Online Measurement of LHC Beam Parameters with the ATLAS High Level Trigger

Advanced Computing and Analysis Techniques in Physics Research

Emanuel Strauss
estrauss@slac.stanford.edu

On Behalf of The ATLAS Collaboration

SLAC NATIONAL ACCELERATOR LABORATORY

September 5th, 2011

Measurement of the LHC Beam Parameters

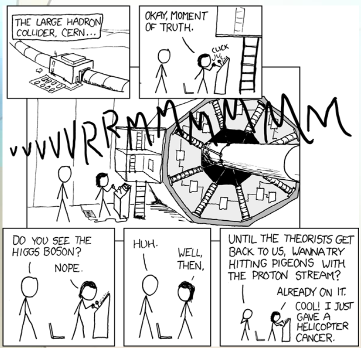

Stolen from XKCD: http://xkcd.com/401/

- Reconstruct the spatial distribution of *pp* interactions, globally and per-bunch-crossing (centroid x,y,z, width σ_x, σ_y, σ_z, and tilt θ_{xz}, θ_{yz})
- Use tracks reconstructed by the Inner Detector Tracking
- Perform online in the High Level Trigger
  - High-rate, large number of usable events which will ultimately be rejected by the trigger system
  - Time and resource limited
- Results are used by the ATLAS trigger system (*b*-jet triggers), and feedback is sent to the LHC operators in near real time.
- Requires live updates of the HLT parameters
  - Position input to tracking and track selection
  - Position and width input to b-jet triggers

LHC and ATLAS

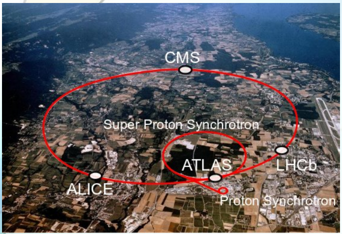

- pp collisions at $\sqrt{s} = 7$ TeV
- Bunch crossings every 50 ns
- Maximum 2808 bunches (at 25 ns)– typically ~ 1300
- Per-Bunch $\sim 1.2 \times 10^{11}$ protons

Precision Tracking and Vertexing:

- 3 cylindrical layers and 2x3 disks of silicon pixels
- 4 cylindrical layers and 2x9 disks of silicon strips
- Begins a few centimeters from beam, extends to radius of 1.2m
- Measures tracks up to $|\eta| < 2.5$

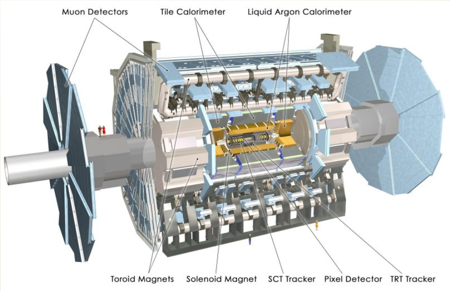

ATLAS TDAQ Design

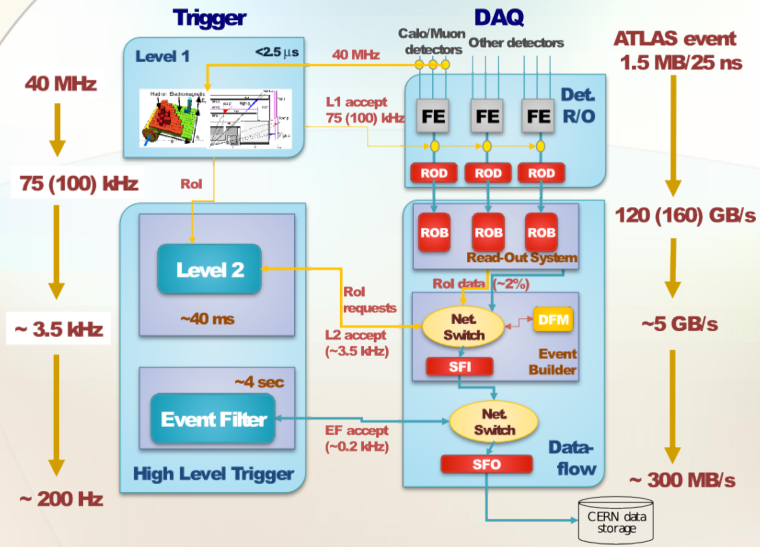

Original parameters from the ATLAS Technical Design Report

Luminous Region

Defined as the region in which the two beams collide

- The luminous region of two Gaussian beams is itself a Gaussian
- Luminous-centroid motion mirrors IP-orbit and RF-phase drifts
- Precise determination of global beamspot parameters using tracks from primary vertices
- Input rate of 1.5 kHz means $\sim$ 100,000 usable vertices (≥ 6 tracks) per minute
  - Means $\sim$ 1Hz per bunch-crossing

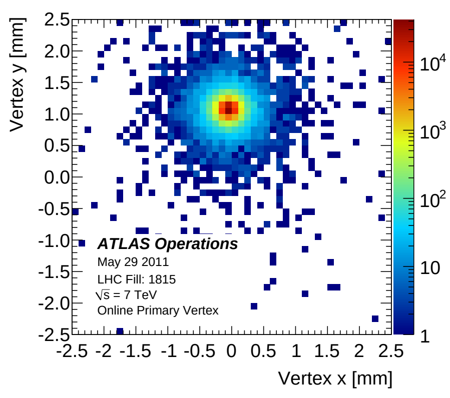

Distribution of primary vertices in the transverse plane for one minute of data-taking.

Luminous Region

Defined as the region in which the two beams collide

- The luminous region of two Gaussian beams is itself a Gaussian
- Luminous-centroid motion mirrors IP-orbit and RF-phase drifts
- Precise determination of global beamspot parameters using tracks from primary vertices
- Input rate of 1.5 kHz means $\sim$ 100,000 usable vertices (≥ 6 tracks) per minute
  - Means $\sim$ 1Hz per bunch-crossing

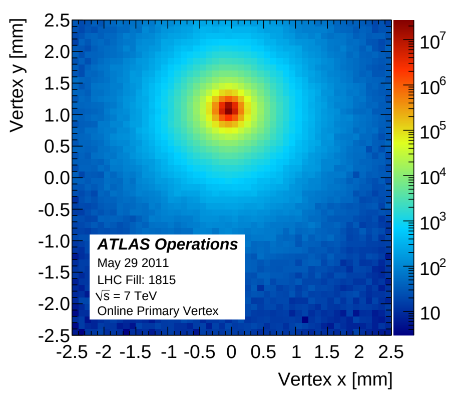

Distribution of primary vertices in the transverse plane for the full run.

Vertex Reconstruction

Beamspot chain seeded from an L1 multi-jet trigger.

- Perform a "full-scan" of the Pixel and SCT detectors
  - Most algorithms only pull smaller Regions of Interest
- Dedicated algorithm reconstructs tracks from hits
- Pre-clustered tracks are fed to a a fast vertex fitter
  - Kalman Filter
  - Uses track covariance directly
  - Avoids costly matrix inversions
  - Requires $\lesssim$ 1ms per vertex
- The results are histogrammed for later calculations

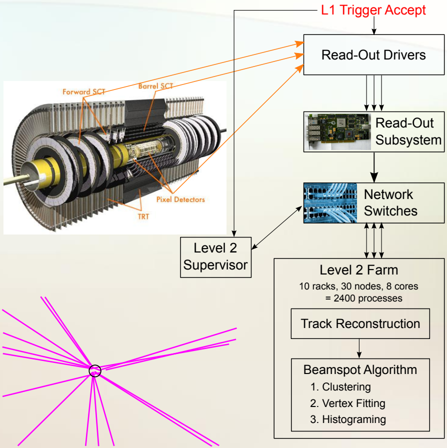

Vertex Reconstruction

Beamspot chain seeded from an L1 multi-jet trigger.

- Primary vertices from multiple collisions (pile-up)
- Multiple measurements of the beamspot per event
- Increases event size and cost
- Possible systematic bias from assigning tracks to the wrong vertex

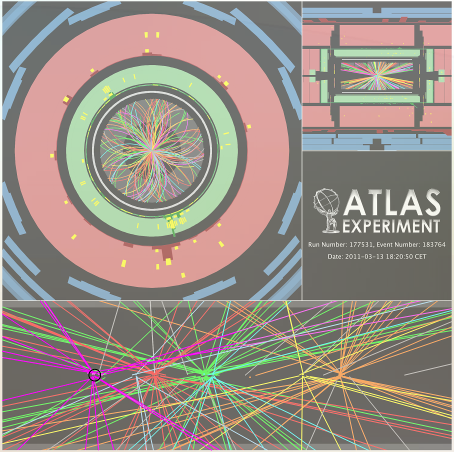

Beamspot Algorithm Environment

At Level 2, a massive farm of CPUs process the candidate events using custom software algorithms

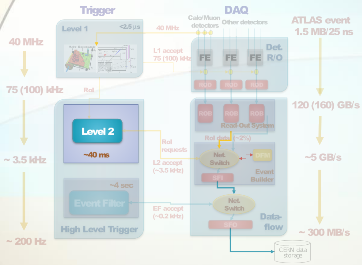

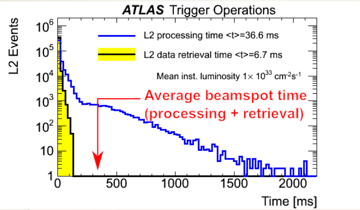

- Partial reconstruction of the events, using only the inner tracking
- Optimized requests: beamspot asks first, caches the data for other algorithms
- Beamspot algorithm takes $\sim 250\mu$s, full-scan tracking takes $\sim$ 300ms
- Future focus: extracting more out of each data-request

L2 imposes many constraints, but the large number of usable events allows for very rapid feedback (μm precision per-minute), and the ability to measure quantities for each bunch crossing (5% precision every 20 minutes).

Beamspot Projections

The positions of the vertices are histogrammed, then collected at regular intervals to perform fits:

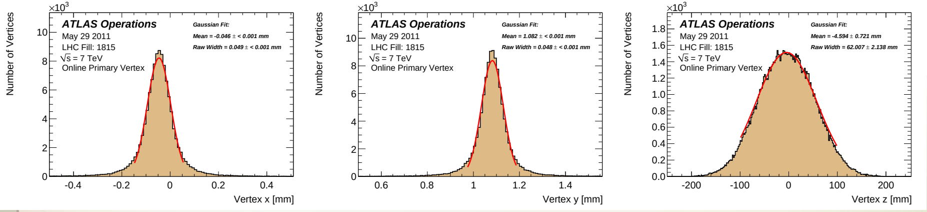

- Position and widths are measured along the x, y, and z directions
- Tilts are measured from the slope of the xz and yz profiles
- x and y true beamspot widths are $\sim 20\mu$m, z is ~ 50mm
- Resolution is $\sim 10\mu$m in xy, $\sim 100\mu$m in z
  x and y are very-sensitive, whereas z is largely impervious

Vertex Resolution

Vertex distribution mirrors the beamspot parameters, convolved with detector resolution

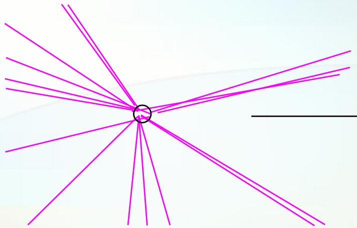

- Sort the tracks in ϕ and split tracks of each vertex alternately into two collections
- Re-fit the two split vertices, storing the difference in x, y, and z
- Estimate the resolution by fitting the distance (Δ_{ij}) between the split vertices
- Subtract the resolution in quadrature from the Gaussian width of the original distribution
- Report the "true" beamspot width

In Situ Resolution Correction

For each split vertex, fill a 2D histogram with Δ_{ij} and # Tracks

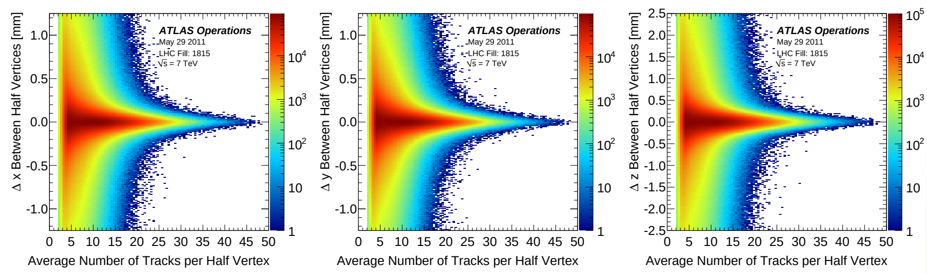

- The resolution is extracted with Gaussian fits in slices of # Tracks ≥ 6
  - Higher track multiplicities lead to sharper resolutions
  - Lower track multiplicities are biased by fakes and mis-matched tracks
- $\sigma_{beamspot} = \sqrt{\sigma^2_{gaus} - \sigma^2_{res}}$
- Resolution (per # Tracks) does not vary as a function of time for a given fill, so events are accumulated as the fill progresses
  - Several updates at start of fill, during boot-strap
  - Few updates during rest of fill, due to emittance growth, drift, orbit corrections

Impact of the Resolution Correction

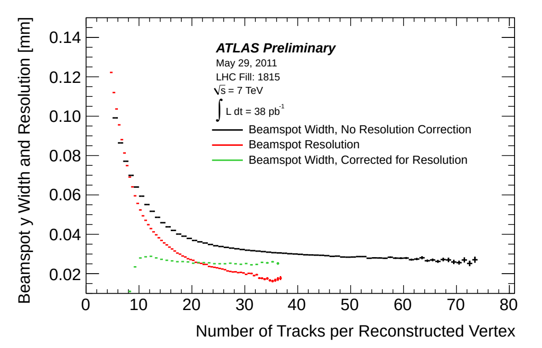

The observed vertical width, measured resolution, and resolution-corrected width versus track multiplicity

Input trigger is carefully selected to provide a sample rich in high track-multiplicities (ie. high-resolution events)

- More tracks lead to better resolutions
- Fewer tracks can lead to systematic biases

  Also more susceptible to effects from background tracks, eg. non-gaussian tails
- Beamspot width is measured by a fit within the stable region

- Method breaks down when $\sigma_{beamspot} < \sigma_{res}/2$

- An n track measurement requires enough statistics with $2n$ tracks

Beam Parameters Versus Time

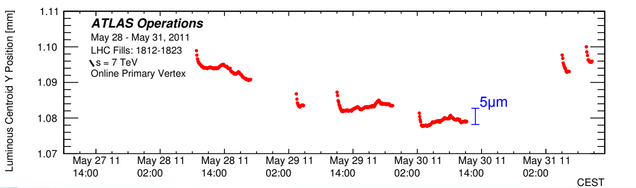

Vertical beam position

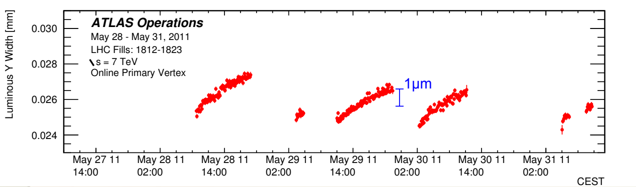

Vertical beam gaussian width

- Track drift over the course of a fill (eg. smooth changes)
- x,y position are sensitive to IP-orbit corrections by the LHC (eg. sizeable jumps)

Highly dynamic values over a broad range

- Very precise width measurement, despite lost statistics from split-vertex method
- Track emittance blow-up of the beam over length of fill in x, y, and z
- Changes in β^* can change the width between fills

σ goes as $\sqrt{\beta^*}$. β^* has evolved from 11m to 1.5m and now to 1m as we speak (18μm luminous size!)

Re-distribution to the Trigger System

It is crucial that the trigger system be given updates of the beamspot parameters

- The track reconstruction is performed in a large window around the stored beamspot position
- Changes in the position, width, and tilt affect the triggers relying on *b*-jet probability
  - *b*-jet triggers use the transverse impact parameter significance $d_0/\sigma(d_0)$ of tracks within a jet
- The width has a significant impact on the accept rate of the triggers.
- Automatic updates are issued, comparing the measured values against the currently stored values

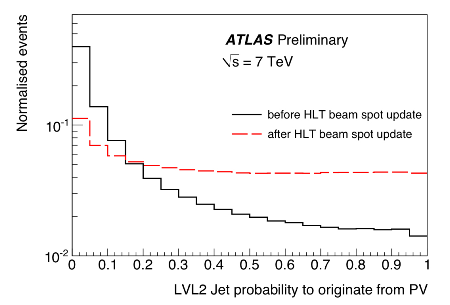

Probability for the L2 tracks to originate from primary vertex before and after the HLT beam spot update

Mis-measurements cause us to waste bandwidth on background events and lose signal events

Pre-fit Optimizations

At $\mathcal{O}(200)$ fits per minute, even simple Gaussians can become costly

- Global fits every minute for monitoring and every 10 minutes for HLT updates
  $\sim$ 60 track bins $\times$3 directions $\times$2 histograms
- Per-BCID fits every 20 minutes
  $\sim$ 1300 bunch crossings $\times$3 directions + 3 resolution histograms

Save some CPU cycles by:

- Requiring $\geq$ 100 events per histogram
- Fitting within $\pm 1.8\times$ RMS
- Re-binning histograms until the average error per-bin is $\leq$ 20%
- Throw out histograms with less than 6 bins in the core after re-binning
  - Too many bins and/or background tails can bias towards wider values
  - Fewer bins speeds up the fit
  - Too few bins wipes out shape details

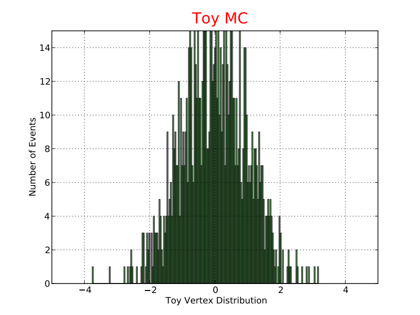

Pre-fit Optimizations

At $\mathcal{O}(200)$ fits per minute, even simple Gaussians can become costly

- Global fits every minute for monitoring and every 10 minutes for HLT updates
  $\sim$ 60 track bins $\times 3$ directions $\times 2$ histograms
- Per-BCID fits every 20 minutes
  $\sim$ 1300 bunch crossings $\times 3$ directions $+$ 3 resolution histograms

Save some CPU cycles by:

- Requiring ≥ 100 events per histogram
- Fitting within $\pm 1.8\times$ RMS
- Re-binning histograms until the average error per-bin is $\leq 20\%$
- Throw out histograms with less than 6 bins in the core after re-binning
  - Too many bins and/or background tails can bias towards wider values
  - Fewer bins speeds up the fit
  - Too few bins wipes out shape details

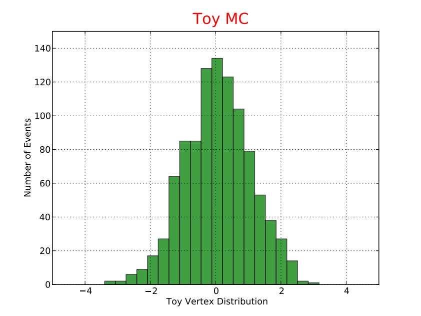

Pre-fit Optimizations

At $\mathcal{O}(200)$ fits per minute, even simple Gaussians can become costly

- Global fits every minute for monitoring and every 10 minutes for HLT updates
  $\sim$ 60 track bins $\times 3$ directions $\times 2$ histograms
- Per-BCID fits every 20 minutes
  $\sim$ 1300 bunch crossings $\times 3$ directions $+$ 3 resolution histograms

Save some CPU cycles by:

- Requiring ≥ 100 events per histogram
- Fitting within $\pm 1.8\times$ RMS
- Re-binning histograms until the average error per-bin is $\leq 20\%$
- Throw out histograms with less than 6 bins in the core after re-binning
  - Too many bins and/or background tails can bias towards wider values
  - Fewer bins speeds up the fit
  - Too few bins wipes out shape details

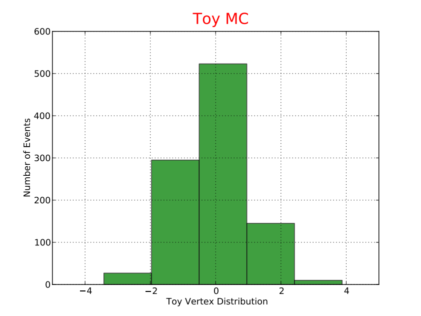

Per-Bunch-Crossing Measurements

Studying the beamspot parameters per-Bunch-Crossing ID reveals important features

- Separate Gaussian fits for $\sim$1300 bunch crossings in x, y, and z.
- Sub-micron precision every 20 minutes for the LHC requires a rate of $\sim$ 1Hz per bunch pair

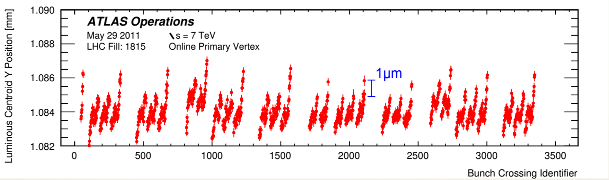

The beamspot position in y of each bunch crossing, measured over an entire fill.

- Variations up to 5 μm in x and y
- Collisions at other IPs cause differences between trains
- Number of long-range collisions (where the beams are $\sim$ 10 RMS apart) causes movement within the train
- Different spread in bunch-to-bunch emittances (or charges) on injection

Per-Bunch-Crossing Measurements

Studying the beamspot parameters per-Bunch-Crossing ID reveals important features

- Separate Gaussian fits for ~1300 bunch crossings in x, y, and z.
- Sub-micron precision every 20 minutes for the LHC requires a rate of ~ 1Hz per bunch pair

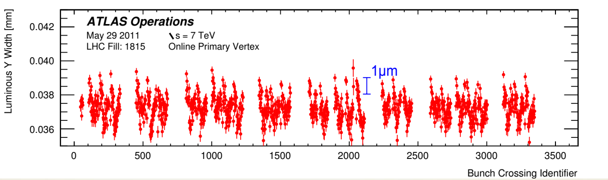

The beamspot width in y of each bunch crossing, measured over an entire fill.

- Variations up to 5 μm in x and y
- Collisions at other IPs cause differences between trains
- Number of long-range collisions (where the beams are ~ 10 RMS apart) causes movement within the train
- Different spread in bunch-to-bunch emittances (or charges) on injection

Per-Bunch-Crossing Measurements

Studying the beamspot parameters per-Bunch-Crossing ID reveals important features

- Separate Gaussian fits for ∼1300 bunch crossings in x, y, and z.
- Sub-micron precision every 20 minutes for the LHC requires a rate of ∼ 1Hz per bunch pair

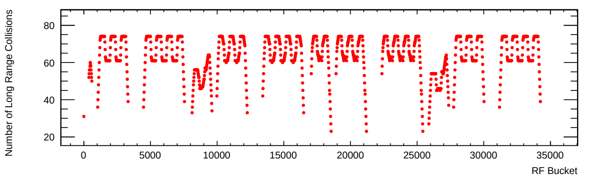

The number of long-range collisions (ie. parasitic collisions where the beams are not fully separated)

- Variations up to 5 μm in x and y
- Collisions at other IPs cause differences between trains
- Number of long-range collisions (where the beams are ∼ 10 RMS apart) causes movement within the train
- Different spread in bunch-to-bunch emittances (or charges) on injection

- Presented an online measurement of the LHC beam parameters in ATLAS using the HLT
- Computational resources are limited, but plenty of statistics
- Measure global beamspot parameters per-minute, corrected for resolution effects
  - Tracking of IP-orbit and RF-phase variations
  - Tracking of emittance blow-up
  - Parameters are fed back to the HLT
- Beamspot parameters are measured per bunch-crossing every 20 minutes
  - Studies of correlation between train structure and long-range collisions/collisions at other IPs
- Rapid monitoring and feedback to the ATLAS trigger system and LHC operators
- https://twiki.cern.ch/twiki/bin/view/AtlasPublic/OnlineBeamSpotPublicResults

Backup Slides

Bunch Crossing Positions

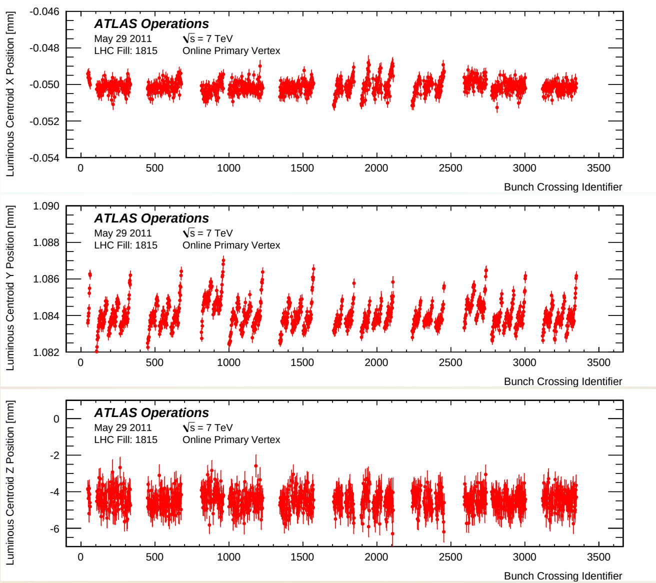

Bunch Crossing Widths

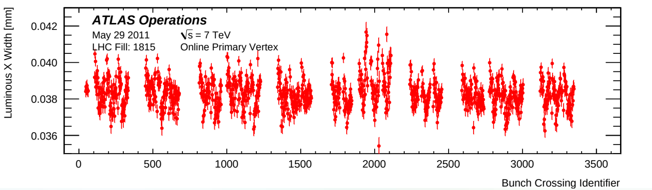

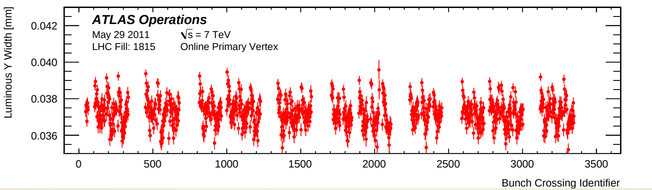

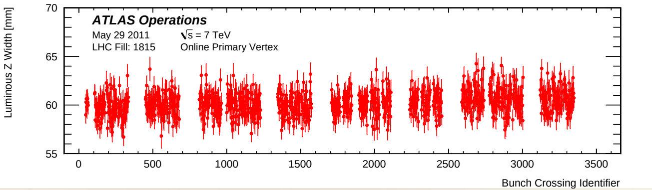

Position For a Single Fill

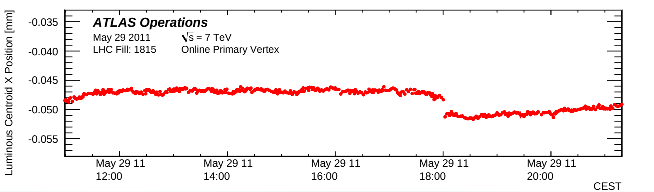

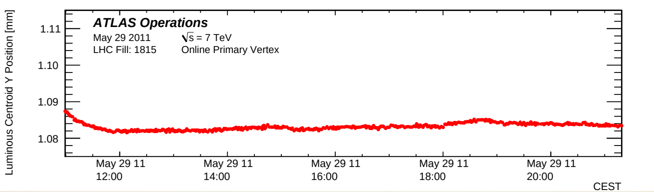

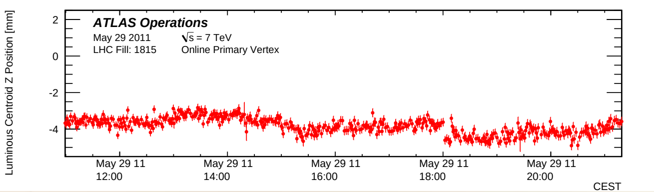

Widths For a Single Fill

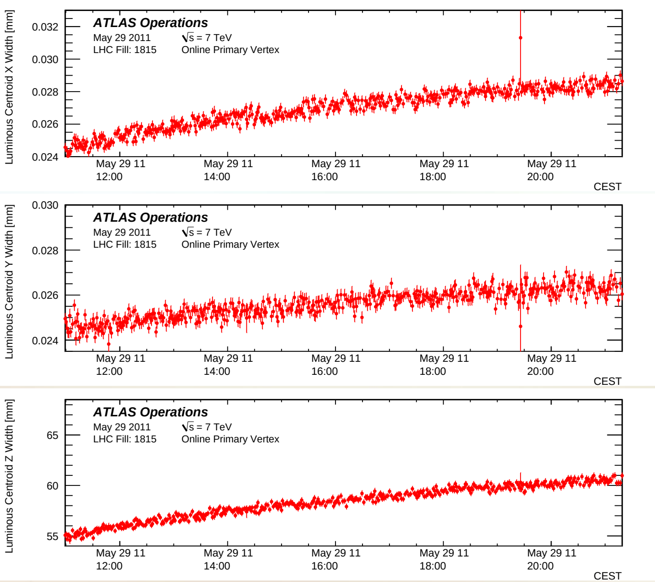

Validating the Resolution Correction

Simulate a smeared gaussian distribution, and the split-vertices, using Toy MC

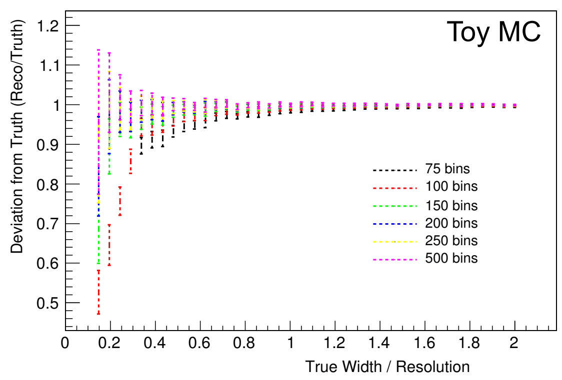

- With too few bins, the shape details are lost and the reported values are too narrow
  - Additional bins for 2D histograms are costly, increase size of transfers from the L2 farm
- To achieve a $\lesssim 10\%$ error, the resolution must be $< 2 \times$ the beamspot width
  - Can get arbitrary resolution by selecting a sample with high enough track-multiplicities, but need enough statistics too!

Resolution Corrections

The observed width, measured resolution, and resolution-corrected width as a function of track multiplicity

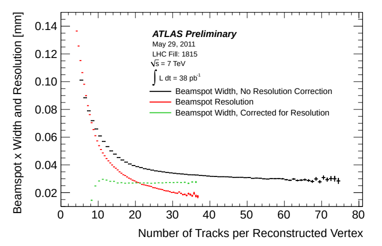

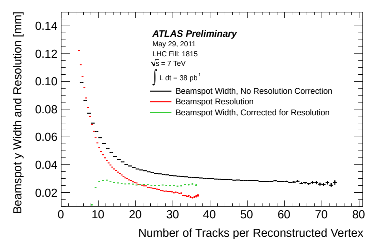

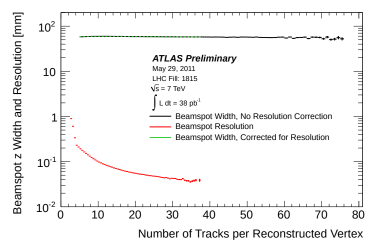

- More tracks lead to better resolutions
- Fewer tracks leads to systematic biases, also more susceptible to effects from background tracks
- Method works down to track multiplicities for which the resolution is not more than twice the true spot size

Re-distribution to the Trigger System

When updating HLT parameters during a run, the values must be passed to each node in the L2 Farm on a sharp time boundary

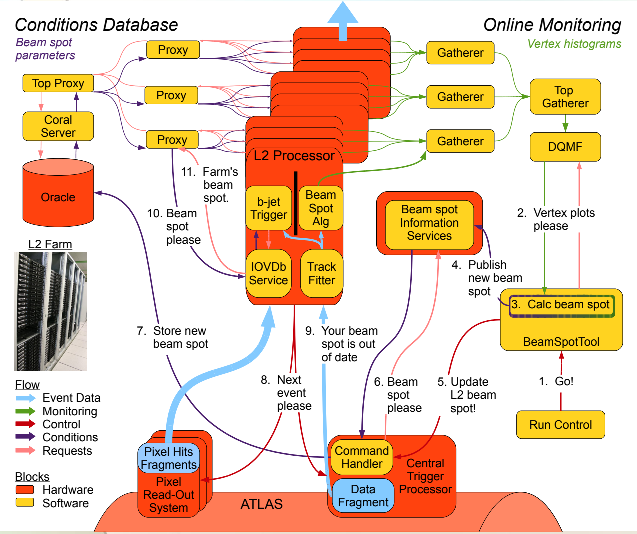

- The parameters in the conditions database are updated
- The Central Trigger Processor notifies the L2 farm nodes
- Each node queries the database for new values
- The values are disseminated via a tree of proxies
- Upon the next luminosity block, all nodes are using the new values

LHC Feedback

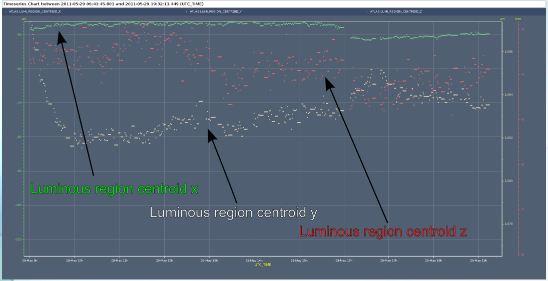

ATLAS beamspot position during Fill 1815, plotted using the LHC Logging Project's *Timber* GUI : http://lhc-logging.web.cern.ch/lhc-logging/software/

LHC Configuration page display with the beam parameters reported by each of the experiments

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

At the end of a fill, detailed values computed in 20 minute intervals are pushed to the LHC operators.

The 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 $\rightarrow$ 4.5 kHz L2 $\rightarrow$ 400 Hz Event Filter (logging to disk)
- L2 does partial reconstruction (First trigger with access to Si-tracker data)
- Currently 10 racks of 30 nodes with 8 cores each (2400 processes) for L2, 27 racks for all of HLT