Evaluation of the **CBM FLES input interface** at 2016 CERN/SPS beam test

Dirk Hutter
hutter@compeng.uni-frankfurt.de

Prof. Dr. Volker Lindenstruth

FIAS Frankfurt Institute for Advanced Studies

Goethe-Universität Frankfurt am Main,
Germany

http://compeng.uni-frankfurt.de

DPG-Frühjahrstagung

Münster, 31. März 2017

The CBM Experiment at FAIR

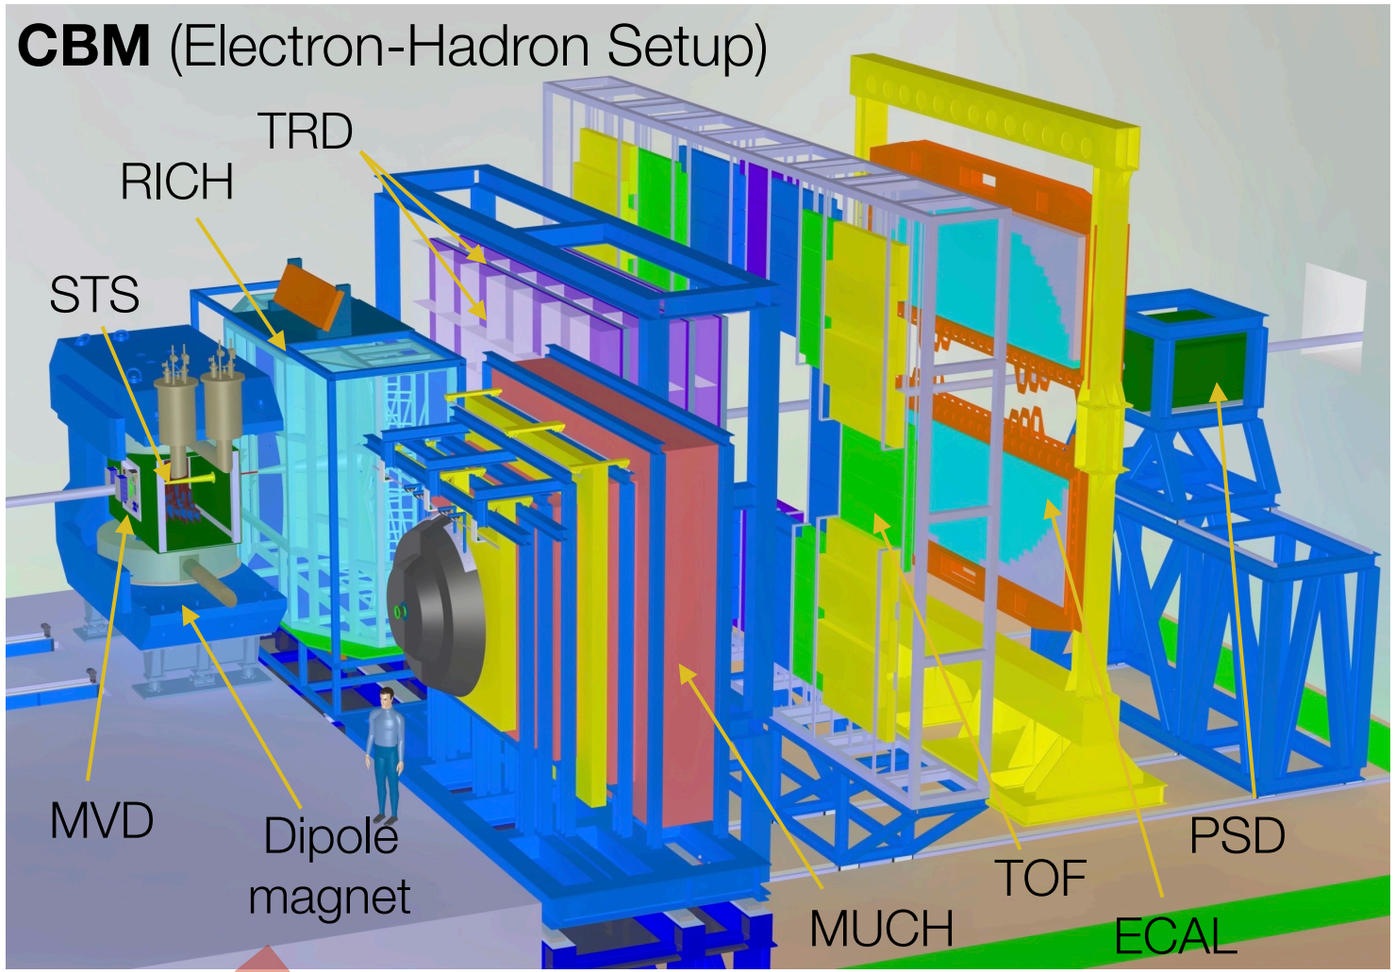

- Fixed target heavy ion experiment at FAIR
- Physics goal: exploration of the QCD phase diagram
- Extreme reaction rates of 10 MHz and track densities of 1000 tracks in aperture
- Conventional trigger architecture not feasible

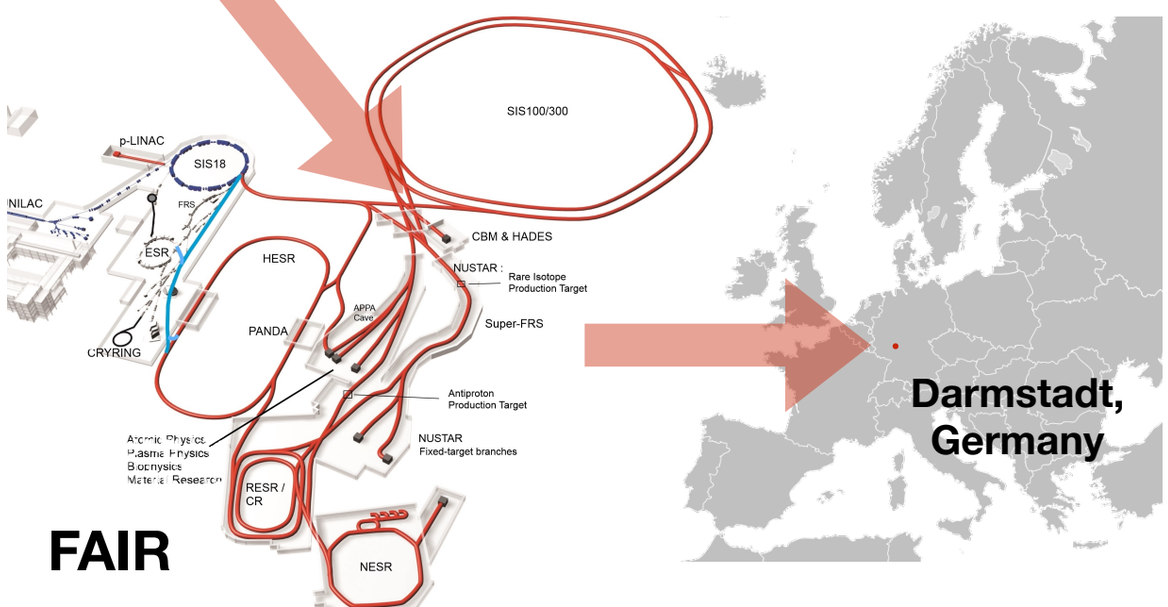

➡ Self-triggering free-streaming readout electronics

➡ Event selection exclusively done in an HPC cluster

CBM Readout Structure

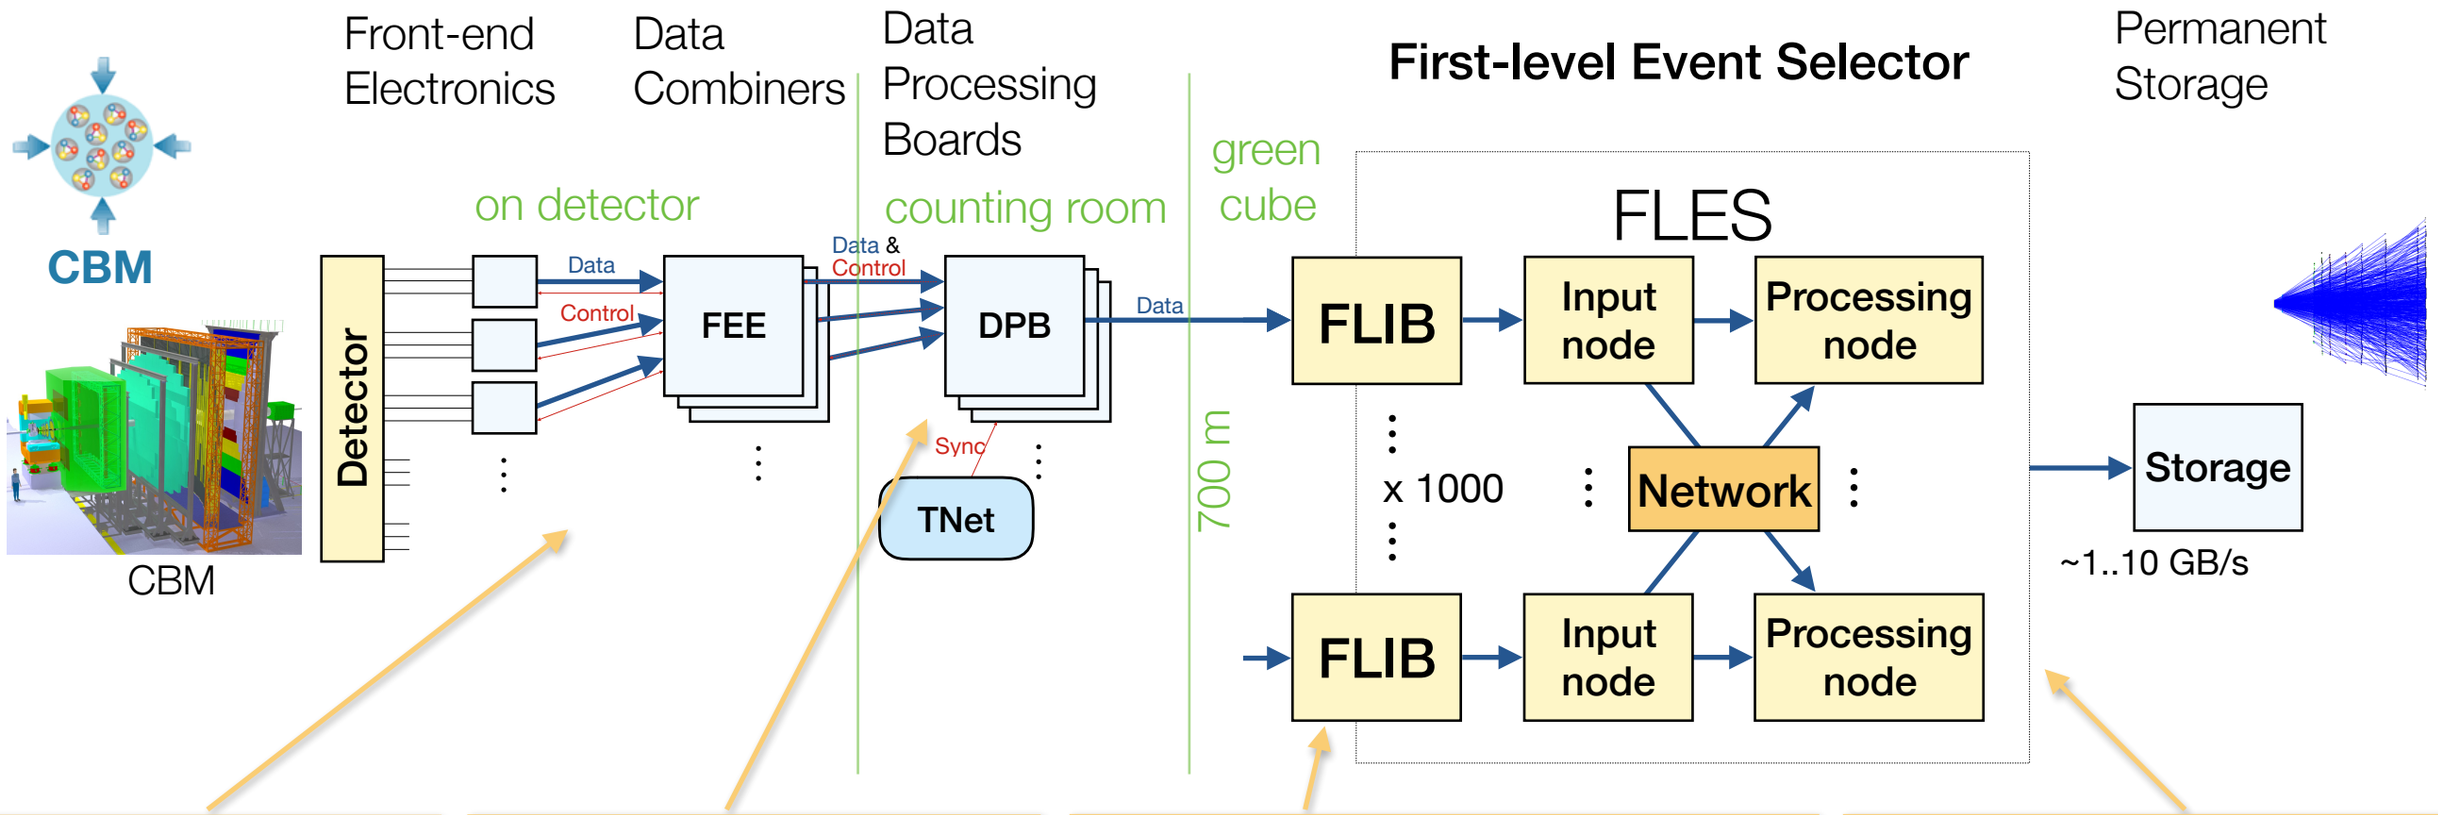

Detector Front-ends

- Autonomous hit detection and **zero-suppression**
- Associate **time stamp** with each hit, aggregate data

Data Processing Board (DPB)

- Local data **preprocessing**: Feature extraction, time sort messages
- Package subsystem date into global container format: **microslices**

FLES Input Interface

- > **1 TByte/s** total input data rate
- FPGA-based PCIe board
- High throughput **DMA engine**
- Long distance link
- Preprocessing and **time indexing** for timeslice building

FLES Cluster

- ~ **60.000 cores**
- RDMA-enabled network
- **Timeslice building**: Deliver global timeslices to reconstruction code
- Full reconstruction, associate hits with events
- Event selection

FLES Interface Board (FLIB)

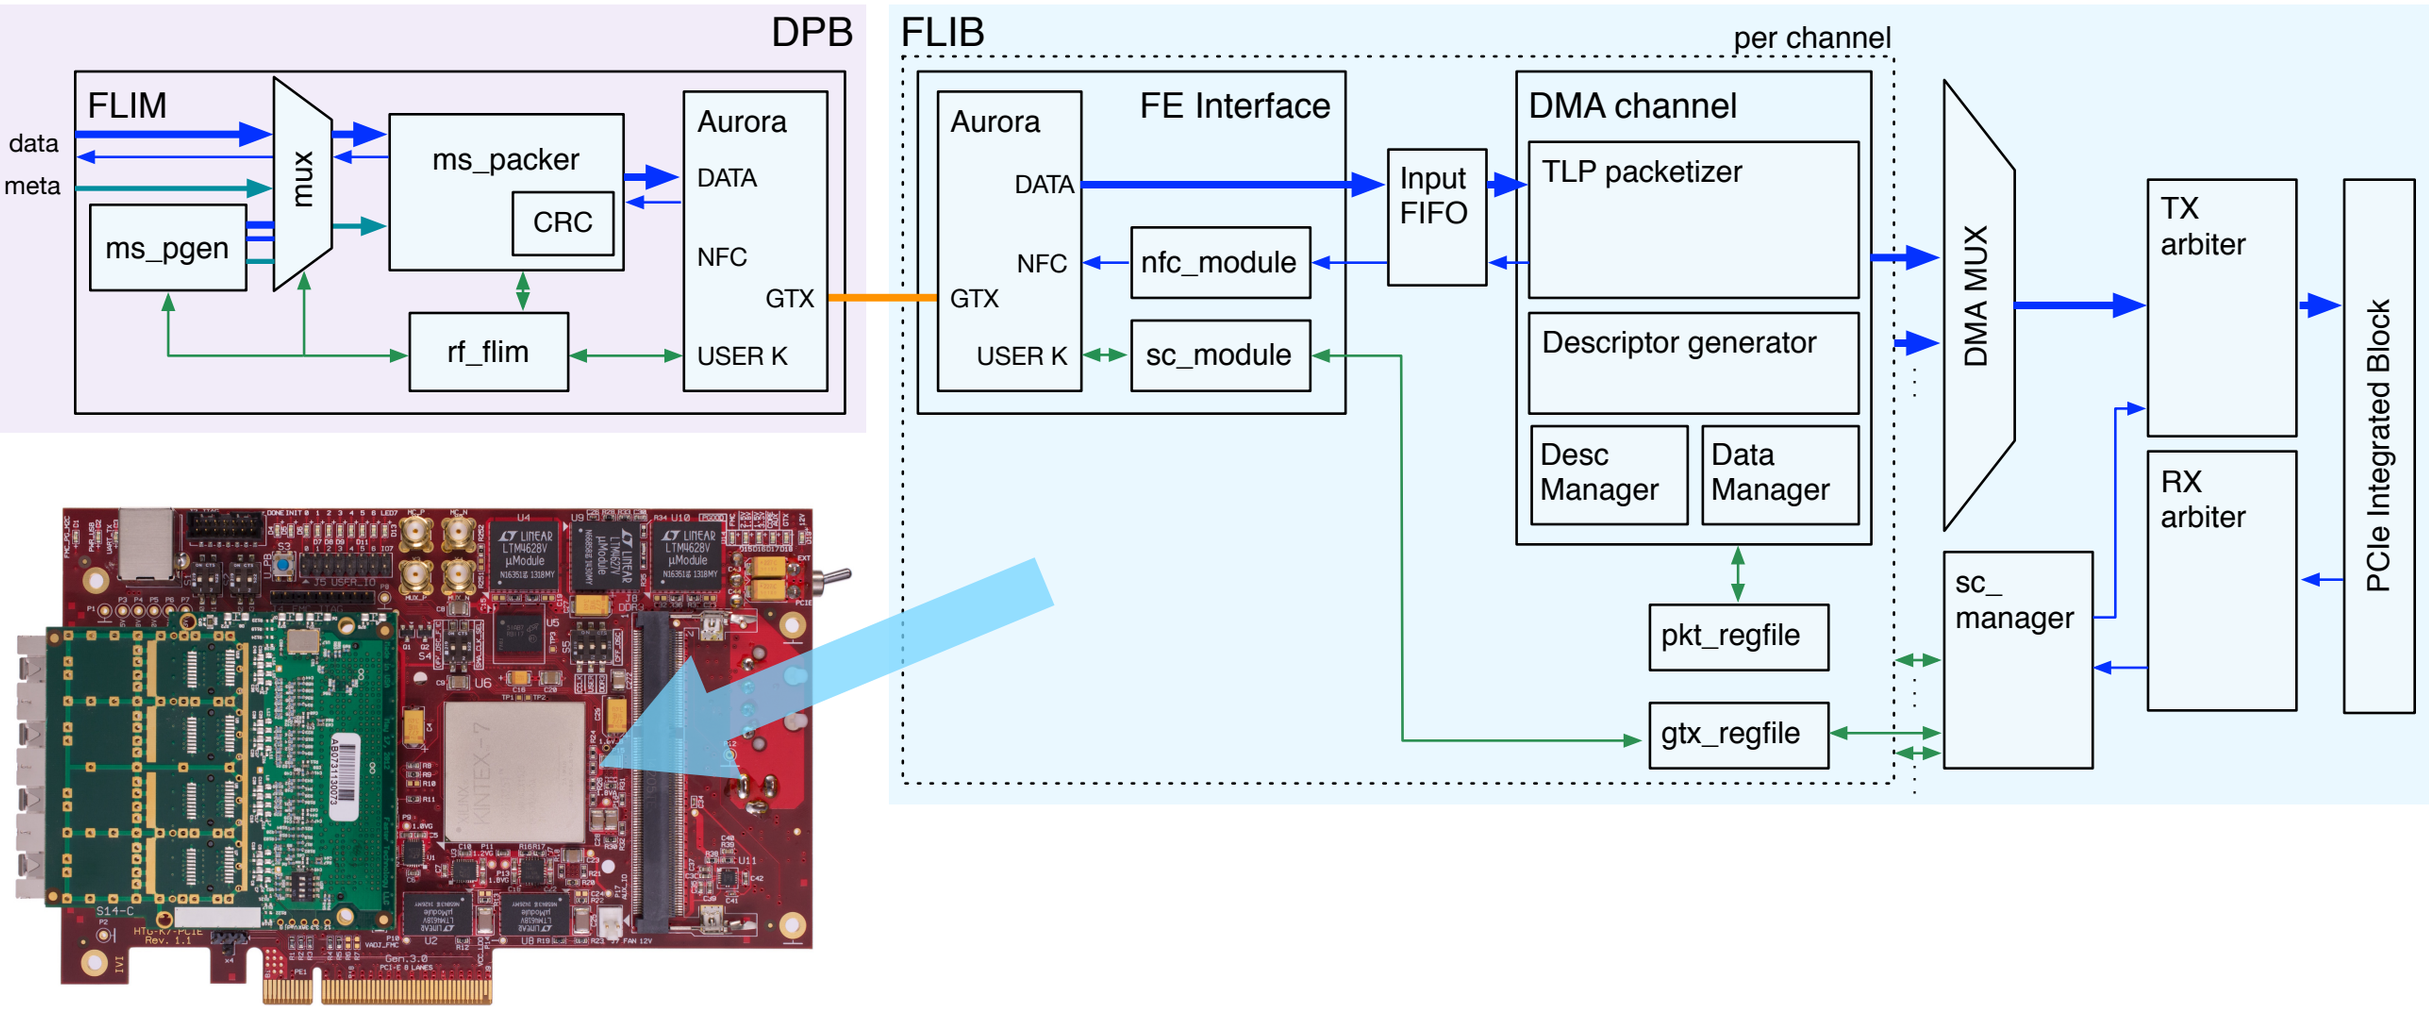

- PCIe FPGA board with custom HDL design
- Dev. platform HTG-K7: PCIe 8x gen2, 8 input links
- Accompanying HDL module implementing link in front-end

DMA Performance

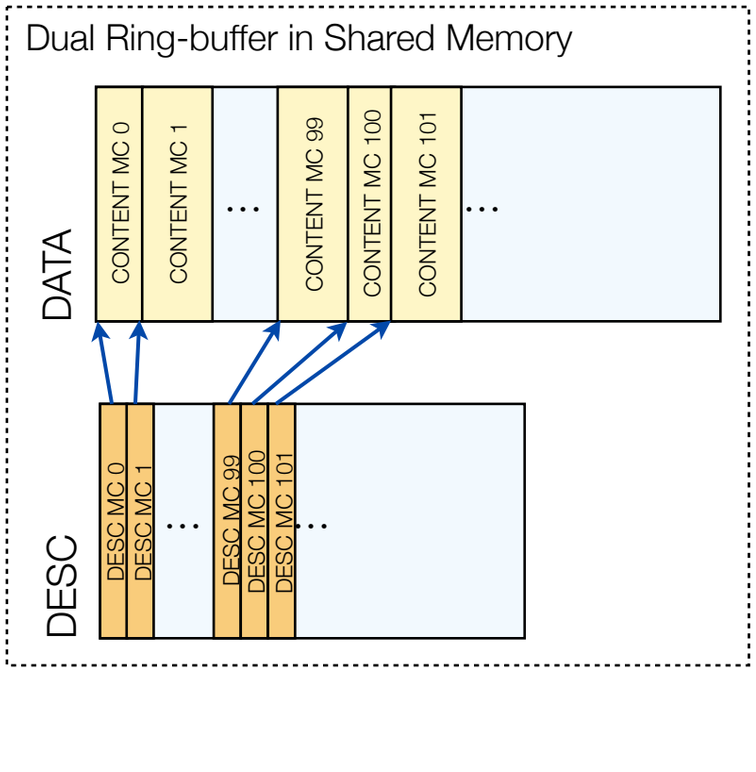

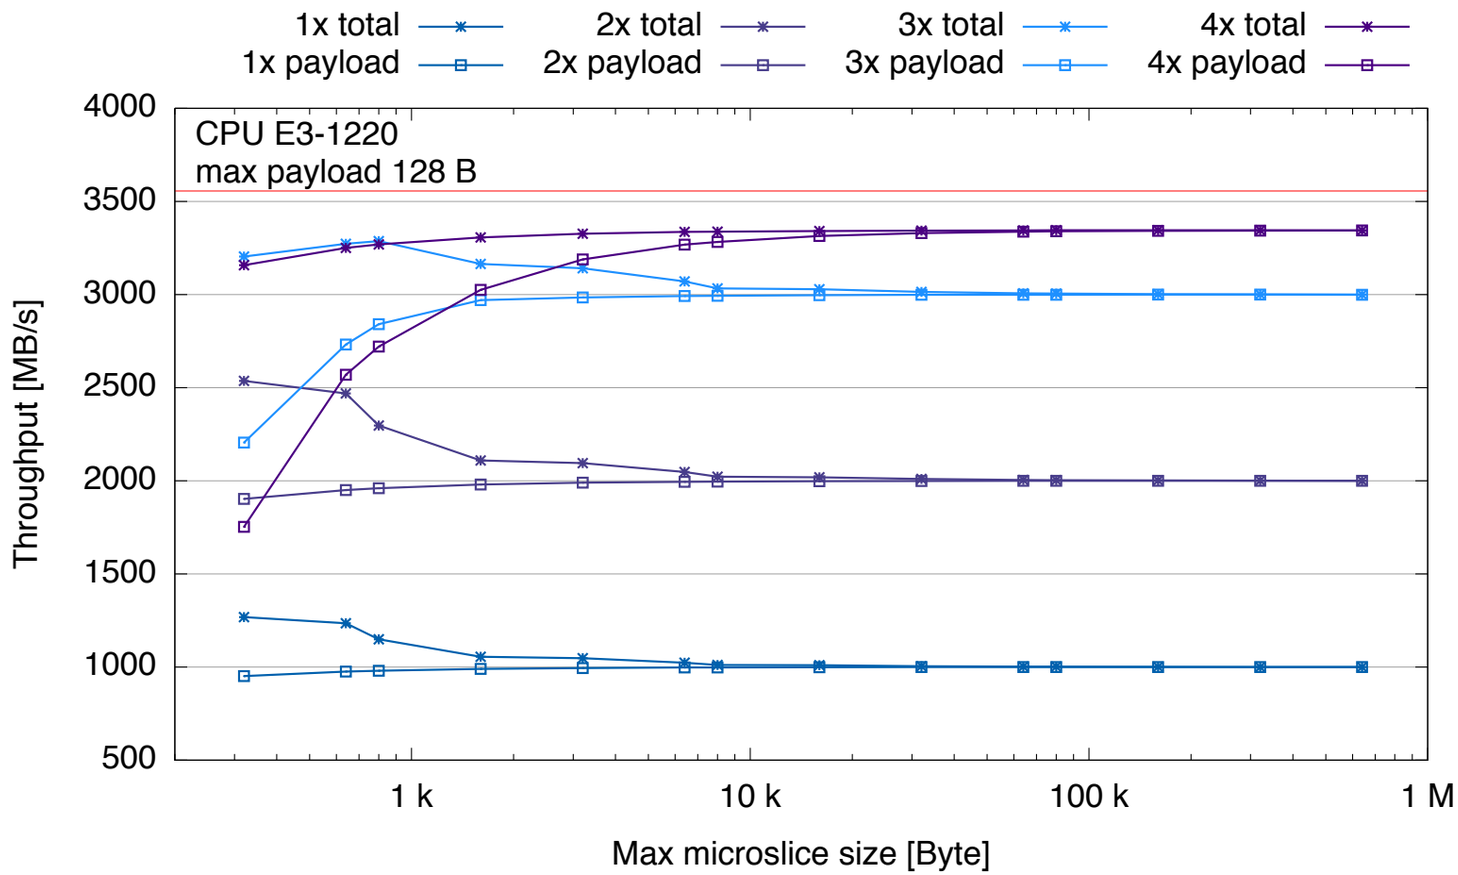

- Full offload DMA engine manages dual ring-buffer structure
  - Data buffer for microslice data content
  - Descriptor buffer for index table and microslice meta data
- Optimised data scheme for zero-copy timeslice building
- Direct DMA to POSIX shared memory / Infiniband RDMA buffers

FLES Data Transport

more info:
J. de Cuveland (HK 15.1)

- 10 GBit/s custom optical link
- HDL front-end module
- Direct DMA to InfiniBand buffers
- Shared memory interface
- Timeslice building
- InfiniBand RDMA, true zero-copy

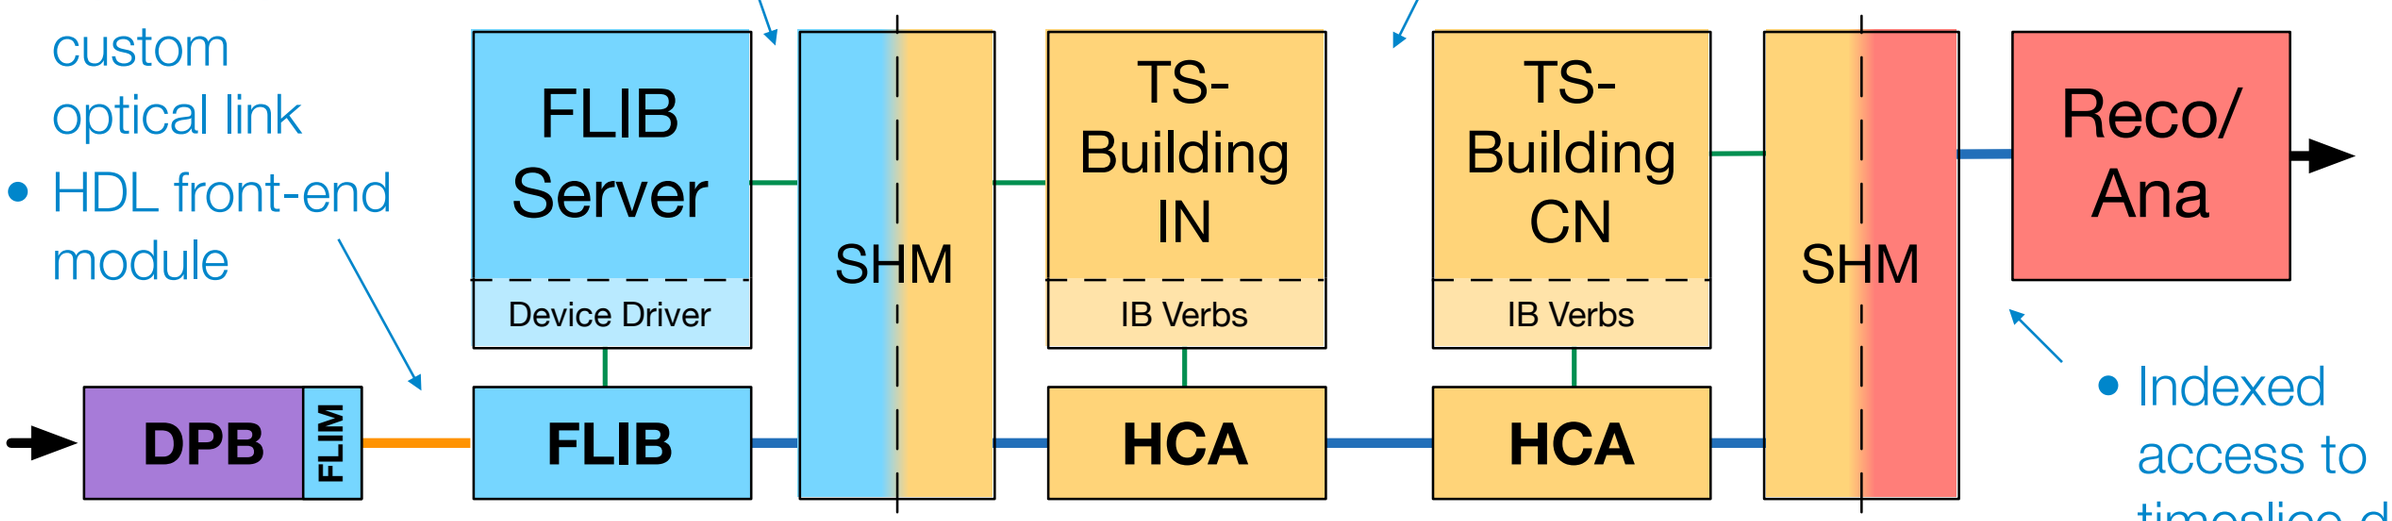

- Indexed access to timeslice data

- Paradigms:
  - Do not copy data in memory
  - Maximize throughput
- Design principle: Work towards final system
  - Create software with the large setup in mind
  - Adopt for small test setups

- Prototype implementation:
  - C++, Boost, IB verbs
- Input Interface integration
  - Publisher process manages FLIB
  - Shared memory interface to timeslice building
  - Holds data and synchronisation structures

FLES Readout at SPS-2016

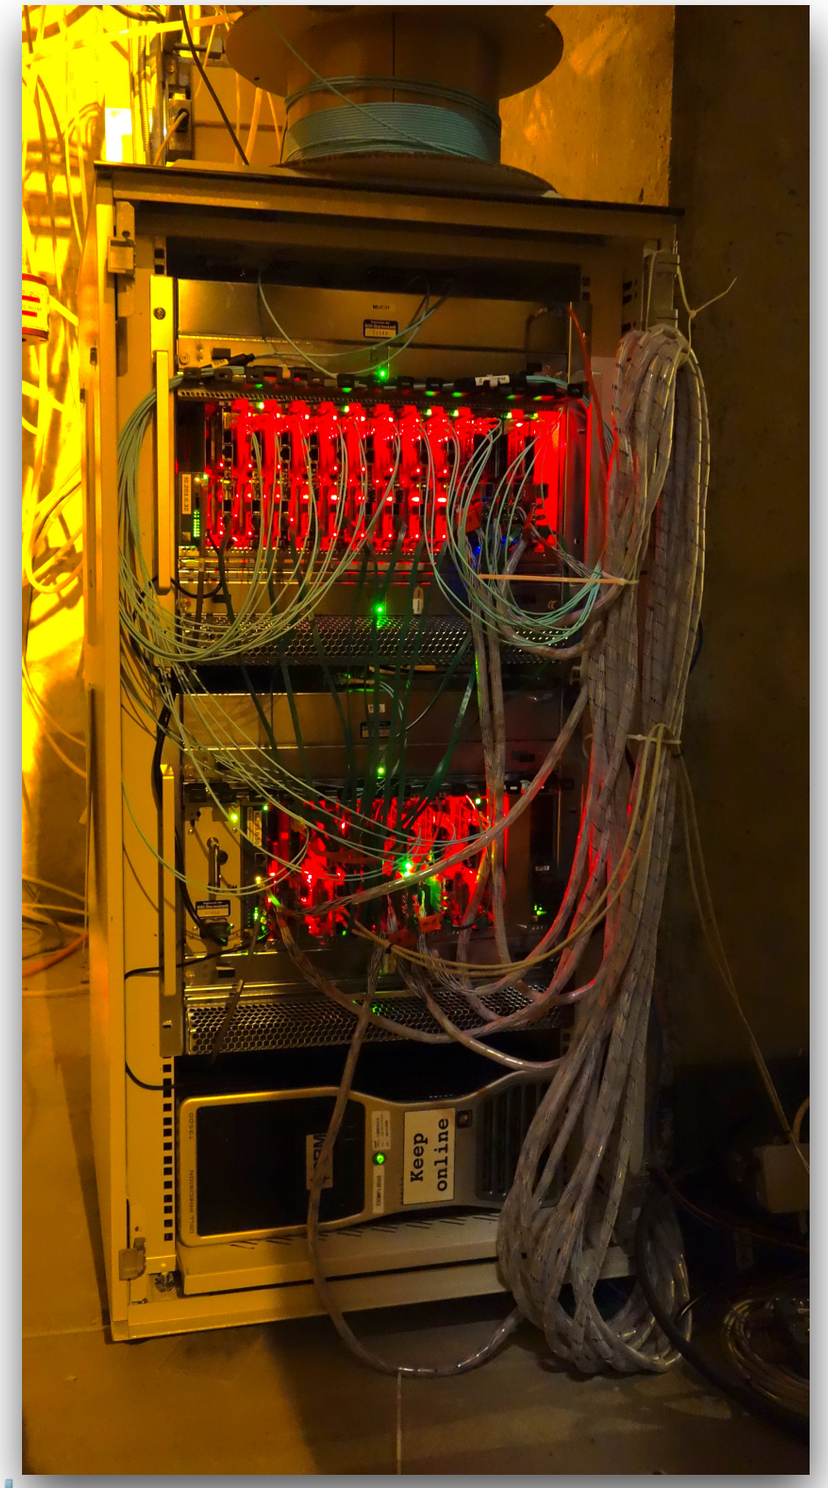

- First common FLES readout of multiple CBM detector systems
- Readout chain from DPB to timeslice like planed for the final system

~50m

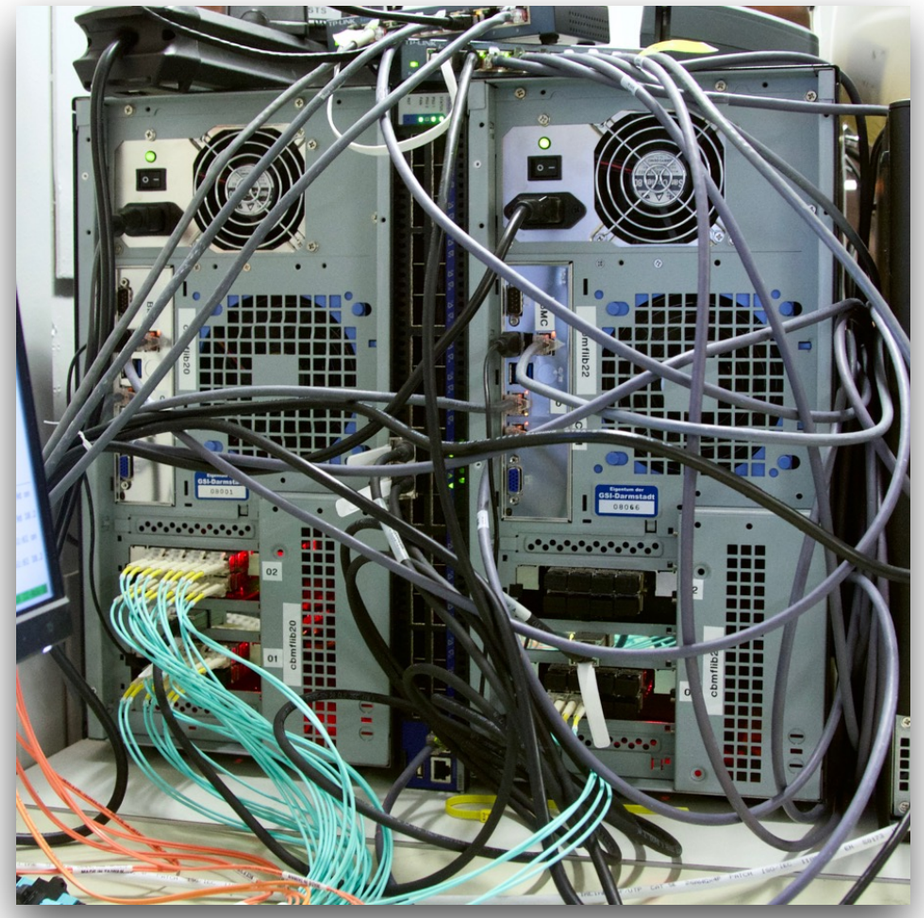

Readout Setup Details

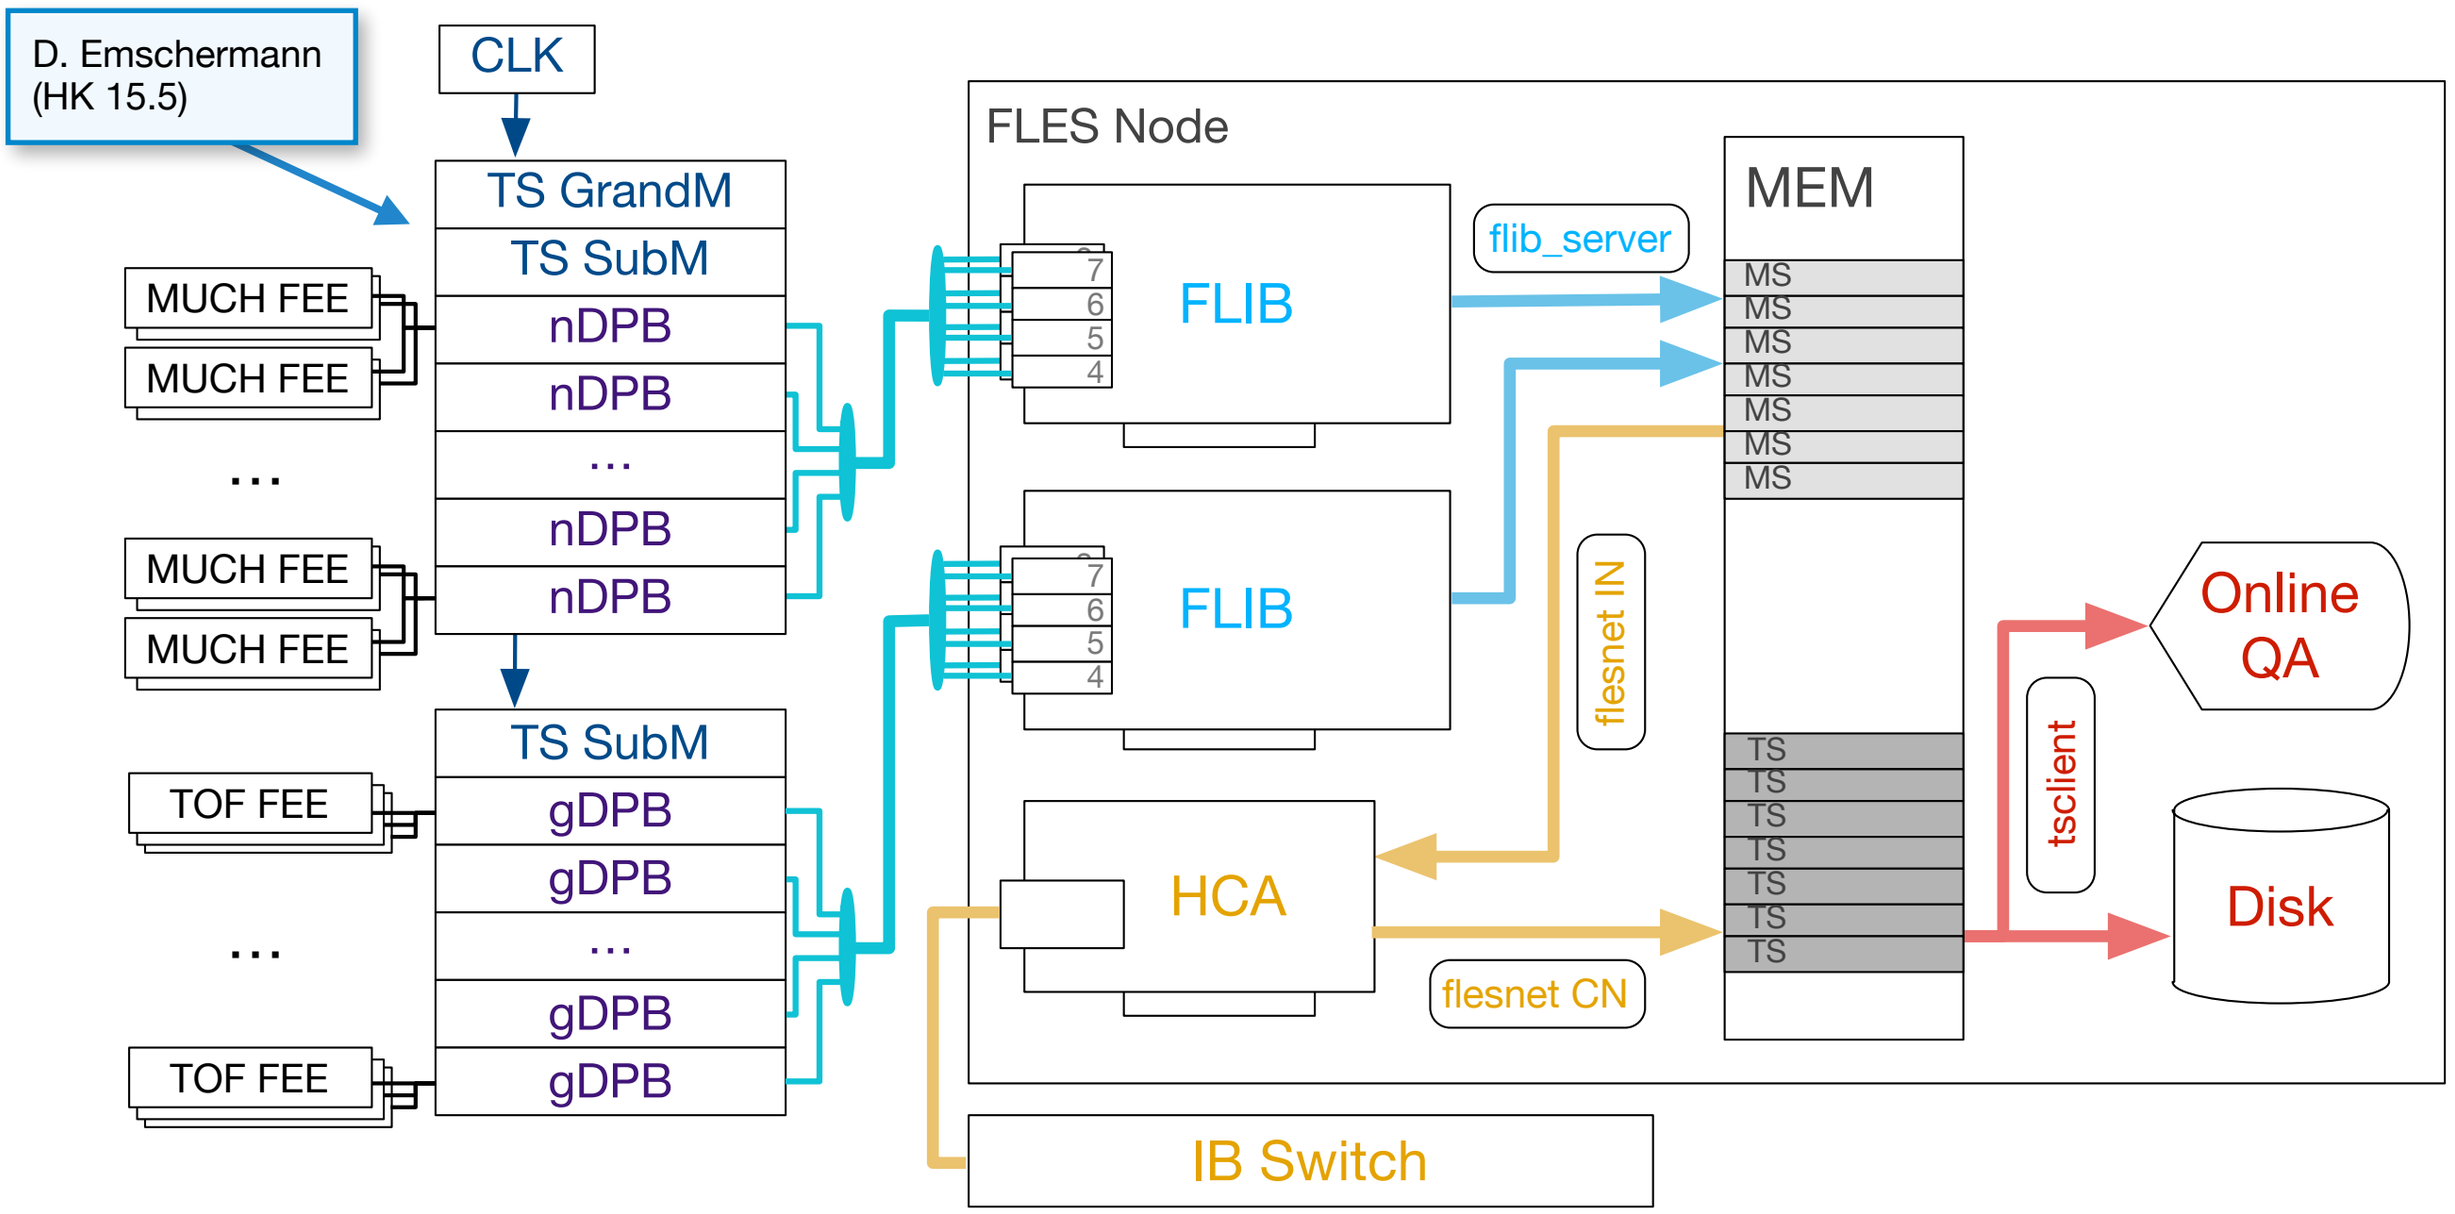

- 1 FLES node with 2 FLIBs running input and compute node processes
  - Up to 16 input links (up to 12 links were used simultaneously)
  - Timeslice building via QDR InfiniBand
  - Timeslice client for checking, archiving and publishing to online analysis

Automated Start of Data Taking

- Experience from lab setups
  - Flexible, but complex system is prone to layer 8 errors
  - Goal: maximise automatisation to simplify usage
- Automated pre-run checks if system is in a clean state
  - Catches common error conditions before anything happens
- Flesnet readout start/stop fully scripted
  - Single command to start all needed processes in correct order
  - All configuration is done in one file
  - Post-run checks for for left over processes and data structures
  - Immediate feedback, very helpful to find rare error conditions

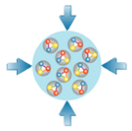

SPS-2016 in Numbers

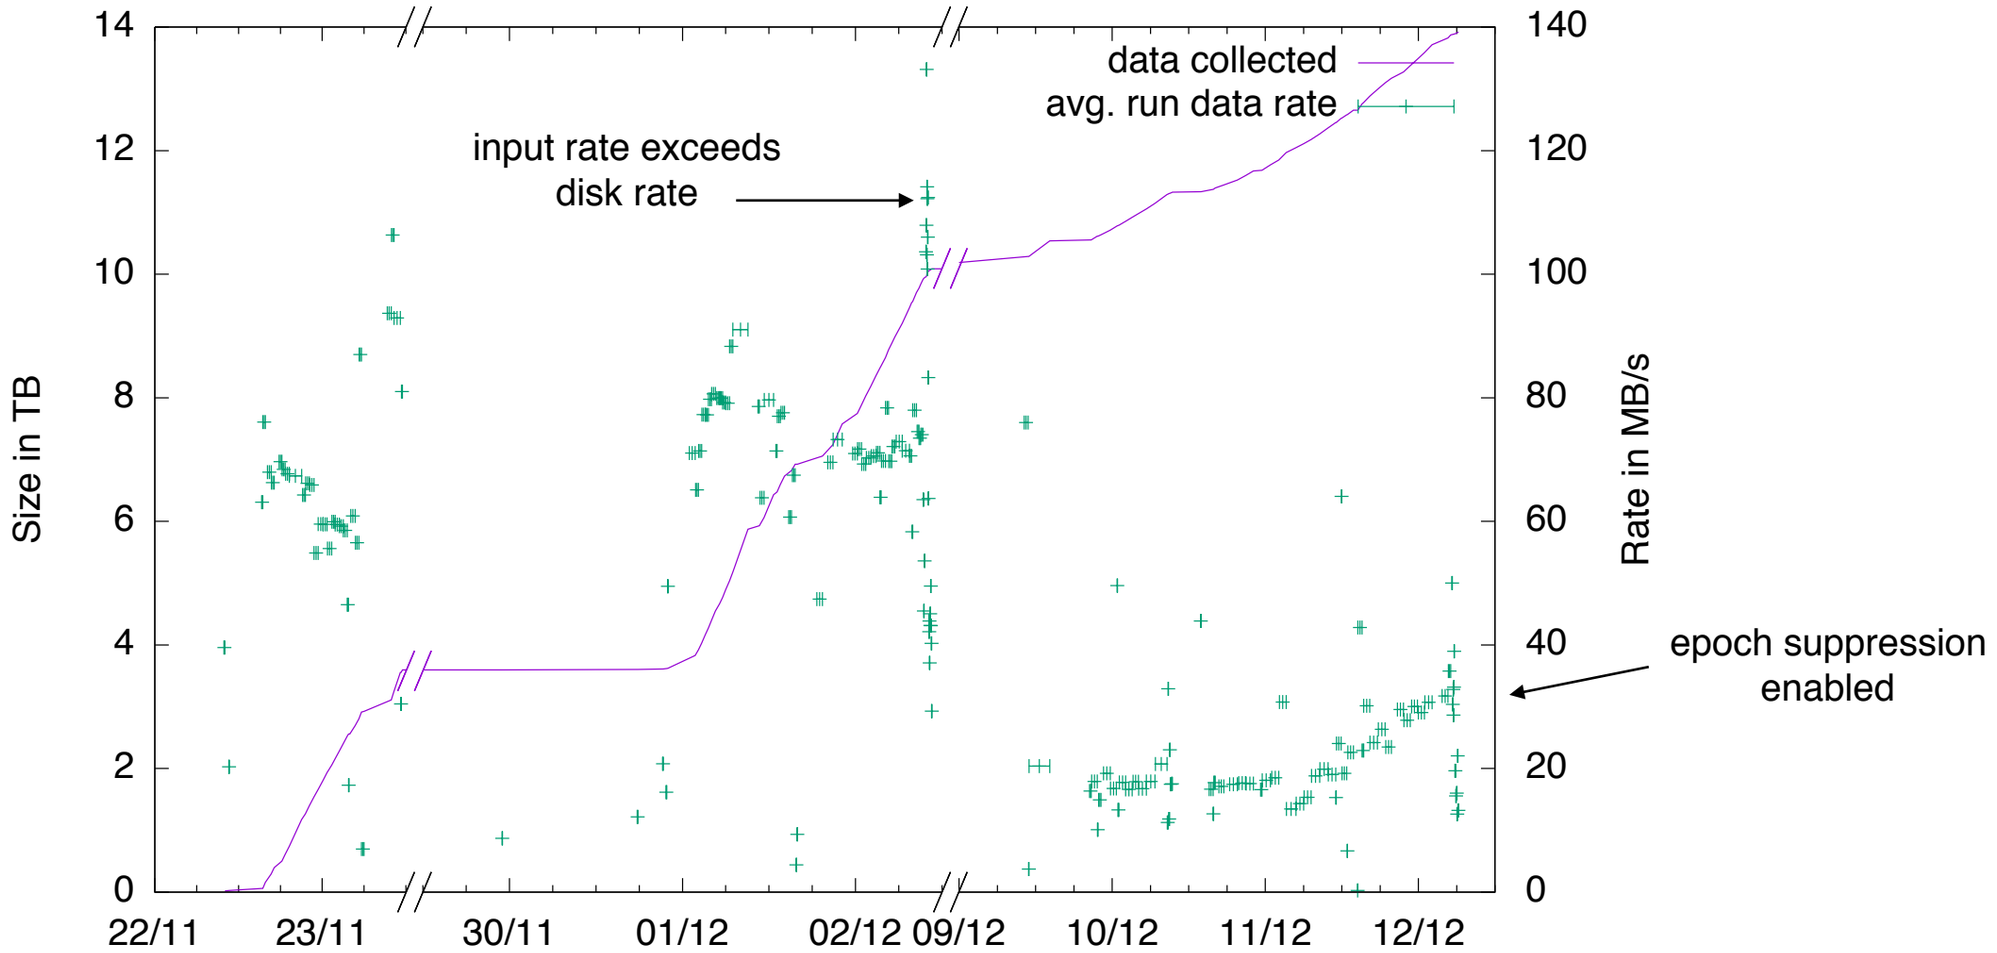

- Data taking divided into in 3 beam periods
- Stable operation of FELS input interface and transport framework over the whole period
- **176 successful runs** with a total of **14 TB timeslice data**
- No run errors related to FLES readout

Run Statistics

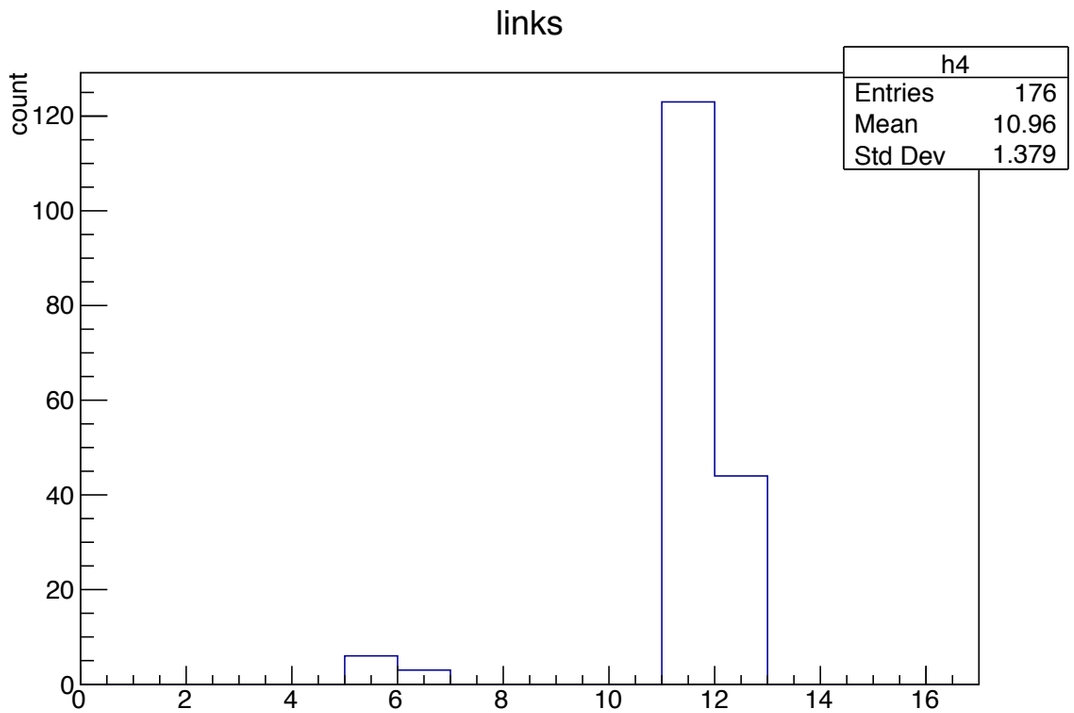

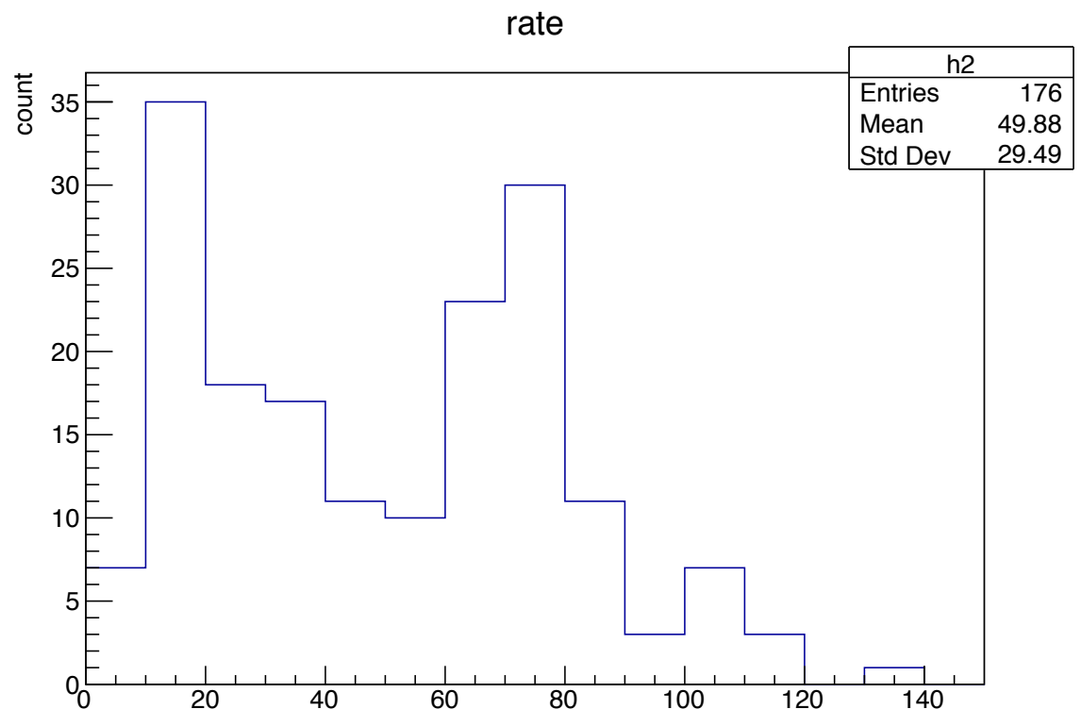

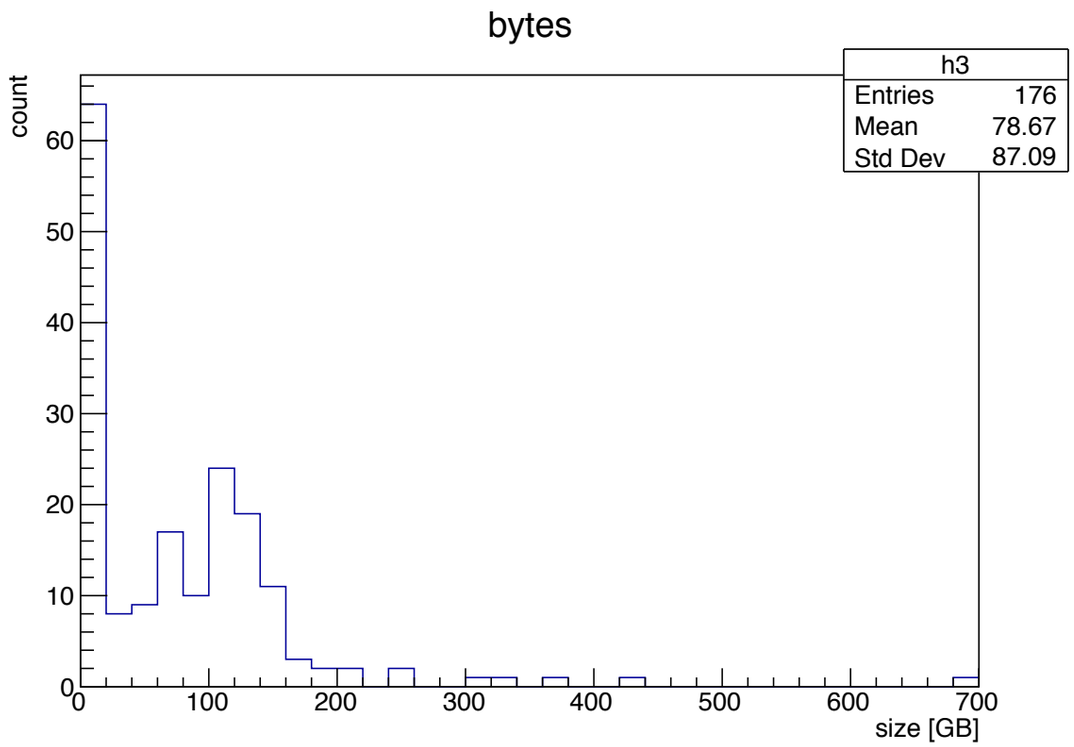

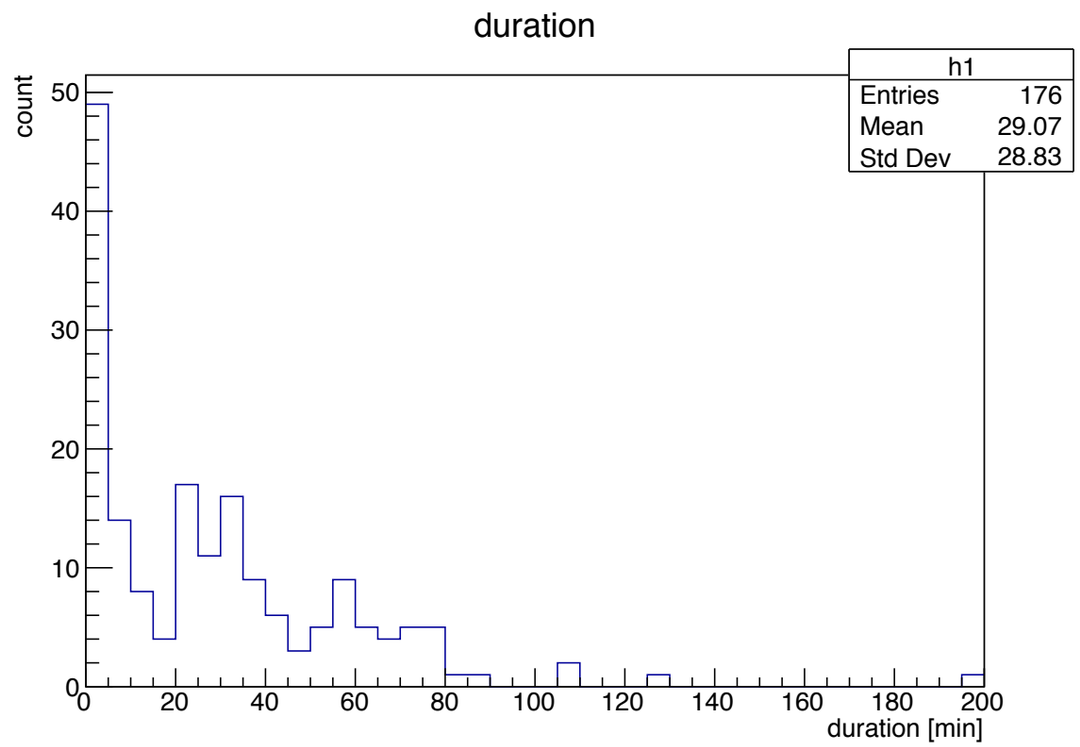

Summary

- Successful evaluation of the FLES input interface under beam conditions
- First common, free-streaming readout of multiple CBM detector systems
- Shows feasibility of input interface concept for final CBM setup
- Good experience with automatisation
- Next steps:
  Extend setups to multiple FLES nodes towards miniCBM

Thanks for your attention

Dirk Hutter
hutter@compeng.uni-frankfurt.de

CBM