other stages are executed on the GPU—and the bottlenecks keep changing. The Zynq UltraScale+ MPSoC, which encapsulates multiple processing engines and the PL, helps the GPU accelerate both 2D and 3D graphics by providing flexibility to perform calculated work distribution between the GPU, PL, and CPU, removing the performance bottlenecks and increasing the overall throughput.

Conclusion

The Zynq UltraScale+ MPSoC is the true heterogeneous multi-processor SoC platform—“true” as it is much more than just a high-performance 64-bit processor coupled with programmable logic. It has dedicated engines for real-time processing, graphics processing, and video encoding and decoding. This allows the user to choose where algorithms are implemented, thus optimizing system performance and power. The Zynq UltraScale+ MPSoC was created with complete flexibility at both the software and hardware levels, while integrating the most diverse number of specialized engines the embedded market has ever seen.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
10/23/2017	1.0	Initial Xilinx release.

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>.

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.