

A pipelined first-level trigger for the H1 forward-muon spectrometer

T. Ahmed ^a, J.D. Dowell ^a, N. Ellis ^a, J. Garvey ^a, P. Jovanovic ^a, P. Newman ^a,
G. Noyes ^a, W. Stokes ^a, J.P. Sutton ^a, L. Jönsson ^b, B. Lundberg ^b, E. Eisenhandler ^{c,*},
B. Claxton ^d, J. Dowdell ^d, M. French ^d, S. Jaroslowski ^d, M. Prydderch ^d

^a University of Birmingham, Birmingham, UK

^b University of Lund, Lund, Sweden

^c Queen Mary & Westfield College, University of London, London, UK

^d Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, UK

Received 11 January 1995

Abstract

We have built a fast, pipelined first-level trigger processor for the H1 experiment at HERA. The trigger finds tracks in the forward-muon spectrometer which point back to the interaction vertex. The inputs to the trigger come from drift chambers with a 6 cm drift space corresponding to a maximum drift time of 1.2 μ s. These chambers consist of two layers of drift cells whose staggered configuration makes it possible to extract the time at which a particle traversed the chambers to better than one HERA bunch-crossing period, 96 ns. The trigger hardware exploits this characteristic in a real-time operation, and is able to find pointing tracks and associate them to production in a specific electron-proton crossing. The trigger processor must be deadtime free. It operates in a pipelined mode with 48 ns steps and has a latency of about 22 HERA bunch-crossing periods. The compact design is based on two semi-custom integrated circuits. Both are field-programmable 32×32 coincidence matrices, one having serial loading of its inputs and the other using parallel loading. The system was installed in the H1 experiment early in 1993 and has run successfully since then.

1. Introduction

HERA is a high-energy colliding beam accelerator operating at DESY, Hamburg. It consists of two intersecting storage rings situated in a tunnel of 6.3 km circumference, one containing electrons of 27 GeV energy and the other protons of 820 GeV energy. The beams of particles collide at two points where experiments are mounted: H1 in the north area and Zeus in the south area.

The circulating beams are not continuous but consist of a stream of particle bunches 96 ns apart. The proton bunches are approximately 1.1 ns long and the electron bunches approximately 0.3 ns long. In the final design configuration of HERA there will be 210 bunches of each type which will effectively populate the complete circumference of the rings, and provide a peak luminosity of $2 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$. Storage rings with colliding bunched beams do not represent a new concept. However, in other colliders the time between bunch crossings is much longer, e.g. the CERN Collider where the time was 3.8 μ s and

LEP where the time is now 11 μ s. At HERA the crossing time of 96 ns puts additional constraints on the design of the detectors, and on the selection of interactions in real time via triggering. It is also a useful step towards the CERN LHC, where the crossing time will be 25 ns.

H1 has adopted a multi-level trigger philosophy involving four levels of triggering: L1, L2, L3 and L4. These triggers examine the information from the various components of the experiment, and select events satisfying pre-set combinations of conditions. They act sequentially and apply more and more selective criteria. L1 is the first selection process and has to examine detector information every bunch crossing, whereas L4 is a farm of computers with a processing capability of about 100 events/s. During the running of 1993 and 1994 H1 operated with triggers L1 and L4 only [1].

The L1 trigger must be sensitive at every bunch crossing if the experiment is to be efficient in its use of the luminosity of HERA. Triggers contributing at this level must either select the interaction of interest in less than 96 ns, or operate in a pipelined mode, accepting collisions every 96 ns and processing them in 96 ns steps. The number of steps in the pipeline determines the latency of the trigger, which is the delay between the crossing and the

* Corresponding author.

L1 decision being produced. The trigger latency has quite significant consequences for the specification of detector front-end electronics. A further requirement of the L1 trigger is that it must be able to identify each interaction which it selects with a unique bunch crossing. We will refer to this process as *t-zero extraction*. This requirement means that in any combination of L1 trigger conditions, at least one element of the combination must be based on a detector which has intrinsic time resolution better than the bunch crossing period of 96 ns. Multiwire proportional chambers and scintillation counters can achieve this resolution easily. The liquid argon calorimeter of H1 cannot, if also required to retain good energy resolution.

The L1 trigger described in this paper is based on drift chambers with a maximum drift time of 1.2 μ s, but with a wire configuration such that *t-zero* can be extracted. It looks in the forward-muon spectrometer of H1 for particles whose tracks point back to the region where the electron and proton beams collide. These particles have penetrated the end-cap iron of the H1 magnet. Many of them will be muons, which constitute a signature for the production of heavy-flavour quarks in the interaction. The interaction region is about 30 cm long and has transverse size of less than 0.1 mm. Particles which enter the forward-muon system come not only from electron–proton collisions, but also from proton beam–gas interactions, proton beam–wall interactions, and synchrotron radiation from the electron beam. The rates from the background processes are much higher than those from genuine electron–proton interactions. The pointing requirement provides a powerful means of rejecting such tracks at L1.

2. Forward-muon spectrometer

The H1 forward-muon spectrometer is shown in Fig. 1 [2]. It consists of six drift chamber planes, three either side of an iron toroidal magnet. Each plane is subdivided into octants in azimuth, with octants 0, 2, 4, 6 being identical, and slightly different from octants 1, 3, 5 and 7 for mechanical reasons. The geometrical arrangement of indi-




Fig. 1. Diagrammatic side view of the H1 forward-muon spectrometer. The end-cap iron, liquid-argon calorimeter and central tracker are to the right.

Table 1
Location of θ drift-chamber planes and number of drift cells per octant

Plane	z-coordinate	No. of cells
θ_1	637 cm	24
θ_2	712 cm	28
θ_3	868 cm	36
θ_4	943 cm	40

vidual drift cells is such that it provides a measurement of the polar angle θ in planes 1, 3, 4 and 6 (denoted θ_1 , θ_2 , θ_3 and θ_4) and a measurement of the azimuthal angle ϕ in planes 2 and 5 (denoted ϕ_1 and ϕ_2). The ϕ chambers play no part in the trigger described in this paper and no further mention will be made of them.

The region covered by the forward-muon drift chambers is in the range of polar angles 3° to 17° . The locations of the chamber planes along the beam direction relative to the nominal interaction point are given in Table 1.

A drift cell has a rectangular cross section, with a depth of 2 cm and a width of 12 cm. The centrally-positioned sense wire defines a maximum drift distance of 6 cm. The lengths of the drift cells can vary from 40 to 240 cm, depending on their distance from the beam-pipe. A drift-chamber plane has two layers of drift cells, with one layer displaced by half a cell width with respect to the other as shown in Fig. 2. The purpose of this arrangement is to resolve the ambiguity between hits on either side of the wire, and to correlate the hits in the spectrometer with the corresponding electron–proton collision in HERA as will be explained in Section 5.1. To reduce the number of electronic readout channels required, sense wires of adjacent cells are connected via a 330 Ω resistor as indicated in Fig. 3. Both ends are read out, and the coordinate along the wire measured by charge division. The drift cells are filled with a gas mixture consisting of 92.5% argon, 5% CO_2 , and 2.5% methane. They are operated with a drift voltage of 4.2 kV and have a maximum drift time of 1.2 μ s. The readout to the data acquisition system uses the




Fig. 2. Configuration of drift cells in one drift-chamber plane.



Fig. 3. Detail of connection and charge division on adjacent drift cells.

same F1001 flash-ADC system [1] as the other drift chambers of H1) to sample and digitise its pulses, which are then reconstructed in time using a $Q-t$ algorithm.

3. Philosophy of the trigger

The forward-muon trigger is required to find tracks in the forward-muon system which point back to the electron-proton interaction region. The trigger must also be able to associate tracks in the drift cells with a particular bunch crossing, which requires a time resolution of better than 96 ns. The contrast between this and the time resolution of a single drift cell, 1.2 μ s, is very striking.

The pointing requirement is imposed for tracks which reach the toroid, and also for tracks which traverse the toroid. The pointing is degraded for both categories. Muons in the region before the toroid which have come from the interaction point have been deflected in the central solenoidal magnetic field, and have undergone multiple coulomb scattering (m.c.s.) in both the liquid argon calorimeter and the endcap iron of the H1 solenoid. Tracks in the region after the toroid have also been bent in the toroidal field and have undergone further m.c.s. in the toroidal iron.

The trigger described here finds tracks that hit the same octant of each drift-chamber plane, but not tracks which cross from one octant to another. This causes a small loss in efficiency. The logical steps of the forward-muon trigger are shown in Fig. 4. The analogue signals from two linked drift cells are fed into a discriminator module, which applies an amplitude threshold to each signal to determine if a track is present, and produces an output logic signal on one of two output lines corresponding to the cell which the track traversed (i.e. which has the larger signal). The discriminator outputs are sent to road-finder modules which select hits corresponding to tracks which point back to the interaction region, within the limits allowed by multiple scattering and bending. This is done separately for tracks in planes Θ_1 and Θ_2 , and for planes Θ_3 and Θ_4 . Finally, tracks found before the toroid and after the toroid are correlated in order to find traversing tracks. The outputs from each octant are combined in the final-decision module to produce topological triggers. Each element of Fig. 4 will be described below.

4. The discriminator module

The pulses coming from the drift cells rise in about 20 ns and fall in about 100 ns, with a maximum pulse height of about 2 V. A complication of the readout is that to reduce cost adjacent chamber wires are connected together by a resistor, as shown in Fig. 3. In the offline analysis the coordinate along the wire is obtained by charge division. The sense wires are 50 μ m-diameter nichrome wires for the longer cells and 40 μ m-diameter for the shorter ones, with resistances of 550 and 859 Ω m⁻¹ respectively. The 330 Ω linking resistor thus corresponds to wire lengths of 60 cm for the longer cells and 38 cm for the shorter ones. The trigger is required to identify which of the two linked cells was struck in real time, and this is achieved by comparing the amplitude of the pulses on the two wires and allocating the hit to the wire which has the bigger amplitude. This is done using the circuit shown in Fig. 5. If either wire signal amplitude is above a threshold, which can be set in 256 equal steps up to a maximum value of 1.5 V, the discriminator fires and its output is shaped to give a 50 ns-long pulse. After a delay of 20 ns, this enables a comparator into which the two wire signals are fed. The two complementary comparator outputs are arranged such that the one which corresponds to the bigger amplitude goes high. The two outputs are strobed out of

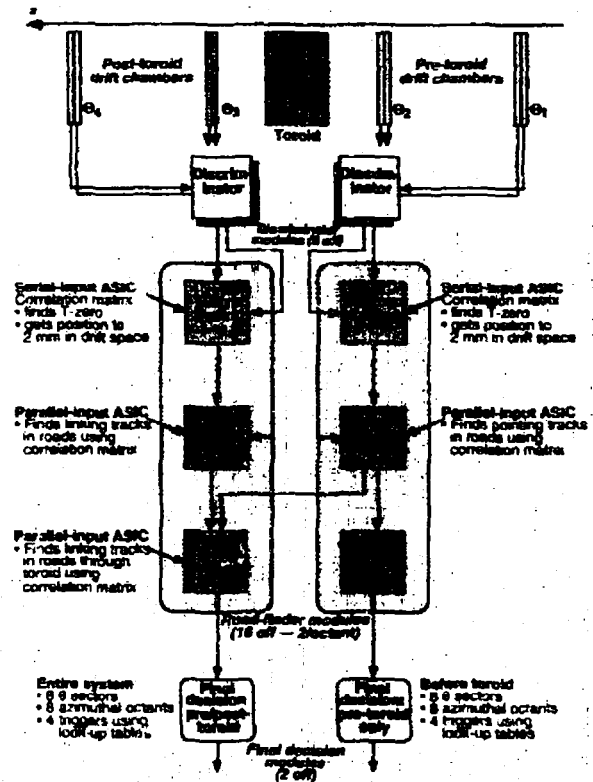


Fig. 4. Block diagram of the forward-muon trigger.




Fig. 5. Circuit diagram of the discriminator.

the circuit with a 100 ns-long strobe. The 20 ns delay to the enable was selected to match the rise time of the pulse, so that the comparator is enabled on the pulse peak rather than the rising edge. The 100 ns output strobe ensures that the discriminator is effectively disabled until the next bunch crossing. The performance of the circuit was optimised on a test rig consisting of two coupled drift cells and a cosmic muon telescope, and was shown to be nearly 100% efficient at wire discrimination for thresholds above a few tens of mV. The results are shown in Fig. 6; the number of wrong assignments to the coupled cell is very small.

The four Θ planes have between 24 and 40 cells per octant, as indicated in Table 1. In order for the trigger logic to be independent of the drift-chamber plane to which it is applied, all of them have been arranged to provide 28 trigger channels. This is achieved by adding four "null" cells to the first plane, and by ORing the outer 16 cells of the third plane into pairs as well as the 24 outer cells of the fourth plane. The ORing is done in the discriminator modules, whose output signals then consist of 24 signals from Θ_1 , and 28 from each of Θ_2 , Θ_3 , and Θ_4 . These signals, which are at this stage asynchronous with the HERA 96 ns clock, are sent out onto the back-plane and into the road-finder modules about to be described.

There is one discriminator module per octant. Each module consists of a mother-board holding 64 plug-in 6-layer printed-circuit boards (2.3 cm x 5.1 cm), each of which contains one of the discriminator circuits shown in




Fig. 6. Efficiency of discrimination as a function of discriminator threshold.

Fig. 5. The discriminator mother-board is an 8-layer printed circuit 9U high x 40 cm deep (1U = 1.75 in.).

5. The road-finder module

These modules, indicated by the large shaded areas in Fig. 4, extract t-zero for each pair of wire signals, find tracks in Θ_1 and Θ_2 , or in Θ_3 and Θ_4 , and finally correlate traversing tracks before and after the toroid.

5.1. T-zero extraction

In Fig. 2 we have superimposed a track on a pair of overlapping drift cells. This has drift times of T_1 and T_2 to the two sense wires. For pulses arriving at the wires, originating from a single track perpendicular to the drift direction which has gone through both drift volumes, the sum of T_1 and T_2 is equal to the drift time over the full 6 cm drift distance of the chambers, T , which is a constant:

$$T_1 + T_2 = T. \tag{1}$$

What appears on the wires is of course only a stream of pulses, relative to some unknown starting time. This unknown starting time identifies the bunch crossing which produced the particle that left the track.

The way in which these pulses are processed by the trigger is shown in Fig. 7. The pulses from a pair of overlapping drift cells are shaped by discriminators as described above, and then are strobed into two shift registers, 32 locations long, which form the x and y axes of a coincidence matrix. The system clock which provides the strobe has a period of 48 ns and is synchronous to the 96 ns HERA clock. An application-specific integrated circuit (ASIC) contains all the logic shown in Fig. 7.

The strobing frequency of the discriminator outputs into the shift registers provides the limiting factor in the time resolution of the trigger. As mentioned in Section 4, the pulses coming from the drift cells rise in about 20 ns. For a fixed discriminator threshold, this means that the




Fig. 7. Block diagram of the t-zero ASIC.

maximum time slew of the discriminator output pulses relative to the pulse centroid is less than 20 ns. The discriminator pulses are asynchronous with the system clock, and so strobing them into the shift registers introduces a maximum additional time slew of 48 ns. The combined time slew, coming from the discriminators and the strobing, will still be less than the time between HERA bunch crossings, and therefore should not result in an incorrect t -zero being extracted. A system clock period of 96 ns is clearly not viable, while a system clock of 24 ns would provide better overall time resolution but would require a 64×64 coincidence matrix. The additional cost of this system was not considered to be justified.

The pulses are clocked along the shift registers with the system clock. Each of the elements of the 32×32 coincidence matrix can be considered as a 3-input AND gate, as shown in Fig. 7. The inputs correspond to:

- an element of the x -axis shift register,
- an element of the y -axis shift register,
- a field-programmable bit.

The outputs of the AND gates are ORed in the y direction and output to a z register as shown. The programmable elements of the coincidence matrix that are activated determine the correlation between the x and y values which will produce an output at z . A diagonal load as illustrated in Fig. 7 defines a condition.

$$x + y = \text{CONSTANT}. \quad (2)$$

When this condition is satisfied by the input pulses, an output pulse appears in the z -register at a value $z = x$. If this CONSTANT is equal to the drift time corresponding to the full drift distance, T , expressed in units of system clock ticks, then we have a constraint which is equivalent to that in Eq. (1). Note that for particles not normal to the drift direction the equation must include an additional constant.

The delay between the particle being produced and the signal appearing in the z -register is the sum of the time for the particle to reach the drift cell, the maximum drift time in the chamber, the time for the wire signal to arrive at the discriminator and produce an output signal, and the delay through the t -zero chip itself. This latter is one 48-ns clock tick. Thus the time of the bunch crossing responsible for the track can be extracted.

Not only does the coincidence matrix allow the determination of the bunch crossing responsible for the track, it also determines the location of the track in the drift volume. Given that the drift time is 1.2 μ s for 6 cm, the ionisation drifts 2.4 mm for every 48 ns clock tick. The distance, d , of the track from the sense wire is given by

$$d = (t_0) - z \times 2.4 \text{ mm}. \quad (3)$$

Thus we get both a measurement, in 48 ns steps, of the time at which the track crossed the chamber (from the time taken for a coincidence in the matrix to occur), and a measurement, in 2.4 mm steps, of the track position in the

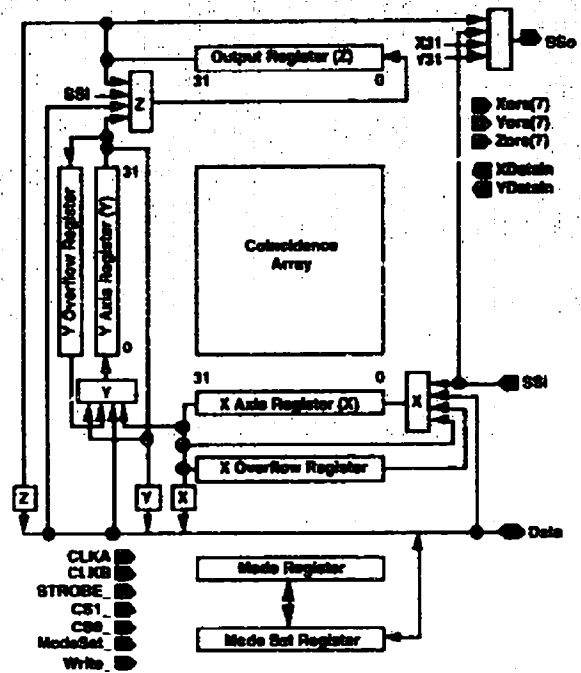


Fig. 8. Logic layout of the t -zero ASIC.

drift volume (from the value of z where the coincidence occurs). The philosophy of this approach represents the heart of the forward-muon trigger. Given hits in space and time in each of the four Θ planes, the problem of finding pointing tracks can be addressed.

The implementation of the logic in Fig. 7 is done in a semi-custom ASIC using 1.5 μ m CMOS technology. The logical layout of the t -zero ASIC is shown in Fig. 8. Here we indicate some of its significant characteristics. The chip can be operated in four modes, controlled by the *mode* register. The modes are:

1) **NORMAL RUNNING** in which the two wires from overlapping drift cells are connected to the x and y registers and the discriminator pulses are clocked using the 48 ns clock. In this mode the x and y registers are each connected to 32-bit long overflow registers, which continue to clock the data along. The purpose of the overflow registers is to allow offline monitoring of the trigger performance. The delay between the interaction and the registration of a level-one trigger in H1 is 24 bunch-crossing periods, or 48 system-clock ticks. When a L1 trigger is declared by the Central Trigger Logic the state of the trigger processor is frozen. By this time the wire hits which contributed to the forward-muon trigger have clocked on along the x and y registers and may even have shifted out of the register. The overflow register retains them for offline analysis.

2) **REWIND mode** is implemented after the Central Trigger Logic has decided to read out an event. The contents of the x and y registers of the t -zero chips have

to be restored to the value they had at the time the trigger was produced, about 10.5 bunch crossings earlier. In this mode, the x and y registers and their overflows are connected to form a closed loop and the contents are clocked back by the correct number of system clocks for hits to occupy the same positions in the shift register that they had at the instant the trigger was produced, and therefore the z register and all subsequent registers in the trigger are also as they were.

3) READOUT mode is implemented to read out the values of the x , y and z registers after the rewind. In this mode the three registers are connected in series and read out serially as part of a module-wide scan path.

4) LOADING mode is used to load the bits enabling the individual cells in the coincidence array. The contents of a horizontal row are loaded into the x register, the address into the y register, and then the load instruction is carried out.

The z register contents change every 48 ns to reflect the correlation between the hits in the x and y registers. The 32 bits of the register correspond to the position of the hit in the drift volume in units of 2.4 mm as indicated in Eq. (3). This coordinate is required as input to the track-finding ASIC that will be described in Section 5.2. However the intrinsic resolution of 2.4 mm, which has come as a result of the time resolution requirement of 48 ns, is better than is needed for tracking. This is because m.c.s. in the end-cap iron of the H1 solenoid and in the liquid argon calorimeter cause a smearing in direction of the tracks emerging from the interaction point. This will be discussed in detail in Section 5.2. Therefore the positional information from the t-zero chips is only provided on seven lines which correspond to

- 1 ORs of 8 bits,
- 2 ORs of 16 bits,
- 1 OR of 32 bits.

giving nominal spatial resolutions in the drift direction of

1.92, 3.84, and 7.68 cm (actually limited to 6 cm by the size of the drift cells).

Also available as outputs are the wire-ORs, i.e. ORs of regions of the x and y registers, which are produced in the t-zero chips with the same granularity as above but without any coincidence requirement. Their use will be described below in Section 5.2.

A full description of the t-zero chip specification is available in Ref. [3].

5.2. Track finding

As indicated in Fig. 4, the track finding is done in several consecutive stages. In Section 5.1 we described what is done by the t-zero chips in each of the Θ planes. This already applies some pointing constraint, as the sum of drift times indicated in Eq. (1) is only true for tracks traversing the chambers at right angles to the drift direction. The sum of drift times depends on the angle of the track, and this is catered for by changing the constant in Eq. (2), that is by moving the position of the diagonal coincidence-matrix load.

These t-zero chip outputs are used to find pointing tracks in the pre-toroid chambers, Θ_1 and Θ_2 , and separately in the post-toroid chambers Θ_3 and Θ_4 . This is done in parallel in two separate electronic modules, as shown in Fig. 4, using track-finding ASICs. Finally, the pre-toroid tracks are correlated with the post-toroid ones to find traversing tracks.

We will now discuss the spatial resolution required at the drift chambers to define a track pointing to the interaction region. The track emerges through the central and forward tracking systems, the liquid argon calorimeter, and the return yoke of the magnet before it reaches planes Θ_1 and Θ_2 . The change in the angle of the track due to m.c.s. varies as $1/p$, where p is the particle momentum. The H1SIM Monte Carlo has been used to investigate the

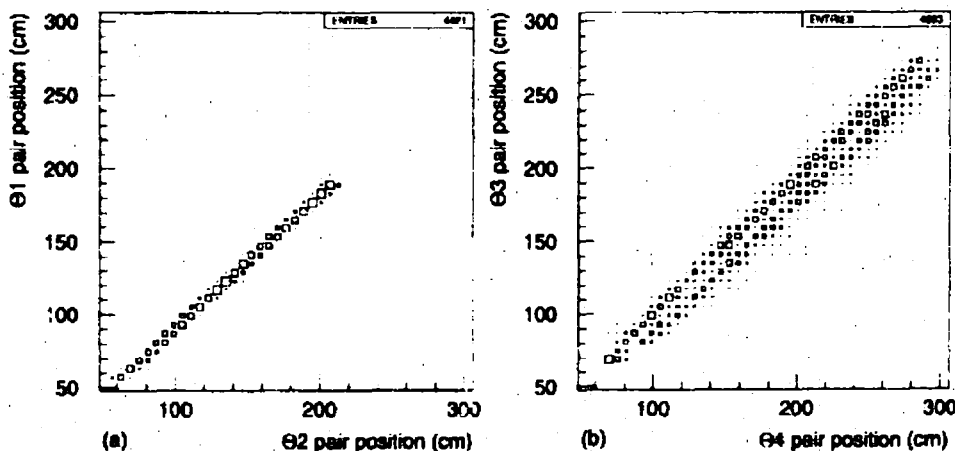


Fig. 9. Monte Carlo correlation between hits in (a) Θ_1 and Θ_2 , (b) Θ_3 and Θ_4 , for tracks corresponding to 5–10 GeV/c particles originating in the interaction region.

effects of m.c.s., and in particular to study the correlation between coordinates in Θ_1 and Θ_2 for tracks coming from the interaction region. The results are shown in Fig. 9 for tracks generated with a uniform momentum distribution from 5–10 GeV/c. It is clear that to obtain 100% efficiency for 5 GeV/c particles the road widths that are required at Θ_1 and Θ_2 are cm wide rather than mm. These roads narrow as the momentum is increased, since the width is dominated by m.c.s. rather than by the length of the interaction region, and by 20 GeV/c are only about 1 cm wide. The much wider correlation between Θ_3 and Θ_4 , in two separate bands, reflects the fact that both signs of particle have to be accepted after bending in opposite directions, and so in order to have good efficiency for 5 GeV/c particles after the toroid very wide roads have to be set.

The purpose of the trigger is to select interactions which produce muons in the forward chambers, and to reject background processes such as beam-wall and beam-gas interactions. These background processes are likely to confuse the pointing trigger, particularly for regions near the beam where the tracks are at very small polar angles. If the rate of false triggers in this region proves high, the trigger must retain the possibility of narrowing the roads to reduce it. This narrowing will lead to a loss in trigger efficiency at 5 GeV/c for particles at small angles to the beam. Moving out in angle towards the limit of 17° , the width of the roads required to reject background processes becomes less critical.

To preserve as much flexibility and safety as possible, while keeping the overall size of the trigger processor within acceptable limits, the "t-zero" road widths given in Table 2 were implemented. This gives a total of 64 roads in each octant of each Θ plane. Extracting these road widths from the t-zero chips is just a matter of selecting the appropriate OR outputs.

Having decided to divide each octant of each Θ plane into 64 "t-zero" roads, pointing tracks can be found in Θ_1 and Θ_2 , and in Θ_3 and Θ_4 , by correlating the signals from these roads in a 64×64 coincidence matrix. However such a procedure would require each of the four drift cells through which the particle passed to give a signal. This would clearly lead to inefficiencies as the cells have walls 0.7 mm thick, and for normal incidence this will lead to a straight geometrical inefficiency of 1/60 per drift volume and a maximum track-finding efficiency of 93% even if the drift cells have 100% signal efficiency.

Table 2
Road widths used for the t-zero chips

T-zero chips	Road width	No. of roads
1–9	1.92 cm	36
10–19	3.84 cm	20
20–27	7.68 cm	8




Fig. 10. The ORing of the wire signals from adjacent drift volumes.

The way around the problem is to accept tracks which have only three out of the four possible chamber hits. Such a solution effectively removes the problem of track-finding inefficiencies. However, it has several bad effects on the integrity of the tracks which are found. Three-bit tracks consist of a t-zero-validated pair of hits in one Θ plane and a single wire hit in the other. If the track is seen in only drift volume the time resolution is clearly the full drift time of 1.2 μ s, and the spatial resolution, because of the ambiguity for hits either side of the wire, becomes twice the full drift distance, i.e. 12 cm. If the single-wire hit rate is high, due to background interactions, the chance of a random correlation being made between a pair of hits in one chamber and a wire hit in the other is high. The random problem can be reduced by dividing the 6 cm drift space up into narrower roads before tracks are looked for. The same road widths as those described above for the t-zero z register are used.

To cater for tracks with only three hits, the wire-hit information from every wire in the system has to be brought to the track-finding chips. What is actually required is the OR of the two wire signals involved. This is shown schematically in Fig. 10, where we show two overlapping drift cells of one of the Θ planes. The ambiguity problem results in the volume covered by this OR being that within the overall outline of the two drift cells. On the t-zero chip, as already mentioned in Section 5.1, various wire-ORs are formed. The 8-bit ORs are used here, as shown in Fig. 10. Matching signals from the two wires are then combined to form wire ORs which correspond to the drift volumes indicated by different patterns. It can be seen that the effective volume represented by one wire OR covers a total of (3×1.92) cm in the θ coordinate.

The extent of the ambiguity problem is shown in Fig. 10. For the wire-OR hits indicated, the ambiguities lead to uncertainties in track position as shown. Only the wire-ORs of every second t-zero chip are required to cover the Θ planes completely. For the wire-OR signals, each Θ plane

Table 3
Roads for wire-OR signals

T-zero chips	Road width	No. of roads
1, 3, 5, 7	1.92 cm	16
9, 11, 13, 15, 17, 19	3.84 cm	12
21, 23, 25, 27	7.68 cm	4

is divided into 32 "wire-OR" roads. These are shown in Table 3; it should be noted that these are not the same as the "t-zero" roads of Table 2.

The track finding between θ planes is implemented as shown in Fig. 11 for θ_1 and θ_2 . The process is based upon a semi-custom ASIC in 1.5 μm CMOS technology which incorporates a 32×32 coincidence matrix similar to the t-zero ASIC. The block diagram is shown in Fig. 12. The inputs to this track-finding chip are 32-bit parallel on both x and y registers, rather than serial as on the t-zero chip. The mode register has equivalent operation to that of the t-zero chip [4]. The system clock takes in the data from the t-zero chips every 48 ns.

Four ASICs are used to perform the track finding in a pair of θ planes, as shown in Fig. 11. Here it should be noted that the inputs from chamber θ_2 always form the x axis of the track-finding chip, while θ_1 is always on the y axis. On the left-hand pair of ASICs it is the wire-ORs of θ_2 and the right-hand pair the t-zero z-register roads of θ_2 , while θ_1 is treated in the opposite manner. The z registers of both configurations of the track-finding chips give the position of the found track in θ_2 . The outputs of both these configurations of track-finding chips correspond to tracks with at least three wire hits, on the left two hits in θ_1 and one in θ_2 , and on the right two hits in θ_2 and one hit in θ_1 . The OR of the two z registers combines both possibilities and gives the protection that is needed against chamber inefficiency. The AND of the two z




Fig. 11. Configuration of track-finding ASICs to preserve the option of demanding four hits OR three hits in θ_1 and θ_2 .

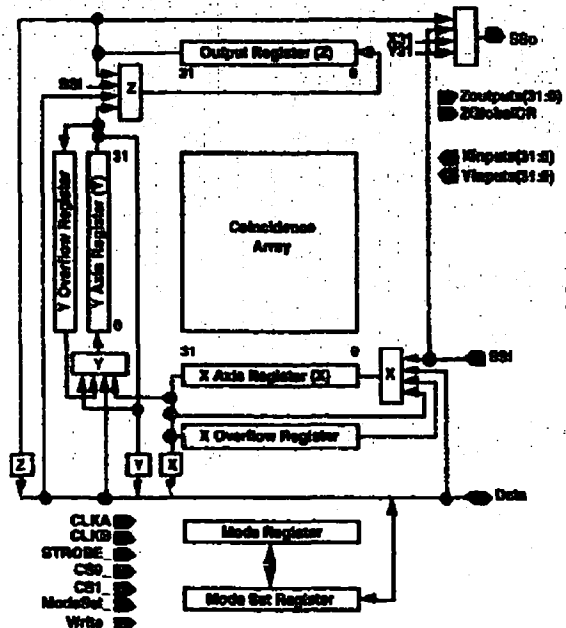


Fig. 12. Block diagram of the track-finding ASIC.

registers requires two hits in θ_1 and two hits in θ_2 . This will have a reduced random track component at the expense of lower efficiency, and may have to be used if the particle flux through the system is high.

The configuration of tracking matrices for chambers θ_3 and θ_4 is a mirror reflection of Fig. 11 in the toroid. The x axes are the wire-ORs and the t-zero z-register roads of θ_3 , and as a result the z register gives the position of the tracks found in θ_3 .

The final operation to find tracks traversing the toroid is to correlate the pre-toroid tracks with the post-toroid tracks. The ANDing and ORing of the z registers indicated in Fig. 11 produces a 32-bit word which is a measure of the position of the track in the drift chamber plane θ_2 for tracks before the toroid, and in θ_3 for tracks after the toroid. These 32-bit words form the x and y axes of a fifth track-finding ASIC, which is called the correlator chip to distinguish it. This determines the combinations accepted as traversing tracks. This is shown in Fig. 4, where it can be seen that the 32-bit word of track information from the pre-toroid road-finder module is sent to the post-toroid road-finder module.

All road-finder modules are identical, and so the 32-bit pre-toroid word also goes to a track-finding chip on the pre-toroid module (see Fig. 4). This is configured by software as a dummy, to pass the word through to the output of the road-finder module. What is finally sent out by the road-finder module is the 8 ORs of 4 bits each from the 32-bit z register of the correlator chip. These signals




Fig. 13: Scan path for the road-finder module, showing how the ASICs are connected for readout.

are available for tracks found before the toroid and for tracks traversing the toroid. They correspond to eight regions in θ for each octant.

5.3. Configuration of the module

The road-finder module (RFM) contains 54 t-zero chips, four track-finding chips and one correlator chip. All the RFMs operate as pre-toroid modules or post-toroid modules, depending on the slot they occupy in the crate. The configuration of x , y , and z registers and their overflows on the t-zero chips and the track-finding chips, which is implemented for readout after a trigger in order to perform offline evaluation of the trigger performance, has been described in Section 5.1. The configuration of the scan path for the complete module is shown in Fig. 13. A photograph of a module is shown in Fig. 14. The RFM also contains RAMs, which can be used to simulate input data from the discriminator modules and thus allow stand-alone testing of module operation. The RFM is implemented as a 12-layer printed-circuit board, 9U high \times 40 cm deep.

6. The final-decision module

The logical flow of trigger information from the outputs of the road-finder modules through final-decision modules and then to the Central Trigger Logic of H1 is shown in Fig. 15. The information coming from each RFM consists of eight bits, which give the approximate location in θ of the trigger tracks. There are two RFMs per octant, one providing these eight bits for pre-toroid tracks, and the

other for tracks traversing the toroid. Thus there are 64 bits of trigger data in all describing tracks before the toroid, and 64 bits of trigger data describing tracks traversing the toroid. (These 128 bits are also sent to the L2/L3 triggers.) There is a separate final-decision module (FDM) for each of these two possibilities. The function of the FDM is to extract from this data the eight trigger bits, or "subtrigger elements", that go to the Central Trigger Logic.

The FDM consists of a set of RAM lookup tables. One RAM takes in the eight bits from one road-finder module as an address and outputs its four-bit contents. These four bits can be programmed to give information on the multiplicity and location of hits in the octant. The first bit for each of the octants is combined to form an eight-bit word which is used to address a second level of RAM lookup tables. This is repeated for each of the four bits. The output from this second level is a single bit per RAM, and as a result the final output of the final-decision module is a four-bit word. The configuration of the lookup tables used at present is:

- bit 1 single muon in the small- θ region (i.e. inner four θ zones).
- bit 2 single muon in the large- θ region (i.e. outer four θ zones).
- bit 3 more than one muon in the same octant.
- bit 4 more than one muon in different octants.

Such a system of lookup tables is of course very versatile, since its contents can be changed easily if other trigger topologies look interesting. Each FDM sends four bits to the Central Trigger Logic.

The FDM allows another mode of operation, that of global ORs. The road-finder module, as well as sending the eight-bit trigger corresponding to the eight regions in θ to the FDM, also sends the OR of these eight bits. If this is set it indicates a track somewhere in the octant. The ORs coming from the eight octants can be directed with minimal delay to the output of the FDM. This option is provided to allow a slightly shorter trigger latency if that proves necessary. It passes the problem of topological triggers on to the Central Trigger Logic.

The FDM performs one final function in the trigger. The system clock used to drive the chamber data through the trigger has a period of 48 ns, i.e. one half of the 96 ns HERA clock period used to latch the trigger bits into the Central Trigger Logic. The final-decision module can strobe trigger results from the first half cycle of the HERA clock, from the second half cycle of the HERA clock, or it can combine the two cycles together. At present we use the latter option.

The FDM is implemented as a 10-layer printed-circuit board, 9U high \times 40 cm deep.

7. Control and readout

The complete forward-muon trigger is housed in four 9U high \times 40 cm deep crates. Each crate contains a crate-




Fig. 14. The road-finder module.




Fig. 15. Logical flow of information through the final-decision module.

controller module, two discriminator modules, four road-finder modules, and caters for two complete octants of the trigger. Two of the crates also contain a final-decision module, one for pre-toroid triggers and the other for toroid-traversing triggers. The overall configuration is shown in Fig. 16. A detailed description of the readout will not be given here but the function of the significant components will be indicated.

As has been described in Section 5.1, on receipt of an LI trigger from the Central Trigger Logic the system state is frozen. Immediately following this the x and y registers of the t-zero chips are rewound to bring the input data back to where it had been in the registers at the instant the trigger was produced. Re-establishing the data in the x and y registers of the t-zero chips results in the contents of all subsequent registers being re-established, including the final trigger bits. These data, the input and the output of all the t-zero, road-finder, and correlator chips, and also the final-decision bits, must be read out and retained with the




Fig. 16. The crate configuration of the trigger, indicating the various busses involved in readout and control.

detector data in order to allow monitoring of the trigger performance offline. Within one road-finder module this is done by connecting all the *x*, *y* and *z* registers into a single serial scan path and reading them out via a shift register on each module onto a 16-way data bus on back of the trigger crate, and into a multi-event buffer in the crate




Fig. 17. Summary of trigger timing.




Fig. 18. Discriminator curves taken with cosmic muons.

controller. The total time allowed in H1 to achieve this and reconfigure the modules for data taking is 800 μ s.

The dialogue with the front-end multi-event buffers and the trigger modules is controlled by a 68000-based CPU07 processor which resides in a VME crate. This processor organises the loading of the modules with register, coincidence-matrix and RAM settings, and the readout of the data from the modules into the multi-event buffer in the crate controller. It also controls the operation of a DSP, which takes the event data out of the multi-event buffer, zero-suppresses it, and passes it on to a memory (DPM) in the VME crate from where it is accessed by a FIC8230 processor for formatting before being sent to the data-acquisition system. Also in the VME crate are modules ("Fast" and "Slow") which handle the dialogue with the Central Trigger Logic.

8. Performance of the trigger

8.1. Trigger timing

The delay between the interaction taking place and a L1 trigger being declared by the Central Trigger Logic must be less than 2.3 μ s (24 bunch crossings). This limit is defined by the length of the drift-chamber readout pipelines. The drift-chamber pulses are digitised by a flash-ADC (FADC) system running with a clock frequency ten times that of HERA, and the digitised pulses are stored in pipeline memories 256 locations deep. All triggers have to report to the Central Trigger Logic within 23 bunch crossings of the interaction, i.e. 2.2 μ s. For the forward-muon trigger, taking into account the drift time in the chambers, and the cables from the chambers to the trigger and from the trigger to the Central Trigger Logic, the time available for the trigger to make its decision is much shorter. This is summarised in Fig. 17, where the total latency of the trigger logic is seen to be about 300 ns.




Fig. 19. Time resolution of the t-zero ASIC inputs.

8.2. Performance of the discriminators

The discriminators for the forward-muon trigger serve two functions. They must produce a logic pulse if either pulse is above a programmable voltage threshold, and they

must determine which of the coupled drift cells the particle traversed by comparing the pulse heights at the two ends of the coupled wires.

In Fig. 18 we show a plateau curve for the wires of the whole of plane Θ_3 , obtained using a cosmic-ray trigger. The cosmic ray was defined as a track seen in planes Θ_1 , Θ_2 and Θ_4 , and the fraction also seen in Θ_3 , is plotted as a function of discriminator threshold value in Θ_3 . Plane Θ_3 has 28 wires. Separate discriminator threshold values can be applied to two sets of 14 wires in each octant, the set nearest the beam (inner) and the set furthest from the beam (outer). Because of geometry the wires closest to the beam are the shortest, and therefore on average give bigger pulses for a fixed ionisation loss in the chamber. This effect can be clearly seen in Fig. 18, where the knee of the plateau for the inner wires is at a higher threshold value than that for the outer wires. Since for physics runs the inner regions of the chambers see the most activity due to background processes, it is important to set the discriminator thresholds as high as possible to reduce the rate of inputs to the trigger logic.

The time resolution of the input to the t-zero chips has been measured for real data by comparing the drift time measured by the arrival time of hits in the x registers of the t-zero chips with the drift time measured offline by the

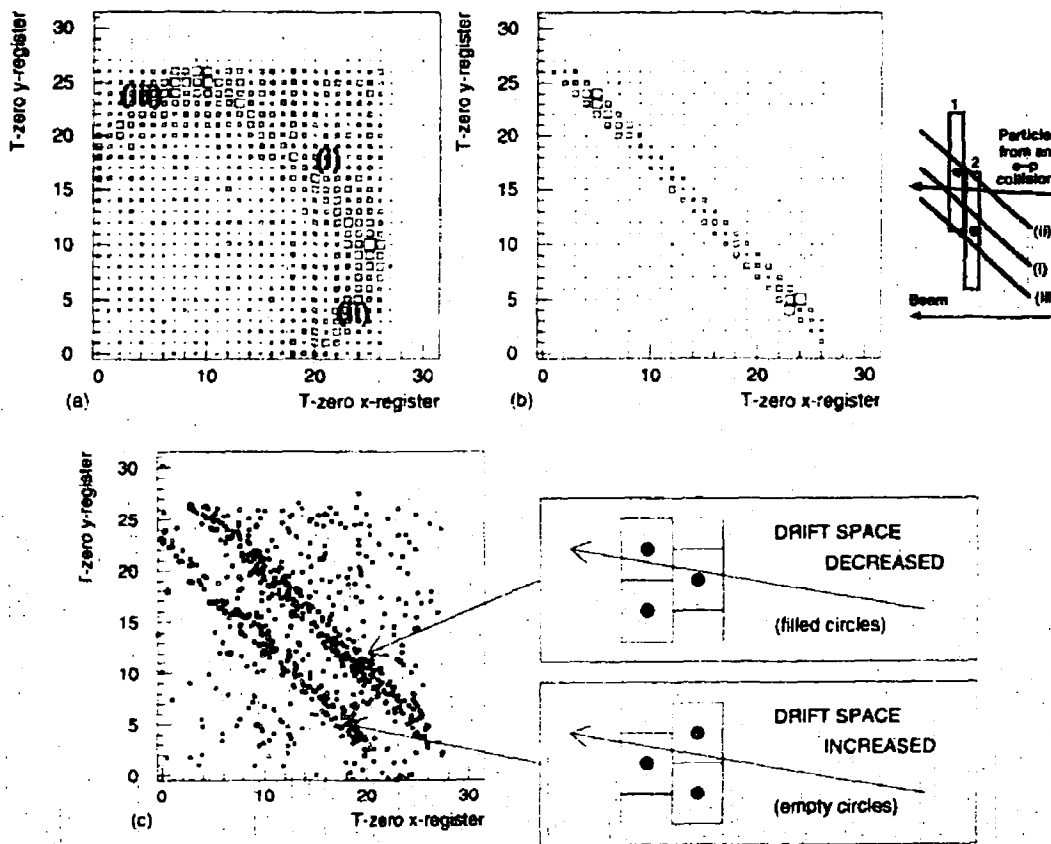


Fig. 20. The performance of the t-zero ASICs.

FADC system used in the chamber readout. The FADC system is designed to give a spatial resolution of $150 \mu\text{m}$, which is equivalent to a time resolution of 3 ns. The results are shown in Fig. 19, where it can be seen that the rms deviation is 17 ns. This resolution should be sufficient to allow the trigger to determine uniquely the bunch crossing responsible for a track, in most cases.

As already stated in Section 5.1, a strobe period of 24 ns rather than 48 ns would be justified to optimise the time resolution, but this would require a much more expensive 64×64 coincidence matrix in the t-zero chip.

8.3. Performance of the t-zero chips

The performance of the t-zero chips is summarised in Fig. 20. To illustrate the environment in which they are operating and their correct functioning, chips at a particular distance from the beam have been selected. In Figs. 20a and 20b the activity in the t-zero chips attached to the two drift volumes of Θ_2 nearest the beam in one octant is shown. In Fig. 20c the activity in the t-zero chips attached to wires 20 and 21 of all octants of Θ_2 is shown. The x and y axes of the plots correspond to the x and y axes of

the t-zero chips. The data plotted represent the state of the shift registers at a time 16.5 HERA bunch crossings after the interaction occurred.

Fig. 20a is very illuminating. It shows several features which allow an understanding of where the particles giving pairs of hits originate. The configuration of the two drift volumes relative to the beam is shown as an insert. The regions of correlated hits in Fig. 20a consist of three lines at right angles to each other. The lines represent tracks coming through the chambers at a steep angle to the beam direction and not originating at the interaction region. The three lines represent tracks from the same origin, but entering the chambers at different positions. For tracks of type (i) the sum of drift times is small compared to a track at normal incidence, which explains why the diagonal is displaced to the top right of the plot. For tracks of type (ii) the correlation is seen at the bottom right hand corner, which corresponds to tracks with a short drift time in drift cell 1 and a long drift time in drift cell 2. For tracks of type (iii) the correlation is at the top left hand corner, which corresponds to tracks with a short drift time in chamber 2 and long in chamber 1. It can be seen that this correlation plot, Fig. 20a, allows the location of the source

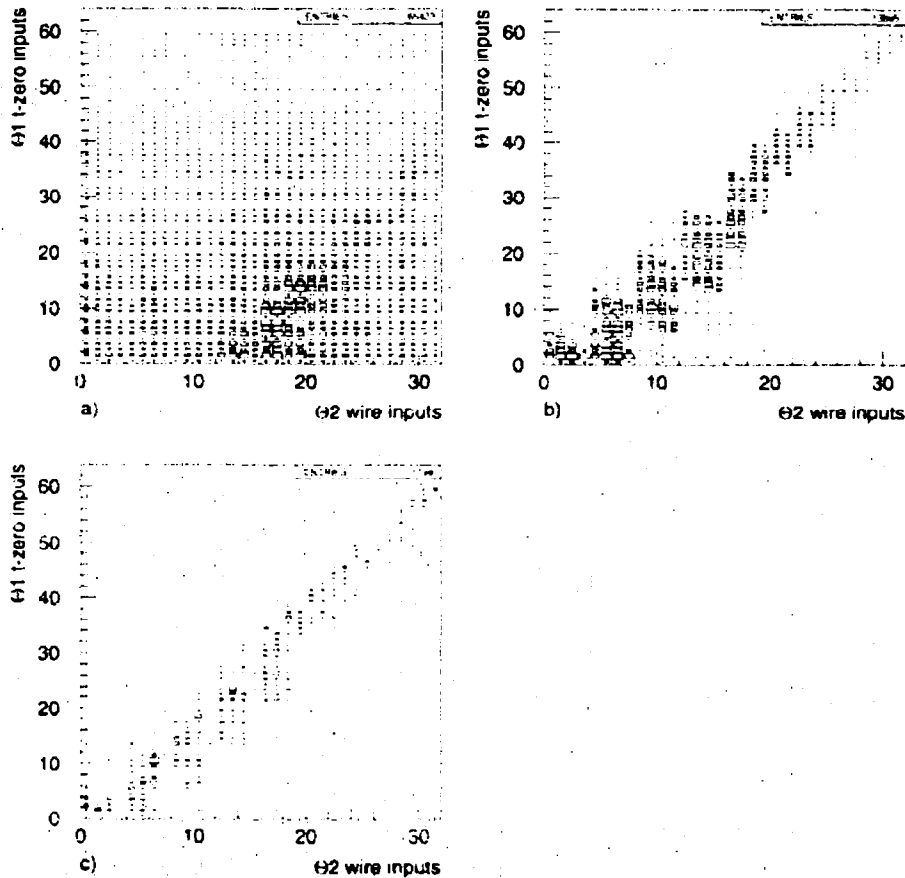


Fig. 21. Performance of the track-finding ASICs.

of tracks to be determined, in principle. Such tracks are background, and are produced by particles interacting with a collimator in the beam pipe of HERA.

We are interested in triggering on tracks originating from the electron–proton interaction region. In Fig. 20b we show the distribution of chamber hits which satisfy our trigger requirement, i.e. events where the t-zero chip has produced an output signal. It can be seen that these hits indicate the coincidence matrix load required to find tracks coming from the interaction region. The events displayed in Figs. 20a and 20b come from the same data sample, and the plots indicate the power of the t-zero chips in rejecting particles which do not come from a given source. By adjusting the load in the coincidence matrix of the t-zero chip the source of tracks which lead to a trigger can be moved along the beam direction.

In Fig. 20c the correlation plots are for t-zero chips attached to drift volumes at large values of θ . The plots show the hits which have produced an output signal. The two diagonal correlations correspond to the configurations shown in the insert, and correspond to the geometrical arrangements of the drift volumes for odd and even-numbered octants of θ_2 at chip no. 20.

Figs. 20b and 20c indicate that the t-zero chips are applying the drift time correlation between adjacent drift volumes in the expected way. Fig. 20a gives an impression of the source of tracks actually traversing the muon chambers.

8.4. Performance of the track-finding chips

The performance of the track-finding chips is summarised in Fig. 21. The data shown are from tracks in planes θ_1 and θ_2 . Once again the data correspond to the register inputs to the track-finding chips at a time 16.5 bunch crossings after the interaction. This is the time at which the trigger decision is made in the track-finding chips. In Fig. 21a the activity in planes θ_1 and θ_2 is shown for all octants. A strong correlation can be seen, which corresponds to tracks with a small θ value in θ_1 and larger in θ_2 . This indicates the same source of particles as shown in Fig. 20a, and is due to tracks emerging from collisions in a collimator. The t-zero correlation has not removed all the pairs of hits corresponding to such tracks, as can be seen clearly in Fig. 20b. Fig. 21b shows the data which survive the tracking correlation, i.e. have given an output signal from the track-finding chips. The accepted data do not fall on a clean diagonal line as one might expect. This is because the coincidence load is computed to accept tracks with a total of only three hits in planes θ_1 and θ_2 . It should be noted that in Fig. 21b the inputs from θ_2 are wire inputs only, while those from θ_1 are the outputs of t-zero chips. For tracks with hits in all four drift volumes of θ_1 and θ_2 , there will be an output also from the track-finding chips taking in θ_1 wire hits and θ_2 t-zero outputs. Fig. 21c shows the correlation that

remains if such a requirement is made offline using software, and therefore indicates the hits accepted by the trigger in the AND mode of operation described in Section 5.2. The numbers of entries in Figs. 21a, 21b and 21c are interesting. They are an indication of the reduction in coincidences between θ_1 and θ_2 which is achieved by requiring pointing tracks.

8.5 Overall system performance

8.5.1. Timing

The forward-muon trigger is organised so that the trigger output signal occurs at a fixed time after the particles have been produced in the electron–proton interaction. This is possible because the t-zero chips can be configured to compensate for the spread in drift times of the particles traversing the 6 cm drift volume of the chambers. The timing resolution of the trigger is determined by the timing resolution of the strobing of the discriminator hits into the x and y registers of the t-zero chips. This was presented in Section 8.2. H1 triggers based on scintillation counters and MWPCs have good intrinsic time resolution, and by comparing the arrival time of the forward-muon trigger bits at the Central Trigger Logic (CTL) with the arrival time of these triggers it is possible to determine with what precision the bunch crossing responsible is being extracted. In Fig. 22 the arrival time of the forward-muon trigger at the CTL is shown for a sample of events which contain a track in the forward-muon chambers. In selecting these events the forward-muon trigger bits were ignored by the CTL, and the selection used well-timed triggers from the forward proportional chambers and the forward end-cap streamer chambers only. The horizontal scale is in units of HERA bunch crossing period, 96 ns. It can be seen that the majority of the forward-muon triggers come at precisely the correct




Fig. 22. Time of arrival of forward-muon triggers.

Table 4
Forward-muon trigger rates

Trigger description	AND-rate [Hz]	OR rate [Hz]	Ratio
FMT - Endcap Muon	0.044	1.9	43
FMT Dimuon	0.006	6.8	1133
FMT - Barrel Muon	0.008	0.3	38

time. The spread is consistent with the time resolution given in Fig. 19.

8.5.2. Trigger rates

During physics runs the rate of hits in the forward-muon chambers is very high. As described in Section 5.2, track finding in the trigger normally requires only three hits out of the possible four. In Θ_1 and Θ_2 for example, tracks can consist of two hits in Θ_1 and one hit in Θ_2 , OR one hit in Θ_1 and two in Θ_2 . Each of the single hits has an effective resolving time of the full drift time in the muon chambers, 1.2 μ s, and the two-plane coincidences will therefore have a considerable random component at high rates. However, the AND of these coincidences requires four hits in Θ_1 and Θ_2 . This had an effective resolving time of about 48 ns during the 1993 running, and therefore a random component which is roughly $48/(1200 \times 2) \approx 0.02$ of the OR rate, which is a big improvement. The

price paid for demanding four out of four possible hits is an inefficiency of a few percent.

Measurements of the forward-muon trigger (FMT) rates for typical running conditions late in 1993 with the trigger in the OR and the AND condition are shown in Table 4. It is clear that the effect of random coincidences is very severe, and in the OR condition the random tracks completely dominate the trigger rate. This is evident from the comparison of single muon and dimuon rates in the two conditions. At present it seems that the trigger will have to be operated in the AND condition with a corresponding loss in efficiency. We are investigating the effect on the OR rates of raising the discriminator thresholds. This is a good thing to do, as quite a large fraction of the discriminator output pulses are due to the discriminators retiring on oscillations on the tails of chamber pulses. These false outputs then lead to production of random tracks. As can be seen from Fig. 18, raising the thresholds will reduce individual wire efficiency. However, requiring the track OR condition, i.e. only three hits required out of four drift volumes, may still allow nearly 100% track efficiency.

8.5.3. Track distributions during physics runs

The forward-muon trigger provides easy access to wire-hit activity in the forward-muon system. For each H1 trigger the contents of all the x and y registers of the t-zero chips are read out. Combining all the hits from Θ_1

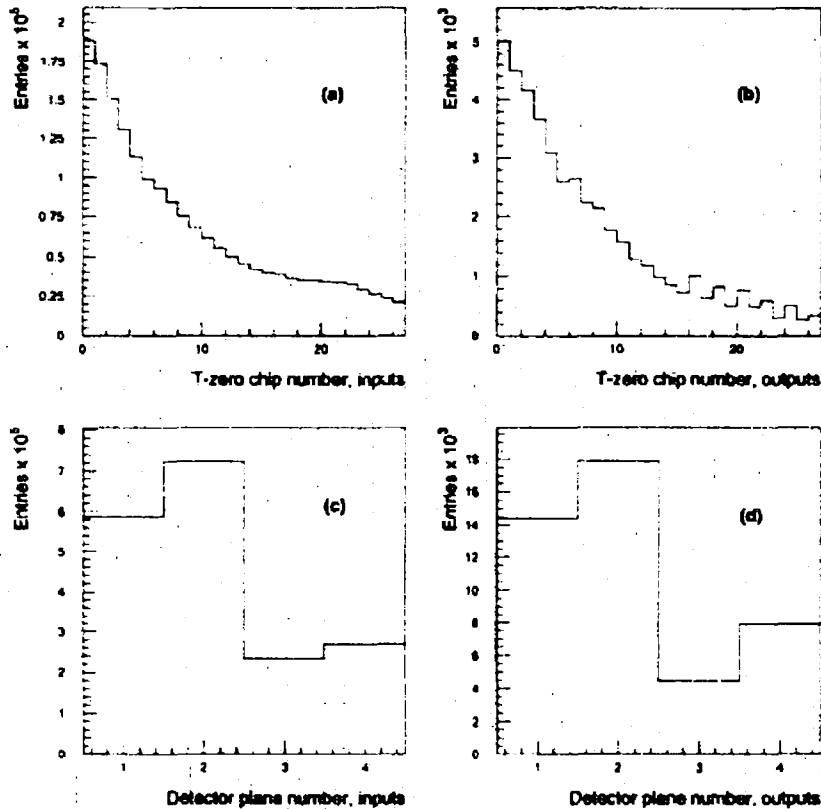


Fig. 23. T-zero ASIC activities for all chamber planes.

to Θ_4 , this distribution is shown in Fig. 23. Fig. 23a shows the profile for all wires hit, and Fig. 23b only for wire hits which lead to a t-zero chip output, i.e. correspond to pairs of hits pointing to the interaction region. Notice the enormous reduction in the number of entries, indicating the level of activity in the chambers overall. In Fig. 23b the profile of the t-zero hits is seen to be structured for t-zero chips near the outer edge of the chambers. This is due to the wire-ORing described in Section 5.2, which has to be done in Θ_3 and Θ_4 to accommodate all the chamber wires in the 27 t-zero chips allocated to each chamber. In Figs. 23c and 23d are shown the total hits in each Θ chamber. The effect of the toroid is obvious.

8.5.4. Muon triggers

The trigger is extremely flexible, and has been configured to trigger on various sources of muons. Alignment of the chambers is done using cosmic muons and beam-halo muons. For triggering on cosmic muons near the horizontal, the trigger is configured to require tracks traversing the four planes for each octant. To achieve this, the coincidence matrices for the t-zero chips are all filled, as well as the coincidence matrices in the track-finding and correlator chips. With the trigger configured in this way the drift cells are effectively continuous detectors.

Beam-halo muons are produced by the coasting proton beam in HERA interacting with residual gas, vacuum chamber or collimators. They travel roughly parallel to the

proton beam and traverse the experiment over quite a considerable area, illuminating the whole forward-muon chamber system. With the toroid switched off they are a source of straight tracks through the system, and as such are useful for geometrical alignment of the chambers. For beam-halo muons which have approximately normal incidence on the muon chambers, all t-zero chips have a simple diagonal load, as do the track-finding and correlator chips. All t-zero chips are loaded identically. Beam-halo data for alignment purposes is regularly taken with the forward-muon trigger configured in this way.

For physics runs, the trigger is configured to select tracks pointing to the nominal electron-proton interaction region. The final-decision modules are configured to require the following trigger elements: a single muon in either of two θ regions, more than one muon in one octant, and more than one muon in different octants. Each of these four is produced for tracks before toroid and also for tracks traversing the toroid, making a total of eight subtrigger elements which are sent by the forward-muon trigger to the Central Trigger Logic.

Fig. 24 shows an example of an interesting event triggered by the forward-muon trigger. It is a dimuon event with a dimuon invariant mass of $3.2 \pm 0.3 \text{ GeV}/c^2$, and is probably an elastic J/ψ event. This event triggered on the single muon trigger and not the dimuon trigger. This is because one muon had only three chamber hits before the toroid due to a non-functioning drift cell.




Fig. 24. Interesting dimuon event in H1.

9. Conclusions

Based on our experience so far, the forward-muon trigger appears to work well. It is able to find, in real time, tracks with a time resolution of better than one HERA bunch crossing, 96 ns, in a drift-chamber system with a drift time of 1.2 μ s. The configuration of the trigger is field programmable. It incorporates coincidence matrices of 32×32 elements which can give a decision every 48 ns. Since these are implemented as part of two semi-custom ASICs, the system is quite compact.

The concept of such a tracking trigger is very relevant to experiments operating on storage rings where the time between bunch crossings is small. In particular, muon triggers for LHC experiments, where the time between crossings will be only 25 ns, could be designed around this flexible, field-programmable concept.

Acknowledgements

We would like to thank the H1 Manchester University group, and especially Phil Biddulph, Mike Ibbotson and

Chris Hilton, for a great deal of help with designing, testing, commissioning and running this trigger in conjunction with their forward-muon drift-chamber system. Derek Newman-Coburn conceived and implemented the very tidy cabling system, and Dennis Grant and Roger Harris also contributed in a number of ways.

This work was funded by the UK Science and Engineering Research Council (now the Particle Physics and Astronomy Research Council), and the Swedish Natural Science Research Council.

References

- [1] Abt et al., The H1 Detector at HERA, DESY 93–103.
- [2] H. Cronstrom et al., Nucl. Instr. and Meth. A 340 (1994) 304.
- [3] RAL Microelectronics Group, Specification for t-zero ASIC, RAL no. 114 (1991).
- [4] RAL Microelectronics Group, Specification for track-finding ASIC, RAL no. 115 (1991).