A Pipelined First-Level Forward Muon Drift Chamber Trigger for H1

T. Ahmed, J. Dowell, N. Ellis, I. Fensome, J. Garvey, P. Jovanovic, G. Noyes, W. Stokes
University of Birmingham, Birmingham, UK
L. Jönsson, B. Lundberg
Lund University, Lund, Sweden
E. Eisenhandler
Queen Mary and Westfield College, University of London, London, UK
J. Dowdell, M. French, S. Jaroslawski, M. Prydderch
Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, UK

Abstract

We have built a fast, pipelined trigger processor that uses drift-time information from the forward muon chambers of the H1 experiment at HERA. The trigger searches for high-momentum tracks coming from the interaction region, and is based on information from two sets of chambers before the H1 forward iron toroid and two sets after it. The bunch crossing is identified using the offset double layers of drift cells that make up each set of chambers. The processor is implemented using two types of custom chips, which are 32×32 field-programmable coincidence arrays with serial and parallel input respectively. The processing is pipelined at 20.8 MHz, twice the HERA bunch-crossing frequency, and the parallel-input chip has been tested to 35 MHz. After the drift chamber data is shifted in, the processor makes its decision in less than 300 ns.

INTRODUCTION

The new electron–proton collider HERA, at DESY in Hamburg, collides 820 GeV protons with 30 GeV electrons. Head-on collisions of particle bunches occur in the detectors every 96 ns. Since first-level triggers cannot operate at that speed, deadtime is avoided by buffering all detector data and pipelining all trigger logic. The situation is complicated further by slow detectors like the forward muon drift chambers of the H1 experiment, where the information can take as long as 14 bunch crossings to reach the readout electronics. First-level triggering is a significant challenge, and provides valuable experience for the even more difficult situation that will arise at LHC/SSC.

In this paper we describe a fast, pipelined track-finding trigger for muons in the H1 forward muon drift chambers. This digital processor relies on two types of fast ASIC to do matrix coincidence logic. The first of these identifies the bunch-crossing (T_0) in which the collision occurred; this essential information must be obtained from some subdetector for all H1 triggers, since several of the subdetectors require many bunch-crossing periods to gather the data from a single collision. The second type of ASIC searches for muon tracks in 'roads' in the chambers. Both chips have field-programmable coincidence matrices. The use of ASICs permits the trigger to be more sophisticated and flexible than would otherwise have been possible, while keeping the number and complexity of the modules themselves relatively low.

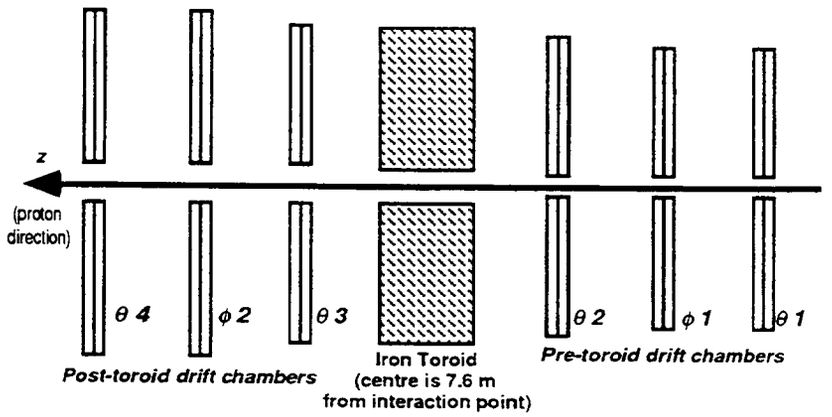

Fig. 1a. Side view of the H1 forward muon detector.

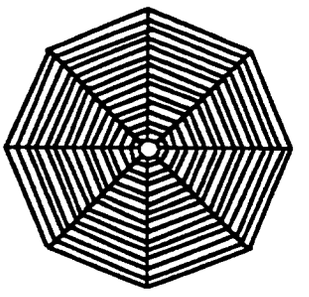

Fig. 1b. Direction of θ chamber sense wires.

THE H1 FORWARD MUON DETECTOR

The H1 forward muon system consists of θ and ϕ drift chambers upstream and downstream of an iron magnetic toroid, as shown in fig. 1a. The forward muon trigger uses only the θ chambers, which measure the distance of the track from the beam axis, and are divided into octants in azimuth (see fig. 1b). There are 1024 θ sense wires.

The H1 first-level triggers and the data from all detectors are pipelined for 22 bunch crossings (2.11 μs) to allow deadtimeless operation while the first-level trigger operates. However, less than 300 ns of that time is available for the forward muon trigger to do its job, because of the long maximum drift time of the drift chambers together with unavoidable cable delays. (The total length of the data acquisition pipelines is 256 bunch crossings.)

Digitizing the wire hit signals is complicated by the fact that the sense wires of two adjacent drift cells are connected by a resistor, and are only read out at one end of each cell. The resistor ensures that the signal from the struck cell is always larger on the wire coming from that cell. The discriminators that digitize the analogue chamber signals incorporate logic to ensure that only the struck wire sends a signal into the processor.

EXTRACTING THE T_0 OF WIRE HITS

A block diagram of the processor is given in fig. 2. After the signals have been discriminated they are sent to the T_0 ASICs shown in the diagram. These ASICs perform the dual function of identifying the bunch crossing in which the track occurred, and also giving the track coordinate within each drift cell with a spatial resolution of about 2 mm. The principle of T_0 identification is shown in fig. 3. The discriminator hits from adjacent drift cells in the two chamber layers are shifted into 32-bit shift registers at twice the HERA bunch-crossing frequency. For a track crossing the chambers at normal incidence, the sum of the drift times in the two layers is a constant. Since tracks are not all perpendicular to the drift gaps, the implementation must allow for deviations from constant drift-time sums. This is done by using a field-programmable 32×32 coincidence matrix between the two shift registers for each pair of drift planes [1]. Good tracks are near the diagonal of the matrix. The coincidence requirement (yes or no) for each cell can be downloaded, and the output axis is the logical OR of each column. In fact, the full 2 mm spatial resolution is not needed because of smearing due to the size of the interaction region and multiple-scattering of muons in the iron magnet yoke. Therefore, depending upon distance from the beam, the output is ORed into groups of either 8 (1.5 cm), 16 or 32 adjacent bits — the finer resolution is used to help sort out the noisy region near the beam.

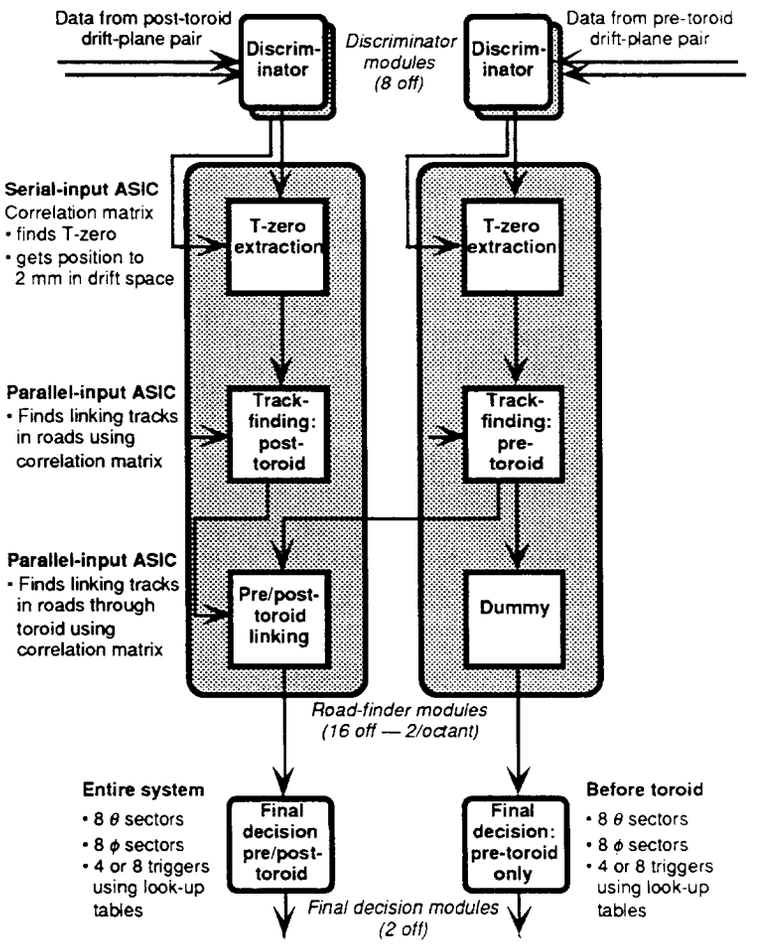

Fig. 2. Block diagram of the forward muon trigger.

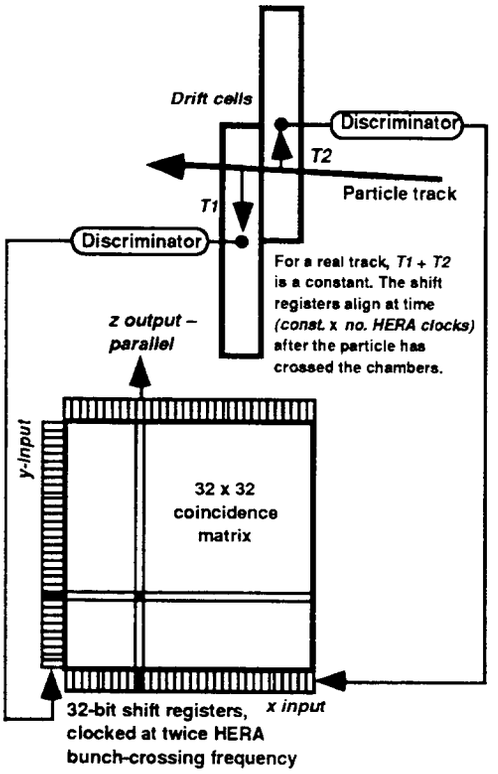

Fig. 3. Principle of the T_0 *ASIC.*

FINDING PRE-TOROID AND POST-TOROID TRACKS

The track elements in each pair of chamber planes are next required to be in valid roads both before and after the toroid. The results from the T_0 ASICs of the pre-toroid planes θ_1 and θ_2, which represent the distance of the track segment from the beam, are fed to parallel-input 32×32 coincidence-matrix ASICs [2]. Here we demand that the tracks point roughly back to the interaction region, consistent with multiple coulomb scattering in the calorimeter and iron yoke, and with bending in the central magnetic field and magnetized iron yoke. Note that multiple coulomb scattering (with deflection inversely proportional to momentum) is the dominant effect in this crude momentum threshold. At the same time, results from the T_0 ASICs of planes θ_3 and θ_4 are also fed to parallel-input ASICs to demand reasonable track candidates after the toroid. To allow for chamber inefficiency, the logic actually requires only one of the two pairs of chambers to have a T_0-validated hit (i.e. both planes firing); in the other a single hit in either of the two planes suffices. Thus, we require 3 out of 4 hits both before and after the toroid.

The third step in the logic correlates the position of tracks before and after the toroid, using the same type of parallel-input coincidence matrix chip as the track finding. This requires that the bending and multiple-scattering in the forward iron toroid are consistent with muons above a momentum threshold (a typical momentum threshold is 5 GeV), and produces the final results for each octant of the detector.

COMBINING THE OCTANTS

The results for each octant, subdivided into eight regions at different polar angles to the beam, are passed to 'final decision' logic that counts muon candidates in order to apply multiplicity requirements as well as allowing topological demands such as muons back-to-back in azimuth. This logic makes extensive use of RAMs to allow great flexibility, first on the individual octants and then on the entire system. A block diagram of this logic is given in fig. 4.

IMPLEMENTATION OF THE TRIGGER PROCESSOR

The processor uses four large crates (9U × 40 cm) for the trigger logic, with a VME crate for control, monitoring and readout. Each large crate handles two octants, and each octant requires one discriminator module and two 'road-finder' modules, one for pre-toroid and one for post-toroid. The final decision logic occupies two modules. There are 54 T_0 ASICs on each of the 16 road-finder modules. These chips are built in 1.5 μm CMOS using a 44-pin package. The pre- or post-toroid road finding is done by four parallel-input ASICs on each road-finder module, while a fifth one on the post-toroid module correlates the pre- and post-toroid tracks. (The pre-toroid module has a dummy ASIC so the two modules can be identical.) The parallel-input chip is a 120-pin package. The system runs at 20.8 MHz, but the parallel-input ASIC has been tested to 35 MHz.

Both types of ASIC implement the coincidence matrix using a full custom cell, similar to one designed for OPAL [3], in order to save space. The control logic is semi-custom. The three matrix axes are linked by a scan path, and this is continued on the road finder modules to link all the ASICs. This allows full readout without using an excessive number of pins on either the chips or the module.

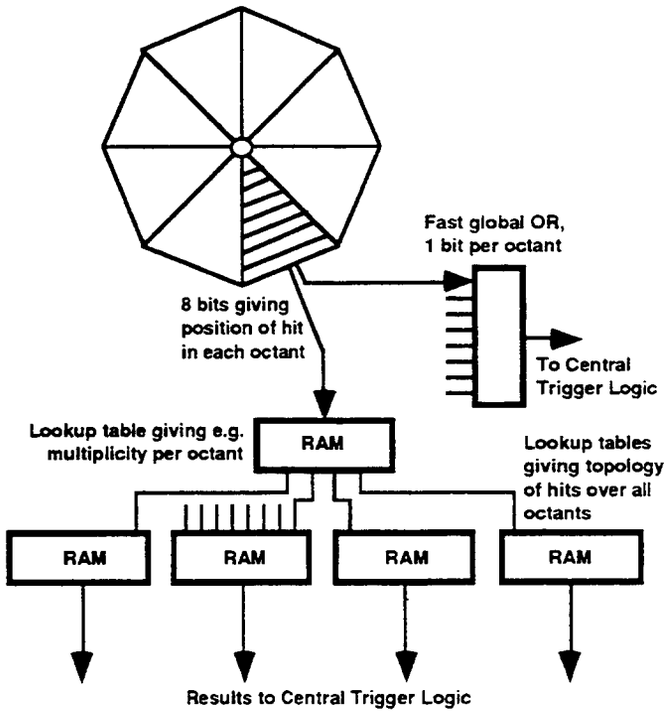

Fig. 4. Block diagram of the final decision logic.

The processor is controlled by a VME-based CPU module linked by a parallel bus to crate controllers in each trigger crate. Readout information to allow the trigger system to be monitored both online and offline is stored in multi-event buffers in the crate controllers before being passed to a VME-based DSP for zero-suppression. It is then sent to the H1 data acquisition system by a second CPU module, which also handles communication with the H1 central trigger logic and does some low-level monitoring tasks. Overall control of the system and high-level monitoring is done by an Apple Macintosh® IIci computer in the H1 control room.

A somewhat more detailed description of the system is available [4].

CURRENT STATUS

At present the trigger crates, VME modules and Macintosh are installed and communicating with the H1 central trigger logic and data acquisition systems. A minimal one-octant pre-toroid trigger using prototype chips and modules has been working for several months. All discriminator and final decision modules are installed. Prototype ASICs of both types were successfully tested and final production is arriving. The final version of the road-finder module has just been tested, and full production will arrive in a few weeks. Much software is ready. We expect to have the entire system running in late Autumn, 1992.

Acknowledgements

We are grateful to many people of the technical staffs at Birmingham, QMW and RAL for expert help on this project. We have also benefitted greatly from the advice and cooperation of too many H1 colleagues to list individually. This work has been funded by the Science and Engineering Research Council (UK) and the Wallenberg Foundation (Sweden).

References

[1] Specification sheet RAL–114, *T-Zero I.C.*, Rutherford Appleton Laboratory, Signal Processing and Microelectronics Group, 1991.

[2] Specification sheet RAL–115, *Track Finder I.C.*, Rutherford Appleton Laboratory, Signal Processing and Microelectronics Group, 1991.

[3] M.J. French and F. Slorach, *An OPAL 32-Bit Coincidence Array Integrated Circuit* in *International Conference on the Impact of Digital Microelectronics and Microprocessors on Particle Physics*, Trieste, 28–30 March 1988, World Scientific Publishers, Singapore.

[4] J. Dowdell, M. French, S. Jaroslawski and M. Prydderch, *Fast Pipelined Trigger Processor for the Forward Muon Chambers of the H1 Experiment at the Hadron Electron Ring Accelerator (HERA) at DESY, Hamburg*, IEEE Nuclear Science Symposium and Medical Imaging Conference, Santa Fe, November 1991.