Low-Power, Networked MIMD Processor for Particle Physics

Outline

- Introduction to High Energy Physics
- Introduction to Triggers
- ALICE TRD Trigger Architecture
- The MIMD Processor
- Tests and Measurements
- Summary

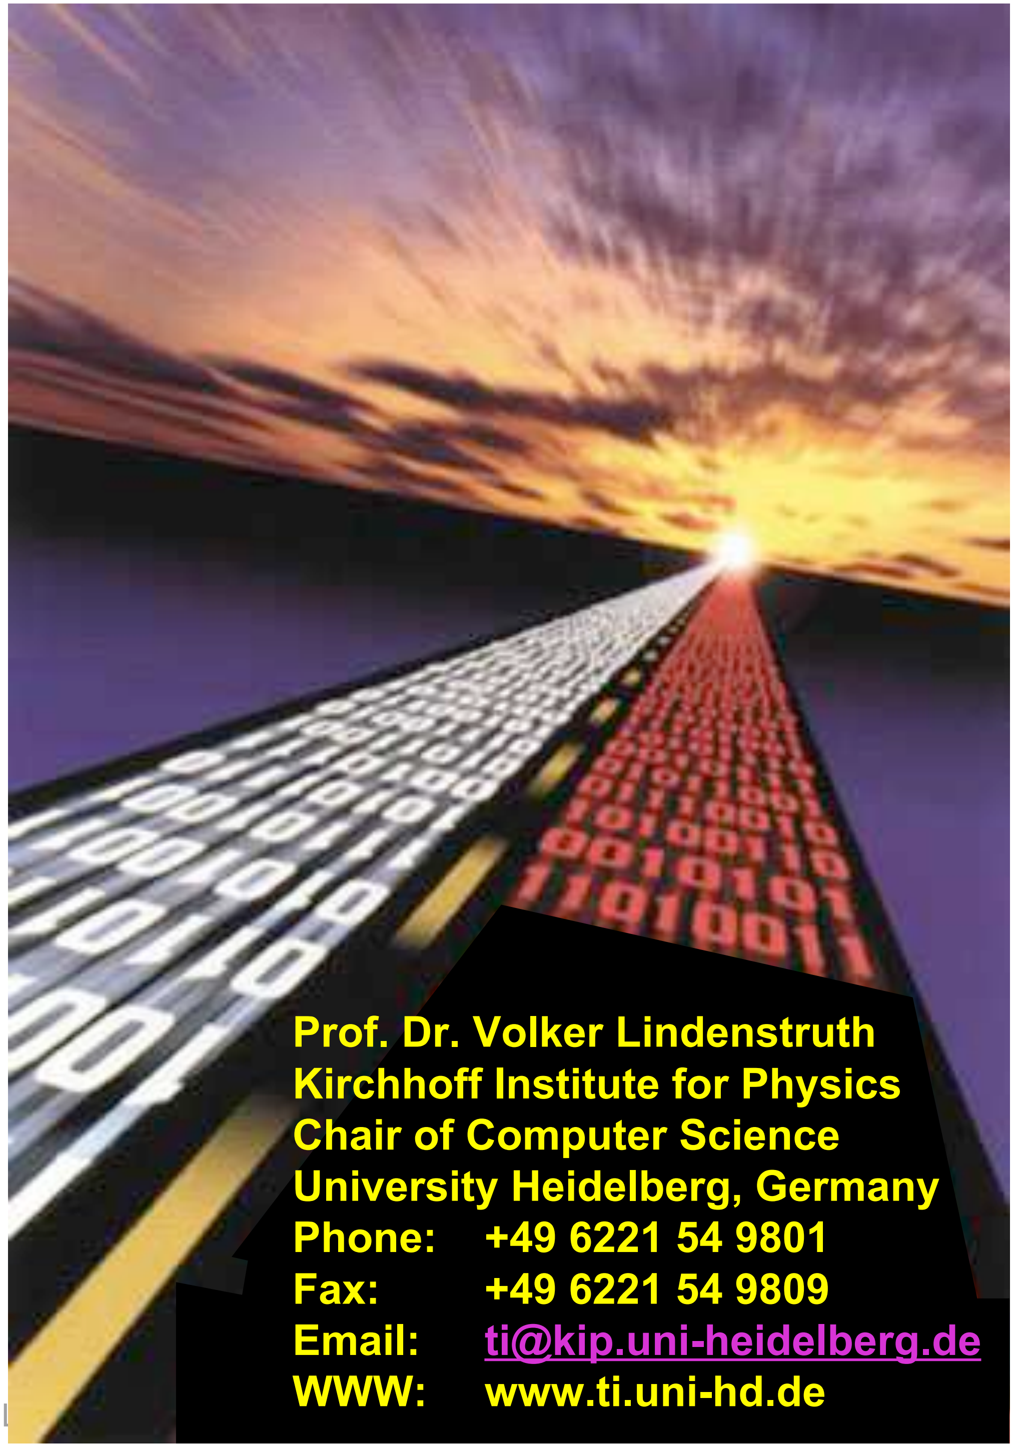

Prof. Dr. Volker Lindenstruth
Kirchhoff Institute for Physics
Chair of Computer Science
University Heidelberg, Germany
Phone: +49 6221 54 9801
Fax: +49 6221 54 9809
Email: ti@kip.uni-heidelberg.de
WWW: www.ti.uni-hd.de

Introduction to High Energy Physics

- **The following slides give a brief overview of the scientific environment of the application for the MIMD Processor**
- **One goal of the experiments under construction is the generation of a quark/gluon plasma, which can only be created by colliding nuclei at very high kinetic energies**
- **The following slides show a schematic overview of an appropriate experiment, capable of detecting the reaction products of such collisions**
- **The photo shows the experiment at its current state. It is scheduled to take first experimental data in Q2 2007**

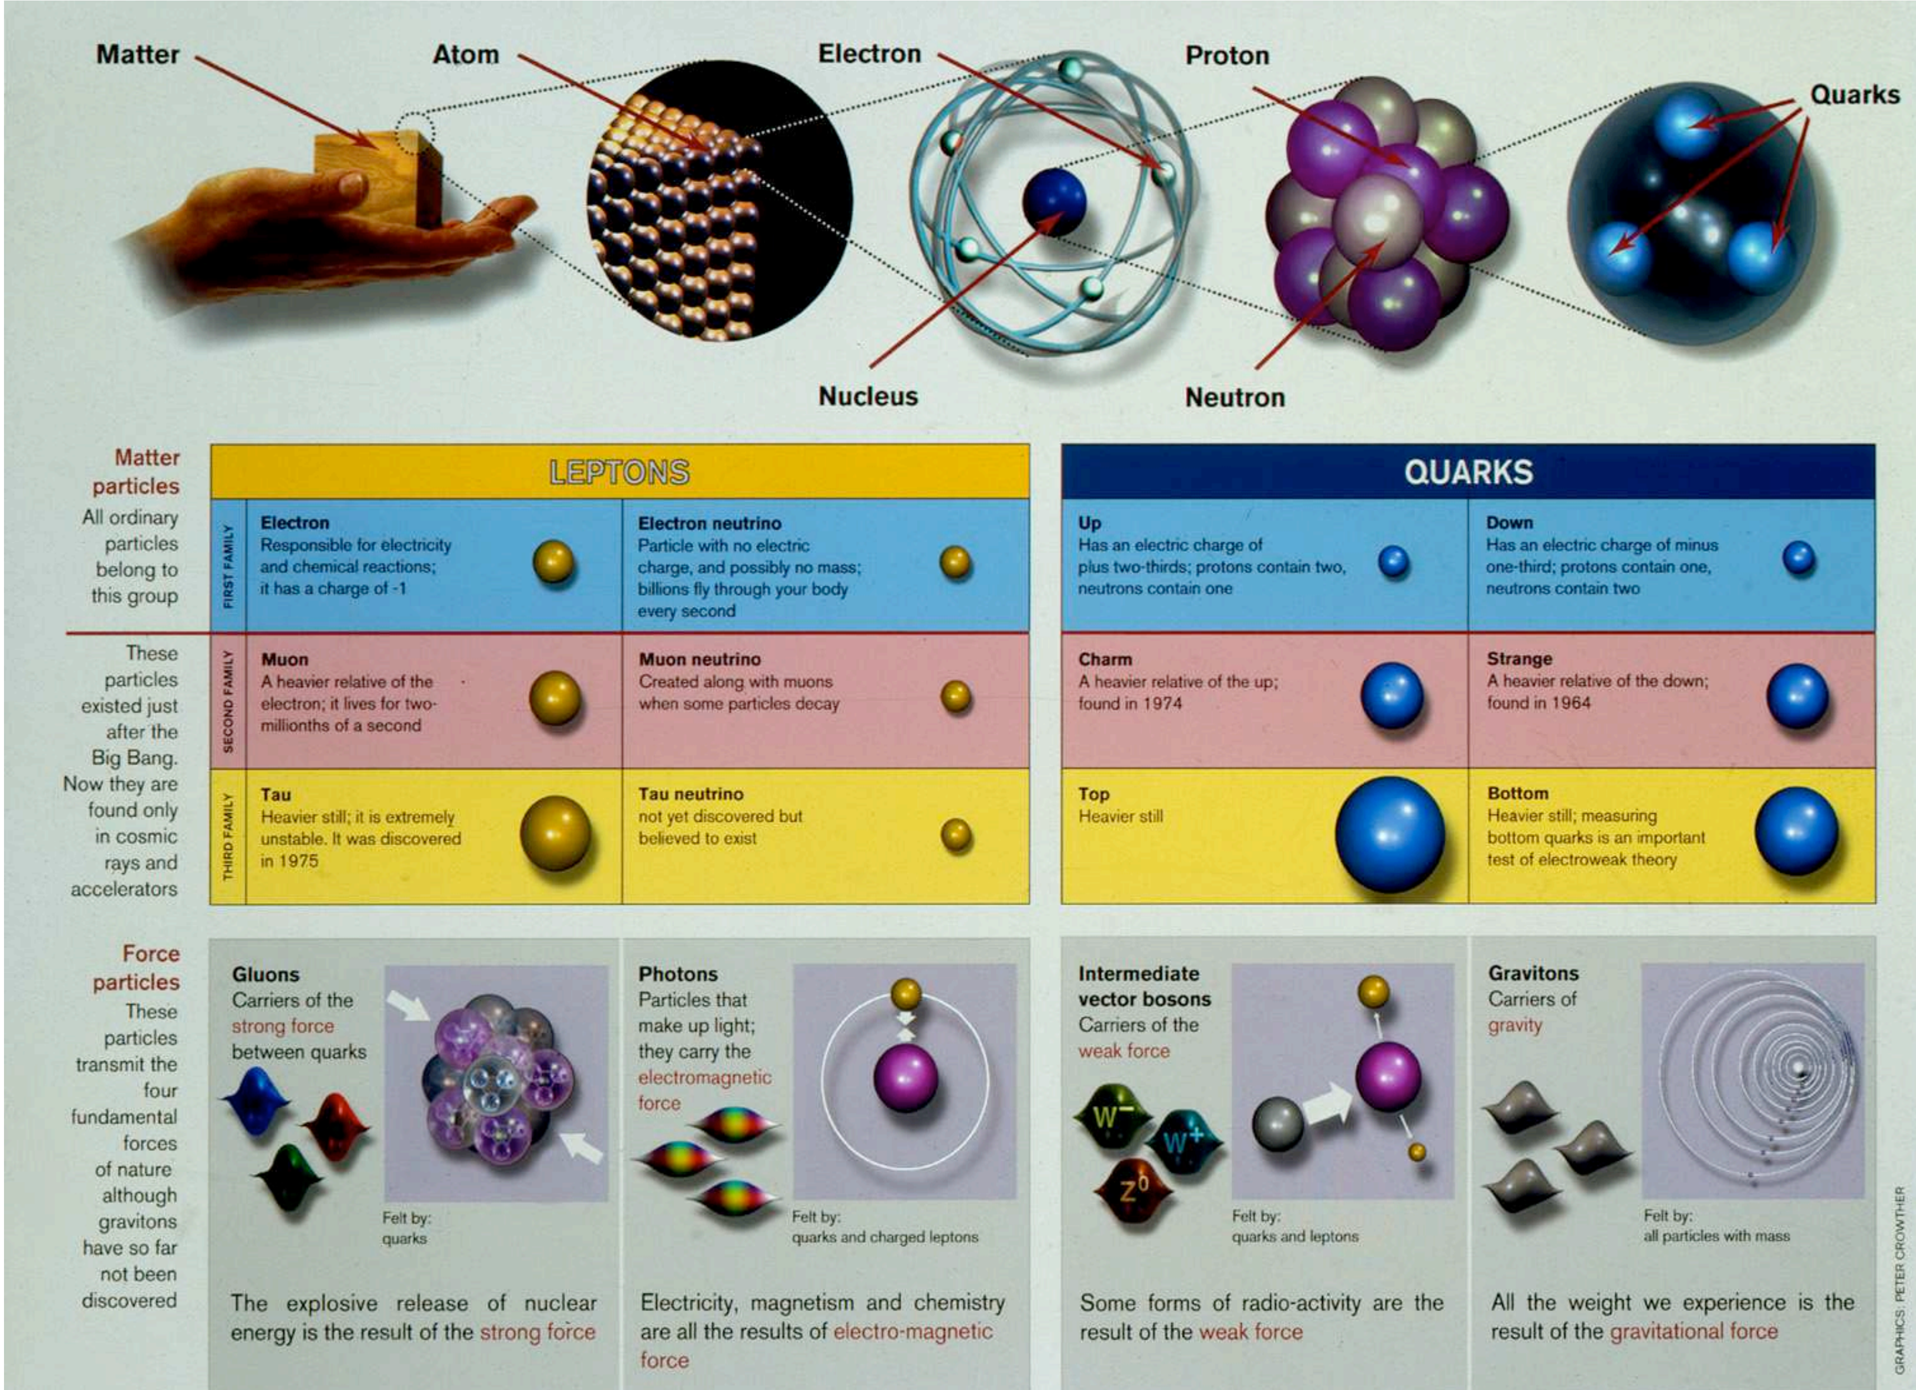
Matter
Atom
Electron
Proton
Quarks
Nucleus
Neutron
Matter particles
All ordinary particles belong to this group
These particles existed just after the Big Bang. Now they are found only in cosmic rays and accelerators
LEPTONS
QUARKS
FIRST FAMILY
Electron
Responsible for electricity and chemical reactions; it has a charge of -1
Electron neutrino
Particle with no electric charge, and possibly no mass; billions fly through your body every second
Up
Has an electric charge of plus two-thirds; protons contain two, neutrons contain one
Down
Has an electric charge of minus one-third; protons contain one, neutrons contain two
SECOND FAMILY
Muon
A heavier relative of the electron; it lives for two-millionths of a second
Muon neutrino
Created along with muons when some particles decay
Charm
A heavier relative of the up; found in 1974
Strange
A heavier relative of the down; found in 1964
THIRD FAMILY
Tau
Heavier still; it is extremely unstable. It was discovered in 1975
Tau neutrino
not yet discovered but believed to exist
Top
Heavier still
Bottom
Heavier still; measuring bottom quarks is an important test of electroweak theory
Force particles
These particles transmit the four fundamental forces of nature although gravitons have so far not been discovered
Gluons
Carriers of the strong force between quarks
Felt by: quarks
The explosive release of nuclear energy is the result of the strong force
Photons
Particles that make up light; they carry the electromagnetic force
Felt by: quarks and charged leptons
Electricity, magnetism and chemistry are all the results of electro-magnetic force
Intermediate vector bosons
Carriers of the weak force
W−
W+
Z0
Felt by: quarks and leptons
Some forms of radio-activity are the result of the weak force
Gravitons
Carriers of gravity
Felt by: all particles with mass
All the weight we experience is the result of the gravitational force
GRAPHICS: PETER CROWTHER

LHC and the Geneva Area

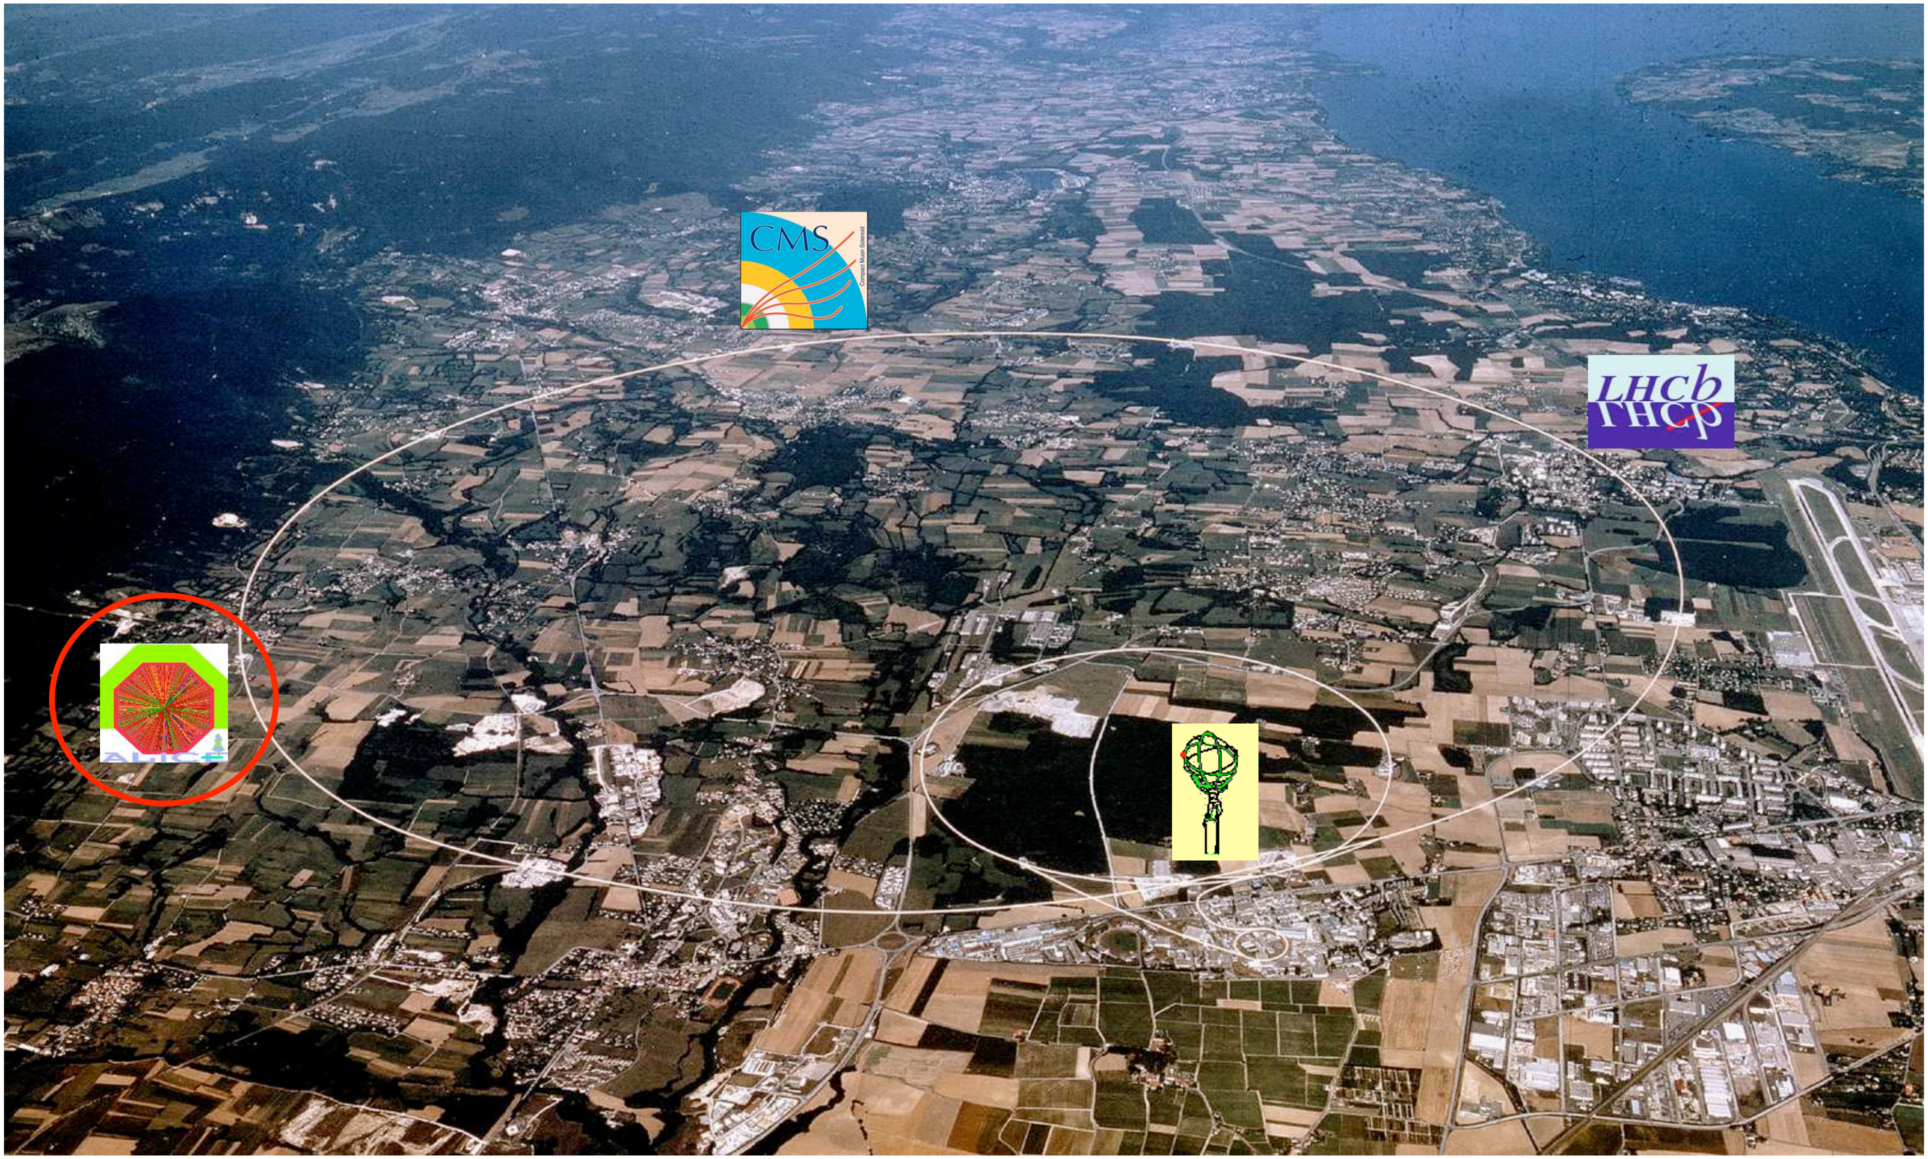

The ALICE Experiment (one of four at LHC)

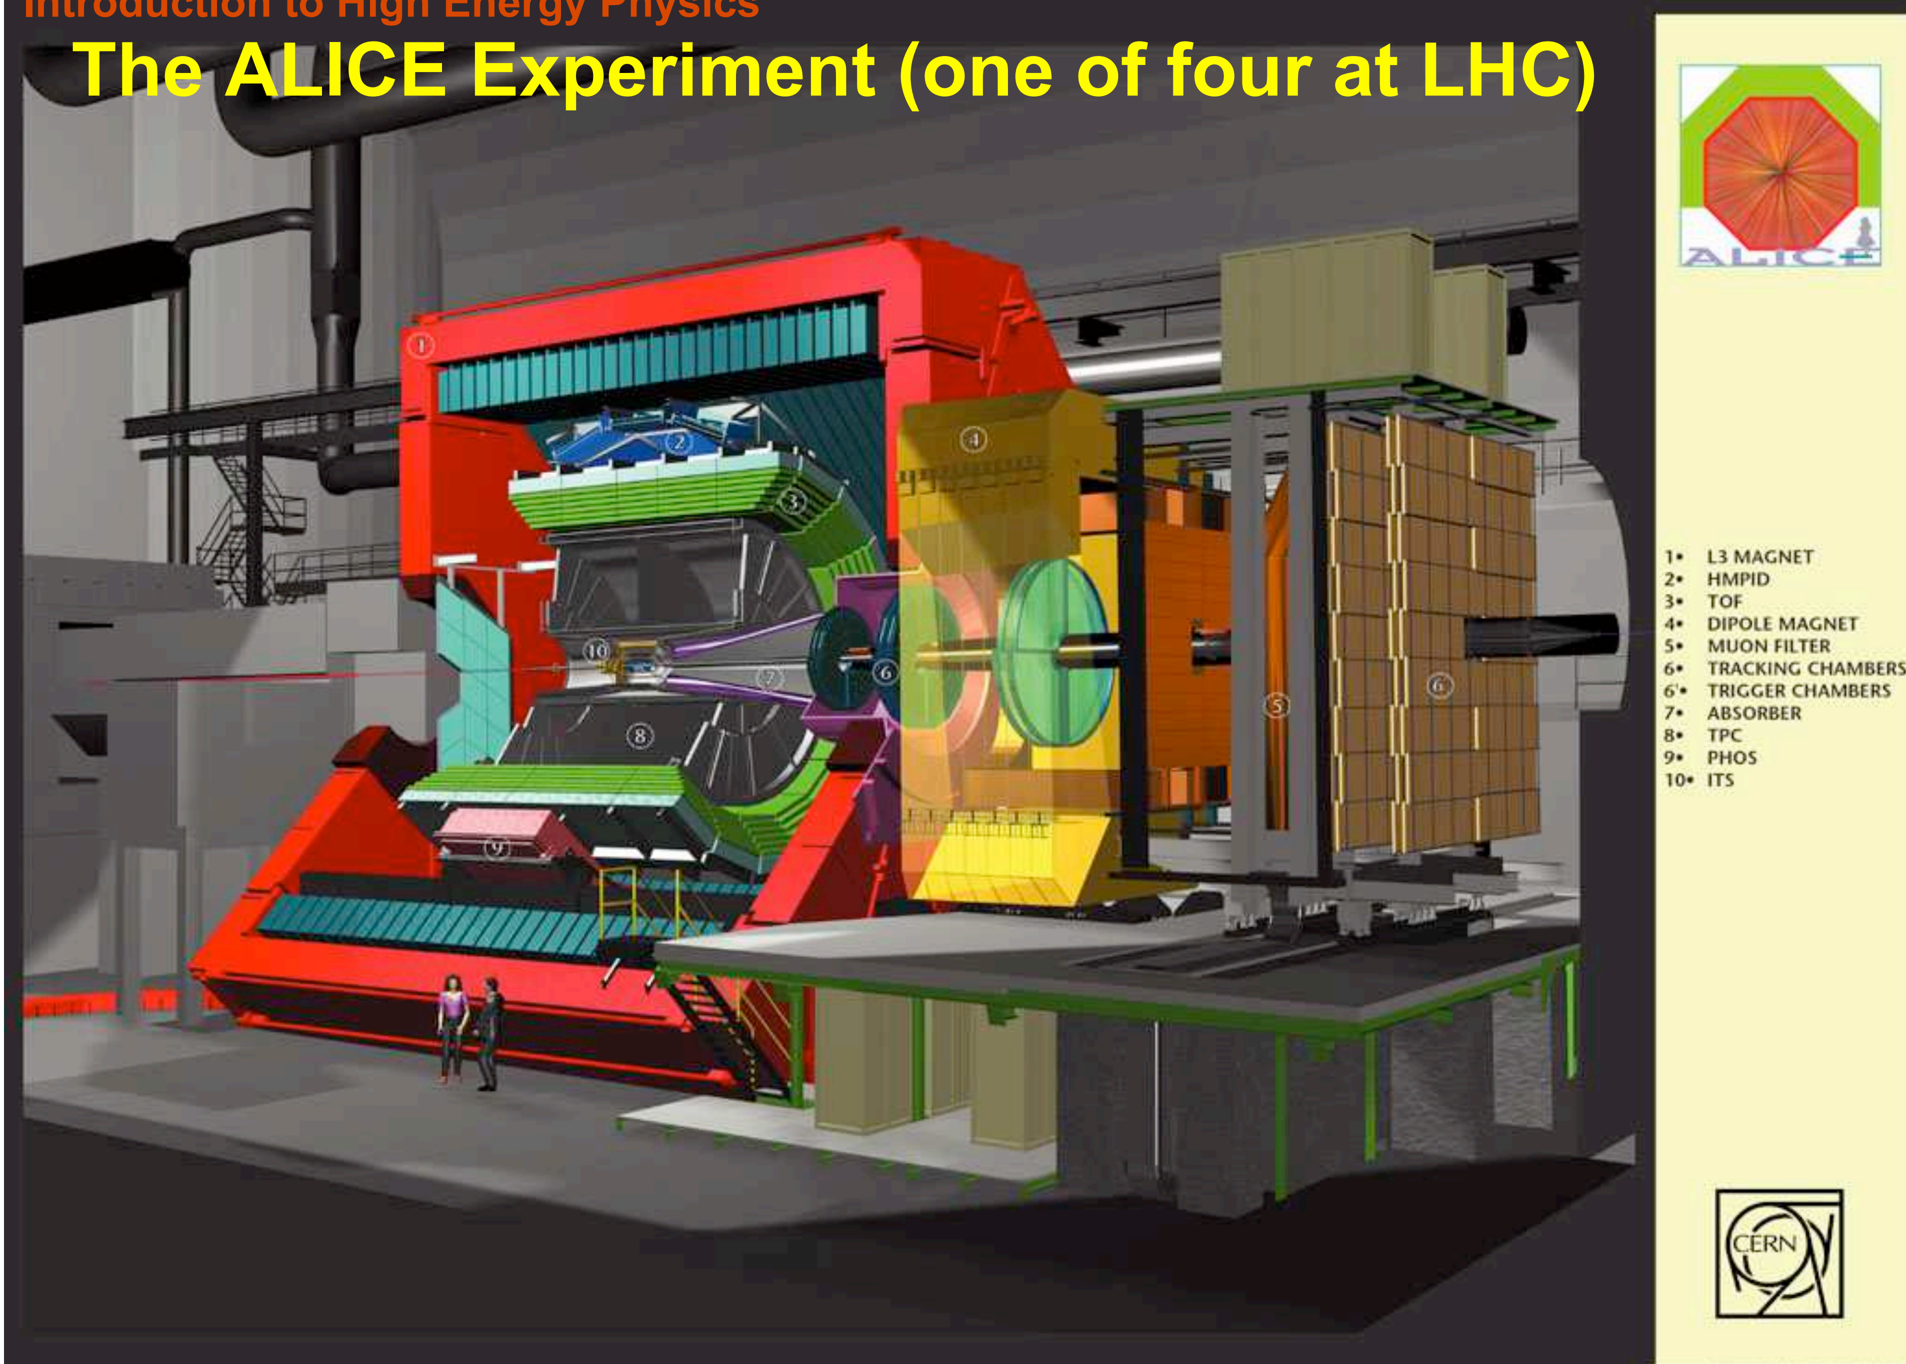

The ALICE Experiment (one of four at LHC)

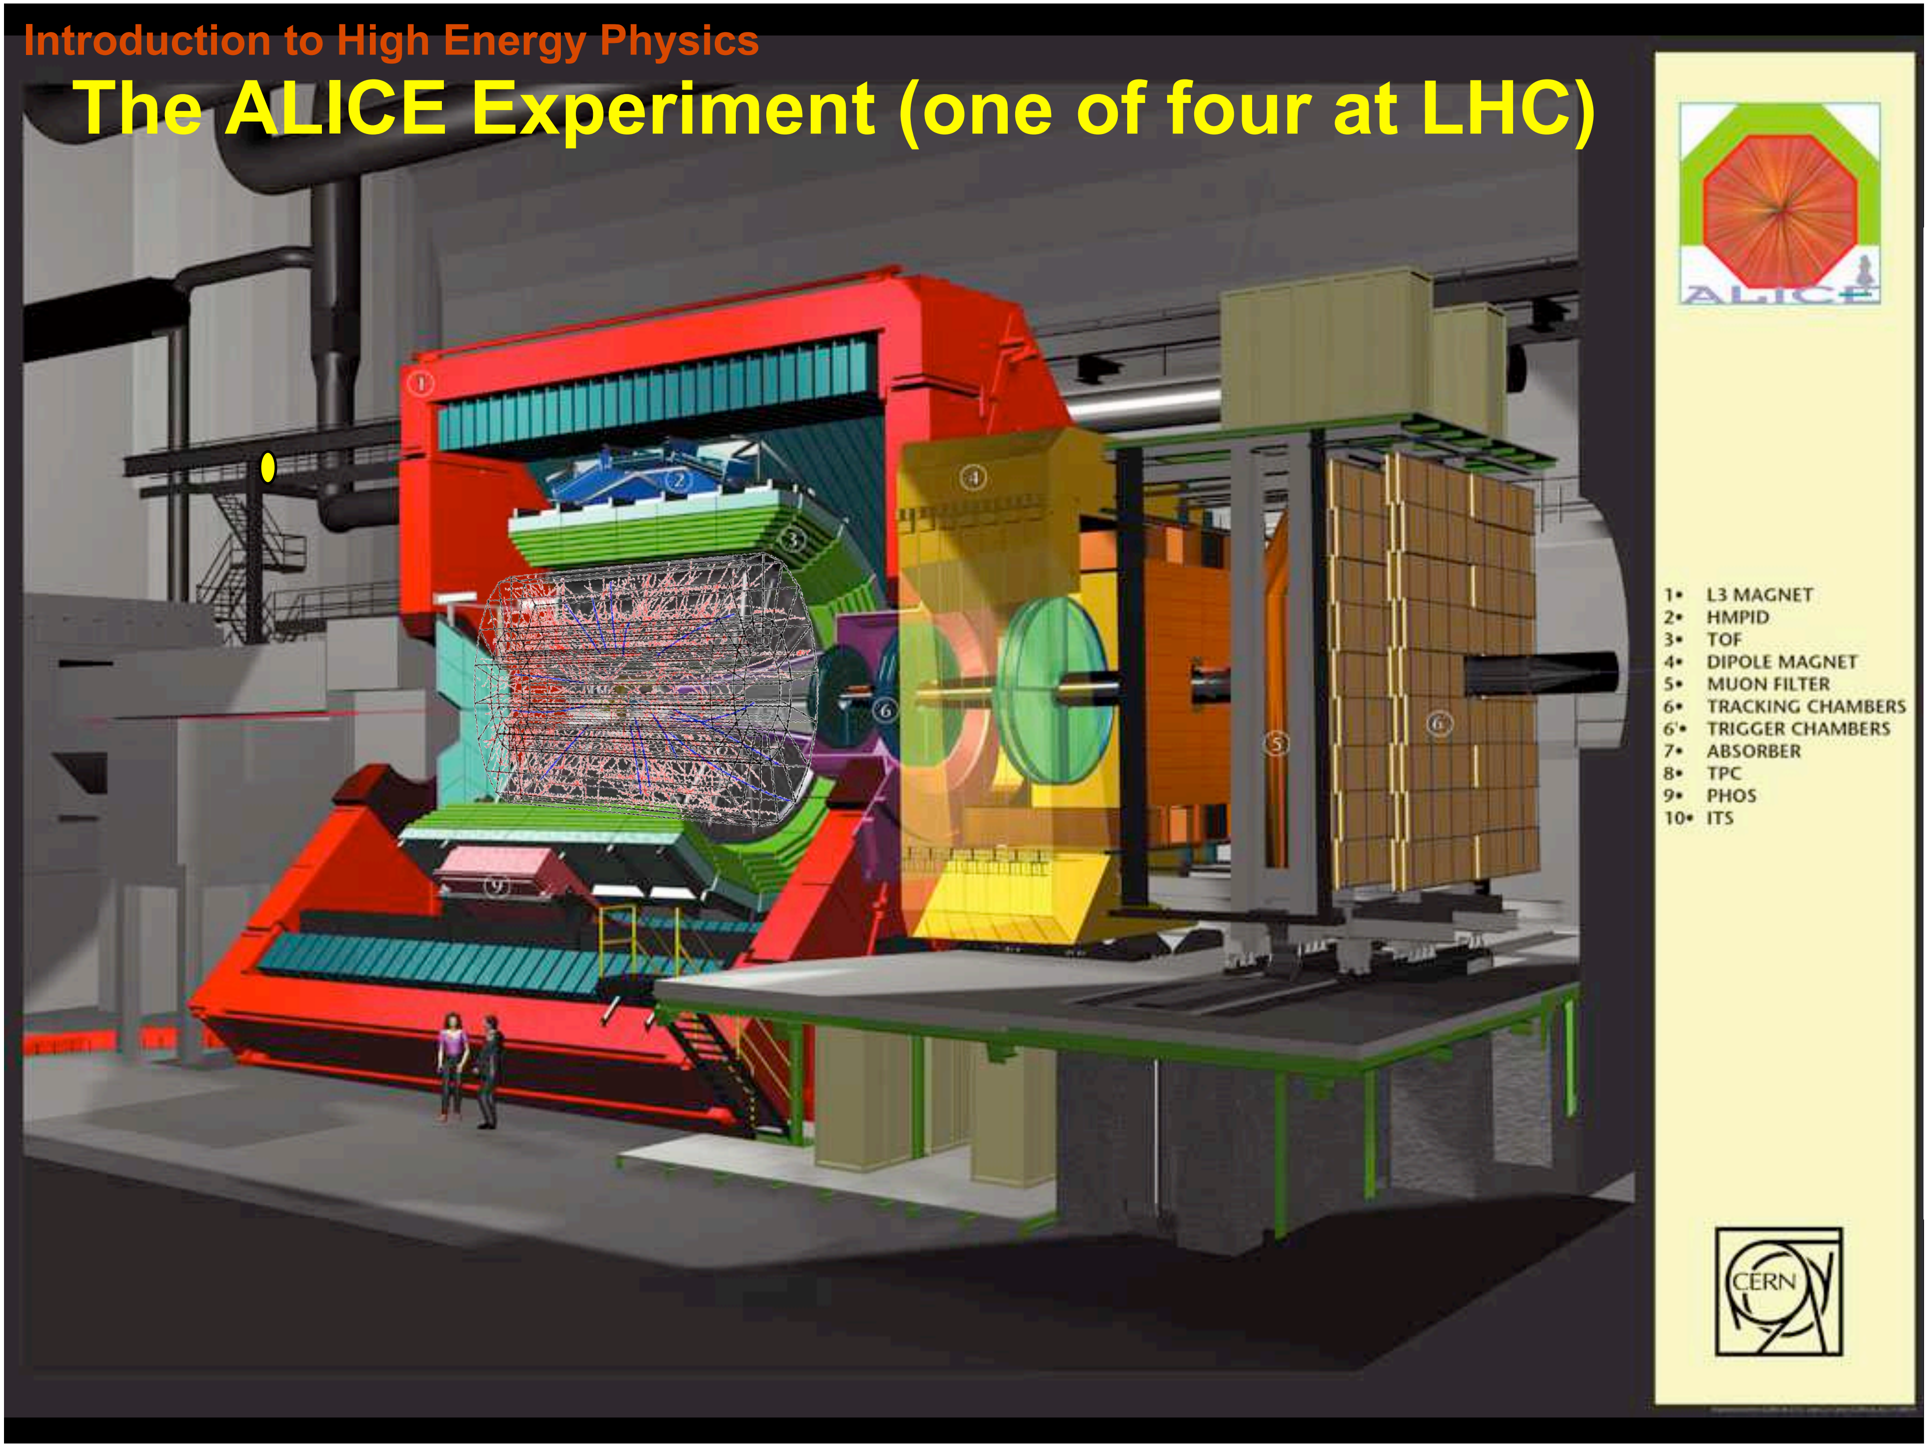

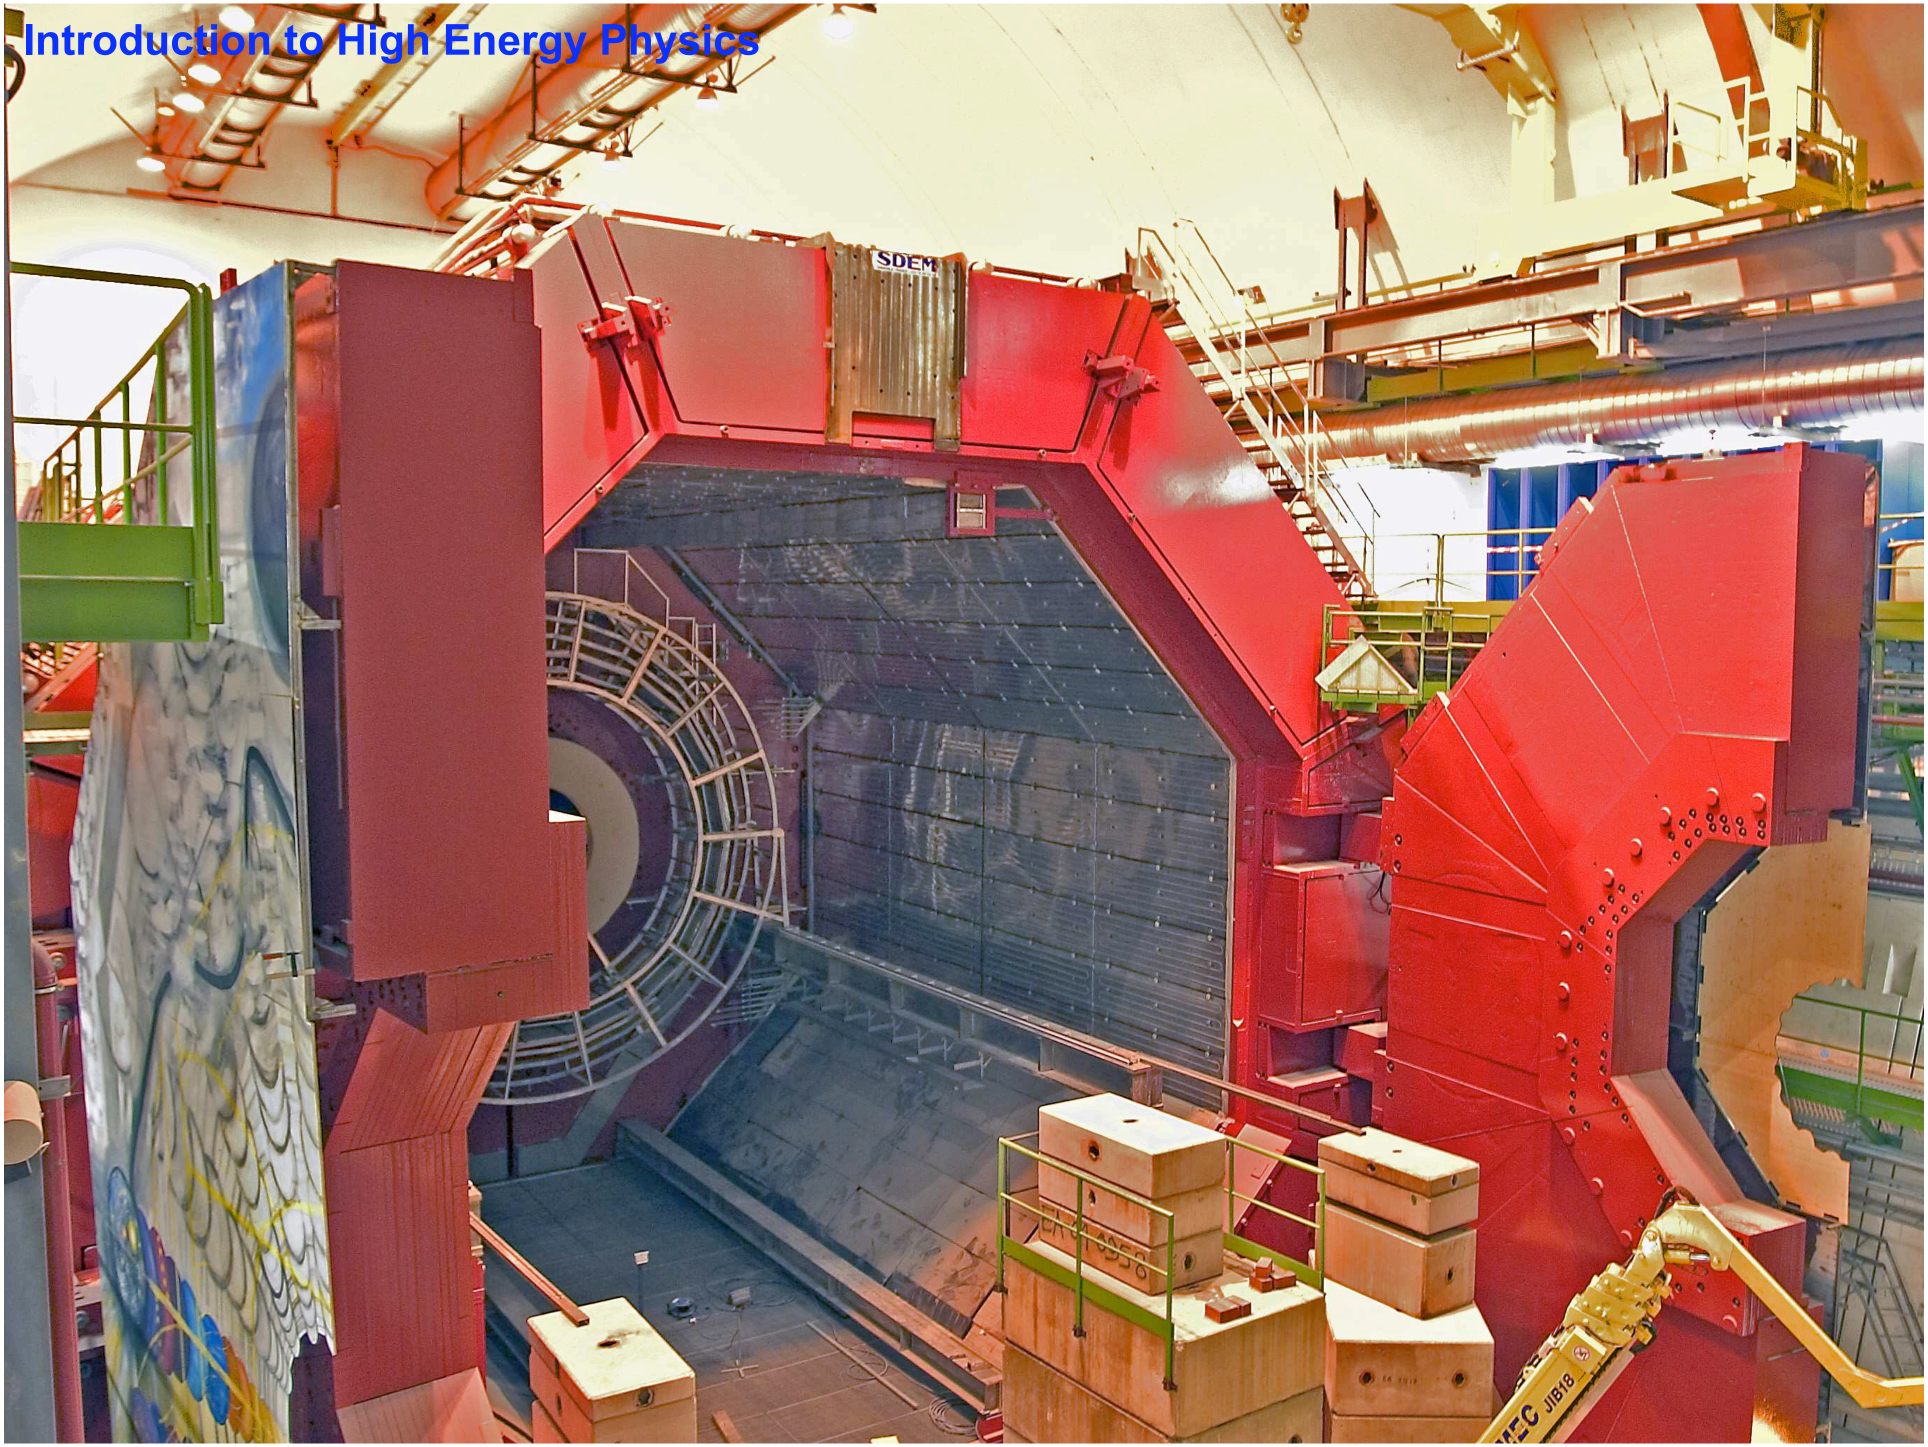
SDEM

Introduction to Triggers

- **The large number of particles, created in every interaction and the statistical nature of the reaction on one hand and the large amounts of data, exceeding 10 TB/sec on the other hand require intelligent selection functionality, allowing to select the specific reactions, the experiment is looking for**
- **Such on-line selection functions, trigger the readout and archival of the data associated to a given event and are therefore called triggers**
- **Triggers can be classified according to the event size and event rate they are capable of processing**

The governing principles

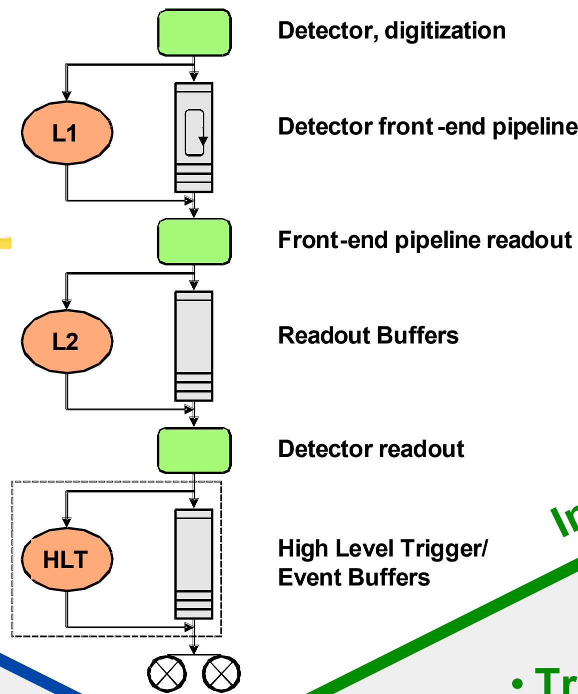

Decreasing with time after event

Increasing with time after event

- Event Rate
- Data Rate

- All Triggers >L0 reject
- Strictly hierarchical
- No undo of reject
- Buffering .vs. Latency
- Isochronicity .vs. Latency
- Pipelining
- Decomposition, assignment

- Trigger Layer
- Information considered
- Complexity
- Event Size
- Trigger latency
- Required buffering
- Available decision time
- Distance to detectors
- Selectivity

microelectronics FPGAs, DSPs PC farms

event L1 L2 HLT

µs 100 µs ms Time

Existing and planned Trigger Systems

Exp.	Acc.	Trigger	Physics Signature	Selectivity 1/x	Channels #	(Inp)Event Rate kHz	Processing Time µs	Event Size kB	URL
ALICE	LHC	L0-2	MinBias, Dimuon						
		TRD	Pt, transition radiation	1000	1200	150	6	30 MB	www.kip.uni-hd.de/TRD
		HLT							
LHCb	LHC	L0 (HW)	Et, Pt, Ninter, Ntracks	10		10000	4		lhcb.web.cern.ch/lhcb/
		L1 (PC farm)	large ImpPar, high Pt	25	170k	1000	<1 ms>	5 kB	
		HLT (PC farm)	B selection	80		40	<50 ms	50 kB	
ATLAS	LHC	L1 (HW)	Calorimeter + Muon	900	7k + 800k	40000	<2.5		atlas.web.cern.ch/Atlas/Welcome.html
		L2	RoI in all detectors	35		75	<10 ms	30 kB	
		L3 (Event Filter)	Full event reco	10		2	~1 s	1.5 MB	
CMS	LHC	L1 (HW)	Cal., Muon	400		40 MHz	<3.8		cmsinfo.cern.ch/Welcome.html/
		HLT (PC farm)	All detectors	1000		100 kHz			
H1	HERA	L1 (HW)	Vertex, Cal, Muon	300000		10 MHz	<2.3		www-h1.desy.de/
		L2 (HW)	Correlation (ANN)	1		30 Hz	20		
		L3 (RISC, Event Filter)		3		30 Hz	<0.8 ms		
		L4 PC Full reco		3		10 Hz	100 ms		
ZEUS	HERA	FLT (HW)	Global properties	20000		10 MHz	~0.7		www-zeus.desy.de/
		SLT (INMOS transputers)	Regions	10		500 Hz	10 ms		
		TLT (Farm)	Full data	10		50 Hz	none		
HERA-B	HERA	FLT HW (RoI)	Muon, ECAL, track	200		10 MHz	10		www.hera-b.desy.de
		SLT PC (RoI) Vertex	Sec. Vertex (RoI)	100		50 kHz	7 ms		
		TLT PC Vertex	Sec. Vertex	10		500 Hz	100 ms		
		L4 PC Full reco	Event selection	~2.5		50 Hz	4 s	150 kB	
PHENIX	RHIC	L1 (HW)			~350k	10kHz/10MHz			www.phenix.bnl.gov/
		L2 (SW)		>6				100 kB	
STAR	RHIC	L0 (HW)	CTB, ZDC, MWPC			10 MHz	~1.5		www.star.bnl.gov/
		L1 (HW)	CTB, ZDC, MWPC				~100		
		L2 (HW)	CTB, ZDC, MWPC				~5 ms		
		L3 (SW)	TPC, SVT, FTPC			5 Hz	~200 ms	10 MB	
PHOBOS	RHIC	L0 (HW)	Paddle counters		150k+2k	200 Hz			www.phobos.bnl.gov/
		L1 (HW)	Cherenkov cnts, ZDC						
BRAHMS	RHIC	L0	ZDC, BBC	10	~4k	~1.2 kHz			www4.rcf.bnl.gov/brahms/WWW/
		L1							
D0	FNAL	L1 (HW)	Fiber tracker (Pt)	1500	~1000k	7.6 MHz	~2.5		www-d0.fnal.gov/
		L2 (HW)	Silicon Track Trigger	5		5 kHz			
		L3 (PC farm)	Full data	20		1 kHz		250 kB	
CDF	FNAL	L1 (HW)	Calo, Muon, Track	150		7.6 MHz	~5.5		www-cdf.fnal.gov/
		L2 (HW)	Calo, Muon, Track	150		50 kHz	20		
		L3 (SW)	On-line reco	4		300 Hz		200 kB	
BTeV	FNAL	L1 (500 FPGA+2500 DSP)	Detached vertex	100	30M	7.6 MHz	330	200 kB	www-btev.fnal.gov/
		L2/3 (2500 PC Linux farm)		~20		4 kHz (out)	300 ms	100 kB	

Trigger Overview

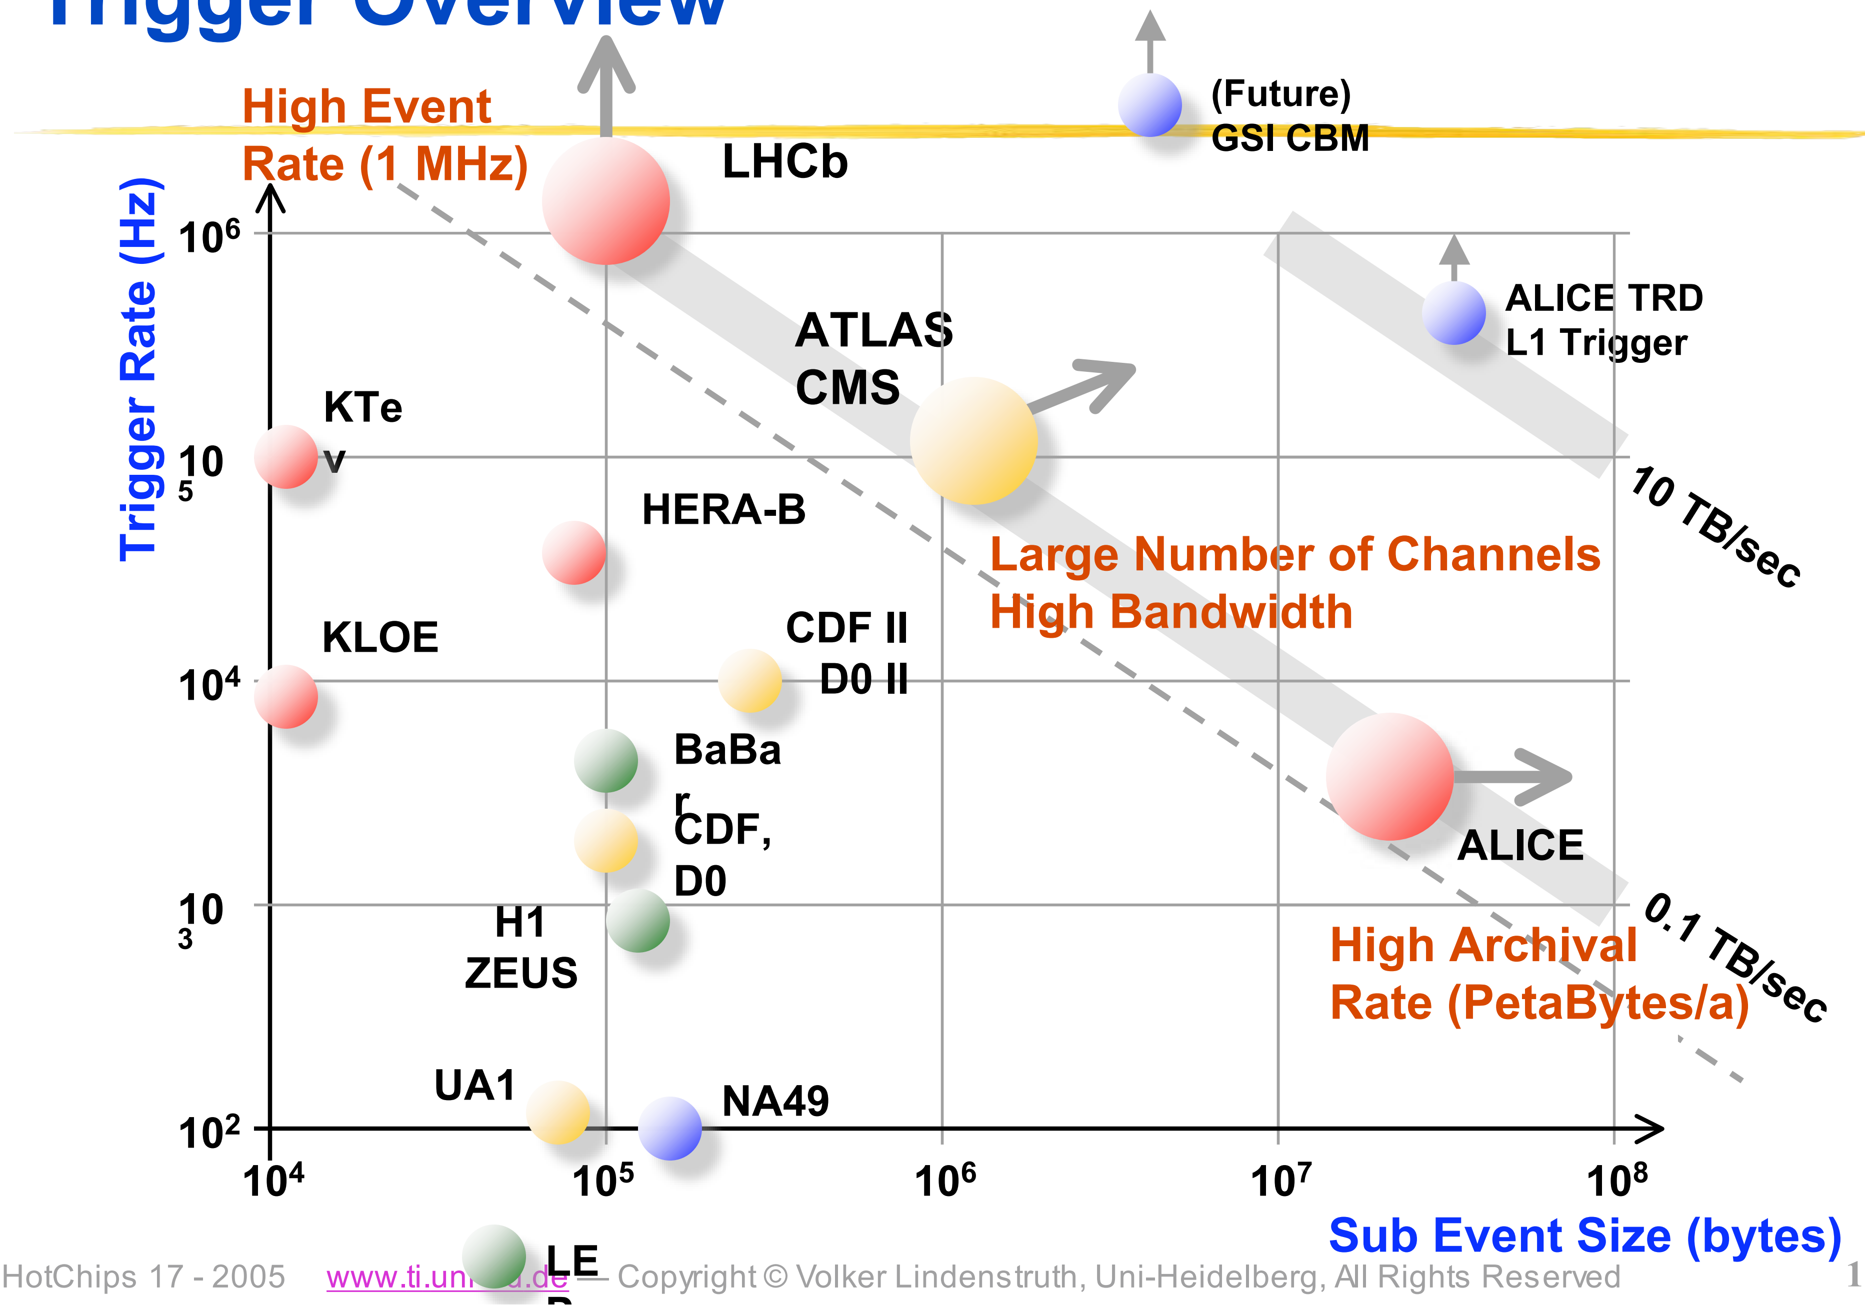

ALICE TRD Trigger Architecture

- **The Transition Radiation Trigger is built to select high energy electron/positron pairs, which are created very rarely in the center of the quark gluon plasma**
- **The over all architecture is a cylindrical detector around the interaction vertex, where all particle trajectories are searched for very stiff tracks.**
- **The following slides show the resulting electronics chain and on-line processing functionality**
- **The very tight time budget is sketched in the timeline**

QGP Simulation
@ 37 fm/c = $1.2 \cdot 10^{-22}$s
after interaction

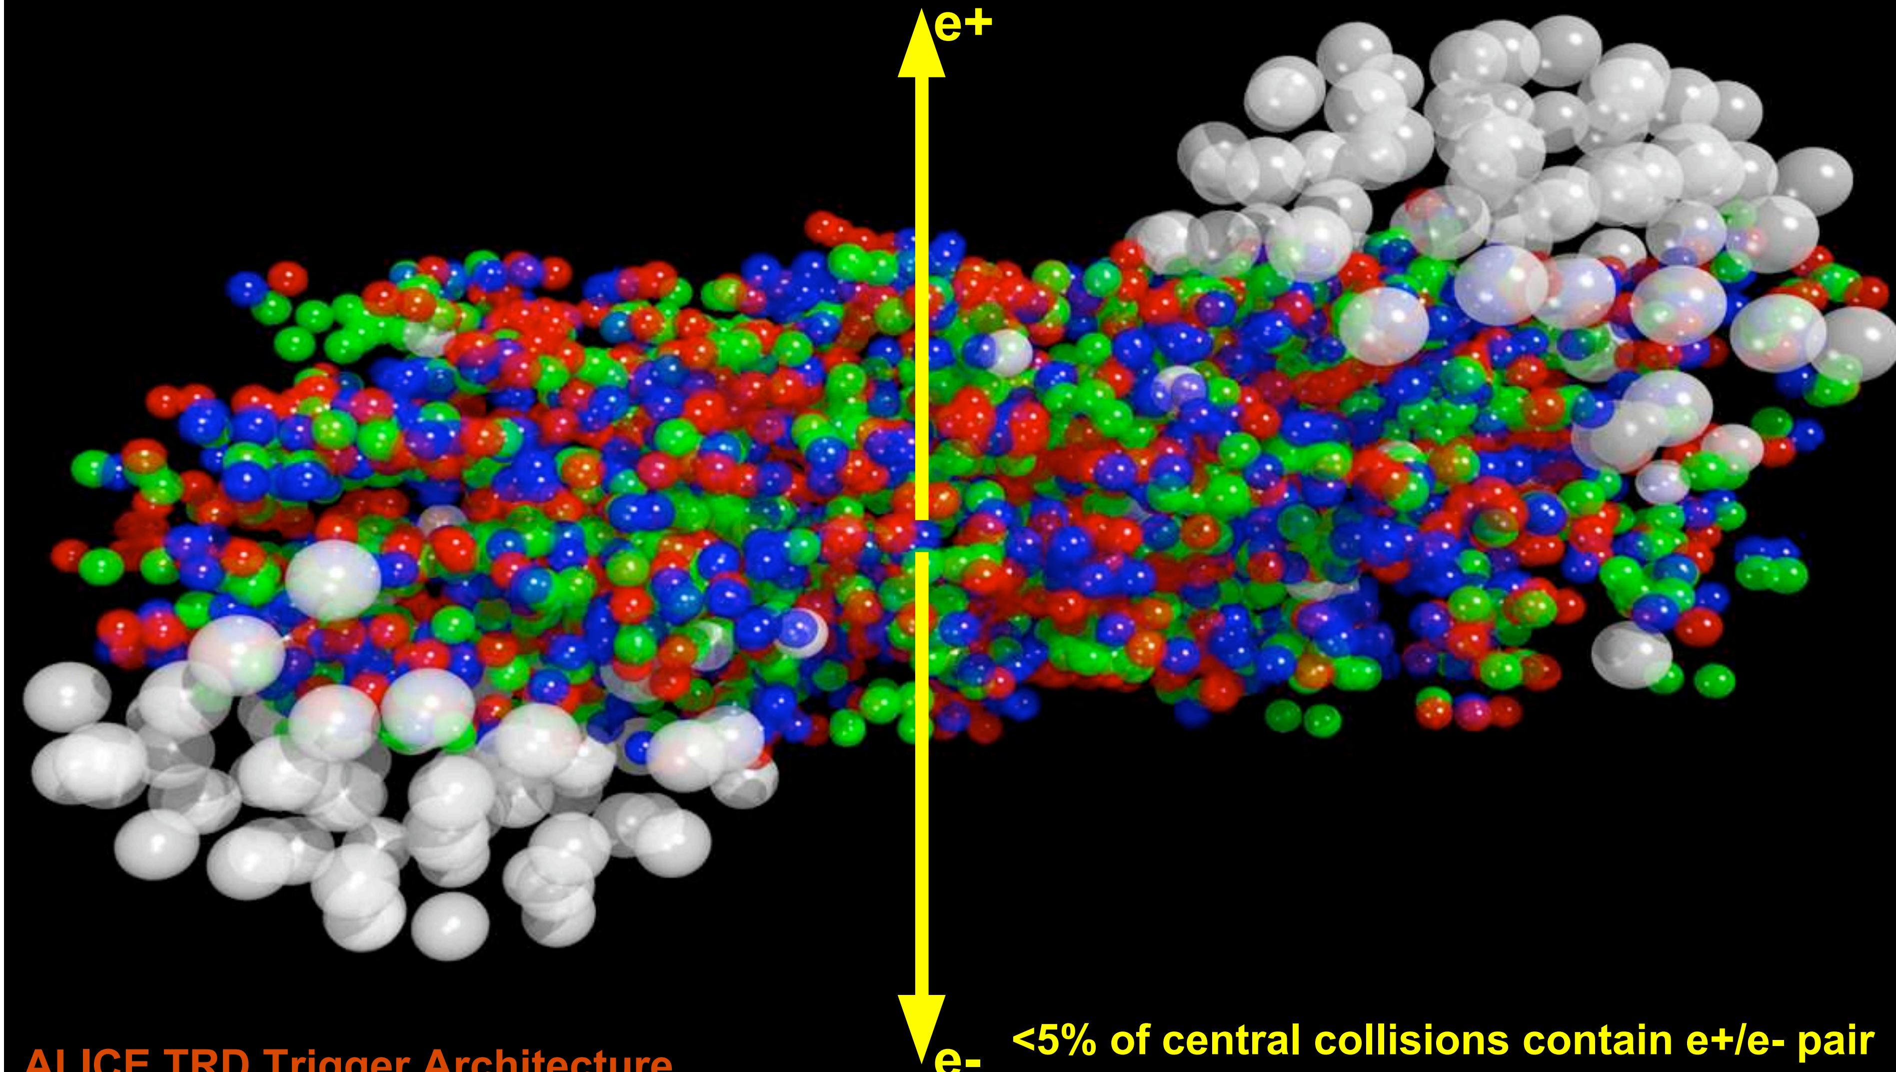

<5% of central collisions contain e+/e- pair

Transition Radiation Detector - TRD

- 1.2 million channels
- 1.425 million ADCs
- peak data rate: 17.8 TB/s
- 64224 MCMs
- computing time: 6 μs

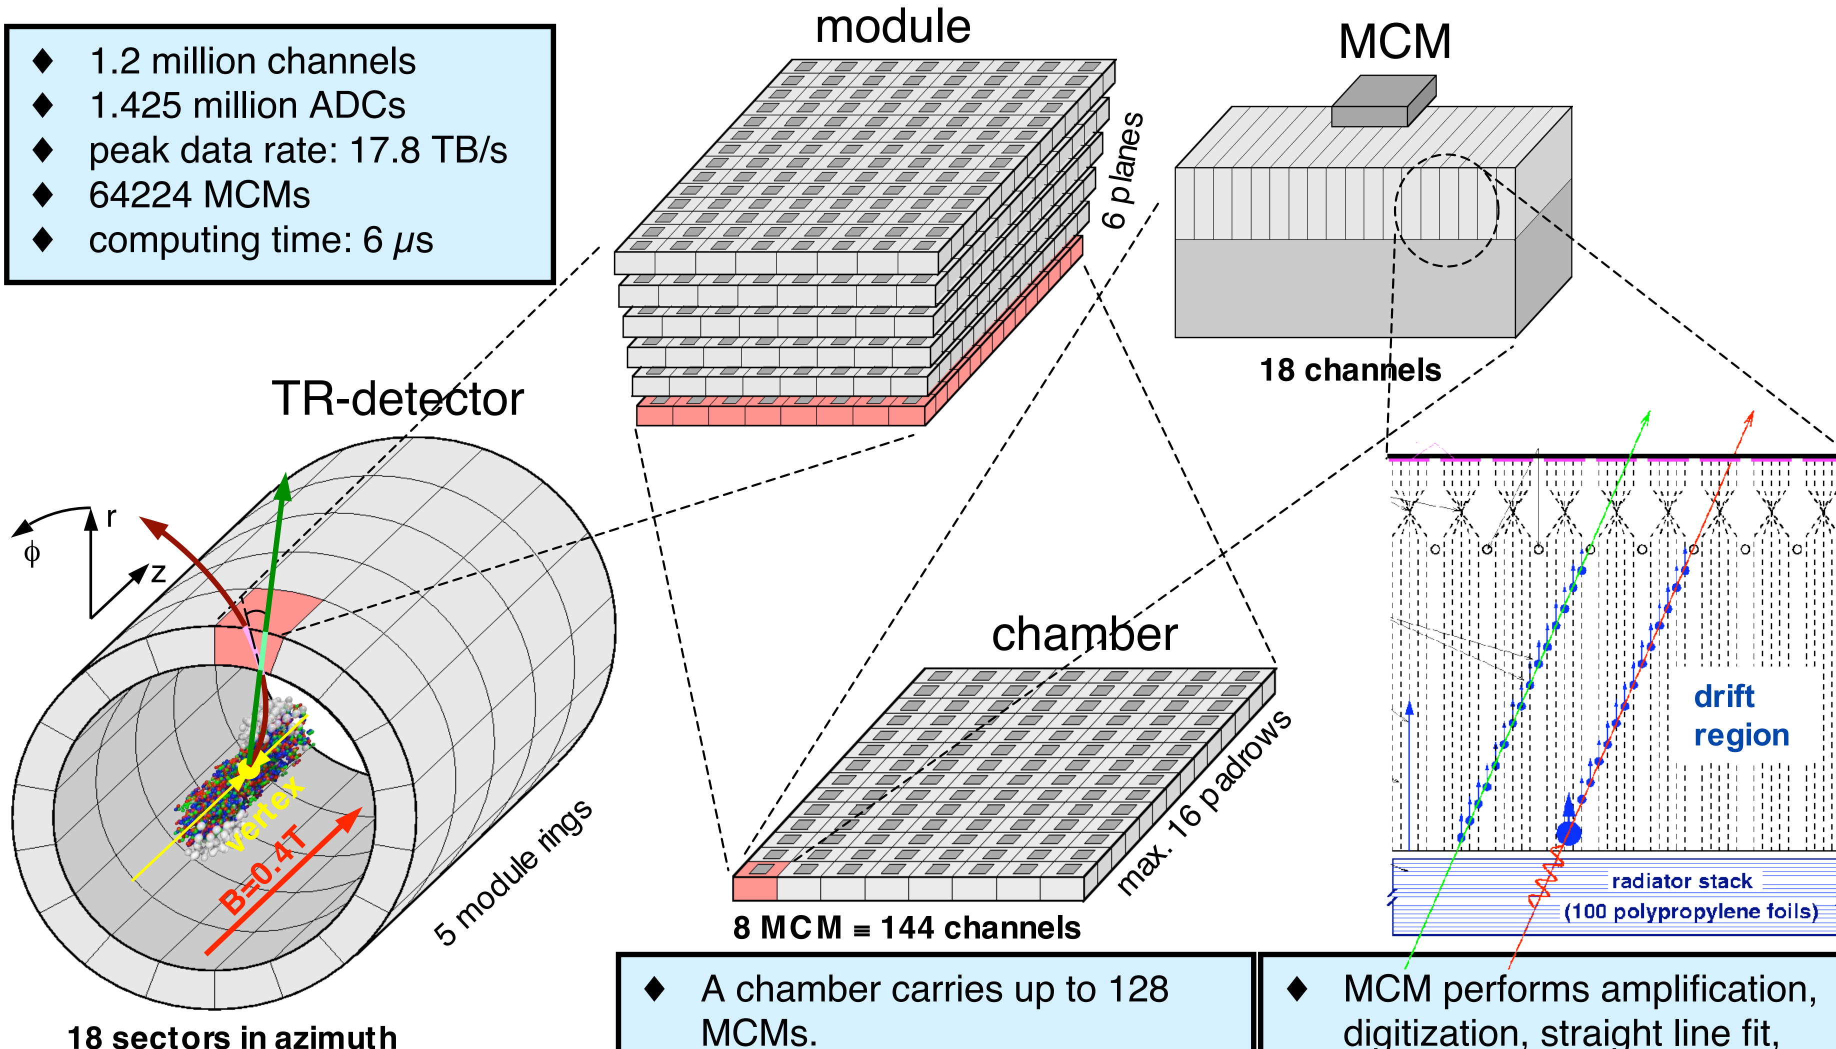

- A chamber carries up to 128 MCMs.
- Each MCM serves 18 channels

- MCM performs amplification, digitization, straight line fit, readout network

TRD Electronics Chain

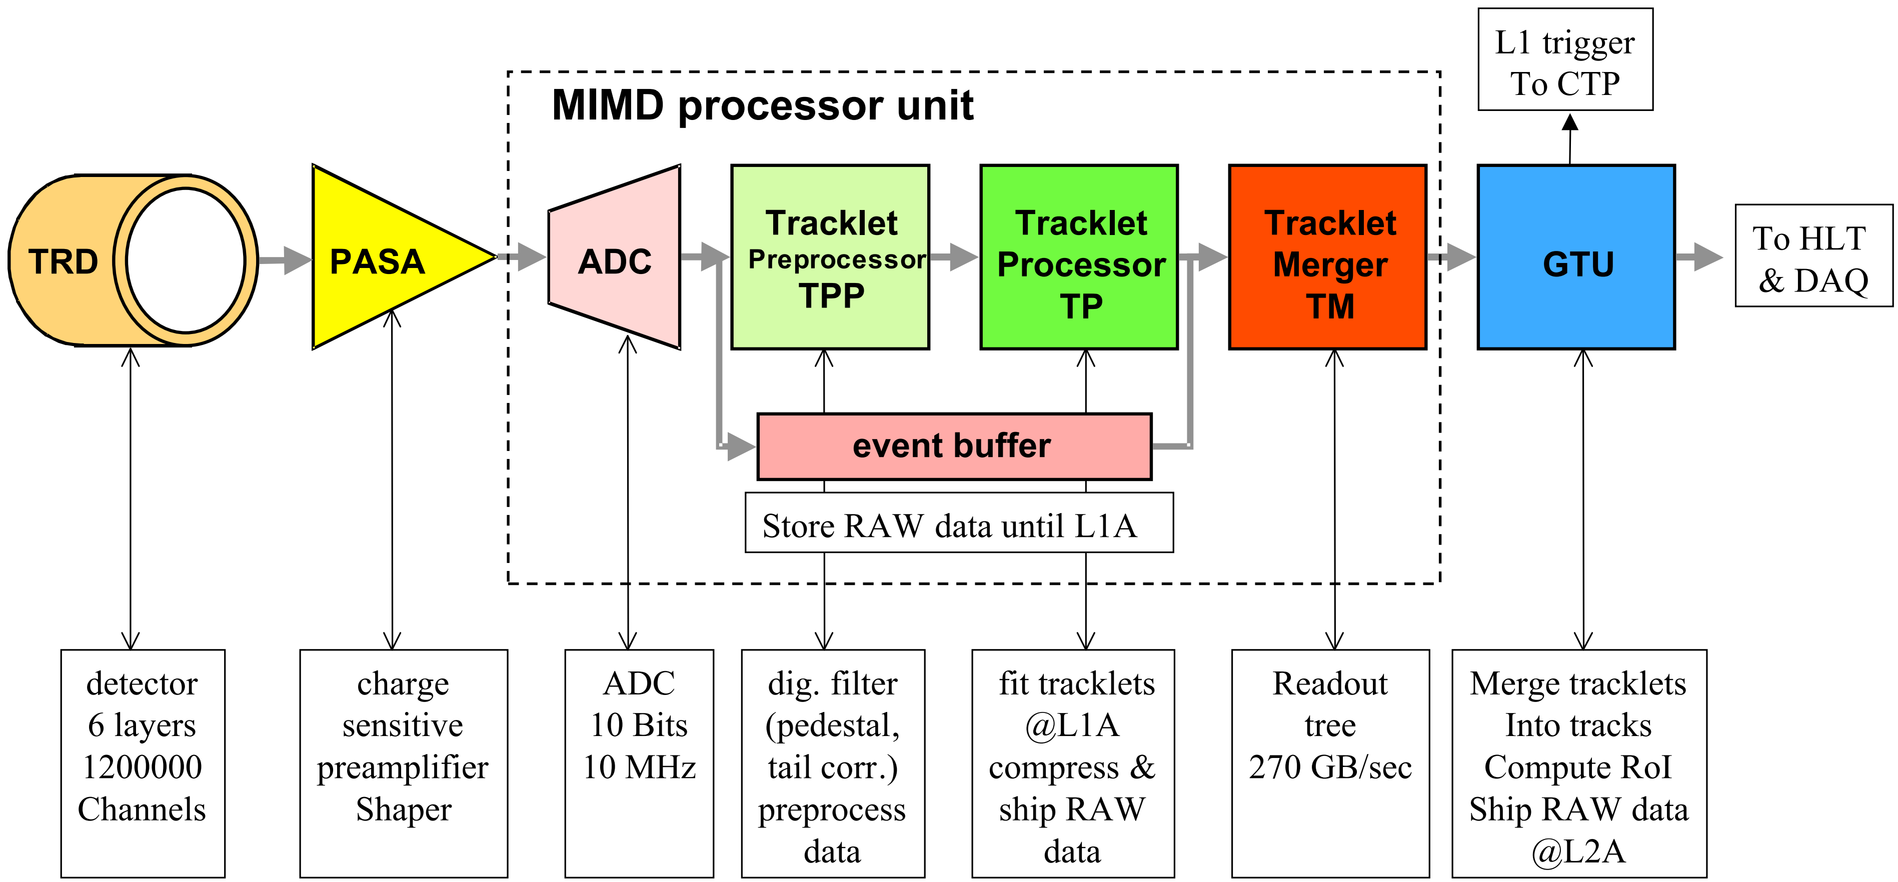

Tracklet Fit (Tracklet Preprocessor)

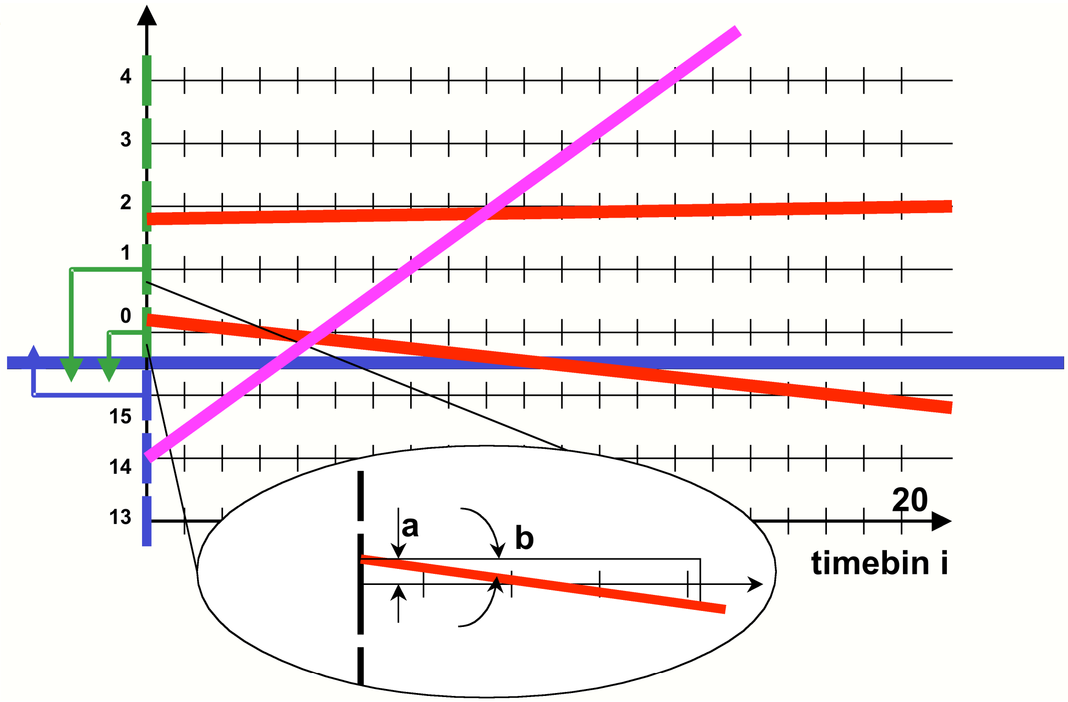

During Drift Time:

N = hit count
y_i = position
$\sum x_i$ = time bin sum
$\sum y_i$ = position sum
$\sum x_i y_i$ = time bin*position sum
$\sum y_i^2$ = position² sum

After Drift Time:

a = intercept

$$a = \frac{\sum x_i^2 \sum y_i - \sum x_i \sum x_i y_i}{N \sum x_i^2 - \left(\sum x_i\right)^2}$$

b = slope

$$b = \frac{N \sum x_i y_i - \sum x_i \sum y_i}{N \sum x_i^2 - \left(\sum x_i\right)^2}$$

χ^2 = track quality

merge track segments in pad row

Pipelined Calculation :

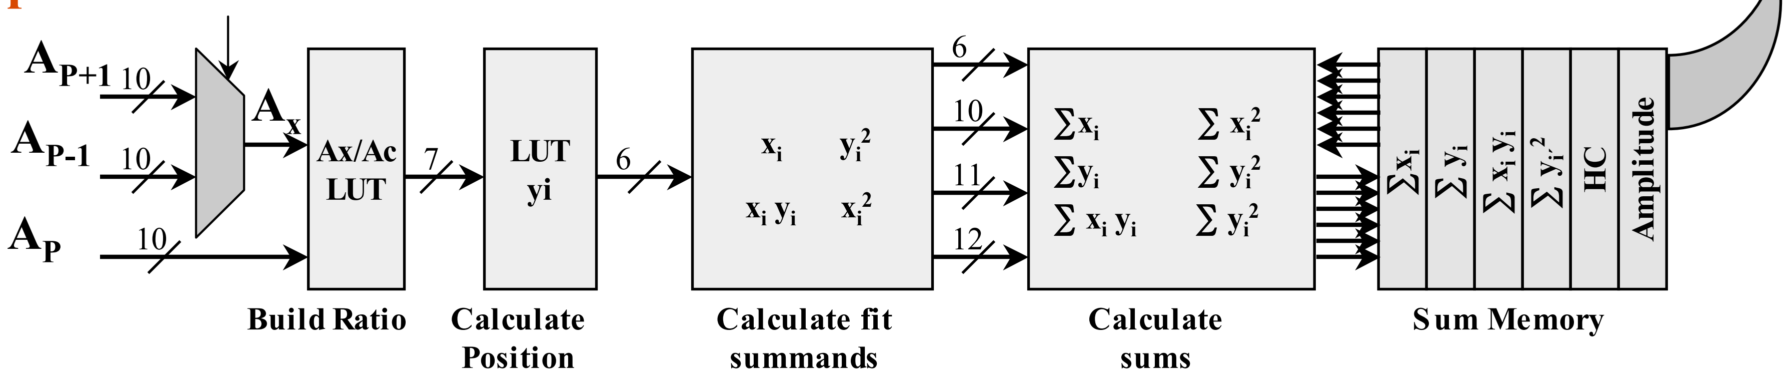

ALICE TRD Trigger Timing

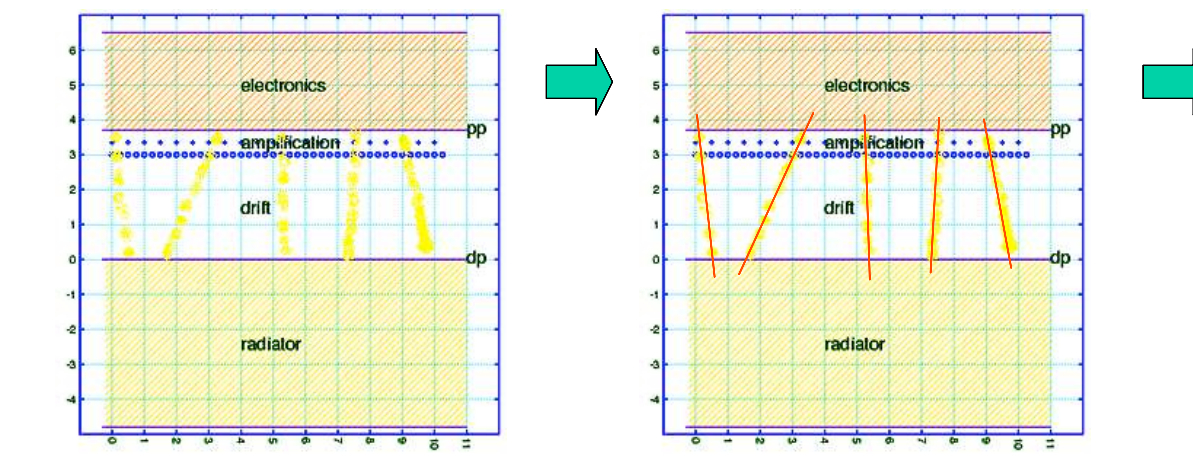

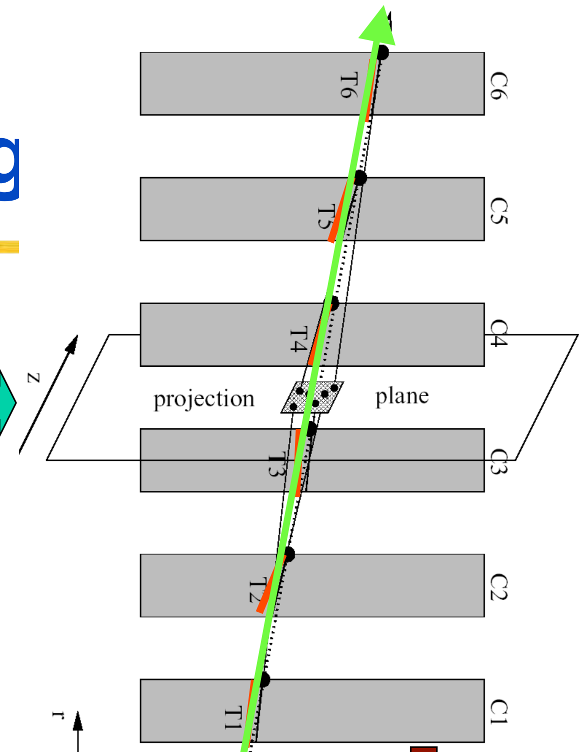

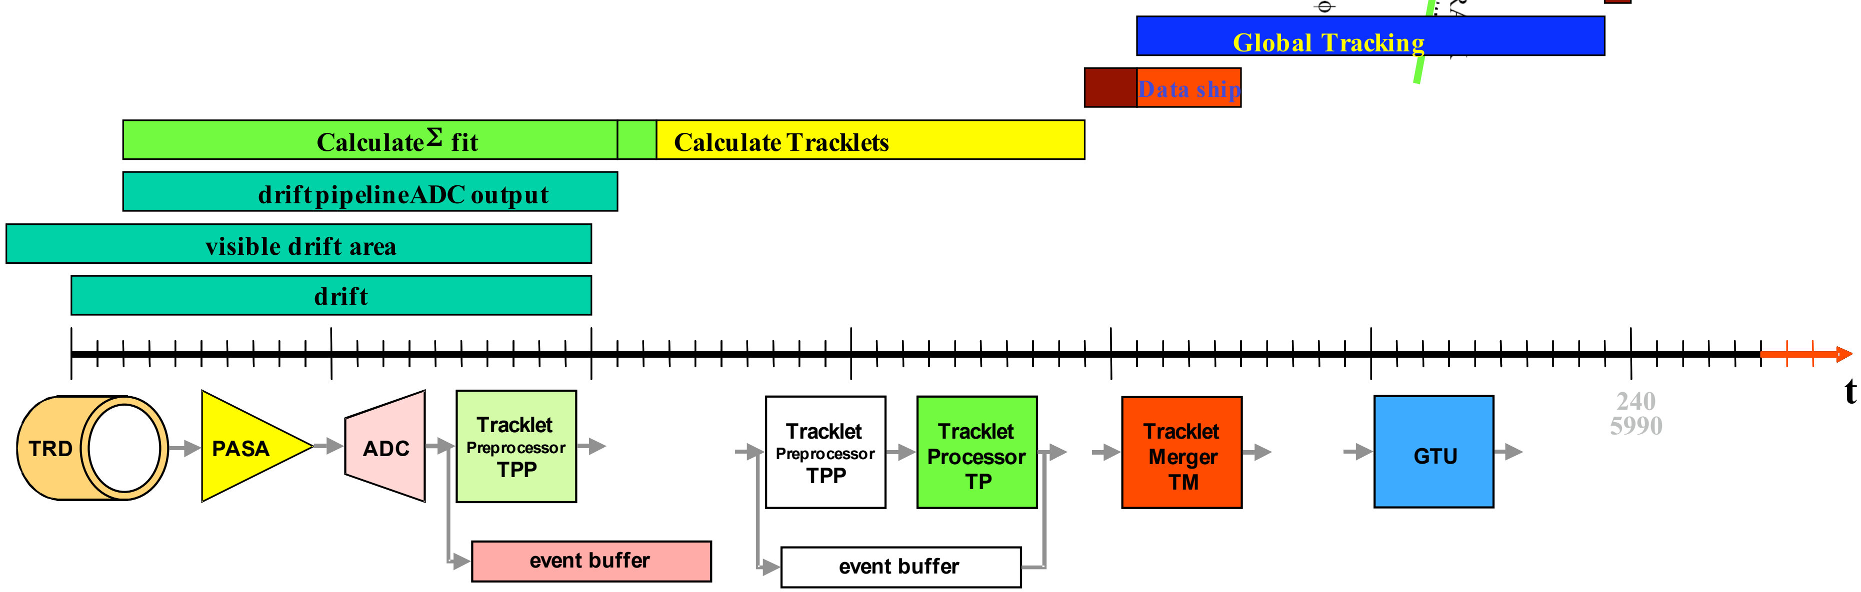

The MIMD Processor

- All required processing is done on the detector in the active area of the experiment, making material budget, power and radiation tolerance essential requirements
- The following slides outline the most important building blocks, like the ADC, digital filters, multi port memory, etc
- The MIMD architecture is outlined. Four CPUs communicate via a global register file and a global quad port memory. In addition each individual processor interfaces to the fitting preprocessor via a private register file, allowing the preprocessor to project its results directly into the processors register file
- The readout of the 260000 processors is performed using a parallel high-speed push architecture
- The last slide in this section is a die photo with an outline of the different functional regions. The chip is in production. So far 30 wafers have been produced in the UMC 180 nm process.

ALICE TRD MCM

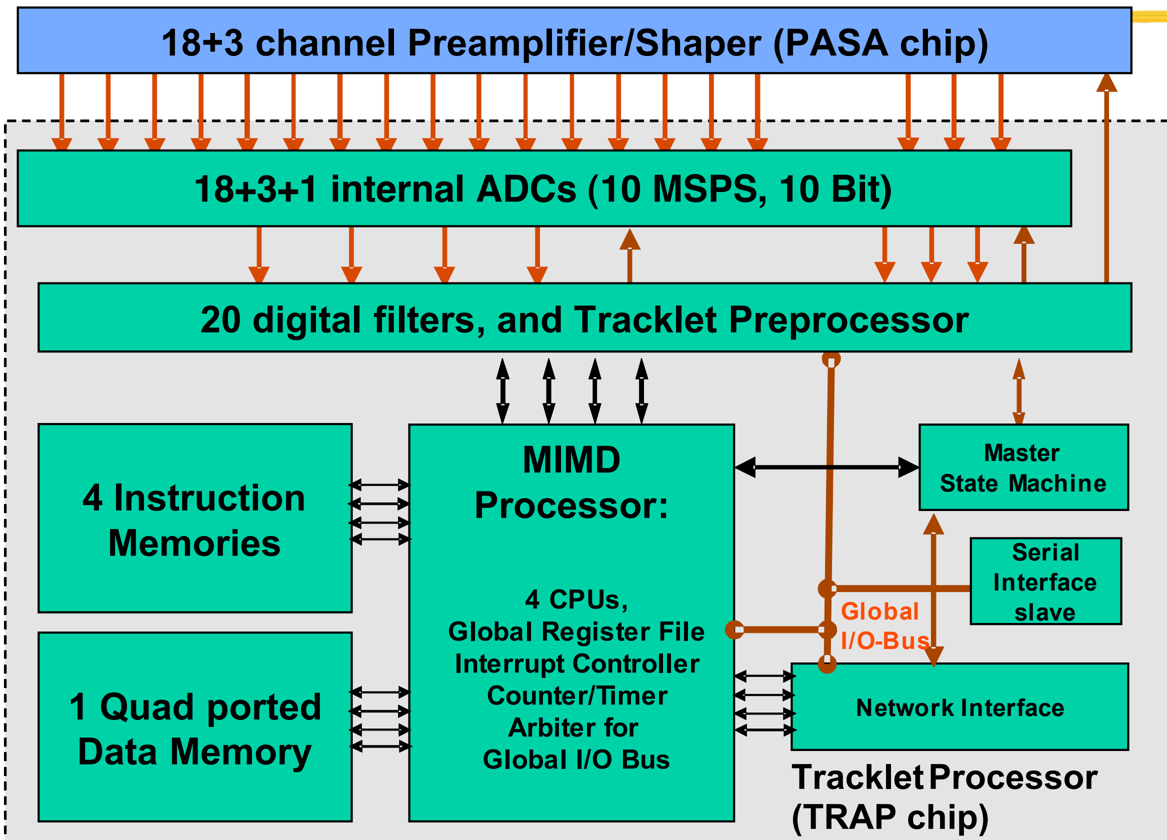

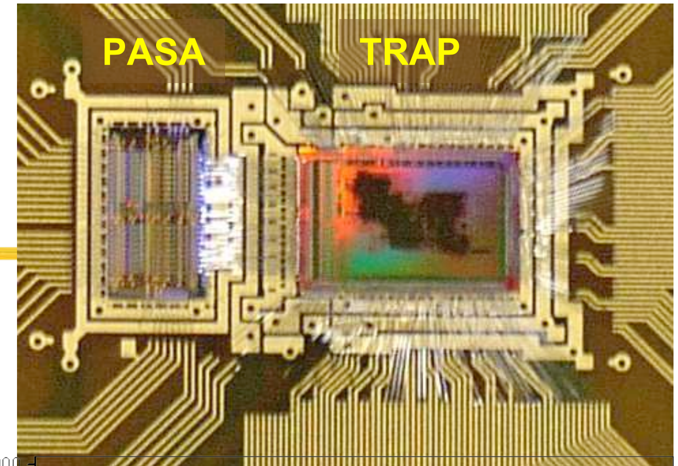

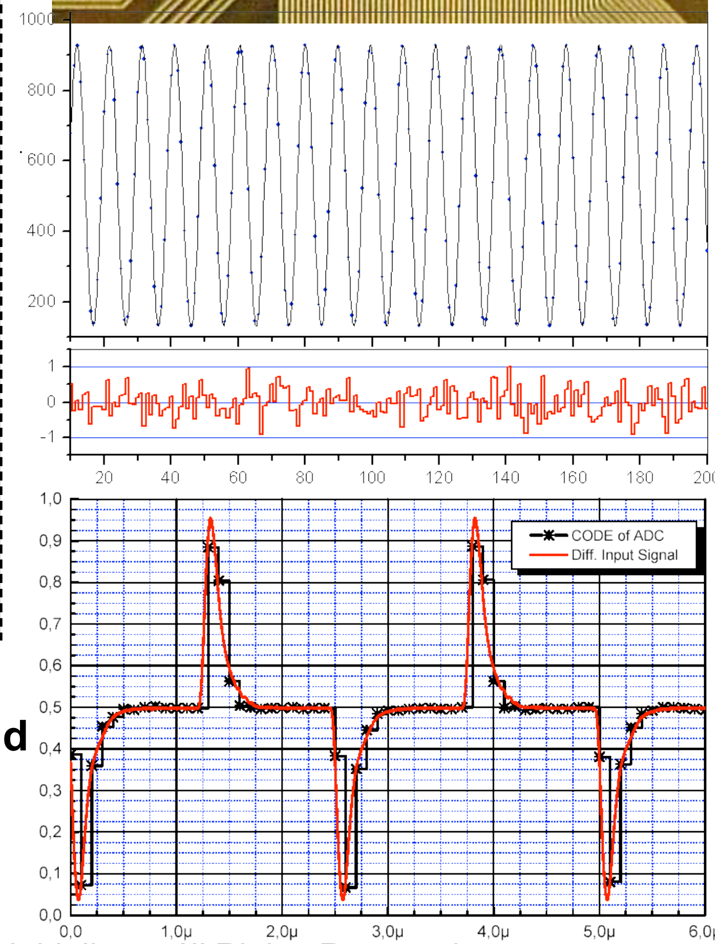

- Deep submicron UMC180nm
- Concurrent operation of 4 120 MHz RISC processors and 21 10-Bit low-power ADCs @9.5 ENOB
- PASA and digitizing/processing chip on same MCM
- Very low cost

ADC

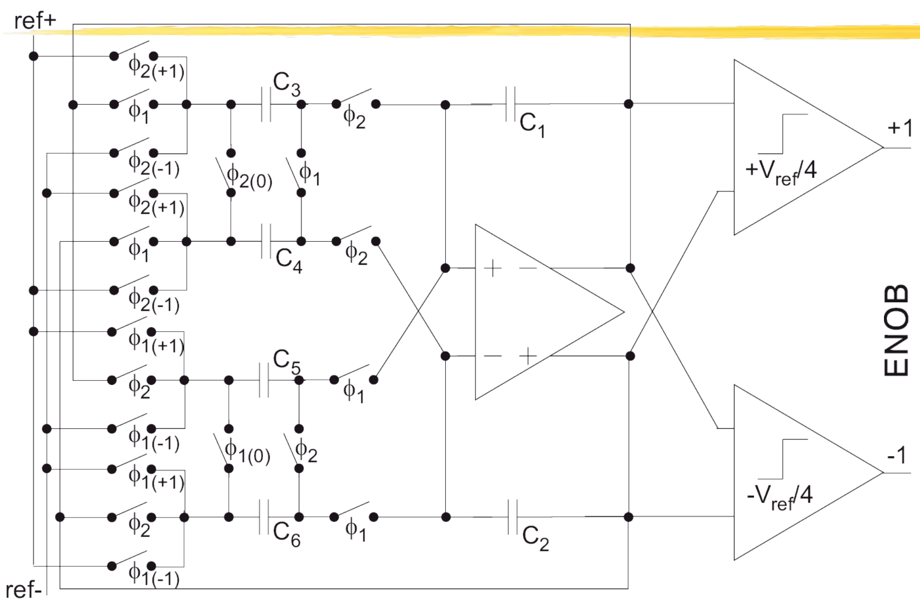

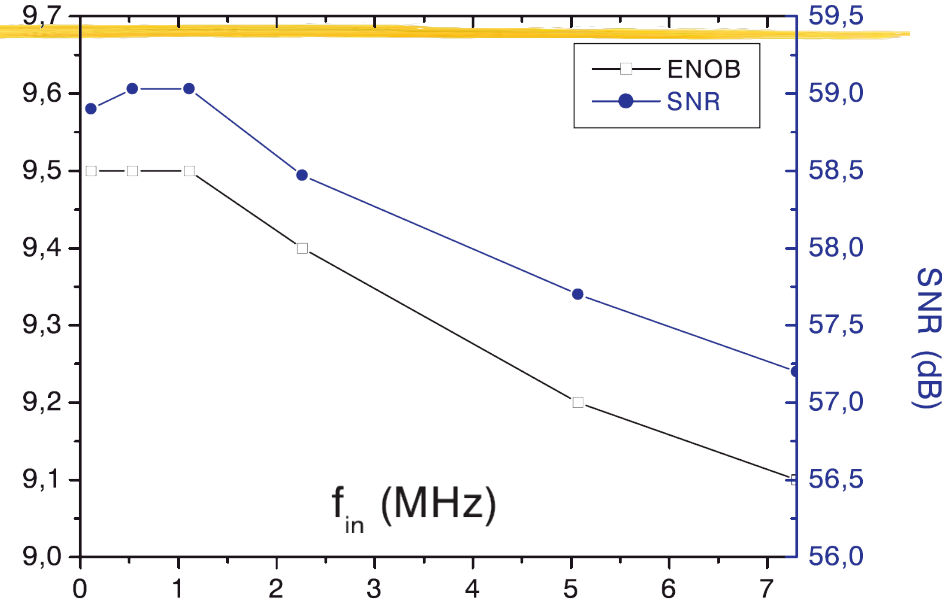

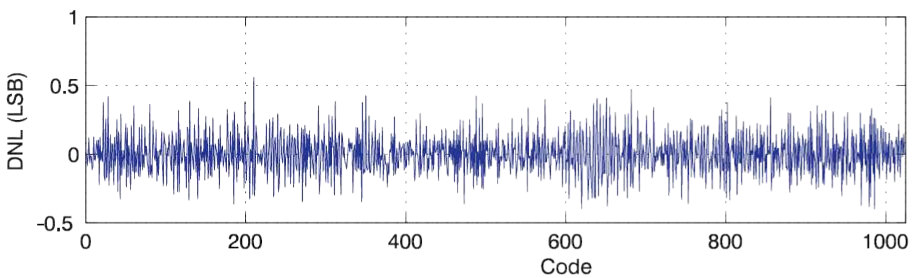

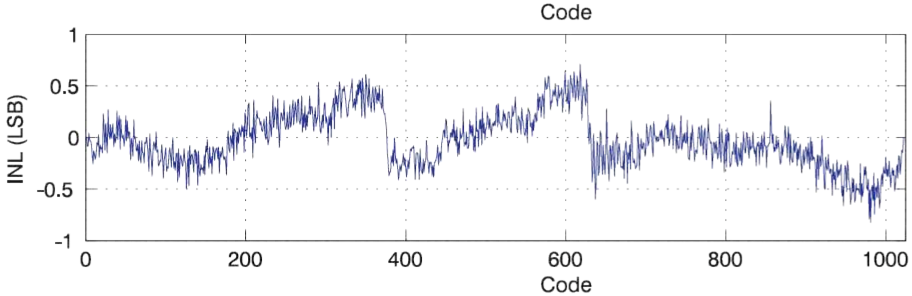

Resolution	**10bit**
Sampling Rate	**10.4MS/s**
Process	**0.18um CMOS + MIMCAPS**
Size (per ADC)	**0.11mm^2**
Power	**9.5mW @ 10.4MS/s**
Input (pro.)	**+/-1V - +/-1.4V**
Supply	**1.8V, 3.3V**
DNL	**-0.4 / +0.6 LSB**
INL	**-0.8 / +0.7 LSB**
ENOB @1MHz Signal	**9.5 bit**
SNR @1MHz Signal	**59.2 dB**
SFDR @1MHz Signal	**73.0dB**
THD @1MHz Signal	**-69.5dB**

TRAP – Digital Filter

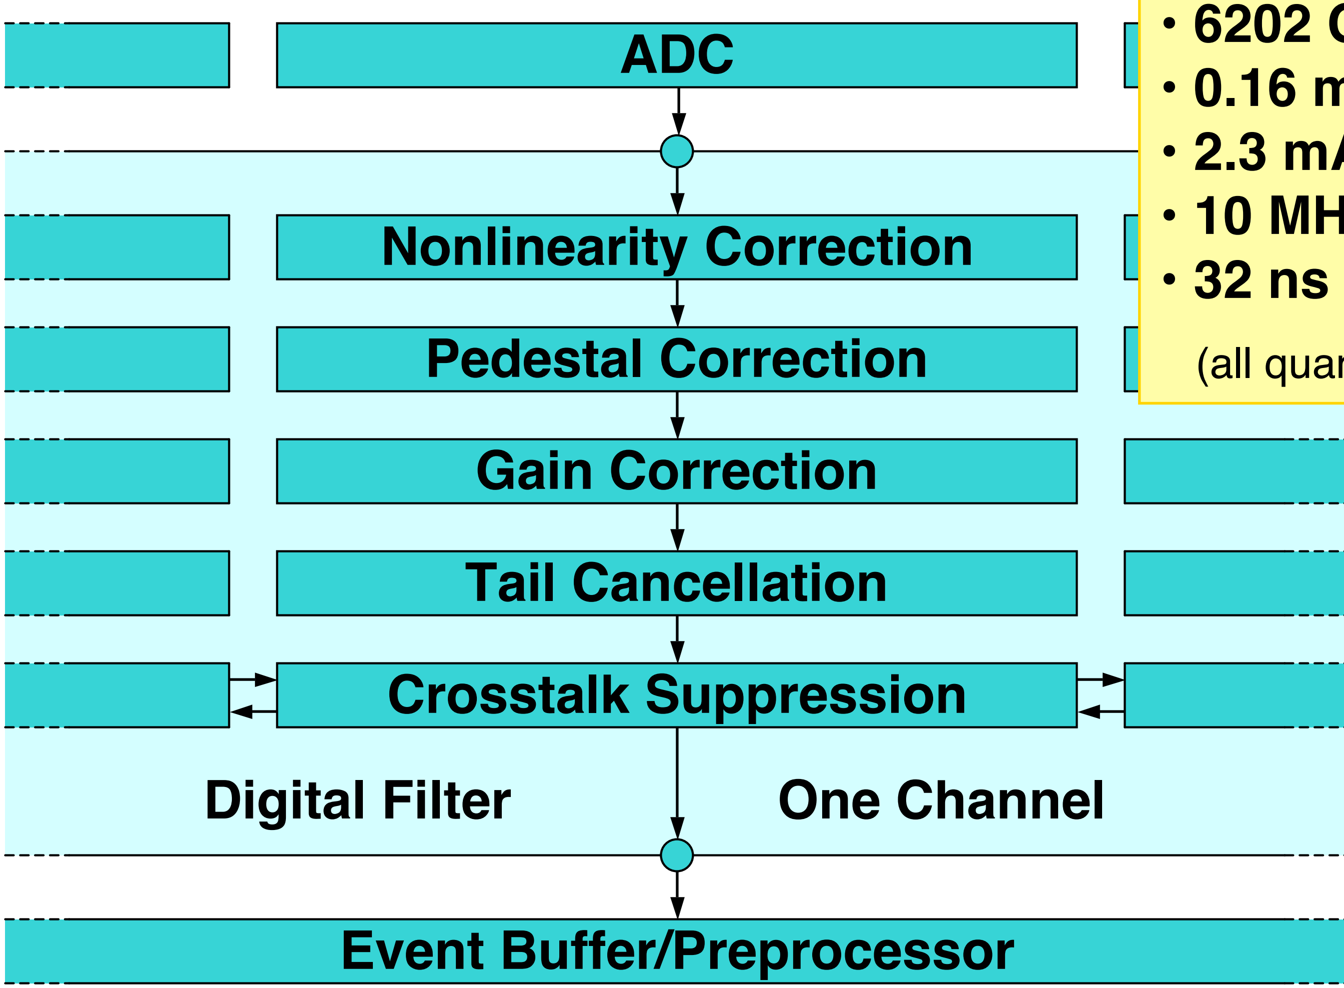

- 6202 Cells
- 0.16 mm²
- 2.3 mA
- 10 MHz
- 32 ns Latency

(all quantities per channel)

Marcus Gutfleisch

TRAP – Tracklet Preprocessor

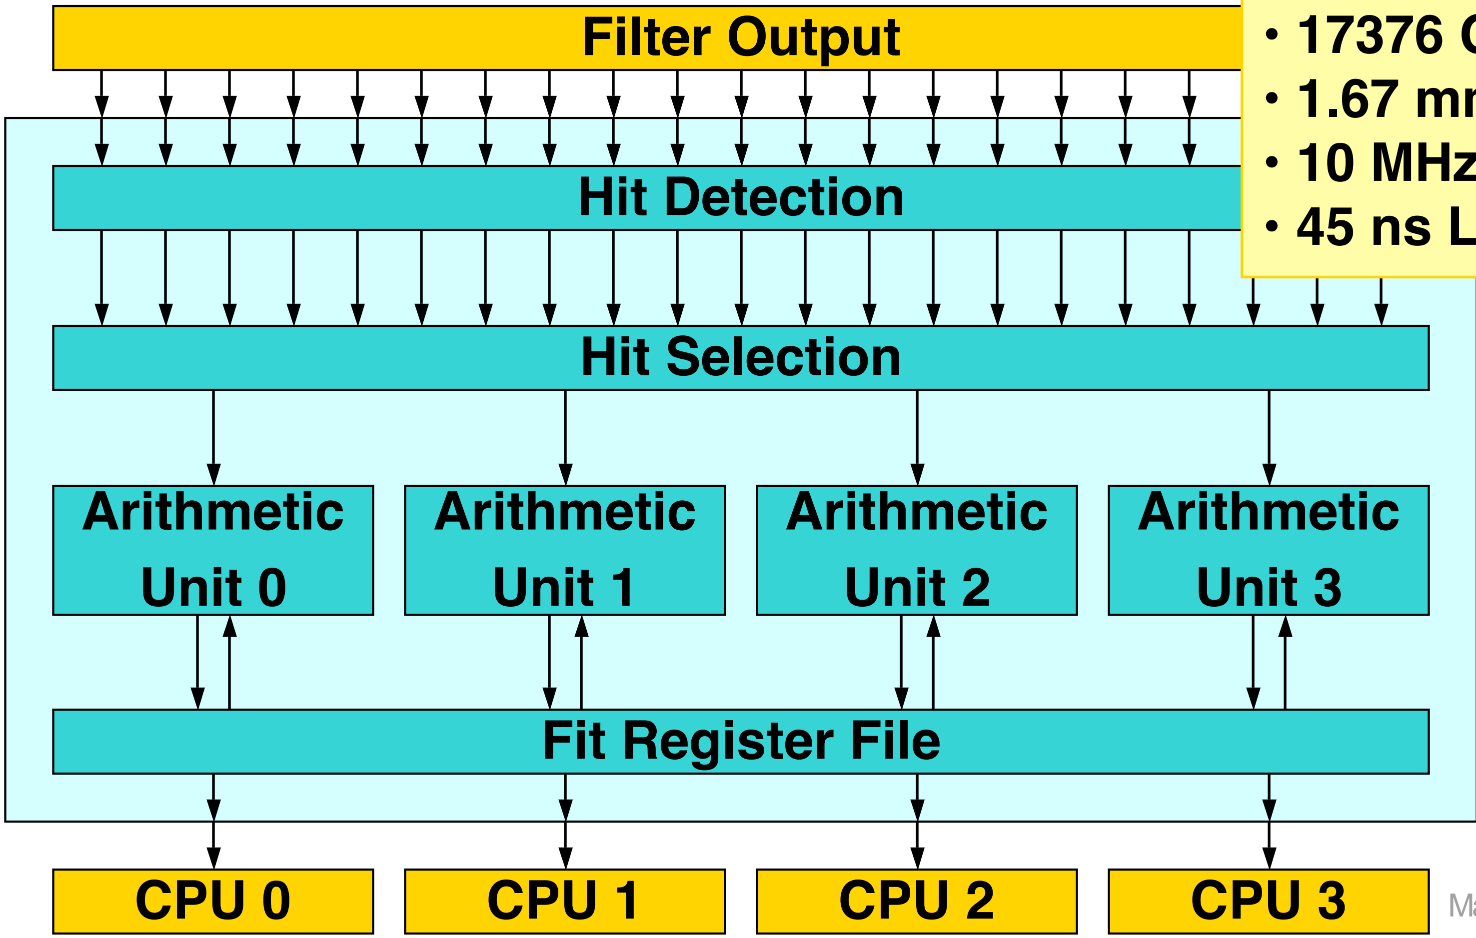

- 17376 Cells
- 1.67 mm^2
- 10 MHz
- 45 ns Latency

Marcus Gutfleisch

The MIMD Architecture

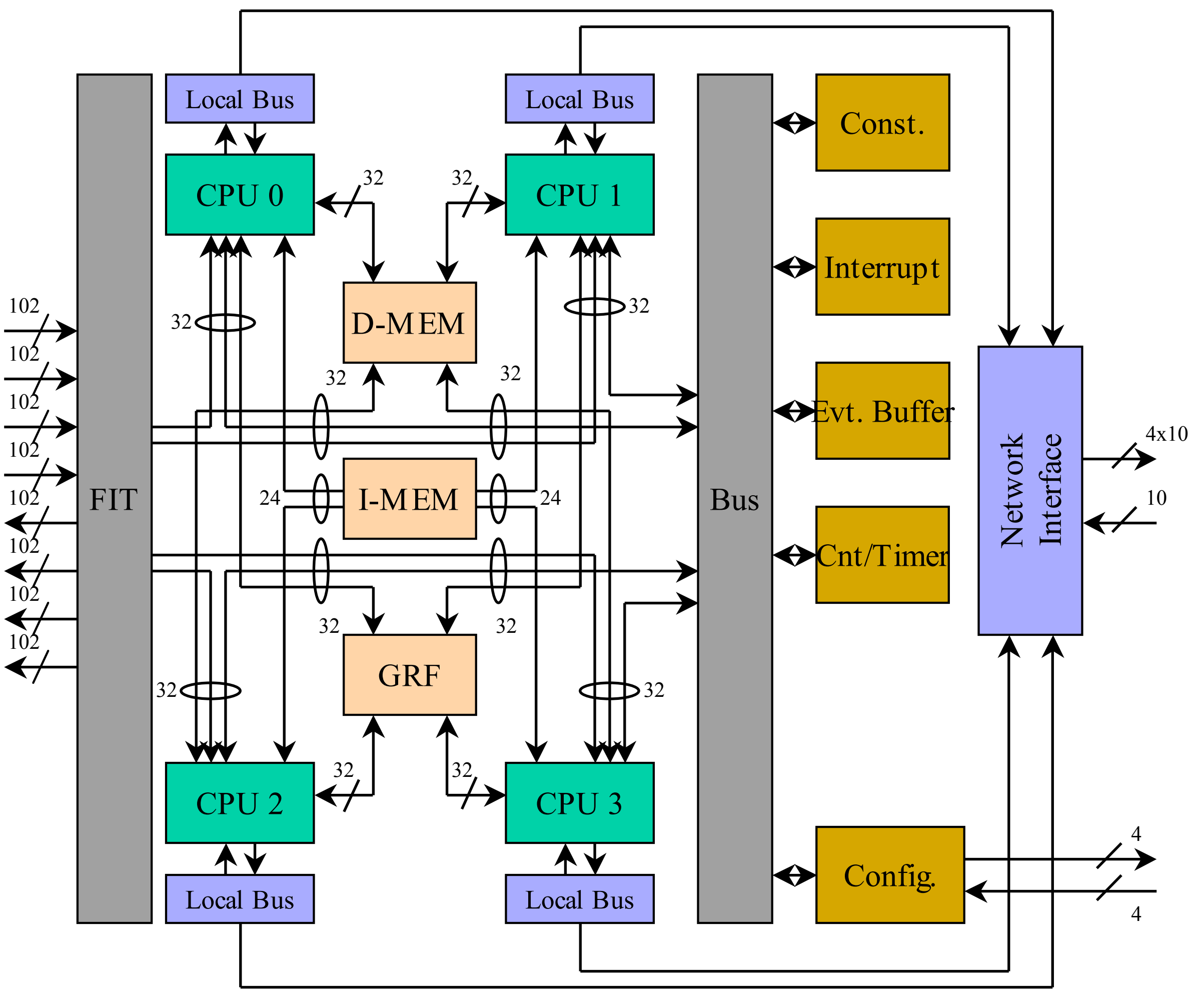

- Four RISC HARVARD CPU's
- Coupled by Registers (GRF) and Quad ported data Memory
- Register coupling to the Preprocessor
- Global bus for Periphery
- Local busses for Communication, Event Buffer read and direct ADC read
- I-MEM: 4 single ported SRAMs
- Serial Interface for Configuration
- IRQ Controller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

Quad Port Memory

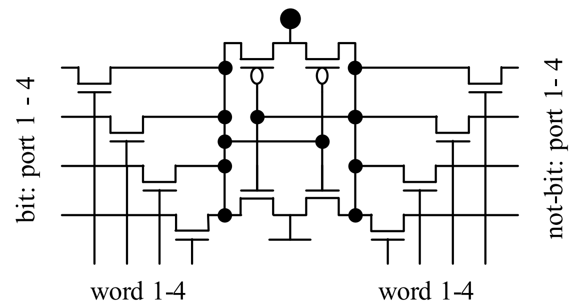

Features

- four independent bi-directional and synchronous ports
- designed for 120 MHz
- Area one bit cell: 5,08 * 4,86 µm
- Overall area less then:
  1900 * 800 µm
  (1024 words with 32+7 bits)

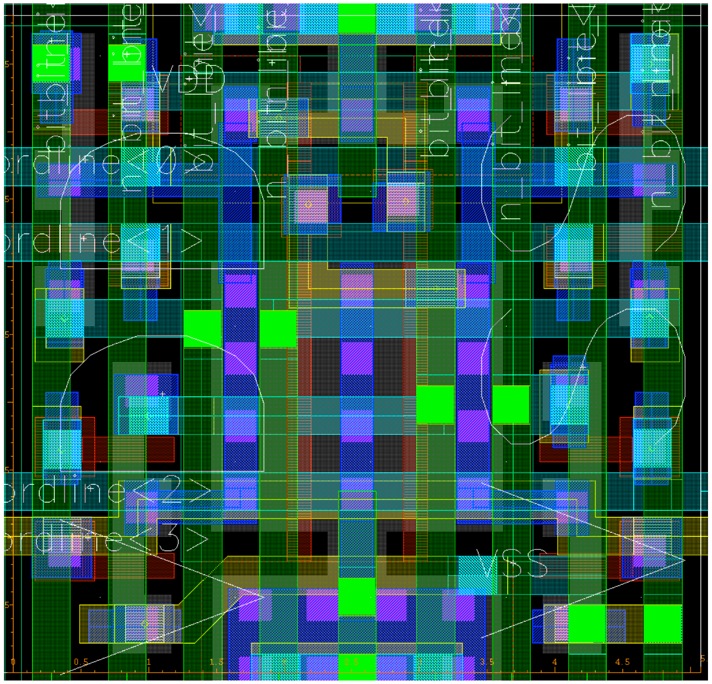

Tracklet→Track→Trigger (Tracklet Merger)

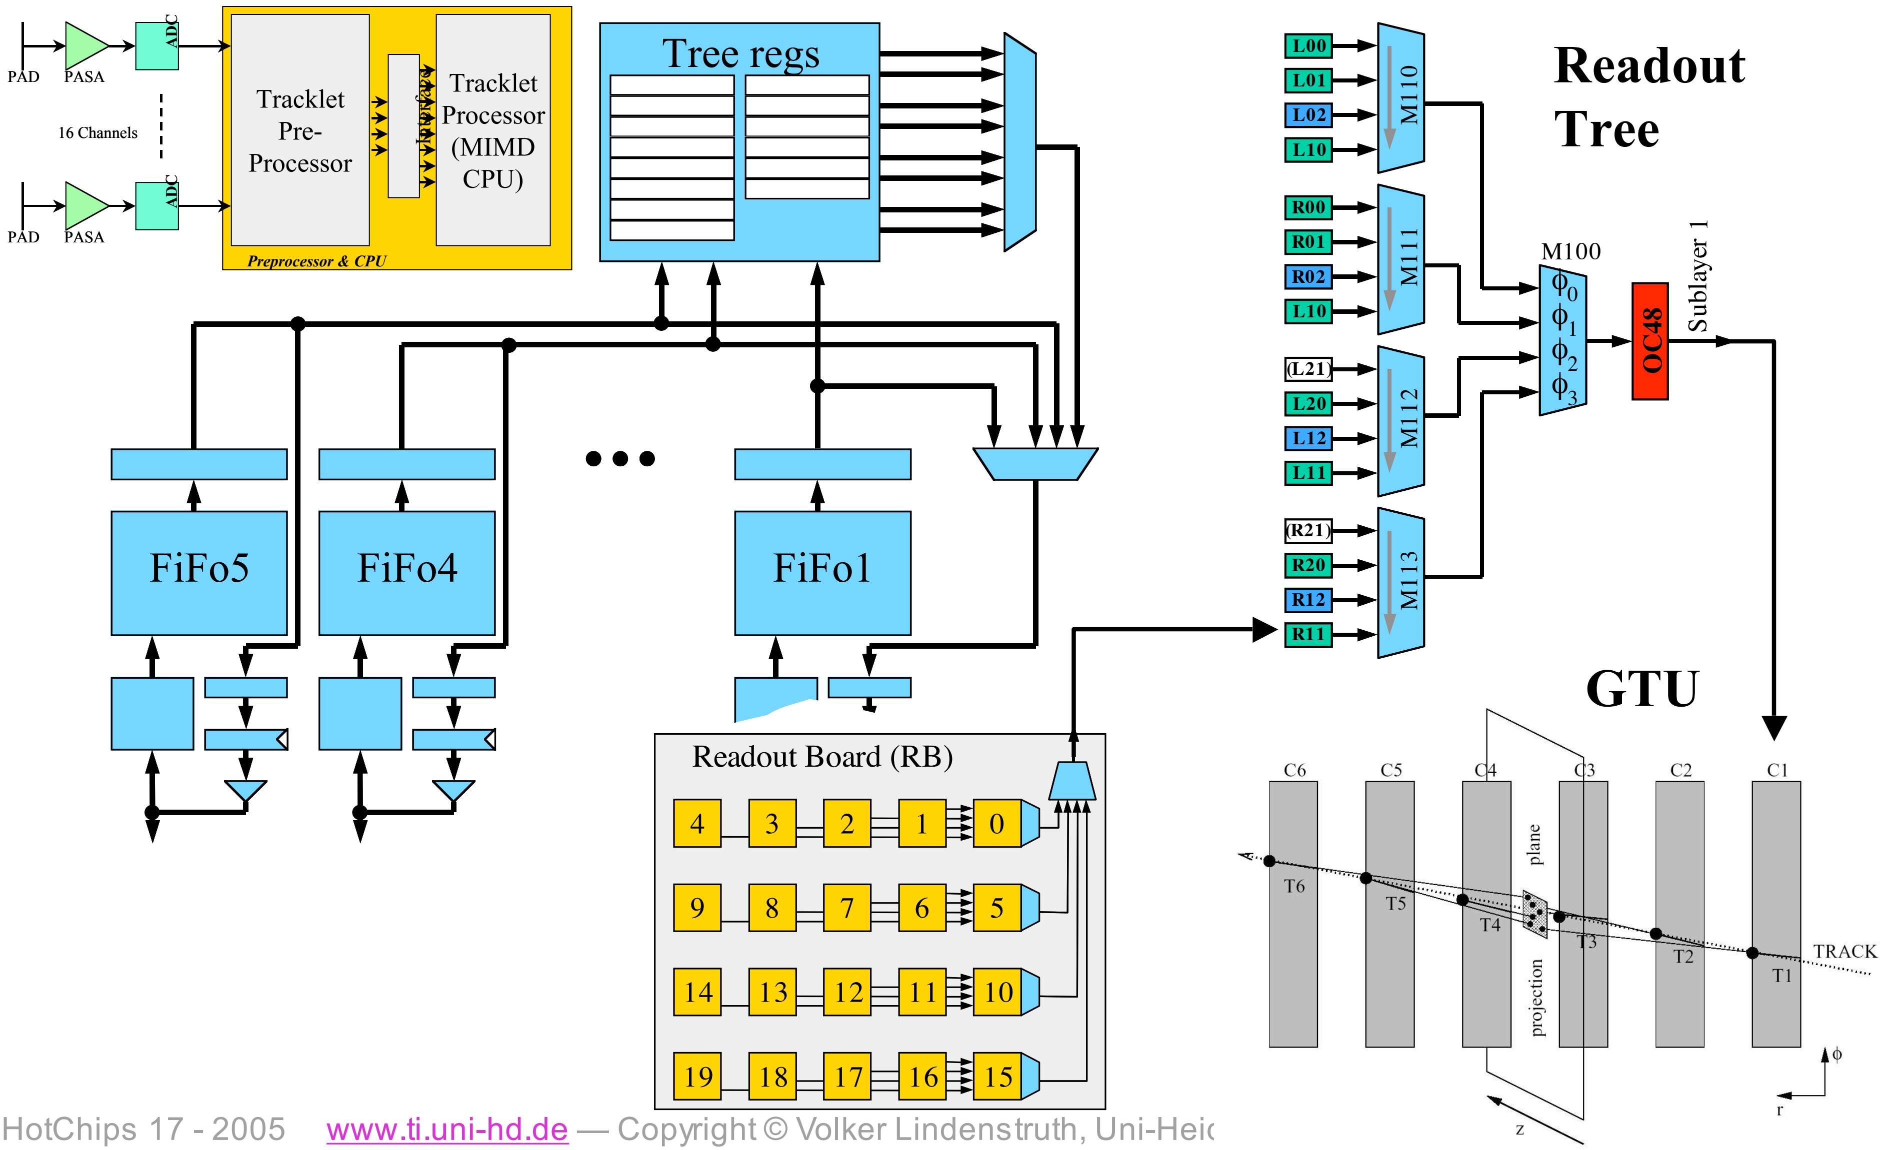

Programming Model

- **Available processing time is 1µs / 120 instructions**
- **All programming in assembler**
- **MIMD simulator developed for software development and performance benchmarking**
- **Initialization and control of the 250000 processors in 65000 chips is performed by 540 single board linux computers, integrated on the detectors**

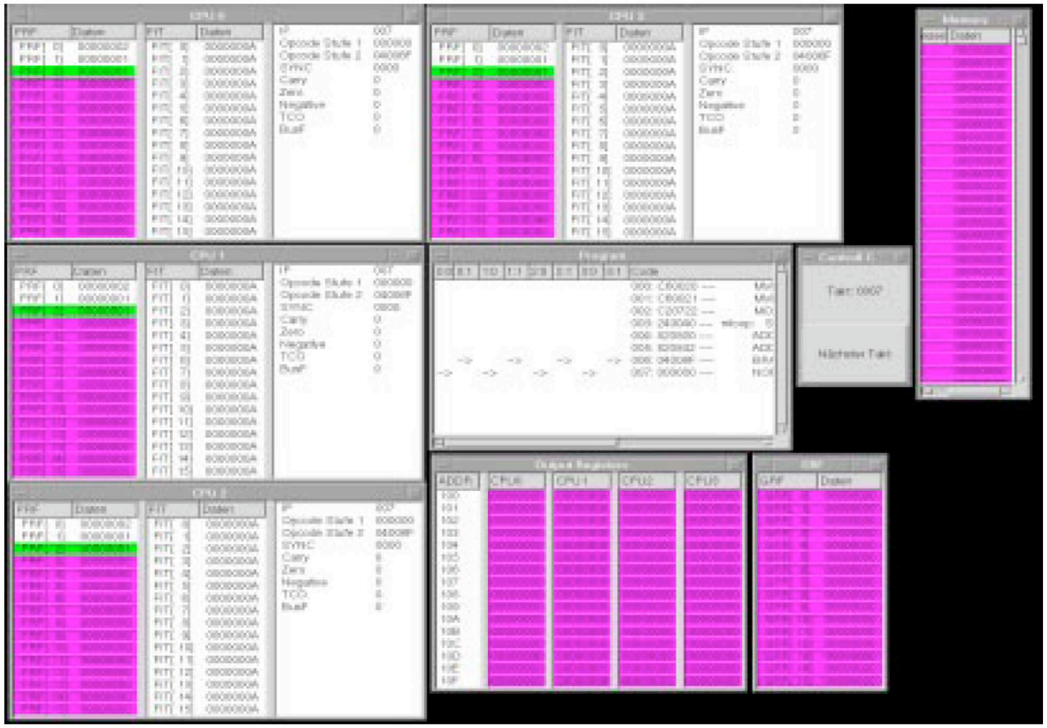

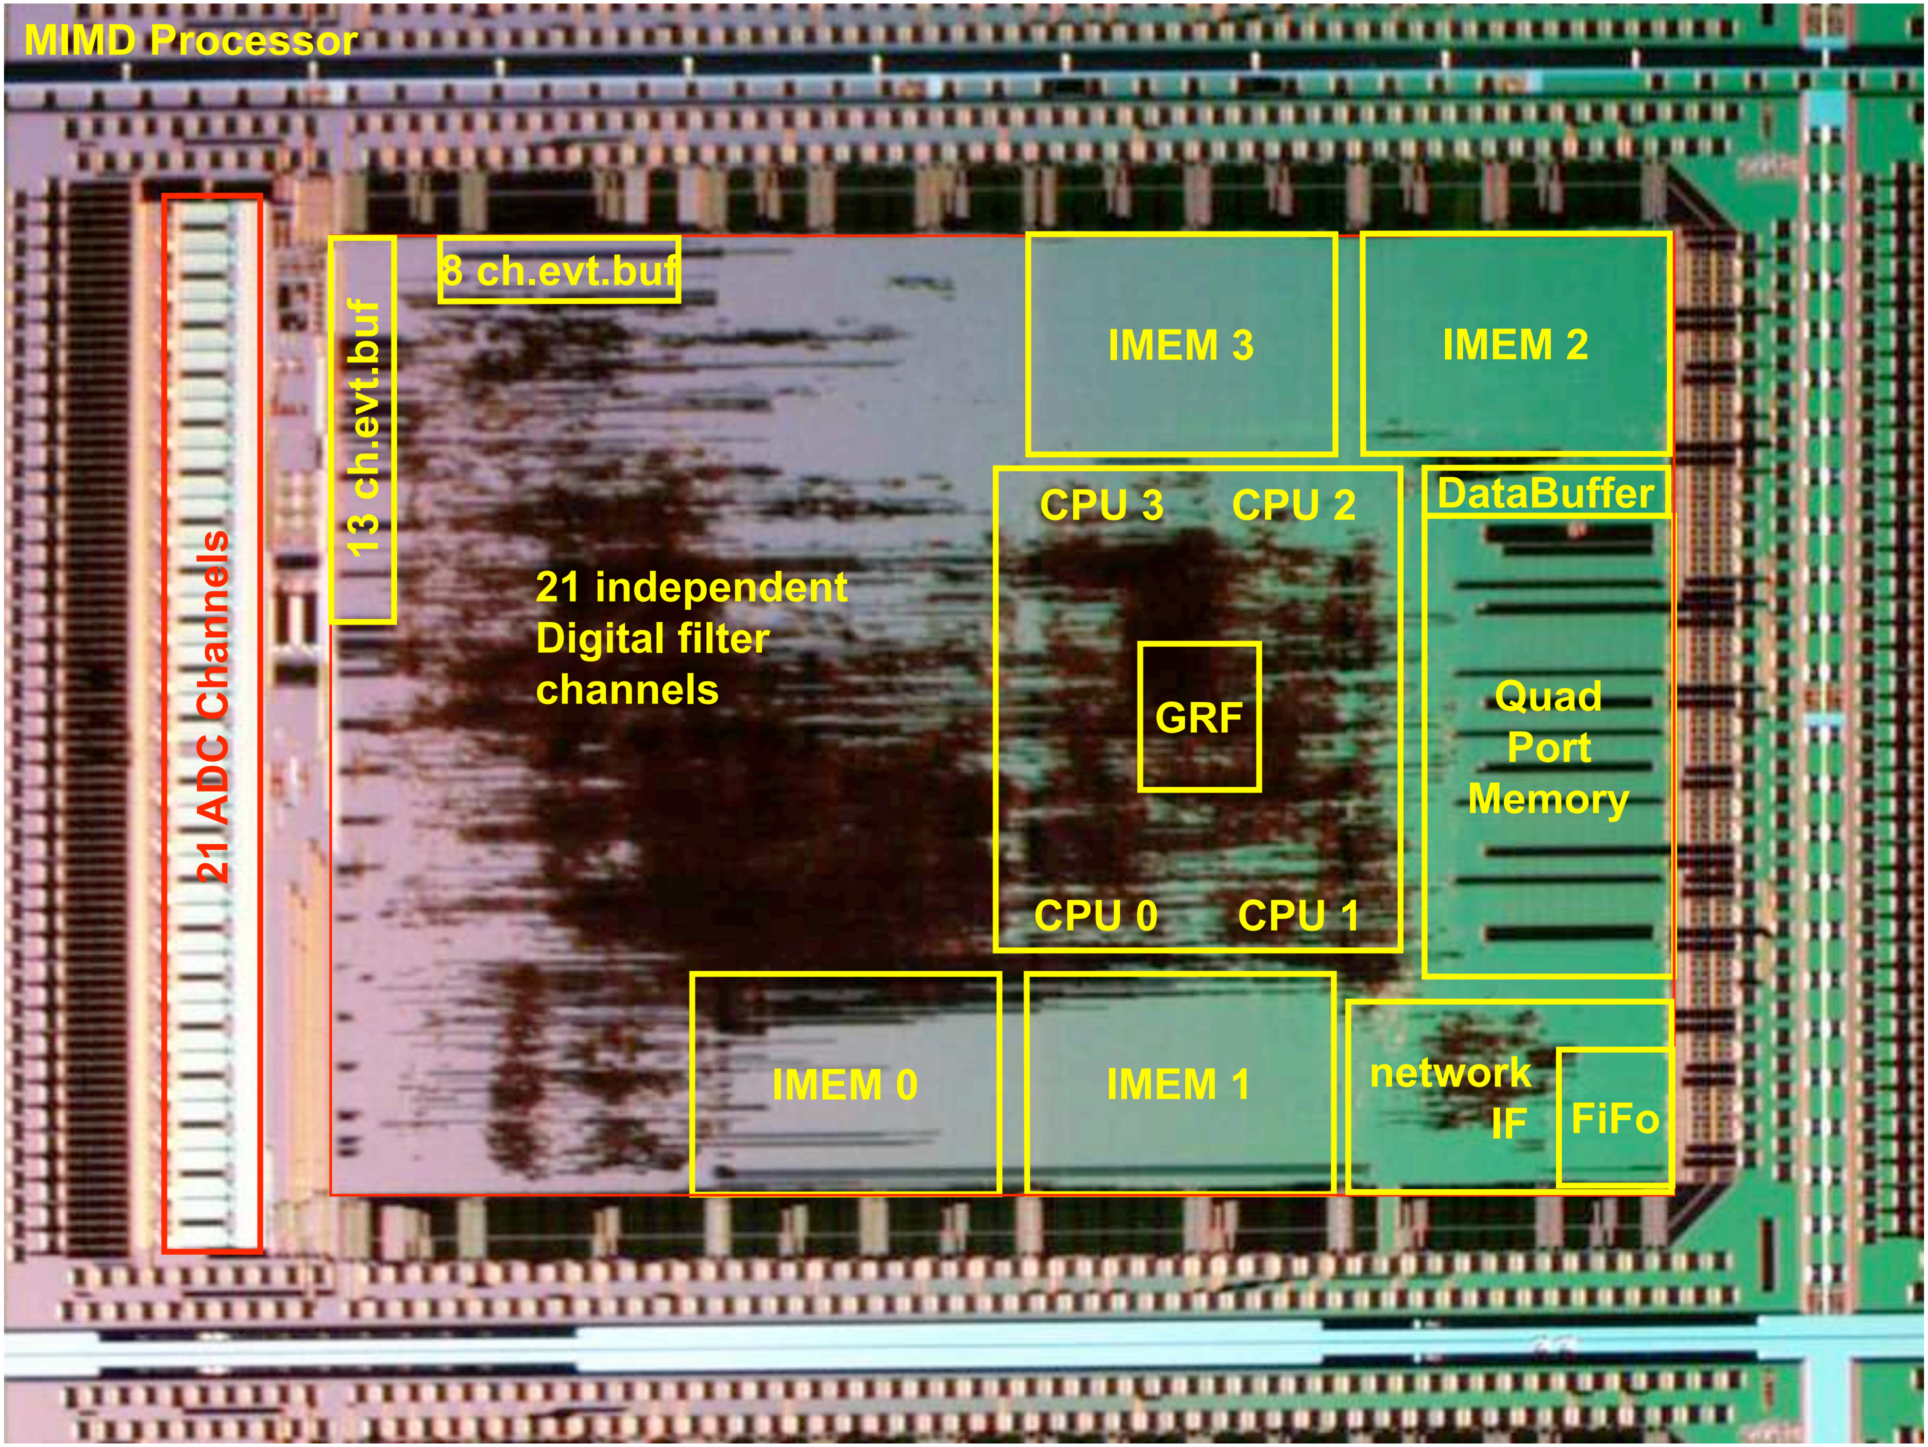
MIMD Processor
21 ADC Channels
13 ch.evt.buf
8 ch.evt.buf
21 independent
Digital filter
channels
IMEM 3
IMEM 2
CPU 3
CPU 2
DataBuffer
GRF
Quad
Port
Memory
CPU 0
CPU 1
IMEM 0
IMEM 1
network
IF
FiFo

ALICE TRD Global Tracking Unit

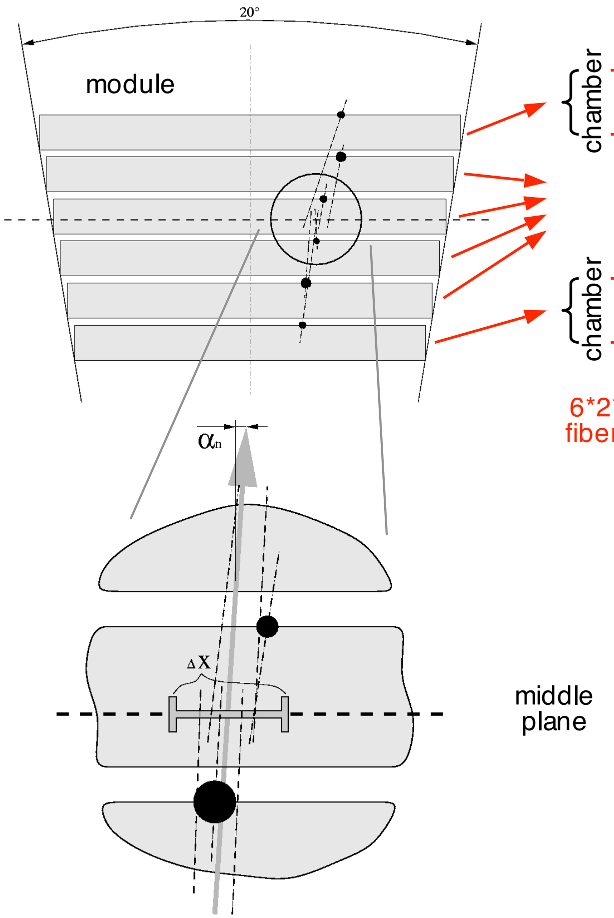

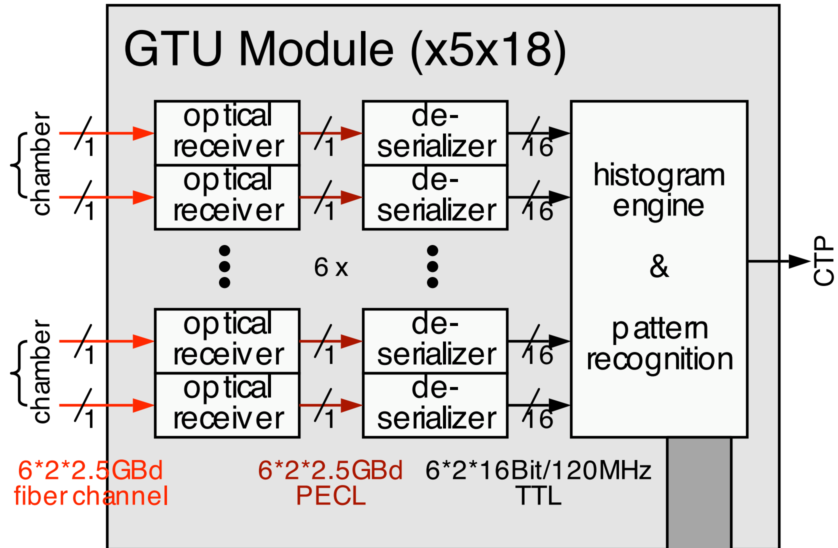

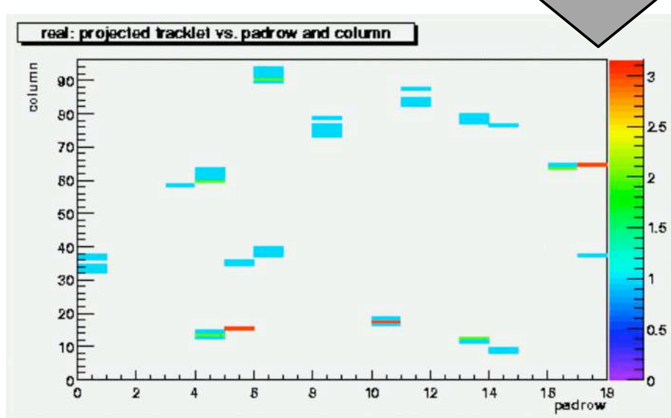

- global tracking is performed module wise (540 modules).
- input of one GTU are 6*2 2.5GBd optical fiber channel links which are converted into 6*2 16Bit/120Mhz data streams.
- readout order of individual chambers are optimized for histogramming and pattern recognition in parallel.
- histogramming engine fills a data field with tracklet candidates regarding their position.
- looking for patterns with
  - a cluster of at least 4 candidates
  - adequate similar angle
  - high momentum
  - adequate quality
- quality and position of found hit is sent to central trigger processor
- GTU operates also for raw data readout and compression

Tests and Measurements

- **The following slides show the integration of the MIMD processor chip with a preamplifier on a multi-chip module, the integration of these MCMs on readout boards and their integration on the chambers.**
- **A setup of one detector stack in a real beam experiment is shown next. The final detector will encompass 90 stacks.**
- **The measurement of a particle track, measured by the Processors is shown next. The correct functionality is successfully verified.**

Putting things together

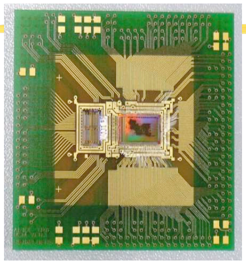

MCM = Preamp+TRAP

TRAcklet Processor (TRAP) - 21 x 10 bit ADCs, digital filters, 4 RISC CPUs, fast readout

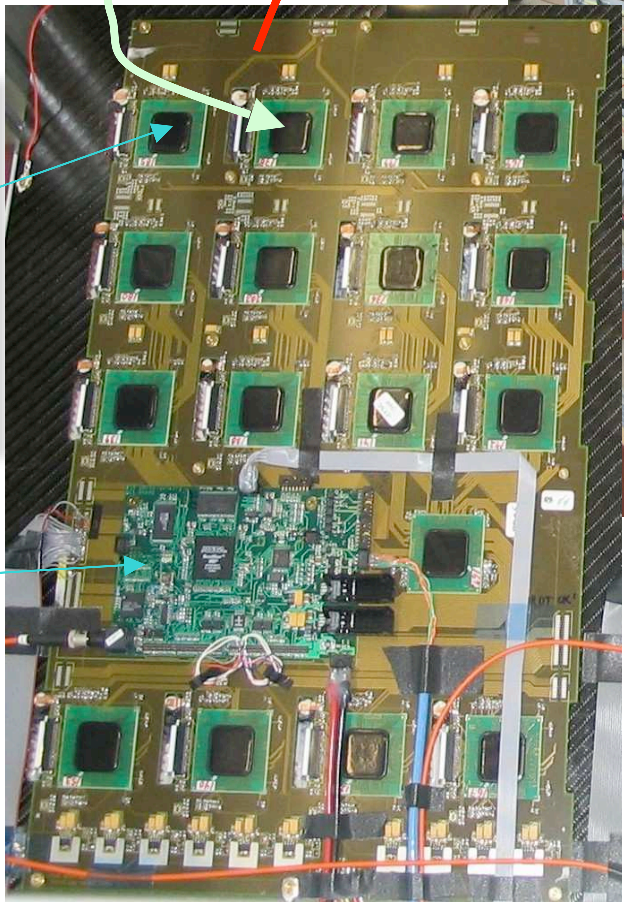

Configuration SBC

Trigger & clock distribution, ARM CPU+FPGA, embedded Linux, Ethernet, serial link to the CPUs

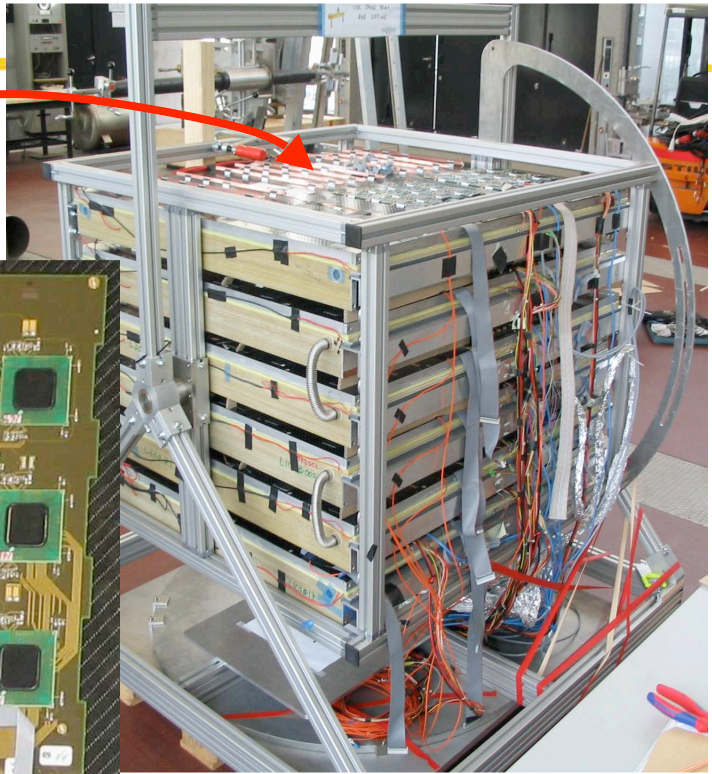

Detector Stack (one of 90)

Tests and Measurements

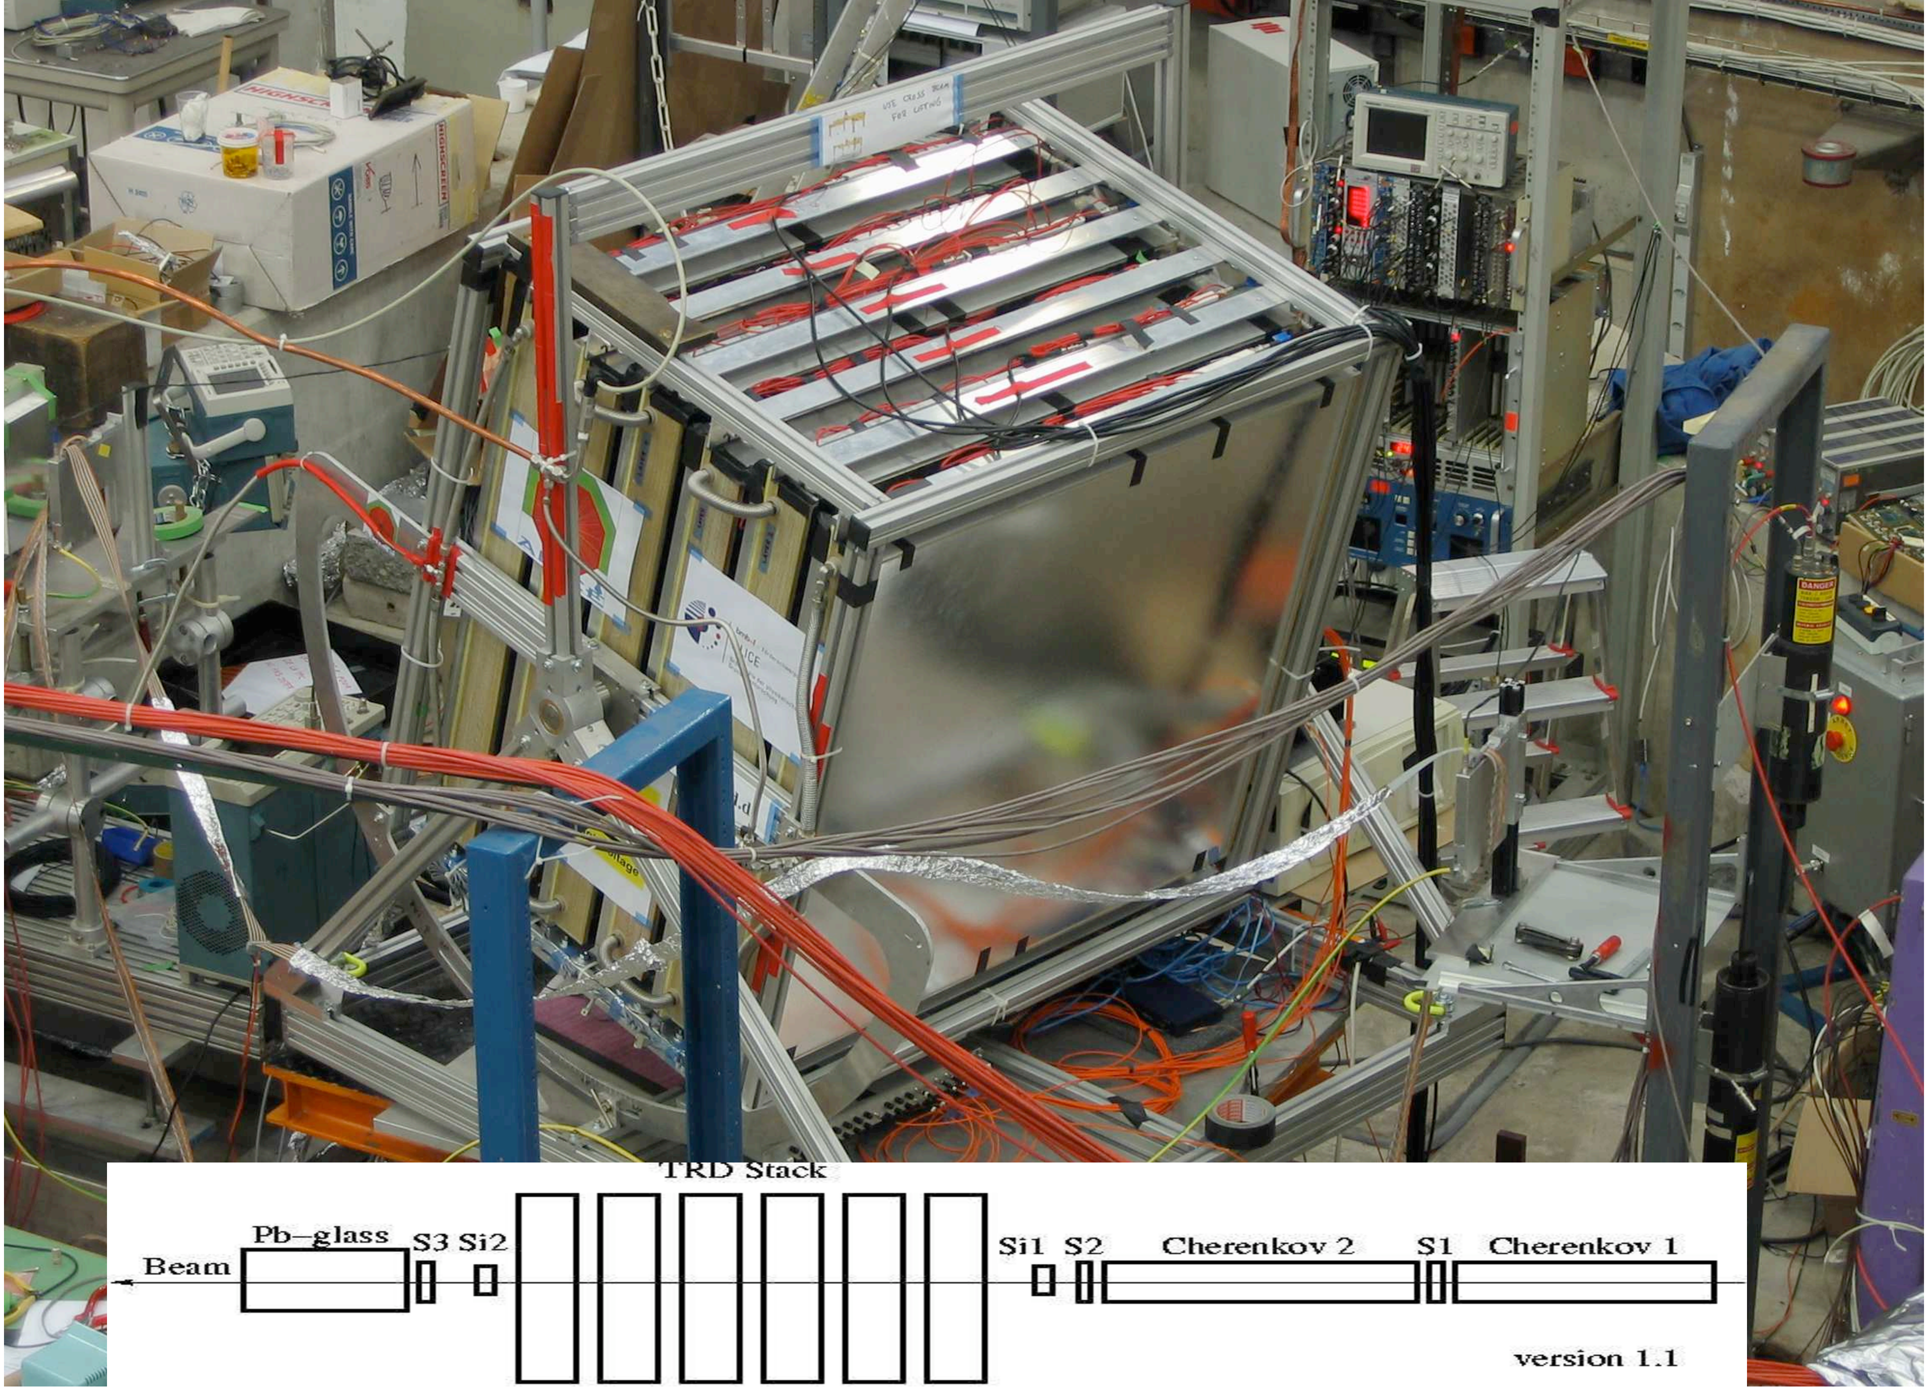

A TRD stack event

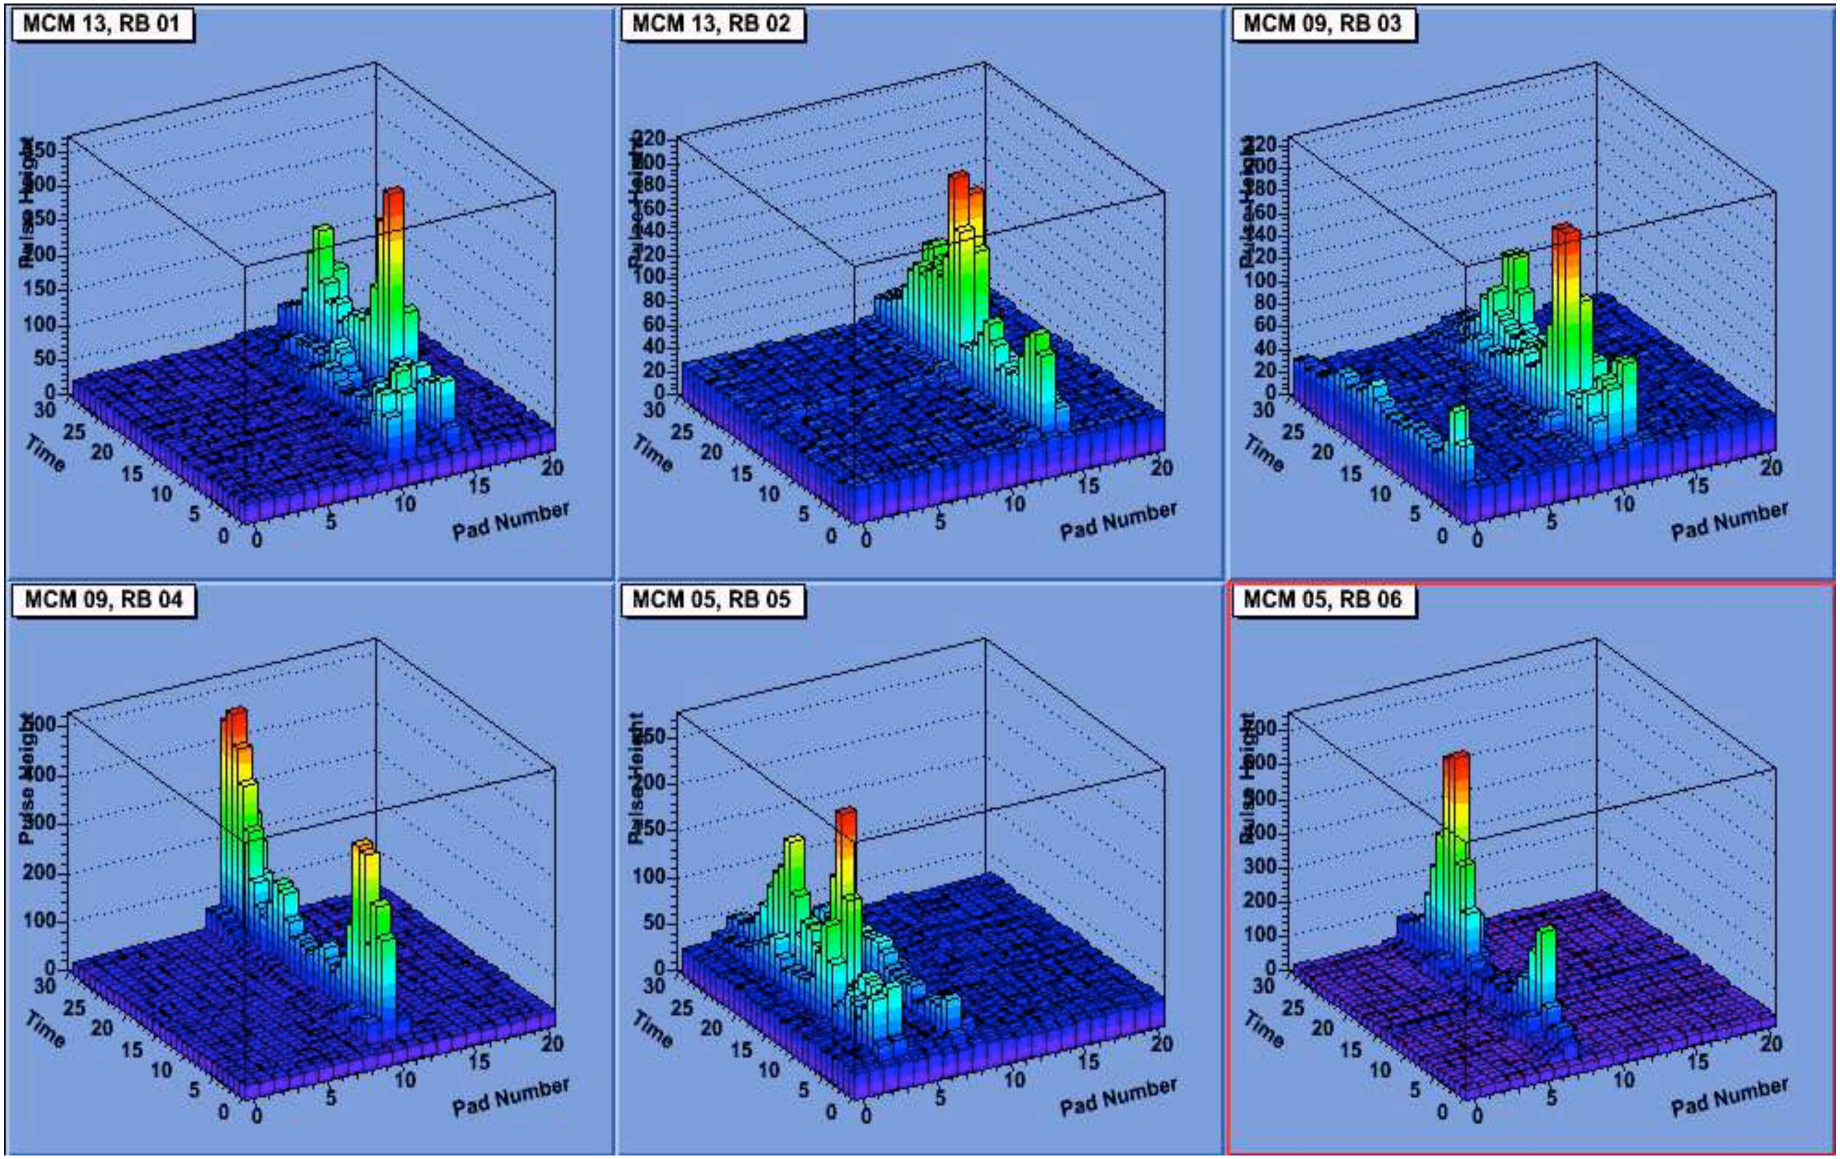

Summary, Conclusions

- The high processing and integration requirements result in the electronics being integrated with the detector
- The tight latency budget results in a MIMD architecture using multi port memories for communication and synchronization
- Iterative processing is performed by a preprocessor, projecting its results into a read-only part of each processors register file, avoiding load instructions
- All critical data structures and memories implement error correction to compensate for SEUs
- The processor with its complex periphery is fully functional in the system
- Next steps are the construction of the entire detector and the development of the control infrastructure

Thank you for your Attention

Questions?