Concept, Design and Performance of the Second Level Triggers of the H1 Detector

T. Nicholls[1], M. Charlet[1], J. Coughlan[3], E. Elsen[1], D. Hoffmann[1], H. Krehbiel[1], H.-C. Schultz-Coulon[1], J. Schütt[2] and F. Sefkow[4]
[1]Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany
[2]II. Institut für Experimentalphysik, Universität Hamburg, Hamburg, Germany
[3]Particle Physics Division, Rutherford Appleton Laboratory, Oxfordshire, UK
[4]Physik-Institut der Universität Zürich, Zürich, Switzerland

Abstract

Since the end of the 1995 HERA running period the second level (L2) trigger system of the H1 experiment has been in operation. The L2 decision is presently derived from a combination of two independent hardware systems; one is based on the output of neural networks while the other decides on the grounds of topological event signatures in the detector. In contrast to the first level trigger (L1), which is fully pipelined and operates at the HERA bunch crossing rate of 10.4 MHz, the decision of the L2 systems is delivered within 20 μs of the preceding L1 keep signal, allowing sufficient time for transmission of trigger information and computation of the L2 trigger response. This scheme significantly reduces the trigger rate, with minimal impact on detector deadtime for L1 rates below 1 kHz. The L2 trigger has proven to be flexible and highly effective in isolating physics processes from background.

I. INTRODUCTION

The HERA accelerator, which collides electron and proton beams at energies of 27 and 820 GeV respectively, has been running for several years, allowing the study of a wide range of physics topics in a previously unexplored regime. Two colliding-beam experiments, H1 and ZEUS, are designed to exploit this unique opportunity.

In order to achieve the high luminosities required for high precision measurements and the investigation of rare processes in electron-proton interactions, HERA delivers multi-bunch beams with a collision frequency of 10.4 MHz. This necessitates pipelining of all data in the front-end readout systems of the experiments, placing severe constraints on the time available for a first level trigger decision. Furthermore, the intense background induced by the proton beam is often difficult to distinguish from interesting physics processes and demands a trigger system with a high rejection power.

Since the increasing luminosity delivered by HERA requires more sophisticated event triggering procedures to separate specific physics channels from background, the second level of the H1 trigger system was incorporated into the existing scheme at the end of the 1995 running period. This report describes the concepts of the L2 system, with special emphasis on its implementation within the central trigger decision. Results on its general performance are reported.

II. THE H1 DETECTOR

The H1 collaboration has constructed a general purpose detector [1], consisting of a number of complex subdetectors. These are designed to provide both complementary and redundant measurements of various aspects of high-energy electron-proton collisions.

Closest to the interaction point, surrounding the beam pipe, is a silicon tracker for precision measurements of charged particle tracks and vertices. Outside this lies a tracking system of interleaved multiwire proportional chambers (MWPCs) and drift chambers, providing the primary measurement of charged particles. A highly segmented liquid argon calorimeter (LAr) with 45,000 channels surrounds the tracking detectors. A superconducting solenoidal magnet of 6 m diameter encloses the inner detectors and generates a magnetic field of 1.2 T. The outer iron yoke, which provides a return path for the magnetic field, is instrumented with limited streamer tube chambers for muon identification.

The centre-of-mass system of electron-proton collisions at HERA is strongly boosted in the outgoing proton beam (forward) direction, and the systems described above are arranged to reflect this. Drift chambers and a toroidal iron magnet in the forward direction provide additional detection and momentum measurement of muons emitted at low angles. In the backward direction a lead scintillating-fibre calorimeter (SPACAL) enables precise measurements of the scattered electron. Additional detectors, situated at some distance from the detector in both directions, provide measurements of particles at very low angles and of the luminosity delivered to the experiment.

The large number of channels in the H1 detector ($\sim$270,000) leads to a raw event data size of approximately 3 Mbytes. After compression of data at the front end, 50 – 100 Kbytes of data per event are logged to tape.

III. TRIGGERING AND DATA ACQUISITION AT H1

The H1 trigger system is divided into four levels, of which two — the first level trigger (L1) and the asynchronous online filtering (L4) — have been in operation since the start of data taking in 1992. Figure 1 shows the hierarchy of the trigger system.

The L1 system is phase-locked to the HERA accelerator clock (HCLK) signal of 10.4 MHz and provides a trigger decision for each bunch crossing after 2.3 μs without causing

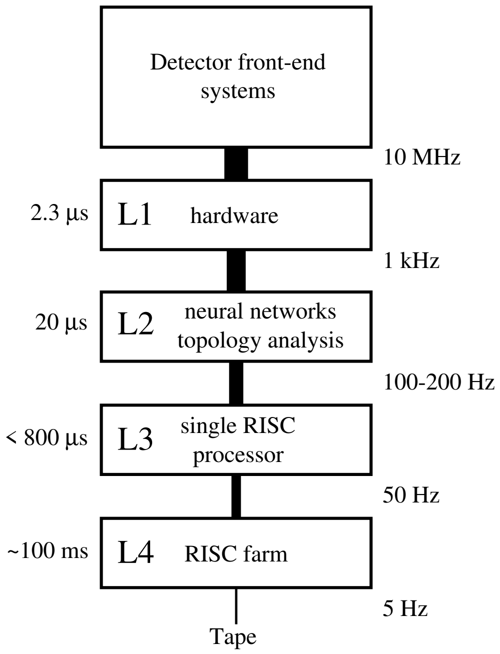

Fig. 1 Hierarchy of the H1 trigger system, showing the decision times and typical input rates at each level. Note that L3 is not used presently, and the L4 input rate of 50 Hz is achieved by L2 alone.

deadtime. The subdetector systems feed data into front-end pipelines. A subset of these data is used by most systems to generate fast information concerning the general properties of the event. This is encoded in Boolean decisions known as trigger elements.

Typical properties used to generate L1 trigger elements are multiplicity and momenta of charged tracks in the drift chambers or localised energy deposits (clusters) in the calorimeters. Trigger information in the highly segmented LAr calorimeter is defined using a geometry based on "big towers" which point to the interaction vertex. The fast signals generated by the MWPC systems can be used to reconstruct the interaction vertex position along the beam direction [2]. This allows a powerful rejection of beam induced background which typically arises from interaction of the proton beam with accelerator hardware or residual gas molecules in the beam pipe upstream of the detector. Time-of-flight information, measured in various scintillator systems and the backward calorimeter [3], is also used to reject background.

The L1 trigger latency, and therefore the front-end pipeline length, reflects the signal formation time in the subdetectors, cable delays and the processing time in the trigger subsystems. The maximum drift time in the tracking chambers is of the order of 11 bunch crossings (BCs). Charge collection and integration in the liquid argon calorimeter requires 13 BCs and the LAr trigger information, including the summation of all energies, is therefore the last to become available for the L1 decision.

Up to 192 trigger elements are received from the trigger subsystems by the level 1 logic of the central trigger (CTL1) [4]. Here they are individually synchronised and aligned to ensure that information from the same bunch crossing enters into the trigger decision. A history of the trigger elements is read out for monitoring and alignment purposes.

The L1 trigger decision is made on the basis of 128 logical combinations of the trigger elements, known as subtriggers. The combinations are performed in RAM lookup tables which drive wired-AND logic to form the subtrigger conditions. This configuration is designed to give maximum flexibility. It is desirable to adjust the rate of some subtriggers in order to achieve a suitable mix of useful data and to restrict the overall rate; this is achieved by prescaling. A count-down scaler is provided for each of the subtriggers, the output of which is set when terminal count is reached. Furthermore, it is possible to gate each subtrigger with beam information such that, for instance, physics triggers only fire when both bunches in a crossing are filled, or certain monitor triggers are enabled in empty bunches only. The final L1 keep (L1K) signal is then simply the logical OR of all 128 prescaled and gated subtrigger conditions.

A number of sets of subtrigger conditions and prescale factors are defined to provide low deadtime and trigger rates suitable for the evolving beam conditions that occur during the lifetime of one accelerator fill. This highly flexible system allows the acceptance for both rare and high rate processes to be optimised simultaneously as conditions allow.

A positive L1K signal is broadcast by the central trigger, which is responsible for distribution of all clock and control signals to the subsystems. At this point the front-end pipelines are frozen and deadtime begins. The L1K is then validated by subsequent levels of the trigger; L2 and L3 operate in more detail on the information from the subdetectors during the primary deadtime, requiring 20 μs and up to 800 μs respectively to complete a decision. The L2 trigger delivers a decision based on two independent hardware systems. An L2 keep (L2K) signal causes readout of the front-end data to begin, while an L2 reject (L2R) allows the pipelines to be re-enabled and deadtime ceases. The L3 system, based on software algorithms running on a RISC processor, refines the L2 decision, potentially causing an abort of the readout if the event is to be rejected. The L3 system has been tentatively used in previous years but is hampered by the restriction on the input rate due to the rather long decision time. It is therefore not currently used for physics selection and will not be discussed further.

The subdetector systems read out and format data in parallel, which are then placed in multi-event buffers. The formatted data are acquired asynchronously by the central data acquisition [5] over an optical fibre ring and the full event is assembled by a dedicated processor, known as the event builder. The use of sufficient buffering prevents this stage from contributing to the deadtime, providing that the total bandwidth is not exceeded. A further level (L4) of triggering operates asynchronously to reduce the data logging rate and volume. This takes the form of a farm of RISC processors [6], each running a reduced version of the offline reconstruction code and operating on a full event in approximately 100 ms. The output of the L4 farm is logged to tape at a typical rate of 5 – 10 events per second.

The trigger and data acquisition systems of H1 are implemented in custom-designed and commercial modules conforming to the ANSI/IEEE VMEbus standard [7]. All trigger control and timing signals are transmitted in a current-driven differential format to reduce pickup in the noisy environment of the experimental hall.

IV. The L2 Trigger Systems

The L2 trigger is provided with all the raw trigger data of the independent L1 trigger subsystems. This allows decisions based on correlations of trigger data at a much higher granularity than the L1 trigger elements and enables a further validation of the L1 subtriggers.

The transmission of data to the L2 systems occurs after the L1K signal has been set. This is achieved by means of the H1 Parallel Quickbus Zero-suppression Processor (PQZP) system [8] which consists of two types of custom VME modules. The Store Card pipelines up to 64 bits per bunch crossing in RAM. After an L1K, information corresponding to the bunch crossing of interest is transmitted to the L2L3 card, which contains FIFO memories and zero-suppression logic, driving 16-bit buses to the L2 (and L3) systems. Multiple store cards can be used to increase the amount of data sent. This system has the advantage of allowing the data sent to L2 to be partially pre-formatted before transmission.

The L2 data are typically segmented into bins in the polar and azimuthal angles (θ and ϕ). Information from the following subdetectors is sent to the L2 systems: MWPCs (vertex position analysis, map in (θ,ϕ) of pointing tracks), central drift chambers (45 ϕ-sectors for positive and negative tracks with high and low momenta), SPACAL (global energy and cluster maps), LAr (global energy sums, energies and thresholds for big towers), central and forward muon systems (hit maps). The central trigger also sends the L1 trigger elements and subtriggers and control information.

The L2 data are processed to produce L2 trigger elements by two independent systems, operating in parallel. These make decisions based on the output of properly-trained neural networks and an analysis of topological quantities respectively. A detailed discussion of the implementation of these systems is beyond the scope of this report; instead we outline the basic principles of operation.

The L2 neural network (L2NN) system [9] generates up to 16 L2 trigger elements. These are realised in pattern recognition modules which consist of several custom and commercial VME cards. A data distribution board (DDB) receives the L2 data and preprocesses them into a suitable format. The outputs are fed as 8-bit quantities to modules containing CNAPS 1064 processors [10] which can implement various types of neural networks. Each processor receives a maximum of 64 such inputs and computes an 8-bit output within 8 μs. The output value is compared to a threshold to define a trigger element, which is fed back to the DDB and sent to the central trigger logic.

The neural networks are trained using previously recorded data samples, consisting of equal mixtures of background and real events from the physics process of interest. In order to check for any bias arising from the trigger conditions applied to select the physics training samples, Monte Carlo simulated events are also used to verify the action of the network. The training is performed either directly on the pattern recognition modules or in computer simulation.

The L2 topological trigger (L2TT) system [11] is based around a 16 $\times$ 16 matrix which represents the geometry of the detector in (θ,ϕ) space. The L2 data are received by custom acquisition modules and preprocessed into a Boolean "projection" onto the matrix. Calorimeter big tower information is compared to a set of thresholds and combined in various manners with track and muon data. A maximum of 150 such projections can be prepared and analysed sequentially during the L2 decision time, at the rate of one per HCLK cycle. Each projection is loaded in turn into the topology module which performs an analysis of the topological properties of the event. The estimators extracted from the analysis, in combination with the maximum energy threshold applied to the calorimeter big towers giving a non-trivial topology, are used to calculate a figure of merit for each projection representing the "distance to background" of the event. The distance variable is evaluated for each projection and mapped by means of lookup tables onto one of 16 "machines" which represent topologies of interest. The sum of the distances accumulated by each machine is compared to a threshold to generate a trigger element. A DSP controls the projection sequencing, machine calculations and trigger element delivery to the central trigger.

The configuration of the L2TT system is determined by the analysis of previous data samples, both for interesting processes and background. An online monitoring system, which is able to capture L2 trigger events at a high rate during running, is used for the development of L2TT strategies.

The complementary approaches of the two systems afford maximum flexibility for the L2 trigger. Multi-layer neural networks are able to deal with arbitrary hyperplanes in a space of continuous input variables such as energies or the number of entries in certain bins of a vertex distribution. A sigmoid transfer function allows detailed adaptation to the resolution of the measured input variables and optimisation for efficiency and small background. On the other hand, the topological approach easily deals with discrete symmetries of the input variables and can be used, for example, when a local spatial match between subdetectors proves beneficial. The two systems therefore have different regions of applicability and together lead to an optimal signal to noise ratio for the various physics processes triggered in the H1 detector.

V. The L2 Decision Logic

The L2 systems described above send their trigger elements to the central trigger level 2 decision logic (CTL2), the purpose of which is to evaluate the final L2 decision (L2K or L2R). Logical combinations of the L2 trigger elements are made, forming subtriggers which are used to validate certain L1 subtriggers. This assignment is vital for the efficiencies of the L2 subtriggers to be well defined. Furthermore, it is necessary to allow some L1 subtriggers to create an L2K

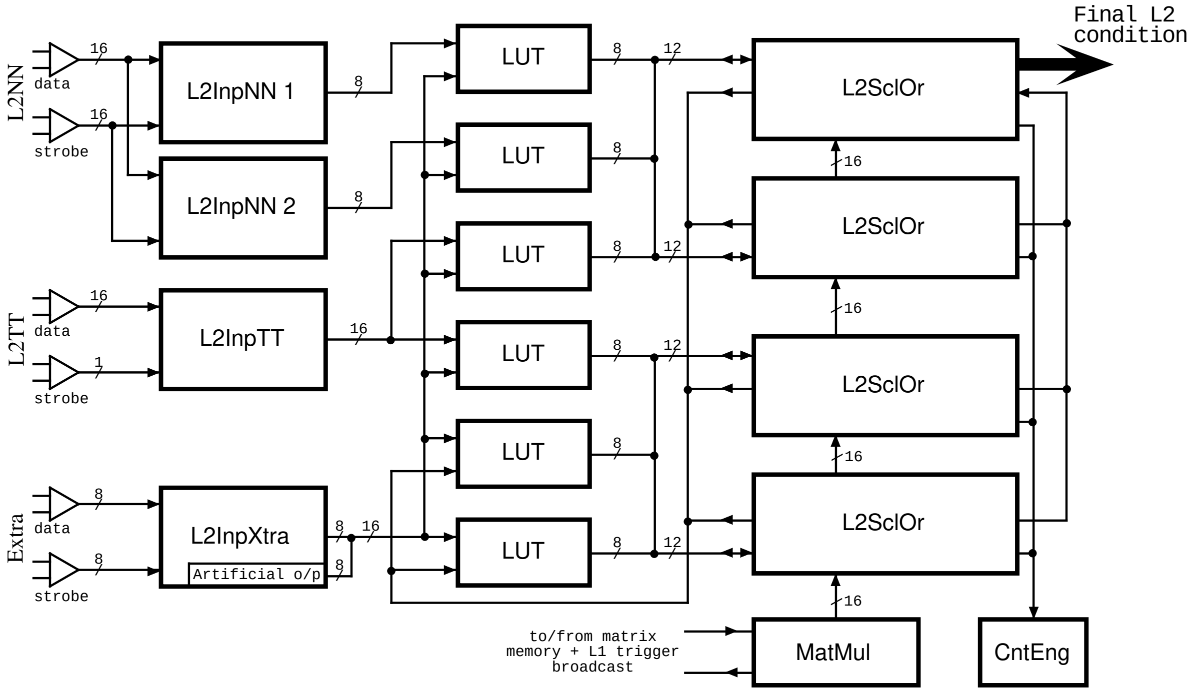

Fig. 2 The L2 decision logic of the central trigger, showing the main flow of data. Control and ancillary data lines are not shown for the sake of clarity.

without validation by L2. The CTL2 therefore has a highly configurable architecture, which allows maximum flexibility in the assignment of L2 conditions to L1 subtriggers.

Figure 2 shows the functional flow of data through the CTL2 logic, which is able to receive up to 16 elements from each of the two systems. In addition, there are eight extra inputs intended to receive signals from a future L2 trigger processor. The CTL2 logic is located on the Fast Card of the central trigger, which coordinates the trigger sequencing for the entire detector. Much of the CTL2 functionality is implemented with Quicklogic programmable ASICs [12]. This approach was chosen for flexibility, ease of configuration and to achieve the speed required for determination of the final L2 decision.

The L2 trigger elements arriving at the central trigger are converted from the current-driven differential format to unipolar by input receivers. The L2 systems are required to actively send either an accept or reject decision for each trigger element, and therefore transmit additional signals to strobe valid data into the CTL2. The L2NN sends one strobe for each element, reflecting the potentially different times required for each neural network to process and transmit a decision. The L2TT elements are sent simultaneously by the DSP and are therefore accompanied by only one strobe signal. The trigger elements are furthermore required to arrive within the pre-defined L2 receive time interval, which begins after L1K and ends at least two bunch crossings before the L2 decision is required, in order to allow sufficient time for the final calculation.

Management of the input signals is performed by a number of similar pASICs of which there are two for L2NN (named L2InpNN1/2), one for L2TT (L2InpTT) and one for the extra inputs (L2InpXtra). These offer the same basic functionality, differing only in details of the handling of the input signals. The input data are latched by these devices when gated by a valid strobe occurring within the L2 receive time. The latched output data are only cleared once L1K becomes false after restart. In the absence of a validated signal, the output is set to a default state specified individually in a register. By setting another register flag, the output can also be forced to the default. These features allow absent or jammed elements to be handled in an appropriate manner. The input pASIC also indicates if more than one valid strobe occurred; this does not affect the subsequent decision (which is based on the first valid strobed element received) but is useful for diagnostic purposes. A barrel shifter, controlled by two inputs driven by other parts of the CTL2 logic, allows the order of the latched data to be rotated by two or four bits.

The L2InpXtra pASIC has eight artificial outputs which are driven directly by the status of internal registers. These outputs allow L1 subtriggers not validated by L2 to cause an L2K decision (see below).

The latched data from the input pASICs are fed into lookup tables (LUT) which allow further logical combinations of the trigger elements to be made. The lookup tables are realised with six SRAM devices each having 13 address and eight data lines. Eight of the address inputs on each SRAM are driven by one byte of the latched data. The remaining five address inputs are connected to a variety of signals including the artificial outputs of the L2InpXtra pASIC and lines driven by later stages of the

L2 logic, enabling significant flexibility in the operation of the lookup tables.

The 48 outputs of the lookup tables form raw L2 subtriggers. Before an L2 decision can be made, these must validate the appropriate L1 subtriggers which is achieved by matrix multiplication. A vector of validation bits, V_i, is calculated using the expression

$$V_i = M_{ij} \cdot S_j \tag{1}$$

where S_j is the vector of L1 subtriggers for the event and M_{ij} is a matrix describing which L1 conditions are to be validated by which L2 subtrigger. The final validation is performed by multiplication of V_i with the vector of raw L2 subtriggers, A_i. Note that, since these are Boolean operations, multiplication and addition are replaced by logical AND and OR respectively.

The matrix multiplication is the most time consuming part of the L2 decision calculation and is therefore performed in parallel with the computation of the L2 trigger elements by a further pASIC, MatMul. The L1 subtriggers are available from the CTL1 logic after L1K and, with the matrix data M_{ij}, are retrieved from local RAM by the pASIC. In order to complete the calculation in a reasonable time, MatMul is clocked at three times the HCLK frequency and implements pipelining of the matrix operations, leading to a total required time of 18.8 μs. The output is the validation vector V_i, which contains 64 bits.

The manipulation of the L2 subtriggers to generate a final L2 decision is performed by the four L2SclOr pASICs. These each receive twelve L2 subtriggers from the lookup tables and sixteen bits of the validation vector from MatMul. The four excess bits of the vector V_i, which is available before the end of the L2 receive time, drive address lines of the lookup tables and can also be used to control the input pASICs. This allows L1 subtriggers to modify the generation of L2 subtriggers.

The L2SclOr pASICs validate the raw subtriggers, A_i, with the transmitted result of the matrix multiplication by making the logical AND of the vectors A_i and V_i. The resulting positive decisions are prescaled in a similar way to the L1 subtriggers by an 8-bit count-down scaler. The reject decisions are also fed to a scaler, allowing a controllable fraction of rejects to be overridden. This enables monitoring of the effect of L2 on L1 triggers. The validated, prescaled, overridden subtriggers are then ORed to form the final L2 condition; in terms of the vector-arithmetic used to describe the L2 operation, the L2K decision is

$$\mathrm{L2K} = \sum_i P_i \left(A_i V_i \right) \tag{2}$$

where P_i is a function describing the action of the prescale and override logic on the subtriggers. The validated, prescaled (but *not* overridden) L2 subtriggers are stored in registers on the L2SclOr pASICs. After each L2K, they are passed via a serial bus to the counting engine (CntEng) device, which implements 32-bit scalers for the 48 subtriggers. The scaled values are stored and can be used to determine, for example, the rate of each subtrigger.

Not all L1 subtriggers require an L2 condition. It is, however, still necessary to generate an L2K for these. For this purpose, one of the artificial outputs of L2InpXtra is used. This is always set, and mapped through the lookup tables onto an L2 subtrigger. The matrix memory load is defined such that all L1 subtriggers not assigned to a real L2 subtrigger are associated with, and validated by, this artificial signal.

Configuration of the CTL2 logic is performed under VME control. The matrix memory is directly accessible from the VMEbus. The loading of the lookup tables is performed with the assistance of the input and L2SclOr pASICs; the data outputs of the input pASICs directly define the lookup table input addresses, L2InpNN sets the write enable of the LUT and the data inputs of the L2SclOr are reversed to load data into the specified addresses. The L2 (and L1) setup is defined using a human-readable syntax which is parsed and compiled into the format necessary for loading into the logic.

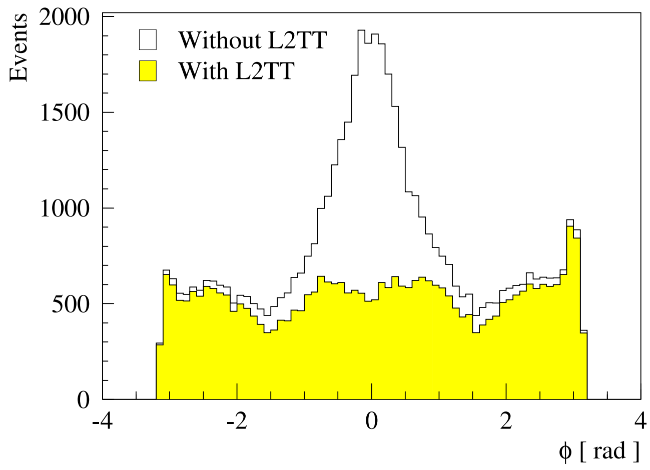

Fig. 3 ϕ distribution of the most energetic cluster in the backward calorimeter (SPACAL), without (open histogram) and with (shaded histogram) the L2TT requirement (see text).

VI. Performance of the L2 Trigger

The level 2 trigger has been in operation since the end of 1995, with the L2TT and L2NN systems being fully commissioned and integrated into the experiment during the 1996 running.

A number of the L2TT machines have been configured to recognise background topologies, providing trigger elements with significant rejection capability in addition to those designed to enhance the efficiency for specific physics channels. A lucid example is the control of an effect in the SPACAL backward calorimeter, in which localised groups of cells close to the beam pipe suffer from a significantly higher rate than is typical. This increased activity, whose source is believed to be the electron beam, has a timing consistent with physics interactions, and so generates a substantial fraction of the L1 trigger rate arising from SPACAL. The L2TT system is able to recognise this activity. This is demonstrated in Figure 3 which shows the ϕ distribution of clusters in SPACAL, with and without L2 rejection. The raw distribution is strongly peaked around $\phi = 0$ while application of the L2TT

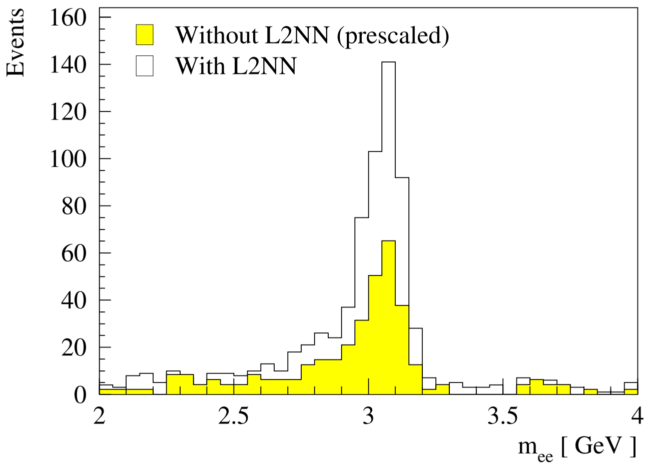

Fig. 4 Invariant mass spectra of e^+e^- pairs showing a peak from the decay of J/ψ mesons. The two histograms, normalised to the same integrated luminosity, show the physics yield before (shaded) and after (open) introduction of L2NN.

produces a significantly flatter distribution. This result directly demonstrates the ability of L2TT to identify specific topologies within the detector.

For the configuration of the L2NN, focus was placed on physics channels which could previously only be triggered efficiently with an intolerable L1 rate. In particular, the decay of heavy quark vector mesons to lepton pairs was targeted. Six pattern recognition modules have routinely been used to select events since the middle of the 1996 data-taking period. Rate reduction factors varying from 2 for some channels to 100 in others, with efficiencies in the range of 60 – 90%, have been achieved. Figure 4 shows an example of the benefits afforded by the use of L2NN in this manner. The channel of interest is the leptonic decay of J/ψ mesons, which is observed as a peak in the invariant mass spectrum of the decay products (electron-positron pairs). Before the introduction of L2, the L1 trigger used to select this channel was prescaled by a factor of two in order to bring the rate down to an acceptable level. After application of a L2NN trigger element, the rate was reduced by a factor of four, allowing the prescale to be removed. In spite of the net rate reduction, a clear increase in the yield of such events for physics analysis is observed.

The introduction of L2 has enabled the design capabilities of the H1 trigger system to be exploited more fully. Whereas the L1 trigger rate was previously constrained to not exceed the required input rate of L4, a significant reduction is now available at L2. Typical rates, which lead to acceptable deadtime values below 10%, are 400 Hz at L1 and 50 Hz at L2, with L4 logging approximately 5 Hz of physics and 5 Hz of monitor triggers to tape. The highly configurable CTL2 logic, with its extended monitoring capabilities, allows the L2 trigger to be operated with little risk of physics losses. Furthermore, its open architecture and flexibility, in combination with the two parallel, independent trigger systems allows operation to be readily improved as experience is gained and as the luminosity delivered by HERA steadily increases.

VII. Conclusion

Since the end of 1995, the L2 trigger system of H1 has been used extensively and proven to be highly effective in isolating various physics processes from background. The flexibility of the chosen implementation, including the use of parallel, independent trigger processors and configurable final decision logic, has enabled rapid and easy response to changing demands on the trigger system imposed by varying beam conditions and physics requirements.

The HERA collider, with a high bunch crossing rate and large beam backgrounds, presents a significant challenge to the experiments and is an ideal testing ground for the environment anticipated at future collider facilities. The concept of an open architecture, multi-level trigger such as that employed by H1 has proven to be very successful and the experience gained will certainly be of benefit to the next generation of high rate experiments.

VIII. Acknowledgment

The authors are grateful for the assistance of J.-C. Bizot, F. Gaede, P. Merkel and A. Wegner in the preparation of this paper.

IX. References

[1] I. Abt et al., "The H1 detector at HERA", *Nucl. Instr. and Meth.*, vol. A386, pp. 310-396, 1997.

[2] S. Eichenberger et al., "A fast pipelined trigger for the H1 experiment based on multiwire proportional chamber signals", *Nucl. Instr. and Meth.*, vol. A323, pp. 532-536, 1992.

[3] E. Eisenhandler et al., "The H1 SPACAL time-to-digital converter system", *IEEE Transactions on Nuclear Science*, vol. 42, pp. 688-692, 1995.

[4] F. Sefkow, E. Elsen, H. Krehbiel, U. Straumann and J. Coughlan, "Experience with the first level trigger of H1", *IEEE Transactions on Nuclear Science*, vol. 42 pp. 900-904, 1995.

[5] W. Haynes, "Experiences at HERA with the H1 data acquisition system", *Proceedings of the Int. Conference on Computing in High Energy Physics 1992, Annecy, France*, pp. 151-161, 1992.

[6] A. Campbell, "A RISC multiprocessor event trigger for the data acquisition system of the H1 experiment at HERA", *IEEE Transactions on Nuclear Science*, vol. 39, pp. 255-258, 1992.

[7] The VMEbus specification, ANSI/IEEE standard 1014-1987, VITA, 10229 N. Scottsdale Rd., Suite E, Scottsdale, AZ 85253.

[8] C. Beigbeder and D. Breton, "The H1 PQZP system", *H1 Internal Note H1-02/93-269*, unpublished, February 1993.

[9] J. Köhne et al., "Realisation of a second level neural network trigger for the H1 experiment at HERA", *Proceedings of AIHENP 96, Lausanne*, World Scientific, 1996.

[10] Adaptive Solutions, Inc., 1400 N.W. Compton Drive, Suite 340, Beaverton, Oregon 97006-1992.

[11] J.C. Bizot, G. Nowak and J. Turnau, “Hardware study for a topological level 2 trigger”, *H1 Internal Note H1-09/92-240*, unpublished, 1992.
J.C. Bizot, G. Nowak and J. Turnau, “Status of simulation for a topological level 2 trigger”, *H1 Internal Note H1-02/92-212*, unpublished, 1992.

[12] QuickLogic Corp., 1277 Orleans Drive, Sunnyvale, CA 94089-1138