THE HYPERCP DATA ACQUISITION SYSTEM[1]

D. M. Kaplan[2], W. R. Luebke, A. Chakravorty, C. G. White, S. L. White
Illinois Institute of Technology, Chicago, IL 60616, USA

Y.-C. Chen, C. Ho, P. K. Teng, C. Yu
Academia Sinica, Taipei, Republic of China

P. Gu, K. B. Luk, W.-S. Choong, B. Turko
University of California and
Lawrence Berkeley National Laboratory, Berkeley, CA 94720

C. James
Fermilab, Batavia, IL 60510

E. C. Dukes, K. Nelson
University of Virginia, Charlottesville, VA 22901

Abstract

For the HyperCP experiment at Fermilab, we have assembled a data acquisition system that records on up to 45 Exabyte 8505 tape drives in parallel at up to 17 MB/s. During the beam spill, data are acquired from the front-end digitization systems at ≈ 60 MB/s via five parallel data paths. The front-end systems achieve typical readout deadtime of $\approx 1\,\mu$s per event, allowing operation at 75-kHz trigger rate with $\lesssim$30% deadtime. Event building and tapewriting are handled by 15 Motorola MVME167 processors in 5 VME crates.

1 Introduction: The HyperCP Experiment

The HyperCP experiment [1] (Fermilab E871) aims to achieve 10^{-4} sensitivity to possible hyperon *CP* violation. This is two orders of magnitude beyond the most sensitive current limit [2]. Since so far only "indirect" *CP* violation[3] has been definitively observed [3], and that only in decays of the neutral kaon, it is of interest to search as sensitively as possible both for "direct" *CP* violation and for *CP* violation occurring outside the K^0 sector. Hyperon *CP* violation qualifies in both of these respects.

While hyperon decays should violate *CP* symmetry in the Standard Model, the signal is expected to be below the HyperCP experiment's threshold of sensitivity [4]. However, mechanisms capable of inducing the observed matter-antimatter asymmetry of the universe [5] require *CP* violation at levels possibly exceeding the Standard-Model prediction [6], and thus possibly detectable in HyperCP.

[1]Presented at the *Sixth Annual LeCroy Conference on Electronics For Particle Physics*, Chestnut Ridge, NY, May 28–29, 1997.
[2]E-mail: kaplan@fnal.gov
[3]i.e., occurring via meson-antimeson mixing.

Hyperon CP violation can be studied most sensitively by comparing the Lee-Yang [7] α parameter characterizing the parity violation in the decay of a hyperon species with that for the corresponding antihyperon species [8]. To do this, HyperCP will reconstruct 10^8 to 10^9 hyperon and antihyperon decays in the cascade decay modes $\Xi^- \to \Lambda^0\pi^- \to p\pi^-\pi^-$ and $\overline{\Xi}^+ \to \overline{\Lambda}^0\pi^+ \to \overline{p}\pi^+\pi^+$. This data sample will also give interesting sensitivity to various rare hyperon decays, and (since the hyperon beam contains more charged kaons than cascades) to possible direct CP violation in charged-kaon decay as well [1].

The HyperCP spectrometer is shown schematically in Fig. 1. An intense unpolarized beam of Ξ^- ($\overline{\Xi}^+$) hyperons is produced at 0° by 800 GeV protons striking a metal target, with the secondaries momentum-selected by means of a curved magnetic channel set to 150 GeV with 25% FWHM momentum bite. Following a 13 m evacuated decay pipe the hyperon decay products are detected in a high-rate magnetic spectrometer using multi-wire proportional chambers (MWPCs). (The needed rate capability is determined by the $\approx$40 MHz of charged particles, dominantly pions and protons, emerging from the channel.) To switch between hyperon and antihyperon running we reverse the polarities of the hyperon and analyzing magnets, selecting either positive or negative beam. This approach results in nearly identical acceptance and efficiency for hyperons and for antihyperons, essential for minimizing biases. To maintain equal rates in the detectors, different target lengths are used in the two modes.

A simple and highly CP-invariant trigger is provided by requiring the coincidence of at least one particle on the right and one on the left of the beam. This is done using the "pion" and "proton" hodoscopes, which are located sufficiently far downstream of the analyzing magnets that the proton and pions from the hyperon decay have all separated from the beam. To suppress low-energy backgrounds, the trigger also requires a minimum energy deposit in the hadron calorimeter. Nevertheless the trigger rate is dominated by background due to interactions of beam particles (and neutrals accompanying the beam) in the spectrometer material.

To provide sensitivity to rare decay modes involving leptons, there is also a muon detection system, consisting of proportional-tube planes interspersed with iron hadron absorber. Muon triggers are formed with signals from four planes of scintillation hodoscopes located behind the absorber.

2 The HyperCP Data Acquisition System

The acquisition of so large a hyperon sample requires a highly capable data acquisition system, designed for 100 kHz trigger rate and 20 MB/s average data rate to tape. Although the actual rate of reconstructable hyperon decays is $\lesssim$1 kHz, this allows sufficient "headroom" to accommodate the rate of background triggers. Our system is in many respects an update of the data-acquisition system used by Fermilab E791 [9], which handled 10 kHz trigger rate and recorded to tape at 10 MB/s. The major differences are

- replacement of the front-end digitization subsystems;

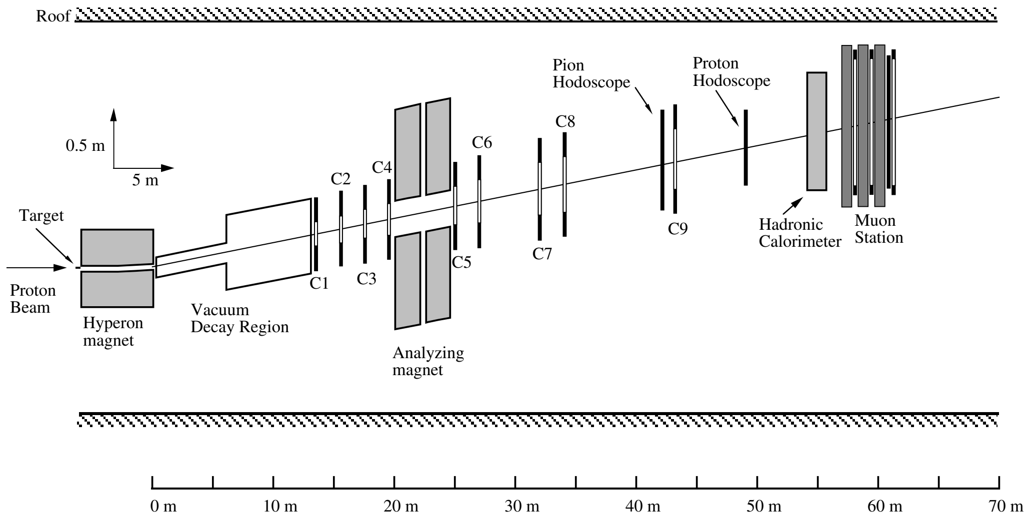

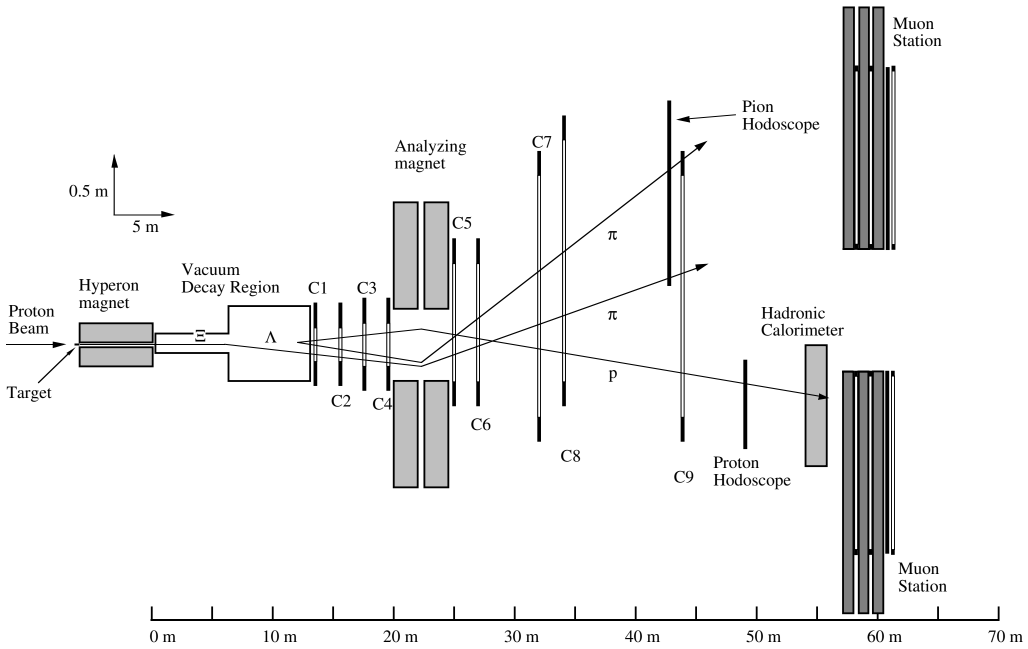

Figure 1: The HyperCP spectrometer.

- increasing the size of the spill buffer from 640 to 960 MB;
- use of (68040-based) Motorola MVME167 processors instead of (68020-based) ACP I processors;
- use of Ciprico RF3573 tape controllers instead of RF3513s;
- use of Exabyte EXB-8505 tape drives instead of EXB-8200s;
- use of Silicon Graphics back-end workstations running Fermilab DART software [10] instead of a VAX.

Fig. 2 shows the structure of the data-acquisition system. Information from the detectors is digitized and sparsified by the front-end systems, then transmitted via optical fiber to the VDAS (video data-acquisition system) spill buffers [11] in the control room. Five event-building systems operating in parallel access the VDAS to assemble events and record them on tape. Each event-building system is housed in a single VME crate and comprises three MVME167 processors, five Event Buffer Interfaces (to provide access to the VDAS), and three Ciprico RF3573 SCSI host adapters [12] which each control three EXB-8505 tape drives [13]. Control of the system is provided via Ethernet links to each "Boss" MVME167 processor and to two booting and monitoring processors (not shown in the diagram) located in the Electronics Hall.

2.1 Front-End Digitization System

The major task of the front-end system is digitization of data from the $\approx$20,000-wire MWPC system. In addition it must read out the trigger hodoscopes, calorimeter, and muon system. To achieve sufficiently low deadtime at up to 100 kHz average event rate, all systems must satisfy a digitization and readout time specification of $\lesssim 1\,\mu$s/event.

2.1.1 MWPC CR System

Readout of the MWPCs and trigger hodoscopes is carried out using an adaptation of the Nevis Laboratories MWPC Coincidence Register (CR) system [14]. Having already built and operated nearly 10,000 channels of this system in Fermilab Experiment 789, we chose for reasons of cost-effectiveness to augment that existing system for HyperCP rather than building a new system from scratch. The CR system (Fig. 3) consists of multiple crates (built using CAMAC-standard card cages) each containing inexpensive two-layer printed circuit boards communicating via a custom ECL backplane. A typical crate accommodates 22 32-channel CR cards, thus instrumenting 704 MWPC wires.

The CRs (Fig. 4) accept differential-ECL inputs via a front-panel card-edge connector and receive gate signals from the backplane. Inputs are conveyed to the CR front panel via Ansley multiconductor ribbon cables [15], which are sufficiently oversized in length to provide trigger delay. An optional test voltage can be applied to

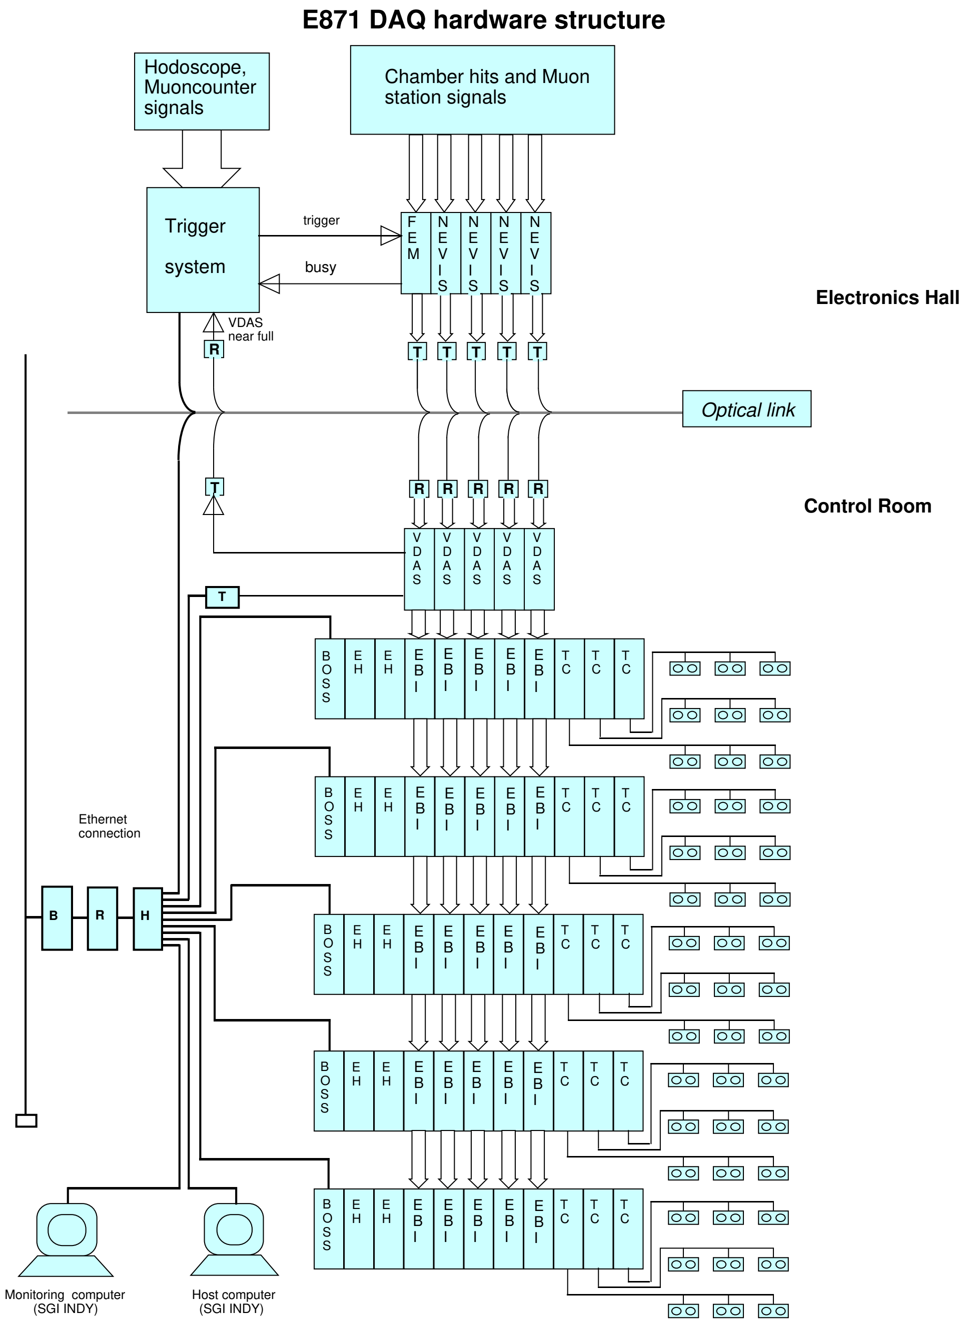

Figure 2: Block diagram of the HyperCP data acquisiton system.

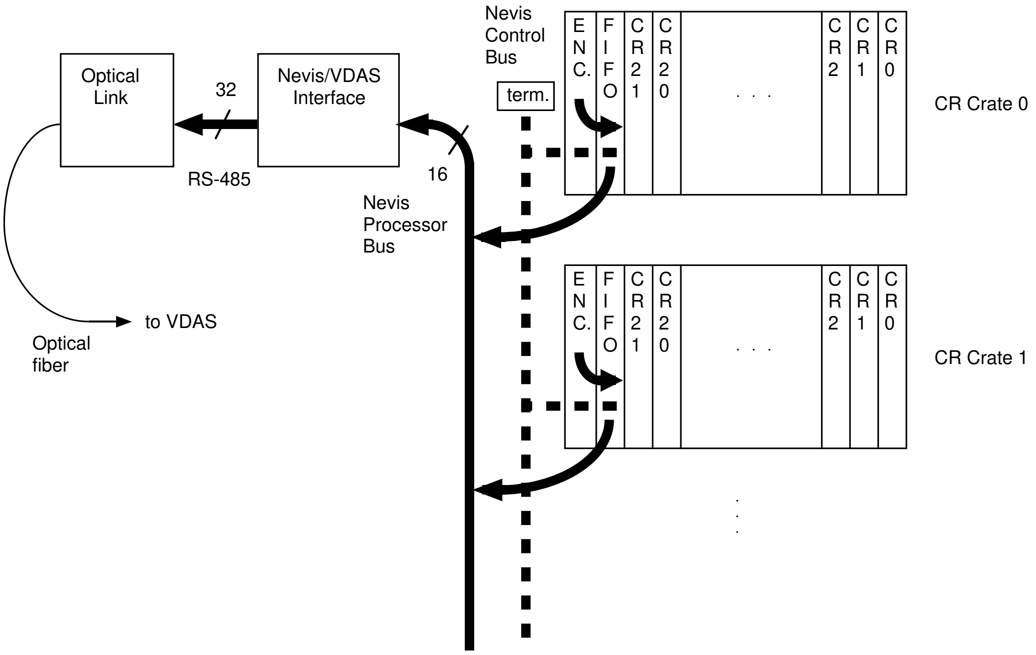

Figure 3: Block diagram of a CR-system data path.

the backplane to force all channels on for diagnostic purposes. A Prompt Reset signal (not used in HyperCP) supports multiple trigger levels: a gate can be generated with a low-level trigger, and all CRs can then be reset rapidly if a higher-level trigger is not satisfied. The Group Reset signal is used to reset one CR at a time during the readout cycle (described below). To allow easy interchange of modules, card addresses are hard-wired into the backplane by extraction of pins from the (card-edge) backplane connectors. Each CR drives all five address signals, but only those whose pins are intact are asserted on the backplane.

Readout and sparsification of data from the CRs is accomplished by the Encoder card, of which there is one per crate. Each CR card containing non-zero data in turn issues its data on the backplane as 32 bits in parallel for latching by the Encoder. The Encoder then employs priority encoders to cycle through only the hit bits (i.e. those that are non-zero), encoding each as a 10-bit address and transmitting it to a buffer memory via a front-panel cable. Unless the buffer becomes full, hit addresses are transmitted at a rate of one per clock cycle. In the mean time, a reset pulse is sent to the CR card, and control of the backplane passes to the next CR having non-zero data. The Encoder also contains a 4-bit counter which is incremented for each event, whose value is inserted into the Complete word (which follows the last word of event data) for use as an Event Synchronization Number (ESN). By comparing ESNs from the various Encoders, it is possible to verify that all data within an event pertain to the same trigger. Readout of the CR system is synchronous to a central clock, and has been shown to operate reliably with clock period as low as 75 ns.

The Encoder output cable and transmission protocol obey the Nevis Processor Bus standard [16]. The (single-ended ECL) Processor Bus cable includes a 16-bit

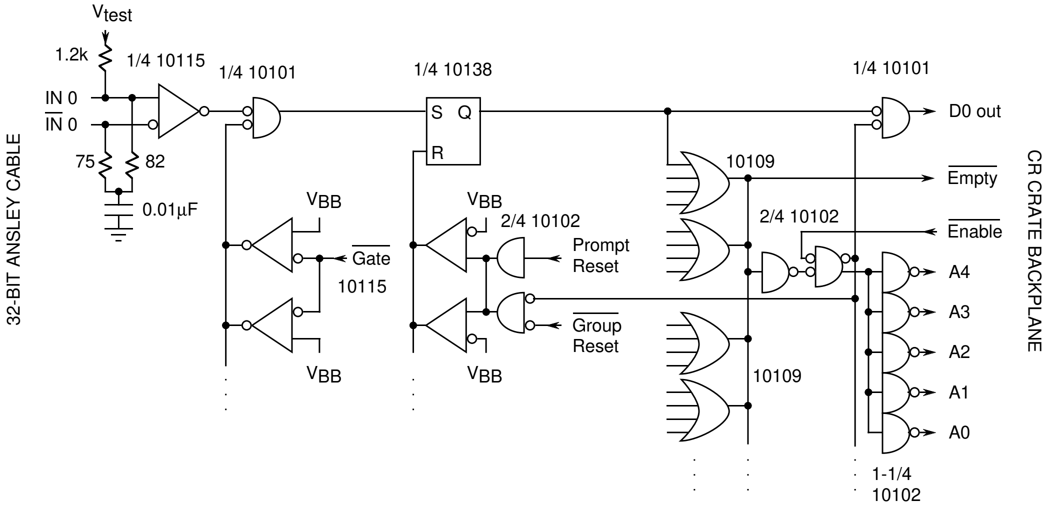

Figure 4: Schematic diagram of CR card (one channel of 32 shown).

data space and a 4-bit name space as well as control lines. The data space is used to transmit the 10-bit wire address, while the name space identifies the Encoder. The Encoder output logic includes a token-passing scheme for control of the Processor Bus. Thus up to 16 Encoders may be chained together on a single cable.

In the original ca.-1970 design, token passing was used within the crate to pass control from one CR to the next. In 1980, a fast readout scheme was implemented for Fermilab E605, in which token passing was eliminated in favor of priority encoding via a Fast Carry card mounted to the crate backplane. The Fast Carry card receives "carry out" ($\overline{\text{Empty}}$) signals in parallel from all CRs having non-zero data and issues a "carry in" ($\overline{\text{Enable}}$) to the first such card.

In addition to the Fast Carry, Gate Fanout cards are mounted to the crate backplane at each end. They receive NIM gate signals and fan them out to the individual CR-card slots. They also serve to terminate the backplane data bus. Provision is made to gate all CRs within half of a crate with a common gate signal, or (in case fine tuning of gate widths or delays is needed) to gate sets of four adjacent CRs in common. Gate signals as short as 10 ns have been used reliably.

Data from each Encoder are sent to a FIFO Buffer card. The FIFO design is based on the Nevis Processor 128-word Buffer module [16]. Our version contains 256 16-bit words of 12-ns ECL RAM which can be both written and read on each clock cycle. An additional memory bit records the state of the Complete signal from the Processor Bus. This bit serves as a marker between events, allowing a variable number of events to be buffered as needed to smoothe out event-rate fluctuations. Like the Encoders, the FIFO cards employ a token-passing scheme on their outputs, allowing multiple FIFOs to share a single Processor Bus cable.

In addition to data readout with sparsification, it is desirable to record certain information (e.g. bits recording the state of the trigger logic) which are not sparse. To handle non-sparse data efficiently, we have developed a "piggyback" card for one Encoder in the system, allowing it to pass bit patterns directly from certain selected

CRs to its output cable. To ensure that the non-encoded CRs send their data onto the bus regardless of the state of their inputs, one input bit out of each 32 must be forced on, leaving 31 bits per CR available. This is necessary since there is no addressing information in the event record for the non-sparsified data, thus they can be identified only by their position and sequence in the event.

2.1.2 FEM System

The muon-detector hodoscopes and proportional tubes are digitized by the FEM system, designed and built by the Academia Sinica group for our experiment. The system is housed in a 9U VME crate. Each front-end module (FEM) receives 128 differential-ECL front-panel input signals, which are converted to TTL and, if coincident with a gate signal, stored in a 128-bit latch. At the same time, an externally-generated 8-bit ESN is latched. The ESN and data latches (organized as 16-bit words) are then read out (synchronously to a 20 MHz clock) into a 16-bit-wide multi-event FIFO 2048 words deep. The FEM thus has 0.45 μs of front-end deadtime, after which it is free to process another event.

Each FEM includes an on-board encoder, which accepts data from the FIFO synchronously to a 10 MHz clock, computes hit addresses for all non-zero bits, and stores them in a buffer memory. This encoding cycle takes 0.1 μs/hit + 1.7 μs.

The FEMs are read out across the backplane at a rate 0.1 μs/hit (plus 0.2 μs of overhead per event) by a FEM Interface (FEMI) module. The FEMI includes pipelined cluster-finding circuitry, so that clusters of up to 7 adjacent wires can be encoded as a single output word. The output from the FEMI is first stacked into 32-bit words and buffered in a FIFO, then sent via an optical link to the VDAS.

Given the typical occupancy of the muon system, the FEM system puts out about 20 hits/event on average. This entails some 2.2 μs/event of data transfer over the FEM backplane to the FEMI. Since readout to the VDAS is 32 bits wide and synchronous to the 10 MHz clock, it is twice as fast as readout into the FEMI.

2.1.3 Calorimeter ADC System

The hyperon decays of interest to HyperCP produce a high-momentum hadron whose electric-charge polarity is opposite to that of the beam. A hadronic calorimeter is employed to exclude both halo muons and interactions of secondary-beam particles in the spectrometer material, both of which could satisfy the simple left-right trigger requirement. The calorimeter is segmented twice laterally and 4-fold in depth, providing eight signals which must be digitized for off-line trigger verification and monitoring.

Although the calorimeter is used primarily as a trigger device, the digitization should have 14-bits of dynamic range, to allow gain monitoring with muon minimum-ionizing signals as well as the measurement of high-energy hadronic showers. The small number of channels involved allows a digitization solution which is simple, e.g. avoiding multiple ranges, and fast. The ADC comprises an integrate-and-hold circuit followed by a 14-bit 3 MHz sampling flash converter. Following the conversion cycle the eight 16-bit data words are sent via Processor Bus to a FIFO buffer module, and thence (along with MWPC CR data) via a Nevis/VDAS Interface and optical link to the VDAS (see below).

2.2 VDAS Spill Buffer and Interfaces

As in E791, the data from the front-end subsystems are accumulated in VDAS buffers [11]. These accept 32-bit input longwords at speeds up to 100 ns/longword using the RS-485 signal protocol; details may be found in the references. To interface between the VDAS system and the Processor Bus, we have built Nevis/VDAS Interfaces for each of the four CR-system data paths. The interface serves three main functions:

1. stack pairs of 16-bit words from the Processor Bus into 32-bit longwords;
2. precede each event fragment with a leading byte count by storing the event data temporarily in a 2048-word FIFO and counting bytes until Complete;
3. compare ESNs to verify that data from all encoders within a data path are from the same event.

The 16-bit data space of the Processor Bus exceeds that required for the 10-bit wire address within each CR crate. To use memory efficiently, the Nevis/VDAS interface packs the 4-bit Encoder name into unused high-order data bits. The Complete word is then suppressed from the output stream, reducing the event length by about 12%.

The E791 collaboration constructed a VDAS system configured as eight 80 MB buffers housed in four VDAS crates, for a total of 640 MB [9]. To accommodate the increased data rate in HyperCP, we have built additional memory cards to increase the total capacity of the system to 960 MB. With our faster front-end systems, we have reduced the number of data paths from eight to five, while doubling the overall throughput. We thus configure our VDAS system as five buffers housed in five crates.

Due to the long distance ($\approx 500'$) between our electronics room and our control room, we transmit the data from the Nevis/VDAS interfaces and the FEM system to the VDAS via high-speed optical links [17]. Five optical fibers (one per data path) transmit data from the electronics room to the control room, and one returns the VDAS Near Full condition (the logical OR of the Near Full signals from the five VDAS buffers) to the electronics hall. VDAS Near Full is used to veto triggers.

2.3 VME Event-Building System

As in E791, the tasks of event building are divided into supervisory tasks, carried out by a Boss processor, and actual data gathering and formatting, carried out by Event Handlers. Using 33 MHz 68040-based Motorola MVME167 modules for these processing nodes, we find that two Event Handlers per crate are sufficient for "optimal" use of the VME bus.[4]

[4]By optimal, we mean utilizing the bus to the maximum capability within the limitations of the chosen architecture.

Event Buffer Interfaces (EBIs) [9] provide VME access to the VDAS data. The Event Handlers transfer data via 32-bit-wide DMA from the VDAS through the EBIs to tapewriting buffers in on-board memory local to each MVME167. The EBI token-passing scheme serves to synchronize access to each VDAS buffer among the five VME crates. If we assume (for the sake of discussion) equal-length event fragments in all data paths, this scheme results in maximal parallelism, i.e. all VDAS buffers constantly being read by Event Handlers. Given the random fluctuations in event-fragment length from event to event and among data paths within an event, VDAS buffers containing shorter fragments are typically kept waiting while longer fragments are being processed, and parallelism is less than maximal.

Since our event fragments are small (typically 20 to 30 longwords), care must be taken in coding the Event Handler inner loop to maximize throughput. Using EBIs from E791, the DMA burst data-transfer rate is measured to be 300 ns/longword. Given DMA startup overhead, we are able to build events at a maximum rate of 3.5 MB/s per crate, and five crates operating in parallel thus give 17 MB maximum throughput to tape.

2.4 System Performance and Anticipated Improvements

Since the April 1997 start of our data run, we have been carefully studying the performance of the detectors as a function of beam intensity, and at our current typical 75 kHz rate of 580-byte events for 19 seconds of beam spill every minute, the throughput just described has been adequate for our needs. At this event rate we experience $\approx$30% deadtime as monitored by the ratio of triggers accepted to triggers generated. The deadtime should be dominated by the front-end deadtime in the most burdened crate. With 14 data words per event, the unencoded CRs have the largest average readout deadtime of 1 μs/event. In addition, there is $\approx 1\,\mu$s overhead in forming the trigger and issuing gates to the CRs, due in part to the unequal delays of signals from the various MWPCs. If the rate of beam spill were perfectly uniform, at 75 kHz, 2 μs/event would imply 15% deadtime. Given Tevatron beam-spill fluctuations, deadtimes are typically about 1.5 times the smooth-spill estimate, so we expect $\approx$ 20% deadtime if the unencoded data dominate. However, as shown in Fig. 5, fluctuations in MWPC occupancy can result in some other CR crate having more data than the unencoded CRs in any given event, which possibly accounts for the larger observed deadtime. These issues are under active investigation, and further optimization is anticipated.

We are investigating several avenues to meet our design goal of 20 MB/s to tape. We are installing new fast EBIs which can transfer data at 200 ns/longword. We are also exploring further optimizations of the inner event-building loop and of the tape controllers' VME data-transfer speed. We anticipate that these improvements will keep pace with our operation of the experiment at higher rates as our understanding of detector and event-reconstruction limitations improves.

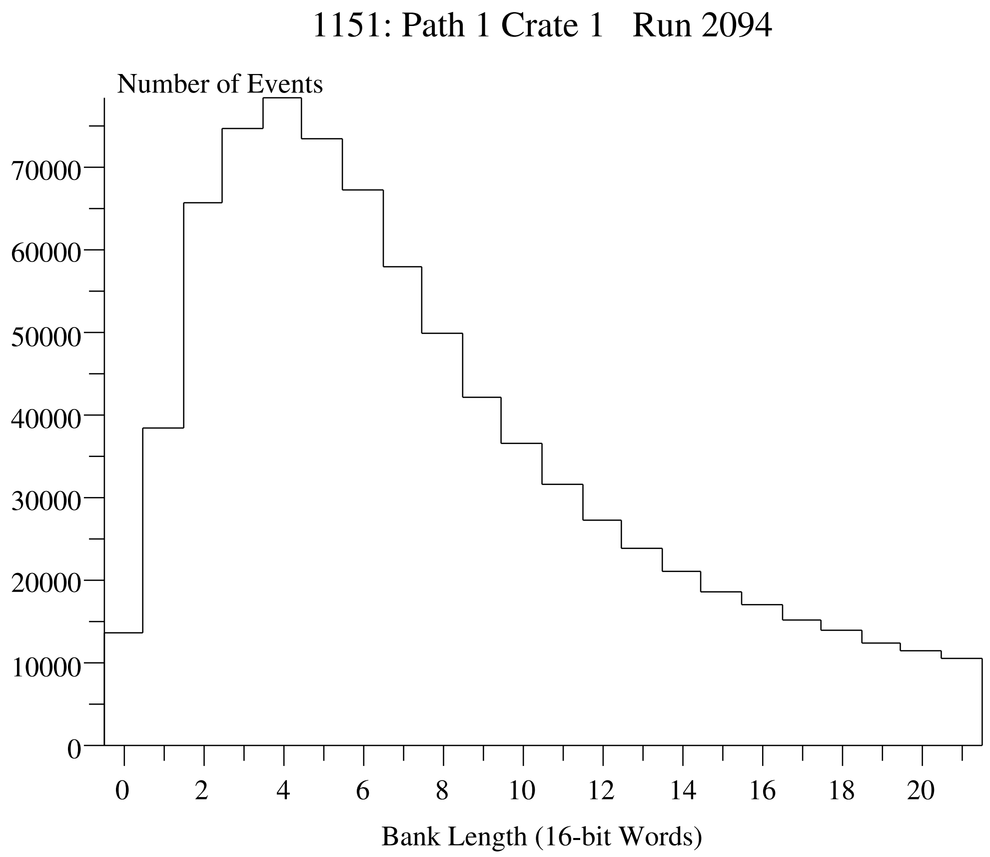

Figure 5: The distribution of number of 16-bit words transferred from a single CR crate per event (note the tail extending well past the 14 unencoded words/event).

3 Acknowledgements

We thank T. Carter, K.-C. Cheng, J. Meadows, and S. Hansen for help and useful advice, and the DART group of the Fermilab Computing Division for their support. This work was funded by the U.S. Department of Energy and the [WHOEVER FUNDS ACADEMIA SINICA].

References

[1] J. Antos *et al.*, Fermilab Proposal 871 (revised), Mar. 26, 1994; see also E. C. Dukes, "A New Fermilab Experiment to Search for Direct *CP* Violation in Hyperon Decays," in *Proc. 11th Int. Symp. on High Energy Spin Physics*, Bloomington, IN, Sept. 15–22, 1994; K. B. Luk, "Experimental Prospects for Observing CP Violation in Hyperon Decay," **Heavy Quarks at Fixed Target**, B. Cox, *ed.*, p. 81; D. M. Kaplan, "Fixed Target *CP* Violation Experiments at Fermilab," IIT-HEP-96-2, to appear in *Proc. 5th Hellenic School and Workshops on Elementary Particle Physics*, Corfu, Greece, 3–24 Sept. 1995, hep-ex/9509009.

[2] P. D. Barnes *et al.*, Phys. Rev. C **54**, 1877 (1996).

[3] B. Winstein and L. Wolfenstein, Rev. Mod. Phys. **65**, 1113 (1993).

[4] For recent reviews see D. Chang and C.-H. Chen, Chin. J. Phys. **34**, 831 (1996), hep-ph/9606424; P. Zenczykowski, Phys. Rev. D **55**, 1688 (1997).

[5] A. D. Sakharov, Pis'ma Zh. Eksp. Teor. Fiz. **5**, 32 (1967) [JETP Lett. **5**, 24 (1967)].

[6] H. Quinn, Nucl. Phys. Proc. Suppl. **50**, 17 (1996); P. Huet and E. Sather, Phys. Rev. D **51**, 379 (1995).

[7] T. D. Lee and C. N. Yang, Phys. Rev. **108**, 1645 (1957).

[8] J. F. Donoghue, B. R. Holstein, and G. Valencia, Int. J. Mod. Phys. A **2**, 319 (1987).

[9] S. Amato *et al.*, Nucl. Instrum. Meth. **A324**, 535 (1993).

[10] R. Pordes *et al.*, "Fermilab's DART DA System," presented at Computing in High Energy Physics (CHEP 94), San Francisco, CA, 21–27 Apr. 1994, FERMILAB-CONF-94-103.

[11] A. E. Baumbaugh *et al.*, IEEE Trans. Nucl. Sci. **NS-33**, 903 (1985); K. L. Knickerbocker *et al.*, IEEE Trans. Nucl. Sci. **NS-34**, 245 (1986).

[12] Ciprico, 1161 Homestead Ln., Chanhassen, MN 55317.

[13] Exabyte Corp., 1685 38th St., Boulder, CO 80301.

[14] Original design ca. 1970 by F. W. Sippach, Columbia University Nevis Laboratories, P.O. Box 137, Irvington, NY 10533, updated for fast readout by F. W. Sippach and H. Cunitz, Columbia University Nevis Laboratories, in 1981. The CR design was updated from ECL II to ECL 10,000 by J. A. Crittenden and D. M. Kaplan in 1984.

[15] Manufactured originally by Ansley Co. (later Thomas & Betts/Ansley) during the 1970s and 1980s but no longer commercially available.

[16] W. Sippach, G. Benenson, and B. Knapp, IEEE Trans. Nucl. Sci. **NS-27** (1980) 578; Y. B. Hsiung *et al.*, Nucl. Instr. and Meth. **A245** (1986) 338; D. M. Kaplan, in *Proc. Symp. on Recent Developments in Computing, Processor, and Software Research for High-Energy Physics*, Guanajuato, Mexico, May 1984, R. Donaldson and M. Kreisler, *eds.*, Fermilab, 1984, p. 31.

[17] J. T. Meadows *et al.*, FERMILAB-CONF-95-110, presented at *9th Conference on Real-Time Computer Applications in Nuclear, Particle and Plasma Physics*, East Lansing, Michigan, 22–26 May 1995.