White Paper

FPGA
Live Video Production

FPGAs for Live Video Production Workflows

Authors

Francisco Perez
Platform Architect
Intel Corporation

Table of Contents

Introduction 1
Challenges 2
Solutions 2
ST2110 implementations in FPGA 2
FPGAs for Video/Audio Processing 3
FPGAs for Camera Applications 4
Solution Ingredients and Roadmap 4
Intel Agilex® 7 and Intel Agilex® 5 FPGAs Enable 8K Video Data Paths 4
VVP as Building Blocks for Video Processing 5
Using Open FPGA Stack (OFS) and Intel® FPGA SmartNIC N6000-PL Platform to Expedite Virtualized FPGA Developments 5
Summary 6
Where to Get More Information . . 6

Introduction

The transition from dedicated serial digital interface (SDI) point-to-point connections to Ethernet IP networks for media content exchange and collaboration continues to be a highly disruptive evolution across the broadcast-and-media production supply chain. IP technology has evolved and matured for file-based production and progressed to include live production, timing, and synchronization.

There are many challenges to overcome to meet the demands of live video production including—but not limited to—interoperability, performance, latency, and connectivity.

The move from proprietary hardware to computer-based commercial off-the-shelf (COTS) systems and devices is maintaining a fast pace. IP media technology is all about networks, servers, storage, and applications.

Live IP video production has introduced new network designs and topologies to meet the mission critical requirements of live media transport over a network. These new requirements are pushing the limits of the current paradigm for COTS hardware architecture, with increased demands for media processing, which necessitates the use of specific hardware accelerators.

Intel, as the Official Processor Partner of the Olympic and Paralympic Games has partnered with the Olympic Broadcasting Services (OBS) to define a reference architecture to meet these new challenges and enable live video production workflows, leveraging diverse technologies across multiple silicon platforms and complete software stacks.

We are showcasing the first trial of this novel architecture in **The Olympic Winter Games Beijing 2022.** We also have a roadmap to expand the reach and functionality of this novel technology over time and at future events as well.

"The ongoing collaboration between Intel and the Olympic Broadcasting Services defines a reference architecture for a software-defined outside broadcast van under the working name the "Virtualized Outside Broadcast Van (vOB)."

The goal is a fully virtualized, software-based architecture based on a common base platform using COTS hardware and retains the user experience familiar to broadcast engineers and operators using traditional broadcast appliances.

The standards-based platform enables multiple software applications from one or more vendors to be deployed on the same physical platform. This approach facilitates simple scalability of physical hardware resources to match the complexity and compute requirements for various broadcast events.

Piloted at the curling event during **The Olympic Winter Games 2022**, the reference architecture mirrored the standard broadcast production made available to international broadcasters.

If you are involved in architecting, designing, or selecting your next media processing infrastructure, read on to find out how Intel® FPGA can be used to implement video and audio processing applications to meet the demanding real-time requirements for complex live video productions.

Challenges

Traditional broadcast infrastructure built with SDI baseband routers, coaxial cables, and BNC connectors is transitioning to Ethernet using IP network switches and software-defined networking (SDN) controllers. The increasing speeds supported by Ethernet switches and the corresponding improvements in aggregate non-blocking switch and network throughput has paved the way for the use of IP/Ethernet to transport critical broadcast applications in a cost-effective manner. This approach still delivers the same robustness and stability of legacy SDI operations–all with greatly increased agility, flexibility, and scalability to meet ever-evolving media formats.

To meet this challenge, a collaborative partnership of many organizations was formed combining a selection of standards, specifications, protocols, and recommendations from the media, internet, and IT communities. As a result, there are a number of new standards that have been approved and ratified for live production. The SMPTE ST-2110 Professional Over Managed IP Networks Suite of Standards and the AMWA NMOS group of specifications are the main ones.

In a ST2110 installation, all the media essences (video, audio, data, control, and sync) are transmitted as packets within the network.

Although this seems to be structurally similar to a regular Ethernet network, we should note a couple of unique requirements for a successful ST2110 working deployment: traffic shaping and hitless protection switching.

Associated with the flow is the requirement to manage, process, and convert several simultaneous high-resolution video and high-quality audio streams, as part of the real-time production process. To complete this task, a high-performance platform is required to provide enough compute power and bandwidth to produce results at the lowest possible latency.

To accomplish the task mentioned in the previous paragraph, flexible media accelerators can be added to general purpose COTS servers to improve processing capabilities while lowering power requirements. These media accelerators are a perfect fit to be implemented by Intel FPGAs as they can offer massive parallel processing and memory bandwidth. Adding flexible media hardware accelerators to your COTS server improves deterministic latency with independence on the software load conditions. This approach also saves precious computing resources that are available for your user applications.

As FPGAs are in-field reprogrammable within milliseconds, different audio/video processing pipelines can be tailored and deployed according to the use case, making the FPGA a perfect solution to add flexible media acceleration to general-purpose compute servers.

Solutions

ST2110 implementations in FPGA

SDI to IP Gateways

While it may be possible to move to an "All-IP" based infrastructure, most systems are required to support a variety of legacy SDI-based devices that are not easily converted to operate in the network world. The challenge is to find the most efficient way to bridge the gap between the two worlds while maintaining the strengths and advantages of both.

SDI IP gateways play an important role in a hybrid SDI or IP studio production systems to transport the audio, video, and ancillary data across the SDI and IP islands. They provide aggregation of one or more essence streams into a 10 GbE, 25 GbE, or even higher bandwidth network segment. For example, 50 GbE and 100 GbE. They also enable signal buffering to ensure proper time alignment and provide clean transitions between IP streams. They may also include mezzanine or intra-frame codecs for network bandwidth savings.

SDI IP gateways provide the necessary capability to connect distributed studios or distributed facilities. They provide the ability to span across different PTP time domains. The IP packetization of SDI or separate audio, video streams, and ancillary data is based on SMPTE ST 2110 standard.

A conceptual block diagram of the SDI IP gateway system is shown in Figure 1.

ST2110 NIC

As we move to all-IP deployments, new requirements to support real-time media transport have arisen. There is a need for video-aware architected COTS ST2110 network interface cards that can deliver the 24/7 reliability and interoperability required of broadcast equipment while still offering the flexibility and scalability expected of IP infrastructures, including standard network-stack functionality.

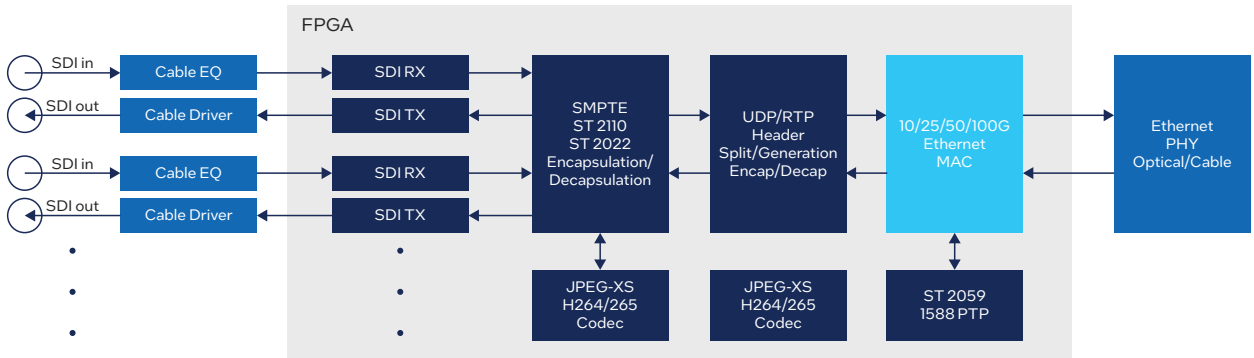

Figure 1. Multichannel SDI to ST2110 gateway implemented in a single FPGA.

These specific network adapters offer complete ST2110 offload allowing OEM software solutions to provide reliable and scalable COTS-based 10, 25, and 100 GbE IP solutions. These cards actually offload all the required packet processing for ST2110 media transport including audio, video, and ancillary data, as well as redundancy and packet pacing built into the network interface card (NIC) to free up system resources. Software solutions can provide guaranteed narrow senders under all conditions independent of CPU load and operating system (OS) interactions.

Figure 2 is a simplified block diagram of a ST2110 NIC based on an FPGA. Different options offer capacity to handle up to 8x 4k60 or 32x 1080p60 streams IN and OUT, for up to 100 Gbps link with almost zero CPU usage.

As these ST2110 NIC cards will be typically deployed in servers supporting virtualized flows—where the end applications would be running within virtual machines and/or containers—SR-IOV for VF support is a mandatory requirement. Having the option to implement direct memory access (DMA) channels in hard silicon, instead of soft logic, attached to the SR-IOV VFs can bring several benefits in terms of logic and power reduction for those applications.

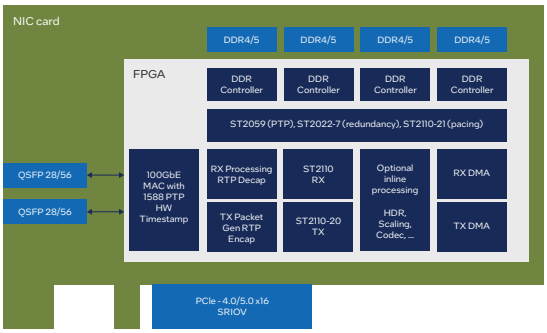

Figure 2. FPGA-based network interface card with full offloading of ST 2110 protocol from the host server CPU.

FPGAs for Video/Audio Processing

Video mixer, replay, multiviewer, scaler, HDR processing, and more

System integrators often need to bridge between disparate protocols or convert video formats between equipment when building solutions for broadcast facilities or pulling together various A/V products for live events. Intel offers a broad portfolio of intellectual property (IP), reference designs, and hardware to accelerate time to market. Intel also offer performance, differentiation, and integration capabilities that are not feasible with ASSPs or graphics processing units (GPUs).

Up/down/cross converters provide the capabilities for video processing and scaling, changing formats between Standard Definition (SD), High Definition (HD), and 4K in different color spaces and aspect ratios. These converters can be integrated with various connectivity interface standards such as SDI, High Definition Multimedia Interface (HDMI), DisplayPort, and IP. In general, video processing cards encompass a large variety of algorithms. However, these cards have multichannel video processing pipelines integrated with connectivity interfaces and often provide PCIe streaming DMA access host processing in production environments.

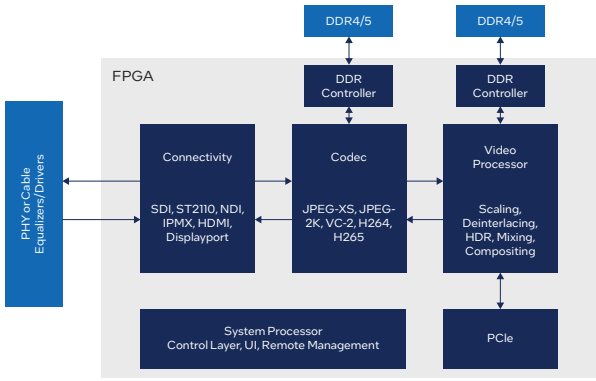

Figure 3. FPGAs allow integration of complete connectivity and video processing pipelines.

Multichannel Audio Processing

Intel FPGAs are unrivaled in providing the digital signal processing (DSP) performance required for audio processing, interfacing, compression, embedding, and conversion. The inherent parallelism of FPGA architectures means that many channels of audio can be processed together using very efficient resources. The high performance of Intel FPGAs means that audio far surpasses the needs of most applications.

Intel offers architectures as cost-effective and flexible alternatives to ASSPs and DSPs, along with partner ecosystem providing audio and speech codecs, echo cancelers, and more. Intel platforms enable engineers to integrate audio functionality into their products rapidly, offering differentiation, channel density, and fast time to market. Designers can take full advantage of the DSP performance, bandwidth, and features of Intel FPGAs to implement system-on-chip designs that eliminate the need for separate components to perform audio processing tasks, thereby reducing costs, particularly for multi-channel audio applications.

It is important to note that FPGAs can hold together connectivity and processing. Therefore, a given audio or video application can benefit from integrating media transport connectivity (like ST2110) to build a compact and reliable solution efficiently.

Figure 4 shows an audio processing card that can integrate several I/O interfacing capabilities for an end-to-end audio solution in a single device.

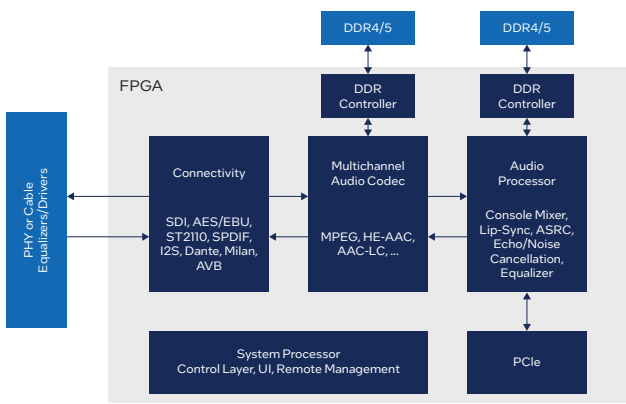

Figure 4. FPGAs allow integration of complete connectivity and audio processing pipelines.

Low-Latency Mezzanine Compression Technologies for Live Production

Captured resolutions and video quality are rising. Besides, the sheer amount of raw data that is produced, collected, and transported is becoming hard to manage.

Transporting compressed video is becoming more popular as the market continues to see increasing demand for higher resolutions, faster frame rates, and more intelligent distribution systems.

Now, more than ever, there is a need for a standardized compression solutions enabling designers to manage more pixels, save cost and power, simplify connectivity, and preserve quality all with low latency and implementation complexity.

FPGA architecture responds to these challenges as it is well-suited for mezzanine codec implementation (JPEG-XS, VC-2, SpeedHQ, and so on), leveraging parallel processing, and internal memory for very low latency and compact solutions without requiring external frame buffering.

FPGAs for Camera Applications

As imaging resolutions and frame rates increase beyond HD and 60 frames per second and require wider dynamic range support, it is essential to perform real-time image and video processing within the space-constrained designs of cameras and camcorders. In addition, new capabilities are required such as video analytics and metadata acquisition to enable efficient workflows and monetize content. The traditional solution is a mix of ASICs, processors, and FPGAs, which complicate the system and drive up power consumption and heat.

Intel FPGAs integrate high-speed connectivity from the latest 4K and 8K sensors with a fully flexible image and video processing pipeline, add lossless or lossy encoding capabilities where needed (typically lossless for studio environments, or lossy for wireless transmission in ENG cameras, or for local storage on SD cards or SATA) and offer support for a variety of output connectivity standards. All these features are now supported on a single device, offering significant space, cost, and power savings.

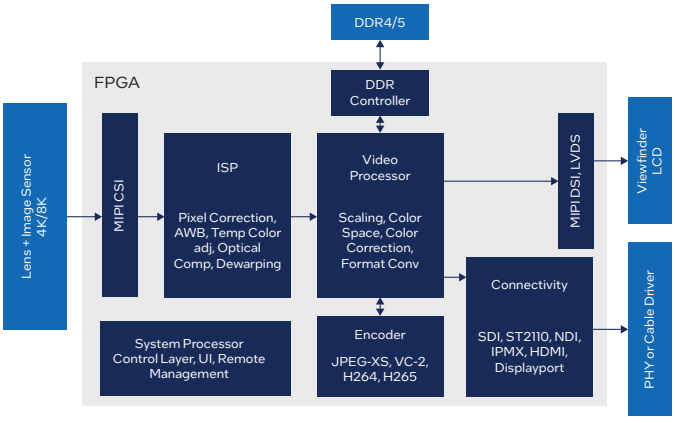

Figure 5. A full-featured high-end camera can be efficiently implemented in a single FPGA.

Equipped with 4K or even 8K image sensors, studio cameras capture and deliver high definition, High Frame Rate (HFR), High Dynamic Range (HDR) and Wide Color Gamut (WCG) video through high-bandwidth cable (12G-SDI, 25/100GbE cable, or optical fiber) for studio broadcast and sports events. A typical studio-camera system head is shown in Figure 5.

Solution Ingredients and Roadmap

Intel Agilex® 7 and Intel Agilex® 5 FPGAs Enable 8K Video Data Paths

Video resolutions have increased over the decades from SD (720x486), through HD (1920x1080) and UHD 4K (3840x2160) to 8K (7680x4320) and beyond. The clock frequency required to handle this increasing bandwidth has likewise increased. The "pixel clock" for SD resolution video was a mere 27 MHz; easily accomplished today but challenging at its introduction in the early 90s. HD video resolutions required clock frequencies of 74.25 MHz or 148.5 MHz, which again were challenging but achievable for their era. Today "4K" resolution requires a pixel clock rate of 594 MHz, a requirement that is now achievable in our latest FPGA family, while "8K" needs 2,376 MHz. These very high clock rates forced a different approach from video engineers.

To cope with this limitation, video IP cores such as scalers or color space converters, were redesigned to process multiple pixels on each clock cycle. In majority of cases, this meant duplicating the entire video pipeline within the IP core. Moving from 1 pixel-in-parallel to 2 pixels-in-parallel (PIP) for "4K" video could result in a doubling of FPGA resources used. The current early adopters of "8K" video designs often rely on a similar technique of processing 8 pixels-in-parallel, with predictable increases in FPGA utilization.

Intel Agilex 7 and Intel Agilex 5 FPGAs, the new family of 10 nm Intel FPGA, is designed to operate at higher frequencies than previous families. This enables FPGA developers to minimize resource usage and power for a given logic function. The ability to reach 600 MHz, often without requiring extensive rewriting of existing Register Transfer Level (RTL), is of particular interest to video designs as it enables "4K" video at 60 frames-per- second to be processed as 1 pixel-in-parallel (PIP).

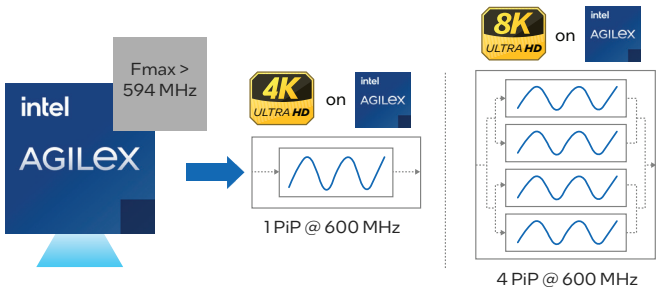

Figure 6. Intel Agilex 7 and Intel Agilex 5 FPGA high-performant fabric allows video processing execution at high clock frequencies, allowing a very optimized footprint saving resources and power consumption.

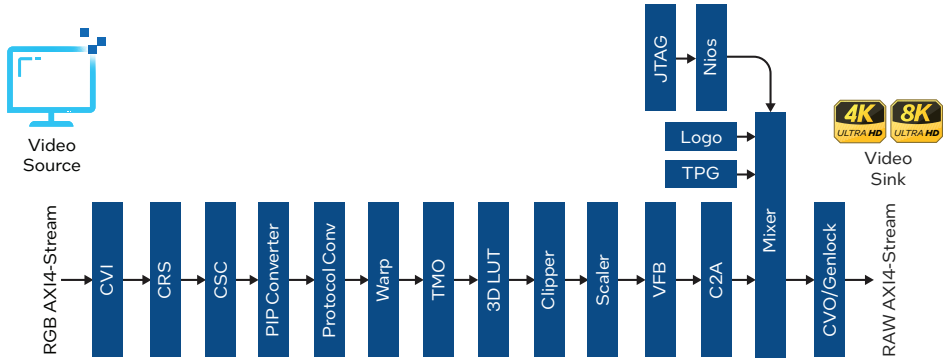

Figure 7. With the VVP library, you can easily create complex video-processing pipelines on Intel FPGA.

VVP as Building Blocks for Video Processing

Video and Vision Processing (VVP) library is a collection of 20+ IP functions with support for all Intel FPGA. It provides a design philosophy for rapid new design creation and easy integration of custom value-add features as well as with video connectivity IP cores, such as: HDMI, DisplayPort, 12G-SDI, SMPTE 2110, and MIPI.

The library supports a wide range of resolutions, frames per second (fps), bits per color (bpc) (1080p/4K/8K, HDR ready, 120+ fps, 16 bpc) offering a visual quality that exceeds most of the ASSPs.

Using Open FPGA Stack (OFS) and Intel® FPGA SmartNIC N6000-PL Platform to Expedite Virtualized FPGA Developments

The OFS is a scalable, source-accessible hardware and software infrastructure delivered via git repositories that enables you to customize your own unique acceleration platform solutions. This second-generation, hardware-and-software infrastructure features Intel® Stratix® 10 FPGA, Intel Agilex 7 FPGA, Intel Agilex 5 FPGA, and future Intel FPGA families. This infrastructure is used by Intel and selected third-party platforms to address the challenges associated with designing FPGA-based acceleration platform solutions.

OFS enables an efficient path to develop a custom FPGA

platform as it:

- Provides a framework of FPGA synthesizable code, simulation environment, and synthesis or simulation scripts.
- Accelerates workload development with industry-standard Arm AMBA 4 AXI interface, workload examples, and AFU simulation.
- Leverages software drivers up streamed to the Linux kernel and OPAE software and libraries.

The Intel® FPGA SmartNIC N6000-PL Platform, formerly known as the Arrow Creek Acceleration Development Platform (ADP), consists of a pair of PCIe card designs featuring Intel Agilex 7 FPGA and Intel Agilex 5 FPGA that are used as the basis for production cards by a partner ODM and/or customers for different markets. The key difference between the designs is that one includes an integrated dual port 100GbE Intel E810 Ethernet Controller.

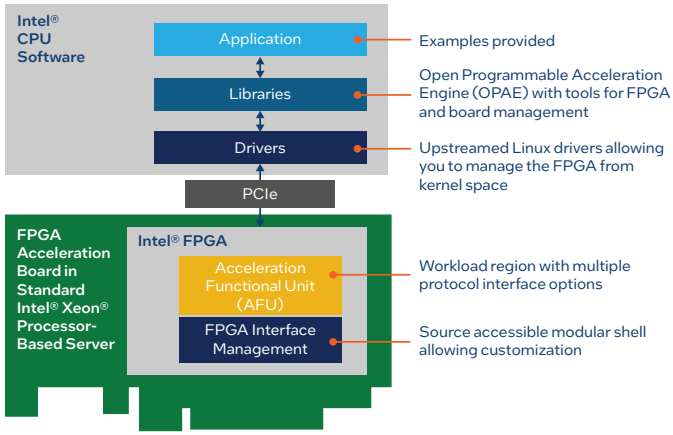

Figure 8. OFS provides a complete set of hardware and software resources to speed up the development of virtualized FPGA video processing applications.

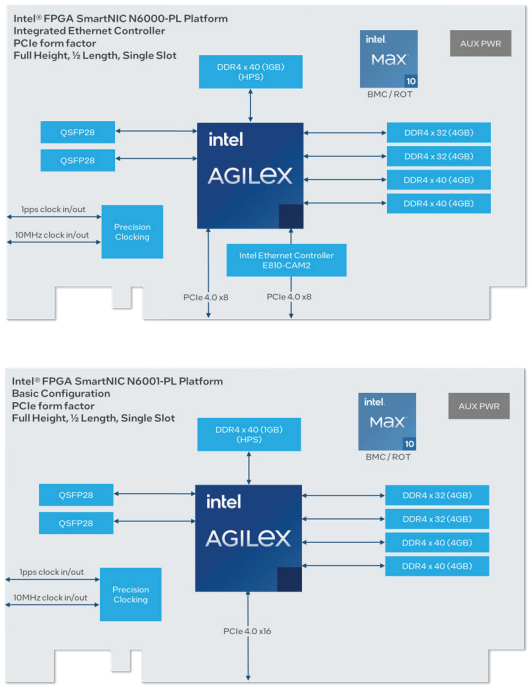

Figure 9. Intel FPGA SmartNIC N6000-PL Platform is a ready-to-use PCIe card with onboard Intel Agilex 7 and Intel Agilex 5 FPGAs and fully supported by OFS.

The Intel FPGA SmartNIC N6000-PL Platform provides a reference hardware platform enabling customers to develop, tune, and optimize specific workloads on real hardware before going to full production. These customers are typically at the point where they want to make production-quality cards, but first want to optimize their IP before committing to a final design.

Figure 10. Physical view of the Intel FPGA SmartNIC N6000-PL Platform board.

The Intel FPGA SmartNIC N6000-PL Platform is the first in a family of development platforms targeting a wide range of market segments and solutions.

The Intel FPGA SmartNIC N6000-PL Platform includes multiple OFS-based FIM designs for the Intel Agilex 7 and Intel Agilex 5 FPGAs supporting several memory, Ethernet, and PCIe system configurations. The key goal of OFS is to provide quicker time to market by offering source-accessible hardware, open source software, and a simulation infrastructure that you can easily customize for your platform application.

Summary

Intel FPGA provides the right architecture, compute capabilities, and power efficiency required to handle the requirements for live video production workflows. Intel FPGA can implement full-featured video processing pipelines by using different connectivity and video/audio processing building blocks and enable low latency implementations by dynamically reconfiguring optimized data paths. Intel provides the right ingredients ranging from silicon (Intel Agilex 7 and Intel Agilex 5 FPGAs), IPs (VVP, connectivity), and complete platforms to enable COTS/ virtualized flows (Intel FPGA SmartNIC N6000-PL Platform and OFS).

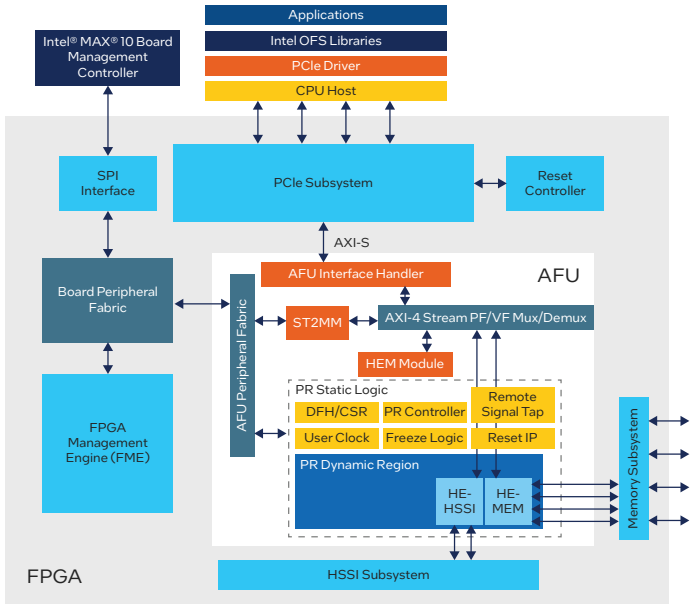

Figure 11. Detailed diagram of different sub-modules included in the OFS package.

Where to Get More Information

- OFS web page: www.intel.com/ofs
- How Intel Agilex FPGA is Enabling Resource and Power Efficient 4K, 8K, Video Processing Solutions White Paper: www.intel.com/content/dam/www/central-libraries/us/en/documents/agilex-fpga-video-processing-solutions-white-paper.pdf
- Video and Image Processing Intel FPGA IP web page: www.intel.com/content/www/us/en/programmable/products/intellectual-property/ip/dsp/m-alt-vipsuite.html

WP-01314-1.1