

Online measurement of LHC beam parameters with the ATLAS High Level Trigger

E Strauss

SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025-7015

E-mail: estrauss@slac.stanford.edu

Abstract. We present an online measurement of the LHC beamspot parameters in ATLAS using the High Level Trigger (HLT). When a significant change is detected in the measured beamspot, it is distributed to the HLT. There, trigger algorithms like b-tagging which calculate impact parameters or decay lengths benefit from a precise, up-to-date set of beamspot parameters. Additionally, online feedback is sent to the LHC operators in real time. The measurement is performed by an algorithm running on the Level 2 trigger farm, leveraging the high rate of usable events. Dedicated algorithms perform a full scan of the silicon detector to reconstruct event vertices from registered tracks. The distribution of these vertices is aggregated across the farm and their shape is extracted through fits every 60 seconds to determine the beamspot position, size, and tilt. The reconstructed beamspot values are corrected for detector resolution effects, measured *in situ* using the separation of vertices whose tracks have been split into two collections. Furthermore, measurements for individual bunch crossings have allowed for studies of single-bunch distributions as well as the behavior of bunch trains. This talk will cover the constraints imposed by the online environment and describe how these measurements are accomplished with the given resources. The algorithm tasks must be completed within the time constraints of the Level 2 trigger, with limited CPU and bandwidth allocations. This places an emphasis on efficient algorithm design and the minimization of data requests.

1. Introduction

In 2011, ATLAS collected over 5 fb^{-1} of data from collisions produced by the Large Hadron Collider (LHC) at center-of-mass energy $\sqrt{s} = 7 \text{ TeV}$. The collisions within the luminous region, or beamspot, follow a Gaussian distribution along the x -, y -, and z -axis¹ determined by the interaction region of the two beams. Interaction Point (IP) parameters typical of the results reported in this paper are listed in Table 1. As the LHC lacks instrumentation close to the IP, the ATLAS detector provides the best measurement of the beamspot parameters. Not only is this valuable monitoring information for the LHC operators, it is crucial for running several ATLAS triggers with high efficiency and purity.

Using tracks measured by the ATLAS Inner Detector (ID), we reconstruct the position of primary vertices in three dimensions. The mean of the vertex positions in x , y , and z mirrors the center of the luminous region as it drifts over the course of a fill. The widths of the

¹ The ATLAS reference system is a Cartesian right-handed coordinate system, with nominal collision point at the origin. The anti-clockwise beam direction defines the positive z -axis, with the x -axis pointing to the center of the LHC ring. The pseudo-rapidity is defined as $\eta = -\ln(\tan(\theta/2))$, where the polar angle θ is taken with respect to the positive z direction.

\sqrt{s} (TeV)	7.0
Typical Number of Protons per Bunch	$1.2 - 1.5 \times 10^{11}$
Half-Crossing Angle of the Beams	120 μ rad
Typical IP β -function (m)	1.0 and 1.5
Typical Number of Bunches	1000 - 1300
Peak Instantaneous Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	3.65×10^{33}
Typical Number of Collisions per Bunch Crossing	10 to 20

Table 1: The typical beam parameters during 2011 LHC operation.

Gaussian reflect the size of the luminous region, and can track the evolution of the transverse and longitudinal emittances. The measurements we provide must therefore be made both precisely, and quickly enough to track the changes in these parameters.

As we will discuss in later sections, the second level (L2) of the ATLAS High Level Trigger (HLT) provides an excellent environment for making these measurements. Its technical constraints require the use of highly efficient algorithms in order to maximally exploit each event. Despite these constraints, the high rate of events and access to tracking information makes the HLT the best place to perform such a complex measurement with the high frequency and low latency we require. As we will discuss in later chapters, the L2 environment allows for very rapid feedback of the beamspot parameters at the ATLAS interaction point, with precisions at the μm level. With longer time intervals it is also possible to measure the beamspot position and width for each of the bunch crossing with approximately five percent precision.

2. The ATLAS Inner Detector

The primary vertices in the event are reconstructed using the ATLAS ID tracking [1]. The ID has three sub-detector components. The algorithm uses information from the two inner-most, the silicon pixel detector, and the silicon strip detector (SCT), for their excellent spatial resolution. The pixel detector has three layers which sit just outside of the LHC beam-pipe, and 3 disks on each end-cap. Its hit resolution is about 10 μm in $r\phi$, and about 115 μm in z . Just outside of this, the SCT is comprised of 4 cylindrical layers in the barrel and 9 disks on each end-cap. The SCT hit resolution is about 17 μm in $r\phi$ and about 850 μm in z . The ID is immersed in a 2 T solenoidal field, which bends the tracks for accurate p_T measurements. All told, the pixel and SCT covers a radius of approximately 600 cm and measure tracks up to pseudo-rapidity $|\eta| < 2.5$.

3. The ATLAS High Level Trigger

The ATLAS trigger system [2] is comprised of three levels; one hardware (L1), and two software (L2 and EF). The rate of events written to tape is maintained at approximately 300 Hz, five orders of magnitude down from the approximately 20 MHz produced by the LHC during the 2011 data-taking period. L2, where we perform our measurement, is capable of accepting a rate of approximately 75 kHz; of this about 1.5 kHz is allocated to our trigger algorithm.

When the Central Trigger Processor determines that a relevant L1 trigger has fired, an instance of the trigger algorithm gets run by one of the CPUs on a farm of over 6,000 processors. The result is histogrammed locally, and aggregated every 60 seconds across the farm. They are then stored, and made available for processing by tools running within the ATLAS Data Acquisition (DAQ) infrastructure [3].

4. Beamspot Algorithm

The beamspot parameters are determined by reconstructing the distribution of primary vertices produced by the hard-scatter of the beams. Events with these primary vertices are collected by the trigger chain at an input rate of approximately 1.5 kHz. The beamspot measurement is made in several steps:

Tracking in which the hits from the pixel and SCT detectors are used to measure the position and momentum of tracks.

Vertexing in which the tracks are clustered into common primary vertices.

Histogramming in which the distributions of the vertices in x and y directions are stored by the trigger farm nodes before being aggregated.

Measuring in which the aggregated distributions are used to extract the beamspot position, width, and tilt versus time, correcting for resolution effects. This is done in multiple passes, with various time windows, to achieve the statistics necessary for monitoring, as well as global and per-bunch-crossing measurements.

Distribution in which the parameters are published for monitoring, distributed to the HLT nodes as input for other triggers, and broadcast to the LHC operators for feedback.

4.1. Tracking and Vertex Reconstruction

Trigger algorithms at L2 typically operate on regions of interest, identified at L1, to keep the data-requests from the detector read out system low. The beamspot algorithm requests information about all tracks in the event and performs a “full-scan” of all the hits in the ID.

Fast custom algorithms are used to create 3D space-points from these hits before fitting the tracks with a Kalman filter. These algorithms have been tuned to perform this operation very quickly to stay within the stringent L2 timing requirements. The L2 tracking manages to maintain efficiencies over 90% compared to the offline track selection.

Although the clustering and vertexing portion of the algorithm is very fast, typically taking approximately 250 μ s per event, the reconstruction of tracks from the ID takes 300 ms on average and requires relatively large transfers of detector data. For this reason, the trigger steering has been optimized to run the beamspot algorithm first, such that the ID data can be cached for other algorithms.

The trigger algorithm first applies some basic quality cuts to the tracks from the previous step. Those which satisfy the cuts are clustered by impact parameters. The true distribution of the impact parameter along the z axis is much larger than the ID z resolution. As the distribution of tracks is much broader than the transverse distributions, this clustering is a quick and efficient way of grouping the tracks, simplifying the vertexing and reducing the mis-assignment of tracks.

After being clustered, the vertices are measured by a rapid vertex fitter [4]. The fitter uses the track covariance directly, avoiding costly matrix inversions and can reconstruct a vertex in under a millisecond.

The luminosity attained by the LHC in 2011 has been at the cost of “pile-up”. As the beam intensities have increased, so have the number of interactions per bunch crossing, averaging around 20 in the last few weeks of the 2011 run. Though this provides the opportunity to make multiple measurements per event, the higher occupancy increases both the event size and the time it takes to reconstruct all tracks. Worse still, the close proximity of multiple vertices introduces possible systematic biases when tracks are assigned to the wrong vertex. The z window in which tracks are clustered has been tightened to cope with the higher density of vertices. Over the course of 2011 the L1 seed was switched from a Minimum Bias trigger, to a three jet trigger. The three jet trigger has a lower rate, but more tracks per vertex on average. Raising the p_T threshold on the tracks fed to the vertex fitter has helped to remove background, and raising the required number of tracks per vertex has improved the average uncertainty on

the vertex position. These tighter requirements help to maintain a high quality of data produced by our algorithm, despite harder operational conditions.

4.2. Histogramming the Luminous Region

The spatial distributions of the vertices are stored locally in histograms by each of the over 6000 core processes on the trigger farm which run the beamspot algorithm. As described in Section 3, these are then collected across the farm. Here, they are available to a tool in the DAQ infrastructure that uses them to determine the beamspot parameters. By summing consecutively produced histograms, the tool can make measurements for arbitrary time periods. The position can be precisely measured as the mean of a Gaussian function, fit independently to the x -, y -, and z -projections of the vertex distribution, as shown in Figure 1. The width of the distribution of vertex positions is the beamspot width, convolved with the detector resolution. We describe in the next section how a resolution-corrected measurement is made. The tilt with respect to the ATLAS coordinate system is measured from the slope of the xy and yz profiles.

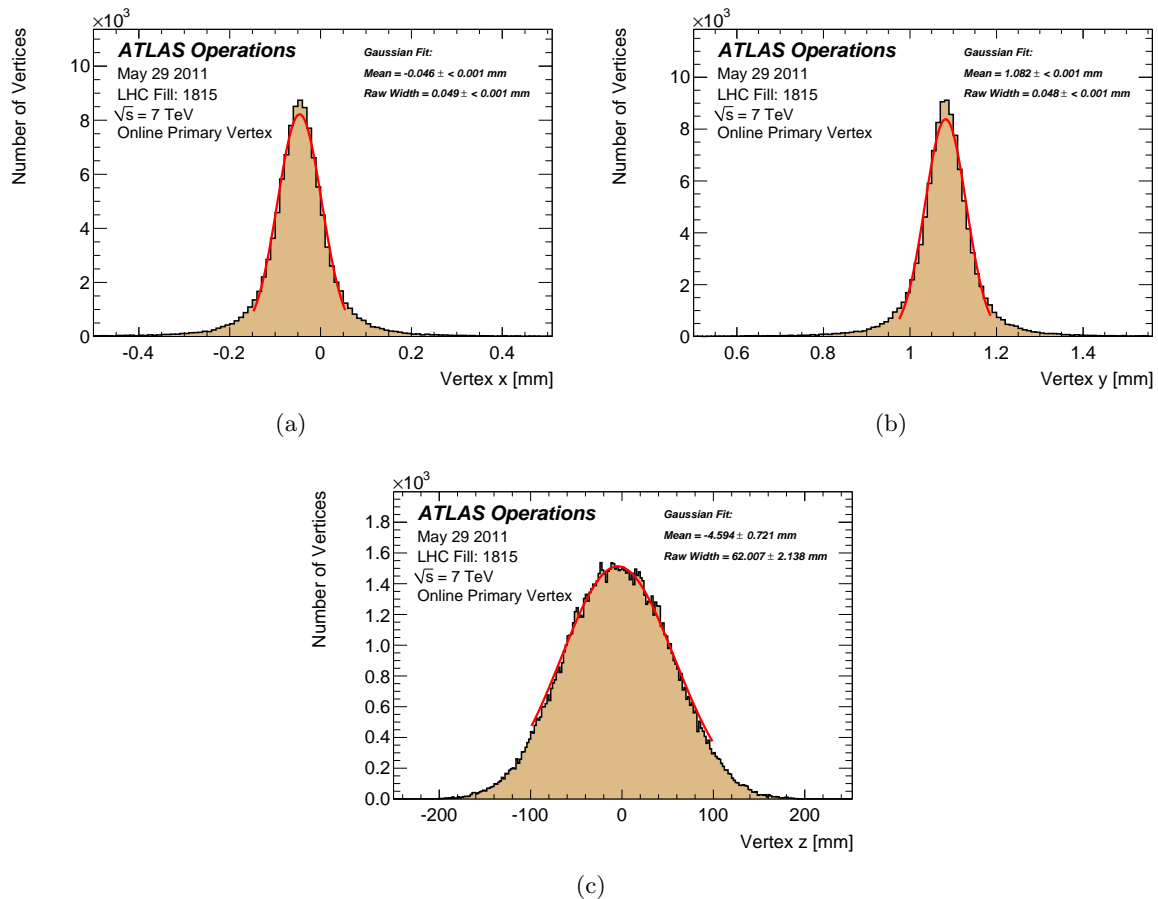


Figure 1: The x (a), y (b), and z (c) distribution of primary vertices reconstructed online with ≥ 6 tracks in the High Level Trigger in 1 minute of data-taking.

4.3. In Situ Resolution Correction

The true beamspot width in 2011 was roughly 15 to 45 μm in x and y , and 25 to 75 mm in z . The z vertex resolution is on the order of 100 μm , so the resolution is not a significant

consideration in the measurement. This is not the case in x and y where the resolution is on the order of $10\ \mu\text{m}$.

A split-vertex method is used to remove the effect of detector resolution on the beamspot width measurement in all three directions. The tracks which have been used to fit the primary vertices are sorted in ϕ , and then assigned alternately into two sub-clusters for a new round of vertex fitting. These tracks are expected to come from the same true vertex, the differences in their reconstructed vertex is assumed to be entirely due to resolution effects. The true beamspot width, σ_{beamspot} , can therefore be measured as $\sigma_{\text{beamspot}} = \sqrt{\sigma_{\text{raw}}^2 - \sigma_{\text{res}}^2}$ where σ_{raw} is the width of the primary vertex distribution and σ_{res}^2 is $1/\sqrt{2}$ times the width of the displacement between the two split-vertices. Not all vertices are created equal though, the more tracks in the vertex, the smaller the uncertainty. The primary vertex distributions and the split-vertex displacements are stored in 2D histograms against the number of tracks per-vertex so that the correction can be applied as a function of track multiplicity. Figure 2a shows the narrowing distribution at higher track multiplicities as well as the rapid drop-off in statistics.

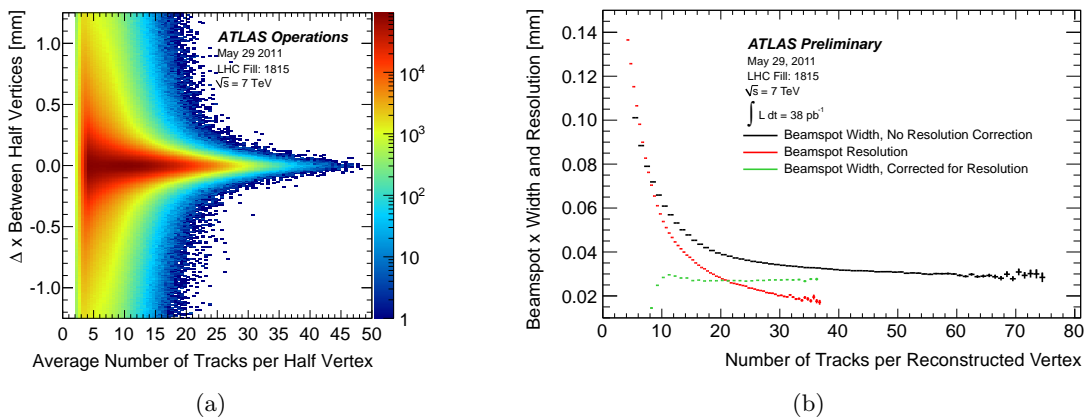


Figure 2: The displacements in x between each pair of separately reconstructed split-vertices as a function of the number of tracks, measured online in the High Level Trigger (a). The observed width of the primary vertex distribution, the measured resolution, and the resolution-corrected width, all as a function of the number of tracks per vertex in x (b).

The split-vertex method has a few limitations. The measured resolution is biased towards larger values when the true beamspot width is less than half the size of the resolution. This can be compensated for by restricting ourselves to vertices with large track-multiplicities. At the beginning of the year we required six tracks per vertex, as the beam widths narrowed, we raised the threshold to ten tracks per vertex.

Unfortunately, the split-vertices probe the resolution for vertices with half as many tracks as the primary vertex. For a sample of reconstructed vertices with up to n_{Max} tracks per vertex, in the best case all those above $n_{\text{Max}}/2$ will effectively be thrown out as we do not have an estimate of their resolution. Figure 2b demonstrates this. It shows the width of the primary vertex width, the resolution measurement, and the resolution-corrected beamspot width as a function of track multiplicity for the x -axis. The primary vertex width distribution extends up to 75 tracks per reconstructed vertex, whereas the resolution and corrected width only extend up to 36.

Finally, the systematic uncertainties and background effects at low track multiplicities can cause large deviations in the resolution-corrected width measurement. As the size of these effects diminish, the corrected width distribution flattens, typically as it approaches the size

of the resolution measurement. At higher track multiplicities, the statistical variations in the resolution measurement become more significant.

Despite these limitations, the values produced by the online algorithm have been found to be in good agreement with those expected based on LHC instrumentation, as well as those measured with an alternate method offline [5]. Figure 3 gives a sense of the small statistical uncertainty achieved even on short time intervals.




Figure 3: Time-variation of the luminous centroid position (a) and resolution-corrected width (b) in y measured in the High Level Trigger every five minutes, for six separate LHC fills recorded over the span of four days

4.4. Measurement and Re-distribution

As the online track reconstruction is performed in a window around the stored beamspot position, incorrect values can significantly decrease the track reconstruction efficiency. Jet triggers which calculate b -jet probability [6] are another major client of the beamspot parameters. These typically use the transverse impact parameter significance of tracks within a jet to identify those which were produced by the decay of a b -quark.

A mechanism was developed to issue automatic updates to the HLT conditions database. The values measured using 10 minutes of integrated data are compared against those currently in the database. If the newly computed values and the database values differ significantly, the new values are published to the database for the the HLT nodes to retrieve. Such an update is issued if any the following conditions are met:

- the position along any of the three axes changes by more than 10%,
- the width along any of the three axes changes by more than 10%,

- or there is a reduction of 50% in the statistical uncertainty of any of the values.

Typically, this means frequent updates at the start of the run, as the resolution is boot-strapped, followed by a period of infrequent updates, dominated by the beam drift and emittance blow-up.

4.5. Measurements Per Bunch Crossing

A third set of measurements is made every 20 minutes for the LHC operators. The longer integration time is needed to accumulate enough statistics to measure the beamspot parameters for each of the approximately 1300 individual bunch crossings (BC) to sub-micron precision on the position. This is a feature which only the online beamspot tools are able to provide due to our high rate of usable events. The large statistics required place it just at the limit of our capabilities, requiring approximately 1 Hz per BC. Even so, there are not enough events to measure the per-BC parameters as a function of track multiplicity. A coarser method is used, in which the vertex parameters are histograms in 1D, summing all track multiplicities. An example of such a measurement can be seen in Figure 4 where distinct structures are visible. In particular the vertical position of the interaction-point shows variations of up to $5 \mu\text{m}$ and repeating patterns across the injected bunch trains [7, 8].




Figure 4: The luminous centroid position (a) and width (b) in y measured in the High Level Trigger farm for each of the 1024 colliding bunch crossings separately.

Using a single Gaussian fit to describe the sample introduces some systematic biases, as the primary vertex distribution is actually a sum of Gaussian distributions where events with fewer tracks dominate but have broader widths. The resolution correction is applied by producing a 1D histogram of split-vertex differences, where each event is given a weight such that the track multiplicity distribution of the split-vertices mirrors that of the primary vertices. This ensures

an equivalent event mixture is used in both measurements, reducing the size of the systematic bias.

5. Pre-fit Optimizations

The beamspot values are computed every minute for monitoring purposes, every 10 minutes for the HLT, and every 20 minutes for the LHC operators. Between these three sets of processes, the beamspot tool must perform on the order of 200 fits per minute. At this rate, even simple Gaussian fits become quite costly. We apply various cost-cutting approaches to keep the processing time reasonable.

- (i) Fits are only performed on histograms containing at least 100 entries. This requirement ensures a reasonable statistical uncertainty on the fit parameters, and saves us from wasting time on histograms whose entries are too sparse.
- (ii) Fits are only performed within ± 1.8 RMS of the histogram mean. The distribution outside of this interval contains a larger fraction of background events which are not Gaussianly distributed.
- (iii) Histograms are re-binned before the fit, such that the average error per bin in the core is less than 20%. In sparsely populated histograms, too-high granularity causes the fit to take longer to converge. The re-binning also smooths out the effects on the RMS of a few bins far out in the tails. A more stringent requirement on the statistical uncertainty would cause further re-binning which might wipe out details about the shape of the distribution.

6. Conclusions

We have used the ATLAS ID and HLT to measure the beamspot parameters in quasi-real-time. This feat required several components to work: optimized data-requests from the ID, fast reconstruction algorithms in the HLT, aggregation of histograms from the trigger farm at fixed time intervals, and rapid measurements by a tool running in the DAQ system. By measuring the parameters every 60s we provide precise monitoring of the beamspot to ATLAS. On a longer time-scale, the parameters are distributed to the HLT to maintain high tracking and b -jet tagging efficiency by continuously adjusting the vertex position. The LHC operators benefit from our tool as well, as the the global beamspot and the beamspot per-bunch-crossing are published for them every twenty minutes. The results from the online determination have been cross-checked against expectations from the parameters of the two beams, as well as the offline beamspot reconstruction results. Data taking in 2011 provided unique challenges and development is ongoing on the algorithm and the data selection to prepare for those which 2012 will bring.

References

- [1] G. Aad, et al., The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003.
- [2] The ATLAS Collaboration, G. Aad et al., Performance of the ATLAS Trigger System in 2010, submitted to EPJC, 2011.
- [3] J. Cogan, ATLAS Online Determination and Feedback of LHC Beam Parameters, 13th International Conference on Accelerator and Large Experimental Physics Control Systems: ICALEPCS 2011, Grenoble, France.
- [4] D. Emeliyanov, A Fast Vertex Fitting Algorithm for ATLAS Level 2 Trigger, PoS ACAT (2007) 058.
- [5] R. Bartoldus, et al., of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC (2010) MOPEC008, 1st International Particle Accelerator Conference: IPAC10, 23-28 May 2010, Kyoto, Japan.
- [6] A. Coccaro, on behalf of the ATLAS Collaboration, Track Reconstruction and b-Jet Identification for the ATLAS Trigger System, these proceedings.
- [7] M. Schaumann, R. Alemany Fernandez, Beam-Beam Induced Orbit Effects at LHC (2011) WEPC081, 2nd International Particle Accelerator Conference: IPAC11, 4-9 September 2011, San Sebastian, Spain.

- [8] W. Herr, R. Alemany, R. Giachino, G. Papotti, T. Pieloni, R. Calaga, E. Laface, M. Schaumann, Observations of Beam-Beam Effects at High Intensities in the LHC (2011) WEODA01, 2nd International Particle Accelerator Conference: IPAC11, 4-9 September 2011, San Sebastian, Spain.