HIGH PERFORMANCE COMPUTING SYSTEM IN THE FRAMEWORK OF THE HIGGS BOSON STUDIES

N. Belyaev, A. Klimentov, R. Konoplich, D. Krasnopevtsev, K. Prokofiev, V. Velikhov

on Behalf of the ATLAS Collaboration

OUTLINE OF THE TALK

After the discovery of the Higgs boson, it is important to measure its properties to verify the Standard Model and look for deviations that could present hints to new physics.

The High performance computing systems (HPC) now become more and more valuable for the successful and qualitative measurements.

This talk will cover the following topics:

- Introduction into LHC experiments
- HPC in the framework of the ATLAS experiment
- Computing: from simulation to reconstruction
- Supercomputers in the Higgs physics
- Performance of the GRID Tier 1 computing clusters
- Conclusions

LHC experiments

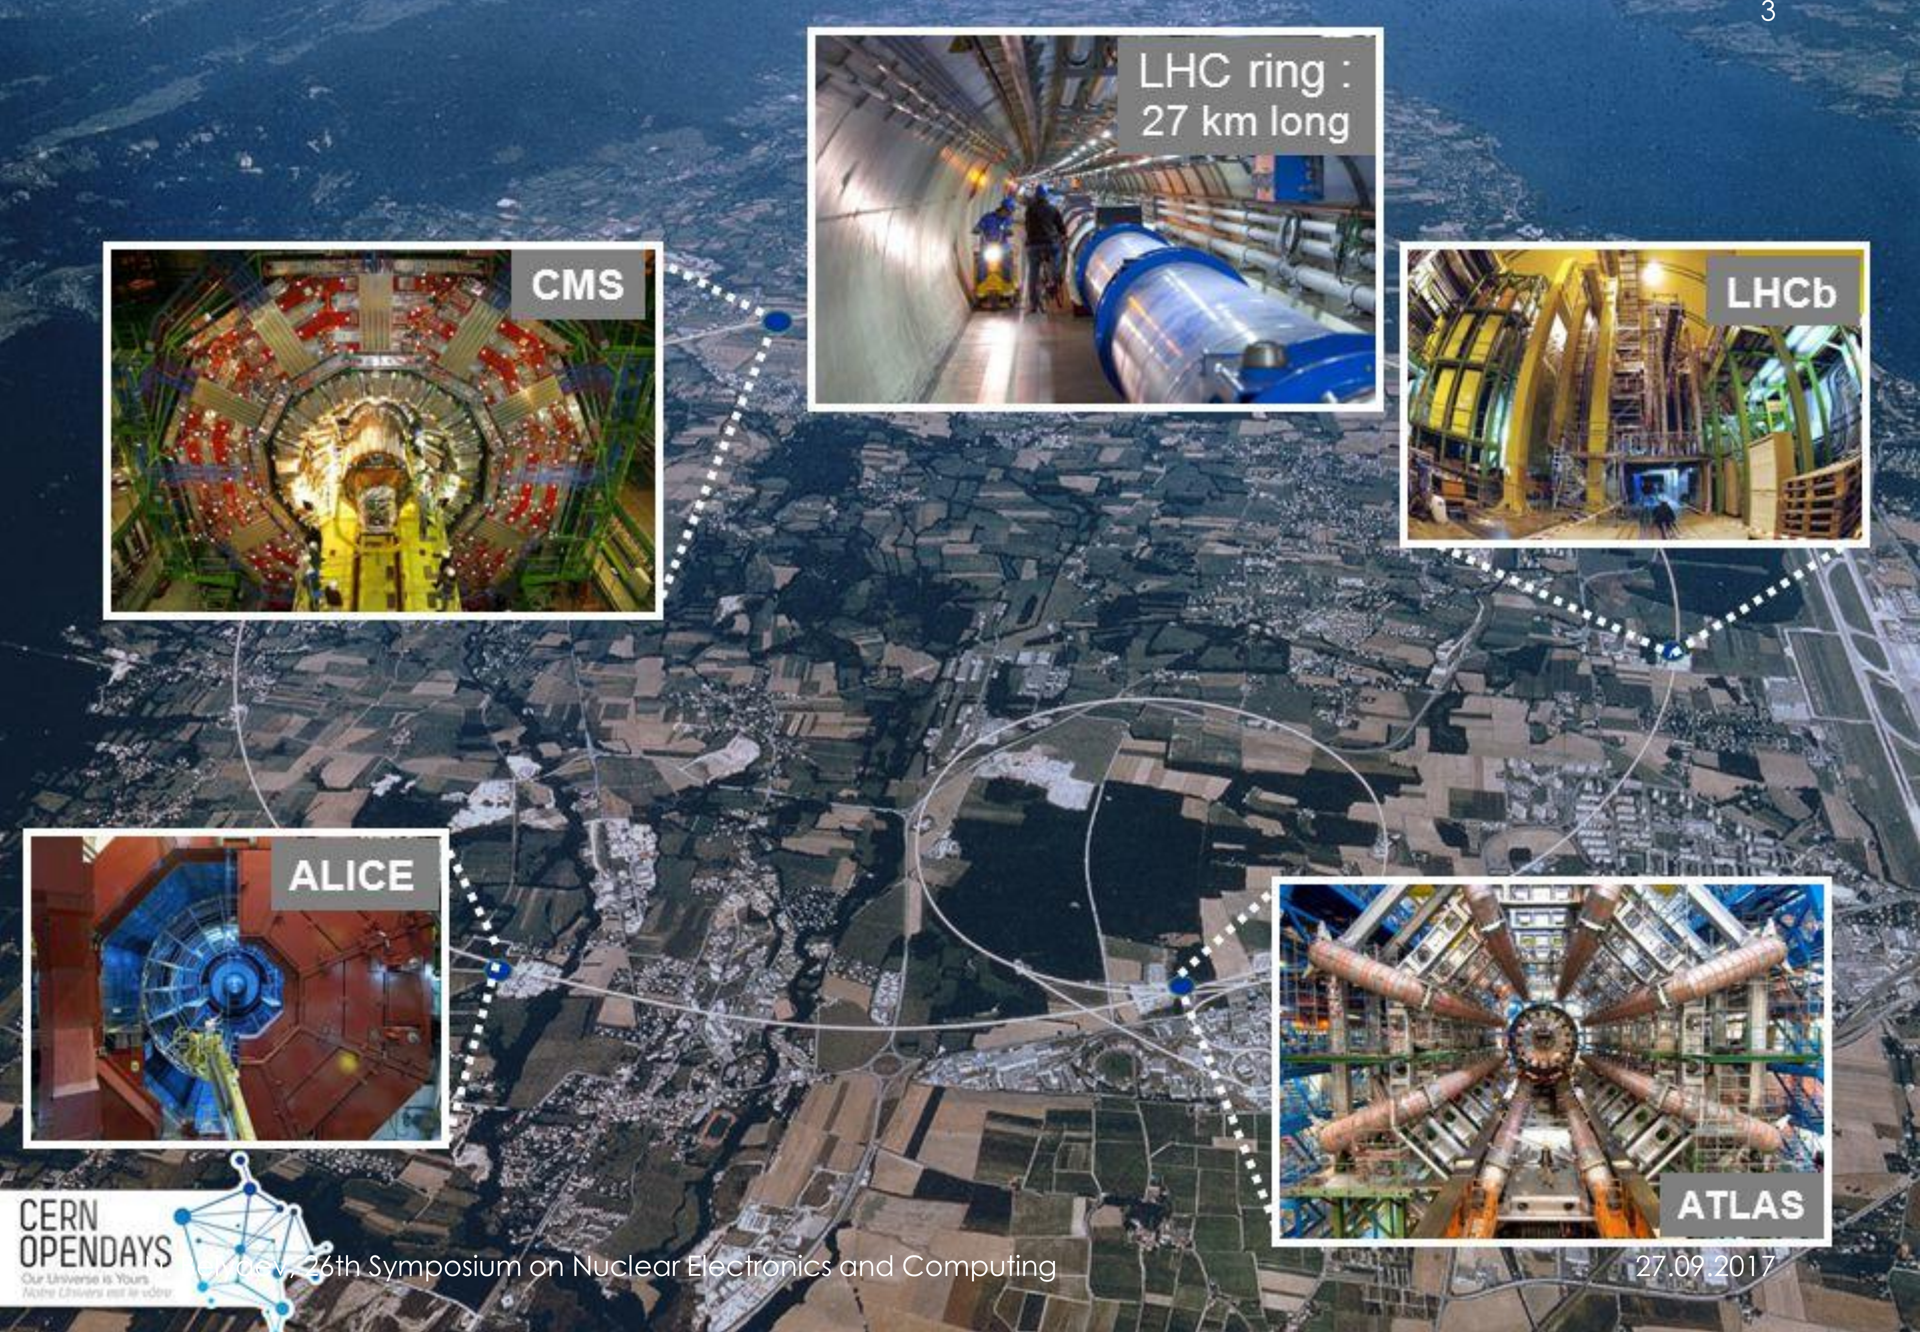

THE ATLAS DETECTOR

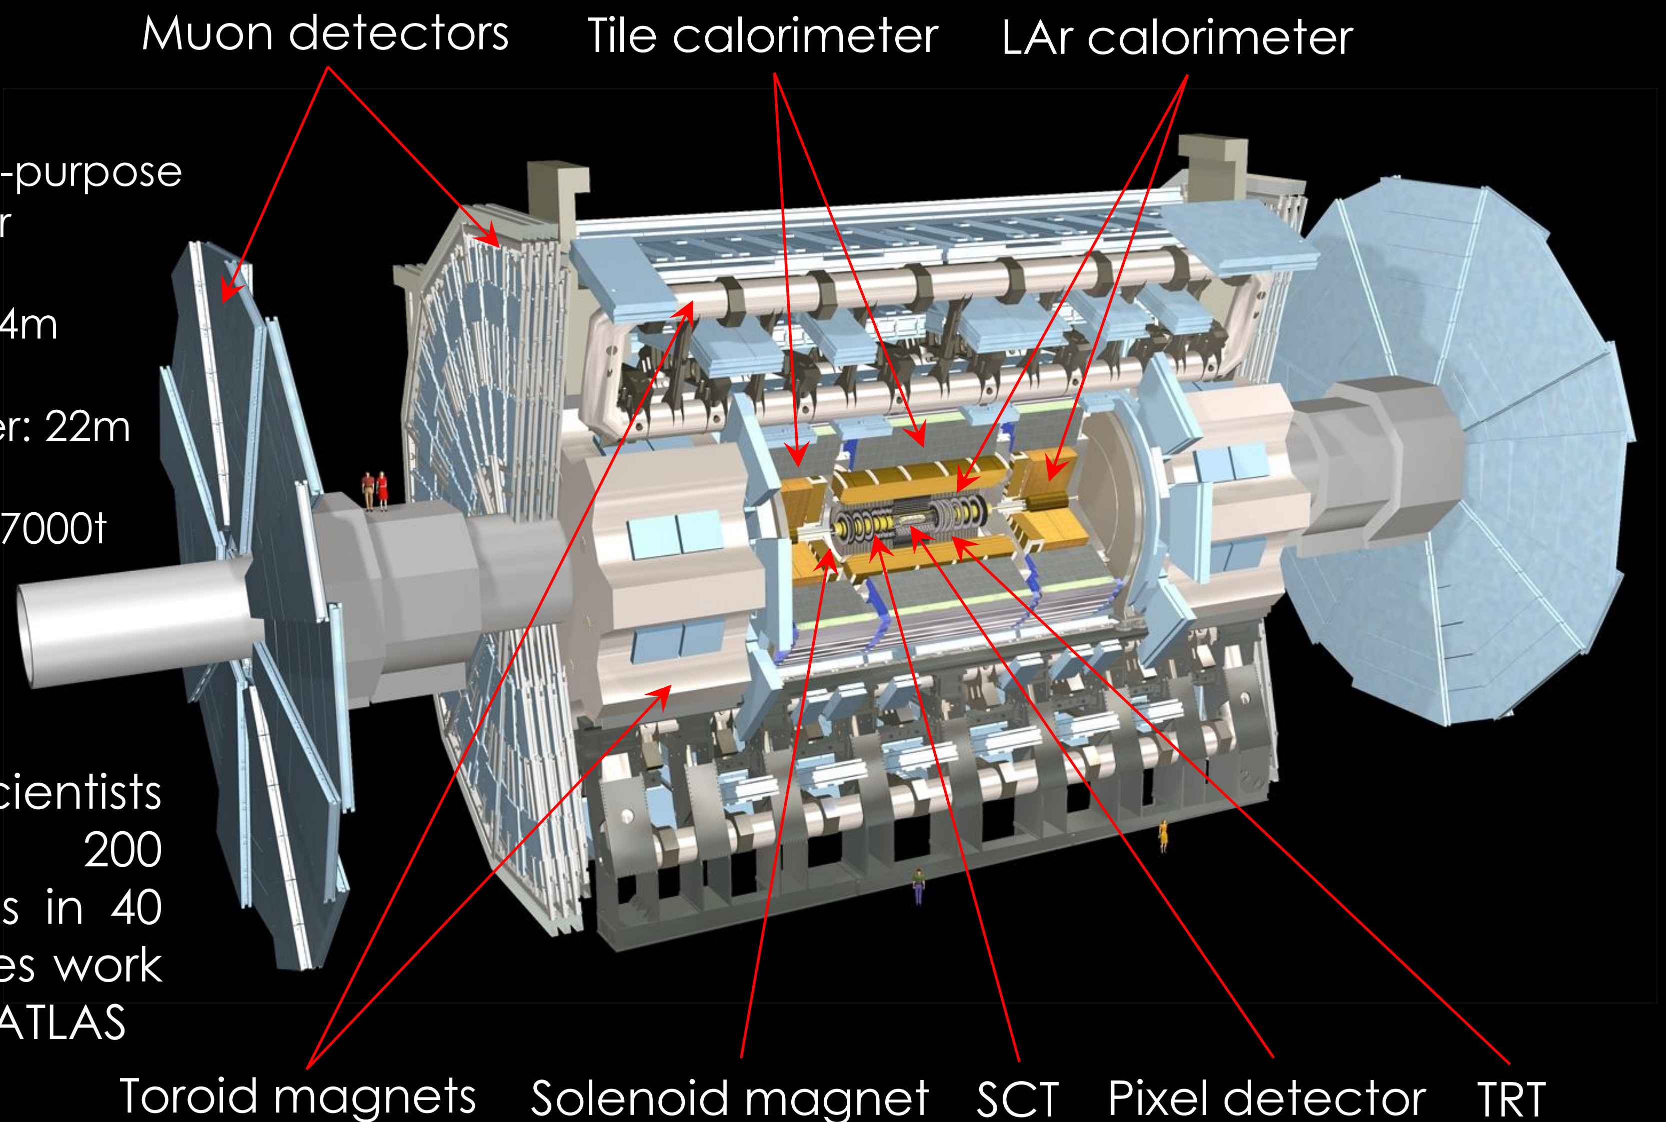

- General-purpose detector
- Width: 44m
- Diameter: 22m
- Weight: 7000t
- 3000 scientists from 200 institutes in 40 countries work on the ATLAS

WHY DO WE NEED HPC

- The main goal of the LHC-based experiments is to address fundamental physics questions, such as validation of the Standard Model (SM) of particle physics, search for new particles etc.
- With the ATLAS detector one can record data from the proton-proton collisions, but how it is possible to perform the measurements and search for new physics?
- The most effective way is to obtain theory predictions from Monte Carlo simulations, which are based on different theoretical models.

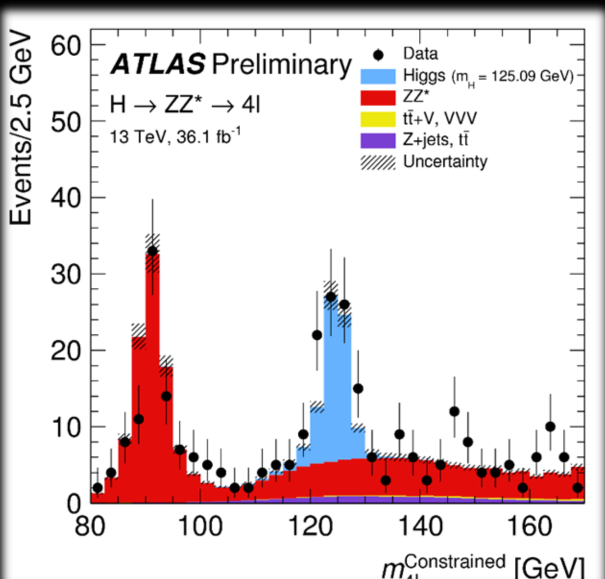

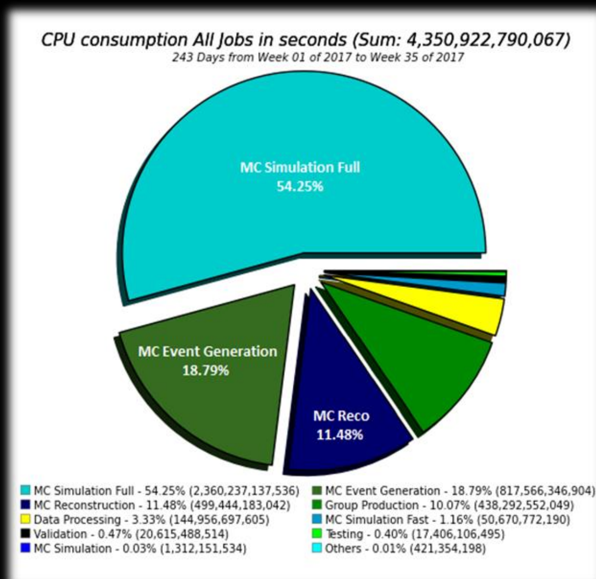

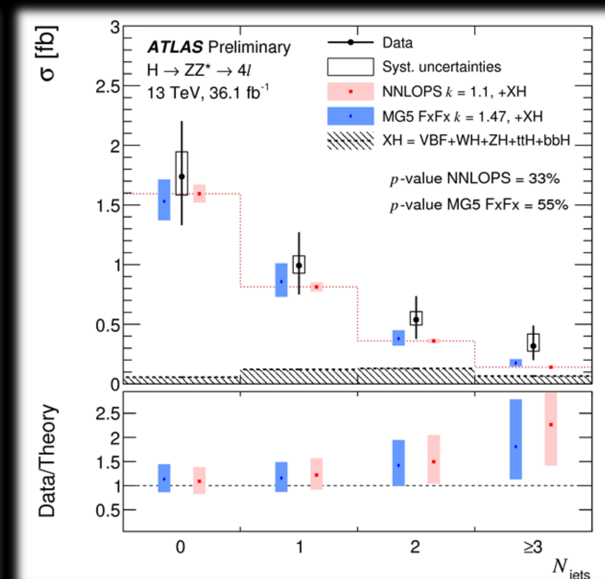

DATA/MC PRODUCTION CHAIN

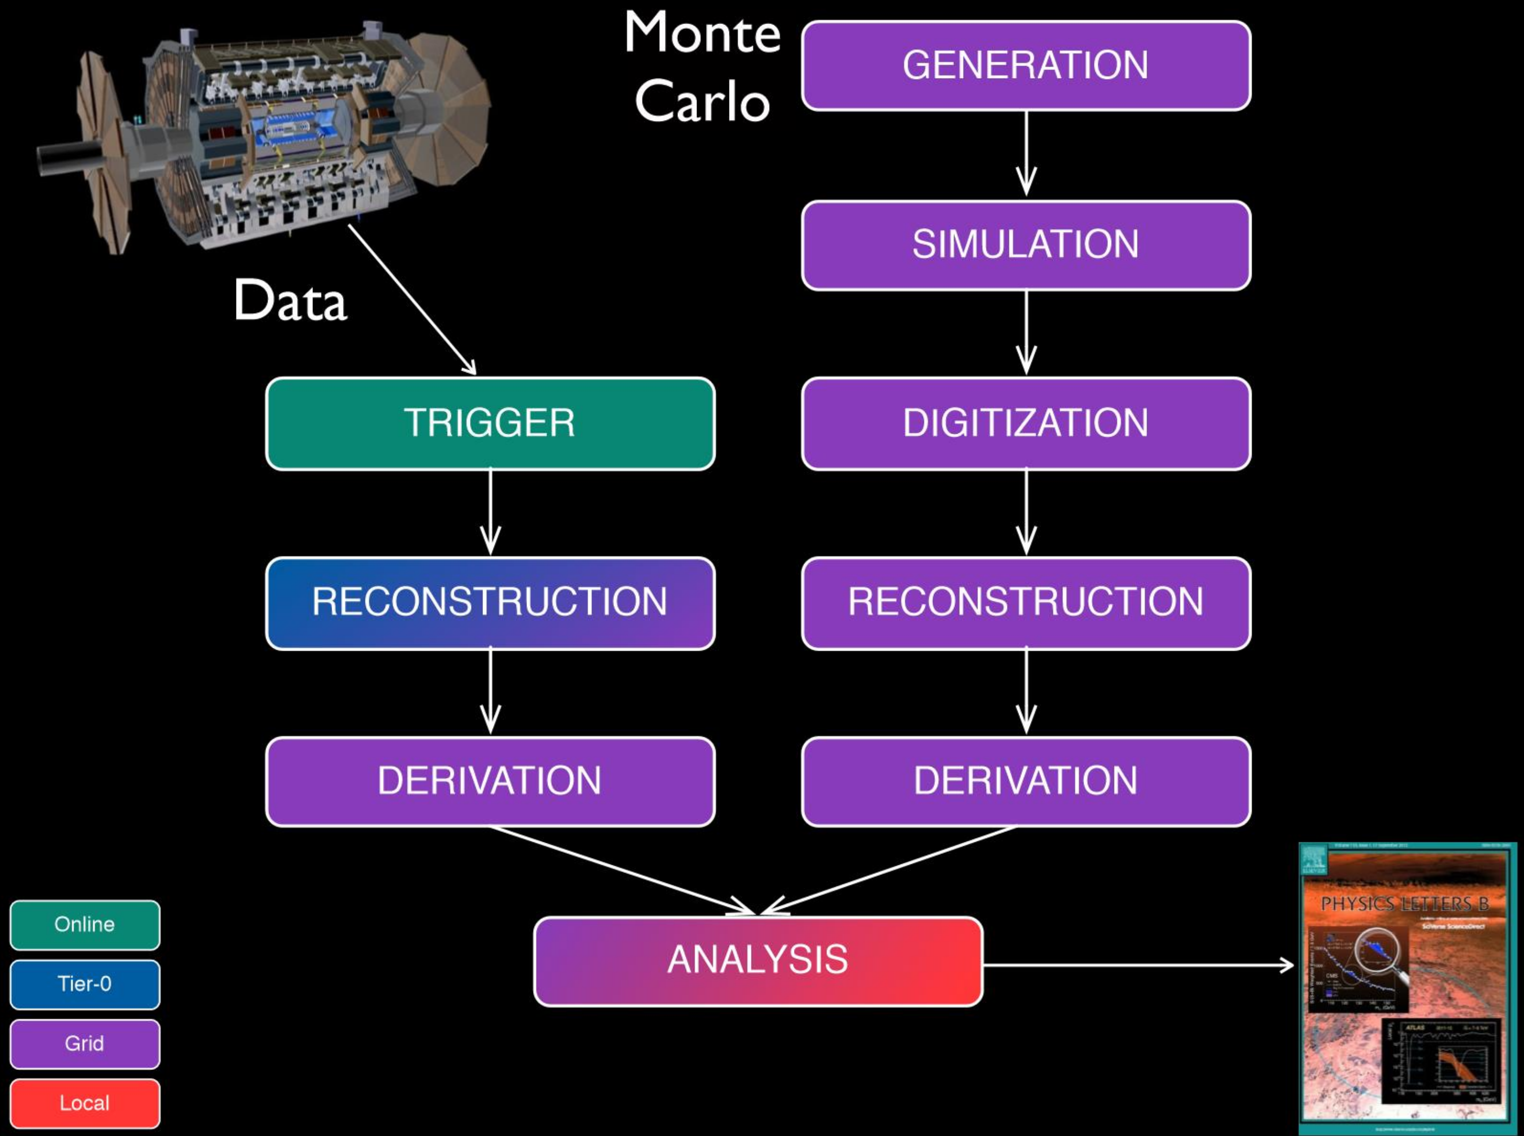

DATA/MC PRODUCTION CHAIN

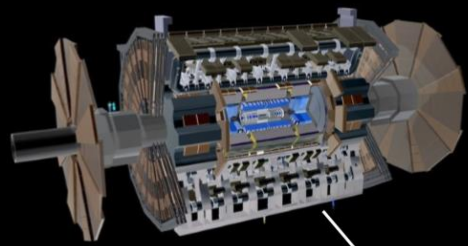

Data

Monte Carlo

GENERATION — *Event generator output (EVNT)*

SIMULATION — *Simulated interaction with detector (HITS)*

DIGITIZATION — *Simulated detector output (RDO)*

RECONSTRUCTION — *Analysis object data (AOD)*

DERIVATION — *Derived AOD (DAOD)*

Raw data (RAW) — TRIGGER

Analysis object data (AOD) — RECONSTRUCTION

Derived AOD (DAOD) — DERIVATION

ANALYSIS

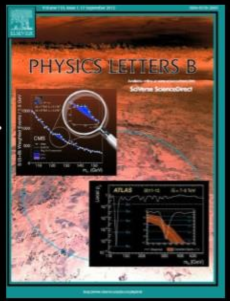

- Online
- Tier-0
- Grid
- Local

DETAILED SIMULATION CHAIN

Event simulation is a complex procedure which include many steps, such as generation, digitization, reconstruction and so on. The full chain is indicated below.

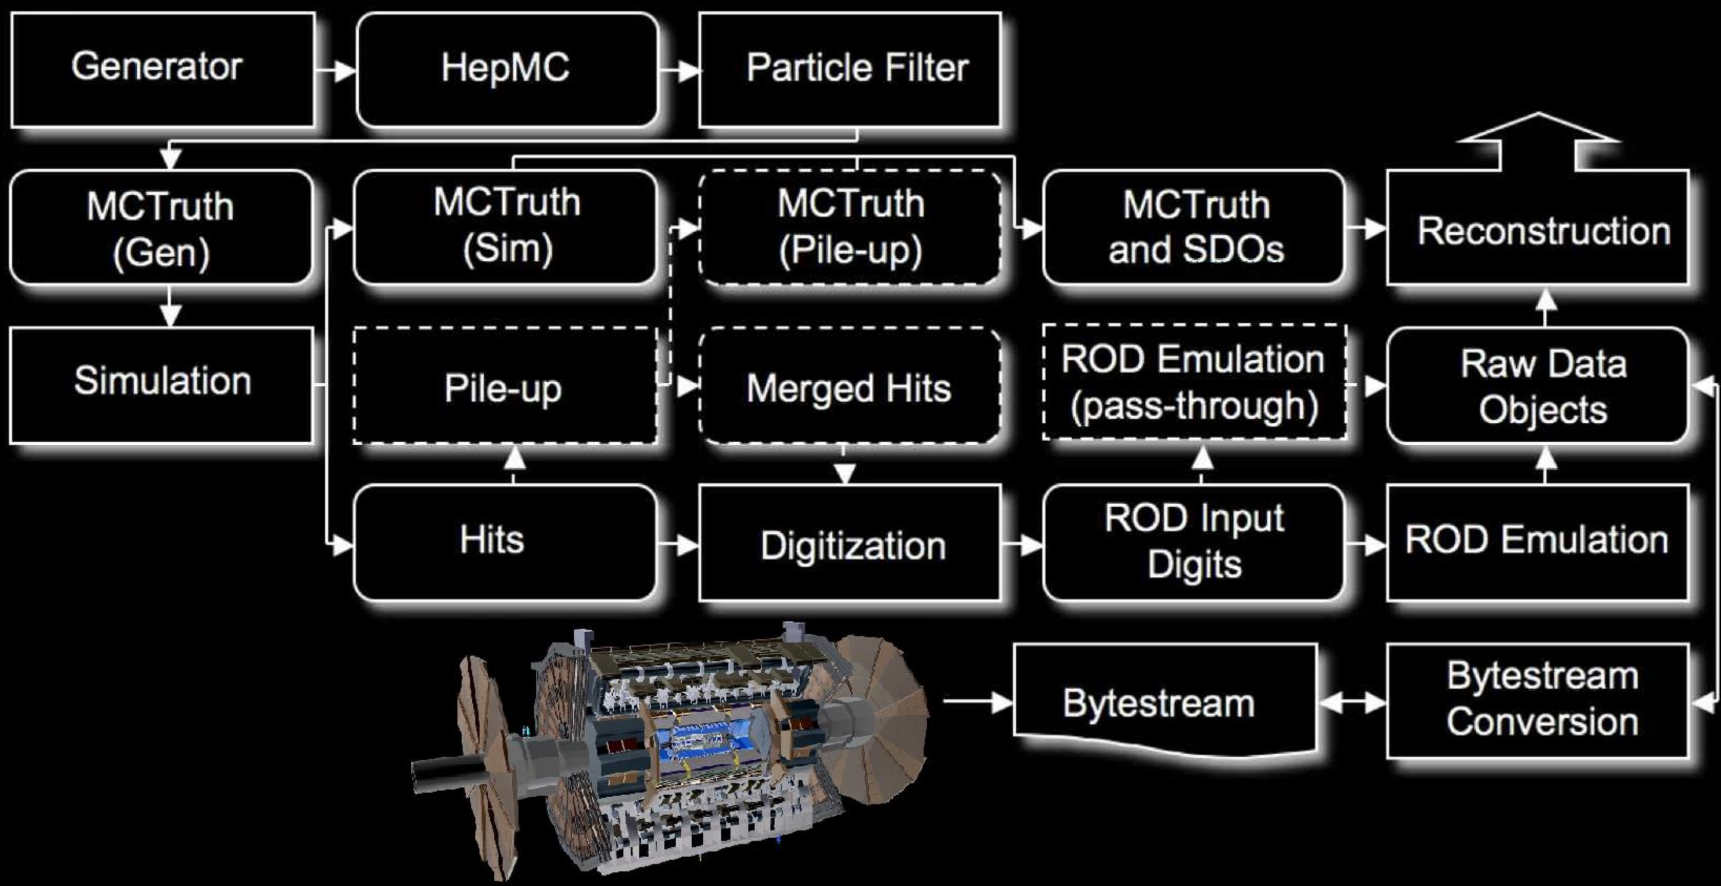

EVENT RECONSTRUCTION

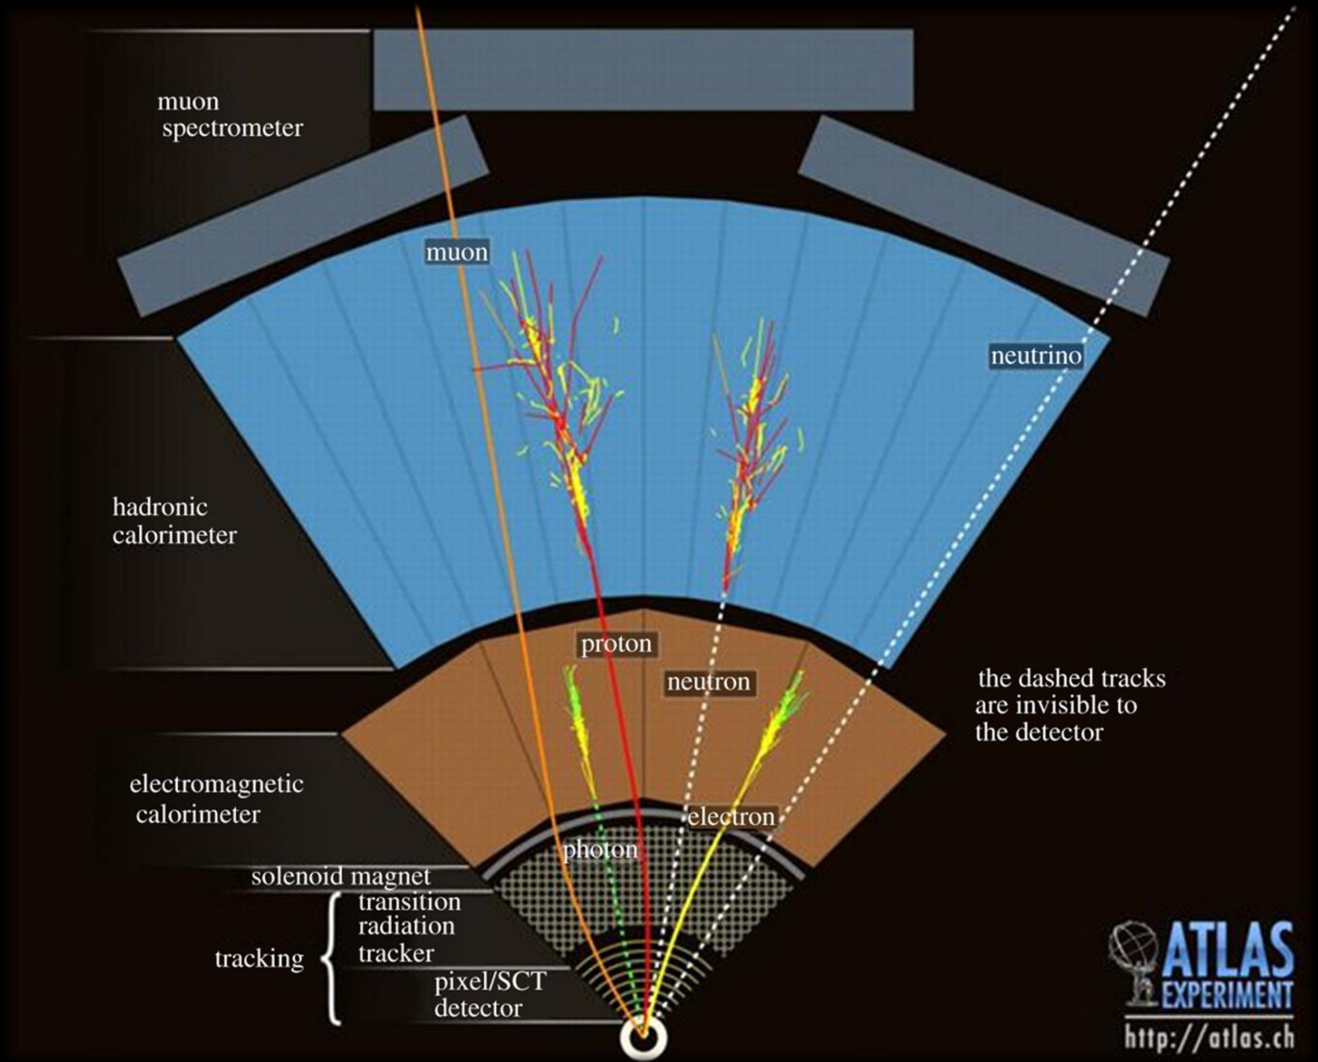

During the Monte Carlo simulation it is required to reconstruct tracks from the particles traveling through the detector volume.

EVENT RECONSTRUCTION

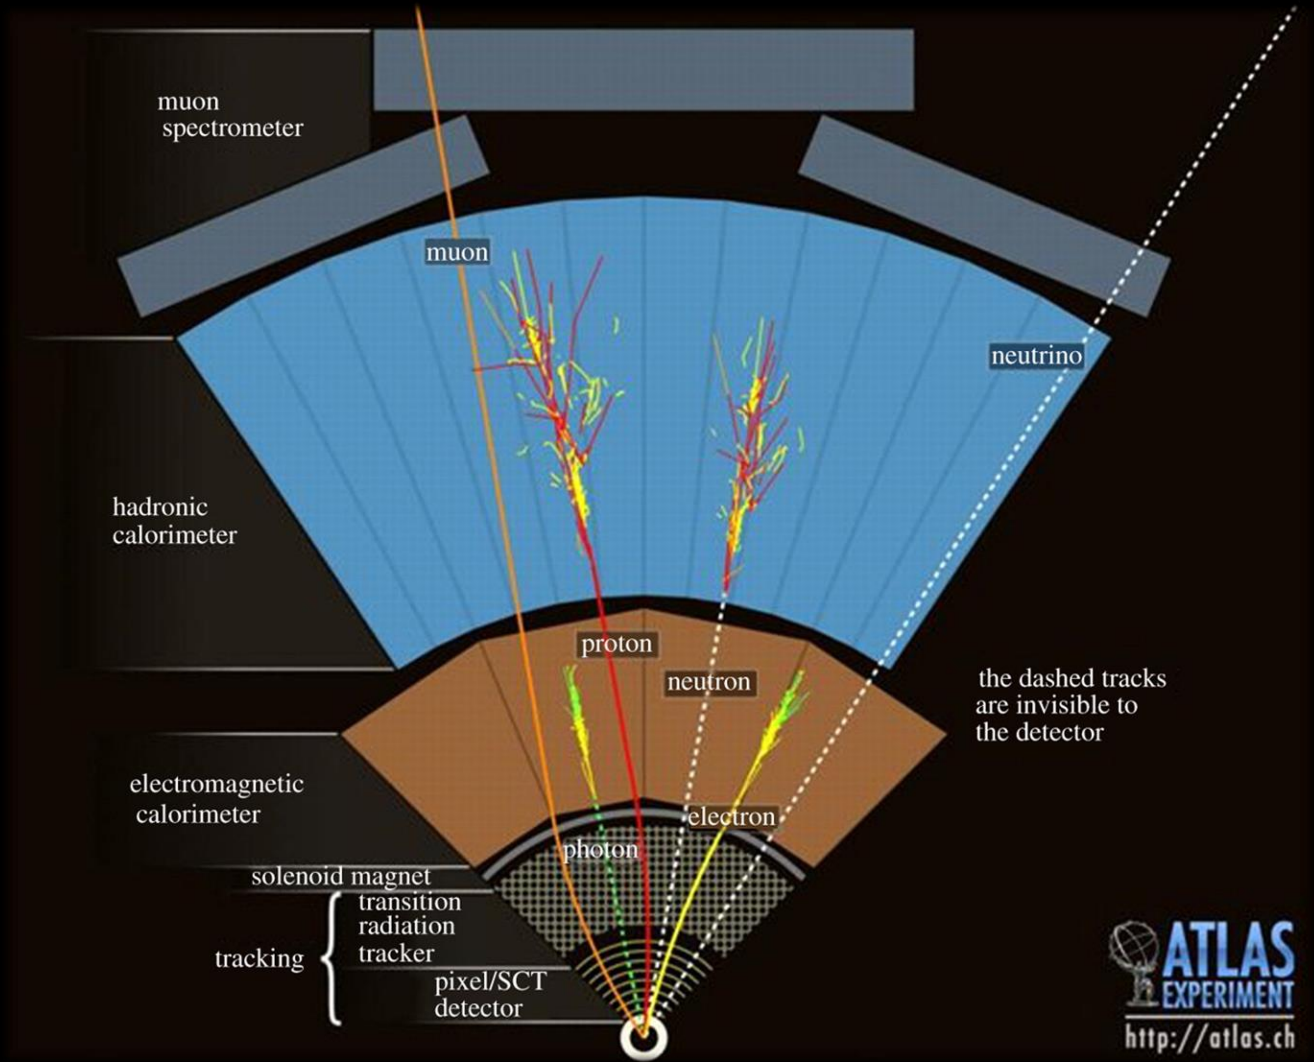

During the Monte Carlo simulation it is required to reconstruct tracks from the particles traveling through the detector volume.

Such simulation is complicated due to the fact, that we have to simulate the detector's response in each subsystem: Inner detector, Calorimeters, Muon chambers etc.

EVENT RECONSTRUCTION

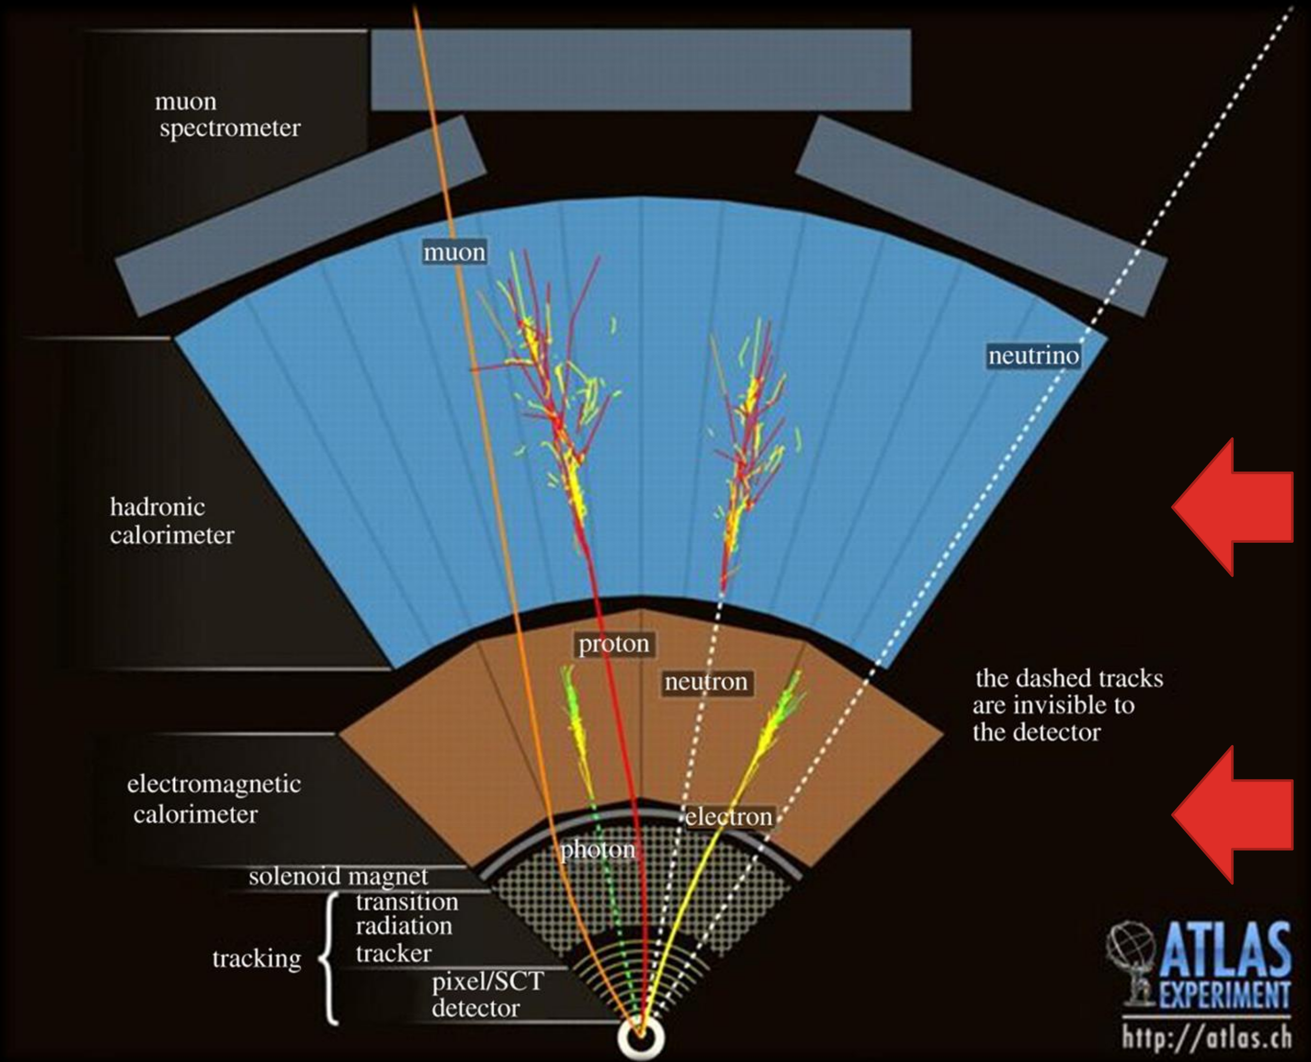

During the Monte Carlo simulation it is required to reconstruct tracks from the particles traveling through the detector volume.

Such simulation is complicated due to the fact, that we have to simulate the detector's response in each subsystem: Inner detector, Calorimeters, Muon chambers etc.

xAOD DATA FORMAT

- During the first data taking period of the LHC (Run 1), ATLAS data format was different for analysis and reconstruction.

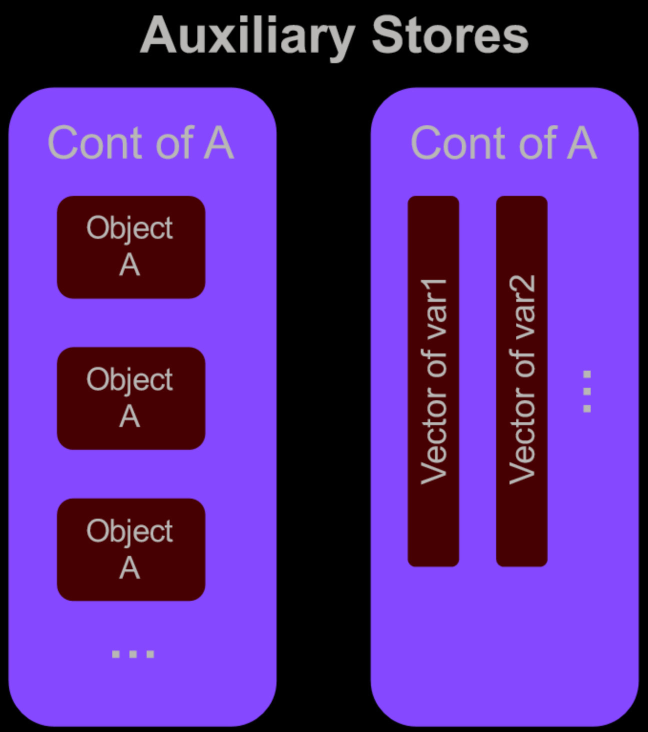

- This format had two aspects which made it unpopular among the analyzers but popular among the reconstruction people.
  1. It was fast to retrieve groups of events.
  2. This format was optimized for space which required object reconstitution for some objects.
- For the analysis other properties were wanted.
  1. Fast retrieval of individual variables.
  2. Direct usability in ROOT with no externals.
- For Run 2 the ATLAS collaboration developed new data file format called xAOD.
- By using an object called an auxiliary store, it is now possible to write data in either format simply by changing the ROOT settings. Some of the advantages of Run 1 were compromised, but the advantage of a single format outweighs them.

ATLAS ANALYSIS SCHEME

- Athena: ATLAS software execution control framework.
- Skimming/slimming: a method to filter out the events to be written in the output file.

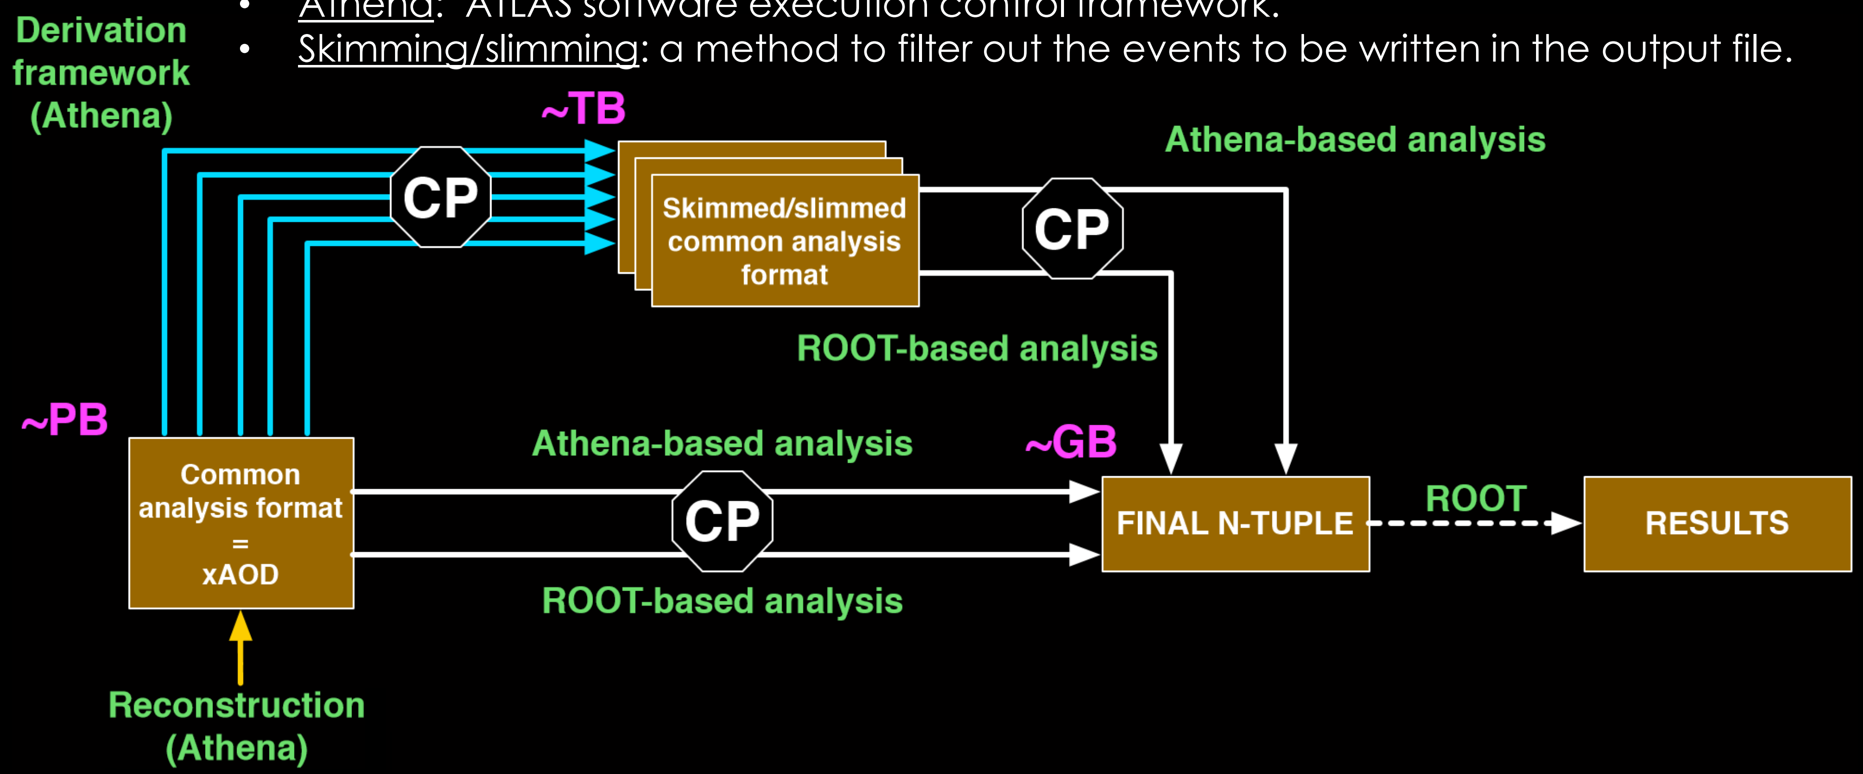

ATLAS ANALYSIS SCHEME

- Athena: ATLAS software execution control framework.
- Skimming/slimming: a method to filter out the events to be written in the output file.

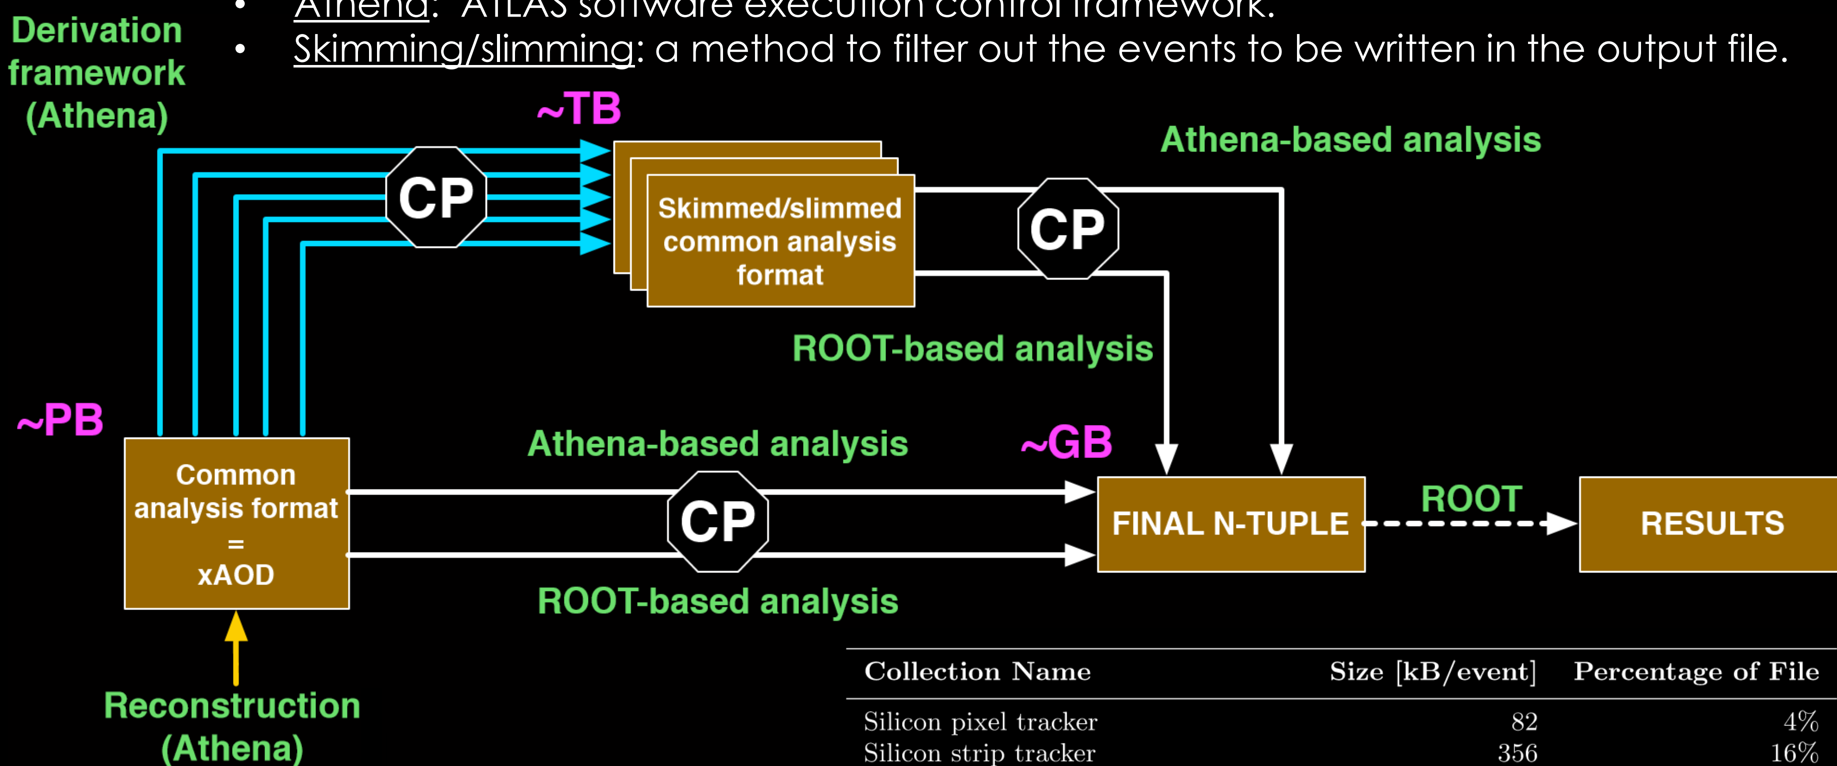

- Size of RAW samples, which represent event data flowing from the trigger system (Run 1).
- In Run 2, size of samples is quite similar (~1.5 MB/event).

Collection Name	Size [kB/event]	Percentage of File
Silicon pixel tracker	82	4%
Silicon strip tracker	356	16%
Transition radiation tracker	921	46%
Electromagnetic Barrel calorimeter	89	4%
Electromagnetic Endcap calorimeter	104	5%
Hadronic Barrel calorimeter	29	1%
Hadronic Endcap calorimeter	22	1%
Forward calorimeter	42	2%
calorimeter calibration hits	243	12%
Muon system (all collections)	3	<1%
Truth (all collections)	134	8%
Total	**1987**	**100%**

JEDI/PANDA SERVICE

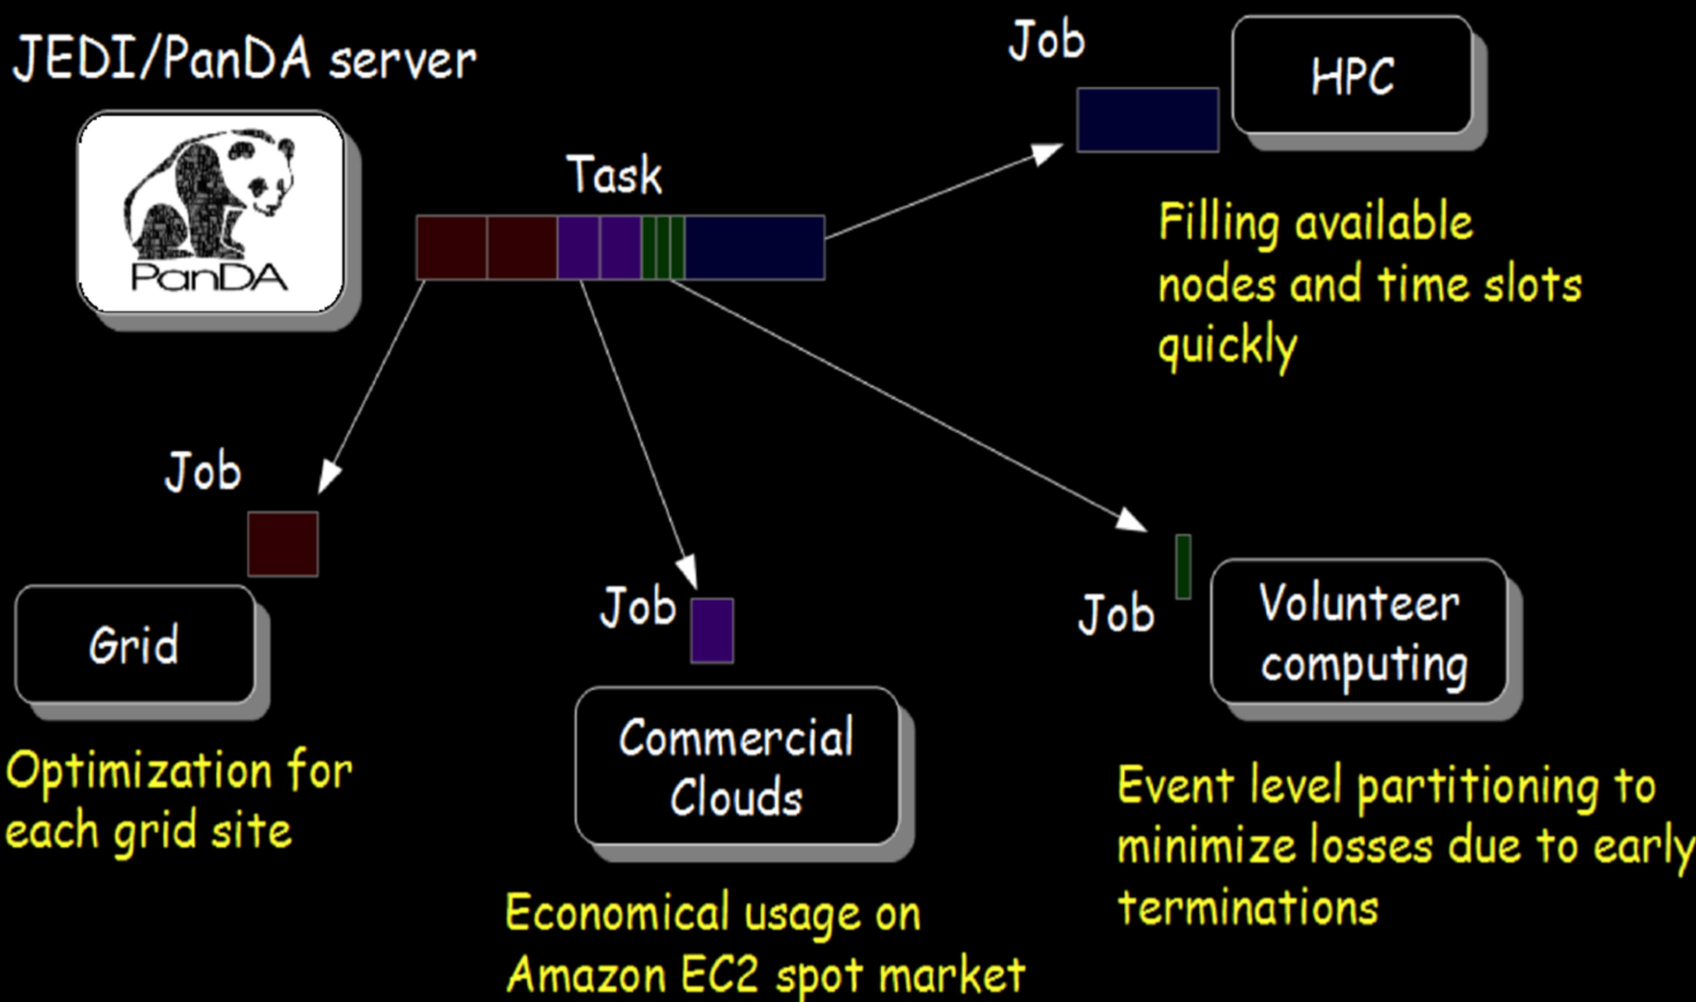

RUNNING GRID JOBS: ALL

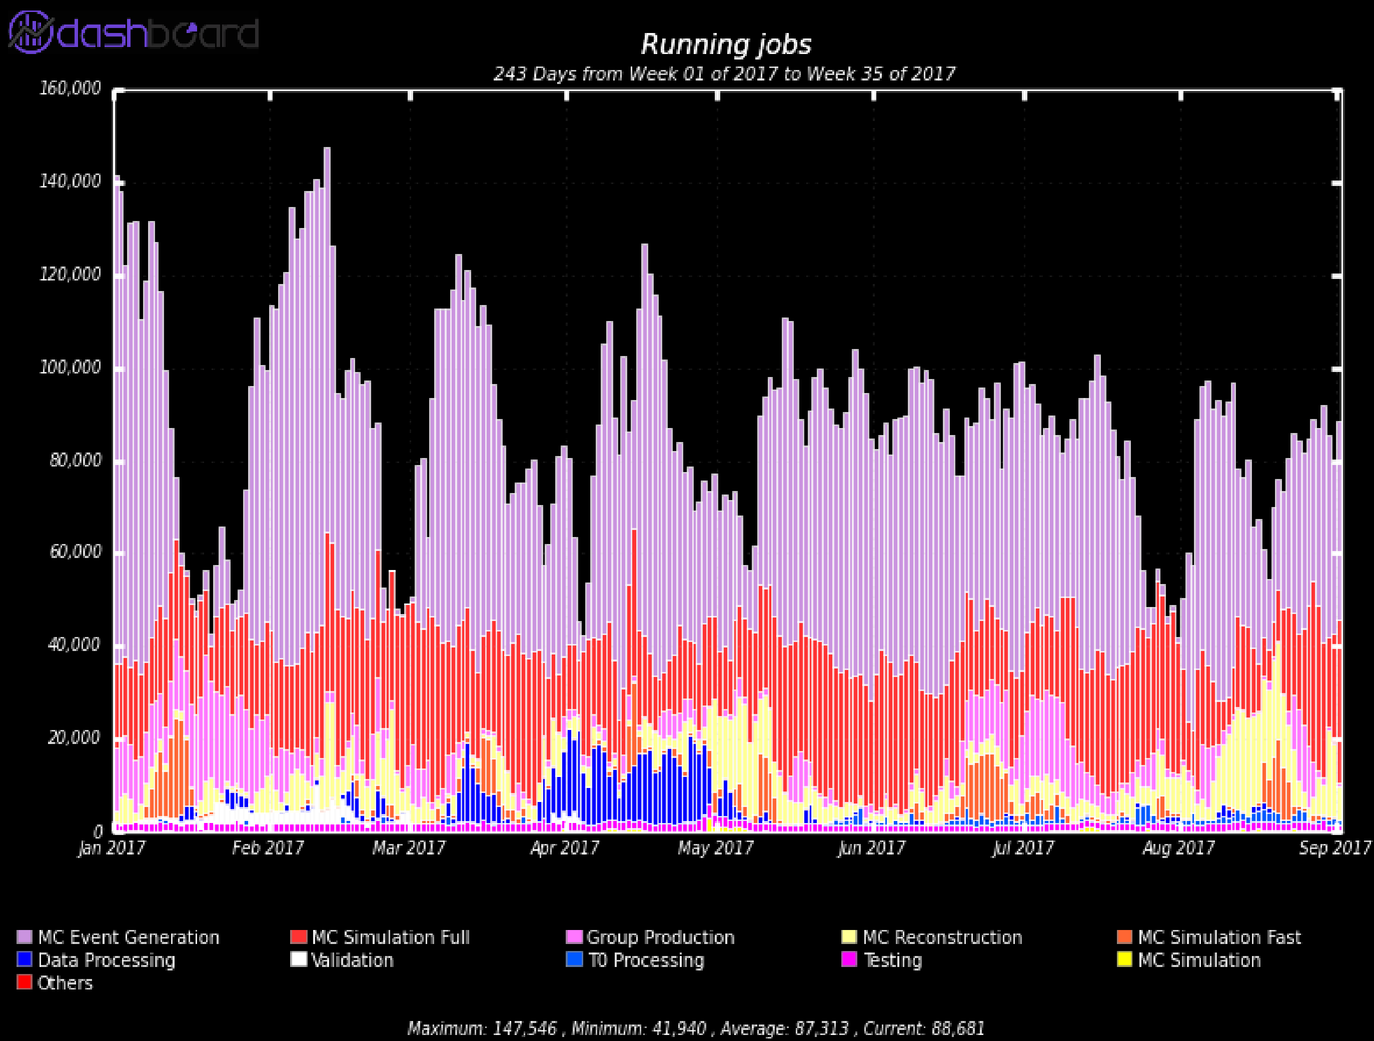

During 2017 the highest fraction of all running jobs was related to:

- MC Event Generation
- MC Simulation Full
- Group Production

Maximum daily amount of running jobs was equal to ~150k.

Number of submitted jobs: 153 520 492

Number of completed jobs: 135 636 142

RUNNING GRID JOBS: HIGGS

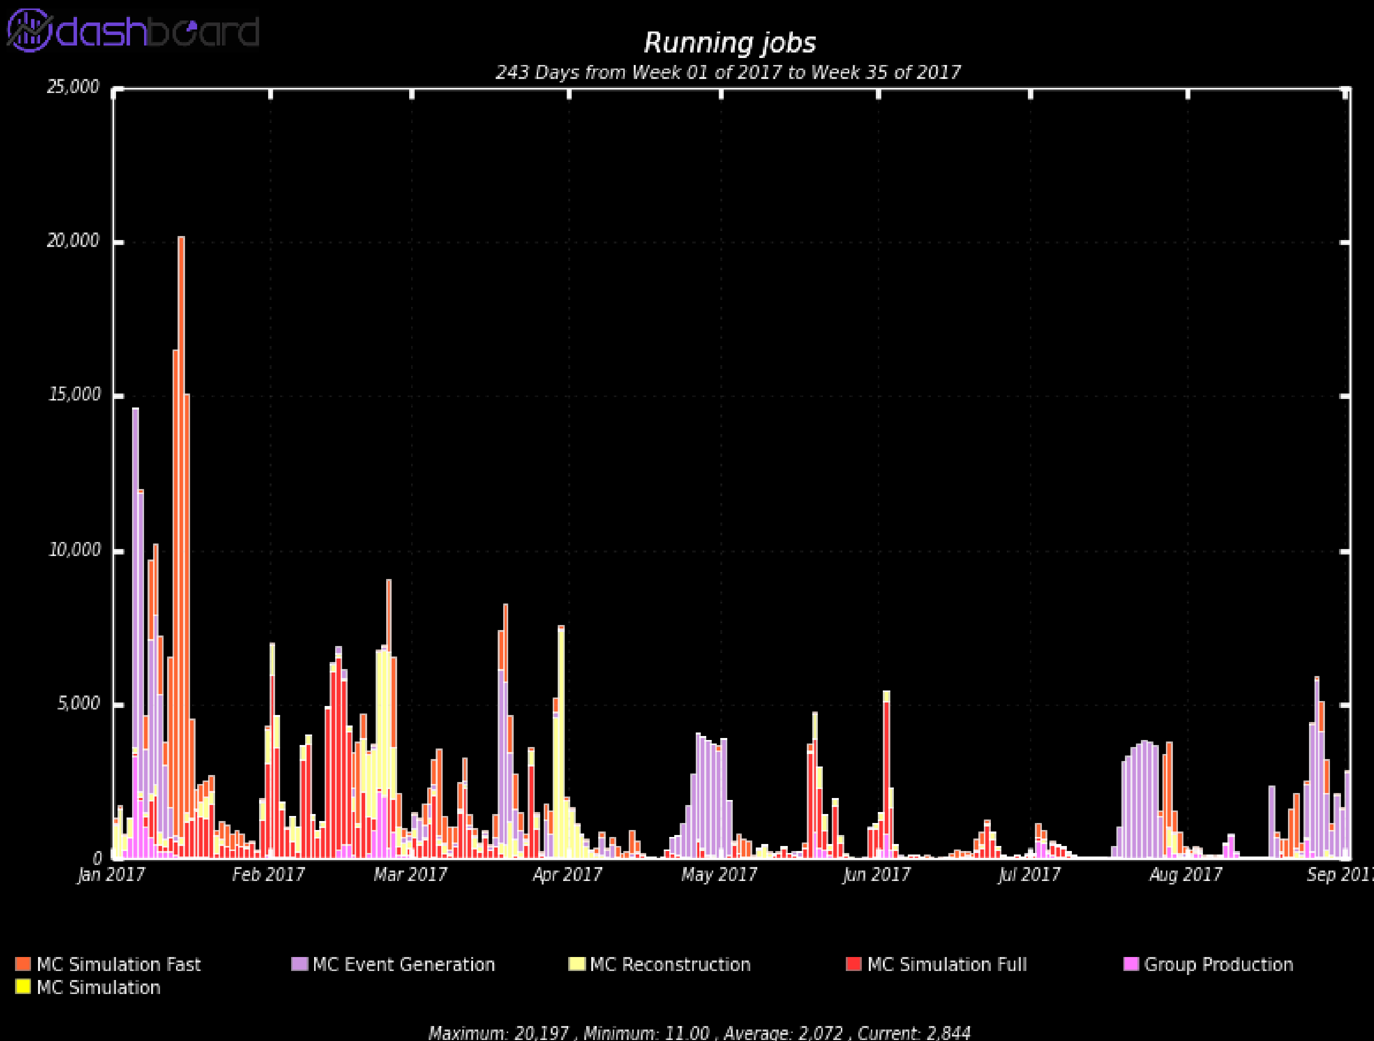

It is possible to estimate the amount of Higgs jobs by checking the corresponding production groups, namely:

- ap_higg
- gp_higg
- phys-higgs

Number of submitted jobs: 3 766 426

Number of completed jobs: 3 144 899

However, these numbers don't include private and unknown tasks.

HIGGS BOSON PHYSICS

- Finding the Higgs boson is by far the most important discovery within particle physics during recent years.

- Following discovery it is important to measure the properties of the Higgs boson to verify the SM and look for deviations that could present hints to new physics.

- LHC provides the access to a large number of production and decay modes.

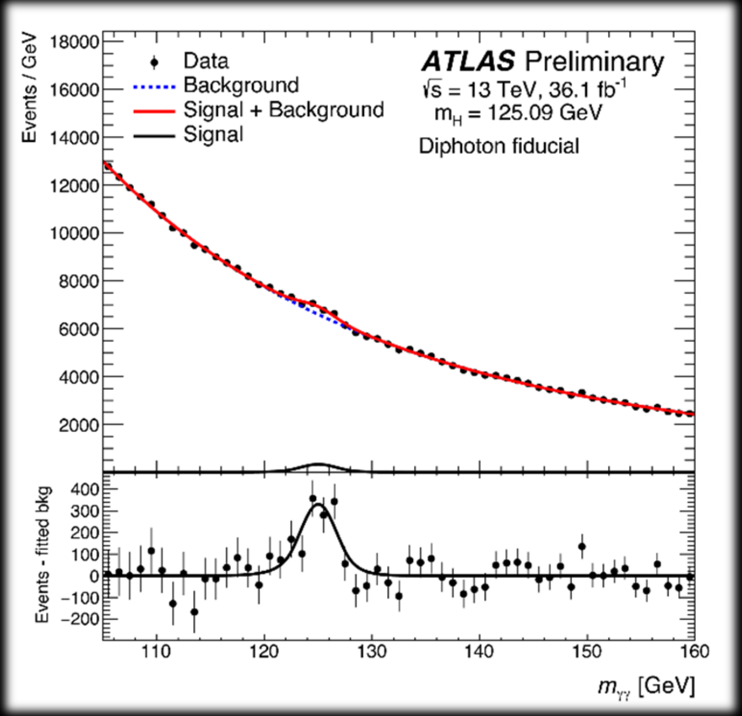

- Statistics taken by the LHC is rapidly increasing in time as well as the number of theories beyond the Standard Model (BSM) that has to be checked.

- <u>All these facts require enormous number of Monte Carlo samples to be produced.</u>

LO AND NLO GENERATION

Leading order (LO) gluon fusion (ggF) diagram

Next-to-leading order (NLO) gluon fusion diagrams

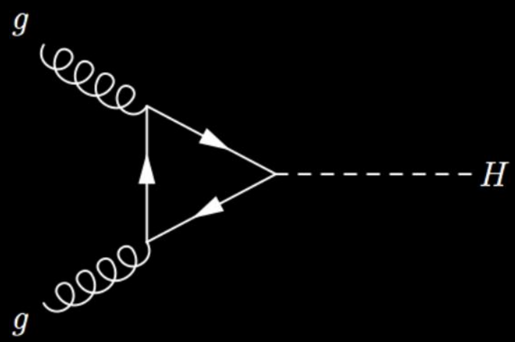

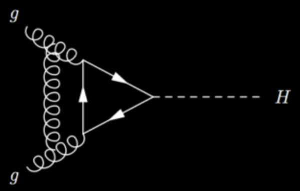

0 jets

0 jets

…
plus many others similar Feynman diagrams!

- NLO generation requires much more CPU time than LO. Generation times of 100k $pp \to H + 0j$ events at CERN LXPLUS machines are indicated in table below.

Number of jets	0		
Number of Feynman diagrams	24		
Generation time	~10 minutes		

LO AND NLO GENERATION

Leading order (LO) gluon fusion (ggF) diagram

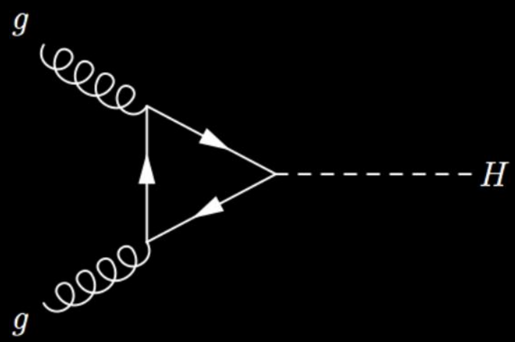

0 jets

Next-to-leading order (NLO) gluon fusion diagrams

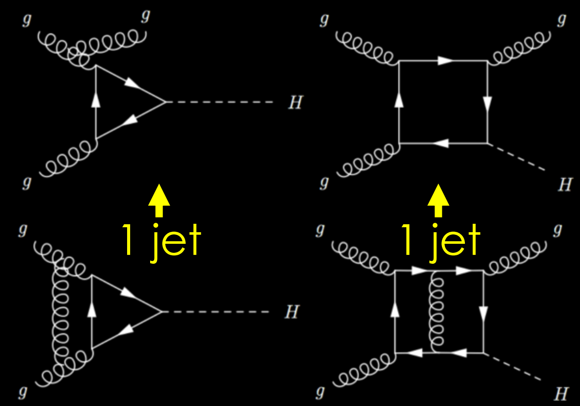

0 jets

1 jet

…

plus many others similar Feynman diagrams!

- NLO generation requires much more CPU time than LO. Generation times of 100k $pp \to H + 0j/1j$ events at CERN LXPLUS machines are indicated in table below.

Number of jets	0	1	
Number of Feynman diagrams	24	1050	
Generation time	~10 minutes	~2 hours	

LO AND NLO GENERATION

Leading order (LO) gluon fusion (ggF) diagram

Next-to-leading order (NLO) gluon fusion diagrams

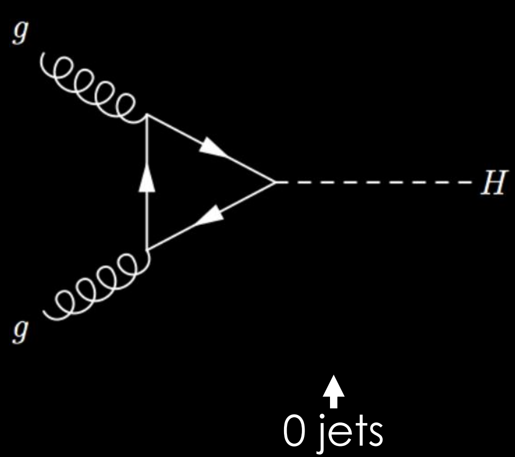

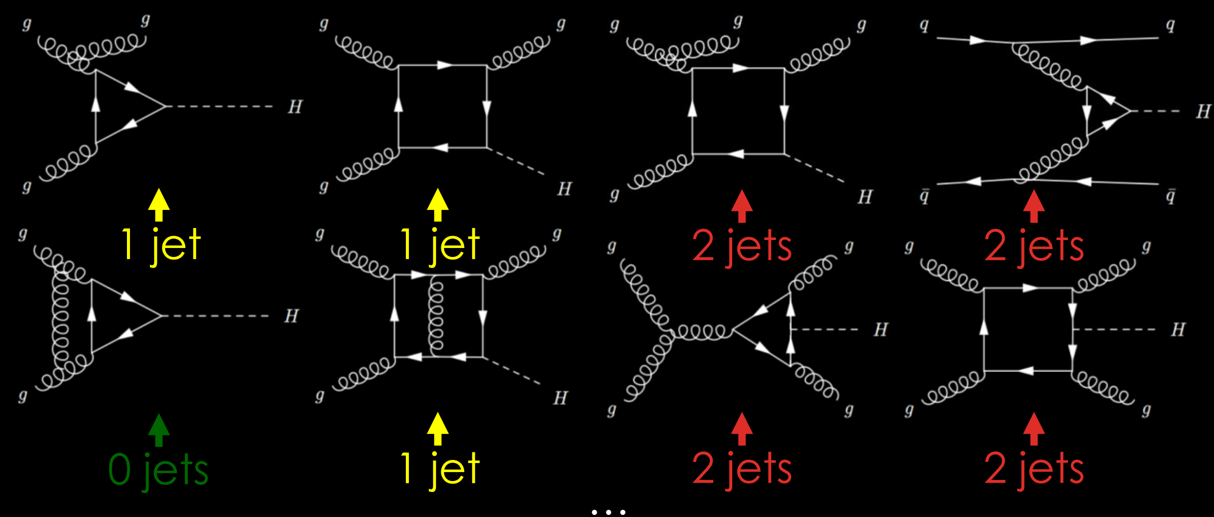

plus many others similar Feynman diagrams!

- NLO generation requires much more CPU time than LO. Generation times of 100k $pp \to H + 0j/1j/2j$ events at CERN LXPLUS machines are indicated in table below.

Number of jets	0	1	2
Number of Feynman diagrams	24	1050	21510
Generation time	~10 minutes	~2 hours	~24 hours

HPC: TITAN SUPERCOMPUTER

OAK RIDGE
National Laboratory

- 27 PFLOPS (Peak theoretical performance)
- Cray XK-7 18,688 compute nodes with GPUs
- 299,008 CPU cores AMD Opteron 6274 @2.2 GHz (16 cores per node)
- 32 GB RAM per node
- NVidia TESLA K20x GPU per node
- 32 PB disk storage (center-wide Luster file system)
- >1TB/s aggregate FS throughput
- 29 PB HPSS tape archive

Titan's performance is higher than power of all ATLAS resources in WLCG!

PANDA INTERFACE WITH TITAN

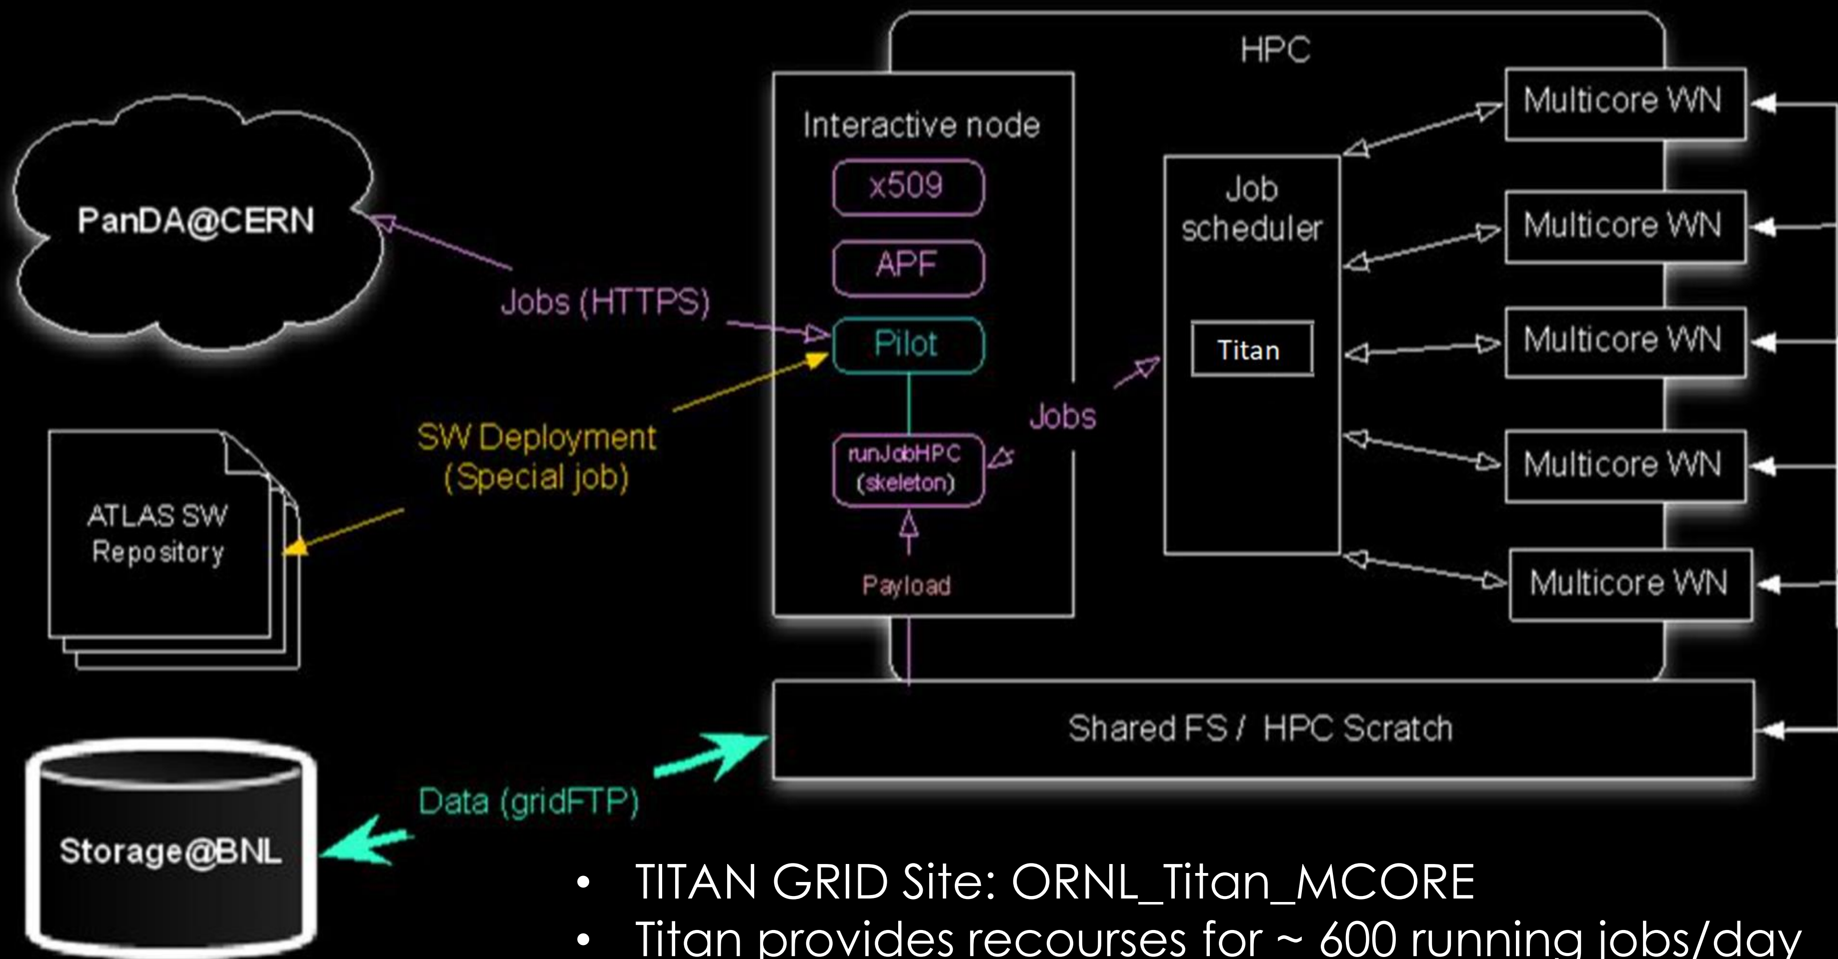

- TITAN GRID Site: ORNL_Titan_MCORE
- Titan provides recourses for ~ 600 running jobs/day
- Almost 4M of completed jobs for 9 month of 2017

More details will be presented by D. Oleinik.

TITAN AND PANDA

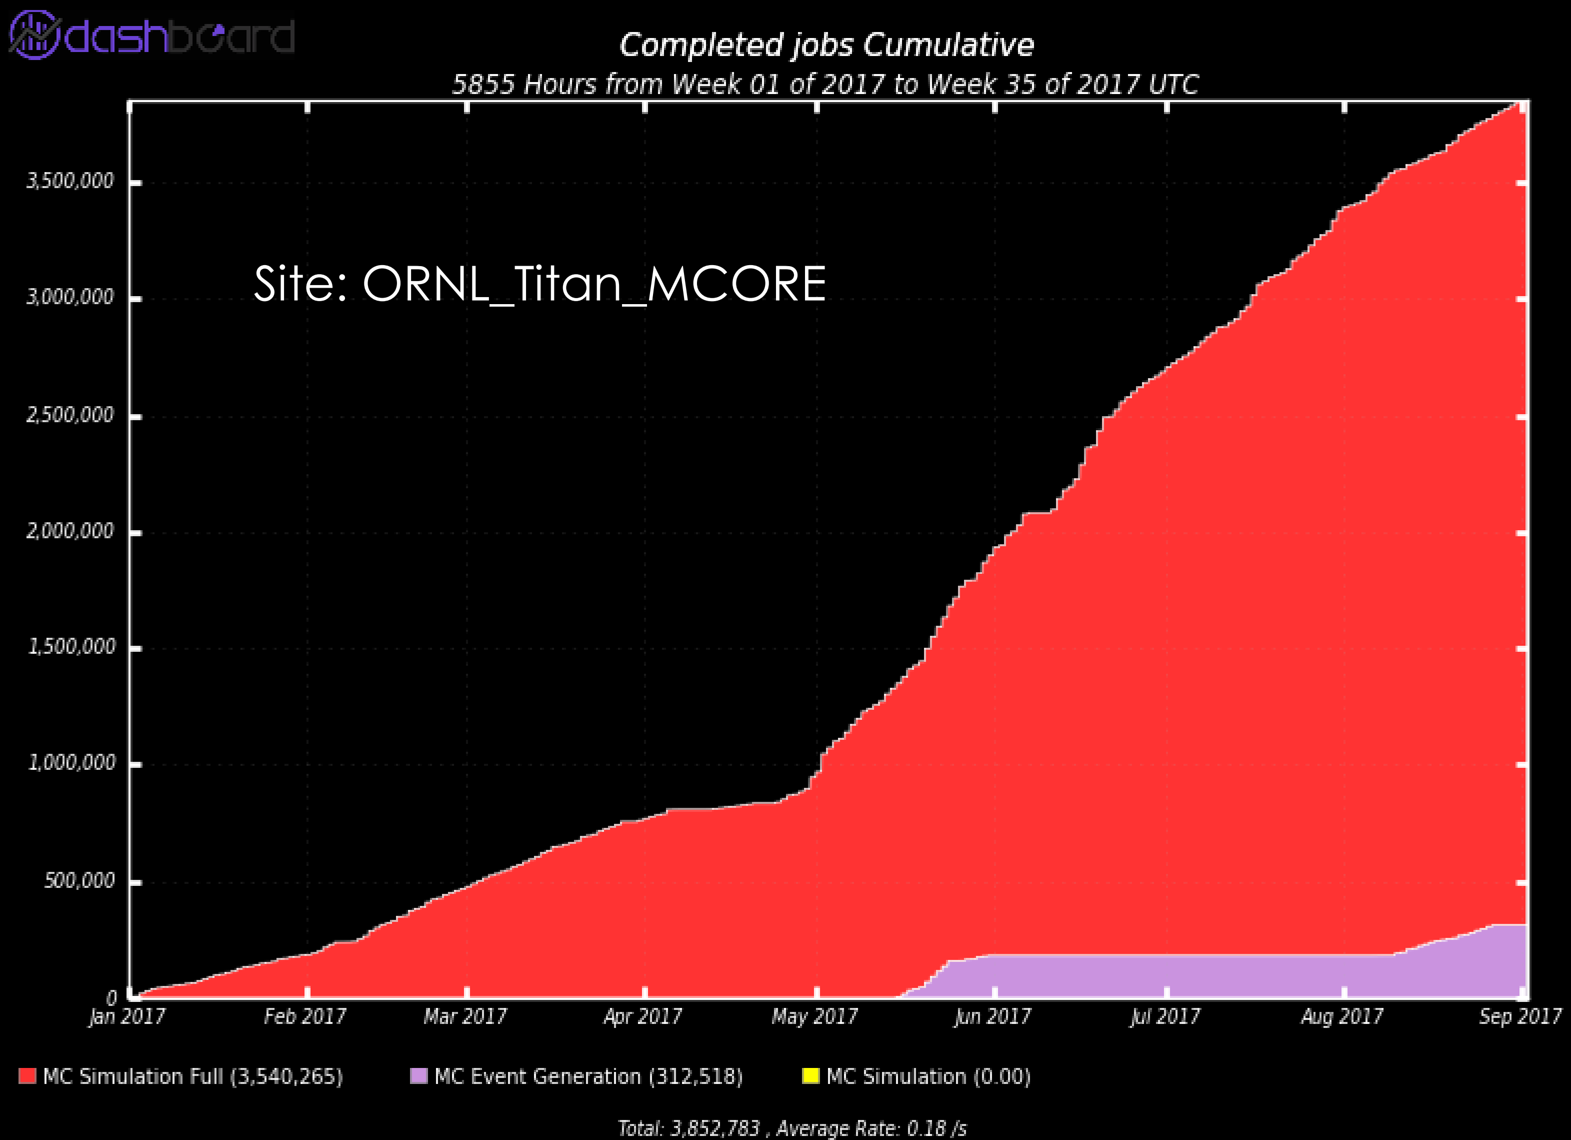

HIGGS PHYSICS SIMULATION

- Titan was used to produce $pp \to H + 0j/1j/2j$ Monte Carlo samples with the PowhegBox/HJJ generator at NLO level. Signature of this process can be easily reconstructed in ATLAS and correlations of the final state particles are sensitive to new physics.
- The results of simulation show, that while with the LXPLUS the generation event rate is about 2.5K events/hour, with the Titan it's about 650K/hour!

Quite impressive!

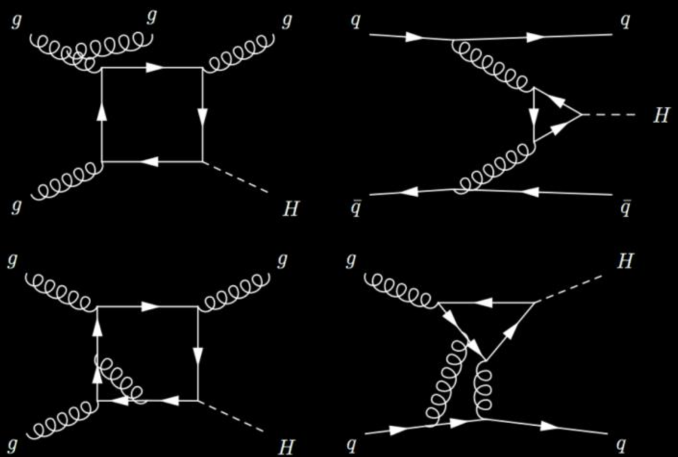

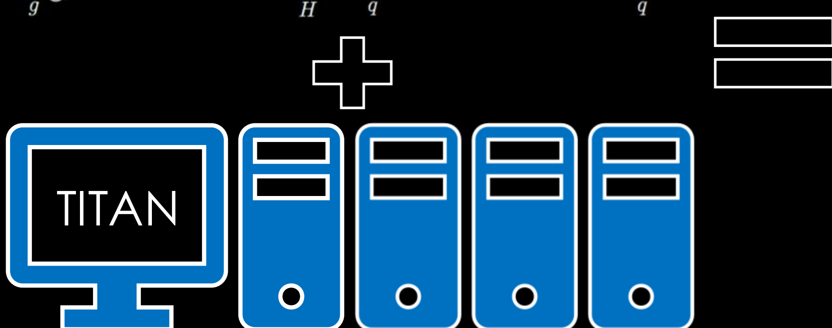

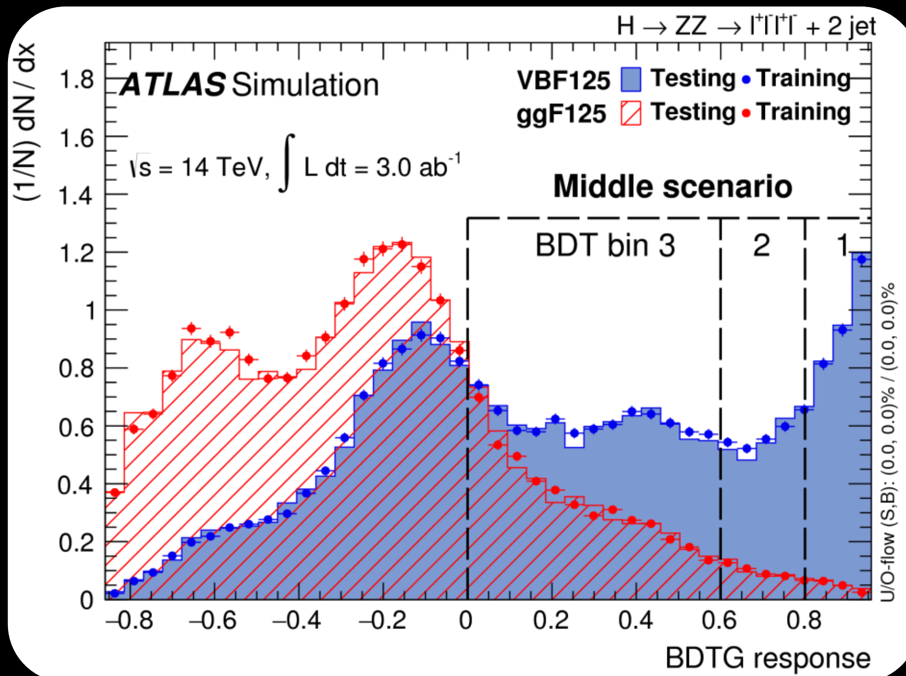

HIGGS PHYSICS SIMULATION

- After discussed performance tests, Titan was used to produce additional $pp \rightarrow Hjj \rightarrow ZZjj \rightarrow 4ljj$ events, but at Full Simulation level. ATLAS Full Simulation consists of many stages (the whole detector should be simulated) and this procedure requires significant CPU time to be completed.
- Total statistics was set to 6M of Higgs events.
- Task: mc15_13TeV.341939.Sherpa_CT10_2DP20_myy_100_165.simul.e4407_s2726.
- Prodsys task parameters:

nEventsPerInputFile	nEventsPerJob	nFiles	CPU Time, HS06 seconds	
			Expected	Total
2000	100	3000	72048000000	2284598809

- Samples produced after this simulation were used for physics analysis. This analysis was related to Higgs physics prospects at the High-Luminosity LHC. The obtained results were published by the collaboration: **ATL-PHYS-PUB-2016-008** and **PoS ICHEP2016 (2017) 426.**

NRC-KI COMPUTING FACILITIES

The second generation HPC with the peak performance equal to 122,9 TFLOPS, №2 in 15-th issue of Russian Top50 supercomputers.

- 10240 CPU cores = 1280 nodes 2x Intel Xeon E5450 3,00GHz 4 core 16 Gb RAM
- User Interface (UI) node allows to run jobs in batch system (SLURM) and compile the code
- Shared FS Lustre for Worker Node's (WN's) and UI
- CVMFS connected to WN's
- Broadband to Tier-1 Storage Element (ANALY_RRC-KI GRID site)

NRC-KI COMPUTING FACILITIES

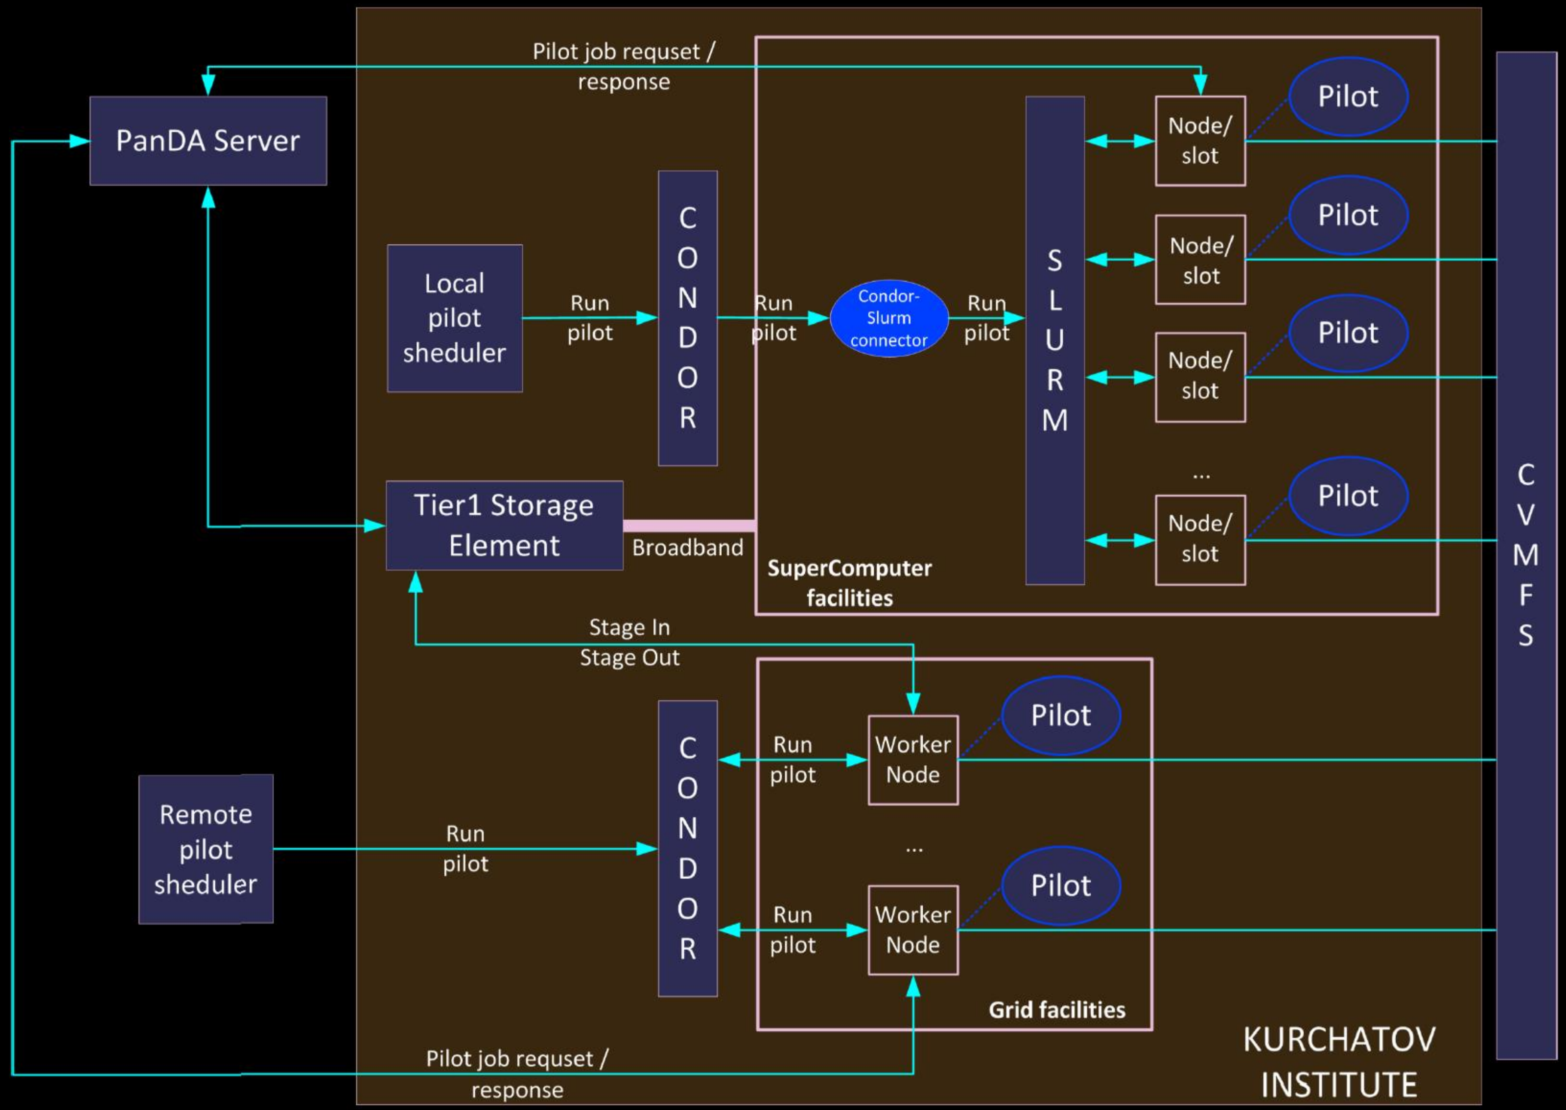

PERFORMANCE TESTS

- For the performance tests, 50k of $pp \rightarrow Hjj$ gluon fusion events were simulated at the generator level with the aMC_NLO Monte Carlo generator.
- The same tasks were submitted to 6 different GRID sites, including NRC-KI Tier 1 GRID site.
- The results of performance tests can vary depends on the condition of each individual machine, its occupancy and some other parameters. However, the performance of considered GRID sites looks more or less the same.

Site	nEventsPerJob	nFiles	CPU Time, HS06 seconds	
			Expected	Total
ANALY_RRC-KI-T1	5000	10	28500000	17197806
ANALY_DESY-HH	5000	10	19550000	11357616
ANALY_INFN-LECCE	5000	10	17648874	10350000
ANALY_GLASGOW_SL6	5000	10	15050000	9020417
ANALY_IN2P3-CC	5000	10	15250000	4677853
ANALY_TOKYO_ARC	5000	10	14000000	9457170

CONCLUSIONS

- Overview of the High performance computing systems in the ATLAS experiment was presented.
- One of the main physics fields where HPCS provides the crucial impact is the Higgs boson studies.
- Overview of the Titan supercomputer was also presented and its capabilities for the Higgs boson physics was demonstrated.
- Performance of NRC-KI Computing cluster was studied and comparison with other ATLAS GRID sites was demonstrated.
- This comparison demonstrates that NRC-KI Computing cluster is capable to handle the most actual tasks in the High energy physics field on a par with other ATLAS GRID sites.

THANKS FOR YOUR ATTENTION!

BACKUP

DERIVATIONS

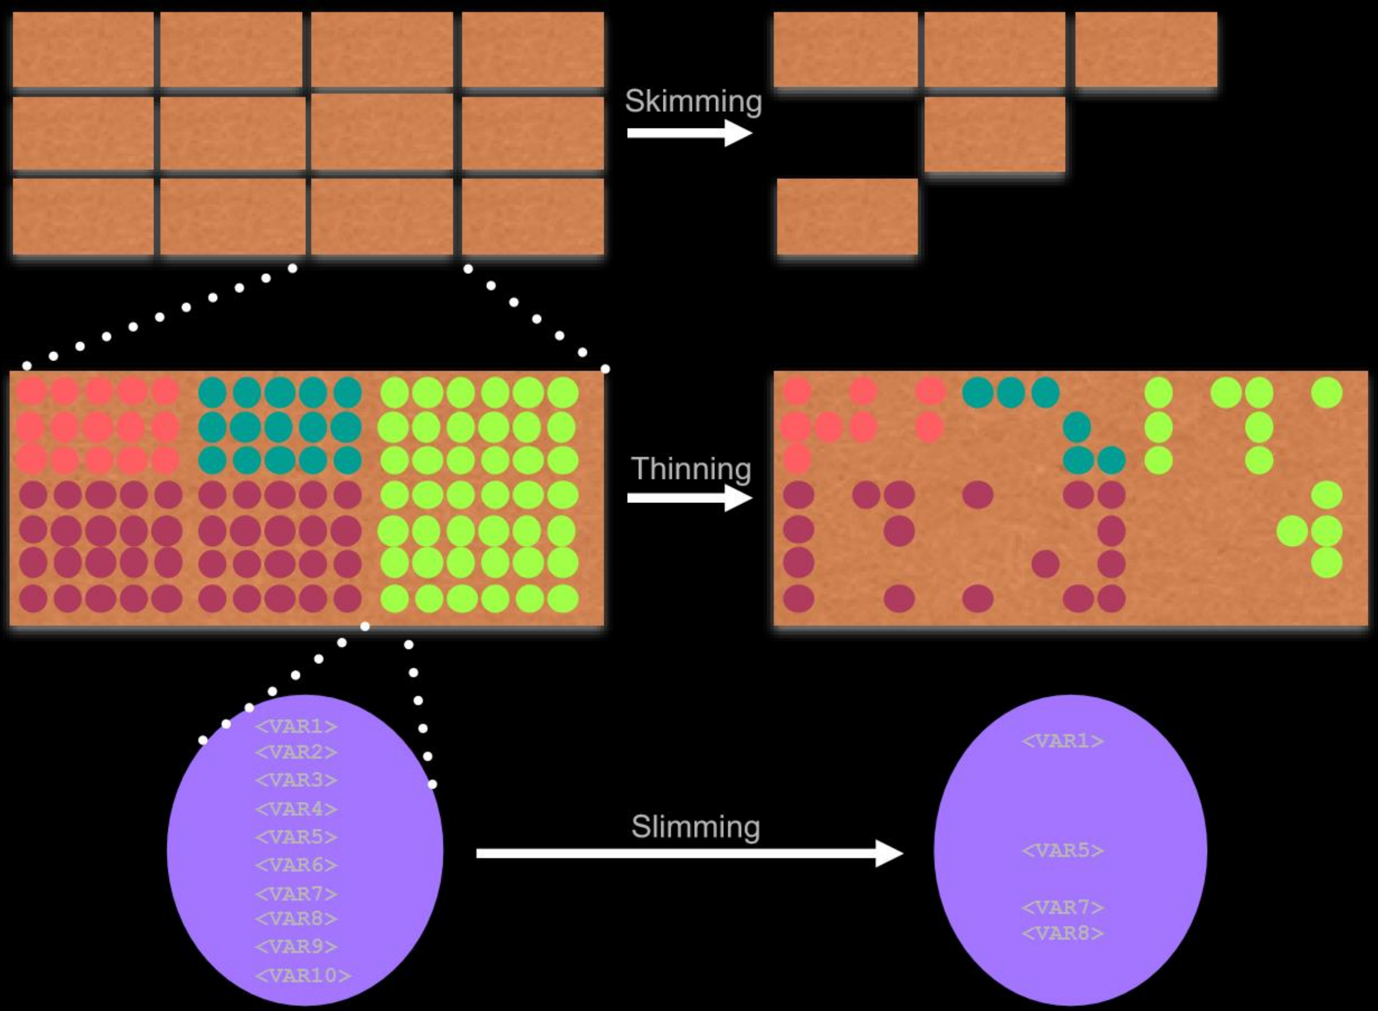

Skimming: removal of whole events based on pre-set criteria

Thinning: removal of whole objects within events based on pre-set criteria

Slimming: removal of variables within objects uniformly across events