

SUMMARY OF MULTIPLE BENCHMARKS ON THE HIGH PERFORMANCE DATA PROCESSOR (HPDP)

**Ioannis Katelouzos⁽¹⁾, Tilemachos Tsiapras⁽¹⁾, Jacques Monnier⁽¹⁾, Kostas Makris⁽¹⁾,
Daniel Bretz⁽²⁾, Simon Klugseder⁽²⁾, Antonios Tavoularis⁽³⁾, Gianluca Furano⁽³⁾, Tim Helfers⁽²⁾,
Constantin Papadas⁽¹⁾, Laurent Hili⁽³⁾**

⁽¹⁾ *Integrated Systems Development SA (ISD SA), 15125 Maroussi, Greece*

⁽²⁾ *Airbus Defence and Space GmbH, 82024, Taufkirchen, Germany*

⁽³⁾ *ESTEC, 2201 AZ Noordwijk, Netherlands*

ABSTRACT

The increasing use of new earth-observation and communication technologies, the augmenting life time far beyond 10 years coupled with rapidly changing customer needs demand for a high performance and flexible processing technology for on-board data processing. The algorithmic and processing requirements for such processing are of a magnitude larger than those which could be successfully handled by classical processors. Therefore, a new High-Performance Data Processor (HPDP) has been developed jointly by ISD (Greece) and ADS (Airbus Defence and Space, Ottobrunn, Germany).

After a brief architecture description, this paper presents typical HPDP applications

Keywords

High Performance Processor, Moon Asteroid strike Detection, Vessel Detection, AES, LDPC, Image Compression

1. INTRODUCTION

Reconfigurable hardware technology has the capability to alleviate drawbacks of dedicated hardware as the payload functionality can be reprogrammed. This means that the adaptability of the satellite to new standards is guaranteed. The available flexibility allows the manipulation of routing and data processing functions, and therefore provides a way to alleviate the rigidity of space suitable processors available today.

The big advantage with respect to FPGAs is the reduced power consumption and the fact that the user should not care about the timing enclosure.

ISD S.A. jointly with Airbus Defence and Space GmbH have undertaken the implementation of the commercially available reconfigurable array processor

IP (XPP from the company PACT XPP Technologies [1]) in a radiation hardened technology. This development is supported by the Greek incentive scheme, the DLR and ESA and is named "High Performance Data Processor" (HPDP).

2. REQUIREMENT FOR A DATA PROCESSOR IN SPACE

In contrast to consumer technology, there are strong constraints due to space semiconductor technologies that severely limit the performance for regular CPU architectures. These include the radiation hardened process and various design aspects, resulting in much lower operational clock frequencies (in the order of 50-300 MHz for digital circuits).

The selection process for a data processor suitable for space is of strategic importance as the turnaround cycle for such a processor is around 10 years. The assessment has to be performed taking the underlying semiconductor technology process into account.

The requirements of a future processor can be summarized as follows:

- Processing Performance: Minimum 1 GigaOps/s with support of floating point capability
- Throughput Performance for IO: Minimum 4 Gbps
- Radiation Hardness: Total dose 20-100 Krad, Latch-up immunity better than 80 MeV/mg/cm²
- Tolerance against Single Event Effects
- Re-programmability required (in orbit, only maintainable per remote access; ability to serve multiple projects)
- Power consumption to be as low as possible (because of power supply and power dissipation constraints)

- Long Term availability without regulations (preferably European source)
- Packages must be suited to survive thermal cycling (during orbiting) and mechanical stresses (during launch phase)
- Space qualification grades QML-Q or QML-V

3. THE HPDP ARCHITECTURE

The HPDP architecture integrates the XPP reconfigurable processing core IP, space suitable peripherals and memory interfaces. Fig.1 depicts the major building blocks of the HPDP architecture.

The HPDP core i.e. XPP IP, consists of 40 ALU-PAEs and 16 RAM-PAEs and typically processes high bandwidth data streams. Two FNC-PAEs (Function PAEs) are coupled to the array's communication channels via versatile crossbars. The FNC-PAEs perform control flow tasks, sequential algorithms and system management. They communicate directly between each other through vertical communication channels.

Several DMA controllers transfer data in the background. The FNC-PAE controls the DMA operation. FIFOs uncouple the DMA channels from system RAM bursts, high speed interfaces and potentially stalled pipelines within the XPP. The specialized DMA controllers (4D-DMA) generate 4-dimensional access patterns for the on-chip buffers (X-RAM) or the off-chip memory (SRAM or SDRAM memory banks).

The HPDP architecture includes the following memory ports:

- Configuration Memory: The configuration memory is made up of PROM/EEPROM or SRAM devices and stores the code (boot code and/or application code).
- Data Memory: The memory port can be connected with either a SRAM or a SDRAM memory bank

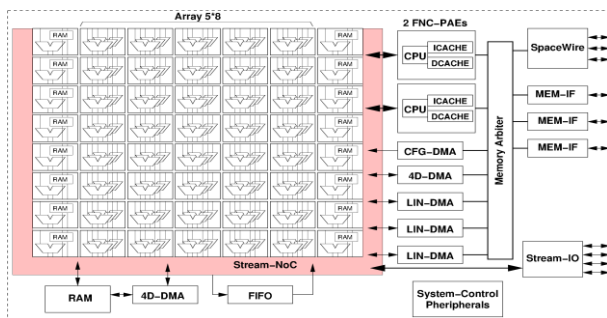


Figure 1. Overview of HPDP architecture

The communication with the external controller is based on the widely used SpaceWire standard.

In summary the following key features are provided by the HPDP architecture and by the realization onto the ST65SPACE technology from STM:

- Based on XPP III Array Processor from PactXPP Technologies providing 40 Giga operations per second (End-of-Life)
 - 40 ALU Processing Array Elements (16-bit) running with 250MHz each
 - 2 Harvard type VLIW 16-bit processor cores (FNCs) running at 125MHz
 - 256 Kbit high speed on-chip RAM
- 4 Mbyte on-chip SRAM and Memory interface for external SRAM or SDRAM devices with EDAC
- 4 x 1.1 Gbit/s Streaming Ports compatible with HSSL
- 3 SpaceWire interfaces for monitor and control operating at 100 Mbps on each channel with routing capability

The architecture includes standard space relevant features like:

- ECC/EDAC and scrubbing function in the external memory interfaces
- Error protection in on-chip memory
- Triple mode redundancy reset and clock logic
- Clock synchronous design
- Space relevant control interface (SpaceWire)

4. REALIZATION

Currently the chip has been fabricated with the ST 65nm technology and sizes 98mm². It is packaged in a dedicated ceramic 625-pin, four deck PGA package design jointly by ISD and Kyocera (D).

The packaging of the dices has been done by SERMA.

5. DEMOKIT FOR HPDP ARCHITECTURE

Being tailored for ground-based laboratory purposes, it enables convenient and full access to all relevant features and signals. Nevertheless, transfer of the basic design to a space-compatible configuration can be achieved with relatively low effort. The full kit consists of four identical boards, each equipped with the key components HPDP chip, SDRAM, EEPROM, SpaceWire, and Channel-Link interfaces. Just like the HPDP chip itself is highly configurable, the demo kit continues this concept as it allows free combinations of up to four boards with direct communication links

between each chip. At the same time a SpaceWire chain, direct SpaceWire and JTAG access to each chip, and two bi-directional Channel-Link interfaces are available for connections to external control interfaces or data sources and receivers, respectively. This allows for versatile test scenarios, both aiming at inter-chip connectivity and parallel real-time data-processing in pipeline fashion. With not much more than a single power supply to get a full stack of four HPDP boards up and running, the demo-kit will be transportable and easy to use. This will make it a valuable tool for the investigation of HPDP use cases and prototyping/optimization of full applications, independent of sophisticated test environments.

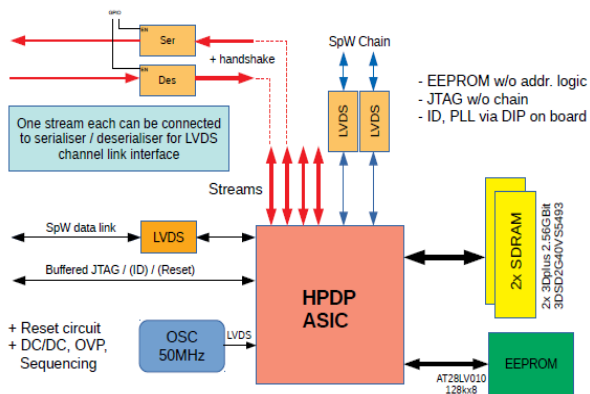


Figure 2. HPDP Demokit Architecture




Figure 3. HPDP Demokit Board

6. HPDP DEVICE STATUS

The HPDP Chip is operating in Airbus Ottobrunn and ISD Athens laboratories. The device is fully functional and detailed power consumption measurements have shown that the device consumes nominally 1.65 W. This emphasizes the low power characteristics of the device

and emphasizes the idea of an eventual re-packaging in a plastic package (lower mass and cost). The device has been subjected to TID campaign and is has been demonstrated that sustains 300Krad without apparent degradation effects. Moreover, extensive heavy ion campaign have shown it is latch-up free up to 72 MeV at 85°C and immune to undesired phenomena induced by the radiations.

7. APPLICATION DEMONSTRATIONS

Several applications have been implemented on the HPDP device by using the simulator to demonstrate its suitability. In a second step the demo-kit has been used in order to execute the application on the device itself and confirm key performance figures.

7.1. Moon Asteroid Strike Detection

The Moon Asteroid Strike Detection has been developed in the frame of a GSTP De-Risk activity.

The idea is to count collision flashes in the moon surface and from their intensity and duration to extract meaningful information about the collision energy and the mass of the colliding asteroid.

Currently we use sequences of images taken from the NELIOTA database. These are sets of 15 images taken around the impact at a frame rate of 30fps. The images have been captured by the Kryoneri Observatory in Greece and they have a resolution of 1080x1280 pixels and 8bit depth per pixel. In future, the algorithm will be modified in order to exploit 12bit images.

The algorithm for the asteroid strikes detection, is divided into two main parts. The first one concerns the temporal noise reduction and the second part is the detection of the strike itself. The flow of the algorithm is shown in Fig.4.




Figure 4. MAS Algorithm diagram




Figure 5. Asteroid impact on true images

The temporal noise reduction is done through a running average over the last 8 images which constructs a background image free of any temporal noise.

The detection of the events is done by taking the difference between the current image and the background that has been calculated as described above. All the differences between the two images are recorded along with the number of the different pixels. Basically, per event the following characteristics are extracted:

- the intensity of the impact over time (multiple of the inter-frame period),
- the size of the impact (in terms of pixels) over time and
- the coordinates of impact.

The algorithm has also the capability to distinguish between true strikes and moving objects (i.e other satellites) using the concept of a “moving strike”.

In the current implementation, the images are stored into the SDRAM of the system which presents a bottleneck on the memory’s throughput. Despite this fact, the performance reached is 116fps at the expense of 1.65W of power consumption. In an eventual future implementation, along with the increase of the bit depth per pixel, a faster DDR interface will be used, which will permit to both increase the performance and the number of the images used for the averaged background.

7.2. Vessel Detection and Identification

During the same De-Risk activity, another algorithm has been developed, where the position and the size of sea vessels are detected from satellite images. In a differed time, the output of the algorithm is compared with AIS data in order to detect suspect vessels.

The algorithm consists of a Sobel convolution step in order to reduce the noise of the image, followed by Kernel convolution in order to detect the actual vessels. The algorithm proves the suitability of the HPDP to

implement image convolution and the actual steps are shown in Fig. 6 and Fig. 7




Figure 6. SOBEL Flow




Figure 7. Kernel convolution Flow

The Sobel filter is actually an edge detection filter, which amplifies the various features of the image and reduces the background noise. The Sobel operator consists of a pair of 3x3 convolution kernels and is designed to perform 2D spatial gradient over an image.

For the Kernel convolution, 6 kernels of 21x21 pixels each are applied to the image and after comparison with thresholds the detected vessels are reported.

This particular implementation is optimized for vessels that have their sizes comparable to the kernels in respect to pixels. In order to make this algorithm more generic, the initial images could be scaled and reprocessed. This generates coordinates for various sizes of vessels and in the end the results are merged.




Figure 8. The 6 kernels for the vessel detection

After the convolution of the image through the kernels, a mask is constructed that represents the position and the sizes of the vessels as shown in Fig.9.




Figure 9. Image Masks Example

Finally, the positions of the vessels are compared with AIS (Automatic Identification System) data and suspect vessels are flagged.

This implementation can process HD images at a rate of 9.6fps at the expense of 1.65W power consumption and exhibits an overall performance of 60%.

Alternatively, another version of the algorithm has been developed which instead of using kernels, uses a neural network constructed and trained with the Tensor flow.

The network consists of 22000 parameters and 5 layers, 3 of which are convolutional and the last 2 are dense. The input to the network are 64x64 images which represent a sliding window of the starting image and 2 outputs providing the confidence and the presence or absence of a vessel in the input. The database used consists of 6000 labeled examples taken from true satellite images from various heights and conditions. The main architecture of the network can be seen in Fig. 10.

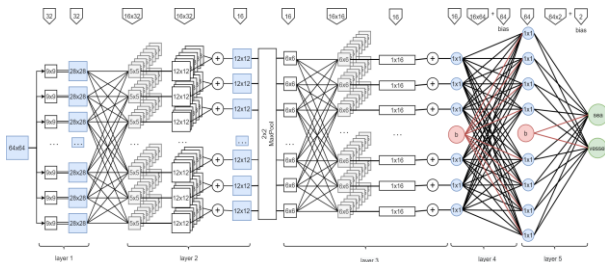


Figure 10. Neural network architecture

The results of the network on the framework using floating point parameters and operations, yield 96% of successful calls. As the systolic array contains mainly ALUs of 16bits, in order to avoid operations overflow, we need to use 8bit integer values. Simulations have shown that we can reach a 88% of correct calls under that assumption.




Figure 11. From left to right: Original image, Kernel convolution, Neural Network

A further comparison has been made between the two algorithms used, that is between the kernel convolution and the use of the Neural network as shown in Fig.11. It is clearly demonstrated that for the same power consumption the neural network implementation outperforms the kernel convolution implementation. Besides, it demonstrates the adequacy of the HPDP device for executing machine learning implementations.

7.3. Data Encryption (AES 256)

Another type of algorithm implemented on the HPDP, is an AES encryption, with key length of 256b and 128b. For the 256b version, two flavors have been implemented, one featuring the CBC (Cipher Block Chaining) mode and the other without it.

First the key expansion is calculated in the FNC0 and is passed via a FIFO to the array. The key expansion process generates the expanded version of the 256b key which is used for the actual encryption. The expanded key is used repeatedly throughout the encryption process and it is stored in the FIFO of the array in order to reduce additional delays. The data to be encrypted are fetched to the incoming stream by the DMA in groups of 4 bytes. From the other side, the cypher sends data to the on-chip SRAM in packets of 4 bytes too. The basic process of the encoding can be seen in Fig. 12.




Figure 12. AES encoding process

Two instances of the AES256 without CBC IP (resp. One instance of the AES256 IP with CDC) can fit in the array giving a total throughput of 11.7MB/s (resp. 5MB/s) at the expense of 1.65W power consumption. Likewise, two instances of the 128it can fit in the array delivering 16.2MB/s encrypted data for the same power consumption.

7.4. Image Compression (CCSDS123)

In the context of the EU H2020-funded HI-SIDE project (776151), the CCSDS 123.0-B-2 image compression algorithm has been ported on HPDP by ISD.

The CCSDS123, is a multispectral standard that is based on predictive methods for the compression part. The encoding is a hybrid between sample-adaptive and block-adaptive approaches. This method is more suitable for near-lossless and lossless compression and in our case, we have implemented the Lossless flavor of the standard.

The implementation of the entire standard was not possible to fit into one configuration of the array, this is why there are two configurations, one for the prediction operations and one for the encoding. The reconfiguration of the array can be done on-the-fly with a very small-time-overhead in the order of 40us per reconfiguration. Overall, the attained throughput is 1Gb/s image data consumption and the power consumption of the device is 1.65W with no external processing capability required.

In Fig. 13 and Tab. 1 there is a presentation of the figures we have achieved by parallelizing the CCSDS123 algorithm and mapping it into the HPDP array.

It is obvious the relation between the performance and the number of bands and also the clear improvement over the serial version of the algorithm.




Figure 13. Performance Gain in relation to number of bands

Table 1. Performance results on various test cases

Test	Nx	Ny	Nz	PC runs (time)	HPDH runs (time)	PC runs (bits/s)	HPDP runs (bits/s)
0	100	1000	3	122.60 ms	19.20 ms	19.57 Mb/s	124.83 Mb/s
1	100	1000	9	367.70 ms	19.22 ms	19.58 Mb/s	374.49 Mb/s
2	100	1000	18	809.57 ms	22.42ms	17.78 Mb/s	642.11 Mb/s
3	100	1000	24	1.025 s	24.82ms	18.73 Mb/s	773.38 Mb/s
4	100	1000	36	1.572 s	30.02ms	18.32 Mb/s	959.16 Mb/s
5	100	1000	45	1.831 s	37.22 ms	19.65 Mb/s	967.06 Mb/s
6	100	1000	72	2.831 s	58.82 ms	20.34 Mb/s	979.15 Mb/s
7	512	2048	45	18.46 s	386.30 ms	20.44 Mb/s	977.18 Mb/s

Currently ISD is working on similar lossless algorithms featuring low entropy encoder targeting the TRUTHS mission

8. CONCLUSION

The High-Performance Data Processor for Space Applications has been developed thanks the support from ESA, DLR and the Greek Incentive Scheme. Typical applications have been implemented by using the whole development flow and then have been validated on the demo-kit. It has been demonstrated that the HPDP device is an ideal platform for the on-board execution of algorithms featuring parallel code execution at the expense of very low power consumption. Besides it has been demonstrated that the HPDP is adequate for executing machine learning applications.

REFERENCES

- [1] Baumgarte, V. et al.: PACT XPP—A Self-Reconfigurable Data Processing Architecture. The Journal of Super-computing, Vol. 26, Issue 2, Sept. 2003, pages 167-184.
- [2] Acher, G. et al.: TU München— A High Performance Reliable Dataflow based Processor for Space Applications, Computing Frontiers'13, May 14–16, 2013, Ischia, Italy.
- [3] The Consultative Committee for Space Data Systems CCSDS. Low-Complexity Lossless Multispectral and Hyperspectral Image Compression, Recommended Standard CCSDS 123.0-B-2, Blue Book, 2019.
- [4] ESA-SRE-NELIOTA-TN-0001 "NELIOTA Detection Algorithms" September 2015