


HPDP-40 SUMMARY OF MULTIPLE BENCHMARKS ON THE HIGH PERFORMANCE DATA PROCESSOR (HPDP)



Ioannis Katelouzos, Tilemachos Tsiapras, Jacques Monnier, Kostas Makris, Daniel Bretz, Simon Klugseder, Antonios Tavoularis, Gianluca Furano, Tim Helfers, Constantin Papadas and Laurent Hilli

Contents

- On the HPDP40 device – current status
- Benchmarks Overview – four applications
- Conclusions

HPDP40 device (1/2)

HPDP40 implements the eXtreme Processing Platform (**XPP**), a runtime reconfigurable data processing engine developed by PACT XPP Technologies AG.

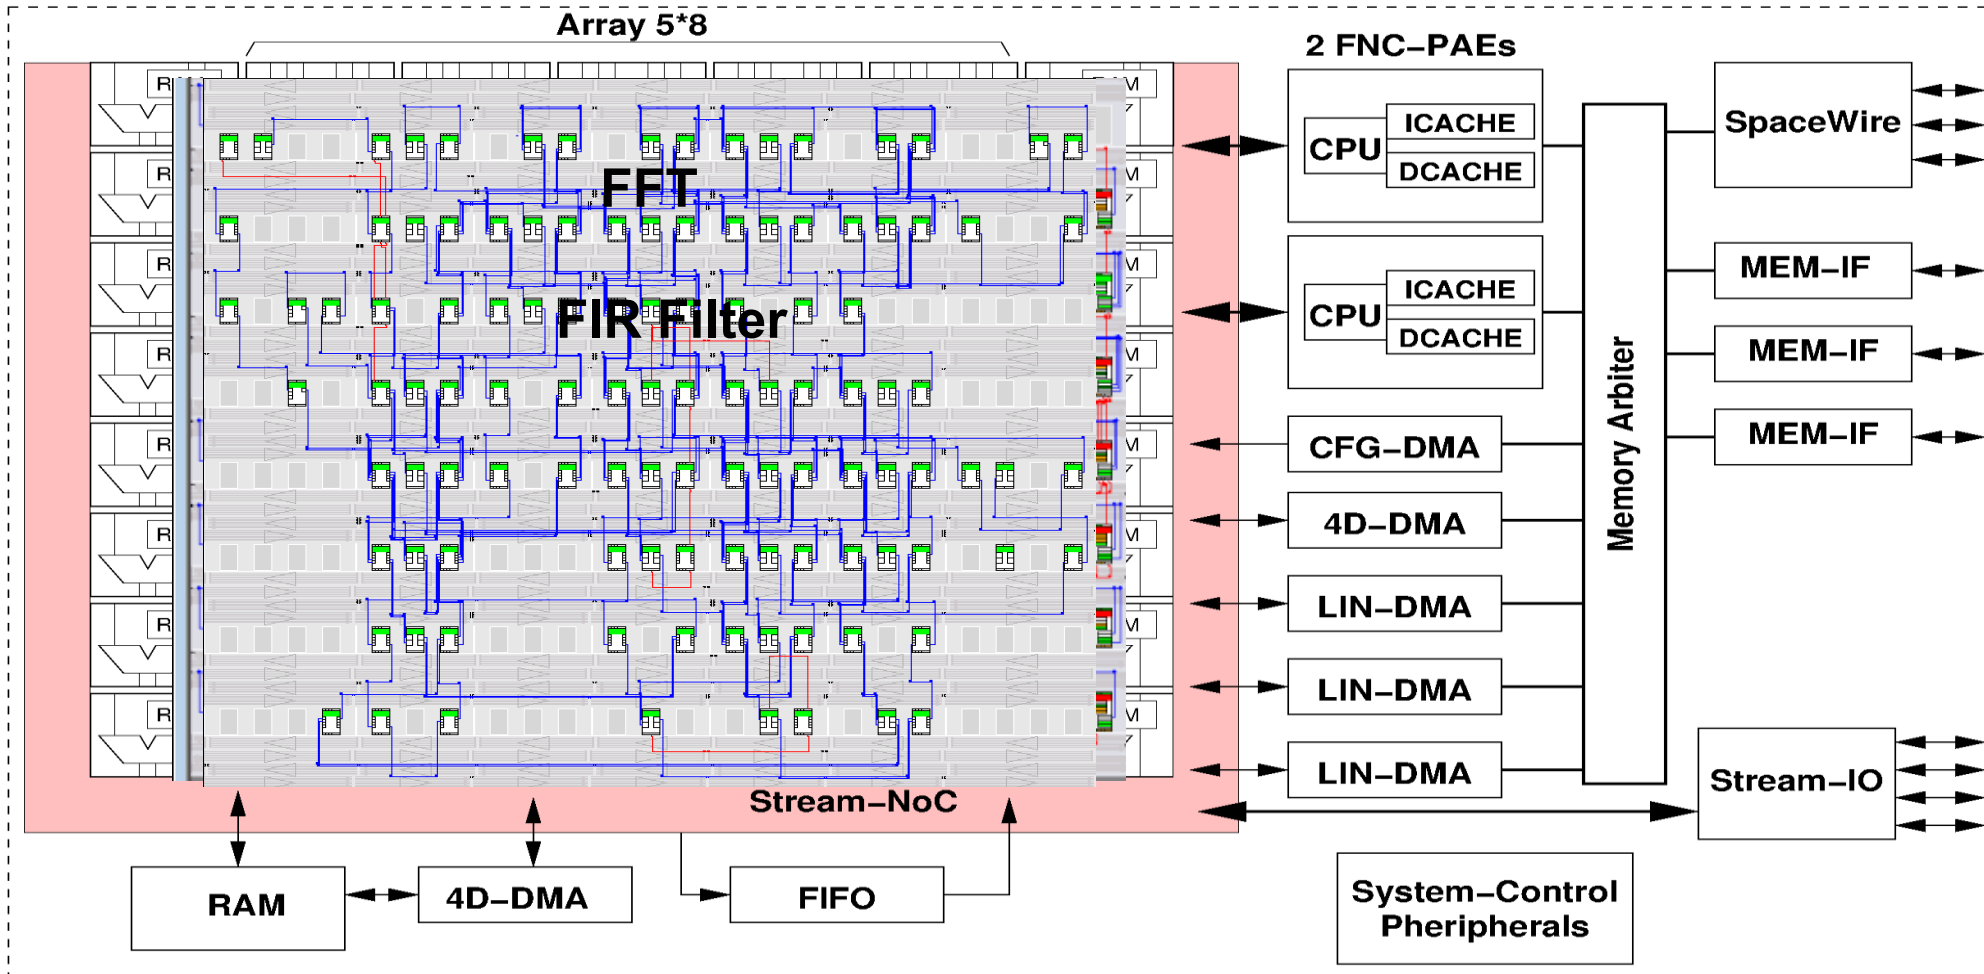
XPP configuration:

- 40 ALU Processing Array Elements (16b) running at 250MHz
- 16 columns RAM blocks for memory
- 2 VLIW processor cores (FNC PAE) running at 125MHz
- Connected by a reconfigurable data and event network

The device provides:

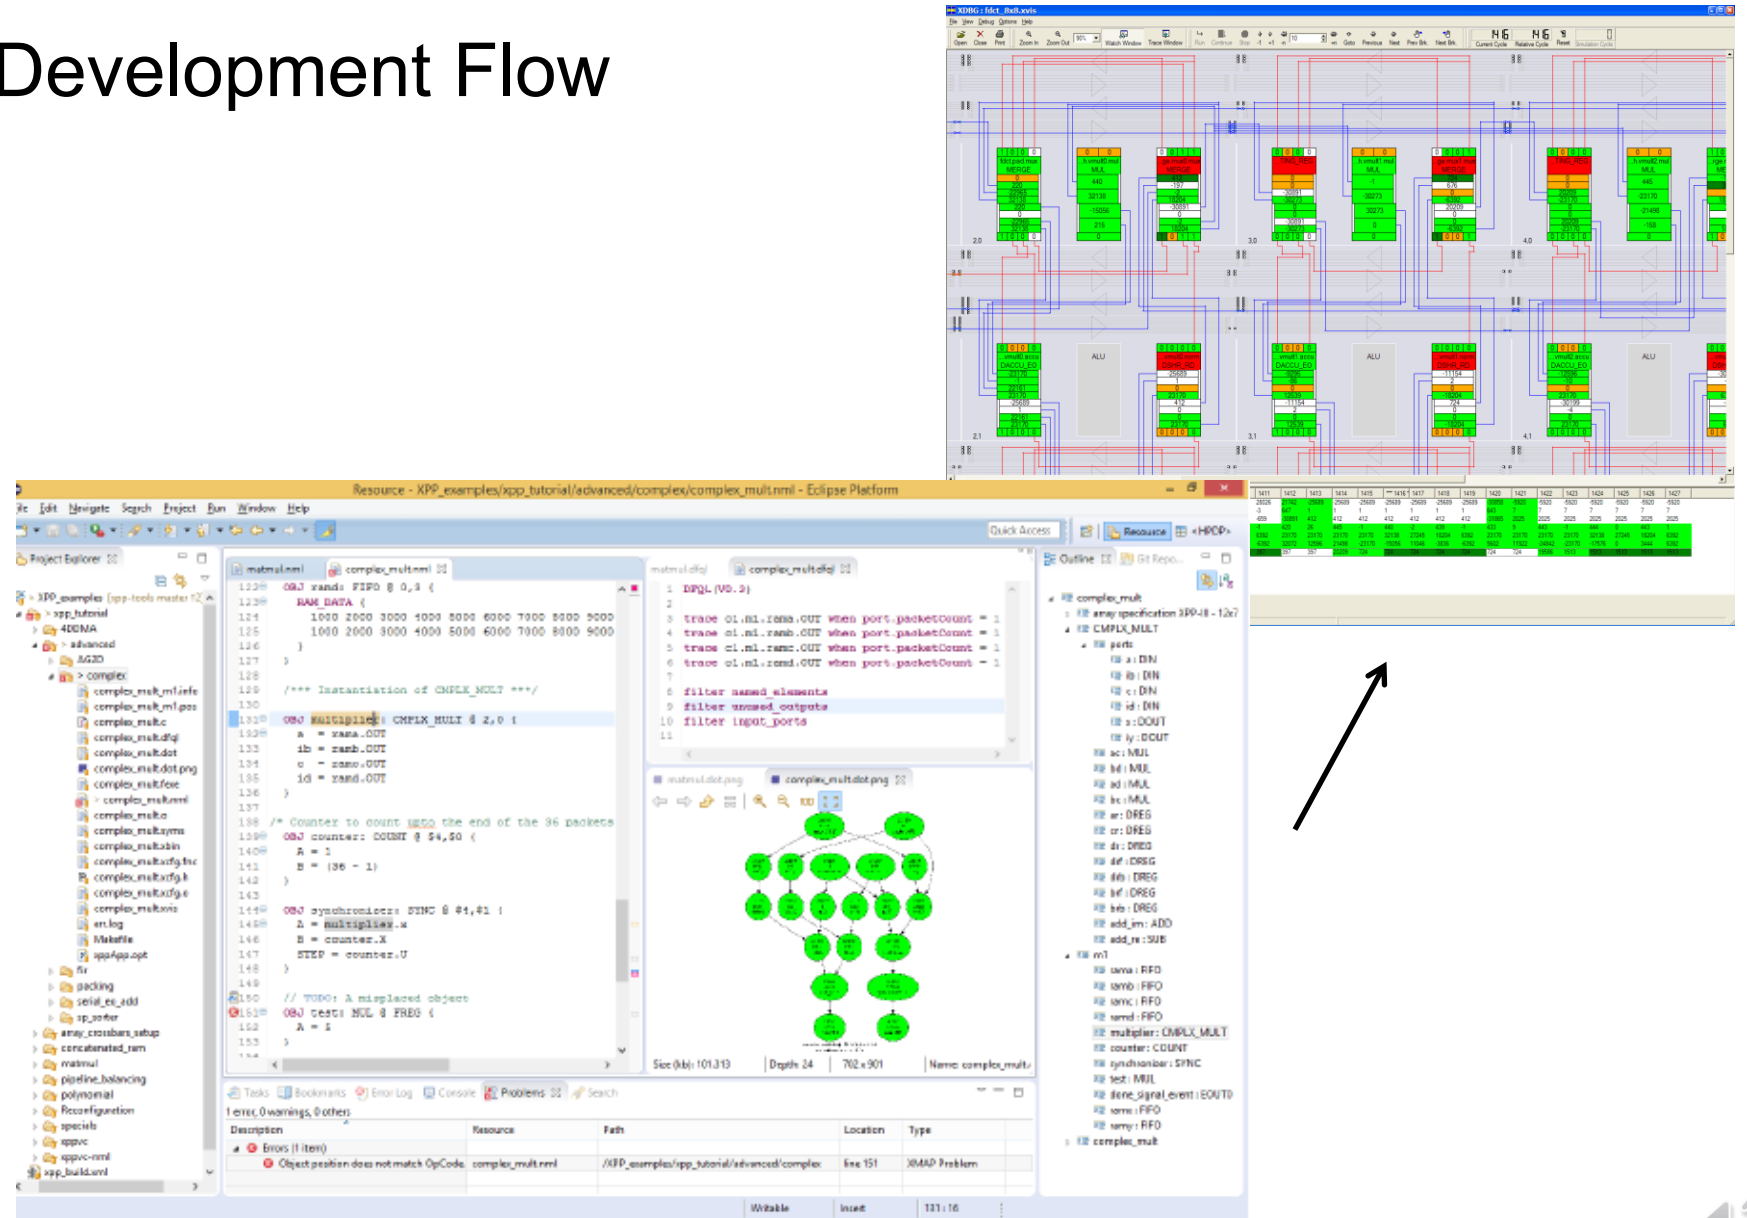

- 40Gops of fixed point arithmetic operations (if required, floating point operations can be emulated)
- 4x 1.1Gbps Streaming Ports
- >4 Mbyte on-chip SRAM
- Memory protections, Watchdog

HPDP40 Device (2/2)

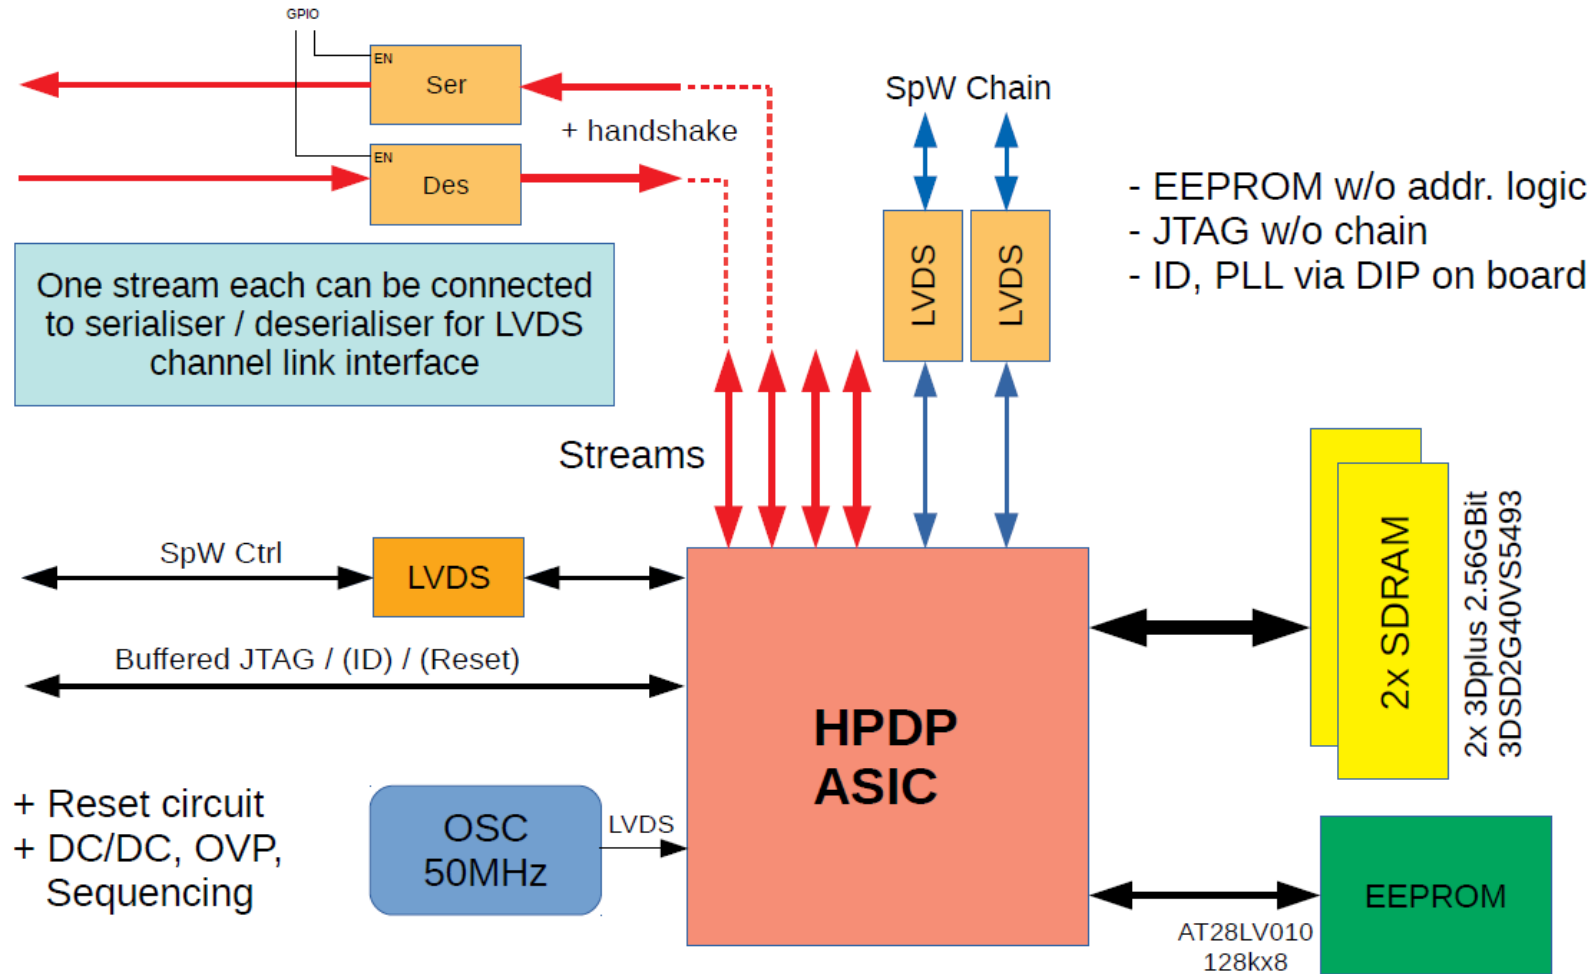


HPDP40 Software Development Flow

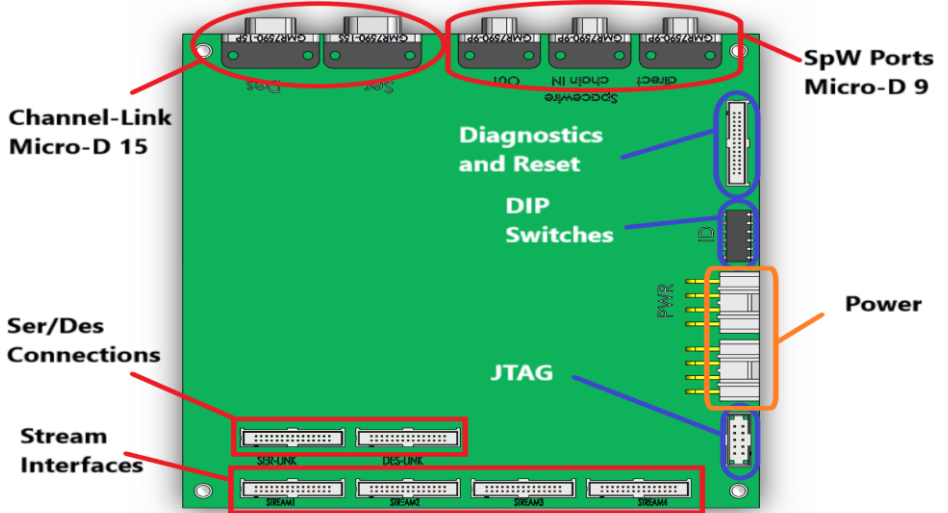

$$y = ax^2 + bx + c.$$




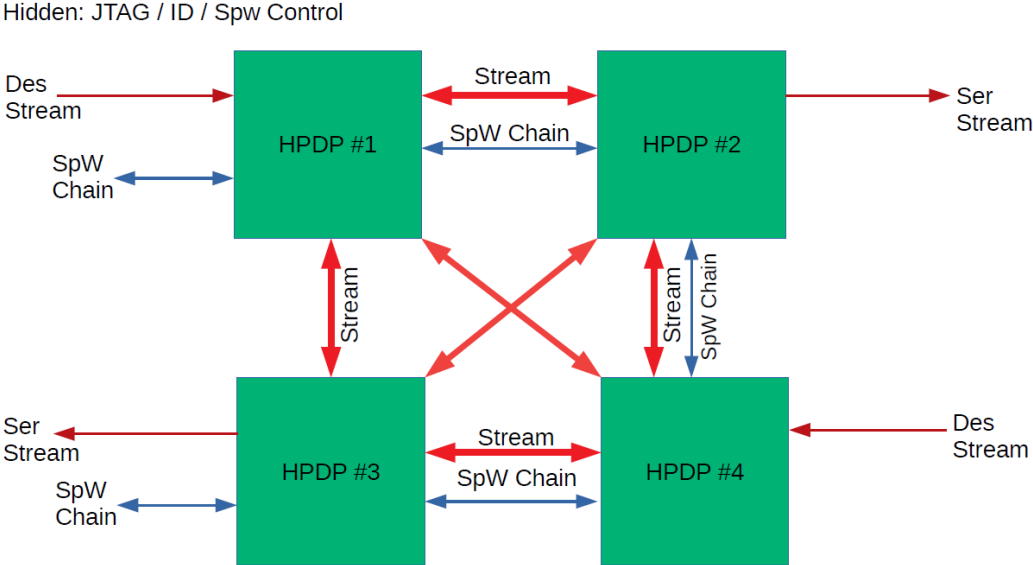
Evaluation board (1/3)



Evaluation board (2/3)



Evaluation board (3/3) - Tiling



Radiation tests

- It has being implemented in the 65nm radiation hardened technology of ST Microelectronics (C65SPACE).
- Total dose
 - Sustaining > 300Krad
- Heavy Ion
 - SEL -> None at LET > 72.2 MeV-cm²/mg @ 90°C junction temperature and maximum supply voltage
 - SEU -> $8.0E-6 < \sigma < 1.4E-5$ for $20.4 < LET < 46.1$ MeV-cm²/mg @ 70°C junction temperature
 - SEFI -> $7E-6 < \sigma < 2.2E-5$ for $20.4 < LET < 46.1$ MeV-cm²/mg @ 70°C junction temperature


Benchmarks overview

- Asteroid feature extraction based on the real time analysis of high resolution, high frame rate images of the moon surface.
 - Count of flashes
 - Intensity and Duration
- Vessel identification for maritime surveillance applications
 - Deck detection and size estimation
 - Coordinates extraction and correlation with AIS data
 - AI expansion of the algorithm
- On-the-fly encryption/decryption
 - AES 256 based implementation
 - Featuring key expansion capability
- Images Compression (CCSDS123)
 - Implementation of prediction and encoder in the array

Moon Asteroid Impacts (1/2)

Objective:

- Find asteroid impacts on moon.
- Measure flash intensity and extent.
- Estimate energy released.



Asteroid impact on true images

Key Features:

- Temporary Noise Reduction implemented by averaging the previous 7 frames.
- Calculation of the Spatial Extension of the flash on each event.
- Extraction of the coordinates of the event.
- Calculation of the Intensity of the event.
- Background static information (i.e. stars, white noise) is removed by the algorithm.
- Identification of High energy particle traces on sensor.

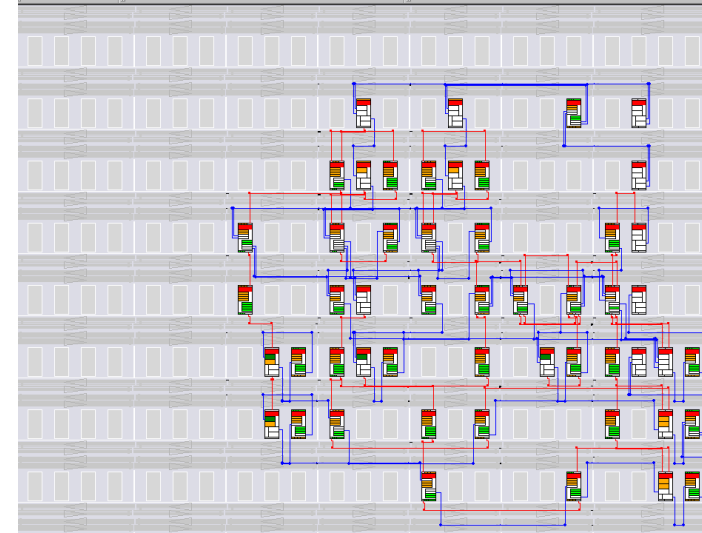
Moon Asteroid Impacts (2/2)

Implementation:

- Mapping of the identification algorithm takes less than 1/2 of the array.
- Current implementation processes 8 pixels per cycle (7 previous frames + current frame).

Performance:

- Images of 1080x1280 (8bit gray scale) at 250MHz require 140Kcycles to be processed.
- Attain performance 116fps.
- Total power consumption of 1.65W.



Mapping of the impact detection and measurement

Improvements under consideration and future applications:

- 16 bit images
- Average on 16 previous images
- Higher resolution images
- Same type of algorithm can be used for high precision star tracker applications.

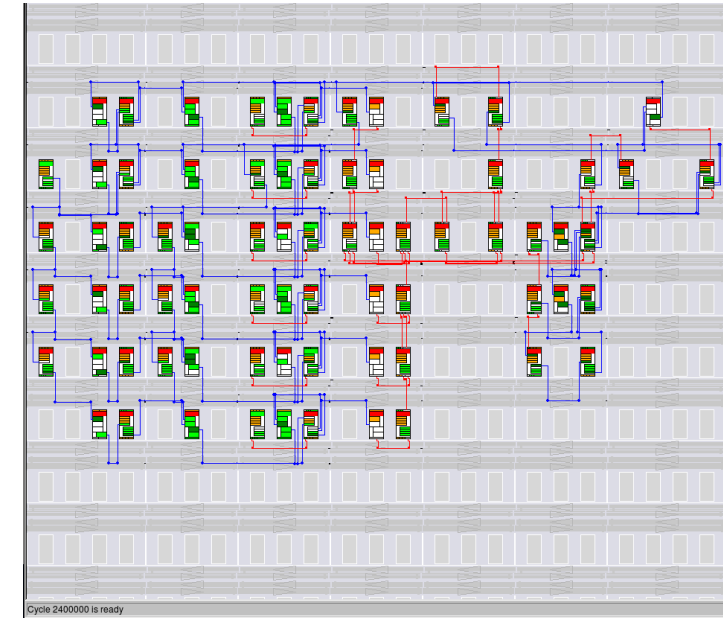
Vessel identification and Positioning – classical approach (1/2)

Objective:

- Identify vessels at sea.
- Cross-check with AIS data
- Flag the suspect vessels

Key Features:

- Optional pre-processing with the Sobel filter for noise reduction.
- Convolution of the image with 6 kernels is performed (Shape detection).
- Eventual multiple detections of the same vessel are eliminated automatically.
- Coordinates and size of each vessel are extracted.
- Array occupancy factor ~50%
- The land mask issue has been partially addressed and the probability to receive false number of vessels is low. Obviously an earth mask is the optimal solution.
- Comparison with AIS data will be done in differed time.




Mapping of the algorithm on the XPP

Vessel identification and Positioning – classical approach (2/2)

Performance:

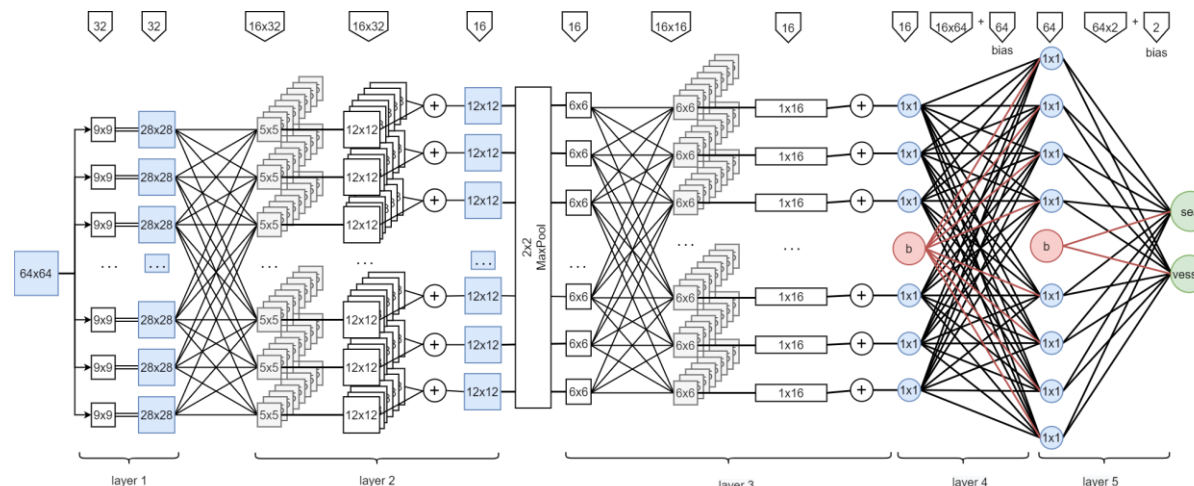
- For the Sobel filter, for images 1024x1024 the attained performance is 94fps
- For the Kernel Convolution filter, for images 1024x1024 the attained performance is 9.6fps
- In all cases the power consumption is 1.65W.
- Overall efficiency ~ 60%



True satellite image and kernels reply

Vessel identification and Positioning – AI approach (1/2)

- Architecture study in Python.
- Use Tensor flow framework to design and train the network.
- Image database consists of 6000 64x64 pre-labeled satellite images.
- Neural network consists of 22.000 parameters and 5 layers.
 - ✓ Each parameter is 32b floating point
 - ✓ 3 convolutional layers (RELU activation) and 2 dense.
- After training this implementation yields 96% correct classifications.
- Porting of the neural network on the HPDP implies:
 - ✓ Scaling of the parameters in order to have 8-bit integer arithmetic
 - ✓ Since the whole network does not fit in the array we have used the unique HPDP feature the on-the-fly reconfiguration. The required delay per reconfiguration is 10Kcycles ~ 40us.




Vessel identification and Positioning – AI approach (2/2)


- Sliding window of 64x64 will be executed on true images.
- Obtained performance of 88% of correct classifications.
- Power consumption of 1.65W.



Original Image



Kernel Convolution - result



NN Result

AES Encryption

Features:

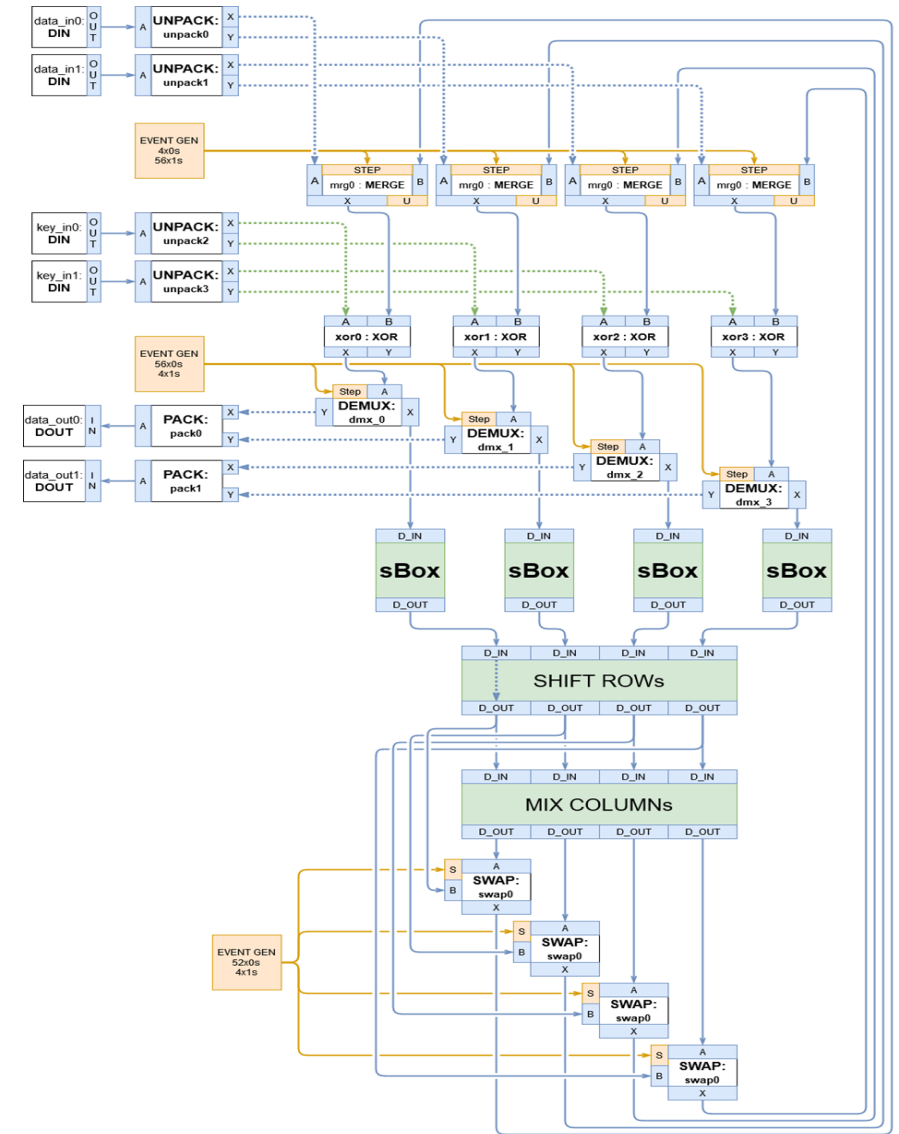
1. Three different version: 128bit, 256 bit and 256 bit with CBC (Cipher Block Chaining)
2. Key expansion done once on the serial CPU.
3. S-Blocks and CBC implemented in the array.
4. Fast Data refill from the internal SRAM.

Performance:

- AES128: 16.2 MB/s @ 250MHz (two instances per HPDP)
- AES256: 11.7 MB/s @ 250MHz (two instances per HPDP)
- AES256cbc: 5 MB/s @ 250MHz (one instance per HPDP)
- Power consumption fixed at 1.65W

Potential improvements:

- Enhance the performance up to a factor 4 by having four (4) HPDP devices interconnected in a tile fashion, all executing in parallel



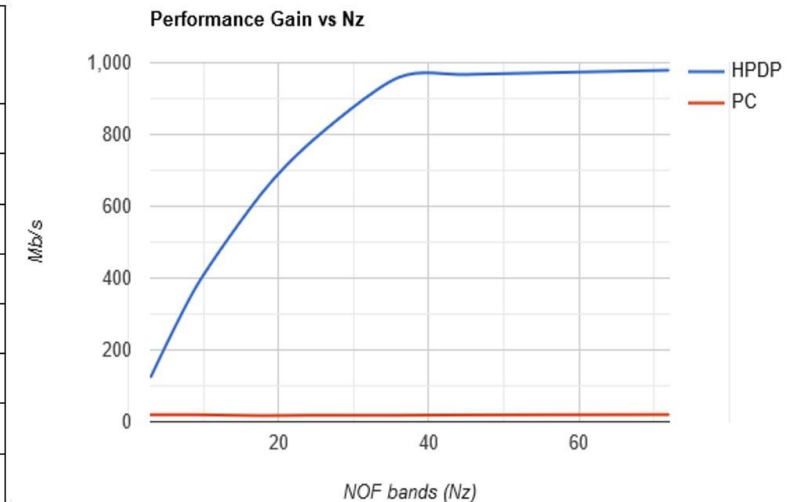
CCSDS123 image compression

Features:

1. Lossless version only
2. Two array configurations are required (prediction, encoding)
3. Reconfiguration on-the-fly takes 40us (10.000 cycles @250MHz)

Performance @ 1.65W power consumption:

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



Conclusion

- The hardened HPDP40 device is available – it is an ITAR free device
- Several applications have been developed during the last 2years
 - the complete flow has been exercised
 - the learning curve has been well understood and dimensioned
 - results have been evaluated on the evaluation board
- The very fast on-the-fly reconfigurability proves to be a unique feature as well as the tilling
- Certainly the HPDP can be used for executing of Neural-Networks
- **Our understanding as of today for on-board processing:**
 - It offers the ultima flexibility (several applications in parallel, several applications on the same dataset)
 - No need for external processor or uC
 - Possibility for running an OS (Linux)
 - Advantages wrt FPGAs
 - **To the best of our knowledge the best compromise reported between mass, power consumption, modularity and computational capability.**