LHC experimental data: From today's Data Challenges to the promise of tomorrow

B. Panzer – CERN/IT,
F. Rademakers – CERN/EP,
P. Vande Vyvre - CERN/EP

Academic Training CERN

Outline

- Day 1 (Pierre VANDE VYVRE)
  - Outline, main concepts
  - Requirements of LHC experiments
  - Data Challenges
- Day 2 (Bernd PANZER)
  - Computing infrastructure
  - Technology trends
- Day 3 (Pierre VANDE VYVRE)
  - Trigger and Data acquisition
- Day 4 (Fons RADEMAKERS)
  - Simulation, Reconstruction and analysis
- Day 5 (Bernd PANZER)
  - Computing Data challenges
  - Physics Data Challenges
  - Evolution

Outline

- Day 1 (Pierre VANDE VYVRE)
  - Outline, main concepts
  - Requirements of LHC experiments
  - Data Challenges
- Day 2 (Bernd PANZER)
  - Computing infrastructure
  - Technology trends
- Day 3 (Pierre VANDE VYVRE)
  - Trigger and Data acquisition
- Day 4 (Fons RADEMAKERS)
  - Simulation, Reconstruction and analysis
- Day 5 (Bernd PANZER)
  - Computing Data challenges
  - Physics Data Challenges
  - Evolution

Day 1

- Outline of this series
- Main concepts

- Requirements of the LHC experiments
  - Trigger @ LHC
  - Data acquisition @ LHC
  - Data storage @ LHC

- Data Challenges
  - Evolutions of online and offline computing fabrics
  - Motivations of data challenges

Day 1

- Outline of this series
- Main concepts

- Requirements of the LHC experiments
  - Trigger @ LHC
  - Data acquisition @ LHC
  - Data storage @ LHC

- Data Challenges
  - Evolutions of online and offline computing fabrics
  - Motivations of data challenges

Overall view of the LHC experiments.

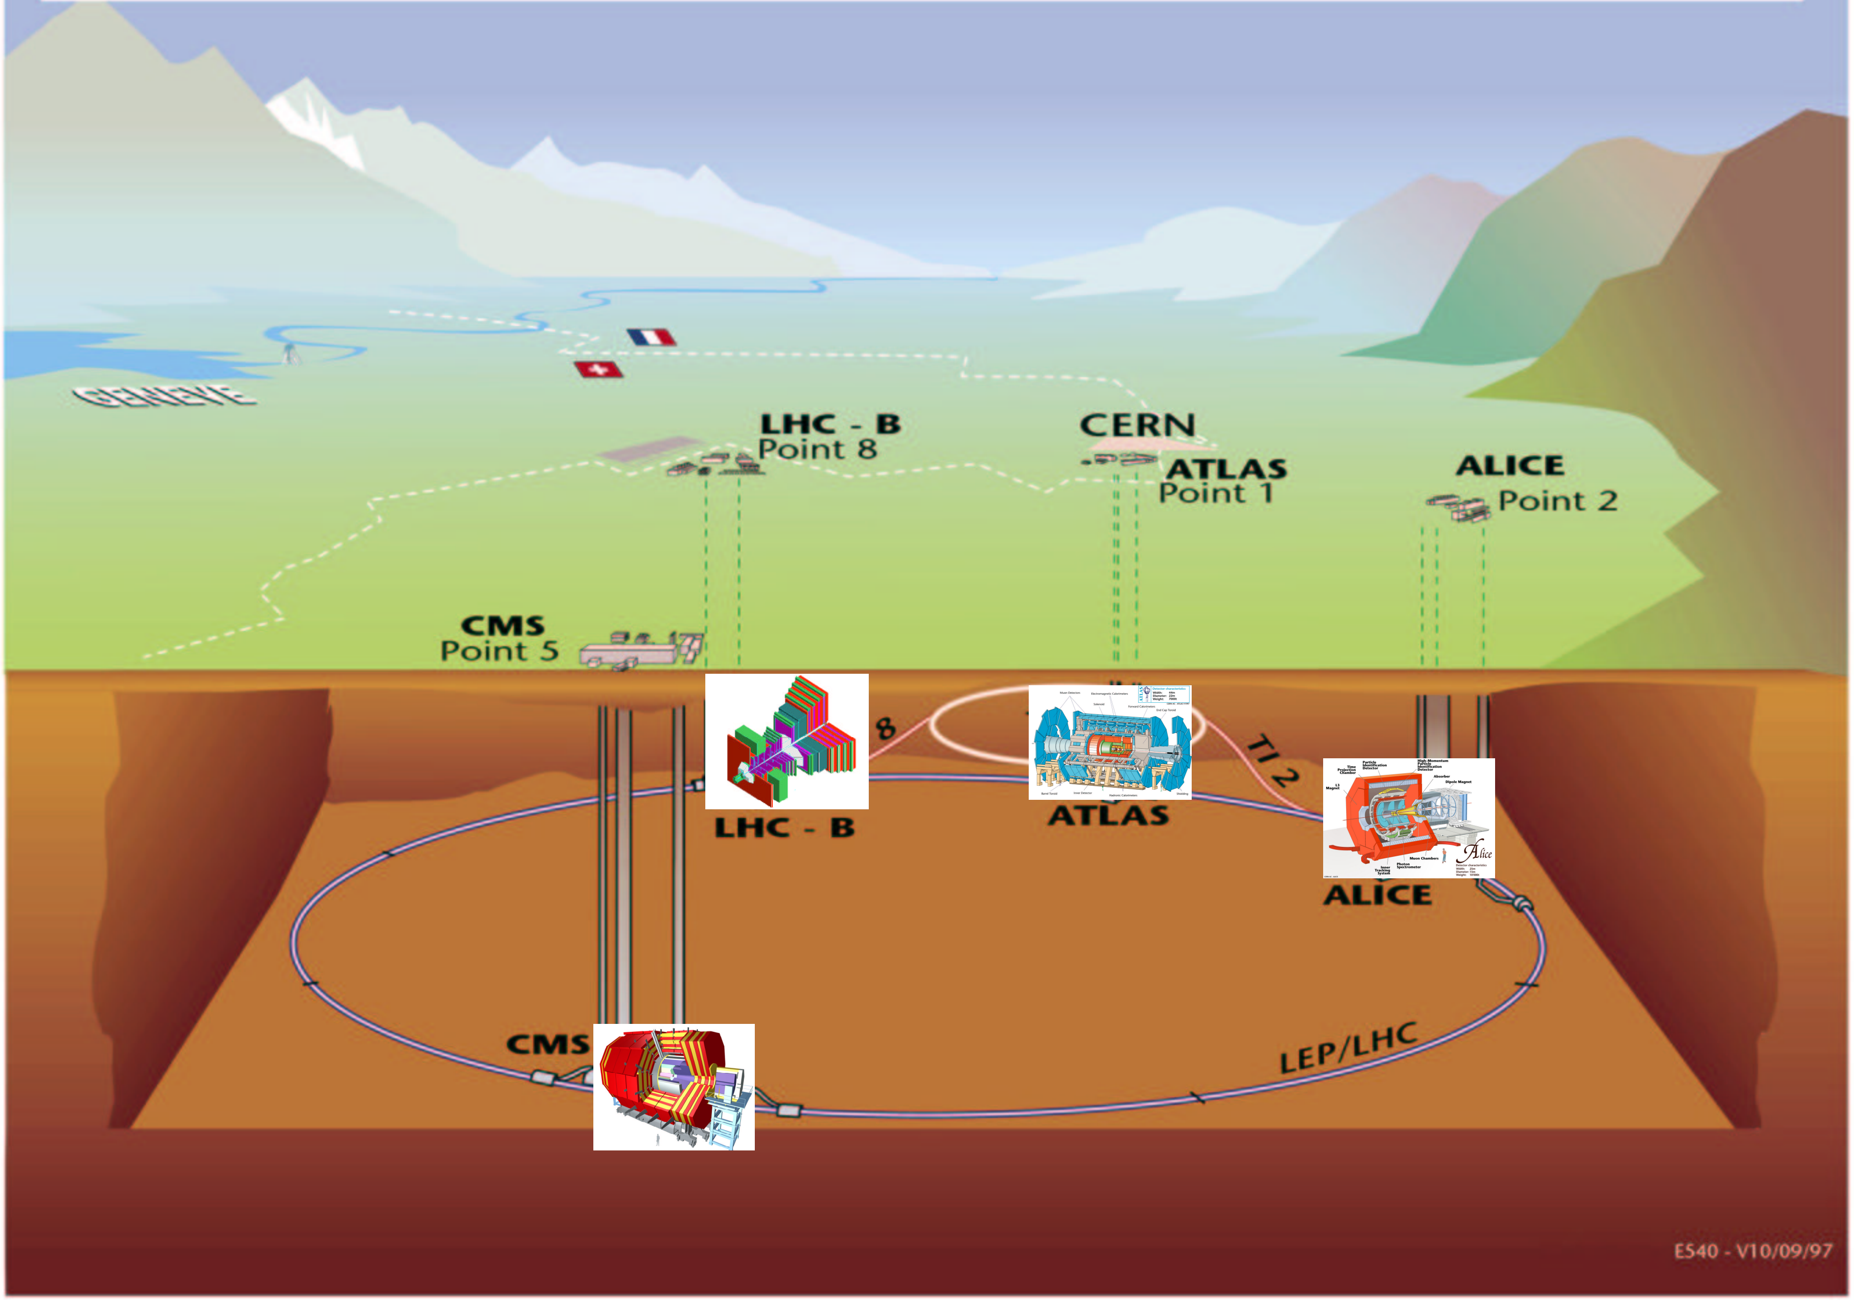

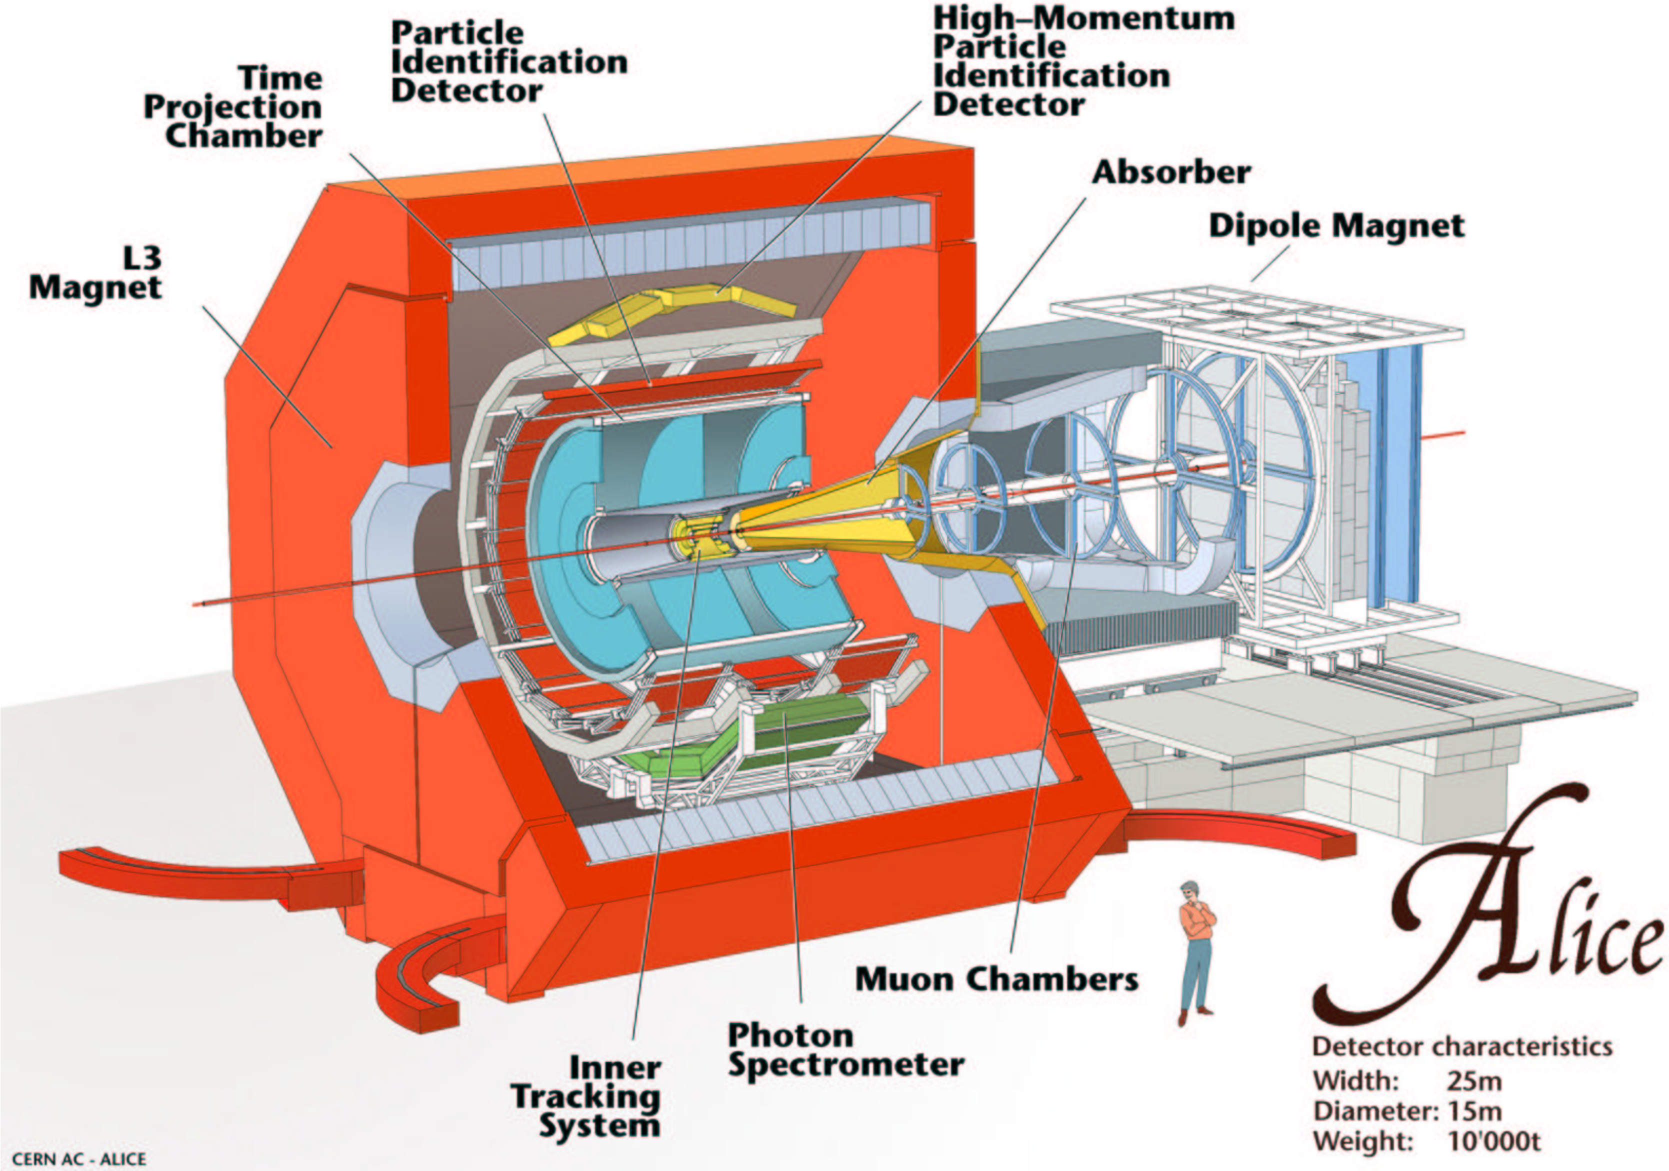
Time Projection Chamber
Particle Identification Detector
High-Momentum Particle Identification Detector
Absorber
Dipole Magnet
L3 Magnet
Muon Chambers
Photon Spectrometer
Inner Tracking System
Alice
Detector characteristics
Width: 25m
Diameter: 15m
Weight: 10'000t
CERN AC - ALICE

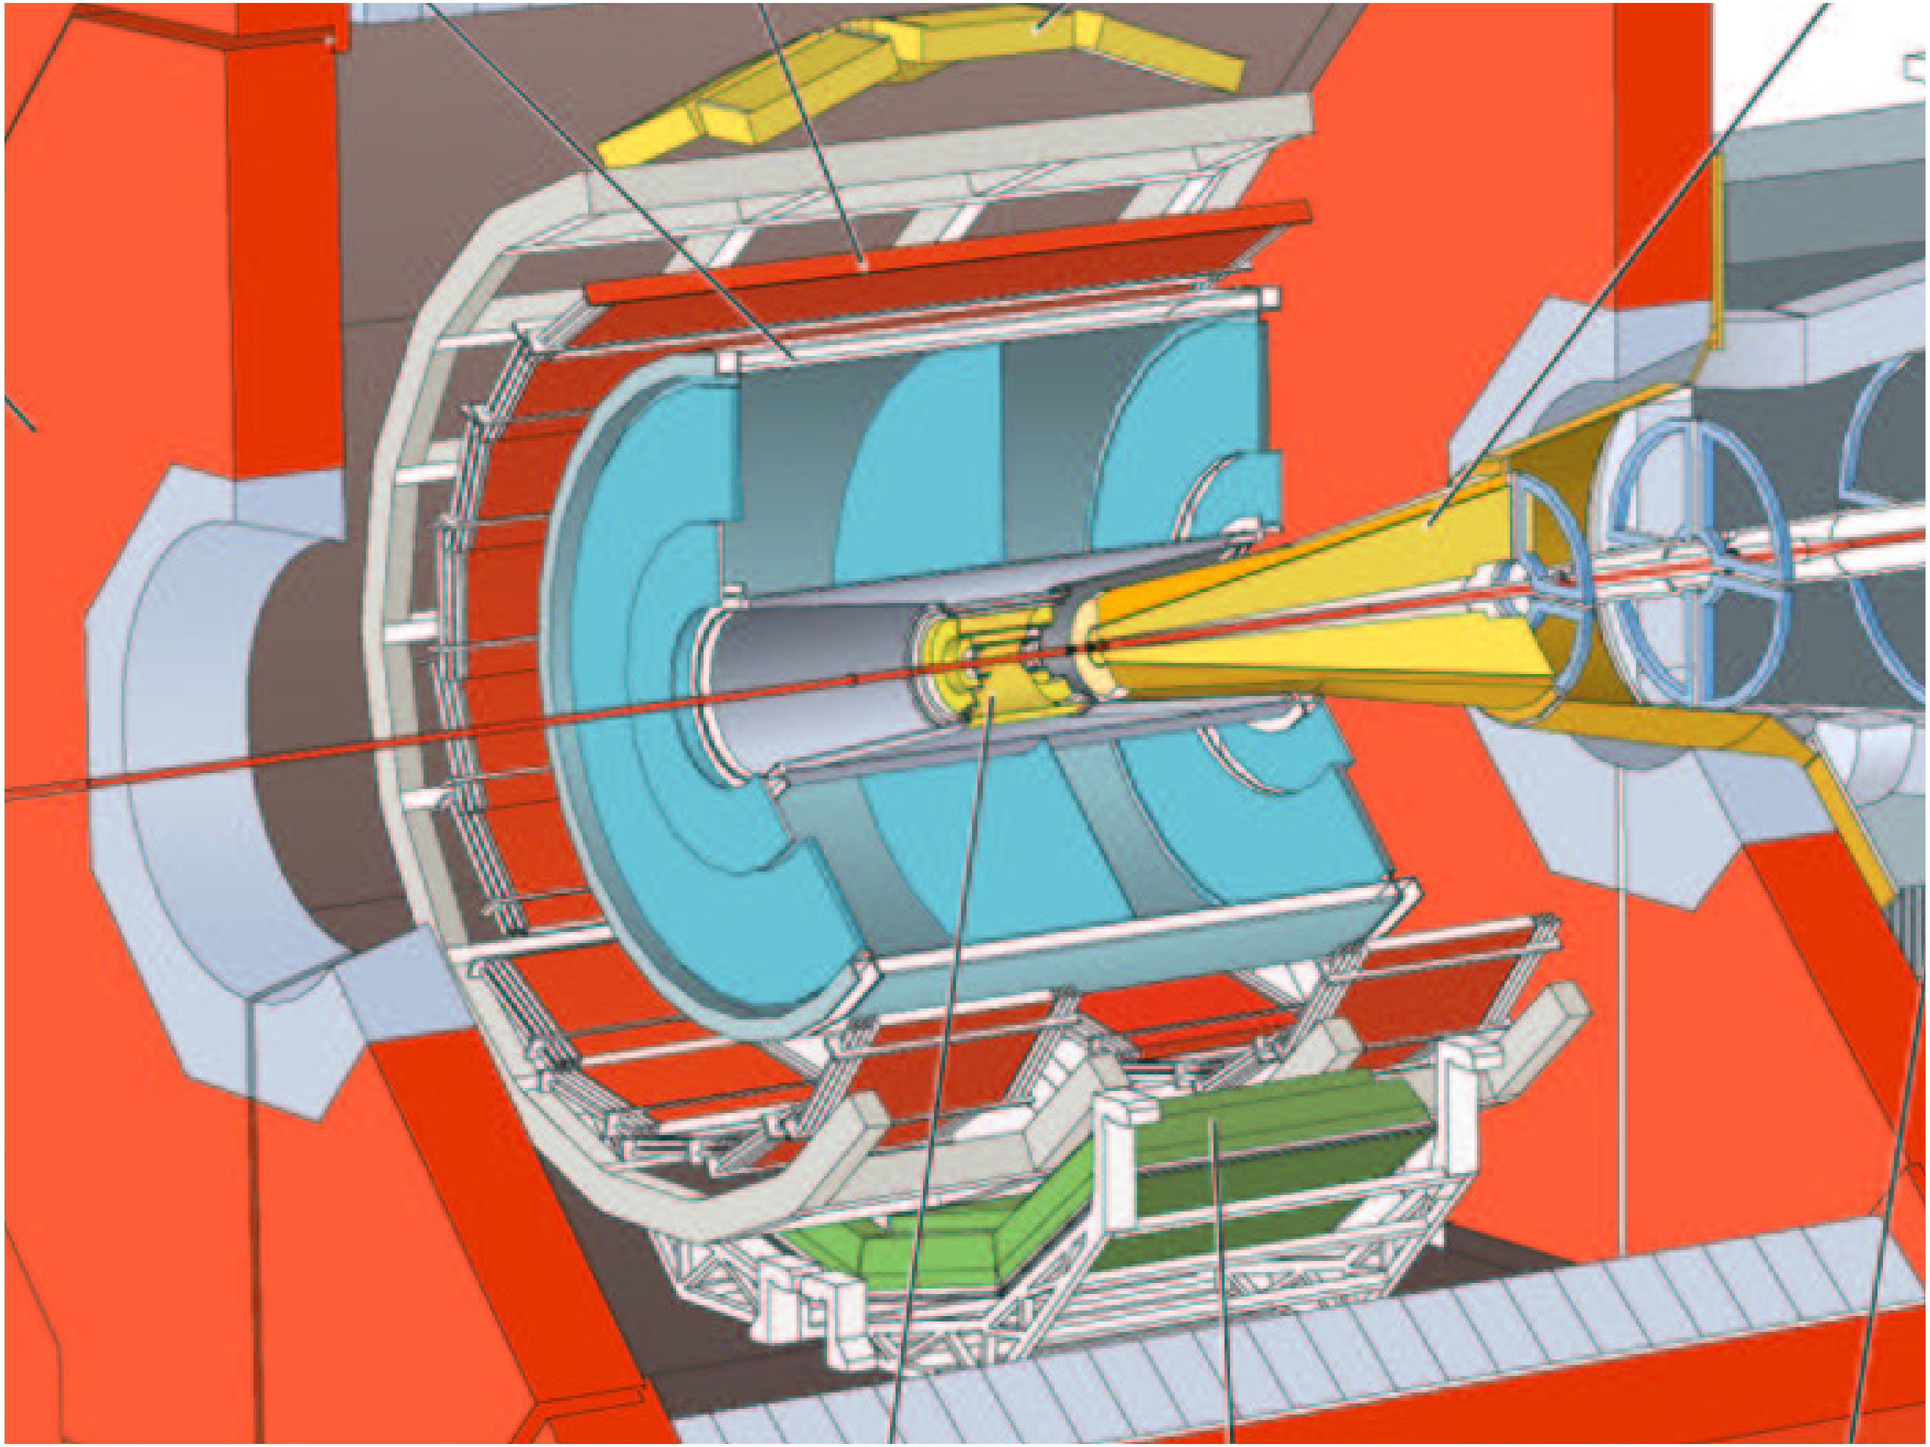

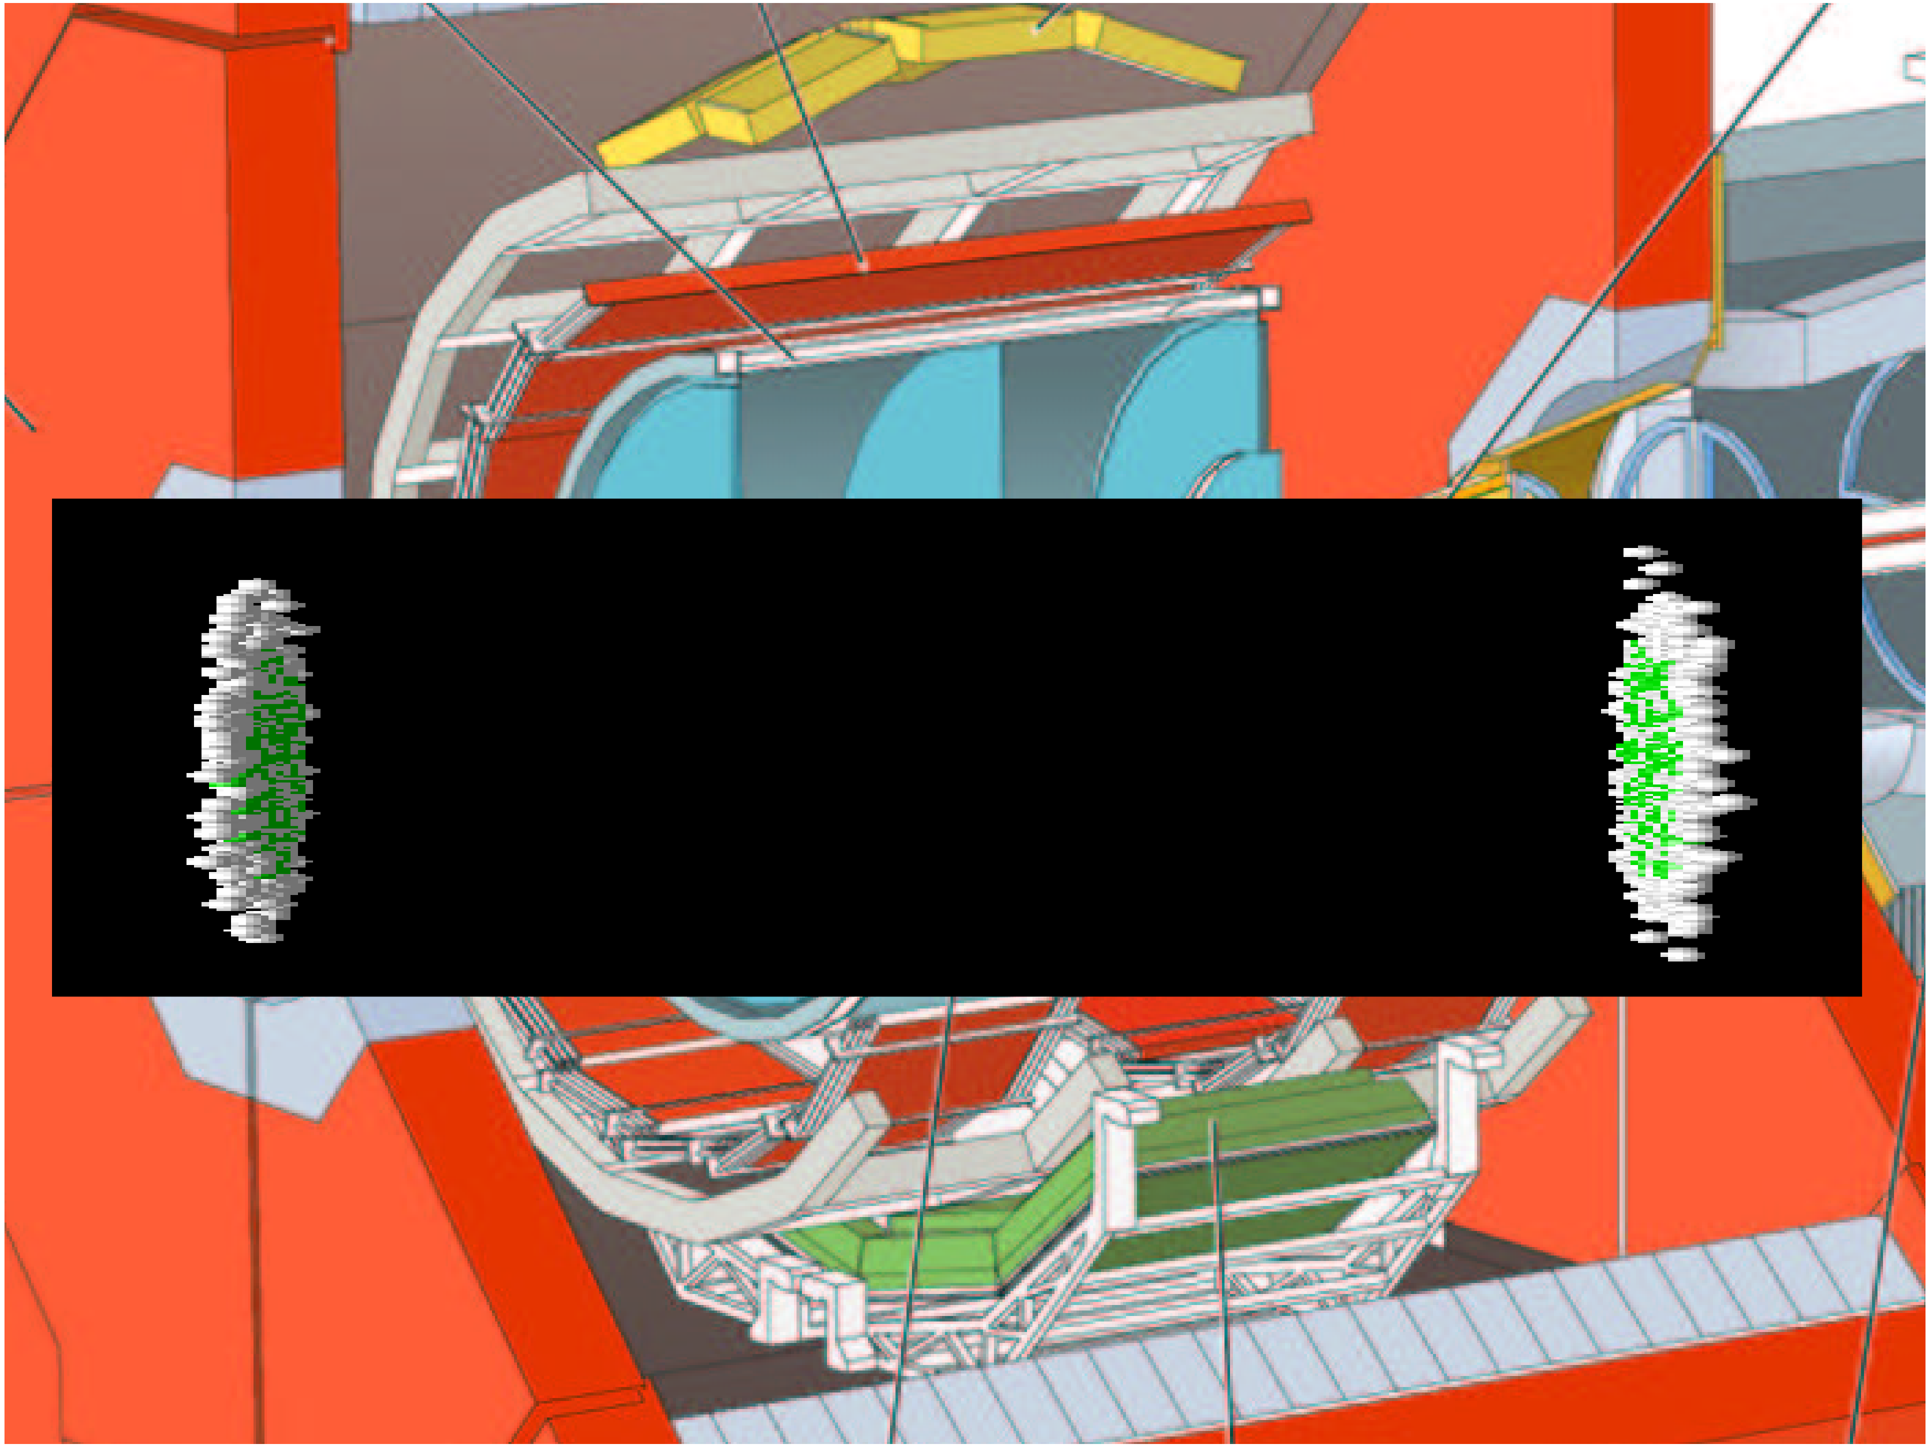

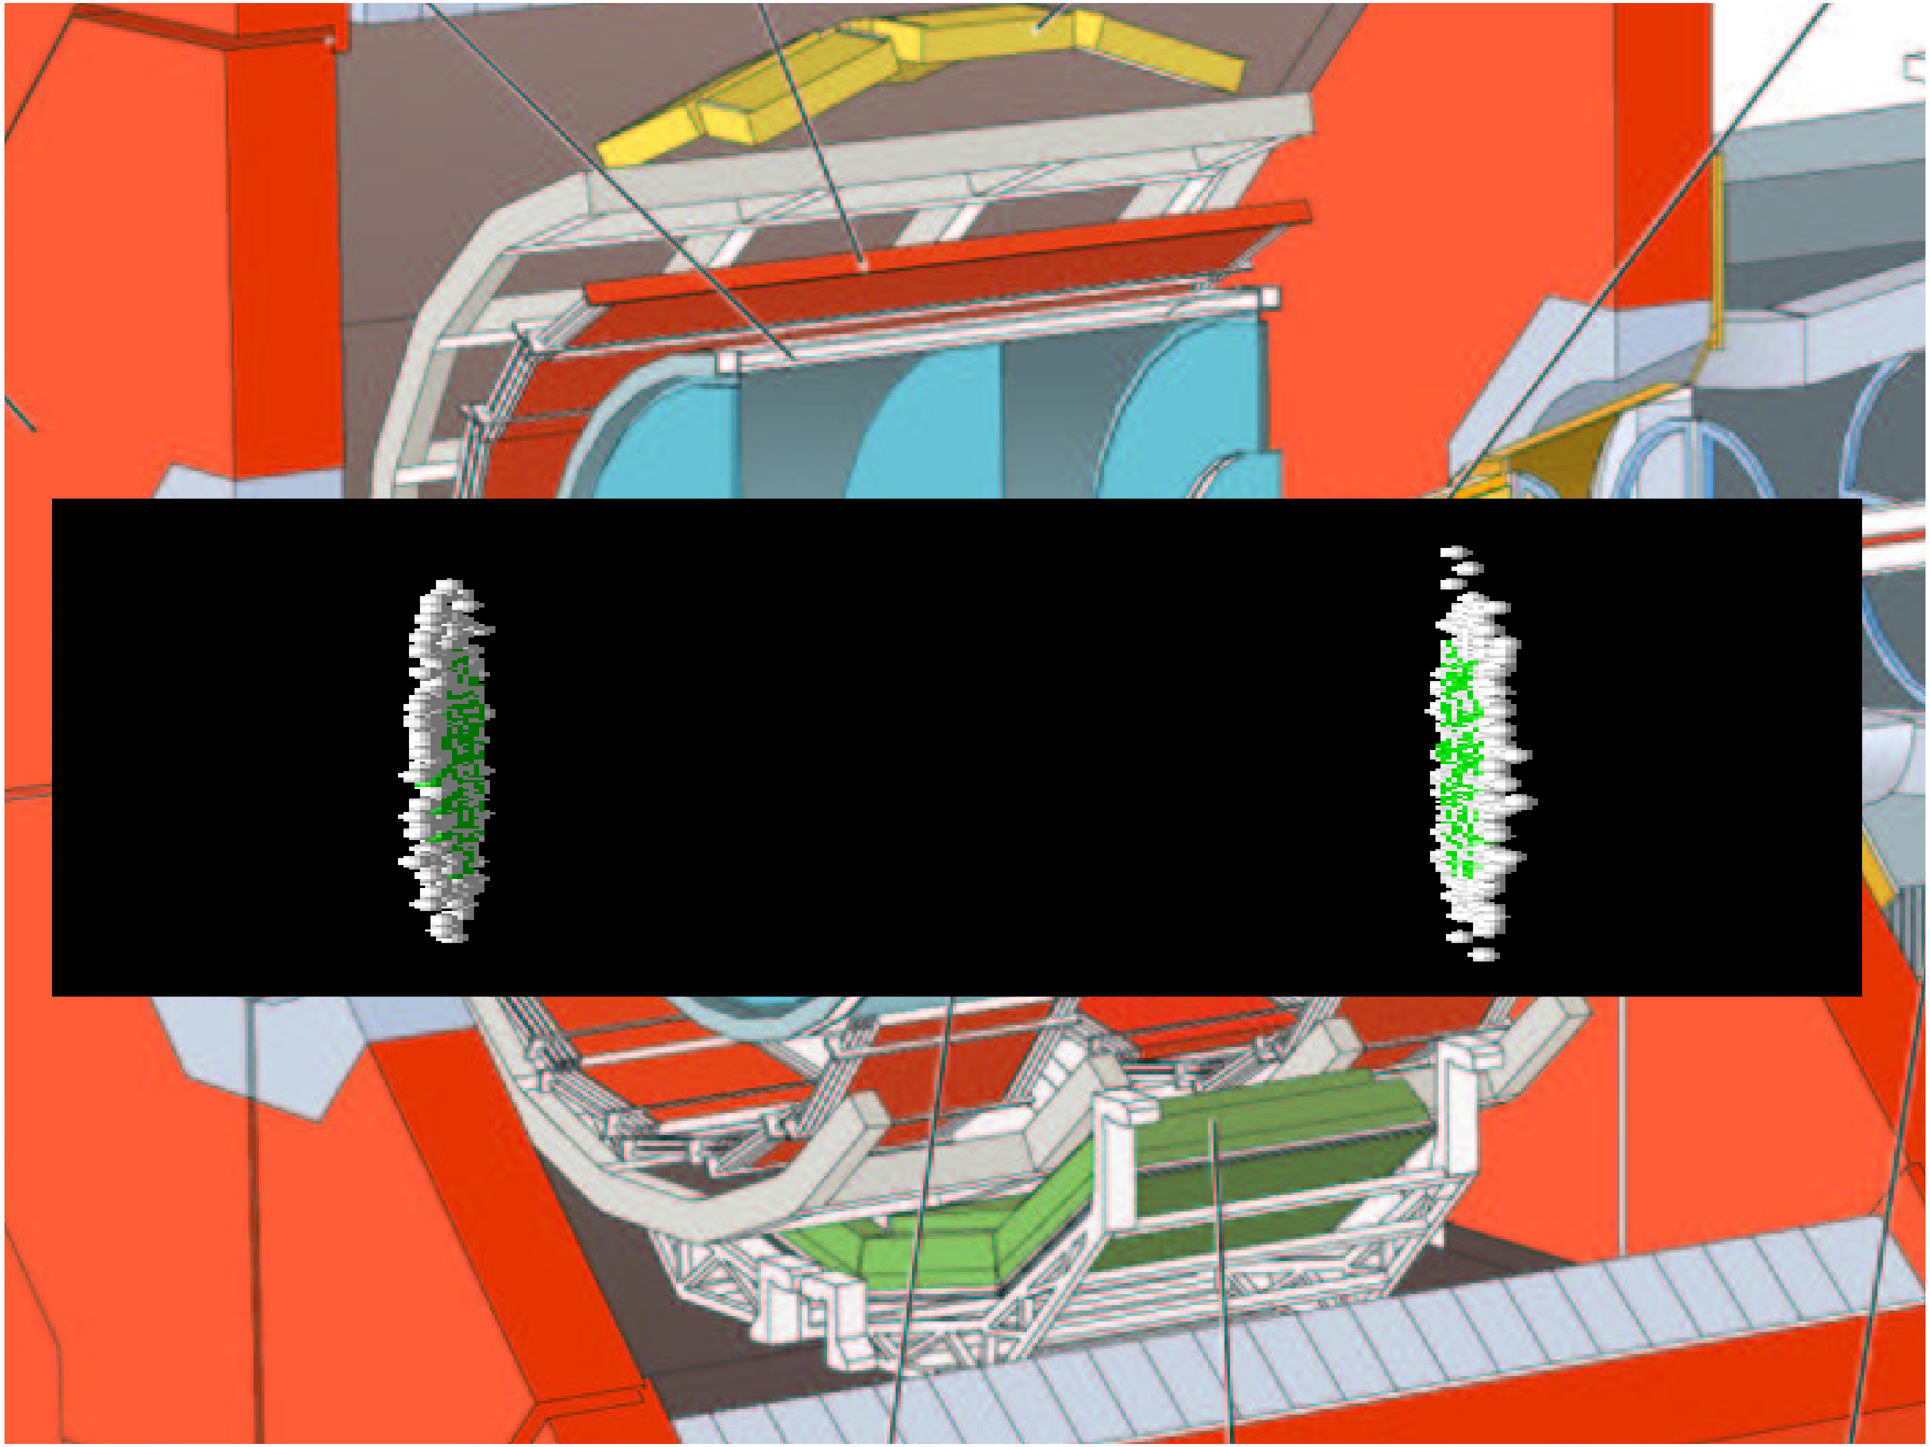

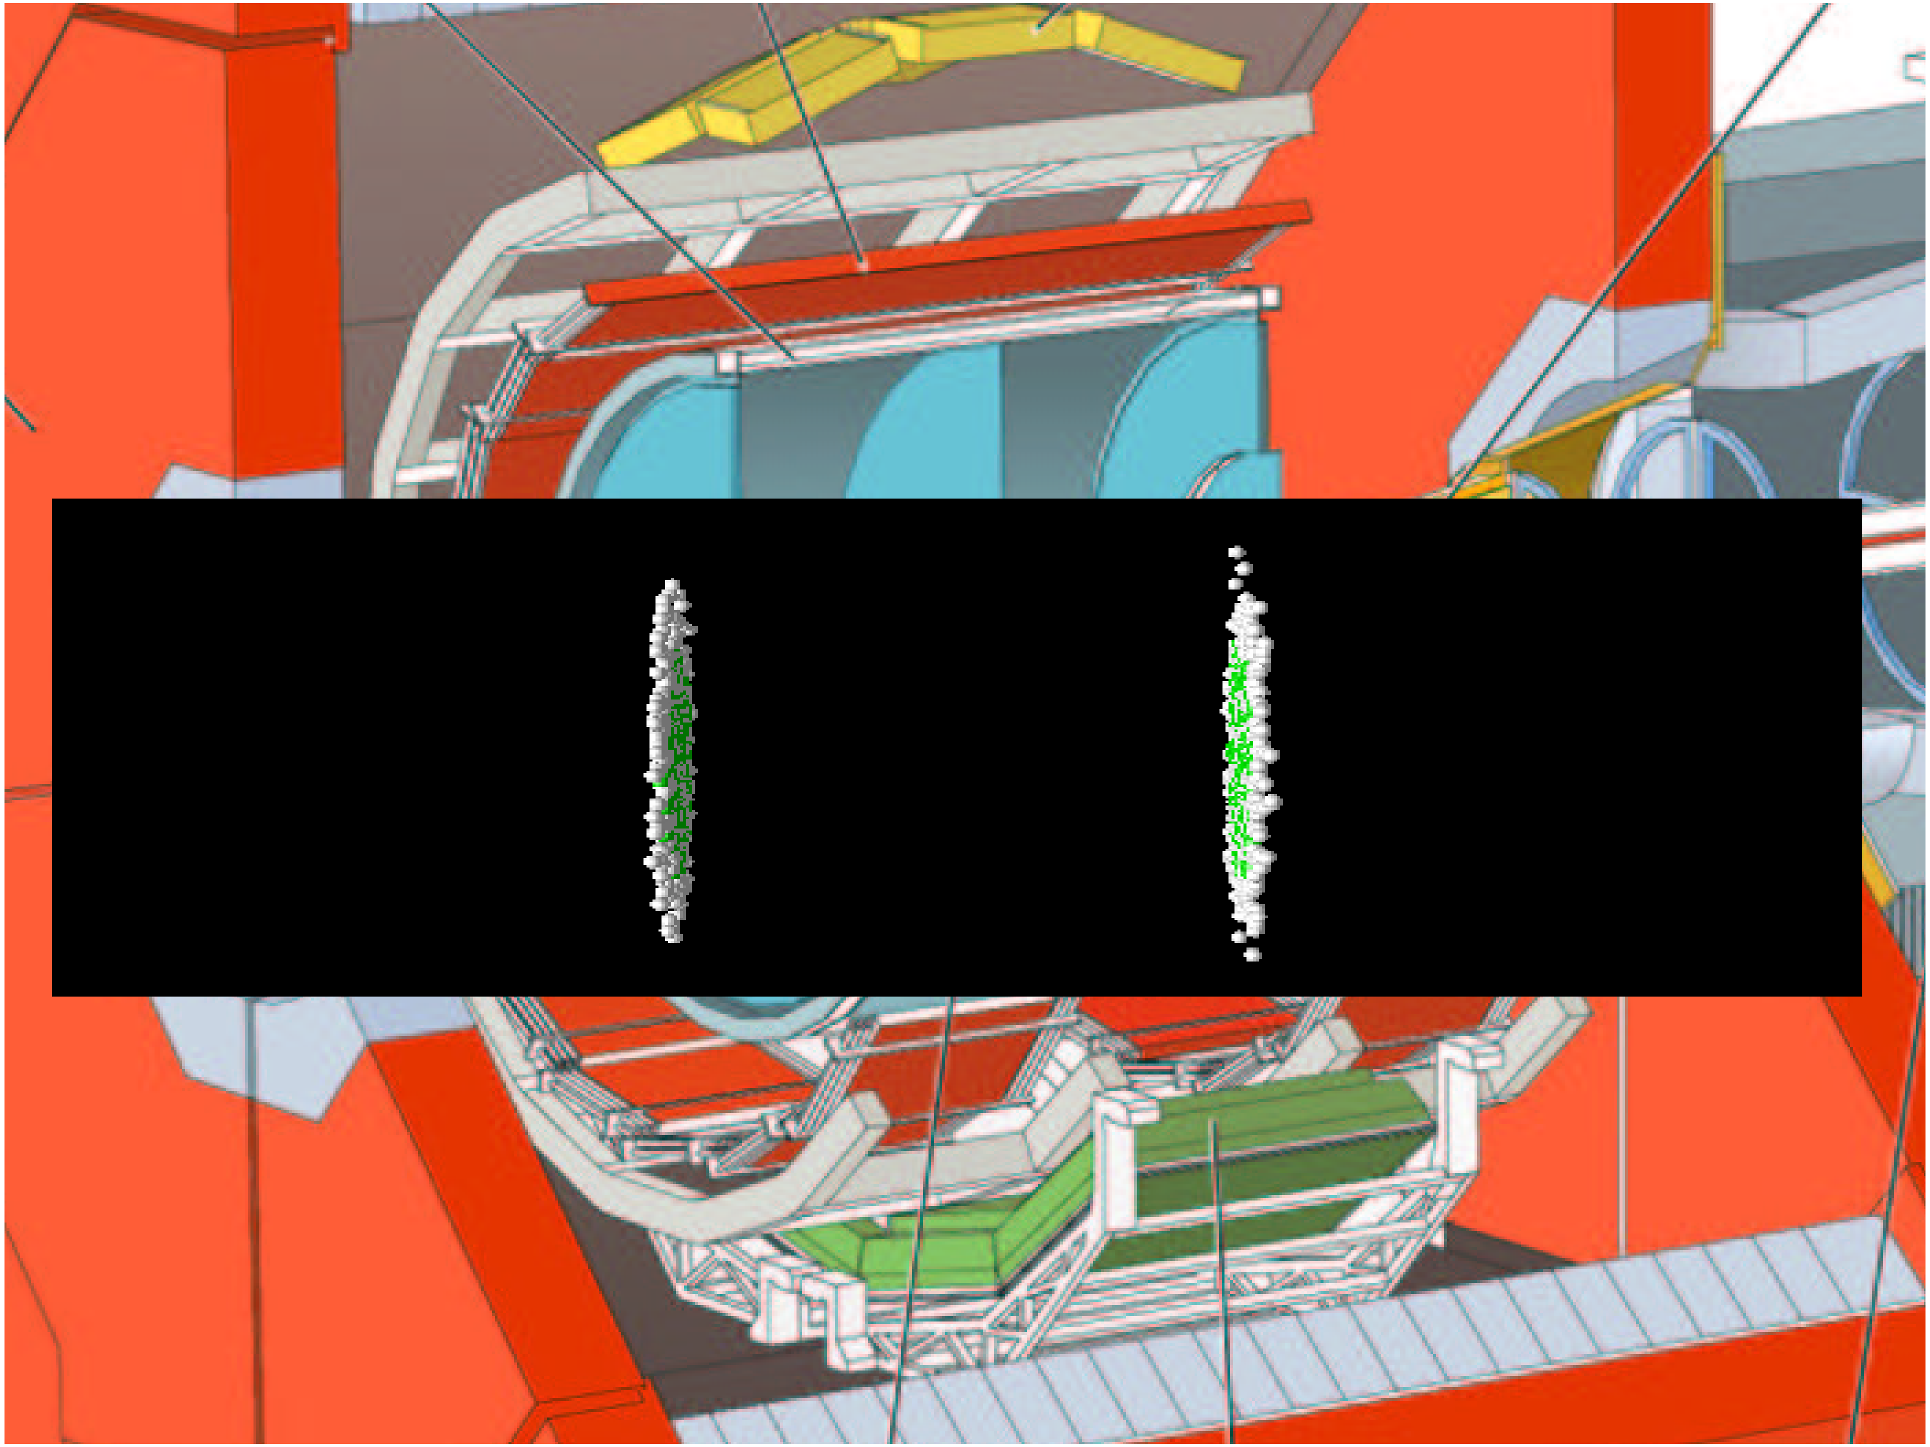

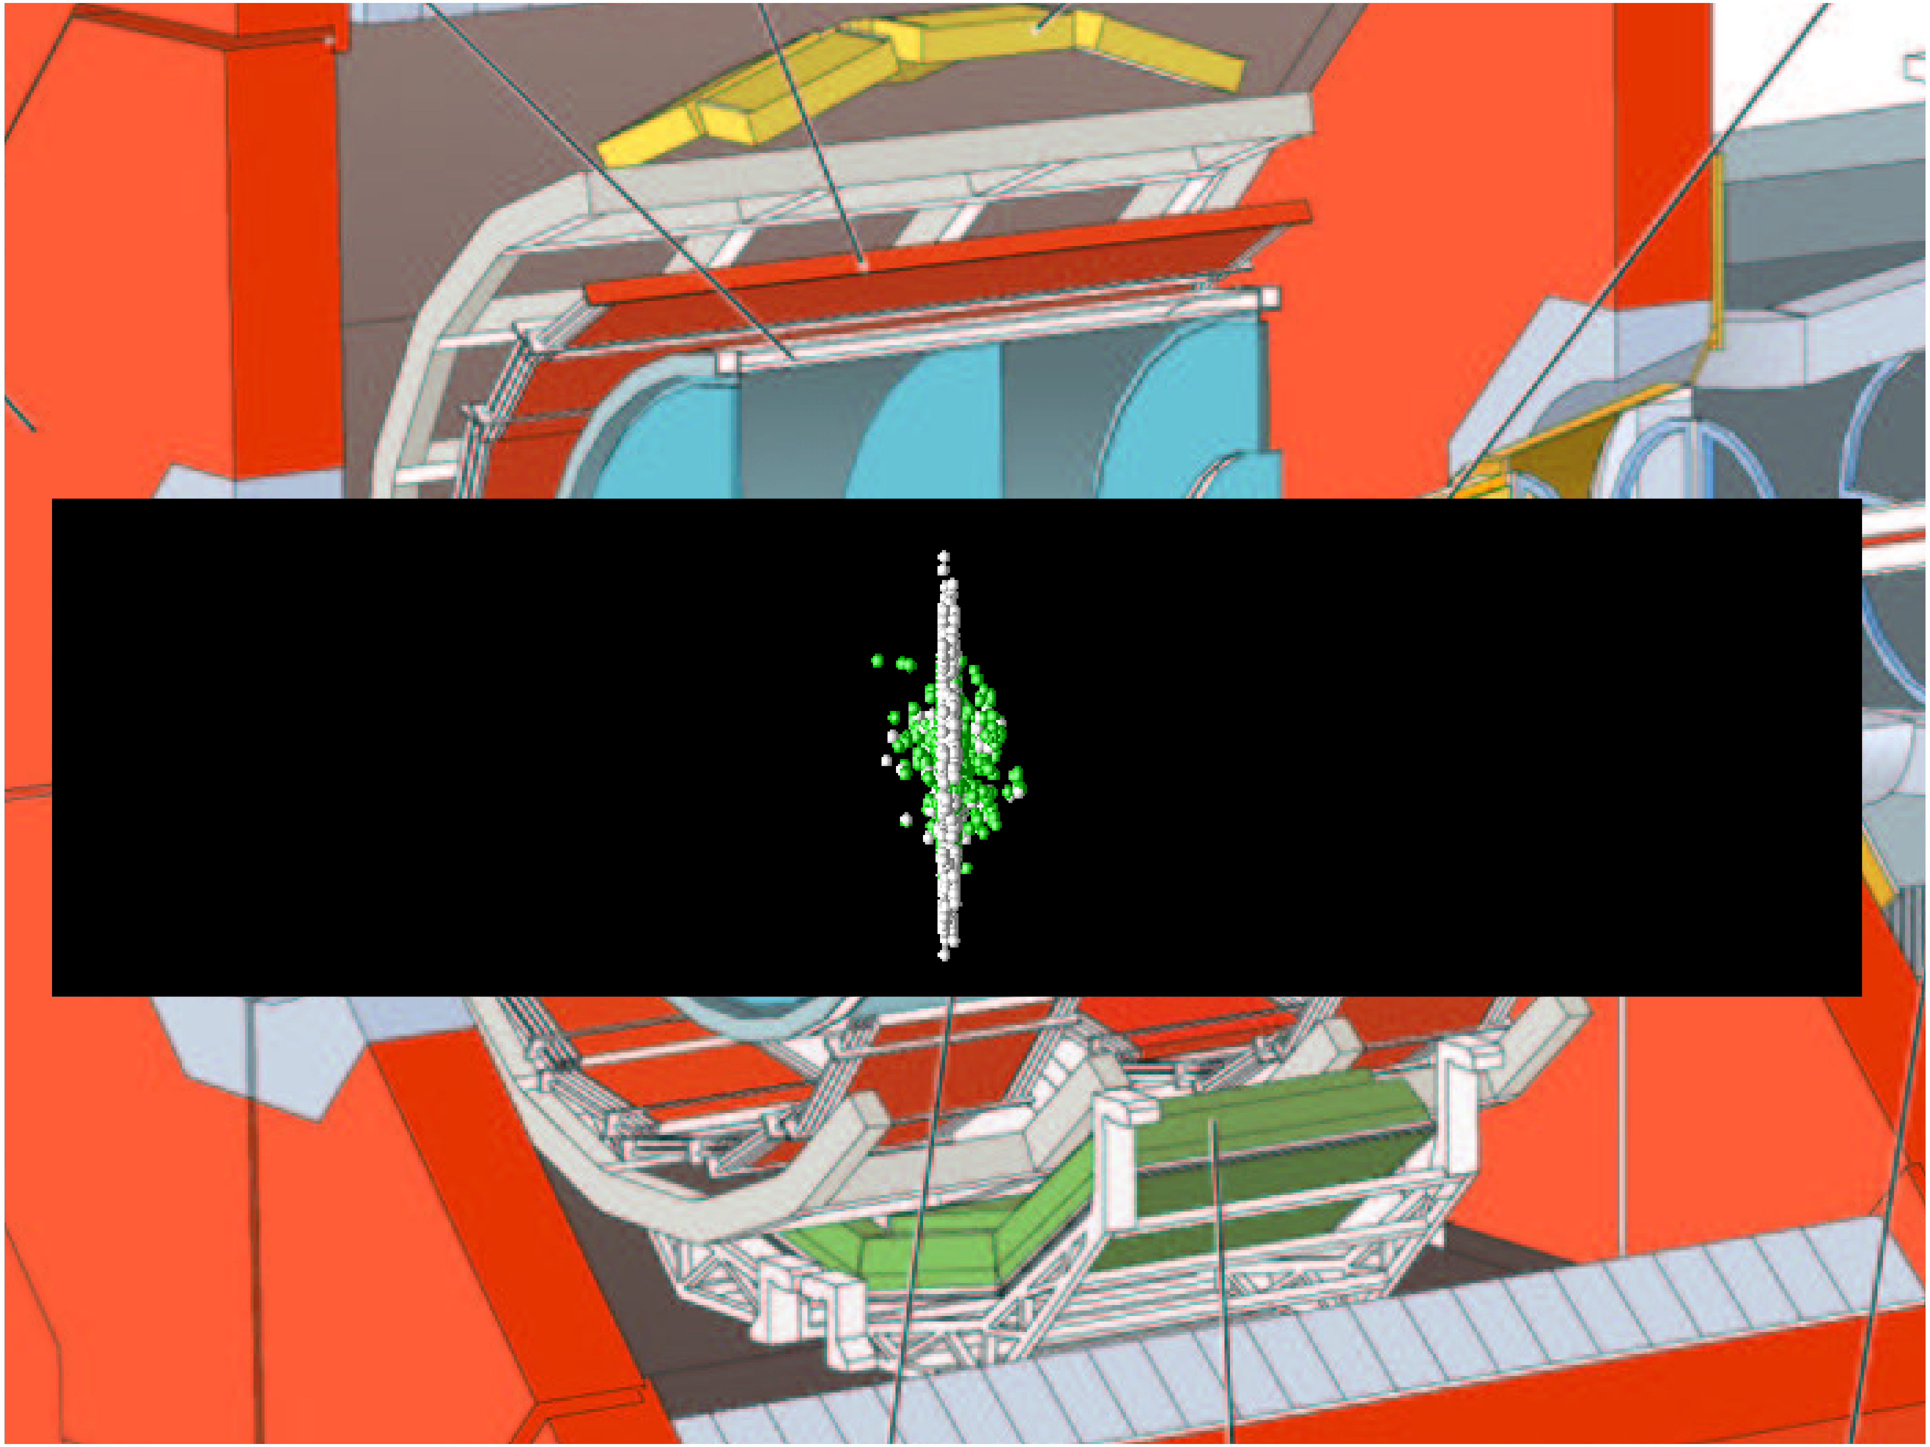

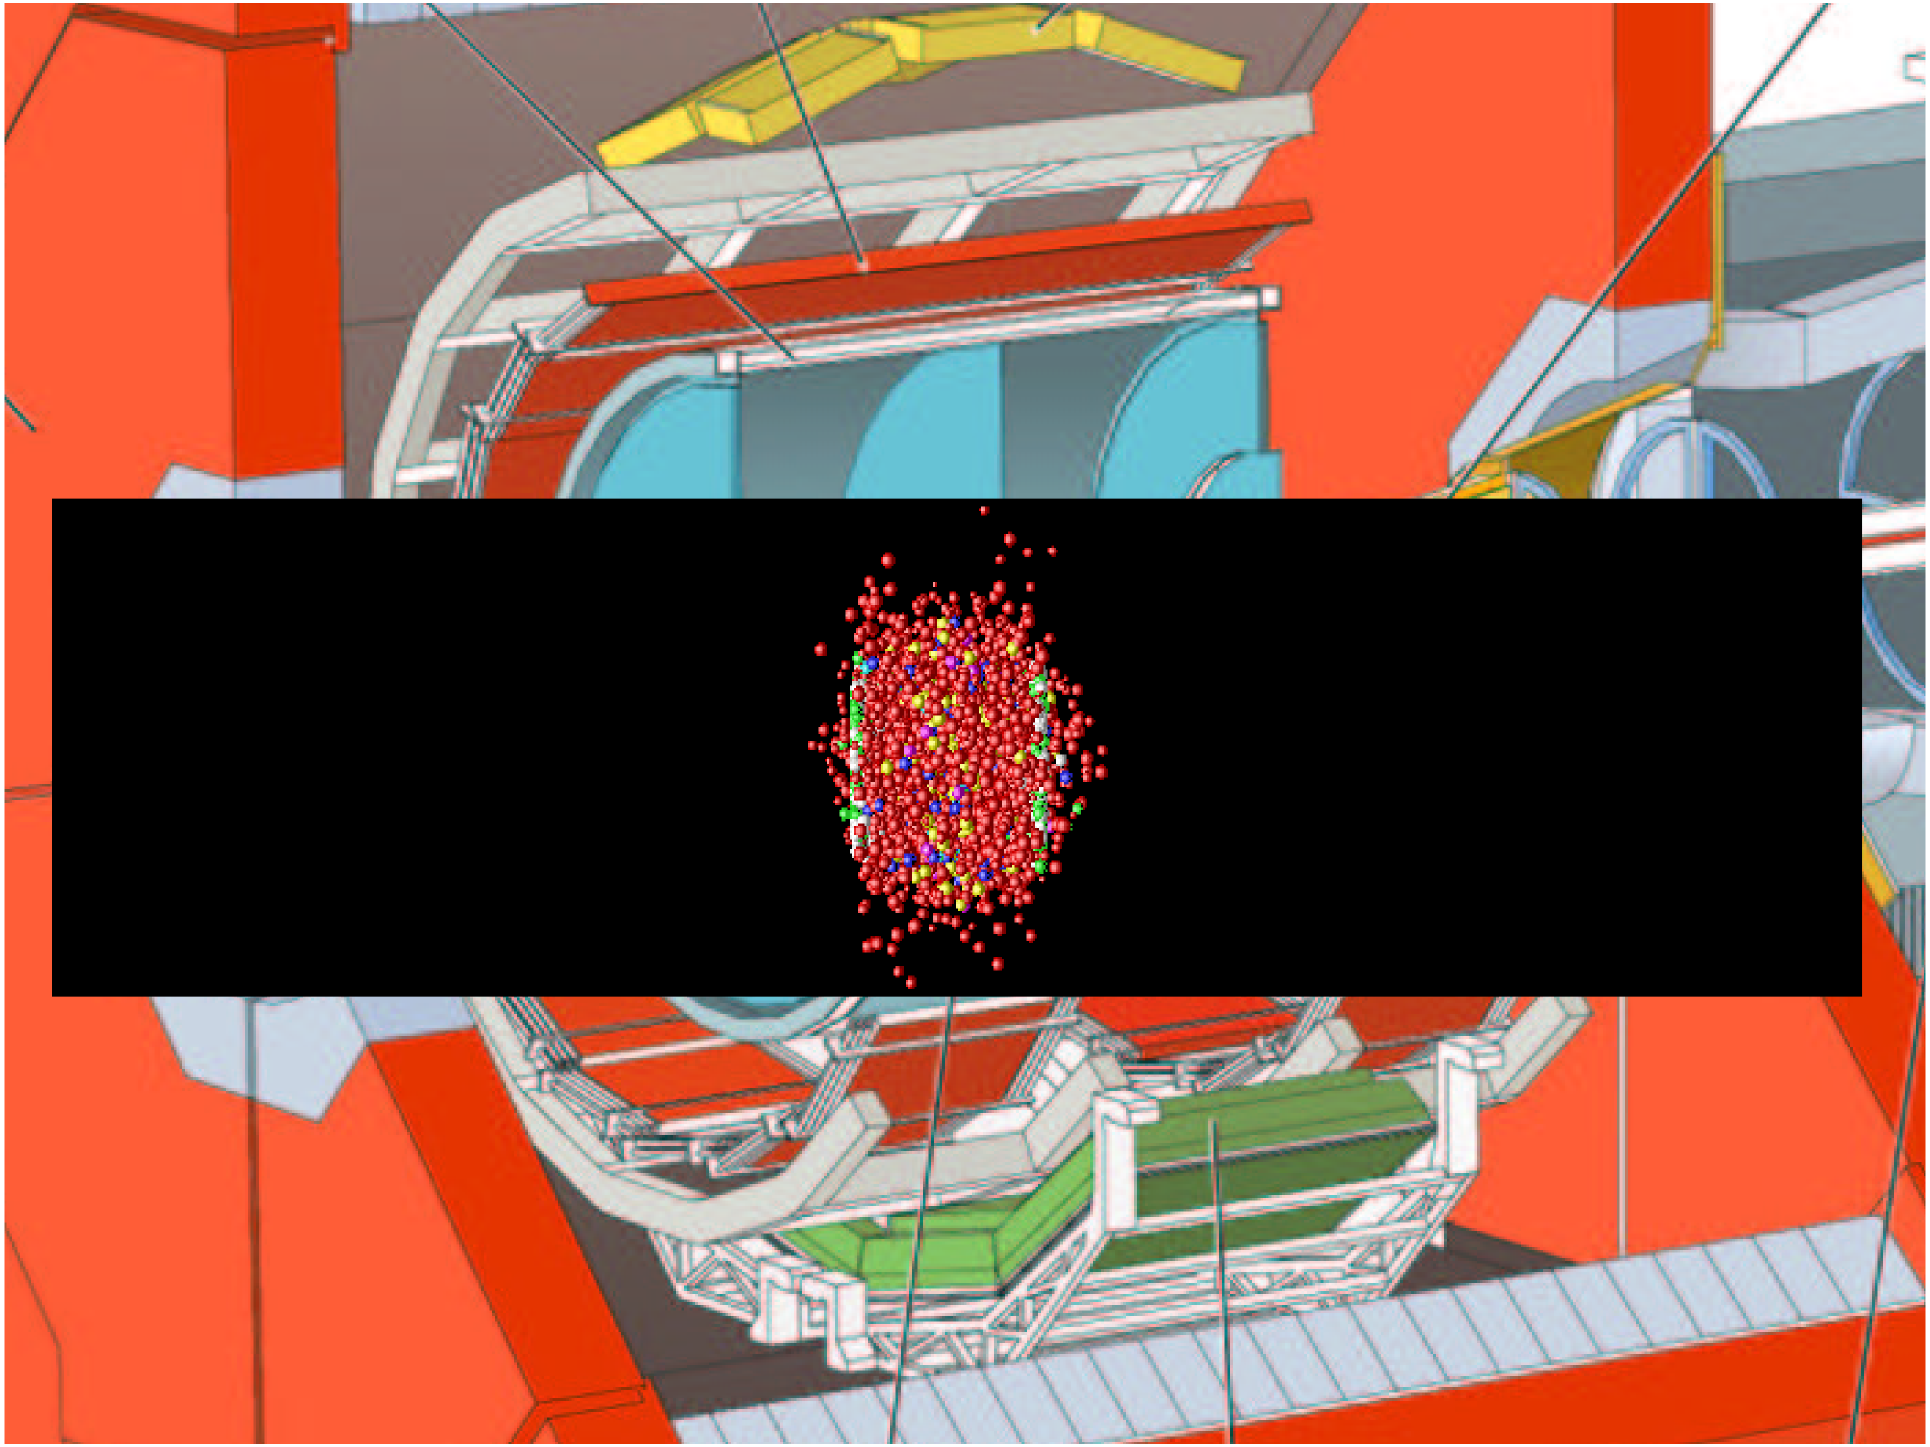

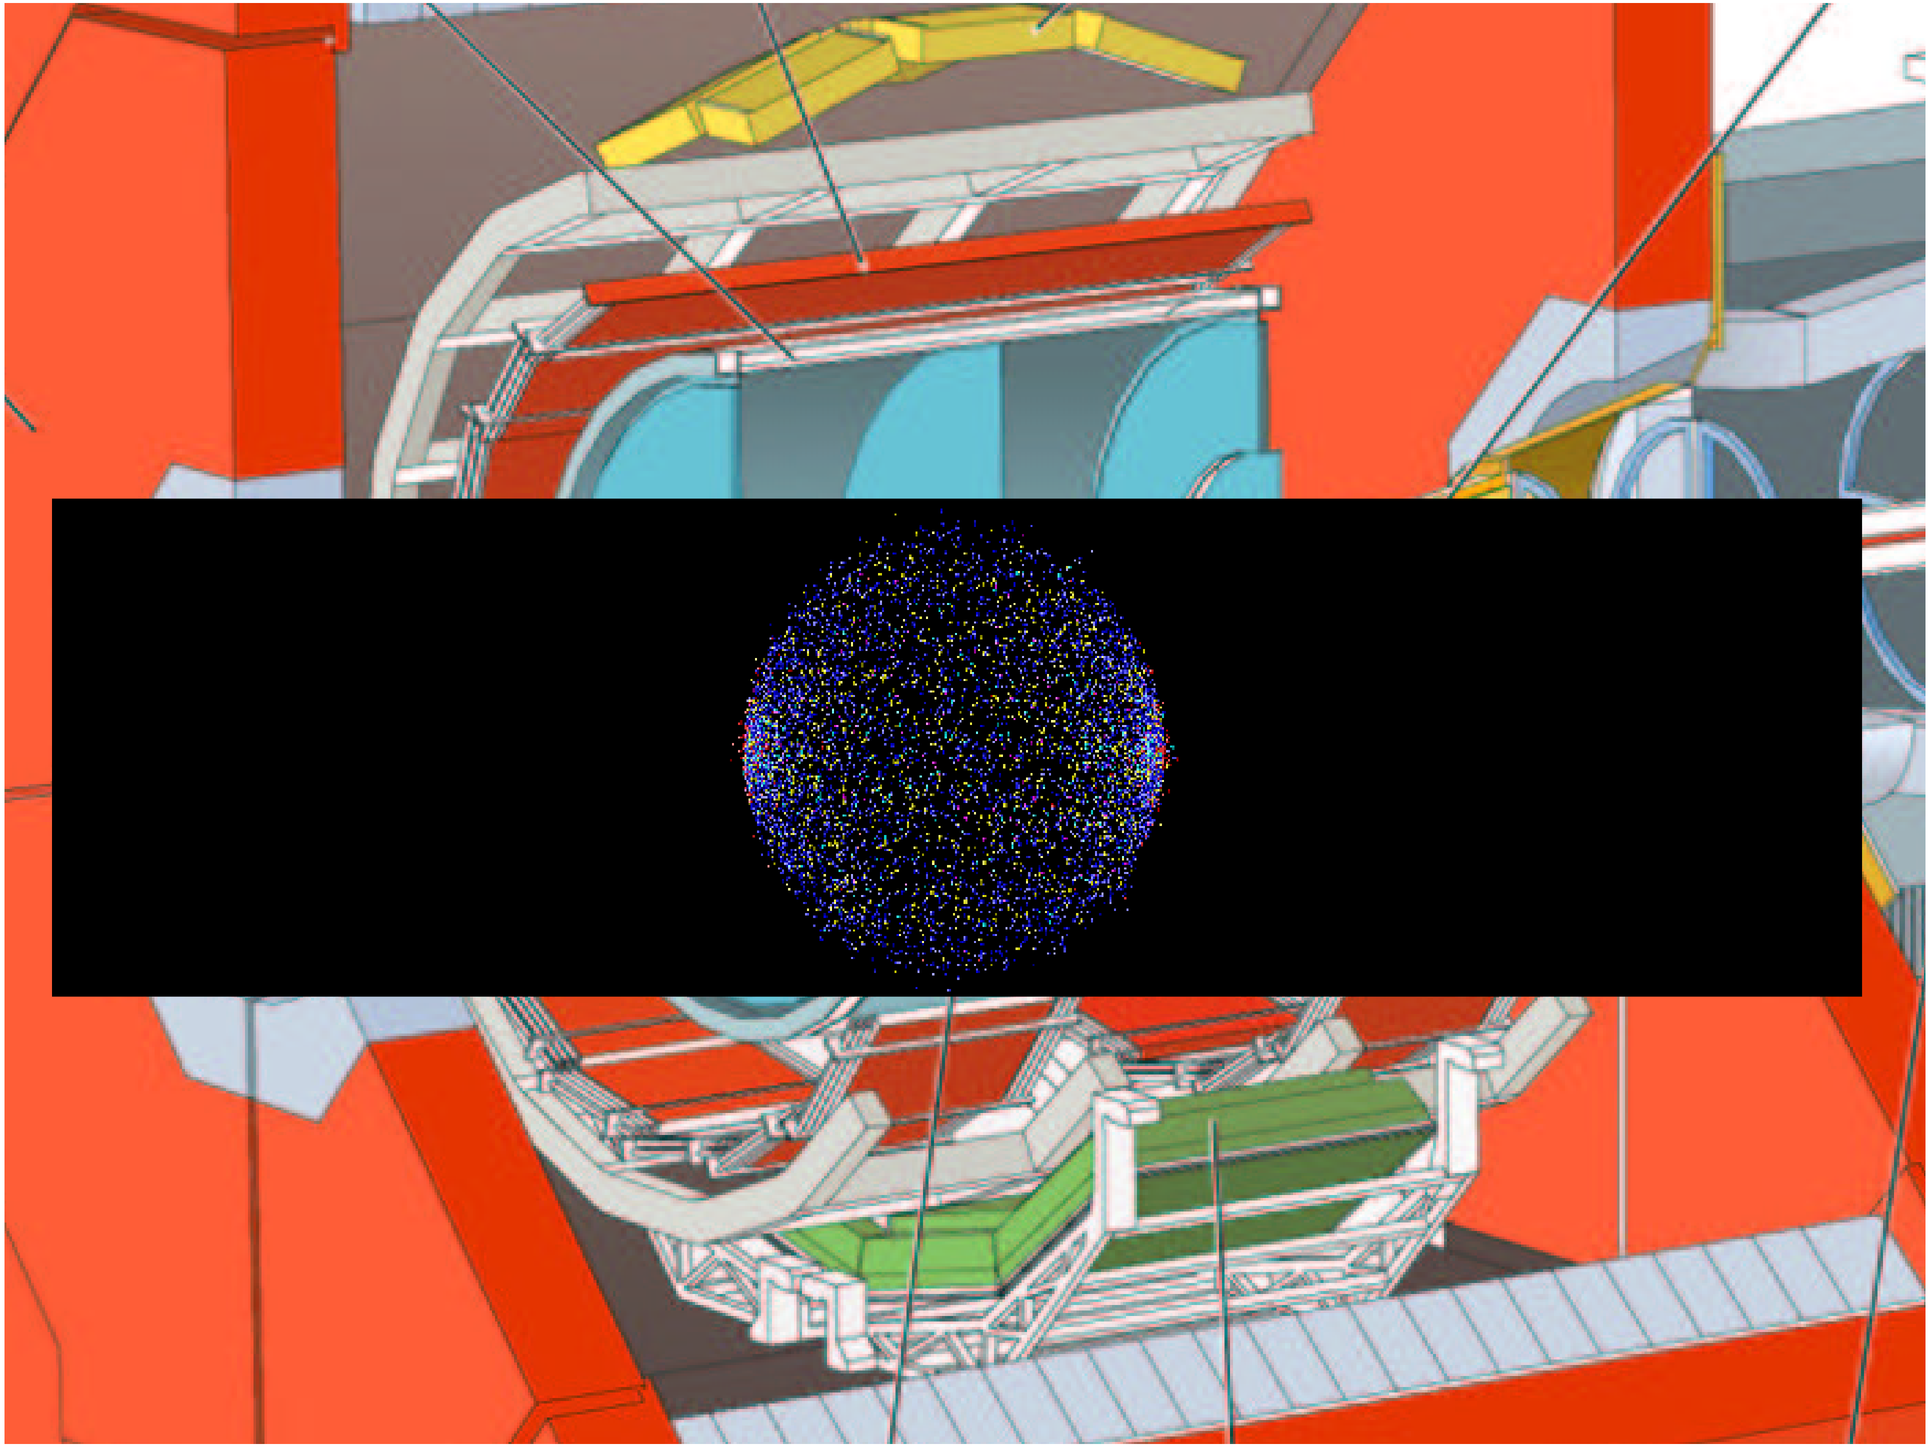

Trigger

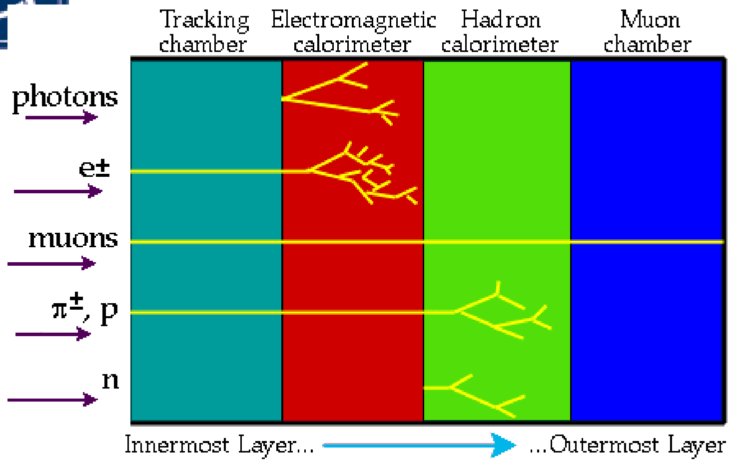

Multi-level trigger system

Reject background

Select most interesting collisions

Reduce total data volume

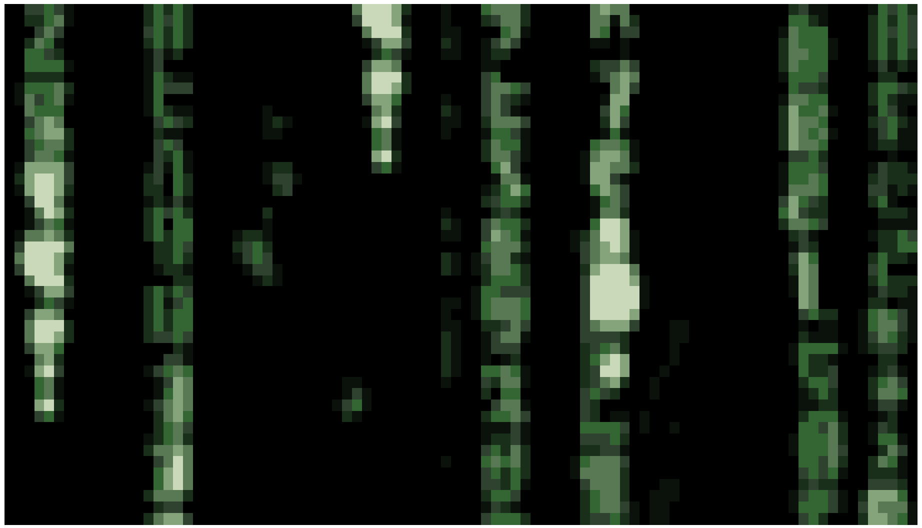

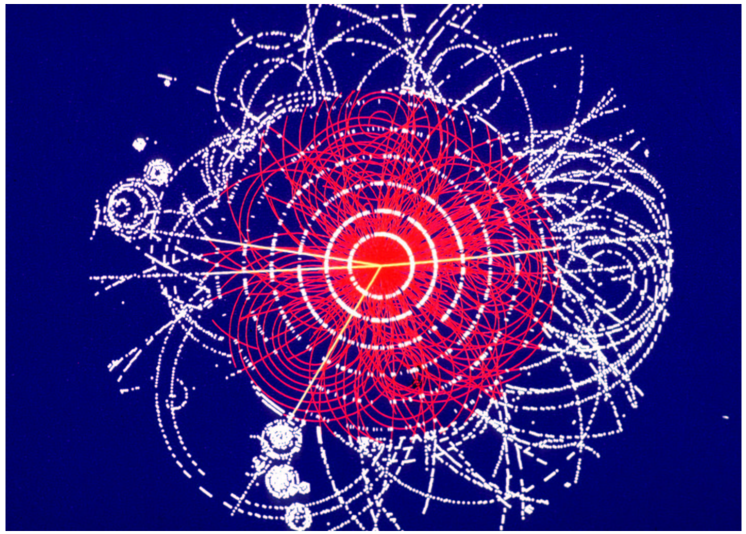

Data acquisition

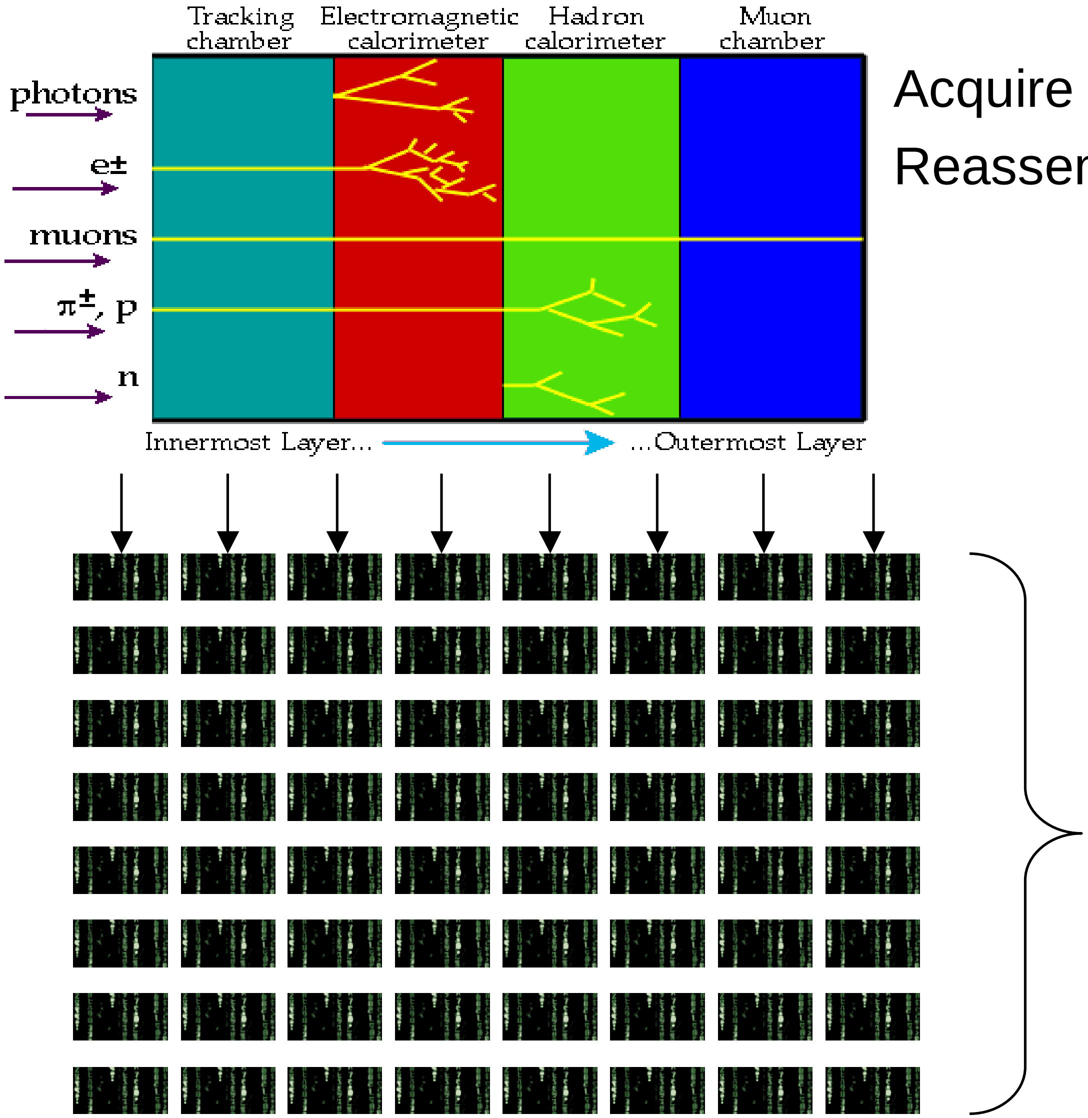

Acquire data from 1000's of sources

Reassemble all the data of same event

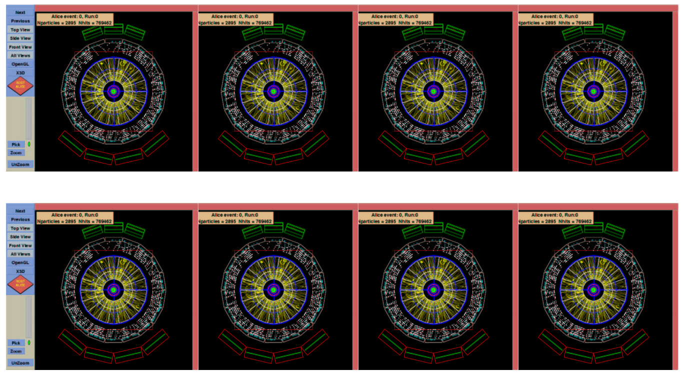

Online dataflow

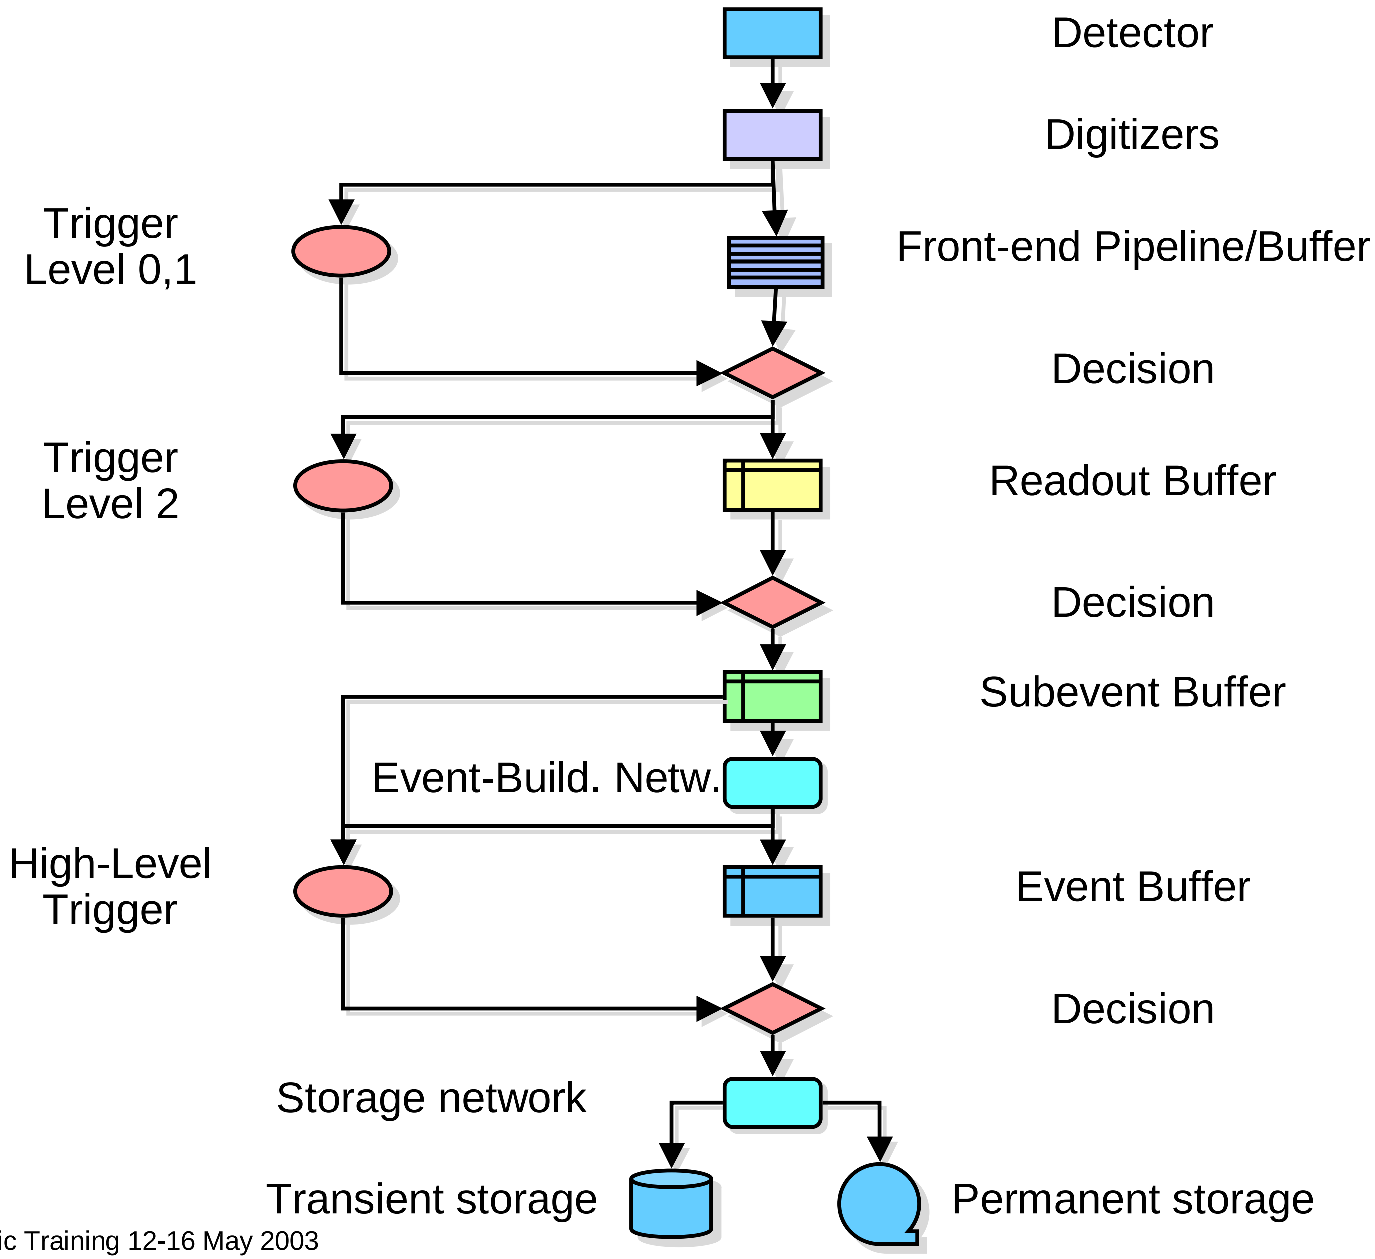

Offline packages

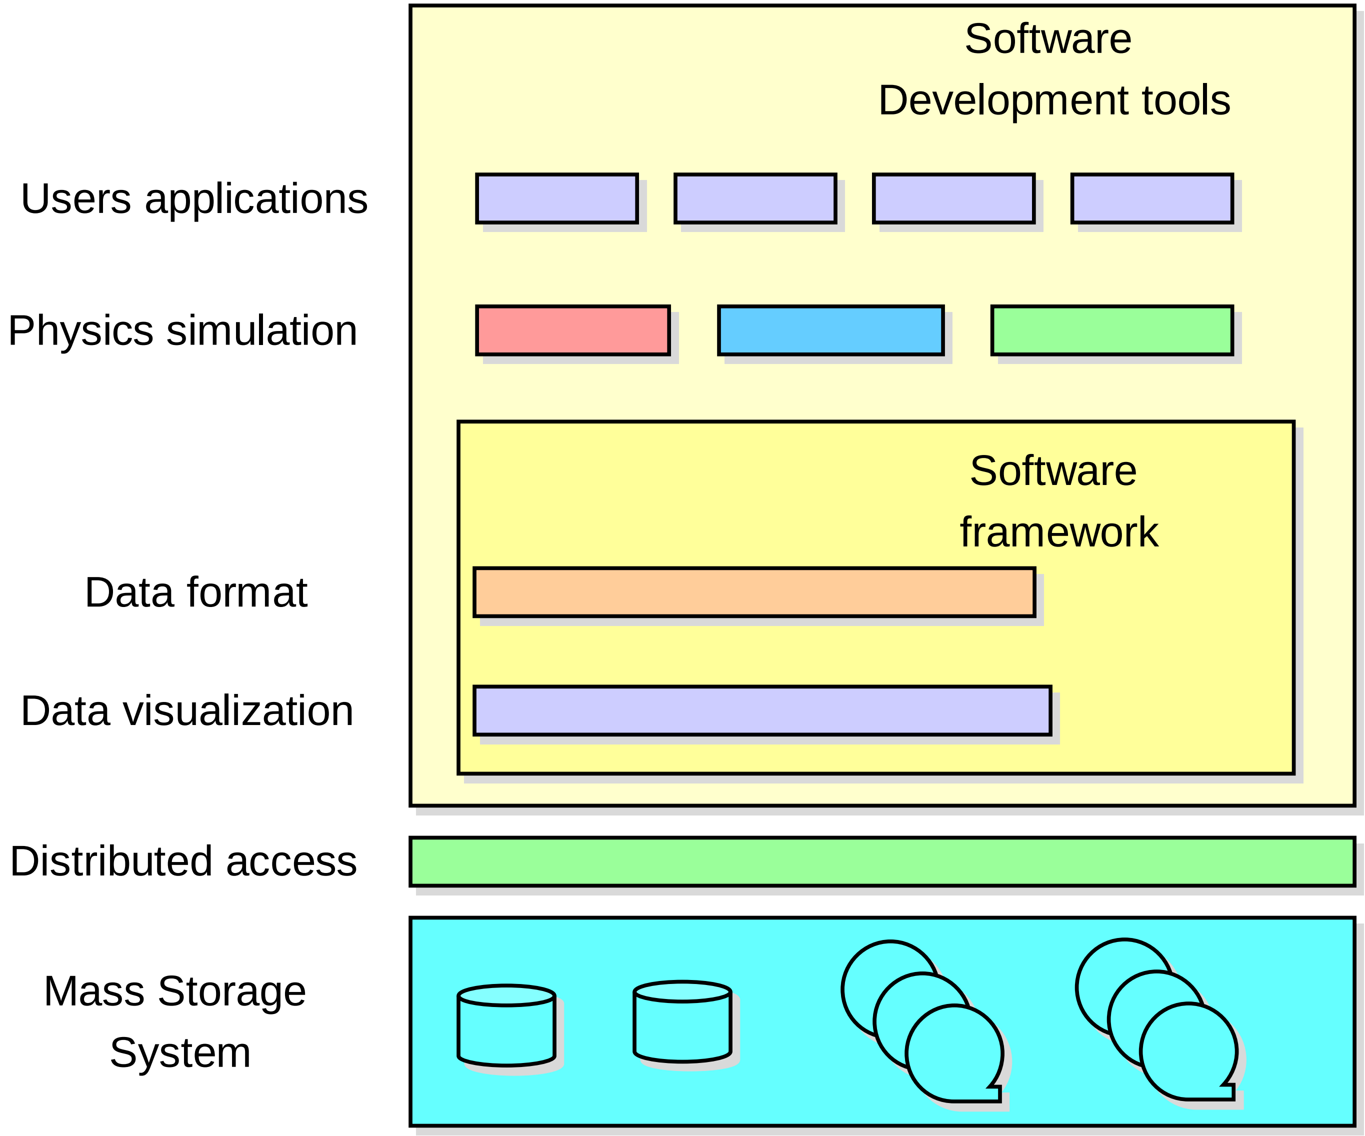

Experiment dataflow

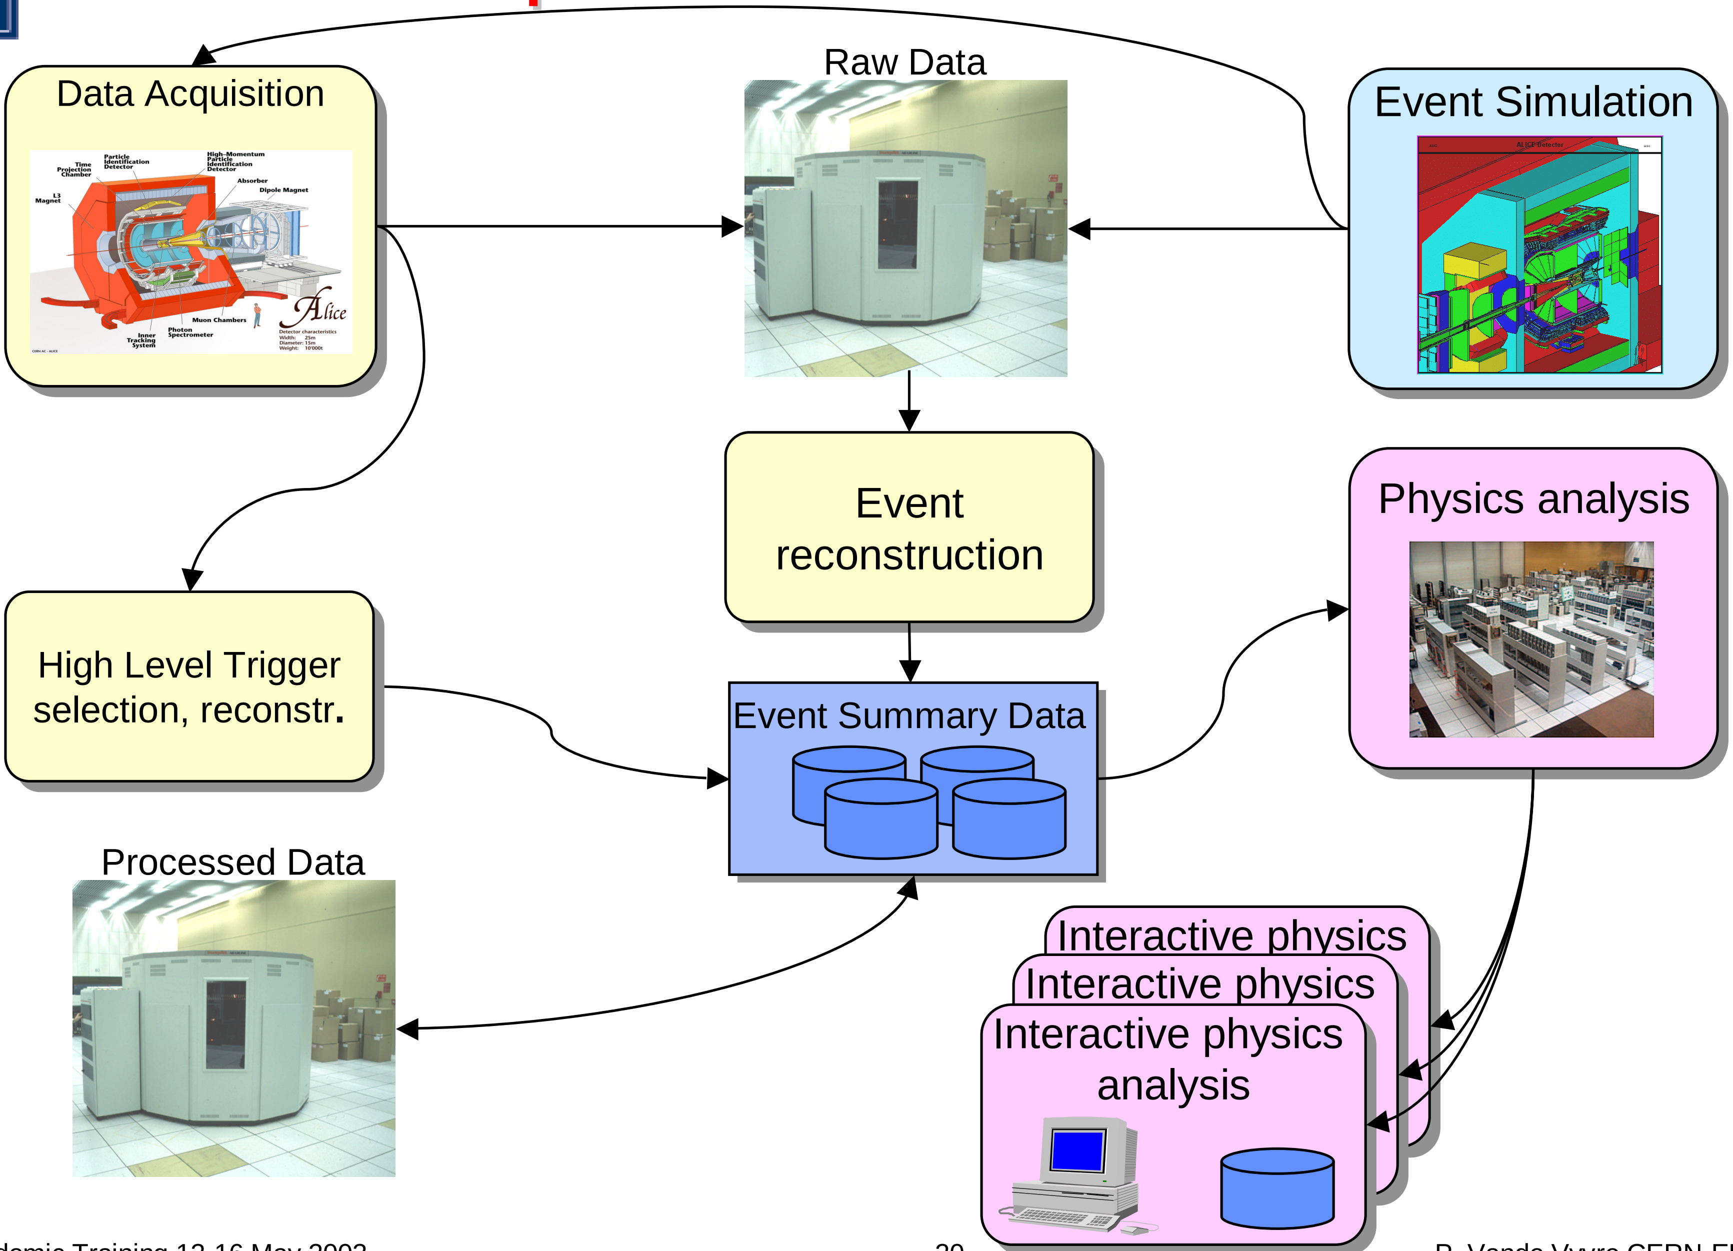

Day 1

- Outline of this series
- Main concepts

- Requirements of the LHC experiments
  - Trigger @ LHC
  - Data acquisition @ LHC
  - Data storage @ LHC

- Data Challenges
  - Evolutions of online and offline computing fabrics
  - Motivations of data challenges

LHC experimental data: from today's Data Challenges to the promise of tomorrow (1)

- The LHC experiments constitute a challenge for electronics, data acquisition, processing, and analysis.

Trigger @ LHC (1)

		# Trigger Levels	Rate First Level Trigger (Hz)	
ALICE	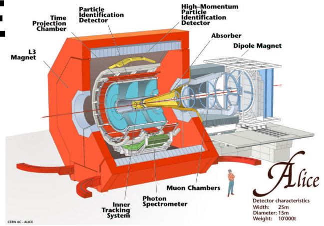	4	Pb-Pb p-p	$6x10^3$ 10^3
ATLAS	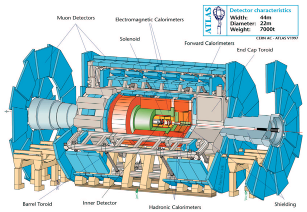	3	L 1 L 2	10^5 $2x10^3$
CMS	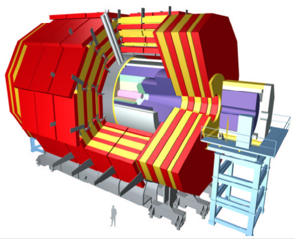	2	L 1	10^5
LHCb	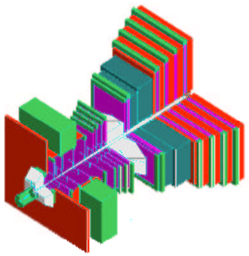	3	L 0 L 1	10^6 $4x10^4$

Trigger @ LHC (2)

ALICE	ATLAS	CMS	LHCb	
40	40	40	40	MHz
0.9/5.2	2.5	2.5	4/<2000	µs
6	75	100	1100/40	kHz
	120			GByte/s
0.08	10			ms
2	2			kHz
~200	~100	~100	~200	Hz

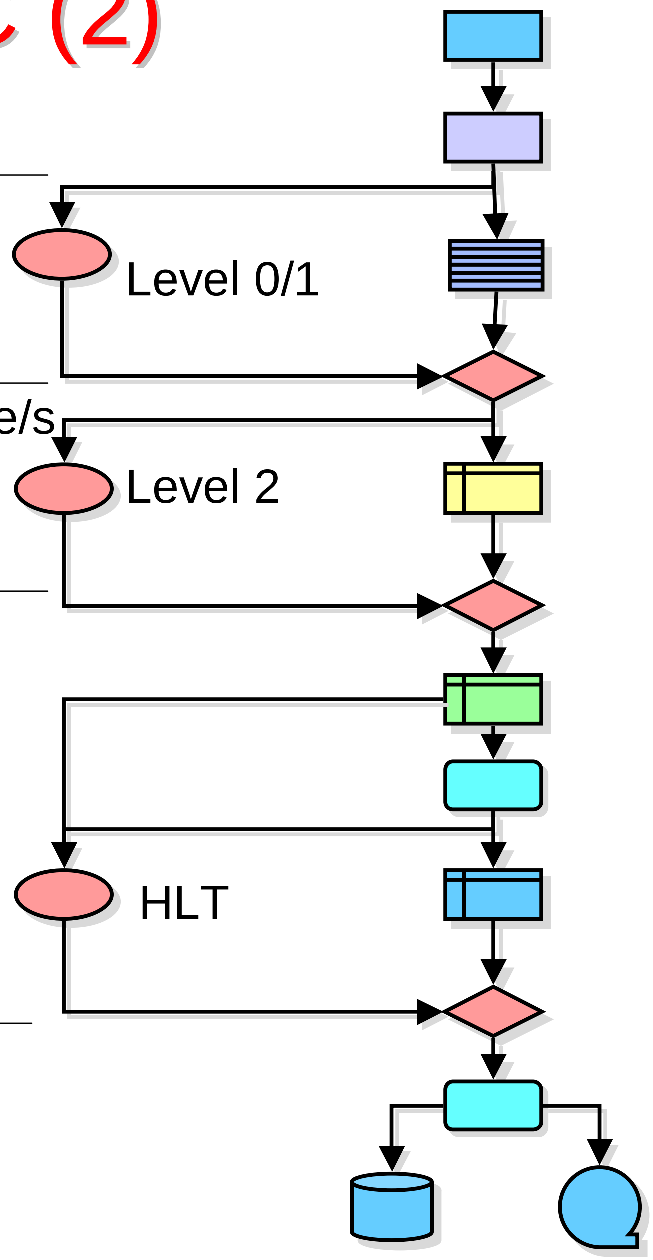

DAQ @ LHC (1)

			Event Size (Byte)	Readout (HLT input) (Events/s.)	Readout (HLT input) (GB/s)
ALICE	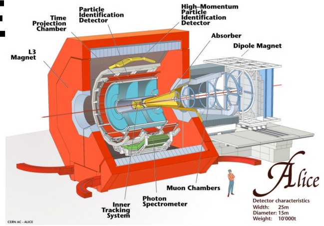	Pb-Pb	5×10^7	2×10^3	25
		pp	2×10^6	10^2	1
ATLAS	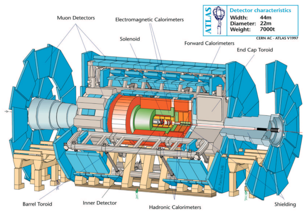		10^6	2×10^3	10
CMS	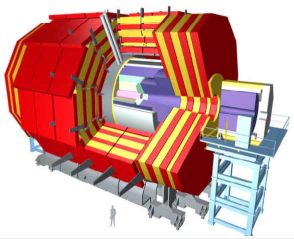		10^6	10^5	100
LHCb	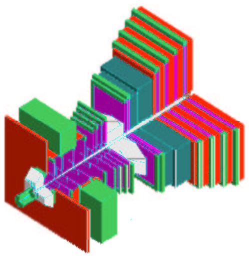		2×10^5	40×10^4	4

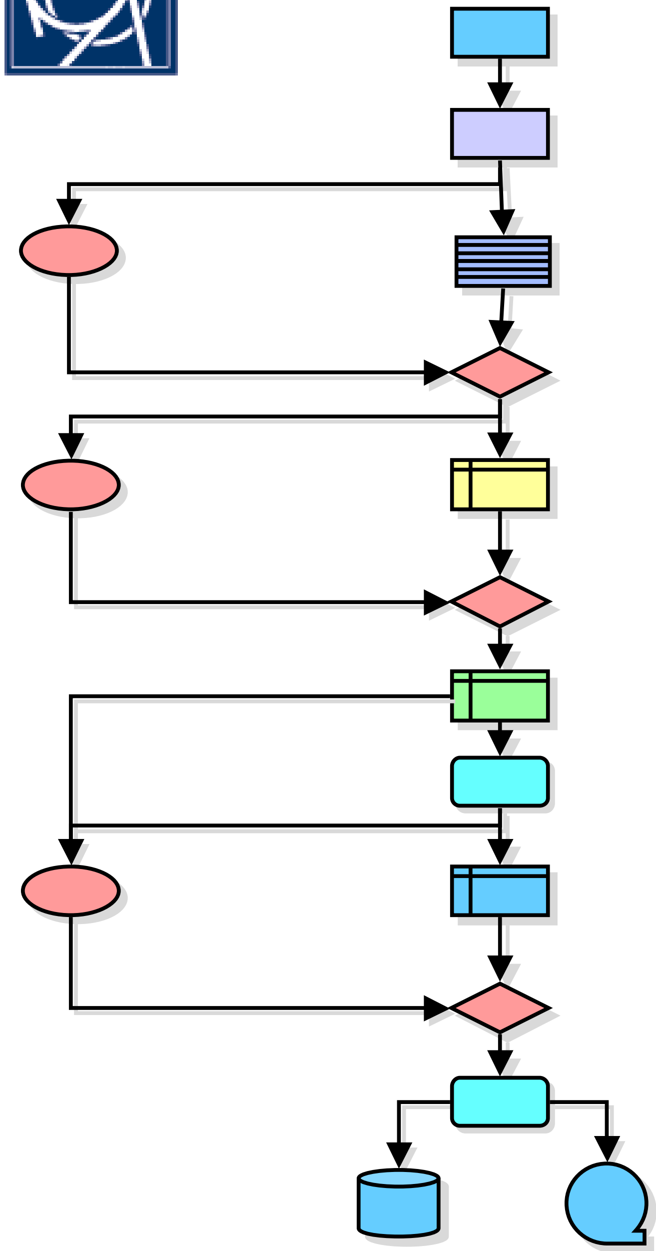

DAQ @ LHC (2)

ALICE	ATLAS	CMS	LHCb	
25	10	100	4	GBytes/s
2.5	6			GBytes/s
200	100	100	40	MBytes/s
1250	300	100		MBytes/s

Mass Storage @ LHC

			Readout (HLT output) (Events/s.)	Readout (HLT output) (MB/s)	Data archived Total/year (PBytes)
ALICE	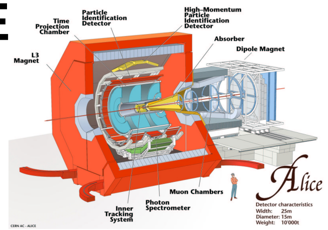	Pb-Pb	$2x10^2$	1250	2.3
		pp	10^2	200	
ATLAS	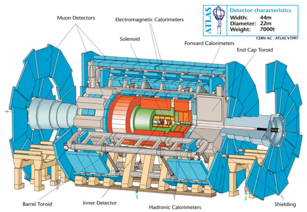	Pb-Pb		300	6.0
		pp	10^2	100	
CMS	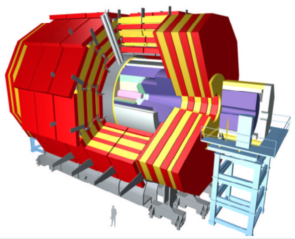	Pb-Pb		100	3.0
		pp	10^2	100	
LHCb	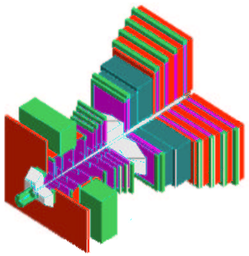		$2x10^2$	40	1.0

Rates & Bandwidths @ LHC

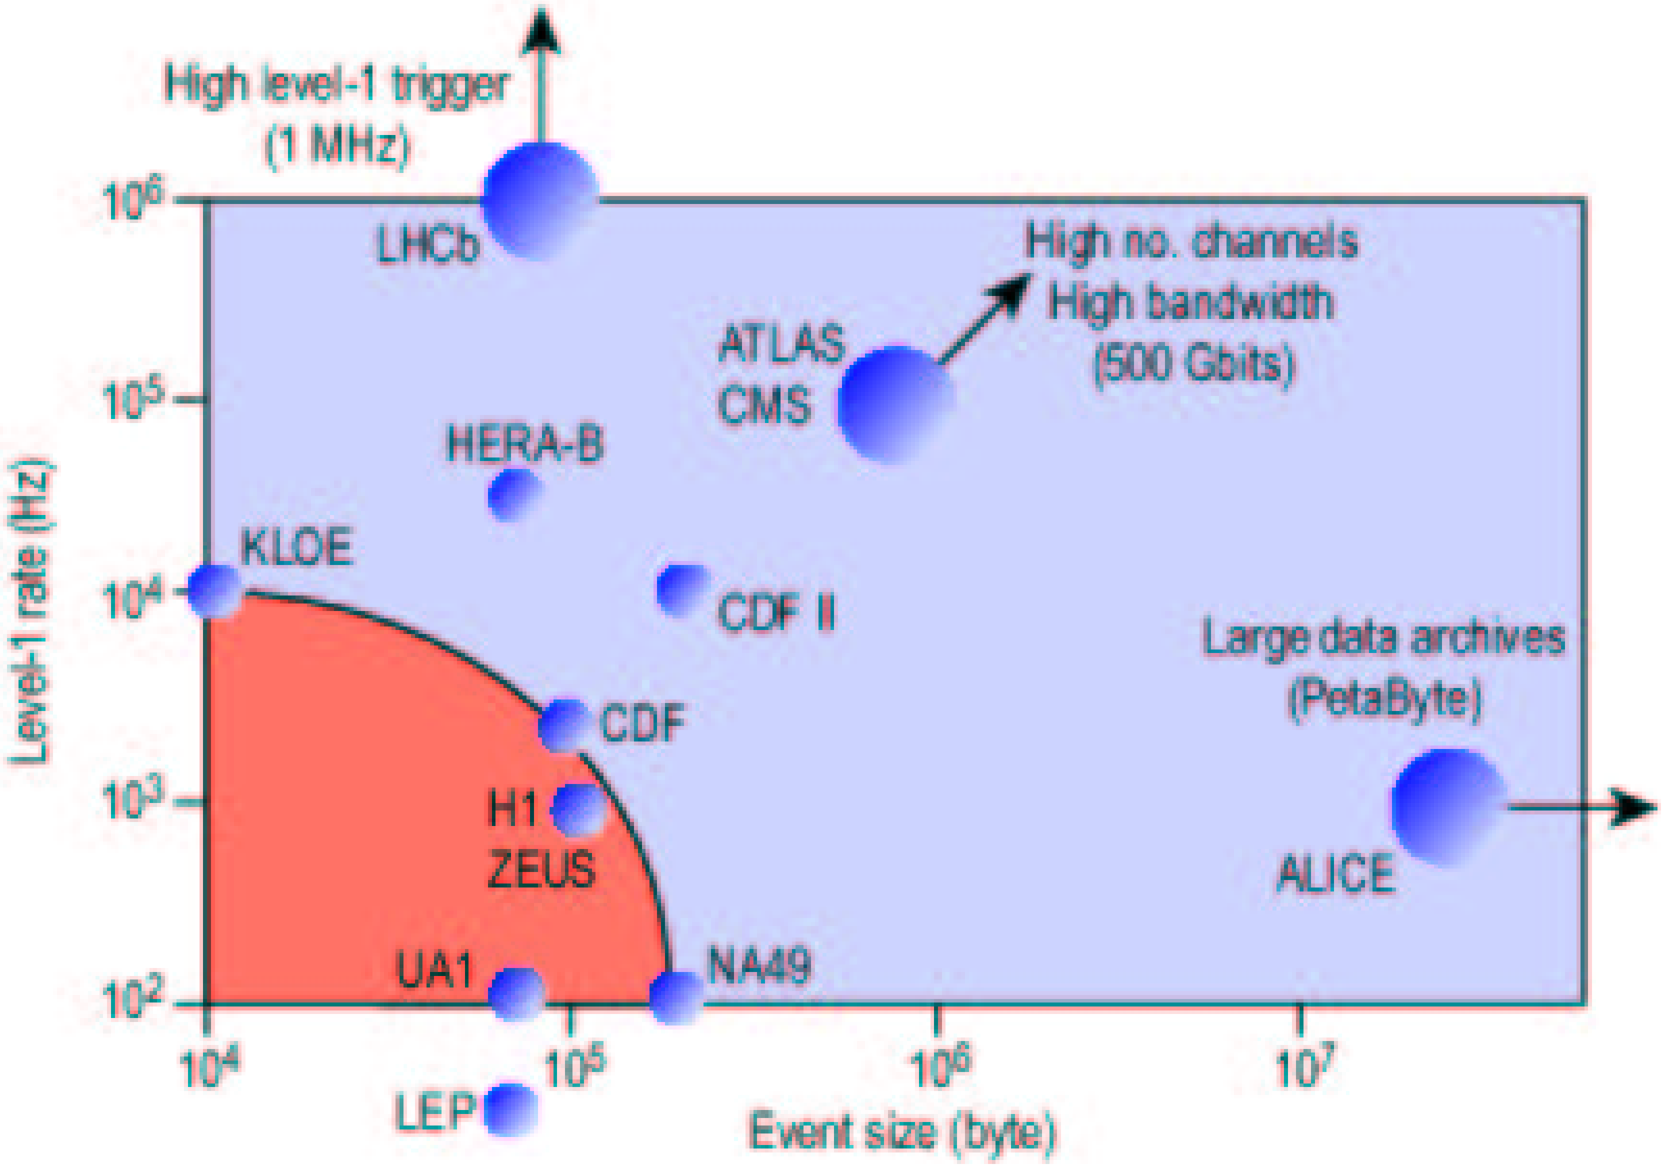

LHC experimental data: from today's Data Challenges to the promise of tomorrow (2)

- The LHC experiments constitute a challenge for electronics, data acquisition, processing, and analysis.
- This challenge has been addressed by many years of R&D activity during which prototypes of components or subsystems have been developed.

“R&D humanum est” (1)

◆RD-27
First-level trigger systems for LHC experiments.

◆RD-11
EAST Embedded architectures for second-level triggering in LHC experiments

◆LCB_005
Event Filter Farm

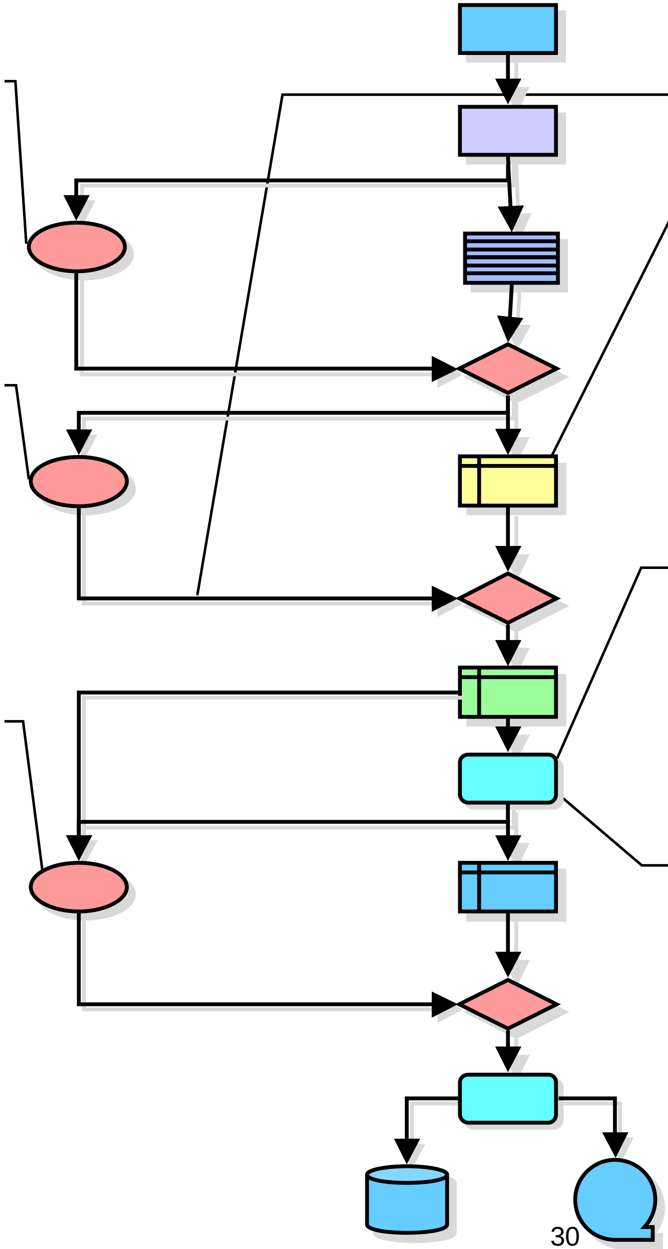

◆RD-12
Readout system test benches.

◆RD-13
A scalable data taking system at a test beam for LHC.

◆ RD-24
Applications of the scalable coherent interface to data acquisition at LHC (SCI).

◆RD-31
NEBULAS: An asynchronous self-routing packet-switching network architecture for event building in high rate experiments (ATM).

"R&D humanum est" (2)

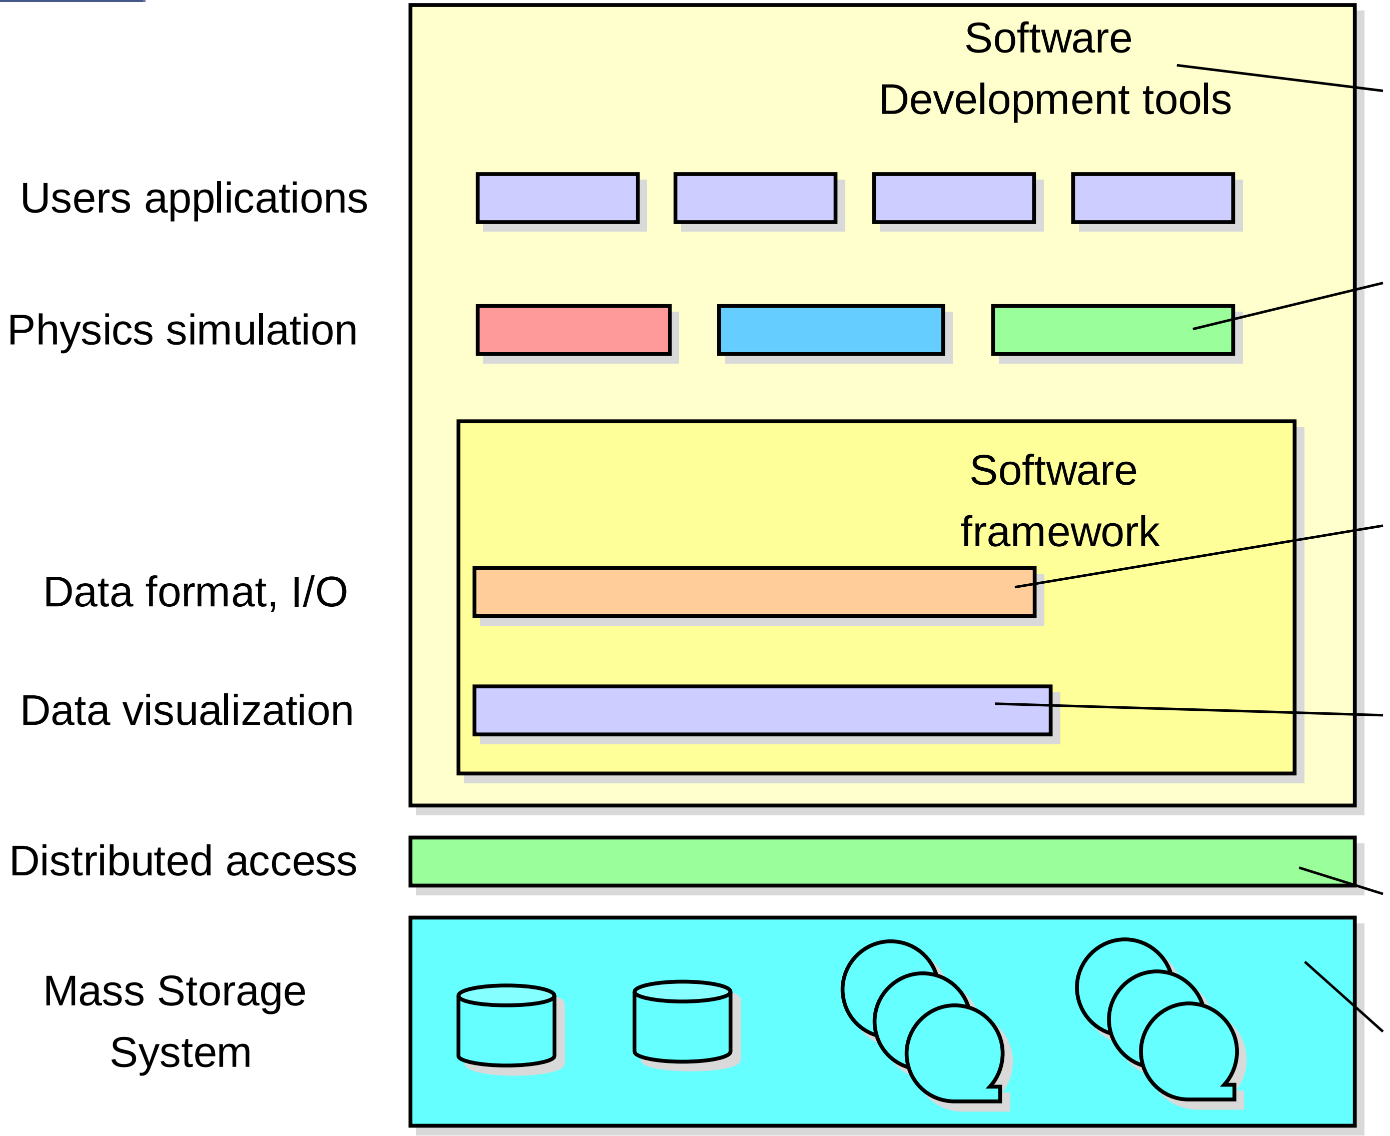

- RD 41 - MOOSE
- LCB_006 – SPIDER

- GEANT 3
- RD44 GEANT 4
- FLUKA

- ROOT I/O
- RD-45 - OODBMS

- ROOT display
- LCB_001 – LHC++

- LCB_003 – MONARC
- GRID projects

- HPSS
- Eurostore
- CASTOR

Outcome of R&D

- Design and implementation of hardware components
  - TTC system for the trigger distribution
- Design and implementation of software packages
  - ROOT package
- Proof of concept of major concepts
  - Positive recommendation of using a communication switch for the event building based on tests with ATM. Different technologies considered today (Gigabit Ethernet, Myrinet).
- Positive recommendation of technologies
  - Object Oriented (OO) programming for the LHC software.
- None or few negative recommendations but some technologies have not been adopted by experiments
  - OO database for the storage of raw data
  - Usage of Windows for physics data processing

LHC experimental data: from today's Data Challenges to the promise of tomorrow (3)

- The LHC experiments constitute a challenge for data acquisition, processing, and analysis.
- This challenge has been addressed by many years of R&D activity during which prototypes of components or subsystems have been developed.
- The present generation of prototypes used for the LHC data acquisition and computing infrastructures are based on commodity components.

Moore’s Law

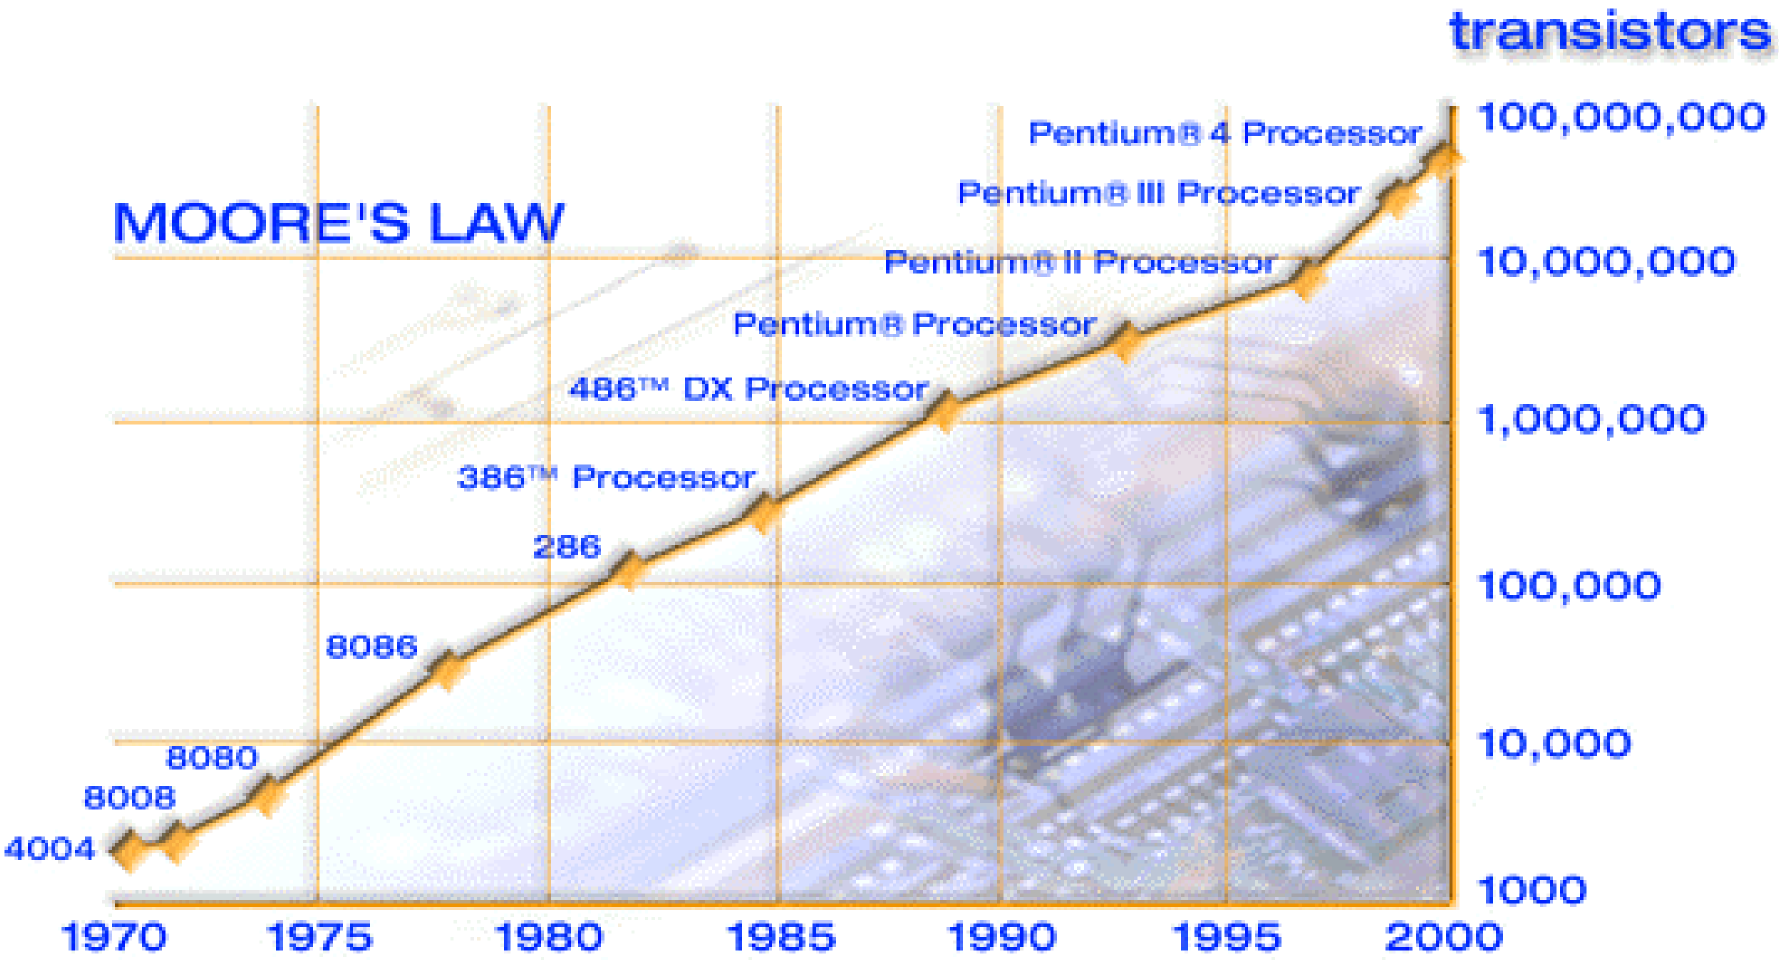
transistors
MOORE'S LAW
Pentium® 4 Processor
Pentium® III Processor
Pentium® II Processor
Pentium® Processor
486™ DX Processor
386™ Processor
286
8086
8080
8008
4004
100,000,000
10,000,000
1,000,000
100,000
10,000
1000
1970
1975
1980
1985
1990
1995
2000

Chip key parameters

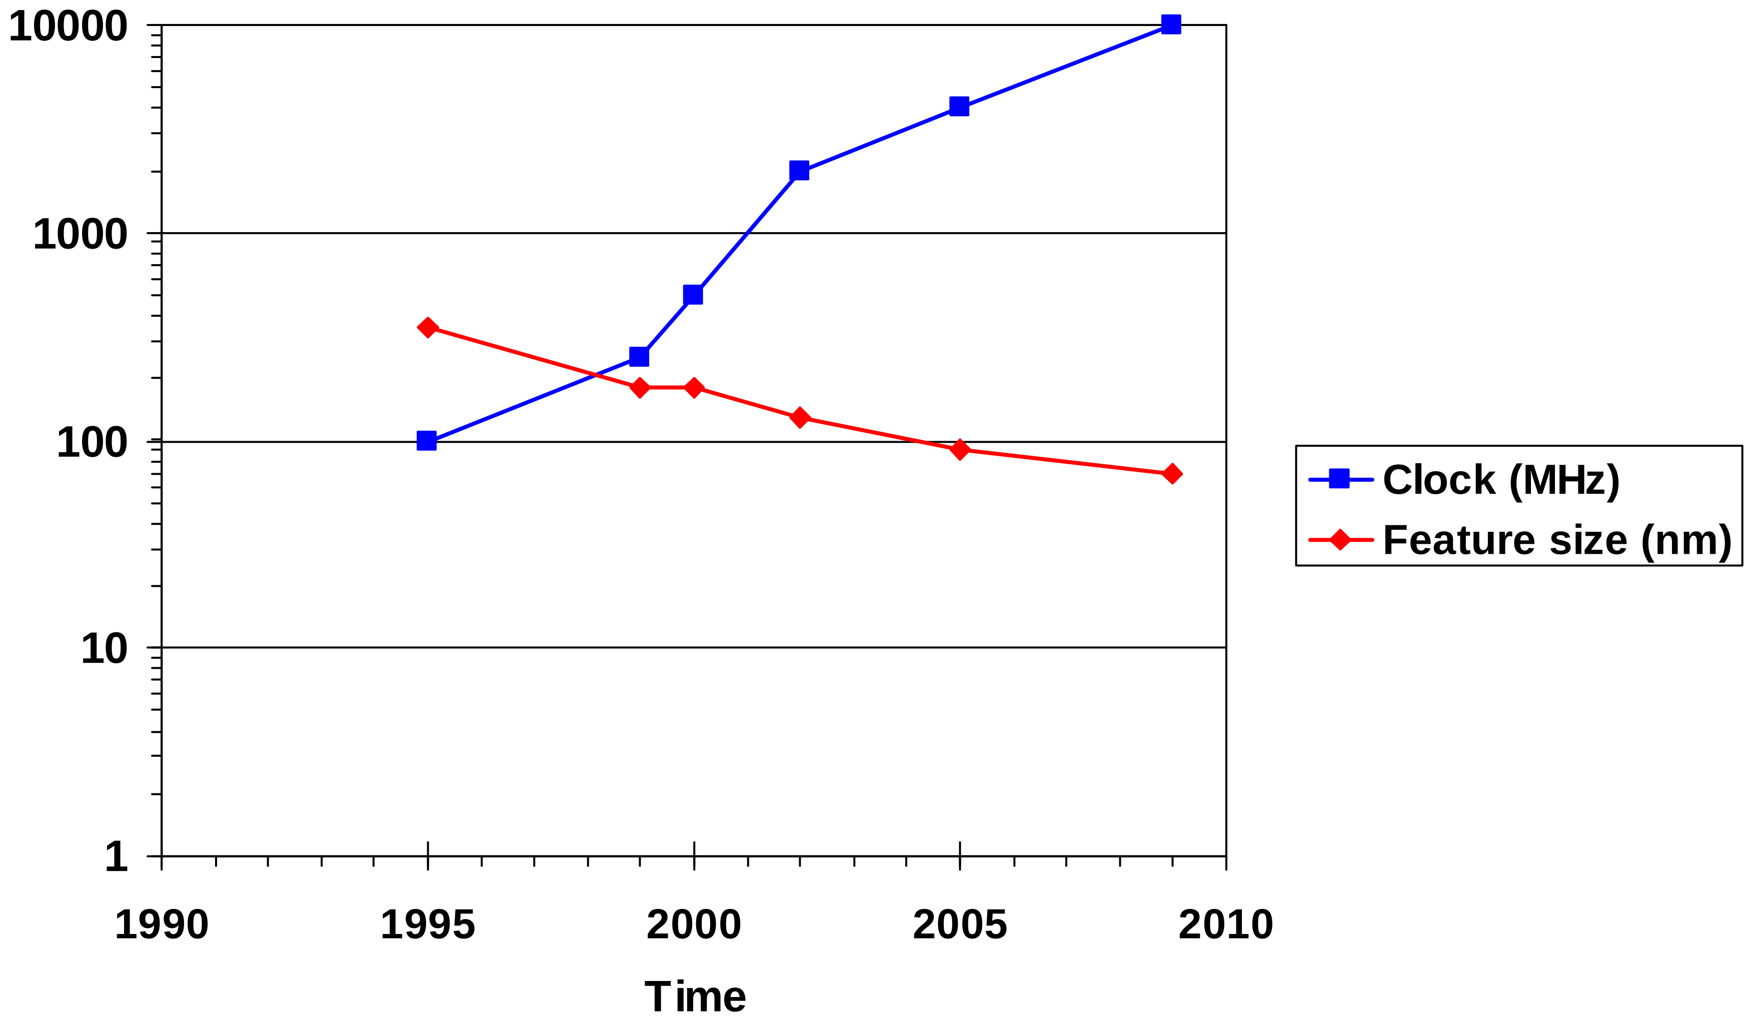

Memory capacity

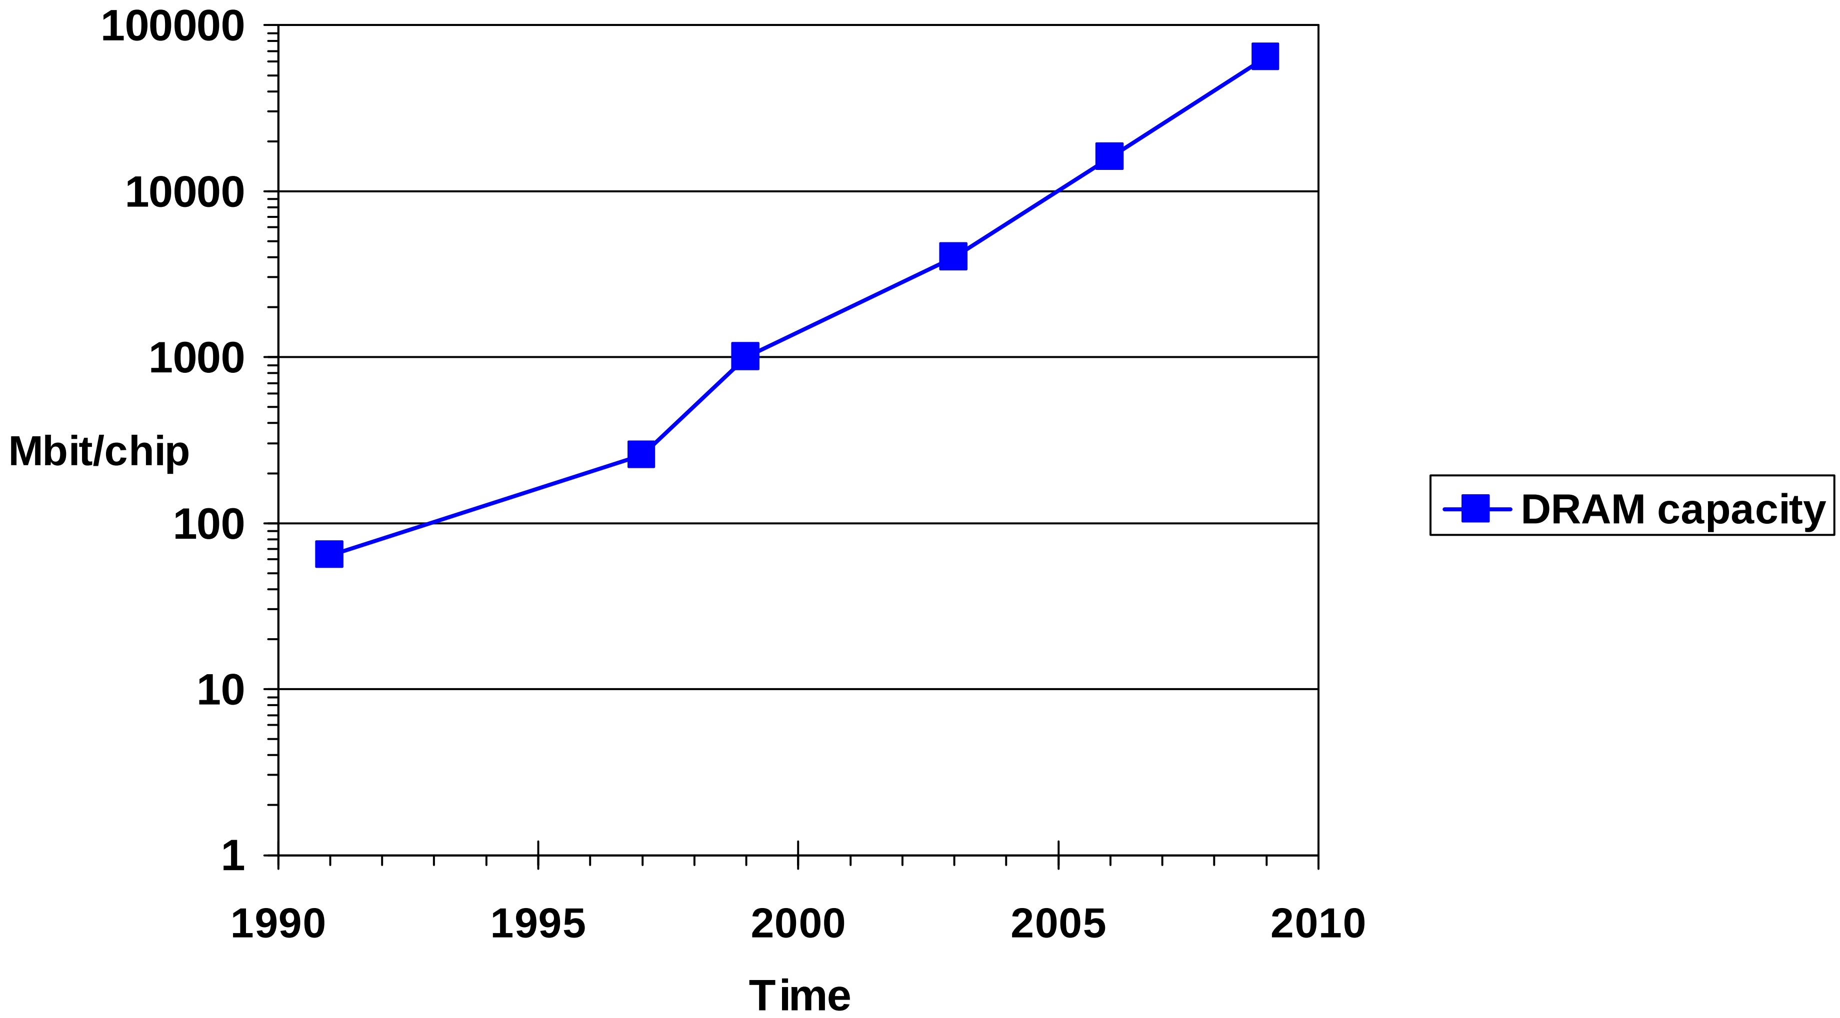

Memory and I/O bus Bandwidth

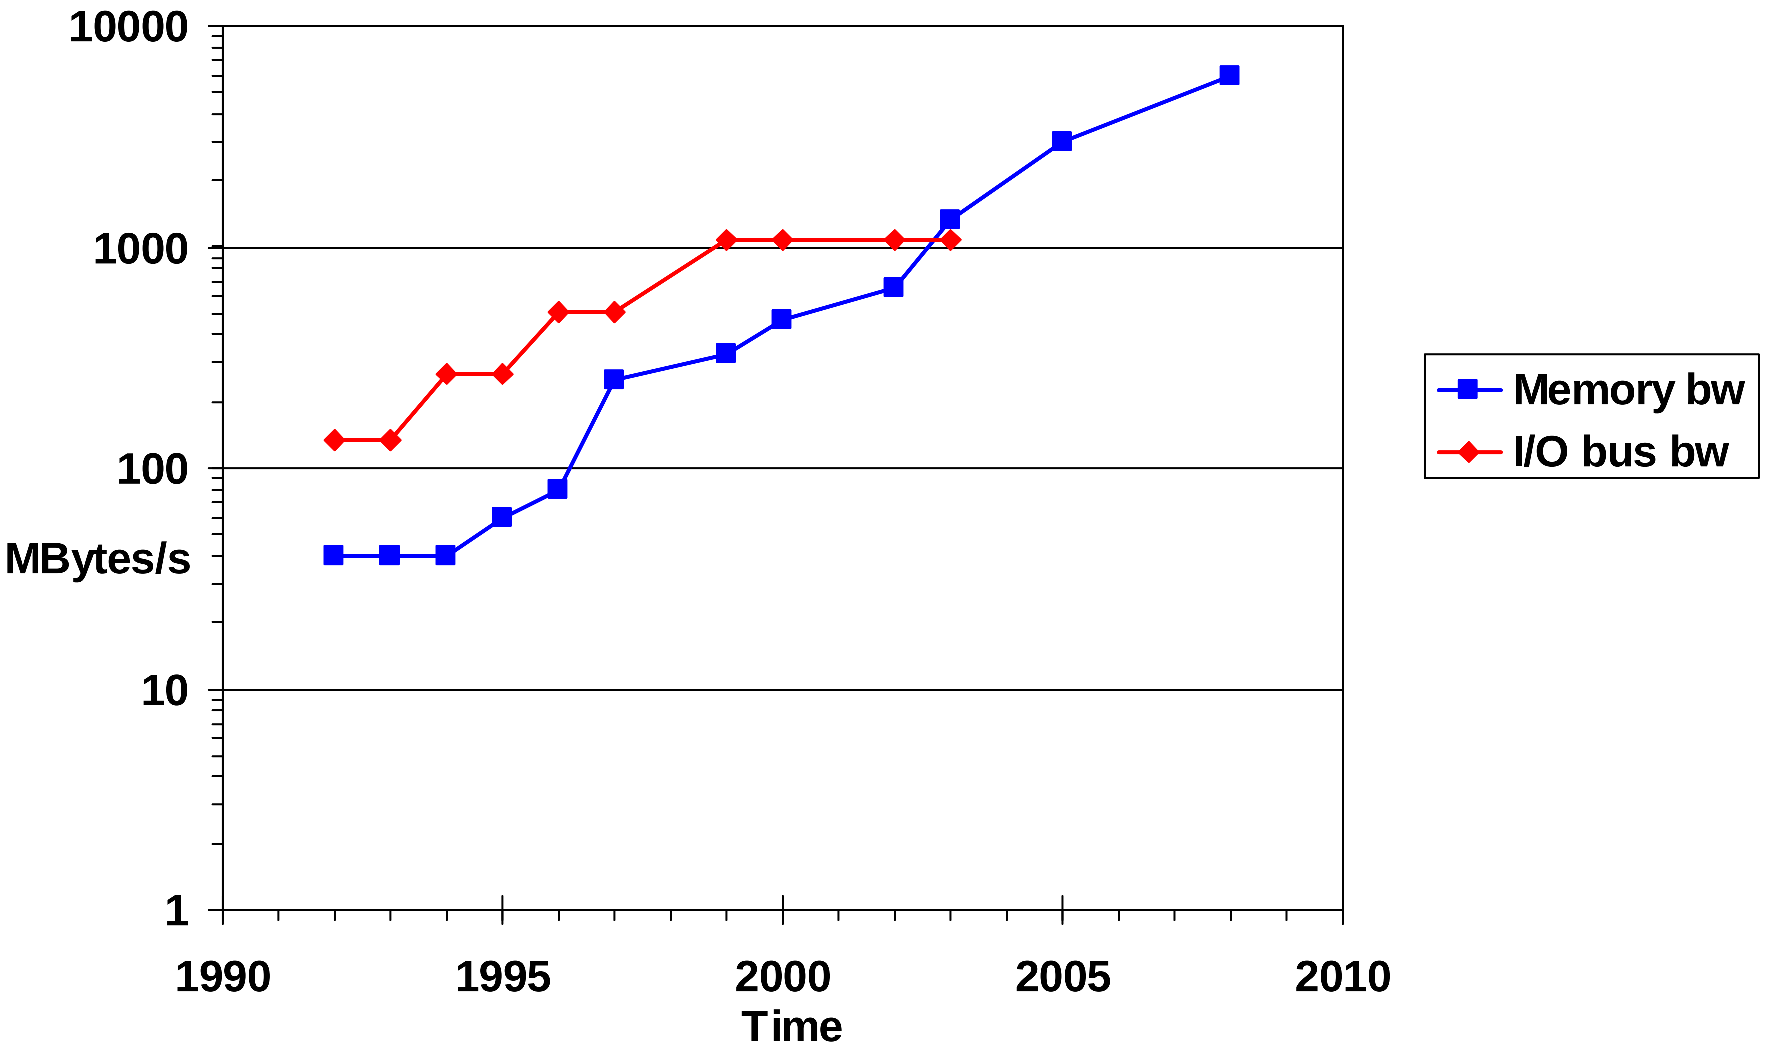

Networking technology

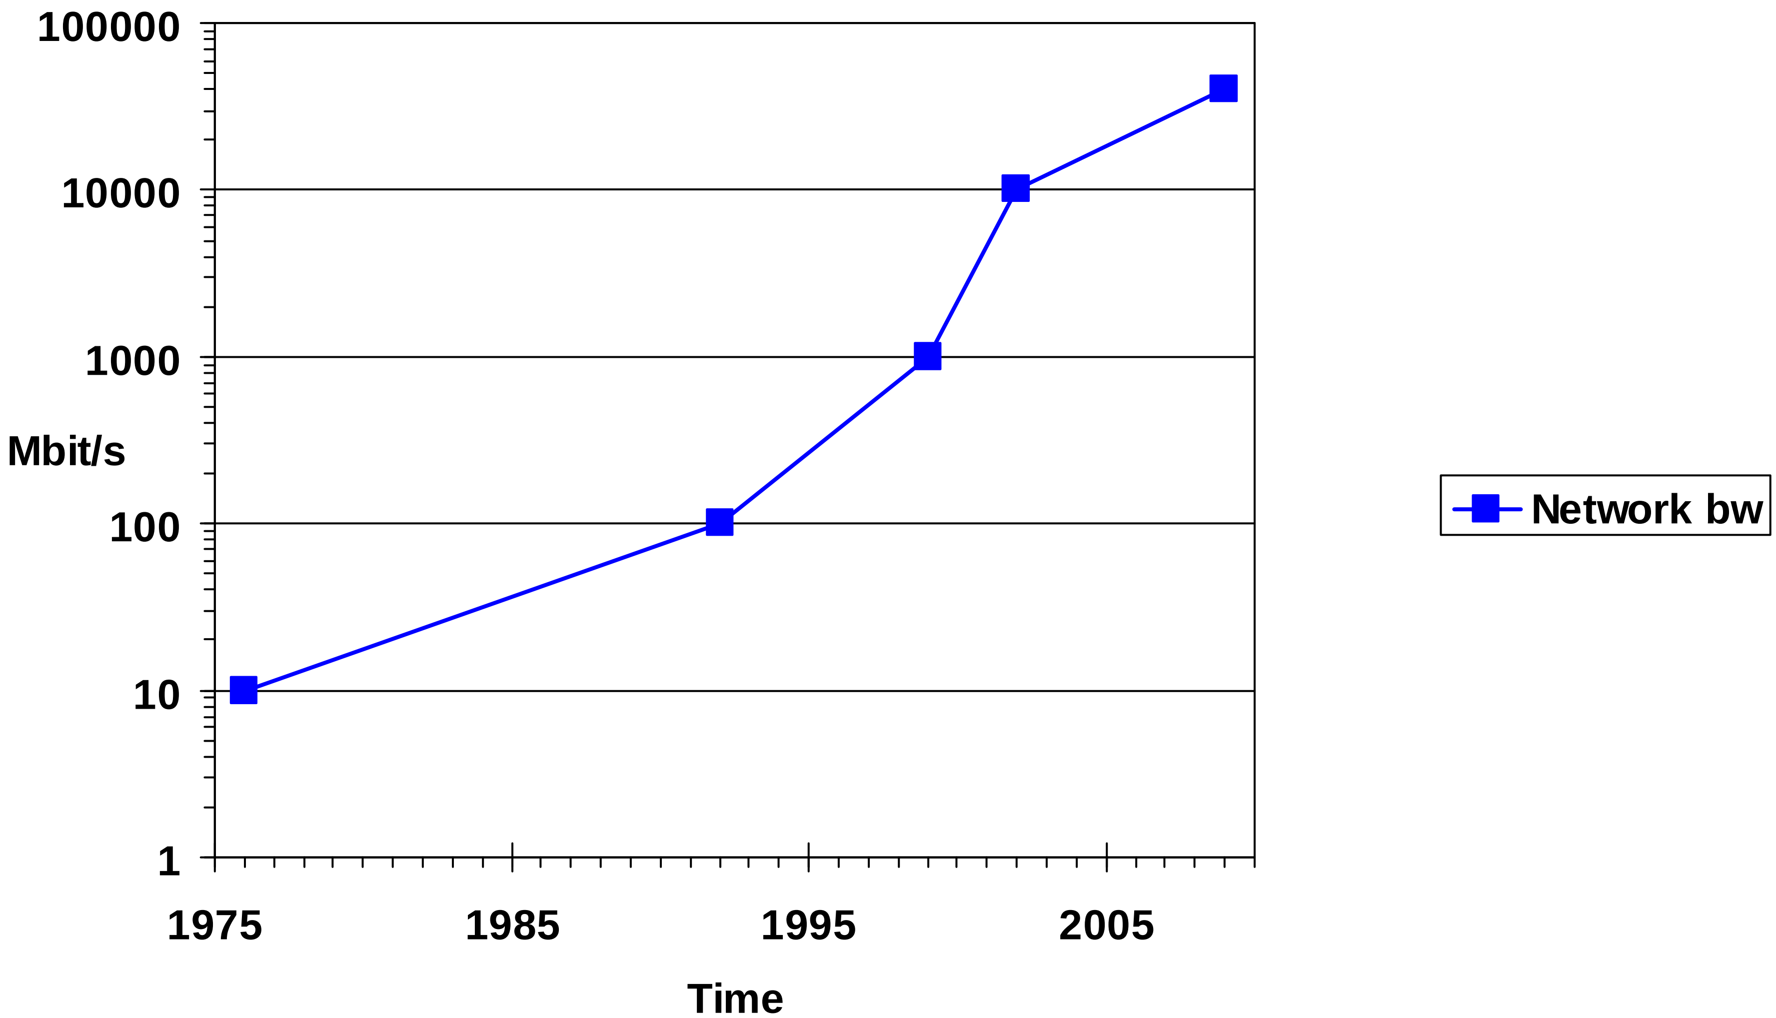

Moore's law: myth and reality (1)

- Observation by G. Moore in 1965 when working at Fairchild
  - "Cramming more components onto integrated circuits", Electronics Vol. 38 Nb 8, April19, 1965
  - "Complexity of minimum cost semiconductor component had doubled every year".
  - Cost per integrated component ≈ 1/number of components integrated

    But yield decreases when components added

    Minimum cost at any point in time
- In 1975, prediction that doubling every 2 years
  - G. Moore co-founded Intel
  - His law became the Intel business model
  - Initially applied to memory chips, then to processors
- Interpretation and evolution of Moore's law
  - In the 1980's: ⇒ doubling of transistors on a chip every 18 months
  - In the 1990's: ⇒ doubling of microprocessor power every 18 months
- Subject of debate in the semiconductor industry. However…
  - Intel: in 1971 the 4004 had 2250 transistors, in 2000 the PIV had 42 Millions
  - Exponential evolution over 30 years

Moore's law: myth and reality (2)

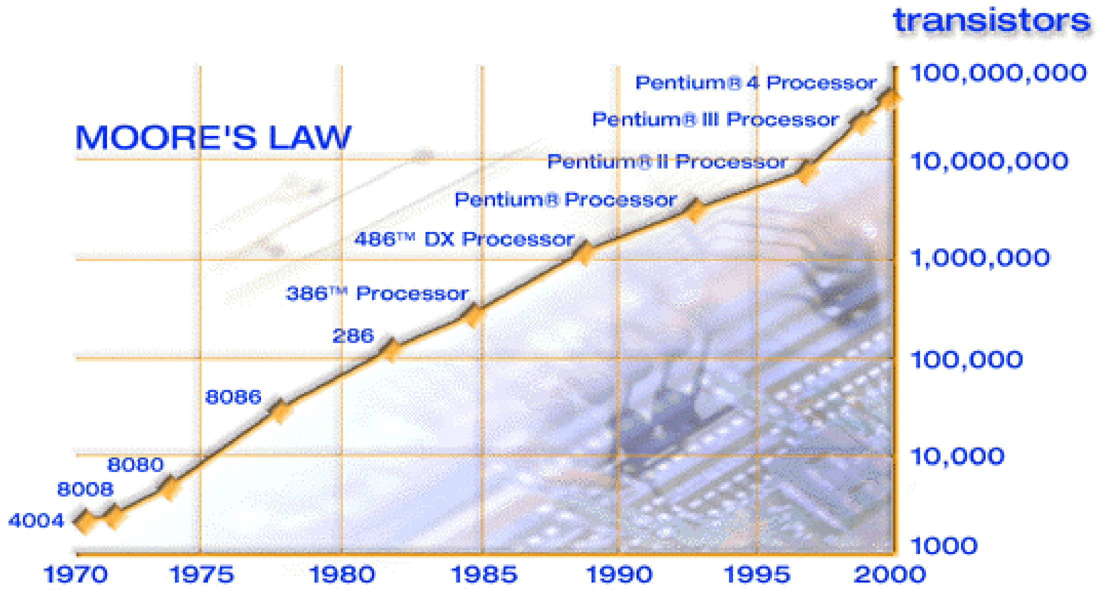

© Intel corp.

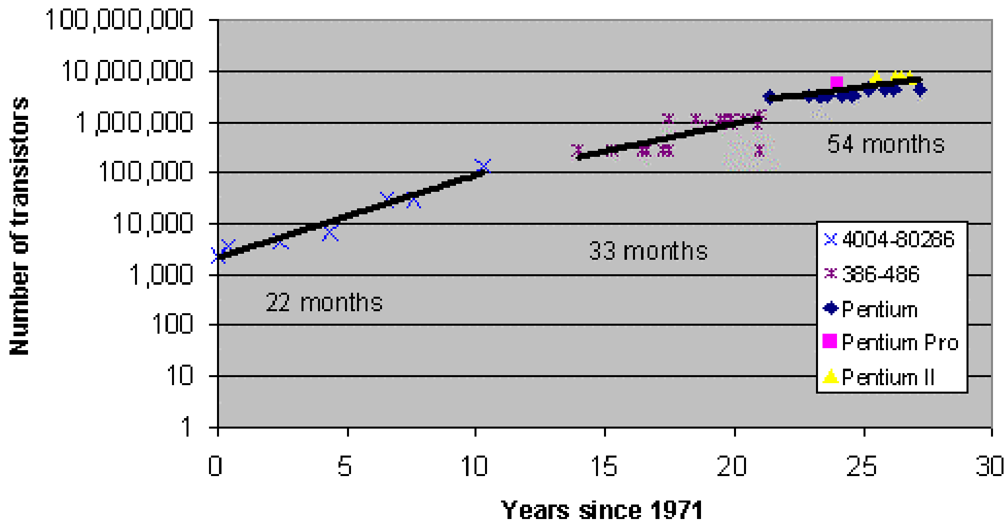

Mr. Illka Tuomi

LHC experimental data: from today's Data Challenges to the promise of tomorrow (4)

- The LHC experiments constitute a challenge for data acquisition, processing, and analysis.
- This challenge has been addressed by many years of R&D activity during which prototypes of components or subsystems have been developed.
- The present generation of prototypes used for the LHC data acquisition and computing infrastructures are based on commodity components.
- This prototyping phase is culminating now with an evaluation of the prototypes in large-scale tests ("Data Challenges").

Online Systems Evolution

- Dramatic evolution thanks to chip integration:
  - Electronics more and more sophisticated and intelligent
  - Data multiplexing, filtering, compression and formatting on chip
  - Electronics migrate from racks to detectors
  - Decrease of number of electronics slots needed in standard racks
- Dramatic increase of the DAQ bandwidth needed
- The rack of the year 2000 is a PC !

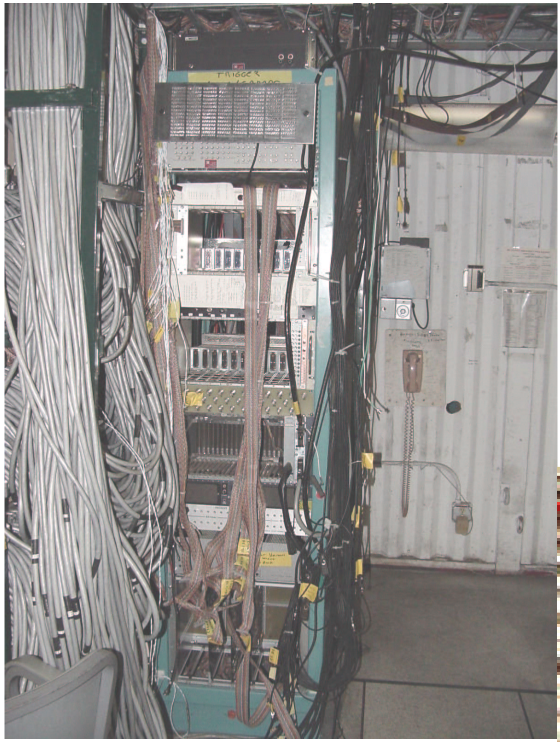

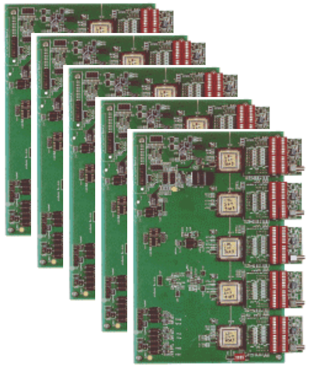

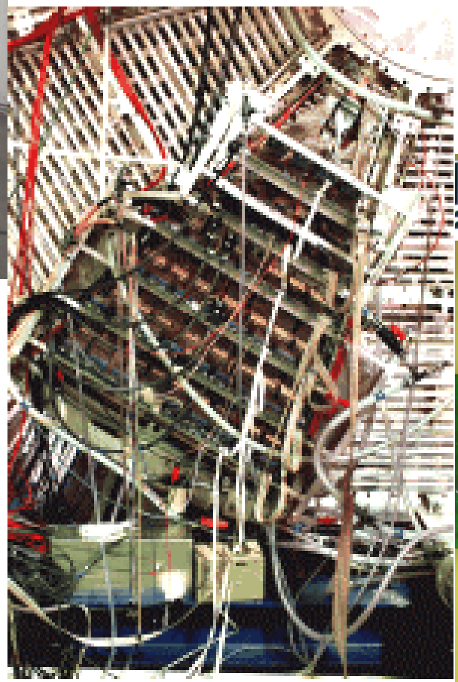

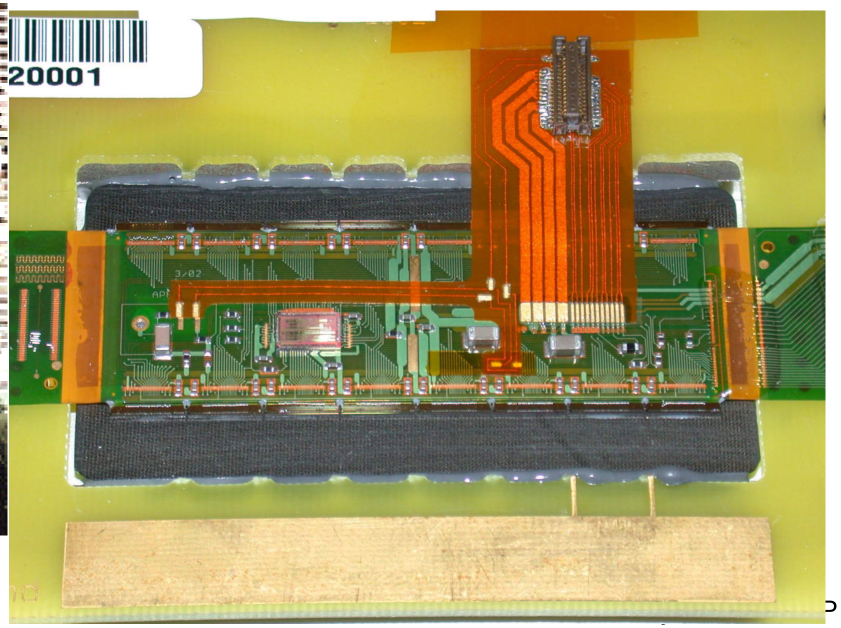
20001

Computing Center Evolution

- Large scientific computing centers:
  - No more mainframes and specialized networks
  - Massive transition to computing farms
- For HEP experiments, the computing centre is providing “online services”
  - Physics data archives
  - Computing power factory
  - File repository
- With the GRID, the online will not be limited to the experimental area: the world will be online !
  - Virtual access to the control room
  - Fast and remote access to the experimental data
- The computing center is not offline any longer

Building Blocks

- Commodity is (almost) unique by definition
- Massive move to commodity in online and offline
  - Identical or similar building blocks to build the fabrics
  - Processing power: PCs based on Intel or compatible processors
  - Operating system: Linux
  - Networking: Gigabit Ethernet
  - Storage
    - **Transient: IDE-based disks**
    - **Permanent: not (yet ?) commodity**
- Opportunity to use the same test bed for several activities

Why do we need Data Challenges ?

- More and more requirements to online systems
  - DAQ and HLT systems becoming larger and larger
  - Similar to a computing center
- System made of 100s of boxes from different manufacturers
  - Integration work transferred from computer manufacturer to farm integration teams
  - Need to test the system at large
- Buy as late a possible
  - Large integration work starting at the installation time $\Rightarrow$ large risk
- System = Hw + Sw
  - Scaling and/or combination effects
  - Combined system testing as early as possible

Data Challenges

- “Challenge”:
  - An accusation, reproach
  - The act of calling to account
  - A summons to fight, to single combat or duel
  - A difficult or demanding task, one seen as a test of one’s abilities
- Data Challenge
  - Yearly exercise
  - Hardware and software
  - Online, offline, computing center
  - “Here and now”

MSS tests

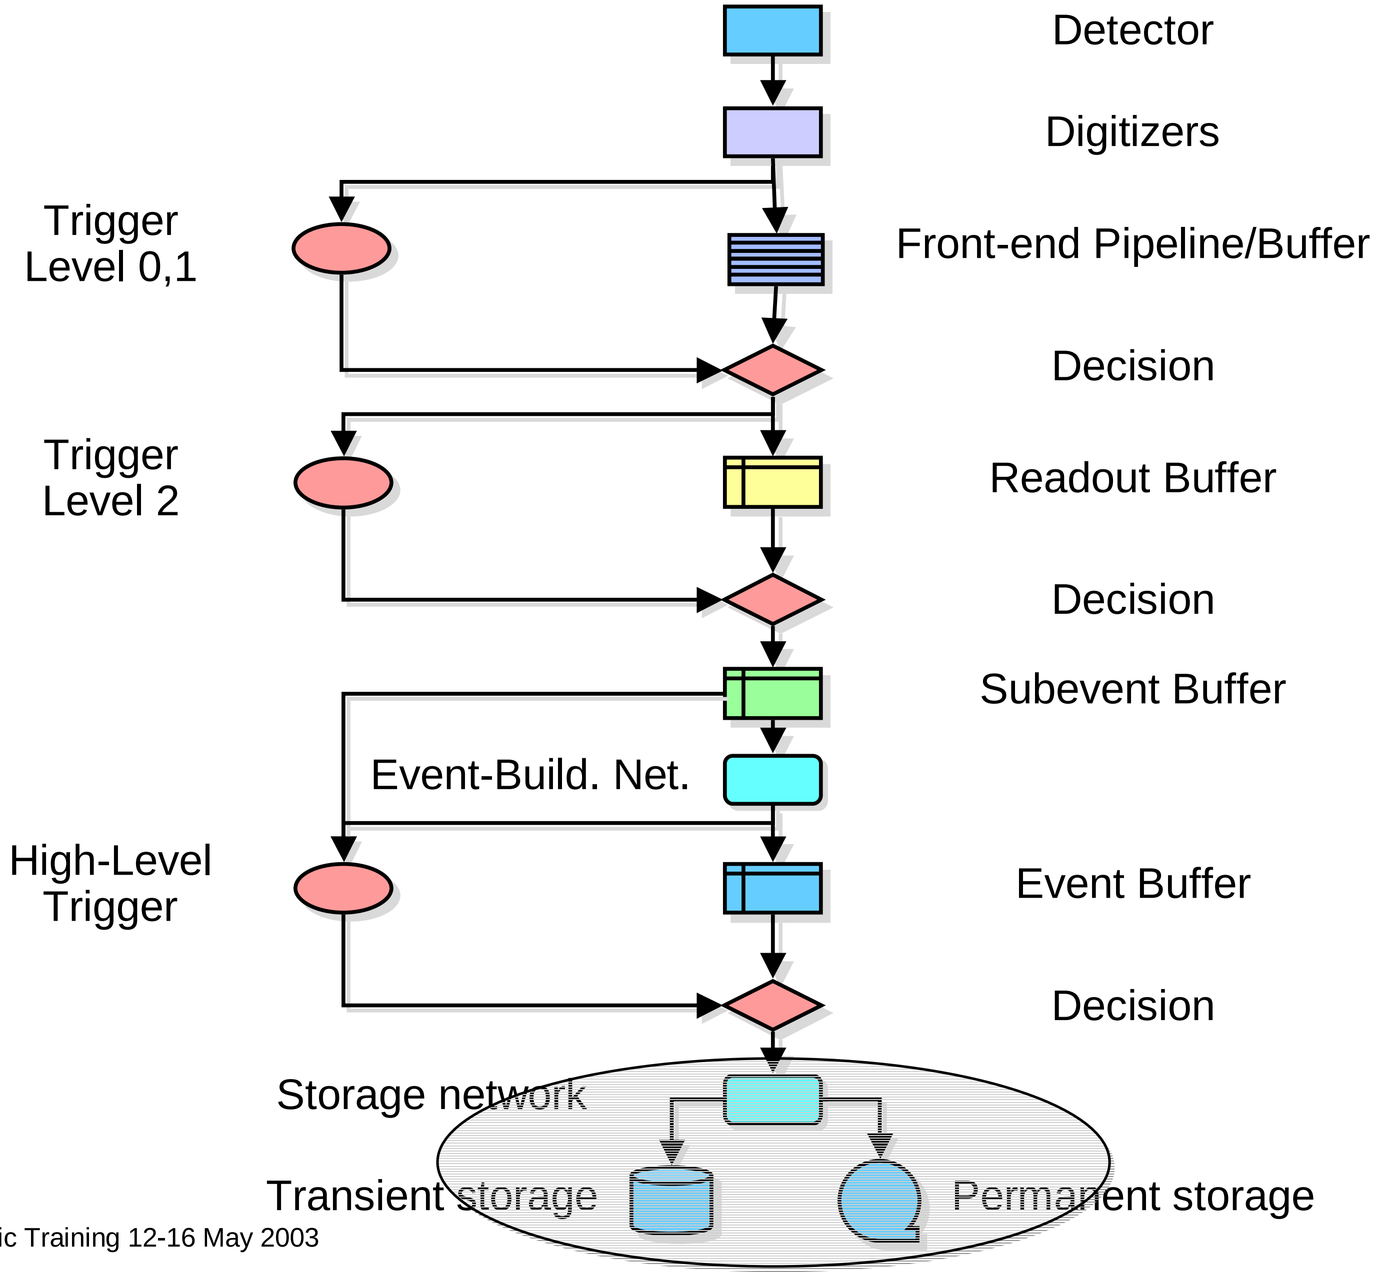

Detector
Digitizers
Trigger
Level 0,1
Front-end Pipeline/Buffer
Decision
Trigger
Level 2
Readout Buffer
Decision
Subevent Buffer
Event-Build. Net.
High-Level
Trigger
Event Buffer
Decision
Storage network
Transient storage
Permanent storage

Computing Data Challenges

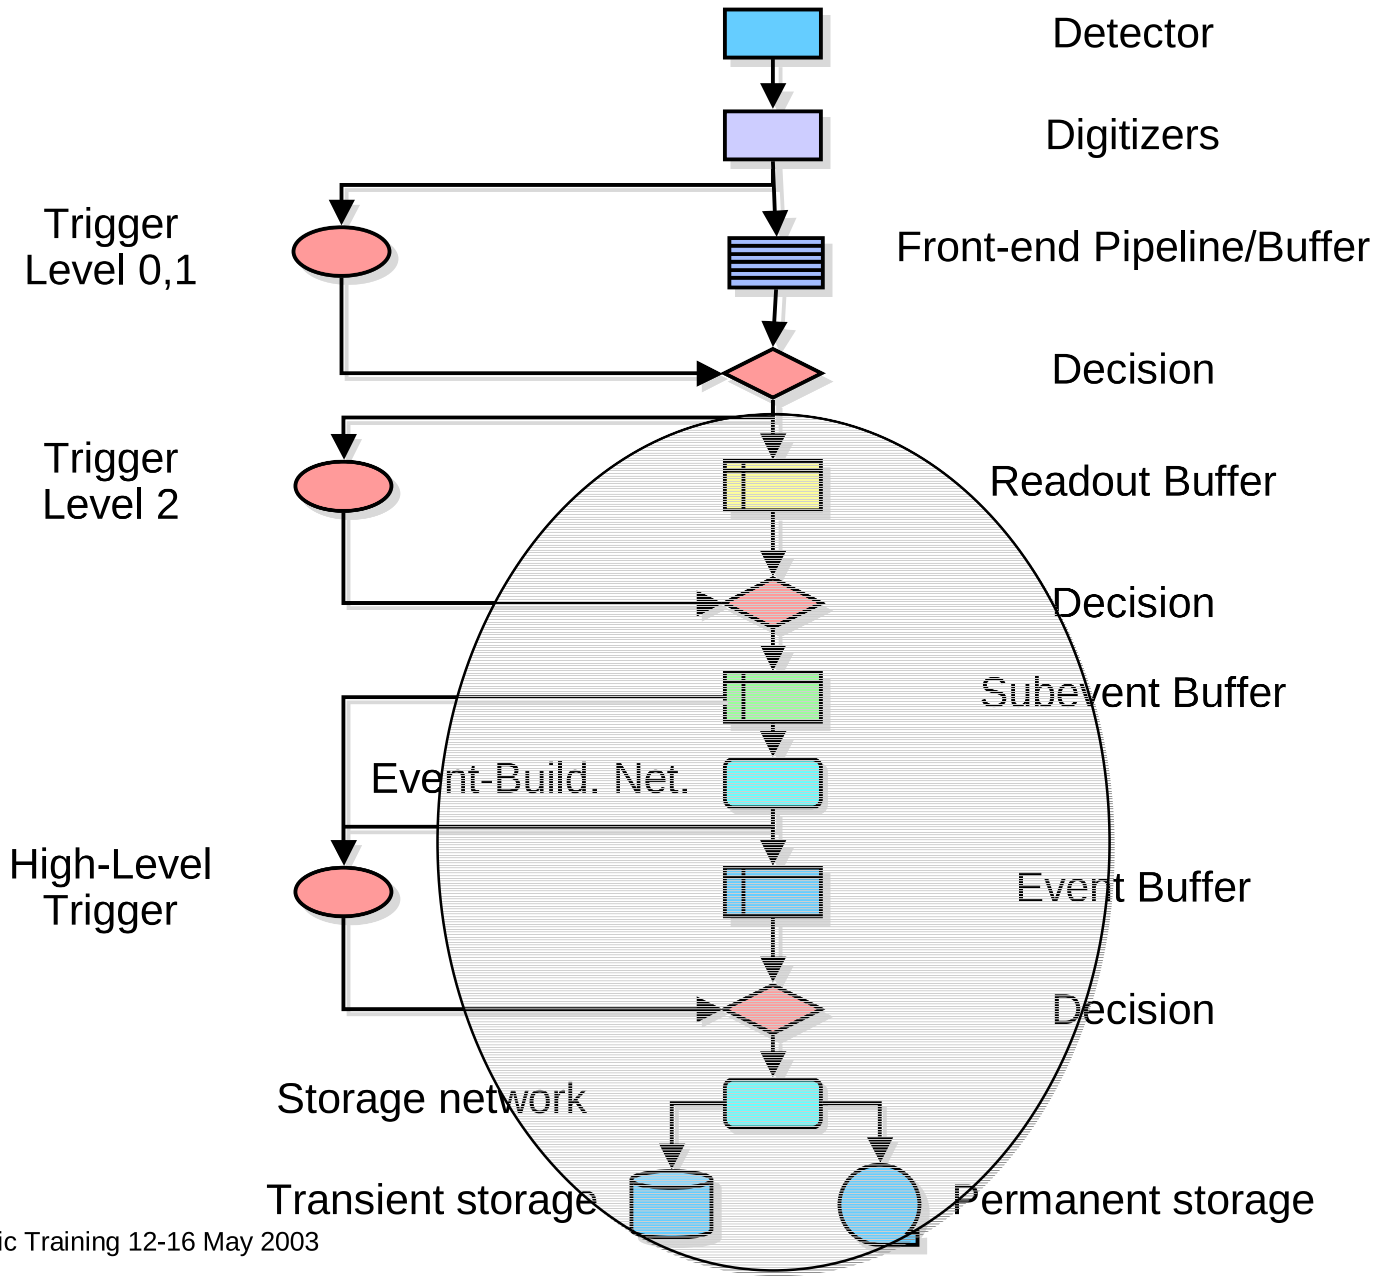

Computing Data Challenges

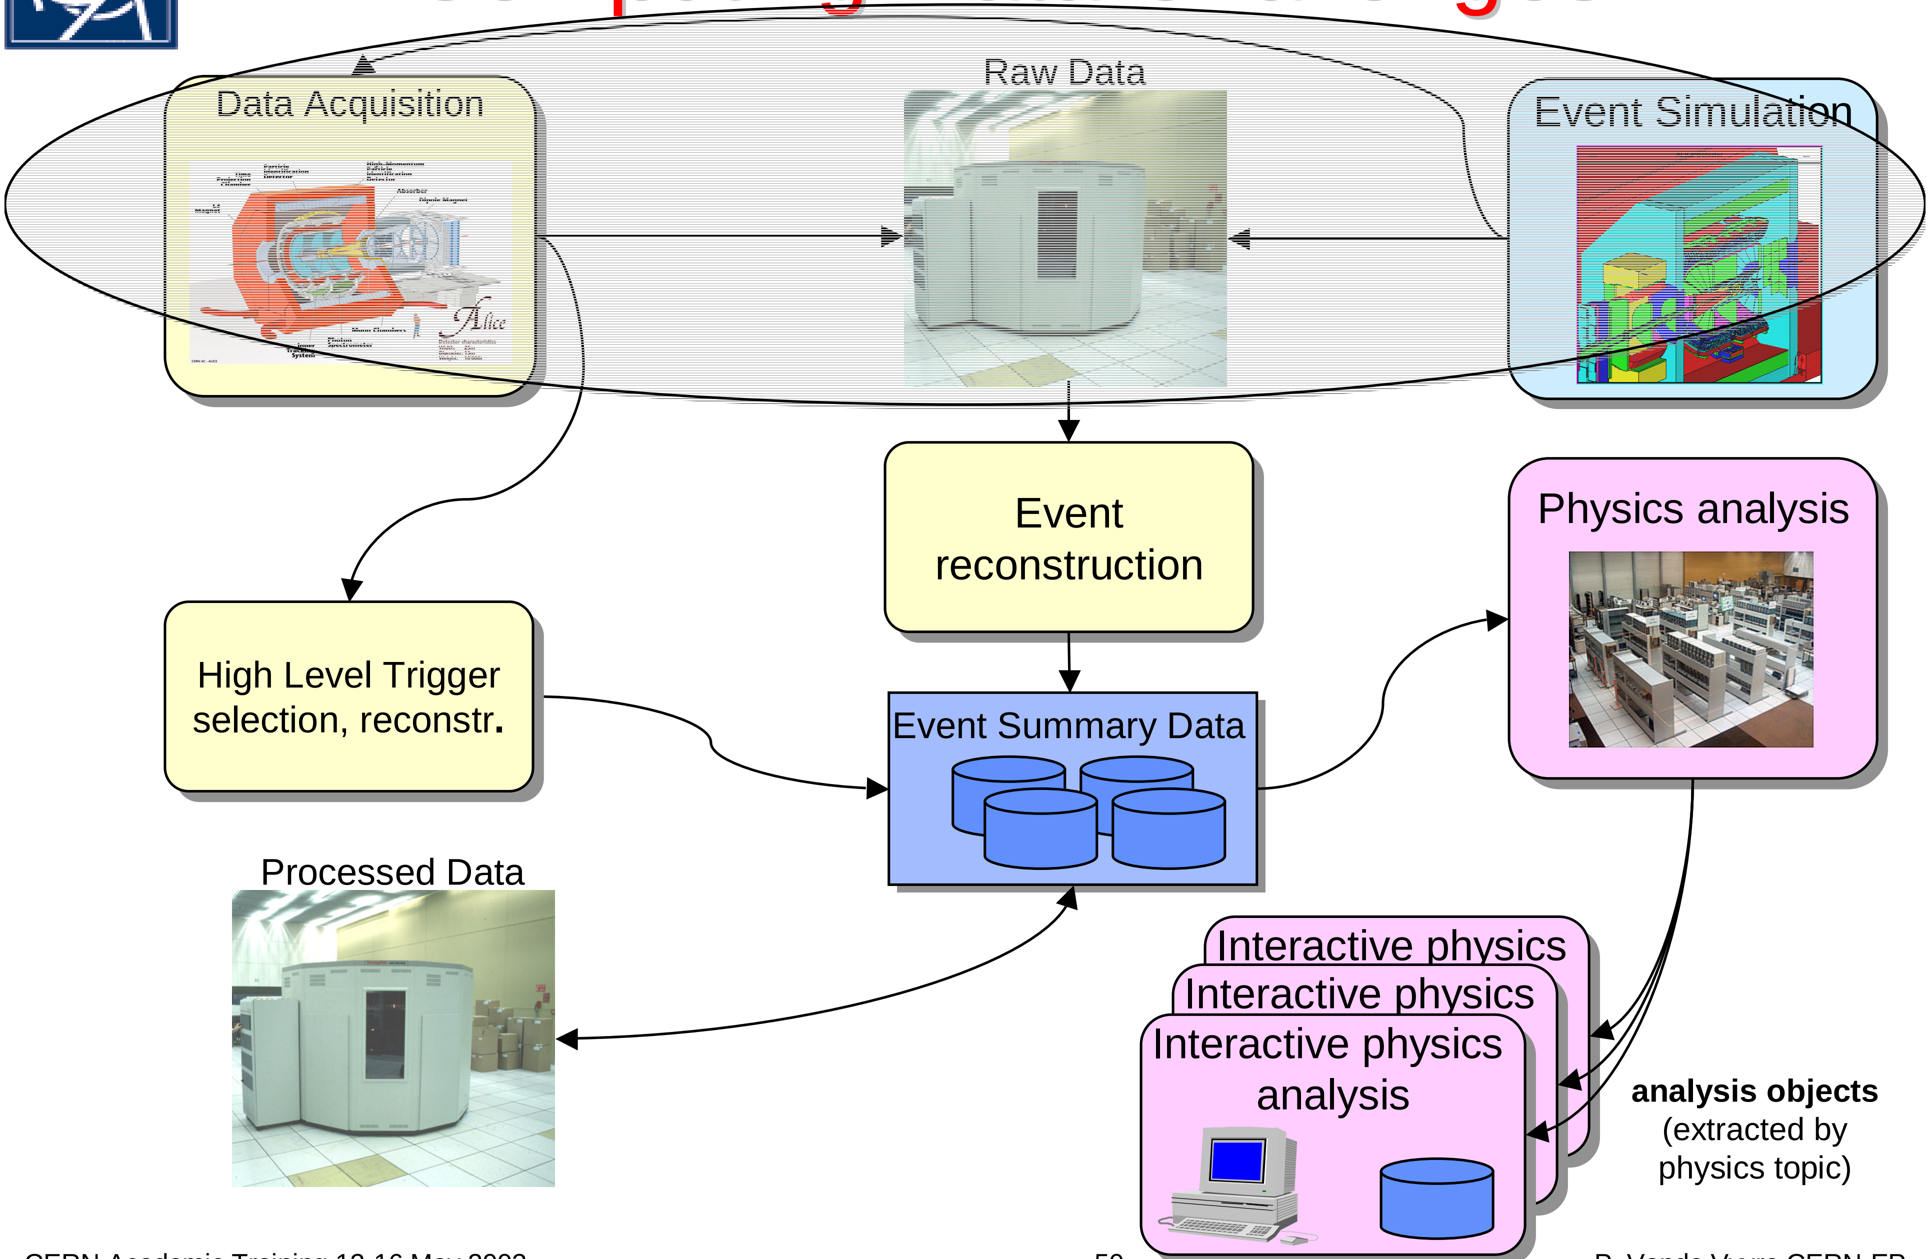

Conclusions

- The LHC experiments constitute a challenge for data acquisition, processing, and analysis.
- Many years of R&D
  - Recommendations
  - Prototypes of components or subsystems have been developed
- LHC data acquisition and computing will massively use commodity components
  - Moore’s law
  - Adequate performances of commodity products
- Combined large-scale tests in “Data Challenges”

Tomorrow

- Day 1 (Pierre VANDE VYVRE)
  - Outline, main concepts
  - Requirements of LHC experiments
  - Data Challenges
- Day 2 (Bernd PANZER)
  - Computing infrastructure
  - Technology trends
- Day 3 (Pierre VANDE VYVRE)
  - Data acquisition
- Day 4 (Fons RADEMAKERS)
  - Simulation, Reconstruction and analysis
- Day 5 (Bernd PANZER)
  - Computing Data challenges
  - Physics Data Challenges
  - Evolution