Introduction to HEP computing

John Apostolakis & Witek Pokorski
EP/SFT
CERN

Outline

- modern High Energy Physics (HEP) experiments
- physics software
  - online processing - triggering, selection
  - offline processing – reconstruction, simulation
- conclusions

HEP experiments - past

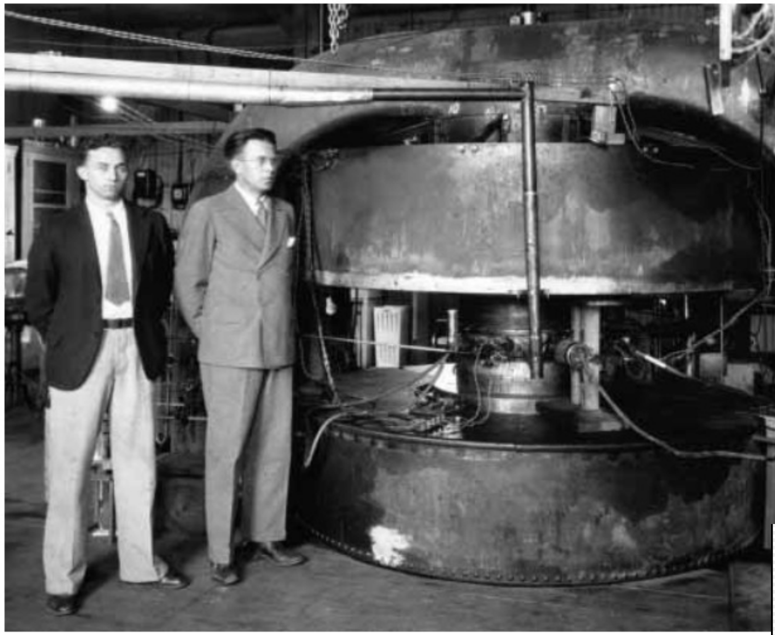

S. Livingstone and E. Lawrence
69cm cyclotron

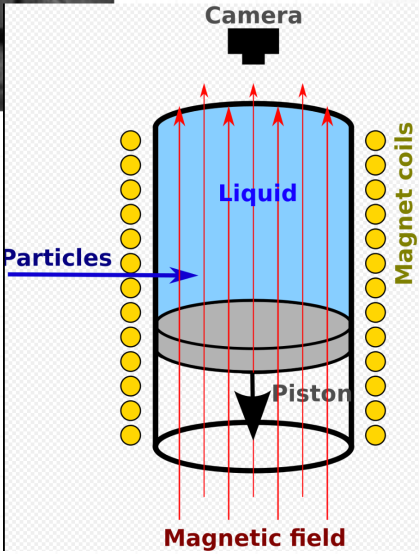

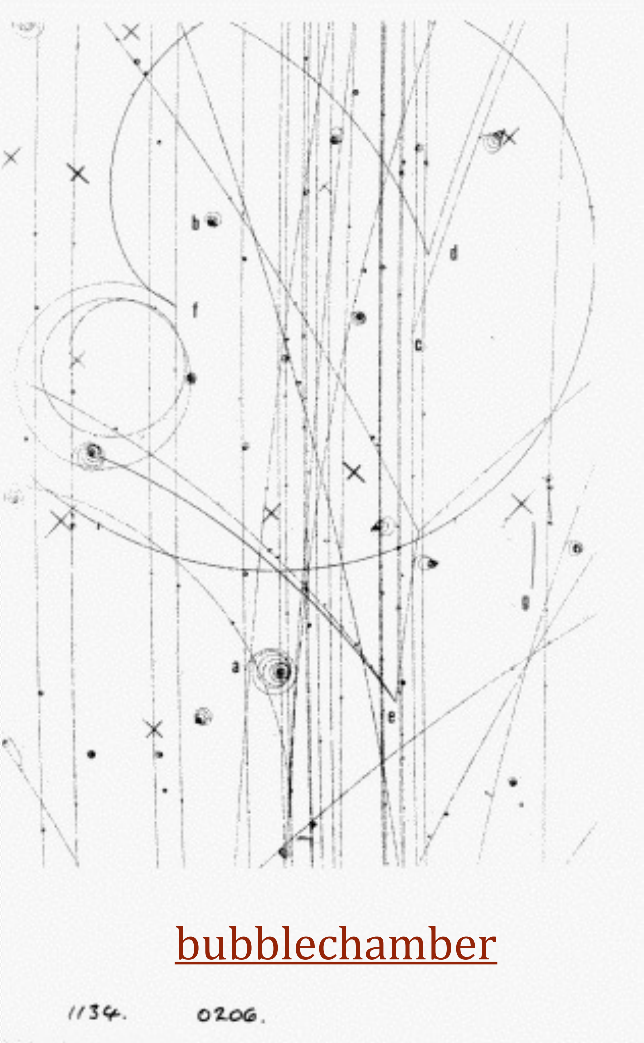

bubblechamber

HEP experiments - present

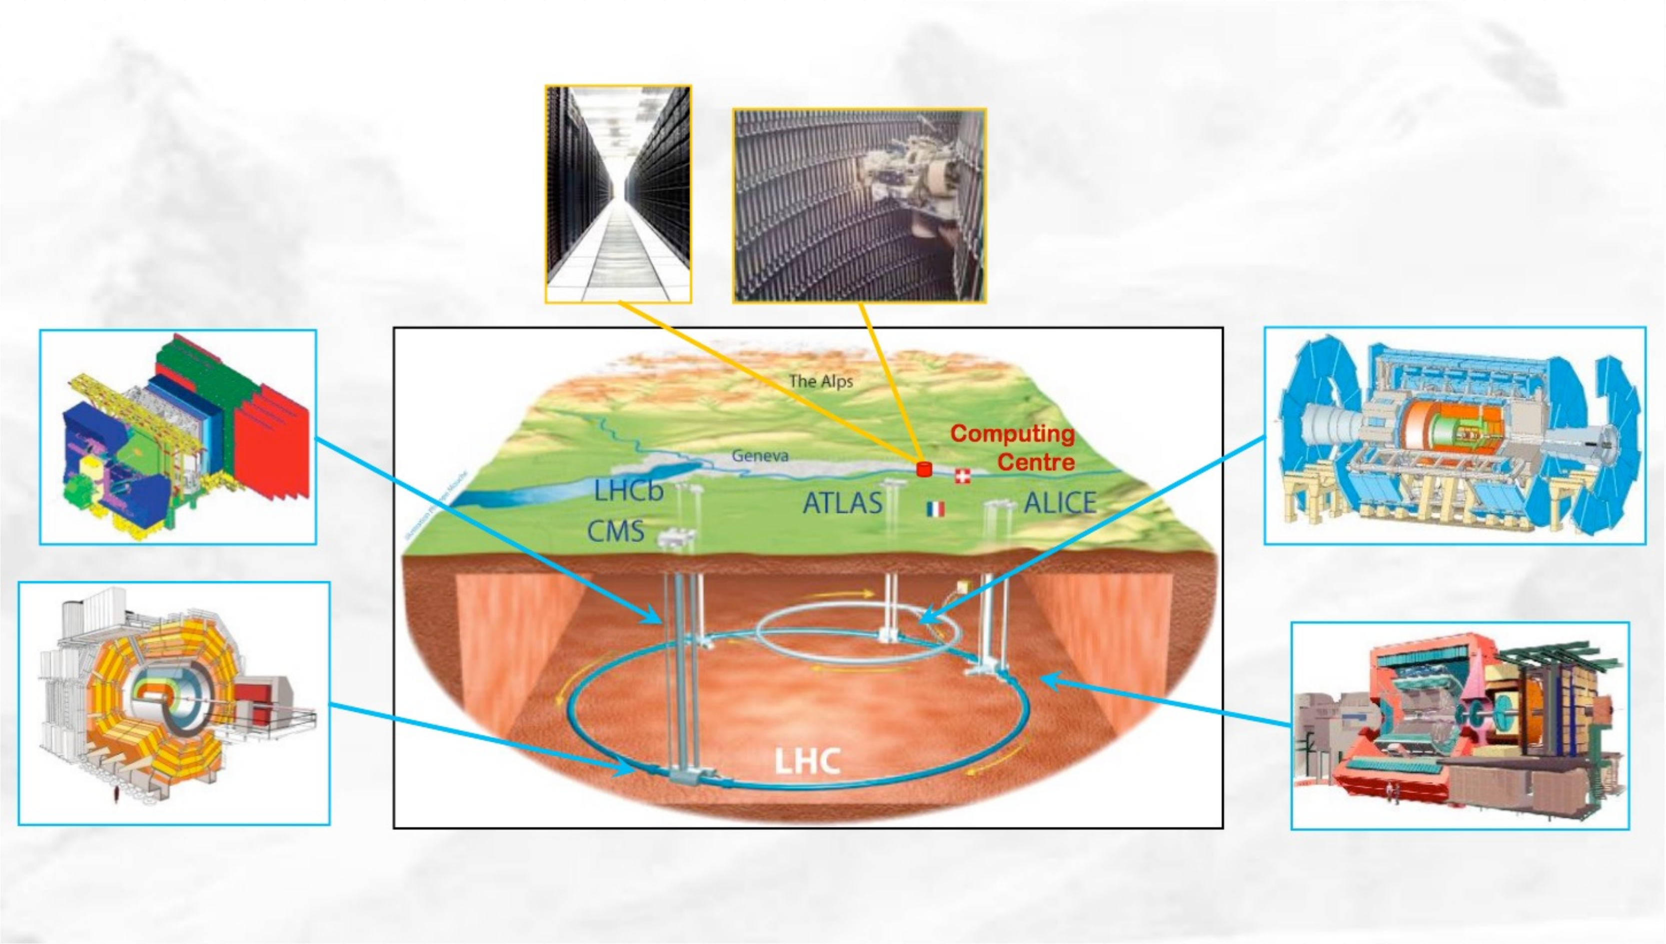

HEP experiments - present

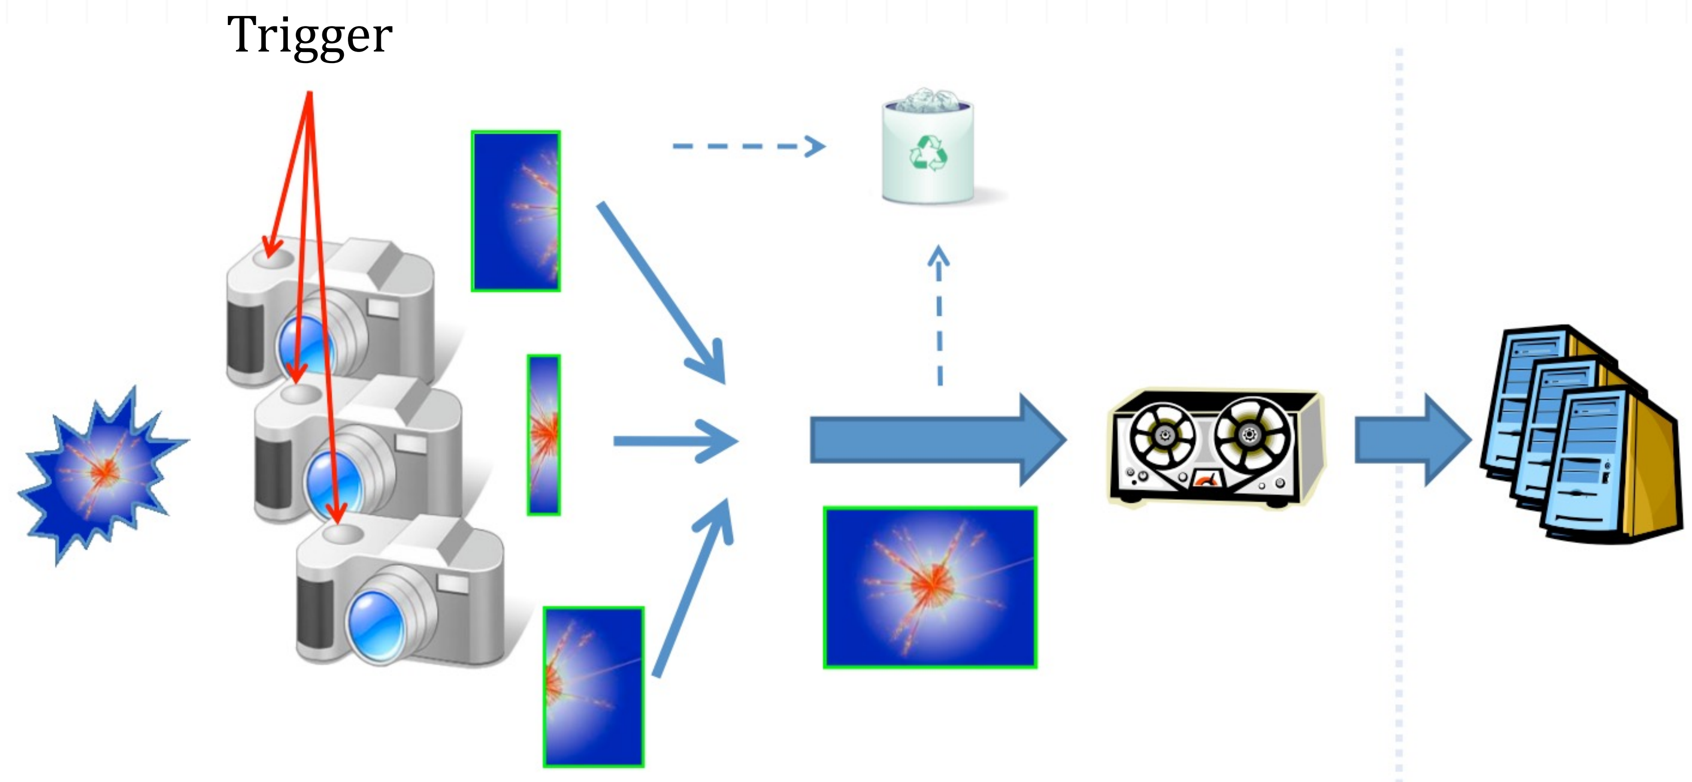

Trigger

Collision

Detectors

Event fragments

Full event

Storage

Offline analysis

Collisions at the LHC

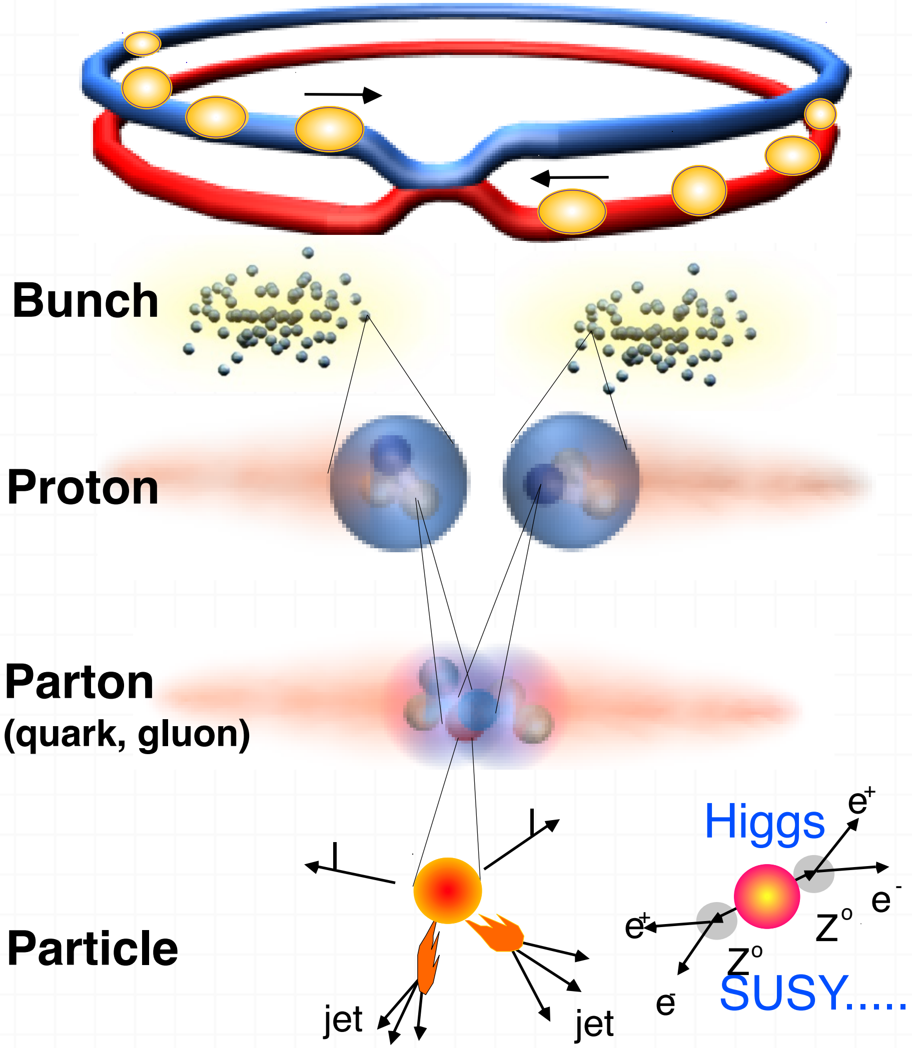

Proton - Proton	**2804 bunch/beam**
Protons/bunch	**10^{11}**
Beam energy	**7 TeV (7×10^{12} eV)**
Luminosity	**$10^{34} cm^{-2} s^{-1}$**
Crossing rate	**40 MHz**
Collision rate ≈	**10^7-10^9 Hz**

Physics selection at LHC

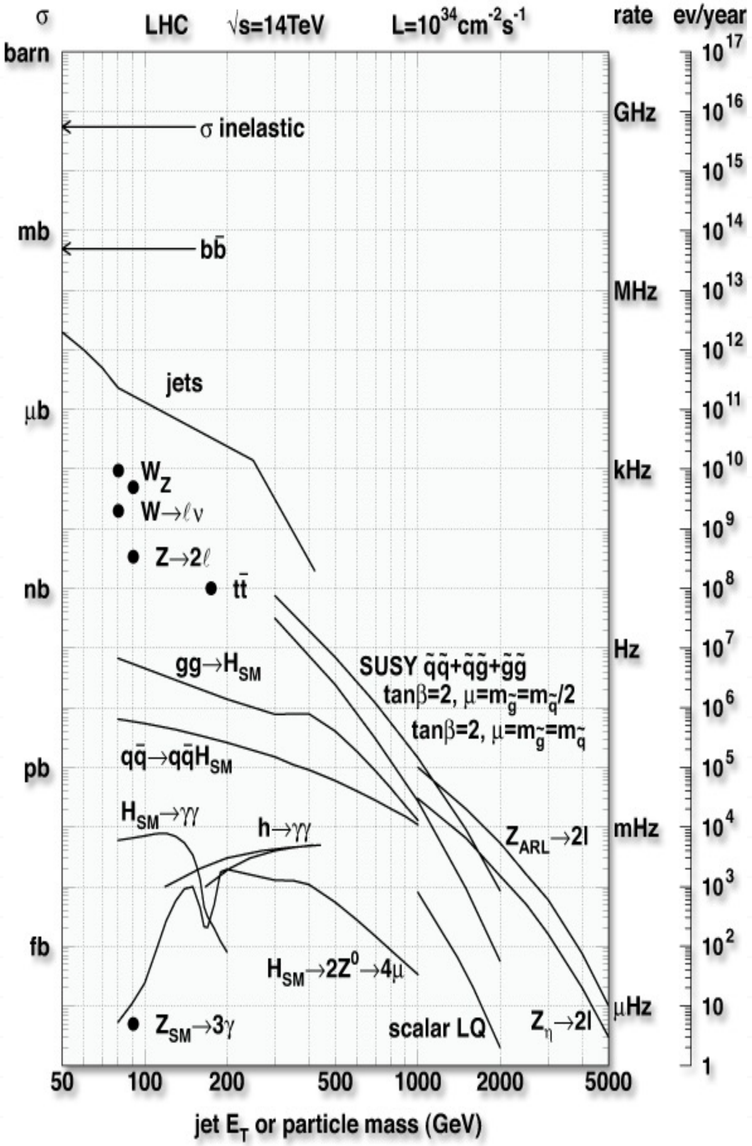

- Cross sections of physics processes vary over many orders of magnitude
  - Inelastic: 10^9 Hz
  - $W \to \ell\, \nu$: 10^2 Hz
  - t t production: 10 Hz
  - Higgs (100 GeV/c^2): 0.1 Hz
  - Higgs (600 GeV/c^2): 10^{-2} Hz
- QCD background
  - Jet E_T ~250 GeV: rate = 1 kHz
  - Jet fluctuations → electron bkg
  - Decays of K, π, b → muon bkg
- Selection needed: $1:10^{10-11}$

Physics Selection at LHC

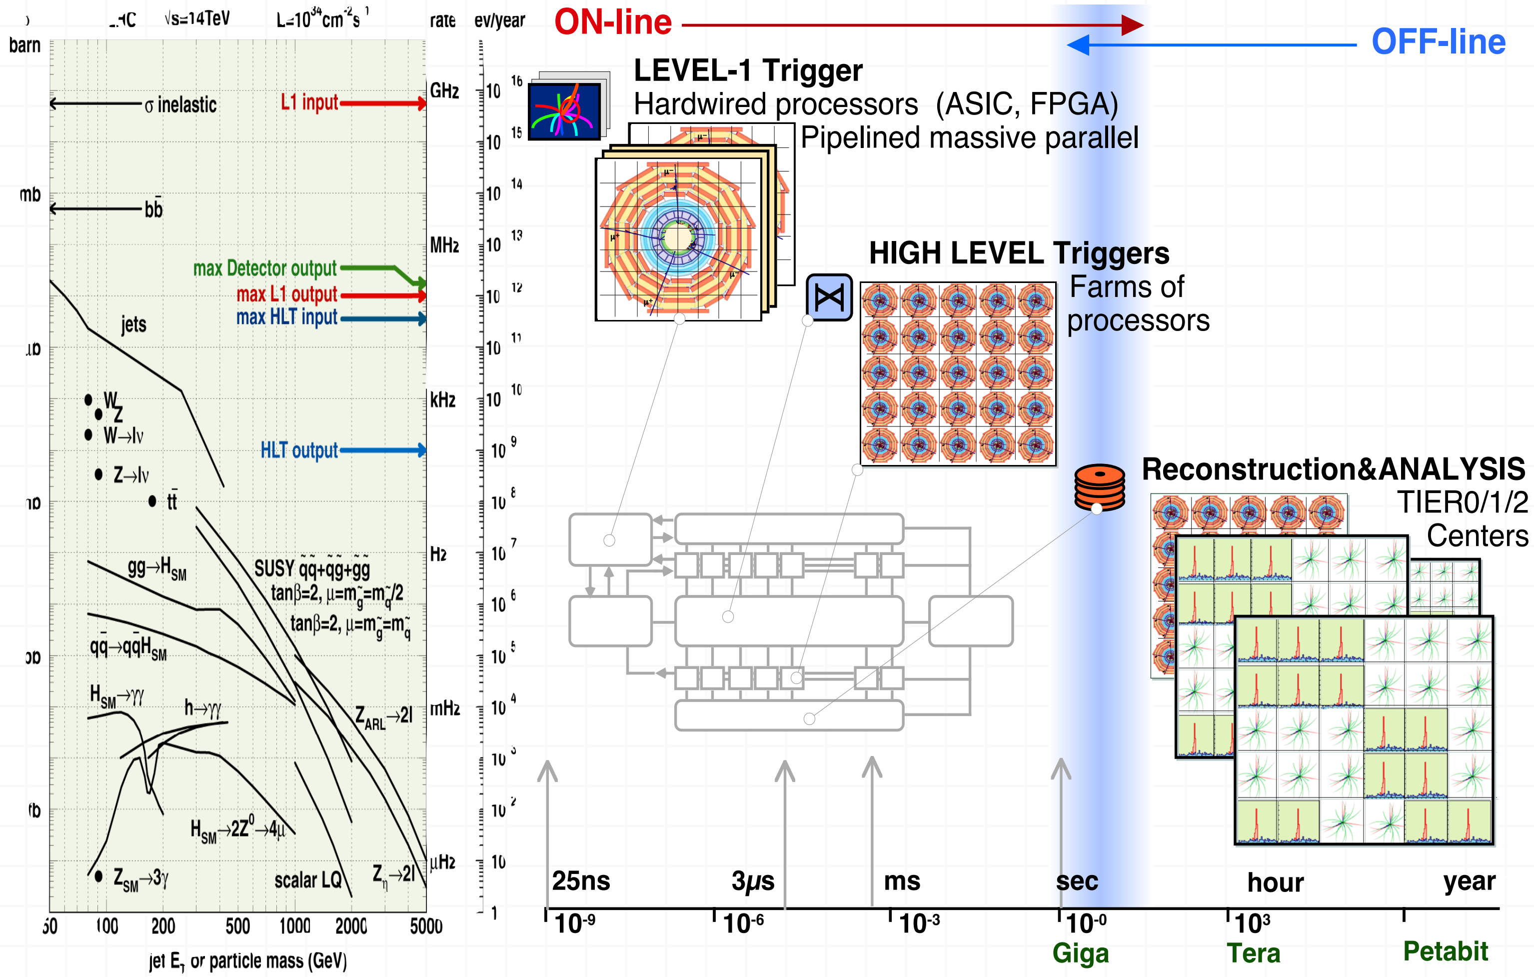

ON-line

OFF-line

LEVEL-1 Trigger

Hardwired processors (ASIC, FPGA)

Pipelined massive parallel

HIGH LEVEL Triggers

Farms of processors

Reconstruction&ANALYSIS

TIER0/1/2 Centers

25ns **3µs** **ms** **sec** **hour** **year**

10^{-9} **10^{-6}** **10^{-3}** **10^{-0}** **10^{3}**

Giga **Tera** **Petabit**

Processing Stages

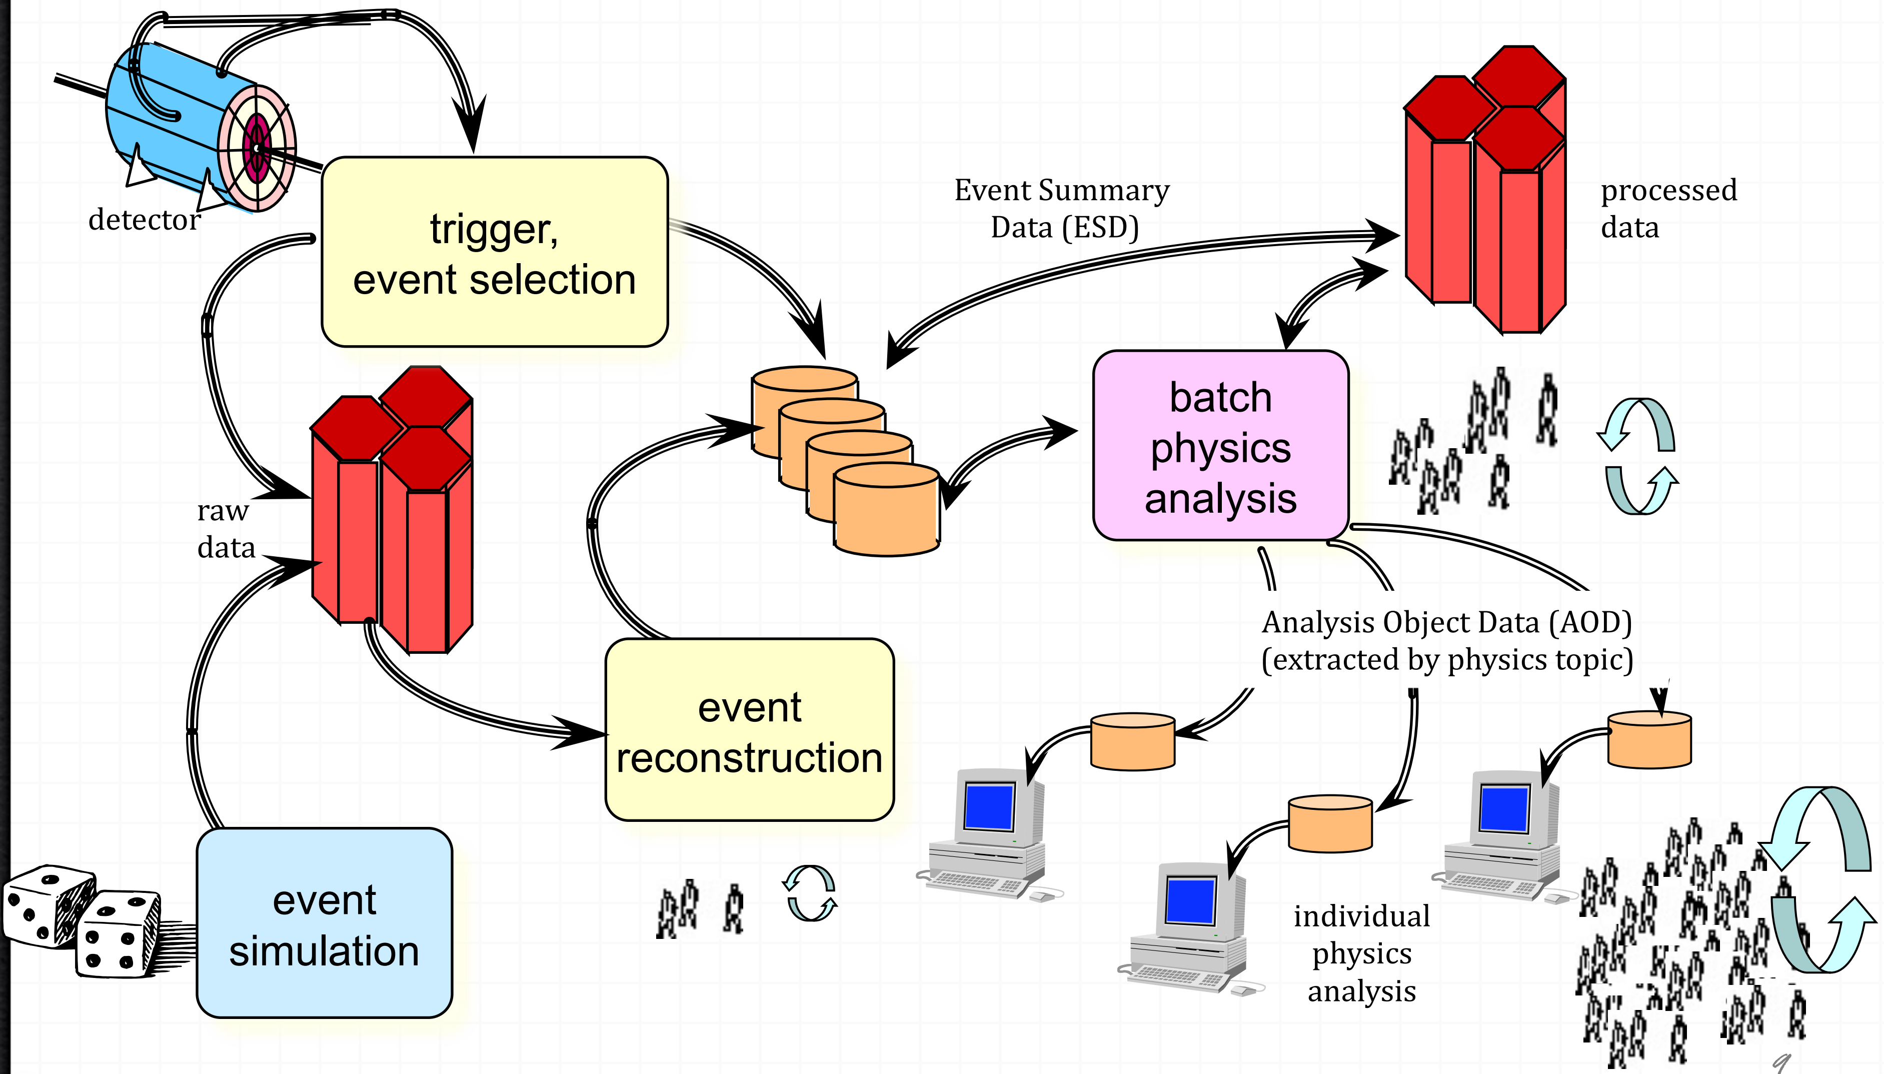

detector

trigger,
event selection

raw
data

event
simulation

event
reconstruction

Event Summary
Data (ESD)

processed
data

batch
physics
analysis

Analysis Object Data (AOD)
(extracted by physics topic)

individual
physics
analysis

Physics software

- The scientific software needed to process this huge amount of data from the LHC detectors is developed by the LHC collaborations
  - Must cope with the unprecedented conditions and challenges (trigger rate, data volumes, etc.)
  - Each collaboration has written millions of lines of code
- Modern technologies and methods
  - Object-oriented programming languages and frameworks
  - Re-use of a number of generic and domain-specific ‘open-source’ packages
- The organization of this large software production activity is by itself a huge challenge
  - Large number of developers distributed worldwide
  - Integration and validation require large efforts

Processing Stages - Trigger

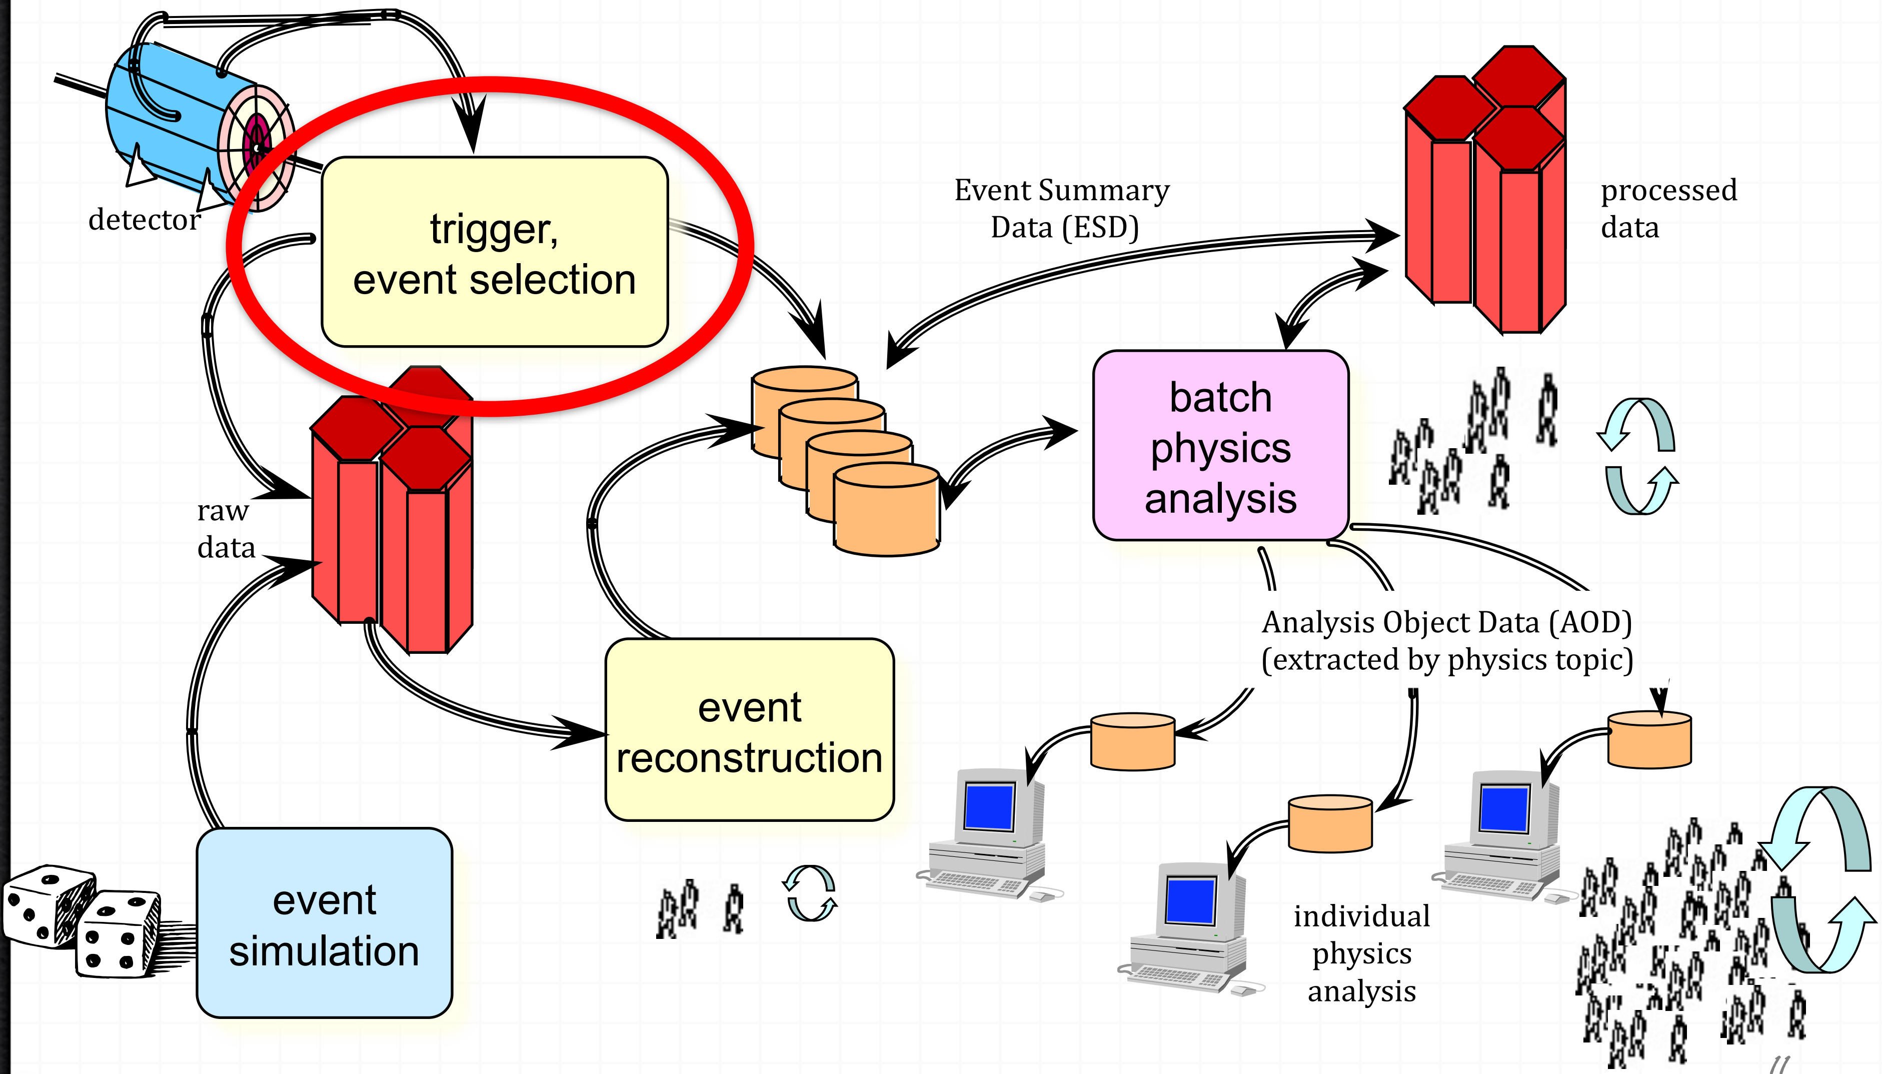

detector

trigger,
event selection

Event Summary
Data (ESD)

processed
data

batch
physics
analysis

raw
data

event
reconstruction

Analysis Object Data (AOD)
(extracted by physics topic)

event
simulation

individual
physics
analysis

Trigger

- Task: inspect detector information and provide a first decision on whether to keep the event or throw it out
- The trigger is a function of event data, detector conditions and parameters

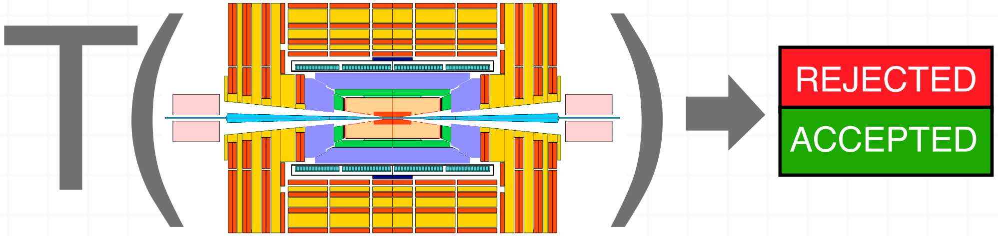

REJECTED

ACCEPTED

- Detector data not (all) promptly available
- Selection function highly complex

$\Rightarrow$T(...) is evaluated by successive approximations

TRIGGER LEVELS

Trigger Levels

- Level-1
  - Hardwired processors (ASIC, FPGA, …)
  - Pipelined massive parallel
  - Partial information, quick and simple event characteristics (pt, total energy, etc.)
  - 3-4 μs maximum latency

~ $1:10^4$

- Level-2 (optional)
  - Specialized processors using partial data

~ $1:10^1$

- High Level
  - Software running in processor farms
  - Complex algorithms using complete event information
  - Latency at the level of fractions of second
  - Output rate adjusted to what can be afforded

~ $1:10^2$

Trigger Levels and Rates

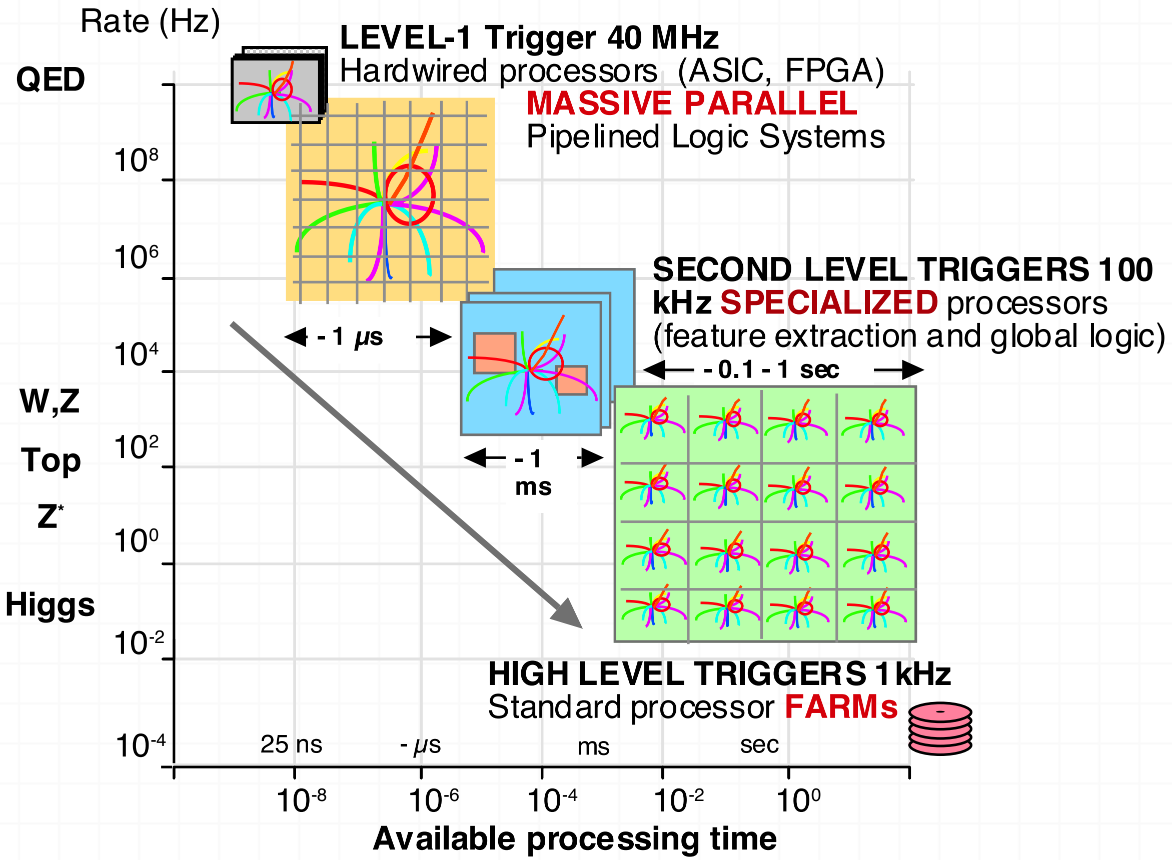
LEVEL-1 Trigger 40 MHz
Hardwired processors (ASIC, FPGA)
MASSIVE PARALLEL
Pipelined Logic Systems

SECOND LEVEL TRIGGERS 100 kHz SPECIALIZED processors
(feature extraction and global logic)

HIGH LEVEL TRIGGERS 1 kHz
Standard processor **FARMs**

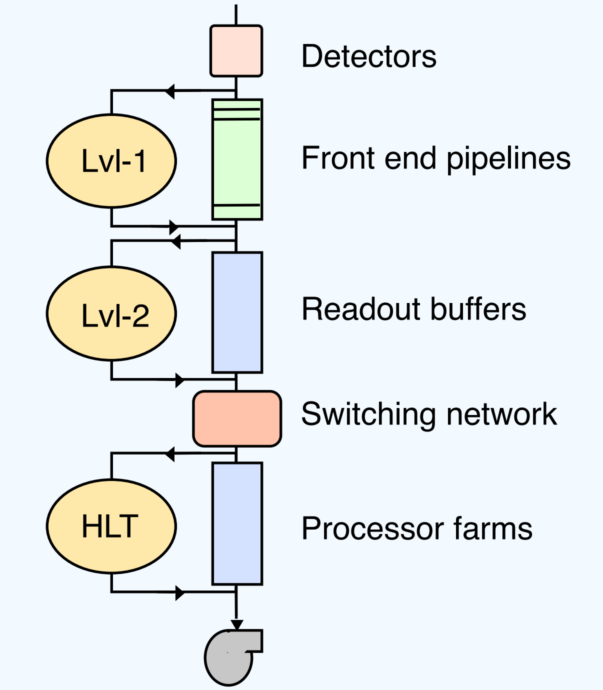

"Traditional": 3 physical levels

Processing Stages - Reconstruction

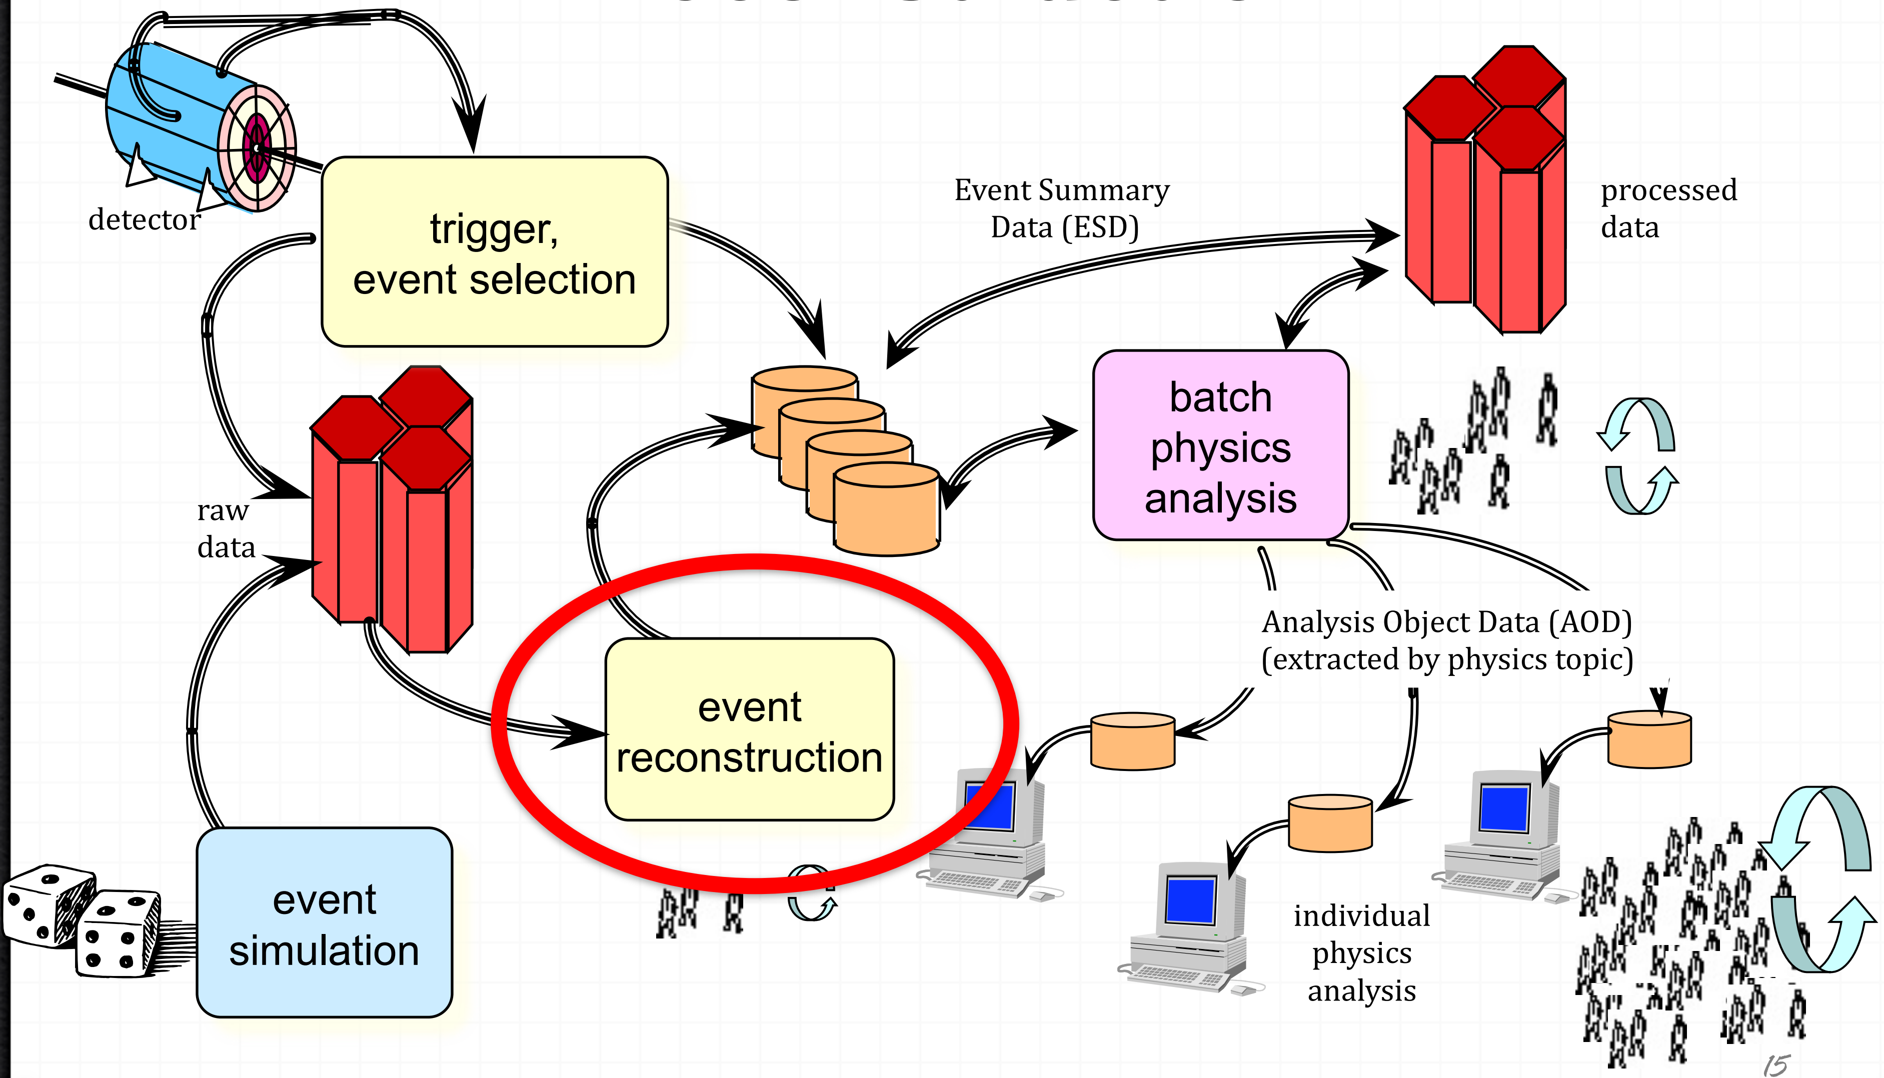

The Reconstruction challenge

Starting from this event

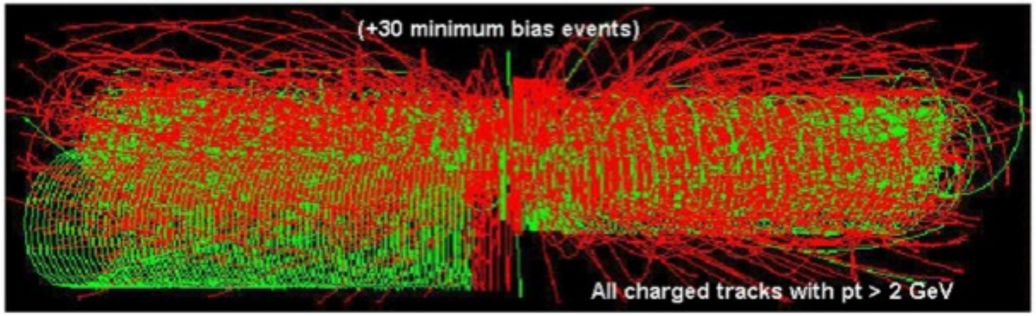

Looking for this “signature”

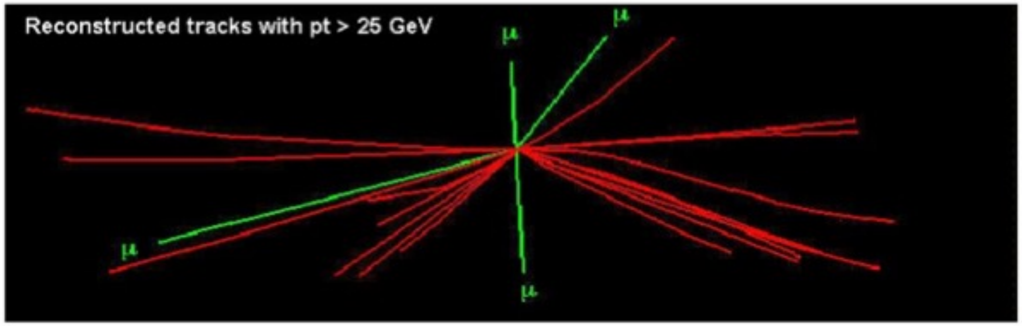

→ **Selectivity: 1 in 10^{13}**

(Like looking for a needle in 20 million haystacks)

What is reconstruction

- Tracker ‘hits’ form a puzzle
  - Which tracks created them?
- Each energy deposition is a clue
  - There are thousands of measurements in each snap-shot
- The experiment’s reconstruction must obtain a solution!
  - In well measured magnetic field
  - Matches the traces to tracks

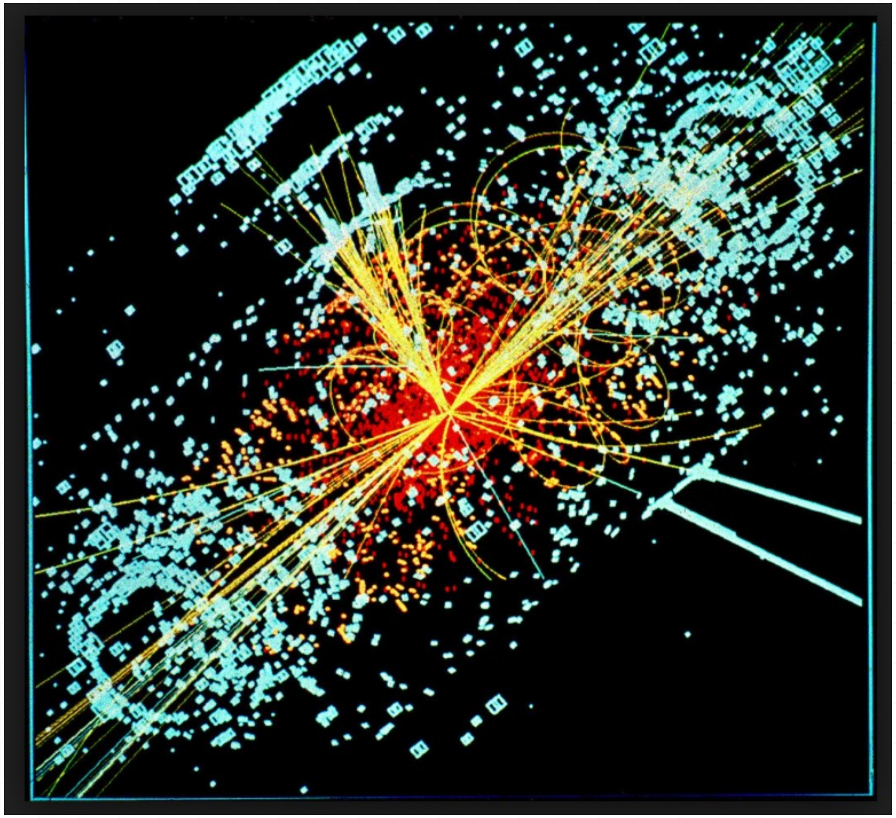

How it works – a simple example

- Start with the locations of the traces on first two planes

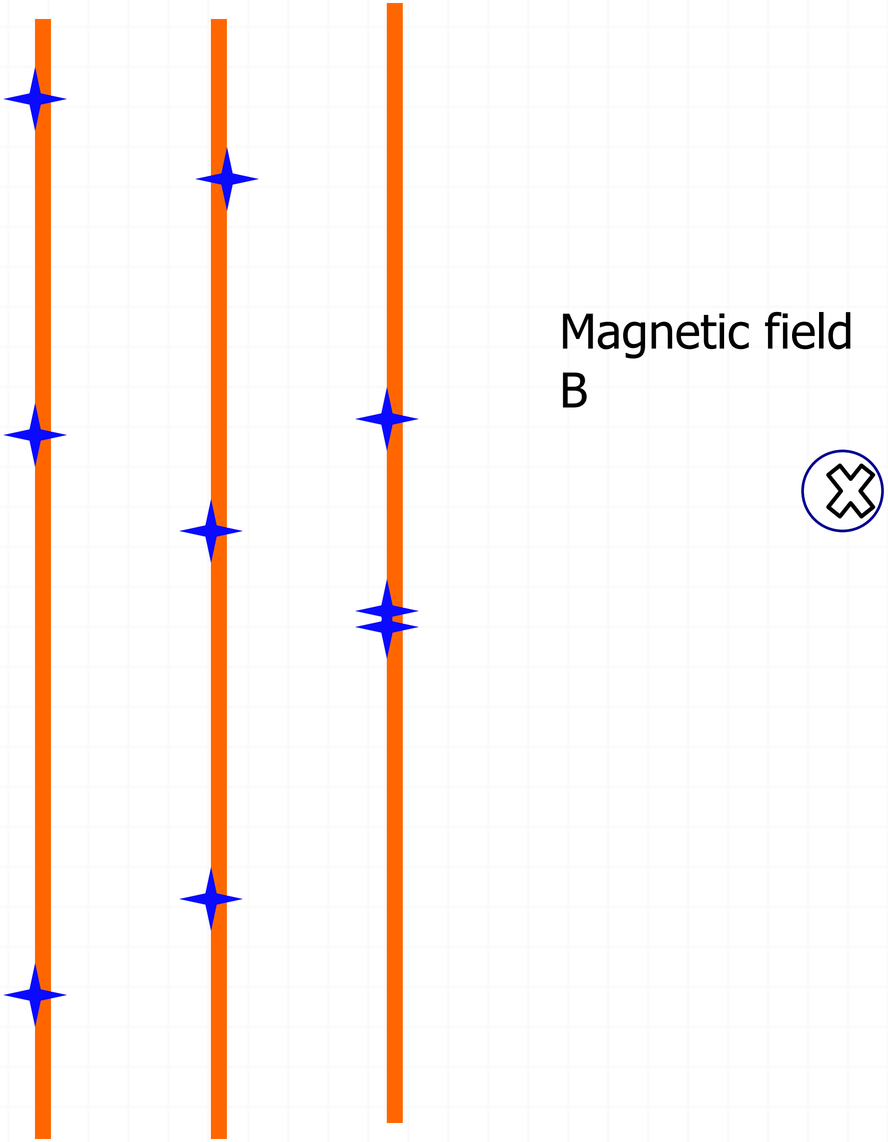

How it works – a simple example

- Start with the locations of the traces on first two planes
- Try different combinations
  - Project to subsequent planes
  - Calculate differences between measured positions and ‘predictions’

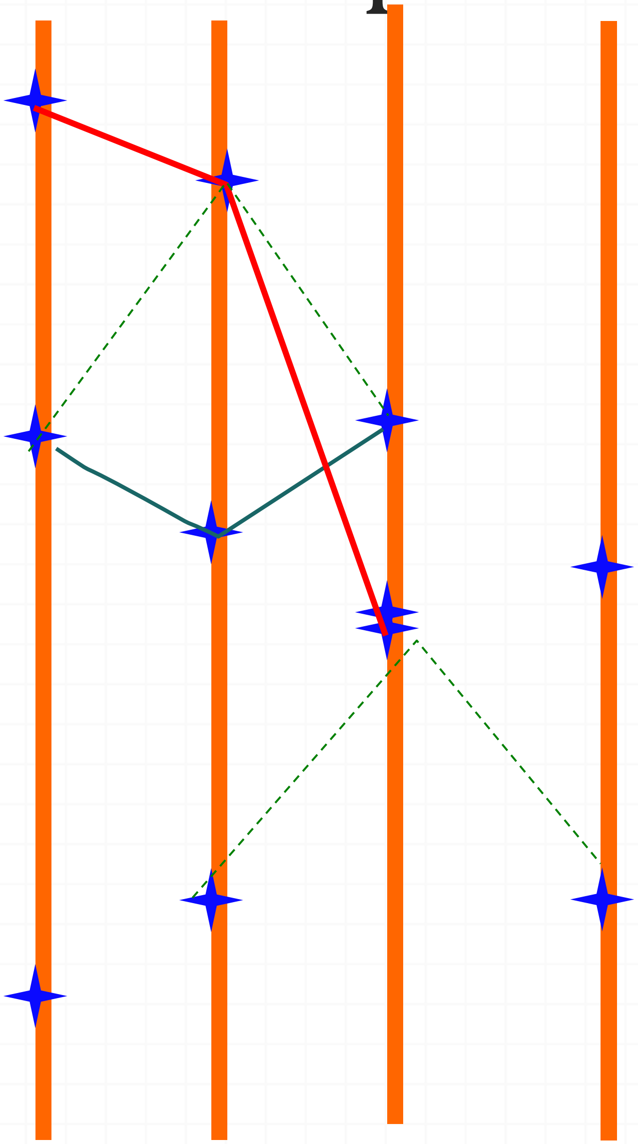

(Kalman filter)

How it works – a simple example

- Start with the locations of the traces on first two planes
- Try different combinations
  - Project to subsequent planes
  - Calculate differences between measured positions and 'predictions'
- Finally the candidate tracks are identified
  - else look 'quickly' for the straight(er) ones – high energy tracks

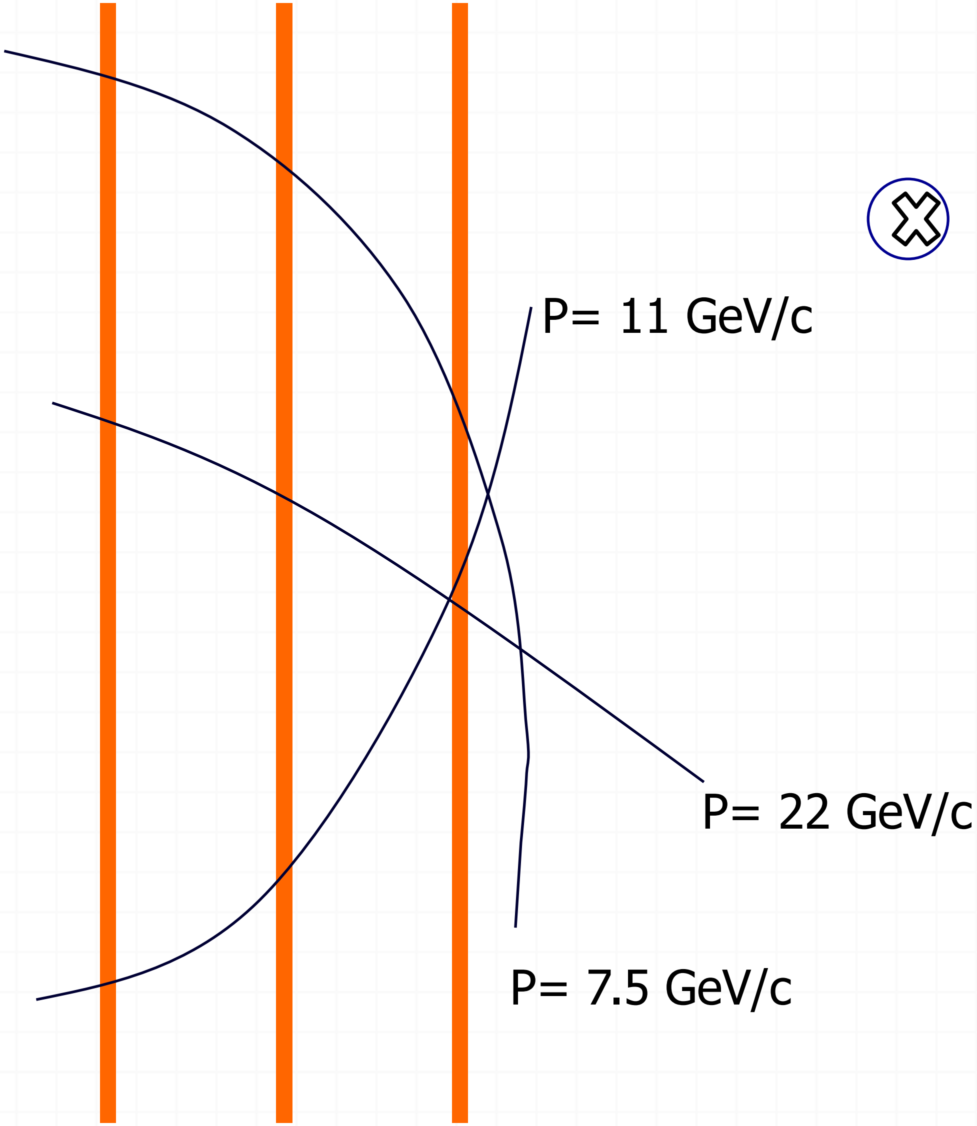

Detector conditions data

- Reflects changes in state of the detector with time
- Event Data cannot be reconstructed or analyzed without it
- Versioning
- Tagging
- Ability to extract slices of data required to run with job
- Long life-time

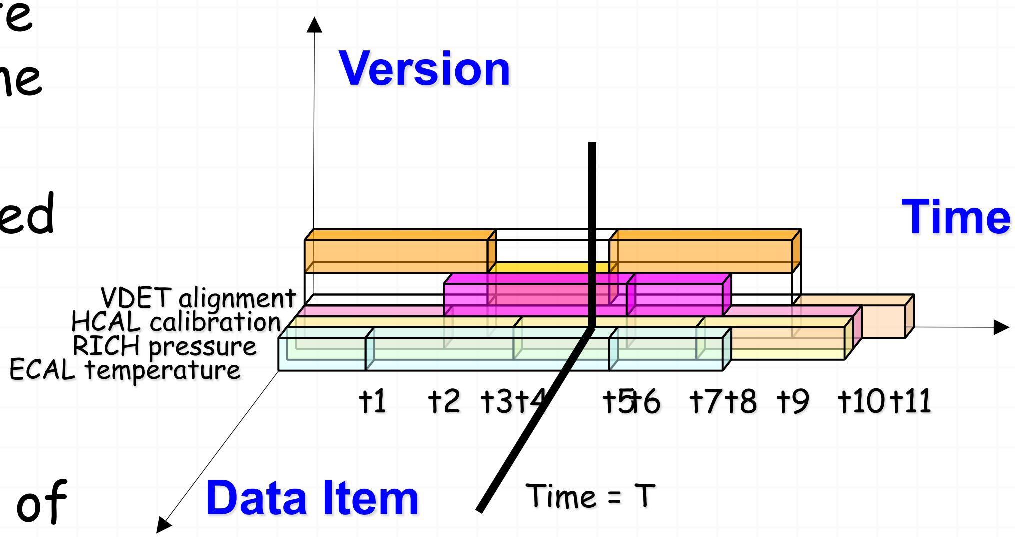

Online and offline reconstruction

- Are collisions first-tagged really interesting enough to keep (given capacity constraints)?
  - Online reconstruction – seek to reconstruct ‘as much as you can’ quickly to enable decision
- Critical part of experiment – collisions which are not recorded are lost
- Later there is more time to reconstruct the contents of a collision – but this is also complex

Processing Stages - Analysis

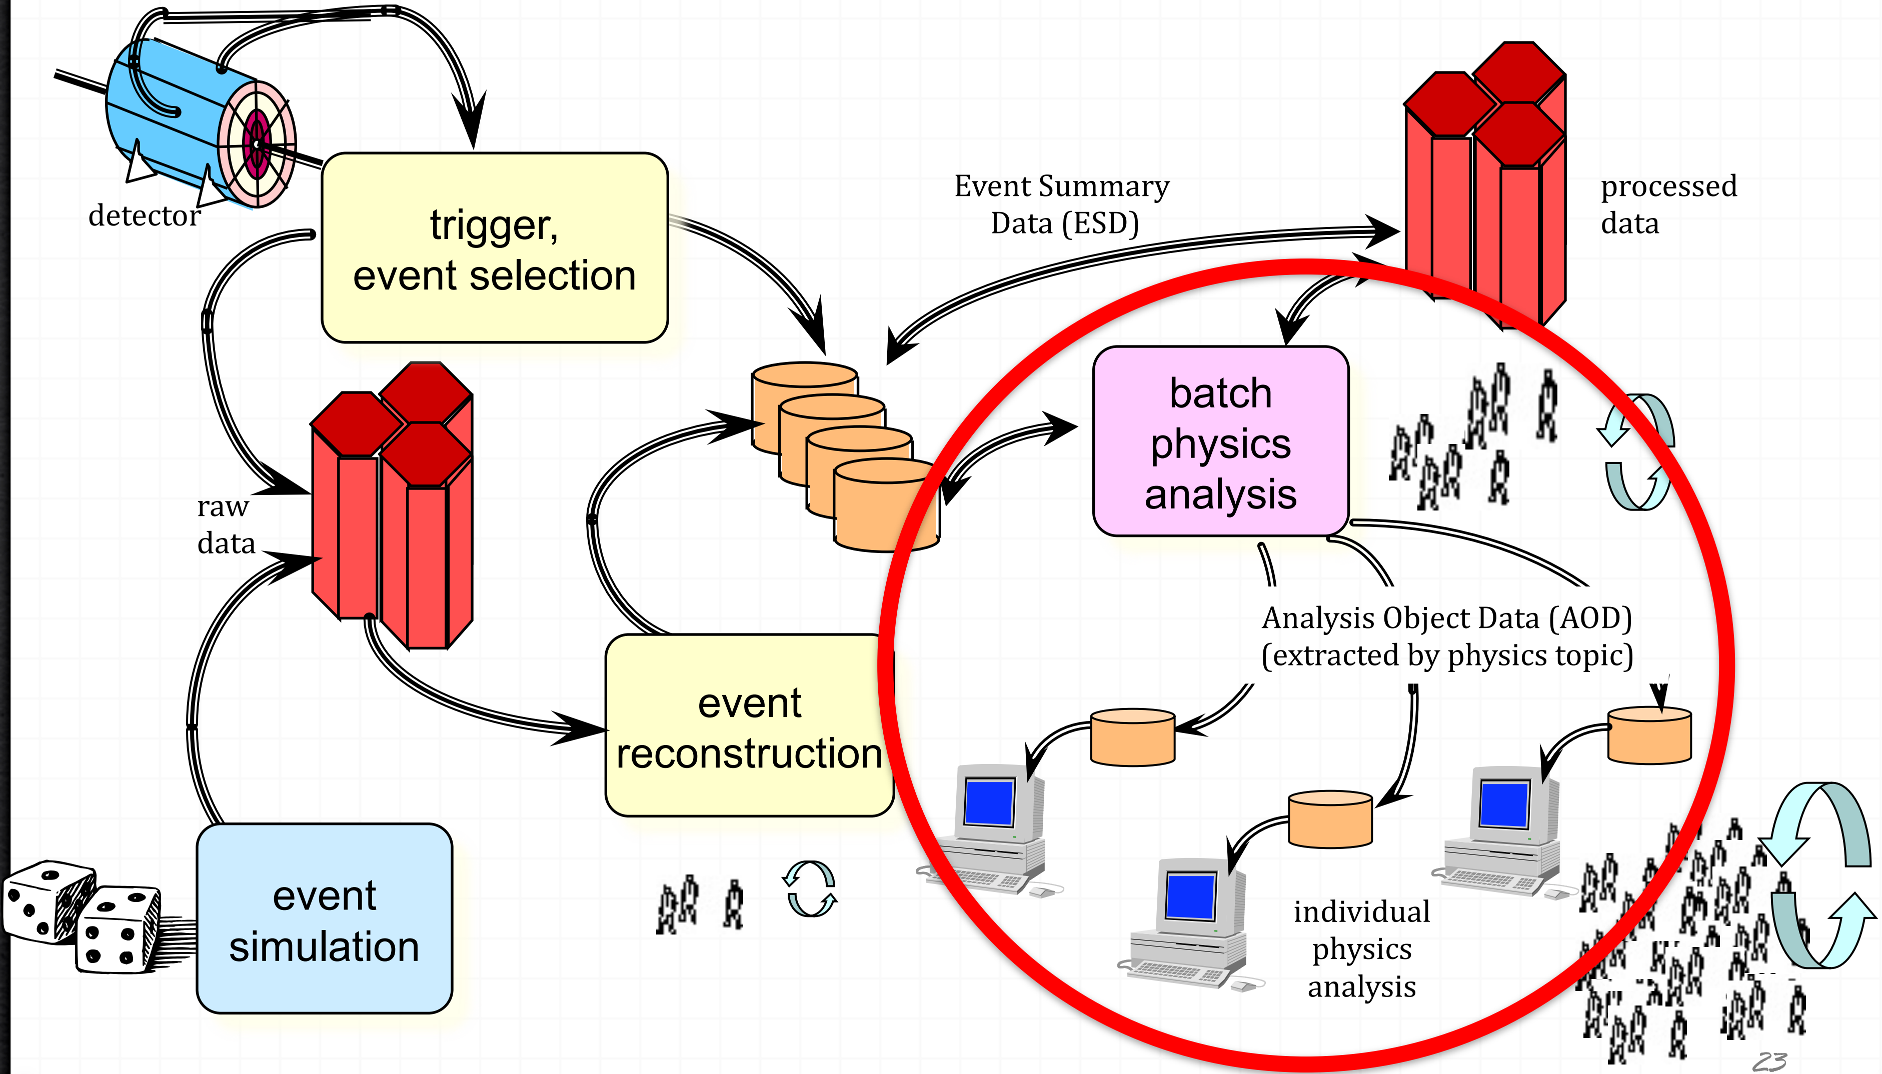

Data analysis

- Uses the results of Reconstruction
  - the products are reconstructed tracks, Energy deposits (calorimeters)
  - Hierarchy of data from original (RAW), to summary (AOD)
- An experiment's physics teams use the (large) pool of data
  - No longer in one central location, but in multiple locations (cost, space of building, computers, disks, network) .... using the GRID

Event bookkeeping

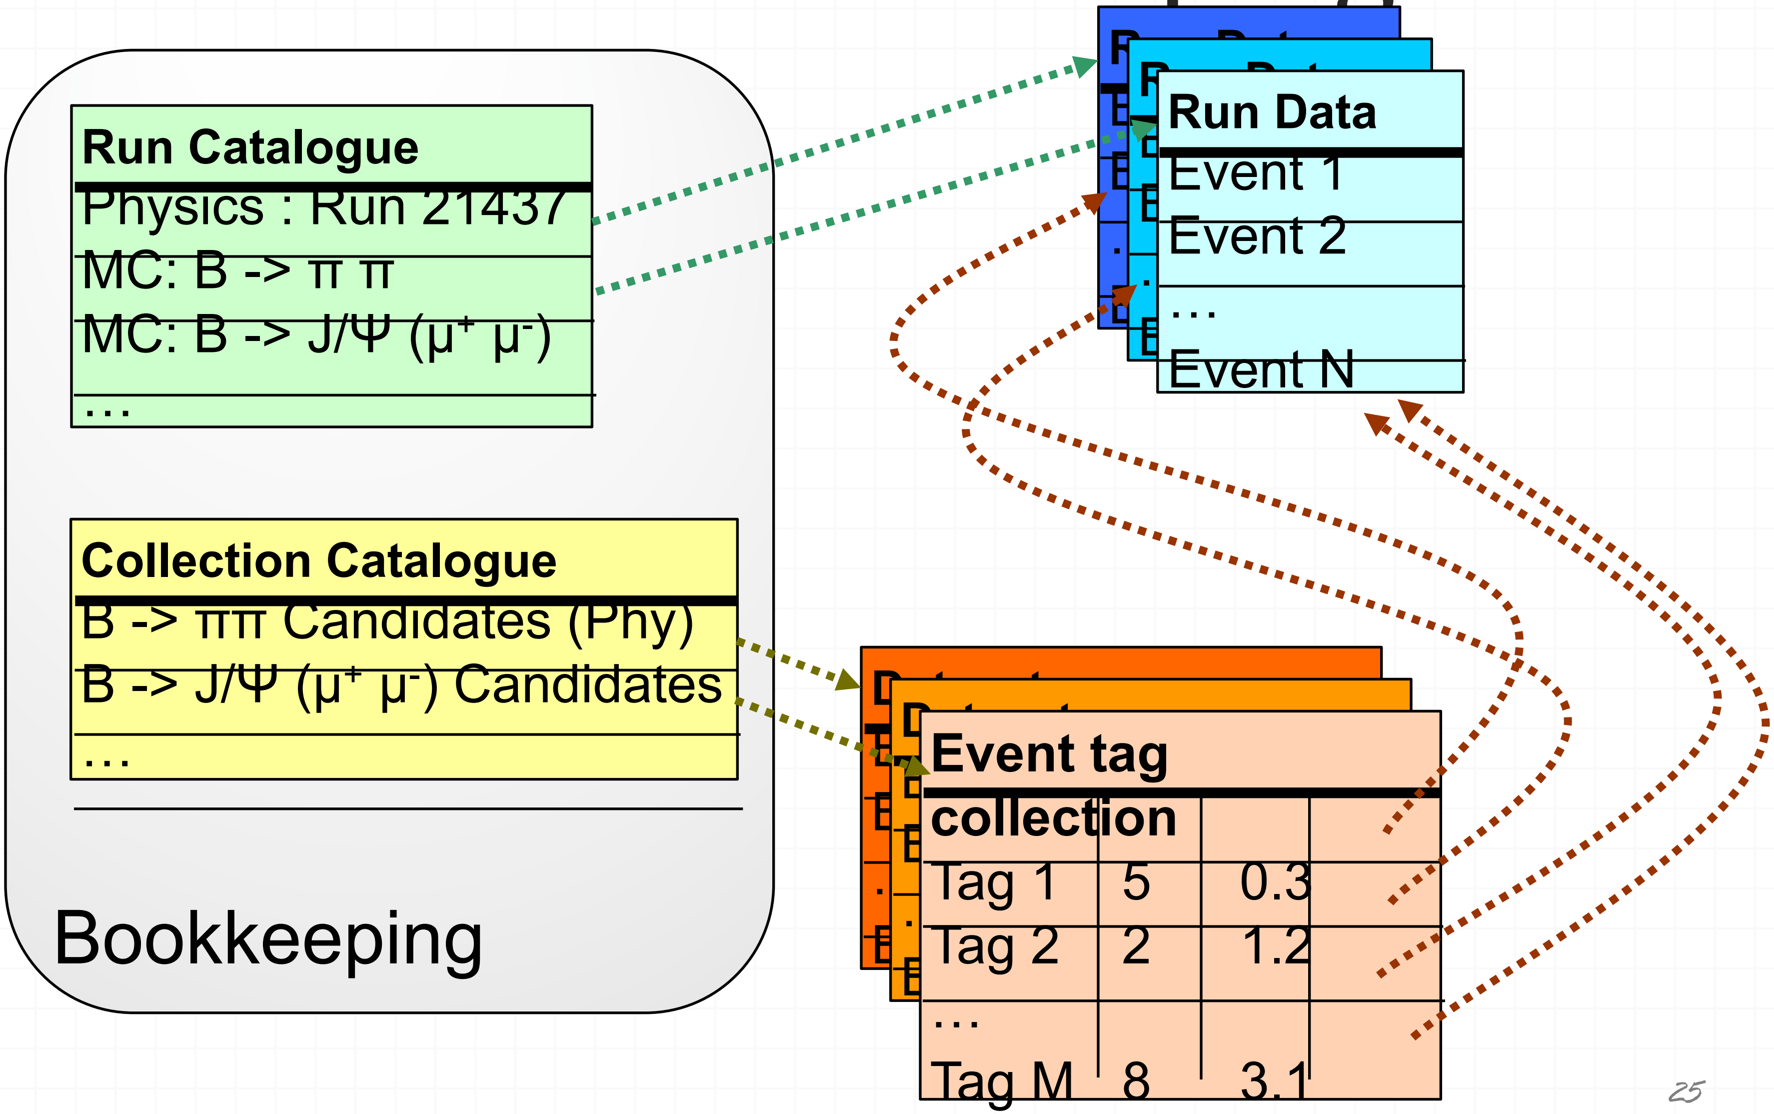

Run Catalogue

Physics : Run 21437

MC: B -> π π

MC: B -> J/Ψ ($\mu^+ \mu^-$)

…

Collection Catalogue

B -> ππ Candidates (Phy)

B -> J/Ψ ($\mu^+ \mu^-$) Candidates

…

Bookkeeping

Run Data

Event 1

Event 2

…

Event N

Event tag collection		
Tag 1	5	0.3
Tag 2	2	1.2
…		
Tag M	8	3.1

ROOT

- "At the root of the experiments", project started in 1995
- Open Source project (LGPL3)
  - mainly written in C++; 4 MLOC
- ROOT provides (amongst other thin
  - C++ interpreter
  - Efficient data storage mechanism; 1 ROOT
  - Advanced statistical analysis algorithms
    - histogramming, fitting, minimization, stat
  - Scientific visualization: 2D/3D graphics, PDF, Latex
  - Geometrical modeler
  - PROOF parallel query engine

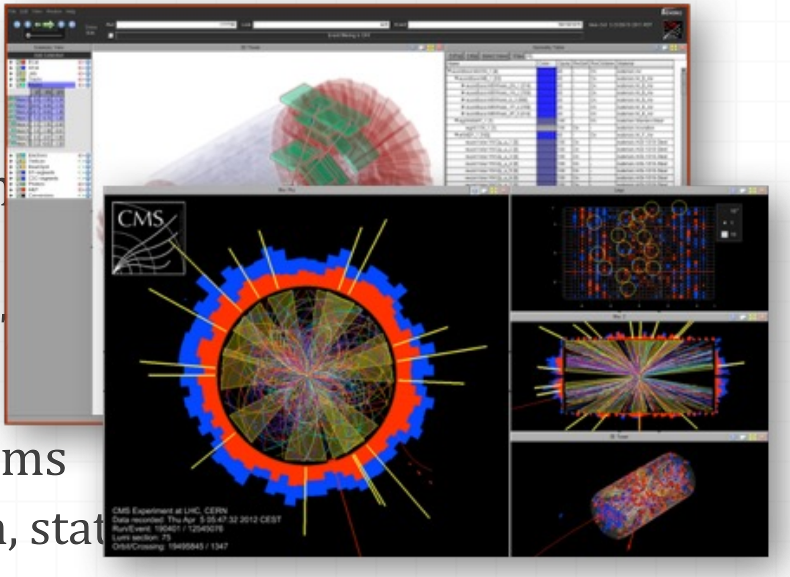

ROOT in plots

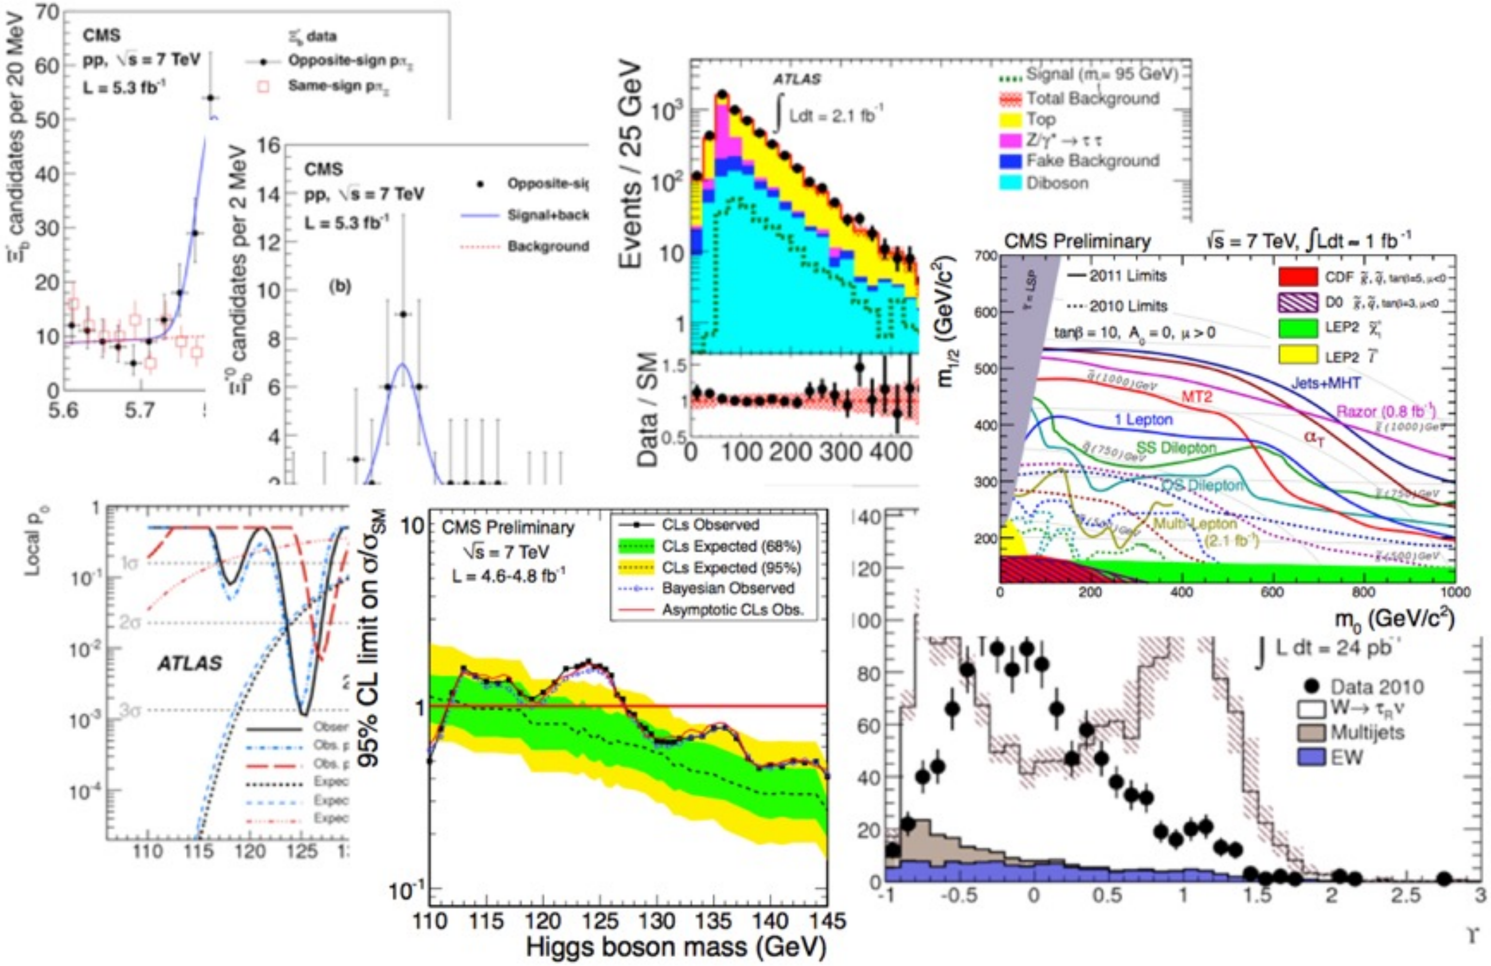

Processing Stages - Simulation

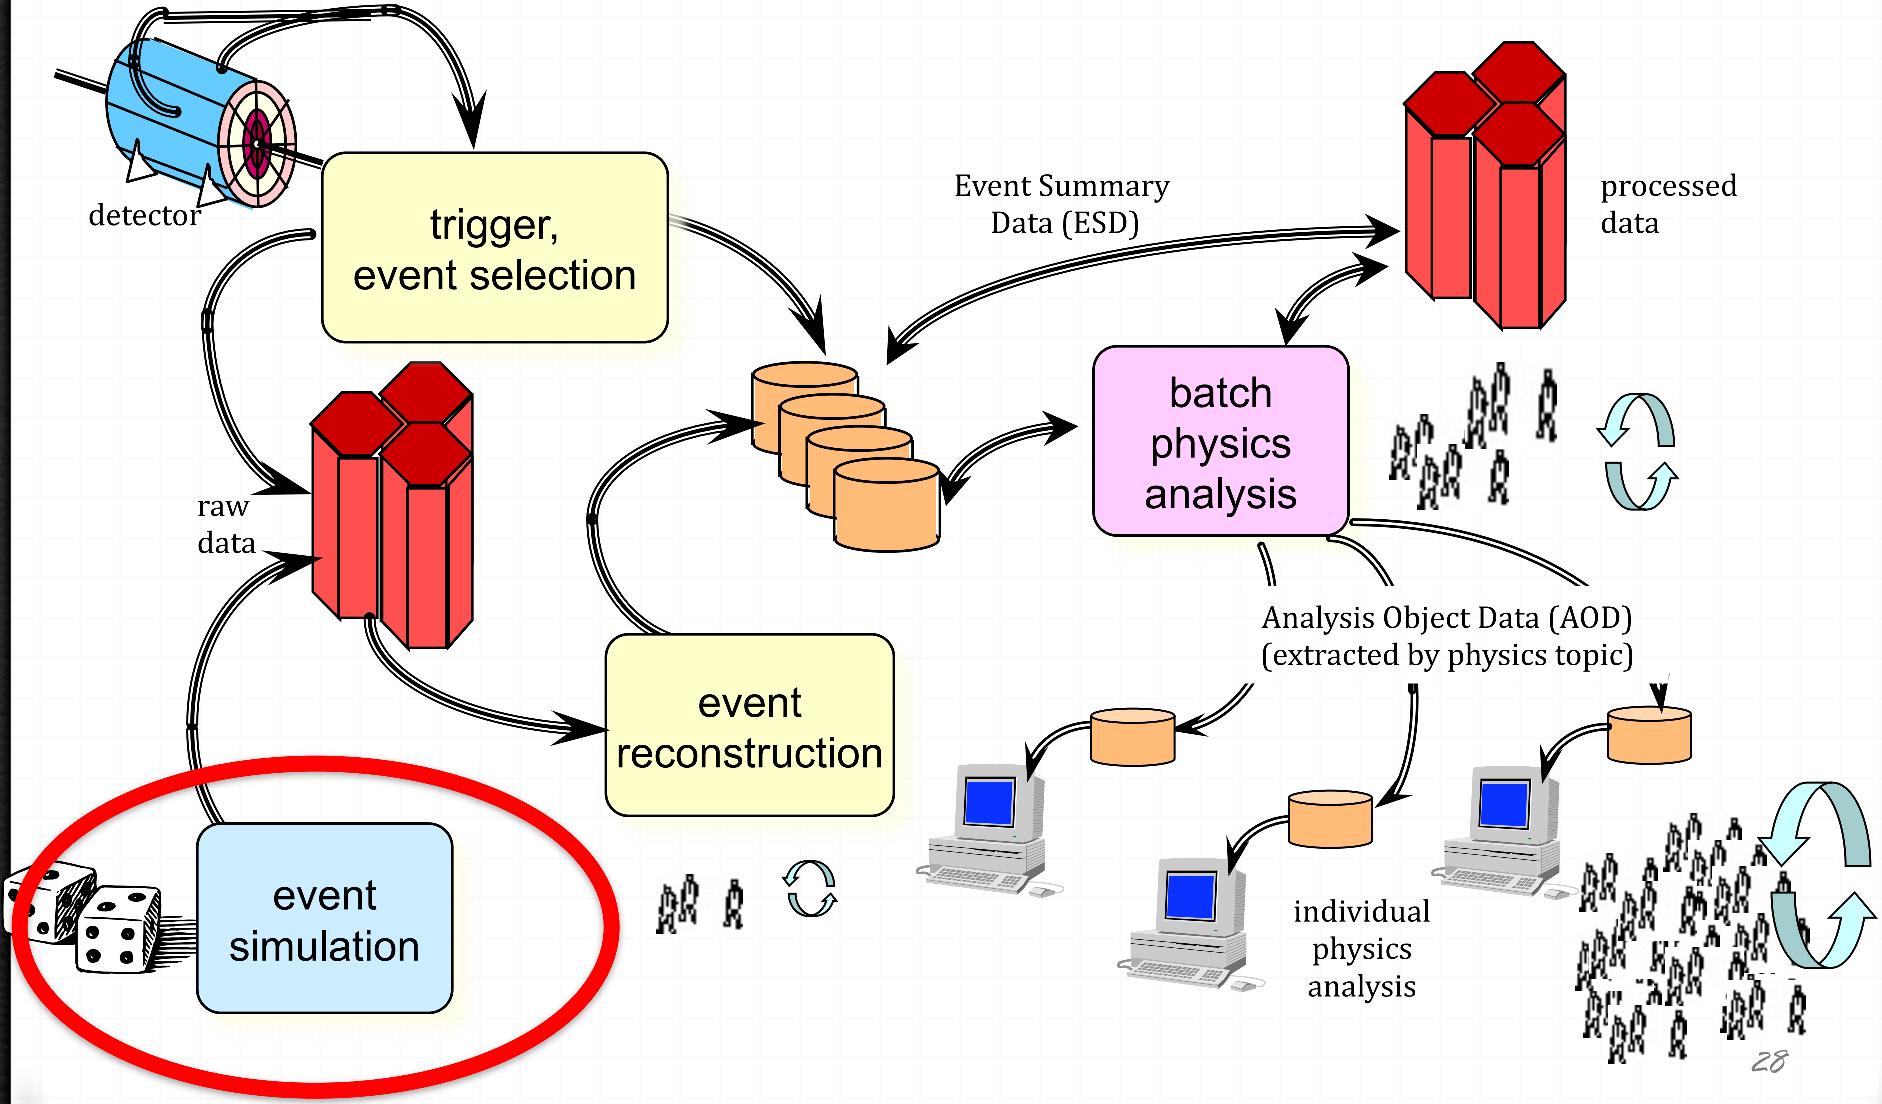

What is simulation?

- simulation = doing ‘virtual’ experiment
- take all the known physics
- start from your ‘initial condition’ (two protons colliding)
- calculate the ‘final state’ of your detector to get the ‘experimental’ results
  - solve equations of motion, etc
- IMPOSSIBLE to be done analytically

Monte Carlo simulation

- What is Monte Carlo?
  - Throwing random numbers
    - to calculate integrals
    - to pick among possible choices
- Why Monte Carlo?
  - complexity of the problem
  - lack of analytical description
  - need of randomness like in nature
    - Quantum mechanics: amplitudes => probabilities
      - Anything that possibly can happen, will! (but more or less often)
    - Want to generate events in as much detail as Mother Nature
      - get average and fluctuations right
      - make random choices, ~as in nature

Buffon’s Needle

- One of the oldest problems in the field of geometrical probability, first stated in 1777.
- Drop a needle on a lined sheet of paper and determine the probability of the needle crossing one of the lines
- Remarkable result: probability is directly related to the value of π
- The needle will cross the line if $x \leq L \sin(\vartheta)$. Assuming $L \leq D$, how often will this occur?

$$P_{cut} = \int_0^\pi P_{cut}(\theta)\frac{d\theta}{\pi} = \int_0^\pi \frac{L\sin\theta}{D}\frac{d\theta}{\pi} = \frac{L}{\pi D}\int_0^\pi \sin\theta\, d\theta = \frac{2L}{\pi D}$$

- By sampling P_{cut} one can estimate π.

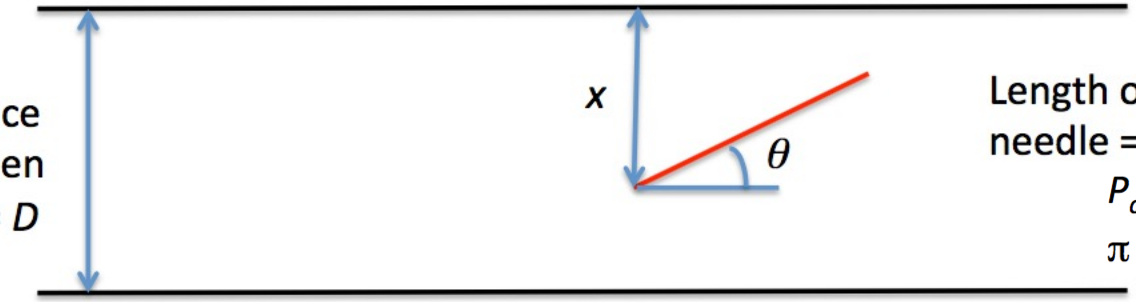

$P_{cut} = h/n = 1$

$\pi = (2\, L / D) * (\, n / h\,)$

Laplace method of calculating π (1886)

- Area of the square = 4
- Area of the circle = π
- Probability of random points inside the circle = π / 4

- Random points : N
- Random points inside circle : N_c

$\pi \sim 4\, N_c / N$

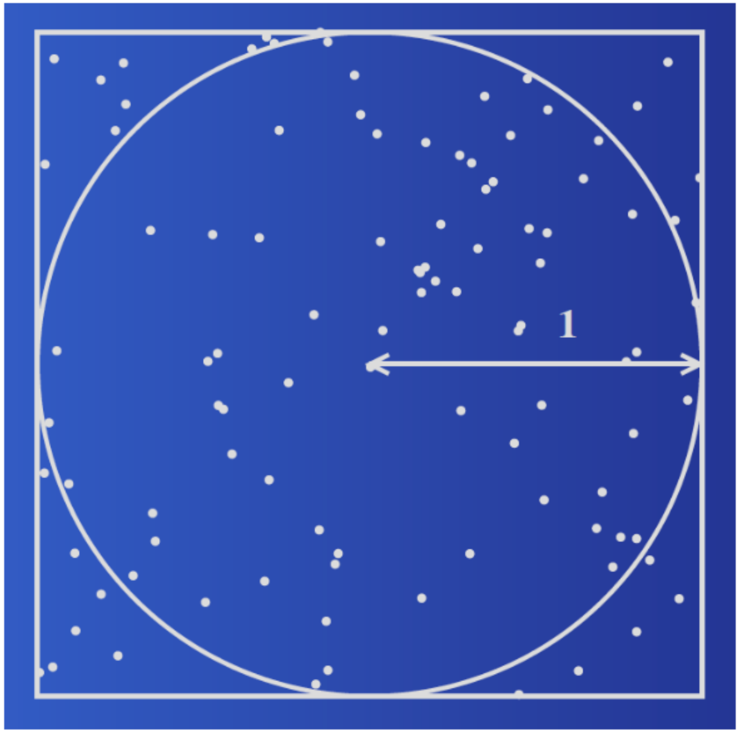

Why do we need simulation?

- to design the apparatus (detector) to fulfill its role
- to prepare the reconstruction and analysis of results
  - training on ‘known’ (simulated) events
- to understand the results
  - we need to know what to expect to
    - verify existing models
    - find new physics

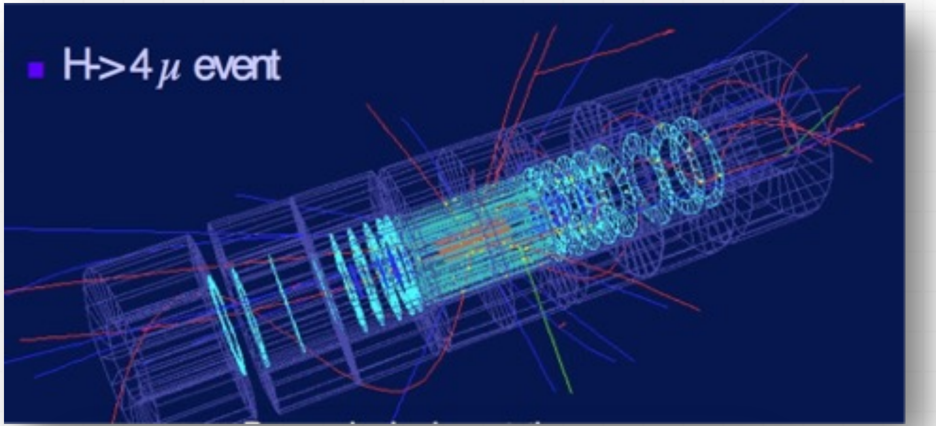

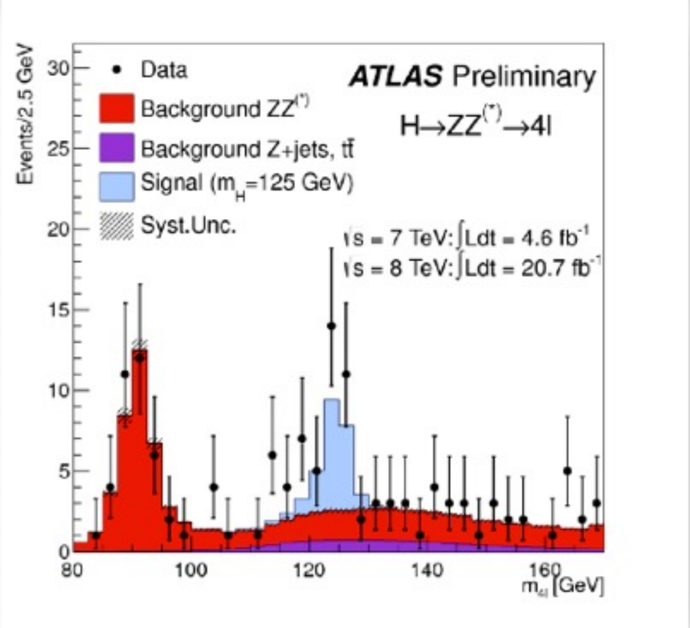

Simulation chain for HEP experiments

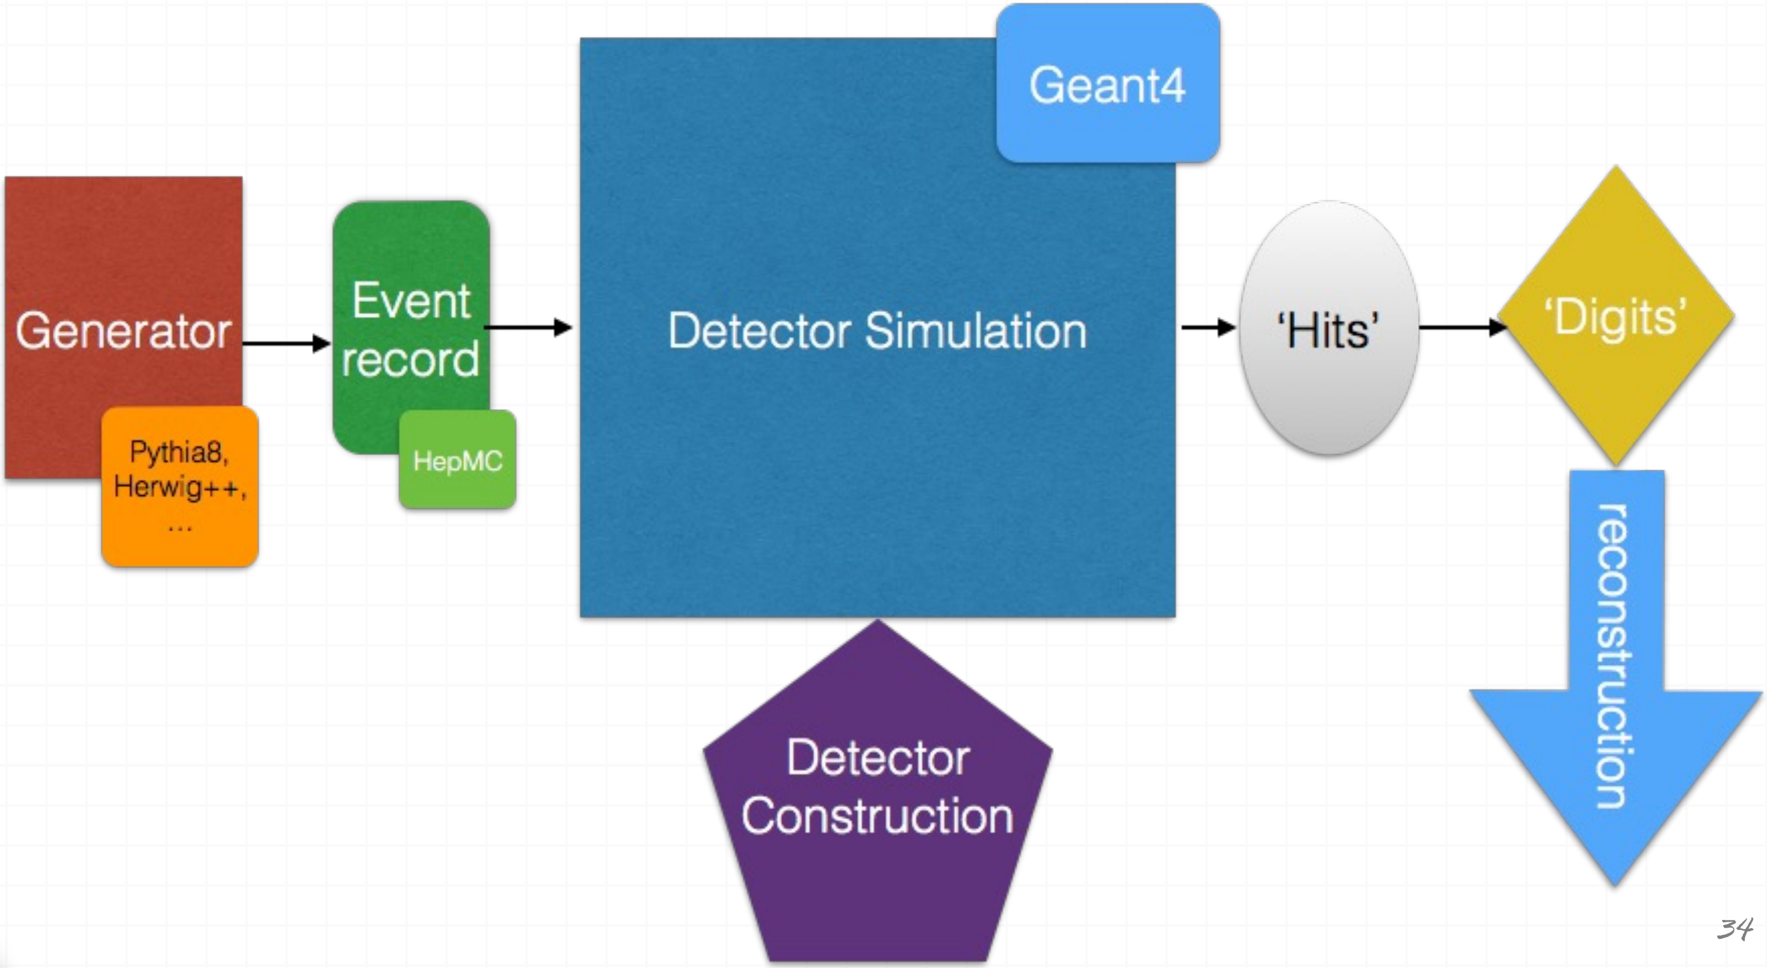

Monte Carlo generators

- simulate particles reaction in vacuum
    - knows nothing about the surrounding detector
- all Standard Model processes are included
- no propagation of particles, just generation of the products of the 'primary' collision
- the output of the 'generators' is the input to the 'transport' code

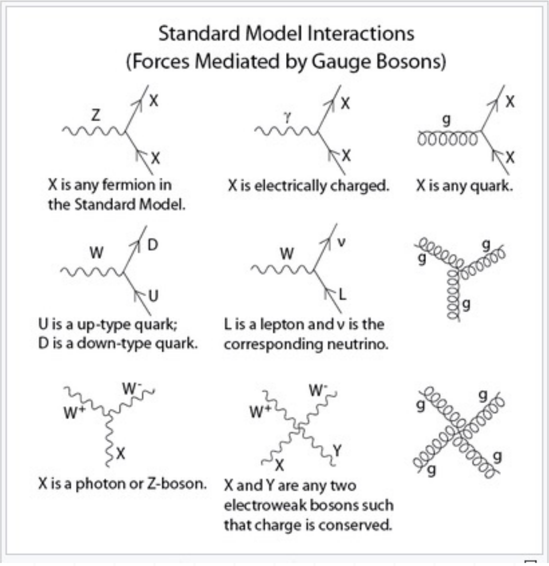
Standard Model Interactions
(Forces Mediated by Gauge Bosons)

X is any fermion in the Standard Model.

X is electrically charged.

X is any quark.

U is a up-type quark;
D is a down-type quark.

L is a lepton and v is the corresponding neutrino.

X is a photon or Z-boson.

X and Y are any two electroweak bosons such that charge is conserved.

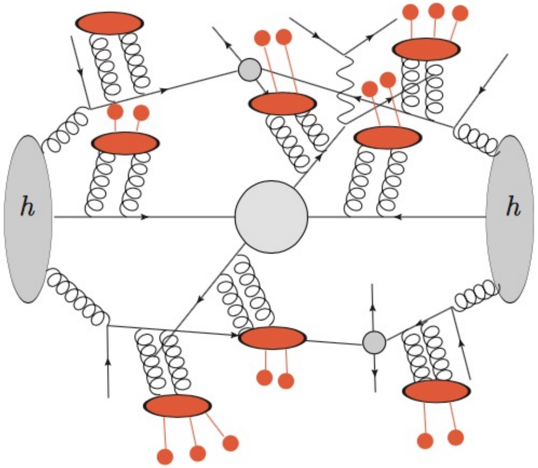

Transport Code: Geant4

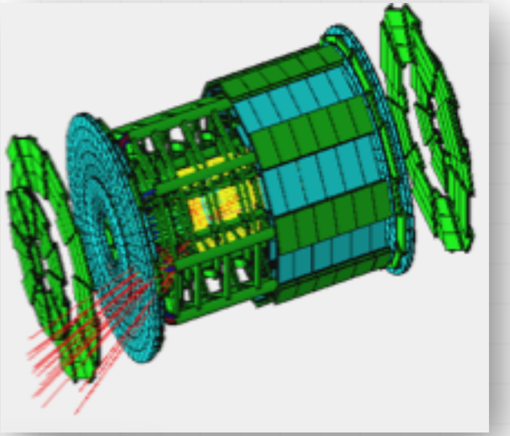

- Geant4 is a toolkit (C++) for the simulation of the passage of particles through matter.
- Its areas of application include high energy, nuclear and accelerator physics, as well as studies in medical and space science
- In HEP has been successfully employed for
  - Detector design
  - Calibration/alignment
  - Data analysis

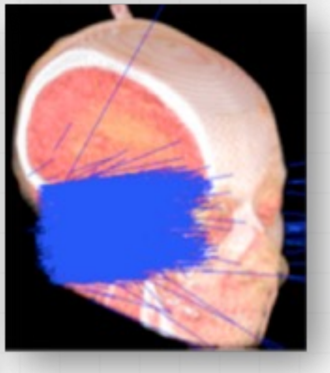

What does Geant4 do?

- ‘propagates’ particles through geometrical structures of materials, including magnetic field
- simulates processes the particles undergo
  - creates secondary particles
  - decays particles
- calculates the deposited energy along the trajectories and allows to store the information for further processing (‘hits’)

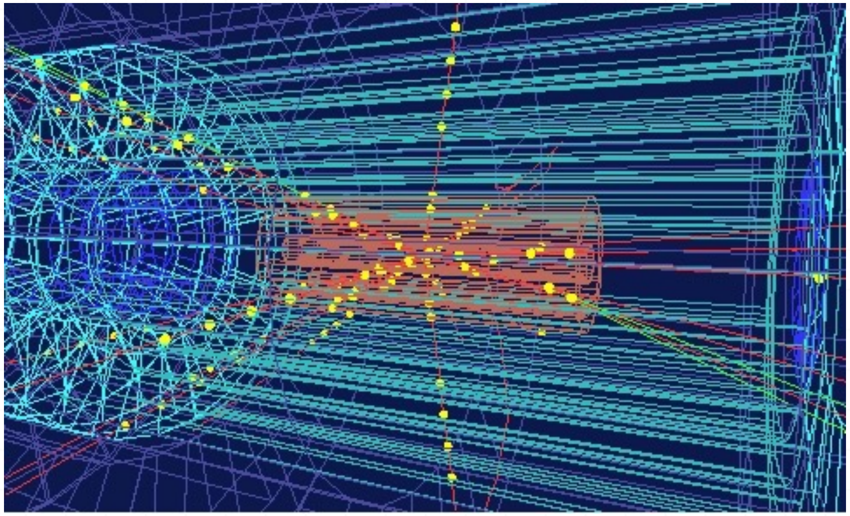

Simulation ingredients

- we model
  - Detector’s Geometry
    - Shape, Location, Material
  - Physics interactions
    - All known processes
      - Electromagnetic
      - Nuclear (strong)
      - Weak (decay
  - we ‘shoot’ particles and ‘propagate’ them through the modeled detector

$$\sigma_{\text{total}} = \Sigma\, \sigma_{\text{per-interaction}}$$

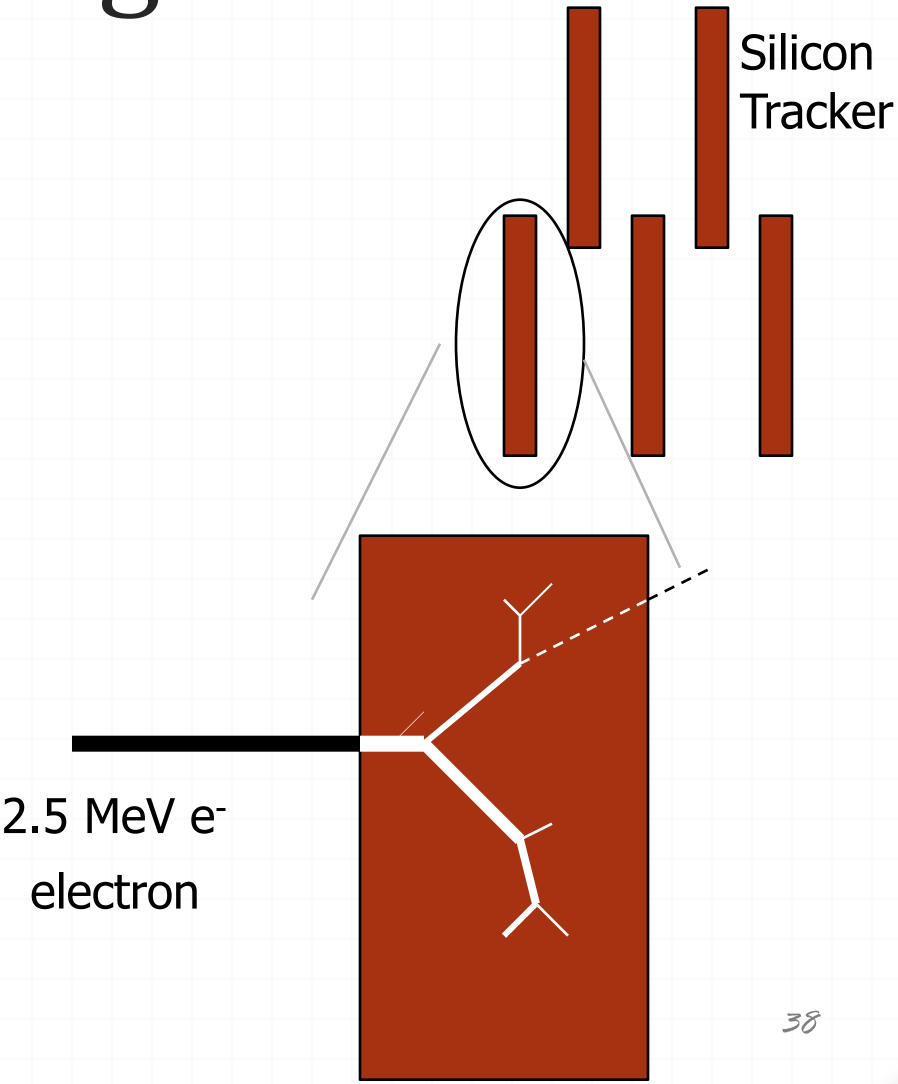

Geant4 application

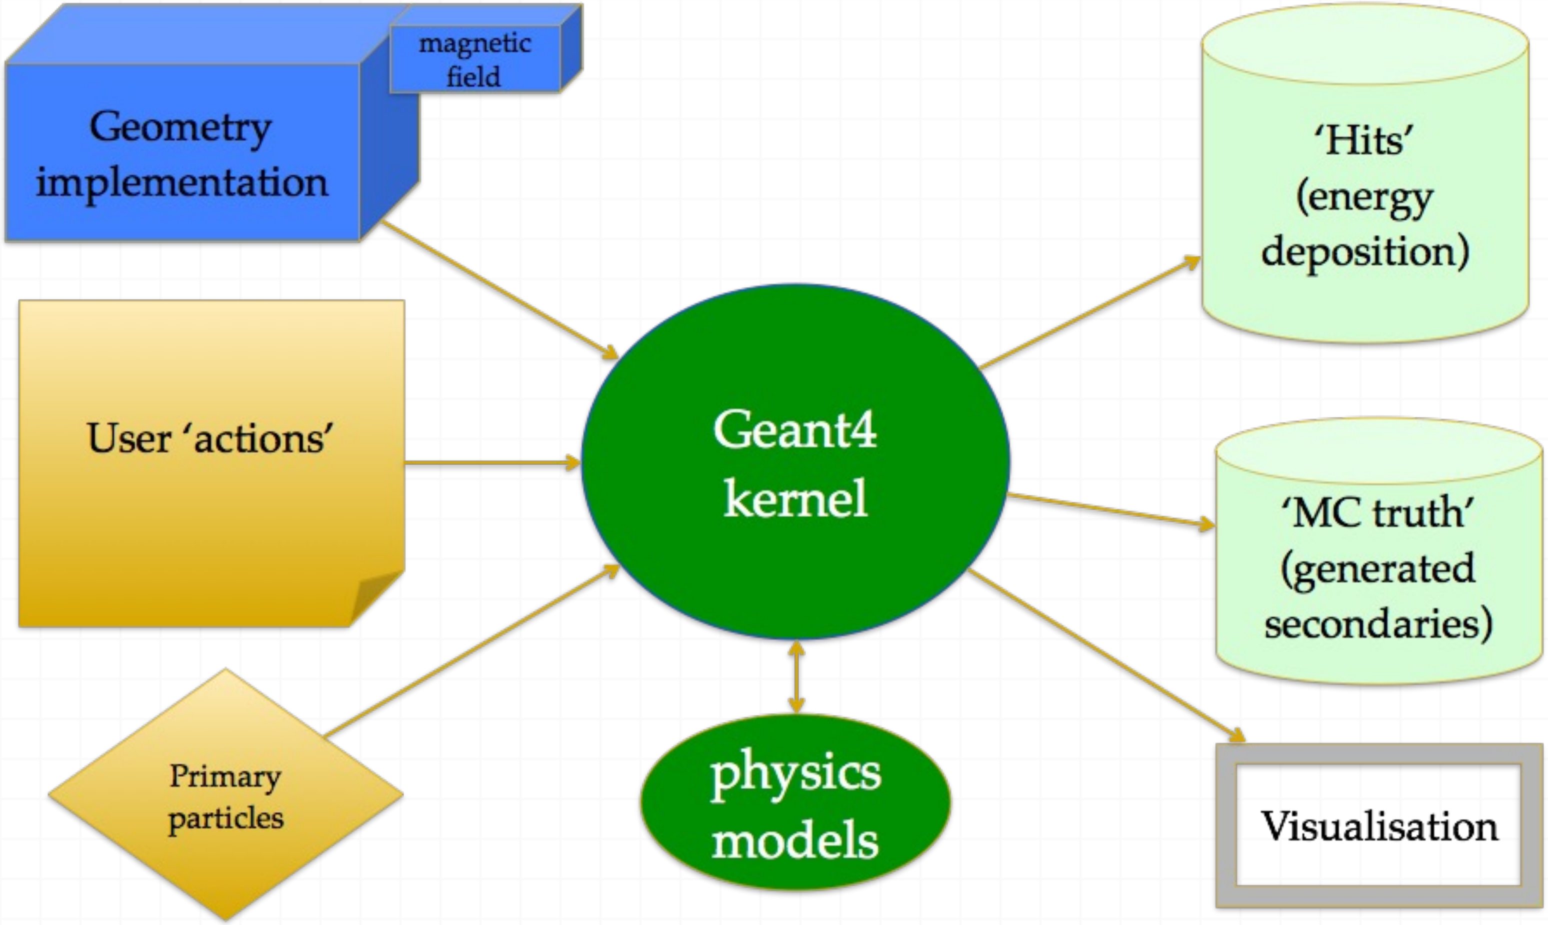

Geometry and Materials

- How to implement (efficiently) this in your computer program?
- You need 'bricks'
  - 'solids', 'shapes'
  - you need to position them
  - you want to 'reuse' as much as possible the same 'templates'
- Database of Materials
  - National Institute of Standards (NIST)
- Magnetic Fields
  - numerical integration of the equation of motion (Runge-Kutta method)

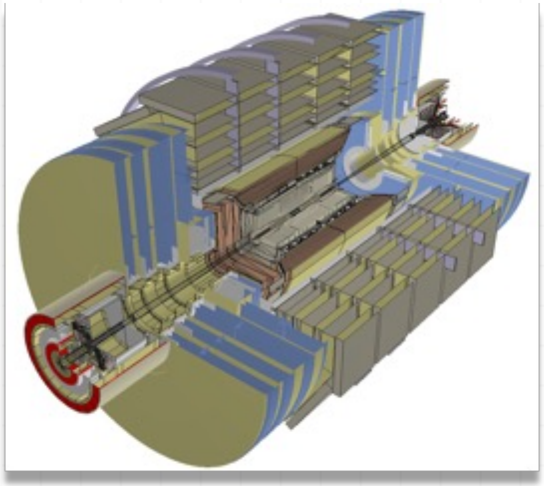

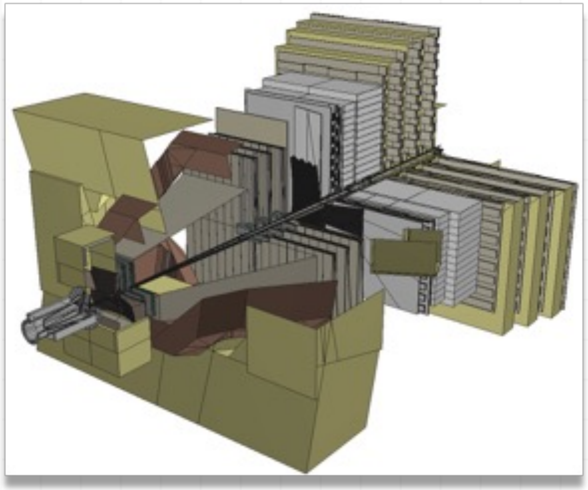

Geometry construction

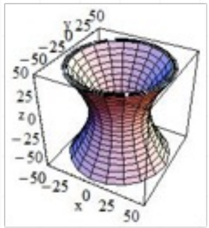
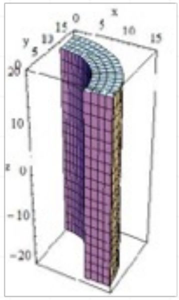
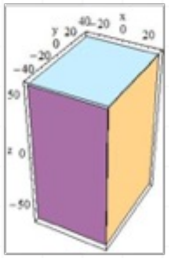
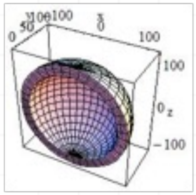
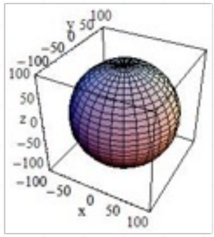
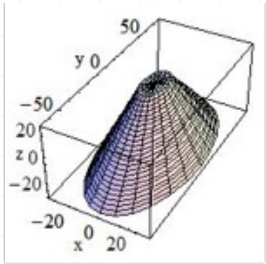
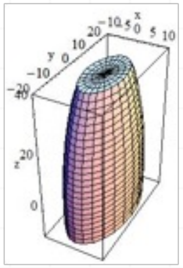
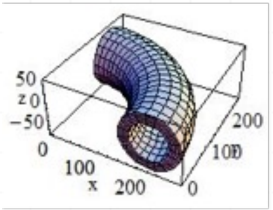

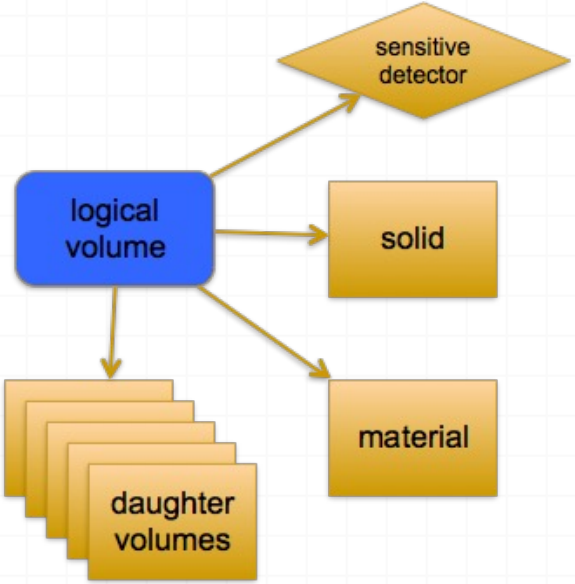
sensitive detector

logical volume

solid

material

daughter volumes

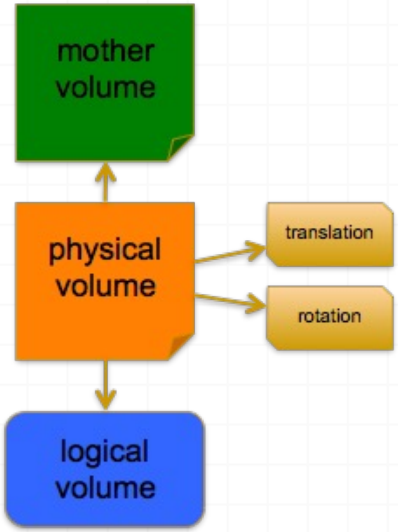
mother volume

physical volume

translation

rotation

logical volume

Physics…

- What happens to particles in matter?
- We want to model the physics we know
  - each possible physics process provides the "interaction length" compared with distance to next geometrical boundary
    - the smallest wins
  - generating a "final state" and secondaries tracks
- Electromagnetic
  - gammas and charged particles
- Hadronic
  - neutrons, mesons (K,π), muons, ...

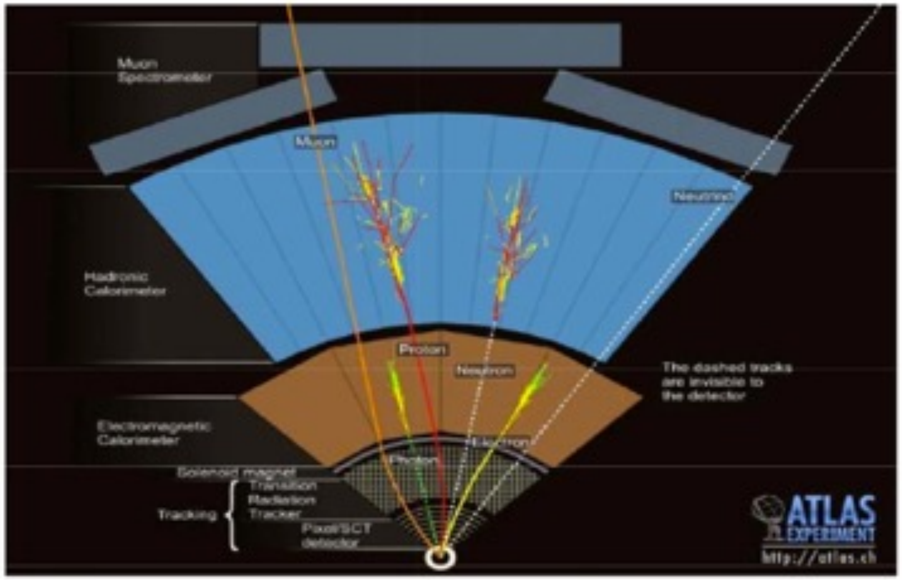

Because of the detailed geometries, the detailed physics and the required precision the simulation is very CPU hungry

Physics processes

- models exist for different physics process
  - none of the model is perfect
    - sometimes different models for the same process cover different energy ranges
    - tradeoff between precision and CPU
- for example Electromagnetic processes include:
  - Gammas:
    - Gamma-conversion, Compton scattering, Photo-electric effect
  - Leptons(e, μ), charged hadrons, ions
    - Energy loss (Ionisation, Bremstrahlung) or PAI model energy loss, Multiple scattering, Transition radiation, Synchrotron radiation,
  - Photons:
    - Cerenkov, Rayleigh, Reflection, Refraction, Absorption, Scintillation
  - high energy muons and lepton-hadron interactions

ATLAS

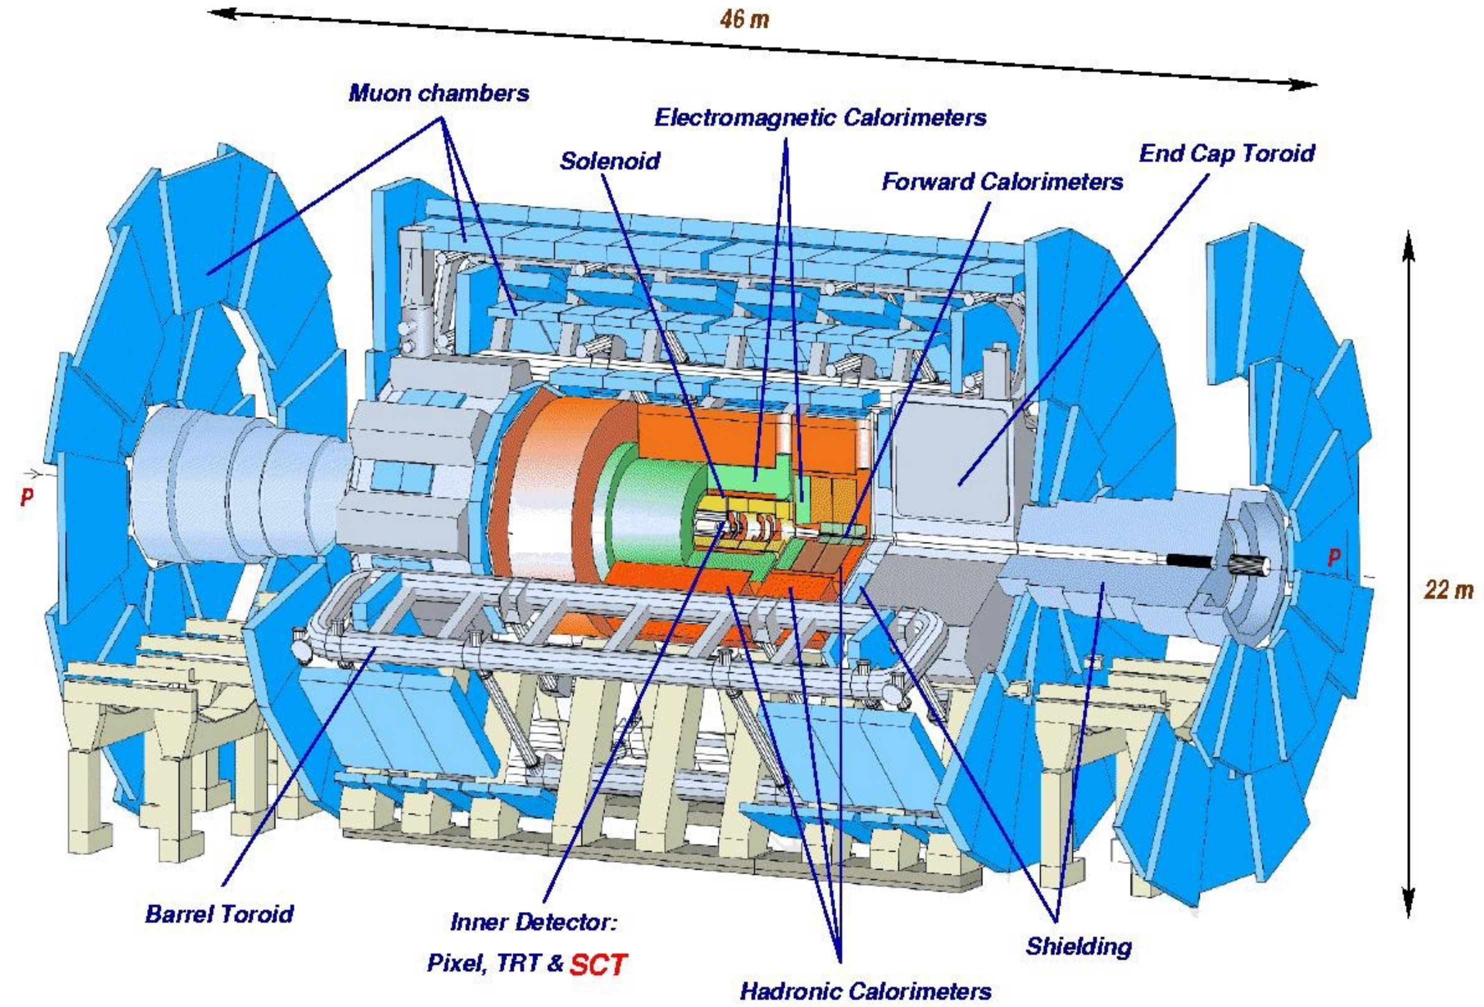
46 m
Muon chambers
Electromagnetic Calorimeters
Solenoid
End Cap Toroid
Forward Calorimeters
P
P
22 m
Barrel Toroid
Inner Detector:
Pixel, TRT & SCT
Shielding
Hadronic Calorimeters

ATLAS Calorimeter (a very, very small part of it)

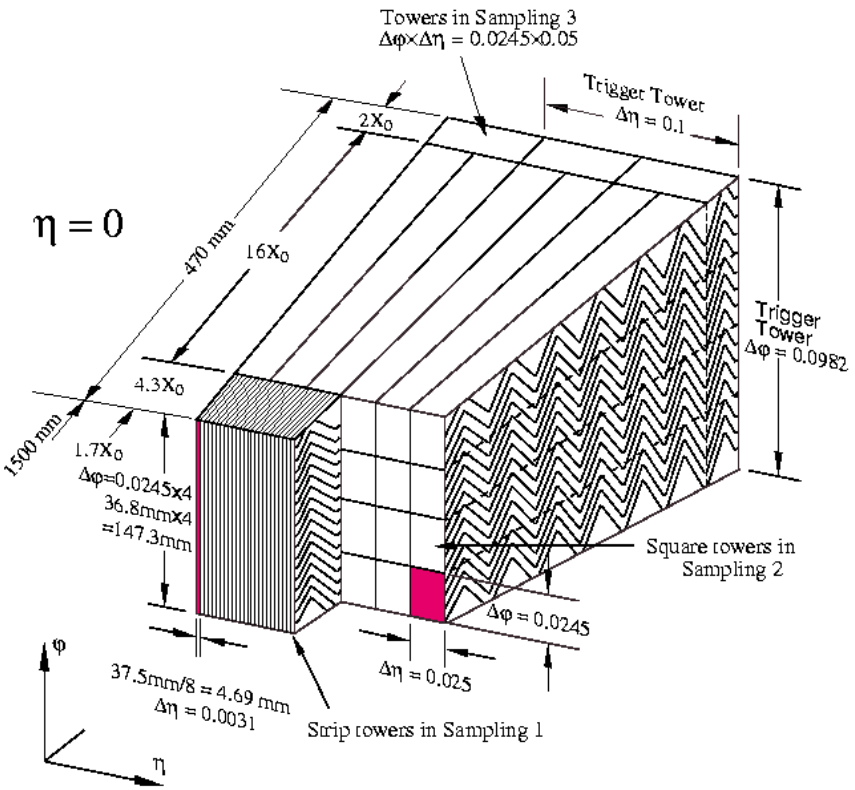

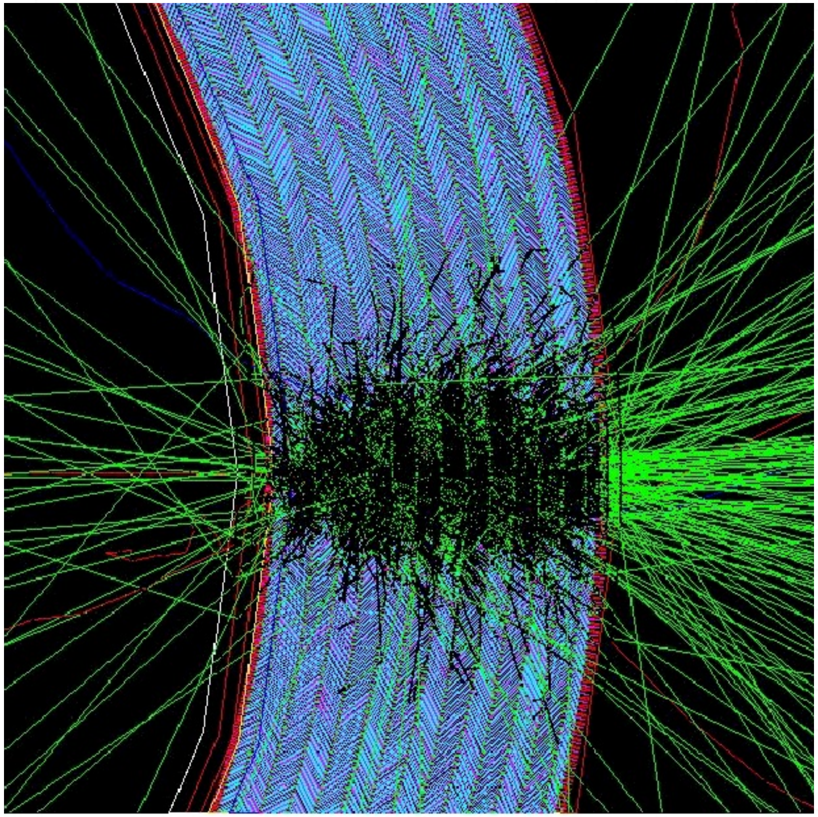

. Kordas "Geant4 for the ATLAS EM calo" — CALOR2000, Annecy, 12 October2000 (4

Data rates

- Particle beams cross every 25 ns (40 MHz)
  - Up to 25 particle collisions per bea
  - Up to 10^9 collisions per second
- Basically 2 event filter/trigger leve
  - Hardware trigger (e.g. FPGA)
  - Software trigger (PC farm)
  - Data processing starts at readout
  - Reducing 10^9 p-p collisions per second to O(1000)
- Raw data to be stored permanently: >15 PB/year

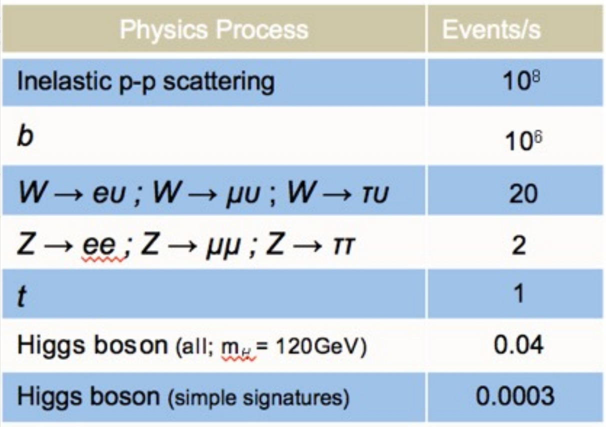

Physics Process	Events/s
Inelastic p-p scattering	10^8
b	10^6
$W \to e\nu$; $W \to \mu\nu$; $W \to \tau\nu$	20
$Z \to ee$; $Z \to \mu\mu$; $Z \to \tau\tau$	2
t	1
Higgs boson (all; m_H = 120GeV)	0.04
Higgs boson (simple signatures)	0.0003

This is our Big Data problem!!

Big Data requires Big Computing

- The LHC experiments rely on distributed computing resources:
  - WLCG - a global solution, based on the Grid technologies/middleware
    - distributing the data for processing, user access, local analysis facilities etc.
    - at time of inception envisaged as the seed fo f the technologies
- Tiered structure
  - Tier-0 at CERN: the central facility for data processin
  - 11 Tier-1s: big computing centers with high quali most complex/intensive processing operations
  - ~140 Tier-2s: computing centers across the w data analysis and simulation.
- So far computing was not a limiting factor fo of the LHC experiments

Capacity:
~350,000 CPU cores
~200 PB of disk space
~200 PB of tape space

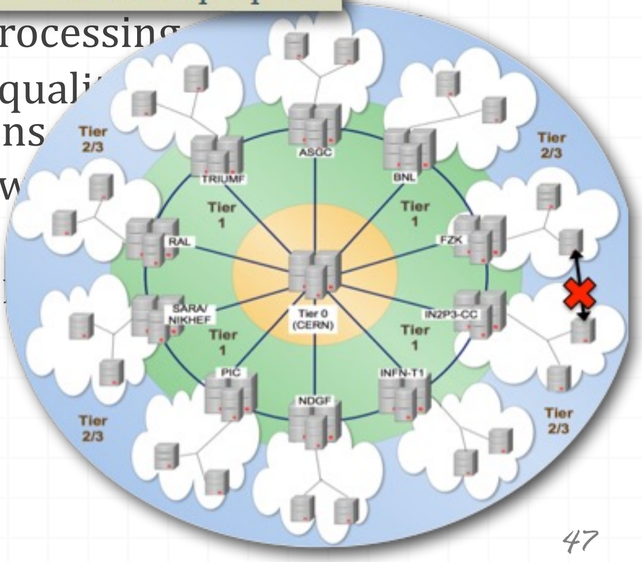

Running jobs on LCG

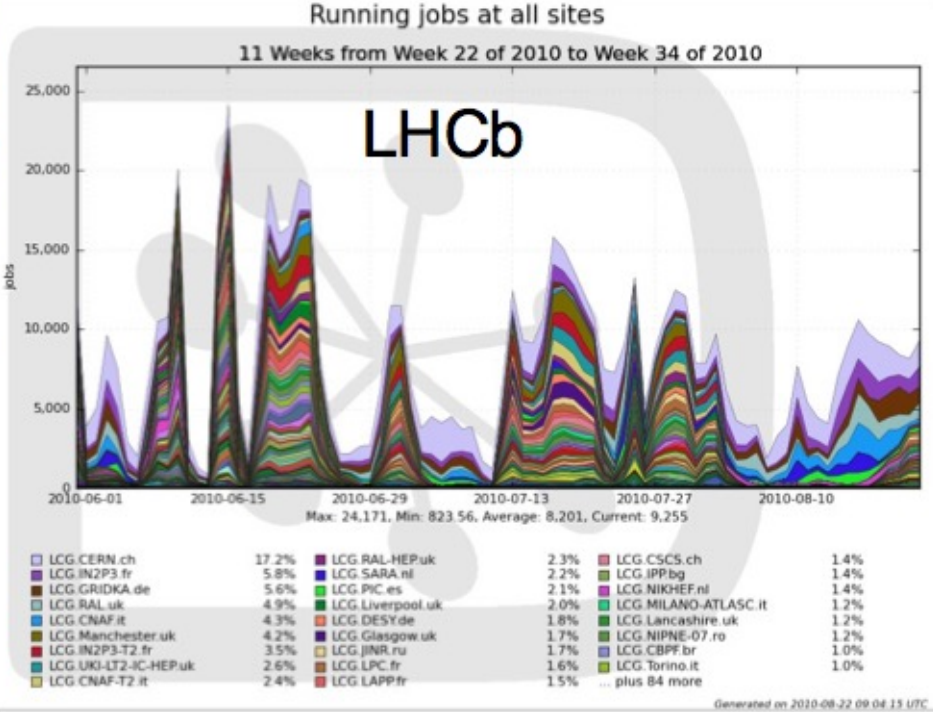

LHCb

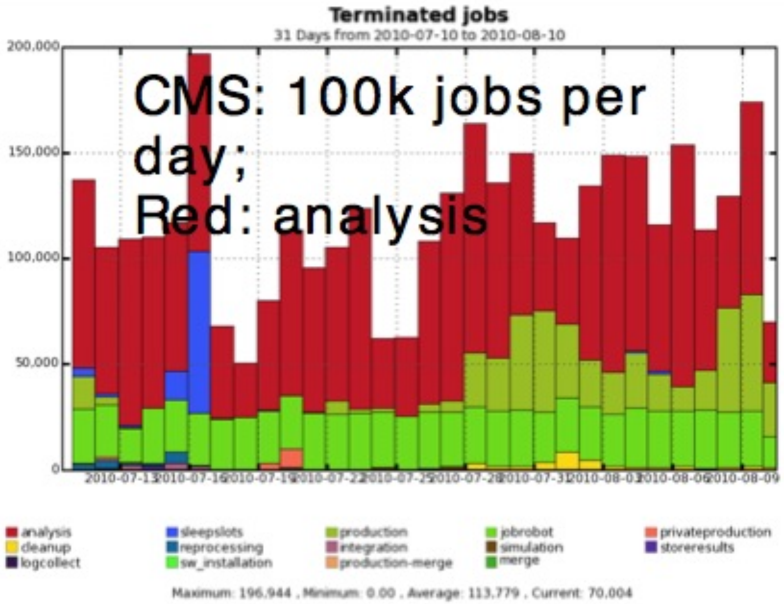

CMS: 100k jobs per day;
Red: analysis

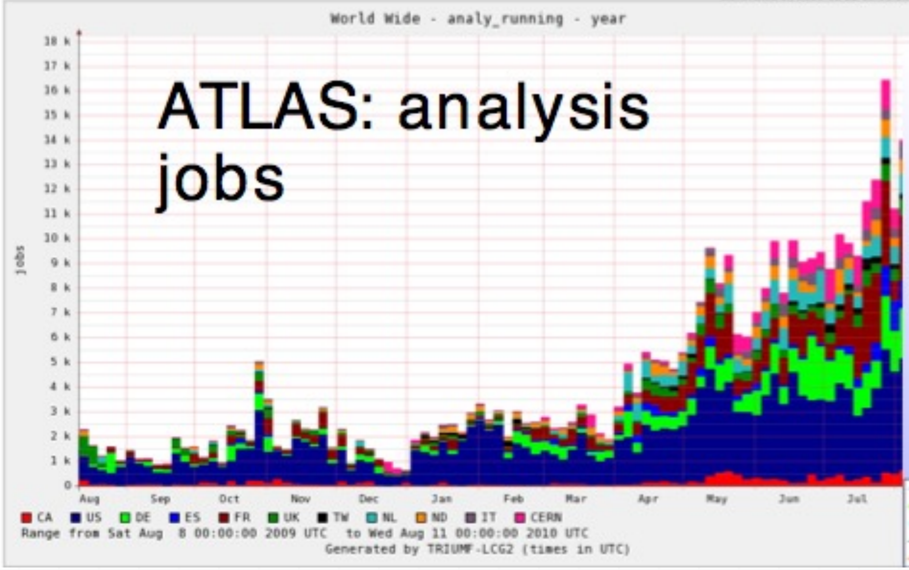

ATLAS: analysis jobs

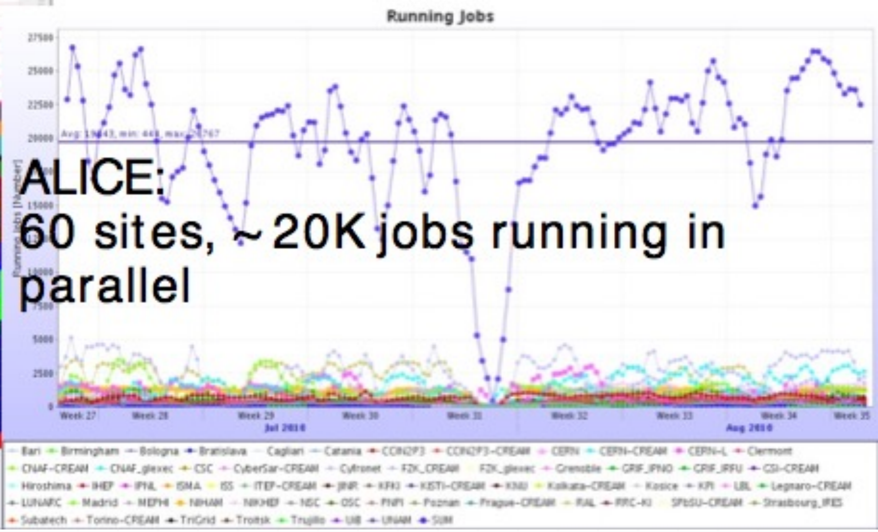

ALICE:
60 sites, ~20K jobs running in parallel

2010 Tier-0 Data Taking

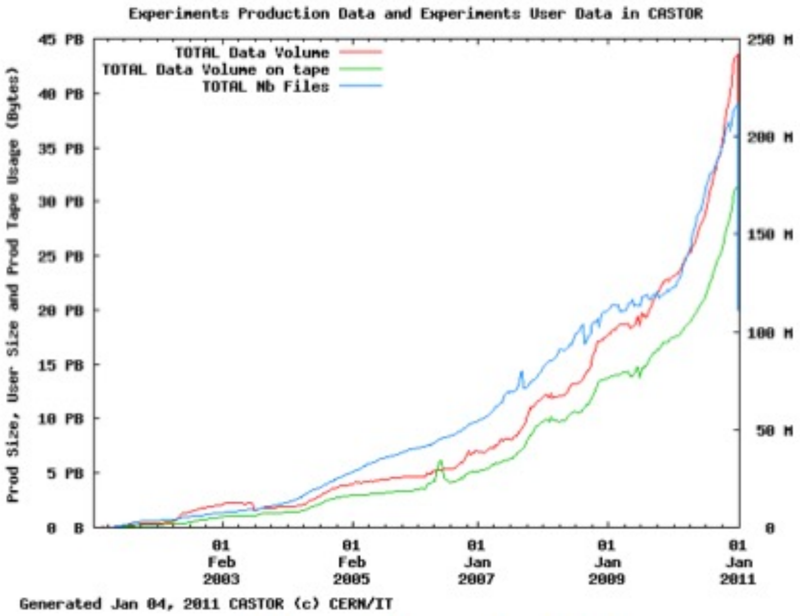

Stored ~ 15 PB in 2010 with peaks at 220 TB/day during Pb+Pb

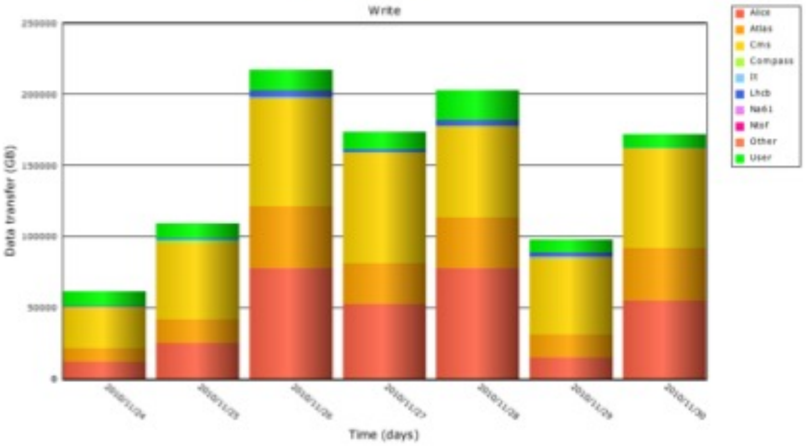

Tier-0 Bandwidth
Average in: 2 GB/s with peaks at 11.5 GB/s
Average out: 6 GB/s with peaks at 25 GB/s

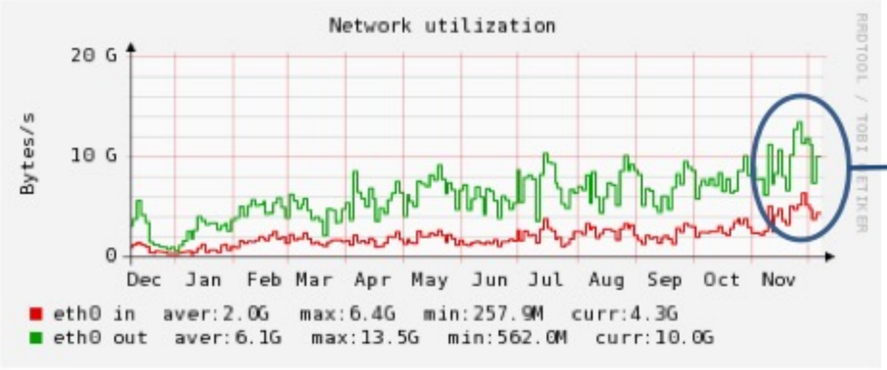

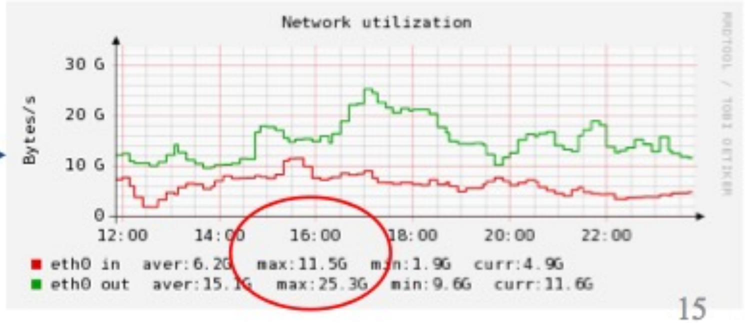

A Success Story!

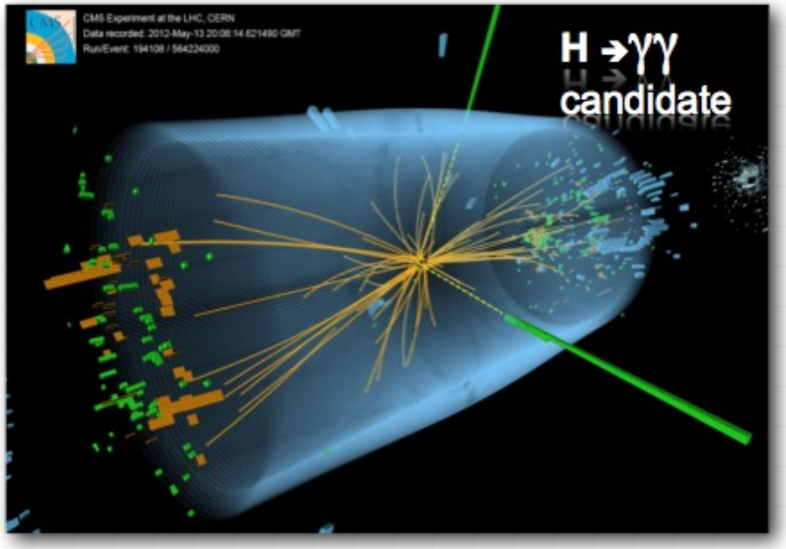

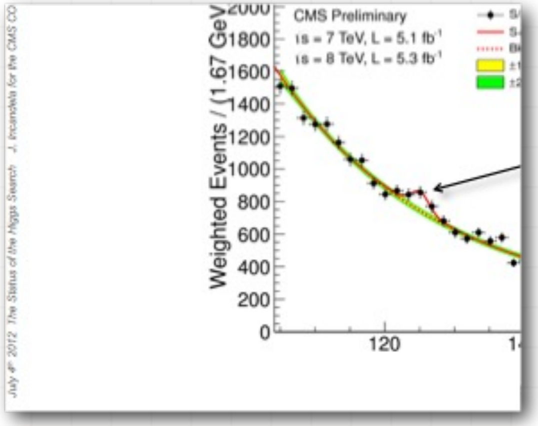

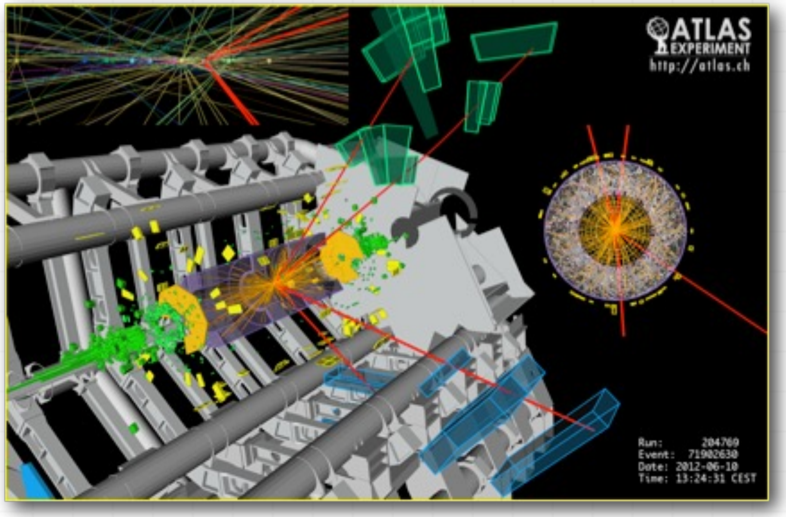

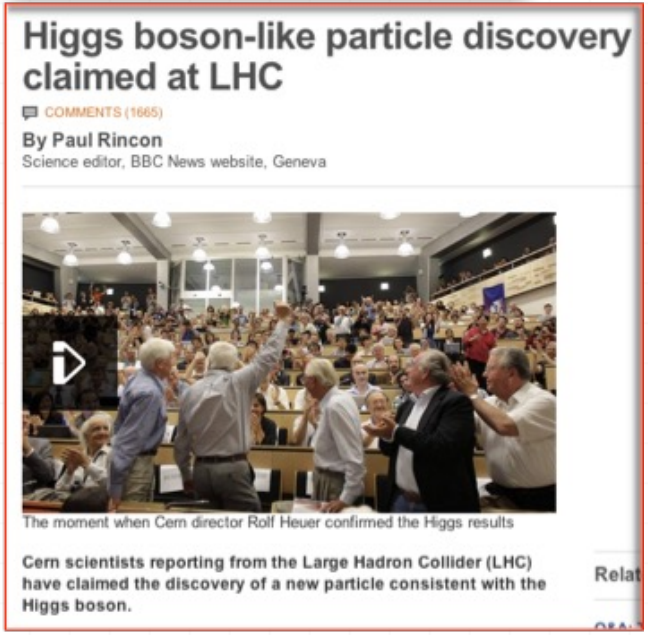

Higgs boson-like particle discovery claimed at LHC

COMMENTS (1665)

By Paul Rincon
Science editor, BBC News website, Geneva

The moment when Cern director Rolf Heuer confirmed the Higgs results

Cern scientists reporting from the Large Hadron Collider (LHC) have claimed the discovery of a new particle consistent with the Higgs boson.

Prospects for HEP Software

- **Potential gains** can be made by exploiting features of todays CPUs' micro architecture
    - using vector registers, instruction pipelining, multiple instructions per cycle
    - improving data and code locality and making use of hardware threading
- Today multi-core architectures employing O(10) cores are well exploited using a multi-process model (1 job/core).
- New architectures to **off-load large computations** to accelerators (GPUs, many-CPU chips, FPGAs)
- However this performance will **not scale** to new generations of many-core architectures employing O(100) cores **due to memory issues**
    - technical issues related to connecting many cores to shared memory
    - memory footprint of HEP code is increasing due to increasing event complexity as the energy and luminosity of the LHC is increased

Conclusion

- modern HEP experiments would be impossible without computing
  - online triggering and selection
  - offline reconstruction, analysis and simulation
- huge data volumes
- distributed processing