Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Nuclear Instruments and Methods in Physics Research A 518 (2004) 544–548
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

Rapid 3D track reconstruction with the BABAR trigger upgrade

S. Bailey[a,*], R. Barlow[b], J. Boyd[c], G. Brandenburg[a], X. Chai[d], N. de Groot[e], N. Felt[a], G. Grenier[d], V. Halyo[f], S. Harder[a], O. Igonkina[g], W. Innes[f], M. Kelly[c], S. Kolya[b], S. Lee[d], U. Mallik[d], D. Mercer[b], M. Morii[a], J. Oliver[a], J. Olsen[f], N. Sinev[g], D. Su[f], E. Torrence[g], E. Won[a]

[a] *Harvard University, Cambridge, USA*
[b] *University of Manchester, UK*
[c] *University of Bristol, UK*
[d] *University of Iowa, USA*
[e] *Rutherford Appleton Laboratory, UK*
[f] *Stanford Linear Accelerator Center, USA*
[g] *University of Oregon, USA*

Abstract

As the PEP-II luminosity increases the BABAR trigger and dataflow systems must accomodate the increasing data rate. A significant source of background events at the first trigger level comes from the beam particle interactions with the beampipe and synchrotron masks, which are separated from the interaction region by more than 20 cm. The BABAR trigger upgrade will provide 3D tracking capabilities at the first trigger level in order to remove background events by distinguishing the origin of particle tracks. Each new z_0 p_T Discriminator board processes over 1 Gbyte of data per second in order to reconstruct the tracks and make trigger decisions based upon the 3D track parameters.

PACS: 07.05.Hd

Keywords: BaBar; Trigger; Data Acquisition; Track Reconstruction

1. Introduction

The BABAR detector records $e^+e^- \rightarrow \Upsilon(4S) \rightarrow B\bar{B}$ events produced at the PEP-II asymmetric B factory. The Level 1 (L1) trigger is a 12 μs fixed latency hardware trigger. Upon a L1 accept, the entire detector is read out for further processing by the Level 3 trigger. As the PEP-II luminosity increases, beam-related backgrounds are also expected to increase, driving up the L1 output rate from its current rate of approximately 1 kHz. The detector readout subsystems are limited to a maximum rate of 3.5 kHz, thus the L1 accept rate must be kept below this rate as well.

The BABAR drift chamber consists of 40 layers of drift chamber cells, organized into 4 axial superlayers and 6 stereo superlayers. The L1 Drift Chamber Trigger (DCT) receives drift chamber information at a rate of 3.7 MHz. It uses this

*Corresponding author.
E-mail address: bailey@slac.stanford.edu (S. Bailey).

0168-9002/$ - see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2003.11.081

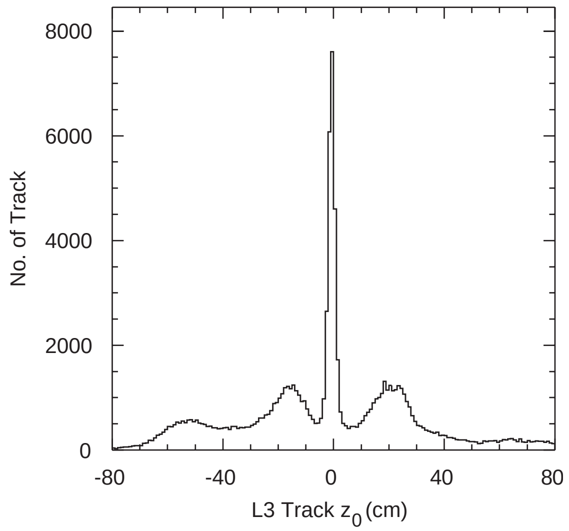

Fig. 1. Origin along the beamline (z-axis) of drift chamber tracks which pass the current DCT. The peak at $z_0 = 0$ is from e^+e^- collisions; the remaining tracks are primarily background tracks from beam particle collisions with beamline components.

information to reconstruct the tracks in order to make trigger decisions. Many of the background events which pass the current DCT are due to beam particle interactions with beamline components which are separated from the interaction region by more than 20 cm along the beamline (the z-axis), as shown in Fig. 1. The current DCT uses a 2D track reconstruction which has no z information and thus it cannot distinguish these background events from those which come from e^+e^- collisions.

The BaBar trigger upgrade will provide 3D tracking capabilities at the first trigger level in order to reject these background events. Eight new z_0 p_T Discriminator (ZPD) boards will reconstruct 3D track information and make trigger decisions based upon the z_0 origin of the tracks. The current BaBar detector and trigger systems are described in Ref. [1].

2. Track Segment Finder modules

Track Segment Finder (TSF) modules [2] receive 1 bit of information from each drift chamber cell every 269 ns. The TSFs use wire hit patterns and time development information to identify track segments in each superlayer with a resolution of ~ 800 μm. These track segments are sent to the ZPD modules which perform 3D track reconstruction. The current L1 DCT uses TSF modules which report detailed segment information only for the axial superlayers. The BaBar trigger upgrade requires 24 updated TSF modules which will send the segment information from both the axial and the stereo superlayers to the ZPDs.

3. ZPD algorithms

Each ZPD attempts to find and fit tracks which come from $z = 0$ and pass through a seed segment in superlayer 7 or 10. The TSFs provide up to 96 seed segments every 269 ns. The combined results from the track fits are used to make trigger decisions.

3.1. ZPD Track Finder algorithm

The ZPD Track Finder algorithm starts with a seed segment and looks for tracks which come from the interaction region and pass through that segment. It uses a Hough transform [3] to test 260 (p_T, $\tan\lambda$) combinations in parallel, where p_T is the transverse momentum and λ is the angle from the perpendicular to the z-axis. For each (p_T, $\tan\lambda$) combination, the Finder counts the number of superlayers which have segments that are consistent with those track parameters. The combination with the most matching superlayers is selected. In the case of multiple combinations having the same number of matching segments, an average of the (p_T, $\tan\lambda$) parameters is used. If no combination has 8 or more segments (including the seed segment), then no track is reported.

After finding the most likely track parameters for a given seed segment, the Finder associates segments with that track. It uses a look-up table to find the expected ϕ positions of segments in each superlayer based upon the p_T and $\tan\lambda$ combination selected. Within each superlayer, the closest segment to the expected position is associated with the track, provided that the closest segment is within one drift chamber cell width from the expected location.

3.2. ZPD Track Fitter algorithm

The Track Fitter begins by performing an r–ϕ fit to refine the curvature measurement. This curvature is used to determine the p_T of the track.

The differences in ϕ values between the seed segment and the other segments are due to the track curvature and the angles of the drift chamber wires in the stereo superlayers. The fitted curvature is used to subtract the ϕ difference due to curvature. The remaining ϕ residuals are related to the z position of the segments by

$$\phi^{\text{residual}} \approx z \frac{\tan \beta}{r} \quad (1)$$

where β is the wire stereo angle and r is the radius of the superlayer at the drift chamber endplate.

Defining d as the 2D track length when projected into the r–ϕ plane, the track forms a straight line in the d–z plane. The final Fitter step performs a straight line fit in the d–z plane to measure $\tan \lambda$ and z_0.

The z_0, error on z_0, p_T, $\tan \lambda$, and number of segments information from each track is used to make trigger decisions based upon cuts on these values.

4. Expected performance

The ZPD algorithms have been developed using real data from the BaBar drift chamber. This allows a realistic estimate of algorithm performance in the presence of actual detector operating conditions. To study the algorithm performance in the presence of larger backgrounds, background events from non-colliding beam runs have been overlaid on e^+e^- collision events. The algorithm degradation is minimal within the forseen operating conditions of the drift chamber. Additionally, Monte Carlo simulations have been used to study the efficiency of the trigger algorithms for specific types of signal events. The ZPD event selection must maintain an efficiency for physics signal events of greater than 99% while rejecting as many background events as possible.

The z_0 resolution of the Fitter algorithm is 4.1 cm. The majority of background tracks come from $|z_0| > 20$ cm.

Fig. 2 shows the efficiency of several types of events as a function of the $|z_0|$ cut applied by the ZPD. Triangles show the selection efficiency for Monte Carlo simulation of generic $B\bar{B}$ signal events. Circles show the selection efficiency for Monte Carlo simulation of the signal $B \rightarrow \pi^0\pi^0$. Since this signal has no charged tracks, the ZPD can only trigger on tracks from the other B. The efficiency is greater than 99.5% for cuts above 8 cm. Squares show the ZPD selection efficiency for background events which pass the current L1 trigger. These events are from a real data run where only the high energy e^- beam was present; Since there were no e^+e^- collisions possible in this run, all events are background events. With a cut of $|z_0| < 10$ cm, half of the background is rejected. The background rejection power is improved by including additional information such as a p_T requirement on the tracks.

The limiting factor for background rejection power of the current algorithms is the probability of creating a fake track with $z_0 \approx 0$. Since the Finder algorithm attempts to find tracks which come from $z_0 = 0$, any fake tracks created from

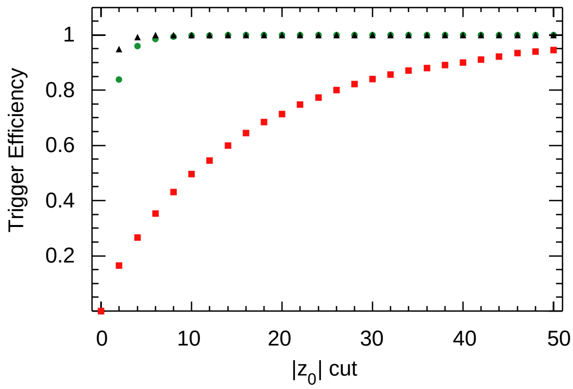

Fig. 2. ZPD selection efficiency for a variety of events types as a function of the $|z_0|$ cut. Triangles show Monte Carlo simulation of generic $B\bar{B}$ signal events; circles show Monte Carlo simulation of the signal $B \rightarrow \pi^0\pi^0$. Squares show background events from a data taking run with only the high-energy e^- beam present.

random segments will have a z_0 distribution peaking around 0. The fit quality of these tracks is worse than for a good track but the latency requirements on the algorithms restrict the quality checks which can be calculated. Algorithms are being developed to better distinguish the quality of tracks and reduce the fraction of fake tracks which are produced.

5. ZPD hardware implementation

The ZPD is a 12-layer 9U Eurocard using 8 Xilinx Virtex-II FPGAs. Each ZPD module will process 12 seed segments, covering 45° in ϕ. In order to reconstruct tracks with p_T down to 200 MeV/c, each ZPD must receive segment information from 135° in ϕ. Segment data arrive serially on 153 signals at 60 MHz from the J1 and J2 backplane connectors. These serial segments are received within a Decoder/Driver FPGA and sent across a 75-bit-pair "MegaBus" to 6 Finder/Fitter algorithm FPGAs.

The Megabus is implemented as a 75-bit-pair LVDS bus operating at 120 MHz. The signal traces are routed in straight lines directly beneath the ball grid array (BGA) Finder/Fitter FPGAs. This allows the signal taps to be very short. Since all Finder/Fitters need all segments, the pinouts for each Finder/Fitter FPGA are identical. Termination resistors are placed between the LVDS pairs on the driver side and after the last Finder/Fitter FPGA.

The Finder/Fitters are Xilinx Virtex-II 4000-ff1152 FPGAs. Approximately 40% of the resources are used by the Hough transform logic of the Finder algorithm. Another 5% is used for the Fitter algorithm and diagnostic memories. The Finder performance scales with the number of p_T and $\tan\lambda$ combinations considered; these can be expanded to use the remaining resources once the rest of the design is completely finished.

6. Testing and commissioning

ZPD testing has been performed using diagnostic memories which hold 64 events of data and can either record the incoming data or send data to the next stage. Diagnostic memories are located at the input and output of each FPGA as well as between the Finder and Fitter algorithms within the Finder/Fitter FPGAs.

The fastest signals on the ZPD are the 75 LVDS pairs of the Megabus, which operates at 120 MHz. Great care was taken to ensure clean routing and good termination of the Megabus signals in order to minimize noise and signal reflections. The Megabus has been tested with 10 Tbyte of data with no bit errors.

The input segment data arrive at 60 MHz on single ended signals which are passed through the backplane from an I/O interface board. These signals have been tested at the level of 10^{10} bits per signal with no errors.

Other electrical signals operate at a lower speed and have been tested with no errors found. Interface testing has also been performed between the various boards of the new DCT system with no bit errors.

The algorithm implementation has been tested by comparing the Finder and Fitter outputs with those expected from the C++ code which was used to study the algorithm performance. The ZPD has a bitwise match for all track results over thousands of events tested.

Integration testing and debugging is ongoing. A partial system test covering 45° in ϕ will be performed in June 2003 before the PEP-II Summer shutdown. In the Fall 2003 the upgraded DCT will be commissioned in parallel to the existing trigger system. A fiber optic splitter will provide drift chamber data to both the current trigger and the upgraded trigger, allowing comparisons between the systems and commissioning of the new trigger without interfering with the ongoing BaBar data taking operations. When the new trigger is fully ready, the global L1 trigger system will be switched to use the upgraded trigger as input.

7. Conclusions

The BaBar trigger upgrade plays an important role in ensuring that BaBar can maximize the excellent luminosity provided by PEP-II. The new

ZPD boards provide rapid 3D tracking in order to reject the dominant L1 background which comes from beam particle interactions with beamline components separated from the interaction region. The algorithms have been tested using real data and have good performance.

References

[1] The BaBar Collaboration, Nucl. Instr. and Meth. A 479 (2002) 1.
[2] A. Berenyi, et al., IEEE Trans. Nucl. Sci. NS-46 (1999) 348.
[3] P.V.C. Hough, A method and means for recognizing complex patterns, U.S. Patent 3069654, 1962.