

Examples

ASIC and FPGA digital designs in physics experiments

Examples

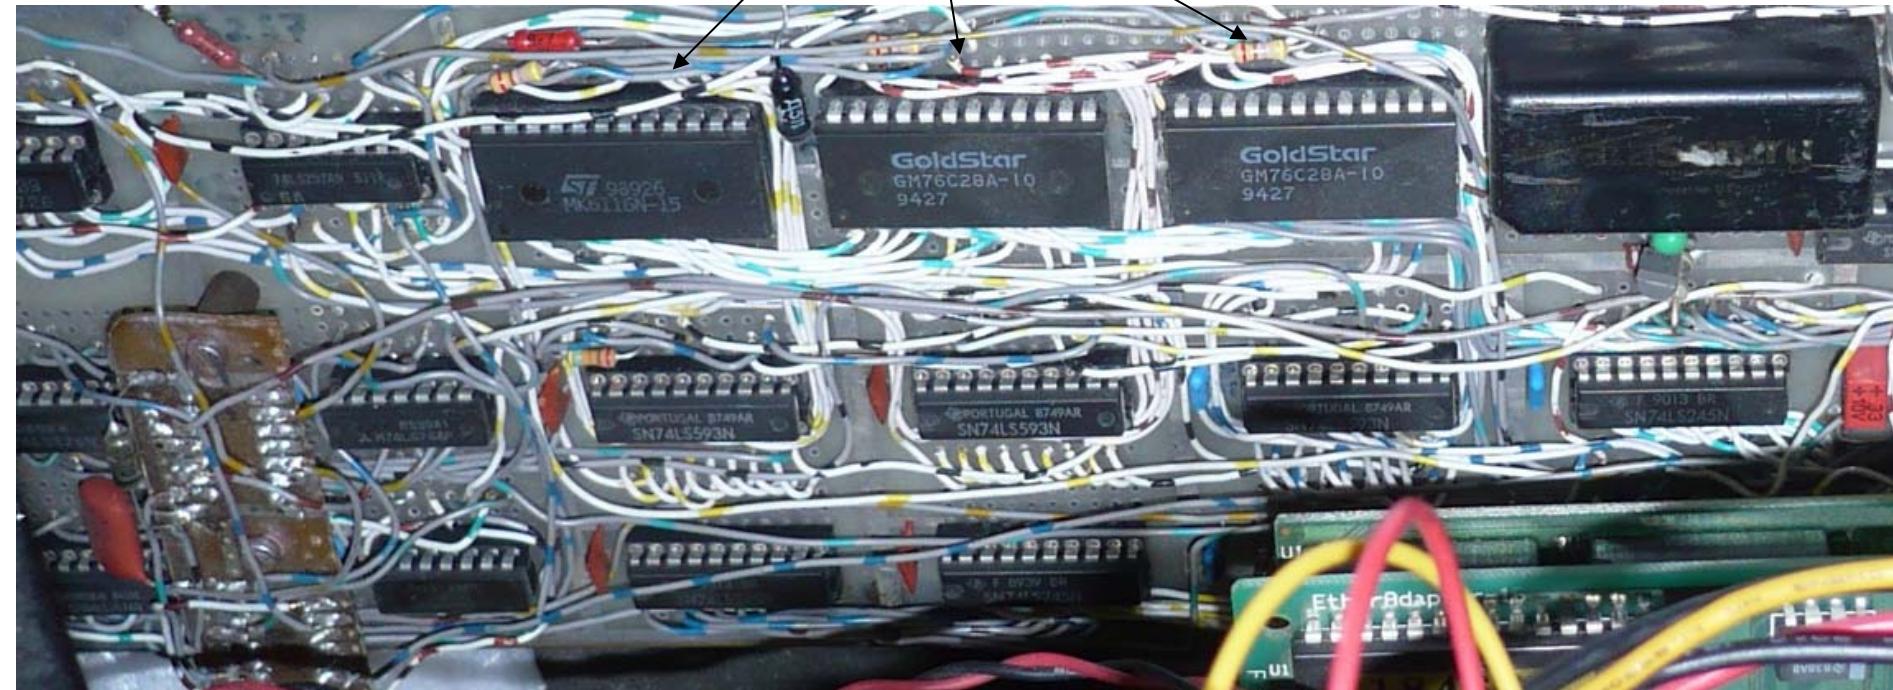
- Prototypes in the old times
- Sweet-16 – a student RISC processor
- TRAP chip for ALICE TRD – mixed mode ASIC
- Optical Readout Interface (ORI) – CPLD
- Detector Control System Board for ALICE (TRD+TPC) – FPGA + ARM CPU core
- Power Distribution Box (PDB) – antifuse FPGA
- Global Tracking Unit for ALICE TRD – large FPGA farm

Prototypes yesterday...

Interface board for PHA (pulse high analysis) with 74'xxx

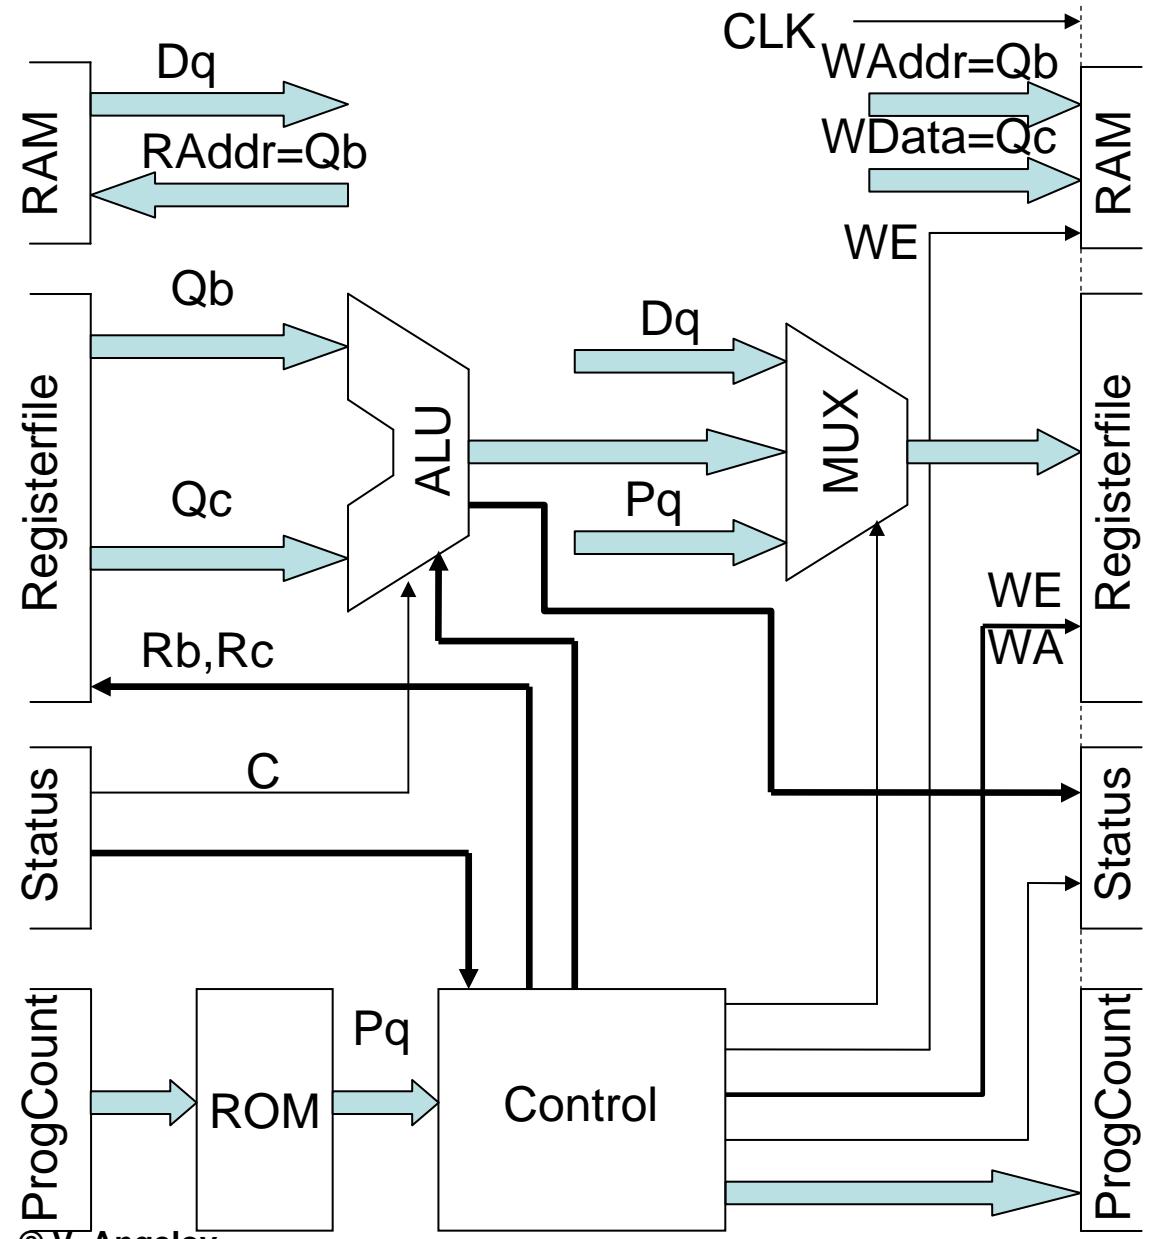
3 x SRAM 2k x 8

Battery for the SRAM



Full size ISA card for IBM/XT/AT

Sweet-16

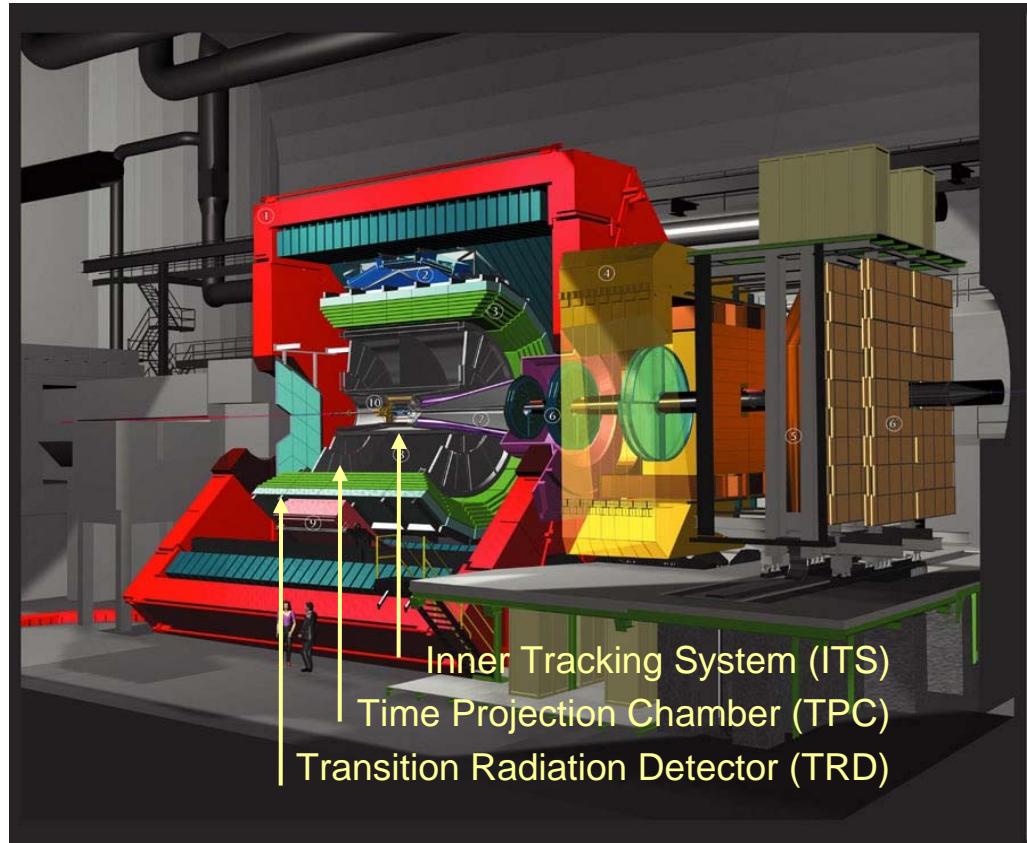


- 16 bit RISC processor
- 1 clock/instruction
- easy to implement by students without experience
- compact and portable to different technologies, used in FPGAs and ASICs

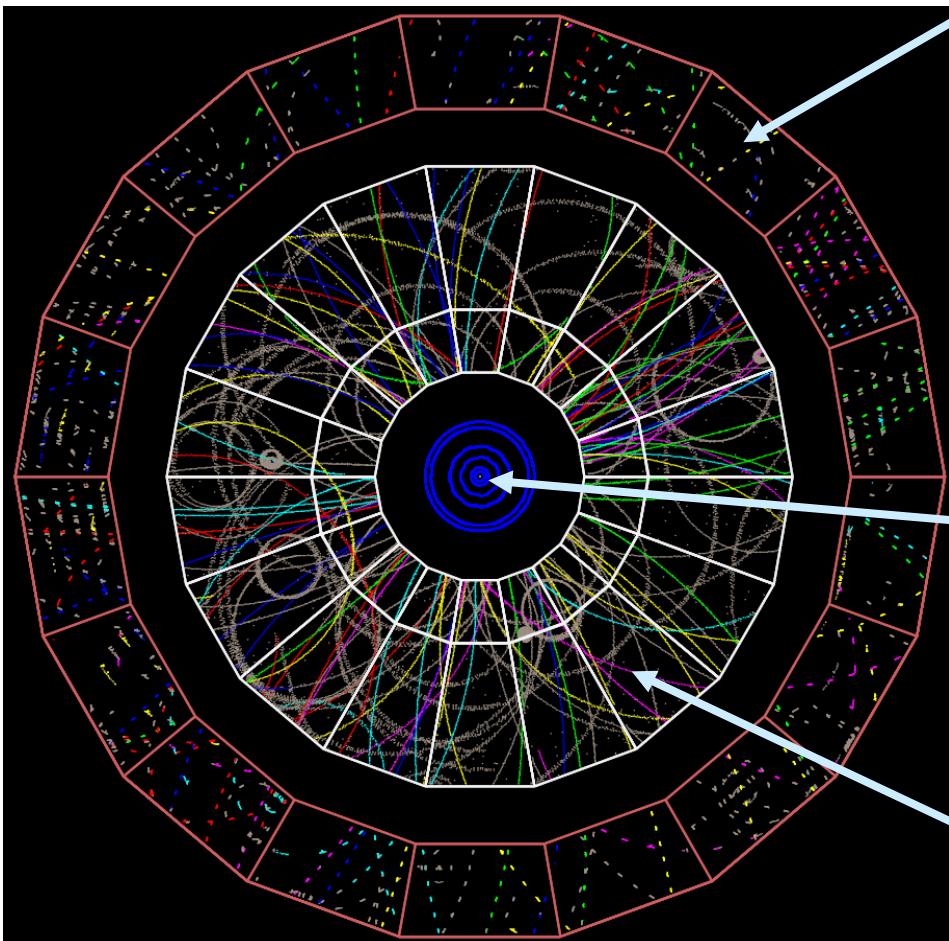


A Large Ion Collider Experiment

- Pb-Pb Collision at 1.1 PeV/Nukleon
- Creation of Quark Gluon Plasma
- TRD is used as a trigger detector due to its fast readout time ($2 \mu\text{s}$):
 - Transversal Momentum
 - Electron/Pion Separation



Transition Radiation Detector



TRD - Transition Radiation Detector

- used as trigger and tracking detector
- > 24000 particles / interaction in acceptance of detector
- up to 8000 charged particles within the TRD
- trigger task is to find specific particle pairs within 6 μs .

ITS - Inner Tracking System

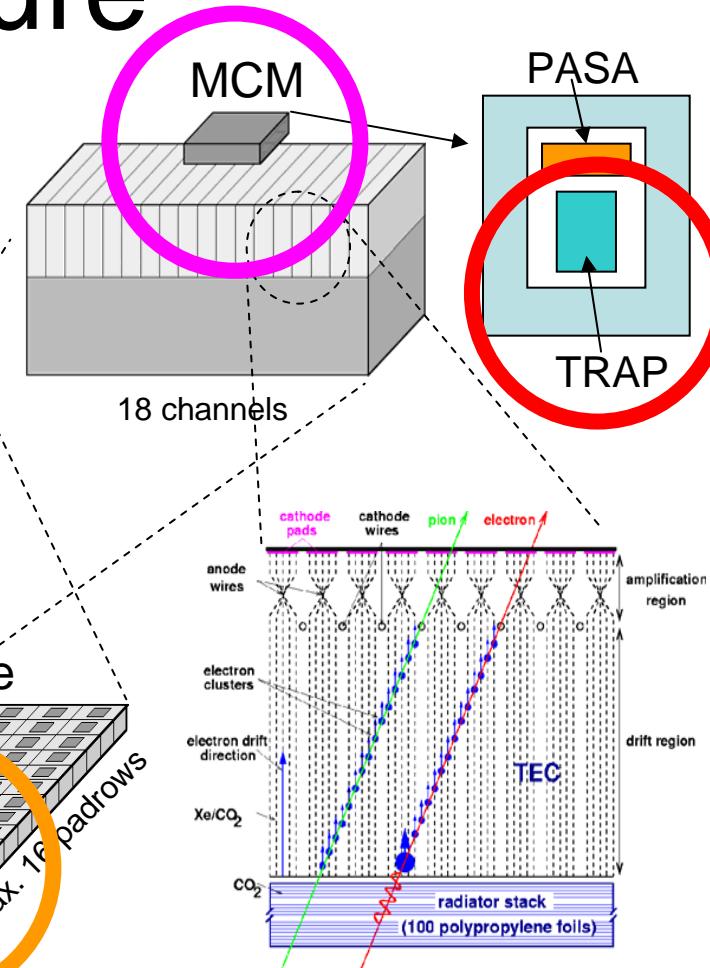
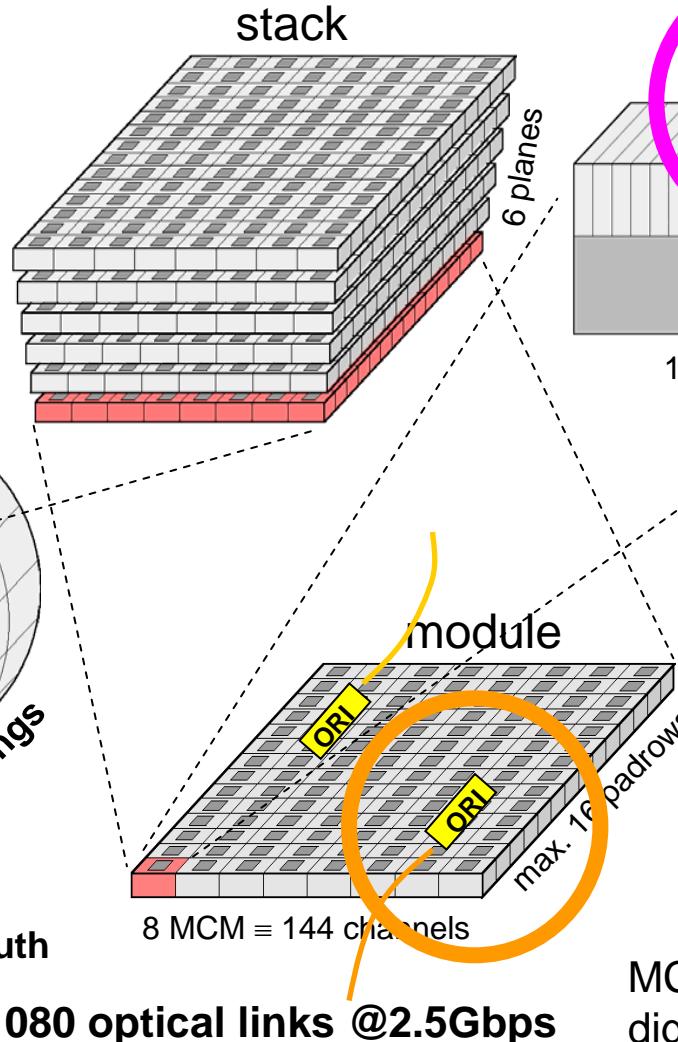
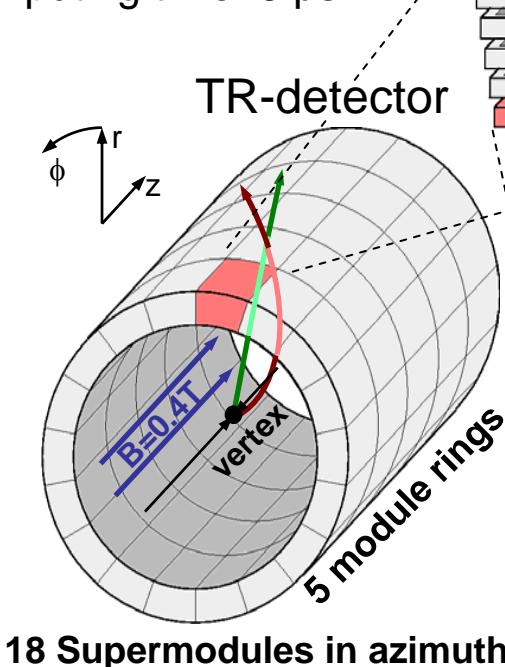
- event trigger
- vertex detection

TPC - Time Projection Chamber

- high resolution tracking detector
- but too slow for 8000 collisions / second

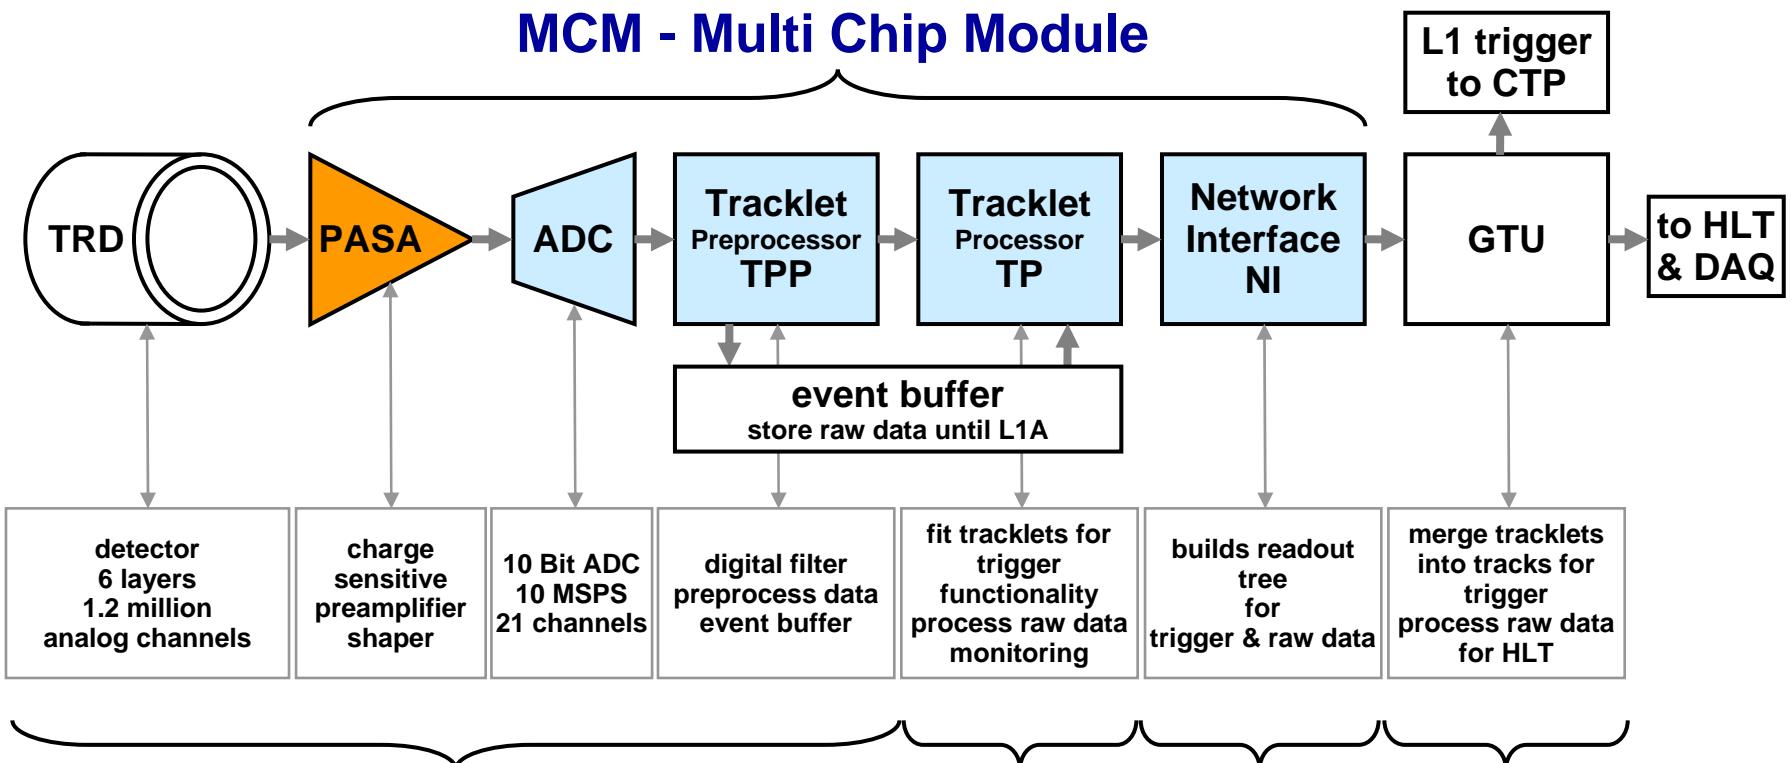
TRD Structure

- 1.2 million channels
- 1.4 million ADCs
- peak data rate: 16 TB/s
- ~65000 MCMs
- computing time: 6 μ s



MCM performs amplification, digitization (10bit 10MHz), straight line fit, readout network

Partitioning, Data Flow & Reduction



time:	during first 2 μ s (drift time)	after 3.5 μ s	after 4.1 μ s	after 6 μ s
data / event:	33 MB	max. 80 KB	600 GB/s	-
peak rate:	16 TB/s	-	-	-
mean rate:	257 GB/s	-	-	-
reduction:	1	~ 400	to trigger decision	

FEE development

How to design the FEE?

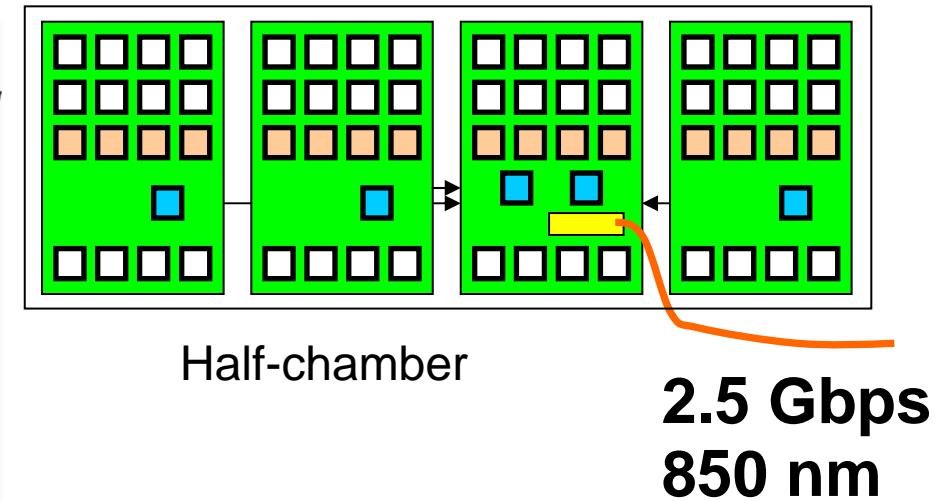
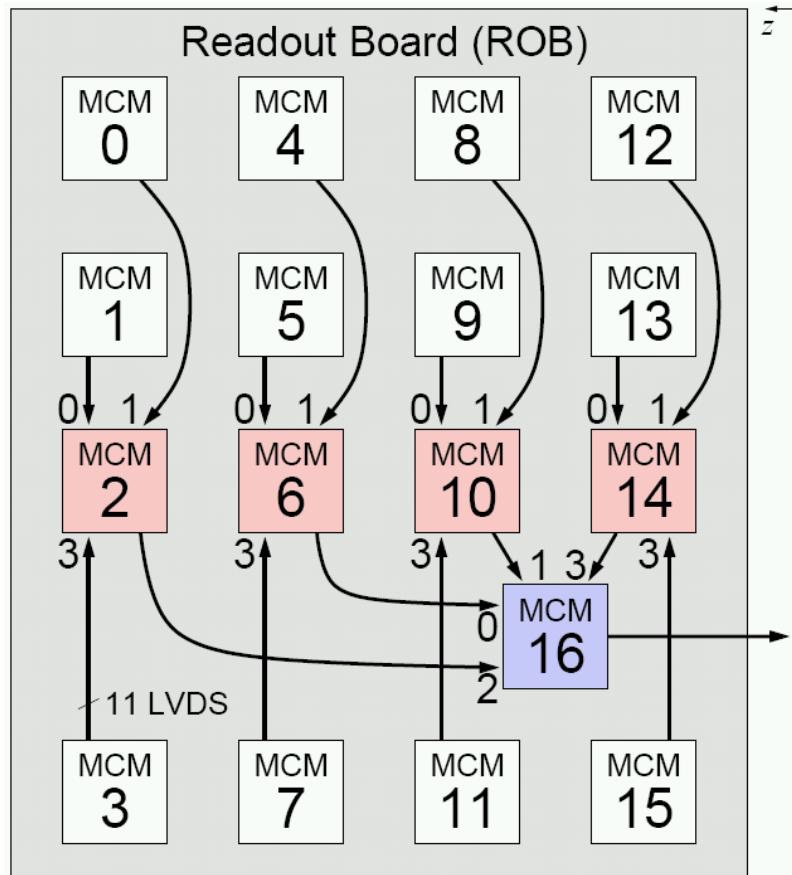
- fast and low latency
- low power - precise power control ($1\text{mW}/\text{channel} \rightarrow 1\text{kW}$)
- low cost
 - avoid using connectors
 - use some simple chip package (MCM + ball grid array)
 - standard components? which process? IP cores? Layout (TPC vs. TRD)?
- flexible, as much as possible, as the exact processing not known
 - make everything configurable
 - use CPU core for final processing
- reliable, no possibility to repair anything later
 - redundancy, error and failure protection
- self diagnostic features

FEE development(2)

Chip Design flow

1. Detector simulations to understand the signals and what kind of processing we need
2. Select PASA shaping time, ADC sampling rate and resolution
3. Behavior model of the digital processing including the bit-precision in every arithmetic operation
4. Estimate the processing time and select the clock speed of the design (multiple of LHC and ADC sampling clock)
5. Code the digital design, simulate it, synthesize it, estimate the timing and area, optimize again...
6. Submit the chip, this is the point of no return!
7. Continue with the simulations, find some bugs and think about fixes
8. Prepare the test setup
And so on, TRAP1, 2, TRAPADC, TRAP3, TRAP3a (final)

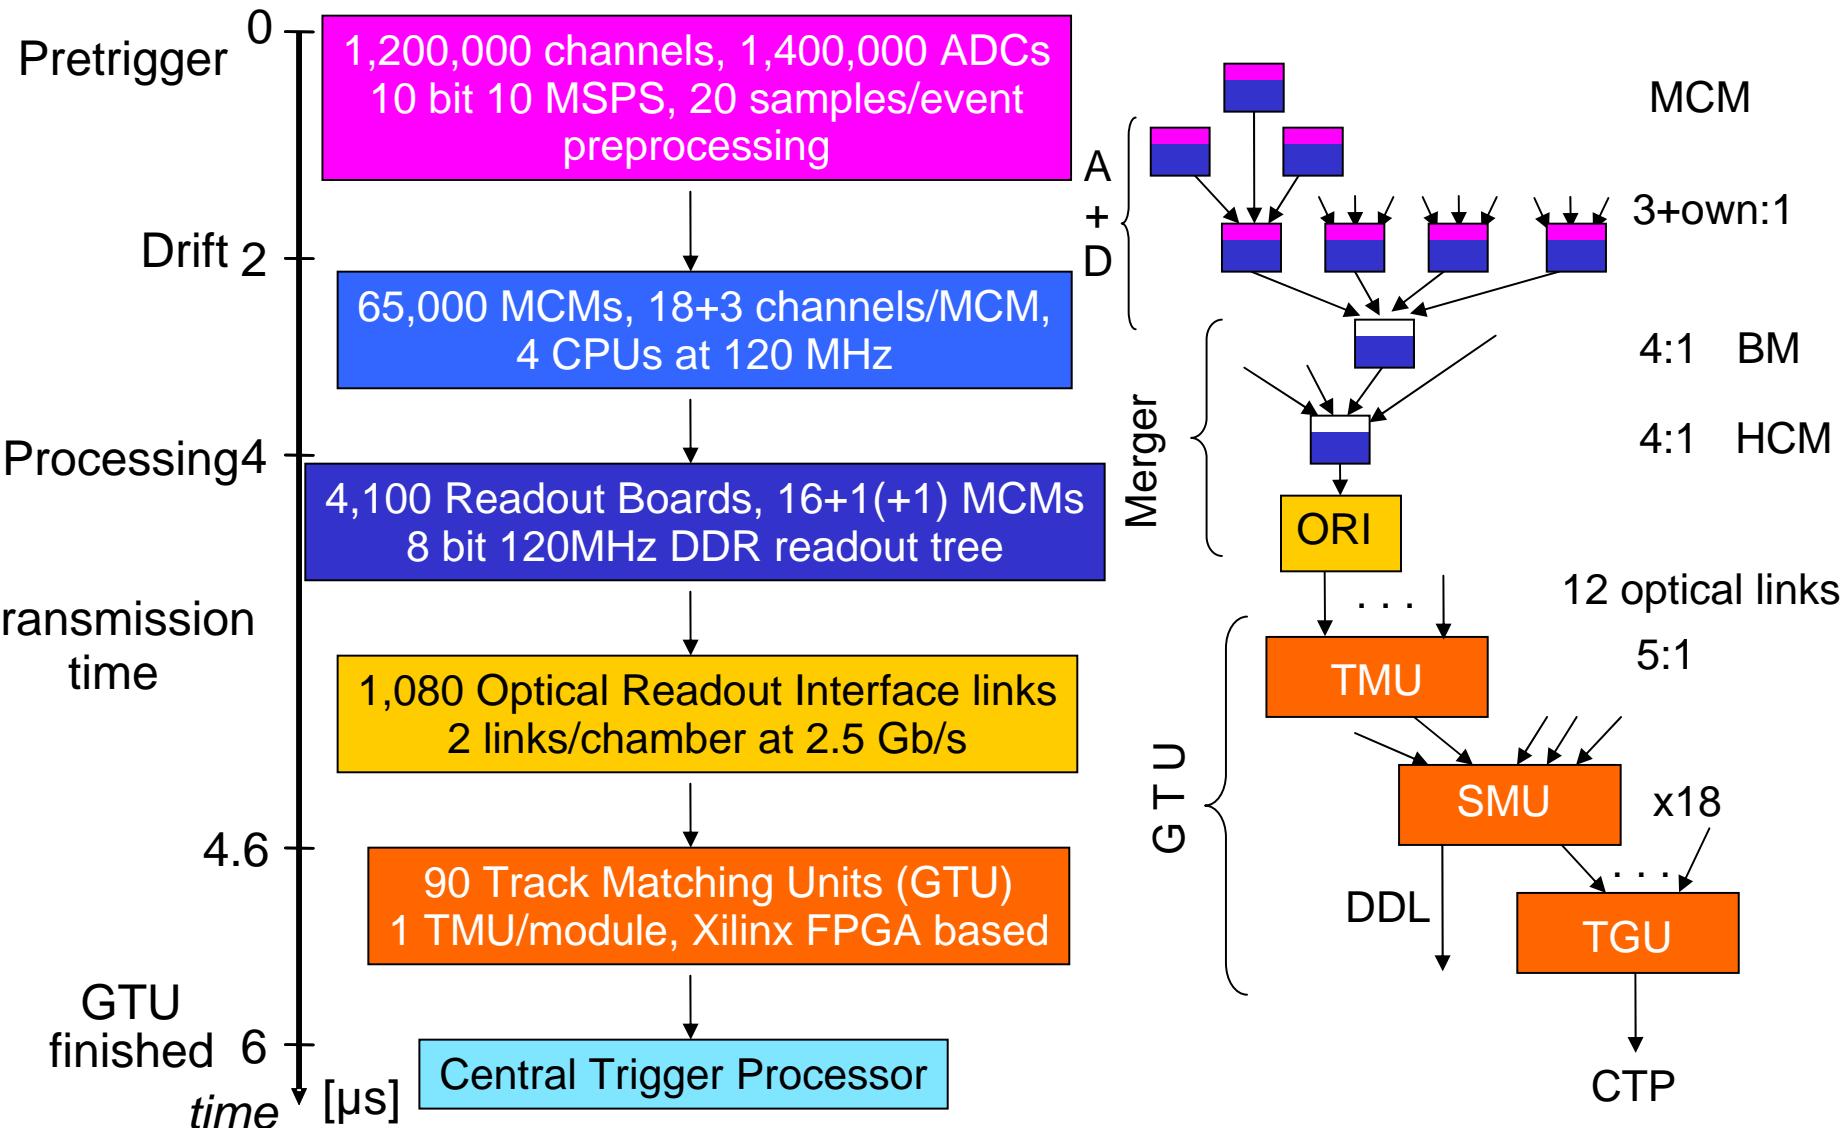
Readout Boards



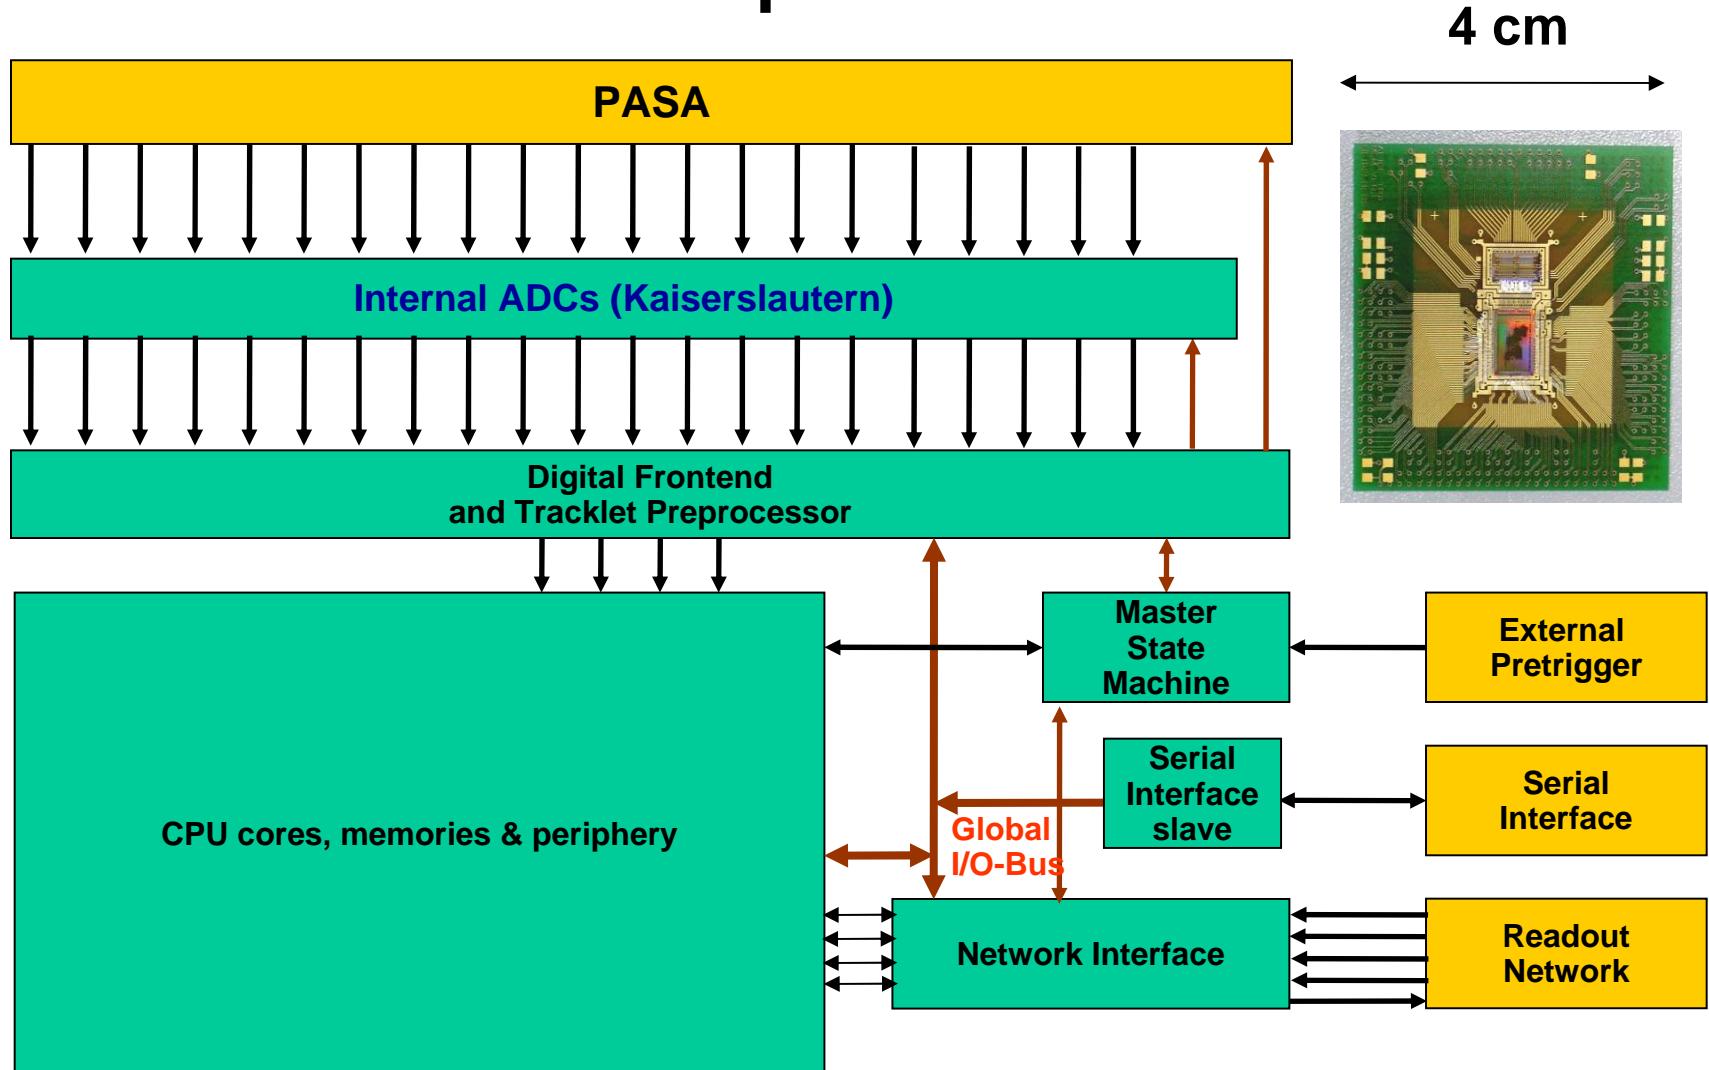
MCM as Readout Network

- data source only
- data source and data merge
- data merge only

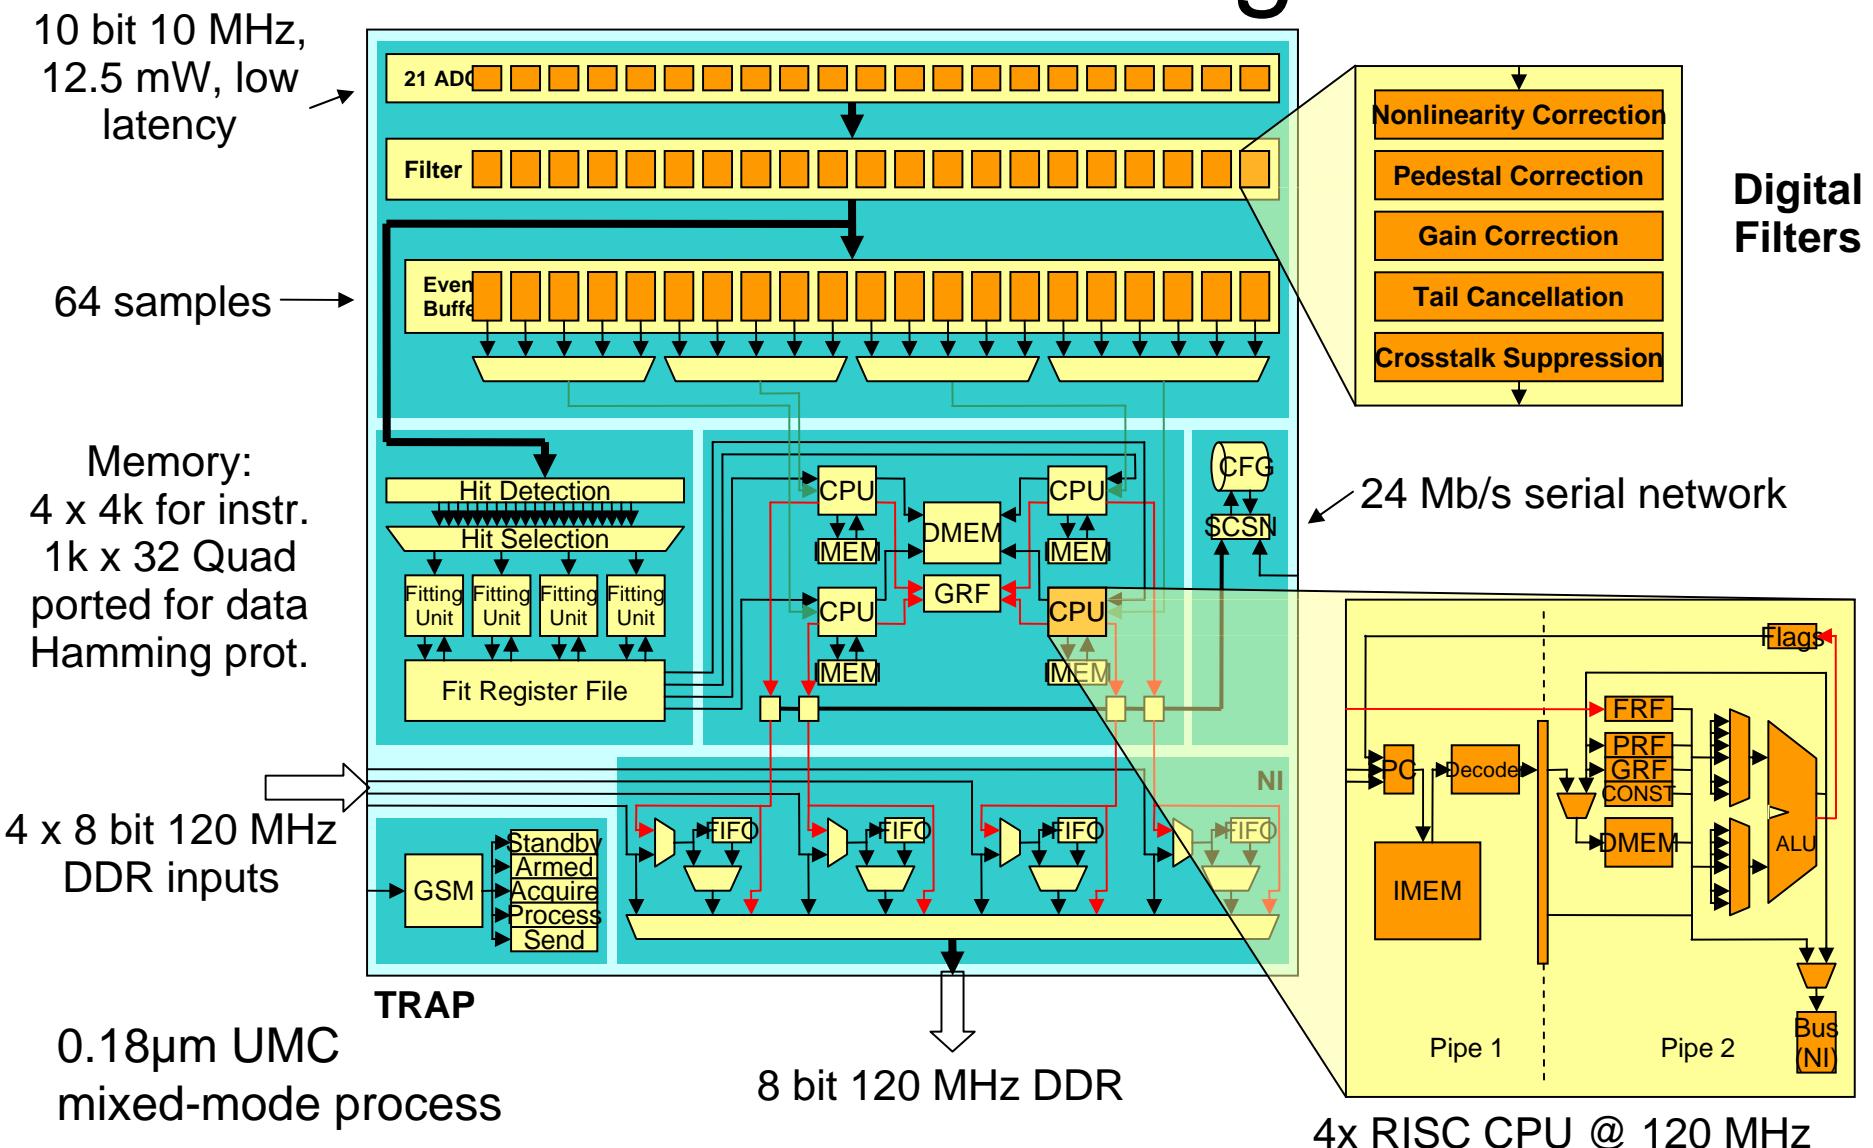
Detector Readout



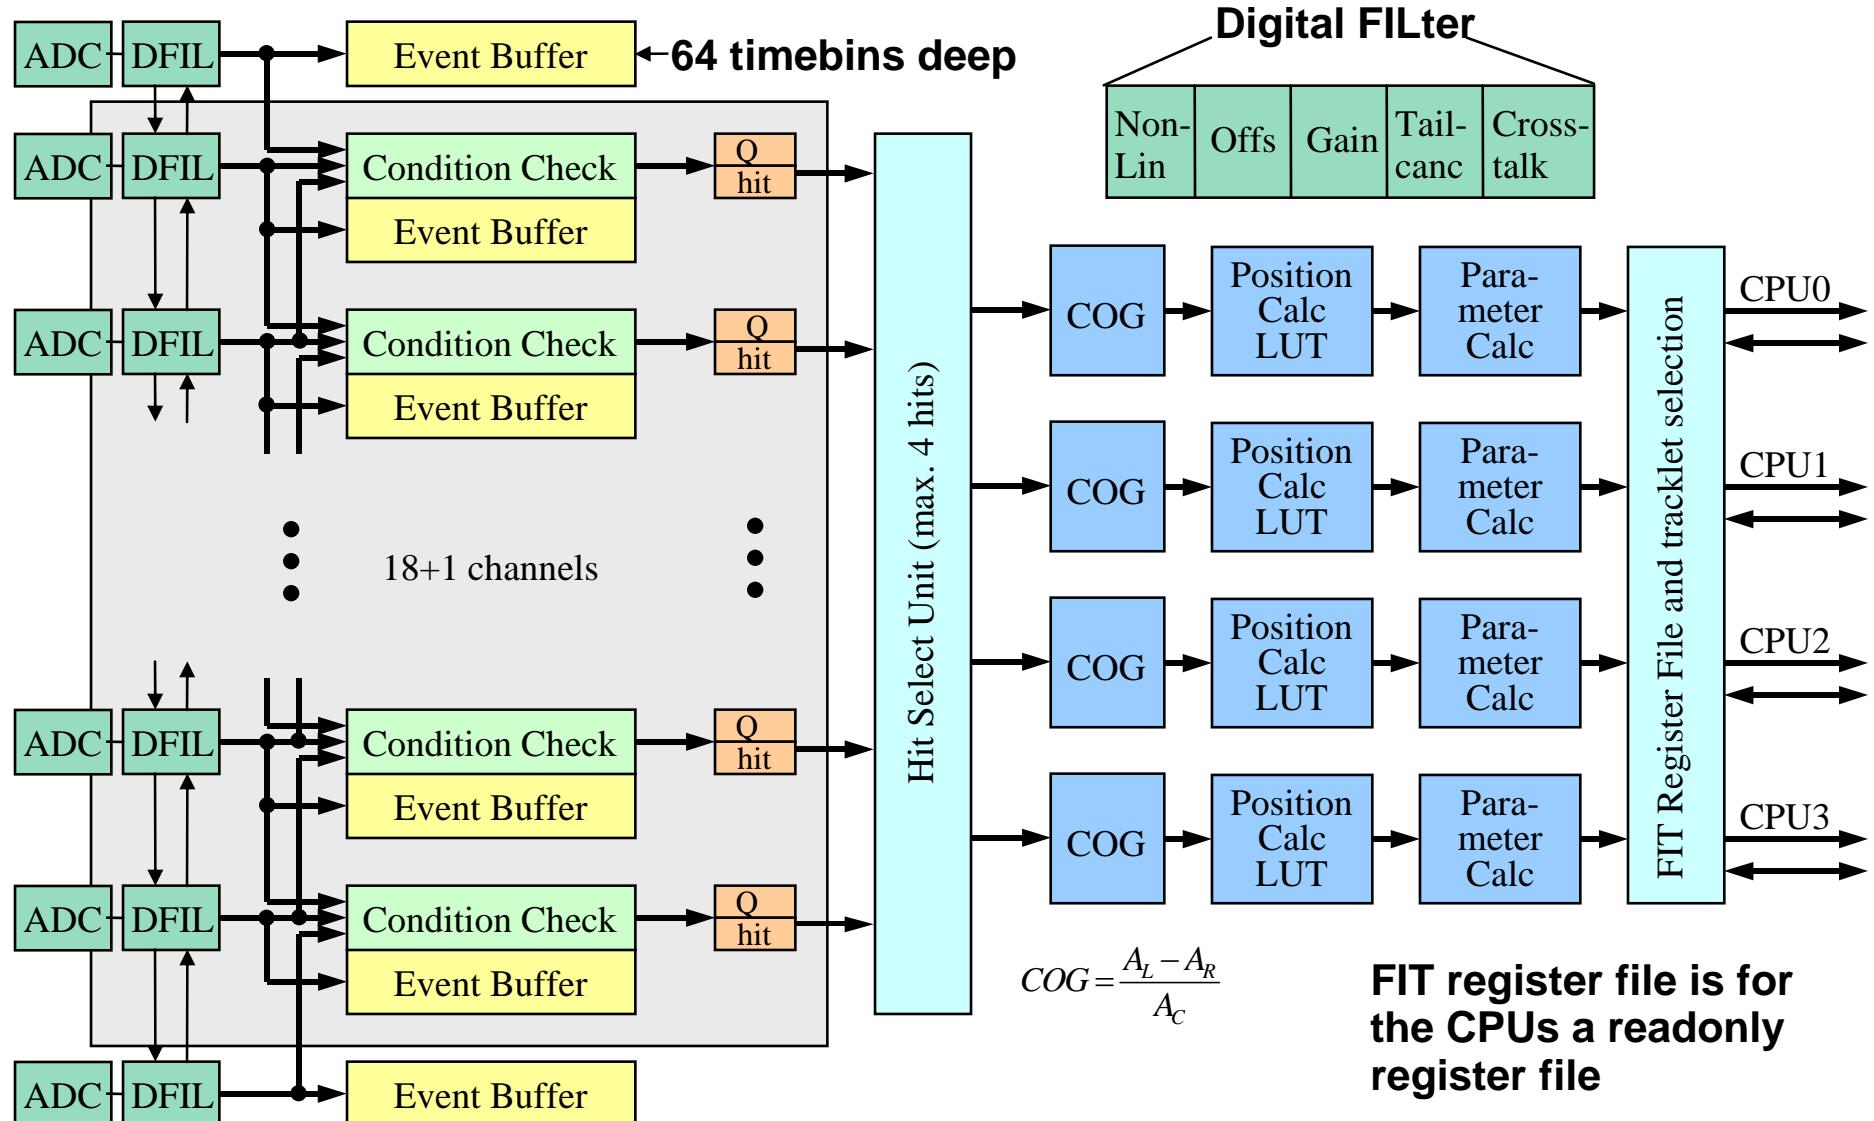
Multi Chip Module



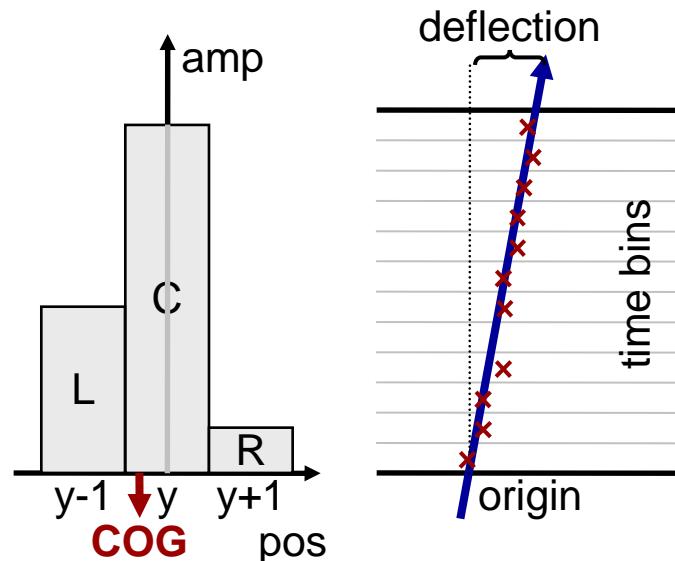
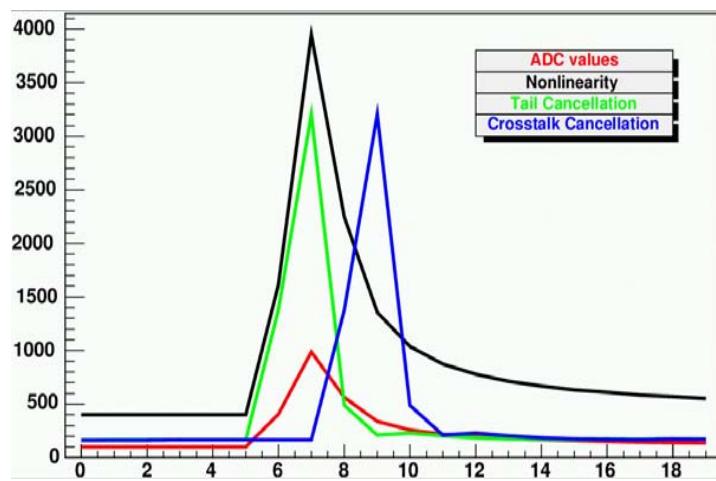
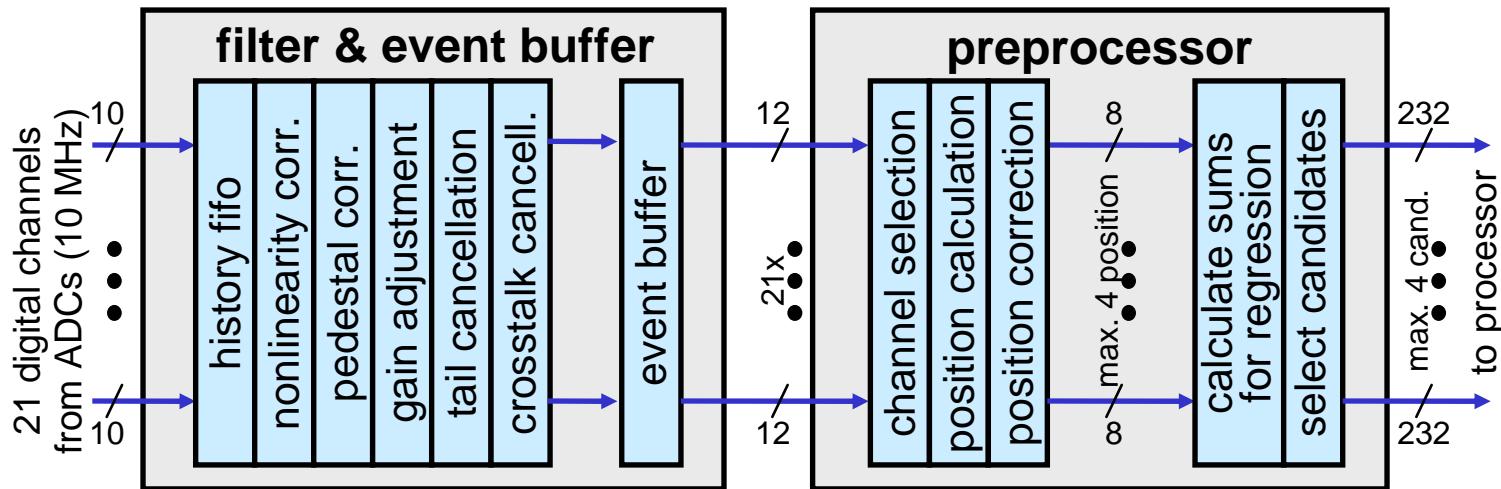
TRAP block diagram



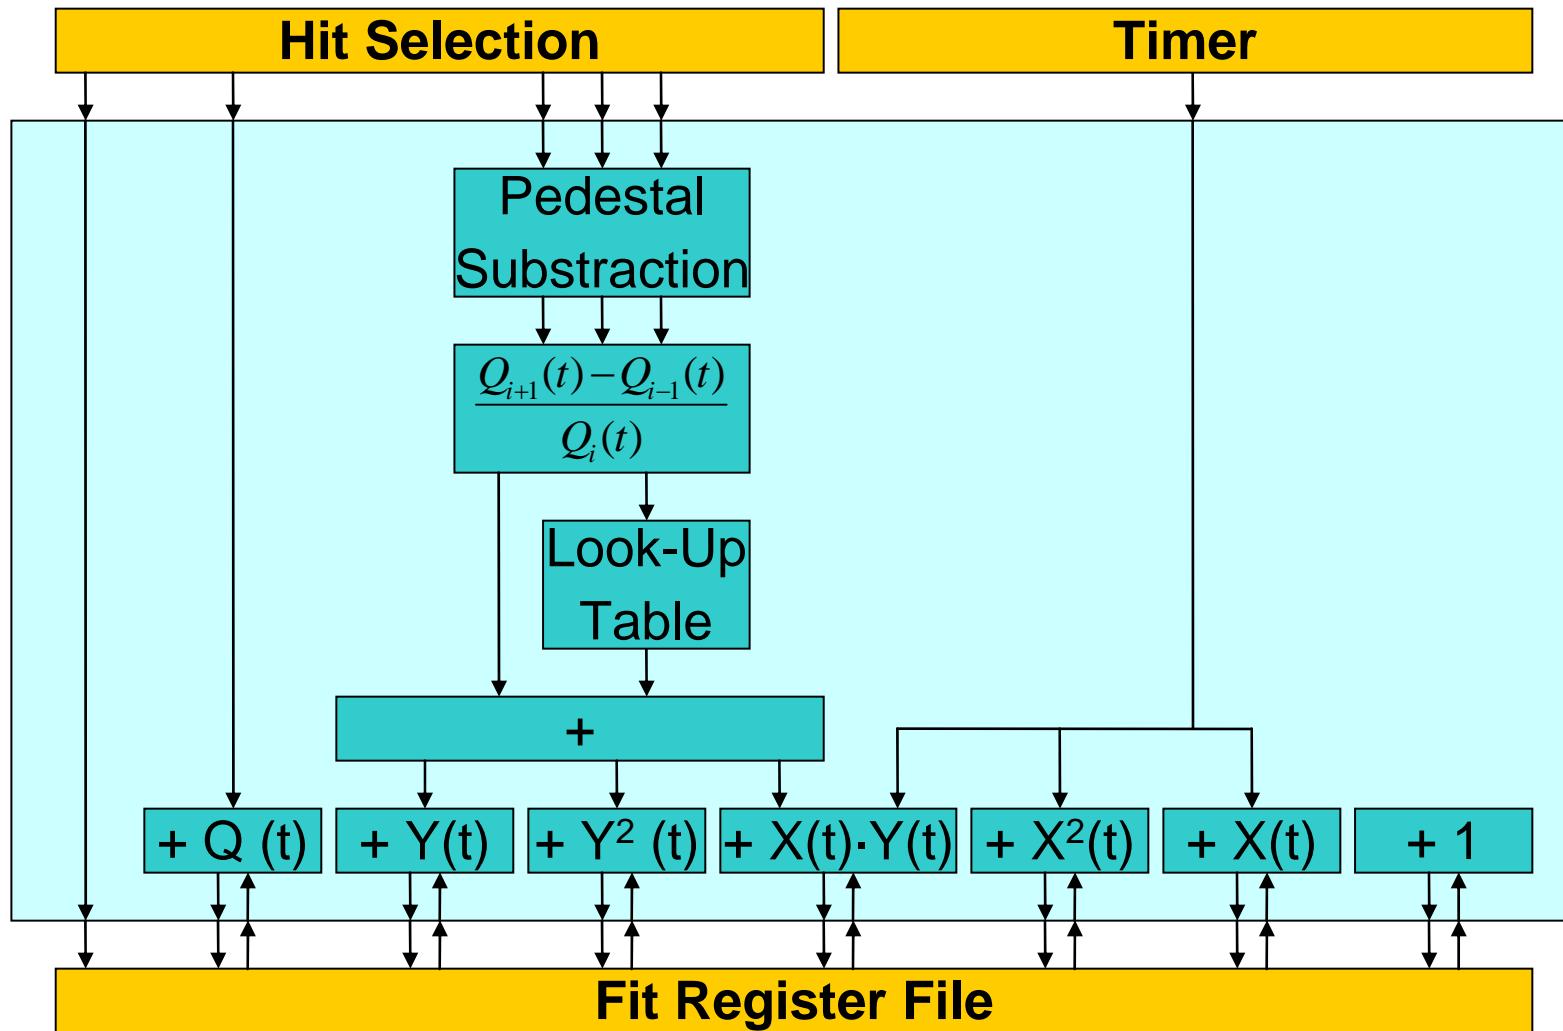
Filter and Tracklet Preprocessor



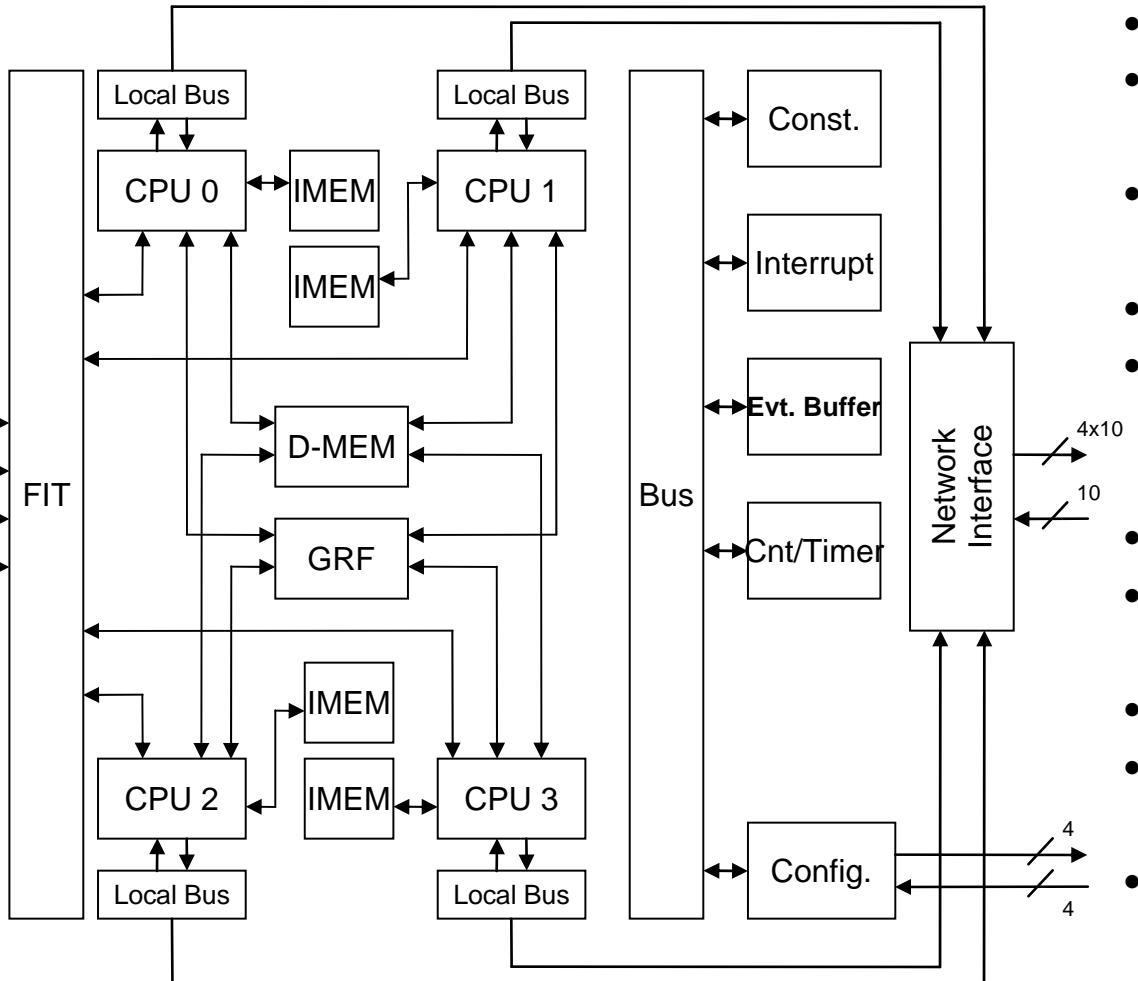
Filter & Preprocessor



Tracking Arithmetics



The MIMD Architecture



- Four RISC CPU's
- Coupled by Registers (GRF) and Quad ported data Memory
- Register coupling to the Preprocessor
- Global bus for Periphery
- Local busses for Communication, Event Buffer read and direct ADC read
- I-MEM: 4 single ported SRAMs
- Serial Interface for Configuration
- IRQ Controller for each CPU
- Counter/Timer/PsRG for each CPU and one on the global bus
- Low power design, CPU clocks gated individually

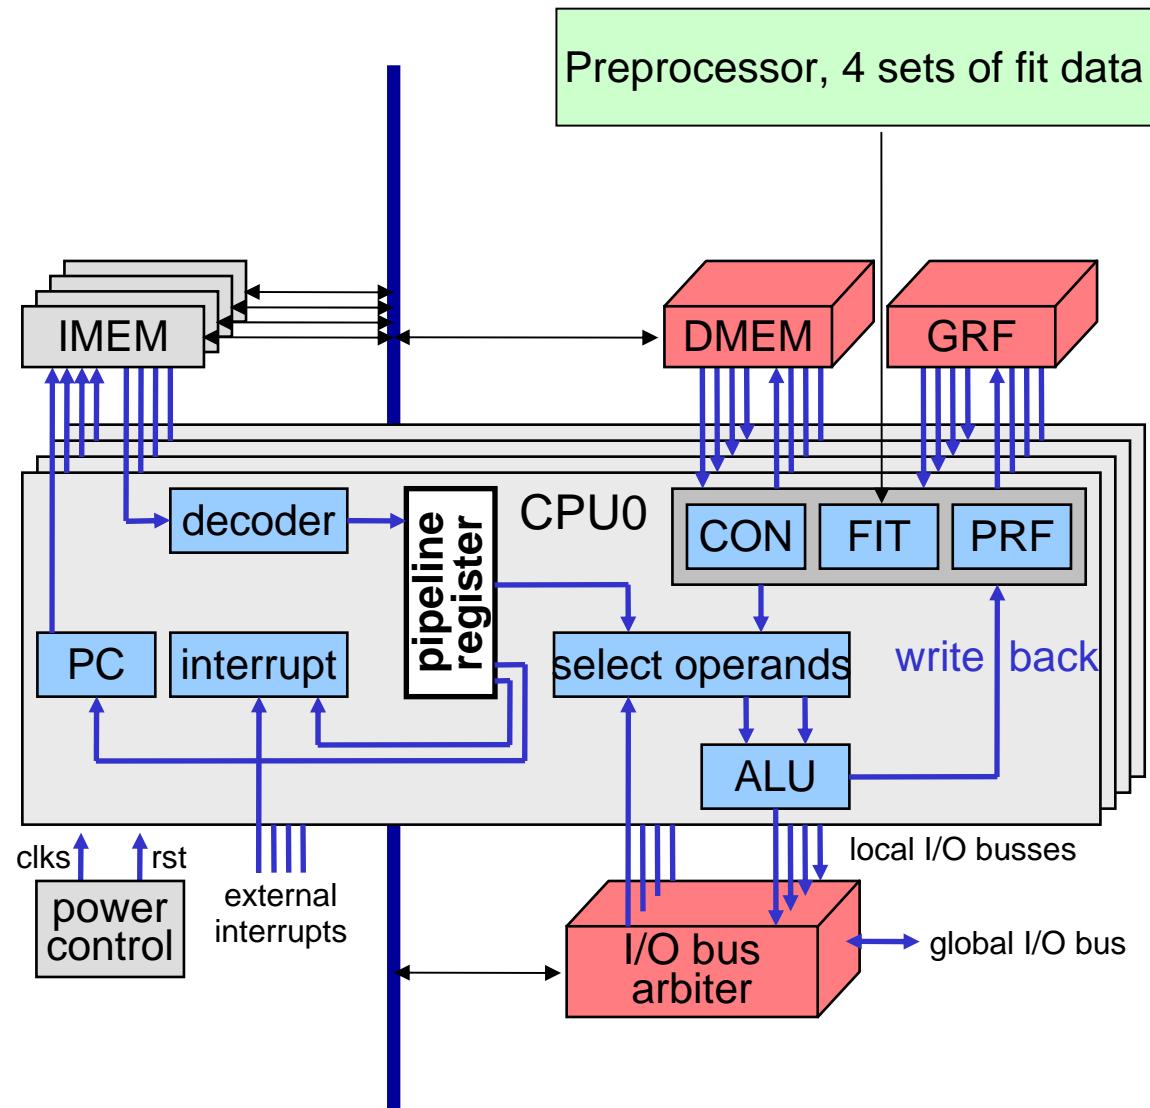
MIMD Processor

MIMD processor

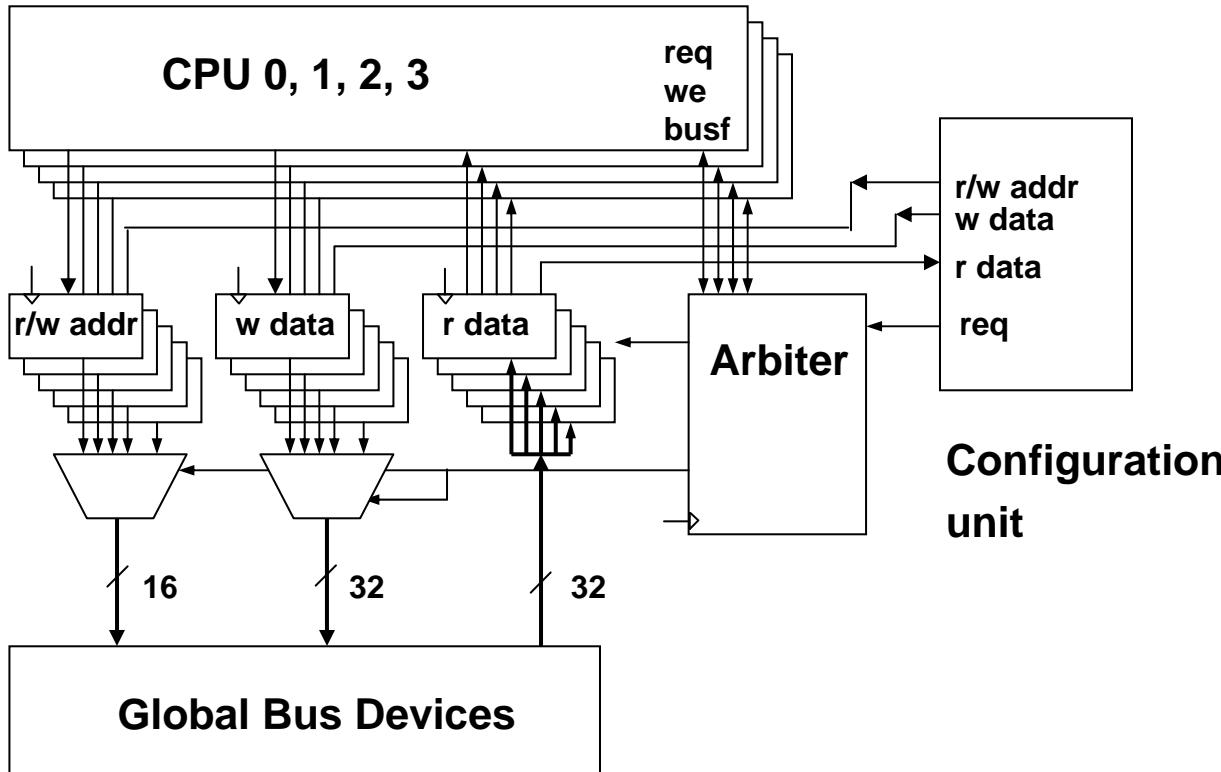
- 4 CPUs
- shared memory / register file
- global I/O bus arbiter
- separate instruction memory
- coupled data & control paths

CPU

- Harvard style architecture
- two stage pipeline
- 32 Bit data path
- register to register operations
- fast ALU
 - 32x32 multiplication
 - 64/32 radix-4 divider
- maskable interrupts
- synchronization mechanisms



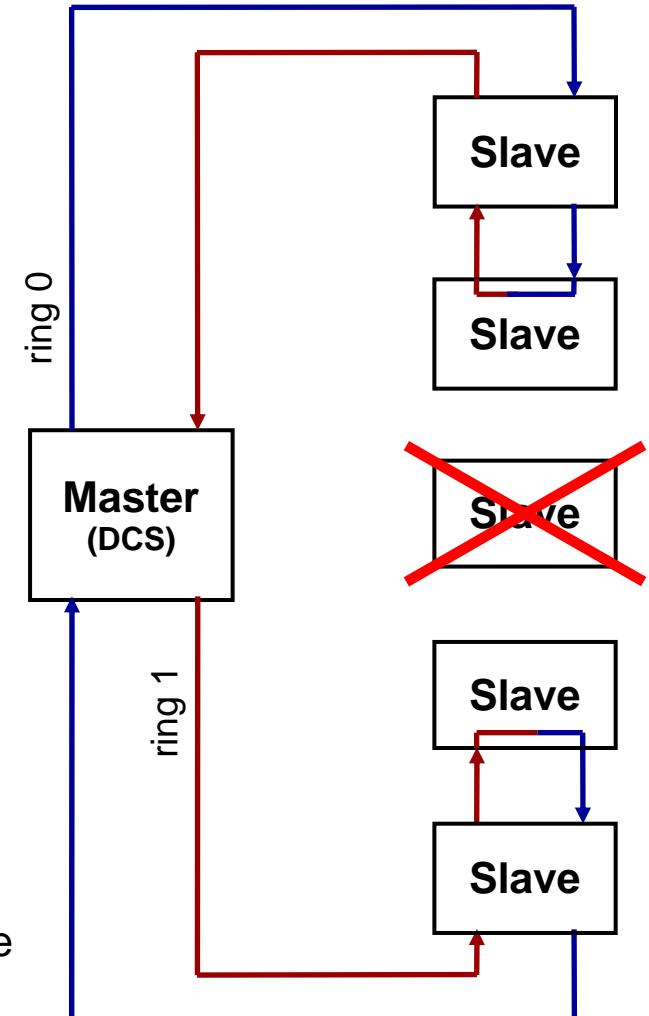
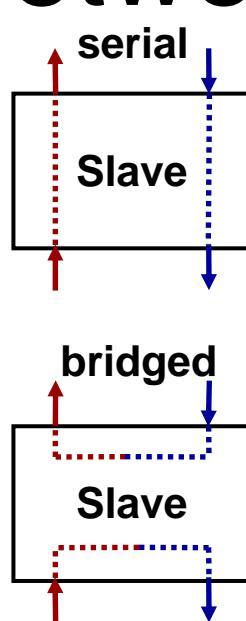
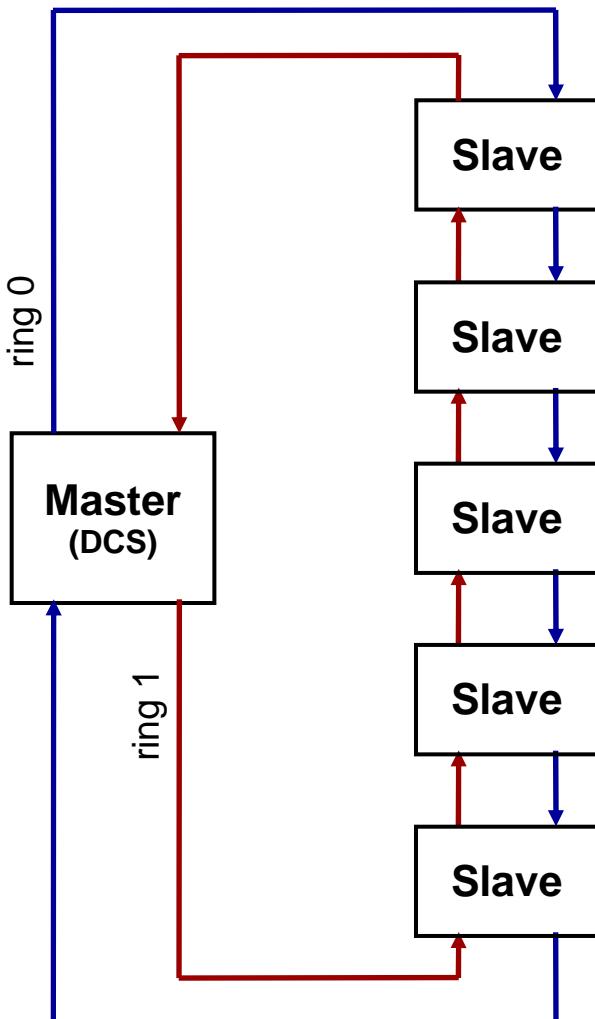
Local and Global IO



- Local bus uses the same r/w address and w_data signals. The read data register on the global bus is a read only device in the local bus.

- Load/Store Instructions
- No tri-state, the output data are ORed, the non-selected devices respond with 0
- Synchronously read/write on the Global Bus (Arbiter), the access time can be programmed.
- Read has priority over write, the configuration unit has priority over CPU 0, 1, 2, 3

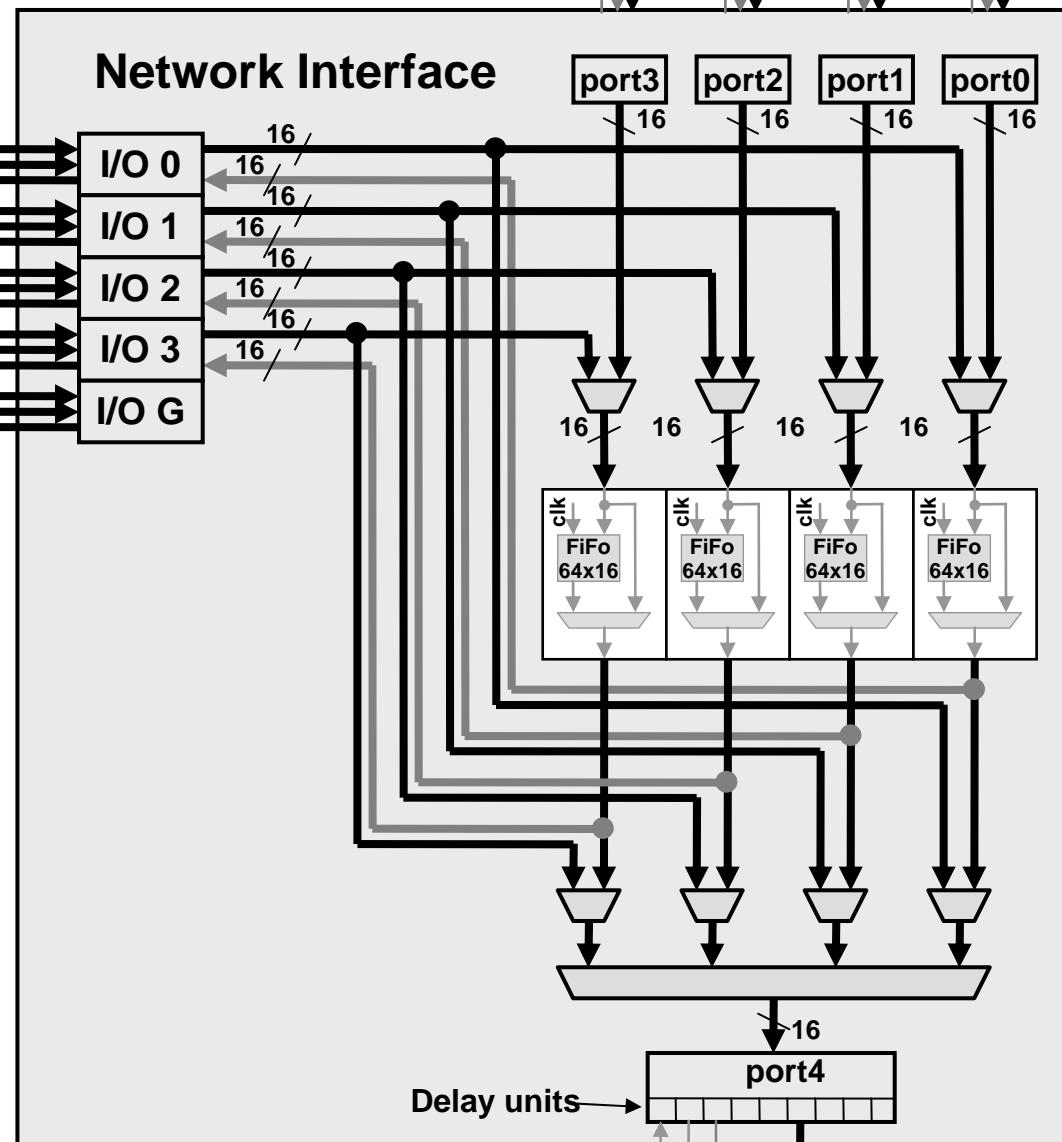
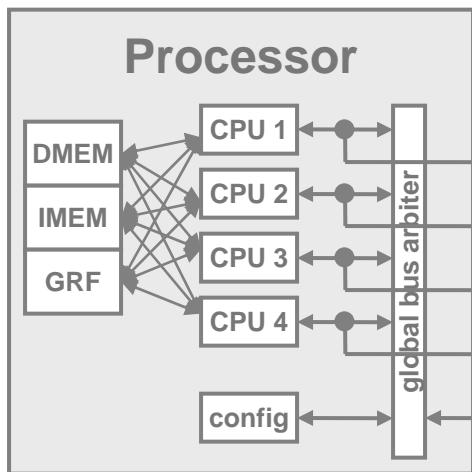
SCSN - Slow Control Serial Network



SCSN

- Up to 126 slaves per ring
- CRC protected
- 24 MBit/s transfer rate
- 16 addr., 32 data-bits/frame

NI Datapath



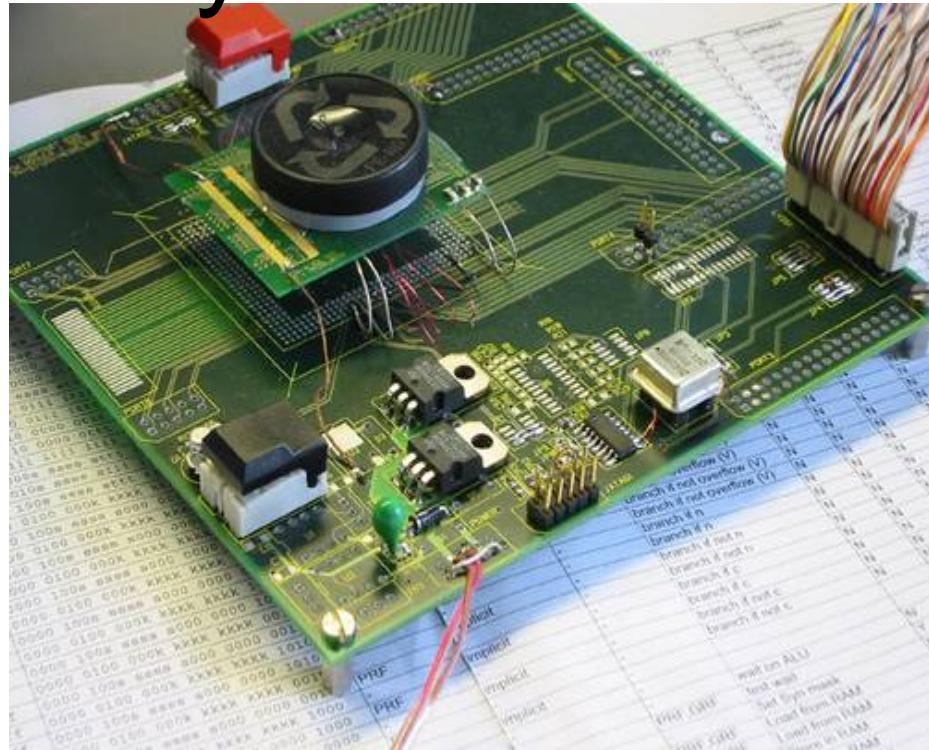
Network Interface

- local & global I/O interfaces
- input port with data resync. and DDR decoding
- input FIFOs (zero latency)
- port mux to define readout order
- output port with DDR encoding and programmable delay unit

TRAP development – a long long way



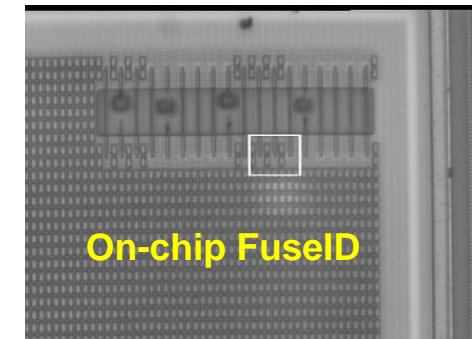
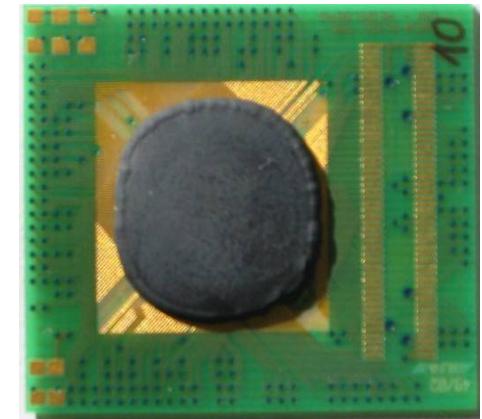
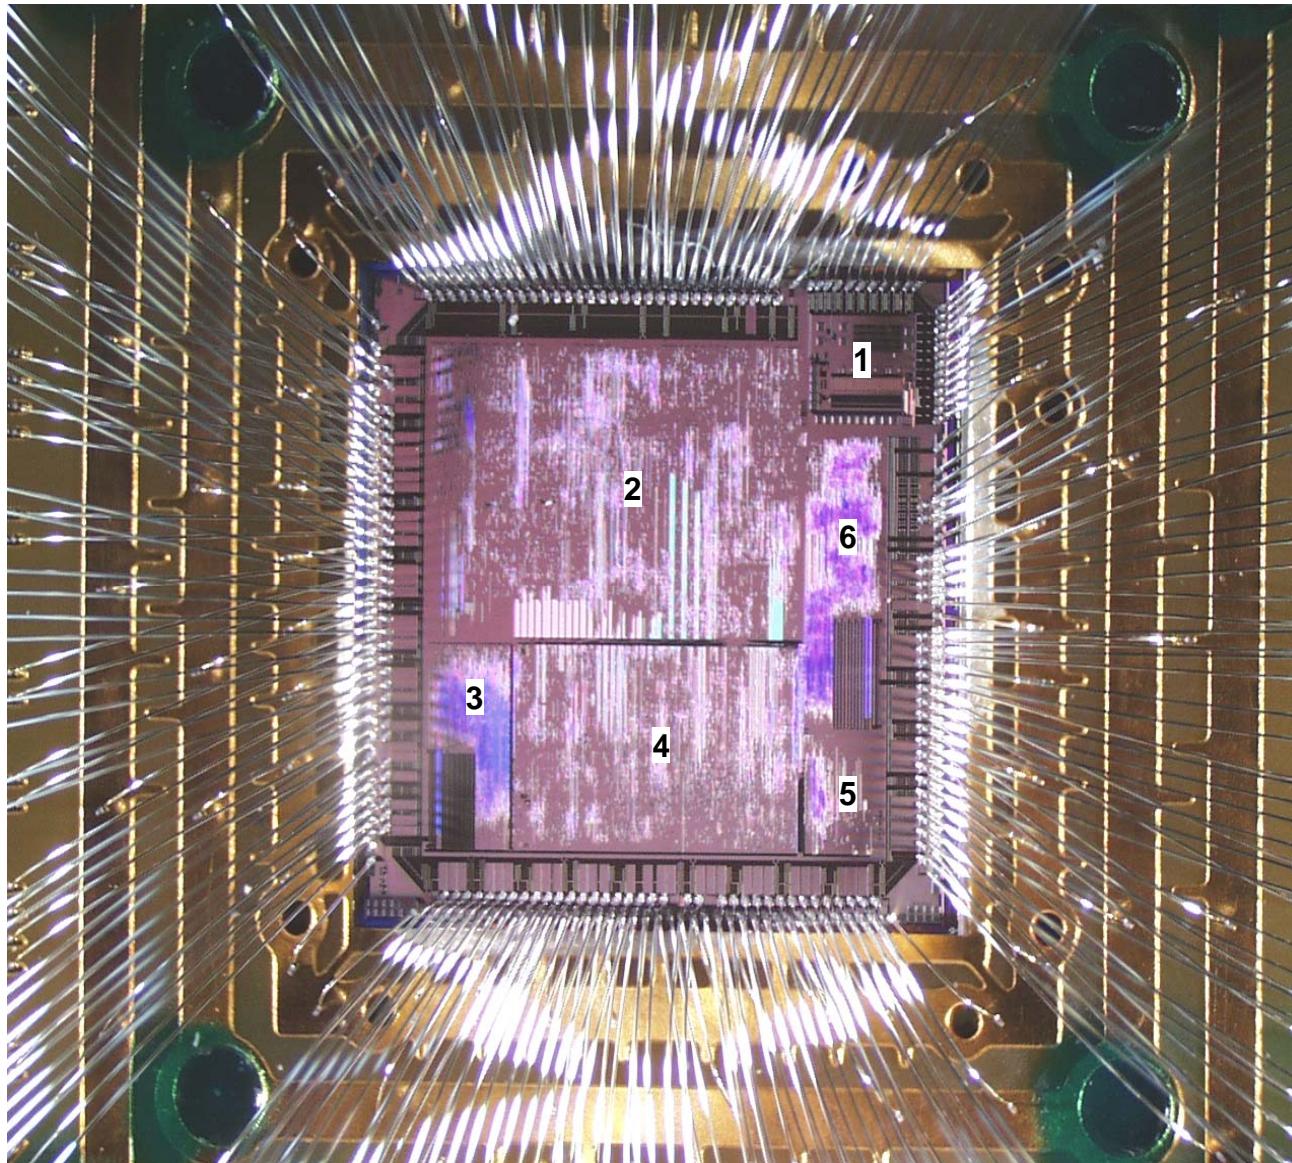
MCM for 8 channels with the first prototypes of the Digital chip (FaRo-1) and Preamplifier, commercial ADCs
Beg 2001



First tested TRAP chip, in „spider“ mode
Summer of 2002

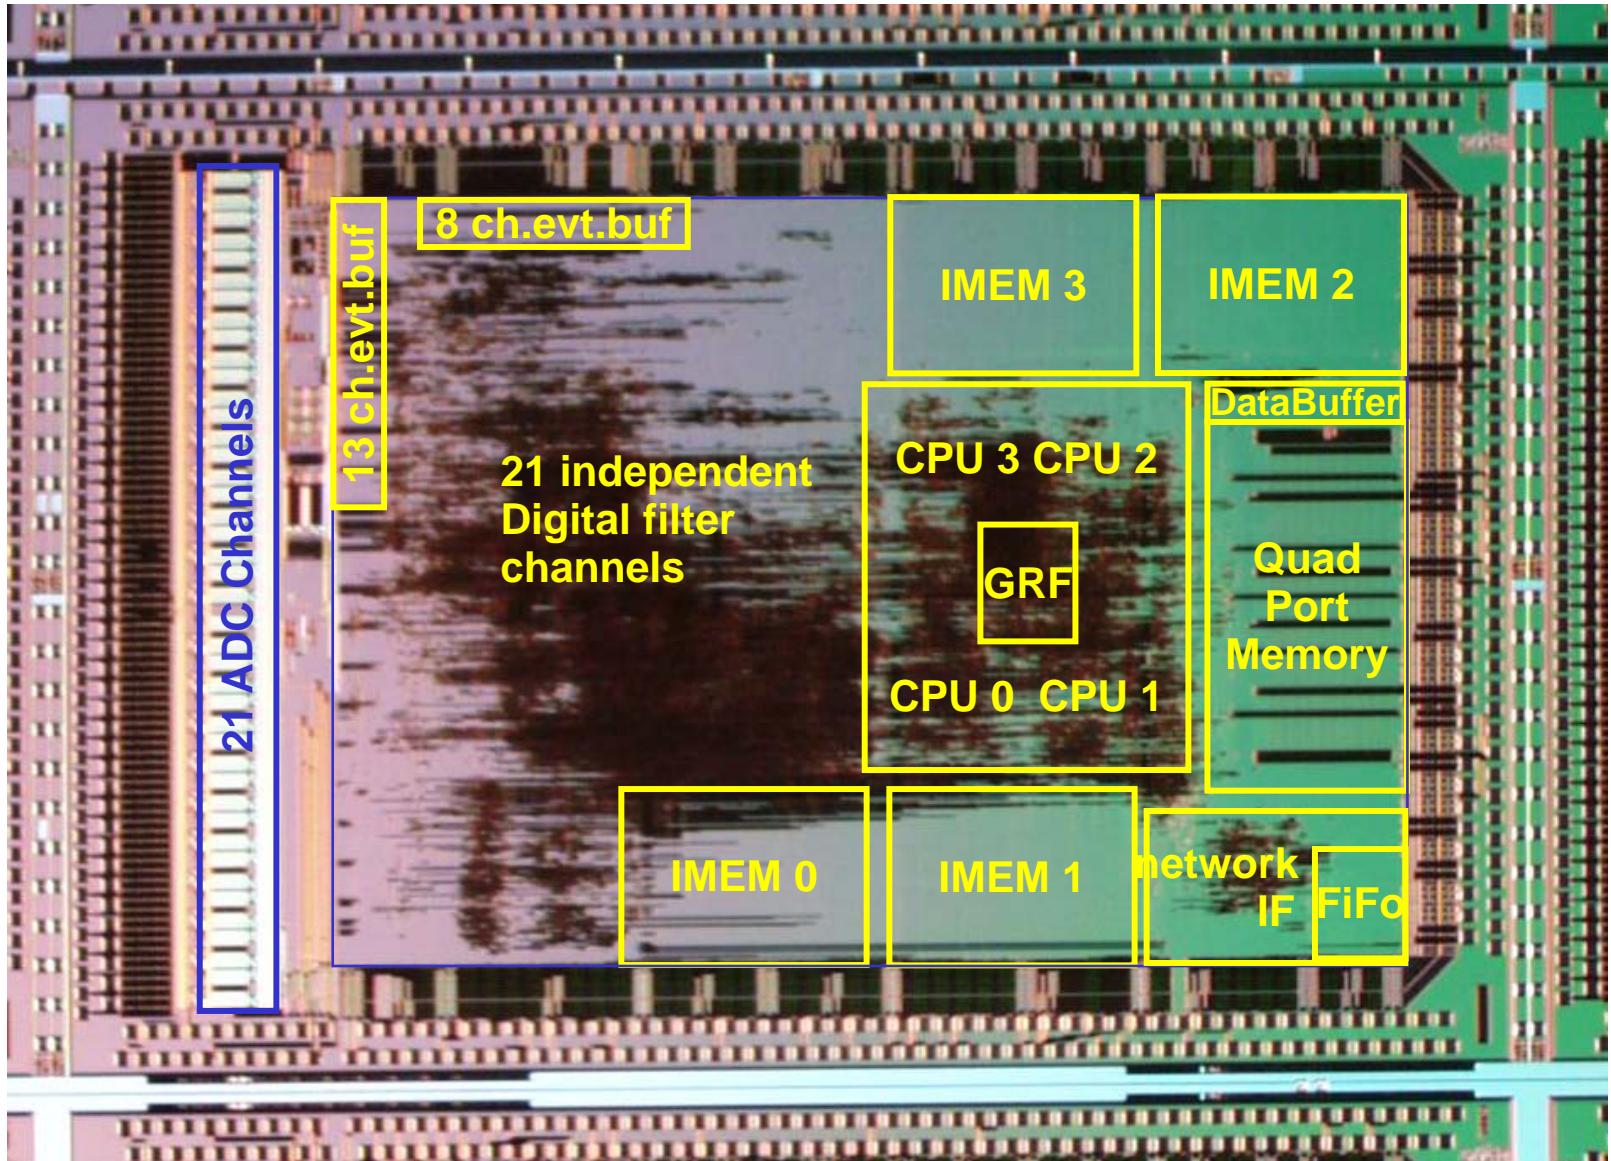
In total \approx 60,000 lines (synthesis) and 18,000 lines (simulation) of VHDL code

TRAP1 bonded on a MCM

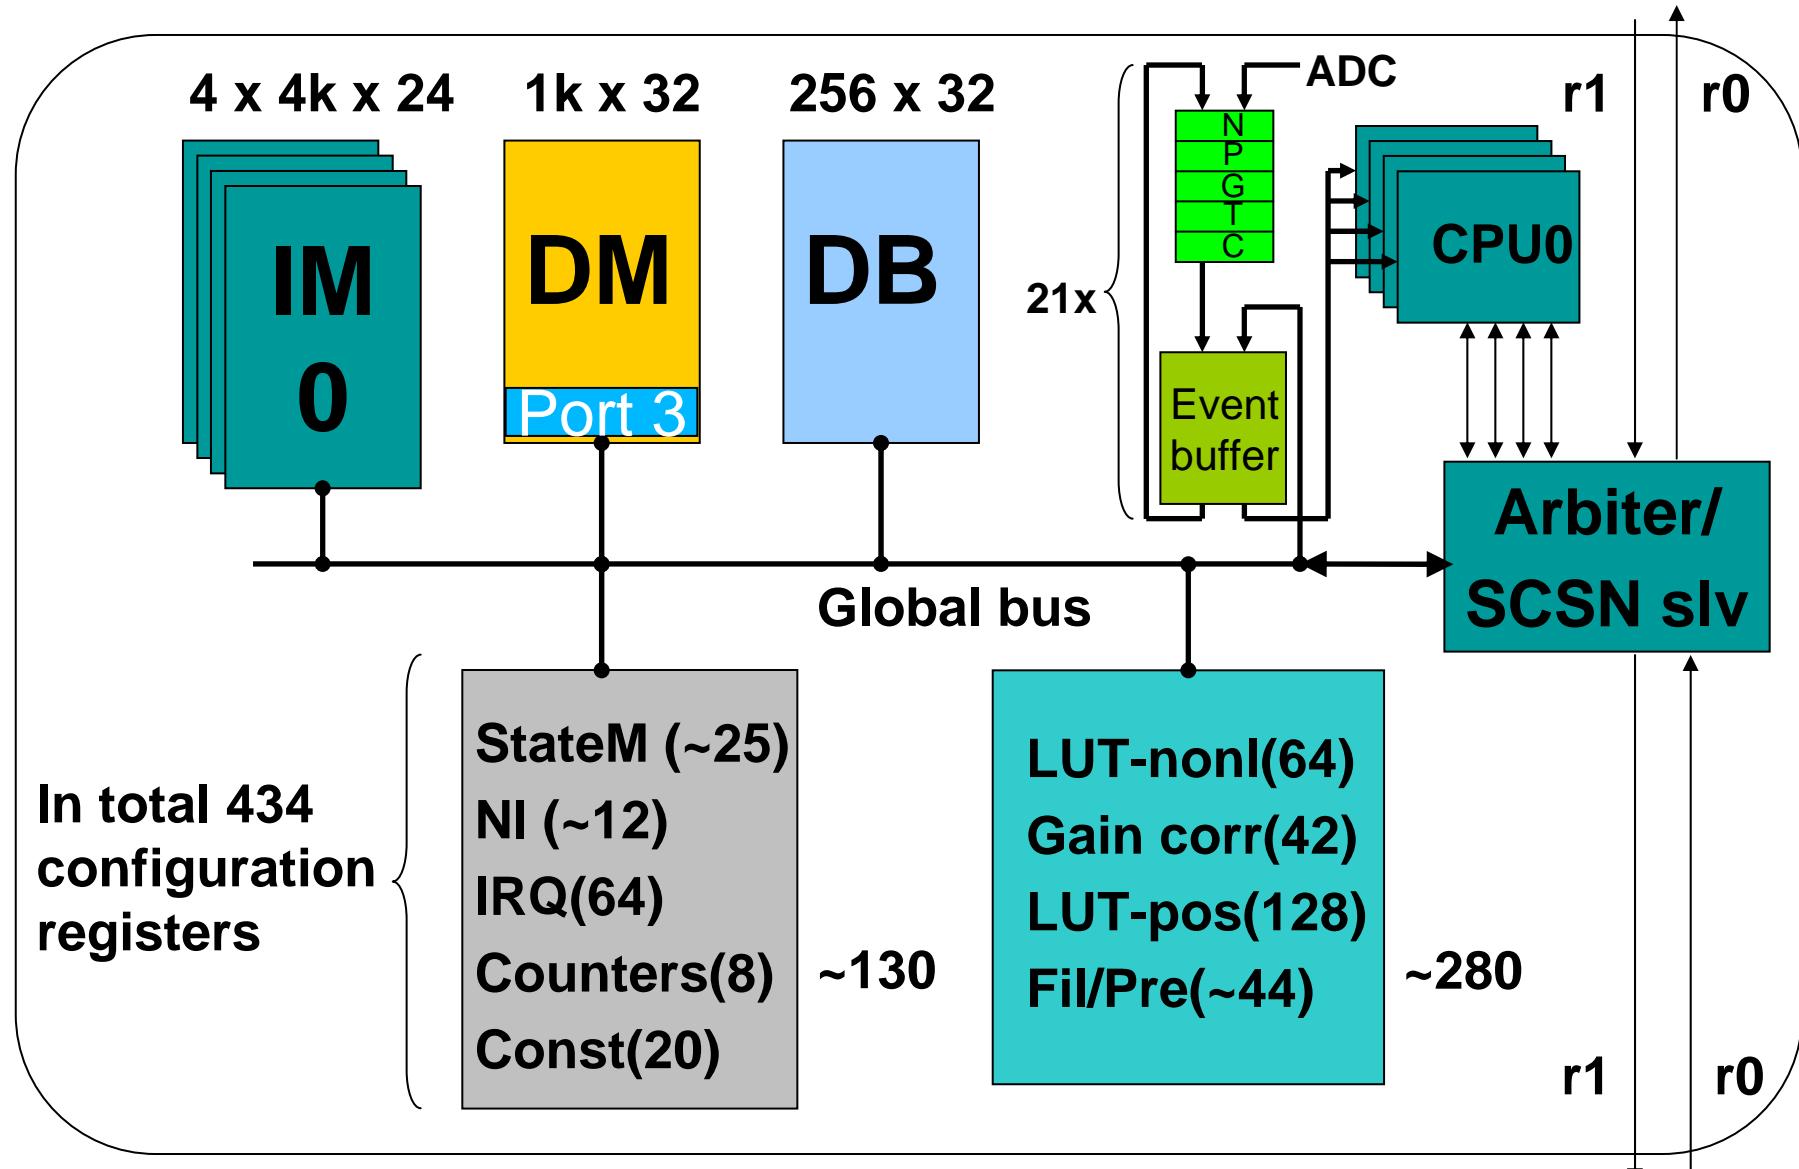


1. ADC
2. Filter/Preprocessor
3. DMEM
4. CPUs
5. Network Interface
6. IMEM

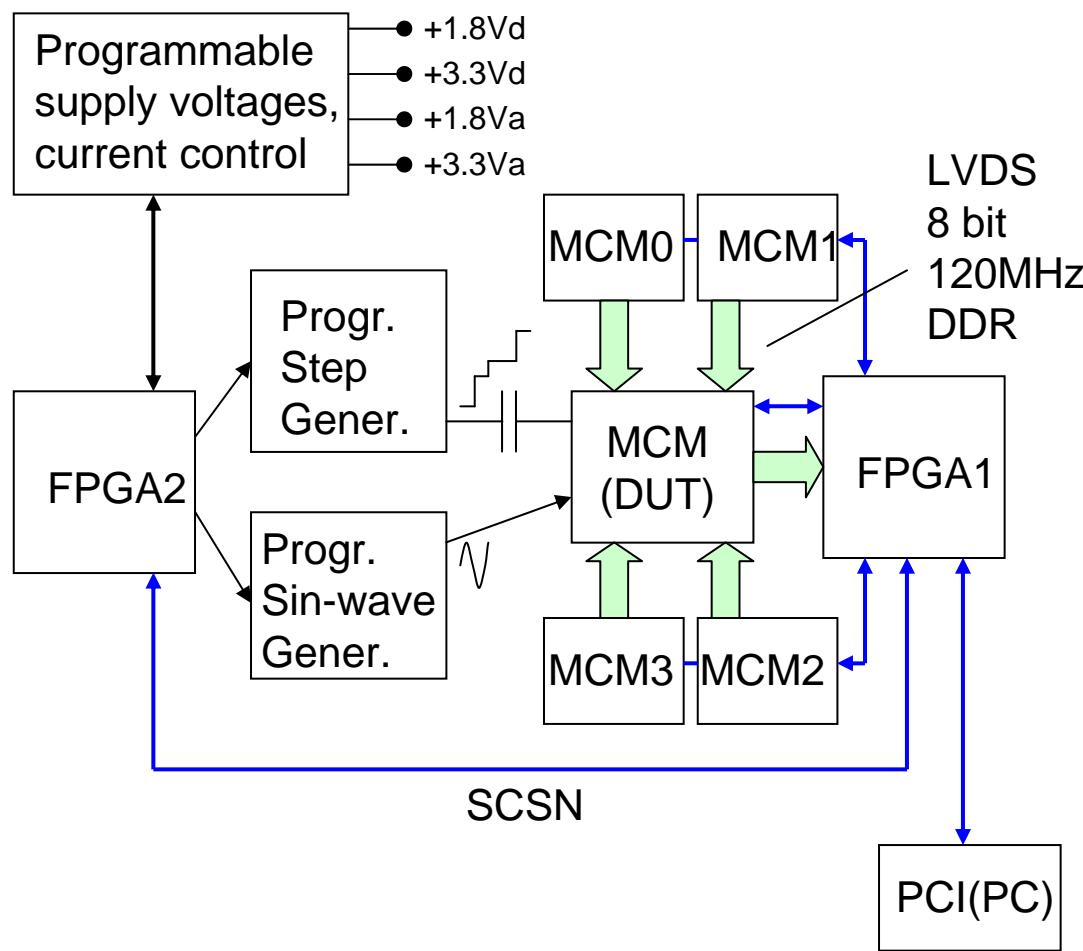
TRAP Layout



TRAP internal tests

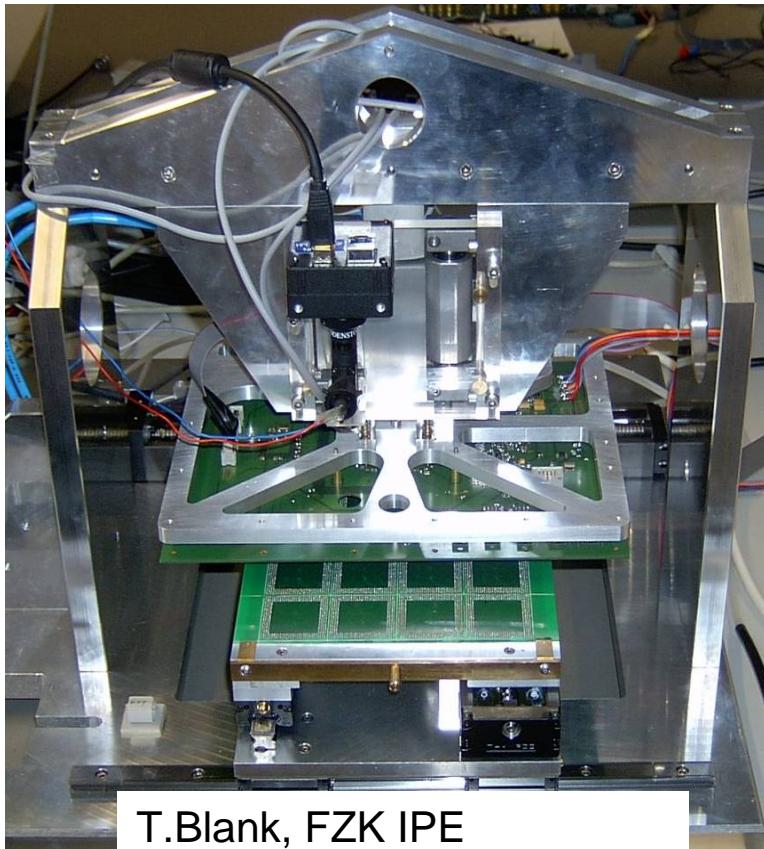


Test flow of the MCM testing



- Apply voltages, control the currents
- JTAG connectivity test
- Basic test using SCSN (serial configuration bus)
- Test of all internal components using the CPUs
- Test of the fast readout
- Test of the ADCs by applying 200 kHz Sin-wave
- Test of the PASA by applying voltage steps through serial capacitors
- Store all data for each MCM in separate directory, store in XML file the essential results
- Export the result for MCM marking and sorting

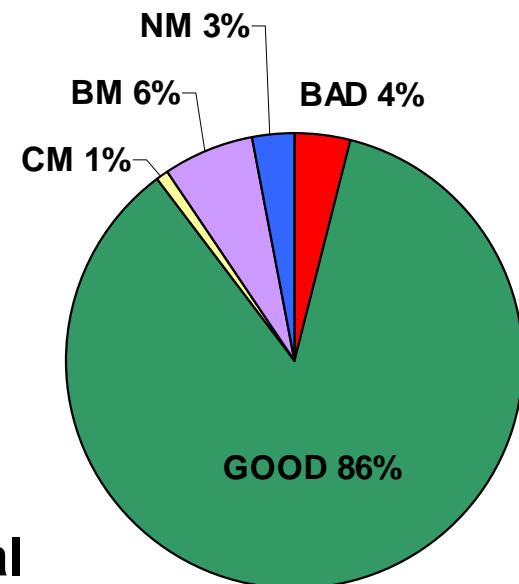
MCM Tester and results



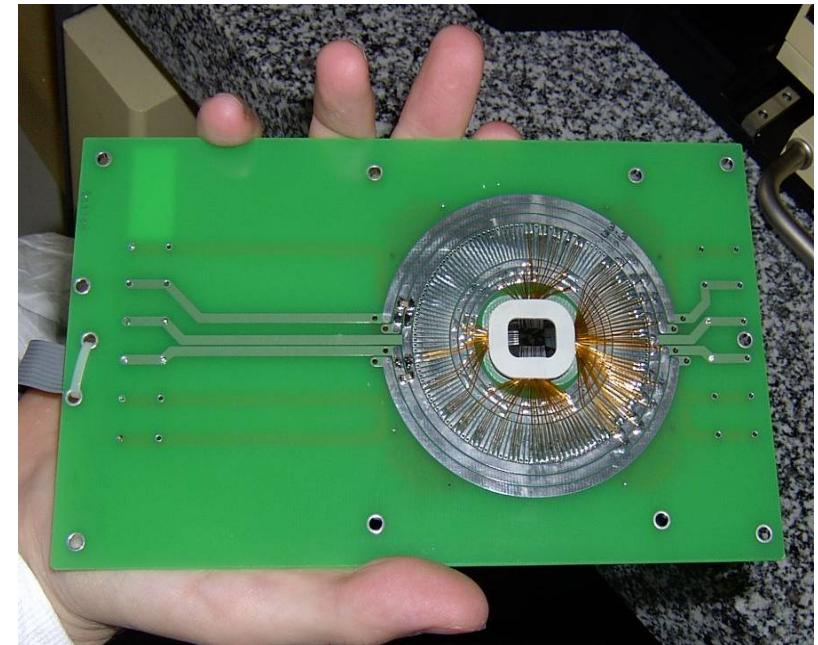
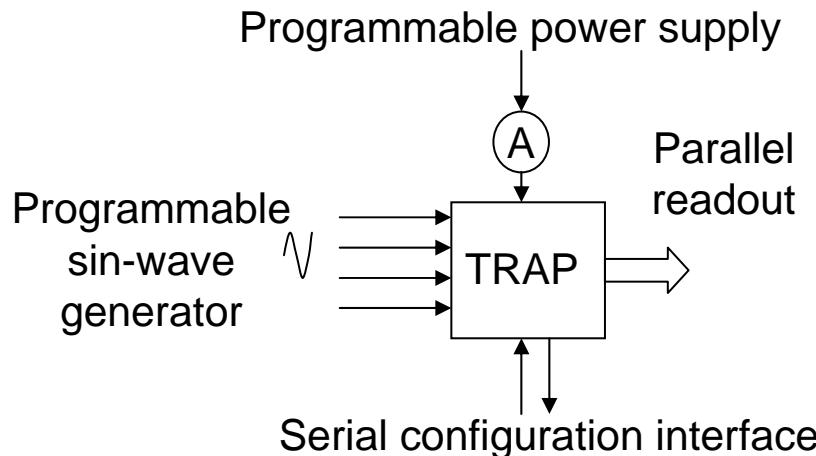
T.Blink, FZK IPE
(Karlsruhe)

- store the result into a DB
- mark later the tested MCMs with serial Nr. and test result code

- test of 3x3 or 4x4 MCMs
- digital camera with pattern recognition software for precise positioning using an X-Y table
- vertical lift for contacting
- about 1 min/MCM for positioning and test



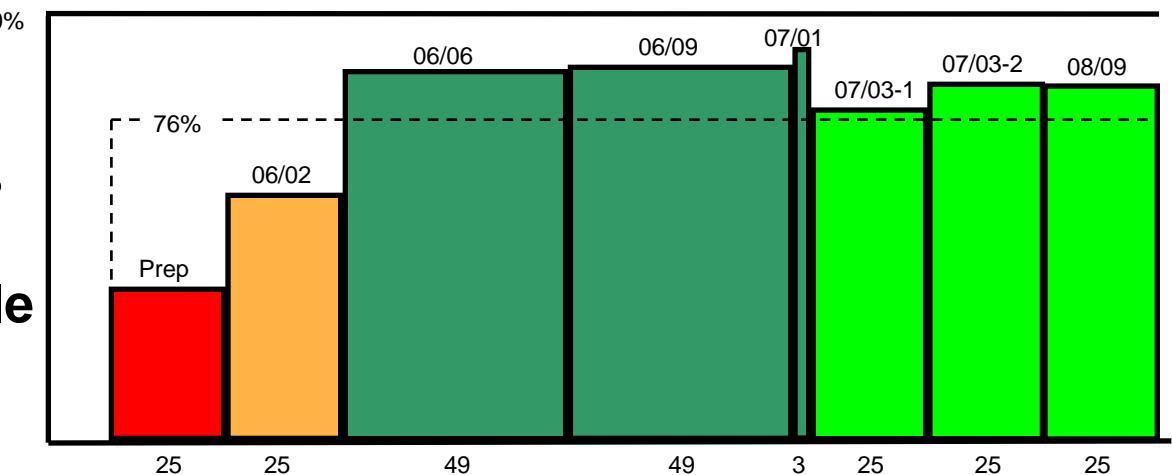
TRAP wafer test and results



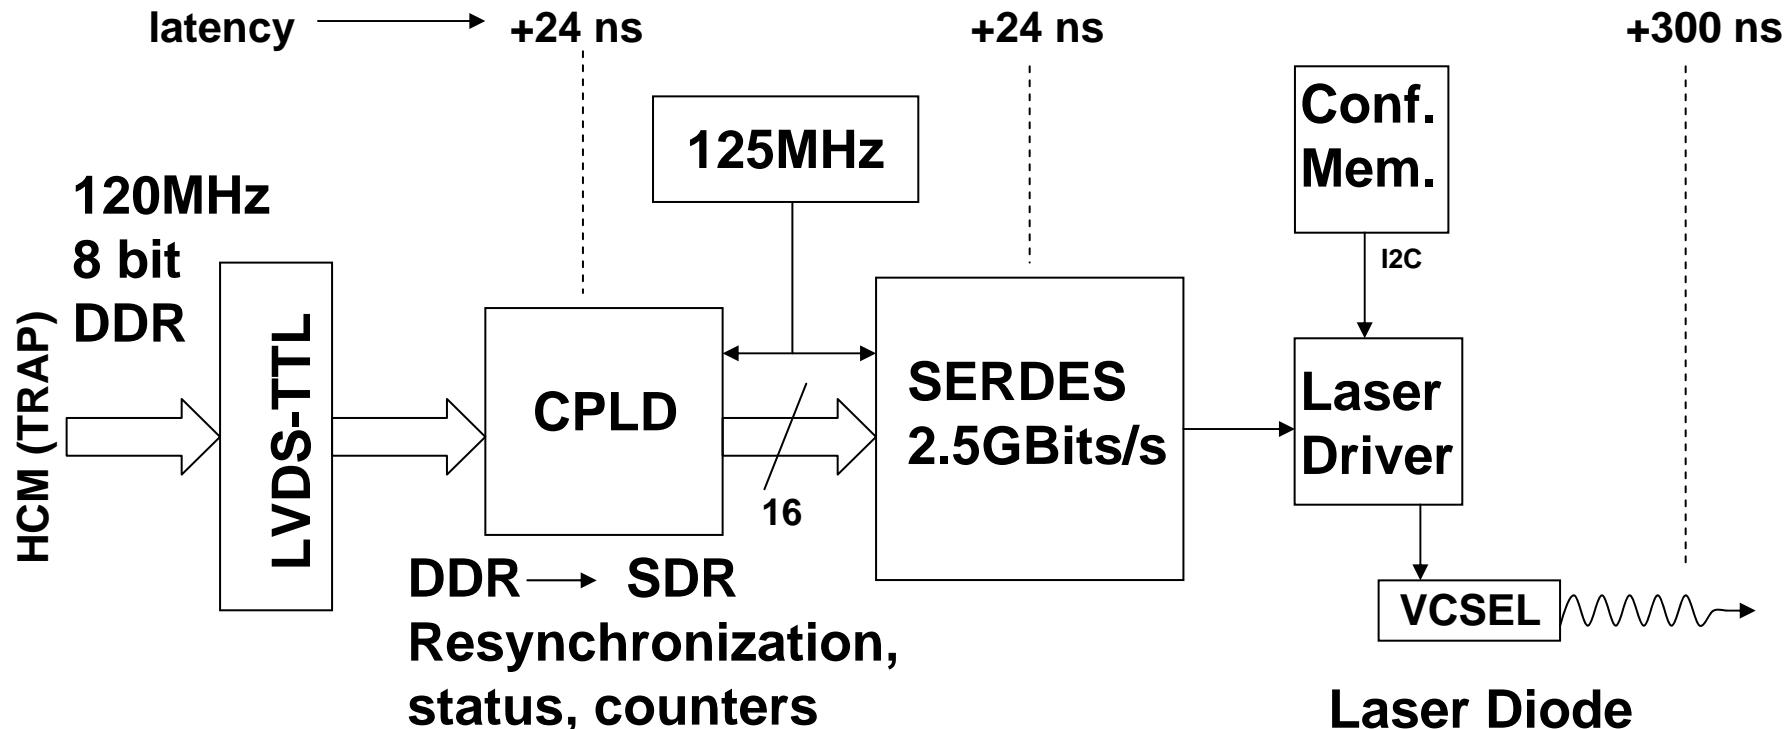
576 TRAP chips/wafer

Fully automatic partial test of the TRAP

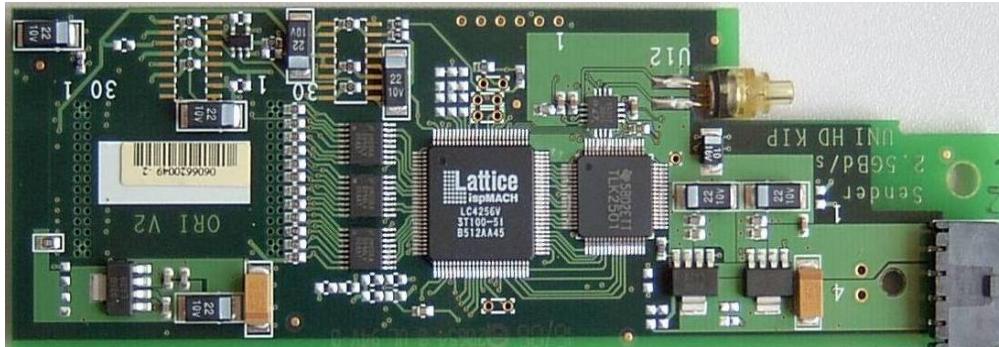
**Up to now produced
and tested 201 wafers
with ~129,000 TRAPs,
of them ~98,000 usable**



Optical Readout Interface (ORI)

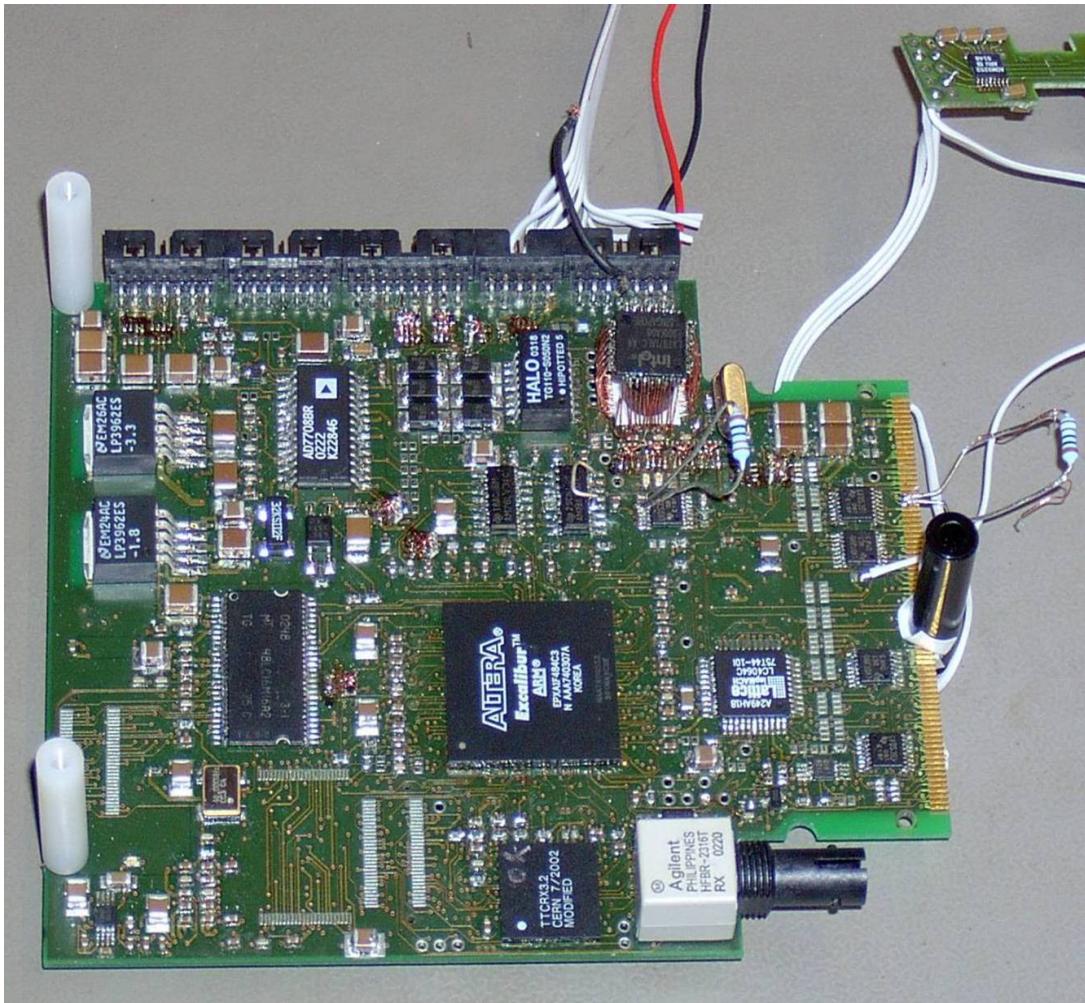


Magnetic field & radiation tolerant!



All 1200 produced and tested, 1199 of them fully functional

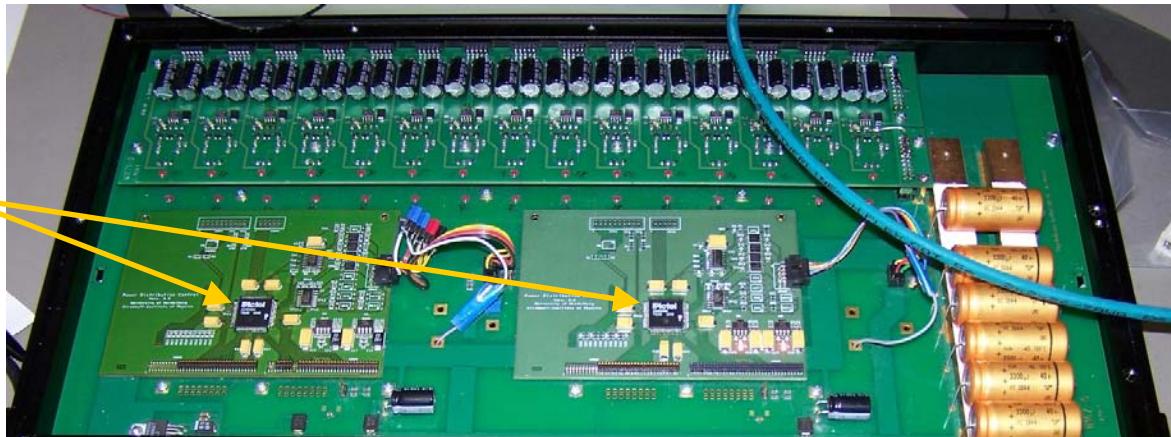
DCS Board



- ARM based technology
- 100k FPGA flexibility
- 32MB SDRAM
- LINUX system with EasyNet

Power Distribution Box

Actel
antifuse
FPGA



Switch on/off
the power
supply to 30
DCS boards

Altera
ARM+FPGA



Control of 9
PDB/2

540 DCS boards in ALICE TRD

Design of VLSI Circuits using VHDL

The ALICE TRD Global Tracking Unit

An Example of a
large FPGA-based System
in High-energy Physics

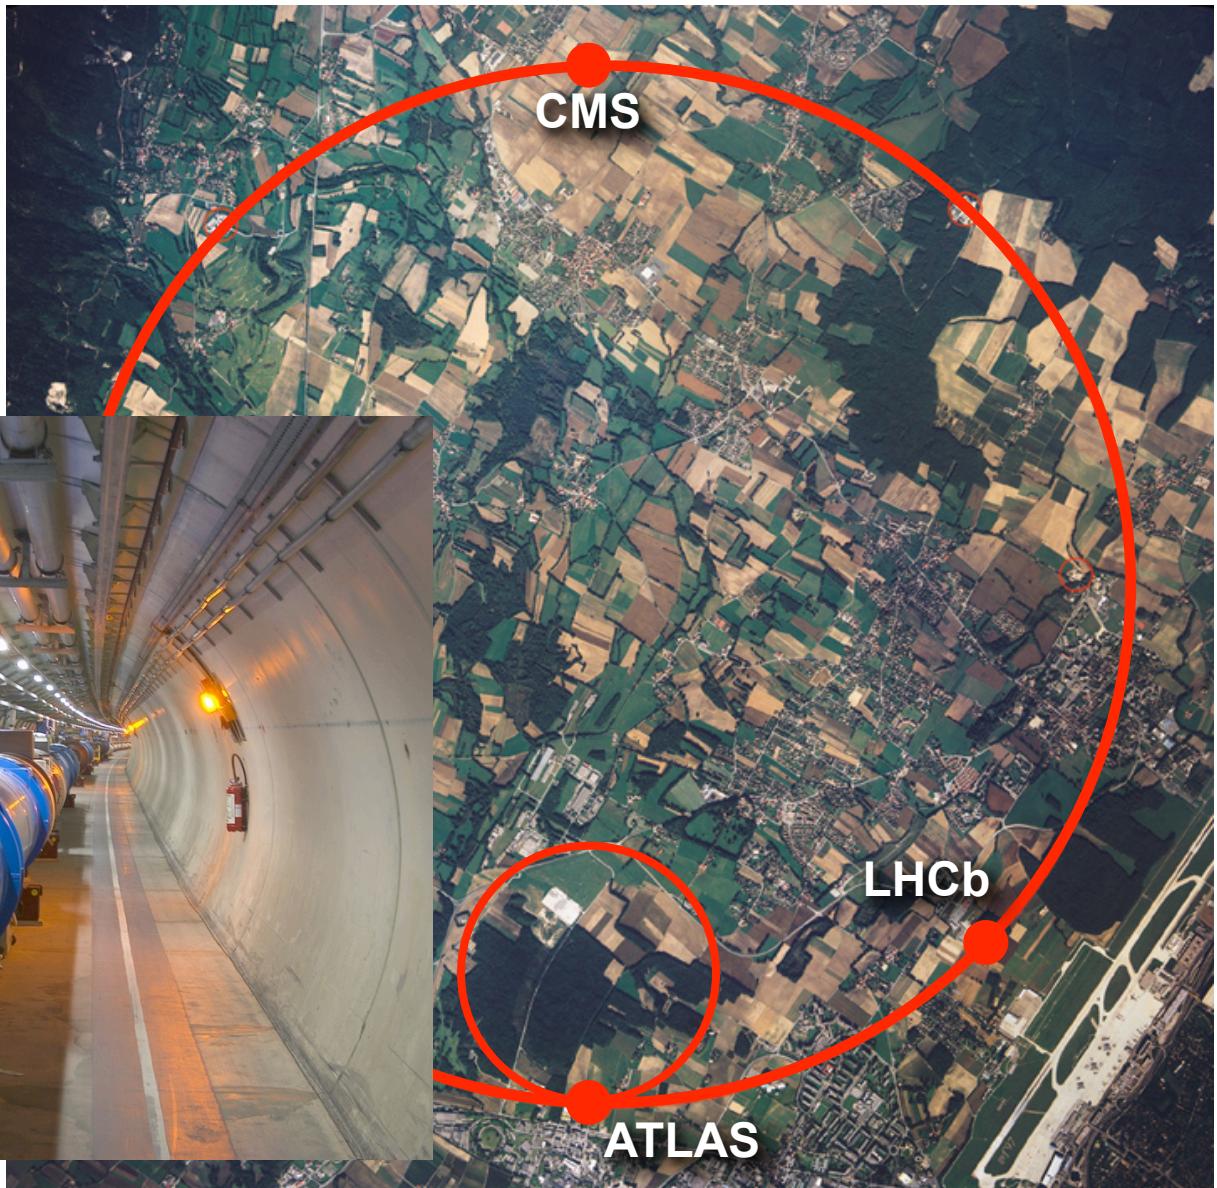
A Typical Example

- High-energy & heavy ion experiments:
 - Huge amounts of data,
selection of interesting events → triggers
 - Performance limited by data processing
power of front-end electronics
- Requirements for electronics:
 - Complex trigger algorithms, very short decision times
→ high-performance & low latency processing
 - Advanced trigger interlacing strategies to minimize
detector dead times (multi-event buffering)
→ high bandwidth data paths
 - Demands change quickly as research advances
→ flexibility

The Large Hadron Collider

LHC

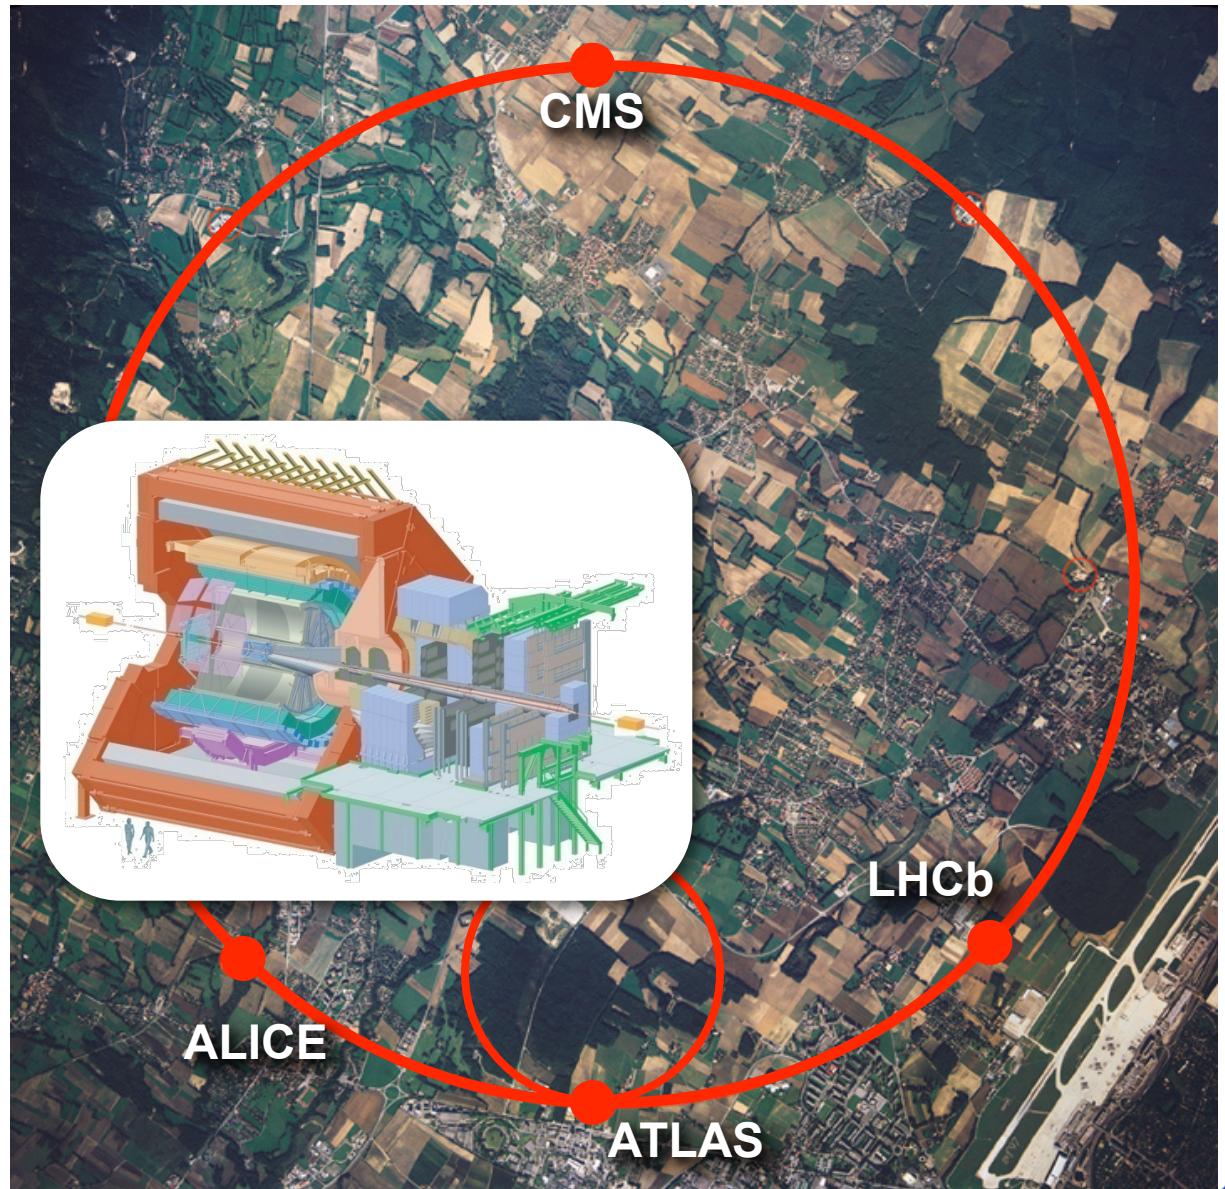
- p-p @ 14 TeV
- Pb-Pb
@ 1150 TeV



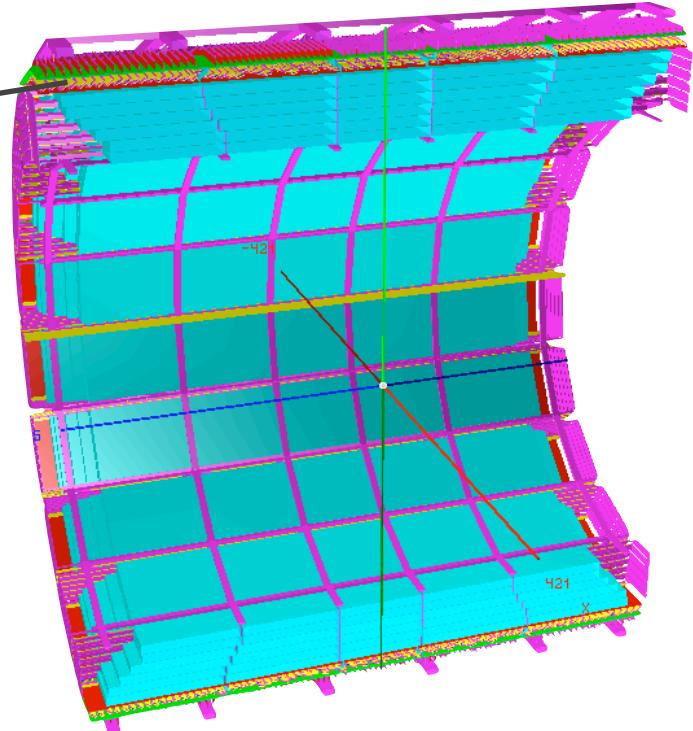
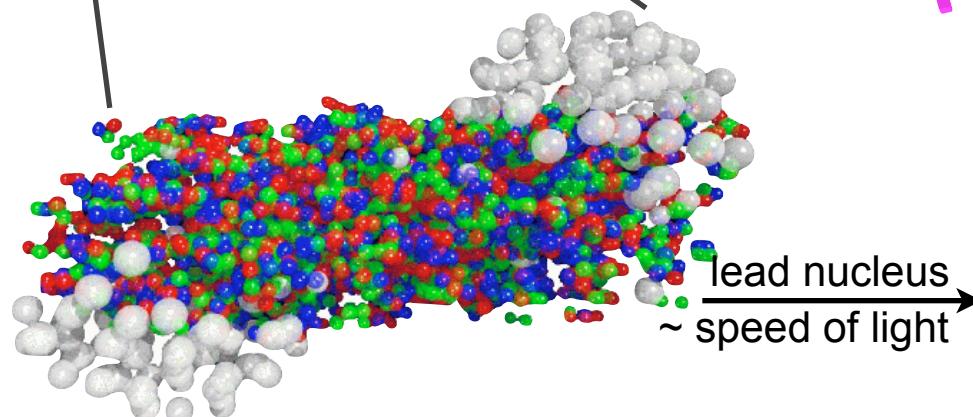
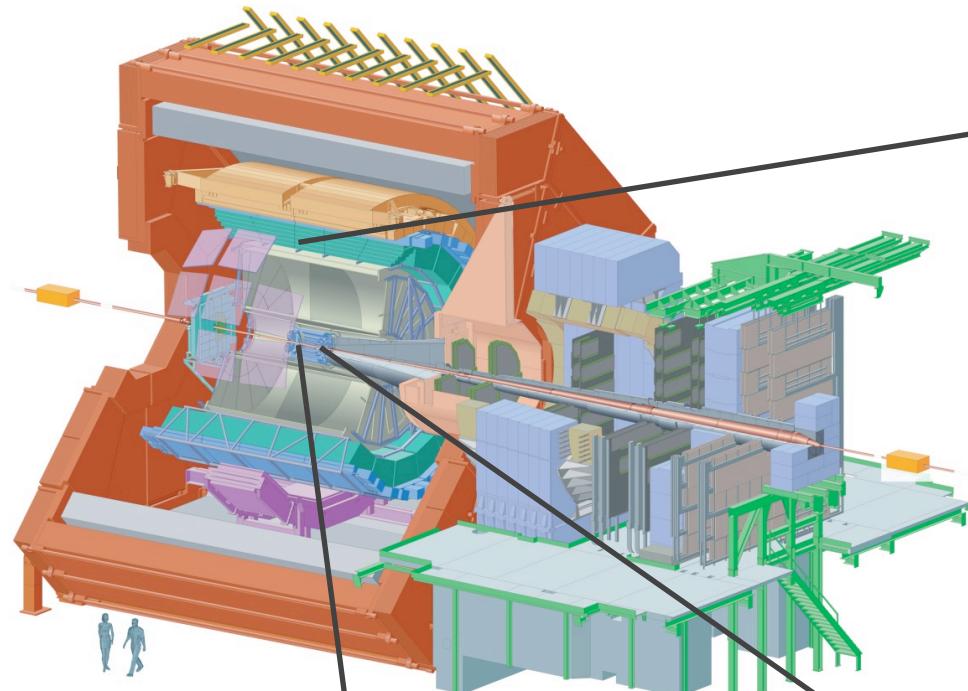
The Experiment ALICE

ALICE

- Research on Quark-Gluon-Plasma
- Many detectors covering a wide momentum range & PID
- Designed for high multiplicity events in Pb-Pb collisions

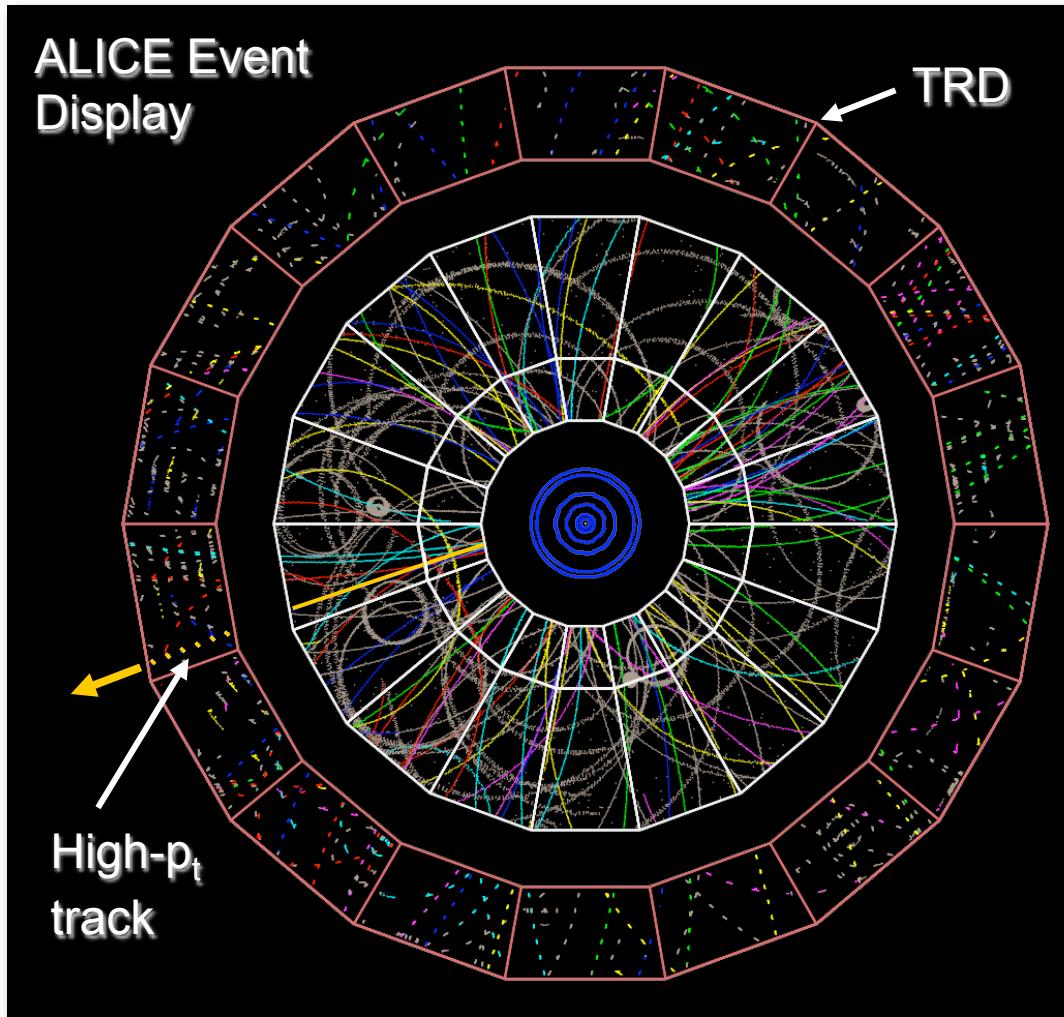


ALICE & TRD



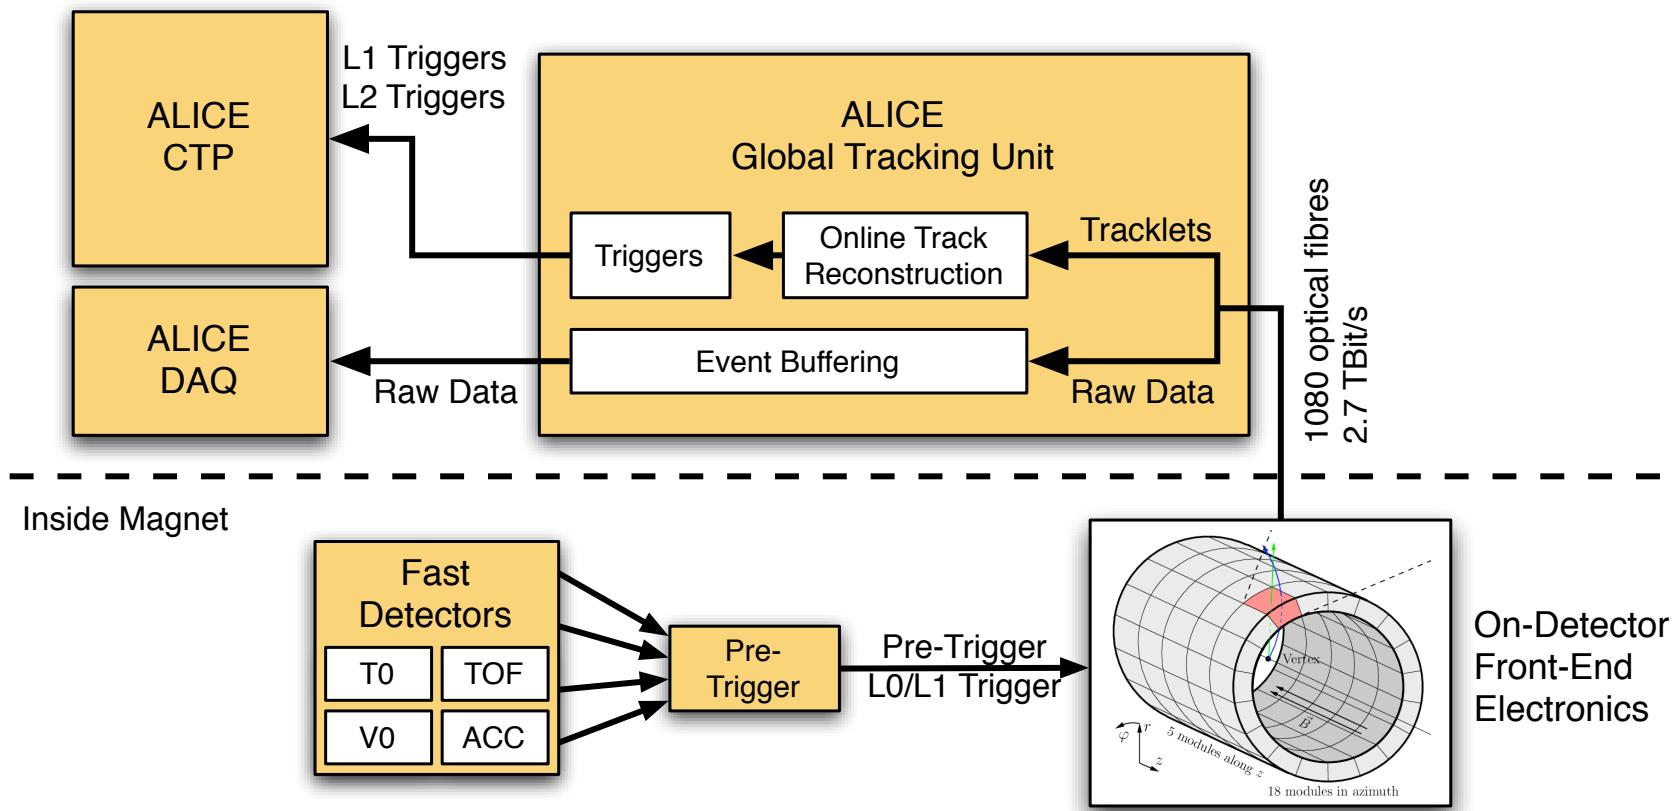
540 drift chambers in
6 layers, 90 stacks, 18
super-modules, $|\eta| \leq 0.9$

Task of the TRD



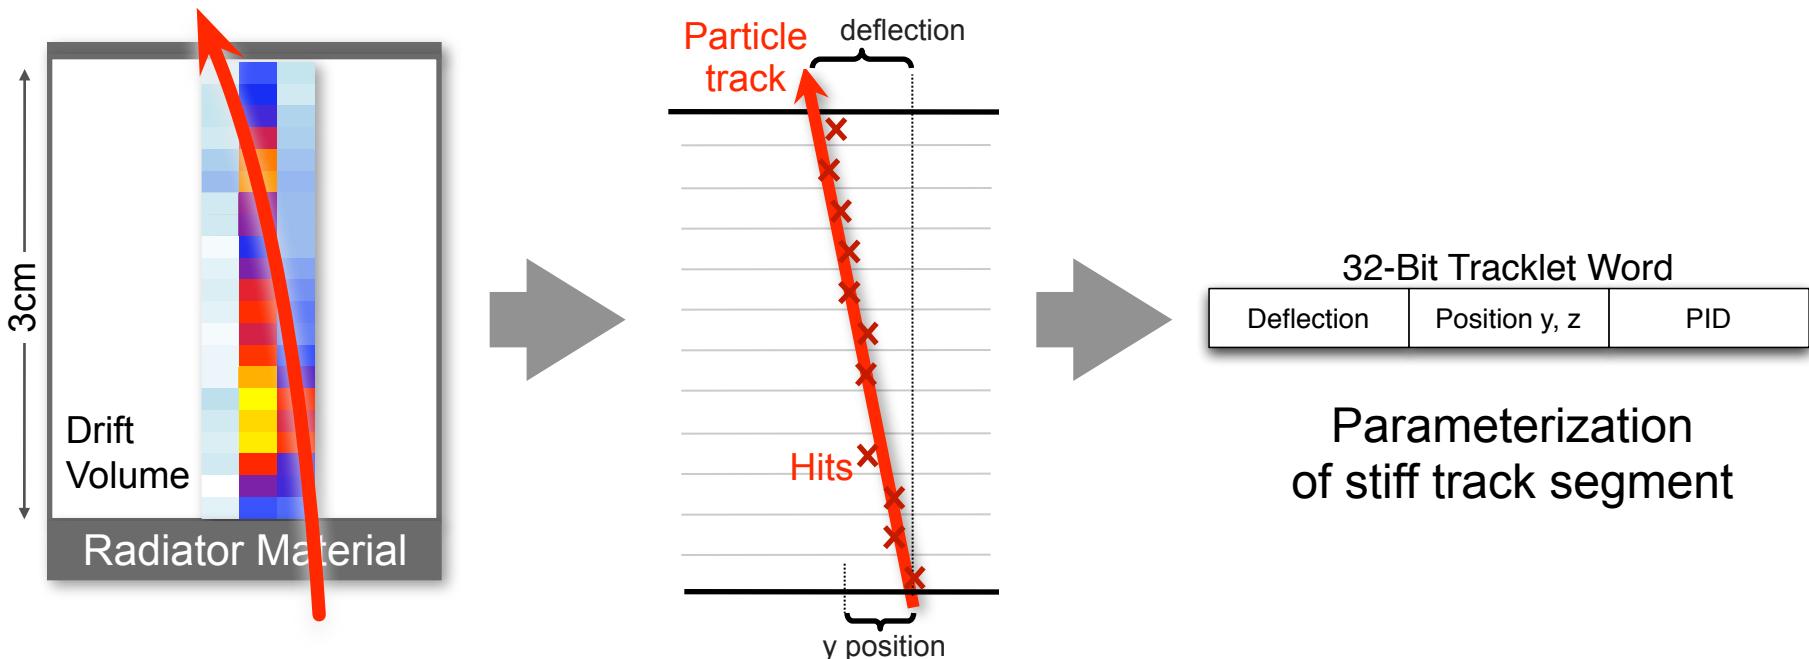
- High multiplicities:
up to 8,000 charged tracks in acceptance
- Fast trigger detector:
L1 trigger after 6.2 μ s
- Barrel tracking
detector: raw data

The TRD Data Chain



- On-Detector Front-End Electronics: 65,564 ASICs
- Global Tracking Unit: 109 FPGAs

On-Detector Data Processing

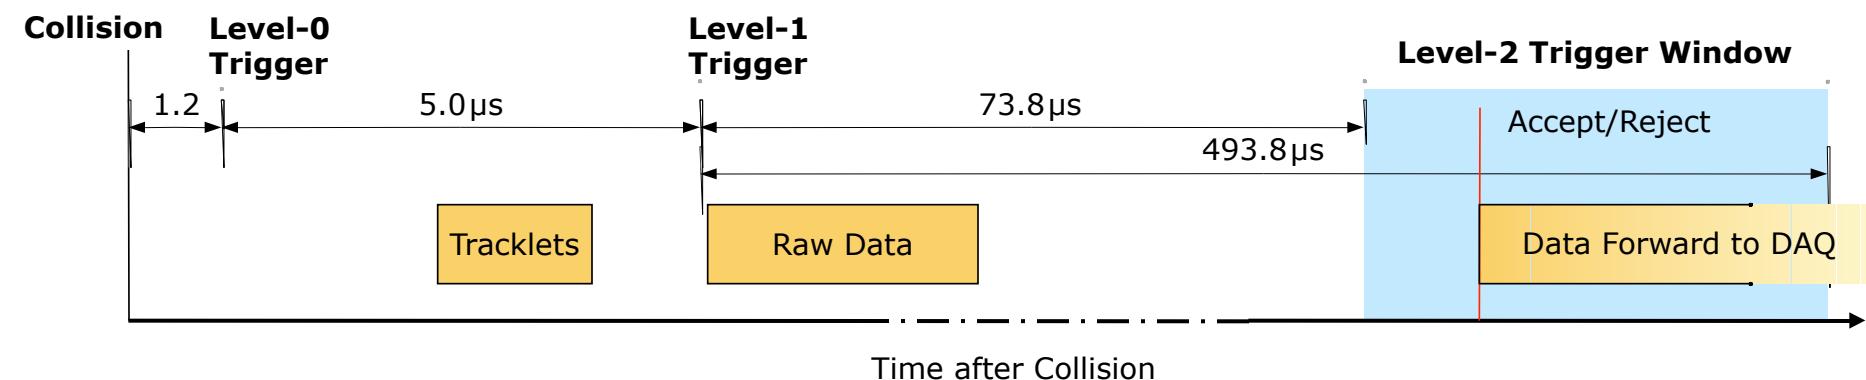
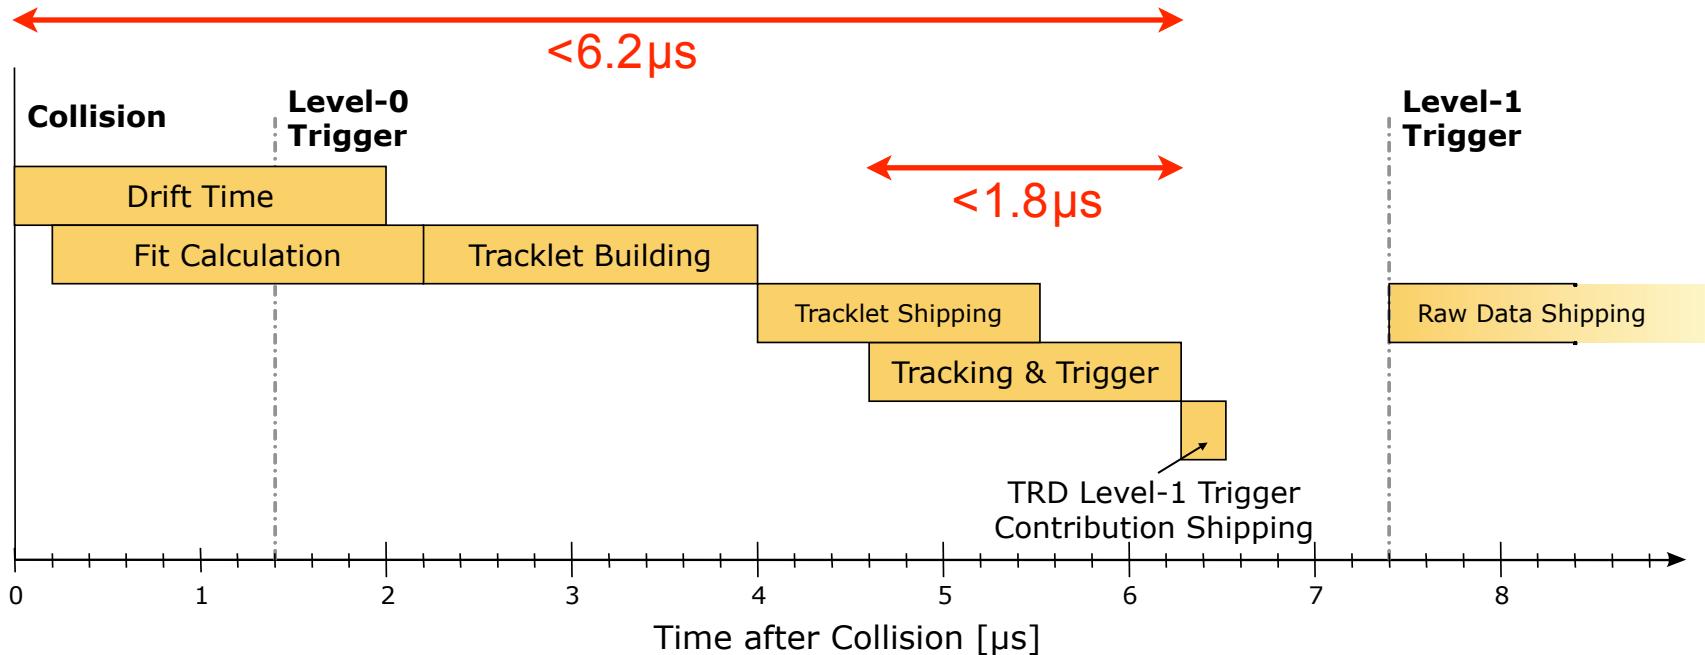


- 540 drift chambers,
6 stacked radially,
18 sectors in azimuth
- 1.4 million analog
channels
- 10 MHz sampling rate

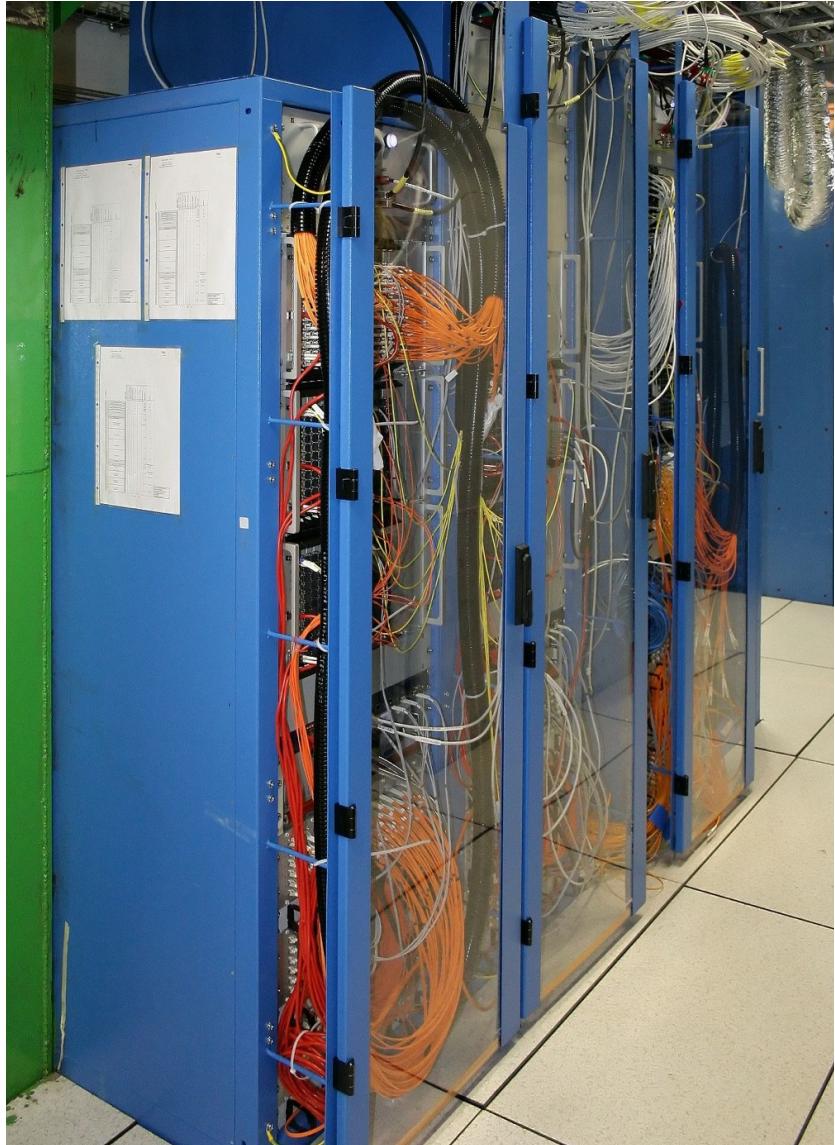
- 65,564 Multi-Chip modules,
262,256 custom CPUs
- Massively parallel calculations:
hit detection, straight line fit,
PID information
- tracklets available 4.5µs
after collision

- Up to 20,000 tracklet words,
32-Bit wide
- Transmission out of magnet
via 1080 optical fibres
operating at 2.5 GBit/s
- 2.1 TBit/s total bandwidth

Tight Timing Requirements!

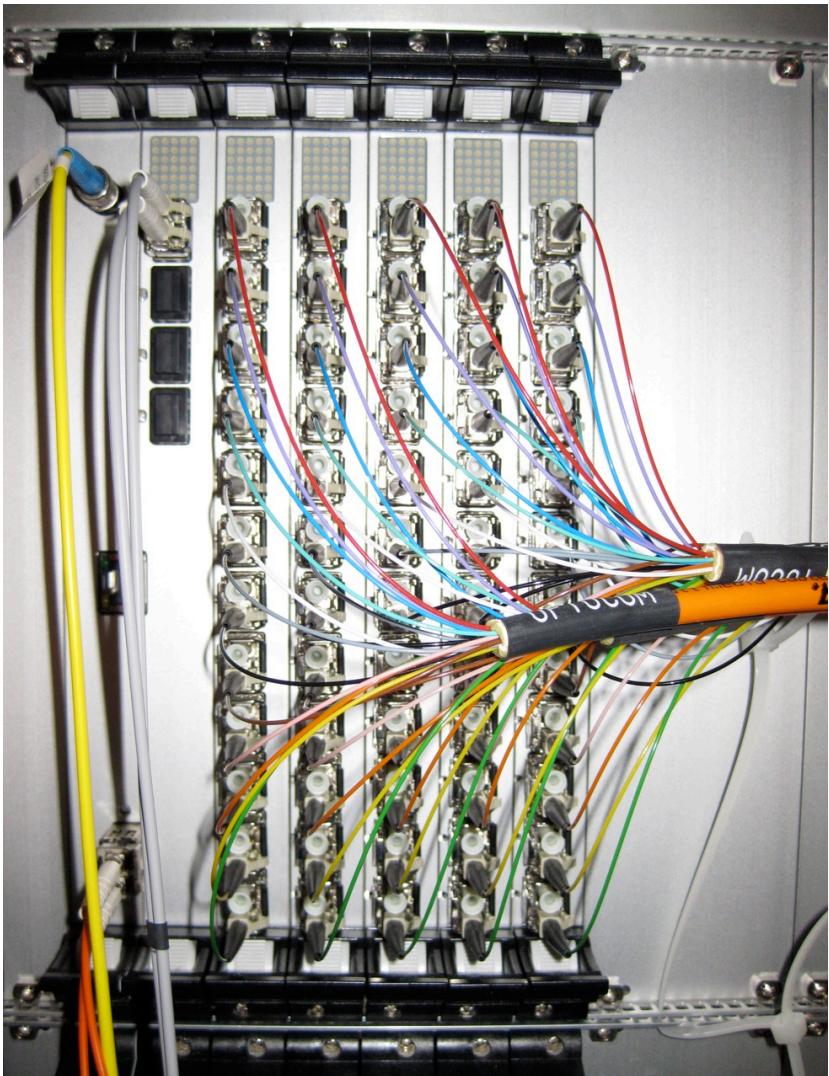


Global Tracking Unit

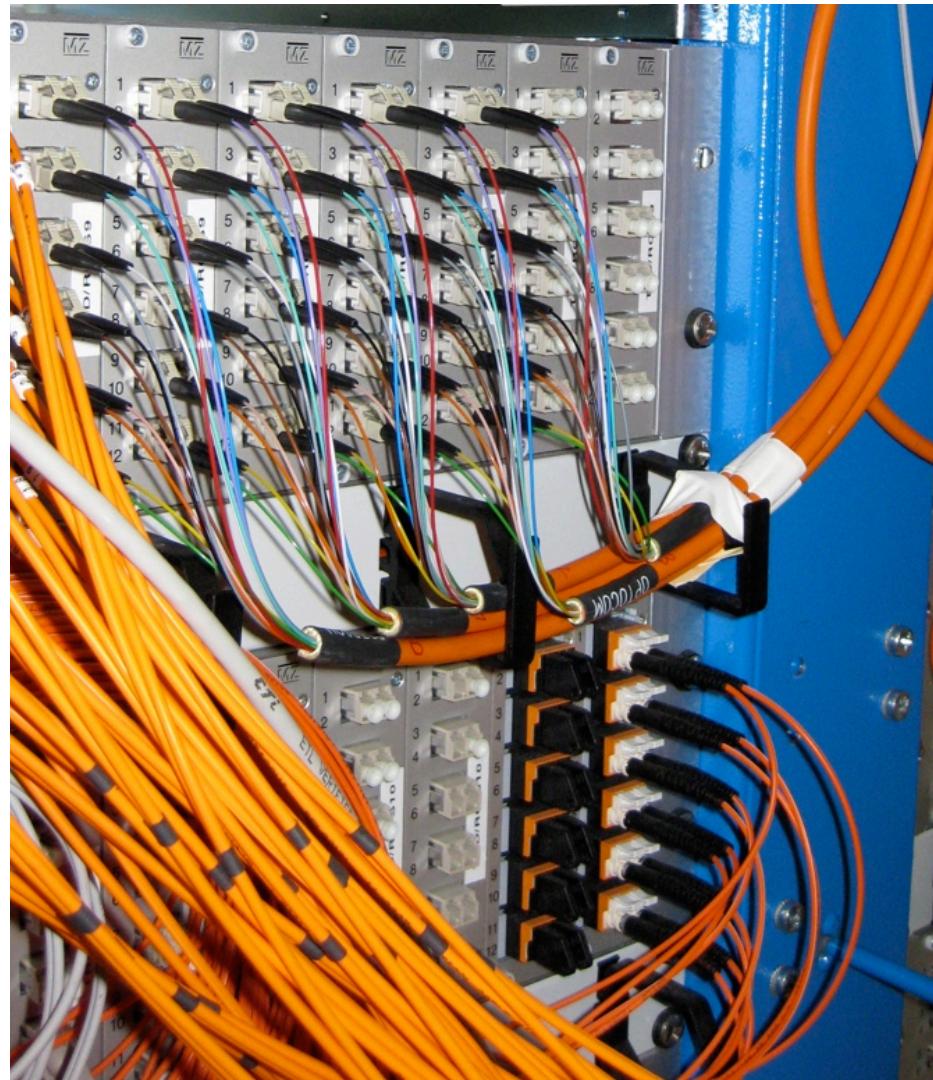


- Fast L1 trigger after $6.2\mu\text{s}$
 - Detection & reconstruction high-momentum tracks
 - Calculation of momenta
 - Various trigger schemes: di-lepton decays (J/ψ , Υ), jets, ultra-peripheral collisions, cosmics
- Raw data buffering
 - Multi-Event buffering & forwarding to data acquisition system
 - Support interlaced triggers, multi-event buffering, dynamic sizes
- 109 boards with large FPGAs in three 19" racks outside of magnet

Global Tracking Unit



GTU segment for one TRD supermodule

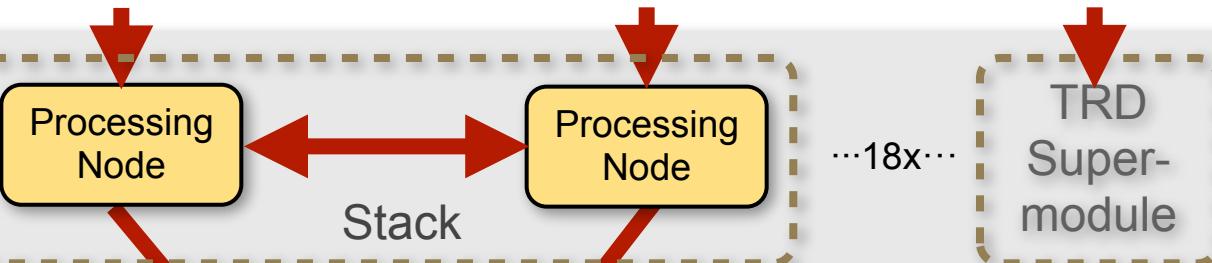


Patch panel with 60 fibres for one TRD supermodule

3-Tier Architecture

Tracklets from Front-End (1,080 optical fibres, 2.1 TBit/s)

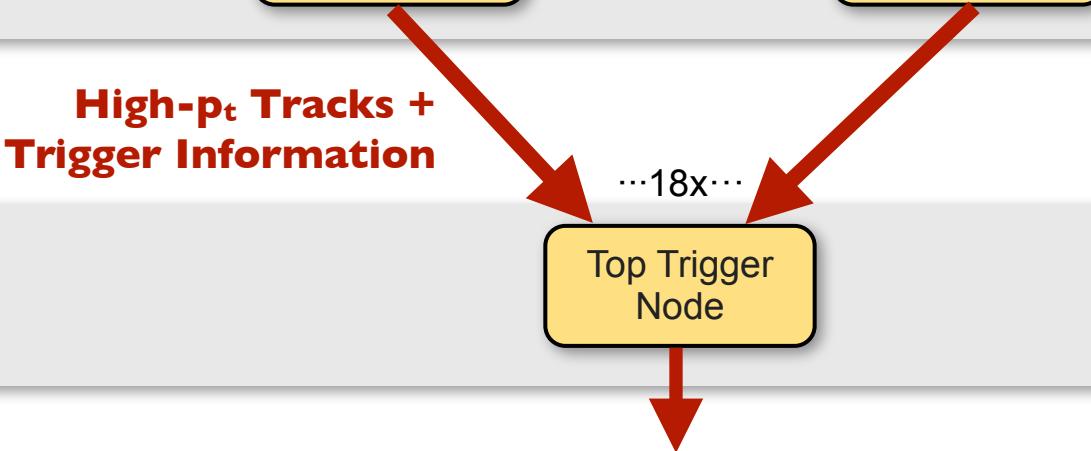
all track segments of one detector stack



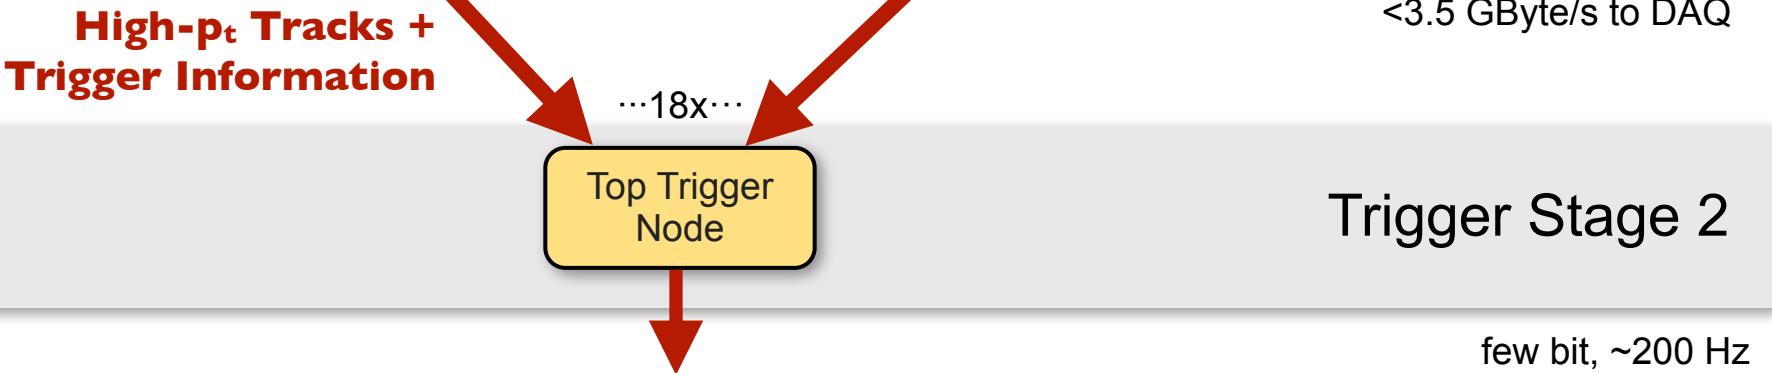
Online Track Reconstruction
90 Boards



2.9 GByte/s into each Node



Trigger Stage 1
Raw Data → DAQ



<3.5 GByte/s to DAQ

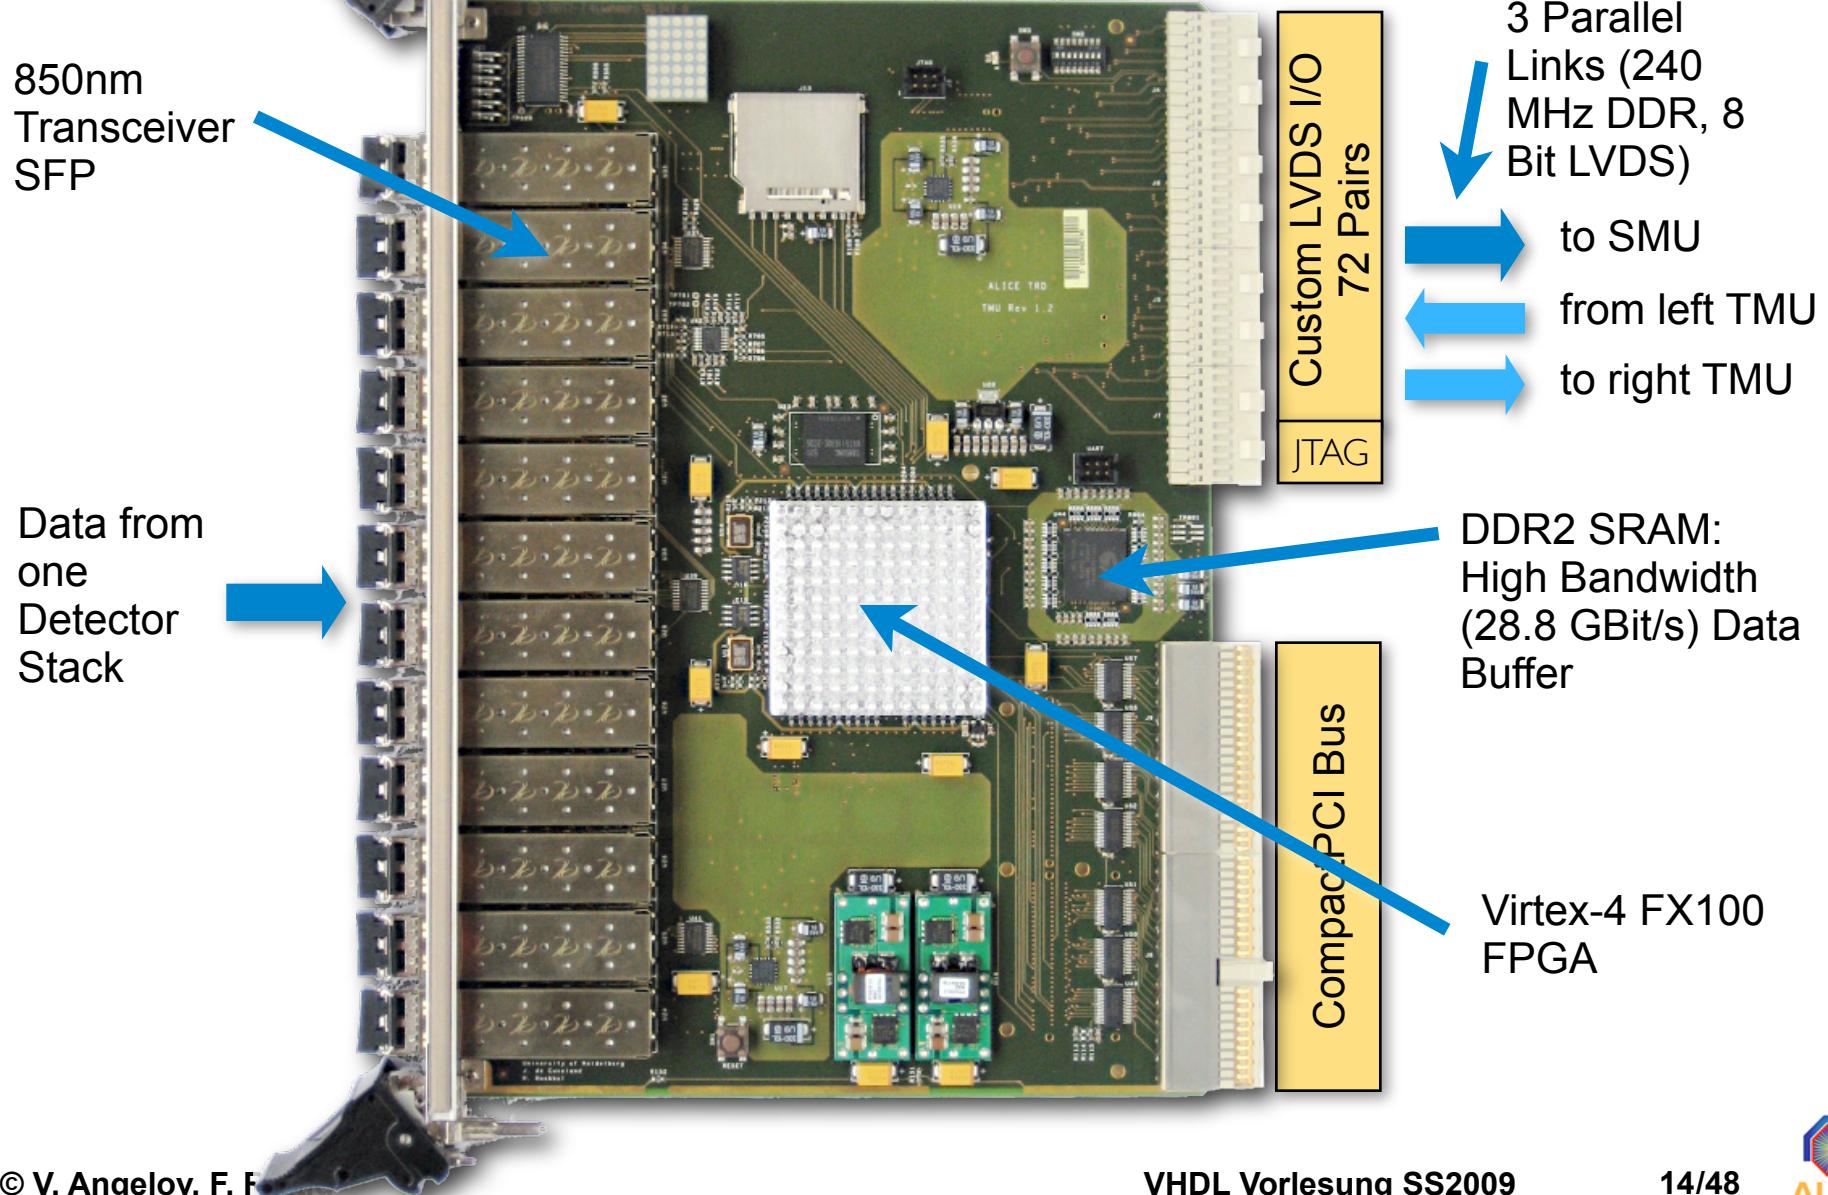
Trigger Stage 2

few bit, ~200 Hz

Processing Node

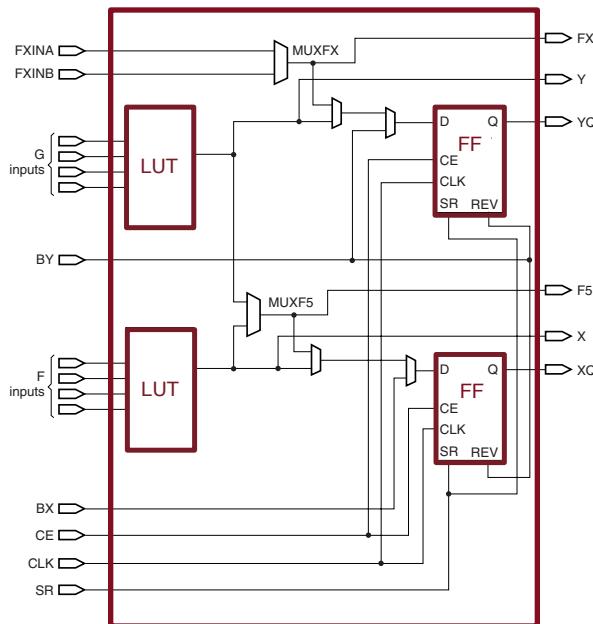
- Inputs: tracklets & raw data via 12 optical data streams at 2.5 GBit/s each
→ 2.9 GByte/s per node,
261 GByte/s total
- Data push architecture
→ capture at full bandwidth of 2.1 TBit/s
- Tasks:
 - Online Track Reconstruction
 - Multi-Event Buffering

Processing Node

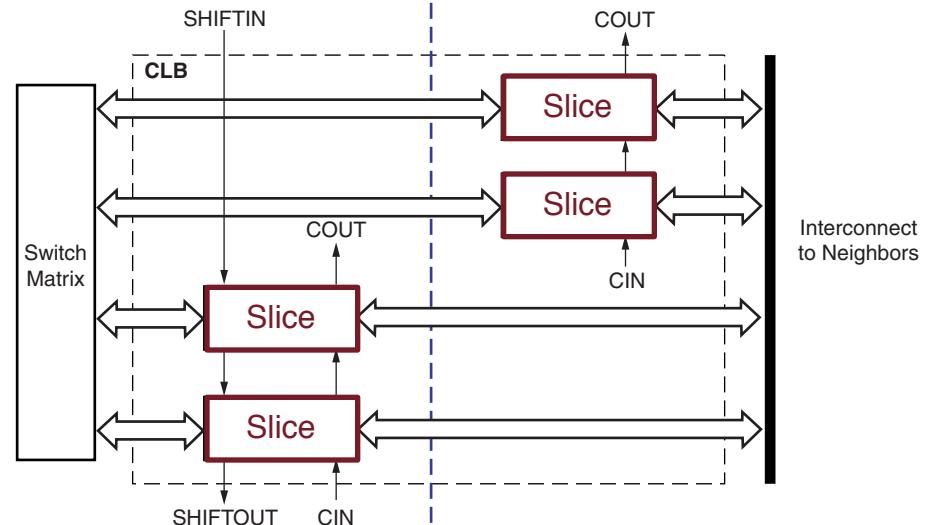


Virtex-4 FX Family

Device	Configurable Logic Blocks (CLBs) ⁽¹⁾				XtremeDSP Slices ⁽²⁾	Block RAM		DCMs	PMCDs	PowerPC Processor Blocks	Ethernet MACs	RocketIO Transceiver Blocks	Total I/O Banks	Max User I/O
	Array ⁽³⁾ Row x Col	Logic Cells	Slices	Max Distributed RAM (Kb)		18 Kb Blocks	Max Block RAM (Kb)							
XC4VFX20	64 x 36	19,224	8,544	134	32	68	1,224	4	0	1	2	8	9	320
XC4VFX40	96 x 52	41,904	18,624	291	48	144	2,592	8	4	2	4	12	11	448
XC4VFX60	128 x 52	56,880	25,280	395	128	232	4,176	12	8	2	4	16	13	576
XC4VFX100	160 x 68	94,896	42,176	659	160	376	6,768	12	8	2	4	20	15	768



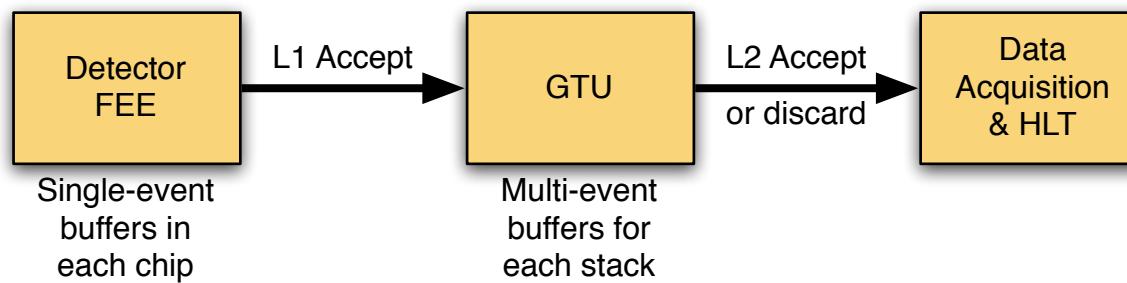
Virtex-4 Slice



Virtex-4 Configurable Logic Block (CLB)

Multi-Event Buffering

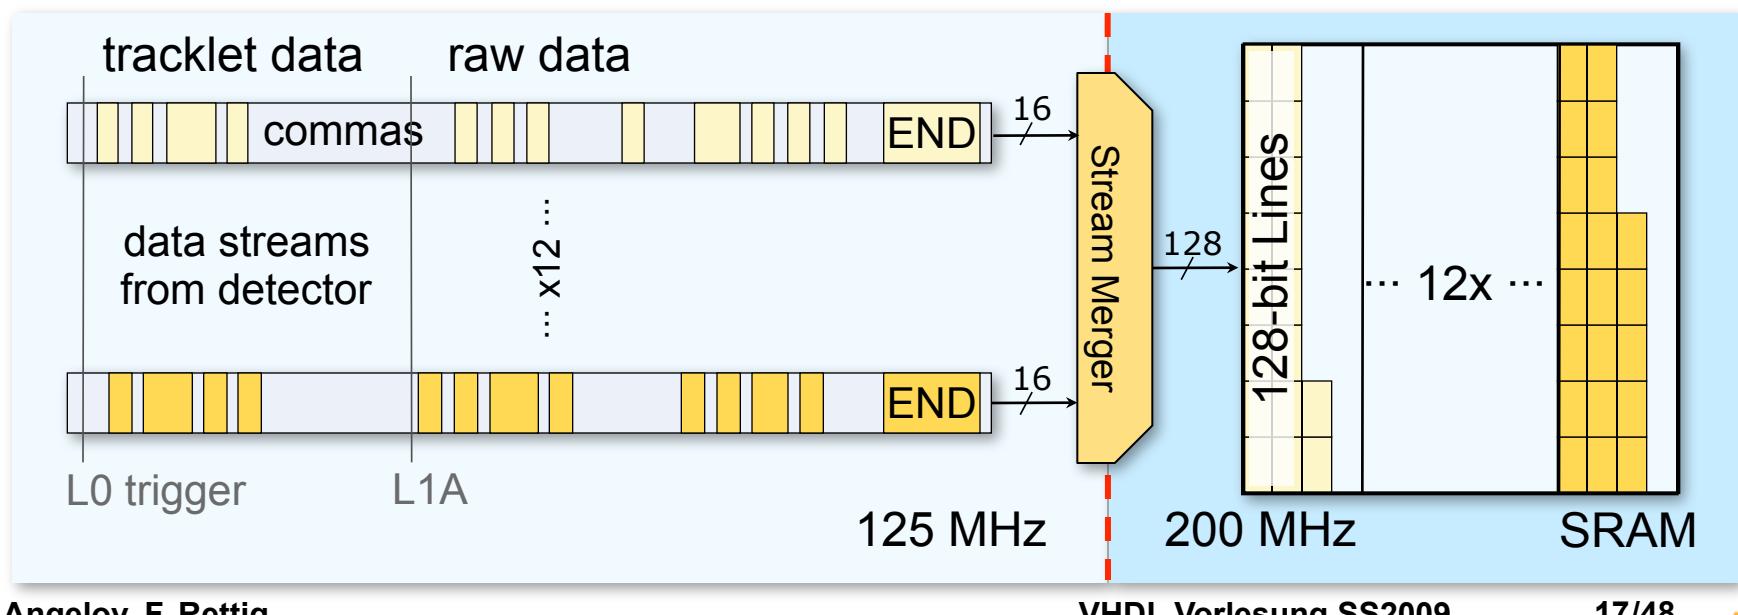
- Allows for significant reduction of detector dead time due to:
 - Interleaved 3-level trigger sequences
 - Decoupling of front-end electronics operation from data transmission to data acquisition, 2-stage readout



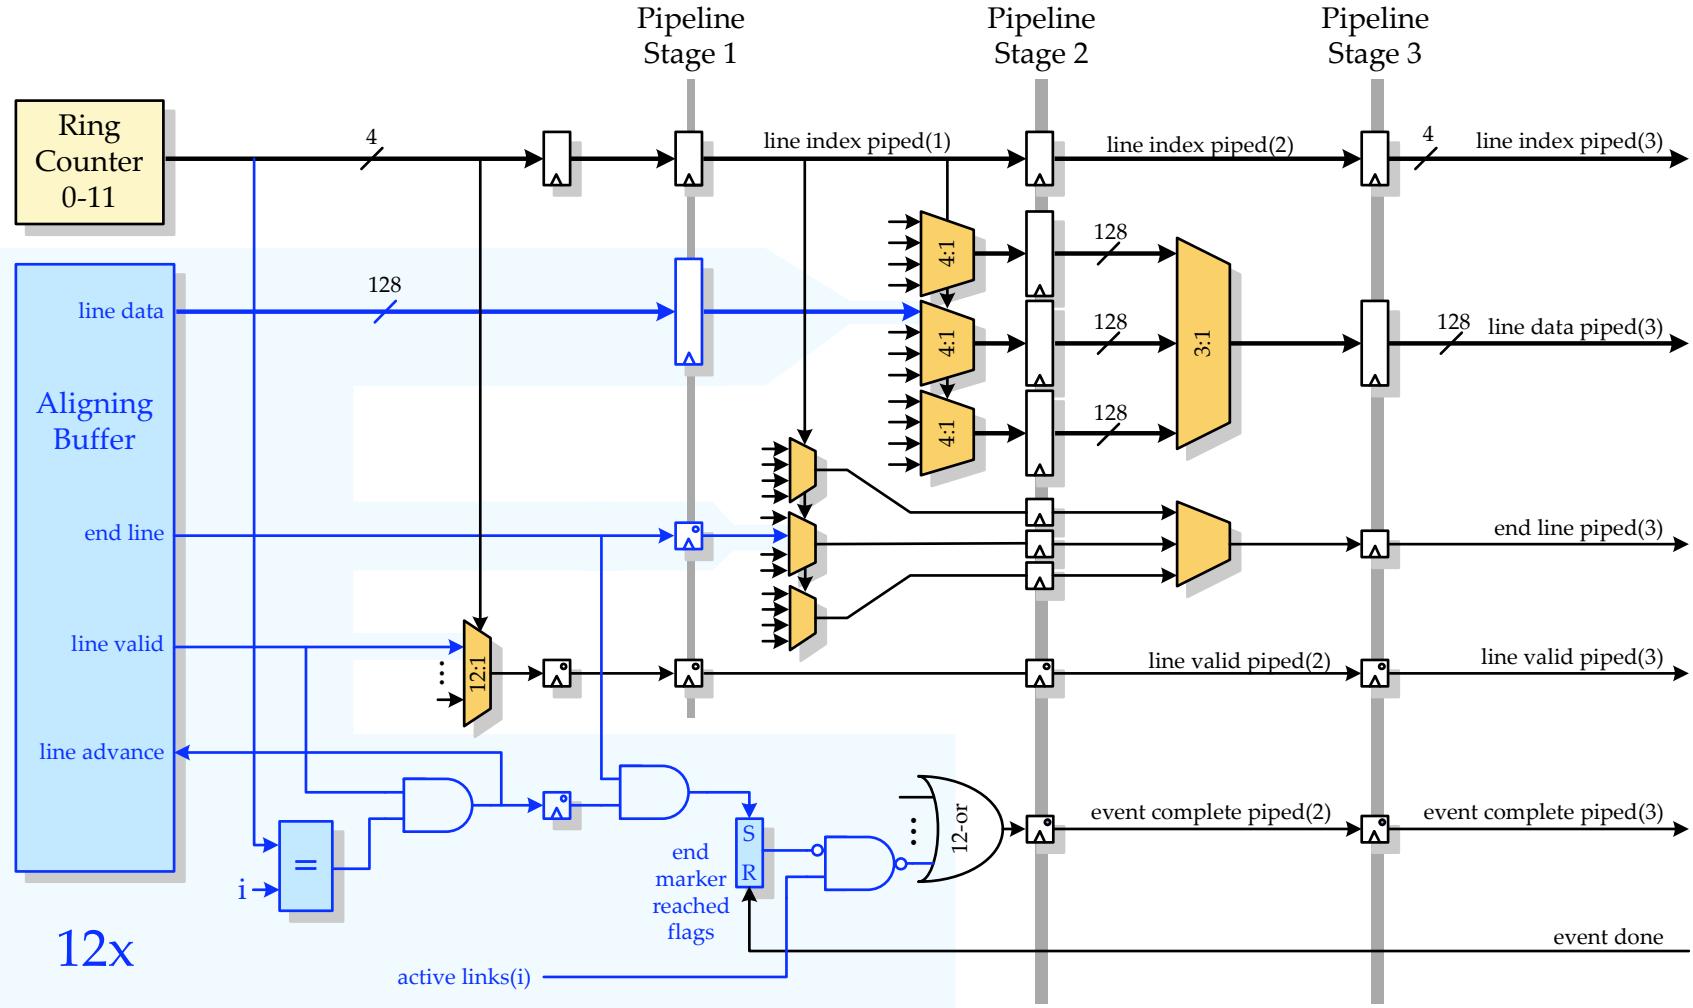
- Dynamic buffer allocation for strongly varying event sizes
- Buffers: 4-MBit SRAMs, 64-bit 200 MHz DDR interface
- 12 independent 128 bit wide data streams at 200 MHz

Multi-Event Buffering II

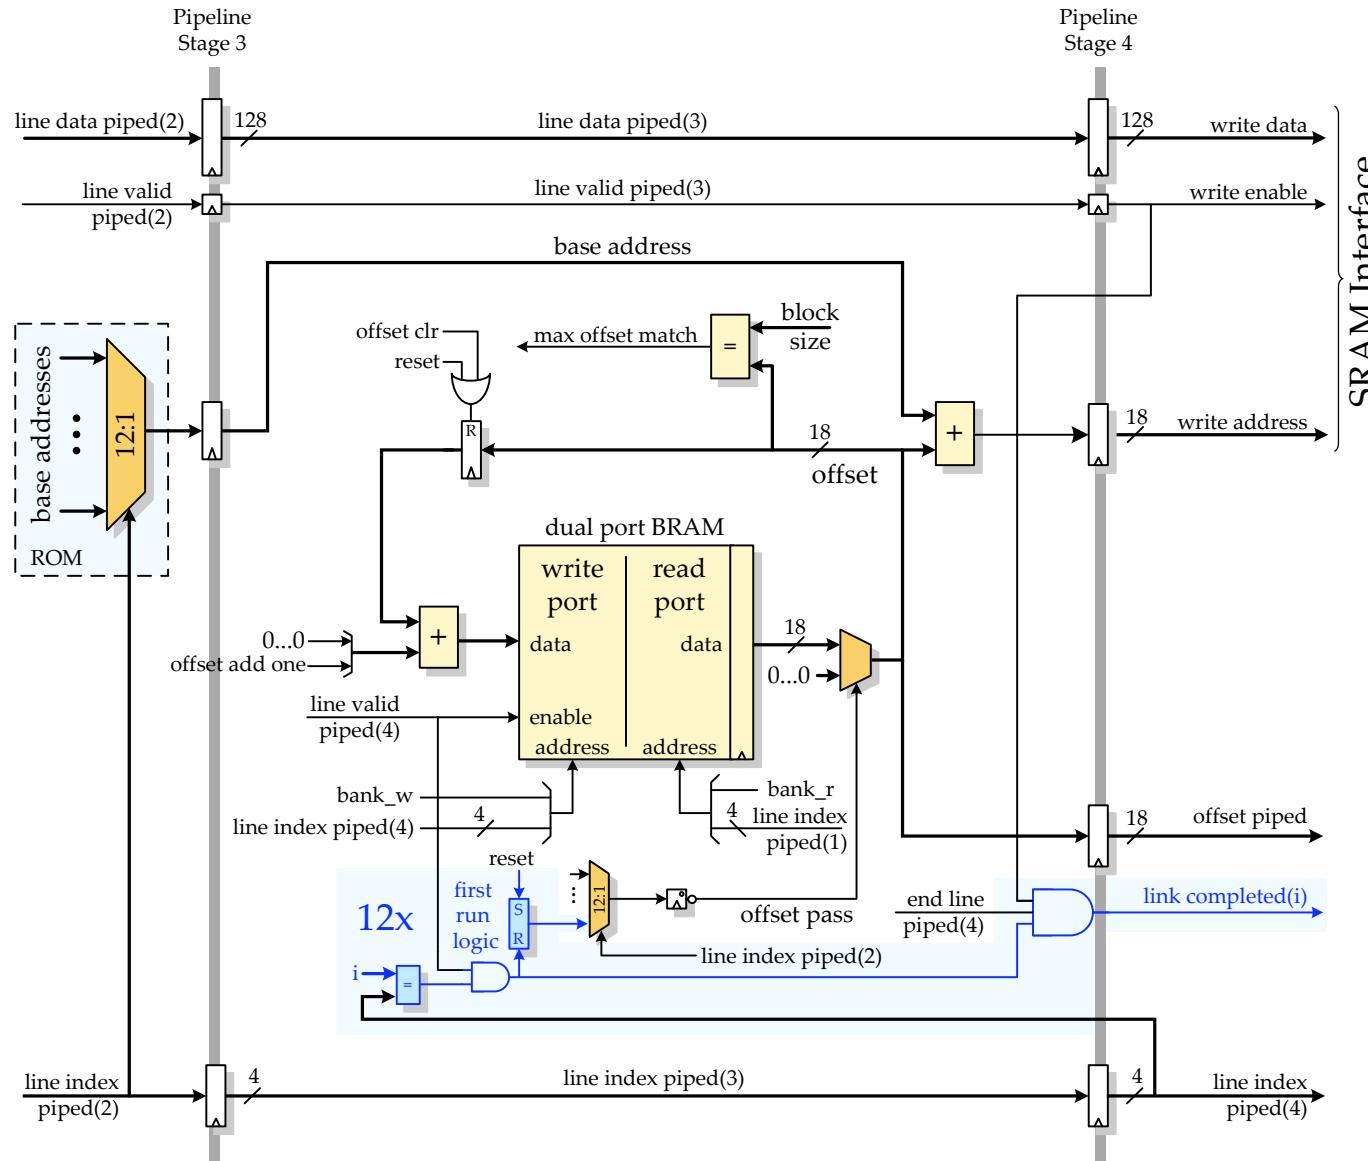
- 12 independent data streams via 2.5 GBit/s links, in fabric as 16-bit streams at 125 MHz (net 1.94 GBit/s)
- De-randomizing/gap elimination, merging to single dense 128-bit 200 MHz data stream to SRAM (>94% of all clock cycles, 23.3 GBit/s)
- Allocation of separate memory regions for each link/ event (12 independent ring buffers, 2 write+1 read pointers)



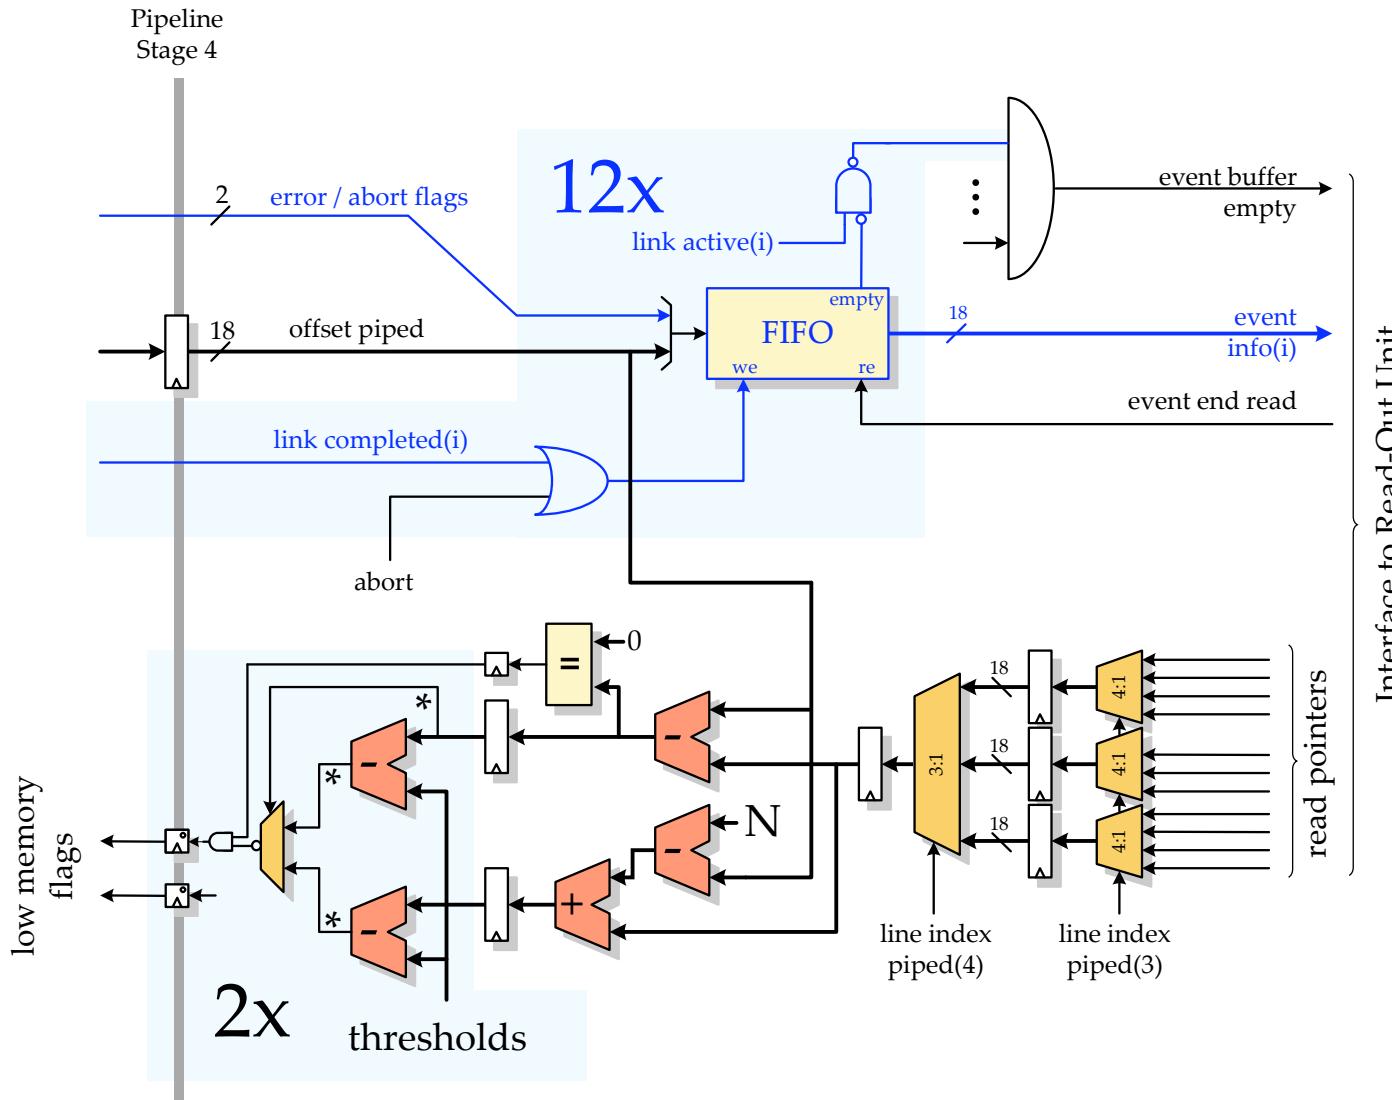
Event Buffering Pipeline I



Event Buffering Pipeline II

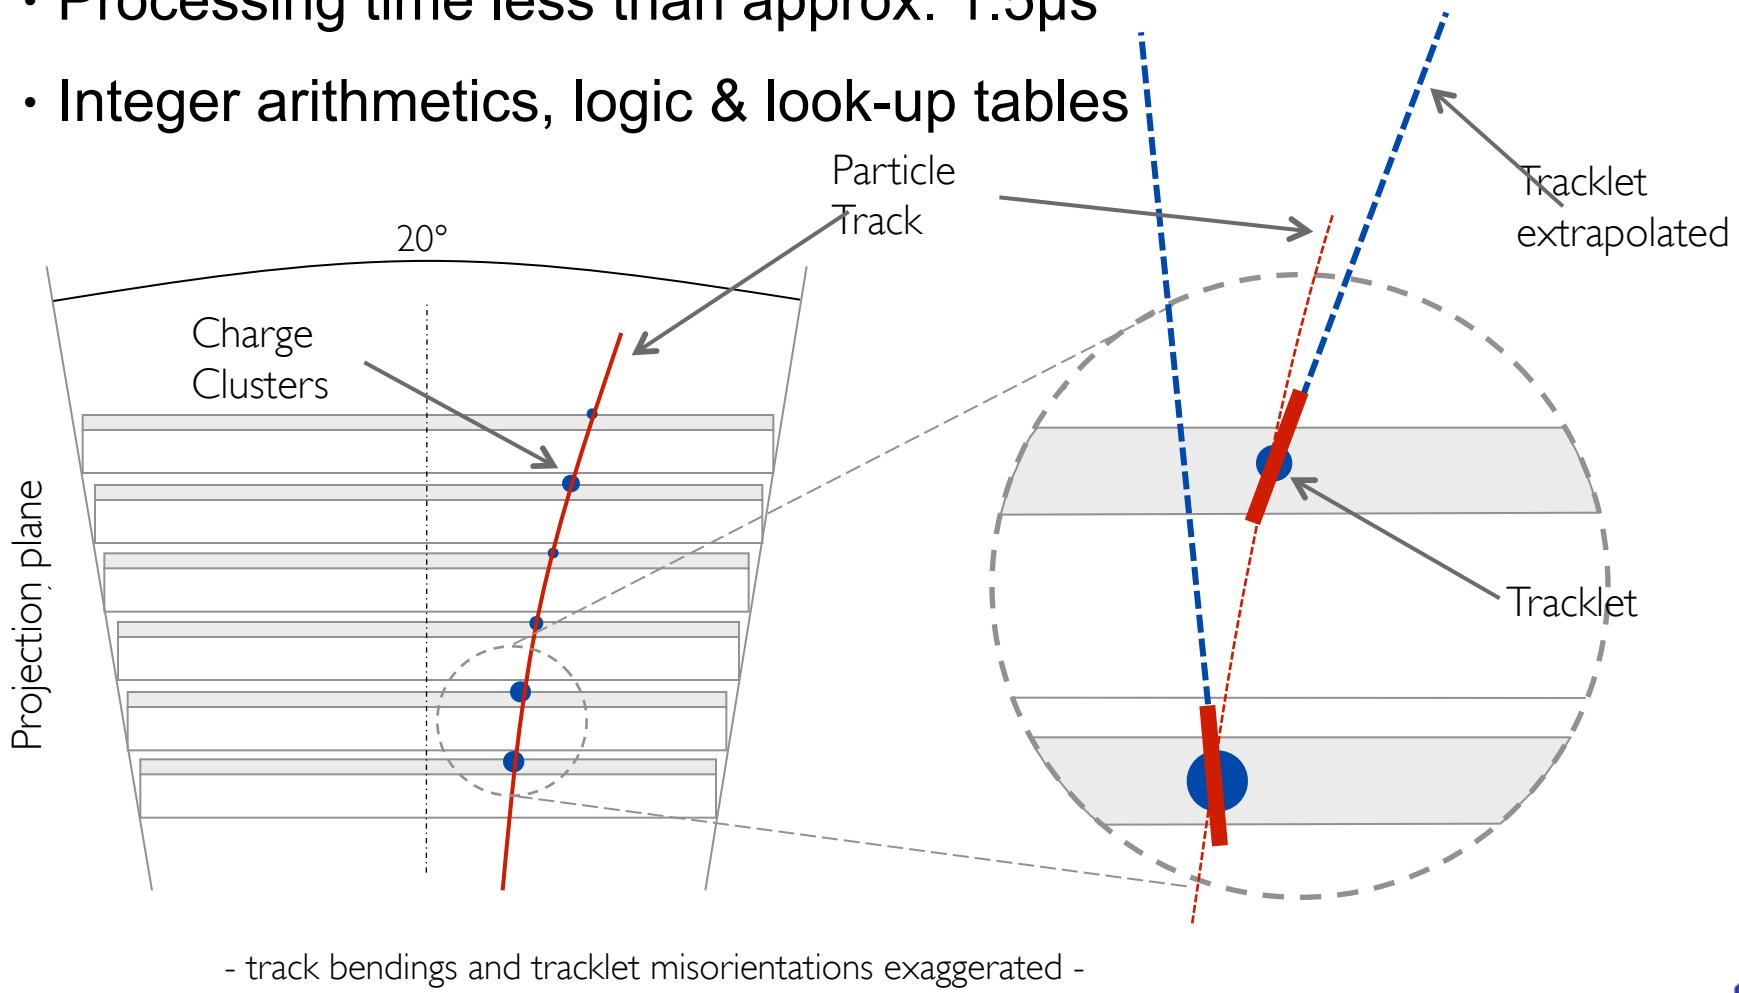


Event Buffering Pipeline III



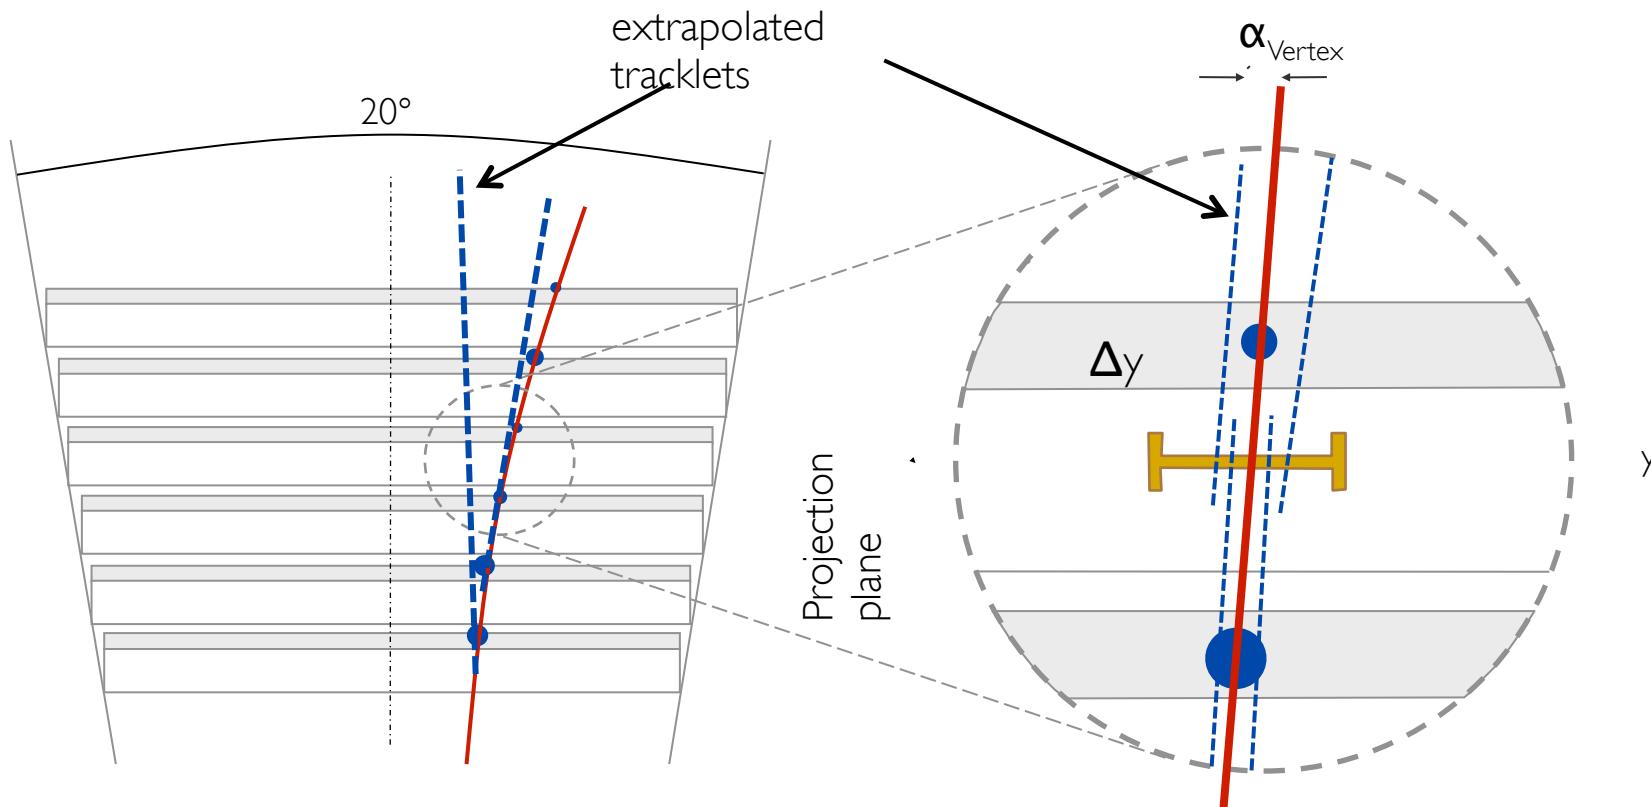
Global Track Matching

- 3D track matching: find tracklets belonging to one track
- Processing time less than approx. $1.5\mu\text{s}$
- Integer arithmetics, logic & look-up tables



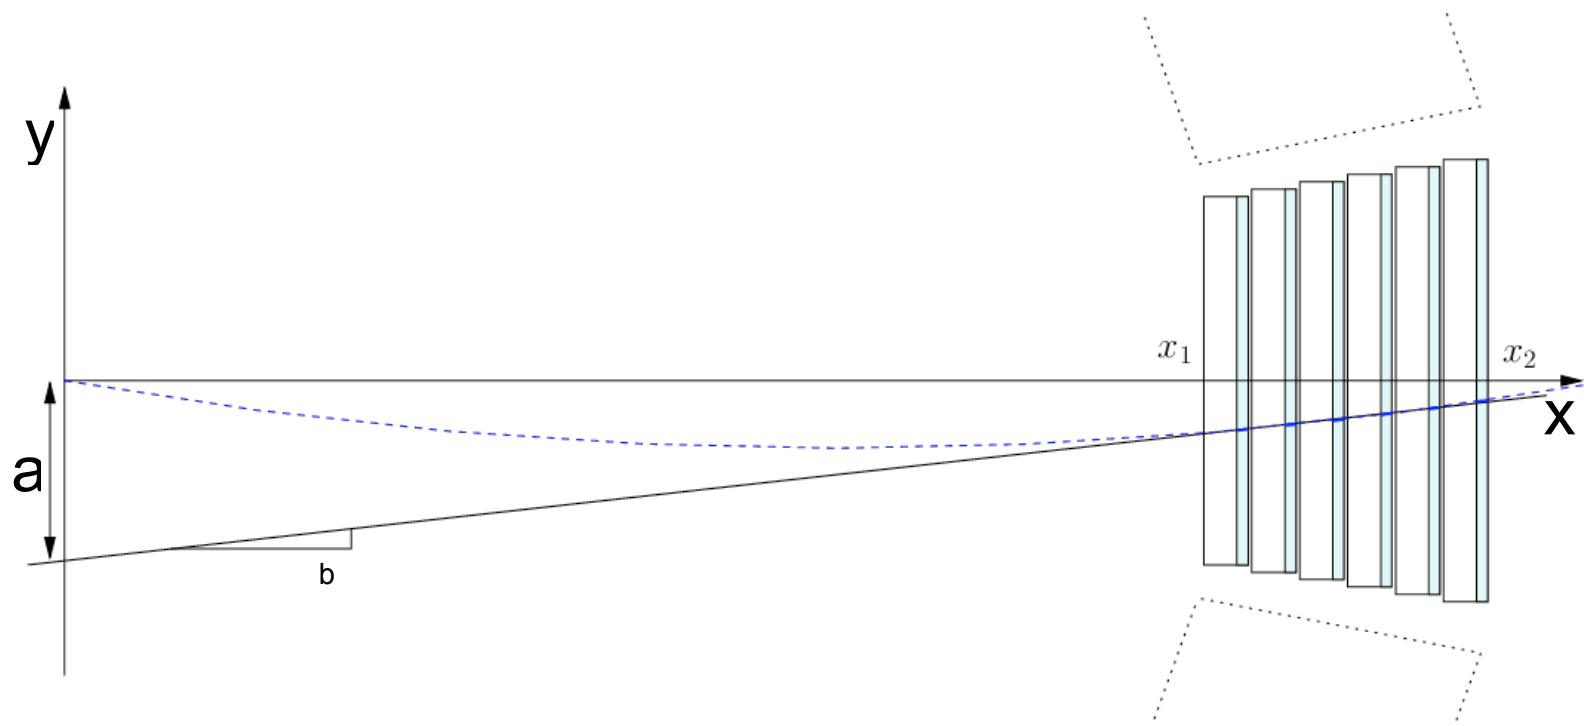
Global Track Matching II

- Projection of tracklets to virtual transverse planes
- Intelligent sliding window algorithm: Δy , $\Delta\alpha_{\text{Vertex}}$, Δz
- Massively parallel hardware implementation

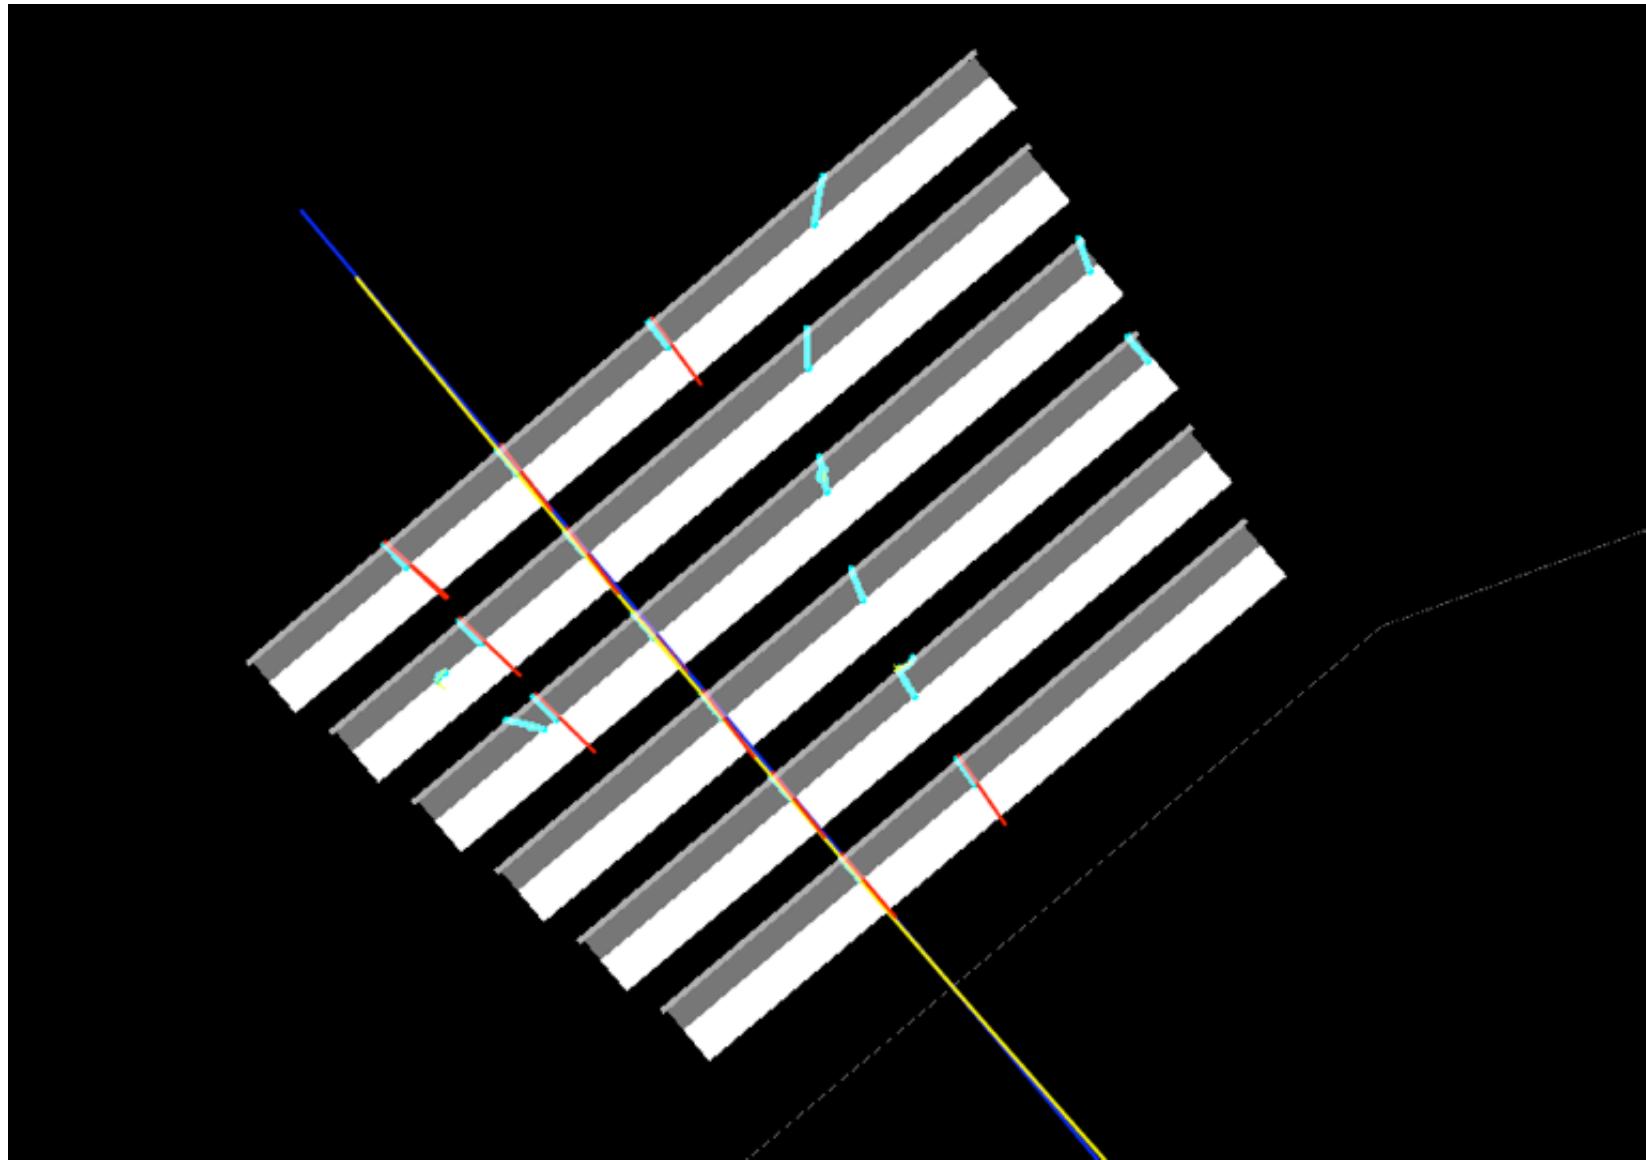


Momentum Reconstruction

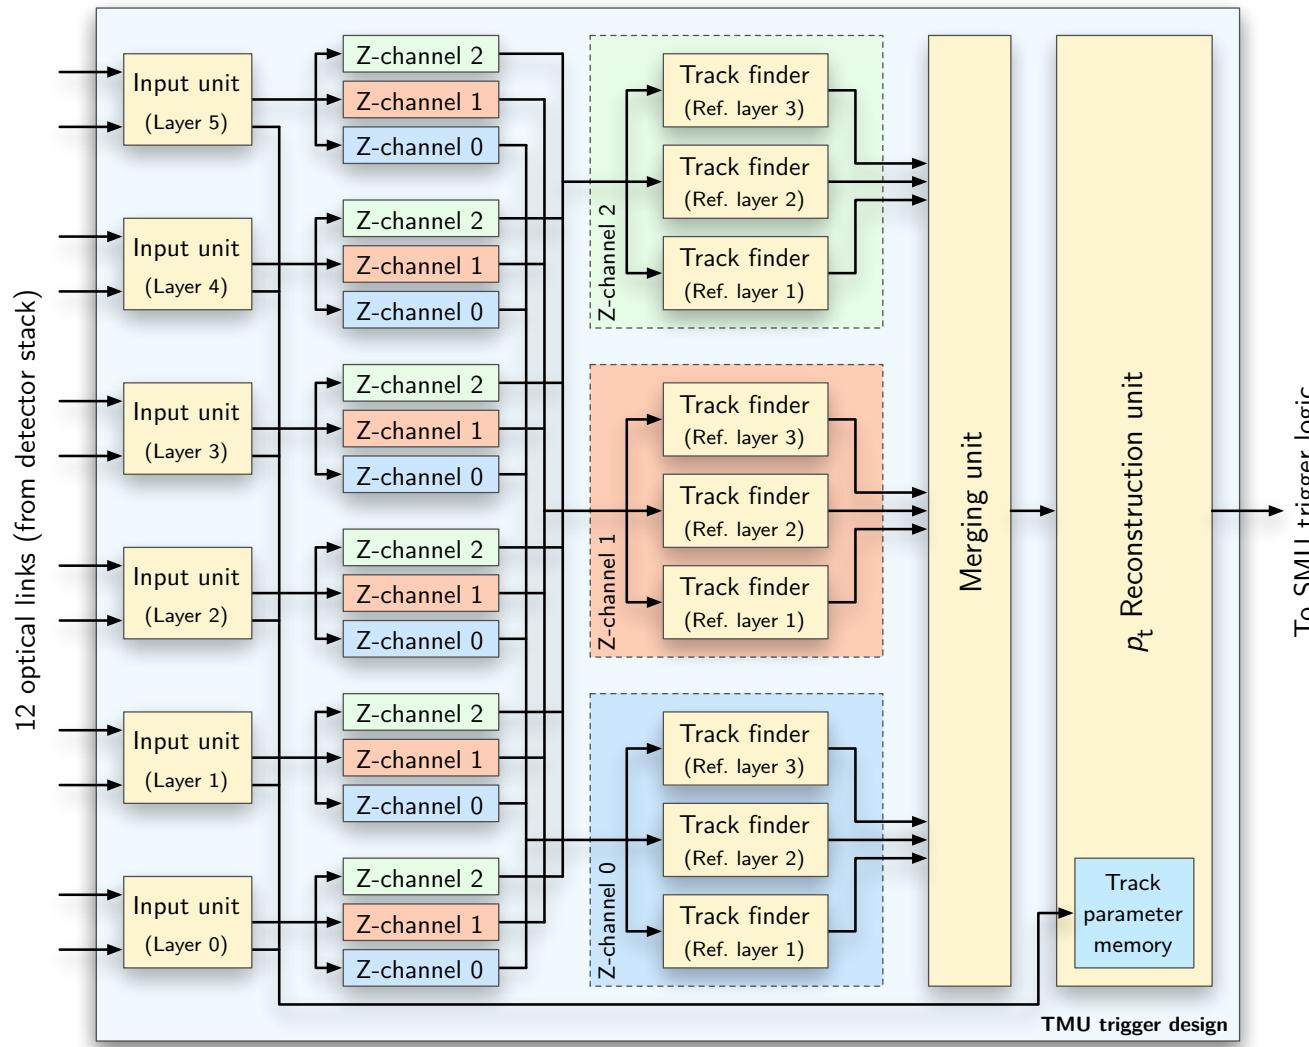
- Assumption: particle origin is at collision point
- Estimation of p_t from line parameter a : $p_t = \frac{const}{a}$
- Fast cut condition for trigger: $const \leq p_{t,min} \cdot a$



An Example...

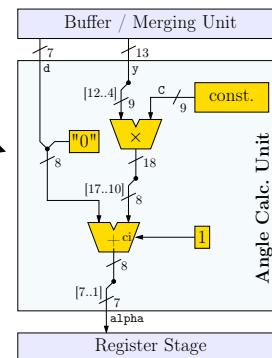
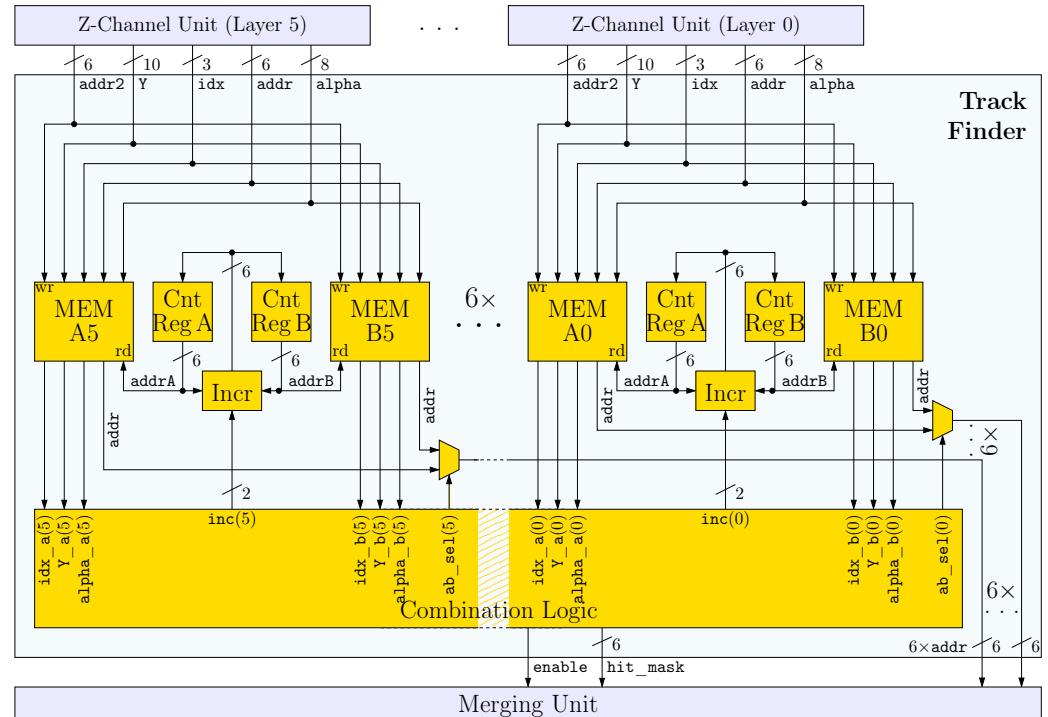
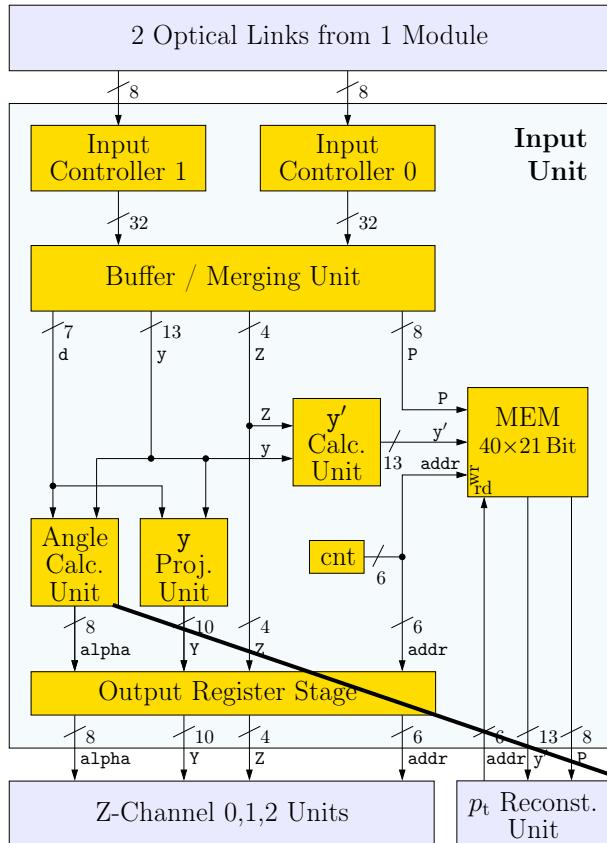


Online Track Matching I

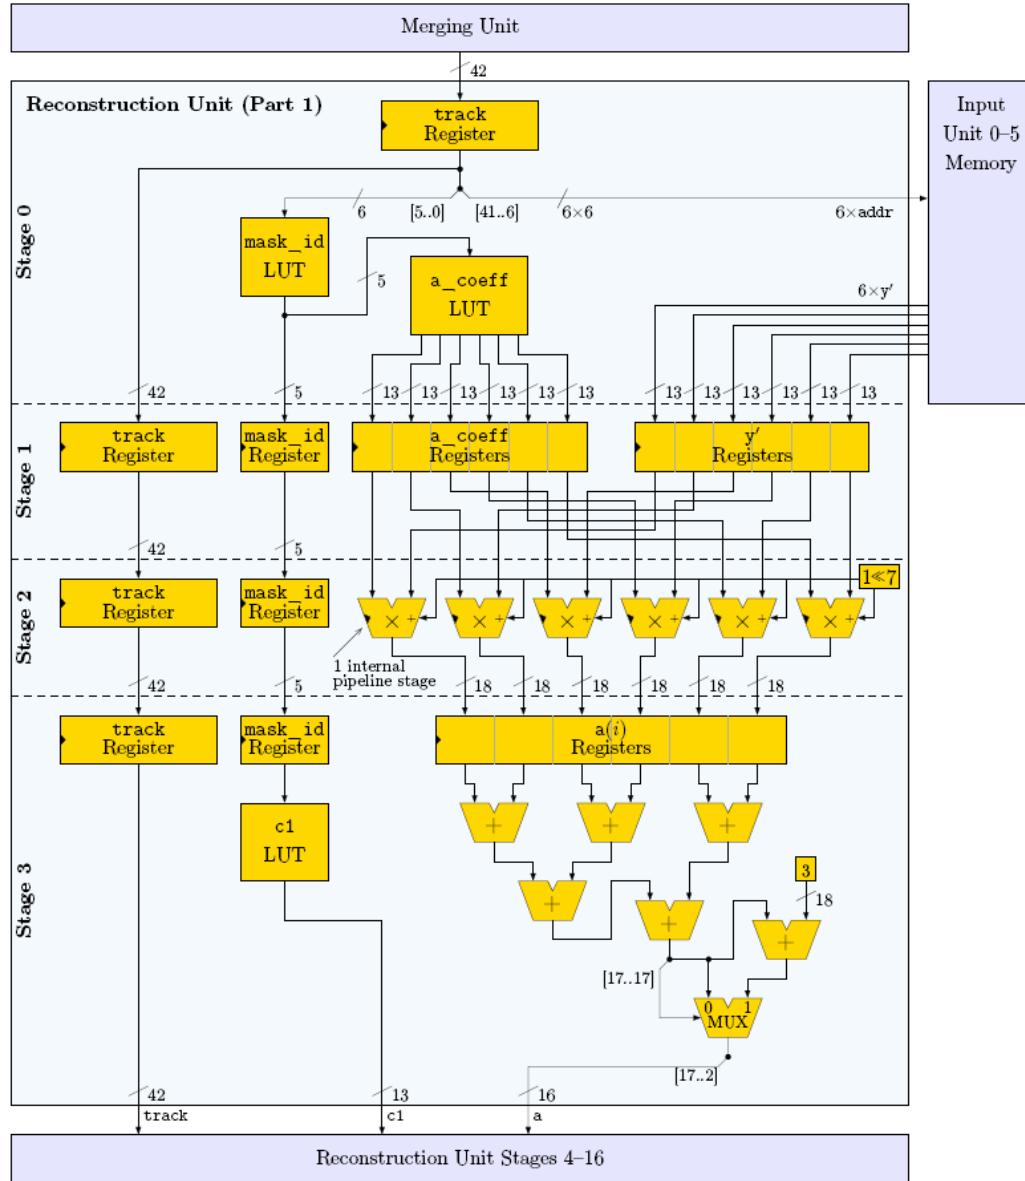


- 18 matching units running in parallel
- Up to 240 track segments/event
- Fully pipelined, data push architecture
- Fast integer arithmetic and pre-computed look-up tables used
- High precision p_t reconstruction $\Delta p_t/p_t < 2\%$
- 60 MHz clock

Input & Track Finder Unit

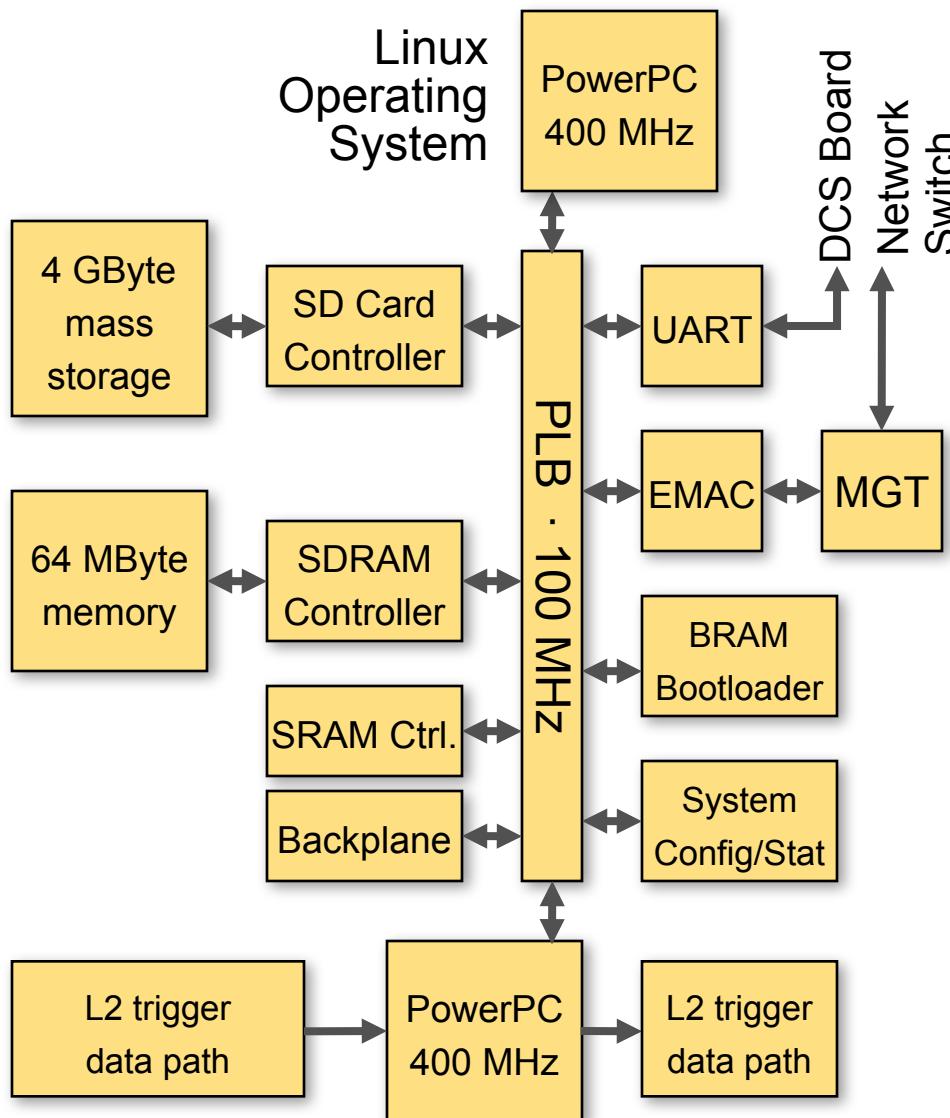


Reconstruction Unit I



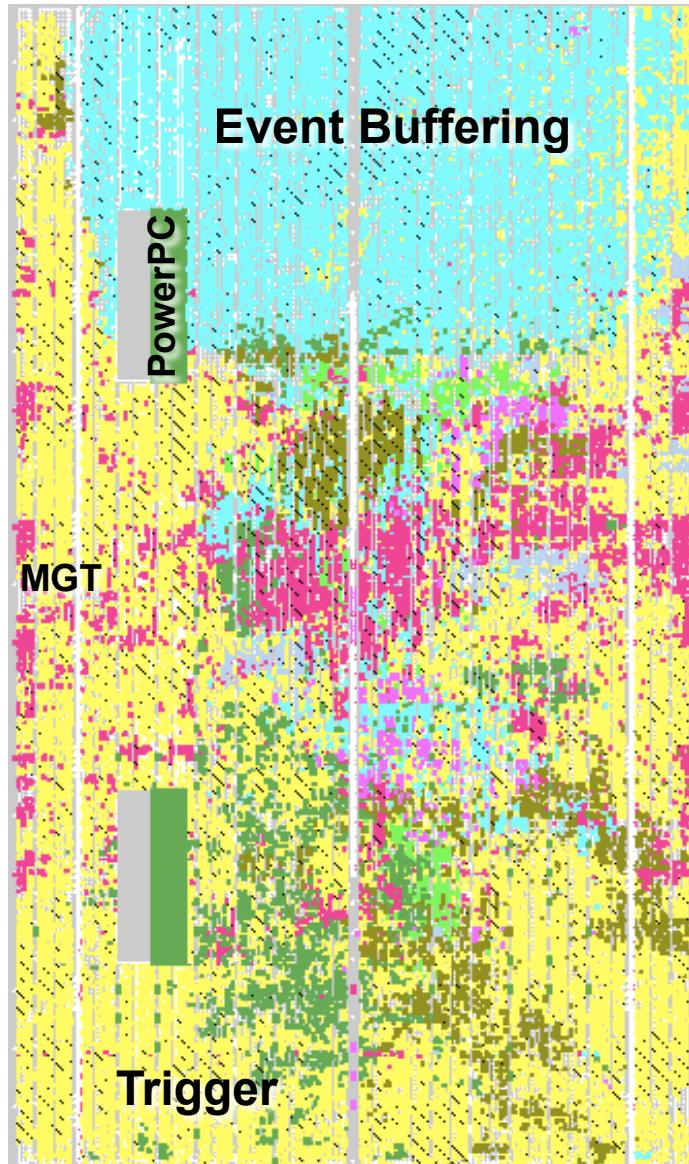
- fully pipelined data push architecture
- optimized for low latency
- High precision pt reconstruction
 $\Delta\text{pt}/\text{pt} < 2.5\%$
- Uses addition, multiplication and pre-computed look-up tables
- 60 MHz clock

Embedded PowerPC System



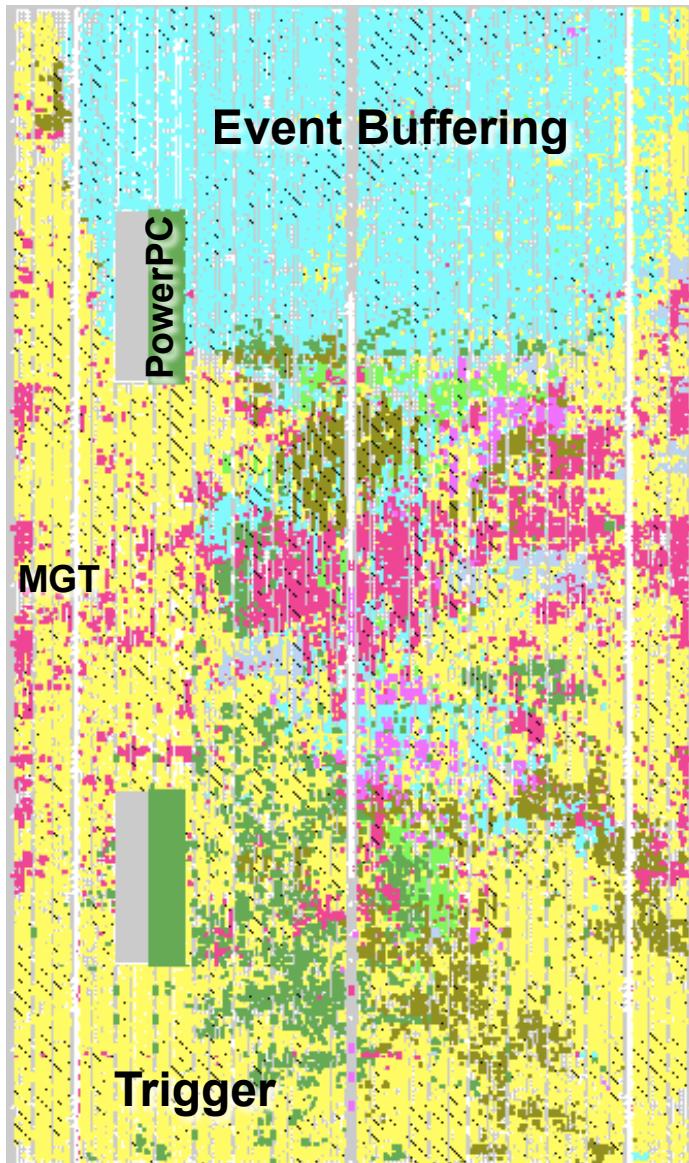
- Bus components
 - DDR2 SDRAM controller
 - UART, Gigabit-Ethernet
 - SD Card controller
 - SRAM controller
 - Configuration & status interface
- Two PowerPC Cores:
 - Linux Operation System: monitoring & control (PetaLinux/Monta Vista)
 - HW/SW-Codesign (planned): Level-2 trigger calculations, real-time monitoring & control

TMU Design Resource Usage



- 38,601 slices occupied (91%)
 - 45,716 logic LUTs (54%)
 - 53,500 LUTs total (63%)
 - 29,936 FFs (35%)
- 4 DCMs (33%), 1 PMCD (12%), 17 BUFGs (63%)
- 165 BRAMs (43%)
- 12 MGTs (60%), 9 DSPs (5%), 2 PowerPCs
- 345 IOB (56%)
- Gate equivalent: 11,625,408

TMU Design Resource Usage



Resource	Event Buffering	Tracking
FF	10,921	8,858
LUT	5,940	24,086
BRAM	14	78
DMEM	19 / 0	93 / 1,128

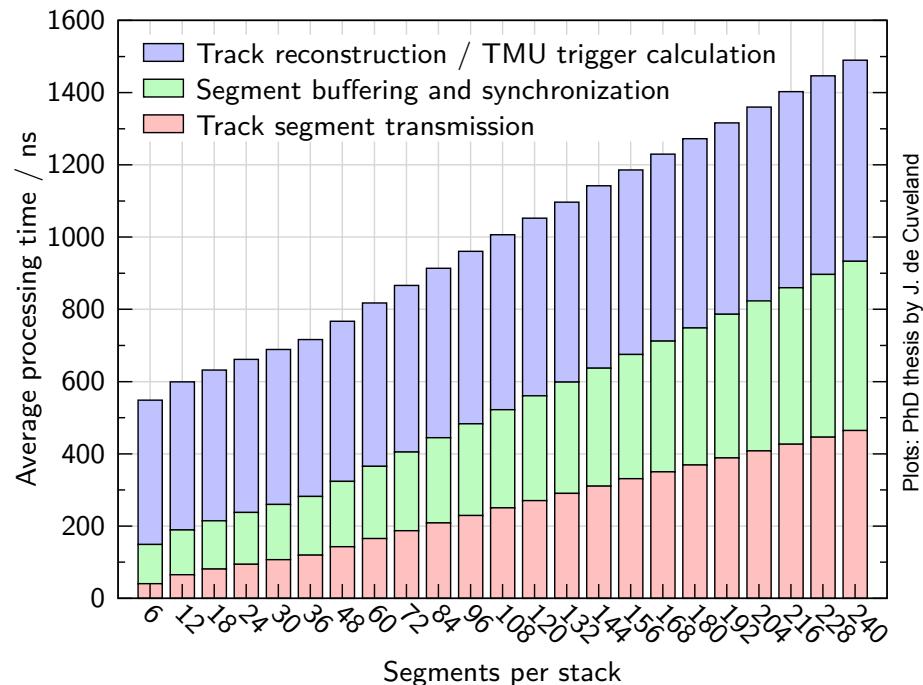
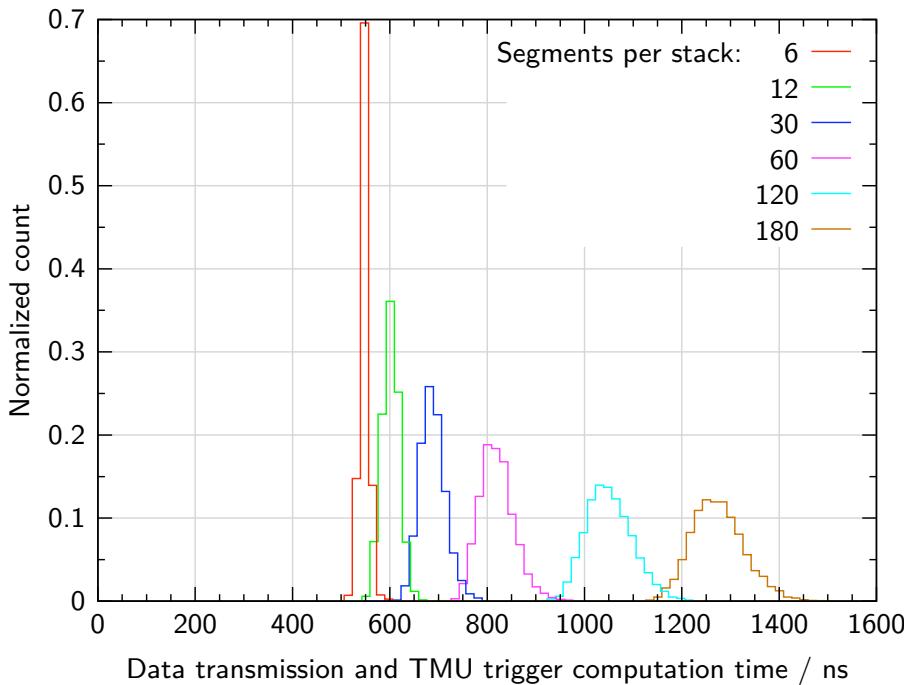
Embedded PowerPC System:
4,003 FFs, 4,068 LUTs, 69 BRAMs

VHDL Code

Design Part	Number of non-blank lines
Total	204,445
TMU	86,693
Synthesis	35,458
Event Buffering	15,127
Tracking	41,144
Simulation	10,023
SMU	40,371
Synthesis	35,458
Simulation	49,130
TGU	16,878
Common/Shared	60,055

GTU Tracking Timing

- Computation latency depending on number and tracklet content - 550 ns offset, rising only slightly
- Total latency depending heavily on number of tracklets
- Full hardware simulation with ModelSim and Testbench



Plots: PhD thesis by J. de Cuveland

TRD Beam Test at CERN

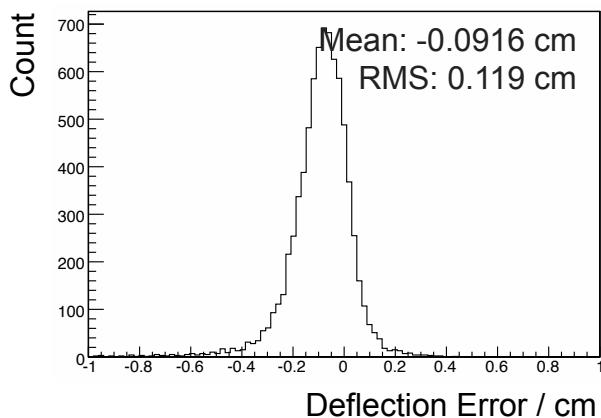


1 TRD Super-module

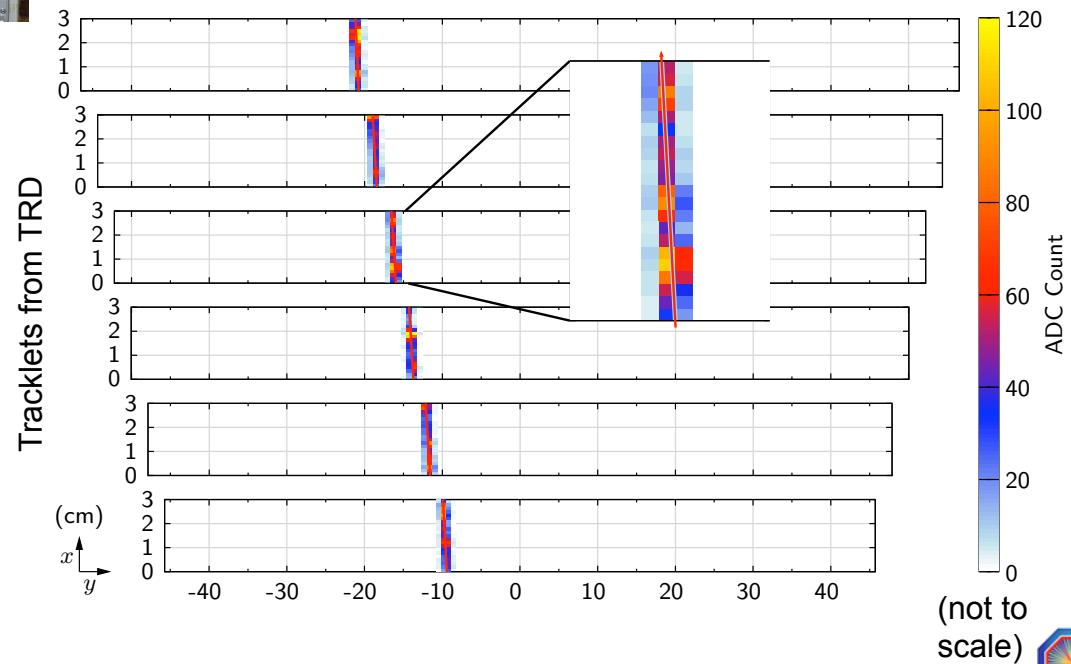
1 GTU Segment

November 2007 Beam Test Setup
at CERN Proton Synchrotron

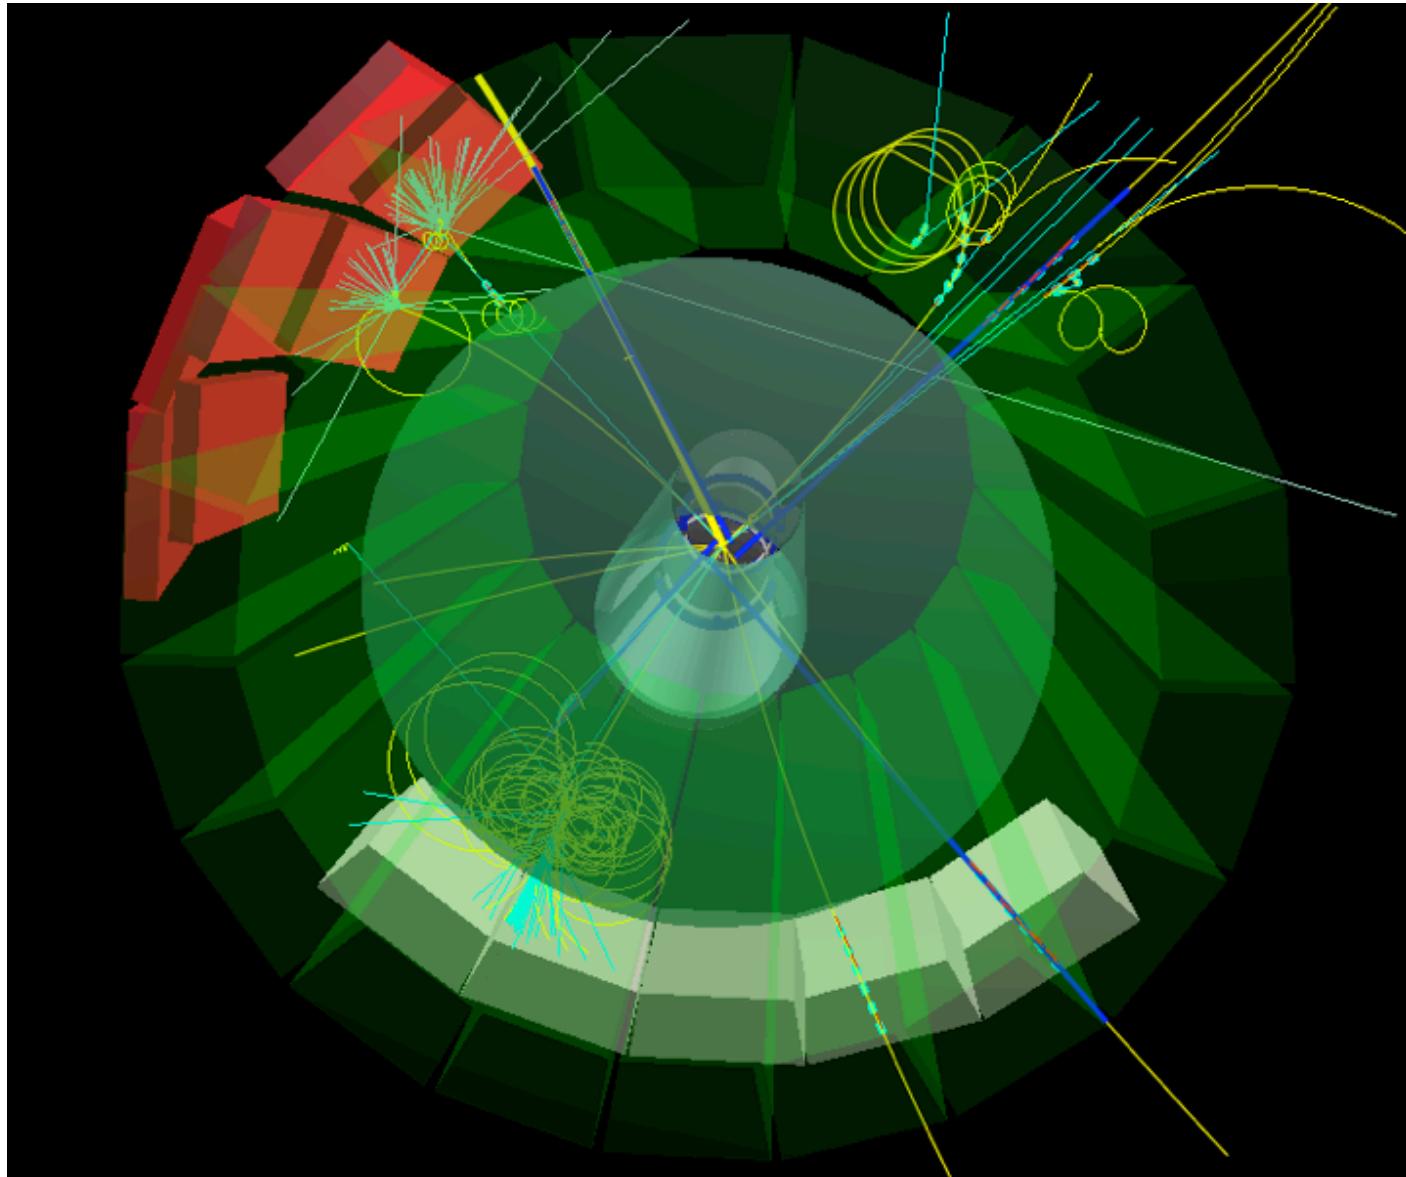
Single Tracklet Deflection Precision



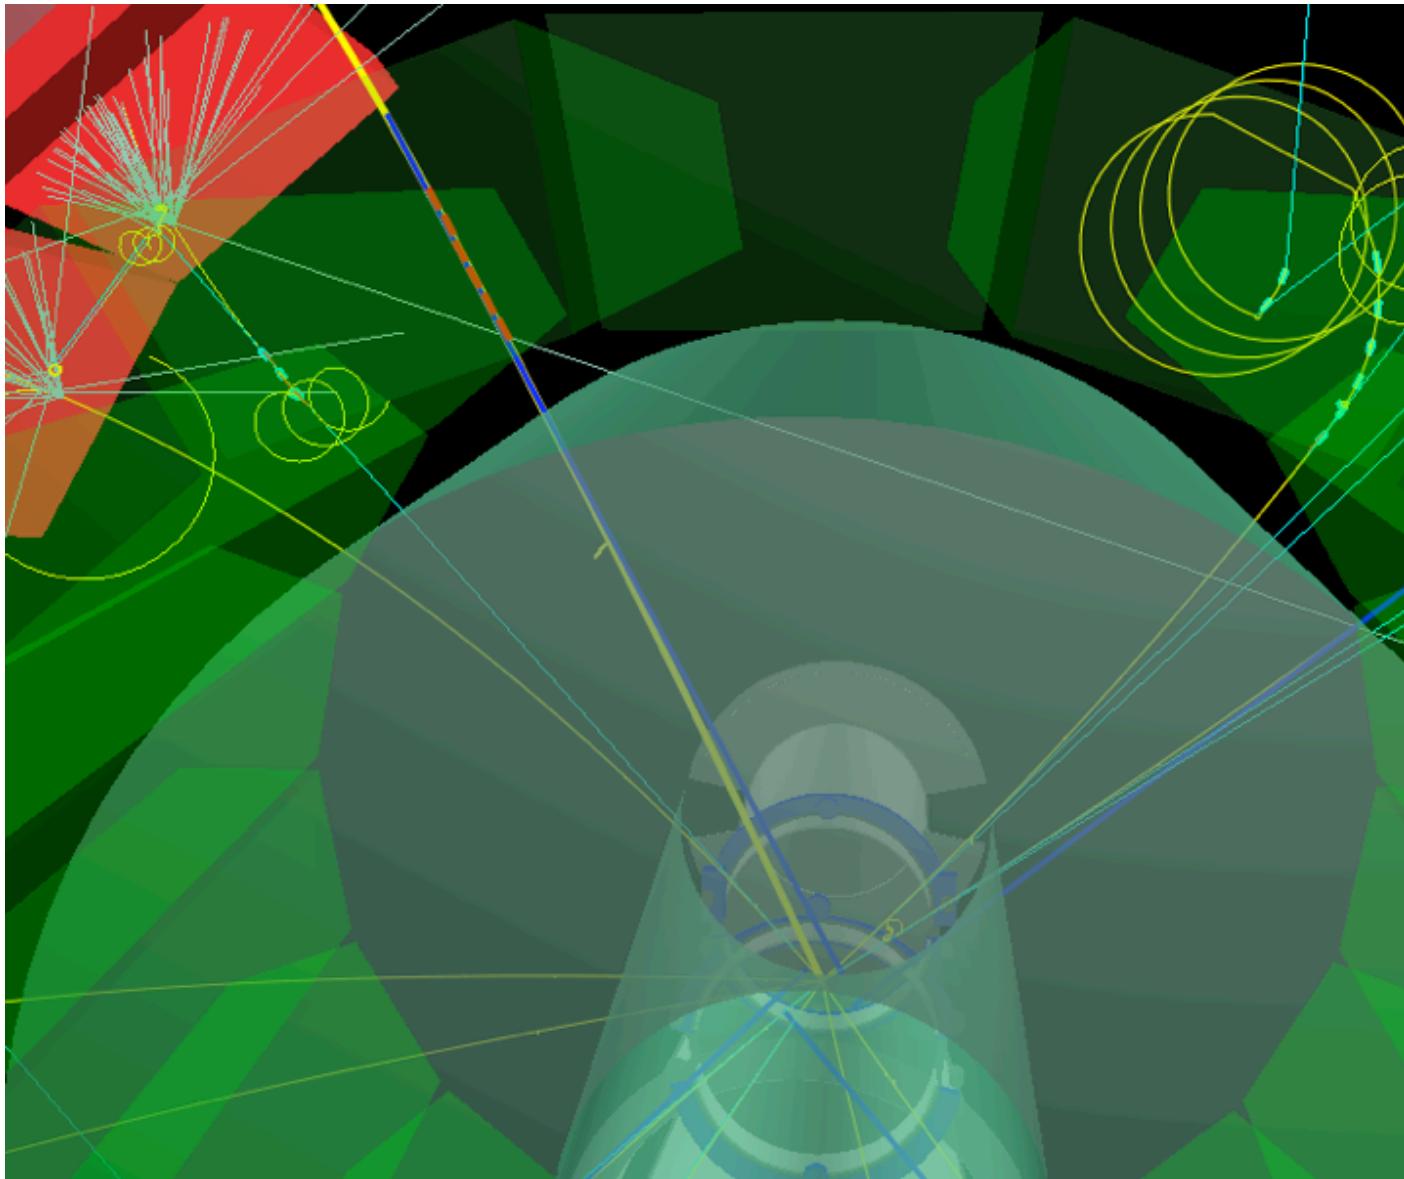
- Accelerator: CERN Proton Synchrotron (PS)
- Particles: Electrons, Pions
(Transverse Momenta: 0.5 – 6 GeV/c)
- Good statistics for detector calibration (More than 1 Mio. events per momentum value)
- 8 days of continuous operation
- First run with tracklets, consistent with raw data



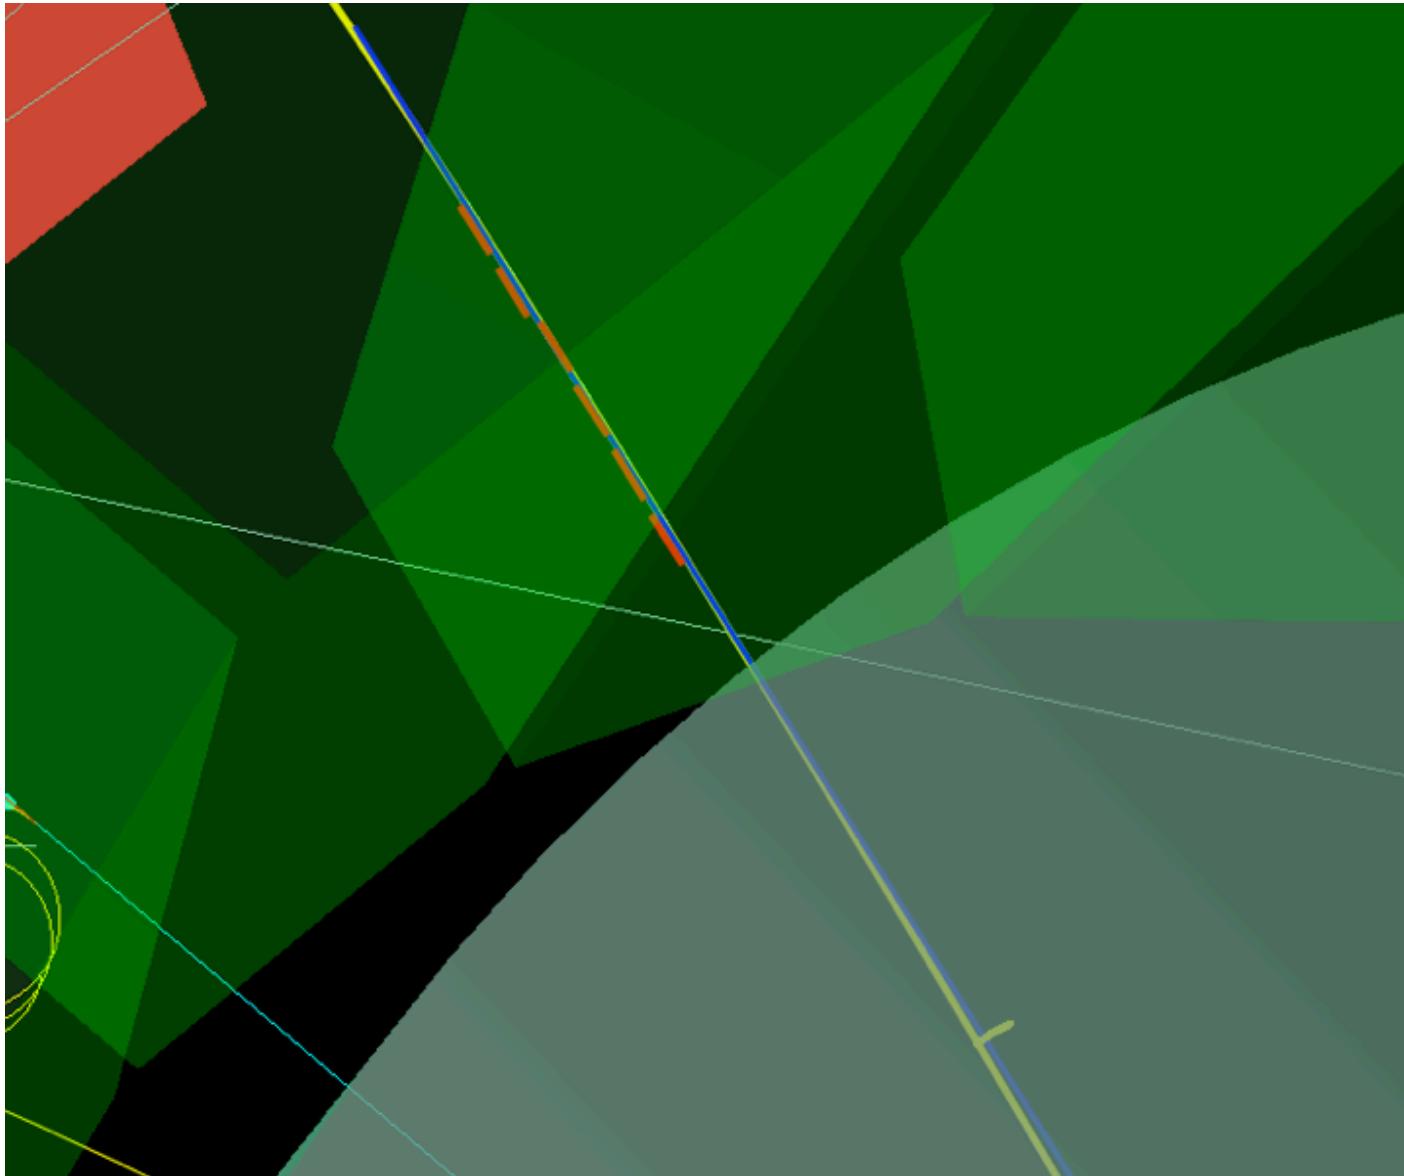
Simplified Event



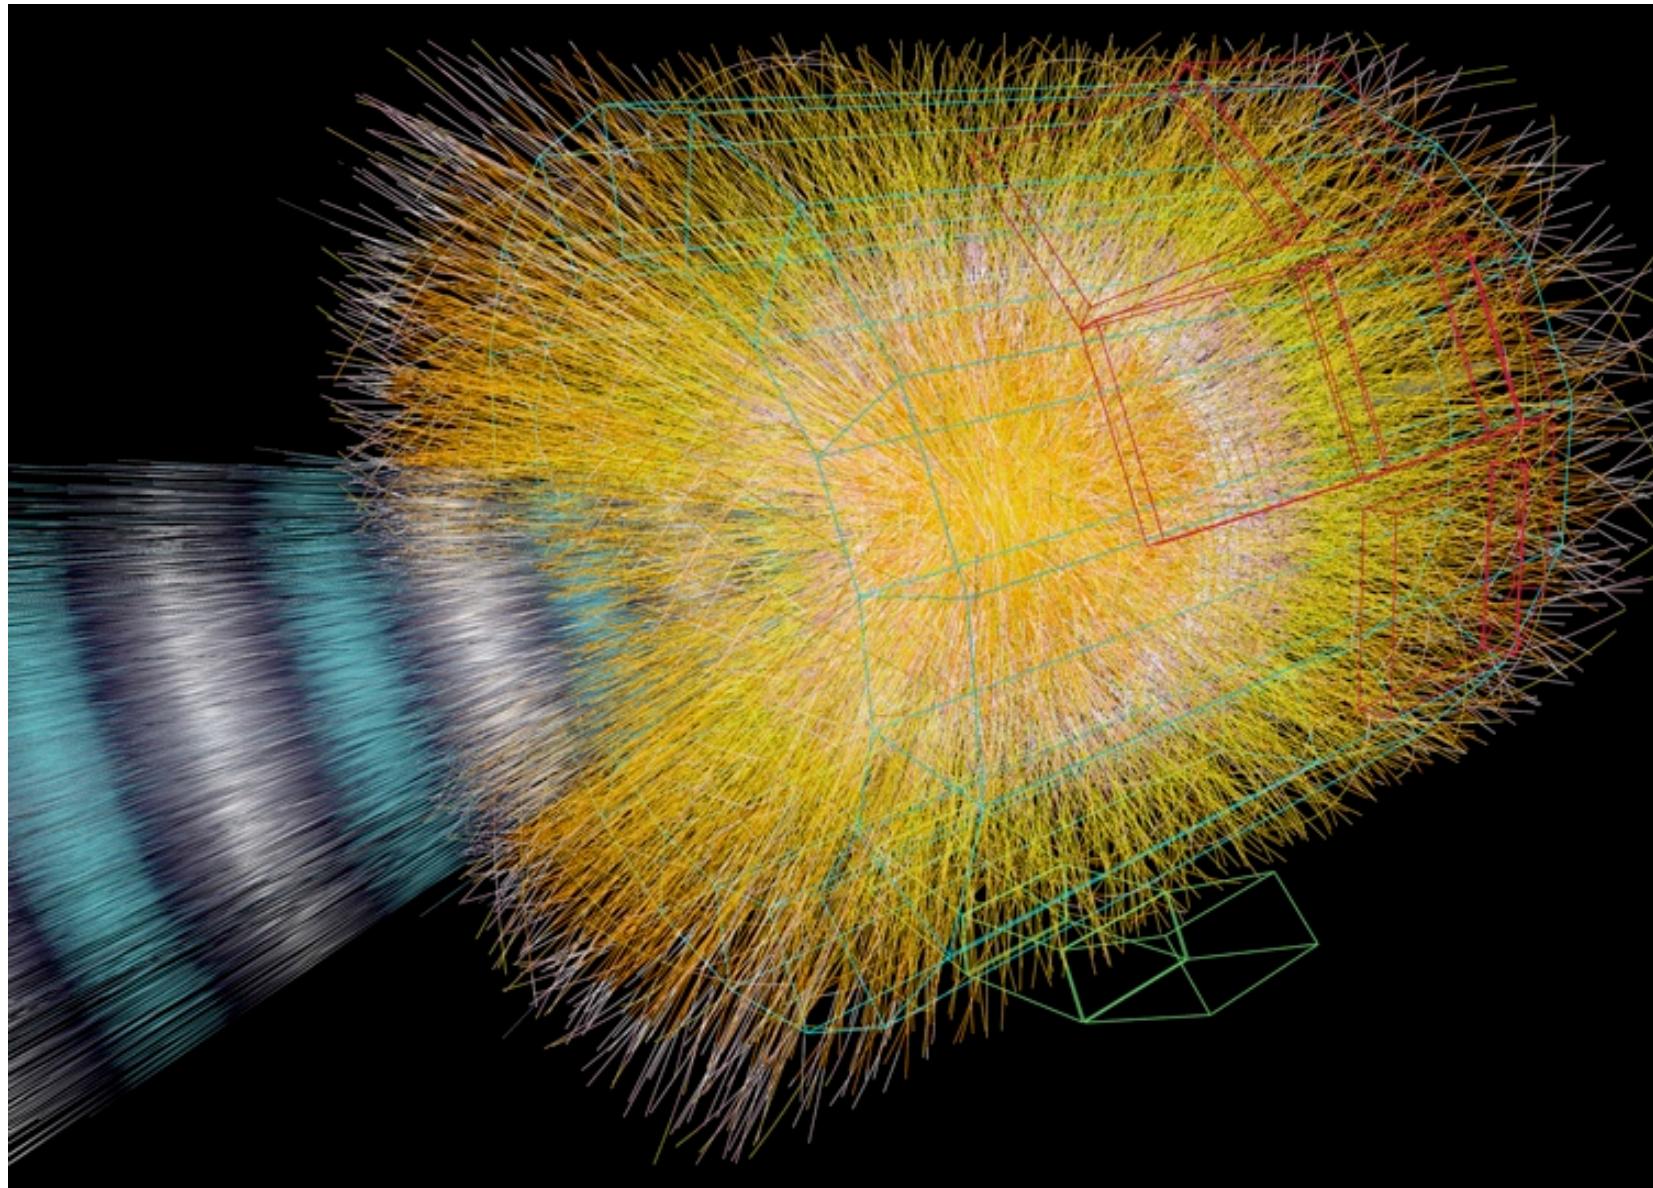
Simplified Event II



Simplified Event III



Realistic Pb-Pb Event



Concentrator Node

- Inputs: reconstructed tracks from first tier & raw data
- Tasks:
 - Apply trigger schemes
 - Interface to data acquisition system, process trigger sequences and read-out

Concentrator Node

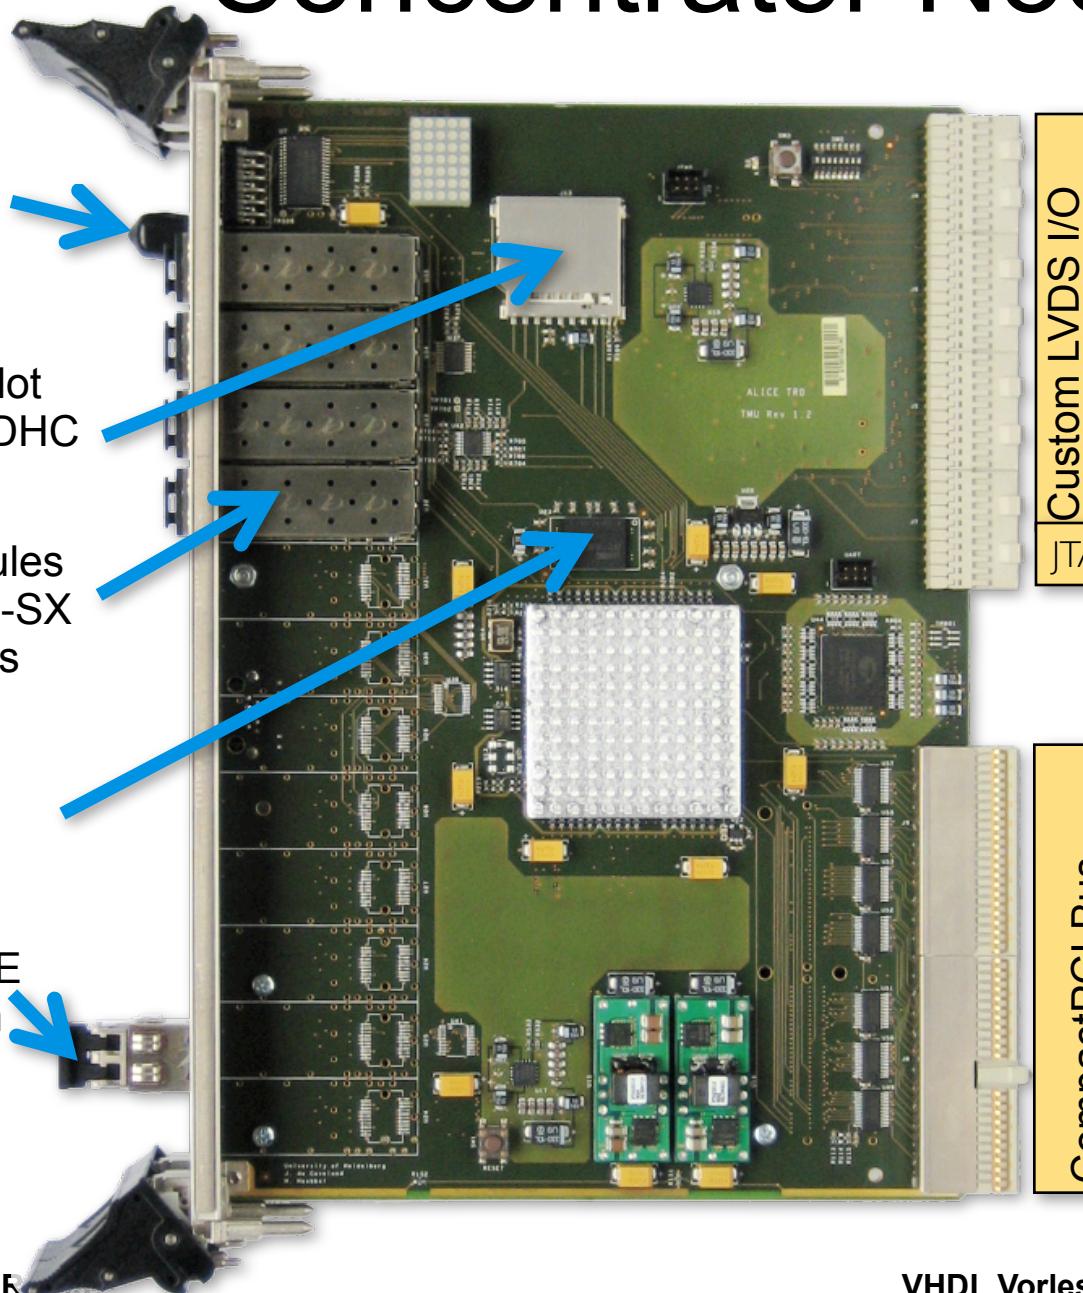
Interface to
ALICE TTC
system

SD Card Slot
4 GByte SDHC
Cards

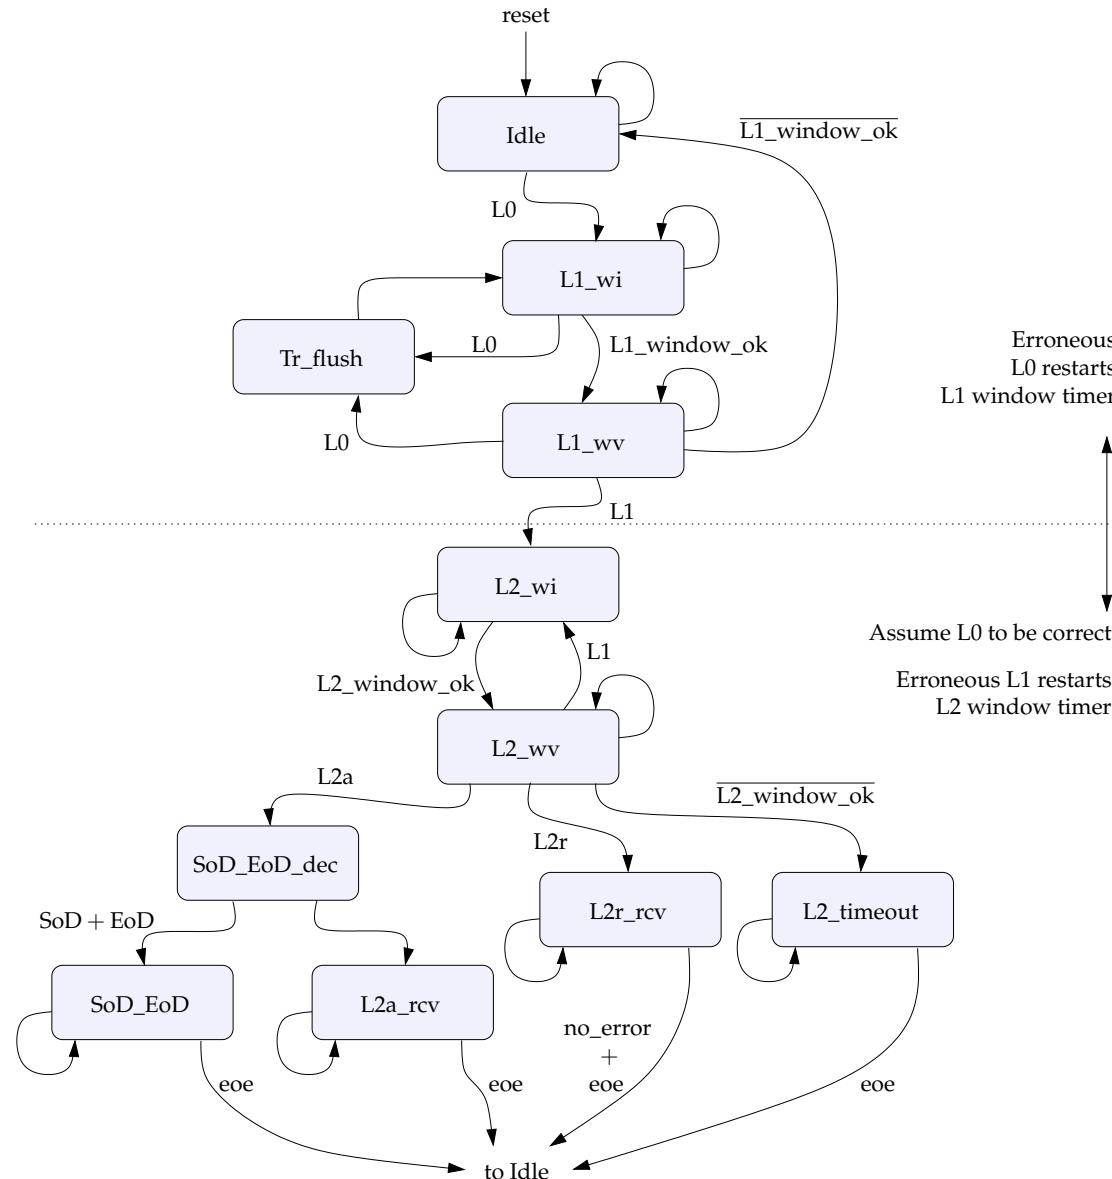
SFP modules
1000Base-SX
to switches

DDR2
SDRAM
64 MByte

Link to ALICE
DAQ system



FSM Example: Trigger Handling

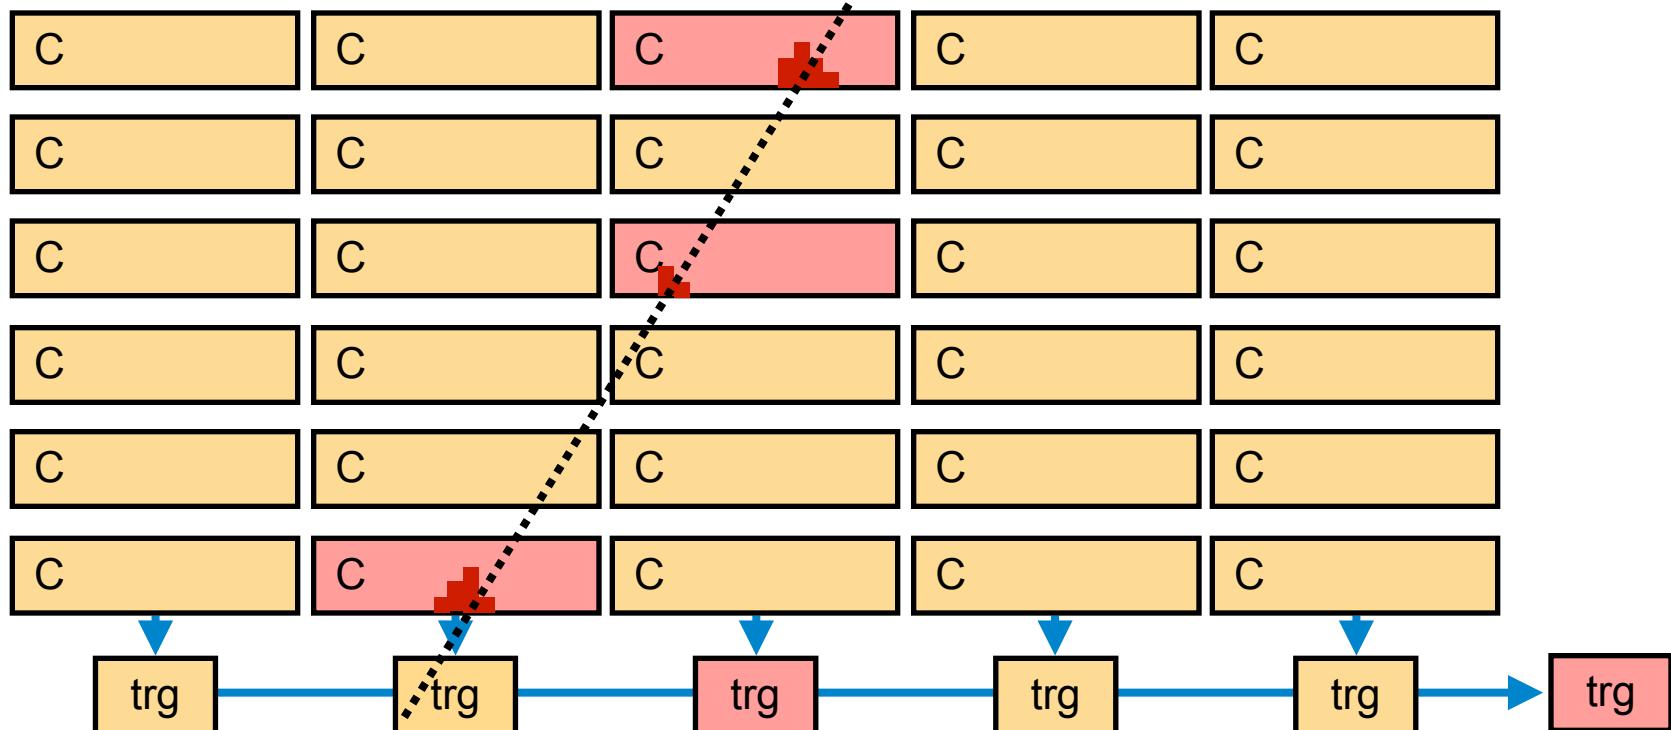


Trigger Schemes

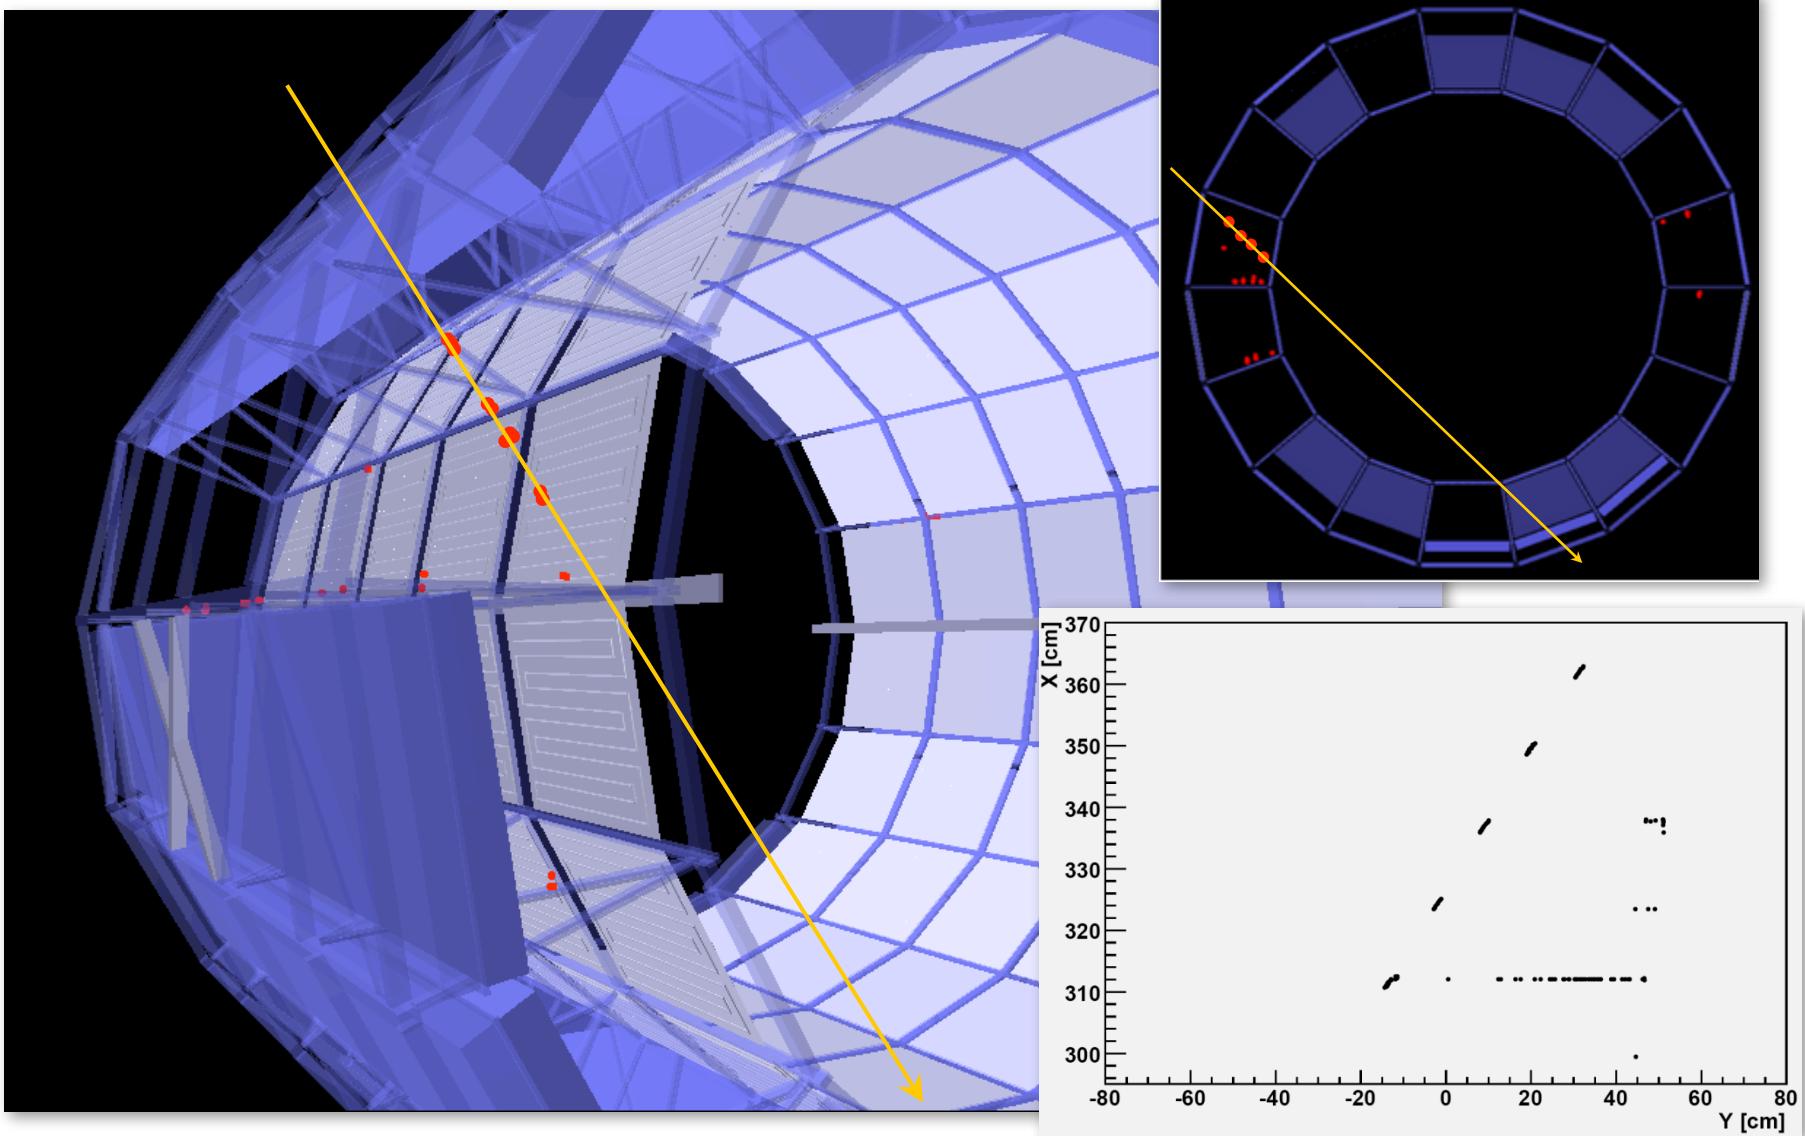
- Cosmic Trigger
- Jet Trigger
 - Simple jet definition: more than certain number of high- p_t tracks through a given detector volume
 - Additional conditions: jet location, coincidences, ...
 - $N_{\text{tracks}}=1$: single high- p_t particle trigger
- Di-Lepton Decay Trigger
 - Coincidence of high- p_t e^\pm tracks
 - Calculation of invariant mass for higher selectivity
- Various Other Schemes
 - Ultra-peripheral collisions

Cosmics Trigger

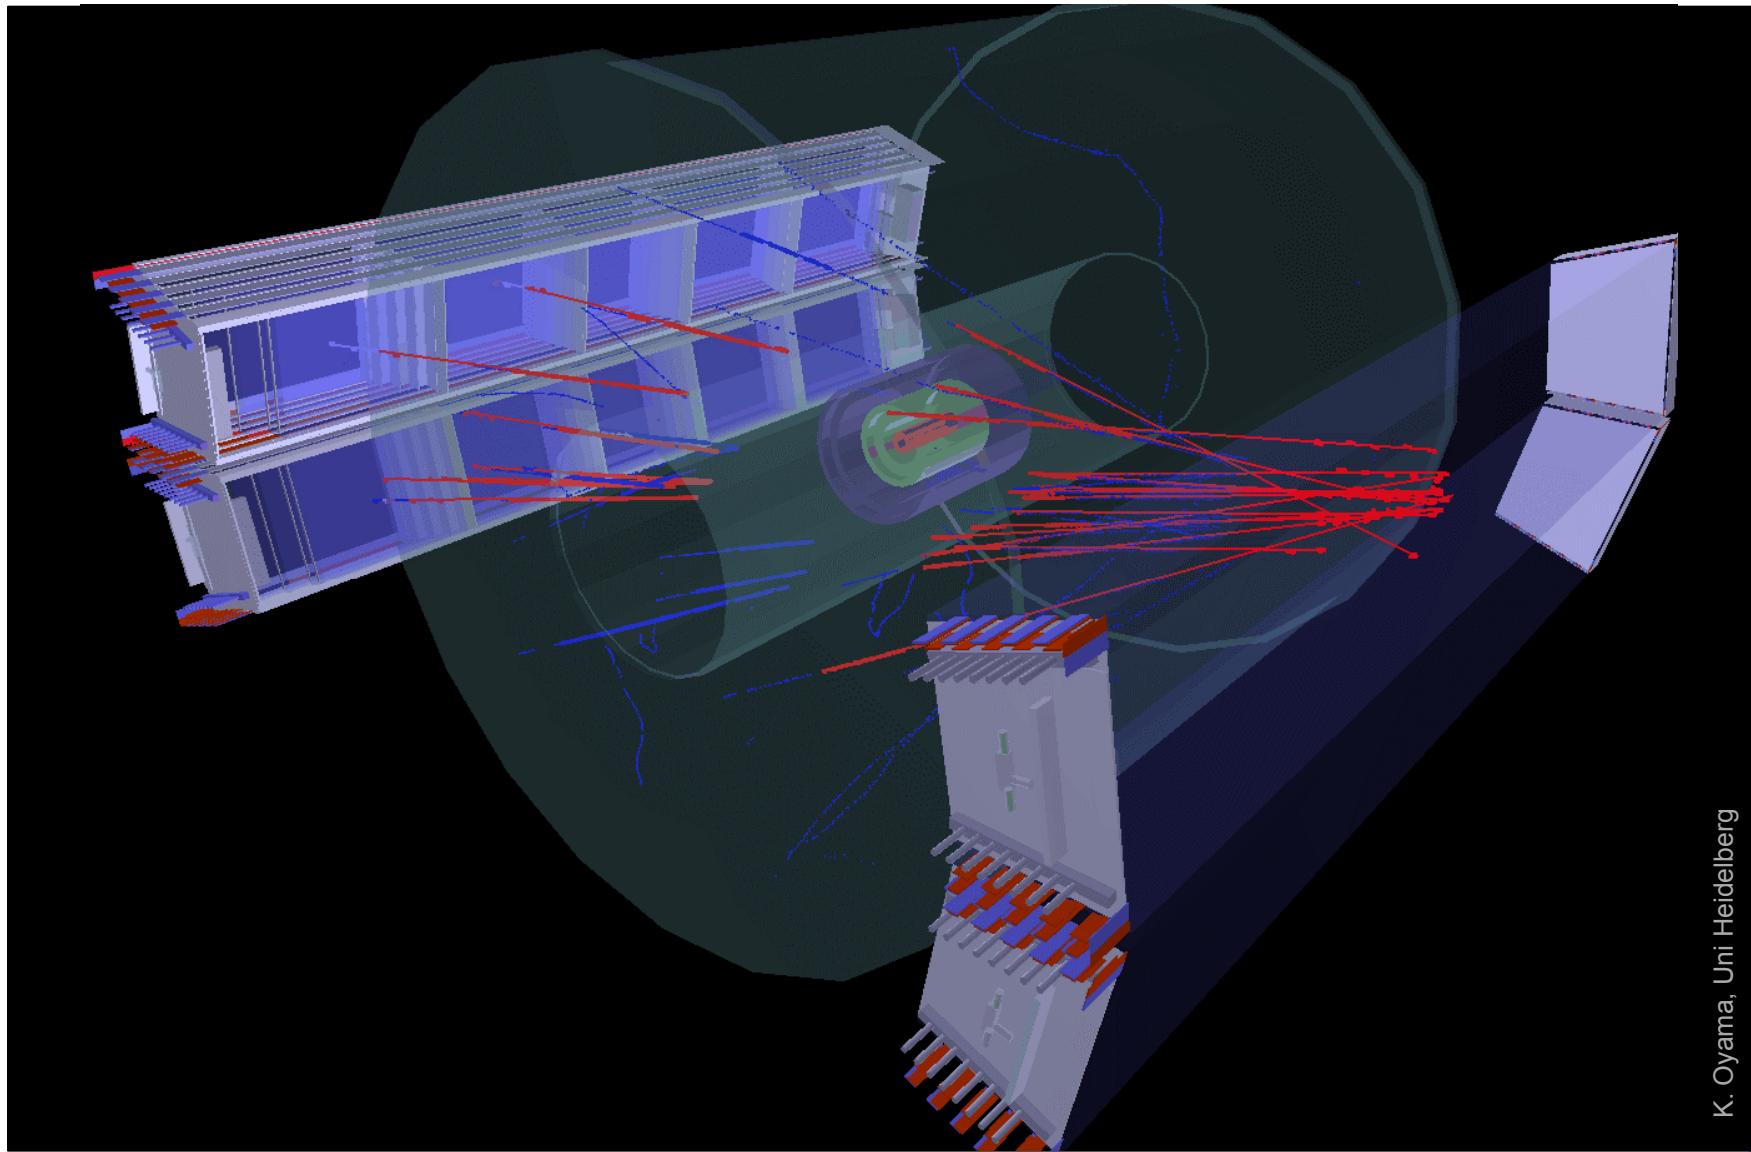
- Chamber: $\text{min} \leq \text{sum of charge/hits} \leq \text{max}$
- Stack: $\text{min} \leq \text{chambers hit} \leq \text{max}$
- Supermodule: $\text{min} \leq \text{stacks hit} \leq \text{max}$
- Detector: coincidence between supermodules



Cosmic Event Triggered



Cosmic Event Triggered II

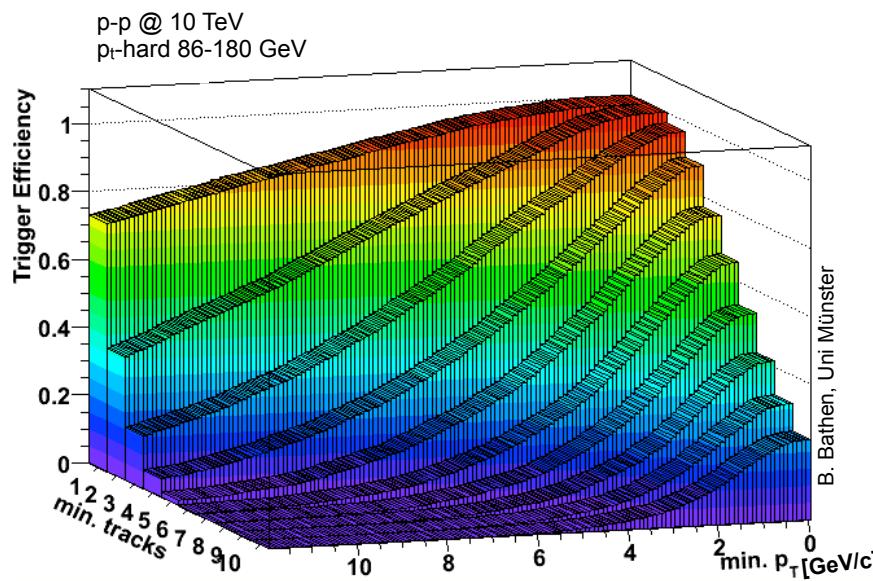
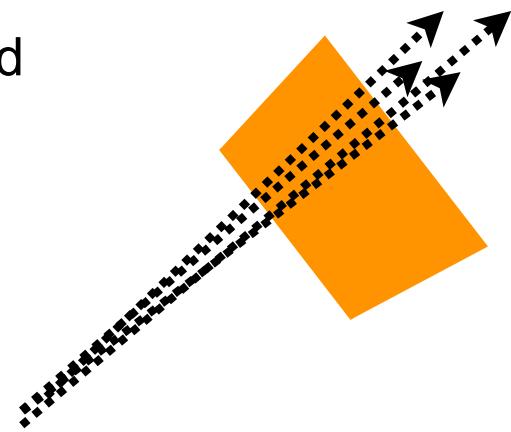


K. Oyama, Uni Heidelberg



Jet Trigger

- Identify tracks with $p_t \geq p_{t,\text{threshold}}$ within certain region
- Threshold conditions:
 - Number of tracks
 - Sum of momenta of tracks

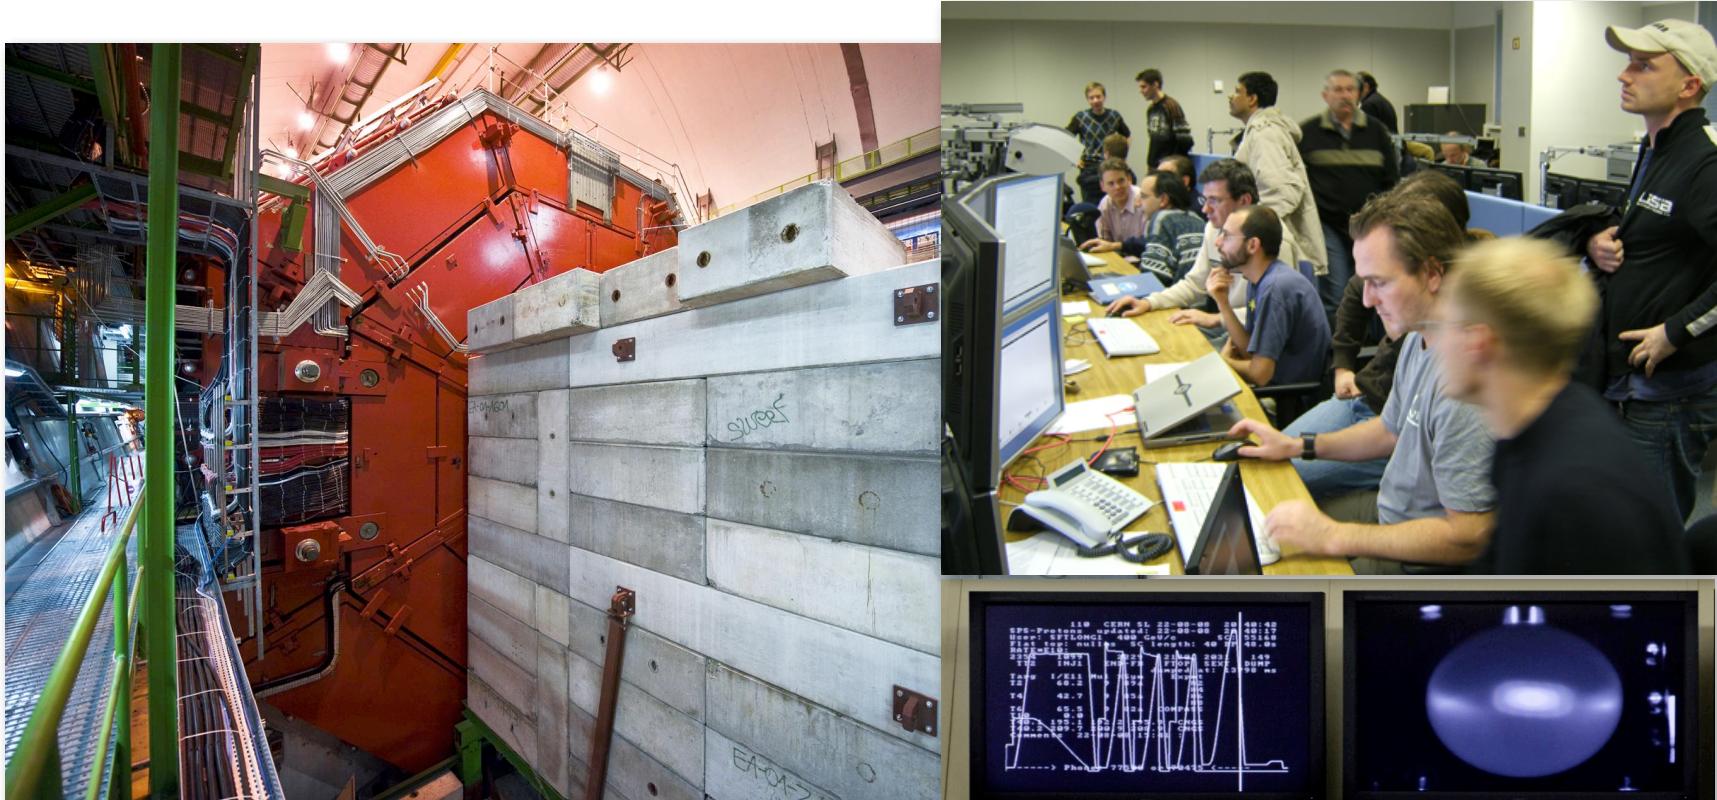


- Granularity: sub-stack-sized areas overlapping in z- and Φ -direction
- Realizable at first trigger stage
- Multi-Jet coincidence at top level

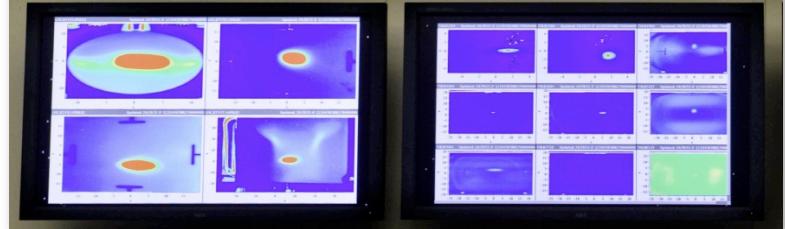
Di-Lepton Trigger

- Find e^+e^- pairs with invariant mass within certain range (J/ψ , Υ , ...)
- Huge combinatorics for Pb-Pb collisions
- Current work:
 - Pre-selection of track candidates, application of sliding window algorithms
 - Massively parallelized invariant mass calculation in FPGA hardware
 - Fast trigger contribution for Level-1 (after 6 μ s) more elaborate decision for Level-2 (80 μ s)

Waiting For LHC Start-Up



LHC restart
scheduled for end of
year...

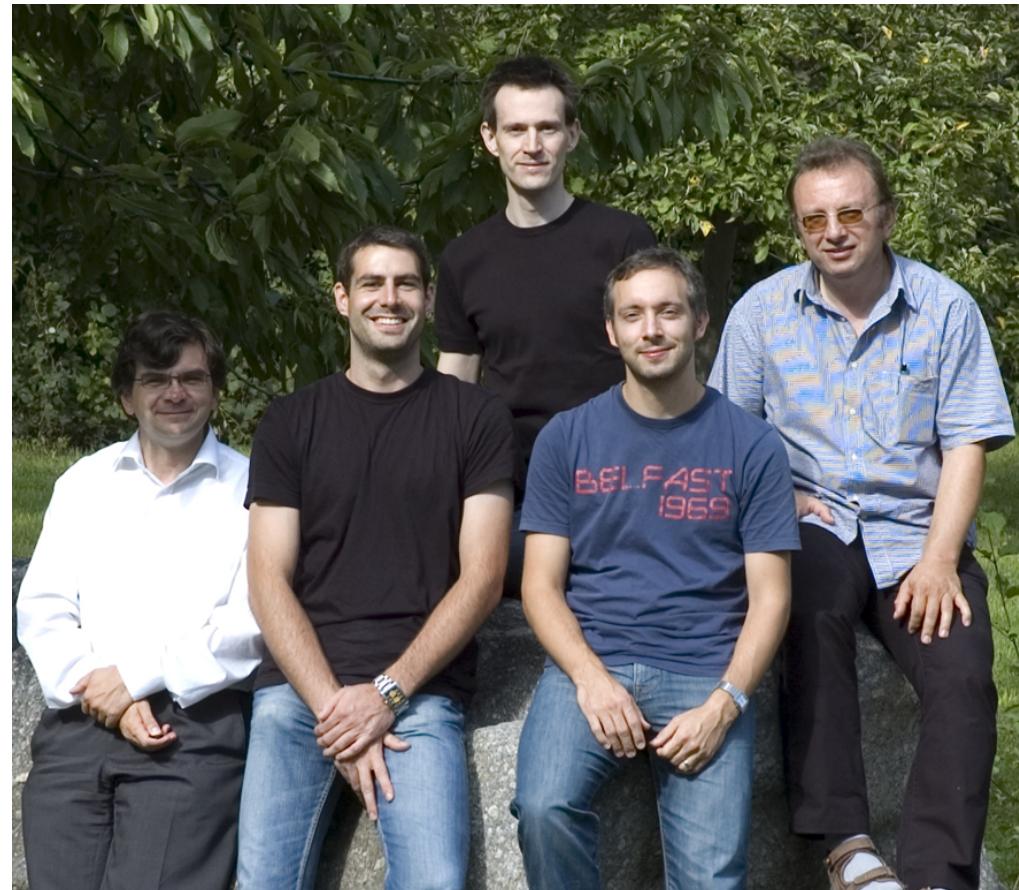


The GTU People

Venelin Angelov,
Jan de Cuveland,
Stefan Kirsch,
and Felix Rettig

Former members:
Thomas Gerlach,
Marcel Schuh

Prof. Volker Lindenstruth
Chair of Computer Science
Kirchhoff Institute of Physics
University of Heidelberg
Germany



<http://www.ti.uni-hd.de>



bmb+f - Förderschwerpunkt
ALICE
Großgeräte der physikalischen
Grundlagenforschung

Deutsche
Forschungsgemeinschaft
DFG

