Online Determination of the LHC Luminous Region with the ATLAS High Level Trigger

Machine Detector Interface & Beam Instrumentation

TIPP 2011

Chicago, June 13, 2011

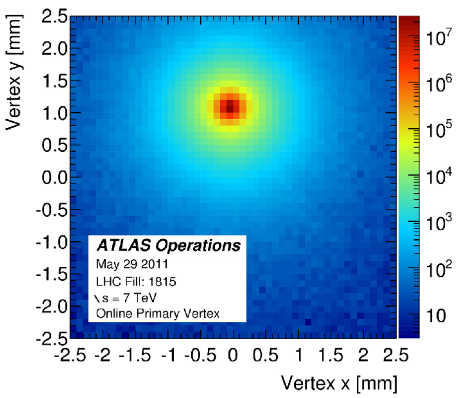

Rainer Bartoldus

SLAC

(on behalf of the ATLAS Collaboration)

ATLAS as "Beam Instrumentation"

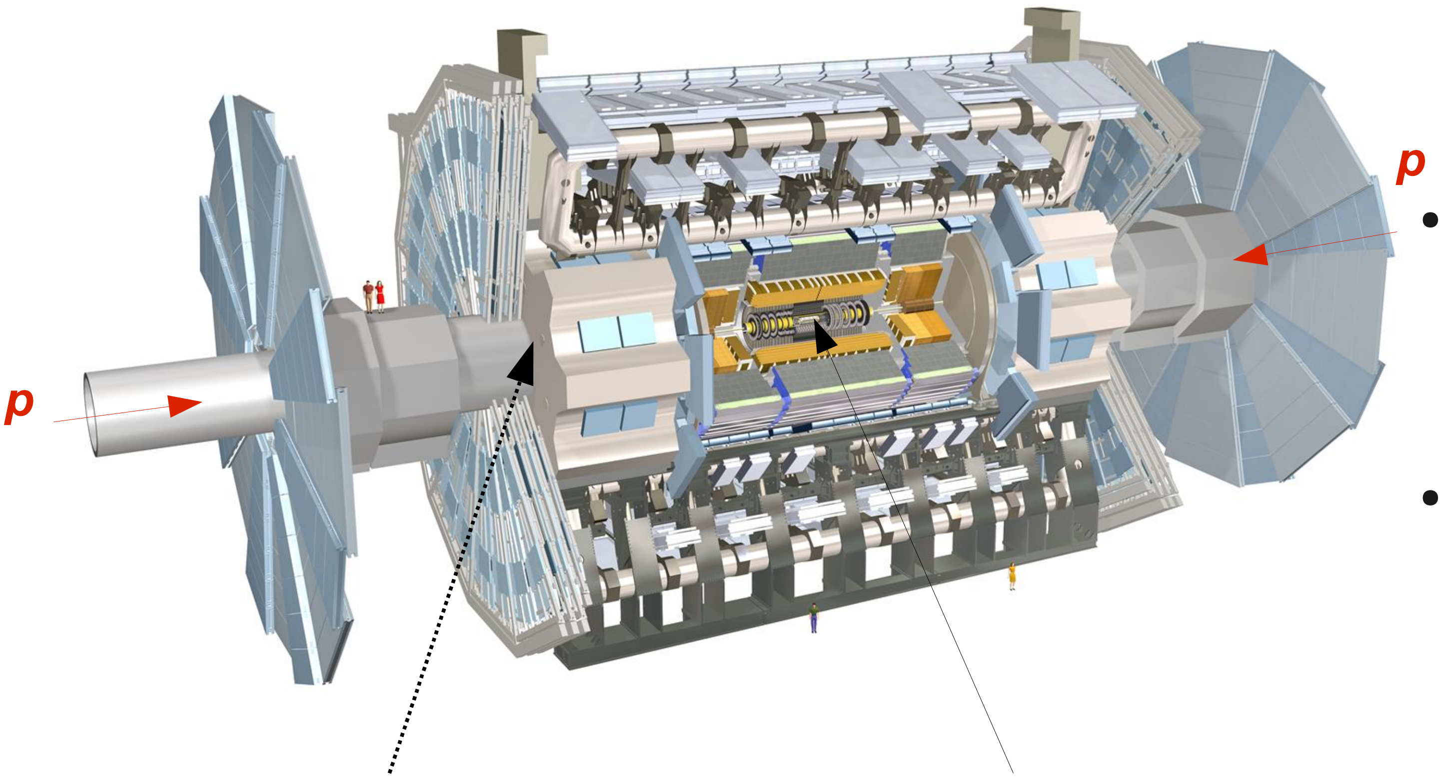

Beam Position Monitor
~ 21.5 m from the IP
2 channels (analog)

ATLAS Inner Detector
~ 0.05 m from the IP
> 80 million channels

- The LHC is an extremely well instrumented machine
- It is amazing how much we know about its beams and how well we know it
- Nevertheless, close to the interaction region the experimental detectors are not only best positioned but overwhelmingly well equipped to characterize it

- Their sophisticated High Level Trigger systems allow to do that in real time

Overview

- The ATLAS High Level Trigger (HLT) provides a unique platform for measuring LHC luminous region parameters
- Doing this in the online environment is particularly challenging in several ways:
  - Tightly constrained CPU and bandwidth budget of the trigger system
  - Massively parallel execution of algorithms that need special infrastructure to be aggregated and fanned out again
  - You get to see every event only once
    - No iterations on conditions, resolution etc: everything is 'at the edge of time'
- On the positive side, it offers unique advantages, too:
  - Unparalleled statistics, taking advantage of the many rejected events
    - Practically the only place with enough rate to see per-bunch time-evolutions
  - Very short latency to give quasi real-time feedback (minute scale)
- In addition, the Trigger itself needs to know very precisely - and adjust to - the position, size and orientation of the luminous region (*e.g. b*-tagging)
- This measurement is part of a bigger feedback loop around the HLT

ATLAS Inner Detector

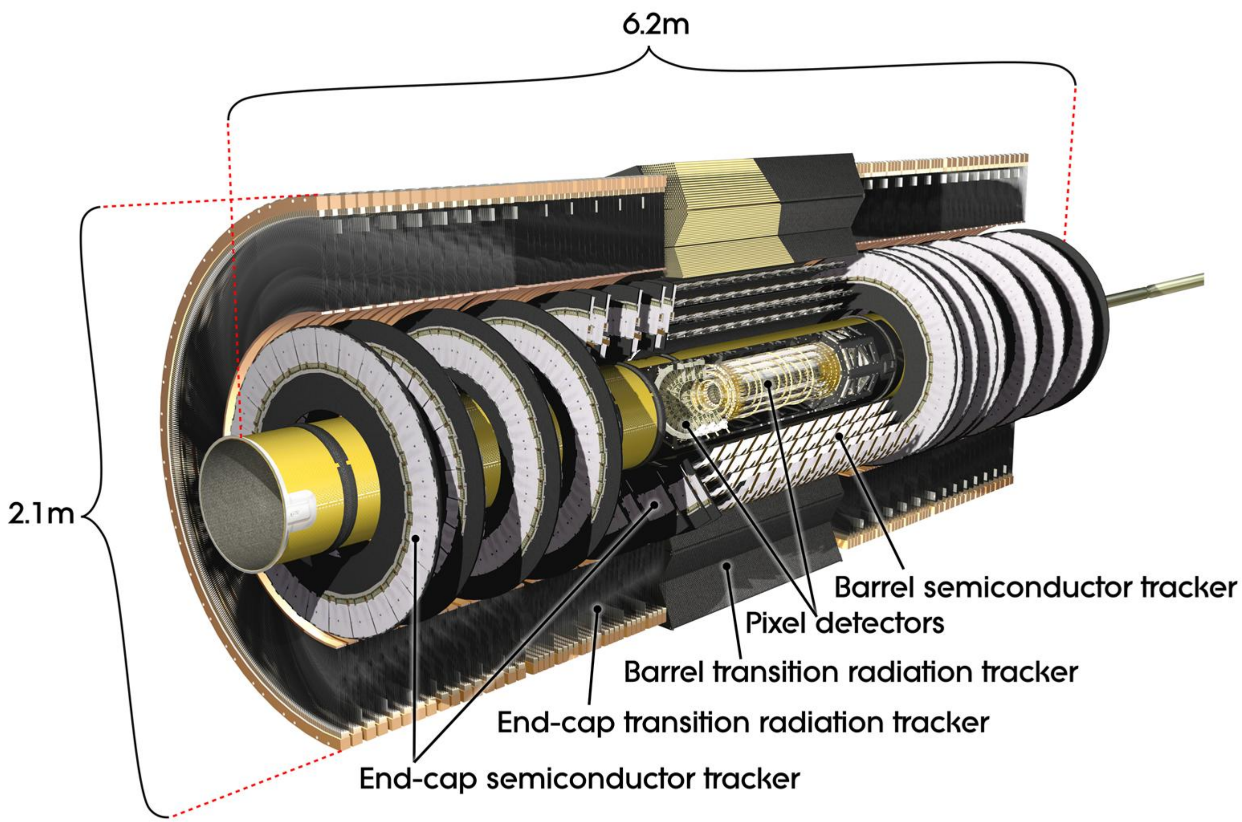

- Silicon Pixels
  - 3 barrel layers
  - 2 x 3 endcap discs
  - $\sigma_{r\varphi} \approx$ 10 μm
  - $\sigma_z \approx$ 115 μm
- Silicon Strips (SCT)
  - 4 barrel layers
  - 2 x 9 endcap discs
  - $\sigma_{r\varphi} \approx$ 17 μm
  - $\sigma_z \approx$ 580 μm

- Only the silicon detectors (Pixel+SCT) needed here for pattern recognition, track finding and fitting
- Vertex resolution is dominated by the Pixel detector
- Only a small fraction (~80 kB) of the total event size (~1.6 MB)
  - But all of the detector needed on every event (not just Region of Interest)

Primary Vertices

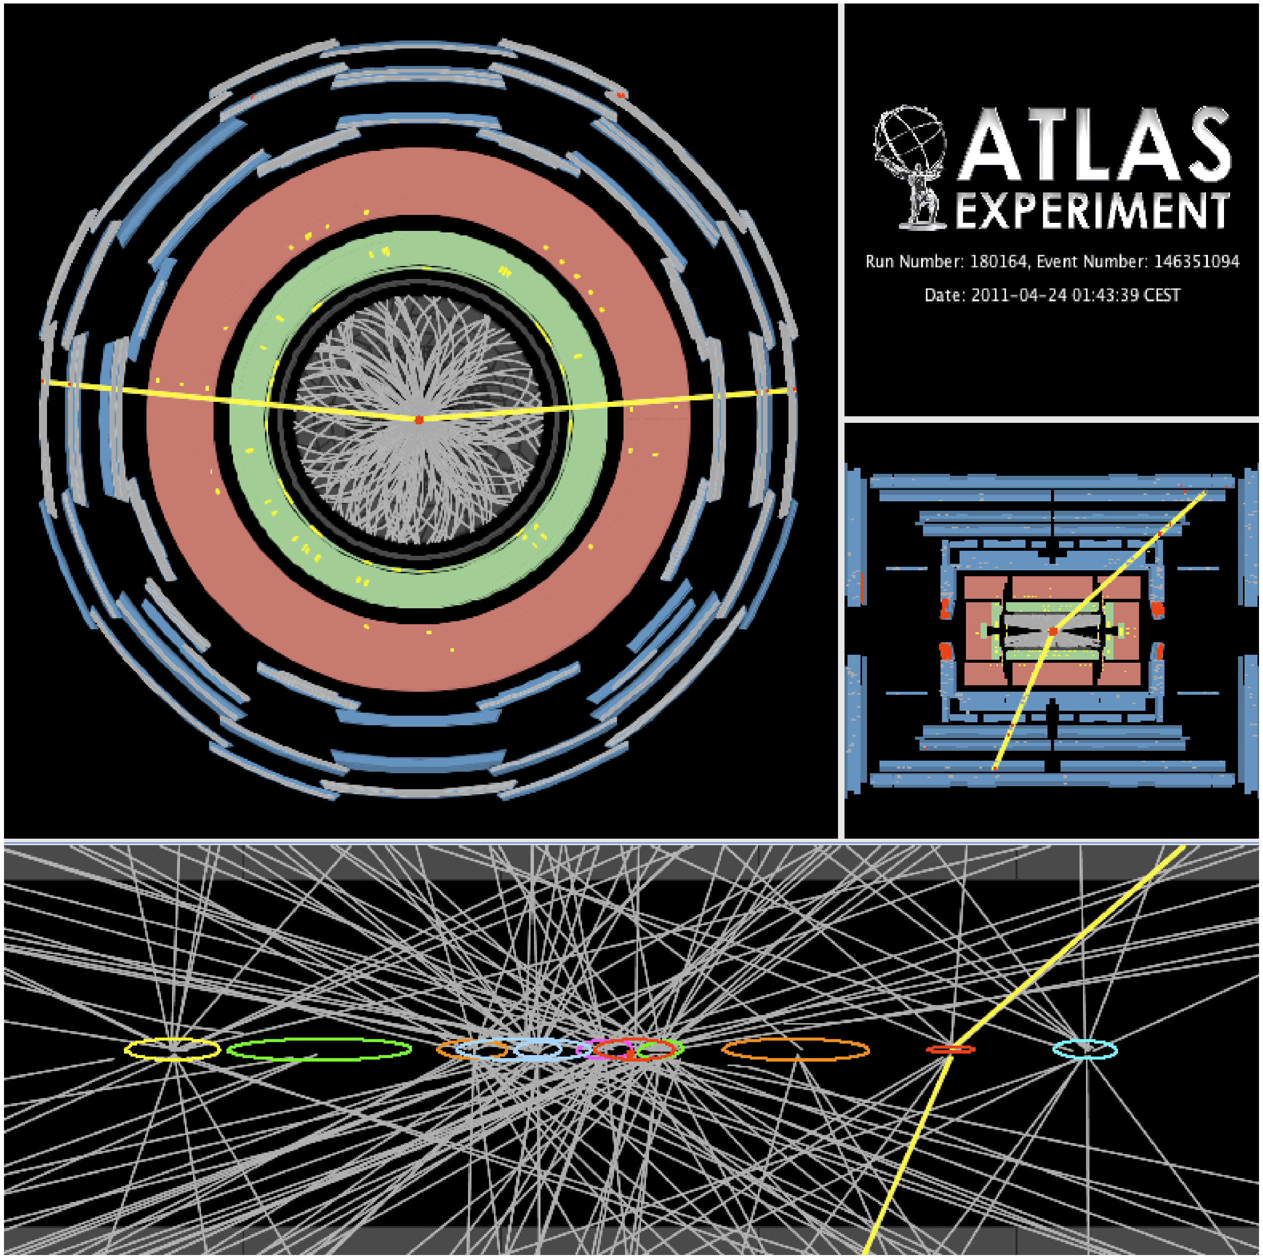

- At present luminosities, just above 10^{33} cm^{-2} s^{-1}, there are on average 6 interactions per bunch crossing
  - Called "pile-up"
- In principle these allow us to make several measurements on each event
- However, they are computationally very **expensive to reconstruct and resolve in real time**

ATLAS High Level Trigger

- Runs after the Level 1 hardware trigger
- Massively parallel, farm of 1000+ nodes
- Two stages: Level 2 (L2) + Event Filter
- Current rates: 50 kHz L1 → 4.5 kHz L2 → 400 Hz Event Filter (logging to disk)
- L2 does partial reconstruction
  - First trigger with access to Si-tracker data

- Chose L2 Trigger to host **beam spot algorithm**
  - Highest available statistics, low latency
  - But: challenge to do full scan of silicon tracker detectors for data transfer and pattern recognition, track reconstruction
  - At the edge of available bandwidth + CPU
- Currently 10 racks of 30 nodes with 8 cores each (2400 processes)

- *For more on ATLAS Trigger/DAQ: c.f. earlier talks by Sergio Ballestrero and Srini Rajagopalan at this conference*

Beam Spot Algorithm

- Dedicated algorithm that executes on the L2 processors (as part of a rich trigger menu of physics and calibration triggers)
- Invokes a "full-scan" track reconstruction using Pixel and SCT, and then employs a fast vertex fitter to reconstruct primary vertices
  - Clusters of tracks are formed along the beam line (z), around the highest transverse-momentum (p_T) seed track
  - In an iterative procedure, the clusters are fitted to a common vertex
- The coordinates (and other relevant properties) of each vertex are histogrammed online
- From the spatial distribution of primary vertices we extract through fits the **position, size, and orientation of the luminous region**
  - Parameters are the centroid x,y,z; widths σ_x,σ_y,σ_z; and title angles θ_{xz},θ_{yz}
  - Together with their errors (significance)
- Challenge: An important ingredient to measuring the true luminous region widths is to correct for the intrinsic vertex resolution
  - Typical vertex resolution is on the same order as the transverse widths
- Special data needs to be acquired to accomplish this

Primary Vertex Distributions

- Projections of the three-dimensional distribution of reconstructed primary vertices are histogrammed and published once per minute
- They are aggregated (“gathered”) across the farm and re-published
- The large amount of available statistics gives rise to very precise determinations of all parameters
- In addition, the vertex count can serve as a measure of luminosity (although this gets increasingly difficult with pile-up)

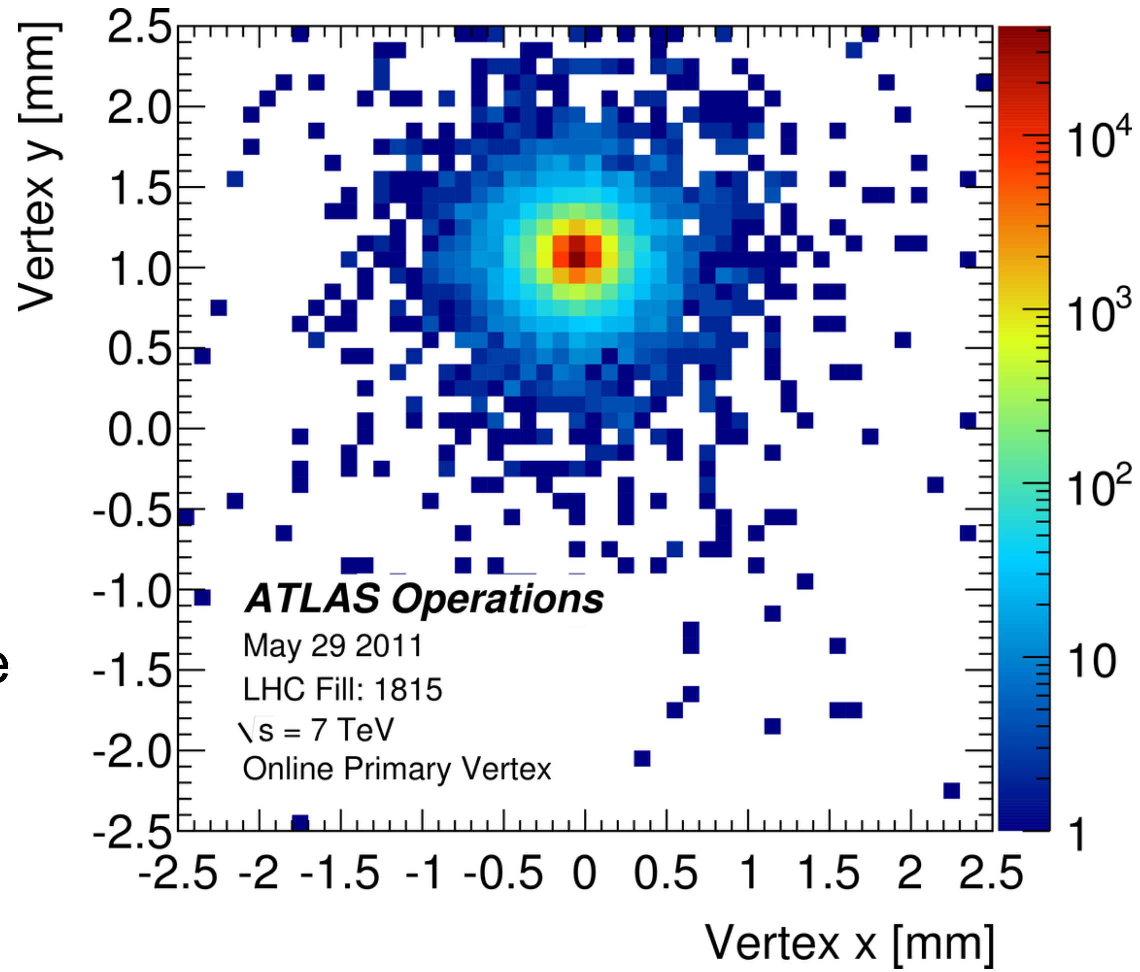

- Transverse distribution of 100,000 vertices reconstructed by the HLT
- One minute of data taking!

Primary Vertex Distributions (cont.)

- 1D-projections of the vertex distributions
- Means of gaussian fits measure centroid position to better than 1% of the width with one minute of data
- Transverse widths (x,y) are not yet corrected for resolution

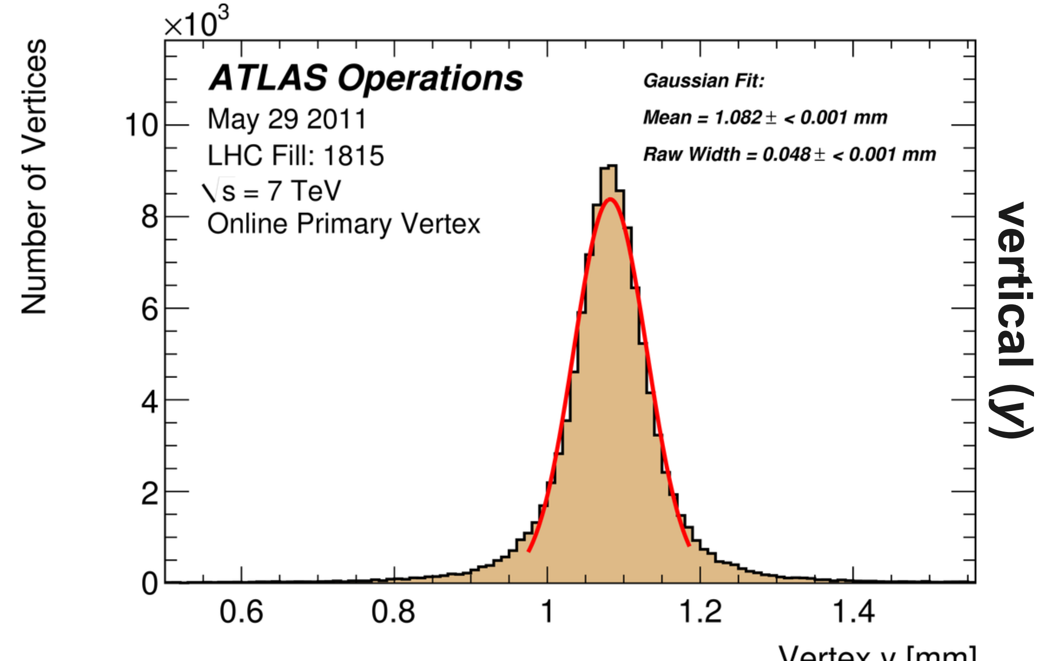

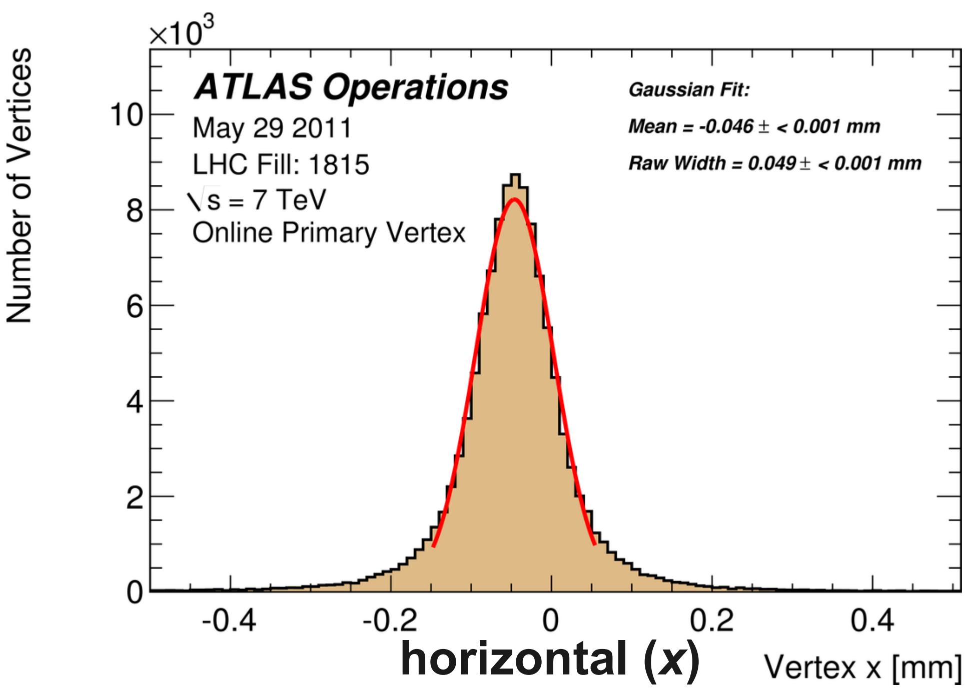

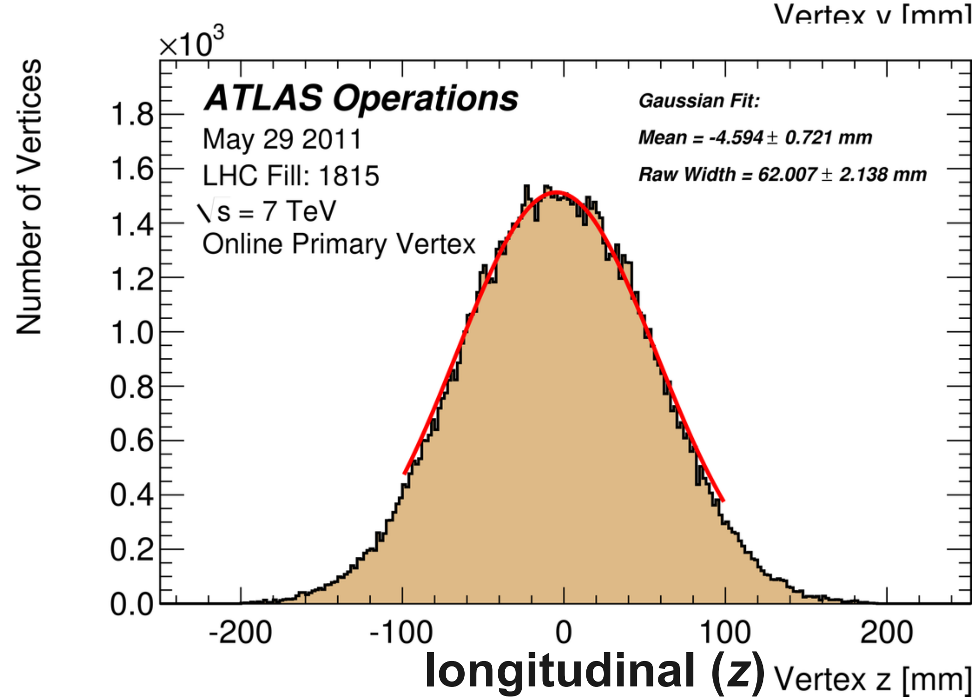

Track Impact Parameters: Cross Check

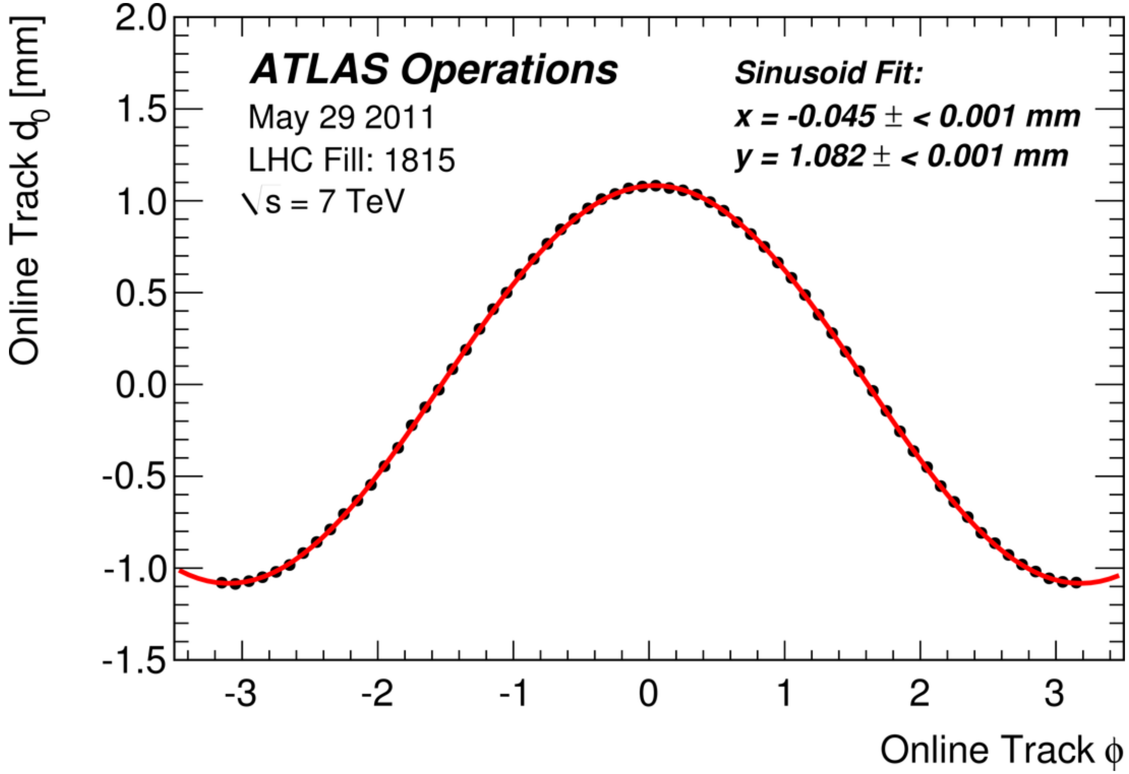

- Sinusoidal fit to track d_0 versus φ
- → amplitude and phase translate to x and y position of **point of closet approach**
- (errors are too small to be visible)

- As an alternative to (and cross check of) the vertex method one can use the track transverse impact parameter d_0 versus azimuth angle φ distribution to extract the centroid position
- The results are in excellent agreement with the vertex method
- The method doesn't provide the widths and has less control over backgrounds or against pile-up (so vertex is the preferred method)

Beam Line Tilt

- Horizontal and vertical angles of the luminous region in the ATLAS coordinate system can (and must) be measured very accurately
- These are mostly a result of the residual global misalignment of the ATLAS pixel detector with respect to the LHC beam line
- Transverse distributions are measured in bins along the beam (z)-axis, then a linear fit is performed to extract the slope
- **Residual tilts are very small and stable in time**. We originally found a 500 μrad horizontal tilt, and decided to compensate for it by rotating the ATLAS global coordinate system

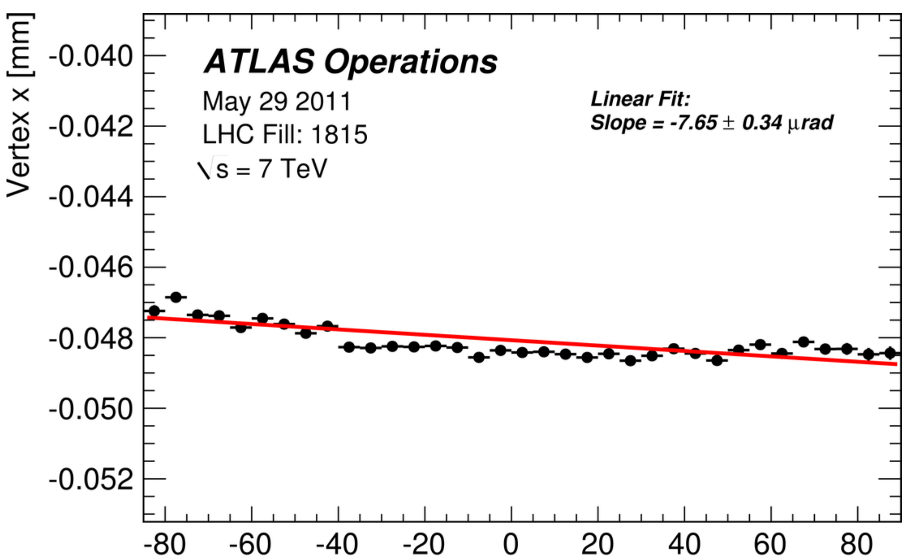

horizontal (x)

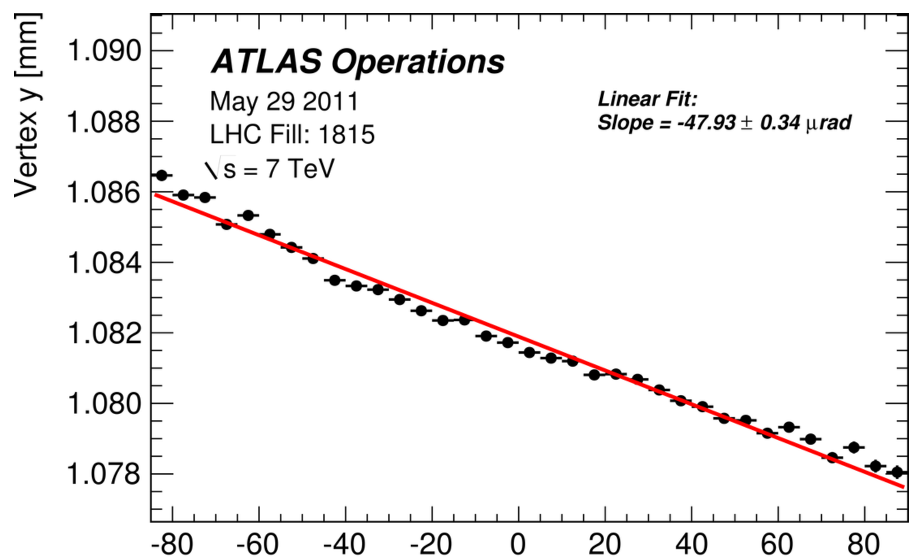

vertical (y)

Time-Variation of Beam Position

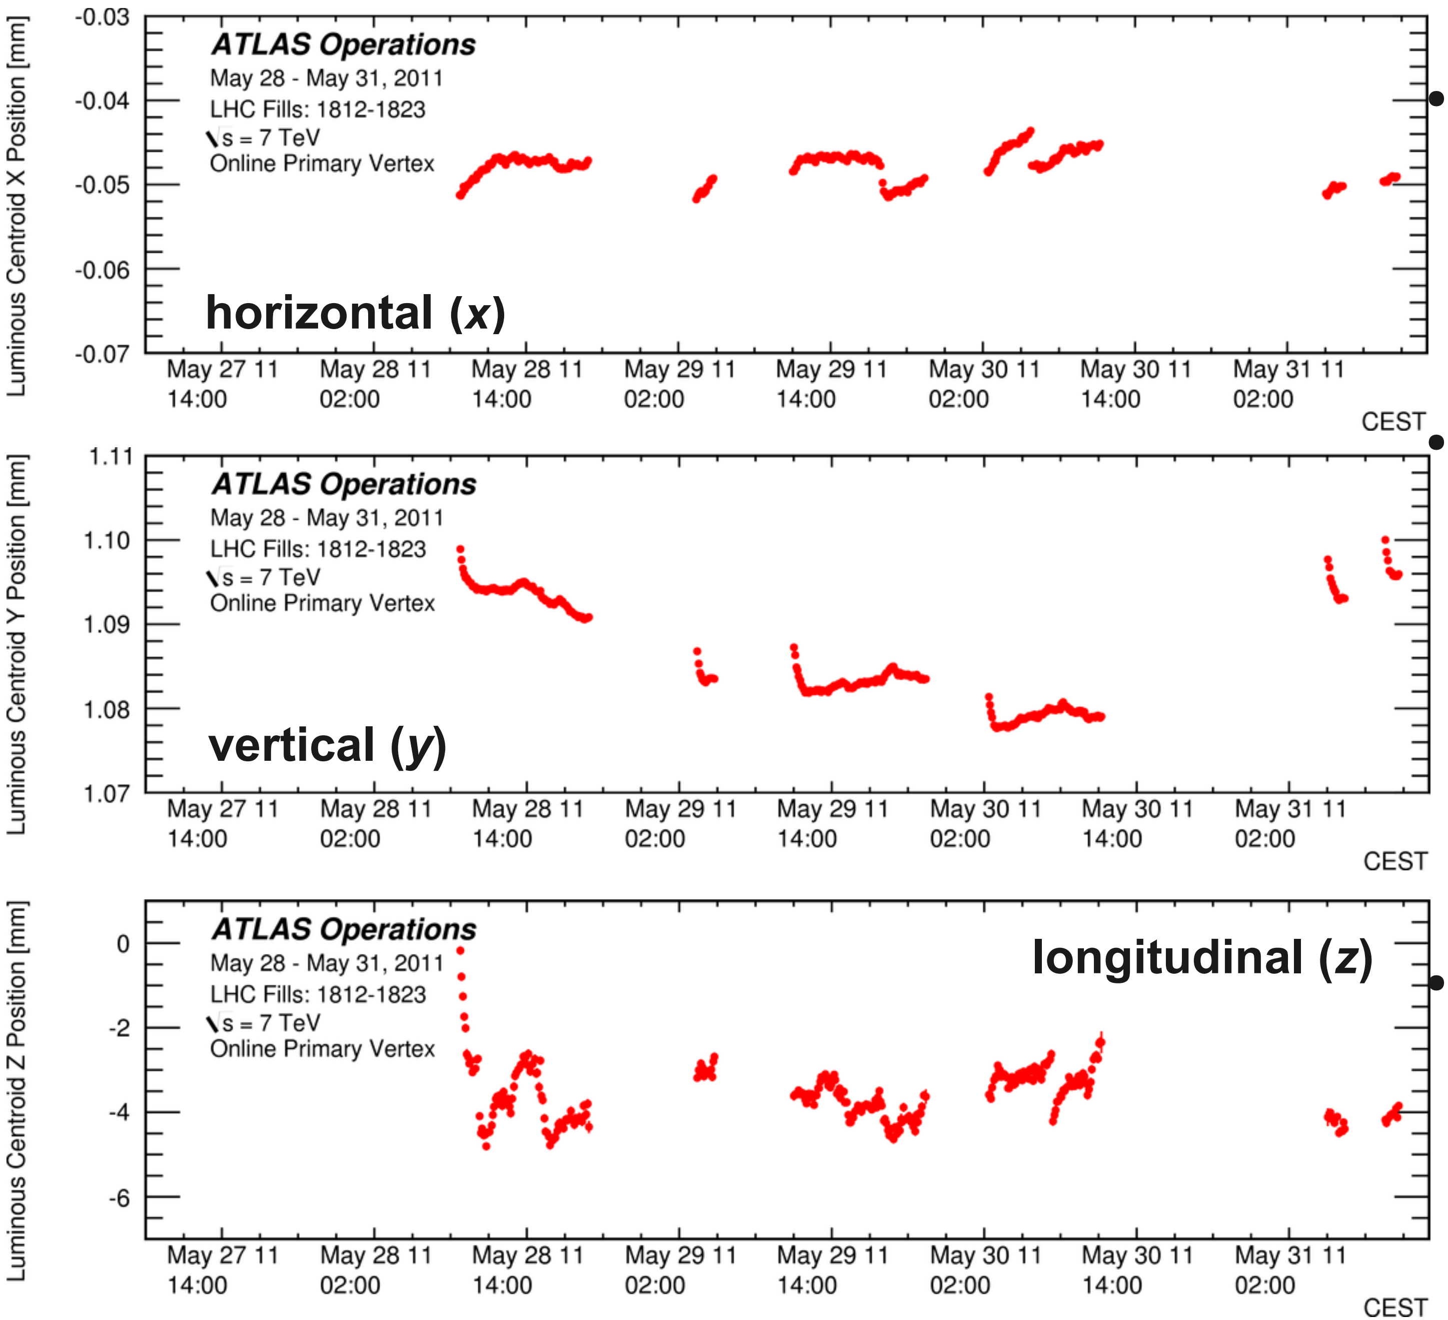

- Three-dimensional **centroid position** is measured over the course of several LHC fills
- Significant variations of the transverse position (from orbit corrections and drifts) are visible **within a fill and from fill to fill**
- Similar significant longitudinal time variations are related to RF-phase changes

Effect on *b*-Tagging

- Algorithms that depend on track impact parameters or secondary vertex significance are sensitive to knowing the beam spot position and size
- These triggers have to be configured with the current beam spot parameters
- Currently the most sensitive trigger is the *b*-jet trigger
  - Shifts on the order of a few σ of the luminous size can kill the efficiency and increase the fake rate to make the trigger non-functional
- The stored beam spot needs to be updated during the run to keep/get *b*-tagging working

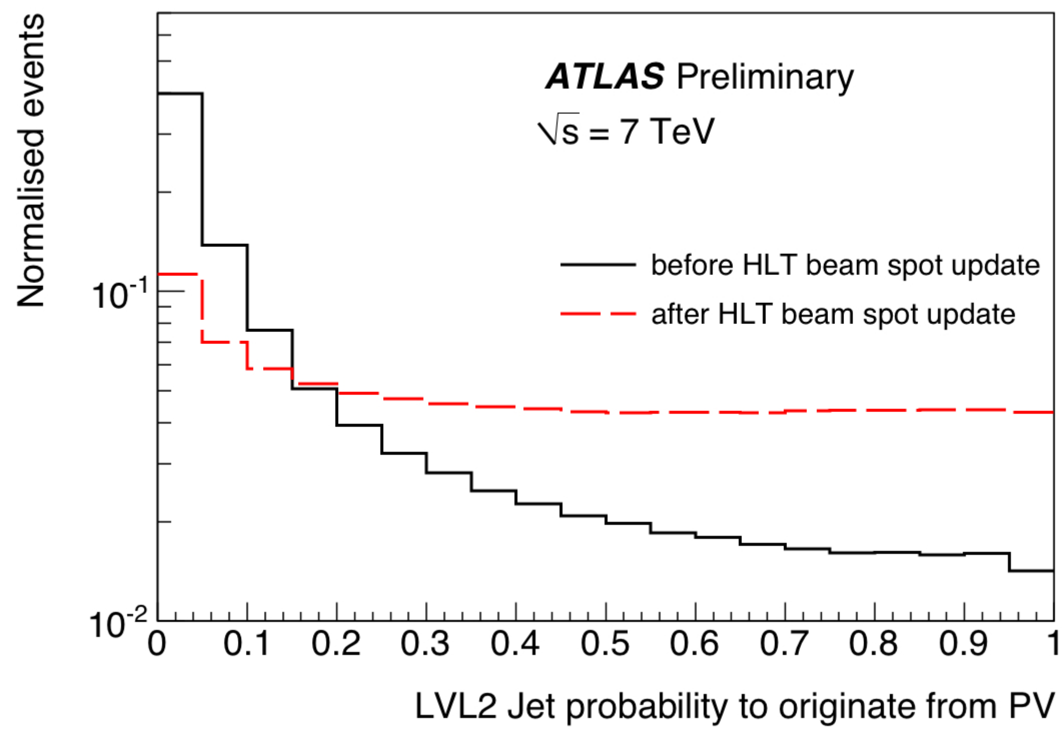

Probability that L2 jets originate from the luminous region

- QCD jets will be flat
- *b*-jets will peak near zero
- An off-set beam spot will fake large impact parameters

Measuring Resolution: Vertex Splitting

- The vertex resolution changes not only with detector conditions but also with the event mix from the (always evolving) trigger menu
- The only way to account for this is to **measure it online**
- We use the method of splitting each primary vertex into two, then fitting the two halves separately and measuring their displacement
  - Represents √2 times the resolution of a vertex of half the multiplicity
- A lot of care must be taken to avoid systematic effects (biases)

Tracks from one vertex are sorted in Φ and then split as: 1,2,1,2...

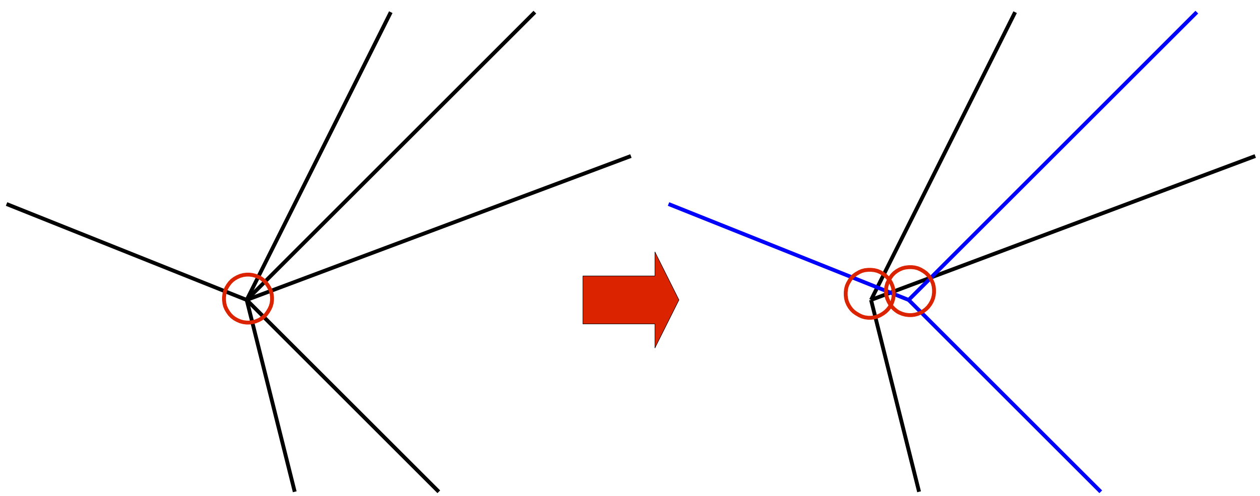

Two vertices from tracks coming from the same interaction

One original vertex...

... split into two half-vertices, fitted separately

Split-Vertex Resolution

horizontal (*x*)

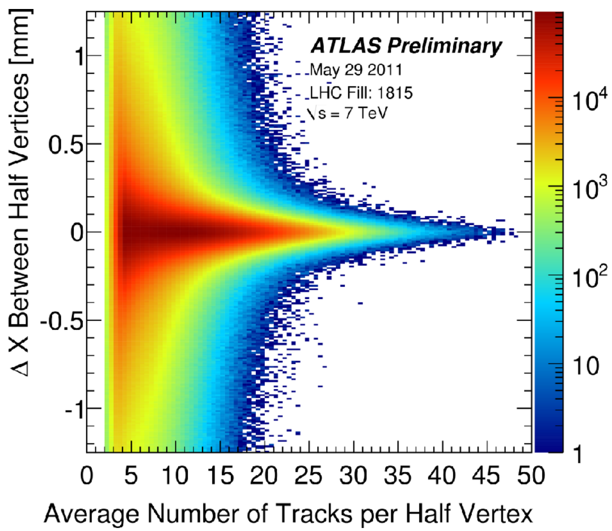

vertical (*y*)

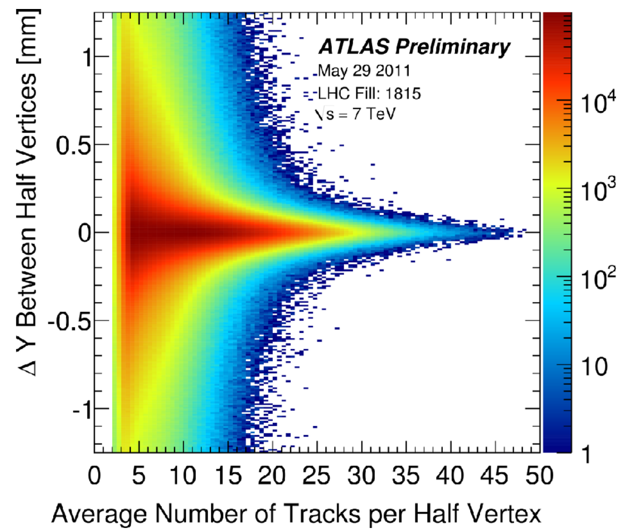

longitudinal (*z*)

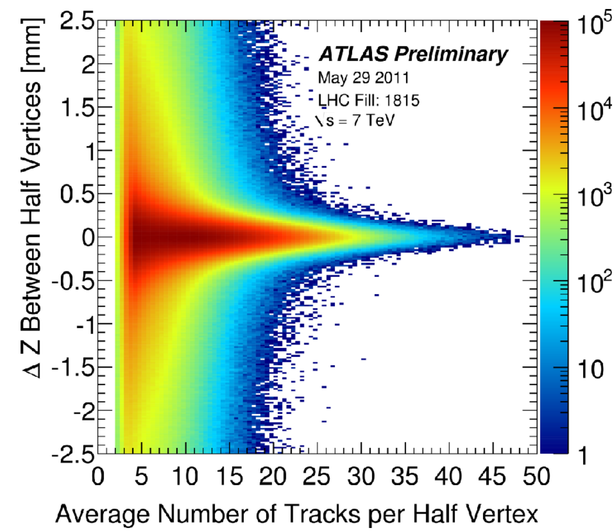

- Vertex resolution is a strong function of track multiplicity
- The two transverse resolutions are the same; longitudinal resolution is about twice that amount; (less different than silicon hit resolutions)
- Split vertex distributions show **two competing effects**:
  - The resolution is highest for vertices with the highest number of tracks
  - Statistics is highest for the low-multiplicity vertices and falls precipitously
- To optimize precision, one has to include the lowest multiplicity that can still resolve the spot size to be measured

Online Resolution Correction

Method works *down to* track multiplicities for which the resolution is not more than twice the true spot size

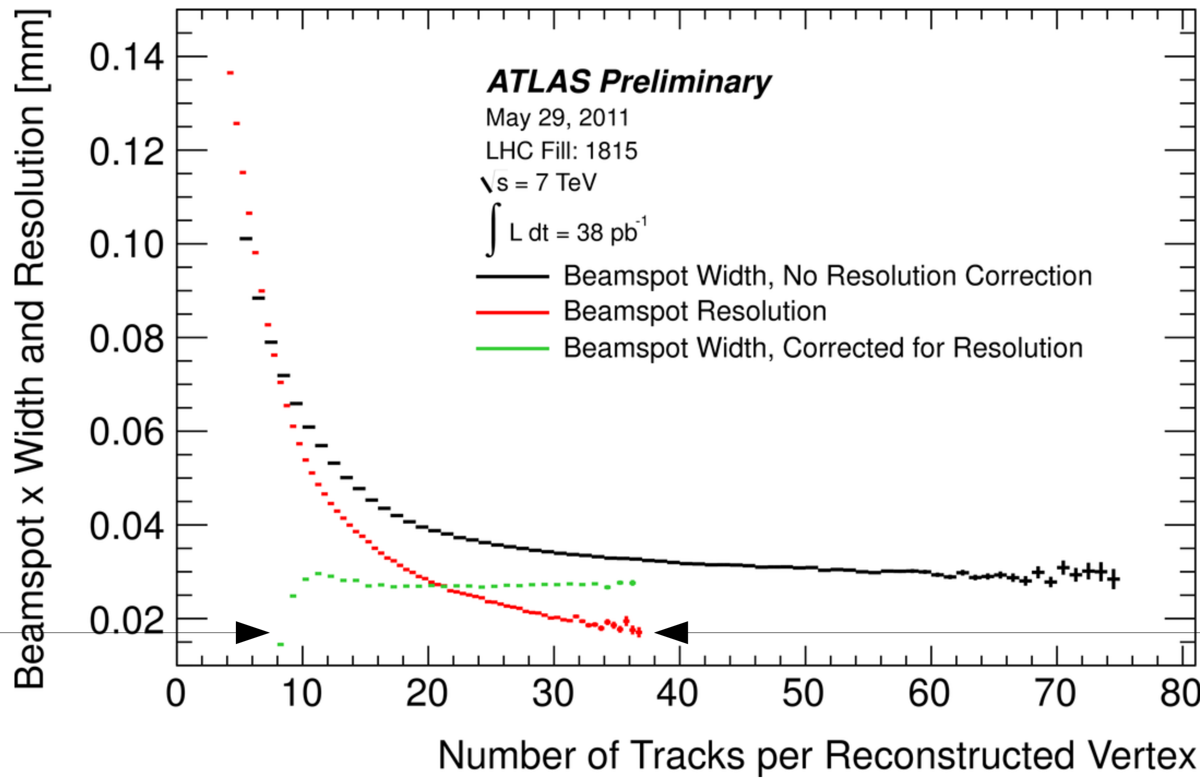

It works *up to* track multiplicities for which there is still statistics from vertices with twice the number of tracks

- The **observed width** and **measured resolution** are strong functions of the number of tracks in a vertex
- → The correction has to be applied in bins of track multiplicity
- **Corrected width** stays approximately flat regardless of multiplicity

Vertex Resolution Along *z*

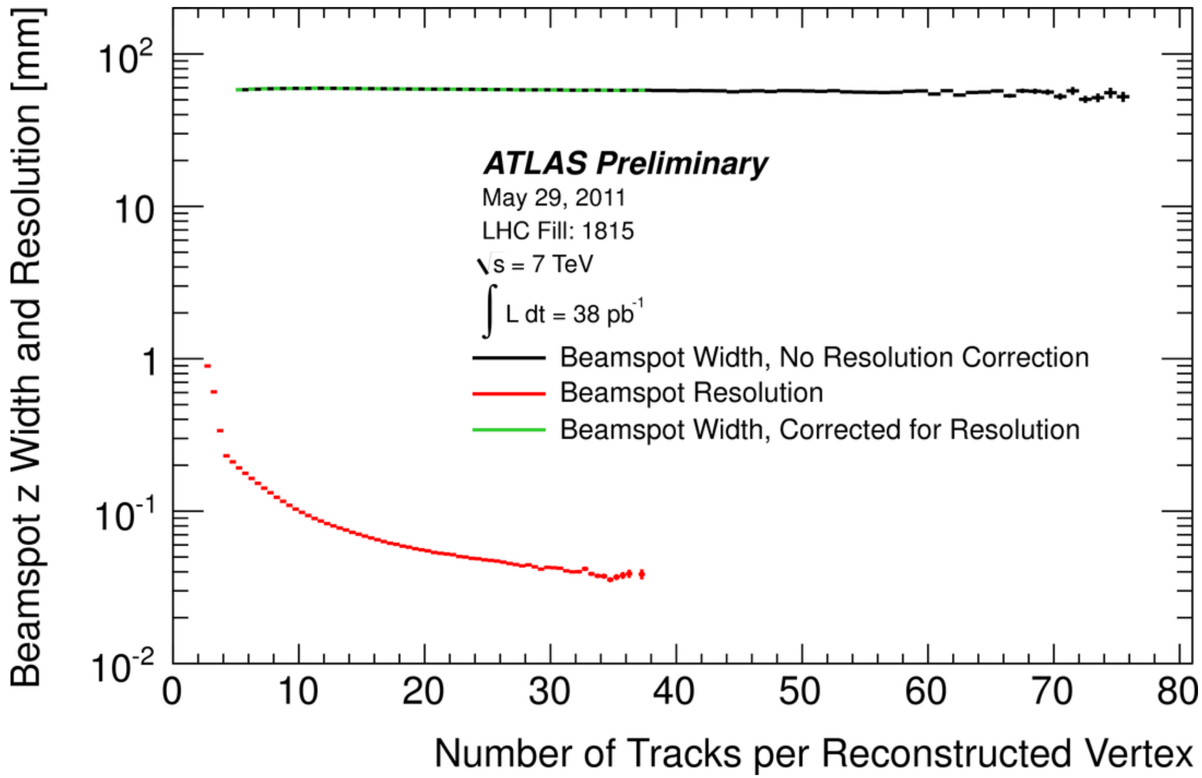

- **Vertex resolution along the beam axis** (as expected) somewhat lower than in the transverse plane
- Shows similar behavior as a function of number of tracks per vertex
- **Correction is entirely negligible** on the scale of the longitudinal size (~55 mm)

Time-Evolution of Luminous Size

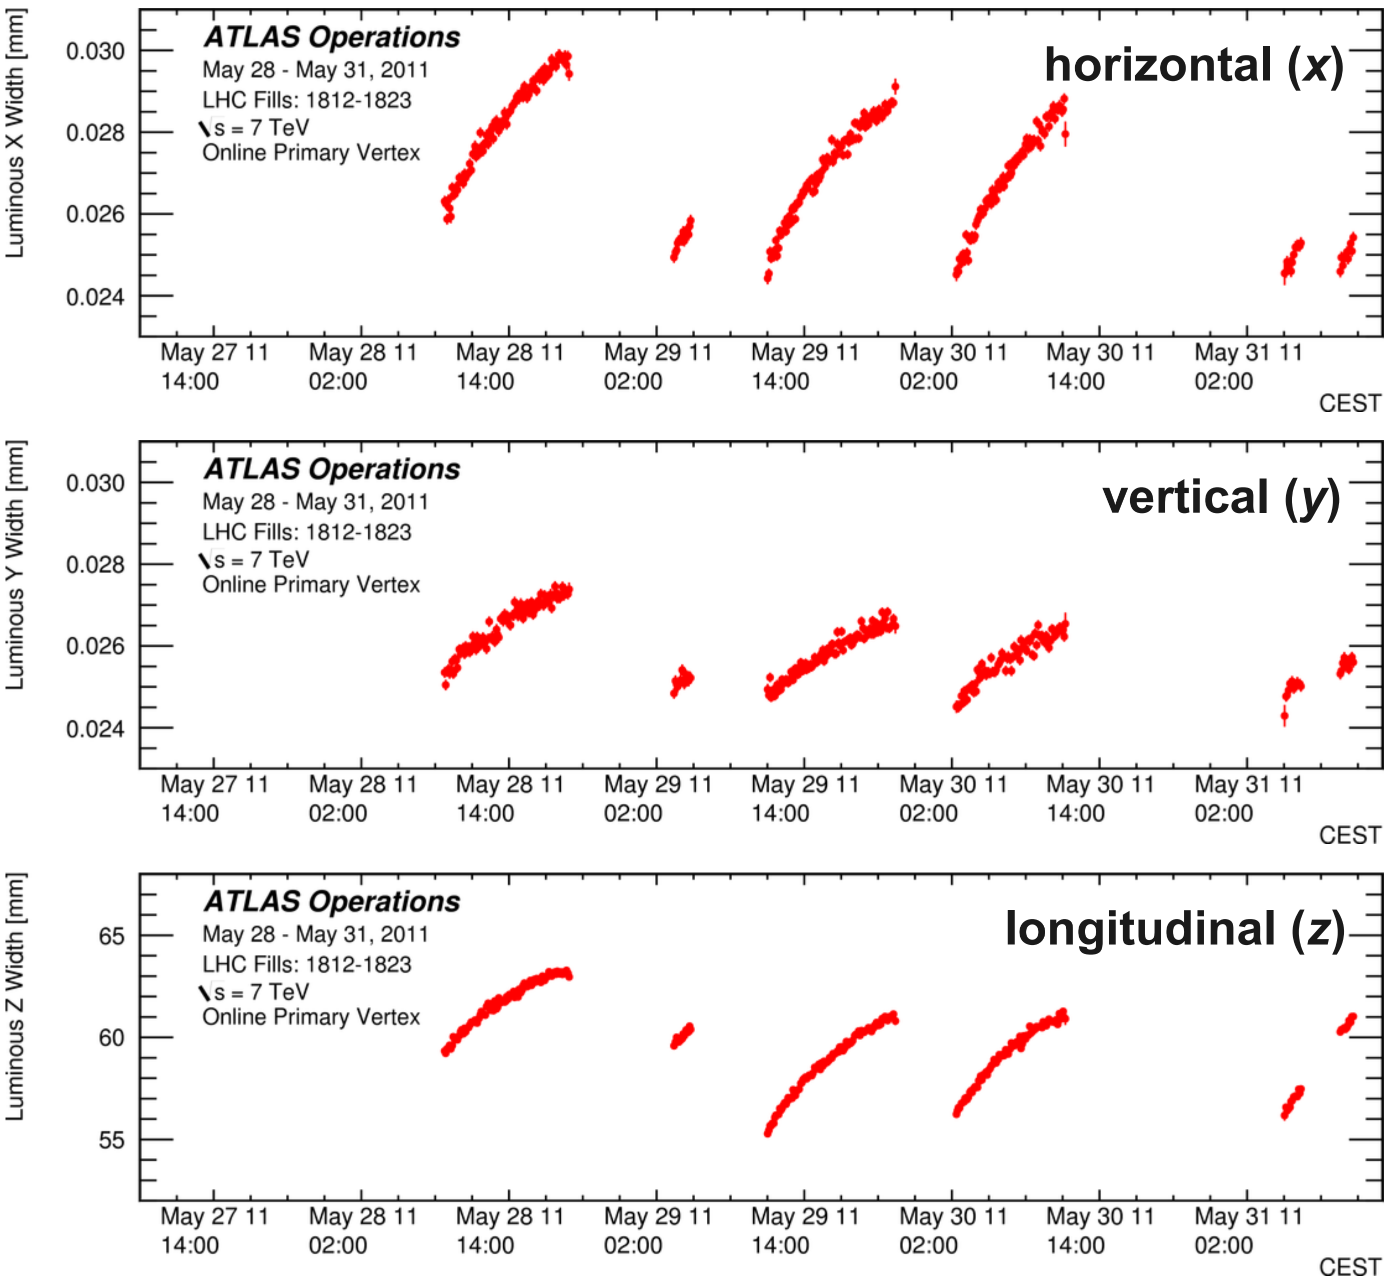

- Effect of transverse **emittance blow-up** during each fill is clearly visible
  - Approximately 15% horizontal, 10% vertical luminous width increase over a 10 hour fill
- Longitudinal emittance growth behaves similarly
- Fill-to-fill variations are comparatively small, but not negligible

LHC Feedback

ATLAS *Operations*

Live histograms are published once a minute

Available in the control room as well as over the web

Give details of the track/vertex distributions and parameters

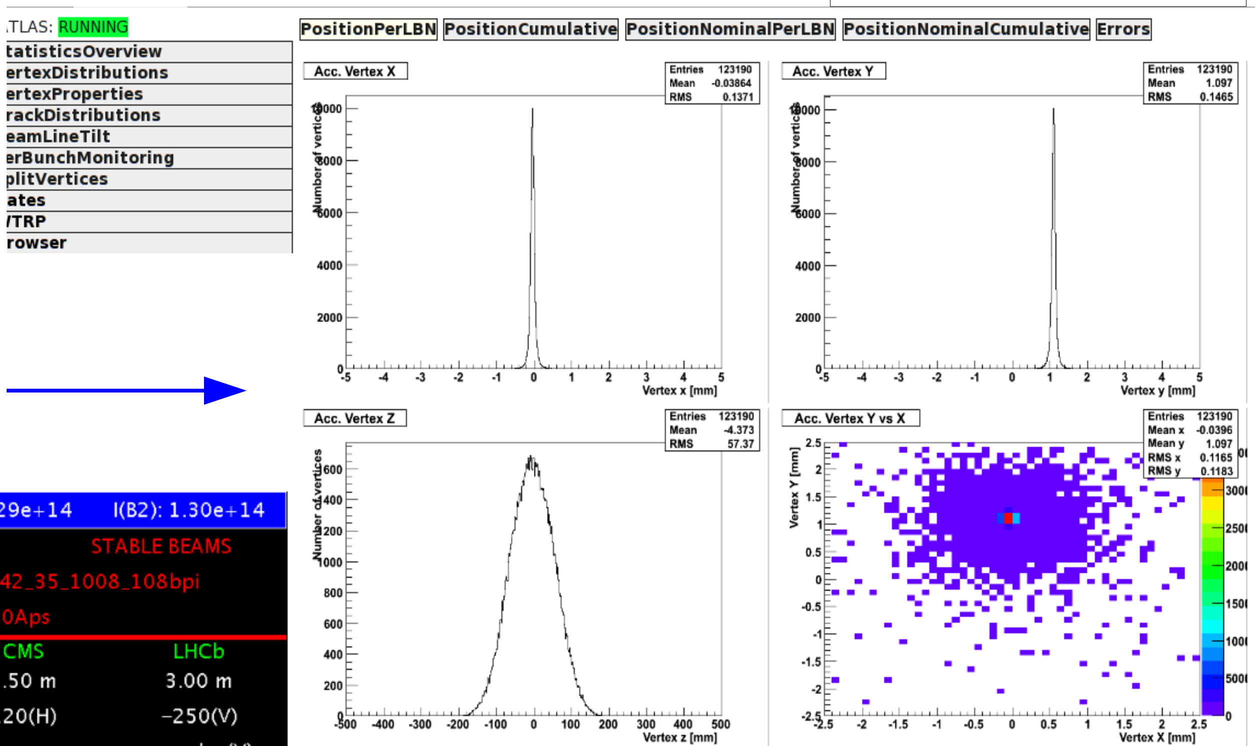

29-May-2011 12:07:26 Fill #: 1815 Energy: 3500 GeV I(B1): 1.29e+14 I(B2): 1.30e+14

Accelerator Mode: PROTON PHYSICS Beam Mode: STABLE BEAMS

Active Filling Scheme: 50ns_1092b+1small_1042_35_1008_108bpi

Active Hypercycle: 3.5TeV_10Aps

	ATLAS	ALICE	CMS	LHCb
Beta*	1.60 m	10.00 m	1.50 m	3.00 m
Crossing Angle (urad)	-120(V)	80(V)	120(H)	-250(V)
Spectrometer Angle (urad)		no_value(V)		no_value(V)
Beam Separation (mm)	0(H)	.3(H)	-.5(V)	-.11(V)
Expected Collisions per turn	1042	35	1042	1008

	ATLAS	ALICE	CMS	LHCb
BPTX: deltaT of IP (B1-B2)	0.02 ns	0.10 ns	-0.01 ns	0.01 ns
Luminous size (x,y) in um	25.0,24.9		25.8,23.6	45.8,44.1
Luminous size (z) in mm	56.1		52.7	41.1
Lumi Centroid (x,y) in um	-46.8,1081.9		179.0,-746.7	462.3,-17.1
Lumi Centroid (z) in mm	-3.6		-6.2	0.1
Luminous Tilt in urads	-12.24,-53.45		59.75,78.53	-30.98,43.54

The luminous region parameters are extracted through fits also once a minute

Corrections are performed and values are sent over to the LHC

Beam Spot Parameter Redistribution

- Time variations make it necessary to update parameters used by the HLT **during running** (including a bootstrap at the beginning of a fill)
- Critical for algorithms sensitive to primary vertex such as *b*-tagging
- The first challenge is how to transmit the parameters back to the many thousands of HLT processes (on a sharp time boundary)
- ATLAS has a proxy-technology for database configurations that makes it possible to do such updates extremely fast also within a run

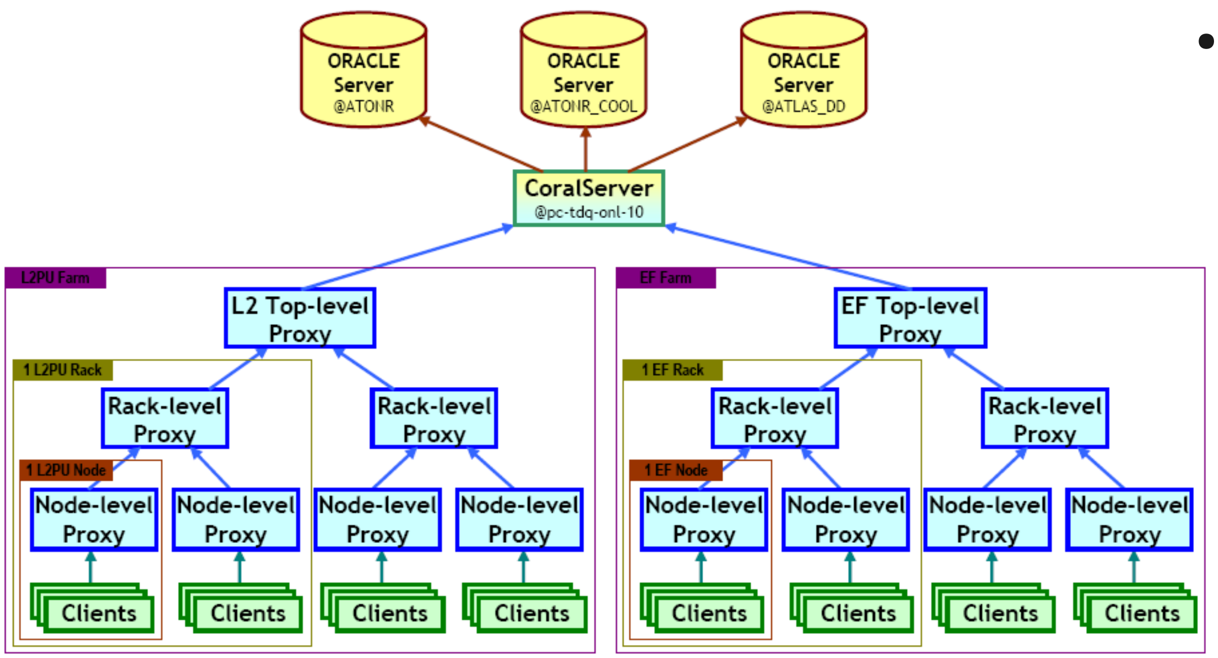

- CORAL Server & Proxy
  - Dedicated database technology developed for ATLAS HLT as example
  - Proxy hierarchy allows simultaneous access of 1000s of clients
  - *Multiplexes* and *caches* queries and responses

(Collaboration with CERN IT + U.Mainz)

High Level Trigger Feedback Loop

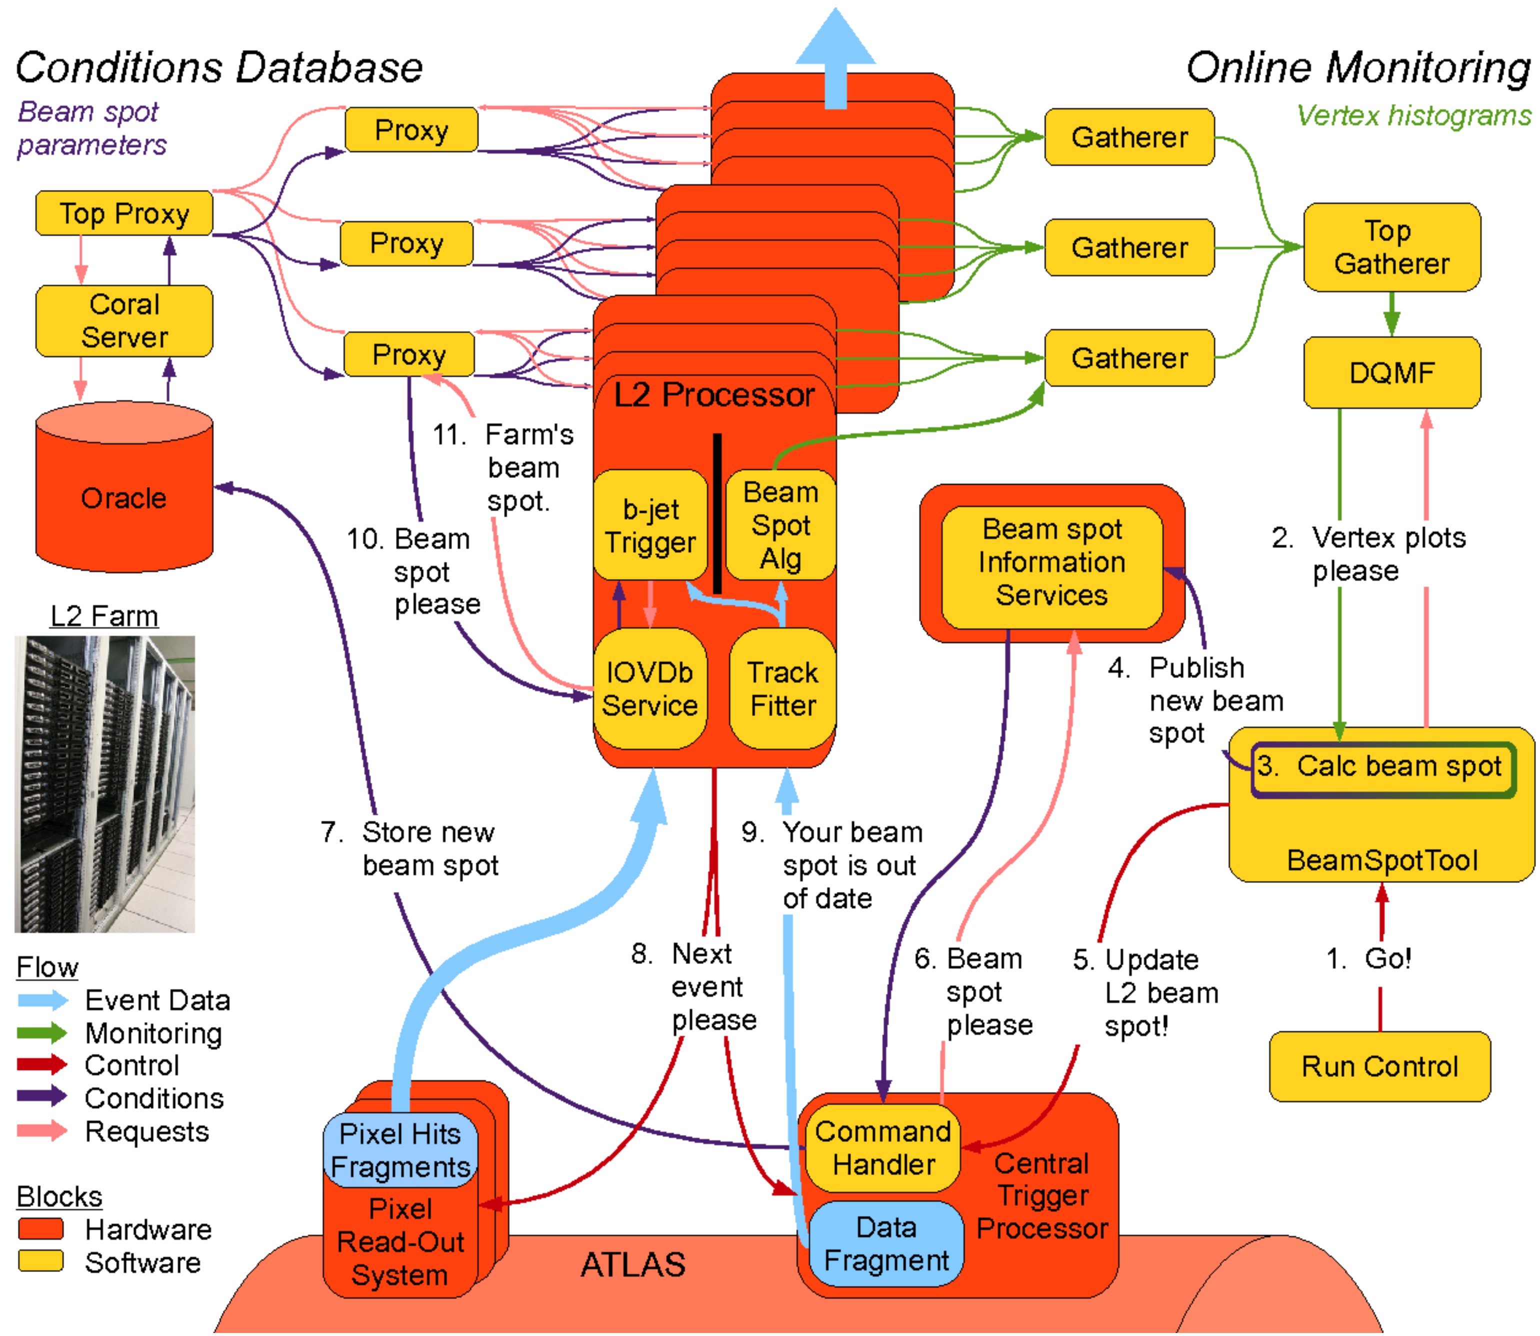

Trigger Auto-Updates

ATLAS *Operations*

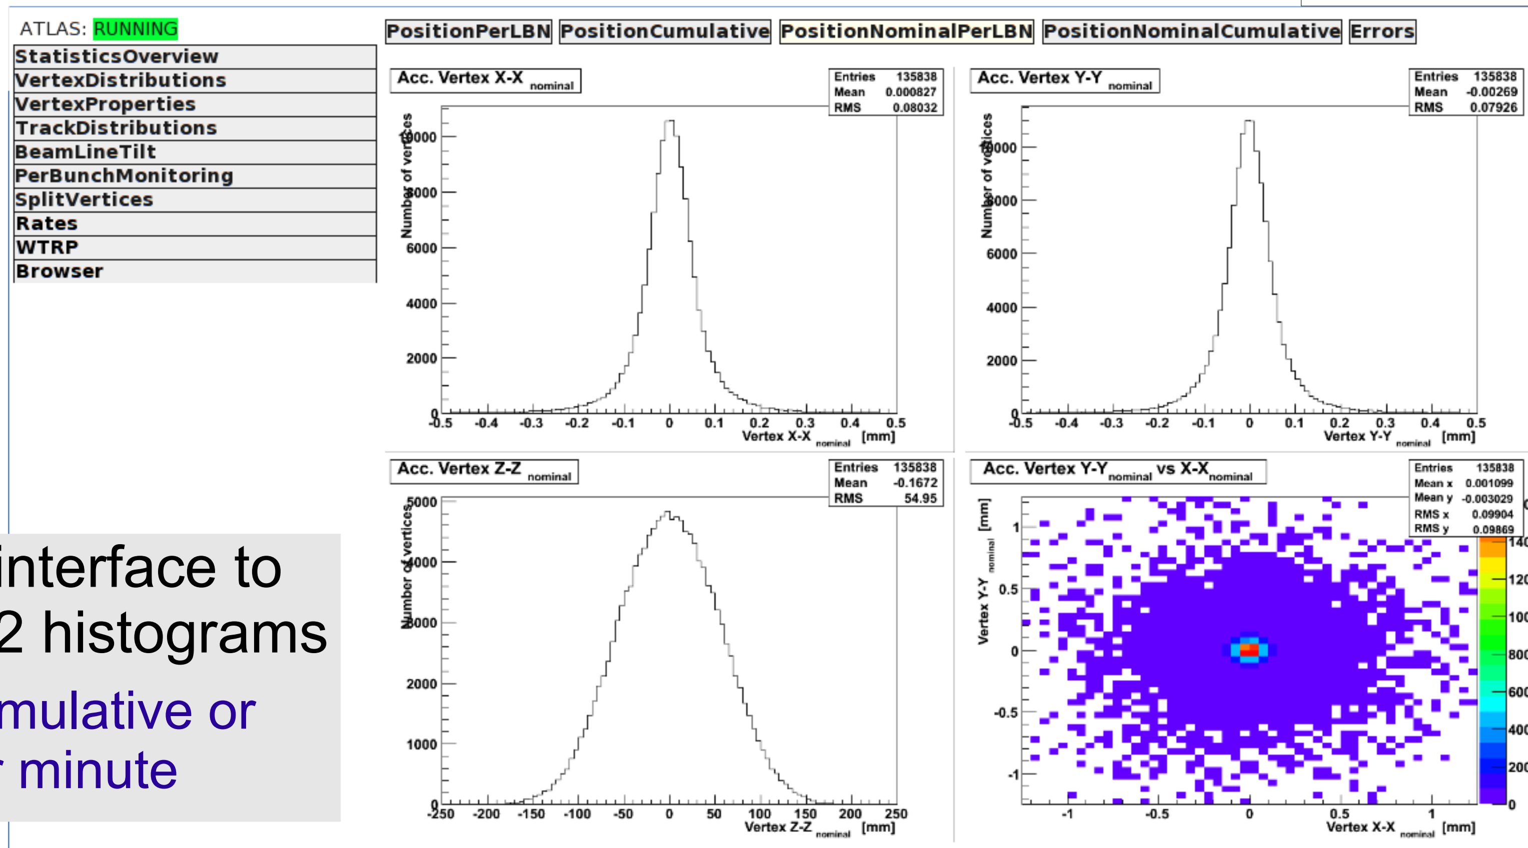

- Web interface to live L2 histograms
  - Cumulative or per minute
- Automatic updates of the farm are monitored online using “delta distributions” that compare observed values with the currently stored “nominal” values
- These are tuned to stay within a narrow range:
  - **Updates are triggered when** positions move by ±10% of the width, or widths change by ±10% from nominal (both with 2σ significance); or errors improve by more than 50%

Online-Offline Comparison

- A few hours after each fill, the more sophisticated **offline beam spot determination** runs on a fraction of the logged events
  - Complementary:
    > resolution, << statistics
  - Unbinned maximum likelihood fit with error scaling
- Automatic **monitoring** shows excellent agreement between these very different methods!
  - 1 µm difference in $\boldsymbol{\sigma_y}$ due to still uncorrected residual tilt
  - Slight difference in $\boldsymbol{\sigma_z}$ due non-gaussian tails

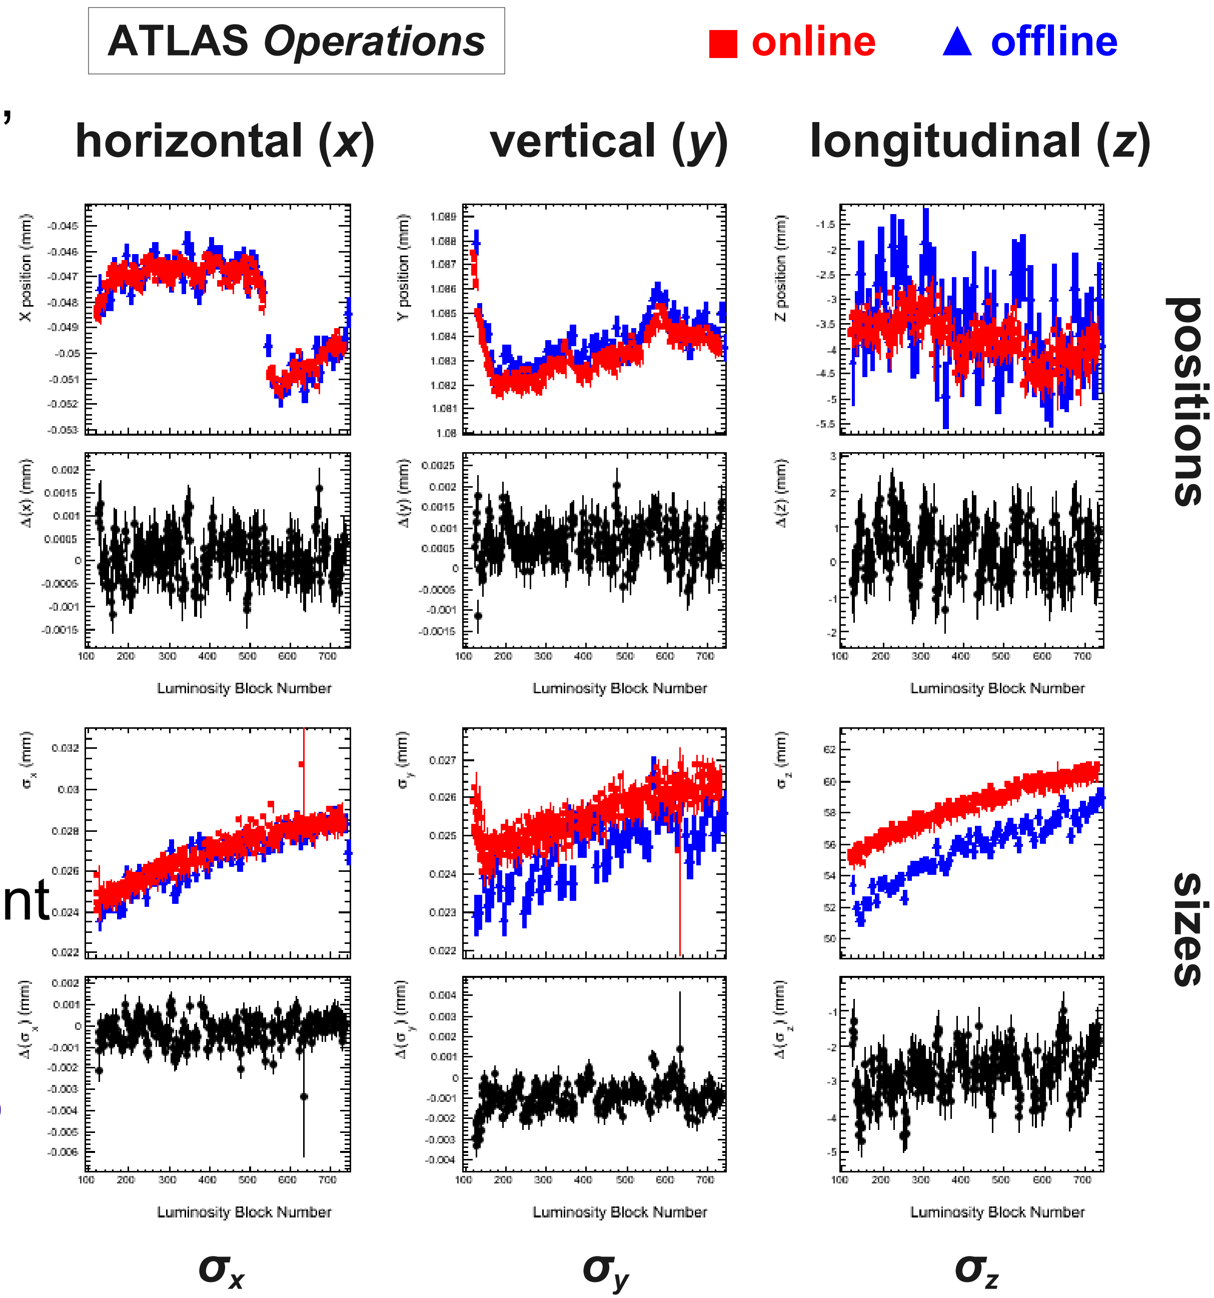

Bunch-by-bunch Monitoring

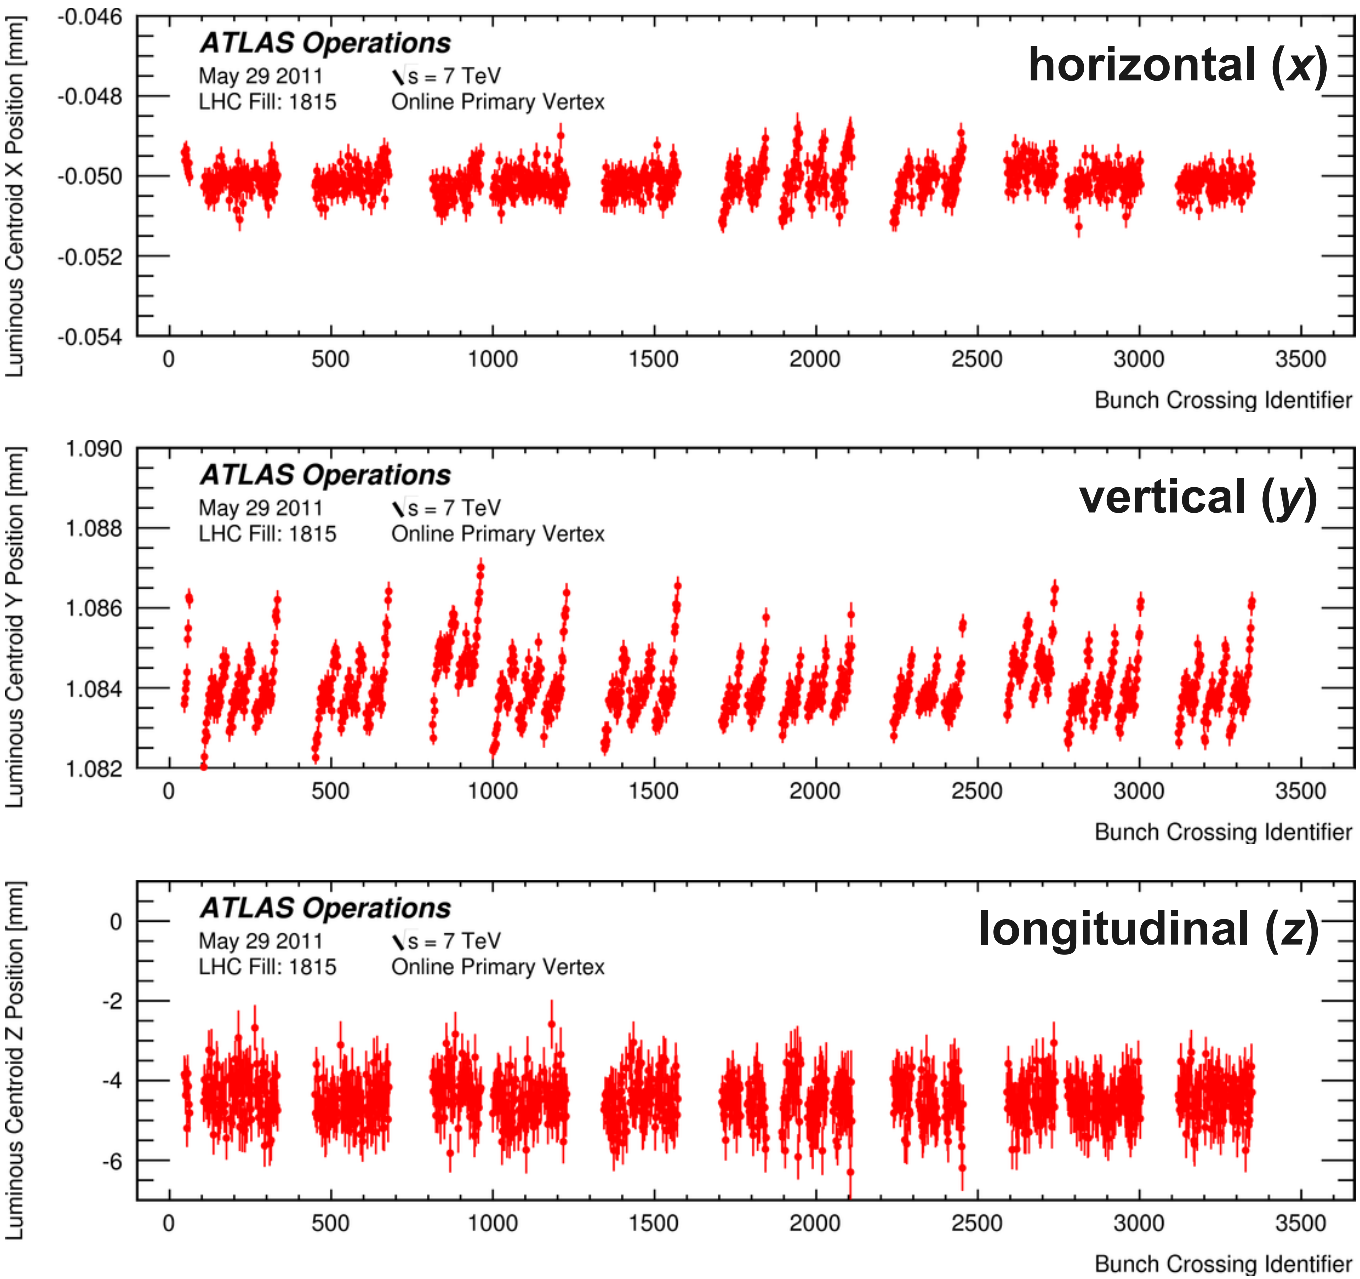

- Luminous centroid positions for 1042 bunches colliding in ATLAS
- Distinct structures are visible across the bunches
  - Particularly in the vertical plane with variation of 5 µm
- Repeating patterns for bunch trains
- Data show entire fill: expect to be able to make a 5% measurement every ~20 min

Conclusions

- The ATLAS High Level Trigger is the first system able to see tracks and vertices, and to allow a reconstruction of the luminous region
- High event rates make this both extremely challenging and precise
- The ability to use (mostly) rejected events is unique to the HLT
- A method to correct for resolution based on vertex splitting has been put in place that is working well down to current spot sizes
- Timelines of luminous region parameters are produced online that provide interesting insights into IP-orbit and RF-phase variations as well as transverse and longitudinal emittance growth
- Owing to the high rate of events, the system is able to do per-bunch measurements even at the current O(1000) colliding pairs
- A feedback system has been developed that transmits live parameters to the LHC and also performs automatic updates of the HLT farm
- In many respects, the ATLAS luminous region measurements can complement LHC beam instrumentation

Outlook

- The online beam spot measurement in the ALTAS High Level Trigger continues to be a challenge
- Just when we managed to adapt to 25 µm spot sizes and 1042 bunches, the LHC keeps pushing the envelope
  - We may expect 1380 bunches in a couple of weeks (and will have to deal with 2808 bunches eventually)
- There is the possibility/hope to further squeeze the emittance while increasing the bunch charges for another factor of ~3 in luminosity
  - This would create on the order or 25 interactions per crossing, making it even harder to reconstruct the even smaller beam spot
- We have to work to turn the higher event rates into an advantage and refine and improve our methods (*i.e.* increase the resolution)
- Most of the instrumentation we added was not only not foreseen in the original design but seemed daring or impossible to do...
  - Yet the flexible architecture of the ATLAS online left room for new ideas
- We just have to rely on measurement capabilities to improve along the way and on new ideas to exploit them!

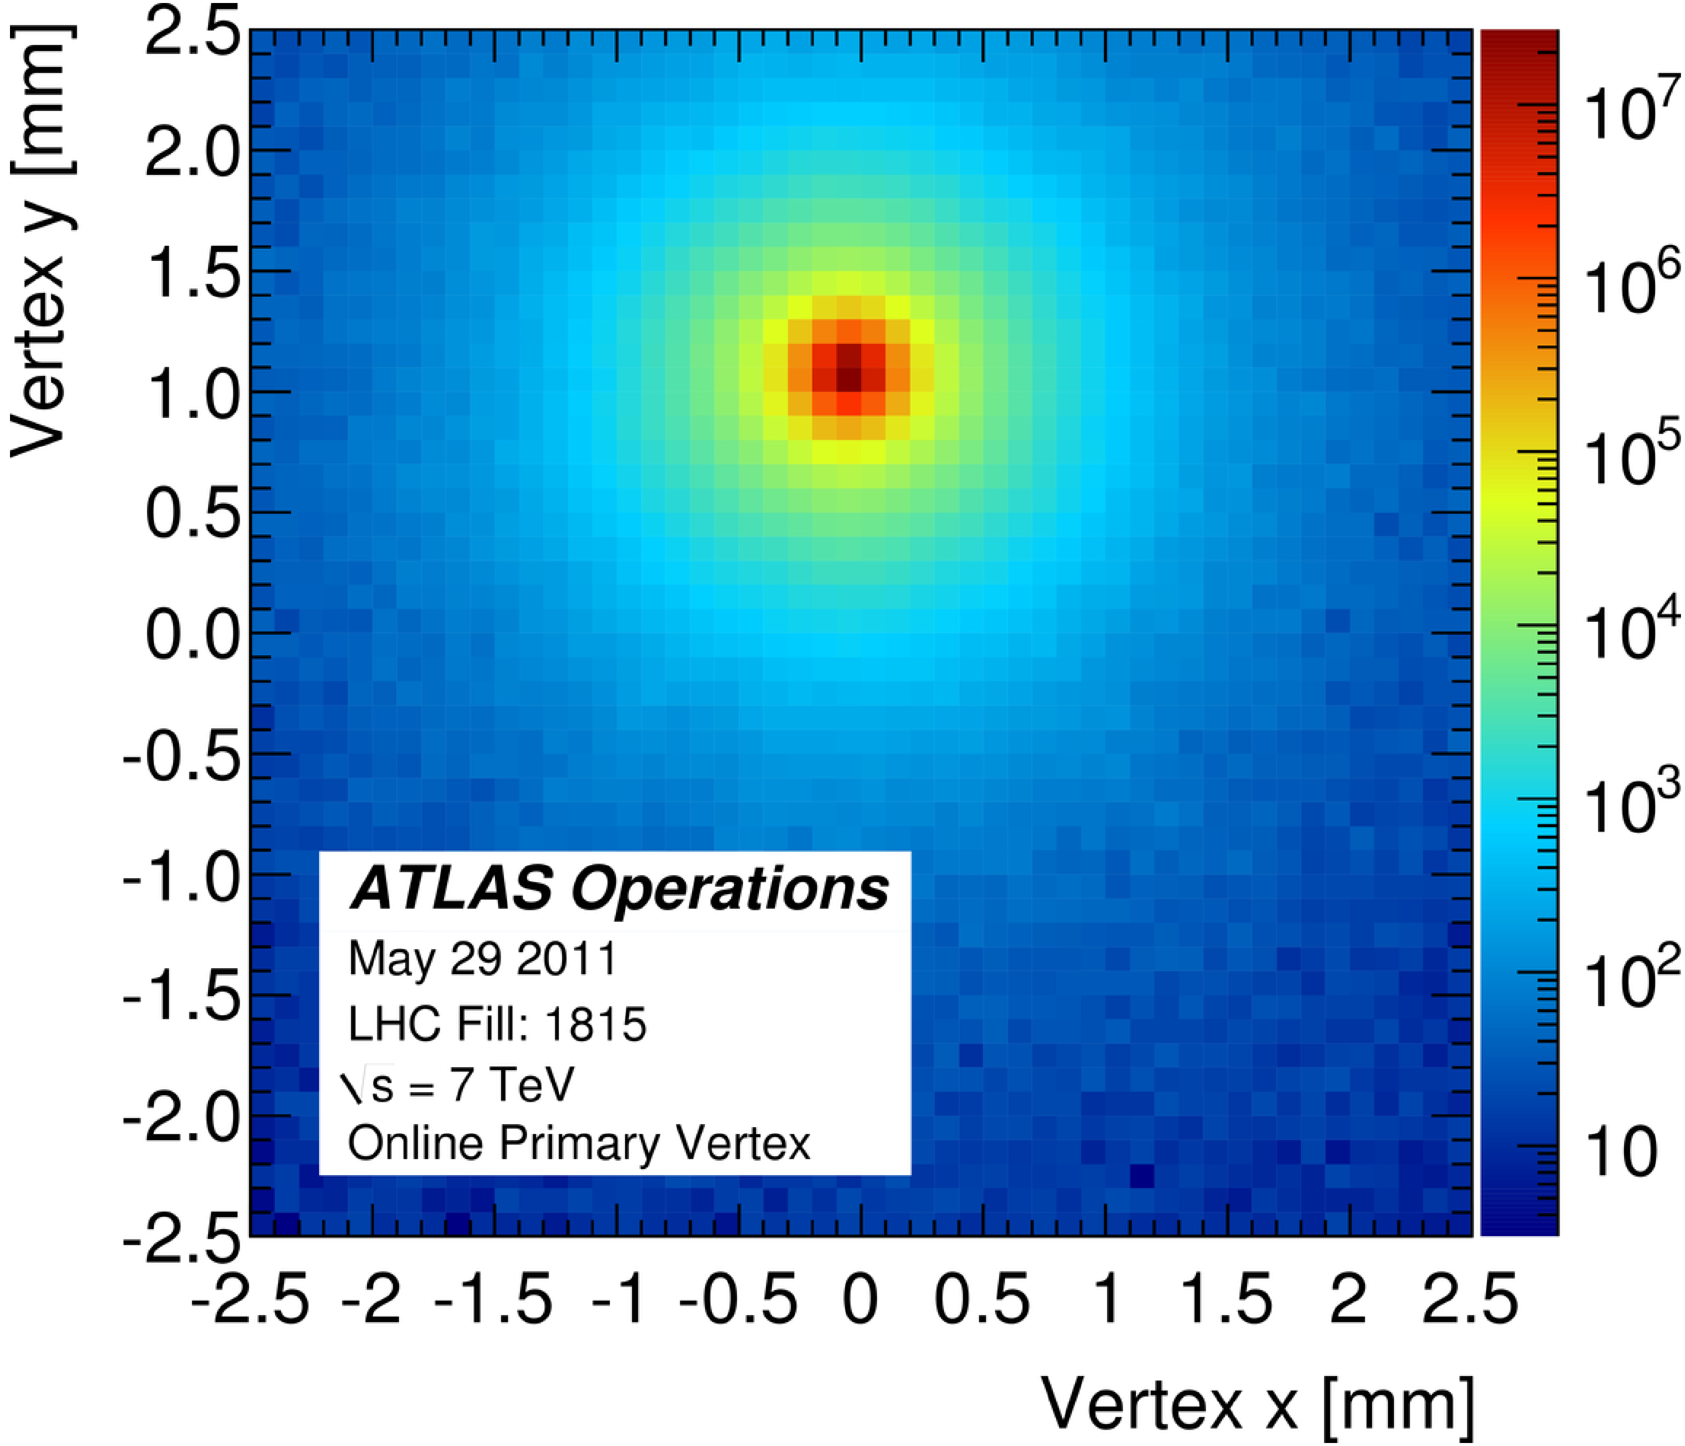
Vertex y [mm]
2.5
2.0
1.5
1.0
0.5
0.0
-0.5
-1.0
-1.5
-2.0
-2.5
-2.5 -2 -1.5 -1 -0.5 0 0.5 1 1.5 2 2.5
Vertex x [mm]
10^7
10^6
10^5
10^4
10^3
10^2
10
ATLAS Operations
May 29 2011
LHC Fill: 1815
√s = 7 TeV
Online Primary Vertex

Backup Material

Primary Vertices