The DØ Upgrade

The DØ Collaboration

xx March 1996

Contents

1 Introduction

The DØ Upgrade was endorsed by the PAC in April 1995, and a baseline cost established by a Director's Review. The scope of the upgrade remains unchanged since then. In this document we present information on three of the subsystems of the upgrade: the forward preshower, the Level 2 trigger and the muon system. Our understanding of these systems has developed considerably over the last year. Regrettably this has led to increases in the baseline cost estimate.

It has always been realized that a Level 2 trigger system would be essential to maintain adequate rejection and achieve reasonable trigger rates in the upgrade. Over the last year the Level 2 system has benefited from simulations of the expected trigger rates including multiple interactions at high luminosities. We have concluded that significant computational power will be needed at Level 2 to meet our bandwidth goals.

This document summarizes the information presented to the PAC in the document of January 1996, and addresses the concerns raised at the January 1996 PAC meeting. We therefore seek the endorsement of the PAC for these three subsystems, so that the review process may be brought to a conclusion and we may embark on the challenging but exciting task of preparing the D0 Upgrade detector to address the physics questions of the next century.

2 L2 Trigger for the upgraded DØ detector

2.1 Introduction

The DØ upgrade trigger system extends upon the existing system in place for Run I, and was described in our April 17, 1995 submission to the PAC. The existing Run I trigger yields a multi-stage set of decisions. Level 0 is an inclusive interaction trigger based on scintillator hodoscopes in the forward directions. Level 1 is a deadtimeless trigger based on calorimeter and muon PDT inputs operating in about 3 μs. Increased power is provided by the so-called Level 1.5 trigger system, based on processors operating in the 100 μs time frame, which refines muon information and allows more flexible construction of electron and jet objects. The final stage of triggering is done in the L2 board-level processors in which a fully assembled event can be analyzed with variants of offline algorithms.

The upgrade trigger system proposed originally in outline in 1990, and approved by the Laboratory in 1991, simply extends the power of these systems to meet the increased event rates and shorter bunch crossing times – and rationalizes the naming convention, with the old Level 1.5 becoming Level 2 (L2), and the old Level 2 becoming Level 3 (L3). The upgrade Level 1 (L1) systems described and approved previously operate in a few μs upon the trigger signals from single detector systems – calorimeter, fiber tracking/central preshower, forward preshower and muon chamber/pixel detectors. The L2 system described here makes decisions on the scale of a few hundred μs, typically combining information from several detectors to give refined information about physics objects present in the event. The L3 system remains a software trigger based on sophisticated algorithms operating on the full event.

The specifications for the input rates at each level of the upgrade triggering are: 10KHz out of L1: $\sim$ 1000 Hz out of L2, and $\sim$ 10 Hz out of L3. The deadtime introduced by the L2 system is less than 1%, while providing a factor of 10 rejection. The combined effect of the trigger system is to reduce the raw inelastic collision rate by a factor of about 10^6 to save the interesting physics candidate events.

The L2 system described below obtains its rejection power by considering the correlations between trigger information from different detectors. We expect the specific correlation algorithms to evolve as experience is gained and new physics is added to the menu. The generic algorithms are those of matching spatial information in different detectors, and in computing correlations (mass, etc.) for multiple objects. Specific examples envisioned include: calorimeter – tracker matches; preshower – calorimeter matches, determination of jet parameters with reasonable jet cone definitions, determination of E/p, and rough mass calculations for dileptons, di-jets, lepton–jet, etc.. We require the L2 trigger to have a high degree of flexibility for implementing new algorithms that are necessary to address new physics. In Run I, this flexibility paid off in many instances – new triggers for rapidity gap events, for specific η-region muons and di-muons, dijet triggers were all introduced rapidly and efficiently as the physics concerns dictated. It was also extremely useful in coping with the large range of instantaneous luminosities in Run I.

Our understanding of the rejection power afforded by the cross-detector algorithms operating in L2 has been developed through simulations and experience with Run I data. These rejection factors have been presented in earlier submissions to the PAC (April 1995) and give good understanding that the overall bandwidth goals are achievable in the face of the high rates and multiple interaction environment. The simulations are based on our understanding of the character of the signals available to the L2 system from L1 or the detector trigger outputs.

Since the last PAC submission, the understanding of the trigger system has been solidified with detailed timing simulations of the overall trigger system, which have led to the incorporation of L2 preprocessors which prepare lists of compacted information for transmission to the event processors. This feature is crucial in allowing affordable data transfer rates. We have adopted the design principle that generic choices for preprocessor and global processor hardware and software should be made to avoid expensive one-of-a-kind design and fabrication.

2.2 Requirements for the L2 Trigger

2.2.1 Outline of the L2 system

Many different arrangements can be thought of for the various components in the L2 Trigger and we have considered several of them. The case of a processor farm has been studied in some detail and based on those studies we have arrived at the current design. This system will consist of a set of preprocessors, which are detector specific and prepare the information to be sent to global processors. The global processors each work on one event at a time. An overview of this system is shown in Figure 1. In the following we describe the system in more detail.

The principal choice made in the L2 Trigger design was to not directly work with the raw detector data but rather to work only with "concentrated" information from the detectors. This choice is dictated by the very limited time available for the L2 system to reach a decision. The concentrated information is available from either the L1 Trigger system for a detector or else from the detector front-end (FE) through a special preprocessor. As shown in Fig. 1, each major detector system: muon, calorimeter, tracking (CFT), preshower and silicon vertex (Si) feeds data into a specific preprocessor. The result of each preprocessor is a list of candidate objects from that specific detector: muons, electrons, jets, tracks and secondary vertices. The number of preprocessors is limited and is determined by the number of detector systems.

2.2.2 Timing studies of the L2 system

A major requirement of the L2 system is to introduce a minimal amount of deadtime ($<$1%). To verify this we have simulated the deadtime in the L2 trigger using a modeling package based on queuing theory. Many different simulations have been performed and here we describe a representative one

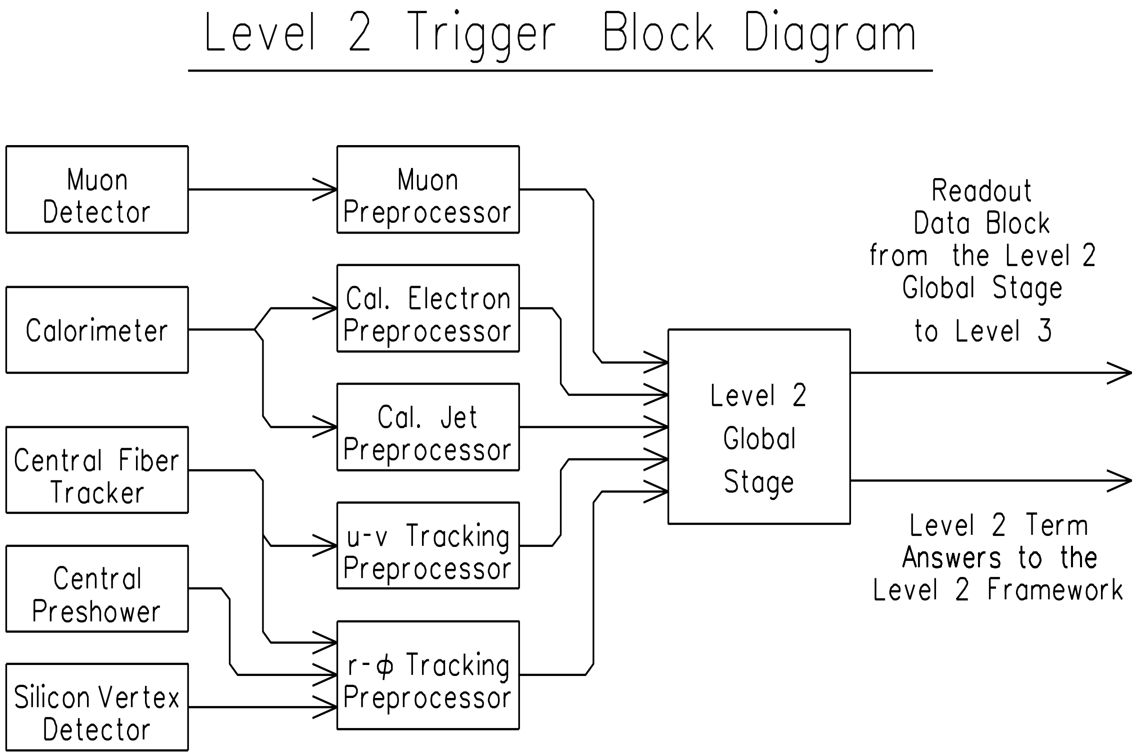

Figure 1: Overview of the L2 trigger components

with realistic choices of parameters. In this simulation we have fixed the preprocessor times for different detectors as shown in Fig. 2; these include a cutoff on the processing time. For each distribution shown, the first time listed is the mean processing time for an exponential time distribution, and the second one is the data latency. The exponential time distribution was chosen because it has a moderately long tail with about 5% of the cumulative distribution above a cutoff of 3 times the mean. This is fairly conservative, and to maintain our conservative stance we charge the data latency against the total time of the preprocessors. By fixing the time of the preprocessors we can explore the requirements on the global processor. For all of the studies we will describe below, we have assumed a L1 accept rate of 10 KHz and a single global processor.

In Fig. 3 we show results from some of our studies. The solid and dashed curves show the effect of the tail on the processing time of the preprocessors. For reference, the deadtime for a mean global processing time of 80μs is $\sim 3.7\%$ if we apply the cutoff on the preprocessor time and $\sim 6.4\%$ if we do not apply a cutoff. As can also be seen from this plot, the two curves are converging as the mean global processing time increases. This is to be expected as increasing the mean processing time in the global processor begins to overshadow the time spent in the preprocessor. The apparent dip in the curves around 60 μs is an artifact of a simple spline fit to the data points. Another parameter we have investigated is the effect of the transfer of the output from the preprocessors to the global processor. We have assumed the

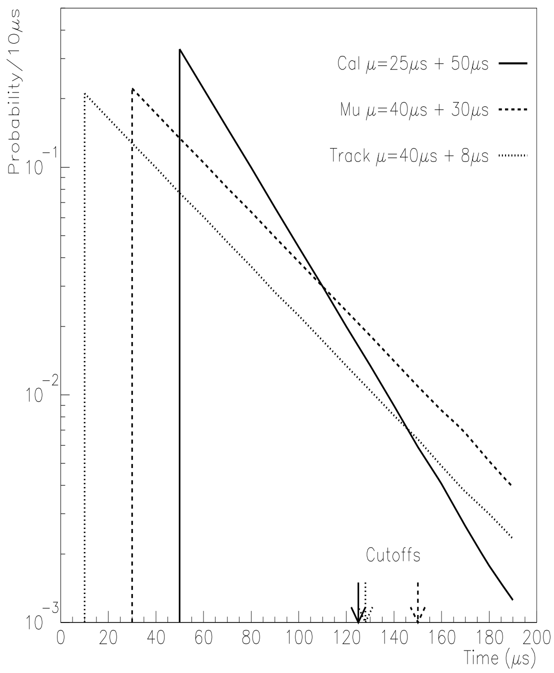

Figure 2: Distribution of the processing time/event in different preprocessors: muon (Mu), calorimeter (Cal) and tracking(Track).

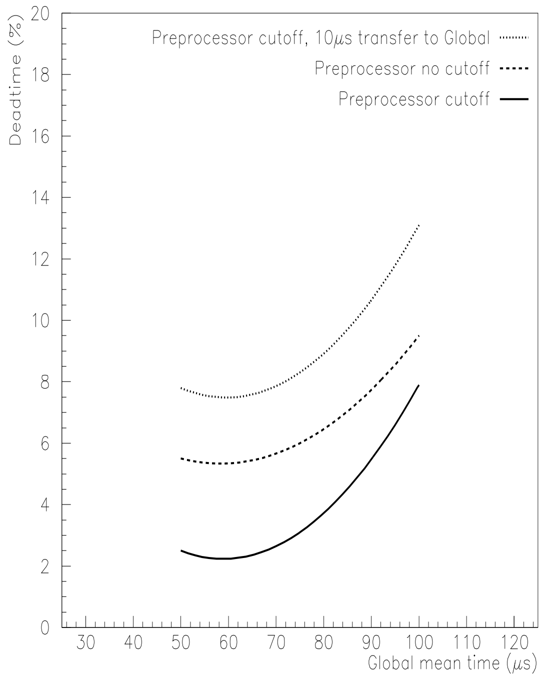

Figure 3: Deadtime as a function of processing time/event in the global processor for different assumptions.

memory for this transfer is local to the global processor and, thus, charge the transfer time against the preprocessors. The dotted curve in Fig. 3 illustrates this effect for an assumed transfer time of 10 μs. At a mean global processing time of 80 μs we find a deadtime of $\sim 8.8\%$ which can be compared directly with the $\sim 3.7\%$ from the solid curve. Again, these two curves are converging as the mean global processing time increases.

In addition to the preceding studies, we have also investigated the effects of processing time correlations between the preprocessor and the global processor. These correlations arise from events that produce many candidates in the preprocessor. For a mean global processing time of 80 μs we find a deadtime of $\sim 5\%$ compared with $\sim 3.7\%$. This effect now grows with mean global processing time so that the curves diverge.

Simulations like this have provided guidelines for the processing time available in the preprocessor and global processor to complete an event. Typically this time is of order 30-40 μs in the preprocessor and about 70 μs in the global processor.

2.2.3 Preprocessor tasks and time estimates

2.2.4 Muon preprocessor

The muon preprocessor is designed to find objects in the muon detector using the digitized data for all the muon detectors, and report segments of muon tracks to the global L2 processor. The muon preprocessor will divide the muon detector into geographically separate pieces consisting of 45 degree segments in ϕ and in either the central or one forward region of the detector. The muon preprocessor is anticipated to require up to 20 μs for serial input from the detector, 20 μs for internal data transfer and buffering, 30 μs for calculations, and 20 μs for serial output to the global L2 processor.

In the central region the muon preprocessor input will be the drift time information from the wide-angle muon system drift tubes (WAMUS) and the time information from the central muon scintillator (CMUSC) and A-ϕ scintillator. The preprocessor will consider groupings of drift times consistent with L1 muon trigger centroid in the WAMUS chambers and then use the drift times to calculate a position in the z-coordinate and an angle in $r - z$ for a straight line segment. For good segments the drift time information will be used to extract a delta-time measure of the position along the wire which will give a coordinate in the ϕ-direction and an angle in r-ϕ. The segment coordinates will point at an element of scintillator, and the segment will be tagged with the time for that element of scintillator, including any correction to the scintillator time for time-of-flight. Separate segments will be found inside and outside the iron toroid magnet.

Studies of algorithms have begun for the central inner segments. Samples of data from run I were scanned and separated into good and bad muons. Events were required to have a L1 trigger centroid before the impact of the L2 preprocessor algorithm was determined. For 95 % efficiency on the good muons, a rejection factor of 7.8 on the background was measured. A study of scintillator timing found that with 99 % efficiency a rejection of 5 was possible.

In the forward region the input to the muon preprocessor will be the forward muon system drift tubes (FAMUS) and the forward muon scintillator (FMUSC). The FAMUS data will come from the L1 input centroids from the MCEN cards and will not include the full drift time, but will include a pattern reference as well as position. The preprocessor will convert the MCEN output into an x- or y-coordinate from the beam and an angle between that coordinate and the z-direction. There is no wire information in the other coordinate, so that coordinate will be found using the position of the FMUSC element with the time most consistent with the interaction. The matching scintillator time(s) will be tagged to the segment with any corrections. As in the central region, separate segments will be found inside and outside the toroid.

The output of the muon preprocessor in all regions is a list of segments from all regions including the position and angle of the segment, quality of fit and scintillator time(s). By providing individual segments the global processor will be able to have the maximum efficiency for keeping muons that may miss either an inner or outer layer while having sufficient information in the calorimeter or fiber tracker to identify it as a good muon.

2.2.5 Calorimeter preprocessor

The L1 calorimeter trigger is based on the transverse energy deposition in trigger towers covering 0.2×0.2 in $\Delta\eta \times \Delta\phi$. The minimum transverse energy(E_T) required in a trigger tower can be set as can the number of trigger towers above the threshold. Electromagnetic objects can be triggered on by putting a threshold on the minimum amount of transverse energy deposited in the electromagnetic portion of a trigger tower. This trigger is implemented out to $|\eta| \leq 4.0$, giving a total of 1280 trigger towers.

L2 will need to trigger on both electromagnetic (EM) objects and on hadronic (HAD) objects. The goal for the L2 preprocessors is to send object information to the global L2 processor within 100 μs. Given these time constraints on the L2 preprocessor, we will use two calorimeter preprocessors, one for electrons/photons and one for jets. Each preprocessor will form its decision based on the E_T in the 1280 trigger towers. In addition, the L1 decision, including which trigger towers (*seed towers*) passed the specific L1 E_T threshold for each trigger, will be available to the L2 calorimeter preprocessors.

Since the L1 trigger information is available to the L2 preprocessors, we can save time by running the algorithm only once over each seed, using the lowest common thresholds of all triggers which passed at L1. The resulting objects (electrons, photons, jets) passing these "least common denominator" requirements will then be passed to the global L2 processor which will make decisions for each specific trigger based on the thresholds relevant to that trigger.

Finding objects in the L2 calorimeter preprocessors consists of several tasks, each of which will take a certain amount of time. The tasks include (i) getting the trigger tower data into the preprocessors (~ 13 μs), (ii) unpacking the L1 trigger information to get the η and ϕ position of the seed towers (~ 3 μs), (iii) running the object algorithm and (iv) creating the block of output data to be sent to the global L2 processor ($\sim 10\ \mu$s).

At the beginning of Run II, the Tevatron will operate with beam crossings spaced 396 ns apart. With this spacing of the beam crossings, the amplitude of the calorimeter trigger tower signals will be proportional to the energy deposited in the calorimeter from a single beam crossing. Later in Run II, when the Tevatron begins operating with beam crossings separated by only 132 ns, there is the possibility that it will take a few beam crossings periods for the trigger tower signals to reach their full response due to the energy deposited in the calorimeter by a single beam crossing. In this case, we anticipate that the L2 calorimeter preprocessor will need to read in and manipulate trigger tower information from several beam crossings in order to extract the signal from the beam crossing that actually caused the L1 trigger. We make the assumption that we will need to use information from four beam crossings around the beam crossing that caused the L1 trigger. Therefore, the time required by the L2 calorimeter preprocessor to complete certain of its event processing steps will increase by a factor of four.

The timing studies for the calorimeter preprocessor algorithms were done in the C programming language running on a Digital 3000/M300 ALPHA machine ($\sim$ 100 MIP's). Algorithms for finding electrons and jets similar to the ones used in Run I were coded and used to determine the processing times required. The total times for the calorimeter preprocessors are summarized in Table 1. The data input time shows a range due to four beam crossings assumed needed at 132 ns. The algorithm time, which is described in detail below, is shown for an average event and the range again indicates the time assuming four beam crossings are needed. The expected preprocessor times are comfortably within our time budget for 396 ns beam crossings. At 132 ns, we will be pushing the system to its limit and will need all the time improvement we can get by using the fastest processor and fastest data input available.

	Electrons	Jets
Data Input	$13 - 50\ \mu$s	$13 - 50\ \mu$s
Unpacking seed info	$3\ \mu$s	$3\ \mu$s
Mean algorithm/event	$5 - 20\ \mu$s	$15 - 60\ \mu$s
Output data creation	$10\ \mu$s	$10\ \mu$s
Totals	$30 - 83\ \mu$s	$41 - 123\ \mu$s

Table 1: Calorimeter preprocessor time estimates.

Calorimeter: Electron/Photon Algorithm

The electron (and photon) algorithm consists of the following steps:

- Form a 1×2 (trigger towers) cluster from each seed using its highest E_T neighbor and compare this sum to the lowest threshold needed by the electron or photon triggers which passed L1.
- Compare the EM fraction of the 1×2 cluster to the lowest common EM fraction threshold. The EM fraction is defined as the ratio of the cluster EM E_T to the cluster total E_T (EM+HAD).
- Compare the isolation fraction of the 1×2 cluster to the lowest common isolation fraction threshold. The isolation fraction is defined as the ratio of the cluster EM E_T to the total E_T sum of a 3×3 cluster.
- Compare the number of EM clusters passing the above requirements to the minimum number of EM clusters required.

As can be expected, the algorithm time is highly correlated with the number of seeds. Since the latter steps in the algorithm are not run unless the EM cluster passes each previous step, there is a characteristic spread in the time distribution. For events with a single seed, the time distribution to run the electron algorithm for typical EM requirements is given in Fig. 4. The requirements were: 7 GeV EM E_T seed at L1, 10 GeV 1×2 cluster E_T, EM fraction of 0.80, and isolation fraction of 0.01. The average time is ~ 2.8 μs. About 1.0 μs of that time is for algorithm overhead. Each additional seed adds ~ 2.1 μs to the average time and EM seed multiplicity (for 7 GeV seeds) is on the order of $1-2$ seeds per event. The seed multiplicity, and hence algorithm time, increases slowly with luminosity.

Calorimeter: Jet Algorithm

The jet algorithm consists of the following steps:

- Form a 5×5 (trigger towers) cluster around each seed tower and compare this E_T sum to the lowest common jet E_T threshold needed by the jet triggers which passed L1.
- Order the clusters in E_T and check for any overlapping clusters (where overlapping is defined as having the center tower in the cluster within another cluster), keeping the higher E_T cluster as a jet.
- Compare the number of jets to the minimum number of jets required.

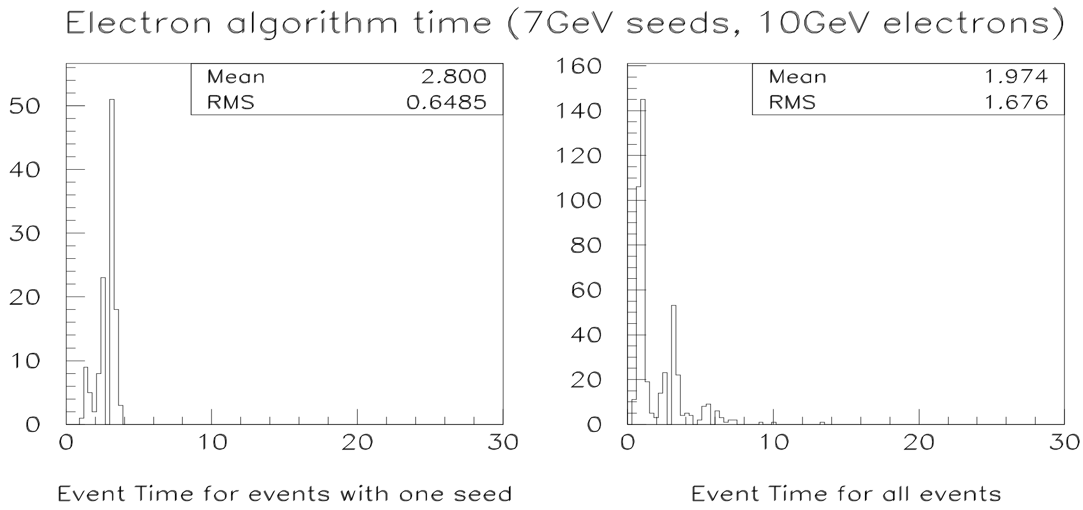

Figure 4: Event time for L2 Calorimeter preprocessor for electrons with a single (7GeV) seed in the event and for all events.

As with the electron algorithm, the jet algorithm time is very dependent on the number of seeds. However, as the number of seeds increases, so does the number of clusters above the jet E_T threshold which then need to be checked for overlaps. Hence, there is also a dependence on the number of clusters. The average algorithm time for an event with a single 3 GeV seed and a cluster E_T threshold of 15 GeV is comparable to the electron algorithm, ~ 3.0 μs. Each additional seed adds ~ 2.2 μs to the average time with another ~ 1.1 μs from the increase in cluster multiplicity. The average seed multiplicity for 3 GeV seeds and low luminosity is ~ 3, growing to ~ 6.0 at high luminosities. Seed multiplicity for jets depends more on luminosity than in the electron case. The lower seed threshold, which applies to total trigger tower E_T (EM+HAD), is more easily reached due to the additional energy in the calorimeter from the multiple minimum bias events underlying the hard interaction. Hence, the total event time for jets is larger and the tail is much broader than for electrons.

2.2.6 Tracking preprocessor

The L1 central fiber tracking (CFT) trigger allows up to 6 track candidates in each of the 80 ϕ sectors of the tracking system. This multiplicity of candidates allows a minimum p_T threshold of 1.5 GeV up to the highest luminosities with high efficiency. For each track candidate the addresses of the axial fibers($r-\phi$) with u-v fiber (4 layers) information will be sent to the L2 tracking preprocessor for each L1 trigger. The u, v layers are at 2 degrees stereo. In addition the preprocessor will receive the p_T thresholds used in the L1 trigger for individual events and will use this information to determine which subset of L1 candidates to process.

For the track candidates surviving these criteria, the information from the u-v fibers crossing the selected axial fibers, will be requested from the front ends. Combining this information with the L1 information from the axial fibers allows to form three dimensional tracks in r, ϕ and z. Simulations of such matching algorithms have shown a reduction of fake tracks of more than a factor of 8. In a feasibility study it has been shown that an algorithm like this, implemented on a digital signal processor (DSP), can be used and results in the above quoted rejection.

The candidates passing such a matching track algorithm will be transferred to the global stage of the L2 trigger. The information sent will consist of η, ϕ and Z vertex for each resulting track candidate and the candidates would be ordered in p_T.

2.2.7 Silicon vertex preprocessor

The feasibility of a secondary vertex trigger(SVT) based on information from the silicon vertex detector and tracking system is being studied actively. At the time of the writing of this document DØ has not yet decided whether the tracking system of the upgraded detector and the possible physics potential require a secondary vertex trigger or not. This is actively being investigated by a group dedicated to this problem and we expect to resolve this question by the time of the next PAC meeting.

In the meantime the design of the tracking trigger is progressing, keeping in mind that it may include a possible SVT in the future. The hardware implementation of the L2 tracking trigger and L2 secondary vertex trigger is expected to be very similar. The processors in each case are likely to be identical or at least very similar. So at this point in time the design of triggers and DAQ is such that all the hooks and connections for a possible SVT are in place. However given the current state of funding we do not envision to start Run II with such a trigger, but feel that it may be a candidate for a staged upgrade after the start of run II.

2.2.8 Global processor tasks and time estimates

The role of the global processor is to combine the object information from the various detector preprocessors to evaluate the L2 terms for each L1 trigger which fired. The goal is to perform those evaluations within a time budget of 100 μs. The general methodology is that the global processor will build *scripts* describing the tasks it needs to perform from the L1/L2 trigger information. At its simplest, a script may consist of the number of each type of object needed (jet, μ, e, γ, etc); the object definition which may require combining information from multiple preprocessors; and the object E_T, η, or ϕ requirements. From our experience in Run I, we can envision quite complex scripts to pick out interesting events. After building the scripts for the fired triggers, the global processor will evaluate each script for each relevant object or set of objects found by the preprocessors to form its decision.

The global processor will therefore have to perform a series of wide-ranging tasks. These tasks include: (i) getting the preprocessor data into the global processor; (ii) event consistency checking; (iii) scriptbuilding; (iv) manipulation of preprocessor information to complete tasks preprocessors are unable to do (e.g., matching muon track segments from different parts of the muon detectors to find muons); (v) script evaluation; and (vi) reporting result to the L2 framework and creating the block of output data to be sent to Level 3.

The average data expected per event from all preprocessors has been estimated to be 528 bytes for data packed into 4 byte words. The expected data through-put rate for the proposed global L2 processor (see following discussion on hardware) is 320 Mbytes/s. Thus the data input is expected to take ~ 1.7 μs/event.

Preliminary timing studies, again on a 3000/M300 ALPHA processor, were done for the script evaluation time. These studies indicate that script evaluation takes $\sim 0.5 - 1.5$ μs per object per trigger. Given object multiplicities on the order of $2 - 3$ per event and an estimate of 5 triggers per event, implies ~ 15 μs for complete script evaluation per event.

The timing is not fully understood for most of the other global processor tasks. It is also highly likely the list of tasks for the global processor will grow both during the upgrade preparations and during Run II. Therefore, the most critical requirements for the global processor are flexibility and speed.

2.3 Proposed hardware

2.3.1 Muon preprocessors

The muon preprocessor consists of two VME crates: one interface crate and one preprocessor crate. The interface crate includes the serial link interface cards and an interface control card. The preprocessor crate includes parallel processor cards, a processor control card, and a VME Buffer-Driver (VBD) card for data transmission to Level 3.

Each serial link interface card (SLIC) includes up to 21 receivers for the 160 Mbit/s serial links used by the muon front ends. The SLIC will have sufficient memory to buffer 8 input events and then transmits the bytes to the parallel processors at 40 Mbytes/s. The buffer control and transmission to the parallel processors is controlled by the interface control card over the user-defined lines of the VME crate.

The parallel preprocessor currently considered is the CNAPS-128 by Adaptive Solutions. It is available as a 6-U VME module and has a 40 MByte/s direct I/O link and 128 parallel processors. The 128 processors are arranged in a single-instruction multiple-data (SIMD) architecture and can handle 20 million multiply-accumulate operations per second. The processor control card would be required to read the results from the CNAPS-128 cards and also includes the connections to the trigger framework, global L2 trigger, and L1 muon trigger.

Each CNAPS-128 requires one SLIC for the input, and to utilize the SIMD architecture each processor would operate on a narrow road within one region. Separate processors would handle each of the four possible segment finding algorithms. The entire muon detector requires 2 central inner, 4 central outer, 2 forward inner, and 4 forward outer SLIC - CNAPS pairs. Other than the CNAPS-128 and VBD all other cards need to be made by DØ.

2.3.2 Calorimeter Preprocessors

The calorimeter L2 preprocessor will be implemented using a RISC ALPHA based trigger processor board. This method of implementation was chosen for a number of reasons. It is fast enough to complete the L2 preprocessor algorithms within the required time budget. It retains much more flexibility in allowing a wide selection of triggering algorithms than a hardwired logic implementation would. The processor based implementation costs less than a L2 preprocessor based solely on custom built specialized hardware.

The ALPHA based trigger processor board that we plan to use is being developed jointly by CDF and DØ. This development is a continuation of the ALPHA based trigger processor that CDF began using during collider Run I. The principal feature of this trigger processor design, that makes it useful in our applications, is its high speed input data path that can feed the same information to multiple processor boards from multiple data sources. This

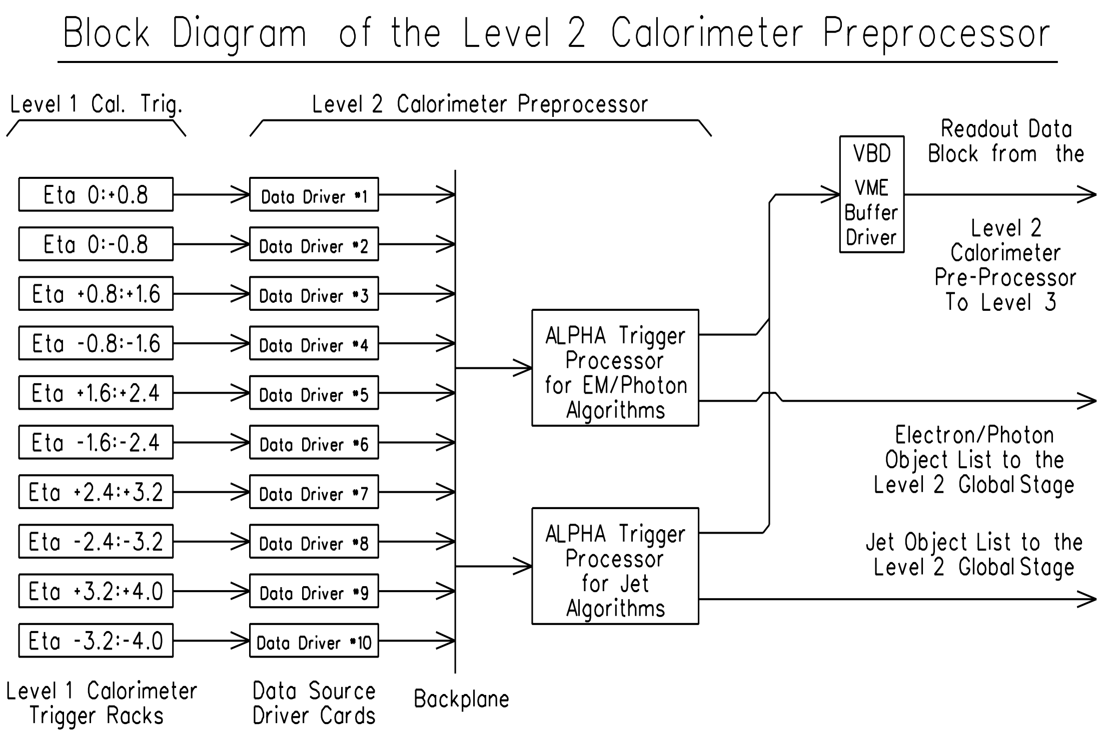

Figure 5: Input and output of the calorimeter L2 preprocessors.

feature is what is missing from most commercial processor board designs.

The input information to the preprocessor comes from the L1 calorimeter trigger. This information consists of electromagnetic and hadronic trigger tower signals (1280 of each) and masks showing which trigger towers were over the various L1 thresholds (4 EM thresholds and 4 Jet thresholds). Each of the 10 racks in the L1 calorimeter trigger delivers its section of this information to the L2 preprocessor over a separate data path. This is illustrated in Figure 5. By using these 10 data paths working in parallel, all of the input information is delivered to the preprocessor in about 12 μs if only a single beam crossings worth of calorimeter data is needed. At the highest beam crossing rates where the calorimeter trigger tower signals may not respond fast enough to separate the effects of individual beam crossings, then the preprocessor may need up to 4 beam crossings worth of information. In this case the 10 data paths from the L1 calorimeter trigger would take about 50 μs to deliver this data to the preprocessor. The design of the data path between the L1 calorimeter trigger and the L2 preprocessor will be similar to the equipment used in Run I to feed data to the Level 1.5 calorimeter trigger.

When the data arrives at the preprocessor it is distributed via the processor crates custom backplane to a number of the ALPHA trigger processor boards. This distribution mechanism is flexible and it will be used to simultaneously deliver the input data to two processor boards. One processor board will use this data with the Electron/Photon algorithm and the other processor board will use this data with the Jet algorithm.

When input data arrives in the preprocessor crate, it is driven onto the crate's backplane by data source driver cards. There can be up to 15 of these data source driver cards in a crate. Memory on each of the ALPHA trigger processor boards receives this data. Up to 7 trigger processor boards may be installed in a crate. Any combination of these trigger processor boards may be enabled to receive the calorimeter data from a particular event. The memory on the trigger processor boards can be organized as an input event FIFO. The information in this memory is directly accessible by the ALPHA processor on each of the trigger processor boards and it is also visible over VME for readout or diagnostic purposes.

There are two outputs from the preprocessor. One output is the Lists of Identified Objects (Electron/Photon and Jet) that the preprocessor sends to the L2 Global Stage. The other output is a readout to the Level 3 DAQ system. The readout to Level 3 DAQ uses the same setup and equipment as is used by all of the other data sources to the DAQ system (e.g. detector front-end crates).

2.3.3 Tracking preprocessors

The information from all the central tracking front end systems is transmitted to the counting house over fast optical serial links. A passive optical splitter near the destination will split each link into two and the second link will be routed to the L2 preprocessor input. This can easily be done for all the central tracking systems (fibers, silicon and central preshower), because they all read out by means of the SVX readout chip.

The L2 system inputs the data through 16 receiver modules, located in two VME crates. Each VME crate contains 7 processor boards for performing the track matching described earlier. The requirements for these boards are well established at this time and they will probably be commercially available DSP based processor boards with high bandwidth I/O capability.

2.3.4 Data transfer

Each of the L2 preprocessors sends its list of identified objects (LIO) to the L2 Global Stage. Each list begins with a header which contains the L1 trigger number of the event and other flags (e.g. the number of entries in the following list). This header allows the Global Stage to verify that each preprocessor has sent data to it from the proper L1 accept.

For some subset of the L1 triggers, the LIO from a particular preprocessor may not be required for making the trigger decisions in the L2 Global Stage or for subsequent analysis of the event in the Level 3 stage. For those events where only L1 triggers from this subset have fired, then this particular L2 preprocessor does not need to actually process the event. However, it still must transfer to the Global Stage a header for this event. The header indicates the L1 trigger number of the event and flags that the preprocessor did not actually analyze the event. Because for every event, every preprocessor sends to the Global Stage a data block consisting of a header plus normally a LIO, it is a simple job for the Global Stage to implement its input buffering. The Global Stages input buffer is setup as an event FIFO. Each of the preprocessors sends its data block for a given event to the Global Stage asynchronously as soon as it has completed its work on that event. When the Global Stage is ready to begin processing of a given event, it pulls in from its side of the FIFO buffers the next data block from each of the preprocessors. The first step in the L2 Global Stage's processing of a new event is to examine the headers from the data blocks from all of the preprocessors to verify that they contain the proper L1 trigger number.

2.3.5 Global processors

The Global Stage of the L2 trigger uses information from the various L2 preprocessors (the LIO) to evaluate L2 terms. For each event, the L2 Global Stage is told by the L2 Framework, which L2 terms the Global Stage will need to evaluate in order for the L2 Framework to make the Accept/Reject decisions for each of the L1 Triggers that fired on a given event. In this way the Global Stage only performs the work that is actually necessary for the L2 Accept/Reject decisions to be made for each event.

The L2 Global Stage uses the same ALPHA based trigger processor hardware as is used in the calorimeter preprocessor. The block diagram arrangement of these two systems looks similar. Each of the L2 preprocessors sends its data to the Global Stage over a link that connects to a Data Source Drive Card(DSDC) on the backplane in the Global Stage's crate. These DSDC's can distribute the information from the preprocessors in a flexible way. If the work performed in the Global Stage is relatively simple then a single processor board can handle it. If this work is more complicated then the same preprocessor information can be sent to two or more processor boards with each of the processor boards working to evaluate a subset of the L2 Terms.

The output from the L2 Global Stage consists of two parts. One output is the L2 term answers. The Global Stage returns to the L2 Framework the pass/fail answers to all of the L2 terms that it was asked to evaluate for each event. The other output is the data block that the Global Stage sends to the Level 3 DAQ system for events that have been passed by the L2 Framework.

This data block from the L2 Global Stage will contain information about why its was able to pass or fail each of the L2 Terms that it evaluated for a given event.

2.4 L2 Budget

The L2 system design will evolve somewhat in the next year. However the system has become much better defined in the last 4 months. Based upon the current system architecture, module counts, our experience with the Level 1.5 trigger, emphasizing the use of commercially available products and developing a high speed VME based processor with CDF we feel reasonably confident about the cost estimates of the system. Table I summarizes the total estimated cost of the L2 The total estimated M&S cost for the full L2 system is summarized in Table 2 and is $1.4M. The estimated need for engineering and design is approximately 15 man-years, most of which we anticipate to come from the staff of collaborating institutions in the U.S. and abroad.

Muon preprocessors	$340K
Calorimeter preprocessors	$270K
Tracking preprocessor for $(r-\phi)$	$170K
Tracking preprocessor for stereo $u-v$	$450K
Global processor stage	$170K
Totals	$1,400K

Table 2: M&S budget summary for the Level-2 trigger system.