Physics experiments at CERN

CERN Meet Up with OCP High Performance Computing Project, Geneva, June 2015

Niko Neufeld, CERN/PH

About CERN

- CERN is the European Organization for Nuclear Research in Geneva
  - Particle accelerators and other infrastructure for high energy physics (HEP) research
  - Worldwide community
    - 21 member states (+ 3 incoming members)
    - Observers: Turkey, Russia, Japan, USA, India
    - About 2300 staff
    - >10’000 users (about 5’000 on-site)
    - Budget ~1100 MCHF / USD / EURO …

- Birthplace of the World Wide Web

The Mission of CERN

- Push forward the frontiers of knowledge
  - The secrets of the Big Bang
  - Origin of mass

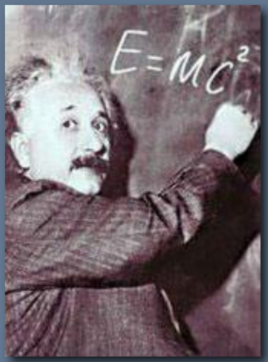

- Develop new technologies for accelerators and detectors
  - Information technology - the Web and the Grid
  - Medicine - diagnosis and therapy

- Train scientists and engineers of tomorrow

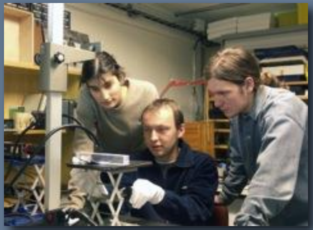

- Unite people from different countries and cultures

An international lab

Distribution of All CERN Users by Nationality on 14 January 2014

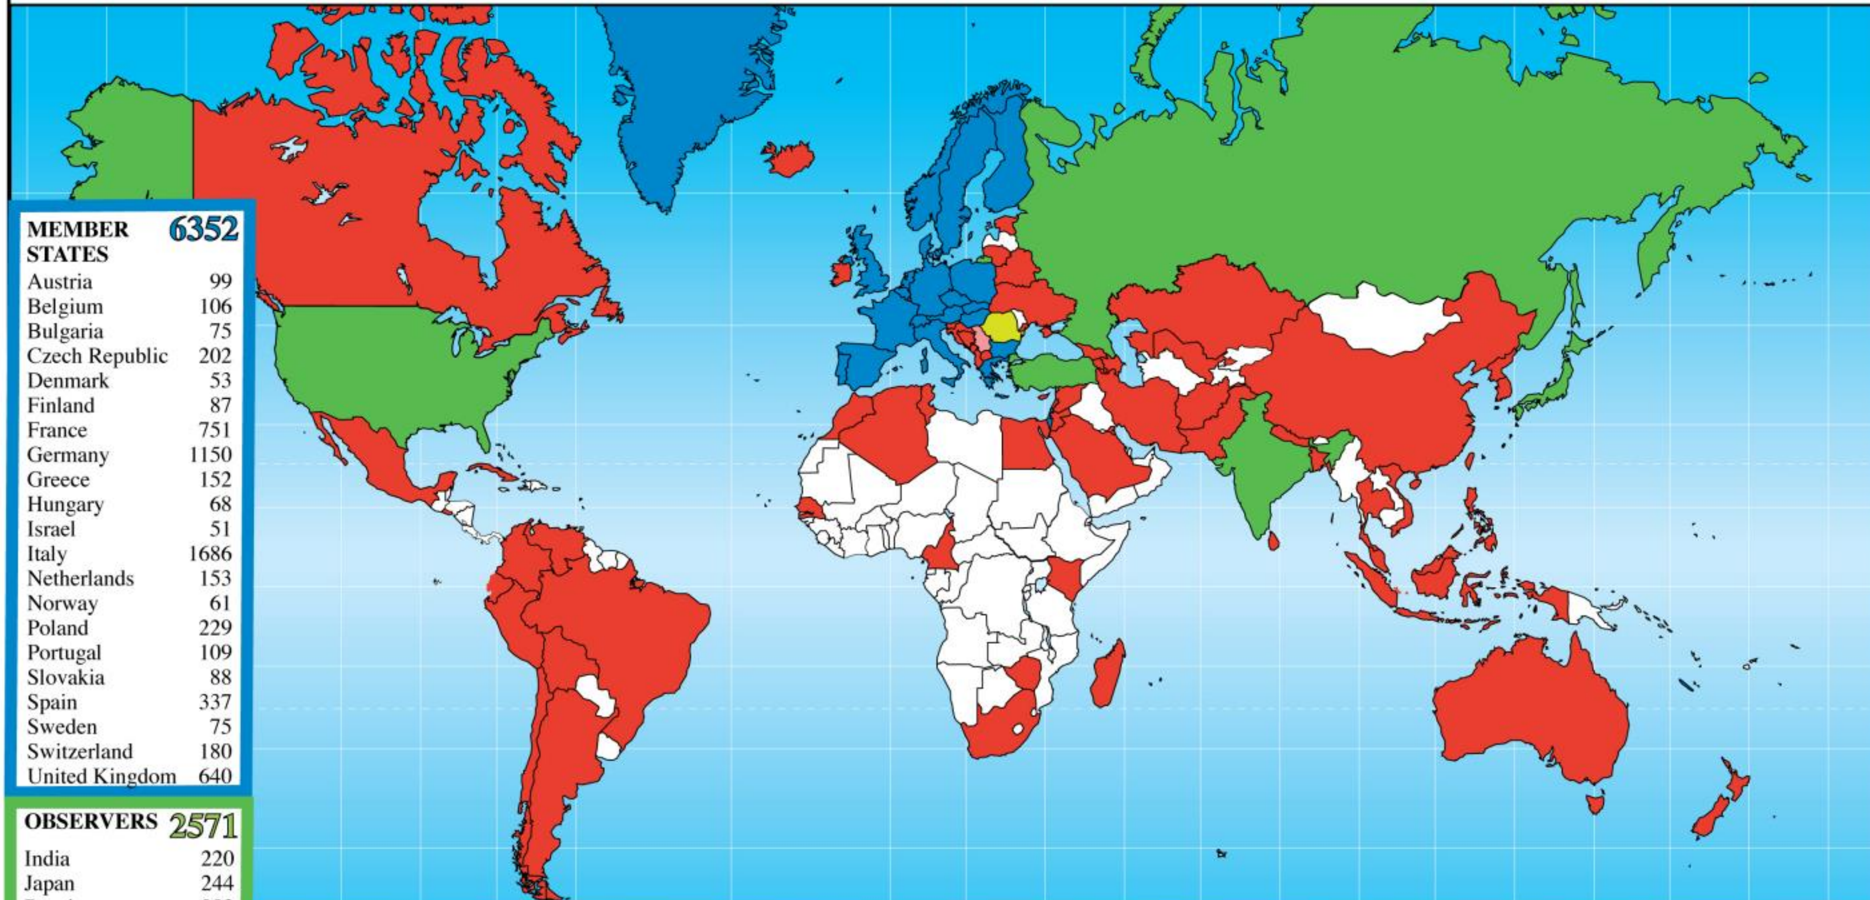

MEMBER STATES	6352
Austria	99
Belgium	106
Bulgaria	75
Czech Republic	202
Denmark	53
Finland	87
France	751
Germany	1150
Greece	152
Hungary	68
Israel	51
Italy	1686
Netherlands	153
Norway	61
Poland	229
Portugal	109
Slovakia	88
Spain	337
Sweden	75
Switzerland	180
United Kingdom	640

OBSERVERS	2571
India	220
Japan	244
Russia	982
Turkey	146
USA	979

CANDIDATE FOR ACCESSION	
Romania	118

ASSOCIATE MEMBERS IN THE PRE-STAGE TO MEMBERSHIP	
Serbia	41

OTHERS	1415
Afghanistan	1
Albania	2
Algeria	8
Argentina	11
Armenia	25
Australia	25
Azerbaijan	8
Bangladesh	4
Belarus	47
Bolivia	3
Bosnia & Herzegovina	1
Brazil	108
Cameroon	1
Canada	134
Cape Verde	1
Chile	12
China	280
China (Taipei)	45
Colombia	30
Croatia	35
Cuba	7
Cyprus	16
Ecuador	3
Egypt	19
El Salvador	1
Estonia	16
Georgia	36
Gibraltar	1
Hong Kong	3
Iceland	4
Indonesia	1
Iran	28
Ireland	22
Jordan	2
Kazakhstan	1
Kenya	1
Korea, D.P.R.	1
Korea Rep.	117
Kuwait	1
Lebanon	12
Lithuania	19
Luxembourg	4
Madagascar	4
Malaysia	15
Mauritius	1
Mexico	64
Montenegro	3
Morocco	12
Nepal	5
New Zealand	7
Pakistan	41
Palestine (O.T.).	4
Peru	8
Philippines	1
Saudi Arabia	3
Senegal	1
Singapore	2
Sint Maarten	2
Slovenia	27
South Africa	16
Sri Lanka	5
Syria	2
Thailand	12
T.F.Y.R.O.M.	1
Tunisia	6
Ukraine	55
Uzbekistan	4
Venezuela	9
Viet Nam	9
Zimbabwe	2

Fundamental questions…

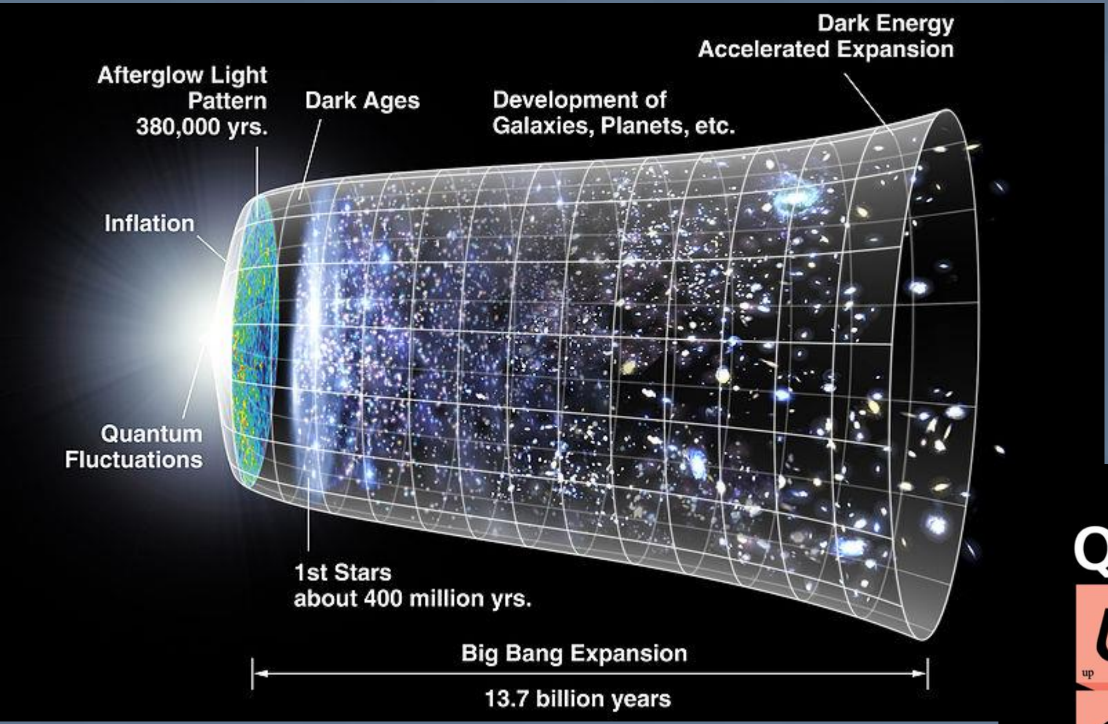

- How to explain that particles have mass?
  - Theories and accumulating experimental data…getting close
- What is 96% of the Universe made of?
  - We only observe 4% of the apparent mass
- Where is all the anti-matter?
  - Why is Nature not symmetric?
- What was the state of matter just after the Big Bang?
  - "Soup" of quarks and gluons before they condensed into matter?

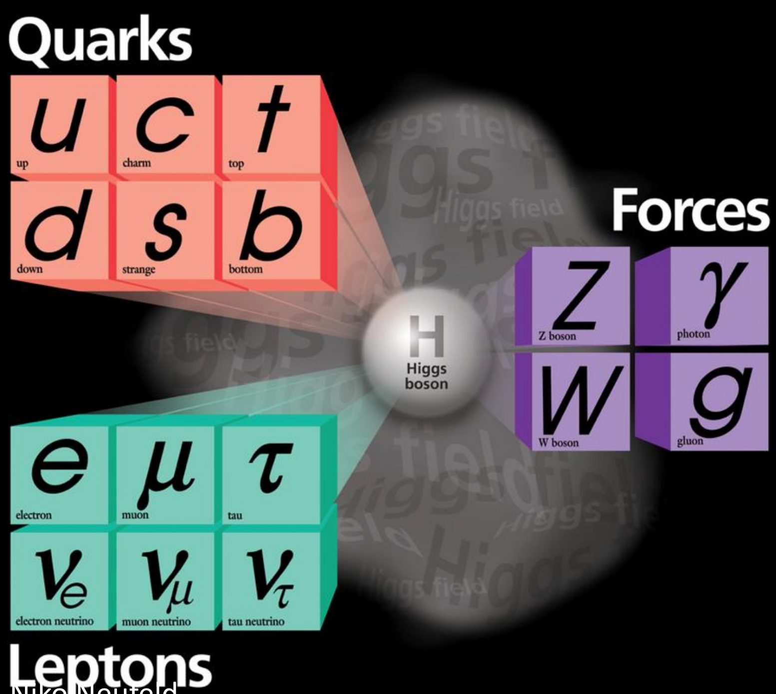

Tools: LHC and detectors

pp, B-Physics, CP Violation (matter-antimatter symmetry)

Exploration of a new energy frontier in p-p and Pb-Pb collisions

General Purpose, proton-proton, heavy ions
Discovery of new physics: Higgs, SuperSymmetry

LHC ring:
27 km circumference

Heavy ions, pp
(state of matter of early universe)

The LHC Experiments today

- ALICE – “A Large Ion Collider Experiment”
  - Size: 26 m long, 16 m wide, 16m high; weight: 10000 t
  - 35 countries, 118 Institutes
  - Material costs: 110 MCHF
- ATLAS – “A Toroidal LHC ApparatuS”
  - Size: 4 6m long, 25 m wide, 25 m high; weight: 7000 t
  - 38 countries, 174 institutes
  - Material costs: 540 MCHF
- CMS – “Compact Muon Solenoid”
  - Size: 22 m long, 15 m wide, 15 m high; weight: 12500 t
  - 40 countries, 172 institutes
  - Material costs: 500 MCHF
- LHCb – “LHC beauty” (b-quark is called “beauty” or “bottom” quark)
  - Size: 21 m long, 13 m wide, 10 m high; weight: 5600 t
  - 15 countries, 52 Institutes
  - Material costs: 75 MCHF
- Regular upgrades … first 2013/14 (Long Shutdown 1)

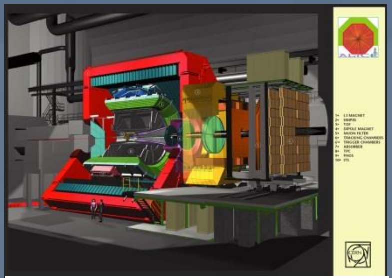

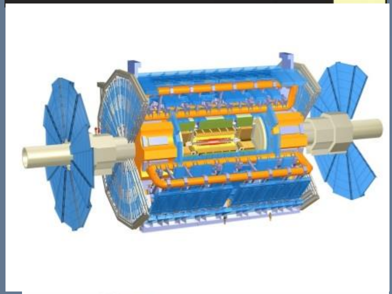

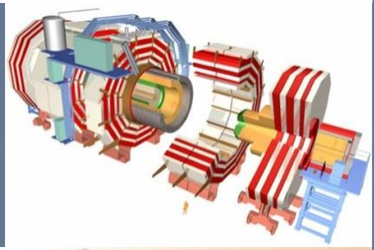

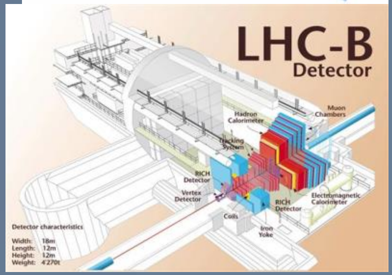

1 CHF ~ 1 USD

Collisions at the LHC: summary

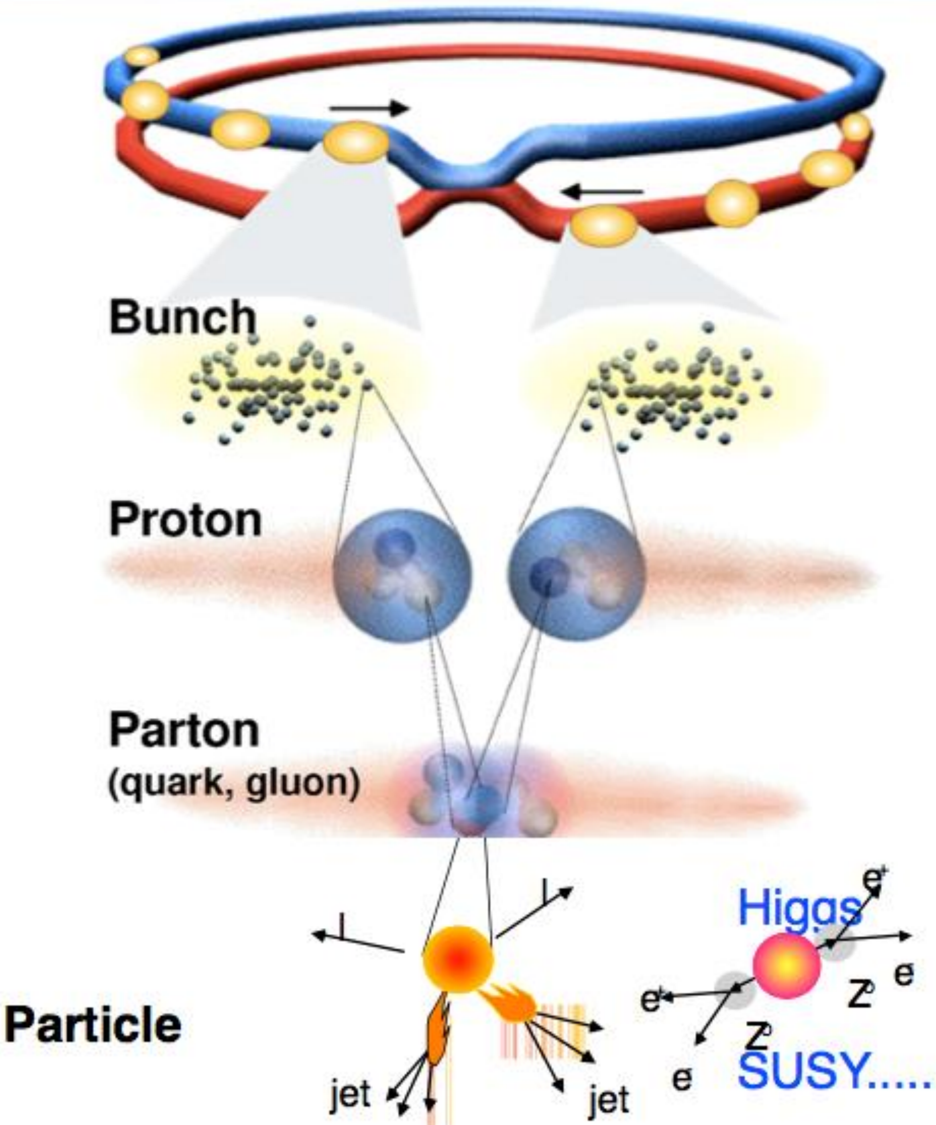

Proton - Proton	2808 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	$10^{34} cm^{-2}s^{-1}$
Crossing rate	40 MHz
Collision rate ≈	10^7-10^9

New physics rate ≈ .00001 Hz

Event selection:
1 in 10,000,000,000,000

Higgs-boson in CMS

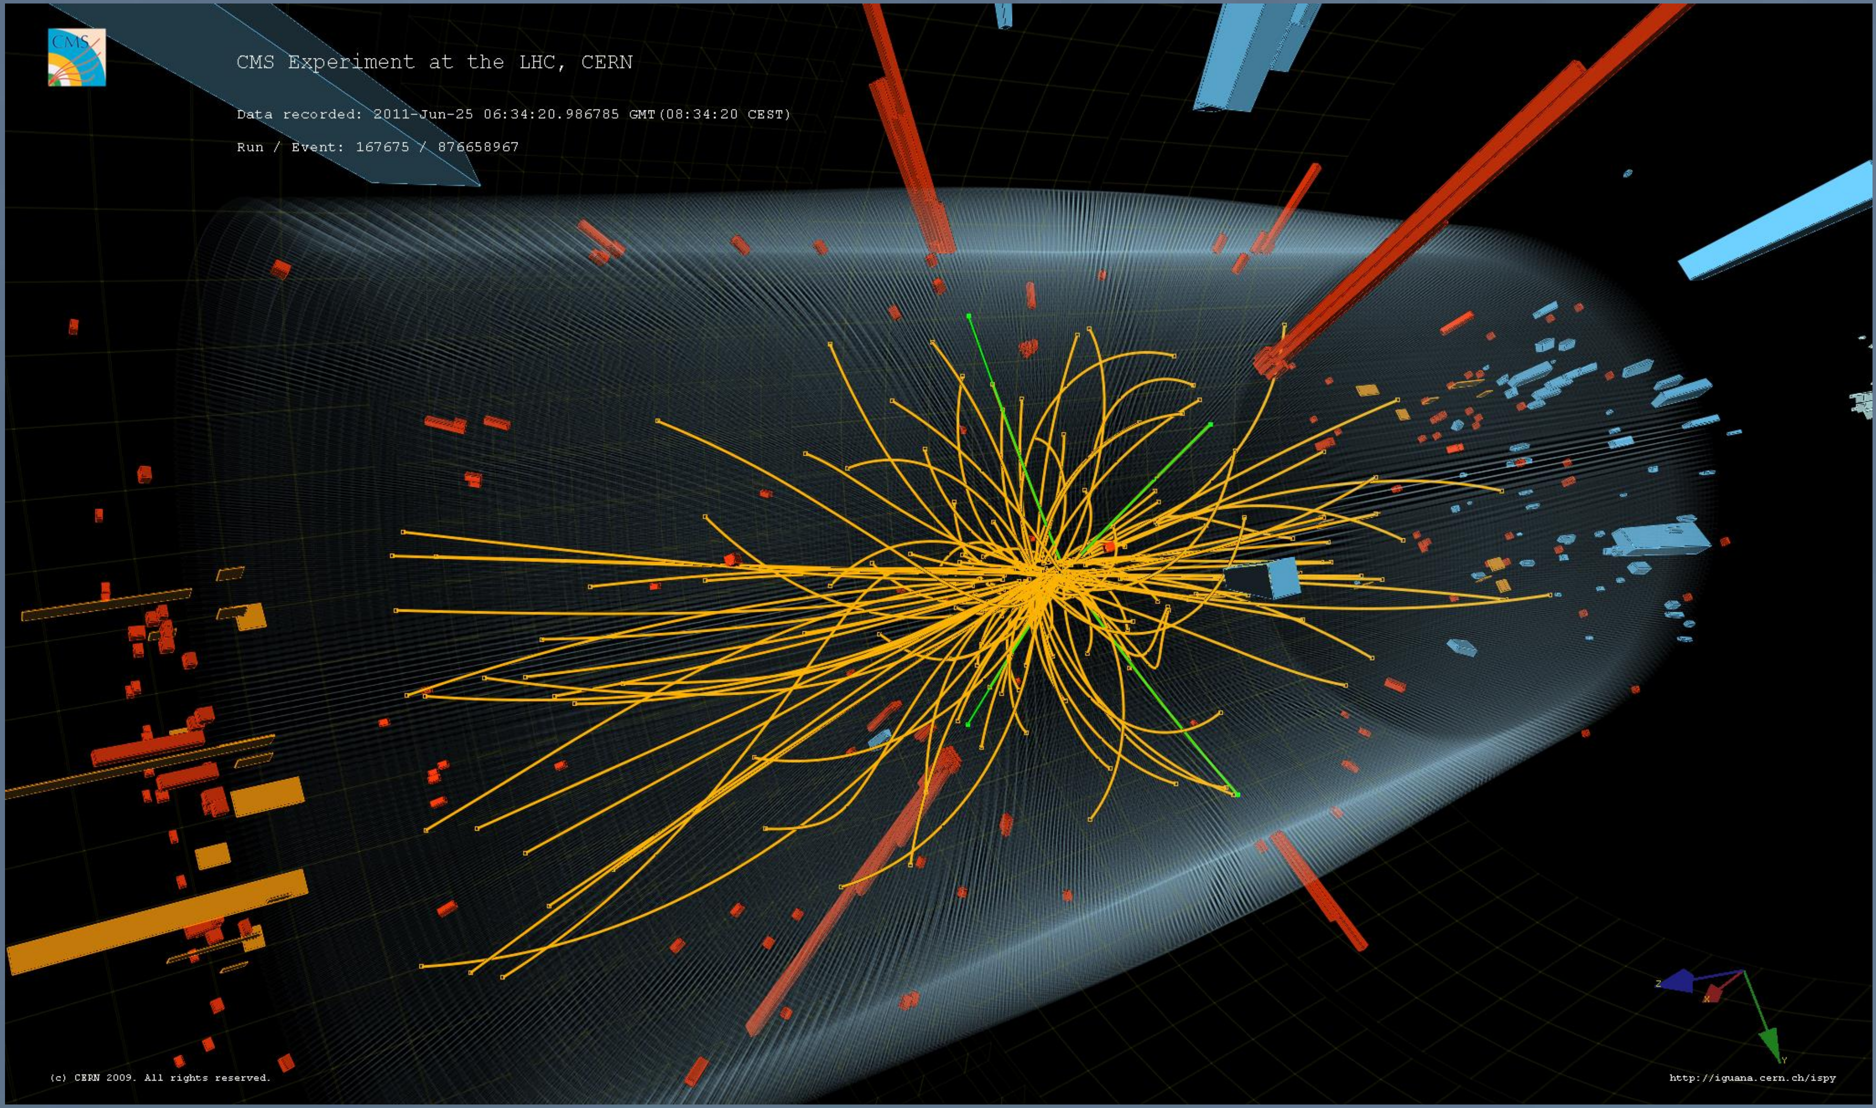

Lead meets lead in ALICE

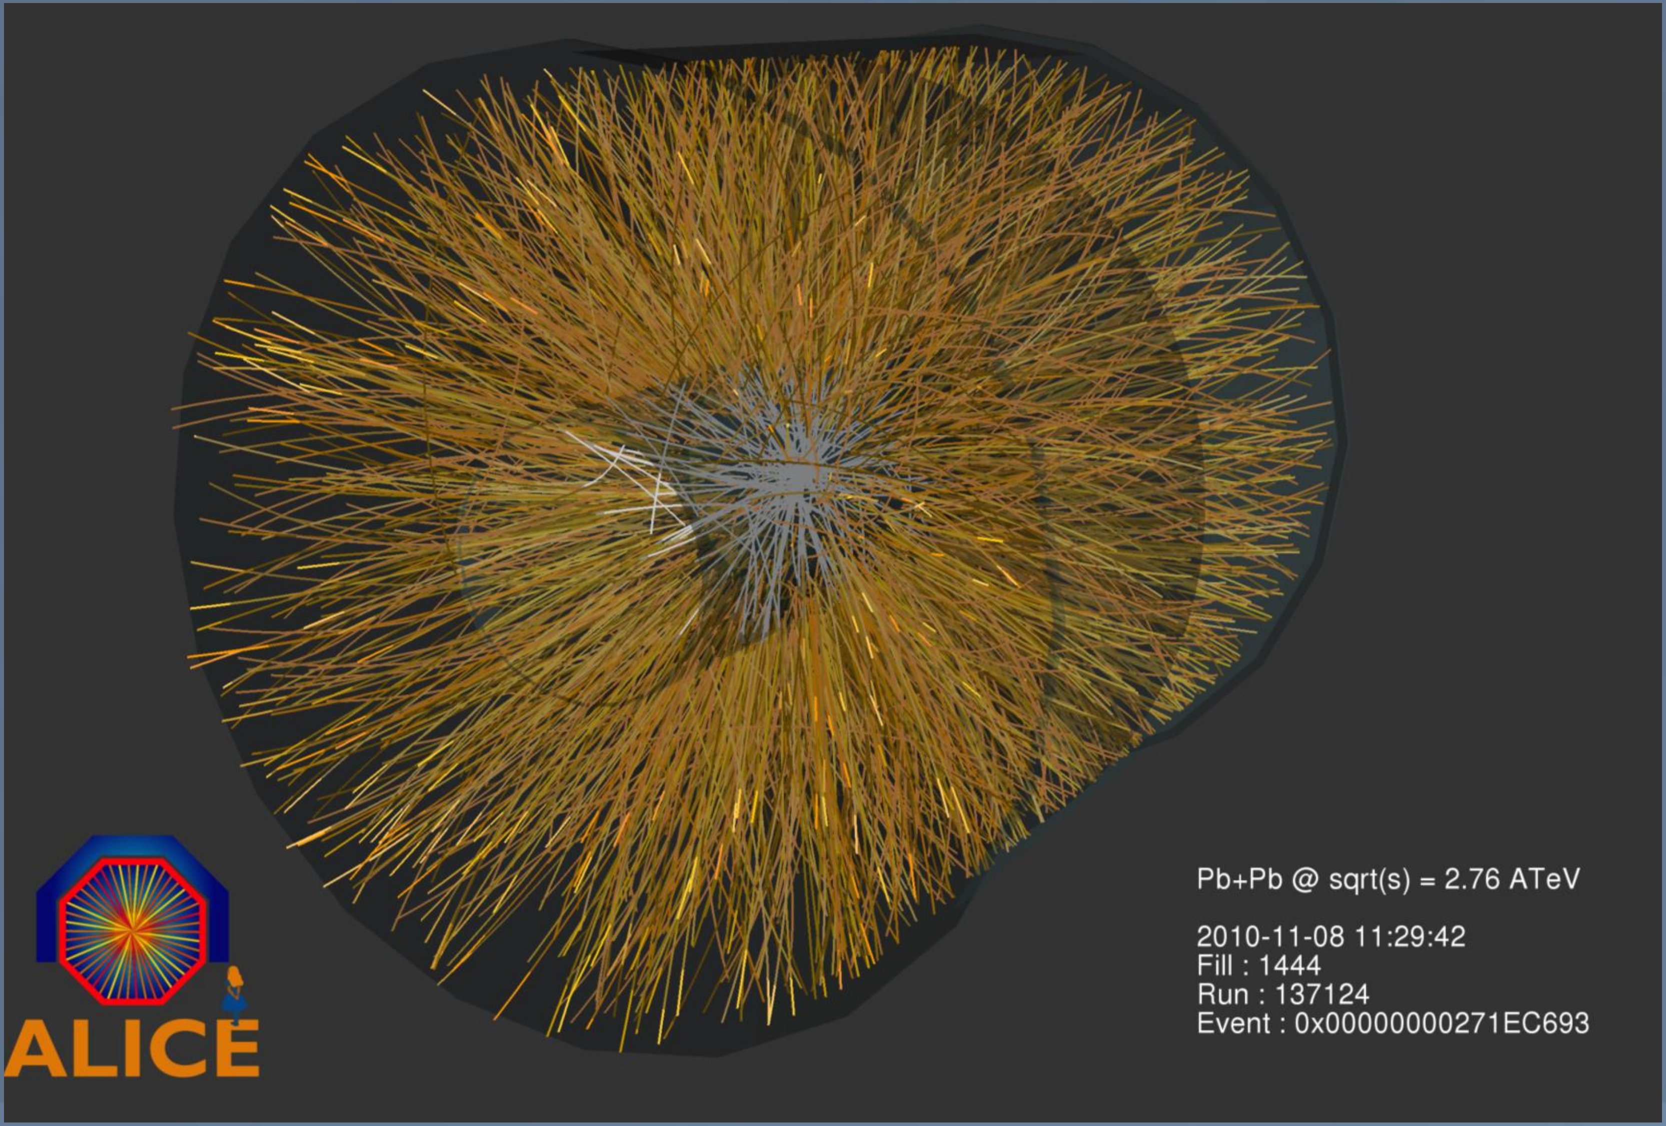

Mr. Higgs'es boson is also in ATLAS

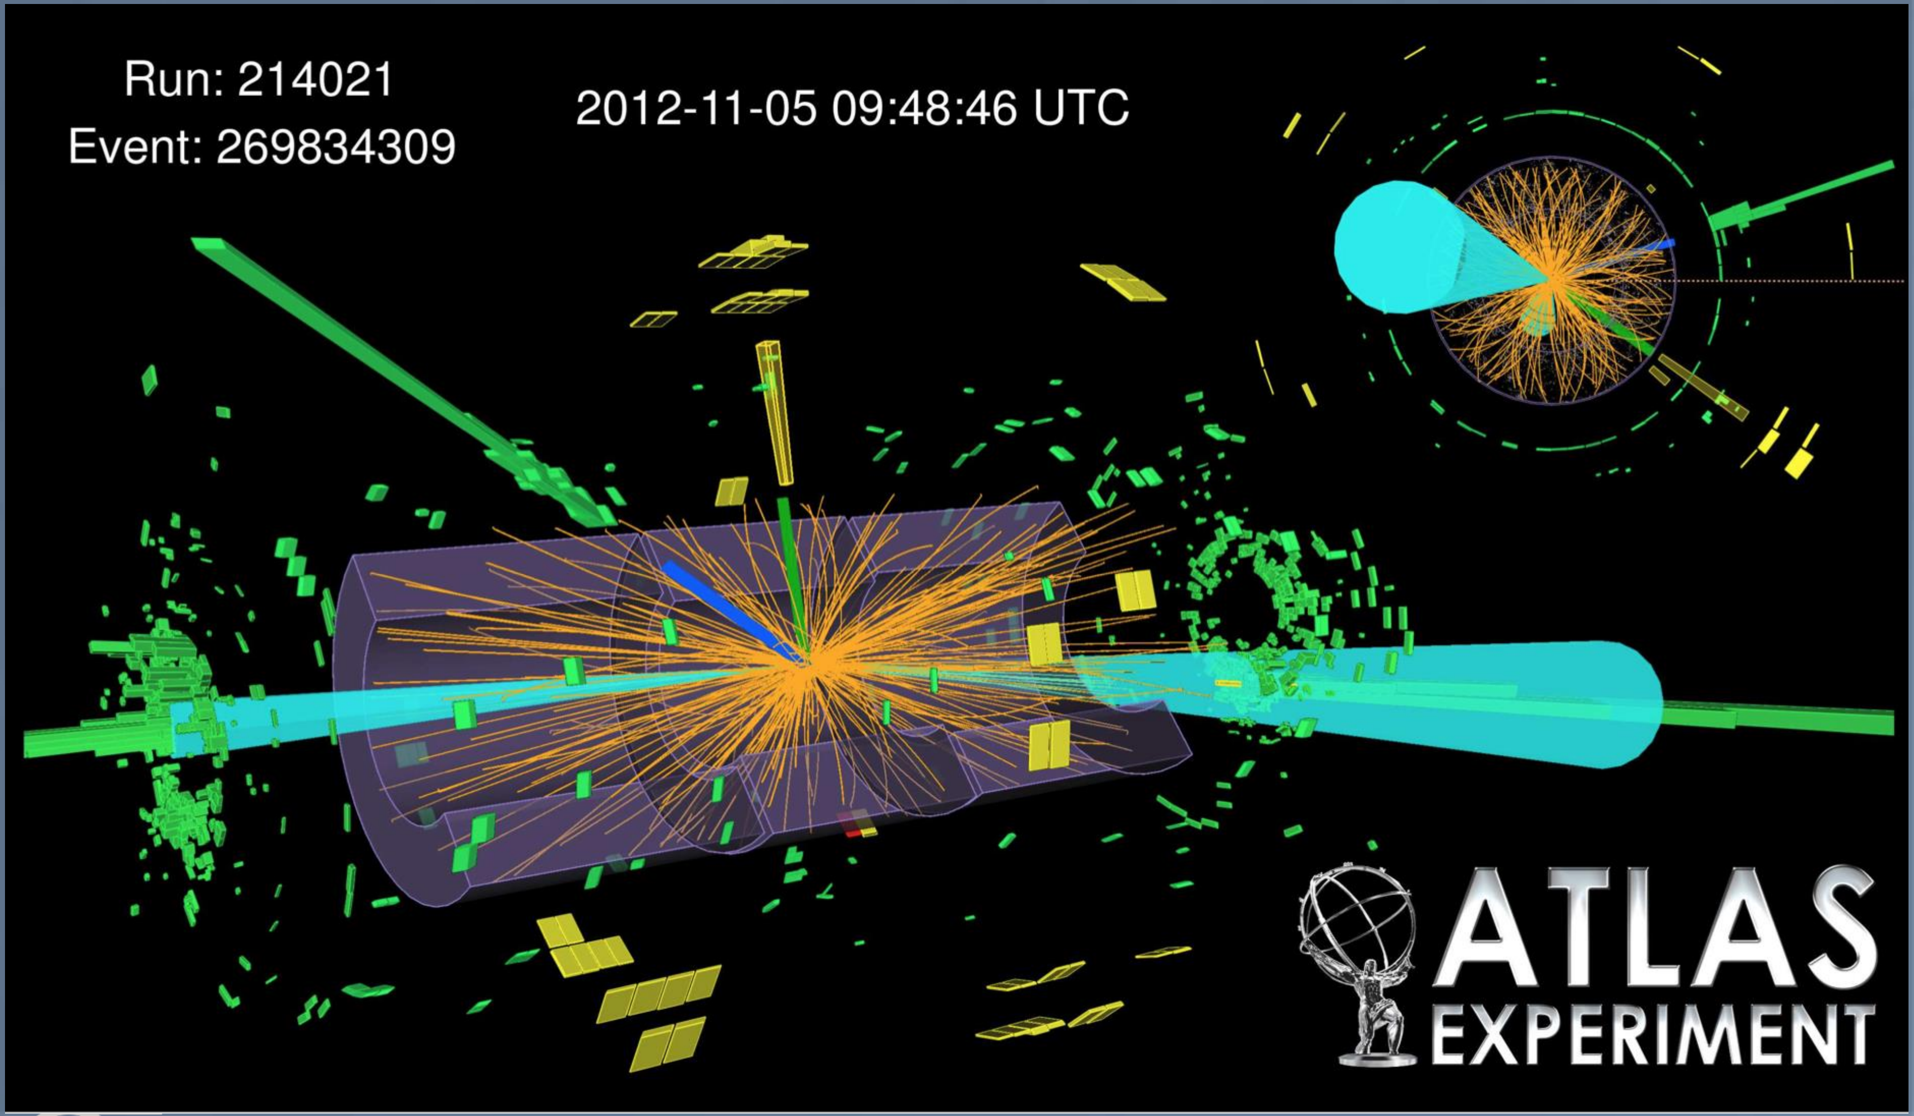

An extremely rare event in LHCb

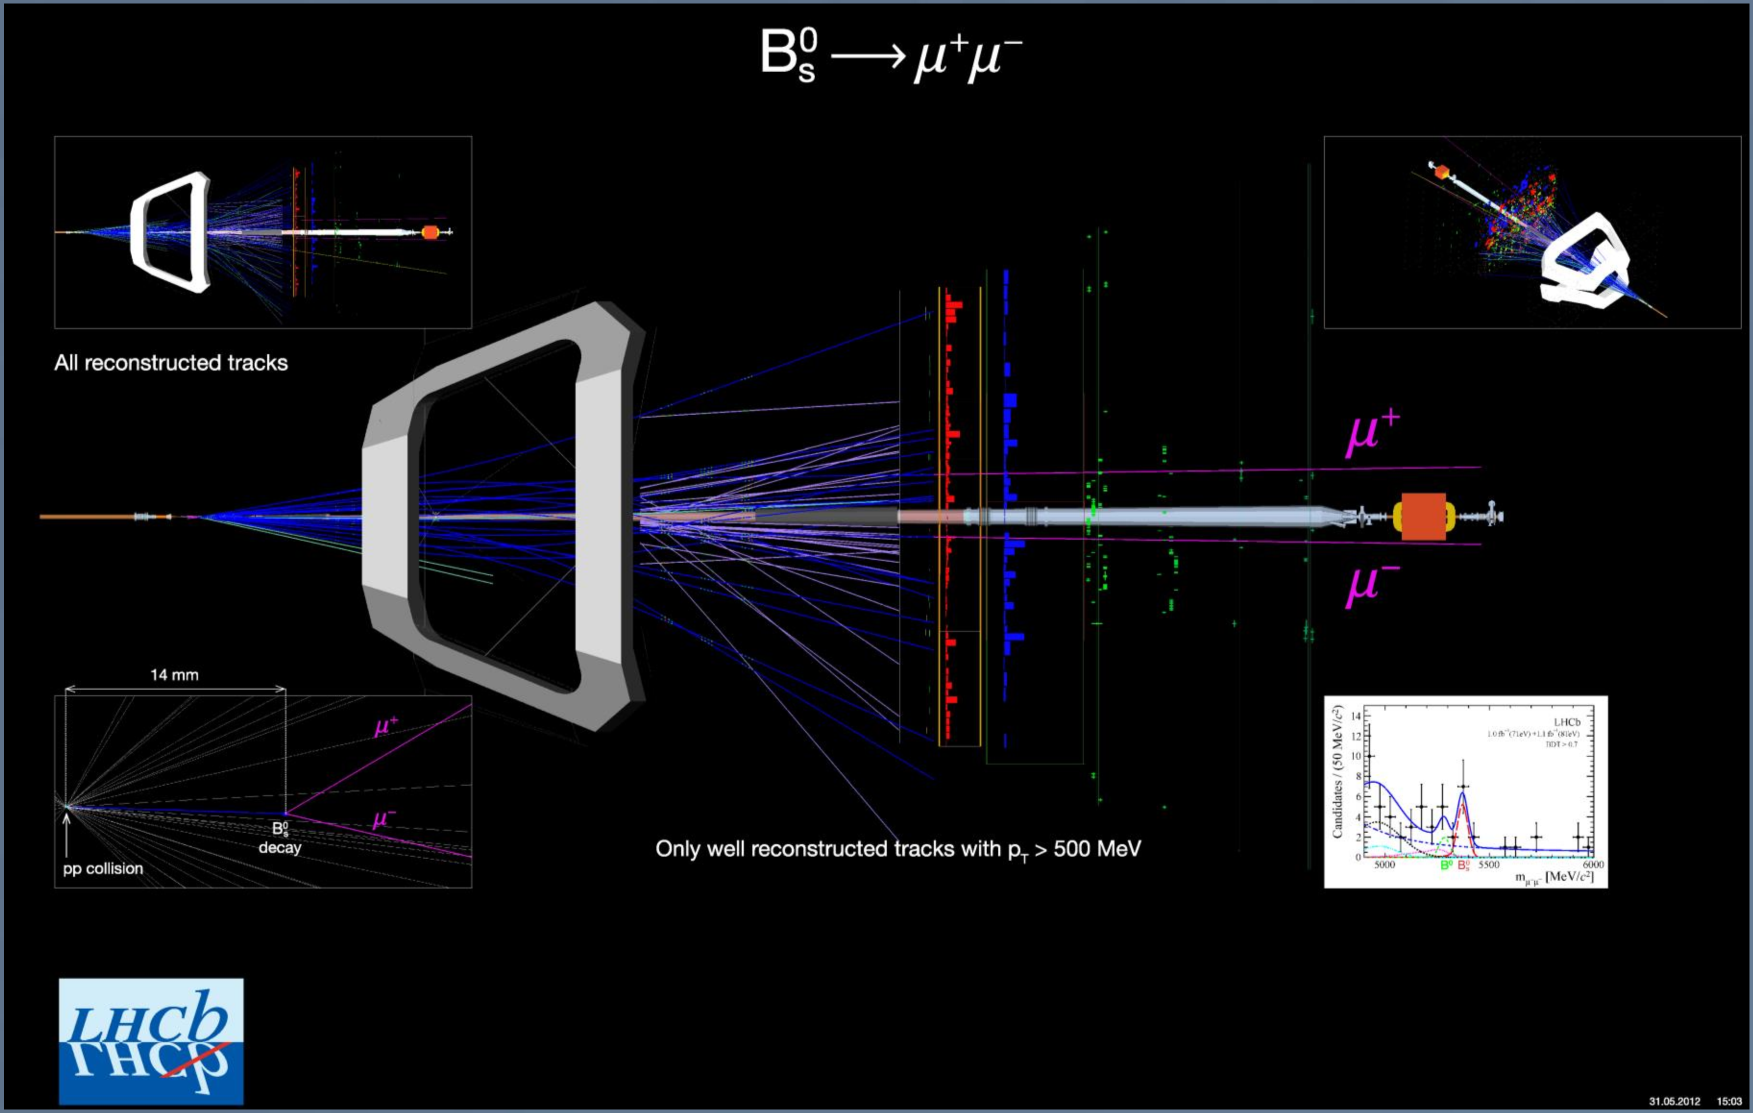

A Complex of accelerators

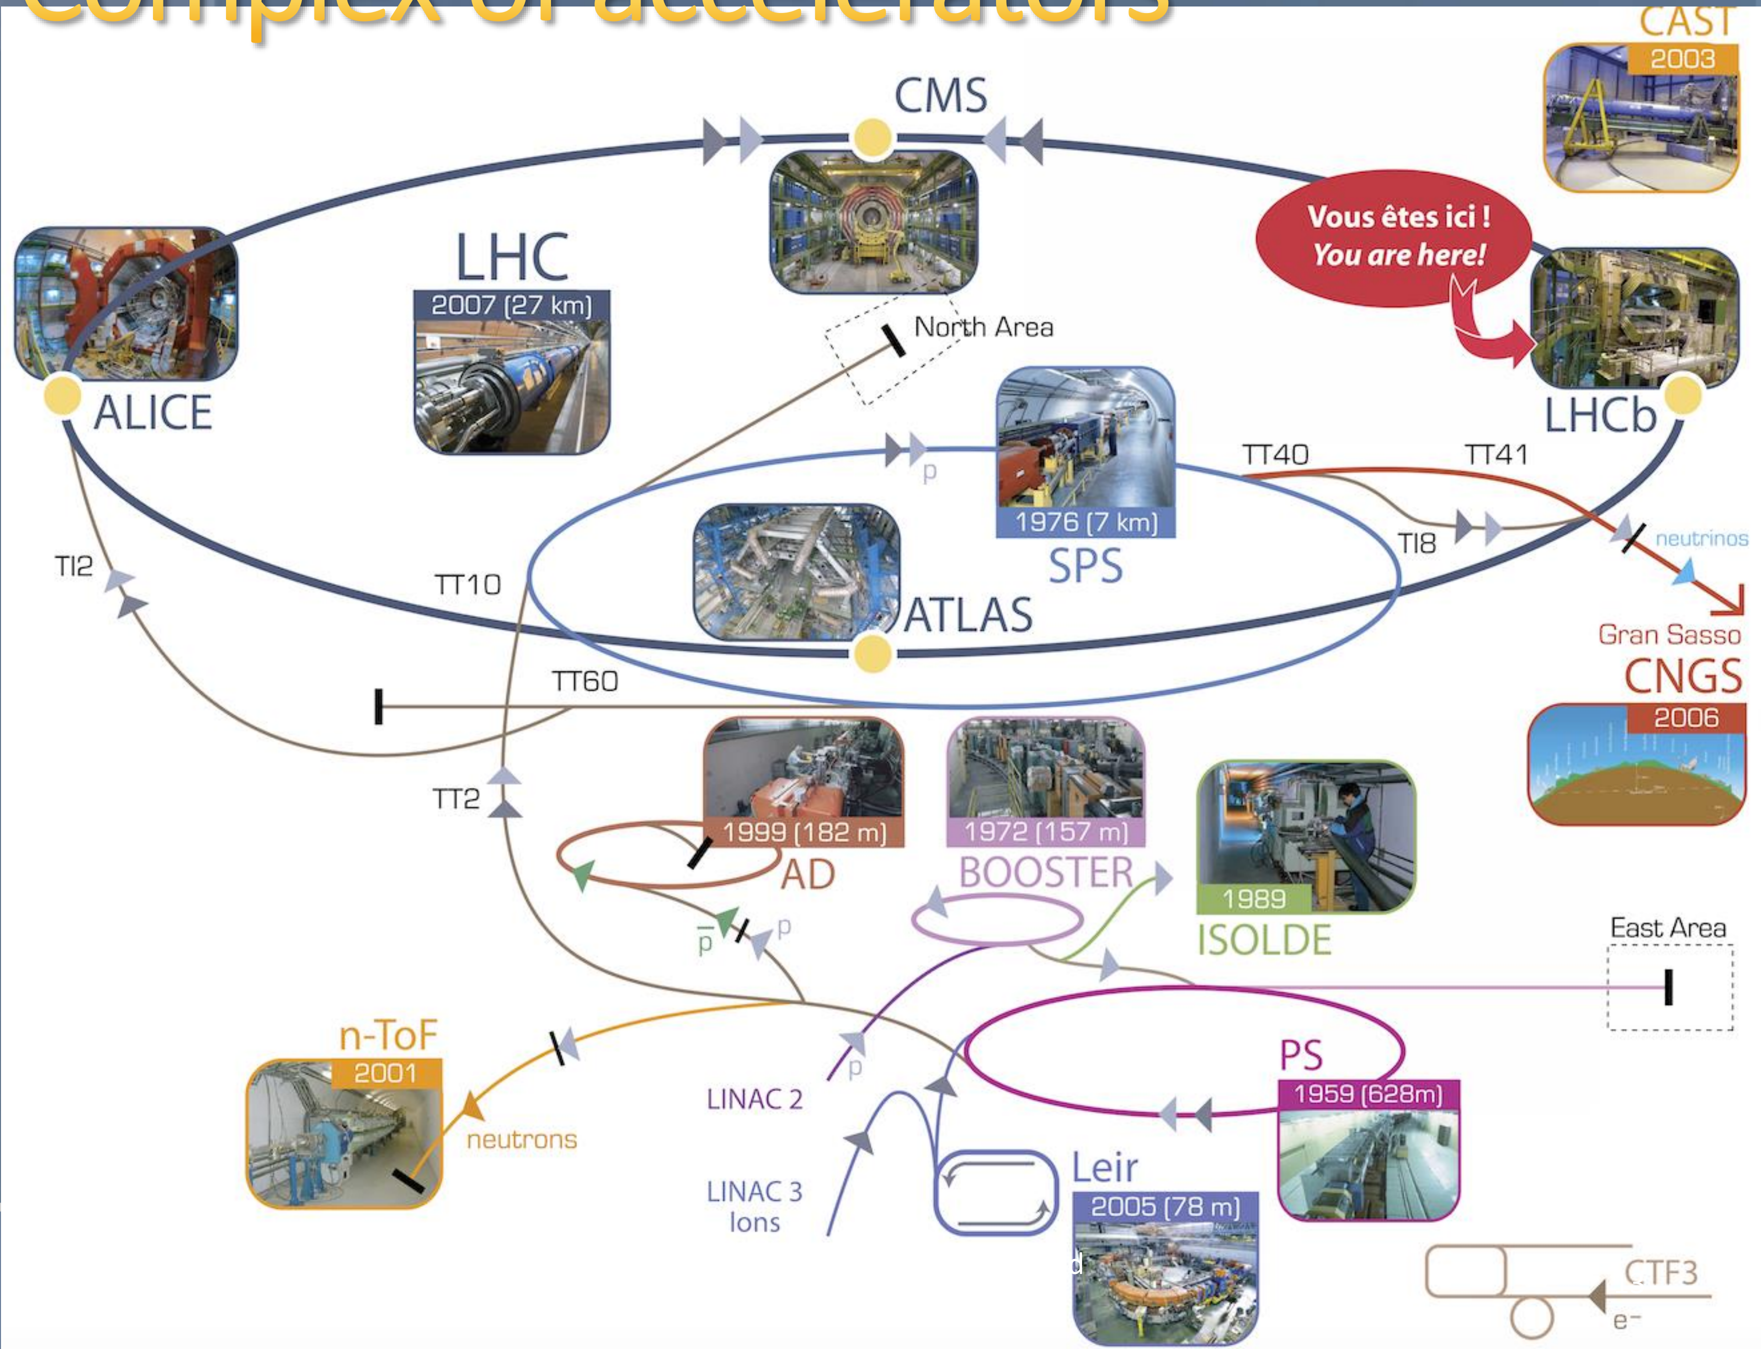

And much more than just the LHC…

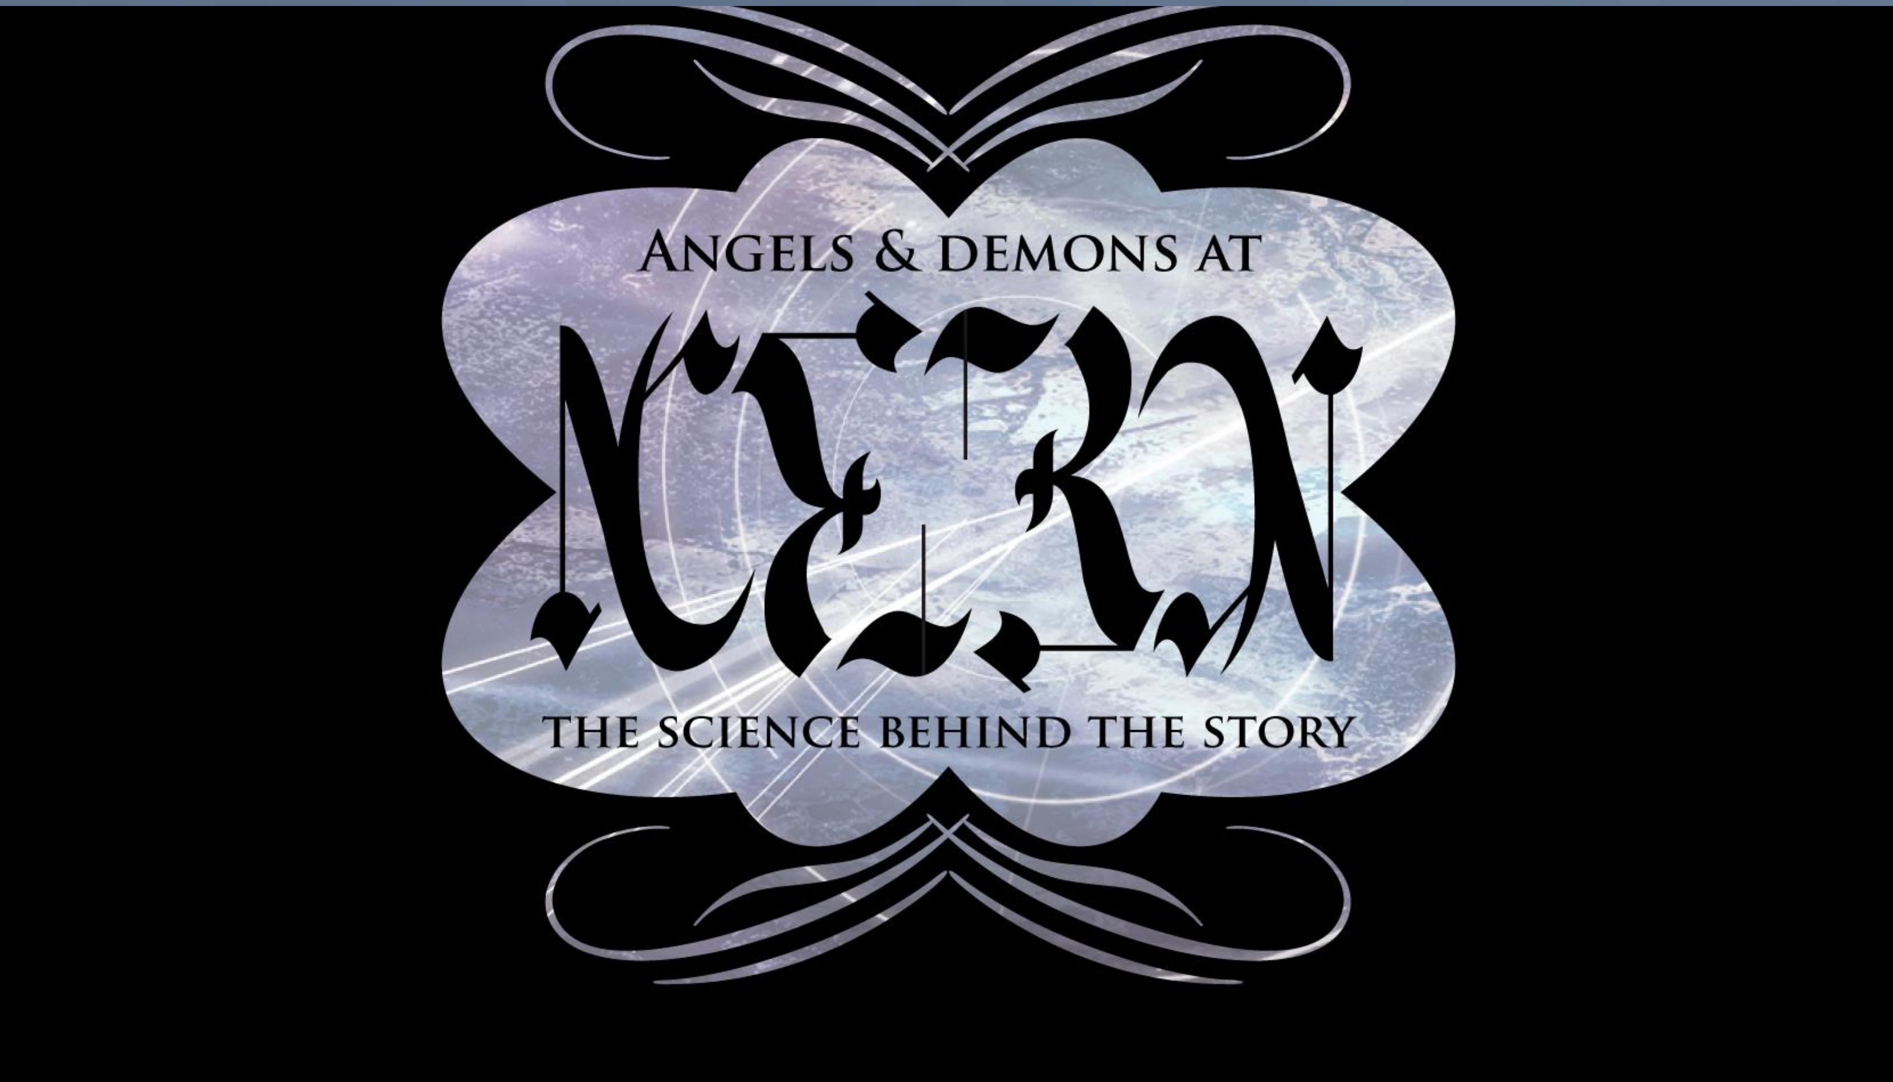

Anti-matter factory

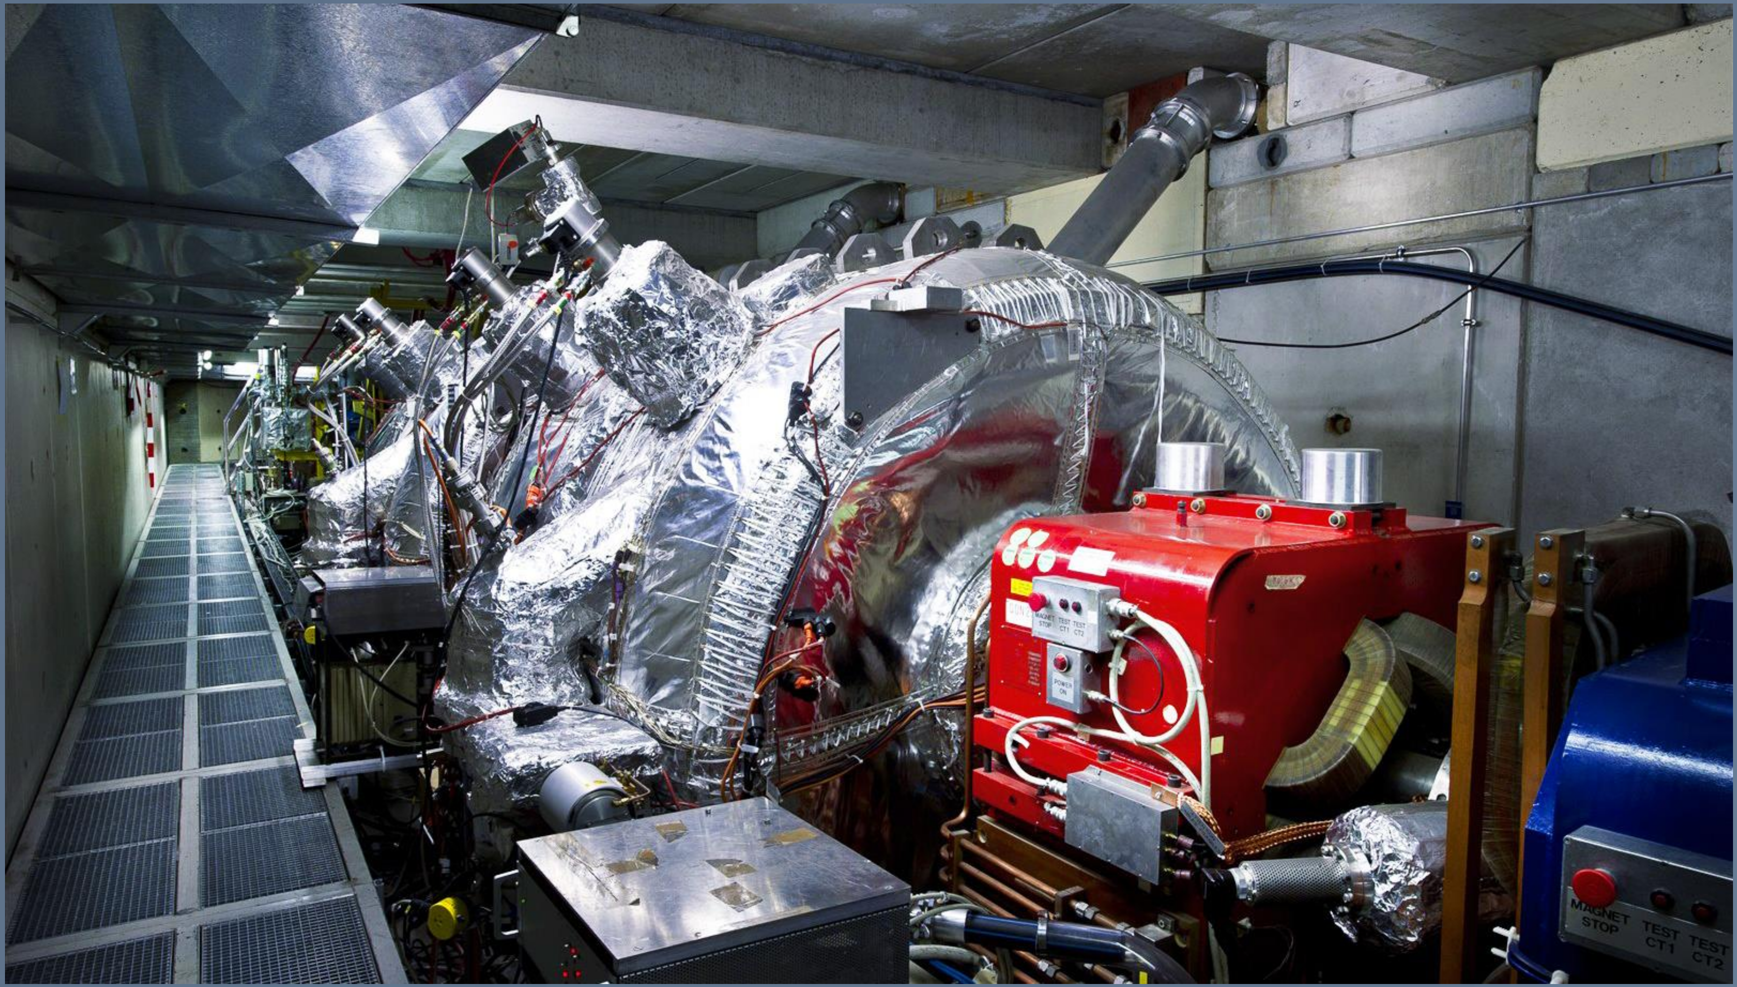

So apart from hunting daemons... What are we doing with the anti-matter?

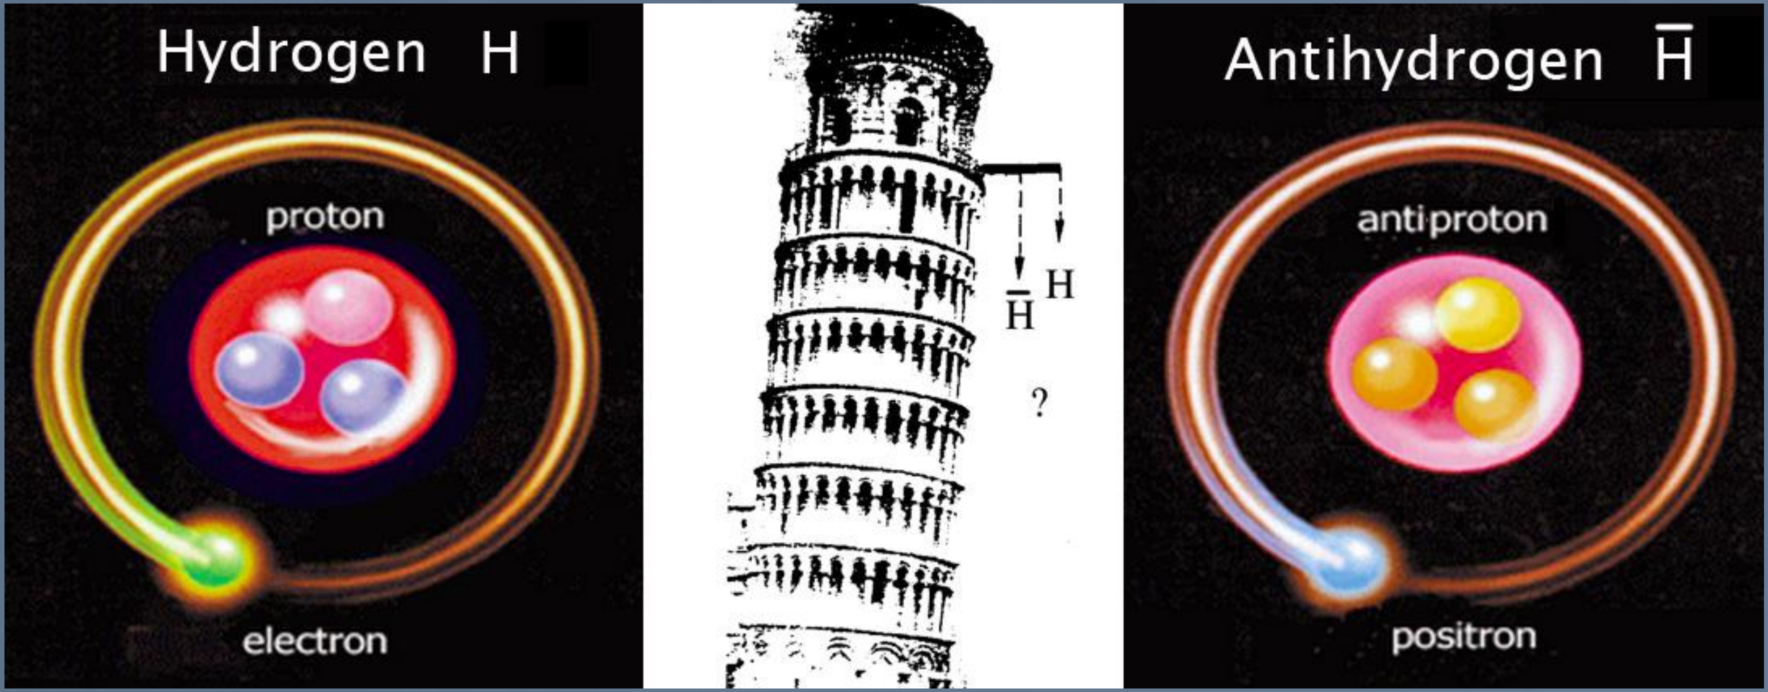

What are these data?

- Raw data:
  - Was a detector element hit?
  - How much energy?
  - What time?
- Reconstructed data:
  - Momentum of tracks (4-vectors)
  - Origin
  - Energy in clusters (jets)
  - Particle type
  - Calibration information
  - …

- 150 Million sensors deliver data … 40 Million times per second
- Up to 6 GB/s to be stored and analysed after filtering

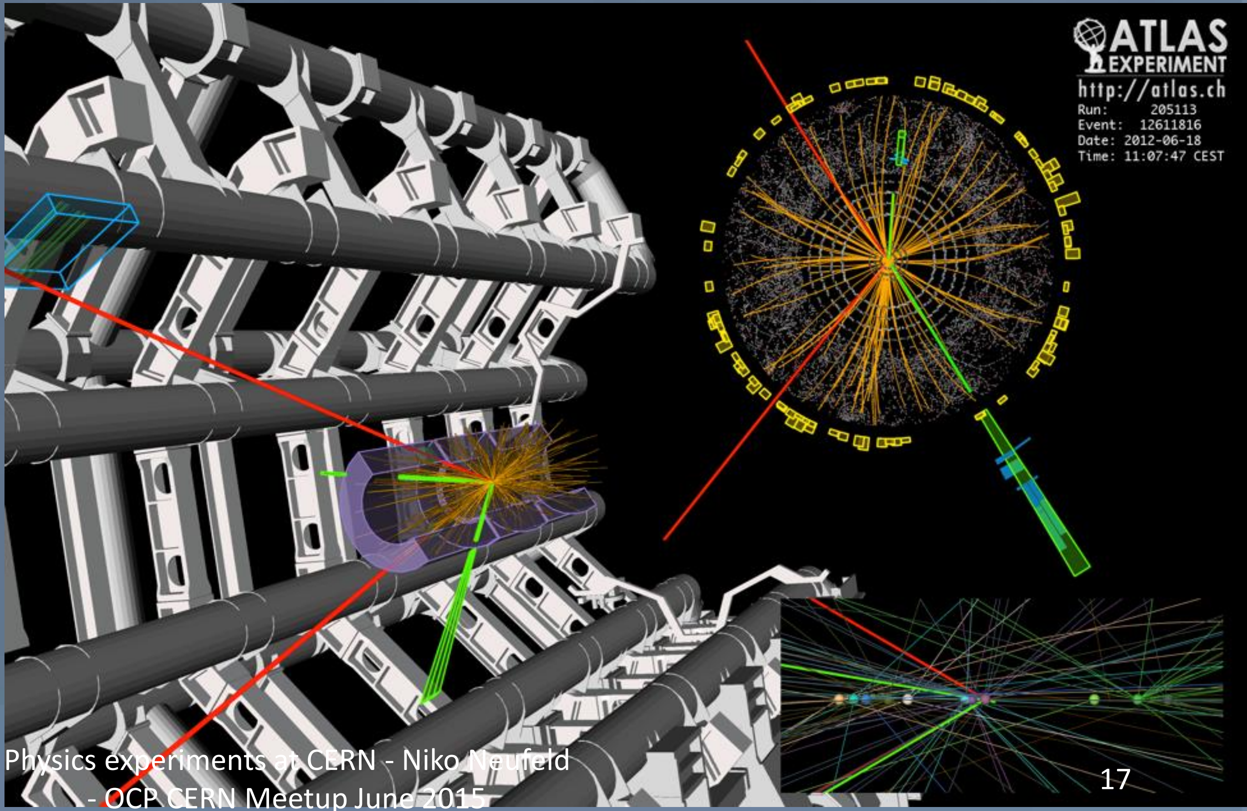

Data and Algorithms

- HEP data are organized as *Events* (particle collisions)
- Simulation, Reconstruction and Analysis programs process “one event at a time”
  - Events are fairly independent → Trivial parallel processing
- Event processing programs are composed of a number of Algorithms selecting and transforming “raw” event data into “processed” (reconstructed) event data and statistics

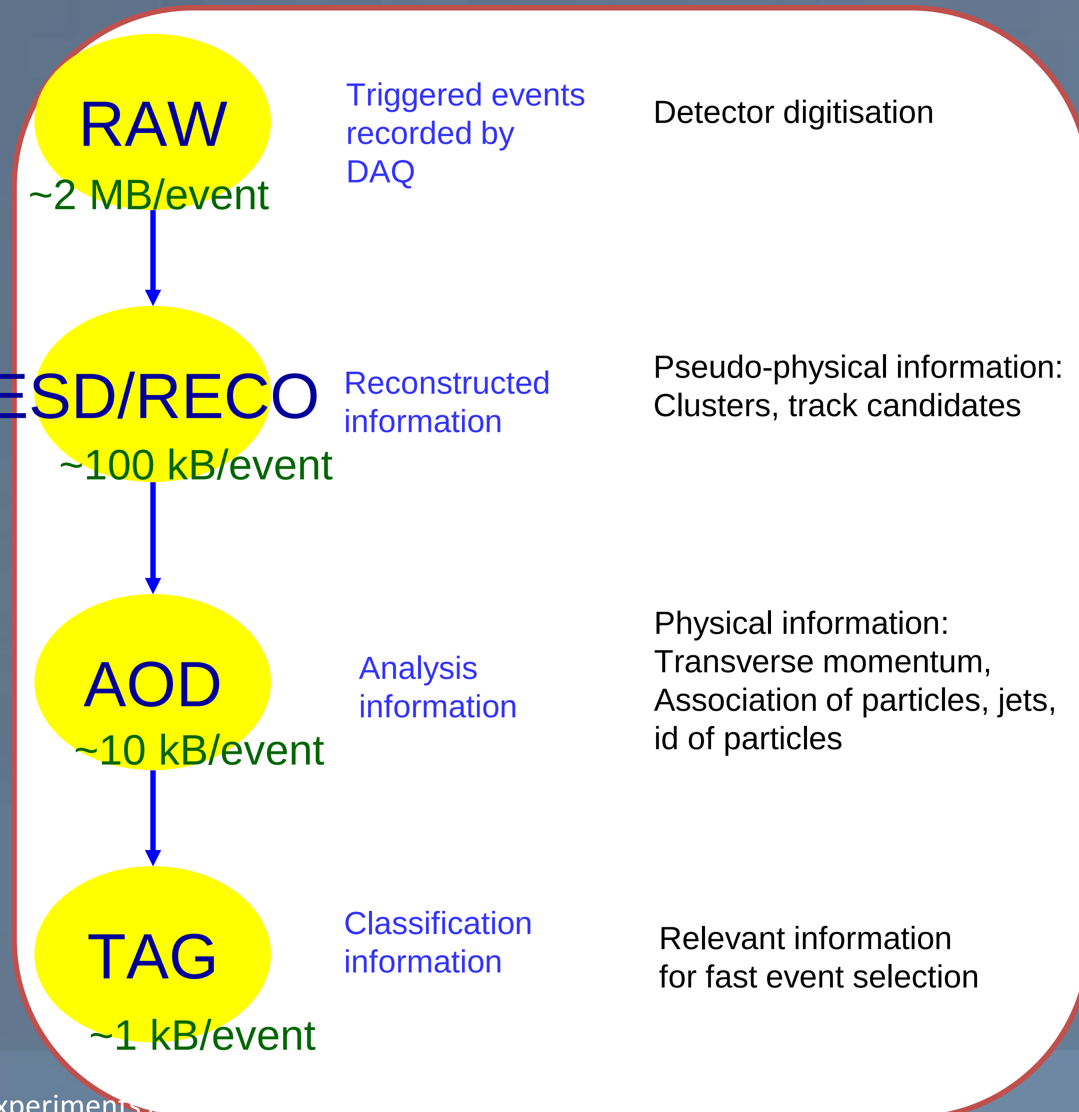

Know Your Enemy: pp Collisions at 14 TeV at 10^{34} $cm^{-2}s^{-1}$

- $\sigma(pp)$ = 70 mb --> >7 x 10^8 /s (!)
- In ATLAS and CMS* 20 – 30 min bias events overlap
- $H \to ZZ$
  $Z \to \mu\mu$
  $H \to$ 4 muons: the cleanest ("golden") signature

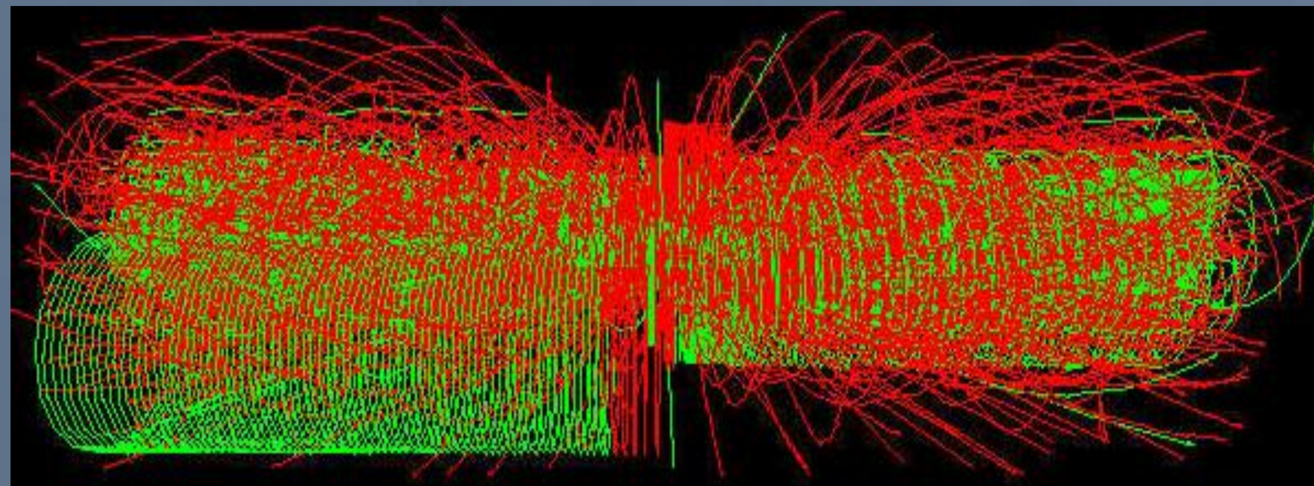

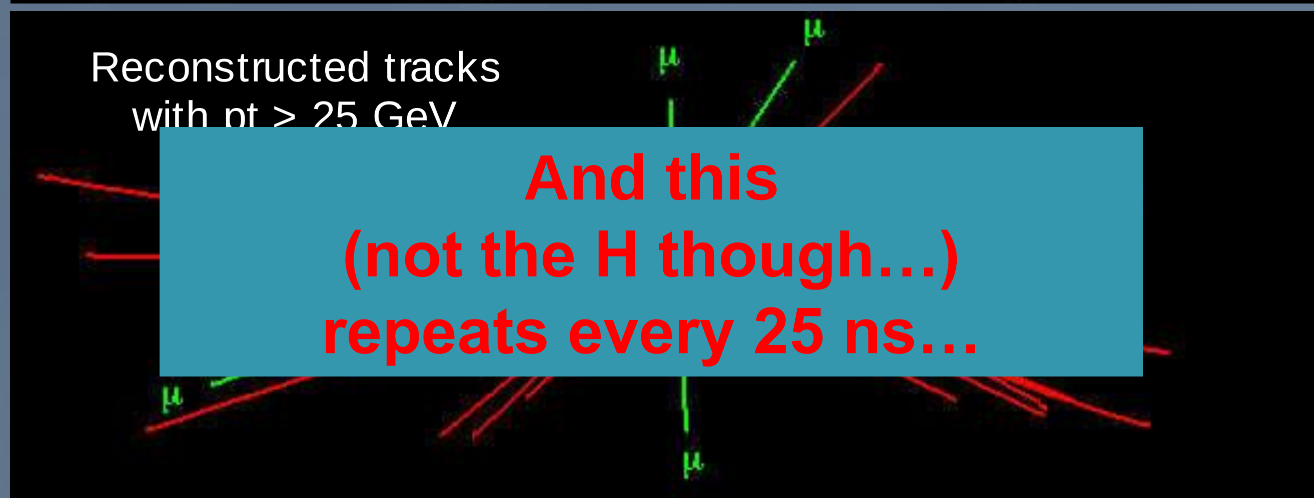

And this
(not the H though…)
repeats every 25 ns…

*)LHCb @4x10^{33} cm^{-2}-1 isn't much nicer and in Alice (PbPb) is even more busy

What's mother nature's menu?

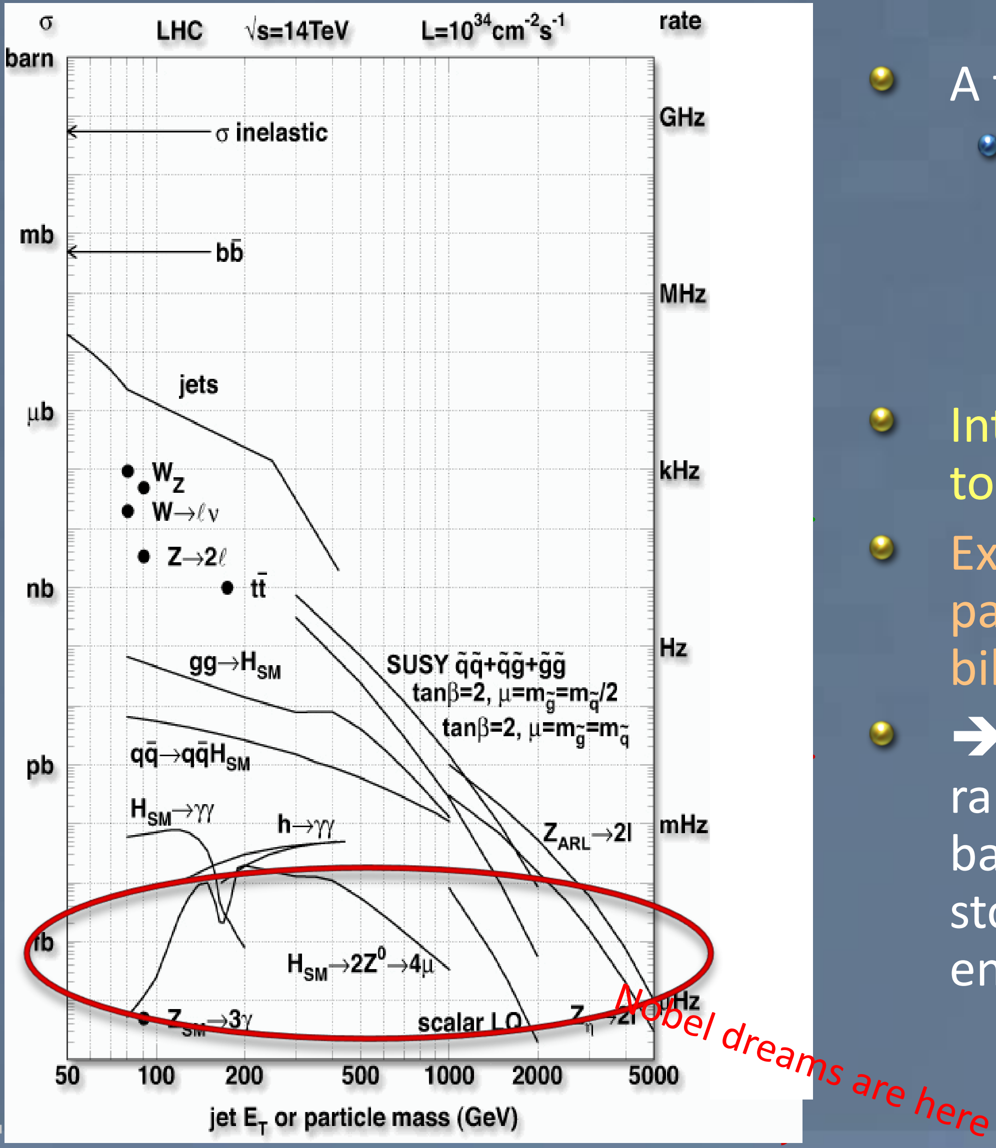

- A typical collision is "boring"
  - Although we need also some of these "boring" data as cross-check, calibration tool and also some important "low-energy" physics
- Interesting physics: one in a million to one in 100 million collisions
- Exciting physics involving new particles/discoveries: one in a billion or less
- ➔ Need to efficiently identify these rare processes in the overwhelming background, such that we need to store only a "smaller" highly enriched subset of the raw data

Data Rates

- Particle beams cross every 25 ns (40 MHz)
  - Up to 25 particle collisions per beam crossing
  - Up to 10^9 collisions per second
- Basically 2 event filter/trigger levels
  - 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
$W \rightarrow e\upsilon$; $W \rightarrow \mu\upsilon$; $W \rightarrow \tau\upsilon$	20
$Z \rightarrow ee$; $Z \rightarrow \mu\mu$; $Z \rightarrow \tau\tau$	2
t	1
Higgs boson (all; m_H = 120GeV)	0.04
Higgs boson (simple signatures)	0.0003
Black Hole (certain properties)	0.0001

	Incoming data rate	Outgoing data rate	Reduction factor
Level1 Trigger (custom hardware)	40000000 s^{-1}	10^5 – 10^6 s^{-1}	400-10,000
High Level Trigger (software on server farms)	2000-1000000 s^{-1}	1000 -10000 s^{-1}	10-2000

LHC planning

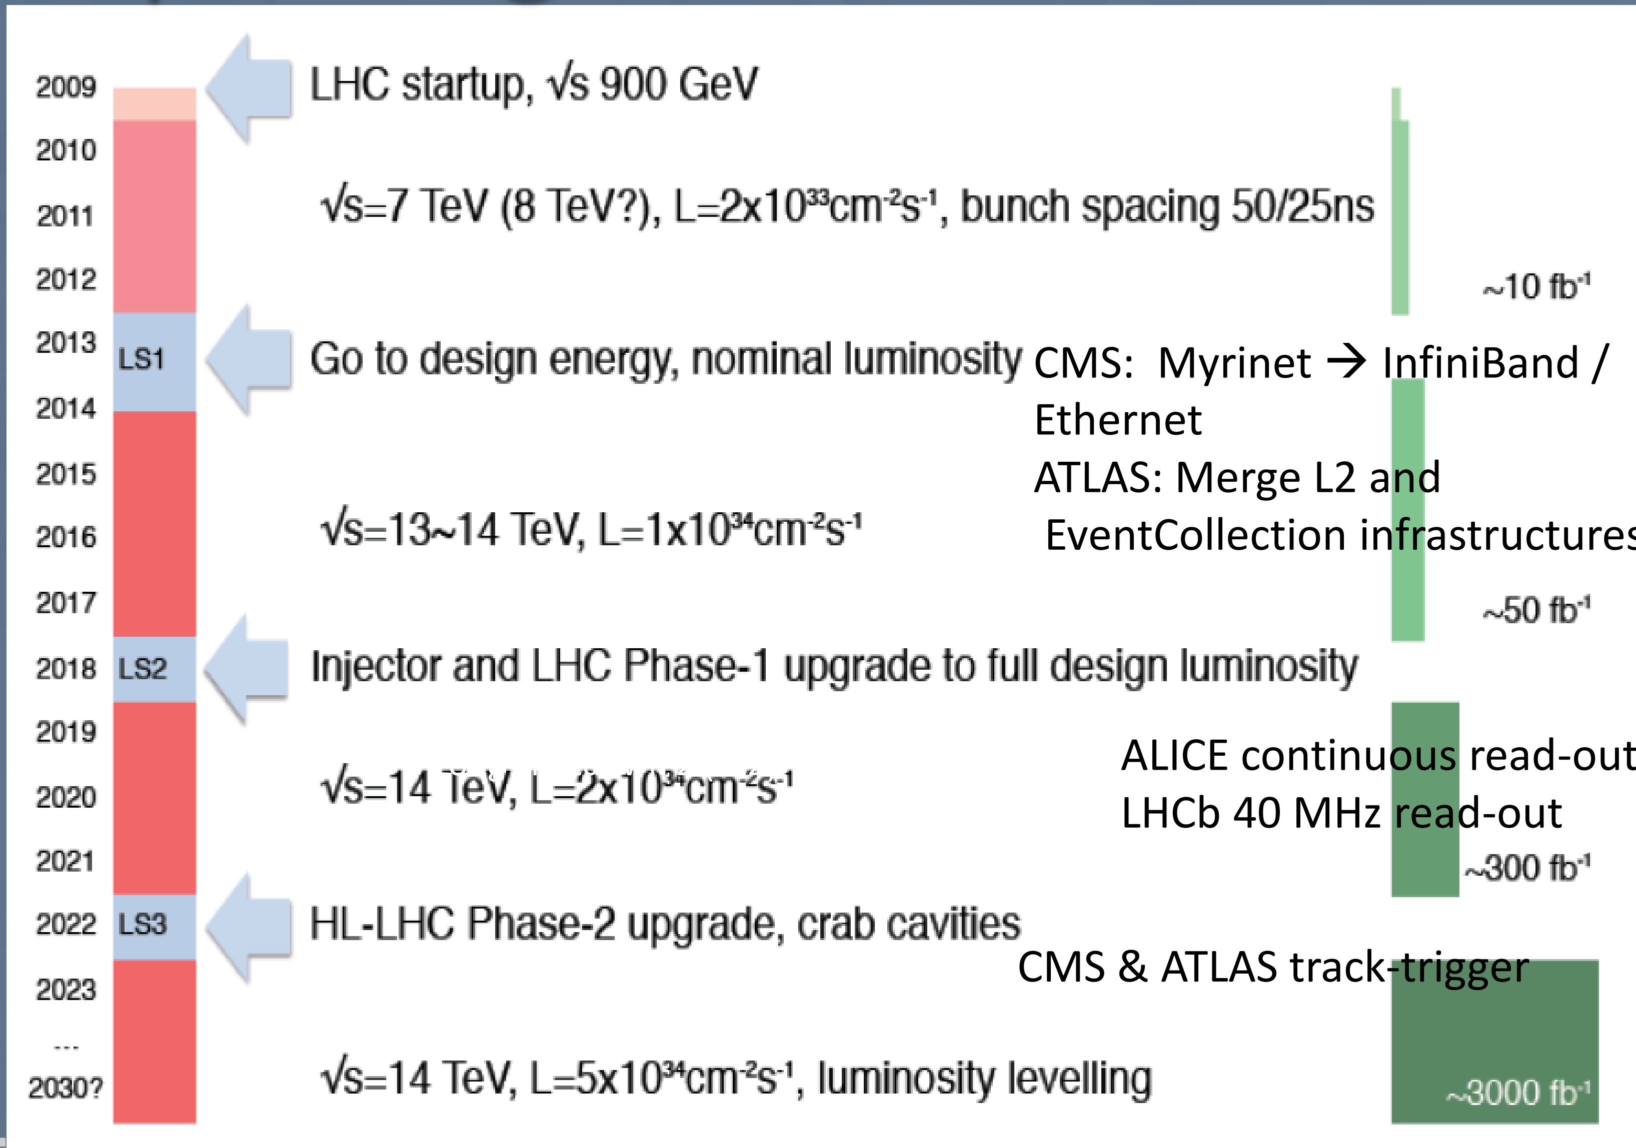

Defeating the data-torrent

1. Thresholding and tight encoding
2. Real-time selection based on partial information
3. Final selection using full information of the collisions

These selection systems are called “Triggers” in high energy physics

Challenge #1
The first level trigger

Selection based on partial information

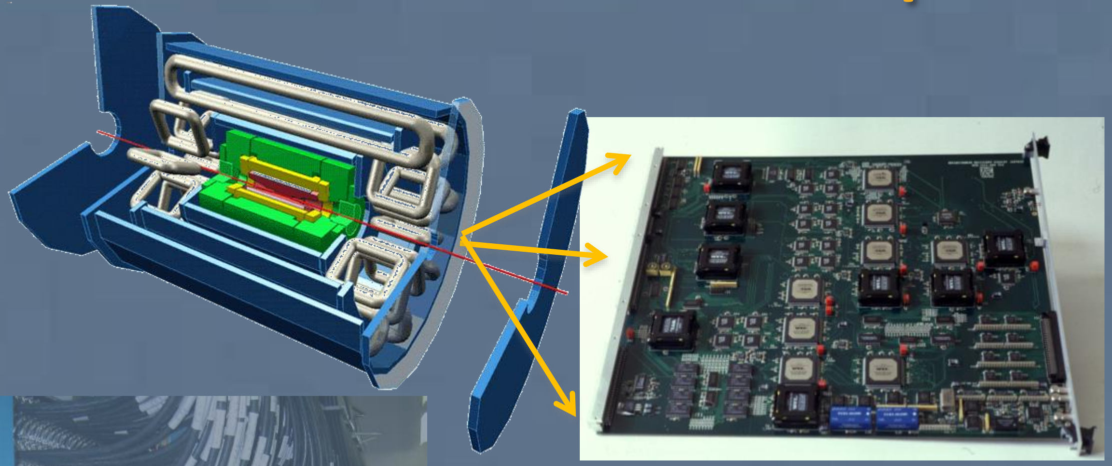

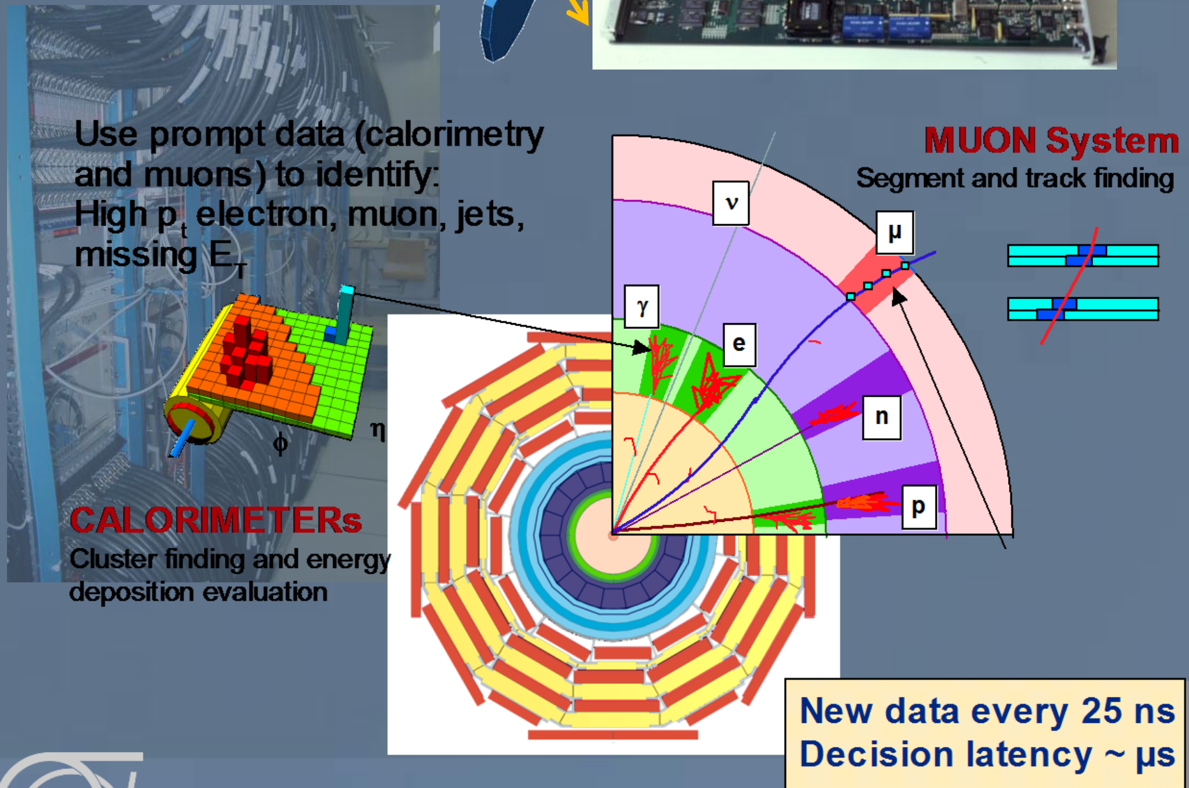

- A combination of (radiation hard) ASICs and FPGAs process data of "simple" sub-systems with "few" O(10000) channels in real-time
  - Other channels need to buffer data on the detector
- this works only well for "simple" selection criteria
- long-term maintenance issues with custom hardware and low-level firmware
- crude algorithms miss a lot of interesting collisions

Level 1 Trigger

- The Level 1 Trigger is implemented in hardware: FPGAs and ASICs → difficult / expensive to upgrade or change, maintenance by experts only
- Decision time: ~ a small number of microseconds
- It uses simple, hardware-friendly signatures → looses interesting collisions
- Each sub-detector has its own solution, only the uplink is standardized →

Finding Muons (2d view)

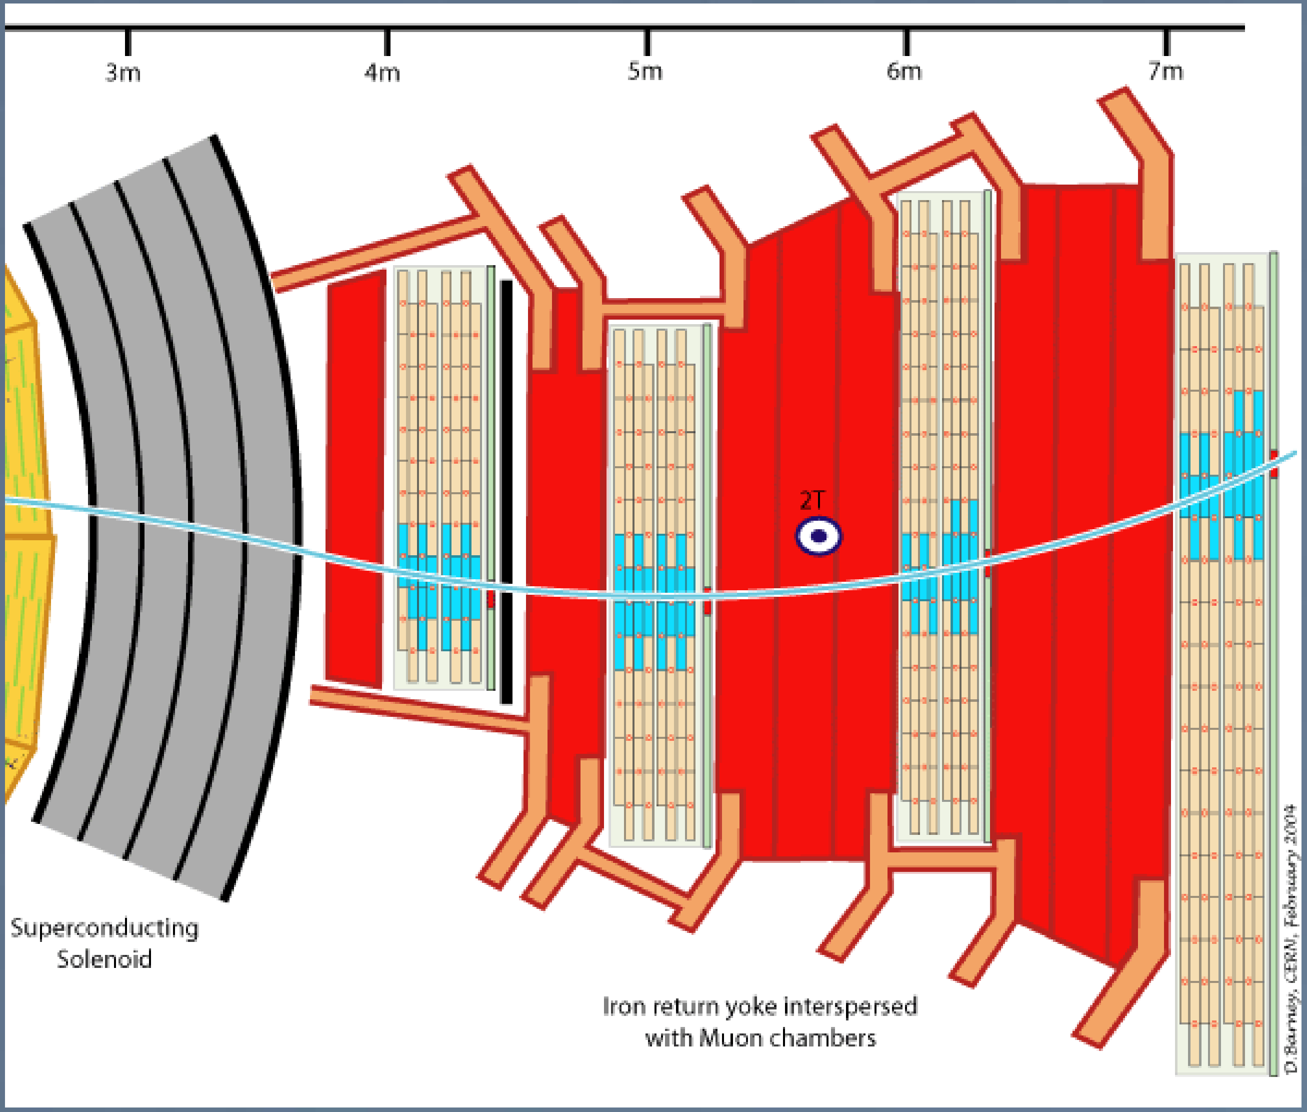

Calorimeter data

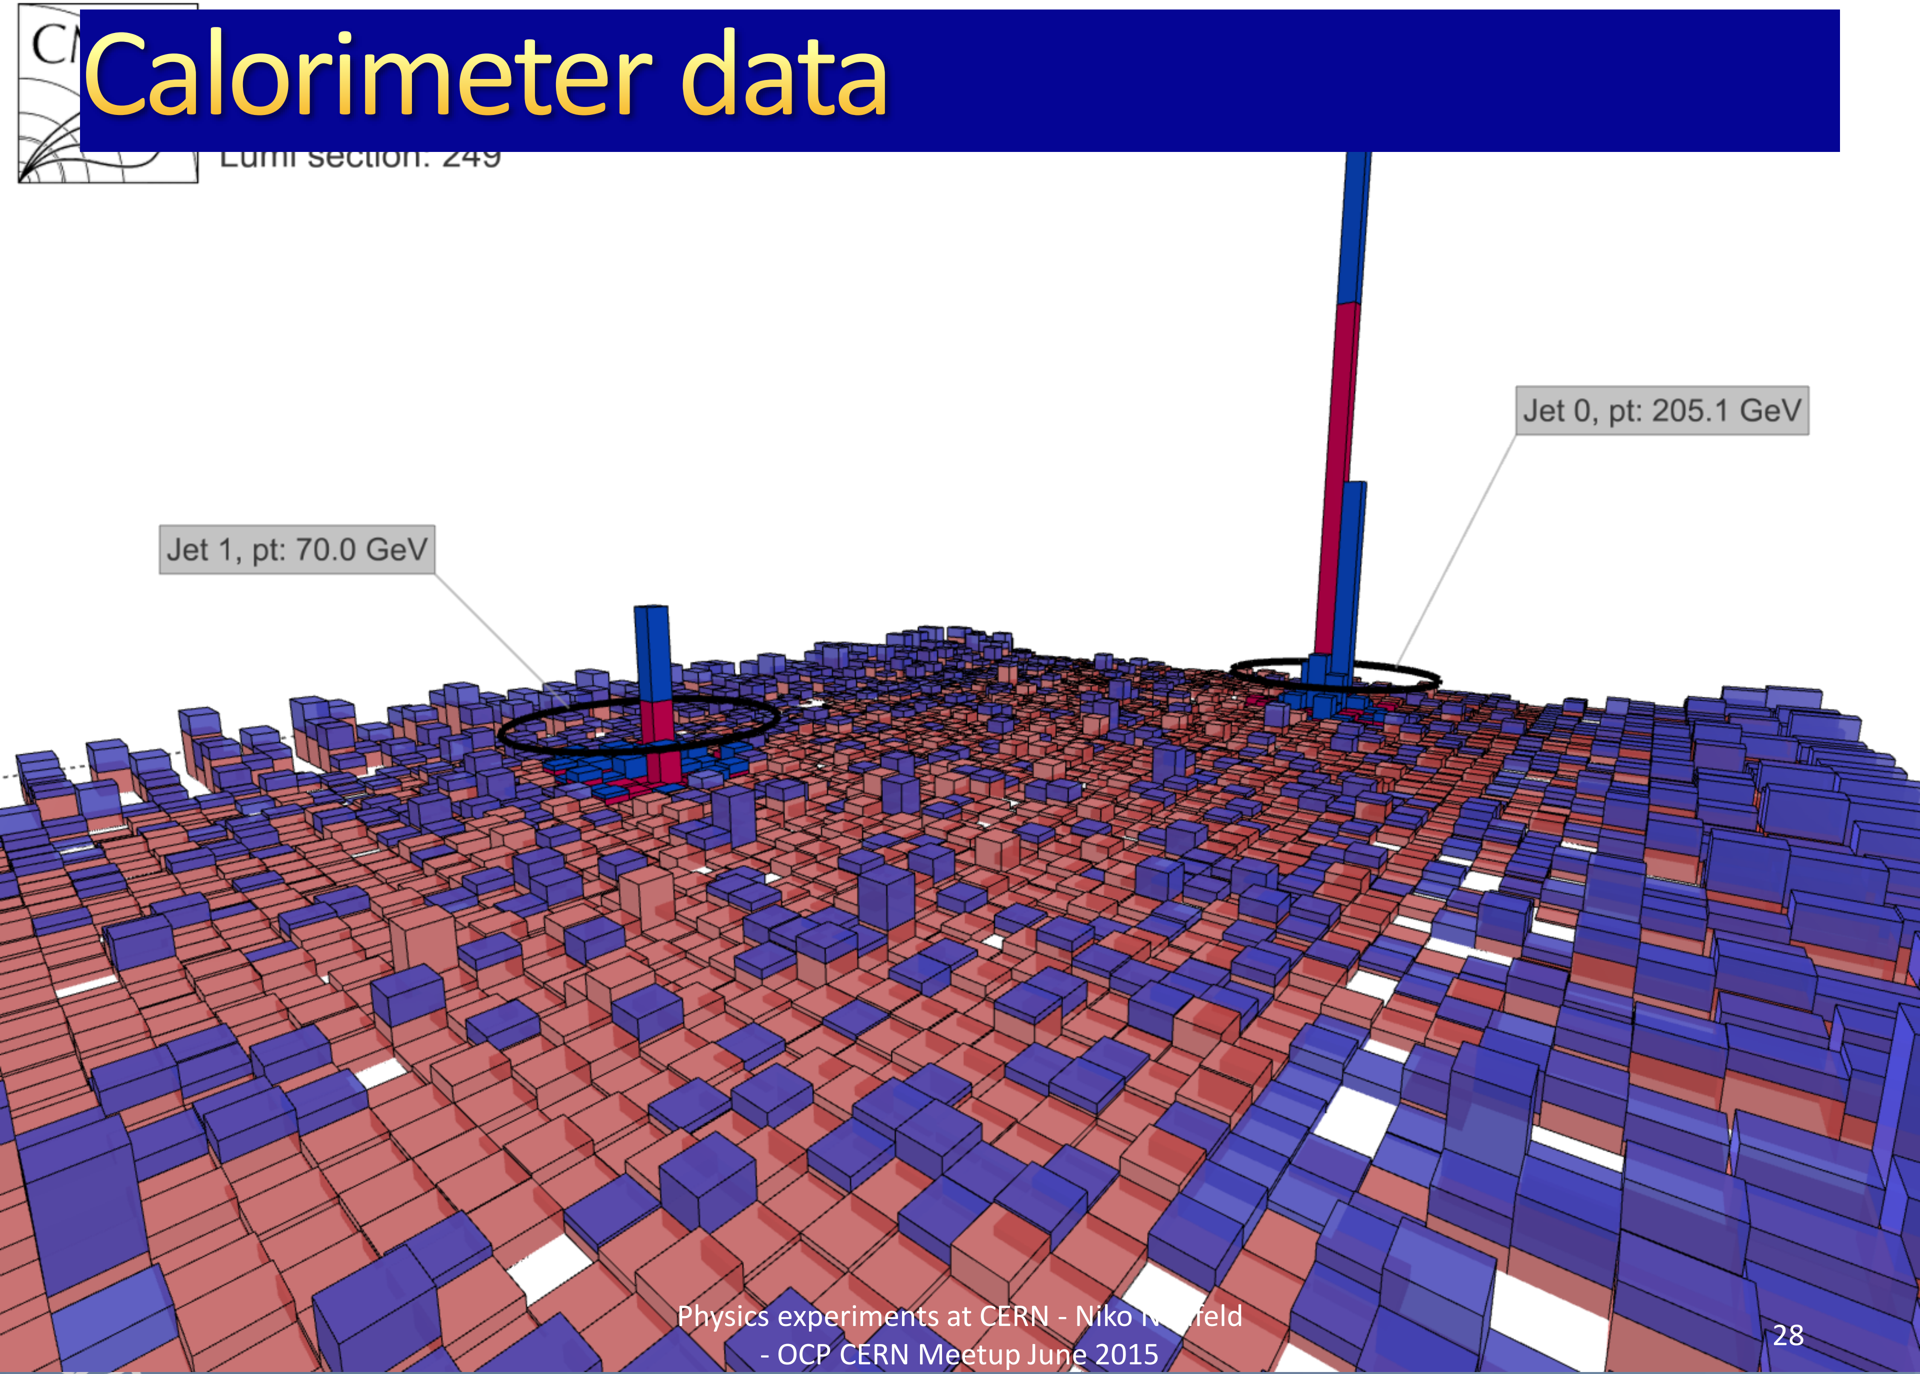

A Track-Trigger at 40 MHz 2020++

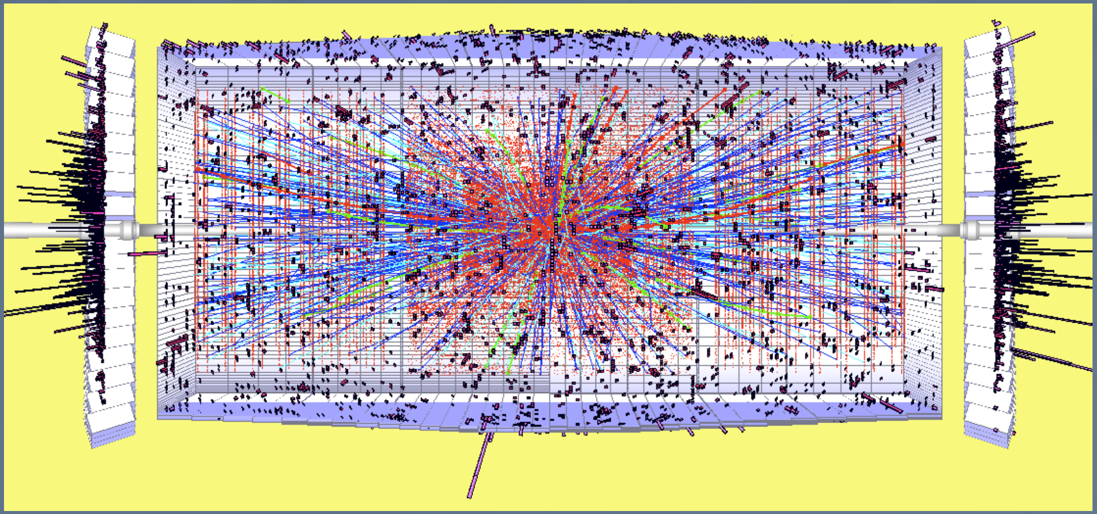

- Goals:
  - Resolve up to 200÷250 collisions per bunch crossing
  - Maintain occupancy at the few % level
  - Maintain overall L1 rate within 100 KHz
  - Keep latency within ~ 6 µs (ECAL pipeline 256 samples = 6.4 µs)
    - The current limit is the Tracker
- L1 tracking trigger data combined with calorimeter & muon trigger data
  - With finer granularity than presently employed.
- Physics objects made from tracking, calorimeter & muon trigger data transmitted to Global Trigger.

Level 1 challenge

- Can we do this in software?
- Maybe in GPGPUs / XeonPhis → studies ongoing in the NA62 experiment
- We need low and – ideally – deterministic latency
- Need an efficient interface to detector-hardware: CPU/FPGA hybrid?
- Or forget about the whole L1 thing altogether and do everything in HLT → requires a lot of fast, low-power, low-cost links, did anybody say Si-photonics?

Challenge #2
Data Acquisition

Data Acquisition (generic example)

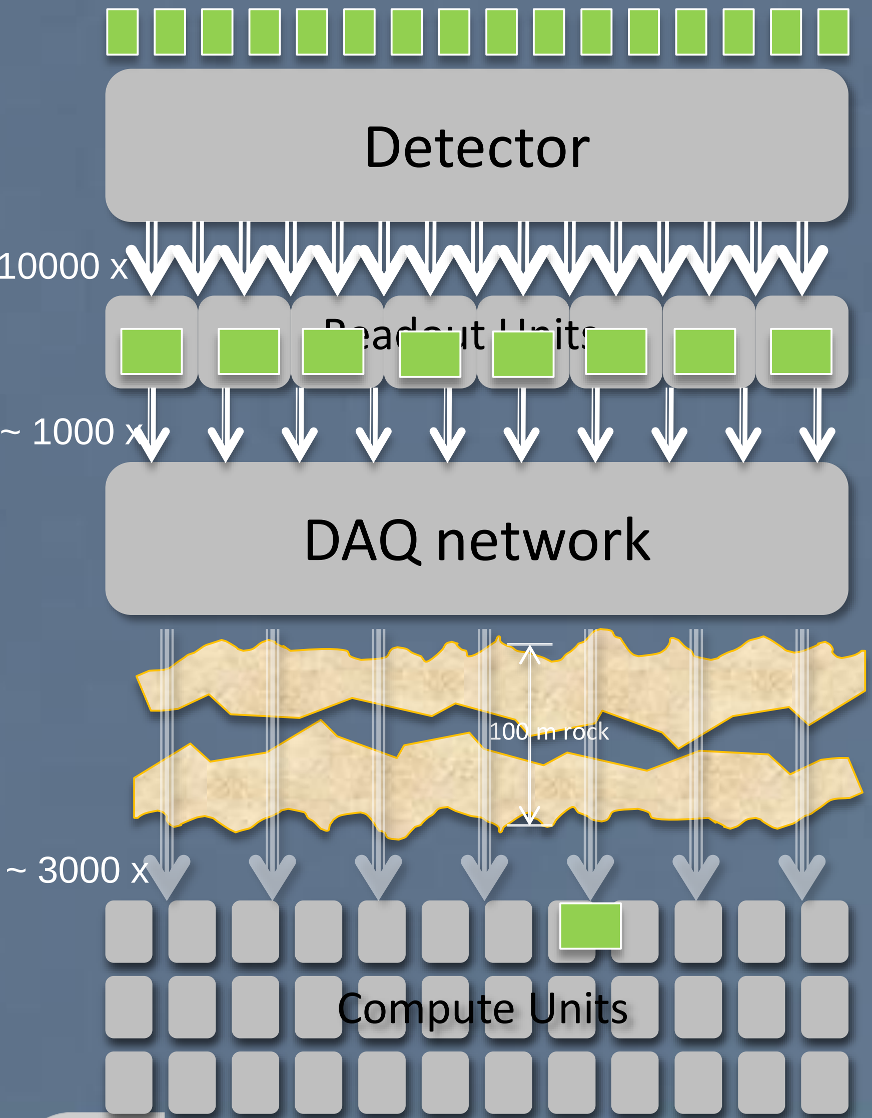

Every Readout Unit has a piece of the collision data
All pieces must be brought together into a single compute unit
The Compute Unit runs the software filtering (High Level Trigger – HLT)

GBT: custom radiation- hard link from the detector 3.2 Gbit/s

DAQ ("event-building") links – some LAN (10/40/100 GbE / InfiniBand)

Links into compute-units: typically 10 Gbit/s (because filtering is currently compute-limited)

Future LHC DAQs in numbers

	Data-size / collision [kB]	Rate of collisions requiring full processing [kHz]	Required # of 100 Gbit/s links	Aggregated bandwidth	From
ALICE	20000	50	120	10 Tbit/s	2019
ATLAS	4000	500	300	20 Tbit/s	2022
CMS	4000	1000	500	40 Tbit/s	2022
LHCb	100	40000	500	40 Tbit/s	2019

Design principles

- Minimize number of expensive "core" network ports
- Use the most efficient technology for a given connection
  - different technologies should be able to co-exist (e.g. fast for building, slow for end-node)
  - keep distances short
- Exploit the economy of scale → try to do what everybody does (but smarter ☺)

DAQ challenge

- Transport multiple Terabit/s reliably and cost-effectively
- Integrate the network closely and efficiently with compute resources (be they classical CPU or “many-core”)
- Multiple network technologies should seamlessly co-exist in the same integrated fabric (“the right link for the right task”)

Challenge #3
High Level Trigger

High Level Trigger

“And this, in simple terms, is how we find the Higgs Boson”

- Pack the knowledge of tens of thousands of physicists and decades of research into a huge sophisticated algorithm
- Several 100.000 lines of code
- Takes (only!) a few 10 - 100 milliseconds ***per collision***

High Level Trigger: Key Figures

- Existing code base: 5 MLOC of mostly C++
- Almost all algorithms are single-threaded (only few exceptions)
- Currently processing time on a X5650 per event: several 10 ms / process (hyper-thread)
- Currently between 100k and 1 million events per second are filtered online in each of the 4 experiments

Pattern finding - tracks

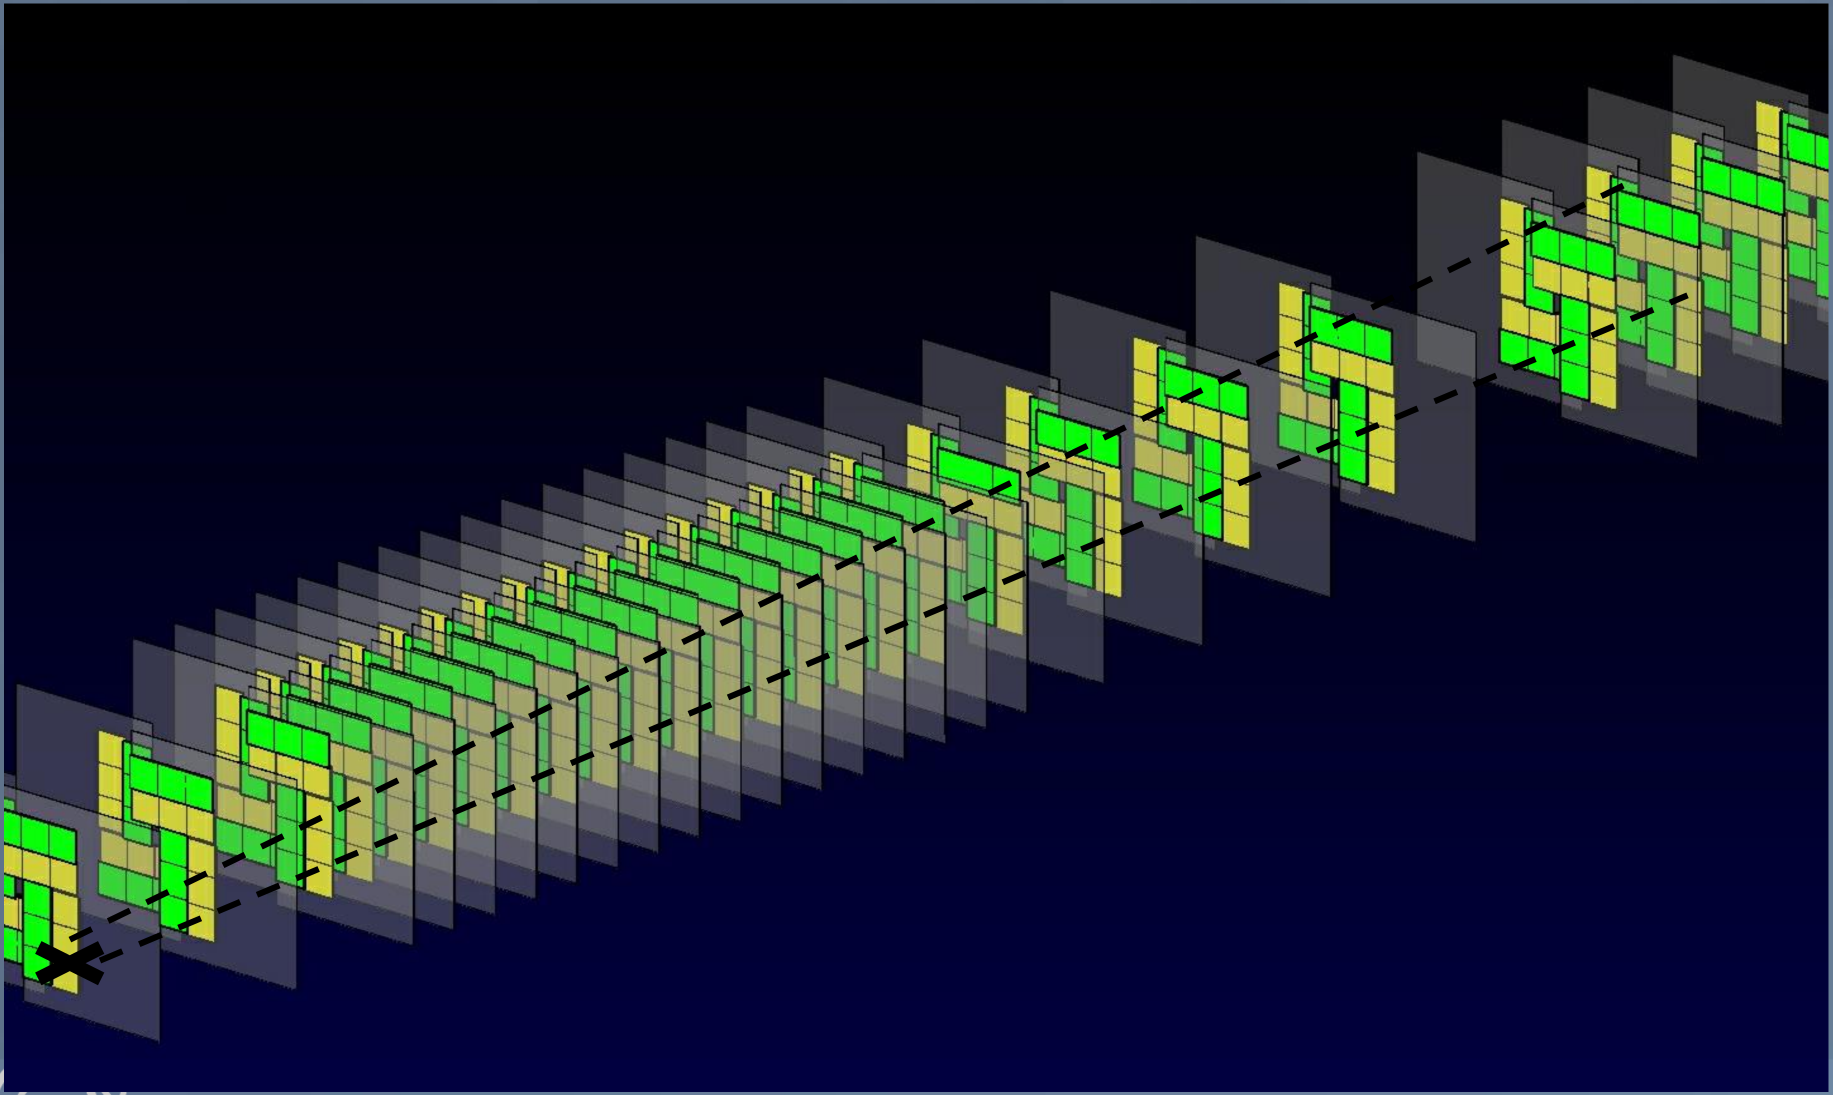

High Level Trigger compared to HPC

- Like HPC:
  - full ownership of the entire installation → can choose architecture and hardware components
  - single “client” / “customer”
  - have a high-bandwidth interconnect
- Unlike HPC:
  - many independent small tasks which execute quickly
    ➔ no need for check-pointing (fast storage)
    ➔ no need for low latency
  - data driven, i.e. when the LHC is **not** running (70% of the time) the farm is idle → interesting ways around this (deferal, “offline usage)
  - facility is very long-lived, growing incrementally

High Level Trigger challenge

- Make the code-base ready for multi/many-core (this is not Online specific!)
- Optimize the High Level Trigger farms in terms of cost, power, cooling
- Find the best architecture integrating “standard servers”, many-core systems and a high-bandwidth network

Summary

- CERN is the leading particle physics laboratory in the world hosting more than 100 experiments
- The flagship programme is the LHC, whose experiments produce an unprecedented amount of data
- Sophisticated IT is needed at all stages to cope with these data
- The IT needs of the experiments are always growing and (often) more interesting physics requires more networking and more computing, while budgets stay constant ☺