Advanced Analysis Software

your

~~Where does ATLAS data come from?~~

Where does ATLAS data come from?

Federica Legger

How do we get from this….

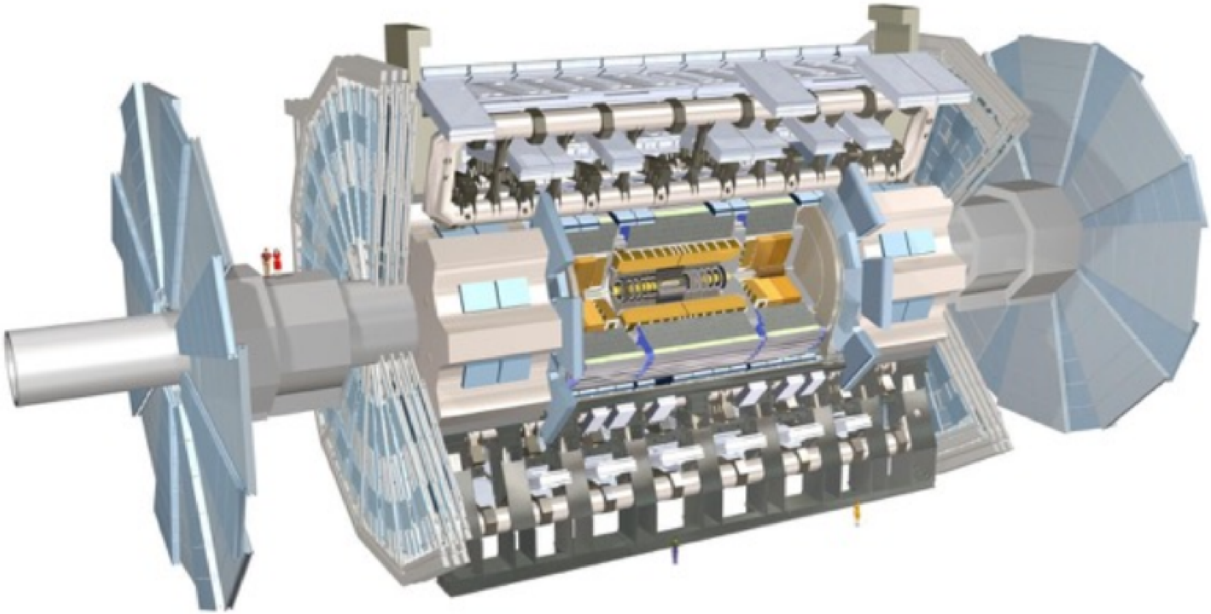

to this…?

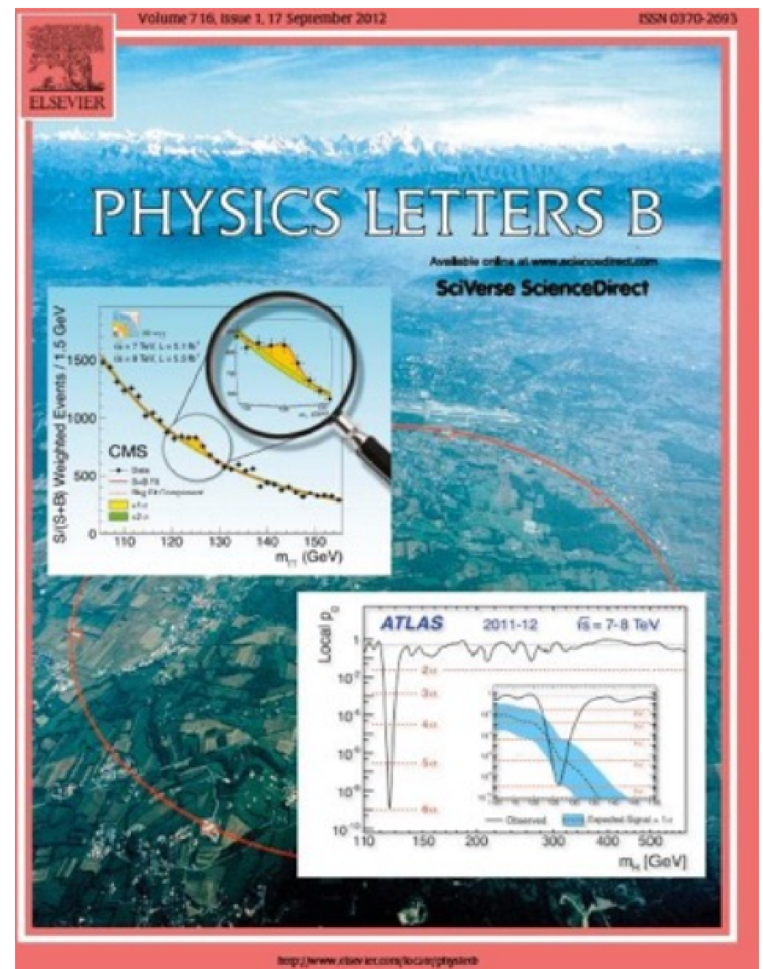

What's in 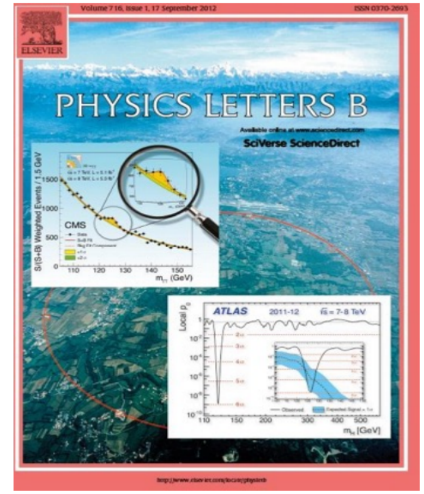 ?

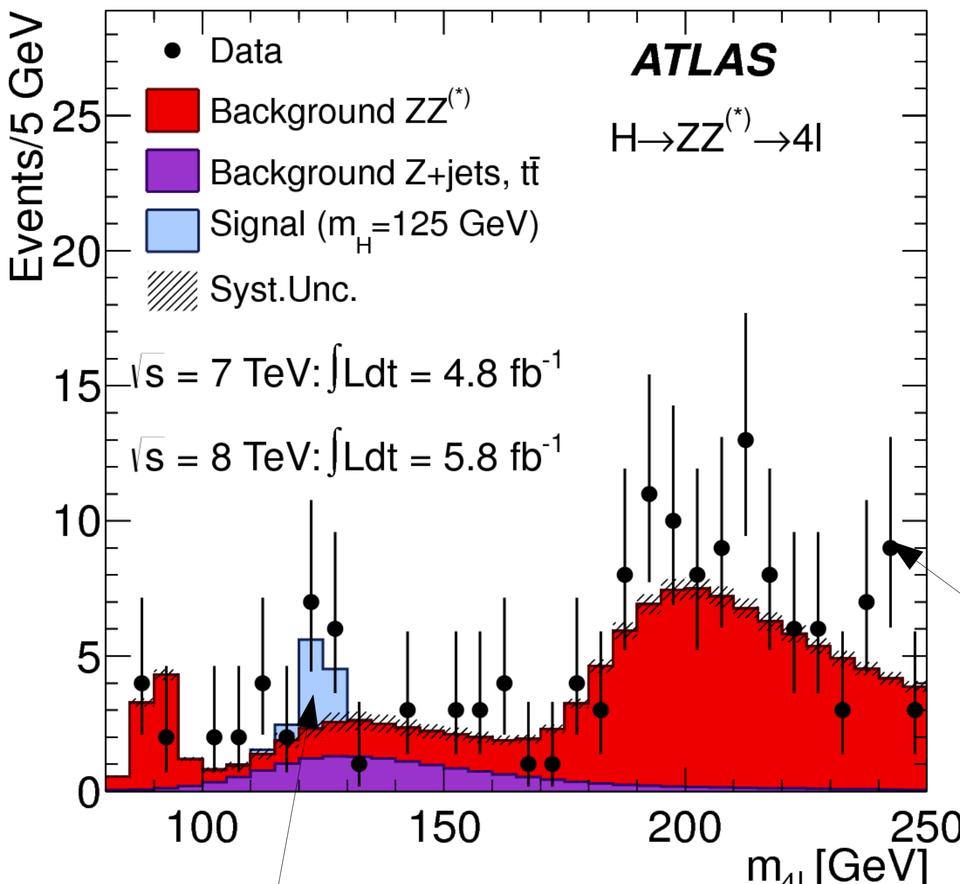

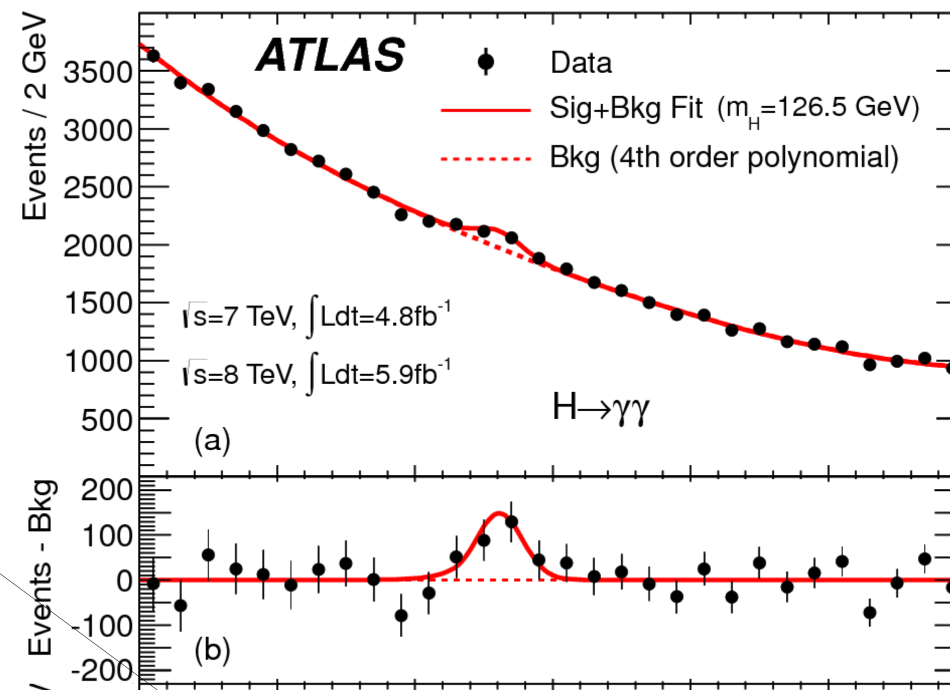

Monte Carlo data

LHC data

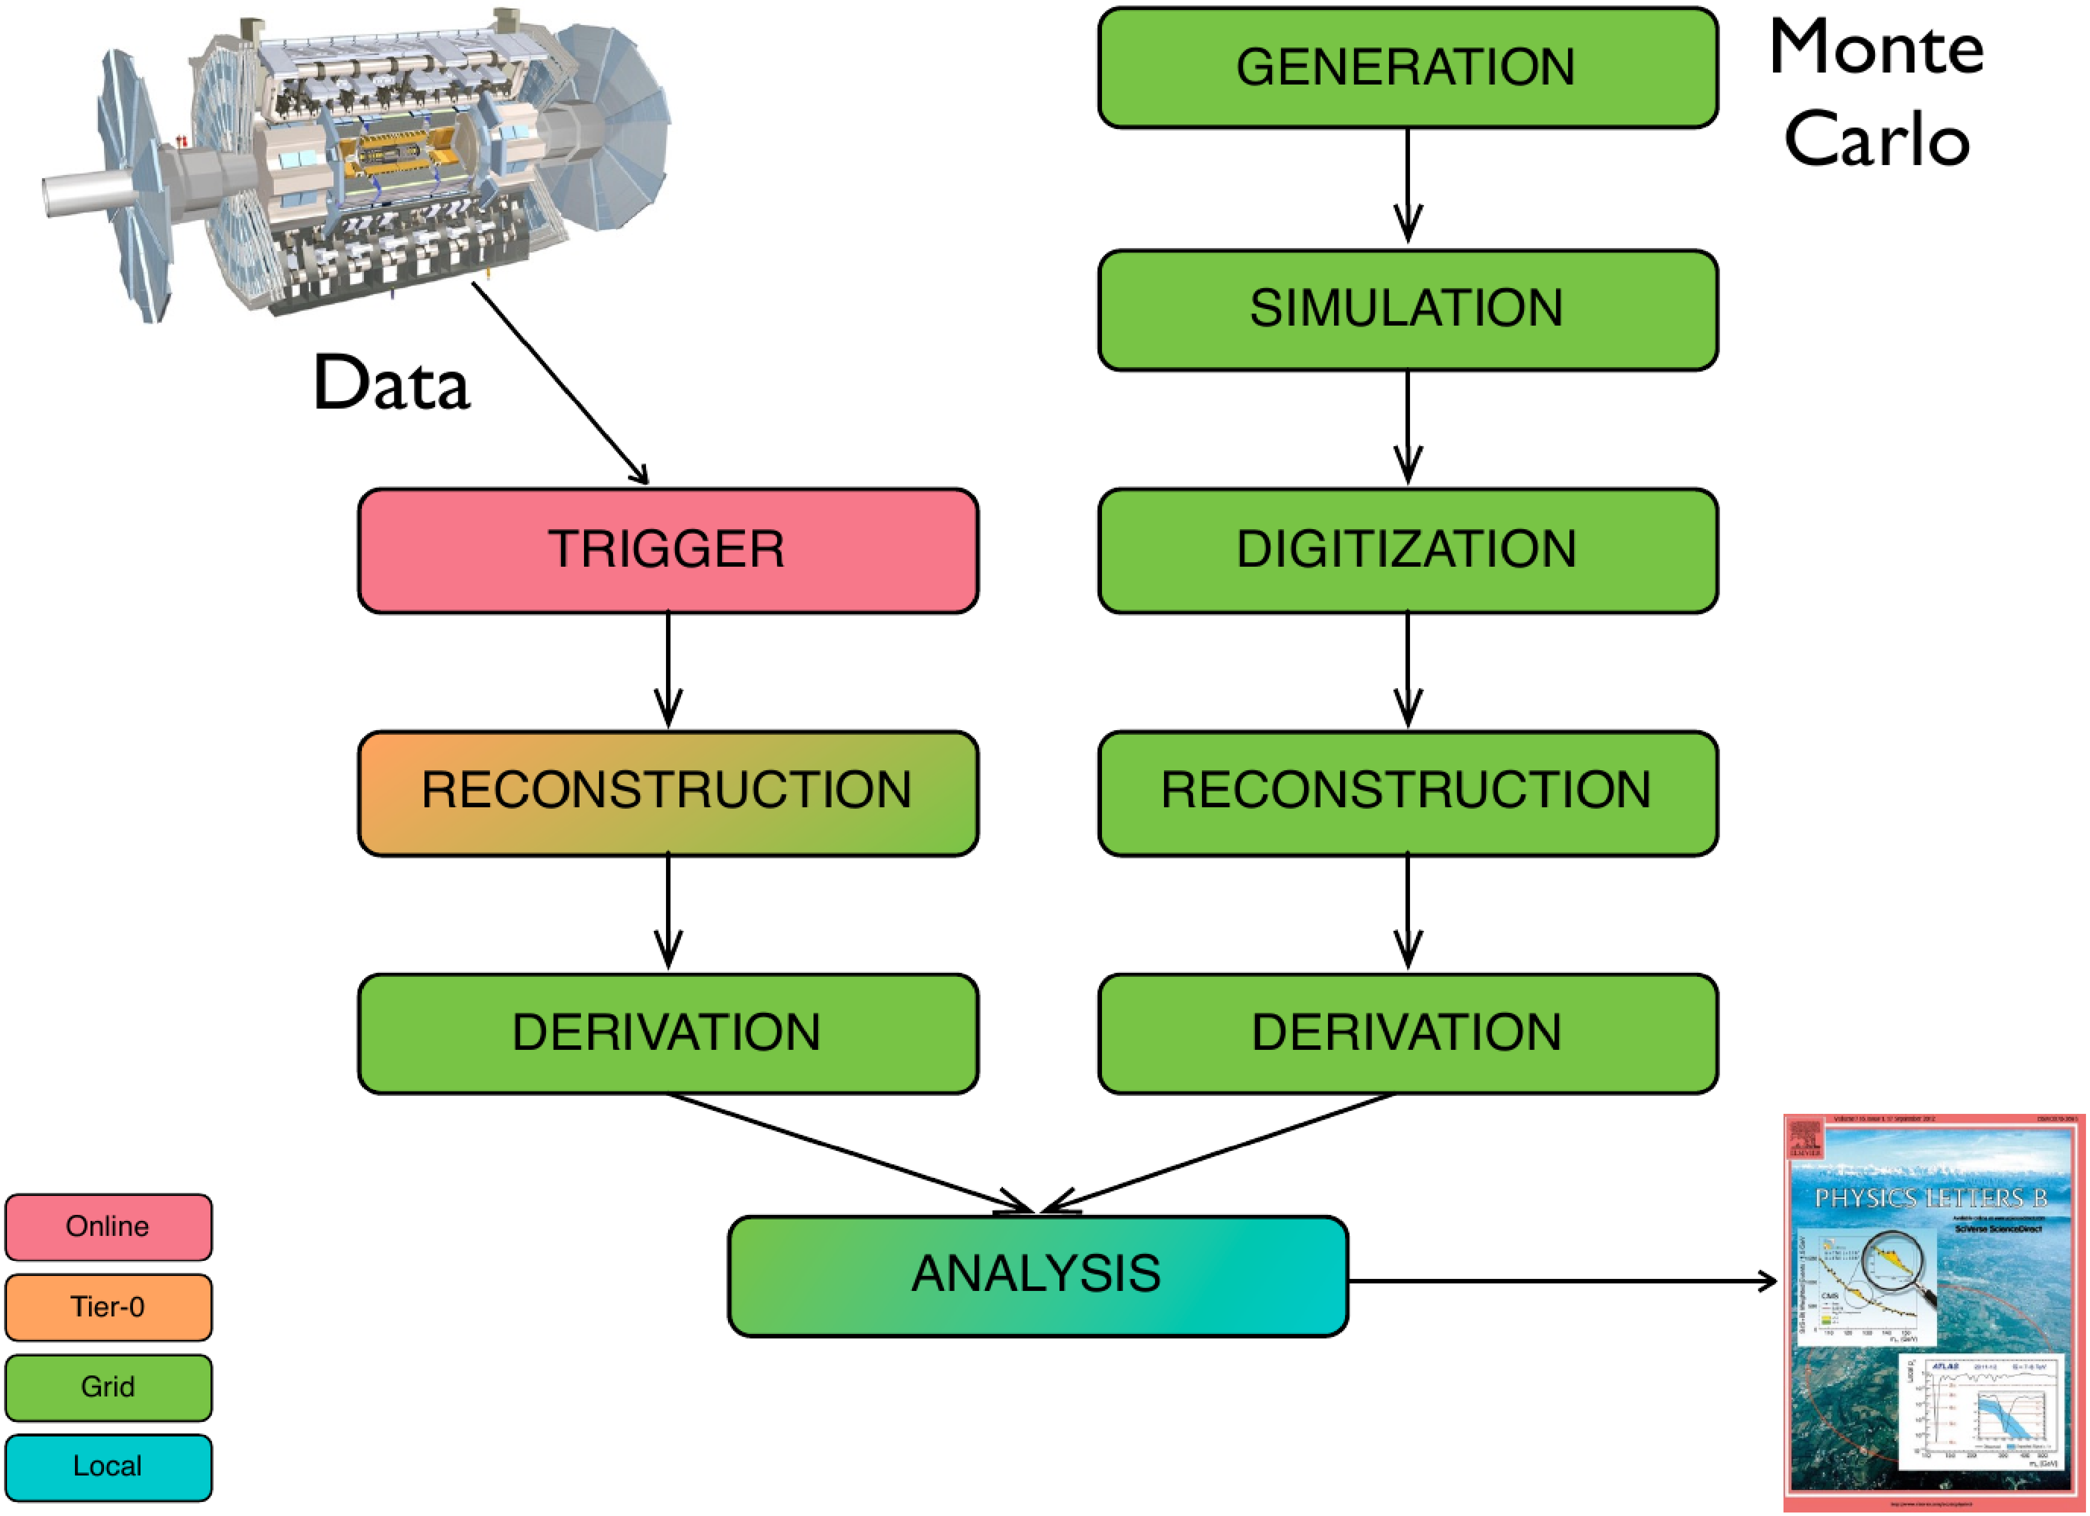

Monte Carlo
GENERATION
SIMULATION
DIGITIZATION
RECONSTRUCTION
DERIVATION
Data
TRIGGER
RECONSTRUCTION
DERIVATION
ANALYSIS
PHYSICS LETTERS B
Online
Tier-0
Grid
Local

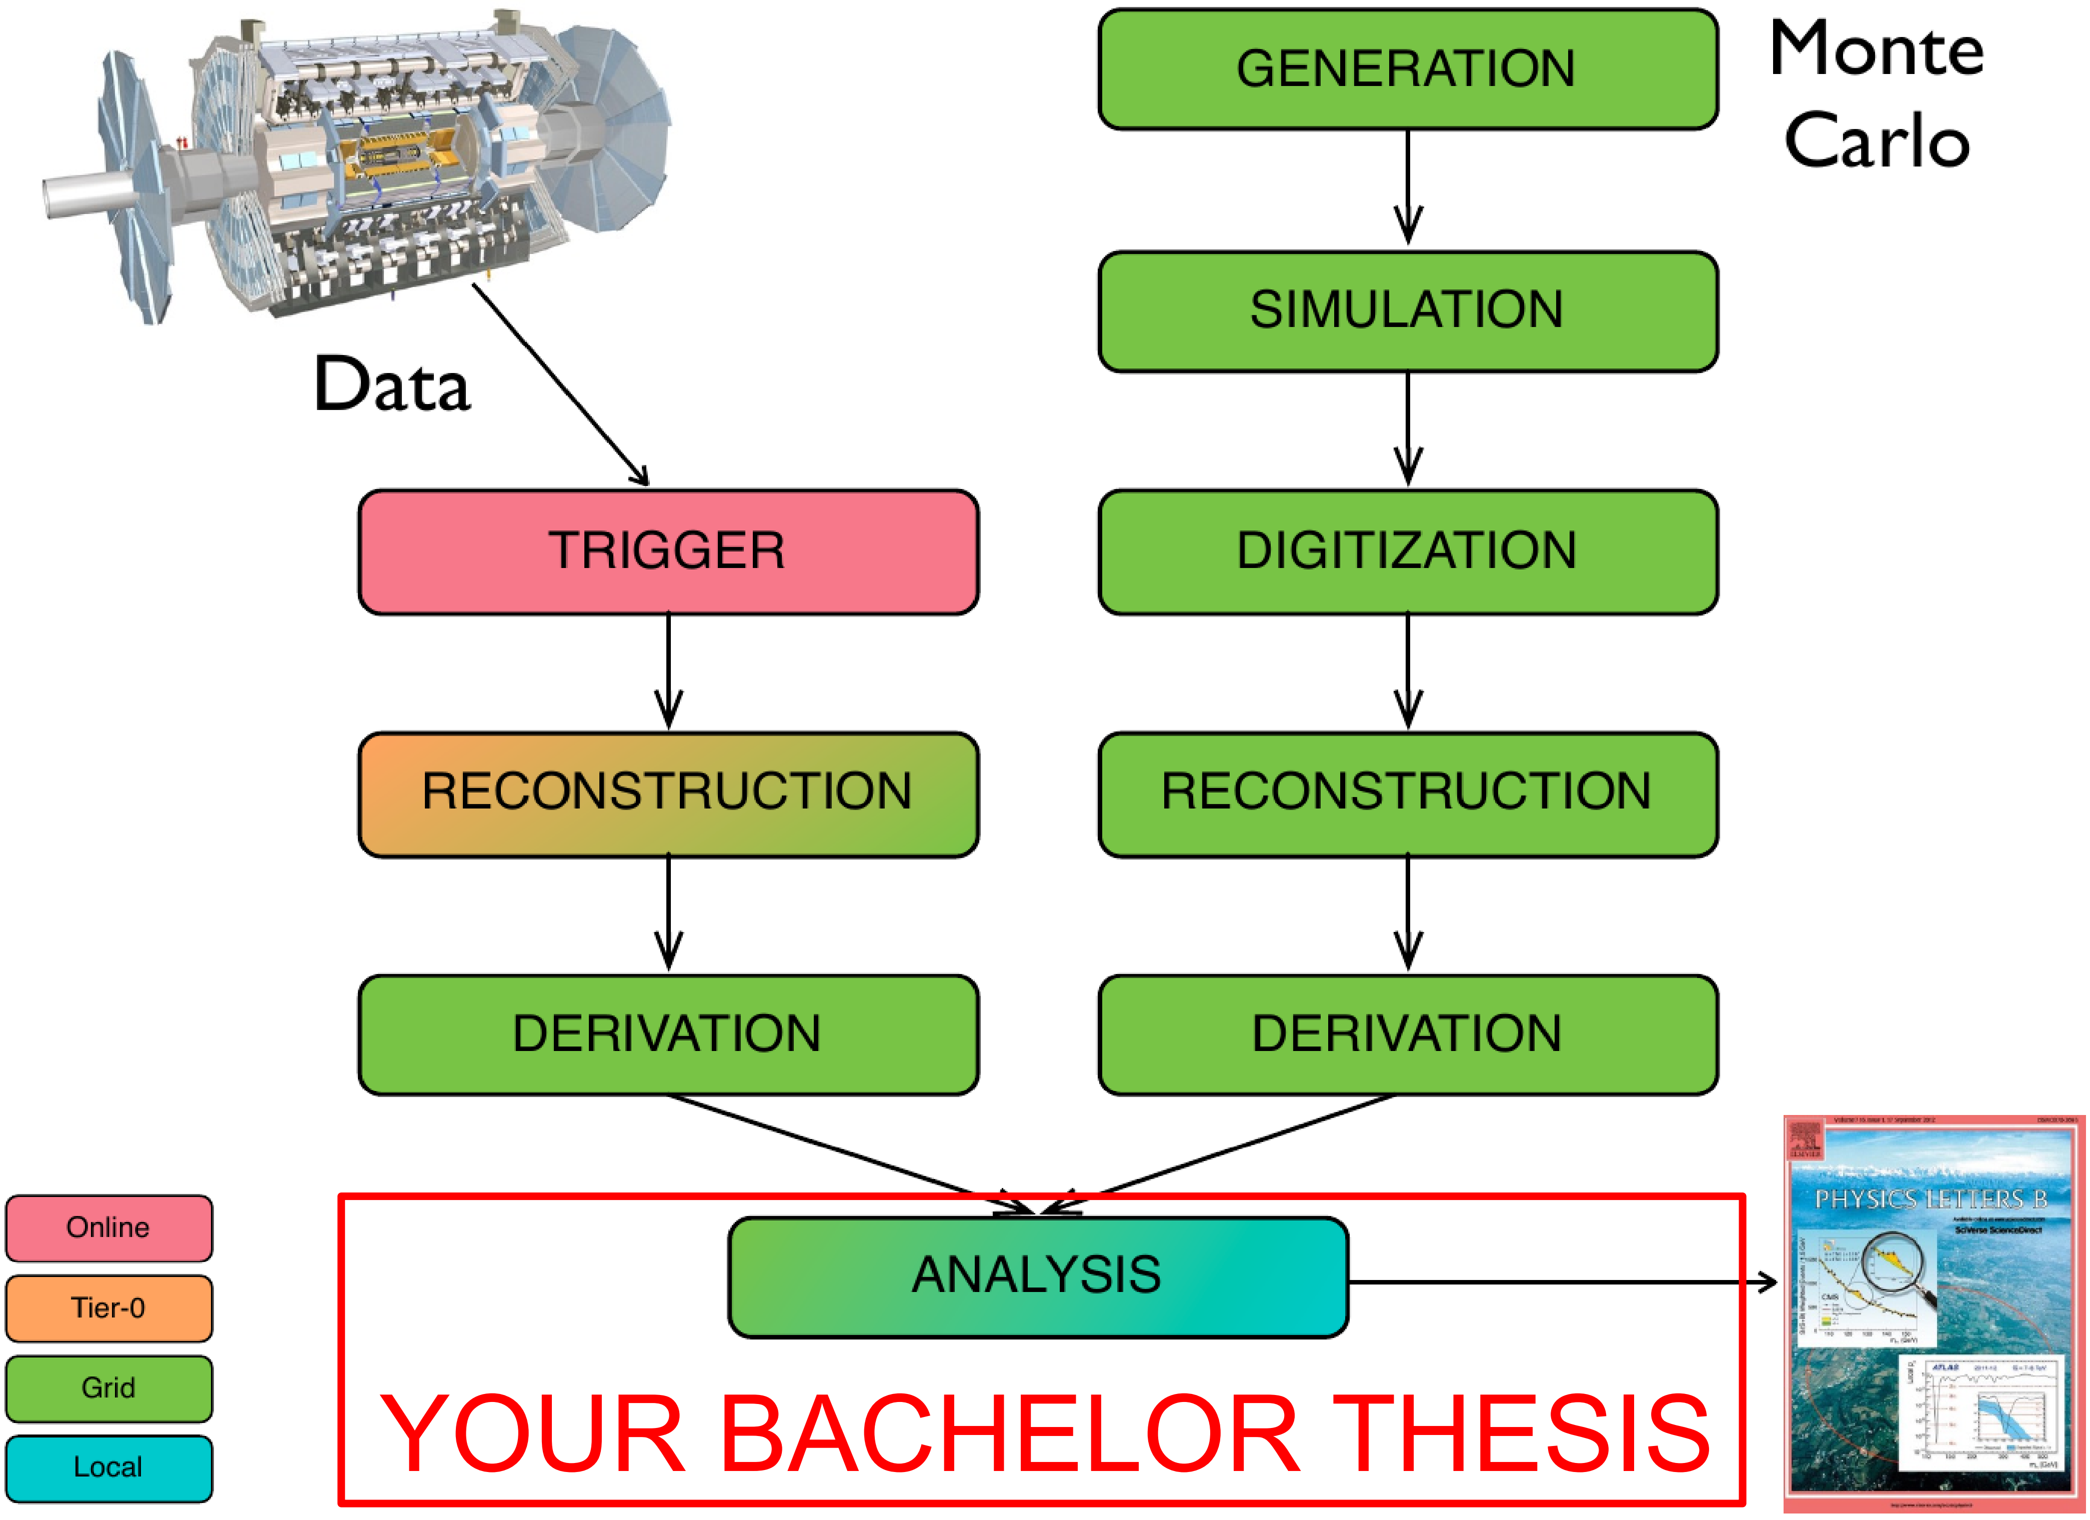

Monte Carlo
GENERATION
SIMULATION
DIGITIZATION
RECONSTRUCTION
DERIVATION
Data
TRIGGER
RECONSTRUCTION
DERIVATION
ANALYSIS
YOUR BACHELOR THESIS
PHYSICS LETTERS B
Online
Tier-0
Grid
Local

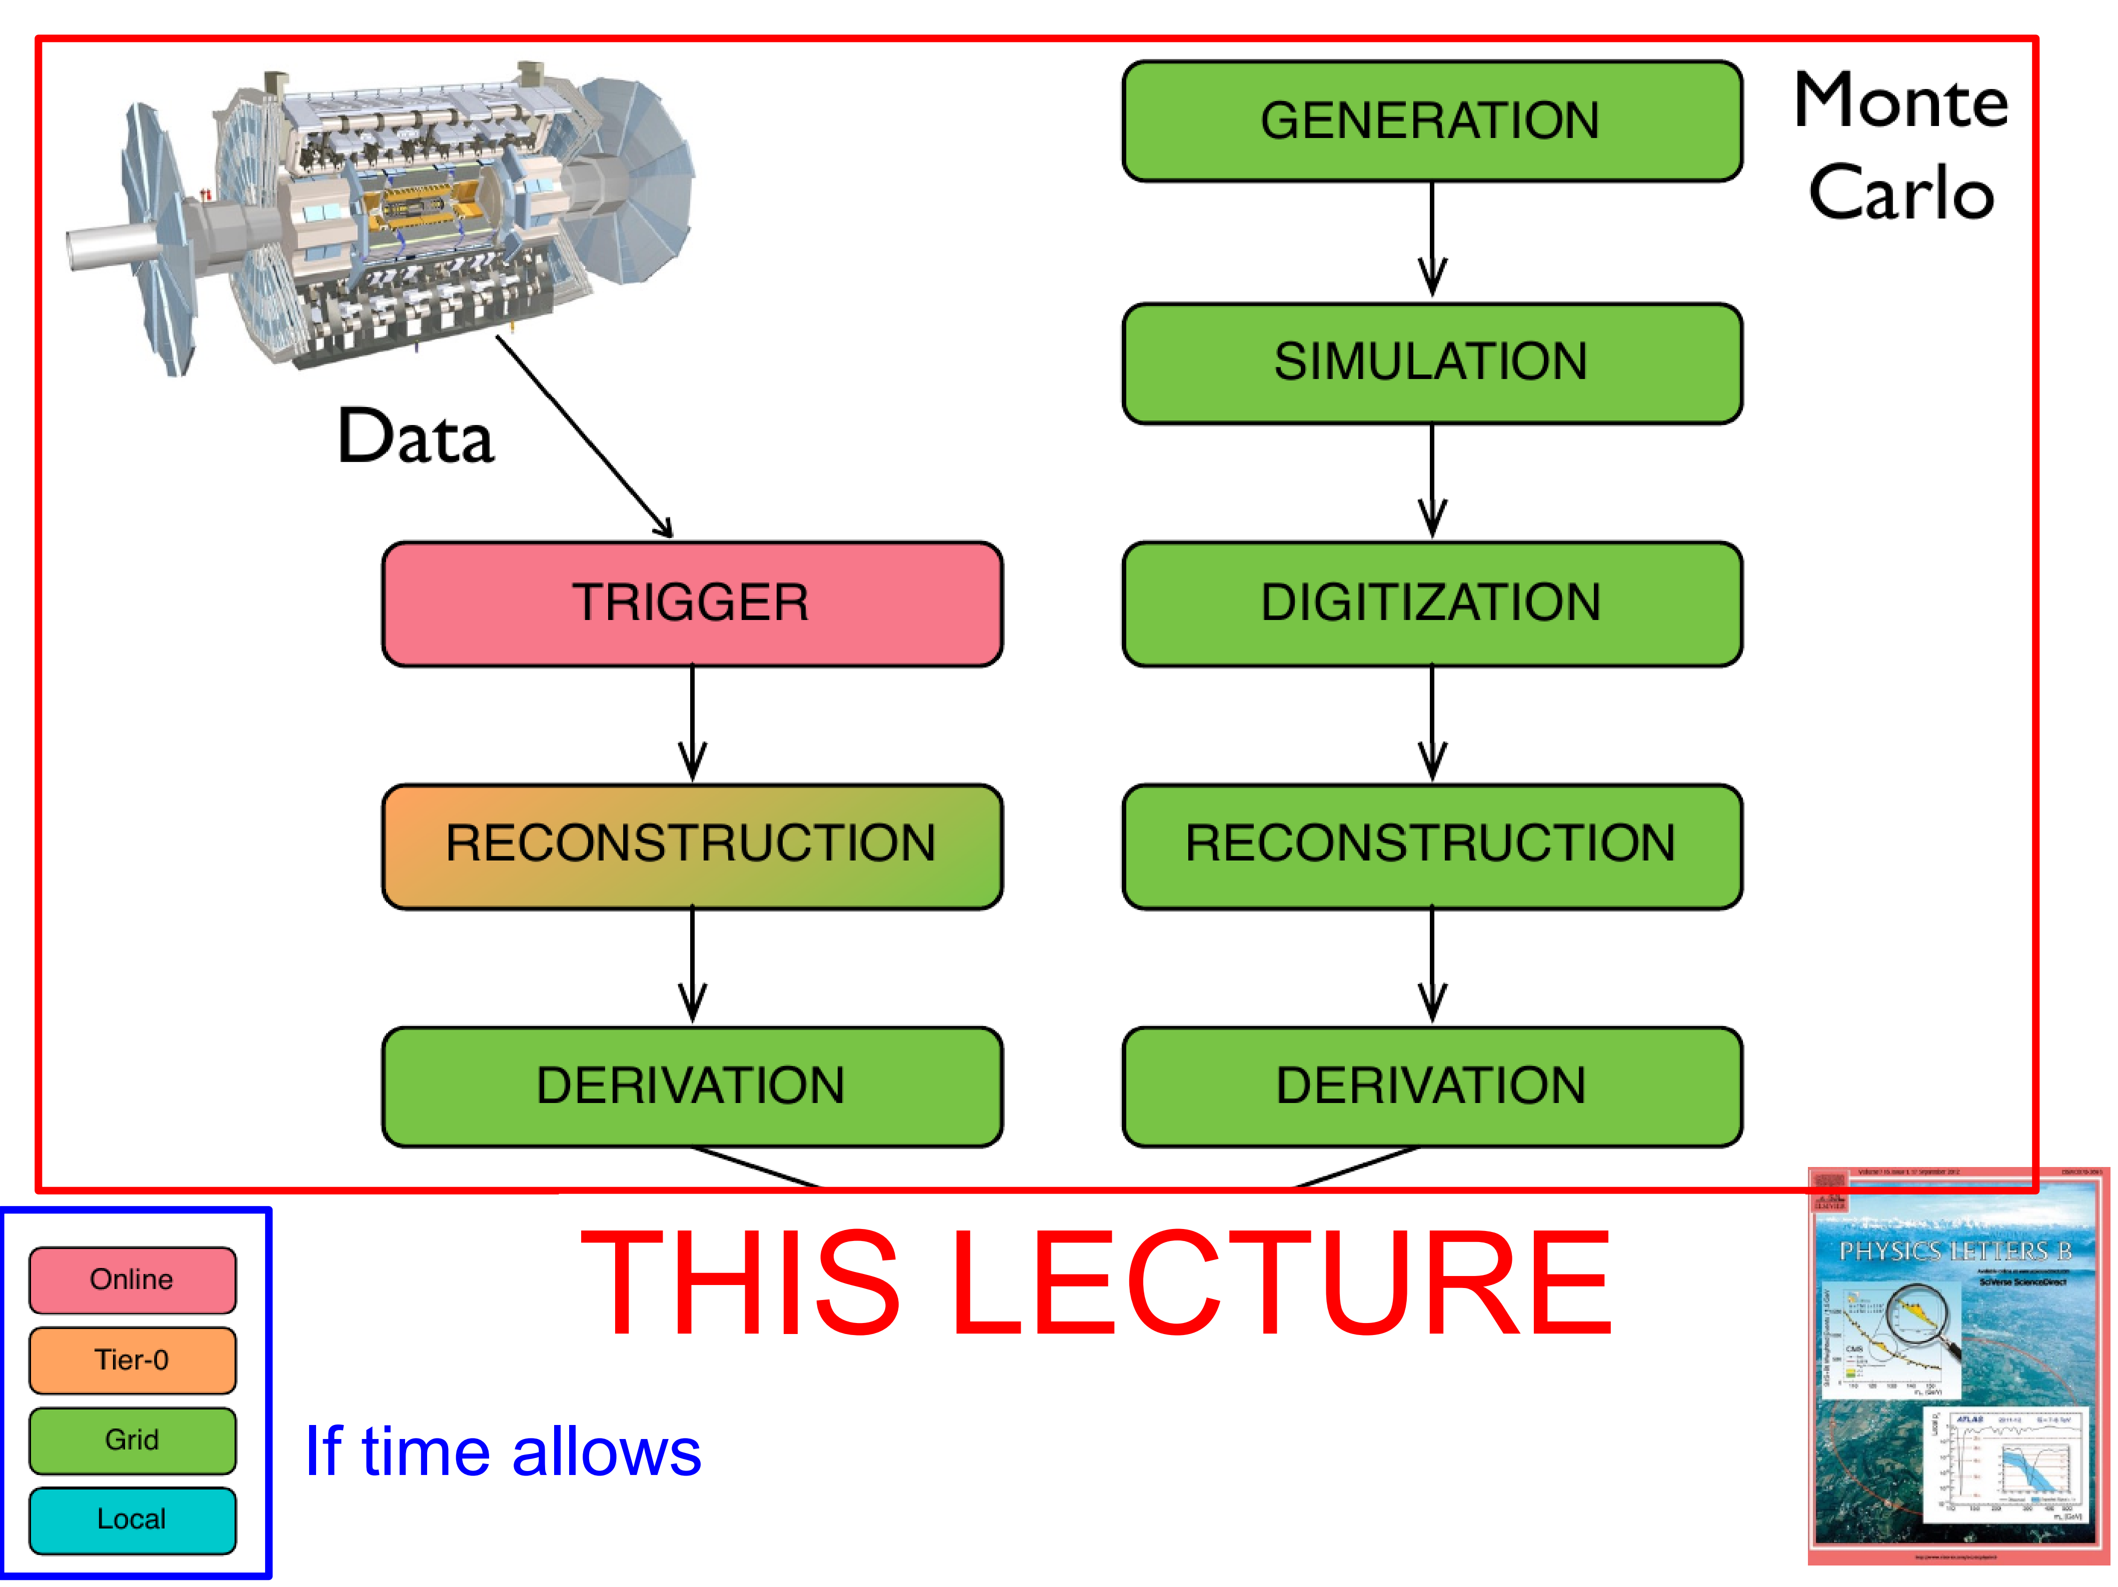

Monte Carlo
Data
GENERATION
SIMULATION
TRIGGER
DIGITIZATION
RECONSTRUCTION
RECONSTRUCTION
DERIVATION
DERIVATION
THIS LECTURE
Online
Tier-0
Grid
Local
If time allows
PHYSICS LETTERS B

Trigger and data preparation

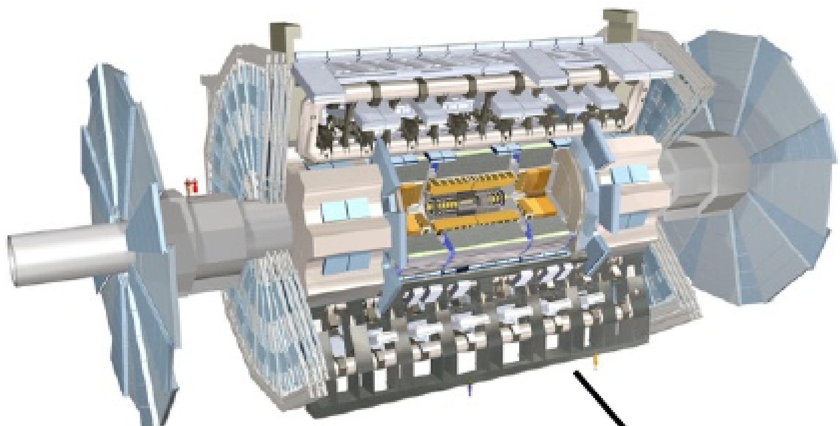

TRIGGER

RECONSTRUCTION

Trigger

- *Trigger*: hardware and software that decides what events get written to disk
  - LHC delivers events at up to 40MHz (1 bunch crossing every 25ns)
  - We can afford to write up to 1KHz
  - This means we can keep one in 40,000 events
- The ATLAS trigger uses a two-step mechanism
  - LEVEL 1: implemented in hardware inside ATLAS - reduces rate from 40MHz to 100KHz
    - Has very little time to decide whether to keep an event or not, so can only use parts of the detector that can be quickly read out → not inner detector, only muons and calorimetry

      Abbreviated as L1
  - HIGH LEVEL: implemented in software running on dedicated machines nearby ATLAS: reduces rate from 100KHz to 1KHz
    - Has a bit more time so can use the full detector information, including the inner detector

      Abbreviated as HLT

Trigger in analysis

- The types of signature that the trigger accepts are encoded in the *trigger menu*
  - Example of a trigger: 2 muons with p_T greater than 40 GeV (**`HLT_2mu40`**)
  - There are several hundred such items in the trigger menu
  - Main trigger page: https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasTrigger
  - Trigger menu: https://twiki.cern.ch/twiki/bin/view/Atlas/TriggerPhysicsMenu
- Some trigger items are *prescaled* (every N events passing kept, rest thrown away) to keep the rate within acceptable limits
  - Prescales must be accounted for in your analysis, especially in the luminosity calculation
  - Ideally you can find a suitable unprescaled trigger for your analysis
- The trigger is not 100% efficient for any signature, and estimating the trigger efficiency is a crucial part of any physics analysis (especially measurement)
  - Often this requires matching reconstructed objects with the equivalent trigger

Some trigger jargon

let's imagine a single muon trigger, HLT_mu40, which selects events where there is a muon reconstructed by the HLT (online with pt > 40 GeV)

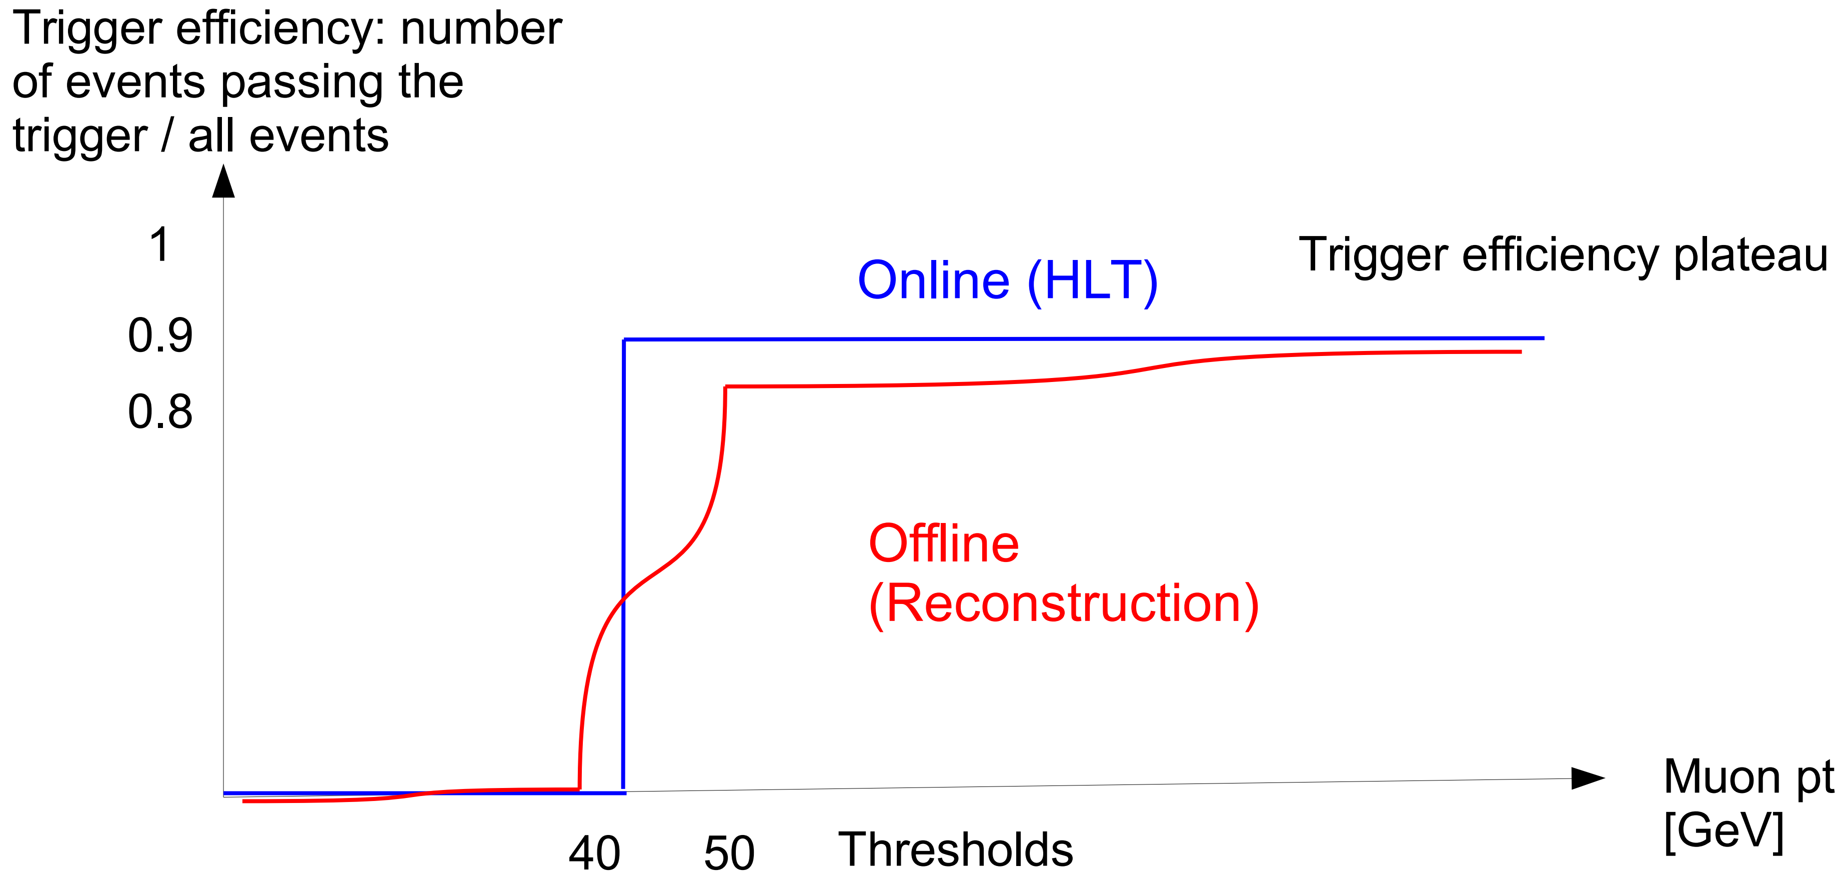

Bulk reconstruction

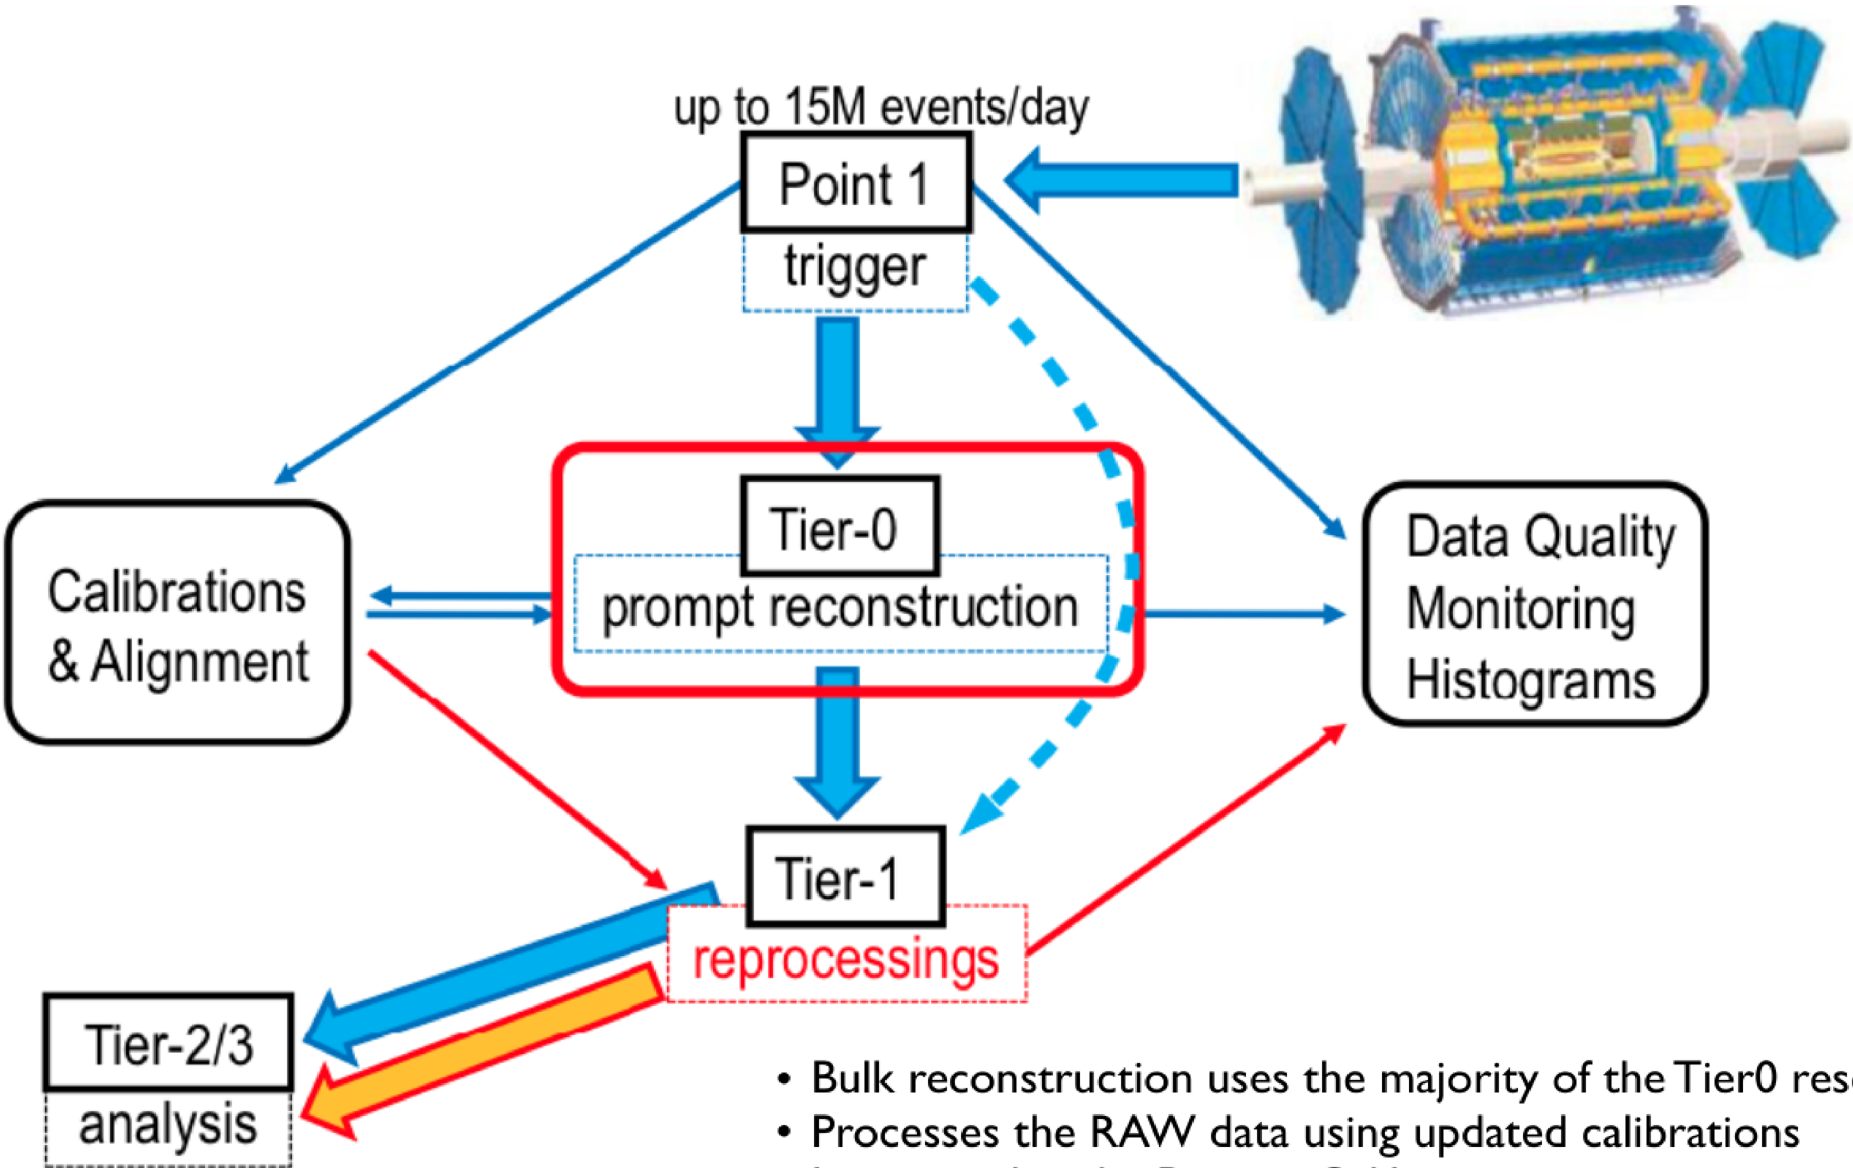

- Bulk reconstruction uses the majority of the Tier0 resources
- Processes the RAW data using updated calibrations determined in the Prompt Calibration
- Produces many outputs used for a variety of purposes, the most important being the AOD

What is Data Preparation?

- Prompt reconstruction of the data:
  - On a Computing Farm at CERN (Tier0 Operations)
  - Using fast calibrations (Beamspot, Calibrations and Conditions)
  - And well-validated software (PROC)
- The Data Preparation group are responsible for which data:
  - How much data? (Luminosity)
  - How good is it? (Data Quality)
  - The Not-data data !!! (Non-collision backgrounds)
  - Data describing the data (Metadata)

Data Quality and Luminosity

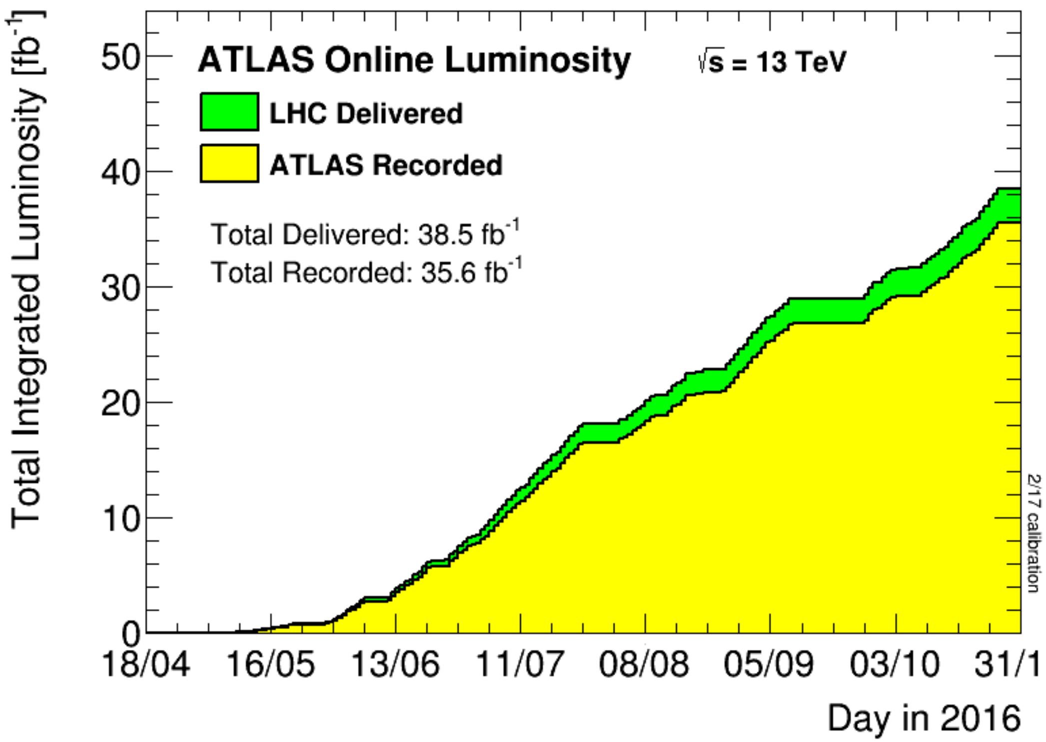

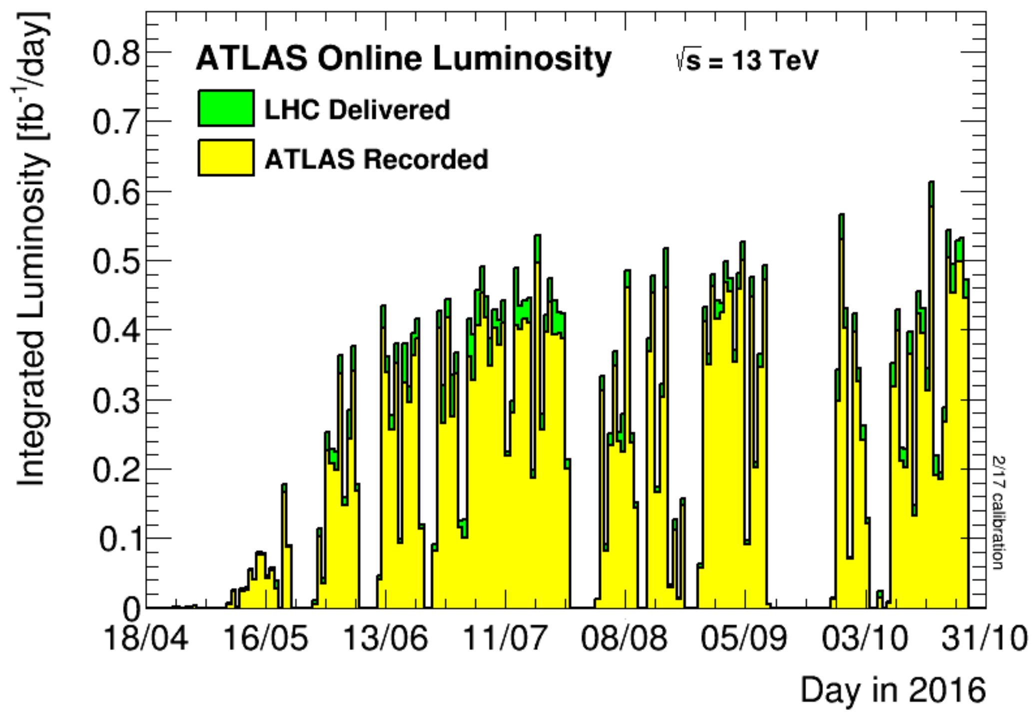

ATLAS p-p run: April-December 2012										
Inner Tracker			Calorimeters		Muon Spectrometer				Magnets	
Pixel	SCT	TRT	LAr	Tile	MDT	RPC	CSC	TGC	Solenoid	Toroid
99.9	99.1	99.8	99.1	99.6	99.6	99.8	100.	99.6	99.8	99.5
All good for physics: 95.5%										
Luminosity weighted relative detector uptime and good quality data delivery during 2012 stable beams in pp collisions at √s=8 TeV between April 4th and December 6th (in %) – corresponding to 21.3 fb⁻¹ of recorded data.										

How data is divided up in ATLAS

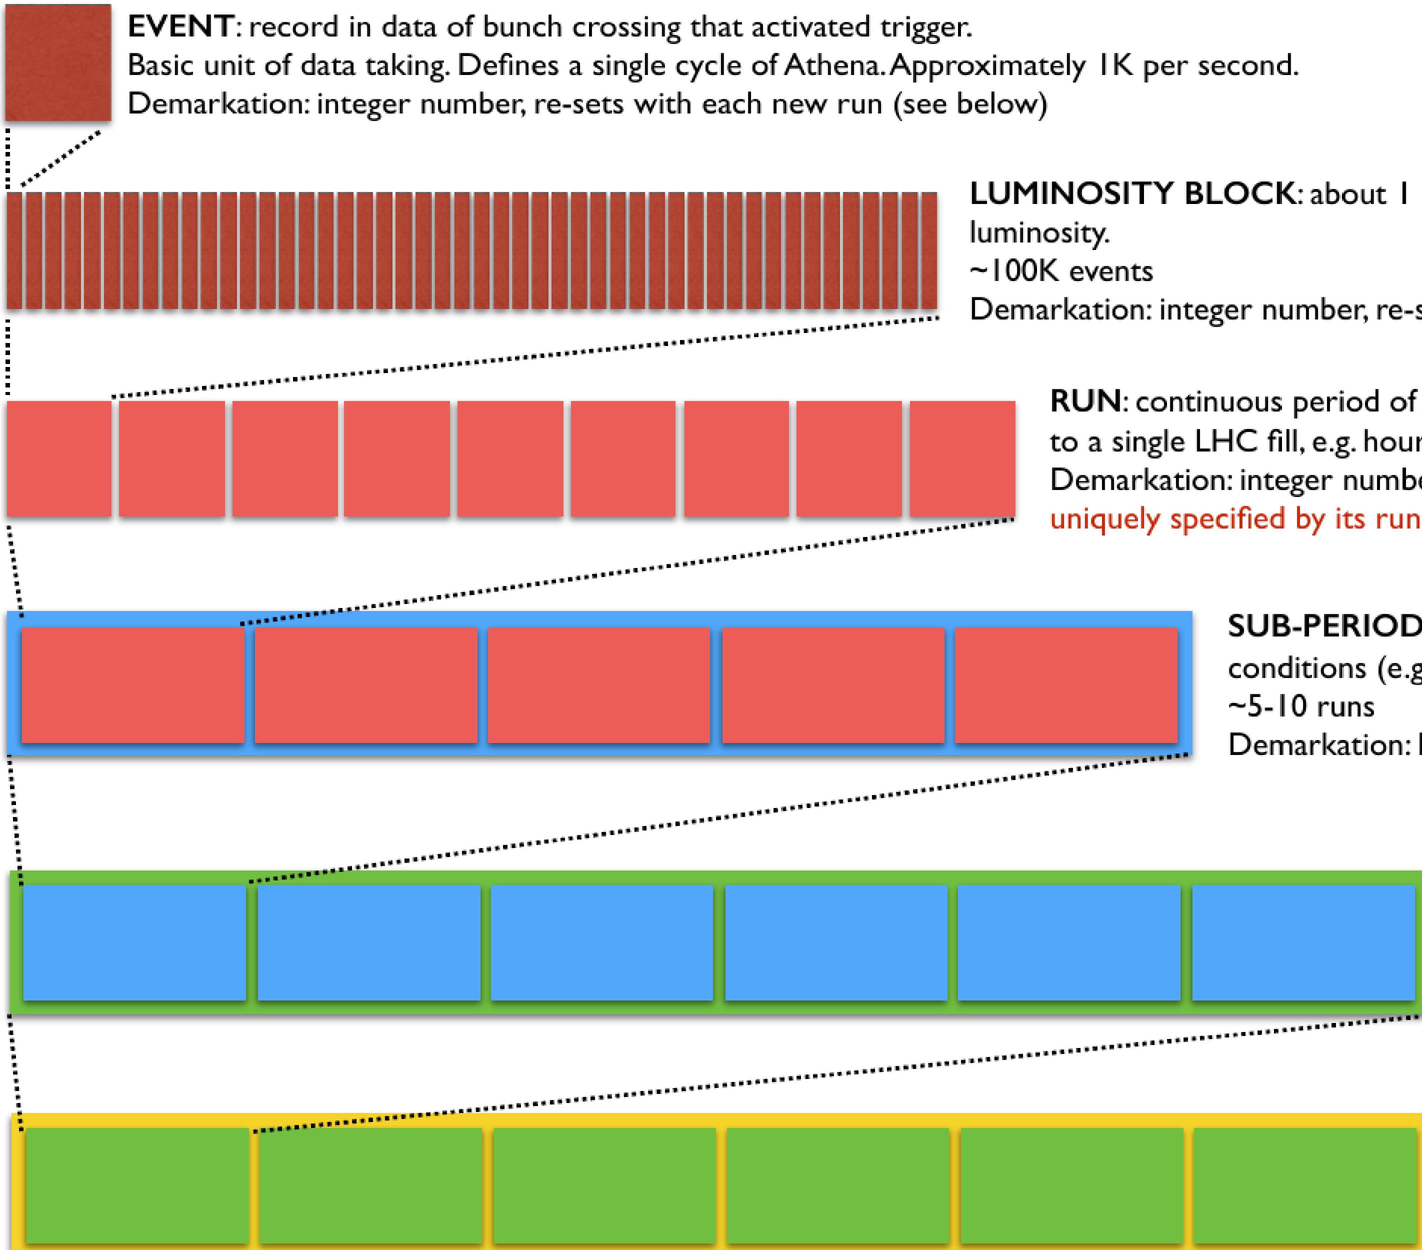

EVENT: record in data of bunch crossing that activated trigger.
Basic unit of data taking. Defines a single cycle of Athena. Approximately 1K per second.
Demarkation: integer number, re-sets with each new run (see below)

LUMINOSITY BLOCK: about 1 minute of data taking. Unit of known luminosity.
~100K events
Demarkation: integer number, re-sets with each new run

RUN: continuous period of ATLAS data recording. Usually corresponds to a single LHC fill, e.g. hours of data taking. ~1000 luminosity blocks.
Demarkation: integer number. Never resets → each event can be uniquely specified by its run and event number

SUB-PERIOD: group of runs taken with very similar conditions (e.g. trigger, LHC).
~5-10 runs
Demarkation: letter and number (e.g. E5)

PERIOD: group of sub-periods taken with similar conditions, or during some period (e.g. increments after technical stop)
~5-10 sub periods
Demarkation: letter (e.g. E); resets each year

LHC RUN: long period of data taking, between shutdowns. 2009-2012 was Run I, we are now in Run II, and Run III will start in 2021. Typically 20-30 periods per LHC run.

Good runs lists

- Remember those luminosity blocks? ~ 1 minute of data taking
  - Several hundred lumi blocks per run; one or several lumi blocks per file
  - Some lumi block ranges will have *bad data quality* and should not be used for analysis

- The data quality flags are set in the data quality monitoring
- They are encoded in a database
- Different versions of the data quality flags get a different name, e.g. DetStatus-v61

Lumi blocks and analysis

- When you do your analysis you need to filter out the bad lumi blocks
- This is done by means of a *good runs list (GRL)*
  - XML file listing each run and the range of good lumi blocks in each: read in by the analysis code
- The luminosity "seen" by your analysis depends on three things:
  - the luminosity blocks processed (from the GRL)
  - the trigger prescales applied
  - the level-1 trigger live-time (fraction of delivered luminosity that ATLAS recorded)

What a GRL looks like inside


```
<LumiRangeCollection>
<NamedLumiRange>
<Name>PHYS_StandardGRL_All_Good</Name>
<Version>2.1</Version>
<Metadata Name="ARQEquivalentQuery">find run data12_8TeV.periodAllYear and dq PHYS_StandardGRL_All_Good DEFECTS#DetStatus-v61-pro14-02 g</Metadata>
<Metadata Name="Query">Period: data12_8TeV.AllYear; Defect: PHYS_StandardGRL_All_Good; Defect tag: DetStatus-v61-pro14-02; Ignoring None</Metadata>
<Metadata Name="RQTSVNVersion">DQDefects-00-01-00</Metadata>
<Metadata
Name="RunList">200842,200863,200913,200926,200965,200967,200982,200987,201006,201052,201113,201120,201138,201190,201191,201257,201269,201280,201289,201489,201494,201555,2
01556,202660,202668,202712,202740,202798,202965,202991,203027,203169,203191,203195,203228,203256,203258,203277,203335,203336,203353,203432,203454,203456,203523,203524,20
602,203605,203636,203680,203719,203739,203745,203760,203779,203792,203875,203876,203934,204025,204026,204071,204073,204134,204153,204158,204240,204265,204416,204442,2044
4,204564,204633,204668,204726,204763,204769,204772,204796,204853,204857,204910,204932,204954,204955,204976,205010,205016,205017,205055,205071,205112,205113,206368,206369
206409,206497,206573,206614,206955,206962,206971,207044,207046,207221,207262,207304,207306,207332,207397,207447,207490,207528,207531,207532,207582,207589,207620,207664,20
7696,207749,207772,207800,207809,207845,207864,207865,207931,207934,207975,207982,208123,208126,208179,208184,208189,208258,208261,208354,208485,208631,208642,208662,208
05,208717,208720,208780,208781,208811,208870,208930,208931,208970,208982,209024,209025,209074,209084,209109,209161,209183,209214,209254,209265,209269,209353,209381,20955
,209580,209608,209628,209629,209736,209776,209787,209812,209864,209866,209899,209980,209995,210302,210308,211620,211670,211697,211772,211787,211867,211902,211937,212034,2
12103,212142,212144,212172,212199,212272,212619,212663,212687,212721,212742,212809,212815,212858,212967,212993,213039,213079,213092,213130,213155,213157,213204,213250,21
359,213431,213479,213486,213539,213627,213640,213684,213695,213702,213754,213796,213816,213819,213900,213951,213964,213968,214021,214086,214160,214176,214216,214388,2143
0,214494,214523,214544,214553,214618,214651,214680,214714,214721,214758,214777,215027,215061,215063,215091,215414,215433,215456,215464,215473,215541,215571,215589,215643
/Metadata>
<LumiBlockCollection>
<Run>200842</Run>
<LBRange Start="27" End="36"/>
<LBRange Start="43" End="47"/>
<LBRange Start="49" End="85"/>
</LumiBlockCollection>
<LumiBlockCollection>
<Run>200863</Run>
<LBRange Start="57" End="74"/>
<LBRange Start="76" End="83"/>
<LBRange Start="86" End="156"/>
<LBRange Start="158" End="165"/>
<LBRange Start="167" End="191"/>
…….
…….
<LumiBlockCollection>
<Run>215643</Run>
<LBRange Start="400" End="460"/>
<LBRange Start="462" End="479"/>
<LBRange Start="481" End="493"/>
<LBRange Start="505" End="509"/>
<LBRange Start="511" End="517"/>
<LBRange Start="522" End="567"/>
</LumiBlockCollection>
</NamedLumiRange>
</LumiRangeCollection>
```

Reprocessing

Reprocessing at Tier-1s

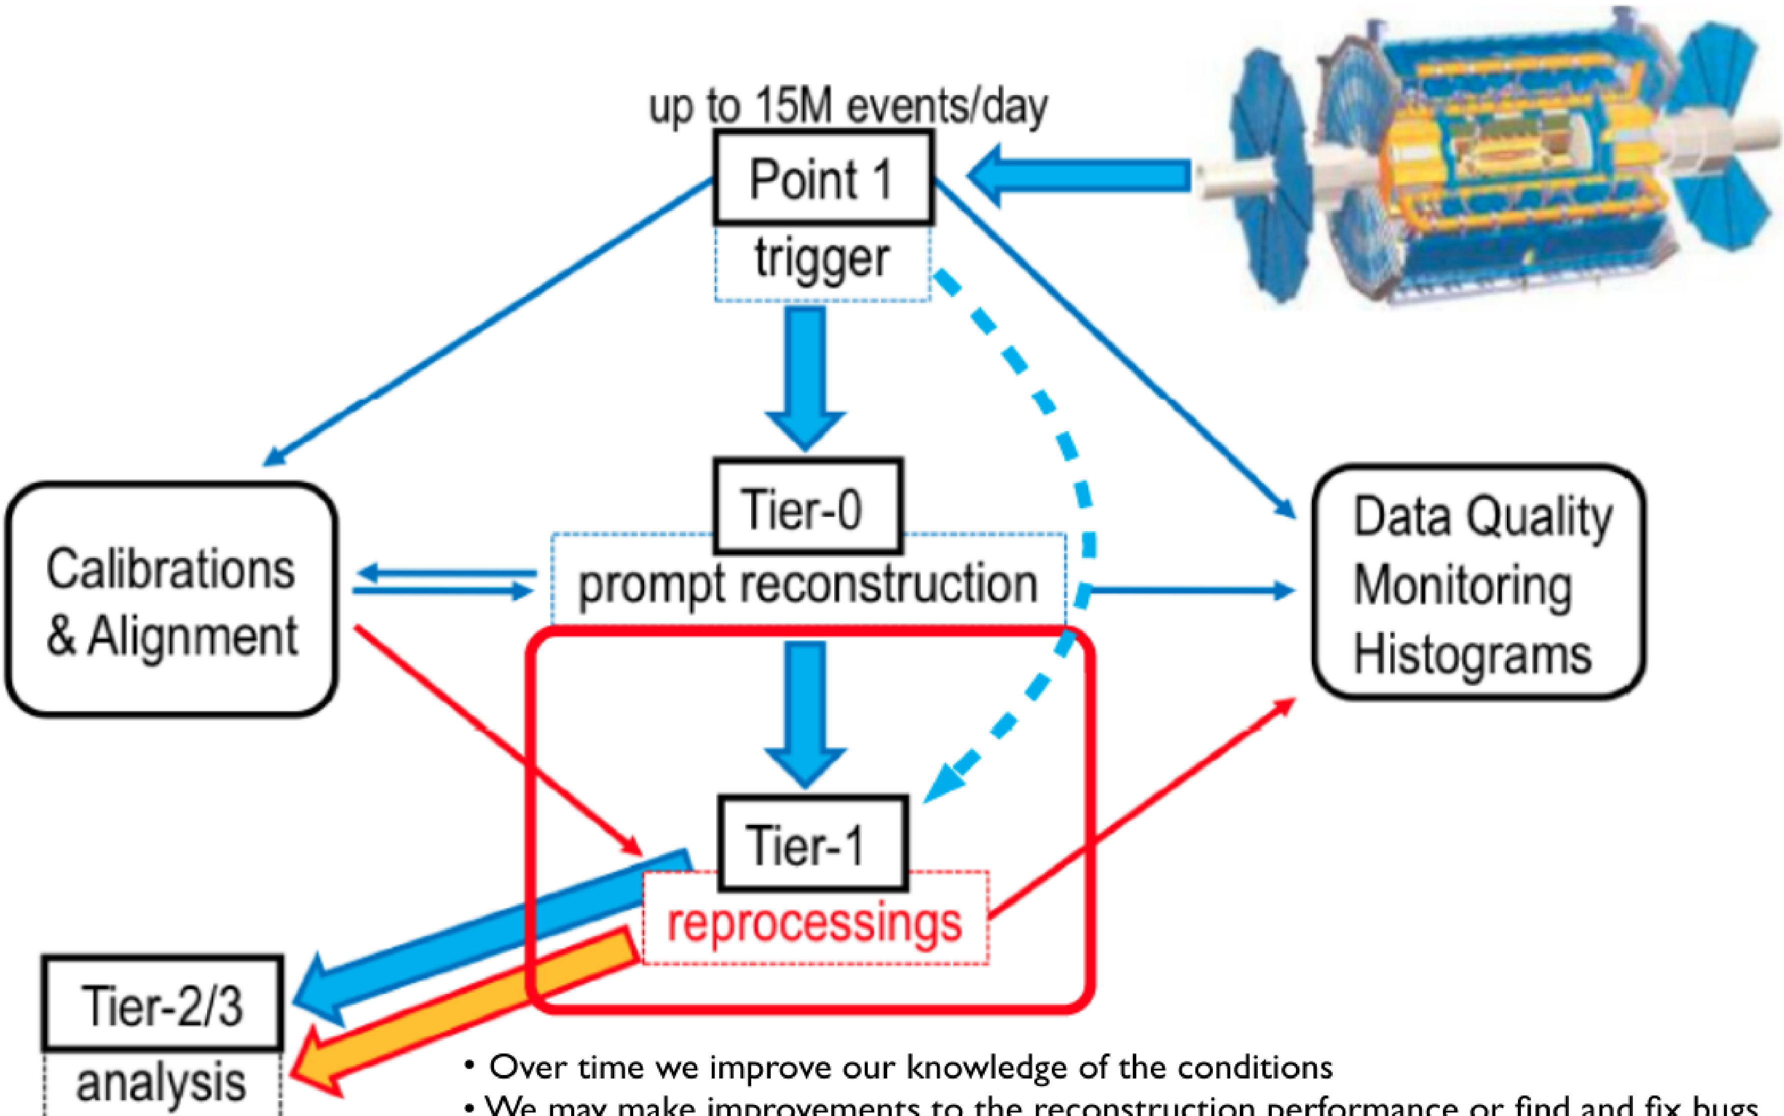

- Over time we improve our knowledge of the conditions
- We may make improvements to the reconstruction performance or find and fix bugs
- After a time, we reprocess the existing raw data with the improved software/conditions
- This has to be done on the Grid since the Tier-0 will be busy with prompt data processing

Monte Carlo Production

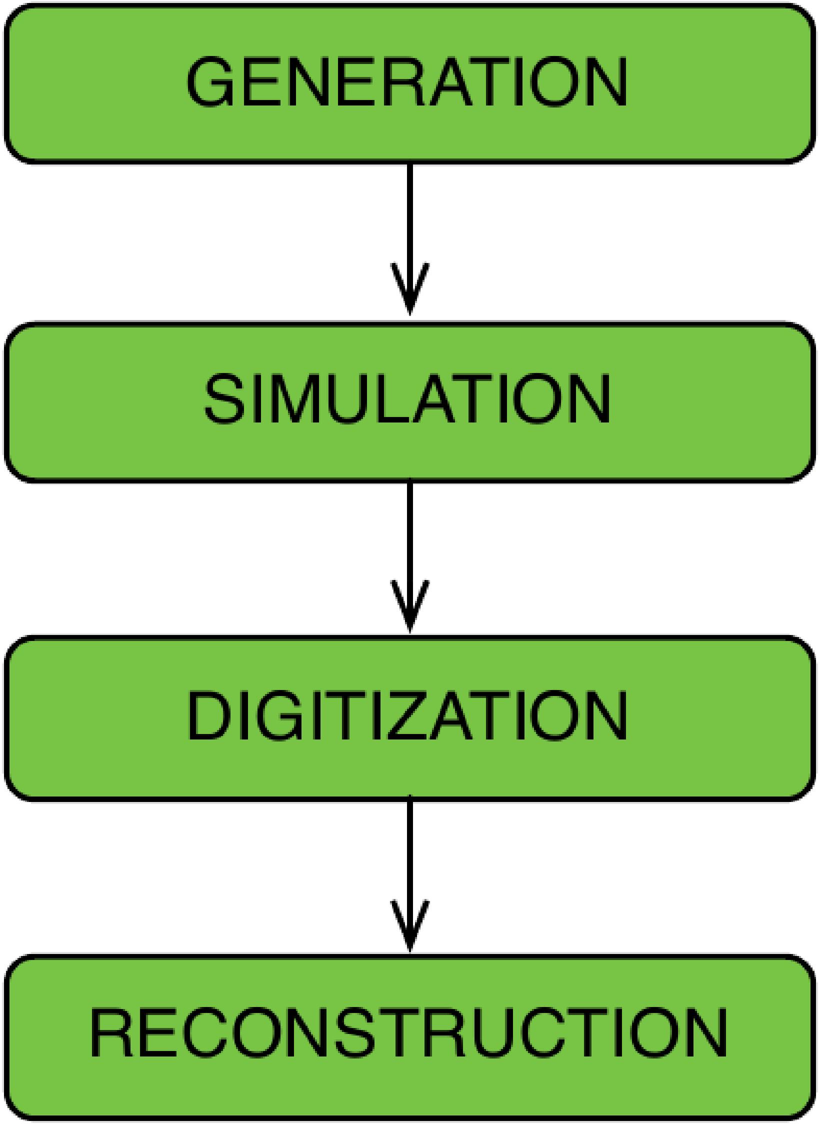

Why do we need Monte Carlo?

- Testing whether we understand the performance of the detector
- Calculating reconstruction efficiencies (usually for validating data-driven methods)
- Modelling the expected background under a process of interest
- Modelling the process of interest
- Training multivariate classifiers (neural networks, boosted decision trees etc)
- Setting systematic uncertainties
- etc etc etc

Monte Carlo production

- All official Monte Carlo production is done on the Grid
- Several-step process
  - *Event generation*: simulation of the interaction between the quarks and gluons in the colliding protons, the subsequent parton showering and hadronization and decays into stable particles
  - *Detector simulation:* calculation of how the particles from the generator interact with the detector material; how they shower into secondaries; how much energy they deposit in each sensitive element
  - *Digitisation*: turning the simulated energy deposits into a detector response that "looks" like the raw data from the real detector
  - After this step, the process is the same as for real data
- The analysis data for MC and real data looks the same, *except* that in MC the original generated events (the "truth") are available as well as the reconstructed objects
- Extra low momentum events must be injected into the chain to simulate the presence of multiple proton collisions ("pile-up") in a given LHC bunch crossing
  - Complication: the average number of collisions per bunch crossing is a function of the LHC parameters, and we typically do not know this when we start the Monte Carlo production
  - Monte Carlo events must be weighted in analysis to account for discrepancies between the real and simulated pile-up.

The ATLAS Analysis Model and the xAOD

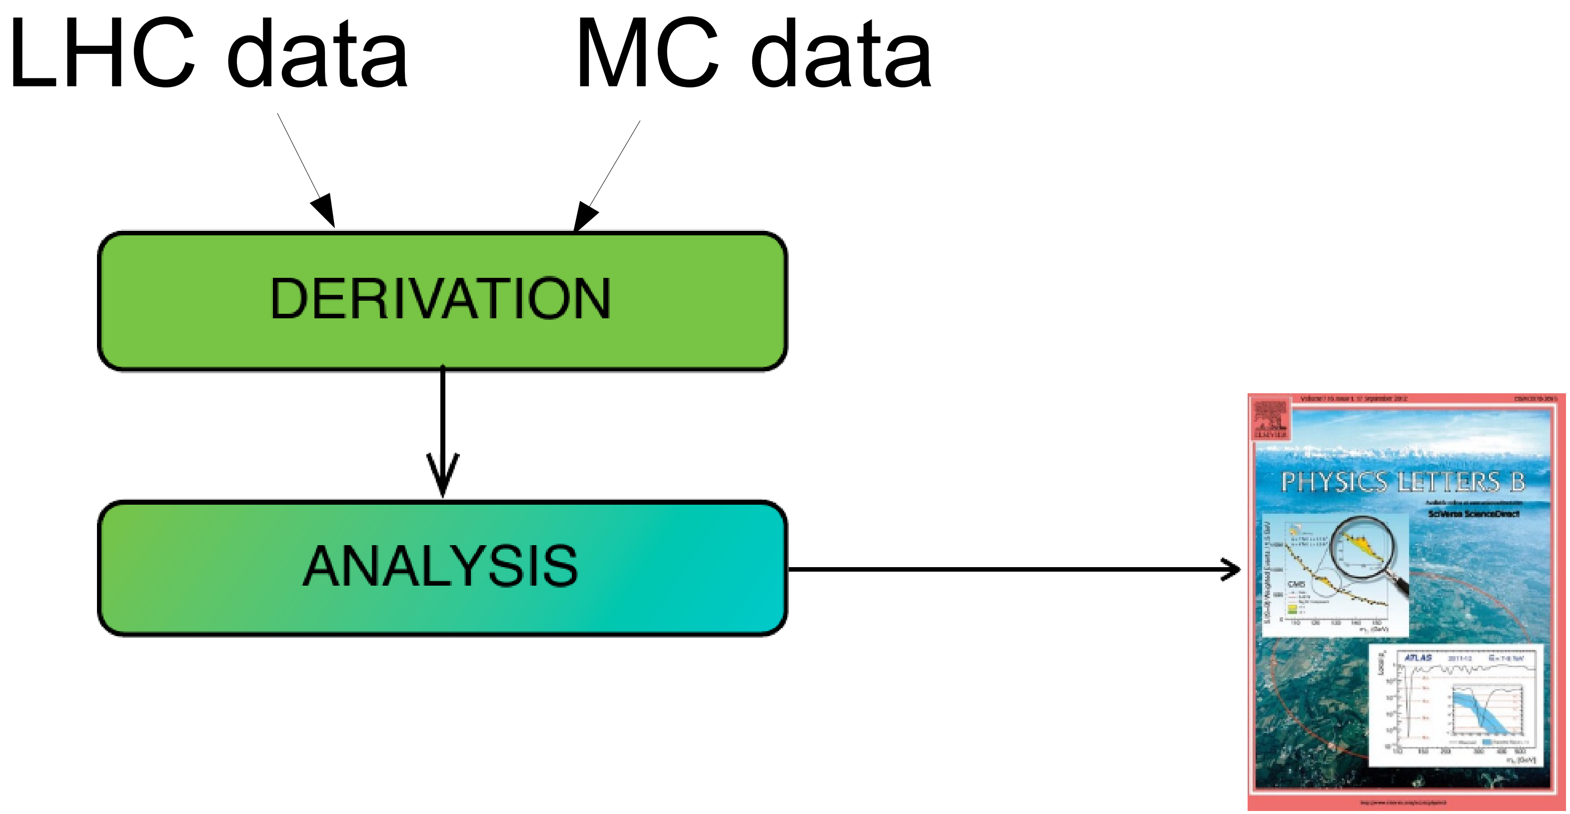

Derivation step

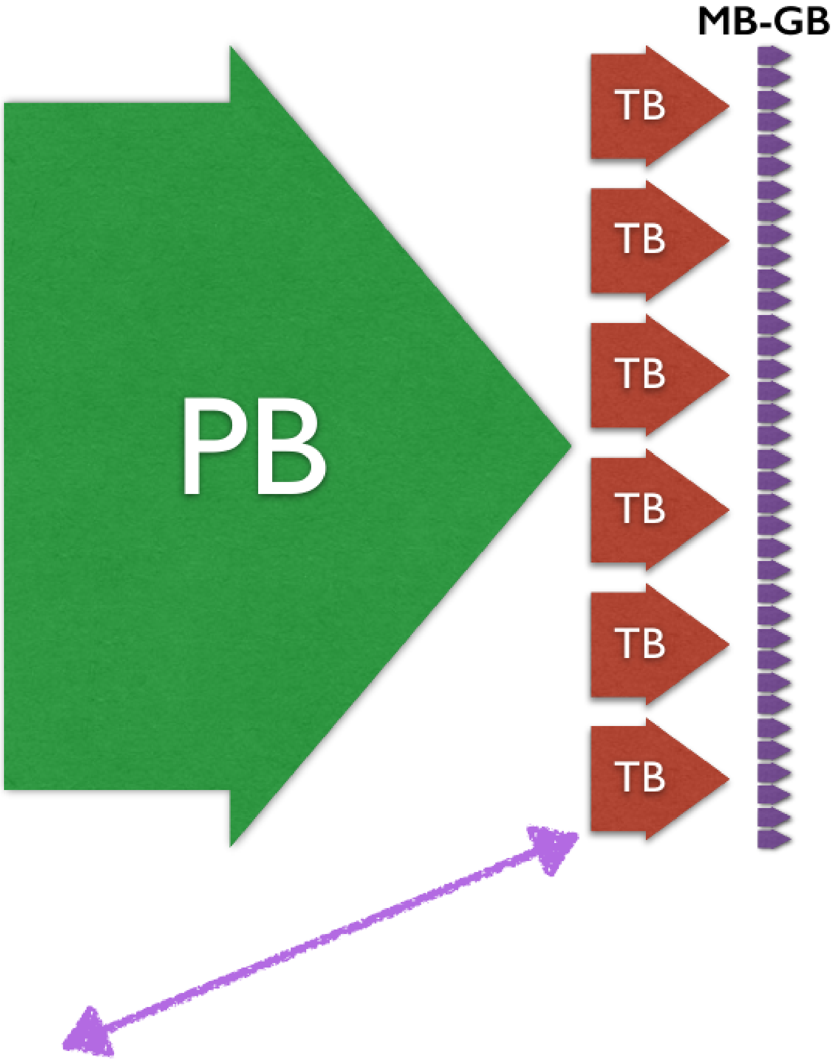

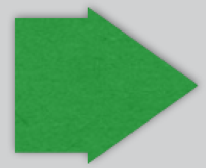	Full output of reconstruction, ~PB size	One format
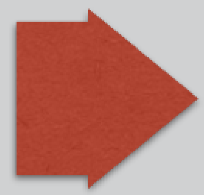	Intermediate analysis format ~TB size	~100 formats
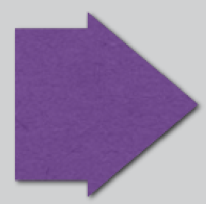	Final n-tuple ~MB-GB size	~1000 formats

- These formats tend to be specific to a single analysis or group of analyses
- Calibrations and common object selections are often applied as they are made
- They generally need to contain all variables needed for calculating systematics

Data formats in ATLAS

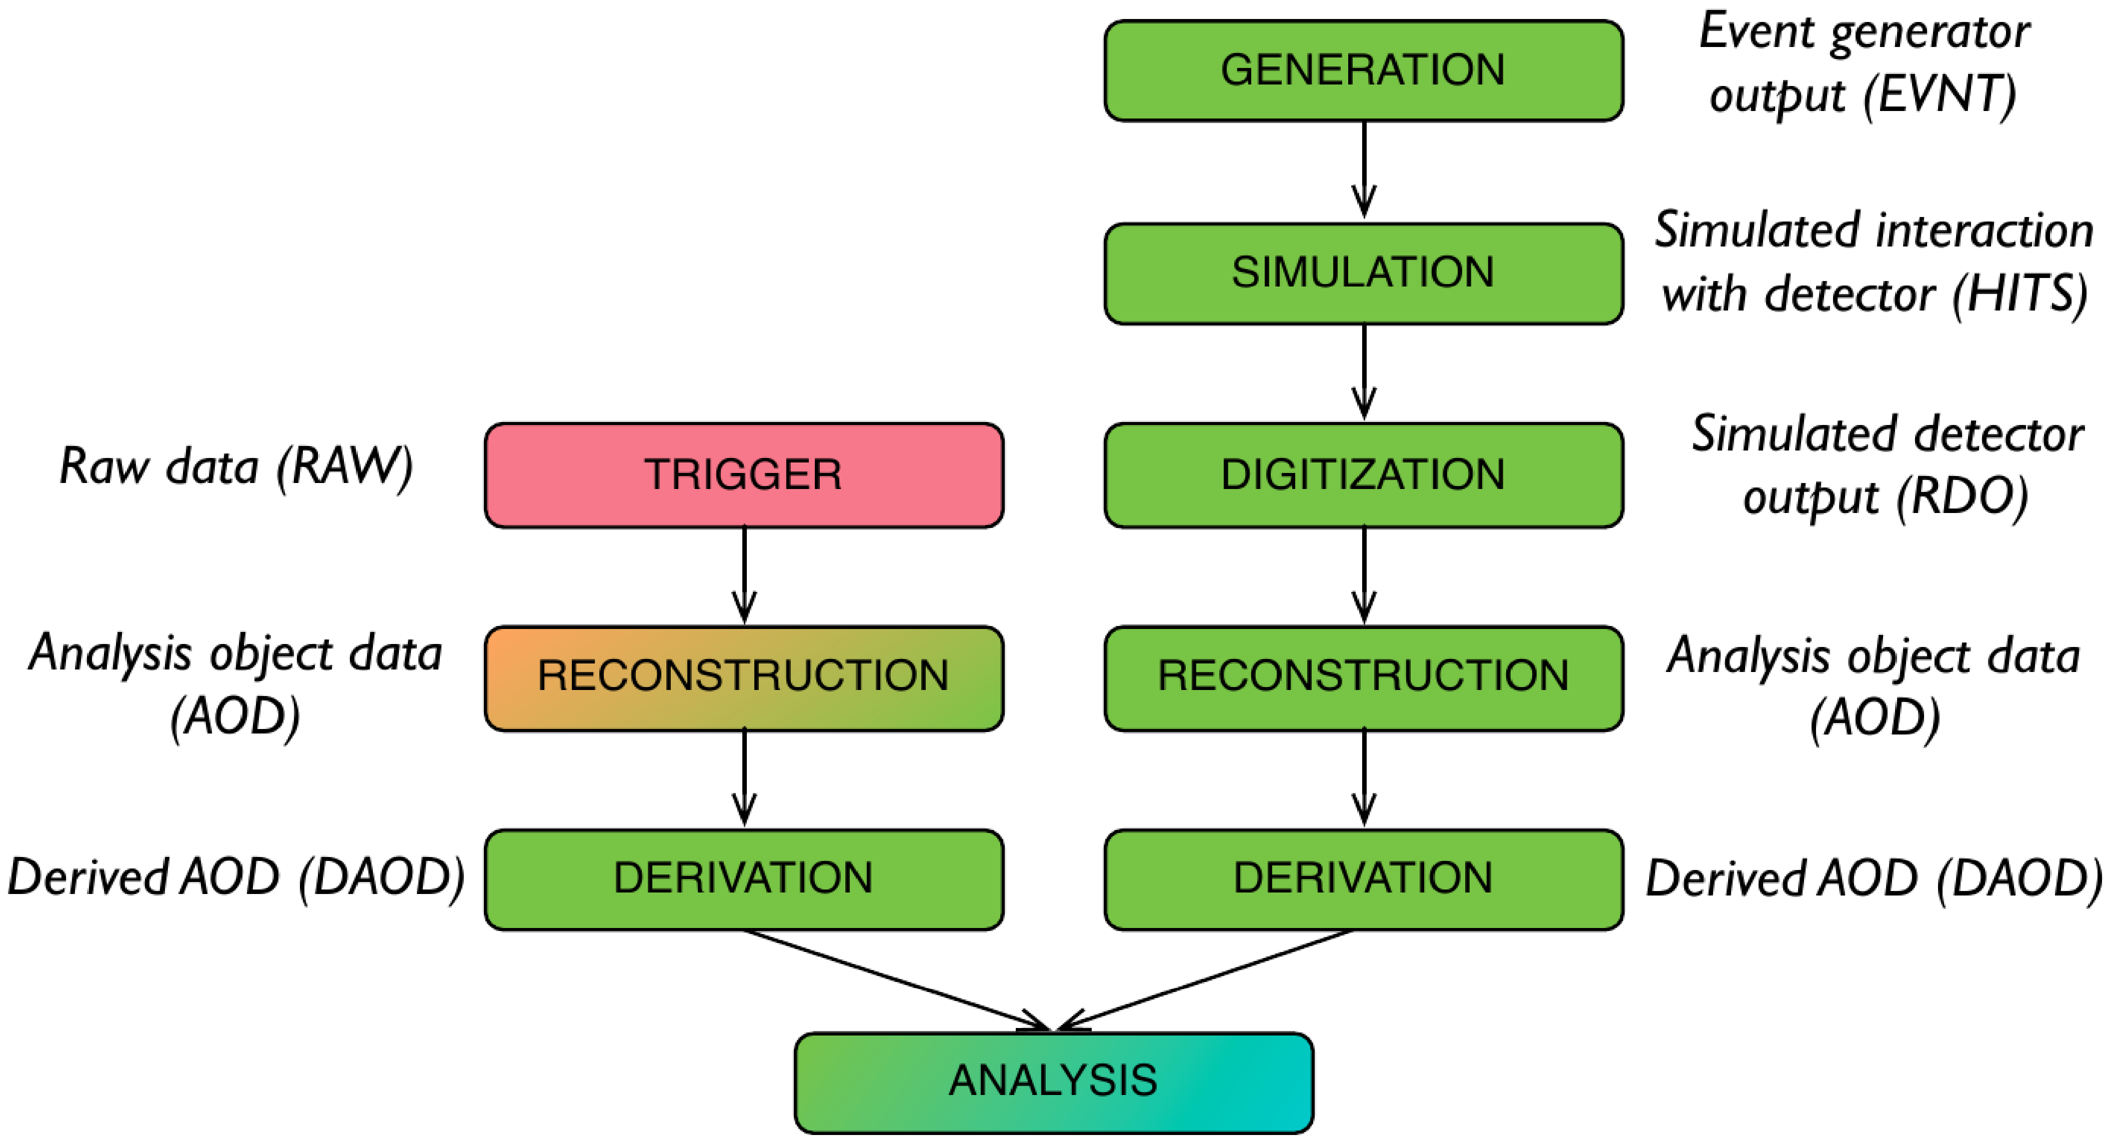

Data formats in ATLAS

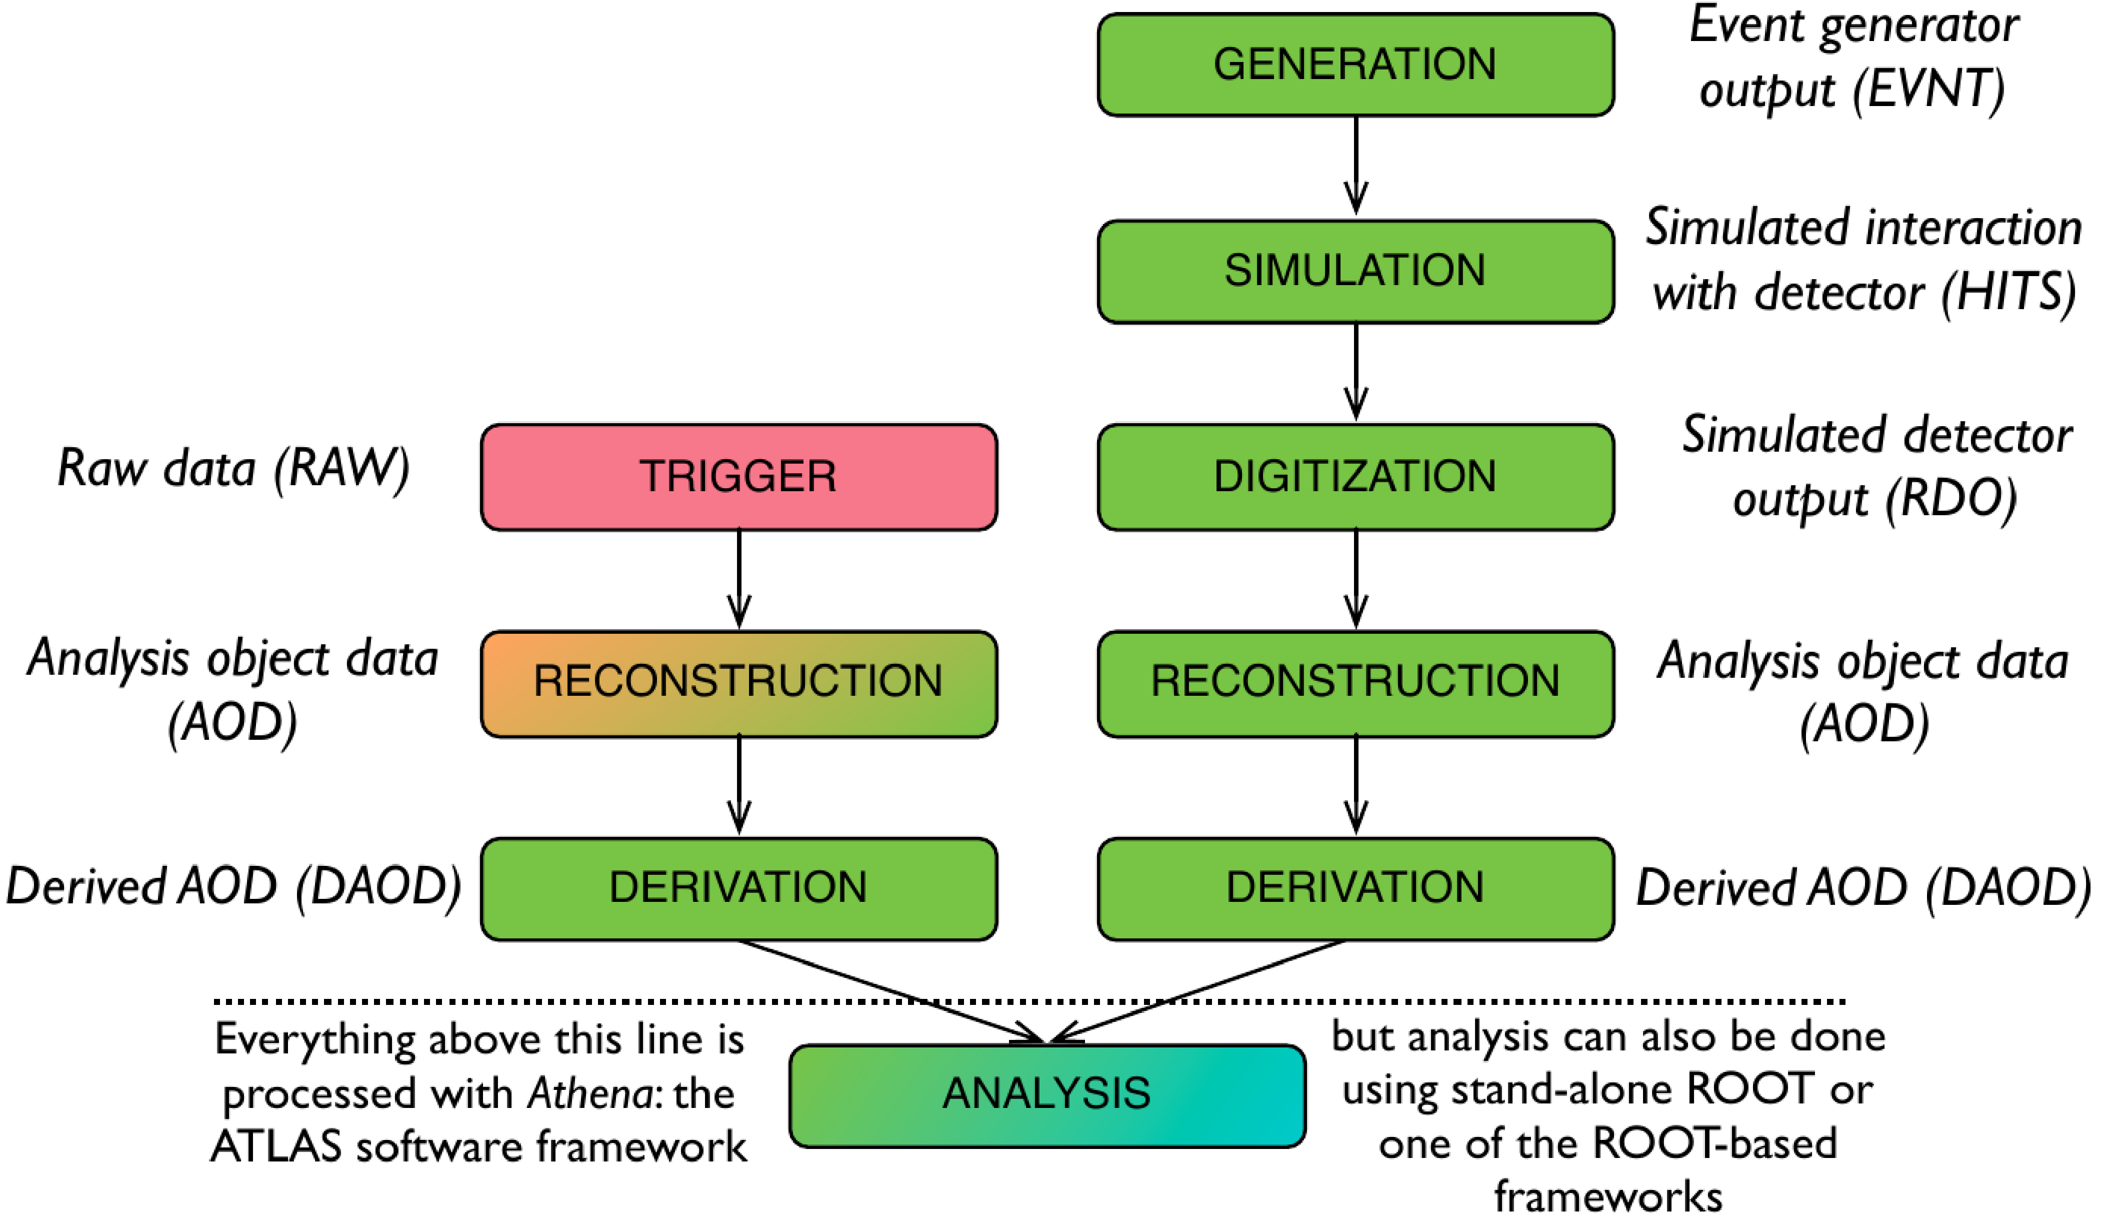

Dataset nomenclature

Data:

project tag: 2012 pp data 8 TeV

Run number

stream

merged files

Data type: AOD

data12_8TeV.00209980.physics_Egamma.merge.AOD.f476_m1223

AMI tag describes configuration of each step (Tier-0 bulk reconstruction f, file merging m)

Simulation:

project tag: "mc12" setup 8 TeV

MC DSID unique # for process

"human-readable" description of MC sample

merged files

Data type: SM W/Z ntuples

mc12_8TeV.119353.MadGraphPythia_AUET2BCTEQ6L1_ttbarW.merge.NTUP_SMWZ.
e1352_s1499_s1504_r3658_r3549_p1328/

AMI tag describes configuration of each step (evt generation e, full simulation s, reconstruction r, D3PD creation p)

/: is a "container" (points to other datasets)

- Datasets are collections of files, form fundamental unit of ATLAS data and MC
- Each processing step changes the data type, and adds the AMI tag used
- Changes with Run 2
  - the merge m-tag has mostly disappeared for reprocessed data and MC
  - the p-tag will in general refer to the Derivation Framework version - DAOD

Dataset Nomenclature & Analysis

- Real data is grouped together into "periods" with ~stable running conditions
  - 2012 data broken down into periods A-L
  - "AllYear" covers the whole year
- Possible to have several "processing campaigns" for a given dataset
  - First processing at Tier0 is e.g. t0pro14
  - A reprocessing campaign at Tier1s is e.g. repro14
  - The combination with overlaps removed pro14
- We create "containers" to collect several datasets that belong together, easier for use in your analysis, e.g.:
  - data12_8TeV.periodAllYear.physics_Muons.PhysCont.AOD.pro14_v2/

Components of the analysis model

- xAOD
  - ROOT-readable data format and event data model
  - Produced by directly by reconstruction
  - Much more on this later
  - Note: we refer to the *files* as AODs, but the *contents* as xAOD. Confusing, I agree.
- Derivation framework
  - AODs are too big to analyse directly so we centrally reduce them according to the needs of the physics groups, producing DAODs which are still made of xAOD objects, but are much smaller
  - Avoids users having to do this themselves on the Grid
  - Currently producing close to 90 DAOD formats for all ATLAS activities
- Analysis framework
  - What you use to do your analysis
  - Capable of reading the xAOD objects and applying all tools from the combined performance groups

From here


IT'S UP
TO YOU
DO SOMETHING!


CAUTION
MEN AT
WORK


ANY
QUESTIONS
?
www.rekruitin.com
ReKruiTIn.com

What's all this grid nonsense then?

- Recognition in the early planning of LHC computing that a single centre would not supply enough resources for the LHC experiments
  - Mainly because no one would fund a central computing facility

- In fact ATLAS uses three grids:
  - EGI in Europe, Asia and Canada (UMD middleware)
  - OSG in USA (OSG stack)
  - NorduGrid in Nordic countries (ARC middleware)
- All badged as "Worldwide LHC Computing Grid compatible"

Computing Model Cartoon

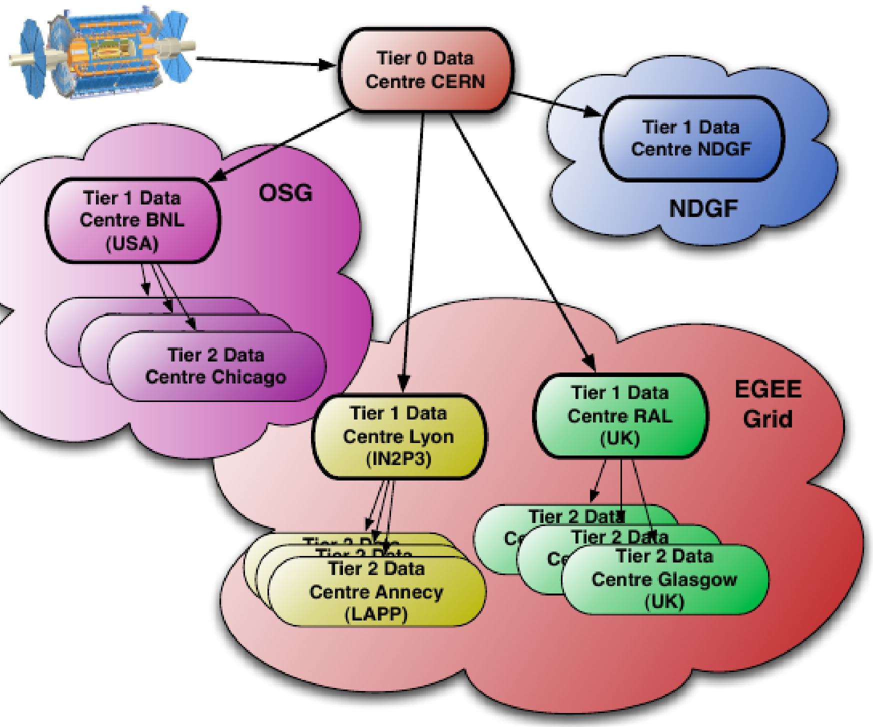

- 1 Tier-0: CERN
- 10 Tier-1: National Computing Centres (BNL, RAL, IN2P3, ...)
- 40 Tier-2: Regional Computing Centres (ScotGrid, Frascati, Toronto, ...)
- Composed of multiple individual sites
  - ~100 Analysis queues in PanDA

- T0: prompt reconstruction
- T1s/T2s: reprocessing, MC production, derivation production, analysis (mainly at T2s)
- T3s: analysis only

ATLAS Grid Architecture

- Distributed Data Management System (DDM): Data movement and catalog system, policy engines

New version for Run 2: Rucio

- Other databases and data book keeping (AMI, GRL, Coma for metadata; conditions data in distributed Oracle + frontier; Event Index; etc.)
- Production System (PanDA)

New for for Run 2: JEDI and DeFT

- Distributed Analysis Interfaces (ganga + pathena)

DDM Rucio

- Data management is at the core of all ATLAS grid activities
  - All data organised as datasets
    - Which contain multiple files
  - Datasets are the ‘work units’ on the grid
    - So the DDM central catalog records the location of datasets on the grid, as well as the content of each dataset
    - Datasets are also the basic input to an analysis
  - DDM also moves data between sites on the grid

https://twiki.cern.ch/twiki/bin/view/AtlasComputing/RucioClientsHowTo

Send in the Jobs!

- To run ATLAS jobs on the grid you need an distributed analysis framework which
  - Understands the ATLAS data model
  - Understands the grid
  - Understands the analysis code
- Distributed Analysis Interfaces (ganga + pathena)

https://twiki.cern.ch/twiki/bin/view/AtlasComputing/DistributedAnalysis

for listening

Any further