ONLINE EVENT FILTERING IN THE JADE DATA ACQUISITION SYSTEM

Howard E. MILLS

University of Manchester, England

The data acquisition system developed for the JADE experiment at PETRA, DESY includes the facility to use software to filter out background events. The design, implementation, testing and experience gained are discussed.

1. Introduction

The JADE detector [1], installed at the PETRA e^+e^- storage ring at DESY, Hamburg, is shown in a sectional view in fig. 1. It consists of a cylindrical drift chamber inside a 0.48 T solenoid. Between the drift chamber and the solenoid are time-of-flight (TOF) counters; outside the solenoid there is an array of lead glass blocks for the electromagnetic shower measurement, and further drift chambers interleaved with absorbers for muon detection (the muon filter). With up to 240 words (of 16 bits) per track, the central drift chamber supplies the largest amount of data. The average event length varies between 2000 and 3500 words: multihadronic events, which occur at a maximum rate of a few per hour, are longer – typically between 4000 and 8000 words. The trigger rate varies between 2 and 6 Hz depending upon the beam conditions. The experiment commenced data taking in Spring 1979.

The JADE Data Acquisition System (JDAS) is described in detail elsewhere [2]. JDAS is implemented on a Norsk-Data NORD-10S/50 dual processor [3] and a Plessey MIPROC-16 microprocessor [4]. This paper describes a further important feature of the online analysis that includes routines to filter out events which are clearly identified as background.

2. Operation and data flow

It is necessary to describe the hardware trigger system and data flow in order to understand the origin of background events and why it is desirable to filter events online.

The electron and positron bunches within the PETRA machine cross at the interaction region surrounded by the detector, approximately every 4 μs. The JADE trigger system [5] reduces this 250 kHz rate to a few per second. For normal operations there are many different trigger sources which request that an event be read out whenever something of interest may have occurred. The trigger system is a multiple level trigger. The first level (T1) decision is based upon the fast analogue signals from the lead glass array which are ready about 350 ns after the beam crossing. The various "T1 accept triggers" are based upon the amount and configuration of energy deposited in the lead glass array.

The second level of the trigger (T2) is based upon the data from the inner drift chamber. If there is no T1 accept condition, but certain combinations of the TOF and/or lead glass counters have been hit, the decision is postponed to the T2 stage of the trigger, otherwise the trigger electronics are reset. In the "postpone" state, the trigger system waits for the inner detector drift times and the results of the track logic (about 2.5 μs). If enough hits have been generated by charged particles passing through the inner detector to satisfy the various T2 track logic requirements then the event is accepted.

If there is no T2 accept condition, but the preconditions for a muon trigger are set, then the next stage of the trigger (T3) is invoked, otherwise the trigger is reset. The T3 level trigger waits for the muon chamber drift times (about 4 μs). If the T3 track logic is satisfied then the event is accepted otherwise the trigger electronics are reset. There are currently 10 "T1 level", 10 "T2 level" and 1 "T3 level" triggers. For a given event a single or combination of trigger bits will have been set. With this system, detectable collisions between electrons and positrons occur at a rate of a few per minute (for a luminosity of 3×10^{30} cm^{-2}s^{-1}). If any of these triggers is set an interrupt is sent to the computer.

The NORD-10S performs tasks such as handling interrupts, CAMAC readout, data logging via a link to the central DESY computer or to magnetic tape, communication with the operator and event and histogram display. The MIPROC-16 is used for part of the online event filtering scheme. The NORD-50 processor is used for a single program which performs event validation,

MAGNETDETEKTOR JADE
MAGNET DETECTOR

1 **Strahlrohrzähler** BEAM-PIPE COUNTERS
2 **Endseitige Bleiglaszähler** END PLUG LEAD GLASS COUNTERS
3 **Drucktank** PRESSURE TANK
4 **Myon-Kammern** MUON CHAMBERS
5 **Jet-Kammern** JET CHAMBERS
6 **Flugzeit-Zähler** TIME OF FLIGHT COUNTERS
7 **Spule** COIL
8 **Zentrale Bleiglaszähler** CENTRAL LEAD GLASS COUNTERS
9 **Magnetjoch** MAGNET YOKE
10 **Myon-Filter** MUON FILTERS
11 **Beweglicher Endstopfen** REMOVABLE END PLUG
12 **Strahlrohr** BEAM PIPE
13 **Vorwärts-Detektor** TAGGING COUNTER
14 **Mini-Beta Quadrupol** MINI BETA QUADRUPOLE
15 **Fahrwerk** MOVING DEVICES

Gesamtgewicht TOTAL WEIGHT: ~1200 t
Magnetfeld MAGNETIC FIELD: 0.5 T

W.Knauf/DESY
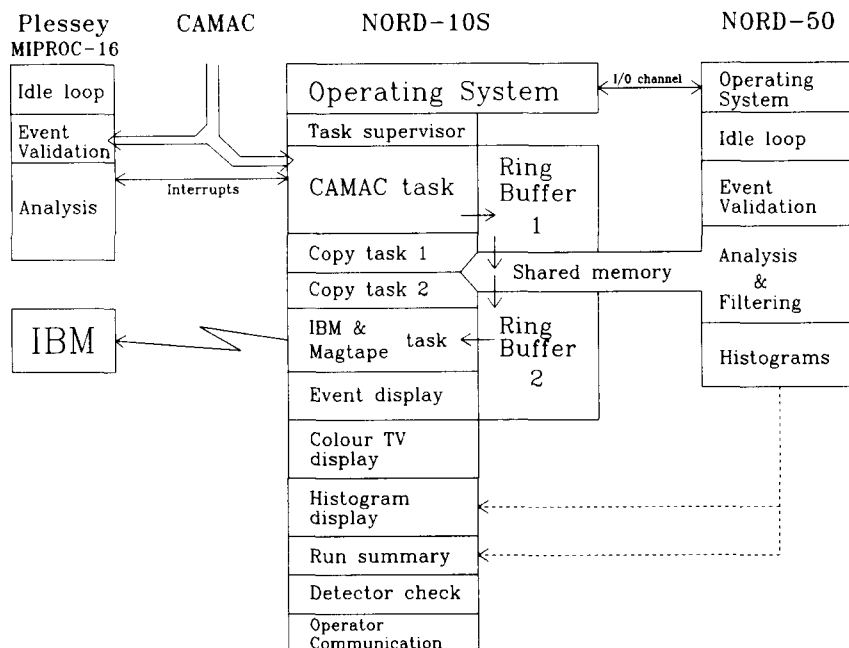
39377

Fig. 1. Sectional view of the JADE detector.

Fig. 2. Organization of JDAS tasks and buffers.

analysis and monitoring, which includes histogram filling.

The organisation of tasks and buffers within JDAS is shown in fig. 2. When a trigger occurs, the event is read out from CAMAC into the Ring Buffer 1 in the NORD-10S. When the event is complete it is passed to the NORD-50 where the analysis program resides. This program first validates the data and reports any serious faults to the operator. The event is then analysed and monitored with histograms being filled and the event classified as one of the final states electron or muon pair, 2-photon, multihadronic, beam–gas or cosmic event etc. The results of the analysis are appended to the event. If the event is not marked for rejection it is passed to the Ring Buffer 2 within the NORD-10S and hence to the central DESY IBM computer or to local magnetic tape. At the central computer the data are stored temporarily on disc and then dumped to magnetic tape when a sufficient amount has accumulated. A number of dump tapes are combined to obtain efficiently filled tapes with complete runs. Each run of 8000 events lasts about an hour: thus in a typical day JADE collects approximately 160000 events, equivalent to seven 2400 foot 6250 b.p.i. magnetic tapes. However only a few thousand of these events come from e^+e^- collisions and are useful for physics. It is therefore clear that a considerable saving in magnetic tapes, link traffic and offline computer time can be achieved if a significant fraction of background events can be safely removed online.

A brief description of the offline processing is given for later reference. It commences with the REFORM job which converts the data to a form more suitable for analysis. The REDUC1 filter program [6] performs the first stage data reduction to remove clear background events. After this stage about 30 to 40% of the original events remain. The REDUC1 output tapes then pass through further selection programs (REDUC2 etc.) to extract the different types of physics events. The data at the REDUC1 stage still contain a considerable proportion of background events which have similar topologies to genuine events.

3. Special considerations for online event filtering

Experience has shown that a proposal to reject events by software within the data acquisition system of an HEP experiment encounters considerable mistrust from colleagues. One argument against filtering out events online is that one does not have a second try if a mistake is later discovered, whereas an offline data reduction program can always be corrected and rerun. However, there is a great reluctance to rerun the JADE REDUC1 filtering program due to the time required and the tape handling involved. A further problem might be if there were some unexpected physics processes that could be rejected by online filtering.

Several current HEP experiments have the capability of filtering events online, but in general these have only

run in a passive mode, marking the events but not actually rejecting them. With JDAS it was foreseen from the start that at an early stage events would be filtered out.

4. Background events

Interactions between electrons and positrons of the PETRA bunches occur at a low rate, since at high energies the cross-section is small. Detectable Bhabha scattering and two-photon events occur at the rate of a few hundred per hour; muon pairs and multihadronic events at a maximum rate of a few per hour. Some rare events may only occur once per year. The typical trigger rate of 2–6 Hz is caused by background events of various types which satisfy the trigger conditions. There are several types of background:

Cosmic particles   which pass through the detector at the same time as the bunches cross and hence may set the trigger ($\sim$ 1 Hz rate).

Beam–gas events   where an electron (or positron) within a bunch collides with a gas molecule within the beam pipe. The beam pipe is pumped down to a very low pressure but some gas remains. When a gas molecule is hit particles can be scattered into the detector (this and the next type of event form the great majority of the triggers).

Beam–wall events   where an off-momentum particle from the beam collides with the synchrotron radiation absorbers attached inside the wall of the beampipe, sending one or more particles into the detector.

Spurious triggers   where signals induced by the beam cause random hits which in turn can cause the track triggers to be set ($\sim$ 0.5 Hz). Synchrotron radiation hits in the inner detector can cause the track trigger logic to be set in this way.

Electronic noise   where electronics within the detector malfunction and cause spurious triggers. In general the rate from such sources is very low but large fluctuations can occur.

The objective of the online event filtering within JDAS to remove as many background events as possible that can be positively identified as not being due to e^+e^- collisions.

5. General method

Within JDAS several event-filtering algorithms have been developed which are applied to events of a particular trigger or group of triggers. The steps involved in developing and implementing such an algorithm are:

1) The algorithm is developed on the central DESY IBM computer. It is then tested on tapes of raw data in order to estimate how many events the algorithm would reject. Then it is tested on REDUC1 (or sometimes REDUC2) data in order to find out how many events are rejected that should not be. Any such events are written to a file and scanned visually. Finally the algorithm is tested on physicist-selected events where none should have been marked for rejection. The algorithm should be as fast as possible – there is little advantage in having a slow algorithm that causes the system to lose events because of increasing dead time.

2) The algorithm is transferred to the online computer system. Once implemented the algorithm is allowed to run for a suitable period where events are marked but are not actually rejected. The events are then checked to ensure that the results agree with those of the offline analysis.

3) The event filtering is activated. The JDAS control program provides a menu to switch on and off the various rejection processes. The number of events rejected is recorded on the run summary sheet that is automatically printed at the end of each data acquisition run. An important feature is that of each category of events to be filtered out, 5% are kept and transferred to the IBM with the rest of the data for later tests. The reason for rejection is appended to the event. Thus if a physicist finds that an event in his final selection has the reject flag set, then he can calculate a correction factor and the problem can be studied and – if possible – corrected.

4) At intervals the "5%" events are checked offline.

6. Processors used for event filtering in JADE

JDAS has two processors which are used to filter out events – the MIPROC-16 and the NORD-50. There is no fundamental reason why a particular filtering algorithm should be implemented in one processor rather than the other. Each has its advantages and drawbacks.

The MIPROC-16 has to be programmed in assembly language. The cross-assembler for the MIPROC-16 runs on the NORD-10S and the resultant program is downloaded. This processor is very fast (200 ns instruction time) and, since it is closely coupled to the readout, any events that are filtered out here do not have to be passed to the NORD-50 for analysis. Hence the effect of event filtering in the MIPROC-16 is to reduce the

deadtime involved in the whole system as well as to reduce the number of events sent to the central IBM for storage and further processing. The lower rate of events through the NORD-50 also permits longer analysis times there.

The NORD-50 is programmed in FORTRAN 77 which makes it far easier to adapt algorithms developed on the IBM. Since the NORD-50 is further along the data flow chain than the MIPROC-16, its effect on possible reduction of deadtime is marginal. However the reduction in volume of output data is valuable.

Note that when memory requirements are mentioned that this does not include code required to unpack and validate the relevant data. This code is required anyway for the validation and monitoring. The complete NORD-50 code and local storage requires around 192 kbytes.

Generally the simplest and most effective algorithms should be installed in the MIPROC-16. Historically the earliest developed algorithms were placed in the MIPROC-16 and the later algorithms into the NORD-50.

The algorithms that are implemented in JDAS are now discussed.

7. Filtering Algorithms in the MIPROC-16

The two event filtering algorithms installed in the MIPROC-16 are designed to clean up the track triggers – the T2 level accept events. These events are defined as having one or more charged tracks which penetrate all three rings of the inner detector and hit a TOF counter. Different T2 accept triggers – there are currently 8 – demand different numbers and configurations of tracks. T1 and T3 level accept events are ignored by the MIPROC-16.

As the trigger and inner detector data are read by the NORD-10S from the CAMAC they are also strobed into the MIPROC-16. The inner detector is divided into three rings and the outer ring (ring 3) is read in first. As the data flow in, the MIPROC-16 starts the validation process. Any fault found during validation inhibits rejection of the event. Once the ring 3 readout is complete, the first algorithm is applied. The T2 trigger hardware counts the number of hits within sectors of the inner detector. If this number exceeds a threshold, then a track candidate has been found. Whilst genuine tracks (above a certain transverse momentum) satisfy the trigger, spurious hits caused by pickup, synchrotron radiation or other photons can also set the trigger. The software also has access to the drift time of each hit and hence can determine if the trigger was set by a genuine track. The method (described in detail elsewhere [7]) is to look for hit triples with similar drift times and to histogram the time differences of successive hits. A peak in the histogram is counted as a track. This is performed for those cells of the inner detector where the T2 trigger has marked a track. The number of found tracks in the event is then compared with the number originally required by the trigger and if it is no longer satisfied, the event is a candidate for rejection.

Fig. 3 is an $r-\phi$ view (orthogonal to the beam axis)

Fig. 3. $r-\phi$ view of a T2 event with a "false" track. Each hit is shown twice – the genuine hit and its mirror hit due to left/right ambiguity.
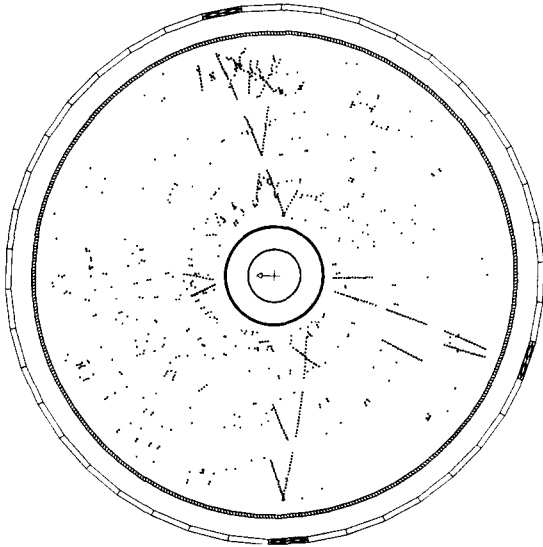
Fig. 4. r-φ view of a beam-wall event.

of such an event. The three rings of the inner detector are shown together with the TOF counters. It is a typical T2 level accept event where a trigger for tracks collinear in r-φ has been set. One genuine track exists in the lower left-hand quarter of the inner detector (in the illustrations the left-right mirrors of each hit are displayed). The second, spurious, track is in the upper right-hand quarter of the inner detector. The eye can see that no genuine track exists here – the hits may be caused by synchrotron radiation and scattered electrons. The algorithm detects only the single track – the trigger demands two – and hence the event is rejected.

The second stage is a quick determination of the event Z-vertex [8] (interaction point along the beam line). The technique is to combine all hits in drift chamber layers 1 to 16 (ring 1) with corresponding hits in layers 17 to 32 (ring 2) and to form from any such two-hit combination the intercept or closest approach to the beam line (in the Z direction). All these positions are accumulated into a histogram, and if an event has a vertex it should show up as a cluster or peak. This is called the Z-vertex. Any event whose closest peak to the centre of the detector (which is coincident with the interaction point) is at least 300 mm away is classified as a background event and rejected. In the offline analysis with better calibration constants the Z-vertex cut is reduced to ± 200 mm.

Figs. 4 and 5 show an event rejected by this method. Fig. 4 shows the r-φ view. Three tracks can be seen which are clearly originating from the wall of the beampipe rather than the centre. Fig. 5 shows the z-y view (parallel to the beam axis) where the tracks * can be seen on the left-hand side of the detector. The MIPROC-16 has found a good Z-vertex at −1080 mm (marked by the cross) and hence the event is a candidate for rejection. The event can be explained as an off-momentum particle in the positron beam striking the beampipe and causing the tracks. Figs. 6 and 7 show another event that is rejected – it is due to a cosmic particle passing through the detector clearly outside the cuts (Z-vertex at −680 mm).

The MIPROC-16 rejects about 55% of the track triggers – the precise fraction depends on the beam conditions – and since the original triggers consist of

---

* In all figures, tracks which are drawn as heavy lines with arrows have been drawn by hand. Tracks drawn as fine lines are computer fits.


Fig. 5. z-y view of the event in fig. 4.

Fig. 6. $r-\phi$ view of a T2 cosmic event.

approximately equal numbers of track (T2) and energy (T1) triggers, about a quarter of all events are rejected by the MIPROC-16.

8. Filtering algorithms in the NORD-50

The filtering algorithms that are active in the MIPROC-16 were also installed in the NORD-50 for use in case the MIPROC-16 malfunctioned.

The first additional filtering algorithms installed in the NORD-50 involved T1 level accept triggers.

9. Low energy neutral trigger – cosmic particles

One of the low energy neutral triggers requires two low threshold energy deposits in the lead glass barrel (with a veto if any TOF counter is hit). These would be due to a reaction producing at least two photons and no detectable charged particles. It is used to study all-photon final states in photon–photon collisions. The trigger is heavily contaminated by cosmic particles outside the TOF timing gate but within that of the lead glass system. If an event has this low energy neutral trigger set and no other trigger, then the muon filter is searched for tracks that would be left by a cosmic particle. The presence of such a track causes the event to be flagged for rejection. A typical example is shown in fig. 8. The two clusters in the lead glass barrel can be seen as well as the hits in the muon filter. The numbers displayed are the block energies summed along the rows of the lead glass barrel. The method [9] is to search through the muon filter in a scheme allied to that used by the T3 trigger hardware. This scheme divides the muon filter into overlapping "streets" pointing towards the interaction region. The algorithm ORs hits within each chamber layer together and then counts how many layers within each street have been hit. The innermost layer is ignored (because of frequent additional hits due to other causes). A track in the muon filter is defined to have 3 out of 4 layers within a street hit (the endwalls must have 4 out of 4 due to noise problems). The actual T3 trigger readout cannot be used for two reasons. In order that the T3 trigger does not miss genuine muons, some layers within streets can be permanently switched "ON" if one or more chambers are not functioning. This can result in false "tracks" which is perfectly


Fig. 7. $z-y$ view of the event in fig. 6.

Fig. 8. Low energy neutral trigger – cosmic particle.

correct for a hardware trigger, but disastrous for a filtering algorithm. The second reason is that good tracks would be missed due to the fact that the T3 gate is shorter than the muon digitiser gate. Hits can be present in the muon chambers but not in the T3 trigger. This does not cause problems for in-time beam-beam events, such as muon pairs where the T3 level trigger is fully efficient, but for out-of-time cosmic particles (which we wish to filter out) the T3 trigger is inefficient – which again is reasonable for a hardware trigger. Hence analysing the actual chamber hits rather than the T3 trigger readout will give a better rejection factor. It is realised that a cosmic particle passing through the detector at the same time as a genuine 2-photon event could cause that event to be rejected. This is acceptable for physics reasons since it occurs at a low rate and can be corrected for.

This algorithm is fast and does not require much memory (4 kbytes). The result is that around 85% of these triggers are rejected online. The ratio of this trigger to all triggers varies – in recent times it has been around 18%. Thus it can be seen that this algorithm has a significant effect on the volume of data and it is considered to be the only way that the low energy neutral trigger can be used, since otherwise its effect would be to swamp the data acquisition system. The events that are retained by the algorithm do still contain a number of cosmic particles and these have to be cleaned up during offline analysis. The algorithm has been used offline to help analyse data taken before it was activated online.

More recently additional triggers for γγ physics have been added. The same cosmic filtering algorithm is applied to them.

10. Lead glass barrel energy trigger – cosmic particles

The lead glass barrel energy trigger is formed by summing the lead glass analogue signals. If the sum exceeds a threshold (normally 3.5 GeV) this causes a T1 accept trigger. The endcaps have a similar trigger and there is an overall lead glass energy trigger which is the sum of the barrel and endcap sums. The barrel energy trigger is set by multihadronic events, electron pairs and events containing high energy photons. It is contaminated by cosmic particles.

If a cosmic particle passes near the centre of the detector in the r–ϕ plane the energy deposited in the lead glass barrel is usually insufficient to set the energy trigger – it may cause a T2 trigger which is dealt with elsewhere in this paper. If the path of a cosmic particle is offset so that it clips the barrel, it will pass through several consecutive blocks and the sum of the energies may exceed the threshold. An example of such an event is shown in figs. 9 and 10. Fig. 9 shows the cosmic particle passing through the muon filter, clipping the lead glass barrel (the numbers are summed energies in each row) and passing out through the muon filter. Fig. 10 shows an r–ϕ perspective view of the lead glass barrel where the characteristic long thin cluster can be seen. An algorithm has been devised to detect and filter out such events.

A scheme had already been developed to identify such events during offline analysis [10]. This was a modification of the standard lead glass cluster finding algorithm which had been extended to try to detect long thin clusters. In the standard analysis, clusters are made up of the hit lead glass blocks that are connected by their edges. The "grazing" clusters consist of blocks that


Fig. 9. Lead glass barrel energy trigger – cosmic particle.

Fig. 10. $r - \phi$ perspective of the event in fig. 9.

are connected by their edges or at the corners. An event is permitted to have either one or two such clusters – two would occur if the cosmic particle passed just inside the barrel or if a block failed to register a hit. The cluster or clusters are checked to ensure that the ratio of the length to the width is sufficiently large (so as to avoid selecting the more circular clusters due to photons originating from the interaction region). The offline algorithm was modified slightly, tested and transferred to the NORD-50. It is applied to events where the

barrel energy trigger has been set. The total energy trigger (lead glass barrel and endcaps > 6 GeV) is also permitted to be set but if any other T1 accept trigger is set the event is not considered for rejection. If the event also has a T1 postpone condition and there is a well-defined Z-vertex (as would be made by a charged track) then the event is retained.

The algorithm requires a moderate amount of computer memory (11 kbytes) and processing time. The time taken increases at a rate faster than proportional to the number of blocks hit. Experience has shown that if a large number of blocks have fired, due to electronic pickup for example, then the algorithm requires an unacceptable amount of time. To guard against this only events with less than 50 hit blocks are treated. Tests have shown that this limit does not cause any reduction in the efficiency of rejection. The overall result is that about 50% of the events caused by the barrel energy trigger are rejected (corresponding to 300–500 events per run). Fig. 11 shows a projection parallel to the beam of an event rejected by this method. It appears to be an almost horizontal cosmic event. The particle can be seen entering through an endwall of the muon filter, passing along a row of blocks in the lead glass barrel and leaving through the opposite end of the muon filter.

11. Lead glass energy trigger – noise

The detector has suffered from occasional sparking and noise problems in the lead glass system. This can


Fig. 11. Barrel energy trigger – horizontal cosmic.

Fig. 12. "Big block" barrel lead glass event.

occur as a single block giving a high energy reading which causes the total energy trigger to be set. The problem can occur in the endcaps or in parts of the barrel. The normally low rate of this noise means it is very difficult to locate the problem at source. Fig. 12 shows an event with a single barrel block with about 6.5 GeV energy. Fig. 13 shows the same event viewed from the side. Fig. 14 shows an event with a single large (16.8 GeV) block in the endcap. To identify and filter out such events online the following algorithm was adopted.

For events where the total energy triggers are set, and no others, the contents of the lead glass blocks are scanned. If there is a single block with more than 1000 ADC counts, which corresponds to just over 6 GeV, and the sum of the remaining blocks is less than 50 counts (300 MeV) then the event is probably due to noise. If a lead glass block has a genuine hit of around 6


Fig. 13. $z - y$ view of the event in fig. 12.

GeV then the adjoining blocks should also have some energy deposited. Accordingly the blocks around the block with the large signal are checked. If none of them has been hit then the event is flagged for rejection.

The effects of this filtering vary from zero to several hundred barrel and endcap events rejected per 8000 event run. A high rate of such events has occurred after damage caused by water leaks in the magnet and cooling systems. The algorithm requires less than 1 kbyte of memory and a negligible amount of CPU time.

12. Track triggers – fast pattern recognition

When examining the T2 events retained by the MIPROC-16 filtering program, it was clear that a large number of "obvious" background events were not being filtered out. One cause was additional spurious hits which resulted in the Z-vertex algorithm not being able to find a sufficiently clear peak in the histogram. Another cause was that hits in the inner ring could sometimes have poor resolution of the Z coordinate. This problem can be due to having several tracks very close together.

If the hits belonging to the individual tracks could be identified then using these hits the Z-vertex would be the best that could be achieved. The full JADE pattern recognition routines would require far too many resources to be used online, but a fast pattern recognition algorithm [11] had been developed with a view to being used to do preliminary track finding in candidate "good" events. The fast pattern recognition algorithm differs from the standard offline one in that hits within the inner two rings of the inner detector are used rather than all three rings and that the track-disentangling algorithm is much simplified – the latter results in the greatest saving of time. The consequence is that if there are many tracks within a cell (such as in a hadronic jet) then it is possible that not all tracks will be found. For online filtering purposes this does not cause problems since such events do not depend on the exact number of tracks. The fast algorithm was employed to work on candidate reject events.

The method first performs track finding in the $r-\phi$ plane. It is only sensitive to tracks originating from close to the centre of the detector in this plane – T2 events are defined as having at least two of these. Track finding efficiency is good for tracks with offsets up to ~ 15 mm (cf. maximum beam offset of 2 mm). The hits on each track are then fitted to a straight line in the $R-Z$ plane and the Z-intercept of the track with the beamline determined. The Z-vertex of the event is defined to be the Z-intercept of the track closest to the interaction region. Events are rejected either because all tracks originate outside the 300 mm online cut or because no tracks at all have been found. The algorithm

Fig. 14. "Big block" endcap lead glass event.

was tested offline and from a sample of T2 events from a REFORM tape 24% were found to have at least one track originating inside the 300 mm Z cut, 29% had tracks with none inside the cuts and 47% had no reconstructed tracks at all. This would give a rejection factor of 76%. Since this is quite a significant percentage it was essential to be extra careful with tests. The algorithm was then tried on REDUC2 data. From 14 000 T2 events (representing about 10 d of efficient running), three events were found which had two tracks from the interaction region where poor Z resolution had caused

Fig. 15. Beam–gas event with 4 tracks.

a bad Z fit. These events were selected for rejection on the grounds that the Z-vertex was outside the cuts. The algorithm was tested on 633 muon pair events and none was selected for rejection. All calculated Z-vertices were within 150 mm of the interaction region. All events had at least two tracks except for 6 where only one track was found. This was due to the second track travelling along an inner detector cell wall giving unreliable hits and delta rays. The algorithm was tested on a further selection of 2-prong events without any resulting losses.

Considering the advantages of the algorithm (greatly reduced data flow) against the drawbacks listed, the collaboration decided to use the cut online.

Even fast pattern recognition requires a significant amount of computer time which increases with the square of the number of hits. Less time is spent on events with genuine tracks than on those without. To prevent large events causing problems in this respect the algorithm is only applied to those T2 events with less than 400 hits in the inner two rings. This reduces the potential rejection factor of T2 events from 76 to 64%. The algorithm requires 22 kbytes of memory.

Figs. 15 and 16 show an event selected for rejection because all tracks originate from outside the cut. Fig. 15 is an r–ϕ view. Four tracks have been found coming from the origin. Fig. 16 shows a z–y view. The four tracks are seen to originate from well outside the 300 mm cut. The event can be explained as an electron (approaching from the right of the picture) colliding with a gas molecule within the beampipe. Although the mirror of a track may be occasionally selected (as here in track 1), the resultant Z-vertex of that track will not deviate significantly from the correct position. The

Fig. 16. $z - y$ view of the event in fig. 15.

MIPROC-16 analysis found three tracks in the outer ring of the inner detector (one does not penetrate fully). The Z-vertex algorithm found a peak in its histogram outside the ±300 mm cuts but the peak was insufficiently clear for the event to be rejected.

Figs. 17 and 18 show two events that have been selected for rejection because no tracks originate close to the origin in the $r-\phi$ plane. The event shown in fig. 17 is probably due to an off-momentum particle in the beam colliding with the wall of the beampipe. As with the event in fig. 15, the MIPROC-16 Z-vertex analysis did not find a sufficiently clear peak in its histogram to classify the event for rejection. The event in fig. 18 is due to a cosmic particle. In this case the MIPROC-16

analysis found a clear peak well within the ±300 mm cuts and so retained the event. Only the more sophisticated analysis can tell that the tracks do not originate from close to the interaction region.

The effectiveness of the algorithm depends upon beam conditions. Poor running conditions produce a higher number of beam-wall and beam-gas events and hence a potential higher number of events to filter out – but such conditions also produce events with a greater number of hits. The fast online pattern recognition filters out around 55% of the T2 trigger events that are passed to the NORD-50. The results of the pattern recognition are included in the bank of results which is appended to each event. This information can poten-


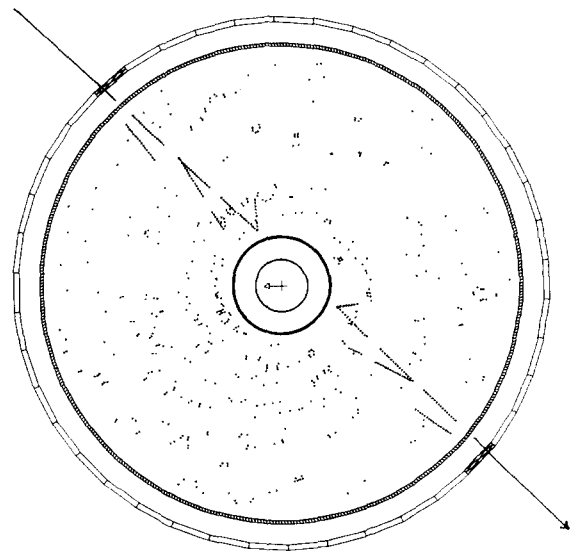Fig. 17. Beam–wall event with no tracks from origin.

Fig. 18. Cosmic event – no tracks from the origin.

Table 1
Summary of event filtering

Algorithm	Trigger	Fraction of all triggers examined (%)	Fraction that is filtered out (varies) (%)
MIPROC-16			
Software track check	T2	50	5
Z-vertex cut		50	50
NORD-50			
Cosmics contaminating the low energy neutral trigger	T1	18	85
"Grazing" cosmic in energy trigger	T1	20	50
LG noise in energy trigger		20	0–30
Pattern recognition	T2	24	55

tially be used as the initial parameters for a more sophisticated offline pattern recognition.

13. Summary

Online event filtering has been operating successfully for several years at JADE. The MIPROC-16 algorithms has operated since the early days of the experiment and the NORD-50 algorithms have been developed since 1982. The effect of the various filtering algorithms is summarised in table 1. The reduction in the amount of data written to the central IBM computer has made a significant saving in computing resources required for offline storage and processing. The overall fraction of events that are filtered out is around 60%, fluctuating with beam and trigger conditions. The actual reduction in volume of data is around 55%. The difference is due to the fact that some of the filtering algorithms tend to cut out the shorter events since these are often due to triggers caused by cosmic particles.

Online event-filtering is closely coupled to the trigger and operation of the experiment. If conditions are changed, such as a decision to alter the setting of the magnetic field significantly, then the calibration constants and cuts have to be updated accordingly. When a new trigger is introduced into the experiment an ad-ditional filtering algorithm may be appropriate.

The commitment to filtering events by online software at an early stage of the experiment has proved to be a worthwhile one.

References

[1] JADE Collaboration, W. Bartel et al., Phys. Lett. 88B (1979) 171; Phys. Lett. 92B (1980) 206; Phys. Lett. 99B (1981) 277.
[2] D. Cords, P. Dittmann, R. Eichler and H.E. Mills, The Data Acquisition System for the JADE Detector, Nucl. Instr. and Meth. A245 (1986) 137.
[3] NORD-10S/50 Hardware Concepts, Norsk Data, Oslo, Norway.
[4] MIPROC-16, Plessey Microsystems, Water Lane, Towcester, Northamptonshire, England. D. Cords, R. Eichler and H. Riege, CERN report 8107, P. 15.
[5] E. Gadermann and H. Krehbiel, JADE Note 30, unpublished; M. Helm, H. Krehbiel and H. Riege, JADE Note 31, unpublished.
[6] W. Bartel et al., Data reduction and data analysis at the IBM, JADE Note 33, unpublished.
[7] P. Dittmann, JADE Note 78, unpublished.
[8] J. Olsson et al., Nucl. Instr. and Meth. 176 (1980) 403.
[9] H.E. Mills, JADE Computer Note 82, unpublished.
[10] S. Yamada and S. Komamiya, private communication.
[11] P. Steffen, private communication.