Design and Implementation of the Level 1 Charged Particle Trigger for the BaBar Detector[1]

A. Berenyi, H.K. Chen, K. Dao, S.F. Dow, S.K. Gehrig, M.S. Gill, C. Grace, R.C. Jared, J.K. Johnson, A. Karcher, D. Kasen, F.A. Kirsten, J.F. Kral, C.M. LeClerc, M.E. Levi, H. von der Lippe, T.H. Liu, K.M. Marks, A.B. Meyer, R. Minor, A.H. Montgomery and A. Romosan
E.O. Lawrence Berkeley National Laboratory, Berkeley, California 94720

Abstract

The environment of the high-luminosity PEP-II machine poses unique design challenges for the trigger system of the BaBar detector. These led to the adoption of a real-time parallel pipelined architecture for the trigger electronics which departs significantly from previous implementations at conventional e^+e^- experiments. One challenge for the trigger designer lies in detecting low multiplicity physics events with high efficiency while keeping the background rate within the data acquisition limits. To achieve this difficult task, creative and innovative high-speed trigger algorithms were designed, simulated and implemented in Field Programmable Gate Arrays, using advanced CAD/CAE tools. The simulation results indicate that these algorithms will be able to perform all required tasks quickly and efficiently. This paper describes the design of the Level 1 Drift Chamber Trigger System of the BaBar detector, including the trigger algorithms, design and test methodology of the implementation, as well as test and simulation results.

I. Introduction

The task of a trigger system in any detector is to select physics events and to suppress background as efficiently as possible. For the BaBar detector [1] at PEP-II, the trigger system requirements for these tasks are very demanding, making the design of a suitable system nontrivial. The environment of the high-luminosity PEP-II machine poses unique challenges to the electronics system (front-end, trigger and data acquisition) for the BaBar detector. The PEP-II beam crossings occur at a rate of 238 MHz, which, in terms of the response time of the electronics, is essentially continuous. Unlike previous e^+e^- colliders, PEP-II may have severe beam backgrounds, resulting in high occupancies. Additionally, due to the lack of a fast detector element for signaling an annihilation, the trigger system must also determine the event time.

However, these are not the only tasks for the BaBar trigger. While the primary mission of the BaBar experiment is the study of CP violation in B decays, the trigger must also be designed to accept a variety of other physics processes as well. Because of their relatively high multiplicities and visible energies, it is straightforward to achieve high trigger efficiencies for Υ and $c\overline{c}$ events. Here the challenge lies in the extra redundancy and robustness against unforeseen beam backgrounds and impaired detector performance. Another challenge for the trigger designer lies in detecting low multiplicity $\tau^+\tau^-$ and $\gamma\gamma$ events with high efficiency while keeping the background rate within the data acquisition limits, so as not to incur deadtime. These low multiplicity events can be mimicked by beam-induced backgrounds, but differ by having: (1) high-P_t or high energy particles in low-multiplicity events, or high multiplicity; (2) back-to-back topology; (3) tag particles that can be matched with the angular locations of calorimeter towers and drift chamber tracks; (4) tracks that originate at the interaction point; (5) particles that have a common event time.

These factors led to the adoption of an innovative system architecture for all front-end, trigger, and data acquisition electronics which departs significantly from previous implementations at conventional e^+e^- experiments. In this architecture [2], the electronics systems collect data from the detector in real time, and process and transport it in parallel. The data are extensively buffered to avoid deadtime losses while making trigger decisions. The trigger system was built with a split-level architecture: a fast hardware processor-based (Level 1) trigger, which provides a front-end strobe, and initiates the feature extraction of pertinent data, and a software-based filter (Level 3) [3], running on commercial processors in the online system, which refines the selection, and further rejects events using software algorithms. Hooks for an optional Level 2 trigger are also built into the system, in case the Level 1 trigger rate exceeds the maximum Level 3 input rate.

The Level 1 trigger system must meet five primary requirements: (1) very high efficiency on all physics events; (2) the event time has to be determined down to the sub-μs level; (3) trigger decision time shorter than the 13 μs depth of the front-end buffers; (4) robust background rejection under all running conditions; (5) trigger rate below 2 kHz due to the limited bandwidth of the DAQ system downstream. The most practical method for delivering a Level 1 strobe in less than 13 μs is to use a hardware processor operating on a reduced representation of a subset of the data (trigger primitives). Ideally, one would use vertex detector information to effectively reject background events. However, simulations have shown that this is not practical since the bandwidth required to supply the necessary information to the Level 1 trigger logic is too high. Hence, to meet the physics requirements, the BaBar Level 1 trigger system has to rely on

[1]This work was supported by the Division of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. In addition, A.B.M. was supported by the Alexander von Humboldt Foundation, Bonn, Germany.

information from the Drift Chamber and Calorimeter only.

The Level 1 trigger system (see Figure 1) consists of three subsystems: the charged particle trigger (Drift Chamber Trigger, or DCT), the neutral particle trigger (Calorimeter Trigger, or EMT) and the Global Trigger (GLT). The DCT and EMT receive information from the Drift Chamber [4] and Calorimeter detectors, process it, and send condensed data to the GLT. The GLT attempts to match the spatial and angular locations of calorimeter towers and drift chamber tracks, and flexibly generates Level 1 triggers to the Fast Control System (FCS), based on the results of the processing with a rate less than 2 kHz. It also uses the Instrumented Flux Return (IFR) to select cosmic ray events. The Level 3 trigger then analyzes the event data from all detector sub-systems and reduces the rate at which events are written to archival storage to less than 100 Hz. This is the expected rate of all physics processes to be recorded at the design luminosity of $3 \times 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

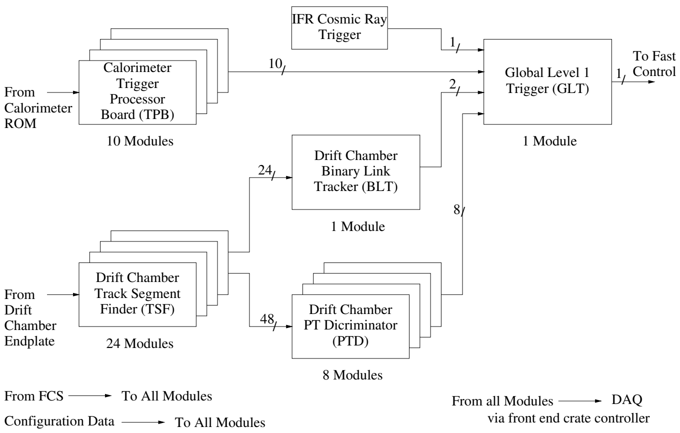

Figure 1: Level 1 Trigger Implementation

II. LEVEL 1 CHARGED PARTICLE TRIGGER CONCEPT

A. Input

A complete hit map including drift time information of the Drift Chamber is provided to the trigger via dedicated data paths every 269 ns. The Drift Chamber is a small-hex-cell design and contains 40 layers, with 96 to 256 cells per layer, for a total of 7104 drift cells. These 40 layers are grouped in 10 radial superlayers with four layers each. There are four axial superlayers (A1, A4, A7 and A10), with each pair of axial superlayers being separated by two stereo superlayers (U2, V3, U5, V6, U8 and V9). A magnetic field B of 1.5 T parallel to the beam axis, forces charged particles on a circular trajectory and thus facilitates the direct measurement of their transverse momenta. The Drift Chamber front-end electronics (FEE) system continuously contributes "real-time" information about particle-related activity in each of its channels to the Drift Chamber Trigger system. The input data to the trigger system consist of one bit of digital data (sampled continuously at 4 MHz) for each of the 7104 cells. The bit conveys time information from an amplitude discriminator whose input is the wire signal for that cell. This configurable discriminator resides in the ELEFANT chip [5] of the Drift Chamber FEE.

B. Trigger Algorithms

From these "hit" signals, the DCT continuously links hits into segments, chains segments into tracks, and simultaneously finds the angular azimuthal position (ϕ) of these tracks and determines whether the tracks represent particles having transverse momentum (P_t) greater than a preset value. This is a fully digital pipelined system, with a latency of approximately 6 μs, which outputs data every 134 ns.

The algorithms of the DCT are executed on three types of modules. First, the track segments (series of hits that are contiguous in space and time within a group) are found. This is done in a set of 24 Track Segment Finder Modules (TSFMs) [6]. The segments are then passed to the Binary Link Tracker Module (BLTM)[7], where they are linked into complete tracks. Detailed azimuthal position information for the segments found in axial superlayers is transmitted to the eight PT Discriminator Modules (PTDMs) [8], which determine if the segments are consistent with tracks of particles having a P_t greater than a user-specified minimum.

1) Track Segment Finder Module (TSFM)

The hits from each of the 7104 channels of the Drift Chamber are sent to 24 hardware-identical TSFMs via fiber-optic Gigalinks (one per module). Each module is responsible for processing a subset of the chamber data and extracting track segments formed from a set of contiguous hits within a group of neighboring cells (a pivot group, see Figure 2). Segments that could be part of a valid track are assigned a weight based on the number of hit cells within a pivot group and the quality of the data in the associated look-up tables (LUT). The contents of the LUT are derived from an offline calibration procedure operating on real or simulated data and consist of high-precision position and event time information. The weight is used mainly to estimate the segment arrival time.

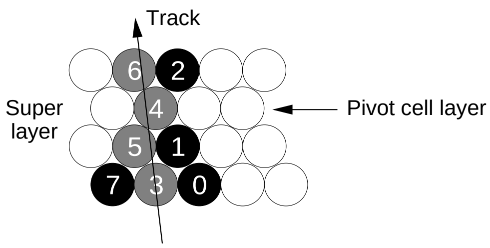

Figure 2: Track Segment Finder pivot cell group.

The search for track segments is organized in terms of pivot groups. For any pivot group, the cells are numbered 0 through 7, with cell 4 being the pivot cell. The shape of a pivot group was chosen such that only reasonably stiff tracks originating from the interaction point can produce a segment. The complete Segment Finder consists of 1776 track segment finder engines, one for each of the 1776 pivot groups. The engine processes data from the eight cells in its assigned pivot group to find valid track segments that pass through its pivot cell and other cells in the pivot group. Depending on where a track passed through a particular cell, the resulting ionization will take 1 to 4 3.7 MHz periods (clock ticks) to drift to the signal wire. It is this time delay, or drift time, that the TSF uses to establish more accurately the position of the track as well as the event time.

Typically, valid segment patterns consist of hits, close together in time and in at least three out of four layers within the superlayer (to account for cell inefficiencies). At the hardware level, this is implemented using a self-triggered 2 bit counter for each of the eight cells in the pivot group. A counter is enabled when a hit is registered. By incrementing the counter every 269 ns a 16-bit address, corresponding to 65536 possible addresses is obtained at any given clock tick. Each non-zero address is then translated by the pre-loaded LUT into a two-bit weight indicating whether there is no segment, a low-quality (un-calibrated) segment, a three-layer segment, or a four-layer segment. The pivot cell is monitored to determine which of the three subsequent clock ticks produces the highest weight, or 'best' pattern. The best recent pattern constitutes a found segment. The position resolution is expected to be better than $1/20$ cell width (0.7 mm) and the per segment event time resolution to be better than 70 ns. The time information is used to align each segment in time and to create a coincidence window for all segments related to the same event. The smaller the window, the better the real-time event time determination of the GLT. The resulting time-adjusted weight and position information are then passed on to the BLTM and PTDMs.

2) Binary Link Tracker Module (BLTM)

The BLTM receives segment hit information from all 24 TSFM at a rate of 320 bits per 7.4 MHz clock tick and links them into complete tracks. Tracks that reach the outer layer of the drift chamber (superlayer A10) are labeled of type "A". Tracks that reach the middle layers of the drift chamber (layer U5) are labeled of type "B". It is possible for some "B" tracks to continue far enough to also become indicated as an "A" track. A single BLTM module processes a 360 MByte/s stream of segment hit data, corresponding to information from the entire Drift Chamber. The segment hits are mapped onto the drift chamber geometry as 10 radial superlayers and 32 sectors in azimuthal direction (supercells). Each bit indicates whether a track segment was found in that sector or not.

The input data to the BLTM is logically "OR"ed with a programmable mask pattern. The masking capability allows the system to activate track segments that correspond to dead or highly inefficient cells so that the track linking efficiency does not degrade. The track linking algorithm uses an extension of a simple algorithm, developed for the CLEO II trigger [10], to link the segments into a continuous track. It starts from the innermost superlayer, A1, and moves radially outward. A track is found if (1) there is a segment hit in every layer, and (2) segments in two consecutive superlayers are within a certain number of supercells (three or five depending on the superlayer type) of each other. This allows for track curvature in the magnetic field and dip angle variations. The extension from CLEO II is to allow for up to two superlayers to be inefficient. The resulting track information is then extended by one clock tick earlier and later and time to account for possible clock jitter in the trigger system. Finally, the data are compressed and output to the GLT in the form of two 32-bit words corresponding to either "A" or "B" tracks respectively. Each bit in a word represents the "track hit" state of a supercell in the designated superlayer.

3) PT Discriminator Module (PTDM)

The PTDMs receive accurate position information on track segments found in the axial superlayers by TSFMs, and determine if the segments are consistent with tracks of particles having a P_t greater than some configurable minimum. Only the axial superlayers are used since their wires are parallel the B field axis, unlike the stereo layers.

The charged particle P_t is determined using the assumed beam vertex position together with a track segment position in one of the 'seed' layers (superlayers A7 or A10), and looking for segments in the other axial superlayers within the limits imposed by the minimum transverse momentum. Using the detailed azimuthal position information for the segments of the four axial superlayers, the eight PTDMs evaluate whether a sufficient number of track segments lie inside the pre-defined envelope and thus are consistent with a particle track of a transverse momentum P_t greater than a specified minimum. The resulting P_t turn-on at $600\,\mathrm{MeV}/c$ is expected to be narrower than 10%.

The processing on each board is subdivided in eight processing engines, one for each of the two superlayers (A10 and A7) in each 1/32 azimuthal wedge. The principal components in each channel are an algorithmic processor engine and programmable lookup memories containing the limits for each individual seed position. The contents of the memories thus specify the envelope of allowed track segment positions for each of the three other axial superlayers and consequently define the effective P_t discrimination threshold. The algorithm allows for inefficiencies.

III. Hardware Implementation

The board control logic on each board is implemented in two FPGAs, a Fast Control unit and a board Operation Controller. The Fast Control FPGA implements a high speed serial link responsible for remote control and configuration of the board. It is a generic chip common to all the Level 1 trigger modules. The Operation Controller interprets the commands from the Fast Control unit and coordinates the functionality of the board by implementing the control and interface logic specific to each type of board. A DAQ memory circular buffer is used to save data related to a particular Level 1 trigger. The data saved can be read out, allowing for later reconstruction of the processing that occurred in the each module during the creation of the particular Level 1 trigger, and also for rapid initial algorithms in the Level 3 trigger. The buffer has sufficient storage to hold the history of local data during the time required for the trigger decision to be made.

Each of the three main boards relies heavily on multiple FPGAs which perform all the on-board control and algorithmic functions. For example, one TSF module has 25 FPGAs which contain 0.5 million gates synthesized from 50,000 lines of VHDL code, and one BLT module has 5 FPGAs which contain 75,000 gates of programmable logic synthesized from about 10,000 lines of VHDL code. For the actual hardware implementation, the ORCA 2C series FPGAs from Lucent Technologies were chosenc̃iteMentor, because in 1995 when the FPGA family was selected the place and route test cases showed superior performance over other technologies. There are 24 TSF Modules, 8 PTD modules and 1 BLT module needed for the entire Level 1 Drift Chamber Trigger System. All cabling is handled by a small (6U) back-of-crate interface behind each main board.

IV. Design and Test Methodology

The hardware implementation for the trigger system was designed using a top-down design methodology.

A. Design Methodology

The design of the trigger architecture and selection algorithms was based on simulation studies of the detailed response of the BABAR detector to physics and background events. A software model of the complete Level 1 Trigger electronics pipeline was developed for this purpose. This model consists of 20,000 lines of Fortran code and describes potential hardware-implementable trigger algorithms. Extensive Monte Carlo studies running millions of physics benchmark and background data were performed. Over 400 possible parameters of the design were varied and tested. The final design was chosen from among those variations.

A major fraction of the design effort was dedicated to the implementation of the trigger algorithms described above in VHDL code to operate within FPGAs. Design and prototyping efforts were greatly speeded up by the use of a fully-integrated Computer Aided Design (CAD) toolset that included the VHDL hardware design language, synthesis, simulation, and automated board layout tools citeMentor. This highly automated design methodology enabled the creation of flawless circuit boards in spite of the nearly one million gate complexity.

B. Test Methodology

Due to the extensive use of FPGAs on the board, our ability to gain physical and electrical access needed for testing purposes is limited. Still, the design of all three main boards emphasized testability.

Input and output memories are attached to the input and output data streams of each module, allowing arbitrarily complex test vectors to be exercised for each board, independent of its external connections. By external control, they can be placed in the Save mode or Playback mode. This design feature provides powerful board-level as well as system-level diagnostics capabilities. In addition, DAQ memories are used to save and output data related to the generation of a particular Level 1 Trigger. This feature can also provide powerful board-level diagnostic capabilities during the debugging and testing stage for all the modules. The DAQ data will be used for monitoring during data taking.

Bit by bit, time frame by time frame, we have triple redundancy in testing the algorithms for each module and the system as a whole. As shown in Figure 3, one can debug and test each module in a systematic way through application of test vectors of real or GEANT simulated event data fed to Monte Carlo simulations of the Level 1 Trigger model. Additionally, during the debugging stage, one can also compare the results against the output from board level simulations using CAD tools. To gain access during the debugging stage to the internal signals inside each FPGA, 16 bit test pins are provided for most of the FPGAs on board. The internal signals (up to 16 at a time) inside the FPGA can be brought out to the test-pins and can be viewed by a Logic-Analyzer.

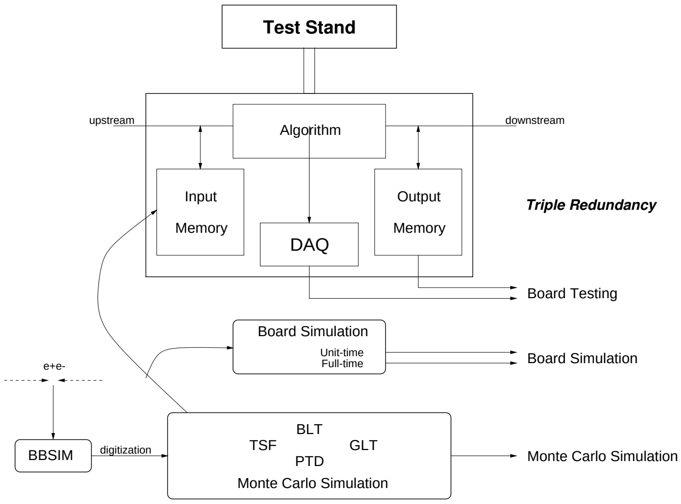

Figure 3: Board Level Test Methodology

C. Test Results

Each prototype module and final production module has undergone extensive testing of its algorithms using the board's diagnostic memory play/record features. Each was fed data from simulations of up to 10^5 Monte Carlo events, and the results were successfully compared against the software model for the Level 1 trigger. The test results were in full agreement with all computer simulations. Each module was also tested at a clock speed of 65 MHz (the nominal clock speed is 59.5 MHz) and no problems were found. No layout or fabrication errors were found on any of the prototype boards, allowing them to serve as the production version for the experiment.

All 24 TSF modules and one BLT module were installed in October 1998. The entire trigger system will be installed and commissioned by the end of 1998.

V. Level 1 Trigger Monte Carlo Simulation Results

Event signatures relevant to the DCT operation are: charged track multiplicity and event topology. Table 1 lists the various trigger selection algorithms along with the criteria used for each one. The multiplicity of each object is represented by its letter, e.g. A is the number of "A" tracks. A star (*) next to a trigger object indicates that a minimum angular separation was required in order to count more than one object (typically 90°). "A′" are "A" tracks for which a minimum p_t is required. Furthermore, a set of physics processes was identified to serve as trigger benchmarks. The trigger benchmarks represent processes that are difficult to distinguish from background.

Table 1
Trigger selection algorithms and their definitions.

D_2 trigger	$(A \geq 1 \ \& \ B \geq 2)$ OR $A \geq 2$
3A trigger	$A \geq 3$
D_2' trigger	$D_2 \ \& \ A' \geq 1$
$D_2'^*$ trigger	$D_2' \ \& \ B^* \geq 2$
2A′ trigger	$A \geq 2 \ \& \ A' \geq 1$

A. *Impact of Cell Inefficiencies on Trigger Efficiencies*

Cell and electronics simulations [11] indicate that the drift chamber efficiency will be above 99%. However, we simulated the trigger efficiency assuming various values for drift chamber efficiency. Table 2 lists trigger efficiencies for $\tau^+\tau^-$ events assuming a default p_t cut of 0.8 GeV/c. The open trigger efficiency (D_2) does not degrade as much as the D_2' trigger efficiency, since the latter uses the PT discriminator which relies on fine position information from the supercells of the drift chamber. If the PT discriminator is used, τ physics is compromised for a cell efficiency below 95%.

Table 2
Drift Chamber Trigger efficiencies for $\tau^+ \rightarrow \pi^+\overline{\nu}$, $\tau^- \rightarrow \pi^-\nu$ with fiducial volume cuts applied. The same 1000 generated events were used for each run.

$\tau^+ \rightarrow \pi^+\overline{\nu}$, $\tau^- \rightarrow \pi^-\nu$	D_2	D_2'	$D_2'^*$
99% Cell efficiency	0.99	0.91	0.88
95% Cell efficiency	0.99	0.88	0.84
90% Cell efficiency	0.95	0.78	0.75

B. *Trigger Efficiencies and Background Rates*

Beam related backgrounds are evaluated at the design luminosity, and due to their variable nature and uncertainties in the predictions, pose more stringent design requirements. Different background conditions have been studied, with a "worst-case scenario" defined as ten times the nominal background rate. Table 3 lists the trigger efficiencies for B and τ benchmark processes for both nominal and $\times 10$ background conditions. Under "worst-case" scenario background conditions, the τ physics is affected while triggering on B physics is still very efficient.

Table 3
Level 1 drift chamber trigger efficiencies for B and τ benchmark processes for both nominal ($\times 1$) and "worst-case scenario" ($\times 10$) background conditions. The same 1000 events were used for each run.

	$\times 1$	$\times 10$		
Process	D_2	$D_2'^*$	3A	$D_2'^* + 3A$
$B^0 \rightarrow \pi^+\pi^-$, $\overline{B}^0 \rightarrow \mu X$	1.00	0.95	0.95	0.99
$B^0 \rightarrow D^+D^-$, $\overline{B}^0 \rightarrow X$	1.00	0.75	0.99	0.99
$B^+ \rightarrow \tau^+\nu$, $\overline{B}^- \rightarrow X$	1.00	0.66	0.91	0.95
$\tau^+ \rightarrow e^+\nu\overline{\nu}$, $\tau^- \rightarrow \mu^-\overline{\nu}\nu$	0.98	0.82	0.04	0.83
$\tau^+ \rightarrow \pi^+\overline{\nu}$, $\tau^- \rightarrow \pi^-\nu$	0.99	0.85	0.14	0.86

These and other extensive simulation studies [12] show that the total Level 1 trigger rate from all sources is below the required 2 kHz, while maintaining an efficiency of more than 99% for most physics benchmark events after detector fiducial cuts. The simulated event time jitter window for 99% of events, even under worst-case benchmarks is less than 200 ns, much less than the required 1 μs. These results validate the conceptual design of the Level 1 Drift Chamber trigger.

VI. References

[1] The BaBar Collaboration, "Technical Design Report", *SLAC-R-95-457 March 1995.*

[2] C.T. Day *et al.,* "The Babar Trigger, Readout and Event Gathering System", *Proceedings of the International Conference on Computing in High Energy Physics, 18-22 September (1995).*

[3] G.P. Dubois-Felsmann *et al.*, "Architecture of the BaBar Level 3 Software Trigger", submitted to IEEE Transactions on Nuclear Science.

[4] J. Albert *et al.*, "Electronics for the BaBar Central Drift Chamber", submitted to IEEE Transactions on Nuclear Science.

[5] S. Dow *et al.,* "Design and Performance of the Elefant Digitizer IC for the BaBar Drift Chamber", submitted to IEEE Transactions on Nuclear Science.

[6] A. Berenyi *et al.*, "Continuously Live Image Processor for Drift Chamber Track Segment Triggering", submitted to IEEE Transactions on Nuclear Science.

[7] A. Berenyi *et al.*, "A Binary Link Tracker for the BaBar Level 1 Trigger System", submitted to IEEE Transactions on Nuclear Science.

[8] A. Berenyi *et al.*, ''A Real-Time Transverse Momentum Discriminator for the BaBar Level 1 Trigger System", submitted to IEEE Transactions on Nuclear Science.

[9] Mentor Graphics Corporation: QuickVHDL Pro (Model Technology), Design Architect, Librarian, Package, Fablink, Layout; Cadence: Spectra PC Board Router; Synopsys Corporation: Design Analyzer, FPGA Shell; and Lucent Technologies: FPGA Foundry.

[10] K. Kinoshita, *Nucl. Instr. Methods A276 (1989)242 .*

[11] W. Innes, "Efficiency and Time vs. Distance in the BaBar Drift Chamber", *BaBar Note #329, October 1996.*

[12] S.K. Gehrig, "Design and Simulated Performance of the Level 1 Trigger System", *BaBar Note #380, Sept. 1997.*