Timm M. Steinbeck
Computer Science/Computer Engineering
Kirchhoff Institute of Physics
Ruprecht-Karls-University Heidelberg, Germany

New Experiences with the ALICE High Level Trigger Data Transport Framework

Outline

- Overview ALICE / High Level Trigger (HLT)

- Overview Framework / Components
- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

Overview - ALICE

ALICE High Level Trigger

- ALICE: A Large Ion Collider Experiment
- Heavy-Ion Mode:
  - Up to 15000 particles in event
  - Max. event size >70 MB
  - Max. input rate from TPC: 200 Hz
  - Input data stream: $\leq$ 25 GB/s
  - Output data stream: $\leq$ 1.2GB/s
- Proton-Proton Mode:
  - Max. input rate from TPC: 1kHz
  - Event size $\approx$3 MB

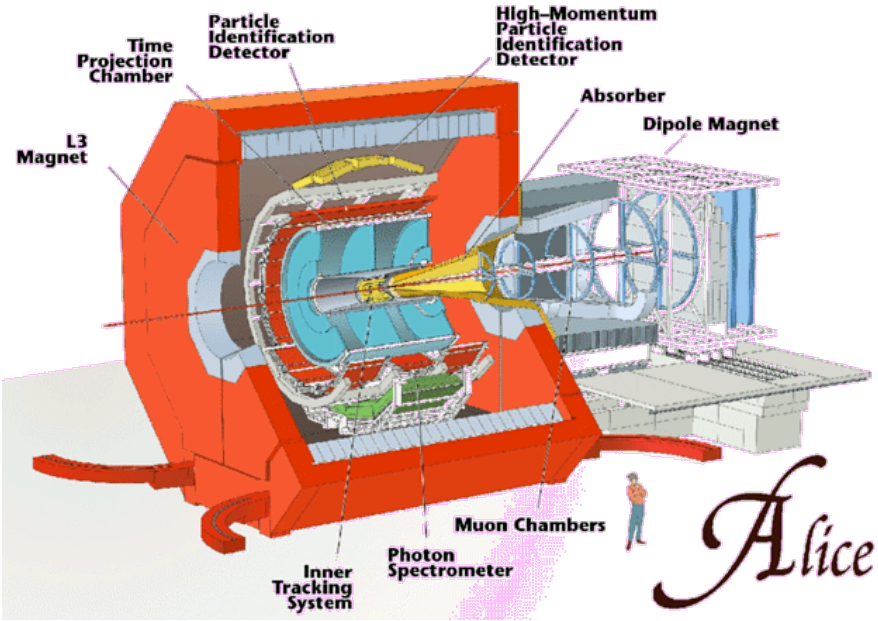

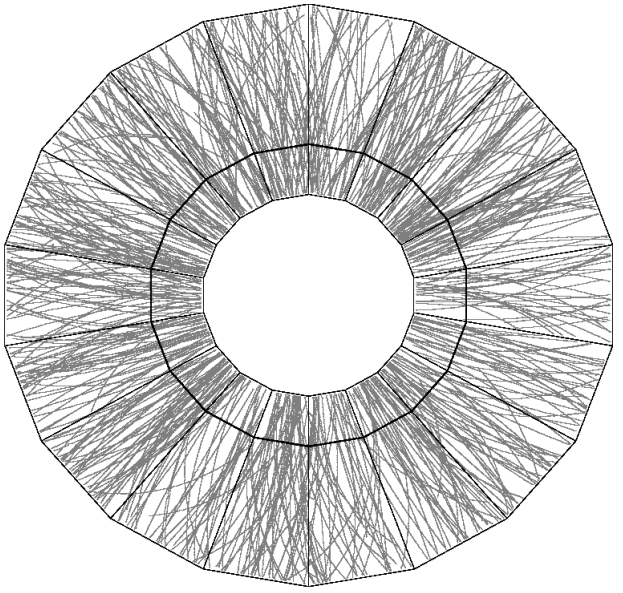

See also Talk by M. Richter

Overview - HLT

ALICE High Level Trigger

- Processing in several steps from raw detector data to full event reconstruction
- Large PC cluster
  - Initially ≈400-500 nodes
  - Arranged in hierarchy levels (HL) that match detector layout and analysis steps
- Exact processing sequence and hierarchy not known
- Flexible and efficient software needed to transport data through cluster
- Framework consisting of independant, communicating components

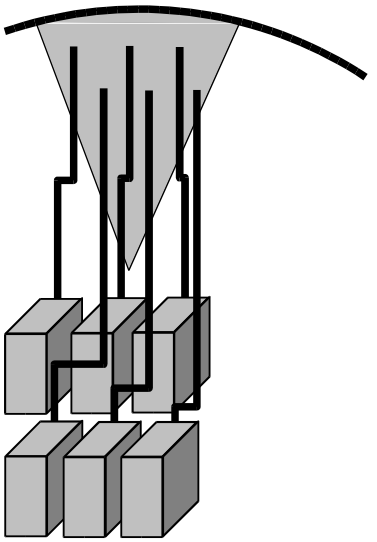

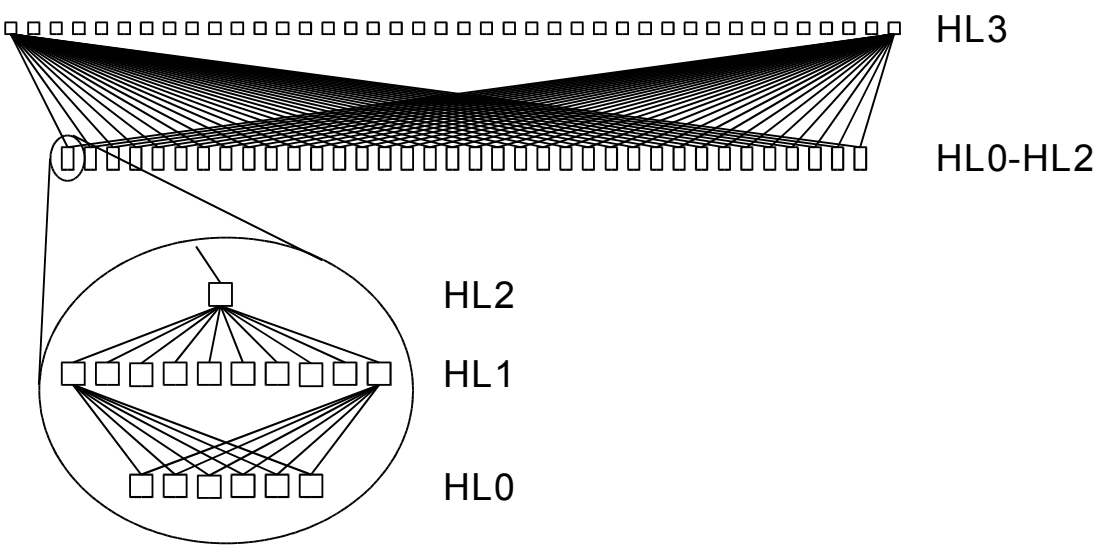

Outline

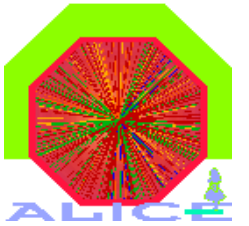

- Overview ALICE / High Level Trigger (HLT)
- Overview Framework / Components
- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

Overview of the Framework

- Framework consisting of independant components
- Communicating via publisher-subscriber interface
  - Data driven architecture
  - Components receive data from other components
  - Process received data
  - Forward own output data to next component
- Named pipes used to exchange small messages and descriptors
- Data is exchanged via shared memory

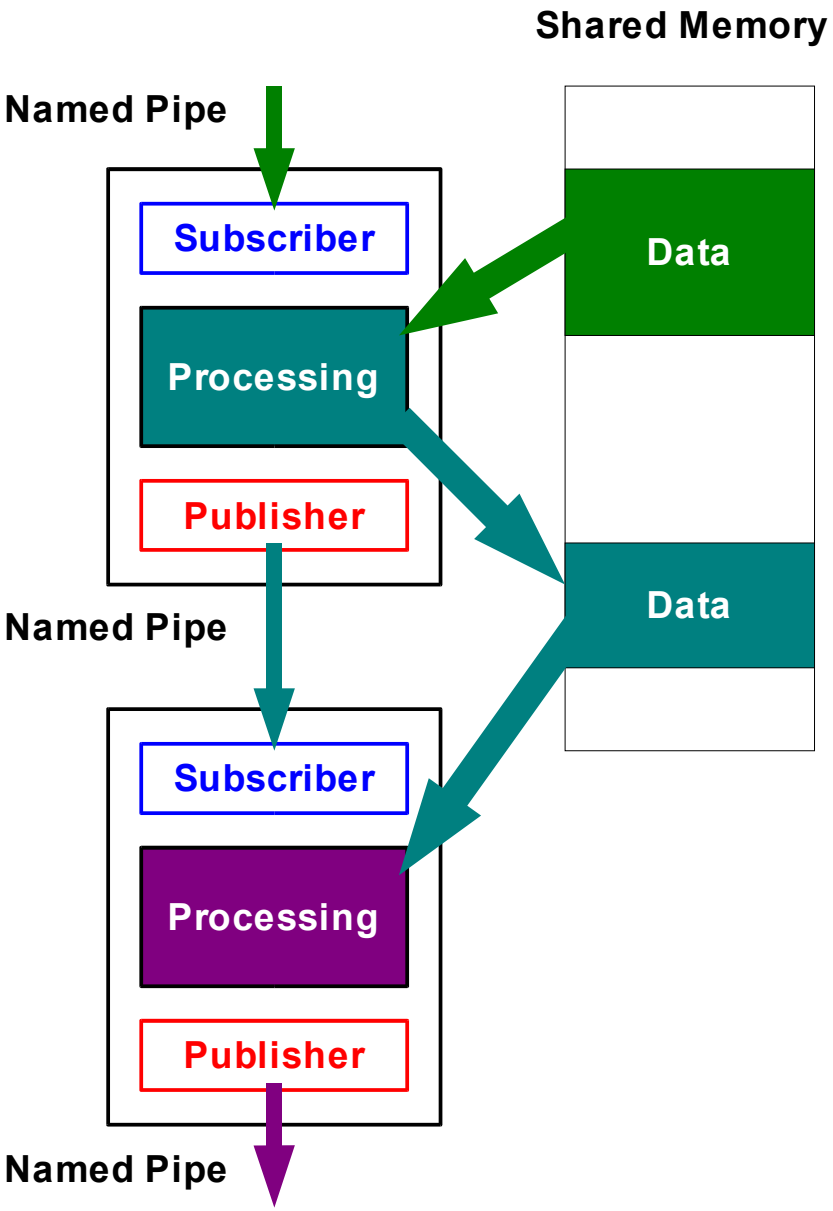

Major design criteria

- Efficiency: CPU Cycles used for analysis
- Flexibility: Different configurations needed
- Fault-Tolerance: PCs are unreliable

Application Components

Templates for application specific components

- Read data from source and publish it (Data Source Template)
- Accept data, process it, publish results (Analysis Template)
- Accept data and process it, e.g. storing (Data Sink Template)

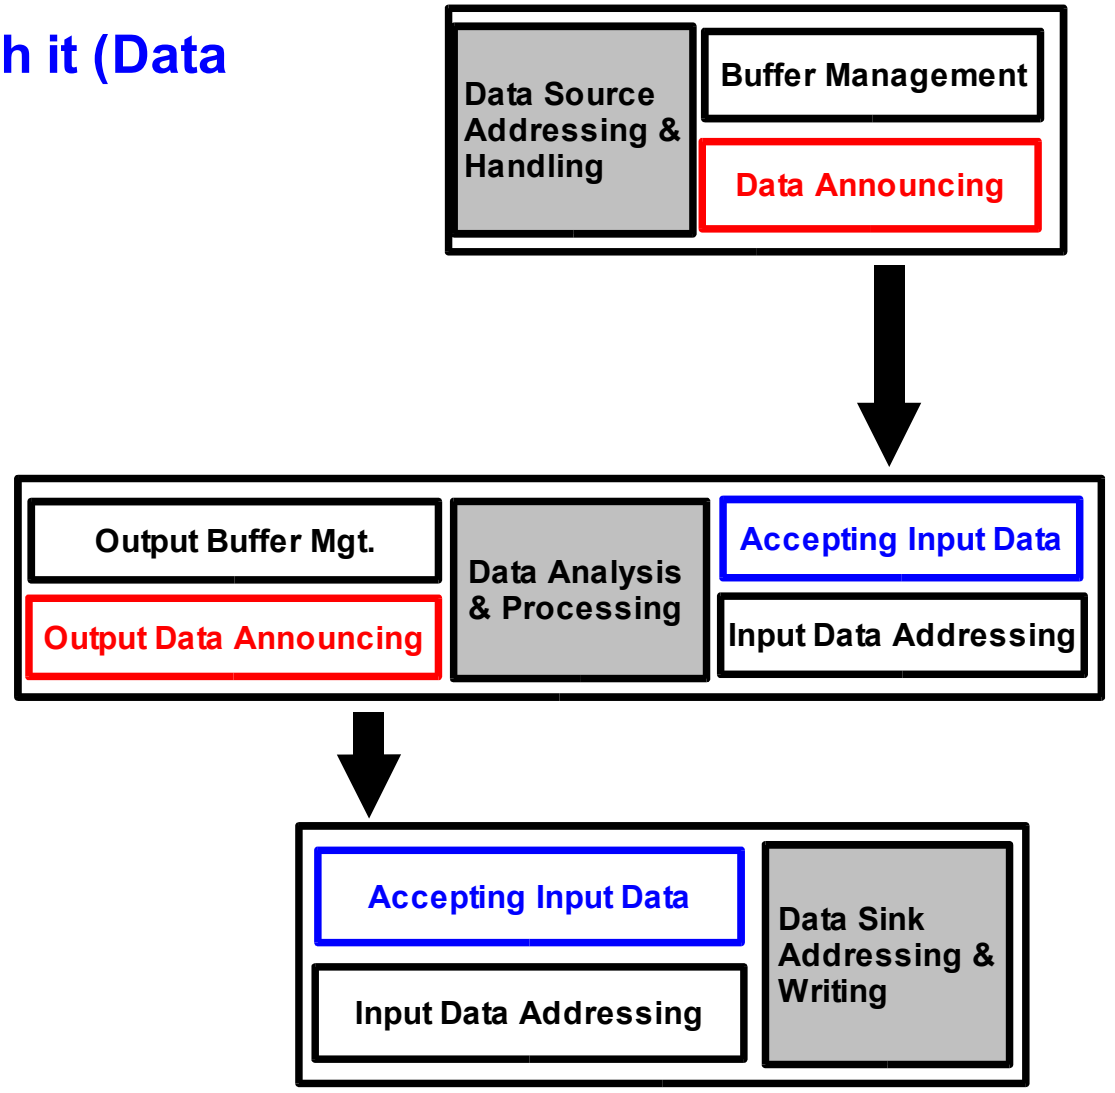

Data Flow Components

Framework contains components to shape the flow of data in a cluster

All components use publisher-subscriber interface

Merge parts of events

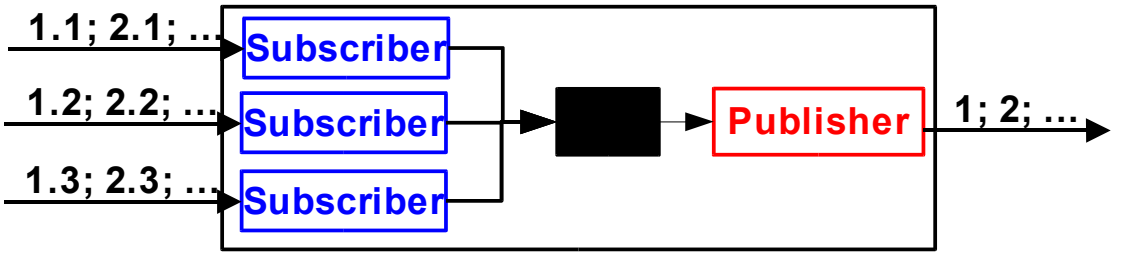

Split and rejoin a data stream (e.g. for load-balancing) (scatter/gather)

Transparently connect components on different nodes (bridges)

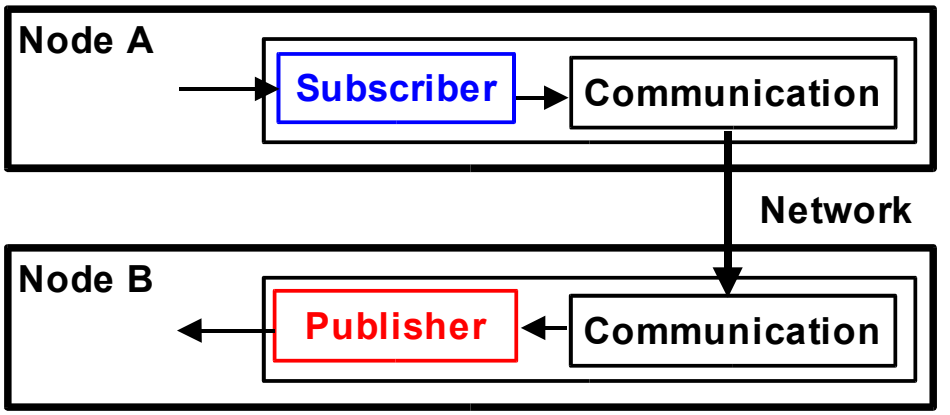

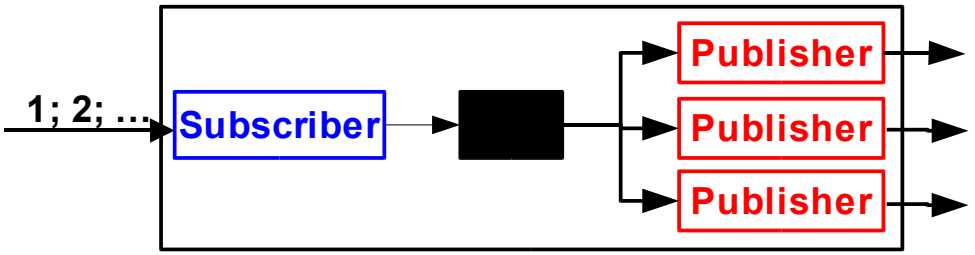

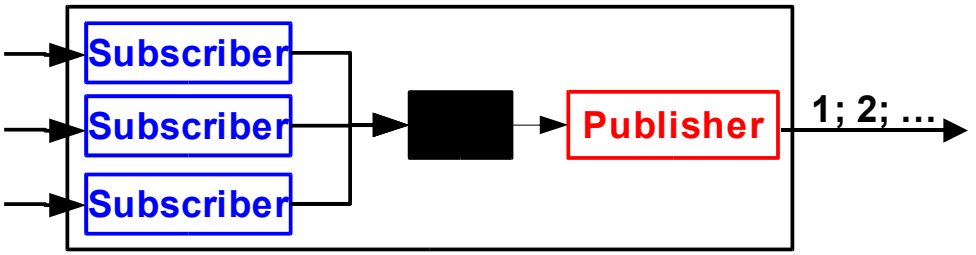

Outline

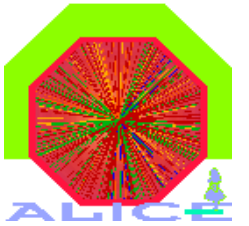

- Overview ALICE / High Level Trigger (HLT)
- Overview Framework / Components

- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

Interface Benchmarks

Benchmark of only Publisher-Subscriber Interface

No shared memory access or anything else

Benchmarks sensitive to L1 cache size

(P4 family only 8 kB)

Maximum rate achieved:

>110 kHz

(Dual Opteron)

Time overhead f. event round-trip: < 18 μs

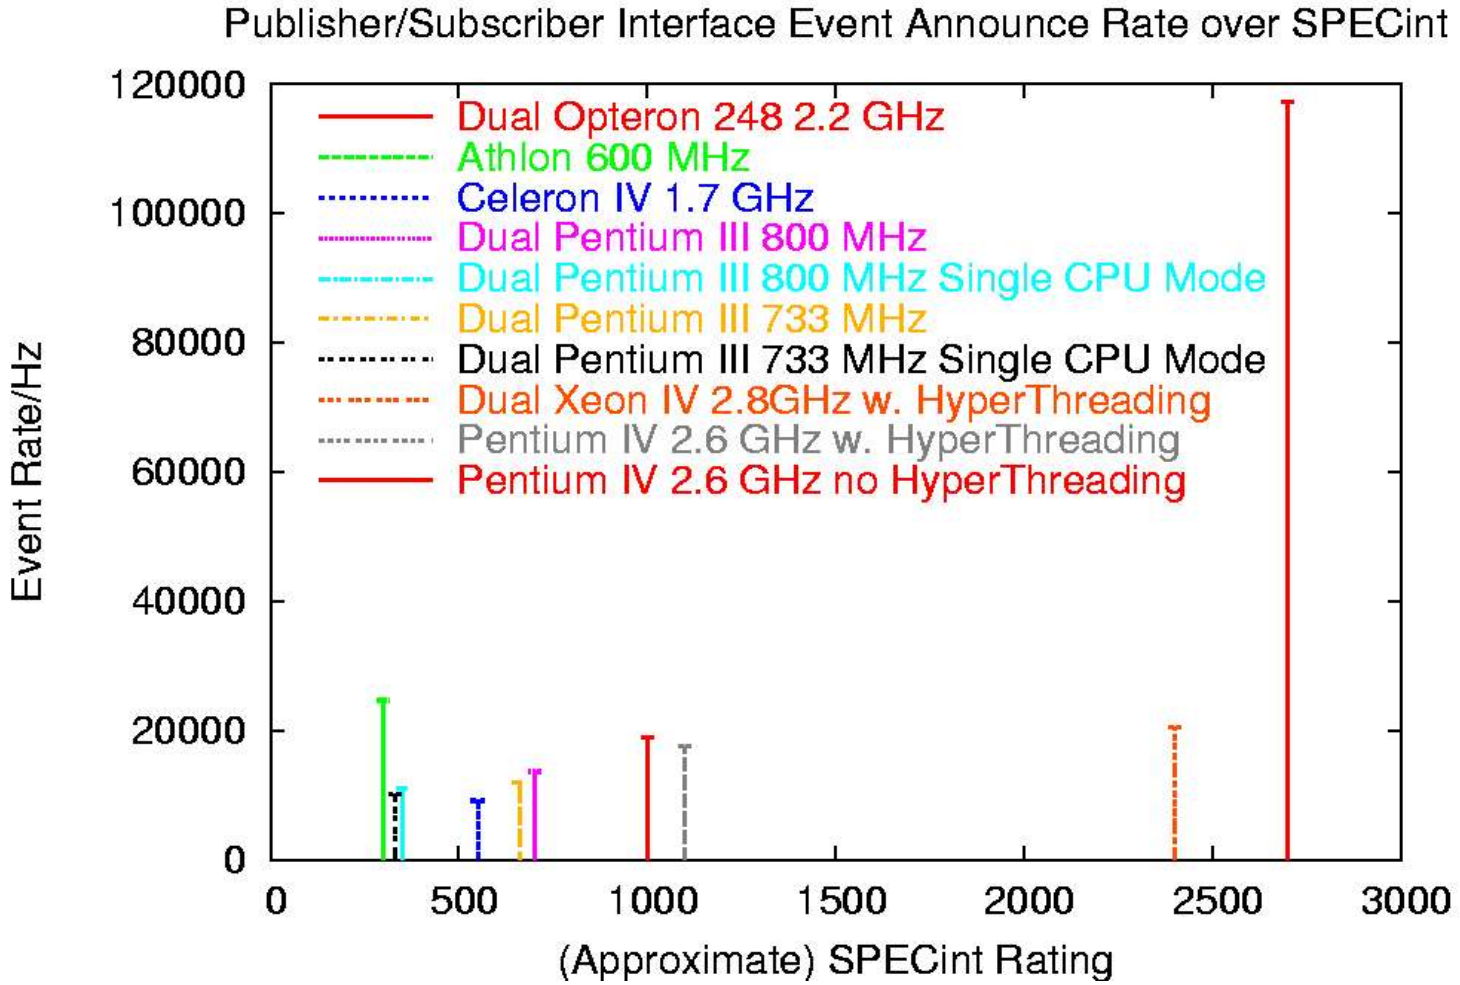

Outline

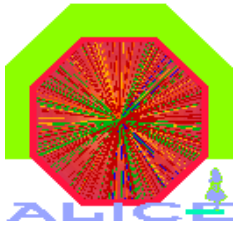

- Overview ALICE / High Level Trigger (HLT)
- Overview Framework / Components
- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

HLT Data Challenge

Long time test of framework running on 7 nodes

800 MHz Dual Pentium 3 nodes

No real analysis done, dummy data and dummy processing components used

All dataflow components utilised

Simulated reduced TPC sector setup (only 4 "readouts" instead of 6, 2 trackers, and 1 global merger)

13 dummy processing components active

Setup running for one month

More than $2*10^9$ events

Achieved rate: > 780 Hz

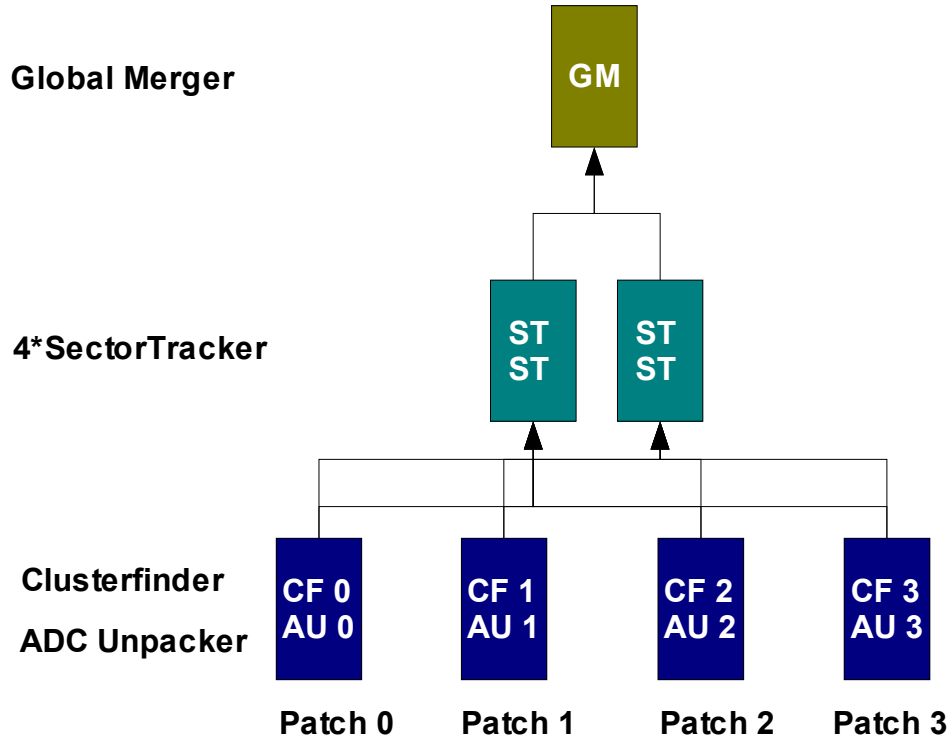

HLT Data Challenge

Event Rate on final merger node measured during runtime of HLT data challenge

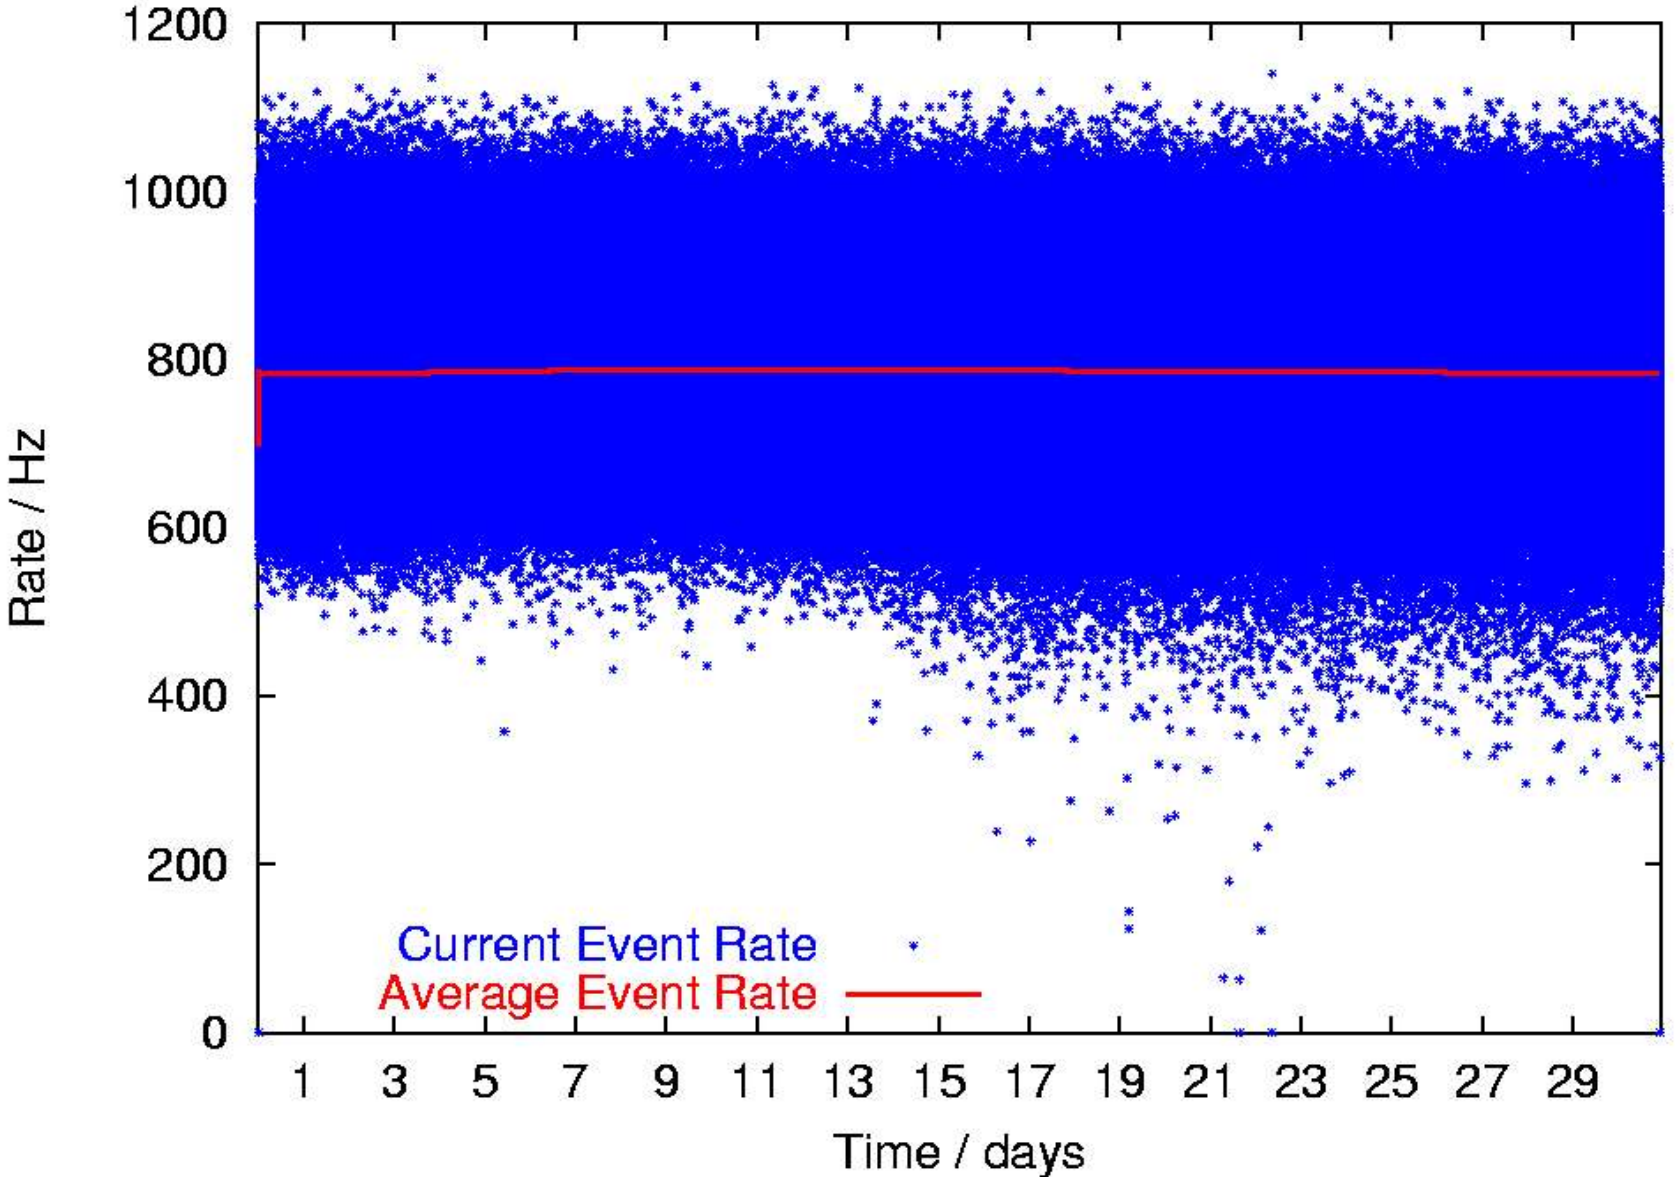

Outline

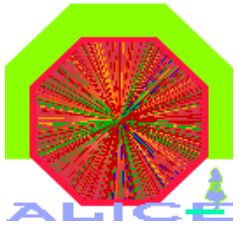

- Overview ALICE / High Level Trigger (HLT)
- Overview Framework / Components
- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

TPC Sector Beamtest

HLT being used during beamtest of TPC sector prototype

Main goals:

- **Test interfaces between DAQ and HLT**
- **Obtain some real world operation experience for HLT**

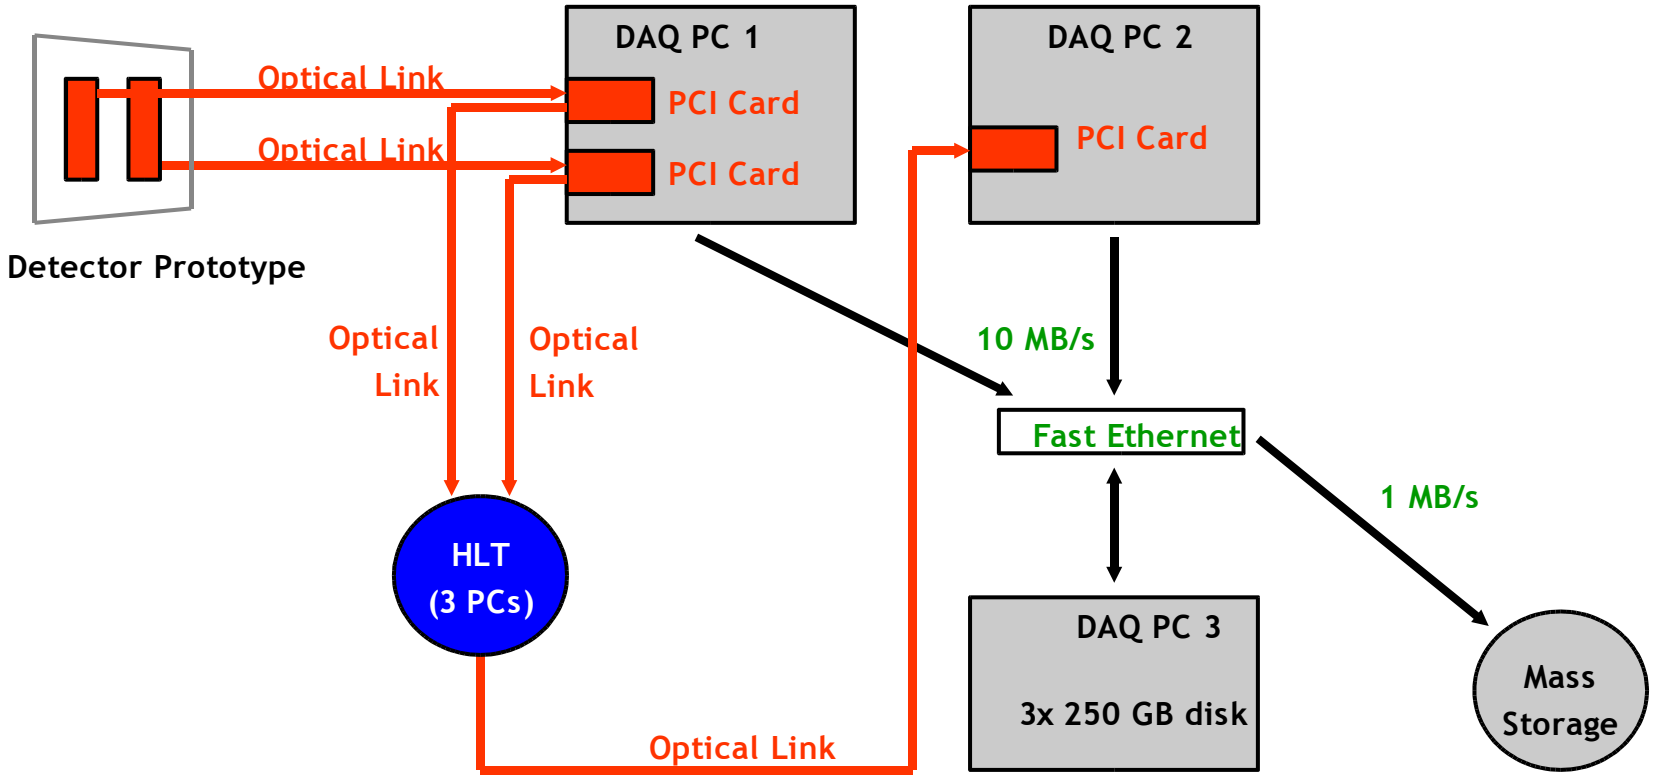

TPC Sector Beamtest

3 Node Mini HLT Cluster used

933 MHz Dual Pentium 3 PCs

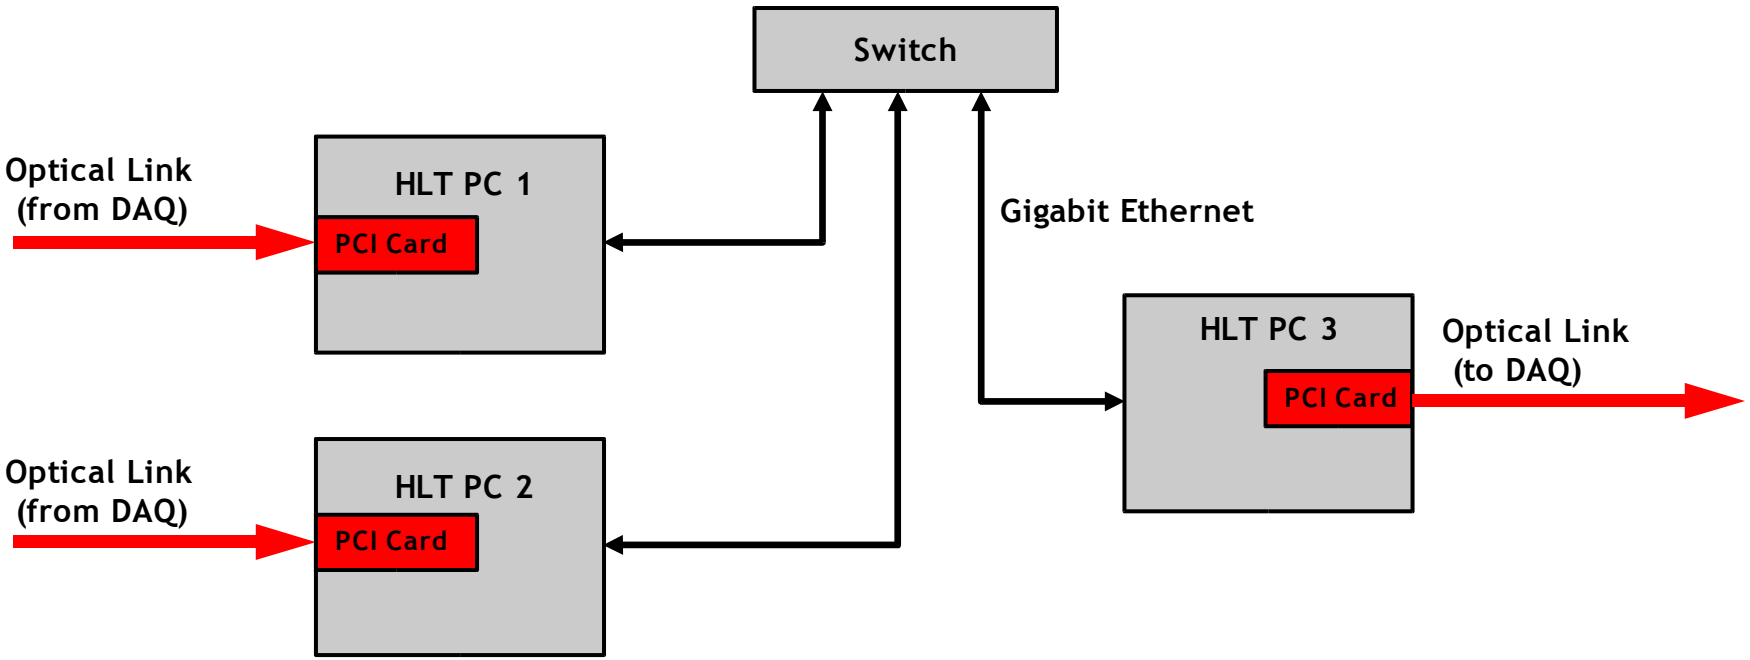

TPC Sector Beamtest

HLT Software Setup

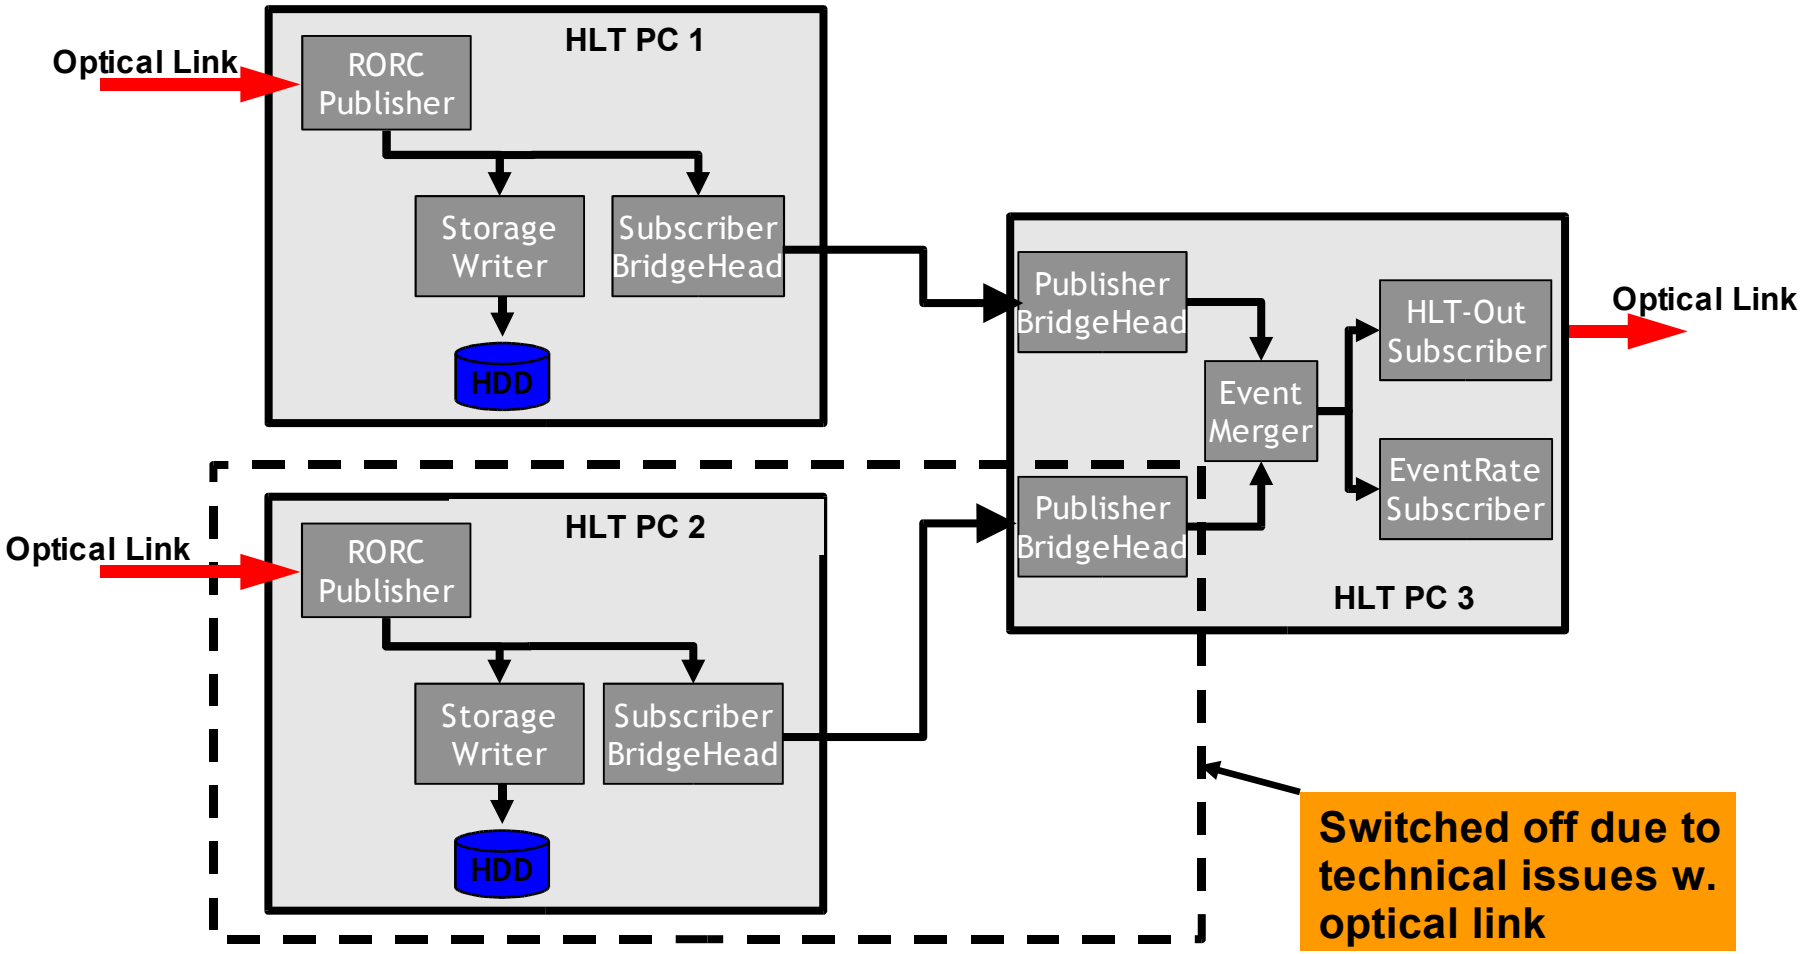

- *EventRateSubscriber* used to measure event rate
- *StorageWriter* used to save data locally on disc

TPC Sector Beamtest

Beamtest HLT Summary

- Software contained only “data flow” components in order to test the interfaces from/to DAQ
- No online analysis done
- Interfaces to DAQ successfully tested, works

Other Issues (outside HLT scope)

- Event rate < 5 Hz due to trigger system
- Event merging not possible due to technical issues w. optical link during HLT active time

Outline

- Overview ALICE / High Level Trigger (HLT)
- Overview Framework / Components
- Benchmarks & Tests
  - Framework Interface Benchmark
  - HLT Data Challenge
  - TPC Beamtest Experience
- Summary & Conclusion

HLT Summary

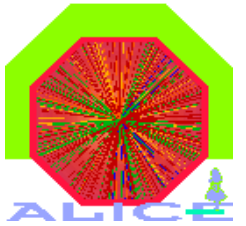

HLT Framework is maturing

Can already be used in projects/testbeams/...

- **Has been used in real beamtest**

Stable

- **Has been running for a month continously**

Performance has improved significantly

- **Enough for Alice HLT already now**

Component approach has already been useful

- **Easy to create configurations for short tests**
- **Flexible configuration changes during testbeam**