The *BABAR* Detector

X. Trigger

X.1. Trigger Requirements

The basic requirement for the trigger system is the selection of events of interest (see Table 1) with a high, stable, and well-understood efficiency while rejecting background events and keeping the total event rate under 120 Hz. At design luminosity, beam-induced background rates are typically about 20 kHz each for one or more tracks in the drift chamber with p_t>120 MeV/c or at least one EMC cluster with E >100 MeV. Efficiency, diagnostic, and background studies require prescaled samples of special event types, such as those failing the trigger selection criteria, and random beam crossings.

The total trigger efficiency is required to exceed 99% for all $B\overline{B}$ events and at least 95% for continuum events. Less stringent requirements apply to other event types, *e.g.*, $\tau^+\tau^-$ events should have a 90-95% trigger efficiency, depending on the specific $\tau^\pm$ decay channels.

The trigger system must be robust and flexible in order to function even under extreme background situations. It must also be able to operate in an environment with dead or noisy electronics channels. The trigger should contribute no more than 1% to dead time.

X.2. Trigger Overview

The trigger is implemented as a two-level hierarchy, the Level 1 (L1) in hardware followed by the Level 3 (L3) in software. It is designed to accommodate up to ten times the expected PEP-II background rates at design luminosity and to degrade slowly for backgrounds above that level. Redundancy is built into the system to measure and monitor trigger efficiencies.

During normal operation, the L1 is configured to have an output rate of typically 1 kHz. Triggers are typically produced within a fixed latency window of 11-12 μs after the e^+e^-collision, and delivered to the Fast Control and Timing System (FCTS). Data used to form the trigger decision are preserved with each event for efficiency studies.

The L3 responds to the entire L1 output, performs a second stage rate reduction for the main physics sources, and identifies and flags the special categories of events needed for luminosity determination, diagnostic, and calibration purposes. At design luminosity, the L3 filter acceptance for physics is ∼90 Hz, while ∼30 Hz contain the other special event categories. The L3 algorithms comply with the same software conventions and standards used in all other *BABAR* software, thereby simplifying its design, testing, and maintenance.

X.3. Level 1 Trigger System

The L1 trigger decision is based on charged tracks in the DCH above a preset transverse momentum, showers in the EMC, and tracks detected in the IFR. Trigger data are processed by three specialized hardware processors. As described below, the drift chamber trigger (DCT) and electromagnetic calorimeter trigger (EMT)

Table 1
Cross sections, production and trigger rates for the principal physics processes at 10.58 GeV for a luminosity of $3 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The e^+e^- cross section refers to events with either the e^+, e^-, or both inside the EMC detection volume.

Event type	Cross section (nb)	Production Rate (Hz)	Level 1 Trigger Rate (Hz)
$b\overline{b}$	1.1	3.2	3.2
other $q\overline{q}$	3.4	10.2	10.1
e^+e^-	∼53	159	156
$\mu^+\mu^-$	1.2	3.5	3.1
$\tau^+\tau^-$	0.9	2.8	2.4

both satisfy all trigger requirements independently with high efficiency, and thereby provide a high degree of redundancy. The instrumented flux return trigger (IFT) is used for triggering $\mu^+\mu^-$ and cosmic rays, mostly for diagnostic purposes.

The overall structure of the L1 trigger is illustrated in Figure 1. Each of the three L1 trigger processors generates trigger *primitives*, summary data on the position and energy of particles, that are sent to the global trigger (GLT) every 134 ns. The DCT and EMT primitives sent to the GLT are ϕ-*maps*. An individual ϕ-map consists of an n-bit word representing a particular pattern of *trigger objects* as distributed in fixed-width ϕ regions from 0 to 2π. A trigger object is a quantity indicating the presence of a particle, such as a drift chamber track or a calorimeter energy deposit. The IFT primitive is a three-bit pattern representing the hit topology in the IFR. The meaning of the various trigger primitive inputs to the GLT are summarized in Table 2.

The GLT processes all trigger primitives to form specific triggers and then delivers them to the FCTS. The FCTS can optionally mask or prescale any of these triggers. If a valid trigger remains, a L1 Accept is issued to initiate event readout. The trigger definition logic, masks, and prescale values are all configurable on a per run basis.

The L1 hardware is housed in five 9U VME crates. The L1 trigger operates in a continuous sampling mode, generating trigger information at regular, fixed time intervals. The DCH FEEs and the EMC *untriggered personality cards* (UPCs) send raw data to the DCT and EMT about 2 μs after the e^+e^-collision. The DCT and EMT event processing times are 4–5 μs, followed by another $\sim$3 μs of processing in the GLT to issue a L1 trigger. The L1 trigger takes approximately 1 μs to propagate through the FCTS and the *readout modules* (ROMs) to initiate event readout. These steps are all accomplished within the 12.8 μs FEE buffer capacity limit.

The DCT, EMT and GLT each maintain a four-event buffer to hold information resulting from the various stages of the L1 trigger. These data are read out by the normal data acquisition system.

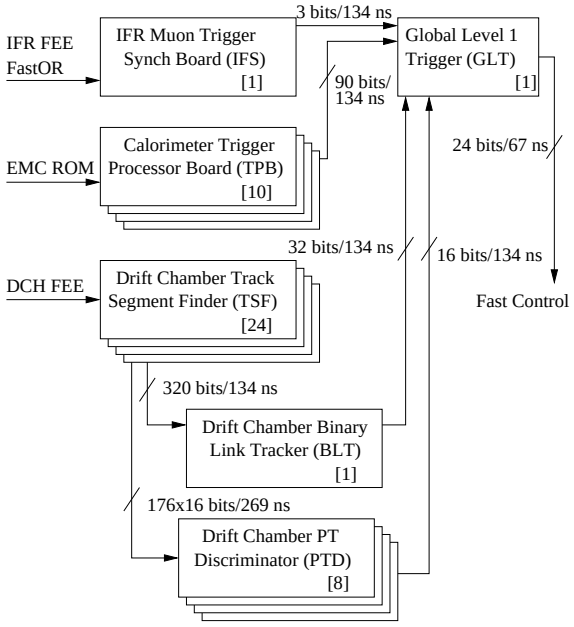

Figure 1. Simplified L1 trigger schematic. Indicated on the figure are the number of components (in square brackets), and the transmission rates between components in terms of total signal bits.

X.3.1. Level 1 Drift Chamber Trigger

The input data to the DCT consist of one bit for each of the 7104 DCH cells, updated every 269 ns. These bits convey time information derived from the sense wire signal for that cell. The DCT output primitives are candidate tracks encoded in terms of three 16-bit ϕ-maps as listed in Table 2.

The DCT algorithms are executed in three types of modules [**?**]. First, track segments, their ϕ positions and drift time estimates are found using a set of 24 Track Segment Finder (TSF) modules [1]. These data are then passed to the Binary Link Tracker (BLT) module [2], where segments are linked into complete tracks. In parallel, the ϕ information for segments found in axial superlayers is transmitted to eight transverse momentum discriminator (PTD) modules [3], which search for tracks above a set $p_{\rm t}$ threshold.

Table 2
Trigger primitives for the DCT and EMT. Most energy thresholds are adjustable; those listed are typical values.

	Description	Origin	No. of bits	Threshold
B	Short track reaching DCH superlayer 5	BLT	16	120 MeV/c
A	Long track reaching DCH superlayer 10	BLT	16	180 MeV/c
A′	High p_t track	PTD	16	800 MeV/c
M	All-θ MIP energy	TPB	20	100 MeV
G	All-θ intermediate energy	TPB	20	250 MeV
E	All-θ high energy	TPB	20	700 MeV
X	Forward endcap MIP	TPB	20	100 MeV
Y	Backward barrel high energy	TPB	10	1 GeV

Each of the three DCT modules (TSF, BLT, and PTD) relies heavily on multiple FPGA's [4] which perform the control and algorithmic functions. All cabling is handled by a small (6U) back-of-crate interface behind each main board.

Track Segment Finder

The TSF modules are responsible for finding track segments in 1776 overlapping eight-cell *pivot groups*. A pivot group is a contiguous set of cells that span all four layers within a superlayer. The pivot group shape is such that only reasonably straight tracks originating from the interaction point can produce a valid segment. Figure 2 shows the arrangement of cells within a pivot group. Cell 4 is called the *pivot cell*; the TSF algorithm is optimized to find track segments that *pivot* about this cell.

The DCH signals are sampled every 269 ns. The passage of a single particle through the DCH will produce ionization that drifts to the sense wires in typically no more than four of these clock ticks. Each cell is associated with a two-bit counter that is incremented at every clock tick for which a signal is present. In this way, a short time history of each cell is preserved. For each clock tick, the collection of two-bit counters for each pivot group forms a 16-bit value used to address a look-up-table. This look-up-table contains two-bit weights indicating whether there is no acceptable segment, a low-quality segment, a three-layer segment (allowing for cell inefficiencies), or a four-layer segment. When an acceptable segment is found, that pivot group is examined to determine which of three subsequent clock ticks produce the highest weight or *best* pattern.

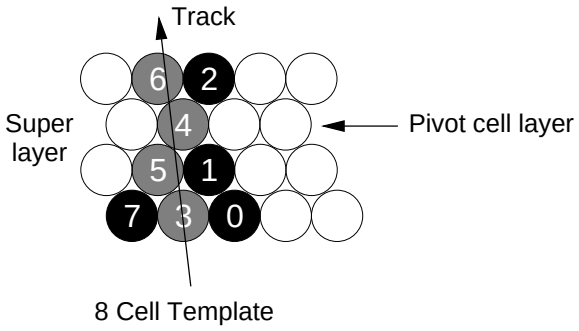

Figure 2. Track Segment Finder pivot group.

The look-up-table also contains position and time information which, along with a summary of cell occupancies, forms the basis of data sent to the BLT and PTD. The TSF algorithm uses the time-variation of the look-up-table weights to refine both the event time and its uncertainty, thus enabling it to output results to the BLT every 134 ns.

The position resolution as measured from the data after calibration, is typically ∼600 μm for a four-layer segment and ∼900 μm for a three-layer segment. For tracks originating from the IP, the efficiency for all TSF segments is 97%, and the efficiency for high-quality three-layer or four-layer TSF segments is 94%.

Binary Link Tracker

The BLT receives segment hit information from all 24 TSF's at a rate of 320 bits every 134 ns and links them into complete tracks. The segment hits are mapped onto the DCH geometry in terms of 320 *supercells*, in 32 sectors in ϕ and ten radial superlayers. Each bit indicates whether a segment is found in that supercell or not. The BLT input data are combined using a logical OR with a programmable mask pattern. The masking allows the system to activate track segments corresponding to dead or highly inefficient cells to prevent efficiency degradation. The linking algorithm uses an extension of a method developed for the CLEO-II trigger [5]. It starts from the innermost superlayer, A1, and moves radially outward.

A track is found if there is a segment in at least eight superlayers and if the segments in two consecutive superlayers fall azimuthally within three to five supercells of each other (depending on the superlayer type). This allows for track curvature and dip angle variations. Tracks that reach the outer layer of the DCH (superlayer A10) are classified as type A. Tracks that reach the middle layer (superlayer U5) are classified as type B. The data are compressed and output to the GLT in the form of two 16-bit ϕ-maps, one each for A and B tracks.

PT Discriminator

The eight PTD modules receive ϕ information of high quality track segments in the axial superlayers (A1, A4, A7 and A10), and determine if the segments are consistent with a track p_t greater than a configurable minimum value. An envelope for tracks above the minimum p_t is defined using the IP, and a track segment position in one of the *seed* superlayers, A7 or A10. A high p_t candidate, denoted as A′, is identified when sufficient track segments with accurate ϕ information from the other axial superlayers lie within this envelope.

Each PTD module searches for seed segments in superlayers A7 and A10, and within a 45-degree azimuthal wedge of the DCH. This search region spans eight supercells, and the processing for each supercell is performed by its own processing engine on the PTD. The principal components

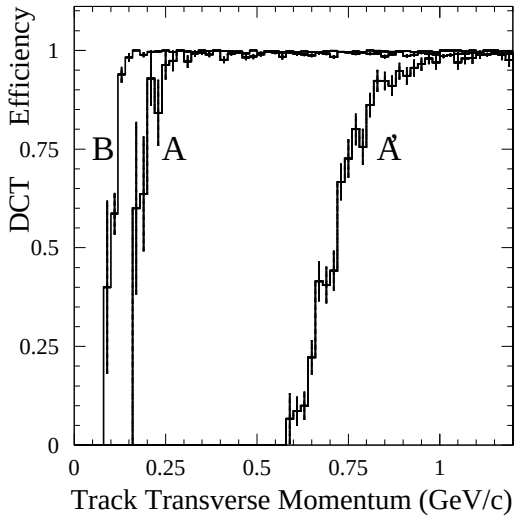

Figure 3. DCT track efficiency versus transverse momentum for A, B, and A′ tracks. The A′ threshold is set to 800 MeV/c.

in each engine are an algorithmic processor and look-up-tables containing the limits for each individual seed position. The contents of the look-up-tables specify the allowed track segment positions for each of the three other axial superlayers and consequently define the effective p_t discrimination threshold. The resulting p_t threshold for the PTD A′ tracks is shown in Figure 3 together with the BLT A, B track efficiency.

X.3.2. Level 1 Calorimeter Trigger

For trigger purposes, the EMC is divided into 280 towers, 7×40 ($\theta \times \phi$). Each of the barrel's 240 towers is composed of 24 crystals in a 8×3 ($\theta \times \phi$) array. The endcap is divided into 40 towers, each forming a wedge in ϕ containing 19–22 crystals. For each tower, all crystal energies above a threshold of 20 MeV are summed and sent to the EMT every 269 ns.

The conversion of the tower data into the GLT ϕ-maps is performed by ten Trigger Processor Boards (TPBs). The TPBs determine energies in the 40 ϕ sectors, summing over various ranges

of θ, compare these energies against thresholds for each of the trigger primitives (see Table 2), estimate the time of energy deposition, correct for timing jitter, and then transmit the result to the GLT.

Each TPB receives data from 28 towers, corresponding to an array of 7×4 in $\theta \times \phi$, or four ϕ-*sectors*. Each of the 40 ϕ-sectors is summed independently. To identify energy deposits that span two adjacent ϕ-sectors, the energy of each sector is also made available to the summing circuit for a single adjoining sector in such a way that all possible pairs of adjacent ϕ-sectors are summed. These energy sums are compared against thresholds to form trigger objects. Each sum is also sent to an eight-tap finite impulse response (FIR) digital filter which is used to estimate the energy deposition time. A look-up-table is used to make an energy-dependent estimate of the timing jitter which, along with the FIR output, is used to time the transmission of any trigger objects to the GLT. Pairs of ϕ-sectors are ORed to form 20-bit ϕ-maps for the M, G, E, and X primitives, while for the Y primitive, groups of four are ORed to form a 10-bit ϕ-map. The complete algorithm is implemented in one FPGA [7] for each ϕ-sector, with four identical components per TPB. Further details of the EMT system can be found in [6].

The basic performance of the EMT can be expressed in terms of the efficiency and timing jitter of the trigger primitives. The efficiency of the primitives can be measured by the number of times a trigger bit is set for a specific energy reconstructed offline in events from a random trigger. Figure 4 shows this efficiency for energies near the M threshold. The efficiency changes from 10% to 90% in the range of 110 to 145 MeV, and reaches 99% at 180 MeV, close to the average energy deposition of a minimum ionizing particle at normal incidence.

The EMT time jitter is measured by comparing the time centroid of ϕ-strip M hits in $\mu^+\mu^-$ events with the DCH track start time, t_0. The difference has an rms width of 90 ns with >99.9% of the matching M hits within a $\pm$500 ns window.

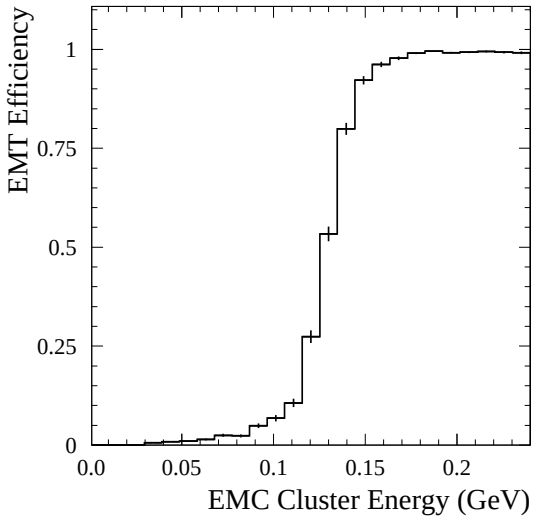

Figure 4. EMT M efficiency vs. EMC cluster energy for an M threshold setting of 120 MeV.

X.3.3. Level 1 IFR Trigger

The IFT is used for triggering on $\mu^+\mu^-$ and cosmic rays. For the purposes of the trigger, the IFR is divided into ten sectors, namely the six barrel sextants and the four half end doors. The inputs to the IFT are the *Fast OR* signals of all ϕ readout strips in eight selected layers in each sector.

A majority logic algorithm defines trigger objects for every sector in which at least four of the eight trigger layers have hits within a time window of 134 ns. The IFR trigger synchronization module processes the trigger objects from the ten sectors and generates the three-bit trigger word (U) encoding seven exclusive trigger conditions, as defined in Table 3. The trigger U$\geq$5, for example, covers all $\mu^+\mu^-$ topologies of interest.

The efficiency of the IFT has been evaluated using cosmic rays triggered by the DCT and crossing the detector close to the IP. For these events, 98% were triggered by the IFT as events with at least one track, and 73% as events with two tracks, inside the geometrical region of the IFR. Most of the IFT inefficiency is concentrated at

Table 3
IFR trigger pattern (U) definition, where μ refers to a signal within a sector.

U	Trigger condition
1	$\geq$2 μ topologies other than U=5-7
2	1 μ in backward endcap
3	1 μ in forward endcap
4	1 μ in barrel
5	2 back-back μ's in barrel + 1 forward μ
6	1 μ in barrel + 1 forward μ
7	2 back-back μ's in barrel

the boundaries between sectors.

X.3.4. Global Trigger

The GLT receives the eight trigger primitives in the form of ϕ-maps as listed in Table 2 along with information from the IFT (Table 3) to form specific triggers that are then passed to the FCTS for the final trigger decision. Due to the different latencies associated with the production of these primitives, the GLT forms a time alignment of these input data using configurable delays.

The GLT then forms some additional combined ϕ-maps from the DCT and EMT data. These maps include matched objects such as BM for B tracks matched to an M cluster in ϕ, back-to-back objects, B* and M*, which require a pair of ϕ bits separated by a configurable angle of typically $\sim 120^\circ$, and an EM* object for back-to-back EM pairs.

All 16 ϕ-maps are then used to address individual GLT look-up-tables which return three-bit counts of trigger objects contained within those maps, e.g., the number of B tracks or number of M clusters. To count as distinct trigger objects, the map bits are typically required to have a separation of more than one ϕ bin. The resulting 16 counts plus the IFT hit pattern are then tested in logical operations. The permissible operations include: always-pass; or a comparison ($\geq$, =, or $<$) with a configurable selection parameter. A trigger line is then set as the logical AND of these 17 operations. This process is performed for each of the 24 trigger lines.

The GLT derives the L1 trigger time from the timing information of the highest priority trigger. This time is derived from a centroid based on its 134 ns interval time profile over a period of 1 μs. Other trigger signals compatible with this time are retained and cached. The average time is calculated to the nearest 67 ns and the 24-bit GLT output signal is sent to the FCTS every 67 ns. The achieved timing resolution for hadronic events has an rms width of 52 ns; and 99% of the events are within 77 ns.

The GLT hardware consists of a single 9U VME module. Most of the logic, including diagnostic and DAQ memories, are implemented in FPGA's [4]. The look-up-table section is implemented as an array of 16 memory chips with 8 Mbytes of configuration data.

X.4. Level 1 Trigger Performance and Operational Experience

The L1 trigger configuration consists of DCT-only, EMT-only, mixed and prescaled triggers, aimed not only for maximum efficiency and background suppression, but also for the convenience of trigger efficiency determination.

Although most triggers target a specific physics source, they often also select other processes. For example, two-track triggers are not only efficient for Bhabha, $\mu^+\mu^-$, and $\tau^+\tau^-$ events, but are also useful for selecting jet-like hadronic events and some rare B decays.

The efficiencies and rates of selected L1 triggers for various physics processes are listed in Table 4. Although triggering on generic $B\overline{B}$ events is relatively easy, it is essential to ensure high efficiencies for the important rare low-multiplicity B decays. For this reason, efficiencies for $B^0 \to \pi^0\pi^0$ and $B^- \to \tau^-\overline{\nu}$ are also listed in Table 4.

The efficiencies listed for the hadronic events are absolute and include acceptance losses based on Monte Carlo simulation, and local inefficiency effects. The efficiencies for τ-pair events are for *fiducial* events, *i.e.,* events with two or more tracks with $p_t > 120\,\mathrm{MeV}/c$ and polar angle θ to reach at least DCH superlayer U5. The Bhabha and μ-pair efficiencies are determined from the data, for events with two high momentum particles, back-to-back in e^+e^- center of mass frame, and within the EMC fiducial volume. The data in

Table 4
Level 1 Trigger efficiencies (%) and rates (Hz) at a luminosity of $2.2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ for selected triggers applied to various physics processes. The symbols refer to the counts for each object.

Level 1 Trigger	$\epsilon_{B\overline{B}}$	$\epsilon_{B\to\pi^0\pi^0}$	$\epsilon_{B\to\tau\overline{\nu}}$	$\epsilon_{c\overline{c}}$	ϵ_{uds}	ϵ_{ee}	$\epsilon_{\mu\mu}$	$\epsilon_{\tau\tau}$	Rate
A≥3 & B* ≥1	97.1	66.4	81.8	88.9	81.1	–	–	17.7	180
A≥1 & B* ≥1 & A′ ≥1	95.0	63.0	83.2	89.2	85.2	98.6	99.1	79.9	410
Combined DCT (ORed)	99.1	79.7	92.2	95.3	90.6	98.9	99.1	80.6	560
M≥3 & M* ≥1	99.7	98.6	93.7	98.5	94.7	–	–	53.7	160
EM* ≥1	71.4	94.9	55.5	77.1	79.5	97.8	–	65.8	150
Combined EMT (ORed)	99.8	99.2	95.5	98.8	95.6	99.2	–	77.6	340
B≥3 & A≥2 & M≥2	99.4	81.2	90.3	94.8	87.8	–	–	19.7	170
M* ≥1 & A≥1 & A′ ≥1	95.1	68.8	83.7	90.1	87.0	97.8	95.9	78.2	250
E≥1 & B≥2 & A≥1	72.1	92.4	60.2	77.7	79.2	99.3	–	72.8	140
M* ≥1 & U≥5 (μ-pair)	–	–	–	–	–	–	60.3	–	70
Combined Level 1 triggers	>99.9	99.8	99.7	99.9	98.2	>99.9	99.6	94.5	970

Table 4 demonstrate that the DCT and the combined EMT/IFT provide fully efficient, independent triggers for most physics processes, although independent triggers for $\mu^+\mu^-$ and $\tau^+\tau^-$ are not individually fully efficient. The efficiencies predicted by the Monte Carlo simulation are generally in good agreement with data when tested using events passing typical analysis selections and based on orthogonal triggers. Prescaled triggers with a very open acceptance of physics events, such as (B≥2 & A≥1) or (M≥2) are also used to measure the trigger efficiencies.

The trigger rates listed in Table 4 are for a typical run with HER (LER) currents at 650 mA (1350 mA) and a luminosity of $2.2\times10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. These rates are stable to within 20% for the same PEP-II configuration, but they are impacted by changes in vacuum conditions, beam currents, and orbits. There are occasional background spikes which can double the L1 rate. However, due to the 2 kHz capability of the data acquisition, these spikes do not induce significant dead time.

For a typical L1 rate of 1 kHz, Bhabha and annihilation physics events amount to ∼130 Hz. There are also 100 Hz of cosmic ray and 20 Hz of random beam crossing triggers. The remaining triggers are due to lost particles interacting with the beam pipe or other components. The distribution of single track z_0 values as reconstructed by L3 for all L1 triggers is shown in Figure 5. The most prominent peaks at $z = \pm 20$ cm correspond to a flange of the beam pipe. The peak at $z_0 = -55$ cm corresponds to a step in the synchrotron mask.

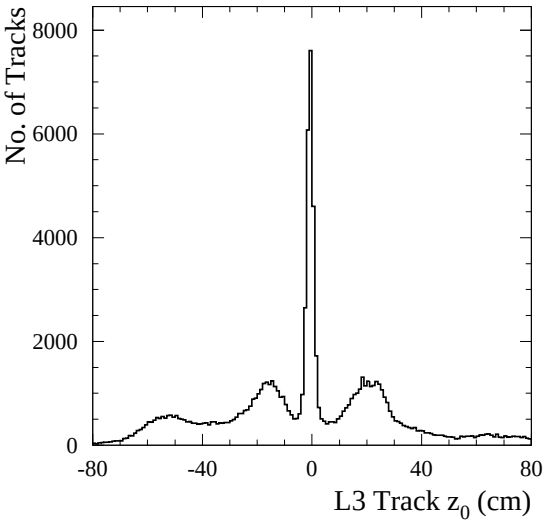

Figure 5. Single track z_0 for all L1 tracks, reconstructed by L3.

The L1 trigger hardware operation has been very stable. For the first one and half years of operation, there have been only four hardware failures in the L1 system, mainly auxiliary or communication boards. Occasional adjustments to the EMT tower mask were used to temporarily suppress noisy channels in the EMC electronics.

X.5. Level 3 Trigger System

The L3 trigger software comprises event reconstruction and classification, a set of event selection filters, and monitoring. This software runs on the online computer farm. The filters have access to the complete event data for making their decision, including the output of the L1 trigger processors and FCTS trigger scalers. L3 operates by refining and augmenting the selection methods used in L1. For example, better DCH tracking (vertex resolution) and EMC clustering filters allow for greater rejection of beam backgrounds and Bhabha events.

The L3 system runs within the Online Event Processing (OEP) framework (see Section **??**). OEP delivers events to L3, then prescales and logs those which pass the L3 selection criteria.

To provide optimum flexibility under different running conditions, L3 is designed according to a general logic model that can be configured to support an unlimited variety of event selection mechanisms. This provides for a number of different, independent classification tests, called *scripts*, that are executed independently, together with a mechanism for combining these tests into the final set of classification decisions.

The L3 trigger has three phases. In the first phase, events are classified by defining L3 input lines, which are based on a logical OR of any number of the 32 FCTS output lines. Any number of L3 input lines may be defined.

The second phase comprises a number of scripts. Each script executes if its single L3 input line is true and subsequently produces a single pass–fail output flag. Internally, a script may execute one or both of the DCH or EMC algorithms, followed by one or more filters. The algorithms construct quantities of interest, while the filters determine whether or not those quantities satisfy the specific selection criteria.

In the final phase, the L3 output lines are formed. Each output line is defined as the logical OR of selected script flags. L3 can treat script flags as vetoes, thereby rejecting, for example, carefully selected Bhabha events which might otherwise satisfy the selection criteria.

L3 utilizes the standard event data analysis framework and depends crucially on several of its aspects. Any code in the form of *modules* can be included and configured at run time. A sequence of these software modules compose a script. The same instance of a module may be included in multiple scripts yet it is executed only once, thus avoiding significant additional CPU overhead.

X.5.1. Level 3 Drift Chamber Tracking Algorithm

Many events which pass L1 but must be rejected by L3 are beam-induced charged particle background that are produced in material close to the IP. L1 does not currently have sufficient tracking resolution to identify these background tracks. The DCH-based algorithm, L3Dch, performs fast pattern recognition (track finding) and track fitting, which determines the five helix track parameters for tracks with $p_{\rm t}$ above 250 MeV/c. To speed up the process of pattern recognition, L3Dch starts with the track segments from the TSF system and improves the resolution by making use of the actual DCH information.

For those TSF segments that have a simple solution to the left-right ambiguity, a track t_0 is determined. The t_0 values for each segment in an event are binned and the mean produced from the values in the most populated bin is used as the estimated event t_0. All events which pass L1 typically have enough segments to form a t_0 estimate. The measured rms resolution on this estimate is 1.8 ns for Bhabha events and 3.8 ns for hadronic events.

The pattern recognition for L3Dch is done with a look-up-table. For this track table, the DCH is divided into 120 ϕ-sectors, corresponding to the number of cells in the innermost layers. The track table is populated with the hit patterns of Monte Carlo generated tracks with a $p_{\rm t}$ above 250 MeV/c and originating within 2 cm of the IP in the x–y plane, and within 10 cm in z. The pattern recog-

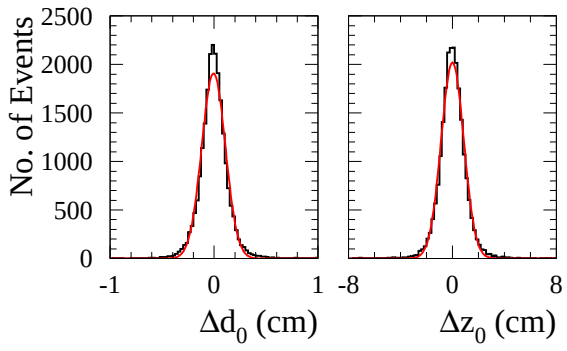

Figure 6. Transverse and longitudinal miss distances between the two tracks in Bhabha events.

nition algorithm searches the table entries looking for matches to segments found by the TSFs. The matched set of segments for a given track is then passed to the track fitting algorithm. The track table allows for up to two missing DCH TSF segments per track.

The track fitting algorithm is provided with both the track segments found in pattern recognition and the individual hits within those segments. From this information the five helix parameters are fitted. The fit is then iterated, adding segments close to the initially fitted track, and dropping hits with large residuals. The final fit does not demand that the track originate from the IP.

The two-track miss distances for Bhabha events are plotted in Figure 6. The resolutions for individual tracks are 0.80 mm and 6.1 mm for d_0 and z_0, respectively. Similarly, the $1/p_{\rm t}$ difference between the two tracks in μ-pair events yields a $p_{\rm t}$ resolution of $\delta p_{\rm t}/p_{\rm t} \sim 0.019 \cdot p_{\rm t}$, with $p_{\rm t}$ in GeV/c.

X.5.2. Level 3 Calorimeter Clustering Algorithm

The all-neutral trigger for L3 is based on information from the EMC. In addition, calorimeter information is a vital complement to the DCH data for the identification of Bhabha events.

The L3 EMC-based trigger, L3Emc, identifies energy clusters with a sensitivity sufficient for finding minimum ionizing particles. EMC data are processed in two steps: first, lists of crystals with significant energy deposits are formed; and second, clusters are identified. The EMC typically sends data for ∼1400 crystals (of 6580 total), the majority of these are caused by electronics noise and beam-induced background. For each crystal, these data include the peak energy and time of the crystal waveform. To filter out noise, L3Emc rejects individual crystal signals below an energy threshold of 20 MeV or which lie outside a 1.3 μs time window around the event time. For the remaining crystals, raw energies and times are converted into physical units and added to the L3Emc crystal list. Clusters are formed using an optimized look-up-table technique requiring only a single pass over the crystal list. Clusters with a total energy above 100 MeV are retained, and the energy weighted centroid and average time, the number of crystals, and a lateral moment describing the shower shape for particle identification are calculated.

X.5.3. Level 3 Filters

Based on the L3 tracks and clusters, a variety of filters perform event classification and background reduction. The logging decision is primarily made by two orthogonal filters, one based exclusively on DCH data and the other based only on EMC data.

The drift chamber filters select events with one *tight* (high $p_{\rm t}$) track or two *loose* tracks originating from the IP, respectively. To account for the fact that the IP is not exactly at the origin, track selection is based on its x–y closest approach distance to the IP, d_0^{IP}, and z_0^{IP}, the corresponding z coordinate for that point. The IP position is a fixed location close to the average beam position over many months. The high $p_{\rm t}$ track is required to have a transverse momentum of $p_{\rm t}>600$ MeV/c and to satisfy a vertex condition defined as $|d_0^{IP}| < 1.0$ cm, and $|z_0^{IP} - z_{IP}| < 7.0$ cm. Two tracks are accepted with $p_{\rm t}>250$ MeV/c and a somewhat looser vertex condition defined as $|d_0^{IP}| < 1.5$ cm, $|z_0^{IP} - z_{IP}| < 10.0$ cm.

Two calorimeter cluster filters select events with either high energy deposits or high cluster multiplicity. Each filter also requires a high effective mass calculated from the cluster energy sums

and the energy weighted centroid positions of all clusters in the event assuming massless particles. The first filter requires at least two clusters of $E_{CM} > 350\,\mathrm{MeV}$ (c.m. system energy) and event mass greater than 1.5 GeV; the second filter requires at least four clusters, and an event mass greater than 1.5 GeV.

At current luminosities, the output of both the DCH and EMC filters is dominated by Bhabha events, which need to be rejected. This is accomplished by a Bhabha veto filter that selects one-prong (with only a positron in the backward part of the detector) and two-prong events (with both e^+ and e^- detected). Stringent criteria on EMC energy deposits are imposed, relying on the track momenta and on E/p. The two-prong veto requires either colinearity between the tracks in the c.m. system or an acolinearity that is consistent with initial state radiation (ISR).

For purposes of calibration and offline luminosity measurements, Bhabha, radiative Bhabha, $\gamma\gamma$ final state, and cosmic ray events are flagged. The output rate of flagged Bhabha events is adjusted to generate an approximately flat distribution of events in polar angle. Radiative Bhabha events are identified by selecting two-prong events with missing energy and requiring an EMC cluster in the direction of the missing momentum. Events with two high energy clusters, back-to-back in the c.m. system select the $e^+e^- \to \gamma\gamma$ process. The cosmic ray selection is DCH-based and requires two back-to-back tracks in the laboratory frame with nearly equal impact parameters and curvature. A significant background from ISR Bhabha events faking this topology is removed using the same kinematic constraints used in the two-prong veto.

The online luminosity monitoring and energy scale monitoring are performed in L3. A track-based lepton-pair selection with a well known efficiency monitors the luminosity. Hadronic filters for selection of continuum and $B\overline{B}$-enriched samples monitor the energy scale. The latter two categories are distinguished by an event shape selection using a ratio of Fox-Wolfram moments [**?**]. The ratio of the $B\overline{B}$-enriched sample to the luminosity is a sensitive measure of relative position on the $\Upsilon(4S)$ peak and thereby monitors the beam energies.

Table 6
Composition of the L3 output at a luminosity of $2.6\times10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Event type	Rate (Hz)
Hadrons, $\tau\tau$, and $\mu\mu$	16
Other QED, 2-photon events	13
Unidentified Bhabha backgrounds	18
beam-induced backgrounds	26
Total physics accept	**73**
Calibration Bhabhas (e^+e^-)	30
$\gamma\gamma$, Radiative Bhabhas ($e^+e^-\gamma$)	10
Random triggers and cosmic rays	2
L1,L3 pass through diagnostics	7
Total calibration/diagnostics	**49**

X.6. Level 3 Performance and Operational Experience

The L3 trigger efficiency for Monte Carlo simulated events are tabulated in Table 5 for events passing Level 1. High efficiencies are independently achieved for the DCH and EMC based filters applied to simulated hadronic events. The comparison between data and Monte Carlo L3 trigger pass fractions for the various filters also show good agreement when requiring tracking, and EMC based hadronic event selections in turn.

An example of the event display used for online trigger monitoring is shown in Figure 7. L3 reconstructed tracks and EMC clusters are shown together with the L1 and L3 trigger line states for the event. The left column lists the L1 trigger lines and their states: on (1); off (0); or on but ignored due to prescale factor (-1). The right column shows the same information for the L3 trigger lines.

For a typical run on the $\Upsilon(4S)$ peak with an average luminosity of $2.6\times10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the L3 event composition is tabulated in Table 6. The desired physics events contribute 13% of the total output while the calibration and diagnostic samples comprise 40%.

The L3 executable currently takes an average processing time of 8.5 ms per event per farm computer. A Level 1 input rate of 2700 Hz saturates

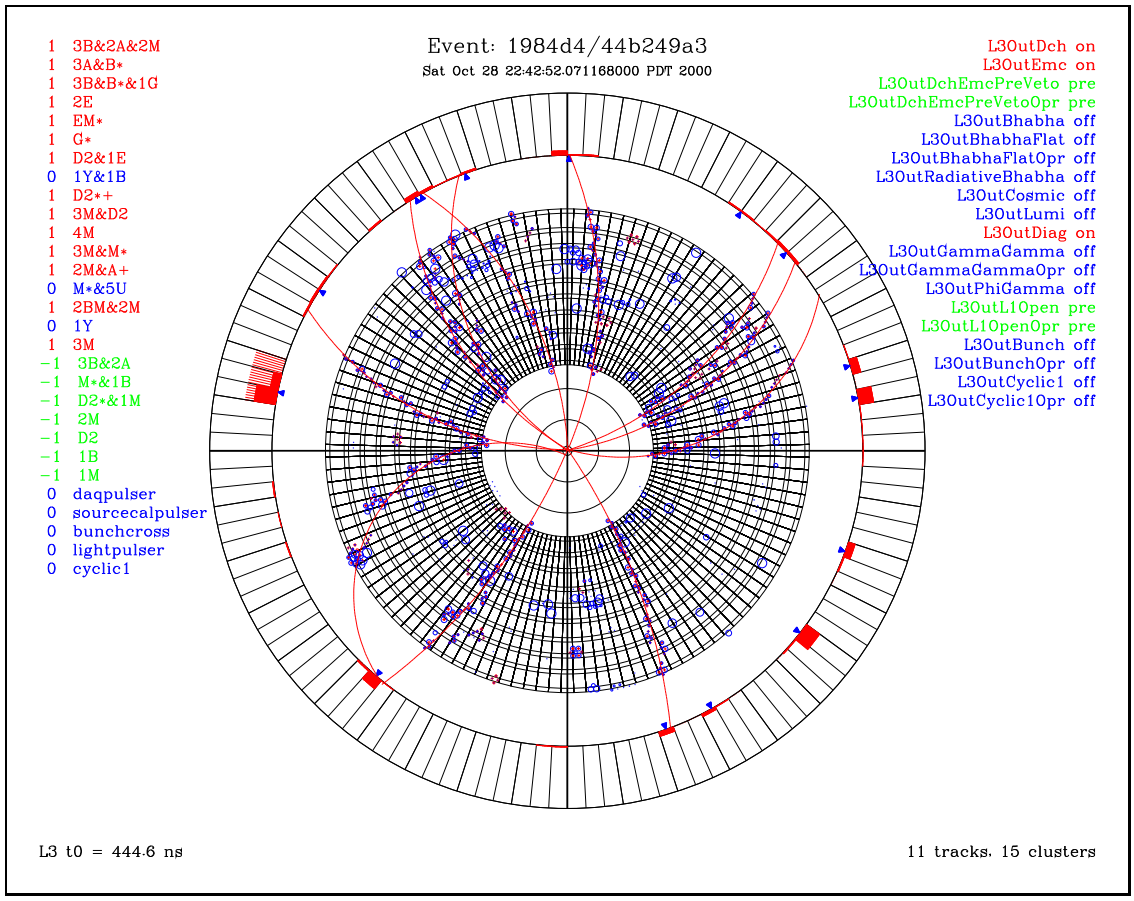

Figure 7. A Level 3 event display. The small circles and small crosses in the DCH volume are DCH hits and TSF segment hit wires respectively. The filled EMC crystals represent energy deposit (full crystal depth=2 GeV) from Level 3 EMC clusters while the small triangles just inside the EMC indicate the location of the cluster centroid.

the Level 3 processors, well above the 2 kHz design requirement. At this input rate the L3 process consumes ∼72% of the CPU time, the rest is spent in OEP, including the network event builder and in the operating system kernel.

X.7. Summary and Outlook

Both the L1 and L3 trigger systems have met their original design goals at a luminosity of $3 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The triggering efficiencies for $B\overline{B}$ events generally meet the 99% design goal for both L1 and L3. The orthogonal triggers based on DCH-only and EMC-only information have successfully delivered stable and measurable overall trigger efficiency. The current system also provides a solid foundation for an upgrade path to luminosities of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ or more.

Short-term L1 trigger improvements will primarily come from further background rejection, afforded by algorithm refinements and upgrades of the DCT. This is essential for reducing the load on the DAQ and L3. The new PTD algorithm will

Table 5
L3 trigger efficiency (%) for various physics processes, derived from Monte Carlo simulation.

L3 Trigger	$\epsilon_{B\overline{B}}$	$\epsilon_{B\to\pi^0\pi^0}$	$\epsilon_{B\to\tau\nu}$	$\epsilon_{c\overline{c}}$	ϵ_{uds}	$\epsilon_{\tau\tau}$
1 track filter	89.9	69.9	86.5	89.2	88.2	94.1
2 track filter	98.9	84.1	94.5	96.1	93.2	87.6
Combined DCH filters	99.4	89.1	96.6	97.1	95.4	95.5
2 cluster filter	25.8	91.2	14.5	39.2	48.7	34.3
4 cluster filter	93.5	95.2	62.3	87.4	85.5	37.8
Combined EMC filters	93.5	95.7	62.3	87.4	85.6	46.3
Combined DCH+EMC filters	>99.9	99.3	98.1	99.0	97.6	97.3
Combined L1+L3	>99.9	99.1	97.8	98.9	95.8	92.0

effectively narrow the track d_0 acceptance window, while new BLT algorithm will narrow the track z_0 acceptance.

For the longer term future, a major DCT upgrade is planned. By adding the stereo layer information, a z_0 resolution of 4 cm is expected, allowing for an efficient rejection of beam-induced background beyond $z = \pm 20$ cm.

Future improvements for L3 will also emphasize background rejection. Improvements in the L3 IP track filter are expected to further reduce beam-induced background to about one third of current levels. The physics filter algorithms will be tuned and improved, primarily for rejecting Bhabha, QED, and two-photon events. Improvements in the L3 tracking algorithms are expected to lower the $p_{\rm t}$ thresholds below 250 MeV/c. A moderate CPU upgrade for the L3 online farm will be sufficient to keep up with luminosities of $\sim 10^{34}\,{\rm cm}^{-2}{\rm s}^{-1}$.

REFERENCES

1. A. Berenyi *et al.*, "Continuously Live Image Processor for Drift Chamber Track Segment Triggering", IEEE Trans. Nucl. Sci. 46 (1999) 348.
2. A. Berenyi *et al.*, "A Binary Link Tracker for the *BABAR* Level 1 Trigger System", IEEE Trans. Nucl. Sci. 46 (1999) 928.
3. A. Berenyi *et al.*, "A Real-Time Transverse Momentum Discriminator for the *BABAR* Level 1 Trigger System", submitted to IEEE Transactions on Nuclear Science.
4. The FPGA's used in the DCT and GLT are from ORCA 2C series manufactured by Lucent Technologies.
5. K. Kinoshita, Nucl. Instr. and Methods A276 (1989) 242.
6. P. D. Dauncey *et al.*, "Design and Performance of the Level 1 Calorimeter Trigger for the BaBar Detector", submitted to IEEE Transactions on Nuclear Science.
7. The FPGA's used for EMT algorithm logic are Xilinx 4020E.