



Lima1 latest developments the kilo-Hertz regime

Laurent Claustre, Samuel Debionne, Alejandro Homs Puron
Sebastien Petitedemange, Emmanuel Papillon
@ BCU / SG / ISDD

Introduction:

- Challenging detectors

Latest features:

- Sideband data & saving
- Core-level improvements
- CPU & NUMA affinity control

Conclusions:

- Lessons learned

Detector capabilities

- 4x 10 Gbit Ethernet links
- Up to 4 GByte/s of decompressed data
- 4 MB in 8-bit \Rightarrow 1 kHz

DAQ context

- Compressed data is received from DCU ZMQ stream
- Pixel depth depends on detector speed, 8-, 16- or 32-bit

Overhead in LImA plugin

- Data is decompressed
- Expanded to 16/32-bit
- Recompressed for saving



PSI/SLS DETECTORS: CHALLENGES

Critical issues

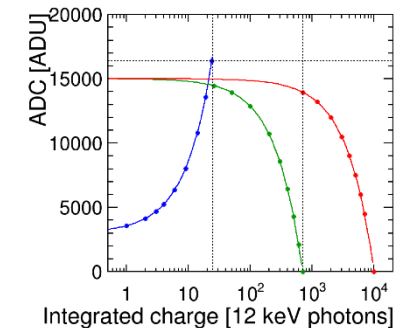

- UDP data streams \Rightarrow packet lost if receiver is not ready
- Linux is not a real-time OS

PSI/Eiger

- One 500k module \Rightarrow 2 GByte/s
- Data transfer speed 8 kHz @ 4-bit (sensor speed: 22 kHz)
- 4-to-8-bit expansion \Rightarrow 4 GByte/s

PSI/Jungfrau


- Adaptive 3-gain per pixel per frame
- CPU-intensive correction
- 32-bit linear scale \Rightarrow 4 GByte/s @ 1 kHz



PSI/Jungfrau-1M

Beamline

1 kHz @ 16-bit
data transfer
Max: 2 kHz



4x 10 Gbit



Data Center

1 kHz @ 32-bit
4 GB/s effective rate
Max: 8 GB/s

4-ch CWDM
MUX/DEMUX



25 Gbit
(~10 Gbit)



Design considerations



- Per-frame binary data
- Travels with image through the processing chain
- Avoid global containers ⇒ better processing scalability



Dectris/Eiger2

- **Dynamic pixel depth** ⇒ avoid data expansion
- Camera plugin inserts blob as sideband data
- Used by reconstruction task ⇒ reduced thread contention
- **Reused for saving compression**

LImA PERFORMANCE: SIDEBAND DATA



Improved frame accumulation

- Added MEAN mode, calculating the average images
- Accumulation is now parallelised for high-throughput plugins like Iris & SIsDetector

More efficient saving subsystem


- Use sideband saving data ⇒ avoid access to global containers
- Improves scalability at high frame rates

Improved processing buffer management

- Allocate auxiliary buffers in advance

CPU & NUMA affinity control

- Assign CPU core/threads to each active task in the system
- Optimize memory access time in NUMA environments



Results

- Jungfrau-1M:
 - Gain & pedestal correction @ 1 kHz
- Eiger-500k:
 - 4-to-8 bit expansion @ 8 kHz
- 4 GByte/s internal data rate
- ~1 GByte/s corrected & compressed output

LESSONS LEARNED

- Computers die, detectors survive!
- Versatile + High Performance control systems
- Limits in HW resources
 - Backend must be over-dimensioned!
- Optimise the code:
 - For scalability
 - Pre-allocate buffers
 - CPU & NUMA affinity control can help
- Additional challenges
 - Integration into IT infrastructure [GPFS]
 - (Continuous) OS upgrade

Thank you very much!



Any question?