HIGH-LEVEL SYNTHESIS OF A SINGLE/MULTI-BAND OPTICAL AND SAR IMAGE COMPRESSION AND ENCRYPTION HARDWARE ACCELERATOR

Paolo Motto Ros, Michele Caon, Tiziano Bianchi, Maurizio Martina, Enrico Magli

Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Turin, Italy

ABSTRACT

Transmitting images from earth observation satellites to ground is a major challenge, and a compression/encryption stage is actually mandatory. Development of hardware accelerators is highly recommended, both to relieve the software from such demanding task, and to improve performance, aiming at quasi-real-time data processing. To this end, we discuss the design, development, deployment and test of a FPGA-based accelerator, featuring a lossless and lossy (near-lossless) compression, including the data encryption too. Its architecture is well suited for different image types, including single- and multi-band optical and SAR images and can be fully run-time configurable. Measured performance showed a throughput of 10 Msamples/s, in agreement with related state-of-the-art works, focused on lossless compression only.

***Index Terms*—** High-level synthesis, Hardware accelerator, FPGA, Optical and SAR imaging, lossy compression, CCSDS-123, Earth observation

1. INTRODUCTION

Satellite Earth Observation (EO) is nowadays rapidly gaining traction and receiving significant ever increasing attention thanks to emergent opportunities and challenges to be exploited, such as environment monitoring and civilian security. To this end, the latency in delivering EO products to the ground surely is among the first key performance indicators. In [1] a latency below five minutes (below one minute in some scenarios) has been demonstrated to be an achievable goal. Among the products to be transmitted to the ground, (multi-band) optical and SAR images represent the most challenging ones, due to their large size and the limited available communication bandwidth. Therefore, a specific compression and encryption sub-system is needed.

The aim of this work is to explore the feasibility, from the design to the deployment on the target platform, and initial performance assessment of a hardware accelerator for compressing and encrypting single- and multi-band optical images and SAR data. The implemented compression algorithm is derived from the CCSDS 123.0-B-2 recommended standard [2], with a more hardware-friendly pre-quantization stage in front of the prediction loop rather than within the loop itself [3]. Encryption has been implemented by randomly flipping the sign of the prediction residuals. To this end, a pseudo-random generator based on the very efficient Keccak cryptographic functions [4] has been integrated.

Considering the specifications of the whole data processing chain for next-generation EO satellites [1], a throughput of 10 Msamples/s has been deemed to be the target performance. This is in agreement with state-of-the-art related works, where, for similar scenarios and designs, the reported throughput is in the range 10–15 Msamples/s [5, 6, 7]. Reported maximum throughput can be up to 150 Msamples/s [6], but it strongly depends on input/output data ordering, implemented features, configurability. Data ordering is a major point, with Band-Interleaved by Pixel (BIP) images (whisk-broom scanners) enabling near one order of magnitude better throughput than Band-Interleaved by Line (BIL, push-broom sensors) or Band Sequential (BSQ) ones (everything else equal) [6]. Similar discrepancies are reported for similar compression algorithms, e.g., Fast Lossless (FL), too [8, 7]. The bottleneck is due to the weights update operations (strictly depending on z): due to the different sample ordering, BIL and BSQ architectures can not exploit pipeline optimizations enabled by a BIP design. This is so critical to eventually pose a constraint on the minimum allowed number of bands in input image [9]. Other throughput bottlenecks are the quantization stage (required in case of lossy compression) [3], and the flexibility provided by enabling run-time configuration (e.g., in [9, 6] image size is fixed at synthesis time to optimize resource usage and resulting performance).

With respect to the related works, in our design we introduce a pre-quantization stage, enabling a lossless and lossy (or, better, near-lossless, i.e., bounding the maximum error between the original and the reconstructed image) compression in the same core, and the encryption support directly integrated in the compression process. All the parameters (including image size, number of bands, dynamic range, encryption key, and initialization vector) can be fully user-configurable at run-time. Performance will be presented with respect to the quantization delta; differently from the related works, both single and multi-band optical images, as well as raw SAR data, will be used as a part of the test/validation dataset. In particular, the need of processing optical images and raw SAR data, both “single-band arrays”, renders the optimization opportunities in BIP architectures ineffective; therefore, an inherent BIL architecture (better suited for the involved sensors) will be developed.

2. IMPLEMENTATION

The resulting compression/encryption algorithm has been first intensively and extensively tested, characterized, and validated through a reference implementation, written in the common C language. It has clearly shown to be a good candidate for a full hardware implementation: considering the complexity of developing and validating a brand new implementation (i.e., a completely new hardware implementation designed from scratch), a High-Level Synthesis (HLS) approach has been preferred. Target platform features a Xilinx Zynq US+ ZU19EG, a Multiprocessor System-on-a-Chip (MPSoC) integrating a Processing System (PS) based on a quad-core ARM Cortex-A53 and a Programmable Logic (PL) based on the Xilinx UltraScale FPGA architecture in a single device.

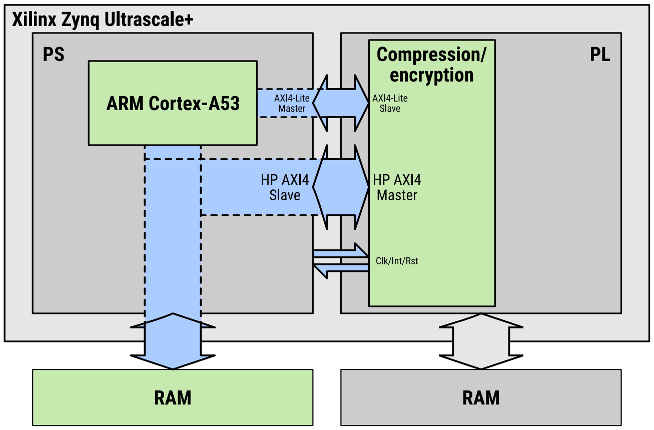

Fig. 1. Overview of the interfaces of the compression/encryption hardware accelerator implemented on the target platform.

The compression/encryption hardware accelerator, from the interface point of view, has a straightforward structure (see Figure 1):

- One (AXI4) master interface (128 bit data width) for reading/writing input/output images directly (acting as a DMA peripheral) from/to the external PS memory.
- One slave (AXI4-Lite) interface to both setup/configure the compression/encryption operations and to control the accelerator itself.
- Few standard control signals: clock, reset, interrupt.

From the software point of view, considering the custom GNU/Linux OS available, the accelerator is viewed as a Userspace I/O (UIO) device, accessible through the standard device subsystem. The use of the hardware interrupt enables a non-blocking management of the compression/encryption implemented in hardware.

As a general rule, all the options of the reference implementation have been kept configurable at run-time; optionally, some of them can be fixed at synthesis time so to allow the synthesis tools to further optimize the resulting architecture. Other parameters can only be set at synthesis time; these are strictly related to hardware details, enabling to eventually implement the designed solution on platforms other than the intended target one (e.g., significantly smaller FPGAs).

Target clock has been set to 200 MHz. Other clock frequencies have been investigated, including 100 MHz, 250 MHz, and 300 MHz. While the last two could be a viable solution, they would offer marginal performance increase at the cost of more critical timings. At the same time, the upgrade from 100 MHz (initial target clock) to 200 MHz has been deemed worthy from the performance perspective without sacrificing tolerance on timings.

Table 1. Resource utilization of the implemented compression/encryption accelerator on the target ZU19EG platform, both at 200 MHz and 100 MHz clock frequency.

Resource	Available	200 MHz	100 MHz
CLB	65340	32915 (50 %)	34487 (53 %)
CLB LUT	522720	196087 (38 %)	202517 (39 %)
CLB FF	1045440	149406 (14 %)	138202 (13 %)
BRAM (tiles)	984	268 (27 %)	253 (26 %)
DSP	1968	330 (17 %)	330 (17 %)

Table 1 reports the resource occupation of the final design first synthesized and then integrated, synthesized, and implemented in the reference hardware design. Reported numbers are the final ones (not estimates) of the designed compression/encryption accelerator only. Two cases have been implemented and tested, i.e., for two different clock frequencies, in order to explore different trade-off and then select the most suited. The source code and the configuration of the designed core is exactly the same in both cases. Considering the implementation running at 200MHz (recommended), overall Configurable Logic Blocks (CLBs) occupation shows a relative large usage (near half of the available resources in the PL), but it still allows for integrating other (smaller) cores. The higher percentage of used CLB LUTs (38 %) with respect to the CLB FFs one (14 %) reflects the computational complexity involved. This is confirmed by a significant usage of DSP blocks (17 %) and by a relative low usage of Block RAMs (BRAMs, 27 %), especially considering that the core has been configured to process images with a size up to 32768 × 16384 and up to 8 bands (16 bit per pixel). By comparing the two solutions, at 200 MHz and 100 MHz, we can see that they show similar numbers, although an increase in the overall CLB occupation (+5 %) and a decrease in the BRAM one

(-6 %). This is reflected in a different balance between CLB LUTs and CLB FFs, showing that, given the same design to be synthesized, targeting 100 MHz, allows, as expected, to "pack" more combinational logic between two successive register stages, thereby enabling shorter pipelines.

3. RESULTS

The developed compression/encryption accelerator has been then deployed on the target platform, and by means of a test program (same interface as the reference one), finally validated through the same test suite used all along the design stages. Runtime performance has been then measured, in software, taking into account the memory transfers (operated by the core) too. All timings thus refer to elapsed time between the start of the compression, with the entire image in the PS DDR external memory, and its end, with the entire encrypted (and compressed) image in the PS DDR memory again.

The overall image test set is complete and sound, entirely representing the intended case scenarios. It can be divided into three main subsets: single-band optical (2 up to $13100 \times 6000 \times 1$), multi-band optical (96 up to $1119 \times 639 \times 5$), and SAR data (4 up to $10254 \times 3425 \times 1$) images. Due to the adaptive nature of the compression/encryption algorithm, indeed, different kind of images could lead, e.g., to different throughputs. Main figure of merit for the assessment of the performance has been deemed to be the average time taken to compress each pixel, measured in ns/pixel, which is independent of the image size. Performance results and comparison are showed as a function of the compression quantization delta, as it is the main parameter controlling output data size (which in turn impacts on the image transmission performance).

Table 2. Median values (ns/pixel) of the compression/encryption running in hardware (200 MHz clock).

Lossless	Quantization delta	Optical images	Multi-band images	SAR raw data
Yes	0	125.6	123.1	128.6
Near	1	108.1	112.2	126.1
Near	2	103.8	105.8	124.5
Near	4	100.4	99.8	121.3
Near	8	98.3	97.7	115.2
Near	16	98.4	97.4	102.0
Near	32	97.1	97.2	100.7

Figure 2 shows box and whiskers plot of all the results summarized, organized per quantization delta (as defined in [3]) and then per image set; Table 2 reports the median values only. By comparing these values, we can further confirm that the estimate of the compression and encryption of an optical multi-band image can be in the same range of a single-band optical one (relative difference in the range between -2 % and 3.8 %), while the same operations on a SAR image can be slower (relative difference in the range between 2.4 % and 20.9 %). Overall median processing time is about 100 ns/pixel, resulting in a median throughput of about 10 Msamples/s (20 MB/s).

Figure 3 shows the box and whiskers plots of the measured performance with the accelerator at 200 MHz (recommended) and 100 MHz. It has to be remarked that performance is not halved, i.e., is not proportional to the ratio between the two clocks. During the development stage this kind of "law of diminishing returns" was even more evident as the clock was further increased to 250 MHz or 300 MHz. Considering the added difficulty to obtain then good timings results, made 200 MHz be the best trade-off between performance, resources occupation, and reliability.

4. CONCLUSION

In this work we presented the design and development of a compression and encryption hardware accelerator for (multi-band) optical images and SAR data. Leveraging a HLS design approach, it has been successfully deployed on a Xilinx Zynq platform, along with the full software stack. Measured performance on datasets comprising different image types and sizes, showed a throughput of about 10 Msamples/s (20 MB/s) with lossless and, most importantly, lossy (near-lossless) compression options. The throughput is in agreement with related state-of-the-art works, focused on lossless compression only, and demonstrates the viability of this approach for achieving very low latency in the delivery of EO products to ground.

5. ACKNOWLEDGEMENTS

This work has been supported by the project EO-ALERT, European Union's Horizon 2020 research and innovation programme under grant agreement No 776311.

6. REFERENCES

[1] M. Kerr, S. Tonetti, S. Cornara, J. I. Bravo, R. Hinz, A. Latorre, F. Membibre, C. Solimini, S. Wiehle, H. Breit, B. Tings, O. Koudelka, F. Teschl, E. Magli, T. Bianchi, A. Migliorati, P. Motto Ros, M. Caon, R. Freddi, M. Benetti, F. Milani, G. Curci, S. Fraile, L. Garcia, C. Marcos, and A. Fiengo, "EO-ALERT: A novel architecture for the next generation of earth observation satellites supporting rapid civil alerts," in *71st International Astronautical Congress (IAC)*, Oct. 2020.

[2] The Consultive Committee for Space Data Systems (CCSDS), *Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression*, vol. Blue Book, Feb. 2019, Recommended Standard CCSDS 123.0-B-2.

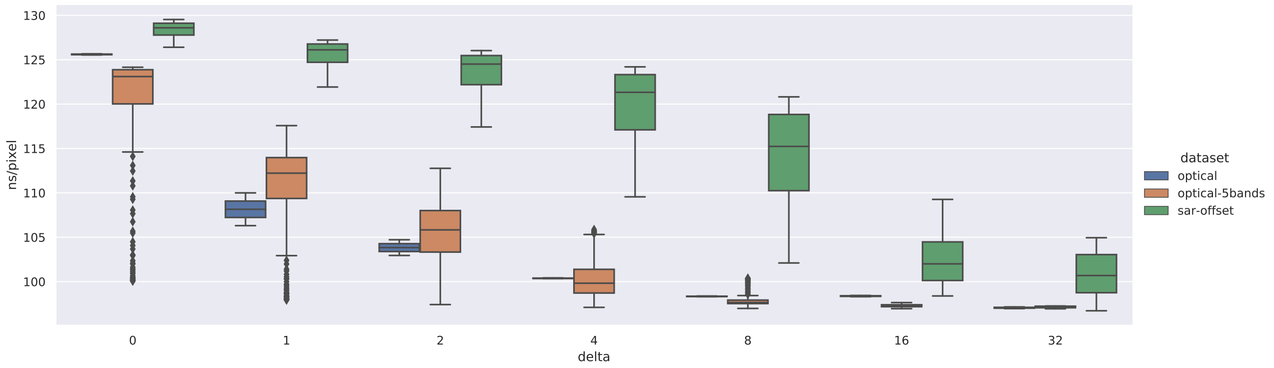

Fig. 2. Performance of the compression/encryption running in hardware (200 MHz clock).

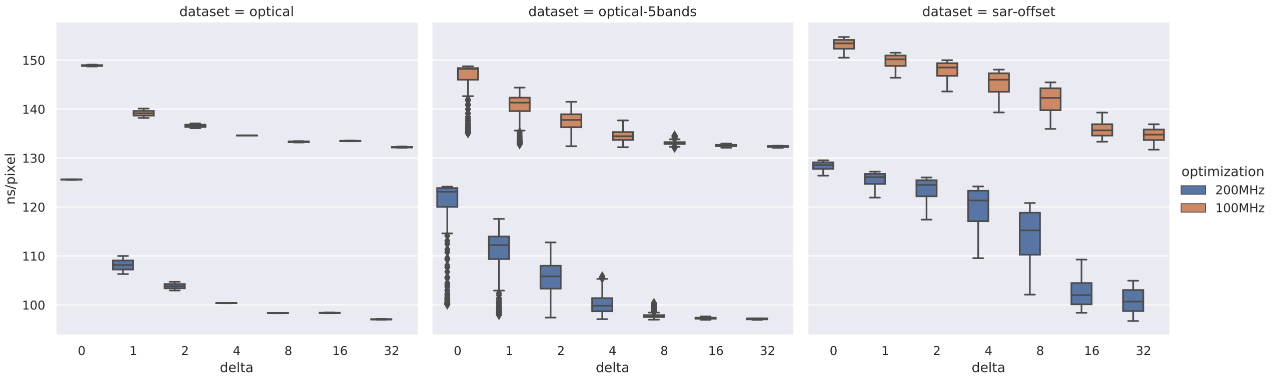

Fig. 3. Performance of the compression/encryption running in hardware at 200 MHz and 100 MHz.

[3] D. Valsesia and E. Magli, “High-throughput onboard hyperspectral image compression with ground-based cnn reconstruction,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9544–9553, 2019.

[4] A. Migliorati, T. Bianchi, and E. Magli, “Selective encryption in the CCSDS standard for lossless and near-lossless multispectral and hyperspectral image compression,” in *Image and Signal Processing for Remote Sensing XXVI*, Lorenzo Bruzzone, Francesca Bovolo, and Emanuele Santi, Eds. International Society for Optics and Photonics, 2020, vol. 11533, pp. 221–228, SPIE.

[5] L. Santos, L. Berrojo, J. Moreno, J. F. López, and R. Sarmiento, “Multispectral and hyperspectral lossless compressor for space applications (HyLoC): A low-complexity FPGA implementation of the CCSDS 123 standard,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 9, no. 2, pp. 757–770, 2016.

[6] Y. Barrios, A. J. Sánchez, L. Santos, and R. Sarmiento, “SHyLoC 2.0: A versatile hardware solution for on-board data and hyperspectral image compression on future space missions,” *IEEE Access*, vol. 8, pp. 54269–54287, 2020.

[7] D. Keymeulen, N. Aranki, A. Bakhshi, Huy Luong, C. Sarture, and D. Dolman, “Airborne demonstration of FPGA implementation of fast lossless hyperspectral data compression system,” in *2014 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, 2014, pp. 278–284.

[8] N. Aranki, D. Keymeulen, A. Bakhshi, and M. Klimesh, “Hardware implementation of lossless adaptive and scalable hyperspectral data compression for space,” in *2009 NASA/ESA Conference on Adaptive Hardware and Systems*, 2009, pp. 315–322.

[9] J. Fjeldtvedt, M. Orlandić, and T. A. Johansen, “An efficient real-time FPGA implementation of the CCSDS-123 compression standard for hyperspectral images,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 10, pp. 3841–3852, 2018.