Enabling a JPEG 2000 Network for Professional Video

by Jean-Marie Cloquet
Manager, Image Processing Division
Barco Silex
jean-marie.cloquet@barco.com

A new reference design from Xilinx and Barco Silex offers JPEG 2000 video transport over Internet Protocol networks.

Because of its superior quality, JPEG 2000 has emerged as the standard of choice for the compression of high-quality video, including the transport of video in the contributing networks of television broadcasters. As a result, suppliers of video equipment have started adding JPEG 2000 encoders and decoders to a variety of transport solutions, supporting various interfaces and sometimes even using proprietary protocols.

This trend, however, has locked video service providers into the products of one or a few vendors. A solution to this dilemma arrived in April 2013 with the publication of the Video Services Forum's (VSF) TR-01 recommendation for transport of video over Internet Protocol (IP) networks, a specification for interoperable equipment. Xilinx and Barco Silex, a certified member of the Xilinx® Alliance Program, quickly joined forces to support the interoperability effort.

Barco Silex has now completed a reference implementation of the VSF TR-01 recommendation. The multichannel video-over-IP with JPEG 2000 solution was recently made public on the Xilinx website. It is based on intellectual-property cores from Xilinx and Barco Silex, and is ready to be customized and integrated by broadcast equipment OEMs. Honoring this effort, the National Academy of Television Arts and Sciences awarded Barco Silex a 2014 Technology & Engineering Emmy Award (Figure 1).

LOOKING FOR SUPERIOR VIDEO COMPRESSION

JPEG 2000 supersedes the older JPEG standard and offers many advantages over its predecessor or other popular formats such as MPEG. By 2004, JPEG 2000 had become the de facto standard format for image compression in digital cinema through the Hollywood-backed Digital Cinema Initiatives (DCI) specification. The possibility of a visually lossless compression makes JPEG 2000 ideal for security, archiving and medical applications.

The broadcasting industry also took notice. Broadcasting and video service companies have huge amounts of live video that has to be transported to post-production and streaming facilities within their so-called contribution networks (Figure 2), without delay or loss of visual quality. Of particular interest for the professional-video industry, therefore, is the possibility of a visually lossless compression—that is, a compression scheme that retains the image quality and still allows efficient storage and transport.

In addition, the other innovations in JPEG 2000 also meant a step forward for the broadcasting industry. Each frame in the video stream is compressed individually as a still frame, in contrast to MPEG formats, which compress frames in groups. This single-frame compression technique results in a low latency but also makes for easy per-frame post-processing and editing. A JPEG 2000 stream may also be partially decompressed and viewed, allowing different applications and viewing experiences from the same stream.

Another big plus is the resilience against transmission errors in the stream. If transmission errors cannot be corrected using forward error correction (FEC), the errors will have a smaller visual impact after decoding compared with other codecs. Finally, JPEG 2000 preserves the image quality even after multiple encoding/decoding processes, which is of capital importance in contribution networks with various stages of video management.

Picking up on this interest, equipment suppliers soon started to implement JPEG 2000 encoders and decoders in their video gear. However, for the transport between locations, they still had a choice among a wide range of implementation options, including proprietary protocols. The drawback for video service providers was that they had to lock into the products of one or a few vendors, instead of setting up the best-matching, cost-efficient infrastructure.

A RECOMMENDATION TO STANDARDIZE VIDEO TRANSPORT

So there was a clear demand from the service providers for a standardized transport to ensure better interoperability between existing and future equipment. What they needed was a transport that could be best organized over IP networks, which were becoming the prevailing network architecture, with standardized equipment ready for high-throughput data transport. Starting in 2007, the Society of Motion Picture and Television Engineers (SMPTE) published a standard for video transport over IP, which has been expanded since. SMPTE 2022 includes, among others, IP protocols for constant-bit-rate video signals in MPEG-2 transport streams (SMPTE 2022 1&2 for compressed video and SMPTE 2022 5&6 for uncompressed video).

Taking these specifications as its basis, the Video Services Forum in 2013 published its VSF TR-01 document, a technical recommendation titled "Transport of JPEG 2000 Broadcast Profile Video in MPEG-2 TS over IP." The VSF is an international association composed of service providers, users and manufacturers dedicated to interoperability, quality metrics and education for video networking technologies.

Any device that adheres to VSF TR-01 will take its input from an SDI (serial digital interface) signal, the legacy standard for uncompressed point-to-point video transport in the broadcast industry. The device will extract the active video, audio and ancillary data (for example, captions) and compress the video in JPEG 2000 format.

The resulting stream is multiplexed into an MPEG-2 transport stream together with the audio and ancillary data. This stream is again encapsulated according to SMPTE2022 in a Real-time Transport Protocol (RTP) stream and transmitted over IP to a receiving device. The receiver will de-encapsulate the RTP/IP stream, demultiplex the MPEG-2 transport stream, decode the JPEG 2000 and place the video, audio and ancillary data onto the output SDI signal.

IMPLEMENTING AN FPGA-BASED REFERENCE SOLUTION

In September 2012, even before the VSF recommendation was published, Xilinx and Barco Silex announced a partnership to develop video-over-IP solutions. The goal was to offer a comprehensive platform of hardware-validated intellectual-property cores, reference designs and system integration services. In this effort, Barco Silex took on the role of system integrator, matching cores from Xilinx (SMPTE 2022, SMPTE SDI, Ethernet MACs) with its own high-performance JPEG 2000 and DDR3 memory controller cores. The goal was to enable OEMs of broadcast equipment to accelerate their product development and to add the latest video-over-IP capabilities to their existing products and those currently in development.

In this framework, the partners have now completed a reference design, composed of a four-channel transmitter-and-receiver platform (Figure 3). The transmitter is able to take up to four SDI high-definition (HD) streams (1080p30), optionally compress them with JPEG 2000 and send them over 1-Gbps (with compression) or 10-Gbps (uncompressed) Ethernet according to the VSF TR-01 standard. The receiver platform, conversely, can receive the IP stream, de-encapsulate and decompress it, and put it on up to four SDI HD links.

In the transmitter platform, Xilinx SMPTE SDI cores receive the incoming SDI video streams. On the uncompressed path, these SDI streams are multiplexed and encapsulated into fixed-sized datagrams by Xilinx's SMPTE 2022-5/6 video-over-IP transmitter core and sent out through the Xilinx 10-Gigabit Ethernet MAC (10GEMAC) and 10G PCS/PMA cores.

On the compressed path, the SDI streams first go to the JPEG 2000 encoder for compression. Next, they are encapsulated into MPEG-2 transport stream packets according to VSF TR-01 by the dedicated TS Engine core implemented by Barco Silex. Last, the SMPTE 2022-1/2 video-over-IP transmitter core packs the streams into fixed-size datagrams and sends them out through the 1G TEMAC. Alternatively, the streams can be multiplexed with uncompressed video channels on a 10-Gbit link using the 10GEMAC and 10G PCS/PMA cores.

On the receiver platform, the Ethernet datagrams of the uncompressed streams are collected at the 10GEMAC. The SMPTE 2022-5/6 video-over-IP receiver core filters the datagrams, de-encapsulates and demultiplexes them into individual streams, and outputs the SDI video through the SMPTE SDI cores. The Ethernet datagrams of the compressed streams are collected at the 10GEMAC, de-encapsulated by the SMPTE 2022-1/2 video-over-IP receiver core and by the TS Engine, and fed to the JPEG 2000 decoder. Its output video is converted to SDI and sent to the SMPTE SDI cores.

For each of the four channels, the uncompressed or compressed path can be chosen independently of what happens on the other channels.

LAYING THE BASIS FOR INTEROPERABLE SOLUTIONS

The companies implemented the reference design in two platforms, one using the Zynq®-7000 All Programmable SoC and the other using the Kintex®-7 FPGA. But the blocks that are used can be integrated in solutions that address the complete range of OEM system requirements, from low-cost, high-volume applications to the most demanding high-performance applications. The intellectual-property cores that were used, such as Xilinx's SMPTE 2022 and Ethernet MAC LogiCORE™ blocks, are available for the full range of Xilinx FPGA systems, up to the UltraScale™ level.

For encoding and decoding, the reference design includes the Barco Silex JPEG 2000 encoder and decoder IP cores. These are silicon-proven, widely adopted single-FPGA solutions for high-performance, simultaneous multi-channel 720p30/60, 1080i, 1080p30/60 and 2K/4K/8K JPEG 2000 encoding and decoding. These cores also support the widest available spectrum of JPEG 2000 options in existence on the market. Essential in stitching multiple video streams together into a smooth, high-data-rate system is

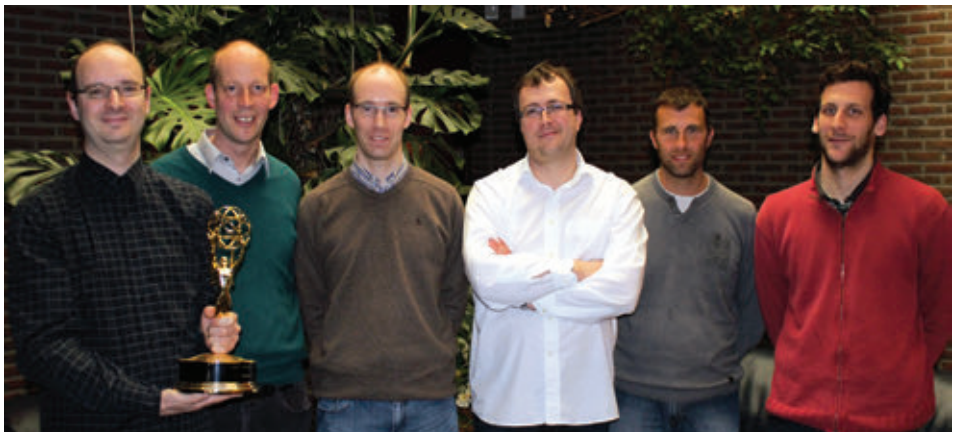

Figure 1 – The Barco Silex video team responsible for the reference design with their 2014 Technology & Engineering Emmy Award for Standardization and Productization of JPEG 2000 Interoperability. From left, they are Luc Ploumhans, Sake Buwalda, François Marsin, Jean-François Marbehant, Jean-Marie Cloquet and Vincent Cousin.

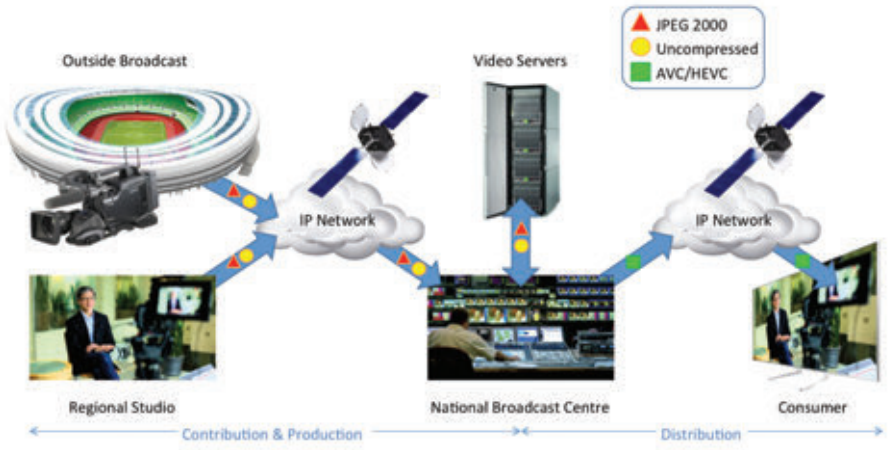

Figure 2 – The contributing networks of broadcasting companies

Barco Silex's DDR3 memory controller. This highly customizable controller is optimized to achieve high bandwidth by reordering accesses and mixing them to the different banks of the SDRAM.

The companies showed a first generation of this reference design in a public interoperability demonstration during the annual VidTrans conference held in February 2014 in Arlington, Va. During this test organized by the VSF, 10 companies (Artel, Barco Silex, Ericsson, Evertz, Imagine Communications, IntoPIX, Media Links, Macnica, Nevion and Xilinx) provided technology and equipment that was interconnected to show live transmission of 720p30 and 1080i60 HD content being compressed in real time using JPEG 2000 encoders and decoders.

A few months later, Barco Silex demonstrated that the reference design could also handle 4K and ultra-high-definition (UHD) signals. As one of the main standards proposed for next-generation video distribution, 4K video carries four times as many pixels as 1080p video, allowing a sharper view and a larger video display. Using the four input channels of the reference design in quad-SDI mode (4K carried over four SDI cables), it is now also possible to take as input a 4K signal and send it over the IP network. This makes the reference design ready for video resolutions up to 4K.

FPGAS ADVANCE THE VIDEO INDUSTRY

The goal of the collaboration between Xilinx and video specialist Barco Silex was to leverage the power and flexibility of FPGA-based platforms in the professional-video market. By combining the JPEG 2000 cores of Barco and the transport cores of Xilinx, OEMs may produce and update standardized broadcast equipment quickly, making their products future proof in the process.

This reference design arrives at a time when the use of IP networks in the video industry is beginning to take off. The ability of OEMs to capture a share of that new market will depend on how fast they can get products out. With reprogrammable solutions based on Xilinx FPGAs, they can launch products even while standards are still evolving.

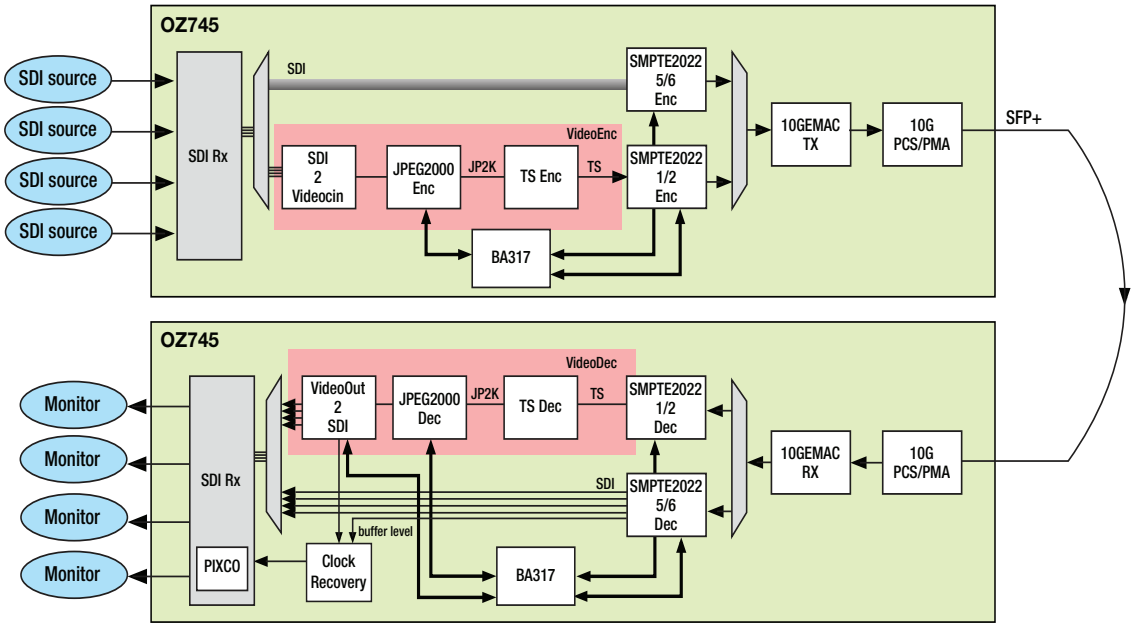

Figure 3 – Schema of the reference design with transmitter and receiver platform