

The BaBar Trigger, Readout and Event Gathering System*

P. Sanders, on behalf of the BaBar Collaboration
Blackett Laboratory, Imperial College, London SW7 2AZ, UK

Abstract

This paper describes the trigger, readout and event gathering system being built for the BaBar experiment at SLAC. It describes the hardware platform through which the data passes and the software tasks running on many processors to guide the data through the platform and collect them together.

*Presented at the 10th IEEE Real Time Conference
Beaune, France
22-26 September 1997*

* Work supported in part by Department of Energy contract DE-AC03-76SF00515.

The BaBar Trigger, Readout and Event Gathering System ¹

Peter Sanders

Blackett Laboratory, Imperial College, London, SW7 2AZ, UK
 Representing the BaBar collaboration

Abstract

This paper describes the trigger, readout and event gathering system being built for the BaBar experiment at SLAC. It describes the hardware platform through which the data passes and the software tasks running on many processors to guide the data through the platform and collect them together.

I. INTRODUCTION

The BaBar detector, situated on the PEP-II asymmetric collider at SLAC, will study the decays of B^0 and \bar{B}^0 mesons to search for CP asymmetry. The PEP-II ring collides beams of 9GeV electrons and 3GeV positrons to produce the $\Upsilon(4S)$ resonance, which can decay to B^0 and \bar{B}^0 . The PEP-II design luminosity of $3 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$ produces approximately 100Hz of interesting events.

The BaBar detector [1] is a typical high energy physics detector consisting of a nested collection of subdetectors with cylindrical symmetry around the PEP-II beam line interaction region. The subdetectors consist of a high resolution silicon vertex detector (SVT), a drift chamber (DCH), a ring imaging Cerenkov device (DIRC), a CsI crystal electromagnetic calorimeter (EMC) and an instrumented iron magnetic flux return (IFR).

To identify interesting events recorded in the detector there exists a trigger comprising of two levels. The first level (Level 1) is a hardware processor whose function is to reduce the event rate to 2kHz. The following level (named Level 3) is a software processor which will reduce the event rate to 100Hz.

Between the subdetector front end electronics, the trigger processors (L1 and L3) and the final event logging medium there is a highly parallelised and pipelined readout and event gathering system. This distributed system is a blend of commercially available processors, custom built electronic modules, standard communication fabric and an object-oriented software system.

¹Presented at the Xth IEEE real time conference, Beaune, France, 22-26 September 1997

II. DATA ACQUISITION SYSTEM PLATFORM OVERVIEW

A block diagram of the upstream part of the data acquisition (DAQ) system is shown in figure 1.

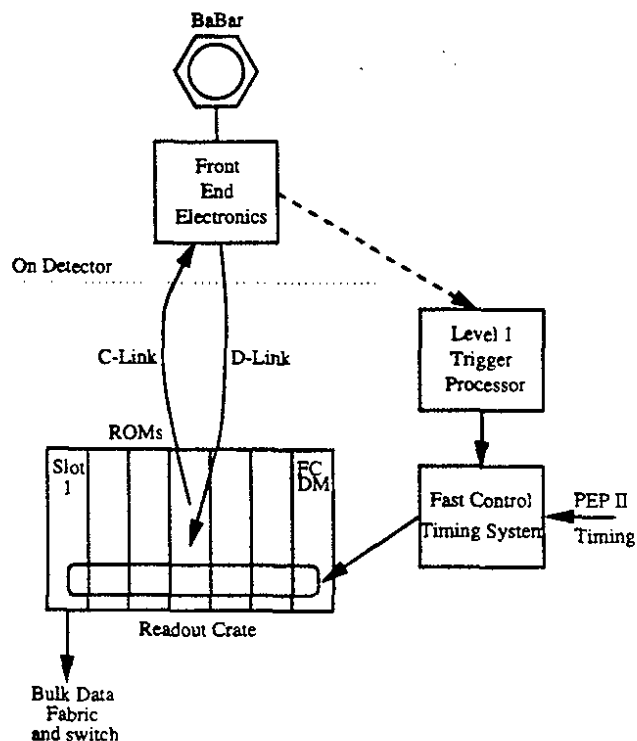


Fig. 1 Block diagram of upstream DAQ system

The system consists of many sets of front end electronics which communicate with readout modules resident in readout crates. The final system will have 150 → 200 readout modules in 20 → 30 readout crates.

The front end electronics from the drift chamber and electromagnetic calorimeter also provide input to the hardware L1 trigger processor. The L1 trigger makes an early trigger decision based on the activity within the detector with a fixed latency of $12\mu\text{s}$.

The L1 decision is sent to the fast control and timing system which distributes any positive result to the readout crate.

Data collected by the readout modules are accumulated by the slot 1 readout module and sent onto the bulk data fabric, where they pass through a switch and arrive at the

online event processing farm. The online event processing farm consists of 20 → 40 UNIX nodes.

III. THE FAST CONTROL AND TIMING SYSTEM

The fast control and timing system (FCTS) creates and distributes common timing and control signals throughout the data acquisition platform.

The FCTS acts as the interface between the PEP-II timing signals and the BaBar DAQ system. It derives the BaBar system clock (60MHz) from the 480MHz RF clock of PEP-II and is capable of synthesising the system clock if required. PEP-II also provides a signal with a frequency proportional to the beam luminosity. This is also distributed through the DAQ system enabling dead time calculations to be performed.

The FCTS receives and distributes the L1 trigger to the BaBar DAQ system. A L1 trigger can occur on any of 24 trigger input lines provided by the L1 trigger processor, which can be masked and scaled. The system also provides the distribution of four auxiliary trigger lines and calibration gates. The system can inhibit L1 triggers if they occur too frequently ($1.5\mu\text{s}$ apart), as the front end electronics will still be BUSY servicing the previous trigger. A 56 bit time stamp, with a granularity of a micro-second, is generated on each trigger to tag every event with a unique time during the ten year lifetime of the experiment.

The trigger tag, consisting of the time stamp and a bit map of the trigger lines, is sent to the DAQ system with the L1 to enable trigger dependent processing.

The DAQ platform can be partitioned into subsystems where the platform is made up of a number of readout crates. The FCTS maintains the partitioning of the DAQ platform by distributing triggers to, and collecting FULL signals from, a set of readout crates within the partition. FULL signals indicate the DAQ pipeline is saturated and the FCTS reacts by throttling the trigger.

IV. THE FRONT END ELECTRONICS

Every subdetector has unique front end electronics (FEE), situated on or within the detector, which digitise the subdetector's response. However all FEE have a common communication fabric and protocol with the DAQ system [3]. This consists of two fibre optic links running at a speed of 1.2Gb/s. The links are $\approx 30\text{m}$ long and comprise of a Hewlett Packard HDMP-1012/1014 (G-link) bit serialisation transmitter/receiver [4] and a Finisar FTR8510 fibre optic transceiver [5].

The links are called the C-link and D-link, their purpose is two-fold:

- To distribute the 60MHz clock to the FEE from which any digitisation clocks can be derived;
- To send commands, both fast commands (trigger, readout, calibration gates) and register read/write commands.

- To receive data, both triggered event data and data from register read requests.

The FEE have several levels of event buffers implemented as FIFOs.

V. THE READOUT CRATE

The readout crate provides an ideal framework for the subdivision of the platform into manageable units and to enable a high level of parallelism. Each readout crate is a 9U VME crate with a custom J3 backplane.

Each crate consists of a fast control distribution module, up to 16 standard readout modules and a slot 1 readout module.

A. The fast control distribution module

The fast control distribution module (FCDM) distributes the signals from the FCTS along the J3 backplane of the crates. These signals consist of the L1 triggers, complete with tag, the system clock and the luminosity signal. The FCDM also returns any FULL signals from the readout modules back to the FCTS.

B. The readout module

The readout module (ROM) is the first level of the DAQ system which is common across all subdetectors. It provides a platform from which applications may control FEE and access event data. It redistributes the FCTS signals and the system clock from the crate J3 backplane to the FEE.

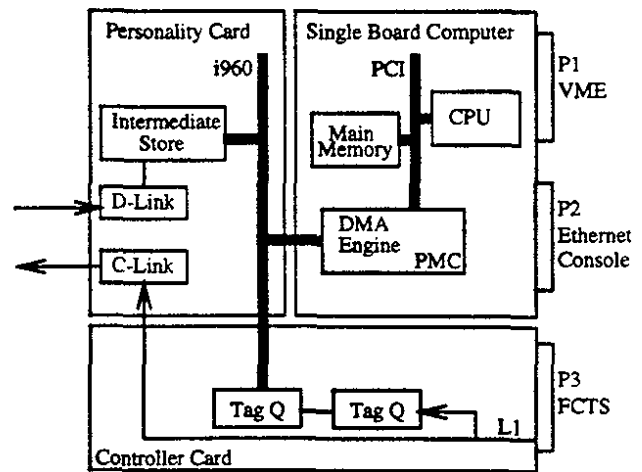


Fig. 2 Block diagram of readout module

A block diagram of the ROM is shown in figure 2. The ROM consists of four separate cards and a front panel held together to form a 9U VME board. It consists of:

- A personality card (PC) housing the C-link transmitter and D-link receiver. Event data received on the D-link are buffered in an intermediate store from which they can be retrieved across the i960;

- Single board computer (SBC) which is a commercially available 6U VME card fitting into the 9U ROM. It contains a PowerPC processor and an external PCI bus slot into which the PMC can be attached, and it supports the VxWorks real time operating system [6]. The main function of the SBC is to receive event data and pass them on down the DAQ chain to the slot 1 ROM. The SBC has the opportunity to process the data, in many cases performing the tasks of feature extraction, sparcification (zero suppression) and reformatting the data. The SBC can itself send commands to the FEE through the PC, enabling FEE setup and testing to be performed;
- A controller card (CC) providing an interface between the ROM and the FCTS signals carried on the J3 backplane. The CC arbitrates the i960 bus and contains a queue to record the tags associated with the data requests from the FEE. If the intermediate store fills the CC will send the FULL signal to the FCTS;
- A PCI mezzanine card (PMC) which forms a bridge between the i960 bus on the PC/CC, and the PCI bus on the single board computer (SBC). The main purpose of the PMC is to use dedicated microcode to move the event data from the PC intermediate store into the SBC.

C. The Slot 1 readout module

The Slot 1 readout module serves as the VMEbus system controller. It is comprised of the same SBC and CC as the standard ROM but has an additional PMC which is an interface between the SBC PCI and the bulk data fabric.

VI. ACTION ON LEVEL 1 TRIGGER

On receiving a L1 trigger from the FCTS the ROM passes this directly onto the FEE through the CC/PC and down the C-link. The FEE stores the event in an event buffer in the front end. Meanwhile the CC puts the trigger tag into a queue to be associated with the event segment later.

After the L1 trigger the CC generates a read event command which causes the event segment to be sent down the D-link into the PC. On completion the data are stored in the intermediate store on the PC and the trigger tag is moved to another queue.

The event segment and associated trigger tag can now be transferred by the PMC DMA engine into main memory on the SBC where the event segment can be processed.

The Slot 1 ROM gathers the processed event segments from the ROMs within the crate to form an event fragment to be sent over the bulk data fabric through the switch to the online event processing farm.

VII. DATAFLOW

Dataflow consists of a set of distributed software processes designed to gather the event together over various media (see table 1). Each node maintains a copy of the BaBar finite state machine which enables it to be configured and attached to a partition within the DAQ system.

When in a running state data are gathered together within the DAQ system. The system is event driven, hence the actual triggering and data acquisition form the stimuli for dataflow.

In the event of the dataflow system becoming overloaded the pipeline backs up to the ROMs where the FULL signal is asserted and the FCTS throttles the L1 trigger.

Table 1
Dataflow hierarchy

Dataflow process	Upstream node	Downstream node	Dataflow fabric
Segment builder	Front end electronics	Readout module	Optical fibre
Fragment builder	Readout module	Slot 1 ROM	VME bus
Event builder	Slot 1 ROM	Online event farm	Bulk data fabric

The software in each node is written almost entirely in C++ with an object oriented methodology.

VIII. THE ONLINE EVENT PROCESSING FARM

The online event processing farm (OEP) is a cluster of processors running the UNIX operating system. The OEP farm receives fully assembled events, where each event is 25kBytes, from dataflow. The farm applies a software L3 trigger to reduce the acquisition rate to 100Hz and stores the events in an object oriented database.

The OEP farm also has the ability to monitor the quality of the data. It has the opportunity to reconstruct some of the events and analyse them to assess the performance of the detector, trigger and DAQ system.

IX. RUN CONTROL

The run control has two major purposes within the context of the DAQ system.

It provides an interface to the BaBar finite state machine. The operator can navigate the various nodes within the platform through the finite state machine to ensure the detector, trigger, electronics and DAQ system are configured and ready for triggers.

It also provides an interface to the FCTS to partition the platform and assign triggers to that partition. The FCTS has the flexibility to partition the system down to

the readout crate level. Once partitioned the FCTS can be programmed to provide certain triggers to that partition and can scale them.

X. CONCLUSION

The BaBar DAQ system described in this paper is presently under construction at SLAC. It will provide a flexible and reliable method of transferring data from the detector to the event store under the guidance of dataflow, imposing minimal deadtime within the system.

PEP-II, the detector and DAQ system will be completed in early 1999 when data will be taken for physics analysis.

XI. ACKNOWLEDGEMENTS

I would like to acknowledge the BaBar trigger, DAQ and electronics groups, whose work is presented in this paper. I thank them for their fine and extensive documentation which has eased the writing of this paper. I would particularly like to thank Mike Huffer who could not attend the conference and for whom I stood in.

XII. REFERENCES

- [1] BaBar collaboration, "BaBar technical design report", SLAC-R-95-457, March 1995.
- [2] R.T.Hamilton, M.E.Huffer, J.L.White, "The dataflow platform users guide", BaBar note 385.
- [3] BaBar DAQ group, "Electronics control and dataflow between the readout module and the front end electronics system", BaBar note 281.
- [4] Hewlett Packard Company. Information available from <http://www.hp.com/>
- [5] Finisar corporation, Mountain View, CA, USA. Information available from <http://www.finisar.com/>
- [6] Wind River Systems Inc. Alameda, CA, USA. Information available from <http://www.wrs.com/>