LHC : from Petabytes to papers

David Rousseau, LAL-Orsay,

LSST week, Lyon 2017

15th June 2017

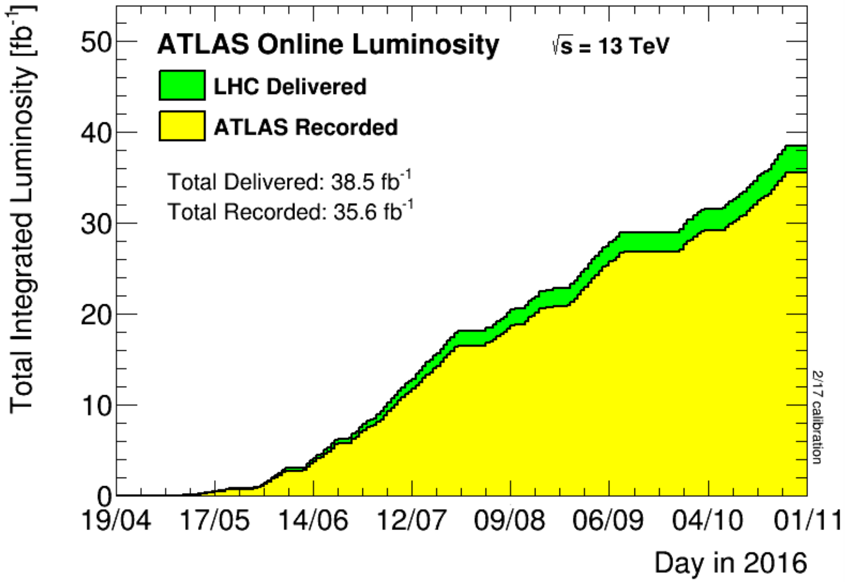

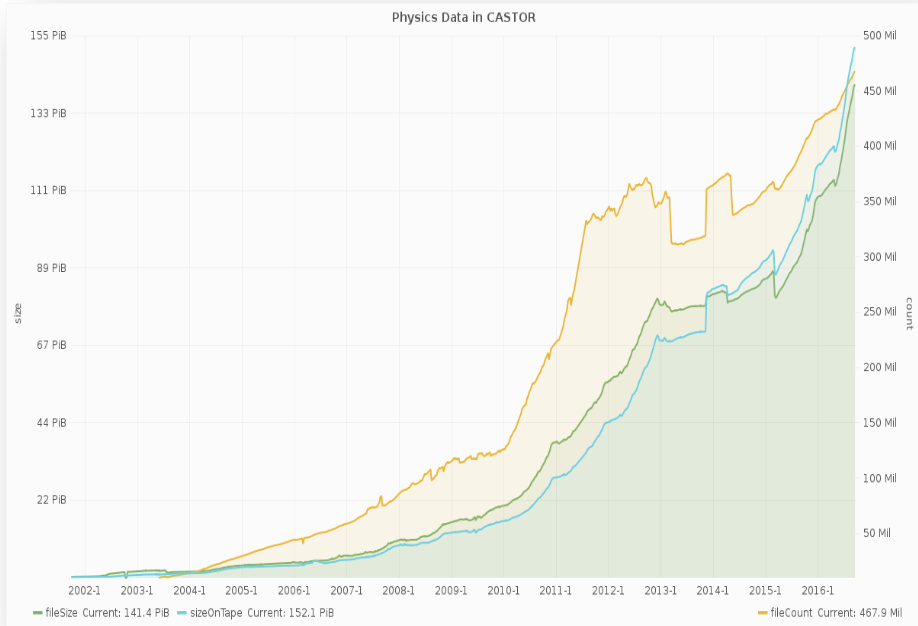

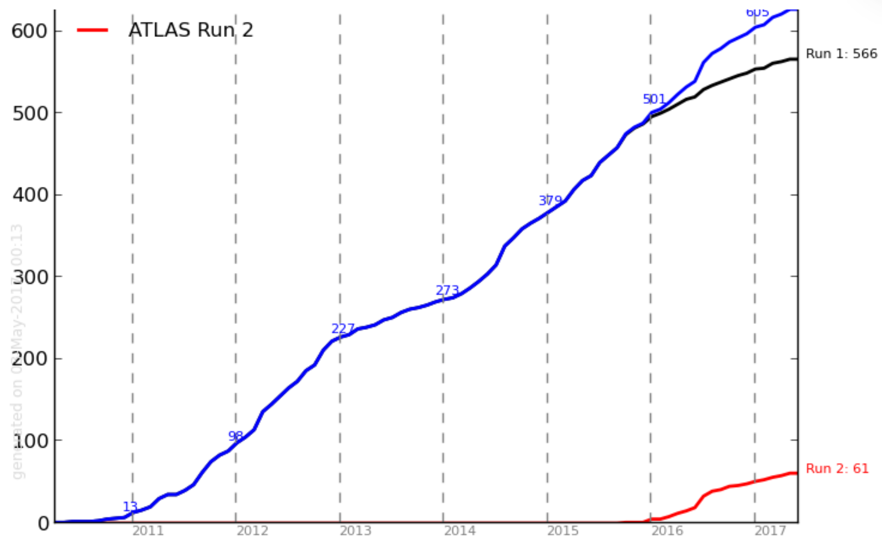

Latest ATLAS 13 TeV results using the full 2015 and 2016 data-set

- Physics Analysis and Performance
- Object Performance and Trigger

This page provides a compact summary of the public results on 13 TeV proton–proton collision data prepared by ATLAS for the 2017 winter, spring and summer conference conferences (from La Thuile to EPS). All results are based on the full Run-II data-set taken in the year 2015 and 20 corresponding to an integrated luminosity of 36.1 fb−1 . For questions, please, contact the ATLAS physics coordinators. See the central publications page for other results from ATLAS.

This page also provides links to ATLAS Physics Briefings for specific results. These are intended to provide a general introduction to ATLAS physics results for a non-specialist audience.

The preliminary results are added only as they are shown by speakers.

Physics Analysis and Performance

Topic	Document type	Short title (link to ATLAS Physics Briefing)	Reference and Link	Release conference
SUSY search	Paper	Search for new phenomena in a lepton plus high jet multiplicity final state with the ATLAS experiment using √s = 13 TeV proton-proton collision data	arxiv-1704.08493	
DM search	Paper	Search for dark matter at √ s=13 TeV in final states containing an energetic photon and large missing transverse momentum with the ATLAS detector	arxiv-1704.03848	
Zprime search	Conference note	Search for new high-mass phenomena in the dilepton final state using 36.1 fb−1 of proton -proton collision data at √s= 13 TeV with the ATLAS detector	ATLAS-CONF-2017-027 (briefing)	DM@LHC
DM search	Conference note	Search for Dark Matter Produced in Association with a Higgs Boson Decaying to bb at √s=13 TeV with the ATLAS Detector	ATLAS-CONF-2017-028	DM@LHC
SUSY search	Conference note	Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in 13 TeV pp collisions with the ATLAS detector	ATLAS-CONF-2017-026	Moriond QCD
SUSY search	Conference note	A search for pair-produced resonances in four jets final states at √s=13 TeV with the ATLAS detector	ATLAS-CONF-2017-025	Moriond QCD
DM search	Conference note	Search for new phenomena in events with missing transverse momentum and a Higgs boson decaying into two photons at √s = 13 TeV with the ATLAS detector	ATLAS-CONF-2017-024	Moriond QCD
Dijet search	Paper	Search for new phenomena in dijet events using 37 fb−1 of pp collision data collected at √s=13 TeV with the ATLAS detector	arxiv-1703.09127 (briefing)	Moriond EW
Wprime search	Conference note	Search for a new heavy gauge boson resonance decaying into a lepton and missing transverse momentum in 36 fb^{-1} of pp collisions at √ s = 13 TeV with the ATLAS experiment	ATLAS-CONF-2017-016 (briefing)	Moriond EW
VLQ search	Conference note	Search for pair production of vector-like top quarks in events with one lepton and an invisibly decaying Z boson in √s=13 TeV pp collisions at the ATLAS detector	ATLAS-CONF-2017-015	Moriond EW
SUSY search	Conference note	Search for squarks and gluinos in final states with jets and missing transverse momentum using 36 fb−1 of √s=13 TeV pp collision data with the ATLAS detector	ATLAS-CONF-2017-022 (briefing)	Moriond EW
SUSY search	Conference note	Search for production of supersymmetric particles in final states with missing transverse momentum and multiple b-jets at √s=13 TeV proton-proton collisions with the ATLAS detector	ATLAS-CONF-2017-021	Moriond EW
SUSY search	Conference note	Search for a Scalar Partner of the Top Quark in the Jets+ETmiss Final State at √s = 13 TeV with the ATLAS detector	ATLAS-CONF-2017-020 (briefing)	Moriond EW
SUSY search	Conference note	Search for direct top squark pair production in events with a Higgs or Z boson, and missing transverse momentum in √s = 13 TeV pp collisions with the ATLAS detector	ATLAS-CONF-2017-019	Moriond EW
SUSY search	Conference note	Search for long-lived charginos based on a disappearing-track signature in pp collisions at √s=13 TeV with the ATLAS detector	ATLAS-CONF-2017-017 (briefing)	Moriond EW
SUSY search	Conference note	Search for Heavy Resonances Decaying to a W or Z Boson and a Higgs Boson in the qq¯(′)bb¯ Final State in pp Collisions at √s=13 TeV with the ATLAS Detector	ATLAS-CONF-2017-018	Moriond EW
SUSY search	Conference note	Search for new phenomena in a lepton plus high jet multiplicity final state with the ATLAS experiment using √s = 13 TeV proton-proton collision data	ATLAS-CONF-2017-013	Moriond EW
Higgs physics	Conference note	Search for the dimuon decay of the Higgs boson in pp collisions at √s = 13 TeV with the ATLAS detector	ATLAS-CONF-2017-014	Moriond EW
B-physics	Conference note	Angular analysis of B0_d -> K*mumu decays in pp collisions at √ s= 8 TeV with the ATLAS detector (NB 8 TeV result)	ATLAS-CONF-2017-023	Moriond EW

Brief LHC introduction

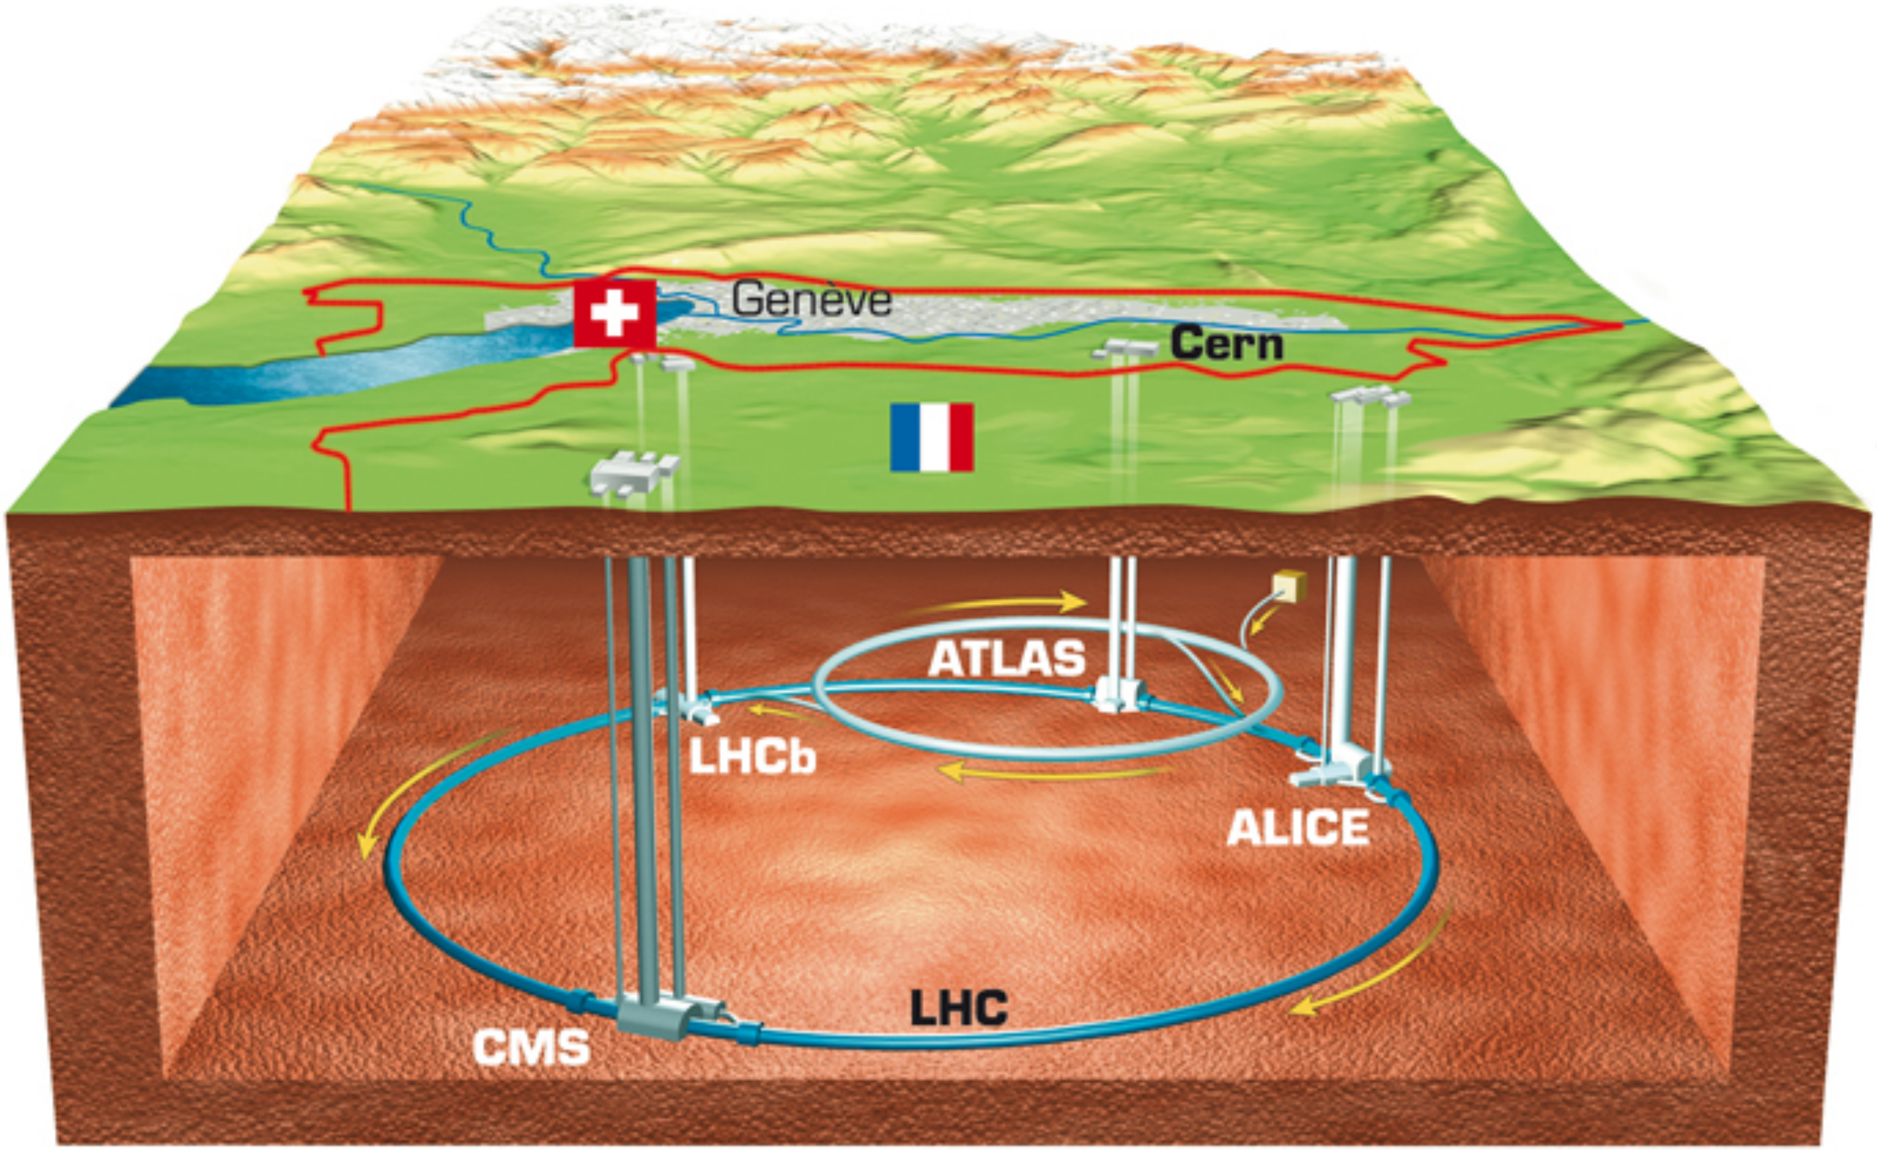

Run 1 : 2010-2012 7-8 TeV Run 2 : 2015-2018 13 TeV

Proton collision

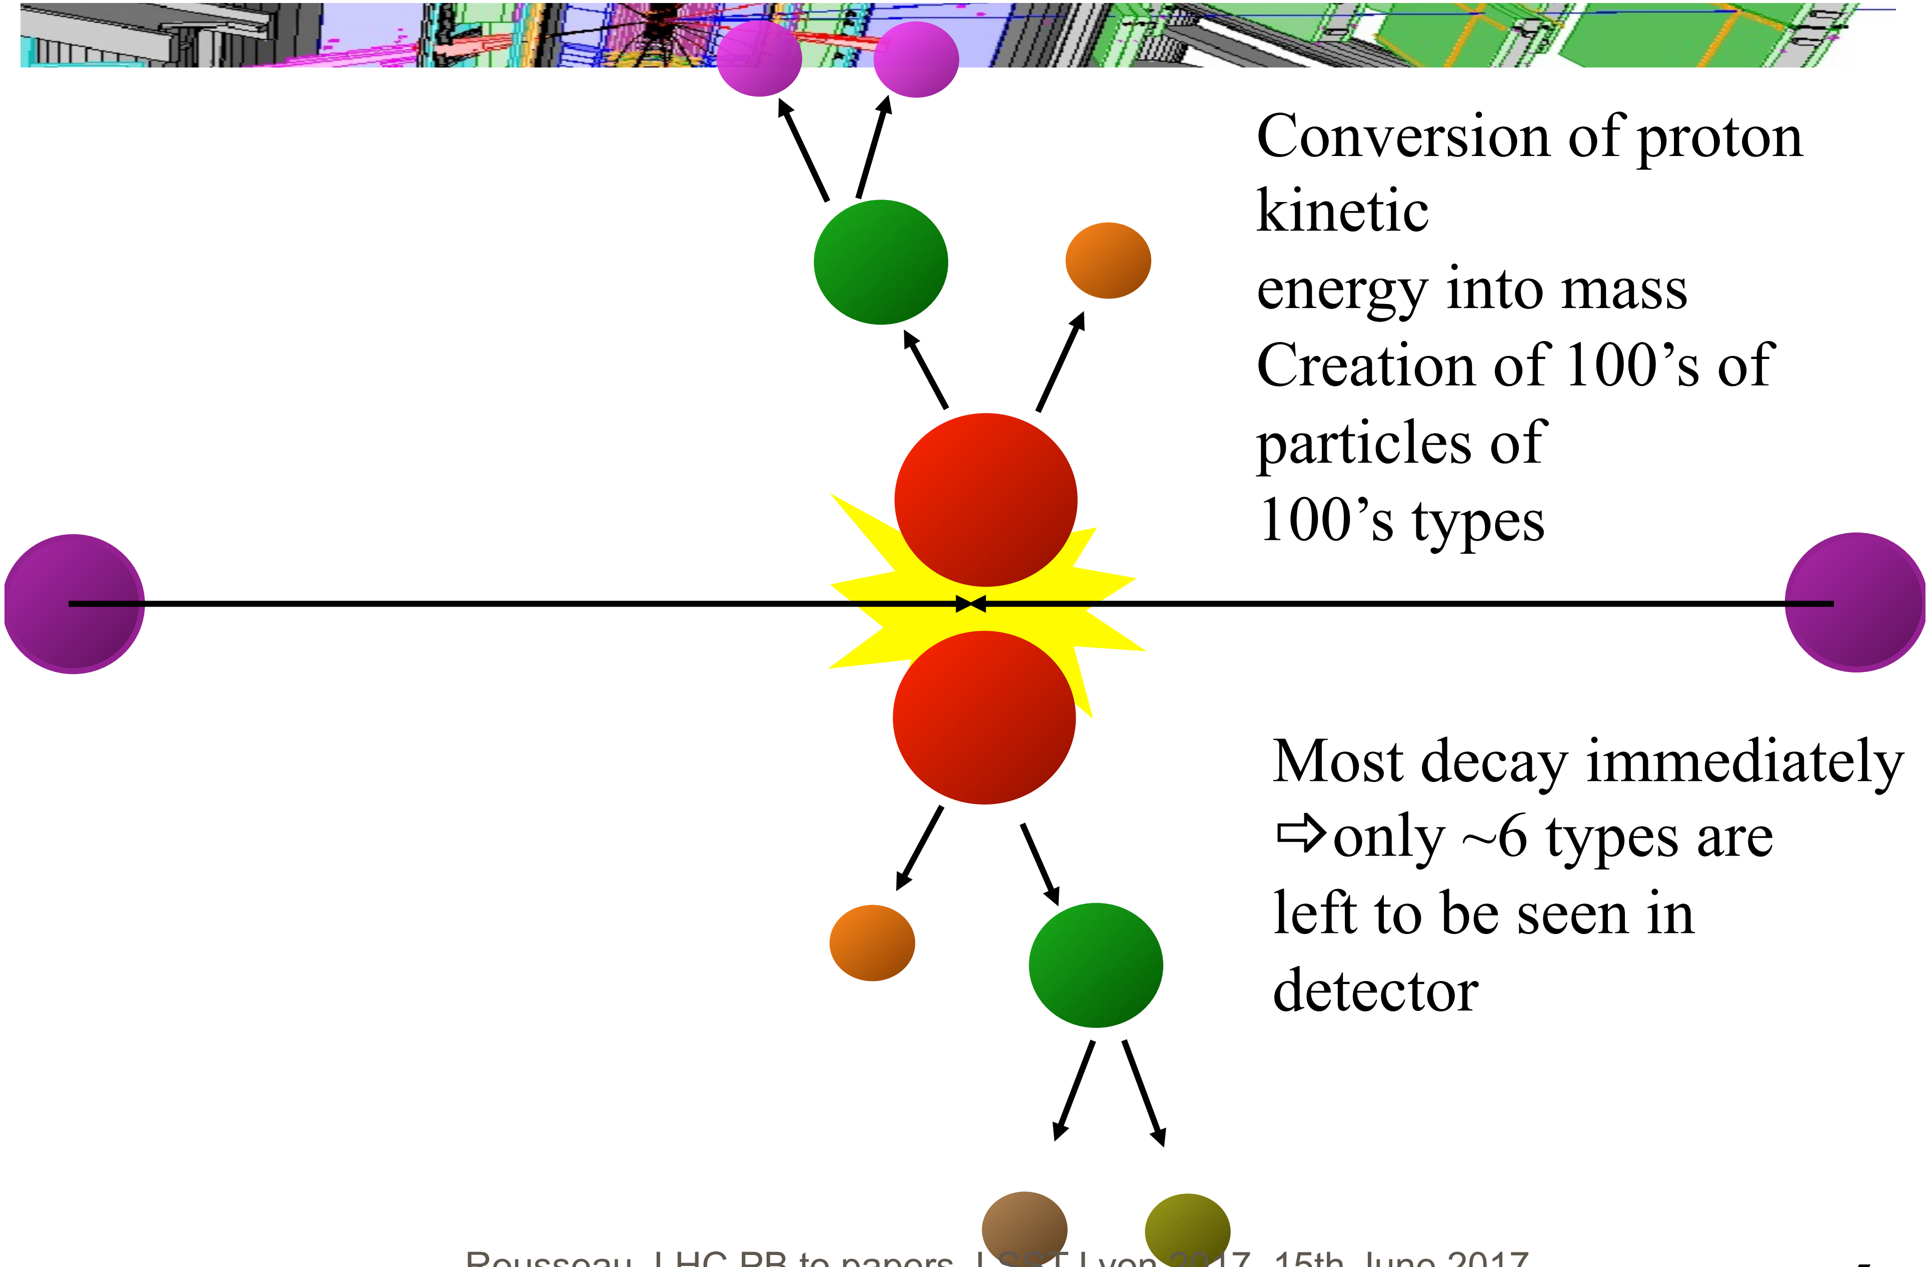

Conversion of proton kinetic energy into mass
Creation of 100's of particles of 100's types

Most decay immediately
⇨only ~6 types are left to be seen in detector

Particles ID

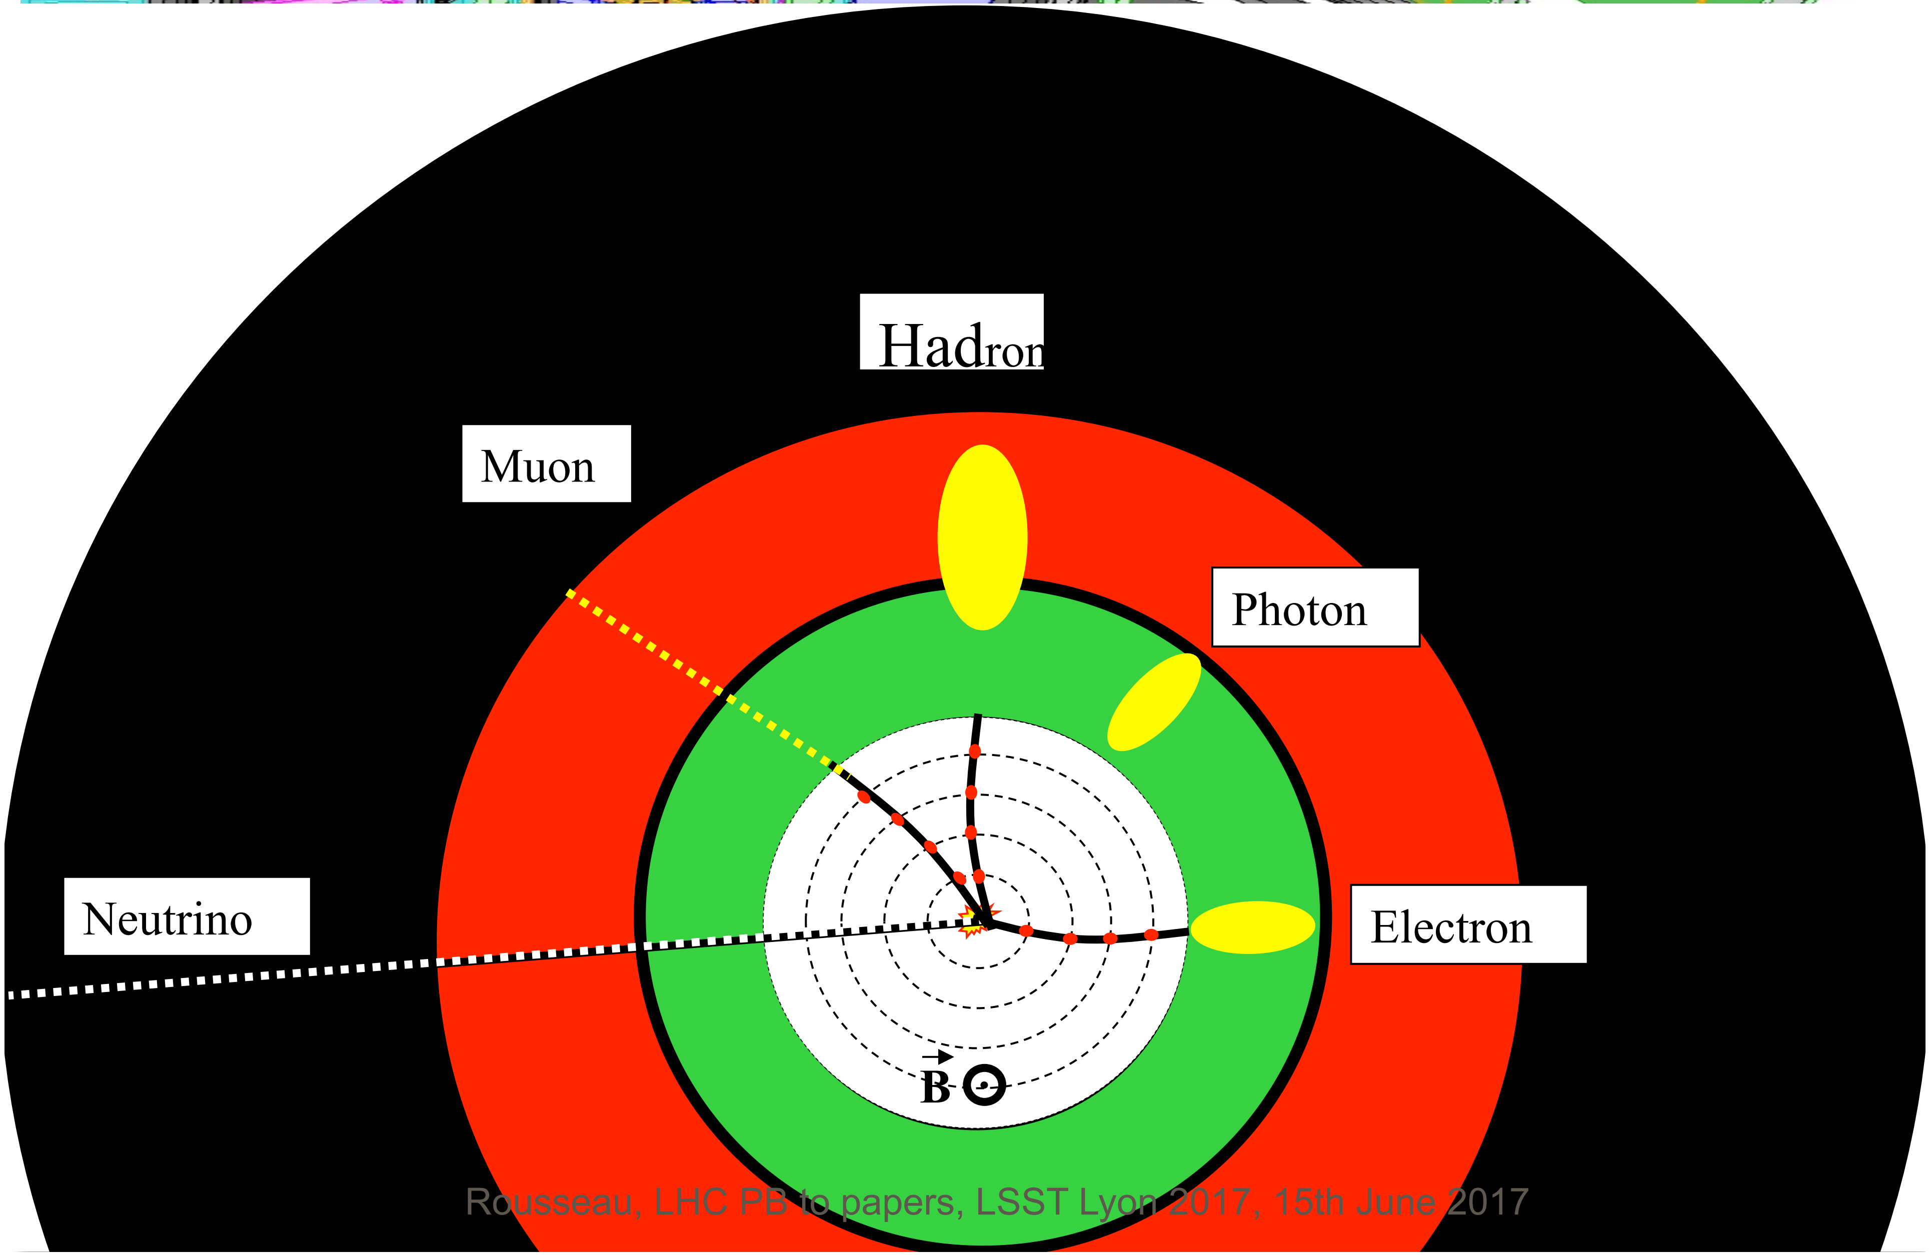

ATLAS detector

Diameter: 25m
Length: 46m
Weight: 7000 tons

3000 km cables
100 millions channels

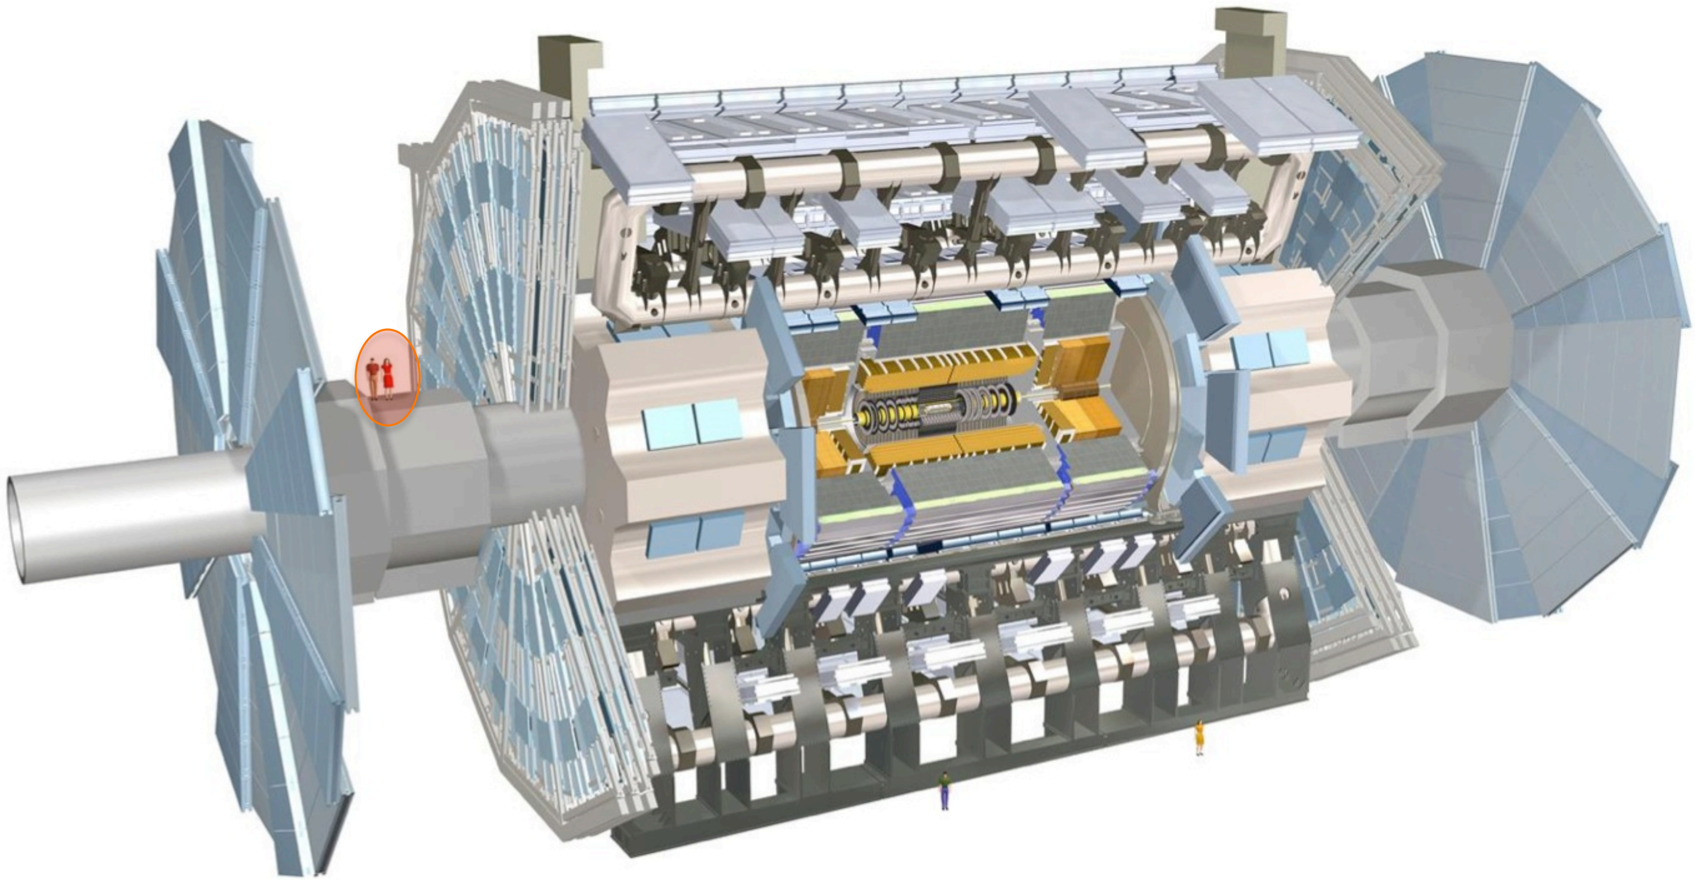

An event

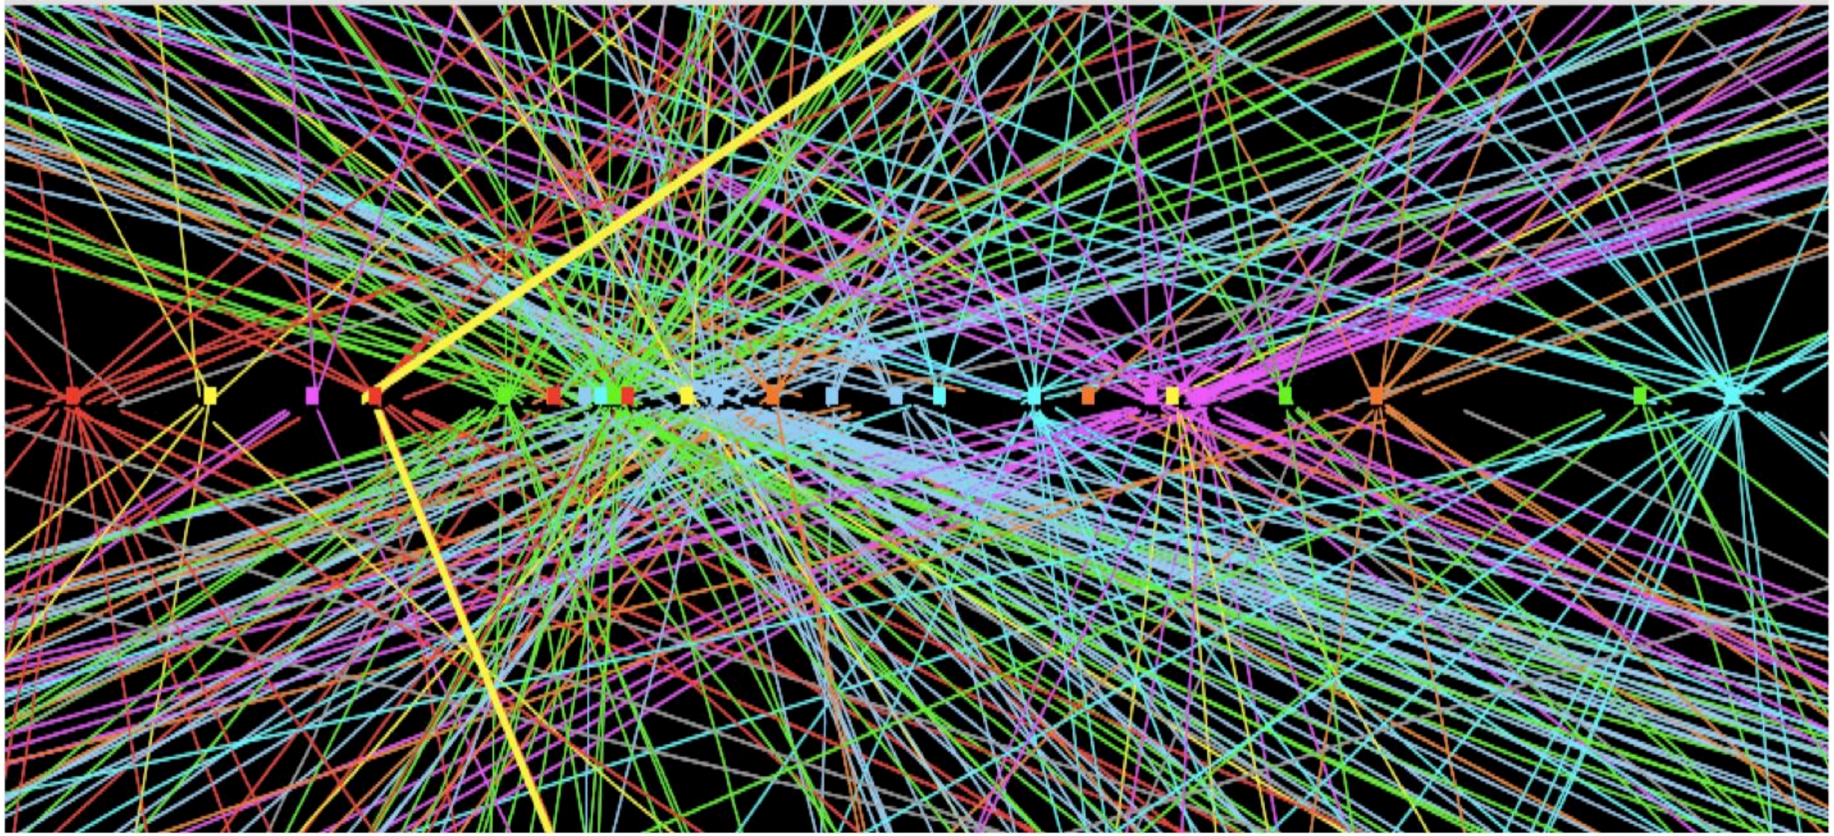

Bunch collision : proton collision of interest accompanied by ~25 parasitic « minimum bias » collision (will reach ~200 @HL-LHC 2015)

One bunch collision = one event, basic unit of treatment

No correlation between events (except for calibration)

Rare events

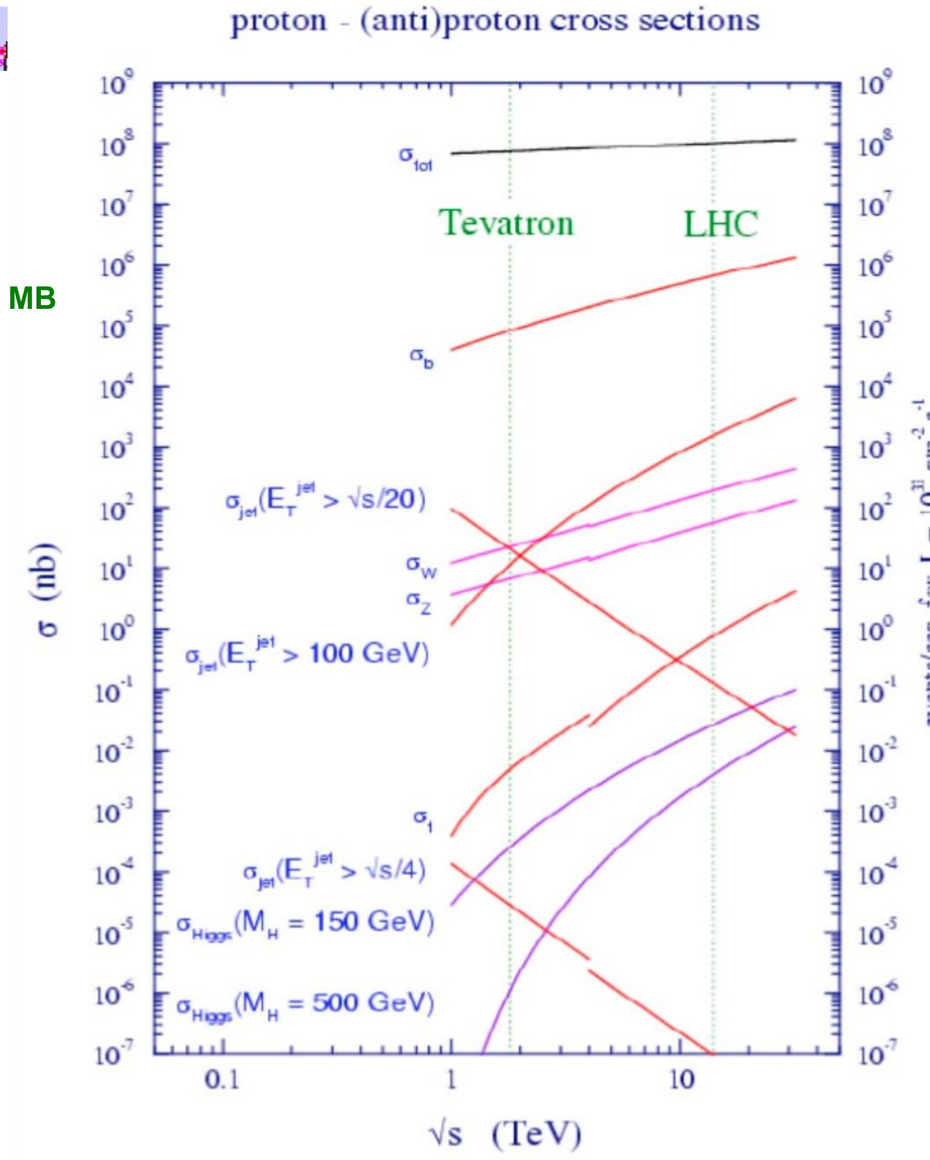
proton - (anti)proton cross sections
MB
σ (nb)
√s (TeV)
Tevatron
LHC
σtot
σb
σjet(ET jet > √s/20)
σW
σZ
σjet(ET jet > 100 GeV)
σt
σjet(ET jet > √s/4)
σHiggs(MH = 150 GeV)
σHiggs(MH = 500 GeV)
0.1
1
10

Trigger

- 40 million proton bunch collision per second
- ~400 / second on disk
- cascade triggering hardware+software, decision latency 1μs-1s, based on ~1000 trigger chains (==sequence of requirements)

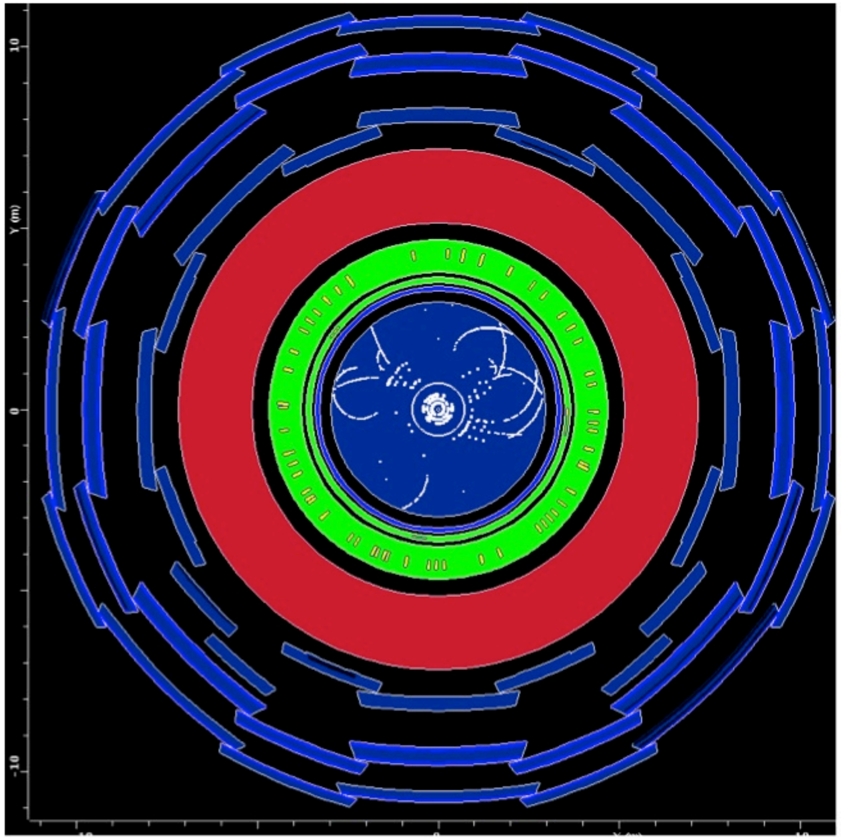

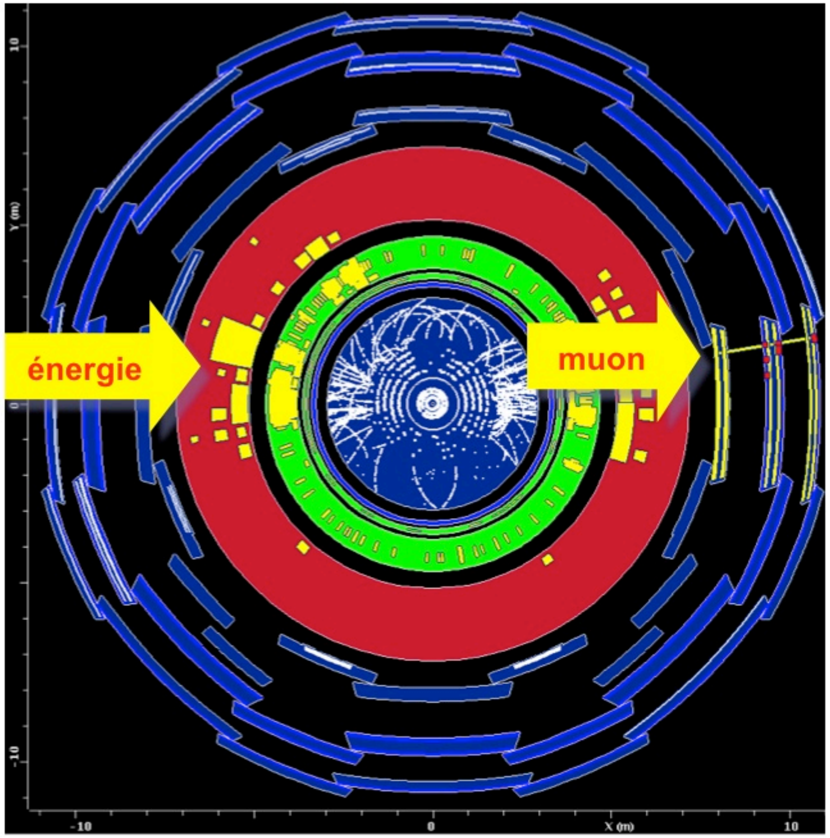

Higgs in a nutshell

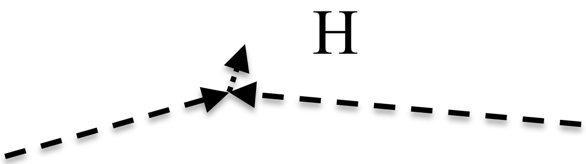

H, before decay

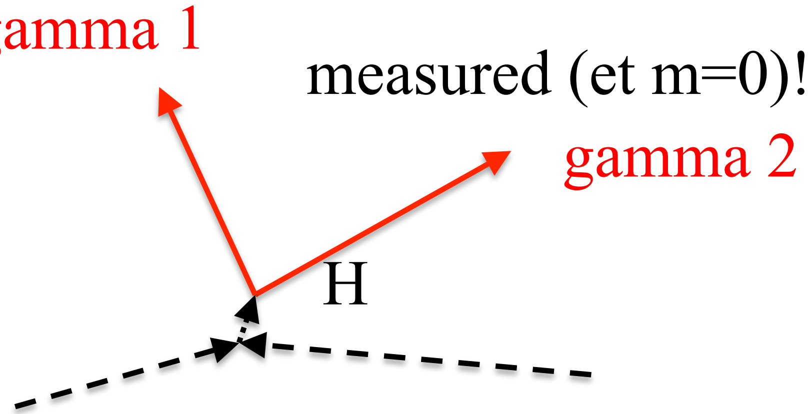

H after decay

$m_H^2 = E_H^2 - p_H^2$

$E_H = E_{g1} + E_{g2}$

$\vec{p}_H = \vec{p}_{g1} + \vec{p}_{g2}$ $\Rightarrow m_H$

Finally

10^{14} collisions

Trigger

10^{9} events on disks

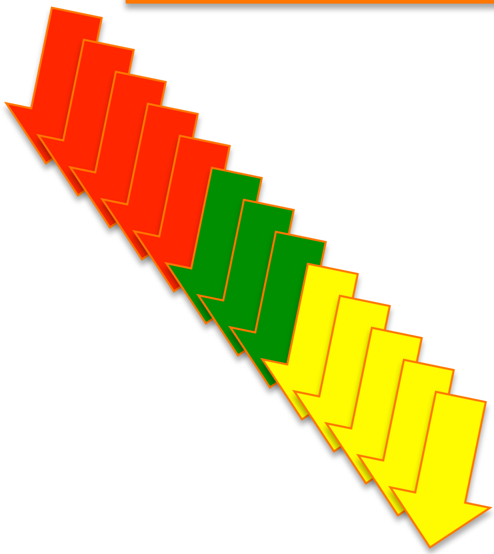

Analysis selection

10^{5} events with 2 gamma

→histogram

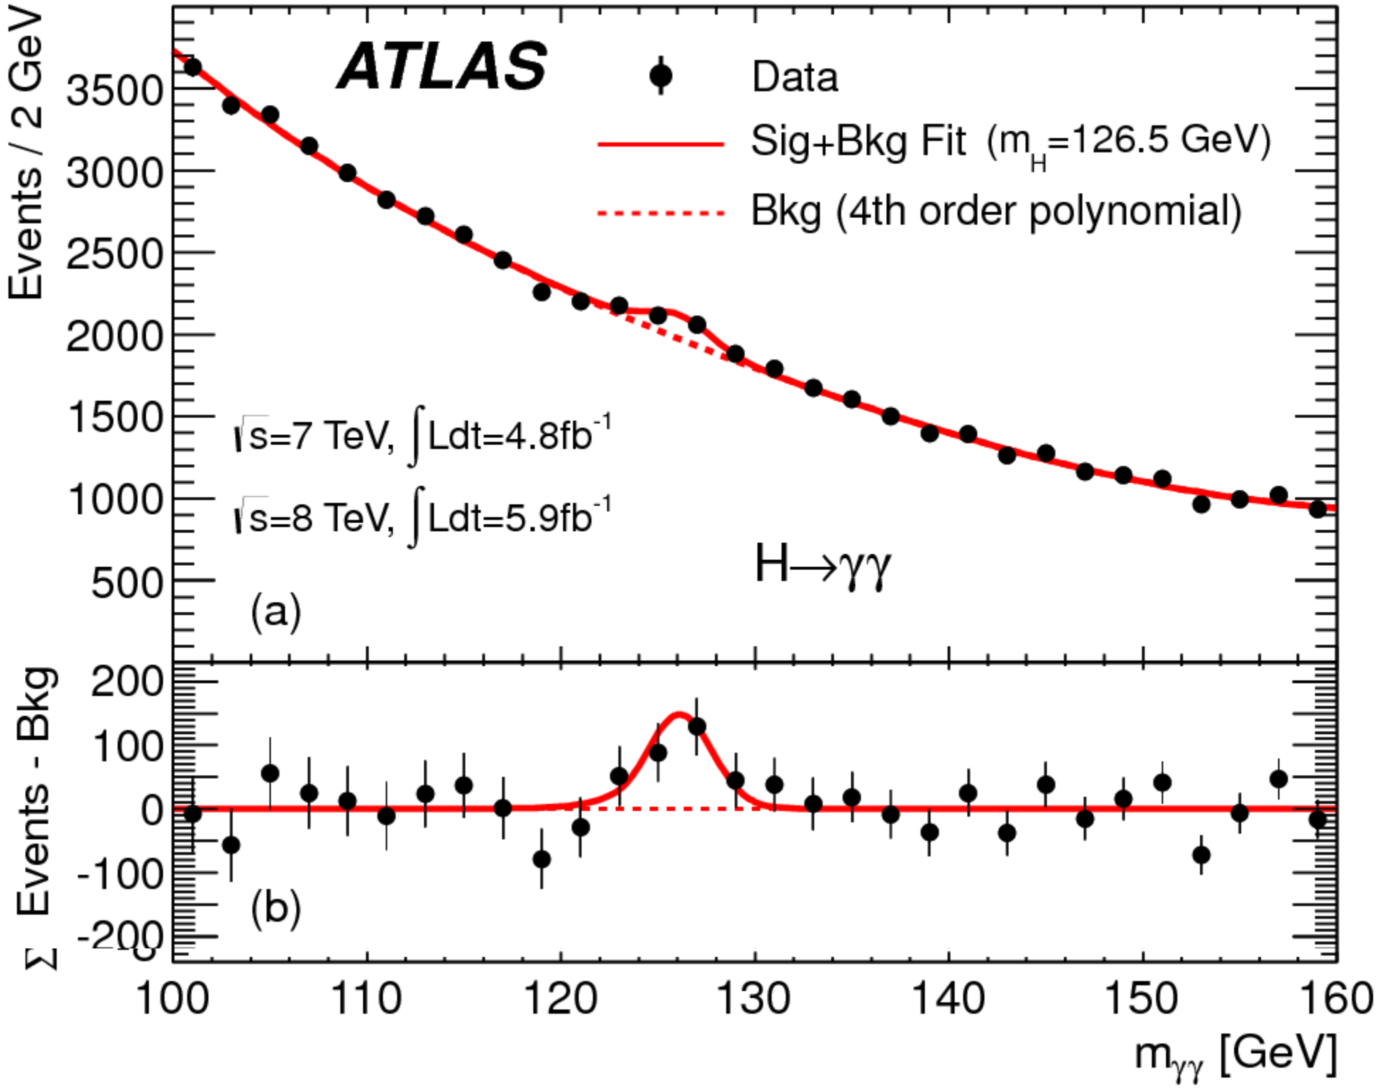
ATLAS
Data
Sig+Bkg Fit (m_H=126.5 GeV)
Bkg (4th order polynomial)
√s=7 TeV, ∫Ldt=4.8fb^-1
√s=8 TeV, ∫Ldt=5.9fb^-1
H→γγ
(a)
(b)
Events / 2 GeV
Σ Events - Bkg
m_γγ [GeV]
3500
3000
2500
2000
1500
1000
500
200
100
0
-100
-200
100
110
120
130
140
150
160

Candidate H→ gamma gamma

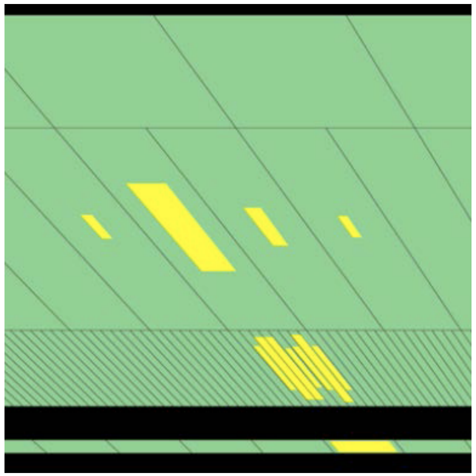

π^0

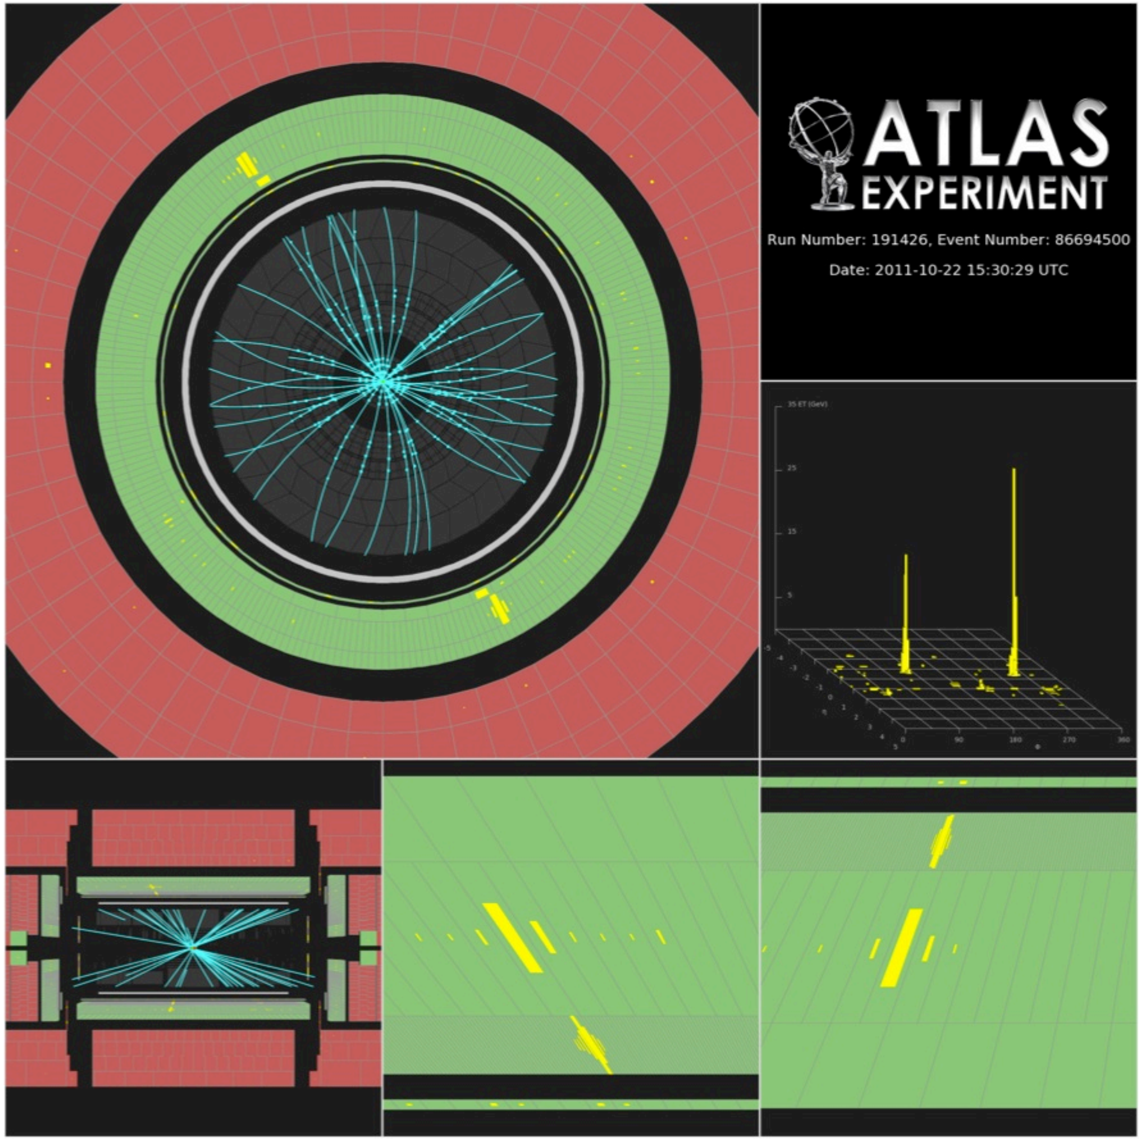

Candidat $H \rightarrow Z(\rightarrow \mu^+\mu^-)Z(\rightarrow e^+e^-)$

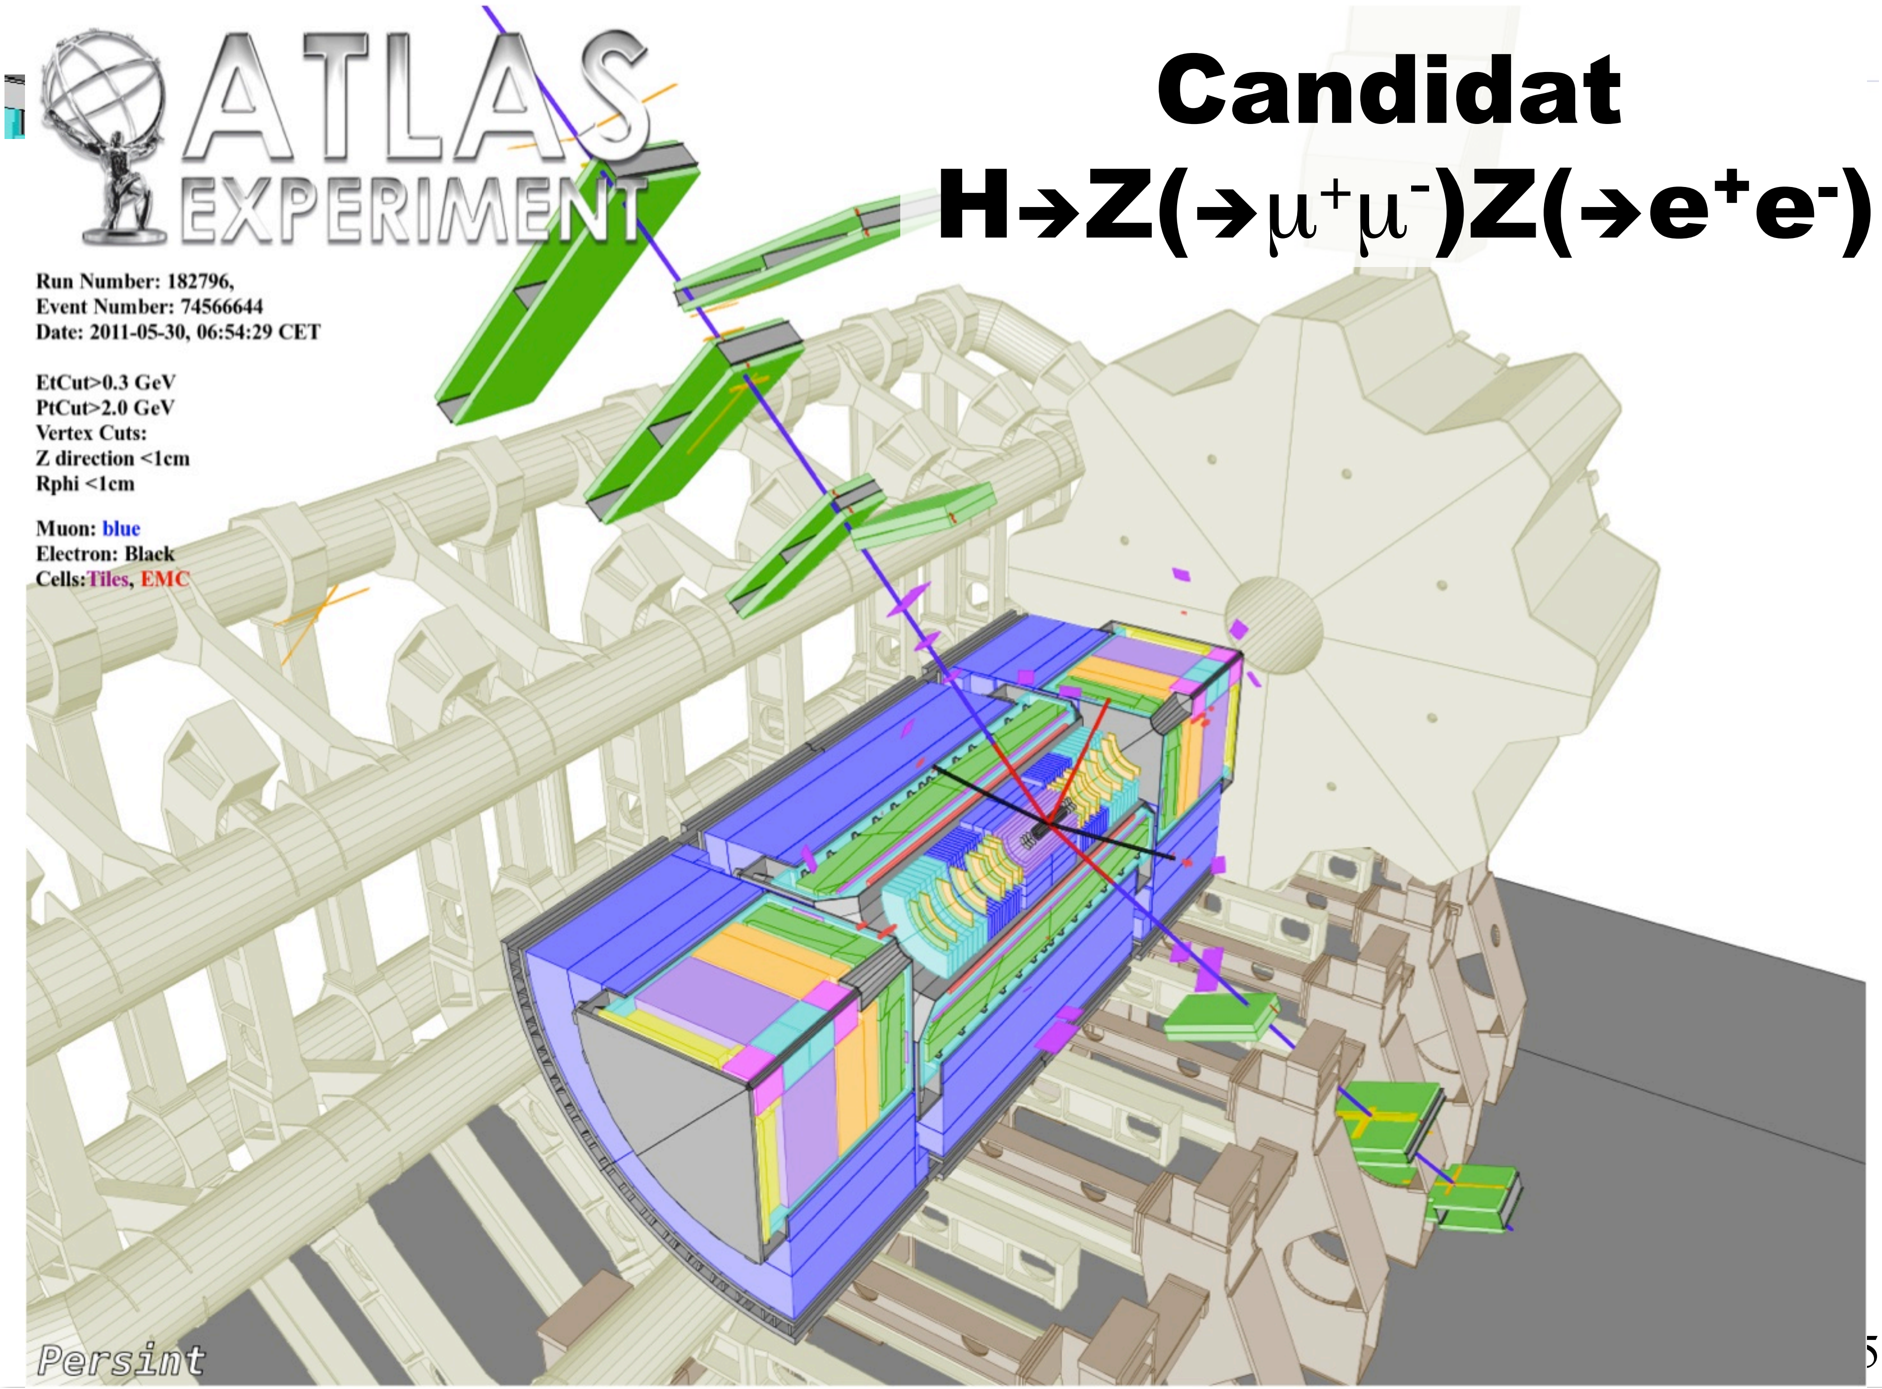

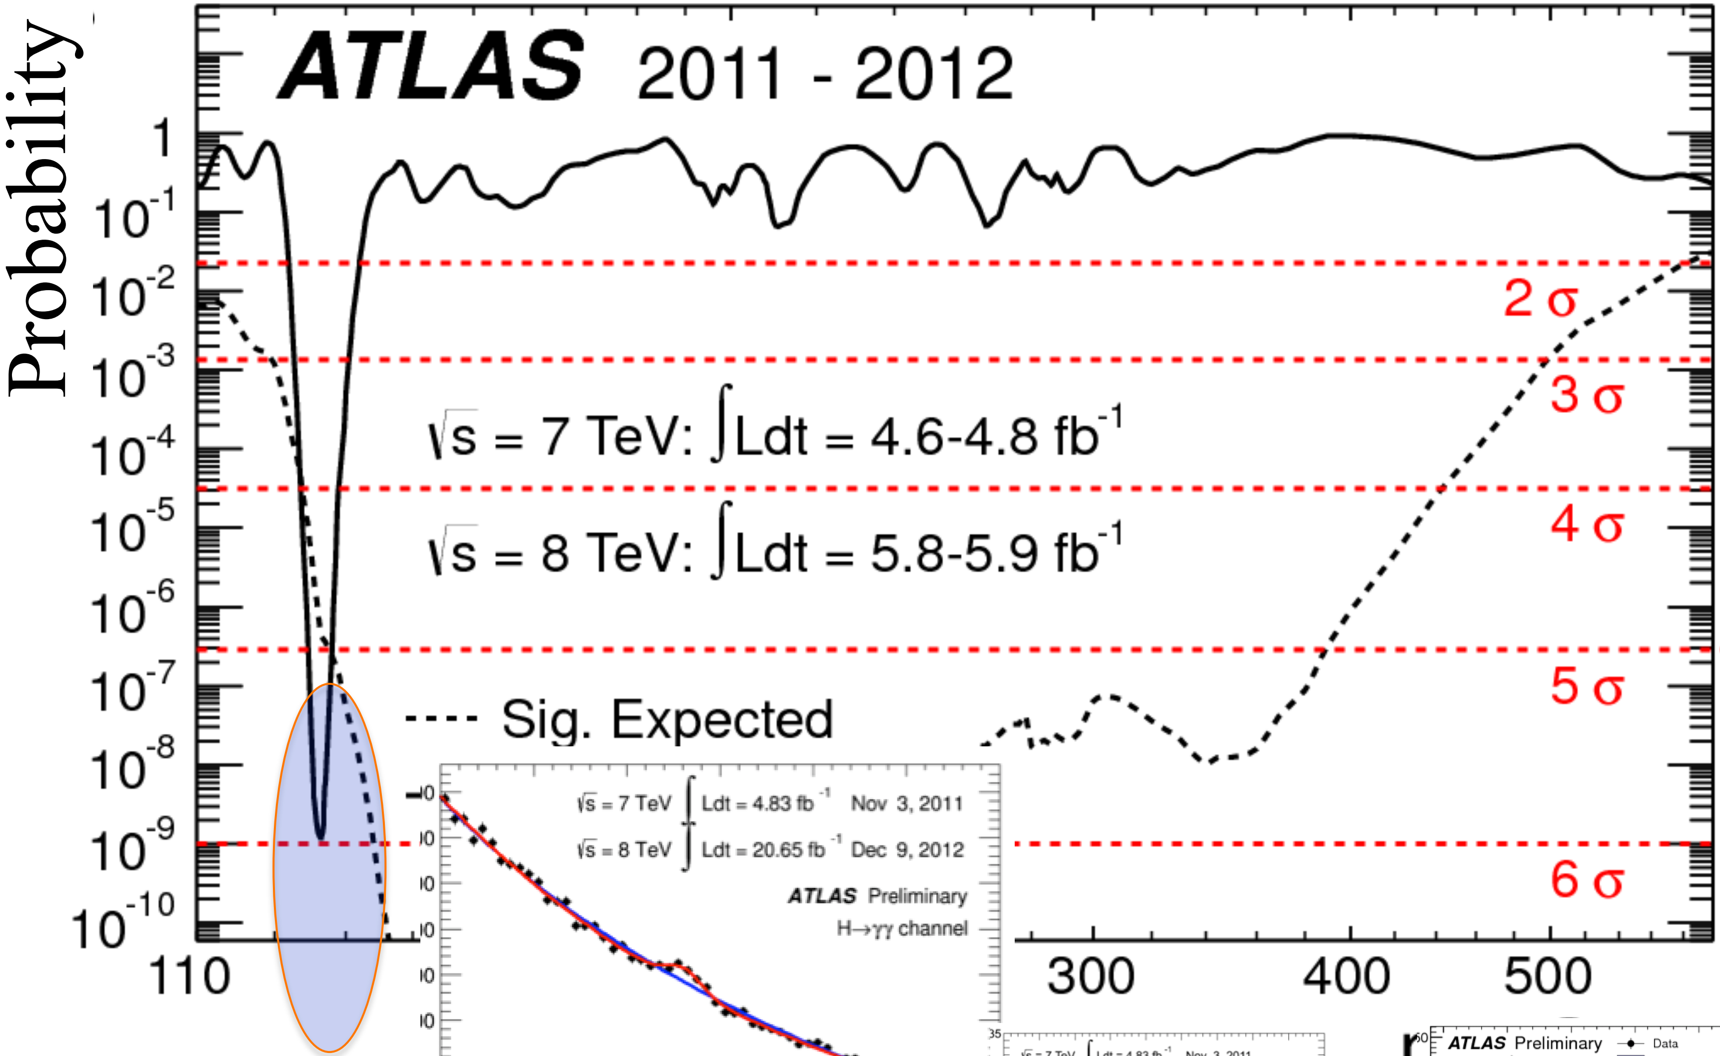

Combine several channels

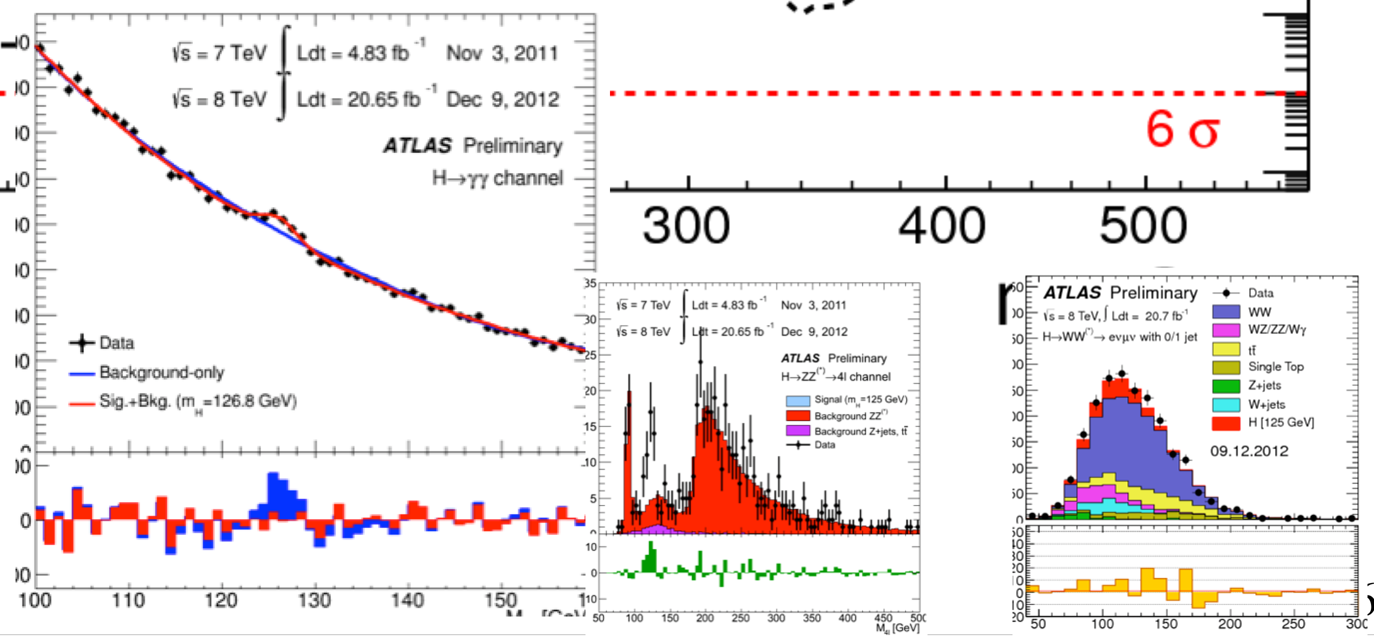

Processing steps

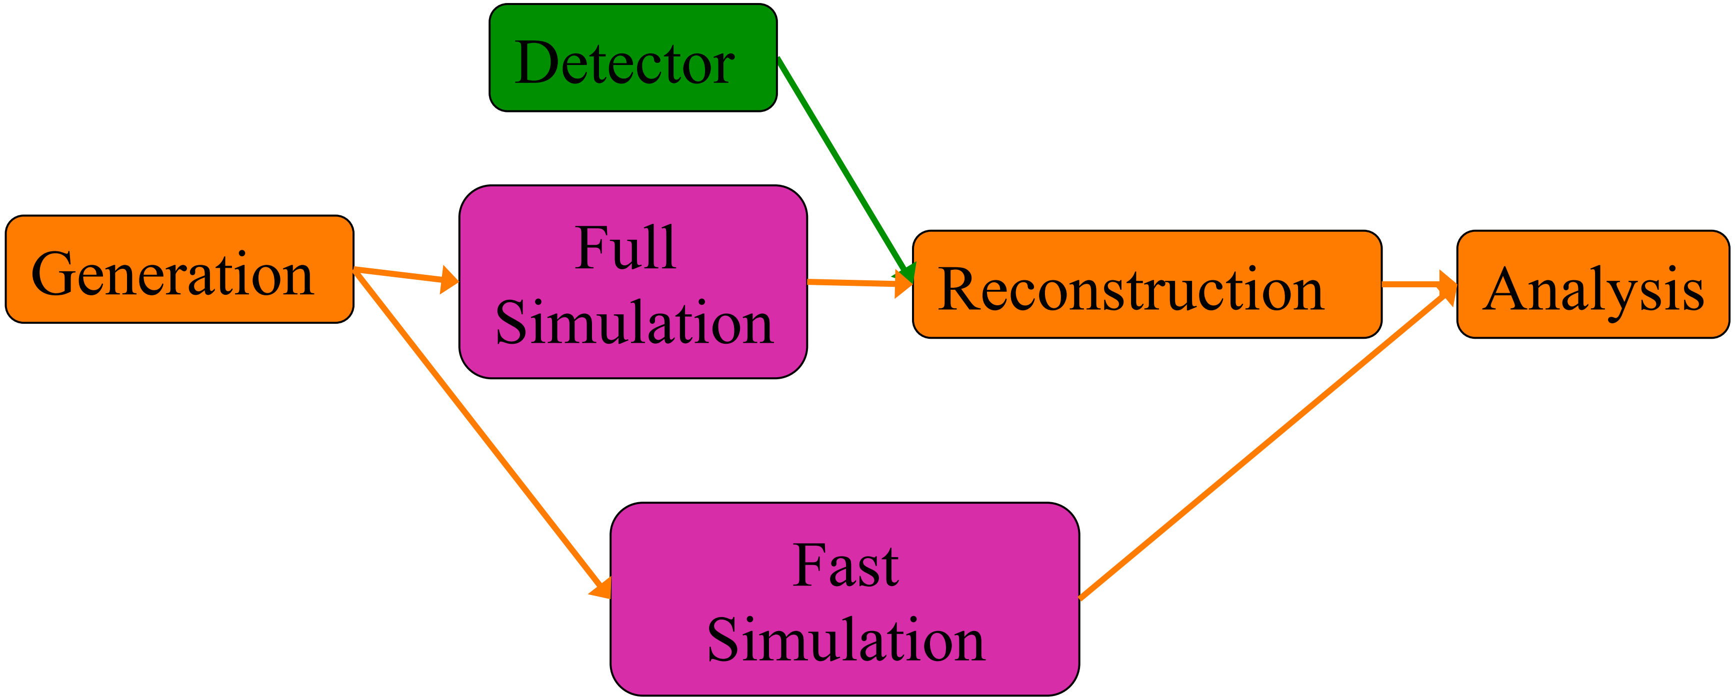

ATLAS CPU usage

300.000 cores

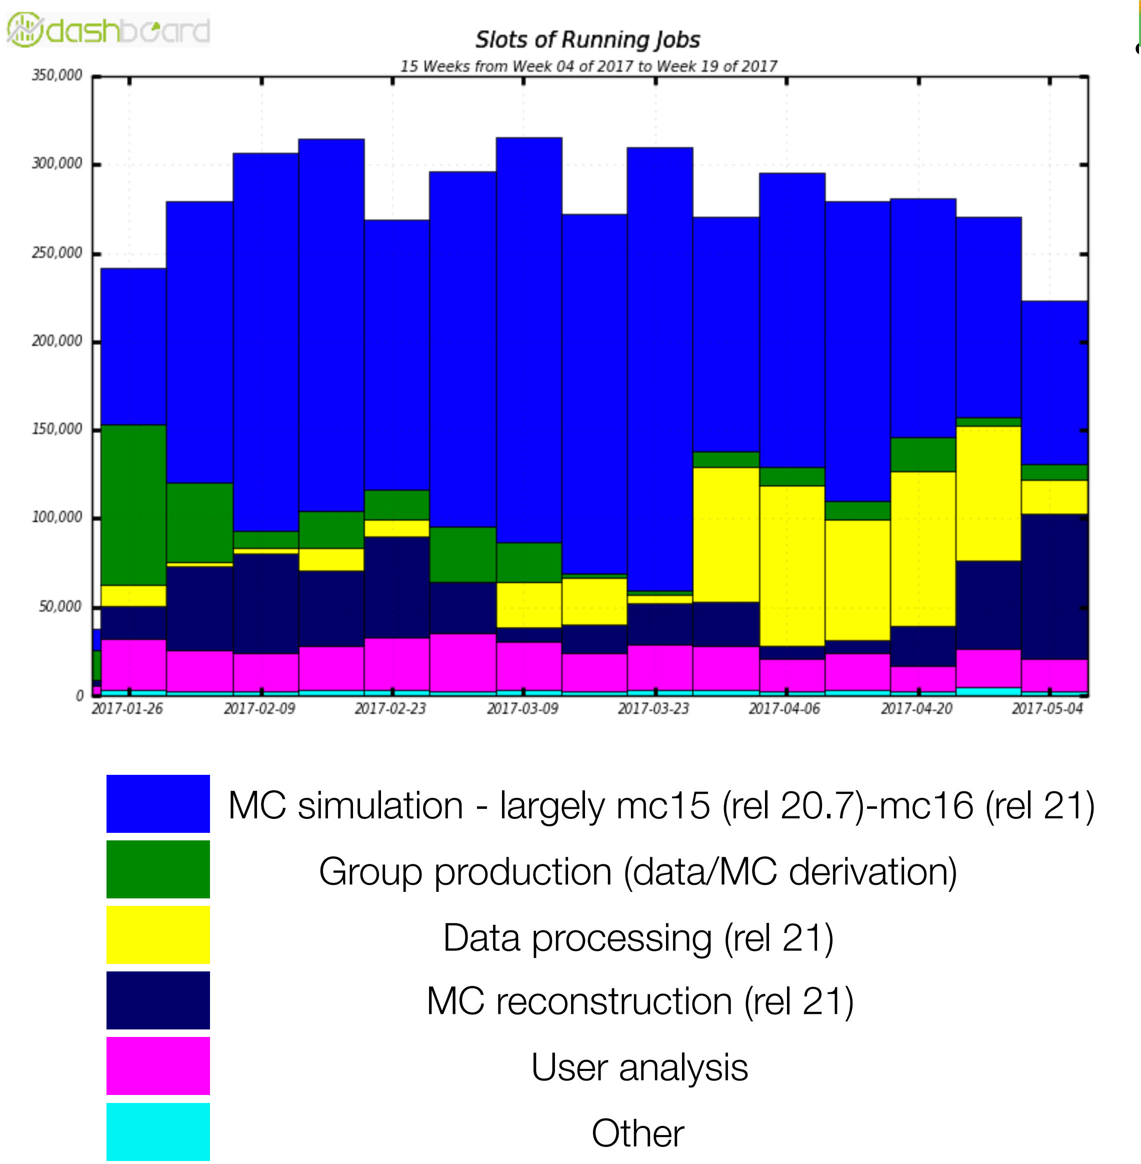

Simulation

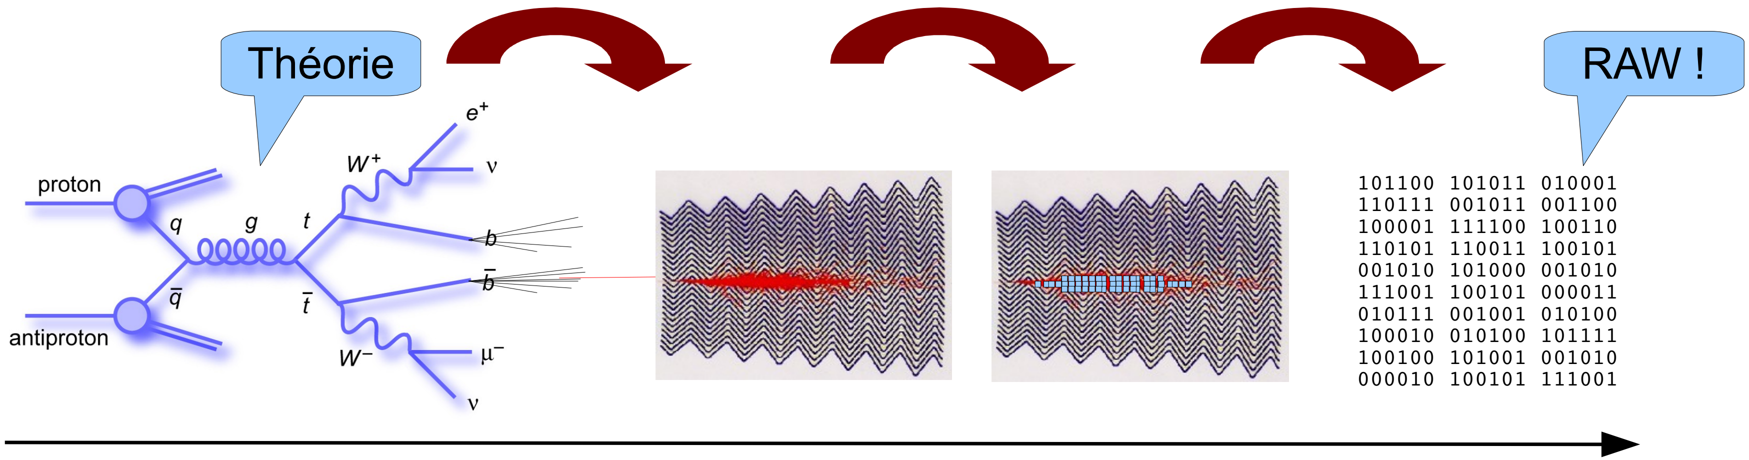

- Rely on blend of G4/Fast sim/Parametric. Challenge : the optimal blend is very analysis dependent. But only one pot of resources.

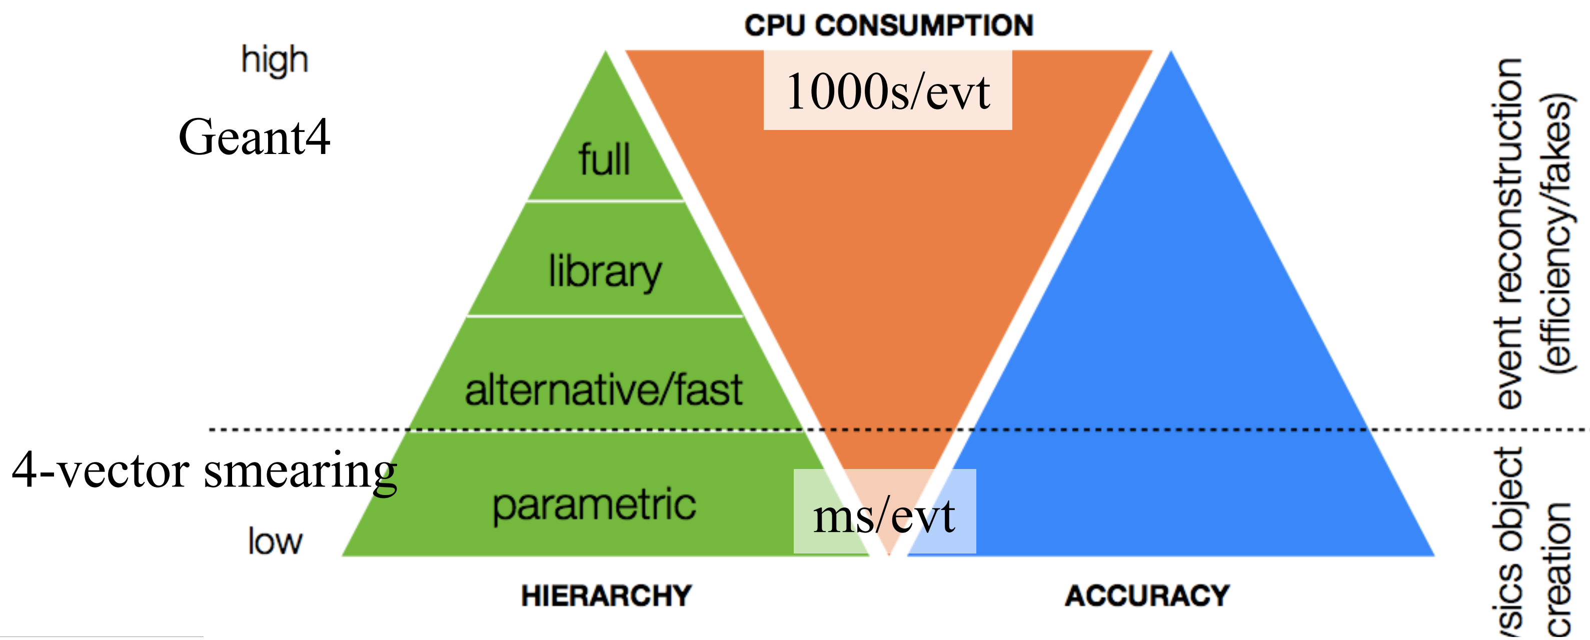

Reconstruction

- Reconstruct analysis objects from raw data
- Budget ~10s ~100kB /event (x ~10E9 data+MC)
- Analysis object to satisfy 99% of analysis
    - At the beginning, we had several alternative algorithm running in parallels
    - ➔after a few years (and big pressure from resource usage, and physics harmonization), only one left
    - ➔still under developments in particular pile-up mitigation, preparation for upgrade, and new ideas (Machine Learning)

Calibration work

Study of known particles : in particular Z and top

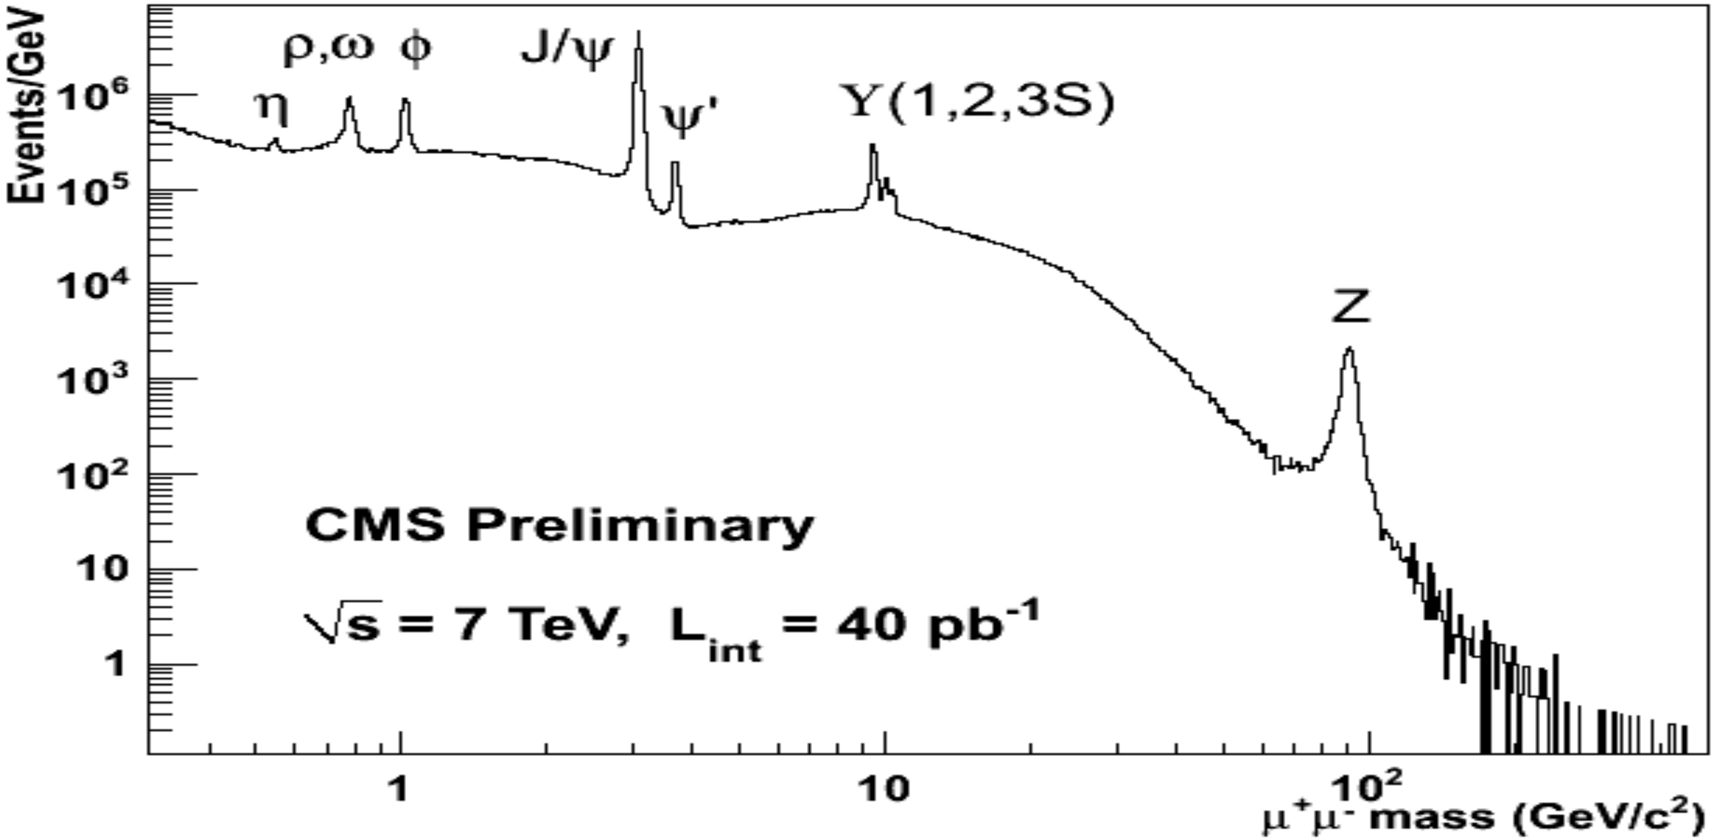

Calibration work

- “Combined Performance” groups
  - Develop ID/reco algorithms (e.g. electron, tau) : impact xAOD
  - Emit “recommendations” (final corrections, and systematics) : to be applied at final stage

Pileup stability

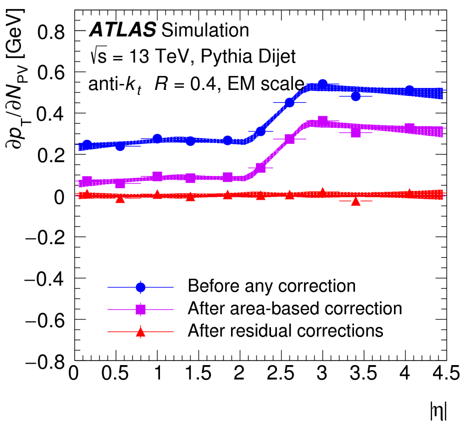

Muon reconstruction efficiency

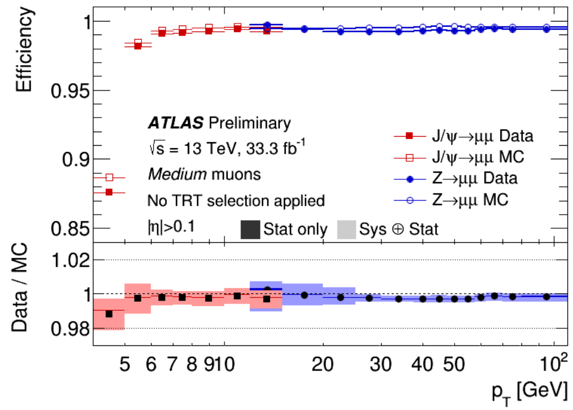

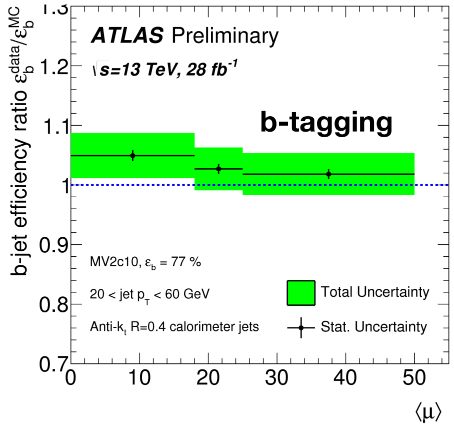

Jet energy scale

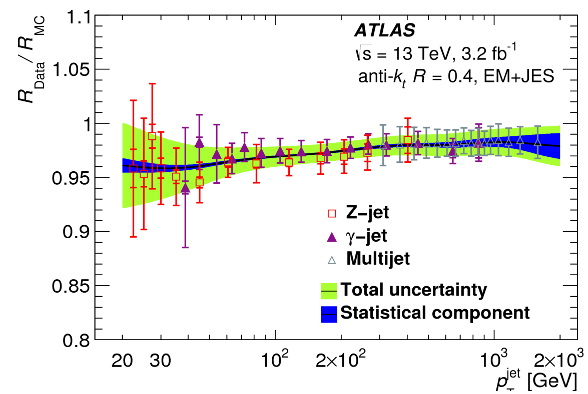

Conditions data

- Hundred of thousands of calibration/alignment parameters needed to reconstruct the data
  - Evolving during data taking (temperature, aging…): march 2012 parameters are not the same as october 2012 ones
  - Evolving during post treatment : march 2012 parameters determined in march 2012 less correct than the one available in october 2013
- "48 hours calibration loop"➔fast computation of calibrations, input to first reconstruction
- More involved calibrations prepared for future reprocessing
- Hierarchical system of tags :
  - one (immutable) master tag uniquely determines all the parameters of the running period (with their evolution)
  - Newer master tag when more precise parameters are made available
- Database (Oracle) keep ~forever data for all past and present tags
  - In practice, regular house cleaning of buggy data

"Frozen Tier0" Policy

- Accuracy vs stability
- "Tier0"==6000 cores in CERN computing center running reconstruction within 48 hours of data taking
- The individual analyst does not want/cannot follow what is happening in the 2500 packages
- Production "Frozen Tier0" release e.g. 17.2.X.Y is defined
  - The release evolved but is guaranteed to give bit by bit identical results (automatic checks to enforce it)
- Allowed changes:
  - Better, less resource hungry code (cpu, memory, disk)
  - Possible crash fixes
  - Bugs are deliberately left unfixed (!!!!)
  - Adding information to output allowed (within reasons)
- →data from the detector and simulation have all consistent content for months, until a major reprocessing is done (once a year now, even less)

Software validation

- It is much more difficult to check something is correct in absolute than to compare it to something which is deemed correct
- Every night the future release is built and and automatically compared to a reference
- A developer is supposed to warn when his new code is supposed to change things, and to check that it did what it was supposed to do, but there are side effects.
- People take shifts to monitor the result of the nightlies, ask experts if something is wrong
- New model being put in place : as above but replace "nightly build" by "git Merge Request triggered build"
- Every ~month the release goes through extended validation:
  - Standard plots built on high statistics
  - Semi-automatic checks
  - Perused by experts with consultation with their community e.g.:
    - Electron energy is slightly higher now ? Yes there is a new correction algorithm enabled
    - We find less muons now ? Strange nothing has changed. (…investigation, then…). This is most likely due to that change in the tracking which was supposed to be inocuitous (…more investigation…). This is fixed by InDetRec-01-02-03

Managed sw development

SVN with TagCollector

==> moving to gitlab and Merge Requests

Status	BundleName	Action	Package tag	Requested by	Your decision	Request Comm
validating			BTagging-00-06-46 Dependencies OK last run :		Accept Pending Postponed Request Feedback r validation tests nightly	(Cuhadar Donszelmann Tula **justification:** Avoid crash **Savannah bug report (edit** **private code changes:** [yes **public interface changes:** [**python changes:** [yes] **requirements changes:** [ye **other changes:** [none]
validating		update new_request	METReconstruction-00-00-14 Dependencies OK last run : 2014-03-22 at 09:12:24 PM CET Test Dependencies current: METReconstruction-00-00-13 SVN - Trac_Diff	Tulay Cuhadar Donszelmann email	Accept Pending Postponed Request Feedback Reject Passes developer validation tests after one successfull nightly Comment Submit	(Cuhadar Donszelmann Tula **justification:** 'addition to tau change in jet EDM' **Savannah bug report (edit** **private code changes:** [nor **public interface changes:** [**python changes:** [none] **requirements changes:** [nc **other changes:** [none]
validating (bundle across projects) Completed bundle	vakho_22March_bundle	update new_request	TauDiscriminant-01-07-46-06 Dependencies OK last run : 2014-03-22 at 07:42:53 AM CET Test Dependencies current: TauDiscriminant-01-07-46-05 SVN - Trac_Diff	Tulay Cuhadar Donszelmann email	Pending Postponed Request Feedback Reject Passes developer validation tests after one successfull nightly Comment Submit	(Coordinator) **justification:** problem following recent cha xAODJet **Savannah bug report (edit** **private code changes:** [yes **public interface changes:** [**python changes:** [none] **requirements changes:** [nc **other changes:** [none] **completed bundle:** [yes]

Impact of pileup on reco

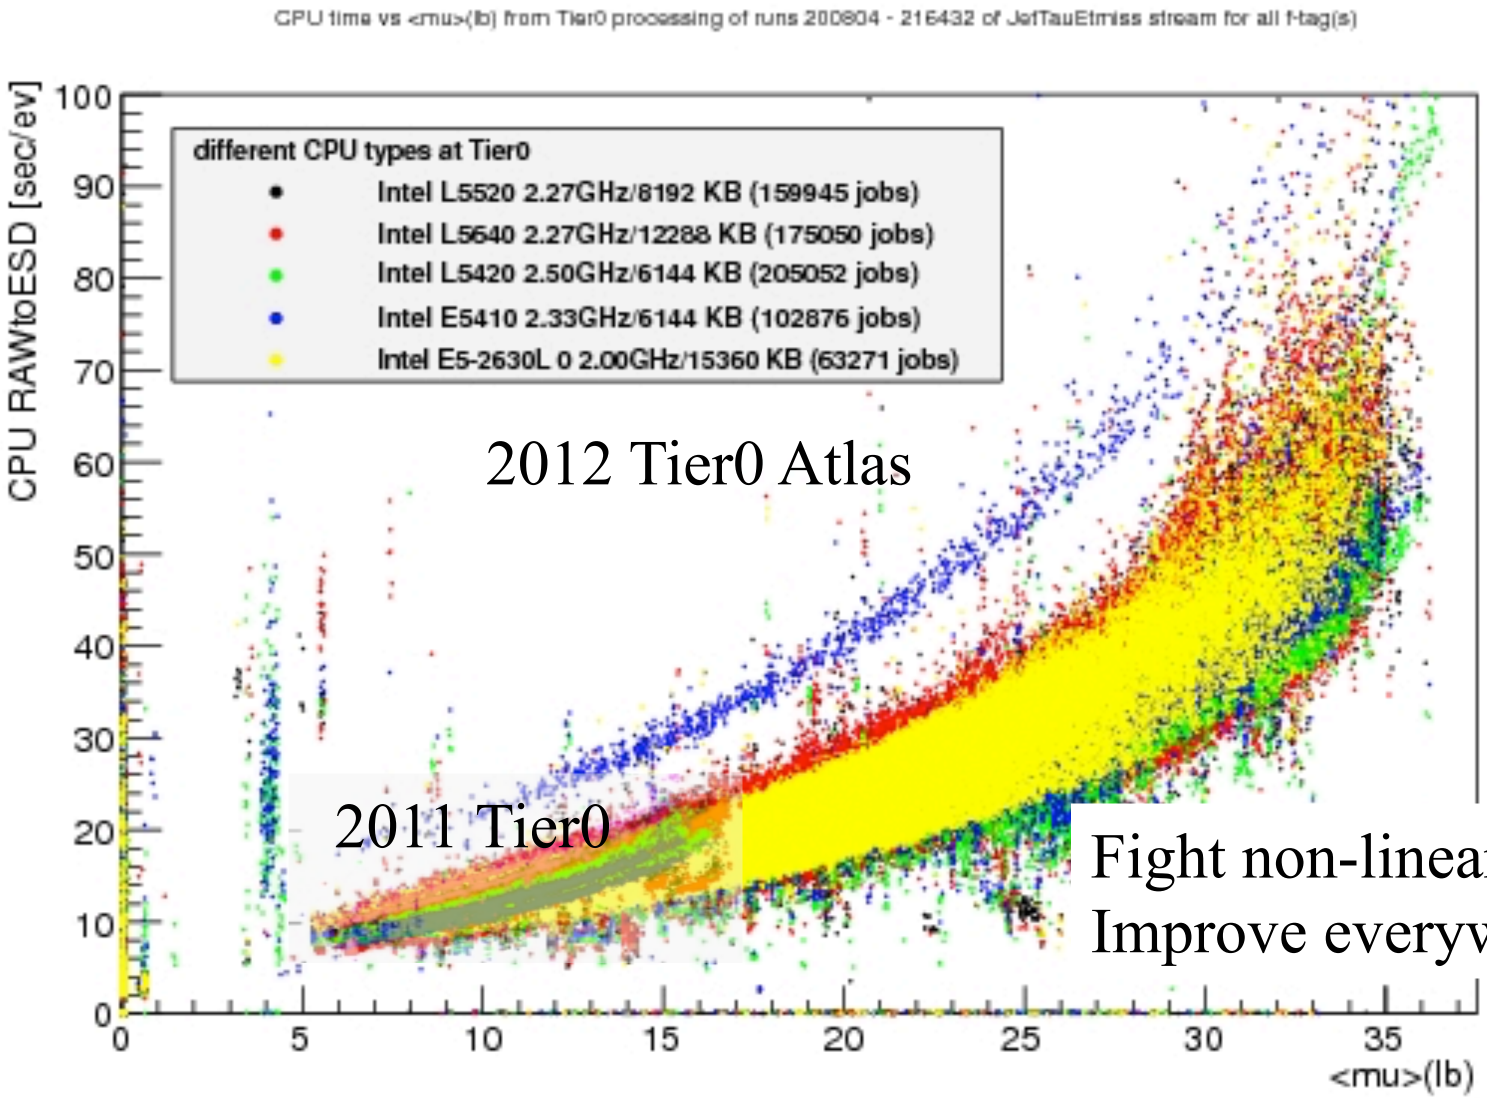

Fight non-linearities
Improve everywhere

200→

Analysis cycle

x Petabytes

Creativity in physics analysis : having a lot of ideas…and…being able to test them quickly

KiloBytes

RAW

How much time does it take to redo a plot ? (new cut, new variable,…)

How much time does it take to redo a full analysis ? (proper reweighted plots and values after all corrections)

plot

1-3 months

Every

4-12 months

Few seconds every minute

root

Petabytes

Analysis objects

Difficulty : ~100 analysis streams/teams sharing same resources

dedicated Ntuple

Giga/MegaBytes

~1000

Unique for 99% analysis

One or more dedicated intermediate datasets with event selection and information reduction. Balance between completeness, **resource** usage, speed to use and speed to reproduce the dataset.

Analysis software

- Two levels of analysis
  - Event combinatorics (e.g. compute Higgs candidate mass per event):
    - Need explicit event loop
    - Go through 100 E6 real and simulated events
    - On shared resource
    - Root or dedicated frameworks CMSSW, Gaudi/Athena
  - Final analysis
    - no explicit event loop
    - On physicist' laptop
    - TTree->Draw
    - histogram manipulations
    - Limit/signal setting RooFit/RooStat
    - Root definitely the framework here

Derivations (new in Run 2)

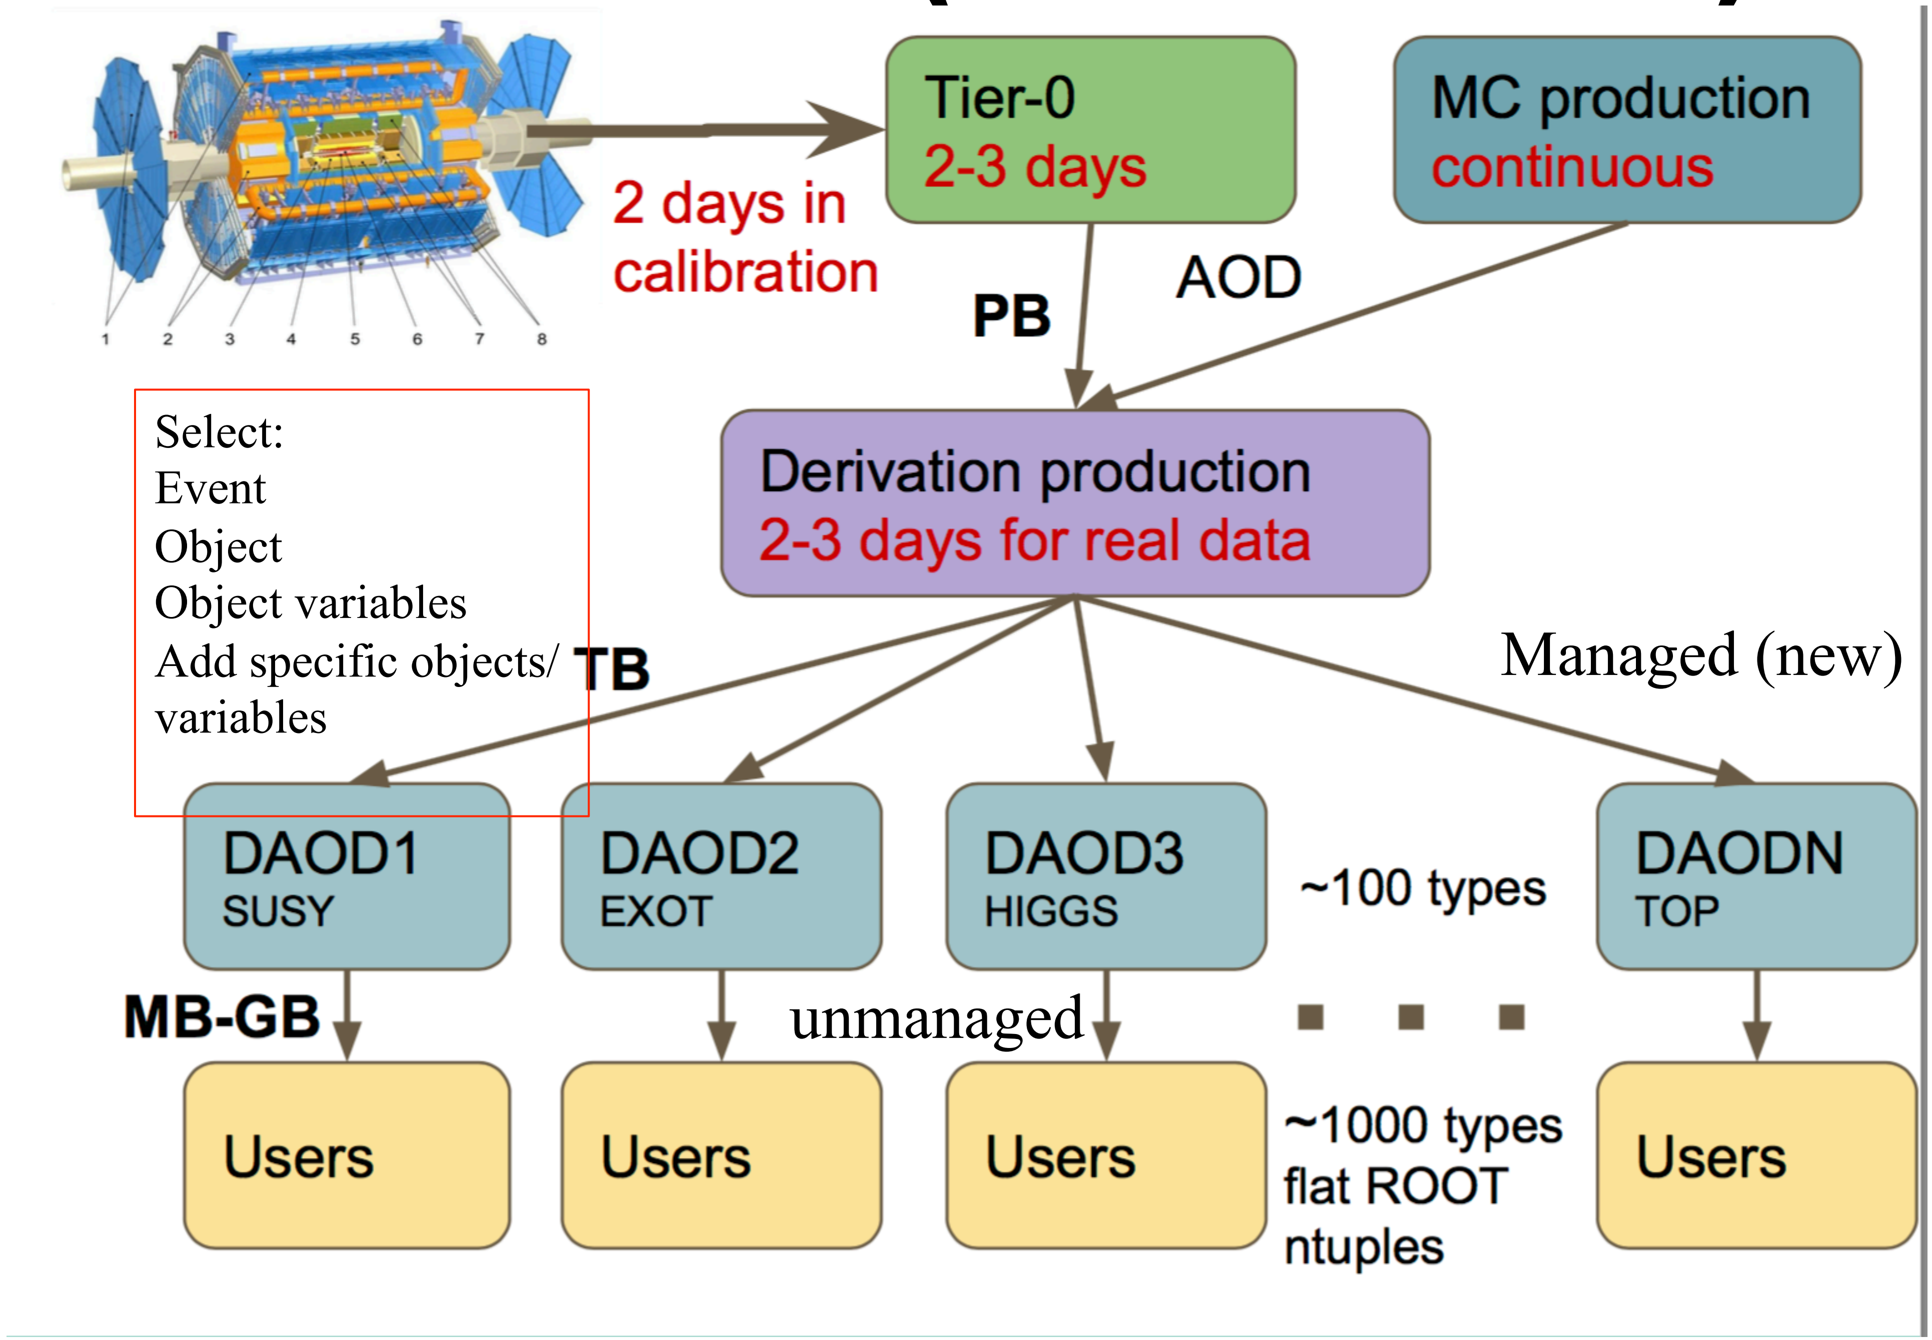

A 1% analysis

- Need improvement of reconstruction algorithm → will take years

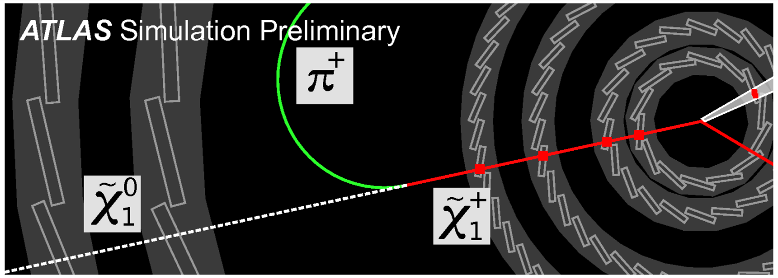

- Very common (in MSSM) SUSY scenario:
  - A **chargino nearly degenerate with a neutralino** (wino-like LSP).
  - The chargino **becomes long-lived** (typical τ = 0.2 ns or cτ~6 cm).
  - Effort to increase to **increase sensitivity at low lifetime.**
  - IBL **key element** for this search.

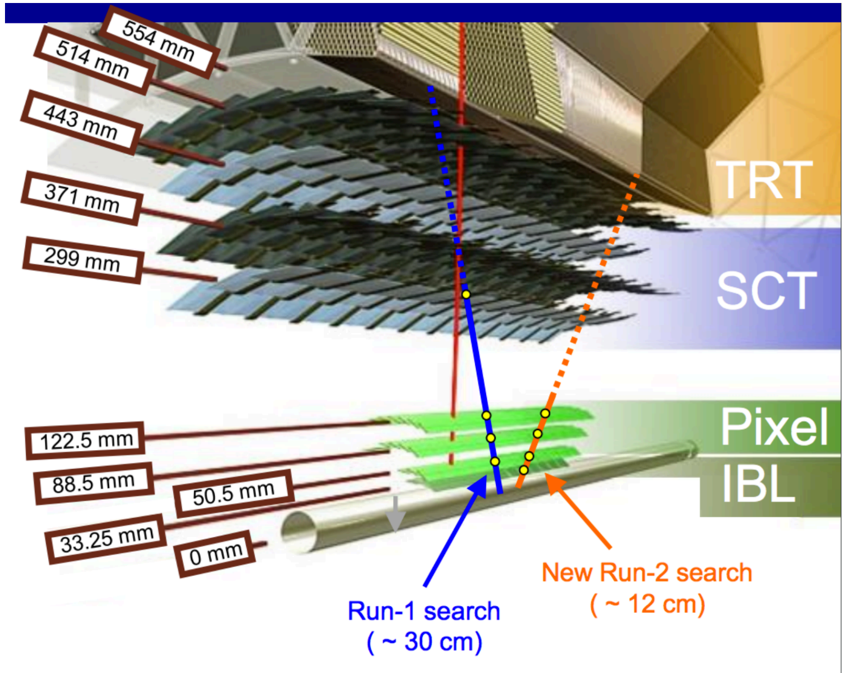

Things we (still) do badly

- Release notes:
  - We never manage to have a reasonably detailed english description of what is changed from one release to the next
  - Rely on expert memory (and detailed list of package changes)
- Final Analysis model:
  - Hundredths of different analysis
  - Each tend to be focussed on their own development
  - We did not manage to enforce common framework
  - Very difficult to enforce sharing of resource
  - Excessive time spent on "acceptance challenge"
    - Good to have 2 or 3 persons to do the same things to make sure there is no mistakes, but 10 ?

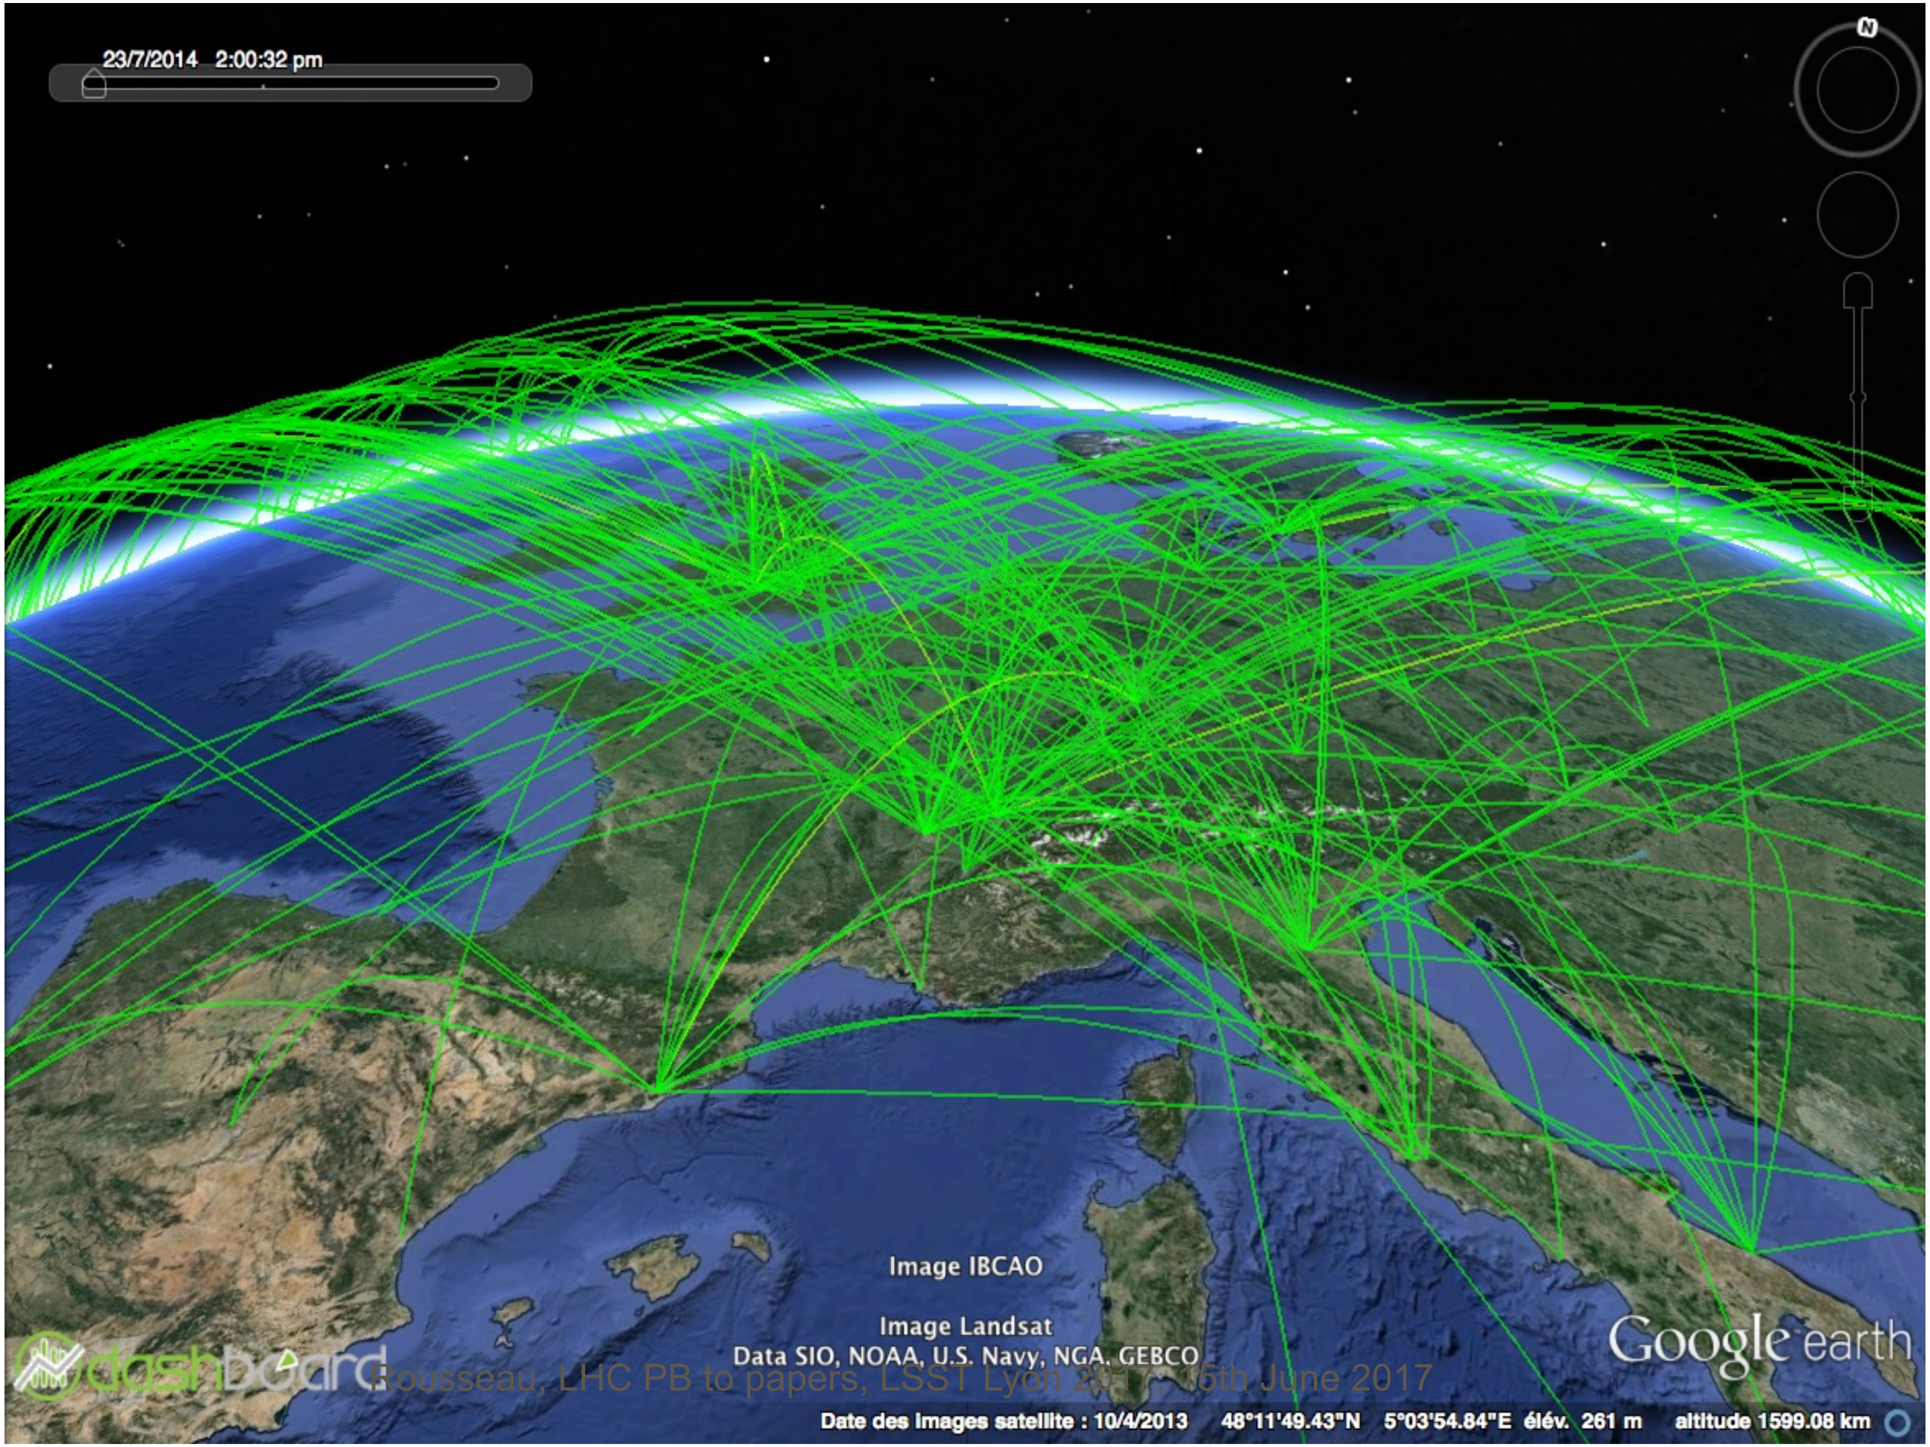
23/7/2014 2:00:32 pm
N
Image IBCAO
Image Landsat
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Google earth
dashboard
Date des images satellite : 10/4/2013 48°11'49.43"N 5°03'54.84"E élév. 261 m altitude 1599.08 km

Data processing key numbers

- Billions events, ~PetaBytes raw data per year
- First reconstruction at CERN, then distributed world-wide
- Simulation in outside centers
- ~PetaBytes derived data per year, billions events and PetaBytes
- ➔total 100PB managed (several years, several copy)
- 300.000 cores running continuously, half for simulation only
- 4 millions line of code (C++, and python) written by 1000 people over 10 years, ~150 active

On going GRID evolution

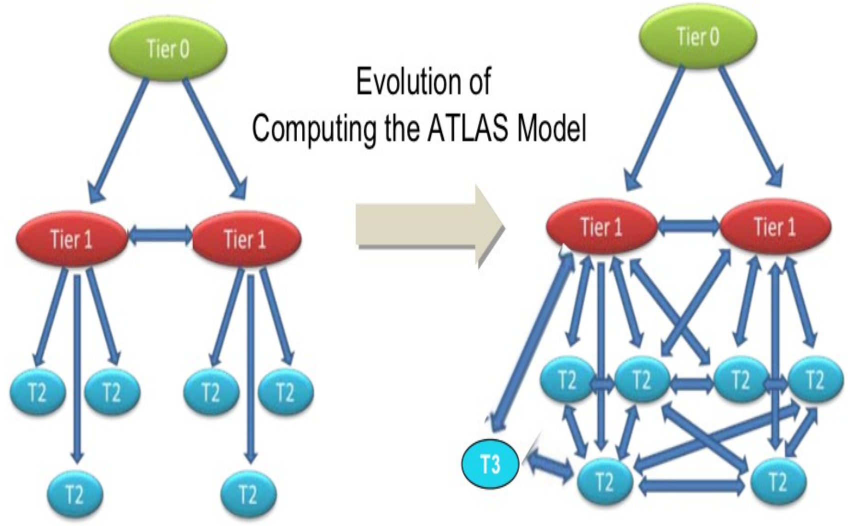

Before: Hierarchical dataflow
Now : fully connected

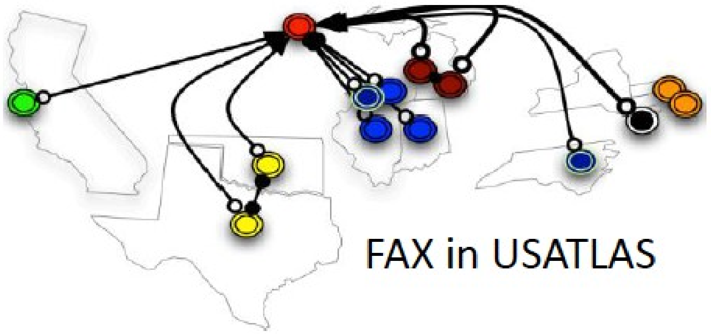

Before : jobs goes to data
Now: data goes to job
(as failsafe for now)

MB/s

ser

Future challenges

Processor evolution

Design of LHC experiments software

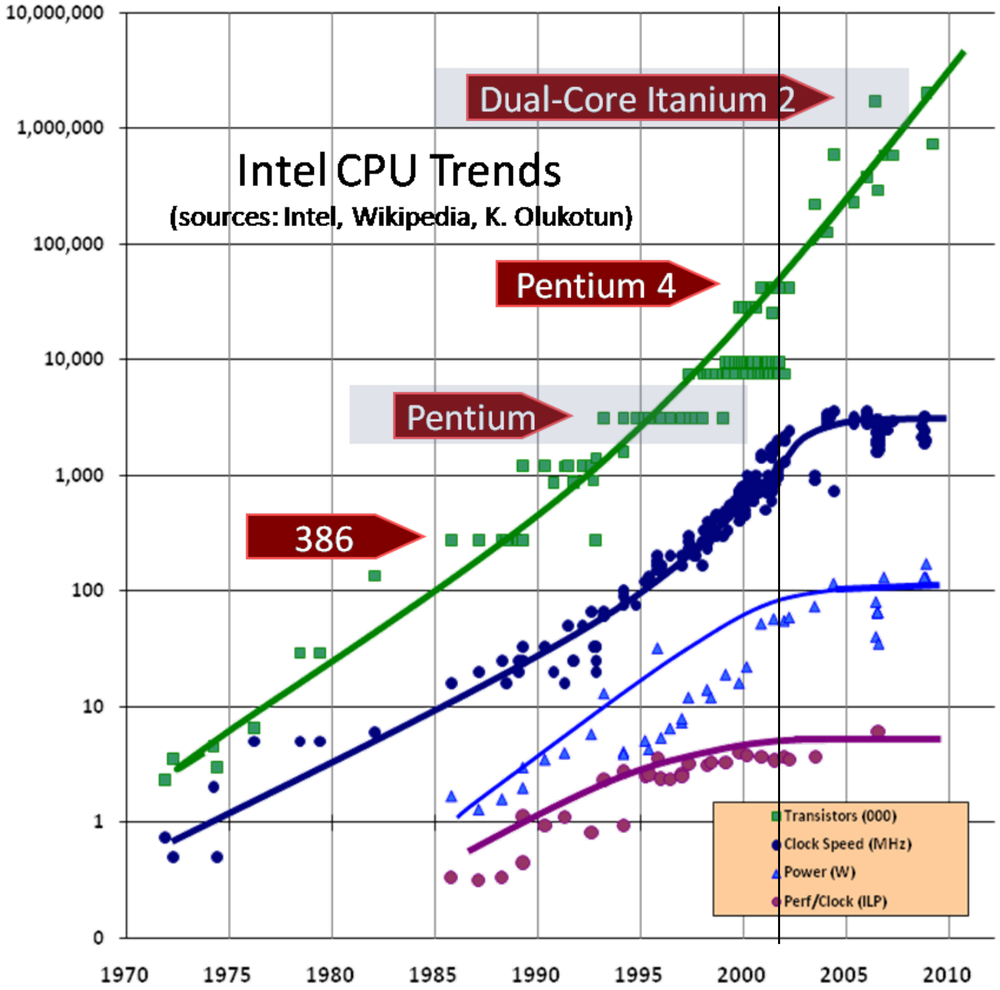

HL-LHC Context

- High-Luminosity LHC events in 2025:
  - more complex (pile-up ~150 rather than ~25 in run 1, also energy x ~2)
    - No reliable estimate today of the impact on CPU, as existing code shows non linear divergence. Indicatively, multiplicity increases by a factor 8.
  - higher read-out rates (factor ~10)
- Flat resources (in euros) and Moore's law give us a factor 10 in CPU power (if and only if we can use the processors as efficiently as today!)
- ➔handling HL-LHC event added complexity, and maintenance/improvement of processor efficiency rely on software improvements. If not, impact on physics.
- ➔we need a factor 10 improvement (CPU and data storage) from better software/organisation

Servers (2)

Preliminary extrapolation of CPU and disk server costs (based on CERN procurements)

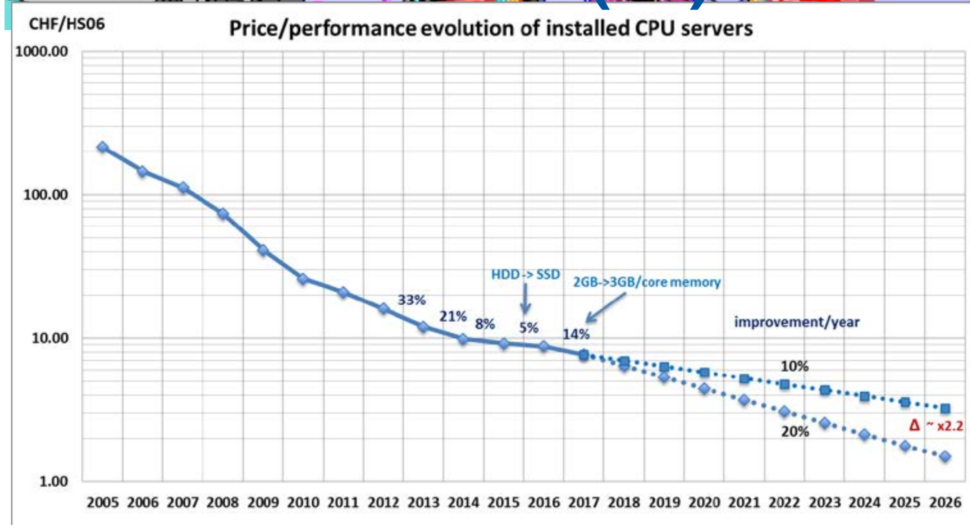

Pessimistic and reasonable improvement extrapolations

Influence of changing software and hardware architecture requirements to be taken into account (programs, data model, data centre,)

- Moore’s Law and Kryder’s Law are slowing down
  - 18 months → >= 3 years
- Real cost/performance evolution driven by financial and market aspects rather than technology

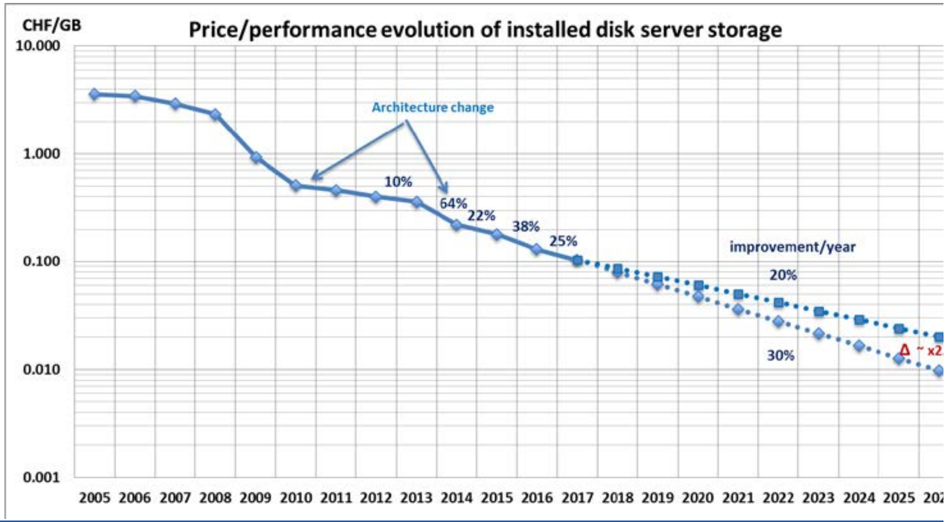

HL-LHC baseline resource needs

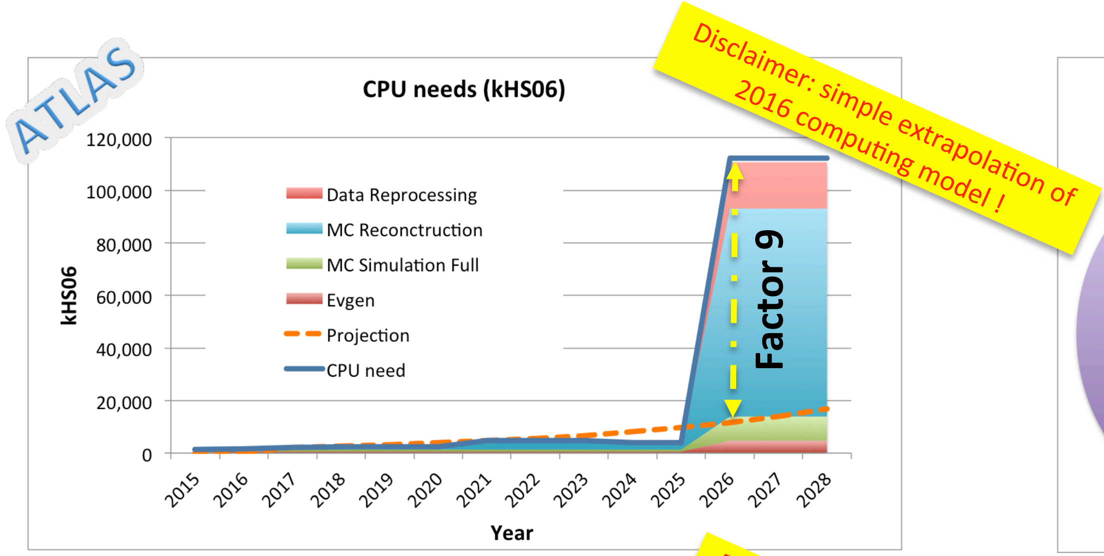

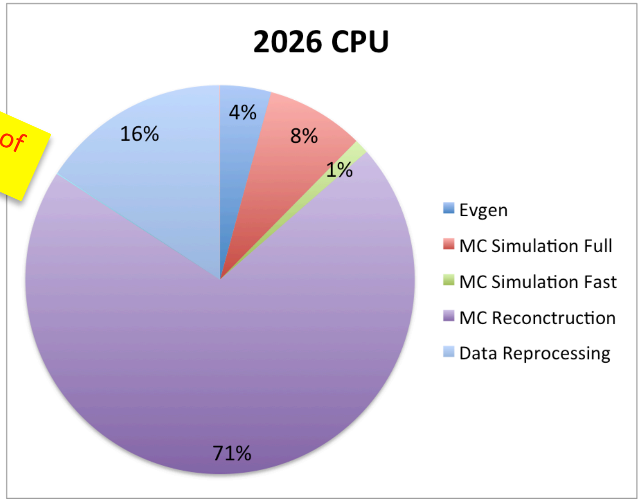

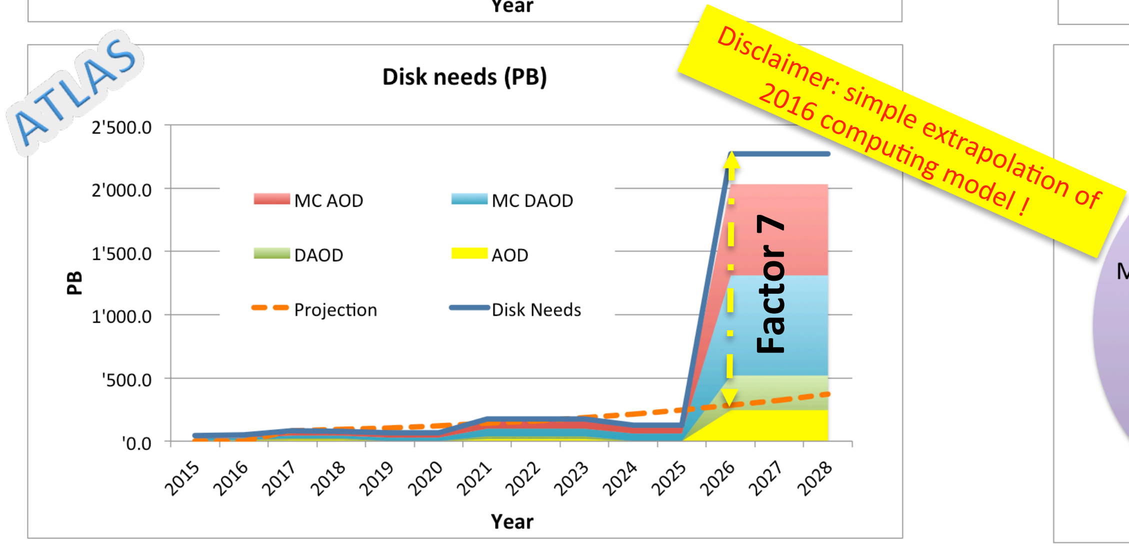

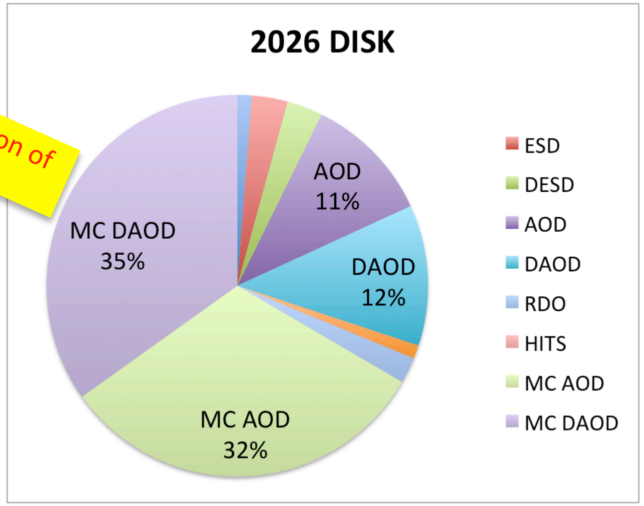

Major crisis ?

Note on data organisation

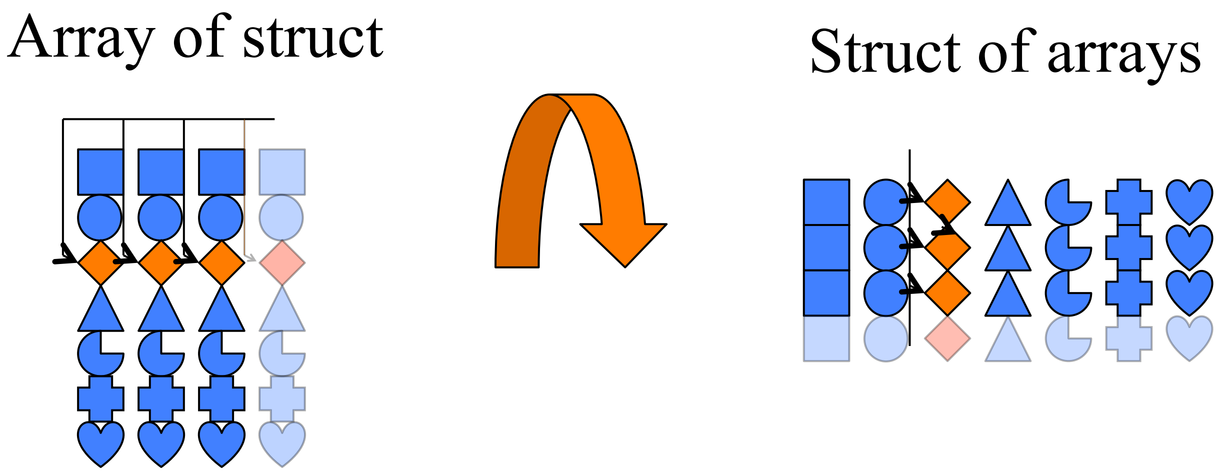

➔ More suitable for vectorisation

- Also object type known at compile time ➔better compiler optimisation➔need to simplify inheritance structure
- Data Object need to be completely revised

One core one job

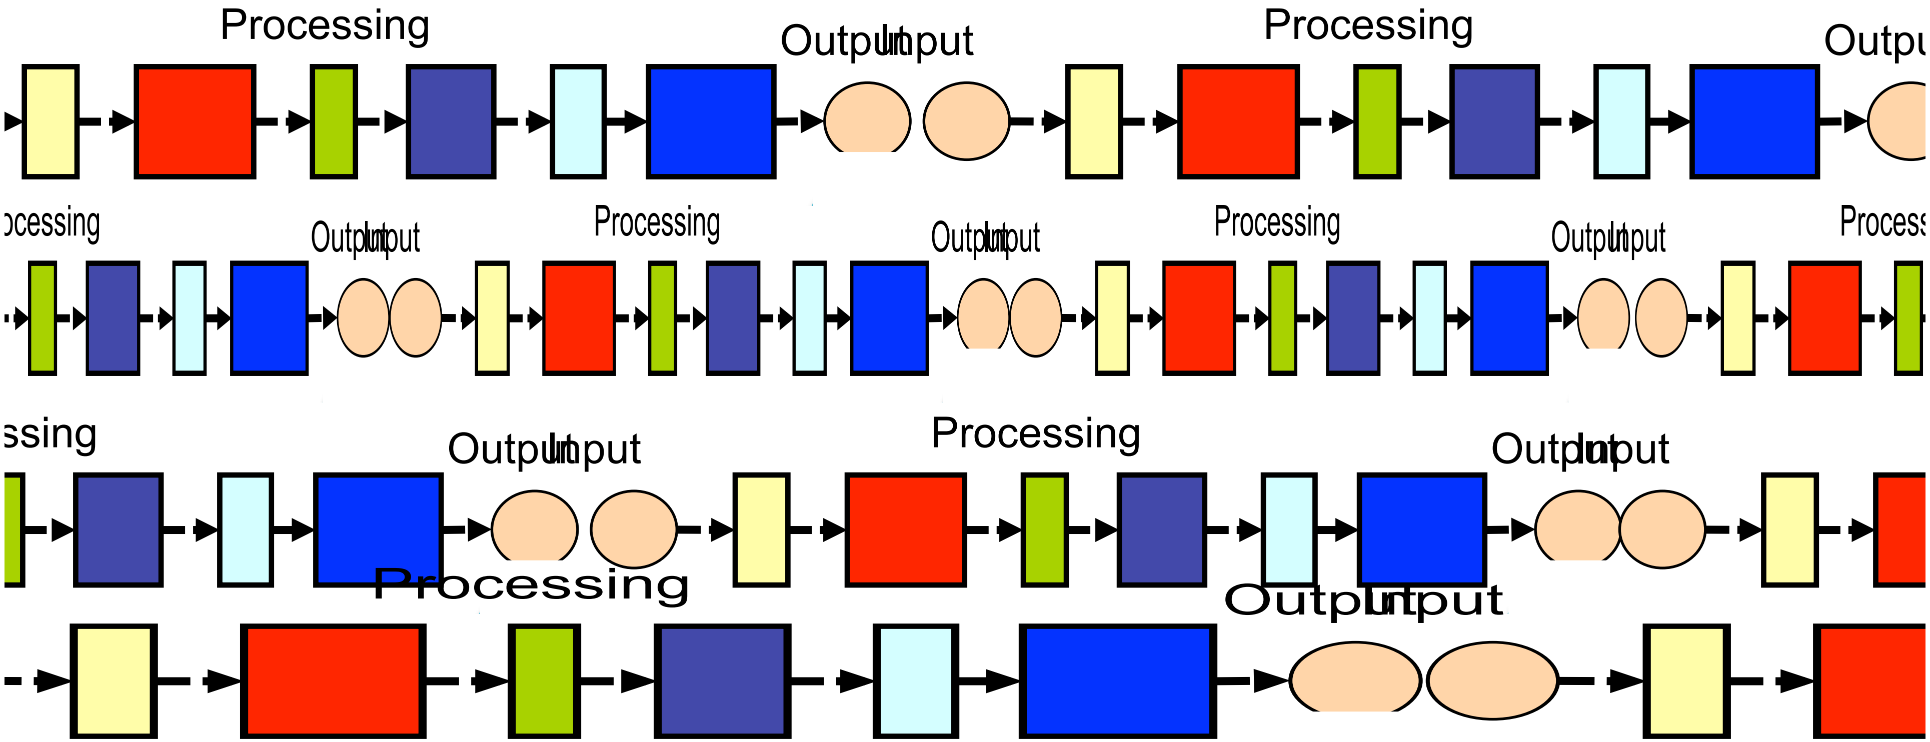

- Today, typical grid workhorse is a 16GB memory, 8 core CPU (2GB/core)
- Each core is adressed by the batch system as a separate processor
- Each job process event one by one, running one by one a finite number of algorithms
- One processor may handle simultaneously e.g. one Atlas reco job, 3 CMS simulation job, and 4 LHCb analysis jobs
- This works (today), however disorganised competition for resources like memory, I/O (Already memory issues for Atlas Run 2 reconstruction, requires 3GB memory per core.)

One processor one job

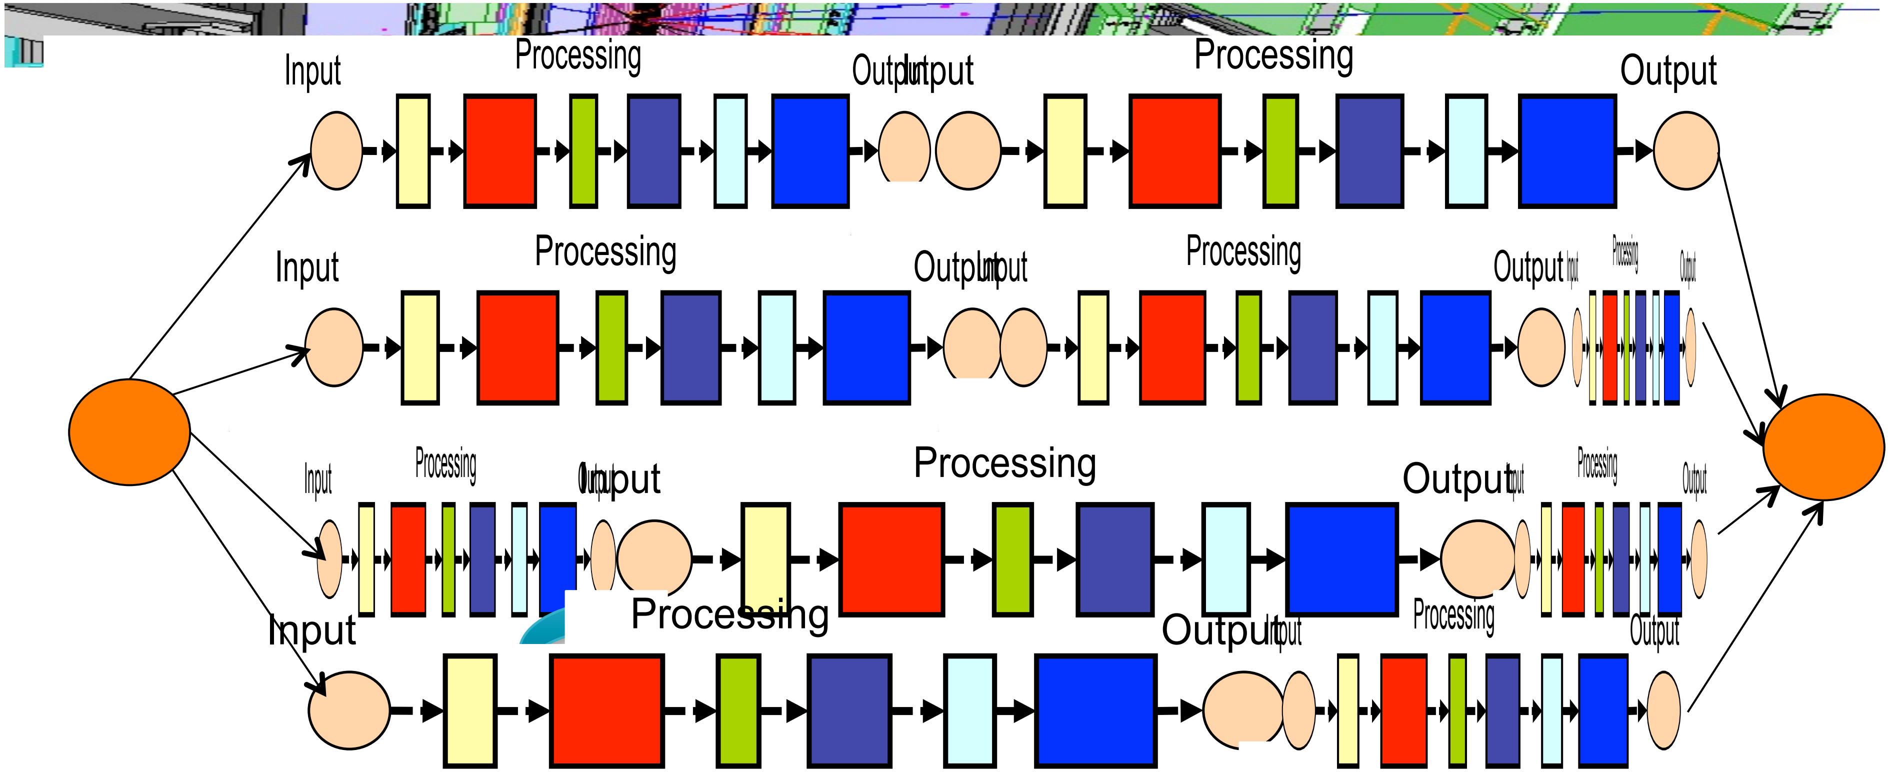

- In production since last year (GaudiMP, AthenaMP)
- One job goes to one processor (which is completely free)
- The framework distributes event processing to all cores, while sharing common memory (code, conditions,…) using Copy-on-Write
- No change to algorithmic code required (in principle, took 5 years of debugging though…)
- ˜50% reduction of memory achieved (w.r.t. independent jobs)

Event level parallelism

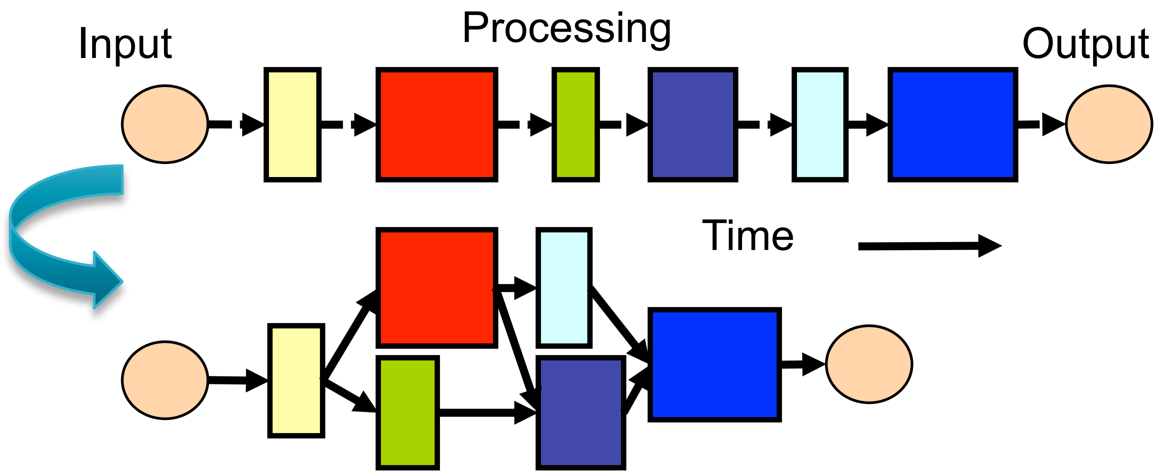

- framework schedules intelligently the algorithms from their dependency graph
- e.g. run tracking in parallel with calorimeter, then electron ID
- in practice too few algorithms can run in parallel
- ➔most cores remain idle

Real life

- Direct Acyclic Graph extracted from real reconstruction job
- Today, algorithms run sequentially

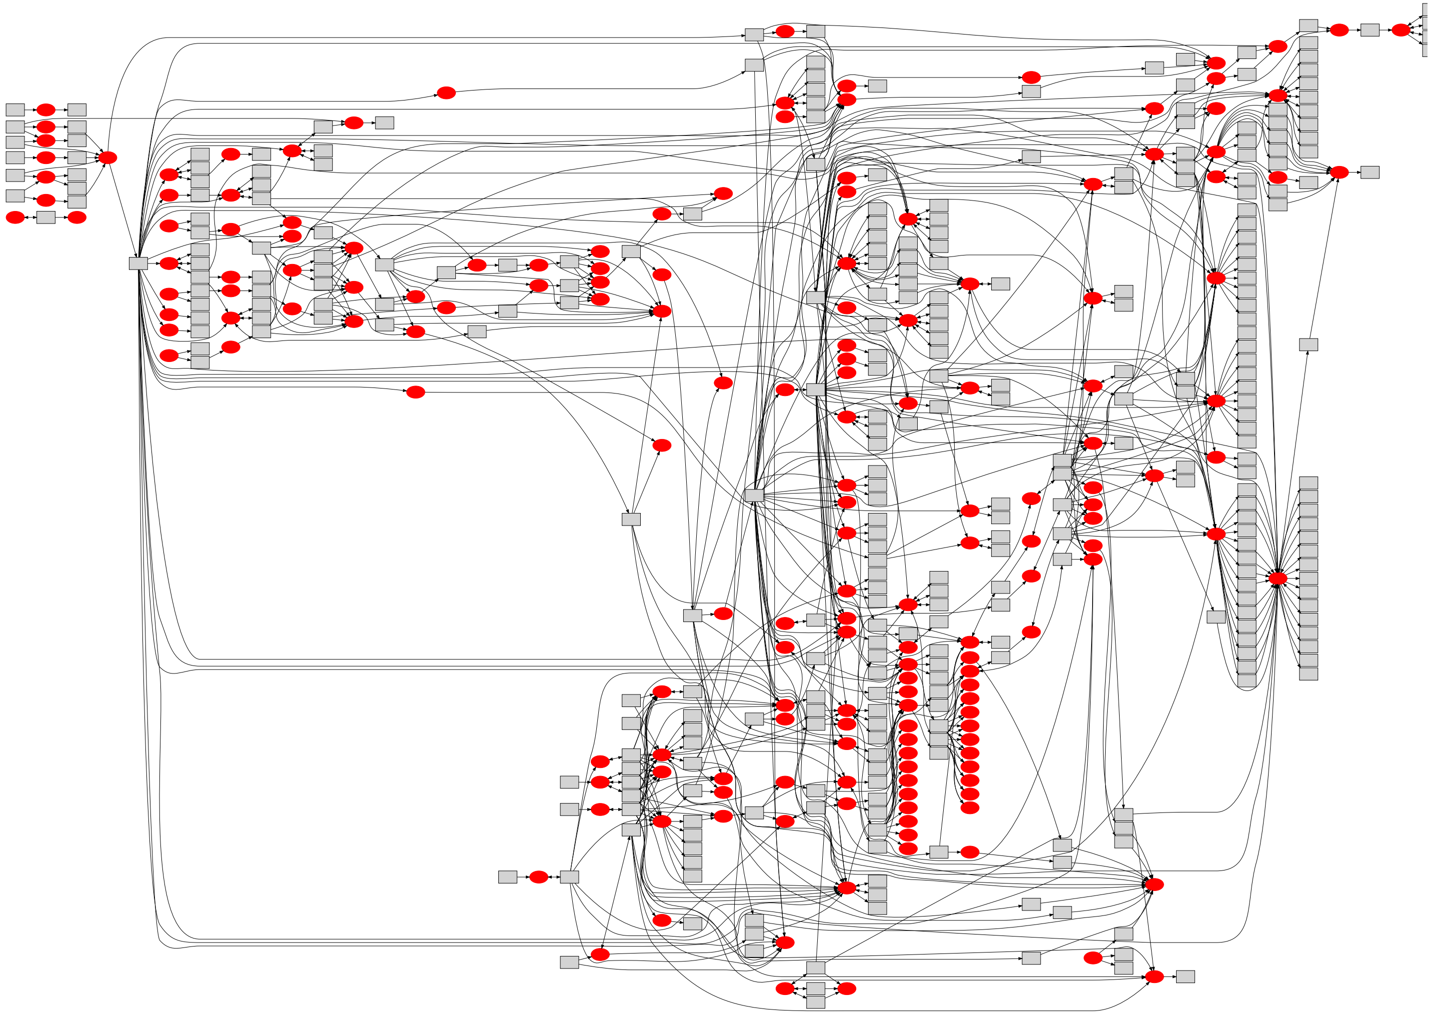

Event level concurrent event processing

a.k.a the Holy Grail

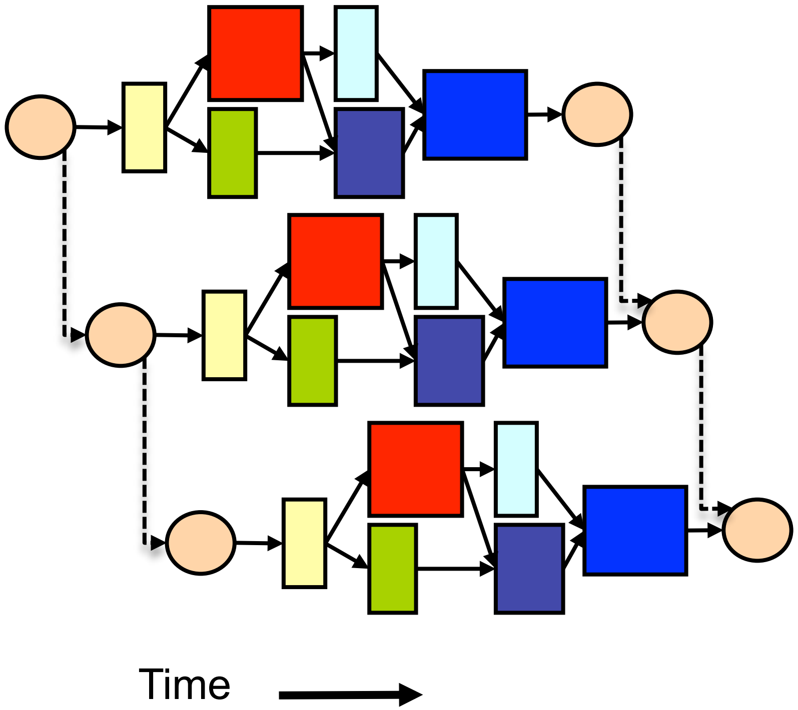

- The framework process several events simultaneously…
- …distributes intelligently algorithms to cores
- can allocate more cores to slowest algorithms
- can optimise use of specialised cores

- In addition to algorithm scheduling, the framework provides services to pipeline access to resources (I/O, conditions, message logging…)
- Algorithms should be thread safe : no global object (except through the framework), only use thread safe services and libraries
- Algorithms do not need to handle threads themselves
- ➔regular software physicist with proper training can (re)write algorithms

Machine Learning

One mention and one example

Tracking with NN

arXiv 1604.01444 Aurisano et al

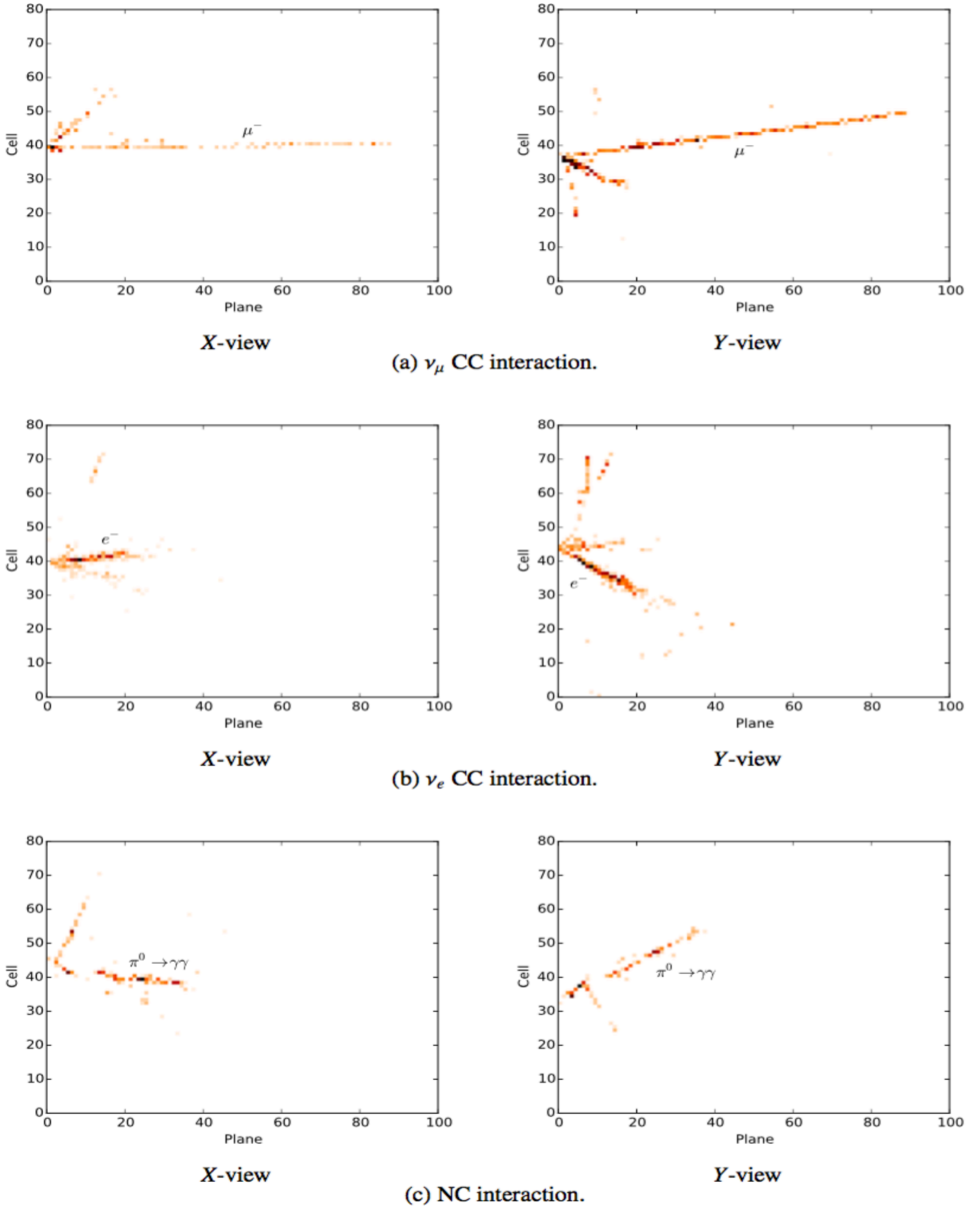

(a) ν_μ CC interaction.

(b) ν_e CC interaction.

(c) NC interaction.

Neutrino interaction classification
Using Convolutional Neural Network (GoogleNet)

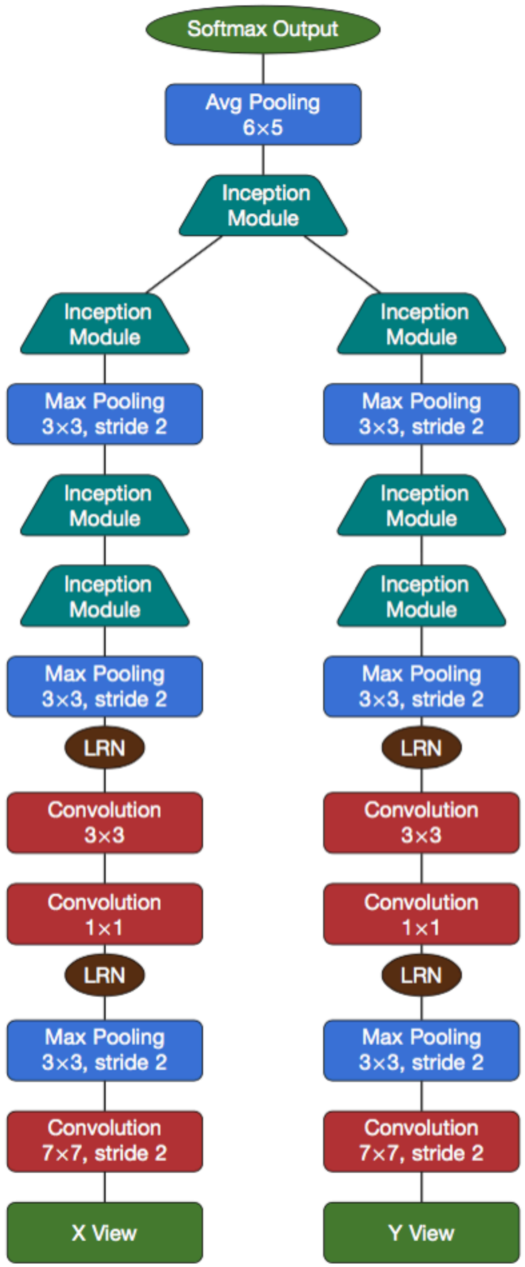

Generative Adversarial Network

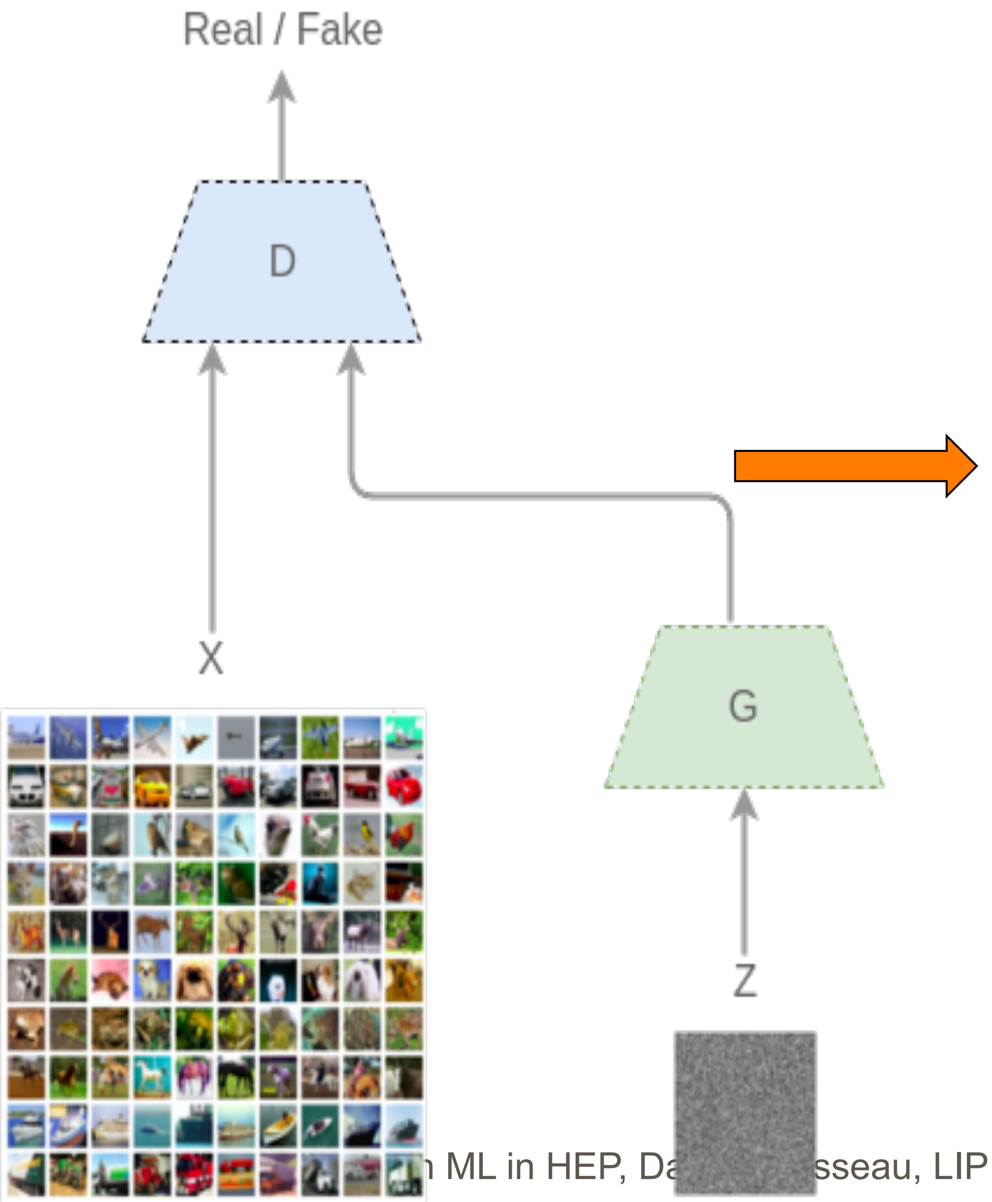

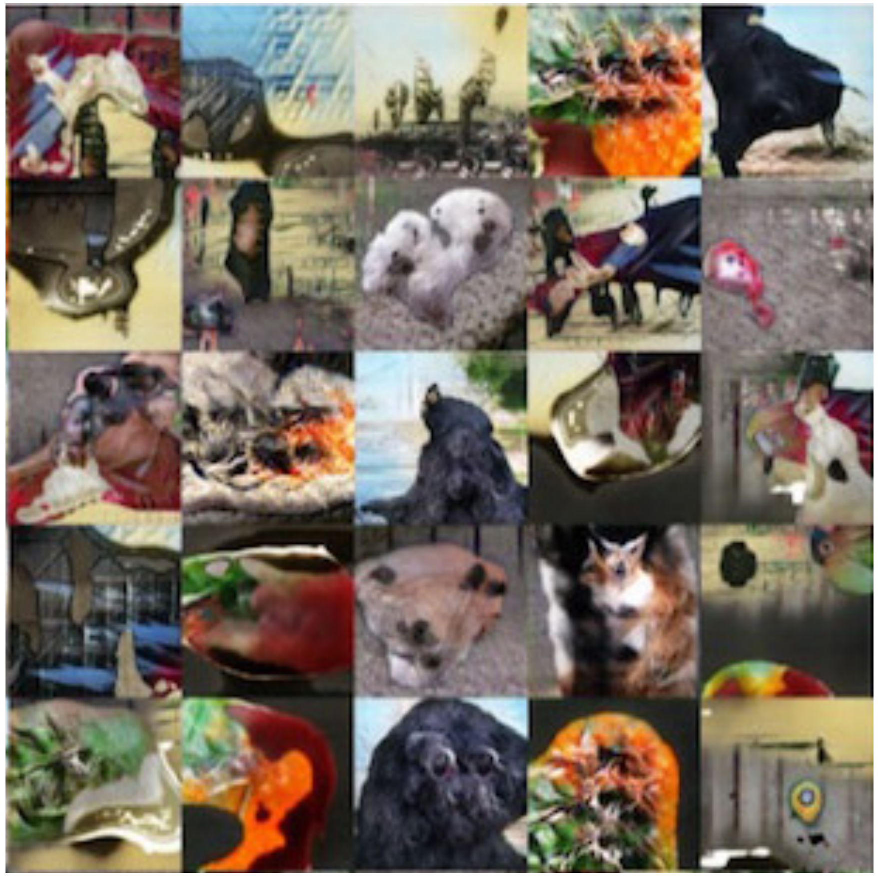

Condition GAN

Text to image

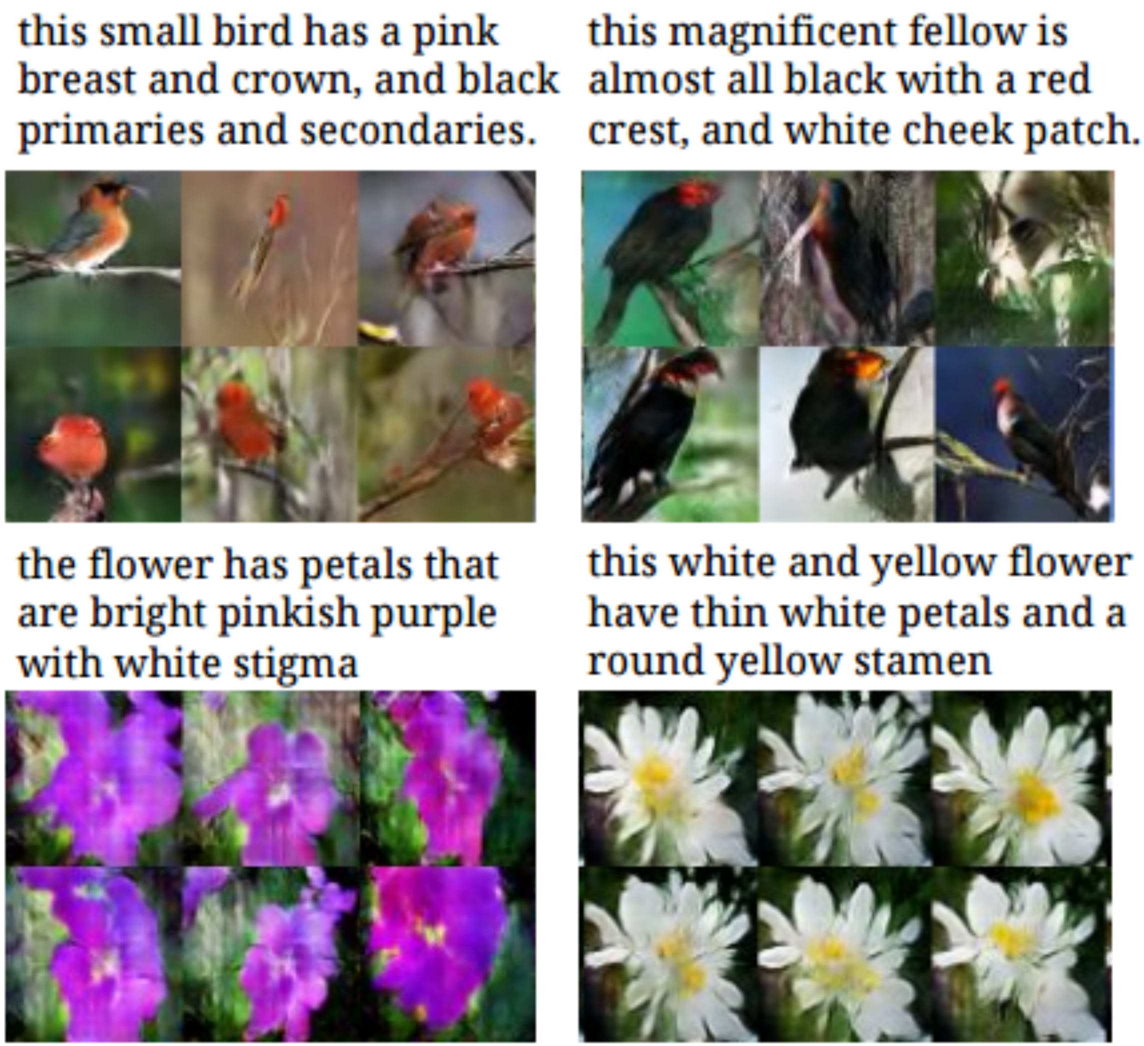

GAN for simulation

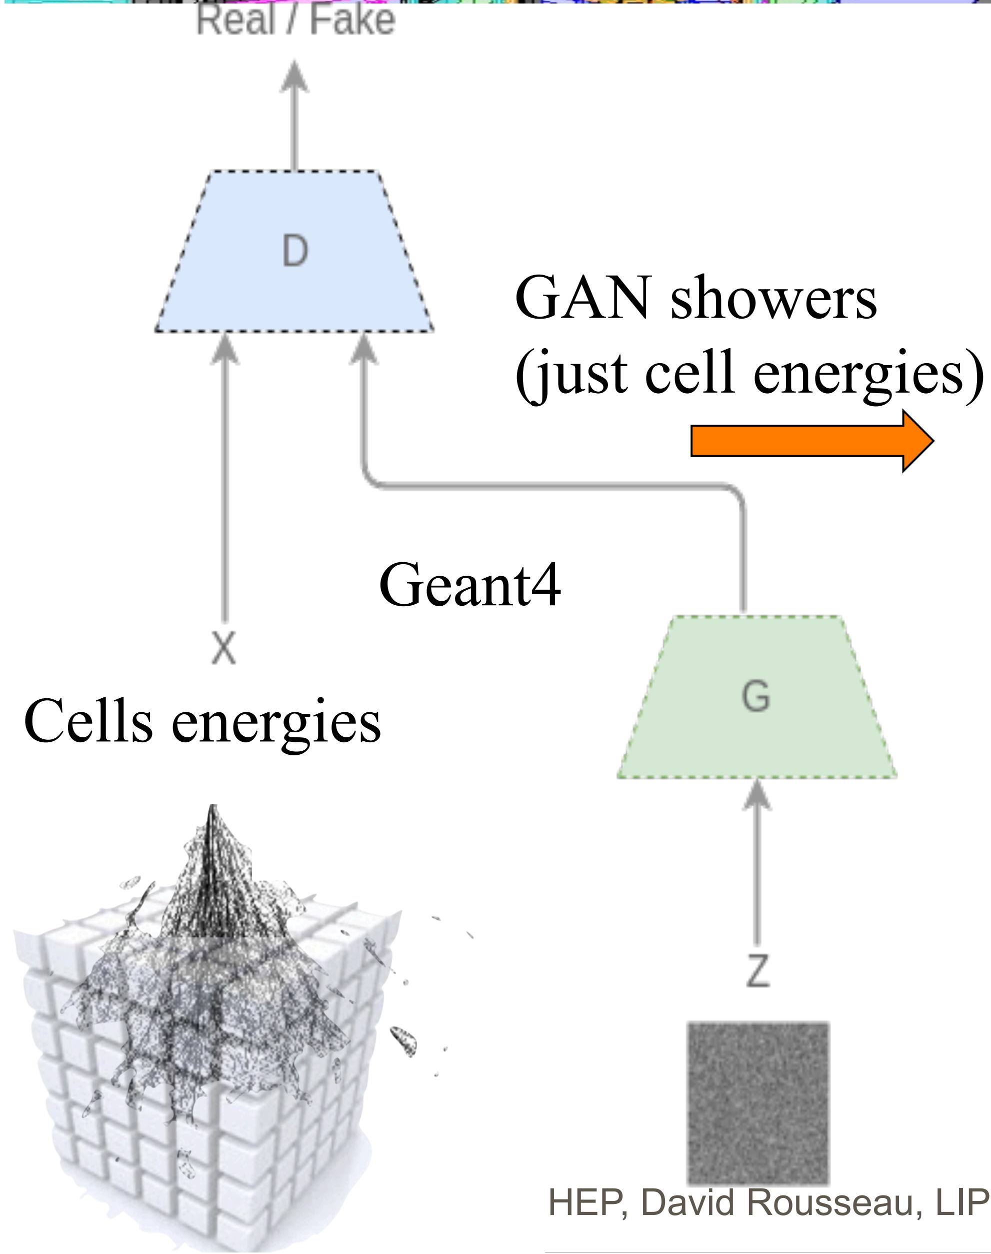

- Half of LHC grid computers (~300.000 cores) are crunching Geant4 simulation 24/24 365/365
- …while LHC experiments are collecting more and more events
- ➔reducing CPU consumption of simulation is very important

- Imagine training a GAN on single particle showers of all types and energies
- Then when an event is simulated it would ask for GAN showers on request (superfast)
- Would replace current fast simulation, frozen shower libraries….

Proof of concept e.m. shower

Avoid edge effects
(0,1 < |η| < 1,1)

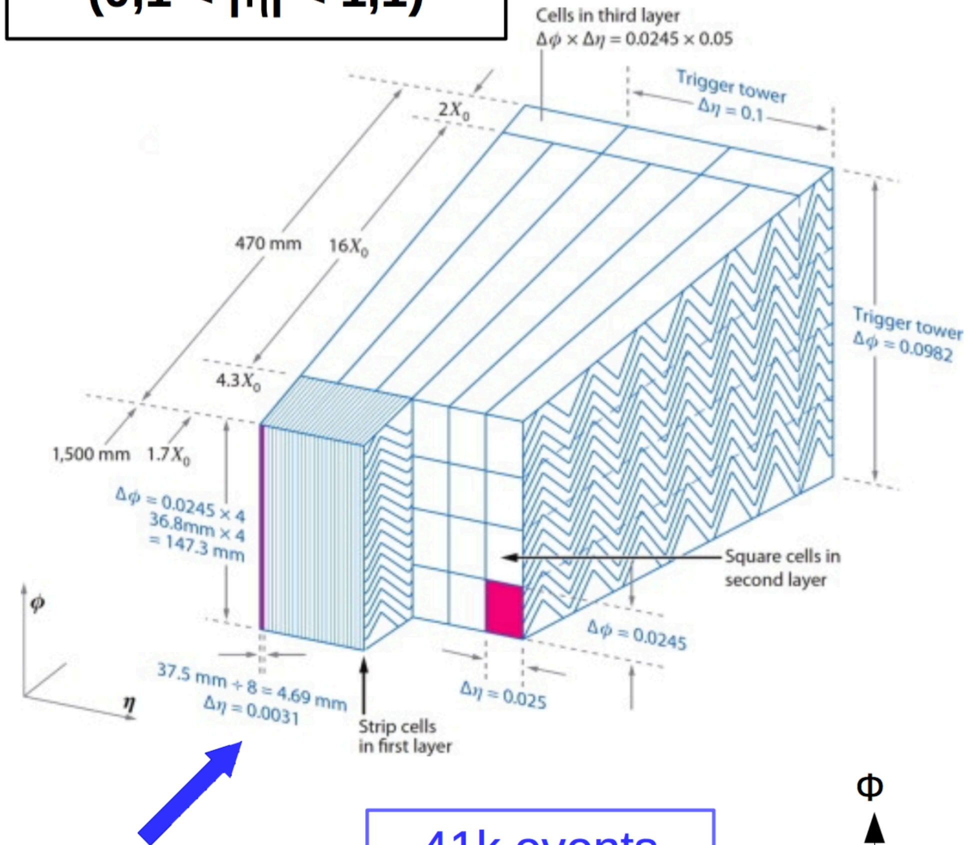

Single Photon

41k events
(<300GeV)

CLUSTER 3*7
(89 cells)

→ centered on the most energetic cell (Middle)

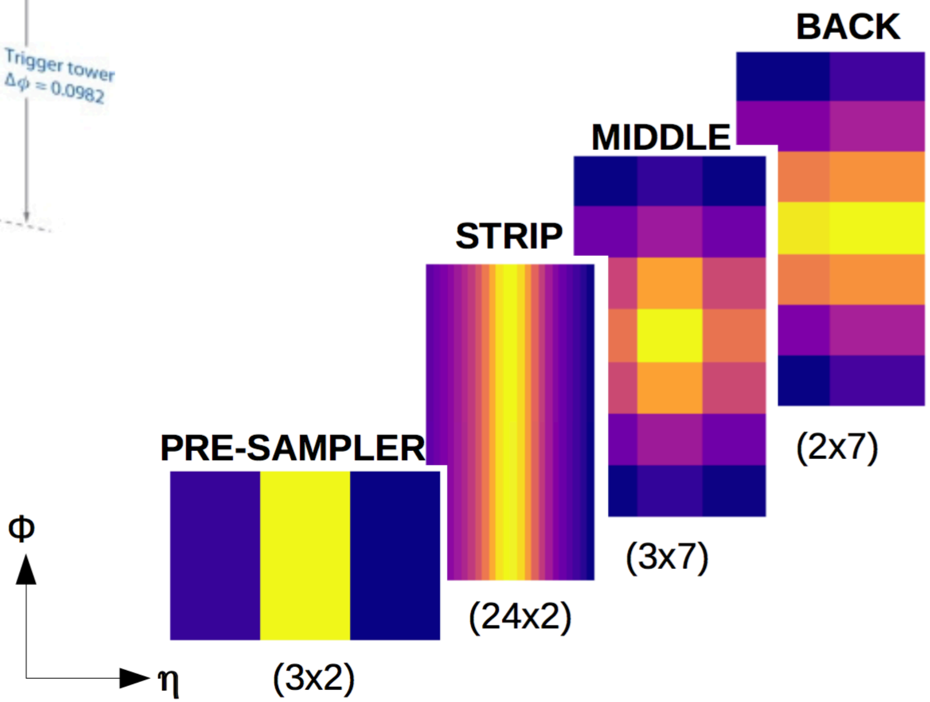

GAN architecture

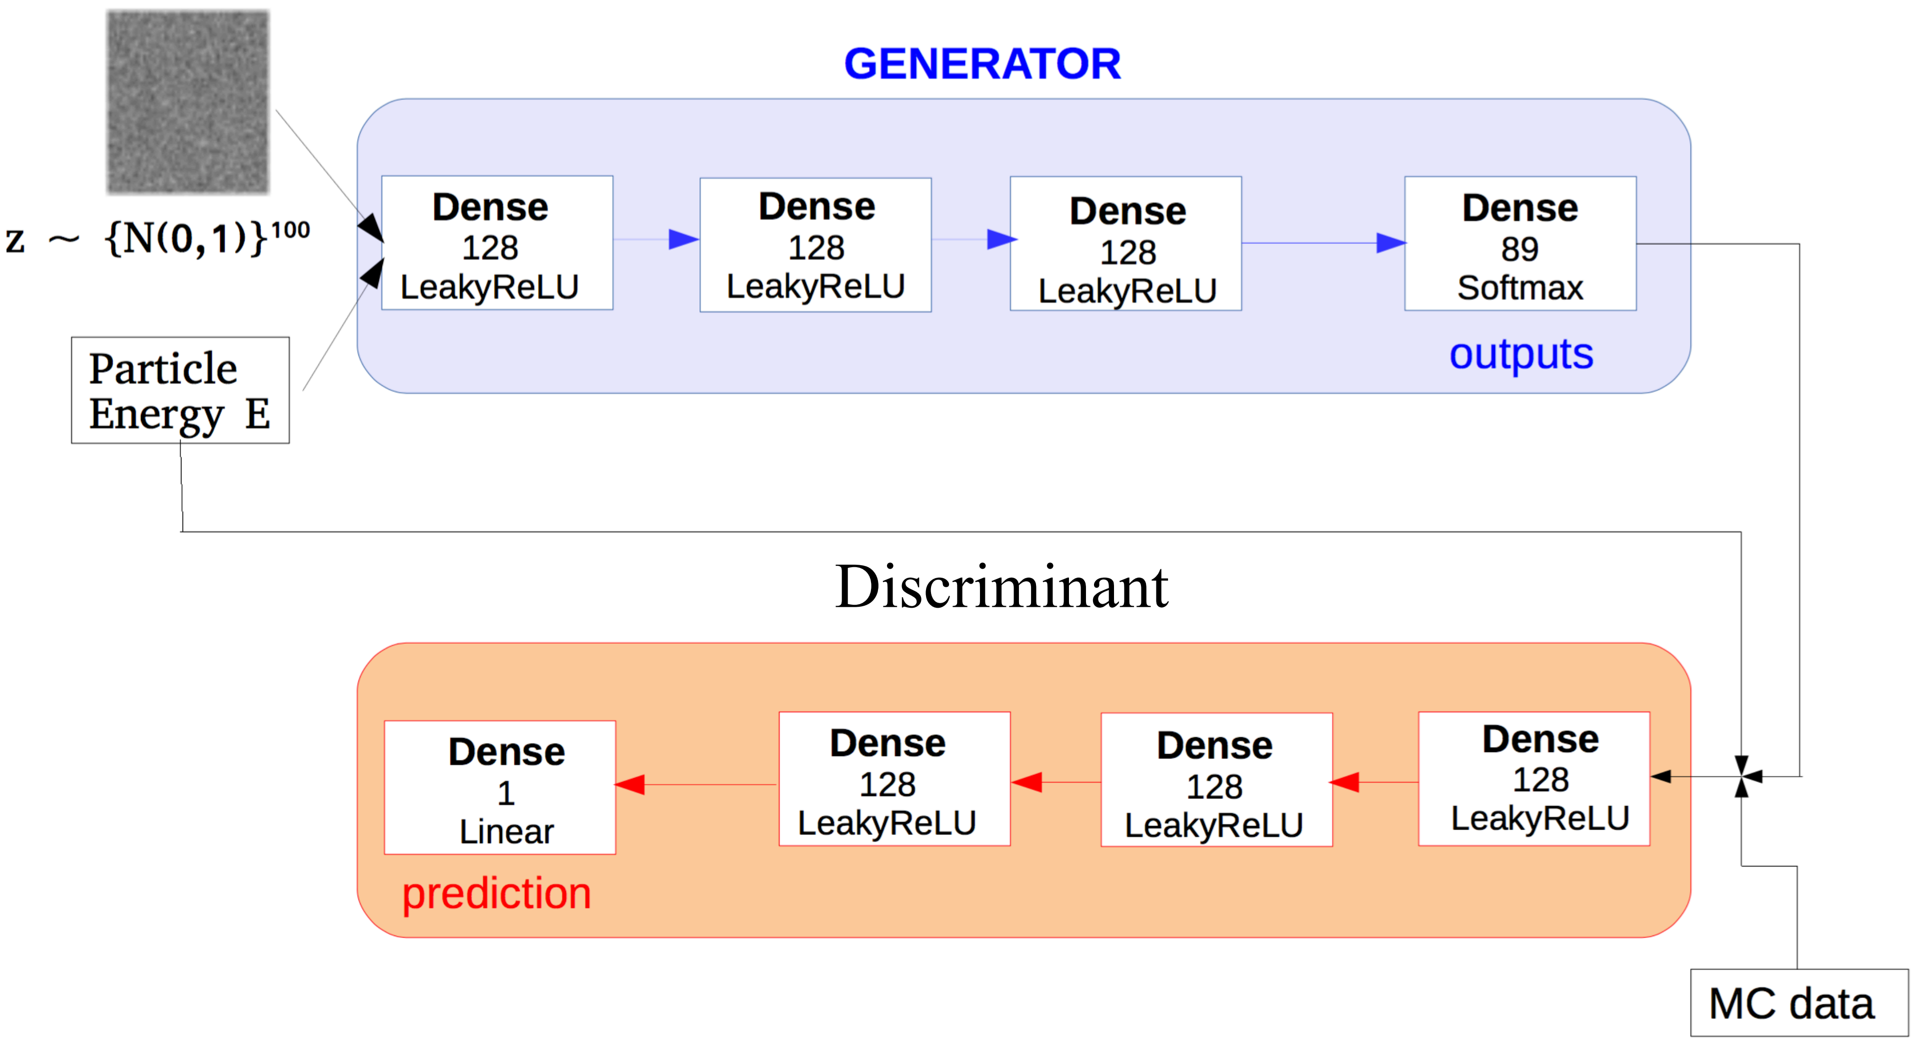

Average cell distribution

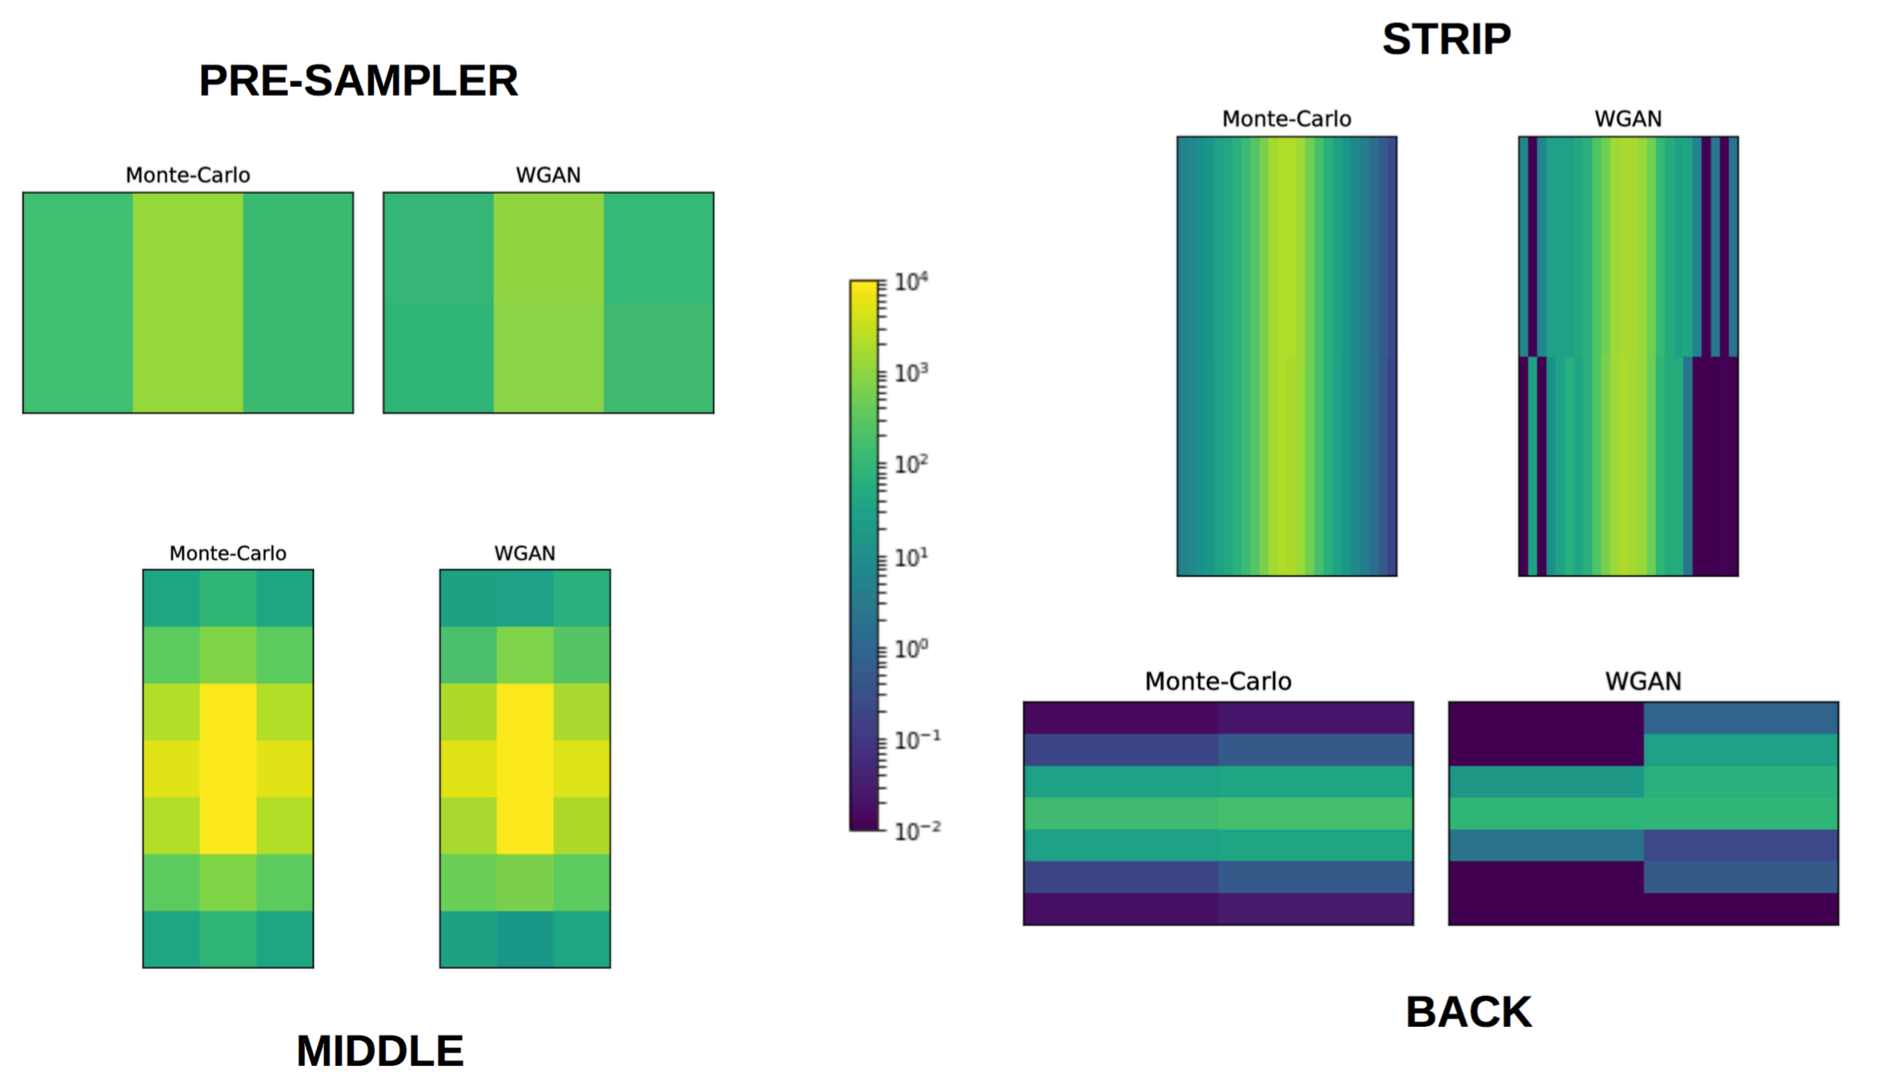

Some distributions

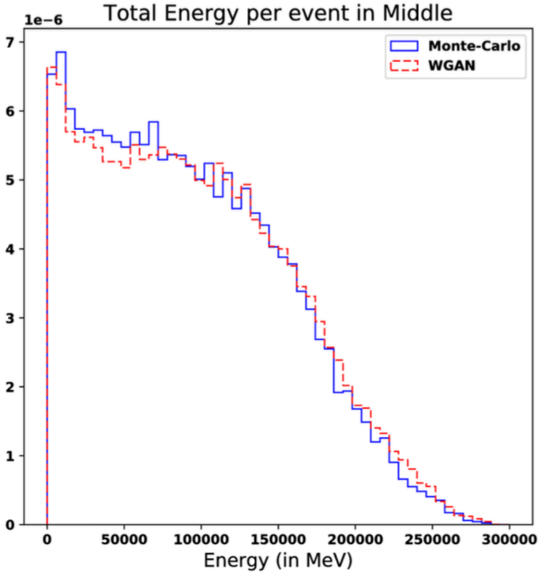

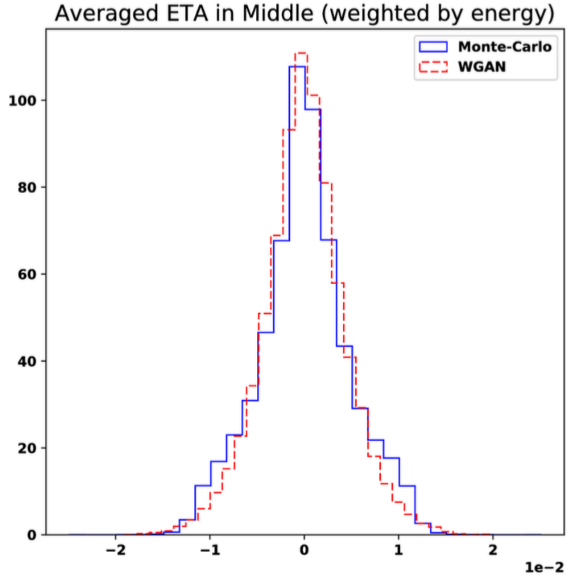

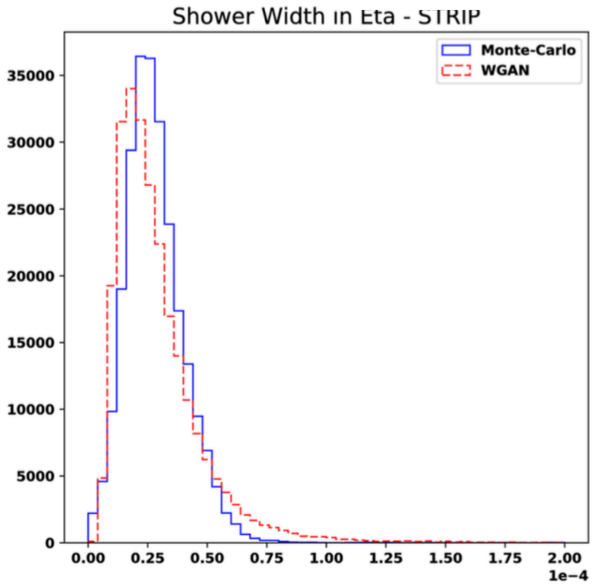

Not bad,
just the beginning

Note : GeantV similar project starting

Summary

- Technical challenge
- Organisation challenge
- Delicate balance
  - Accuracy vs stability
  - Fast analysis turn around vs access to full information
- Change of paradigm on going:
  - Parallel processing
  - Machine Learning