THE BABAR TRIGGER, READOUT AND EVENT GATHERING SYSTEM

C.T. DAY, R.G. JACOBSEN, J.F. KRAL, M.E. LEVI

Lawrence Berkeley National Laboratory,
1 Cyclotron Road, Berkeley, CA 94720, USA

J.L. WHITE

Stanford Linear Accelerator Center,
P.O. Box 4349, Stanford, CA 94309, USA

FOR THE BABAR TRIGGER, DAQ AND ELECTRONICS GROUP

The design of the trigger and data acquisition architecture for the *BABAR* experiment is described, with particular emphasis on readout modules and event assembly.

1 Introduction

Future high-luminosity high energy physics experiments at fixed target and colliding beam experiments are facing heavy demands on rapid processing and transfer of large volumes of data. To meet the stringent requirements posed by operation at a high-rate facility, the *BABAR* trigger and data acquisition architecture[1] has been designed to take advantage of recent advances in field programmable gate arrays (FPGAs), microprocessors and network switches.

The e^+e^- beam crossings at the PEP-II collider will occur essentially continuously (at 238 MHz) when compared to the electronics response time. Under severe conditions, beam-induced backgrounds are expected to produce very high occupancies (e.g. hundreds of MHz per layer in the silicon vertex tracker), as well as large trigger rates (e.g. several kHz of interactions with two or more charged tracks reaching the calorimeter). At the design luminosity of $3 \cdot 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the rate of physics to be recorded is about 100 Hz, of which 30 Hz are from $\Upsilon(4S)$ decays.

The rates, predicted by calculations of detector response to both physics and background, were used to model candidate systems. The vertex detector will sustain up to about ten times the predicted occupancies. It is relatively straightforward to select B physics events with greater than 99% efficiency, even during severe running conditions. The trigger and DAQ system is designed to accommodate up to ten times nominal conditions with $< 10\%$ deadtime. The *BABAR* experiment is expected to take colliding beam data in 1999.

2 Trigger and DAQ Architecture

An innovative architecture based on a philosophy of early data processing and data reduction is proposed to handle the high occupancies and large trigger rates. The trigger is comprised of two levels: a hardware processor (L1) that reduces the

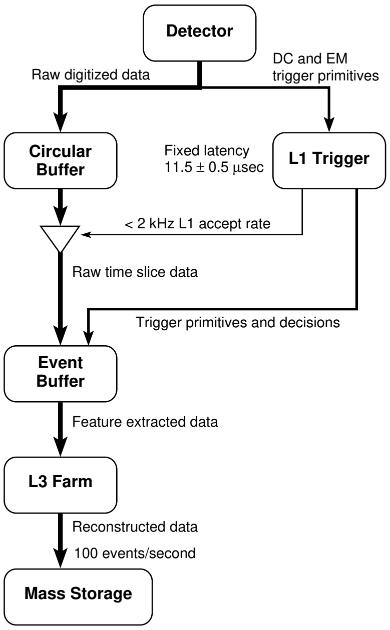

Figure 1: Trigger and DAQ event flow.

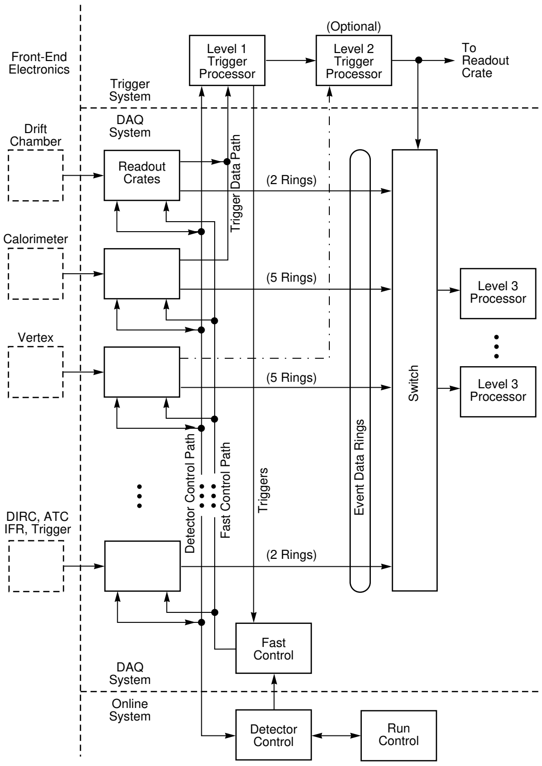

Figure 2: DAQ system block diagram.

trigger rate to a maximum of 2 kHz, followed by software algorithms executing in commercial CPUs (L3). As is seen in Figure 1, during the L1 latency period of 12 μs, data are stored digitally, in circular buffers either on front-end ICs or in readout cards. After an L1 strobe, time slices are transferred via fast links to back-end buffers, where they are available for feature extraction and data sparsification by commercial microprocessors. These will be located on the sections of the VME-based readout modules that are uniform across all detector systems. Processed data from each event (about 25 kb each for a maximum sustained rate of 50 Mb/s) are assembled into an L3 workstation via a commercial high-speed network. The parallel nature of the electronics system is illustrated in Figure 2. The total rate to mass storage is 2.5 Mb/s.

If the L1 accept rate is greater than either the rate for feature extraction or event assembly, the first software filter will be implemented in Level 2 processors. Ongoing work focuses on (1) refined L1 trigger algorithm simulation studies to show that the trigger rate can be reduced to $<$ 2000 Hz while retaining low-multiplicity τ and $\gamma\gamma$ physics modes, (2) prototype work on the readout modules and benchmarks of feature extraction algorithms for systems like the calorimeter, and (3) demonstration through benchmarks and simulation that event assembly can proceed at 2000 Hz.

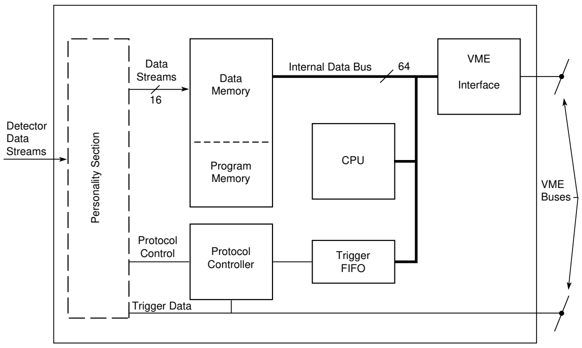

Figure 3: Readout card.

3 Trigger Processor

During normal running, two tracks *or* two clusters consistent with minimum ionizing particles form an orthogonal trigger condition with high efficiency for low-multiplicity physics (e.g. 89% for $\tau^+\tau^-$ events where both τs decay to pions). At ten times nominal background occupancies, the rate is estimated to climb from an acceptable 1200 Hz to 8500 Hz. Requiring one track to have $> 600\,\mathrm{MeV}/c$ or one energy deposit to have $> 600\,\mathrm{MeV}$ brings the trigger rate down to 1400 Hz, with only modest loss in efficiency (80% for the above process).

All the L1 trigger logic will implemented in FPGAs, for primitive-formation on the drift chamber and calorimeter readout modules and for tracking and clustering logic in central trigger VME crates. Using track-segment and energy-tower primitives, the central triggers provide counts of tracks and energy clusters above adjustable thresholds to the global trigger.

4 DAQ Readout Modules

Readout modules housed in VME crates serve as the interface between front-end electronics and the rest of the *BABAR* system. They have have two sections: a system-specific ‘personality section’ and a generic readout card section, shown in Figure 3. The readout card contains common interfaces to the fast control, detector control and event assembly systems. Since all detector systems use the same readout cards, the buffer management is uniform. A PowerPC running the VxWorks real time operating system provides processing for feature extraction.

The PowerPC 604 was chosen because of its 64-bit bus, large cache, single-cycle integer and floating-point multiply and high speed. Current PowerPC 604 100 MHz processors rate 160 SpecInt, with a doubling in speed to become available (604E). Deep dual-ported buffers are implemented as Video RAMs. All control and interface logic is housed in a single ATT Orca 2C15–3S304 FPGA. The VxWorks operating system makes code portable between the readout module and readout controller. *BABAR* will use only VxWorks and Unix, which fit with the choice of EPICS for detector control. A board support package for the readout card is being developed.

Detector System	Bytes per Event	Number of Sources	Bytes per Source per Event	Source Data Rates	Number of Rings	Data Rate per Ring
Vertex	8400	5	1680	3.4 Mb/s	5	3.4 Mb/s
Drift Ch.	3600	18	200	0.4 Mb/s	2	3.6 Mb/s
DIRC	1300	1	1300	2.6 Mb/s	1	3.2 Mb/s
Aerogel	300	1	300	0.6 Mb/s	see DIRC	
Calorimeter	7700	18	425	0.9 Mb/s	5	3.1 Mb/s
IFR	800	1	800	1.6 Mb/s	1	3.2 Mb/s
Trigger	800	3	270	0.6 Mb/s	see IFR	
Totals	22900	47			14	

Table 1: Data rates at ten time nominal backgrounds at a trigger rate of 2 kHz.

5 Event Gathering

The *BABAR* data acquisition system assembles events from about 600 readout modules in 45 VME crates from seven detector systems, with rates shown in Table 1, into about ten L3 farm nodes. As shown in Figure 1 of the contribution by G.S. Abrams et al., "*BABAR On-line Prototype Development*" in these Proceedings, the first major phase of event assembly is the gathering of data from the readout modules into a readout controller. The second major phase is the collection of event fragments from the readout controllers into a farm workstation.

The *BABAR* system is based on the use of VME-based PowerPCs as readout controller processors running the VxWorks real time operating system. The data transfer will use the TCP/IP protocol and a switched network. The Technical Design baseline fabric of FDDI was chosen for cost estimates and early tests. Fibre Channel and ATM are also being considered, with a decision by the end of 1997. Major system control issues have not yet been investigated.

We have initiated a program of research and development with field measurements and system simulation. Preliminary VME—FDDI benchmark tests were performed to investigate the use of TCP/IP in a simple topology. These have been performed with an isolated ring defined by a Cisco Workgroup Stack concentrator. The VME source has an Interphase V/5211 FDDI/CDDI interface and a Motorola MVME167 33-MHz CPU board. An Alpha 250 4/266 workstation was used as sink.

Transmission under VxWorks 5.1.1 and using the Sparta driver, involves a single copy of the user data into network buffers managed by the TCP/IP protocol stack. The jagged structure of the throughput curves in Figure 4 is probably due to the CPU time involved in managing these network buffers. An older version of the driver did a second copy for odd-sized buffers, shown for comparison. Under VxWorks 5.2, the single data copy can be eliminated and the network buffers bypassed. Also, we plan to use readout controllers with the PowerPC 604 processor, instead of the 68040 processor, which will also likely increase the crate data rate.

The latency of a model system was investigated with IBM RESQme package, using the *BABAR* rate estimates of Table 1, with multiple rings to reduce cost. The

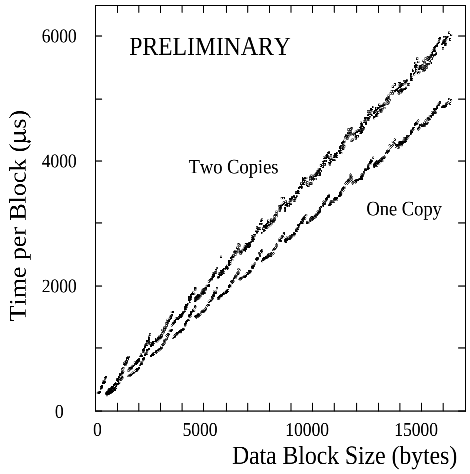

Figure 4: VME throughput with TCP/IP.

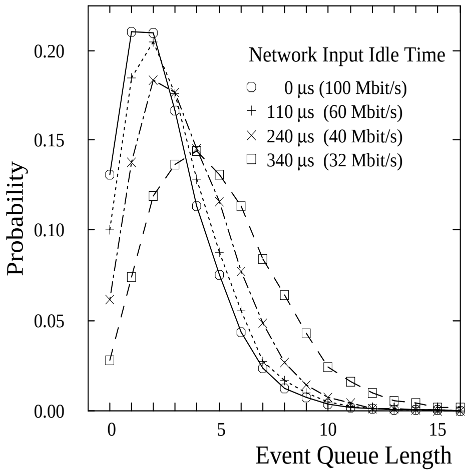

Figure 5: System latency simulation.

FDDI medium runs at a fixed effective rate of 100 Mbit/s. If the crate systems could deliver event data at least that fast, then the maximum effective output from any crate would be limited by the FDDI medium. However, the readout controllers will will inevitably be busy some (currently unknown) amount of time due to formatting and other tasks.

The distributions in Figure 5 show the number of events in the crate-resident event buffers as the data delivery rate of the crates is reduced by overheads. With no overheads, FDDI is adequate using only a moderate number of buffers. When the mean network input idle time grows to the point where only about 30% of the ring capacity can be supplied by a crate, corresponding to the median time between events, the buffer queue begins to grow without bound. This is close to the current measurements. However, we expect improvements from using VxWorks 5.2 and PowerPC processors.

6 Conclusion

While this architecture was designed primarily for the *BABAR* experiment, the adopted strategies may be applicable to many detectors at future high-luminosity machines. Since the recent completion of our technical design, we are focussing on validating the simple two-level trigger architecture that has a maximum trigger rate of 2000 Hz.

Reference

1. The *BABAR* Collaboration, D. Boutigny et al., Technical Design Report, SLAC–R–95–457, March 1995. See also references therein.