KTeV Data Acquition System and Online Software

Theo Alexopoulos [d], Ed Blucher [a], Vivian O'Dell [b]
Tsuyoshi Nakaya [a] Taku Yamanaka [c]

[a] *University of Chicago, The Enrico Fermi Institute, 5640 South Ellis Avenue, Chicago, IL 60637, USA*

[b] *Fermilab, PO Box 500, Batavia, IL 60510, USA*

[c] *Osaka University, Physics Department Toyonaka, Osaka 560 Japan*

[d] *University of Wisconsin-Madison, 1150 University Avenue, Madison, WI 53706, USA*

This paper describes the KTeV Data Acquisition system (DAQ), designed for the new KTeV fixed target experiment at Fermilab. The KTeV DAQ is capable of a maximum average throughput rate of 160MB/sec into the level 3 trigger, has 3600 MIPs online for full online reconstruction and Level 3 trigger, and can output to DLT tapes at 18MB/sec.

Key words: data;acquistion;online

1 Introduction

KTeV(FNAL E799/E832), a new fixed target program at Fermilab, is designed to measure the direct CP violation parameter (ϵ'/ϵ) as well as measure rare decays in the neutral kaon system. The KTeV program is a follow up to the two previous Fermilab experiments E731 and E799-I. In E731, (ϵ'/ϵ) was measured to one part in 10^{-3}, and a limit on the branching ratio of 4.3×10^{-10} was placed on the rare kaon decay $K_L \to \pi^0 e^+ e^-$. In this paper, we describe the requirements, design and implementation of the KTeV DAQ[1], which is part of the DART project at Fermilab[1][2].

2 Overview of the DAQ Architecture

KTeV has three levels of trigger. Levels 1 and 2 are hardware triggers using specifically designed electronics. The main components of Level 1 are signals from scintillator counters, a calorimeter total energy sum and a logical OR of drift chamber signals and its design rate is 100 kHz. The main components of Level 2 come from drift chamber hit counting and calorimeter cluster counting and is designed to run at 20 kHz with an average event size of 7.5kB. The Level 3 is a software filter running on three multiprocessing SGI Challenge L deskside servers and is designed to run at 20kHz input and 2kHz output.

The Fermilab accelerator delivers 800 GeV protons to the fixed target program in a 60 spill structure. Each spill protons are delivered in 19 ns contiguous rf buckets for 23 seconds, and no beam is delivered for the remaining 37 seconds in the cycle. In order to use the entire 60 second spill cycle time, data from an entire spill must be buffered until the Level 3 has analyzed it. Given the above design parameters, the DAQ should be able to handle a total data flow of 20 kHz event rate $\times$ 7.5 kB event size or 150 MB/sec and buffer a total of 150 MB/sec $\times$ 23 seconds or 3.45 GB. The level 3 filter is designed to save only 10% of the data coming in , and thus the DAQ must be able to write out saved events at an average rate of $\frac{3450\text{MB}}{60\text{seconds}} \times 10\%$ or 5.8 MB/sec. Figure 1 shows a

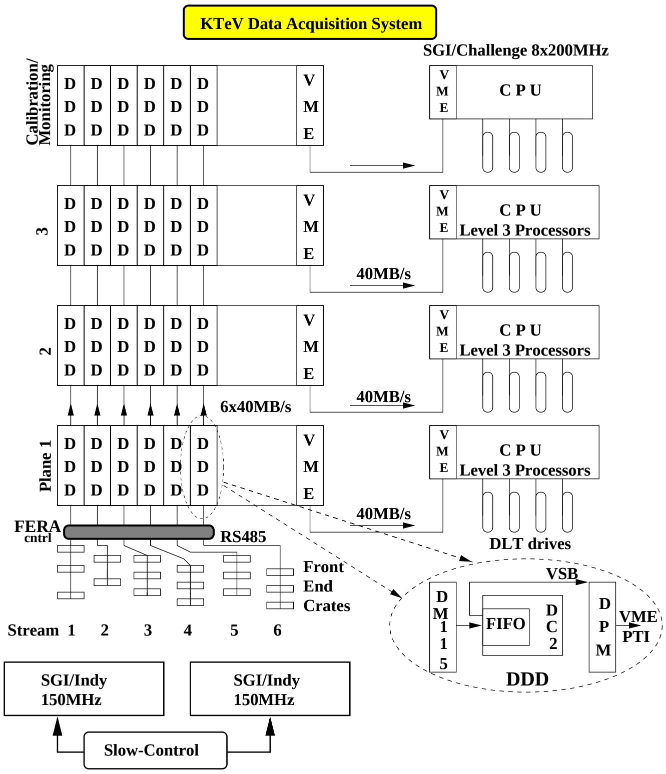

Fig. 1. Block diagram of the KTeV DAQ.

block diagram of the KTeV DAQ. Data from sub-detector systems flow into the memory nodes along one of 6 “streams” of RS485 cables, with a maximum bandwidth of 40 MB/sec/stream. Thus the maximum total data flow along the streams is 240 MB/sec. The data from each stream are buffered in dual ported VME/VSB memory while being shipped to a Challenge L deskside server via a VME to VME interface. There are a total of 4 sets (or “planes”) of VME crate/Challenge machine pairs, giving a total of 24 memory nodes (4.62 GB total memory) in the DAQ matrix. The RS485 cables into each plane are daisy chained, allowing data on the cable to be sent to any specific plane or combination of planes.

Figure 1 also shows a block diagram of one stream on one plane of the KTeV DAQ. There are many different types of front end crates and modules in the KTeV experiment, but each crate is read out over a FERA bus by a FERA readout module designed and build by Fermilab (DYC3 or CTIRC). On receipt of a level 2 trigger, the data from each crate is read into the FIFO of this module, and the data from all crates in a stream are pushed serially onto the RS485 cable. Each RS485 stream is daisy chained to a memory node on each plane, connecting four nodes in total. Each memory node consists of 3 VME/VSB dual ported memory modules, a board containing a DM115 receiver device and a DC2 controller device.

Data flow up the RS485 cables to the memory node, consisting of a DM115/DC2 pair, a triple width VSB bus, and 3 VSB/VME Dual Ported memories. The combination DM115/DC2 is a commercial board, made by Access Dynamics, which plugs into the VSB bus at the back of the VME crate. The DM115 accepts RS485 data and pushes them into an 8 kb FIFO. The DC2 does a data DMA over the VSB bus from the FIFO into the dual ported memory, made by Micro Memory, Inc. In addition to controlling the DMA, the DC2 attaches the subevent size to the top of each event, makes a table of event VME address in the memory and keeps a counter of the total number of subevents processes. The slowest step in this process is the DMA step and access to the memory, which goes at about 22MB/sec. Because of this limitation, in order to push each stream up to its maximum of 40 MB/sec, it is important to be able to switch planes on an event by event basis. By doing this, each DC2 sees only one out of three events and thus the maximum aggregate rate per stream is limited only by the RS485 rate.

The data are pulled from the VME crates into the SGI Challenges over a VME to VME interface (PTI940). In order to achieve high VME DMA rates, subevents are blocked in groups of roughly 1000 (“super-events”) for pulling into the Challenge. The six super-events, one from each stream, are then sent to one of 8 Level 3 filtering jobs running in parallel. The individual events are fully reconstructed and selected for offline analysis according to the online level 3 criteria. The computing power for this filtering corresponds to about 3600 MIPs with three SGI Challenges. The fourth plane (VME crate + SGI Challenge) is used for processing calibration events and for detailed online monitoring of a percentage of the physics data.

After the level 3 selects an event, it builds the event (*i.e.* copies it to contiguous memory space) and sends it to a separate process which handles the tape media I/O. The built event is then written to a specific tape drive depending on trigger type.

The KTeV DAQ is commissioned and started taking data in July 1996. Because of very clean kaon beams in our new secondary beam line, the actual rates during E832 running are about 35 kHz L1 (compared to the design rate of 100 kHz), 10 kHz L2 (design rate 20 kHz) and 1.7 kHz to tape (design rate 2 kHz). The average event size going into level 3 is about 6 kB, giving a total data flow into level 3 of about 60 MB/sec. The DAQ has been tested at 150 MB/sec, and this is needed for the E799.

3 Online Software

The level 3 processes analyze events in parallel. Histograms of relevant parameters for the filtering are made independently on each CPU with HBOOK, and stored in shared memory. A separate process dumps all the contents of the shared memory, plus a sum off all histograms in shared memory, to a file on an NFS mounted disk. Thus monitoring of the histograms can be done on a CPU by CPU basis, plane by plane basis or summed over all planes, to check for inefficiencies in any one process or global problems in detectors. In addition to physics parameters, DAQ parameters such as CPU usage per process per spill and event sizes per crate or stream are also histogrammed.

In addition the DAQ we use a Slow-Control system which is responsible for the reliability and safety of the KTeV detector.

4 Conclusions

KTeV(E832) has written over one billion events to tape. The DAQ has been tested at design rates and is running smoothly at about 1/2 the design parameters during E832 operations.

References

[1] T. Nakaya, et al., *Proceedings for the 1995 Computing in High Energy Physics* (1995).

[2] G. Oleynik, et al., *Proceedings for the 1995 Computing in High Energy Physics*