The ALICE Data-Acquisition Read-out Receiver Card

C. Soós *et al.*
(for the ALICE collaboration)

LECC 2004
13-17 September 2004, Boston

Outline

- Introduction
- Hardware
- Firmware
- Software
- Performance
- Applications
- Summary

Introduction: ALICE DAQ

- ALICE requirements with Pb-Pb beams
  - 25 GB/s aggregate BW from detectors
  - 2.5 GB/s aggregate event building BW
  - 1.25 GB/s aggregate BW to tape
- About 400 Detector Data Link (DDL)
  - Full-duplex optical link
  - Data rate up to 200 MB/s
- The Readout Receiver Card (D-RORC)
  - PCI compatible adapter
  - Integrates two DDL interfaces
- ALICE data-acquisition software: DATE

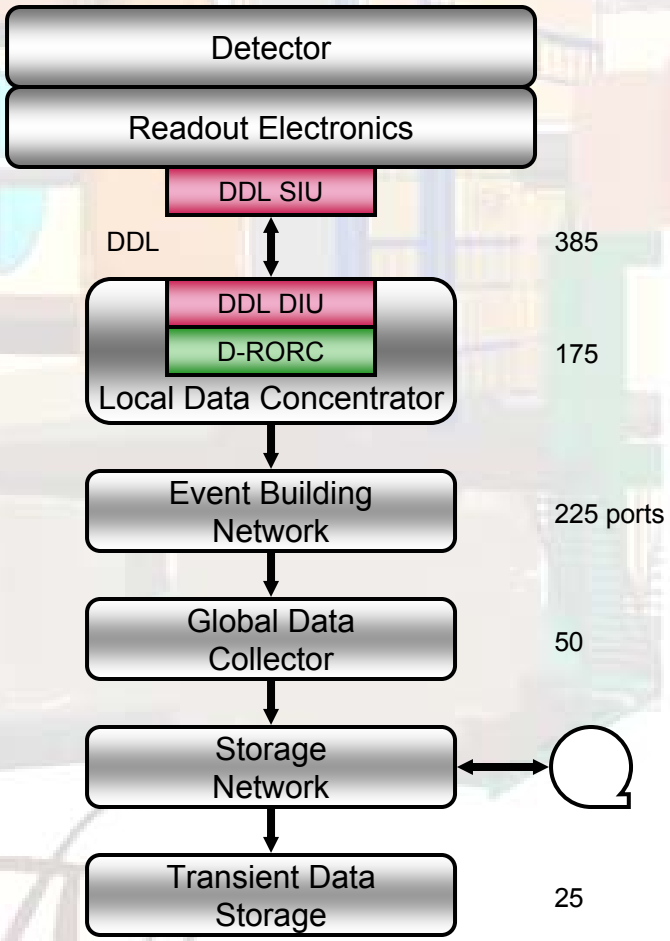

Introduction: DAQ and HLT

- Some of the detectors require filtering or data reduction
- DAQ and HLT interface
  - Standard DDL links
  - Data splitter in the D-RORC
- Each D-RORC hosts two DIUs
  - First DIU receives detector data
  - Second DIU transfers the copy of the raw data
- HLT data is transferred back using DDL and D-RORC

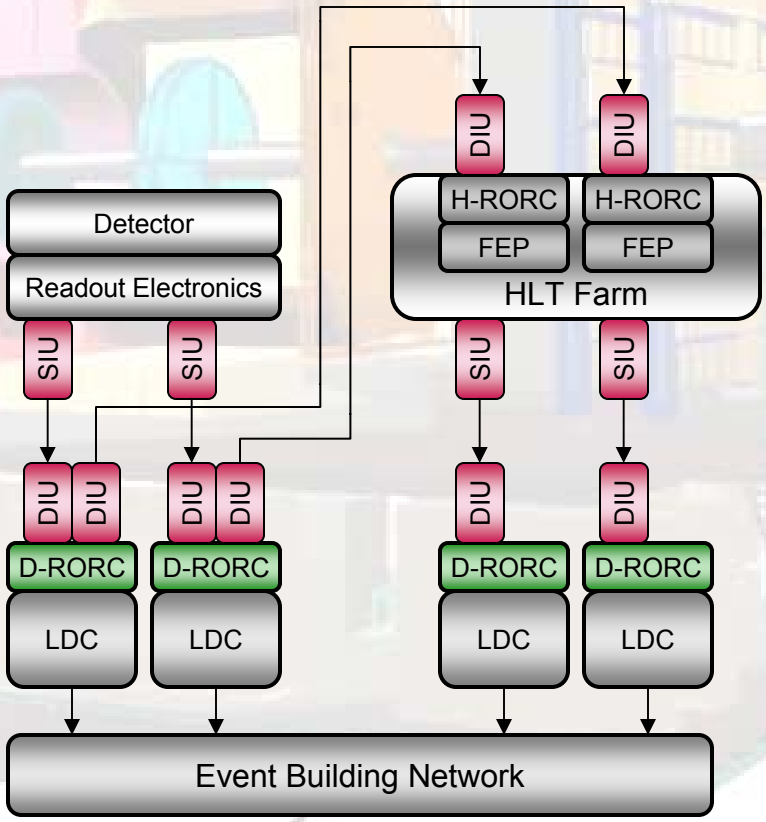

Hardware

JTAG interface
- FPGA programming and debugging
- Flash memory programming

Conf. Flash
- EPC4
- Programmable via internal or external JTAG chain

Electrical transceivers
- Multi-rate transceivers
- Serial-Parallel-Serial converters
- Integrated 8B/10B endec
- Clock recovery

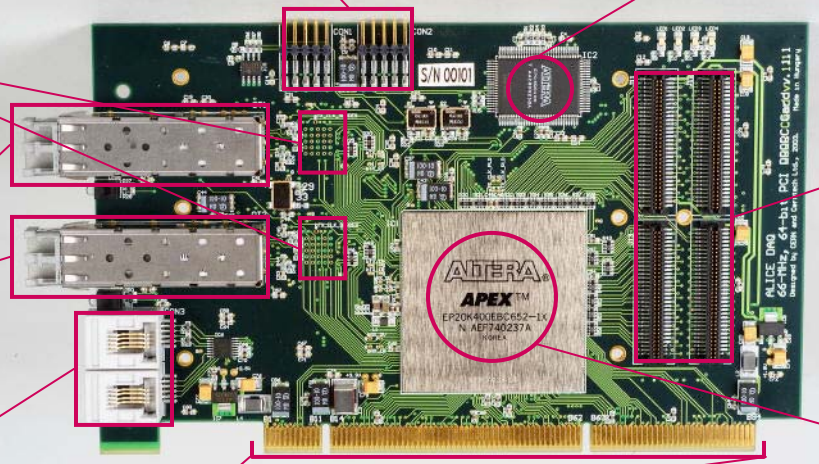

CMC interface
- Standard extension I/F
- About 180 user I/O

Optical transceivers
- 850 nm VCSEL laser
- 2.125 Gbit/s
- Pluggable modules

Altera FPGA
- APEX-E device family
- EP20K400E

LVDS interface
- High-speed serial I/F
- 2 inputs + 2 outputs
- Purpose: detector busy

PCI 64-bit/66 MHz
- +3.3V compatible signals
- Bus master enabled

Firmware

- PCI interface core
  - Handles PCI transactions
  - Performs PCI mastering
- Receiver and Transmitter
  - Buffer data and initiate the DMA
- Control registers
  - Mapped to PCI memory
  - Provide control interface
  - Provide status information
- DDL interface
  - Performs DDL transactions
  - Provides DDL status

High-speed I/F to the transceivers

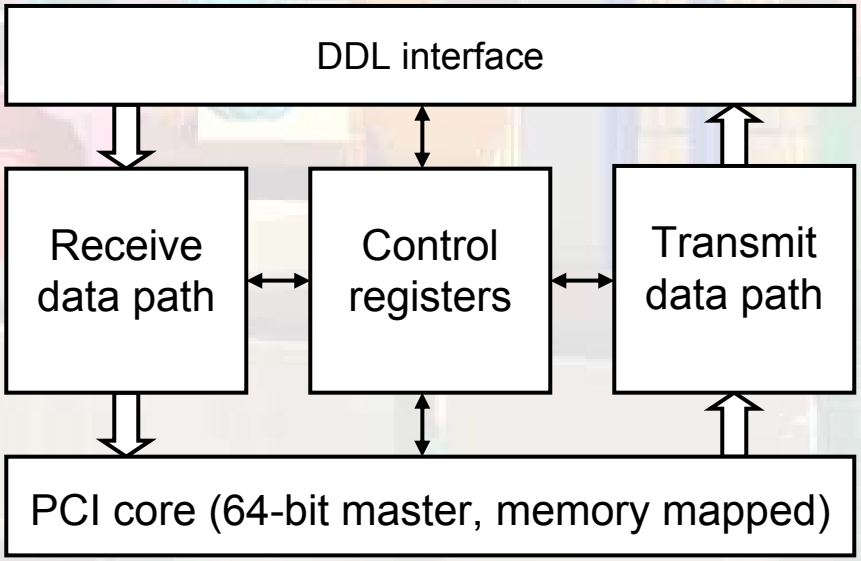

64-bit PCI or PCI-X bus, 3.3V signaling (!)

Receiver DMA

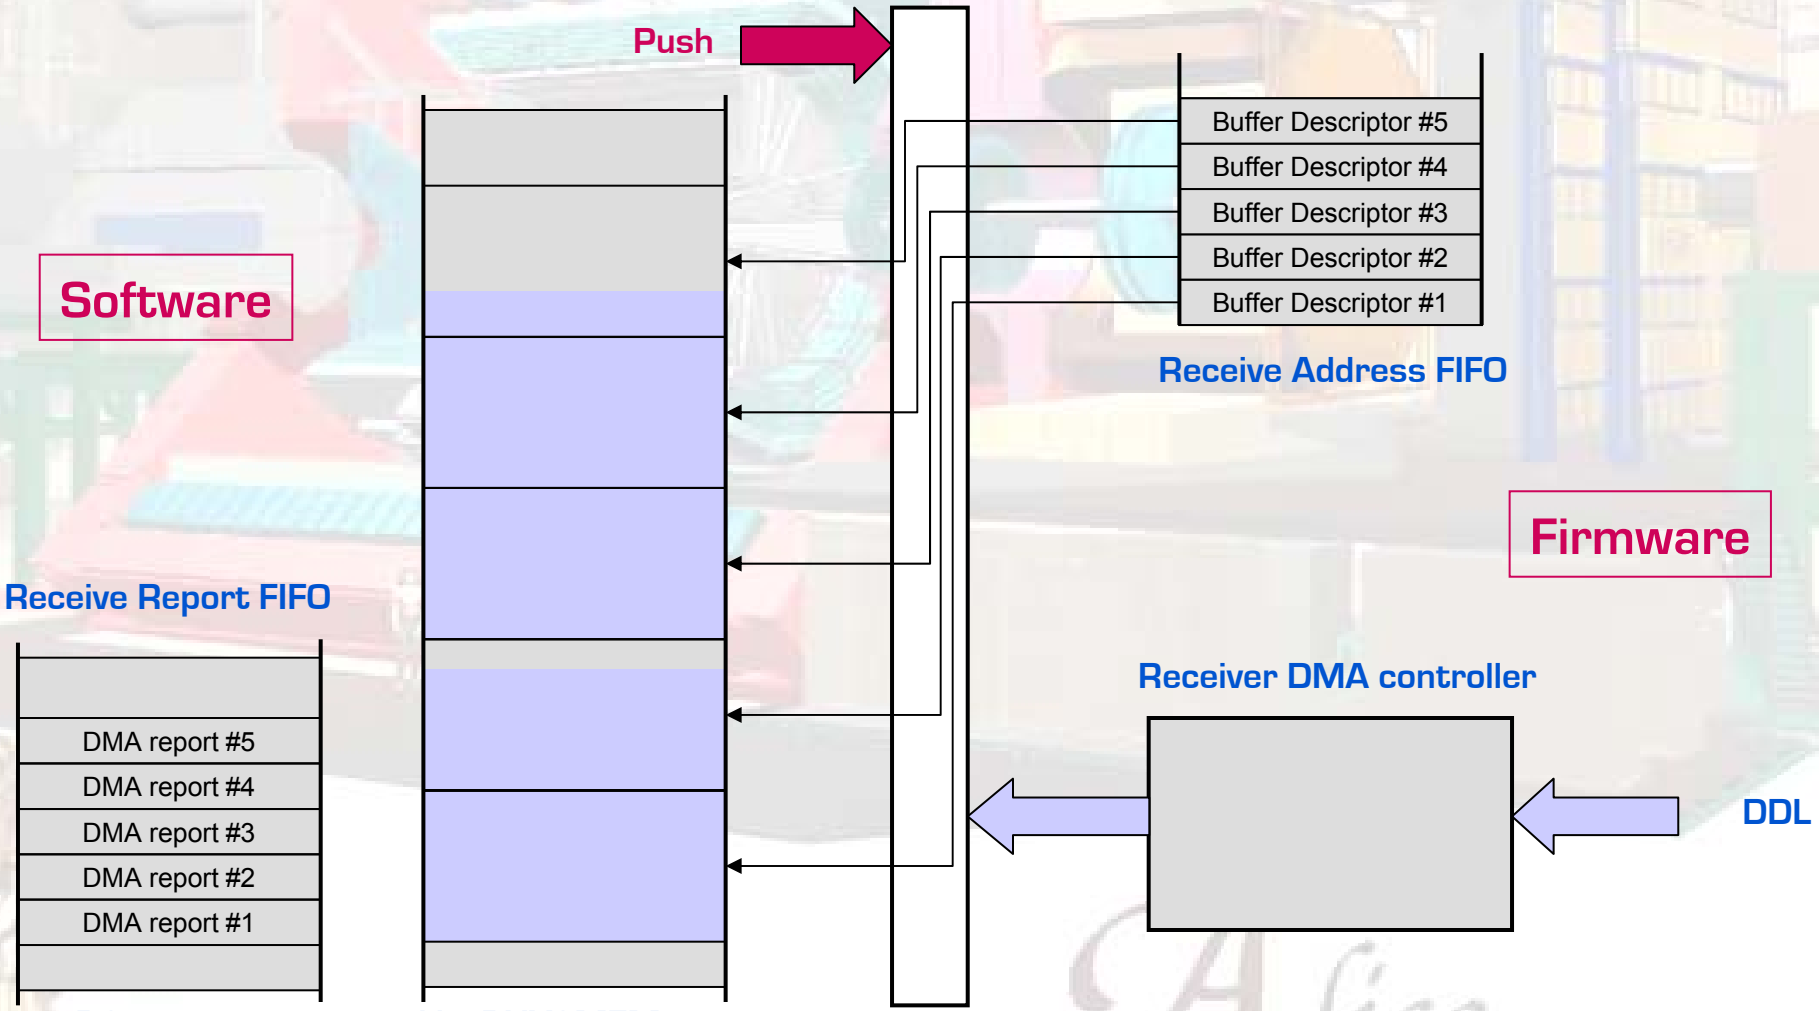

Software

- D-RORC driver: Linux device driver, runtime loadable module
  - Finds the D-RORC cards on the PCI buses
  - Maps the registers into the user memory space
- D-RORC API layer: collection of library routines written in C
  - Ensures exclusive access to the hardware using device locking
  - Provides simple programming I/F for higher level applications
- Command line executables
  - Hardware identification
  - Reset components (DIU, SIU etc.)
  - Send commands, reads status
  - Send data blocks
  - Receives and checks data blocks

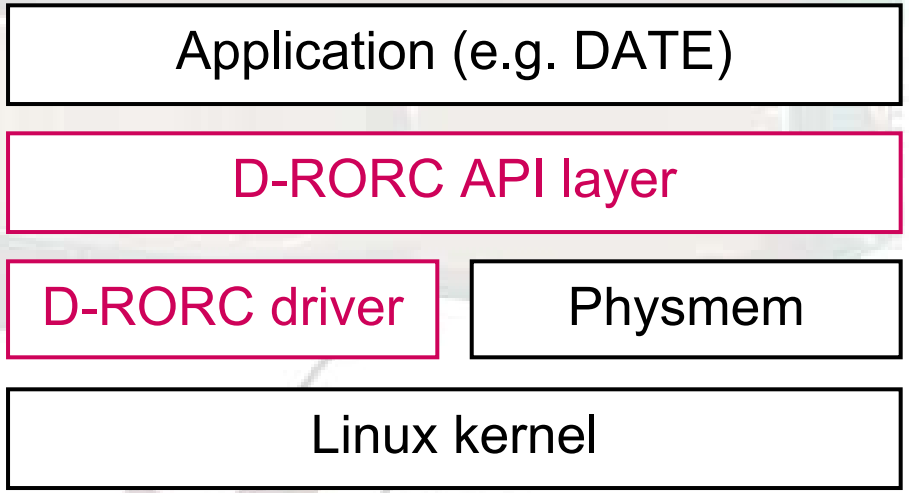

Performance: Test bed

- Supermicro server motherboard with dual Xeon CPUs @ 2.4 GHz
- Six PCI-X slots, 4 bus segments (3+1+1+1)
- Linux OS
- ALICE Data-Acquisition software (DATE)

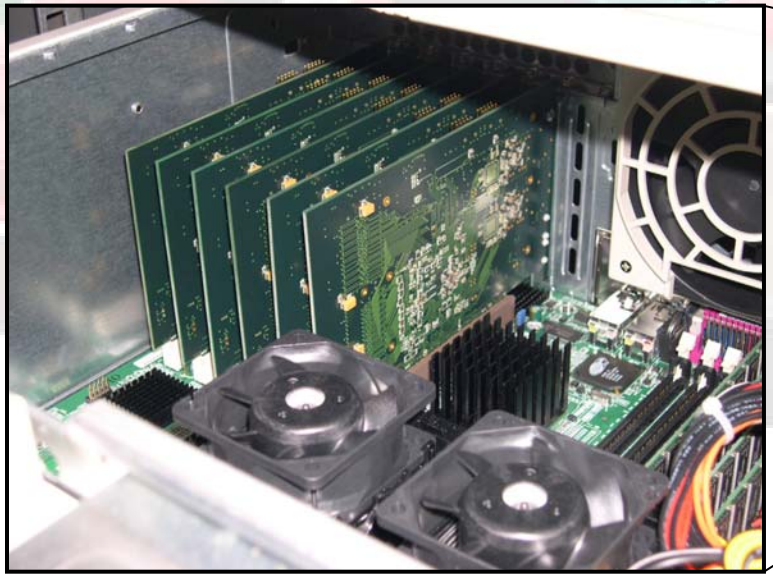

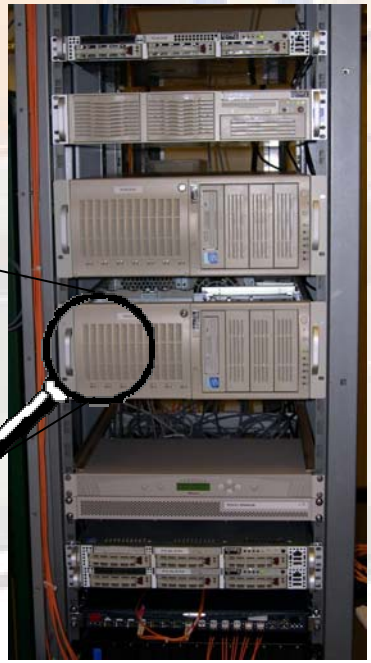

Performance: Single channel

- Bandwidth vs. block size measurements with internal and external (DDL) data source using one D-RORC channel
- Steady increase until the maximum bandwidth is reached
  - Internal: BWmax = Fpci [MHz] x 4 [Bytes] = 264 MB/s
  - External: BWmax = BWddl = 206 MB/s

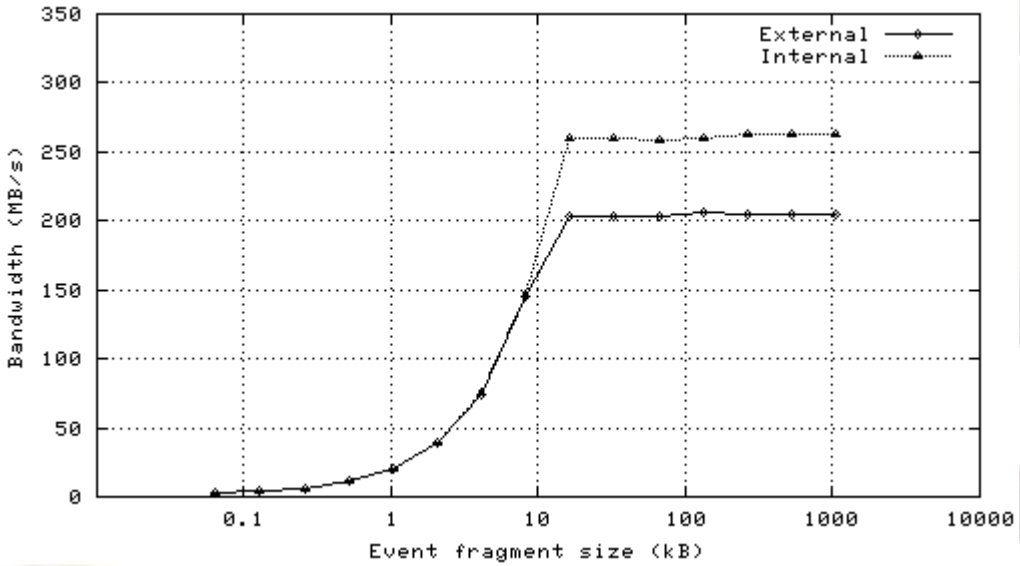

Performance: Dual channel

- Bandwidth vs. block size measurements with internal and external (DDL) data source using two D-RORC channel
- Steady increase until the maximum bandwidth is reached
  - Internal: BWmax = Fpci [MHz] x 4 [Bytes] x 2 – Loss = 484 MB/s
  - External: BWmax = BWddl x 2 = 412 MB/s

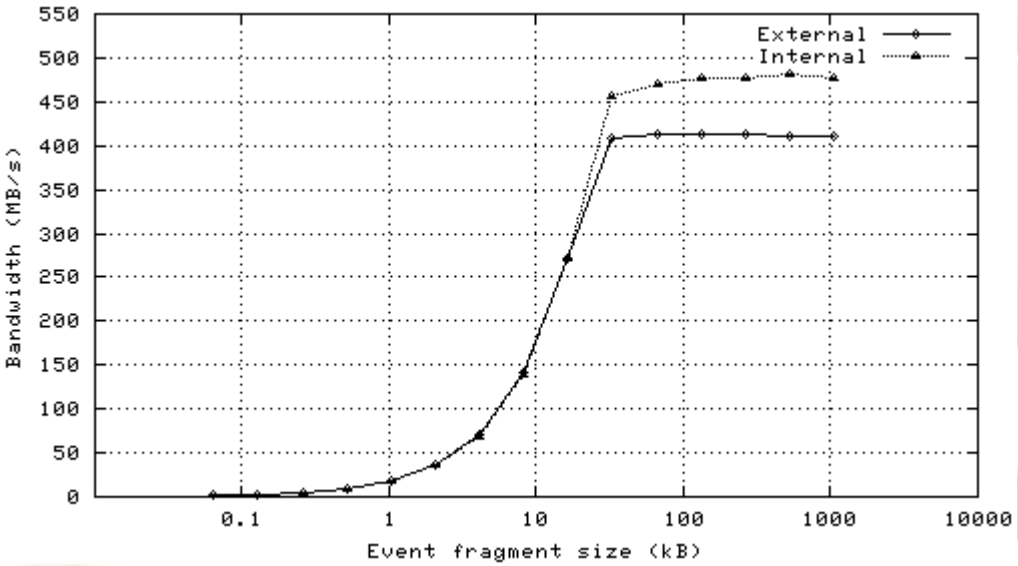

Performance: Dual vs. 2 Single

- Bandwidth vs. block size measurement with internal data source
- Different maximum (different arbitration)
  - Dual-channel D-RORC: BWmax = 484 MB/s
  - Two single-channel D-RORC: BWmax = 464 MB/s

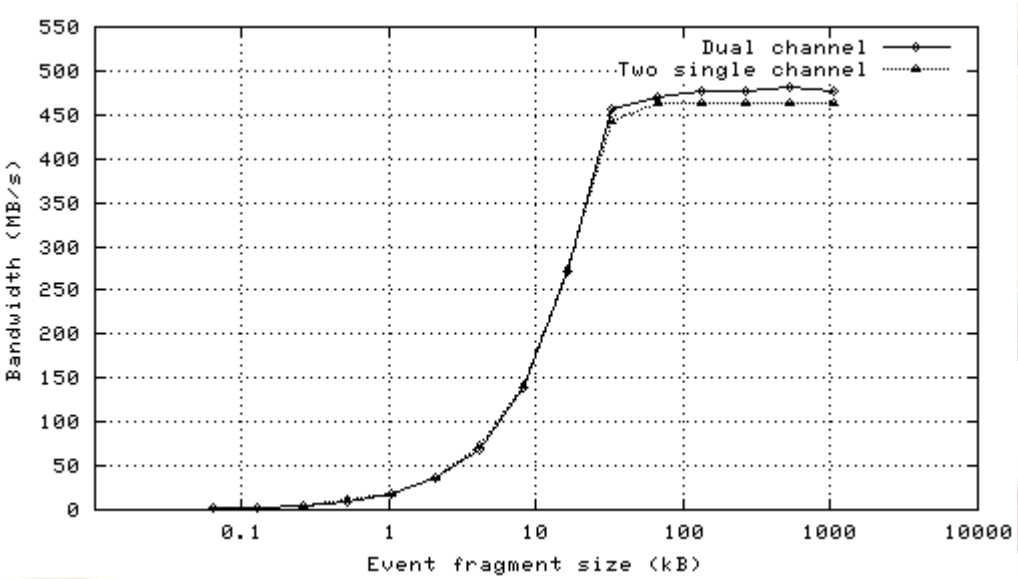

Performance: Six D-RORCs

- Testing the fully populated PC using internal data source
  - Interoperability test
  - Measure the maximal input bandwidth

#2	#4	PCI #6				1 Ch	1 Ch
	#3	PCI #5			1 Ch	1 Ch	1 Ch
Controller #1	#2	PCI #4		1 Ch	1 Ch	1 Ch	1 Ch
	Segment #1	PCI #3	1 Ch				1 Ch
		PCI #2	1 Ch	1 Ch			1 Ch
		PCI #1	1 Ch	1 Ch	1 Ch	1Ch	1 Ch	1 Ch
Bandwidth [MB/s]			264	464	424	528	792	1045	840
Normalized Bandwidth [MB/s/Ch]			264	232	141.3	264	264	261.3	140

Applications: TPC sector test

- Test beam of one complete Inner Read-out Chamber (IROC) of the ALICE TPC detector (May 2004)

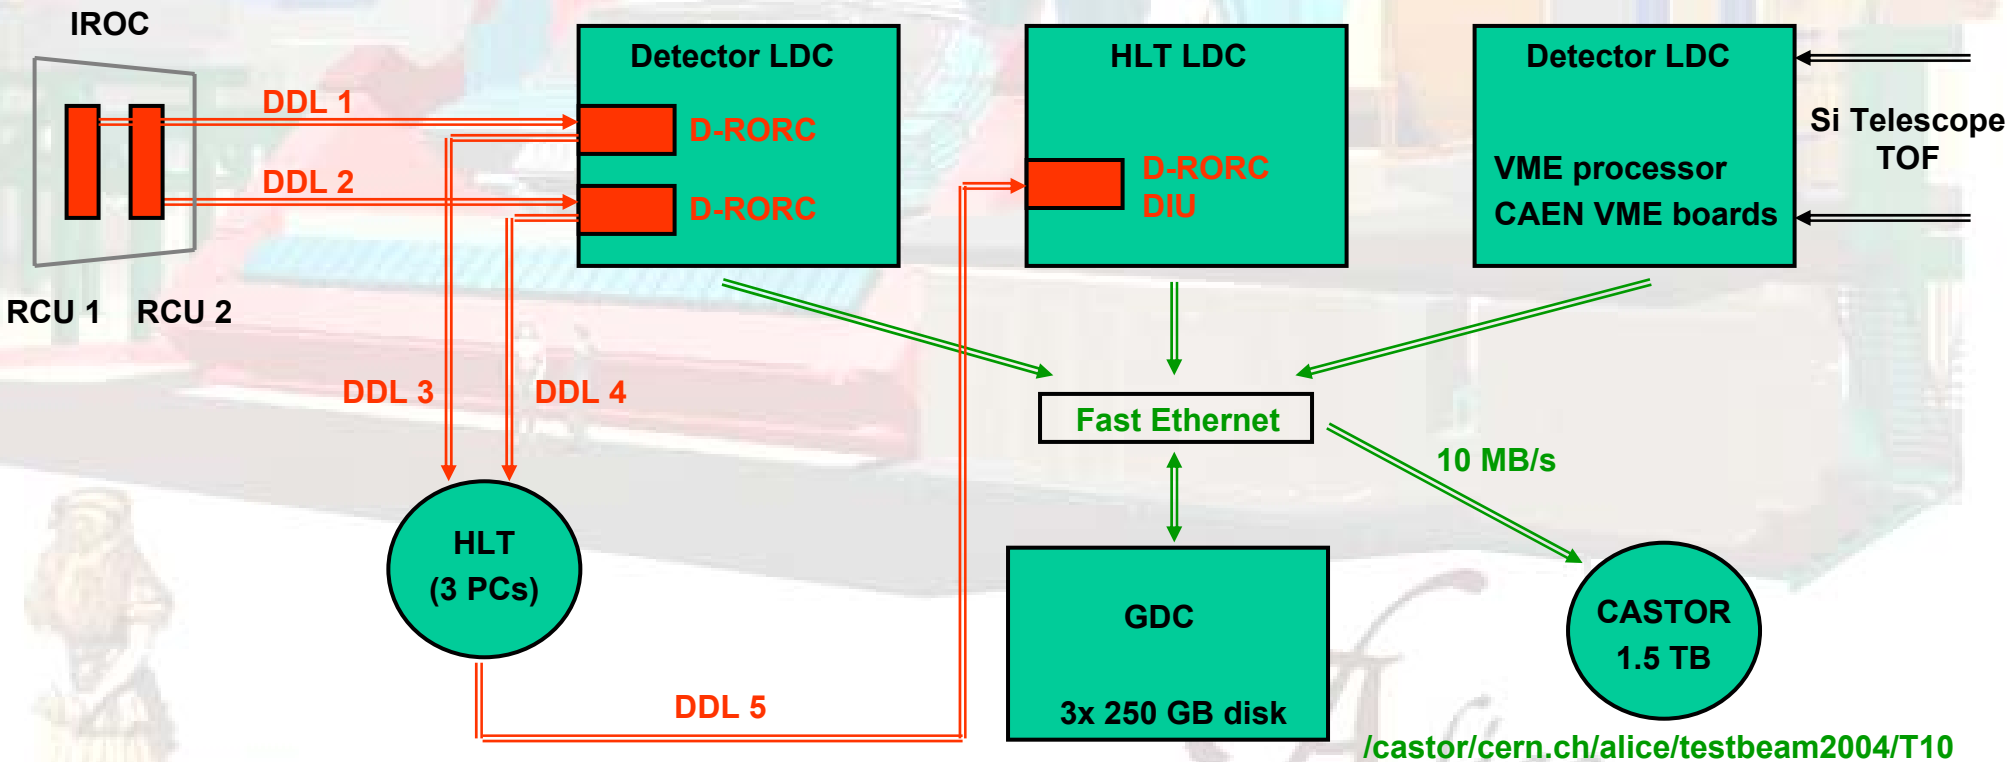

Summary

- D-RORC card has been developed as the high-speed interface between the DDL and the PCI bus
- Using two integrated DDL channels
  - reduces the number of PCI slots
  - offers data paths from the detectors to the DAQ and HLT systems
- Linux device driver and API library based on standard C, as well as Linux executables are available
- The card has been tested thoroughly in the lab
  - 1 CH bandwidth = 264 MB/s
  - 2 CH bandwidth = 484 MB/s
  - 4 D-RORC on different PCI segments = 1045 MB/s
- Real applications show the stability and reliability of the card


Thank you!
Alice