The BTeV Trigger System [1]

Daniel M. Kaplan [2]
Illinois Institute of Technology, Chicago, IL 60616, USA

for the BTeV Collaboration

Abstract

BTeV is a dedicated beauty and charm experiment proposed for the Fermilab Tevatron. The broad physics program envisaged for BTeV requires a trigger that is efficient for a wide variety of heavy-quark decays, including those to all-hadronic final states. To achieve this, we plan to trigger on evidence of detached vertices at the very first trigger level, taking advantage of fast-readout pixel detectors to facilitate fast pattern recognition. Simulations show that 100-to-1 rejection of light-quark background events can be achieved at Level 1 using specialized trackfinding hardware, and that an additional factor of 10 to 100 in data reduction can be achieved using general-purpose-processor farms at Levels 2 and 3. This is adequate to allow data-taking at luminosities in excess of $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

---

[1] Invited talk presented at the *6th International Conference on B-Physics at Hadron Machines (BEAUTY 99)*, 21–25 June 1999, Bled, Slovenia.

[2] kaplan@fnal.gov

1 Introduction

The goal of the proposed BTeV experiment [1] (Fig. 1) at the Tevatron Collider is to perform a comprehensive study of B-hadron mixing, CP violation, and rare decays during Tevatron Runs II and III. In addition BTeV could amass the world's largest sample of charm decays, thus carrying out a "multi-pronged" assault on the Standard Model's predictions in the heavy-quark sector. These goals are made possible by the Tevatron's large production cross sections for beauty and charm, estimated at $\approx 100\,\mu$b and $\approx 500\,\mu$b respectively, and by an ambitious Level-1 triggering scheme based on detached vertices. In marked contrast to other experiments [2], no Level-0 pretrigger is to be used.

At the proposed BTeV operating luminosity of $2\times10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, given the Tevatron's planned 7.6 MHz bunch-crossing rate there will be an average of 2 interactions per crossing. To search *every* crossing for evidence of detached vertices requires novel fast-readout pixel vertex detectors [3] coupled to a massively-parallel trigger system based on specialized trackfinding hardware. Moreover, the vertex detector puts out ~ 100 GBytes of data per second. At this data rate it is impractical to assemble a complete vertex-detector event for each beam crossing. The Level-1 trigger must therefore employ *sub-event* parallelism: the data from each event must be sliced up in some way, with the slices processed by separate dedicated processing systems. As described below, we are exploring various ways in which the data might be sliced.

2 Trigger System Overview

Fig. 2 shows the conceptual scheme we envision. Level 1 runs at the full crossing rate and incorporates conventional high-p_t lepton triggers as well as the vertex trigger. The vertex trigger is the main physics trigger, but the ancillary lepton triggers are simple to implement and will provide a cross-check of the efficiency of the vertex trigger. Events satisfying Level 1 will be sent to a farm of Level-2 processors, which will be general-purpose CPUs. The Level-2 processors will request event data from additional detectors besides the vertex detector in order to refine the trigger decision, with still more event data requested as successive cuts are passed. Once a complete event has been built the event is by definition at Level 3, although the Level-2 and Level-3 code may well reside on the same processors. At Level 3 a stripped-down version of the offline analysis can be performed, and the event can be compressed into a summary format to reduce the needed tapewriting bandwidth.

Although separate buffers are shown in Fig. 2, the Level-1, -2, and -3 buffers will most likely be implemented as a single buffer, i.e., the pieces of the event will reside in the Level-1 buffers for the individual detector subsystems until called forth for each level of the trigger algorithm. This arrangement (also used by CMS [4]) serves to reduce the event-building bandwidth needed, since the event is built up gradually as it progresses to higher and higher trigger levels, with complete events assembled only at the Level-2 output rate. Fig. 3 is a more detailed block diagram showing these features.

3 Baseline Level-1 Vertex Trigger

The centerpiece of the BTeV trigger system is the Level-1 displaced-vertex trigger. Its "baseline" version (Fig. 4) is based on work by the Penn-Fermilab group [5] and seeks to eliminate combinatorial processing time by an ingenious design of the vertex detector. The baseline vertex detector (Fig. 5) consists of 31 stations each containing three pixel planes, with pixel size $50 \times 400\,\mu\mathrm{m}^2$. The pixel planes are arranged in triplets, with the pixel narrow dimension of the outer two planes measuring the bend (y) view, and that of the inner plane measuring the nonbend (x) view. A triplet of overlapping pixels within a station thus determines not only a space point accurate to $\approx \pm 5\,\mu$m in both x and y, but also a "mini-vector" in space accurate to $\approx \pm 1$ mr in y and ± 50 mr in x that thus projects into a very small fraction ($\approx 2 \times 10^{-5}$) of the area of the next station, within which the occupancy due to other tracks or random hits is negligibly small. This largely eliminates the combinatorial processing time characteristic of trackfinding algorithms.

The most compute-intensive portion of the baseline algorithm is identifying the hit triplets within each station ("segment finding"). A design for the "station-hit" processors that perform this computation has been studied based on the TI TMS320C6X family[3] of integer DSPs. Sub-event parallelism is achieved by logically subdividing each pixel plane into 32 azimuthal sectors ("ϕ-slices") and sending the hits from each ϕ-slice to a separate dedicated station-hit processor. A Monte Carlo estimate of the mean occupancy per interaction gives $0.23/\phi$-slice.

[3] Rated at 1600 MIPS with 200 MHz clock.

The timings for segment finding are shown in Table 1. For simplicity, four representative cases are presented, in which the three planes (A, B, and C) of a station each have 1, 2, 3, or 4 hits. These are averaged in the bottom row of the table, weighted by the probability of occurrence of each case for single-interaction crossings. Note that the average is substantially less than the time for the (1,1,1) case since the most likely occupancy in a ϕ-slice is zero, for which case the TMS320C6X algorithm is assumed to be bypassed in zero time. The four columns of timings are for the TI C compiler with optimization turned off, C compiler with optimization on, hand optimization in assembly language, and hand optimization with the 3D geometry of the pixel plane [6] taken into account.

While the timing simulation has not yet been performed at 2 interactions/crossing, based on Poisson statistics the average time will approximately double, remaining comfortably below the 132 ns beam-crossing period. A nice feature of the TMS320C6X DSPs is their internal parallelism: as the detector occupancy (and hence the complexity of the calculation) increases, the parallel-processing efficiency increases also, so that whereas the (4,4,4) case might be expected to take 64 times as long as the (1,1,1) case, in fact the time goes up by a factor less than 20. We also see that hand optimization in assembler is essential to achieving the performance potentially available in the TMS320C6X.

The hit segments for each ϕ-slice are brought together in a set of 32 track farms, based on AD SHARC DSPs[4] in our design study. (Segments near ϕ-slice boundaries are sent to the track farm for the adjacent slice as well.) In the track farms, segments are linked together into tracks, which are then sent to vertex farms to be associated into primary vertices. Simulation shows that the Level-1 primary-vertex resolution may be characterized by a Gaussian-like core with $\sigma \approx 270\,\mu$m; non-Gaussian tails broaden the overall r.m.s. to $\approx 760\,\mu$m. This resolution is independent of luminosity up to at least $2.5 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The good performance for multiple-interaction crossings is made possible by the Tevatron's long bunches (30 cm r.m.s.): at our average of 2 interactions/crossing the probability of two primary vertices occurring sufficiently close to each other to be confused is well under 1%.

The Level-1 trigger decision is based on requiring n tracks to miss the primary vertex by at least $m\ \sigma$. Table 2 shows the Level-1 efficiency for various processes of interest requiring at least 2 tracks to miss the primary

[4] 120 MFLOPS (peak)/80 MFLOPS (sustained) with 40 MHz clock.

by 4σ. In addition, to improve the point-back resolution [7], displaced-vertex track candidates are required to exceed a transverse-momentum threshold of 0.5 GeV. (This is feasible using only pixel information since the vertex detector is located within the analysis magnet.) Efficiencies for typical beauty decays are $\gtrsim$50% while light-quark events are rejected by more than 100-to-1. Although it is more difficult to trigger on charm than on beauty (since charm lifetimes are shorter and secondary p_t's lower), we nevertheless achieve a respectable 6% efficiency for $D^0 \to K\pi$. Given the $\approx$order-of-magnitude higher charm cross section and the larger branching ratios, BTeV's large samples of beauty decays will be accompanied by much larger samples of charm decays, giving access to possible new physics via charm mixing, *CP* violation, and rare decays [8].

4 Levels 2/3

The Level-2 algorithm refines the tracks found at Level 1 by adding any pixel hits that the Level-1 algorithm may have missed, as well as the first few points from the forward tracking (straw tubes plus silicon strips), and performing a Kalman-filter track fit. This improves the momentum resolution from $\approx$ 10% to $\approx$ 2%, and the vertex-fit Gaussian core to $\sigma \approx 185\,\mu$m. With a 2-tracks-at-5σ detachment requirement, the result is a joint light-quark rejection of 1000-to-1 per crossing for Level 1 and Level 2 combined, with 40% overall $B^0 \to \pi^+\pi^-$ efficiency (efficiencies for other modes have not yet been checked but are expected to be comparably high).

Assuming existing magnetic-tape technology, the 7-kHz output rate of Level 2 is about an order of magnitude too high for recording of 200-kByte events. Level 3 thus needs to provide a factor $\approx$10 in bandwidth reduction. This can be achieved in various ways, for example by tighter vertex cuts, use of particle-ID information, or event compression (summarizing the raw data). The scenario of Fig. 2 assumes factors of about 5 each from event compression and from physics cuts.

5 Cost Estimate

We have carried out a fairly detailed cost estimate for the trigger and DAQ system, as shown in Table 3. While the system is not inexpensive, the cost is reasonable on the scale of a collider experiment.

6 Beyond the Baseline

The baseline trigger scenario presented above has the drawback of requiring 3 pixel planes per station, even though for offline analysis 2 planes per station are entirely sufficient. Given the total thickness of the baseline vertex detector (about 1 radiation length) as well as its cost, a 2-plane-per-station solution would clearly be desirable, and we are exploring other trigger approaches that are compatible with 2 planes per station.

Two-plane stations require giving up the mini-vector approach; the key trigger problem then becomes the trackfinding combinatorics. At 2 interactions/crossing, a pixel-plane quadrant has 8 hits on average; the probability of >16 hits is less than 1%. We are studying a combinatorial linefinder based on Xilinx gate arrays that can process the resulting large number of combinations in $<132\,\mathrm{ns}$/crossing. The linefinder looks for hit triplets in each view in three successive stations, using linear interpolation since track curvature may be neglected over 6 cm in z. This approach might be viewed as replacing detector slicing in ϕ with slicing in z: a separate linefinder is provided for each station, all of which run in parallel, and a subsequent step links the track segments together into tracks, finds the primary vertices, and identifies tracks that miss the primary.

The two-plane algorithm just described finds each track many times since it looks in all stations in parallel. We are also exploring alternative algorithms that could reduce the processing power needed by finding each track only once near the point of optimal vertex resolution. For example, this could be achieved by classifying pixel chips as "inner" (closest to the beam) and "outer" (farther from the beam) and looking for track segments that are composed of an "inner" pixel plus two "outer" pixels. Such an algorithm has been simulated and finds about 90% of tracks.

7 Conclusions

In summary, BTeV's goal of a heavy-quark hadroproduction experiment triggered on displaced vertices at luminosity up to $2\times10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ appears to be feasible. We are beginning tests of a Level-1 prototype board for the baseline algorithm and continuing our design studies of alternative algorithms. Our plan is to select the 2-plane or 3-plane approach by the end of this year and prepare a Proposal to Fermilab by May of 2000.

References

[1] A. Kulyavtsev *et al.*, "Proposal for an Experiment to Measure Mixing, CP Violation and Rare Decays in Charm and Beauty Particle Decays at the Fermilab Collider – BTeV," May 1999, http://www-btev.fnal.gov/public_documents/ptdr/ptdr.html.

[2] For example, HERA-B and LHC-b (see talks by E. Gerndt and O. Schneider, this Conference).

[3] L. Moroni, presented at this Conference.

[4] S. Cittolin *et al.*, in *Workshop on Recent Developments in High Energy Physics,* NCSR Demokritos, 9–11 April, 1998, p. 87.

[5] D. Husby *et al.*, *Nucl. Instrum. Meth.* **A383** (1996) 193.

[6] As discussed by Moroni [3], tiling the pixel plane with sensor chips requires them to be overlapped on both sides of a substrate, thus measurements within a station are at two different z values, complicating the segment-finding algorithm.

[7] W. Selove, in *Proceedings of the Workshop on B Physics at Hadron Accelerators*, P. McBride and C. S. Mishra, *eds.*, Fermilab-CONF-93/267 (1993), p. 617.

[8] D. M. Kaplan and V. Papavassiliou, to appear in *Proc. Workshop on CP Violation*, 3–8 July 1998, Adelaide, Australia, IIT-HEP-98/3, hep-ph/9809399 (1998); D. M. Kaplan, in *Proc. Symposium on Flavor Changing Neutral Currents: Present and Future Studies*, Santa Monica, CA, 19–21 Feb 1997, D. B. Cline, ed., World Scientific, 1997, p. 81.

hits/ϕ-slice (A,B,C)	$t_{\text{non-opt.}}$ (ns)	$t_{\text{comp.-opt.}}$ (ns)	$t_{\text{ass'y}}$ (ns)	t_{3D} (ns)
(1,1,1)	1915	1630	135	163
(2,2,2)	5705	3670	315	585
(3,3,3)	14015	7560	1340*	1055
(4,4,4)	28915	13840	2885*	3060
wt. avg.	530	395	37	49

*only partially optimized

Table 1: Timings for segment finding per single-interaction crossing.

Process	Eff. (%)
Min-bias	0.9
$B^0 \to \pi^+\pi^-$	55
$B_s \to D_s K$	70
$B^- \to D^0 K^-$	60
$B^- \to K_s \pi^-$	40
$B^0 \to K^+\pi^-$	54
$B^0 \to J/\psi K_s$	50
$B_s \to J/\psi K^*$	69
$B^0 \to K^*\gamma$	40
$D^0 \to K^-\pi^+$	6

Table 2: Level-1 trigger efficiencies for various processes of interest, evaluated for heavy-quark events passing all off-line cuts, or (for min-bias events) per crossing.

Item	Quantity	Unit cost ($)	Total ($)
L1 buffers	2000	1.3k	2.6M
L1 supervisor			0.1M
Quadrant proc. bd.	124	4.7k	0.6M
Farm proc. bd.	36	17.9k	0.6M
Data switch	200	2.0k	0.4M
Router/Evt. mngr.			0.5M
L2/3 buffers	500	1.0k	0.5M
Infrastructure			0.3M
L2/3 processors			2.2M
Total			7.8M

Table 3: Estimated costs for baseline trigger/DAQ components.

Figure 1: Elevation and plan views of the BTeV spectrometer.

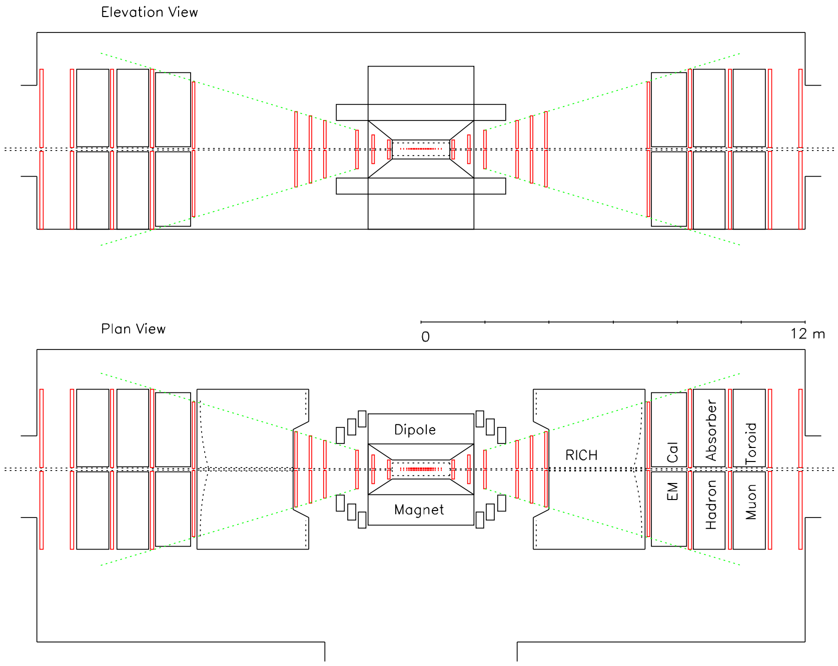

Figure 2: Conceptual diagram of BTeV 3-level trigger system.

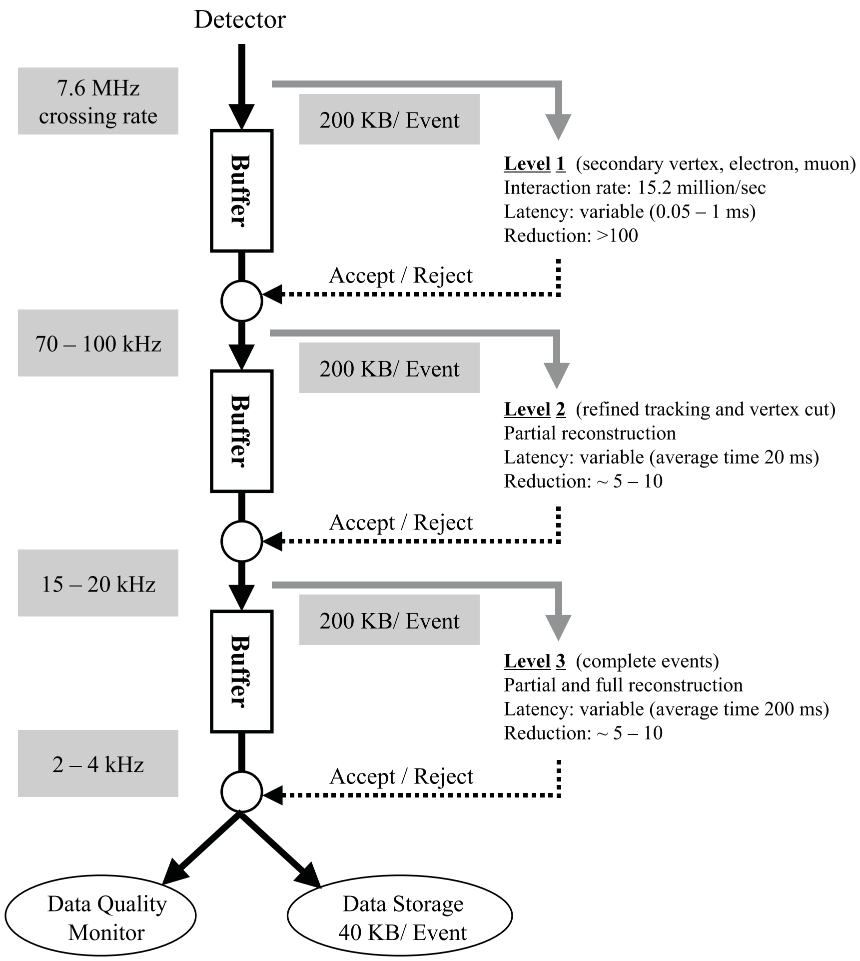

Figure 3: Block diagram of proposed DAQ system indicating design bandwidths before and after Level-1 trigger: up to 1500 GBytes/s of data are digitized from the detector subsystems and stored in Level-1 buffers, and up to 25 GBytes/s are passed to Levels 2 and 3.

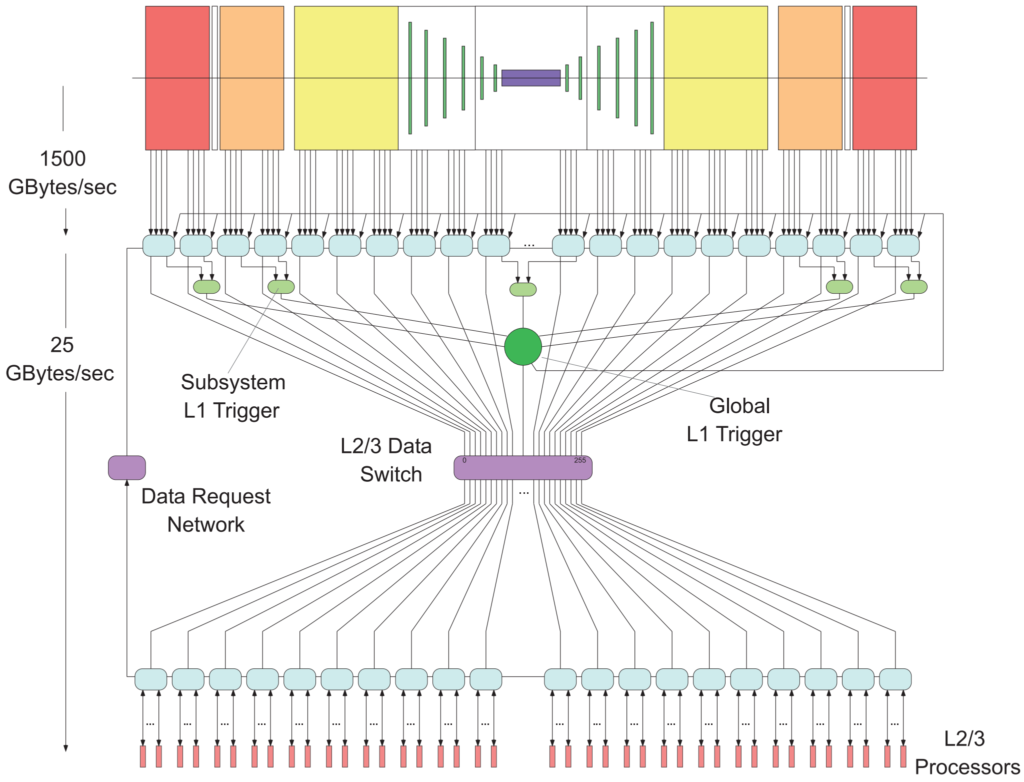

Figure 4: Block diagram of "baseline" Level-1 trigger.

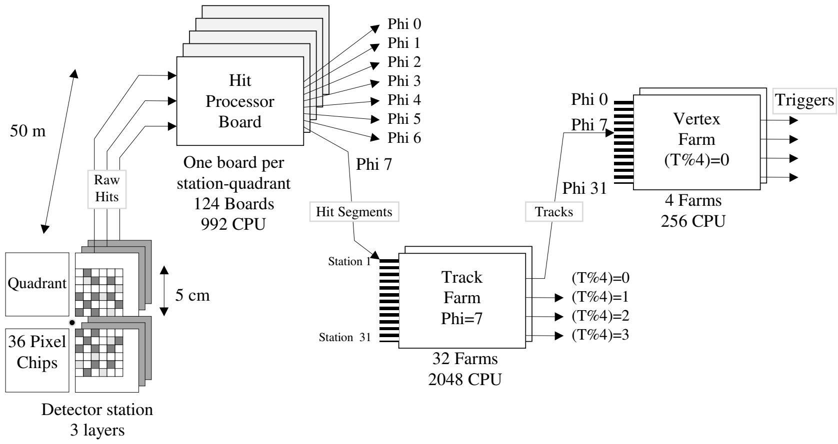

Figure 5: Layout of “baseline” pixel vertex detector.