LHCb Trigger and Data Acquisition System

Beat Jost
Cern / EP

Presentation given at the
11th IEEE NPSS Real Time Conference
June 14-18, 1999 Santa Fe, NM

Outline

- Introduction
- General Trigger/DAQ Architecture
- Trigger/DAQ functional components
  - Selected Topics
    - Level-1 Trigger
    - Event-Building Network Simulation
- Summary

Introduction to LHCb

- Special purpose experiment to measure precisely CP violation parameters in the $B\overline{B}$ system
- Detector is a single-arm spectrometer with one dipole
- Total b-quark production rate is ~75 kHz
- Expected rate from inelastic p-p collisions is ~15 MHz
- Branching ratios of interesting channels range between 10^{-5}-10^{-4} giving interesting physics rate of ~5 Hz

LHCb in Numbers

Number of Channels	900000
Bunch crossing rate	40 MHz
Level-0 accept rate	1 MHz
Level-1 accept rate	40 kHz
Readout Rate	40 kHz
Event Size	100 kB
Event Building Bandwidth	4 GB/s
Level-2 accept rate	~5 kHz
Level-3 accept rate	~200 Hz
Level-2/3 CPU Power	$2 \cdot 10^6$ MIPS
Data rate to Storage	20 MB/s

LHCb Detector

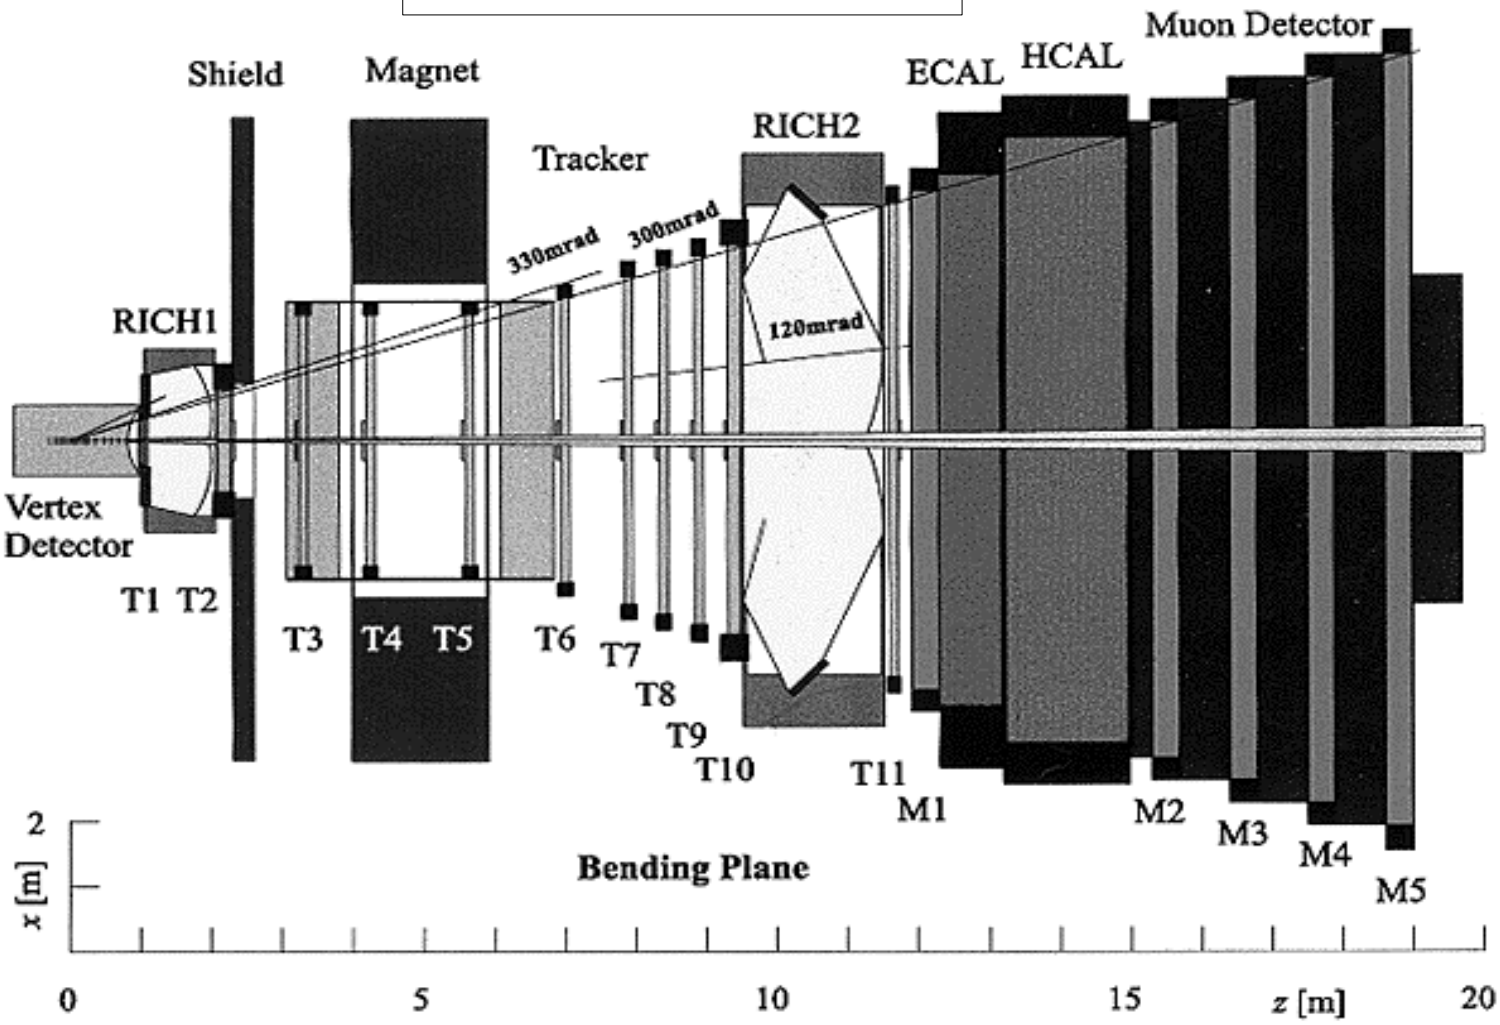

Typical Interesting Event

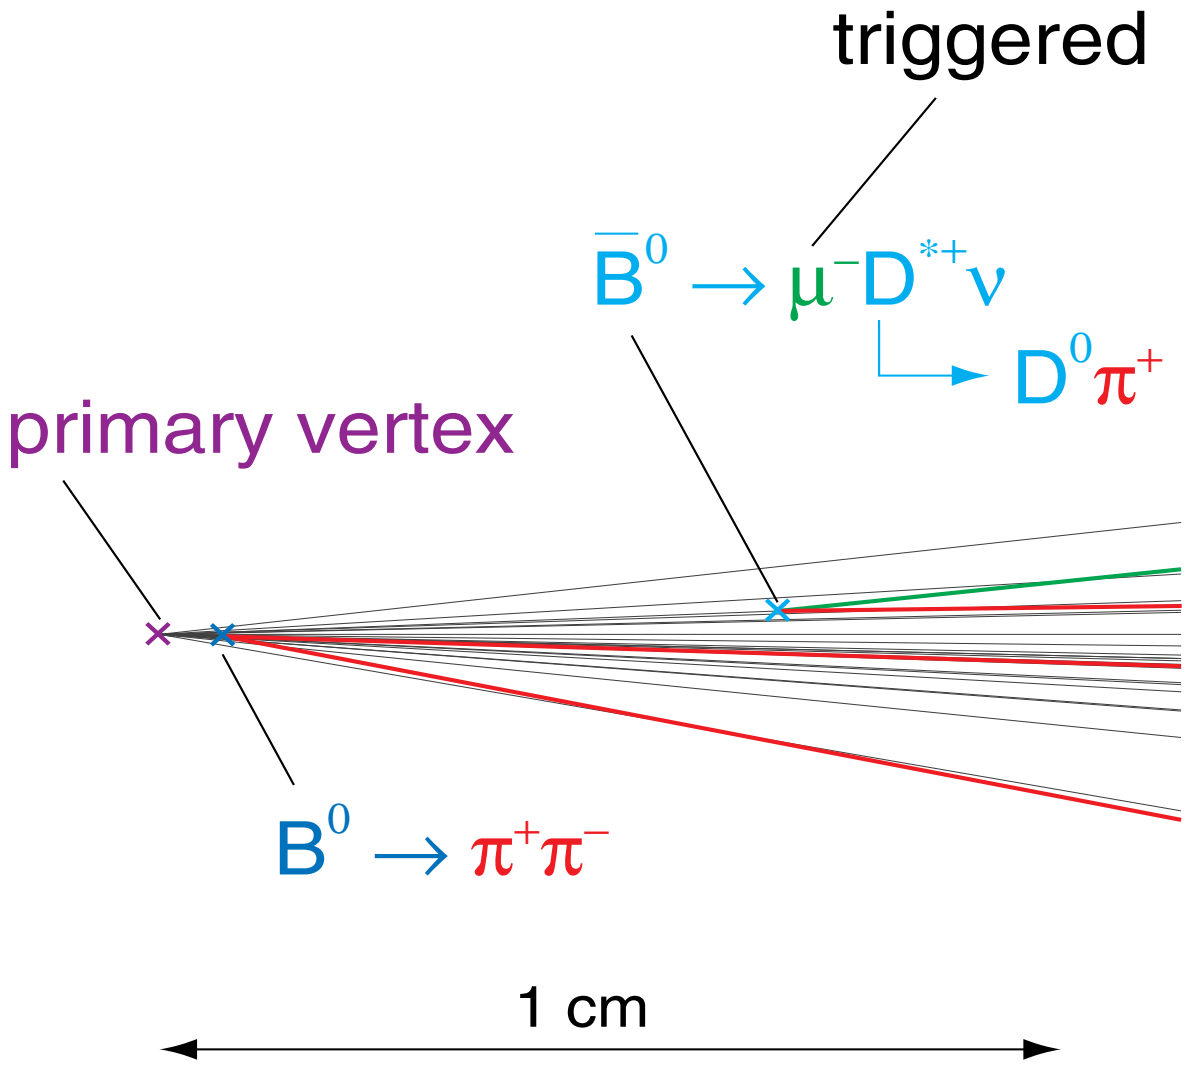

Trigger/DAQ Architecture

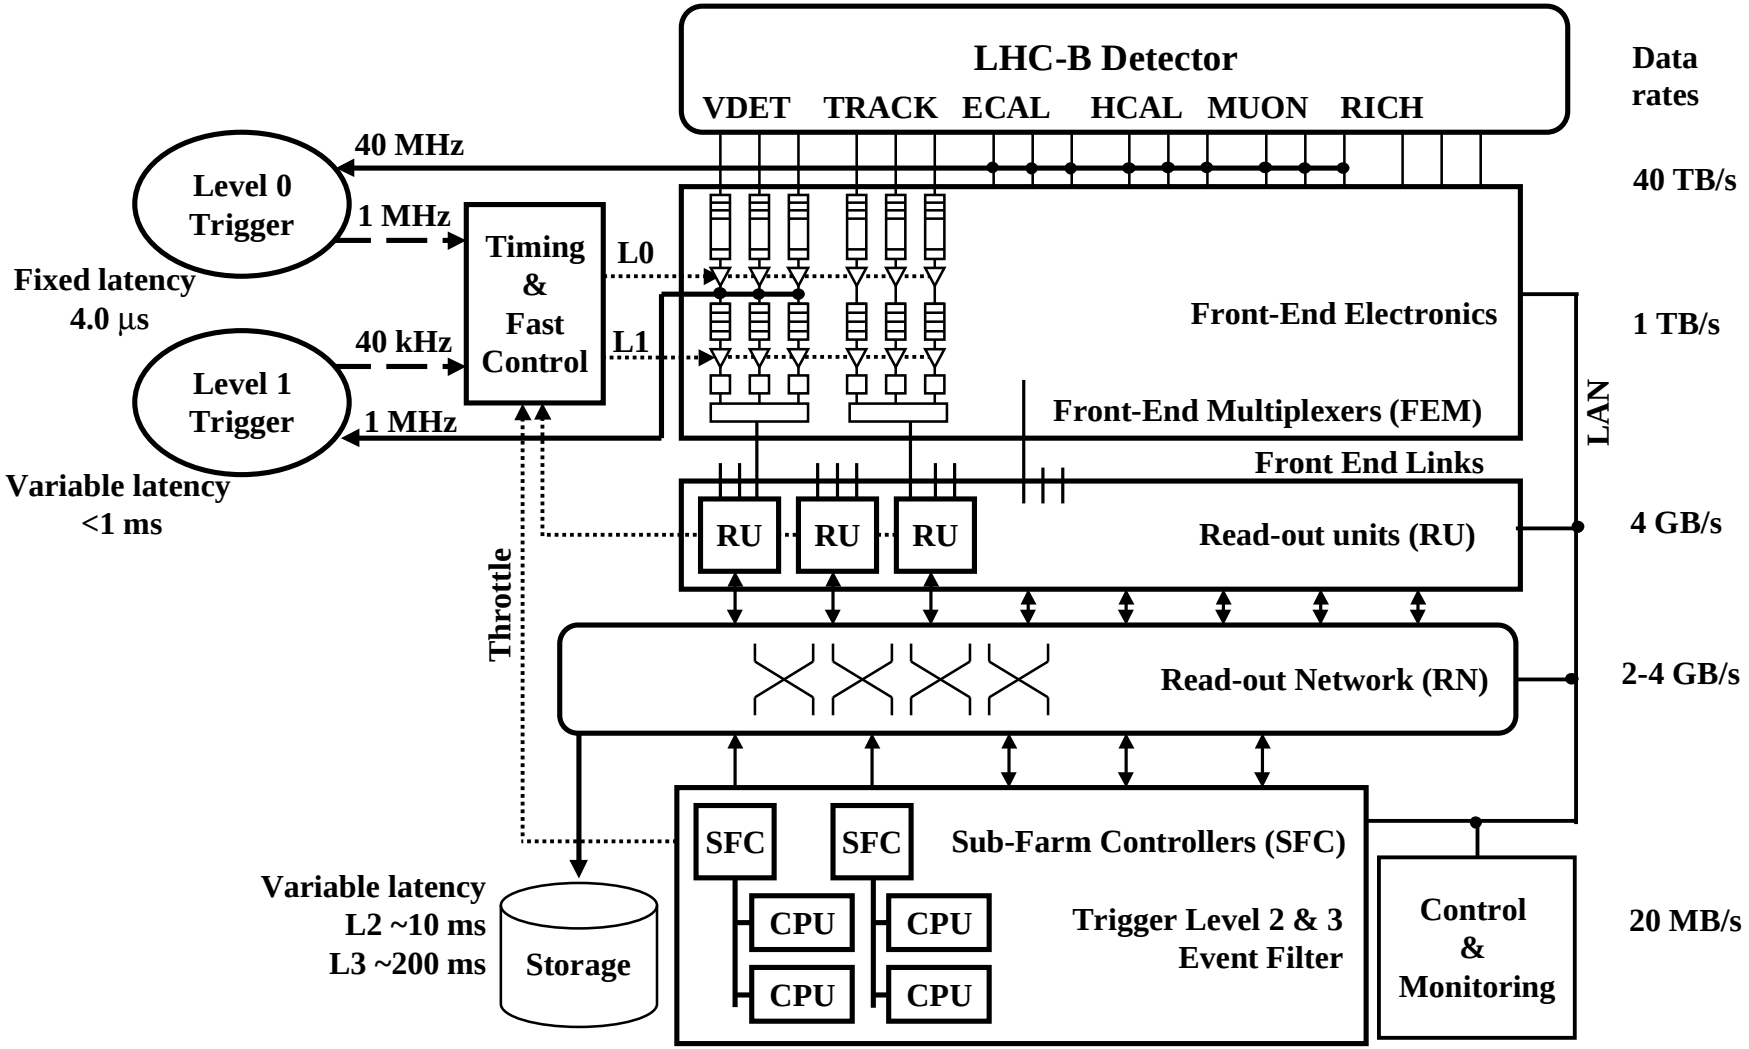

Front-End Electronics

- Data Buffering for Level-0 latency
- Data Buffering for Level-1 latency
- Digitization and Zero Suppression
- Front-end Multiplexing onto Front-end links

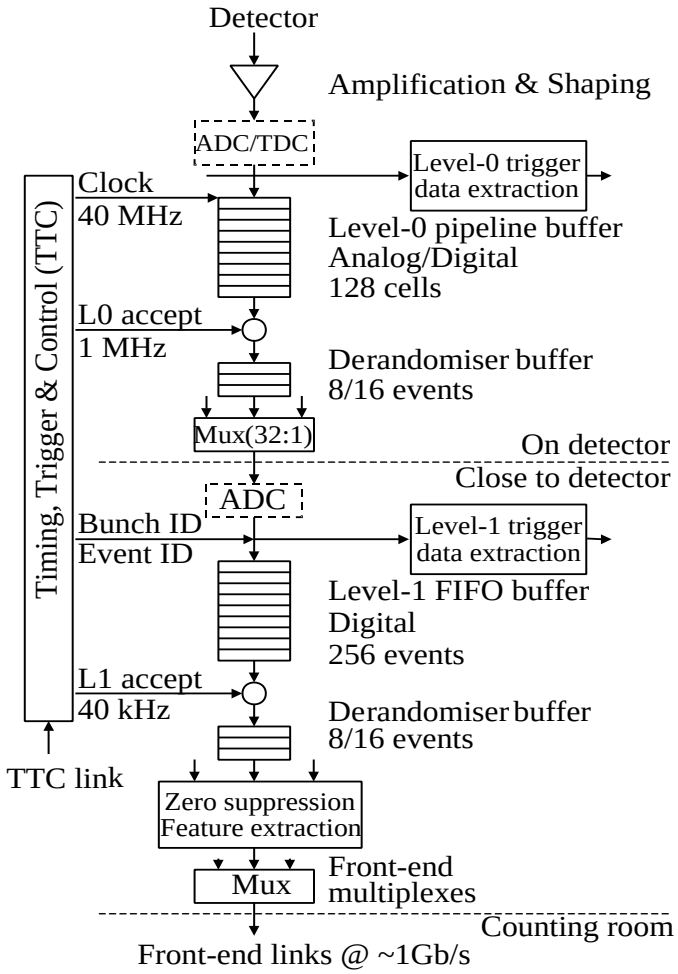

Timing and Fast Control

- Provide common and synchronous clock to all components needing it
- Provide Level-0 and Level-1 trigger decisions
- Provide commands synchronous in all components (Resets)
- Provide Trigger hold-off capabilities in case buffers are getting full

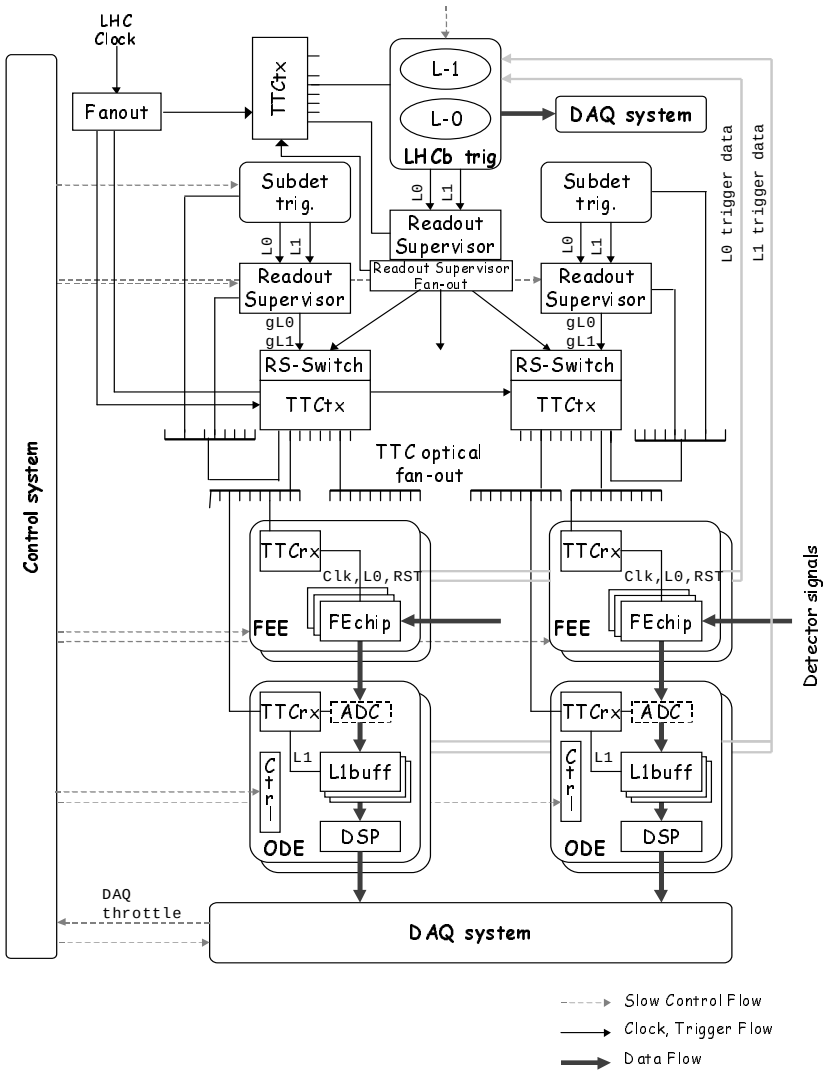

Level-0 Trigger

- Large transverse Energy (Calorimeter) Trigger
- Large transverse momentum Muon Trigger
- Pile-up Veto
- Implemented in FPGAs/DSPs basically hard-wired

Input rate: 40 MHz
Output rate: 1 MHz
Latency: 4.0 μs (fixed)

DAQ Functional Components

- Readout Units (RUs)
  - Multiplex Front-end links onto Readout Network links
  - Merge input fragments to one output fragment
- Subfarm Controllers (SFCs)
  - assemble event fragments arriving from RUs to complete events and send them to one of the CPUs connected
  - Load balancing among the CPUs connected
- Readout Network
  - provide connectivity between RUs and SFCs for event-building
  - provide necessary bandwidth (4 GB/sec sustained)
- CPU farm
  - execute the high level trigger algorithms
    - Level-2 (Input rate: 40 kHz, Output rate: 5 kHz)
    - Level-3 (Input rate: 5 kHz, Output rate: ~100 Hz)
  - ~2000 processors (à 1000 MIPS)

Note: There is no central event manager

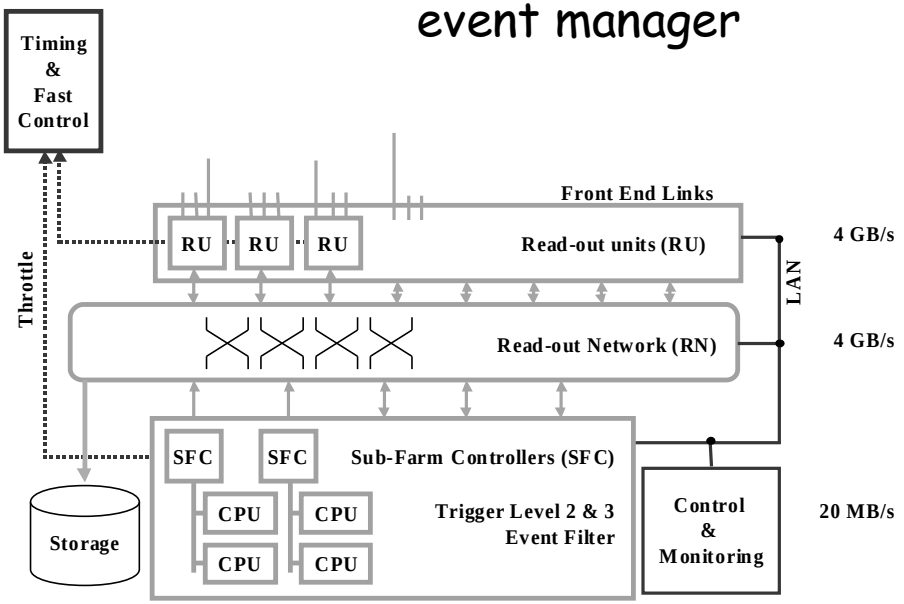

Control System

- Common integrated controls system
  - Detector controls (classical 'slow control')
    - High voltage
    - Low voltage
    - Crates
    - Temperatures
    - Alarm generation and handling
    - etc.
  - DAQ controls
    - Classical RUN control
    - Setup and configuration of all components (FE, Trigger,DAQ)
    - Monitoring
  - **Same system for both functions**

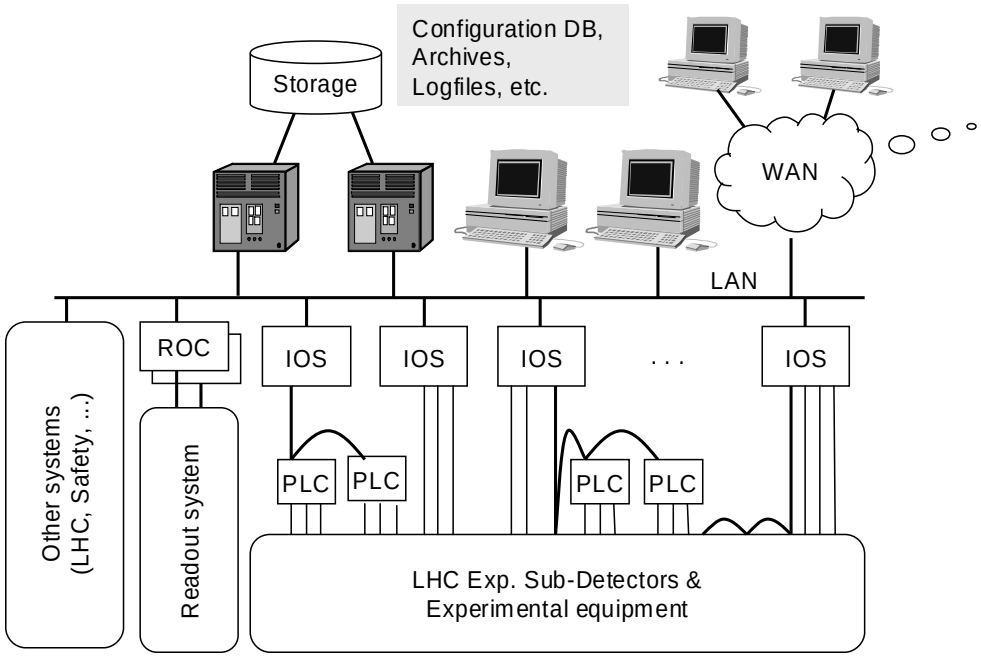

Level-1 Trigger

- Purpose
  - Select events with detached secondary vertices
- Algorithm
  - Based on special geometry of vertex detector (r-stations, φ-stations)
  - Several steps
    - track reconstruction in 2 dimensions (r-z)
    - determination of primary vertex
    - search for tracks with large impact parameter relative to primary vertex
    - full 3 dimensional reconstruction of those tracks
  - Expect rate reduction by factor 25
- Technical Problem: 1 MHz input rate, 3 GB/s data rate, small event fragments, Latency

Level-1 Trigger (2)

- Implementation
  - ~32 sources to switching network
  - Algorithm running in processors (~200 CPUs)
  - Basic idea is to have a switching network between data sources and processors
  - In principle very similar to DAQ, however the input rate of 1 MHz poses special problems.

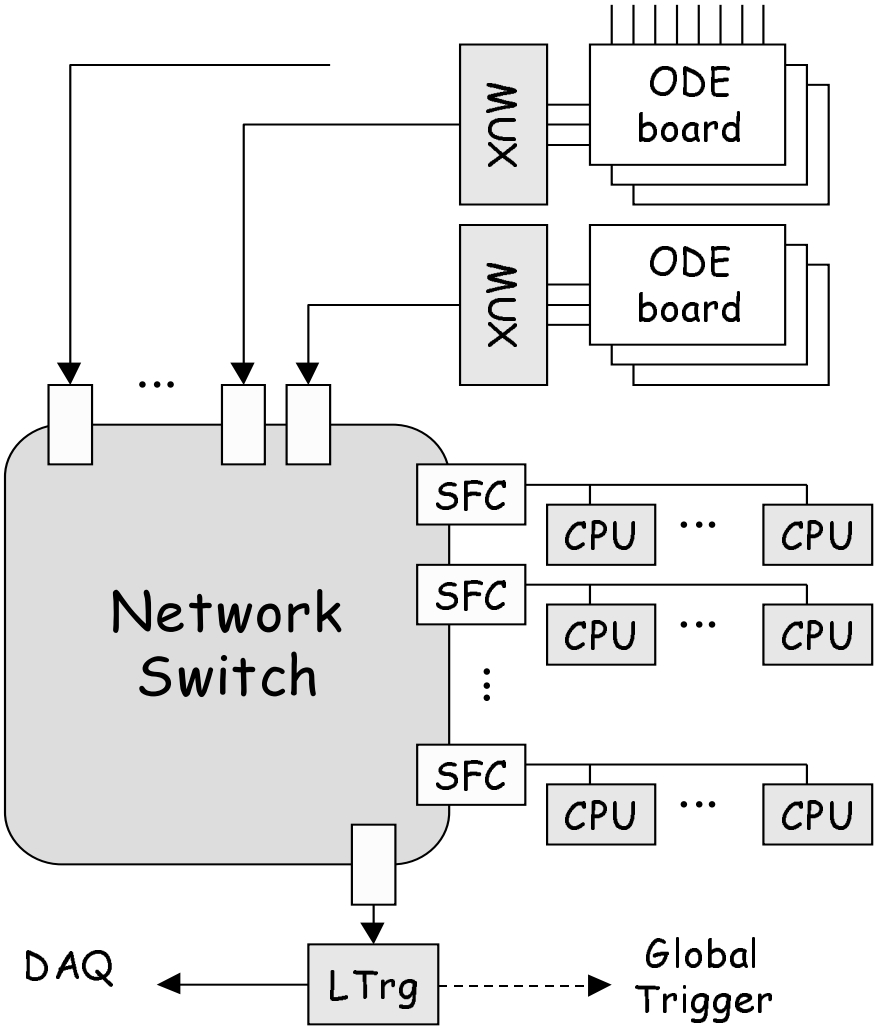

Event-Building Network

- Requirements
  - 4 GB/s sustained bandwidth
  - scalable
  - expandable
  - ~100 inputs (RUs)
  - ~100 outputs (SFCs)
  - affordable and if possible commercial (COTS, Commodity?)
- Readout Protocol
  - Pure push-through protocol of complete events to one CPU of the farm
    - Simple hardware and software
    - No central control → perfect scalability
    - Full flexibility for high-level trigger algorithms
    - Large bandwidth needed
    - Avoiding buffer overflows via 'throttle' to trigger

Event-Building Network Simulation

- Simulated technology: Myrinet
  - Nominal 1.28 Gb/s
  - Xon/Xoff flow control
  - Switches:
    - ideal cross-bar
    - 8x8 maximum size (currently)
    - wormhole routing
    - source routing
    - No buffering inside switches
- Software used: Ptolemy discrete event framework
- Realistic traffic patterns
  - variable event sizes
  - event building traffic

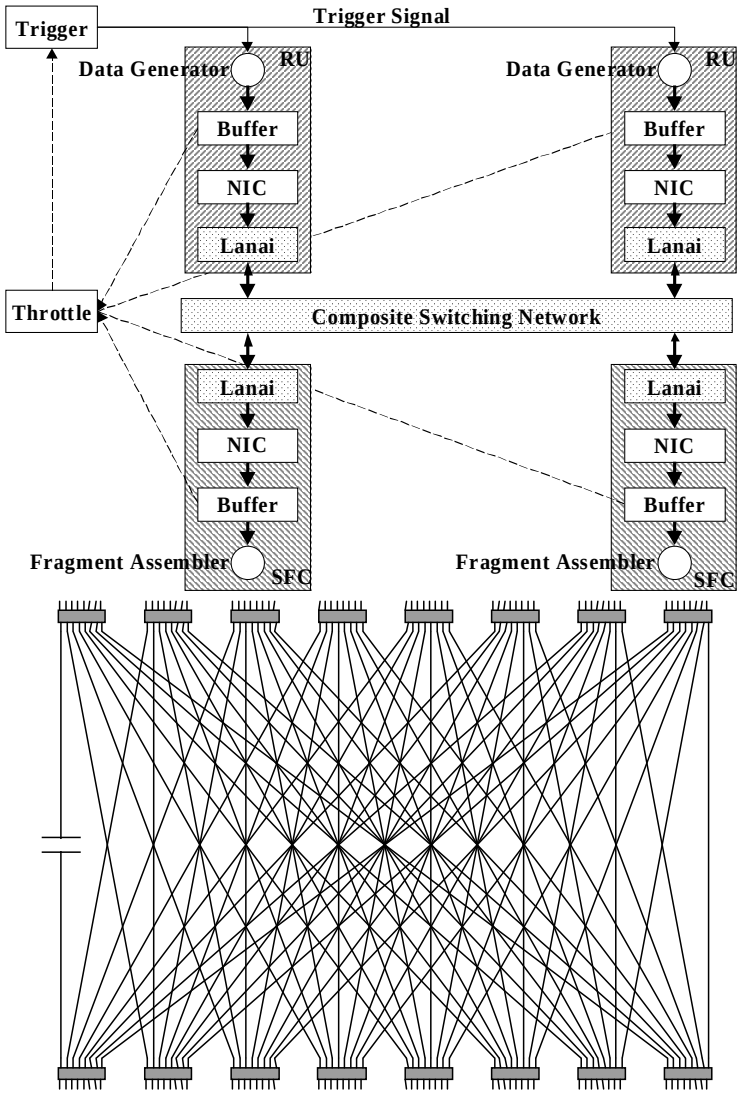

Network Simulation Results

Results don't depend strongly on specific technology (Myrinet), but rather on characteristics (flow control, etc)

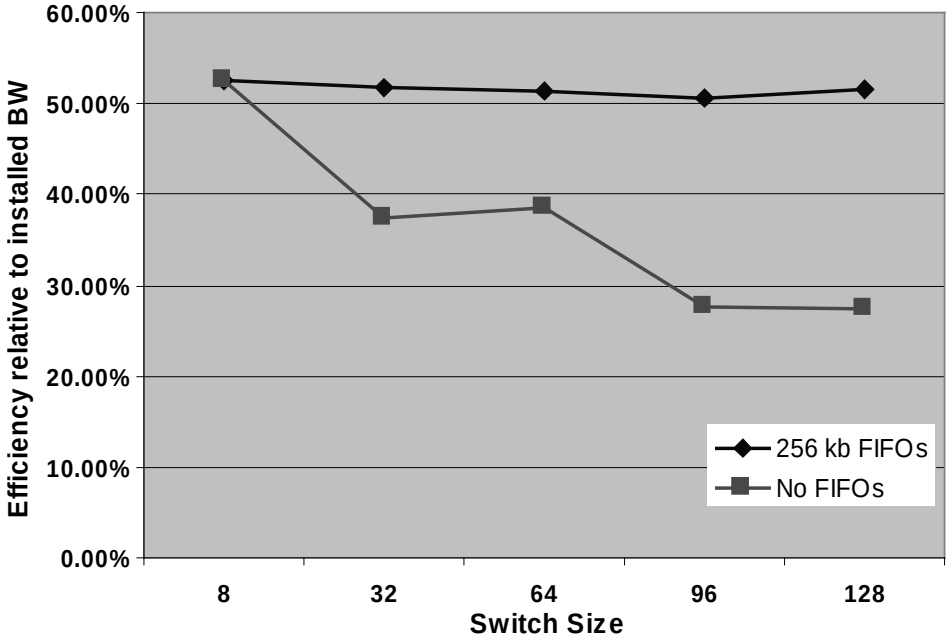

Switch Size	Fifo Size	Switching Levels	Efficiency
8x8	NA	1	52.5%
32x32	0	2	37.3%
32x32	256 kB	2	51.8%
64x64	0	2	38.5%
64x64	256 kB	2	51.4%
96x96	0	3	27.6%
96x96	256 kB	3	50.7%
128x128	0	3	27.5%
128x128	256 kB	3	51.5%

- FIFO buffers between switching levels allow to recover scalability
- 50 % efficiency "Law of nature"

Summary

- LHCb is a special purpose experiment to study CP violation
- Triggering poses special challenges
  - Similarity between inelastic p-p interactions and events with B-Mesons
- DAQ is designed with simplicity and maintainability in mind
  - Push readout protocol $\rightarrow$ Simple, e.g. No central event manager
  - Harder bandwidth requirements on readout network
  - Simulations suggest that readout network can be realized by adding FIFO buffers between levels of switching elements
- Unified approach to Controls
  - Same basic infrastructure for detector controls and DAQ controls
  - Both aspects completely integrated but operationally independent