Code optimization in High Energy Physics

- challenges at the LHC -

Nuno Castro (nuno.castro@cern.ch)
António Onofre (antonio.onofre@cern.ch)

ATLAS

LIP

20th November 2012

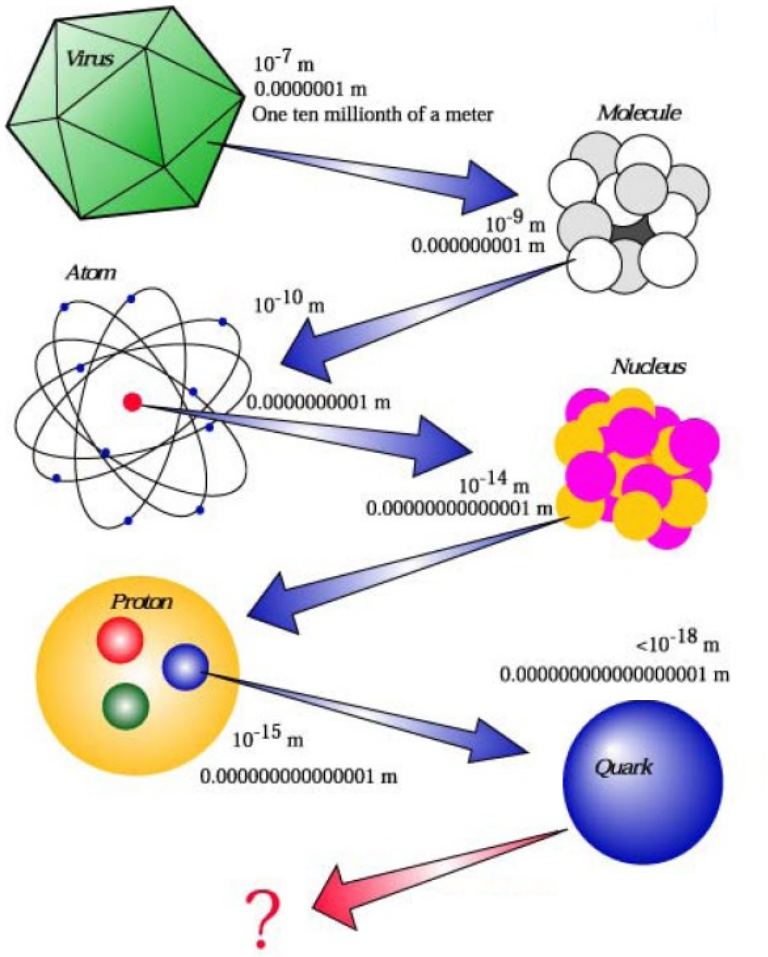

$$\lambda = \frac{h}{p}$$

See at small scales

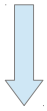

High p

High Energy Physics

Particle Physics

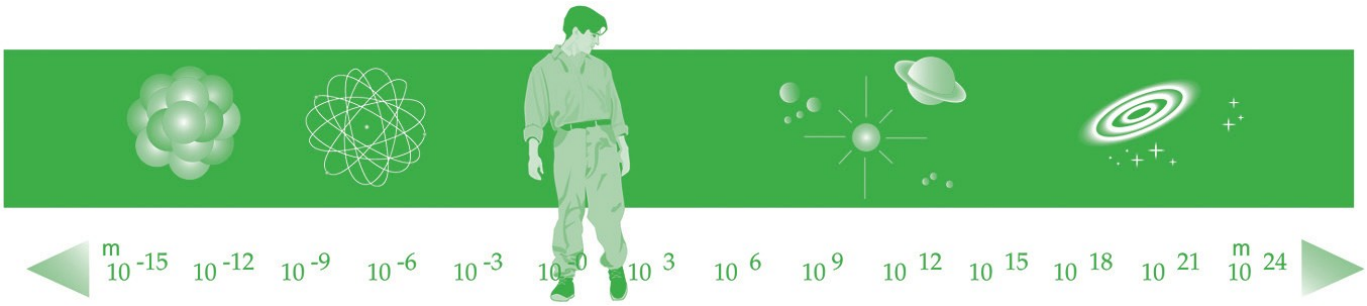

the 2 frontiers:
very big and very small

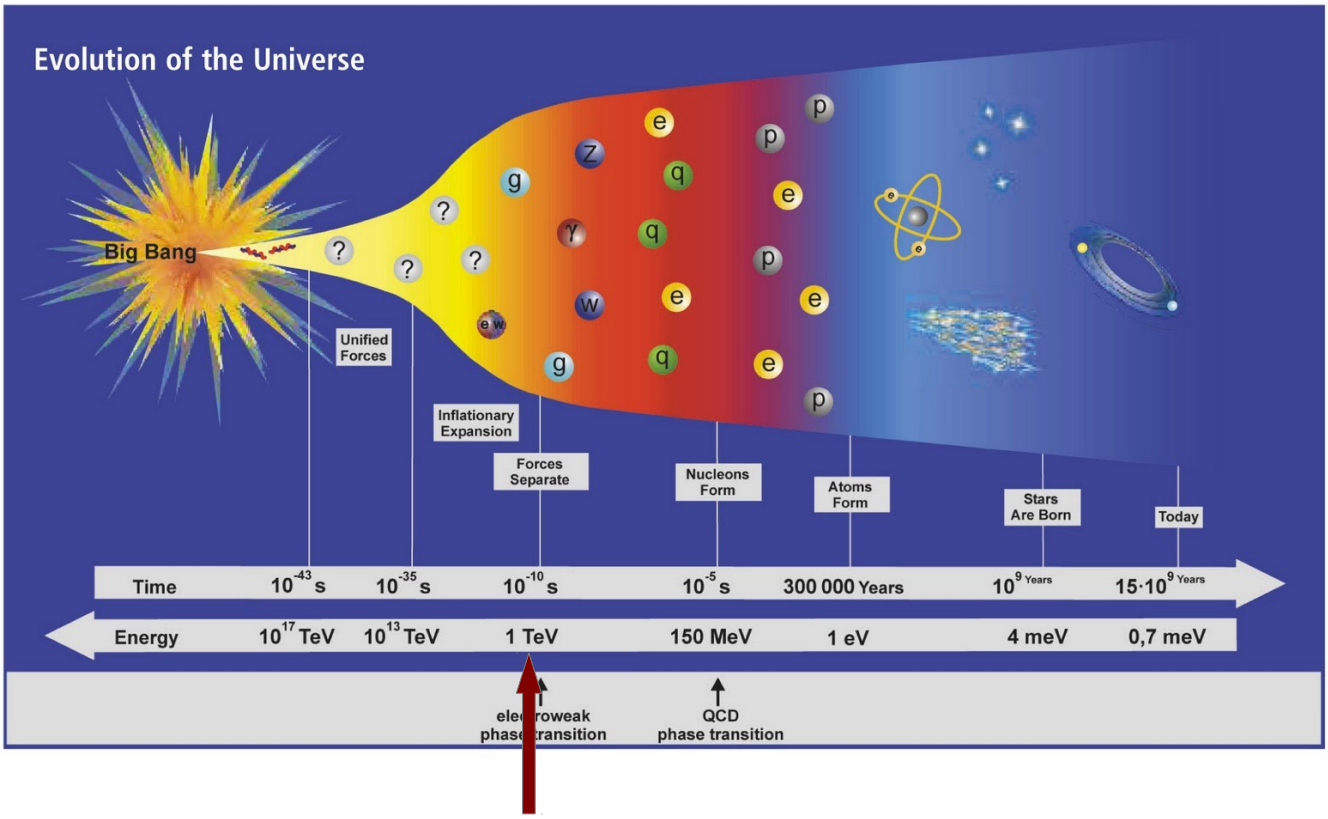

Large Hadron Collider (LHC)

Standard Model of Particle Physics

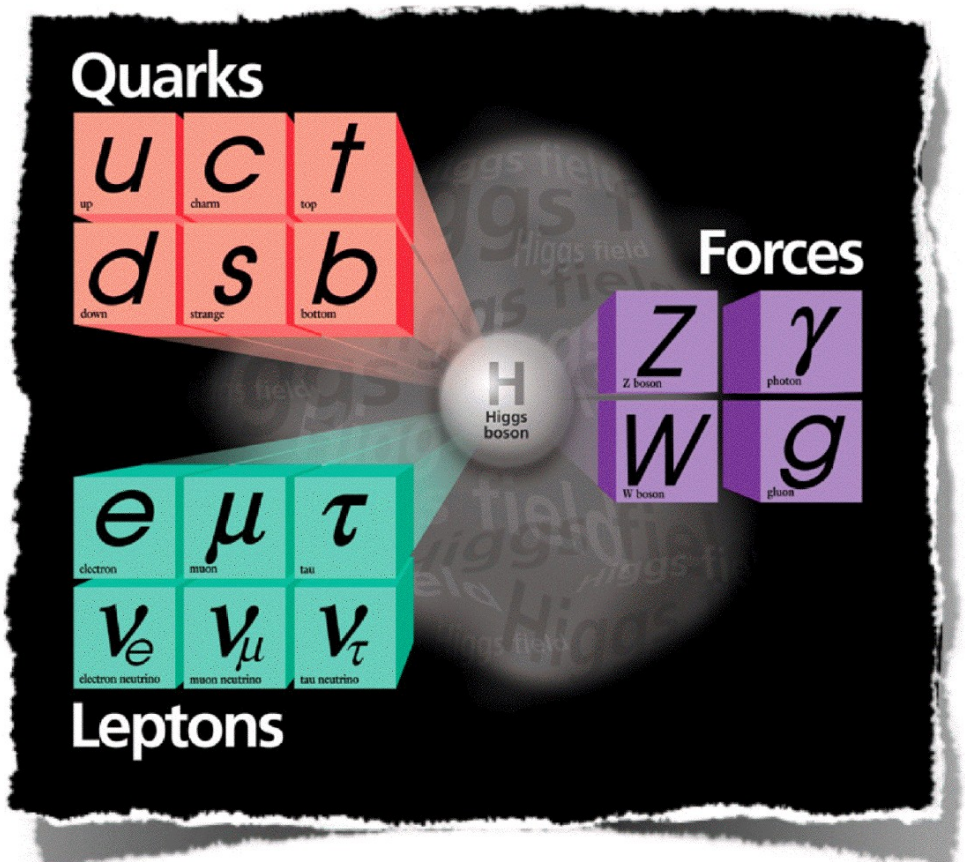

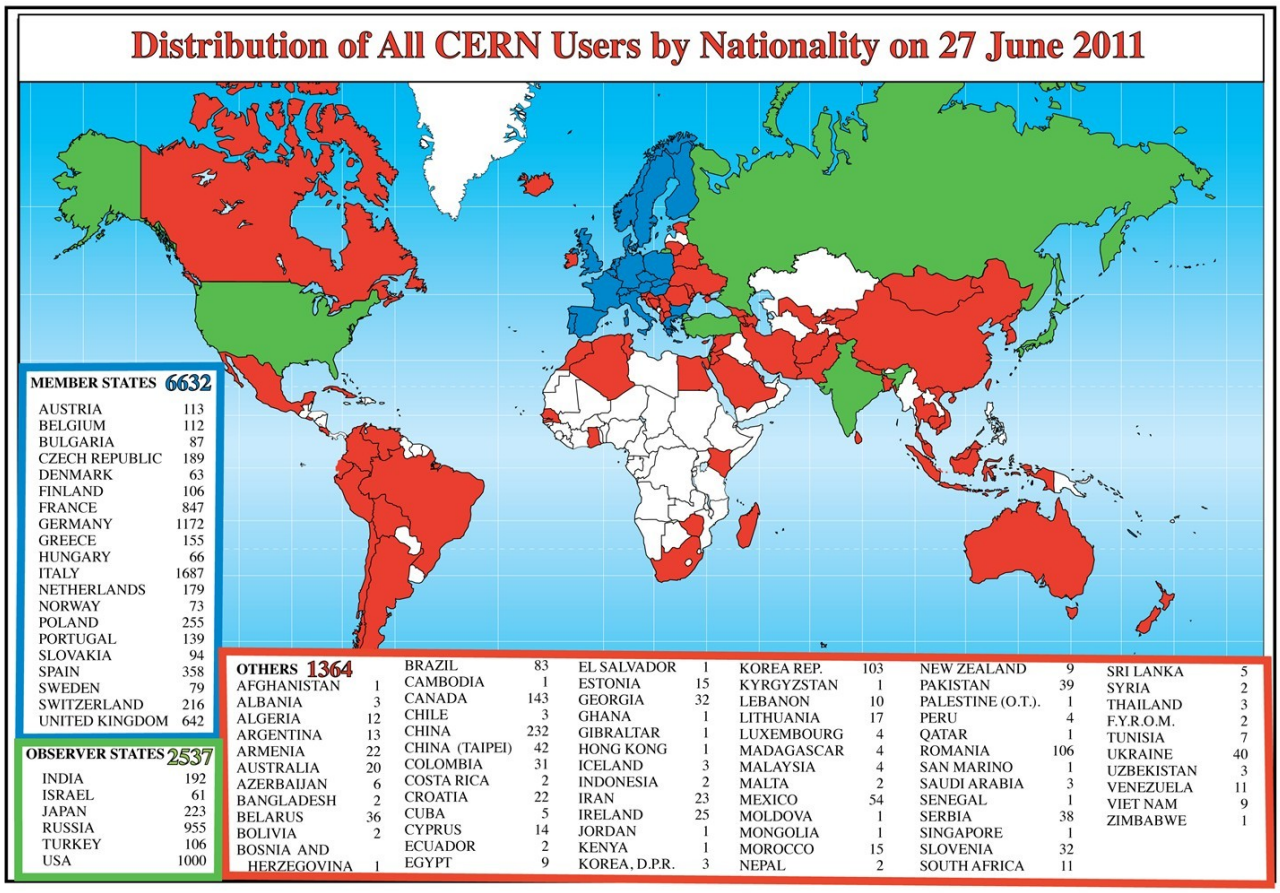
Distribution of All CERN Users by Nationality on 27 June 2011
MEMBER STATES 6632
AUSTRIA 113
BELGIUM 112
BULGARIA 87
CZECH REPUBLIC 189
DENMARK 63
FINLAND 106
FRANCE 847
GERMANY 1172
GREECE 155
HUNGARY 66
ITALY 1687
NETHERLANDS 179
NORWAY 73
POLAND 255
PORTUGAL 139
SLOVAKIA 94
SPAIN 358
SWEDEN 79
SWITZERLAND 216
UNITED KINGDOM 642
OBSERVER STATES 2537
INDIA 192
ISRAEL 61
JAPAN 223
RUSSIA 955
TURKEY 106
USA 1000
OTHERS 1364
AFGHANISTAN 1
ALBANIA 3
ALGERIA 12
ARGENTINA 13
ARMENIA 22
AUSTRALIA 20
AZERBAIJAN 6
BANGLADESH 2
BELARUS 36
BOLIVIA 2
BOSNIA AND HERZEGOVINA 1
BRAZIL 83
CAMBODIA 1
CANADA 143
CHILE 3
CHINA 232
CHINA (TAIPEI) 42
COLOMBIA 31
COSTA RICA 2
CROATIA 22
CUBA 5
CYPRUS 14
ECUADOR 2
EGYPT 9
EL SALVADOR 1
ESTONIA 15
GEORGIA 32
GHANA 1
GIBRALTAR 1
HONG KONG 1
ICELAND 3
INDONESIA 2
IRAN 23
IRELAND 25
JORDAN 1
KENYA 1
KOREA, D.P.R. 3
KOREA REP. 103
KYRGYZSTAN 1
LEBANON 10
LITHUANIA 17
LUXEMBOURG 4
MADAGASCAR 4
MALAYSIA 4
MALTA 2
MEXICO 54
MOLDOVA 1
MONGOLIA 1
MOROCCO 15
NEPAL 2
NEW ZEALAND 9
PAKISTAN 39
PALESTINE (O.T.). 1
PERU 4
QATAR 1
ROMANIA 106
SAN MARINO 1
SAUDI ARABIA 3
SENEGAL 1
SERBIA 38
SINGAPORE 1
SLOVENIA 32
SOUTH AFRICA 11
SRI LANKA 5
SYRIA 2
THAILAND 3
F.Y.R.O.M. 2
TUNISIA 7
UKRAINE 40
UZBEKISTAN 3
VENEZUELA 11
VIET NAM 9
ZIMBABWE 1

The LHC

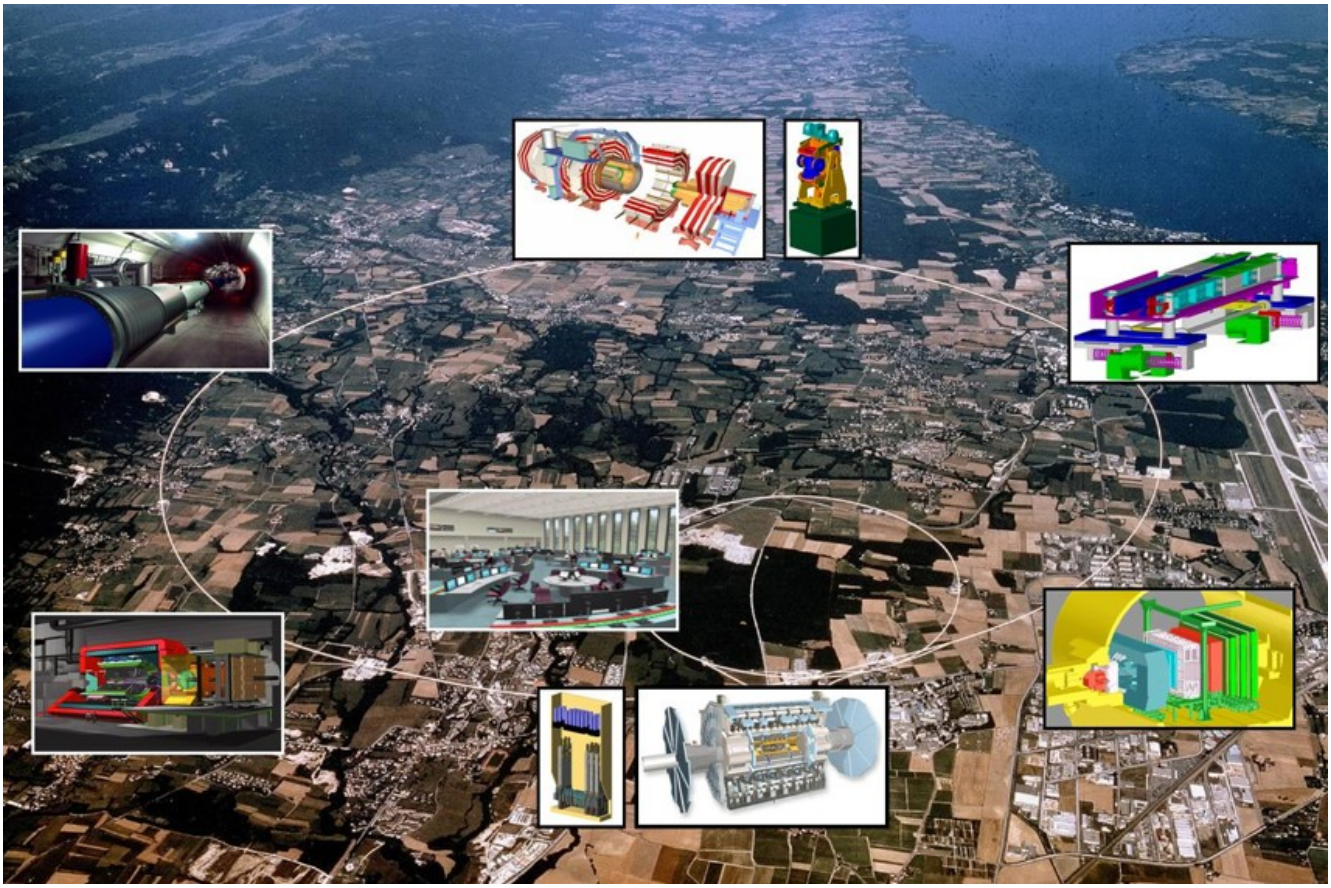

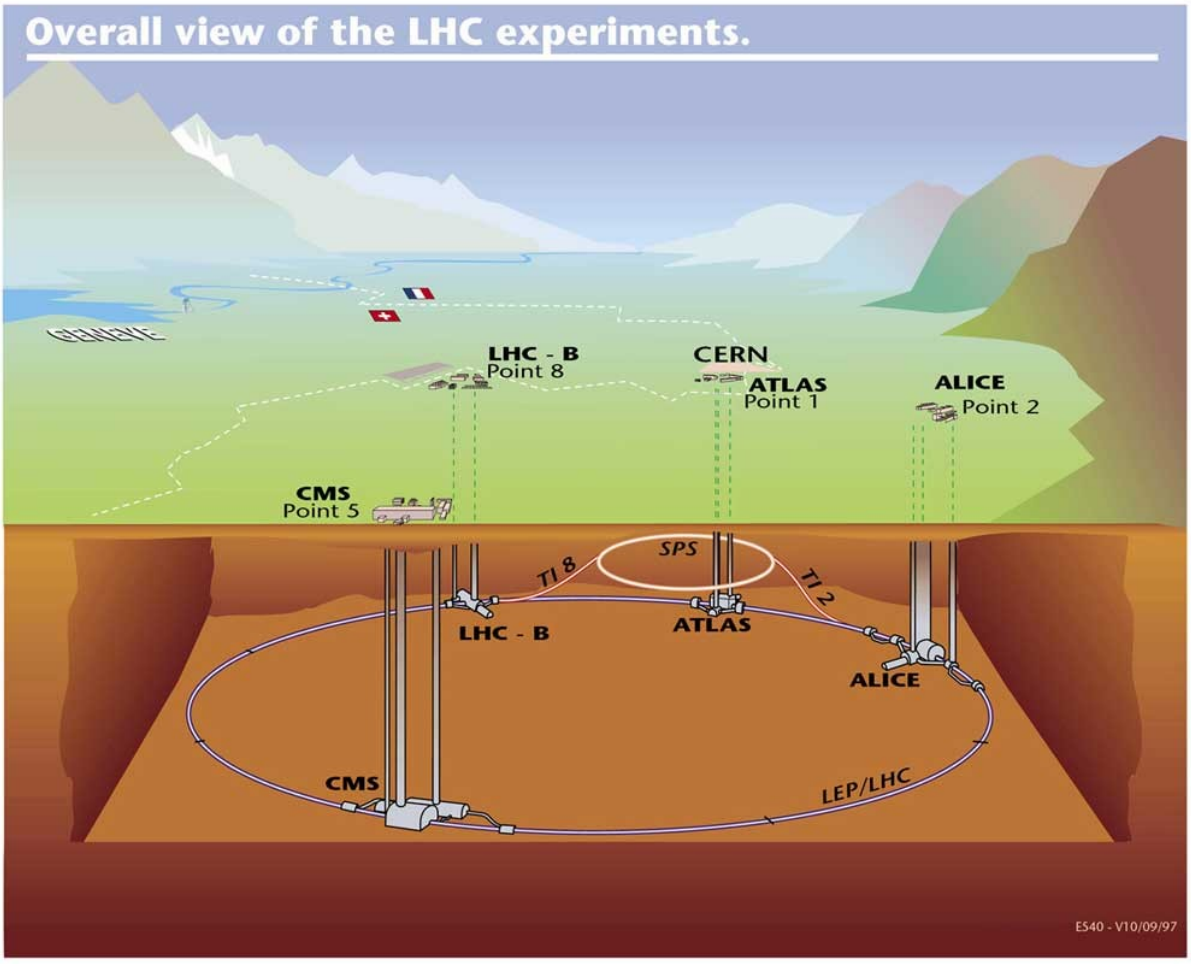
Overall view of the LHC experiments.
GENEVE
LHC - B
Point 8
CERN
ATLAS
Point 1
ALICE
Point 2
CMS
Point 5
SPS
TI 8
TI 2
LHC - B
ATLAS
ALICE
CMS
LEP/LHC
E540 - V10/09/97

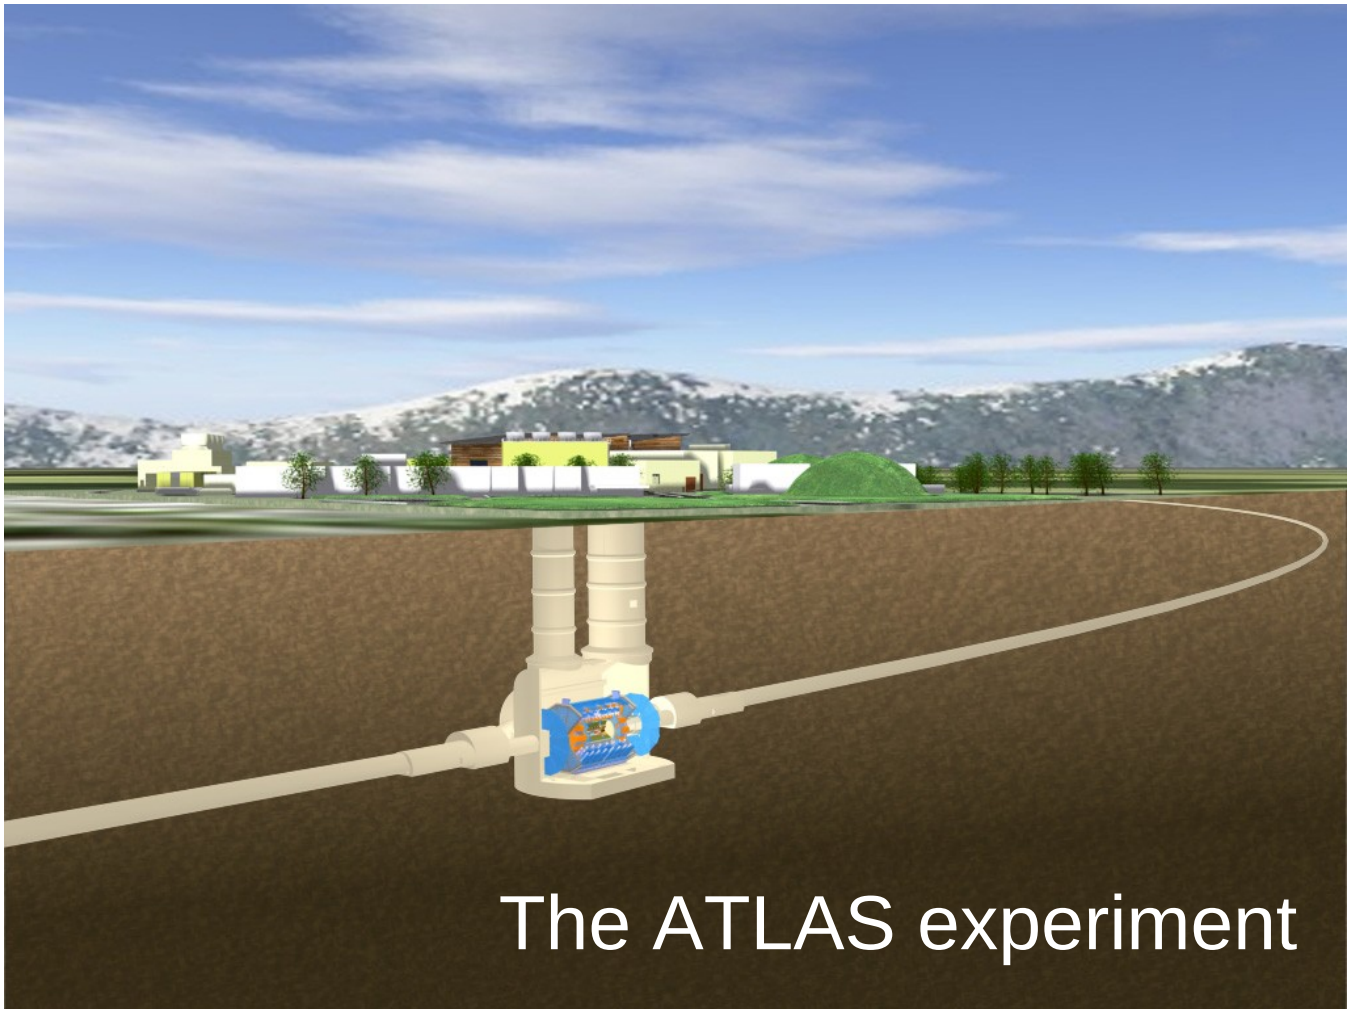
The ATLAS experiment

The ATLAS detector

~ 7000 t
~ 10^8 channels
~ 3000 km of cables
~ 40 x 10^6 collisions / s
~ 1 PB of data / s

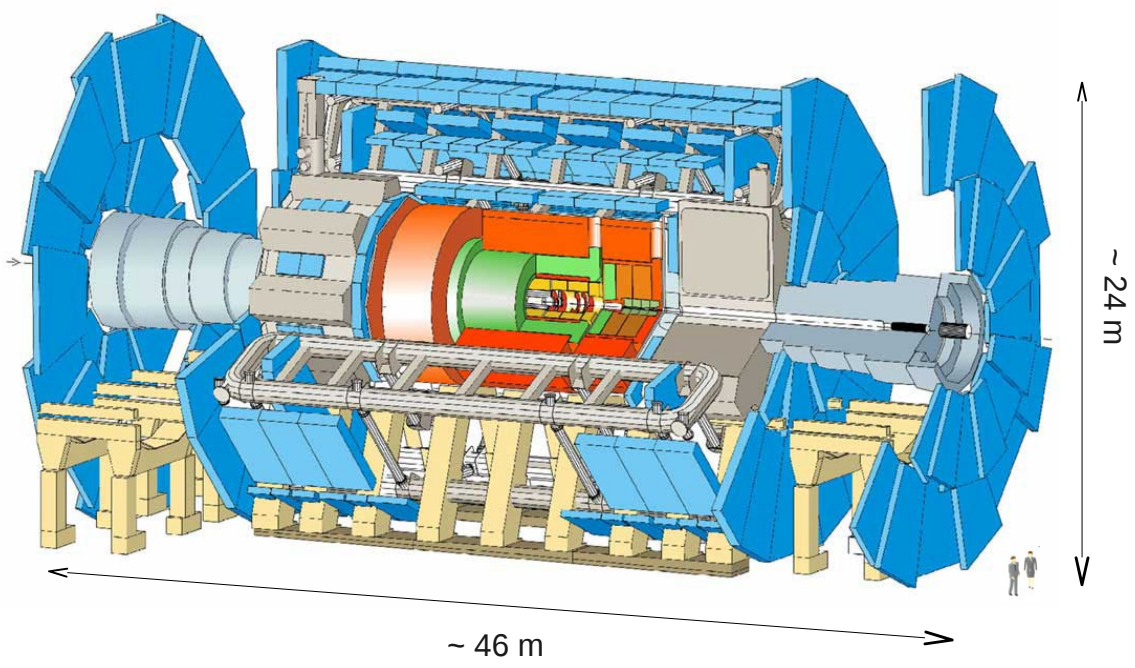

Inner detector

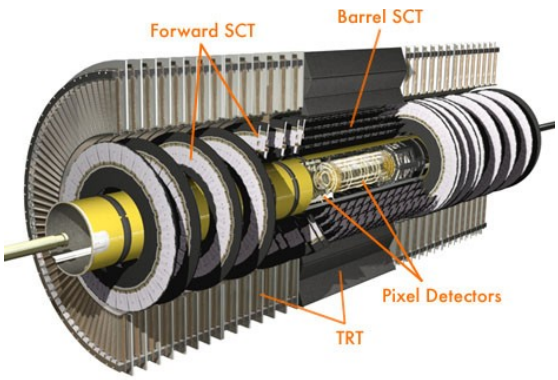

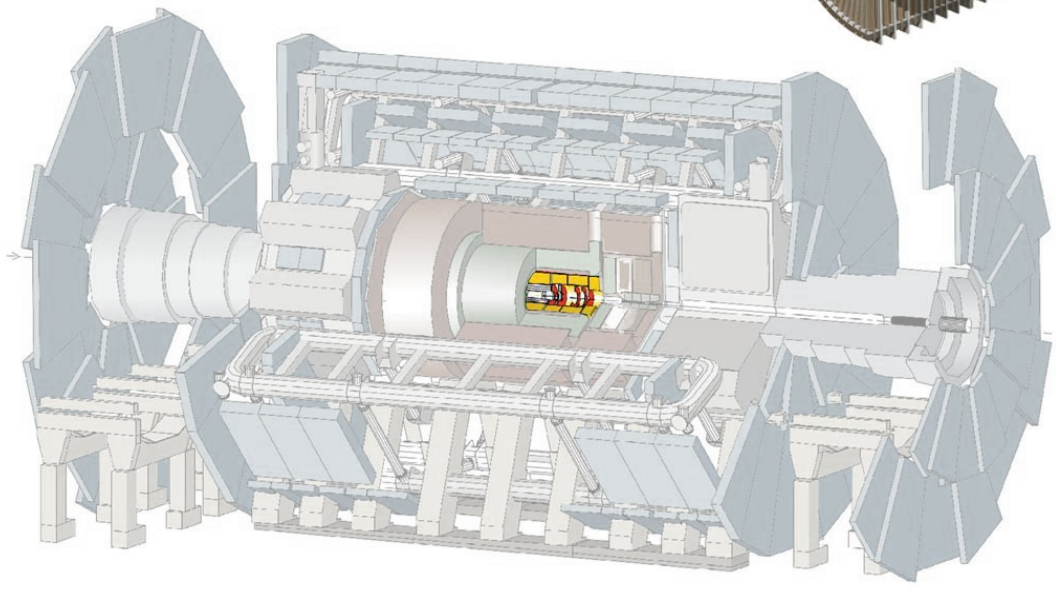

$\sigma/p_T \sim 3.8 \times 10^{-4}\ p_T\ (\text{GeV}) \oplus 0.015$
i.e. σ/p_T <2% for p_T < 35 GeV

Electromagnetic calorimeter

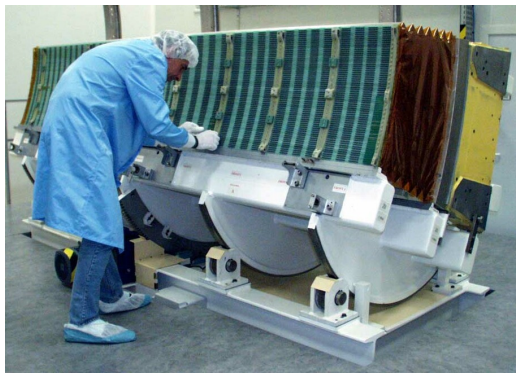

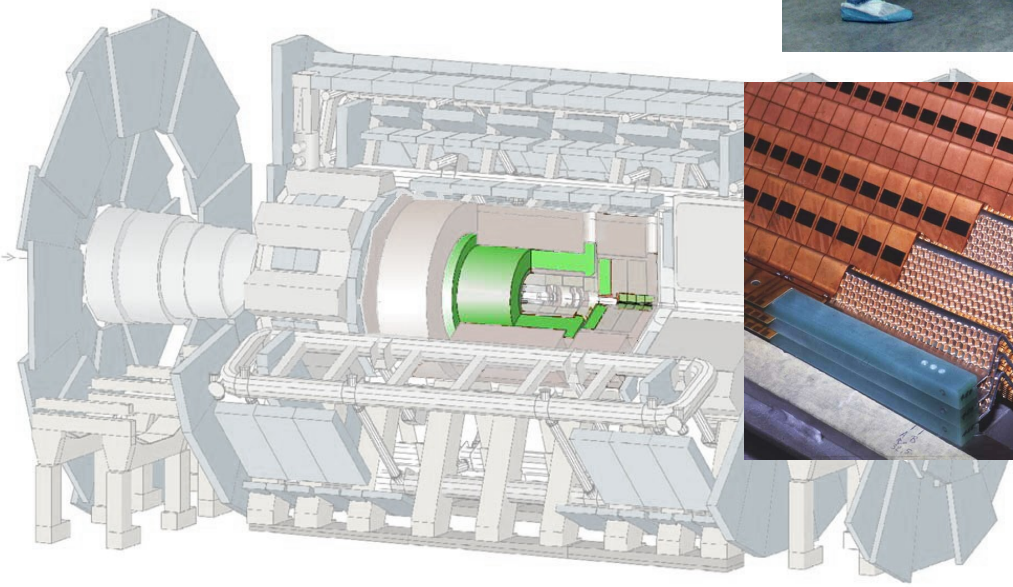

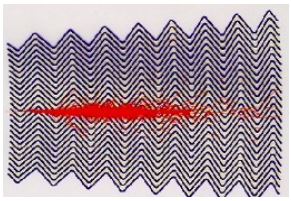

$\sigma/E \sim 10\%/\sqrt{E}$

Hadronic calorimeter

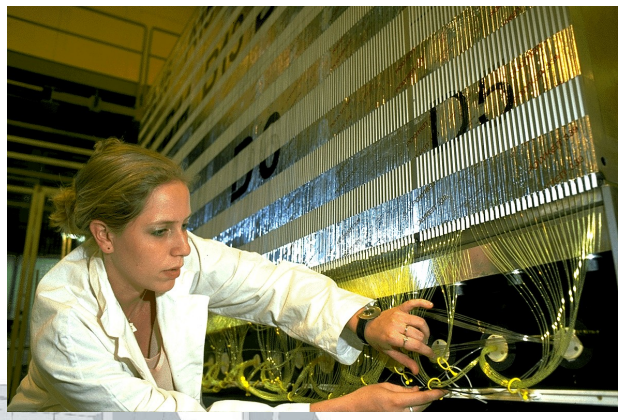

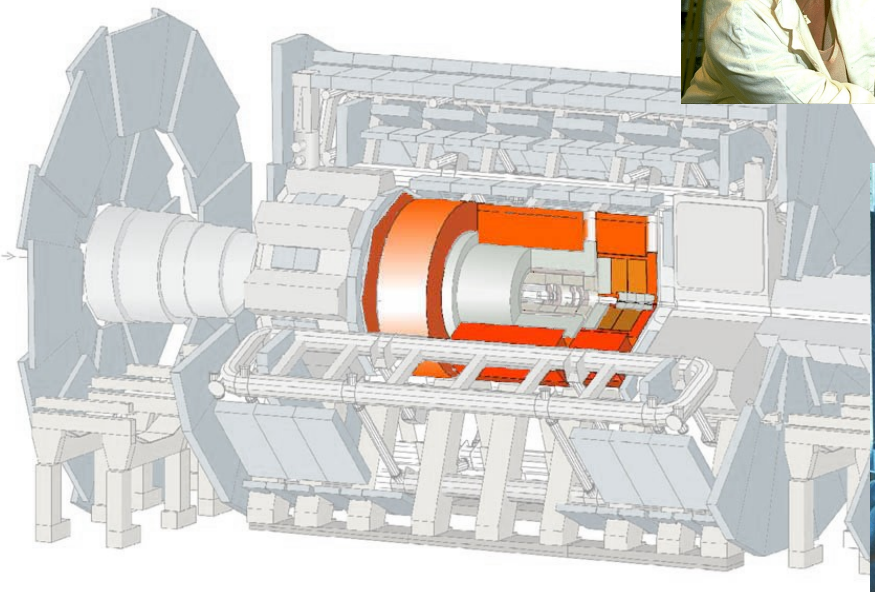

$\sigma/E \sim 50\%/\sqrt{E} \oplus 0.03$

Muons detector

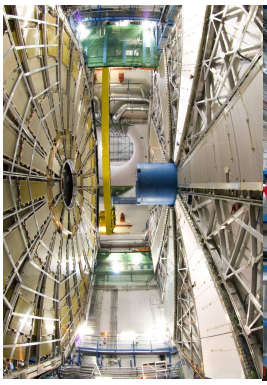

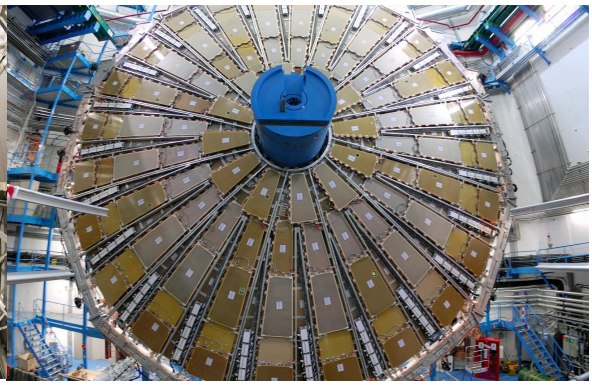

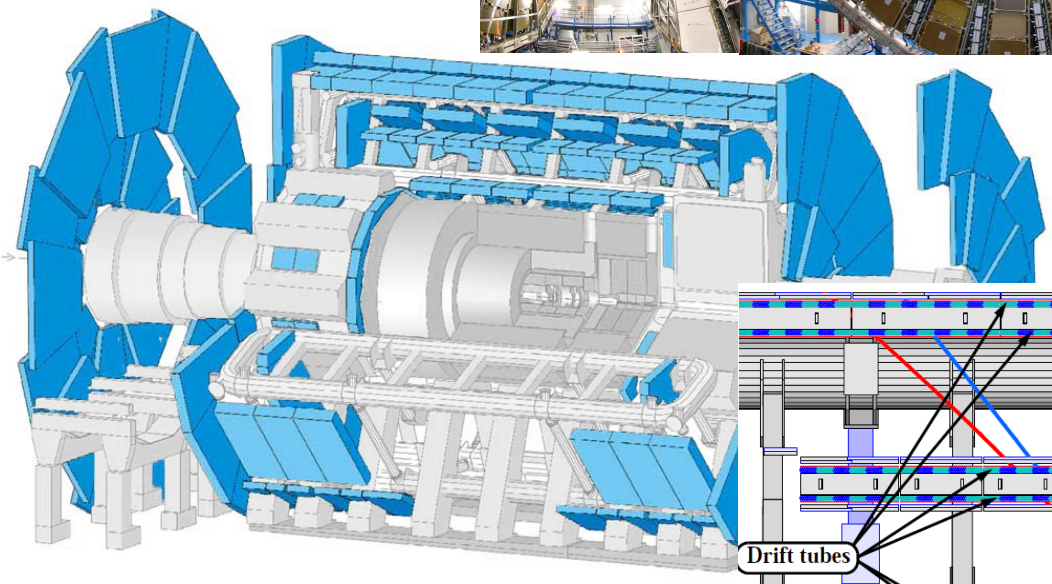

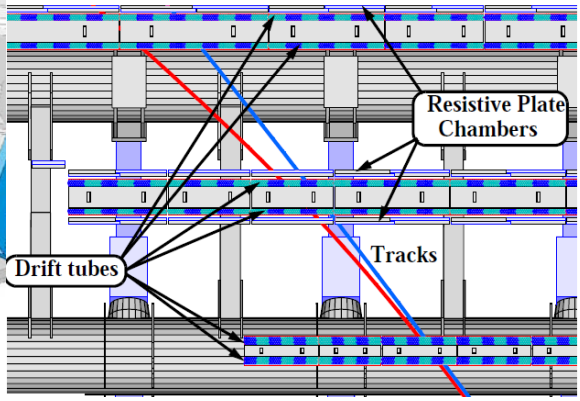

< 10% up to E_μ ~ 1 TeV

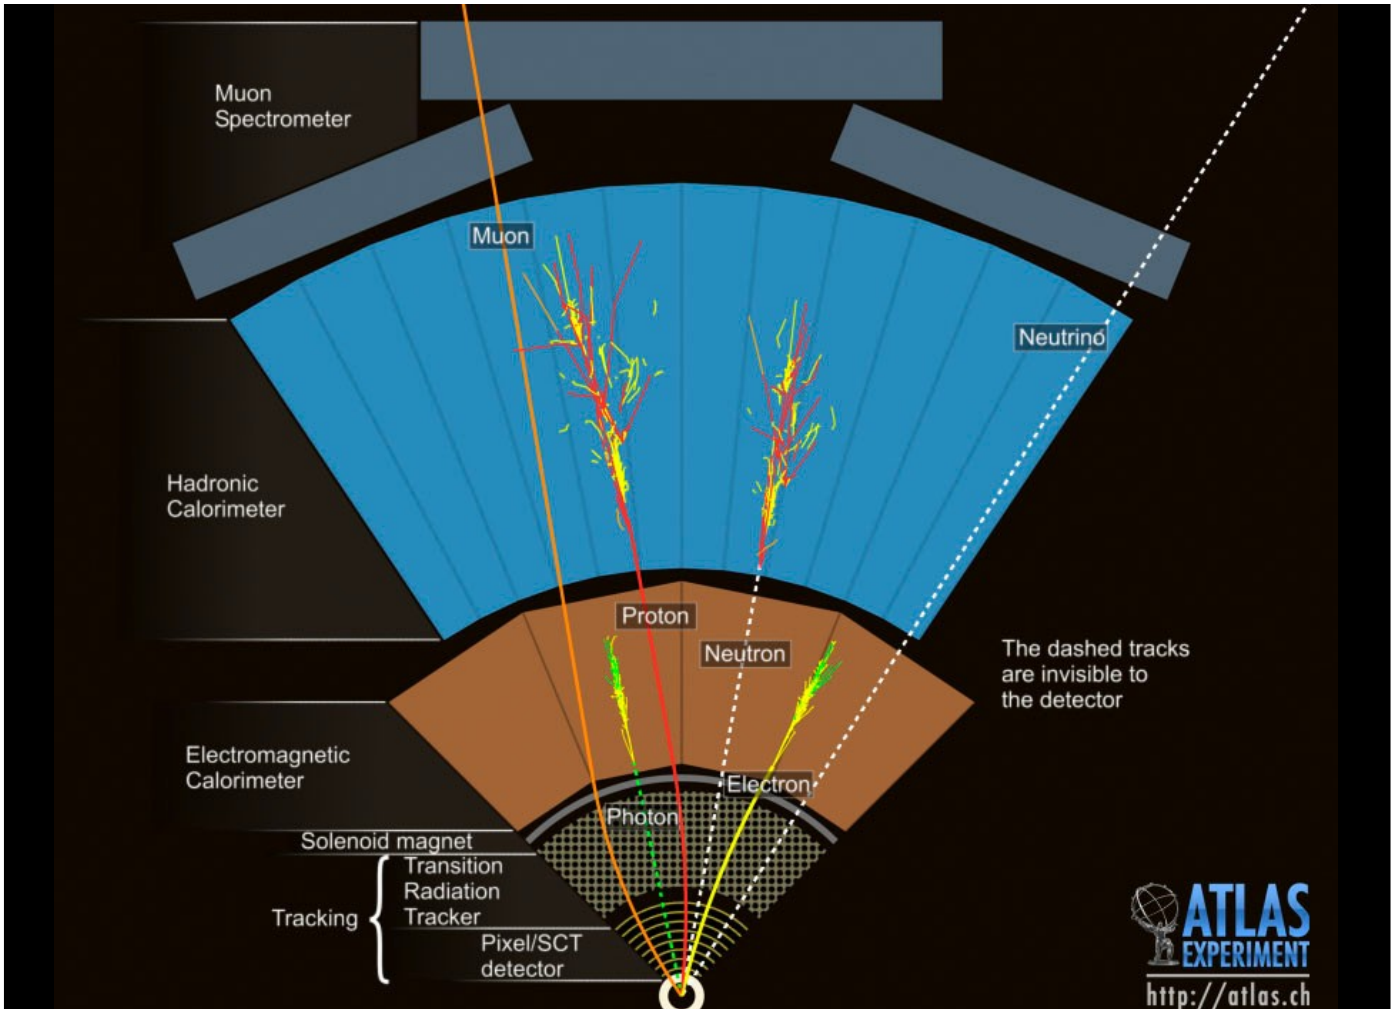

What happens in a collision?

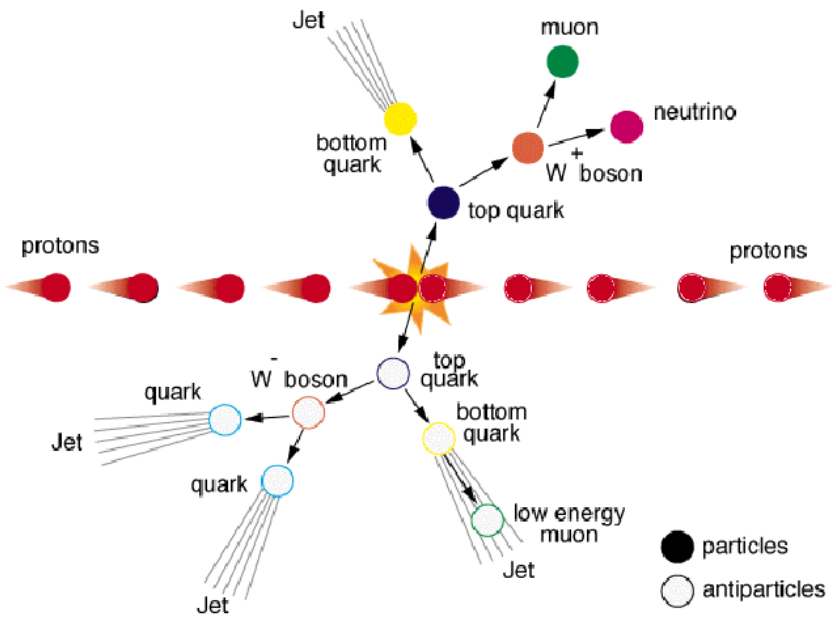

What do we see in the detector?

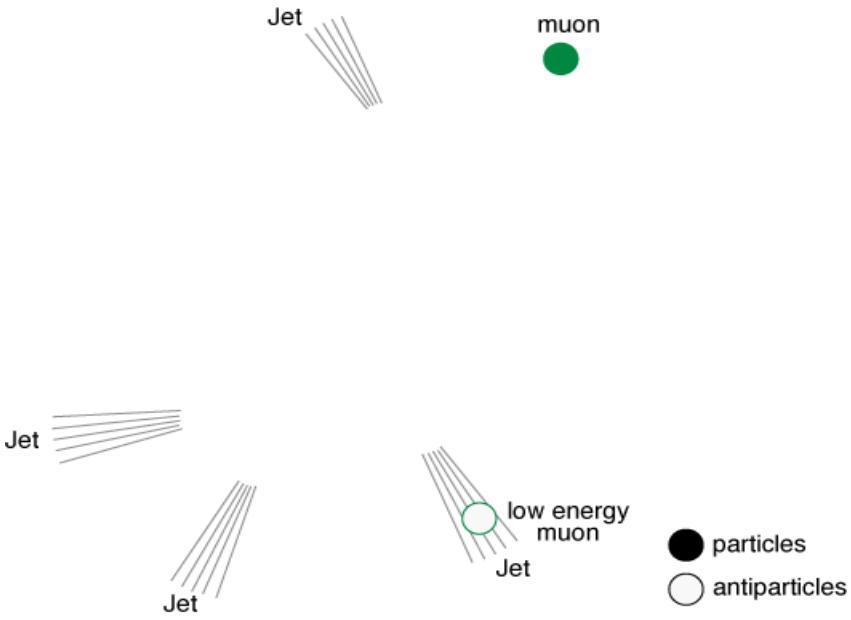

Data collected by ATLAS

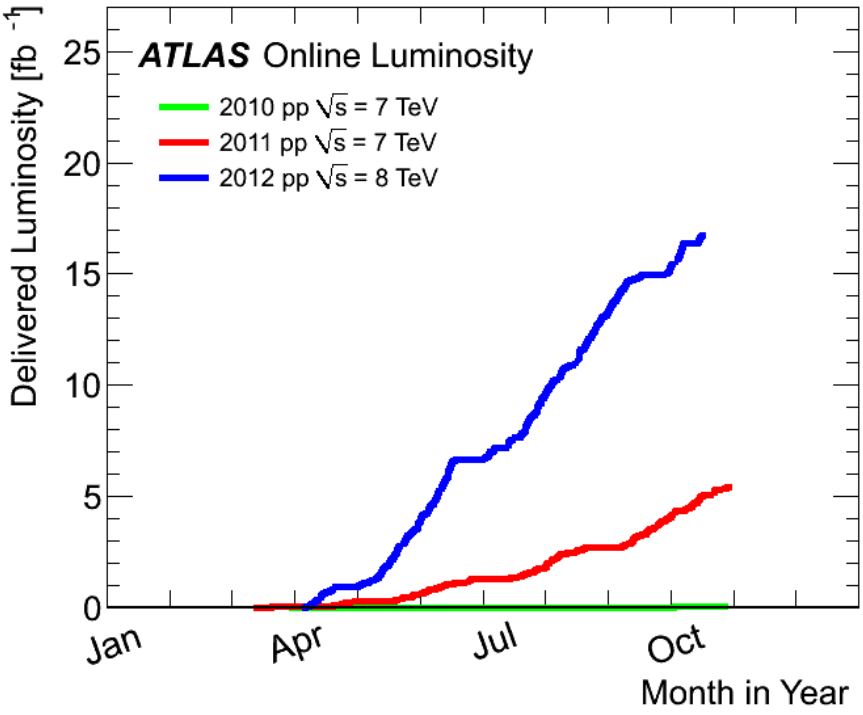

Summary of *pp* collisions results

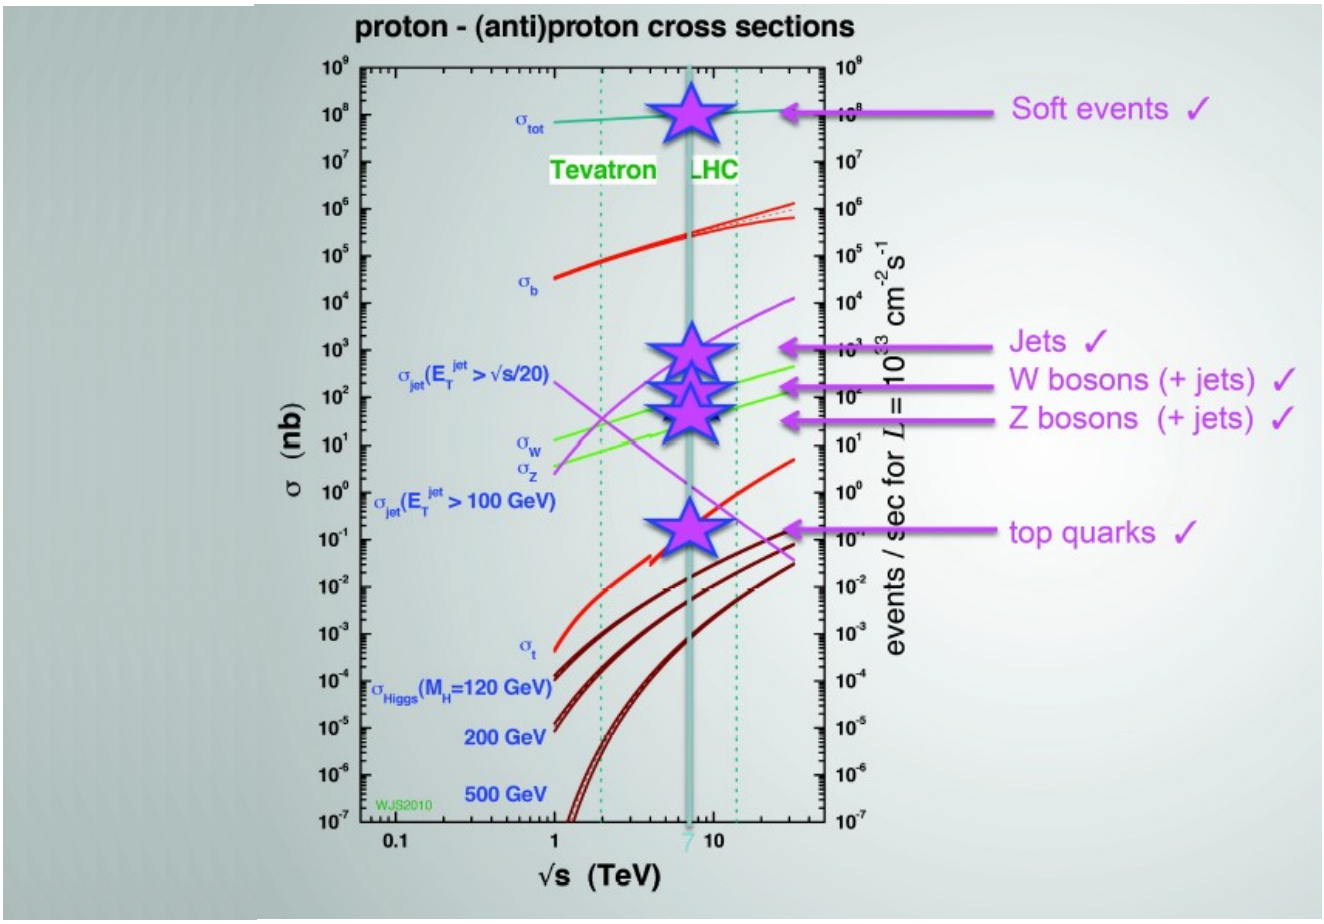

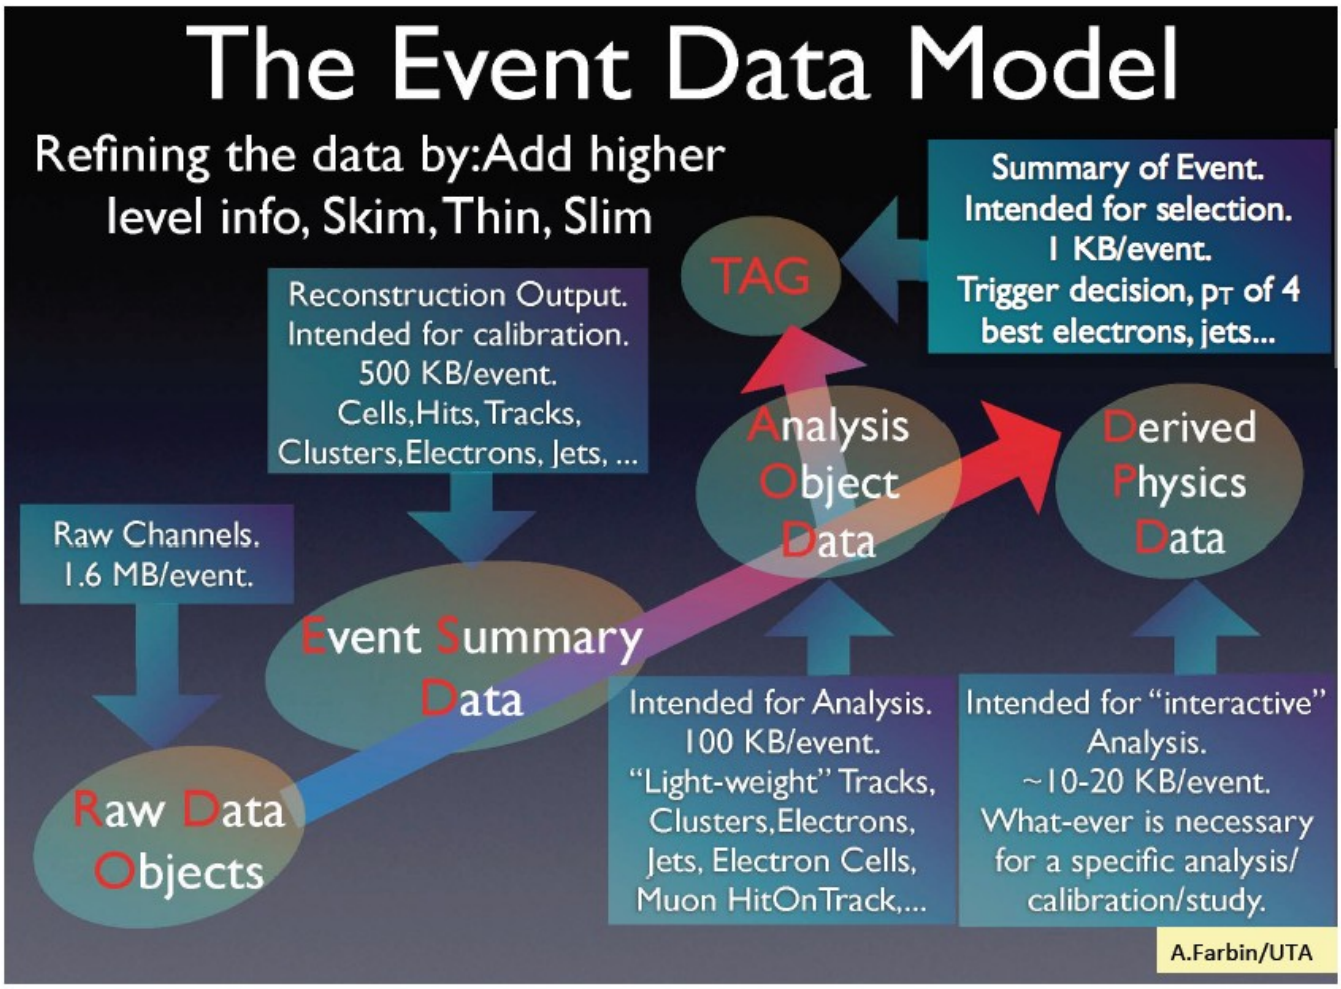

Grid Data Processing

- ATLAS Grid Data Processing (GDP) uses Grids with three different interfaces split in ten "clouds" organized as large computing centres with tape data storage (Tier-1 sites) each associated with 5-6 other computing centres (Tier-2 sites)
  - ATLAS clouds evolve from the hierarchy (left) to the mesh (right)

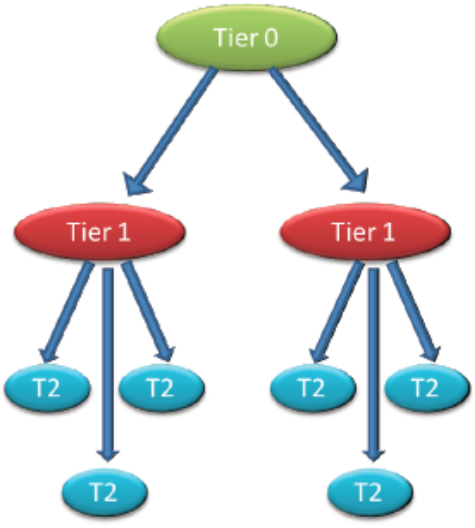

15%: Tier-0 center at CERN
- Data recording
- First-pass data reconstruction
- Data distribution

40%: Tier-1 centers
- Permanent storage
- Reprocessing
- Group data processing

45%: Tier-2 centers
- Simulation
- End-user analysis

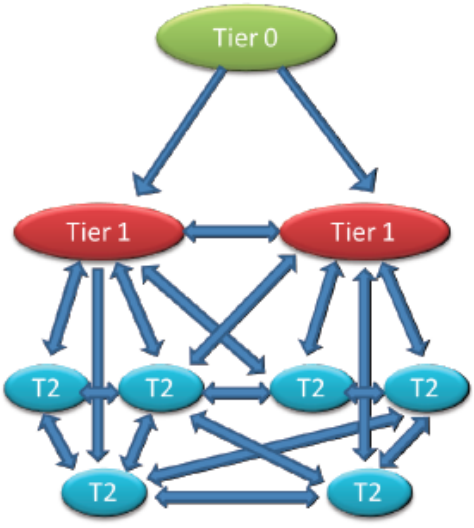

- GDP system empowered further data processing steps performed by dozens of ATLAS physics groups with coordinated access to computing resources worldwide

From Reprocessed Data to Physics Analysis

- In contrast with the major reprocessing campaign that are conducted only few times per year, the centrally managed production for physics groups process full available dataset once every few weeks, providing further improvements in the data used for ATLAS physics analysis shortly after the reprocessing or data taking

- Physics analysis of the detector data relies on the results of the "Monte Carlo experiments" produced by the CPU-intensive detector response simulations tasks mostly running on the Grid Tier-2 sites

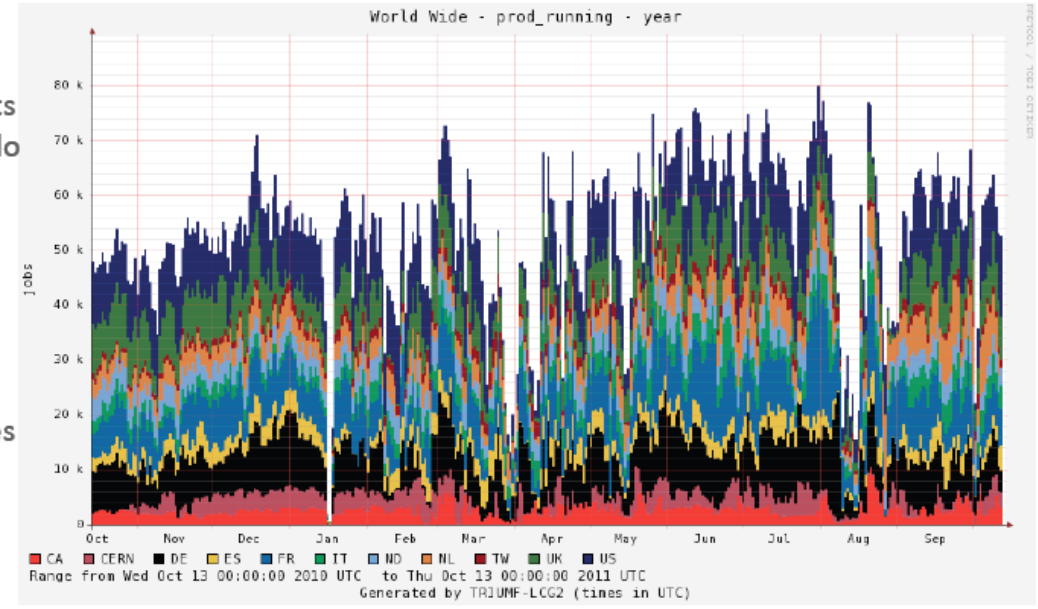

Data placement – original

Available to all users:

- Only data that are on disks at the tier-2s is readily available for user analysis.
- Only the AOD (and the TAG) is shipped to the tier-2s.
- Very detailed information that resides only in the ESD is NOT accessible to users.
- The AOD is still very large for frequent analyzes by users.

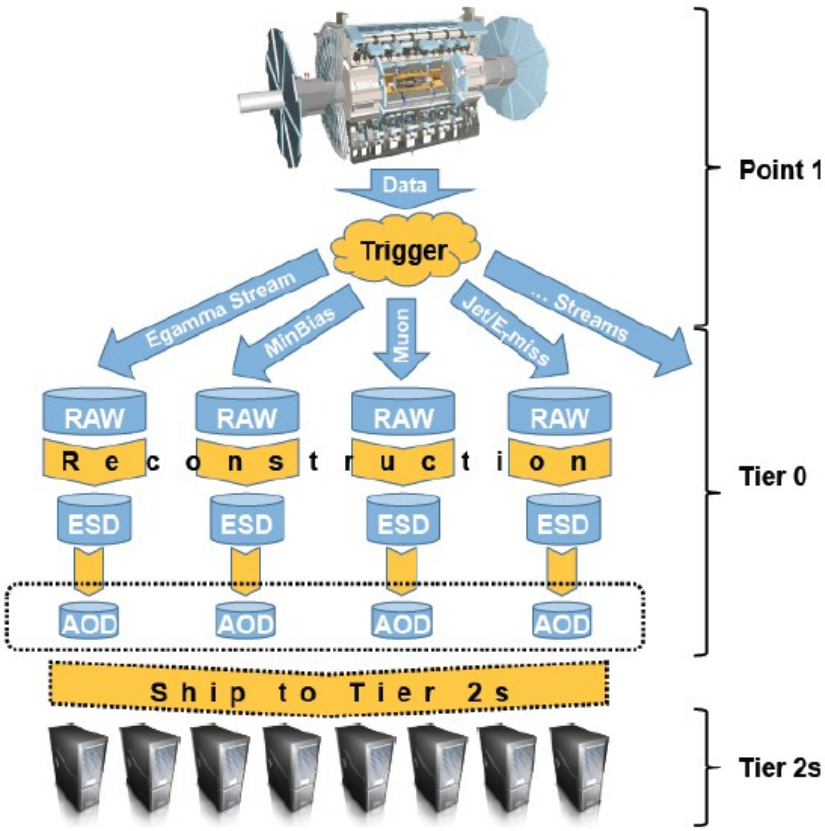

Data production chain – adding DAODs

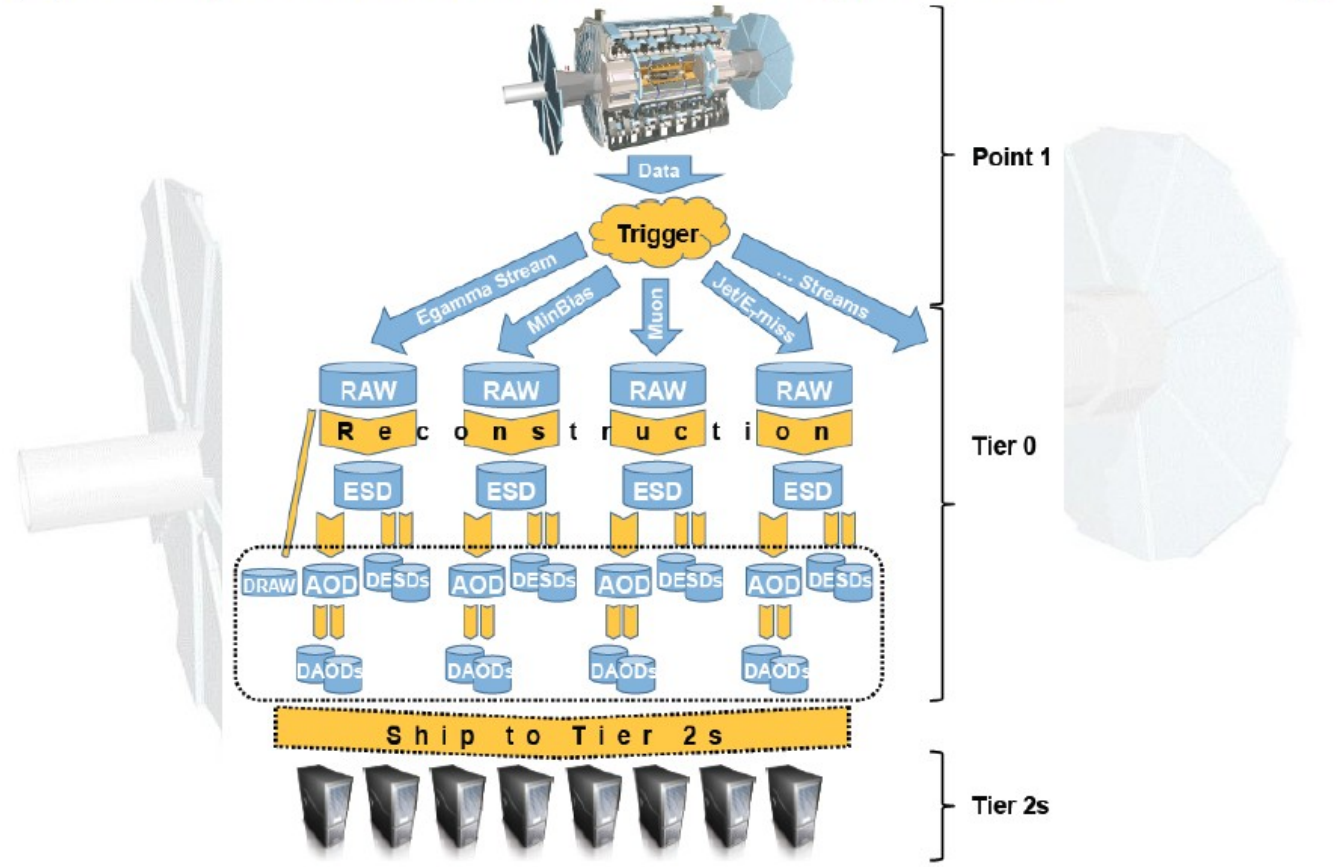

Group ntuples

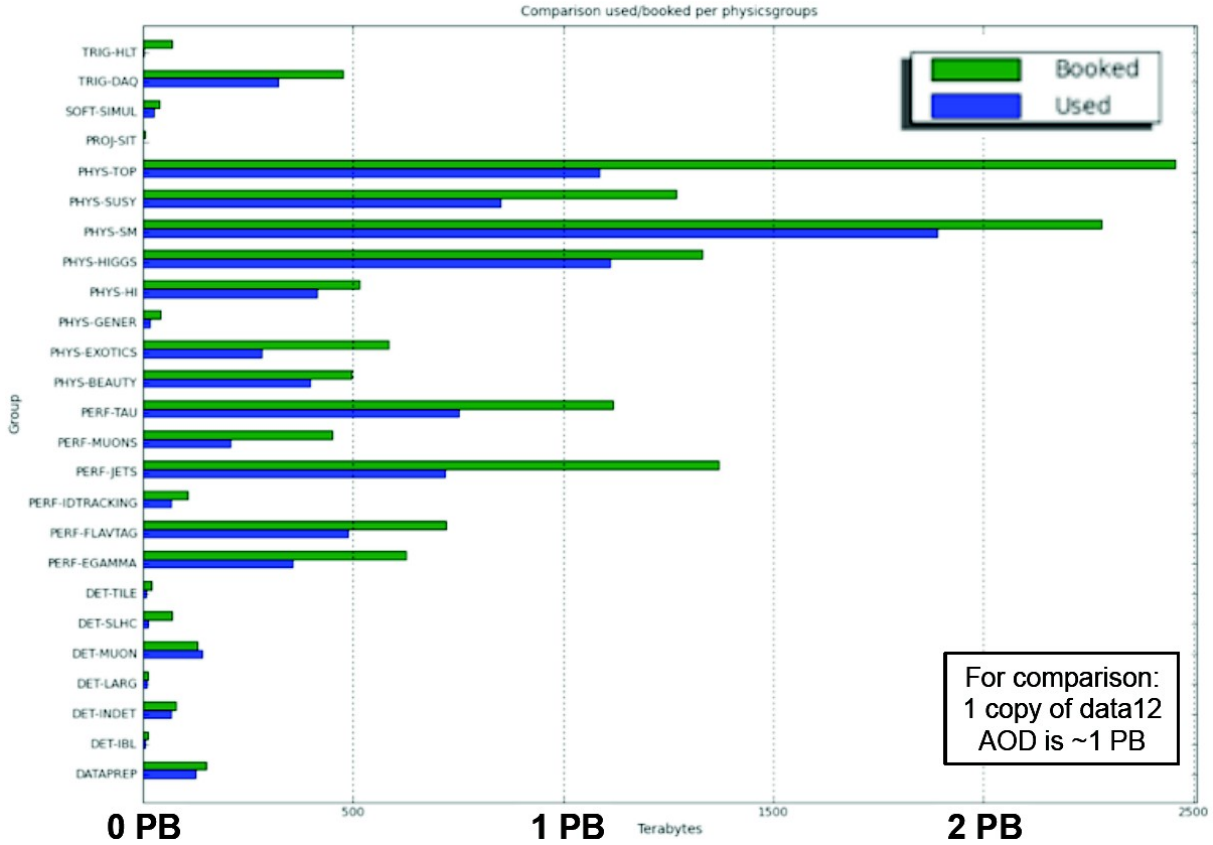

What do we do with the ntuples?

data
(either real data or MC)
analysis

LipCbrAnalysis

LipCbrAnalysis

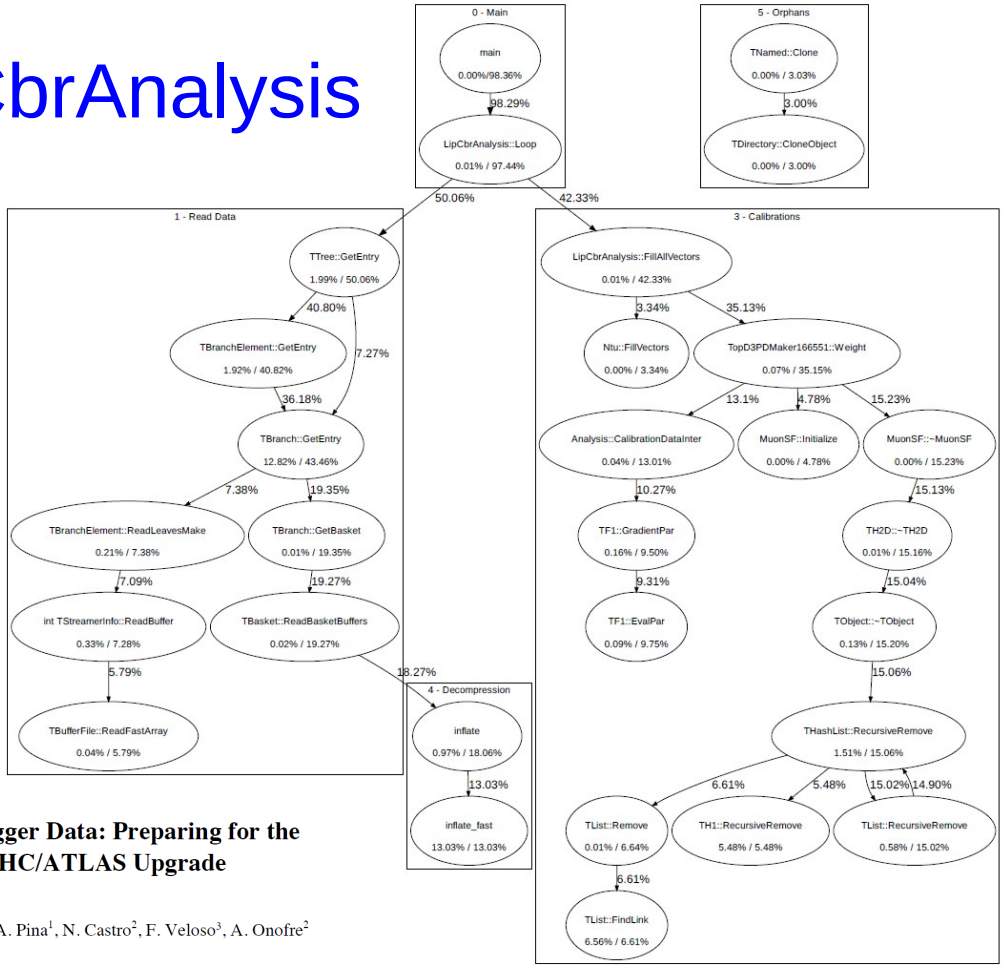

Even Bigger Data: Preparing for the LHC/ATLAS Upgrade

V. Oliveira[1], A. Pina[1], N. Castro[2], F. Veloso[3], A. Onofre[2]

What do we get in the end?

- (small) txt files

- root files with histograms (small comparetd to the input ntuples)

$$H \rightarrow \gamma\gamma$$

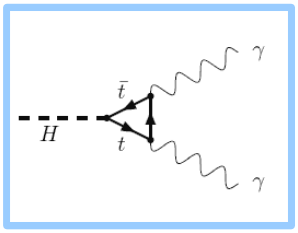

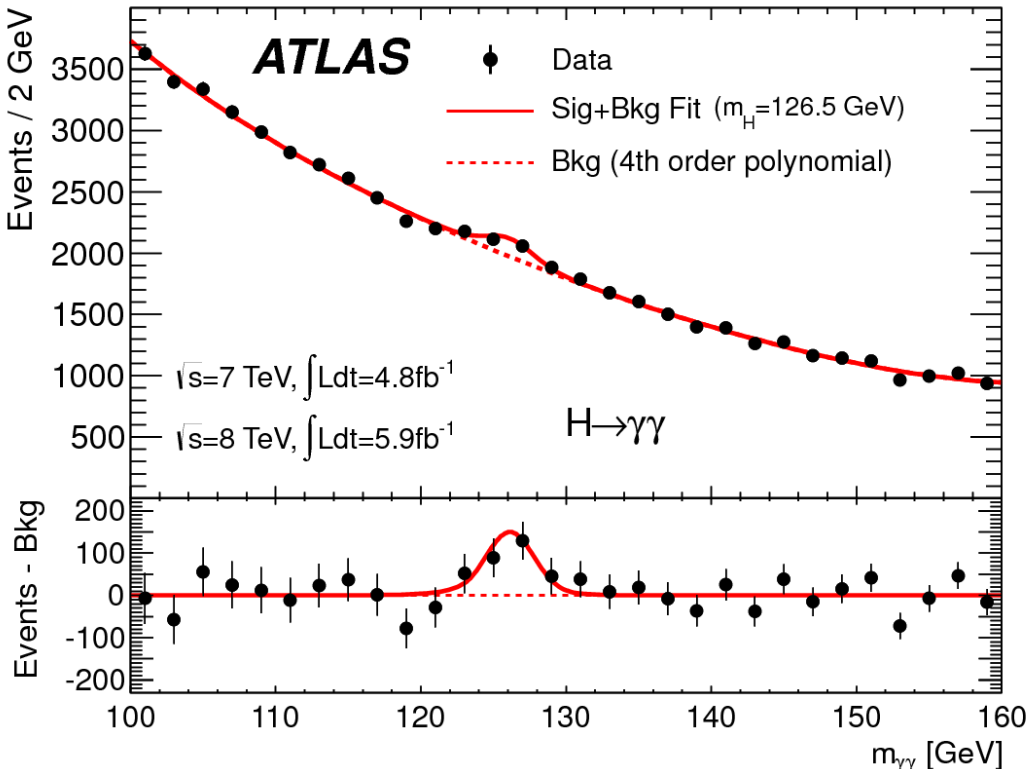

Outstanding Challenges

- **Analysis Code Optimization**
  - Use current analysis code (LipCbrAnalysis), identify possible inefficiencies (way data is read, identify redundant operations, identify obsolete code, etc.) and correct them, without major changes in the underlying framework
  - We need to gain a significant factor in speed in order to remain competitive with increasing flow of data

Going Even Further:

- **I/O Optimization**
  - Is it possible to have a more efficient I/O handling within our framework (we read on average 5TB of data/per analysis)?
  - We crucially depend on ROOT package (not trivial to optimize!) http://root.cern.ch/drupal/
- **Local Resources Optimization:**
  - test the use of Parallel ROOT Facility (PROOF)
  - ...