

Hadron Collider Triggers With High-Quality Tracking at Very High Event Rates

Alberto Annovi, Antonio Bardi, Mario Campanelli, Roberto Carosi, Pierluigi Catastini, Vincenzo Cavasinni, Alessandro Cerri, Allan Clark, Mauro Dell'Orso, Tarcisio Del Prete, Andrea Dotti, Giampiero Ferri, Stefano Giagu, Paola Giannetti, Giuseppe Iannaccone, Michele La Malfa, Fabio Morsani, Giovanni Punzi, Marco Rescigno, Chiara Roda, Mel Shochet, Franco Spinella, Stefano Torre, Giulio Usai, Laurent Vacavant, Iacopo Vivarelli, Xin Wu, and Lucia Zanello, *Member, IEEE*

Abstract—We propose precise and fast-track reconstruction at hadron collider experiments, for use in online trigger decisions. We describe the features of fast-track (FTK), a highly parallel processor dedicated to the efficient execution of a fast-tracking algorithm. The hardware-dedicated structure optimizes speed and size; these parameters are evaluated for the ATLAS experiment. We discuss some applications of high-quality tracks available to the trigger logic at an early stage, by using the LHC environment as a benchmark. The most interesting application is online selection of b -quarks down to very low transverse momentum, providing interesting hadronic samples: examples are $Z^0 \rightarrow b\bar{b}$, potentially useful for jet calibration, and multi- b final states for supersymmetric Higgs searches. The paper is generated from outside the ATLAS experiment and has not been discussed by the ATLAS collaboration.

Index Terms—Parallel processing, particle tracking, pattern matching, triggering, very large scale integration.

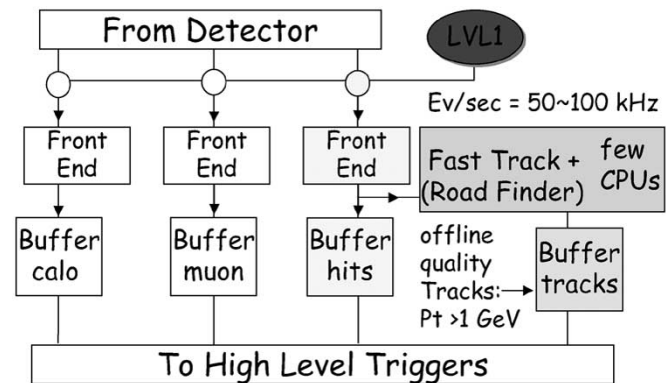


Fig. 1. The FTK processor has access to the tracker data of events selected by the LVL1. The FTK performs substantial data reduction by selecting high P_T track candidates and organizing them with their hits into standard DAQ buffer memories, ready for the high-level triggers (level-2 and level-3).

I. INTRODUCTION

WE PROPOSE the use of a dedicated hardware processor, fast-track [1] (FTK), for online pattern recognition of tracker detector data. FTK is an evolution of the Silicon Vertex Tracker (SVT) [2], the tracker now running at the Collider Detector at Fermi Lab (CDF) experiment. FTK is a powerful processor that, in combination with a few standard CPUs, reconstructs high-quality tracks for all detector-fiducial particles of transverse momentum (P_T) above 1 GeV or even less. This work can be performed at the very high event rates accepted by the level-1 trigger (LVL1), i.e., up to 50–100 kHz.

Fig. 1 shows how FTK could be integrated in a modern LHC [3] data acquisition (DAQ) system. The trigger selection is organized into a multilevel architecture. The LVL1 provides a first-rate reduction to 50–100 kHz. The level-2 and level-3 selections (high-level triggers) reduce the rate to a few hundred hertz of events to be written on tape. Tracking data are collected at the LVL1 rate in the front end, then they are stored into large memory buffers. These buffers are interfaced to a large CPU farm for higher level triggers. FTK, without interfering with the operation of the DAQ system, “sniffs” the tracker data flowing to the memory buffers and filters out interesting high-quality tracks, storing complete information in an additional memory buffer that CPUs can access at high rate.

The advantage of this kind of implementation is a dedicated high-input bandwidth for FTK. The proposed system is almost totally independent and its interference with the DAQ is min-

Manuscript received June 8, 2003; revised December 9, 2003.

A. Annovi, P. Catastini, V. Cavasinni, M. Dell'Orso, A. Dotti, G. Ferri, M. La Malfa, C. Roda, and G. Usai are with the Dipartimento di Fisica, Università di Pisa, 56100 Pisa, Italy (e-mail: alberto.annovi@pi.infn.it; pierluigi.catastini@pi.infn.it; vincenzo.cavasinni@pi.infn.it; mauro.dellorso@pi.infn.it; andrea.dotti@pi.infn.it; giampiero.ferri@pi.infn.it; michele.lamalfa@pi.infn.it; chiara.roda@pi.infn.it; giulio.usai@pi.infn.it).

A. Bardi, R. Carosi, T. Del Prete, P. Giannetti, F. Morsani, and F. Spinella are with the Istituto Nazionale di Fisica Nucleare, Pisa, Italy (e-mail: antonio.bardi@pi.infn.it; roberto.carosi@pi.infn.it; tarcisio.delprete@pi.infn.it; paola.giannetti@pi.infn.it; fabio.morsani@pi.infn.it; franco.spinella@pi.infn.it).

M. Campanelli, A. Clark, and X. Wu are with the Departement de Physique Nucleaire et Corpusculaire, Université de Genève, CH-1211 Genève, Switzerland (e-mail: mario.campanelli@mail.cern.ch; allan.clark@mail.cern.ch; xin.wu@mail.cern.ch).

A. Cerri and L. Vacavant are with the Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA (e-mail: acerri@lbl.gov; lvacavant@lbl.gov).

S. Giagu and L. Zanello are with the Dipartimento di Fisica, La Sapienza, 00185 Rome, Italy (e-mail: stefano.giagu@roma1.infn.it; lucia.zanello@roma1.infn.it).

G. Iannaccone is with Dipartimento di Ingegneria dell'Informazione, Università di Pisa, I-56126 Pisa, Italy (e-mail: g.iannaccone@iet.unipi.it).

G. Punzi and I. Vivarelli are with Scuola Normale, 56126 Pisa, Italy (e-mail: giovanni.punzi@pi.infn.it; iacopo.vivarelli@pi.infn.it).

M. Rescigno is with Istituto Nazionale di Fisica Nucleare, 00185 Rome, Italy (e-mail: marco.rescigno@roma1.infn.it).

M. Shochet is with the Enrico Fermi Institute and the Department of Physics, University of Chicago, Chicago, IL 60637 USA (e-mail: shochet@hep.uchicago.edu).

S. Torre is with the Istituto Nazionale di Fisica Nucleare, Pisa, Italy, and also with the Dipartimento di Fisica, Università di Siena, 53100 Siena, Italy (e-mail: stefano.torre@pi.infn.it).

Digital Object Identifier 10.1109/TNS.2004.828639




Fig. 2. Cross-sectional view (extracted from [5, Fig. 1-1]) through the beam axis of one quarter of the ATLAS SCT and pixel detectors. The most forward SCT disks (7-9) are not used in this proposal and not shown in the figure. Each patterned line corresponds to a different logical layer.

imal. It can be added even after the baseline has been built, as an upgrade, if the possibility of adding a bypass to spy on the events is included in the DAQ from the beginning. Technical details of the proposed system can be found elsewhere [1].

In Section II, we describe the structure and size of a possible application to a real experiment and its performance.

It should be emphasized that the algorithm for every particle used in the level-2 selection (level-2 object) could take advantage of high-quality full track reconstruction. In order to show the potential of FTK, we discuss in Sections III and IV some possible level-2 applications to the LHC environment.

II. A REALISTIC FTK APPLICATION

We define the structure of the proposed system in a realistic case to show that the insertion of the hardware dedicated processor FTK into hadron collider experiments is not very complex. We choose the ATLAS experiment [4] for this exercise.

A. The Data Flow to FTK

The silicon detectors are necessary for a good impact parameter measurement, and they should be used for the FTK on-line-pattern recognition. Since the quality of the tracking obtained only with silicon detectors is very good (the best impact parameter resolution is $17 \mu\text{m}$ in [6] to be compared to $12 \mu\text{m}$, the best value that can be obtained if the whole tracking detector is exploited), we can plan, at least at the beginning, to use only these detectors. Fig. 2 shows a cross section of the engineering layout through the beam axis of one quarter of the silicon detectors [except the most forward semiconductor tracker (SCT) disks].

FTK can perform pattern recognition on up to 12 logically defined tracking layers [1]. A possible layer definition is shown in Fig. 2. This is an example that uses only six layers (4 R- ϕ SCTs and 2 Pixels) and could work at the very beginning, when a pixel layer will be missing (it can be easily added to FTK later). FTK is a scalable processor. A small system, appropriate for the initial luminosity running of the LHC, would serve as an R&D prototype for a full high luminosity system. More layers can be defined using also the SCT stereo layers, for a total of 11 layers (8 SCT + 3 pixels). In the simplified example of Fig. 2, barrel and disk layers are linked together to guarantee full polar angle

TABLE I
EVENT SIZE FOR DIFFERENT DETECTOR LAYERS. CONSERVATIVE ESTIMATE OF THE AVERAGE NUMBER OF CLUSTERS PER EVENT PER DETECTOR LAYER FOR THE LOW LUMINOSITY OCCUPANCY. CLUSTER RATES FOR DISKS REFER TO BOTH END CAPS

Pixel Layer	# Cluster /Event	SCT Layer	# Cluster /Event
Barrel 0	1286	Barrel 0	635
Barrel 1	854	Barrel 1	652
Barrel 2	902	Barrel 2	655
Disk 1	216	Barrel 3	665
Disk 2	216	Disk 1	430
Disk 3	216	Disk 2	323
Disk 4	108	Disk 3/4/5/6	430

coverage; detectors with the same patterned lines are combined to build a complete layer. Each patterned line corresponds to a single layer for a total of six.

Is it possible to make such a large amount of data available to FTK in time? The input bandwidth sets an upper limit either on the event rate or on the size of the detector connected to it. The FTK processor increases the data flow rate with respect SVT at CDF by exploiting the parallel readout of the detector layers on six buses. Nevertheless, in order to sustain very high event rates, it is necessary to organize FTK as a set of independent engines (input bandwidth of 4 Gbit/s), each working on a different sector of the silicon tracker. Let us imagine dividing the detector into azimuthal (ϕ) sectors. This segmentation generates some inefficiency at sector boundaries that can be removed by allowing a small overlap region at the boundaries.

We perform a conservative estimate of the cluster rates for each patterned layer to determine the appropriate number of sectors for ATLAS. It is assumed that two strips, or two pixels belong, on average, to a cluster. We evaluate first the average number of clusters per event in each detector layer. We use the occupancies described in [5] to evaluate the contribution due to a single minimum bias event. We calculate the correct occupancy at low and high luminosity using the contribution of 5 and 25 minimum bias events, respectively, and three other events to add a conservative estimate of the hard interaction contribution ([5, Figs. 2-6] shows that this is a very conservative estimate). For the SCT disks we consider the worst-case barrel occupancy. On top of this we add a noise contribution [5]: 10^{-5} for the pixels and 10^{-4} for the SCT. Table I shows the results for low luminosity (five minimum bias events). Table II shows the cluster rates for the logical layers defined by the patterned lines in Fig. 2. We use the rates expected at low luminosity with a reduced level-1 output rate (50 kHz) and the rates expected at high luminosity to define the minimum and maximum numbers of the FTK processors or ϕ sectors.

Since each logical layer is loaded in parallel at a frequency of 40 MHz, at the beginning we can segment the ATLAS detector in 2ϕ sectors. At high luminosity, we need 8ϕ sectors. These numbers are appropriate for a six-layer configuration. The addition of extra layers is possible but they would be received on the same six buses, so the cluster rates of Table II will be higher. Should the number of sectors be considered too high, it

TABLE II
EVENT SIZE FOR DIFFERENT DETECTOR LAYERS. CONSERVATIVE ESTIMATE OF THE CLUSTER RATES PER LOGICAL LAYER FOR LOW AND HIGH LUMINOSITY. THE LEVEL-1 OUTPUT RATE IS CHOSEN TO BE 50 AND 75 kHz RESPECTIVELY FOR LOW AND HIGH LUMINOSITY

Logical Layer	Cluster Rate (MHz) 50 kHz LVL1 rate $L=2 \times 10^{33} \text{ cm}^{-2} \text{ sec}^{-1}$	Cluster Rate (MHz) 75 kHz LVL1 rate $L=10^{34} \text{ cm}^{-2} \text{ sec}^{-1}$
Layer 0	64	308
Layer 1	61	202
Layer 2	50	225
Layer 3	65	297
Layer 4	61	296
Layer 5	64	305

is possible to transfer to FTK only a subset of the LVL1 triggers. For example, only multijet and jet + lepton triggers could be analyzed to search for b-jets, interesting for low-mass Higgs physics.

A single 9U VME crate would contain the engine core for a detector sector, as described in [1]. The core size is dominated by the large *bank* of pre-calculated hit *patterns* (*pattern bank*) used to perform pattern recognition [1]. The *pattern bank* size has to be calculated to confirm the FTK size for the ATLAS experiment (see Section II-C and D). The FTK core, in conclusion, should grow from an initial size of 2 VME crate to a final dimension of 8 crates. The FTK connection to the detector is excluded in this computation. In ATLAS the whole silicon tracking data are collected by roughly 180 devices, (RODs), working in parallel. Each ROD should provide an output to FTK. Design is under way of a modified S-link output mezzanine board that provides a second copy of the track data being sent from the RODs to the buffer memories (ROBs). If the detector is divided into 2 (8) sectors, each processor would receive 90 (23) links, find clusters and organize them into the 24 inputs received by the processor core. In fact there are four inputs [1] for each one of the six FTK input buses.

B. The FTK and Detector Simulation

The performance of the system has been studied using only a part of the silicon ATLAS detector: the ATLAS central detector [5] (barrel). Only seven cylindrical layers are used to find tracks: three pixel layers linked to four R- ϕ silicon layers. A stand-alone simulation program has been used to generate tracks in the ATLAS detector. It takes into account effects such as multiple scattering, ionization energy loss, detector noise, detector inefficiencies, and resolution smearing. Even if the chosen detector layout is different from the one described in Section II-A, the simulation results are important to show the FTK capability and size.

C. Pattern Bank Size

We estimate the bank size for the barrel of the ATLAS experiment to show that the necessary system size is modest. We generate tracks in the whole detector (no detector symmetries are exploited, to prevent alignment problems) and we store new patterns corresponding to the generated tracks, until the bank

TABLE III
DEFINITION OF ROAD SIZES. THREE OPTIONS CONSIDERED FOR THE ROAD WIDTHS

Road Type	Pixel Detector Granularity (R ϕ x z)	SCT Detector Granularity (R ϕ x z)
Large	5 mm x 6.5 cm	10 mm x 12.5 cm
Medium	2 mm x 6.5 cm	5 mm x 12.5 cm
Thin	1 mm x 6.5 cm	3 mm x 12.5 cm

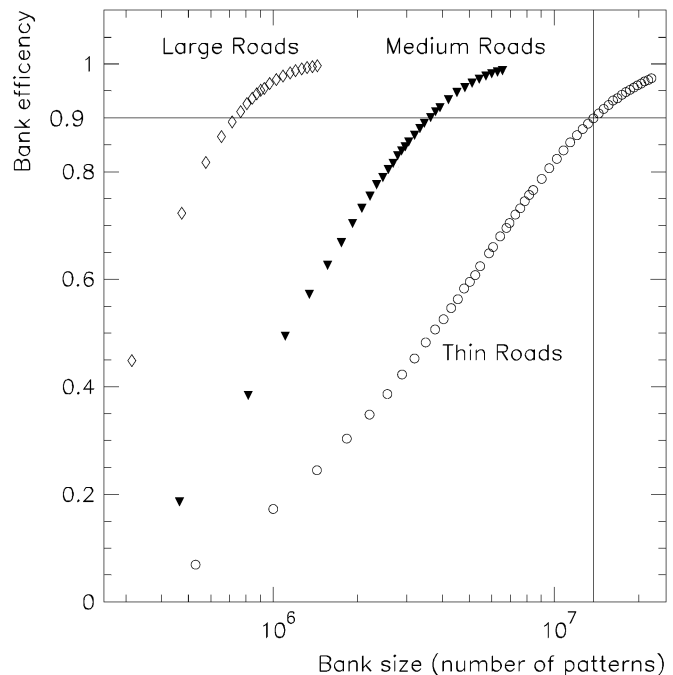


Fig. 3. Bank efficiency as a function of the bank size for P_T thresholds of 1 GeV/c and for different sizes of roads: large (open diamonds), medium (full triangles), and thin (circles). The bank size is for 1/4 of the whole barrel.

reaches the desired efficiency. A reference bank efficiency has been conventionally fixed at 90%.

The generated track typology also affects the bank size. It is convenient to restrict the bank to include only tracks we care about.

For this purpose we limit the region where the tracks come from (luminosity region) to those values relevant for the physical processes to be studied. We assume a cylindrical luminosity region, circular on the transverse plane with a radius of 1 mm and ± 15 cm long in the longitudinal direction. This restriction is compatible with B-meson decay products, whose impact parameters are a few hundred microns. We include all tracks with $P_T > 1$ GeV. The road size is another critical parameter for both the processor performance and the pattern bank size. We use different values for Silicon and Pixel detectors. Three options, described in Table III, have been considered for the road widths. The size of the precomputed pattern bank has been evaluated for these options of road size. Fig. 3 shows the bank efficiency versus the bank size for the three road options.

In order to choose the optimal pattern bank, one must also evaluate the residual computing power needed at the LVL2 to complete the full resolution tracking.




Fig. 4. Left plot: average number of hit combinations per road as a function of the silicon detector road size. The different curves are for various QCD jet parton thresholds in low luminosity running. The third point of the curve $P_T > 200$ GeV corresponds to 1400 combinations/road. Right plot: average ratio of the number of matching roads to the number of real tracks as a function of the silicon detector road size.

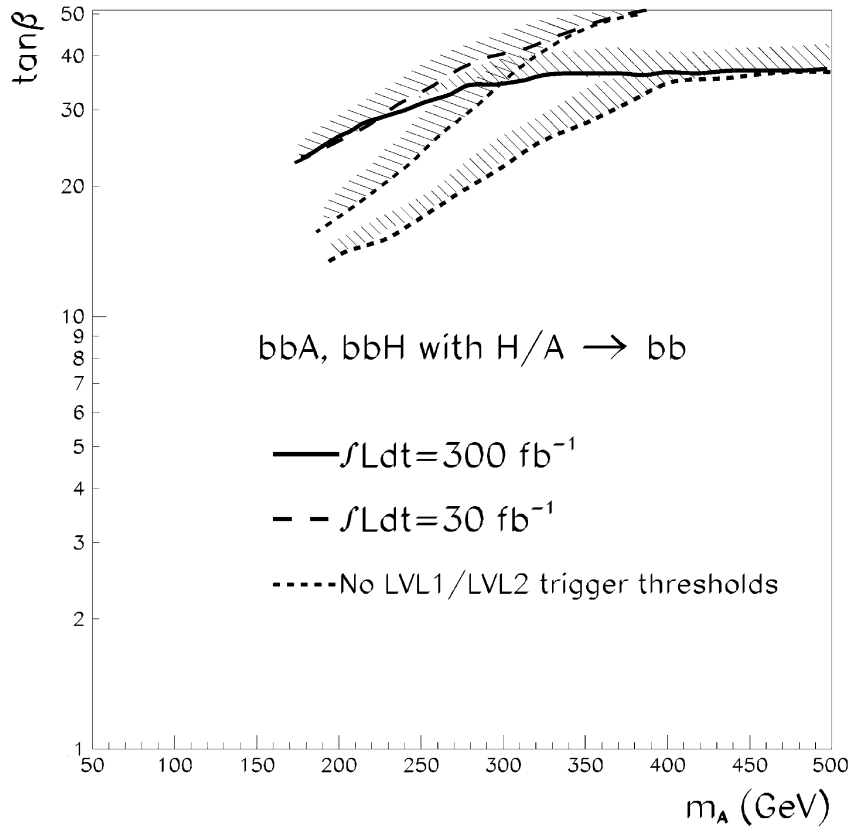


Fig. 5. In [12, Fig. 19–69]. 5σ -discovery contour curves for the processes $b\bar{b}H/A \rightarrow 4b$, in the $(m_A, \tan\beta)$ plane. The integrated luminosities are 30 fb^{-1} (dashed line) and 300 fb^{-1} (solid line). The dotted lines show the expected extension in coverage if the level-1 and level-2 trigger thresholds are not applied.

D. Finding Tracks at Full Resolution

The simulation program performs two subsequent steps to reconstruct the event, in order to reproduce the behavior of the hardware procedure (see [1] for details):

- coarse track (road) finding: all roads are found by simulating the FTK processor;
- track fitting: all found roads are processed to find the best track parameter values (high quality [6], [7]) and to re-

ject the fake ones. This is achieved using linear approximations for the track constraints and principal component analysis [8].

Because of the finite size, a road may contain physical hits belonging to different particles. Also, depending on the road size and on the event hit density, there is a level of combinatorial background consisting of fake roads, i.e., track candidates that will be rejected at full resolution. Therefore, the complexity of the track fitting strictly depends on the following quantities:

TABLE IV

FIT RATES. NUMBER OF FITS NECESSARY TO RECONSTRUCT DIFFERENT QCD EVENTS. THE TABLE COLUMNS SHOW, IN ORDER, QCD JET ENERGY, NUMBER OF FITS PER REAL TRACK, NUMBER OF REAL TRACKS PER EVENT, LEVEL-1 RATE BUDGET THAT COULD BE ASSIGNED TO THE SAMPLE AT A LUMINOSITY $L = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$, AND THE FIT RATE REQUIRED FOR EACH SAMPLE

Event Type	# fit /track	# tracks /event	# fit /event	Level-1 rate (kHz)	Fit rate (MHz)
QCD P_T 200	658	17	11186	0.2	2.2
QCD P_T 100	74	16	1184	2	3
QCD P_T 40	14	10	140	5	0.75
QCD P_T 10	6	8	48	20	1.5

- $\langle \text{roads}/\text{track} \rangle$: the average ratio between the number of matching roads and the number of actual tracks per event;
- $\langle \text{combinations}/\text{road} \rangle$: the average number of hit combinations per road.

Fig. 4 shows these critical numbers as functions of the road size for the low luminosity run (5 minimum bias events are used as pile-up). The P_T threshold for tracks is fixed at 1 GeV. Results are reported for different QCD samples generated with increasing P_T thresholds of the outgoing partons: $P_T = 10, 40, 100$, and 200 GeV. Detector hits are produced from tracks in events generated with Pythia [9]. The event primary vertex is smeared using a Gaussian distribution of $\sigma = 5.6 \text{ cm}$ along z . We allow for one missing point, requiring six fired layers out of seven to find a track, thus reducing the effect of detector inefficiency.

We observe that the events with more energetic jets are characterized by larger combinatorics and a larger number of fakes, as is expected because of the jet fragmentation. We also observe that the thinnest road minimizes the differences between the samples. We choose this road size.

The average number of fits per real track (calculated as the product of the two quantities plotted in Fig. 4) is given in Table IV for the different QCD samples, together with conservative estimates of the level-1 bandwidths for those events and the corresponding fit rate. The number of fits to be executed is strongly dependent on the jet fragmentation. For the LVL1 triggers reported in Table IV, the total level-2 fit rate is 8 MHz. Since a Pentium III with an 800-MHz CPU can execute the linear fits at a rate of 1 MHz, eight such CPUs are enough to reconstruct these critical events. We plan to fit these CPUs in the FTK crates. If the detector is divided into two sectors, we can assign four CPUs to each sector. Since the fits will be distributed to the eight CPUs, even the most energetic QCD sample has an average latency of 1–2 ms.

In conclusion, a few level-2 processors can reconstruct the input events.

The bank size for an FTK processor working on half of the ATLAS barrel amounts to 30×10^6 patterns (see Fig. 3) assuming (a) a conservative choice of the thinnest roads, to safely handle the worst conditions, (b) a P_T threshold of 1 GeV/c, (c) two FTK processors working in parallel on two detector ϕ sectors. Taking into account the pattern densities estimated in [1] for the year 2005 for ASIC implementation (5×10^6 patterns per board), we conclude that such a bank will fit in six slots of a VME crate.

It is interesting to compare these results with the performances reported in [1] for a similar FTK system applied to similar experimental conditions. The study reported in [1] used overly conservative assumptions, producing a much too powerful road finder, that performs the whole pattern recognition job by itself (very few track fits have to be performed by CPUs). The more realistic estimates in this paper allow a better balance between the road finder work and the CPU work. As a result, with similar dedicated hardware and a larger number of CPUs the pattern recognition can be applied to tracks with a minimum track P_T down to 1 GeV, which is a clear improvement compared to the threshold of 2 GeV of [1]. This comparison shows the importance of system parameter tuning.

Summarizing, a complete processor for one half barrel is estimated to fit in a 9 U VME crate. For details of the type and number of boards see [1]. More studies are necessary to evaluate the bank size for a complete detector including the silicon disks.

III. B-TAGGING JETS IN HIGH P_T EVENTS

One of the major trigger applications for tracks is secondary vertex reconstruction for heavy flavor identification. Since we expect many of the new and exotic particles to be strongly coupled to heavy flavors, the possibility of selecting b-quark events (for both high- P_T physics and low- P_T B-physics) is very desirable. CDF has shown that it is possible to trigger on tracks at both level-1 and level-2. Full resolution and impact parameter selection are provided at level-2. Successful results of the CDF trigger based on on-line impact parameter reconstruction include the study of large samples of D and B hadronic decays and the selection of $Z^0 \rightarrow b\bar{b}$ hadronic events. It is possible to implement similar triggers at the LHC experiments. Preliminary rates and efficiencies have been calculated using the ATLAS software, but the same ideas could be part of the CMS trigger program [10]. In [1], the possible FTK insertion in the CMS DAQ is discussed.

We propose new trigger strategies, based on the FTK potential, to collect samples rich in low- P_T b-quarks. We consider two physics cases that will take advantage of these new triggers. First, we analyze a high cross-section process $Z^0 \rightarrow b\bar{b}$. Then, we consider a rare process $b\bar{b}H/A$. Finally, we compare our b-jet selection proposal to the CMS trigger program [10].

A. New Trigger Strategies for b-Quark Events

In order to set reference goals for the new triggers, we consider as acceptable the following output rates.

- A few kilohertz for level-1 selections. These rates are not too high, since FTK will produce high quality track information for these events allowing a minimum level-2 computing power to reduce the rates as necessary;
- A few hundred hertz for level-2 selections.
- A few hertz for the level-3 selections.

At this stage, the precise values of these rates are not critical. To estimate trigger rates we use Pythia [9] and ATLASfast [11]. Actual trigger algorithms and threshold turn-ons were not simulated, so the rate estimates will be refined in the future with a detailed trigger simulation.

Jets are localized energy depositions in the calorimeters which are constructed in a tower geometry. The jet energy E and momentum (P_x , P_y , P_z) are the scalar and vector sums, respectively, of calorimeter tower energies inside a cone of radius $R^2 = \Delta\eta^2 + \Delta\phi^2$ centered in the jet direction. ϕ is the azimuthal angle in the plane perpendicular to the beam, and the pseudorapidity is defined as $\eta = -\ln \tan(\theta/2)$, where θ is the polar angle. The jet clustering uses a cone radius $R = 0.4$, with appropriate P_T -dependent jet corrections applied at each level. More specific corrections are needed because jets produced by b-quarks are rich in soft leptons. While the electron is well measured by the calorimeter, a large fraction of the muon (μ) energy is lost: in ATLASfast the reconstructed μ 4-momentum is added to the jet 4-momentum, if the μ is not isolated and its distance from the jet axis is smaller than $R = 0.2$. A b-jet is defined as a jet containing a b-quark within a cone of radius $R = 0.2$ with respect to the jet axis.

In order to evaluate the contributions to the rate from mistagging and from b-tagging inefficiency, we simulate generic QCD (with all b production mechanisms included). We use two sets of values for the tagging parameters: (a) b-tagging efficiency of 100% and a mistagging probability of 0% (*perfect tagging*) and (b) b-tagging efficiency of 60% and a mistagging probability of 1% for u,d-quarks and of 10% for c-quarks (*realistic tagging*). We define two different level-1 selections for two different categories of events: $\mu 6 + \text{soft jets}$ and *all jet*.

1) *Level-1: $\mu 6 + \text{Soft Jets}$* : High cross-section $b\bar{b}$ processes in which the b-jet E_T are not very large present problems at level-1. A good example of such a process is $Z^0 \rightarrow b\bar{b}$. At least two jets must be within the tracking acceptance region ($|\eta| < 2.5$) in order to be tagged as b-jets at level-2. However, a two jet requirement at low P_T would give a level-1 rate that is much too high. We thus add a soft-lepton trigger requirement as follows:

- μ with $P_T > 6$ GeV and $|\eta| < 2.5$;
- soft jet with $P_T > 25$ GeV (i.e., sharp cut at 25 GeV) and $|\eta| < 2.5$;
- second soft jet with $P_T > 10$ GeV and $|\eta| < 2.5$.

2) *Level-1: All Jet*: Events with more than 2 jets in the final state can exploit a pure calorimetric level-1 selection. A good example is final states with four b-jets, making b-tagging at level 2 very efficient. The ability to significantly reduce the trigger rate at level 2 makes an all jet level-1 trigger with moderate jet P_T thresholds quite plausible. We require at least three soft jets and a total transverse energy in the event (ΣE_T) above a certain threshold. We consider two alternative sets of P_T thresholds. The (a) set of cuts is as follows:

- jet with $P_T > 50$ GeV and $|\eta| < 2.5$;
- jet with $P_T > 20$ GeV and $|\eta| < 2.5$;
- jet with $P_T > 15$ GeV and $|\eta| < 2.5$;
- $\Sigma E_T > 150$ GeV.

The (b) set of cuts is as follows:

- jet with $P_T > 70$ GeV and $|\eta| < 2.5$;
- jet with $P_T > 50$ GeV and $|\eta| < 2.5$;
- jet with $P_T > 15$ GeV and $|\eta| < 2.5$;
- $\Sigma E_T > 200$ GeV.

3) *Level-1: Rates*: We used generic QCD events generated with Pythia ($P_T > 10$ GeV for the outgoing partons, $\sigma =$

TABLE V
LEVEL-1 SELECTIONS AND RATES

Cut (GeV)	$\sigma_{\text{QCD}}(\text{nb})$	R(kHz)	$\varepsilon_{\text{H/Abb}}(\%)$	$\sigma_{\text{H/Abb}}(\text{pb})$
$\mu 6 + \text{soft jets}$	1290	2.6	18.8	20.8
all jet-a	7500	15	44.5	49.6
all jet-b	2000	4	24.6	27.1

6.53 mb) to study the $\mu 6 + \text{soft jets}$ selection. To study the *all jet* selection we generated a QCD sample characterized by higher P_T threshold (MSEL = 1, $P_T > 20$ GeV, $\sigma = 0.64$ mb).

Table V summarizes the rates. For each case, the table shows the selection requirements, the QCD cross section σ_{QCD} and the rate R for an instantaneous luminosity of $L = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$; the efficiency $\varepsilon_{\text{H/Abb}}$ and cross section $\sigma_{\text{H/Abb}}$ are shown for SUSY Higgs ($m_{\text{H/A}} = 200$) events that, having four b-jets in the final state, can exploit all of these level-1 selections. The signal efficiencies are quite large. The rates are underestimated, since the slow trigger threshold turn-ons of muons and jets were not simulated. In fact muons and jets below the nominal threshold (corresponding to the value where the trigger is 95% efficient) contribute significantly, despite their low trigger acceptance. Part of the missing rate is however recovered since we apply jet corrections at level 1. We see that the level 1 rate approximately doubles when the corrected jet energy scale is used in the selection.

However, the final rates of Table V could be even higher if a physics channel becomes so interesting to gain a large fraction of the total level-1 budget (75 kHz). Whatever the level-1 rate is, the use of FTK allows level 2 rates as low as 100 Hz or even few Hz, as shown in the following.

The natural selection criterion at level-2 is the presence of b-jets inside the tracking acceptance. Again we define two different selections for different level-2 categories: $M_{b\bar{b}}^{\text{cut}}$ and *multi-b-jet*.

4) *Level-2: $M_{b\bar{b}}^{\text{cut}}$* : This selection is for events characterized by two b-jets only in the final state. At level-2 we require that both jets are b-tagged. In addition, we cut on the two b-jet mass:

- at least two b-jets with $|\eta| < 2.5$;
- $M_{b\bar{b}} > M_{\text{cut}}$.

5) *Level-2: 3-b*: This selection is for events with more than 2 b-jets in the final state. We require the three highest P_T jets to be tagged as b-jets (3-b leading).

6) *Level2: Rates*: We can execute each level-2 selection on both level-1 triggers. Table VI shows the resulting rates and signal efficiencies. To study the efficiency of this selection on the Z^0 and Higgs signals, we simulate $10^5 Z^0 \rightarrow b\bar{b}$ with Pythia (MSEL = 11, $P_T > 1$ GeV, $\sigma = 7.6$ nb). Higgs signal efficiencies are evaluated with 2×10^5 events ($m_A = 200$ GeV, $\tan \beta = 30$, $\sigma = 111$ pb).

For statistical reasons the b-tagging is taken to be perfect for the 3-b selection (b-tagging efficiency 100% and mistagging probability 0 for all other quarks). This gives a conservative estimate of real trigger rates, since the rate losses due to a realistic b-tagging efficiency (60%) are only partially balanced by the rate increase due to mistagging (1%). We know also that a real b-tagging efficiency reduces the signal efficiency by roughly a factor $1/4 (\sim 0.6^3)$ when requiring three real b-jets.

TABLE VI

LEVEL-2 SELECTIONS AND RATES. QCD AND SIGNAL EVENTS PASSING THE LEVEL-2 TRIGGER REQUIREMENTS. THE TABLE COLUMNS SHOW, IN ORDER, DIFFERENT SELECTION CUTS, THE CROSS SECTION FOR THE SELECTED QCD EVENTS, AND LEVEL-2 OUTPUT RATES FOR AN INSTANTANEOUS LUMINOSITY $L = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$, THE EFFICIENCY FOR SIMULATED SIGNAL EVENTS AND THE CROSS SECTION FOR THE SELECTED SIGNAL EVENTS

Cut (GeV) (level-1)	$\sigma(\text{nb})$	R(Hz)	$\epsilon_{Z^0}(\%)$	$\sigma_{Z^0}(\text{PB})$
realistic b-tagging				
$M_{bb} > 50$ ($\mu 6 + \text{soft-jet}$)	83	165	1.2	87
$M_{bb} > 50$ (all jet-b)	17	35	0.29	22
perfect b-tagging				
Cut (GeV) (level-1)	$\sigma(\text{nb})$	R(Hz)	$\epsilon_{H/Abb}(\%)$	$\sigma_{H/Abb}(\text{pb})$
3-b ($\mu 6 + \text{soft-jet}$)	4.5	9	8.3	8.9
3-b (all jet-a)	9.8	19	27	28
3-b leading (all jet-a)	2.8	5.6	18	19
3-b (all jet-b)	4.4	9	18	19
3-b leading (all jet-b)	0.9	1.7	13	14

B. High Cross-Section $b\bar{b}$ Processes: $Z^0 \rightarrow b\bar{b}$

In many important LHC physics channels, such as Higgs decay into b 's and Standard Model top decay, accurate mass reconstruction requires the energy scale and resolution of b-jets to be well understood. For light quark and gluon jets, there are a number of ways to set the energy scale, e.g., $\gamma + \text{jet}$ and $Z^0 + \text{jet } P_T$ balancing. These scales are modified for b-quark jets largely because of the energy lost to neutrinos and unidentified muons in the large fraction of b-jets containing a primary or secondary semileptonic decay. At CDF, the correction for b-jet response has been derived in the past from simulated events. For precision measurements and M_{bb} resolution improvement, this is less than satisfactory; a direct experimental measurement is preferable. The observation of a known resonance would provide an independent check of the calibration procedure described in [12]. The obvious candidate is the Z^0 decay.

Note that the M_{bb}^{50} trigger accepts one $Z^0 \rightarrow b\bar{b}$ event every few seconds. The significance for an integrated luminosity of 20 fb^{-1} , estimated for the M_{bb}^{70} selection is

$$S = \sigma(Z^0) \times \sqrt{L_{\text{int}}} / \sqrt{\sigma_{\text{QCD}}} = 46, 23$$

respectively, for the $\mu 6 + \text{soft jets}$ and *all jet* samples.

If we look for the Z^0 signal in the optimized window $80 < M_{bb} < 100$, the significance is $\sim 60(30)$, with a signal-to-noise ratio $(S/N) \sim 1/250(1/370)$ for the $\mu 6 + \text{soft jets}$ (*all jet*) sample.

Of course, the level-2 output rate must be reduced further before recording standard size events. The FTK processor allows an interesting possibility: the ability to store large numbers of selected high rate calibration triggers by writing *mini-events* with FTK tracks but without all of the full raw tracker data. If level-2 b-tagging is of high quality, then the double-tagged dijet sample will be almost pure $b\bar{b}$. In that case, mini-events could be written, allowing a larger rate to be output.

If mini-events are not written out, the level-2 M_{bb}^{50} output rate must be further reduced by a factor of 30–100 by the level-3 trigger. In principle, the $\mu 6 + \text{soft jets}$ sample could simply be prescaled: a prescale as high as 100 would give a significance of six in the 80–100 GeV window. However, it is clearly desirable

to use quality cuts to select golden events to be written on tape. This will be the subject of future investigation.

We emphasize that this channel is not accessible if b-tagging is not used in the trigger. It might be reasonable to think that b-tagging would not help much to reduce the $\mu 6 + \text{soft jets}$ sample that mainly comes from real b 's (74% of the events have a b inside). However, most of these events have a single b-jet (flavor excitation and g-splitting) and only $\sim 10\%$ of the $\mu 6 + \text{soft jet}$ sample pass the M_{bb}^{cut} of 50 GeV.

C. Low Cross-Section $b\bar{b}$ Processes: $b\bar{b}H/A$

For low cross-section $b\bar{b}$ processes that contain a distinctive and relatively rare signature, high efficiency trigger strategies can be developed. The SUSY Higgs final state studied here contains four b-jets, making b tagging at level 2 very efficient. The level-1 *all jet-a* selection is 45% efficient for the signal (see Table V) and produces the highest efficiency also at level-2 (26% in Table VI). However the level-1 output rate is 15 kHz. The rate for the *all jet-b* is an acceptable 4 kHz and the 18% efficiency at level-2 is greater than that of the μ selection. Of course, if a $\mu 6 + \text{soft jets}$ LVL1 is used for other purposes, the level-2 3-b trigger could be fed by both the *all jet* and $\mu 6 + \text{soft jet}$ LVL1, significantly increasing the overall efficiency for this process.

The analysis for this signal can be found in [12]. It requires the four most energetic jets to be tagged as leading b-jets. However, such a selective cut is unnecessary for trigger purposes.

The rate of events with three or more b-jets predicted by Pythia has very large uncertainties. However, it is interesting to note that at CDF the predicted rates seem to be much larger than those experimentally observed [13]. This is contrary to the common belief that shower Monte Carlos always underestimate QCD multijet production. It is known that, for final state multiplicities larger than two, most b-jets are generated either by gluon splitting (often gluons radiated from the initial state partons) or by flavor excitation. Pythia has been extensively tested against LEP data, but the processes of gluon splitting from initial state radiation and flavor excitation are not present in the LEP events.

In conclusion, we have hints from CDF that actual rates could be much smaller than those predicted by Pythia. Table VI shows that for the *3b leading* selection applied to *all-jet (b)* level-1 criteria the rates are reasonable and the signal efficiency is 13%. This is even better than the optimal one reported in [12] for the supersymmetric Higgs signal. We again note that the signal efficiency would be even larger if in addition to the *all-jet* trigger we used the $\mu 6 + \text{soft jet}$ trigger at level 1. Moreover, if the Pythia QCD cross section prediction is indeed overestimated as suggested by the CDF data, we could require three b-tags out of the four leading jets rather than demanding that all three of the leading jets be tagged. This would increase the overall signal efficiency in the trigger by approximately a factor of two by reducing the overall b-tagging inefficiency.

In order to make the physics case clearer, we cite [12, Fig. 5], which shows 5σ -discovery contours in the $m_A, \tan\beta$ plane for the $H/A \rightarrow 4b$ analysis and for integrated luminosities of 30 and 300 fb^{-1} . The figure compares the discovery contours due to trigger inefficiencies assumed in 1999 (solid and dashed

curves) to those potentially within the detector reach (dotted curves). Trigger inefficiencies basically come from high P_T jet thresholds and, in the up-to-date scenario of budget cuts, have been hardened since 1999, pushing the discovery curves to even higher values of $\tan\beta$.

Conversely, the trigger strategy just described here is fully efficient for this analysis, in the sense that it would correspond to the dotted lines in Fig. 5. As a result of the rate reduction from b-tagging in the level 2 trigger, we can record on tape events selected with low P_T thresholds (70, 50, and 15 GeV for the first three jets, respectively). These thresholds are significantly smaller than those presently set in the jet menus of the LHC experiments. For instance, the 4-jet sample is expected to have jets with P_T above 100 GeV, and the 3-jet sample has even higher thresholds. Currently, such high P_T thresholds are not so much due to a problem of the level-1 rate as to that of rate reduction in the level-2/3 trigger, since the total level-3 output budget for the jet triggers is in the range 10–25 Hz.

D. The CMS b-Jet Selection

It is interesting to see that an inclusive b-quark selection appears in the recent CMS High Level Trigger TDR [10]. However, only an inclusive b-jet selection is added to the trigger menu, with very high P_T thresholds (240 GeV).

Multi-b-jet final state selections do not appear in the menu yet. For example a particular selection for hadronic $H_t\bar{t}$ events that are very energetic is proposed in [10], with no use of a multi-b-jet request. Preliminary studies have shown [10] a valuable increase in discovery potential by using in addition the fully hadronic $H_t\bar{t}$ decays. In the CMS level-1 trigger four jets with $P_T > 50$ are required with a signal efficiency 60% (background rate of 200 Hz) and one jet is b-tagged at level-2 with a final rate of 20 Hz and a final signal efficiency of 55%. For comparison, we note that the all-jet selections we propose at level-1 would be more than 90% efficient for $H_t\bar{t}$. High quality multiple b-tagging at level-2 (3 b selection) reduces the rate to a few hertz with the advantage of overlapping the level-2 and offline analysis cuts: it will reject events that anyhow would be rejected offline.

The regional track reconstruction used at the CMS level-2 helps to reconstruct in a short time events characterized by few and small interesting regions. The multi-bjet selection instead is not easily implemented by this kind of architecture since it requires analysis of a very large fraction of complex events.

IV. ELECTRON AND TAU TAGGING

Level-2 b-tagging is not the only possible FTK application. The algorithm for every level-2 object could take advantage of high-quality full track reconstruction. Each level-2 algorithm at LHC experiments will start with the use of calorimeter and muon data in order to decrease the rate before proceeding with the time intensive track reconstruction. The use of FTK tracks in the first step of these algorithms would significantly reduce level-2 execution times and thus allow lower object thresholds and consequently larger acceptance for important physics processes. We review in the following an ATLAS and a CMS study that are very good examples.




Fig. 6. In [14, Fig. 8.3]. Different rate reduction paths in the Level 2-EventFilter selection chain for electrons at design luminosity as a function of the HLT steps. The efficiencies are given with respect to level-1 for the calorimeter selection at level-2 (LVL2-Calor) and level-3 (EF-Calor) and for different tracking selections at level-2 (LVL2-TRT and LVL2-Prec using, respectively, the TRT and silicon detectors) and level-3 (EF-ID).

A. High Quality Tracks for the Electron Selection

We cite Fig. 6 from [14, Sec. 8.4.1.3], which shows the results of a very interesting ATLAS study motivated by the attempt to minimize the use of system resources. The study is old, but it underlines the importance of having a level-3-quality (the ATLAS level-3 is named Event Filter, EF) tracking in the electron selection.

The study reports: “As an example, Fig. 6 shows that an increase in efficiency can be obtained, with a modest increase in the total HLT output rate, by moving the whole level-2 tracking selection to the EF. However, in this case, the input rate to the EF would increase by a factor of about eight, with important consequences on the computing load on the EF.” The EF tracking provides also the best rejection power on the background (more than a factor 10 in Fig. 6).

We underline that FTK would strengthen this beautiful and flexible architecture, providing EF quality tracks before the level-2 algorithms start, so that tracks can be used even at the electron level-1 output rate (~ 20 kHz) with very little computing power. The use of FTK tracks at level-2 will be faster than the pure calorimetric selection, since the level-2 track algorithm will be a simple loop over a list of precalculated objects.

B. High Quality Tracks for the Tau Selection

We report here an interesting CMS study [10] showing the importance of high quality tracking for the τ tagging. Fig. 7 shows some of the CMS results. The $H \rightarrow \tau\tau$ efficiency is plotted as a function of QCD efficiency, i.e., background rate. A good quality tracking algorithm (TRK) is directly applied to the first jet for events selected at level-1. This algorithm finds seeds in the pixels (Px1) and uses six points to reconstruct the tracks. Jets

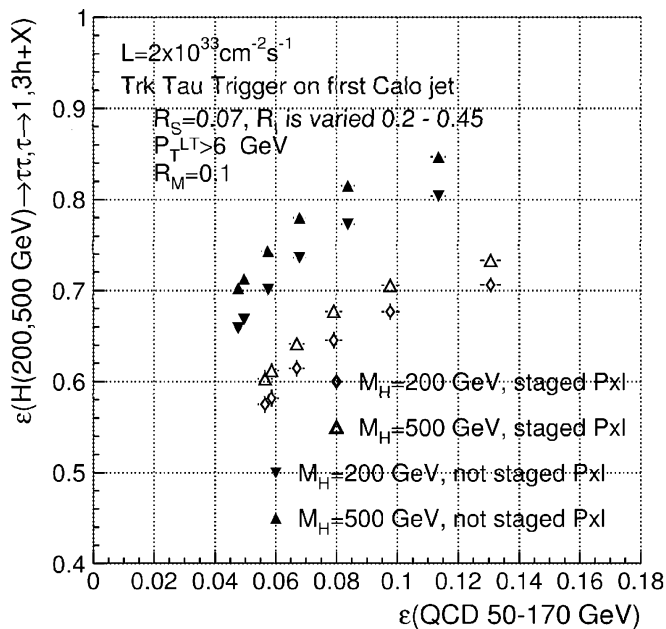


Fig. 7. In [10, Fig. 15–52]. Efficiency of TRK τ selection for the first calorimeter jet in $A/H \rightarrow \tau\tau \rightarrow 2\tau$ -jet for two higgs masses, $M_H = 200$ and 500 GeV, versus the efficiency for QCD di-jet background events at an instantaneous luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

are tagged as τ -jets if tracks belonging to the signal vertex are found in a signal cone ($R_s = 0.07$) centered on the leading signal track, that is the track with the highest P_T found in a matching cone ($R_m = 0.1$) centered on the jet axis. A larger cone (free parameter R_i) around the jet define the isolation region: no other tracks with $P_T > 1$ GeV should be found there except the few already found inside the signal cone. The details of the algorithm can be found in [10].

The double τ -jet tagging performed with the TRK tracking provides signal efficiencies of $\sim 60\%$ for QCD rejections of ~ 4 – 7% ([11, Fig. 15–54]). CMS has also studied the capability of a calorimetric algorithm. It is interesting to note that the same efficiency for Higgs ($\sim 60\%$) is obtained with a calorimeter algorithm that rejects QCD by only a factor of 10 ([10, Fig. 15–49 left]).

The calorimeter trigger causes efficiency losses that can be avoided if a high-quality tracking algorithm is used first in the high rate level-2 selection. In the CMS TDR it is concluded that a calorimeter step is necessary to decrease the level-1 output rate by a factor of 3 in order to reduce the total τ -tagging algorithm execution time. A pure tracking algorithm, even a simplified pixel-only algorithm, is predicted to be too slow to be executed at the level-1 output rate. The pixel algorithm will probably be used because it is faster, although this choice decreases the signal efficiency.

The FTK processor would provide high quality tracks before level-2 execution begins. The use of these tracks at level-2 will be even faster than the calorimeter step. As we plan to do for b-jets (Sections II and III), we could require at level-2 that two jets are identified as τ -jets, then we calculate the two τ -jets mass $M_{\tau\tau}$ at level-3. A cut on $M_{\tau\tau}$ can be used to further reduce the rate at level-3.

V. CONCLUSION

A hardware track finder for hadron collider experiments would provide rapid track reconstruction for all events passing the LVL1. By producing high quality tracks early in the level-2 process, FTK would speed up the level-2 algorithms for most objects, reducing the number of necessary CPUs and allowing higher LVL1 rates. In particular, the fast track finder would extend the trigger acceptance for interesting events with low P_T objects such as b-quarks or τ s, by separating them from the huge QCD background.

We have studied new trigger strategies to select b-quarks, using the LHC environment as a benchmark; more work must be done for τ s and inclusive electrons. The level-1 and level-2 criteria presented here are simple, well defined and widely applicable. The level-1 selection is based on two main strategies. For high energy events, a purely calorimetric trigger requires at least three jets, with much lower P_T thresholds than those possible without high quality level-2 tracking. For low energy events, the requirements consist of a soft μ plus a pair of low- P_T jets. Level-2 selection is based on “high-quality” b-tagging. A pair of b-tagged jets and a cut on its invariant mass, small enough to accept Z^0 events, reduces the rate to acceptable values. Event Filter selection is straightforward for multiple b-jet events, since a requirement of three b-tagged jets reduces the rate to few hertz. More studies aimed at specific physics signals will be carried out in the future.

Finally, FTK offers the possibility of compacting the event size (minievent) for selected calibration data samples. FTK track data could substitute for the full set of inner detector hits. The resulting large reduction in event size would allow a much larger number of events to be written to tape for the same bandwidth.

REFERENCES

- [1] A. Annoni *et al.*, “The fast tracker processor for hadron collider triggers,” *IEEE Trans. Nucl. Sci.*, vol. 48, pp. 575–580, June 2001. [Online] Available: <http://www.pi.infn.it/~orso/ftk>.
- [2] A. Bardi *et al.*, “SVT: An online silicon vertex tracker for the CDF upgrade,” *Nucl. Instrum. Methods*, vol. A409, pp. 658–661, May 1998.
- [3] “Design Study of the Large Hadron Collider (LHC),” The LHC Study Group, CERN 91-03, May 1991.
- [4] “ATLAS Technical Proposal,” ATLAS Collaboration, CERN/LHCC/94-43, LHCC/P2, Dec. 1994.
- [5] “ATLAS Inner Detector Technical Design Report,” ATLAS Collaboration, 1, 2 CERN/LHCC/97-16 and 17, ATLAS-TDR-4, and 5, Apr. 1997.
- [6] M. Cervetto. Use of CDF linear fit in track reconstruction. presented at ATLAS Pesa Software Workshop. [Online]. Available: <http://agenda.cern.ch/askArchive.php?base=agenda&categ=a021820&id=a021820s14t21/transparencies>
- [7] R. Carosi and G. Punzi, “An algorithm for a real time track fitter,” in *Proc. Conf. Rec. 1998 IEEE Nuclear Science Symp.*, Toronto, ON, Canada, Nov. 1998.
- [8] H. Wind, “Principal component analysis and its applications to track finding,” in *Formulae and Methods in Experimental Data Evaluation*, R. Bock, K. Bos, S. Brandt, J. Myrheim, and M. Regler, Eds: Eur. Phys. Soc., 1984, vol. III, pp. k1–k16.
- [9] T. Sjöstrand, P. Edén, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna, and E. Norrbin. (2001). *Computer Phys. Commun.* [Online], p. 238, vol.135. Available: <http://www.thep.lu.se/~torbjorn/>; The simulated processes are MSEL = 1 for QDC and ISUB = 186, 187, 181, 182, for the supersymmetric Higgs, with the Pythia version 6.136.

- [10] "The Trigger and Data Acquisition Project Technical Design Report, Volume 2, Data Acquisition and High-Level Trigger," CMS Collaboration, CERN/LHCC 02-26, CMS TDR 6, Dec. 2002.
- [11] E. Richter-Was, D. Froidevaux, and L. Poggioli, "ATLASFAST 2.0 a Fast Simulation Package for ATLAS," report, [Online] Available: <http://root.cern.ch/root/hepvis98/atlfast.html>, Nov. 1998.
- [12] "ATLAS Detector and Physics Performance: Technical Design Report," ATLAS Collaboration, 1, 2 CERN/LHCC/99-015, ATLAS-TDR-15, May 1999.
- [13] J. A. Valls, "CDF run II discovery reach for neutral MSSM higgs bosons via $pp \rightarrow b\bar{b}\phi \rightarrow b\bar{b}b\bar{b}$ with $\phi = h, H, A$," in *Fermilab Conf. 99/164-E*, 1999. Compare Tables 4 and 5.
- [14] "ATLAS High-Level Triggers, DAQ and DCS Technical Proposal," ATLAS Collaboration, CERN/LHCC/2000-17, Mar. 2000.