# Data Preparation in ATLAS

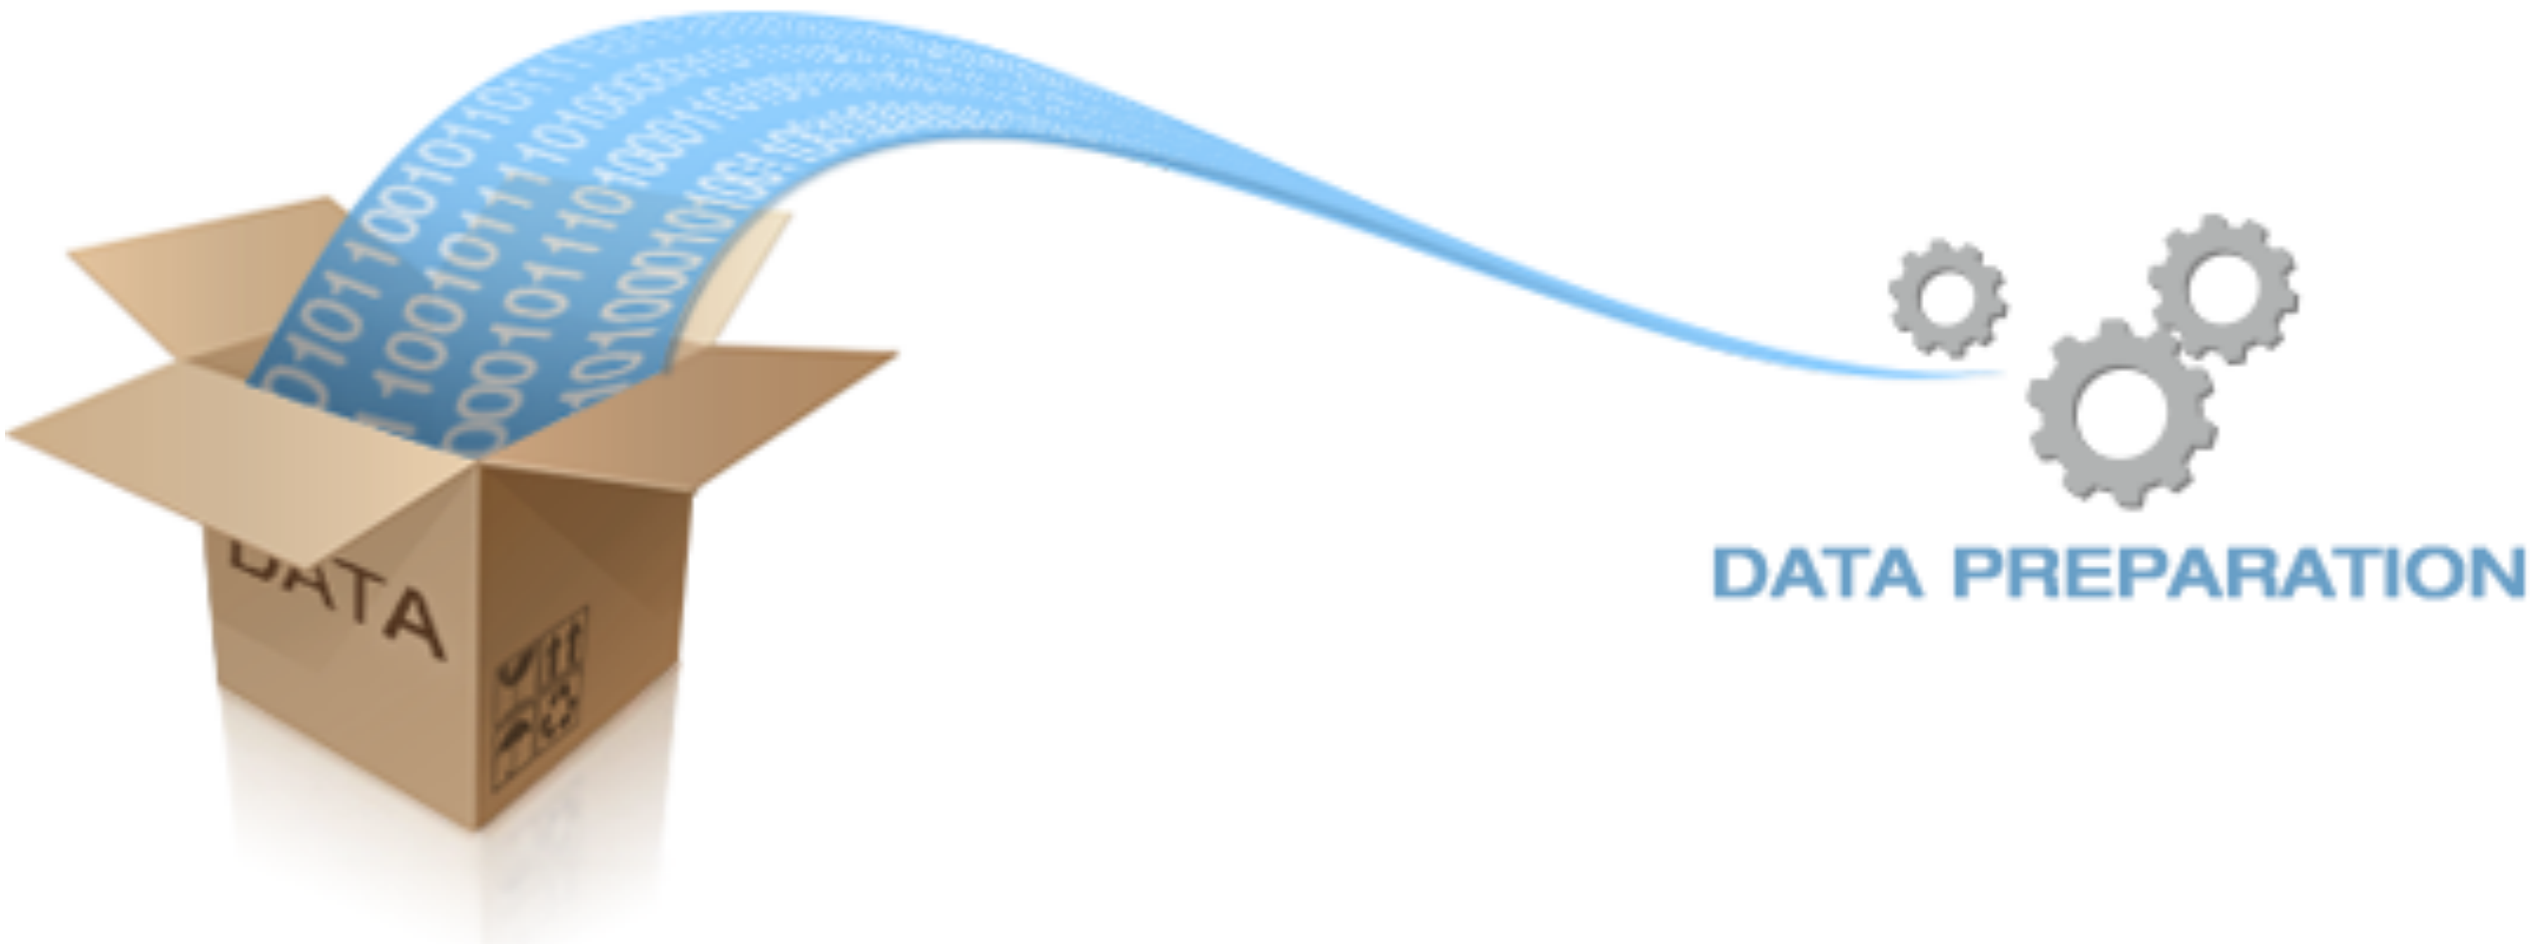

# Introduction

- Will try to give an introduction to the area we call data preparation, where I am currently the deputy coordinator.
  - Will take the opportunity to also discuss other things that are tangential, and which can be interesting to understand.
  - Aimed at Ph. D. students and non-ATLAS members, so some of you will find a lot of material familiar.
- The ATLAS collaboration consists of over 3000 physicists, from institutes all over the world.
  - Over 1000 of them are Ph. D. students.
  - Some are stationed at CERN, most are not.

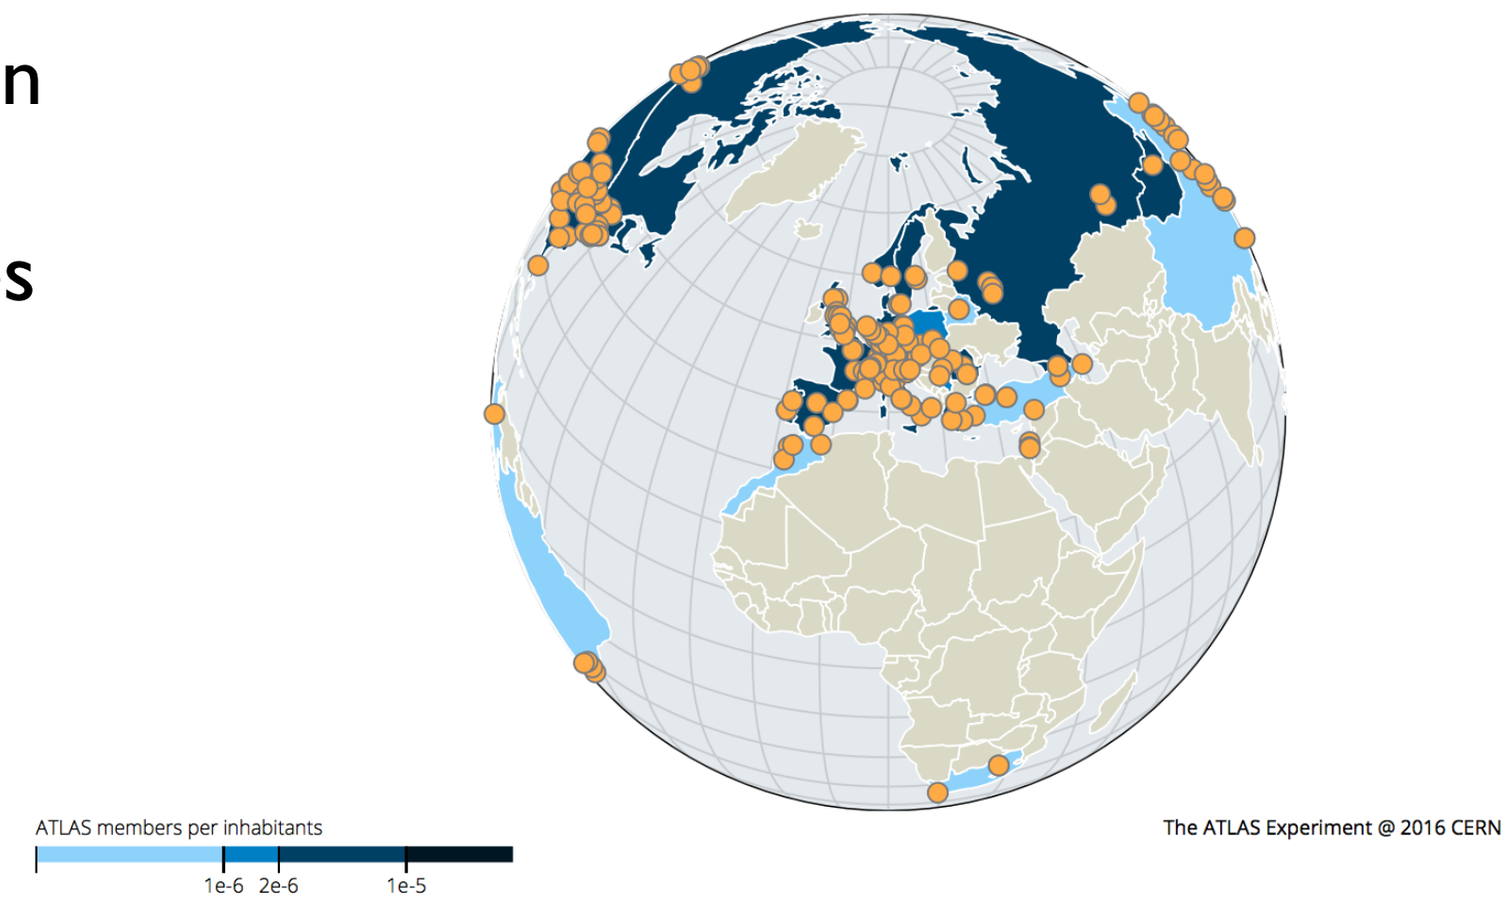

# Introduction

- In such a big collaboration, the work needs to be organised.
  - Most of what happens in ATLAS is divided into five "Activities", which (if you are a ATLAS member) you recognise from our home page where they are the big buttons at the top.
  - In general terms, data preparation is what happens to the data after it leaves the detector and before it is used in physics analysis.
  - It acts a little as a middle layer between detector, trigger and comp-uting, and the physics activity.

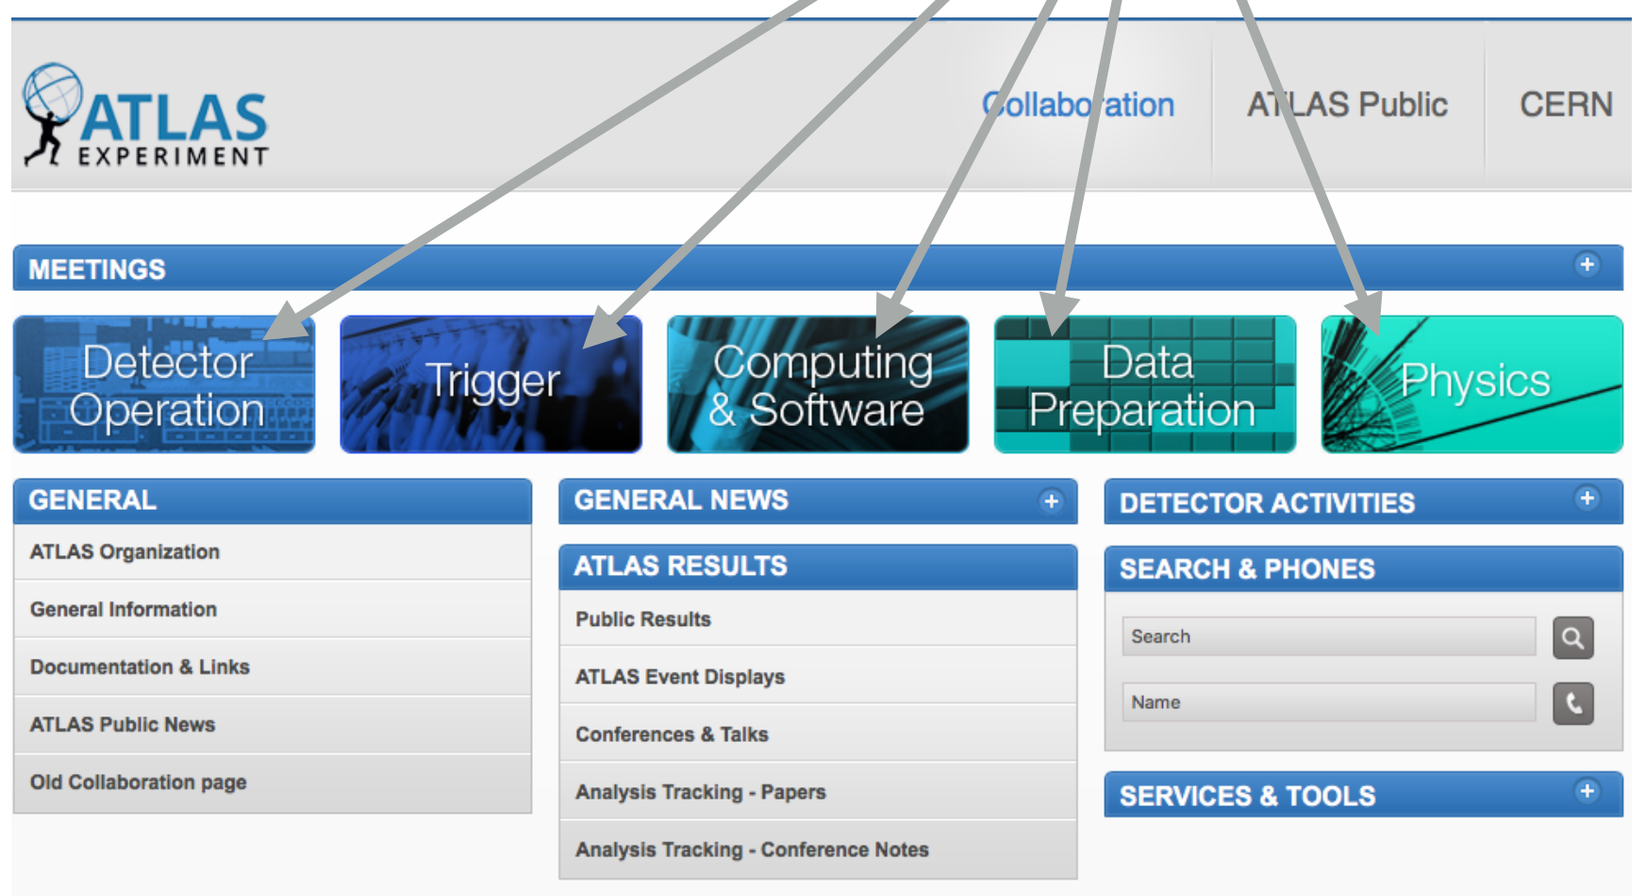

# Introduction

- The tasks of the data preparation group includes the data quality assessment, both online and offline, the reconstruction of the data, the low-level calibration, the luminosity and so on.
  - Before explaining some of these efforts, I would like to start by giving some background to why we create such massive datasets.

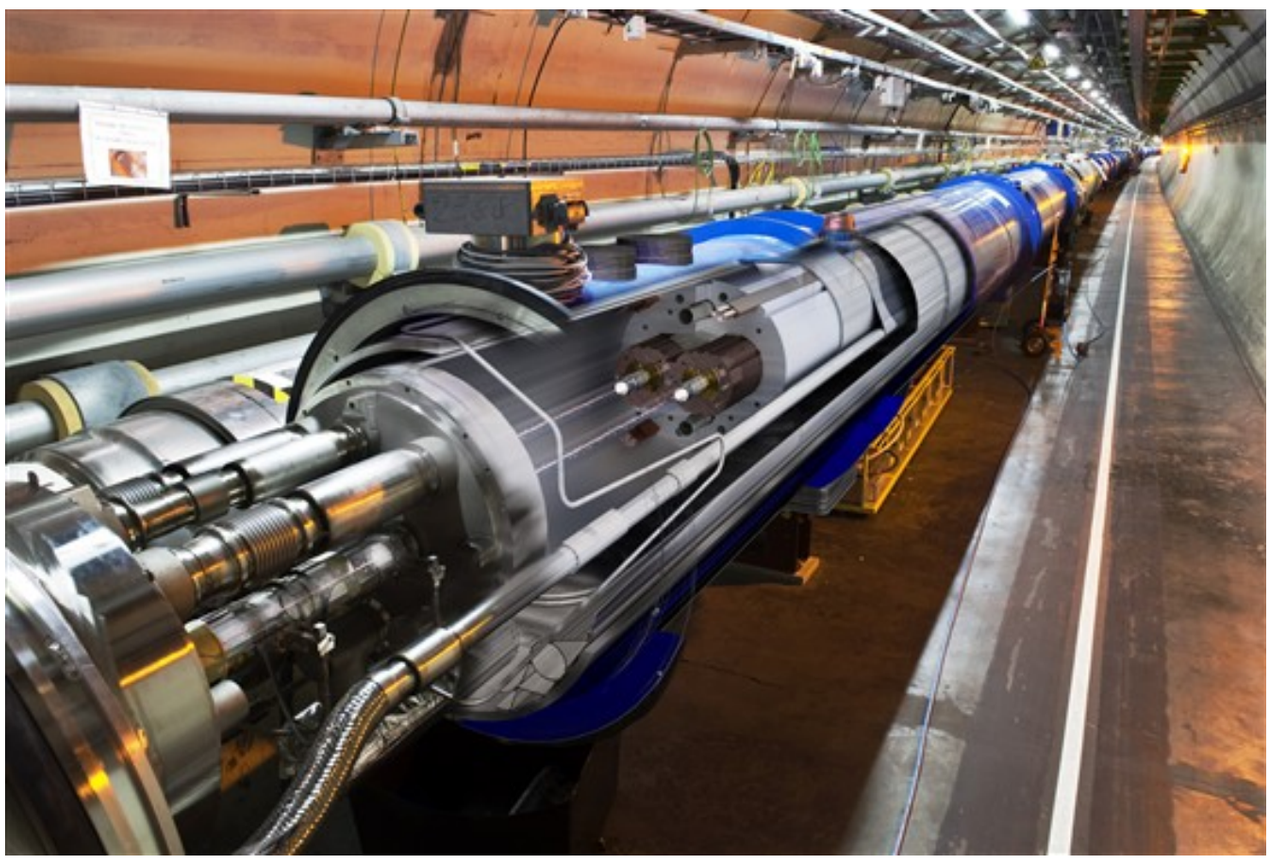

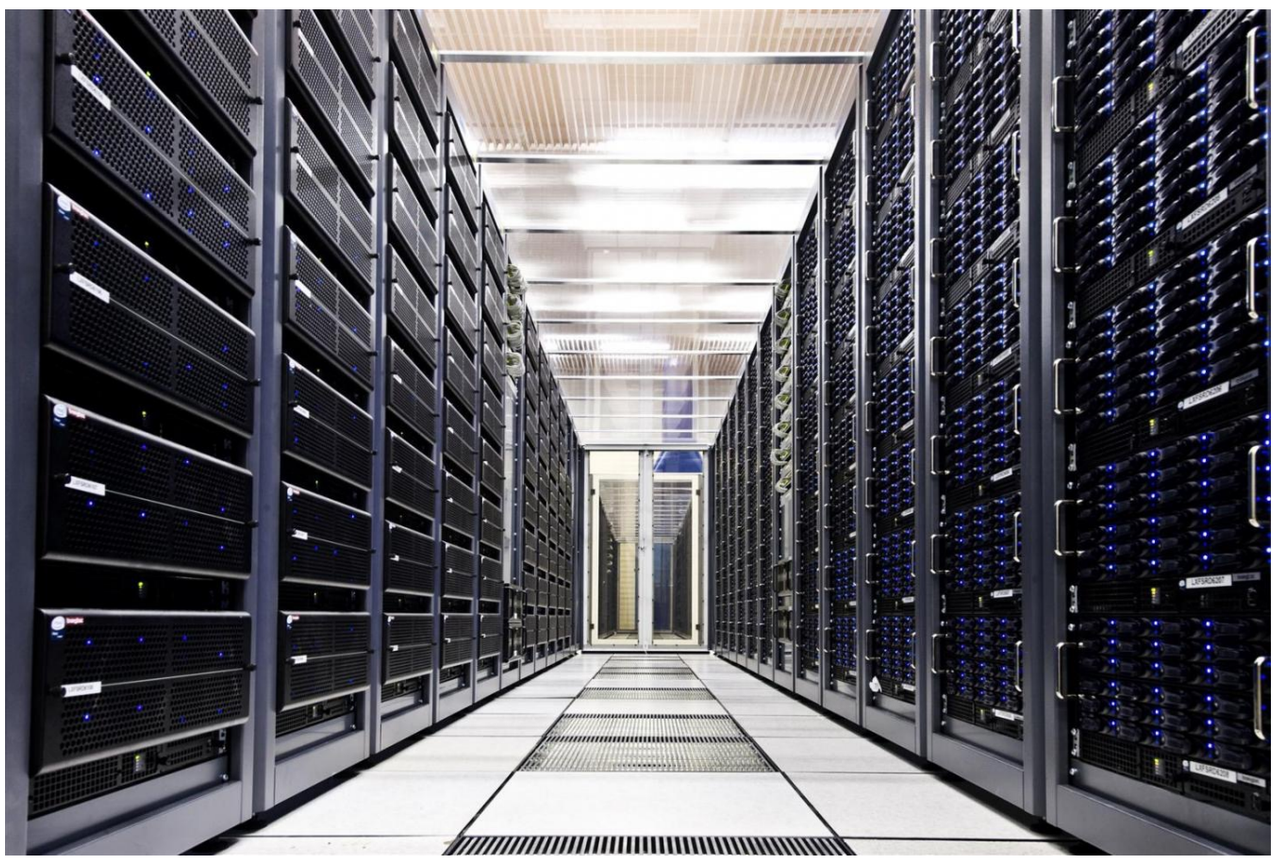

# Accelerators

- The potential of an accelerator is determined by two factors:
  - The *collision energy*, which determines how heavy particles can be that are created in the collisions, and influences the probabilities (cross sections) for interesting things to happen.
  - The *luminosity*, which is a measure of "how many" collisions the accelerator delivers, per unit of time or integrated over some time period.
- The LHC is the highest energy accelerator ever, but it is also (I think) the highest luminosity accelerator ever.
  - Generates large dataset.

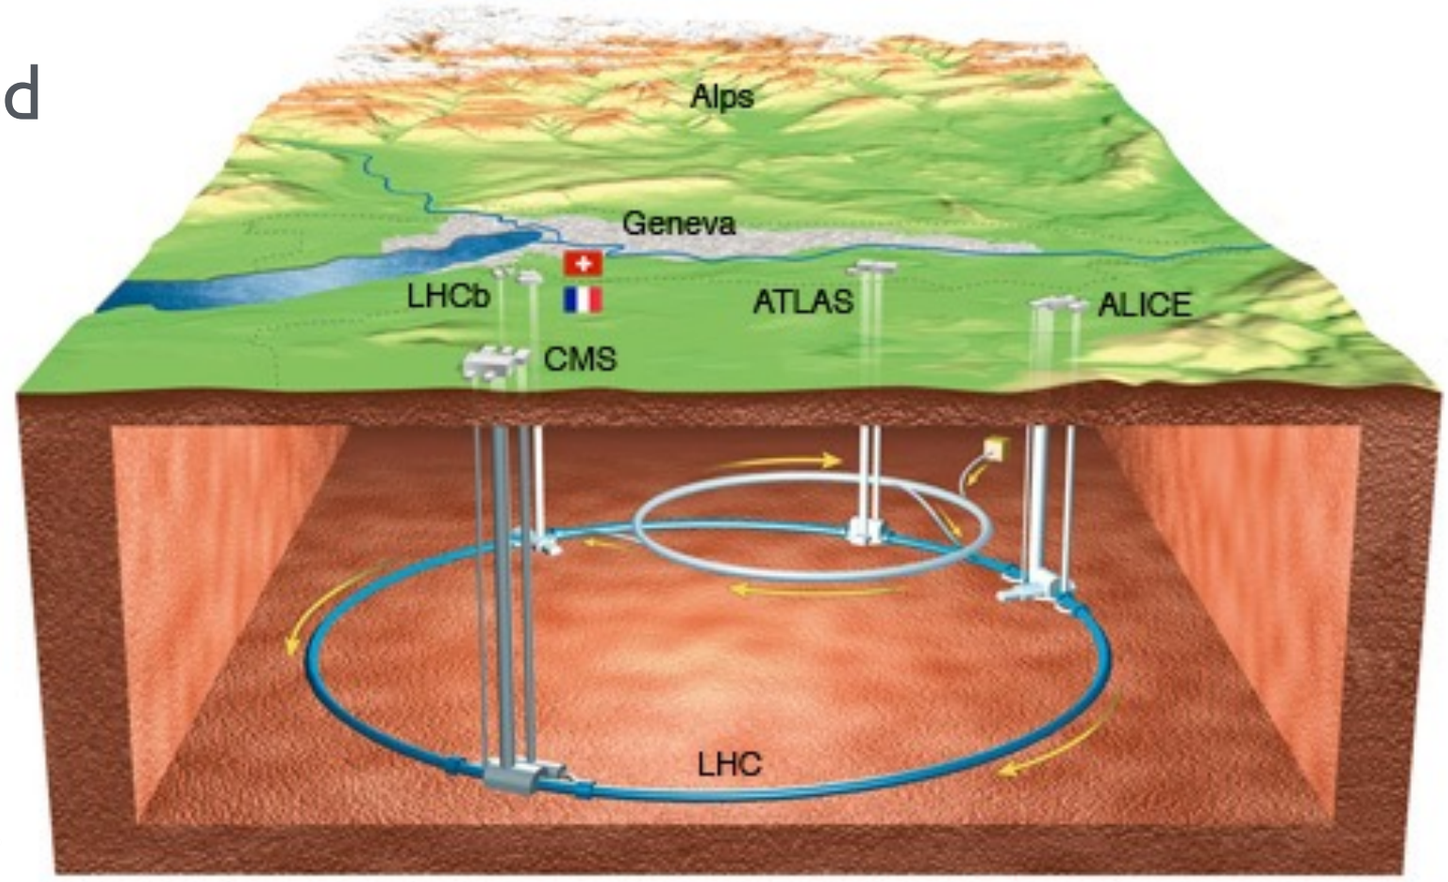

# Collision Rate

- To reach the highest intensity of collisions, the protons in the LHC are accelerated in bunches of about $10^{11}$ protons each.
  - The beams are brought into collision at the 8 interaction points.
- At an energy of 6.5 TeV, the speed of the protons is 0.999997828 c. The spacing between two bunches is 25 ns.
  - The distance between bunches is $3 \cdot 10^{8} \times 2.5 \cdot 10^{-8} = 7.5$ meters.
  - Given that the ring is 27 km long, a maximum of about 3600 bunches fit.
  - In reality that number is slightly smaller, we aim for about 2400 bunches per beam this year.
- 25 ns means 40 MHz collisions!

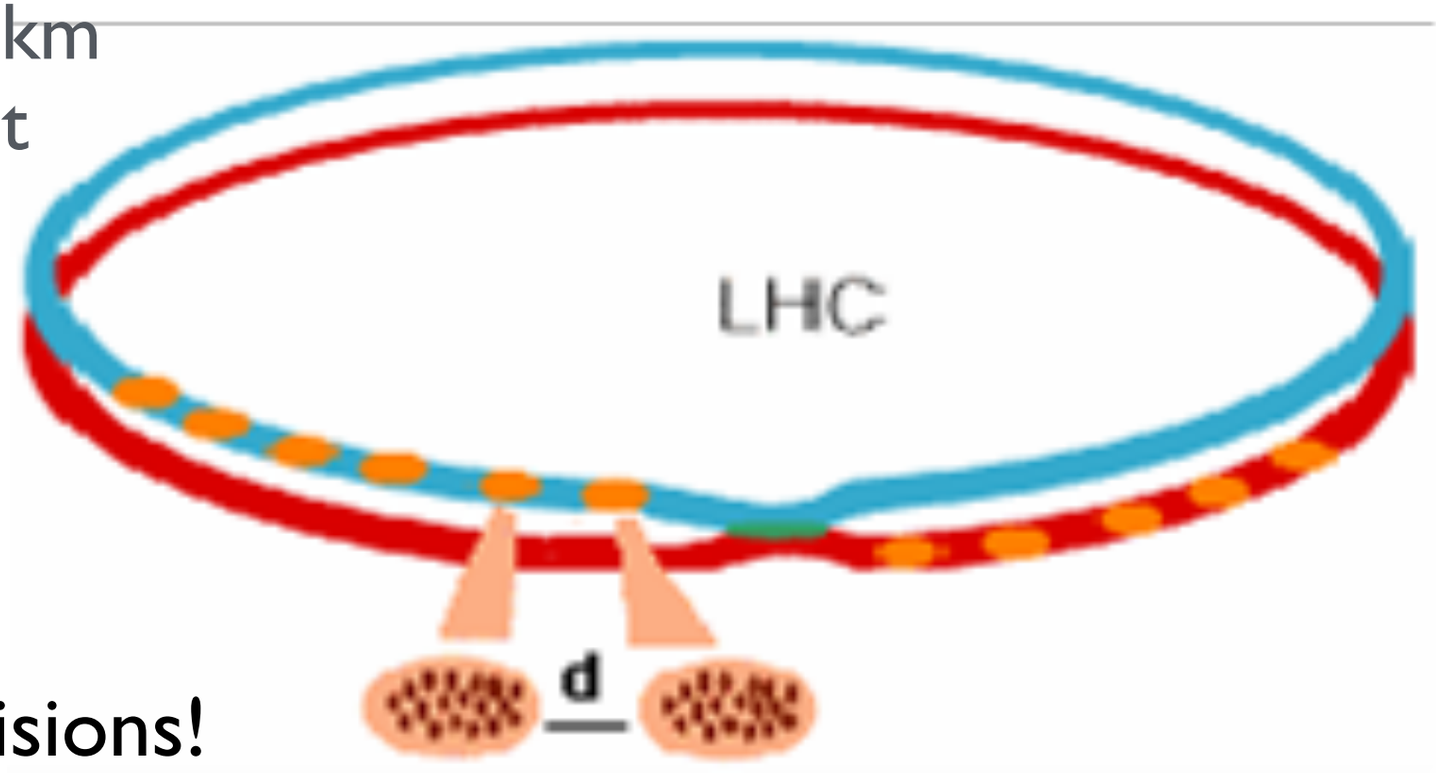

# Multiple Interactions

Jonas Strandberg

- Other measures that can be taken to increase the luminosity, besides increasing the collision rate, are for example:
  - Increasing the number of protons in each bunch.
  - Making sure that the bunches are tightly squeezed in space.
- As a result each time two proton bunches collide, several p-p pairs can interact with each other.
  - Here you see a zoom to the interaction region, and the reconstructed charged particles in an event. 7 collisions seen.
  - Such simultaneous interactions are called pile-up, measured by $\mu$.
  - This year $\mu$ will be around 30.

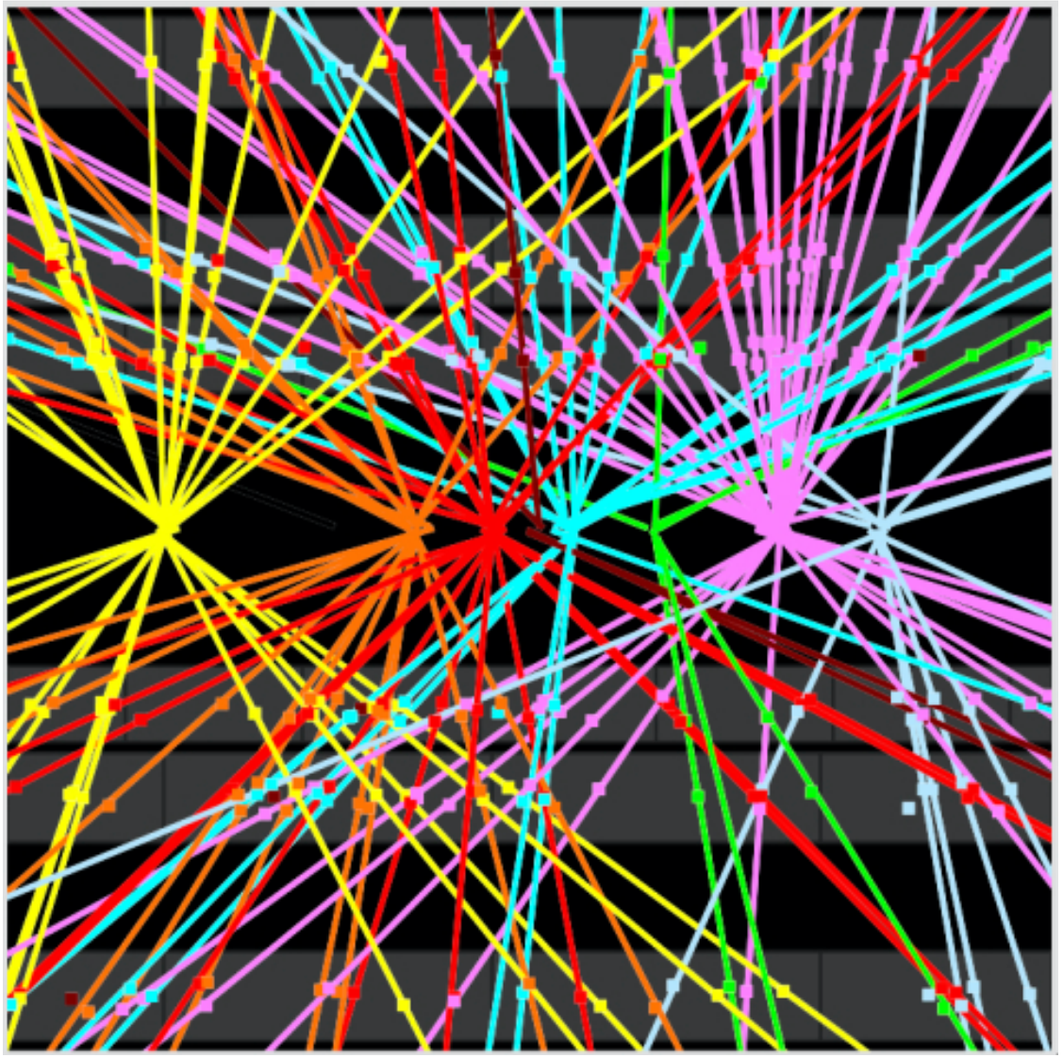

# Same Event, More Pile-up…

Jonas Strandberg

- CMS Simulation: 300 GeV H→ZZ→eeμμ at various luminosities

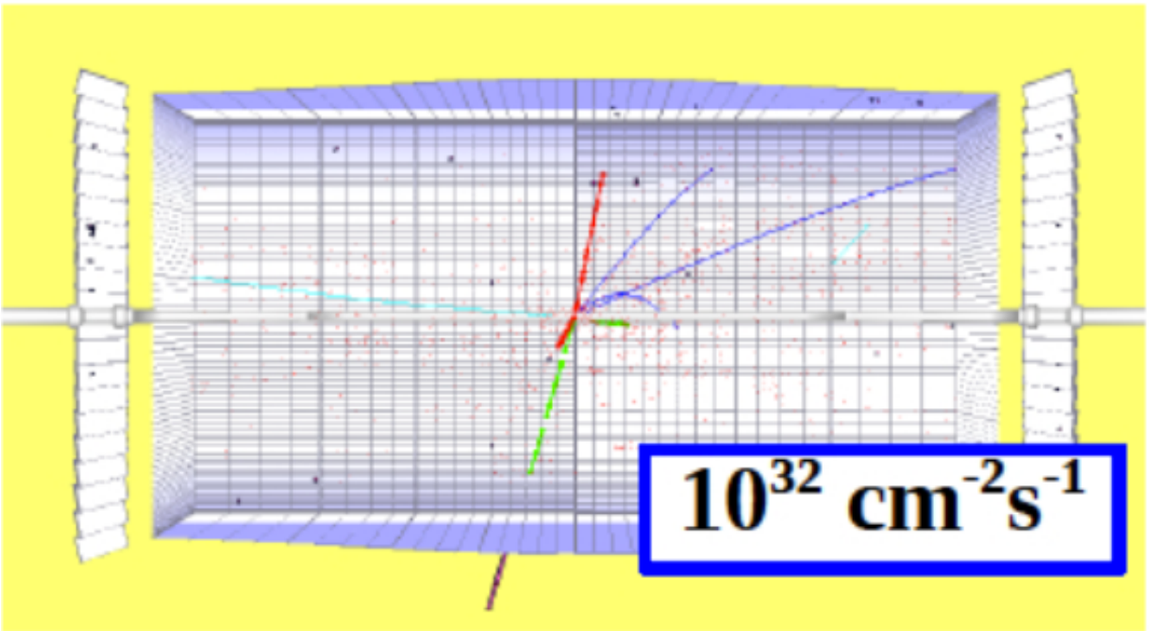


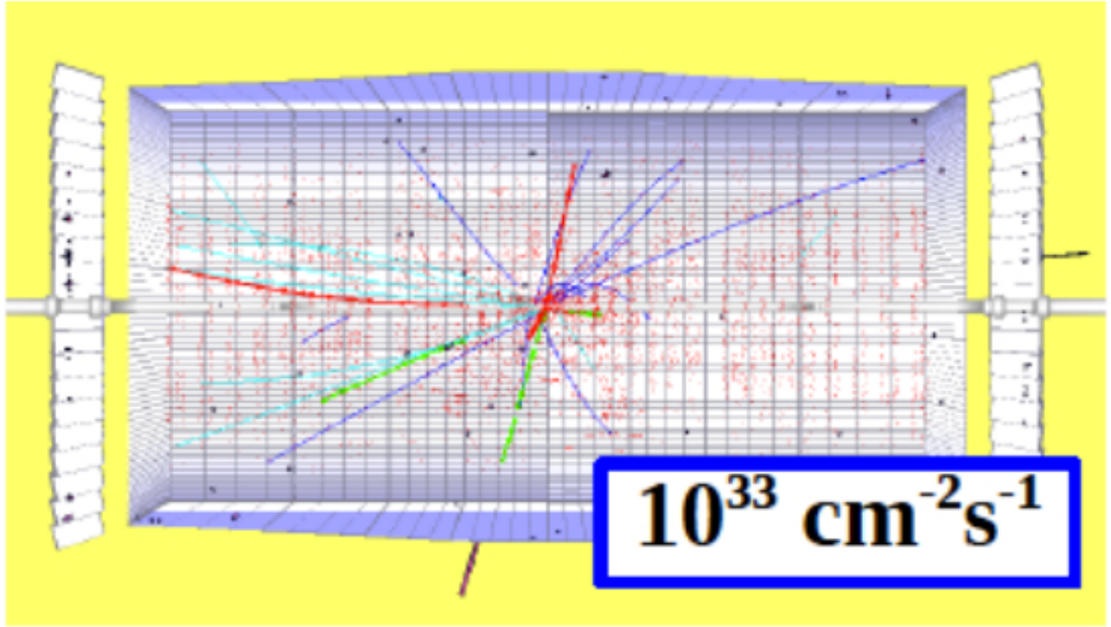


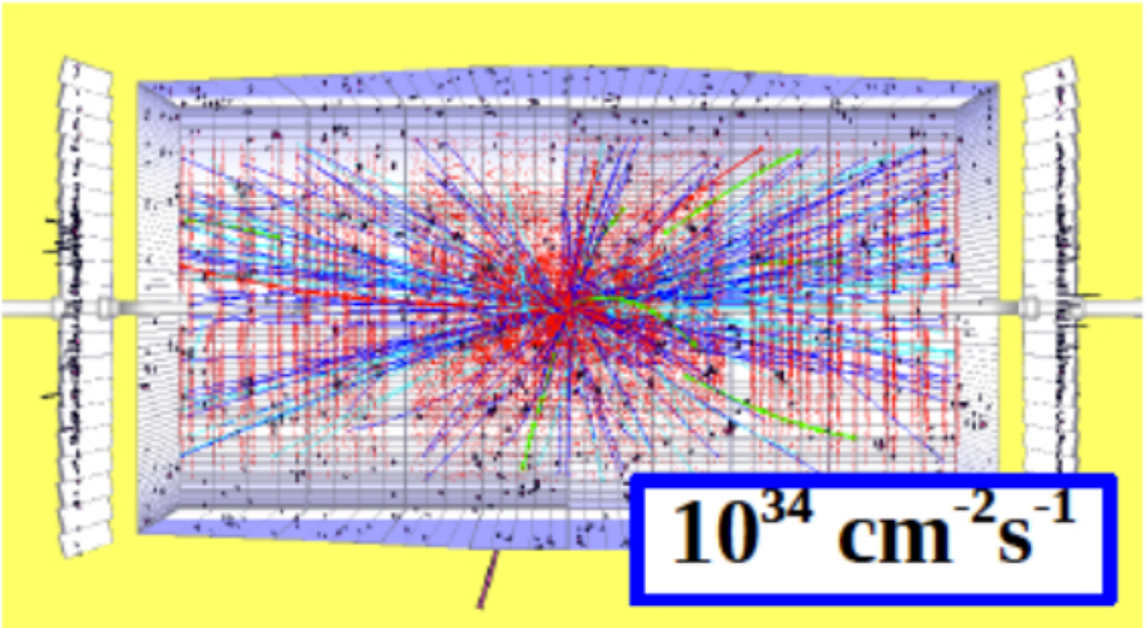


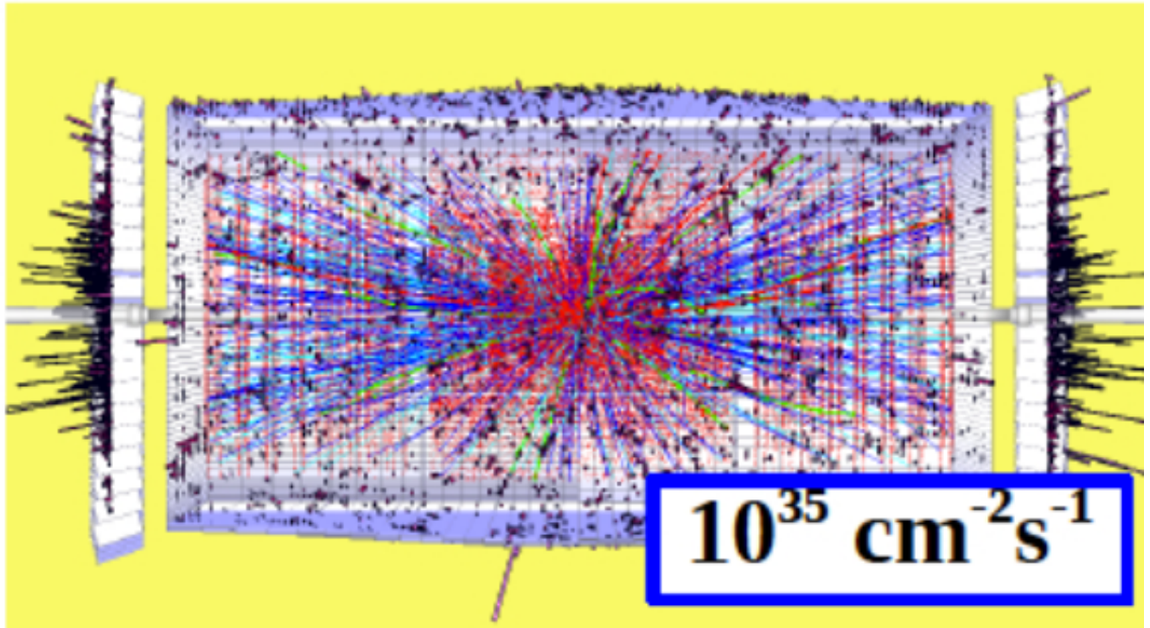

# Why the High Collision Rate?

Jonas Strandberg

- Production of new and heavy particles happen rarely when two protons collide at the LHC.
  - E.g. the discovery of the Higgs boson was one of the main goals.
  - Only one out of ten billion pp collisions create a Higgs boson.
  - 40 MHz times 30 interactions times $10^{-10}$ gives ~1 Higgs per10 s.
  - Then only a fraction of the Higgs boson decays can be found amongst the large backgrounds.
- New BSM particles will be created even less often than the Higgs.

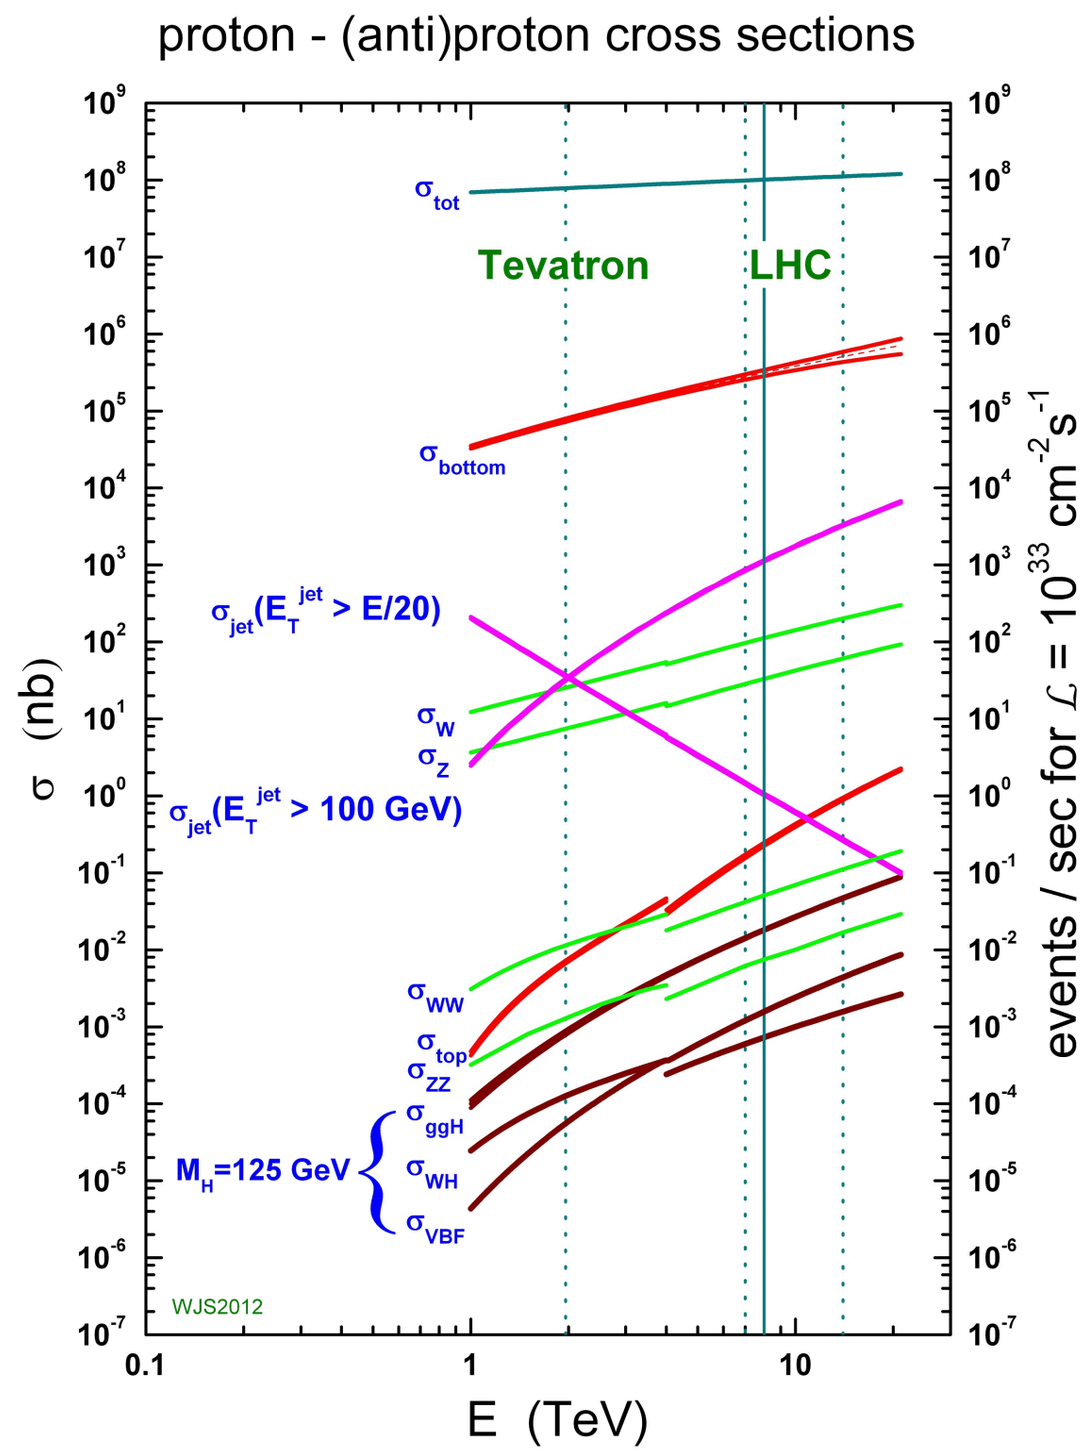

# Measured Cross Sections

Jonas Strandberg

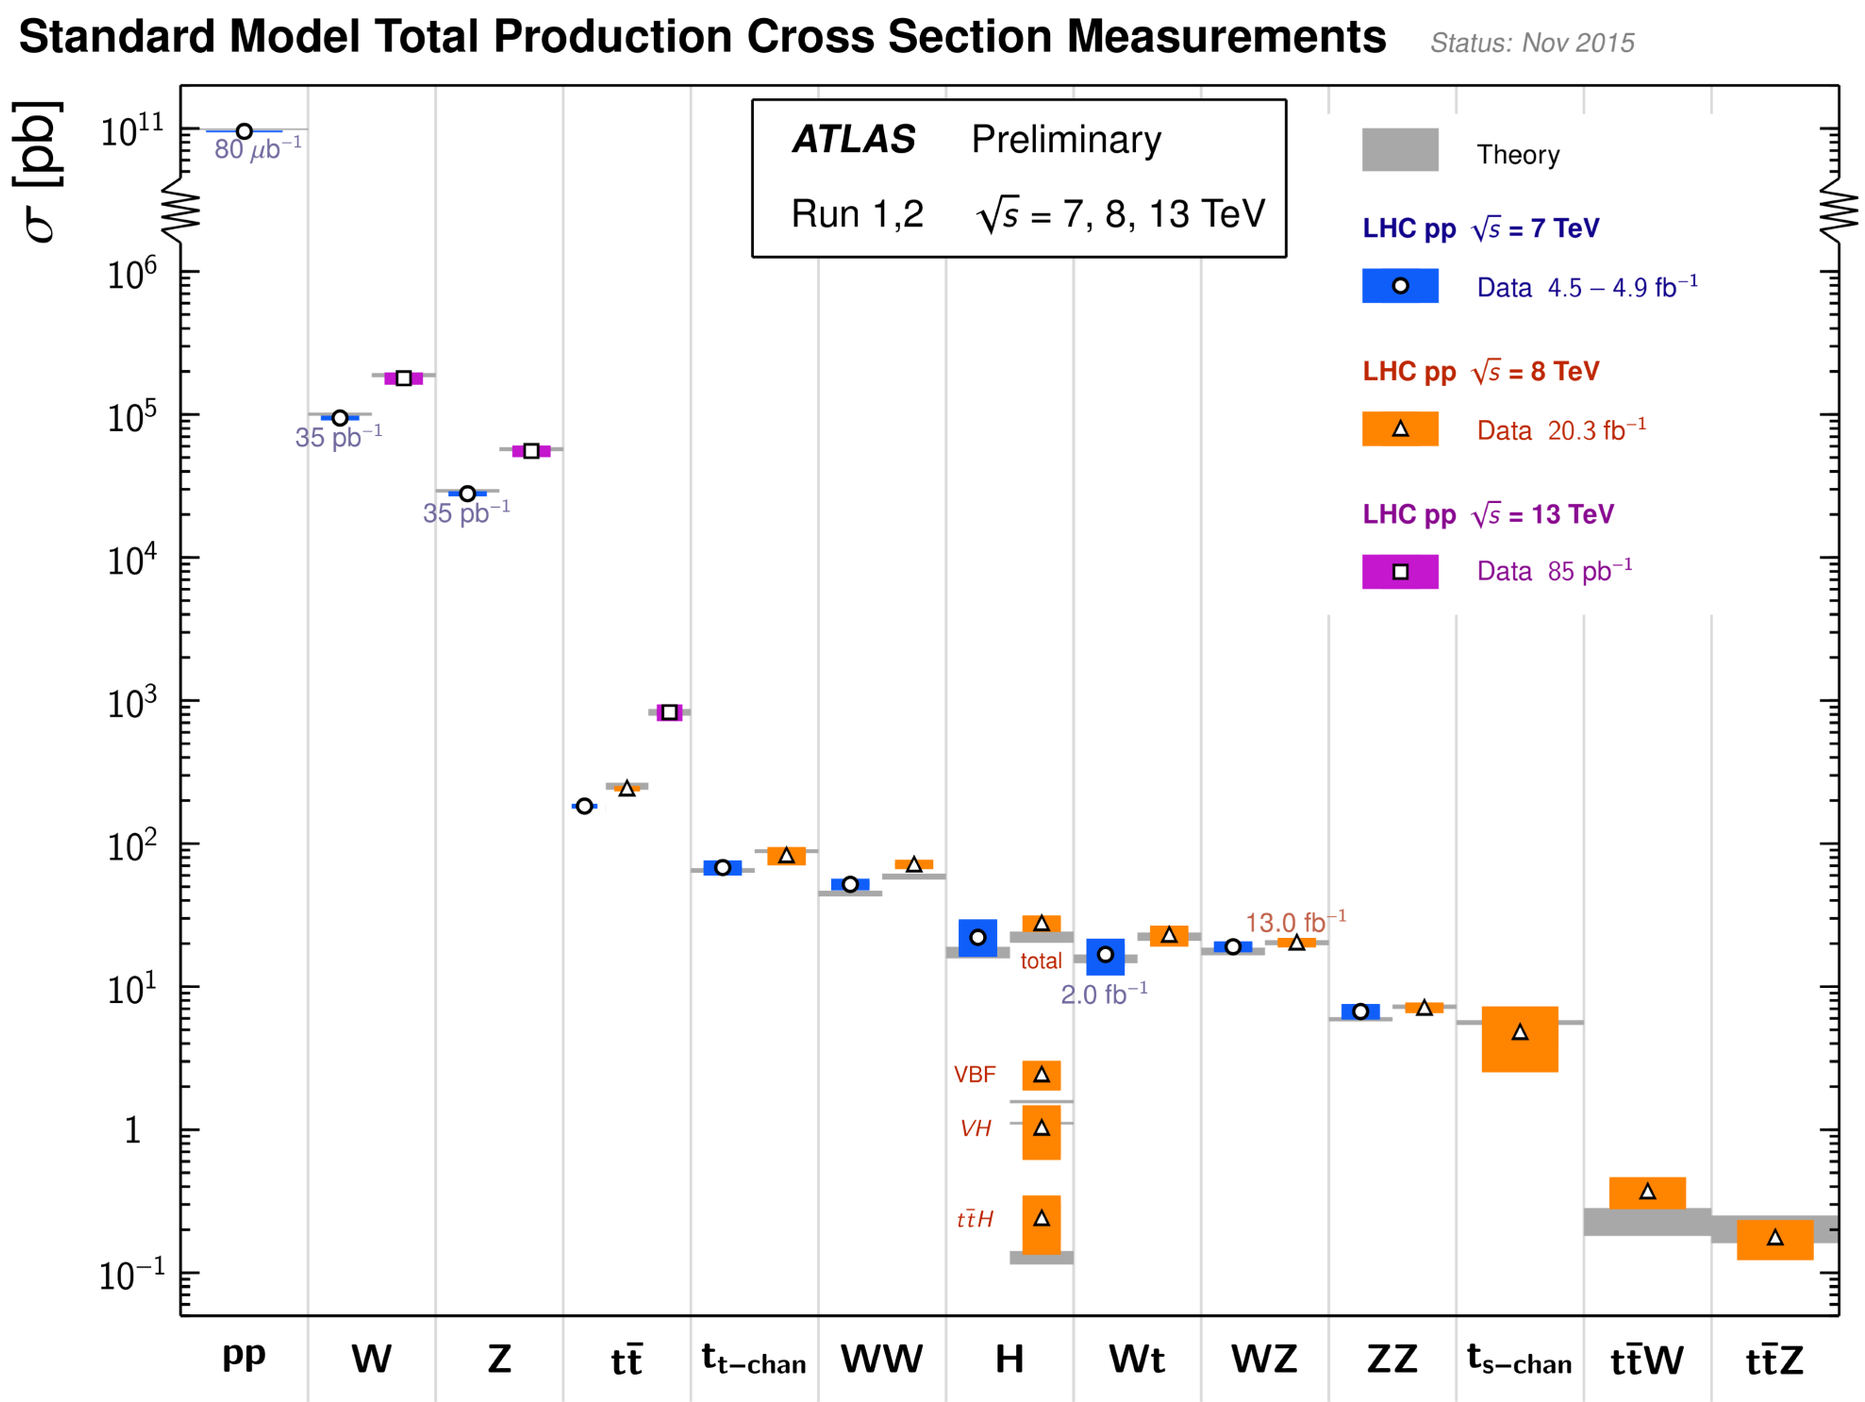


- Cannot resist showing one slide on the measured σ too…

# Data Size

- So we need a lot of collisions to create a statistically significant sample of Higgs bosons, or eventual new BSM particles.
- The RAW data contains all the hits from the detector.
  - A *hit* typically means an energy measurement from a passing particle, and an identifier for the detector channel.
  - The ATLAS detector has about 100 million readout channels.
- Even if we are not reading out empty detector channels in the RAW data, the size is still about 1.5 MB per event.
  - 40 MHz times 1.5 MB is 60 TB of data per second! Even if we could handle 60 TB/s data flow, we cannot afford the disk to store that much data. Disk cost is actually the limiting factor.
  - The budget, and our computing resources, are sized such that we should be able to cope with a 1 kHz readout rate to disk.

# The ATLAS Trigger

Jonas Strandberg

- The selection from 40 MHz to 1 kHz is done by the trigger.
  - Filters out interesting event for storage and offline analysis.
- The trigger system in ATLAS is divided into two levels:
  - Level 1, fast electronic inside the detector cavern, selects one event out of ~400, reduces rate from 40 MHz → 100 kHz.
  - The High-Level Trigger (HLT), fast reconstruction of the whole event on dedicated computer farm. Reduces rate to targeted 1 kHz.
- The trigger selects on things like presence of muons, electrons or high-momentum jets.
  - The first stage (selection) in any physics analysis.

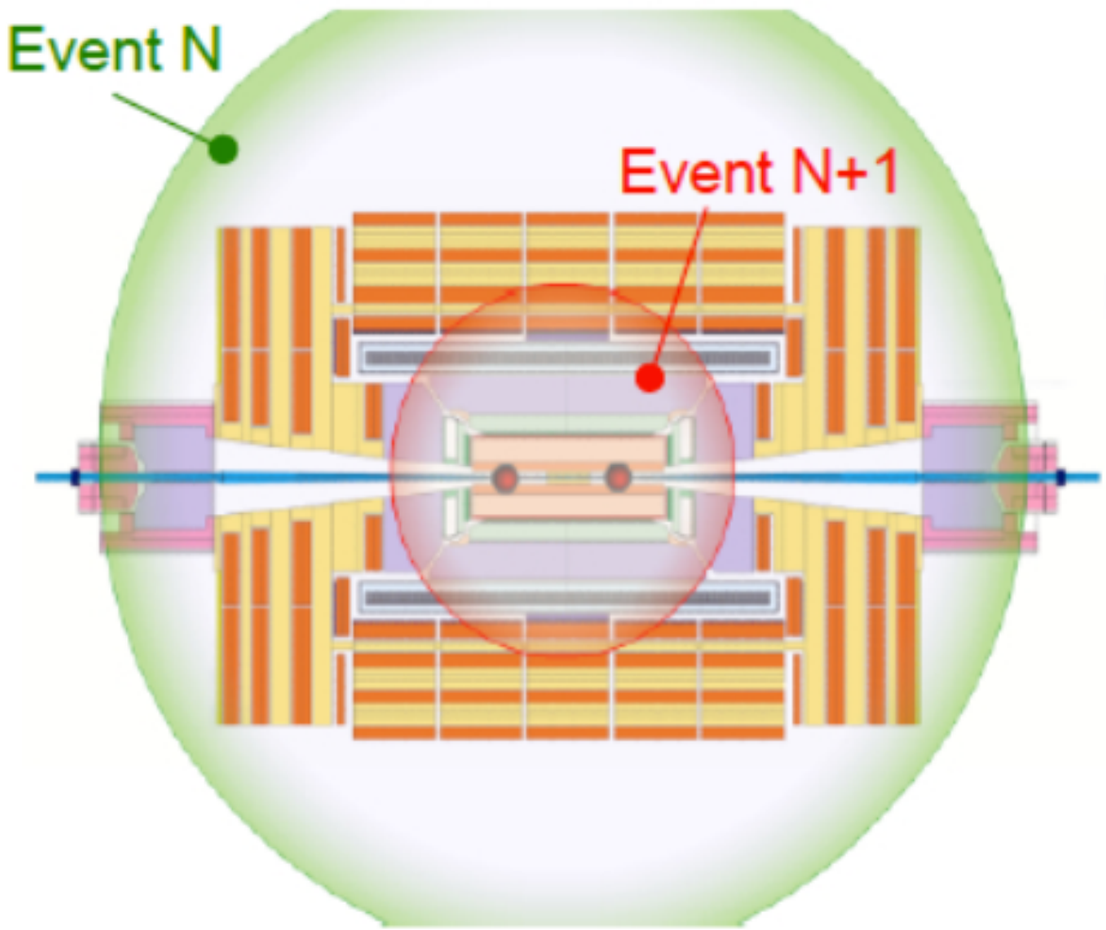

# Decision Time for the Trigger

Jonas Strandberg

- Since 25 ns is much too short to make any kind of decision, the data has to be buffered in pipelines and processed in parallel.
  - Gives the L1 trigger system about 2.5 $\mu$s to make a decision. If an event passes, it is sent along to the HLT.
  - HLT is a massive networked computer farm, parallel processing of events allows for a couple of seconds to make a decision.
  - If an event passes the HLT trigger it is sent off for storage and offline reconstruction at another computer farm.

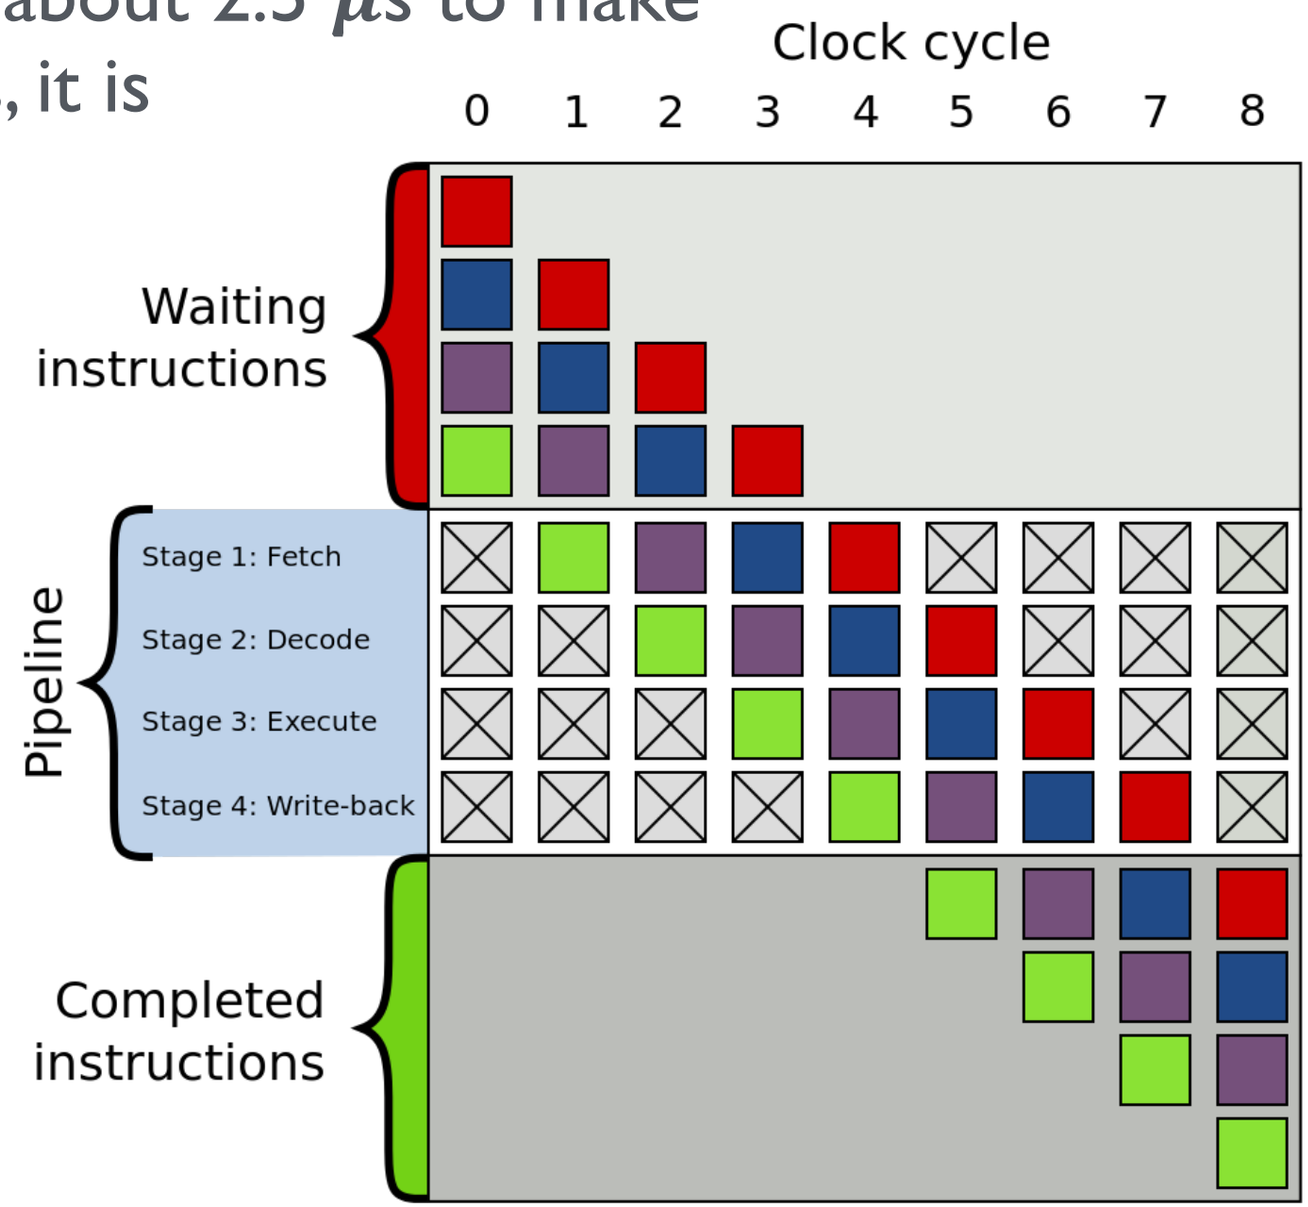

# The ATLAS Trigger

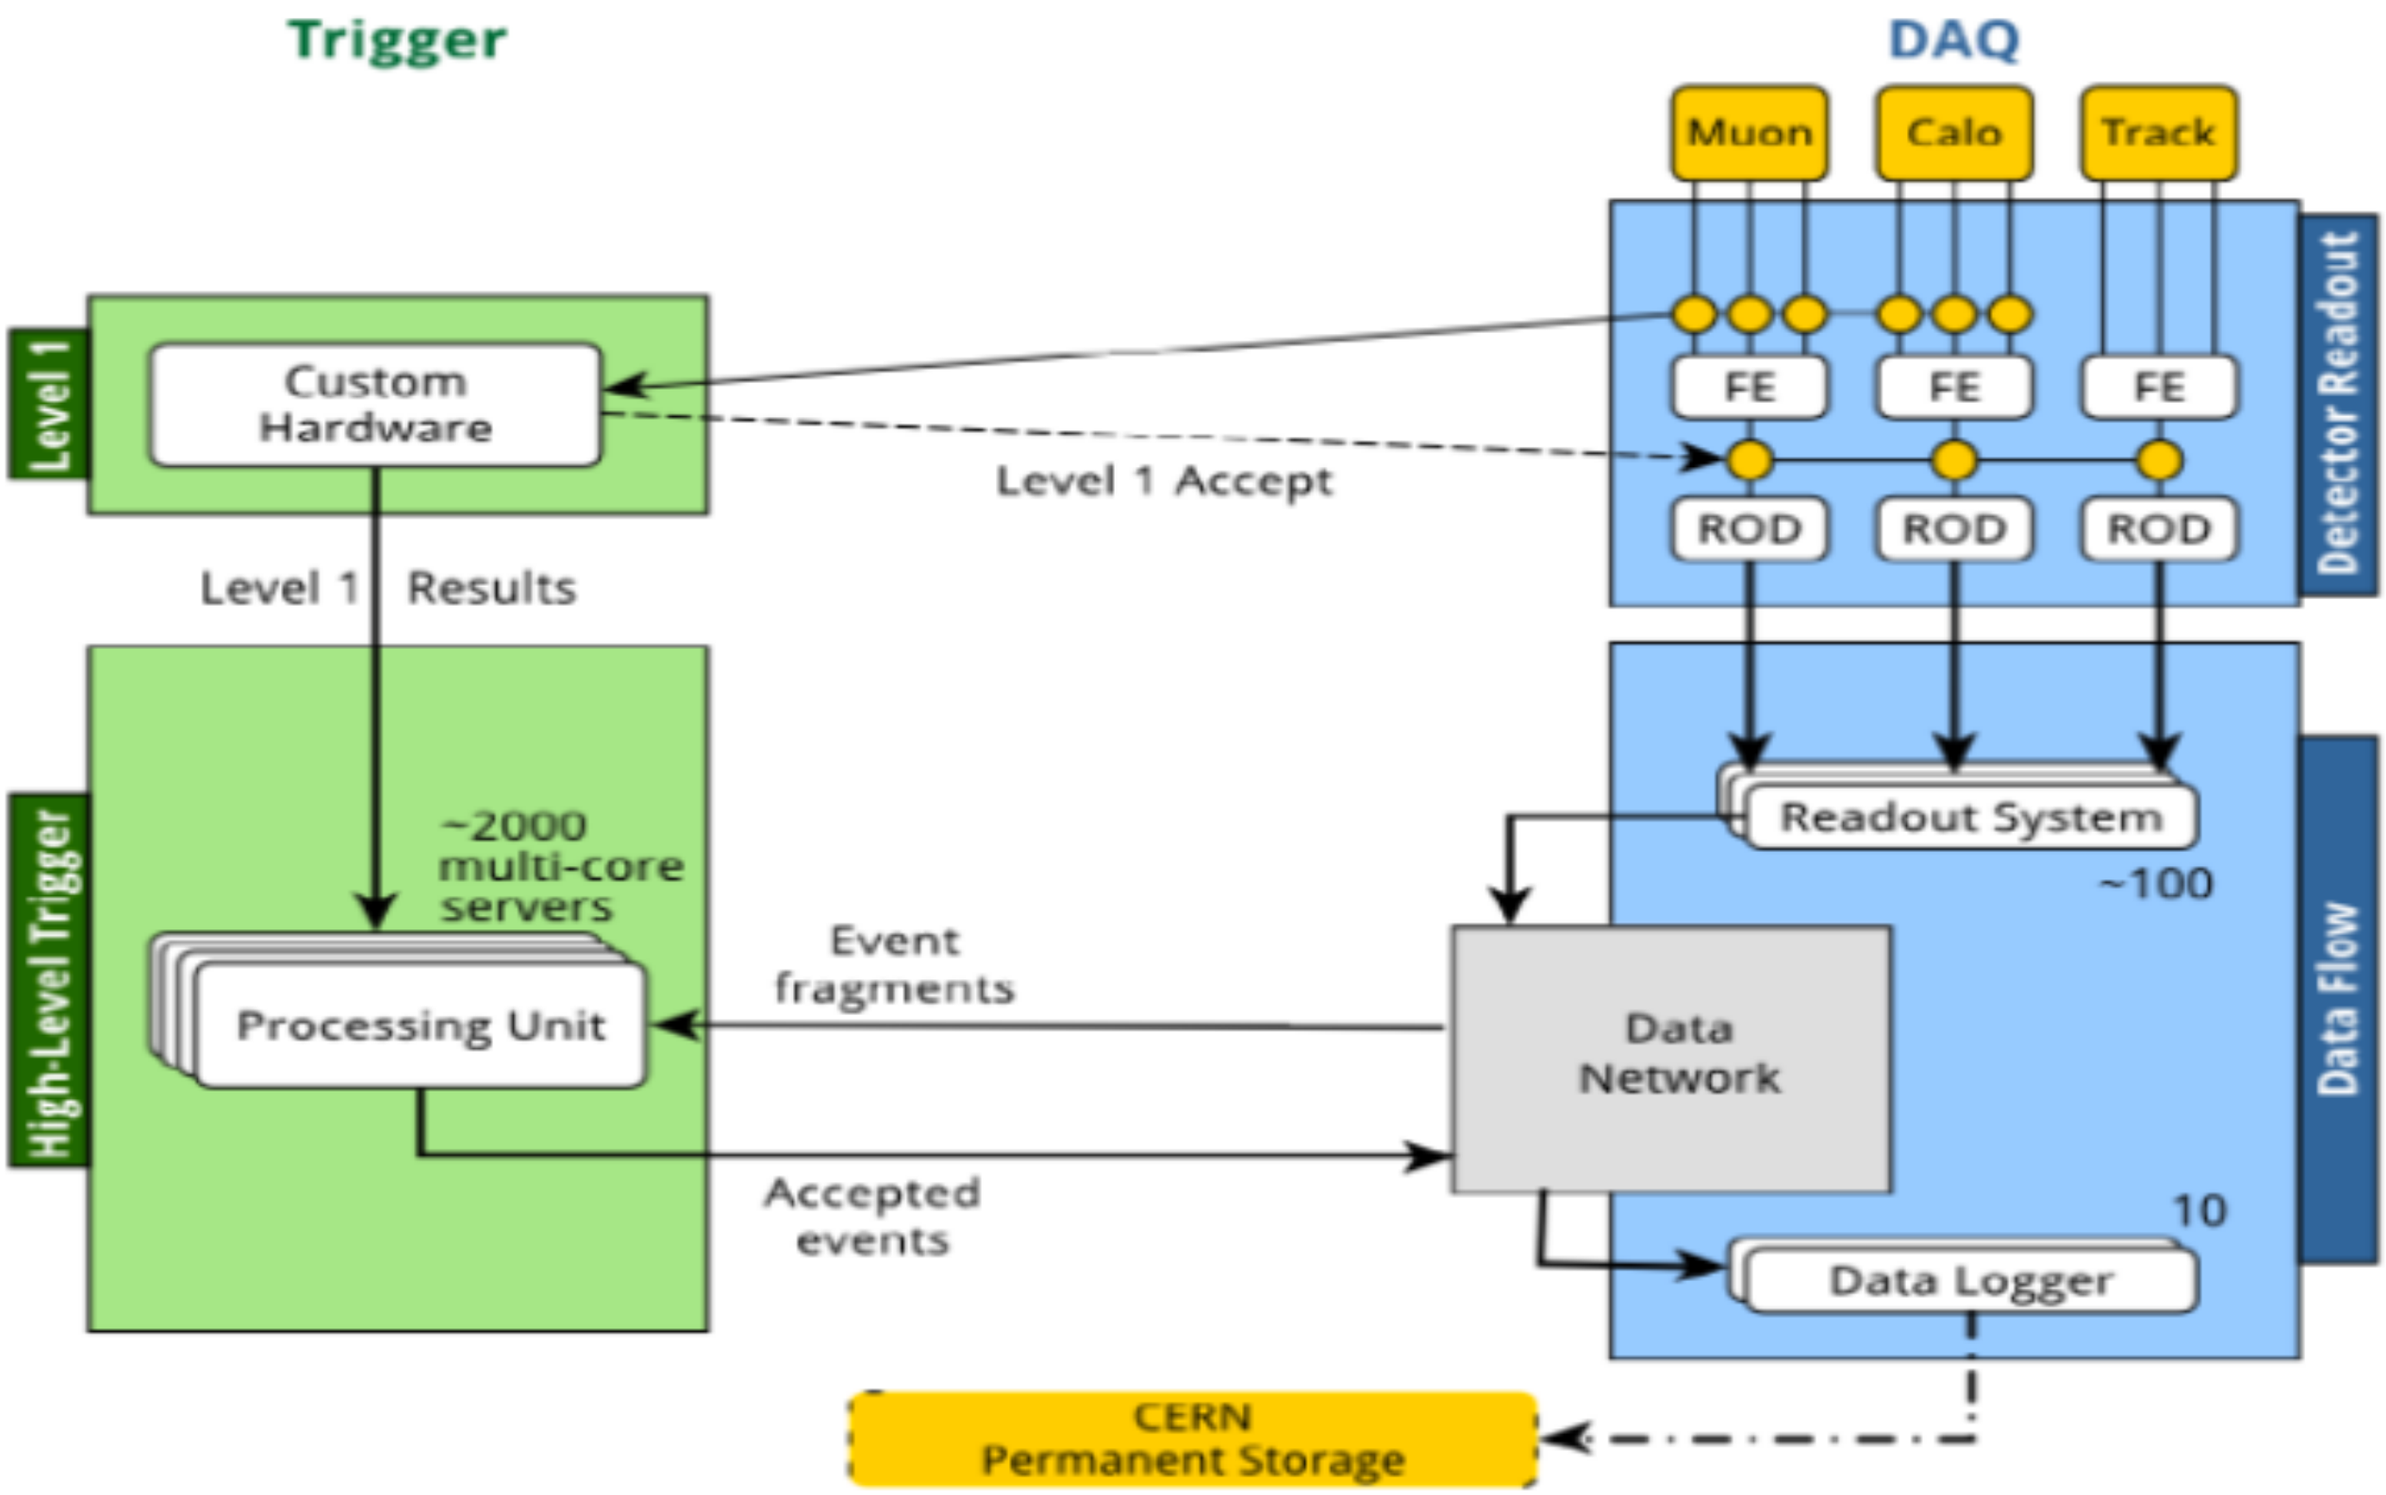

# What if Something is Missed?

Jonas Strandberg

- So we are limited to 1 kHz output rate for offline analysis:
  - To reach that we only save one part in forty thousand.
  - Aren't we worried that we throw away some interesting events?
- If we would record every event, we would get one Higgs boson every 116 days instead of one every 10 seconds!
  - We wouldn't be able to find the rare physics processes.
- We do save some small fraction of events that are "unbiased".
  - For example, we can devote 1 or 2 Hz of bandwidth to events that are just passed through the trigger, randomly sampled.
- A lot of the work with the trigger is deciding how to divide the available 1 kHz bandwidth among the different signatures.
  - E.g. 100 Hz electron triggers, 100 Hz muon triggers …

# Data Streams

- For ease of use, the trigger can also place events that are saved offline in different data *streams*, I will mention three here:
  - The *main* physics stream (called physics_Main) which collects the events for almost all physics analyses.
  - The *express* stream, which is a sampling of events of all types from the Main stream, typically about 2% of the events.
  - There are many other types of streams, for calibration, for checking the integrity of the detector and for special analyses. One example of the latter is the data scouting stream.

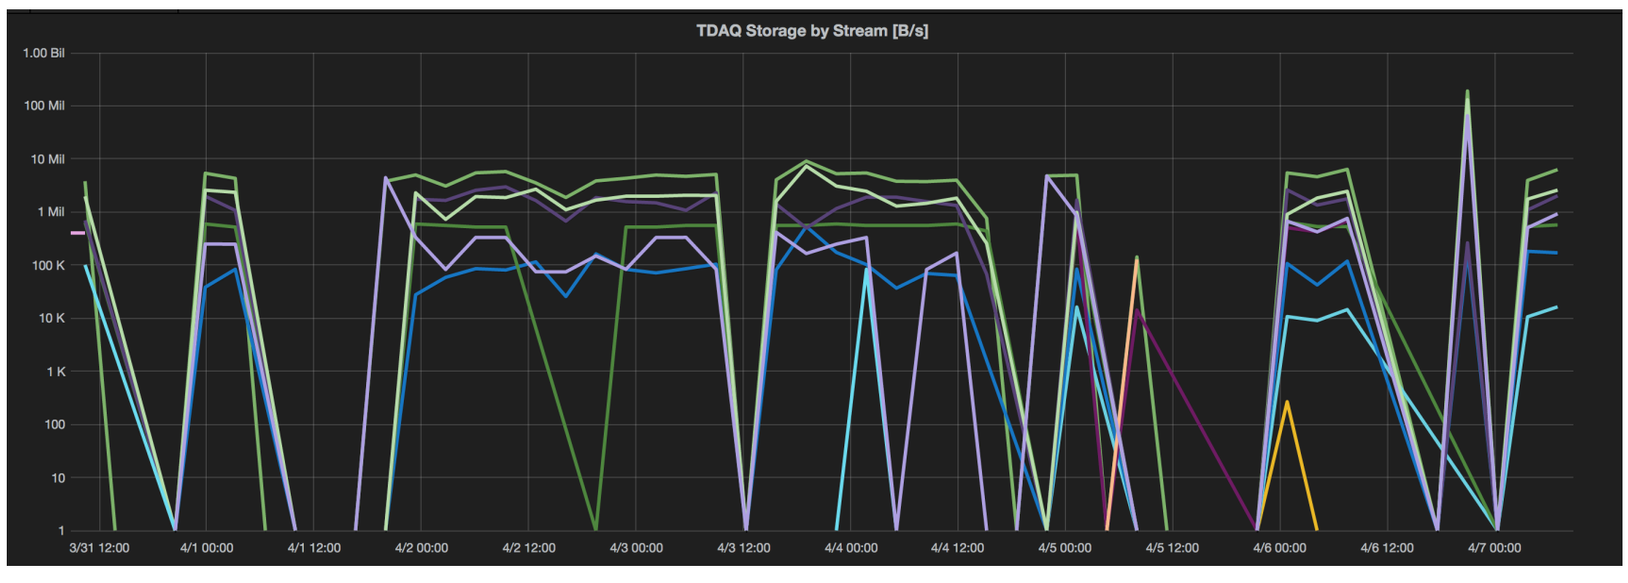

# Prompt Reconstruction

Jonas Strandberg

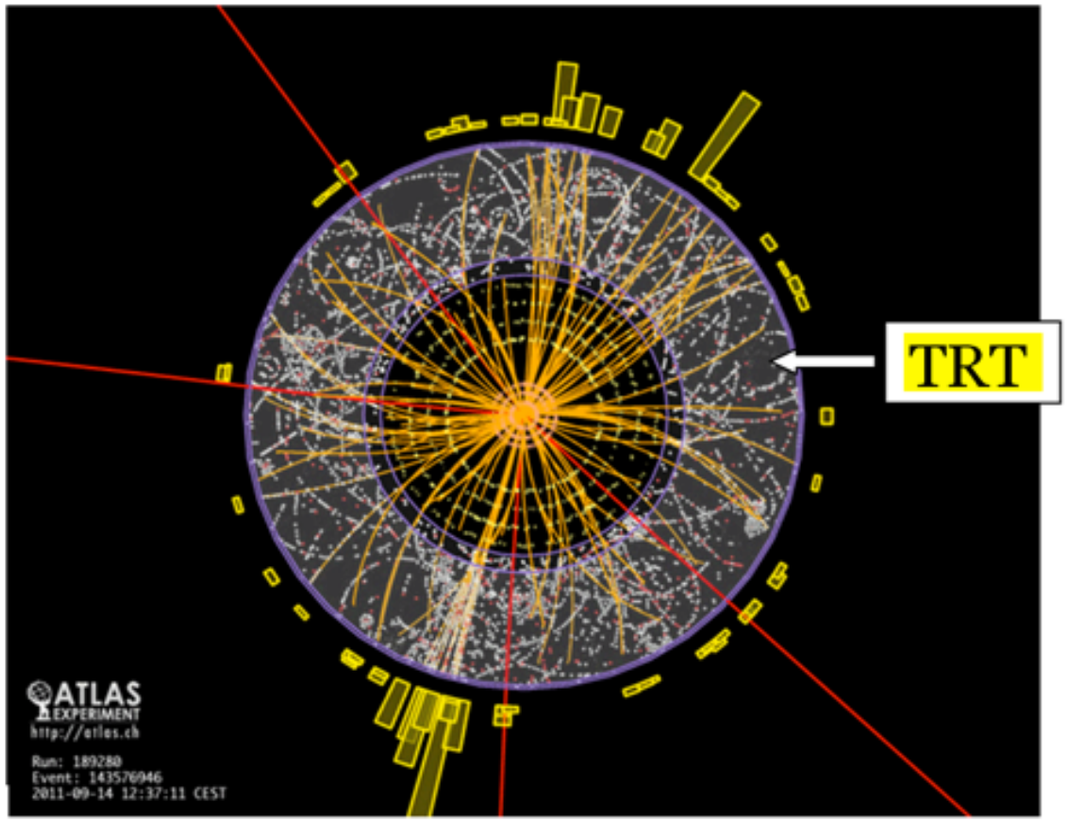


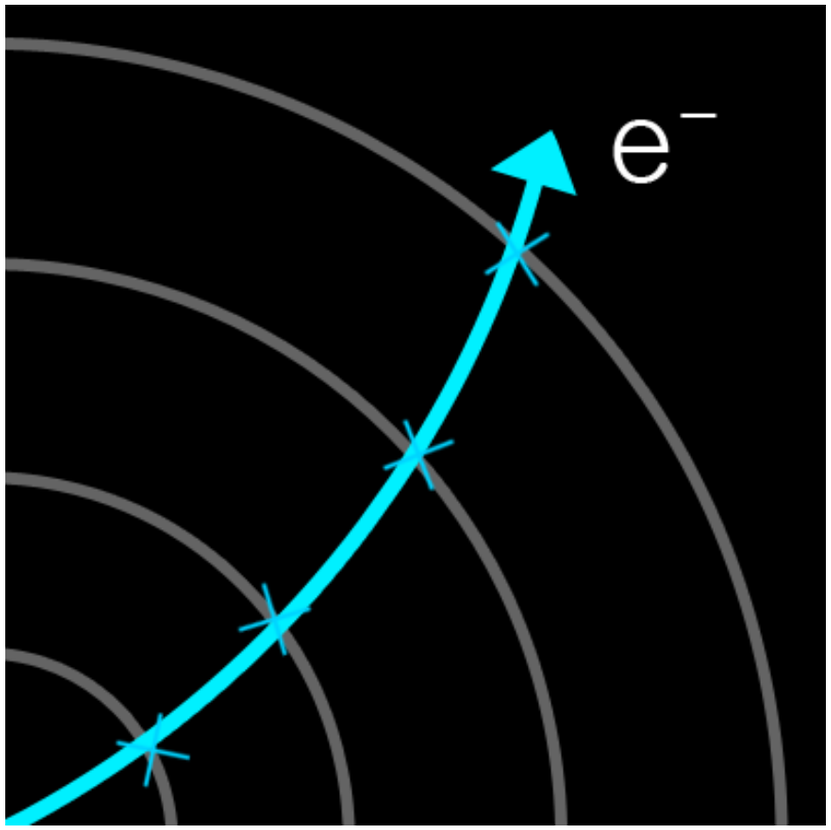


- When a data taking run is over (typically about 8 hours of data taking), the reconstruction of the express stream starts.
  - The reconstruction happens at the so called Tier-0 data centre.
  - Configuring the reconstruction, and following up on the jobs, is handled by the Prompt Reconstruction Operations Coordinators (PROCs), and the Tier-0 team.

# Calibration Loop

- After the reconstruction of the express stream, several activities commence on the express stream output:
  - Histograms of various reconstructed quantities are uploaded to a web server, where the quality of the data can be assessed.
  - If needed, a new alignment of the innermost tracker is derived.
  - The beamspot is determined (with the new alignment).

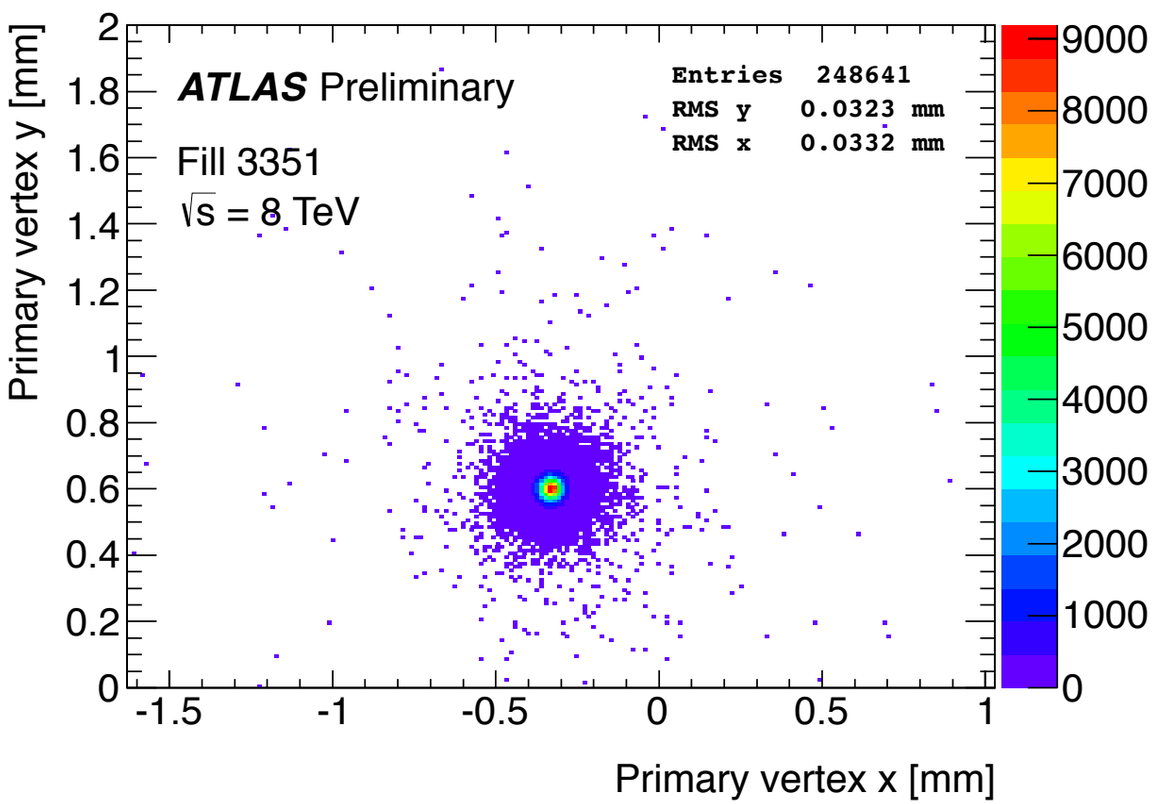


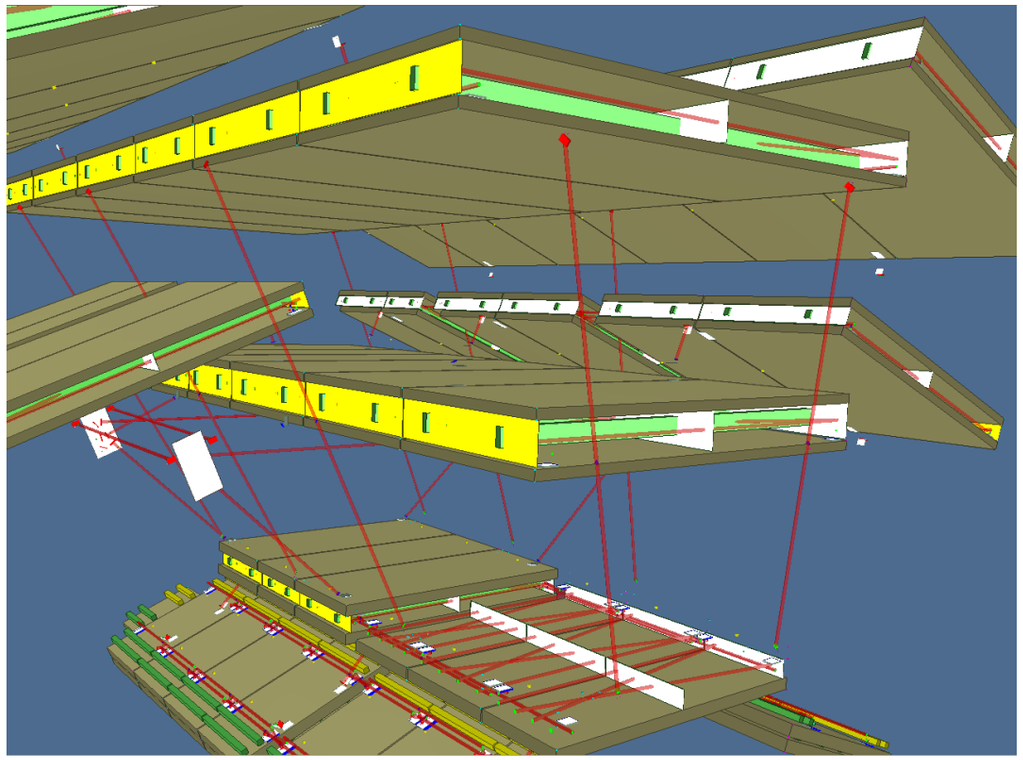

# Calibration Loop

Jonas Strandberg

- All these activities are collectively called the calibration loop, and the time-frame allocated to this is 48 hours.
  - Special data formats for checking e.g. the noise in the detector are produced, and experts create maps of noisy modules.
- After the calibration loop has finished, the so called *bulk* reconstruction starts, on the Main stream.
  - Small detector defects (noisy modules for example) identified in the calibration loop are masked in the bulk reconstruction.

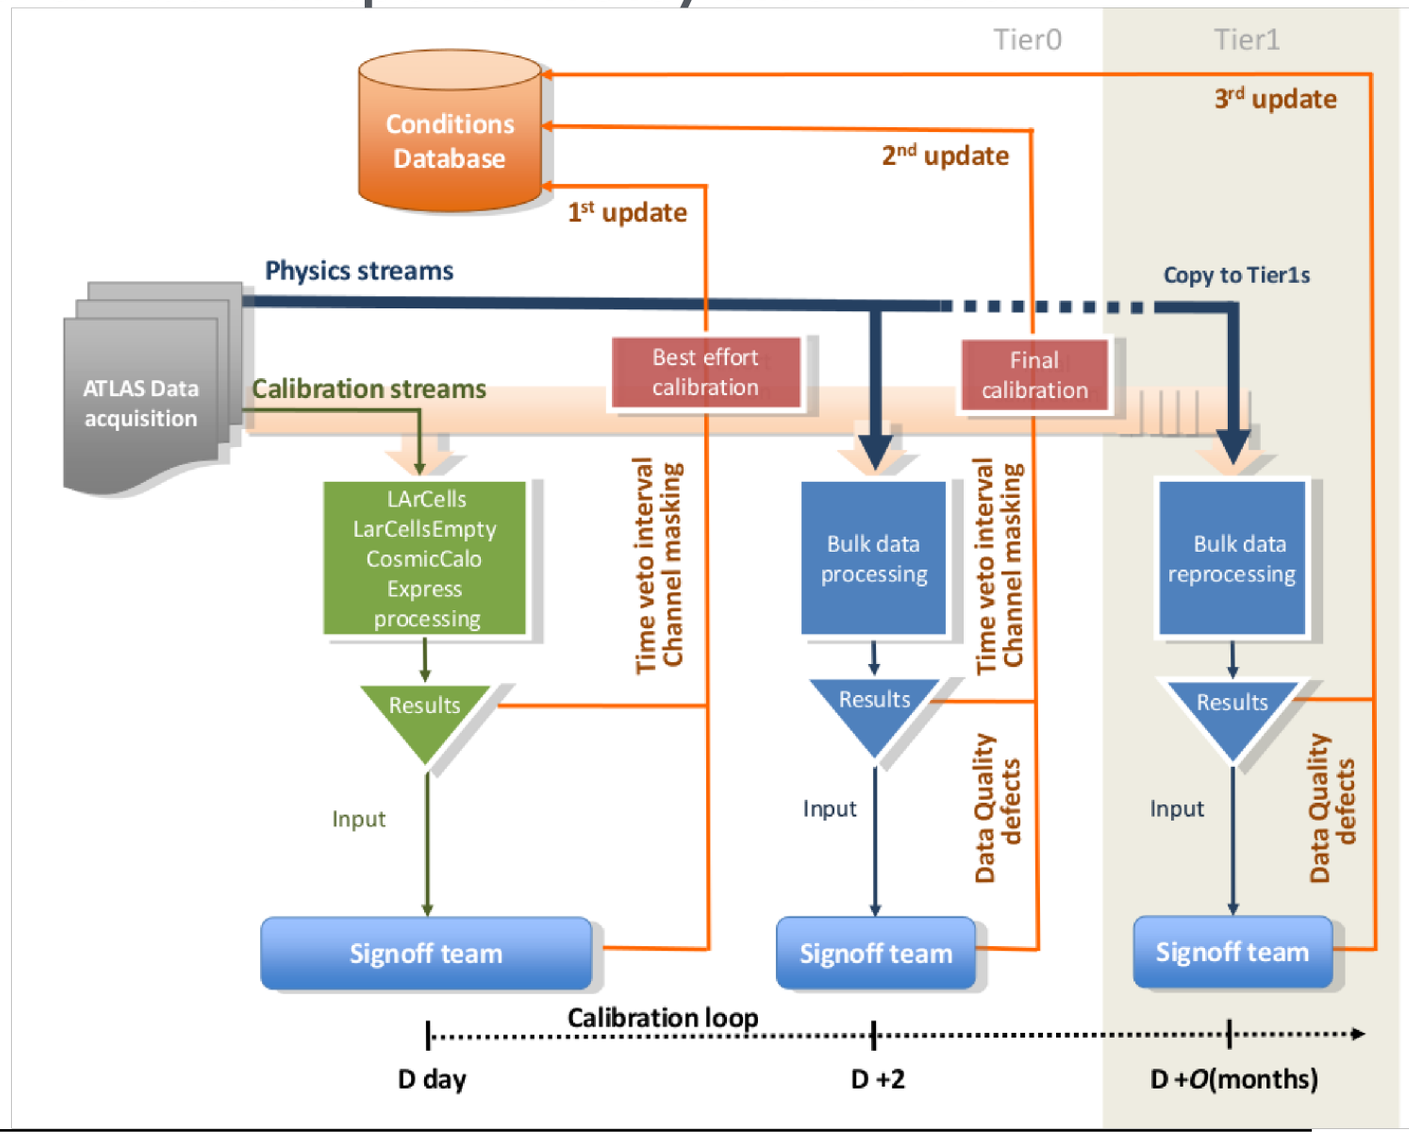

# Calibration Loop

Jonas Strandberg

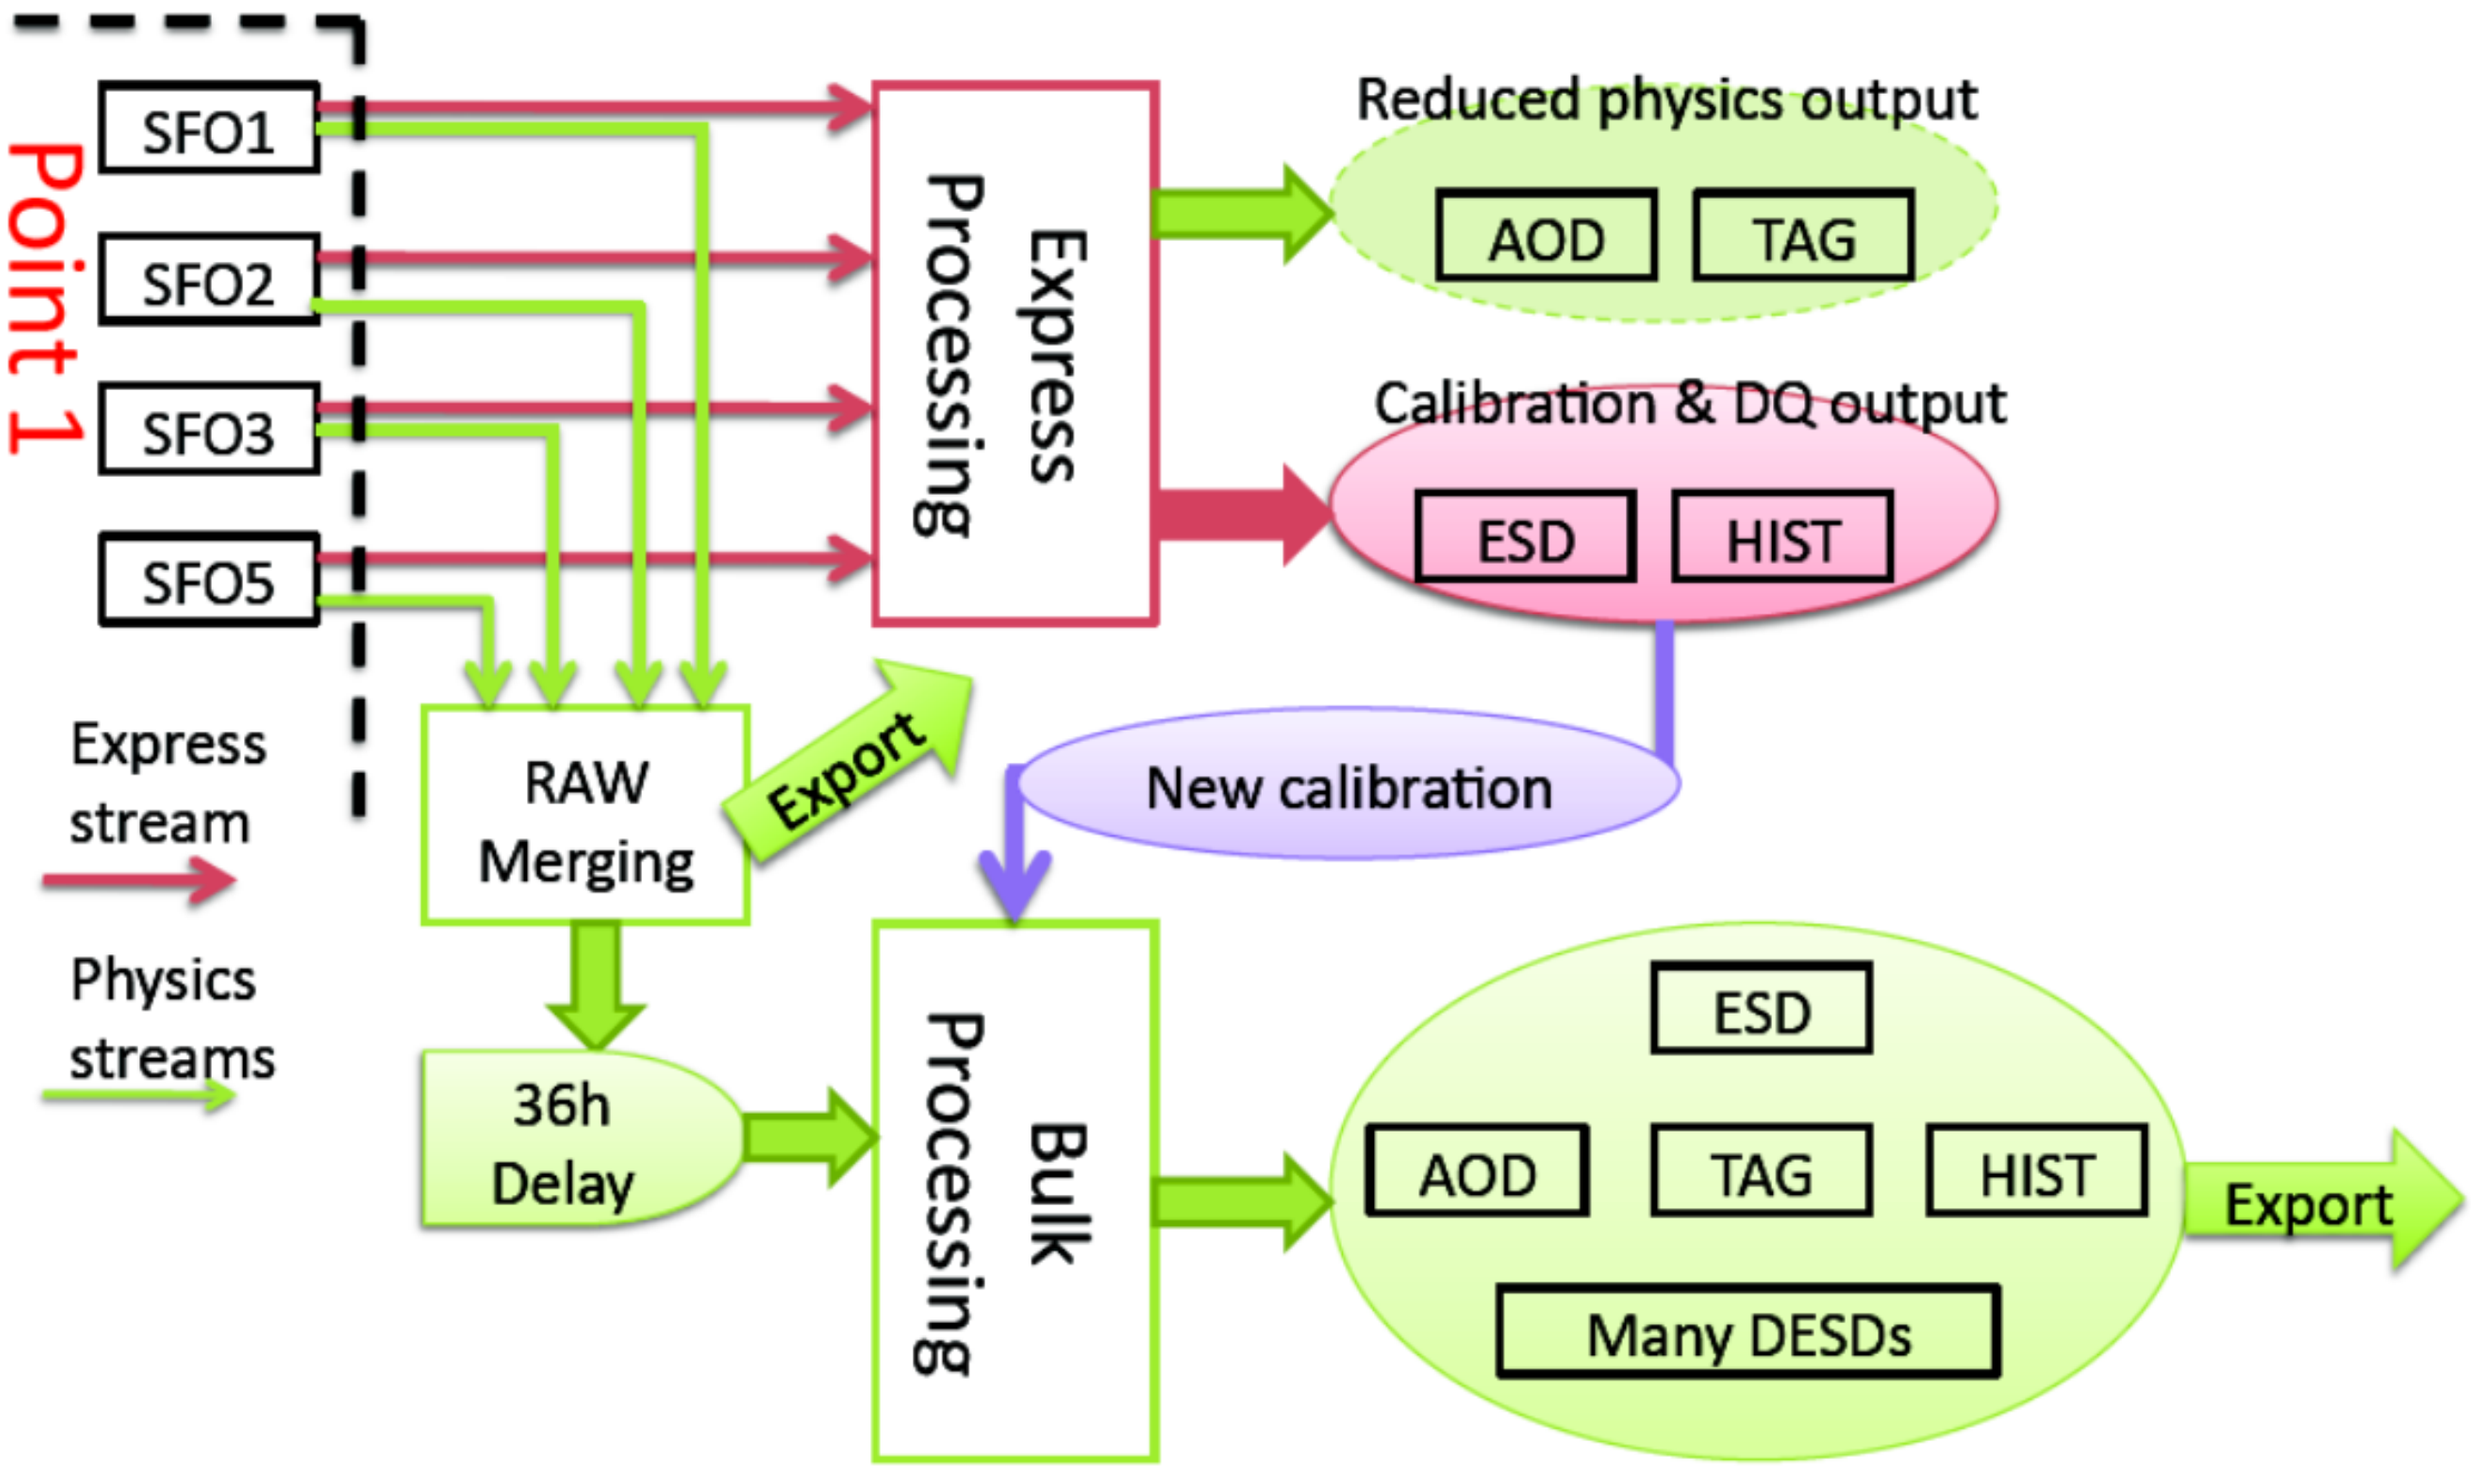

# Data Quality

- The Data Quality group checks the histograms produced in the express stream processing and helps identify problems.
- In the bulk reconstruction, an additional set of histograms are produced with the full statistics of the main stream.
  - This includes high-level quantities such as particle masses and other things that need the full statistics to be monitored.
  - A report is made by each (sub-)detector community in ATLAS, as well as by the combined performance group contacts.

| ATLAS pp 25ns run: August-November 2015 | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Inner Tracker | | | Calorimeters | | Muon Spectrometer | | | | Magnets | |
| Pixel | SCT | TRT | LAr | Tile | MDT | RPC | CSC | TGC | Solenoid | Toroid |
| 93.5 | 99.4 | 98.3 | 99.4 | 100 | 100 | 100 | 100 | 100 | 100 | 97.8 |
| **All Good for physics: 87.1% (3.2 $fb^{-1}$)** | | | | | | | | | | |

Luminosity weighted relative detector uptime and good data quality (DQ) efficiencies (in %) during stable beam in pp collisions with 25ns bunch spacing at $\sqrt{s}$=13 TeV between August-November 2015, corresponding to an integrated luminosity of 3.7 $fb^{-1}$. The lower DQ efficiency in the Pixel detector is due to the IBL being turned off for two runs, corresponding to 0.2 $fb^{-1}$. Analyses that don't rely on the IBL can use those runs and thus use 3.4 $fb^{-1}$ with a corresponding DQ efficiency of 93.1%.

# Good Runs List

Jonas Strandberg

- The lowest granularity of data which can be flagged as good or bad is one luminosity block, corresponding to ~60 s of data.
  - The bulk reconstruction is used to flag data which is unusable for physics analysis because of detector or reconstruction issues.
  - The full set of good/bad decisions is called a Good Runs List (GRL). Every analysis has to apply a GRL to their dataset.
- Once you have the list of luminosity blocks (LB) that are used in the analysis, you also need to know the luminosity per LB.
  - The instantaneous luminosity varies during the run, we consider it to be constant on the timescale of 60 s.
  - The integral of the luminosity is calculated as a sum of LBs.

# Luminosity

- The instantaneous luminosity is measured by several detectors, some which are dedicated to the luminosity measurement.
  - For example the LUCID detector.
- All luminometers are calibrated to the absolute scale in special Van-der-Meer runs, that determines the beam parameters.

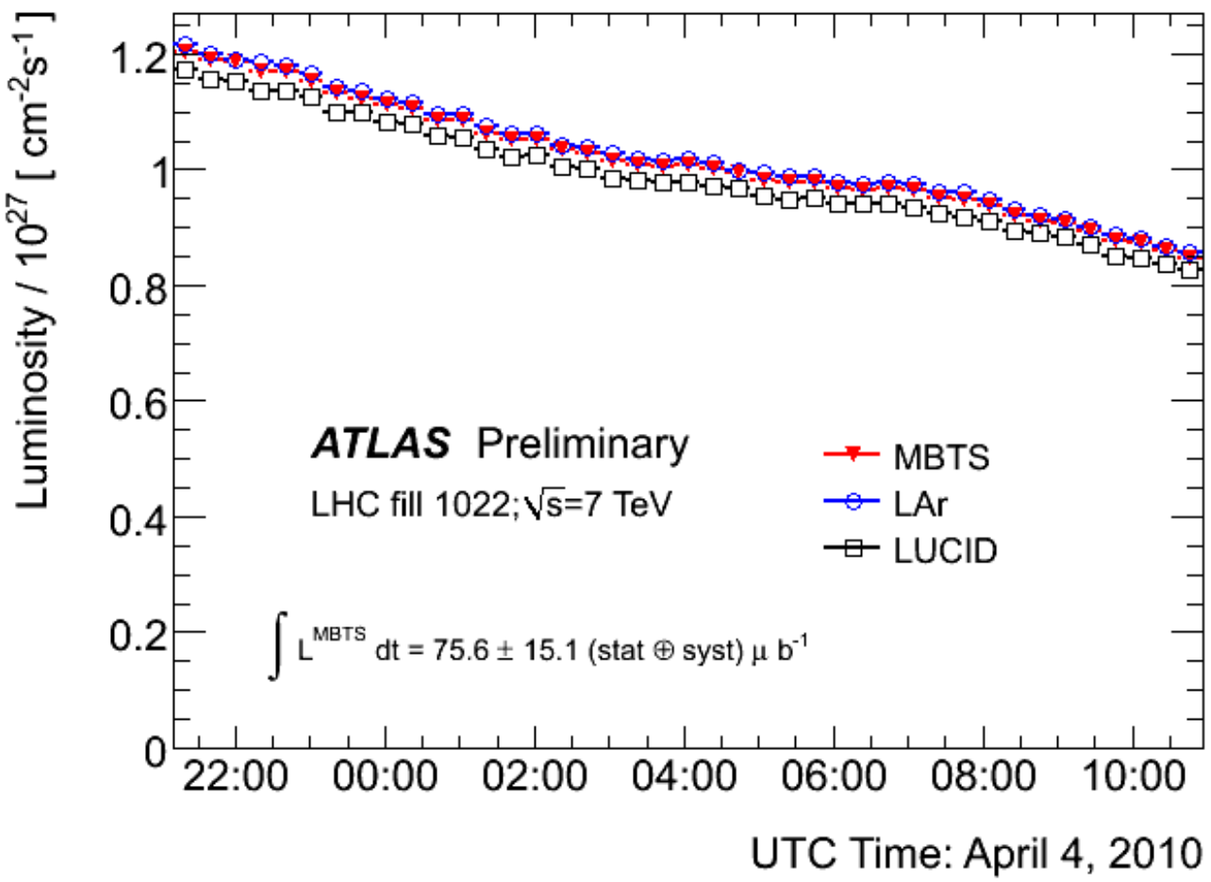


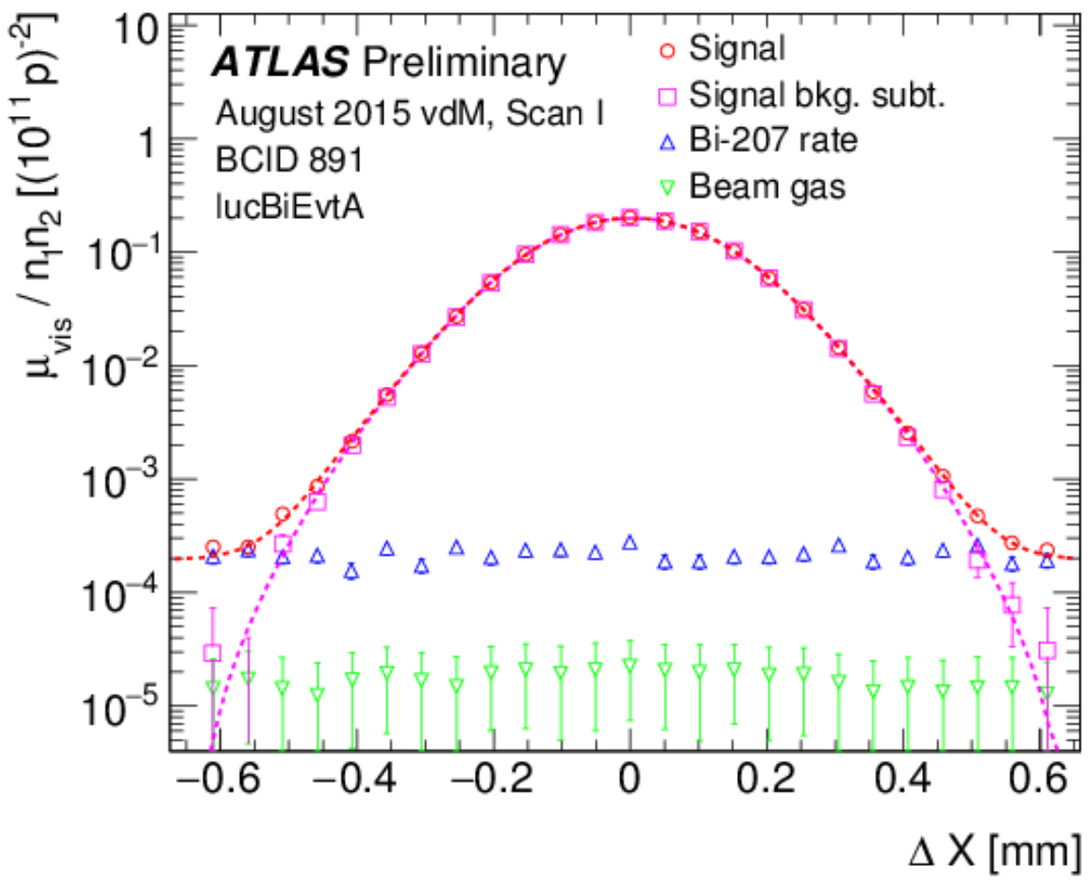

# Luminosity 2015

Jonas Strandberg

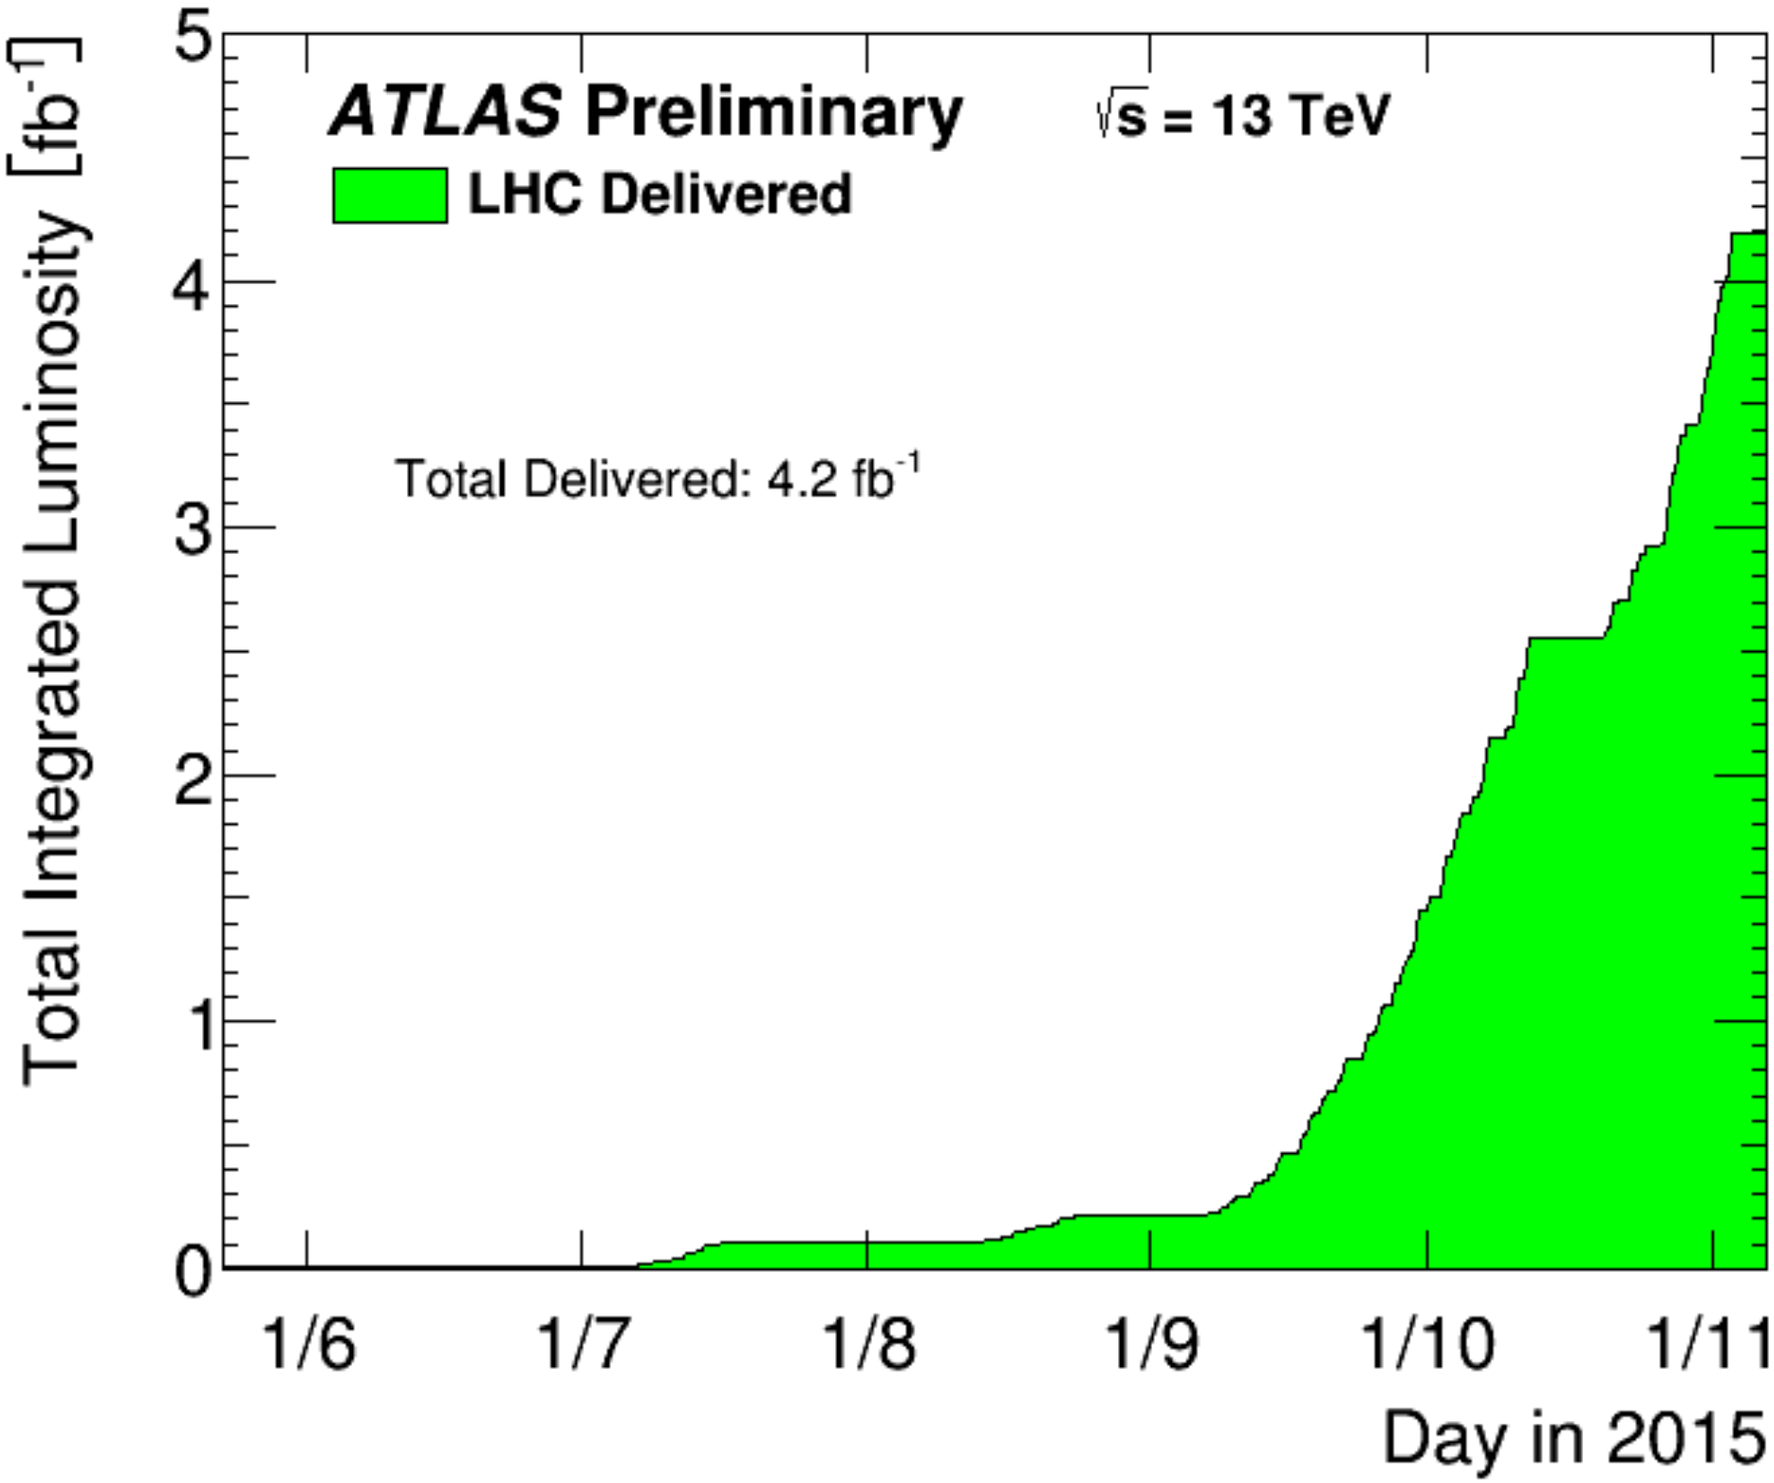

# Luminosity 2015

Jonas Strandberg

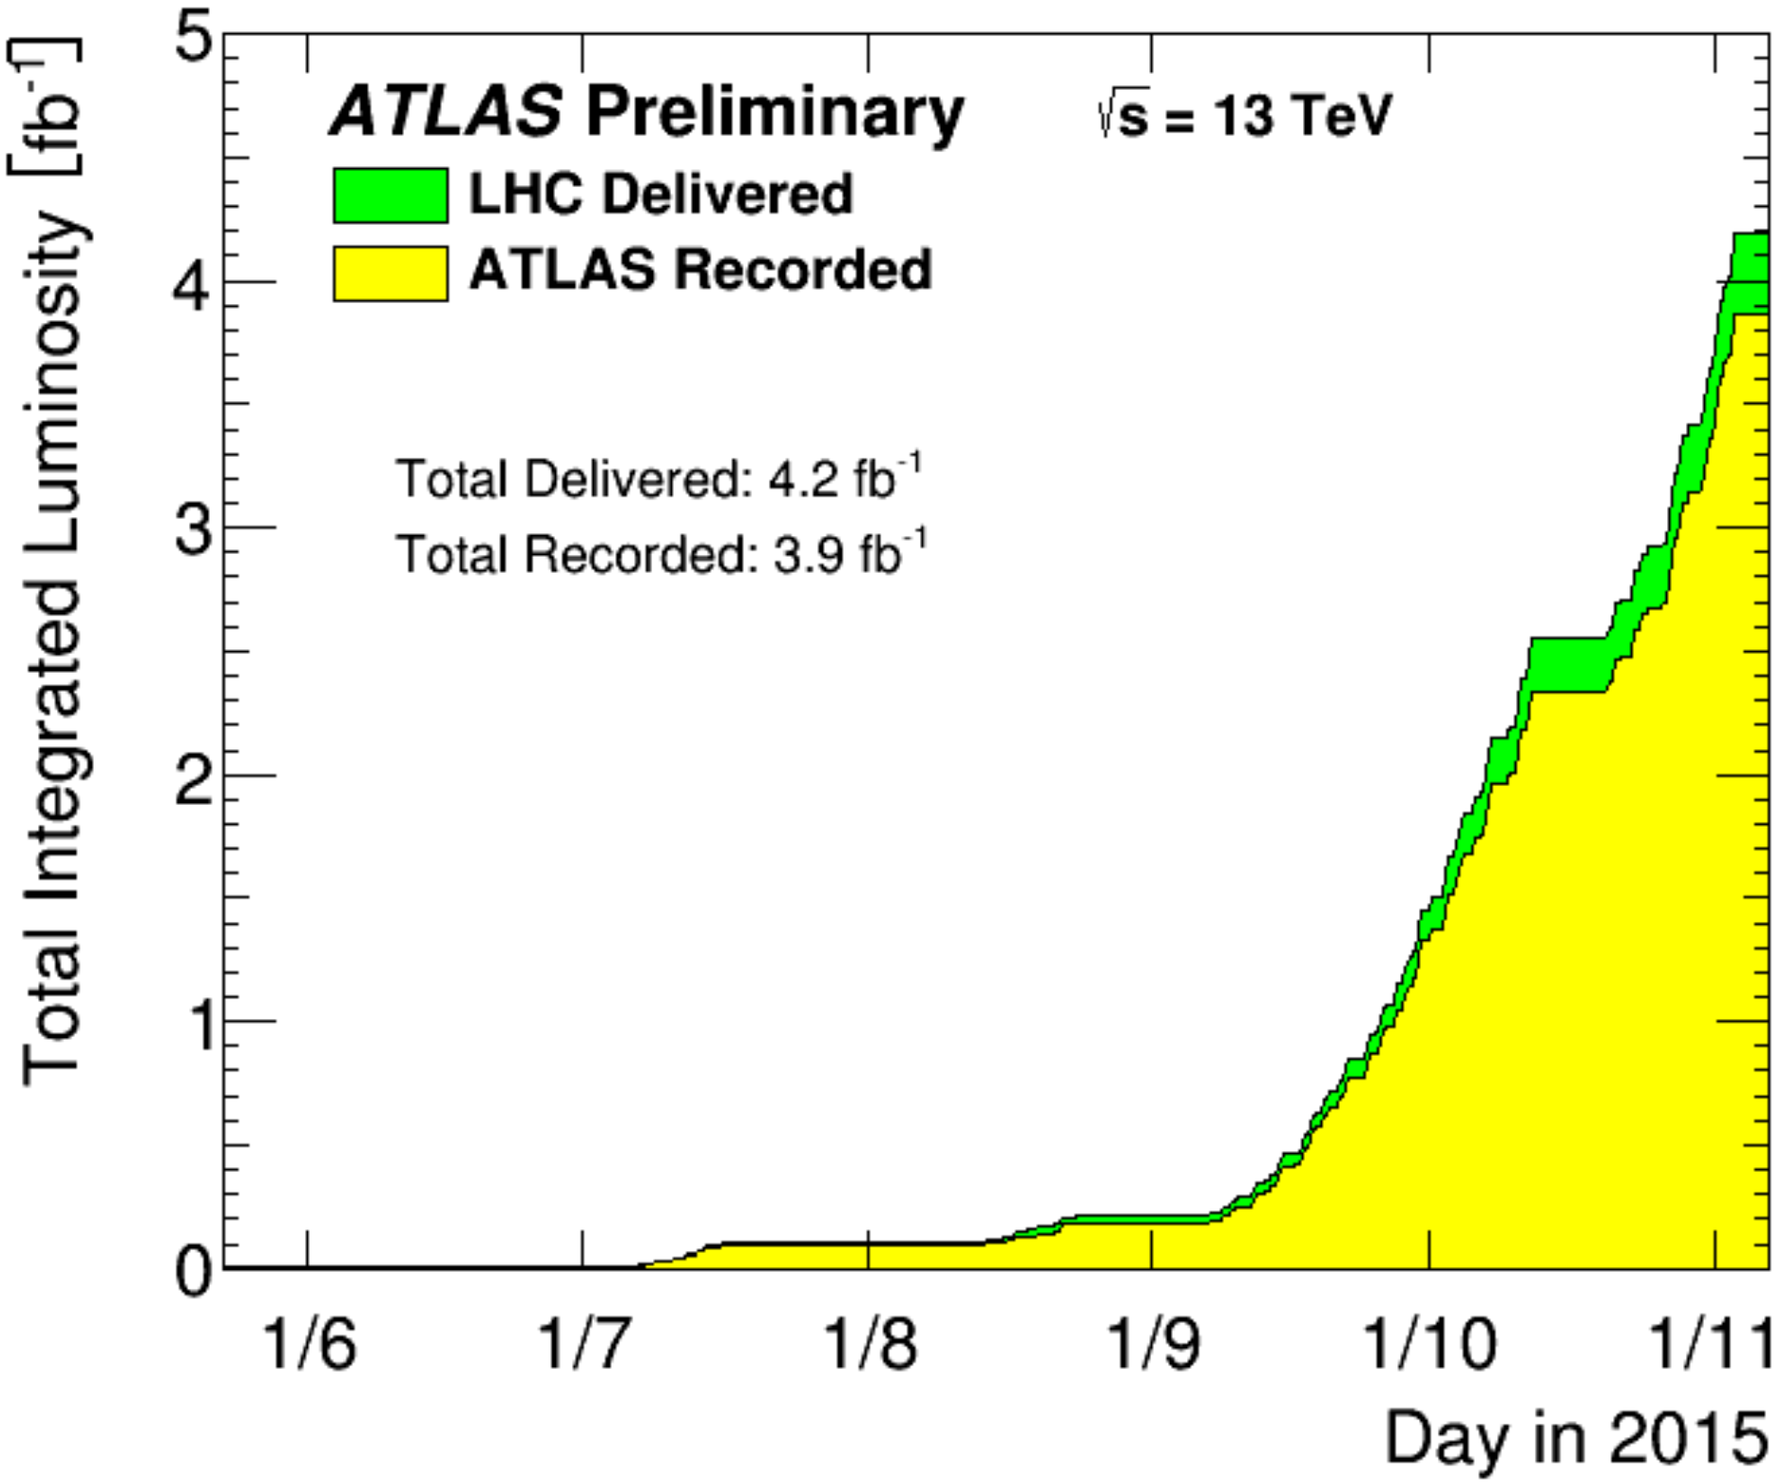

# Luminosity 2015

Jonas Strandberg

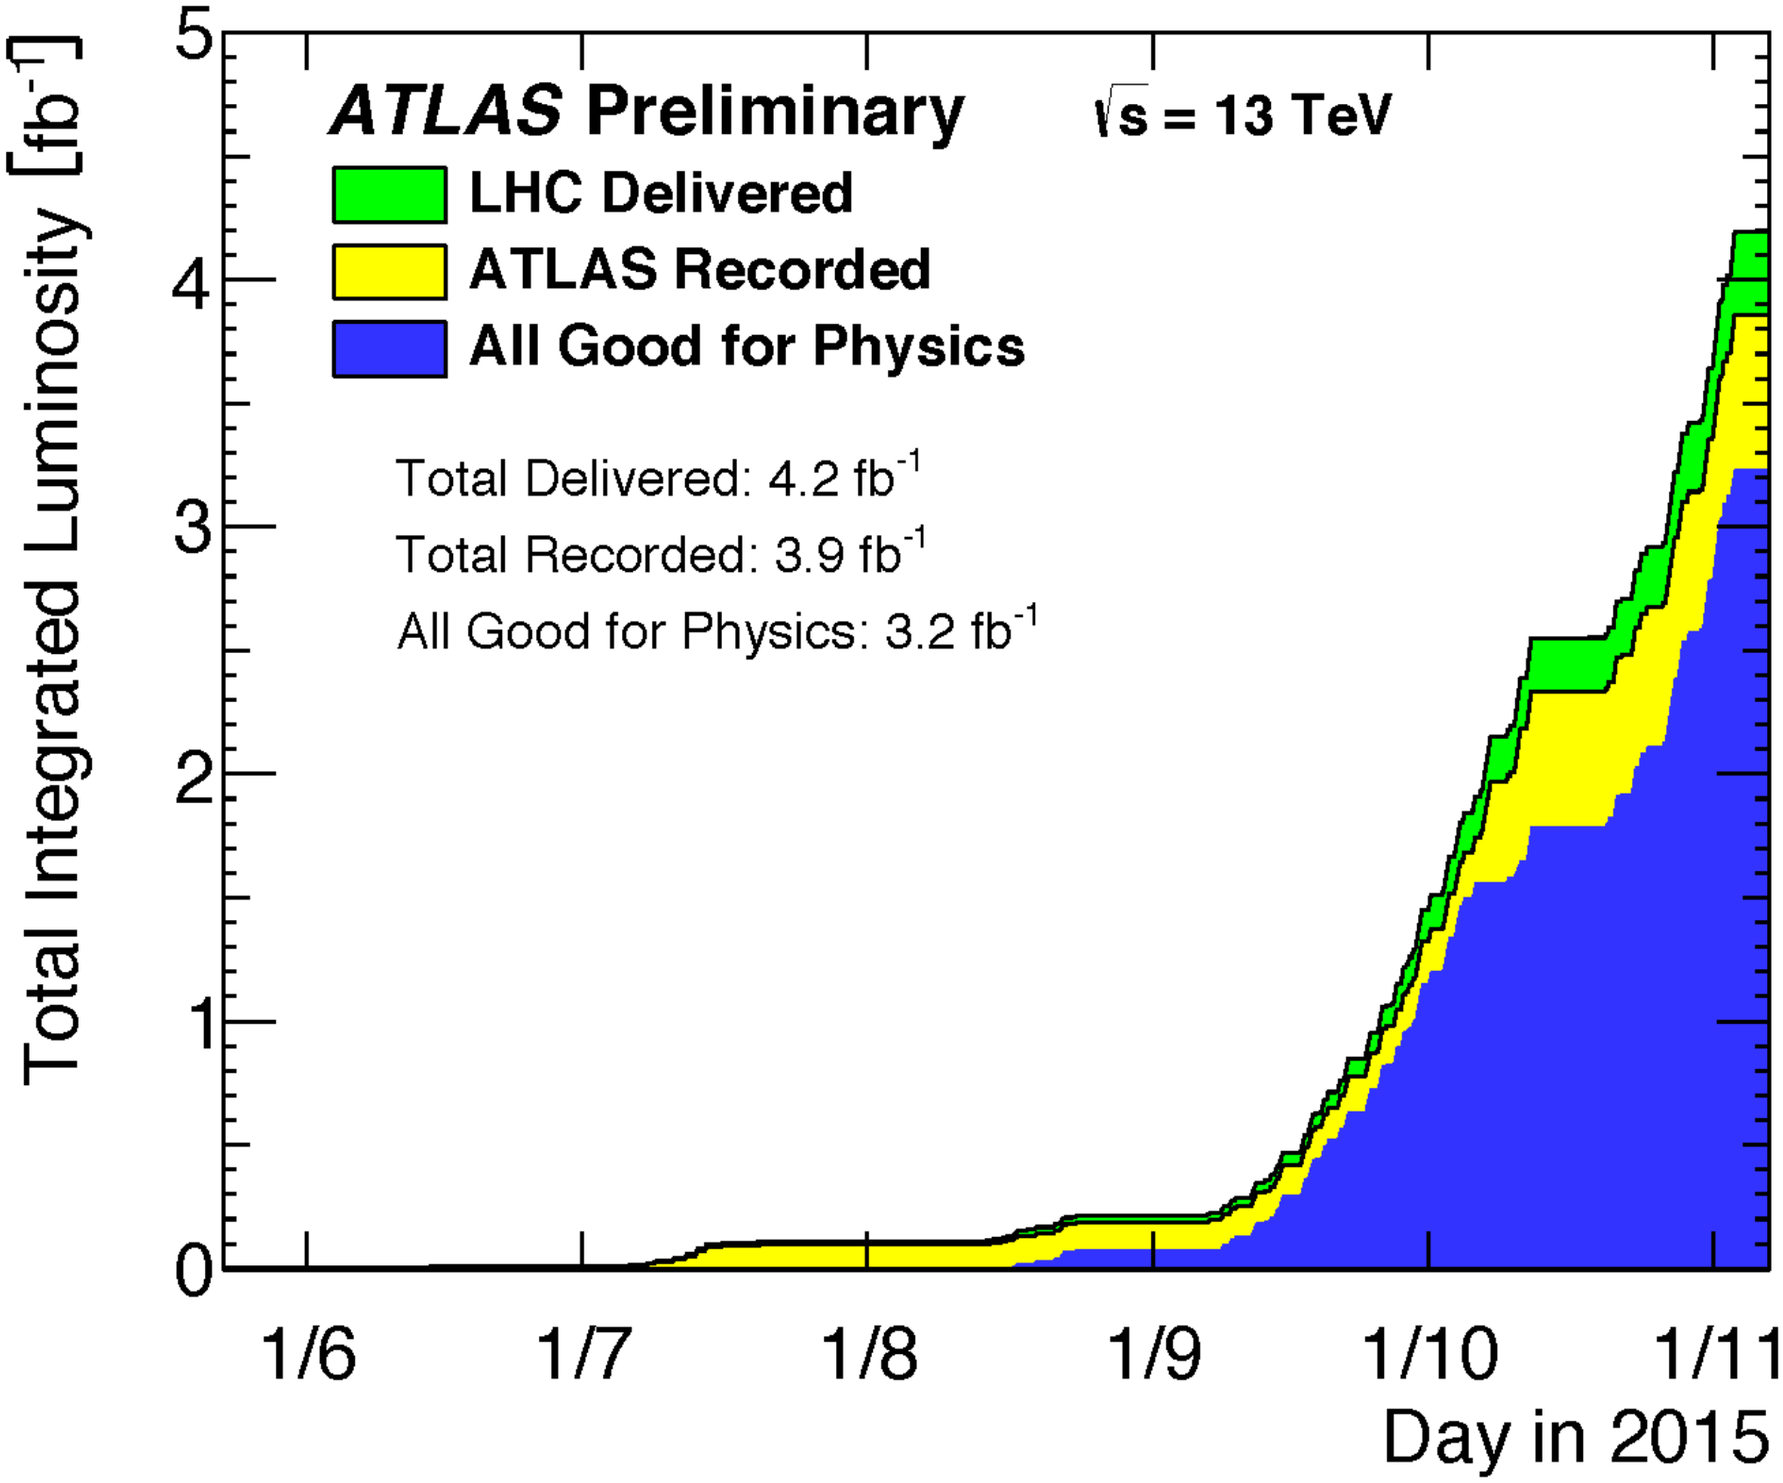

# Distributed Computing

Jonas Strandberg

- After the Tier-0 bulk reconstruction has finished, the data gets distributed on the Grid.
  - First to the Tier-1 centres, then further on to the Tier-2 and Tier-3 centres.
  - At this point the so called *derived* data formats are produced (DAODs), which are then inputs to the various analyses.
  - Sometimes analyses also make final, small, ntuples from the DAODs.

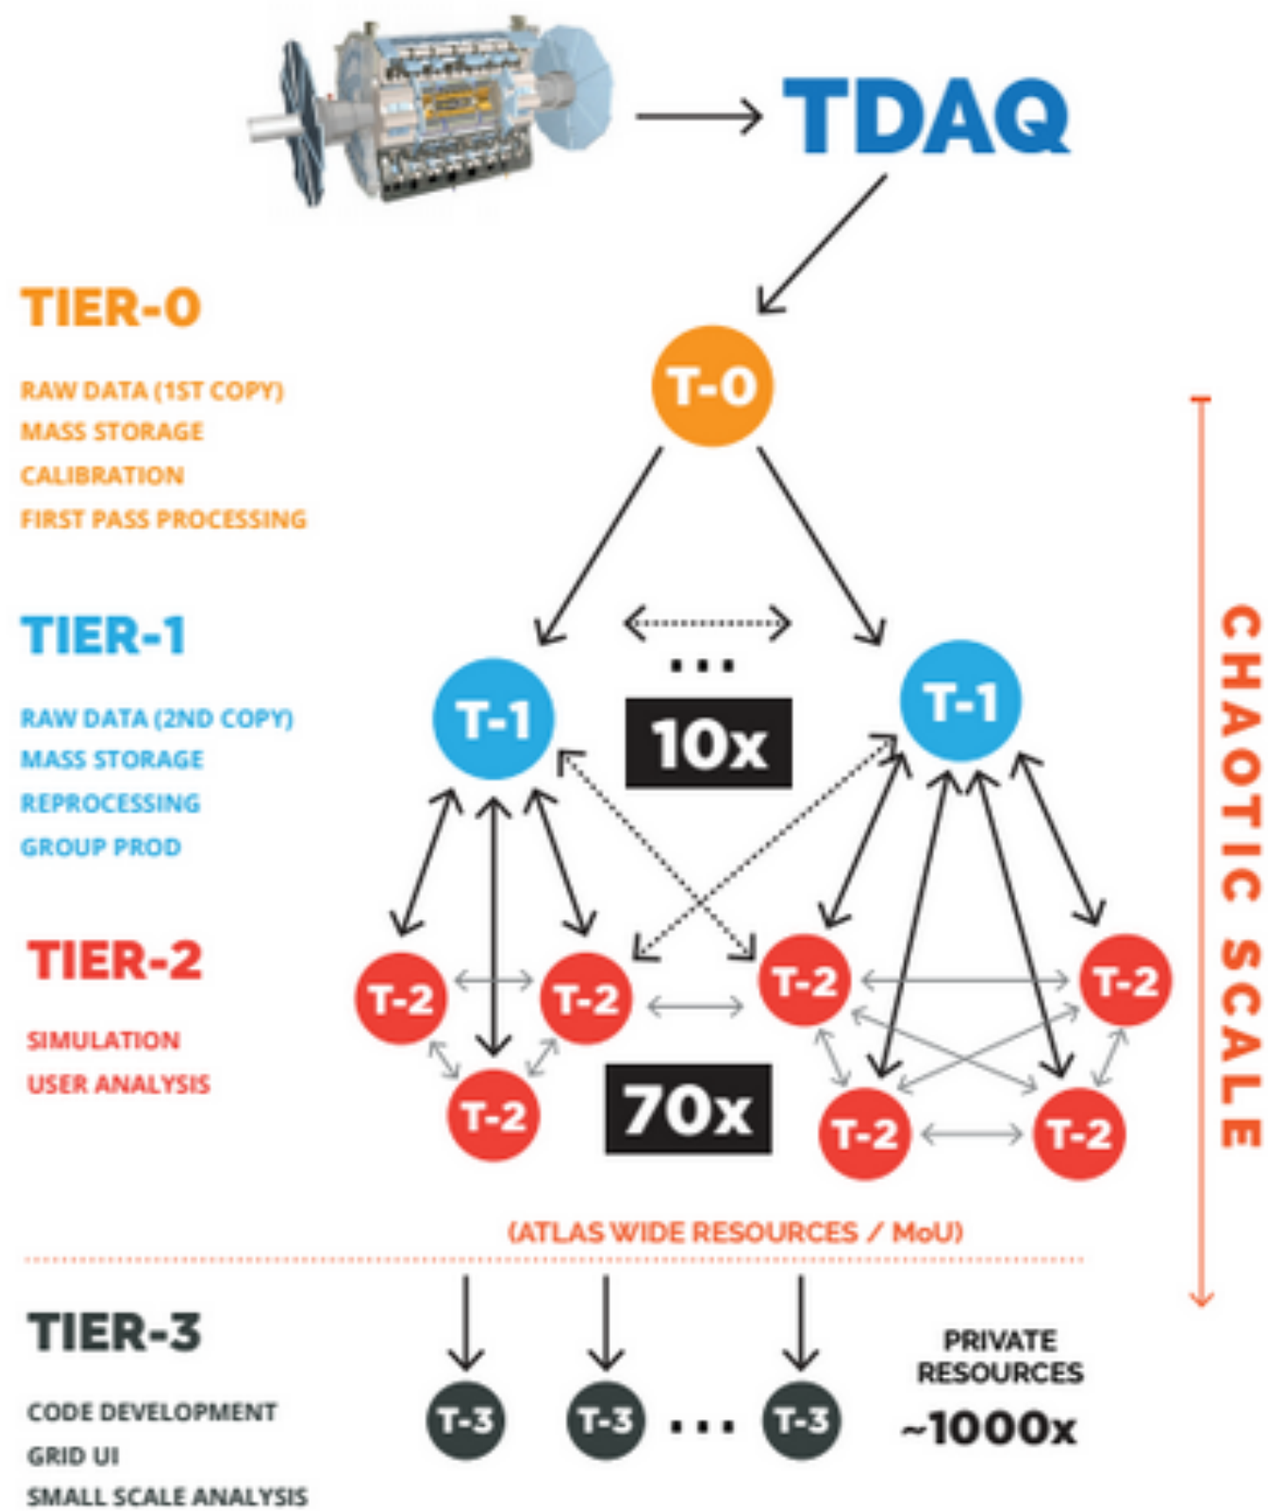

# Distributed Computing

Jonas Strandberg

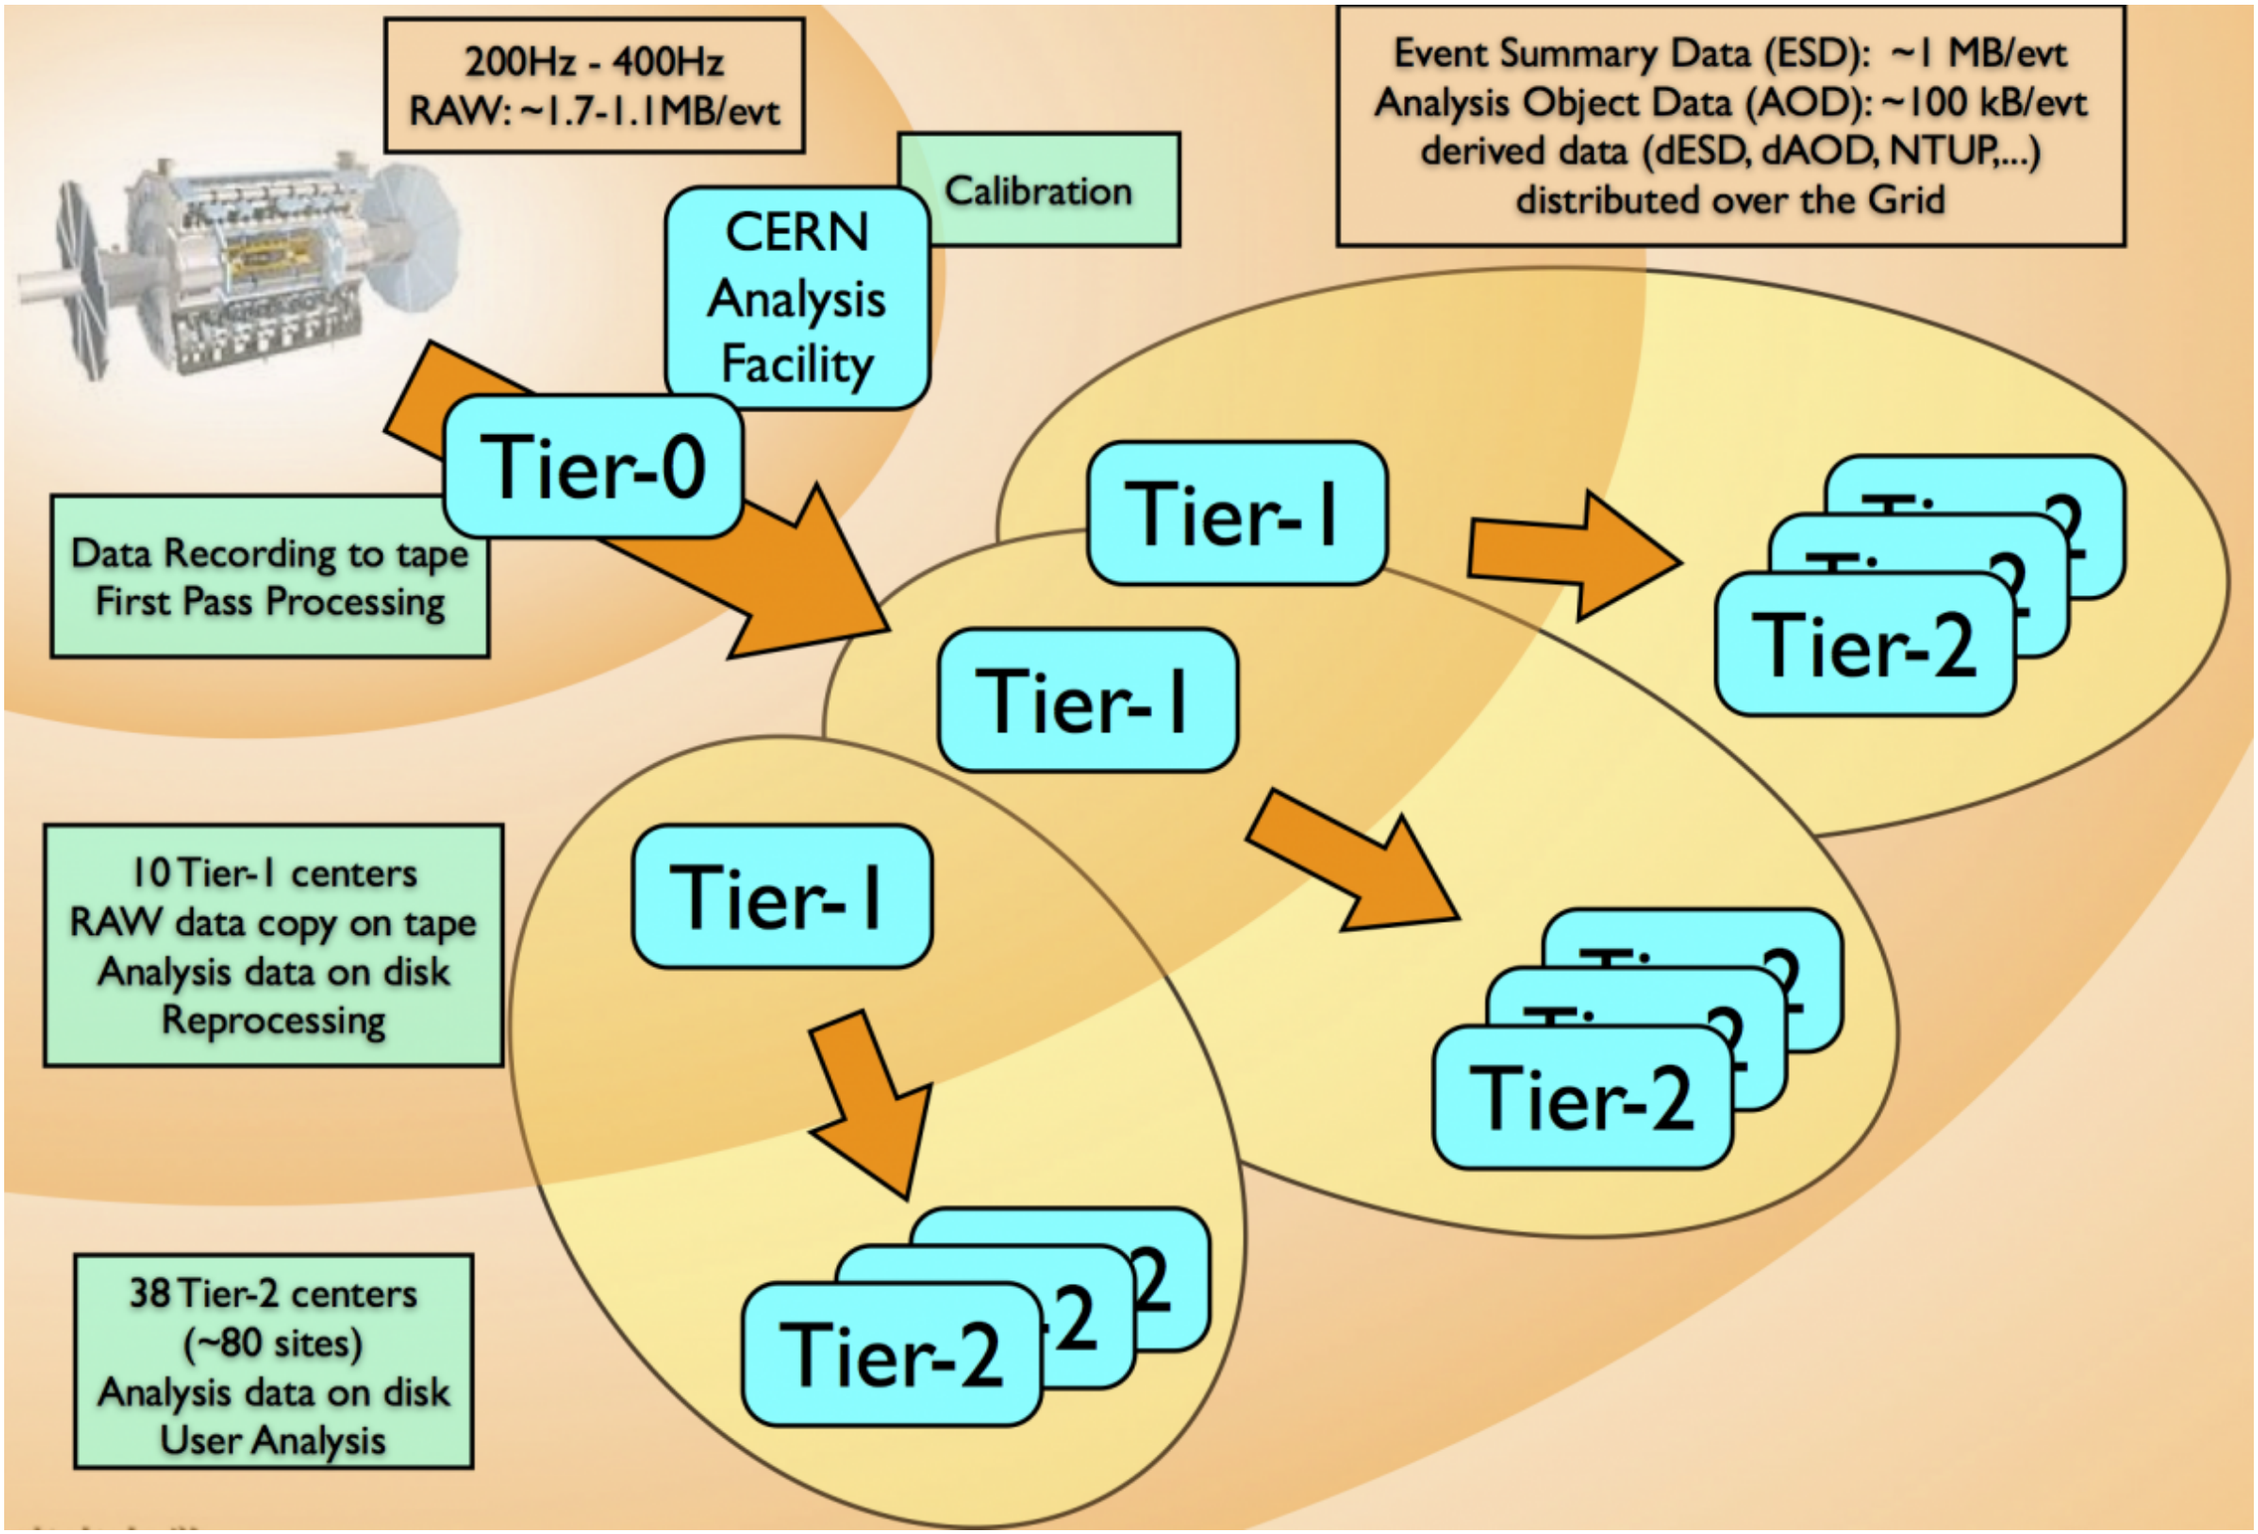

# Reprocessing

Jonas Strandberg

- To make analysis possible, we keep the performance of the software stable during the data taking for one year.
  - All developments are collected into the next software release.
- At the end of the year, we typically *reprocess* all our data.
  - Reconstruct all the RAW data again with the new release.
  - This is also an opportunity to recover some data, that were lost due to fixable problems.
- Maintaining and validating our releases is a major effort!


# Summary and Conclusions

Jonas Strandberg

- Data preparation is concerned with all steps necessary to go from the RAW data from the detector to physics analysis.
  - At least that is how it is supposed to work, more times than not it is more like what these people are doing: fighting fires, and trying to fix the latest thing that broke!

- The trigger, detector calibration, reconstruction and luminosity determination are things that affect all ATLAS analysis.
  - The data preparation work is the input to all our physics results.
- We are about to resume data taking soon, at a higher luminosity and with several intriguing things in data to check.
  - Another big effort this year will be preparing the next release.

**Another exciting year ahead of us!**

# Backup

# The IBL Bowing

Jonas Strandberg

- In 2015, had to start aligning the innermost tracker every run.
  - Increased LV currents caused temperature change, cause bowing of the IBL staves. Introduce an extra ES step to correct for it.

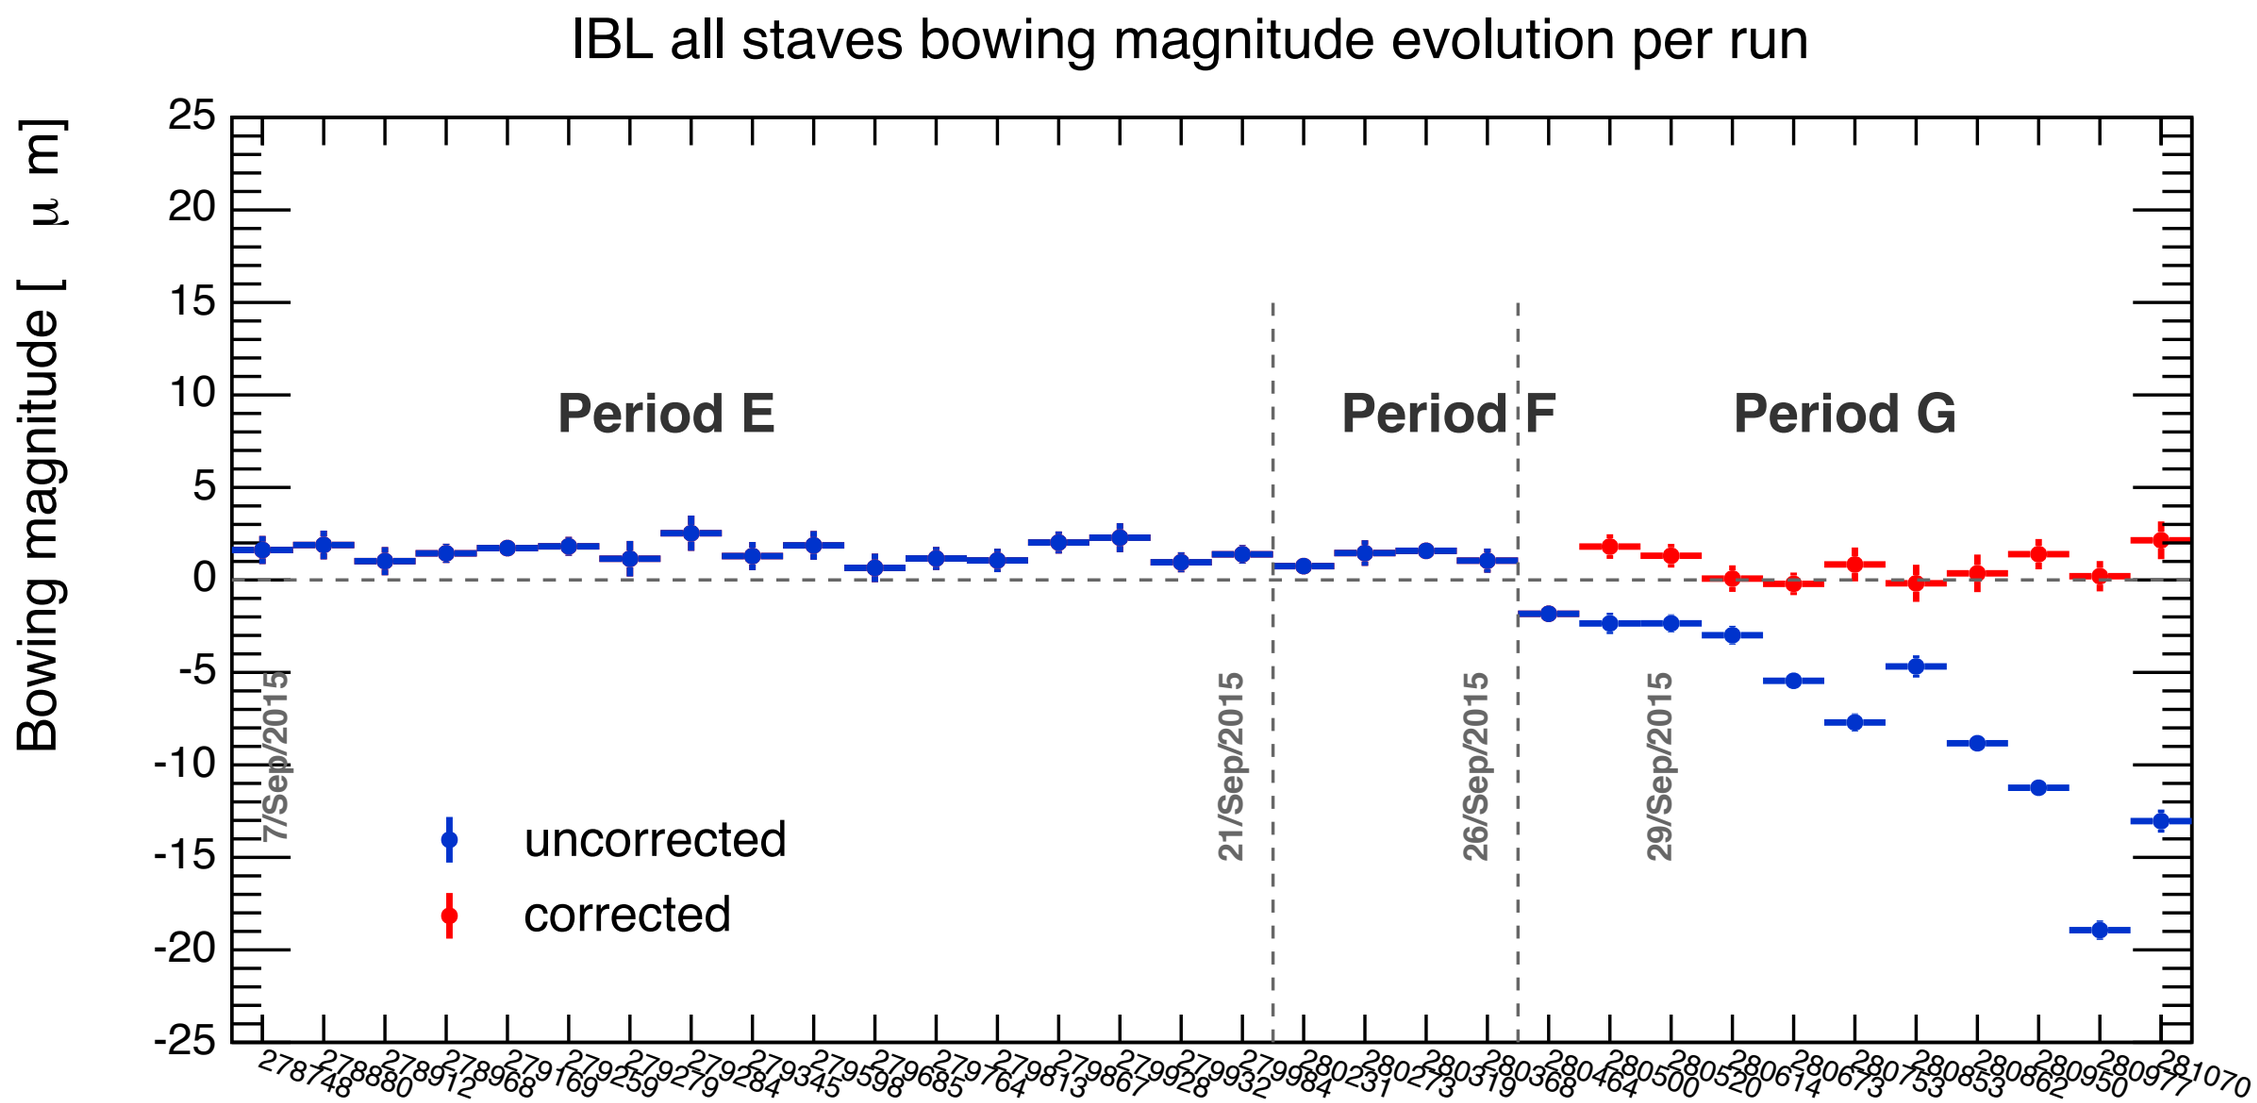

# Data Preparation Subgroups

Jonas Strandberg

**Prompt Reconstruction (PROC)**
Magda Chelstowska, Tulay Donszelmann

**Conditions**
Susumu Oda, Monica Verducci

**Data Quality**
Rosy Nikolaidou, Anyes Taffard, Lisa Shabalina

**Tier0 Operations**
Armin Nairz, Luc Goossens

**Data Preparation**
Paul Laycock, Jonas Strandberg

**Luminosity**
Witold Kozanecki, David Salek, Vincent Hedberg

**Non-collision backgrounds**
Saverio D'Auria, Stephen Gibson

**Reprocessing**
David South, Kostas Nikolopoulos

**Metadata**
Aurelio Juste, David Malon

**Beamspot**
Anthony Morley, Rainer Bartholdus