Available online at www.sciencedirect.com
ScienceDirect
Nuclear Instruments and Methods in Physics Research A 582 (2007) 761–765
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

Tracking at level 2 for the ATLAS high level trigger

M. Sutton

Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

Available online 10 October 2007

Abstract

A set of algorithms for fast pattern recognition and track reconstruction using 3D space points for the ATLAS Level 2 (LVL2) Trigger at the LHC is described. The presented algorithms work by first identifying the z position along the beamline of any interesting interaction. This allows selection of groups of spacepoints pointing back to this position using a histogramming technique that greatly reduces the combinatorial overhead and execution time before moving on to fitting the tracks and vertices. Results on the timing and performance of the tracking at LVL2 are presented using both simulated data and real test beam data and cosmic data.

PACS: 29.40.Gx; 29.40.Wk; 07.05.Kf; 07.05.Hd

Keywords: Tracking; Trigger

1. Introduction

The large event rate of the Large Hadron Collider (LHC) makes the online selection of interesting physics events an essential and challenging ingredient required to achieve the physics goals of the LHC. The 40 MHz bunch crossing rate of the LHC means that a rejection factor of more than five orders of magnitude is required to reduce the output rate of the ATLAS High Level Trigger (HLT) to the 200 Hz that can be written and stored offline.

For the initial "low" luminosity ($2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$) phase of LHC operation each bunch crossing will produce around five separate pp interactions while the later "high" luminosity phase ($10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$) will produce around 25 pp interactions per bunch crossing. Consequently, interesting interactions will usually be overlaid by some number of uninteresting, minimum bias, or "pileup" interactions from which they must be disentangled. These pileup interactions increase the occupancy of the detector, a particular problem for the silicon tracking detectors where occupancy is highest. Hits from the overlying pileup interactions increase the probability of incorrectly assigning hits from uninteresting interactions during the online pattern recognition stages. Furthermore, the larger data volume increases the readout and reconstruction latencies so it is important that the trigger algorithms are both fast and robust to ensure that rare and interesting events are not lost.

1.1. The ATLAS trigger and DAQ systems

To achieve the high level of event rejection, the ATLAS trigger [1] has been designed as a three level trigger system. The first level (LVL1) trigger is a hardware pipelined trigger based on custom electronics with reduced granularity information from the calorimeter and the muon system. This is clocked at the bunch crossing rate and delivers a decision after a fixed latency of 2.5 μs. During LVL1 processing regions of interest (RoIs) within the detector are identified that may contain features of interest. On a LVL1 accept the data are read out into custom Readout Buffers (ROBs) where they are stored for access by the level 2 (LVL2) system.

The LVL2 system consists of a farm of fast commodity CPU's running custom algorithms. During the LVL2 processing the algorithms have access to the full granularity detector data but only for the regions of interest identified by LVL1. This reduces the data volume and processing required at LVL2 since the time budget for the LVL2 algorithms is around 10 ms and it is required to

E-mail address: sutt@mail.desy.de

0168-9002/$ - see front matter
doi:10.1016/j.nima.2007.10.005

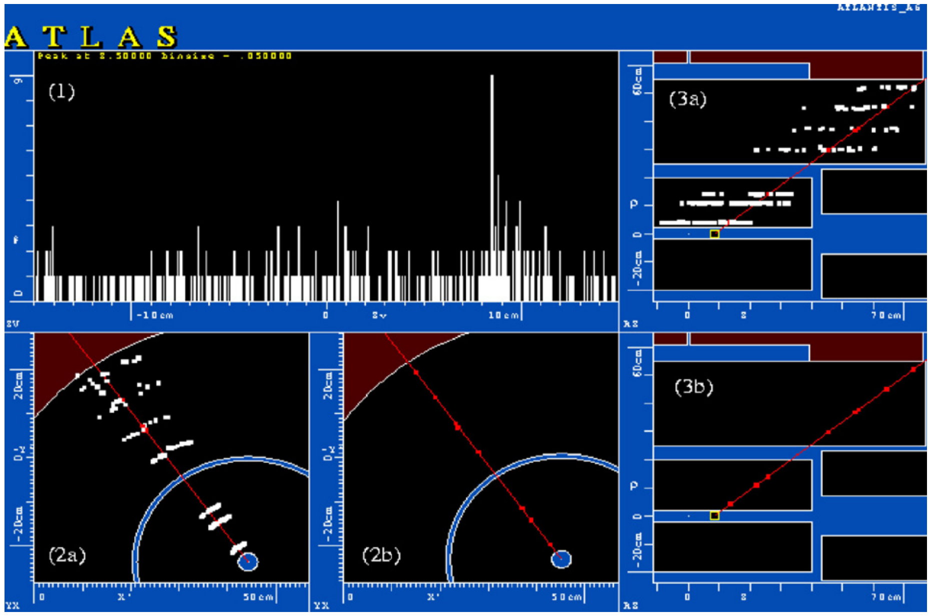

Fig. 1. The ZFinder algorithm for a single electron region of interest.

operate with an input rate of around 75 kHz. The maximum output rate is envisioned to be around 1 kHz.

On a LVL2 accept, the full event is read out and constructed by the Event Builder which passes the data onto the third level Event Filter (EF). This is also a CPU farm, but running seeded versions of the offline reconstruction algorithms with access to the complete detector data and the full alignment and calibration.

The LVL2 trigger is the first stage where information from the tracking detectors is available. At LVL2 it is also possible to combine tracking information from the Inner Detector with calorimeter or muon system data to enable improved particle reconstruction, essential for the identification of signatures, such as high transverse energy electrons, photons and muons or final states containing b-quarks, which may signal new or rare physics processes.

2. The IDScan level two tracking algorithm

The IDScan Algorithm [2] attempts to reconstruct, with high efficiency, high momentum tracks using hits from the silicon layers—Pixel and Semi-Conductor Tracker (SCT)—of the Inner Detector.

Each pp interaction within a bunch crossing will have a different z-position along the beamline. Identifying hits consistent with tracks from vertices at different z will allow the number of hit combinations used for subsequent pattern recognition to be reduced.

The ZFinder [3] divides the RoI into many small slices in ϕ. Lower momentum, curved tracks will in general produce small numbers of hits in several, different ϕ slices, whereas those from high momentum straight tracks will produce most of their hits in a small number of neighbouring ϕ slices. Within each narrow slice all possible pairs of spacepoints within the slice are used to calculate the z position of their intersection with the beamline assuming their projection in the plane of ρ–z is approximately linear.[1] The z positions for each pair is used to fill a one dimensional histogram with the histogram maximum providing the z position of the interaction with the largest number of high momentum tracks. This is illustrated in Fig. 1 for a fully simulated single high $p_T = 25\,\text{GeV}$ electron overlaid with several low momentum pileup interactions without misalignment. The complete set of spacepoints in the RoI are shown in Fig. 1(2a) and (3a). The electron spacepoints are shown in Fig. 1(2b) and (3b) showing hits in each of the seven silicon layers of the Inner Detector barrel. This means that for a single electron there are around 21 hit pairs that could contribute to the z vertex in the histogram. This is shown in Fig. 1(1) showing the identified vertex at $z = 8.5\,\text{cm}$. In practice, the innermost spacepoint of each pair is required to come from the Pixel detector, where the spatial resolution is higher, so that around 15 pairs would be more usual.

For RoIs with more than one track, for example a jet RoI, many more hit pairs pointing to the correct vertex would be expected. By keeping the ϕ slices small, most slices have few pairs, so the algorithm execution time scales quadratically with the small number of pairs within a slice, but approximately linearly with the overall number of spacepoints in the RoI.

For single 25 GeV electrons, the z-vertex resolution is around 150 μm increasing with increasing pseudorapidity. For low luminosity events the efficiency is between 95% and 97%. Fig. 2 shows the resolution of the reconstructed vertex position for single muons as a function of the muon transverse momentum for high luminosity events with

[1] In the ATLAS frame, z lies along the beam direction with ρ being the radius transverse to the beamline.

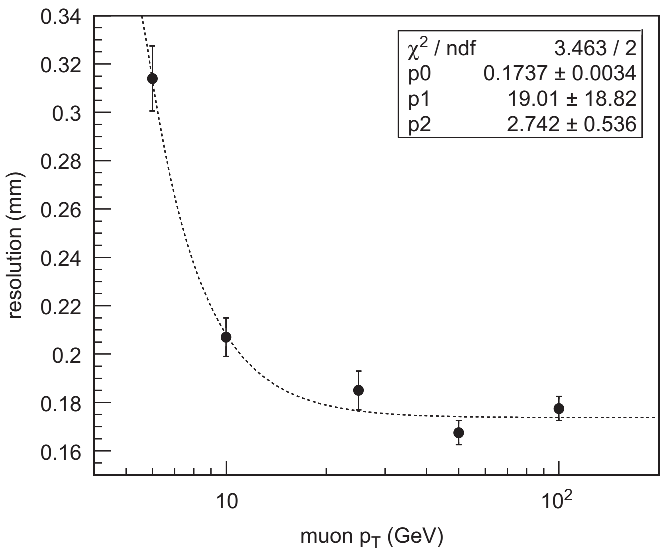

Fig. 2. The ZFinder resolution for muons as a function of muon transverse momentum.

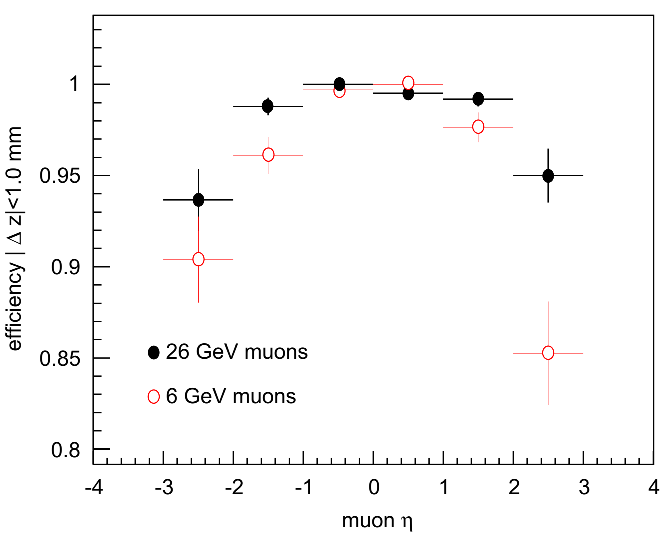

Fig. 3. The ZFinder efficiency for muons of 6 and 26 GeV.

pileup. Fig. 3 shows the efficiency as a function of the muon pseudorapidity. For muons with momentum greater than 6 GeV the efficiency for reconstructing the vertex position within 1 mm of the true vertex position is greater than around 90–95% for $|\eta(\text{muon})|<2$. The fall of the efficiency with the pseudorapidity is largely a resolution effect from the requirement that the reconstructed vertex must be within 1 mm of the true vertex since the resolution worsens at larger pseudorapidity.

Once the z position of the interaction vertex has been identified, the HitFilter algorithm calculates the pseudorapidity of each spacepoint with respect to this z. Hits from a given high-momentum track tend to be at approximately the same pseudorapidity with respect to their vertex. All the hits in a given RoI are binned in a two-dimensional histogram in $\eta-\phi$ space so that hits consistent with a given track from that z position would cluster in neighbouring bins of η and ϕ, whereas hits from a track originating at different z would be distributed between many bins at different η.

The book-keeping monitors the number of different layers that have hits, so that the hits in a cluster of adjacent bins are only considered to belong to a track candidate if they are observed in at least four out of the expected seven layers. This helps to remove fake candidates. If the bin size in η–ϕ is small the occupancy of each bin will be low. As with the ZFinder algorithm the execution time scales linearly with the number of spacepoints. The efficacy of the HitFilter is predicated on the correct z position being used for the pseudorapidity calculation. If the correct position is identified the efficiency of the HitFilter approaches 100% for $|\eta|<1.5$.

After the HitFilter has identified groups of hits which may be consistent with tracks, the GroupCleaner is used to help separate hits in those groups which may contain multiple tracks or contain additional hits from unwanted tracks. This is done by taking all possible triplets of hits within the group for a given cluster of hits and using the property that all triplets of hits from the same track should share the same simple track parameters, ϕ_0 and $1/p_T$, in the transverse plane. The extracted values are then binned in a two-dimensional histogram in $\phi_0 - 1/p_T$ space, with combinations containing hits in more than four layers considered as initial sets of hits on tracks for input to the Track Fitter.

2.1. The track fit and vertex algorithm

For the Track Fitter, two alternative track fitting methods are available, both using a distributed Kalman Filter to take account of energy loss and multiple scattering.

The first—the perigee parameter fit [4]—uses a uniform magnetic field and a helical track approximation, estimating the track parameters at the point of closest approach transverse to the beamline. Surface-to-surface extrapolation is not required. The second—the full LVL2 track fit—uses the non-uniform ATLAS magnetic field map and extrapolates the tracks between surfaces using a parabolic track approximation. The second fit gives better p_T and ϕ_0 resolution for tracks with large polar angle. However the execution time on a 2.4 GHz Xeon CPU is 0.06 ms per track for the perigee fit and 0.3 ms for the full LVL2 fit for 20 GeV muons. The fractional transverse momentum resolution, $\sigma(p_T)/p_T$, for the LVL2 track fit for 20 GeV muons is 0.036 compared to 0.032 for the full offline track fit and the online impact parameter resolution is 21.4 μm compared to 18.8 μm for the offline fit.

Tracks found in the Pixel detector and the SCT can then be used as seed tracks for extrapolation into the Transition Radiation Tracker (TRT) where TRT hits are added and the track fit updated using a Probabilistic Data Association Filter [5].

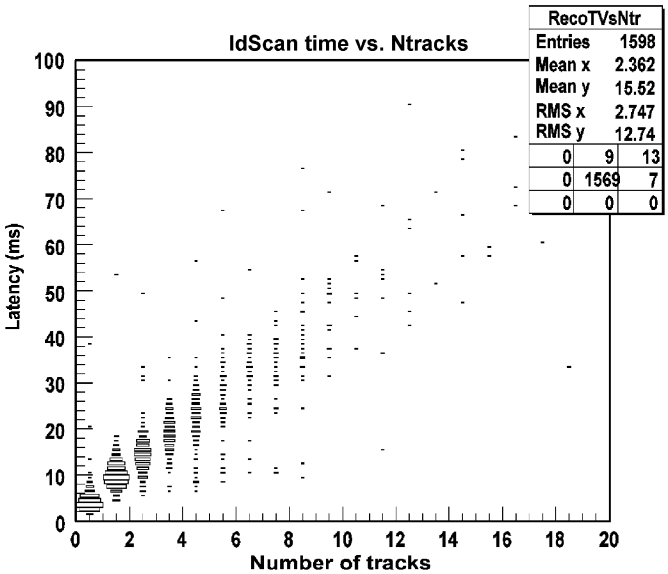

Fig. 4. The execution time, measured on a 2.4 GHz Xeon PC, of IDScan including the track fit for *b*-jet RoIs versus the number of tracks in the RoI.

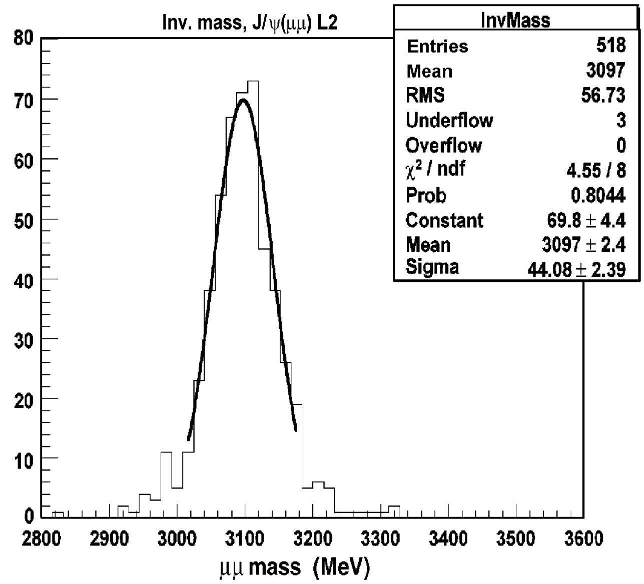

Fig. 5. The mass residual for reconstructed J/ψ decays.

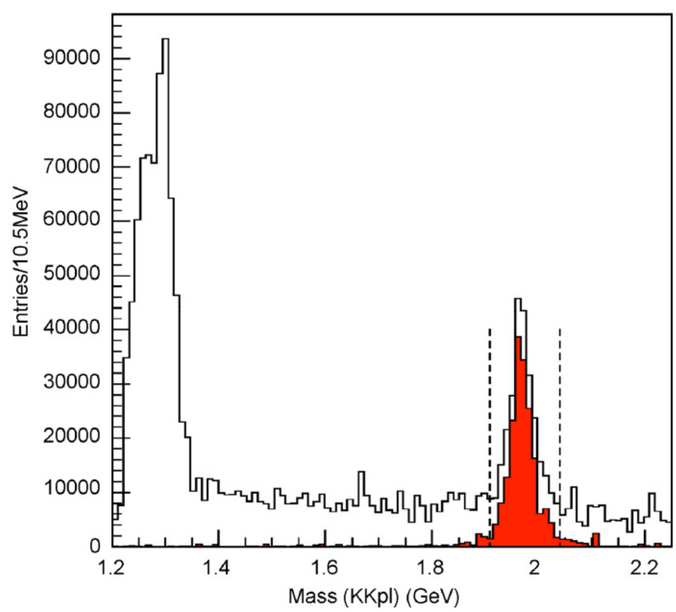

Fig. 6. The D meson mass from completely reconstricted $D \rightarrow \phi\pi \rightarrow KK\pi$ decays from *b*-physics events.

Fig. 4 shows the execution time for track reconstruction in *b*-jet RoIs as a function of the track multiplicity measured on a 2.4 GHz Xeon PC and illustrates that the dependence of the execution time on track multiplicity is linear to a good approximation with a mean of around 16 ms. For *b*-physics RoIs the efficiency for track finding rises sharply as a function of the track transverse momentum and is approximately 80% for tracks around 1 GeV rising to better than 90% for tracks with $p_T > 2$ GeV. The efficiency is approximately flat for $|\eta| < 2$ and falls off rapidly towards $|\eta| \sim 3$.

For vertex finding at LVL2 a fast iterative algorithm based on the Kalman Filter is used. A geometrical fit without mass constraints—equivalent to Billoir's "full" vertex fit [6]—is used after a linear transformation of the track parameters which reduces the track covariance matrix to a block diagonal form. This reduces the size of the vectors and arrays needing to be transformed during the iterations and leads to a significant improvement in the execution time. The fit results in the vertex position, track parameters and the full track covariance matrix for each track at the vertex.

Fig. 5 shows the invariant mass distribution for fully reconstructed $J/\psi \rightarrow \mu\mu$ decays using the LVL2 vertex algorithm. For perfectly aligned Monte Carlo data when including the TRT extrapolation the reconstructed width is 44.1 MeV compared to 43.4 MeV for the offline algorithm. The execution time for the online LVL2 algorithm is 0.06 ms whereas for the offline algorithm it is 0.16 ms. Without including the TRT extrapolation, the online width increases to 63 MeV.

To illustrate the applicability of the tracking and vertexing in practice, Fig. 6 shows the mass spectrum for D mesons reconstructed from a simulated *b*-physics sample with instantaneous luminosity of $1 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The D mesons are reconstructed from LVL2 tracks for the decay $D \rightarrow \phi\pi$ where the ϕ mesons are in turn reconstructed from secondary vertex tracks for the decay $\phi \rightarrow KK$ by assuming that each of the tracks from the secondary vertex is a K meson and taking only those vertices where the reconstructed invariant mass is within a window around the true ϕ mass. The filled distribution shows the mass combinations from the true D mesons suggesting that at this luminosity, the combinatorial background is manageable.

3. Surface cosmic data

Besides these, and many other, studies using simulated data, the performance has also been studied using data from the combined test beam of 2004 and more recently using data from the surface cosmic test from 2006. Until such time as actual luminosity data is available, cosmic data represents the best real data for the evaluation of the tracking algorithms. For the cosmic test, several sectors of the completed barrel SCT and TRT were cabled and instrumented in conjunction with a scintillator cosmic trigger.

For these data, no magnetic field was present so no p_T measurement was possible. In addition the cosmic muons do not originate from the beam line and so can have a large impact parameter with respect to the nominal beamline position in the transverse plane. Since the LVL2 tracking is optimised for tracks originating from the beamline this can have a significant impact on the track finding efficiency. In order to study the performance of IDScan running in a mode as close as possible to its operation during physics data taking, rather than retune the windows and cuts to allow for this large impact parameter, a naïve shift of the coordinate frame is used to shift the space points so that the cosmic tracks appear to originate from the beamline.

Fig. 7 shows the IDScan track finding efficiency as a function of the number of spacepoints. For large numbers of spacepoints, suggestive of more than one cosmic track passing through the detector, the efficiency falls and can be attributed to a failure of the naïve shift rather than to the IDScan algorithm.

A complimentary tracking approach, SiTrack [7], which uses lookup tables to define which hit combinations to consider sees a similar efficiency for low hit occupancy and no fall at higher occupancies, although in this case SiTrack uses a special set of lookup tables inappropriate for normal physics running to account for the potentially large impact parameter.

Since no magnetic field is present, the momentum of the tracks is unknown and the sample may contain a large fraction of very low momentum tracks where multiple scattering will be large which would adversely affect the track reconstruction and efficiency.

A cosmic test using the full detector in place in the ATLAS hall is envisioned for the summer of 2007 and the LVL2 tracking algorithms will be tested in the full ATLAS DAQ and trigger chain with the ATLAS magnetic field switched on.

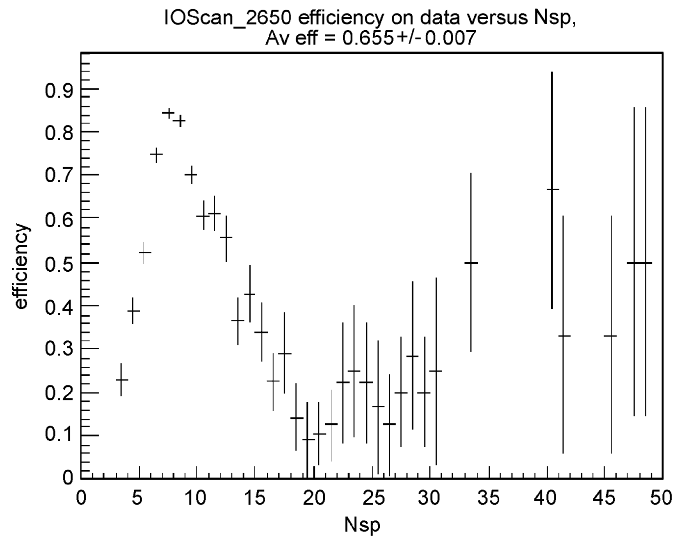

Fig. 7. The efficiency as a function of the number of spacepoints.

4. Summary and outlook

A set of algorithms for the fast reconstruction of high momentum tracks in the demanding high occupancy environment of the silicon tracking detectors in the ATLAS level 2 trigger has been presented.

The performance in terms of the processing time for *b*-jet events—the most time consuming events at LVL2—is seen to be within the time budget and the reconstruction efficiency is greater than 95% for tracks above 2 GeV produced centrally in the detector. Further development is underway and the algorithms presented have the desired flexibility to cope well with conditions when ATLAS data taking commences in 2007 and contribute to the successful completion of the ATLAS physics program.

References

[1] The ATLAS HLT-DAQ-DCS Technical Design Report (2003), ⟨http://atlas-proj-hltdaqdcs-tdr.web.cern.ch⟩.
[2] The ATLAS HLT-DAQ-DCS Technical Design Report (TDR) (2003) p. 228 and references therein.
[3] H. Drevermann, N. Konstantinidis, ATLAS Note ATL-DAQ-2002-014 (2002).
[4] D. Emeliyanov, ATLAS internal communication ATL-COM-DAQ-2004-012 (2004).
[5] D. Emeliyanov, Nucl. Instr. and Meth. A 566 (2006) 50.
[6] P. Billoir, S. Qian, Nucl. Instr. and Meth. A 311 (1992) 139; P. Billoir, S. Qian, Nucl. Instr. and Meth. A 350 (1994) 624.
[7] The ATLAS HLT-DAQ-DCS Technical Design Report (2003) p. 229 and references therein.