PROCEEDINGS OF SCIENCE

Novel Methods and Expected Run 2 Performance for ATLAS Track Reconstruction in Dense Environments

Gabriel Facini[*†]
University of Chicago
E-mail: gabriel.facini@cern.ch

Novel algorithmic developments to the ATLAS track reconstruction software targeting the core of high p_{T} objects is discussed explaining improvements made for Run 2 resulting in:

- 5% increase of efficiency to reconstruct all three charged daughters of a hadronically decaying τ with p_{T} above 600 GeV
- 10% more pixel clusters found on tracks in the jet core,
- 17% efficiency gain is seen for charged particles created at a radius of 30 mm,
- 10% (14%) efficiency improvement in the core of high p_{T} light (b) jets,
- 7-13% increase in b-tagging efficiency for a given light-jet rejection.

24th International Workshop on Vertex Detector -VERTEX2015-
1-5 June 2015
Santa Fe, New Mexico, USA

[*]Speaker.
[†]On behalf of the ATLAS Collaboration.

PoS(VERTEX2015)041

1. Introduction

Detailed understanding and optimal track reconstruction performance of ATLAS in the core of high p_{T} objects is paramount for a number of techniques such as jet energy and mass calibration, jet flavor tagging, and hadronic tau identification as well as measurements of physics quantities like jet fragmentation functions and jet charge. With the increase of center of mass energy for Run 2 of the LHC, the potential to produce new heavy resonances is considerable. Decays of the new particles, and the associated background, can be dominated by high p_{T} objects where algorithm performance falters. Therefore to fully exploit the discovery potential of the full Run 2 dataset, detailed performance studies are needed. The dense environments of the cores of high p_{T} jets and τ-leptons are characterized by charged particle separations on the order of the resolution of the ATLAS Inner Detector. With the recent insertion of a new innermost layer in this tracking detector, which allows measurements closer to the interaction point, and the increase in the center of mass energy, these difficult environments will become even more challenging and relevant in Run 2. Novel algorithmic developments to the ATLAS track reconstruction software targeting these topologies is discussed

Track reconstruction in dense environments (TIDE) is aided by an artificial neural network (NN) based approach was introduced in 2011 to identify clusters created by multiple charged particles [7]. The usage of this information was reconsidered during a targeted overhaul of the ambiguity solving stage of the ATLAS track reconstruction chain.

Section 2 contains a brief description of the detector focusing on the Inner Detector. An overview of the track reconstruction algorithm chain is given in Section 3. Section 4 states the Monte Carlo (MC) generators and samples utilized. In Section 5, changes to the pixel cluster and track reconstruction strategy are presented. Delaying the usage of the pixel cluster NN to a later stage where its ability to resolve merged clusters is enhanced due to the available track information is a major component of the change. Section 6 presents distributions of tracks-in-jet parameters summarizing the changes resulting from the optimization. This note concludes with Section 7 showing the positive impact on the performance of flavor tagging algorithms for jets with p_{T} above 100 GeV.

2. The ATLAS Inner Detector

The ATLAS experiment, a multi-purpose particle detector at the LHC, is described in detail in Ref. [2].[1] The inner detector provides position measurements for charged particles in the range $|\eta| < 2.5$ by combining information from three sub-detectors. Starting from the interaction point and focusing on the barrel region, the high-granularity silicon pixel Detector segmented in R–ϕ and z covers the vertex region and typically provides four measurements per track, including the new innermost layer (layer 0), the insertable b-layer (IBL) [11] added for Run 2. The IBL has a

[1] ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the center of the detector and the z-axis along the beam pipe. The x-axis points from the IP to the center of the LHC ring, and the y-axis points upwards. Cylindrical coordinates (R,ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the z-axis. The pseudorapidty is defined in terms of the polar angle θ as $\eta = -\ln\tan(\theta/2)$. Angular distance is measured in units of $\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$.

mean radius of 33 mm and a typical pixel size of 50 μm (250 μm) in the transverse (longitudinal) direction with a thickness of 200 and 230 μm for the planar and 3D sensors, respectively. For the remaining three layers (1–3) of the pixel system, located at mean radii of 50.5, 88.5, and 122.5 mm respectively, a typical pixel has a size of 50 μm (400 μm) in the transverse (longitudinal) direction with a thickness of 250 μm. The charge collected on a pixel sensor is obtained by measuring the pulse height using the time–over–threshold (ToT) technique [12]. Outside the pixel volume, the silicon micro-strip detector (SCT) has four double-sided strip layers. On one side, the strips are parallel to the beam direction and at a stereo angle of 40 mrad on the other. The information on both sides is combined on each layer to provide an average of four three-dimensional measurements. The SCT sensors are connected to binary readout chips, which do not provide information on the collected charge. The silicon detectors are complemented by the Transition Radiation Tracker (TRT), which extends track reconstruction radially and provides R–ϕ information. for charged particles within $|\eta| = 2.0$. The following section provides an overview of how tracks are created starting with information from the silicon detectors.

3. The Track Reconstruction Algorithm

The following section describes the baseline track reconstruction algorithm used during Run 1. Charged particle reconstruction begins with the conversion of the raw data from the Pixel and SCT detectors into three-dimensional measurements referred to as space-points. In the Pixel Detector, each cluster equates to one space-point, while in the SCT, clusters from both sides of a strip layer must be combined to obtain a three-dimensional measurement.

A single isolated particle will typically create charge on multiple pixels of a given sensor with the total predominantly determined by the incident angle. A connected component analysis (CCA) [17] groups pixels into a cluster. When the spacial separation of charged particles traversing the module is only a few pixels, charge deposits overlap and the CCA algorithm reconstructs a single merged cluster. An artificial neural network (NN) is thereafter used to identify merged clusters [7] with those above a tunable limit split into two or three identical copies. All clusters are then assigned a position and error from additional NNs.

The primary track reconstruction algorithm utilizes iterative track-finding seeded on combinations of space-points from the silicon detectors. A staged pattern recognition approach, first introduced in the LEP experiment [19] era, is used: a loose track candidate search, which allows a number of combinatorial track candidates, is followed by a stringent ambiguity processor that compares and rates the individual tracks by assigning a relative track score to each track. This section outlines the seed finding and ambiguity solving stages of the primary track reconstruction chain. Further details, including a description of TRT track extensions, can be found in Ref. [3].

3.1 Iterative Combinatorial Track Finding

Seeds are formed from sets of three space-points. A first, crude estimate of the perigee parameters of a track seed is performed with respect to the centre of the interaction region using a perfect helical trajectory in a uniform magnetic field.

Four possible types of combinations can be made from the Pixel and SCT detector space-points. SCT-only seeds are considered first having the highest purity or the fraction of seeds that

result in good quality tracks. A number of criteria are placed on the seeds to maximize purity and the use of space-points in multiple seeds is carefully controlled. Purity is further improved by requiring one additional space-point to be compatible with the seed before a combinatorial Kalman filter [16] builds the track candidates. A track candidate is created for each compatible space-point extension.

The result is a very high efficiency for reconstructing primary particles[2] and removing tracks created from purely random collections of space-points. Obtaining this level of performance with reasonable CPU resources is partly due to the seed purity requirements mentioned above. From $\sim$13 space-point combinations created for an isolated charged particle traversing the entire Inner Detector, the time intensive combinatorial Kalman filter is, on average, called in its entirety 1.1 times. However, as all reasonable combinations of tracks have been made, there are a number of track candidates with incorrectly assigned space-points. This necessitates an ambiguity solving stage.

3.2 Ambiguity Solving of Track Candidates

The ATLAS ambiguity processor performs the final selection of reconstructed tracks. Track candidates are considered in descending order of a track score based on simple measures of the track quality such as the number of clusters and holes as well as the χ^2 of the track fit.

Shared measurements must be limited as they are also a strong indicator of poorly measured tracks. Unfortunately, shared measurements are also expected for merged clusters. This fact motivates the pixel NN [7] tool introduced in Run 1.

To count shared clusters, a track candidate is only compared to those tracks previously accepted by the ambiguity processor. A track can have two shared clusters and a cluster can be shared by two tracks. A cluster is removed from a track candidate if either rule is broken. The track candidate is then re-scored and returned to the ordered list of candidates.

For qualifying track candidates, a full resolution fit is performed and the track is added back into the list of candidates. If an already fitted track passes through the ambiguity processor without any interference, it will be added to the final track collection.

3.3 Truth-Based Reconstruction

An important tool for studying tracking performance in dense environments is truth-based reconstruction in simulated events [14]. For a given charged MC particle, the pattern recognition process is skipped and the hit collection is gathered using MC truth information. If the cluster content requirements are satisfied the particle parameters are obtained from a full resolution fit. The final set of tracks has the ideal cluster content and gives the maximum reconstruction efficiency given the ATLAS detector geometry and performance as well as the minimal hit requirements.

3.4 Truth-Based Track Quality

In simulation, tracks are classified using a truth score determined by the fraction of measurements originating, at least in part, from the same simulated particle. A properly reconstructed track

[2]For example muon reconstruction efficiency is in excess of 99% [9].

is required to have a truth score above 70%. Such a requirement is imposed for all reconstruction efficiencies presented in this note.

Fake tracks are those which have a truth score below 70%. Due to the careful pruning of seeds, the majority of reconstructed fake tracks are from the mis-allocation of clusters to a track and not purely random combinations of clusters.

4. Samples

In order to perform detailed studies, high statistics simulated samples in focused topologies were created with several Monte Carlo generators interfaced with a full ATLAS detector simulation [4] based on the GEANT4 program [13]. These samples, with a single particle decaying into a set of close-by charged particles, are ideal for algorithm development, as is shown in Section 5. For this purpose parent particles were generated with a flat transverse momentum (p_{T}) spectrum from 10 GeV to 1 TeV within $|\eta| < 1$. A two track sample ($\rho \to \pi^+\pi^-$) was used to study the double track efficiency. A sample consisting of a single τ-lepton decaying to three charged hadrons ($\tau^\pm \to \pi^+\pi^-\pi^\pm\nu_\tau$) was used to test the performance when more than two particles contribute to a cluster and the parent particle has a non-zero lifetime. Finally a single B^+ sample was used to study a high multiplicity of tracks from a long-lifetime parent particle, in particular charged particles created after the innermost pixel layer. The performance was then studied in a fully simulated physics process as described in Section 6. Events with a 3 TeV Z' decaying to a pair of top quarks were fully simulated by the PYTHIA 8 [18] generator using the AU2 tune [6] and the MSTW 2008 L0 [15] PDF sets. This sample includes the effect of multiple pp interactions in the same and neighboring bunch crossings (pile-up) by overlaying minimum bias events simulated with PYTHIA 8 on each generated event. The number of overlaid events is chosen so that the distribution of the average number of interactions per pp bunch crossing in the simulation matches that expected in the data at $\sqrt{s} = 13$ TeV. On average this is 41 interactions per bunch crossing.

5. Ambiguity Solver Optimization

A revision of the ambiguity solver was performed to optimize track reconstruction performance in dense environments (TIDE). The following section describes the changes made to the baseline reconstruction. One particular aspect of the track reconstruction chain that was improved is the treatment of the clusters formed from multiple particles. A summary of the changes is given in Section 5.3.

5.1 Merged cluster identification

A measure of collected charge per pixel via the ToT [12] technique and the physical location of the pixels provides the NN with enough information to identify merged clusters efficiently. However, the inherent randomness of charged particle interactions with thin silicon layers prevents the NN from ever performing perfectly. Also, the emission of δ-rays causes difficulties as they can lead to bigger clusters and larger charge deposits compared to expectations from a single particle. The position determination can also be biased by δ-rays. For collimated charged particles, merged

clusters that are not identified can lead to exaggerated shared cluster and hole rates and thereby impact the fake rate and reconstruction efficiency.

The Run 1 track reconstruction chain evaluated the pixel cluster NN and performed the subsequent splitting immediately after cluster reconstruction. However, the incident angle of the particle can be used to improve the NN performance. After cluster reconstruction, this is estimated via the physical position of the pixel module with respect to the beam-spot position. Figure 1 shows the rate of incorrectly split one-particle clusters versus the rate of non-split multi-particle clusters [7]. For the same rate at which one-particle clusters were split in Run 1, a 15% increase in efficiency for correctly identifying merged clusters can be obtained if a precise incident angle measurement is used.

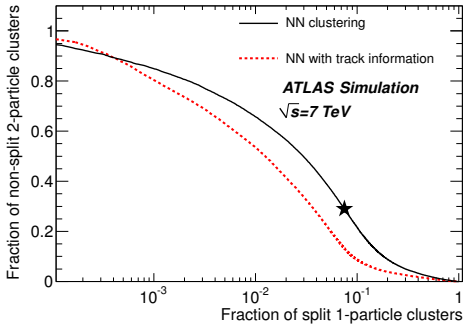

Figure 1: The fraction of split one-particle clusters versus the fraction of non-split two-particle clusters in simulation. The cut value of the NN varies along the solid and dashed lines: a tight cut corresponds to the lower right corner of the plot, and a loose cut corresponds to the upper left corner. The distribution is shown both for the NN using only the cluster and beam-spot information (solid line) and for the NN additionally including the track information (dashed line). The chosen working point for Run-I for the setup without track information is indicated with a star. The usage of this tool is discussed in Section 5.1 and full details are given in Ref. [7]. This figure can be found in Ref. [1].

Within the ambiguity solver, the track candidate properties are precise enough to produce the improved NN performance. This motivates delaying the NN evaluation until the ambiguity solver. Clusters are then no longer split or physically copied. With fewer clusters entering the seed finding stage, 10% fewer track candidates are produced, thus saving CPU time.

After this change, the NN is consulted only when a cluster is contained in multiple track candidates. Those passing the pixel NN cut are used by the track candidates without penalty and additional NNs are used to determine the position and error of each contribution. This leads to the notion of shareable clusters which replaces the obsolete idea of split clusters. Clusters which are not shareable can still be shared but under the penalty described above.

Even with the improved performance from the precise incident angle, some merged clusters on multiple track candidates will not be identified by the NN. This can be overcome by correlating information on consecutive layers of the Pixel Detector. In general, the separation between collimated charged particles increases as they travel outward through the Inner Detector. Therefore given a merged cluster-on-track on a given layer, the next layer inward should also contain a merged cluster. Furthermore, both clusters should be used by the same track candidates. To utilize this information, if the same track candidates compete for clusters on two consecutive layers, the

cluster on the inner layer will be marked shareable if the cluster on the outer layer is identified as merged by the pixel NN tool.

These changes motivate a re-optimization of the requirements on the NN outputs. With the goal of maximizing efficiency against fake rate, these were loosened to 0.35 (0.4) for the two (three) particle NNs. In the Section 5.2, the optimization of four additional parameters used to control the fake rate is discussed.

The average number of truth-identified merged clusters on truth-based tracks (Section 3.3), is compared to the average number of both split and shareable clusters in Figure 2 for the single ρ and single τ samples. The average separation of the two (three) charged particles from a ρ (τ) decay decreases with an increase in the parent particle momentum leading to more merged pixel clusters as shown in the points labeled ideal. Both the average number of merged and shareable clusters falls to zero at the lowest parent particle p_{T} shown, while the baseline reconstruction has too many split clusters due to the NN splitting single-particle clusters. The imperfect efficiency of the NN is apparent at high p_{T} but is a lesser problem for the TIDE optimized chain. The inefficiency is larger for the τ sample where clusters created by more than two particles are more prevalent. The trend of the number of shareable clusters on tracks reconstructed in the TIDE optimized chain follows the true number of merged clusters nicely implying it has more physical meaning.

Cluster assignment efficiencies are shown in Figure 3 for the first two layers of the pixel detector. They measure the fraction of clusters created by a particle that are then used on the reconstructed track of said particle. With the closest truth particle separated by 400 μm at the innermost layer, the cluster assignment efficiency at this layer increased from 81% (56%) to 98% (82%) per track for the ρ (B^{+}) sample. A significant improvement was also achieved at all other pixel layers. Inefficiencies evident in the B^{+} sample are a consequence of the particle's lifetime. By decaying closer to the innermost layer, several daughter particles of the B^{+} are likely to contribute to a single merged cluster.

Figure 4 shows the total effect of the changes described in this note on the average number of pixel clusters on tracks in the ρ and three-prong τ sample as a function of the parent particle p_{T}. With the improved identification of merged clusters, as shown in Figure 2, fewer clusters must be removed to keep track candidates below the shared cluster requirements applied during track reconstruction. This leads to more clusters on tracks. In the ρ sample of Figure 4, the TIDE configuration mimics the ideal content. In 3-prong τ decays, the ideal number of clusters decreases due to the τ's decaying after the innermost layer. The TIDE configuration is an improvement over the baseline, however more inefficiencies remain in events with clusters created by more than two charged particles. Changes introduced by the TIDE optimization add up to 0.3 (0.5) pixel clusters on tracks in the ρ (τ) sample.

5.2 Sharing Measurements: A Trade Off

The pixel NN approach minimizes the negative impact of shared pixel clusters on track reconstruction performance by helping to identify merged clusters. In the previous section, the NN was motivated as a tool used to identify merged clusters; it can also identify one-particle clusters. Clusters with a two-particle NN score less than 0.05 are explicitly forbidden from being shared. This is less a stringent requirement compared to Run 1.

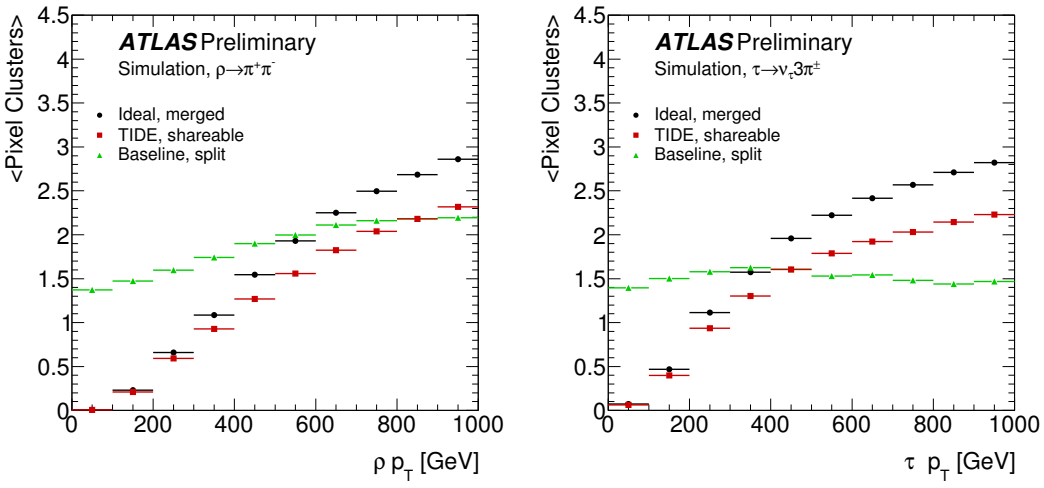

(a) single ρ sample (b) single 3-prong τ sample

Figure 2: A comparison of the average number of merged pixel clusters on truth-based reconstruction tracks (Section 3.3) and split (shareable) pixel clusters is shown as a function of the ρ and τ transverse momentum. The trend of the average number of shareable clusters on tracks reconstructed by the TIDE optimized chain follows the true average number of merged clusters as shown by the points labeled ideal, better than the baseline reconstruction chain. Section 5.3 lists the full set of differences between the baseline and TIDE reconstruction chain. This figure can be found in Ref. [1].

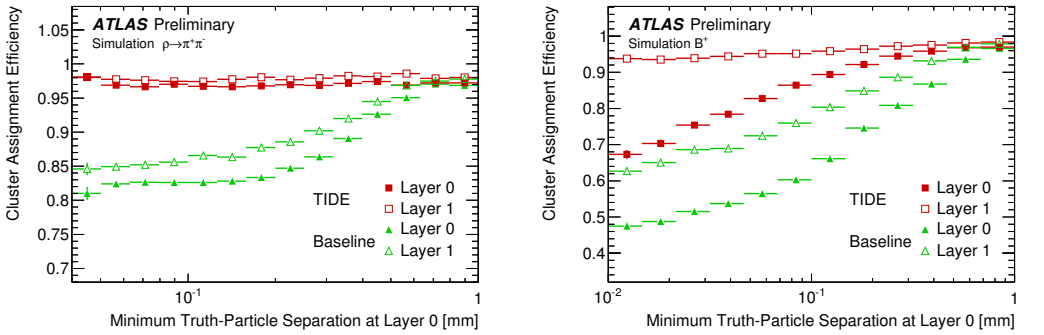

(a) single ρ sample (b) single B^+ sample

Figure 3: The efficiency by which reconstructed clusters are properly assigned to a track is shown for the two innermost pixel layers (layer 0 and layer 1) as a function of the minimum truth particle separation at layer 0. A vertical slice shows the expectation for the number of clusters on the first two layers for a given pair of tracks and their separation at layer 0. The TIDE optimized setup outperforms the baseline on both layers. Section 5.3 lists the full set of differences between the baseline and TIDE reconstruction chain. This figure can be found in Ref. [1].

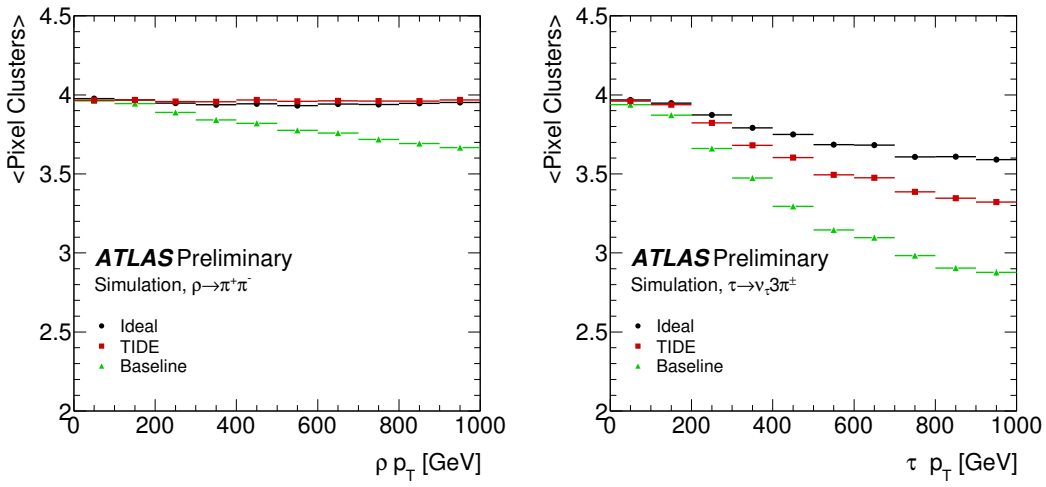

(a) single ρ sample

(b) single 3-prong τ sample

Figure 4: The average number of pixel clusters on track is shown in the ρ (a) and τ (b) sample for the ideal, baseline and TIDE reconstruction as a function of the parent particle p_{T}. With the improved identification of merged clusters, as shown in Figure 2, less clusters are removed to keep track candidates below the requirements on shared clusters applied during track reconstruction leading to more clusters on track. In the ρ sample, the TIDE configuration mimics the ideal content. In 3-prong τ decays, the ideal number of clusters decreases due to the τ's decaying after the innermost layer. The TIDE configuration is an improvement over the baseline, however more inefficiencies remain in events with clusters created by more than two charged particles. Changes introduced by the TIDE optimization add up to 0.3 (0.5) pixel clusters on track in the ρ (τ) sample. Section 5.3 lists the full set of differences between the baseline and TIDE reconstruction chain. This figure can be found in Ref. [1].

A multivariate strategy cannot be applied to the SCT clusters, as no information on charge deposition is available. Therefore, the requirements on shared clusters, which restrict the fake rate, can also limit the overall track reconstruction efficiency accessible in dense environments. Additional techniques to control the fake rate were introduced and tuned to help with this problem. Fake rates will increase if short tracks from secondary particles are extended through incorrectly sharing clusters on the inner layers. To control this, a number of silicon clusters must be on a track before any sharing is allowed. Because of the addition of the IBL, the minimum was raised from eight to nine clusters. Two requirements were introduced to mitigate fakes. Tracks with p_{T}less than 1 GeV cannot have shareable clusters and a track is required to have a minimum of four unique SCT clusters.

Figure 5 shows the reconstructable or physics efficiency for the ρ and τ samples as obtained from the truth-based approach described in Section 3.3. This is the maximum achievable efficiency given the cluster content requirements, including the limit on the number of shared clusters. The merged pixel clusters are assumed to have been identified and so only a cut on the number of shared SCT clusters is performed. Currently this limit is two shared clusters, as the efficiency improvement obtained from loosening this cut is not sufficient to justify the associated increase in the proportion of fake tracks.

The efficiency for reconstructing all primary tracks in events in which none are expected to

have more than two shared SCT clusters is shown in Figure 6. A clear improvement in the reconstruction efficiency can be seen in the optimized setup. The improvement increases with momentum, which reflects the improved handling of multi-particle clusters in the Pixel Detector. In events without inelastic interactions, such as hadronic interactions, photon conversions, etc., the track reconstruction efficiency in the single ρ sample is very close to 100%. The efficiency is lower in the more complex τ sample. This decrease in efficiency with momentum is caused by pattern recognition errors due to NN inefficiencies. This can reduce the truth score and cause a track to be classified as fake. The effect is amplified when looking at events where additional clusters in the detector are generated by secondary particles.

After the relaxation of the NN output requirement for clusters to be shared and other changes described in the previous section, more shared pixel clusters are expected. The effect of the changes is counterbalanced by a tightening of the way shared clusters are counted in order to keep the fake rate down at Run 1 levels. Previously, only the track in question was evaluated for the cluster cuts. Accepted tracks were not reconsidered even if the shared cluster count changed, as more tracks were accepted. For Run 2, a track will not be accepted if it fails the cluster pattern cuts or causes an accepted track to fail the same cuts. This increases the preference for tracks with higher score.

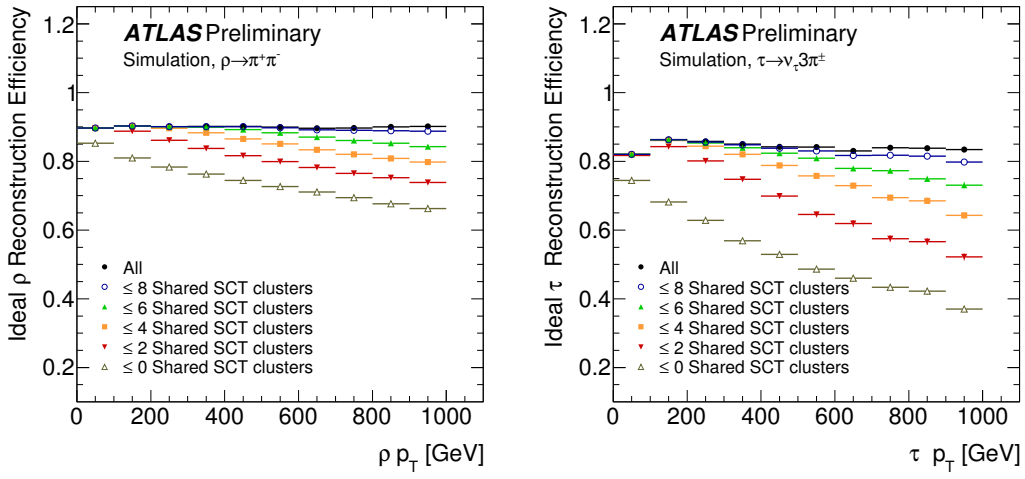

(a) single ρ sample

(b) single 3-prong τ sample

Figure 5: The truth-based reconstruction (Section 3.3) efficiency is shown for reconstructing all the decay products of the ρ and the three prong τ with various limits on the number of shared clusters allowed on a track candidate assuming all the merged pixel clusters were properly identified. The number of shared clusters is a limiting factor on the efficiency for charged particles with little spatial separation. This figure can be found in Ref. [1].

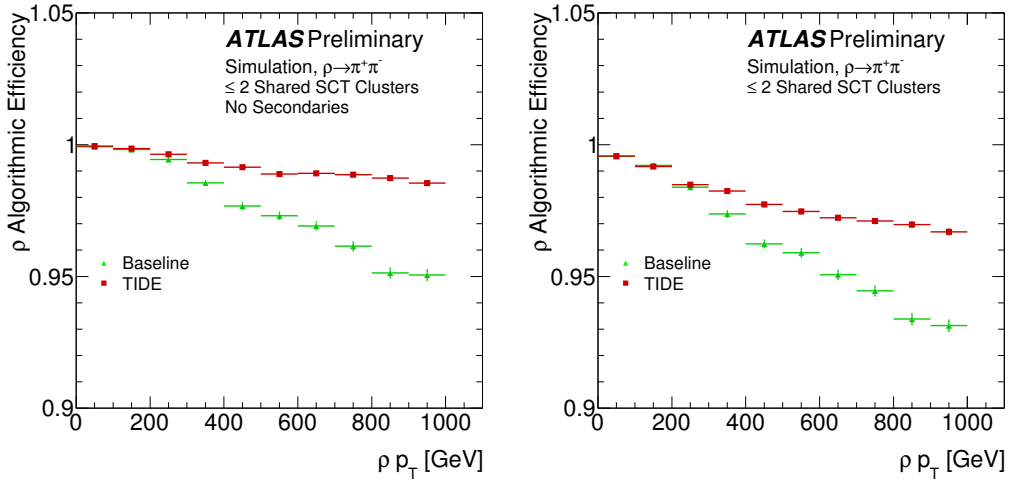

(a) single ρ sample

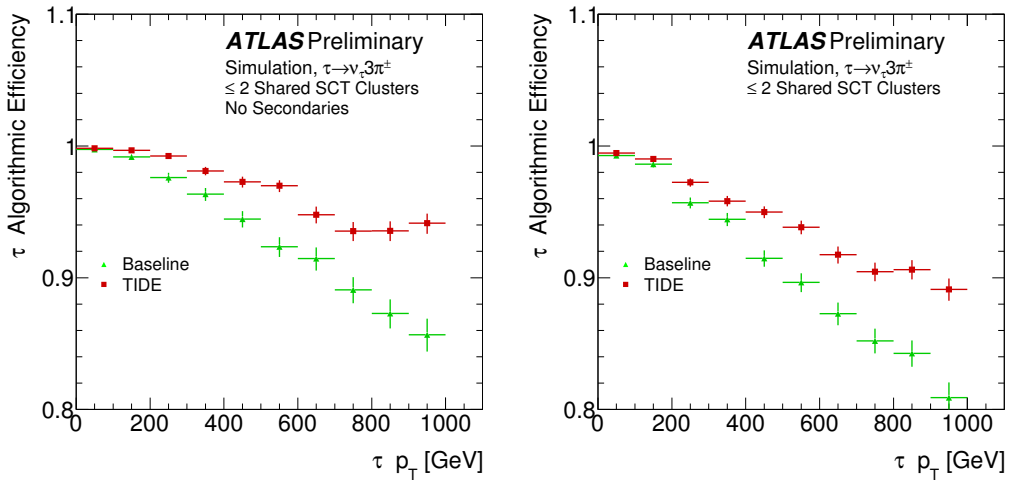

(b) single 3-prong τ sample

Figure 6: The efficiency to reconstruct all decay products of a ρ or 3-prong τ in events where truth-based tracks do not share more than two clusters in the SCT is shown as a function of the parent truth particle p_{T}. Left: events are restricted not to contain any secondary particles through nuclear interactions of the decay products with the detector material. Right: same set of plots, but no restrictions on the number of secondaries in the event. Section 5.3 lists the full set of differences between the baseline and TIDE reconstruction chain. This figure can be found in Ref. [1].

5.3 Summary of Changes

The major driving force of the changes is delaying the decision on how to use the information encapsulated within the NN. A similar strategy of taking decisions as late in the track reconstruction chain as possible was also the theme of Ref. [19]. The following is a summary of all the changes made:

1. The evaluation of the pixel cluster NN was delayed from cluster reconstruction to the track ambiguity solver.
2. The notion of split clusters was replaced by the shareable clusters definition:
   - Pixel clusters used by multiple tracks and passing the NN selection are shareable.
   - Pixel clusters used by multiple tracks and failing the NN selection are shareable if the same tracks all use a shareable cluster on the next layer.
   - Shareable clusters can be shared without penalty (split clusters were not allowed to be shared at all).
3. Tuning of existing cuts to optimize efficiency and rejection:
   - NN output for two particle hypothesis, threshold for being shareable: 0.35
   - NN output for more than two particle hypothesis, threshold for being shareable: 0.40
   - NN output for two particle hypothesis, minimum to be allowed to be shared: 0.05
   - Minimum number of silicon clusters required to allow sharing of track's clusters: 9
4. Cuts introduced to control fake rate:
   - Minimum number of unique SCT clusters: 4
   - Minimum $p_{\mathrm{T}}$of the track candidate to allow clusters to be shareable: 1 GeV

6. Track Reconstruction Performance in Jet Cores

In the previous sections, the changes to the track reconstruction chain have been described, and the impact on the performance in simple samples has been shown. These samples were crucial in understanding the problem at hand and testing alternate strategies, but they are insufficient for fine tuning and validation. This section shows the changes in cluster content of tracks within jets in fully simulated $Z' \to t\bar{t}$ events. Truth jets are constructed from generator-level particles with the anti-kt [10] algorithm using a distance parameter of $R = 0.4$. Only jets with p_{T} greater than 100 GeV and within $|\eta| < 2.5$ are used.

Figure 7 shows the change in the average number of clusters on the innermost pixel layer as a function of the track distance to the jet axis. A clear increase can be observed with the new (TIDE) reconstruction. The rate of having a split (for the original setup) or shareable (for the new setup) cluster on track is also shown. The number of shareable clusters increases with the particle density.

The resulting track reconstruction efficiency as a function of the jet's p_{T} is shown in Figure 8 for the old and the new setup. An increase in efficiency, especially for high jet momenta, is observed.

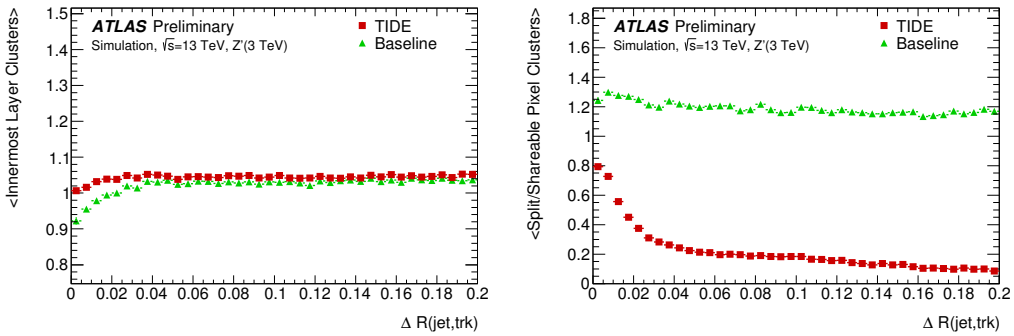

Figure 7: The average number of innermost pixel layer clusters on primary tracks (with a production vertex before the first layer) and number of split/shareable pixel clusters on these tracks are shown on the left and right respectively, as a function of the angular distance of the track from the axis of jets with $p_{\mathrm{T}} > 100$ GeV. Two track reconstruction algorithms are shown: green triangles label the baseline reconstruction (Section 3) and red squares label the TIDE optimized reconstruction (Section 5). This figure can be found in Ref. [1].

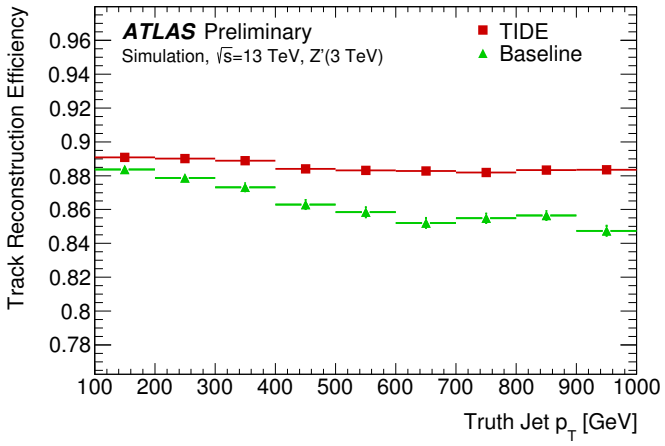

Figure 8: The average efficiency to reconstruct primary tracks with a production vertex before the first layer in jets as a function of jet p_{T}. The same sample generation, with limited statistics, is used for both reconstruction algorithms resulting in correlated features. Two track reconstruction algorithms are shown: green triangles label the baseline reconstruction (Section 3) and red squares label the TIDE optimized reconstruction (Section 5). This figure can be found in Ref. [1].

Figure 9 shows the charged-primary particle reconstruction efficiency dependence on the angular distance of a particle to the jet axis and the transverse distance from the interaction vertex to where the particle was created. Jets were selected to have a p_{T} between 450 GeV and 750 GeV. Charged particles were required to have been created within a radius ($R^{\mathrm{part}}_{\mathrm{prod}}$) of 100 mm and transverse the entire SCT detector ($R^{\mathrm{part}}_{\mathrm{decay}} > 600$ mm). A $p_{\mathrm{T}} > 2$ GeV requirement was also used. Several particles from a displaced decay can create a single merged cluster. Clusters created from more than two particles have been shown to reduce the cluster assignment efficiency (Figure 3) and track reconstruction efficiency (Figure 6) in samples created from a single B^{+} and single τ-lepton decaying to three charged hadrons respectively. The same effect drives the reconstruction efficiency decrease towards the core of a jet in fully simulated samples. The problem is exacerbate for b-jets

due to the displaced decay of heavy-flavor quarks. In all cases the new (TIDE) reconstruction provides an improvement. An approximate 17% efficiency gain is seen for charged particles created at a radius of 30 mm as well as a 10% (14%) improvement in the core of high p_{T} light (b) jets.

7. Impact on Flavor Tagging

Jet-flavor tagging exploits the lifetime of b-quarks via track impact parameters or the identification of displaced vertices [8]. Several taggers optimized for the different means of identification are combined in a multivariate technique. The impact parameter resolution is likely to degrade if the innermost measurement is missing or shared. In Run 1, impact-parameter-based taggers considered only tracks with a cluster on the innermost layer. Therefore, this type of tagger will profit from increasing the number and precision of innermost clusters on track. Secondary vertex taggers employ multivariate discriminants using vertex properties. Such variables include the secondary vertex mass, the vertex energy fraction or the momentum of the tracks in the vertex compared to all tracks considered by the discriminant, and the secondary vertex momentum. Increasing the efficiency for highly collimated track pairs improves the secondary vertex efficiency. Also, since the collimated tracks carry a considerable fraction of the charged-particle momentum, the vertex energy fraction and vertex momentum become more discriminant.

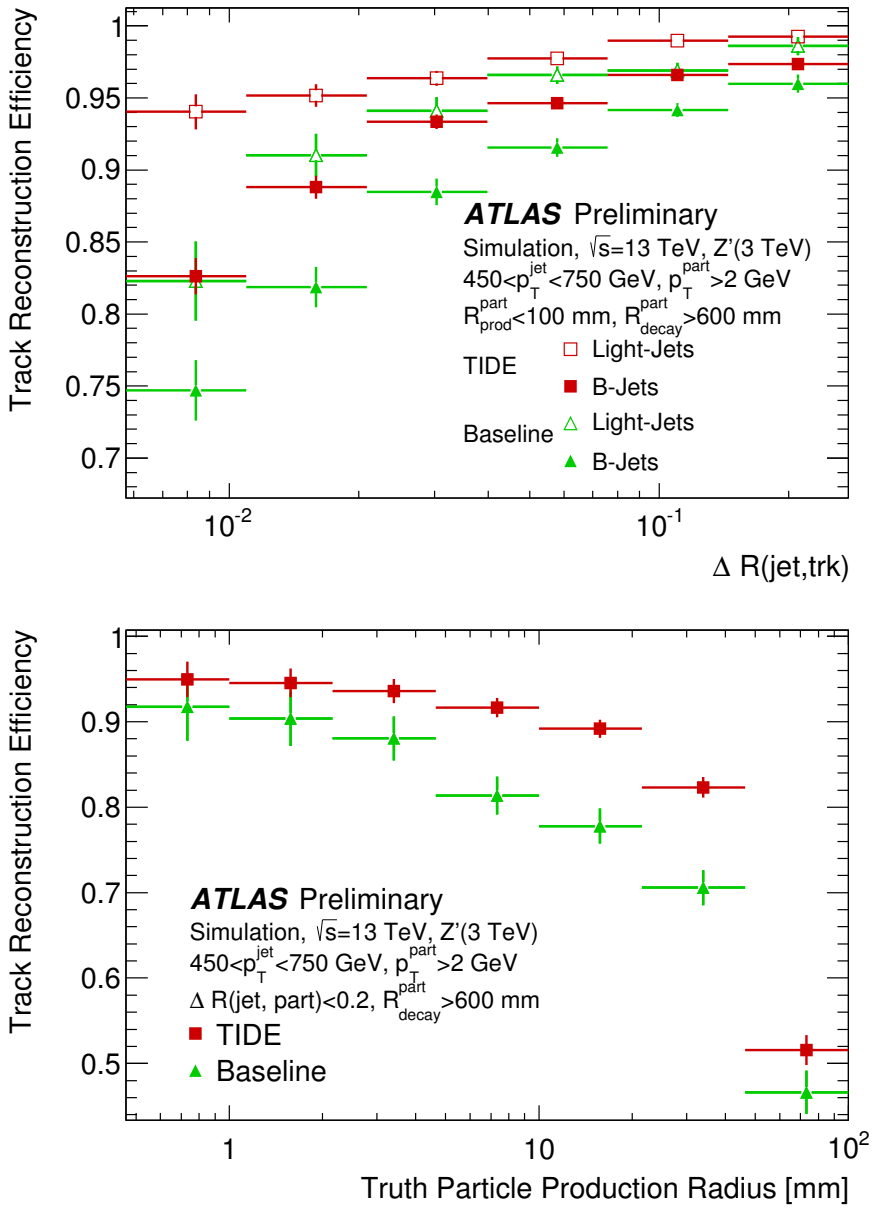

Figure 9: The average efficiency to reconstruct charged primary particles is shown as a function of the angular distance of the particle from the jet axis (top) and as a function of the production radius of the particle (bottom). Jets were selected to have a p_{T} between 450 GeV and 750 GeV. Charged particles were required to have been created ($R^{\mathrm{part}}_{\mathrm{prod}}$) within 100 mm of the beam line and transverse the entire SCT detector ($R^{\mathrm{part}}_{\mathrm{decay}} > 600$ mm). A $p_{\mathrm{T}} > 2$ GeV requirement was also used. The decrease in track reconstruction efficiency in the jet core and for displaced tracks is driven by clusters created from more than two particles. Also highly displaced tracks can lack a sufficient number of clusters to be reconstructed. The new (TIDE) reconstruction provides an approximate 17% efficiency gain for charged particles created at a radius of 30 mm as well as a 10% (14%) improvement in the core of high p_{T} light (b) jets. Two track reconstruction algorithms are shown: green triangles label the baseline reconstruction (Section 3) and red squares label the TIDE optimized reconstruction (Section 5). This figure can be found in Ref. [1].

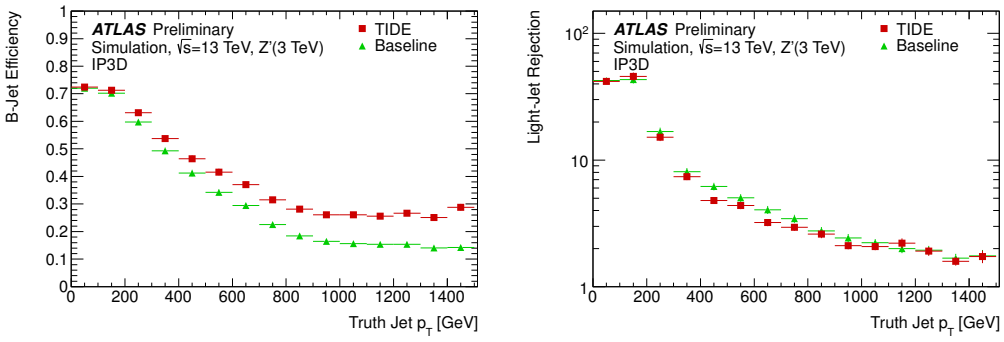

Figure 10: *B*-jet efficiency (left) and light-jet rejection (right) at the baseline 70% working point of the IP3D algorithm as a function of truth-jet transverse momentum. The truth jets are reconstructed from generator-level particles in Z' events, using the anti-k_t algorithm with R=0.4, and are required to be within $|\eta| < 2.5$. The performance with the new (TIDE) reconstruction (red squares) is compared with that using the baseline reconstruction (green triangles). This figure can be found in Ref. [1].

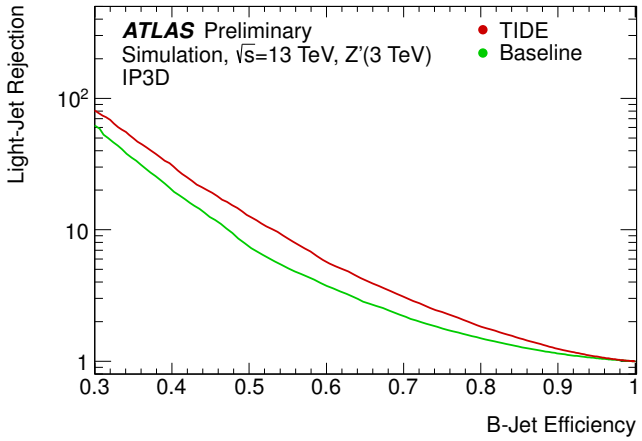

Figure 11: *B*-jet efficiency versus light-jet rejection for the IP3D algorithm with the new (TIDE) configuration as well as the baseline setup for anti-k_t R=0.4 jets with p_T greater than 100 GeVand $|\eta| < 2.5$ in a Z' sample are shown. For the baseline light-jet rejection of the commonly used 50, 60, 70, and 80% *b*-jet identification working points, an efficiency increase of 13, 11, 9, and 7%, respectively, is observed without retraining the likelihood. This figure can be found in Ref. [1].

Figure 10 shows the improvements to the impact parameter based tagger IP3D [5] for b-jet tagging efficiency as well as the effect on light-jet rejection as a function of the truth jet p_{T} for the Z' sample using anti-k_{t} R=0.4 jets with $|\eta| < 2.5$. The gain in efficiency versus rejection is shown in Figure 11. A relative 13 (7)% increase in efficiency is obtained for a constant light-jet rejection rate equal to that of the 50 (80)% b-jet identification working point. A slightly reduced improvement is seen if tracks without a measurement on the innermost layer are considered in the comparison as well. The enhanced performance is driven by the increased reconstruction efficiency of displaced charged particles as shown in Figure 9. The likelihood on which the IP3D tagger is based was not retrained with tracks built using the TIDE track reconstruction for the results shown here.

8. Conclusion

This note presents recent changes in the ATLAS track reconstruction chain derived from detailed studies of track reconstruction in dense environments, such as the cores of high p_{T} jets and τ-leptons that are characterized by charged particle separations comparable to the Inner Detector pixel cluster dimensions. The ambiguity processor stage of the reconstruction chain was optimized, including an improvement in the usage of a NN based approach to identify pixel clusters created by multiple charge particles. The changes were demonstrated with simulated samples of a single particle decaying to a set of collimated charged particles. The full impact in high p_{T}jets is up to 10% more IBL clusters on tracks in the jet core, a flat track-reconstruction efficiency as a function of jet p_{T}, and a more meaningful pixel cluster classification. A 17% efficiency gain is seen for charged particles created at a radius of 30 mm. Similarly a 10% (14%) track reconstruction efficiency improvement is shown in the core of high p_{T} light (b) jets. For heavy-flavor jet identification, a relative 7–13% increase in b-tagging efficiency for a given light-jet rejection rate for jets with p_{T} greater than 100 GeV is also obtained. A 7-13% increase in b-jet identification of jets with p_{T} greater than 100 GeV is shown for a given light-jet rejection rate and a 5% improvement in the reconstruction of three-prong τs with with p_{T} above 600 GeV has been demonstrated.

References

[1] The Optimization of ATLAS Track Reconstruction in Dense Environments. Technical Report ATL-PHYS-PUB-2015-006, CERN, Geneva, Mar 2015.

[2] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003, 2008.

[3] ATLAS Collaboration. Performance of the ATLAS Silicon Pattern Recognition Algorithm in Data and Simulation at $\sqrt{s} = 7$ TeV. (ATLAS-CONF-2010-072), Jul 2010.

[4] ATLAS Collaboration. The ATLAS Simulation Infrastructure. *Eur.Phys.J.*, C70:823–874, 2010.

[5] ATLAS Collaboration. Commissioning of the ATLAS high-performance b-tagging algorithms in the 7 TeV collision data. (ATLAS-CONF-2011-102), Jul 2011.

[6] ATLAS Collaboration. Summary of ATLAS Pythia 8 tunes. (ATL-PHYS-PUB-2012-003), Aug 2012.

[7] ATLAS Collaboration. A neural network clustering algorithm for the ATLAS silicon pixel detector. *JINST*, 9:P09009, Jun 2014.

[8] ATLAS Collaboration. Calibration of *b*-tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment. (ATLAS-CONF-2014-004), Feb 2014.

[9] ATLAS Collaboration. Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data. *Eur. Phys. J.*, C74(11):3130, 2014.

[10] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 0804:063, 2008.

[11] M. Capeans et al. ATLAS Insertable B-Layer Technical Design Report. (CERN-LHCC-2010-013. ATLAS-TDR-19), Sep 2010.

[12] G. Aad et al. ATLAS pixel detector electronics and sensors. *JINST*, 3:P07007, 2008.

[13] GEANT4 Collaboration, S. Agostinelli et al. GEANT4–A Simulation toolkit. *Nucl.Instrum.Meth.*, A506:250–303, 2003.

[14] Roland Jansky. Truth Seeded Reconstruction for Fast Simulation in the ATLAS Experiment. (CERN-THESIS-2013-194), 2013.

[15] A.D. Martin, W.J. Stirling, R.S. Thorne, and G. Watt. Parton distributions for the LHC. *Eur.Phys.J.*, C63:189–285, 2009.

[16] R. Frühwirth. Application of Kalman Filtering to Track and Vertex Fitting. *NIM*, A262:444–450, 1987.

[17] RA. Rosenfeld and J. Faltz. Sequential operations in digital picture processing. *J. ACM*, 13(4):471–494, October 1966.

[18] Sjöstrand, Torbjörn and others. An Introduction to PYTHIA 8.2. 2014.

[19] Daniel Wicke. A New algorithm for solving tracking ambiguities. (LC-TOOL-1999-007-TESLA), 1999.