Available on CMS information server **CMS NOTE 2006/030**

The Compact Muon Solenoid Experiment

CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

B-tagging in the High Level Trigger

M. Vos, F. Palla

INFN - sezione di Pisa, Italy

Abstract

A number of techniques exist in the CMS reconstruction framework that reduce the computing cost of track reconstruction. In this note, the b-tagging potential of these algorithms is discussed. The performance of an inclusive b-jets trigger is investigated. It is found that use of the b-tagging information can significantly improve the HLT efficiency of the inclusive jets stream for fully-hadronic final states. The most relevant physics case is given by top pairs with fully hadronic decays, where the b-jets streams yield a signal efficiency of 15 %.

1 Introduction

The beam energy and luminosity of the Large Hadron Collider will give rise to an unprecedented QCD production rate. The online selection of rare signal events among the background will be even more important than it has been in previous experiments. This is especially clear for fully hadronic final states. The large background rate will drive up the threshold in purely hadronic triggers like the inclusive jets trigger. The high thresholds may have severe implications for the physics programme for final states without isolated high p_T leptons.

As an illustration, consider the case of fully hadronic $t\bar{t}$ events, i.e. top pairs with the following decay chain: $t\bar{t} \to W^-bW^+\bar{b} \to q\bar{q}bq\bar{q}\bar{b}$. During Run I in the Tevatron these events were efficiently triggered by a 4-jet trigger with a transverse energy threshold of 10 GeV [1]. As the LHC luminosity ramps up, the equivalent 4-jet threshold will have to be raised considerably to control the background rate. In the default low-luminosity (L = $2 \cdot 10^{33} cm^{-2} s^{-1}$) trigger table in the trigger and DAQ TDR [2] an E_T threshold of approximately 90 GeV [1] is envisaged for the fourth jet. In that case, the trigger efficiency for hadronically decaying top pairs would amount to only a few per cent.

Many interesting fully hadronic final states contain one or several b jets: the top pairs already mentioned, but also hadronic decays of bbH, SUSY decays, etc. The selection of such final states can greatly benefit from the experimental sensitivity for the presence of b jets. Efficient b-tagging allows the background to be significantly reduced.

It seems natural to exploit the b-tag signature to reach the best possible trigger performance. In this note, the possibility to use b-tagging in the HLT is explored in some detail.

In the HLT environment severe constraints are posed by the available CPU time. Several techniques are employed in the CMS HLT to reconstruct high-quality tracks while minimizing the computing load. In section 2 we discuss in detail the reconstruction efficiency, fake rate, impact parameter resolution, and CPU time requirement of two of these techniques.

Another issue of concern is the robustness of the algorithms under online conditions. The b-tagging algorithms based on lifetime require the 3-dimensional position of the primary vertex to be accurately known. The validity of the nominal beam position in the transverse view is discussed. The z-coordinate (along the beam line) of the primary vertex has to be reconstructed for each event. The performance of the pixel-only primary vertex finding algorithms [5] is briefly revisited. Both issues are discussed in section 3.

Once the quality of tracks and the primary vertex is established, the b-tag discriminant can be determined. The performance of the chosen algorithm on a benchmark sample is discussed in section 4.

Finally, the implementation of an inclusive b-jets trigger is discussed in section 5. The background rate is studied in detail. As a benchmark signal, $t\bar{t}$ events with a fully hadronic final state are considered.

In section 6 the most important findings are briefly summarized.

2 Fast track reconstruction in the HLT

The CMS High Level Trigger operates on events from the Event Builder: i.e. the full granularity of the detector is available. Moreover, the Filter Farm offers the full flexibility of a software environment to develop triggering algorithms. The most restrictive constraint is therefore posed by the computing power available to process each event.

The full-fledged Kalman filter based algorithm intended for the offline track reconstruction [3] in the tracker presents a heavy computing load. Therefore, in the following sections several techniques are investigated which speed up track reconstruction, optimizing the efficiency and fake track rate while still retaining as much resolution performance as possible.

2.1 Regional tracking

This solution makes use of the offline algorithm [3], but in this case it is applied only in certain regions of the detector that are likely to produce interesting additional information (regions of interest, typically identified by the

---

[1] This threshold on the reconstructed object corresponds to 95 % efficiency point for generated partons of 113 GeV as in table 15-24 of reference [2]

previous trigger level). Regional tracking is an established technique in the CMS reconstruction software. It is very well adapted to several important steps in the HLT: an important gain is made in the timing, while the track quality does not degrade with respect to the full offline reconstruction. In reference [2] it was shown that thanks to this technique, b-tagging in the High Level Trigger is possible.

The regional seeding algorithm implemented in the CMS reconstruction software [4] allows only those seeds in the pixel detector to be created that are compatible with the selected region. The region limits the allowed values for the seed parameters in several ways. The direction of the seed is required to lie within a certain distance from the jet axis. A maximum is set on the azimuthal angle between track and jet axis and on their difference in pseudo-rapidity. In the current implementation these are set to $\Delta\phi = 0.25$ rad and $\Delta\eta = 0.25$. Further, the track vertex is constrained. Here, the track ought to originate within 2 mm of the primary vertex, both in the transverse ($R - \phi$) and longitudinal ($R - z$) plane. Finally, only those seeds compatible with a minimum transverse momentum - of 1 GeV/c - are kept.

Trajectory building and fitting then proceeds in the same fashion as for global tracking. There is no further check that the resulting tracks are compatible with the region.

As these tracks are very similar to the offline tracks, the canonical track quality requirements applied by the b-tagging algorithms can be applied:

- at least 8 hits in the silicon (pixel+strip) tracker
- at least 2 hits in the pixel detector
- a minimum transverse momentum of 1 GeV/c
- a maximum transverse impact parameter of 2 mm

2.2 Pixel-only tracking

A more radical approach is to reconstruct tracks based only on the information of the pixel detector. Over virtually all of the tracker acceptance, tracks cross three pixel layers, each of which provides a three-dimensional space point with excellent ($\sim 12\mu$ m) resolution. The pixel detector granularity (pixels measure $150 \times 100\mu m^2$), allows the occupancy to be kept at the level of 10^{-4} even in the innermost region of the CMS tracker. It is due to these characteristics that the information from three pixel layers suffices to constrain the pattern recognition.

The pixel reconstruction algorithm implemented in the CMS reconstruction software [5] is based on a relatively simple combinatorial approach: all triplets of hits in consecutive layers are created [2]. In practice, the main difference with the offline track reconstruction algorithm consists in the fact that no attempt is made to recover tracks where one of the layers was inefficient: three hits are required in three layers. Thus, the track finding efficiency is expected to scale as the hit efficiency to the third power.

Of course, the requirements on the number of hits have to be adapted. The track quality cuts for pixel-only tracks read:

- 3 pixel hits (inherent in reconstruction)
- a minimum transverse momentum of 1 GeV/c
- a maximum transverse impact parameter of 2 mm

In the following section, the performance of this algorithm is compared to that of the regional Kalman-filter based algorithm.

2.3 Comparison of tracking performance

To investigate the potential of the regional and pixel-only track reconstruction for online b-tagging, the tracking performance inside high E_T jets is studied in detail. To this end, reconstruction is performed in the two leading jets (ordered in order of decreasing transverse jet energy) in hadronically decaying $t\bar{t}$ events.

---

[2] Recovery of hit pairs is possible, but will not be considered in this note

In the CMS HLT environment, algorithms of arbitrary complexity can in principle be implemented. The most severe constraint is posed by the available CPU time. Several techniques are employed in the CMS HLT to reconstruct high-quality tracks at a minimum computing load.

A considerable speed-up is obtained by a regional application of the algorithm. Typically, regions-of-interest are defined on the basis of the result of the preceding trigger level (a calorimetric cluster, the extrapolation of a track in the muon chambers, etc.) Seeds are reconstructed only in a multi-dimensional region (constraining the η ϕ direction of the tracks, the origin in the $R\phi$ and Rz planes, and the minimum transverse momentum). Trajectory building and fitting then proceeds in the same fashion as for global tracking.

The time performance of the regional approach allows the default offline algorithm - the combinatorial Kalman filter based track finder - to be used in the later stages of the high level trigger. The track reconstruction thus obtained is of comparable quality as the offline reconstruction. Only at the very edges of the region parameters is a minor degradation of the efficiency observed (due to the limited precision of seeds). The fake rate can be controlled to the level of 1 % by an adequate choice of track quality requirements.

For the earlier stages of the high level trigger, speed becomes even more of a concern. An extremely fast reconstruction based on hit triplets in the pixel detector has been developed. The algorithm is described in detail in reference [5]. Pixel-only reconstruction is sufficiently fast that global reconstruction of all tracks with transverse momentum greater than 1 GeV can be performed in the high level trigger.

The simplified pattern recognition has to rely on three hits out of three pixel layers, thus posing a severe requirement on the single layer efficiency. The simulation of the CMS tracker contains a detailed description of a long list of efficiency loss sources [6]. Based on this simulation, the effect on the efficiency is expected to be quite limited [7]. But, of course the robustness of the algorithm against defective components is much reduced.

The fake rate is rather well controlled by the three-out-of-three requirement (to the level of 10 %, see reference [3]) and can be further improved by requiring compatibility with the primary vertex.

The track parameter resolution suffers from the small lever arm (the pixel barrel layers are located at radii of 4, 7 and 10 cm). For a 10 GeV track, the resolution of the full tracker is better than 1 %, whereas the pixel-only tracks have a resolution of approximately 25 %.

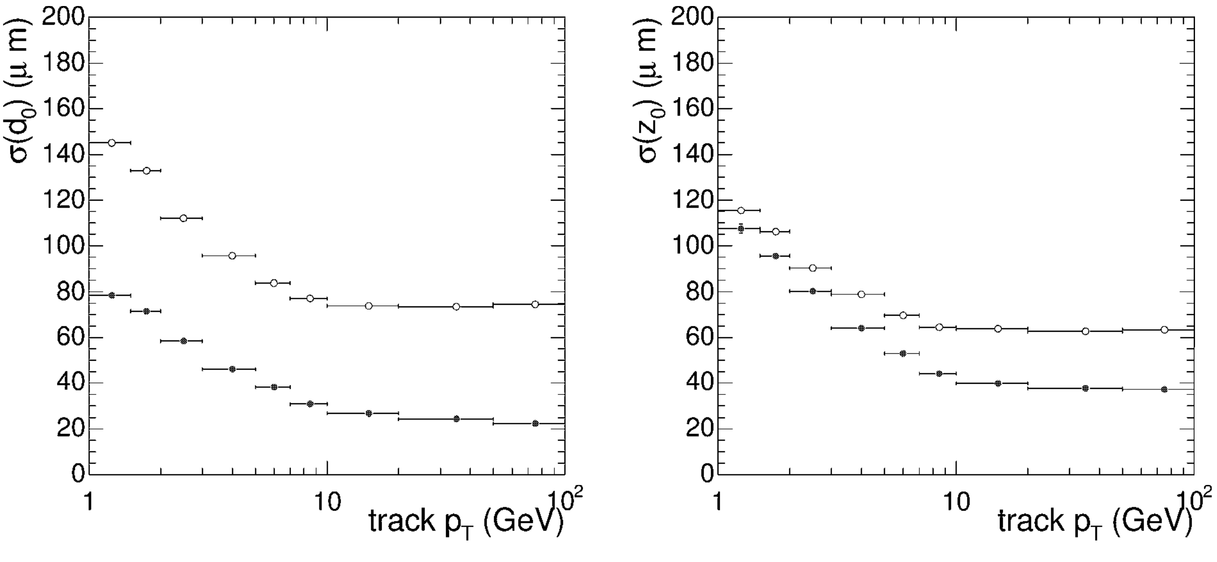

Figure 1: Resolution - measured as the width of a Gaussian fit to the residual distribution - for the transverse (leftmost figure) and longitudinal (rightmost figure) impact parameter versus transverse momentum of the associated simulated track. Closed markers represent regional track reconstruction. Open markers correspond to pixel-only tracks.

In figure 1 the resolution on the transverse and longitudinal impact parameter for tracks in large E_T-jets in the hadronic $t\bar{t}$ sample is shown. Open markers correspond to pixel-only tracks, while closed markers represent the full tracker performance. The momentum uncertainty affects the precision of the impact parameter estimate, leading to a significant degradation of the pixel-only d_0 resolution with respect to the full-tracker performance. In the longitudinal plane, the resolution is much less affected. The performance shown here for both the full tracker and the pixel-only tracks is in good agreement with that found for single-track performance in reference [2] and [5].

The high-occupancy environment in high E_T jets does not significantly degrade the performance.

For this note, the effect of tracker mis-alignment has not been studied. CMS has defined several mis-alignment scenarios, all of which assume that a preliminary pixel alignment to 10 μm is available. Of course, for a reasonably accurate alignment to be available in the online environment, the return time to produce an alignment should be small with respect to the time-scale of the tracker deformations.

In reference [9], the degradation of the impact parameter resolution for the various scenarios is studied in detail in reference. For the performance of the track counting b-tagging algorithm a small but clear degradation is observed for the various misalignment scenarios [10]. The level of the performance loss is of minor importance for b-tagging in the online environment.

Most b-tagging algorithms based on impact parameter rely on the significance of the impact parameter measurement with respect to the nominal vertex. Therefore, an accurate estimate of the track impact parameter errors is as important as the resolution on the measured value itself.

The starting point for the error estimate is given by the space point error on the hits in the tracker. The hit errors are estimated carefully during RecHit creation. For low-momentum tracks, the hit error is dominated by multiple scattering in the tracker material. A parameterization of this contribution to the error in all tracker layers is available in the CMS reconstruction software.

The requirements on the pixel-only algorithm - intended for use in the HLT - are rather different from those of an offline algorithm. Generally, the time constraint is more important; and moreover, not all use cases require accurate errors. To cope with the different needs, the pixel track finder can be configured to use one of several track parameter evaluation algorithms.

For the purpose of b-tagging, a detailed numerical error calculation has been implemented. Naturally, the more detailed calculation might lead to a loss of algorithm speed. In the current use case, however, the track parameter evaluation does not lead to a significant degradation of the time performance.

In the pixel detector, knowledge of the track incidence angle allows the resolution and error estimate of the individual hits to be improved [6]. Therefore, the track parameters are evaluated in two iterations. A first track fit uses the hit positions and errors obtained with just the assumption that the track comes from the nominal collision region. In the second iteration, the track direction is given to the algorithm that creates the RecHits.

The combined effect of hit error and multiple scattering contribution for the three hits on the impact parameter is evaluated by a numerical determination of the derivatives $\delta d_0/\delta x$, where for d_0 one should read the transverse, longitudinal and 3D impact parameter, and x denotes either of three coordinates of the hit.

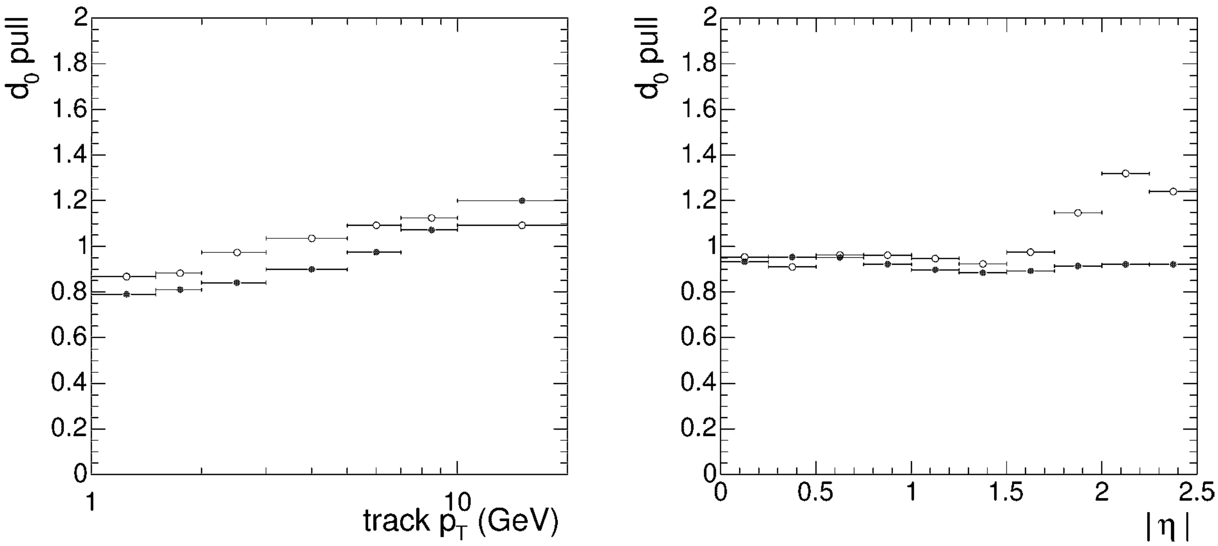

Figure 2: Width of the pull distribution for the transverse (leftmost figure) and longitudinal (rightmost figure) impact parameter versus pseudo-rapidity of the associated simulated track. Closed markers represent regional track reconstruction. Open markers correspond to pixel-only tracks.

In figures 2 the stability of the width of the pull distribution over the transverse momentum and pseudo-rapidity range of interest are shown. Except for pixel-only algorithm in the disk-region, the errors are found to be accurate to better than 20 % throughout the parameter space.

2.4 Summary

As expected, regional application of the combinatorial track finder yields excellent tracks. Except for a slight loss of efficiency toward the edges of the region, the performance is comparable to that of the offline algorithm. The time performance of the algorithm would allow it to be used in the latest stages of the HLT.

The performance of the pixel-only track finder is found to be very satisfactory. Compared to the offline or regional algorithm, the efficiency is slightly reduced. The transverse impact parameter resolution is significantly degraded, but the fake rate is very well controlled and the loss in resolution on the longitudinal impact parameter is very limited. Accurate error estimates can be obtained without deteriorating the time performance.

Considering the gain in the timing performance, pixel-only tracks become a very interesting alternative in the early stages of the HLT.

3 Vertex

In this section, the characteristics of the interaction region are discussed in the light of the online b-tagging. Detailed expectations for the beam spot and variations therein on various time scales have been drawn up by the interface group:

- The beams are quite narrow: 16 μm.
- Interactions take place over a long distance along the beam axis: 90 % of collisions ought to occur within 12 cm of (0,0,0)
- During the physics coast (which should take of the order of 10 hours) a very small variation is expected: smaller than 20 % of the beam width.
- The position of the center of the collision area is expected to be reproducible from one fill to the next only at the level of 1 mm.
- After the initial period of very low luminosity where the beams collide head-on, the crossing angle will be fixed to 285 μ rad. This angle is expected to be extremely stable throughout the physics coast and even between fills.

Given the small beam size, the position of the center of the beam line can be used by online algorithms as the primary vertex position in the transverse plane without significant loss of precision. During the fill, the variations in beam position are negligible. From one fill to the next, however, the beam spot will not be reproducible to better than O (mm). The reconstruction of the beam spot is discussed in detail in section 3.1.

Depending on the luminosity, several to several 10s of collisions will take place each bunch crossing. The primary vertex has to be reconstructed and identified on an event-by-event basis. The performance of the so-called fast divisive algorithm on pixel hit-triplets is discussed in section 3.2.

3.1 Beam spot reconstruction

To correct for the variations in the beam spot position at each fill, a very fast (order of a second) reconstruction of the beam spot can be performed on the first events using the information from the tracker. The precision of this procedure should be comparable to the beam size (i.e. 15 μm).

With the speed requirement mentioned above, it is crucial that the beam spot determination give a precise, unbiased result with a limited number of events and tracks. Moreover, to render the beam spot determination independent of the correct functioning of the high-level triggers, or even the L1 triggers, the method should be independent of the exact type of events fed into it. Of course, the beam spot determination should be efficient over the full range of positions, i.e. even for macroscopic displacement of the order of several mm.

The $d_0\phi$ algorithm has been used successfully in previous experiments. For each track the transverse impact parameter d_0 is determined. This distance is then signed according to the angular momentum convention. The displaced beam spot will result in a net offset in the impact parameter distribution which depends on the orientation of the track. If the track momentum is along the line linking the nominal and true beam positions in the transverse plane, no net impact parameter is found. For tracks perpendicular to this line, the displacement is maximum. Plotting d_0 versus $/phi$, as in figure 3.1, the typical shape is observed. A fit with a sine allows the two polar coordinates of the beam spot that fix its position in the transverse plane to be determined. If the detector symmetry axis and the beam line are not parallel, the $d_0\phi$ method may be applied in narrow slices along the z-axis.

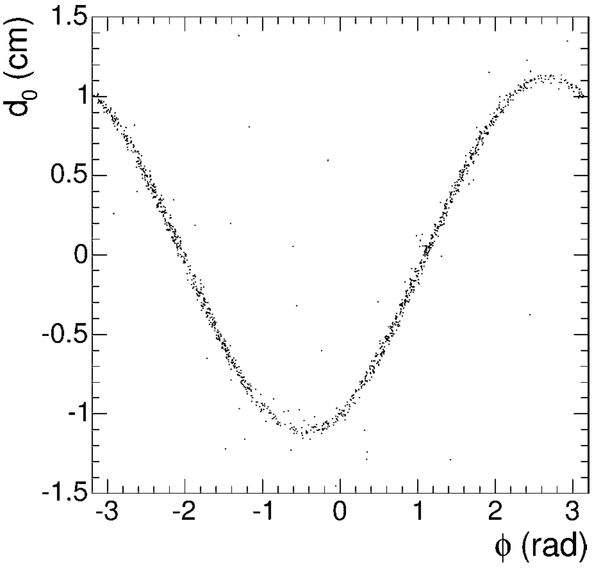

Figure 3: Scatter plot of transverse impact parameter with respect to the nominal beam spot and signed by the convention versus track azimuthal angle at impact point. Only tracks satisfying the track quality requirements discussed in the text are taken into account. Result for 1000 minimum bias event without pile-up.

To study the performance of the $d_0\phi$ algorithm, a small sample of minimum bias events with displaced beam spot has been generated. The displacement amounts to 1 cm in X and 5 mm in Y. No pile-up is included in the current analysis. After simulation and digitization, the tracks are reconstructed using the standard CMS combinatorial track finder [3]. Only one parameter has been changed with respect to the default operation of this track finder: the maximum distance to the nominal (0,0) beam spot allowed in the pixel seed generator is increased to 2 cm.

A number of track quality requirements are applied:

- at least 8 hits in the silicon tracker
- at least 2 hits in the pixel detector
- track fit $\chi^2 < 5$
- transverse impact parameter error $\sigma_{d_0} < 150\mu$m

These cuts are needed to reduce the fake rate to the level of a few per cent. The track reconstruction efficiency is slightly reduced for the displaced beam spot. Where an average of 2.8 tracks per event remained after quality cuts for nominal beam spot, this number drops to 2.3 when the beam spot is displaced by more than 1 cm.

The polar coordinates of the beam spot are inferred from a simple least-squares fit with a sine function to the shape in figure 3.1. Outliers - due to fake tracks or conversions - might gain a large weight in the fit and significantly bias the result. Therefore tracks giving entries that are far ($> 3mm$) from the initial fit are removed. It is found that only two or three iterations are needed to reach a stable result.

To evaluate the performance of this simple algorithm several hundred small samples are created. For each of these the reconstructed beam spot is compared to the simulated value. The resolution of the algorithm is evaluated as the width of the residual distribution of polar coordinates d_0 and ϕ. For samples of 500 events (containing of the order of 1000 tracks) a resolution of 8 μm is found for d_0 and of 0.7 mrad for ϕ, corresponding to an error of approximately 8 μm in the transverse plane.

We conclude that the $d_0\phi$ method yields a sufficiently precise beam spot determination using only 500 minimum bias events.

3.2 Primary Vertex

The maximal b-tagging performance is obtained using the 3D impact parameter, which combines the information in the transverse and longitudinal plane. In that case, however, a measurement of the position along the beam line of the primary vertex of the event is needed.

Two primary vertex finding algorithms based on pixel triplets are available in the CMS reconstruction software [5]. Vertices are found by a histogramming or a divisive method on the intersection of the pixel tracks with the ideal beam axis. The primary vertex is selected by choosing the vertex with the highest sum of track transverse momentum (the quantity that this selection is based on is actually $\sum p_T^2$, where p_T is the track transverse momentum, truncated at 10 GeV/c). As these algorithms are primarily intended for use in the High Level Trigger they combine excellent timing performance to high efficiency. The reconstruction and selection efficiency and the spatial resolution are listed in table 1. Results are given for a benchmark signal sample (hadronically decaying $t\bar{t}$ events) and the background relevant for the proposed b-jets trigger. In all cases the detector response is simulated in full detail in the GEANT4 based simulation package [8]. Hits from pile-up corresponding to a luminosity of $2 \cdot 10^{33} cm^{-2} s^{-1}$ are added. Finally, the events are digitized and reconstructed using the standard CMS reconstruction software [4].

Table 1: Performance of the divisive method [5] for primary vertex finding in some relevant low-luminosity (L = $2 \cdot 10^{33} cm^{-2} s^{-1}$) samples. The first and second columns show the reconstruction and selection efficiency (the definitions are explained in the text). The resolution of the z-coordinate measurement is listed in the third column.

QCD $\hat{p}_T$ bin (GeV/c)	*reconstruction eff. (%)*	*selection eff. (%)*	*resolution σ (μm)*
50-80	99.0	94.7	36
80-120	99.5	97.6	31
120-170	100.0	98.8	29
170-230	100.0	99.7	25
230-300	100.0	100.0	24
300-380	100.0	100.0	23
hadronic $t\bar{t}$	100.0	100.0	27

The reconstruction efficiency, i.e. the fraction of events for which at least one vertex is reconstructed is very close to 100% for all samples considered here. The probability to select the correct vertex is measured as the fraction of events for which the z-coordinate of the selected primary vertex lies within 200μm of the true primary vertex. As b-tagging requires a very precise primary vertex determination, a window of 400μm is chosen here, rather than the $1mm$ in reference [5]. For the softer di-jet events in the first three bins, a significant efficiency loss (up to 5%) is found.

In the third column the resolution of the $z-$coordinate determination is listed, determined as the σ of a Gaussian fit to the residual distribution. As expected, the best resolution is for the hardest events. In any case the error is small compared to the error in the longitudinal impact parameter measured by a single track.

The efficiency and resolution of a fast primary vertex finding algorithm based on tracks reconstructed in the pixel detector is satisfactory for the type of events relevant to the b-jets trigger.

3.3 Summary

B-tagging algorithms rely on an accurately known primary vertex.

The beam spot position in the transverse plane is expected to be sufficiently stable during the LHC physics coast. From one fill to the next, however, the position is expected to be reproducible only to the level of O (mm). An algorithm based on reconstructed tracks that would provide a fast and sufficiently precise beam spot measurement on the first events has been identified.

With a minor change of one of its parameters, the standard combinatorial track finder, seeded using pixel hits, provides efficient reconstruction of tracks that originate from a vertex displaced by up to 1 cm. The $d_0\phi$ method yields a beam spot measurement with an error that is sufficiently small compared to the beam size using only 1000 tracks.

In the longitudinal ($R-z$) view, the primary vertex has to be reconstructed on an event-by-event basis. The performance of the fast divisive algorithm [5] on pixel hit triplets is found to be sufficient for the purpose of b-tagging.

4 b-tagging @ HLT

In this section, the b-tagging performance of the fast tracking algorithms described above is discussed.

The essential observable for life-time based b-tagging algorithms is the signed impact parameter significance. Traditionally, the impact parameter is available in two orientations, i.e. measured separately in the $R-\phi$ and $R-z$ planes. If the point of closest approach between track and jet axis can be found accurately, the 3-dimensional impact parameter becomes available. In any case, the significance is determined as the ratio of the measured impact parameter and its error. Provided the errors are accurate, this measure indicates the compatibility of the track with the (primary) vertex. Tracks from the primary vertex form a narrow, symmetric distribution around 0. Finally, the impact parameter significance is signed depending on the projection of the point of closest approach on the jet axis. If this projection lies in the same direction as the jet, a positive sign is assigned, while if it lies in the opposite direction it is negative. With this sign convention, tracks stemming from decays of relatively long-lived particles form a tail toward large positive values.

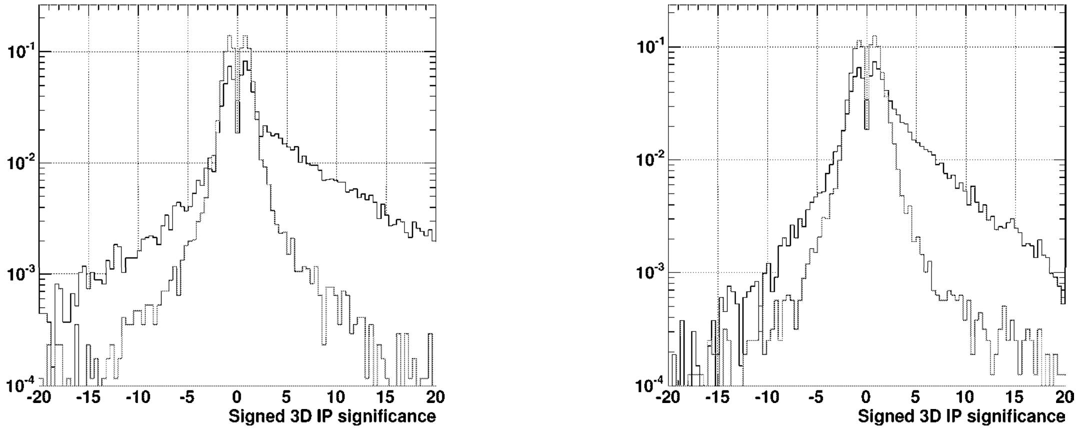

Figure 4: A comparison of the 3D impact parameter distribution for regional reconstruction (leftmost figure) and pixel-only reconstruction (rightmost figure). The 3D impact parameter distribution for tracks in b jets (filled histogram) and in light jets (neither b nor c, empty histogram) is shown. The results are obtained on reference (b-enriched) di-jet samples with jet E_T from 50-80 GeV. The jet flavor is determined from Monte Carlo truth information by associating the jet to the heaviest parton (after FSR) in an association cone of $\Delta R = 0.5$.

In figures 4 the signed 3D impact parameter significance distribution is shown for tracks inside b jets (filled histogram) and all tracks inside the leading two light jets, i.e. jets that are not associated to b or c partons (empty histogram). The tail due to the true lifetime is clearly visible in both distributions. The excess in the positive tail for tracks from b-decay is slightly reduced in the case of pixel-only reconstruction, an effect of the larger errors.

The timing overhead from the b-tagging algorithm is negligible. Therefore, standard b-tagging algorithms can be used in the HLT. The choice to apply standard, offline algorithms for the HLT b-tagging has a number of important implications. It requires, however, that the software objects returned by the regional and conditional version of the offline track reconstruction and the pixel-only algorithms implement a common set of methods; typically they should inherit from a common base class in the CMS reconstruction software.

The track counting algorithm, described in reference [10], combines excellent performance with a great conceptual simplicity. Moreover, the algorithm does not rely on any kind of calibration other than the track error parameterization. Therefore, this algorithm is considered most appropriate for HLT b-tagging.

In figures 5 the b-tagging performance of the track counting algorithm is shown for tracks reconstructed regionally in the full tracker and for pixel-only tracks. In both cases the primary vertex is reconstructed using the divisive algorithm based on pixel hit-triplets described in section 3.2.

The two curves in each figure represent the b-tagging efficiency versus mis-tag probability. The lower curve (round markers) displays the result for b jets versus light jets in the fully hadronic $t\bar{t}$ sample. The upper curve (square

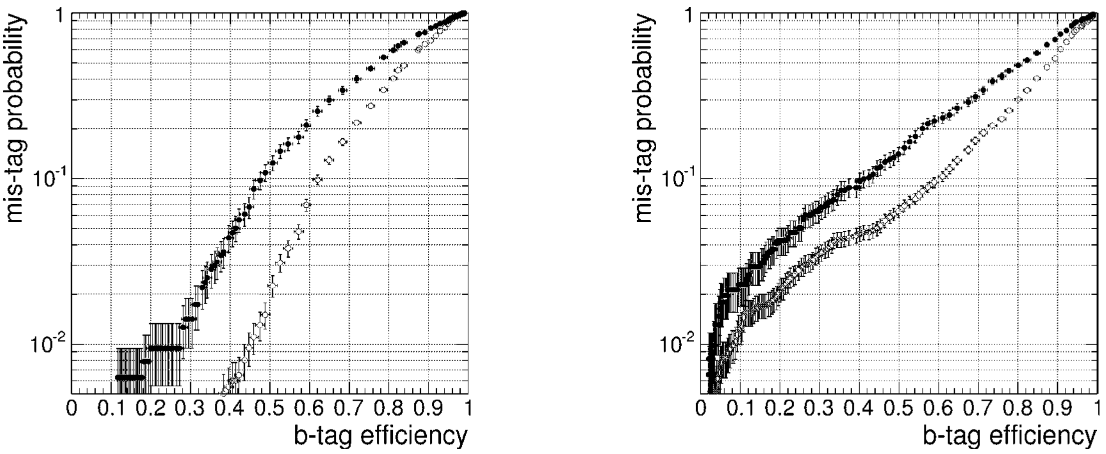

Figure 5: Regional (leftmost figure) and pixel-only (rightmost figure) reconstruction b-tag efficiency versus mis-tag rate the leading jets in hadronic $t\bar{t}$ events. The curves are obtained by scanning over the signed 3D IP significance for the N^{th} track in the jet (ordered in order of decreasing significance), where $N = 2$ for pixel-only reconstruction and $N = 3$ for regional reconstruction. The curve for b-jets versus light jets is shown with open markers, that for b-jets versus c jets with filled markers. All jets are classified according to the heaviest parton found in the cone. Typically, b jets in this sample stem from top decay (i.e. $E_T \sim m_t/2$), while light and c jets are mostly from W decay (i.e. $E_T \sim m_W/2$).

markers) shows the c-jet rejection.

The b-tagging performance measured on the $t\bar{t}$ sample is valid only for a limited interval of jet energies (and thus $\hat{p}_T$ range of di-jet events). For harder jets the performance of the track counting algorithm with the parameters chosen here is much degraded [10]. However, the jet energy distribution in the $t\bar{t}$ sample is representative for the dominant QCD di-jets contribution to the b-jets trigger rate.

For regional reconstruction the discriminant b-tagging variable is chosen as the signed 3-dimensional impact parameter significance of the third track of the jet (tracks are ordered in order of decreasing signed 3D IP significance). To accomodate for the lower efficiency in the case of pixel reconstruction, the second track is considered in that case.

Clearly, the performance of the track counting algorithm on regionally reconstructed approaches that of the offline algorithm in the relevant region. In the HLT, improving the light jet rejection beyond a factor 50 (i.e. a mis-tag probability of a few %) does not lead to a better trigger performance, as the $b\bar{b}$ production becomes dominant at that level. The observed (small) difference in performance is caused by the lower efficiency for low p_T tracks.

The pixel-only b-tag performs very well when a rather more modest rejection of the order of 5-10 (i.e. a mis-tag probability of 10-20 %) is required. For this range a very good efficiency can be obtained.

An HLT b-tag may very well benefit from the complementary performance of both algorithms by applying both algorithms in series. The pixel-only algorithm allow events that most certainly have no b-content to be rejected extremely fast. The time thus gained may be reinvested in performing a very precise b-tag for the remaining events. Moreover, the pixel-only algorithm may define the region-of-interest for the second step. The much slower, but much more precise regional reconstruction is then applied to a single jet in a limited fraction of events.

Based on figures 5 the working point for the track counting algorithm is chosen. For regionally reconstructed tracks, the signed 3D IP significance of the 3rd track is required to be greater than 3. Thus, the b-tag efficiency is 48 %, for a light jet mis-tag probability of 1.4 % and a c mis-tag probability of 10 %.

A very different working point is chosen for the b-tag based on pixel-only tracks in agreement with its is envisaged role of a fast "preselection". To limit the efficiency in this step, only 2 tracks with a signed 3D IP significance of 2.5 are required. Thus, the b-tagging efficiency remains high at 74 %, while a mis-tag probability of 20 % for light jets and 34 % for c jets is obtained.

For events with true b-content, the results of the two b-tagging algorithms are extremely correlated. Jets that do not make the soft requirement of the pixel-only b-tag, will almost certainly not be tagged by the algorithm based on regionally reconstructed tracks. The b-tag efficiency of both steps executed in series is not significantly lower than the efficiency of the regional b-tag alone. Thus the pixel-only "preselection" step can dramatically speed up rejection at virtually no cost in efficiency.

4.1 Summary

Due to its excellent performance and conceptual simplicity, the track counting algorithm seems most appropriate for use in the online environment.

As expected from the tracking performance found in section 2 the regionally reconstructed tracks yield a b-tagging performance similar to that of the offline algorithm.

A b-tagging algorithm based on pixel-only tracks offers a very decent tagging power, especially for light-quark rejections in the range from 5 to 20.

The two tracking algorithms complement each other very well. Pixel-only tracks can be run early in the HLT to reject a large fraction of background events. Jets that would seem most promising at this level can be investigated in full detail using the regional algorithm.

5 Inclusive b-jet trigger

In this section, the implementation of an inclusive b-jet trigger in the CMS reconstruction software [4] is discussed.

The principal limitation to the use of b-tagging in the HLT comes from the large $b\bar{b}$ cross-section (55 mb). For reasonably hard scatters, approximately 5 % of events contain genuine b jets. Therefore, rejections greater than 20 are doomed to be very inefficient on the signal. On the other hand, the HLT stage ought to reduce the relevant L1 rate (10s of kHz) to the bandwidth allocated on persistent storage (assume several Hz for the inclusive b-jets stream), i.e. a rejection by more than 3 orders of magnitude.

From these rough estimates it is clear that the b-tag information can only be put to use in conjunction with some other requirement that provides an additional rejection of a factor 50.

In the following sections, the different trigger levels 3) are described in detail. Then, the performance of the algorithm on the background and a benchmark signal - hadronically decaying top pairs - is discussed.

5.1 L1

Several L1 streams may contribute to the signal efficiency. For hadronic final states, the most obvious candidates would be the multi-jet triggers, but unfortunately the E_T thresholds - 135 GeV for a single jet, 57 GeV for the third and 45 GeV for the fourth jet at low luminosity (L = $2 \cdot 10^{33} cm^{-2} s^{-1}$) [2] - are already prohibitive at this level. The quoted numbers are the threshold on the reconstructed objects. They correspond to a 95% efficiency for a parton E_T of 177 GeV, 85 GeV and 70 GeV, respectively.

A very significant contribution may be expected from the (non-isolated) lepton triggers (the muon and electron stream, or combinations of lepton+jet, lepton + MET) through semi-leptonically decaying b jets. Therefore, all muon, electron and jet first level triggers are accepted in the b-jets trigger. The signal efficiency of the first level trigger is discussed in detail in section 5.7.

5.2 HLT - L2

In this implementation, the b-tagging requirement is combined with an inclusive jets trigger. The application of an E_T threshold on the n-th jet gives a large rejection of soft di-jet events, the dominant source of background. Moreover, jet reconstruction from the calorimetry information is relatively fast. It is natural to apply this requirement before the b-tagging is invoked.

For this study, the iterative cone algorithm is applied to so-called ECALPlusHCALTowers, projective towers that sum the energy deposits in 5 × 5 crystals in the electromagnetic and one cell hadronic calorimeters. The jet transverse energy is calibrated using the EGamma calibration. The thresholds at low luminosity on the E_T of the 1st, 2nd, and 4th jet are set to 350, 150, and 55 GeV, respectively. If any of these thresholds is reached the second level is passed, yielding a total rejection of a factor 50 with respect to the L1 rate.

---

3) To distinguish trigger level in the CMS HLT architecture is merely a convention to identify the different steps.

It is instructive to compare these jet thresholds to those of the inclusive jet HLT stream envisaged in reference [2]: for the first, third, and fourth jets, E_T thresholds of 571 GeV, 209 GeV, and 90 GeV are foreseen, respectively for the first, third, and fourth jet. Expressed in 95 % efficiency parton E_T the thresholds read: 657 GeV, 247 GeV, and 113 GeV.

5.3 HLT - L2.5

At this stage is the tracker information accessed to establish whether the event contains b jets.

The input rate at this level is expected to be still quite high: several 100s of Hz to 1 kHz. Therefore, full, global track reconstruction is deemed too time consuming. Instead, a fast "preselection" based on pixel-only tracks is performed. The b-tag determined here is used to preselect those events that are likely to be tagged in the following level. Thus, a fast rejection is achieved of those events that have clearly no b-content or whose b jets are not (easily) taggable for various reasons: it may have produced too few charged particle tracks, the B-hadron flight distance may be very small, the primary vertex may be wrongly reconstructed, etc.

The fact that b's are generally produced in pairs implies that the signal efficiency may be greatly increased by considering more than one jet in an OR. As a compromise between speed and efficiency, the leading and next-to-leading jet are considered in this implementation.

The parameters of the b-tag algorithm are chosen so that a rejection of a factor 5 is obtained with respect to L2.

5.4 HLT - L3

For the events that pass the initial selection based on pixel tracks the b-tag based on regionally reconstructed tracks is invoked. The reconstruction at this level is guided by the result in the previous level: tracks are reconstructed in regions-of-interest around the jets tagged at L2.5. Thus, generally only one jet needs to be reconstructed at this level.

This step is optimized to achieve a rejection of a factor 4 with respect to the previous level.

The b-tagging in two stages improves the robustness of the algorithm. In case the performance of either of the two b-tagging stages is degraded (i.e. if for example the pixel triplet efficiency becomes too low) the weight of this trigger level can be reduced (at a cost in CPU load).

5.5 Trigger performance: timing

The average time required per event of the b-jets trigger stream as a whole is a function of the time requirements of each of the algorithms (jet reconstruction, pixel-only track reconstruction and regional reconstruction) and of the rejection in each step.

An order-of-magnitude estimate can be obtained by the following example: assume a L1 background rate of 100 kHz, and rejections (R_N) and CPU time t_N for each of the algorithms L2, L2.5, L3 described previously. Then, the following total CPU time required each second is given by:

$$t_{TOT} = 100kHz \cdot (t_2 + \frac{t_{2.5}}{R_2} + \frac{t_3}{R_2 R_{2.5}}) \tag{1}$$

With a rejection of 50 in the jet requirement (L2) and of 5 in the pixel-only b-tagging step (L2.5), it is clear that the timing constraint on the L3 reconstruction can be relaxed.

5.6 Background rejection

The background rejection of the inclusive b-jets trigger - implemented as described previously - is studied on a large sample of unbiased di-jet events. The response to the generated events is simulated in full detail using the CMS GEANT4-based simulation software [8]. The hits are merged with those from the minimum bias pile-up corresponding to a luminosity of $2 \cdot 10^{33} cm^{-2} s^{-1}$. Finally, the events are digitized and reconstructed using the reconstruction software [4].

To reduce the number of events that needs to be simulated the QCD sample is split in a large number of $\hat{p}_T$ bins [4].

---

[4] $\hat{p}_T$ represents the momentum transfer in the principal hard interaction

For the lowest bins - up to $\hat{p}_T < 50$ GeV/c not enough statistics is available to study the trigger rates - even after the first stages no events remain. Given that the jet transverse energy thresholds are much higher than the transverse momentum of the leading parton, it is expected that these soft events do not contribute significantly to the output rate of the L2 stage.

For extremely hard scatters - $\hat{p}_T > 380$ GeV/c - the production rate becomes small compared to the HLT output rate. These bins can therefore safely be ignored: even if the trigger efficiency were 100% they would not contribute significantly to the total rate.

Table 2: Background rates for low-luminosity (L = $2 \cdot 10^{33} cm^{-2} s^{-1}$) running at the LHC. The first column shows the production rate for each $\hat{p}_T$ bin. In the second column, the L1 output rate is shown. The output rate of the HLT b jet and inclusive jets triggers is given in the third and fourth columns. In the fifth column the rate of events that are only triggered by the b-jets trigger is listed.

$\hat{p}_T$ bin (GeV/c)	*production (Hz)*	*L1 output (Hz)*	*b-jets (Hz)*	*N-jets (Hz)*	*only b-jets (Hz)*
50 - 80	$4 \cdot 10^4$	$2 \cdot 10^3$	< 2	< 2	< 2
80 - 120	$6 \cdot 10^3$	770	1.1	0.0	0.9
120 - 170	$1 \cdot 10^3$	370	3.5	0.6	2.7
170 - 230	200	140	4.3	1.6	2.2
230 - 300	48	44	3.5	1.8	0.9
300 - 380	12.8	12.7	3.8	1.3	0.5
total	$4.9 \cdot 10^4$	$3.3 \cdot 10^3$	16.2	5.3	6.7

From table 2 it is clear that the inclusive b-jets trigger with the set of parameters described above is able to reduce the total background rate to 16 Hz. Many of the events selected by the b-jets trigger - especially in the high $\hat{p}_T$ range - would also have been selected by other triggers. For these very hard scatters even the higher thresholds of the inclusive jets trigger are generally reached.

To study the correlations between the b-jets stream and the remaining triggers, the table in the Trigger DAQ TDR [2] is assumed. For the events with largest $\hat{p}_T$ the b-jets and inclusive jets trigger are extremely correlated. Virtually all events that trigger the b-jets stream are also accepted by the inclusive jets trigger.

If the events that would have been triggered by other streams in the HLT are not considered, i.e. only counting the additional rate due to the b-jets trigger, the results listed in the column labeled *b-jets only* in the table are obtained. For hard di-jet events the fraction of events that are only selected by the b-jets trigger is actually quite small compared to the total rate of the b-jets stream. Thus, even though the total b-jets trigger rate amounts to 16 Hz, the total trigger rate grows by only 7 Hz when the b-jets trigger is added to the remaining streams.

5.7 Signal selection

As a benchmark physics channel hadronically decaying $t\bar{t}$ events are chosen. This channel is considered the most interesting example for a number of reasons:

- the trigger efficiency obtained using the existing HLT table [2], without inclusive b-jets trigger, is very small.
- this topology contains two central, hard b jets from the top decay and is thus relatively easily tagged.
- these events are certainly relevant in the CMS physics program. Previous experiments have made very significant measurements on this channel [1]. Preliminary studies show that the offline isolation of this signal from the background should be feasible also in the CMS detector [11].

The chosen topology is certainly not the only one of relevance. In general, all physics channels that satisfy the above criteria could benefit greatly from an inclusive b-jets trigger along the lines of the one described in this note. The all-hadronic final state, for example, is the one most severely curtailed in the existing HLT trigger table, but the addition of the b-jets trigger will also reinforce the trigger performance for other final states having a high p_T lepton, by adding redundancy.

The trigger efficiency has been studied in detail on the inclusive top pair sample. It is simulated and reconstructed along the same lines as the di-jet sample described above. The events with fully hadronic decays are selected using the Monte Carlo truth of the events: only events where the W's from the $t \rightarrow Wb$ and $\bar{t} \rightarrow W\bar{b}$ processes decay to quarks are considered.

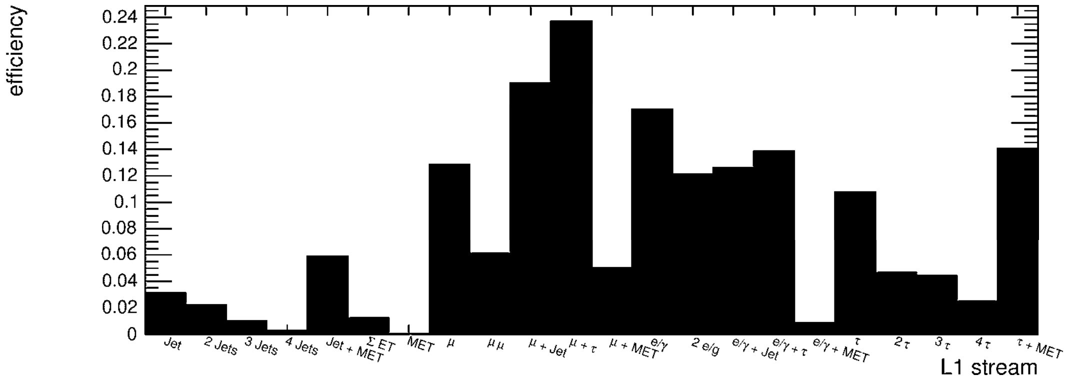

Figure 6: The efficiency of several L1 streams for hadronically decaying $t\bar{t}$-pairs. The streams correspond to those defined in the trigger DAQ TDR [2].

In figure 6 the efficiency to select these signal events at the first level of the trigger (L1) are listed for several L1 streams. The efficiency in the jets streams is artificially low as the thresholds are applied to the raw uncalibrated transverse energy of the jet. It is clear that most events are triggered by the muon + central jet or τ-jet and electron/photon + central or τ-jet. Even the single lepton streams yield signal efficiencies of 12 % and 17 %. A fraction of the events triggered by the electron photon streams certainly correspond to misidentified jets, and probably there is a small fraction of fake muons among the events triggered by the muon streams. The majority of these events are triggered by true leptons from the semi-leptonic decay of b-hadrons.

The total efficiency at the first (hardware) trigger level for fully hadronic $t\bar{t}$ events is 57 %.

In figure 7 the efficiency of the High Level Trigger for fully hadronic $t\bar{t}$ events is shown. The efficiency is normalized with respect to the events that pass the L1 trigger. Only the most relevant streams in the trigger-DAQ Technical Design Report [2] are shown. The photon HLT stream triggers in nearly 10 % of the events. The 4-jet trigger is efficient in 7.5 % of events. Otherwise, the efficiency is very low.

In the two rightmost bins in figure 7 the signal efficiency due to the first and second jets (in an E_T ordered list) considered in the b-jets trigger is shown. Each of these jets contributes an efficiency of approximately 15 %. Combining the efficiencies from both jets (i.e. combining both bits in OR) the signal efficiency amounts to 26 % of all events that passed the first level. These numbers show that the correlation between the results from the first and second jet is rather limited. Often, the leading jet (remember there are at least six hard partons in these events) is not a b-jet , or is not tagged as such.

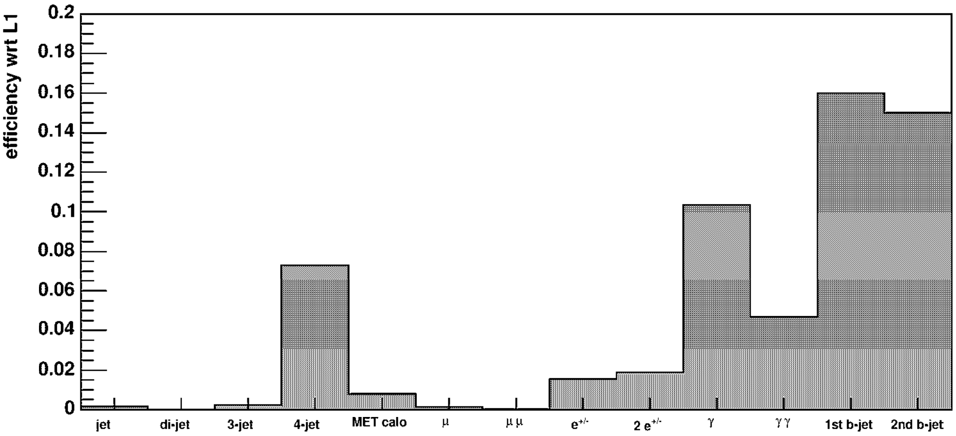

Figure 7: The efficiency of several HLT streams for hadronically decaying $t\bar{t}$-pairs. The streams correspond to those defined in the trigger DAQ TDR [2]. The efficiency is normalized to all events passing the L1 trigger.

The total trigger efficiency (L1 and HLT) of the b-jets stream for fully hadronic $t\bar{t}$ events is 15 %. An analysis of the offline potential of the events triggered by b-jets trigger is currently ongoing [11].

5.8 Summary

An inclusive b-jets High Level Trigger has been implemented in the CMS reconstruction software [4]. In the first stage events with large jet activity are selected. Thresholds of 350 GeV, 150 GeV and 55 GeV are set on the transverse energy of the first, third, and fourth jet. Then the 3D track counting [10] b-tagging discriminant based on pixel-only tracks is calculated for the leading and next-to-leading jet. If any of these jets satisfies the b-tag requirement, Kalman filter track reconstruction is performed in a region around this jet.

This implementation is found to reduce the background effectively. A total rate of 16 Hz is estimated. More than half of these events (9 Hz) triggers at least one other HLT stream listed in the DAQ TDR [2] table. Only 7 Hz is uniquely triggered by the b-jets stream.

The introduction of the b-jets trigger would yield a trigger efficiency of 15 % for hadronically decaying top pairs.

6 Conclusions

Several algorithms that provide fast and high performance track reconstruction exist in the CMS reconstruction framework. In this note, the b-tagging potential of two of these has been investigated.

Regional track reconstruction yields virtually offline-quality tracks and b-tagging. The CPU time required, however, does not allow a jet to be reconstructed until the latest stages of the high level trigger.

Pixel-only reconstruction provides a simpler, much faster algorithm that relies on triplets of hits reconstructed in the three pixel layers. The algorithm provides excellent efficiency and fake rate. The impact parameter resolution is significantly degraded in the transverse plane (with respect to offline-quality tracks), but considerable tagging power remains.

The $d_0\phi$ algorithm based on tracks reconstructed in the tracker, and can provide a fast and precise beam spot reconstruction required at start-of-fill.

The z-coordinate along the beam line of the primary vertex needs to be reconstructed on an event-by-event basis if the 3D impact parameter is to be used. The divisive algorithm based on pixel hit-triplets [5] provides an adequate efficiency and resolution for this purpose.

An inclusive b-jets trigger has been implemented in the CMS reconstruction software. It accepts events from a large number of L1 streams. The HLT decision is reached in three steps. In the first step the calorimeter information is accessed to reconstruct jets. Thresholds on the calibrated jet transverse energy of 350 GeV, 150 GeV and 55 GeV are applied to the first, third, and fourth jet. In the consecutive step, pixel-only tracks are reconstructed in the leading and next-to-leading jets. If at least one of these jets is tagged, tracks are reconstructed regionally. The final b-tag is calculated on the basis of this more precise information.

The total background rate of the inclusive b-jets trigger is estimated to be 16 Hz. Over half of these events (9 Hz) are also triggered at least one of the other envisaged HLT streams: only 7 Hz is added by the introduction of the b-jets trigger.

The efficiency for fully hadronic final states may be significantly increased by an inclusive b-jets trigger. The implementation described in this note yields an efficiency of 15 % for the benchmark channel considered here - fully hadronic $t\bar{t}$ events.

The authors would like to thank the Marcin Konecki and Susanna Cucciarelli for their help on the pixel-only reconstruction package and Andrea Rizzi and Christian Weiser for their willingness to adapt the b-tagging framework.

References

[1] F. Abe et al., First observation of the all hadronic decay of $t\bar{t}$ pairs, Phys. Rev. Lett. 79 (1997) 1992.

D. Wicke for the CDF and D0 collaborations, Top pair production cross-section measurement in the all-hadronic channel at CDF and D0, Int. J. Mod. Phys. A20 (2005) 3183.

[2] The CMS collaboration. The Trigger and Data Acquisition project, volume II: Data Acquisition & High-Level Trigger, Technical Design Report, CERN/LHCC 2002-26

[3] W.Adam et al., Track reconstruction in the CMS tracker, CMS note in preparation

[4] The object-oriented reconstruction framework. Version 8.7.3 is used for all results in this note. Further information can be found under this URL: http://cmsdoc.cern.ch/orca

[5] S. Cucciarelli, M. Konecki, D. Kotliński, T. Todorov, Track parameter evaluation and primary vertex finding with the pixel detector, CMS Note 2003/026

[6] S. Cucciarelli, D. Kotlinski, T. Todorov, Position determination of the pixel hits, CMS note 2002/049

[7] S. Cucciarelli et al., Seed generation, Track Finding and Vertex Reconstruction with the Pixel Detector, CMS note in preparation

[8] The GEANT4-based full simulation package for the CMS detector. Version 2.4.5 is used. Further information can be found under this URL: http://cmsdoc.cern.ch/oscar

[9] L. Barbone et al., Impact of Silicon Tracker Misalignment on Track and Vertex Reconstruction, CMS note in preparation

[10] A. Rizzi, F. Palla and G. Segneri, Track impact parameter based b-tagging, CMS note in preparation

[11] C. Ciocca and A. Castro (supervisor), Study of $t\bar{t}$ production in multijet events at LHC with CMS, Ph.D. thesis, Bologna University.