

Wall Screen: an Ultra-High Definition Video-Card for the Internet of Things

Miguel Costa, Ricardo Moreira, Jorge Cabral, José Dias*, Sandro Pinto
Centro ALGORITMI, Universidade do Minho, Portugal

*MOG Technologies, Portugal

{miguel.costa, ricardo.moreira, jcabral, sandro.pinto}@dei.uminho.pt, *jose.dias@mog-technologies.com

Abstract—8k ultra-high definition (UHD) is paving the way for the next-generation video systems. In the audiovisual industry, besides delivering a more immersive experience, it is a mean to smooth spatial artifacts during video sampling. In the medical industry, it may provide surgeons with increased reality in surgeries such as endoscopy. Nevertheless, researchers are struggling to meet the high throughput required by this resolution and hardware solutions miss the flexibility required for on-demand updates. In this context, we propose an 8k video-card based on a hybrid platform endowed with "soft" programmable logic and "hard" processors. The former accelerates the video pipeline from capture to encoding and playback - via SDI or PCIe. The later connects with a cloud to stream video and interfaces with the user. Results showed that our platform outputs 8k UHD video in YUV 4:2:2 10-bits and H.264 formats at 60 and 30 frames per second, respectively.

Index Terms—ultra-high definition, UHD, 8k, video-card, video, Wall Screen.

I. INTRODUCTION

8k ultra-high definition (UHD) was recently proposed as the standard for the next-generation digital television broadcasting [1]. Japan planned to pioneer in this area by performing the first large-scale broadcasting of 8k content in the 2020 Olympics [2]. Moving towards 8k UHD video systems raises a set of new technological challenges, which stands out the ability to process 8k video data in real-time. In terms of pixel resolution, 8k is four times bigger than 4k and sixteen times greater than full high definition (FHD) formats, thus involving tremendous computation and data bandwidth [1]. Even with advanced compression methods such as H.264 or H.265, 8k video still requires large bandwidths [3,4].

These engineering challenges is leading industry players and the scientific community to design systems for the acquisition and playback of 8k video based on platforms with dedicated hardware [1,3,5]. If conventional hardware solutions pose a bottleneck to on-demand updates, reconfigurable platforms, such as field-programmable gate array (FPGA), satisfies both the huge computational requirements of 8k video and the required flexibility for algorithm updates.

Empowered with such emerging challenges, we propose a hardware-accelerated system capable to capture and playback, in real-time, video with resolutions up to 8k (7680 x 4320). The proposed platform, based on the Zynq UltraScale+ SoC, employs a hybrid architecture whereby an FPGA accelerates the video processing pipeline, while ARM application processors manages the video-card connection with the cloud.

More specifically, the FPGA implements the logic to interface with serial digital interface (SDI) and PCIe connections, for video capture and transmission, and contains hardware blocks for video encoding and playback through HDMI. The processing system enables the video data publication on the Internet, on-demand user configuration, and monitors the video-card performance, notifying the user about critical errors. Those features explore the current IoT trends to interconnect the video domains in the video production environment. Results show that our video-card can process 8k video in YUV 4:2:2 10-bits and H.264 formats at 60 and 30 frames per second (fps), respectively.

II. MOTIVATION

To cope with the relevance of 8k video for the industry, our platform aims at tackling the main challenges in the processing of 8k video and explore the ongoing IoT trends to address some problems in the professional video domain.

A. Why we need 8k?

It is believed that improved resolution translates into better visual differentiation and better depth perception. Fueled by this conception, the scientific community is putting some effort to uncover the impact of 8k video in different application domains, ranging from the multimedia industry to the healthcare industries.

In the audiovisual industry, 8k may empower anti-alias filters to smooth the spatial artifacts resultant from the sampling of images into discrete pixels, such as edge "jaggies" and "moiré" patterns [6]. Anti-alias filters are unfeasible in resolutions near to the human visual acuity, including FHD, as they blur the image at a point that is perceptible to a common viewer. In this context, some authors [6] suggest that increasing the image resolution beyond the human visual acuity and then apply an anti-alias filter can be the solution to deliver images with no spatial artifacts while maintaining the sharpness of the edges.

Sports networks are particularly interested in 8k technology as they believe it brings increased levels of realism making the public experience the full impact of live events [2]. Japan Broadcasting Corporation (NHK) planned to broadcast the 2020 Japan Olympics in 8k, having already launched the world's first satellite for television broadcasting [2].

There are also heightening expectations for the contribution of 8k technology in the healthcare industry, particularly on endoscopy [7]–[10]. As the successful outcome of this surgery depends heavily on the intra-operative image of the organs of interest, the clinical application of 8k is already being evaluated. In November 2014, the first-ever cholecystectomy using an 8k endoscope was performed in Japan [7]. From that date on, the 8k camera used in this endoscope was improved and adapted to use in other clinical cases, including gynecologic [8] and ophthalmic [9] surgeries, and colorectal cancer treatment [10]. Surgeons reported that 8k images were excellent in reproducing appearances of solidity and reality, giving a sense of looking onto the original field of view and causing less eye strain [7]–[10].

B. Why the capture and playback of 8k video is challenging?

In terms of pixel resolution, 8k is four times bigger than 4k and sixteen times greater than FHD [1]. Therefore, the main challenge in the migration of current video systems to 8k is in the handling of the respective volume of data. 8k video in YUV 4:2:0 8-bits at 30 fps generates a minimum data rate of 1.40 GB/s (Table I). Although this chroma subsampling leads to the loss of the third color channel, current high-end video-cards are still struggling on supporting such data rate. Moving to a video production environment rises even more daunting challenges. 8k video at 30 fps with no chroma subsampling (YUV 4:4:4) and 10-bits of color depth requires a minimum data rate of 3.48 GB/s (Table I), which is not supported by any video-card on the market. This analysis does not consider the inevitable padding between pixels, which will push even further the already high data rates required. Section V lists the most powerful video cards available for the professional video domain and highlights their limitations in the processing of 8k video.

The ongoing trend of broadcasting video content for domestic consumption also raises challenges concerning network bandwidth. To reduce the effort put on the network, researchers are exploring advanced video compression techniques, such as H.264 or H.265. However, even the most advanced encoder (H.265) still demands a bandwidth between 80 Mb/s and 100 Mb/s [4]. Considering the limitations of the current 4G LTE networks, this would limit the availability of 8k content to a niche of users. The mass broadcast is expected to become more practical in the future when 5G technology becomes widespread [11]. The broadcast of 8k video also puts a lot of effort into the computational power of the receptor devices due to the complexity of the video compression formats [3].

TABLE I: Minimum data rate for 8k video (GB/s)

	8-bits			10-bits			16-bits		
	4:2:0	4:2:2	4:4:4	4:2:0	4:2:2	4:4:4	4:2:0	4:2:2	4:4:4
8k@30	1.40	1.86	2.79	1.74	2.32	3.48	2.79	3.71	5.57
8k@60	2.79	3.71	5.57	3.48	4.64	6.96	5.57	7.42	11.13
8k@120	5.57	7.42	11.13	6.96	9.27	13.91	11.13	14.84	22.25

C. The role of IoT in professional video-cards

The current lack of integration between the video acquisition and distribution tools makes maintenance tasks, such as installations, upgrades, anomalies verification, and configurations to be often performed on the entire installed equipment fleet. This task raises frequent non-conformity and unintended behaviours, which are only detected when the equipment is submitted to the execution of certain operations. Exploring the current information systems infrastructure to enable the management, monitoring, and collection of statistical data about the video equipment performance and error trace will bring definitely compelling advantages.

Furthermore, current content acquisition and distribution tools are normally developed without taking as a priority requirement the integration and adaption between the video acquisition equipment and the underlying needs of the web streaming infrastructure. Considering that 15.2% of the world’s population in 2019 paid for a video streaming service, a market with a revenue of almost 23 B€ [12], tackling this issue is paramount. As detailed in Section V, even the most powerful video cards on the professional video domain does not provide facilities for video web streaming.

III. WALL SCREEN VIDEO-CARD

The proposed video-card is built upon the XCZU19EG - FFVC1760-2 SoC, which is endowed with a processing system (four ARM Cortex-A53) and programmable logic (FPGA fabric). Fig. 1 depicts the overall video-card architecture, highlighting its main hardware blocks and the data-flow between them. The communication is handled by AXI4 buses, being AXI4-Full used for high-performance memory access and AXI4-Stream used for high-speed video data streaming.

Uncompressed video frames can be captured via SDI or PCIe. When it comes to SDI, data flows through the video-card via a 12G-SDI quad-link connection, which implies the use of four high-speed GTH transceivers for serial data transfer. These are interfaced with hardware blocks included in the SMPTE UHD-SDI Subsystem, where the information is interpreted and transformed in AXI4-Stream Video. Streams coming from or heading the four SDI channels are serialized or de-serialized before being forwarded to other modules.

The read and write of video frames in memory are controlled by the Memory Manager. This (in-house) module attests to the validity of every read video frame, employing mechanisms to detect whether a frame is corrupted. As memory access is the most critical part of the video-card, it is equipped with a memory chip only dedicated to SDI-PCIe transactions. The DMA/Bridge PCIe Subsystem has direct access to this memory chip either for reading and writing operations. Nevertheless, the integrity of the frames is still ensured by the Memory Manager. To make the video streaming over Wi-Fi feasible, the video-card employs an H.264 Encoder. Compressed frames are placed by the Memory Manager in a distinct memory region.

The last element that accesses memory is the Processing System. The Processing Systems runs (a supervised version of) Linux, and provides an easy-to-use channel from the cloud

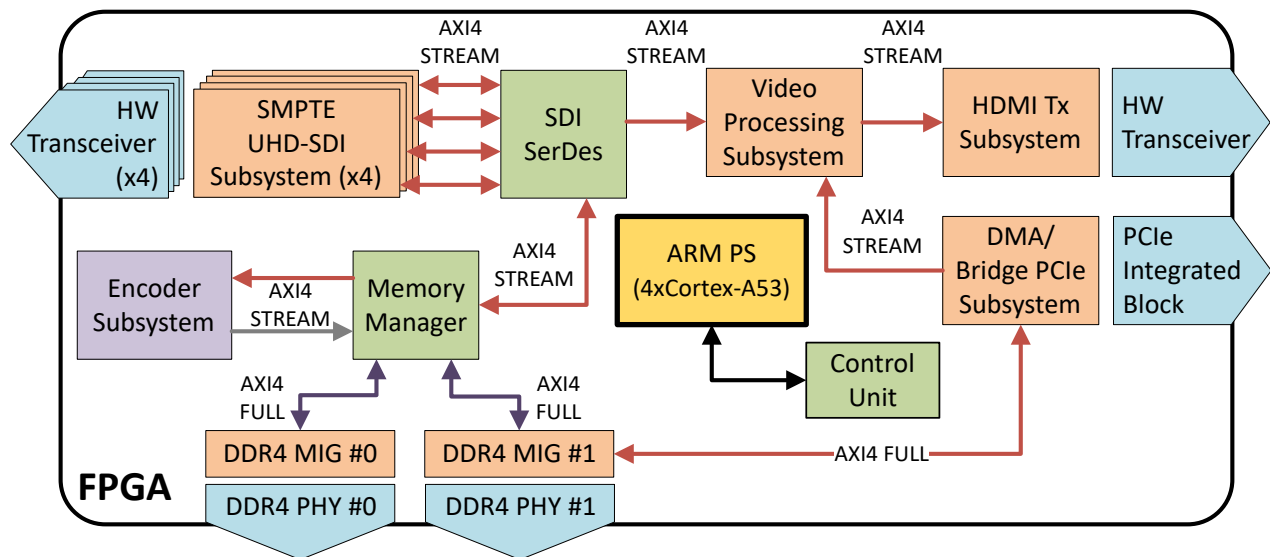


Fig. 1: Wall Screen video-card architecture. Orange: Xilinx IP Cores; purple: open-source IP Cores; green: in-house IP Cores.

to the system. Besides streaming encoded video to the cloud, it is also an interface for the configuration of the video-card, performed in the last instance by the Control Unit. The Control Unit also collects data relating to the video-card status, sending it to the Processing System, which interfaces with the user.

The HDMI Transmission Subsystem enables the real-time playback of video received from PCIE or SDI via HDMI. The conversion of the video format (e.g. resolution, chroma subsampling, etc.) to meet the configurations of the HDMI controller is performed by the Video Processing Subsystem.

A. UHD-SDI Subsystem

This subsystem performs the interface between the physical quad small form-factor pluggable (QSFP) transceivers and the remaining components of the video board. For each serial line of the QSFP transceiver, this subsystem delivers an interface to two AXI4-Stream buses, one for data transmission and another for data reception. It supports a myriad of Society of Motion Picture and Television Engineers (SMPTE) standards, including the most recent SMPTE ST 2082-1 (12G-SDI). Notwithstanding, the chroma subsampling is limited to 4:2:2, with a maximum of 10 bits of depth-color.

As two pixels require 5 bytes of memory, 8k video (7680x4320), in YUV 4:2:2 10-bits format, at 60 fps requires a minimum data rate of 37.08 Gb/s. This data rate requires, at least, an octa-link 6G-SDI or quad-link 12G-SDI setup, both achieving a maximum data rate of 48.00 Gb/s. To reduce the engineering effort in the serialization and de-serialization of SDI video channels, we decided for the later setup.




Fig. 2: Stream issued by SDI Ser/Des - 2 pixels per clock.

B. SDI Serializer/De-serializer

When capturing video via SDI, this module serializes the four video streams received, per clock, into a single video signal. The merging process follows the same image mapping standard implemented by the video source - Quad-Link SDI 2SI (two-sample interleaved). In this format, the first two links carry the odd video lines, while the third and fourth links carry the even lines. The odd and even lines are firstly combined and then it is carried out the switch between the odd and even lines. Data streams issued by the SDI Serializer/De-serializer follow the UYVY standard (Fig. 2).

C. HDMI Transmitter and Video Processing Subsystems

The HDMI Transmitter Subsystem, a Xilinx IP Core, incorporates the logic to interface with the HDMI physical layer. It parses the input video and audio streams, coming from the SDI and PCIe interfaces, transforming them in an HDMI stream which is then forwarded to the physical layer. This subsystem supports HDMI 2.0 and allows the user to configure a myriad of settings for HDMI video playback. The conversion from the video coming from the SDI and PCIe channels to the specific format for HDMI playback is performed by the Video Processing Subsystem, also developed by Xilinx. These two IP Cores enable the video playback via HDMI with:

- Resolutions up to 4K (4096 x 2160) at 60 fps;
- 8, 10, 12, and 16-bit depth-color;
- Color space for RGB, YUV 4:4:4, 4:2:2, and 4:2:0.

D. Memory Manager

The Memory Manager regulates the reading and writing of video frames in memory. It considers the memory organized as a circular linked list, where each slot is dedicated to a single frame. Raw video frames are saved in memory following the same data structure used in the video data streams issued by the SDI Serializer/De-serializer (Fig. 2).

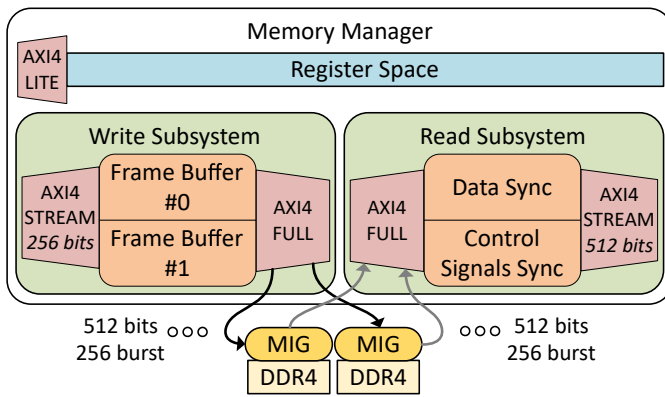


Fig. 3: Memory Manager architecture.

To ensure the throughput required for the flow of 8k video in real-time, the video-card is equipped with two memory chips - (i) one dedicated to SDI-PCIe transactions and (ii) another dedicated to the remaining modules. The former only contains uncompressed video and it is directly accessed by the PCIe interface. The Memory Manager must still ensure the validity of each read frame, by signaling the occurrence of race-conditions in the memory region being read. The later chip contains uncompressed video frames to be encoded by the H.264 Encoder and the respective encoded frames to be further streamed by the Processing System. Fig. 3 depicts the Memory Manager architecture.

Write Subsystem: This subsystem converts the data frames received by its AXI4-Stream video ports into AXI4-Full data frames to be written in memory by the memory controller. To ensure maximum throughput and lessen the impact of additional latency of AXI4-Full over AXI4-Stream protocol, this subsystem relies on a buffering mechanism to take advantage of the maximum data rate supported by the first protocol. The Memory Manager buffers all incoming data until it can ensure a write burst of 256 words containing 512 bits. Since the transmission of buffered data to memory occurs in parallel with the buffering of new data, the employed buffering mechanism follows a double-buffer topology to avoid the occurrence of race conditions.

The Write Subsystem is also equipped with mechanisms to detect a series of critical errors: (i) early end of line, (ii) late end of line, (iii) unexpected start of frame (early or late), and (iv) data loss. Early/late end of line errors are issued whenever the number of bytes received within a frame line is inferior/superior to the expected. The unexpected start of frame occurs whenever the number of lines received between two consecutive assertions of the start-of-frame signal differs from the frame height programmed in the register space. As these errors occur due to misconfiguration of the video-card, the Memory Manager halts, waiting for a reset signal and a new configuration from the Control Unit to restart its operation.

Data loss occurs when the memory controller (DDR4 MIG) can not receive data for too long, leading to the re-write of the buffer (Buffer #0 or Buffer #1) that has just been filled up. This error does not halt the Memory Manager as it does not result from an erroneous configuration.

Read Subsystem: A frame reading process starts with the configuration of the AXI4-Full bus, established between the Memory Manager and the memory controller using the start address of the desired frame. As the respective data packets arrive on this AXI4-Full bus, the Memory Manager generates a set of AXI4-Video streams that are sent to the modules that previously issued read requests. To optimize the throughput and reduce the latency disparity between AXI4-Full and AXI4-Stream protocols, the former was set to transfer 256 words of 512 bits (maximum configuration supported) towards the Memory Manager for each validated address.

Anytime a new frame is received in the Memory Manager, the Write Subsystem replaces the oldest frame in memory, even if the corresponding address is being read. However, the Memory Manager detects such situations, updating its status register anytime a race condition occurs and signaling it to the module reading the frame.

E. Encoding

The Encoding Subsystem emerges from the need of compressing the raw video stored in memory to enable its streaming, over Wi-Fi, to the connected IoT devices. This encoder is based on an open-source solution for the H.264 standard¹, which implements a non-interlaced base profile with no limit to the number of streams or video resolution. Notwithstanding the chroma subsampling of the video to be encoded is limited to YUV 4:2:0 with 8 bits of depth-color.

It is believed that the new H.265 encoder may potentiate the web streaming of 8k content due to its doubled compression ratio as well as improved video quality, comparing with its predecessor [3]. Nevertheless, the time to market and development costs constraints led us to select the H.264 standard. H.265 ASIC encoders, as of now, has associated a prohibitive cost to achieve competitive video-card prices. In turn, H.264 is a well-established technology with ready-available open-source HDL IP that we could leverage to extend and achieve a functional prototype in a time window. In the near future, we have all the intention to add support for H.265 in the form of ASIC or HDL.

The encoding process starts with the prediction of luma (16x16) and chroma (8x8) macroblocks. The output of the prediction modules then flows to the transformation loop, which is divided into 5 steps: (i) core transform, (ii) quantize, (iii) dequantize, (iv) inverse transform, and (v) reconstruction. As we are using intra-encoding, the reconstructed pixels are fed back into the prediction modules to be used in the prediction of the next macroblock. The output of the quantization process is then buffered and reordered to be submitted to context-adaptive variable-length coding (CAVLC). The stream returned by CAVLC is then combined with the header data and converted to a NAL stream.

CAVLC is the most used profile due to its baseline compatibility ranging from low-end to high-end devices [13]. For an acceptable video quality towards the broadcast application, the CABAC entropy coding can achieve a bitrate saving between 9% and 14% over CAVLC [14]. However,

¹<https://github.com/bcattle/hardt264>

it also requires additional hardware utilization and increased complexity [13]. Although the intra-encoding achieves a lower compression ratio compared to the inter-encoding approach [15], it is a key factor for enabling the real-time 8k video-encoding by excluding the dependency on older frames.

The encoding subsystem has undergone a multi-core design strategy to overcome the limited clock frequency of ZU19EG-FFVC1760-2 SoC (933 MHz). The encoding subsystem relies on an octa-core architecture to achieve the throughput required for the encoding of 8k video at 30 fps. Each core is dedicated to the encoding of a single frame. As the chroma subsampling of the raw video frames stored in memory (YUV 4:2:2 10 bits) is not supported by the encoder, each incoming video stream needs to be converted before the encoding process. After the chroma subsampling conversion, the video data is buffered and reorganized in macroblocks used as input of the H.264 Encoder.

F. DMA/Bridge Subsystem for PCIe

To interface with the PCIe physical layer, Xilinx offers a DMA Subsystem designed for high-performance data movement between the host address space and the FPGA. Considering the data structure of the raw video frames in memory (Fig. 2), the flow of a single second of 8k (7680x4320) video at 60 fps requires a minimum data rate of 7.42 GB/s. Either PCIe Gen2x16 ($\approx 8GB/s$), Gen3x8 ($\approx 8GB/s$) or Gen3x16 ($\approx 16GB/s$) meet this requirement. However, the later empowers future updates on the video settings (e.g. resolution, chroma-subsampling, etc.), and, consequently, we decided for the latter.

G. Connectivity and Security

The connectivity with the cloud is ensured by a quad-core Cortex-A53 running a Linux operating system (OS) (in symmetric multiprocessing configuration), which includes the TCP/IP stack required for Wi-Fi communication (Fig. 4).

To prevent a "malicious" cloud from compromising the hardware logic, the FPGA communication services are isolated from the Linux OS that communicates with the cloud. A TrustZone-assisted hypervisor developed in-house (LTZVisor [16]) is used as a key-mechanism to provide such a degree of isolation. This strategy prevents the system software running on the processing system from causing micro-architectural interference in the real-time nature of hardware accelerators. Taking into account the well-known vulnerabilities and limitations of TrustZone technology [17], we plan to modify the software stack to use an in-house static partition hypervisor [18] while leaving TrustZone to provide the few primitives for establishing secure authentication and video encryption.

On top of Linux runs the Web Engine and a Streaming Process. The Web Engine synchronizes the video-card with the cloud server to enable the querying of the video-card status and its configuration. This process is also responsible for asynchronously notifying the webserver of any errors and periodic sending of telemetry on video playback statistics. The web streaming of video data implies the the encapsulation of video frames through the FFmpeg library, which are then streamed by a Real-Time Messaging Protocol (RTMP).

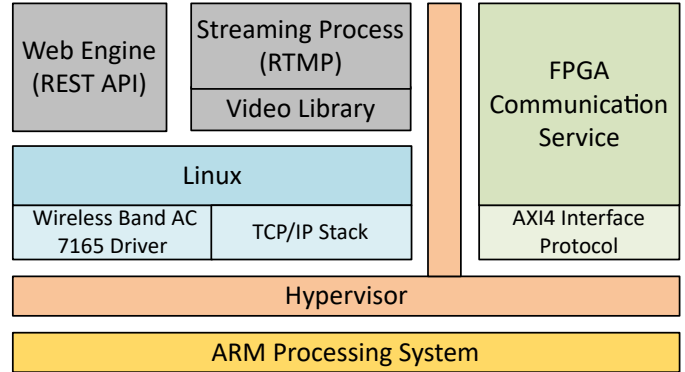


Fig. 4: Software stack. Blue: OS and drivers; grey: applications for the cloud services; orange: security infrastructure.

IV. PRELIMINARY RESULTS

The proposed video-card was evaluated in terms of memory throughput, programmable logic resources utilization and bandwidth required for video web streaming.

A. Memory Throughput

To verify if our video-card meets the throughput required by 8k video, we measured the memory throughput in different scenarios. Four AXI4 Traffic Generators were used to stress the AXI4 bus interfacing the memory controller. The number of AXI4 Traffic Generators performing reading and writing operations varied along with the test scenarios. The speed of write and read operations was measured using a performance monitor, also designed by Xilinx. Before analyzing the results, it must be remembered that the flow of 8k video, in the format detailed in Fig. 2, at 60 fps requires a minimum data rate of 7.42 GB/s. In this context, from the results exposed in Fig. 5, we can conclude that:

- **There is a slight decline in performance as concurrent memory accesses increase.** This is noticed by comparing the test cases involving only one access to the memory (test cases 1 and 2) against the test cases where there are four concurrent accesses (test cases 8, 10, and 12);
- **The maximum throughput is equally divided between write and read operations.** For every test case involving concurrent write and read operations (test cases [5-8], 10, and 12), the value of the bars representing the total write throughput ("Total WR") and the total read throughput ("Total RD") is the same;
- **The maximum measured throughput was near to 16.8 GB/s (87.5% of the estimated theoretical throughput).** This throughput was achieved in test cases with no write operations and with no more than two concurrent reads (test cases 2 and 4);
- **The throughput required for 8k video is only ensured with a maximum of two concurrent accesses.** For test cases with three or more concurrent accesses (test cases [6, 12]), the throughput of each write ("WR" bar) and read ("RD" bar) operation is under 7.42 GB/s. This conclusion led to the addition of a RAM only dedicated to transactions with the PCIe.

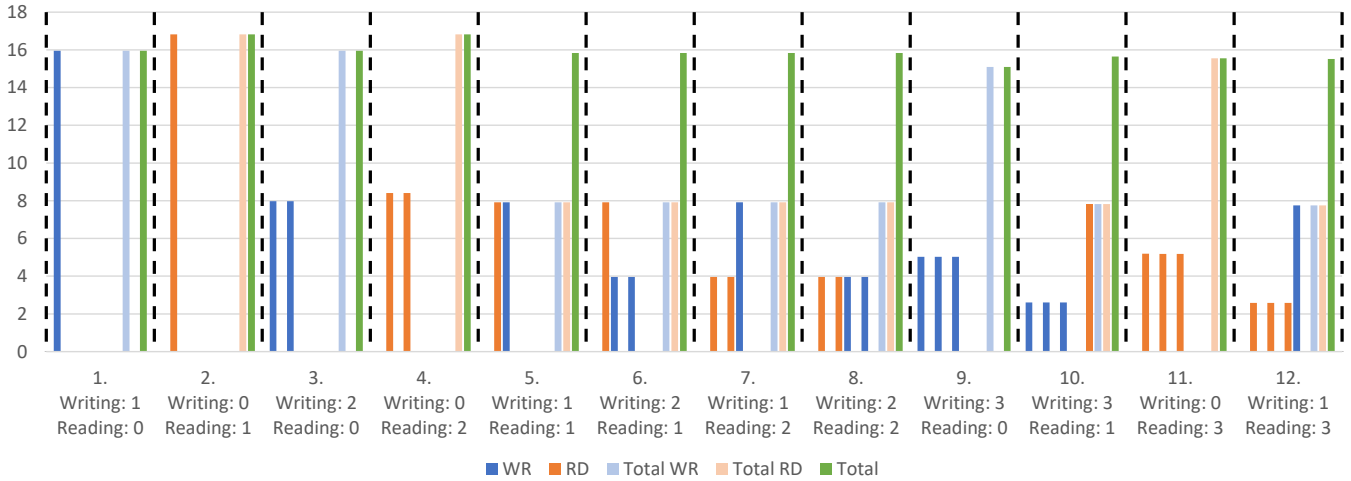


Fig. 5: Memory throughput in 12 different scenarios, involving different combinations of concurrent readings and writings.

B. Hardware Resources

The results detailed in Table II were extracted from the resource utilization report delivered by the Vivado 2019.1 framework for the ZU19EG-FFVC1760-2 SoC in the post-implementation phase. As observed, only a small fraction of the resources available were used. Nevertheless, the usage percentage of Block Ram Tiles (BRAMs) and GTH transceivers is noticeable, as their utilization ascends near to 61% and 64%, respectively. The use of BRAMs is a consequence of the frequent buffering needs that emerge in almost every hardware block. There was almost no LUTs as memory used by the Memory Manager and by the SDI Serializer/De-serializer. The use of LUTs instead of BRAMs for buffering large data blocks may lead the design to fail in the meet of the timing requirements.

The underutilization of the programmable logic resources disrupts the idea that the performance target of the design could have been higher. However, updating the video specification is still limited by the memory throughput, which is the bottleneck in this design and video-cards in general. Nevertheless, the results shown in Fig. 5 and in Table II

TABLE II: Programmable logic resources utilization.

IP Core	#	LUTs Logic	LUTs Mem.	BRAMs	DSP Slices	GTH Tx
UHD-SDI	4	10.28%	9.64%	11.40%	-	9.08%
SDI SerDes	1	0.34%	0.27%	8.94%	-	-
Video Proc. Sub.	1	5.73%	3.66%	11.43%	6.20%	-
HDMI Tx Sub.	1	1.10%	0.82%	0.15%	-	18.00%
Memory Man.	1	0.30%	0.08%	0.82%	-	-
DDR4 MIG	2	3.00%	1.58%	5.18%	0.30%	-
Encoding Sub.	1	1.34%	0.22%	8.90%	0.20%	-
DMA PCIe Sub.	1	13.42%	6.45%	13.41%	-	36.36%
Control Unit	1	0.34%	0.28%	-	-	-
Total	-	35.85%	23.00%	60.23%	6.70%	63.44%

suggest that the raw video specification could be improved from YUV 4:2:2 (10-bits) to YUV 4:4:4 (10-bits). However, such update would demand a tremendous engineering effort, namely for the SDI connection. YUV 4:4:4 (10-bits) video at 60 fps requires a minimum data rate of 55.62 Gb/s, which implies the update of the UHD-SDI Subsystem from a quad-link topology to an octa-link. Concerning the PCIe, the requirements will remain the same as long as the pixels arrangement follow the Y410 standard.

Regarding the video encoding, an update on the specifications of the output video would require a deep redesign of the encoding process, as the current encoder only supports YUV 4:2:0 (8-bits). Furthermore, increasing the already high number of encoding cores to support higher frame rates would probably disrupt serious synchronization concerns.

C. Web Streaming Bandwidth Evaluation

An evaluation study was performed to analyze the required bandwidth for transmitting the encoded video in some of the resolutions supported by the video-card. Once the compression ratio relies on the scenario and in the differences in neighborhood pixels, the results shown in Table III were taken from four different video environments during a minute.

Despite the relatively large bandwidth required for 8k video, we were able to perform real-time web streaming in a local network. Considering that 4G LTE networks achieve a peak speed of 150 Mbps [11], theoretically it is only possible

TABLE III: Bandwidth required for video @30 fps

Resolution	Scenario				
	City	Foggy Mountain	Rocky Green Mountain	Sea Islands	Mix Average
8k-7680x4320	47.0 MB/s	32.4 MB/s	67.5 MB/s	34.5 MB/s	45.3 MB/s
4k-3840x2160	15.5 MB/s	8.0 MB/s	20.2 MB/s	10.5 MB/s	13.5 MB/s
FHD-1920x1080	5.2 MB/s	2.0 MB/s	6.2 MB/s	3.3 MB/s	4.2 MB/s
HD-1280x720	2.6 MB/s	0.9 MB/s	2.9 MB/s	1.6 MB/s	2.0 MB/s
SD-640x480	0.8 MB/s	0.3 MB/s	1.1 MB/s	0.6 MB/s	0.7 MB/s

to support real-time web streaming up to 4k. Even for the most advanced encoding standard (H.265), which demands a bandwidth between 80 Mbps and 100 Mbps, will struggle 4G LTE networks, which typically have a continuous operation around 20 Mbps. On the other hand, 5G networks, which achieve a maximum speed of 10 Gbps [11], promises to easily support the web streaming of 8k video encoded by our platform. Nevertheless, we are planning a future update to the H.265 encoding standard.

V. DISCUSSION AND GAP ANALYSIS

As 8k is still an emerging technology, there are a few 8k-ready video-cards available in the market. Table IV puts into perspective the most powerful commercial UHD video cards available to the date for the professional video domain.

The Pro Capture HDMI 4k Plus and the DELTA 3G-elp-tico-d 4C are the only video-cards whose PCIe technologies can not handle the throughput required to transfer 8k video, in YUV 4:2:2 10-bits format, at 60 fps. The flow of video with such characteristics requires a minimum data rate of 7.42 GB/s, which requires, at least, the PCIe Gen2x16 or Gen3x8 technologies. DekTec DTA-2179 provides the fastest PCIe technology (Gen3x16), the same PCIe solution used in our platform. It achieves a maximum data rate of 15.75 GB/s, more than two times the required for 8k video.

In terms of SDI interfaces, none of the video-cards reviewed can handle 8k video in YUV 4:2:2 10-bits format at 60 fps. With a minimum data rate of 37.08 Gb/s, this configuration requires, at least, an octa-link 6G-SDI connection, whose throughput can go up to 47.52 Gb/s. Our platform uses a more recent and fastest standard (12G-SDI), achieving this same data rate with only a quad-link topology. This approach reduces the overhead and engineering effort in the synchronization of SDI channels.

Similarly to our platform, the platforms available in the market do not support HDMI 2.1 technology, the only HDMI version that supports the data rates required for 8k video. This arises as a consequence of this standard novelty, which was only introduced in November 2017.

In what concerns the integration with IoT technologies, our video-card pioneers in the web streaming of video content without requiring the assistance of third-party systems, such as a host computer. Furthermore, to the best of the authors' knowledge, this is the first video-card in the professional video domain to allow real-time transfers of 8k video between PCIe and SDI interfaces, with a chroma subsampling higher than YUV 4:2:0 8-bits and at a frame rate higher than 30 fps.

VI. CONCLUSION

In this paper, we proposed a hybrid platform designed for the capture and playback of 8k video in the professional domain. The FPGA side provides a series of hardware blocks that accelerate the video processing pipeline from the video capture to encoding and playback. The processing system is responsible for establishing a communication channel with a cloud infrastructure, enabling the web streaming of video

TABLE IV: Industry video-cards vs. Wall Screen.

video-card	PCIe	SDI	HDMI	Web Streaming
DELTA 3G-elp-tico-d 4C	Gen2x8 (4.00 GB/s)	Quad-link 3G-SDI (12 Gb/s)	-	-
Pro Capture HDMI 4k Plus LT	Gen2x4 (2.00 GB/s)	Single-link 6G-SDI (6.00 Gb/s)	HDMI 2.0	-
Bluefish444 Kronos Optikos	Gen3x8 (7.88 GB/s)	Dual-link 3G-SDI (6.00 Gb/s)	HDMI 2.0	-
DekTec DTA-2179	Gen3x16 (15.75 GB/s)	Octa-link 3G-SDI (24.00 Gb/s)	-	-
Wall Screen	Gen3x16 (15.75 GB/s)	Quad-link 12G-SDI (48.00 Gb/s)	HDMI 2.0	H.264

content, the on-demand configuration of the video-card and the report of critical errors back to the user.

This hybrid platform pioneers in the transmission of 8k video, in YUV 4:2:2 10-bits format, at 60 fps, between PCIe and SDI interfaces. For web streaming, the platform relies on an octa-core H.264 encoder, which delivers a stable frame rate of 30 fps. Future work encompasses the update to a single-core encoder compliant with the H.265 standard.

ACKNOWLEDGMENT

This work is supported by: European Structural and Investment Funds in the FEDER component, through the Operational Competitiveness and Internationalization Programme (PORTUGAL 2020) Project no **017891**; Funding Reference: **POCI-01-0247-FEDER-017891**. This work has also been supported by FCT within the Project Scope: UID/CEC/00319/2019.

REFERENCES

- [1] Y. H. Park, J. Kim, M. Kim, W. Lee, and S. Lee, "Programmable multimedia platform based on reconfigurable processor for 8K UHD TV," *IEEE Transactions on Consumer Electronics*, vol. 61, no. 4, pp. 516–523, Nov. 2015.
- [2] S. Hara, A. Hanada, I. Masuhara, T. Yamashita, and K. Mitani, "Celebrating the launch of 8k/4k uhdtv satellite broadcasting and progress on full-featured 8k uhdtv in japan," *SMPTE Motion Imaging Journal*, vol. 127, no. 2, pp. 1–8, Mar. 2018.
- [3] D. Zhou, S. Wang, H. Sun, J. Zhou, J. Zhu, Y. Zhao, J. Zhou, S. Zhang, S. Kimura, T. Yoshimura, and S. Goto, "An 8K H.265/HEVC Video Decoder Chip With a New System Pipeline Design," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 1, pp. 113–126, Jan. 2017.
- [4] A. Ichigaya and Y. Nishida, "Required bit rates analysis for a new broadcasting service using hevcc/h.265," *IEEE Transactions on Broadcasting*, vol. 62, no. 2, pp. 417–425, 2016.
- [5] M. Corrêa, B. Zatt, M. Porto, and L. Agostini, "High-throughput hevcc intrapicture prediction hardware design targeting uhd 8k videos," in *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2017, pp. 1–4.
- [6] E. Reuss, "Beyond the limits of visual acuity: The real reason for 4k and 8k image resolution," *SMPTE Motion Imaging Journal*, vol. 126, no. 2, pp. 33–39, Mar. 2017.
- [7] H. Yamashita, H. Aoki, K. Tanioka, T. Mori, and T. Chiba, "Ultra-high definition (8k uhd) endoscope: our first clinical success," *SpringerPlus*, vol. 5, Dec. 2016.
- [8] Y. Aoki, M. Matsuura, T. Chiba, and H. Yamashita, "Effect of an 8K ultra-high-definition television system in a case of laparoscopic gynecologic surgery," *Videosurgery and other miniminvasive techniques*, vol. 12, no. 3, pp. 315–319, 2017.
- [9] H. Yamashita, K. Tanioka, G. Miyake, I. Ota, T. Noda, K. Miyake, and T. Chiba, "8k ultra-high-definition microscopic camera for ophthalmic surgery," *Clinical Ophthalmology*, vol. Volume 12, pp. 1823–1828, Sept. 2018.

- [10] S. Ohigashi, T. Taketa, G. Shimada, K. Kubota, H. Sunagawa, and A. Kishida, "Fruitful first experience with an 8k ultra-high-definition endoscope for laparoscopic colorectal surgery," *Asian Journal of Endoscopic Surgery*, vol. 12, no. 3, pp. 362–365, 2019.
- [11] M. Agiwal, A. Roy, and N. Saxena, "Next generation 5g wireless networks: A comprehensive survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 3, pp. 1617–1655, 2016.
- [12] Statista. Video Streaming (SVoD) - worldwide — Statista Market Forecast. (2020, May 1). [Online]. Available: <https://www.statista.com/outlook/206/100/video-streaming--svod-/worldwide?currency=eur>
- [13] M. Asghar, M. Ghanbari, M. Fleury, and M. Reed, "Sufficient encryption based on entropy coding syntax elements of h.264/svc," *Multimedia Tools and Applications*, vol. 74, July 2014.
- [14] D. Marpe, H. Schwarz, and T. Wiegand, "Context-based adaptive binary arithmetic coding in the h.264/avc video compression standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 620–636, 2003.
- [15] M. U. K. Khan, J. M. Borrmann, L. Bauer, M. Shafique, and J. Henkel, "An h.264 quad-fullhd low-latency intra video encoder," in *2013 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2013, pp. 115–120.
- [16] S. Pinto, J. Pereira, T. Gomes, A. Tavares, and J. Cabral, "LTZVisor: TrustZone is the Key," in *Euromicro Conference on Real-Time Systems*, ser. Leibniz International Proceedings in Informatics, vol. 76. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, pp. 4:1–4:22.
- [17] D. Cerdeira, N. Santos, P. Fonseca, and S. Pinto, "Sok: Understanding the prevailing security vulnerabilities in trustzone-assisted tee systems," in *2020 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2020, pp. 636–652.
- [18] J. Martins, A. Tavares, M. Solieri, M. Bertogna, and S. Pinto, "Bao: A Lightweight Static Partitioning Hypervisor for Modern Multi-Core Embedded Systems," in *Workshop on Next Generation Real-Time Embedded Systems (NG-RES 2020)*, ser. OpenAccess Series in Informatics (OASISs), M. Bertogna and F. Terraneo, Eds., vol. 77. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020, pp. 3:1–3:14. [Online]. Available: <https://drops.dagstuhl.de/opus/volltexte/2020/11779>



Dr. Jorge Cabral received his PhD degree in Electrical Engineering from Imperial College London, UK. His research interests include Embedded Systems, Instrumentation Systems and Micro-electromechanical systems. Currently, he is an Assistant Professor at the Department of Industrial Electronics and Deputy Director of Algoritmi Research Centre. He leads several National and European R&D projects in the field of Cyber-Physical Systems and Embedded Applications. He authored or co-authored around 100 indexed articles and papers. Contact him at jorge.cabral@dei.uminho.pt.



José Dias is Head of Electronics at MOG Technologies, with a strong background in Electronics, Software Development, and Linux. In the last years, he also worked as Software Developer and Field Application Engineer, at MOG Technologies, and as Hardware Developer at XLPartner Technologies. José received his B.Sc in Electrical, Electronic and Communications Engineering from Polytechnic Institute of Viana do Castelo and a certified specialization on Product Design and Development from Polytechnic Institute of Cávado and the Ave. Contact him at jose.dias@mog-technologies.com.



Miguel Costa is a Ph.D. student of Electronics and Computer Engineering at the University of Minho, Portugal. From this same institution, Miguel received an M.Sc. also in this academic topic, with a special focus on embedded systems and information systems and technologies. His research interests focus on artificial intelligence and system-on-chip design. In the last 3 years, he worked as an R&D engineer in an AI-based automotive HMI and 8k UHD video solutions. Contact him at miguel.costa@dei.uminho.pt.



Ricardo Moreira holds an M.Sc. in Electronics and Computer Engineering from the University of Minho, Portugal. His academic education focused on embedded systems and micro/nano technologies, and his M.Sc. thesis addressed HDL design for memory management and video encoding in 8k video systems. His research interests are system-on-chip design and embedded operating systems. Ricardo is an R&D Embedded Systems Engineer at Centro ALGORITMI, in the University of Minho. Contact him at rvmoreira.5@gmail.com



Dr. Sandro Pinto is a Research Scientist and Invited Professor at the University of Minho, Portugal. He holds a Ph.D. in Electronics and Computer Engineering. During his Ph.D., Sandro was a visiting researcher at the Asian Institute of Technology (Thailand), University of Wurzburg (Germany), and Jilin University (China). Sandro has a deep academic background and several years of industry collaboration focusing on operating systems, virtualization, and security for embedded, cyber-physical, and IoT-based systems. He has published several scientific papers in top-tier conferences/journals and is a skilled presenter with speaking experience in several academic and industrial conferences. Contact him at sandro.pinto@dei.uminho.pt.