Experience with the First Level Trigger of H1

F. Sefkow, E. Elsen, H. Krehbiel, and U. Straumann
Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany
J. Coughlan
Rutherford Appleton Laboratory, Oxfordshire, UK

Abstract

H1 is a high energy physics experiment at the electron proton collider HERA at DESY, routinely taking data since 1992. Calorimetric and tracking systems record the properties of the particles produced in the collisions in a total of 270,000 electronic channels. The collision frequency of 10.4 MHz necessitates pipelining of all data in the centrally clocked trigger and front end readout systems. This report outlines the architecture of the trigger and data acquisition system and describes the experience gained in the first years of running the experiment. A crucial aspect is the clock phase adjustment and stability of all trigger signals.

I. Introduction

HERA is a unique particle accelerator in the world which collides electrons and protons with each other. The beam energies of 27 and 820 GeV, respectively, open up a yet unexplored regime for the study of electron proton (ep) interactions. The beam energy asymmetry frequently leads to strongly forward directed events in the detector which are hard to distinguish from the intense proton beam induced background. Consequently, strong selection power is demanded of the trigger.

The detection of rare processes like deep inelastic scattering with large momentum transfers requires high luminosity that can only be obtained with multi-bunch beams. The chosen bunch crossing interval of 96 ns allows to anticipate the situation at next generation colliders.

H1 has a clocked and pipelined trigger system operated at 10.4 MHz. Phase stability of the signals from all parts of the detector proves to be a vital requirement. In this report, special emphasis is therefore put on the methods used to determine, adjust and monitor the relative timing of the trigger signals from the various detector components.

II. H1 Detector

The detector has been constructed by a collaboration of 350 physicists from 12 countries. It consists of a number of subdetectors that are complex systems by themselves [1]. A particle emitted from the ep interaction point subsequently traverses the beam pipe, the inner tracking system, consisting of multiwire proportional chambers (MWPCs) and drift chambers for charged particle measurement, a liquid argon calorimeter with 45,000 physical channels, placed inside a superconducting coil of 6 m diameter that provides a field of 1.2 T. The iron return yoke is instrumented with a 160,000 channel streamer tube system for muon identification. The detector systems are arranged asymmetrically to account for the boost of the center-of-mass system, and are completed by an electromagnetic backward calorimeter and a forward muon spectrometer. The raw data size of one event amounts to 3 MByte, and after front end data compression 50 to 100 kByte per event are logged.

III. Trigger and Data Acquisition System

The detector subsystems feed data into front end pipelines and send fast information to the central trigger, where the decision on acceptance for readout of the whole event is made. The full data from all components allow a detailed reconstruction of the physics quantities of the ep collision final state, i.e. the precise energies and flight directions of particles and jets. The fast trigger data are typically less accurate. However, after preprocessing in the trigger subsystems, information about complex properties of the event is encoded in 128 so-called trigger elements and sent to the central trigger.

In the MWPC system, for example, the interaction vertex coordinate along the beam direction is reconstructed. This allows powerful rejection of beam induced background. Other properties are the multiplicity and momenta of drift chamber tracks, or topological quantities like the imbalance of transverse momentum observed from the calorimetric measurement. Even some combined information from several detector components is also available in the form of "big rays" from the MWPC system locally matched to energy deposition given in "big tower" calorimeter granularity. The individual trigger systems have been described in ref. [1, 2]. Here the focus is on aspects of the central trigger [3].

The pipelined first level trigger runs deadtime free at 10.4 MHz and is phase locked to the RF signal of the accelerator. It provides for each bunch crossing a delayed decision after 2.5 μs. This time span determines the minimum pipeline length needed to store the full detector data at the front end. It mainly reflects the size of the experiment resulting in cable delays of several bunch crossings and the signal formation time of some detectors. Typical processing times in trigger subsystems are of order 5 bunch crossing intervals (bc). However, the longest drift time in drift chambers, for example, corresponds to 11 bc, and the calorimeter preamplifier integration time is 13 bc. This can be tolerated since the interaction probability per bunch crossing is small ($< 10^{-3}$ at design luminosity), and pile-up poses no rate problem. Most trigger subsystems are able to deliver a trigger element called t_0 bit that identifies the bunch crossing which triggered the event.

The pipelines are held and deadtime accumulates only in the case of a level 1 'accept' decision. A completely deadtime free system would have implied a substantial investment in storage of data and parallel processors. Instead, the real effort was put into a sophisticated level 1 system at the cost of a few percent deadtime.

Data acquisition proceeds in several stages, partly in parallel. Primary readout takes about 1-2 ms, then the front end pipelines are released again. Further subsystem readout and formatting is done in a parallel and asynchronous manner [4]. Data are placed in multi event buffers; a dedicated processor, the 'event builder', collects the data from each subsystem over an optical fiber ring and assembles the full event. As enough buffers are used, this does not contribute to the deadtime, as long as the total bandwidth can sustain the data rate. After this event de-randomization, time constraints are therefore given by average trigger rates and event sizes rather than by peak values.

The system is composed of three logical components that are shown in Fig.1. The scheme exhibits the twofold role of the central trigger: to provide the decision on acceptance of each event on the basis of the "trigger elements" sent by the subdetectors, and to distribute the synchronized timing signals that control the subsystem readout such that data in the pipelines are properly assigned to the correct bunch crossing.

The primary event rate can be reduced by two intermediate trigger levels. With decision times of 20 μs and 800μs, respectively, they work within the primary deadtime. A level 2 decision, based on more detailed trigger data, is made using special purpose processors and actually starts the front end readout processor. In parallel, a flexible level 3 system with software algorithms running on a commercial microprocessor will refine the trigger and potentially abort the readout. At fourth level, a processor farm acts asynchronously as a filter on the full data to reduce the logging volume [5]. Typically, 5 events/second are sent to permanent storage.

The IEEE VMEbus standard [6] is employed throughout the system. For trigger and control signal transfer in the electronically noisy environment of the H1 hall differential signal standards have been defined and strictly adhered to. The subsystems can also be operated autonomously, or in intermediate, partially submissive modes for development and test purposes.

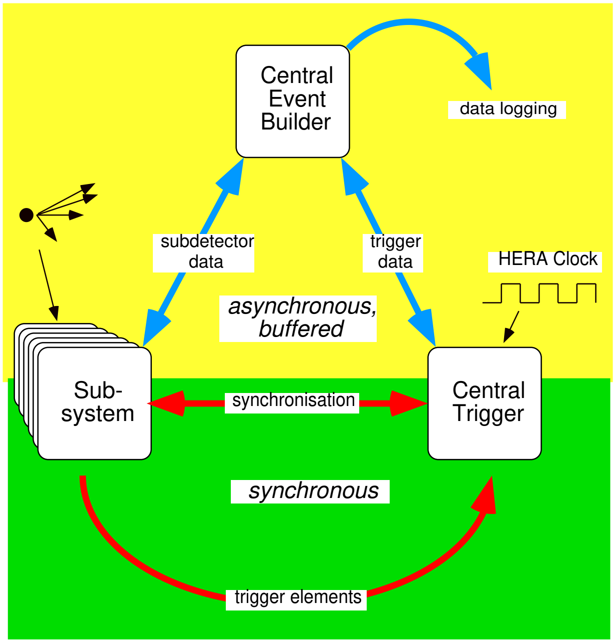

Figure 1: The three logical components of the H1 trigger and data acquisition system. While synchronization between central trigger and subsystems is vital, the event building is allowed to proceed asynchronously, since the event parts have already been marked with a common identification from the central trigger.

Since the beginning of H1, trigger levels one and four have been operational. The interaction rates delivered by the HERA machine, while constantly increasing, have not yet reached their design potential. In 1994, the proton beam current was about 1/3 of its design value, the peak luminosity amounts to 1/5 of the envisaged limit. Background rates were around 10 kHz. The physics selection power of the first level trigger information alone has made it possible to limit the acquisition rate such that second order deadtime effects could be avoided. On average, the system runs with 30..40 Hz level 1 input rate and deadtime around 10-15% whilst maintaining efficiency for physics.

IV. Central Trigger

The tasks of the central trigger are accomplished by in-house developed VME boards. One set of boards synchronize the incoming subdetector trigger data and derive the level 1 decision. A second set performs the level conversion of the signals that are exchanged with the subsystems.

The decision on acceptance of an event is made on the basis of logical combinations of the 128 trigger elements. Every condition must contain at least one t_0 bit that allows to uniquely assign the correct bunch crossing. The combinations are formed using RAM look-up tables, thus providing a maximum degree of flexibility.

The RAMs have 11 address inputs and contain 8 bit data. They are arranged in 16 tables, each consisting of 16 RAMs. The address inputs of a RAM are connected to 8 trigger elements of one input port – usually belonging to one subsystem – and in addition 3 of the t_0 bits. The key role t_0 bits are distributed in such a way, that each is made available in 4 different RAMs of a table (see Fig. 2).

The outputs of the tables drive plain AND logics to form the so-called subtrigger conditions. Any condition of the form

$$s_i = \prod_{0}^{15} f_j(\vec{t}_j) \quad (i = 0,\, 127) \tag{1}$$

can be realized, where the f_j are arbitrary Boolean expressions of 11 trigger elements $\vec{t}_j$. f_j can, for example, be a threshold requirement, imposed on a topological energy sum encoded in trigger bits, and then put in conjunction with a t_0 signal from the MWPC system. A dedicated compiler has been developed that translates Boolean expressions into look-up tables, thereby automatically checking for consistency with the hardware constraints of the logic.

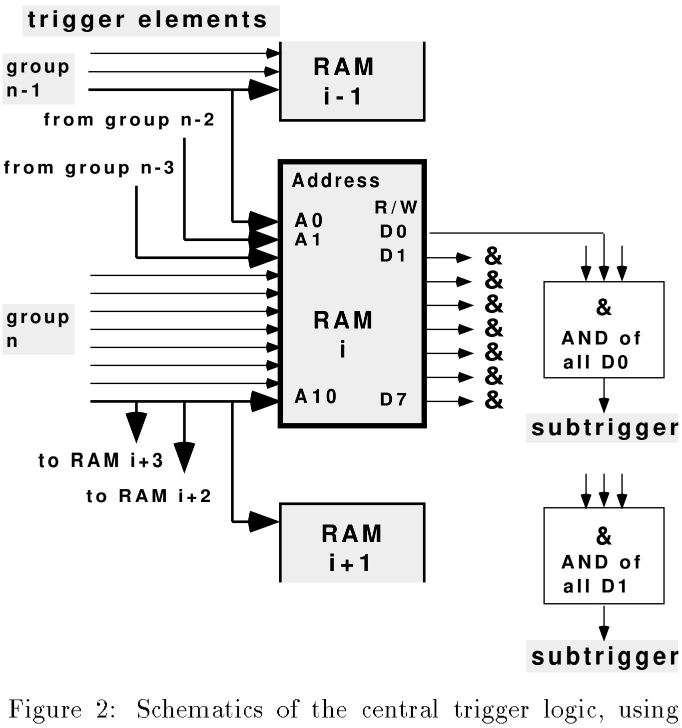

Figure 2: Schematics of the central trigger logic, using RAM lookup tables. The t_0 bits (thick lines) of each trigger element group are distributed also to the RAMs of the next 3 groups; the pattern is repeated in a cyclic manner. Each RAM output goes to exactly one logical AND that forms the subtrigger condition. The trigger elements are connected in an identical way to 16 look-up tables, thus allowing 128 subtrigger conditions.

A valid condition should not necessarily trigger the readout. Triggers with high rates, but nevertheless necessary for monitoring purposes, need to be downscaled. Even some desired physics processes can only be recorded with prescaled rate. A possibility to gate the logic with beam information is also useful; less restricted cosmic triggers, e.g., should only fire when both bunches are empty, physics triggers only, if both are filled. Only the rising edge of a subtrigger condition – if not suppressed by prescaling or gating – leads to subtrigger. The logical OR of the 128 subtriggers constitutes the level 1 trigger accept:

$$\text{level 1 keep} = \sum_{0}^{127} s_i \tag{2}$$

The optimization of this set of conditions is an iterative procedure that needs analysis of the results from actual running. The background rejection power of trigger conditions, the robustness with respect to varying beam conditions and the implications to the system load in terms of trigger rate or average event size are hard to predict. A self-documenting loading scheme with automatic bookkeeping is essential for physics analysis.

The flexibility in the central trigger setup has also proven to be very useful in adapting the data acquisition to beam conditions that vary over the duration of an electron or positron fill generally coasting for a few hours. Three different sets of subtrigger conditions and prescale factors have been routinely used, thereby ensuring low deadtime and optimum acceptance for rare processes during high luminosity, high background periods, as well as allowing less robust triggers to accept high rate processes as beam conditions become less demanding.

V. Synchronization of Trigger Signals

Prior to the formation of subtriggers, the incoming trigger elements are synchronized at the input to the central trigger logic. The synchronizing unit is realized using dual ported RAMs, arranged as circular buffers. It is schematically shown in Fig. 3. Data are written with input addresses controlled by individual scalers, so that offsets between scalers account for relative integer bunch crossing delays. Fractional offsets are adjusted for by shifting the write strobe. The output addresses are driven by a common scaler. The DPRAM thus records the time development of all trigger elements around the time of the event that stopped the pipeline.

An incoming reference clock from the subsystems accompanies each group of 8 trigger elements and automatically accounts for cable delays. It is however not used to strobe the data into the logic, but serves to measure their phases with respect to the central clock, using a 360^o phase-to-digital converter. This is possible since the phase relation of the data with respect to the accompanying clock can in general be kept stable.

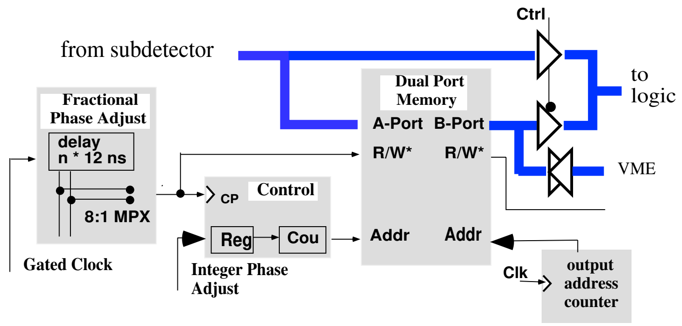

Figure 3: Sketch of the synchronization circuit at central trigger input. Fast subsystems can be delayed by integer multiples of bunch crossing intervals by means of an address offset register. Fractional phase adjustments control the position of the WRITE strobe. 32 addresses are used in a cyclic manner. For fastest transfer of late systems the DPM may be bypassed.

Since the central trigger input circuit only performs the last in a series of synchronization steps in a centrally clocked system, the stability of the phase of the data within a clock cycle is critical. Note that the subsequent logic is of the ripple-through type: subtrigger conditions are formed without further synchronizing steps.

The trigger element timing is required to be stable within margins that are much smaller than a clock cycle. This is especially true for the t_0 bits as obtained, for example, from the MWPC system, from wire plane crossing drift chamber tracks or using constant fraction methods for calorimeter signals. Due to finite resolution, small phase instabilities in a clocked system can lead to sizable fractions of signals shifted by an integer multiple of the bunch crossing period. In the case of a t_0 bit this can cause triggers in too early bunch crossing intervals and would imply that the efficiency for any other coincidence cannot be controlled. Similar requirements hold also for the longer signals, as those derived from mask criteria in the central track chamber or from energy thresholds. The thresholds cannot be well defined without stable time relations.

Several procedures are at hand to determine the correct delay settings. Measured coincidence rates as a function of global or individual delays, at the operators disposition via control of the central trigger, provide fast setup information. On the other hand, readout of pipeline contents has proven very powerful for many online and off-line monitoring purposes, allowing dedicated studies of the time structure of trigger and detector data.

In offline analysis, precise timing measurements can be done even in the case of broad signals. For example, the energy signal in the calorimeter trigger has a FWHM corresponding to 6 bunch crossing intervals. It is read out for nine successive bunch crossings around the nominal t_0. From the maximum of the distribution, the time of the energy deposition can be determined sufficiently precise to assign a proper interval. Thus, in principle all offsets can be determined at once.

Furthermore, phase stability of earlier synchronization steps, within the subsystems, is crucial. The internal time spread of subdetector data has to be kept small. By correlating the timing of different detectors with each other, the internal synchronicity of subsystems can be cross-checked. The track fit in the drift chamber provides timing information with 5 ns accuracy. This has been used, for example, to probe the timing of the calorimeter signals using cosmic rays producing drift chamber tracks and calorimeter showers at the same time. The drift chamber time measurement of cosmic muon tracks has also been used to check the timing of individual channels of the instrumented iron muon system, where large cable lengths have to be accounted for.

Already with present beam intensities, the probability of multiple events has proven to be non-negligible – once rare processes are being looked for. Signals originating from particles from different sources may be overlayed within the recording time of the detector. Although this is no problem for off-line reconstruction, clear trigger decisions require that the information can be properly assigned to the bunch crossing in consideration.

In Fig. 4 such a rare background event in the H1 detector is shown together with the readout of the trigger element pipeline. A cosmic muon entering from the top produced a shower in the LAr calorimeter and triggered the readout of the event. This defines time slice 0. The full data additionally contain the simultaneously recorded tracks of particles originating from an upstream beam gas interaction and signals from a horizontally traversing off-momentum particle. The display of trigger elements however demonstrates that the information related to the three processes has been presented to the logic in a well separated manner: the track-related bits are set in time slice 3, and the off-momentum particle produced a veto signal in a rear scintillator together with the bits from the muon system in time slice 4.

The time-dependent detector information has finally entered into physics analysis [7]. Background pile-up like the example shown, with tracks originating from beam gas interactions, overlayed with cosmic muon induced shower, can fake the signature of a charged current ep interaction, that is characterized by an imbalanced energy deposition and some charged track activity. To the measurement of these rare interactions such events represent important background that can be recognized by its time structure.

VI. Conclusion

Experience with the H1 front end trigger and readout system at HERA with its beam crossing rate of 10.4 MHz has given a taste of conditions at next generation colliders. The centrally clocked and pipelined trigger system has been successfully and reliably operated. Internal monitoring of all timing parameters has proven to be essential.

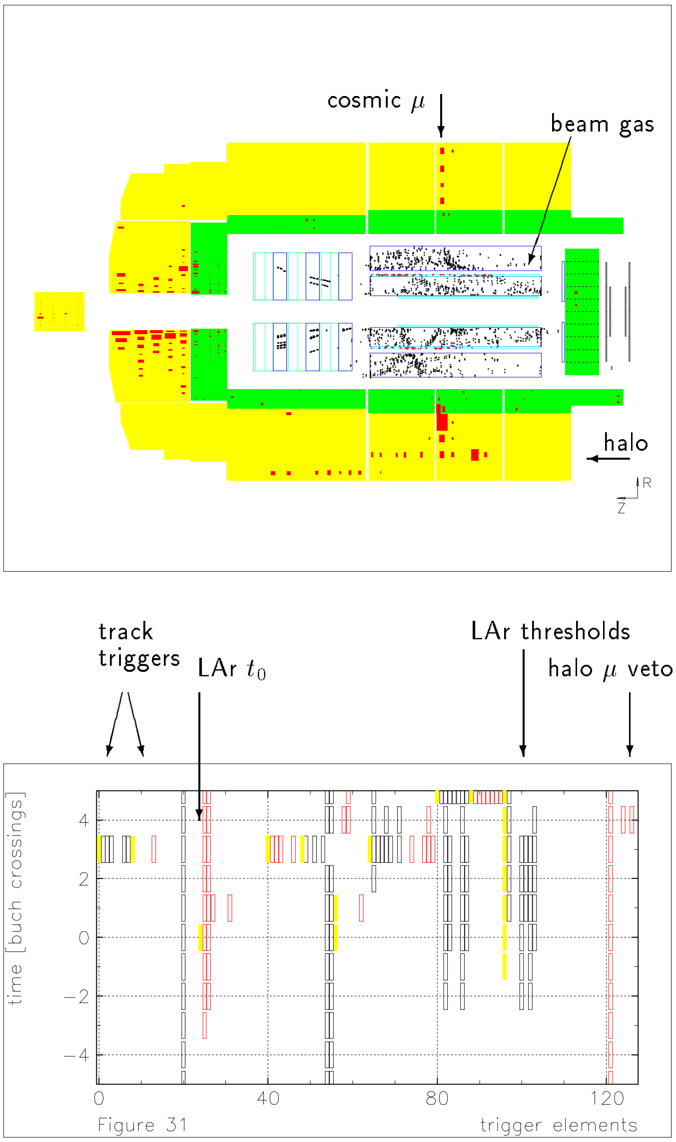

Figure 4: Rare overlay event recorded in the H1 detector (top) and trigger element history (bottom). The trigger, formed through coincidences according to Eq. 1 and 2, defines time slice 0. Positive time slices correspond to bunch crossings (bc) in the future of the event. The cosmic muon traversing the detector from the top to the bottom has triggered the event. The tracks visible in the central chambers have entered 3 bc later, the horizontal track seen in the calorimeter (shaded) shows up in time slice 4.

VII. References

[1] H1 Collaboration (I. Abt *et al.*), The H1 Detector at HERA, DESY preprint 93-103, July 1993.

[2] S. Eichenberger *et al.*, Nucl. Instrum. Meth. **A323**(1992)532.
T. Wolff *et al.*, Nucl. Instrum. Meth. **A323**(1992)537.
B. Andrieu *et al.*, Nucl. Instrum. Meth. **A336**(1993)460.
T. Ahmed *et al.*, A pipelined first-level forward muon drift chamber trigger for H1, *Proc. Int. Conf. on Computing in High Energy Physics 1992*, Annecy, France, 1992, pp. 218-221.

[3] H. Krehbiel, The H1 trigger decider: from trigger elements to L1-keep, H1 internal note 239 (1992), unpublished.

[4] W.J. Haynes, Experiences at HERA with the H1 data acquisition system, *Proc. Int. Conf. on Computing in High Energy Physics 1992*, Annecy, France, 1992, pp. 151-161.

[5] A. Campbell, IEEE Trans. **NS-39**(1992)255.

[6] The VMEbus specification, IEEE standard 1014.

[7] T. Ahmed *et al.*, Phys. Lett. **B324**(1994)241.