5 Trigger

5.1 Introduction and trigger requirements

The task of the trigger system is to select out of the flow of signals registered in the multitude of detector channels those for permanent recording that originate from a given $e-p$ interaction of physics interest and to reject background events. The common background sources of other accelerator experiments are also present at HERA: synchrotron radiation from the electron beam, proton gas interaction in the beam pipe vacuum of about 10^{-9} hPa and stray protons, which produce particle showers by hitting the beam tube and other materials around the accelerator. Beam halo muons and muons from cosmic radiation play a rôle, too.

The variety of physics processes under study in $e-p$ collisions covers a wide range of cross-sections and of rates, respectively. It extends from photoproduction, where the visible $e-p$ cross section of several μb implies an event rate of 20–30 Hz at design luminosity ($\mathcal{L} = 1.5 \cdot 10^{31}\mathrm{cm}^2\mathrm{s}^{-1}$), towards W production expected to occur a few times per week. (See Table 5.)

beam gas interactions			50	kHz
cosmic μ in barrel			700	Hz
tagged γp	1.6	μb	25	Hz
$c\bar{c}$ total	1	μb	15	Hz
DIS low Q^2	150	nb	2.2	Hz
DIS high Q^2 (e in LAr)	1.5	nb	1.4	min^{-1}
Charged current DIS p_T >25 GeV	50	pb	3.0	h^{-1}
W production	0.4	pb	0.5	d^{-1}

Table 5: Cross sections and rates (at design luminosity).

The high luminosity mandatory for the measurement of the rare processes could only be achieved with a large number of proton and electron bunches in the machine (see section 2.4). With a basic proton RF frequency of 52 MHz (208 MHz) 210 bunch positions have been foreseen along the circumference of the storage ring; consequently the bunch crossing interval is 96 ns.

This time span has to be compared to typical signal formation times in the detector: the largest drift times in the central jet chamber (CJC) amount to 1 μs and the calorimeter preamplifier has an integration time of about 1.5 μs. (see Table 6.)

The probabilities per bunch crossing for accidental tracks, calorimeter energy deposits or electron tagger signals are at or below the percent level. However, when integrated over the

width of proton bunch	1.4 ns
distance to next satellite bunch	5 ns
flight time to backward ToF	6 ns
flight time to barrel muon system	20 ns
bunch crossing interval	**96 ns**
longest drift time in CJC	1 μs
integration time of LAr preamplifier	1.5 μs
delay of first level trigger	2.5 μs
front end readout time	$\approx$1 ms

Table 6: Time scales at HERA and H1.

detector sampling times, the pile-up probability is no longer negligible. Therefore it is mandatory that the trigger identifies the bunch crossing ("t_0") associated to the event under consideration.

Most subdetectors of H1 generate signals that can be used for the first level trigger. Some examples are: tracks in certain ranges of curvature, energy depositions with their topology, interaction vertex positions, particle arrival times. Local coincidences between MWPC 'rays' and calorimeter 'towers' are also formed. While fast detectors like MWPCs and scintillators intrinsically provide time resolution within a bunch crossing, "t_0" information can be obtained from calorimeter signals using constant fraction methods, or from the drift chambers by requiring special track configurations. In contrast, threshold-like signals extend over several bunch crossings and have to be put into coincidence with external "t_0" bits.

The first level trigger thus provides a decision for each bunch crossing. The fully pipelined system runs deadtime free at 10.4 MHz and is phase locked to the RF signal of HERA. The decision delay of 2.5 μs determines the minimum pipeline length needed to store the full detector data at the front end.

The asymmetric beam energies of HERA give rise to strongly forward-directed event topologies that are hard to discriminate against background from proton interactions with residual gas atoms. Already at first level, the trigger system has to provide a sophisticated identification of the characteristics of an event and to trigger on the complex signatures of physics processes. To this end, the trigger subsystems of H1 deliver information on distinct properties of the event. In several different ways, use is made on the track origin information that uniquely distinguishes $e-p$ interactions from the beam gas background. Early arrival times in the ToF system indicate that the origin of the event is upstream the proton beam, a large distance of closest approach to the beam axis of CJC measured tracks in the (x, y) plane, or along the beam axis (from the z drift chambers) distinguishes against events not produced in the fiducial volume of the interaction region. In addition, a fast estimate of the z vertex position along the beam axis based on histogramming techniques of signals from the central and forward proportional chambers, constrains the vertex even further.

For some $e-p$ events the vertex information may not be the best requirement at the first trigger level. The signature of the scattered electron can be exploited where appropriate. Depending on the Q^2 region under consideration, it is detected in the electron tagger of the luminosity system, in the backward electromagnetic or in the main LAr calorimeter. Next, the topology of the hadronic final state is widely used. Charge current (CC) events are identified by the imbalance of transverse component of the hadronic energy deposition. Jet signatures spot hard photoproduction events, and the requirement of a positive muon signal serves as a tag for heavy quark production. It has even been possible to trigger on heavy quark production events without detection of the scattered electron, relying just on the topology of the registered charged tracks.

An obvious possible extension of the system is to make use of local correlations between the trigger "objects" of different subsystems. "Some track" together with "some muon signal" is more likely being randomly produced than a "track pointing to a muon signal" in space. Probabilities for accidental coincidences are reduced with increasing granularity.

While at first trigger level only the MWPC and LAr calorimeter data are correlated in such a way, full use of the detailed, high granularity trigger data of most subsystems is being made in the intermediate trigger levels 2 and 3. With decision times of 20 μs and 800 μs, respectively, they work within the primary dead time. The level 2 decision is made using purpose built hardware. A *KEEP* decision at level 2 initiates the readout of the front end under processor control. In parallel, a flexible level 3 system with software algorithms running on a RISC microprocessor can refine the trigger and potentially abort the readout.

5.2 Front end pipelining

The time interval between two consecutive bunch crossings of 96 ns is used as the time unit (1 BC) in the following. The time needed to run trigger signals even through a few electronics circuits performing simple logical calculations is usually longer than that. Moreover the large size of the experiment introduces cable delays of several BC. Finally certain subdetectors have a long detector response time which means that the information of these detectors is only available some BC after the event (LAr: 13 BC due to long integration time of the preamplifiers, central drift chamber: 11 BC due to a longest drift time of 1 μs). Such long response times can only be tolerated because, due to the relatively low event rate (compared to a $p-p$ collider), the probability for an $e-p$ interaction per bunch crossing is small (of order 10^{-3}).

The L1 trigger decision (called L1keep signal) is available centrally 24 BC after the real $e-p$ interaction. Further time is needed to distribute this signal to stop the various subdetector pipelines, otherwise the information which belongs to the relevant BC would be lost. The pipeline length varies between 27 and 35 BC depending on the subdetector, which is just long enough to operate the system. For future system designs we would advise to increase this pipeline length to gain more flexibility in the timing of such a system or - even better - to perform signal processing and zero suppression before entering the pipelines and store the information dynamically.

This concept of a pipelined front end system also avoids a huge amount of analog cable delays and allows to reconstruct offline the history of the event over several BC for timing studies and to identify signal pile up.

H1 uses four different types of pipelines (For a more detailed description of the electronics see the chapters of the respective detectors):

- Fast random access memory (RAM) is used to store the digitized information of the drift chambers (central and forward tracker, forward muon system) as well as of the LAr calorimeter for trigger purposes. The analog to digital converters operate at a frequency of 10.4 MHz for the LAr calorimeter and 104 MHz for the drift chambers. They are synchronized to the HERA clock which is also used to increment the 8-bit address of the RAM, operating as a circular buffer. At L1keep time writing into the RAMs is disabled to save the information for the readout process. In addition RAM buffers are used to pipeline information in the drift chamber trigger and in the central trigger.
- Digital shift registers are used to store the single bit information, generated by threshold discriminators followed by a special HERA clock phase synchronization circuit, in the instrumented iron system, the multiwire proportional chambers, the drift chamber trigger branch, the BEMC trigger branch and the ToF and Veto wall scintillator systems. A custom designed gate array contains the clock synchronization circuit and a 8 channel 32 stage pipeline realized as D-Flip-flops. See the section 4.1.4 of reference [9] for a description of this gate array.
- Analog delay lines are used to store the pulse height of the BEMC. The L1keep signal is used to stop a sample and hold circuit from which the digitization takes place during the dead time of the readout.
- Signal pulse-shaping of the LAr and TC is adjusted such, that the signal's maximum occurs at L1keep time. The same type of sample and hold and digitization is used as in the BEMC case.

The timing of the synchronization step and the analog to digital conversion clocks is critical. The information has to be uniquely attributed to the BC the event originated from. The adjustment of the HERA clock phase, which defines this synchronization time in each subsystem, is therefore a major task, which required several iterations of delay curves combined with offline analysis, and it turned out to be a much more complex operation than expected. Online bunch and clock phase monitors ensure stable operation at the correct timing values.

5.3 Trigger level 1

The trigger level 1 system [31] consists of nine different trigger systems each based on the information of a subdetector. The outputs of these systems are called trigger elements (TE). These TE are fed to the CTL where they are combined to various so-called subtriggers. Each single subtrigger suffices to produce a L1keep signal to stop the pipelines and initiate the event readout.

In the following paragraphs the nine systems and the CTL are described.

5.3.1 Vertex position oriented trigger systems

The geometrical origin of the event is the main handle to suppress background at a HERA experiment. Vertices which lie outside the nominal $e-p$ interaction region identify uniquely background events. These TE are therefore in one or the other way used for almost all subtriggers, with the exception of the higher threshold triggers of the calorimeters.

5.3.1.1 The backward time-of-flight system Beam-wall and beam-gas events originating from the proton beam produce showers, which mostly penetrate both scintillator walls (see Section 2.7 of [9]) behind the BEMC. A background (BG) and an interaction (IA) timing window (derived from the HERA clock signal) define for each scintillator whether the hits belong to particles arriving directly from upstream or via the longer path from the nominal interaction region. The signals from the single scintillator sheets of each wall are logically 'ORed' together to form a signal for each plane, and the two planes are then put into coincidence forming the trigger elements `ToF-BG` and `ToF-IA`.

The `ToF-BG` signal is the simplest and most effective background rejection criterium and is therefore applied to most of the physics subtriggers as a veto condition. The logic is realized in conventional NIM electronics.

5.3.1.2 The $z-$vertex trigger The central and the first forward proportional chamber are used to estimate the event vertex position along the beam axis ($z-$axis) (for a more detailed description see [32]). A particle originating from the beam passes four layers of chambers (either the double layers of CIP and COP or CIP and first forward proportional chamber). The first step of the vertex estimator, the so-called rayfinder, therefore combines the four cathode pad signals (see Section 2.5.2 of [9]) which lie on a straight line into an object called ray. In the plane perpendicular to the beam a 16 fold segmentation (ϕ - sectors) is used, such that the rays of each segment are treated separately. A total of 34'400 different rays are examined for each bunch crossing simultaneously.

A histogram with 16 bins along z with a bin width of 5.4 cm is filled according to the $z-$coordinate of the origin of each ray. The rays formed by the correct combinations of pads all enter in the same bin and consequently form a significant peak above the background generated by rays from wrong combinations of pads which are more or less randomly distributed (see Figure 12(a)). Events which have their vertex far outside from the nominal interaction region do not develop significant peaks, in this case the histogram contains only the background from accidentally formed rays.

From this histogram various TE are derived. The `zVTX-t0` TE is activated if there is at least one entry in the histogram. `zVTX-t0` indicates a minimum of activity in the central region of H1 and also serves for bunch crossing identification in combination with low threshold energy triggers of the calorimeters. The two TE `zVTX-sig1` or `zVTX-sig2` are activated if the histogram peak exceeds a given significance threshold (two different threshold settings are available in parallel). For events with a few tracks only, a special TE indicates that there were only few entries in the histogram albeit in a single cluster. The histogram analysis is fully programmable, such that the meaning of the TE can easily be changed.

The rayfinder is based on a custom designed gate array (1.5 μm CMOS technology), of which 2112 pieces are in use. It contains the logic to examine the signals of 45 cathode pads and to form 30 rays (allowing a flexible 3 out of 4 logic), the adder tree to count the active rays, the grouping of these rays into big rays (see below) and some pipeline structure for delays. For the final histogram building and the peak analysis programmable logic cell arrays (XILINX family 3000 [33]) and a 22 bit look up table realized with 4 MByte of fast static RAM are used (see Figure 12(b) for a block diagram).

5.3.1.3 The forward ray trigger The cathode pad signals of the FPC and the CIP are fed into a logic circuit which finds rays originating from the nominal interaction region and pointing in the forward direction (for more detailed description see [34]). A ray here is a set of impacts on three or four chambers, compatible with a track coming from the interaction region in one $2\pi/16$ ϕ-sector. Out of the 6 or 8 chamber planes (2 planes per chamber) which are hit by the track, only one is allowed to be absent. For each of the 16 ϕ-sectors there are 32 possible rays corresponding to radial bands for impacts on the second FPC. The width of these bands increases in geometrical progression, the lowest radius is 21 cm, and the largest one 75 cm, corresponding to angles of 5.63° and 24.72° respectively.

These rays are counted and a TE is activated if at least one road is found. Other TE indicate active rays in adjacent ϕ-sectors (the xy plane is divided into 16 ϕ-sectors as in the $z-$vertex trigger above). Furthermore, special topology conditions in the 16 ϕ-sectors can be used to activate a TE, e.g. a back to back of all rays.

This system is realized by a total of 320 RAMs, which are used as hierarchically organized look up tables.

5.3.1.4 Other MWPC triggers Additional TE are derived from the CIP to trigger on cosmic rays passing through the beam pipe.

If more than three sectors in the backward quarter of the CIP are set, the TE `CIP-backward` might indicate an upstream p beam-gas event. The TE is used as a veto for certain subtriggers which are derived from tracking information only.

5.3.1.5 Big rays The rays found by the forward ray trigger and the $z-$vertex trigger (in the latter case only the rays originating from the highest peak in the $z-$vertex histogram, fig. 12(b)) are extrapolated and combined to 224 "regions of interest" so-called 'big rays', which have the same geometrical division as the 'big towers' of the liquid argon calorimeter trigger (see below and fig. 15). They 'rays' can both be put into coincidence with the 'towers' (fig. 14) but also be used in a stand alone manner by demanding topological conditions (e.g. back-to-back configuration or minimal multiplicity in backward region) on the pattern of MWPC rays.

5.3.1.6 The central jet chamber trigger

This trigger [35] finds tracks in the CJC, which have a distance of closest approach to the nominal beam line of less than 2 cm from the nominal beam axis and therefore suppresses beam-wall events as well as synchrotron radiation background. To keep the number of channels at a manageable level without degrading the performance 10 selected layers out of 56 radial signal wire layers of the CJC are used in the trigger.

In a first step (fig. 13) the signals from the CJC are digitized by a threshold comparator (independent of the normal readout branch) and synchronized to the HERA clock. This way the drift time information is kept with an accuracy of 96 ns corresponding to about 5 mm of position resolution. For the innermost two drift cells this resolution is improved by doubling the clock frequency.

In a second stage the hits are serially clocked into shift registers. Hit masks are defined according to track position in drift space and track curvature in the magnetic field. A total of 10000 different masks are applied in parallel to the outputs of the shift registers to mark the active roads. Tracks with low or high transverse momentum can be distinguished as well as the charge of low momentum tracks (< 1 GeV/c). The number of roads found in each of the 15 ϕ-segments and in the two momentum bins for each charge are counted separately in 3 bit numbers.

In the final step these track counts are processed to generate the TE. Three different thresholds on the total number of tracks are implemented. In addition a topological analysis in the xy plane is performed, for instance back to back tracks in ϕ can be recognized.

Most of the digital logic is programmed into about 1200 programmable logic cell arrays (XILINX [33] family 3000).

5.3.1.7 The $z-$chamber trigger

The $z-$chamber trigger [36] uses the signals of the drift chambers CIZ and COZ in a way similar to the central $r\phi$ jet chamber trigger, utilizing the relatively high spatial resolution obtained from the drift chambers. Signals are synchronized with twice the HERA bunch frequency and stored in shift register pipelines. Their parallel outputs are fed into coincidence circuits used as look up tables for all possible tracks coming either out of the interaction region of 50 cm length (vertex tracks), or from the upstream proton beam region with $20° \leq \theta \leq 90°$ with respect to the beam line (background tracks).

In a first step tracks are found with a granularity of 2.4 mm within the individual CIZ and COZ systems. For tracks originating from the vertex signals from CIZ and COZ are combined in a 96 bin vertex histogram using an analog technique. By gating the system with an OR of the central proportional chambers the search for a peak in this histogram can optionally be further restricted. A resolution of 5 mm for the vertex reconstruction is achieved. The drift cells associated with a vertex track (both left and right solutions in drift space) are stored thus providing a fast θ measurement of a track.

Tracks not associated with the fiducial volume of the interaction region are counted as upstream background, a signal that can be used as an effective veto.

The shift registers and the look up tables are configured in 1060 logic cell arrays (XILINX [33] family 3064 and 3090).

5.3.2 Calorimetric triggers

The selection of deep inelastic $e - p$ reactions is based primarily on calorimetric triggers. These events are characterized by large transverse energy depositions originating from jets and eventually the primary scattered electron. Certain physics channels “beyond the standard model” are expected to lead to similar event signatures. Correspondingly, the calorimeter triggers have to cope with a wide spectrum of trigger observables, from narrow, localized energy depositions (e.g. electrons) to global energy sums such as transverse or missing transverse energy. Even photoproduction events with heavy quarks, with energy deposits as small as 1-2 GeV, should be efficiently detected, when the calorimeter signal is locally put into coincidence with big rays from the MWPC system.

5.3.2.1 The liquid argon calorimeter trigger The liquid argon trigger [37] system is designed to calculate the energy deposited in various parts of the calorimeter as well as the total energy and other global energy sums which can be weighted by position-dependent weighting factors.

The realization of this system contains an analog and a digital part (fig. 14). In the analog part the signals from the calorimeter cells are split off the readout chain after the preamplifier and are separately amplified and shaped to a pulse width of about 600 ns FWHM. Signals within a Trigger Tower (TT) are summed. The TTs are approximately pointing to the vertex and are segmented in 23 bins in θ and in 32 bins or less in ϕ. While the electromagnetic and hadronic signals are still separated in the TTs, the sum of the two is fed into an analog discriminator which turns both signals off for later summing, if the level is below an adjustable threshold ("AGM threshold") as determined by the electronic noise. The same signal is used to determine the time of the signal (t_0) by a constant fraction type method. The boolean number of $t_0 > 0$ is available as a TE (`LAr-T0`). Using the trigger for low threshold photoproduction physics requires a precise adjustment of thresholds already at the analog stage. This is possible by cross calibrating the trigger to the more precise ADC readout [38].

Depending on the θ region, either one, two or four TTs are summed up to 'big towers' (BT), providing finer granularity in the forward direction (fig. 15). A total of 240 such BT's are formed from the $\approx$ 45000 LAr cells, another 12 BT's are derived in a similar way from the signals of the BEMC and the PLUG calorimeter.

The electromagnetic and hadronic signals of each BT are digitized separately by analog to digital converters running at the HERA clock frequency of 10.4 MHz. The digital outputs are calibrated by a RAM look up table and two threshold discriminators are used to look in each BT for a potential electron signature defined by high electromagnetic and low hadronic energy in the respective sections of the tower. Another discriminator look up table marks all BT's to be transferred to the higher trigger levels. The electromagnetic and hadronic parts of each BT are summed and the total BT energies are then available for further processing. A threshold is set on the total BT signal in coincidence with the 'big rays' as derived from the MWPC triggers (see Section 5.3.1.5). The number of these towers is counted, discriminated and provided as a trigger element to the CTL.

The total BT energy is fed into a set of look up tables producing the weighted energy of this big tower for the global sums: For the total transverse energy a weight of $\sin\theta$ is used, the transverse components E_x and E_y are obtained by multiplication with $\sin\theta\cos\phi$ and $\sin\theta\sin\phi$, respectively, to build the missing energy signal. A further channel can be used for arbitrary purposes, e.g. for generating a trigger with uniformly distributed rates over the polar angle by weighting each BT with the inverse of the background rate of its angular region. A 'topological' channel allows to sum up separately the energies in five predefined topological regions (BEMC, central barrel (CB), forward barrel (FB), inner forward (IF), and PLUG). For these regions, the total energy as well as the energies in each of the contributing quadrants are calculated. Furthermore the total energies are calculated for the total barrel (CB+FB) and the forward (IF+PLUG) and backward (FB+CB+BEMC) regions. For the summing of the weighted BT energies, custom specific gate arrays are used with 8 bit accuracy (7 bit plus sign for E_x, E_y).

In the last step further RAM-based look up tables are used to encode the various global and topological sums into two-bit threshold functions provided as TE to the CTL. One such look up table also builds the missing energy from the two signed components E_x and E_y with subsequent discrimination and encoding into two bits.

More details about this trigger and its performance in 1994 can be found in reference [39].

5.3.2.2 The BEMC single electron trigger The purpose of the BEMC Single Electron Trigger (BSET) (see [40, 41]) is to identify scattered electrons from DIS processes in the angular acceptance of the BEMC. The basic concept of this trigger is to provide cluster recognition and to place energy thresholds on the sum of all energy clusters in the BEMC. The trigger is based on energies deposited in BEMC stacks which have a granularity well matched to the transverse spread of electromagnetic showers.

Analog signals from preamplifiers of single wavelength shifters are first added to form stack sums representing a high granularity TE. The summed analog stack signals are then adjusted in their gain and timing in order to provide an equal analog response to energy and to uniquely assign the energy depositions in all stacks to a single HERA bunch crossing. Two thresholds are applied to these signals. A low threshold just above noise level and a medium threshold to be used as a cluster seed. Typical values are 1.3 GeV for the low threshold and 2.3 GeV for the medium one. A cluster identification module then detects the cluster seeds and assigns neighbouring stacks to define clusters. Two TE reflect the cluster multiplicity (one or more clusters, exactly one cluster). The energy of all clusters is summed up and three possible thresholds (CL1, CL2 and CL3) can be placed on this sum, which activate the respective TE. The total energy summed over all stacks exceeding the low threshold is also compared to a threshold and activates another trigger element. Finally the cluster energy and the total energy sum is digitized into an eight-bit number to be used for correlations with other quantities at the CTL.

The trigger has been able to operate in coincidence with a timing veto from the time-of-flight system (see Section 5.3.1.1) and a cluster threshold of 2.5 GeV during the first data taking period. The threshold curve for the cluster energy is shown in figure 16.

5.3.3 Muon triggers

Both the instrumented iron system and the forward muon spectrometer deliver level 1 trigger information, as described below.

5.3.3.1 The instrumented iron muon trigger The instrumented iron system is logically divided into 4 subdetectors (front end cap, forward barrel, backward barrel and backward end-cap) [42]. Each subdetector consists of 16 modules. The wire signals of layer numbers 3, 4, 5, 8 and 12 of each module are used for the level 1 trigger. The pipeline gate arrays used in this system have this output available after the synchronization step, but before the entry into the pipeline. The "OR" of 16 wires of these signals is called an element and all elements of one chamber are again ORed together to form a single layer signal. Any condition on the 5 trigger layer signals of one module can be requested by means of RAM look up tables (e.g. 3 out of all 5) for each module independently. Two signal lines per module are used. One of these contains the system's t_0 information, the second is used for the coincidence condition. The latter is transferred to the level 1 trigger after its timing has been set correctly by means of the t_0 signal from the first hit trigger layer in the module. This determines the correct BC with a resolution of about 20ns. On level 1 only 8 TE are left distinguishing between exactly one or more than one module trigger in different regions of the polar angle. The 64 coincidence bits as well as additional information on the multiplicity per module is then ready for use in the level 2 and level 3 trigger systems.

The instrumented iron muon trigger has been used up to now for two main purposes: in combination with other H1 triggers (especially from the tracking system) to trigger on muons from heavy quark decays or from $\gamma\gamma$ interactions or as stand alone trigger for monitoring : either on cosmic muons or on muons in the proton beam halo.

5.3.3.2 The forward muon trigger The signals from the drift chambers of the forward muon spectrometer are discriminated and fed into an electronic system which extracts the appropriate t_0 and the coordinate within the drift space of chamber hits which correspond to tracks [43]). This is achieved by making use of the staggered arrangement of two adjacent drift cells as shown in fig. 17. For correct t_0 the sum of the drift times of the two cells is defined for any track coming from the collision point, while the difference of the two drift times gives the coordinate within the drift space. A field-programmable 32×32 coincidence array with serial-load shift register axes, built in a 1.5 μm CMOS custom specific chip is used to make the correlations between the two drift times and thus identify bunch number and track position.

A similar coincidence matrix with parallel-loaded axis is then used to define pairs of track segments which correspond to tracks in the chambers before the toroid and separately in the chambers after the toroid. A variable width road may be applied to require that the tracks point to the interaction vertex. In the third step of this trigger processor pre-toroid and post-toroid tracks are fed into a further coincidence matrix which finds pairs of tracks which have traversed the toroid. The momentum selection of the trigger is made by varying the widths of the roads.

The trigger deals with each octant of the forward muon chambers separately. The track candidates found in each octant are allocated to eight regions at different polar angles relative to the beam. The 8-bit hit patterns from all eight octants are fed into a RAM based look up table which counts the number of muon candidates and allows programmable topological correlations to be made. Eight bits of trigger information are then sent to the CTL as TE.

5.3.4 Triggers derived from the luminosity system

The luminosity system runs with an independent data acquisition and triggering system as described previously in section 4.3. However the trigger signals derived from the three detectors of this system are available also to the main trigger system: independent thresholds can be set on the electron energy, the photon energy and the calibrated sum of both. Together with the signals of the veto counter located in front of the photon detector this information is fed into a look up table to form logical combinations. The outputs are connected to the CTL as TE.

So far mainly the electron signal was used to tag photoproduction events.

5.3.5 Central trigger level 1 decision

The information generated by the subdetector trigger systems described above consists of 16 groups of 8 TE, which are connected to the CTL inputs [31, 44]. The 16 groups of TE are fed into pipelines realized as dual ported RAM based circular buffers, which allows to adjust the delays of all incoming signals to the proper BC. In addition the RAMs can be used to study the time evolution before and after the actual event.

The TE are logically combined to generate a level 1 trigger signal. Up to 128 different subtriggers are formed by applying coincidence and threshold requirements. (Lookup tables are used to form 16 fold coincidences of arbitrary logic expressions from up to 11 predefined input bits). A compact trigger description language (TDL) has been developed to keep up with the ever changing demands for new subtriggers and to properly log the logic and the status of the triggers loaded.

The subtriggers are assigned to a given physics event class (physics trigger), to experimental data needed e.g. for measuring the efficiency of a given detector (monitor trigger) or to cosmic ray events for calibration purposes (cosmics trigger).

The possibility to gate the logic with beam information is also useful; cosmic or calibration triggers should only fire when both bunches are empty, physics triggers only, when both are filled.

The rate of each subtrigger is counted separately and can be prescaled if needed. The final L1keep signal is defined by the logical OR of all subtriggers after prescaling and is distributed to the front end electronics of all subsystems to stop the pipelines. Only at this point the primary dead time starts to accumulate.

5.4 Intermediate trigger levels

The two intermediate trigger levels 2 and 3 operate during primary dead time of the readout and are therefore called synchronous. The calculations which are performed in these systems and the decision criteria applied depend on the subtrigger derived in the level 1 system, which acts in this way as a rough event classification.

During the level 2 latency time (a fixed time of typically 20 μs) the level 2 trigger system evaluates a larger number of detector signal correlations typically in a mixed serial/parallel manner. Depending on the outcome of the analysis a fast *KEEP* or *REJECT* signal is issued at decision time. For the level 2 decision various hardware solutions are under construction including a complex topological correlator [45] and a neural network approach to exploit the correlations between the trigger quantities from the various subsystems in a multidimensional space [46]. The massively parallel decision algorithm of these systems makes them ideally suited for fast trigger applications.

Only if the event is accepted at level 2, the time consuming readout tasks such as zero-suppressing of the drift chamber digital signals and the calorimeter analog to digital conversion and DSP processing are initiated. During this time the trigger level 3 system based on a AM 29000 RISC processor performs further calculations [47]. The level 3 decision is available after typically a few hundred μs, in case of a reject the readout operations are aborted and the experiment is alive again after a few μs.

The calculations of both level 2 and level 3 triggers are based on the same information prepared by the trigger level 1 systems described in the previous section: MWPC (Big Rays, $z-$vertex histogram), central drift chamber (number of tracks found in each ϕ-sector by the level 1 trigger), BEMC (individual stack information, energy sums built in the trigger level 1 system), LAr calorimeter (individual big tower electromagnetic and hadronic energies and the global sums, as built in the trigger level 1 system) and intermediate information generated by the main and forward muon triggers. Topological and other complex correlations between these values are the main applications for the intermediate trigger systems.

The events which survive the level 2 and 3 triggers are taken over by the central data acquisition system with a typical maximum rate of 50 Hz (see chapter 7 for upgrading state of this limit). Since this system works asynchronous to the primary trigger system, there is no further dead time involved as long as the level 3 accept rate stays safely below these 50 Hz. The decision times of these systems (level 2: 20 μs, level 3: assume 100 μs average) and the total primary dead time of about 1.5 ms for fully accepted events imply, that the level 1 (level 2) trigger accept rate must not exceed 1000 Hz (200Hz) to be able to run the experiment with an overall dead time below 10 %.

Fig. 18 indicates the front end response time as a function of the instantaneous input rate. The time is independent of the "distance" between two events until at high input rates 2^{nd} order limitations come into play. One such effect has been identified in the drift chamber readout and been cured by introducing deeper buffering at the front end.

At present the level 2 and 3 processing systems have not yet been in use. (However, all accept and reject control signals were available and the system can be run by arbitrarily forcing level 2 and level 3 accept/reject decisions for each event.) The level 1 accept rate had therefore to meet the central data acquisition limitation of 50 Hz. Since the "physics content" of the TE information is already appreciable this can typically be done in a straight forward manner. Special caution, however, is appropriate when the background conditions of the beams change dramatically, which may happen from time to time.

5.5 The level 4 filter farm

The level 4 filter farm is an asynchronous software trigger based on fast MIPS R3000 based processor boards [48]. It is integrated into the central data acquisition system and has the raw data of the full event available as a basis for its decision making algorithms. This allows for online trigger selections with the full intrinsic detector resolution. In 1993/1994 thirty two processor boards ran in parallel. Each board processes one event completely until a decision is reached. The hardware layout is described in detail in chapter 7.

In order to reach a decision in the shortest possible time, the L4 algorithm is split into various logical modules, which are run only if a quantity calculated by the respective module is needed to reach this decision. The L4 modules use either fast algorithms designed specifically for the filter farm, or contain parts of the standard offline reconstruction program. The execution of the modules is controlled by a steering bank containing text in a steering language written explicitly for this purpose. The final decision is based on statements containing logical combinations of numerical or logical values. The event is either accepted or rejected if the statement is true.

Execution of the statement is terminated and the next statement is executed as soon as a sub-condition is false. It is possible to run any statement in test mode without influence on the actual decision. This allows the evaluation of the effect of new statements with high statistics prior to activation and the flagging of particularly interesting events, e.g. for the online event display. In most cases, the first condition in a statement is simply a mask of L1 trigger bits. This scheme allows for high flexibility without changes in the program code and facilitates book-keeping as the steering bank is automatically stored in the H1 database.

A small fraction - typically 1 % - of all rejected events is always kept for monitoring purposes.

In the 1993 runs the filter farm was mainly used to reject events with vertices outside of the nominal interaction region along the beam axis, which survived the relatively weak conditions set in the first level trigger. The largest reduction was achieved using charge division information from the trackers: In a first statement, a histogram of $z-$intercepts is evaluated. These are derived from the projection of straight lines defined by pairs of well separated wire hits in cells of the central jet chamber (CJC). If more than 50 % of all entries are below $z = -75$ cm, the event is rejected. This condition rejects upstream beam-gas and beam-wall interactions. A similar algorithm evaluates multiple hits on neighbouring CJC wires. If the ratio of hit pairs with an absolute angle of less than 10° vs the beam direction to all hit pairs is greater 1/2, the event is rejected. Events with this topology are predominantly due to electron induced background originating from synchrotron radiation. These algorithms are very fast: both tests combined need an average of 8 ms on the processor boards and rejected 50-80 % of the events triggered by tracker triggers in 1992. The full track reconstruction in the CJC is run on events passing the above criteria. Events with only positive charged tracks or no tracks at all in the beam interaction region are rejected. Three or more well reconstructed tracks originating from $z < -1$ m also lead to rejection of the event.

The BSET trigger described in Section 5.3.2.2 suffers from a large background due to particles hitting a single photo diode (mainly synchrotron radiation related). Events with a disproportionate fraction of the total stack energy in just one of the four photo diodes are rejected. In addition false BSET triggers from upstream proton beam background are rejected by the tracker cuts described previously. Muon triggers in the instrumented iron are verified by requiring a reconstructed track in the iron.

As can be seen in the above examples, the L4 filter farm has been rejecting background events based primarily on technical quantities. It has rejected an average of 70% of the input events. At higher luminosity rejection of beam-gas events in the nominal interaction region will become necessary. In addition, the increased number of filled bunches will necessitate the rejection of cosmic ray events.

The filter farm is not only used for event filtering purposes. As reconstructed data of the whole detector merges there for the first time, it is also well suited for monitoring and calibration. The reconstruction modules fill monitor histograms which can be inspected online. Warning messages can also be sent to the central control console, informing the shift crew immediately of potential problems. Calibration data are sent to the data base for immediate use by the online reconstruction process.

5.6 Performance and outlook

With the above described system we reduced the high background rate at trigger level 1 to typically 26 Hz at the present luminosities delivered by HERA. 70 % of these events were rejected by the level 4 event filter, so we ended up with a tape recording rate of 8 Hz at a overall dead time of the system of 8 %.

The total rate of background events into the experiment due to proton beam losses from upstream was expected to be of order 50 kHz. In fact we measured at a total proton beam current of 0.8 mA (0.5 % of the design value) an event rate of 225 Hz in the ToF scintillator wall, which scales with this expectation.

A proton-residual gas interaction rate of about 1 kHz/m can be calculated from the observed vacuum of 10^{-9} hPa and at design luminosity. This scales to 5 Hz within the interaction region (± 25 cm) and at the reduced luminosity condition mentioned above. We effectively observed a rate of about 2 Hz of such beam-gas events triggered with the central tracking triggers.

The rate of synchrotron radiation background turned out to be a bigger problem than expected. However by using the DC track trigger these events can easily be suppressed. For the low Q^2 deep inelastic neutral scattering events with the scattered electron observed in the BEMC this background is a more serious problem, only spatial correlation between BPC and BEMC and relative signal distributions in the BEMC photo diodes (see previous section) give a significant reduction of this type of background.

Offline studies based on the data taken during the first year of HERA operation show that it is possible by setting more stringent coincidence requirements for the level 1 subtriggers to run the system in the same manner (only level 1 and level 4 triggers) at 10 times higher beam currents without cutting significantly into physics acceptance and with the same performance concerning rates and dead times. However at design luminosity (100 times higher than our first year experience) we will have to tighten the requirements on the events in level 2 and 3 trigger systems. But it looks still possible to trigger with high acceptance for physics events, perhaps with the exception of the heavy quark production of which a large fraction of events have not enough transverse energy for the calorimetric triggers to fire and no or only low energy electrons or muons in the final state. Therefore they have to be triggered in the first level by central tracking information alone resulting in a high beam gas background from the nominal $e - p$ interaction region. We will have to use topological criteria in the level 4 or even in level 2 or 3 to recognize these events.