

The FastTrack Real Time Processor and Its Impact on Muon Isolation, Tau and b-Jet Online Selections at ATLAS

A. Andreani, A. Andreatza, A. Annovi, M. Beretta, V. Bevacqua, M. Bogdan, E. Bossini, A. Boveia, F. Canelli, Y. Cheng, M. Citterio, F. Crescioli, M. Dell'Orso, G. Drake, M. Dunford, J. F. Genat, P. Giannetti, F. Giorgi, J. Hoff, A. Kapliy, M. Kasten, Y. K. Kim, N. Kimura, A. Lanza, V. Liberali, T. Liu, A. McCarn, C. Melachrinos, C. Meroni, A. Negri, M. Neubauer, M. Piendibene, J. Proudfoot, G. Punzi, M. Riva, F. Sabatini, I. Sacco, L. Sartori, M. Shochet, A. Stabile, F. Tang, A. Todri, R. Tripiccione, J. Tuggle, V. Vercesi, M. Villa, R.A. Vitullo, G. Volpi, J. Wu, K. Yorita and J. Zhang

Abstract— As the LHC luminosity is ramped up to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and beyond, the high rates, multiplicities, and energies of particles seen by the detectors will pose a unique challenge. Only a tiny fraction of the produced collisions can be stored on tape and immense real-time data reduction is needed. An effective trigger system must maintain high trigger efficiencies for the physics we are most interested in, and at the same time suppress the enormous QCD backgrounds. This requires massive computing power to minimize the online execution time of complex algorithms. A multi-level trigger is an effective solution for an otherwise impossible problem. The Fast Tracker (FTK) is a proposed upgrade to the current ATLAS trigger system that

will operate at full Level-1 output rates and provide high quality tracks reconstructed over the entire detector by the start of processing in Level-2. FTK solves the combinatorial challenge inherent to tracking by exploiting massive parallelism of associative memories that can compare inner detector hits to millions of pre-calculated patterns simultaneously. The tracking problem within matched patterns is further simplified by using pre-computed linearized fitting constants and leveraging fast DSPs in modern commercial FPGAs. Overall, FTK is able to compute the helix parameters for all tracks in an event and apply quality cuts in less than 100 μs . The system design is defined and studied with respect to high transverse momentum (high- P_T) Level-2 objects: b-jets, tau-jets, and isolated leptons. We test FTK algorithms using ATLAS full simulation with WH events up to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity and comparing FTK results with the offline tracking capability. We present the architecture and the reconstruction performances for the mentioned high- P_T Level-2 objects.

Manuscript received April 3, 2010. This work was supported in part by the U.S. Department of Energy, the U.S. National Science Foundation, Istituto Nazionale di Fisica Nucleare and the European Commission with the FP6-2005-Mobility-6 Programme (L.Sartori is a Marie Curie fellow).

G. Drake, J. Proudfoot, J. Zhang are with Argonne National Laboratory, Argonne, Illinois, USA (e-mail: jinlong@mail.cern.ch).

M. Bogdan, A. Boveia, Y. Cheng, M. Dunford, J. F. Genat, A. Kapliy, C. Melachrinos, M. Shochet, F. Tang, J. Tuggle are with University of Chicago, Chicago, Illinois, USA (e-mail: shochet@hep.uchicago.edu)

J. Hoff, T. Liu, A. Todri, J. Wu are with Fermi National Accelerator Laboratory, Batavia, Illinois, USA (e-mail: thliu@fnal.gov).

F. Canelli and Y. K. Kim are with Fermi National Accelerator Laboratory and University of Chicago, Chicago, Illinois, USA.

L. Sartori and R. Tripiccione are with University of Ferrara, Ferrara, Italy.

A. Annovi, M. Beretta and E. Bossini are with INFN LNF, Frascati, Italy.

M. Citterio, C. Meroni, F. Sabatini are with INFN, Milan, Italy, (e-mail: mauro.citterio@mi.infn.it, chiara.meroni@mi.infn.it)

A. Andreatza, A. Andreani, V. Liberali, M. Riva and A. Stabile are with INFN and Università degli Studi di Milano, Milan, Italy (e-mail: valentino.liberali@unimi.it, alberto.stabile@unimi.it).

A. Lanza and V. Vercesi are with INFN, Pavia, Italy (telephone: +39-0382 987430, e-mail agostino.lanza@pv.infn.it, valerio.vercesi@pv.infn.it)

A. Negri is with INFN and University of Pavia, Pavia, Italy

P. Giannetti and R.A. Vitullo are with INFN, Pisa, Italy (telephone: +39-050-2214000, e-mail: paola.giannetti@pi.infn.it)

V. Bevacqua, F. Crescioli, M. Dell'Orso, M. Piendibene, G. Punzi, G. Volpi are with the INFN and University of Pisa, Pisa, Italy (e-mail: mauro.dellorso@pi.infn.it).

I.Sacco is with Scuola Superiore Sant'Anna, Pisa, Italy.

N. Kimura and K. Yorita are with University of Waseda, Tokyo, Japan (e-mail: kohei.yorita@waseda.jp).

M. Kasten, A. McCarn, M. Neubauer are with University of Illinois at Urbana-Champaign, Urbana, Illinois, USA (e-mail: msn@illinois.edu)

I. INTRODUCTION

THE most interesting processes at hadron colliders are very rare and hidden in an extremely high level of background. Implementing the most powerful selections in real time is therefore essential to fully exploit the physics potential of experiments where only a very limited fraction of the produced data can be stored on tape. Enormous real-time data reduction must be realized. A multi-level trigger [1] (LEVEL1, LEVEL2 and LEVEL3) is an effective solution for an otherwise impossible problem.

Real time track reconstruction can be an important element in triggering at CERN's Large Hadron Collider (LHC), and even more so after the Super LHC (SLHC) luminosity upgrade. There are numerous examples of the importance of tracking in the trigger. The source of electroweak symmetry breaking couples in proportion to mass. Thus heavy fermions are likely in the final state, in particular b quarks and τ leptons. High trigger efficiency for these processes requires sensitivity to the generic hadronic decays of the heavy fermions. The challenge comes from the enormous background from QCD

produced light quark and gluon jets, which can be suppressed using tracking. Tracks coming from a secondary vertex or not pointing to the beam line identify b quark jets, while τ jets can be separated from background using the number of tracks within a narrow “signal cone” and the number in a larger “isolation region”.

Electron and muon triggers can also be improved at high luminosity using track information. Traditionally background is suppressed by applying an isolation requirement using the calorimeters. At SLHC luminosity the energy added by the 75 additional collisions results in either decreased lepton efficiency or increased background contamination. The effect can be greatly ameliorated with a track-based isolation only using tracks pointing to the lepton candidate at the beam line.

The Fast TracKer (FTK) [2] will enable early rejection of background events and thus leave more execution time for sophisticated algorithms by moving track reconstruction into a hardware system with massively parallel processing that produces global track reconstruction with near offline resolution. FTK is based on the very successful CDF Silicon Vertex Trigger (SVT) [3], [6] .

The system is described in section II. For the first time FTK performance is estimated with the expected pile-up at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity consisting on average of 75 minimum bias interaction.. Resolution and efficiency for individual tracks and identification capabilities for physics objects, together with timing performance, are presented in section III. We present results obtained with a full simulation of the associate production of the W and the Higgs bosons. This channel contains all the signatures we are interested in: b-jets, taus and muons.

II. SYSTEM DESCRIPTION

FTK is an electronics system that rapidly finds and fits tracks in the ATLAS inner detector silicon layers for every event that passes the LEVEL1 trigger. It uses all 11 silicon layers over the full rapidity range covered by the barrel and the disks. It receives the pixel and silicon strip data (SCT layers) at full speed as it moves from the front end to the event buffers following a LEVEL1 trigger. After processing data FTK fills equivalent buffer memories with the helix parameters of all tracks with P_T above a minimum value, typically 1 GeV/c. The LEVEL2 processors can request the track information in a Region of Interest or in the entire detector.

FTK is a scalable processor. The system presented below will work well up to a luminosity of $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ but it can be expanded to operate at higher luminosities and staged for lower luminosities. The algorithm consists of two sequential steps. In step 1, pattern recognition is carried out by a dedicated device called the Associative Memory (AM) [4] which finds track candidates in coarse resolution roads. When a road has silicon hits on all or all-except-one layers, step 2 is carried out in which the full resolution hits within the road are fit to determine the track helix parameters and a goodness of fit. Tracks that pass a χ^2 cut are kept.

The first step rapidly carries out what is usually the most CPU intensive aspect of tracking by massive parallelism –

processing hundreds of millions of roads nearly simultaneously as the silicon data pass through FTK. The road width must be optimized. If it is too narrow, the needed size of the AM and hence the cost is too large. If roads are too wide, the load on the track fitters can become excessive due to the large number of uncorrelated hits within a road. This increases both the number of roads the track fitter must process and the number of fits within the road due to the hit combinatorics.

The FTK input bandwidth sets an upper limit on the product of the LEVEL1 trigger rate and the number of silicon hits per event. We minimize this limitation by running the silicon hits on multiple 100 MHz buses within FTK. Nevertheless, in order to sustain a 100 kHz LEVEL1 trigger rate, it is necessary to organize FTK as a set of independent engines, each working on a different region of the silicon tracker. The first step is to divide the detector into azimuthal regions. This segmentation generates some inefficiency at region boundaries removed by allowing a small overlap at the boundaries. At luminosities up to $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, we plan to use 8 ϕ “regions” (45° wide) with an overlap of 10° to guarantee high efficiency for tracks with P_T of 1 GeV/c and above. Each region will have its own “core processor” contained in a 9U VME crate (“core crate”), for a total of 8 engines working independently. We have recently modified the architecture of the FTK [2]. Previously, data flowed serially through the boards within each FTK crate. All of the silicon hits passed through all of the pattern recognition boards, and track candidates passed through a single track fitter board. That scheme worked well at low luminosities simulated at that time. With the simulation of large pile-up we find the data flow problem to be much more challenging.

FTK is now much more highly parallel, with regions segmented into η - ϕ towers, each with its own AM and track fitters. A tower receives substantially less silicon hits, and the track fitters have substantially less candidates to process. Each region and its processor are now subdivided into sub-regions, with enough overlap to maintain high efficiency. The η range is divided into four intervals, and the region’s ϕ range is divided again by two (22.5° plus 10° overlap).

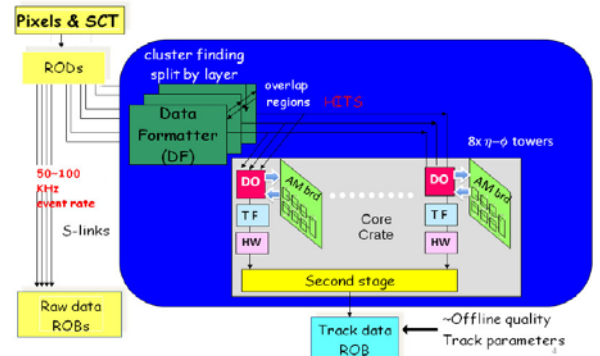


Fig. 1. Functional sketch of an FTK core crate plus its data connections.

The overlap in η takes into account the size of the beam’s luminous region in z. With this detector segmentation, we can distribute the data, from the Pixel and strip (SCT) detector, on the 8 parallel buses at the full 100 kHz LEVEL1 trigger rate with the detector occupancy expected for $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

The pixel and strip data are transmitted from the front-end on S-LINK fibers and received by the Data Formatters (DF) which perform cluster finding (see Fig.1). The barrel layers and the forward disks are grouped into logical layers so that there are 11 layers over the full rapidity range (see Fig.2). The cluster centroids in each logical layer are sent to the Data Organizers (DO). The DFs are not partitioned into regions. They organize the detector data into the FTK η - ϕ tower structure for output to the core crates taking the needed overlap into account.

The DO boards are smart databases where full resolution hits are stored in a format that allows rapid access based on the pattern recognition road ID and then retrieved when the AM finds roads with the requisite number of hits. In addition to storing hits at full resolution, the DO also converts them to a coarser resolution (Super-Strips, SS) appropriate for pattern recognition in the AM.




Fig. 2. The assignment of barrel layers and forward disks to 11 logical silicon layers.

The AM boards contain a very large number of preloaded patterns, or roads, corresponding to the possible combinations for real tracks of a superstrip in each silicon layer. These are determined in advance from a full ATLAS simulation of single tracks using detector alignment extracted from real data[5]. The AM is a massively parallel system in that all roads see each silicon hit nearly simultaneously. When a road has found the requisite number of hit layers, the AM sends that road back to the DOs. They immediately fetch the associated full resolution hits and send them and the road number to the Track Fitter (TF). Because each road is quite narrow, the TF can provide high resolution helix parameters using the values for the center of the road and applying corrections that are linear in the actual hit position in each layer. Fitting a track is thus extremely fast since it consists of a series of integer multiply-and-accumulate steps. In a modern FPGA, approximately 10^9 tracks per second can be fit.

Both the pattern matching and track fitting functions produce duplicates or “ghosts”, duplicate roads in the former and duplicate tracks in the latter.

Duplicate roads occur because we use majority logic in the pattern matching stage. FTK can require N fired layers among the total of M layers (N/M) to declare a road successfully matched. We plan to allow one missed layer in order to keep track-finding efficiency high. The use of the $(M-1)/M$ matching criterion generates duplicate roads. For each real track, it is possible to find a single M/M road (a hit found on

each layer) and/or a large number of $(M-1)/M$ roads (a hit is missing on one layer). We find duplicate roads negligible compared to the large number of fake roads in high detector occupancy conditions.

Duplicate tracks on the other hand are a source of fake tracks. A pair of duplicate tracks shares most but not all of their hits. They occur when a real track has non-associated nearby hits, either noise or coming from another track. If the replacement of real track hits by other hits still produces a satisfactory χ^2 , then there will be duplicate tracks. The Hit Warrior function (HW) is applied after track fitting and reduces the duplicate track rate by keeping only the best χ^2 track among those that share at least a specified number of hits.

An FTK R&D program has been carried out for a number of years, and prototypes have been built. Details of their design and performance are given in [7] and [2] and in the references in those papers. Important R&D for FTK was also carried out in the context of the SVT detector upgrade of the CDF experiment. A published paper [8] describes the upgrade and its relevance to future applications.

The major change since the $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ simulation became available is to replace the serial data flow in each core crate by data flowing in parallel through η - ϕ towers. Only those silicon hits within the tower pass through it. A tower consists of a DO, AM, TF, and HW. Only the tracks exiting the HW are sent out of the tower hardware.

The single DO, TF, and HW in the old architecture are now duplicated in each tower, but that does not result in a cost increase. In fact, the rapid advancement in FPGA technology allows the DO, TF, and HW functions to each fit into a single chip. A tower, including the DO, AM, TF, and HW functions, fits into a single core crate slot, with a main board and an auxiliary card on the back of the crate. We call the logic in a single slot a “Processor Unit”. As a result, the number of crates in the system does not increase. The ability to place multiple functions in a single slot provides a significant advantage for the data flow between functions. These high rate transfers now occur on short PCB lines rather than across a complex high-speed custom backplane as needed with the old architecture [7]. The cost also does not increase very much because the distributed functions can often be implemented on inexpensive FPGAs (Xilinx Spartan or Altera Cyclone families) or smaller sections of the very expensive top-of-the-line devices we needed in the old architecture (Virtex and Stratix respectively).

There is one potential disadvantage in the new architecture. Since the hardware is now partitioned among the η - ϕ towers, it could be more sensitive to large fluctuations in the numbers of hits, roads, and tracks in a tower. Previously the DO, TF, and HW were resources that were shared across the crate; now they are dedicated within each tower. However this is not a problem at high luminosity because the processing is dominated by pile-up hits rather than those from the hard scattering process. Fluctuations are reduced since the large pile-up produces a

quite uniform occupancy in the detector. The computing power within each tower is large enough to handle the fluctuations at high luminosity and is larger than needed at lower luminosity.

III. PERFORMANCE AT THE SLHC PHASE I LUMINOSITY

A. Single track efficiency and resolution

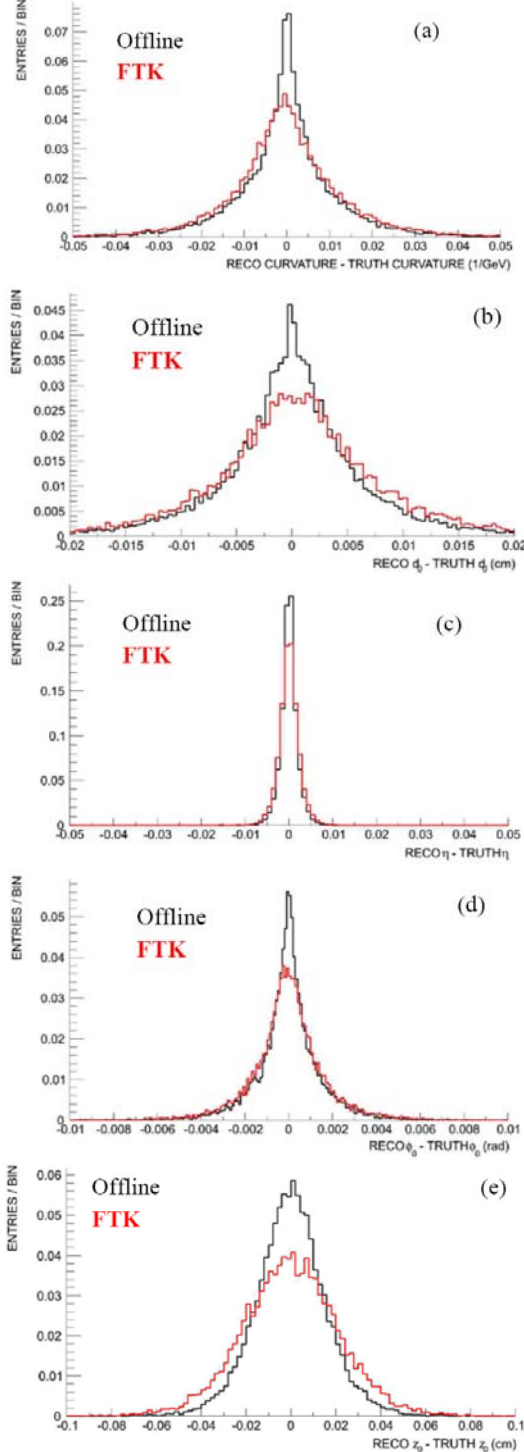


Fig. 3. Comparison of FTK and offline helix parameter resolutions in the barrel region: (a) curvature, (b) d_0 , (c) η , (d) ϕ_0 , and (e) z_0 .

If FTK tracking is to be useful in e , μ , b , and τ triggering, it must perform well on single tracks. In this section, we show

the resolution and efficiency for single-muon events without pile-up, but using the FTK settings that would be used at high luminosity. Comparison is made with offline tracking. Our requirements include a minimum of 9 silicon hits for the offline track reconstruction and a minimum of 10 hits for FTK.

FTK helix parameter resolutions (the transverse impact parameter, d_0 , the z coordinate, z_0 , the azimuth of the momentum direction, ϕ_0 , the pseudorapidity η , and the curvature, defined as the inverse of the particle P_T multiplied by the charge) are compared to the offline ones in Fig. 3 for the barrel region. Performance is only slightly degraded with respect to the offline reconstruction.

For efficiency, the matching of a reconstructed track with a generated track requires that at least 70% of the silicon hits used in the fit be hits that were generated by the original particle. The plots show results for tracks with $P_T > 1 \text{ GeV}/c$. The track efficiency is shown in Fig. 4 as a function of pseudorapidity and transverse momentum. We are currently implementing an improvement that nearly eliminates the efficiency dip near $|\eta|=1.2$, at the transition between the barrel and forward regions.

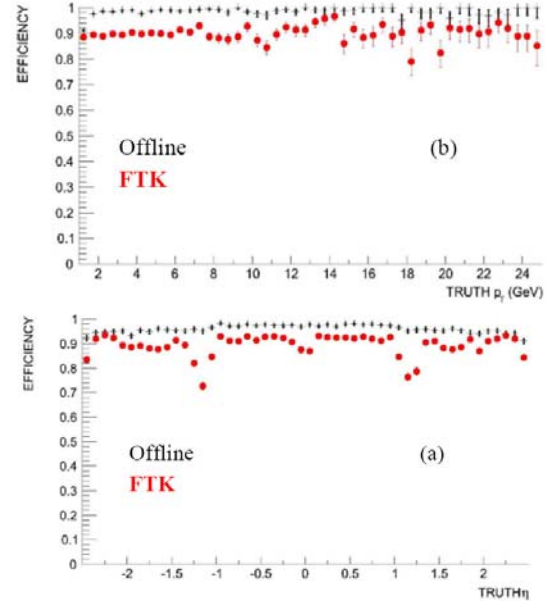


Fig. 4. Single muon tracking efficiency as a function of (a) pseudorapidity and (b) transverse momentum, comparing FTK with offline tracking.

B. b -tagging

New physics that couples to heavy fermions may be rich in final-state b quarks but, depending on the process, not necessarily leptons or very high P_T jets. Given the enormous QCD production of light quarks and gluons, it is important that the ATLAS trigger is able to efficiently select b -jets while providing a large rejection factor against other jets. Secondary vertex b taggers do well at separating b jets from light jets offline because they are able to obtain very high quality track information and to fully exploit it. Time constraints at LEVEL2 make both the tracking and tagging difficult, and one must make compromises to perform both for all Regions of Interest within the 20 ms LEVEL2 decision time. Since FTK tracks have

near-offline efficiency, helix parameter resolution, and fake rates, the immediate availability of high quality FTK tracks following a LEVEL1 trigger would allow the entire 20 ms to be used for more sophisticated tagging algorithms. This could increase the correlation between LEVEL2 and offline taggers and thus allow LEVEL2 to have an operating point close to that of the offline, increasing the light jet rejection factor at LEVEL2. This improvement in background rejection may prove crucial at very high luminosities where the tracking environment will be much more complex than at lower luminosity. In the following, we compare FTK and Offline tracks (reconstructed with the NewTracking algorithm [9]) in aspects key to b tagging.

To compare the performance of a tagger using FTK tracks with that of a tagger using offline tracks, we fully simulate samples of WH production ($M_H = 120 \text{ GeV}/c^2$ and H forced to decay into either a pair of b quarks or a pair of u quarks) to provide samples of signal and background jets, b -quark and u -quark jets respectively. To suppress the substantial rate of fake tracks in offline reconstruction at $3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, we apply offline cuts recommended for very high luminosity: the tracks must leave hits on a minimum number of 9 layers, there must be hits on all the pixel layers, and hits on a maximum of 2 SCT layers can be missing, while the distance between the track and the primary vertex must be less than one millimetre in the transverse plane and 15 cm in the z direction.

Jet b tagging starts by identifying tracks associated with a particular jet. Existing LEVEL1 and offline calorimeter clustering algorithms will have to be retuned for $3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ pile-up. To remove the effects of unreasonable clustering performance, which are separate from tracking performance, we begin with 0.4-cone truth jets, associating with the jet all tracks having momenta within a distance of 0.4 in the η - ϕ plane from the truth jet centroid. We then apply a Δz_0 cut relative to the highest P_T track to reduce impact of pile-up and recompute the jet direction as the P_T -weighted track average η and ϕ . A jet is labeled a b -quark jet if, within a 0.3 radius of the jet direction, there is a b quark from the Higgs decay. If a jet is not labeled as a b -quark jet, it is labeled as a light jet.

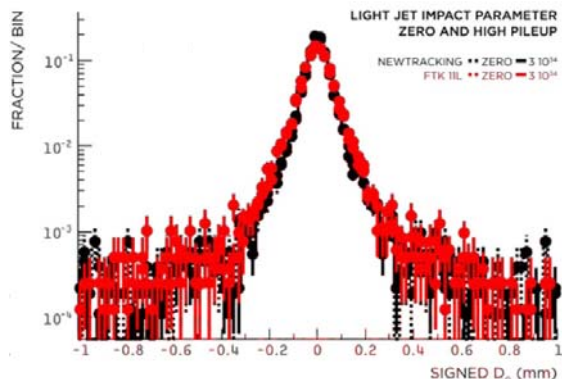


Fig. 5. The signed impact parameter distributions from offline (black) and FTK (red) for tracks in light-quark jets at 0 (dotted), and 3×10^{34} pile-up (solid lines). The distribution is stable as luminosity increases

Fig. 5 shows the distribution of transverse impact parameter significance for tracks in light quark jets with 0 (dotted lines) and 3×10^{34} pile-up (solid lines). The FTK d_0 resolution, while about 30% larger than that of offline, is still well suited for discriminating prompt tracks from those of b -jet decay (see Fig. 6), and both offline and FTK resolutions are stable up to $3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ pile-up. Fig. 7 plots the number of tracks in light jets with $d_0/\sigma_{d_0} > 3$ at high luminosity for the algorithm executed on both FTK and offline tracks, showing that the rate of spuriously high impact parameter tracks is roughly the same for both types of tracking.

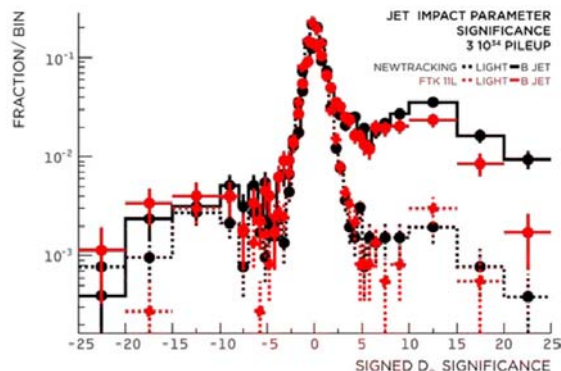


Fig. 6. The signed impact parameter distributions at 3×10^{34} for b jets (solid) and light jets (dotted) for both NewTracking (black) and FTK (red). The increase at significance between 5 and 15 is due to the variable size bins.

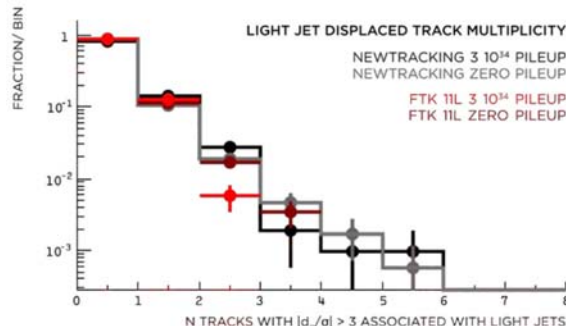


Fig. 7. The number of light-quark tracks with signed impact parameter significance greater than 3 for NewTracking (black) and FTK (red) at 0 pile-up and 3×10^{34} .

With similar single track efficiencies, resolutions, and fake rates, we expect b -tagging performance using FTK tracks to be similar to that using NewTracking. To test this, we compare FTK with NewTracking using a tagging algorithm equivalent to the baseline ATLAS LEVEL2 d_0 likelihood tagger. The algorithm builds likelihood functions from the signed d_0 significance distributions of tracks in a sample of b and light jets (labeled as described above). The ratio, L_b/L_{light} , of the product of the likelihood functions evaluated for each associated track is used as a discriminant between the two types of jets, and one can tune the performance of the algorithm by varying a cut on this ratio.

Fig. 8 shows the light-quark rejection power, the inverse of the probability to tag a light-quark jet as a b jet. It is plotted as a function of the b -tagging efficiency for a statistically-independent sample of WH events. The likelihood functions were built for each track type (either FTK or NewTracking)

and for each luminosity. FTK-based tagging performs well at all luminosities.

We are also studying a secondary vertexing tagger that uses the same underlying constrained vertex fitting code as the offline secondary vertexing.




Fig. 8. The light-quark rejection vs. b-jet efficiency for NewTracking (black) and 11-layer FTK simulation (red) at 3×10^{34} pile-up.

C. τ -identification

As the heaviest charged lepton, τ 's are often among the decay products in electroweak symmetry breaking scenarios. To maintain good sensitivity to these as well as other new physics processes that produce τ leptons, high trigger efficiency for τ 's down to the lowest possible τP_T is important. Of necessity, this means good efficiency and high background rejection for hadronic τ decay, i.e. τ jets.

In hadron collider experiments, τ -jet identification has usually been track-based. Signal and the large QCD jet background are separated by the presence of 1 or 3 tracks in a very narrow cone with little or no track activity in a surrounding isolation cone. When tracking is not available in early trigger levels, calorimeter-based selection requires a narrow isolated jet. This is the essence of the ATLAS LEVEL1 τ trigger. Here we propose rapid rejection of the QCD background by using FTK tracks at the beginning of the LEVEL2 τ selection process. We show that FTK tracking does nearly as well as offline tracking when applied to τ selection.

Tau leptons almost always decay to 1 or 3 charged particles (plus neutrals). The latter presents a greater challenge because of its smaller branching ratio, the size of the QCD background, and the requirement that track reconstruction be efficient for tracks very close to each other. But for some processes, important τ polarization information can be extracted from 3-prong decays.

The τ -tagging algorithm is based on the offline algorithm which utilizes conical η - ϕ regions to search for tracks around the LEVEL1 τ cluster and count tracks within signal and isolation regions.

Tracks with $P_T > 1.5$ GeV/c and within $\Delta R = 0.35$ of the LEVEL1 τ cluster are considered. Within this cone, the FTK track with the highest P_T is found. If it has $P_T > 6$ GeV/c, a signal cone of $\Delta R = 0.13$ is defined around it. Within the signal cone, there must be exactly 1 track in the case of single-prong τ 's, and 2 or 3 tracks in the case of triple-prong τ 's. An isolation cone of $\Delta R = 0.26$ is defined around the highest P_T track, and there must be no tracks of $P_T > 1.5$ GeV/c found between the signal and isolation cones. The cone sizes and

kinematic ranges are chosen to maximize S/\sqrt{B} , where S is the τ signal and B is the background from dijets.

Figures 9 and 10 show the τ reconstruction efficiency for 1-prong and 3-prong τ 's respectively. The data sample used is vector-boson-fusion Higgs production at 1×10^{34} $\text{cm}^{-2}\text{s}^{-1}$ luminosity with the Higgs decaying into two τ 's each of which decays hadronically. The efficiency denominator contains truth τ 's that are successfully matched to a LEVEL1 τ cluster. FTK and NewTracking give similar efficiencies for both 1-prong and 3-prong τ 's.




Fig. 9. Single-prong τ efficiency vs. rapidity and transverse momentum, the latter for $|\eta| < 0.8$. FTK tracking (black) is compared with offline (red) at 1×10^{34} $\text{cm}^{-2}\text{s}^{-1}$.




Fig. 10. Three-prong τ efficiency vs. rapidity and transverse momentum, the latter for $|\eta| < 0.8$. FTK tracking (black) is compared with offline (red) at 1×10^{34} $\text{cm}^{-2}\text{s}^{-1}$.

The 1-prong and 3-prong fake probabilities for FTK and the offline are at the level of few per thousands. The numerator contains jets passing the τ reconstruction, the denominator is all jets.

The τ identification at 3×10^{34} $\text{cm}^{-2}\text{s}^{-1}$ has the same dependence as a function of η and P_T but it is lower than at 1×10^{34} $\text{cm}^{-2}\text{s}^{-1}$ since tracking efficiency drops and pile-up tracks in the isolation cone causes τ 's to be lost. It is roughly 30% for both 1 and 3-prong decays.

D. μ triggering at high luminosity

Traditionally, single lepton triggers rely on calorimeter isolation to suppress backgrounds from real or fake leptons in hadronic jets. This works well at low luminosity. However at high luminosity, when there is substantial calorimeter energy due to the pile-up interactions, this strategy deteriorates. If the isolation threshold is kept low to maintain background rejection, the efficiency drops for leptons of interest. If the isolation threshold is raised, lepton efficiency can be maintained, but at the price of decreased background rejection.

An alternative strategy for high luminosity is to apply isolation based on reconstructed tracks. If all tracks in the event are used, pile-up remains a serious problem. However for tracking, unlike calorimeter deposition, the pile-up and

hard-scatter particles can be separated. Here we analyze a track-based isolation using only those FTK tracks that have a z_0 close to that of the lepton candidate.

The selection of isolated muons is critical to the searches for new physics such as SUSY cascades involving high P_T muons and a high mass Z' decaying into two muons, as well as the study of Standard Model processes such as $W \rightarrow \mu\nu$ or $Z \rightarrow \mu\mu$. In the ATLAS HLT, both calorimeter-based and tracking-based isolation are used to reject major backgrounds like bb events where one b -quark decays to a muon and light-quark jets where the jet produces tracks in the muon spectrometer, faking a muon. Since the rates for bb and light-quark jet production are significantly higher than those for isolated muon processes, it is critical to suppress the background rate in the trigger while maintaining high efficiency for isolated muons.

At high luminosity, calorimeter-based isolation loses its effectiveness in selecting real isolated muons. When the number of pile-up collisions increases, the amount of energy in the calorimeter, especially the EM calorimeter, also increases. With early tracking information available from FTK, calorimeter-based isolation in the trigger can be replaced with a tracking-based isolation that uses only tracks pointing to the z_0 of the muon track. This tracking-based isolation removes any requirements on the calorimeter cell energies, while maintaining high efficiency for isolated muons in environments with large pile-up.

In the current implementation of the HLT, the isolated muon trigger requires that the muon candidate pass isolation cuts based on the energy in EM and hadronic cells and the P_T of inner detector tracks. EM isolation is defined as the sum of all cells with energy above a 60 MeV threshold that lie in a ring of ΔR between 0.07 and 0.4 around the muon track. For hadronic isolation, the cell energy threshold is also 60 MeV and the isolation ring extends from R of 0.1 to 0.41

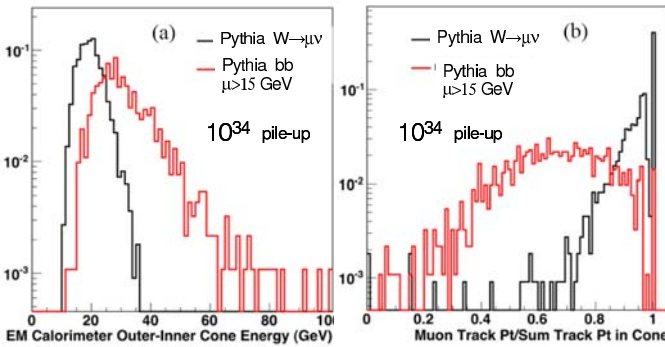


Fig. 11. (a) EM calorimeter isolation and (b) track isolation for signal and background with $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ pile-up.

The EM calorimeter isolation energy is shown in Fig. 11a for isolated muons from $W \rightarrow \mu\nu$ and non-isolated muons from bb jets at $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity. Tracking isolation is defined as the P_T of the muon track as measured by the inner detector divided by the P_T sum of all inner detector tracks within $\Delta R < 0.2$ of the muon track, including the muon track.

As seen in the Fig. 11, isolated muons tend to have smaller energy in the EM isolation ring. However the EM isolation distribution for muons from the events has significant overlap with the isolated muon distribution. For the same events, tracking isolation is shown in Fig. 11b. For isolated muons, this distribution is centered close to one. In contrast to EM calorimeter isolation, the track isolation distribution for bb events is more separated from the isolated muon peak.

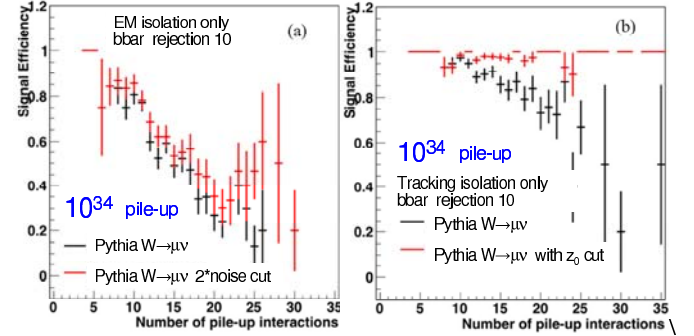


Fig. 12. Isolated muon efficiency using (a) EM calorimeter with two cell energy thresholds or (b) tracking isolation without (black) or with (red) a Δz_0 cut as a function of the number of pile-up interactions in the event. The isolation cut is selected to provide a bb rejection factor of 10.

To quantify the trigger efficiency for isolated muons at high luminosity, the calorimeter and tracking isolation cut values are set so that the trigger rejection factor for bb events is 10. Fig. 12a shows the isolated muon trigger efficiency as a function of the number of pile-up interactions using only EM calorimeter isolation. The trigger efficiency quickly deteriorates with increasing pile-up. Also shown in the figure is the trigger efficiency when the EM cell energy threshold is increased by a factor of two. The efficiency degradation with increasing pile-up is still clearly visible. Fig. 12b shows the isolated muon efficiency when instead tracking isolation is used. Here inner detector tracks, selected using offline reconstruction, must have $P_T > 1 \text{ GeV}/c$ and have at least one hit in either the pixels or SCT. In contrast to calorimeter isolation, tracking isolation is less sensitive to pile-up. The trigger efficiency can be further improved by using only inner detector tracks in the cone that have z_0 within 10 mm of the muon z_0 . As seen in the figure, the trigger efficiency for isolated muons using this track selection is insensitive to pile-up at this luminosity.

Using FTK tracks to calculate the tracking-based isolation yields similar trigger efficiencies for $W \rightarrow \mu\nu$ events. As seen in Fig. 13, the efficiency as a function of pile-up is constant for large numbers of pile-up events using FTK tracking, even for events with 100 pile-up interactions. The plots include both the $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ samples. The track-based isolation cut was tuned to give a bb rejection factor of 10. The overall efficiency for isolated muons is approximately 80% for both luminosities.

For high luminosities, when the number of pile-up interactions increases, the isolated-muon trigger efficiency deteriorates dramatically when calorimeter-based isolation is used. In contrast, tracking-based isolation with fast tracking

from FTK is insensitive to the amount of pile-up. The use of fast tracking information in the isolated-muon trigger results in high trigger efficiency for signal events while still maintaining good rejection of the bb background.

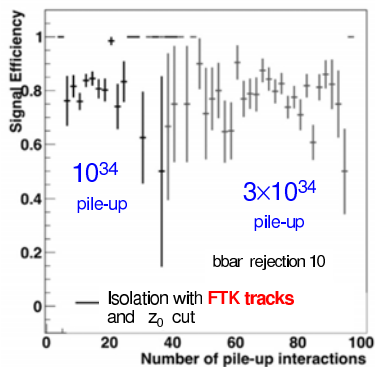


Fig. 13. The $W \rightarrow \mu\nu$ trigger efficiency with FTK track isolation as a function of the number of pile-up interactions in the event at $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (black) and $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (gray). The isolation cut is selected to provide a bb rejection factor of 10.

E. FTK timing

A simulation tool for calculating FTK execution time was developed for tuning the system architecture and parameters and to ensure that FTK can handle a 100 kHz LEVEL1 trigger rate at high luminosity. The system is divided into functional blocks: DF, DO write mode (receiving hits from DF and sending SSS to AM), AM, DO read mode (receiving matched roads from AM and sending roads and hits to TF), TF, and HW. The focus so far has been on the most time consuming steps, from DO write mode through TF. The DF should add little to the overall latency since each cluster found is immediately sent to the DO. Thus the DF and DO execution times almost completely overlap. Similarly, the HW has a relatively short latency for each track that enters before the track is either sent to the output or discarded. Both the DF and HW functions will be added in the near future.




Fig. 14. The FTK execution time for 100 WH events at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$

For each functional block, the time of the first and last words into and out of the block are calculated. Since each core crate operates independently, the FTK event execution time ends when the last word exits the busiest crate for that event. The execution time for a block depends on the number of input words, the processing time per word, and the number of output words. We estimate the processing time/word for each block type from the architecture, our experience with prototypes, and

the available chips on the market. The numbers of input/output words for each block come from FTKSim events. The results use WH events with $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ pileup.

The solid angle covered by a core crate is divided into η - ϕ towers. Each tower has its own unique set of AM roads. However the hits from the DFs often have to be sent to more than one tower due to curvature in the R- ϕ plane of $P_T \geq 1$ GeV/c tracks, the length in z of the LHC luminous region, and multiple scattering. The amount of hit duplication was determined from FTKSim and used by the timing simulation.

These studies assume very large input FIFO buffers. During the engineering design phase, we will do queuing studies to determine the needed depth of each input buffer.

To get high statistics on event execution times and see whether FTK can handle the 100 kHz LEVEL1 trigger rate, we analyzed a larger sample of events. If this rate was too large for our system, we would see the event execution time (from input hits available to event completion) steadily increase as FTK falls behind, working on a stack of previous events before getting to the current one. This does not happen. Some events take longer time than others to do global tracking, but after such an event the execution time quickly returns to the typical range. Fig. 14 has the timing histogram for FTK. FTK operates well for a 100 kHz LEVEL1 trigger rate at $3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

IV. CONCLUSIONS

Events rich in heavy quarks, tau-leptons, and isolated muons and electrons are particularly important for testing the limit of the SM and studying its possible extensions. Tracking devices play an essential role in this, in particular the silicon devices that are becoming the preponderant tracking technology. The electronics used to process the signals from the detectors must be state of the art to obtain the required large real-time data reduction. FTK is very successful in addressing the most difficult technological challenges for online tracking at SLHC Phase I luminosity and above.

REFERENCES

- [1] W. Smith, "Triggering at LHC Experiments", *Nucl. Instr. and Meth. A* vol. 478, pp. 62–67, 2002.
- [2] A. Annovi et al. "Hadron Collider Triggers with High-Quality Tracking at Very High Event Rates" *IEEE Trans. Nucl. Sci.*, vol.51, pp.391, 2004.
- [3] J. Adelman et al., "The Silicon Vertex Trigger upgrade at CDF", *Nucl. Instr. and Meth. A* vol. 572, pp. 361–364, 2007.
- [4] M. Dell'Orso, L. Ristori "VLSI Structures for Track Finding", *Nucl. Instr. and Meth. A* vol. 278, pp. 436, 1989.
- [5] E. Brubaker, et al., "Performance of the Proposed Fast Track Processor for Rare Decays at the ATLAS Experiment" *IEEE Trans. Nucl. Sci.* vol. 55, pp.145-150, 2008
- [6] A. Annovi et al., "A VLSI Processor for Fast Track Finding Based on Content Addressable Memories", *IEEE Trans. Nucl. Sci.* vol. 53, pp. 2428, 2006
- [7] A. Annovi et al., "The Fast Tracker Processor for Hadron Collider Triggers", vol. 48, pp. 575, 2001.
- [8] J. Adelman et al., "On-line tracking processors at hadron colliders: The SVT experience at CDF II and beyond", *Nucl. Instrum. Meth. A*, vol 581, pp. 473, 2007.
- [9] T. Cornelissen et al., "Concepts, Design and Implementation of the ATLAS New Tracking (NEWT), ATLAS Note, ATL-SOFT-PUB-2007-007.