How Performance and Programmability can Power Video Mixer Display Applications

Read About

Mission Display Computer

Video Mixer

Sensor Input Processing

Graphical Processing Units (GPUs)

Intel Integrated GPUs

Introduction

Mission display computers play an important role in an increasing number of imaging applications such as digital moving maps, 360° situational awareness, persistent surveillance, embedded training, and degraded visual environments. These systems are designed with advanced graphics capability to drive multiple displays independently with video from multiple sources. The display video can be a combination of input from multiple sensors, generated digital map video, symbology, and metadata information from a variety of other sources. All of this information is combined and overlaid to provide multiple operators each an instantaneous independent view of the battlefield.

Requirements of a mission display application are a large number of ports to handle the number of video sources, high performance and bandwidth to handle the increasing resolution of sensors and displays, low latency for immediate feedback, and flexibility to handle the variety of sensor and display interface formats. The video mixer is the heart of the mission display computer and generally required an FPGA to combine the multiple inputs into multiple display outputs. The issue with this approach was the difficulties inherent to FPGA programmability. FPGA-based video mixers were generally uniquely tailored to each platform, lacked floating-point arithmetic precision, were difficult to modify to deal with changes in video interfaces or processing requirements, and required lengthy development schedules.

Figure 1: Surveillance Display System

Info

curtisswrightds.com

Email

ds@curtisswright.com

Advances in embedded graphics technology make a product such as the commercial GPUs, with its higher performance and bandwidth, a feasible and attractive alternative to replace the FPGA as the video mixer for some mission display applications. This is further augmented by the video processing and generation capabilities of Intel's built-in integrated graphics in their multi-core CPUs.

This white paper will discuss how the combination of high performance standalone and integrated GPUs combine with FPGA interface flexibility to provide increased programmability and versatility to the video capabilities required of a mission display computer.

Mission Display Computer Challenge

Mission display computers include multiple sensor inputs, multiple internally generated video sources, and multiple independent video displays. Characteristics of heads up versus heads down displays, high resolution versus legacy displays, and the variety of display monitors and imaging sensors make building a mission display computer applicable to multiple systems a challenging problem. Versatility and programmability for a variety of image display platforms as well as changing programmatic requirements are an important feature of a mission display architecture.

The video mixer in these systems require real-time pixel-level performance, bandwidth, and interface flexibility. These requirements have driven the design of yesterday's mission display systems to FPGA-based video mixing solutions. Fully capable mission display applications have been developed in the past based around FPGAs that provided the requisite performance, bandwidth, and interface flexibility.

What has been lacking is the programmability to deal with evolving program requirements. Consider the incorporation of an additional video source in a FPGA-based architecture. The FPGA firmware must now deal with an N+1 to M display interconnect with the design modified to deal with the additional input and its effects on the video mix capabilities. Logic timing and functionality must be reverified over the range of command and video functionality. FPGA HDL programming is not simple and the development schedule is often lengthy. With FPGA performance and flexibility, it may be possible to handle the changes, but advances in GPU capabilities can provide a much more versatile and programmable multiple video display solution with reduced development cost and schedule.

GPU Video Mixer

GPUs are increasingly being found in embedded systems. They are a natural fit for graphics processing acceleration and image rendering for high-resolution sensors and displays. Their size, weight, and power (SWaP) performance measured in floating point operations has significantly increased to make them viable for a number of defense applications.

Standalone GPUs were previously limited in embedded systems because of their inability to meet the power and thermal requirements of embedded platforms and logistical concerns in supporting the extended timelines of defense applications. These limitations have been reduced by continued advancements in both memory and processing with power reductions, their packaging in the common MXM mezzanine format, and the support of graphics suppliers extending device lifecycles for embedded applications.

The advantages for GPU video mix programmability are substantial. The development of custom graphic designs can be costly in terms of dollars, schedule, and risk. The result is often a proprietary solution that needs to be redone for evolving requirements or new applications. GPU programmability using industry standard OpenGL provides the developer with an easier platform to build a graphics system, to adapt it for requirement changes, and to exploit COTS graphic roadmap product development to handle next generation embedded video applications.

The latest GPU products now being incorporated into VPX modules is based on the NVIDIA Quadro Pascal P5000, an appreciable advancement in commercial GPU technology. The P5000's 2048 cores provide 6.2 TFLOPs of single precision compute performance. Its 16 Gigabytes of GDDR5 256-bit graphics memory provide 192 Gigabytes per second memory bandwidth. Its PCIe interface is up to sixteen-lane Gen3 supporting up to 128 GB/s data transfer performance. It supports four high-resolution 4K DisplayPort 1.4 digital video outputs at 120 Hz with 10-bit color depth.

FPGA Video Interface Flexibility

Embedded video systems used in both tactical and surveillance applications must deal with sensor systems in the field that can stream high bandwidth video data in a variety of formats. Hand-in-hand with the GPU video performance is the need to provide the processing engine the input video and then distribute the output.

The Pascal GPU can provide the processing and displays, but it needs to be augmented with FPGA technology to interface to the video inputs. In previous systems, the FPGA contained all the logic for interfaces and video display processing. With the GPU now performing the video display processing, the FPGA processing is simplified to sending captured video to the GPU and reformatting its output to the displays.

The FPGA I/O capabilities are useful in converting a variety of video inputs and routing them to the GPU memory via DMA transfers over the high bandwidth PCIe interface. Likewise, its versatility is also invaluable for being able to output in a variety of video formats. The FPGA can directly interface with GPU video outputs to convert to the required format.

The basis for the FPGA video interface capabilities is configurable and modular IP logic building blocks for the different video formats and the communications interface to the GPU. The IP facilitates development of image I/O tailored to program specific image capture and display requirements.

The Xilinx® Kintex®-7 has been the FPGA choice because of its performance per watt and I/O flexibility in handling multiple video interfaces. The Kintex-7 family provides up to 478K logic cells, 1920 DSP slices capable of 2,845 GMACs, 270 block RAMs with 4.8MBytes total, 32 transceivers to @12.5 Gbps, 500 I/O pins, and has hard IP or soft IP support for a single PCIe Gen2 or Gen3 fabric interface respectively. It can accept multiple digital video inputs and produce multiple video outputs in a variety of interface formats simultaneously.

Integrated GPU Video Processing

Avionic and other defense applications are not driven just by the display of sensor data. There is often the simultaneous requirement to generate video and symbology to operators from a variety of sources. Digital maps and other pre-flight information are displayed to provide operators with position and mission information. Processed and reformatted sensor input provide for surveillance and targeting information. Cues from electronic warfare (EW), target acquisition, and operator actions are another source of inputs.

Intel CPUs with integrated GPUs provide both multi-core and graphics video processing in energy efficient products. The multiple cores, high bandwidth memory, integrated GPU graphics, and fabric and other interfaces combine to make it a powerful processing package.

An integrated graphics processor is an ideal solution for embedded applications video graphics. It combines the processing intelligence and interface capabilities of a single board computer with accelerated floating point and image rendering GPU graphic capabilities. Its floating-point precision allows for the generation of high-resolution video. Application processing can be optimally allotted to either the CPU for decision-making or the GPU for stream processing without the complexity and latency required of off-chip interfaces. The integrated GPU can route multiple DVI and DisplayPort outputs for display.

The capability to drive multiple digital video outputs is particularly compelling for a mission display computer. Generation of video can be allotted to independent processors for the variety of system video generation requirements. This simplifies the development and eases the reuse porting of video applications. For example, one SBC could be tasked with digital moving map processing, another embedded training, a third processing and reformatting an external sensor input, a fourth generating target or mission specific information, etc. The SBCs can be augmented with terabytes of removable external solid state disk (SSD) Flash to be utilized in the video generation process. A video frame grabber XMC provides the capability to receive input for an SBC performing sensor processing.

The 7th Generation Intel “Kaby Lake” Xeon is a multi-core CPU with integrated graphics. An Intel Xeon E3-1505L v6 processor provides four cores (8 threaded) operating at 2.2 GHz (turbo to 3.0 GHz) with 282 GFLOPs AVX2 vector performance. It supports a SATA interface to Flash SSD and a dual channel memory controller to DDR4 ECC SDRAM at 2,400 MT/sec transfer rate. It has 8-lane PCIe Gen3 interfaces (configurable as 8/4/2-lanes) with NTB capabilities for interface to a local XMC and the backplane. It is paired with a mobile Platform Controller Hub (PCH) that provides up to four Gigabit Ethernet ports and USB, SATA, and serial external interfaces.

The CPU also includes a tightly coupled Intel HD Graphics P630 24-core GPU that delivers a new generation of integrated GPU graphics with total peak performance up to 384 GFLOPs. The GPU can drive three independent high-resolution DVI or DisplayPort video outputs simultaneously. With DisplayPort interfaces, it is possible to reduce the 4-lane display interface down to a 2-lane or 1-lane interface in order to reduce wiring interconnect complexity.

Example Next Generation Mission Display Computer System

The mission display computer shown in Figure 2 includes four Curtiss-Wright VPX3-1220 single board computer modules with a single standalone VPX3U-P5000-SDI-8IO GPU graphics processor from Wolf Advanced Technology. The system is packaged in an open standard 5-slot 3U OpenVPX chassis.

The VPX3U-P5000-SDI-8IO graphics processor is powered by the NVIDIA Quadro Pascal P5000 combined with two Xilinx Kintex-7 FPGAs. The module can accept up to eight 1-lane DisplayPort video inputs as well as analog inputs. On the output side, it can support up to four 3G-SDI video outputs while also providing analog outputs. It is capable of independently reformatting each display output to alternative video formats. The SBC CPUs are the Intel "Kaby Lake" Xeon quad-core CPU with 16 GBytes of DDR4 and 32 GBytes of SSD Flash with an integrated HD Graphics P630 GPU.

Each SBC hosts an XMC site. Including the two slots on the Curtiss-Wright PSU3-THOR power supply that also provides a PCIe switch and hosts an XMC site, there are six total XMC slots to enhance system flexibility available for multi-IO, SSD Flash, video input, GigE switch, or other capabilities. For applications that required additional or mission-specific data storage, the SBC can host a mezzanine SSD Flash XMC-554. The XMC cards shown in Figure 2 are example options.

The four SBCs each drive two generated single-lane DisplayPort video inputs to the GPU utilizing their graphics processing and multiple video outputs capabilities. Each SBC is responsible for a specific dedicated video processing function. The general-purpose performance of the Intel CPU and its video graphics capabilities make it an optimal choice for this task.

Instead of generating video, the VPX3-1220 SBC can be utilized for processing and display of live sensor input. It can host an available COTS video frame grabber XMC to provide incoming sensor data. The input data stream would be processed by the SBC for target track, acquisition, recognition, and other applications. The processed video would be reformatted with overlaid target symbology for output to the GPU video mixer.

The eight single-lane DisplayPort video outputs generated by the four SBCs of the mission display computer go to the GPU that functions as an eight by four crossbar. The eight inputs are combined and reformatted into four video outputs. The GPU is a command driven resource with instantaneous capability (at the video framing level) to combine any of these inputs into any of the four video outputs.

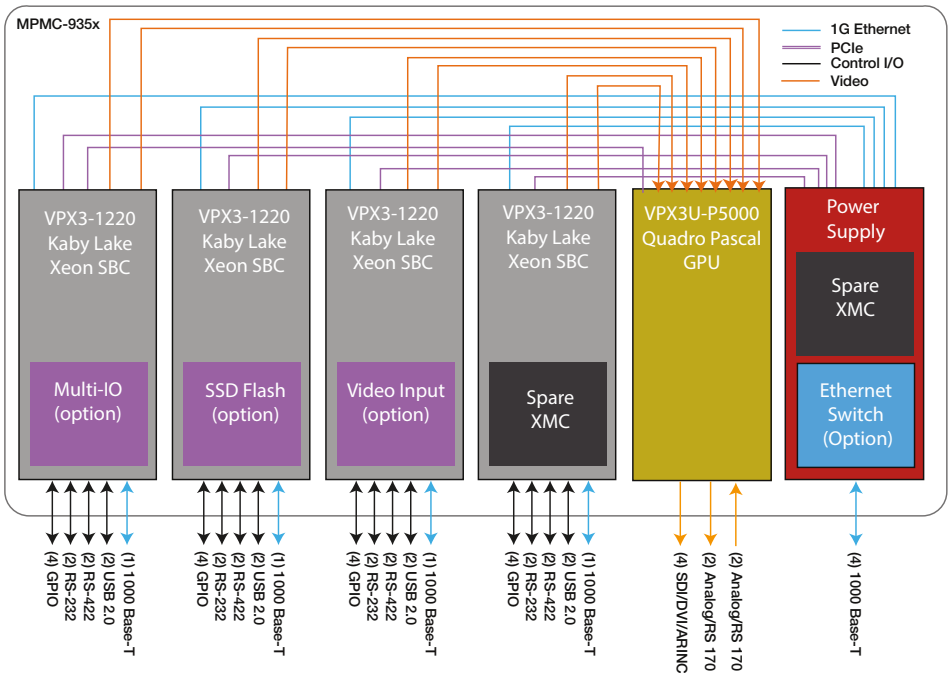

Figure 2: Mission Display Computer

The video mixer application, which simply combines video sources, does not fully utilize all the processing capabilities of a GPU. Video mixer commanding could be as simple as parameters providing the details of the input window, intensity scaling, rotation, zoom, and output window location for each of the eight video sources overlaid on one of the four outputs. Whatever the command paradigm, it should be possible to program a versatile video mixer portable for a number of mission display applications and adaptable with only minor modifications or command level tweaking to changes in video and display requirements.

The two FPGAs on the video graphics processor also bring another dimension of adaptability to the system. The FPGAs allow the flexibility of input in a variety of interface formats. Likewise, for video output, the GPU default DisplayPort could be output or the FPGA could transcode output to 3G-SDI, ARINC-818, or other legacy interface standards to meet program requirements.

GPU Video Mixer Performance

How does that performance relate to the requirements of a mission display computer? For this sizing exercise, assume a video display system with eight video inputs and four video outputs.

The four Xeon integrated graphic SBCs each provide two video inputs though there is capability for the video mixer to accept external sensor input. The video inputs from the SBCs or external sensors are in single-lane DisplayPort format with 1920 x 1080 color resolution at a 60 Hz frame rate.

The two Kintex-7 FPGAs have no issues handling four inputs each. The bottleneck though is the PCIe input bandwidth to the GPU. The P5000 GPU can support up to a 16-lane PCIe Gen3 interface, but each Kintex-7 FPGA is only providing four lanes of Gen2 PCIe. The loading on the PCIe interface is 70% for this example.

Memory storage requirements for this system with input, output, and processing buffers is about a gigabyte and a half of storage, about 10% of the P5000 16 GB GDDR5 memory capabilities respectively. For processing performance, a very conservative 500 operations per pixels for each of the four DisplayPort video outputs would require about one teraflop (15% of a Quadro Pascal P5000 peak).

Memory bandwidth sizing is application independent. The GPU like all parallel devices optimally does a large number of arithmetic operations per pixel per memory access. An optimal number for the P5000 GPU based on its TFLOPs to memory bandwidth is 120 operations per pixel memory access. With the worst case loading of 15%, memory bandwidth issues would only occur in this case if there were less than 20 operations per every pixel memory access.

The video is output directly from the GPU to the FPGA over its four display ports. The Kintex-7 has no issue handling the high resolution (4x the video input) and reformatting the outputs to the four displays.

Conclusion

The number of video inputs and outputs, the processing required for high resolution and high bandwidth sensors and displays, and system latency requirements getting the information to the operator have driven designs to custom FPGA-based graphic systems. The drawback to these FPGA-based architectures has been cost, schedule, and adaptability to requirement changes with many of these solutions custom to the requirements of a single program. This paper offers a COTS alternative for the mission display computer that reduces system development cost and risk by replacing the custom FPGA module with an embedded COTS GPU, paired with an interface FPGA.

Advances in both integrated and standalone GPU capabilities combined with FPGA flexibility can provide an attractive Mission Display Computer architecture. A small, lightweight 3U VPX open architecture system can be designed capable of providing a COTS Mission Display Computer adaptable to a variety of image display applications. The GPU replaces the FPGA for the video mixer functionality. The system is inherently more adaptable to changing requirements with a GPU video mixer using OpenGL programmability. Additionally the GPU can provide much more complex image processing such as enhancement, stabilization, data analysis and visualization, and other applications. The FPGA is still incorporated in this design for its flexibility for video interfaces to and from the GPU graphics engine.

The advanced graphics capabilities of Intel multi-core CPUs with integrated GPUs provides dedicated independent image processing and video generation. The multi-core integrated graphics augmented with SSD Flash can support digital

Author

Chad Augustine,

Product Manager, Integrated Systems

Curtiss-Wright Defense Solutions

Robert Cody,

Embedded Systems Architect

Curtiss-Wright Defense Solutions

moving maps, embedded training, and other video generation or used for digital video record. The XMC site can also be used for the input of live video that can be processed and reformatted for display by the video mixer.

The analysis of the mission display application with the GPU graphics processor showed the current system adequate for the video processing it was designed to support. A more optimal PCIe interface to the GPU capabilities could provide up to four times the DisplayPort inputs or the same number of higher resolution 2-lane or 4-lane DisplayPort interfaces. The GPU processing capabilities were not being fully utilized for simple video mixing functionality. Its processing capabilities were not challenged in this sizing scenario, which was not surprising as the Pascal GPU TFLOPs numbers are higher than the GMACs of the highest performance Xilinx Virtex-7 FPGA.

The embedded GPU video processor will work for this mission display application and for others of this scale or greater as technology advances. Its current limit of eight digital video inputs does not allow it to scale to the higher crossbar connectivity of FPGA-based video mixers. Upgrading the two FPGAs to a Xilinx UltraScale FPGA would provide higher bandwidth PCIe Gen3 interfaces alleviating the I/O bottleneck.

The mission display computer reviewed in this paper is an open architecture 3U OpenVPX product. Its programmability makes it adaptable for a variety of applications. It is able to keep up with evolving video capabilities based on the continued COTS technology advancements of NVIDIA, Xilinx, Intel and other GPU, FPGA, and CPU suppliers. Obsolescence is mitigated by the incorporation of industry standard MXM mezzanine modules and the increased longevity support GPU silicon vendors are providing for embedded markets. A programmable architecture with the capability to mix eight video sources and drive video output to four independent displays is an attractive alternative for a large number of mission display applications.

Learn more

Product: Embedded Video Processing System

Product: VPX3-1220 SBC

Product: VPX3U P5000-SDI-8IO GPU

White Paper: GPUs in Defense: Performance, Productivity, and Longevity

White Paper: Tackling Defense Apps by Harnessing the Intel Core i7's Integrated GPU

Case Study: 3U VPX High Performance Display Processing Platform