

BUS-BASED ARCHITECTURES IN THE H1 DATA ACQUISITION SYSTEM

W. J. HAYNES
2 August 1992

Presented to the International Conference "Open Bus Systems '92"
Zürich, Switzerland, 13 - 15 October, 1992

BUS-BASED ARCHITECTURES IN THE H1 DATA ACQUISITION SYSTEM

W. J. HAYNES

DESY Deutsches Elektronen-Synchrotron, Hamburg, FRG
SERC Rutherford Appleton Laboratory, Oxfordshire, UK

ABSTRACT

The recently commissioned HERA electron-proton collider at DESY has presented a challenging environment in coping with huge volumes of data in real-time. With bunch-crossing intervals of 96 nanoseconds, the H1 experiment can generate nearly 3 Megabytes of raw digital information from several hundred thousand electronic channels. The open VMEbus standard is exploited to provide the necessary modularity and system flexibility for embedding fast readout electronics in a multiprocessor design array. The architectural concepts of the data acquisition structure are discussed together with many of the experiences gained in implementing a large multi-crate system.

INTRODUCTION

The Hadron-Elektron Ring Anlage (HERA) is a large particle physics proton-electron storage complex buried beneath the suburbs of Hamburg. Nearly 4,000 magnets are used to collide 30 GeV electrons with 820 GeV protons in two separated rings 7 km in circumference. The main dipole and quadrupole magnets of the proton ring are all superconducting with the 422 dipoles having a field strength of 4.68 Tesla. Tunnelling commenced in 1985 and culminated with the first electron-proton collisions in October 1991.

H1 is one of two large particle physics experiments currently exploiting this unique facility, involving over 300 scientists from more than 30 institutes worldwide [1]. The detector, itself, consists of drift and proportional chambers (charged particle "trackers") surrounded by a liquid argon electromagnetic and hadronic calorimeter of lead and iron absorber plates, outlined in Figure 1. This inner region is enclosed by a 6 m diameter, 1.2 Tesla superconducting coil. The coil provides a large field volume for the calorimeter, with a minimal amount of inert material in front, ensuring a strong, homogeneous, solenoidal field for the tracking chambers. Surrounding it are a set of instrumented iron plates providing both the return yoke of the magnet and a muon filter/tracker system together with up to three layers of external muon chambers. Various smaller subdetector elements then complement the backward and forward areas to provide near hermetic coverage of the interaction point. The completed detector, weighing nearly 3,000 tons, was commissioned during cosmic-ray tests in Spring 1991 followed by full operation in HERA since May 1992 at centre-of-mass energies of 314 GeV.

From the real-time computational point of view, a total of over a quarter of a million analogue channels are read-out and digitised. This results in nearly 3 Mbytes of raw digitised information for a "triggered event". Since the time between successive electron-proton bunch crossings is just 96

nanoseconds, various levels of hardware triggering, software filtering and digital compression are employed before reducing final data sizes to acceptable storage-media recording rates. To accomplish this, several hundred processing elements are embedded largely within the IEEE VMEbus standard [2], a choice encompassing the advantages of open specification for design and the commercial availability of advanced competitive products. Several descriptions of the system during staging have already been presented elsewhere [3, 4].

SYSTEM OVERVIEW

The HERA collider poses an interesting technical challenge for physics extraction in real-time. Each of up to 210 proton bunches can contain 10^{11} particles; simultaneously a maximum number of 210 bunches can each contain 3.5×10^{10} electrons. The two sets of beams can cross at a rate of 10.41 MHz yielding a bunch crossing every 96.064 ns and a possible e-p collision. At this design luminosity of 1.5×10^{31} $\text{cm}^{-2} \text{s}^{-1}$ the dominant physics is photoproduction at 120 Hz, with the interesting structural processes expected to add just a fraction more to that. It would therefore appear at first sight that the problem of trigger selection is trivial. However, to avoid betatron oscillations and "beam blow up", HERA has been designed so that the electrons and protons collide "head-on" with a zero crossing angle. This necessitates electron bending magnets close to the collision point, or "interaction region", which results in huge fluxes of synchrotron radiation. Much of this is reduced by means of absorbers and collimators, but again close to the interaction point. Together with proton beam halo and beam-gas, all this unfortunately results in a huge background rate of some 10^4 Hz around the interaction region.

Triggering and readout stages

Figure 2 summarises the overall data acquisition system and the key rates at full HERA design luminosity. Information is read-out in parallel from many subdetector partitions before being finally merged. Because of the high background rate any triggering and filtering of data has to be highly sensitive and selective, since the 270,000 analogue channels generate nearly 3 Mbytes of equivalent raw digitised information during the time of interaction, or "event". In total, four main levels of hardware triggering and software filtering can be enabled on all or part of the various detector partitions, each level introducing a higher degree of sophistication in event selection. In parallel data compression and formatting further reduce event sizes to between 50 Kbytes and 100 Kbytes so that final data recording rates are restricted, at present, to a maximum 1.2 Mbytes/s.

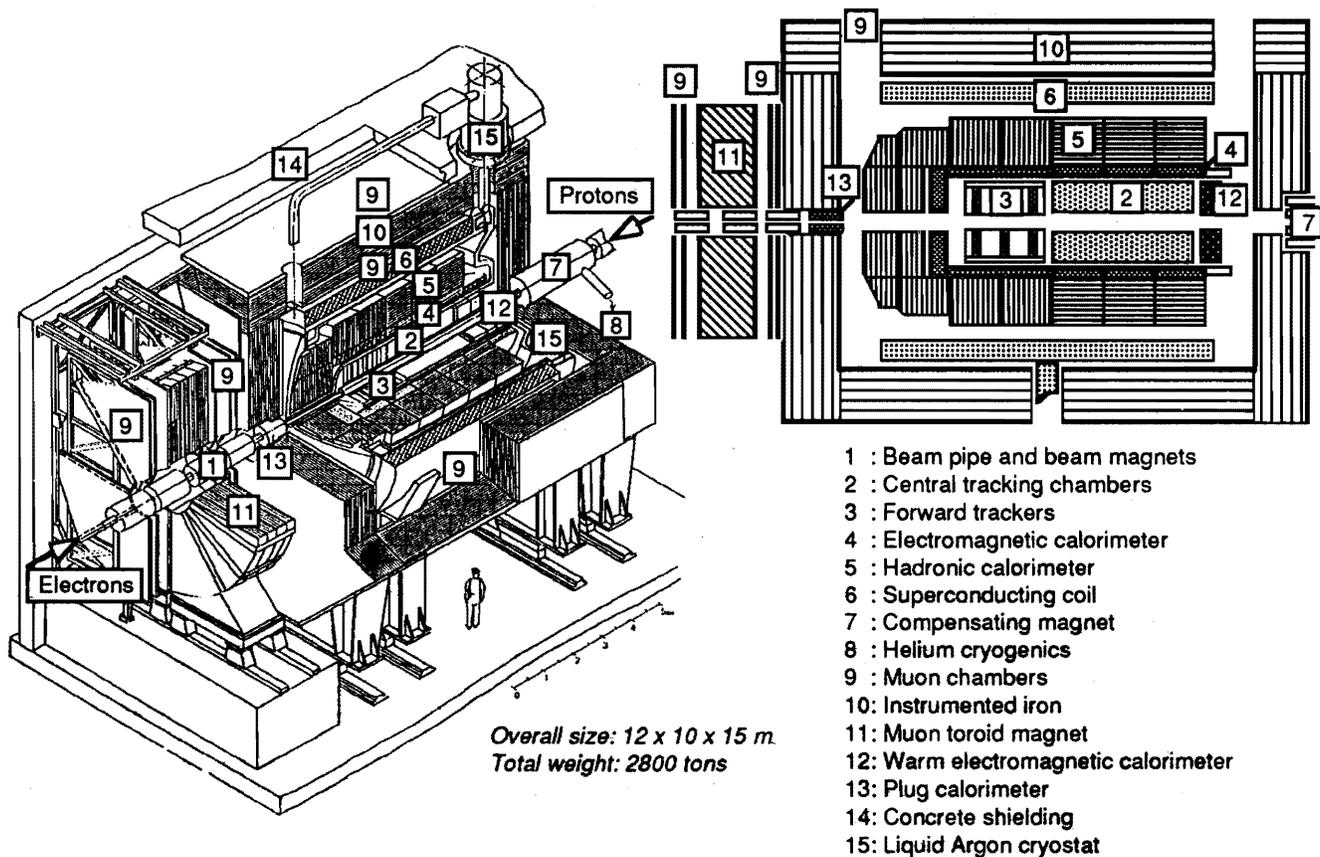


Figure 1 : The HERA detector H1

A purpose-built "pipelined" triggering system (22 bunch-crossings long, dead-time free) is used to select initial candidates for data processing from the background of 10^4 events/s [5,6]. This "Level-1" comprises many triggering elements from different parts of the detector, such as FADC energy sums from the calorimeter, fast vertex-position reconstruction from the proportional chambers using XILINX logic cells and veto-wall and time-of-flight subsystems against off-momentum beam particles. A particularly interesting level-1 element is the fast track finder which again capitalises on XILINX cells, this time to find and count charged tracks in the r - ϕ projection from a sample of the central drift chamber wires [7].

After a Level-1 accept the front-end pipelines are held. Dead-time then begins. Currently, at luminosities nearly 3 orders of magnitude below design, the first-level rate varies up to 20 Hz depending on the elements selected. With this in mind there is the option of making a more-refined hardwired Level-2 triggering decision based on combined information within a further $20 \mu\text{s}$ to further reduce rates. Solutions based on neural networks and ASICs are presently being prepared in readiness for full HERA luminosity. Front-end electronics readout is not initiated until an early-level triggering decision is made whereupon each component has a maximum digitisation time of $800 \mu\text{s}$ before the pipelines are re-enabled. This governs a significant fraction of the dead-time of the complete system; with a hadron-hadron collider such a predicament would not be

tolerable. A large amount of the data compression also takes place during this phase. However a third level of filtering can be enabled based on calorimeter information alone. Therefore, should readout commence, the pipelines are immediately reactivated if a Level-3 reject occurs.

Logical structure and data control

In logical structure H1 merges its subdetector components, in parallel, into individual subsystem VMEbus crates which each contain a readout controller and memory buffer plus a fibre optic link to a coordinating event management task. After the digitisation, compression and formatting of front-end data, full-event records are built over the fibre optic links. A parallel array of RISC processors provides the fourth level of software-coded filtering once all of the full-event information is available [8].

At all stages workstation intervention provides for data monitoring and detector control. To avoid saturation on global VMEbusses, extensive use is made of the local VSB specification [9]. To minimise any further dead-time contributions, memory buffering is provided at the key data processing stages.

Finally, in addition to the main data flow, a Local Area Network (Ethernet) caters for the relatively slow exchange of general operating conditions and files between different subsystems. This also serves a wider international community

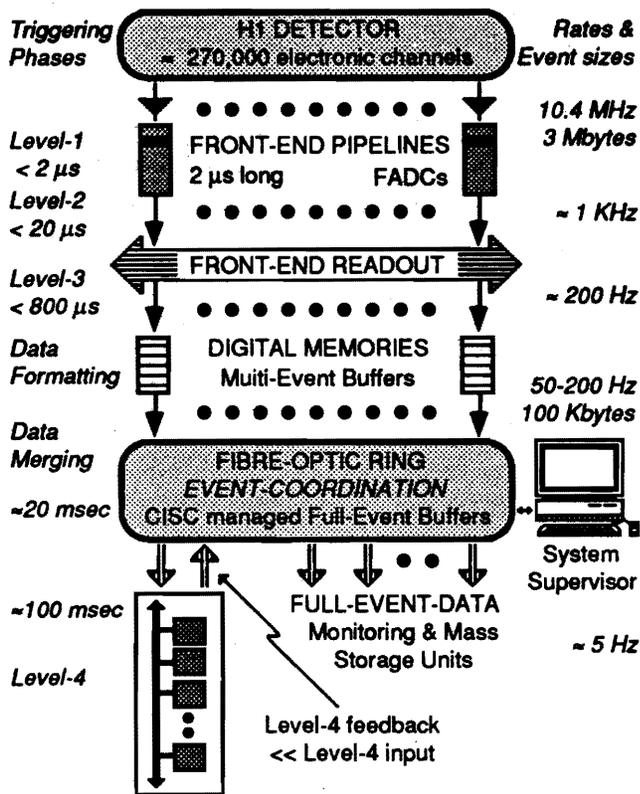


Figure 2 : Overview of the H1 Data Acquisition System

with current status and event information. As a consequence the complete system can be monitored and expertly controlled from any external site with a protected amount of jurisdiction.

FRONT-END DIGITISATION

The digitisation of the 270,000 front-end signals includes a vast array of digital signal processors and custom readout electronics.

Calorimetry

The 45000 channels of the liquid argon calorimeter consist of readout pads sandwiched around absorber plates. These are multiplexed onto 55000 inputs into 14-bit ADCs (25% of the channels have double gain to increase the dynamic range) and read out by over 50 Digital Signal Processors (1024 channels/56001 DSP) in VMEbus [10]. Each DSP performs zero suppression, pedestal subtraction, gain correction and calibration sums [11]. The digitisations are buffered in stages, both before and after DSP processing to minimise dead-times, before being read out over VMEbus (Figure 3). The whole system can be managed by either CISC or RISC processors. In parallel, the calorimeter trigger information (from analogue channel sums) is read out with 8-bit FADCs and 56001 DSPs via a separate VMEbus branch. In each branch DSP encoding takes place in 800 μ s. Single-width VICs are used to interconnect blocks of 15 crates with nominal transfer rates of 8 Mbytes/s [12].

The instrumented iron is also read-out through the calorimeter system.

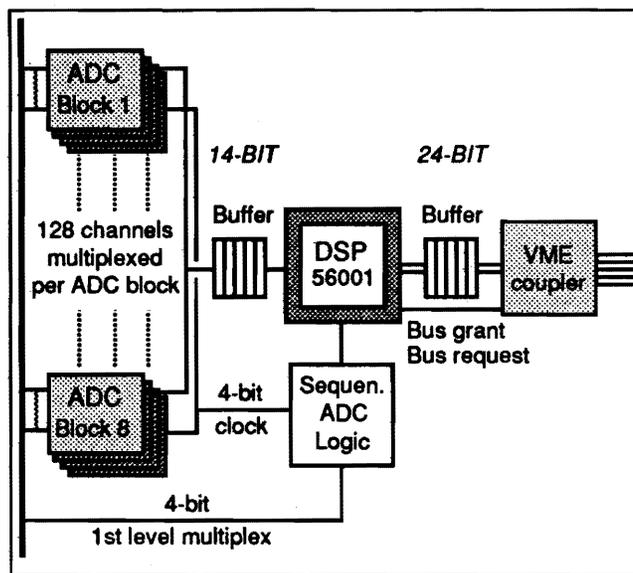


Figure 3 : Schematic of the H1 calorimeter ADC readout

Charged particle tracking and muon readout

A purpose-built solution is used for digitising and compressing the analogue information from the 10,000 sense wires of the inner drift and forward-muon chambers [13]. Each digitising channel consists of a 100 MHz, 8-bit, nonlinear Flash ADC (Figure 4). A total of 45 triple-height Euro crates (based around VMEbus) each enable 16 channels to be read-out per board with 16 boards in a crate. There are 256, 8-bit words per FADC channel, giving a total of 64 Kbytes of raw data information in each crate. A fast scanner module reads out all 256 channels, within a crate, into memory at 80 Mbytes/s within the maximum allowed readout time of 800 μ s [14]. The leftmost four slots of each crate have been designed to accommodate standard double-height VMEbus cards with VSB. This enables each crate to be equipped with commercial 68030-based processor boards (to be upgraded to 68040-based cards) in order to compress the total of over 2 Mbytes of tracker information into a few tens of kilobytes for each triggered event [15]. The data from the separate crates are merged initially over VIC-bus [12]. The luminosity counters, which confirmed the first collisions from HERA in 1991, also use the same FADC system [16].

The wires of the muon system contribute a large number of channels sparsely filled with data, so an encoding system is used [17]. It is contained within a total of 6 standard VMEbus crates, with a division of 64 modules (within each module there are 16 chamber planes with a maximum of 512 channels per plane) each equipped with its own readout controller. These controllers work in parallel, independent of each other. Each wire is connected to a comparator whose output is fed into shift register pipelines (depth of 32 stages). On trigger accept these registers are dumped into memories within their assigned readout controllers at frequencies between 1 and 10 MHz, where the data are zero suppressed and encoded.

The multi-wire proportional chambers (MWPCs) have around 4000 pads. All front-end electronics are based on a user-defined

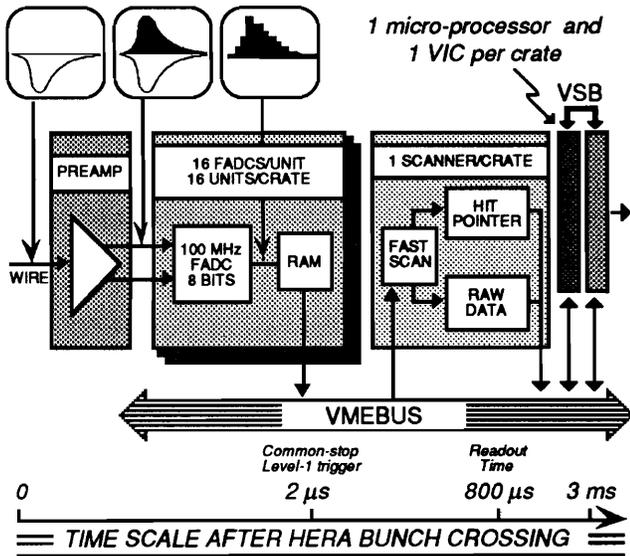


Figure 4 : FADC readout system for the H1 trackers

backplane split into 4 subbranches where 280 receiver cards are distributed between 16 crates [18]. Each receiver card consists of a discriminator which feeds a pipeline, keeping synchronisation with the HERA clock. Controller cards interface the front-end crates to VMEbus via branch drivers each equipped with a DMA controller and on-board FIFO.

Silicon trackers

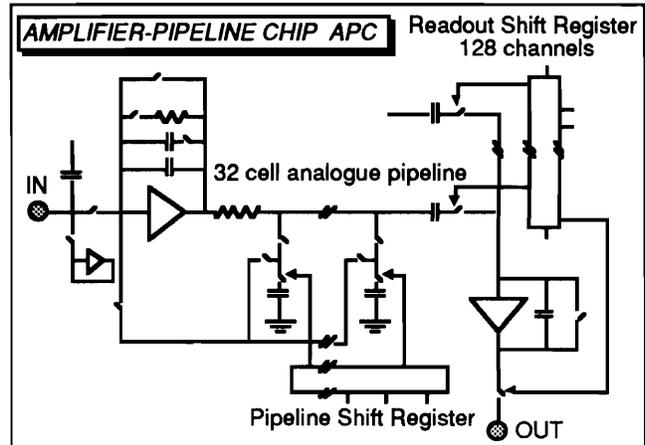
In order to further improve the physics potential of the H1 experiment a proposal is now under way to implement silicon strip detectors by the end of 1994 [19]. A total of 214,000 electronic channels are to be read out by purpose-built front-end amplifier chips which contain switched capacitor analogue-pipeline buffers (Figure 5). A single chip has 128 low-noise channels with a pitch of $44\mu\text{m}$ per channel storing 32 consecutive bunch crossings. Pipeline control will be synchronised to the 96 ns bunch crossing period. Each chip is fabricated in a $2\mu\text{m}$ CMOS process and prototypes have already been tested with a readout speed of 2.5 MHz, thus allowing a 2000:1 multiplexing within the 800 μs readout time. The total power consumption of a single chip, including the pipeline running at 10 MHz, is 38mW. An external readout unit, with a data reduction processor for hit-finding and zero-suppression, is now under the early stages of development with a VMEbus interface.

SYSTEM ESSENTIALS

Altogether there exist nearly 200 electronics crates, the majority VMEbus, spread around the H1 online complex with a multitude of processing tasks executing in parallel. In order to provide a coherently managed system, from many different sources, a certain set of standards need to be adhered to.

Basic hardware components

On H1 an early decision was made to "standardise" on certain modules to ease maintenance and software development



Each VMEbus card can read out & control 4 x 2048 channels

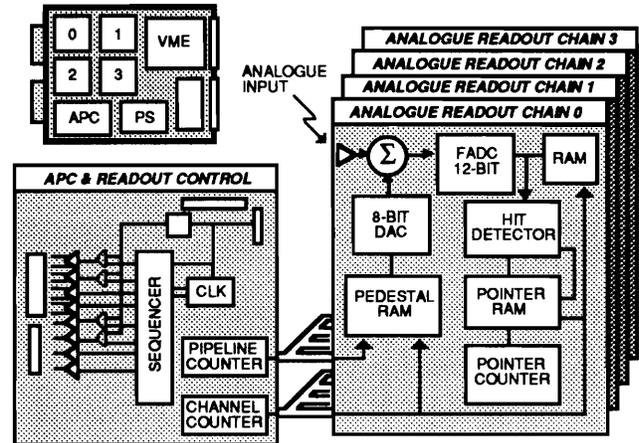


Figure 5 : Silicon tracker readout electronics

overheads. For general-purpose real-time control within VMEbus, a common 68020/30/40 processor card was selected [20]. On-board memory is *dual-ported* so that it can be accessed externally from the VMEbus (for communication) as well as internally. VSB local bus access can be extended to memory boards in adjacent crates but is *never* used as an inter-crate connection.

For even more calculation-intensive applications, such as final-level filtering, RISC processor-based modules are employed. The RAID 8235 is a single-width VMEbus/VSB board based on the 25 MHz MIPS R3000 processor and the R3010 Floating-Point Accelerator with up to 32 Mbytes of onboard DRAM [21]. To obtain the necessary memory bandwidths and reduce instruction access times, 128 Kbytes of independent data and instruction cache ensure that each board has an equivalent computing power of around 50% of an IBM 3090 mainframe.

The DPM 8242 is the basic memory module and can be equipped with up to 16 Mbytes of 70 ns static ram, having equal priority in arbitration from the VME and VSB ports [22]. A broadcast mode selection allows any VMEbus master to write to several memories simultaneously. Both the VMEbus and VSB memory base addresses are *software programmable* via VME accessible registers, permitting a

managing Event Coordinator task (see later) to switch units between different events.

No system would be complete without the provision for software development and operator control. On H1, Macintosh computers are used extensively for this purpose, due to their graphics-user-interface and open NuBus architecture. They are mapped onto VMEbus via MacVEE and Micron, MacVEE Interface Card Resident On NuBus [23]. The consequence for programmers is that they can have memory mapped access directly into a maximum of 24 VMEbus crates from one NuBus-based Macintosh. No sophisticated protocol needs to be initialised, as is the case for many other DMA-based interconnections.

In addition to these standard boards, there exist the usual debugging instruments and analysers necessary for successful system integration.

VMETaxi

In many of the front-end subsystems the standard VICbus is used to interconnect crates. However for coordinating the fast data acquisition transmission protocol, between the different readout subsystems, VMETaxi modules are employed [24]. These can interconnect VMEbus crates up to several kilometres with fibre-optics. Figure 6 illustrates the basic philosophy of the VMETaxi. All boards are interconnected with multimode optic fibres to form a ring, using AMD 7968/7969 TAXI-chip transmitter/receiver pairs. Because of this, there is theoretically no limit on the number of such devices which can be interconnect within a single ring; in practice one is limited by the software overhead in setting up transfers. Each single-width board has a cpu controller and single level VMEbus arbiter. Thus the on-board cpu has a dual functionality; combining processing ability with DMA control on one chip. The software protocol, discussed further in the following sections, is purpose-written and optimised for speed and efficiency in a data acquisition environment; no FDDI or similar networking protocol is adopted. In setting up transfers between any two crates, an initial description packet is sent around the ring; those modules not engaged in the subsequent activity can enable by-pass registers so as to establish a direct connection between the two VMETaxis that are involved, analogous with SCI [25].

Currently a 25 MHz 68020 based board is used with 125 MHz taxi chips; later this year upgraded "Mark-2" modules will be able to exploit 33-50 MHz 68030 processors and 250 MHz taxi chips. With the current 020 boards, the system is capable of sustained 32-bit DMA block transfers of 12.5 Mbytes/s under software control independent of distance; using double-links, on the upgraded 030 cards, up to 50 Mbytes/s will be achievable. The *link* reliability has been tested to a bit error rate of less than 1 in 10^{13} . Program memory is provided for by 128 Kbytes of on-board dual-ported static ram (expandable to 2 Mbytes on the upgraded system), with external memory access through either the VMEbus or VSB ports. On-board EPROM and EEPROM provide for firmware storage and configuration parameters. VME64 and VSB block transfer modes will be provided for on the Mark-2 modules, realised with XILINX gate arrays presently under evaluation with 175 MHz taxi chips.

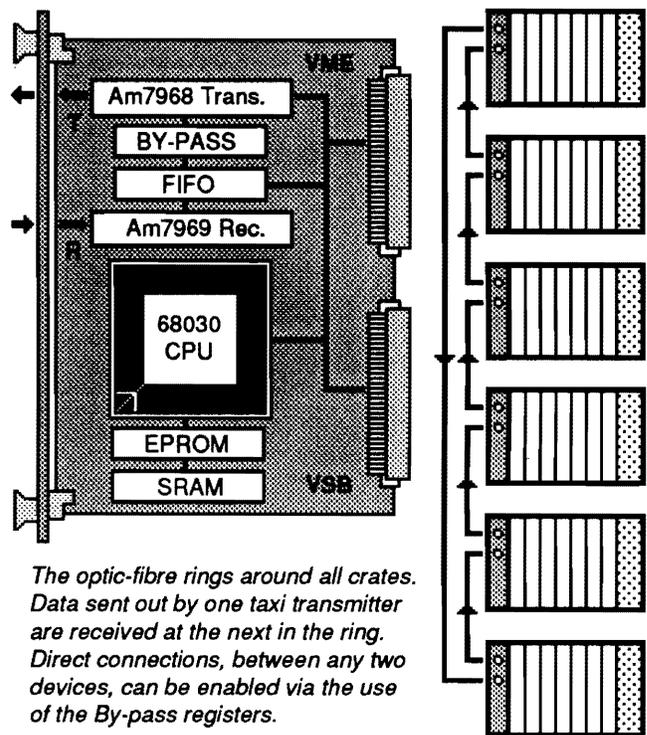


Figure 6 : VMETaxi essentials

Software

There is a well-ordered structure defined by the architecture of the hardware, discussed in the following section, which defines how the software is written and developed. Dedicated tasks run on dedicated processors throughout the VMEbus system with operator intervention provided for via graphics-orientated computers such as Macintosh. Different languages and approaches are used depending on the application being developed. Much of the online data acquisition systems code which executes on 68000 series microprocessors, within the VMEbus, is written in either C or Assembler. Final-level filtering algorithms, which execute on the R3000 boards, have large code resources originating offline and written in Fortran [8]. Control programs which run on the Macintosh computers exploit the latest object-orientated and graphics-based packages currently on the market [26]. It can be seen therefore that the Macintosh provides a convenient integral component in both software development and operator-control. In order to encompass the framework of VMEbus, extra tools have been developed to ease the integration into the VMEbus [27]. As a result, no formal operating system is required to run on the VMEbus boards; all basic functions can be controlled from a Macintosh via external directives through dual-ported mailbox memories.

In order to enhance information exchange between control computers a dedicated package has been written based on international networking protocols [28]; this has the added advantage of enabling the complete system to be monitored from external laboratories and institutes.

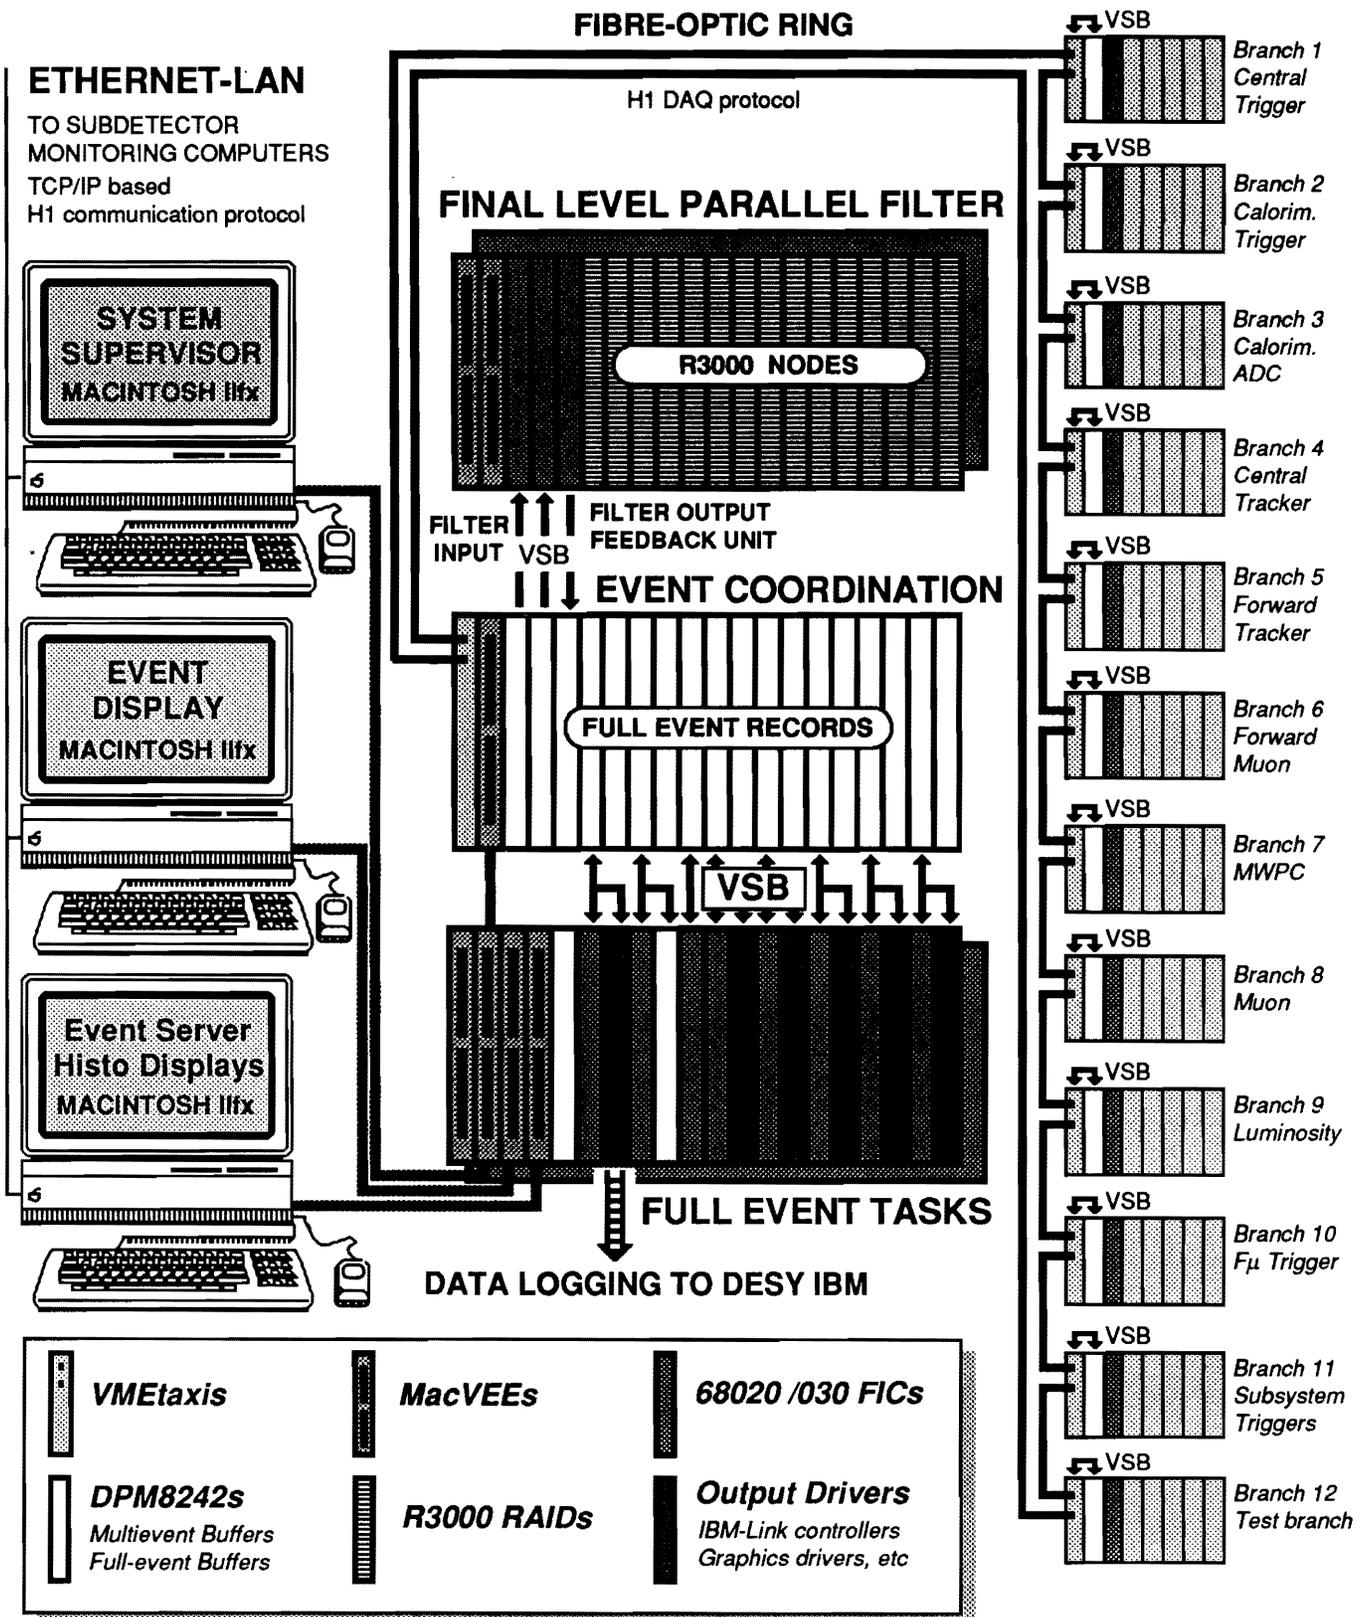


Figure 7 : Physical layout of the central part of the H1 Data Acquisition System

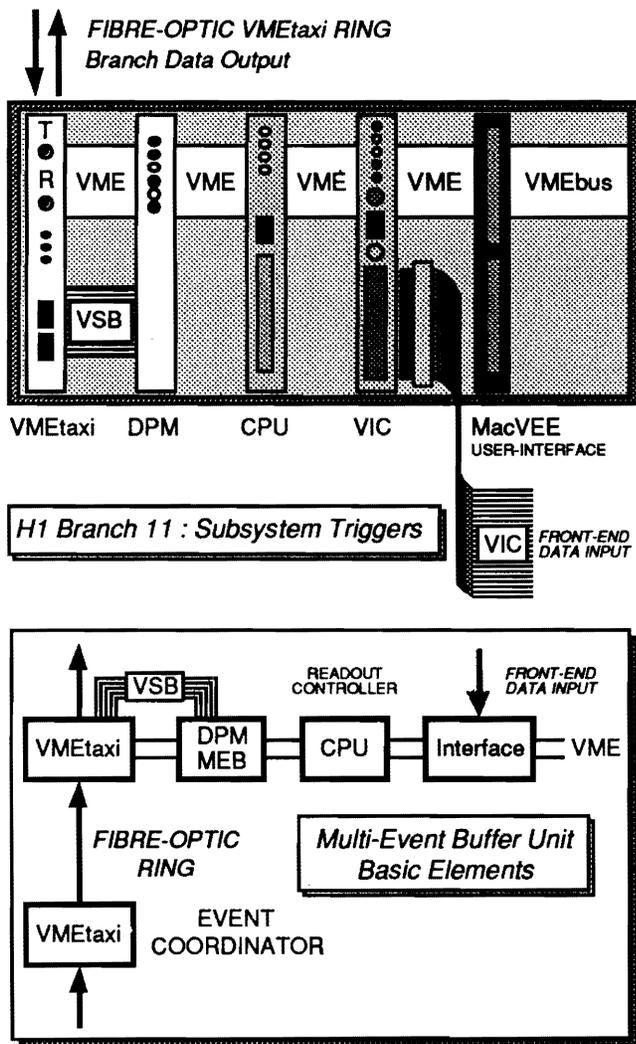


Figure 8 : Multi-event buffer unit and example

DATA ACQUISITION ARCHITECTURE

By convention the data acquisition system breaks down into three main categories for the purpose of central coordination and management. First there are the front-end "producers" of data; the end result of extensive electronics digitisation from many subcomponent parts. Second, the data must be merged and distributed to several full-event "consumers"; subsystems which monitor and record the data on permanent storage media. Third, the system has to be initiated and controlled via external processes for operator intervention.

Central management and coordination

Figure 7 shows the physical layout of the complete system, managed centrally through several layers of software protocol purpose-written around the VMEtaxi fibre-optic ring [29]. Data are read into the ring from each subdetector division (called "branch") via local bus VSB subsystems. They are then merged, to form full-event records, over the ring as with a conventional vertical interconnect. Finally they are broadcast

over a VMEbus to several independent units. At this stage further data processing, monitoring and filtering can be done prior to final data recording. Common, shared, memory blocks provide for communication between all processors involved and allow external computer stations to supervise the complete system and cater for operator intervention. The master of the ring executes the "Event Coordinator" task, controlling the whole sequence of management processes. A separate object module provides a software library to interface individual subsystem elements at different levels, with features ranging from full module testing and control through to basic buffer management and system coordination.

The architecture of the front-end producers

Each subdetector component readout system is autonomous up to and including a central subsystem VMEbus crate. Each central subsystem crate contains a dedicated supervisory readout controller, a multi-event buffer (MEB), a VMEtaxi and any input drivers necessary to access the data (Figure 8). Consequently the design allows a particular component to be decoupled from the rest of the system, during installation and test phases, and does not exclude the use of other busses for front-end digitisation. Data are placed into the multi-event buffer over VMEbus and then extracted, by the VMEtaxi, through VSB. The central coordinating event management task, the "Event Coordinator", runs on the master VMEtaxi of the ring and interacts with each subsystem readout controller via shared memory blocks. Before commencing data acquisition, the Event Coordinator is responsible for hardware recognition and initialisation of the various subsystem components. During acquisition, when it is free for event building, it searches the subdetector multi-event buffers for the next event. When all components are ready, with the same event number, then all of the separate banks are transferred via the optical ring into a full-event buffering system (next section). On completion, each corresponding multi-event buffer is released.

A subsystem crate *may* also contain additional processors for data reduction, an interface to a subdetector monitoring computer and a subsystem trigger controller (STC). All STCs are interfaced to a dedicated central trigger controller which, in turn, coordinates the sequence of all hardware triggering levels together with obtaining information from the HERA machine [30]. Each STC provides signal ports for the component electronics and communicates with the readout through VMEbus read-write cycles and interrupts. A consequence of this is that the Event Coordinator needs no hardwired connection with the central trigger controller itself; all management sequences are handled by software, so providing a portable solution to any large multi-crate VMEbus system.

At this subsystem level, the software library [29] provides multi-event-unit routines catering for initialisation, accessing status information, requesting buffers and signalling when event data are ready for merging. Readout error codes are normally indicated as standard parameters, but messages can be sent through the system in standard ASCII format to describe subsystem-specific anomalies. Figure 8 also shows how the protocol has been transparently mapped from a main branch through to a sub-branch structure by incorporating standard VIC links to merge small subsystem components.

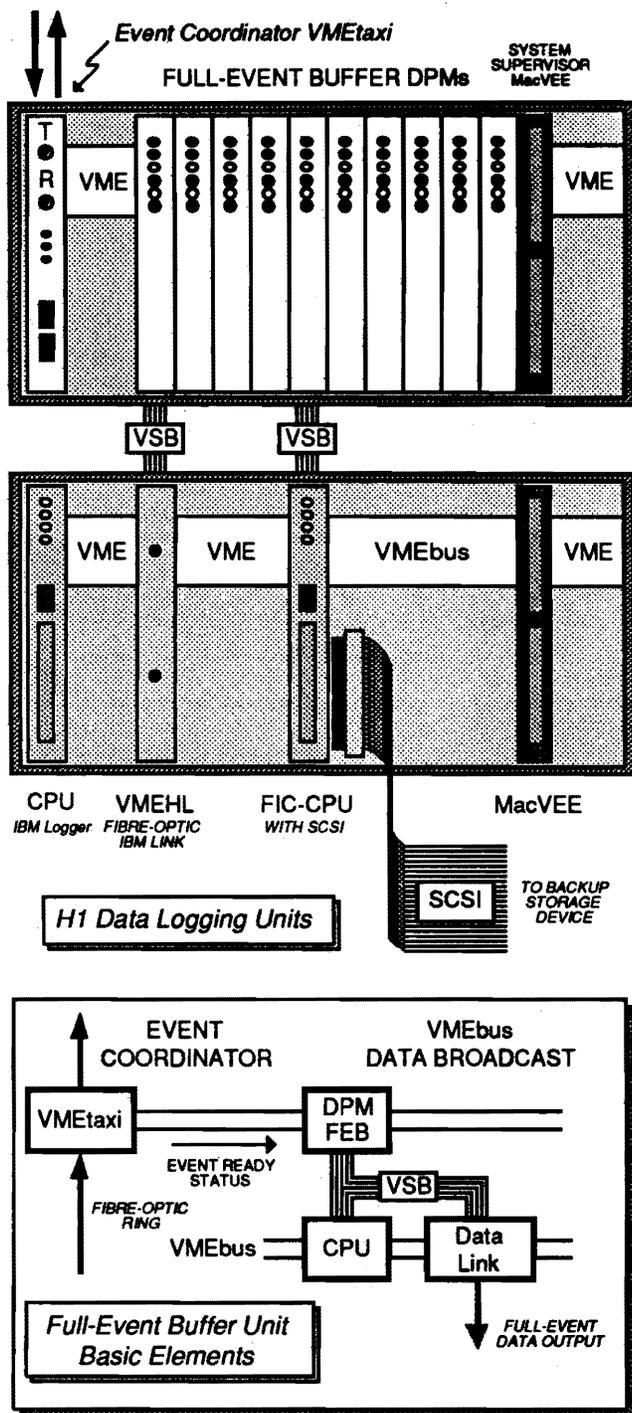


Figure 9 : Examples of full-event buffer units

The management of full-event consumers

As the master Event Coordinator VMEtaxi builds full-event records it simultaneously broadcasts them along its VMEbus to dual-ported memories associated with parallel sets of "full-event units". Event tasks are performed by connecting the unit

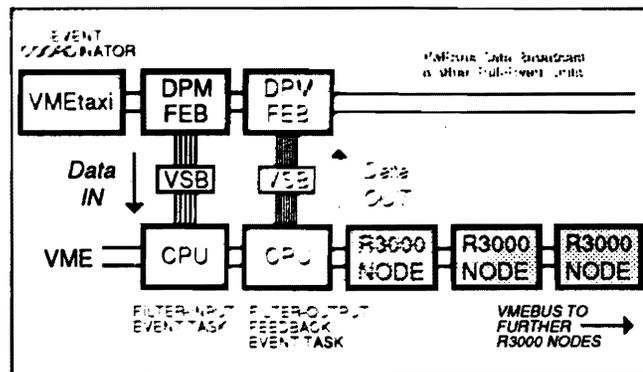


Figure 10 : Filter "farm" full-event buffer unit subsystem

processors with their respective full-event buffer (FEB) memories via VSB in order to minimise any bandwidth overheads (Figure 9). Since the memory has read/write access from both the event task processors and the Event Coordinator, full-event units need not only be "consumers" of data. An event unit can also become a "producer" of data, "feeding-back" information into the system. Consumers are able to determine which data they wish to receive, ie directly built events from the front-end subsystems or data fed-back into the system from, for example, a parallel filter farm. All this is handled modularly by the software protocol running on the VMEtaxi system and its associated library of routines.

Typical full-event tasks are data-logging, event display, data monitoring and histogramming. Final event records are sent to the central IBM facility at DESY, some 3 km distant, at rates of up to 7 Mbytes/s [31]; at present the disc writing rate limits the system to 1.2 Mbytes/s. The data logging task requests each event to be recorded and so there must always be a free buffer associated with it. In case of link failure a backup event task, that drives a storage device directly from VMEbus, can be enabled (Figure 9).

A more sophisticated form of unit is the fourth level filter "farm" consisting of many R3000 processor "nodes" working in parallel (Figure 10). Each node executes the same algorithm independently on different events in order to provide a final level of triggering and reconstruction. Event data are passed to the different nodes under control of a filter input controller, which communicates with the Event Coordinator as a normal consumer full-event task. When a node is free it signals this to the filter input task which searches for an event ready in its full-event buffer. The data are then extracted over the VSB and sent to the free node, releasing that particular buffer and enabling a further node to be serviced. When a node has finished processing an event it signals completion to a filter output controller (a "feedback" event task). The latter then decides whether to pass the event back into the central system managed by the Event Coordinator and free the node for further processing. As a full-event unit, the parallel filter can be easily enabled or disabled. Furthermore, the software protocol of the Event Coordinator allows for several banks of filters to be implemented, each with their own input and output controllers. The reason that the single Event Coordinator task can accommodate the handling of the filter-output is due to the relatively low acceptance rates at Level-4; however additional processors can be introduced to share this task if required.

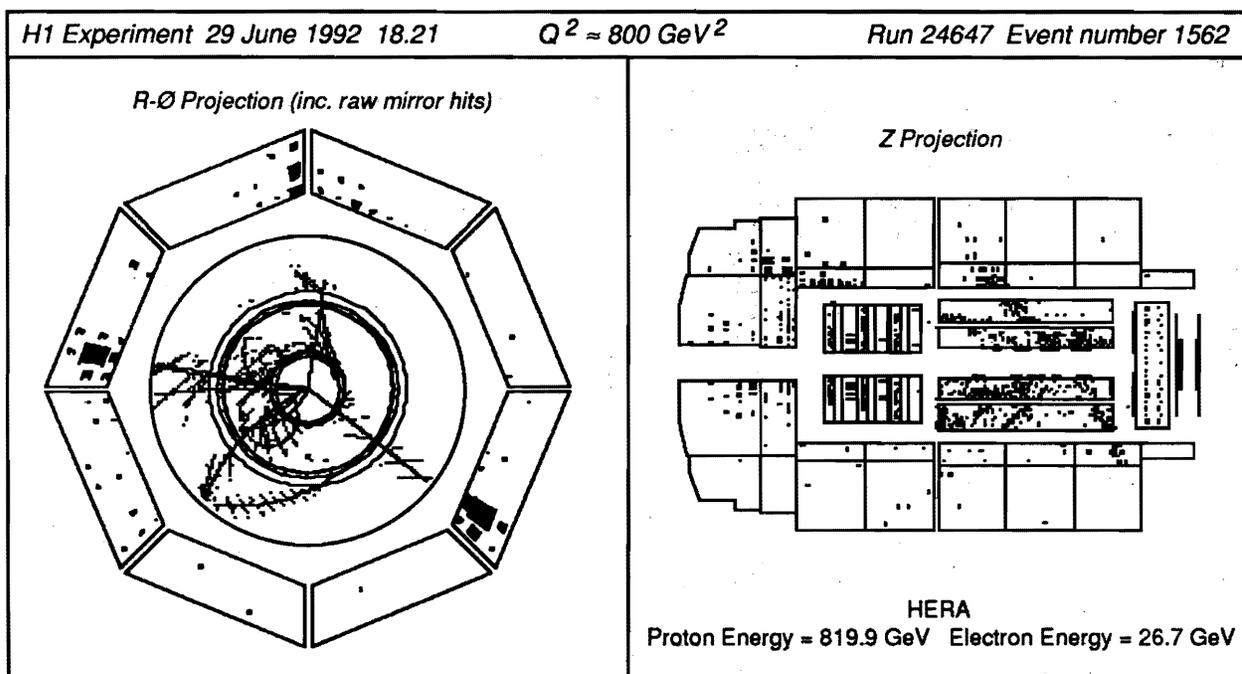


Figure 11 : One of the first e-p deep inelastic events seen from HERA and reconstructed in real-time at H1

System supervision and operator-control

The H1 system is capable of running entirely in VMEbus; dedicated processors run dedicated tasks so that the net result is one of a conventional multi-tasking system. This leads to a natural division of tasks and processes. The primary tasks of data acquisition are performed by the subsystem readout controllers, event builders, filter control tasks and event units, all managed by the Event Coordinator. All fast processing and readout is done within VMEbus. An exhaustive software library caters for external control by providing a full set of routines for system configuration, testing, system-status and monitoring [29]. Macintosh II series computers take over the responsibility of program development and run-time operator-control. The central System Supervisor Mac [32] initialises the readout and checks the status of all elements either directly or indirectly by initiating VMEbus tasks such as the Event Coordinator. To provide a central focus of monitoring it also communicates with event tasks and subsystems either through VMEbus or over Ethernet. Additionally it serves the network with status information for external monitoring. The graphics-based philosophy of the Mac ensures that the operator is presented with, arguably, one of the most human-interactive interfaces, available today, to a complex real-time system.

REMARKS AND CURRENT STATUS

It is because of the architectural structure defined in VMEbus that nothing more than a Macintosh can sit at the top of the protocol and provide the overall supervision. From the master VMEtaxi the system can manage up to 32 branches and 16 full-event units, including the parallel filter farm, providing status information and accepting control command sequences from the System Supervisor. Presently 60 Kbyte-sized events

can be coordinated with 12 branches on H1 at rates greater than 50 Hz; with the Mark-2 VMEtaxi upgrades, 200 Hz will be surpassed. The primary sources of dead-time remain at the front-end, where future injections of greater processing power and refined triggers are made easier by working within an internationally accepted bus standard. The availability of networking standards has led to a proliferation of online usage, including the ability to control and monitor the system externally even from the latest generation of portable notebook computers.

The H1 data acquisition system is now enjoying the first rewards from HERA. The choice of an open-bus specification contributed much to it being prepared well on-time; the online-reconstructed event of Figure 11 was recorded during first HERA operation [33]. Moreover, it has performed with a relatively high level of stability when one considers the amount of sophisticated electronics and software involved.

CONCLUSION

A real-time data acquisition system of the complexity and size of H1 has to be autonomous, flexible, upgradable and modular. Design-engineers and software-programmers have to come together from widely different backgrounds to implement a coherent framework working in unison. An open-bus standard, providing it can accommodate the bandwidths required and is correctly "architected", is therefore a natural choice in the current evolutionary time-slot.

ACKNOWLEDGEMENTS

Many engineers, programmers, physicists and technicians have contributed towards the H1 data acquisition system. The systems software for the central data acquisition and filter farm has been written mainly by A.Campbell, E.Defeur, P.Fuhrmann, W.J.Haynes, P.Hill, J.Olsson and M.Zimmer.

Engineering and technical support has been provided by R.Hedgecock, T.Külper, H.Quehl and A.Sirous. The hardware and software for the front-end subsystems and trigger integration has been implemented mainly by J.Coughlan, F.Descamps, D.Düllmann, E.Banas, G.Eckerlin, E.Elsen, E.Evrard, A.Fomenko, J.Huber, P.Huet, S.Kolya, H.Krehbiel, F.Martin, D.Mercer, G.Noyes, M.Savitski, I.Sheviakov, U.Straumann, J.Tutas, A.Usik, C.Vallée, E.Wünsch and W.Zimmermann.

Appreciation must be expressed to the excellent cooperation of the companies and individuals involved in the development and implementation phases, notably Micro-Research of Helsinki, C.E.S. of Geneva and to B.G.Taylor of CERN for his continued support of MICRON/MacVEE. The basic architectural concepts, of using VMEbus in large High-Energy Physics experiments, were pioneered by UA1 back in the early 80s and much is owed to the foresight and ideas of S.Cittolin. Finally, the author is grateful for the support provided him by the UK Science and Engineering Research Council, the Commission of the European Communities Directorate General for Science Research and Development together with the generous hospitality of the DESY laboratory.

REFERENCES

- [1] H1 Collaboration, *Technical proposal for the H1 detector*, DESY, Hamburg, FRG, March 1986.
- [2] *The VMEbus specification*, IEEE standard 1014.
- [3] W.J.Haynes, *The H1 data acquisition system*, Int. Conf. Computing in High Energy Physics '91, Tsukuba, Japan, March 1991, Universal Academy Press, Inc.
- [4] W.J.Haynes, *The data acquisition system for the HERA H1 experiment*, Rutherford Appleton Laboratory, UK, RAL 90-039. Workshop on VMEbus in Physical Research, Dubna, June 1990.
- [5] E.Elsen, *The H1 trigger and data acquisition system*, Int. Symp. on Electronic Instrumentation in Physics, Dubna, May 1991.
- [6] U.Straumann, *First experience with HERA beams at H1*, Stanford Summer Institute Topological Conference, Stanford, USA, July 1992.
- [7] Th.Wolff et al, *A drift-chamber track finder for the first-level trigger of the H1 experiment*, 6th Wire Chamber Conf., Vienna, Austria, February 1992.
- [8] A.Campbell, *A RISC multiprocessor event trigger for the data acquisition system of the H1 experiment at HERA*, Conf. Real Time '91, Jülich, FRG, June 1991.
- [9] *The VME Subsystem Bus (VSB) specification*, IEEE standard 1096.
- [10] Creative Electronics Systems SA, *DSP 8150 Digital Signal Processor VMEbus module*, Geneva, CH, 1989.
- [11] F.Descamps, F.Martin, C.Vallée et al., *The H1 calorimeter readout system*, Int. Conf. Computing in High Energy Physics '92, Annecy, France, September 1992.
- [12] *VICbus, VME Inter-Crate Bus*, ISO/IEC JTC 1/SC26.
- [13] W.Zimmermann et al., *A 16 channel VME Flash ADC system*, DESY, February 1989.
- [14] D.Mercer et al., *The H1 Scanner M1070*, Univ. of Manchester, UK, July 1990.
- [15] S.Kolya, *H1 drift chamber data acquisition system*, DESY, October 1992.
- [16] A.I.Lebedev et al., *Electron tagger trigger and data acquisition*, DESY, May 1988.
- [17] J.Tutas, *The limited streamer tube system of H1*, 26th Int. Conf. on High Energy Physics, Dallas, USA, August 1992.
- [18] E.Evrard et al., *MWPC readout system based on Easybus*, DESY, 1990.
- [19] H1 Collaboration, *Technical proposal to build Silicon tracking detectors for H1*, DESY, PRC 92/01, July 1992.
- [20] Creative Electronics Systems SA, *Fast Intelligent Controller FIC 8230/8232*, Geneva, CH, 1987/1992.
- [21] Creative Electronics Systems SA, *RAID 8235, R3000 VMEbus controller*, Geneva, CH, 1990.
- [22] Creative Electronics Systems SA, *VME/VSB Dual Ported Memory DPM 8242*, Geneva, CH, 1989.
- [23] B.G.Taylor, *MICRON: VMEbus and CAMAC access from Macintosh II*, ESONE Conf. VMEbus in Research, Zürich, CH, October 1988.
- [24] Micro-Research Oy, *VMEtaxi user manual*, Helsinki, SF, July 1990.
- [25] *The Scalable Coherent Interface (SCI)*, IEEE standard 1596.
- [26] J.Coughlan, D.Düllmann, M.Savitski, M.Zimmer, *Object orientated programming for online systems at H1*, Int. Conf. Computing in High Energy Physics '92, Annecy, France, September 1992.
- [27] W.J.Haynes, *VME_TOOLS, VMEbus interaction from the Macintosh Programmer's Workshop shell*, DESY, August 1989.
- [28] E.Deffur, P.Fuhrmann, M.Zimmer, *Communication on the H1 Local Area Network*, DESY, February 1991.
- [29] W.J.Haynes, *VMEXI_SSP : VMEtaxi System Software Package*, DESY, Version 3.9, June 1992.
- [30] H.Krehbiel, *The H1 trigger control system*, DESY, September 1988.
- [31] G.Hochweller et al, *VME-HSSL VMEbus High Speed Serial Link*, DESY, May 1990. Int. Conf. Open Bus Systems '92, Zürich, CH, October 1992.
- [32] M.Zimmer, *Graphics-orientated operator interfaces at H1*, Int. Conf. Computing in High Energy Physics '91, Tsukuba, Japan, March 1991, Universal Academy Press, Inc.
- [33] P.Hill, *MacEvLook Online version of H1 event display*, DESY, October 1991.