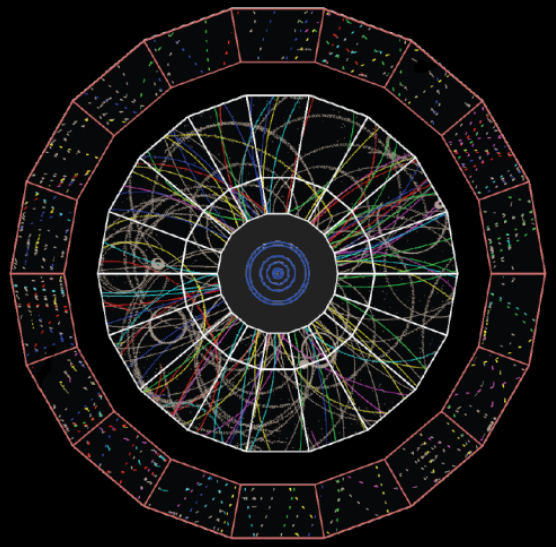

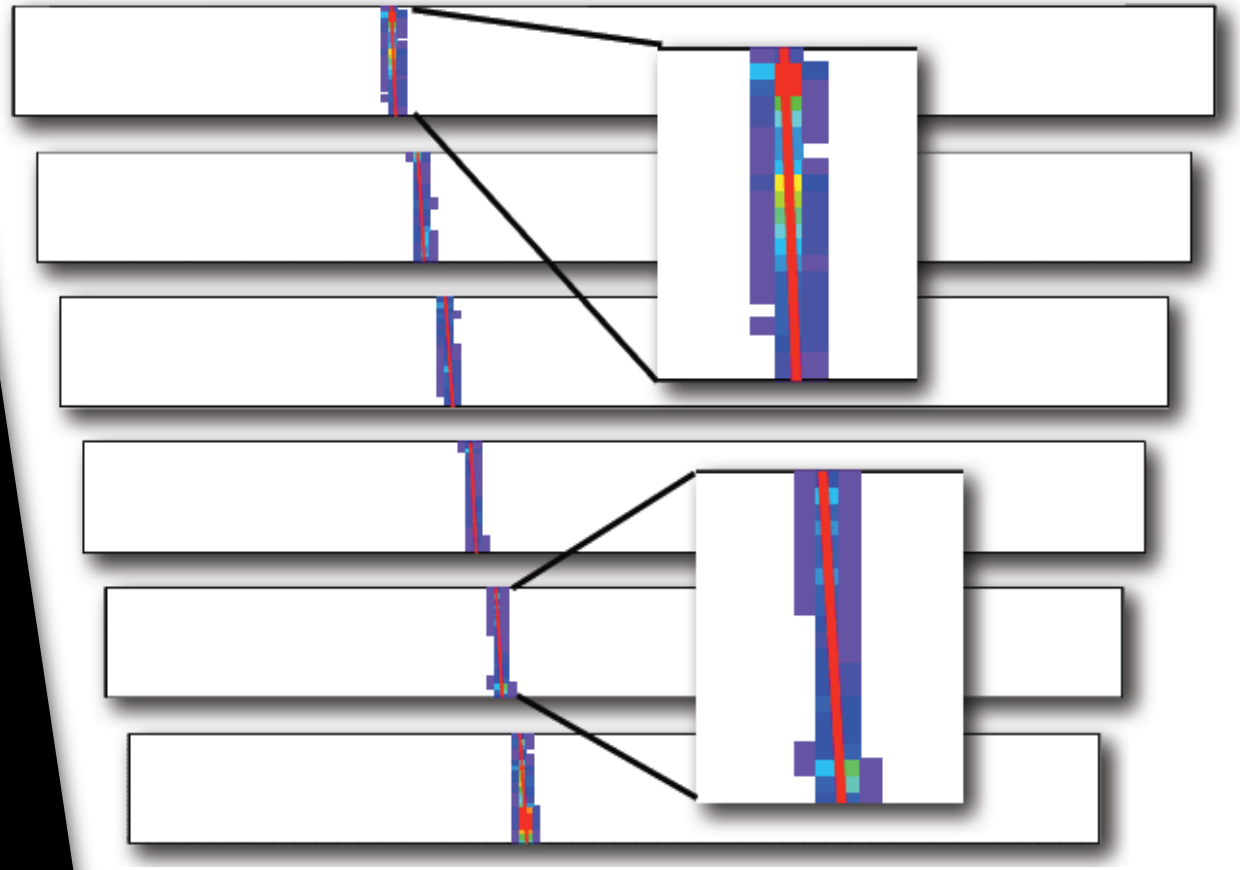

The ALICE TRD Global Tracking Unit

Felix Rettig
University of Heidelberg
Kirchhoff Institute of Physics

SEI Workshop, KIT

September 29th 2010

The Experiment ALICE

CERN

- pp @ 14 TeV
- PbPb @ 1150 TeV

ALICE

- Research on QGP in PbPb collisions
- Many detectors covering a wide momentum range & PID
- High multiplicity events in Pb-Pb collisions

ALICE & TRD

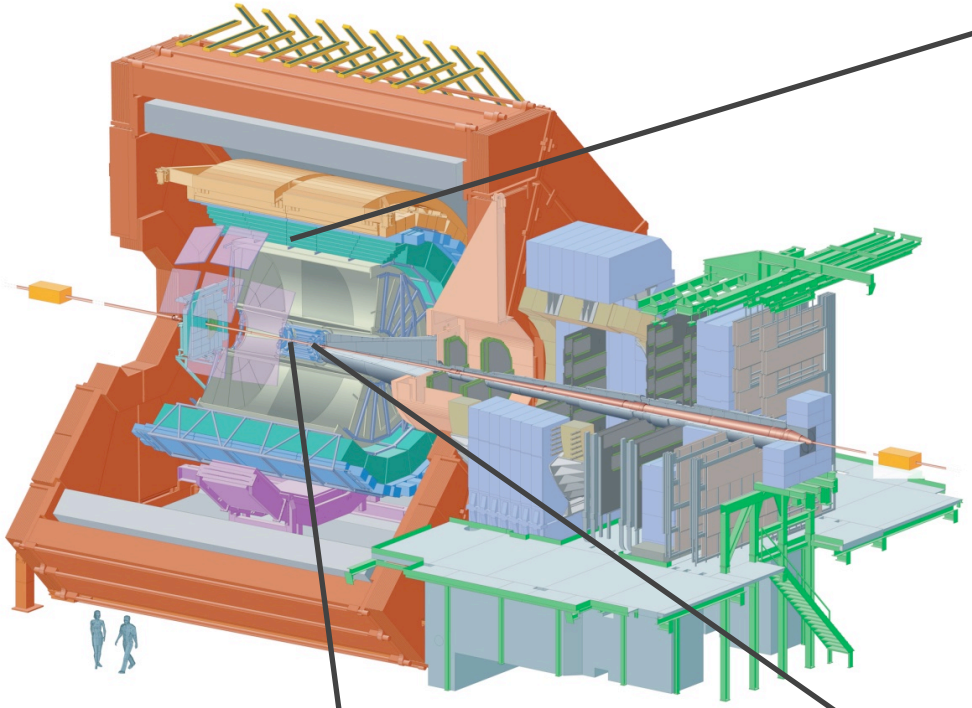

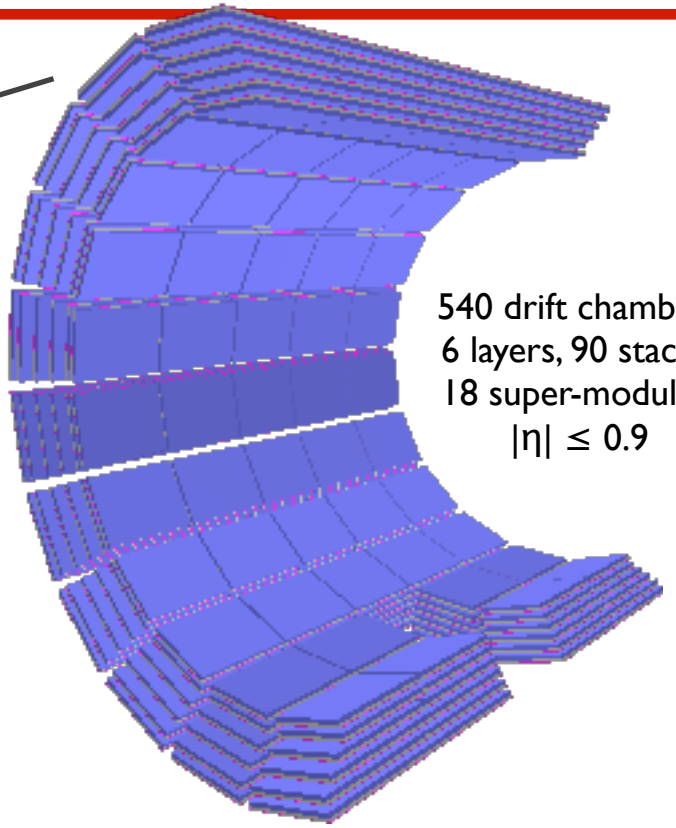

540 drift chambers
6 layers, 90 stacks,
18 super-modules,
$|\eta| \leq 0.9$

Transition Radiation Detector (TRD)

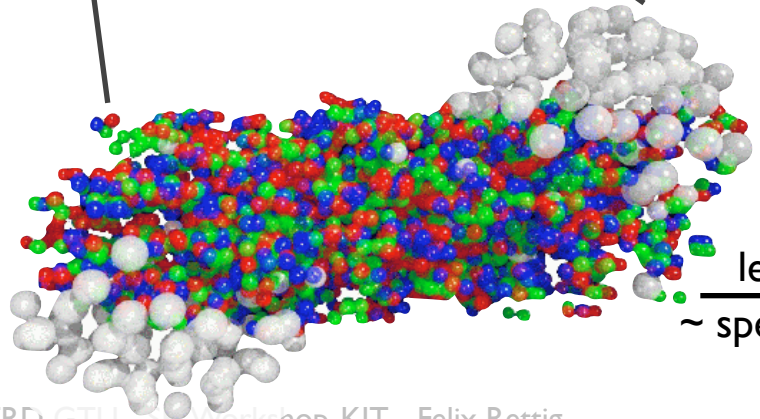

lead nucleus
~ speed of light

lead nucleus
~ speed of light

Task of the TRD

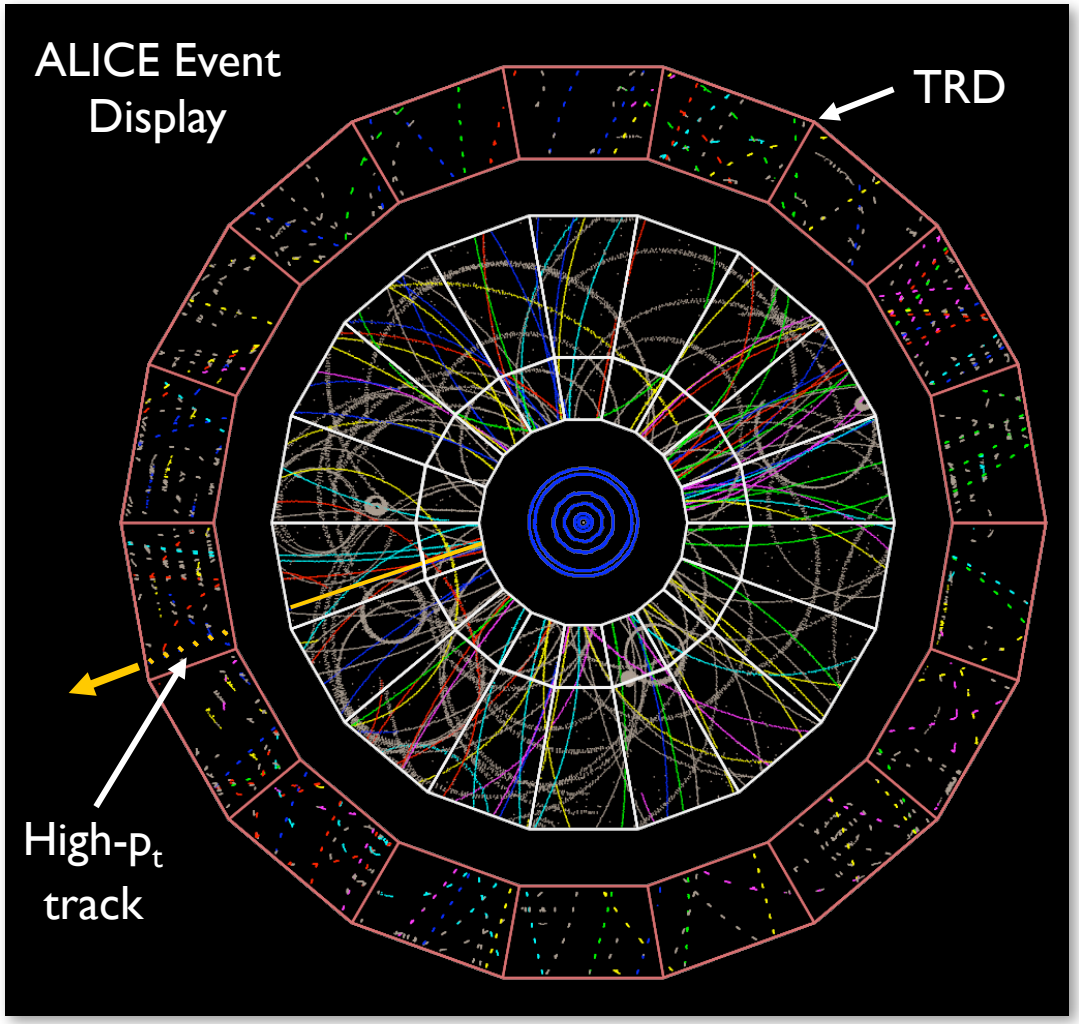

- High multiplicities: few thousand charged particle tracks in acceptance per event
- Fast trigger detector: L1 trigger after 6.2µs
  - Online reconstruction of high-p_t tracks, calculation of p_t
  - Various trigger schemes
- Barrel tracking detector: raw data for offline analysis
  - Raw data buffering & forwarding to data acquisition system
  - Support interlaced triggers and multi-event buffering, interface to ALICE central trigger

TRD Data Chain

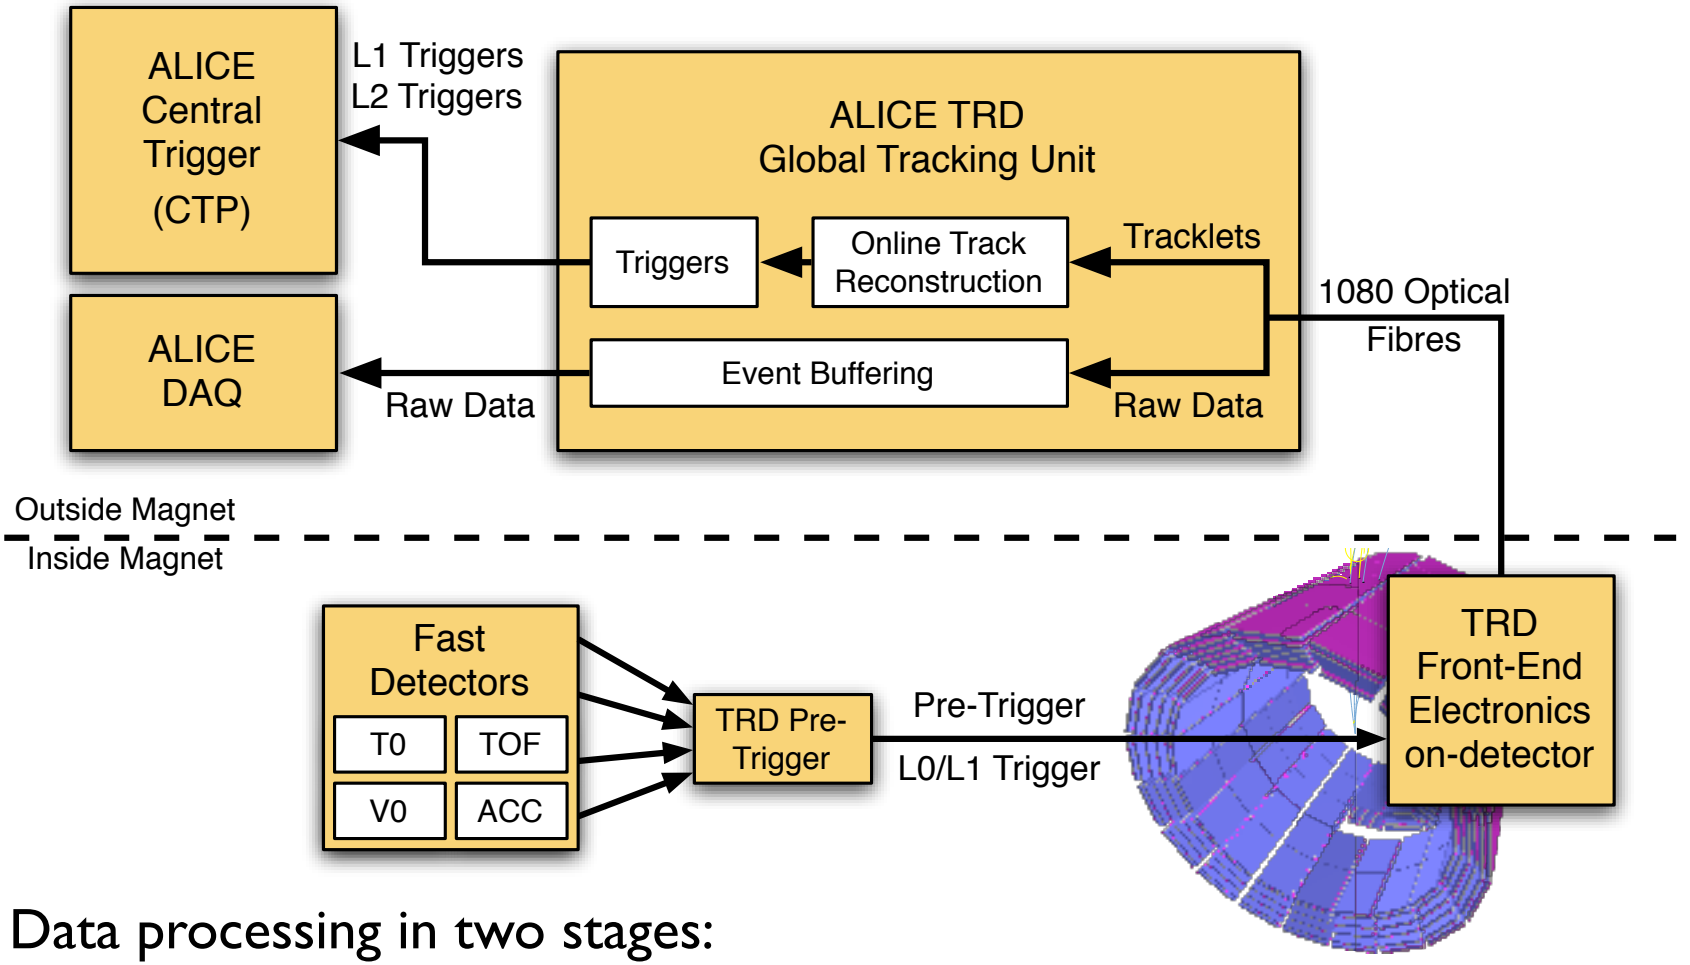

Data processing in two stages:

- On-Detector Front-End Electronics: 65,564 ASICs with 262,256 CPUs
- Global Tracking Unit: 109 FPGAs

TRD Front-End Data Processing

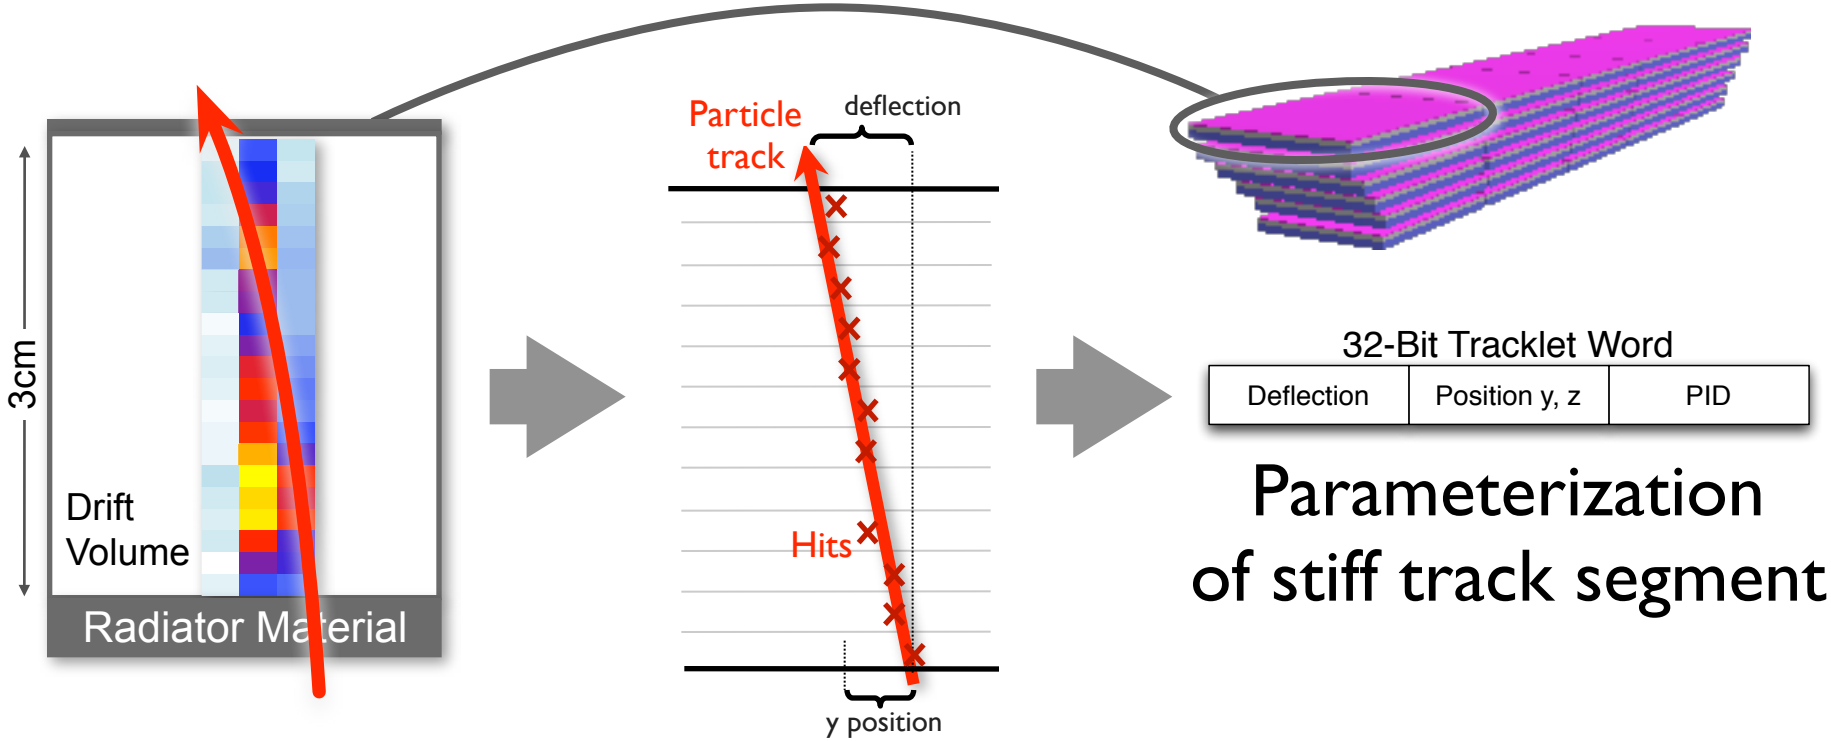

- 540 drift chambers, 6 radial layers, 18 sectors azimuthal
- 1.4 million analog channels
- 10 MHz sampling rate, 10 bit
- Buffer for one single event

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

- Up to 20,000 tracklet words, each 32-bit wide
- Transmission out of magnet via 1080 optical fibres operating at 2.5 Gbit/s
- 2.1 Tbit/s total bandwidth

Simplified Example: Front-End Tracklets

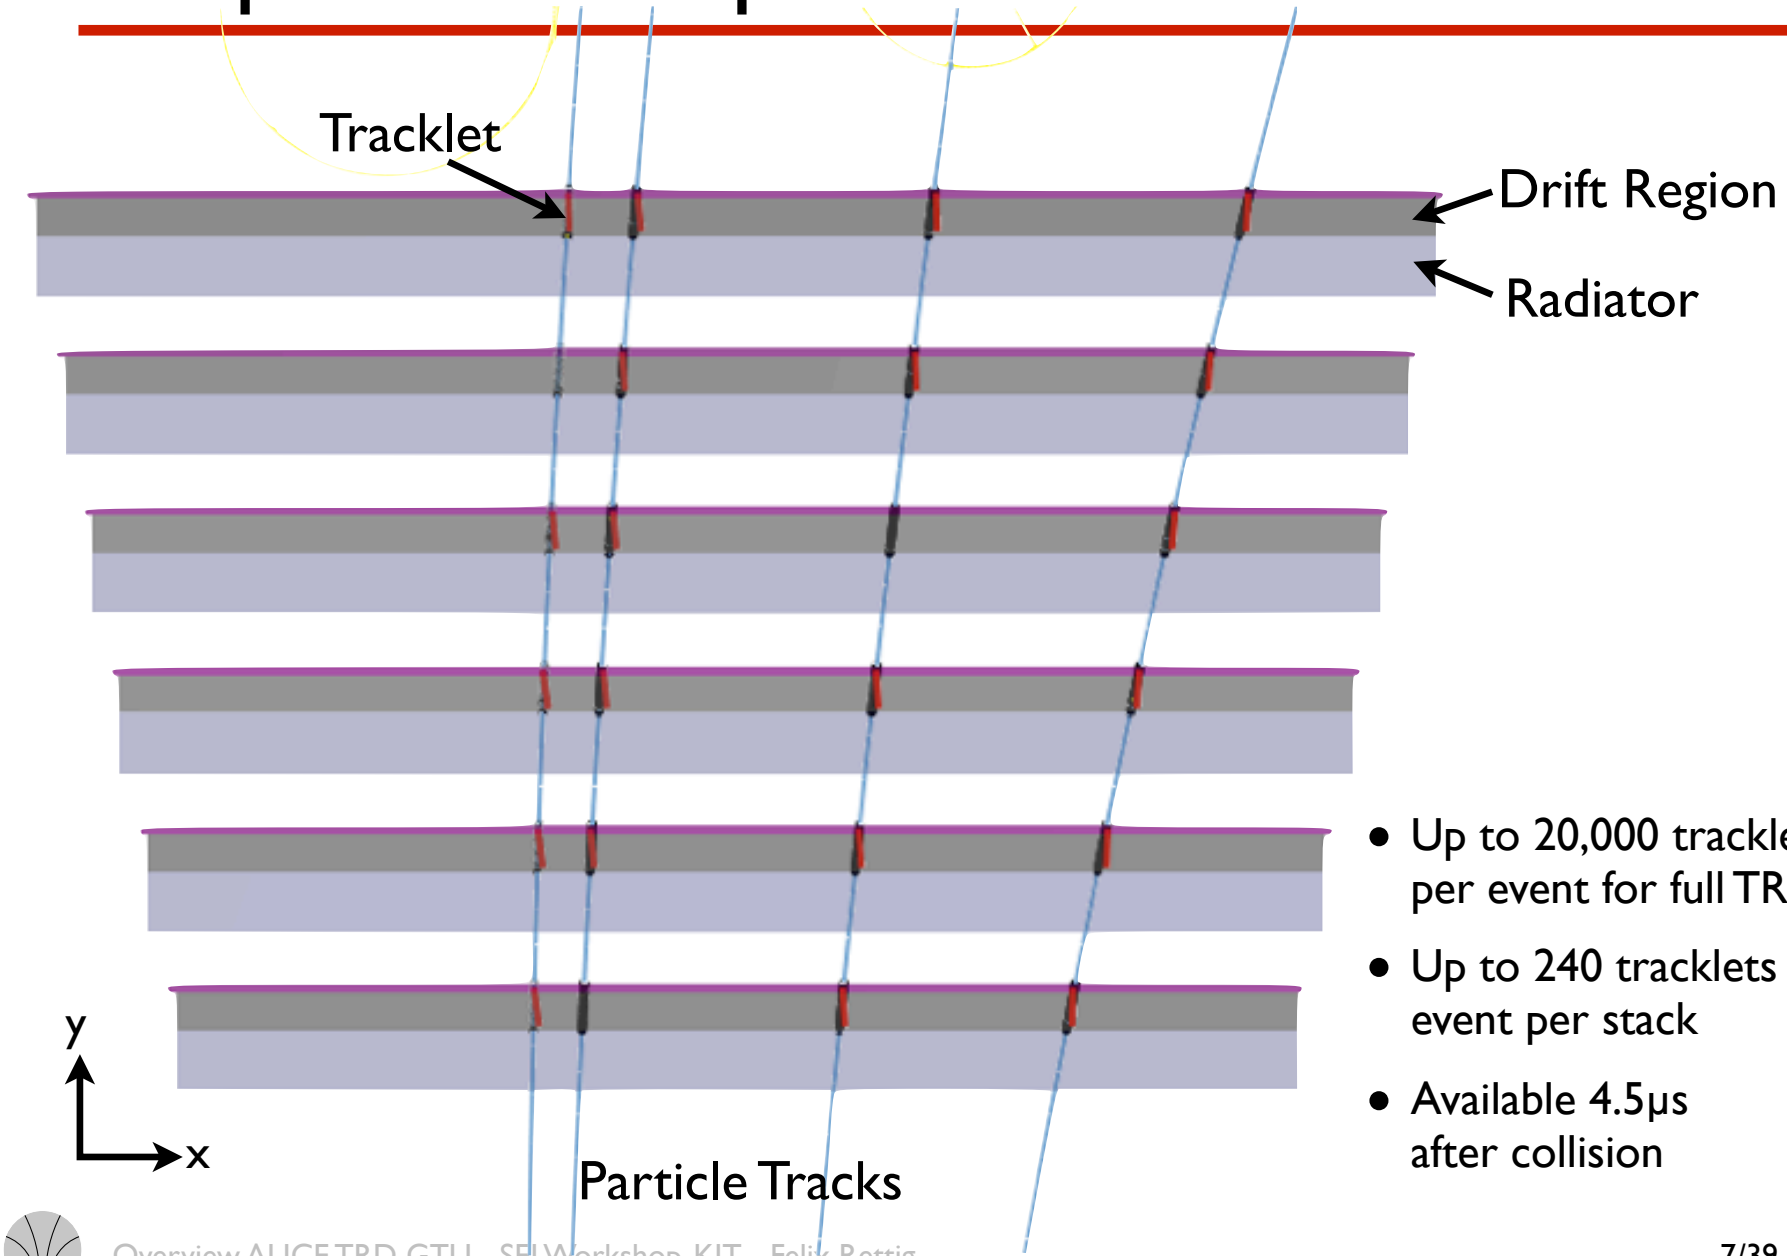

- Up to 20,000 tracklets per event for full TRD
- Up to 240 tracklets per event per stack
- Available 4.5µs after collision

Simplified Example: Reconstructed Tracks

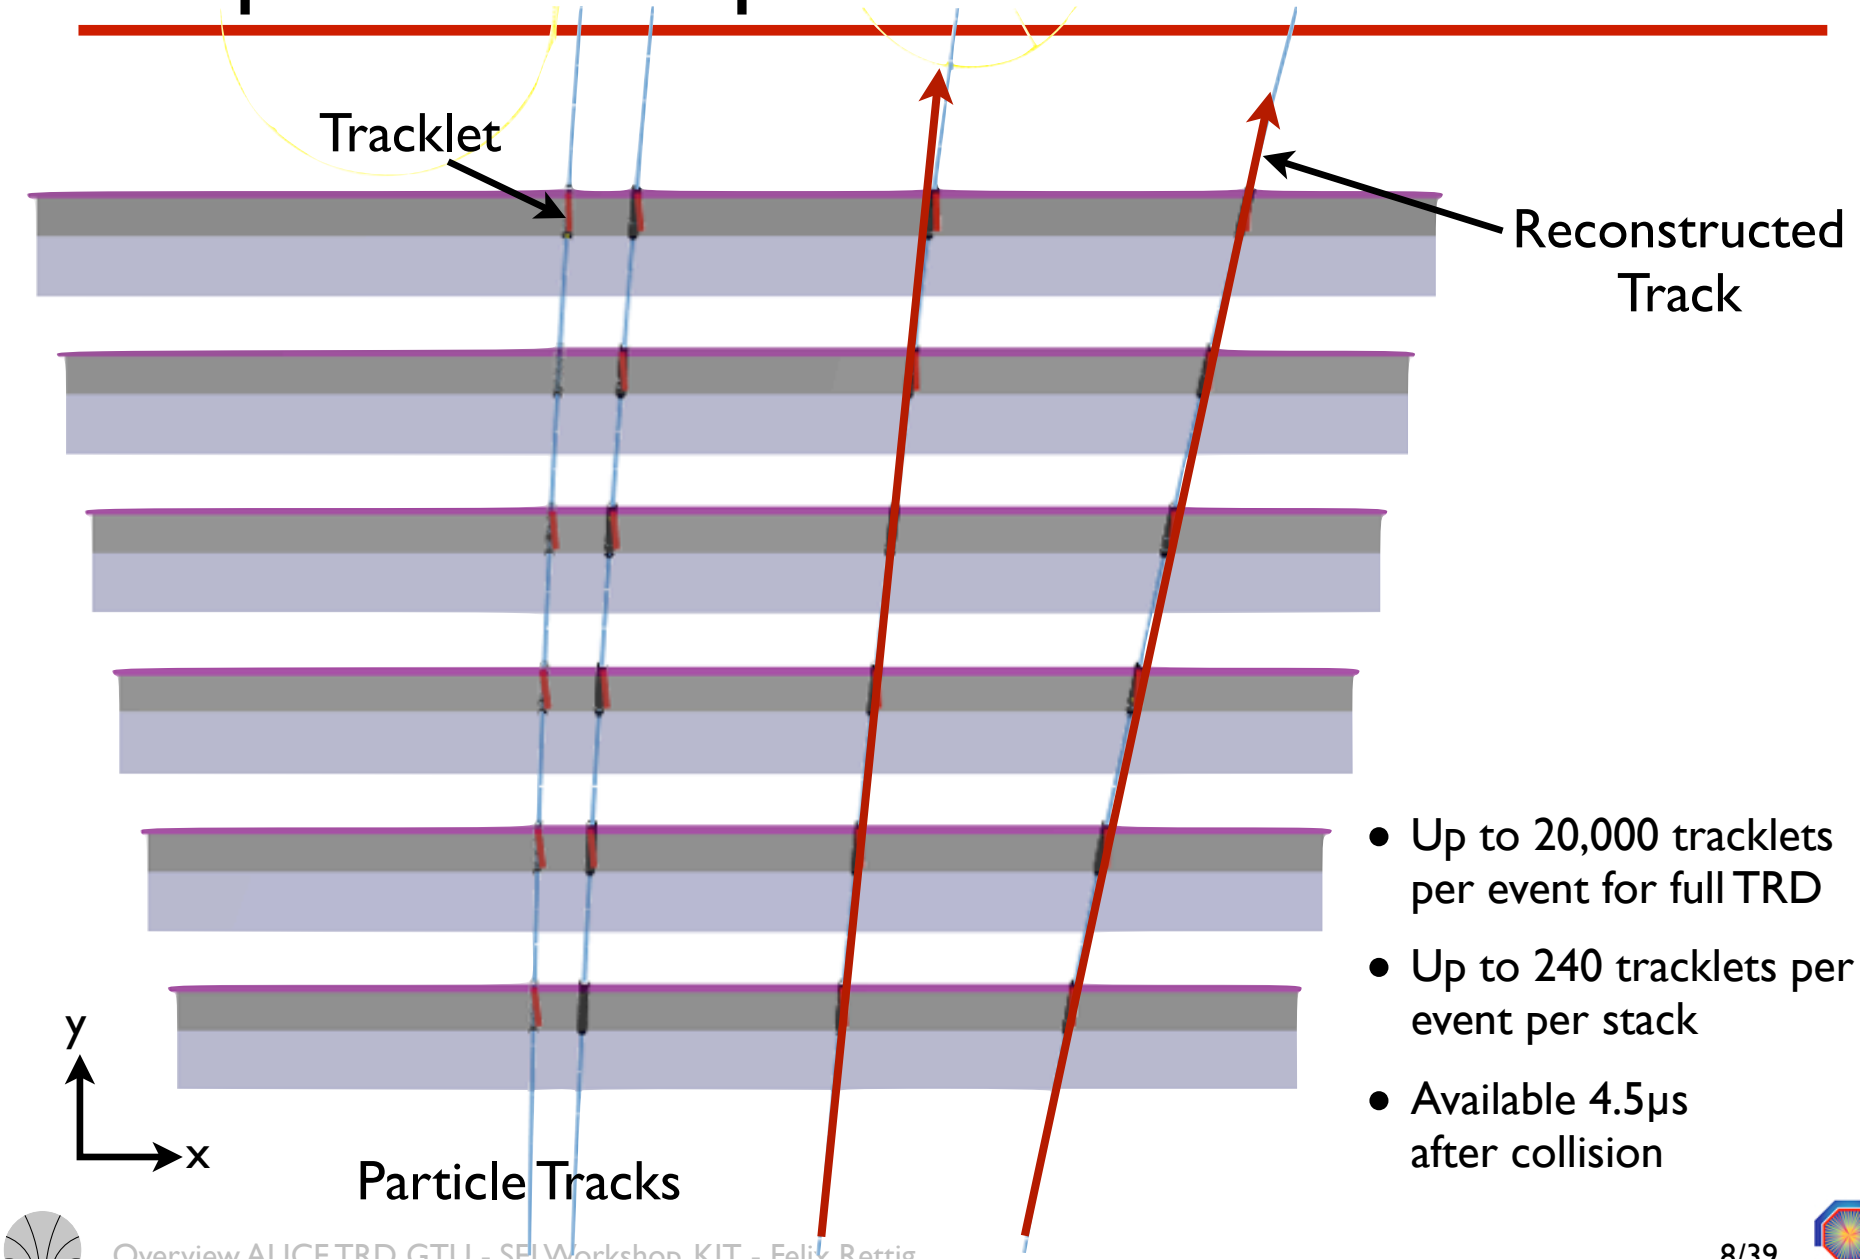

- Up to 20,000 tracklets per event for full TRD
- Up to 240 tracklets per event per stack
- Available 4.5µs after collision

Tight Timing Requirements for Trigger

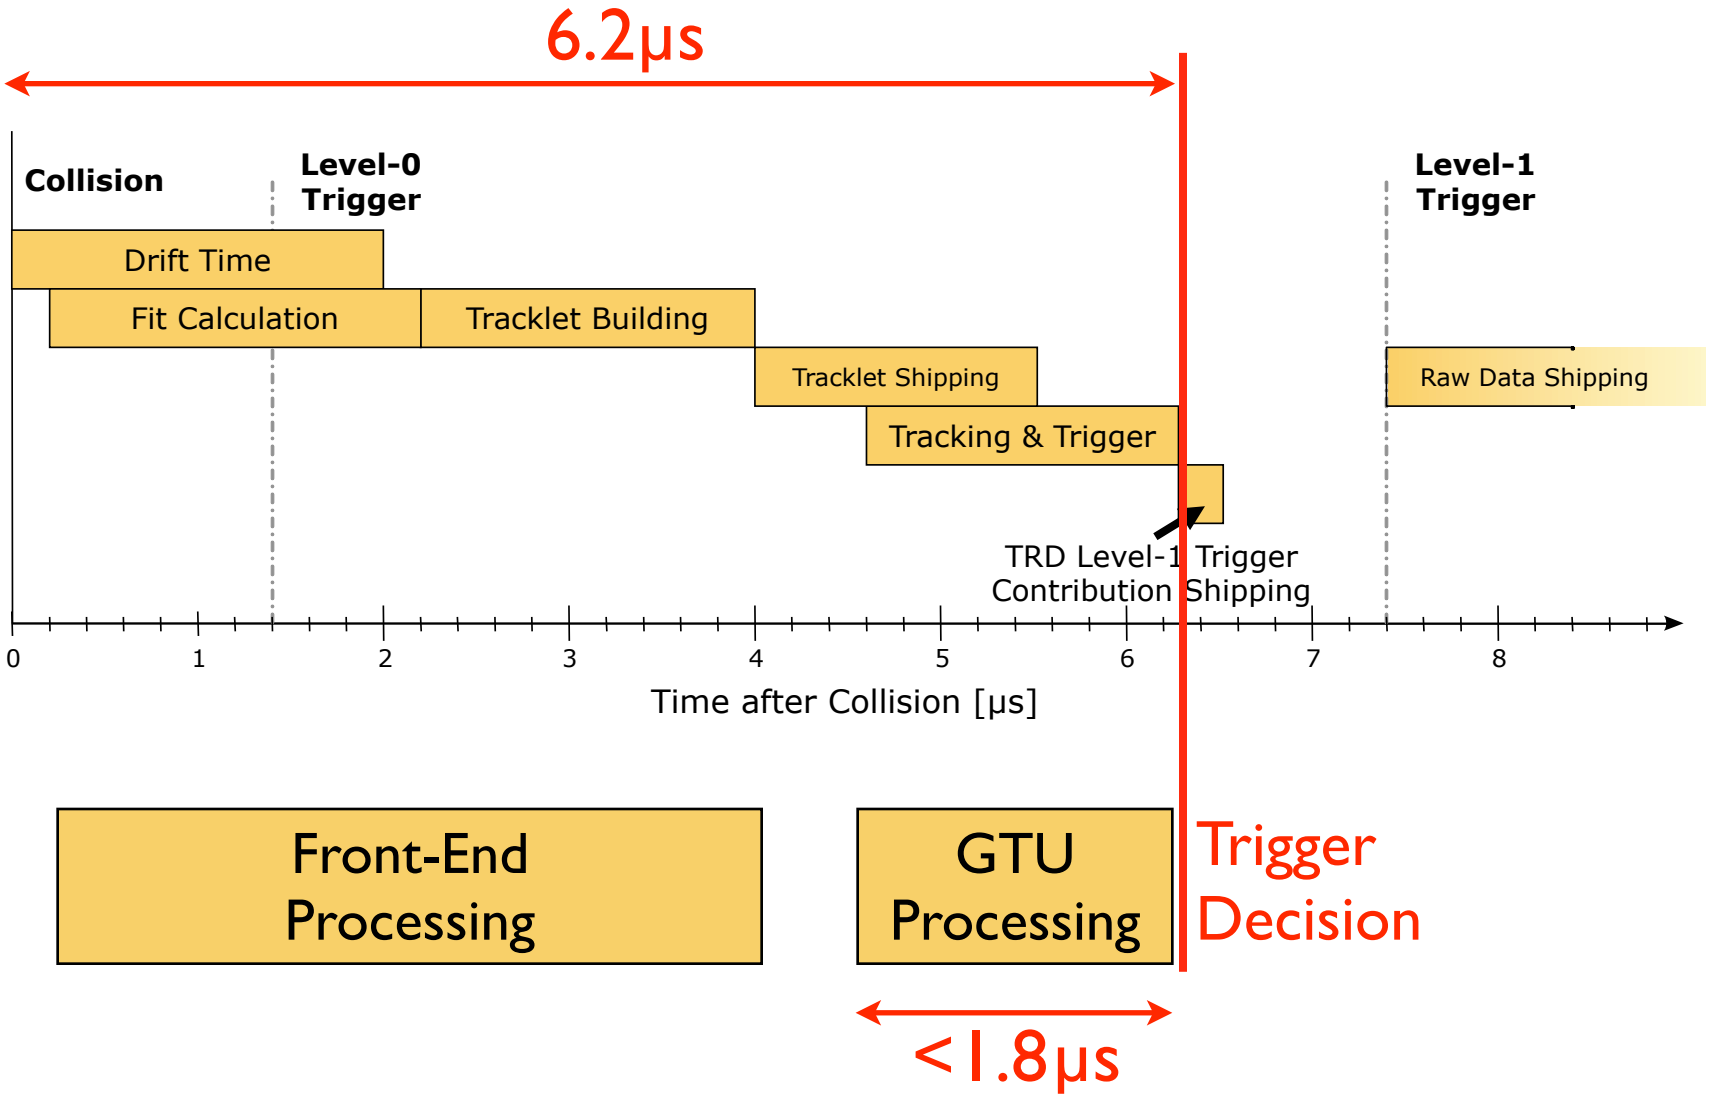

Global Tracking Unit

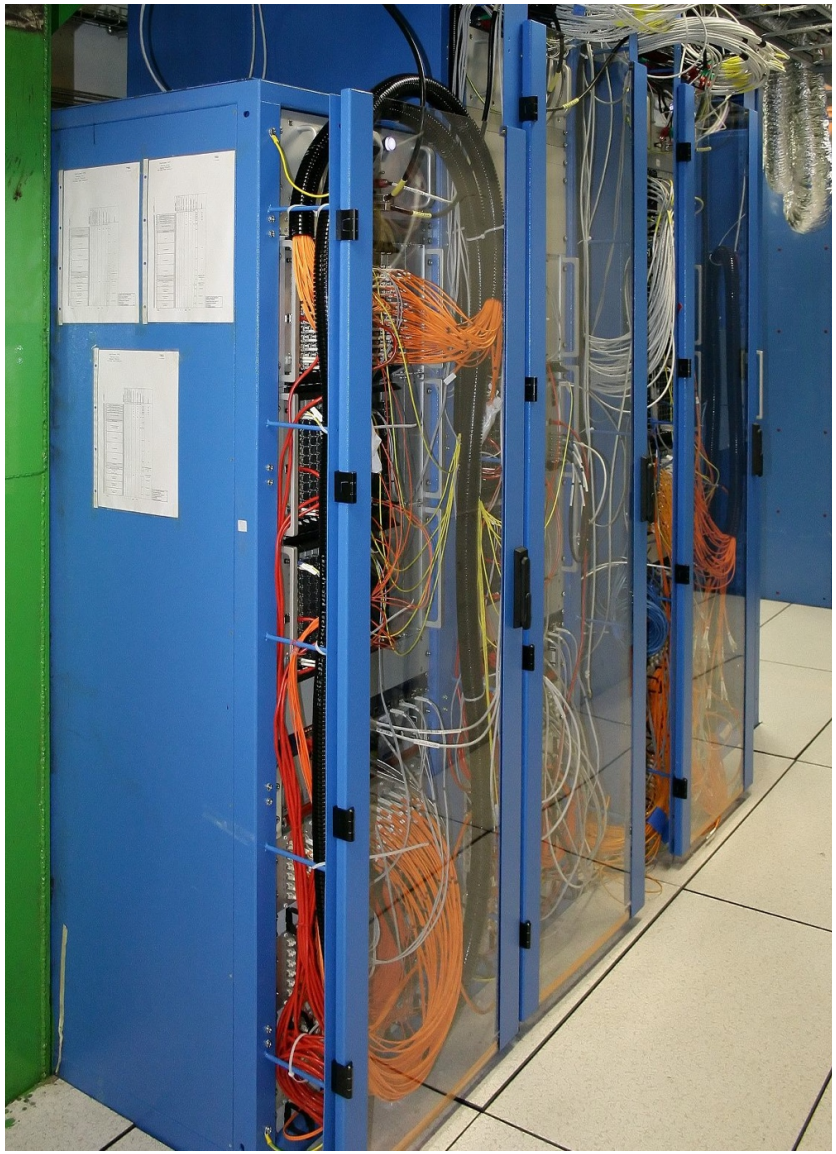

- GTU: Second Processing Stage
  - Three 19" racks outside L3 magnet
  - 109 custom PCAs with large FPGAs
- Level-1 Trigger Contribution
  - Detection & full 3D reconstruction of high-p_t tracks based on tracklets
  - Calculation of transverse momenta
  - Provides various trigger schemes: di-lepton decays (J/ψ, ϒ), jets, ...
- Raw Data Buffering
  - 2.1 Tbit/s via 1080 links from detector
  - Multi-event buffering & interface to ALICE DAQ system
  - Interlaced trigger sequences & extended error handling

3-Tier Architecture

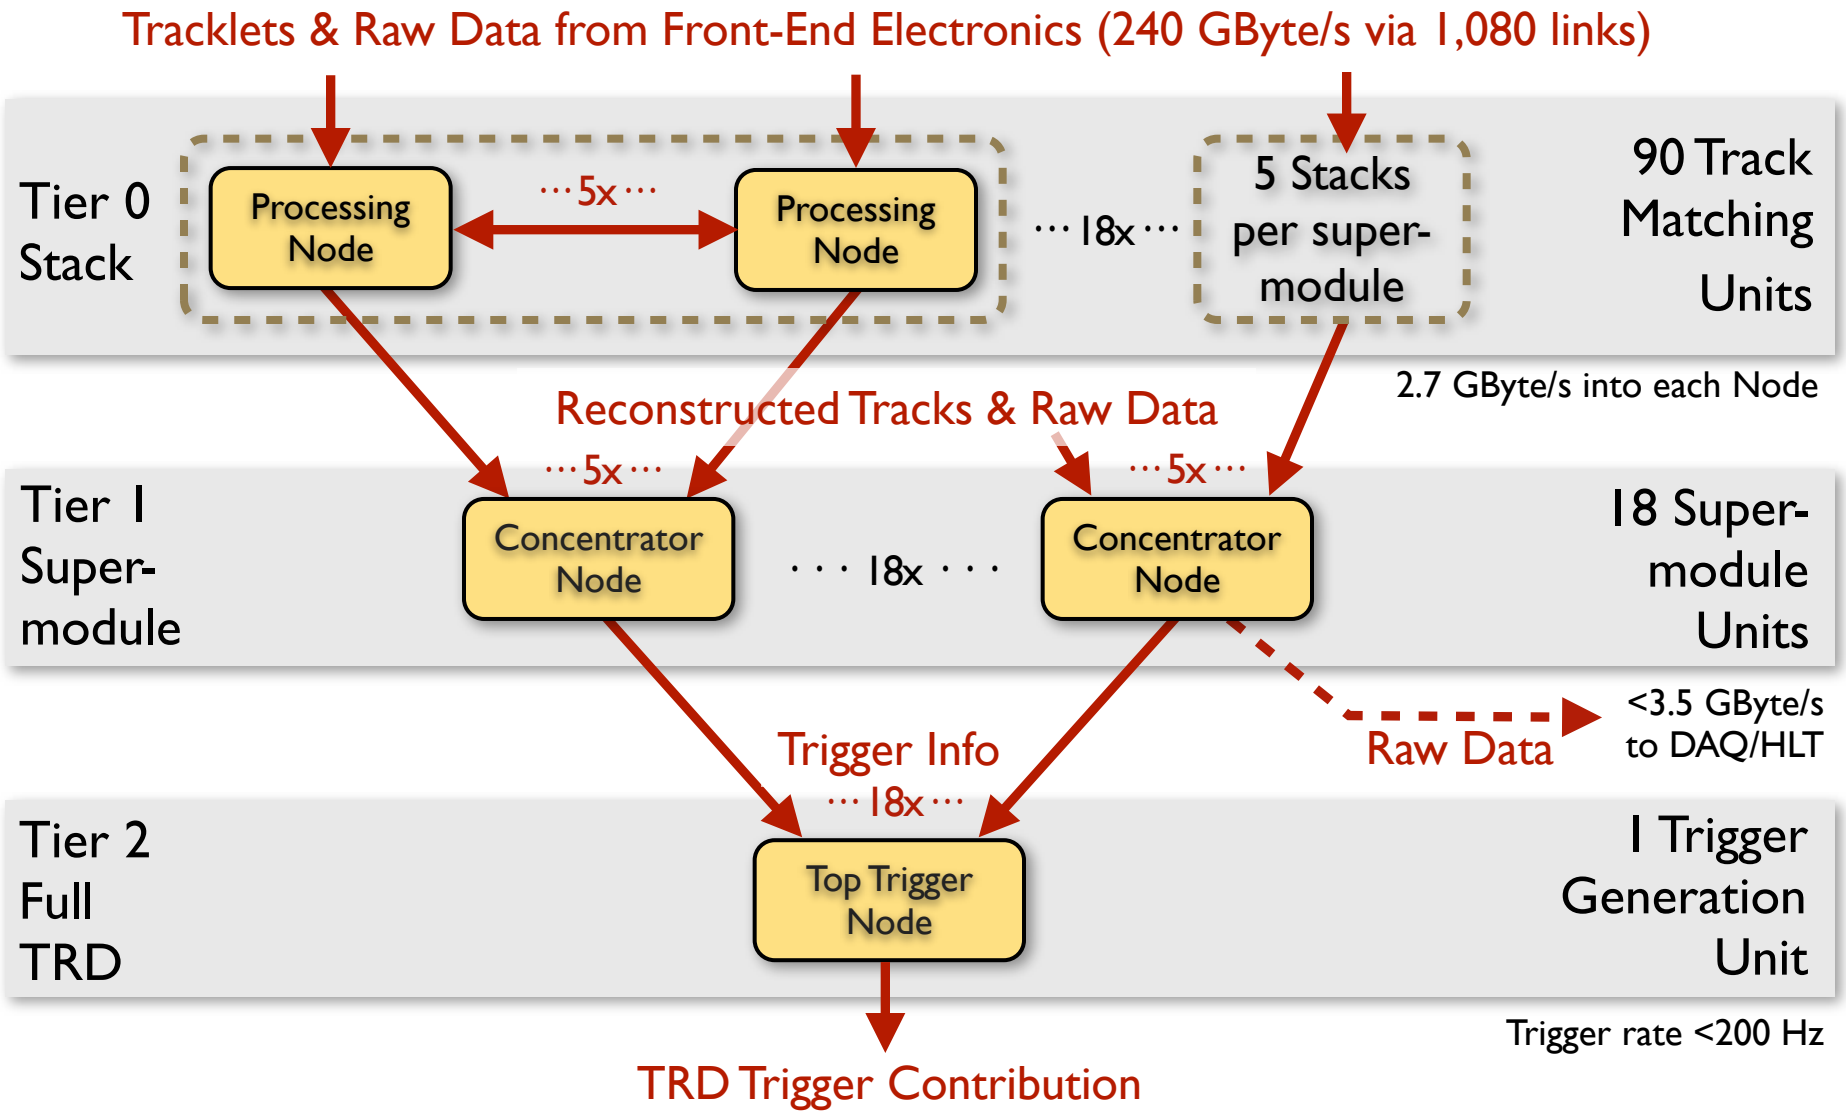

One GTU Segment

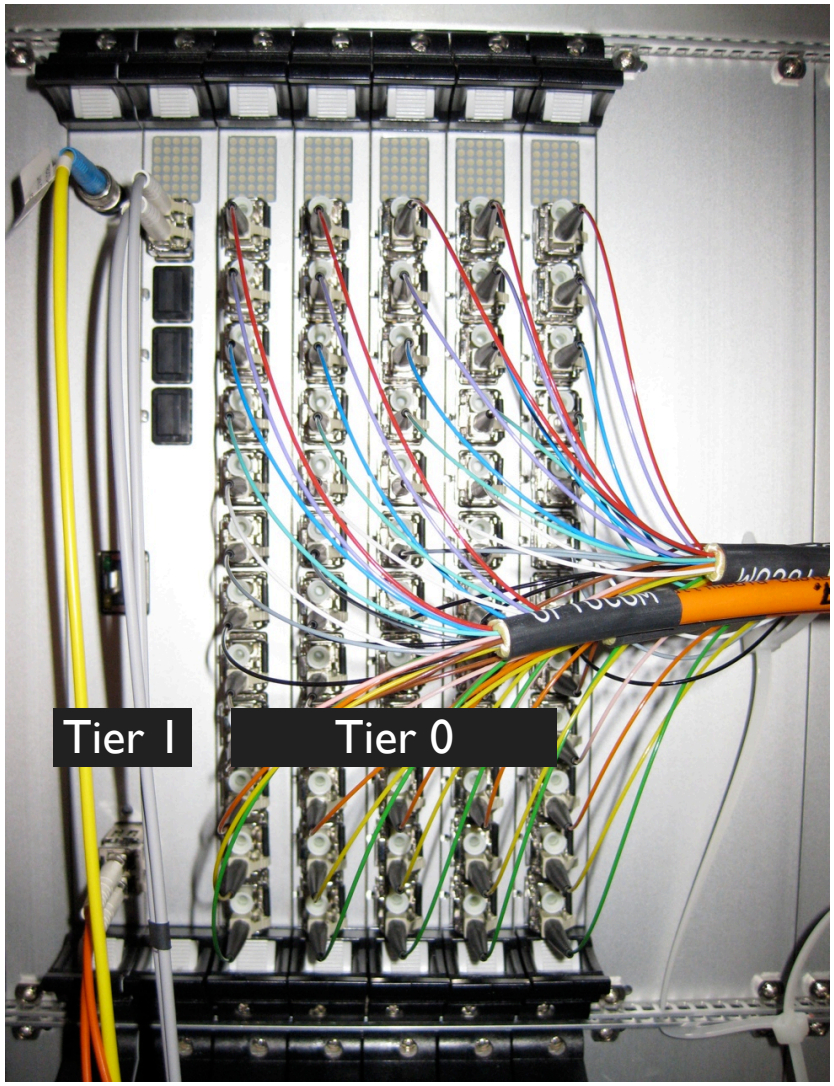

GTU segment for one TRD supermodule

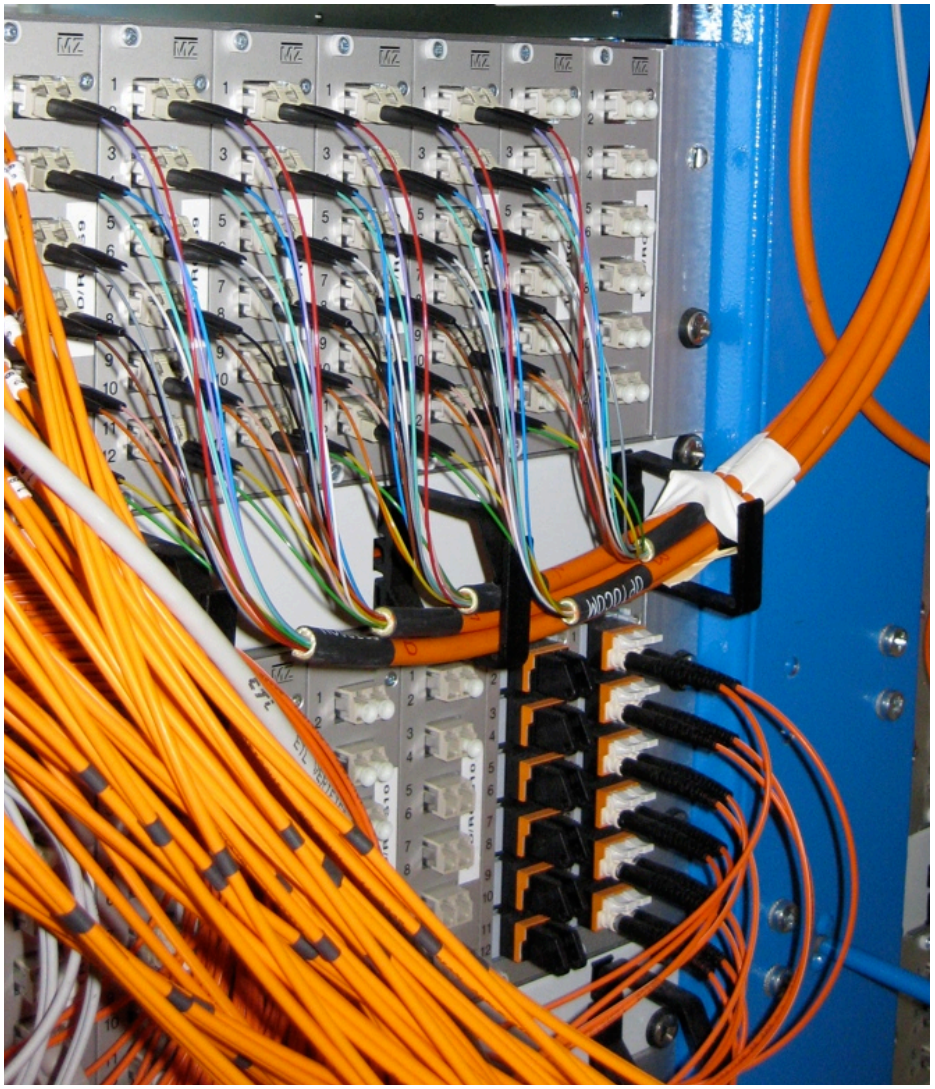

Patch panel with 60 fibres for one TRD supermodule

GTU Processing Node

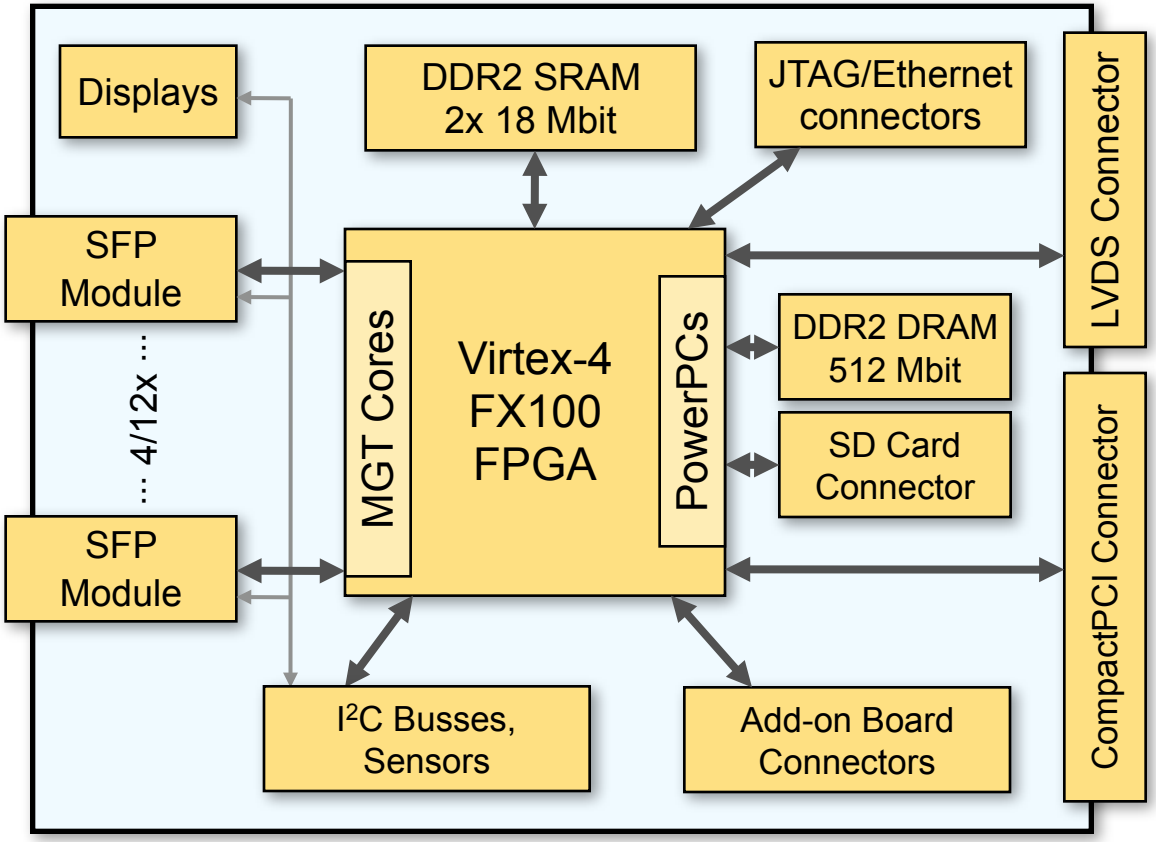

- CompactPCI card, 6U height
- 14 layer PCB
- Tier-specific assembly & add-on cards
- Virtex-4 FX100 FPGAs
- 2 Embedded PowerPC cores
- 64 MByte DRAM
- SDCard and ethernet connectors

GTU Processing Node - Tier 0

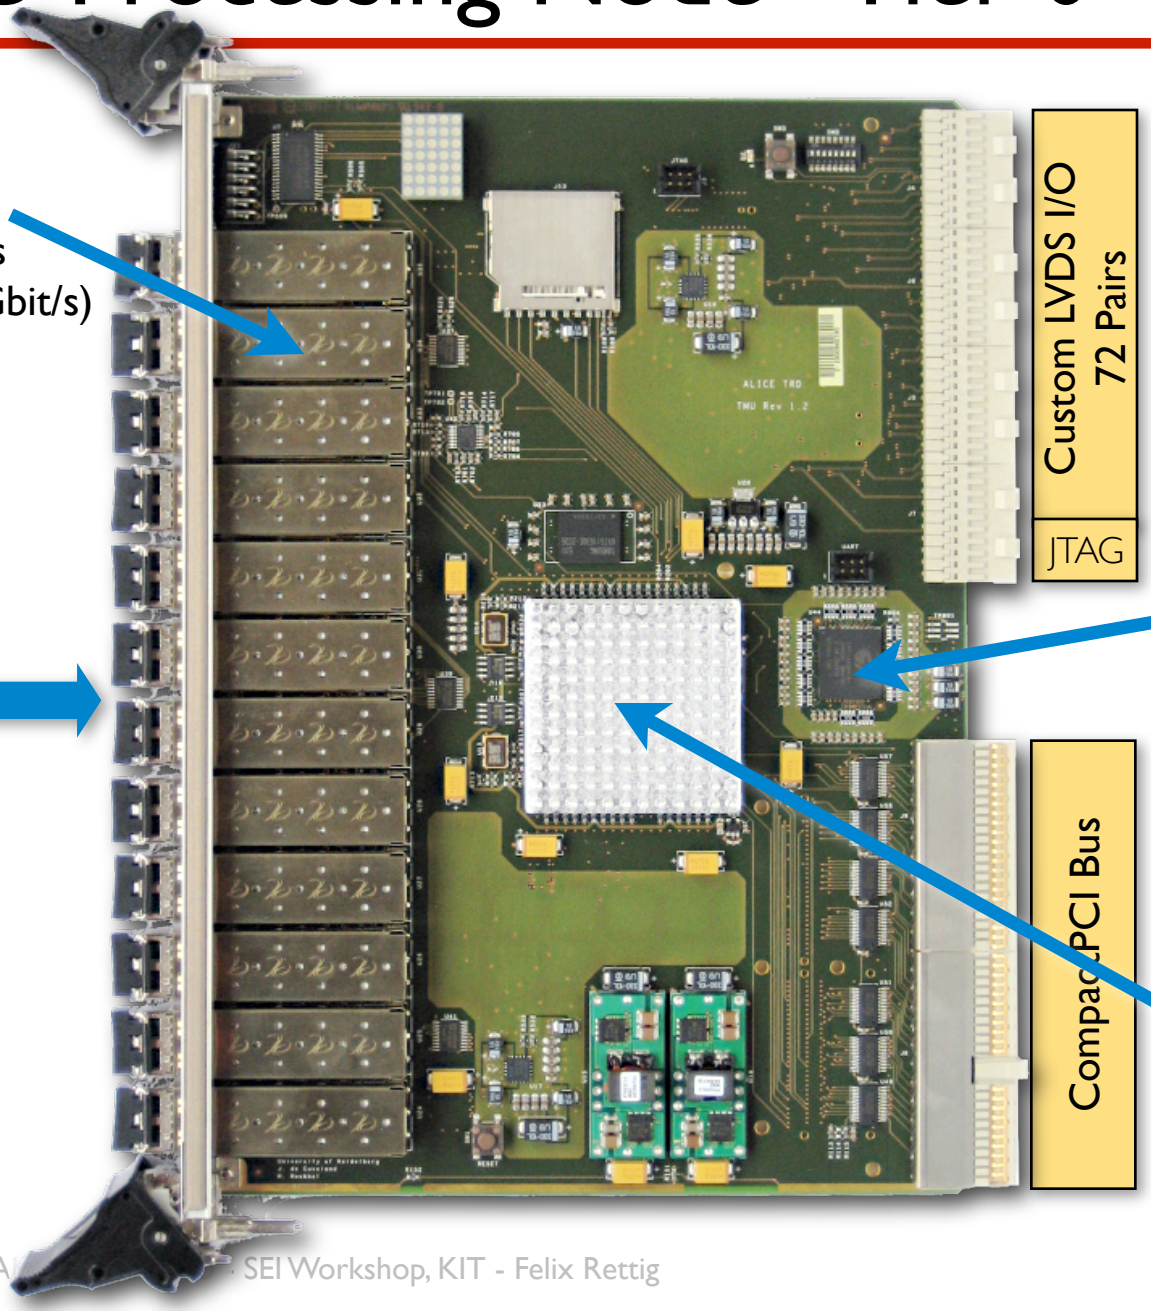

12 SFP Modules (850nm, 2.5Gbit/s)

Data from one Detector Stack

3 LVDS Links (240 MHz DDR, 8 Bit) 3.8 Gbit/s each

To tier 1

From left node

To right node

DDR2 SRAM 4 MByte High Bandwidth Data Buffer (28.8 GBit/s)

Virtex-4 FX100 FPGA 95k LCs, 768 I/Os, 20 MGTs, 2 PowerPC Cores

GTU Concentrator Board - Tier 1

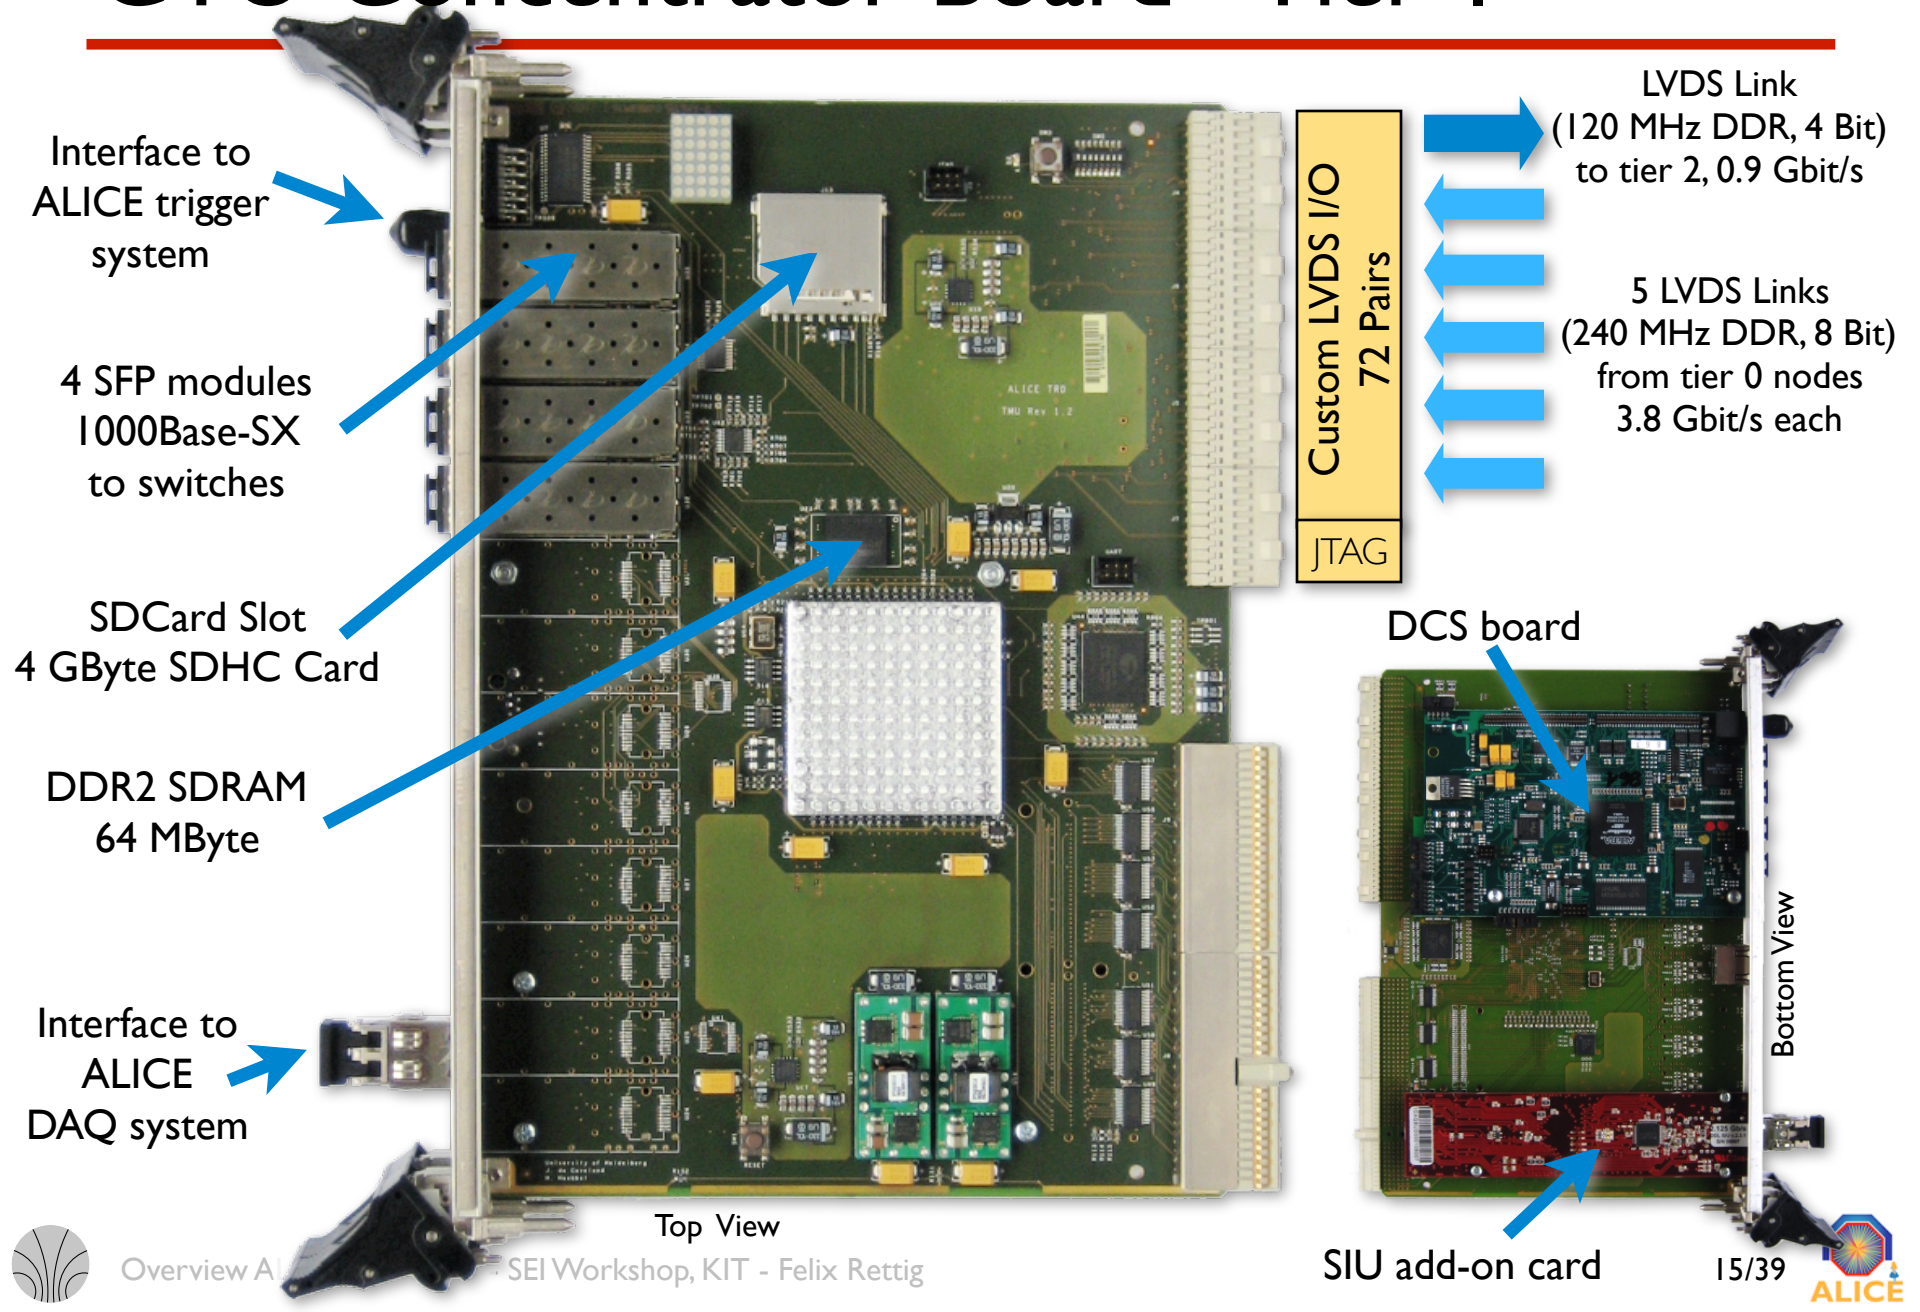

Tracking - Track Matching

- 3D track matching: find tracklets belonging to one track
- Processing time less than approx. 1.5µs

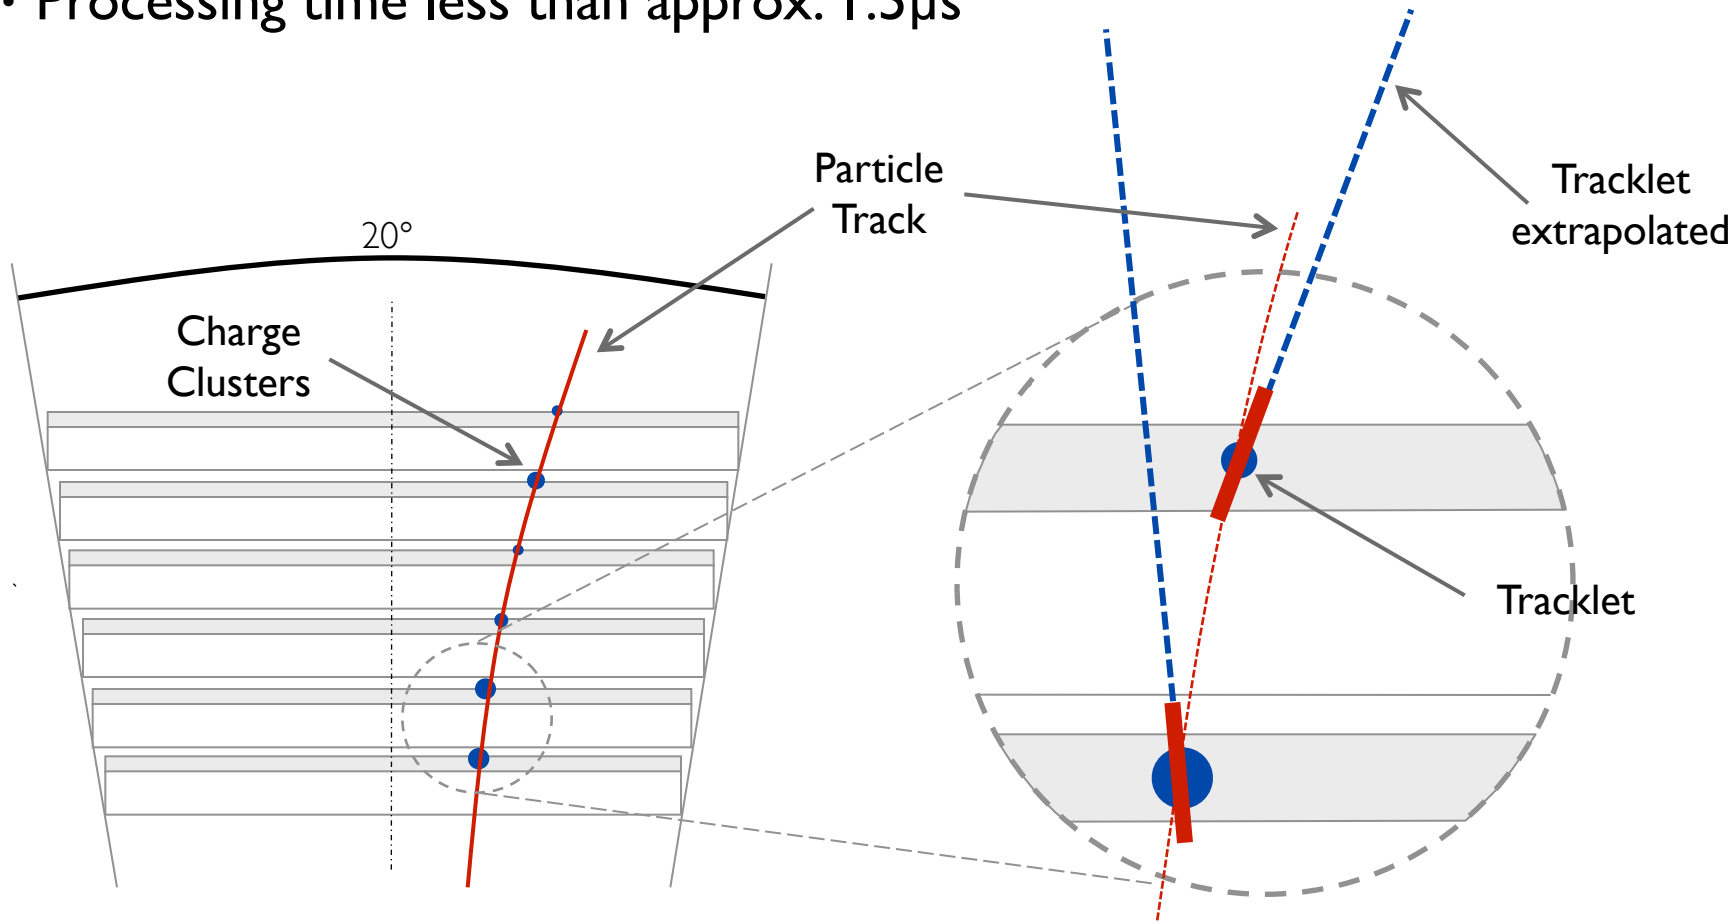

- track bendings and tracklet misorientations exaggerated -

Tracking - Track Matching II

- Projection of tracklets to virtual transverse planes
- Intelligent sliding window algorithm: Δy, $\Delta\alpha_{Vertex}$, Δz
- Track: ≥4 tracklets from different layers inside same window

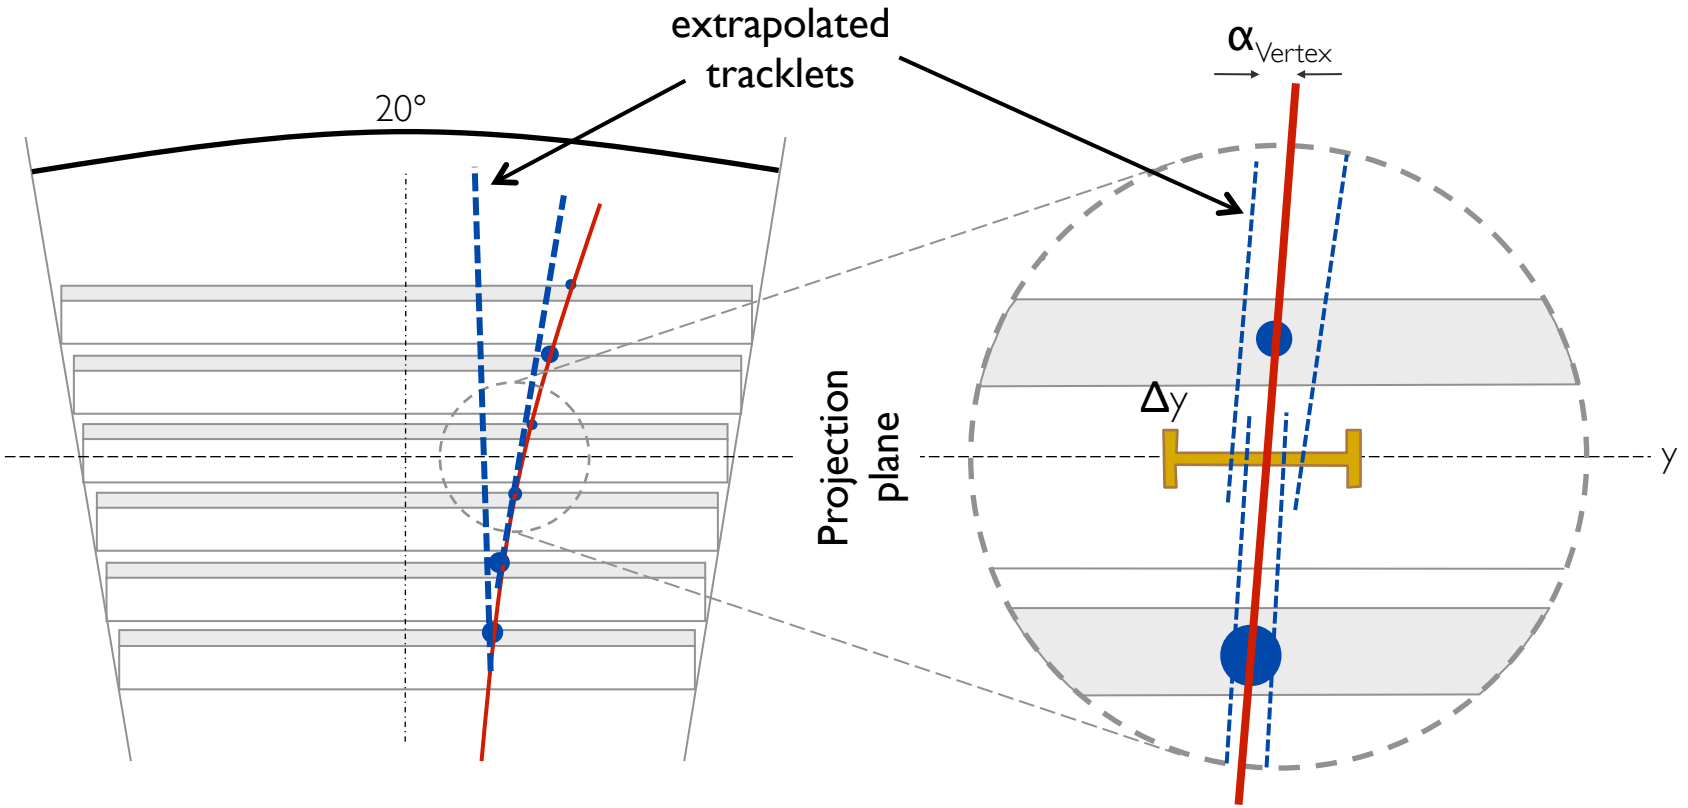

Tracking - Track Reconstruction

- Linear fit on matching tracklets: line parameter a, sum of tracklet PID
- Primary vertex assumption
- Estimation of p_t from a: $p_t = \frac{const}{a}$, $\Delta p_t / p_t < 1\%$
- Fast p_t cut decision: $const \leq |p_{t,min} \cdot a|$

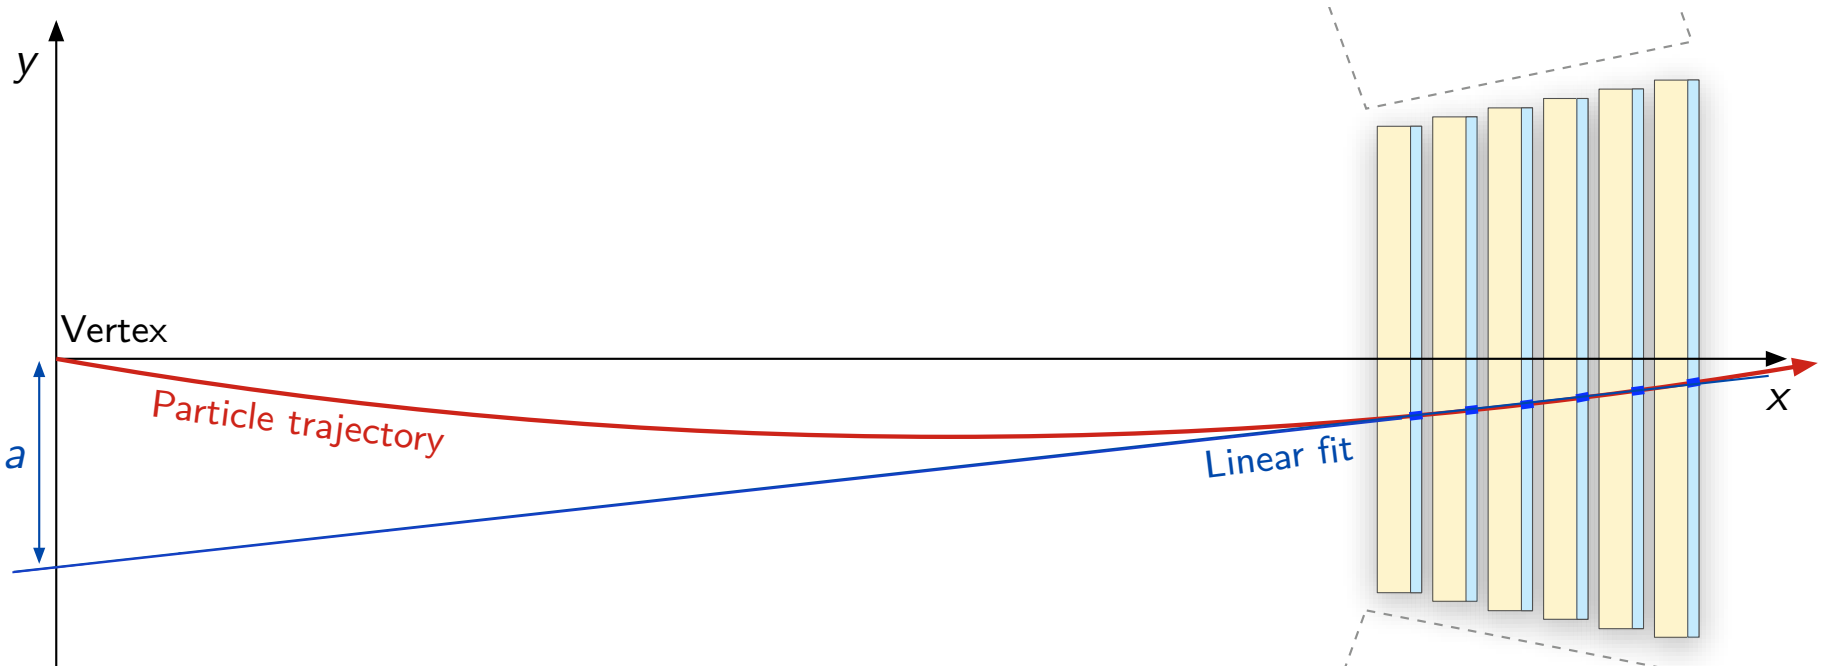

Tracking - Design Overview

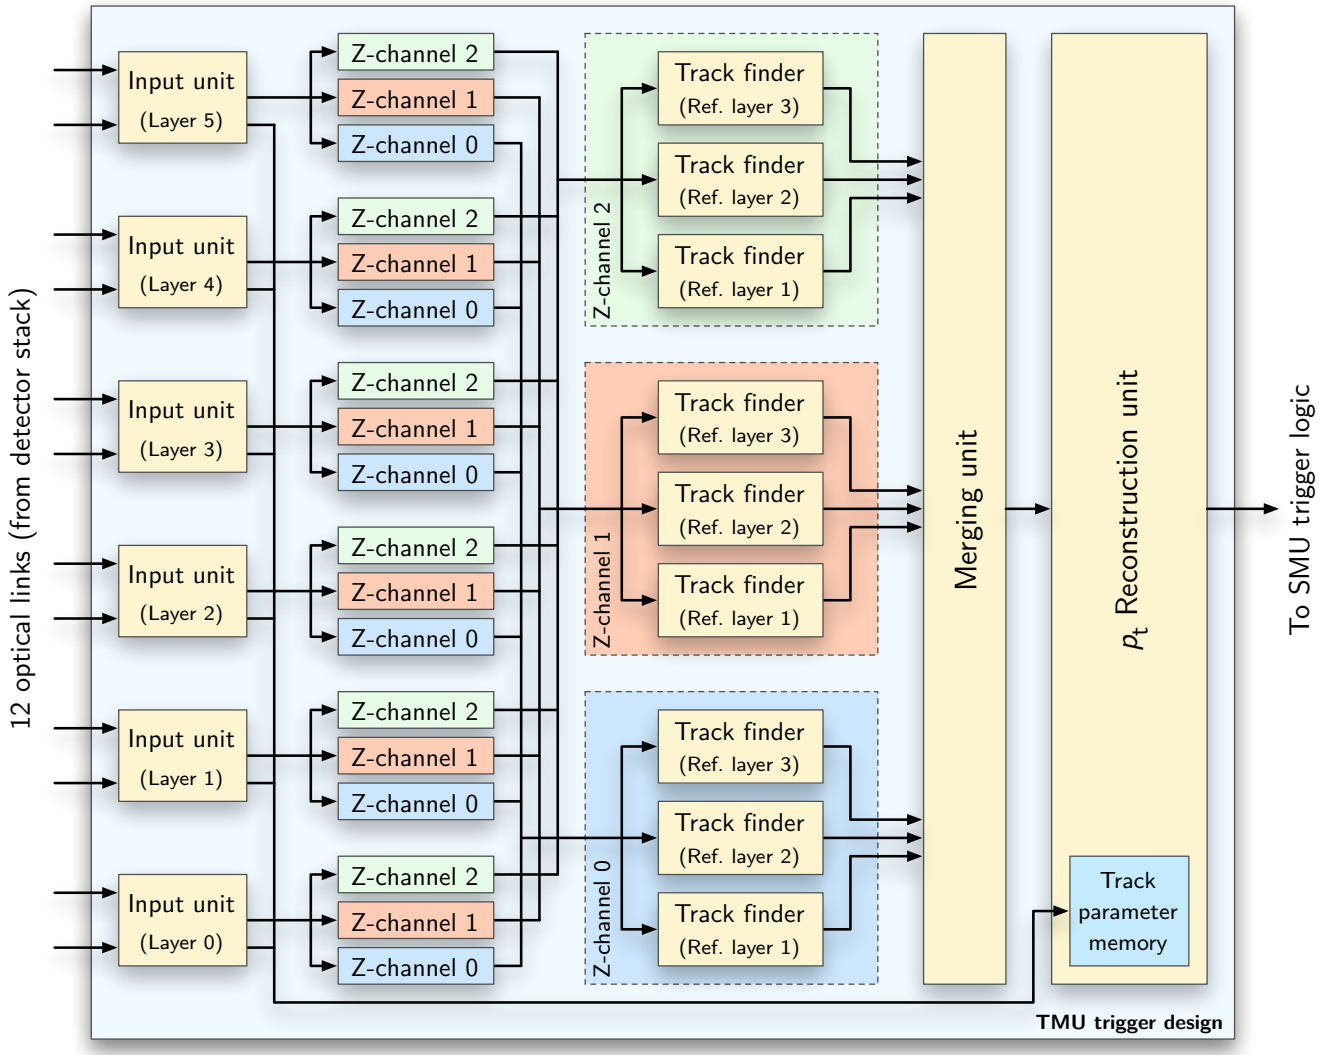

- Fully pipelined data push architecture for minimal latency
- 18 matching units running in parallel, 9 track finders
- 18 bit fixed point arithmetics, pre-computed look-up tables and DSP blocks used
- Critical path: deep combinatorial paths BRAM read data → read addr, 12 logic levels @ 60 MHz

Tracking - Momentum Resolution

November 2007 Beam Test Setup at CERN PS

Beam Test 2007 Results:

- Accelerator: CERN Proton Synchrotron
- Electrons, Pions with p_t 0.5 – 6 GeV/c
- 8 days of continuous operation, few million events
- GTU algorithm: $\Delta p_t/p_t < 1\%$
- TRD total: $\Delta p_t/p_t < 3\%$

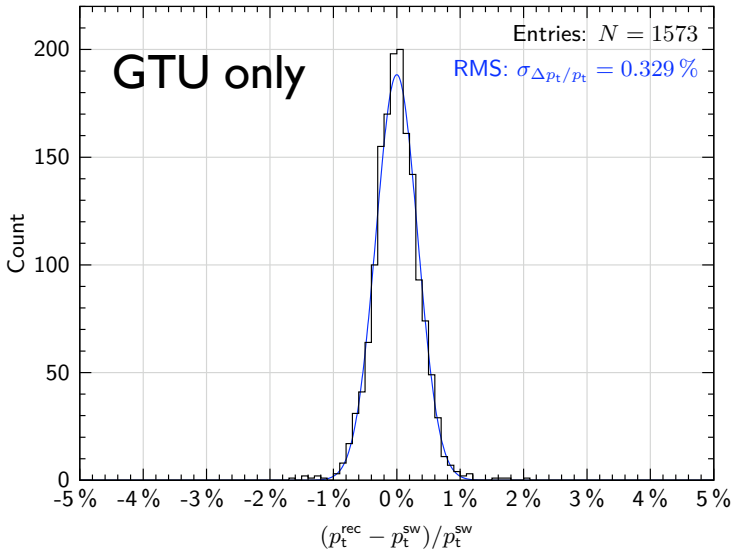

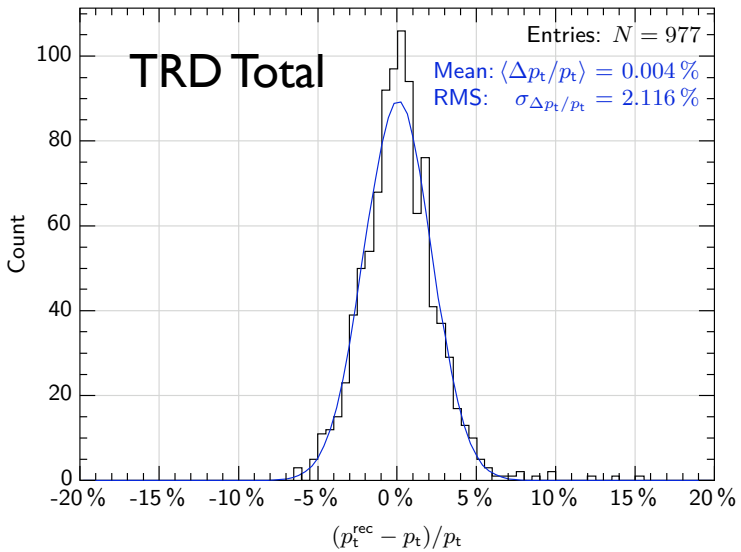

Tracking - Processing Time

- Minimum latency of about 550 ns
- Slow nearly linear rise with number of tracklets
- Total latency depending significantly on number of tracklets

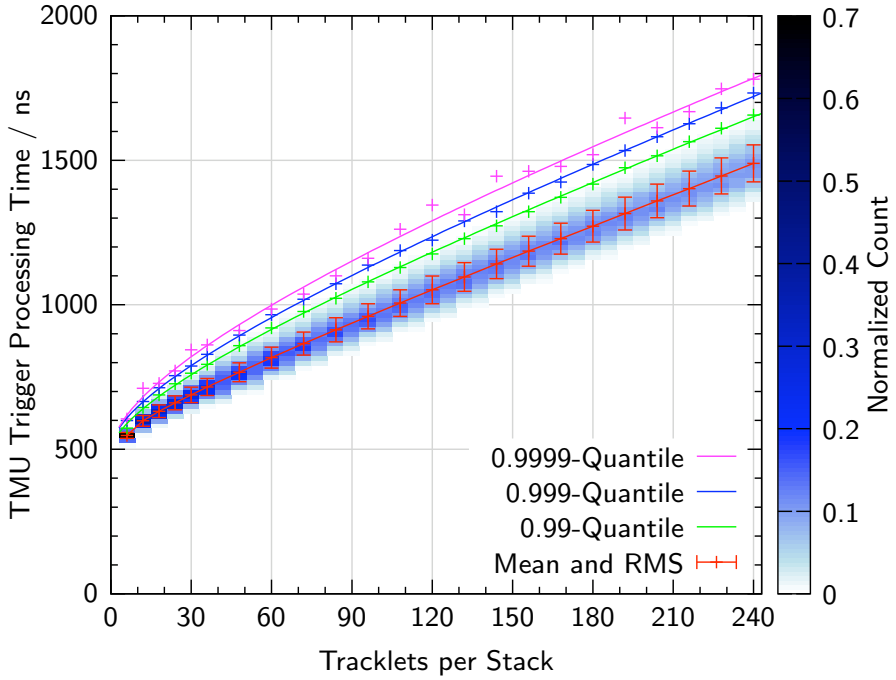

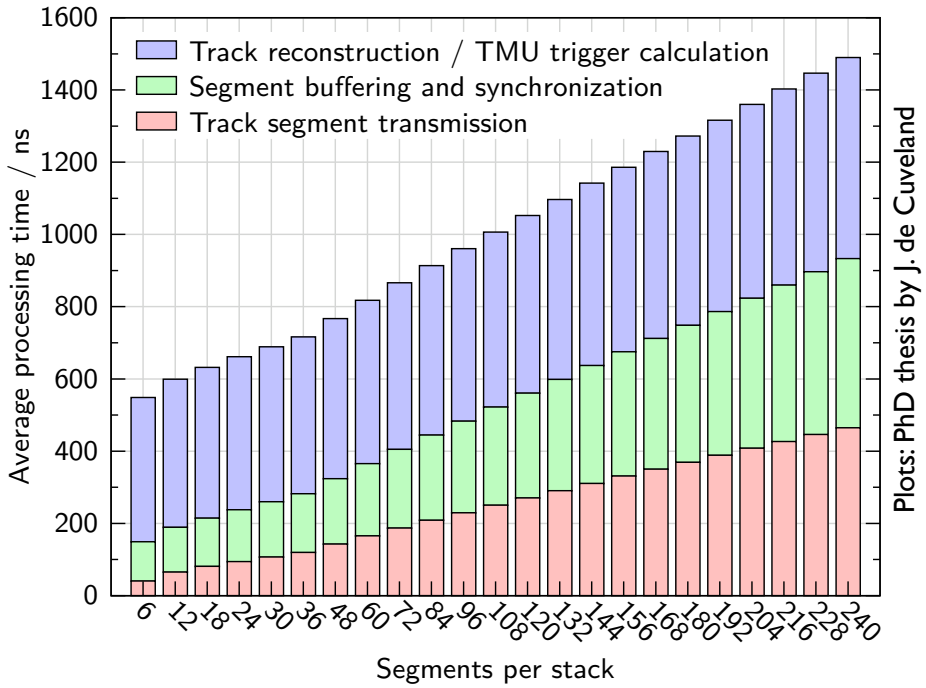

Tracking - Efficiency

TRD + GTU simulation

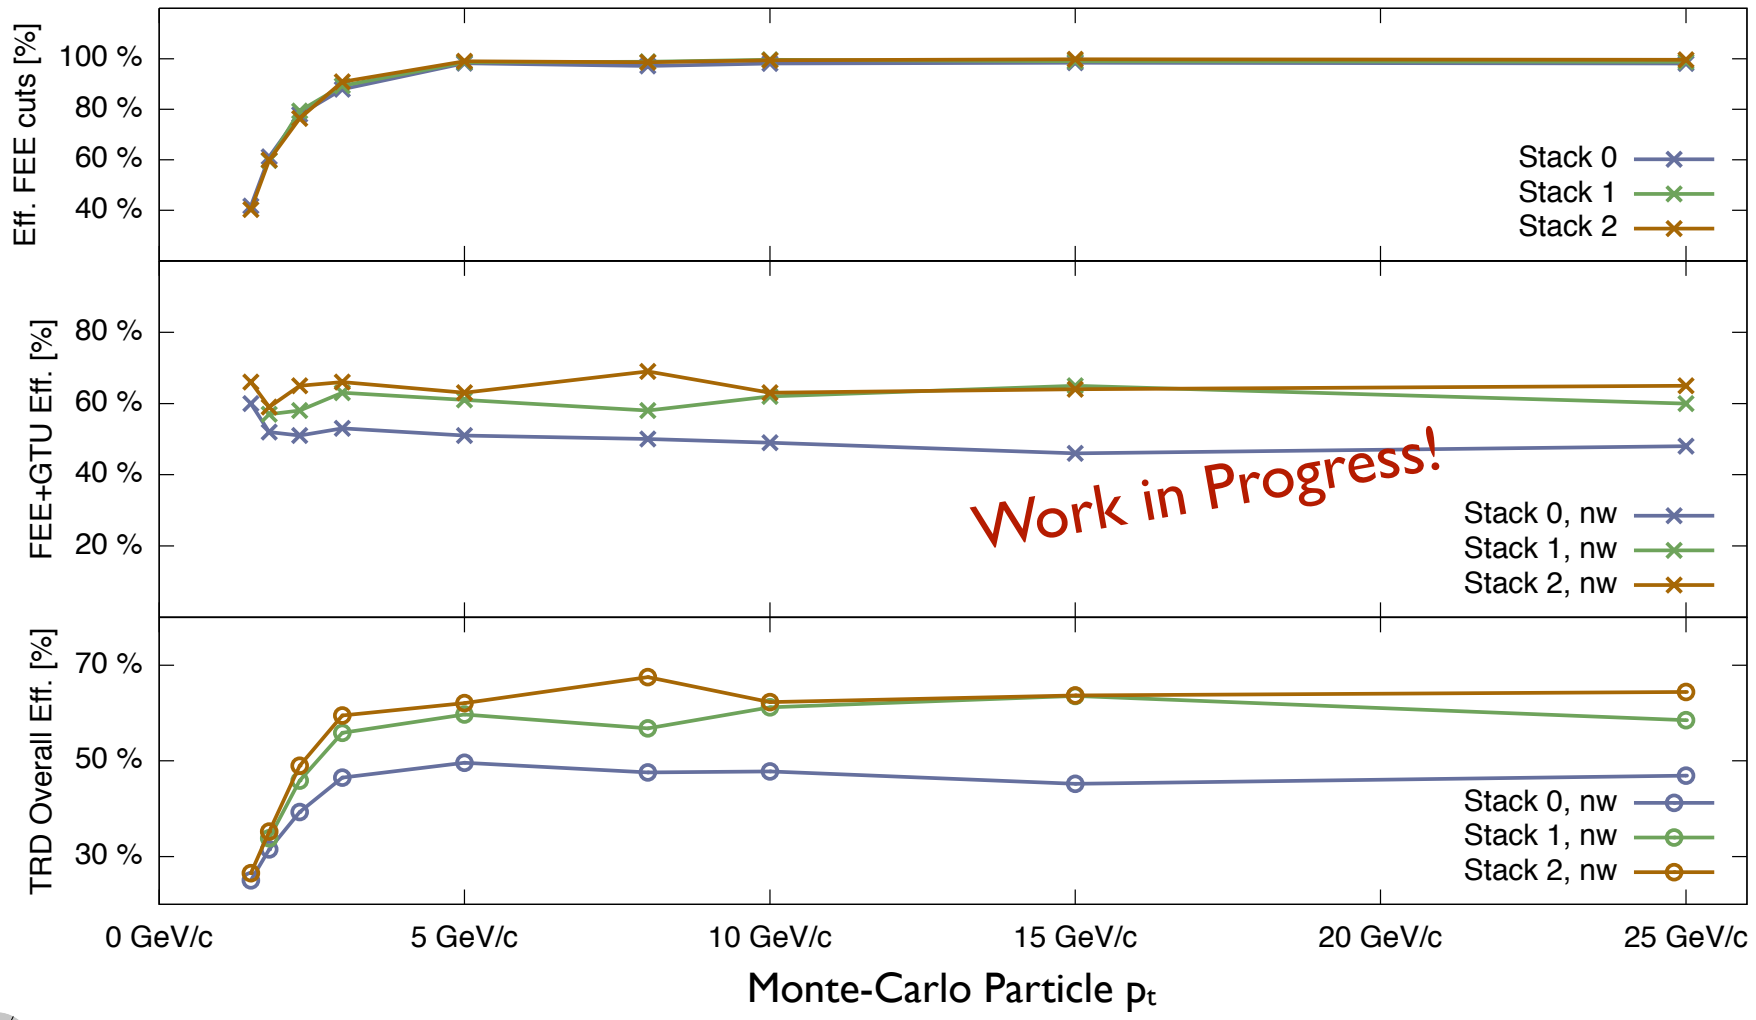

Tracking - Latest Collisions at LHC

- 7 supermodules installed, data taking with collisions
- Tracklet tuning ongoing: resolution + availability time
- Latest GTU tracking results:

```
Supermodule 01, Stack 2:
Valid L0 sequences (341787 of 341789=99%):
Tracking in time (<6us):   99% (341654/341787)
Tracking duration:         0.6us: 339928   0.7us: 1723    0.9us: 3     4.5us: 133
Tracking done after L0:    5.5us: 186686   6.0us: 154620  6.5us: 348   8.0us: 133
Tracks (num/cnt):          0: 340842   1: 933   2: 12
```

Run 124886 2010-07-02 - 17m, 341787 events, 3.5 TeV p-p, Trigger 333 Hz

Tracks per Stack	STACK0	STACK1	STACK2	STACK3	STACK4
SEGMENT 00	561	957	142	890	961
SEGMENT 01	615	863	760	636	1011
SEGMENT 07	488	576	262	96	504
SEGMENT 08	593	481	441	570	442
SEGMENT 09	475	607	328	692	817
SEGMENT 10	910	937	710	554	906
SEGMENT 17	865	373	540	192	552
TRD total	21307 tracks				

- 21,307 tracks in 341,787 events (6%)
  21 GTU tracks/s

Triggers - Scopes

- Tracking in tier 0
  - transverse momentum
  - y & z position
  - particle type

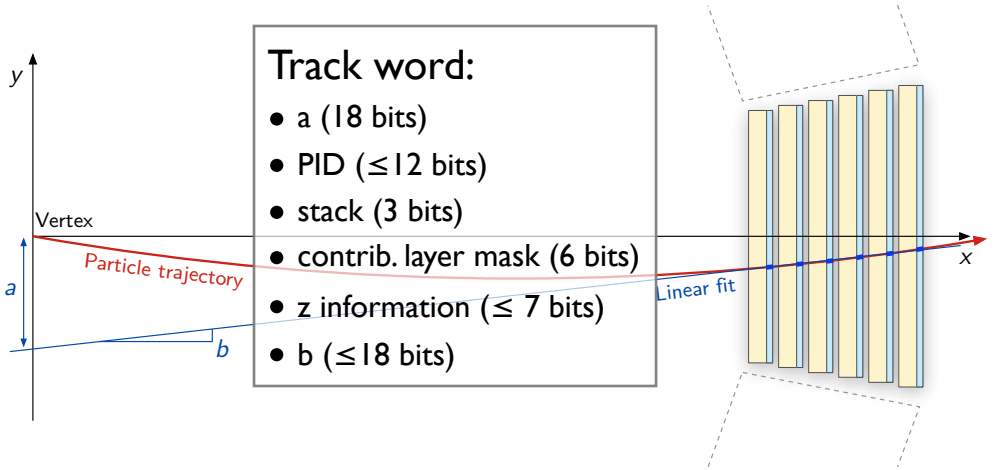

- Trigger in tier 1 - segment level
  - single high-p_t trigger
  - jet trigger (full z coverage)

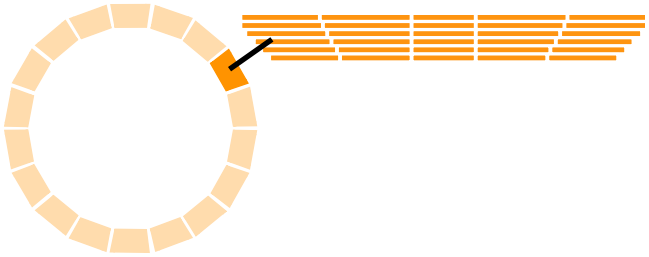

- Trigger in tier 2 - detector level
  - jet trigger (full Φ coverage), multi-jets
  - di-lepton decay trigger

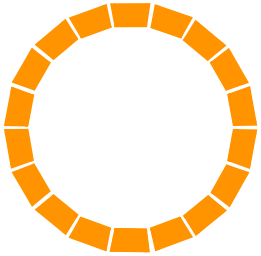

Triggers - Cosmics

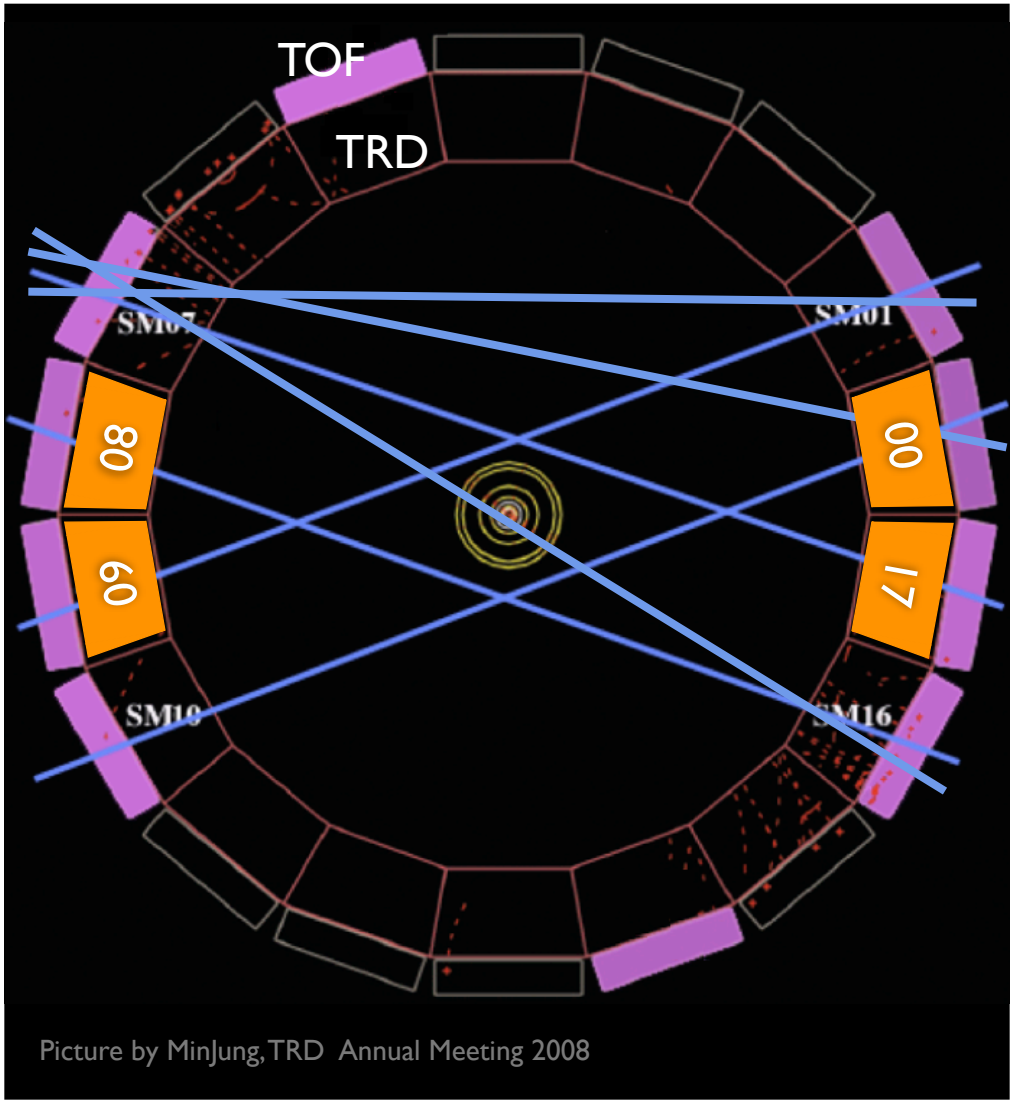

Picture by MinJung, TRD Annual Meeting 2008

Trigger operating only on tracklets, without tracking

- 12/2007-10/2008
- First L1 trigger running in ALICE
- 4 TRD supermodules
- L0 triggers by TOF or random pulser
- Purity: >93 %
- $L1/L0_{TOF}$ ratio: ~1/20
- L1 rate: 0.05 - 1 Hz
- 55,000 events taken

Triggers - Cosmic Example

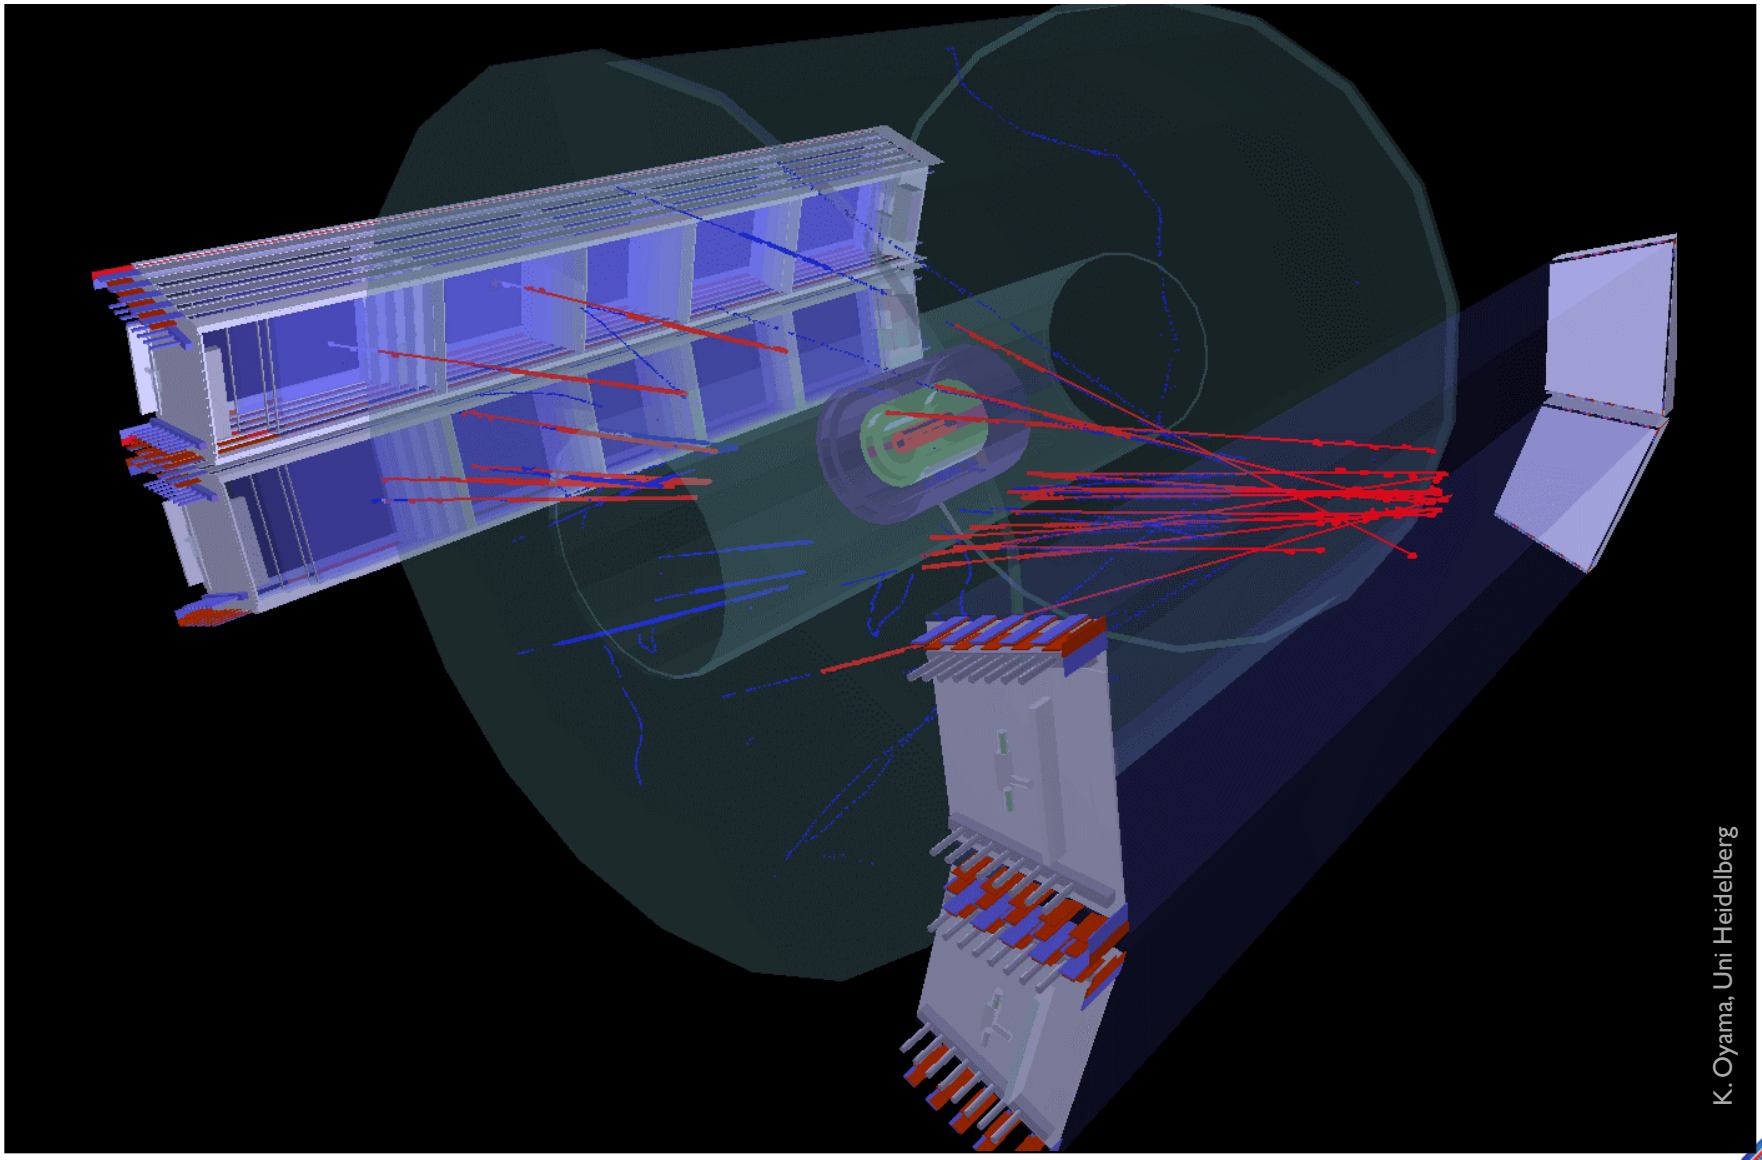

Triggers - Jet Trigger

- Consider tracks within fixed geometric regions
- Threshold conditions:
  - Number of tracks above p_t threshold
  - Sum of momenta for those tracks
- Variations:
  N tracks above p_{t1} and M tracks above p_{t2}, ...

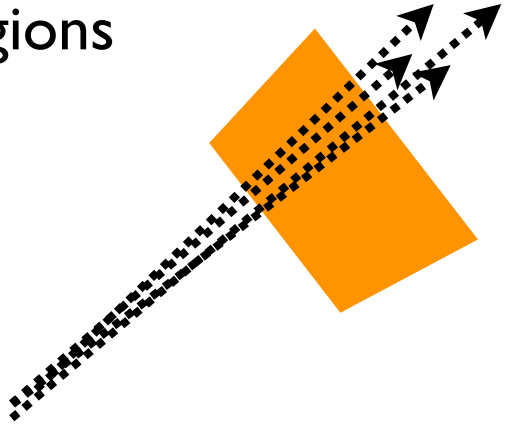

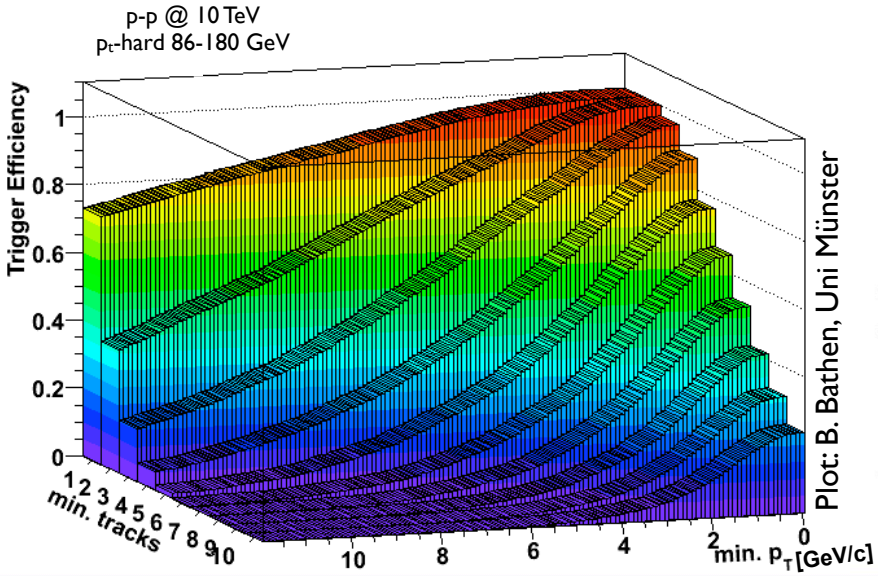

- Tier 1: jet detection
  - overlapping areas in z-direction

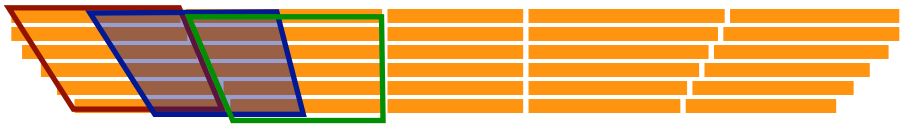

Tier 2: multi-jet coincidence

- jet detection with overlapping areas in Φ-direction

Triggers - Di-Lepton Decay

- Simple version:
  detect e^+ & e^- with p_t > threshold
- Advanced version:
  find e^+e^- pairs with invariant mass
  within certain range (J/ψ, ϒ, ...)
- Huge combinatorics for Pb-Pb collisions
- Current study:
  - Pre-selection of relevant track candidates feasible?
  - Application of sliding window algorithms
  - Massively parallelized invariant mass calculation with full utilization of computing resources provided by the Virtex-4 FPGA

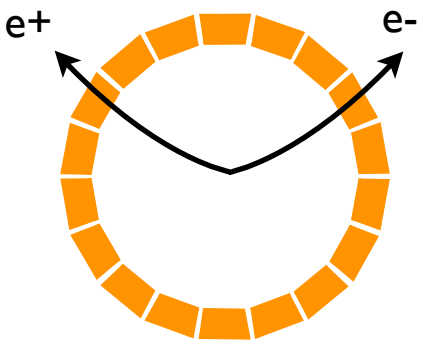

Multi-Event Buffering

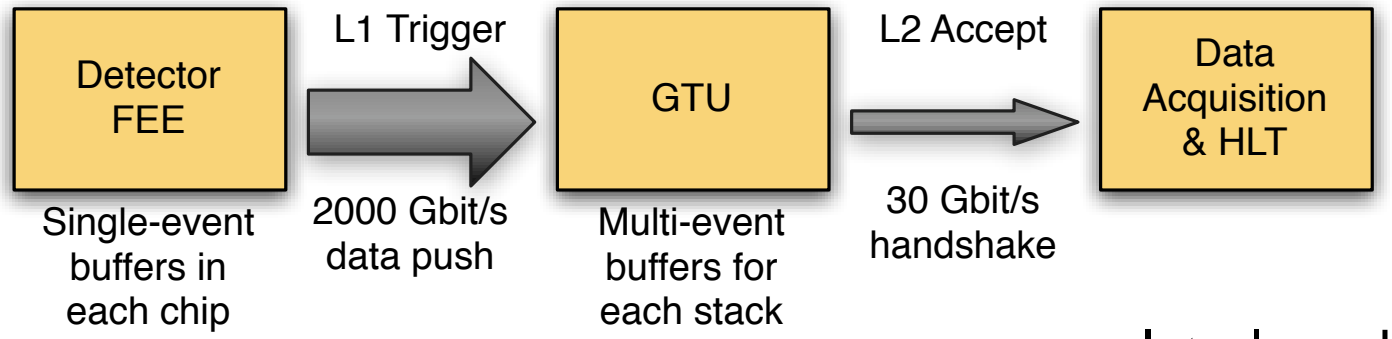

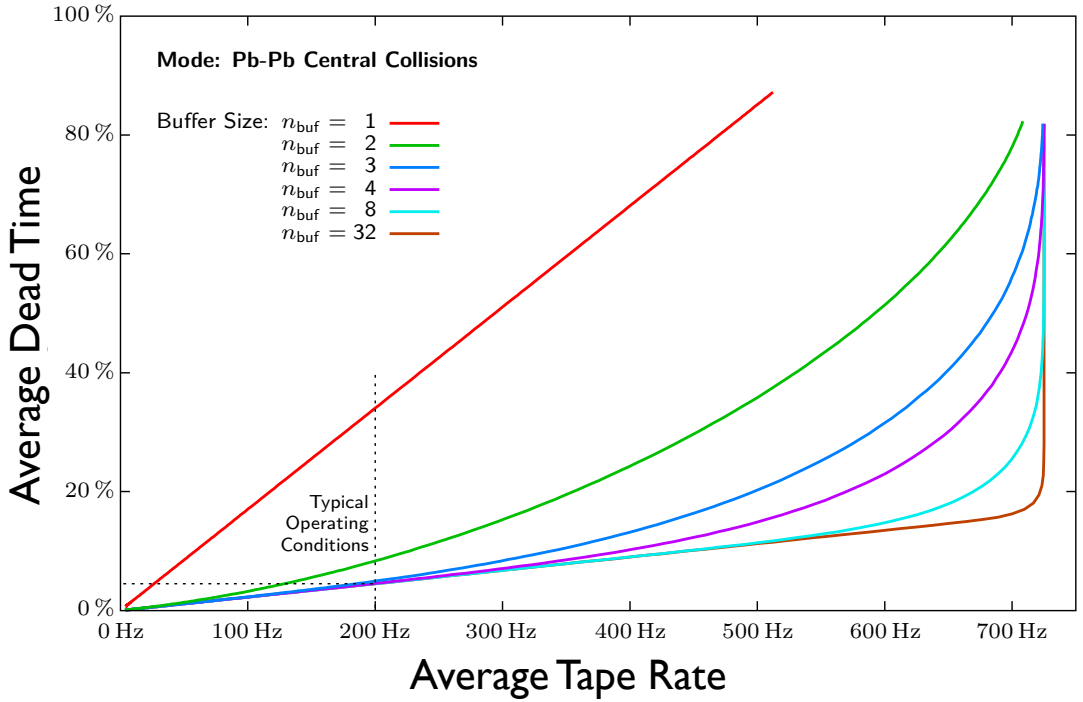

- Interleaved 3-level trigger sequences
- Single-event buffering in 65,564 FEE chips
- Multi-event buffering in 90 GTU boards
- Data taking decoupled from 2nd level decision & readout
  → reduction of dead time

Multi-Event Buffering - Tier 0

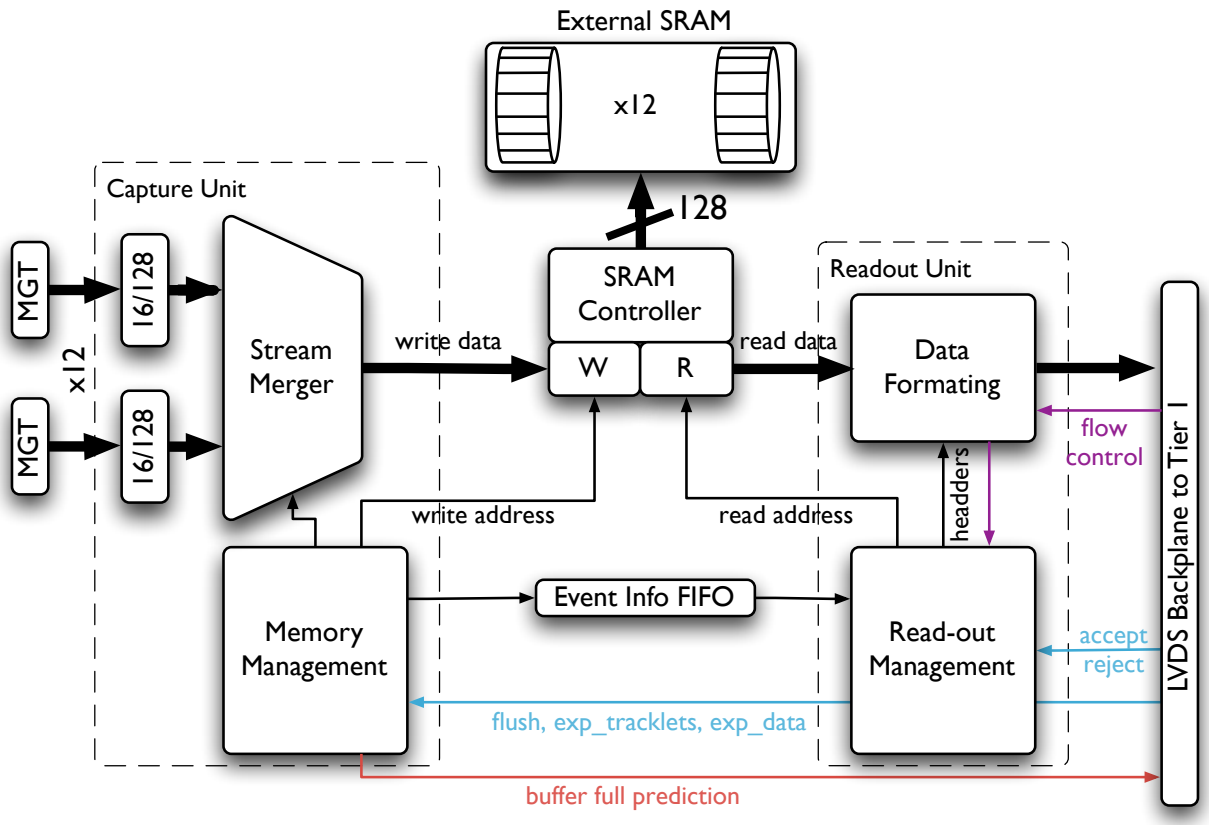

- 16→128 bit collation & 125→200 MHz crossing via dual-port BRAMs
- Wide data streams at high frequencies → many pipeline stages
- 128-bit wide 12:1 muxs at 200 MHz needed
- Dedicated high-performance SRAM controller with optimal write/read arbitration
- Storage of 2 write and 1 read pointers for each link in one BRAM
- Wide counters and arithmetics at 200 MHz

Multi-Event Buffering - Tier 1

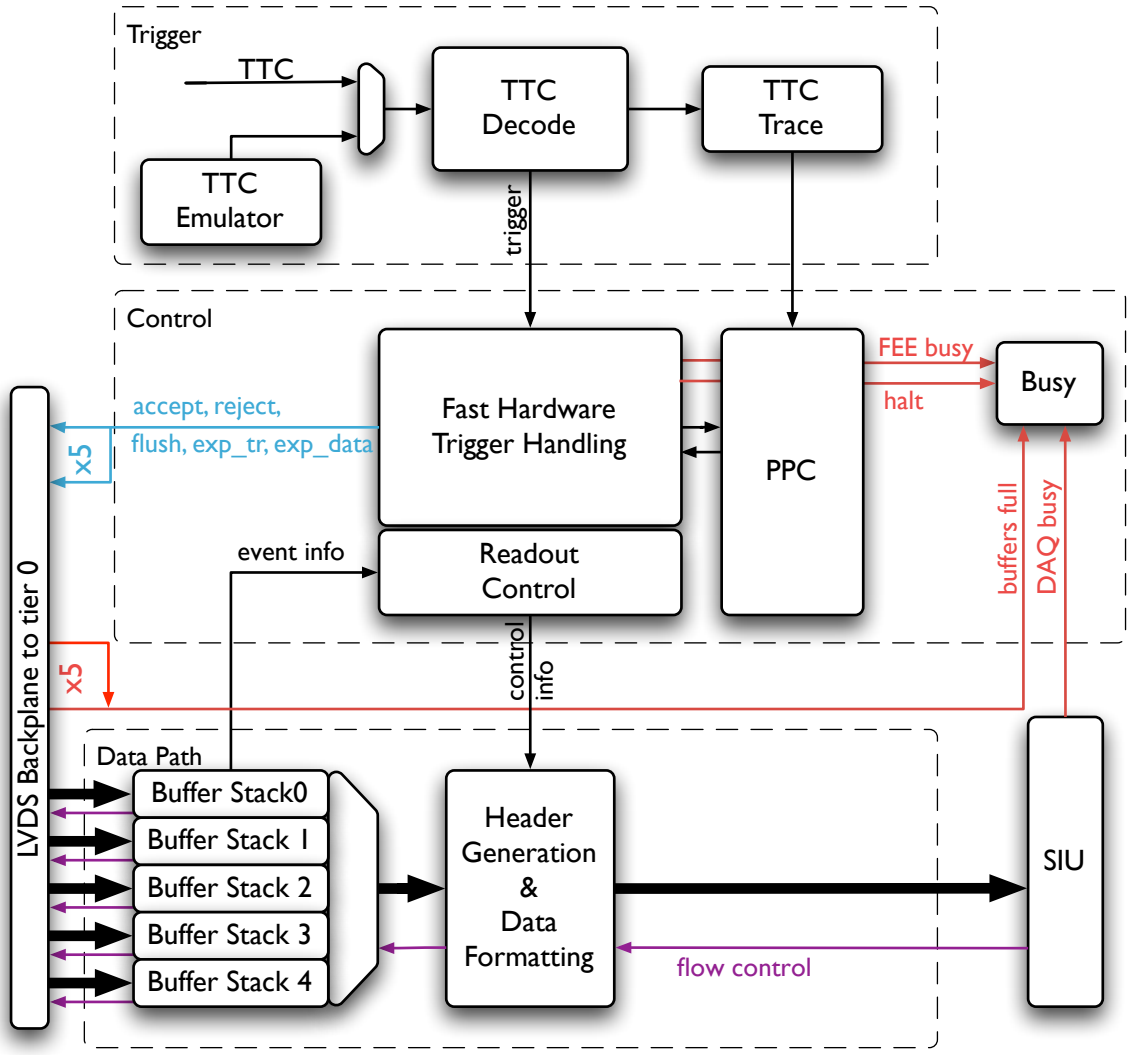

- Control & read-out 5 tier 0 buffers
- Single-Event Buffering used in ALICE since 2007
- Multi-Event Buffering under development: HW/SW co-design:
  - fast control in normal operation in fabric
  - complex error recovery in software
- Low latency PPC needed

Dual PowerPC System

- 109 Virtex-4 FX100 → 218 embedded PowerPC cores
- High-Level PowerPC
  - Running Linux for system control purposes
  - Xilinx 2.6.30 kernel + BusyBox (later Gentoo)
  - Tier 1 & 2: optical gigabit ethernet with MGT/EMAC planned, tier 1 as "switch" for tier 0 using PPP
- Low-Level PowerPC
  - Real time operations with tight time requirements
  - Extended multi-event buffering control
  - Low-level monitoring, statistics gathering
  - Designed for minimum latency & resource usage
- Interconnection between PowerPCs
  - Shared BRAM memory with hardware mutex support

PowerPCs System

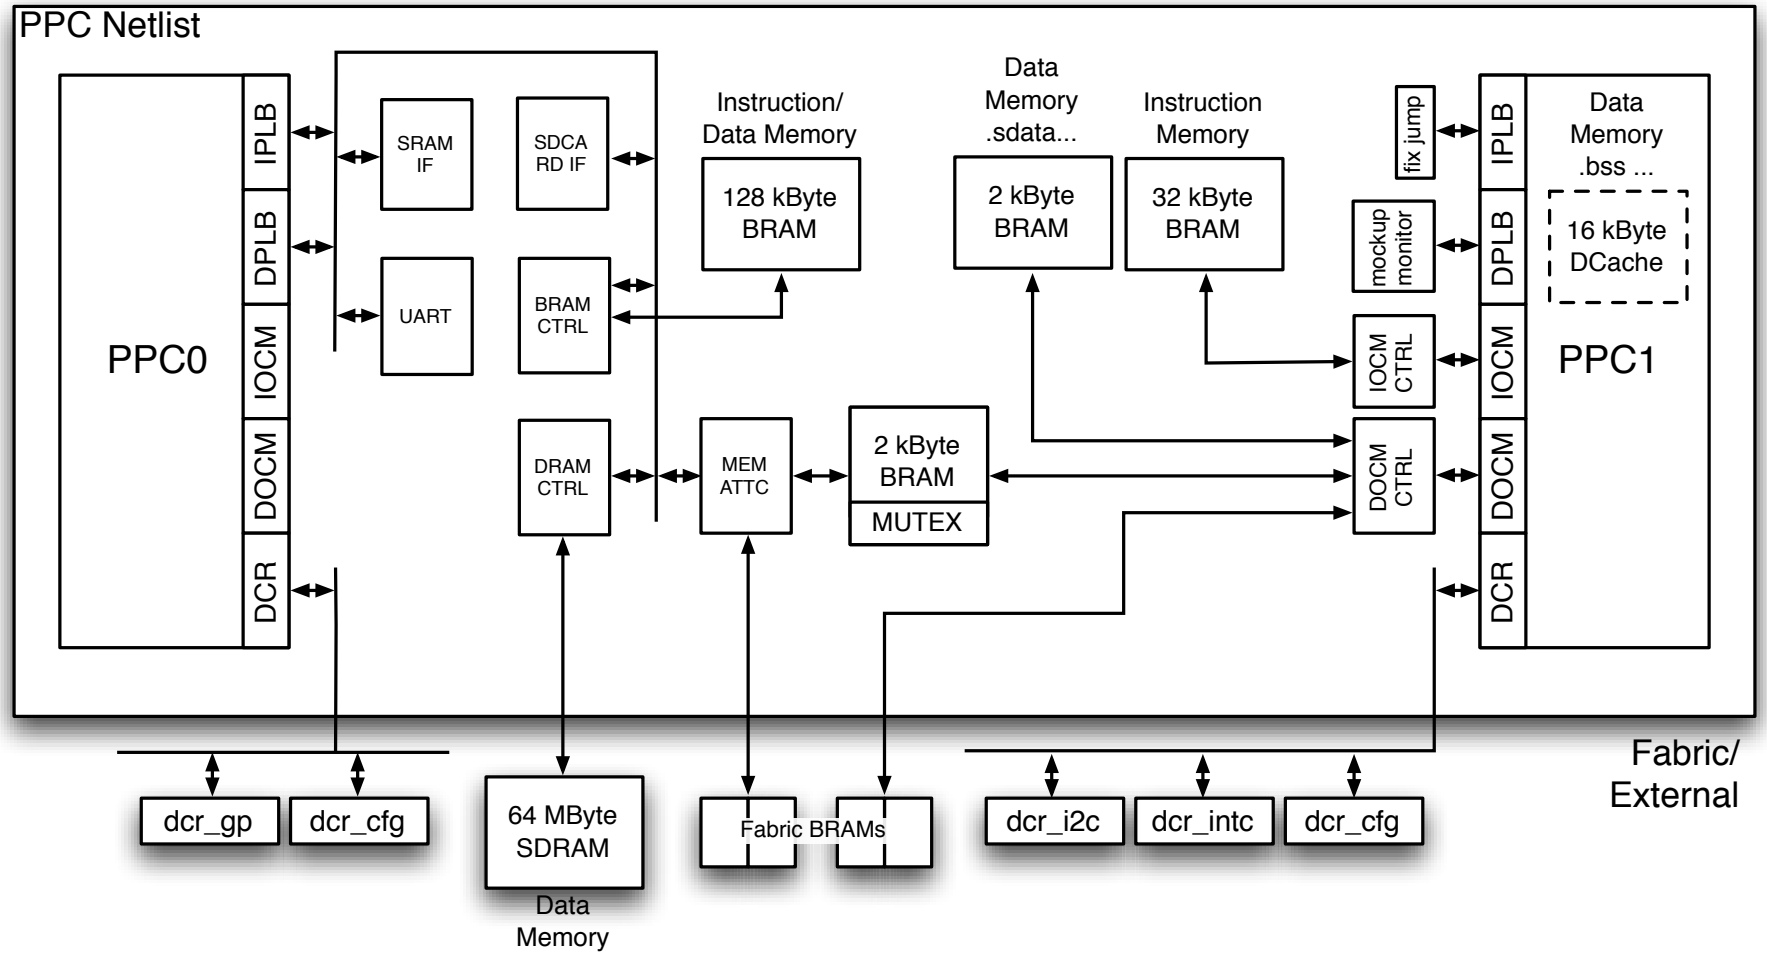

- PPC0: 200 MHz, typical 32-bit PLB layout with 100 MHz
- PPC1: 400 MHz, 200 MHz 64-bit OCM busses,
  DCache only used as data memory → interrupt latency 370ns

Tier 0 Design, Rev. 1712

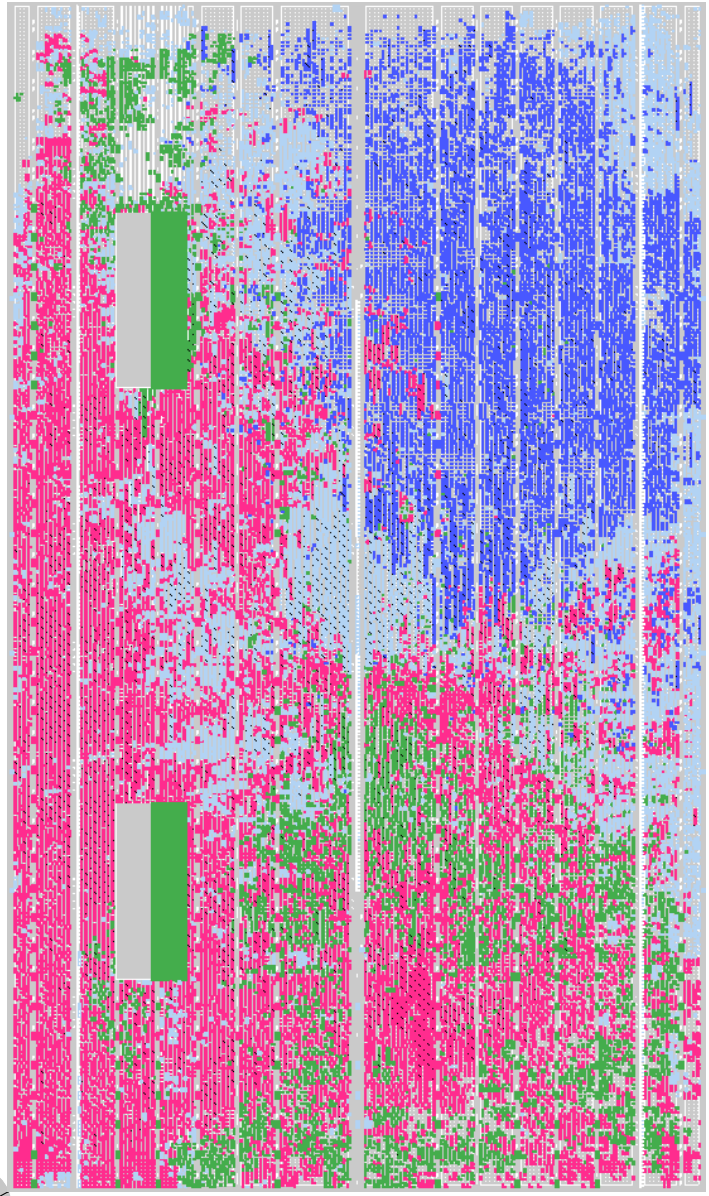

Res.	Event Buffering	Tracking	PPCs
FF	10,945	8,863	3,733
LUT	5,925	23,463	4,086
CY / DSP	2505	8,507 / 9	346
BRAM	14	132	88
Dist. mem	32	1221	187

Total slices: 38,910 (92%)

Total LUTs: 53,406 (63%)

Global clocks: 16 (50%)

Block RAMs: 248 (65%)

Tracking Event Buffering PowerPC

Tier 1 & 2 Designs, Rev. 1712

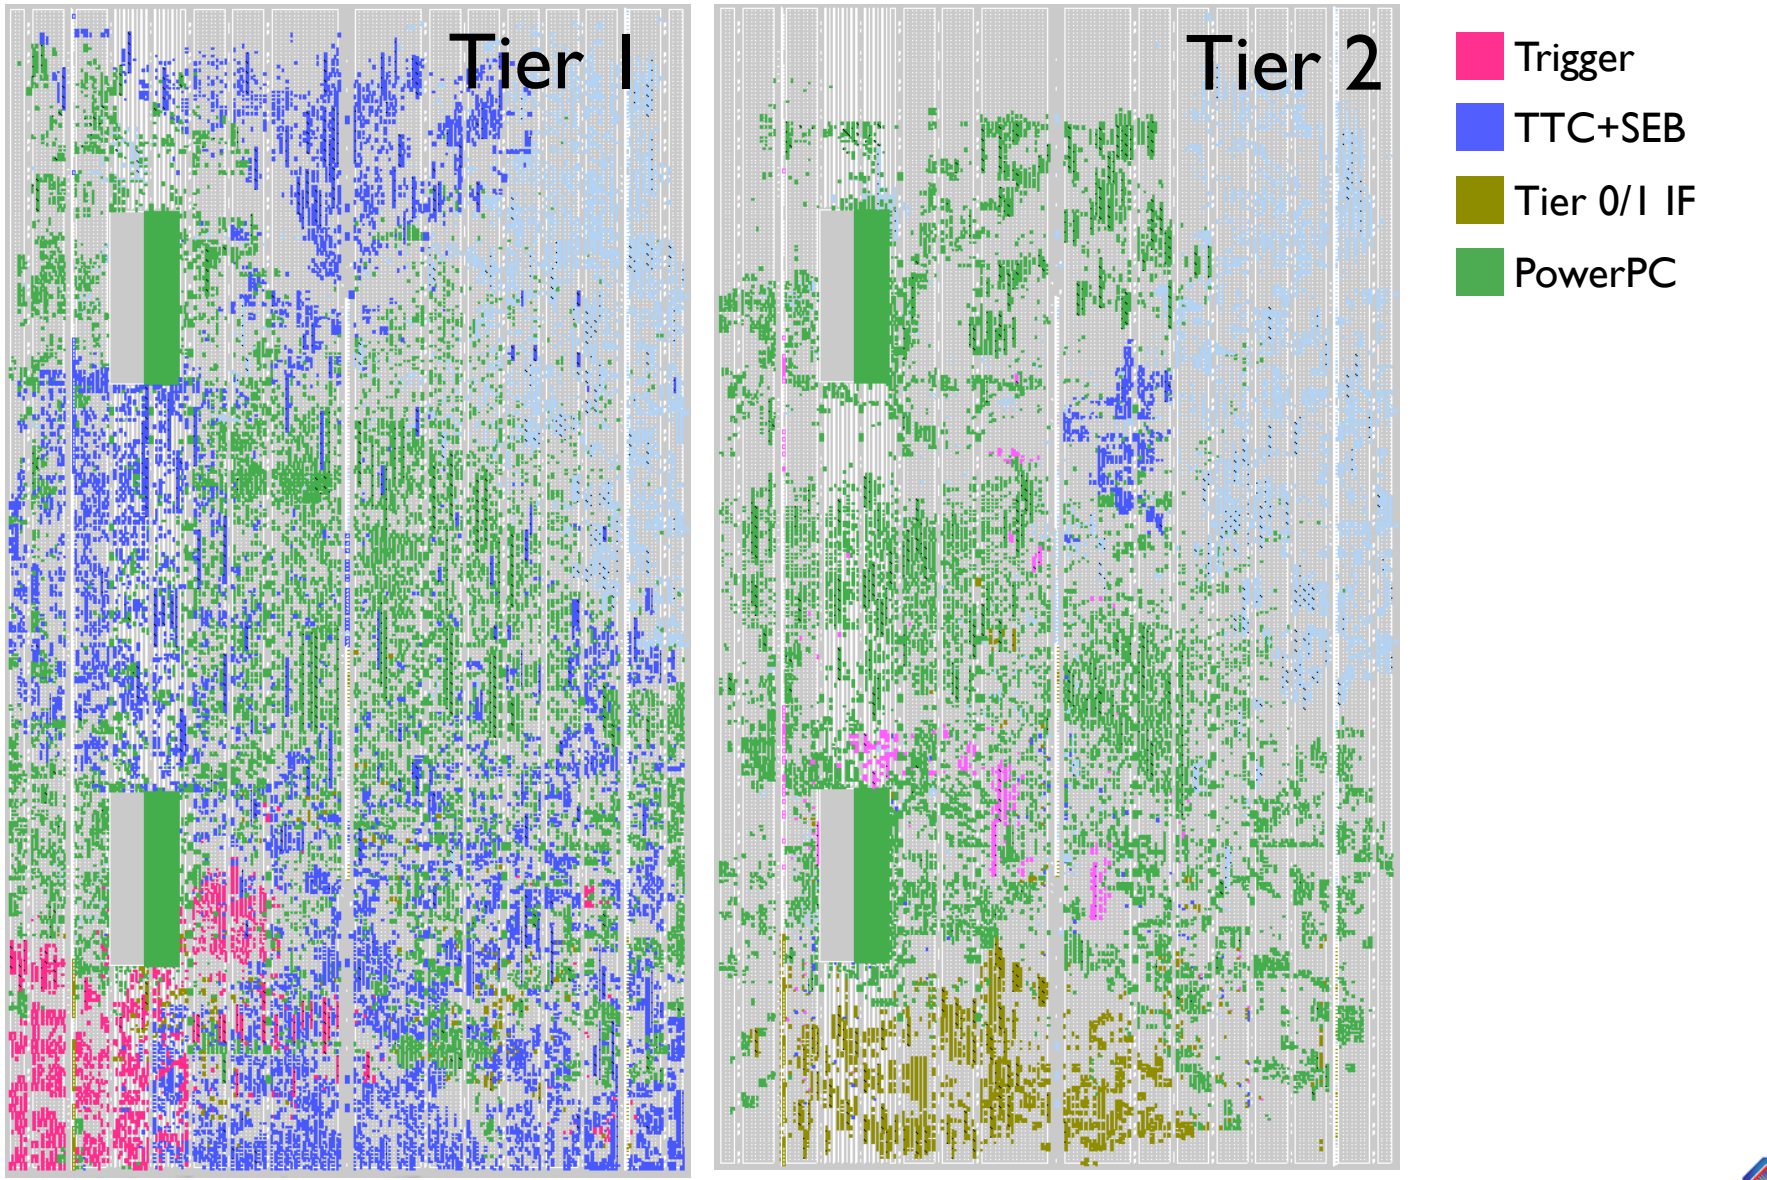

Build and Simulation Flow

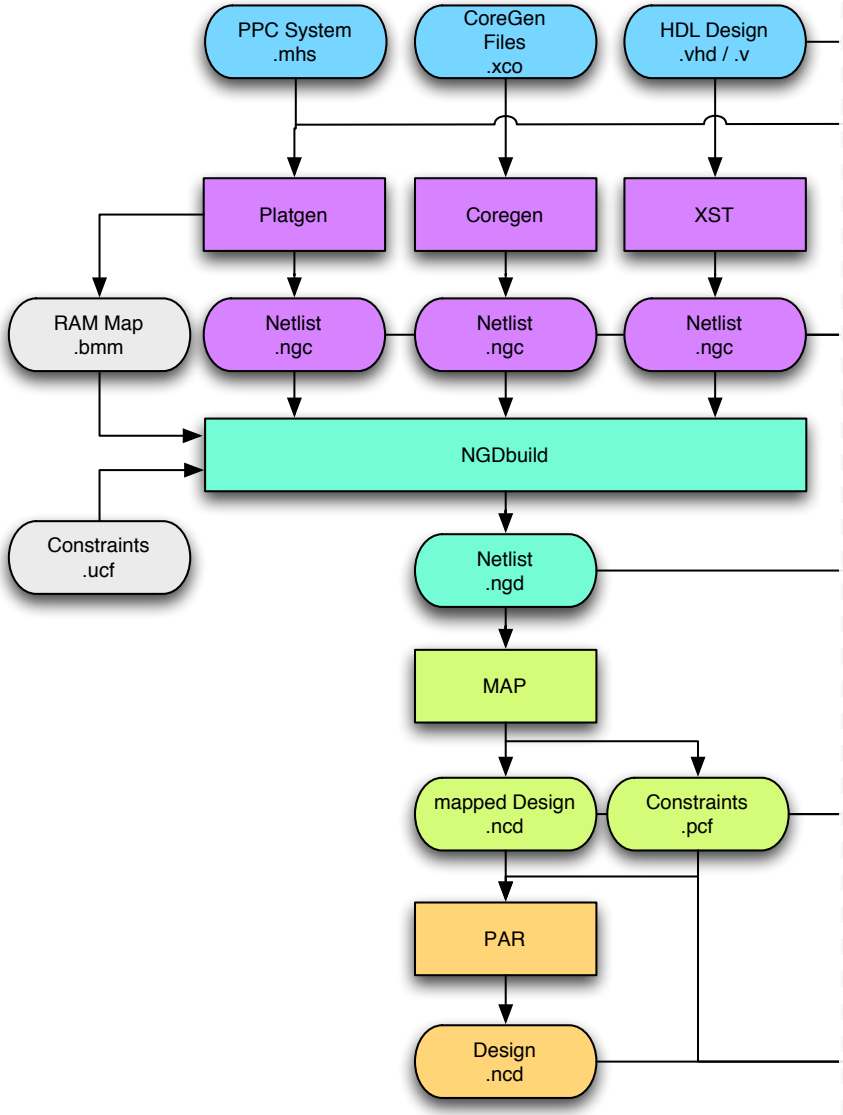

- Makefile-based Flow
- Nightly Build System
- Further Development while system is in continuous use at LHC
  → need for detailed verification
- Three tiers, many interfaces to external systems
  → complex sim. models & interaction between multiple FPGA designs

Build and Simulation Flow

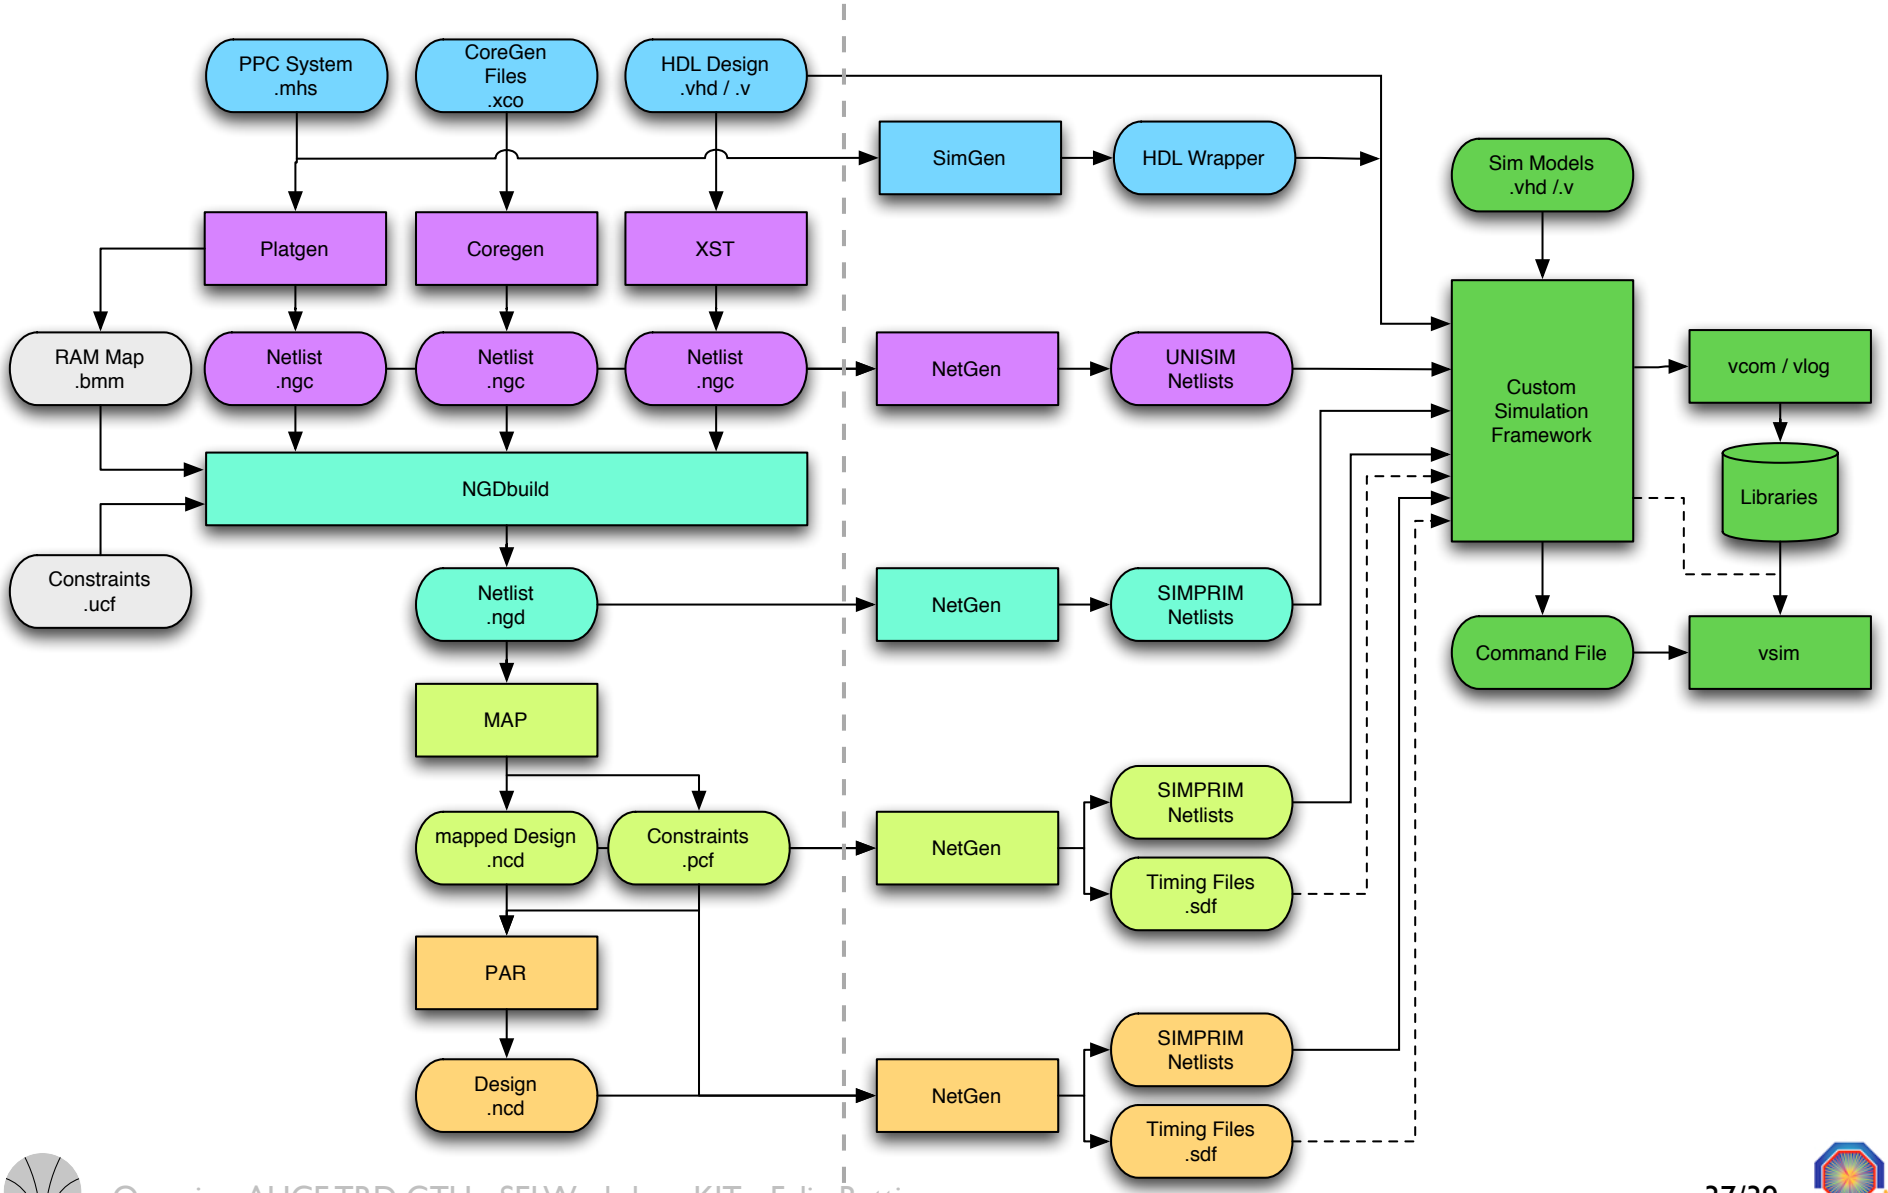

Outlook

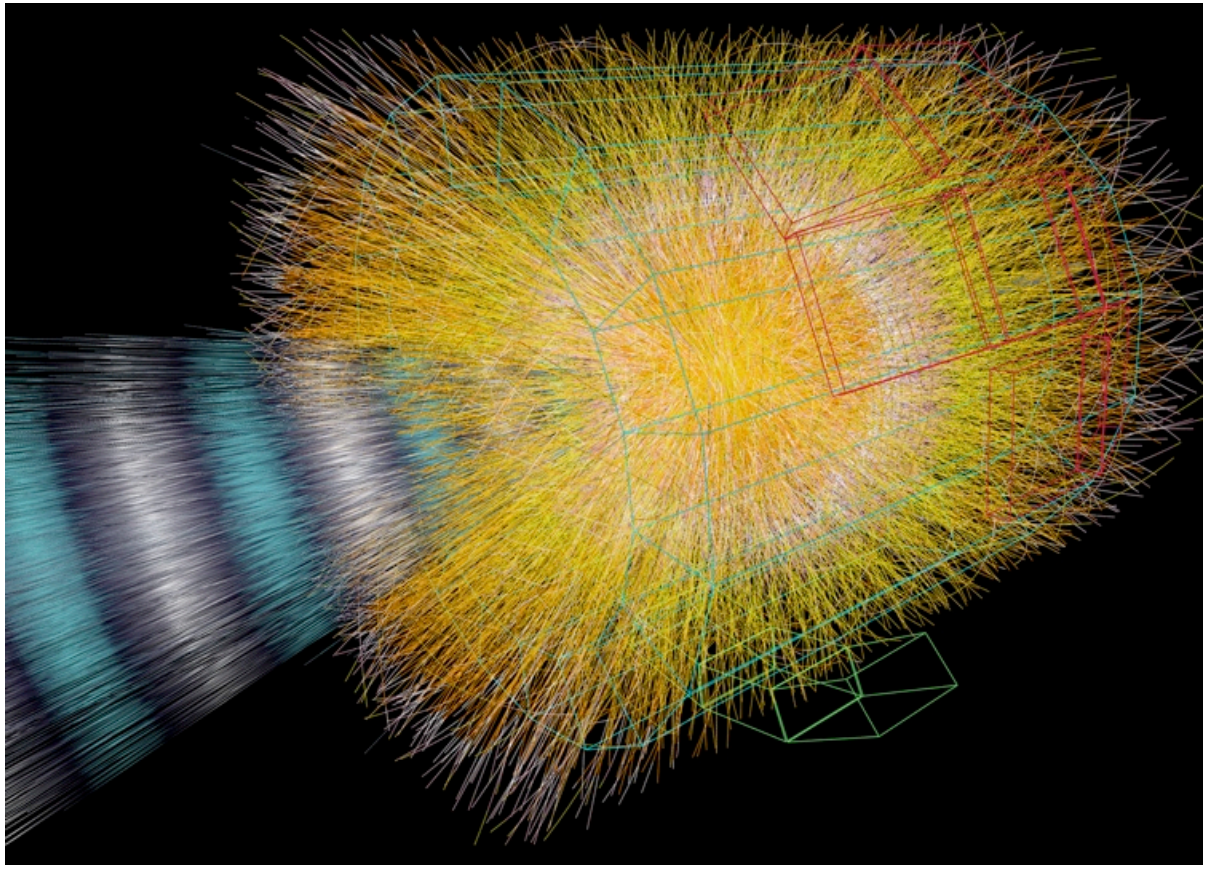

First Pb-Pb collisions in November...

High multiplicities, combinatorics explode

- Tracking almost consumes all time available for trigger
- Jet and high-p_t trigger in operation
- Ideas:
  - Tuned L1 triggers (6-8μs)
  - Elaborate L2 triggers (80μs)

Thank You for Your Attention!

Contact:

Felix Rettig

rettig@kip.uni-heidelberg.de

Prof. Dr. Volker Lindenstruth

Chair of Computer Science
Kirchhoff Institute of Physics
University of Heidelberg

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

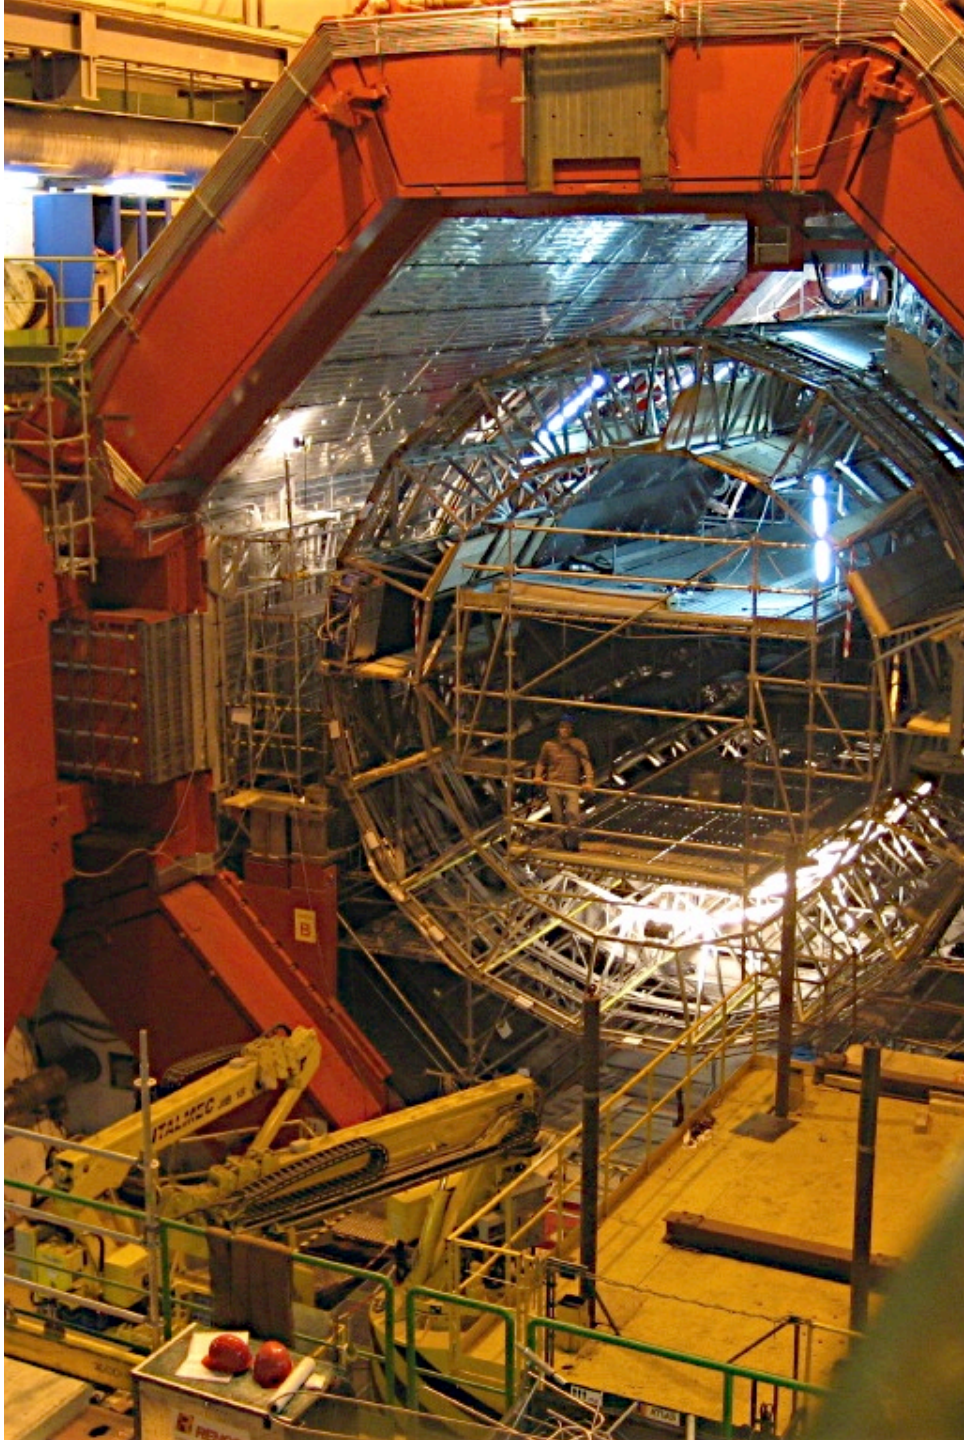

Current TRD Supermodule Status

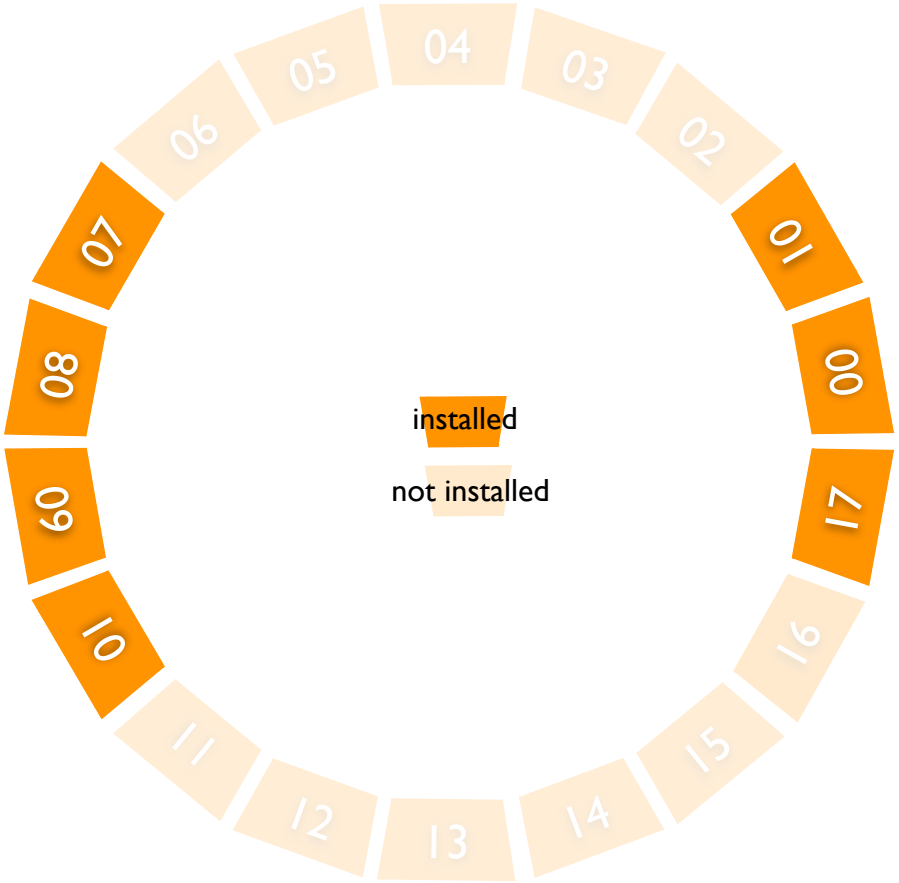

- 7/18 Supermodules installed
- 18/18 GTU segments installed

C16		C17		C18	
00	09	03	12	06	15
01	10	04	13	07	16
02	11	05	14	08	17

final cabling

no cabling

PowerPCs - Performance

Parameter	Core: 200 MHz IOCM: 100 MHz DOCM: 100 MHz	400 MHz 200 MHz 100 MHz	400 MHz 200 MHz DCache
Main loop period (I2C, ...)	57ms		
Interrupt → first handler instruction	1.13µs (226 icycles)	1.05µs (420)	0.37µs (148)
second handler instr.	1.25µs (250)	1.13µs (452)	0.45µs (180)
third handler instr.	1.37µs (274)		
4. if condition check	3.43µs (686)		1.54µs (580)
if (true condition) {	540ns (108)		
if (false condition) {} else {	540ns (108)		

ALICE Trigger Hierarchy

Trigger	Pre-Trigger	Level-0	Level-1	Level-2	High-Level
Time after Interaction	0.3 µs	1.2 µs	6.5 µs	~ 88 µs	> 1 ms
Average Rate (Pb-Pb)	~ 5000 Hz	~ 5000 Hz	~ 400 Hz	~ 200 Hz	~ 100 Hz
Description/ Use	TRD Specific Wake-Up	Strobe to Sampling Electronics	Major Rate Reduction	TPC Past-Future Protection	Software Trigger, Data Compression
TRD Contribution	generated for TRD	TRD contributes to L0 via Pre-Trigger	TRD contributes to L1 via GTU	–	–