Filip Leonarski :: MX Data :: Swiss Light Source

Jungfraujoch: kilohertz data acquisition and on-the-fly image analysis for MX beamlines

LEAPS Innov WP7 Seminar
13 April 2022

MX development is driven by X-ray detectors

Detector developments:

- Lower noise
- Wider X-ray energy range
- Higher frame rate (throughput/time resolution)

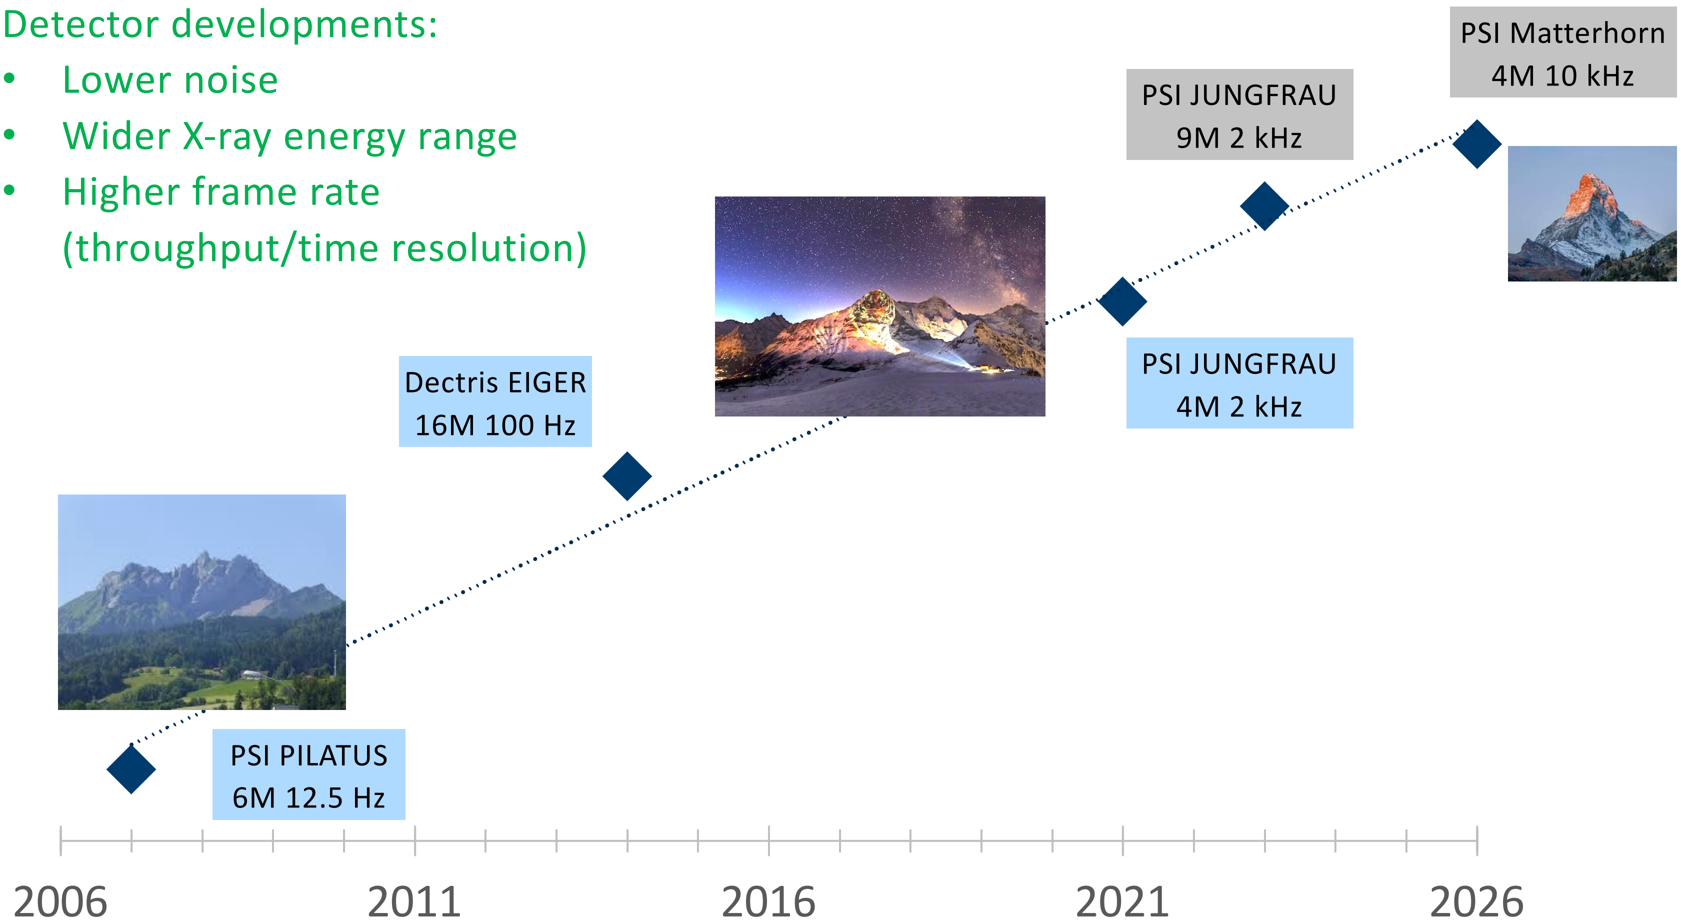

Detector development is challenging for IT infrastructure

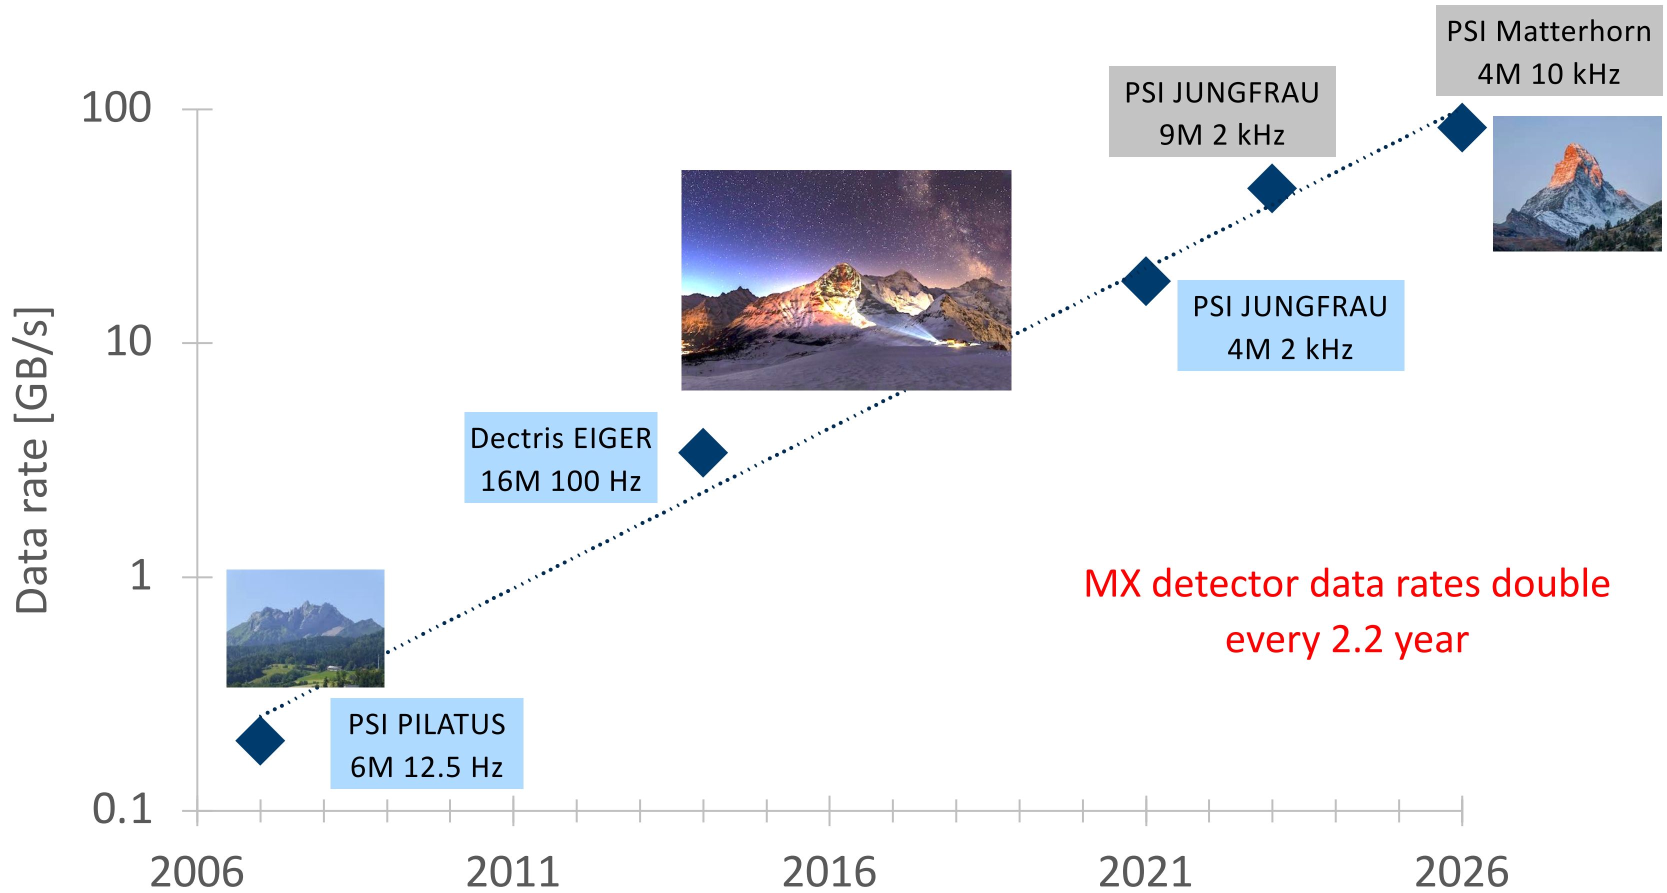

"JUNGFRAU detector for brighter x-ray sources: Solutions for IT and data science challenges in macromolecular crystallography"
Leonarski F., *et al, Structural Dynamics* (2020)

Solution to data challenge

- <u>Computing challenge</u>
  - Mainstream computing is not growing exponentially anymore
  - Massive parallelization (multiple servers) of software-based solution possible, but not sustainable in long run
  - Data movement (memory/network) is the main issue
- <u>Development needed</u>
  - **Reduce data close to source**
  - Combine acquisition and on-the-fly processing
  - Use task-specific computing architectures

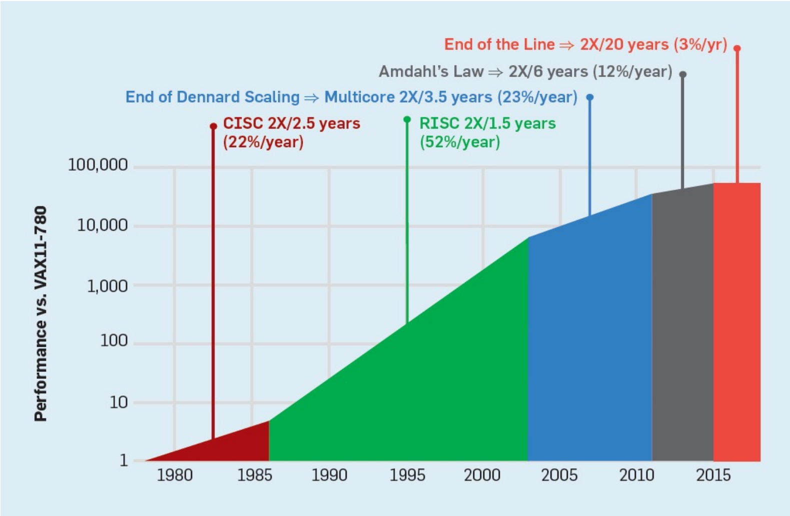

Hennessy & Patterson
doi:10.1145/3282307

Science application: JUNGFRAU detector for synchrotron MX

- JUNGFRAU applications
  - Native-SAD (3-6 keV)
  - High-flux (time resolved/high throughput)
  - Room temperature

- JUNGFRAU operational requirements
  - To reach >95% duty cycle at a synchrotron source, one needs to operate JUNGFRAU at kHz frame rates (at least internally)
  - Raw ADUs need to be converted to photons for analysis and compression

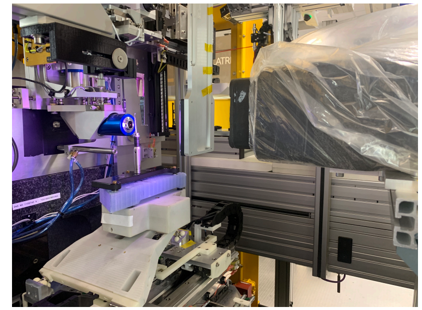

Test at **VMXi Diamond Light Source** (UK)

nature methods | ARTICLES
https://doi.org/10.1038/s41592-018-0143-7

Fast and accurate data collection for macromolecular crystallography using the JUNGFRAU detector

Filip Leonarski[1], Sophie Redford[1], Aldo Mozzanica[1], Carlos Lopez-Cuenca[1], Ezequiel Panepucci[1], Karol Nass[1], Dmitry Ozerov[1], Laura Vera[1], Vincent Olieric[1], Dominik Buntschu[1], Roman Schneider[1], Gemma Tinti[1], Erik Froejdh[1], Kay Diederichs[2], Oliver Bunk[1], Bernd Schmitt[1*] and Meitian Wang[1*]

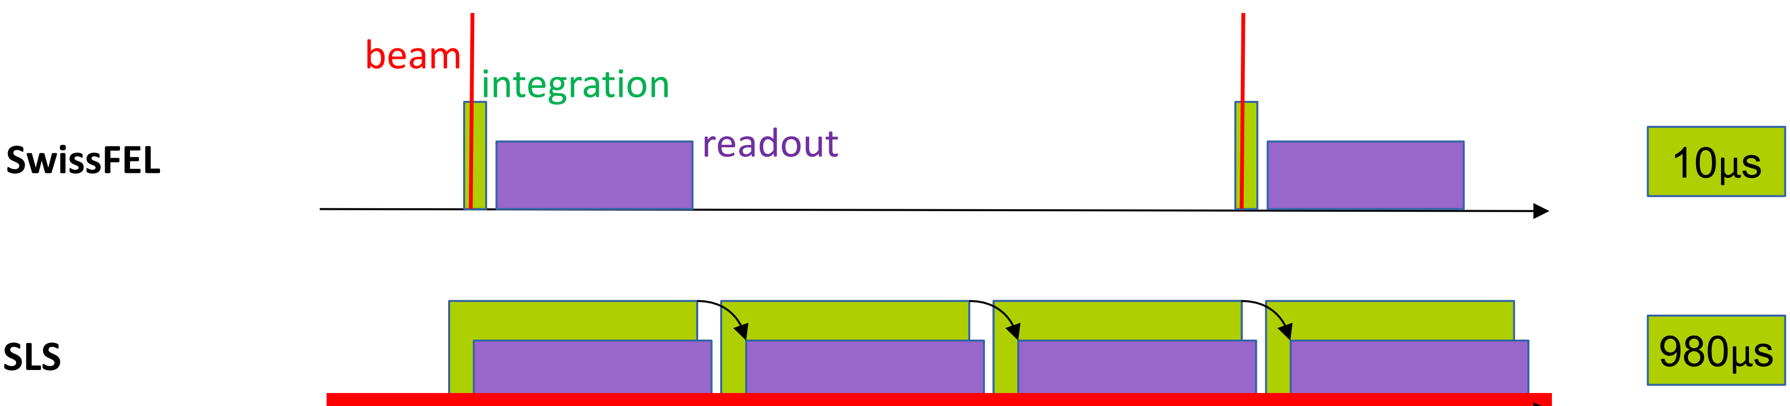

A. Mozzanica

Jungfraujoch development story

- 2017: JUNGFRAU detector demonstration for synchrotron MX
- 2018: Jungfraujoch backend development started
  (technology exploration)
- 2019: Collaboration with IBM started
- 2020: First operation after 6 months of development at X06SA
- 2020: Tested JF 4 Mpixel 1 kHz at BL-1A, Photon Factory KEK
  (native-SAD in helium at 4 keV; rotation crystallography)
- 2021: Tested JF 4 Mpixel 2 kHz at BioMAX, MAX IV
  (time resolved serial crystallography)

Test at **BL-1A Photon Factory KEK** (JP)

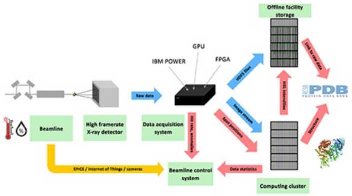

JUNGFRAU detector for brighter x-ray sources: Solutions for IT and data science challenges in macromolecular crystallography

Filip Leonarski, Aldo Mozzanica, Martin Brückner, Carlos Lopez-Cuenca, Sophie Redford, Leonardo Sala, Andrej Babic, Heinrich Billich, Oliver Bunk, Bernd Schmitt and *Meitian Wang*

Structural Dynamics **7**, 014305 (2020); https://doi.org/10.1063/1.5143480

Jungfraujoch general architecture

Eiger Mönch Jungfrau

Detector

Smart network card (FPGA)

CPU memory ring buffer

Online analysis (CPU/GPU)

Work in progress

Web console

(Lossless) Compression

HDF5 Writer

Preview on user console

1 card = 12 GB/s
(~2.5 Mpixel JF)
4 cards possible
(up to 48 GB/s)

Jungfraujoch server (IBM IC922/POWER9)

Jungfraujoch FPGA smart network card

FPGA board with OpenCAPI interface

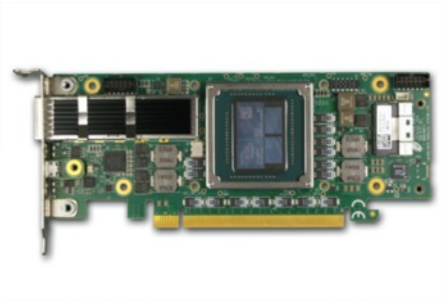

- Data acquisition
- Initial data analysis
- Pre-compression

(2.5 Mpixel/board for JF)

Functionality described in C++ and transcompiled with high-level synthesis to hardware description language

Jungfraujoch FPGA smart network card

- Network
  - Ethernet/IP/UDP stack
  - ARP and ICMP for integration
  - Can be reused for different UDP data source
- Pixel conversion for JUNGFRAU
  - 460 GB/s high-bandwidth memory is crucial (memory BW problem on CPU)
  - Pedestal tracking possible
- Pre-compression/compression
  - Bitshuffle filter trivial on FPGA
  - Any transformations need to be done before (multipixel/spot finding)
  - Compression must happen on CPU (at the moment)

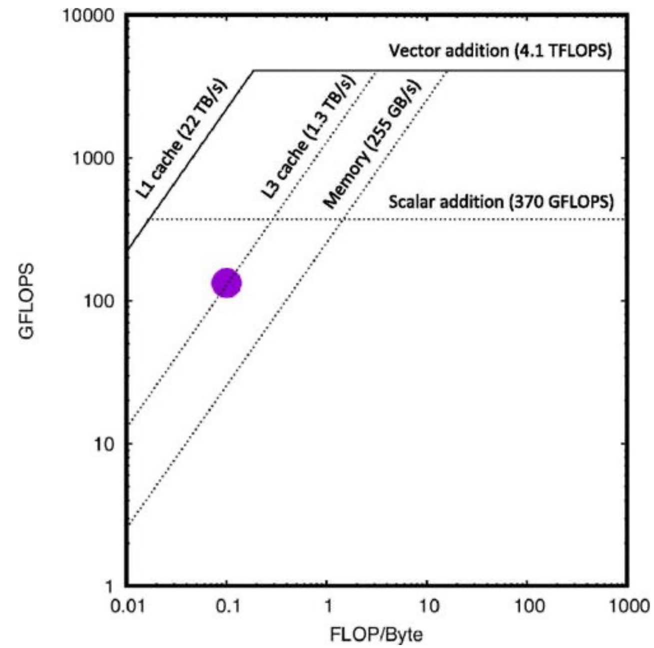

CPU performance (roofline plot)
See: "JUNGFRAU detector for brighter x-ray sources: Solutions for IT and data science challenges in macromolecular crystallography" Leonarski et al. Structural Dynamics (2019) https://doi.org/10.1063/1.5143480

OpenCAPI interface

- Coherent interconnect => FPGA has access to virtual memory, no kernel level development necessary
- Cheaper development / simpler testing / better safety as compared to PCIe
- Requires IBM POWER hardware (or Intel CXL soon)
- Experimental version of Jungfraujoch works with PCIe

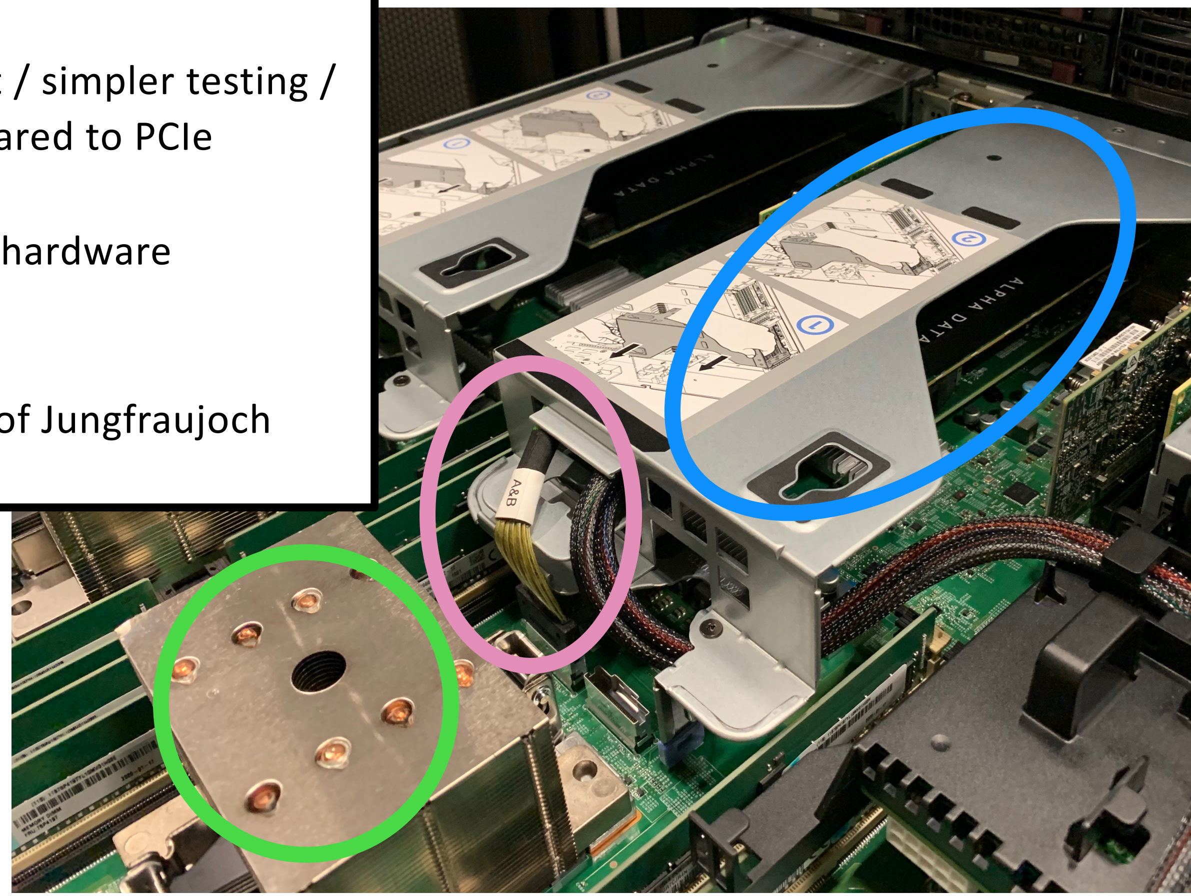

POWER9 CPU **OpenCAPI cable** **FPGA board**

Jungfraujoch FPGA, CPU and GPU

- Host memory writing
  - Work request/work completion mechanism
  - Unit of transfer 1 frame/module: 500k pixel = 1 MiB
  - For OpenCAPI: standard user-space mmap buffer, no special HW allocation, just 128b alignment
  - For PCIe: kernel-space physically continuous buffer

IBM POWER9

- Spot finding
  - COLSPOT algorithm from XDS
  - CPU, GPU and FPGA implementations
  - FPGA is possible for real-time full frame rate performance, but constrained and inflexible
  - GPU is optimal for 100-200 Hz (online feedback)
  - CPU for even lower frame rate

Nvidia Tesla T4

High-level synthesis => continuous integration

- For HLS code written in C++ initial tests can be done on CPU
  - GCC + few extra headers
  - 20 minutes to get wide coverage
  - Possible FPGA HLS code + receiver CPU code in single test

- Hardware simulation for deadlocks and (rare) HLS compiler issues
  - Cadence Xcelium
  - IBM provides OpenCAPI simulation environment => host code is the same as for real hardware test
  - 2-3 hours for simple data acquisition tests (3 images)

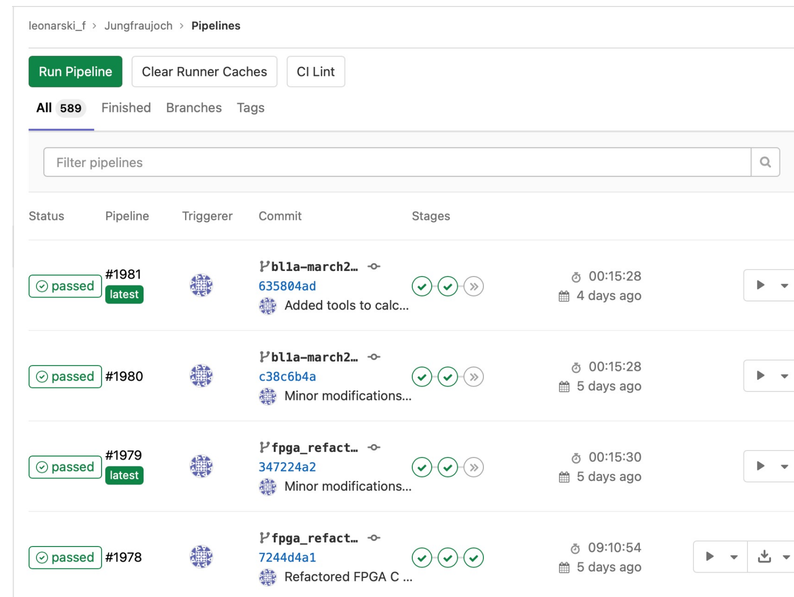

Performance milestone – JF4M 2 kHz @ BioMAX

- Collaboration between MAX IV and PSI (MX + Detector + Bio)
- JF 4Mpixel from PSI + MAX IV computing infrastructure with Jungfraujoch
- Continuous measurements up to 2 kHz (17 GB/s) frame rate on protein targets
  - Consecutive 8 minute runs of 1 million frames at 2 kHz (not tried longer)
- 1 ms time resolution for biological KR2 sample with laser trigger
- Online conversion and compression reduced initial 17 GB/s to below 4 GB/s (GPFS limit)

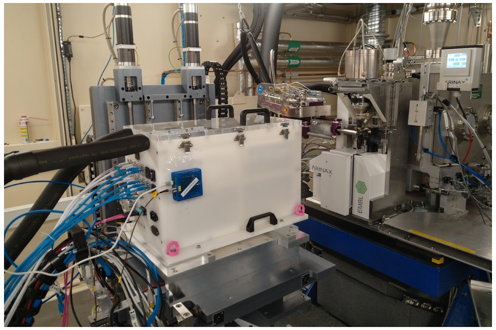

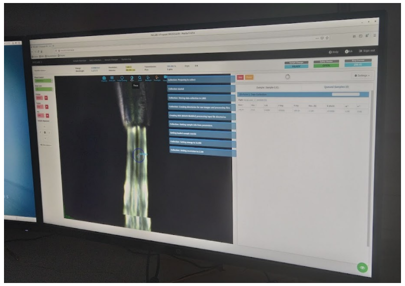

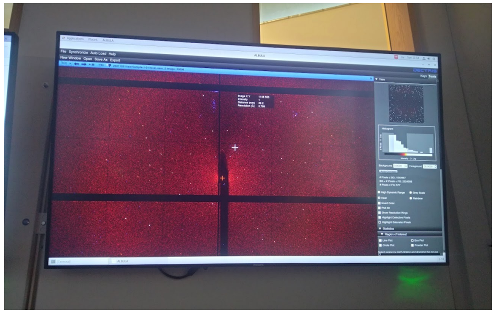

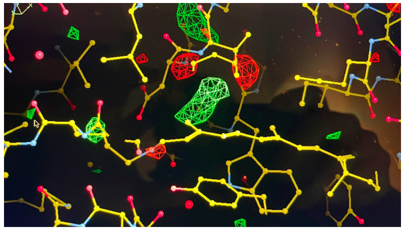

Conclusions

- Task-specific architectures (like FPGA or GPU) are important tools to handle high data rates in coming years
- Design with FPGA is not difficult with proper tools (coherent interconnects & high-level synthesis)

Acknowledgements

MX Group (PSI)
- Vincent Olieric
- Florian Dworkowski
- Takashi Tomizaki
- Chia-Ying Huang
- Sylvain Engilberg
- Sylvain Aumonier
- Greta Assmann
- Justyna Wojdyła
- Meitian Wang

Detector Group (PSI)
- Aldo Mozzanica
- Martin Brückner
- Carlos Lopez-Cuenca
- Viktoria Hinger
- Bernd Schmitt

Science IT (PSI)
- Leonardo Sala
- Hans-Christian Stadler
- Marie Yao Xingxing
- Piero Gasparotto
- Markus Janousch

SLS management (PSI)
- Oliver Bunk

Laboratory of Biomolecular Research (PSI)
- Quentin Bertrand
- Michal Kępa
- Tobias Weinert
- Jörg Standfuss

Photon Factory, KEK
- Naohiro Matsugaki
- Yusuke Yamada
- Masahide Hikita

MAX IV
- Jie Nan
- Zdeněk Matěj
- Thomas Ursby

Uni Konstanz
- Kay Diederichs

IBM Systems (France)
- Alexandre Castellane
- Bruno Mesnet