

<u>Zdenek Matej</u>¹, Kenneth Skovhede^{1,2}, Carl Johnsen², Artur Barczyk¹, Clemens Weninger¹, Andrii Salnikov¹ and Brian Vinter²



¹⁾ MAX IV Laboratory, Lund, Sweden
 ²⁾ Niels Bohr Institute, København, Denmark



How one can get to HDF5 data processing with FPGAs?



Zdenek - scientific sw, material science & diffraction methods



MAX IV synchrotron laboratory ... few years ago

(Thought): How one can do it?

32 x 590 = ~18880 Mbps. I.e. 18880 x 33 = ~623040 Mbps = ~623 Gbps

We are writing a grant application for dealing with synchrotron and XFEL data and we will use GPUs.

Hmm, GPUs. Are you using also FPGAs for these applications ?

No, bla, bla, bla ... (many excuses).

• Xilinx Kintex Ultrascale+ xcku5p at 590 MHz at ~3 % utilization, resulting in an inner processing speed of



Carl ... another FPGA genius from NBI

Brian ... professor at NBI

github.com/bh107/SME-Binning

"bincounting"

Artur – physics, CERN, MAX IV networks and infrastructure



MAXIV

languages, "inventor" of SME*

Kenneth – expert in FPGA

Processing HDF5 data with FPGAs

HDF5 Hug 2023

Processing photon science data with FPGAs

HDF5 and heterogenous hw accelerated compute

- frame corrections
 - conversions
 - mask
 - flat field
- simple processing
 - image binning
 - roi-counting
- spot finding
- GPU like applications
 - azimuthal integration
 - phase retrieval
 - full field tomography reconstructions
- (de)compression
- ML inference
 - image classification
- frame & event filtering



Credit: The HDF Group - maintainers of HDF file format

- all x-ray scattering and imaging data acquired by MAX IV DAQ are stored in the HDF5 format
- bslz4 is the default compression for 2D detector image data (frames)

FPGAs pros

- real time
- parallel, high performance and throughput
- integrated networking couple with detectors
- power effective
- common hw with many other applications (like GPUs)



2010–2015: *800 FPGA application reported* by J. Romoth et al. (2017), doi: 10.13140/RG.2.2.16364.56960 + *increasing variety*

cons

- difficult to program
- difficult to install
- not so rich scientific sw ecosystem (?)



Existing FPGA applications

in Photon and Neutron (PaN) sciences

• Maxime Martelli et al. (CNRS, Paris)

3D Tomography Back-Projection Parallelization on Intel FPGAs Using OpenCL

- doi: <u>10.1007/s11265-018-1403-6</u>
- Filip Leonarski et al. (SLS/PSI)
 Fast and accurate data collection for macromolecular crystallography using the JUNGFRAU detector

 doi: 10.1038/s41592-018-0143-7
- César González et al. (Barcelona)
 High Performance Computing PP-Distance Algorithms to Generate X-ray Spectra from 3D Models
 - doi: <u>10.3390/ijms231911408</u>







Processing HDF5 data with FPGAs

Field Programmable Gate Arrays

- arrays of logic gates
 - O(10) million gates today in largest FPGAs
- 'programmed' by creating connections between logic blocks
- arranged in functional blocks
 - DSP, logic, memory, I/O
- O(10) Tflops, comparable to GPUs
- (ordinary) programmed in specialized languages, adapted to electronics development (Verilog, VHDL)
 - intellectual property cores (~sw module) can be purchased
- ideal for bit-level data manipulation
- FPGA boards: integrated network interfaces



Credit: Kovačec, doi: 10.1007/978-3-319-14346-0_40



Processing HDF5 data with FPGAs



Processing HDF5 data with FPGAs

MAXIV

FPGAs as a flexible factory

Credit: Kovačec, doi: 10.1007/978-3-319-14346-0_40

Interconnection I/O Cell

Resources

Block

very parallel hardware ~ 100k of logical units





Access to FPGAs

One size does not fit all

- cost and power effective FPGAs ... for professional applications and hobbyists ... or low power (outer space applications)
 - 100k gates, O(0.5) GB mem, 1 Gbs Ethernet, O(100) MHz, O(2) W
 - O(200) USD (development kit or module/minicomputer)
 - shop: distributors of electronic components
 - popular development kits: PYNQ for Zynq SoC or DE10-Nano Kit
 - sign for academic programmes
- medium size FPGAs
 - 1M gates, O(16) GB mem, 40-160 Gbs, O(300) MHz, O(50) W
 - O(4-6) 1000 USD
- large and extra large FPGAs
 - O(10)M gates, O(32) GB mem, 400 Gbs 1.2 Tbs, O(500) MHz, O(150) W
 - > O(6-14) 1000 USD
- DevCloud: github.com/intel/FPGA-Devcloud \$\approx\$



Data processing with FPGAs

Definitions

Fast

- low latency or high throughput ?
- processing independent data elements: frames, pixels, bytes, bits:
 - 4 Mpix detector, 2 kHz frame rate
 - => 4Mpix * 2kHz = 8 **Gpix/s**
 - (16 bits int) 2 Bytes/pix => 16 GB/s
- fast ... high throughput
 - ... many frames per sec (Hz>>1)
 - ... many pixels/s (Gpix/s>>1)
 - ... many bytes/s (GB>>1)

Real time

- predicable, fixed and low latency
- FPGA clocks O(200) MHz -> 5 **ns**
- FPGA clock can be synchronized with external timing system (accelerator/machine) with sub-pico-second *precision*
- processing a diffraction image may be many clocks ... 32 pix/clock, 4 Mpix image => 4 Mpix / 32pix * 5 ns = 0.625 ms

Curiosity:

- CPU servers ... O(5) Tflops
- GPUs ... O(10) Tflops
- Detectors: O(2) Gpix/s ... factor of >2000 ?



Programming FPGAs for scientific data processing tools

- highl level tools for FPGAs
 - Synchronous Message Exchange (SME): github.com/kenkendk/sme
 - Kenneth Skovhede & Carl Johansen (Niels Bohr Institute)
 - C#/C++/Python -> SMEIL (SME intermediate language) -> VHDL -> bitstream
 - vendor agnostic
 - OpenCL or SYCL/OneAPI



OpenCL

- "C/C++ with pragmas", OpenCL is an older standard used also for GPUs
- FPGAs in DevCloud or in production at MAX IV
- High Level Synthesis (HLS)
 - "C/C++ with pragmas"
- orchestration on host (CPU)
 - <u>Python</u> or C++
- HDF5
 - libhdf5 or <u>h5py</u>



language: OpenCL/C code is hw agnostic alg. core: 36 lines host: PyOpenCL





 Diagnostic utilities API for C++ or Python

- No need of extensive knowledge of FPGA-tools ecosystem for electrical engineers
- No need of "writing Linux/Windows drivers"
- Advantages:
- OpenCL
- OneAPI (C++) + SYCL
- OpenCAPI + HLS

Simplified FPGA (re)programming

ADC waveforms

Frameworks

Hardware and language is not all that matters

- Using a "productivity" framework may accelerate the FPGA project, ... especially if one is starting with programming hardware.
- Popular frameworks:
 - PYNQ: HLS + Python and Jupyter (notebooks)





DAC ADC example **Downloading overlays Downloading Overlays**

1. Instantia	ting an overlay
To instantists on o	overlay a bitstream file name is passed to the Overlay
The laboreran file on the Linux file sy	does not need a full path if it resides in the pyrop pa ptem. Two examples of overlay instantiation are show
e trying base, bi from pying impor al - Overlay("b	lt Lounted in yong package 4 Overlay msc-bit")
In the second case	e, users can use absolute file path to instantiate the o
# Bulley the And from pyres impor of = Sherlay("/	n bititroom, but with full path 4 Overlay Yome/x111nu/pymg/bititroom/Base Alt')

sl.dmsload() ol.kittress.timectamp



Grove ADC





HDF5 data with FPGAs

key components



Credit: Unknown Authors, licneced under <u>CC BY-SA</u>







h5py.h5d. read_direct_chunk

bitshuffle. decompress_lz4

pyopencl. enqueue_copy pyopencl. enqueue_nd_range_kernel

HDF5 direct chunk read

is the "key" ingredient that allows this work.

decompression on CPU is optional ... should be removed in future ... 2nd part of the talk



Azimuthal integration

Data processing example: 2D image -> 1D pattern

- for each image pixel (n,m->i_{pix} ... linearized index):
 - 1. trigonometric calc. -> scattering angle ($2\Theta_{n,m}$)
 - 2. binning $(2\Theta_{n,m}) \rightarrow single index (i_{out-bin})$
 - 3. intensity corrections
 - 4. histogram, bincount or sparse matrix multiplication





 $2\Theta_{max}$ Processing HDF5 data with FPGAs

Basic algorithm on FPGAs

bincount - first implementation

• bincount: simple algorithm similar to computing histogram

- result kept in "local" memory (FPGA Block or Ultra RAM)
- If position_old == position_new:
 - True: accumulate (sum values)
 - False: store old acc, load new acc, accumulate
- 1 pixel per clock (per processing unit)



- by Carl Johansen (NBI) using Synchronous Message Exchange (SME) in 1 day
- initially only for integers
- performance numbers:
 - small FPGA at 100 MHz: 1 Gpix/s (10% util. per processing unit)
 - large FPGA at 590 MHz: 20 Gpix/s (3% util. per unit)
- ref: github.com/bh107/SME-Binning
- OpenCL/OneAPI/HLS:
 - one of many FPGA programming "tricks" often called "Shannonization" (Claude Shannon)



Basic algorithm on FPGAs

bincount + bisort & resort for floating point data

- initially only for integers -> real-numbers corrections -> floats(32)
- need for pixel "reordering" (bisort + resort)
- again 1 pix/clock/pipeline





Bitonic sorter

Credit: <u>by Octotron - own work,</u> <u>CC BY-SA 3.0</u>

- Performance prediction:
 - 1 pipeline ... 1.25% of FPGA resources
 - 1 decompression pipeline ... 1%
 - 20% for the framework + 32 pipelines * (1.25+1)% = 92% FPGA utilization
 - 1 pix/s * 32 pipelines * 360 MHz = 11.5 Gpix/s

can be expressed as systolic array -> very suitable for FPGA



Performance figures

AZINT bincount – FPGAs with OpenCL, no decompression



	Medium	Large	comment
size	medium	large	
FPGA	Aria 10 GX	Stratix 10 MX	
process	20 nm	14 nm	User configurable
memory	2 x DDR3	2 x HBM2	QSFP+ Cage 4x A10 GX115 (NF40) A10 GX115 (NF40) A10 GX115 (A10 GX105 (NF40) A10 GX105 (NF40) (NF40
QSPF	2x10/40 Gbs	4 x 100 Gbs	QSFP+ Cage
framework	OpenCL	OpenCL	O Pole 63/28 Credit: Bitt
processing pipelines	32	32	pix/s
ALUTs utilization	45%	40%	
RAMs utilization	60%	25%	fp32, 8k bins
frequency / ideal (MHz)	205 / 240	360 / 480	
host-to-device bandwidth	4.7 GB/s	5.6 GB/s	x8 PCIe Gen3, can handle 4.5M x 500 Hz 🔽
processing (virtual) pixel rate	5.7 Gpix/s [*]	8.9 Gpix/s	allows pixel-splitting = 3 🔽
*comparable to NVIDIA V100 (a)	Gniv/s 12 nm process		

HDF5 Hug 2023

Processing HDF5 data with FPGAs





[azint-cl] OpenCL will use AOCX binary image: ./build/pac_a10/bincount2_intel.aocx fpga integration (n=734), exec. time: 7.605 sec fpga bincout-intel (16.777 Mpix): 96.5 fps, 1619.2 Mpix/s [168. 178. 189. 192. 200.]

FPGA result in 7.6 sec

Conclusions

In this notebook it was demonstrated that

• azimuthal integration for powder diffraction data can be done with bincount

- bincount procedure can be evaluated effectively at FPGA
 - mainly metadata describing the geometry of the experiment are needed
 the procedure can effectively orchestrated from Python
- FPGA based data crunching can be done with Jupyter notebooks in Cloud-like environement
- notebooks can be executed in particular on DeviceCloud with help of papermill

papermill examples/DyCo2.ipynb DyCo2-result.ipynb -k bincount-jupdemo -p PYOPENCL_CTX 0 -p board_name pac_a10



Decompression on FPGAs with OpenCL

bitshuffle & Iz4

oneAPI-samples / DirectProgramming / C++SYCL_FPGA / ReferenceDesigns / decompress /



- key component: LZ77
 decoder
 - history buffer
 - multi-byte in / out
 - oneAPI -> OpenCL retrofitting done in few hours :-)
 - Snappy reader being replaced by Iz4 reader
 - bitshuffle stage needs to be added



رب



Iz4 sequence

each byte adds max 255 to literals length ↓

Token: literals length (4 high bits), match length (4 low bits)



- multiple literals
- single match-copy operation



Iz4 – statistics

bitshuffle & Iz4 in pure Python and numpy

©Dectris Eiger, 32 bit, 16 Mpix, MX experiment



Note: each byte adds only max 255 to length



Iz4 – statistics

bitshuffle & Iz4 in pure Python and numpy



85.0 percentile is 5.0
85.4% entries have max len 5 bytes
82.2% entries have max len 4 bytes
33.8% entries have max len 3 bytes

- e.g. 3 bytes are pure match-copy ops
- optimizations
 - speculate 3-5 bytes sequences



Summary

- we can do data processing like AZINT on FPGAs with the standard scientific precision
- we use direct chunk read to get uncompressed data
- having decompression on FPGAs is close
- AZINT & bincount & bslz4: gitlab.com/MAXIV-SCISW/compute-fpgas





This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101004728



MAX IV Laboratory, Lund Artur Barczyk Andrii Salnikov Clemens Weninger Niels Bohr Institute, Copenhagen Kenneth Skovhede Carl Johnsen Brian Vinter Mads Jörgensen

Thank you for your attention

