

DØ Run 2b Trigger and Online Technical Design Report

The DØ Collaboration

Draft Version 4.3

April 8, 2002

Contents

Contents	2
1 Introduction	7
1.1 Overview	7
1.2 Trigger Upgrade Motivation	7
2 Triggers, Trigger Terms, and Trigger Rates.....	10
2.1 Overview of the DØ Run 2a Trigger System	10
2.2 Leptonic Triggers	12
2.3 Leptons plus Jets	12
2.4 Leptons/Jets plus Missing E_T	12
2.5 Triggers for Higgs Searches.....	13
2.6 Trigger Menu and Rates.....	14
3 Level 1 Tracking Trigger	16
3.1 Goals.....	16
3.1.1 Track Triggers	16
3.1.2 Electron/Photon Identification	16
3.1.3 Track Matching	16
3.2 Description of Current Tracking Trigger	17
3.2.1 Tracking Detectors.....	17
3.2.2 CTT Segmentation.....	17
3.2.3 CTT Electronics	19
3.2.4 CTT Outputs	19
3.2.5 Tracking Algorithm.....	21
3.3 Performance with the Run 2a Tracking Trigger.....	22
3.3.1 Simulations of the Run 2a trigger	22
3.4 Options for Track Trigger Upgrades	24
3.5 Singlet Equations in Track Finding.....	25
3.5.1 Concept	25
3.5.2 Rates and Rejection Improvements.....	26
3.5.3 Tracking Equations	27
3.5.4 Implementation	29

3.5.5 Cost & Schedule	32
3.6 L1 Tracking Trigger Summary and Conclusions	33
4 Level 1 Calorimeter Trigger	35
4.1 Goals.....	35
4.2 Description of Run 2a Calorimeter Electronics.....	36
4.2.1 Overview.....	36
4.2.2 Trigger pickoff.....	37
4.2.3 Trigger summers.....	38
4.2.4 Trigger sum driver.....	39
4.2.5 Signal transmission, cable dispersion.....	40
4.3 Description of Current L1 Calorimeter Trigger.....	41
4.3.1 Overview.....	41
4.3.2 Global Triggers	43
4.3.3 Cluster Triggers	43
4.3.4 Hardware Implementation.....	44
4.3.5 Physical Layout	46
4.4 Motivations for Upgrading the Current System.....	46
4.4.1 Bunch Crossing mis-Identification.....	47
4.4.2 Background Rates and Rejection	47
4.4.3 Conclusions/implications for high luminosity.....	50
4.4.4 Overview of Options for Improvement	51
4.5 Digital Filtering	52
4.5.1 Concept & physics implications	52
4.5.2 Pileup rejection	52
4.5.3 Input data and simulation tools.....	53
4.5.4 Algorithm evaluation parameters	53
4.5.5 Algorithm studied	54
4.5.6 Conclusions	63
4.6 Clustering algorithm options and simulation results	63
4.6.1 Jet algorithms	64
4.6.2 Electron algorithms.....	74
4.6.3 Tau algorithms.....	74
4.6.4 Global sums.....	76

4.7 Topological Considerations	82
4.7.1 Concept & physics implications (acoplanar jets).....	82
4.7.2 Efficiency	83
4.7.3 Rates and rejection improvements	86
4.7.4 Comments	87
4.8 L1 Calorimeter Trigger Implementation.....	87
4.8.1 Constraints	87
4.8.2 L1 Calorimeter Trigger Architectural Overview	87
4.8.3 Trigger Tower Mapping.....	89
4.8.4 ADF Implementation.....	90
4.8.5 ADF to TAB Data Transfer.....	93
4.8.6 TAB implementation	94
4.8.7 GAB Implementation.....	104
4.8.8 TAB-GAB Crate	105
4.8.9 Control and Timing Signals.....	105
4.8.10 Output to the Cal-Track Match System.....	105
4.8.11 Output to L2 and L3.....	106
4.8.12 Downloading, Testing and Monitoring.....	107
4.8.13 Milestones and cost estimate	107
4.8.14 Cost estimate.....	108
4.9 L1 Calorimeter Summary & Conclusions.....	109
5 Level 1 Calorimeter-Track Matching	110
5.1 Overview	110
5.2 Simulation	110
5.2.1 Improving Calorimeter EM Rates Using the L1CTT.....	110
5.2.2 Improving L1CTT Rates Using Calorimeter Jets	112
5.3 Implementation.....	114
5.3.1 L1Muo System.....	114
5.3.2 Some Technical Details	118
5.4 L1CalTrack Trigger	122
5.5 L1CalTrack Cards	130
5.5.1 Serial Link Daughter Board (SLDB).....	130
5.5.2 Muon Trigger Card (MTCxx).....	131

5.5.3 Muon Trigger Flavor Board (MTFB).....	131
5.5.4 Muon Trigger Crate Manager (MTCM)	131
5.5.5 Muon Splitter Cards (MSPLIT).....	131
5.5.6 Muon Trigger Test Card (MTT).....	132
5.5.7 Timing.....	132
5.6 Summary.....	133
6 Level 2 β Trigger	136
6.1 Motivation.....	136
6.2 L2 β Architecture	136
6.3 Cost & schedule	138
6.4 Performance Options	139
7 Level 2 Silicon Track Trigger	141
7.1 Goals.....	141
7.2 STT Upgrade.....	141
7.2.1 Motivation	141
7.2.2 Brief description of Run 2a STT architecture	142
7.2.3 Changes in tracker geometry and implications for STT	145
7.2.4 Simulation studies of Run 2a geometry at Run 2b luminosity....	147
7.2.5 Implementation description for primary STT options.....	148
7.2.6 Summary of cost estimates and major milestones.....	149
8 Online Computing	151
8.1 Introduction	151
8.1.1 Scope	151
8.1.2 Software Architecture	151
8.1.3 Hardware Architecture	152
8.1.4 Motivations	153
8.1.5 Interaction with the Computing Division.....	154
8.1.6 Comments on Run 2a configuration	155
8.2 Plan.....	155
8.2.1 Operational Parameters.....	156
8.2.2 Online Network	157
8.2.3 Level 3 Linux Filter Farm	158

8.2.4 Host Online Systems	158
8.2.5 Control Room Systems.....	160
8.2.6 Data Monitoring Systems.....	161
8.2.7 Database Servers	161
8.2.8 File Servers	162
8.2.9 Slow Control Systems.....	162
8.3 Online Computing Cost Summary	168
8.4 Online Computing Summary	169
9 Summary and Conclusions	170

1 Introduction

1.1 Overview

The primary feature driving the design of the Run 2b trigger elements is the higher rates associated with the approximately factor of 2.5 increase in instantaneous luminosity that will be delivered to the experiments. The rate of events accepted by the Level 1 trigger still must be limited to ~5 kHz to maintain acceptable deadtime levels, so the overall aim of the Level 1 upgrade is to increase the rejection by a factor of at least 2.5.

This report follows a Conceptual Design Report¹, released in October, 2001. That CDR, in turn, was largely based on the report of the DØ Trigger Task Force². The 29-member Task Force was appointed on June 25, 2001 by the DØ Technical Manager (J. Kotcher). These earlier documents explored a large number of trigger upgrade options, many of which were not pursued due to schedule risk, cost, or limited benefits. This Technical Design Report describes the trigger upgrades that we are pursuing, and gives much more detail about the implementations, schedules, and costs.

1.2 Trigger Upgrade Motivation

A powerful and flexible trigger is the cornerstone of a modern hadron collider experiment. It dictates what physics processes can be studied properly and what is ultimately left unexplored. The trigger must offer sufficient flexibility to respond to changing physics goals and new ideas. It should allow the pursuit of complementary approaches to a particular event topology in order to maximize trigger efficiency and allow measurement of trigger turn-on curves. Adequate bandwidth for calibration, monitoring, and background samples must be provided in order to calibrate the detector and control systematic errors. If the trigger is not able to achieve sufficient selectivity to meet these requirements, the capabilities of the experiment will be seriously compromised.

A number of ground rules were established for the design of the Run 2b trigger upgrade. These reflect the expected Run 2b environment: we anticipate operating at a peak luminosity of $\sim 5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ in Run 2b, which is a factor of 2.5 higher than the Run 2a design luminosity. The higher luminosity leads to increased rates for all physics processes, both signal and backgrounds. Assuming ~100 bunches with 132 ns bunch spacing, we expect an average of ~5 non-diffractive “minbias” interactions superimposed on each hard scattering. The increased luminosity also increases occupancies in the detector, leading to a substantial loss in trigger rejection for some systems. Thus, triggers sensitive to pileup or combinatorial effects have rates that grow more rapidly than the growth in luminosity.

We will retain the present trigger architecture with three trigger levels. The Level 1 (L1) trigger employs fast, deterministic algorithms, generating an

¹ http://d0server1.fnal.gov/projects/run2b/Meetings/PAC/Nov01/D0_Rn2b_Trigger_CDR.pdf.

² http://d0server1.fnal.gov/projects/run2b/Trigger/TTF/TTF_Report.pdf.

accept/reject decision every 132 ns. The Level 2 (L2) trigger utilizes Digital Signal Processors (DSPs) and high performance processors with variable processing time, but must issue its accept/reject decisions sequentially. The Level 3 (L3) trigger is based on high-performance processors and is completely asynchronous. The L1 and L2 trigger rely on dedicated trigger data paths, while the L3 trigger utilizes the DAQ readout to collect all event data in a L3 processing node.

We cannot accommodate the higher luminosity by simply increasing trigger rates. The L1 trigger rate is limited to a peak rate of ~ 5 kHz by readout deadtime. The L2 trigger rate is limited to a peak rate of ~ 1 kHz by the calorimeter digitization time. Finally, we have set a goal of ~ 50 Hz for the L3 trigger rate to limit the strain on (and cost of) data storage and offline computing.

The above L1 and L2 rate limits remain essentially the same in Run 2b as in Run 2a. Thus, we must accommodate the higher luminosity in Run 2b by increasing the L1 trigger rejection by a factor of 2.5 and maintaining the current L2 rejection factor of 5. Since Run 2b will focus primarily on high- p_T physics processes, we expect some bandwidth will be freed by reducing the trigger rate devoted to low- p_T processes. However, this reduction is not sufficient to meet our rate limitations, nor does it address the difficulties in triggering efficiently on some important high- p_T processes. Only by upgrading the trigger will we have a reasonable level of confidence in our ability to acquire the data samples needed to carry out the Run 2b physics program.

Potential Run 2b trigger upgrades are further limited by the relatively short time available. Any such upgrade must be completed by the start of high-luminosity running following the installation of the Run 2b silicon tracker in 2005. This goal is made all the more challenging by the need to simultaneously complete and commission the Run 2a detector, acquire physics data, and exploit the resulting physics opportunities. Thus, it is essential that the number and scope of the proposed Run 2b trigger upgrades not exceed the resources of the collaboration.

In the sections below, we describe the technical design of the Run 2b trigger upgrade. Section 2 provides an overview of the trigger architecture and some of the triggering challenges that must be overcome for Run 2b. Section 3 describes the design of the L1 track trigger, which generates track-based triggers and provides tracking information to several other trigger systems. Section 4 describes the design of a new L1 calorimeter trigger that will replace the current trigger (one of the few remaining pieces of Run 1 electronics in DØ). The calorimeter upgrade will employ digital filtering to associate energy with the correct beam crossing in the Run 2b environment and provide the capability of clustering energy from multiple trigger towers. It will also allow improved $e/\gamma/\tau$ triggers that make fuller use of the calorimeter (HAD/EM, cluster shape/size, isolation) and tracking information. Section 5 describes the calorimeter-track matching system, that is based on the existing muon-track matching system. These improvements to the L1 trigger will significantly reduce the rate for multijet background by sharpening trigger thresholds and improving particle identification.

Section 6 describes the upgrade of the L2 β processors to provide additional computational power at L2. Section 7 describes the changes to the L2 Silicon Track Trigger needed to accommodate the new silicon tracker being built for Run 2b. Section 8 describes upgrades to the Online Computing systems needed for Run 2b. Lastly, Section 9 summarizes the content and conclusions of this Technical Design Report.

2 Triggers, Trigger Terms, and Trigger Rates

At 2 TeV, the inelastic proton-antiproton cross section is very large, about 50 mb. At Run 2 luminosities, this results in interaction rates of ~ 25 MHz, with multiple interactions occurring in most beam crossings. Virtually all of these events are without interest to the physics program. In contrast, at these luminosities W bosons are produced at a few Hz and a few top quark pairs are produced per hour. It is evident that sophisticated triggers are necessary to separate out the rare events of physics interest from the overwhelming backgrounds. Rejection factors of nearly 10^6 must be achieved in decision times of a few milliseconds.

The salient features of interesting physics events naturally break down into specific signatures which can be sought after in a programmable trigger. The appearance in an event of a high p_T lepton, for example, can signal the presence of a W or a Z . Combined with jets containing b quark tags, the same lepton signature could now be indicative of top quark pair production or the Higgs. Leptons combined instead with missing energy is a classic SUSY discovery topology, etc. The physics “menu” of Run 2 is built on the menu of signatures and topologies available to the trigger. In order for the physics program to succeed, these fundamental objects must remain un-compromised at the highest luminosities. The following paragraphs give a brief overview of the trigger system and a sampling of the physics impact of the various combinations of trigger objects.

2.1 Overview of the DØ Run 2a Trigger System

The DØ trigger system for Run 2 is divided into three levels of increasing complexity and capability. The Level 1 (L1) trigger is entirely implemented in hardware (see Figure 1). It looks for patterns of hits or energy deposition consistent with the passage of high energy particles through the detector. The calorimeter trigger tests for energy in calorimeter towers above pre-programmed thresholds. Hit patterns in the muon system and the Central Fiber Tracker (CFT) are examined to see if they are consistent with charged tracks above various transverse momentum thresholds. These tests take up to $3.5 \mu\text{s}$ to complete, the equivalent of 27 beam crossings. Since $\sim 10 \mu\text{s}$ of deadtime for readout is incurred following a L1 trigger, we have set a maximum L1 trigger rate of 5 kHz.

Each L1 system prepares a set of terms representing specific conditions that are satisfied (e.g. 2 or more CFT tracks with p_T above 3 GeV). These hardware terms are sent to the L1 Trigger Framework, where specific triggers are formed from combinations of terms (e.g. 2 or more CFT tracks with p_T above 3 GeV AND 2 or more EM calorimeter clusters with energy above 10 GeV). Using firmware, the trigger framework can also form more complex combinations of terms involving ORs of hardware terms (e.g. a match of preshower and calorimeter clusters in any of 4 azimuthal quadrants). The Trigger Framework has capacity for 256 hardware terms and about 40 firmware terms.

The Level 2 trigger (L2) takes advantage of the spatial correlations and more precise detector information to further reduce the trigger rate. The L2 system consists of dedicated preprocessors, each of which reduces the data from one detector subsystem (calorimeter, muon, CFT, preshowers, and SMT). A global L2 processor takes the individual elements and assembles them into physics "objects" such as muons, electrons, or jets. The Silicon Track Trigger (STT) introduces the precise track information from the SMT to look for large impact parameter tracks from b quark decays. Some pipelining is necessary at L2 to meet the constraints of the 100 μ s decision time. L2 can accept events and pass them on to Level 3 at a rate of up to 1 kHz.

The Level 3 (L3) trigger consists of a farm of fast, high-level computers (PCs) which perform a simplified reconstruction of the entire event. Even within the tight time budget of 25 ms, this event reconstruction will allow the application of algorithms in the trigger with sophistication very close to that of the offline analyses. Events that satisfy desired characteristics will then be written out to a permanent storage medium. The maximum L3 output for Run 2a is 50 Hz and is largely dictated by downstream computing limits.

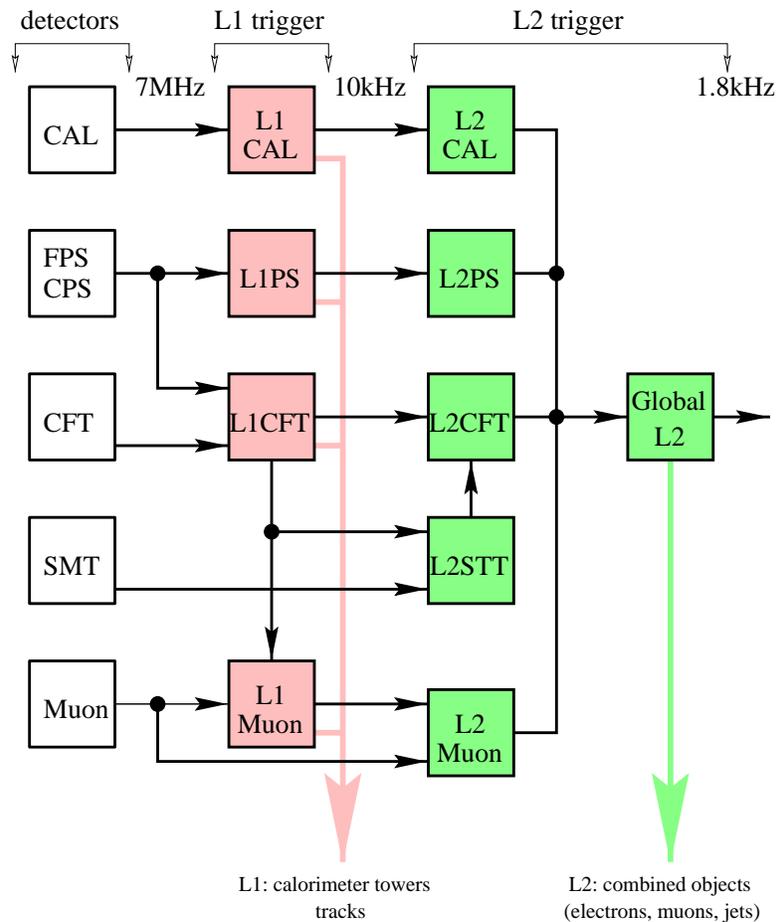


Figure 1. Block diagram of Level 1 and Level 2 triggers, indicating the individual trigger processors that comprise each level.

2.2 Leptonic Triggers

Leptons provide the primary means of selecting events containing W and Z bosons. They can also tag b quarks through their semileptonic decays, complementing the more efficient (but only available at Level 2 through the STT) lifetime selection. The impact of the purely leptonic tag is seen most strongly in the measurements of the W mass, the W and Z production cross sections, and the W width, since the events containing W and Z bosons are selected solely by requiring energetic leptons. The increased statistics provided by Run 2b should allow for a significant improvement in the precision of these measurements, complementing the direct searches in placing more stringent constraints on the Standard Model.

In addition to their inherent physics interest, leptonic signals will play an increasingly important role in the calibration of the energy and momentum scales of the detectors, which is crucial for the top quark and W mass measurements. This will be accomplished using $Z \rightarrow e^+e^-$, $Y \rightarrow e^+e^-$, and $J/\Psi \rightarrow e^+e^-$ for the electromagnetic calorimeter energy scale and the corresponding muon decays for the momentum scale. Since the trigger bandwidth available for acquiring calibration samples must be non-zero, another set of constraints is imposed on the overall allocation of trigger resources.

2.3 Leptons plus Jets

During Run I, lepton-tagged decays of the W bosons and b quarks played an essential role in the discovery of the top quark and were exploited in the measurements of the top mass and production cross section. The new capability provided by the STT to tag b quark decays on-line will allow the collection of many thousands of $t\bar{t}$ pairs in the channel $t\bar{t} \rightarrow \ell\nu + \text{jets}$ with one b -tagged jet. This will be sufficient to allow the study of top production dynamics as well as the measurement of the top decay branching fractions. The precision in measuring the top quark mass will ultimately be limited by our ability to control systematic errors, and the increase in statistics for Run 2b will allow the reduction of several key systematic errors for this channel as well as for the channel $t\bar{t} \rightarrow \ell\nu\ell'\nu + \text{jets}$. One of these, the uncertainty in the jet energy scale, can be reduced by understanding the systematics of the direct reconstruction of W or Z boson decays into jets. The most promising channel in this case is the decay $Z \rightarrow b\bar{b}$, in which secondary vertex triggers can provide the needed rejection against the dominant two-jet background.

2.4 Leptons/Jets plus Missing E_T

Events containing multiple leptons and missing energy are often referred to as the “gold-plated” SUSY discovery mode. These signatures, such as three leptons plus missing energy, were explored in Run I to yield some of the most stringent limits on physics beyond the Standard Model. These investigations will be an integral part of the search for new physics in Run 2. Missing energy is characteristic of any physics process where an invisible particle, such as an energetic neutrino or a massive stable neutral particle, carries away a large

fraction of the available energy. Missing energy combined with leptons/photons or jets can be a manifestation of the presence of large extra dimensions, different SUSY configurations, or other new physics beyond the Standard Model.

2.5 Triggers for Higgs Searches

One of the primary goals of the Run 2b physics program will be to exploit the delivered luminosity as fully as possible in search of the Higgs boson up to the highest accessible Higgs mass³. Since even a delivered luminosity of 15fb^{-1} per experiment may not lead to a statistically significant discovery, the emphasis will be on the combination of as many decay channels and production mechanisms as possible. For the trigger, this implies that flexibility, ease of monitoring, and selectivity will be critical issues.

Coverage of the potential window of discovery is provided by the decay channel $H \rightarrow b\bar{b}$ at low masses, and by $H \rightarrow W^{(*)}W$ at higher masses. In the first case, the production mechanism with the highest sensitivity will probably be in the mode $p\bar{p} \rightarrow WH$. For leptonic W decays, the leptons can be used to trigger on the events directly. If the W decays hadronically, however, the four jets from the $q\bar{q}b\bar{b}$ final state will have to be pulled out from the large QCD backgrounds. Tagging b jets on-line will provide a means to select these events and ensure that they are recorded. Of course, three or four jets with sufficient transverse energy are also required. Another decay mode with good sensitivity is $p\bar{p} \rightarrow ZH$, where the Z decays to leptons, neutrinos, or hadrons. From a trigger perspective, the case where the Z decays hadronically is identical to the WH all-hadronic final state. The final state $ZH \rightarrow \nu\bar{\nu}b\bar{b}$, however, provides a stringent test for the jet and missing E_T triggers, since the final state is only characterized by two modest b jets and missing energy.

Recently, the secondary decay mode $H \rightarrow \tau^+ \tau^-$ has come under scrutiny as a means of bolstering the statistics for Higgs discovery in the low mass region. A trigger that is capable of selecting hadronic tau decays by means of isolated, stiff tracks or very narrow jets may give access to the gluon-fusion production mode $gg \rightarrow H \rightarrow \tau^+ \tau^-$ for lower Higgs masses. This mode can also be important in some of the large $\tan\beta$ SUSY scenarios, where the Higgs coupling to $b\bar{b}$ is reduced, leaving $H \rightarrow \tau^+ \tau^-$ as the dominant decay mode for the lightest Higgs.

The higher Higgs mass regime will be covered by selecting events from $p\bar{p} \rightarrow H \rightarrow W^{(*)}W$ with one or two high-energy leptons from the $W \rightarrow \ell\nu$ decay. This decay mode thus requires a trigger on missing E_T in addition to leptons or leptons plus jets. Supersymmetric Higgs searches will require triggering on final states containing 4 b -quark jets. This will require jet triggers at L1 followed by use of the STT to select jets at L2.

³ Report of the Higgs Working Group of the Tevatron Run 2 SUSY/Higgs Workshop, M. Carena *et al*, hep-ph/0010338.

2.6 Trigger Menu and Rates

As even this cursory review makes clear, the high- p_T physics menu for Run 2b requires efficient triggers for jets, leptons (including taus, if possible), and missing E_T at Level 1. The STT will be crucial in selecting events containing b quark decays; however, its rejection power is not available until Level 2, making it all the more critical that the Level 1 system be efficient enough to accept all the events of interest without overwhelming levels of backgrounds.

In an attempt to set forth a trigger strategy that meets the physics needs of the experiment, the Run 2 Trigger Panel suggested a preliminary set of Trigger Terms for Level 1 and Level 2 triggers⁴. In order to study the expected rates in Run 2b, we have simulated an essential core of triggers which cover the essential high- p_T physics signatures: Higgs boson produced in association with W and Z bosons with Higgs decays to b - \bar{b} , Higgs production and decay to tau leptons, top quark decays in leptonic and semi-leptonic channels, inclusive W and Z boson decays into lepton and muons. The simple triggers we have currently implemented at Level 1 for Run2a will not be able to cope with the much higher occupancies expected in Run2b without a drastic reduction in the physics scope of the experiment and/or prescaling of important physics triggers. Our rate studies have used QCD jets samples in order to determine the effects of background, including multiple low- p_T minimum bias events superimposed on the dominant processes. The results shown in Table 1 indicate the rates expected for the full design luminosity of Run 2a and Run 2b, with the fully implemented Run 2a trigger system with no additional upgrades. (The Run 2b rates in this table neglect the problems associated with running the Run 2a calorimeter trigger at 132ns. This is discussed later in this report.)

⁴ The report of the Run 2 Trigger Panel can be found at http://d0server1.fnal.gov/projects/run2b/Trigger/Docs/Trigger_Panel_Report.ps.

Table 1. Trigger rates for an example trigger menu representing a full spectrum of Run 2 physics channels. (Note: this table is currently a placeholder until more detailed studies are finished.) A representative physics channel for each of the triggers is indicated. The rates for each of these triggers with the design Run2a trigger and the Run2b upgraded trigger are also shown.

Trigger	Example Physics Channel	Rate (kHz) (no upgrade)	Rate (kHz) (with upgrade)
EM Trigger (1 trigger tower > 10 GeV)	$W \rightarrow e\nu$	9	0.5
Jet Trigger (2 trigger towers > 4 GeV)	$ZH \rightarrow \nu\bar{\nu}b\bar{b}$	2	0.5
Track Trigger (2 isolated tracks > 10, 5 GeV matched with EM energy)	$H \rightarrow \tau\tau$	60	0.7
Muon Trigger (muon > 10 GeV)	$W \rightarrow \mu\nu$	6	2

We now turn to describing the upgrades to the trigger system that will enable us to cope with the large luminosities and high occupancies of Run 2b.

3 Level 1 Tracking Trigger

The Level 1 Central Tracking Trigger (CTT) plays a crucial role in the full range of L1 triggers. In this section, we outline the goals for the CTT, provide an overview of the implementation and performance of the present track trigger, and describe the proposed Run 2b CTT upgrade.

3.1 Goals

The goals for the CTT include providing stand-alone track triggers, combining tracking and preshower information to identify electron and photon candidates, and generating track lists that allow other trigger systems to perform track matching. This is a critical part of the L1 muon trigger. We briefly discuss these goals below.

3.1.1 Track Triggers

The CTT provides various Level 1 trigger terms based on counting the number of tracks whose transverse momentum (p_T) exceeds a threshold. Track candidates are identified in the axial view of the Central Fiber Tracker (CFT) by looking for hits in all 8 fiber doublet layers within predetermined roads. Four different sets of roads are defined, corresponding to p_T thresholds of 1.5, 3, 5, and 10 GeV, and the number of tracks above each threshold can be used in the trigger decision. For example, a trigger on two high p_T tracks could require two tracks with $p_T > 5$ GeV and one track with $p_T > 10$ GeV.

Triggering on isolated tracks provides a complementary approach to identifying high- p_T electron and muon candidates, and is potentially useful for triggering on hadronic tau decays. To identify isolated tracks, the CTT looks for additional tracks within a 12° region in azimuth (ϕ).

3.1.2 Electron/Photon Identification

Electron and photon identification is augmented by requiring a significant energy deposit in the preshower detector. The Central Preshower (CPS) and Forward Preshower (FPS) detectors utilize the same readout and trigger electronics as the fiber tracker, and are included in the discussion of tracking triggers. Clusters found in the axial layer of the CPS are matched in phi with track candidates to identify central electron and photon candidates. The FPS cannot be matched with tracks, but comparing energy deposits before/after the lead radiator allows photon and electron candidates to be distinguished.

3.1.3 Track Matching

Track candidates found in the CTT are important as input to several other trigger systems. CTT information is used to correlate tracks with other detector measurements and to serve as seeds for pattern recognition algorithms.

The Level 1 muon trigger matches CTT tracks with hits in the muon detector. To meet timing requirements, the CTT tracks must arrive at the muon trigger on the same time scale as the muon proportional drift tube (PDT) information becomes available.

The current Level 1 trigger allows limited azimuthal matching of tracking and calorimeter information at the quadrant level (see section 2.1). Significantly increasing the flexibility and granularity of the calorimeter track matching is an integral part of the proposed modifications for Run 2b (see section 5). This option requires sending track lists to the calorimeter trigger.

The L2 Silicon Track Trigger (STT) uses tracks from the CTT to generate roads for finding tracks in the Silicon Microstrip Tracker (SMT). The precision of the SMT measurements at small radius, combined with the larger radius of the CFT, allows displaced vertex triggers, sharpening of the momentum thresholds for track triggers, and elimination of fake tracks found by the CTT. The momentum spectrum for b-quark decay products extends to low p_T . The CTT therefore aims to provide tracks down to the lowest p_T possible. The Run 2a CTT generates track lists down to $p_T \approx 1.5$ GeV. The CTT tracks must also have good azimuthal (ϕ) resolution to minimize the width of the road used by the STT.

In addition to the track lists sent to the STT, each portion of the L1 track trigger (CFT, axial CPS, and FPS) provides information for the Level 2 trigger decision. The stereo CPS signals are also sent to L2 to allow 3-D matching of calorimeter and CPS signals.

3.2 Description of Current Tracking Trigger

We have limited our consideration of potential track trigger upgrades to those that preserve the overall architecture of the current tracking trigger. The sections below describe the tracking detectors, trigger segmentation, trigger electronics, outputs of the track trigger, and the trigger algorithms that have been developed for Run 2a.

3.2.1 Tracking Detectors

The CFT is made of scintillating fibers mounted on eight low-mass cylinders. Each of these cylinders supports four layers of fibers arranged into two doublet layers. The innermost doublet layer on each cylinder has its fibers oriented parallel to the beam axis. These are referred to as Axial Doublet layers. The second doublet layer has its fibers oriented at a small angle to the beam axis, with alternating sign of the stereo angle. These are referred to as Stereo Doublet layers. Only the Axial Doublet layers are incorporated into the current L1 CTT. Each fiber is connected to a visible light photon counter (VLPC) that converts the light pulse to an electrical signal.

The CPS and FPS detectors are made of scintillator strips with wavelength-shifting fibers threaded through each strip. The CPS has an axial and two stereo layers mounted on the outside of the solenoid. The FPS has two stereo layers in front of a lead radiator and two stereo layers behind the radiator. The CPS/FPS fibers are also read out using VLPCs.

3.2.2 CTT Segmentation

The CTT is divided in ϕ into 80 Trigger Sectors (TS). A single TS is illustrated schematically in Figure 2. To find tracks in a given sector, information

is needed from that sector, called the home sector, and from each of its two neighboring sectors. The TS is sized such that the tracks satisfying the lowest p_T threshold (1.5 GeV) is contained within a single TS and its neighbors. A track is 'anchored' in the outermost (H) layer. The ϕ value assigned to a track is the fiber number at the H layer. The p_T value for a track is expressed as the fiber offset in the innermost (A) layer from a radial straight-line trajectory.

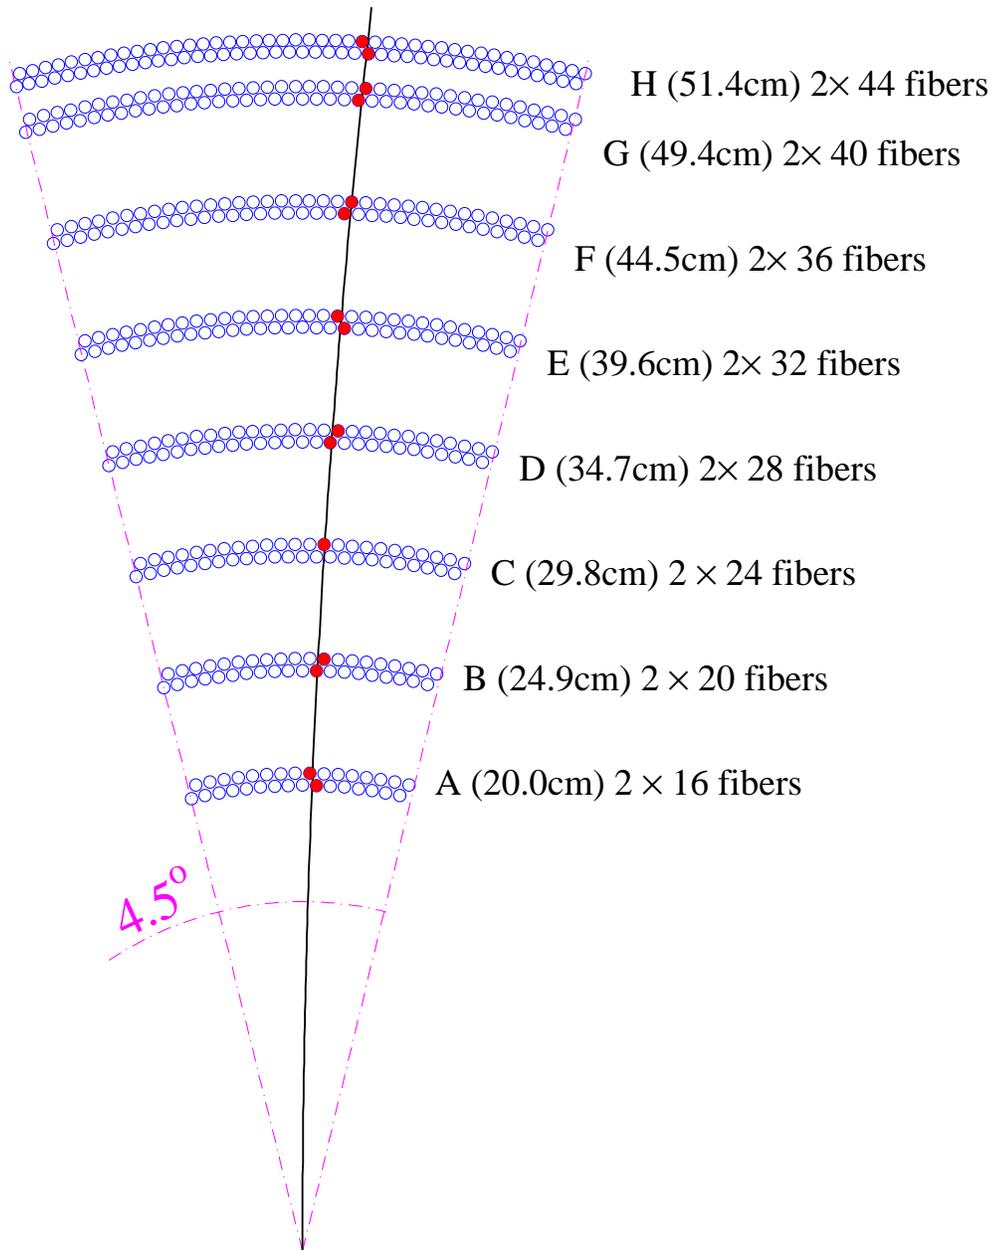


Figure 2. Illustration of a CTT trigger sector and the labels assigned to the eight CFT cylinders. Each of the 80 trigger sectors has a total of 480 axial fibers.

The home sector contains 480 axial fibers. A further 368 axial fibers from each of the neighbors, the 'next' and 'previous', sectors are sent to each home sector to find all the possible axial tracks above the p_T threshold. In addition, information from 16 axial scintillator strips from the CPS home sector and 8 strips from each neighboring sector are included in the TS for matching tracks and preshower clusters.

3.2.3 CTT Electronics

The tracking trigger hardware has three main functional elements. The first element consists of the Analog Front-End (AFE) boards that receive signals from the VLPCs. The AFE boards provide both digitized information for L3 and offline analysis as well as discriminated signals used by the CTT. Discriminator thresholds should be set at a few photoelectrons for the CFT and at the 5 – 10 MIP level for the CPS and FPS. Discriminator outputs for 128 channels are buffered and transmitted over a fast link to the next stage of the trigger. The axial layers of the CFT are instrumented using 76 AFE boards, each providing 512 channels of readout. The axial CPS strips are instrumented using 10 AFE boards, each having 256 channels devoted to axial CPS readout and the remaining 256 channels devoted to stereo CFT readout. The FPS is instrumented using 32 AFE boards. Additional AFE boards provide readout for the stereo CPS strips and remaining stereo CFT fibers.

The second hardware element is the Mixer System (MS). The MS resides in a single crate and is composed of 20 boards. It receives the signals from the AFE boards and sorts them for the following stage. The signals into the AFE boards are ordered in increasing azimuth for each of the tracker layers, while the trigger is organized into TS wedges covering all radial CFT/CPS axial layers within 4.5 degrees in ϕ . Each MS board has sixteen CFT inputs and one CPS input. It shares these inputs with boards on either side within the crate and sorts them for output. Each board then outputs signals to two DFEA boards (described below), with each DFEA covering two TS.

The third hardware element is based on the Digital Front-End (DFE) motherboard. These motherboards provide the common buffering and communication links needed for all DFE variants and support two different types of daughter boards, single-wide and double-wide. The daughter boards implement the trigger logic using Field Programmable Gate Array (FPGA) chips. The signals from the Mixer System are received by 40 DFE Axial (DFEA) boards. There are also 5 DFE Stereo (DFES) boards that prepare the signals from the CPS stereo layers for L2 and 16 DFEF boards that handle the FPS signals.

3.2.4 CTT Outputs

The current tracking trigger was designed to do several things. For the L1 Muon trigger it provides a list of found tracks for each crossing. For the L1 Track Trigger it counts the number of tracks found in each of four p_T bins. It determines the number of tracks that are isolated (no other tracks in the TS or its neighbors). The sector numbers for isolated tracks are recorded to permit triggers on acoplanar high p_T tracks. Association of track and CPS clusters provides the

ability to recognize both electron and photon candidates. FPS clusters are categorized as electrons or photons, depending on an association of MIP and shower layer clusters. Finally, the L1 trigger boards store lists of tracks for each beam crossing, and the appropriate lists are transferred to L2 processors when an L1 trigger accept is received.

The L1 CTT must identify real tracks within several p_T bins with high efficiency. The nominal p_T thresholds of the bins are 1.5, 3, 5, and 10 GeV. The L1 CTT must also provide rejection of fake tracks (due to accidental combinations in the high multiplicity environment). The trigger must perform its function for each beam crossing at either 396 ns or 132 ns spacing between crossings. A list of up to six found tracks for each crossing is packed into 96 bits and transmitted from each of the 80 trigger sectors. These tracks are used by the L1 Muon trigger and must be received within 1000 ns of the crossing. These track lists are transmitted over copper serial links from the DFEA boards.

The L1 CTT counts the number of tracks found in each of the four p_T bins, with subcategories such as the number of tracks correlated with showers in the Central Preshower Detector, and the number of isolated tracks. Azimuthal information is also preserved so that information from each ϕ region can be correlated with information from other detectors. The information from each of the 80 TS is output to a set of 8 Central Tracker Octant Card (CTOC) boards, which are DFE mother boards equipped with CTOC type double wide daughter boards. During L1 running mode, these boards collect the information from each of 10 DFEA boards, combine the information and pass it on to a single Central Track Trigger Terms (CTTT) board. The CTTT board, also a DFE-type mother board equipped with a similar double wide daughter board, assembles the information from the eight CTOC boards and makes all possible trigger terms for transmission to the Trigger Manager (TM). The TM constructs the 32 AND/OR terms that are used by the Trigger Framework in forming the L1 trigger decision. For example, the term "TPQ(2,3)" indicates two tracks associated with CPS hits were present in quadrant 3. Additional AND/OR terms provide CPS and FPS cluster characterization for use in L1. The Trigger Framework accommodates a total of 256 such terms, feeding them into a large programmable AND/OR network that determines whether the requirements for generating a trigger are met.

The DFEA boards store lists of tracks from each crossing, and these lists are transferred to the L2 processors when an L1 trigger accept is received. A list of up to 6 tracks is stored for each p_T bin. When an L1 trigger accept is received, the normal L1 traffic is halted and the list of tracks is forwarded to the CTOC board. This board recognizes the change to L2 processing mode and combines the many input track lists into a single list that is forwarded to the L2 processors. Similar lists of preshower clusters are built by the DFES and DFEF boards for the CPS stereo and FPS strips and transferred to the L2 processors upon receiving an L1 trigger accept.

3.2.5 Tracking Algorithm

The tracking trigger algorithm currently implemented is based upon hits constructed from pairs of neighboring fibers, referred to as a “doublet”. Fibers in doublet layers are arranged on each cylinder as illustrated in Figure 3. In the first stage of the track finding, doublet layer hits are formed from the individual axial fiber hits. The doublet hit is defined by an OR of the signals from adjacent inner and outer layer fibers in conjunction with a veto based upon the information from a neighboring fiber. In Figure 3, information from the first fiber on the left in the upper layer (fiber 2) would be combined by a logical OR with the corresponding information for the second fiber from the left on the lower layer (fiber 3). This combination would form a doublet hit unless the first fiber from the left in the lower layer (fiber 1) was also hit. Without the veto, a hit in both fiber 2 and fiber 1 would result in two doublet hits.

Doublet Layer

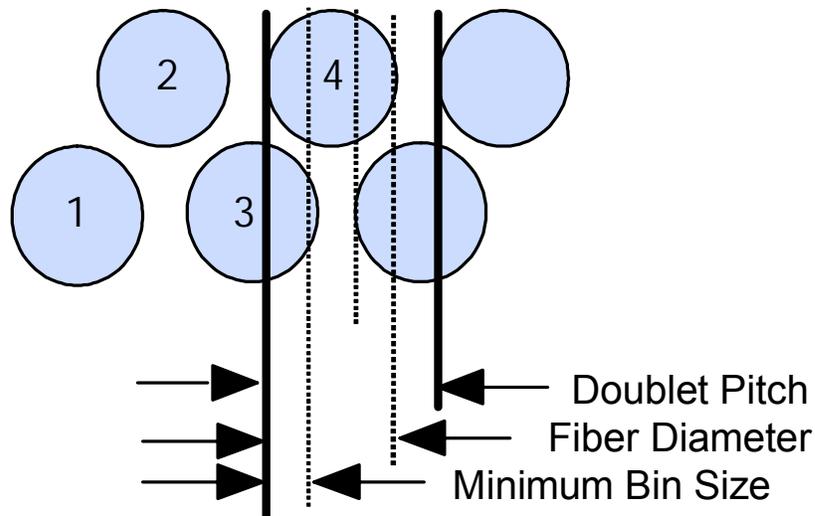


Figure 3. Sketch illustrating the definition of a fiber doublet. The circles represent the active cross sectional areas of individual scintillating fibers. The boundaries of a doublet are shown via the thick black lines. The dotted lines delineate the four distinguishable regions within the doublet.

The track finding within each DFEA board is straightforward. Each daughter board has 4 large FPGA chips, one for each of the four p_T bins. Within each chip the track roads are represented by equations which correspond to a list of the doublets that are hit by a track with a given p_T and ϕ . For each possible road the eight fibers for that road are combined into an 8-fold AND. If all the fibers on that road were hit then all 8 inputs of the AND are TRUE and the result is a TRUE.

The FPGA chips are loaded with the equations for all possible real tracks in each sector in each p_T range. Each TS has 44 ϕ bins corresponding to the 44 H layer doublets in a sector and 20 possible p_T bins for about 12 different routes through the intermediate layers with fixed ϕ and p_T . This results in about 17K equations per TS.

The individual track results are then OR'ed together by ϕ bin and sorted by p_T . Up to six tracks per TS are reported to the trigger. This list of 6 tracks is then sent to the fifth or back-end chip on the daughter board for all the remaining functions.

The FPGA chips have a very high density of logic gates which lends itself well to the track equations. Within these chips all 17k equations are processed simultaneously in under 200 ns. This design also keeps the board hardware as general as possible. The motherboard is simply an I/O device and the daughter boards are general purpose processors. Since algorithms and other details of the design are implemented in the FPGA, which can be reprogrammed via high level languages, one can re-download different trigger configurations for each run or for special runs and the trigger can evolve during the run.

3.3 Performance with the Run 2a Tracking Trigger

We have simulated the rates to be expected for pure track triggers in Run 2b, taking into account the additional minimum bias events within the beam crossing of interest due to the increased luminosity.

3.3.1 Simulations of the Run 2a trigger

Under Run 2a conditions, the current track trigger performs very well in simulations. For example, for a sample of simulated muons with $p_T > 50$ GeV/c, we find that 97% of the muons are reconstructed correctly; of the remaining 3%, 1.9% of the tracks are not reconstructed at all and 1.1% are reconstructed as two tracks due to detector noise. (As the background in the CFT increases, due to overlaid events, we expect the latter fraction to get progressively higher). Since the data-taking environment during Run 2b will be significantly more challenging, it is important to characterize the anticipated performance of the current trigger under Run 2b conditions.

To test the expected behavior of the current trigger in the Run 2b environment, the existing trigger simulation code was used with an increased number of overlaid minimum bias interactions. The minimum bias interactions used in this study were generated using the ISAJET Monte Carlo model. Based on studies of detector occupancy and charged track multiplicity in minimum-bias events, we expect that this should give a worst-case scenario for the Run 2b trigger.

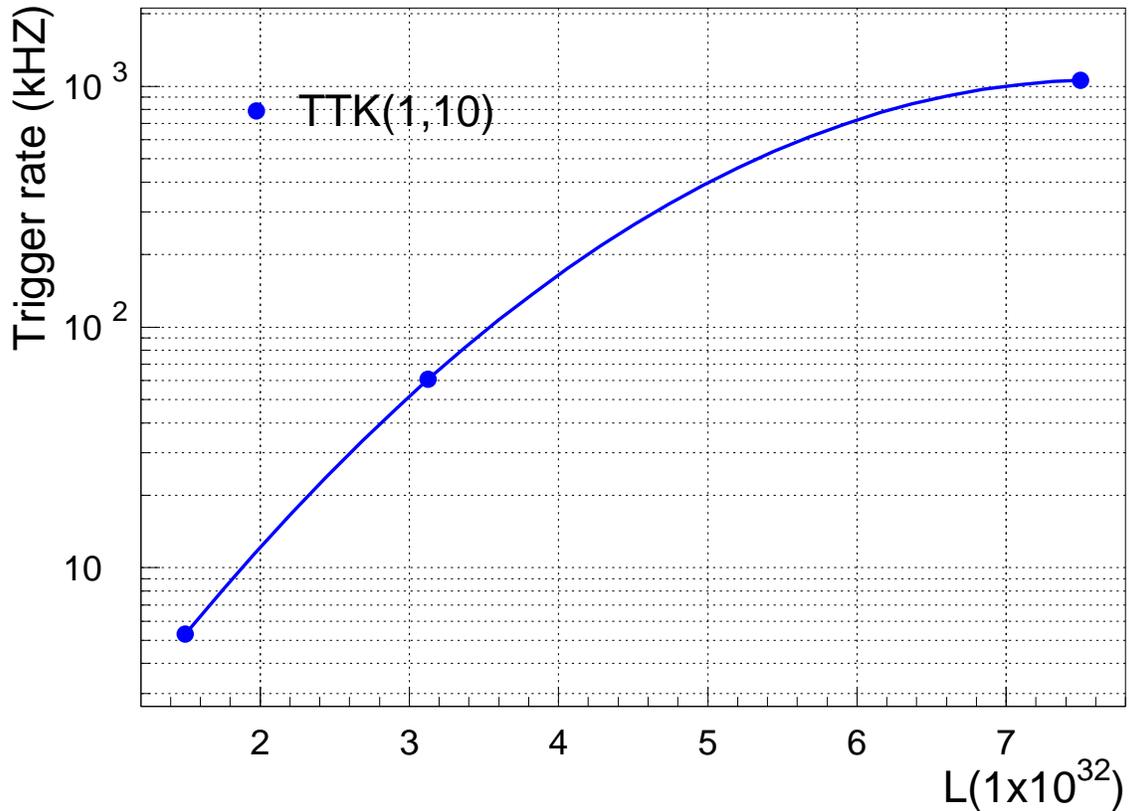


Figure 4. Track trigger rate as a function of the number of underlying minimum bias interactions. TTK(2,10) is a trigger requiring 2 tracks with transverse momentum greater than 10 GeV.

Figure 4 shows the rate for a trigger requiring two tracks with $p_T > 10$ GeV as a function of the number of underlying minimum bias interactions and hence luminosity. During Run 2b, we expect that the mean number of underlying interactions will be about 5. Figure 4 shows that the tracking trigger rate for the current trigger version is expected to rise dramatically due to accidental hit combinations yielding fake tracks. This results in an increasingly compromised tracking trigger.

Figure 5 shows the probability for three specific track trigger terms to be satisfied in a given crossing. They are strongly dependent upon the number of underlying minimum bias interactions. These studies indicate that a track trigger based upon the current hardware will be severely compromised under Run 2b conditions. Not shown in the figure, but even more dramatic, is the performance of the 5 GeV threshold track trigger. This is satisfied in more than 95% of beam crossings with 5 minbias interactions. It will clearly not be possible to run the current stand-alone track trigger in Run 2b. But much worse, the information available to the muon trigger, electron trigger, and STT becomes severely compromised by such a high rate of fake high- p_T tracks.

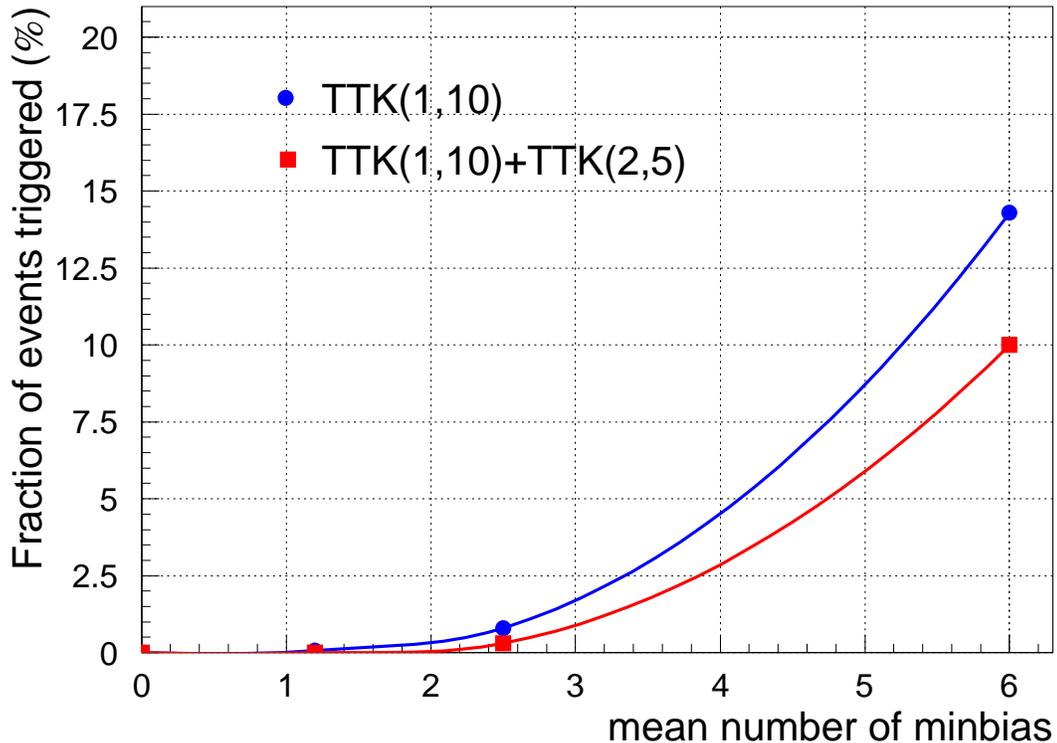


Figure 5. The fraction of events satisfying several track term requirements as a function of the number of minimum bias events overlaid. $TTK(n, p_T)$ is a trigger requiring n tracks with transverse momentum greater than p_T .

Based upon these simulations, we believe that the significant number of multiple interactions in Run 2b and the large CFT occupancy fractions they induce, will compromise performance of the current tracking trigger.

3.4 Options for Track Trigger Upgrades

As demonstrated above, the primary concern with the track trigger is the increase in rate for fake tracks as the tracker occupancy grows. Since the current track trigger requires hits in all 8 axial doublet layers, the only path to improving trigger rejection is to improve the trigger selectivity by incorporating additional information into the trigger algorithm. The short timescale until the beginning of Run 2b and resource limitations conspire to make it impossible to improve the physical granularity of the fiber tracker or to add additional tracking layers to the CFT. Out of a number of options studied, an upgrade to the DFEA boards incorporating much more powerful FPGAs has emerged as the most promising. Essentially, more processing power allows the use of the full granularity and resolution of the individual CFT fibers rather than doublets in Level 1 track finding. Studies of this approach are presented below.

3.5 Singlet Equations in Track Finding

3.5.1 Concept

The motivation behind the use of singlet equations is illustrated in Figure 3, which shows a fragment of a CFT doublet layer. The thick black lines mark the area corresponding to a doublet hit, the current granularity of the L1CTT. As one can see from Figure 3, the doublet is larger than the fiber diameter. Since the hits from adjacent fibers are combined into the doublets before the tracking algorithm is run, this results in a widening of the effective width of a fiber to that of a doublet, decreasing the resolution of the hits that are used for track finding. In particular, the doublet algorithm is such that if fibers 1 and 4 shown on Figure 3 are hit, the trigger considers the doublet formed by fibers 1 and 2 *and* the doublet formed by fibers 3 and 4 to be hit. As the single-fiber occupancy grows, the application of this doublet algorithm results in a disproportionate increase in the hit occupancy seen by the trigger.

Track-finding roads based instead on single fibers will be inherently narrower and will therefore have a reduced probability of selecting a random combination of hits. We have simulated different trigger configurations which include the all-singlet case (16 layers), as well as mixed schemes where some CFT layers are treated as pairs of singlet layers and the rest as doublets. In order to label the schemes we use the fact that the 8 layers of the CFT are labeled from A to H (see Figure 2). We use upper case letters to indicate that hits in this layer were treated as doublets; lower case letters indicate singlets. In this notation “ABCDEFGH” indicates the Run 2a CTT scheme with 8 layers of doublets and “abcdefgh” indicates 16 layers of singlets. Equations specifying which fibers should be hit as a function of momentum and azimuthal angle were generated for all configurations. Note that, in the results reported here, the equations have been generated specifying only which fibers should be hit, and not using vetoes on fibers that should not be hit. This will be discussed more completely in Section 3.5.3.

Because of the space-filling structure of the CFT shown in Figure 3, the number of fibers hit by a track passing through all 8 layers of the CFT varies with position. This is shown in Figure 6, where the probability that a track will have ≥ 8 , ≥ 10 , ≥ 11 , ≥ 12 and 13 hits out of 16 possible for the 16 layer singlet trigger scheme (abcdefgh) is plotted as a function of track sagitta. Here, it is assumed that fibers are 100% efficient.

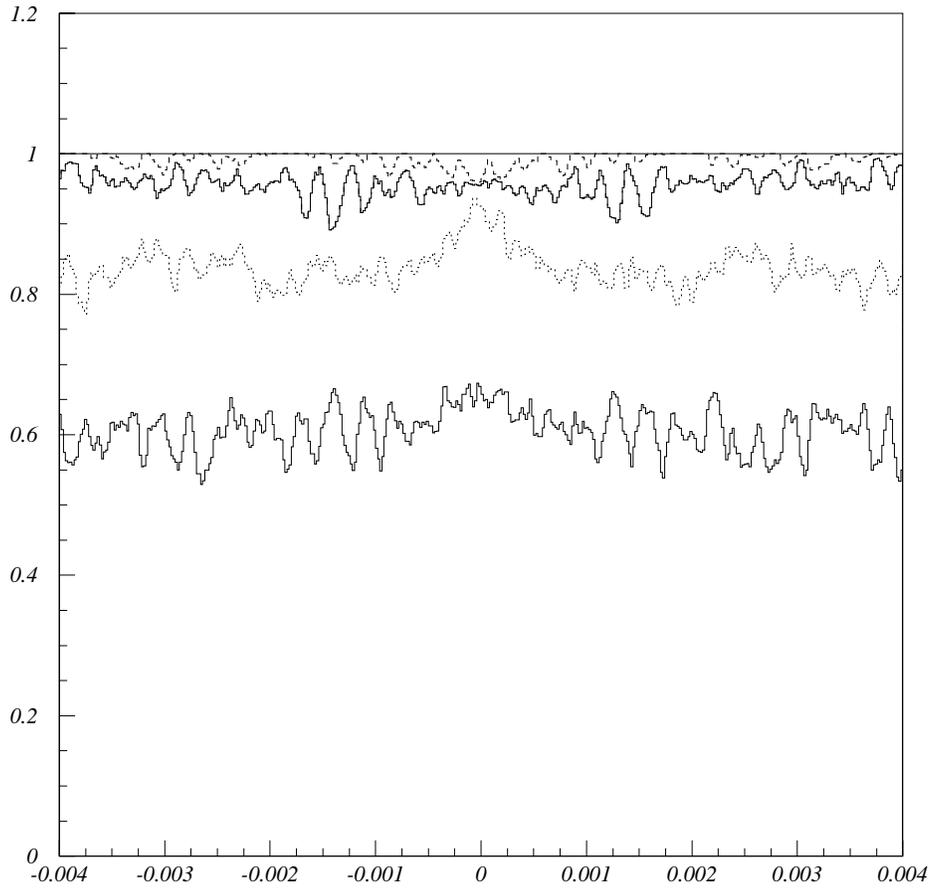


Figure 6. Geometrical acceptance for a charged particle to satisfy $a \geq 8$ (solid line), ≥ 10 (dashed curve), ≥ 11 (upper solid curve), ≥ 12 (dot-dashed curve) and 13 (lower solid curve) hit requirement in the 16-trigger layer configuration “abcdefgh”, versus the particle track sagitta, $s = 0.02 \cdot e / p_T$, for a track starting at the center of a CFT trigger sector.

3.5.2 Rates and Rejection Improvements

The existing trigger simulation was adapted to make a realistic estimate of the trigger performance. Single muons were generated, overlaid on events containing exactly six (ISAJET) minimum bias interactions and put through the detailed $D\bar{0}$ simulation. They were then put through the modified trigger simulator. At this stage in the simulation, single fiber efficiency is still assumed to be 100%. The fraction of events in which a trigger track matched the muon is defined as the trigger efficiency. A separate sample containing six minimum bias events only is used to measure the fake rate, since there are no high- p_T tracks in this sample.

The results of the procedure described are summarized in Table 2. For the case of 12-layer equations, the background is reduced by a factor of about 4

without significant loss of efficiency. For the 16-layer case the improvement is larger, and is more than a factor of 10 for high p_T tracks.

Note also, that the fraction of mis-reconstructed muons, *i.e.* muons which give a trigger in the wrong p_T bin is also reduced when going to singlet equations, especially for the 16-layer case. This is very important for STT, which depends on the quality of the seed tracks from L1CTT.

Table 2. Fractions of events (in %) with 6 minimum bias interactions that satisfy various track trigger requirements. The default Run2a L1CTT “ABCDEFGH” is compared with an implementation of the tracking trigger using different schemes. We consider an all singlet configuration (16 Layers, “abcdefgh”), and three mixed schemes, “abcdEFGH” (singlets for the axial fibers on inner four cylinders and doublets of the axial fibers on the outer four cylinders), “ABCDefgh” (singlets for the axial fibers on outer four cylinders and doublets of the axial fibers on the inner four cylinders), “abcdefGH” (singlets for the axial fibers on inner six cylinders and doublets of the axial fibers on the outer four cylinders). TTK(n, p_T) is a trigger requiring n tracks with transverse momentum greater than p_T .

	Default Doublet Equations	16-Layer Singlet Equations	12-Layer Equations “abcdEFGH”	12-Layer Equations “ABCDefgh”	14-Layer Equations “abcdefGH”
Efficiency for $p_T > 10$	96.9	99.3	98.6	97.3	99.2
Efficiency for $5 < p_T < 10$	91.1	97.8	92.8	90.8	91.6
Efficiency for fake $p_T > 10$	5.8	0.4	1.6	1.4	0.7
Efficiency for fake $5 < p_T < 10$	8.0	0.7	2.4	2.4	1.6
Fake TTK(1,10)	5.8	0.4	1.6	1.4	0.7
Fake TTK(2,10)	0.7	0	0.13	0	0.03
Fake TTK(1,5)	12.1	1.1	3.7	3.7	2.2
Fake TTK(2,5)	2.2	0.05	0.4	0.08	0.13

3.5.3 Tracking Equations

In order to study the FPGA resources necessary to implement the new algorithms, we have generated the track trigger equations for a CFT Trigger Sector for different combinations of doublet and singlet hits. The results are summarized in Table 3. The first column specifies which layers were considered as doublets and which as singlets. We use the same notation for describing the various trigger configuration scheme as before. In this notation “ABCDEFGH” indicates the Run 2A scheme with 8 layers of doublets and “abcdefgh” indicates 16 layers of singlets. The second column gives the increase in the number of equations per CFT trigger sector relative to the Run 2A scheme. The number of terms per equation is the average number of hits on a track and represents the average number of inputs the FPGA has to compare with the equation template

for each equation. The product of number of equations and terms per equation is a measure of the FPGA resources required. We find that the new algorithms require between 4 and 15 times more resources than the Run 2A scheme. As expected, a 16-layer all-singlets algorithm is the most demanding in terms of computing resources and memory.

Table 3. Resource evaluation for different trigger algorithms (see text).

Singlet/Doublet Scheme	Relative number of equations	# terms/equation
ABCDEFGH	1.0	8
Abcdefgh	15.3	12.6
abCdeFgh	10.5	11.4
Abcdefgh	10.0	11.4
abcdefGH	7.7	11.4
ABCDefgh	5.7	10.3
aBcDeFgH	5.6	10.2
abcdEFGH	4.2	10.2

The table represents a first set of resource estimates. Refinements and optimizations of the algorithms can affect the actual FPGA requirements.

The equations used to perform the simulations were generated to include only fibers that should be hit, and not the additional information given by adjacent fibers which did not fire. Implementing an explicit veto would increase the number of terms in each of the equations to the total number of layers in the trigger (i.e., the 16 singlet trigger would have 16 terms in each of its equations.).

It may be necessary in the equations to allow explicitly for missed fibers due to inefficiencies. This would dramatically increase the number of equations. On the other hand, about 90% of the acceptance of the trigger is represented by only 5% of these equations. Some of the equations have an extremely low probability of firing. Thus a fairly high trigger efficiency can be maintained with a much smaller number of equations. The highest p_T bin is made up of only 15% of all equations, while the lowest p_T bin is made up of half the equations. Thus the resource requirements could be reduced by retaining all equations in the higher p_T bins and eliminating the low acceptance equations in the lowest p_T bin. Alternatively, it is possible to count the number of hit coincidences on each equation and to require a minimum number of hits. This requires a different estimate of resource allocation.

We are studying these factors to determine the optimal algorithm to balance high efficiency and acceptance against the resource requirements. We are also

looking into different algorithms that do not use sets of fixed hit patterns (equations) to compare hits to, but rather take a more computational approach.

To estimate the cost of the L1CTT upgrade we define a strawman proposal, which achieves a significant improvement in the performance of the L1CTT within the achievable increase in resources of about a factor 10. This may however not be the optimal algorithm and work on optimizing the algorithm is continuing. For this proposal, we use p_T thresholds of 1.5, 5, 10 and 20 GeV. We simulate the Poisson statistics of the number of photoelectrons. Typically we expect 8 photoelectrons per fiber per minimum ionizing particle. Low fluctuations in this number result in inefficiencies since we require at least two photoelectrons for a hit. We maximize the background rejection of the higher p_T thresholds by using 16 singlet layers. For the lower p_T thresholds, we use less singlet layers to reduce the number of equations required. The parameters are summarized in Table 4. Our cost estimate below is based on the 10 fold increase in resources from Table 4.

Table 4. Strawman proposal for L1CTT upgrade.

pT threshold	Efficiency	Doublet/singlet scheme	Resources relative to total Run 2A resources
pT > 20 GeV	98%	abcdefgh	$28 \times 1.5 \times 0.075 = 3.15$
pT > 10 GeV	98%	abcdefgh	$28 \times 1.5 \times 0.075 = 3.15$
pT > 5 GeV	95%	abcdeFGH	$6.2 \times 1.4 \times 0.2 = 1.3$
pT > 1.5 GeV	95%	abcdEFGH	$3.0 \times 1.2 \times 0.5 = 2.5$
total			10.1

3.5.4 Implementation

The implementation, cost and schedule depends largely on the algorithm chosen and what FPGA resources the algorithm requires.

A block diagram of the existing L1 tracking trigger is shown in Figure 7. The entire track finding logic is included on 80 daughter boards located on 40 mother boards. The current design already brings the singlet hits onto these daughter boards so we do not anticipate any changes to the mother boards or any of the upstream electronics and cabling. The current system publishes the 6 highest p_T tracks in each of 4 momentum bins (24 tracks). The new design will do the same so no changes are needed in the output daughter cards or downstream cabling. The crate controller card stores the logic in flash memory for local downloading to the other cards in the crate. The present design uses 1.6 Mbytes and the controller can hold 512 Mbytes, giving an expansion factor of more than 250. Larger gate array chips do not use much more power so the power supplies and cooling are also adequate.

The present daughter board houses five Xilinx Virtex-I chips. These consist of one Virtex 600, three Virtex 400's and one Virtex 300. They are housed in pin ball grid array packages. The PC board required 10 layers to interconnect these chips. The present Virtex 600 has an array of 64x96 slices with each slice containing 2 four-input look up tables (LUT) giving a total of 12,288 LUT's.

The algorithms that we are considering require at least 10 times more resources, which the present daughter boards cannot provide. Xilinx does not intend to produce more powerful chips that are pin-compatible with the Virtex-I chips we are using now. If we want to use newer, more powerful FPGAs, new daughter boards have to be designed.

The Virtex-II series FPGAs have 8 to 10 times larger logic cells than the largest chips that we are currently using. The Virtex-II series offers chips which have 2M to 8M system gates and about 25K to 100K logic cells. These chips come in a ball grid array as well as Flip Chip packaging similar in size to the existing parts. Thus, we will be able to fit five of these chips on new daughter boards of the same size as the present daughter boards. Due to the denser parts the PC boards may require 2 or 4 additional layers.

In addition, the speed of the Virtex-II chips is in the range of 200-300 MHz. We are looking into the gains we may achieve by utilizing this increased speed and similarities of "sub-units" of different equations. By running the chips at higher speeds, we may be able pipeline some of the processing allowing possible reuse of similar "sub-units" of equations stored in different logic cells, and therefore accommodate a larger number of equations.

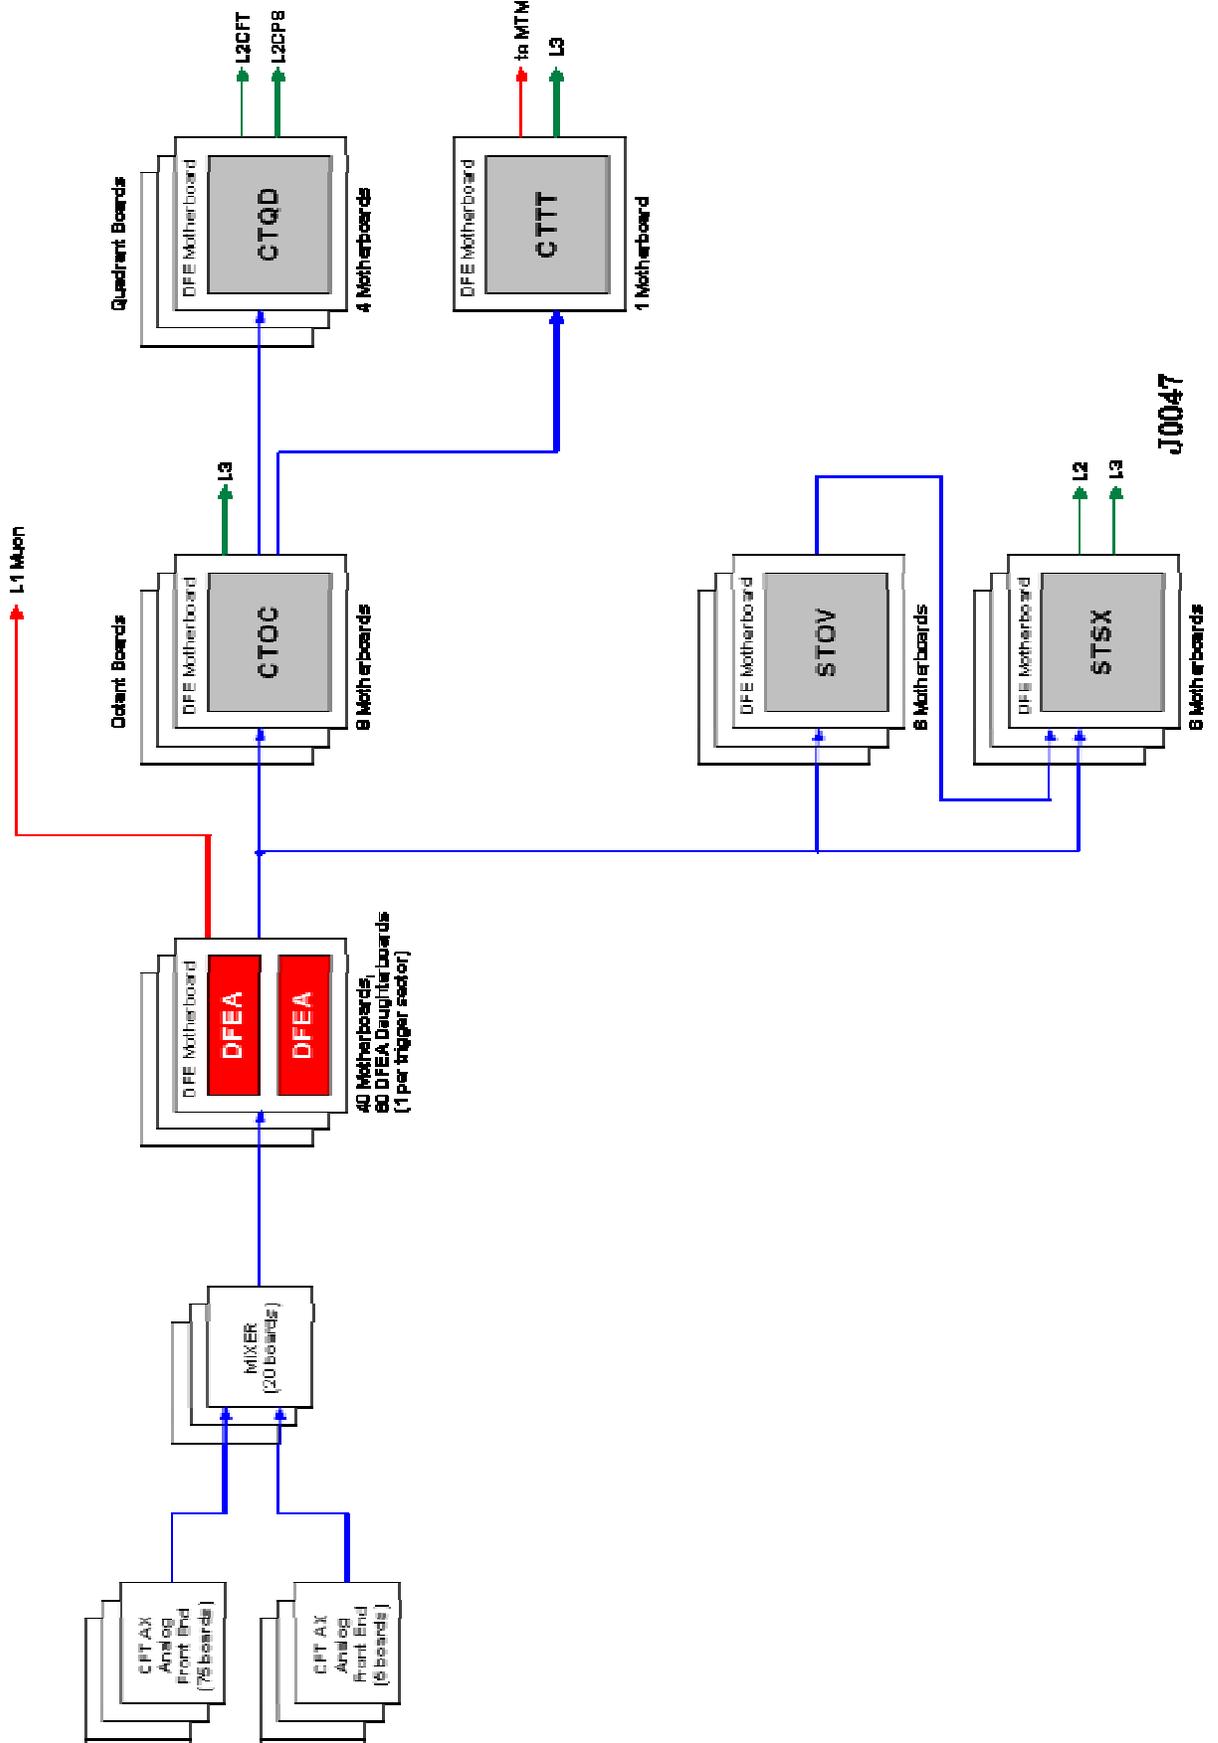


Figure 7. Block Diagram of L1CTT Electronics.

3.5.5 Cost & Schedule

Preliminary cost estimates for this option, based on our experience fabricating and instrumenting the current boards, are shown in Table 5 below. It should be pointed out that this upgrade affects only the 88 DFEA daughter boards. All of the AFE, MB and other DFE type boards are not changed or modified in any way.

The costs of the Virtex-II series FPGA are based on current price of the XC2V8000 and XC2V4000 chips and include a projected reduction in price of 10% per quarter over the next two years. The production quantity purchase of these chips is scheduled to occur in April 2004.

Table 5. Preliminary cost estimate for upgrade to the track trigger associated with the handling of fiber singlets that includes replacement of the daughter boards. A contingency of 50% is applied.

Item/process	Unit Cost (\$)	# Required	Total Cost (\$k)	Total Cost + Contingency (\$k)
Fabricate/stuff new Daughter Boards	500	88	44	66
Purchase new FPGA for 2 low pT bins	900	176	158	238
Purchase new FPGA for 2 High pT bins	2000	176	352	528
Test stand for prototypes	15k	1	15	22.5
Engineering			210	263
<i>TOTAL</i>			<i>780</i>	<i>1117</i>

A preliminary schedule for replacing the DFEA daughter boards is given in Table 6. The engineering time associated with this FPGA upgrade is of limited scope and the production is limited to a single board series; much of the effort and resources will necessarily be focused on the algorithm logic and FPGA programming. We include two rounds of prototype board production, as we do not have much experience with the Virtex II series chips. While these chips were announced last year, the first ones of this series were available for commercial use only starting Jan 2002.

Table 6. Preliminary schedule for the fiber singlet DFEA upgrade.

Description of Task	Completion Date
Prototype algorithm coded and simulated using FPGA simulation tools	11/02
Target algorithm coded and simulated using FPGA simulation tools	6/03
Layout Prototype I boards	8/03
Develop test procedures	7/03
Assemble and test prototype I	12/03
Layout prototype II boards	11/03
Assemble and test prototype II	2/04
Test prototype II at FNAL with the full test chain	1/04
Design, Layout and Fabricate production boards	4/04
Daughter boards tested and ready for installation	10/04
Install and commission the trigger	5/05

3.6 L1 Tracking Trigger Summary and Conclusions

Based upon current simulation results, it is clear that the L1 CTT needs to be upgraded in order to maintain the desired triggering capabilities as a result of the anticipated Run 2b luminosity increases. Because of the tight timescales and limited resources available to address this particular challenge, significant alterations to the tracking detector installed in the solenoid bore are not considered feasible.

Improving the resolution of the L1 CTT by treating at least some fraction of the CFT axial layers as singlets rather than doublet layers in the L1 trigger should improve the background rejection of an upgraded L1 CTT by a significant amount. Simulation studies that treat the hits from fibers on all axial layers as singlets in the trigger yield improvements in the fake rejection rate by more than a factor of ten.

Our studies conclude that the performance of the Run 2b detector will certainly be enhanced if the FPGAs are upgraded to allow for a significant increment in the number of equations that can be handled. The FPGA upgrade provides a significant enhancement in flexibility of the track finding algorithms that may be implemented, and consequently should be given serious consideration. Different options of doublet and singlet layer combinations or all singlets are being considered. This particular upgrade will facilitate a substantial improvement in the background rejection rate at a moderate cost. Our final decision for the exact algorithm and combination of layers to use will be critically coupled to the FPGA resources available two years from now. With the options of

FPGAs (e.g. Xilinx Virtex II series) available at the moment, this upgrade is technically feasible.

4 Level 1 Calorimeter Trigger

4.1 Goals

The primary focus of Run 2b will be the search for the mechanism of electroweak symmetry breaking, including the search for the Higgs boson, supersymmetry, or other manifestations of new physics at a large mass scale. This program demands the selection of events with particularly large transverse momentum objects. The increase in luminosity (and thus increasing multiple interactions), and the decreased bunch spacing (132ns) for Run 2b will impose heavy loads on the Level 1 (L1) calorimeter trigger. The L1 calorimeter trigger upgrade should provide performance improvements over the Run 2a trigger system to allow increased rejection of backgrounds from QCD jet production, and new tools for recognition of interesting signatures. We envision a variety of improvements, each of which will contribute to a substantial improvement in our ability to control rates at the L1 trigger. In the following sections we describe how the L1 calorimeter trigger upgrade will provide

- An improved capability to correctly assign the calorimeter energy deposits to the correct bunch crossing via digital filtering
- A significantly sharper turn-on for jet triggers, thus reducing the rates
- Improved trigger turn-on for electromagnetic objects
- The ability to make shape and isolation cuts on electromagnetic triggers, and thus reducing rates
- The ability to match tracks to energy deposition in calorimeter trigger towers, leading to reduced rates
- The ability to include the energy in the intercryostat region (ICR) when calculating jet energies and the missing ET
- The ability to add topological triggers which will aid in triggering on specific Higgs final states.

The complete implementation of all these improvements will provide us with the ability to trigger effectively with the calorimeter in the challenging environment of Run 2b.

4.2 Description of Run 2a Calorimeter Electronics

4.2.1 Overview

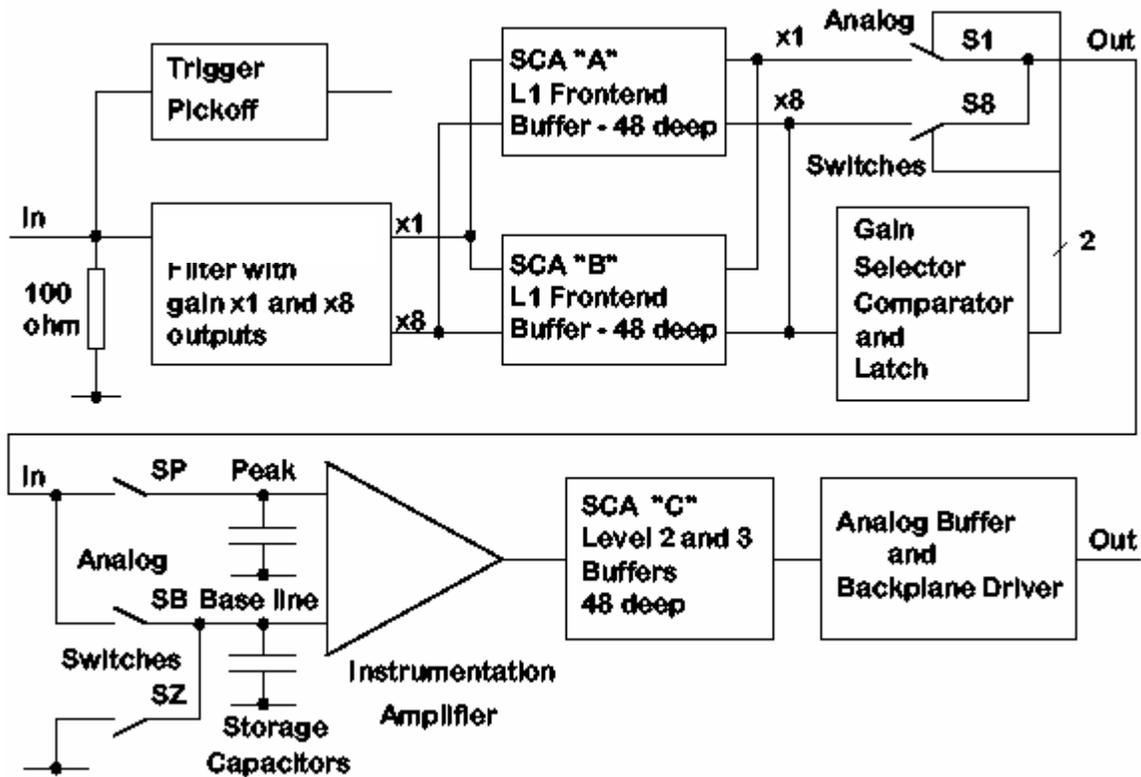


Figure 8. Functional diagram of the BLS system showing the precision readout path and the location of the calorimeter trigger pickoff signal.

The charge from the calorimeter is integrated in the charge sensitive preamplifiers located on the calorimeter. The preamplifier input impedance is matched to the $30\ \Omega$ coaxial cable from the detector (which have been equalized in length), and the preamplifiers have been compensated to match the varying detector capacitances, so as to provide signals that have approximately the same rise time (trace #1 in Figure 9). The fall time for the preamp signals is $15\ \mu\text{s}$. The signals are then transmitted (single ended) on terminated twisted-pair cable to the baseline subtractor cards (BLS) that shape the signal to an approximately unipolar pulse (see Figure 8 for a simple overview). The signal on the trigger path is further differentiated by the trigger pickoff to shorten the pulse width, leading to a risetime of approximately $120\ \text{ns}$ (trace #2 in Figure 9). The signals from the different depths in the electromagnetic and hadronic sections are added with appropriate weights to form the analog trigger tower sums. These analog sums are output to the L1 calorimeter trigger after passing through the trigger sum drivers. The signals are then transported differentially (on pairs of $80\ \Omega$ coaxial cable) $\sim 80\text{m}$ to the L1 calorimeter trigger (the negative side of a differential pair is shown in trace #4 in Figure 9). The key elements of the calorimeter trigger path are described in more detail in the following sections.

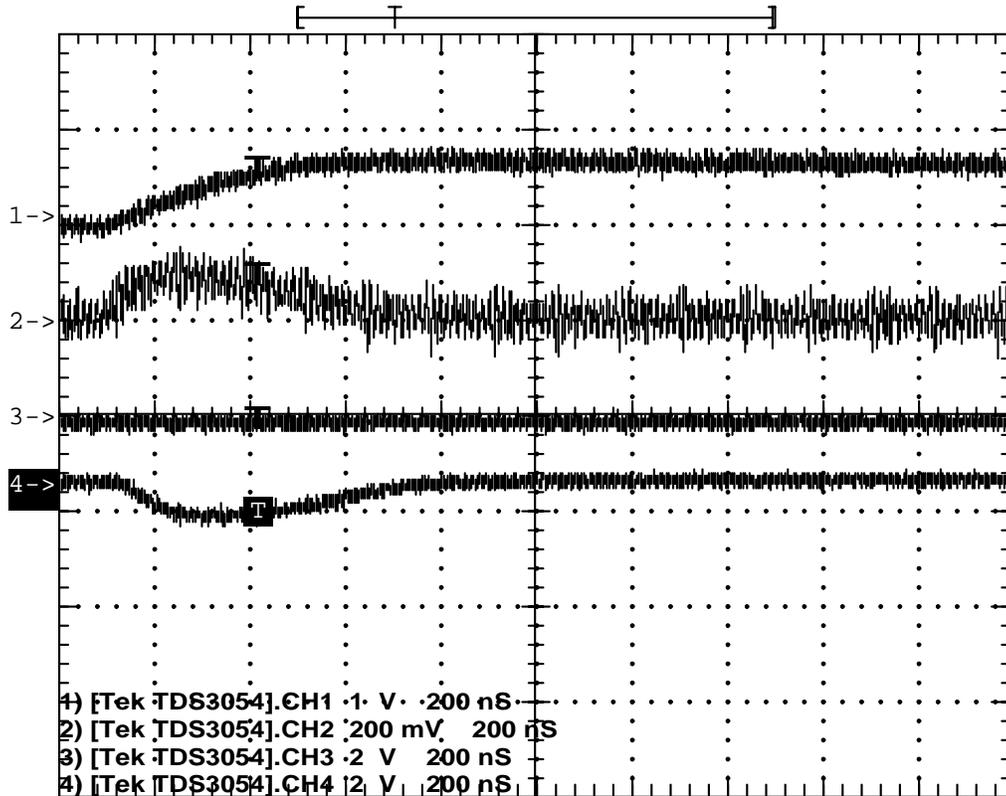


Figure 9. Scope traces for actual detector signals for an EM section. The horizontal scale is 200ns/large division. The top trace (#1, 1V/div) is of a preamp output signal as seen at the input to the BLS. The second trace (#2, 200mV/div) is of the trigger pickoff output on the BLS card (the large noise is due to scope noise pickup, so is not real). The fourth trace (#4, 2V/div) is the negative side of the differential trigger sum driver signal at the BLS that is sent to the L1 calorimeter trigger.

4.2.2 Trigger pickoff

The trigger pickoff captures the preamplifier signal before any shaping. A schematic of the shaping and trigger pickoff hybrid is shown in Figure 10 (the trigger pickoff section is in the upper left of the drawing). The preamplifier signal is differentiated and passed through an emitter follower to attempt to restore the original charge shape (a triangular pulse with a fast rise and a linear fall over 400 ns). This circuitry is located on a small hybrid that plugs into the BLS motherboard. There are 48 such hybrids on a motherboard, and a total of 55,296 for the complete detector.

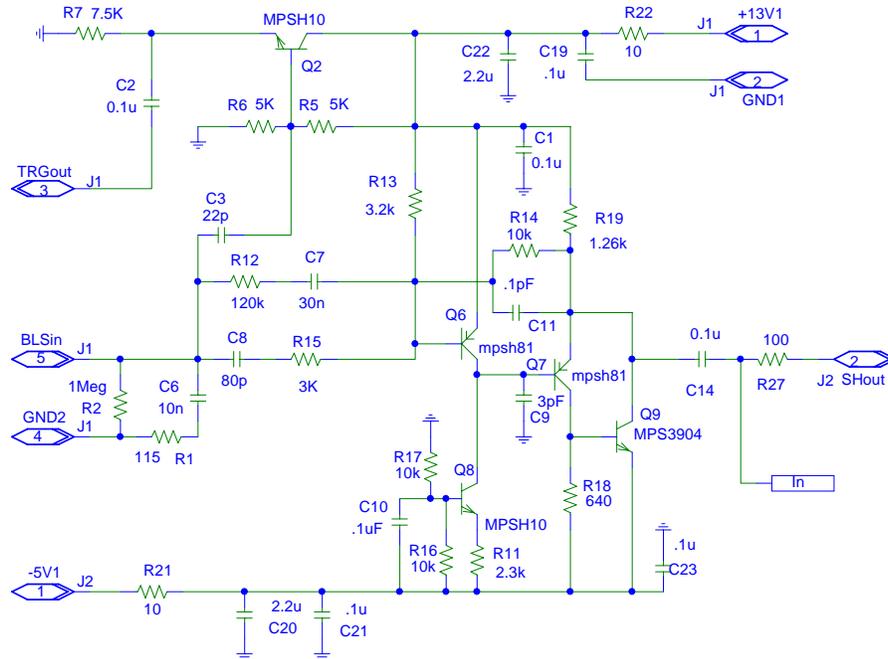


Figure 10. Schematic of the trigger shaper and trigger pickoff (upper left of picture). Pin 5 is the input, pin 3 is the trigger pickoff output, and pin 2 is the shaped precision signal output.

4.2.3 Trigger summers

The trigger pickoff signals for EM and HAD sections in individual towers (note these are not the larger trigger towers) are routed on the BLS board to another hybrid plug-in that forms the analog sums with the correct weighting factors for the different radial depth signals that form a single tower. The weighting is performed using appropriate input resistors to the summing junction of the discrete amplifier. A schematic for this small hybrid circuit is shown in Figure 11.

A single 48 channel BLS board has 8 trigger summer hybrids (4 EM towers and 4 HAD towers). There are a total of 9,216 hybrid trigger summers made up of 75 species. Since they are relatively easy to replace, changes to the weighting schemes can be considered. Recall, however, that access to the BLS cards themselves requires access to the detector as they are located in the area directly beneath the detector, which is inaccessible while beam is circulating.

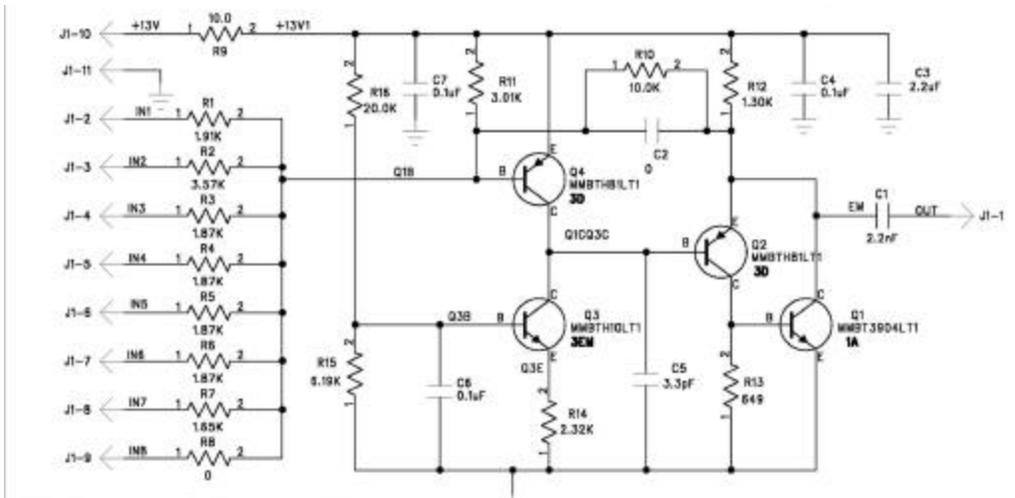


Figure 11. Schematic of the trigger summer hybrid. Up to 8 inputs from the various layers in a single tower can be summed with varying gains determined by the resistors to the summing junction (shown at left).

4.2.4 Trigger sum driver

The outputs of the 4 EM trigger summers and the 4 HAD trigger summers on a single BLS board are summed separately (except at high η) once more by the trigger sum driver circuit (see the schematic in Figure 12) where a final overall gain can be introduced. This circuit is also a hybrid plug-in to the BLS board and is thus easily replaceable if necessary (with the same access restrictions discussed for the trigger summers). In addition the driver is capable of driving the coaxial lines to the L1 Calorimeter trigger. There are a total of 2,560 such drivers in 8 species (although most are of two types).

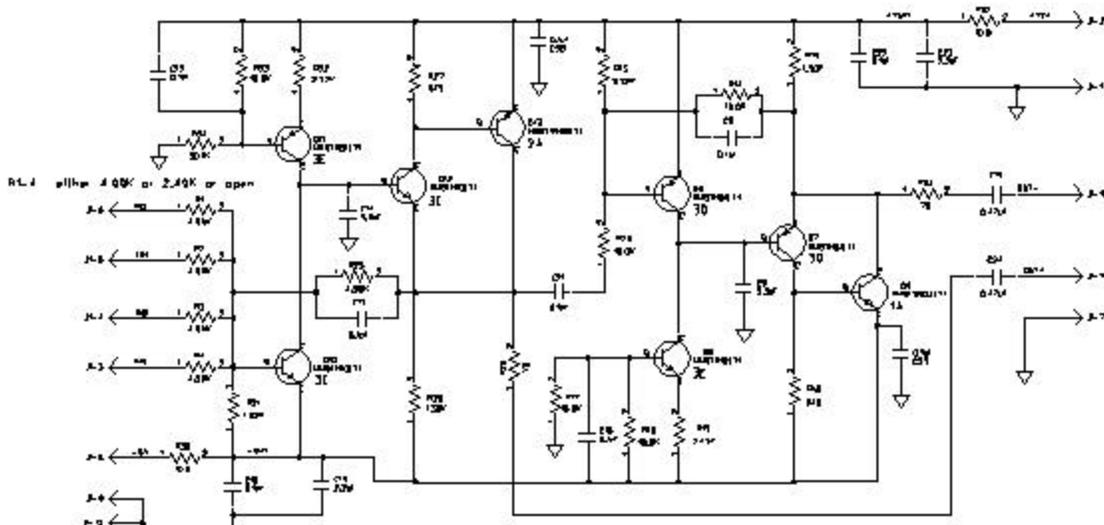


Figure 12. Schematic of the trigger sum driver hybrid. This circuit sums the outputs of up to 4 trigger summer outputs of the type shown in Figure 11.

4.2.5 Signal transmission, cable dispersion

The signals from the trigger driver circuits are transmitted differentially on two separate miniature coax (0.1") cables. The signal characteristics for these cables are significantly better than standard RG174 cable. However first indications are that the signals seen at the end of these cables at the input to the L1 calorimeter trigger are somewhat slower than expected (an oscilloscope trace of such a signal is shown in Figure 13 for EM and Figure 14 for HAD). The cause of the deviation from expectations is not presently known and is under investigation. It is possible that the signal dispersion in these coaxial cables is worse than expected. In any case, we must deal with these pulses that are over 400ns wide (FWHM) and thus span a few 132ns bunch crossings. While there are possible intermediate solutions to deal with this signal shape for 132ns bunch crossings, the most effective treatment calls for further processing of the signal through digital filtering to extract the proper bunch crossing. This option is described in more detail in later sections.

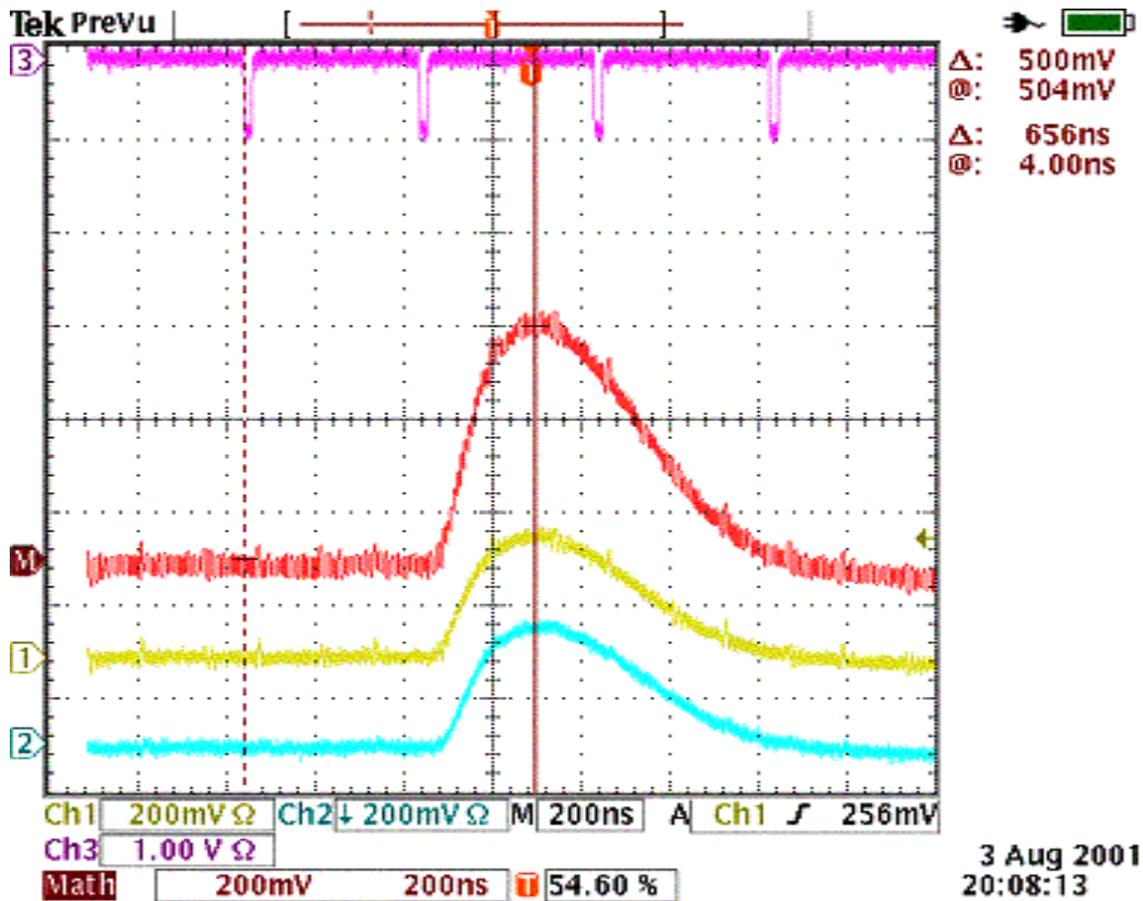


Figure 13. Actual traces of EM trigger tower ($\eta=+1$, $\phi=17$) data from the trigger sum driver signal as measured at the input to the L1 calorimeter trigger. The top trace (#3) shows the time of the beam crossings (396ns). The second trace (M) shows the addition of the two differential signals after inversion of the negative one. The third trace (#1) is the positive side of the differential pair. The fourth trace (#2) is the inverted trace for the negative side of the differential pair.

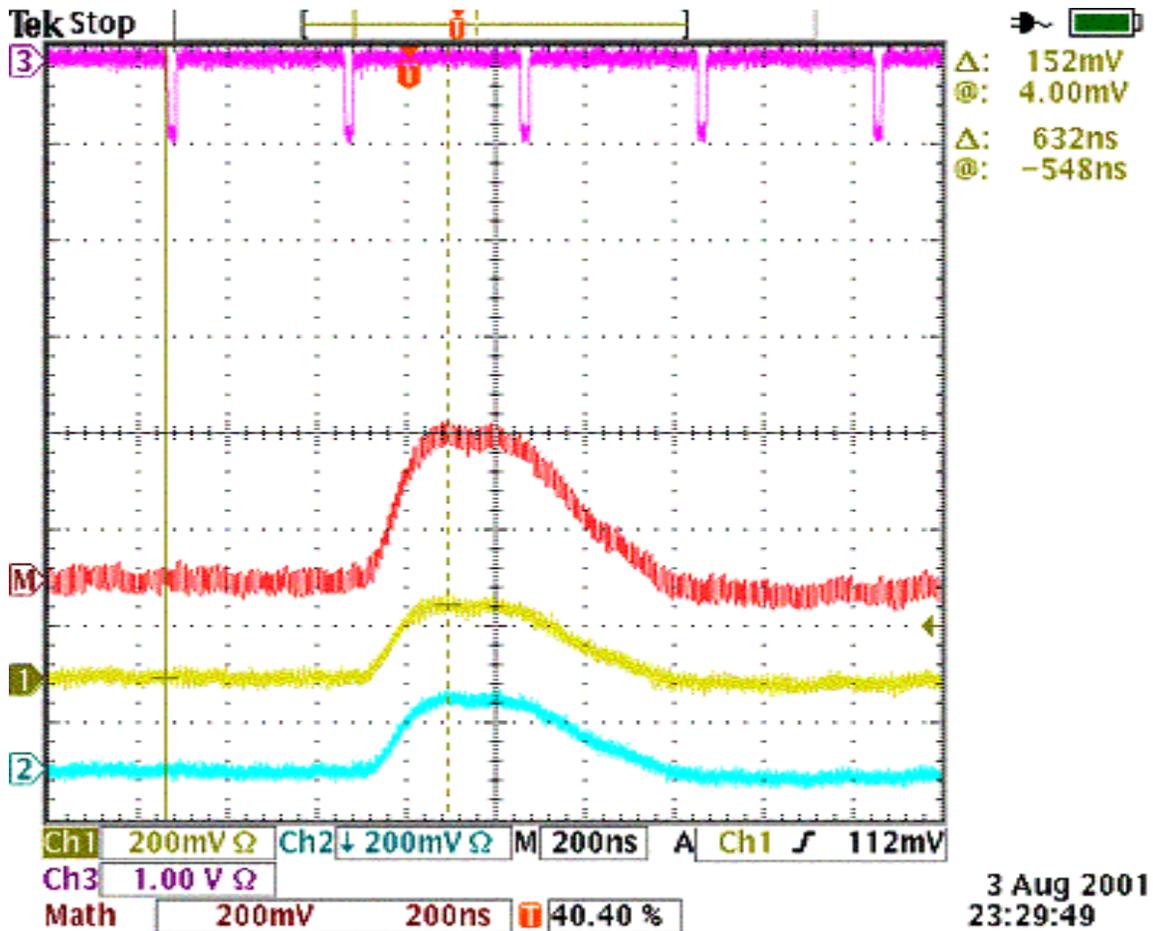


Figure 14. Actual traces of HAD trigger tower ($\eta=+1$, $\phi=17$) data from the trigger sum driver signal as measured at the input to the L1 calorimeter trigger. The top trace (#3) shows the time of the beam crossings (396ns). The second trace (M) shows the addition of the two differential signals after inversion of the negative one. The third trace (#1) is the positive side of the differential pair. The fourth trace (#2) is the inverted trace for the negative side of the differential pair.

4.3 Description of Current L1 Calorimeter Trigger

4.3.1 Overview

The DØ uranium-liquid argon calorimeter is constructed of projective towers covering the full 2π in the azimuthal angle, ϕ , and approximately 8 units of pseudo-rapidity, η . There are four subdivisions along the shower development axis in the electromagnetic (EM) section, and four or five in the hadronic (H) section. The hadronic calorimeter is divided into the fine hadronic (FH) section with relatively thin uranium absorber, and the backing coarse (CH) section. In the intercryostat region $0.8 < |\eta| < 1.6$ where the relatively thick cryostat walls give extra material for shower development, a scintillator based intercryostat detector (ICD) and extra 'massless gap' (MG) liquid argon gaps without associated absorber are located.

The calorimeter tower segmentation in $\eta \times \phi$ is 0.1×0.1 , which results in towers whose transverse size is larger than the expected sizes of EM showers but, considerably smaller than typical sizes of jets.

As a compromise, for triggering purposes, we add four adjacent calorimeter towers to form trigger towers (TT) with a segmentation of 0.2×0.2 in $\eta \times \phi$. This yields an array that is 40 in η and 32 in ϕ or a total of 1,280 EM and 1,280 H tower energies as inputs to the L1 calorimeter trigger.

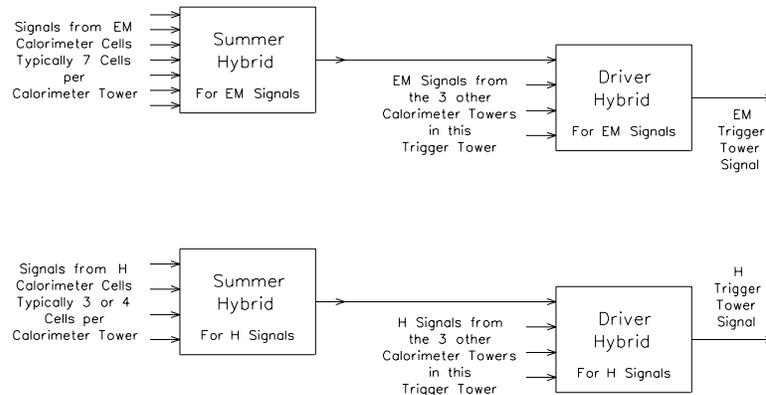


Figure 15. Trigger tower formation.

The analog summation of the signals from the various calorimeter cells in a trigger tower into the EM and H TT signals takes place as described on page 38. This arrangement for summing the calorimeter cells into trigger towers is shown schematically in Figure 15.

Long ribbons of coaxial cable route the 1280 EM and H analog trigger tower signals from the detector platform through the shield wall and then into the first floor of the moving counting house (MCH) where the Level 1 calorimeter trigger is located. The first step in the Level 1 calorimeter trigger is to scale these signals to represent the E_T of the energy deposited in each trigger tower and then to digitize these signals at the beam-crossing rate (132ns) with fast analog to digital converters. The digital output of these 2560 converters is used by the subsequent trigger logic to form the Level 1 calorimeter trigger decision for each beam crossing. The converter outputs are also buffered and made available for readout to both the Level 2 Trigger system and the Level 3 Trigger DAQ system.

The digital logic used in the Level 1 Calorimeter Trigger is arranged in a "pipe-lined" design. Each step in the pipe-line is completed at the beam crossing rate and the length of the pipe-line is less than the maximum $D\phi$ Level 1 trigger latency for Run 2a which is $3.3 \mu\text{sec}$ (driven by the calorimeter shaping times, cables lengths, drift times etc). This digital logic is used to calculate a number of quantities that are useful in triggering on specific physics processes. Among these are quantities such as the total transverse energy and the missing transverse energy, which we will designate as "global" and information relating to

"local" or cluster aspects of the energy deposits in the calorimeter. The latter would include the number of EM and H-like clusters exceeding a set of programmable thresholds.

4.3.2 Global Triggers

Interesting global quantities include:

the total transverse energies:

$$\left(E_T^{EM}\right)_{Total} = \sum_{i=1}^{1280} \left(E_T^{EM}\right)_i$$

$$\left(E_T^H\right)_{Total} = \sum_{i=1}^{1280} \left(E_T^H\right)_i$$

and

$$\left(E_T\right)_{Total} = \left(E_T^{EM}\right)_{Total} + \left(E_T^H\right)_{Total}$$

the missing transverse energy:

$$Mp_T = \sqrt{(E_x^2 + E_y^2)}$$

where:

$$E_x = \sum_{i=1}^{1280} \left[\left(E_T^{EM}\right)_i + \left(E_T^H\right)_i \right] \cos(\phi_i)$$

and

$$E_y = \sum_{i=1}^{1280} \left[\left(E_T^{EM}\right)_i + \left(E_T^H\right)_i \right] \sin(\phi_i)$$

Any of these global quantities can be used in constructing triggers. Each quantity is compared to a number of thresholds and the result of these comparisons is passed to the Trigger Framework where up to 128 different Level 1 triggers can be formed.

4.3.3 Cluster Triggers

The DØ detector was designed with the intent of optimizing the detection of leptons, quarks and gluons. Electrons and photons will manifest themselves as localized EM energy deposits and the quarks and gluons as hadron-like clusters.

Energy deposited in a Trigger tower is called EM-like if it exceeds one of the EM E_T thresholds and if it is not vetoed by the H energy behind it. Up to four EM E_T thresholds and their associated H veto thresholds may be programmed for each of the 1280 trigger towers. Hadronic energy deposits are detected by calculating the EM $E_T + H E_T$ of each Trigger tower and comparing each of these 1280 sums to four programmable thresholds.

The number of Trigger towers exceeding each of the four EM thresholds (and not vetoed by the H energy behind it) is calculated and these four counts are compared to a number of count thresholds. The same is done for the four EM $E_T + H E_T$ thresholds. The results of these count comparisons on the number of Trigger towers over each threshold are sent to the Trigger Framework where they are used to construct the Level 1 Triggers.

4.3.4 Hardware Implementation

4.3.4.1 Front End Cards

The analog signals from the calorimeter, representing energies, arrive at the Calorimeter Trigger over coaxial differential signal cables and are connected to the analog front end section of a Calorimeter Trigger Front End Card (CTFE). A schematic diagram of one of the four cells of this card is shown in Figure 16.

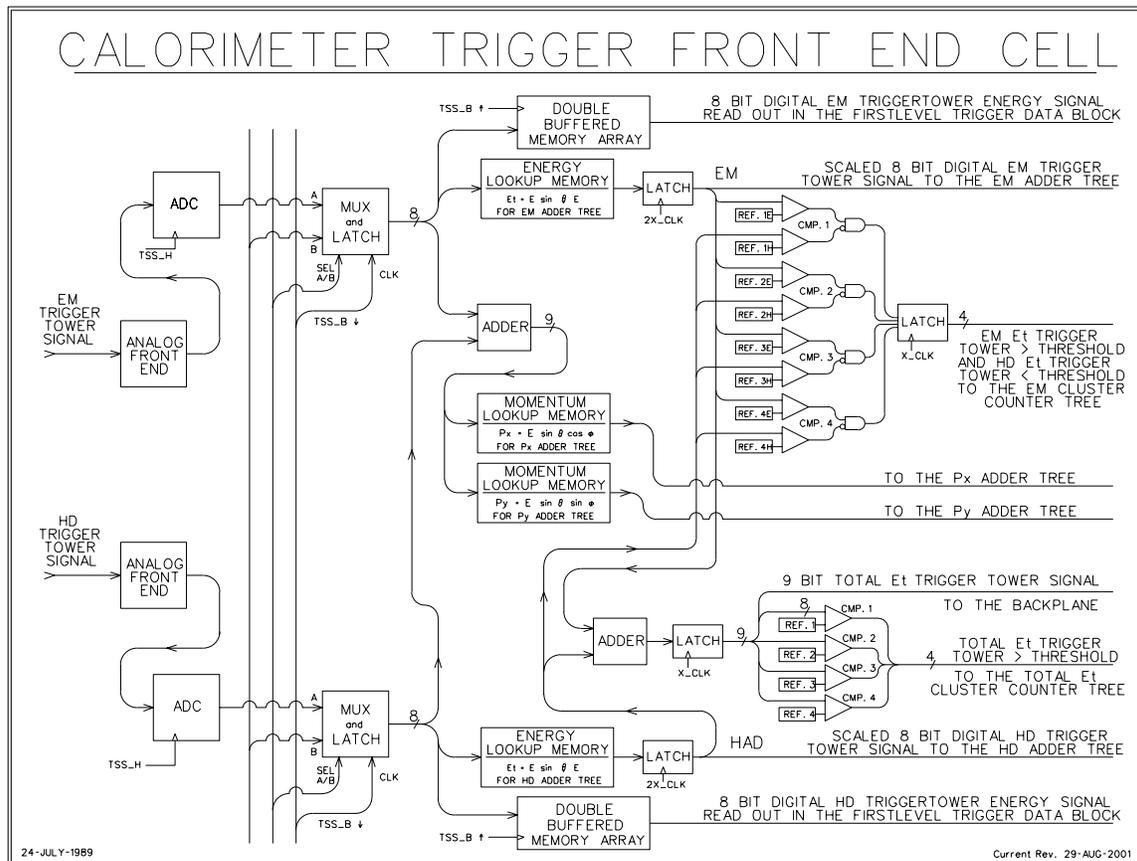


Figure 16. Calorimeter trigger front end cell (CTFE).

The front-end section contains a differential line receiver and scales the energy signal to its transverse component using a programmable gain stage. The front end also contains digital to analog circuitry for adding a positive bias to the tower energies in accord with downloaded values.

Immediately after the analog front end, the EM or H signal is turned into an 8 bit number by fast (20 ns from input to output) FADC's. With our current choice of 0.25 GeV least count this gives a maximum of 64 GeV for the single tower transverse energy contribution.

The data are synchronized at this point by being clocked into latches and then follow three distinct parallel paths. One of these paths leads to a pipeline register for digital storage to await the L1 trigger decision and subsequent readout to the Level 2 Trigger system and the Level 3 Trigger DAQ system.

On the other two paths, each 8-bit signal becomes the address to a look up memory. The content of the memory at a specified address in one case is the transverse energy with all necessary corrections such as lower energy requirements etc. In the other case, the EM + H transverse energies are first added and then subjected to two look-ups to return the two Cartesian components of the transverse energy for use in constructing MP_T . The inherent flexibility of this scheme has a number of advantages: any energy dependent quantity can be generated, individual channels can be corrected or turned off at this level and arbitrary individual tower efficiencies can be accommodated.

The CTFE card performs the function of adding the E_T 's of the four individual cells for both the EM and H sections and passing the resulting sums onto the Adder Trees. In addition it tests each of the EM and EM+H tower transverse energies against the four discrete thresholds and increments the appropriate counts. These counts are passed onto the EM cluster counter trees and the total E_T counter trees, respectively.

4.3.4.2 Adder and Counter Trees

The adder and counter trees are similar in that they both quickly add a large number of items to form one sum. At the end of each tree the sum is compared to a number of thresholds and the result this comparison is passed to the Trigger Framework. A typical adder tree is shown in Figure 17.

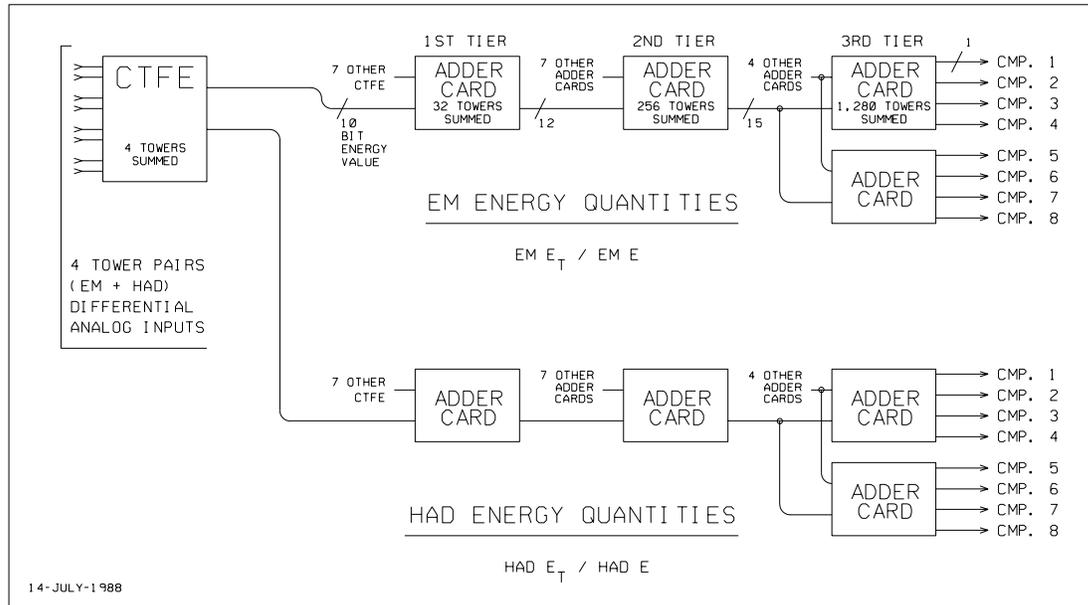


Figure 17. Adder tree for EM and Had quantities.

4.3.5 Physical Layout

Ten racks are used to hold the Level 1 Calorimeter Trigger, which is located in the first floor moving counting house. The lower section of each rack contains the CTFE cards for 128 Trigger towers (all 32 ϕ 's for four consecutive η 's). The upper section of each rack contains a component of one of the Adder or Counter Trees.

4.4 Motivations for Upgrading the Current System

The current L1 calorimeter trigger, which was built in 1988, and was used in Run 1 and Run 2a, has a number of features that limit its usefulness in Run 2b.

- 1) Trigger tower analog signals have rise times that are slightly longer than the 132 ns bunch spacing foreseen in Run 2b. The fall time of the signals, ~ 400 ns, is also significantly longer than the time between collisions. This makes it impossible for the current L1 calorimeter trigger to reliably assign calorimeter energy to the correct beam crossing, resulting in L1 trigger accepts being generated for the wrong beam crossing. Since information about the correct (interesting) beam crossing would be lost in these cases, finding a solution to this problem is imperative.
- 2) The fixed size trigger towers used in the current L1 calorimeter trigger are much smaller than the typical lateral size of a jet, resulting in extremely slow "turn-on" curves for jet and electron triggers. For example, a 6 GeV single tower threshold becomes $\sim 100\%$ efficient only for jets with transverse energies greater than 60 GeV. This poor resolution, convoluted with the steeply falling jet E_T spectrum, results in an overwhelming background of low energy jets passing a given threshold at high luminosity.

- 3) Total E_T and missing E_T resolution is significantly degraded because signals from the ICR detectors are not included in the trigger sums.

To run efficiently under Run 2b conditions, the problem of triggering on the wrong bunch crossing must be resolved. Beyond that, the limited clustering capabilities in the current system result in unacceptably high rates for the triggers needed to discover the Higgs and pursue the rest of the $D\bar{D}$ physics program. Each of these issues is discussed in more detail in the rest of this section, while our solutions are presented in the following sections.

4.4.1 Bunch Crossing mis-Identification

Because the width of the shaped analog TT signals is >400 ns, the current system will experience difficulties, as mentioned previously, when the spacing between bunches in the Tevatron is reduced from 396 ns to 132 ns. The main issue here is identifying energy deposited in the calorimeter with the correct bunch crossing. This is illustrated in Figure 13 and Figure 14, which show representative TT analog signals. The calorimeter readout timing is set such that the peak of the analog signal (~ 200 ns after it begins to rise) corresponds to the bunch crossing, n , where the relevant energy was deposited. For large amplitude signals, however, one or more of the TT E_T thresholds may be crossed early enough on the signal's rise to be associated with bunch crossing $n-1$. Additionally, the signal may not fall below threshold for several bunch crossings ($n+1, n+2, \dots$) after the signal peaks due to the long fall time. Because no events are accepted after an L1 trigger accept is issued until the silicon detector is read out, triggering on bunch crossing $n-1$ would cause $D\bar{D}$ to lose the interesting event at bunch crossing n in such a case.

4.4.2 Background Rates and Rejection

4.4.2.1 *Simulation of the Current System*

In order to assess the physics performance of the present L1 calorimeter trigger, the following simulation is used. The jet performance is studied using a Monte-Carlo sample of QCD events (PYTHIA, with parton p_T cuts of 5, 10, 20, 40 GeV and 0.5 overlaid minimum bias events). A cone algorithm with a radius of 0.4 in $\eta \times \phi$ is applied to the generated stable hadrons in order to find the generated jets and their direction. The direction of each generated jet is extrapolated to the calorimeter surface; leading to the "center TT" hit by the jet. The highest E_T TT in a 3x3 trigger tower region (which is 0.6×0.6 in $\eta \times \phi$ space) around this center is then used to define the "trigger E_T " corresponding to the jet.

4.4.2.2 *Energy measurement and turn-on curves*

In the present L1 calorimeter trigger, the trigger towers are constructed using fixed $\eta \times \phi$ towers. Thus we expect that a trigger tower only captures a small fraction of the total jet energy since the size of the 0.2×0.2 trigger towers is small compared to the spatial extent of hadronic showers. This is illustrated in Figure 18, which shows, for simulated 40 GeV E_T jet events, the ratio of the E_T observed by the trigger to the generated E_T . It can be seen in Figure 18 that this transverse energy is only 25% of the jet E_T on average. Therefore we must use

low jet trigger thresholds if we are to be efficient even for relatively high energy jets. Moreover the trigger E_T has poor resolution, as can be seen in Figure 18. As a result, the trigger efficiency (the efficiency for having at least one TT with E_T above a given threshold) rises only slowly with increasing jet E_T , as shown in the turn-on curves in Figure 19. A similar effect occurs for the EM triggers as well; even though a typical EM shower can be reasonably well contained within a TT, often the impact point of an electron or photon is near a boundary between TTs.

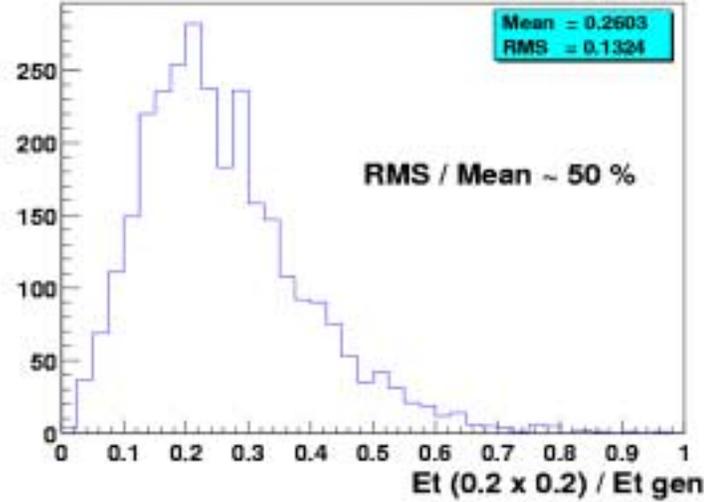


Figure 18. Ratio of the trigger E_T to the transverse energy of the generated jet. Only jets with $E_T \approx 40$ GeV are used in this figure.

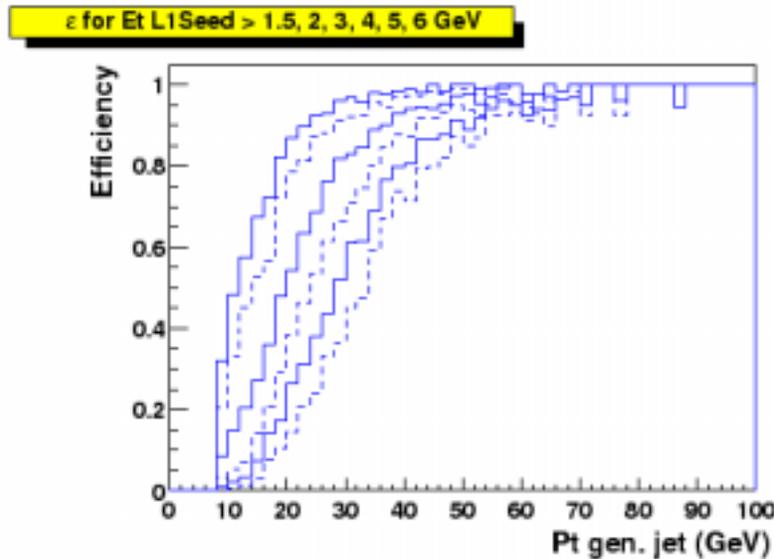


Figure 19. Trigger efficiency as a function of the transverse energy of the generated jet. The curves correspond to thresholds of 1.5, 2, 3, 4, 5 and 6 GeV (respectively from left to right).

4.4.2.3 Trigger rates

The trigger E_T resolution, convoluted with the steeply falling p_T spectrum of QCD events, leads to, on average, the “promotion” of events to larger E_T ’s than

the actual E_T . The number of QCD events which pass the L1 trigger is thus larger than what it would be with an ideal trigger E_T measurement. Due to the very large cross-section for QCD processes, this results in large trigger rates⁵. For example, as shown in Figure 20, an inclusive unprescaled high E_T jet trigger, requiring at least one TT above a threshold defined such that the efficiency for 40 GeV jets is 90%, would yield a rate for passing the L1 calorimeter trigger of at least 10 kHz at $2 \times 10^{32} \text{ cm}^2 \text{ s}^{-1}$. Maintaining this rate below 1 kHz would imply an efficiency on such high E_T jets of only 60%. Trigger rates increase faster than the luminosity due to the increasing mean number of interactions per bunch crossing. Trigger rates are shown in Figure 21 as a function of the mean number of minimum bias events which pile up on the high p_T interaction. These are shown for two multi-jet triggers: the first requiring at least two TT above 5 GeV (indicated as CJT(2,5)); the second requiring at least two TT above 5 GeV and at least one TT above 7 GeV (indicated as CJT(1,7)*CJT(2,5)). These triggers correspond to reasonable requirements for high p_T jets because, as can be seen in Figure 20, a threshold of 5 GeV leads, for 40 GeV jets, to an 80 % efficiency. The rates in Figure 21 are shown for a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. For the higher luminosity of $5 \times 10^{32} \text{ cm}^2 \text{ s}^{-1}$ expected in Run 2b, the L1 bandwidth of 5kHz could be saturated by such dijet conditions alone, unless large prescale factors are applied.

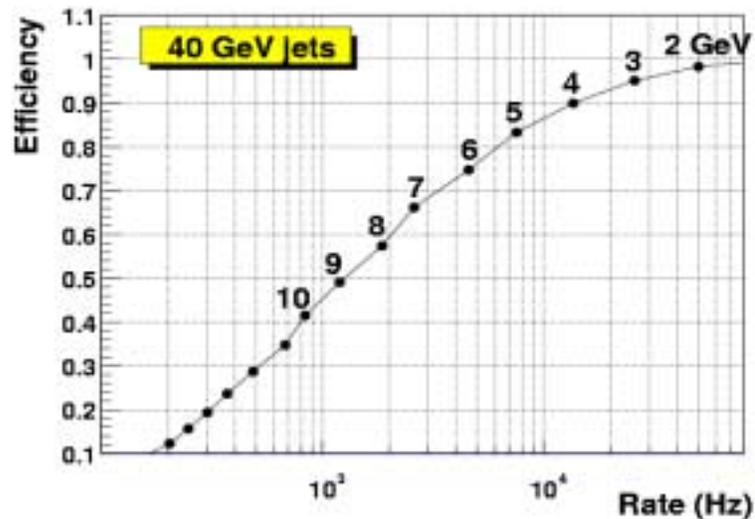


Figure 20. The efficiency to trigger on 40 GeV jets as a function of the inclusive trigger rate when one TT above a given threshold is required. Each dot corresponds to a different threshold (in steps of 1 GeV), as indicated. The luminosity is $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

⁵ These rates are estimated here from samples of PYTHIA QCD events with parton $p_T > 2 \text{ GeV}$, passed through a simulation of the trigger response.

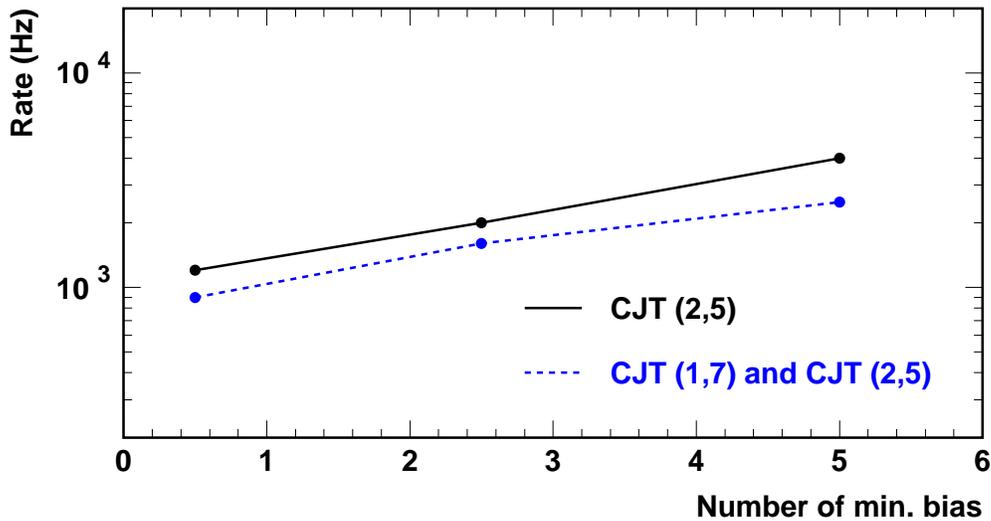


Figure 21. The inclusive trigger rate as a function of the mean number of minimum bias events overlaid on the high p_T interaction. The rates are shown for two di-jet trigger conditions corresponding to two TTs above 5 GeV (CJT(2,5)) and two TTs with above 5GeV and at least one above 7 GeV (CJT(1,7)*CJT(2,5)). The luminosity is $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

A more exhaustive study of the evolution of the L1 trigger rate with increasing luminosity has been carried out⁶. In that document a possible trigger menu was considered, in which ~75 % of the L1 bandwidth is used by multijet triggers. The results are shown in Table 7. It can be seen that, at the luminosity foreseen for Run 2b (corresponding to the 4th row), the trigger rates should be reduced by at least a factor of four in order to maintain a reasonably small dead time. We note that the need to preserve jet triggers is required by some of the Higgs boson physics (see for example section 4.7).

Table 7. The overall level 1 trigger rates as a function of luminosity.

Luminosity	High Pt L1 rate (Hz)	Total L1 rate (Hz)
$1 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$	1,700	5,000
$2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$	4,300	9,500
$5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$	6,500	20,000

4.4.3 Conclusions/implications for high luminosity

Clearly, the bunch crossing mis-identification problem must be resolved for Run 2b or the L1 calorimeter trigger will cease to be effective. The physics studies presented above also show that there is a need to significantly improve the rejection of the L1 calorimeter trigger (while maintaining good efficiency) if we

⁶ B. Bhattacharjee, "Transverse energy and cone size dependence of the inclusive jet cross section at center of mass energy of 1.8 TeV", PhD Thesis, Delhi University.

are to access the physics of Run 2b. One obvious way to help achieve this is to migrate the tools used at L2 (from Run 2a) into L1. In particular, the ability to trigger on “objects” such as electromagnetic showers and jets would help significantly. The “clustering” of TT’s at L1, could reduce the trigger rates by a factor 2 to 4 as will be shown later. The principal reason for this gain comes from the improvement in the quality of the energy cut, when applied to a cluster of trigger towers. Transferring to level 1 some of the functions that currently belong to level 2 would also permit the introduction of new selection algorithms at the L1 trigger level. So while it is clear that there are additional gains to be made through EM trigger tower shape cuts and missing E_T filtering, they will require further study to quantify the specific gains. These studies remain to be done.

From a conceptual viewpoint, an important consequence of selecting physics “objects” at level 1 is that it allows a more “inclusive” and hence less biased selection of signatures for the more complicated decays to be studied in Run 2b. Thus we expect that the trigger menus will become simpler and, above all, less sensitive to biases arising from the combinations of primary objects.

4.4.4 Overview of Options for Improvement

We have examined various possibilities for the changes necessary to address the incorrect bunch crossing assignment problem and the trigger energy resolution problem. The age and architecture of the current system prohibit an incremental solution to these issues. We therefore propose to design and build an entirely new L1 calorimeter trigger system, which will replace all elements of the current trigger downstream of the BLS cables. The status of the design of this new system is given in the following sections. A partial list of improvements provided by the new system is given below.

- Necessary hardware improvements in filtering to allow proper triggering on the correct bunch crossing.
- Implementation of a “sliding window” algorithm for jets and electrons.
- The addition of presently unused calorimeter energy information from the intercryostat detector (ICD) and massless gaps (MG) in the L1 trigger.
- Optimization of trigger tower thresholds.
- Topological cuts.
- The ability to better correlate tracks from the fiber tracker to calorimeter clusters.

Studies of these improvements are discussed in the following sections with the exception of the correlation between tracks and calorimeter clusters, which is described in the Cal-Track Match system section. The status of the implementation of the new system is then outlined in the remaining parts of the chapter.

4.5 Digital Filtering

Digital filtering offers a way to reduce the effect of unwanted triggers due to collisions in close proximity to the desired trigger.

4.5.1 Concept & physics implications

The pulse shape, and particularly the rise time, of the trigger pickoff signal is not optimized for 132ns beam bunch crossing operation (see Figure 13 and Figure 14). Since the trigger pickoff pulse width significantly exceeds the 132ns bunch spacing time of Run 2b, the ability to correctly identify the correct trigger bunch crossing is compromised. There may be intermediate solutions to address this problem at the lower luminosities, but a long-term solution must be developed. This could be done by means of an analog filter with shorter shaping, but this is only achieved with a further loss in signal. A digital filter is a better solution because it is much more flexible for a similar cost.

The trigger pickoff signal is at the end of the calorimeter electronic chain described above. The ideal energy deposition shape is a "saw-tooth" pulse (infinitely fast rise and a linear ~400ns fall) from energy deposited in the cells of the calorimeter at each beam crossing. This is modified by the transfer function of the electronics. The inverse transfer function will transform the pickoff signal back to original energy deposition pulse shape. Digital filtering would be implemented at this stage. The inverse function can be implemented by a FIR (Finite Impulse Response) digital filter. In the presence of noise, the digital filter offers an additional advantage: one can use the theory of optimal filtering to minimize the noise contribution.

In order to define the exact form of a digital filter best suited to the task, a measurement of noise in the trigger pickoff signals is needed. As such measurements become available, a refined design will be undertaken.

4.5.2 Pileup rejection

Two different "pile-up" effects arise with increasing luminosity, the first is due to extra collisions in the crossing of interest (and thus unavoidable), and the second is due to collisions in neighboring crossings that contribute to the crossing of interest because of signal shapes.

In the first case, we find that as the luminosity increases, then for each triggered beam crossing there are several minimum bias events that appear in that same beam crossing. The number of such additional events is Poisson distributed with a mean proportional to the luminosity. The energy added by these events has a distribution close to that of a double exponential (Laplacian). It is possible to minimize the contribution of this noise by using an appropriate digital filter (Matched Median Filter).

In the second case, because the width of the trigger pickoff signal extends over several beam crossing (6 at 132ns on the positive side of the signal), then when two such pulses are close in time, there is some overlap and thus the shape of the pickoff signal becomes more complicated than that of a single isolated pulse. The inverse filter will extract from this signal the two original

pulses. Consequently, the problems caused by overlapping pulses are minimized if one uses digital filtering.

4.5.3 Input data and simulation tools

A series of measurements on several EM and HAD trigger pickoff channels was performed to provide the necessary input to digital filter algorithm studies. Oscilloscope traces and raw data files have been recorded. A chain of programs has been developed to generate training sets based on measured pulses, simulate the analog to digital conversion stage, study digital filter algorithms and compare results with the expected outputs. All programs are standalone and use ASCII files for input and output to provide an increased flexibility and the widest choice of tools for visualization and post-processing.

A typical pulse on an EM channel is shown on the left side of Figure 22. A 4096-point Fast Fourier Transform of this signal is shown on the right side of Figure 22 (the DC component was removed for clarity). It can be seen that most of the energy of the signal is located in frequency components below ~ 10 MHz. In order to remove the high frequency noise that can be seen, we suggest that an analog low-pass filter is placed on each channel before the analog to digital converter. Different filters were investigated by numerical simulation. As shown on the figure, a 2nd order low-pass Butterworth filter with a cutoff frequency of 7.57 MHz seems adequate to remove high frequency oscillations on the signal while preserving the shape of its envelope. Such low-pass filter will avoid the potential problems of spectrum aliasing in the digital domain.

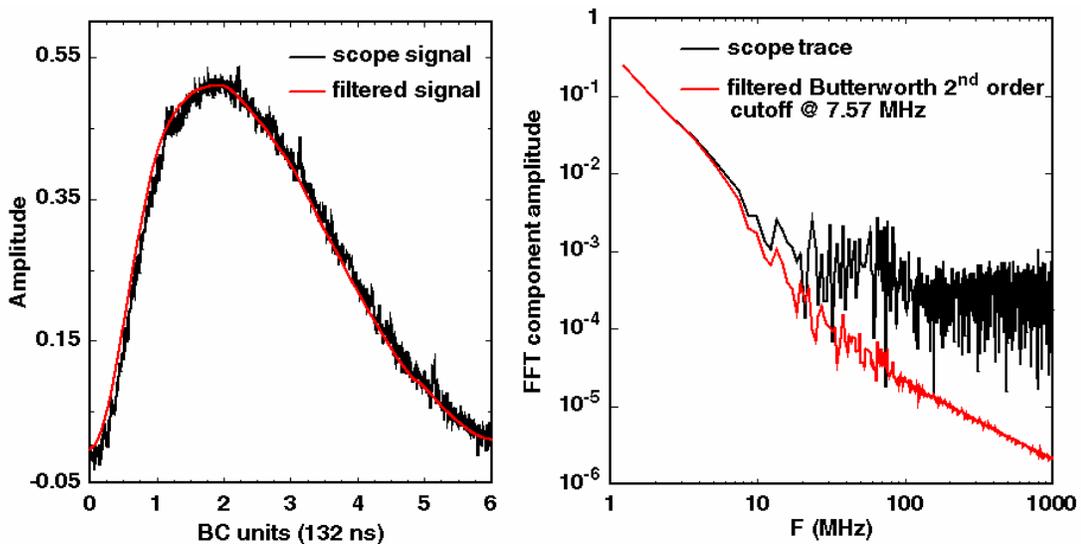


Figure 22. Scope trace of a typical EM pulse and corresponding spectrum. Pulse shape and spectrum after an anti-aliasing filter.

4.5.4 Algorithm evaluation parameters

In order to investigate and compare different options for the digital filter algorithm, several criteria have been defined. A first set is related to the features of the algorithm itself: irreducible latency, number of parameters to adjust and

channel dependency, procedure for parameter determination and tuning, operating frequency, behavior under digital and analog saturation... A second family of criteria relates to the quality of the algorithm: precision on the estimated Et value for the beam-crossing of interest and residual error on adjacent beam-crossings, time/amplitude resolution, ability to separate pulses close in time, probability of having pulses undetected or assigned to the wrong beam-crossing. Several criteria are related to the sensitivity of an algorithm: robustness against electrical noise, ability to reject pileup noise, sensitivity to signal phase and jitter with respect to a reference clock, dependence on pulse shape distortion, performance with limited precision arithmetic, influence of coefficient truncation and input quantization, etc. The last set of comparison criteria concerns implementation: amount of logic required and operating speed of the various components, effective latency.

Defining and selecting the algorithm that will lead to the best trade-off between all these – sometimes contradictory – criteria is not straightforward. Some compromises on performance and functionality will necessarily be done in order to fit in the tight, non-extensible, latency budget that can be devoted to this task while keeping the system simple enough to be implemented with modern, industrial electronic devices at an affordable cost. Algorithm definition and test by computer simulation, electronic hardware simulation and validation with a prototype card connected to real detector signals are among the necessary steps for a successful definition of the digital filter.

4.5.5 Algorithm studied

At present, three types of algorithms have been proposed and investigated. These are:

- A Finite Impulse Response (FIR) deconvolution filter;
- A peak detector followed by a weighed moving average filter;
- A matched filter followed by a peak detector.

4.5.5.1 *FIR deconvolution*

The deconvolution filter is designed to implement the inverse transfer function of the complete calorimeter pickoff chain. When driven with a typical trigger pickoff saw-tooth shaped pulse, the output of the filter is the original pulse. In order to produce a meaningful output for each beam crossing, the filter must have a bandwidth equal at least to the beam crossing frequency. Hence, input samples must be acquired at least at twice the beam-crossing rate (Shannon's sampling theorem). However, only one output value per beam crossing is computed. Coefficient count must be sufficient to ensure that the output of the filter remains null during the falling edge of the input pickoff signal. The determination of coefficients can be made using a set of input training samples that include noise, pileup, time jitter and pulse shape distortion. The differences between the expected values and the actual filter outputs are accumulated and a least mean square minimization is performed to determine the set of coefficients that provide the optimum solution.

The deconvolution filter is linear; parameter tuning can lead to the optimum linear solution for the set of input constraints that is given. This filter performs well to separate pulses close in time as illustrated in Figure 23. A series of pulses of constant height separated by a decreasing amount of time were generated and a simulated trigger pickoff signal was calculated. It can be seen in Figure 23 that the deconvolution FIR filter is able to identify correctly adjacent pulses, even when these occur on two consecutive beam-crossings (i.e. 132 ns apart). However, a non-null residual error is present for some beam-crossings.

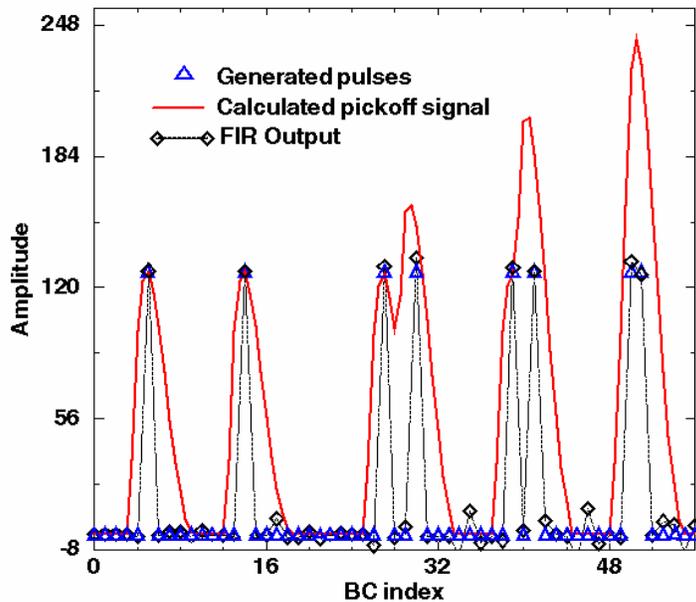


Figure 23. Deconvolution of pulses overlapping in time. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 12-tap FIR is used; 32-bit floating-point arithmetic is used for coefficients and computations.

Various tests were performed to investigate the behavior and the performance of the FIR deconvolution algorithm. An example is shown in Figure 24. In this test, filter coefficients are optimized for a given pulse shape (no noise and no time jitter in the training set), with the peak of the signal precisely phased-aligned with the analog to digital converter sampling clock. A train of pulses of constant amplitude (128 on an 8-bit range) with a phase varying in $[-1/2 \text{ BC}, 1/2 \text{ BC}]$ with respect to the sampling clock is generated. Two sets of observations are distinguished: the value of the output for the beam-crossings that correspond to a simulated deposition of energy and the residual error for the beam-crossings where a null response is expected. For a null phase, it can be seen in Figure 24 that the output of the filter corresponds to the expected output for the beam-crossing of interest and is null for adjacent beam-crossings. When the phase is varied, not only a growing error is made on the energy estimated for the correct BC, but also a non-null output for adjacent BC's is observed. The algorithm is somewhat sensitive to sampling clock phase adjustment and signal jitter. A possible improvement would be optimize the filter coefficients with a training set of samples that include time jittered pulses.

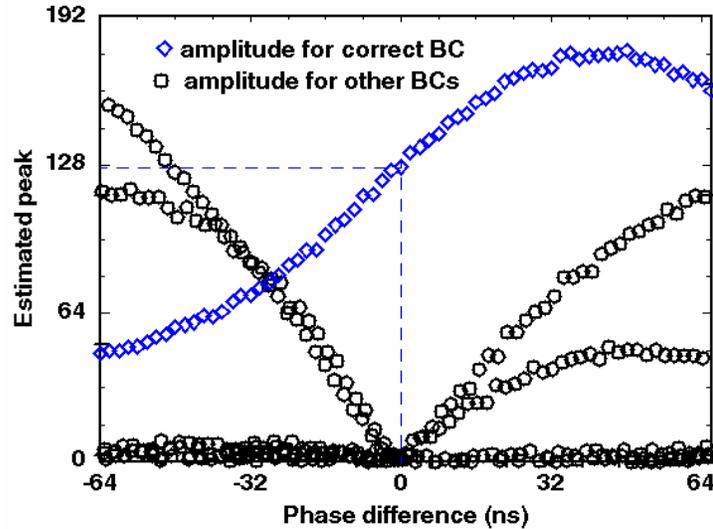


Figure 24. Operation of a deconvolution FIR filter when the phase of pulses is varied. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 12-tap FIR is used; 32-bit floating-point arithmetic is used for coefficients (signed) and computations.

Other difficulties with the deconvolution FIR filter include its sensitivity to limited precision arithmetic and coefficient truncation, degraded performance when signal baseline is shifted and misbehavior when the input is saturated. Implementation is also a potential issue because a filter comprising over 12-tap is needed (assuming input samples are acquired at BC x 2). Although only one convolution product needs to be calculated per BC, a significant amount of resources would be needed to compute the corresponding $7.57 \times 12 = 90$ million multiply-accumulate operations per second per channel. Although an independent study of this algorithm could bring better results, linear deconvolution is not seen as a satisfactory solution.

4.5.5.2 Peak detector + weighed moving average

This non-linear filter comprises two steps: detecting the presence of a peak and calculating its height. The first step is accomplished by comparing the magnitude of ~3-4 successive samples. There is no specific method to pick up the minimum sets of conditions that these amplitudes need to satisfy to characterize the presence of a peak. Let $E(kT)$ be the amplitude of input sample k . A possible set of conditions for peak detection can be expressed as follows:

A peak is present at $t=(k-1)T$ IF

$E(kT) < E[(k-1) T]$ AND

$E[(k-1) T] \geq E[(k-2) T]$ AND

$E[(k-2) T] \geq E[(k-3) T]$

This set of conditions determines the presence of a peak with an irreducible latency of one period T . Empirical studies were performed to determine a viable set of conditions and a satisfactory sampling period T . The conditions mentioned above were retained; sampling at BC x 3 was chosen.

The second part of the algorithm consists in assigning 0 to the output if the peak detector did not detect the presence of a peak, or calculate the weighed average of several samples around the presumed peak. To simplify computations, the sum can be made over 4 samples with an identical weight of $\frac{1}{4}$ for each of them. A common multiplicative scaling factor is then applied.

One of the tests performed with this algorithm is shown in Figure 25. A series of pulses of growing amplitudes is generated. It can be seen that small pulses are not well detected. It should be also observed that, as expected, the output is null between pulses.

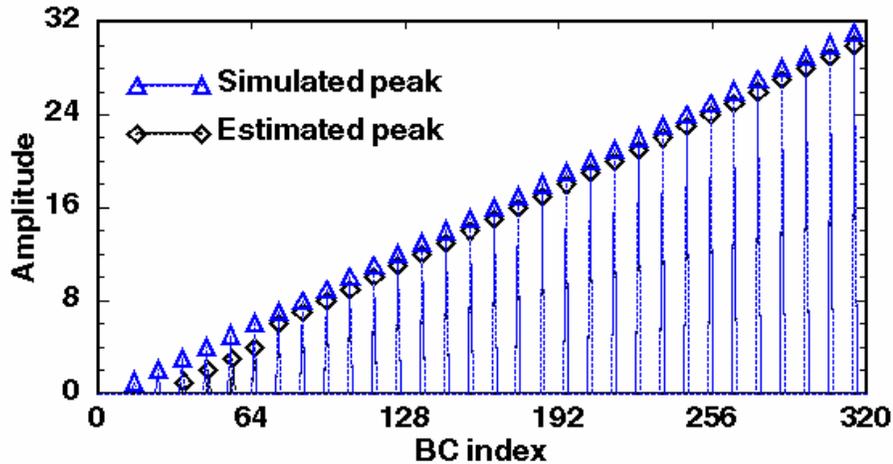


Figure 25. Operation of a peak detector + moving average. Sampling rate is 22.71 MHz (BC x 3), ADC precision is 8 bit; average is made on 4 samples; each weight is $\frac{1}{4}$; a common 8-bit multiplicative factor is applied; fixed-point arithmetic is used.

Other tests show that this algorithm is rather tolerant to signal phase and jitter, does not depend too much on pulse shape (except for double peaked HAD channels), is very simple to implement and has a low latency. Its main limitations are the empirical way for parameter tuning, low performance for small signals, misbehavior in case of pileup, the assignment of energy to the beam-crossing preceding or following the one of interest in some cases, the possibility that a pulse is undetected in some other cases. Although this algorithm is acceptable in some cases, it does not seem sufficiently robust and efficient.

4.5.5.3 Matched filter + peak detector

This algorithm comprises two steps. The matched filter is designed to best optimize the Signal to Noise Ratio when detecting a signal of a known shape degraded by white noise. In this case, it can be shown that the optimal filter for a signal $E(kT)$ is the filter whose impulse response is:

$h(kT) = E(T_0 - kT)$ where T_0 is a multiple of the sampling period T and is selected to cover a sufficient part of the signal. Because T_0 has a direct influence on the irreducible latency of the algorithm, the number of filter taps and the operating frequency of the corresponding hardware, its value should be carefully chosen. The parameters to determine are: the sampling period T , the number of

samples during T_0 , and the phase with respect to the training pulse of the temporal window used to determine the impulse response of the matched filter. It should also be mentioned that the peak produced at the output of a matched filter occurs at $(nT+T_0)$ and that this irreducible latency does not correspond to a fixed delay with respect to the occurrence of the peak in the signal being detected when some of the parameters of the filter are changed. When designing a series of filters running in parallel, care must be taken to ensure that algorithm latency is identical for all channels.

The second step of the algorithm is peak detection. A possible algorithm is the 3-point peak detector described by the following pseudo-code:

Peak present at $t=(k-1)T$ IF $E(kT) < E[(k-1) T]$ AND $E[(k-1) T] > E[(k-2) T]$

This peak-detector adds one period T of irreducible latency. If the conditions that characterize a peak are not satisfied, the output is set to 0, otherwise it is assigned the value of the matched filter.

The various tests that were done to investigate the two algorithms previously described were repeated. An example is given in

Figure 26 where pulses of growing amplitudes (up to $1/8^{\text{th}}$ of the full 8-bit scale) have been generated. It can be seen that the algorithm performs well in that range. All pulses but the smallest ones have been correctly detected and have been assigned to the correct beam crossing. The output is exactly zero for the beam-crossings around those of interest. Intuitively, one can easily understand that the capability to produce minimal width pulses (one sample width) surrounded by strictly null outputs is more easily achieved with a non-linear filter than with a linear algorithm.

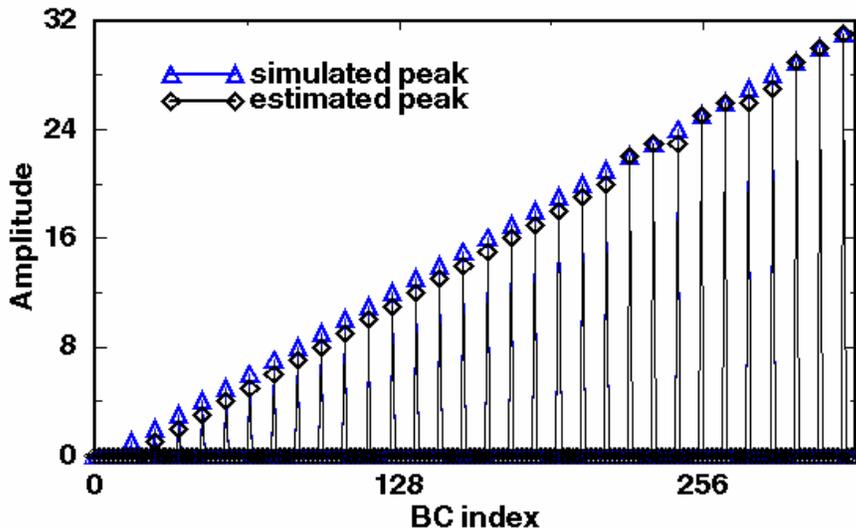


Figure 26. Operation of a matched filter + peak detector. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; 6 6-bit unsigned coefficients are used; fixed-point arithmetic is used.

The sensitivity of the matched filter and peak detector to signal phase shift was studied. Pulses of constant amplitude ($1/2$ full scale) and variable phase

were generated. The relative error on the reconstructed amplitude for the beam crossing of interest is plotted in Figure 27. It can be seen that the relative error is confined within 5% when the phase shift is in the interval [-32 ns, 32 ns]. For larger phase shift values, the pulse is undetected and the output of the filter is null. This is a case of severe failure for this algorithm. For the beam-crossings surrounding that of interest, the output of the filter remains null over the range of phase shifts simulated; no erroneous assignment to the preceding or following beam crossing were observed.

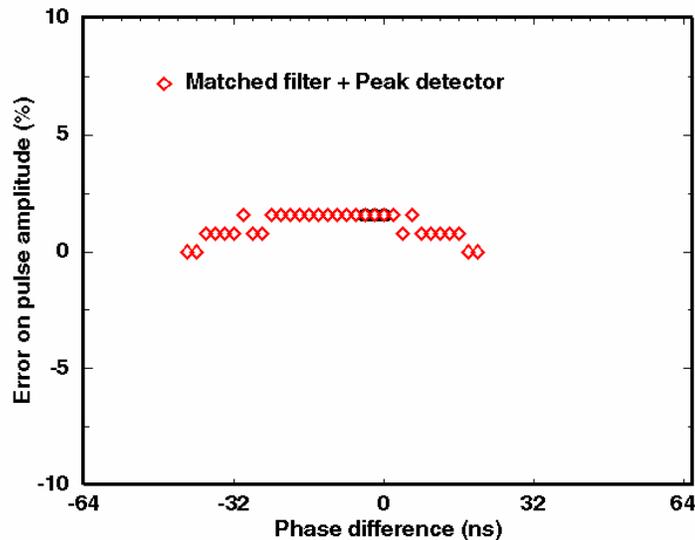


Figure 27. Operation of a matched filter + peak detector when signal phase is varied. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; a 6-tap matched filter with 6-bit unsigned coefficients followed by a 3-point peak detector are used; all computations are done in fixed-point arithmetic.

By comparing these results with that of the FIR deconvolution shown in Figure 24 (where the absolute value of filter output is plotted), it can be concluded that the matched filter algorithm is much more tolerant to signal phase and jitter.

Determining the number of taps for the matched filter requires making a compromise between the quality of the results, the latency of the algorithm and the amount of resources needed for implementation. A test was made to investigate the influence of the number of filter taps. A series of pulses of growing amplitudes (full 8-bit range) were generated. The reconstructed amplitude is shown in Figure 28 for a matched filter with 8-taps and 5-taps respectively. No significant degradation of performance was observed as long as the number of coefficients is greater or equal to 5. The difference in latency between the 8-tap version and the 5-tap version is 1 BC; the amount of computation to perform is increased by 60% when the number of taps is changed from 5 to 8.

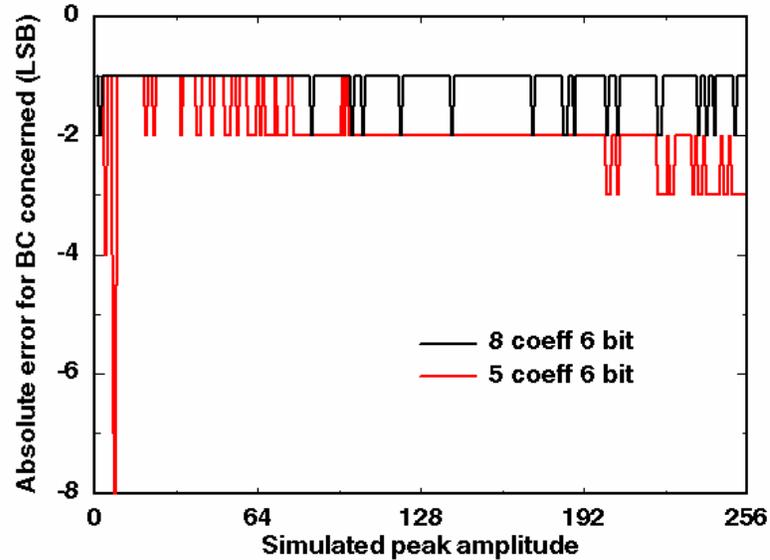


Figure 28. Operation of a matched filter + peak detector with different number of taps. Sampling rate is 15.14 MHz (BC x 2), ADC precision is 8 bit; coefficients are 6-bit unsigned; fixed-point arithmetic is used.

Algorithm behavior in case of saturation is also an important parameter. A series of pulses with amplitude that goes up to twice the range of the ADC (8-bit in this test) was generated. A comparative plot for the 3 algorithms studied is shown in Figure 29. The FIR deconvolution filter has two annoying features: the amplitude estimated for the BC of interest decreases, and the estimation on adjacent BC's grows rapidly as the level of saturation is increased. The peak detector has a satisfactory behavior under moderate saturation, but the peak of energy is assigned to the wrong beam crossing when the saturation level is increased. The matched filter has a smoothly growing output, and still assigns the energy value to the correct beam-crossing under a high level of saturation. Although in real experimental conditions, the combined effects of analog and digital saturation will be much more complex than what was simulated, the matched filter clearly appears to be superior to the two other algorithms.

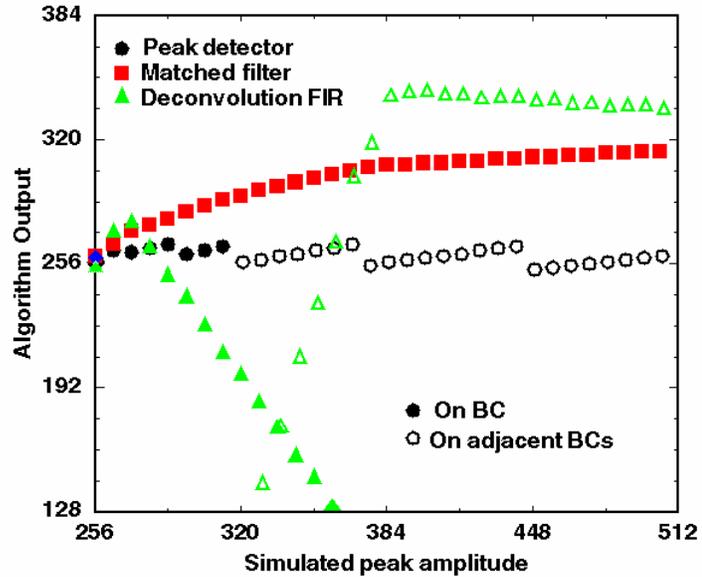


Figure 29. Algorithm behavior under digital input saturation.

Electronic noise reduction and pileup rejection are other properties that need to be considered to select the proper algorithm for digital filtering. At present, almost no studies have been made in these fields but a few simple tests. A series of couple of pulses of constant height (1/2 full range) separated in time by 10, 3, 2 and 1 beam-crossings have been generated. The output for the 3 algorithms studied is shown in Figure 30. As previously mentioned, the deconvolution FIR filter is able to correctly identify pulses that are close in time. On the other hand, both the peak detection scheme and the matched filter algorithm fail to identify the two pulses and their amplitude when pickoff signals overlap. One of the two pulses is systematically dropped and the energy of the remaining pulse is overestimated by a large factor. This configuration corresponds to a case of failure for these two algorithms. Detailed studies are needed to determine what will be the noise and pileup conditions in the real experiment and decide if the level of algorithm failures observed is below an acceptable limit or not. Tests with in the experiment with real signals are also crucial.

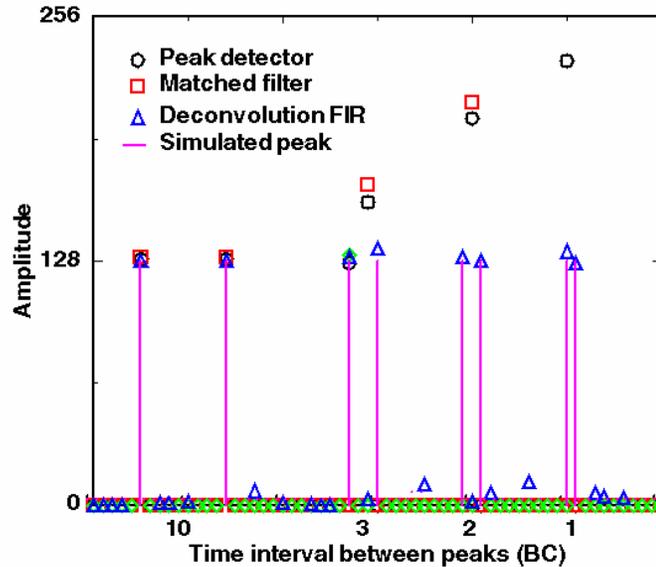


Figure 30. Behavior of the 3 algorithms with pulses close in time.

In order to compare the 3 algorithms studied, 8 criteria of merit were selected and subjective marks between 0 and 5 (0 is worse, 5 is best) were given for each algorithm. The resulting diagram is plotted in Figure 31. While none of the algorithm performs best in all fields, the optimum algorithm is the one whose polygon covers the largest area. Clearly, the matched filter is the algorithm that offers the best trade-off between all criteria. This algorithm is therefore the baseline for the prototype that is being designed and that will be tested in-situ.

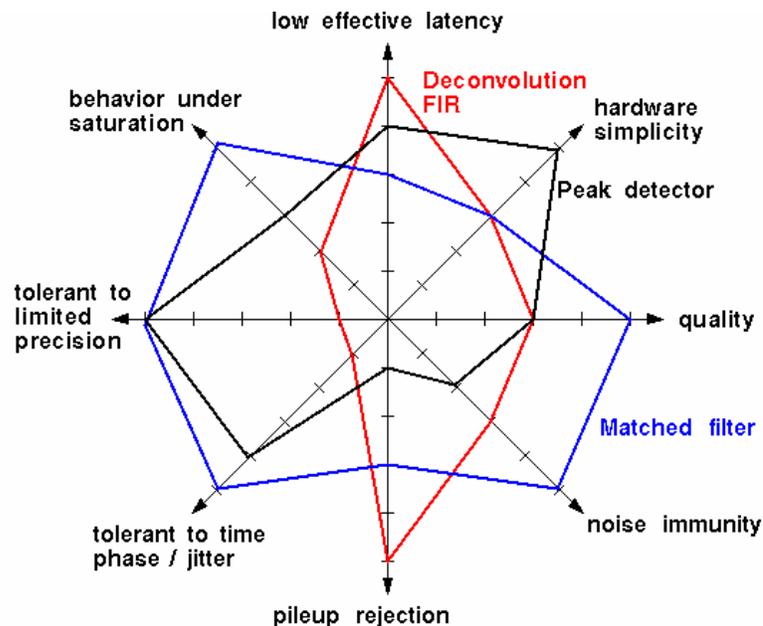


Figure 31. Comparison of the 3 algorithms proposed against 8 criteria of merit.

A number of studies still need to be done to confirm that this algorithm is the most appropriate. These include noise and pileup studies, the possibility to run

computations at the beam-crossing rate (instead of running at BC x 2), the development of a scheme and program for the automatic determination of filter coefficients, etc. Tests on the detector must also be done and analyzed. While simulation offers a very good mean of investigation, real signals shapes, noise, time jitter, pileup conditions and many other effects cannot be taken into account without performing a real test. At present, studies on the digital filter allowed to select a candidate algorithm. In the prototype implementation, some flexibility will be allowed at that level, but algorithm changes will be confined to the capability of the available programmable logic.

4.5.6 Conclusions

Given the relatively slow trigger sum driver pulse shapes observed in Figure 16 and Figure 17, we believe that a digital filter is required to suppress the contributions from signals in nearby bunch crossings to that containing a high pT trigger. Details of the implementation of the digital filter are given in section 4.8.4.

4.6 Clustering algorithm options and simulation results

Algorithms relying on “sliding” trigger towers (TTs) can significantly improve the trigger performances, compared to the actual calorimeter trigger based on single 0.2 x 0.2 TTs, by better identifying the physical objects. Such algorithms have been extensively studied for the Atlas experiment, as described in the Trigger Performance Status Report, CERN/LHCC 98-15.

Various algorithms can be used to cluster the trigger towers and look for “regions of interest” (R), i.e. for regions of fixed size, S , in $\eta \times \phi$ in which the deposited E_T has a local maximum. To find those regions of interest, a window of size S is shifted in both directions by steps of 0.2 in η and ϕ . By convention each window is unambiguously (although arbitrarily in the 2 x 2 case) anchored on one trigger tower T and is labeled $S(T)$. Examples are shown in Figure 32.



Figure 32. Examples of (a) a 3x3 and (b) 2x2 sliding window $S(T)$ associated to a trigger tower T . Each square represents a 0.2 x 0.2 trigger tower. The trigger tower T is shown as the shaded region.

The sliding tower algorithm aims to find the optimum region of the calorimeter for inclusion of energy from jets (or EM objects) by moving a window grid across

the calorimeter η, ϕ space so as to maximize the transverse energy seen within the window. The window of towers so found, together perhaps with a specified set of neighbors, is called the region of interest, R , and is referenced by a specific TT within R as indicated in Figure 32 for a 3x3 or a 2x2 window. The total E_T within R and in the defined neighbor region is termed the trigger E_T relevant to the jet or EM object. A specific example of how the local maximum could be defined is shown in Figure 33. This process, which avoids multiple counting of jet (or EM object) candidates, is often referred to as “declustering”.

>	>	>	>	>
>	>	>	>	>
>	>	S(T)	>	>
>	>	>	>	>
>	>	>	>	>

Figure 33. An illustration of a possible definition of a local E_T maximum for a R candidate. The cluster $S(T)$ is accepted as a R candidate if it is more energetic than the neighboring clusters marked as “>” and at least as energetic as those marked “>”. This method resolves the ambiguities when two equal clusters are seen in the data. In this example, the declustering is said to be performed in a window of size 5x5 in $\eta \times \phi$.

Such “sliding window” algorithms are thus defined by, at least:

- The size in $\eta \times \phi$ of the sliding clusters $S(T)$;
- The size in $\eta \times \phi$ of the window in which the declustering is performed (this defines the minimal distance between two neighboring regions of interest);
- The set of neighbor TTs whose energy is added to that of R , to define the trigger E_T .

Specific parameters which enter in the definition of the electromagnetic or tau triggers only will be detailed in the relevant sections below.

4.6.1 Jet algorithms

4.6.1.1 *Energy resolution and turn-on curves*

The choice of the size of the areas which determine the “trigger jets” has first been studied by looking at the energy resolution achieved, on samples of simulated events, with the following algorithms:

- a) The R size is 0.6×0.6 (Figure 32a) and the trigger E_T is the E_T contained in the RoI.
- b) The R size is 0.4×0.4 (Figure 32b) and the trigger E_T is the E_T contained in the 0.8×0.8 region around the RoI.
- c) The R size is 1.0×1.0 and the trigger E_T is the E_T contained in the RoI.

In each case, the algorithm illustrated in Figure 33 is used to find the local maxima R . For each algorithm, the transverse energy seen by the trigger for 40 GeV jets is shown in Figure 34. This is to be compared with Figure 18, which shows the E_T seen by the current trigger. Clearly, any of the “sliding window” algorithms considerably improves the resolution of the trigger E_T . For the case of the 40 GeV jets studied here, the resolution improves from an rms of about 50% of the mean (for a fixed 0.2×0.2 $\eta \times \phi$ trigger tower) to an rms of 30% of the mean (for a sliding window algorithm), and the average energy measured in the trigger tower increases from $\sim 26\%$ to 56-63% (depending on the specific algorithm).

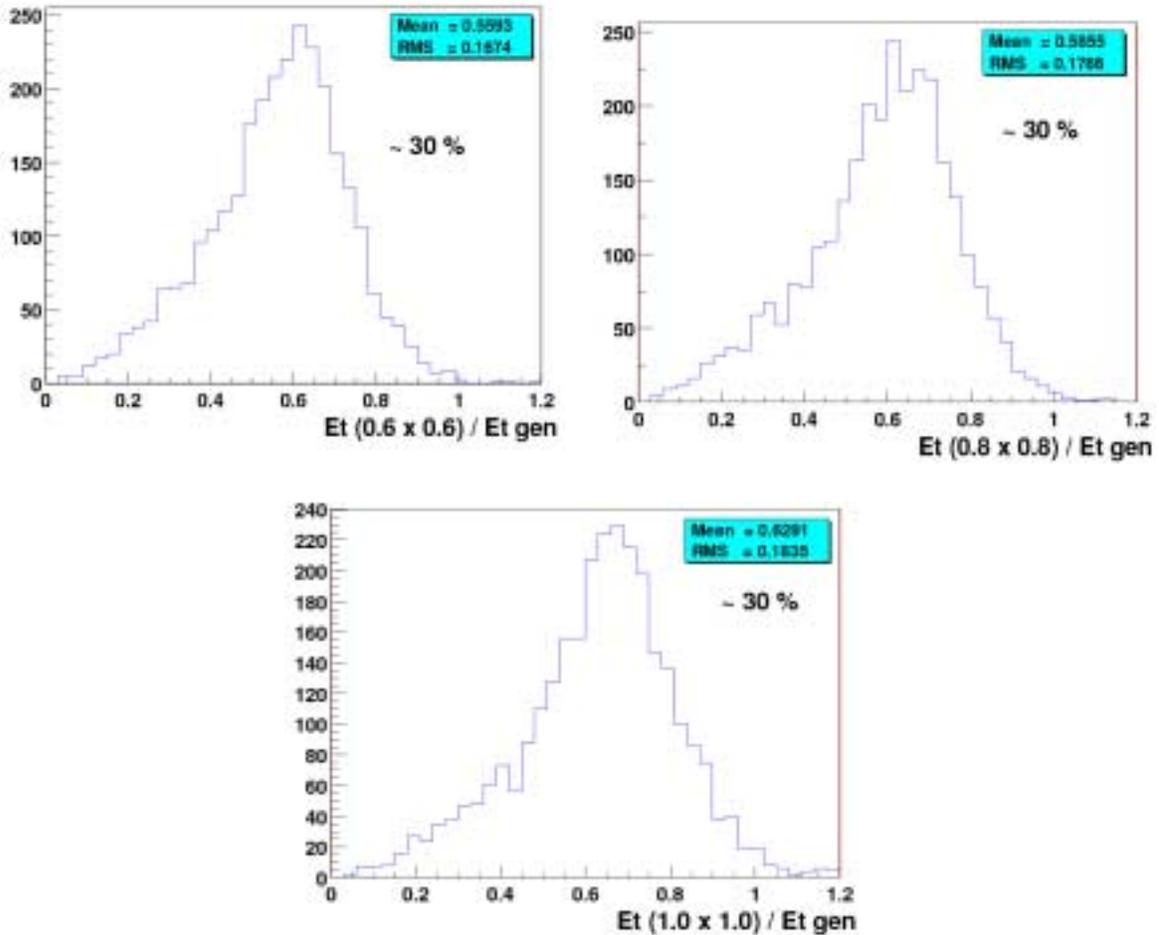


Figure 34. Ratio of the trigger E_T to the transverse energy of the generated jet, using three different algorithms to define the trigger jets. Only jets with $E_T \approx 40$ GeV are used here. The ratio of the rms to the mean of the distribution, the value 30%, is written on each plot.

Since the observed resolution is similar for all three algorithms considered, then the choice of the R definition (*i.e.* of the algorithm) will be driven by other considerations including hardware implementation or additional performance studies.

The simulated trigger efficiency for the (b) algorithm, with a threshold set at 10 GeV, is shown as a function of the generated E_T in Figure 35. The turn-on of the efficiency curve as a function of E_T is significantly faster than that of the current trigger, also shown in Figure 35 for two values of the threshold. With a 10 GeV threshold, an efficiency of 80% is obtained for jets with E_T larger than 25 GeV.

In order to understand which part of these new algorithms are providing the improvement (the sliding window or the increased trigger tower size), we have studied the gain in efficiency which is specifically due to the sliding window procedure by considering an algorithm where the TTs are clustered in fixed 4×4 towers (*i.e.* 0.8×0.8 in $\eta \times \phi$), without any overlap in η or ϕ . The comparison of the “fixed” and “sliding” algorithms is shown in Figure 36. One observes a marked improvement for the “sliding” windows compared to the “fixed” towers, indicating that the added complexity of implementing sliding windows is warranted.

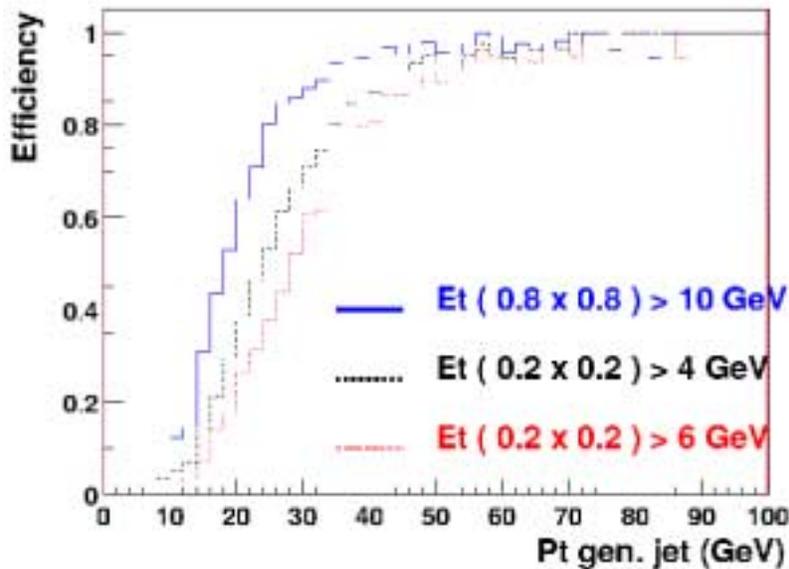


Figure 35. Trigger efficiency as a function of the transverse energy of the generated jet, for the (b) algorithm for $E_T > 10$ GeV (the solid line) and for the current trigger (fixed trigger towers with thresholds of 4 and 6 GeV shown as dashed and dotted lines respectively).

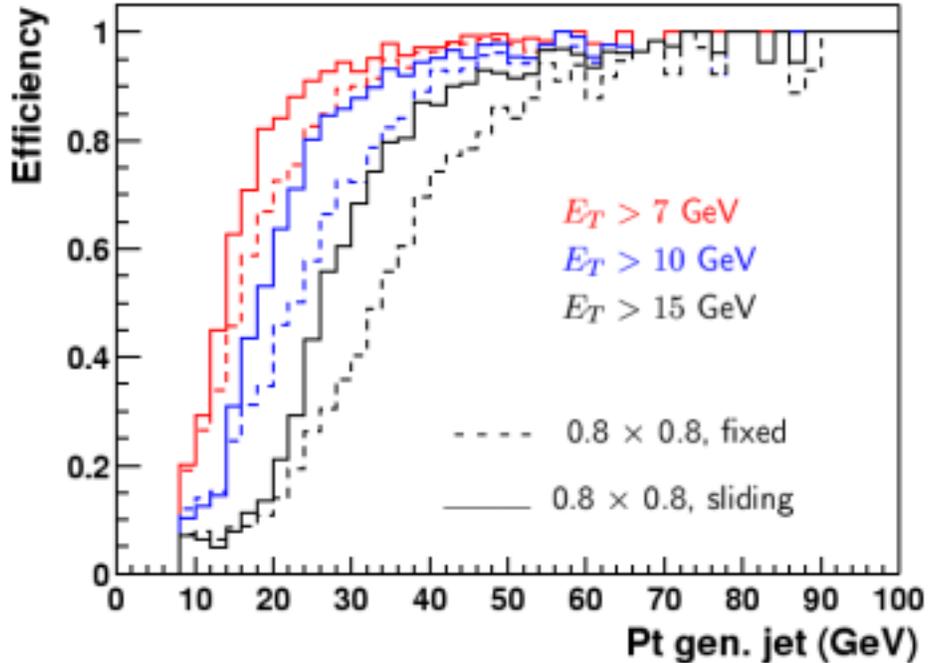


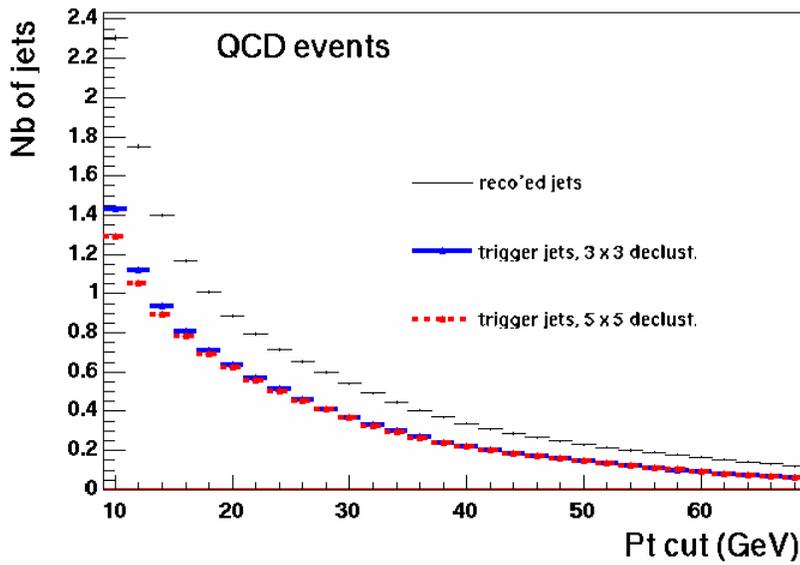
Figure 36. Trigger efficiencies as a function of the generated jet p_T for trigger thresholds $E_T > 7\text{ GeV}$, 10 GeV and 15 GeV (curves from right to left respectively). The solid curves are for the 0.8×0.8 “sliding window” algorithm, and the dashed curves are for a fixed 0.8×0.8 trigger tower in $\eta \times \phi$.

4.6.1.2 Trigger jet multiplicities

The number of jets above a given threshold in E_T will be an important ingredient of any trigger menu. The declustering procedure explained above should not lead to largely overestimated jet multiplicities, while maintaining a reasonable complexity for the TAB implementation. In this section, we compare the jet multiplicities obtained on simulated events using two algorithms. In both cases, the size of the regions of interest is 0.6×0.6 in $\eta \times \phi$ (i.e. 3×3 TTs). The trigger E_T is obtained by adding to the R E_T the energies of the closest neighboring TTs (i.e. trigger jets have a size of 1.0×1.0). In the first (second) algorithm, the declustering is made in a 5×5 (3×3) window: the signals from 7×7 (5×5) TTs are needed to determine whether a R candidate $S(T)$ is a local maximum.

The mean number of jets with E_T above a given threshold is shown in Figure 37, for a sample of simulated QCD events (upper plot), and for pair-produced top quarks which decay fully hadronically (lower plot) leading to high E_T jets. Both trigger algorithms lead to comparable multiplicities, especially when high E_T trigger jets are considered. The multiplicity of jets found by an offline cone algorithm of radius 0.5 is also shown in Figure 37 as the thin line. It is larger than the trigger jet multiplicity, as expected since the trigger jet E_T is not 100% of the reconstructed E_T .

Number of triggered and reco'd jets with $P_t > P_{tcut}$



Number of triggered and reco'd jets with $P_t > P_{tcut}$

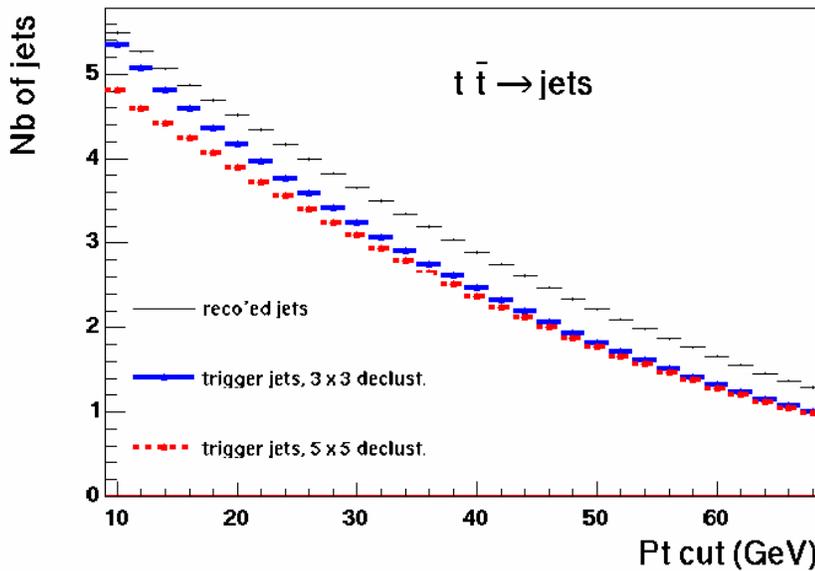


Figure 37: Multiplicities of jets with E_T above a given cut, as found by two trigger algorithms differing in the declustering procedure. The multiplicity of reconstructed jets is also shown.

From this study it seems sufficient to perform the declustering in a rather small (3x3) window.

4.6.1.3 Rates and rejection improvements

In this section, we compare the performance of the sliding window and the existing trigger algorithms. We compare both of these algorithms' trigger efficiencies and the associated rates from QCD jet events as a function of trigger E_T .

a) Rates versus trigger efficiency on hard QCD events

In these studies we require that for the sliding window (b) algorithm there be at least one region of interest with a trigger E_T above threshold which varies from 5 to 40 GeV in steps of 1 GeV. Similarly, for the current trigger algorithm, we require at least one TT above threshold which varies from 2 GeV to 20 GeV in steps of 1 GeV. For both algorithms and for each threshold, we calculate the corresponding inclusive trigger rate and the efficiency to trigger on relatively hard QCD events, *i.e.* with parton $p_T > 20\text{GeV}$ and $p_T > 40\text{GeV}$ respectively. To simulate high luminosity running, we overlay additional minimum bias events (a mean of 2.5 or 5 additional minimum bias events) in the Monte Carlo sample used to calculate the rates and efficiencies. While the absolute rates may not be completely reliable given the approximate nature of the simulation, we believe that the relative rates are reliable estimators of the performance of the trigger algorithms. Focusing on the region of moderate rates and reasonable efficiencies, the results are plotted in Figure 38 where lower curves (open squares) in the plots are for the current trigger algorithm and the upper curve (solid circles) corresponds to the sliding window (b) algorithm. It is apparent from Figure 38 the sliding window algorithm can reduce the inclusive rate by a factor of 2 to 4 for any given efficiency. It is even more effective at higher luminosities (*i.e.* for the plots with 5 overlaid minimum bias events).

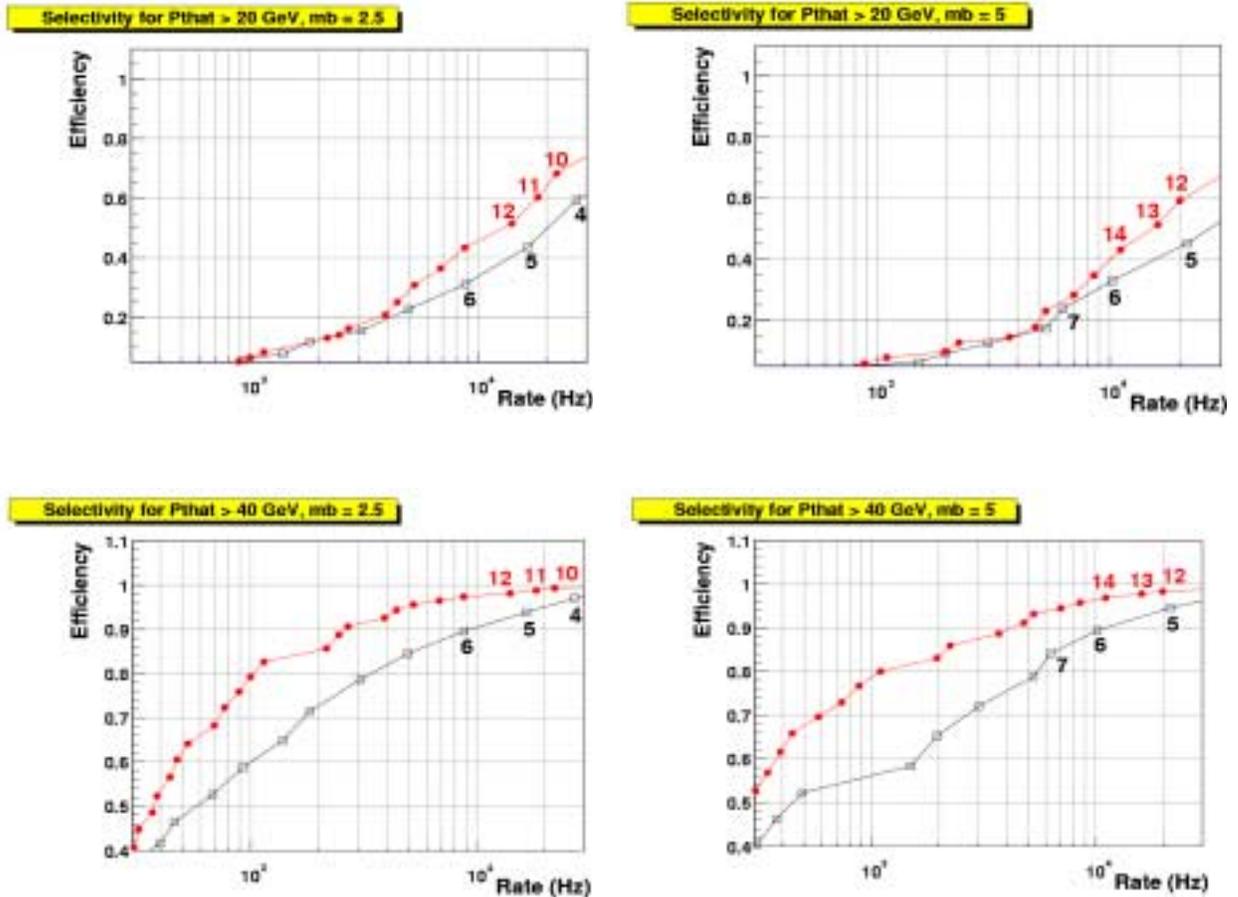


Figure 38. Trigger efficiency for events with parton $p_T > 20$ GeV (upper plots) and parton $p_T > 40$ GeV (lower plots) as a function of the inclusive trigger rate, for the (b) algorithm (solid circles) and the current algorithm (open squares). Each dot (solid circle or open square) on the curves corresponds to a different trigger threshold; the first few are labeled in GeV, and they continue in 1 GeV steps. The luminosity is $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ and the number of overlaid minimum bias (mb) events follows a Poisson distribution of mean equal to 2.5 (left hand plots) or to 5 (right hand plots).

b) Rates versus trigger efficiency on events with a large hadronic activity

In this section we study the performances of sliding algorithms on events which have a large number of jets in the final state. As an example we consider the case of pair produced top quarks which both decay fully hadronically. Other topologies with large jet multiplicities could arise from the production of squarks and/or gluinos.

Three sliding algorithms have been considered here:

- i. The size of the regions of interest is 0.6×0.6 (i.e. 3×3 TTs); the trigger E_T is that of R ; the declustering is performed in a 5×5 window. This algorithm is labeled 3_0_0.
- ii. As (i) but the trigger E_T is obtained by summing the E_T of R and the E_T of the closest neighboring TTs. This algorithm is labeled 3_0_1.
- iii. As (ii) but the declustering is performed in a 3×3 window. This algorithm is labeled 3_m1_1.

In each case, the trigger condition requires that there be at least three trigger jets with E_T above a varying threshold. In addition, the E_T of the highest E_T jet should be above 40 GeV. A similar trigger condition has also been applied using the 0.2×0.2 TTs instead of the trigger jets; in this latter case the highest E_T TT should have $E_T > 15$ GeV. The inclusive QCD rate has been obtained as before, using QCD Monte Carlo events where a mean number of 7.5 minimum bias events has been overlaid. Figure 39 shows the resulting efficiencies and rates. Inclusive rates are shown here for a luminosity of $5 \times 10^{32} \text{ cm}^2 \text{ s}^{-1}$.

It can be seen that the three sliding algorithms considered lead to very similar performances. In particular, no noticeable difference is seen between algorithms 3_0_1 and 3_m1_1 (which differ by the declustering procedure only), as was seen in section 4.6.1.2. The figure also shows that the performances of sliding algorithms are better than those of the current trigger system, also for events with many jets in the final state.

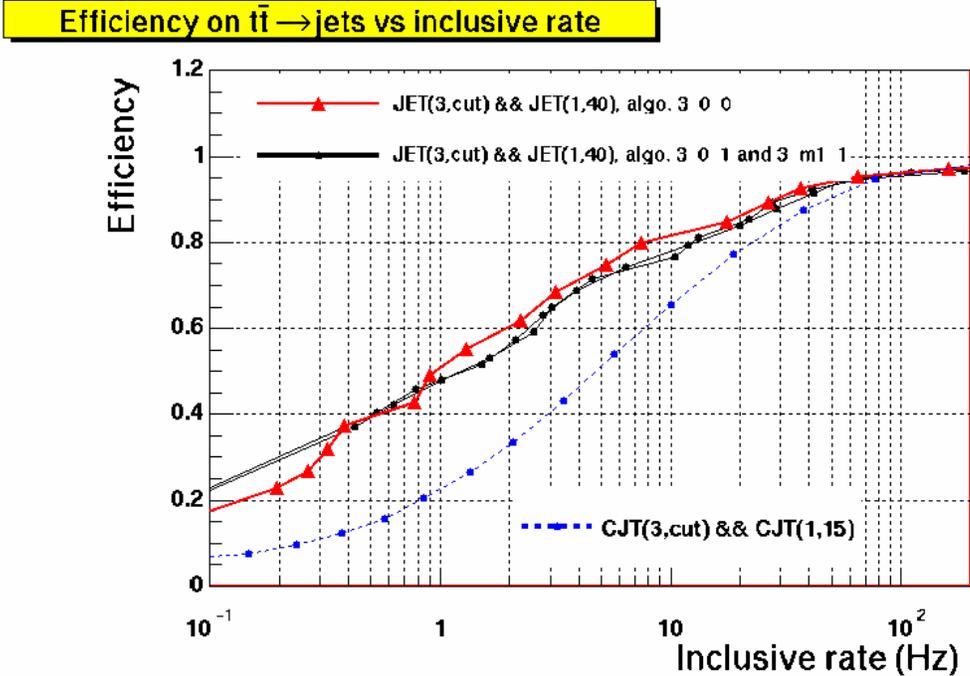


Figure 39. Trigger efficiency for simulated pair produced top quarks which both decay hadronically, as a function of the inclusive trigger rate, for various sliding window algorithms (full curves, solid circles and triangles), and using the current trigger towers (dashed curve, solid circles). The trigger condition for the sliding (current) algorithms requires at least three jets (TTs) with E_T above a varying threshold; the highest E_T jet (TT) must moreover satisfy $E_T > 40$ GeV ($E_T > 15$ GeV).

c) Rates versus trigger efficiency on “difficult” topologies

The improvement in jet triggering provided by the proposed algorithm is important for those physics processes that do not contain a high p_T lepton which in and of itself offers considerable rejection. Since the sliding window algorithm would be implemented in FPGA-type logic devices, it opens up the possibility of including further refinements in the level of trigger sophistication, well beyond simple counting of the number of towers above threshold. We have studied the trigger for two processes which demonstrate the gains to be expected from a sliding window trigger over the current trigger:

- The production of a Higgs boson in association with a $b\bar{b}$ pair. This process can have a significant cross-section in supersymmetric models with large $\tan\beta$, where the Yukawa coupling of the b quark is enhanced. Thus when the Higgs decays into two b quarks this leads to a $4b$ signature. The final state contains two hard jets (from the Higgs decay) accompanied by two much softer jets. Such events could easily be separated from the QCD background in off-line analyses using b -tagging. But it will be challenging to efficiently trigger on these events while retaining low inclusive trigger rates.

- The associated production of a Higgs with a Z boson, followed by $H \rightarrow b\bar{b}$ and $Z \rightarrow \nu\bar{\nu}$. With the current algorithm, these events could be triggered on using a di-jet + missing energy requirement. The threshold on the missing energy could be lowered if a more selective jet trigger were available.

Figure 40 shows the efficiency versus inclusive rate for these two processes, where three different trigger conditions are used:

1. At least two fixed trigger towers of 0.2×0.2 above a given threshold (dotted curves, open squares).
2. At least one TT above 10 GeV and two TT above a given threshold (dot-dash curve, solid stars).
3. At least two “trigger jets” whose summed trigger E_T 's are above a given threshold (solid curve, solid circles).

The algorithm b) has been used here. It can be seen that the third condition is the most efficient for selecting signal with high efficiency but low rates from QCD jet processes.

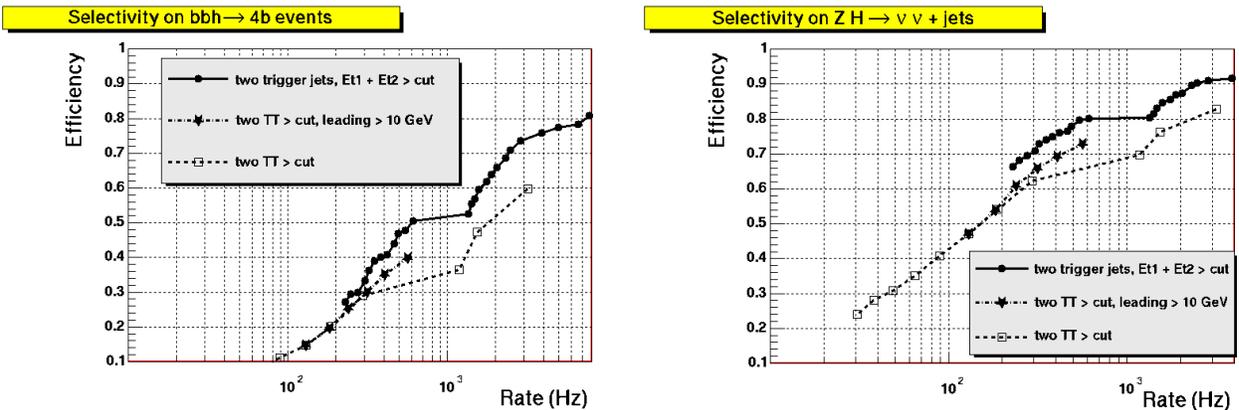


Figure 40. Efficiency to trigger on $bb\bar{h}$ (left) and ZH (right) events as a function of the inclusive rate. The three conditions shown require: at least two TT above a threshold (dotted, open squares), at least one TT above 10 GeV and two TT above a threshold (dot-dash, solid stars), at least two trigger jets such that the sum of their trigger E_T 's is above a given threshold (solid circles).

4.6.1.4 Including ICR in trigger jets algorithms

- E_T (trigger jet) / E_T (reco jet) versus η_{jet} , with and without ICR : this should show that ICR brings some improvement in the relevant η range.
- Efficiency (eg on hard QCD events) versus inclusive rate, when the trigger requires at least one jet with $E_T > cut$, with and without ICR.

This should show what the real physics effect will be, when integrated over a reasonable eta range.

Plots should be available during next week.

4.6.2 Electron algorithms

4.6.2.1 Transverse isolation and electromagnetic fraction

4.6.2.2 Energy resolution and turn-on curves

4.6.2.3 Rates and rejection improvement

4.6.3 Tau algorithms

With some refinements, the sliding window algorithms presented in section 4.6.1 could lead to some sensitivity to the process $gg \rightarrow H \rightarrow \tau^+\tau^-$ at the trigger level. This could be achieved by exploiting the fact that τ jets are more narrow than “standard” jets.

4.6.3.1 Transverse isolation of τ jets

We consider here the sliding algorithm (b) described in section 4.6.1.1, where the size of regions of interest is 0.4×0.4 (2x2 TTs), while the size of trigger jets is 0.8×0.8 (4x4 TTs). We compare the ratio of the $R E_T$ to the trigger E_T , for τ jets coming from $gg \rightarrow H \rightarrow \tau^+\tau^-$ and for jets coming from QCD processes. As shown in Figure 41, QCD jets become more and more collimated as their E_T increases, but the ratio of the “core E_T ” to the trigger jet E_T (called the “core fraction”) remains a powerful variable to discriminate between τ jets and QCD jets.

Core of τ and QCD jets

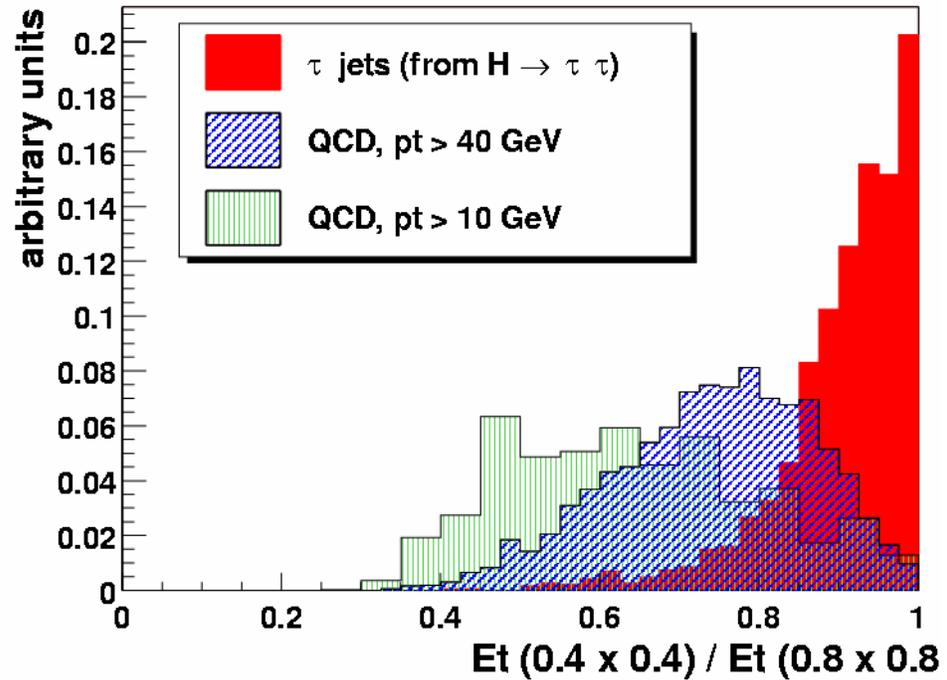


Figure 41: Ratio of the $R E_T$ to the trigger E_T , for the sliding window algorithm (b). The ratio is shown for τ jets coming from a Higgs decay (full histogram), and for jets coming from QCD processes (hashed histograms).

4.6.3.2 Rates and rejection improvement

This can be exploited by defining a specific trigger condition, which requires at least two jets whose summed trigger E_T 's is above a threshold, and for which the core fraction is above 85%. As can be seen in Figure 42, it seems possible to have a reasonable efficiency on the signal (70 %) while maintaining the inclusive rate below 300 Hz. The figure also shows that such an algorithm reduces the inclusive rate by a factor of about 3, compared to the current trigger system.

Selectivity on $H \rightarrow \tau\tau \rightarrow$ dijets (mb=7.5)

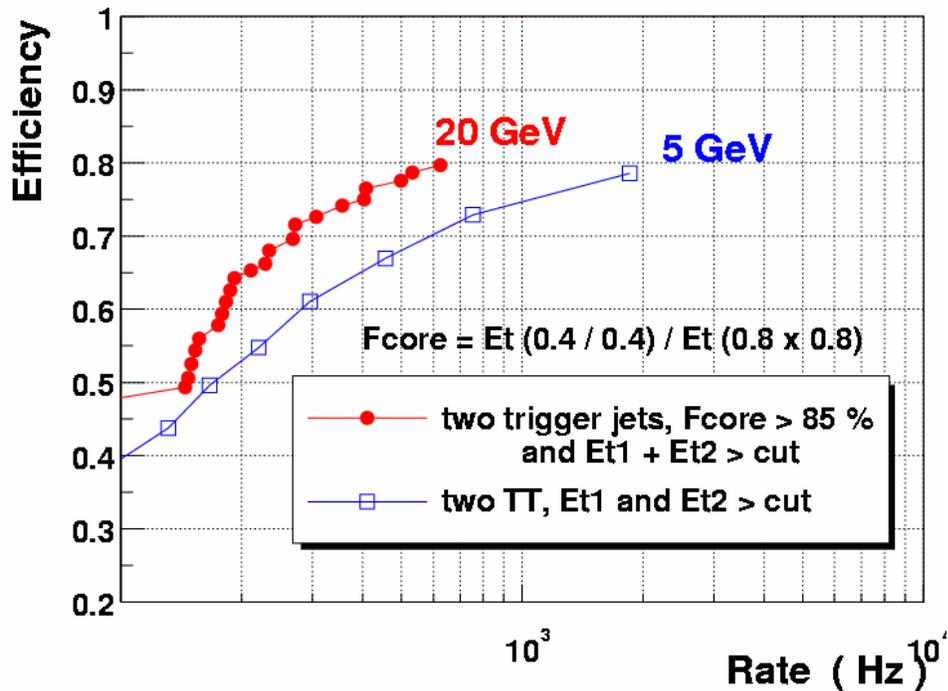


Figure 42: Efficiency to trigger on $gg \rightarrow H \rightarrow \tau\tau$ events as a function of the inclusive QCD rate, for: (closed circles) the sliding window algorithm (b), when requiring at least two jets whose summed E_T is above a varying threshold, and whose core fraction is above 85%; (open squares) the current trigger system, requiring two TTs whose summed E_T is above a varying threshold. The inclusive rates shown here correspond to a luminosity of $2 \times 10^{32} \text{ cm}^2 \text{ s}^{-1}$.

4.6.4 Global sums

The region around $0.8 < |\eta| < 1.5$, known as the inter cryostat region (ICR) encompasses the transition from showers contained within the CC and showers contained within the EC. There is a gap in EM coverage and a major thinning of FH coverage in this area. Since these are the layers which comprise standard trigger towers, there is a major degradation in Level 1 calorimeter response and resolution in this region. This is exacerbated by the presence of significant dead material in the solenoid in Run2. To aid in recovering the energy losses in the ICR region, we use the intercryostat detectors (ICD), which consists of scintillators located in the gap between the calorimeter cryostats, and the “massless gaps” (MG) which consist of the front sections of the endcap calorimeter that have no absorber in front. In this section we study ways to improve the energy measurement at the trigger level.

4.6.4.1 Concept & performance

Global tower E_T sums such as missing E_T or scalar E_T , while very useful, suffer from several significant problems at the L1 trigger. There are two significant issues: first is that the ICR sampling layers are not available in the

calculation at Level 1; second is that the imprecision of the tower E_T 's gets compounded for global sums, resulting in significantly degraded effectiveness. This is particularly true in a multiple interaction environment. There are two possible solutions to these problems. First we can take advantage of work done for Run2a to make the ICR layers available at Level 2 and add these towers back into the global sums at Level 1 in Run2b. Second, we can develop a scheme which discriminates towers which are from multiple interactions and avoids adding them into the sum.

Simulations of single pions and jets in this region indicate that the energy scale in this region goes as low as 40% of the CC/EC scale (as shown in Figure 43), and the resolution is as bad as 6 times worse than in the CC or EC (as shown in Figure 44). These results are very consistent with findings from Run1 Level 1 missing E_T analyses (see Figure 45). One of the major results of this deficiency is that the efficiency and rejection of a Level 1 missing E_T selection are noticeably degraded. These simulations also indicate that adding ICD and MG cells into trigger towers can improve the scale by a factor of 2, while optimizing the resolution by a factor of 3.

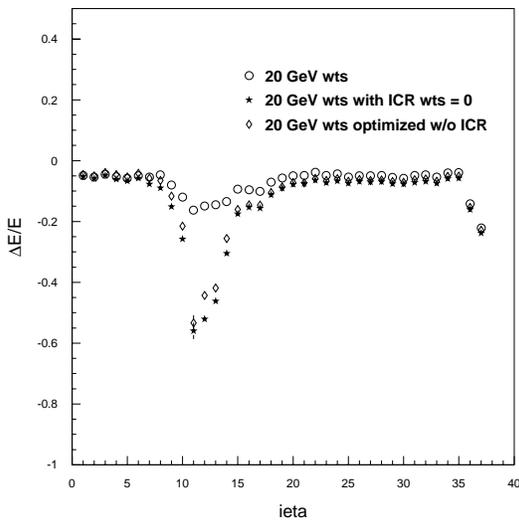


Figure 43. The relative calorimeter energy response in the ICR region for incident 20 GeV pions as a function of $\eta \times 10$. The stars are the response if the ICR weights are set to zero, the open diamonds are the response if the ICR energies are ignored and the remaining calorimeter weights are re-optimized, and the open circles are the response when the ICR region is included and the weights are optimized.

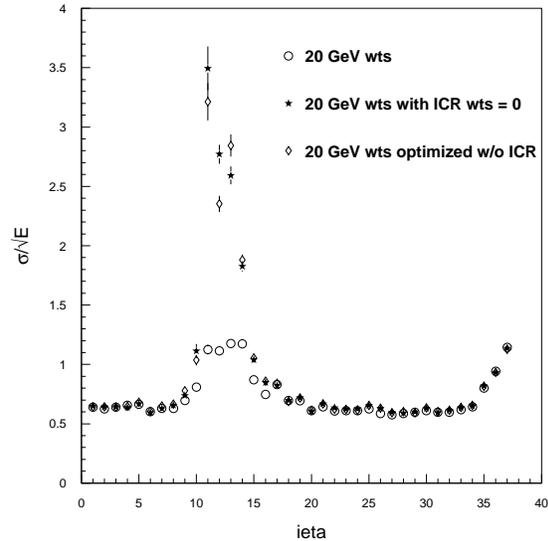


Figure 44. The calorimeter energy resolution in the ICR region for incident 20 GeV pions as a function of $\eta \times 10$. The stars are the response if the ICR weights are set to zero, the open diamonds are the response if the ICR energies are ignored and the remaining calorimeter weights are re-optimized, and the open circles are the response when the ICR region is included and the weights are optimized.

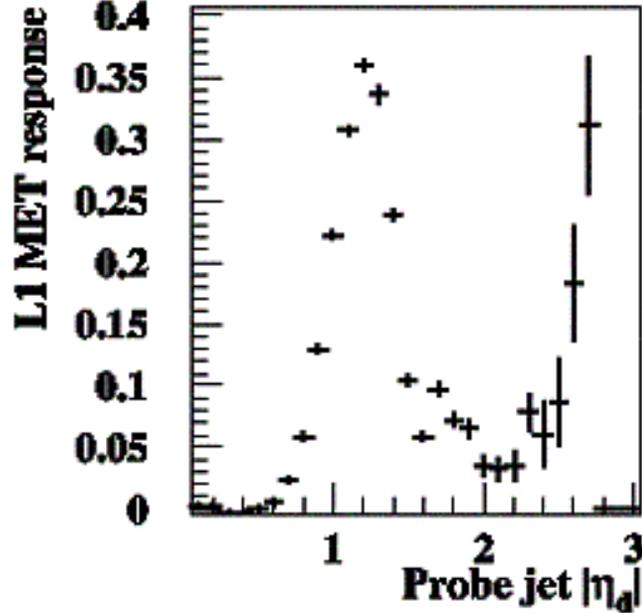


Figure 45. The L1 missing E_T response as a function of $|\eta_d|$ for 85 GeV jets using the Run 1 DØ detector simulation.

4.6.4.2 Simulation results

In principle, it is straightforward to estimate the effect of the ICD and MG to the missing E_T calculation. However our present simulations do not yet fully address a number of issues (including a proper treatment of trigger tower sampling weights, the verification of the modeling and calibration for the ICR, and the proper mapping of calorimeter cells in the MC). The last of these problems is easily solved, but the first two present a larger problem, and so until such time as we have resolved these problems, we will estimate the expected improvement based on other studies.

To estimate the effect of adding the ICR detectors into the missing E_T , we consider the fact that in the region of $1.0 < |\eta| < 1.4$, the sampling weight simulations indicate approximately half of the energy will be deposited in the EM+FH, and the other half in the ICD+MG. As a crude estimate of the magnitude of the effect of adding the ICR layers, we will merely consider the missing E_T measurement with and without the EM+FH layers in this region and assume the ICR improvement will be similar. Although the sample used for this calculation is a QCD sample with jet $p_T > 20$ GeV and 0 minimum bias events overlaid, for historical reasons it is a different sample than that mentioned in the rest of this document with the same specifications. The missing E_T mean and rms in this sample behave as follows:

if remove all ICR TTs: $\mu/\text{rms} = 6.7 \text{ GeV} / 4.8 \text{ GeV}$

if only use EM+FH TTs: $\mu/\text{rms} = 5.5 \text{ GeV} / 3.9 \text{ GeV}$

The number of events passing various Level 1 missing E_T cuts in this sample are shown in Table 8.

Table 8. Events passing L1 missing E_T cuts when the ICR energy is included and when it is removed from the trigger towers.

L1 ME_T	Without ICR	With ICR
> 5GeV	948	766
> 10 GeV	337	185
>15 GeV	95	40
> 20 GeV	37	11
> 25 GeV	9	4

Thus, the region is important to the missing E_T calculation and the rates of passing 15 or 20 GeV selection can change by factors of around 2.5-3. A proper treatment of the gains from adding in the ICD and MG, however, will have to await a satisfactory treatment of the relative weights of various layers.

4.6.4.3 Improving Missing E_T for Multiple interaction Events

Our experience in Run1 indicated the Level 1 missing E_T to be very sensitive to the number of multiple interactions. This results from several factors, including the fact that the fundamental trigger tower fractional energy resolution is poor, especially for very low E_T towers, and the numbers of these towers increases substantially with the number of multiple interactions. As a result, we have explored three ways in which we might improve the missing E_T resolution to reduce this problem in Run2b.

First, we varied the low threshold on the E_T of towers going into the global sum. In Run1, this threshold was 0.5 GeV and was not studied in detail in the light of multiple interactions. Again, we have used the QCD $p_T > 2$ GeV and $p_T > 20$ GeV samples with 0 minimum bias (mb) events overlaid, a 5mb overlay, and a 10mb overlay. We have used the $t\bar{t}$ sample with 2.5 mb overlays for the signal. If we calculate the missing E_T mean and the rms in these samples for various E_T thresholds, we find the results shown in Table 9.

Table 9. Change in the means and rms for the missing E_T for background (QCD) and signal ($t\bar{t}$) samples as a function of the trigger tower (TT) threshold. A selection of 1.5 GeV on trigger towers removes most of the multiple interaction variation for the QCD samples, while having little effect on the signal top sample.

ME _T calc	2 GeV QCD (μ /rms) in GeV	20 GeV QCD 0mb (μ /rms) in GeV	2 GeV QCD 10 mb (μ /rms) in GeV	20 GeV QCD 5mb (μ /rms) in GeV	ttbar (μ /rms) in GeV
TT>0.5GeV	1.0/1.0	5.1/3.8	3.1/2.2	6.5/4.2	35.9/25.4
TT>1GeV	0.6/0.9	5.2/3.9	2.3/1.9	5.8/4.0	35.4/24.7
TT>1.5GeV	0.3/0.7	5.3/4.1	1.6/1.9	5.6/4.0	35.0/24.1
TT>2GeV	0.1/0.6	5.2/4.2	1.0/1.7	5.4/4.2	34.6/23.6

The error on the mean and RMS for the QCD samples is approximately 0.1 GeV. The cut of 2GeV reduces the mean of the QCD sample noticeably. If we consider the 20 GeV sample, the trigger tower cut of 1.5 GeV provides a 20% to 30% lower pass rate for moderate missing E_T selections. Although scalar E_T is generally considered a poor variable at Level 1 because of its sensitivity to multiple interactions, we have studied its mean and rms (see Table 10) for the same thresholds to see what is happening:

Table 10. Change in the means and rms for the E_T scalar sum for background (QCD) and signal (ttbar) samples as a function of the trigger tower (TT) threshold.

Sum E _T calc	2 GeV QCD 0.7 mb(μ /rms) in GeV	QCD 0 mb (μ /rms) in GeV	2GeV QCD 0.7 mb(μ /rms) in GeV	QCD 5mb (μ /rms) in GeV	ttbar (μ /rms) in GeV
TT>0.5GeV	2.9/3.3	23.5/13.0	21.2/18.1	57.7/39.3	179.7/68.8
TT>1GeV	0.8/1.5	17.9/11.9	6.5/7.1	26.6/15.8	161.1/66.4
TT>1.5GeV	0.3/1.1	14.7/11.4	2.8/4.2	18.0/12.5	151/64.9
TT>2GeV	0.2/0.8	12.5/11.1	1.5/3.1	14.2/11.6	143.6/63.8

Comparison of the two QCD samples indicates that low thresholds let in an enormous amount of energy which has nothing to do with the hard scatter interaction.

Because the typical low p_T QCD event E_T is distributed flat in η , we might not expect a degradation in the global sum behavior from including forward trigger towers in the calculation of these quantities. In fact, when looking in simulated events even with large numbers of multiple interactions, one finds very little transverse energy in this region. However, our experience in Run1 indicated strongly that use of forward towers (i.e. those around $|\eta| \sim 3$ or more) substantially degraded the missing E_T behavior. This was especially true in a multiple interaction environment. As a result, we suspect strongly that there is a benefit from being able to easily select what the range is for the calculation, or

perhaps include the η parameter into a weighting scheme with the trigger tower E_T . This requires further study only possible once data is available.

Another concern for the missing E_T measurement involves the truncation of trigger tower E_T 's into 0.5 GeV bins. Since one to two hundred towers are typically added into the Missing E_T , this resolution loss can start to be noticeable. Taking the QCD $p_T > 20$ GeV sample with minimum bias overlay of 0 and 1648 events, we can use the simulator described above in the ICR discussion and toggle truncation on and off. The results are shown in Table 11.

Table 11. Comparison of the effect of TT truncation on the missing E_T . The table lists the number of events (out of a sample of 1648, QCD with $p_T > 20$ GeV and no minimum bias overlaid events) that pass the listed missing E_T thresholds.

Missing E_T	no truncation	no truncation, TT>0.5GeV	with truncation
>5 GeV	947	868	766
>10 GeV	309	261	185
>15 GeV	76	51	40
>20 GeV	22	17	11
>25 GeV	7	5	4

The first column indicates truncation turned off and no threshold applied to trigger towers. The second column also has no truncation and zeros out all towers with $E_T < 0.5$. The third column employs the normal 0.5 GeV truncation. Since truncation lowers tower E_T 's only to the next lowest 0.5 GeV increment, it effectively deweights all of the poorly measured E_T in low E_T towers. In fact, if we consider the QCD $p_T > 20$ GeV sample with 5mb already discussed, the missing E_T mean and rms are mildly improved over the straight 1.5 GeV threshold by a simple weighting scheme. If we choose weights of 5%, 25%, and 75% for $E_T = 0.5, 1.0,$ and 1.5 GeV, respectively, we find the results shown in Table 12.

Table 12. Comparison of simple TT threshold vs. weighting scheme for 20GeV QCD jet sample.

Scheme	μ (GeV)	rms
TT $E_T > 1.5$ GeV:	5.41	4.20
Weighted TT:	5.41	3.96

If the capability exists in an FPGA to enforce a weighting scheme, then one might devise a scheme which does better than this.

Because the trigger tower threshold seems to be the simplest solution that shows progress, and the weighting also seems to help, one might ask whether rejecting low E_T towers unless they are near significant neighbors might help.

Looking again in the 5mb QCD sample at missing E_T means and sigmas, we find the results shown in Table 13. These results point to a significant degradation in missing the E_T mean and resolution.

Table 13. Comparison of effect of rejection low ET towers unless they are near trigger towers (NN) with significant energy deposits.

Cut	μ (GeV)	rms (GeV)
None :	6.45	4.17
NN $E_T > 0.5$ GeV:	6.45	4.37
NN $E_T > 1.5$ GeV:	6.56	4.37
NN $E_T > 3.0$ GeV:	6.72	4.86
NN $E_T > 10$ GeV:	5.62	4.57
NN $E_T > 1k$ GeV:	5.41	4.20

4.6.4.4 Conclusions

In this section, we have explored several different ways to improve the calorimeter missing E_T measurement at Level 1. Studies leading to the optimization of the Run2a trigger have indicated a large improvement in the scale and resolution of jets in this region if the ICD and MG are used. Although our current simulation samples do not have a proper treatment of this region, a crude estimate indicates that this amount of energy should have a noticeable improvement on the missing E_T resolution.

Several attempts were also made to improve the behavior of missing E_T in a multiple interaction environment. The most promising appears to be a simple tightening of the E_T threshold on a trigger tower to around 1.5 GeV which would reduce the background by around 20% in our QCD sample. The actual degradation in the real data may be larger than we see here, however, and the corresponding gain may also increase. We will be in a better position to evaluate this when we have reliable data at various luminosities. There is some evidence that a weighting scheme would provide further benefits.

4.7 Topological Considerations

4.7.1 Concept & physics implications (acoplanar jets)

The search for Higgs boson is the central element of the Run 2b physics program. The Run II Higgs workshop report⁷ concluded that the channel $p\bar{p} \rightarrow HZ \rightarrow b\bar{b} \nu\bar{\nu}$ was critical. This final state poses a difficult topology, two relatively soft jets ($p_T < M_{H/2}$) with modest missing E_T . For a Run 1 style di-jet plus missing E_T trigger, the nominal calorimeter trigger tower and missing E_T thresholds are such that the efficiency for the $b\bar{b} \nu\bar{\nu}$ channel is compromised.

⁷ M. Carena, J.S. Conway, H.E. Haber and J.D. Hobbs, hep-ph/0010338.

The trigger efficiency is driven by the allowable Level 1 rate. While b-tagging can be used at Level2 to control the rate, it is important to note that b tagging will not be possible at Level 1. Thus, it is clear that this channel relies on alternative triggering techniques at Level 1.

4.7.2 Efficiency

To efficiently trigger on the HZ channel one can exploit the unique topology: the Higgs boson recoils from a massive particle decaying invisibly, thereby leading to an acoplanar jet topology. From Monte Carlo based studies, it has been demonstrated that the L1 central track trigger (L1CTT) can be used to identify acoplanar topologies using the fiber tracker. The algorithm is based on identifying the 4.5° wide sector having the highest track p_T sum within the two 45° wide octants having the highest track p_T sum. In Figure 46, the opening angle between the leading partons is shown; the binning reflects the central fiber tracker (CFT) azimuthal segmentation. The red histogram represents the true opening angle and the blue is the corresponding angle reconstructed from charged tracks in the CFT using the above algorithm. Note the QCD background is predominately back-to-back (i.e. the most probable opening angle is 40 sectors or 180°) whereas the Higgs boson signal has a substantial acoplanar component. Figure 47 shows the correlation between $\Delta\phi$ and the missing E_T of the event for signal and representative QCD backgrounds. The figures demonstrate that combining an acoplanar topology cut ($N_{\text{sector}} < 35$) with a looser missing E_T requirement can maintain good signal efficiency while still suppressing most of the QCD background.

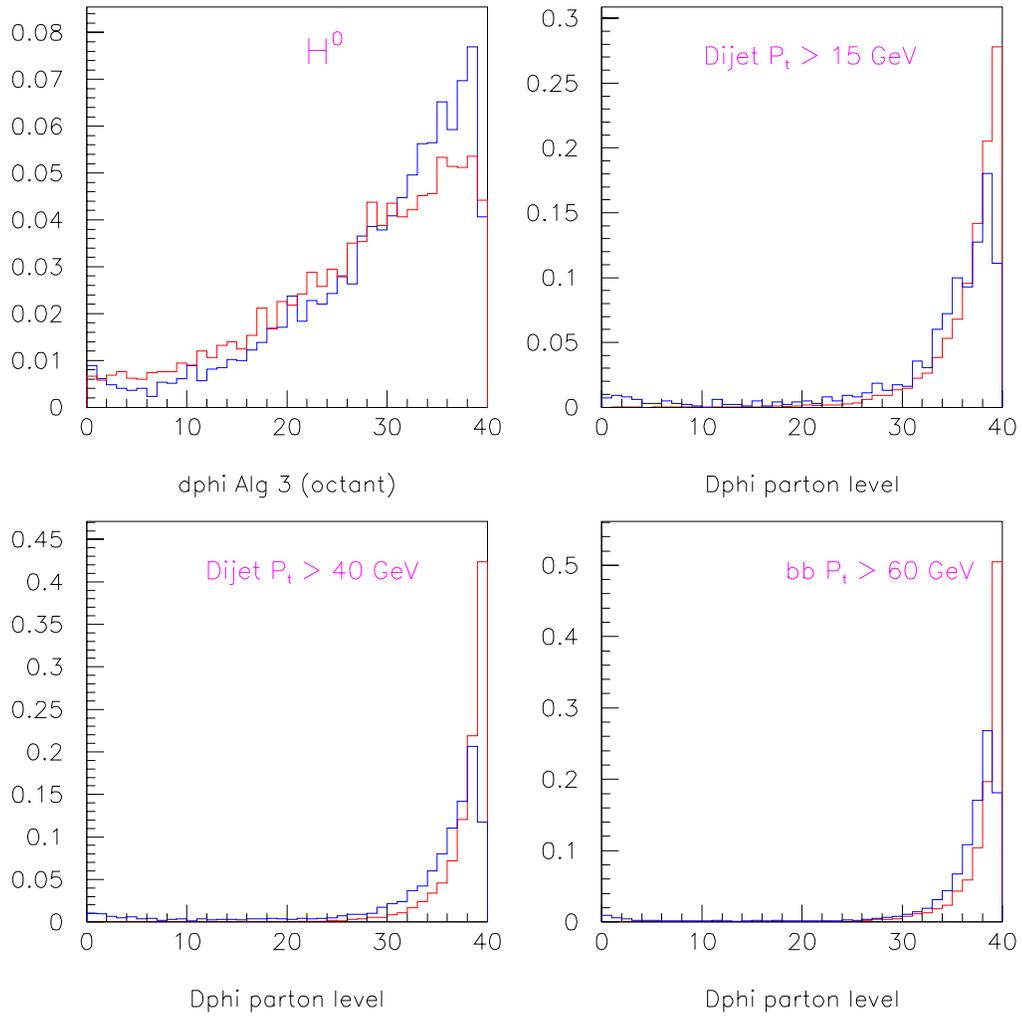


Figure 46. The opening angle between the leading partons.

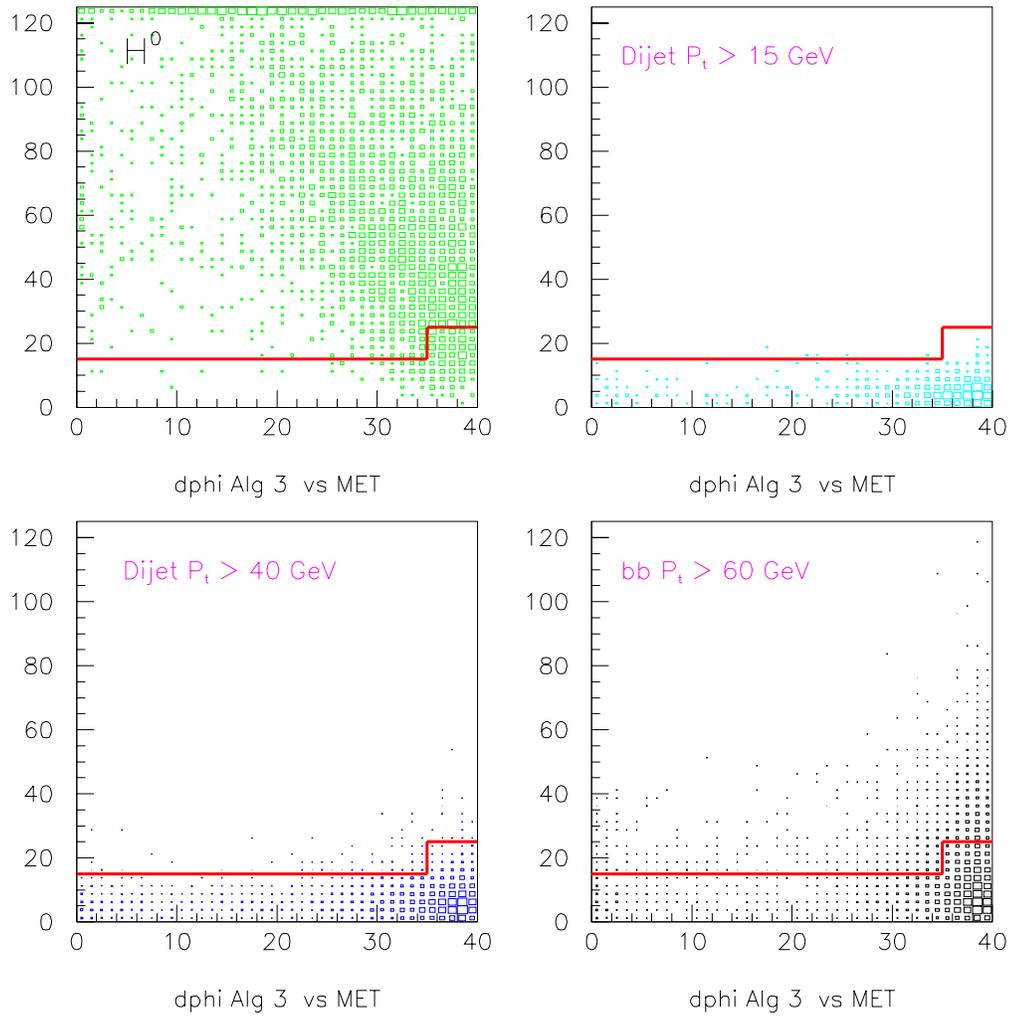


Figure 47. The correlation between $\Delta\phi$ and the missing E_T of the event for signal and representative QCD backgrounds

4.7.3 Rates and rejection improvements

The use of the CFT ϕ correlations becomes compromised at high instantaneous luminosity, as shown in Figure 48. Only for relatively high p_T jets does the correlation remain. At high luminosity one would benefit from the use of the calorimeter to confirm the CFT jets. Modest trigger thresholds are able to reduce the rate from low p_T scatters.

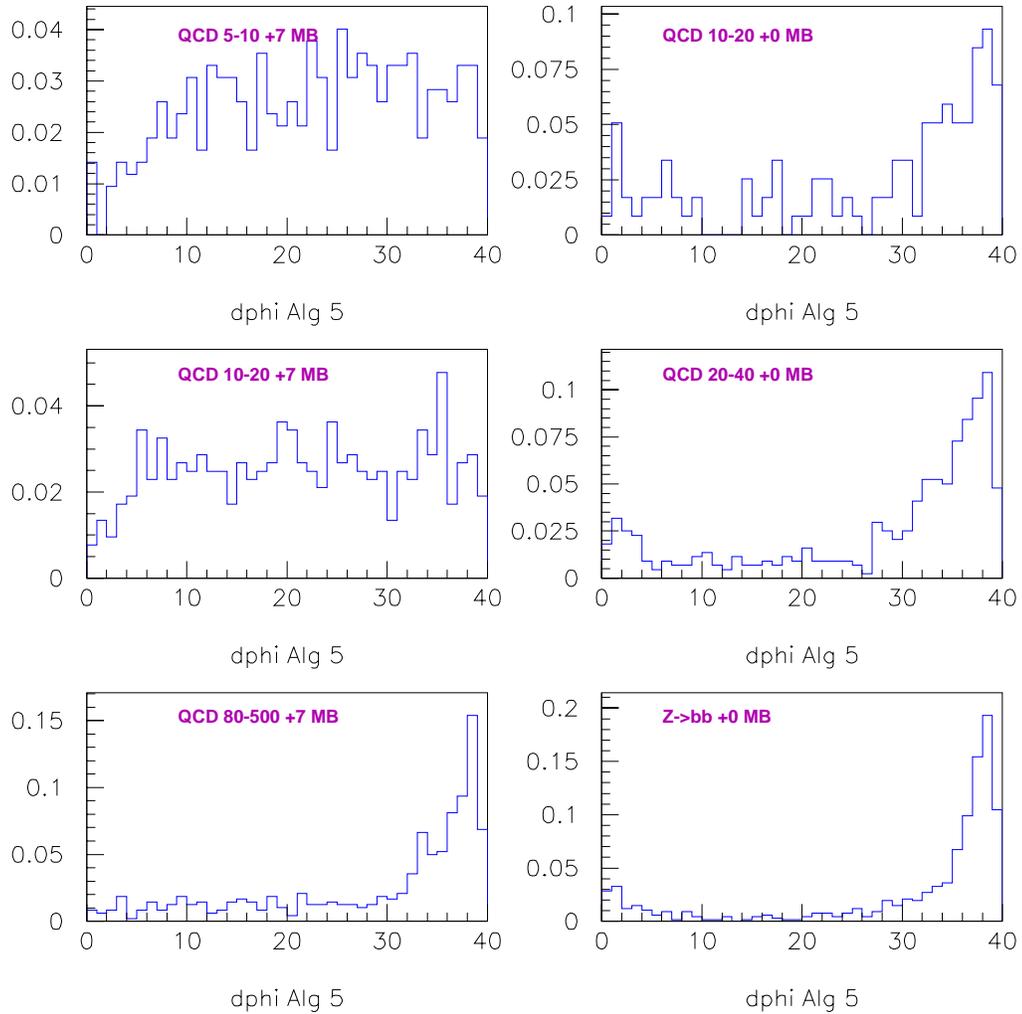


Figure 48. Correlation of $\Delta\phi$ in high luminosity conditions (left hand plots with 7 minimum bias events) and low luminosity (right hand plots with zero additional minimum bias events).

With an improved Level 1 calorimeter trigger that allows correlating CFT and calorimeter based jets these backgrounds can be further suppressed.

4.7.4 Comments

Substantial gain in trigger efficiency for channels such as $Z(\nu\bar{\nu})H$ through imposing cuts on topological correlations between jets. The calorimeter correlations are important extensions of the existing CFT correlations at high luminosity. These calorimeter correlations should be straightforward in the implementation described below using FPGA's in the upgrade calorimeter detector.

4.8 L1 Calorimeter Trigger Implementation

4.8.1 Constraints

Because the L1 calorimeter system needs to be integrated into the existing DØ DAQ system it must obey several constraints.

4.8.1.1 *Existing interfaces*

The interfaces of the new system to the existing hardware should be compatible. In particular the new system must interface to the input pickoff signals, the L1 framework, the L2 and L3 data, the clock, and the timing and control systems.

The layout of the existing input signal cables places a special constraint on the new system. Moving these cables from their current locations in the Moving Counting House would require an enormous effort. Physical locations of the boards in the new system will have to be adapted to the locations of the existing cables.

4.8.1.2 *L1 Latency*

The total L1 trigger latency is 4.2 μsec . After accounting for all transit times and front end processing, the maximum time remaining for complete signal processing, which includes digitization, filtering and the processing of the cluster algorithms is less than 2.7 μsec .

Additional, constraints are placed on the latency of the new system by the requirement of transmitting calorimeter clusters to the Cal-Track Matching system in time for it to meet the total L1 latency requirement. To match the arrival time of tracks from the L1 track trigger at the Cal-Track Match cards (900 ns after the bunch crossing of interest), the calorimeter trigger must send out its clusters within 1.0 μsec of the beam crossing.

4.8.1.3 *Adiabatic integration*

The installation and integration of the new system should be designed and built in such a way as to minimize the effect on data taking.

4.8.2 L1 Calorimeter Trigger Architectural Overview

A block diagram of the new L1 calorimeter trigger system is shown in Figure 49.

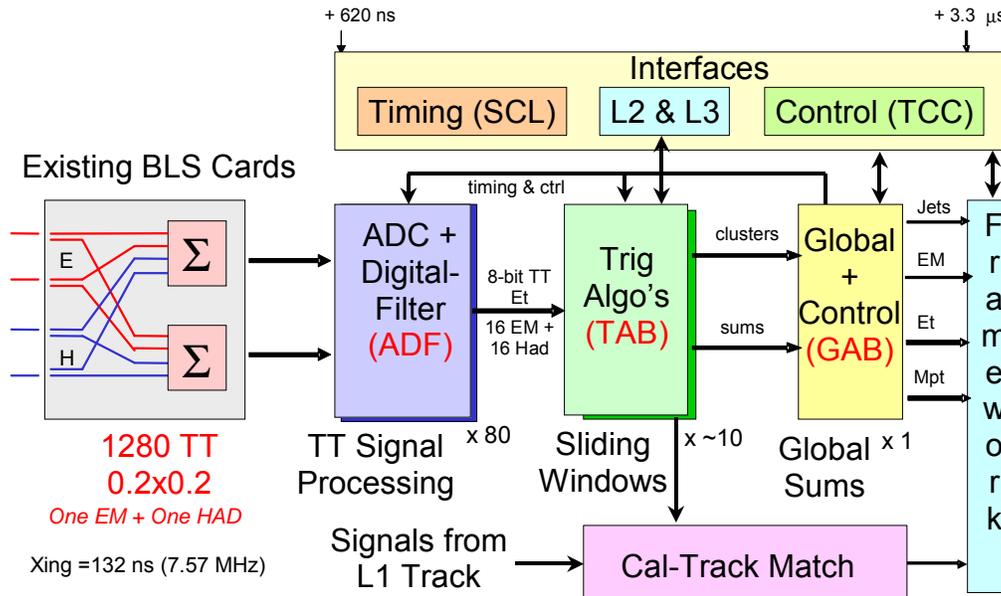


Figure 49. Block diagram of L1 calorimeter trigger, with the baseline subtractor (BLS) cards shown at left

The main elements of the system are listed below.

- ADC-Digital-Filter Boards (ADF) that receive analog TT signals from the BLS cards, digitize them, convert from energy to transverse energy (E_T) and perform the digital filtering to associate energy with the correct bunch crossing. Each of these boards deals with signals from 16 EM TTs and 16 H TTs, meaning that 80 such boards are necessary in the entire system.
- Trigger Algorithm Boards (TAB) that receive TT transverse energies from the ADF boards, produce electron and jet cluster E_T 's using the sliding windows algorithm and begin the global summing process that will yield scalar summed transverse energy ($E_{T,total}$) and missing transverse energy (Mp_T). Outputs will also be provided at this level (or at the GAB level) for data transmission to L2/L3 and to the Cal-Track Match system. Of order 10 of these boards are required
- A Global Algorithm Board (GAB) that receives data from the TABs and produces the final $E_{T,total}$ and Mp_T , as well as providing an interface to the DØ Trigger Framework and a timing fanout. One GAB is required for the system. It will be housed in the same crate as the TABs to facilitate communication between them.

These new electronics and associated communications cards will be housed in VME crates in two to three racks, replacing the present ADC and logic cards located in 10 racks.

Design work has started on all elements of the L1 calorimeter trigger system and a reasonably firm view of the general architecture has been reached. In particular, the chain of firmware for several versions of the sliding windows algorithm has been written and simulated on candidate devices. However, details of the implementation of all components are not complete. In the following we will describe the current state of the design, highlighting those areas in which substantial progress has already been made and pointing out parts of the system where a full implementation has yet to be developed. We will also mention design options that were considered, but rejected, where appropriate.

4.8.3 Trigger Tower Mapping

The 2560 EM and Had Trigger Towers (TT) map to an $\eta \times \phi$ grid of 40x32 cells with 0.2x0.2 extent. This grid extends from -4.0 to 4.0 in η and from 0 to 2π in ϕ . Cells are referenced by η - and ϕ -indices (*ieta* and *iphi*) which run from -20 - 20 (excluding 0) and from 1-32, respectively. This grid has been modified, for Run 2, from a strict geographic mapping at the extreme values of *ieta* (± 19 and ± 20) to allow inclusion of ICR detectors (the ICD and the MG) in the trigger.

The mapping of TTs in η is given in Table 14. All mapping in ϕ corresponds to the physical location of the TT in azimuthal angle. In the table, the column marked “Cable” signifies the labeling scheme for the cable carrying the data, while in the “Comments” column, the *EM* or *Had* refers to whether the relevant TTs appear physically in the EM or Hadronic parts of a shower.

Table 14: Trigger Tower mapping scheme.

TT <i>ieta</i>	Cable	Detector η and ϕ	Comments
all	EM & Had	$\phi = (iphi - 0.5) \times 2\pi/32$	
≤ 16	EM & Had	$\eta = (\text{sign } ieta) \times (ieta - 0.5) \times 0.2$	<i>EM</i> : include only EM cells <i>Had</i> : include only FH cells
17	EM	$\langle \eta \rangle = 3.45$	<i>EM</i> : include only EM cells
	Had	$\langle \eta \rangle = 3.45$	<i>Had</i> : include only FH cells
18	EM	$\langle \eta \rangle = 3.9$	<i>EM</i> : include only EM cells
	Had	$\langle \eta \rangle = 4.1$	<i>Had</i> : include only FH cells
19	EM	$\langle \eta \rangle = 0.95$	<i>Had</i> : MG – Cal <i>ieta</i> = 9,11
	Had	$\langle \eta \rangle = 1.25$	<i>EM</i> : ICD&MG – Cal <i>ieta</i> = 13
20	EM	$\langle \eta \rangle = 0.75$	<i>Had</i> : MG – Cal <i>ieta</i> = 8
	Had	$\langle \eta \rangle = 1.05$	<i>Had</i> : ICD&MG – Cal <i>ieta</i> = 10,12

4.8.4 ADF Implementation

In the Conceptual Design Report, two options were proposed for the design of the ADC boards. The "dumb ADC" option, where the ADC boards would just make the conversion of analog signals to digital format, was not retained. The ADC cards will perform analog to digital conversion and digital filtering of trigger pickoff signals to output for each channel and for each beam crossing an 8 bit E_T calibrated energy value.

4.8.4.1 ADC precision, rate and analog section

Analog signals will be digitized with 10-bit precision and an identical input voltage range for all channels. This scheme is adequate to guarantee 0.25 GeV of resolution and a 62 GeV scale in E_T for all η values. The proposed conversion frequency is 30.28 MHz (i.e. BC x 4). This rate is a good trade-off between the aim of a short conversion latency and the cost of a faster analog to digital converter. A possible choice for the ADC is Analog Devices AD9218 (dual 40 Mbps 3V 10-bit ADC). Conversion latency is 165 ns when operated at 30.28 MHz.

The analog section of the ADC card will include:

- A differential receiver whose input impedance matches that of the driving cable and that provides an attenuation factor to bring trigger pickoff signals in the desired range;
- Circuitry for baseline subtraction to adjust the offset of the previous signal; a digital to analog converter will produce a static correction voltage in a programmable way;
- The analog to digital converter previously mentioned;
- An anti-aliasing filter (e.g. a 1st or 2nd order Butterworth filter). The proposed cutoff frequency for this filter is 7.57 MHz given the fact that the spectrum of the signals to process does not have much content above that frequency. Tests will be conducted on the prototype and filter parameters will be modified if needed.

A candidate device for the differential receiver is Analog Devices low distortion differential ADC driver AD8138. An alternate design based on traditional operational amplifiers is also possible. A suitable device for the zero-offset digital to analog converter is Maxim octal 12-bit serial DAC MAX 5306. The anti-aliasing filter could be implemented with passive RC components placed on the feedback loops of the differential receiver and at the input of the analog to digital converter.

4.8.4.2 Digital filters

The current scheme for digital processing of trigger pickoff signals is a matched filter followed by a 3-point peak detector. Final E_T conversion including clipping and saturation will most likely be implemented in a look-up table. It is not foreseen to estimate the baseline dynamically and subtract it. Only static

adjustment will be made. The matched filter will operate at a rate of $BC \times 2$ or at the BC rate. Decimation by a factor of 2 or 4 on the stream of values converted by the ADC will be done prior to filtering. Selecting on a per tower basis what samples to keep for processing will allow to implement some coarse compensation of the delay due to cable length mismatches. While running the digital filter at $BC \times 2$ will produce better results than running at the beam-crossing rate, it also places more demands on the logic. Assuming that the convolver logic of the filter runs at 90.84 MHz (i.e. $BC \times 12$), the number of filter taps will be limited to 5 if the filter is operated at $BC \times 2$, but could grow to 11 if filtering is done at the BC rate. More simulations and testing are needed at that level to determine the optimum scheme, but these changes will not have a major impact on the design. Coefficient precision is still being studied. Simulations show that 5-7 bit unsigned coefficients can be adequate, but, depending on implementation, there could be a trade-off between coefficient and accumulator widths and the desired operating speed. No choice has been made yet on the target FPGA vendor or family, and several devices are being considered. Xilinx Virtex II devices comprise hardwired 18-bit x 18-bit multipliers that make this family attractive for filter design. Depending on bus widths, speed could be a limitation. Altera Apex II devices do not include hardwired multipliers, but have a large number of embedded memory blocks that make the implementation of convolvers in distributed arithmetic very compact and efficient compared to their multiplier-based counterparts. Implementation studies are being conducted to determine which device will best implement the target algorithm. Technical constraints and cost optimization will determine if implementation is made with a small number of very high density FPGAs or with a larger number of smaller devices.

4.8.4.3 Digital output to Trigger Algorithm Board

In order to simplify the design of the TAB, it seems preferable to make the necessary duplication of data at the level of each ADC card, although this doubles the number of cables between the two sub-systems. Each ADC card shall therefore include two identical output links. Assuming that an ADC card handles 32 channels and outputs 8 bit Et values, the net bandwidth per output link is 242.24 MByte/s. Several possible implementations have been proposed. These are discussed in more detail in section 4.8.5.

4.8.4.4 Clocks and miscellaneous signals

The ADC cards need a common clock and a few control signals. A Serial Command Link (SCL) interface will be needed to provide at least:

- The 7.57 MHz beam-crossing clock,
- The L1 accept signal
- 1 or 2 trigger "L1 qualifiers",
- A reset signal,
- A busy / error output signal.

Whether the ADC boards and TABs will share the same SCL interface is not yet understood. Given the low number of signals required by the ADC cards and the fact that the crates housing these cards would be fully populated in the proposed design, the SCL interface could be shared between the two sub-systems and should be placed on the TAB side. Signal fanout to the crates of ADCs would be done by this SCL interface; intra ADC crate signal distribution would be done by one ADC card via the back plane bus of the crate.

4.8.4.5 Downloading, Calibration, Monitoring and Data Acquisition

Each ADC card will include a slave A24/D16 VME interface (proposed device is Cypress CY7C964A) on a standard J1 VME connector. A VME Vertical Interconnect module will be placed in each crate housing ADC cards to make the interface to the Trigger Control Computer (TCC). This VME path will be used to download FPGA configuration, to control each ADC card (adjust zero-offset DAC, program digital filter coefficients, download lookup tables...), to run tests programs, and to read-back the raw ADC samples that will be captured following a monitoring request or following a self trigger when the ADC card runs in standalone data acquisition mode.

Calibration comprises two steps: coefficient determination and gain scaling. The first step consists in recording for each channel series of pulses in order to have sets of typical pulses shape. At this level there is no need to correlate measurements with the energy that is effectively measured by the precision readout. Once filter coefficients have been determined, these need to be scaled to give the correct transverse energy value. The exact mechanism to correlate the raw ADC data that will be captured via the TCC to the output of the precision readout is still being discussed. In calibration mode, each ADC channel shall be able to record all raw samples during a few bunch crossings. The system will be either self-triggered (e.g. freeze recording if the ADC value on a channel is above a certain threshold), or L1 triggered.

Monitoring will be done by capturing all ADC data on random beam crossings and for a fraction of L1 accept. This will allow verifying that interesting events are not missed and that digital filters operate properly. Each ADC card shall include sufficient buffering to keep all ADC raw samples and intermediate results during L1 latency. However, no buffering beyond L1 latency will be incorporated. No link between the ADC cards and the rest of DØ data acquisition is foreseen. The only fast data path of the ADC cards will be the links to the TAB. Following a L1 accept, the ADC boards will not re-send to the TABs the corresponding 2560 energy values and will probably not send either the series of raw ADC values that correspond to the event that caused the trigger. Both raw samples and filtered energy values will only be available for a very small fraction of L1 accepts, via TCC readout. The possibility to send raw ADC values to the TAB after each L1 accept is still being studied. The dead time that follows a L1 accept (which is introduced by other sub-systems in DØ) should be sufficient to make the transmission of many raw ADC samples to the TABs, but the motivation for doing so, and the large increase of bandwidth that would be required at the TAB to

L2/L3 level to capture all data is still being debated. A possibility would be to send raw samples for a fraction of L1 accepts.

4.8.4.6 Board implementation and crate layout

There are 1280 trigger towers and each trigger tower comprises an EM and an HAD channel. There are 2560 analog channels in total. The target density for an ADC card is 32 channels; that is 80 ADC cards in total. Each ADC card may be composed of a main carrier card that contains all the digital logic and several (e.g. 4) mezzanine cards that include the analog sections and ADC chips. Although this split will drive the cost up, it offers the advantage of modularity and eases parallel developments and independent tests of the analog and digital parts. The other option is to place both the analog section and digital circuitry on the same printed circuit board. Preliminary studies show that a 32-channel ADC card could fit in single width 6U format. All the analog and digital I/O connectors would be placed on full height 6U rear side transition cards at the back of housing crates. This design allows connecting 16 cables from the BLS to each ADC card without replacing the existing connectors or using patch panels. Because 16 such connectors occupy more than 6U in height, these are arranged in two interleaved rows of 8. The use of a standard VME64x backplane is foreseen; an alternative solution would be to use a standard 3U VME back plane on P1/J1 and design a custom back plane on P2. Each 21-slot crate shall contain 20 ADC cards and a VME controller. This density leads to 4 fully populated crates. If the target density of 32 channels per ADC card cannot be reached, double-width 6U modules are a possible option. In this case, 8 crates would be needed. Because full height 6U cards are foreseen for rear modules, the power supply of each crate must not obstruct this space. A candidate crate that satisfies this requirement is Wiener VME 6023 9U Plenum Bin crate.

4.8.5 ADF to TAB Data Transfer

As discussed above, each of the 80 ADF cards will send two copies of data from 16 EM TTs and 16 H TTs to the TABs bit-serially in 8-bit words. Data transmission between the ADFs and the TABs will be accomplished using LVDS links. This solution is the simplest and most cost effective way of transmitting the large amount of data in the system. Several options were considered to implement the LVDS technology: solutions with LVDS drivers and receivers implemented in FPGAs or the use of custom LVDS driver/receiver chipsets. Both schemes are acceptable at the level of the ADC cards. The FPGA solution offers the advantages of elegance, flexibility and lower component count but requires more coordinated engineering than a ready-to-use chipset. Constraints at the TAB level seem to favor the Channel Link solution.

4.8.5.1 LVDS Serializers and Deserializers

Unless unexpected issues are found or an alternate design is shown to be superior, ADC cards to TAB links will be based on 48-bit Channel Link chipset DS90CR483/484, which multiplexes its 48 parallel inputs to 8 data pairs. Data transport format, the possible inclusion of error detection, protocol and synchronization are still under discussion. A possible scheme to transmit the 32

8-bit values of an ADC card is to send 6 8-bit values at a time and repeat the operation 6 times. The link would be clocked at 45.42 MHz. Another option is to serialize each 8-bit value on one line; using effectively 32 of the 48 inputs of the Channel Link. The link would be clocked at 60.56 MHz. Both schemes are acceptable at the level of the ADC card. Although 33% of the link bandwidth would be wasted with the second option, simpler logic is needed and dropping 2 of the 9 cable pairs that are normally needed could lead to some savings in cable size.

4.8.5.2 Cables

As is evident from Table 19, a large number of cables will have to be plugged into each TAB. In the LVDS scheme outlined above, each of these cables would contain at least 7 wire pairs: 6 data pairs + 1 clock pair. These cables will have to plug into the front panel of the TABs because the back will be occupied by the VME bus and, possibly, a custom bus used to distribute timing and control signals. We are still searching for cables/connectors that best meet these requirements, but preliminary indications are that commercially available AMP cables with 2 mm hard-metric connectors may be usable.

4.8.6 TAB implementation

The Trigger Algorithm Boards (TAB) are the main processing component of the L1 calorimeter trigger. They perform the following tasks for a specific $\eta \times \phi$ region of the calorimeter.

- Find EM and Jet clusters and estimate their E_T using sliding windows algorithms.
- Send a list of EM and Jet clusters passing several different E_T thresholds to the GAB.
- Possibly send a list of EM and Jet clusters over threshold to the Cal-Track Match system.
- Calculate a scalar E_T sum and E_x and E_y sums for all TTs in the TAB's region for use in $E_{T,\text{total}}$ and M_{p_T} .
- Send the E_T , E_x and E_y sums to the GAB.
- Format and buffer data for use in L2 and L3 and send this out on receipt of L1 and L2 trigger accepts.
- Allow insertion of fake data, by computer, directly after the inputs to test system functionality.

4.8.6.1 Overall TAB Architecture

A block diagram of the TAB is given in Figure 50. The architecture is driven by three main considerations: the large number of TTs whose signals must be sent to a given TAB to find E_T cluster local maxima, the latency requirement and the desire to use small, low-speed FPGAs to keep costs down.

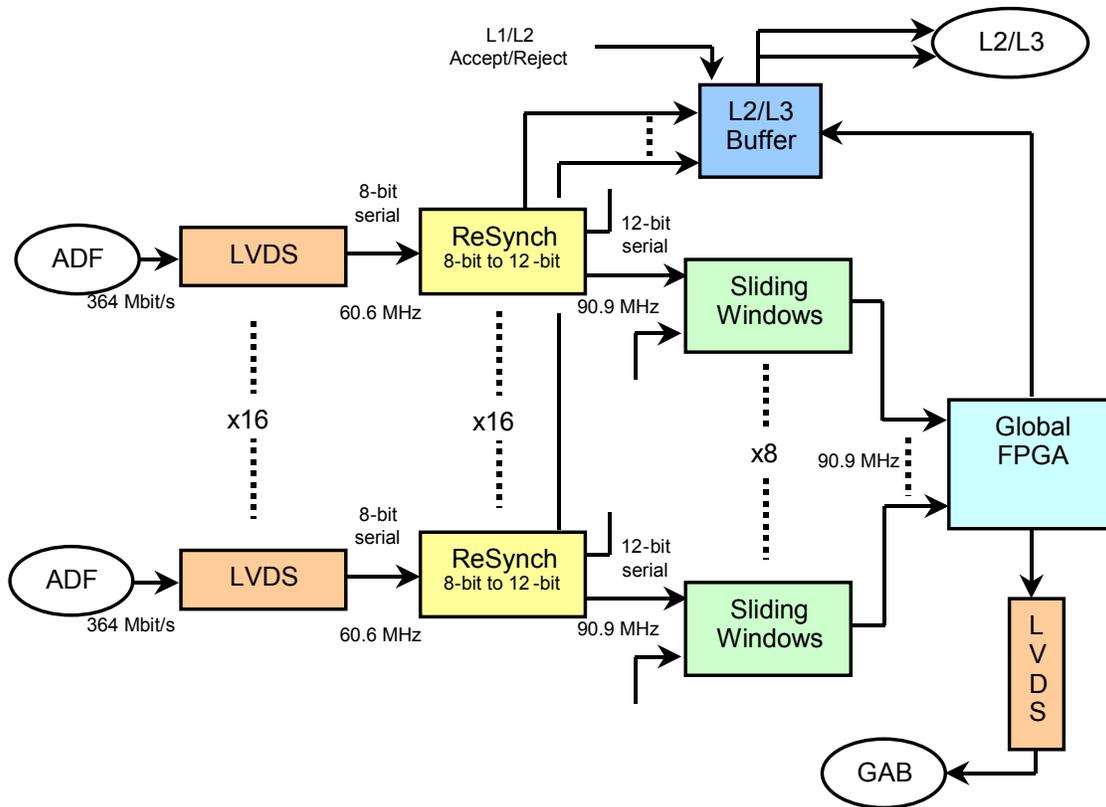


Figure 50: Block diagram of the TAB.

The functional elements of the TABs are the following.

- LVDS receivers for input signals from the ADFs (more details are given in the next section). These produce as output 32 streams of 8-bit serial data corresponding to the EM and H E_T 's for each bunch crossing (BC) from the ADF at 60.6 MHz (8xBC).
- Resynchronizers, which align the serial data streams and retransmit them as 12-bit serial streams, padded with zeroes, at 90.9 MHz (12xBC). The 12-bit range of these outputs is necessary to deal with carries in the serial adders described below. The resynchronizers also fanout the signals for use in several sliding windows chips.
- Pre-summers that add ICR energies to the calorimeter EM or H TTs from the region to which they belong. This option, which is still under study, would enhance energy resolution for jet clusters found in the ICR.
- Sliding Windows FPGAs that implement the sliding windows algorithms for both EM and Jet clusters for a sub-set of the data considered by the board and also perform the first step of global summing.
- A Global FPGA that gathers information from the Sliding windows FPGAs, constructs the list of Rols passing EM and Jet thresholds, does another step in the global summing chain and sends out information from

this TAB to the GAB. This chip may also prepare data for transmission to the Cal-Track Match system.

- An LVDS transmitter that receives data from the Global FPGA at 90.9 MHz and sends it to the GAB.
- At various stages in the processing data is buffered for transmission to L2 on the receipt of an L1 accept and for transmission to L3 on an L2 accept.

Preliminary versions of firmware have been written for all relevant elements of this chain and candidate devices have been chosen.

4.8.6.2 Serial Arithmetic Building Blocks

The sliding windows algorithm is basically a set of sums and compares. These operations can be performed efficiently in a bit-serial manner that minimizes FPGA resources required, particularly the number of input data pins. This serial structure also meshes well with the serialized data arriving from the ADFs. The two basic arithmetic/logic elements required for the sliding windows algorithm are the addition of two numbers and the compare of two numbers. All elements of the algorithm can be constructed from trees of these primitives. Diagrams of a bit-serial adder for two data lines and for a bit-serial comparator, as implemented in first versions of the sliding windows firmware, are shown in Figure 51 and Figure 52.

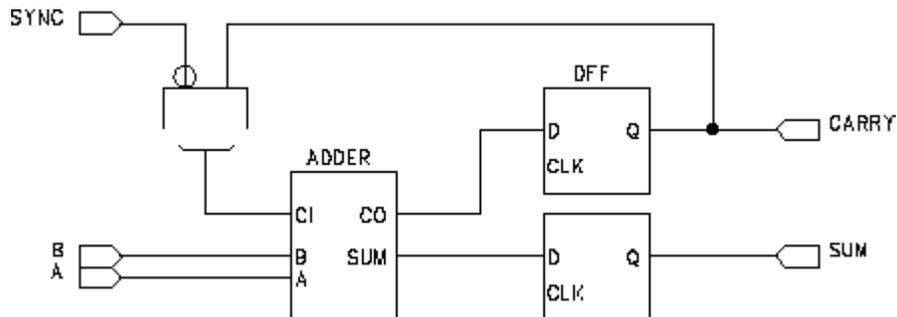


Figure 51 : Serial adder for data A and B. SYNC is a signal that separates one serial word from the next.

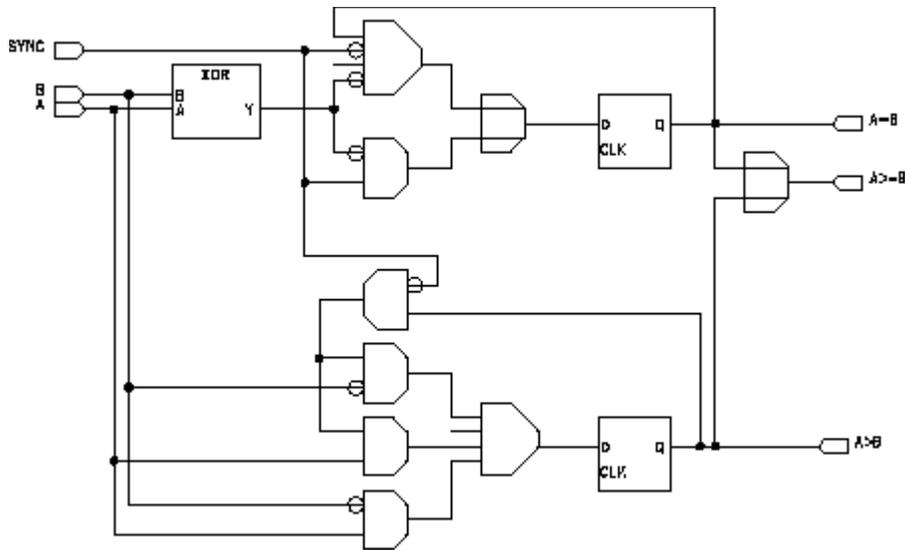


Figure 52: Serial comparator for data A and B. Outputs "A>B" and "A>=B" are both required to unambiguously assign local maxima.

4.8.6.3 Sliding Windows Algorithms Implementation

EM and Jet cluster finding algorithms are implemented on a set of FPGAs on the TAB. Each FPGA finds both EM and Jet clusters in a sub-set of the data considered by the TAB. Data sharing is discussed in more detail below. As mentioned previously, the sliding windows algorithm is a method of finding local maxima in E_T deposition on the $\eta \times \phi$ grid. EM clusters are defined using EM energy, while Jet clusters use EM+H energies for each TT. Local maxima are found by comparing the energies of regions of interest (Rols) consisting of several TTs. Each Rol, labeled by one of its constituent TTs, is compared with a group of nearby Rols to determine if its E_T is the highest in the local area. Care must be taken to avoid double counting of local maxima by using either ">" or ">=" for the comparisons. Additional isolation criteria are then applied to EM Rols that are local maxima. Finally, EM and Jet cluster energies are formed by adding TT E_T 's in a region around the Rol. In the EM case, the cluster region is taken to be the same size as the Rol.

The decision as to exactly which EM and Jet trigger algorithms to use in the final system has yet to be made. It will be based on simulation results that are ongoing. Preliminary studies, presented in section 4.6, indicate that the sliding windows concept provides good results, however, the specific algorithms discussed below are still being simulated. Even after baseline algorithms have been chosen, though, experience gained in Run 2a and 2b may indicate that other choices would give better performance. Flexibility in the algorithms implemented in the TAB firmware is therefore crucial to the success of the new L1 calorimeter trigger.

EM Cluster Algorithms

Even the most ambitious EM cluster finding algorithm is significantly less complex than the Jet algorithms, so we have only examined one possibility there.

It is summarized pictorially in Figure 53. It uses an RoI of 2x2 TTs in EM energy, which is also used as the cluster E_T region. EM clusters are vetoed by significant E_T in the EM TTs directly surrounding the RoI (the 4x4 rim) or by E_T in the Had TTs in a 4x4 region directly behind the RoI that is larger than a fixed value or larger than a fixed fraction of the RoI E_T . Modifications to this should be possible without strain on the system resources.

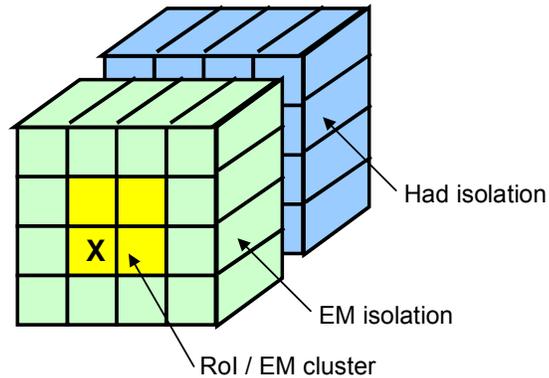


Figure 53: A diagram of the EM sliding windows algorithm.

A schematic of the EM algorithm (without the local maximum finding) shown in Figure 54. It includes the RoI sum, EM isolation sum and Had isolation sum, a comparison of the EM isolation sum with a downloaded threshold and a comparison of the Had isolation sum with $\text{RoI-sum}/8$ (shifted by 3 bits). Also included are threshold comparisons for the RoI for five thresholds. The local maximum finding schematic is shown in Figure 55 where the RoI labeled “E[1][1]” is the candidate RoI.

em/emwindow

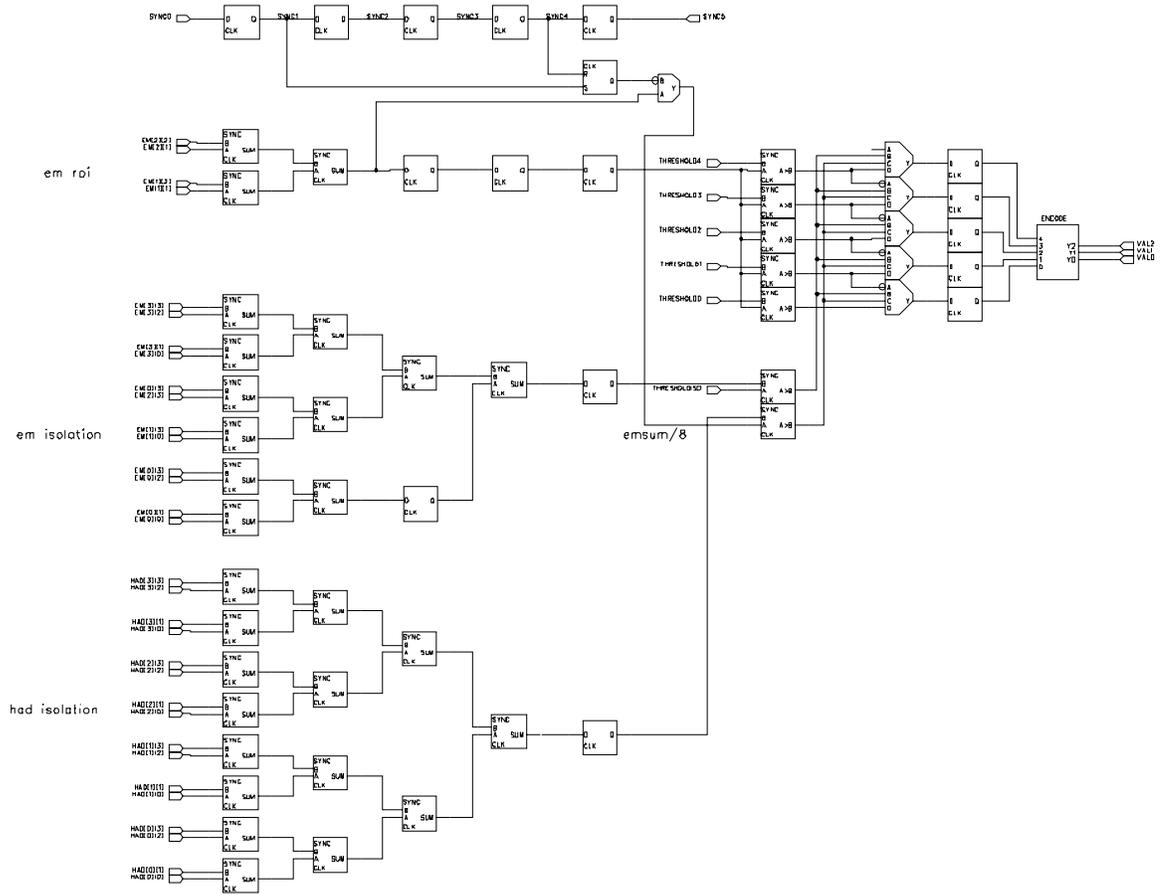


Figure 54: Schematic for the EM sliding windows algorithm.

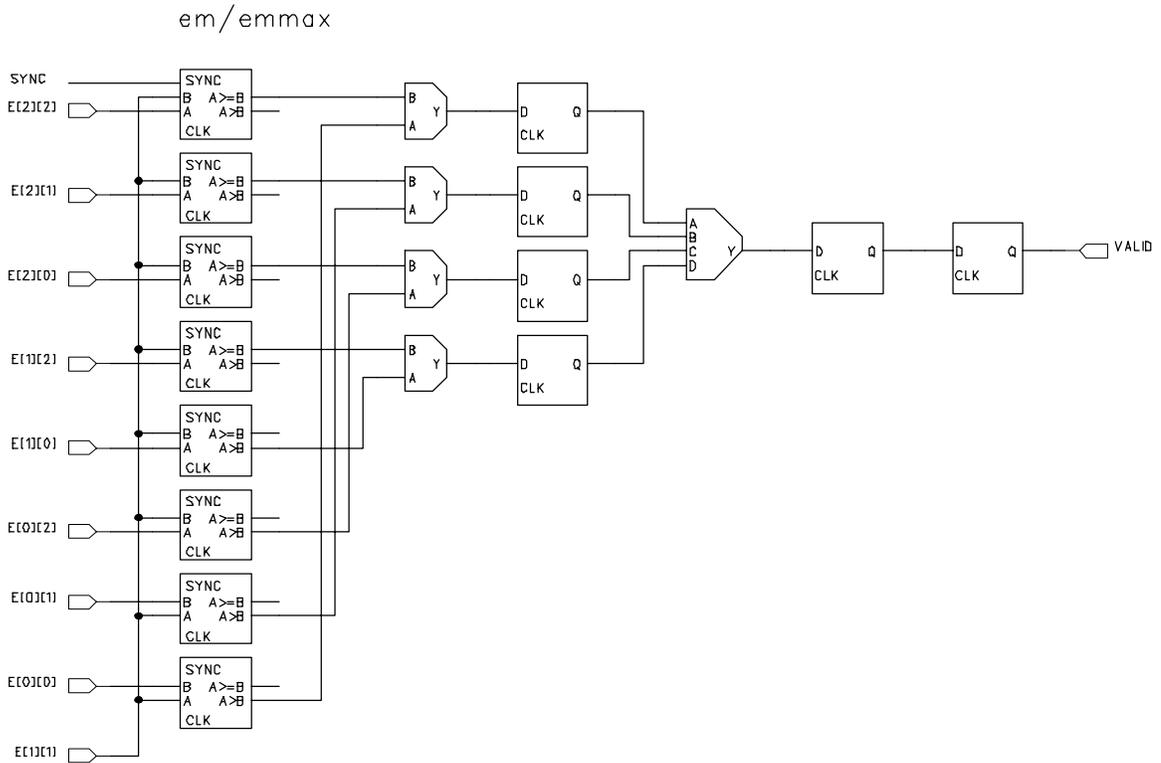


Figure 55: Schematic of local maximum-finding algorithm for the EM sliding window.

Jet Cluster Algorithms

Because of its potential complexity we have considered several possible variants of the sliding windows algorithm for finding Jet clusters with different data and FPGA resource requirements. The parameters we have examined here are: region of interest (RoI) size, region over which compares are made to determine if a given RoI corresponds to a local E_T maximum (decluster region) and size of region over which cluster energy is calculated (cluster E_T region). The E_T sum in an RoI is used as the basis for local maximum finding. The RoI should therefore be chosen to match the approximate size of jet energy deposition cores. The E_T sum in a given RoI is compared to that in neighboring RoI's, with data contained in the declustering region. Larger declustering regions mean smaller energy sharing between nearby local maxima. Finally, the cluster E_T region should be chosen as a tradeoff between including all of the jet energy in the cluster and excluding as much noise as possible. A summary of the options studied is given in Table 15 and a diagram of option (b) is given in Figure 56.

Table 15: Jet algorithm options. Sizes are given in units of TTs in $\eta \times \phi$.

Jet Algo	RoI Size	Decuster Region	Cluster E_T Region
(a)	3x3	5x5	5x5
(b)	3x3	7x7	5x5
(c)	2x2	6x6	4x4

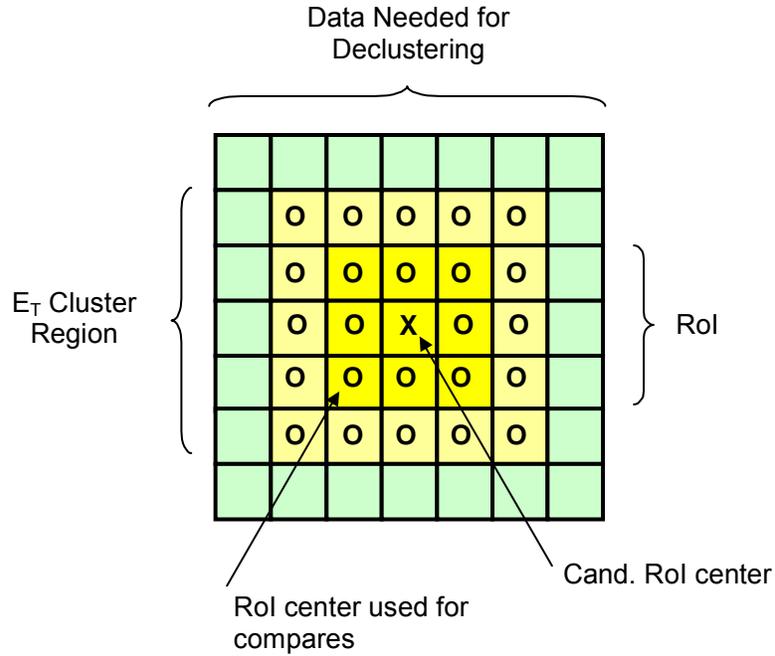


Figure 56: Diagram of Jet algorithm (b).

Threshold Comparisons

The E_T 's of EM and Jet clusters that are found to be local maxima by the sliding windows algorithm are compared with two sets of up to seven E_T thresholds, which can be downloaded at initialization. The results of these comparisons are numbers 1-7 encoding the highest threshold passed. The Sliding Windows chip then sends out information about all RoIs that it considers, as 3 bits for each RoI packed with 4 RoIs a in 12-bit word. The 3 bits for a given RoI are either 0, if the RoI is not a local maximum, or 1-7, corresponding to the highest threshold it has passed if it is a local maximum.

E_T Sums

Preliminary sums for use in calculating $E_{T,\text{total}}$ and Mp_T are also computed on the Sliding Windows chips. Each chip sums all TT EM and EM+H E_T 's over all η for each ϕ considered in the chip.

Sliding Windows Chip Outputs

The outputs of the Sliding Windows chips sent to the Global chip are then a set of 12-bit words sent serially at 90.9 MHz as shown in Table 16.

Table 16: Data sent serially from the Sliding Windows chips to the Global chip as 12-bit words.

No. 12-bit serial words from chip	Description
$N(\text{Rols in chip}) / 4$	Encoded EM-cluster data (3-bits per Rol)
$N(\text{Rols in chip}) / 4$	Encoded Jet-cluster data (3-bits per Rol)
$N(\phi \text{ segments in chip})$	EM E_T sums over η
$N(\phi \text{ segments in chip})$	EM+H E_T sums over η

4.8.6.4 Global FPGA

The Global FPGA receives information from the TABs outlined in Table 16. It collects the encoded EM and Jet data for output and calculates E_x and E_y , derived from the E_T ϕ -sums bit-serially using x, y weights stored in ROM as a Look-Up-Table. The Global chip could also “summarize” the EM and Jet data to reduce output data volume. For example, counts of EM/jet objects above threshold could be produced.

It is also likely that the Global chip will provide data to the Cal-Track match system. This is discussed further in section 4.8.10.

4.8.6.5 Output to the GAB

Data produced in the Global FPGA is sent as 12-bit words serially to an output LVDS driver. The same National Channel Link chipset as used for the LVDS inputs to the TABs could also be used here (see the following section), allowing up to 48 words to be sent bit-serially in parallel.

An example of the data that could be passed to the GAB is given in Table 17. Sending E_T sums by ϕ would allow the possibility of topological E_T triggers and could be used as input to the Cal-Track Match system.

Table 17: Possible list of data transferred serially as 12-bit words from the TABs to the GAB.

No. of 12-bit words	Description
$N(\text{Rols in TAB}) / 4$	Encoded EM-cluster data (3-bits per Rol)
$N(\text{Rols in TAB}) / 4$	Encoded Jet-cluster data (3-bits per Rol)
$N(\phi \text{ segments in chip})$	EM E_T sums over η
$N(\phi \text{ segments in chip})$	EM+H E_T sums over η
2	E_x and E_y

4.8.6.6 Latency

A summary of the latency introduced at various stages of the TAB processing is given in Table 18. The pipelined, serial nature of the algorithms implemented in firmware allow the latency to remain very small while still permitting the use of relatively modest FPGAs.

Table 18: Latency in the TAB in units of bunch crossing times (132 ns).

Stage	Latency (BC)	Comments
Resynch	1.0	changing serialization frequency
Sliding Windows	<1.0	pipelined logic (~10 stages @ 11 ns / stage)
	1.0	serializing output streams
Global	1.0	$E_{x,y}$ calculations
	1.0	serializing output streams
Total	~5.0	= 660 ns

4.8.6.7 Impact of Data Sharing on the TAB Architecture

The single most important consideration driving the structure of the TABs is the sharing of data, both between boards and between elements on a single board. Since any local maximum finding algorithm requires access to data from a region around the element being considered, the same data element (in this case a TT E_T) must be used in several places. In a scheme where different physical devices check if different sets of Rols correspond to local maxima, some TT data must be duplicated and sent to several devices.

We have decided to use data flow architectures such that two copies of the TT data are made on the ADF boards and sent out on two cables to different TABs. We have also considered a scheme where only a single output cable is used for each ADF, with data duplication being done at the TAB level. While this idea halves the number of ADF-to-TAB cables required, it was rejected because the large density of signals that would have to be passed between the TABs would require costly transmission methods and because of the increased latency that would result from receiving and retransmitting input data.

The requirement of two-fold data duplication at the ADF sets the minimum number of Rols that a single TAB must check as possible local maxima. We are considering schemes where regions of n TTs in η by 32 TTs in ϕ (a full ϕ -slice) or 40 TTs in η by n TTs in ϕ (a full η -slice) are sent to each TAB. The full ϕ -slice option eliminates “wrap-around” effects by sending all relevant ϕ information to a single board. The full η -slice option makes ICR data available for inclusion into jet cluster energies. In both schemes ICR energies are included in $E_{T,total}$ and M_{p_T} . The number of Rols checked as local maxima, the total amount of data required to perform the local maximum checks for the Rols considered and the

number of cables from the ADFs into each TAB for each of the Jet sliding windows algorithms discussed above are given in Table 19. Also shown in this table are the number of FPGAs required on each board to process the data for that board's region and candidate Altera devices for which firmware has been written to implement the algorithm and its input and output. For all cases, both Jet and EM algorithms fit on a single chip and the lowest speed class of the device family can be used.

Table 19: Data required by one TAB for the Jet algorithms shown in Table 15.

Jet Algo	Rols/1 TAB ($\eta \times \phi$)	Data/1 TAB ($\eta \times \phi$)	No. TTs/1 TAB (EM+H)	# Cables	# chips per TAB	Cand. Altera Device
(a)	4x32	8x32	512	16	8	20K160-3
	40x4	40x8	640	20	10	
(b)	6x32	12x32	768	24	6	20K300-3
	40x6	40x12	960	30	7	
(c)	5x32	10x32	640	20	8	
	40x5	40x10	800	25	10	

In the schemes shown above, data density and FPGA size increase rapidly with increasing Jet algorithm complexity. However, the number of chips and the number of boards required decrease as more data is sent to a single board.

4.8.7 GAB Implementation

The Global Algorithm Board (GAB) is the final step in the L1 calorimeter chain before the Trigger Framework. Its exact functionality is still under discussion, however, it seems sensible to push as much communication as possible with systems outside the L1 calorimeter trigger off on this board, since only one of them will be required. An incomplete list of the tasks that the GAB may perform is given below.

- Calculates $E_{T,total}$ and E_x , E_y from the TAB partial sums and transmits them to the Trigger Framework
- Reformats the lists of EM and Jet clusters above threshold and transmits these to the Trigger Framework
- Receives timing and control information from the Trigger Framework over the Serial Command Link (SCL) and fans this out to the ADFs and TABs as necessary.
- Sends L1 calorimeter information to the Cal-Track Match system (if this is not done by the TABs).
- Sends data to L2 and L3.

4.8.7.1 *Output to the Trigger Framework*

Options for the format of the output to the Trigger Framework as well as the transmission method are still being explored. All quantities currently sent to the Trigger Framework by the L1 calorimeter trigger (see section 4.3) will continue to be sent though.

4.8.8 TAB-GAB Crate

The TABs and GAB will be housed in one 9U VME crate. This crate will have standard VME connectors on P0, P1 and P2 but will probably require a custom J3 backplane. A VME CPU will also be included in this crate for use in downloading, monitoring and testing.

4.8.9 Control and Timing Signals

The TABs and ADFs need to receive global control signals such as a bunch crossing clock and some sort of event number. This information is available from the Trigger Framework via the Serial Command Link (SCL) using custom SCL mezzanine cards. However, the ADFs and TABs need only a small subset of the full SCL information. It is therefore more efficient to receive this information centrally in the GAB and then fan out the required sub-set to the ADFs and TABs. The precise method by which this fanout will be accomplished is still being discussed, however, using a custom backplane is a likely candidate in the TAB-GAB crate. In the ADF crates, one ADF will receive timing/control information from the GAB and distribute this to the rest of the crate over the backplane.

4.8.10 Output to the Cal-Track Match System

We are still considering possibilities are for transmission of data from the L1 calorimeter trigger to the Cal-Track Match system. The structure of the L1 calorimeter trigger and of the Cal-Track Match system impose several constraints on this output. Some of these constraints are listed below.

- **Latency:** Signals from the L1 calorimeter trigger must be sent out within approximately 1 μ sec of the bunch crossing to match the arrival time of data from the L1 Tracking trigger, \sim 900 ns after the bunch crossing.
- **Data Volume:** The MTCxx cards used in the Cal-Track Match system (see section 5) are designed to receive 16-bit data words at 53 MHz. A maximum of 7 such words can be received in one bunch crossing interval, setting a limit on how much data can be transmitted per event per cable of 112 bits.
- **Cabling:** Although the MTCxx boards have 16 serial input connections, most of these are used to transport data from the L1 track trigger and possibly from the L1 preshower trigger. This leaves 2-4 open inputs for data from the L1 calorimeter trigger.
- **Data Selection:** The serial architecture of the TAB boards puts constraints on how data can be sent out from the boards. In particular, sorted lists are

difficult using the current system, but the highest E_T object in a list, a list of objects over threshold or counts of objects should be possible.

A possible scheme, which meets all of these requirements, would have the Global chip on each TAB send the E_T of the highest E_T EM and Jet cluster in each of the ϕ -slices considered by that board. A count of EM and Jet objects over threshold in each of the ϕ -slices could also be included in the data word for that slice. To be implemented effectively, this would require each TAB to consider the entire η -span of the detector (*i.e.* the $40 \times n$ options in Table 19). Either one or two copies of this data could be sent to the Cal-Track Match boards depending on whether data sharing is required by the matching algorithm used. Detailed estimates of the latency involved in this scheme have yet to be made.

The data transmitted from the L1 calorimeter trigger to the Cal-Track Match system in the tentative scheme is shown in Table 20.

Table 20: Possible data transmission to the Cal-Track Match system for two of the TAB configurations given in Table 19.

TAB data region ($\eta \times \phi$)	# Cables to Cal-Trk	Data packed in 8-bit words per ϕ -slice	# bits per BC
40x4	8/16	EM: E_T of highest cluster + count clust's over threshold	32
		Jet: E_T of highest cluster + count clust's over threshold	32
		<i>Total</i>	64
40x6	6/12	EM: same as above	48
		Jet: same as above	48
		<i>Total</i>	96

4.8.11 Output to L2 and L3

The L1 calorimeter trigger currently sends the same data to L2 and L3 using G-Link fiber optic cables. Ten such cables are optically split to provide data for L2 and L3 from each of the ten racks in the system, corresponding to data from an $\eta \times \phi$ region of 4×32 TTs. The data structure is given in Table 21. In the table, "seed masks" have a bit set if the corresponding TT has E_T above a run-time defined threshold (typically 2 GeV). The "TT E_T 's" are the 8-bit E_T from the digitization for each TT.

Table 21: Run 2a L1 calorimeter data to L2 and L3.

Starting Byte #	# of Bytes	Data
1	12	L2 Header
13	16	EM TT seed mask
29	16	EM+H TT seed mask
45	128	EM TT E_T 's
173	128	EM+H TT E_T 's
301	4	L2 Trailer

A similar data set will be transmitted to L2 and L3 in the new system, with the TT seed masks being replaced by encoded cluster data. However, more information may be desirable. The necessity of passing along all of the TT E_T 's will probably require that this L2/L3 data be sent from the individual TABs.

4.8.12 Downloading, Testing and Monitoring

4.8.13 Milestones and cost estimate

Since a detailed design has not yet been made, the details of the schedule, resources and cost estimate are necessarily rather fluid. However, we are in the process of creating a resource loaded schedule to identify critical areas. We have also prepared a list of major milestones in the project, as shown in Table 22.

Table 22. Major milestones for L1 calorimeter trigger project.

Milestone date	Task
03/02	Performance specification defined, major design choices made
10/02	Prototype design complete
02/03	Prototyping complete (ADC, TAB,...)
06/03	Prototype testing complete
12/03	Preproduction complete
04/04	Preproduction testing complete
11/04	Production complete, testing begins
03/05	Installation & commissioning starts
09/05	Commissioning complete

4.8.14 Cost estimate

A preliminary cost estimate for the Level 1 calorimeter trigger upgrade is presented in Table 23 below. Most of these M&S funds will be needed in FY03 and FY04.

As described in the above sections, much progress has been made in the design of the L1 calorimeter trigger upgrade system. Since complete designs are not available for any of the elements of the system, however, the base cost for most of the items below has been estimated from previous projects in Run 2a that required similar boards, power supplies, backplanes, and other elements.

Table 23. Preliminary cost estimate for the Level 1 calorimeter trigger only. A contingency of 50% has been applied to equipment. Prototypes are estimated to cost a factor of three more than the production components.

Item/process	Unit Cost (\$)	# Required (+ spares)	Total Cost (\$k)	Total Cost + Contingency (\$k)
ADF System				
ADF cards	3500	80+10	315	472
ADF prototypes	10500	5	52	79
Crates & backplanes	6500	8+1	58	88
Power supplies			21	32
VME masters	3250	8+1	29	44
<i>Subtotal</i>			<i>475</i>	<i>713</i>
TAB System				
TABs	7500	10+4	105	158
TAB prototypes	22,500	2	45	68
GABs	4000	1+1	8	12
GAB prototype	12000	1	12	18
Crates, backplanes	6500	1+1	13	20
Power supplies	1000	1+1	2	3
VME Masters	3250	1+1	7	10
<i>Subtotal</i>			<i>192</i>	<i>288</i>
Cables (ADF-TAB + Cal-Trk)	75	300+30	26	39
Engineering			650	650
TOTAL			1344	1,690

4.9 L1 Calorimeter Summary & Conclusions

The high luminosity of Run 2b presents a significant challenge to the L1 calorimeter trigger. The L1 calorimeter trigger upgrade proposed in this section addresses these challenges.

We will need to be able to effectively identify calorimeter energy depositions with the correct bunch crossing – this is addressed by digital filtering techniques. Since almost 80% of the L1 rate is calorimeter based, the importance of sharpening the p_T cut (and thus reducing background rates) as well as the possibility of triggering on real objects such electromagnetic clusters and jets is clear, and being addressed by a “sliding window” technique.

The improvement in global variables such as missing E_T can also be improved with the addition of the energy from the ICR region at L1. The ability to do that has been provided in the present front-end electronics.

Finally, the additional power provided by current FPGA's will allow the migration to L1 of more sophisticated algorithms and topological cuts presently available at L2.

This set of tools provided by the proposed L1 calorimeter trigger will allow us to make the optimal use of the Run 2b luminosity.

The hardware implementation of these improvements has been explored, leading to an overall architectural design. Preliminary detailed designs for several of the most important elements of the system have been made, with a full first-pass design anticipated by summer 2002.

5 Level 1 Calorimeter-Track Matching

5.1 Overview

The goal of the L1CalTrack trigger is to exploit matches in the ϕ position of tracks from the L1CTT trigger with that of EM and jet objects from the L1Cal trigger in order to reduce the L1 trigger rates of EM and track triggers. Information from the Central Preshower (CPS) and Forward Preshower (FPS) detectors is also used. Monte Carlo studies show that the improvement in the reported ϕ position of EM objects at the trigger level from 90° to 11.25° can reduce medium P_T electron triggers by a factor of 2-3. Additionally, large factors of rejection (10-70) can be achieved by matching track triggers with calorimeter towers of modest energy. This latter is important in triggering on hadronic tau decays such as in $H \rightarrow \tau^+ \tau^-$.

The implementation of the L1CalTrack trigger uses the **existing** L1Muo architecture with small modifications. This is sensible since the L1Muo trigger matches the ϕ position of tracks from the L1CTT trigger with that of muon objects derived using muon scintillation counter hits, a similar function to the L1CalTrack trigger. The huge advantage of this implementation is that the L1Muo trigger has been successfully running since the start of Run 2. Thus issues such as synchronization, buffering, outputs to L2 and L3, electronics testing, monitoring, power supplies, and rack infrastructure have proven, working solutions.

5.2 Simulation

5.2.1 Improving Calorimeter EM Rates Using the L1CTT

The L1CTT trigger is not yet operational on DØ hence we rely on Monte Carlo studies at present to estimate the gains of an L1CalTrack trigger. The simulation effort must be improved (by including CPS and FPS information for example) and cross-checked with collider data. Nevertheless, the existing Monte Carlo studies indicate that Run 2a electron and track trigger rates can be reduced by the addition of the L1Cal Track trigger. The reasons for this rejection are the improved ϕ granularity of EM and jet objects from L1Cal and the fact that the fake rates in the calorimeter and central fiber tracker are relatively uncorrelated.

This latter point is shown in the following studies that match EM objects from the calorimeter with tracks from the L1CTT. The calorimeter EM objects are found with different E_T thresholds. The L1CTT tracks have $P_T > 1.5$ GeV/c unless otherwise noted. A calorimeter-track match is defined by matching the calorimeter EM trigger towers ($\Delta\phi=11.25^\circ$) with tracks from the three overlapping L1CTT track sectors (each $\Delta\phi = 4.5^\circ$). QCD jet events were used to simulate the background.

Results are shown in Table 24. The left-hand column gives the E_T threshold for the calorimeter EM objects. The denominator in the subsequent columns is the number of EM objects (trigger towers) exceeding each E_T threshold. The numerator is the number of EM object-track matches. Results at

two different luminosities and two different QCD regimes are shown. For nearly an order of magnitude increase in luminosity, the rate of correlation between trigger towers of significant energy and high- p_T tracks increases by less than 10% in each case. This suggests that track-calorimeter matching will continue to be a powerful tool for background rejection at the highest luminosities.

Table 24 Trigger-tower-track occupancy for 2 GeV and 20 GeV QCD jet k_T and different tower E_T thresholds for low ($4 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$) and high luminosity conditions ($5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$).

EM E_T (GeV)	Jet $k_T > 2$ GeV $4 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$	Jet $k_T > 20$ GeV $4 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$	Jet $k_T > 2$ GeV $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$	Jet $k_T > 20$ GeV $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$
>0.5	9k/197k (4.6%)	42k/161k (26%)	200k/1520k (13%)	92k/291k (33%)
>2	69/297 (23%)	4k/7506 (53%)	1100/3711 (30%)	2130/3482 (61%)
>5	5/9 (50%)	920/1587 (58%)	52/132 (39%)	480/703 (68%)
>10	--	157/273 (58%)	--	96/125 (77%)

The huge numbers of real (and fake) low-momentum tracks in minimum bias events make it impractical to use a track P_T threshold of only 1.5 GeV/c for electron identification. More reasonable values will be found in the range 3-10 GeV/c. Since the rate of fake tracks at these higher momentum thresholds also increases with luminosity, the rate of correlation as a function of track P_T must also be considered.

Table 25 shows such a study, where the fraction of EM object-L1CTT track matches is given as a function of L1CTT track P_T for low k_T jet events at a luminosity of $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. These results show that additional rejection is possible by increasing the track P_T and by requiring that the EM object E_T and track P_T match.

Table 25 Trigger tower-track occupancy for a sample of jets with $p_T > 2$ GeV at $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. The rate at which tracks of varying P_T are matched to calorimeter trigger towers of increasing E_T thresholds is shown. The entries in the Table are the same as in Table 24.

EM E_T (GeV)	Track P_T >1.5GeV	Track P_T >3GeV	Track P_T >5GeV	Track P_T >10GeV
>0.5	200k/1520k (13.2%)	70k/1520k (4.6%)	30k/1520k (2%)	10k/1520k (0.7%)
>2	1100/3711 (30%)	600/3711 (16.2%)	211/3711 (6%)	60/3711 (2%)
>5	52/132 (39%)	34/132 (26%)	19/132 (14%)	11/132 (8%)
>10	4/12 (30%)	4/12 (30%)	2/12 (20%)	2/12 (20%)

The above studies clearly demonstrate a potential reduction in the EM trigger rate by exploiting the correlations in ϕ and P_T/E_T between L1CTT tracks and calorimeter objects. As mentioned above, the ϕ granularity of EM objects will improve in Run 2b by a factor of 8 (90° quadrants versus 11.25° towers). The increased rejection of the improved ϕ granularity is estimated in Table 26. The EM object – track match fraction is given for two track P_T thresholds and compares quadrant and trigger tower matching. Low k_T jet events at a luminosity of $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ are used as the background sample. We use this study to estimate the increase in background rejection for EM triggers at high luminosity using the L1CalTrack trigger to be an important factor of 2-3.

Table 26 Trigger-tower-track occupancy for a sample of jets with $p_T > 2$ GeV at $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. The table presents a comparison of the rate at which tracks of $p_T > 1.5$ GeV or $p_T > 10$ GeV are matched to an individual trigger tower or a calorimeter quadrant containing an EM tower above a given threshold. Each line in the table contains the number of matches divided by the total number of quadrants or towers above that E_T threshold.

EM E_T	Track $p_T > 1.5$ GeV (quadrants)	$p_T > 1.5$ GeV (towers)	Track $p_T > 10$ GeV (quadrants)	$p_T > 10$ GeV (towers)
2 GeV	2470/3711	1100/3711	225/3711	60/3711
5 GeV	103/132	52/132	21/132	11/132
10 GeV	8/12	4/12	2/12	2/12

5.2.2 Improving L1CTT Rates Using Calorimeter Jets

The previous section presented evidence that the addition of track information can improve the rejection of an electron trigger by requiring a track

close in ϕ to the EM trigger tower. In this section we explore the equally useful converse, namely that the calorimeter can be used to improve the selectivity and background-rejection of tracking triggers. Isolated high P_T tracks are signatures of many types of interesting events. However, the triggers that select these tracks suffer from a large background of fakes, even for a track $P_T > 10$ GeV/c. As has been indicated elsewhere in this document, this problem worsens substantially as the number of multiple interactions increases. The matching of these tracks to signals in the calorimeter has the ability to confirm the existence of the tracks themselves, and also to verify their momentum measurement.

In this study, our matching algorithm considers individual L1CFT sectors ($\Delta\phi=4.5^\circ$) with at least one track of a given minimum P_T , and matches them in ϕ to whatever trigger towers they overlap. By doing this, we avoid double counting some of the redundant track solutions that cluster near to each other. In about one third of the sectors, these tracks will overlap two different trigger towers in ϕ ; each match is counted separately. The results of track-trigger tower matching are given in Table 27 using the low k_T , high luminosity QCD sample as representative background. Note that for this study, the E_T in the table is the Total E_T (EM+EH), not just the EM E_T . Given that most tracks are hadrons, this is more representative of the true energy that should be matched to a given track.

Table 28 Trigger-tower-track matching for a sample of jets with $k_T > 2$ GeV at 5×10^{32} cm⁻²s⁻¹. The number of CFT trigger sectors containing at least one track above a given p_T threshold is shown, both without and with matching to calorimeter trigger towers of increasing total E_T .

track p_T	# sectors with tracks	Tot $E_T > 1$ GeV	> 2 GeV	> 5 GeV	> 10 GeV
> 1.5 GeV	52991	16252	3218	200	13
> 3 GeV	12818	5188	1529	144	13
> 5 GeV	4705	1562	476	73	9
> 10 GeV	2243	655	141	31	5

With this algorithm we find substantial rejections from even mild trigger tower thresholds. For example, a 10 GeV/c track matching to a 5 GeV trigger tower provides a factor of ~ 70 rejection against fakes. Matching any track to a 2 GeV tower provides approximately a factor of 10 rejection. The rejection shown in this table is essentially sufficient to allow the high- p_T single and di-track triggers to function at the highest luminosities. At this preliminary stage this is a very promising result.

Clearly further simulation work is needed and is in progress. Additionally we must include the CPS and FPS elements, which should provide additional rejection. Finally, simulation results must be cross-checked with results from

collider data. These latter studies will be carried out in the near future when the L1CTT trigger is operational.

5.3 Implementation

Our working assumption is that the L1CalTrack trigger architecture can be made identical to that of the existing and operational Level 1 Muon (L1Muo) trigger with only small and straightforward modifications. Specifically, we will use minimally modified muon trigger (MTCxx) cards with a new flavor board (MTFB) that performs the calorimeter-track match algorithms. However even the new flavor board will be a straightforward upgrade of the existing flavor board that contains the muon detector-track match algorithms. The L1CalTrack trigger crate manager (MTCM) and trigger manager (MTM) will be duplicates of those used for the L1Muo trigger. Most importantly, the engineering effort on traditionally time-consuming details such as synchronization, buffering, messages to L2 and L3, electronics testing, monitoring, power supplies, and crate infrastructure is then virtually nil. A key (and unresolved) question is whether 16 serial link inputs will suffice for the L1CalTrack trigger. Nevertheless, given all the advantages of using L1Muo trigger hardware and the fact that the L1Muo trigger is being successfully operated, we continue on this path at present.

5.3.1 L1Muo System

A brief description of the L1Muo trigger is given here. This is followed by a few technical details on serial links, synchronization, and buffering. A brief description of how the L1CalTrack trigger uses the L1Muo hardware and possible need modifications.

The L1Muo trigger satisfies the following requirements:

- * Delivers an L1Muo Trigger decision to the TF at 3.3 μ s after Bunch Crossing (BC)
- * Transmits an L1Muo decision for every BC not occurring in the Synch Gap
- * Operates with 132 or 396 ns BC times
- * Synchronizes inputs to each Muon Trigger Card (MTCxx)
- * Provides buffering for input and output data pending an L1 decision from the TF
- * Provides 16 buffers for data pending a Level 2 (L2) decision from the TF
- * Provides 8 buffers for data pending readout to Level 3 (L3)
- * Deadtimeless operation
- * Field programmable trigger logic
- * Online and offline monitoring
- * Complete documentation

Thus the L1CalTrack trigger satisfies the same requirements.

A block diagram of the L1Muo Trigger is shown Figure 57. There are three custom VME crates of Muon Trigger Cards (MTCxx's) corresponding to the central (CF), north (EFN), and south (EFS) geographic regions of the DØ detector. There is one custom VME crate that serves as a Muon Trigger Manager (MTM). The VME crates reside on the detector platform and are thus inaccessible during data-taking.

For each crossing, data is transmitted from muon detector front-end cards and the L1CTT trigger to the Muon Trigger Cards (MTCxx's). The information is transmitted at 1060 Mbits/s over coaxial cable using the AMCC S2032/2033 serial link chip set implemented on Serial Link Daughter Boards (SLDB's). Within each geographic region, the MTCxx cards form local trigger decisions for each octant. The actual trigger decision logic is implemented in Altera ACEX series FPGA's contained on a Muon Trigger Flavor Board (MTFB) that plugs into each MTCxx. Currently we have two types of MTFB's called 05 and 10. The first matches tracks from the L1CTT with hits in the muon detector scintillation counters while the second finds tracks using the muon detector wire chambers. A photo of the MTCxx card is shown in

Figure 58.

The octant trigger decisions are sent over a custom VME backplane to the Muon Trigger Crate Manager (MTCM) which subsequently forms trigger decisions for each geographic region. The regional trigger decisions are then transmitted by the MTCM's using Gbit/s serial links to the Muon Trigger Manager (MTM) which forms the global muon trigger decision that is sent to the Trigger Framework (TF). There are 256 L1Muo trigger terms that the user can choose from at begin run time. L1MUO triggers can be chosen based on geographic region (including $1.5 < |\eta| < 2.0$ where there is no L1CFT Trigger coverage), multiplicity (0-3), PT threshold (presently 2, 4, 7, and 11 GeV/c), and quality (called Loose and Tight). Presently there are 32 user defined AND-OR terms sent to the TF representing the global L1Muo trigger decision. A photo of the MTCM card is shown in

Figure 59.

On receipt of an L1 Accept from the TF, the L1Muo Trigger sends its trigger decision and additional information to the L2 trigger system (via the Serial Link Interface Cards (SLIC's)). On receipt of an L2 Accept, the L1Muo trigger sends its trigger decision and additional information to the L3 trigger system (via the Muon Readout Cards (MRC's)). Additionally, the L1Muo trigger may send all of its input data to the L3 trigger system for 1 of N beam crossings (where N is user-defined).

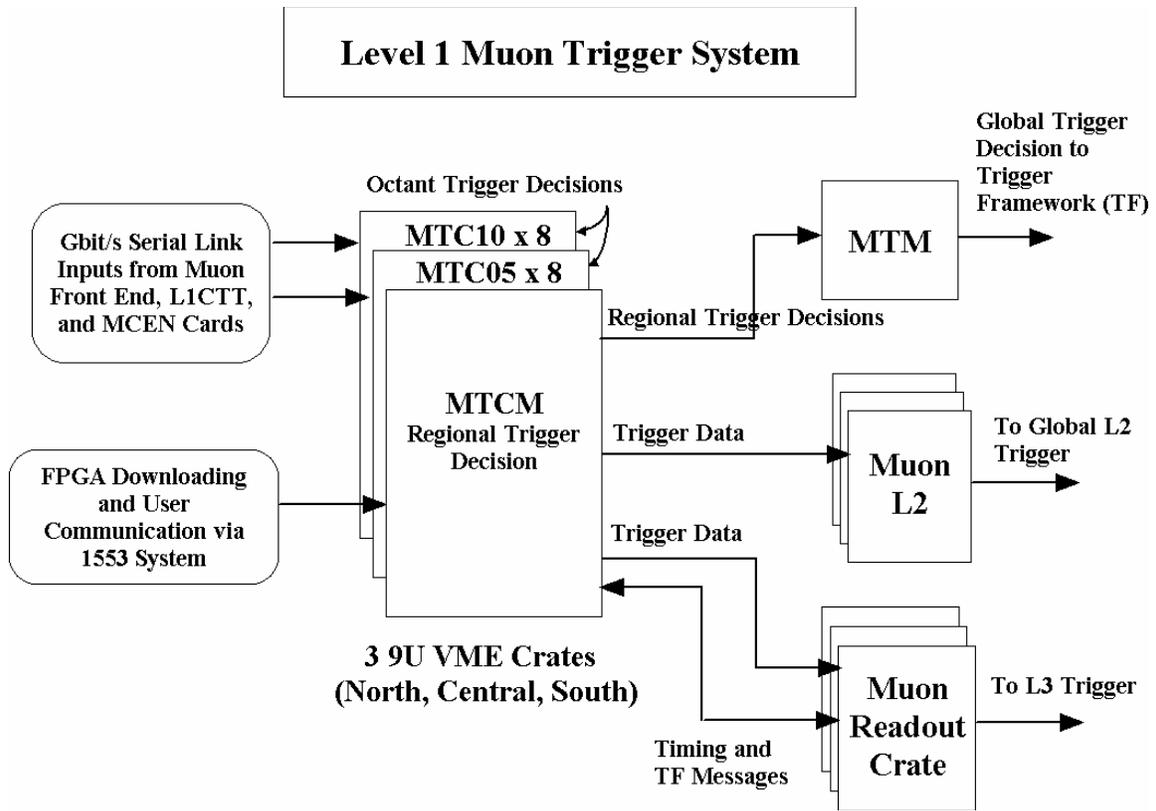


Figure 57. Block diagram of the L1Muo trigger system.

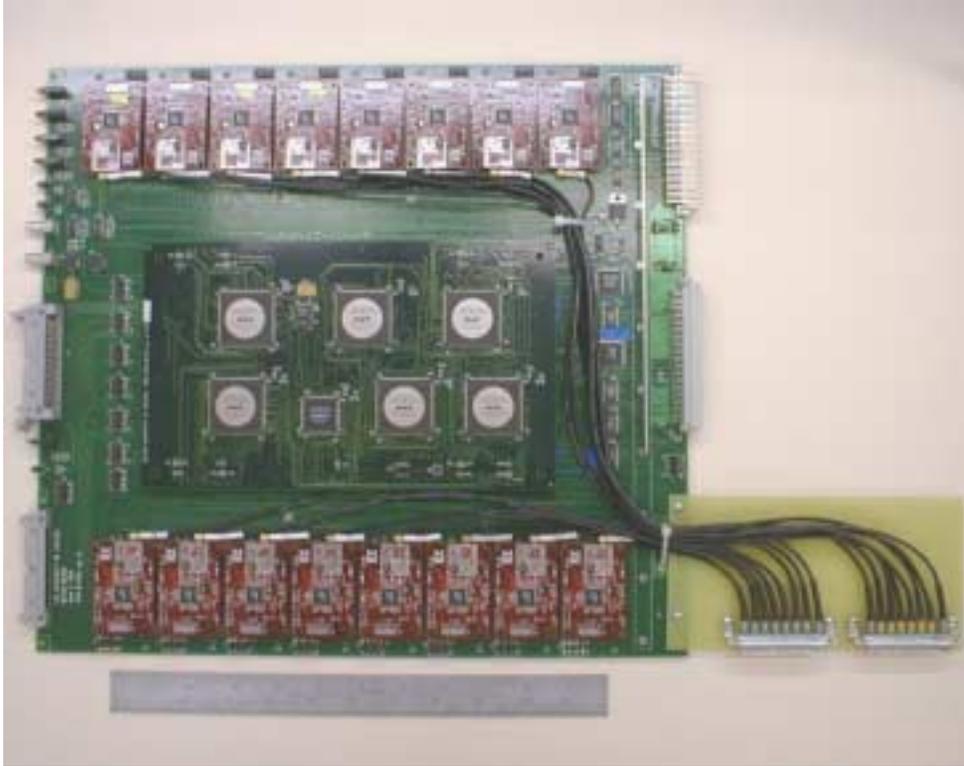


Figure 58. Photo of the MTCxx card for the L1Muo trigger system.

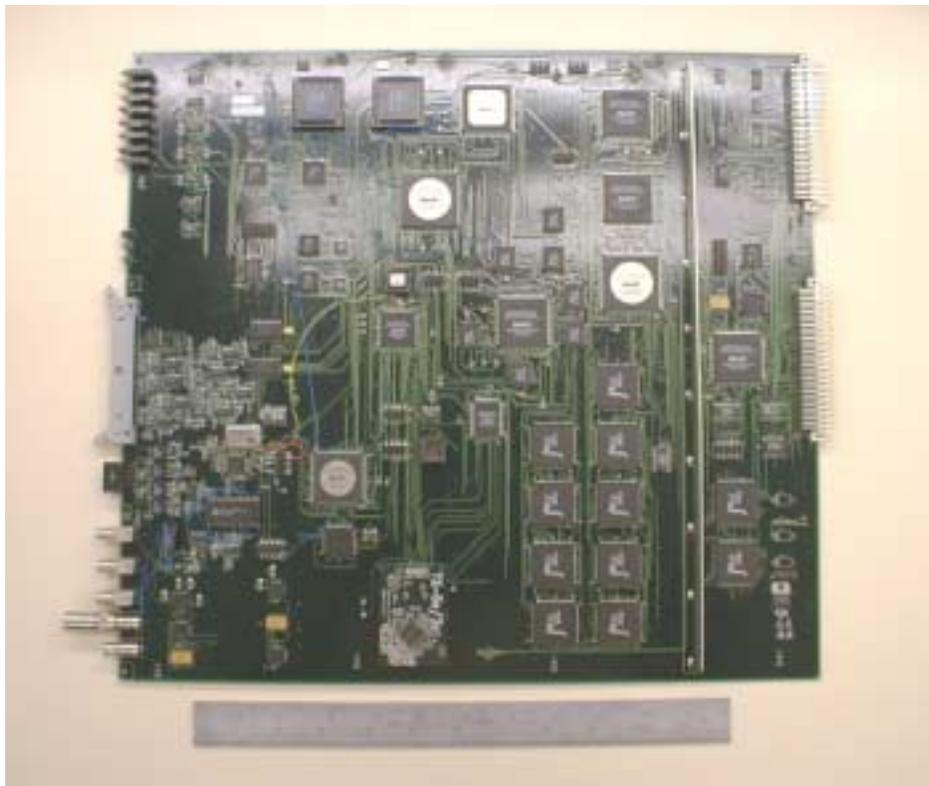


Figure 59. Photo of the MTCM card for the L1Muo trigger system.

Timing and trigger information from the Trigger Framework (TF) takes the following path. The TF sends this information to the Muon Fanout Card (MFC) via a Serial Command Link (SCL). The MFC distributes this information to MRC cards over the VME backplane. An MRC subsequently sends this information to each Muon Trigger Crate Manager (MTCM). The MTCM distributes this information to the MTCxx cards over the VME backplane. Users may communicate with the L1Mu0 Trigger system via two paths: one is the MIL-STD-1553B system and the other a UART between MTCM and MRC.

5.3.2 Some Technical Details

This section is meant to include a few details of interest to wireheads.

5.3.2.1 *Serial Links*

One of the key elements of the L1Mu0 Trigger system are the Serial Link Daughter Boards (SLDB)'s. Gbit/s serial transmission over coaxial cable was chosen to maximize the amount of information that could be brought onto each MTCxx card, to minimize the cable plant, and to minimize the cost. At the time of the MTCxx design and even today, transmission over optical fiber would have been prohibitively expensive given our need for 768 serial links for the L1MU Trigger and 768 serial links for the MCEN system.

The chipset chosen is the AMCC 2042/2043 which is Fiber Channel compatible. The SLDB also contains an Altera 7K series EPLD which handles 8b/10b encoding and parity calculation on the transmitter and 8b/10b decoding and parity checking on the receiver. The MTFB receiver also contains an equalizer circuit and amplifier (HP IVA-05208) needed for error-free transmission over ~150 feet of coaxial cable (LMR-200). Block diagrams of the SLDB transmitter and receiver are shown in Figure 60 and Figure 61. Eye patterns before and after equalization are shown in Figure 62 and Figure 63.

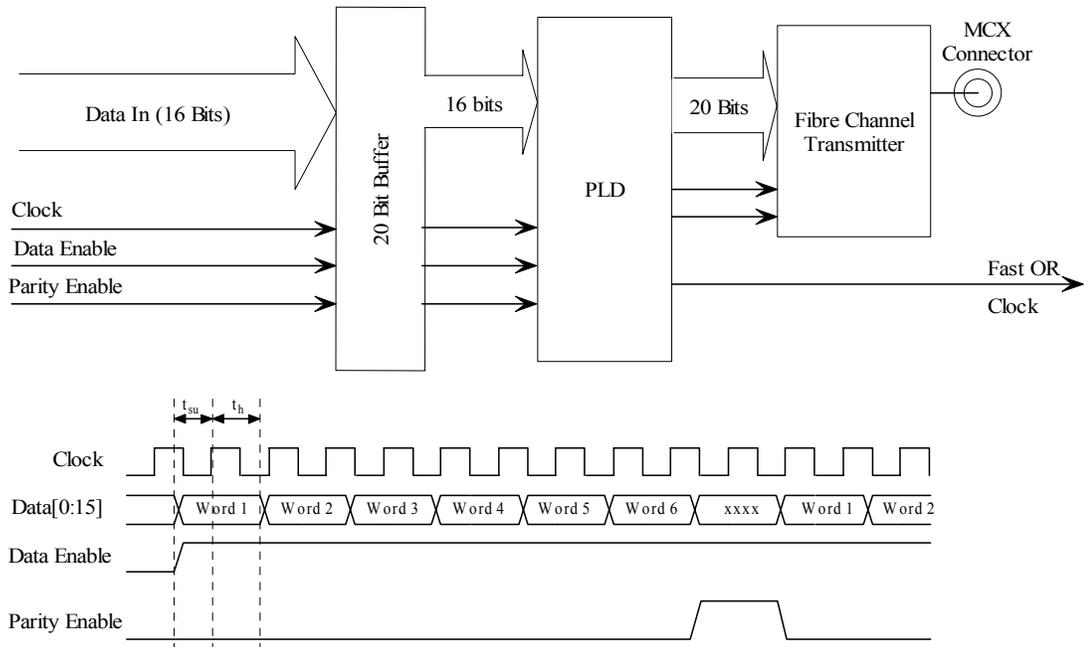


Figure 60. Block diagram of the SLDB transmitter.

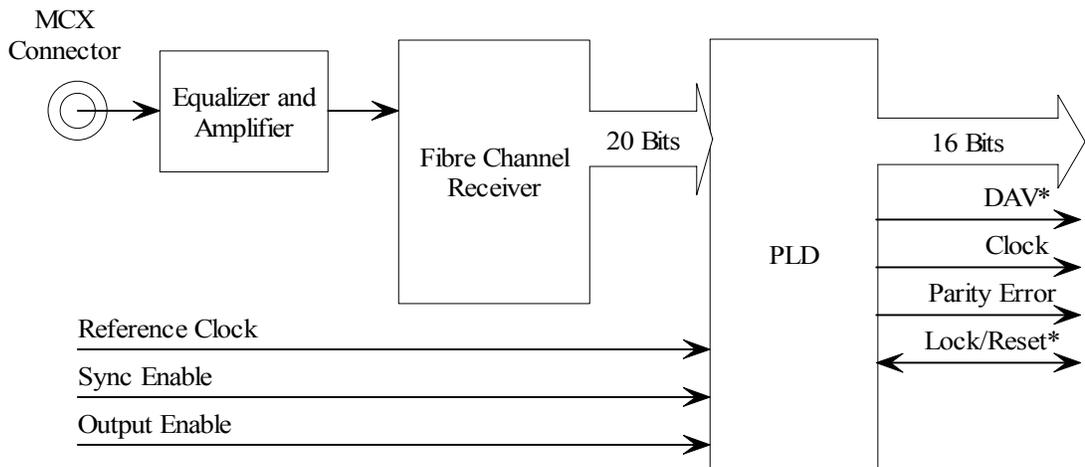


Figure 61. Block diagram of the SLDB receiver.

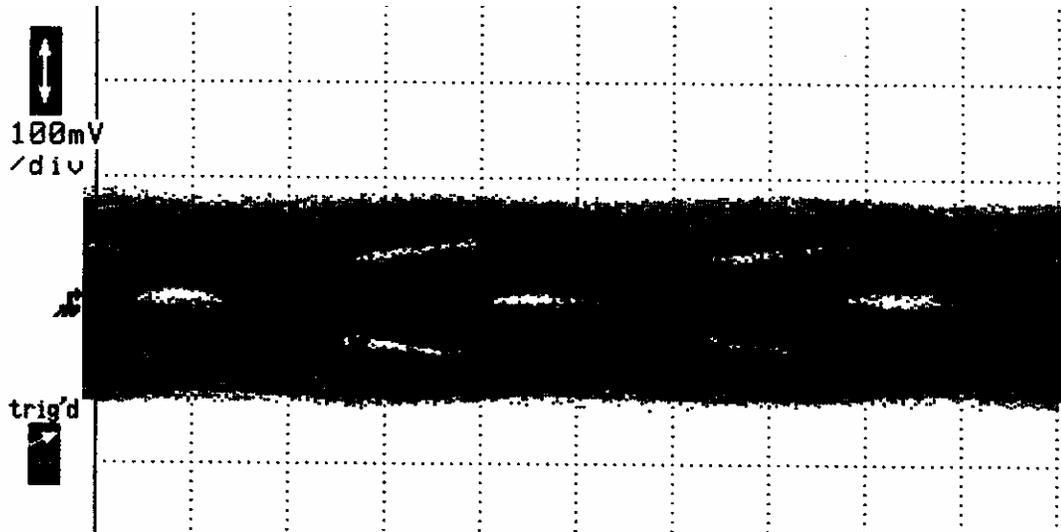


Figure 62. Eye pattern for transmission over 150 feet of LMR-200 coaxial cable without equalization and amplification.

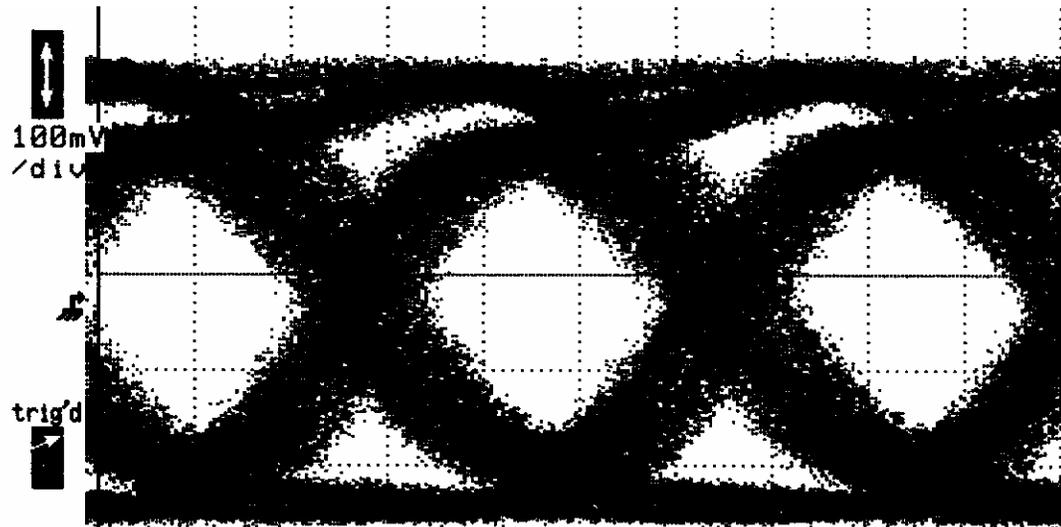


Figure 63. Eye pattern for transmission over 150 feet of LMR-200 coaxial cable with equalization and amplification

The SLDB's have been running in the L1Muo trigger system at DØ for well over a year and have experienced no problems. Additionally the TF uses these SLDB's to transmit SCL information to all Geographic Sectors. Again, their operation has been problem free.

5.3.2.2 Synchronization

A common question is how is the data from different input cables to the MTCxx cards synchronized. The answer is by the use of FIFO's. After INIT, all input FIFO's (Figure 64) are reset. In addition, all active SLDB receivers should be receiving K28.5 Idle characters that are not written to the FIFO's. Thus all FIFO's remain empty.

After INIT, data transmission begins. As data is received at each of the sixteen SLDB receivers, it is immediately written to the input FIFO's. When all FIFO's have data, the MTFB is told to begin processing. Also, an Input Ready signal is sent to the MTCM alerting it that trigger processing has begun. So when all of the Input FIFO's have gone non-empty it is guaranteed that all the input data from the first beam crossing has been synchronized and assembled.

The trigger decision output of the MTFB's are written to FIFO's as well. Each MTCxx sends a Data Ready to the MTCM alerting it that trigger processing is complete. When all MTCxx's have asserted Data Ready, the MTCM returns Start Processing to the MTCxx's which send their MTFB trigger decision data to the MTCM.

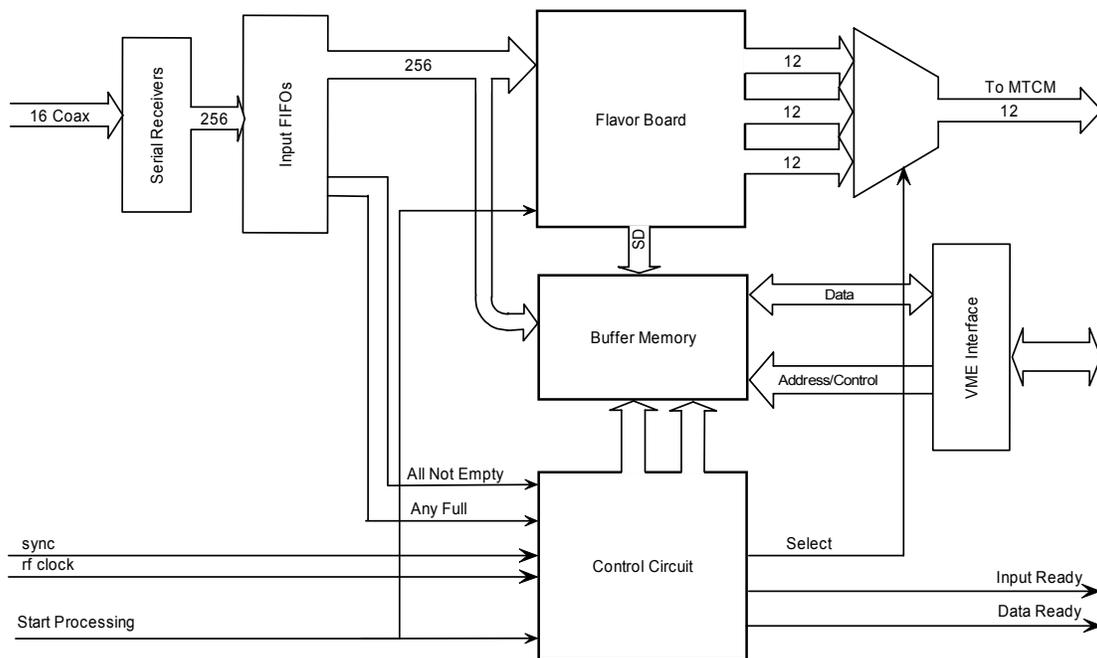


Figure 64. Block diagram of the MTCxx trigger processing.

5.3.2.3 Buffering

Buffering of data pending L1 trigger, L2 trigger, and L2 transfer decisions is achieved through the use of dual port memory (DPM) and a set of FIFO's that keep track of the DPM addresses (see Figure 65). Input data, octant trigger decision data, regional trigger decision data, and other (status, error) information is the data that must be buffered. After Initialization, the appropriate DPM addresses are listed in an Empty Buffer FIFO. With the arrival of the first event, input data, trigger decision data, and other information are written to DPM. The starting address of this DPM data is written to the L1 Pending Buffer FIFO. On receipt of an L1 Accept, the DPM address is written to the L2 Pending Buffer FIFO. A mirror of this FIFO is used by the MTCM to read the DPM in forming the L2 Data message. For L1 Rejects, the DPM address is returned to the Empty Buffer FIFO. On receipt of an L2 Accept, the DPM address is written to the

Buffer Transfer FIFO. This is the address used by the MTCM to read the DPM in forming the L3 Data message. For L2 rejects, the DPM address is returned to the Empty Buffer FIFO.

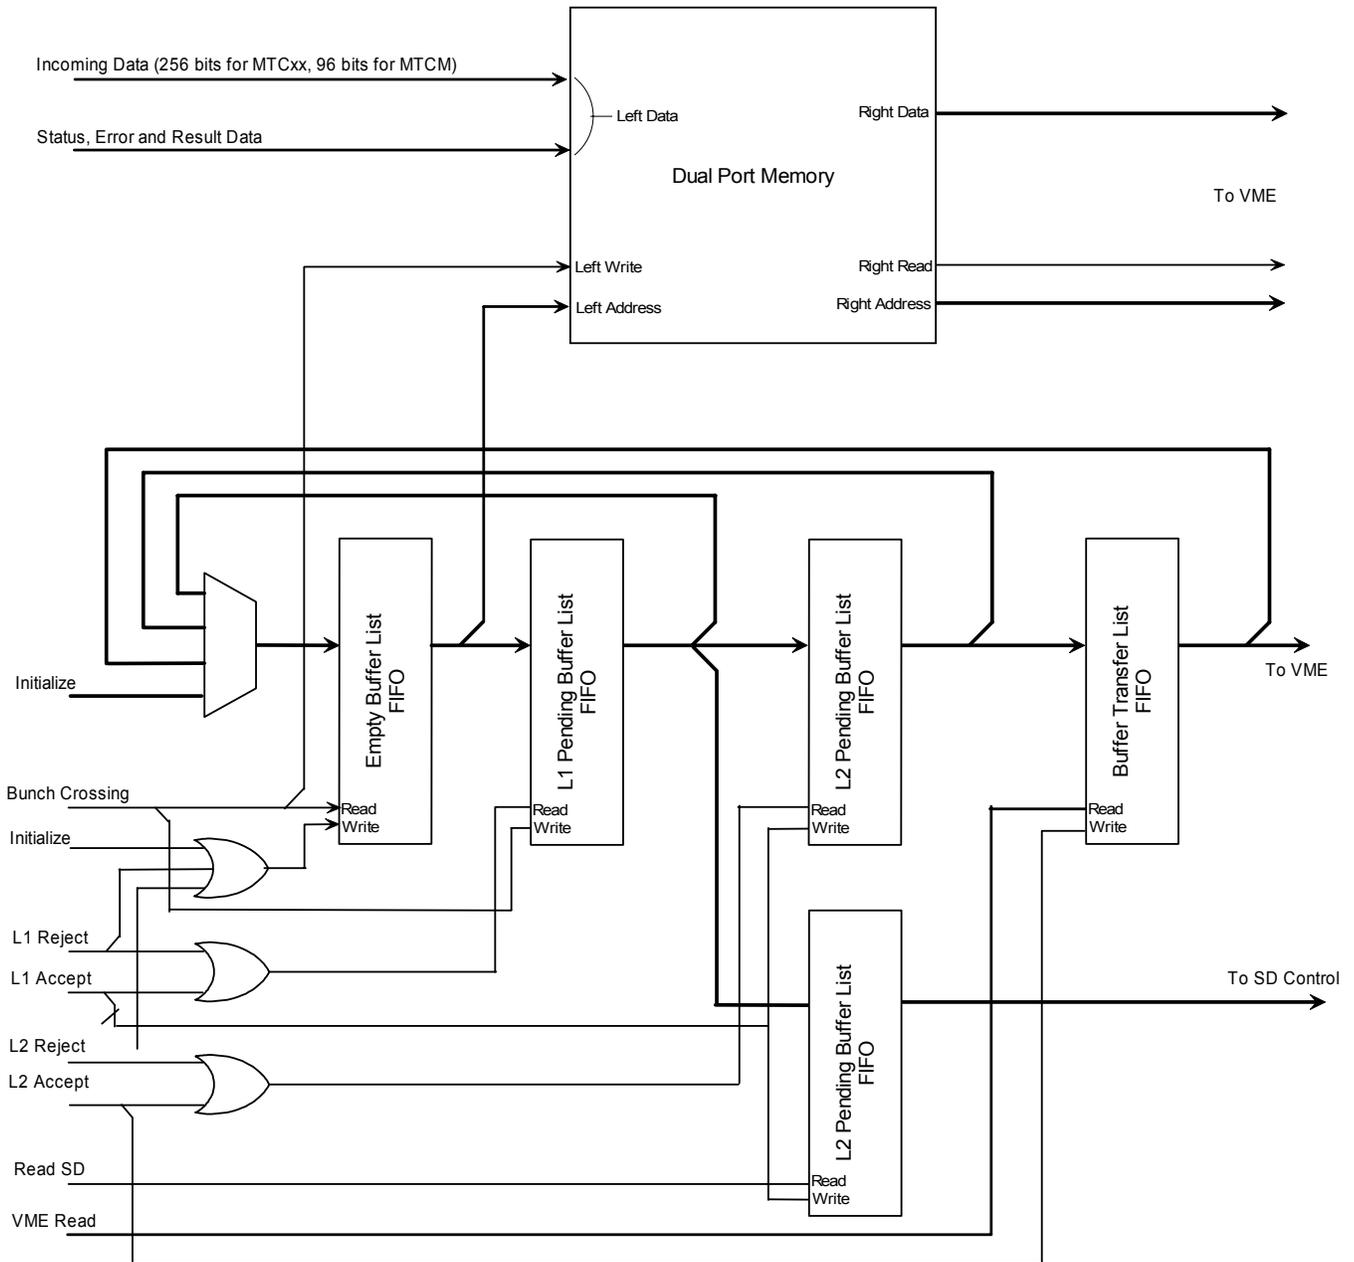


Figure 65. Block diagram of the L1Muo buffering scheme.

5.4 L1CalTrack Trigger

A block diagram of the L1CalTrk trigger is shown in Figure 66. Note the many similarities to Figure 57. Differences between the L1CalTrk and L1Muo

triggers include details and source of the inputs, the MTFB and algorithms that implement the calorimeter-track matching, timing, and the location of the L1CalTrk crates.

Each MTCxx card can accept up to sixteen serial inputs. A preliminary definition of these inputs for the L1CalTrack trigger is given in Table 29. For each crossing, each input effectively contains up to 96 bits (six 16 bit words at 53 MHz) of information. The bit assignment for the L1CTT tracks is defined. The bit assignments for the L1Cal and L1FPS triggers are not yet defined. Note that if additional inputs are needed on the MTCxx card, two to four additional inputs could be accommodated. This would involve changes to the existing MTCxx design and layout however. A requirement of the Run 2b L1CTT trigger is to include the CPS information in the tracks sent to L1Muo. This is because the tracks sent to the L1CalTrack trigger are simply sent on the one of the dual outputs of the L1CTT's SDLB's. Preliminary discussions with the relevant engineers shows including this information to be straightforward.

Table 29. Cable inputs for MTCcal cards.

Cable #	Bits / Cable	Definition	Source
1-10	6 tracks x 16 bits / track	L1CTT trigger tracks (includes CPS)	L1CTT
11-12	96	EM and Jet objects	L1Cal
13-16	96	FPS shower objects	L1FPS

The octant decision formed by MTCcal cards consists of 36 bits that are subsequently sent over the backplane to the MTCM card. The definition of these bits is presently undefined. Note that 36 bits is a hard limit so we must ensure this is sufficient for the variety of triggers produced by the MTCcal cards.

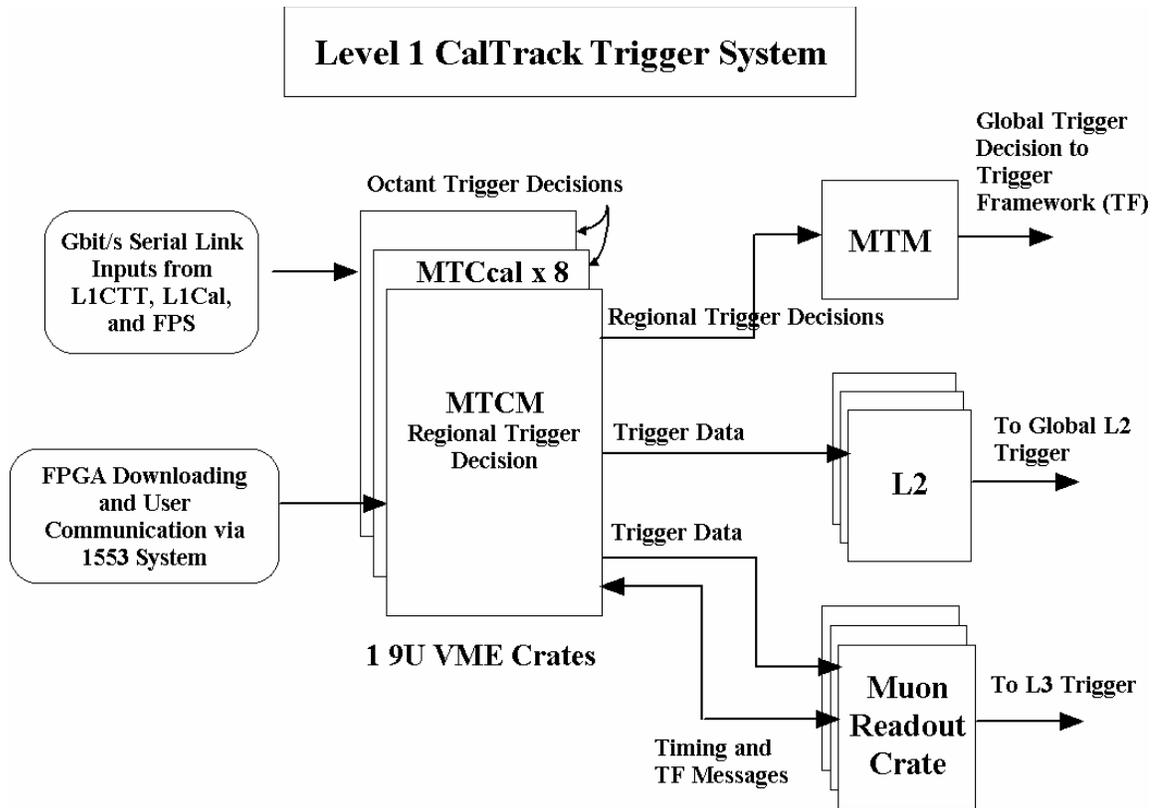


Figure 66. Block diagram of the L1CalTrack system.

Each MTCM receives timing and trigger information from the trigger framework (TF) via a Muon Readout Card (MRC). The detailed path is from TF via a Serial Command Link (SCL) to an Muon Fanout Card (MFC) that distributes the information over a custom VME backplane to the MRC's. The connections between MRC and MTCM are given in Table 30 and Table 31. Because these signals will not be transported over a standard 24-wide Astro cable, a specialized cable between the MTCM and MRC will be needed.

Table 30. Timing and data signal definitions between the MTCM and MRC. The signals are sent over coaxial "Astro" cable.

Pin	Definition	Pin	Definition
1	L3 Data +	2	L3 Data -
3	RF Clock +	4	RF Clock -
5	Encoded Timing +	6	Encoded Timing -
7	L2 Data +	8	L2 Data -

Table 31. Trigger information definitions between the MTCM and MRC. These signals are transmitted over 50c twist and flat cable.

Pin	Definition	Pin	Definition
1	BC Number 1 +	2	BC Number 1 -
3	BC Number 2 +	4	BC Number 2 -
5	BC Number 3 +	6	BC Number 3 -
7	BC Number 4 +	8	BC Number 4 -
9	BC Number 5 +	10	BC Number 5 -
11	BC Number 6 +	12	BC Number 6 -
13	BC Number 7 +	14	BC Number 7 -
15	BC Number 8 +	16	BC Number 8 -
17	INIT +	18	INIT -
19	L1 Accept +	20	L1 Accept -
21	L2 Error +	22	L2 Error -
23	L2 Accept +	24	L2 Accept -
25	L2 Reject +	26	L2 Reject -
27	UART Transmit +	28	UART Transmit -
29	Buffer Available +	30	Buffer Available -
31	Strobe +	32	Strobe -
33	UART Receive +	34	UART Receive -
35	L1 Error +	36	L1 Error -
37	L1 Busy +	38	L1 Busy -
39	L2 Busy +	40	L2 Busy -
41	GND	42	GND
43	GND	44	GND
45	GND	46	GND
47	GND	48	GND
49	GND	50	GND

Each MTCM is connected to the Muon Trigger Manager (MTM) via a serial link. As an aside, the MTM is made of an MTCxx card with an MTM MTFB. The format of this (presently) bitwise transfer is given in Table 32.

Table 32. MTCM serial link data sent to the MTM.

Word #	Definition
0	Spare
1	MTCM trigger decision bits 0-11
2	MTCM trigger decision bits 12-23
3	MTCM trigger decision bits 24-35
4	Spare
5	Parity

On receipt of an L1 Accept, the MTCM transmits data found in Table 33 to the L2 Muon Trigger. The data block size is the same for all events. Presently sixteen-bit words are sent with the lower byte first. Finally note there is a similar but not identical block of data from the MTM crate.

Table 33 MTCM output to L2.

Byte #	Definition
	BeginTransmission (K28.0)
1-2	Word Count (=35)
3-4	Module ID
5-6	Local Crossing Number
7-8	Local Turn Number
9-10	Event Status (undefined)
11-12	Event Status (undefined)
13-14	MTCM Status/Control Register
15-16	MTCM Event Error Register
17-22	MTCM Regional Trigger Decision
23-28	MTCcal Octant 0 Trigger Decision
29-34	MTCcal Octant 1 Trigger Decision
35-40	MTCcal Octant 2 Trigger Decision
41-46	MTCcal Octant 3 Trigger Decision
47-52	MTCcal Octant 4 Trigger Decision
53-58	MTCcal Octant 5 Trigger Decision
59-64	MTCcal Octant 6 Trigger Decision
65-70	MTCcal Octant 7 Trigger Decision
	End Transmission (K23.7)

On receipt of an L2 Accept, the MTCM transmits data found in Table 34. to the MRC where it is subsequently read by the VBD. There are two possible data block sizes. Recall the L1Mu0 trigger system has the option of transmitting all of its input data for 1 of N events (where N is user defined). The two block sizes correspond to whether or not the input data block is included. Presently sixteen-bit words are sent with the lower byte first. Finally note there is a similar but not identical block of data from the MTM crate.

Table 34. MTCM output sent to L3.

Word #	Definition
	BeginTransmission (K28.0)
	Upper 16 bits of VBD Word Count = 0
	VBD Word Count = 246 or 1142 decimal
1	Word Count = 492 or 2284 decimal
2	Module ID
3	Local Crossing Number
4	Local Turn Number
5	Event Status = MTCM Event Error Register from DPM
6	Event Status (undefined)
7	MTCM Control Register
8	MTCM Event Error Register from DPM
9	MTCM Latched Error Register
10	MTCM Trigger Logic FPGA Program ID
11	MTCM Message Builder FPGA Program ID
12	MTCM L1_Accept Divider Number
13	MTCM L2_Accept Divider Number
14	MTCM MTCxx Readout Mask
15	MTCM MTCxx Trigger Mask
16	MTCM L1 Error Mask
17-19	MTCM Regional Trigger Decision
20	MTCxx #0 MTCxx Serial Number
21	MTCxx #0 Flavor Board Type/Serial Number
22	MTCxx #0 MTCxx Status Register
23	MTCxx #0 MTCxx Flash Memory Status
24-34	MTCxx #0 Various Error Registers (82,86,8a,8e,92,94, pad rest with zeros)
35-38	MTCxx #0 Various Mask Registers (06,08,0a,0c)
39-46	MTCxx #0 FPGA Program ID's (total of 8)
47-73	MTCxx #1 Info Block
74-100	MTCxx #2 Info Block
101-127	MTCxx #3 Info Block
128-154	MTCxx #4 Info Block
155-181	MTCxx #5 Info Block
182-208	MTCxx #6 Info Block

209-235	MTCxx #7 Info Block
	MTCcal Octant 0 Trigger Decision
	MTCcal Octant 1 Trigger Decision
	MTCcal Octant 2 Trigger Decision
	MTCcal Octant 3 Trigger Decision
	MTCcal Octant 4 Trigger Decision
	MTCcal Octant 5 Trigger Decision
	MTCcal Octant 6 Trigger Decision
	MTCcal Octant 7 Trigger Decision
	MTCxx Input Data
	End Transmission (K23.7)

Finally, the MTM MTFB (which is used in concert with an MTCxx mothercard) takes the regional trigger decision data from the three MTCM cards and forms a variety of global trigger decisions. The specific triggers sent to the Trigger Framework are downloaded during the physics trigger download of a data acquisition run. The data sent to the Trigger Framework from the MTM MTFB is given in Table 35.

Table 35. MTM data sent to the Trigger Framework.

Pin	Definition	Pin	Definition
1	Specific Trigger 0 +	2	Specific Trigger 0 -
3	Specific Trigger 1 +	4	Specific Trigger 1 -
5	Specific Trigger 2 +	6	Specific Trigger 2 -
7	Specific Trigger 3 +	8	Specific Trigger 3 -
9	Specific Trigger 4 +	10	Specific Trigger 4 -
11	Specific Trigger 5 +	12	Specific Trigger 5 -
13	Specific Trigger 6 +	14	Specific Trigger 6 -
15	Specific Trigger 7 +	16	Specific Trigger 7 -
17	Specific Trigger 8 +	18	Specific Trigger 8 -
19	Specific Trigger 9 +	20	Specific Trigger 9 -
21	Specific Trigger 10 +	22	Specific Trigger 10 -
23	Specific Trigger 11 +	24	Specific Trigger 11 -
25	Specific Trigger 12 +	26	Specific Trigger 12 -
27	Specific Trigger 13 +	28	Specific Trigger 13 -
29	Specific Trigger 14 +	30	Specific Trigger 14 -
31	Specific Trigger 15+	32	Specific Trigger 15 -
33	Gap +	34	Gap -
35	Ground	36	Ground
37	Strobe +	38	Strobe -
39	Ground	40	Ground

5.5 L1CalTrack Cards

In this section we briefly summarize the functions of VME cards and daughter boards comprising the L1CalTrk Trigger.

5.5.1 Serial Link Daughter Board (SLDB)

The Serial Link Daughter Boards (SLDB's) are used to transmit data over coaxial cable between the L1CTT, L1Cal, and L1FPS triggers (transmitters) and the MTCxx cards (receivers) in the L1CalTrack trigger. There are typically one or two SLDB's on the L1CTT, L1Cal, and L1FPS triggers and sixteen SLDB's on the MTCxx cards. The serial links used are the Fiber Channel compatible AMCC S2042/S2043 chipsets. Altera 7K series EPLD's are used for 8b-10b encoding/decoding and parity checking. Seven sixteen-bit words are transmitted at 53 MHz. With 8b-10b encoding this gives a serial transfer rate of 1060 Mbits/s. Error free data transmission over 150 foot lengths of coaxial cable

(LMR-200) is achieved by using an equalization circuit and high-speed amplifier (HP IVA-05208) on the SLDB receiver boards.

5.5.2 Muon Trigger Card (MTCxx)

The Muon Trigger Cards (MTCxx cards) are the primary trigger cards in the L1CalTrack system. The MTCxx card is a generic VME card that accepts a Muon Trigger Flavor Board (MTFB) as a daughter card. The actual calorimeter-track match trigger logic is implemented on the MTFB. The MTCxx/MTFB combination is used form an octant trigger decision. The MTCxx card also contains sixteen SLDB receivers that are used to accept data from the various input sources. The primary functions of the MTCxx cards are to receive and synchronize sixteen serial inputs, to buffer input data (which is subsequently sent to the MTFB) and to buffer input and supplemental trigger decision data pending L1, L2 and L3 accepts.

5.5.3 Muon Trigger Flavor Board (MTFB)

The Muon Trigger Flavor Board (MTFB) is a daughterboard that is used in concert with the MTCxx card. For the L1CalTrack trigger a new flavor board will be used called MTCcal. This flavor board will contain the calorimeter-track match trigger logic. Note the present MTC05 MTFB matches tracks from the L1CTT with hits in the muon detector scintillation counters. The MTCcal MTFB will match tracks from the L1CTT with EM and jet objects from L1Cal. While we must simulate all calorimeter-track match algorithms in both the C++ trigger simulator and MAXPLUS2 FPGA simulator before a final choice of FPGA's can be made we feel the size and number required will be similar to that on the MTC05 MTFB's.

5.5.4 Muon Trigger Crate Manager (MTCM)

The Muon Trigger Crate Manager (MTCM) reads the octant trigger decisions from each MTCxx card and uses this data to form a regional trigger decision that is subsequently sent to the Muon Trigger Manager (MTM). The MTCM also serves to buffer the octant trigger decision data from each MTCxx card and the regional trigger decision pending L2 and L3 accepts. L2 data and L3 data are sent to their respective systems using the Cypress Hotlink (CY7B923/CY7B933) chipset. The MTCM accepts timing and trigger information from the Trigger Framework (TF) and reports error and busy conditions to the TF also via the MRC. The MTCM maintains two paths by which a user can communicate to the L1MU system: 1553 and UART. The UART path is a link between the MTCM and MRC. Either path can be used for downloading, monitoring, and testing.

5.5.5 Muon Splitter Cards (MSPLIT)

The MSPLIT cards are used as fanouts for many of the Gbit/s serial input signals to the MTCxx cards. They are mentioned here because the location of the L1CalTrack crates is now in MCH1 rather than collision hall. This means we must send the SLDB Gbit/s data over 180-200 feet of coaxial cable. This is at the upper limit of our demonstrated length of guaranteed error free transmission. Should an additional signal boost be required, we can use the MSPLIT cards just

inside of MCH1. The splitters are implemented on 9U VME cards but without any VME functionality. They consist of a high-speed amplifier (HP IVA-05208) and high bandwidth transformer splitter. Each card has 8 inputs and up to 24 outputs. Hence they can be used as cheap repeaters.

5.5.6 Muon Trigger Test Card (MTT)

The purpose of the Muon Trigger Test (MTT) card is to test and debug the L1CalTrack trigger cards. These include the MTCxx and MTCM cards. The MTT simulates front-end data for the MTCxx cards using sixteen SLDB transmitters with user-defined input. The MTT serves as an MRC in that it generates trigger and timing data used by the MTCM cards. The MTT also contains three Hotlink receivers so that the L2 and L3 data outputs of the MTCM can be tested. It is mentioned as an another example of how engineering is greatly minimized by using existing hardware and/or designs.

5.5.7 Timing

A detailed spreadsheet of the total latency of the L1CalTrack trigger is not yet available. However we can make an estimate based on measured numbers for the L1Muo trigger and the Run 2a calorimeter trigger plus estimates for the remaining elements. This estimate is given in Table 36.

Table 36. Trigger latency of the L1CalTrack trigger.

Element	Time (ns)
BC to ADF card (measured)	650
ADF processing and serialization (estimate)	400
L1Cal to SLDB transmit enable (estimate)	700
SLDB transmit enable to MTCcal data ready (measured assuming MTC05 SLDB logic)	861
MTCcal data ready to MTCM SLDB transmit enable (measured)	544
MTCM SLDC transmit enable to MTM decision at TF (measured)	406
Total	3561
TF L1 Decision Time	3300
Difference	+261

The first number is taken from the measured time from BC (Bunch Crossing) to the time the calorimeter signals presently reach the Run 2a calorimeter trigger cards (Figure 13). The last three numbers are taken from measurements in the L1Muo trigger system. They basically include MTCcal processing, MTCM processing, and MTM processing including all synchronization, serialization, and cable delay times. Here we assumed 528ns for the MTCcal logic believing it to be similar to the L1Muo MTC05 logic (both process L1CTT tracks). The second and third numbers are estimates of the ADF and L1Cal processing times that

also must include synchronization, serialization, and cable delay times. Any L1 trigger decision must arrive at the Trigger Framework $3.3\mu\text{s}$ after Bunch Crossing thus the L1CalTrack trigger is estimated to be 261ns over this time budget. Work to understand the ADF and L1Cal latencies is in progress.

If, once better numbers for the ADF and L1Cal latencies are known, the L1CalTrack trigger remains over time budget, an alternative solution must obviously be found. One such solution is as follows. Because there is only one L1CalTrack crate of MTCCal trigger cards, the need for the MTCM crate manager to be involved in the trigger decision is diminished. One can envision sending octant trigger decisions directly from each MTCxx trigger card to the MTM trigger manager card. This would save several hundred ns and result in no loss of functionality. Each MTCcal flavor boards would have to contain a serial link transmitter either in daughter board or FPGA form but this does not present a problem given that the MTCcal flavor boards are new in any case.

5.6 Summary

In this section we give the cost estimate, schedule, and milestones. We also summarize the outstanding issues in the L1CalTrack trigger system as follows: continued simulation of the L1CalTrack system especially including the CPS and FPS information and the use of jets to match tracks, detailed specification of the inputs, and detailed timing estimate. Work on the latter two issues is in progress. Work on the first issue is getting underway.

Assuming the L1CalTrack trigger is based on the L1Muo trigger, the cost estimate is given in Table 37. Most of the costs are taken from the actual costs of the L1Muo trigger cards. The contingency is large because of overall L1Cal system uncertainties and in case additional inputs need to be added to the MTCxx card. While the modifications to the MTCxx card would be straightforward, a prototype version would be prudent. The numbers include two MTCxx trigger card spares, two MTCcal flavor board spares, and one MTCM crate manager spare.

Table 37. L1CalTrack cost estimate.

Item or process	Unit cost (\$)	Number required	Total Cost (\$k)	Total Cost + contingency (\$k)
MTCxx Trigger Cards	2300	11	25.3	32.9
SLDB Rransmitter Boards	115	128	14.7	19.1
SLDB Receiver Boards	100	128	12.8	16.6
MTCcal Flavor Boards	1200	11	13.2	19.8
MTCM Crate Manager	3700	3	11.1	14.4
Splitter (Repeater)	705	3	2.1	2.8
LMR-200	0.5	25600	12.8	19.2
Connectors	10	256	2.6	3.8
LMR-100	0.3	256	0.1	0.1
Connectors	8	256	2.0	3.1
VME crates	4000	2	8.0	12.0
Processors	2200	2	4.4	6.6
Power Supplies	2500	2	5.0	7.5
Power Supply Cases and Electronics	1200	2	2.4	3.6
Prototype MTFB	1500	2	3.0	4.5
Prototype MTCxx	3500	1	3.5	5.2
AZ technician	15000	1	15	15
AZ engineering	36400	1	36.4	36.4
TOTAL			176	225

In Table 38, we give a schedule for the L1CalTrack trigger. Assuming the L1CalTrack trigger is based on the design of the L1Muo trigger system, most of the engineering is already completed. A new MTFB must be designed but this will be based on existing MTFB's. Modifications to the MTCxx card may be necessary if more than 16 serial link inputs are required.

Table 38. Schedule for the L1CalTrack trigger project.

Item	Date Completed
Simulation and algorithm development	10/02
MTCM fabrication and assembly	05/03
MTCxx design modifications	1/03
MTFB design (includes algorithm simulation)	08/02
MTCxx fabrication and assembly	06/03
MTFB fabrication and assembly	03/03
Test all cards with MTT	08/03
Procure and assemble cables	03/03
Procure and assemble power supplies	08/03

The schedule contains a large amount of time for simulation in order to understand in detail all the possibilities for the L1CalTrack trigger. Typical fabrication times have been approximately 3-4 weeks in our experience. Typical assembly times were quoted as 3-4 weeks but in reality were 2-3 times this. Assuming no major modifications are needed on the MTCxx cards, we believe we can have the L1CalTrack trigger at DØ by fall of 2003. The schedule above is very comfortable.

Suggested milestones are listed in Table 39.

Table 39. Milestones for the L1CalTrack project.

Milestone	Date
Detailed design review at DØ	10/02
MTCxx submitted for fabrication	03/03
MTFB submitted for fabrication	01/03
L1CalTrack trigger at DØ	08/03

6 Level 2 β Trigger

6.1 Motivation

An overview of the Level 1 and Level 2 trigger system is given in Section 2.1. At Level 2, preprocessors analyze output from the Level 1 trigger of each detector system in parallel. (The L2 Muon and L2 Silicon Tracking Trigger (L2STT) preprocessors also receive fully digitized detector data.) Event selection takes place based on the high level (physics object) information available in a Global stage after detector preprocessing is complete. The preprocessors instrumented in the Run2b trigger include calorimeter, preshower, silicon tracker, and muon components.

The input rate to the L2 trigger is limited by the SMT digitization deadline, and the output rate is limited by the calorimeter precision readout deadline. Since both limits are constant from Run 2a to Run 2b, the primary charge for Level 2 will be to maintain Run2a rejection levels (factor of ~ 5) within the same time budget (to fully realized the advantages from our L1 enhancements). Maintaining Level 2 rejection in the Run2b trigger will be more challenging as some algorithms used in the Run 2a Level 2 trigger move upstream to Level 1 for Run 2b. To accomplish its goal, Level 2 must make better use of the time budget by using more powerful processors. This project is already under way with the construction of the 'Level 2 β ' processors, initially conceived to deal with lower than expected production yields in the Run 2a Alpha processors and to offer a clear upgrade path for increases in future performance.

6.2 L2 β Architecture

All L2 processors occupy 9U VME64 for physics crates. These crates provide dual backplanes: a standard VME bus, and a custom-built 128-bit "Magic Bus" or MBus (a handshaking bus capable of data transfer rates up to 320 MB/s). Each crate contains a number of devices for communication with the experiment's front end and trigger systems and at least two processor cards for analysis of detector subsystem data. The processors are configured for Administrator or Worker functions. Where appropriate, additional specialized hardware for data conversion or processing can be added. A Worker node applies trigger algorithms to its input data. The Administrator does all event processing and local trigger control tasks that do not involve the application of the trigger algorithm. These include verifying data integrity, controlling communication with the trigger framework, controlling the output of monitoring data, and controlling the readout of events to the higher trigger levels.

The L2 β processors⁸ rely on commercially produced single board computers (SBCs). Each SBC resides on a 6U CompactPCI (cPCI) card providing access to a 64-bit, 33/66 MHz PCI bus via its rear edge connectors. Such cards are

⁸ A detailed L2 β TDR is available at <http://galileo.phys.virginia.edu/~rjh2j/l2beta/>

currently available “off the shelf” from several vendors including Advantech-nc⁹, VMIC¹⁰, Diversified Technology Inc.¹¹, and Teknor¹². The remaining functionality of the board is implemented in a large FPGA and Universe II¹³ VME interface mounted on a 6U-to-9U VME adapter card as shown in Figure 67 - Figure 68.

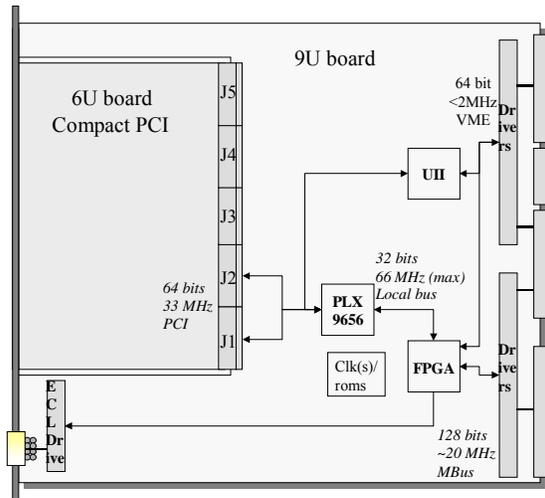


Figure 67. Model for the L2β processor card. Connectors J1 and J2 provide a 64-bit cPCI connection to the CPU.

Figure 68. Level 2 beta prototype. The SBC and 9U cards are shown with the mechanical assembly. The 9U card is shown without components installed.

The adapter card contains all DØ-specific hardware for Magic Bus and trigger framework connections. Custom I/O functions on this card will be implemented in a single FPGA (Xilinx XCV405E) plus assorted logic converters and drivers. This device is particularly suited to our application, because of its large amount of available Block RAM. 70KB of RAM (in addition to >10K logic cells) is used to implement internal data FIFOs and address translation tables for broadcasting data from the Magic bus to CPU memory, reducing the complexity of the 9U PCB. A hardware 64-bit, 33MHz PCI interface to the SBC is implemented with a PLX 9656 PCI Master chip. The SBC, in the adapter, has its front panel at the face of the crate and is easily removable for upgrade or repair. The modular design provides a clear path for CPU performance upgrades by simple swapping of SBC cards.

⁹ <http://www.Advantech-nc.com>.

¹⁰ <http://www.vmic.com>

¹¹ <http://www.dtims.com>.

¹² <http://www.teknor.com>

¹³ Tundra Semiconductor Corp., <http://www.tundra.com>.

Given comparable I/O capabilities, the amount of time required to run complex algorithms should be inversely proportional to the processor speed; more complicated algorithms can be used to process the available data if the processors are faster. However, an increase of processing power is more useful when supplied in the form of a single processor than in the form of a second identical processor working in parallel. This is because load balancing among multiple nodes is difficult in the Level 2 system due to constraints imposed by the front-end digitization. The front-end digitization holds buffers for 16 events awaiting Level 2 trigger decisions. A critical restriction in the system is that L2 results (accept or reject) must be reported to the front-end buffers in the order in which the triggers were taken at L1. While one processor works on an event with a long processing time, other events will arrive and fill the 16 front-end buffers. Other processors working on these events will go idle if they finish processing them quickly, since they cannot receive new events until the pending decision on the oldest event is taken. In other words a farm model is not appropriate for processing events at Level 2. Faster processing for each event in turn is thus more desirable than adding additional β processors, once a baseline level of parallelism is established.

The quality of the Run 2b physics program will depend in large measure on effective rejection of background events in this more demanding environment. The Level 2 β upgrade will provide more resources needed to keep Level 2 in step with these demands and to further improve on background rejection from an upgraded Level 1. A subset of the most heavily-loaded processors should be replaced with higher-performance processors. Assuming that processors in the format used by the L2 β s increase performance by Moore's law, a purchase near the start of Run2b could gain another factor of 4 in processing power over the first L2 β processors.

6.3 Cost & schedule

For Run 2b, we are proposing a partial upgrade of the Level 2 β system by allocating sufficient funds to replace the processors on 12 boards. This is in anticipation of the potential increase in computing power that could at that time be used to implement more sophisticated tracking, STT, and calorimeter/track matching algorithms at Level 2 in response to the increased luminosity.

The cost for upgrading 12 processors is estimated to be \$98K, including \$10K for possible firmware modifications and 36% contingency. Although we have firm quotes for our first SBC purchases, market surveys show approximately 30% spread in SBC prices among manufacturers and, given the speed of evolution in the computer market, we cannot anticipate which models will be preferable several years from now.

The 12 new SBC cards will be used primarily to upgrade worker processor cards and will be distributed as follows: calorimeter (2-3), global (1-2), tracker (1-2), muon (2), Preshower (2), spare/test stand (2-3). An overview of the schedule for the L2 β upgrade is shown in Table 40.

Table 40: L2 β eta upgrade schedule.

Prototype Development	Jun 2004
Prototype Testing	Sep 2004
Assemble Upgrade Processors	Dec 2004
Installation	Jan 2005

6.4 Performance Options

We have begun to consider other algorithms that might profit from additional CPU power. Multi-track displaced vertices could be searched for with the tracks output by the Level 2 Silicon Track Trigger (L2STT). This is beyond the original projected work of the Level 2 Central Tracking Trigger (L2CTT) preprocessor, and would be more CPU intensive. On another front, a sophisticated neural-net filter may search for tau events in the L2 Global processor. The effectiveness of such improvements depends on the actual mix of triggers chosen for Run 2b physics, so these should only be considered as examples. We have not yet studied which algorithms can be imported from Level 3 and applied to the lower-precision data available in Level 2.

An obvious use of additional CPU power would be in the global processor, which does the work of final Level 2 trigger selection by combining the results of preprocessors across detectors. More powerful CPUs will allow us to break the present software restriction of one to one mapping of Level 1 and Level 2 trigger bits (128 each at this point). This would allow more specific trigger processing to be applied to individual L1 trigger conditions at Level 2, as we currently do in Level 3. In addition to channels with inherent physics interest, many signals will play increasingly important roles in the calibration of the detector and efficiency measures for the main physics menu's selection criteria. Added trigger branching will greatly facilitate the collection of these data. It is at times impossible to simulate a dataset with the necessary accuracy to calculate efficiencies and acceptances for complex trigger conditions, especially when hardware calibration effects have exceedingly strong bearing.

In addition to the primary upgrade path of adding higher power CPU cards, a further upgrade avenue may include equipping the cards with dual processors that share the card's memory and I/O. This upgrade is attractive because its incremental cost is low, but it will require a substantial software effort to turn it into increased throughput, even if it is possible to build code that takes advantage of the dual processors without writing thread-safe code. However, a dual-processor upgrade might be attractive for reasons other than performance. One processor could keep the Linux operating system active for debugging of problems in algorithms run in the second processor. Or one could run a production algorithm in one processor and a developmental version in the second processor. This second processor might even be operated in a "shadow" mode (as in Level 3), processing events parasitically, but skipping events if the developmental algorithm gets behind, or is being debugged. These possibilities

will be studied prior to Run2b, though dual CPU cards are not intended as a substitute for higher power upgrade processors.

7 Level 2 Silicon Track Trigger

7.1 Goals

An overview of the Level 1 and Level 2 trigger system is given in Section 2.1. In Level 2, preprocessors analyze output from the Level 1 trigger of each detector system in parallel. The L2 Silicon Track Trigger (L2STT) preprocessor also receives fully digitized detector data. Event selection takes place based on the high level (physics object) information available in the Global stage. The preprocessors developed for the start of Run 2 are for the calorimeter, preshower system, tracking system, and muon system; the L2STT will be added during Run 2A, allowing triggering on tracks with high impact parameters.

The input rate to L2 is limited by the time required to digitize and read out the SMT data, and the output rate is limited by the calorimeter precision readout deadtime. The primary charge for Level 2 will be to maintain the current rejection, with the same time budget. This means the level 2 algorithms for Run 2B have to be refined relative to Run 2A because some of the algorithms that helped reject events in the Run 2A Level 2 trigger will be moved into Level 1 for Run 2b. Upgrading the L2STT to optimize its rejection power by using all of the information from the new Run 2b SMT is an important part of achieving this goal.

7.2 STT Upgrade

7.2.1 Motivation

The $D\bar{0}$ Level 2 Silicon Track Trigger (L2STT) improves the resolution in momentum and impact parameter, and the rejection of fake tracks, compared to the central track trigger alone. The STT matched to the Run 2A SMT detector is being constructed with NSF and DoE funds for delivery in the summer of 2002. An upgrade for Run 2B, however, will be necessary in order to match the new geometry of the Run 2B Silicon Tracker.

Tracks with large impact parameter are indicative of long-lived particles (such as b-quarks) which travel for several millimeters before they decay. The L2STT thus provides a tool to trigger on events with b-quarks in the level 2 trigger. Such events are of particular importance for the physics goals of Run 2. The Higgs boson decays predominantly to $b\bar{b}$ pairs if its mass is less than about 135 GeV/c^2 . The most promising process for detection of a Higgs boson in this mass range at the Tevatron is associated production of Higgs bosons with W or Z bosons. If the Z boson decays to neutrino pairs, the b-quarks from the Higgs decay are the only detectable particles. In order to trigger on such events (which constitute a significant fraction of associated Higgs production) the L2STT is essential to detect at the trigger level jets that originate from b-quarks. The L2STT will also allow the collection of a large enough sample of inclusive $b\bar{b}$ events to see the decay $Z \rightarrow b\bar{b}$. Such a sample is important to understand the mass resolution and detection efficiency for $b\bar{b}$ resonances, and to calibrate the calorimeter response to b-quark jets. The latter will also help to drastically reduce the uncertainty in the top quark mass measurement, which is dominated

by the jet energy scale uncertainty. Detailed descriptions of the physics benefits of STT are written up as DØ Notes^{14,15}.

7.2.2 Brief description of Run 2a STT architecture

The STT is a level-2 trigger preprocessor, which receives inputs from the level 1 central track trigger (L1CTT) and the silicon microstrip tracker (SMT). The STT filters the signals from the SMT to select hits that are consistent with tracks found by L1CTT. The L1CTT uses only the axial fibers of the CFT to find track patterns. No z-information is available for level-1 tracks and SMT hits are filtered based only on their r - ϕ coordinates. Then the L2STT fits a trajectory to each level-1 track and the associated selected hits. In the fit, only axial information is used. Matching axial and stereo hits from the SMT is too complex a task to complete in the available time budget. In the selection of the SMT hits, however, the constraint is imposed that they originate from at most two adjacent barrel sections. The distribution of the hit pattern over the two barrel-sections must be consistent with a track. The fit improves the precision of the measurements of transverse momentum and impact parameter, compared to the level 1 trigger. It also helps reject fake level-1 tracks for which there are no matching SMT hits.

The STT processes these data for 12 azimuthal sectors independently. Each sector consists of 36 detector elements in four radial layers and six barrel segments. The geometry of the SMT in Run 2A provides enough overlap between adjacent detector elements that each detector element can be uniquely associated with one of these sectors without significant loss of acceptance due to tracks that cross sectors.

There are three distinct functional modules in the STT. The fiber road card (FRC) receives the data from L1CTT and fans them out to all other cards that process hits from the same sector. The silicon trigger card (STC) receives the raw data from the SMT front ends and filters the hits to associate them with level-1 tracks. The track fit card (TFC) finally fits trajectories to level-1 tracks and SMT hits. Each of these modules is implemented as a 9Ux400 mm VME card, based on a common motherboard. The main functionality is concentrated in large daughter boards, which are distinct for the three modules. Communication between modules is achieved through serial links. The serial links use low voltage differential signaling (LVDS) at 132 MB/s. We designed PC-MIP standard mezzanine boards that accommodate either 3 LVDS transmitters or 3 LVDS receivers. The motherboard has six slots to accommodate these boards. Data are transferred between the VME bus, the daughter cards and the link mezzanine boards over three interconnected 32-bit/33 MHz PCI busses. Figure 69 shows a block diagram and photograph of the motherboard.

¹⁴ "A silicon track trigger for the DØ experiment in Run II – Technical Design Report", Evans, Heintz, Heuring, Hobbs, Johnson, Mani, Narain, Stichelbaut, and Wahl, DØ Note 3510.

¹⁵ "A silicon track trigger for the DØ experiment in Run II – Proposal to Fermilab", DØ Collaboration, DØ Note 3516.

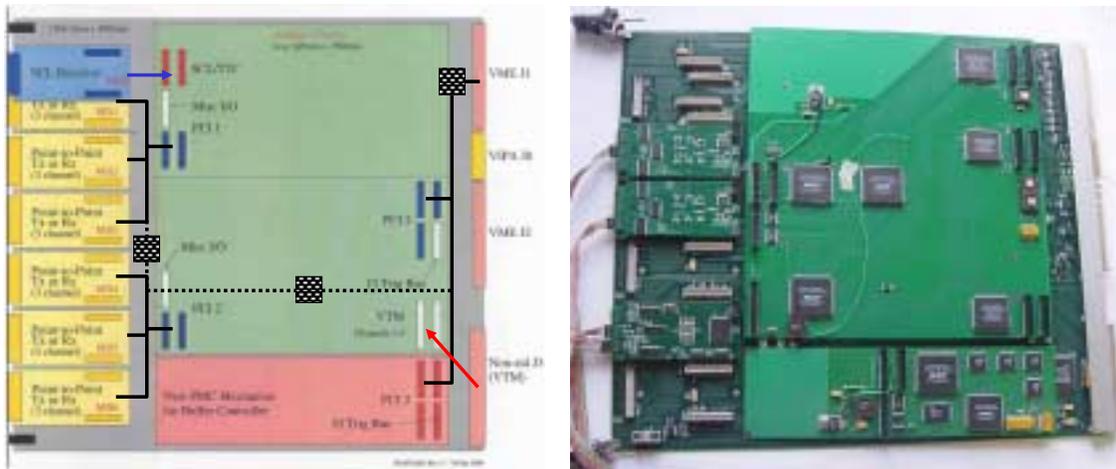


Figure 69. Block diagram and photograph of motherboard. The drawing shows the three PCI busses (solid lines) and the bridges that connect them (squares and dashed lines). The photograph also shows the FRC daughter board, the buffer controller mezzanine board and two LTBs, and one LRB.

The FRC module receives the data from the L1CTT via one optical fiber and a VTM in the rear cage. The FRC also receives information from the trigger control computer via the serial command link (SCL). This information contains the level-1 and level-2 trigger information and identifies monitor events for which all the monitor counters have to be read out. The FRC combines the trigger information with the road data and sends it to all other modules in the crate via serial links. The motherboard can accommodate up to six PC-MIP mezzanine boards. One is used to receive the SCL; the remaining five can be used for LVDS transmitter boards to fan out the L1CTT data, which provides up to 15 links. The FRC also performs arbitration and control functions that direct the flow of data for accepted events to the data acquisition system. The buffer controller mezzanine board (BC) holds a multiport memory in which events are stored until a level-2 trigger decision has been taken. There is one BC on each motherboard. The FRC manages the buffers on all BCs in the crate.

Each SMT module has eight channels, which each process the data from one silicon detector element. The signals from the SMT front ends are transmitted over a 106 MB/s serial link using the HP G-link chips and optical fibers from the electronics platform below the detector to the 2nd floor of the moveable counting house, where the STT is located. Passive optical splitters create two data paths, one to the SVX data acquisition system and another into the STT. The optical signals are received and parallelized in VME transition modules (VTM) sitting in the rear card cage of the crates that accommodate the STT modules. The VTMs are an existing Fermilab design, used by both D0 and CDF. The SMT signals are passed through the J3 backplane to the STC module sitting in the main card cage in the same slot as the VTM. Each VTM has four optical receivers and each fiber carries the signals from two detector elements.

The STC module receives the L1CTT data from the FRC over an LVDS serial link and the SMT signals via optical fibers and a VTM in the rear cage.

Each level-1 track is translated to a range of strips (a “road”) in each of the eight detector elements that may contain hits from the particle that gave rise to the level-1 track using a look-up table. The SMT data are clustered to combine adjacent strips hit by the same particle. These hits are then compared to the roads defined by the level-1 tracks. The hits that are in one or more roads are queued for transfer to the TFC module over an LVDS serial link. The main logic of the STC module is implemented in a single large field programmable gate array (FPGA).

Each TFC receives the hits from one azimuthal sector that were associated with at least one road. Because of the way SMT detector elements are mapped onto the optical fibers, three STC modules receive hits from both sectors in the crate. The outputs of these three STC modules go to both TFC modules in the crate. The remaining six STC modules receive hits from only one sector and their outputs go to only one TFC module. Thus each TFC module has six incoming LVDS serial links. The hits that come in over these links are sorted according to the level-1 track they are associated with. Then all data that is associated with one level-1 track is sent to one of eight DSPs that perform a linearized chi-squared fit. The results of the fits and the L1CTT data are sent via a Cypress hotlink to the level-2 central track trigger (L2CTT). The L2CTT acts as a concentrator for the 12 hotlink inputs from the six STT crates.

The number of crates required for the entire system is driven by the number of STC modules required to instrument all barrel detectors. Each SMT module can process the data from eight detector elements. Each azimuthal sector consists of 36 detector elements. Thus, each azimuthal sector requires 4.5 STC modules. We can accommodate two such sectors in one VME crate, so that one crate contains one FRC module, nine STC modules, and 2 TFC modules (one per azimuthal sector). In addition, each STT crate also houses a power PC and a VME buffer driver (VBD) card. The former controls the VME bus, and is used to download data tables and firmware into the STT modules and to monitor the performance of the STT. The VBD transfers the data of accepted events to the data acquisition system. Figure 70 shows the layout of one STT crate.

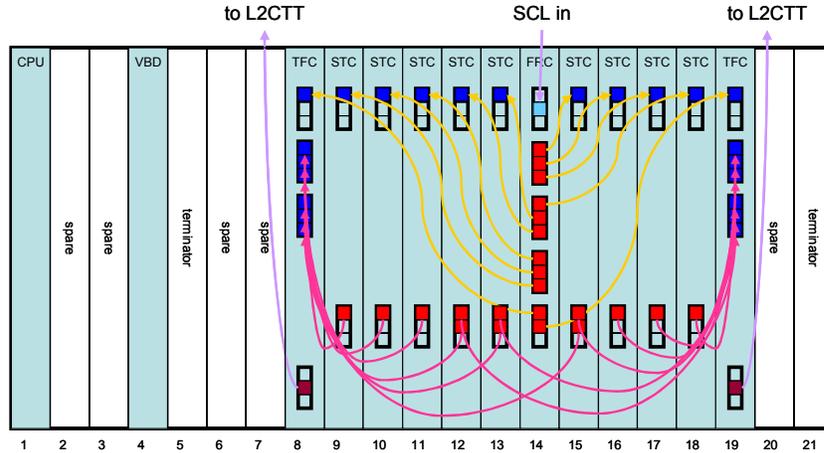


Figure 70. Layout of L2STT crate for Run 2A. The groups of three squares on the front panels indicate PC-MIP boards and the colored squares indicate used channels. The light blue square at the top of the FRC indicates the SCL receiver, and the brown squares at the bottom of the TFCs indicate the hotlink transmitters. Arrows indicate cable connections and are directed from LTBs (red) to LRBs (blue).

7.2.3 Changes in tracker geometry and implications for STT

The design of the silicon microstrip tracker for Run 2B¹⁶ foresees six concentric layers of detector elements, compared to four for the Run 2A design. The inner two layers consist of twelve 78 mm long sensors along the beam direction. Layers 2 and 3 consist of ten 100-mm long sensors and the outermost layers consist of twelve 100-mm long sensors. Figure 71 shows two views of the design.

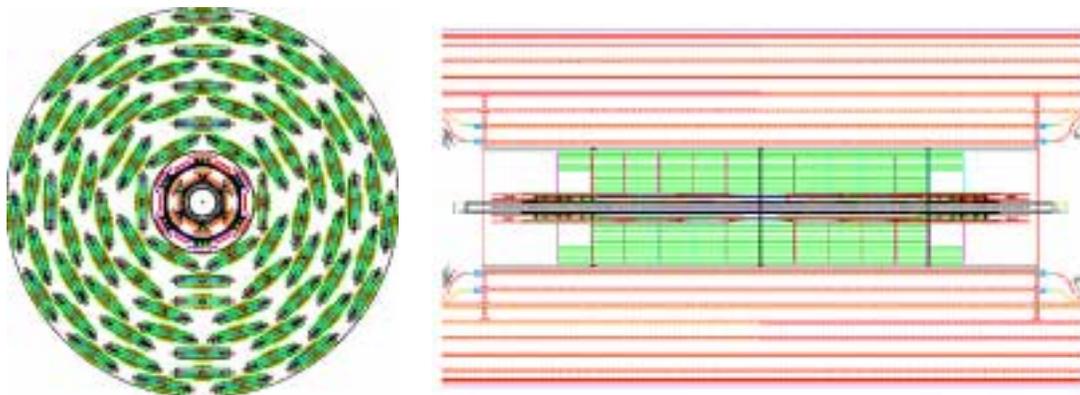


Figure 71. Axial and plan views of the Run 2B silicon microstrip tracker design.

Some sensors are ganged and in layers 1-5 every cable reads out the signals from two sensors to reduce the number of readout units (i.e. cables). Table 41 lists the number of readout units for axial strips in every layer, which

¹⁶ “DØ Run 2B Silicon Detector Upgrade - Technical Design Report”, DØ Collaboration, 2001.

determines the number of STC modules required to process their hits in the STT. The detector elements in each layer alternate between the two radii listed in the table such that adjacent detectors overlap slightly. Readout units for stereo strips are not used in the STT and are therefore not listed here. The number of readout units with axial strips increases from 432 in the Run 2A design to 552 in the Run 2B design.

Table 41 Parameters of Run 2B silicon microstrip tracker design.

Layer	Radius (axial strips)	Strip pitch	Strips	Readout units in ϕ	Readout units in z
0	18.6/24.8 mm	50 μm	256	12	12
1	34.8/39.0 mm	58 μm	384	12	6
2	53.2/68.9 mm	60 μm	640	12	4
3	89.3/103 mm	60 μm	640	18	4
4	117/131 mm	60 μm	640	24	4
5	150/164 mm	60 μm	640	30	4

The data must be channeled into TFCs such that all hits from a track are contained in one TFC. In layers 0, 1, and 2 the overlaps between adjacent detector elements are large enough so that each sensor can be uniquely associated with one TFC. This divides the detector into 12 azimuthal sectors as indicated by the shaded regions in Figure 72. To maintain full acceptance for tracks with $p_T > 1.5$ GeV/c and impact parameter < 2 mm, the data from some sensors in layers 3, 4, and 5 must be channeled into two TFCs, which are in some cases located in different crates. This is not the case in the current configuration, but should not present any problems. We are limited to 8 STC inputs into each TFC, which is sufficient for the Run 2b detector geometry.

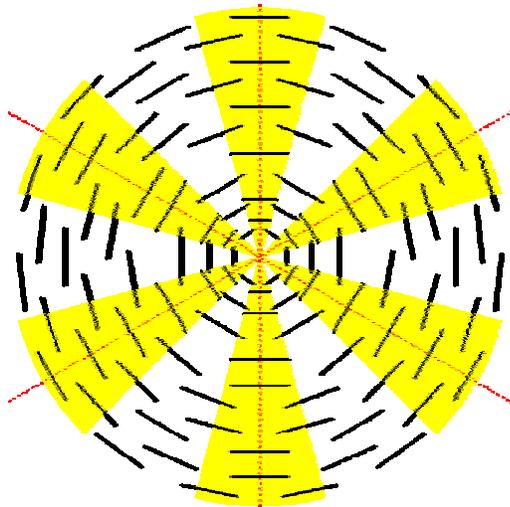


Figure 72. Azimuthal sector structure in the Run 2B silicon microstrip tracker. The yellow wedges indicate the coverage of the 12 azimuthal sectors. All detector elements in layers 3-5, which cover parts of two sectors are assigned to two TFC modules.

The hardware of the Run 2A STT is only sufficient to instrument four of the six layers of the Run 2B SMT. To include five or all six of the layers in the trigger, additional modules need to be acquired. For the most part this amounts to building additional copies of the Run 2A STT modules. None or very little new hardware design is required. The following section explains the upgrade options in detail.

7.2.4 Simulation studies of Run 2a geometry at Run 2b luminosity

Studies of the STT for Run 2B are in progress. The STT has been studied using the Run 2A geometry and trigger simulation with varying number of proton-antiproton interactions per beam crossing. At Run 2B luminosities of $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ and crossing intervals of 132 (396) ns, we expect to see on average 4.6 (13.9) soft proton-antiproton interactions in every triggered hard scattering event. In Run 2B, the hit multiplicity in the silicon detector will increase due to higher interaction rate, the larger number of axial silicon strips, and the decreased radius of the innermost layer of silicon detector elements. The STT will need to handle these conditions while maintaining a high acceptance for b-tagged events and a high rejection of QCD events.

The Run 2a STT is still under construction. Its performance will depend upon numerous factors including the performance of the SMT and L1CTT. The STT algorithm contains tunable parameters such as thresholds, road widths, and errors. The simulations quoted here have not yet been benchmarked against the performance of the detector and nominal values for the tunable parameters have been used. Further studies to maximize performance are in progress.

We find that the hit cluster multiplicity increases linearly with number of additional proton-antiproton interactions. The exact quality cuts used by Level 2 Global need tuning. Nominal values of $\chi^2/\text{dof} < 10$ and impact parameter significance $S_b = b/\sigma > 2.0$ are used in the simulation.

Our studies compare the response of the STT to QCD background events and signal events (namely, W+Higgs and Z+Higgs). QCD events are generated with $p_T > 2.0$ GeV. We have used samples with 0.5, 2.5, 7.5 and 10.0 additional interactions per event. The primary signal topology is Z + Higgs ($Z \rightarrow \nu\bar{\nu}$, $H \rightarrow b\bar{b}$) with 2.5 and 7.5 additional interactions.

We find that the signal efficiency varies little with interaction multiplicity. The rate at which QCD events are accepted increases significantly with interaction multiplicity. From this we conclude that increased luminosity results in constant signal acceptance but increased background rate. This will require tighter selection criteria or improved track reconstruction.

Offline simulations carried out using the Run 2B SMT geometry¹⁷ show that a 6-layer detector achieves better fake rejection and higher b-tagging efficiency. Requiring 4 hits in 5 layers results in a fake rate three times higher than requiring 5 hits in 6 layers. Without layer 4 the double b-tagging efficiency drops by about 10%.

The increasing track multiplicity will also increase the load on the DSPs in the TFC modules that perform the track fitting. We are presently performing timing studies with the TFC prototypes to determine whether we have to increase the number of TFC modules per crate.

7.2.5 Implementation description for primary STT options

We have considered three upgrade options for the STT in Run 2B:

- Option A includes all axial strips in the trigger and doubles the number of TFC modules to keep up with the increased number of tracks per event expected at higher luminosities. This requires three additional STC modules per crate and two additional TFC modules. To accommodate all these, we have to modify the J3 backplane in the STT crates and we have to design a fan-out module for the serial links.
- Option B includes all axial strips in the trigger but maintains the current number of TFC modules. This requires 3 additional STC modules per crate. The STT crate layout for this option is shown in Figure 73.
- Option C includes only the axial strips from five SMT layers in the trigger and maintains the current number of TFC cards. This requires one additional STC module per crate and can probably be done with the spares from Run 2A. However, we will then have to replenish our spare inventory.

¹⁷ "Evaluation of alternate designs of the silicon tracker", T. Bolton, E. Chabalina, R. Demina, A. Khanov, A. Nomerotski, F. Rizatdinova.

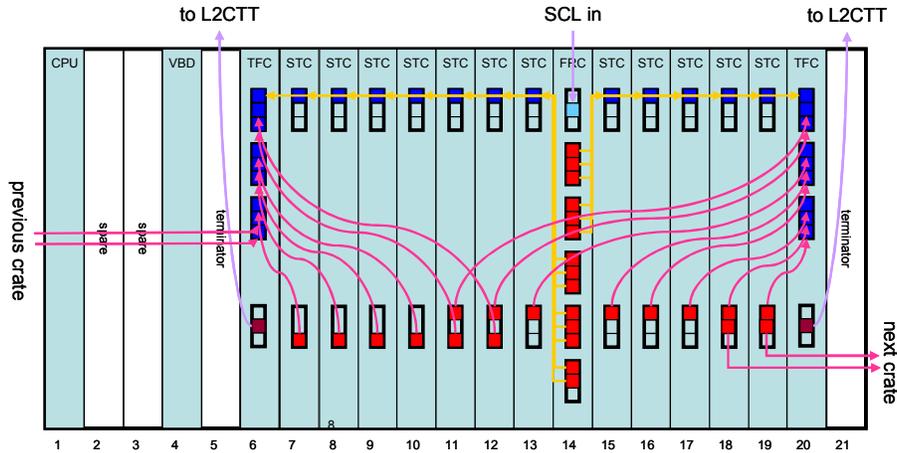


Figure 73. Layout of STT crate for upgrade option B.

Regardless of the option chosen, the fitting algorithm in the TFCs has to be modified to reflect the different number of layers and new coordinate conversion tables will have to be computed. The firmware in the STCs will have to be modified to handle the modified inputs from the L1CTT and new road look-up tables will have to be computed.

7.2.6 Summary of cost estimates and major milestones

Table 42 summarizes the cost of the three options. Quantities include approximately 10% spares. For each option, the first column indicates the number of units needed, the second is the cost estimate, in most cases based on cost estimates for the Run 2A construction (without contingency). The third column is contingency. Engineering cost is based on half a year of an engineer to modify the firmware for options B and C plus an additional half year of engineering to design the LVDS fan out and the cost of redesigning the J3 backplane for option A.

The most effective way to acquire this hardware would be at the time the production of STT modules for Run 2a takes place. Combining production runs for Run 2a and 2b, as well as purchasing many of the components before they become obsolete, would save much time, manpower, and money. Since the Run 2a STT module manufacturing is scheduled for the beginning of CY02, we will need the funds for the Run 2b STT upgrade in FY02. Parts are being ordered now for these and without additional funds we will not be able to increase the number of modules that we are planning to produce. If funds become available later, additional production runs will have to be made.

Major milestones that drive the schedule for the STT upgrade are:

- Completion of Run 2A STT will make manpower available to work on the upgrade. Additional production cycles can start at this time. The lead time for ordering parts, the production of PC boards, the assembly of the boards and testing is about six months. Thus, there is plenty of time to produce the boards and set up test systems before installation at D0.

Funds to procure parts and produce additional boards must be available at this time.

- After the Run 2B input data formats are finalized, work on firmware changes can start. Funds to cover engineering for this purpose must be available at this time.
- After the end of Run 2A, installation of the additional fiber optic splitters and the new boards at D0 can begin. Commissioning can start with test inputs and cosmic ray data from the detector.
- After the resumption of accelerator operation for Run 2B, commissioning will be completed with collider data to make the STT fully operational.

Table 42. Cost estimate for the STT upgrade options discussed in the text.

Component	Cost		Option A			Option B			Option C	
Motherboard	\$2065	32	\$55200	\$16560	20	\$34500	\$10350	7	\$12075	\$3623
BC	\$1015	32	\$31744	\$9523	20	\$19840	\$5952	7	\$6944	\$2083
STC	\$2565	20	\$61300	\$18390	20	\$61300	\$18390	7	\$21455	\$6437
LTB	\$265	52	\$20696	\$6209	26	\$10348	\$3104	7	\$2786	\$836
LRB	\$615	60	\$34140	\$10242	20	\$11380	\$3414	7	\$3983	\$1195
VTM	\$2565	18	\$43542	\$13063	18	\$43542	\$13063	-	-	-
TFC	\$5065	8	\$40520	\$12156	-	-	-	-	-	-
Hotlink	\$665	14	\$9310	\$2793	-	-	-	-	-	-
LVDS fanout	\$565	8	\$4520	\$2556	-	-	-	-	-	-
J3 backplane	\$1013	8	\$7428	\$2228	-	-	-	-	-	-
Cables	\$30	170	\$5100	\$3570	46	\$1380	\$966	20	\$600	\$420
Splitters	\$125	80	\$10000	\$3000	80	\$10000	\$3000	26	\$3250	\$975
Fibers	\$25	160	\$4000	\$1200	160	\$4000	\$1200	52	\$1300	\$390
Engineering			\$74200	\$61020		\$33600	\$33600		\$33600	\$33600
Total			\$401700	\$162510		\$229890	\$93039		\$85993	\$49558

8 Online Computing

8.1 Introduction

8.1.1 Scope

For the purposes of this document, the DØ Online systems will be defined to consist of the following components:

- DAQ and Online network,
- Single Board Computers (SBCs) in VME readout crates,
- Level 3 Linux software filter farm,
- Host Online system,
- Control room computing systems,
- Data monitoring computing systems,
- Database servers,
- File servers,
- Slow control system, including VME processors,
- plus the associated software for each of these elements.

8.1.2 Software Architecture

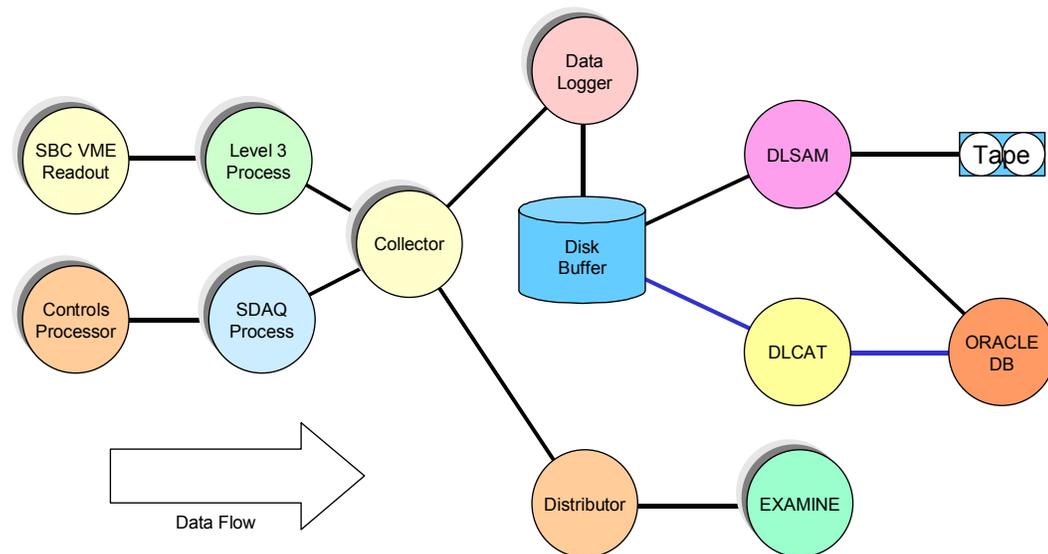


Figure 74. DAQ system software components

The software architecture of the Run 2b Online system is unchanged from that of Run 2a. Some components will need updating, but there are no structural differences planned. The major software components in the event data path are illustrated in Figure 74. The slow control system components are not illustrated in the figure, nor are the non-event monitoring systems.

8.1.3 Hardware Architecture

The hardware architecture of the Run 2b Online system is also unchanged from that of Run 2a. The current architecture is illustrated in Figure 75 and Figure 76. At the center of the system are two high capacity network switches (Cisco 6509). The event data path includes the Single Board Computers in the VME readout crates, the Level 3 Linux filter nodes, the Host Online systems on which reside the Collector, Distributor, and Data Logger processes, and the final data repository in the Feynman Computing Center (FCC). The EXAMINE processes on the Monitor system nodes provide real-time event data analysis and monitoring functions. Some of the Slow Control system nodes also participate in the Secondary Data Acquisition (SDAQ) path. An ORACLE database serves for configuration control and recording of run parameters. Also included in these figures are the Control Room, File Server, and Slow Control system nodes.

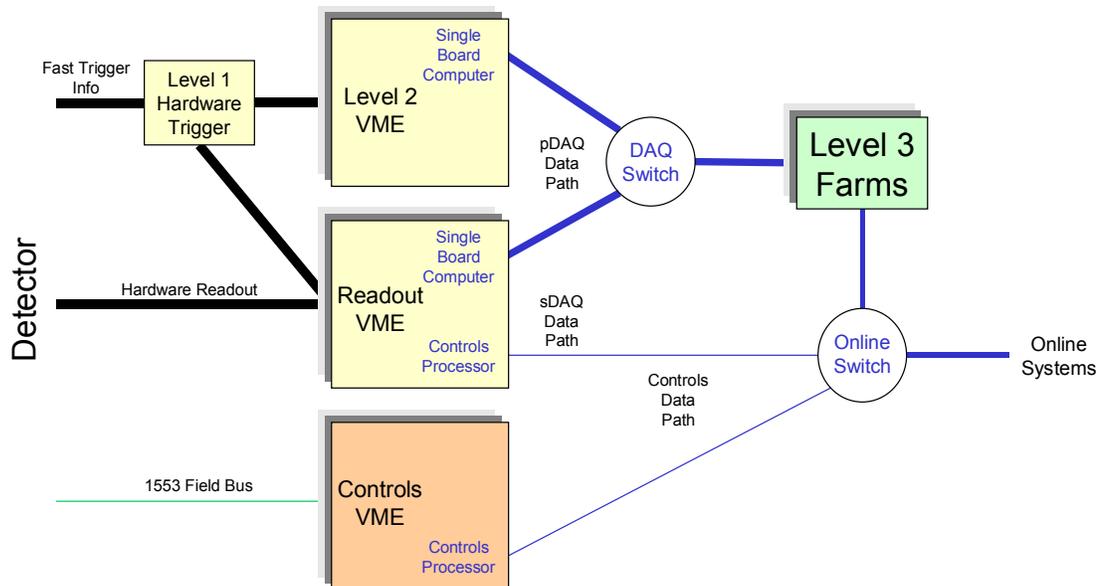


Figure 75. DAQ system hardware components.

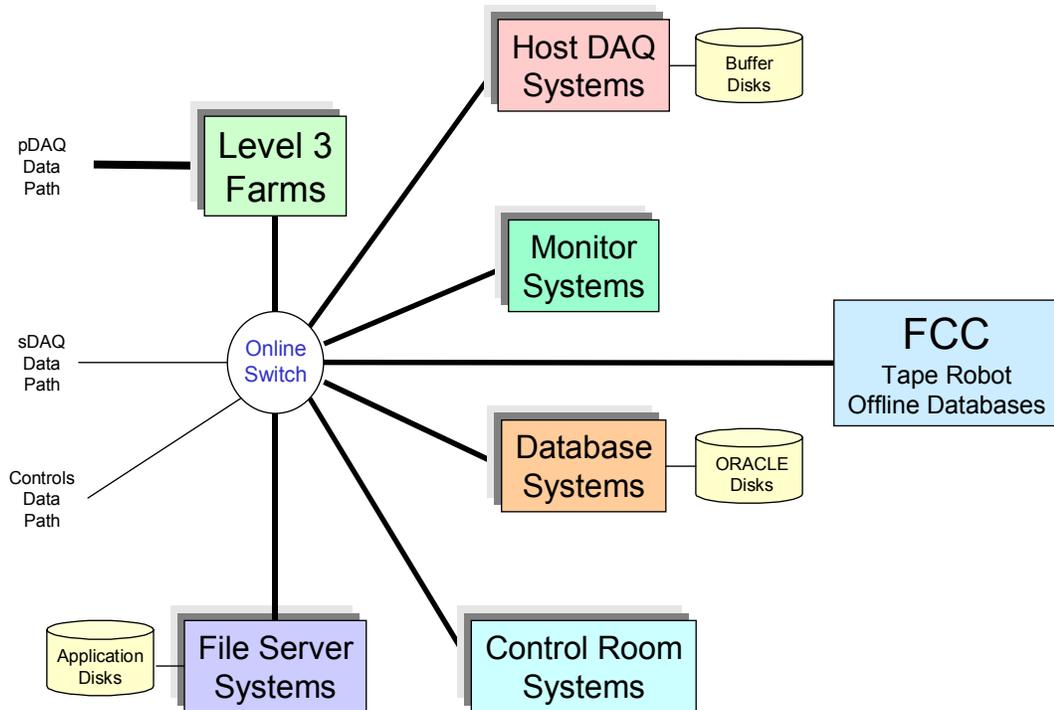


Figure 76. Online system hardware components.

For Run 2b many components of these computer systems will need to be updated or replaced.

8.1.4 Motivations

The primary considerations governing the development of the DØ Online system for Run 2b are supplying the enhanced capabilities required for this running period, providing hardware and software maintenance for the (by then) five-year old hardware, and supplying the required software support. We expect the requirements for the Online data throughput to at least double, largely driven by the ability of the Offline analysis systems to absorb and analyze the data. Many of the existing Online computing systems will reach the end of their viable lifetime in capability, maintainability, and software support by the Run 2b era. The gradual replacement of many of these component systems will be essential.

8.1.4.1 Enhanced Capabilities

The rate at which DØ records data to tape has been limited by the cost of storage media and the capability of the Offline systems to analyze the data. Assuming five years of improvements in computing capability, it is reasonable to expect that the Offline capacity for absorbing and analyzing data will more than double. The Online system must be capable of providing equivalent increased data throughput.

After five years of experience in analyzing events, more sophisticated software filters will be running on the Level 3 trigger farm. These more complicated codes will increase execution time. The resulting increased

computing demand in Level 3 will need to be met by either an increase in the number of processors, replacement of these units by more capable processors, or both.

It is also expected that data quality monitoring software will be vastly improved by the Run 2b era. These capabilities again are likely to come at the cost of increased execution time and/or higher statistical sampling requirements. In either case, more numerous and more powerful monitoring systems will be required.

8.1.4.2 Hardware and Software Maintenance

By the time of Run 2b, the computing systems purchased for Run 2a will be more than five years old. In the world of computing hardware, this is ancient. Hardware maintenance of such old equipment is likely to be either impossible or unreasonably expensive. Experience shows that replacement by new (and under warranty) equipment is more cost effective. Since replacement of obsolete equipment not only addresses the maintenance question, but also issues of increased capability, it is the most effective course of action.

The DØ Online system is composed of several subsystems that have differing hardware components and differing maintenance needs. Subsystem specific issues will be addressed in the following sections.

8.1.4.3 Software Support

Several different operating systems are present in the Online system, with numerous custom applications. We have tried, wherever possible, to develop software in as general a fashion as possible so that it can be migrated from machine to machine and from platform to platform. However, support of certain applications is closely tied to the operating system on which the applications run. In particular, ORACLE database operations require expertise that is often specialized to the host operating system. By the time of Run 2b, there is expected to be a consolidation in ORACLE support by the Laboratory that will not include the existing DØ Online database Compaq / Tru64 Unix platform. These platforms will thus need to be replaced.

8.1.5 Interaction with the Computing Division

The Run 2a Online system was developed through an active partnership with the Computing Division's Online and Database Systems (CD/ODS) group. It is essential that this relationship be maintained during the transition to the Run 2b system. While the level of effort expended by CD/ODS personnel has already decreased relative to what it was during the height of the software development phase of the Run 2a Online system, the continued participation of this group will be needed to maintain the system and to migrate the existing software to new platforms as these are acquired. Computing Division assistance and expertise will be particularly critical in the area of database support since the Oracle consultant who led the design of the current system is not expected to be involved in maintaining the system. The continued involvement of the CD in the Online effort, which will presumably be described in a future MOU, will be left

mostly implicit in later sections of this document, but will nevertheless continue to be crucial to the success of the effort.

8.1.6 Comments on Run 2a configuration

For Run 2a many of the Online functions have been combined on a single cluster of large servers. The Data Logger, Database, and File Server operations are all performed on one or more of a set of three Compaq AlphaServers configured as a TruCluster. The cluster configuration allows each of the three nodes to share disk resources, which include the event data buffer disks, database disks, and general user disks. The cluster acts as an NFS server for the control room and monitoring nodes, and as the NIS master for Online accounts. High availability and reliability for these functions are provided by the cluster configuration, allowing one of the three nodes to go down with the remaining nodes assuming the critical functions.

The Run 2a cluster includes a dual-processor AlphaServer 4000 (purchased in 1997), a second dual-processor AlphaServer 4000 (1998), and a quad-processor AlphaServer GS80 (2000). Disk storage shared among the cluster members includes 1.1 TB in a fibre channel RAID array, 2.8 TB in fibre channel JBOD, and 0.6 TB in a shared SCSI RAID array. Backups are performed on a single DLT4000 tape drive on each cluster member.

The remaining Online Host, Control Room, and Monitoring nodes are single or dual processor Linux systems. These nodes are mostly interchangeable, with only the Control Room systems being slightly unusual with their configuration requiring multiple graphics cards and monitors. Other than system and scratch disks, these systems get all of their storage resources from the AlphaServer cluster.

The cluster configuration has proved both reliable and efficient. However, the I/O performance of the AlphaServer components, because of both age and architecture, is not optimal. The existing Run 2a components are maximally utilized in order to provide the Run 2a target event data rate of 50 Hz. Concentrating a large number of functions in a small number (3) of machines leads to potential congestion. The addition of resources to this specialized configuration is possible but costly.

The Run 2b architecture addresses the performance and expandability / flexibility issues by replacing the central cluster with a larger number of dedicated function Linux systems. There is a philosophy of one machine per function, with redundant systems where necessary. Where possible, functions are spread across multiple parallel machines. The details of the architecture are described in following sections.

8.2 Plan

A description of planned upgrades follows for each component noted in the Introduction. The philosophy and architecture of the Online system will not change, but components will be updated. Note that most changes are best

achieved by a continuous, staged approach, while others involve large systems that will need to be replaced as units.

8.2.1 Operational Parameters

Table 43 summarizes the Run 2 parameters relevant to the Level 3 and Online systems. The impact of each value upon the system configuration will be discussed in the following sections.

Table 43. Run 2 Level 3 and Online parameters

<i>Parameter</i>	<i>Run 2a</i>	<i>Run 2b</i>
Level 3 farm nodes (dual processors)	64	128
Level 3 processing time budget	~ 125 msec/event	~ 250 msec/event
Average event size	~ 250 Kbytes	~ 300 Kbytes
Level 3 input rate	1000 Hz	1000 Hz
Peak Level 3 accept rate	50 Hz	100 Hz
Peak logging rate	12.5 Mbytes/sec	30 Mbytes/sec
Detector duty factor	> 99%	> 99%
Accelerator duty factor	~ 75%	~ 75%
Online system availability	> 99%	> 99%
Local data buffer	48 hours	48 hours
Local data buffer	~ 2 Tbytes	~ 4 Tbytes

8.2.1.1 *Level 3 parameters*

The required number of Level 3 farm nodes is a function of the event rate into Level 3 and the required processing time for filtering. These are highly tunable numbers. The input rate can be adjusted with thresholds in Level 1 and Level 2. The processing time depends upon the choice of software filters. If a required Level 3 rejection rate is not possible in the available time, then thresholds can be adjusted at Level 3. The choice in Table 43 of 128 Level 3 farm nodes for Run 2b results from the best estimate of the required analysis time with the target input rate.

8.2.1.2 *Event rate parameters*

The rate at which events are logged determines the required capabilities of the Host Online system. The rate is bounded by the maximum trigger rejection ratio and the maximum Offline storage, reconstruction, and analysis capacities. The Online system must be designed to cope with the largest rate otherwise allowed. This limitation is from the Offline computing systems, which see a Run 2b rate of 100 Hz with a 75% overall duty factor as the maximum allowable.

8.2.1.3 System availability

The DØ experience in Run1 was that an uptime of > 99% for the computing systems is achievable. In Run2 there is a similar goal, such that lost beam time from computing problems is small compared to other sources of detector down time (the Detector duty factor). Since the DØ event data is transferred to the Feynman Computing Center for logging, problems in this step can potentially contribute to system down time. To decouple possible problems with the tape robot and database systems, which are principally Offline systems with competing priorities, we require a local disk buffer to retain data for a period of 48 hours. This period should be sufficiently long to recover from any Offline system disruptions.

8.2.2 Online Network

8.2.2.1 Description

The backbone of the DØ Online computing system is the switched Ethernet network through which all components are interconnected. The Run 2a network is based on a pair of Cisco 6509 switches, one for the event data path from the SBCs in the VME readout crates to the Level 3 filter farm nodes, and the other which services the Level 3 event data output as well as general network traffic. Each switch is composed of a chassis with an interconnecting backplane and various modules that supply ports for attaching the DAQ and Online nodes. The total capacity of each switch is determined both by the chassis version and the number and versions of the component modules. The DAQ switch has a fabric-enabled backplane capable of a total throughput of 128 Gbps. The Online switch has a backplane capable of a total throughput of 64 Gbps.

The Cisco 6509 for the DAQ network can currently support 16 Gb fiber connections from the SBCs (via Cisco 2948G switches in the counting house) and 96 100Mb connections to Level 3 farm and support nodes. There are approximately 32 ports available for additional farm nodes. If the number of farm nodes is to be increased beyond that level, then additional 100Mb blades will need to be purchased for the switch.

The Online Cisco 6509 switch currently directly supports over 100 nodes, including 48 Level 3 nodes. There are approximately 48 available 100Mb ports, but more 100Mb blades would be needed to expand beyond that level. An increase in the number of high bandwidth host system nodes will require more gigabit ports beyond the current fully-utilized 10 ports, necessitating the addition of a Gb capable blade.

8.2.2.2 Run 2b Upgrade Plan

To support a total of 128 Level 3 nodes, one additional 100baseT module is required on each switch. To support the expanded Host Data Logging systems (page 158), one additional 1000baseT module is required in the Online switch. If a 1000baseT module is not available, existing 3COM 4900 switches (with 24 available 1000baseT ports) can be used with 1000baseSx uplinks. In this case an additional 1000baseSx module would be needed.

8.2.3 Level 3 Linux Filter Farm

8.2.3.1 Description

The Level 3 trigger consists of two principle elements: a high speed data acquisition system that provides readout of the entire detector at rates expected to exceed 1 kHz, and a processor farm that utilizes software filters written to select events that will be permanently recorded. Since the required Run 2b data acquisition bandwidth is expected to be made available once the Run 2a Level 3 hardware is fully commissioned, the most likely need for Level 3 upgrades will be to provide increased processing power in the farm.

The Run 2b Level 1 and Level 2 triggers will have increased selectivity and there will be a matching increase in selectivity of the Level 3 filter. This requires the use of more complex algorithms that necessitate the need for faster processors in the Level 3 nodes. Historically, DØ has equipped the Level 3 farm with the fastest processors on the market within this chosen processor family and this approach must be continued. At the time of the Run 2b upgrade, Moore's law would lead us to expect a four-fold increase in processing speed over what is currently available. Thus, a significant increase in Level 3 processing power will be obtained by enhancing the Run 2a Level 3 processors with the latest technology available in 2004.

There are currently 48 dual-processor Linux nodes in the Run 2a Level 3 filter farm. An additional 16 are to be purchased with funds allocated to the Run 2a DAQ enhancement.

8.2.3.2 Run 2b Upgrade Plan

The computing capacity of the Linux filter farm will be stressed at desired Level 3 input rates with only the existing hardware. As Offline analysis software improves, some algorithms will move into Level 3. As Level 2 filter algorithms are improved, the complexity of the Level 3 algorithms will increase in tandem. All of these efforts to enhance the capability of the Level 3 trigger will come at the expense of processing time. More and improved filter nodes will therefore be required. In order to allocate a processing time budget of 250 msec per event, a total of 128 dual-processor nodes are required. An additional 64 dual-processor nodes must then be acquired for Run 2b operations.

The Level 3 farm is very similar to the processing farms used in the Offline event reconstruction and analysis. A standard configuration is a 2U rack-mounted system with dual processors and 1 Gbyte of memory. The Level 3 nodes benefit from dual 100 Mb network connections to segregate input and output event data networks. The evaluation and purchase steps for Level 3 farm nodes will closely follow Offline farm activities.

8.2.4 Host Online Systems

8.2.4.1 Description

The Online Host nodes perform the functions of the Collector, Data Logger, Data Distributor, and the processes for shipping event data to the Feynman

Computing Center (FCC) for recording on tape. The Collector (one or more processes acting in parallel) receives event data from the Level 3 filter farm nodes and routes it to Data Logger and Data Distributor processes. The Data Logger (one or more processes acting in parallel) writes the event data to the buffer disks. The Data Distributor maintains an event pool for monitoring applications and transmits the events to the various monitoring nodes. The DLSAM and DLCAT processes read the event data and metadata from the buffer disks and ship the information over the network to the FCC.

The Host system must be able to absorb over the network the maximum recorded data rate of 30 MB/s. This rate must be sustained as events are logged to the buffer disk, read back from the buffer disk, and shipped over the network to the FCC. Additionally, a fraction of the event stream must be routed to the event monitoring tasks. There is additional network traffic internal to the Host system as events are routed among parallel tasks.

The Host system is an integral and required component of the experiment. As such, it must be a highly available system. The target availability is greater than 99%.

8.2.4.2 Run 2a Host Online system

The current DØ Online Host system is centered upon three Compaq / Digital AlphaServers in a cluster configuration. Two of the machines are AlphaServer 4000s (purchased in 1997 and 1998) and the third is an AlphaServer GS80 (purchased in 2000). These machines mount disks in the form of two RAID arrays, ~500 GB in a Compaq/Digital HSZ50 unit and ~800 GB in a Compaq/Digital HSG80 unit, and an additional 2.8 TB in Fibre Channel JBOD disk. This cluster supports data logging, the ORACLE databases, and general file serving for the remainder of the Online system.

The long-term maintenance of these systems is a serious concern. While they can be expected to still be operational in the Run 2b era, the high availability required for critical system components may be compromised by the inability to obtain the necessary maintenance support. Maintenance costs for these systems, particularly 7x24 coverage, will increase with age. By the time of Run 2b, maintenance costs are likely to exceed replacement costs.

These systems currently run Compaq Tru64 UNIX, previously known as Digital UNIX, or Digital OSF1. With the pending purchase of Compaq by Hewlett Packard, long-term support for this operating system may be problematic.

8.2.4.3 Run 2b Upgrade plan

All applications developed for the data acquisition system that currently run on the Host systems were written with portability in mind. In particular, all will work under Linux. The proposed upgrade to the Host systems is therefore to replace them with Linux servers. Since the existing Host system provides Data Logging, Database support, and File Serving functions, each of these needs must be accommodated by the replacement system. These requirements will be addressed individually in this and following sections.

The Data Logging system must, with high (> 99%) availability, be capable of absorbing data from the Level 3 filter systems, distributing it to logging and monitoring applications, spooling it to disk, reading it from disk, and dispatching it to tape-writing nodes in the FCC. The maximum input data rate is 30 Mbytes/sec. A local disk buffer of ~4 Tbytes is required to retain event data if the Offline tape robots are unavailable. The event data on the buffer disks is normally read and transferred to the FCC at the 30 Mbytes/sec input rate, but a 2x higher rate is necessary to both take sustained new data and unload stored event data following any outage. The high availability requirement, satisfied in the current system by using a cluster of three machines, precludes the use of a single machine. Currently the cluster members share the disk buffers, but this is not a strict requirement.

The proposed upgrade solution is for a set (two or three) of Linux servers (dual or quad processors) to act as the new Data Logging nodes. The data acquisition applications can run in parallel to distribute the load at full bandwidth, but a single node should be capable of handling nearly the entire bandwidth for running under special conditions. Each system will require gigabit connectivity to the Online switch, thereby raising the number of gigabit ports required.

Some R&D effort is needed to test such a configuration. The possibility of clustering the Linux nodes and the possibility of sharing the disk storage will be examined. The purchase of the complete Data Logging system can be staged, as not all members need to be identical (as noted above, the current Host system was purchased in three increments). It is expected that the R&D unit to be purchased could suffice as the third server in a three-server configuration.

A possible configuration consists of three high-end servers with a SAN JBOD disk array. A typical server would be configured with four CPUs, 4 Gbytes of memory, a local RAID controller for system disks, gigabit network cards, and available 64-bit PCI slots. The buffer disk array can be built from commodity disks in a JBOD Fibre Channel crate, as currently employed for the Run 2a system.

Other components of the Host Online system – the Collector and Distributor – will be housed on Run 2a Linux systems which will remain sufficiently capable for these functions in Run 2b.

8.2.5 Control Room Systems

8.2.5.1 *Description*

The current DØ Control Room system is composed of 12 Linux nodes (single and dual processor) that manage 27 monitors. These systems range in age from one to five years. Many of the monitors are already showing the effects of age. It is expected that we should replace some fraction of the Control Room nodes and monitors each year.

8.2.5.2 *Run 2b Upgrade plan*

The Control Room systems will be gradually upgraded and replaced throughout the lifetime of the DØ experiment. Assuming a 5-year viable lifetime

for these components, we should expect to replace 20% each year. This implies that 2 to 3 systems supporting 5 to 6 monitors will need to be replaced annually. This will be done with purchases of moderate-performance (dual CPU, multiple graphics cards) graphics systems or reuse of systems otherwise made available.

8.2.6 Data Monitoring Systems

8.2.6.1 Description

Real-time monitoring of event data is accomplished by a scheme in which representative events are replicated and distributed to monitoring nodes as they are acquired. The monitoring ranges from examination of low-level quantities such as hit and pulse height distributions to complete event reconstruction. In the latter case, the environment and the code are similar to that of the Offline reconstruction farms. There are one or more monitoring applications for each detector subsystem, and for the trigger, luminosity, and global reconstruction tasks.

The rate at which the monitoring tasks can process events, as well as the complexity of monitoring, is limited by the processing capabilities of the monitoring nodes. The Control Room systems and several rack-mounted Linux nodes currently share this load. Much can be gained by upgrading the experiment's monitoring capability. As more sophisticated analysis software becomes available, these improved codes can be run in the Online environment to provide immediate feedback on data quality.

8.2.6.2 Run 2b Upgrade plan

The monitoring nodes, rack mounted Linux systems, will need to be continually updated. Such upgrades can occur gradually. As with the Control Room nodes, a useful lifetime of 5 years implies that 20% of the monitoring systems should be upgraded each year. A typical monitoring node configuration is a 2U rack-mounted dual-processor system with 1 GB of memory. These systems are very similar to the Level 3 or Offline farm nodes, and purchases will be made in conjunction with these larger acquisitions. Four or five nodes will be upgraded per year.

8.2.7 Database Servers

8.2.7.1 Description

The ORACLE databases currently run on the AlphaServer cluster, with the database files residing on the attached RAID arrays. As mentioned above, long-term support for this hardware is questionable. Additionally, ORACLE database and application support from the Computing Division no longer includes the Tru64 UNIX platform.

The principal requirement for the database server is high availability (> 99%). Support needs include maintaining the hardware, the operating system, and the application software (ORACLE). User application development also benefits from having independent production and development database instances.

8.2.7.2 Run 2b Upgrade plan

The planned replacement of the database servers is by two redundant SUN or Linux systems with common access to RAID disk arrays. The Computing Division supports both of these systems. The purchase of these systems is best staged over two years, with early purchase of the development machine and later purchase of the production machine. The performance required of the database machines is not expected to be extremely demanding. A mid-range server system should be sufficient.

The RAID array for the database will likely be shared with the File Server systems, as neither application has particularly demanding disk I/O performance. The production database instance is expected to require no more than 400 Gbytes, and 100 Gbytes should suffice for the development instance (Run 2a sizes are 80 Gbytes and 35 Gbytes). An additional 200 Gbytes will be necessary as spooling space for the ORACLE backups.

The Database system will be supported with a backup system, in conjunction with the File Server systems of the next section.

8.2.8 File Servers

8.2.8.1 Description

The Host cluster currently provides general-purpose file serving. Linux nodes within the Online system access the Host file systems by NFS. Approximately 500 GB of RAID disk is currently available. Files stored include the DØ software library, Fermilab software products, DAQ configuration files, detector subsystem application data, and user home areas. Since the existing file servers are the AlphaServers, replacement is necessary, for reasons already delineated.

The requirement for the file server system is primarily one of high reliability (> 99%) of both system and disks. The needed network and disk I/O rate is moderate, with 5 Mbytes/sec being sufficient and easily achievable.

8.2.8.2 Run 2b Upgrade plans

The proposed solution is a pair of redundant Linux servers with common access to both RAID and JBOD disk arrays. The RAID array will be shared with the Database systems. Assuming a 1.3 TB RAID array and the previously stated database needs, then 600 GB will be available for general system use. The JBOD disk is typically provided from otherwise unused disk space on local disks of the distributed nodes, so no explicit purchase is necessary. A tape stacker system of moderate capability is needed for backups. Acquisition of these systems can be staged.

8.2.9 Slow Control Systems

8.2.9.1 Description

The Slow Controls system consists of host-level console computers that run the controls application programs and that communicate with Input/Output

Controller (IOC) computers over an Ethernet LAN. The IOC's, in turn, communicate with detector hardware components (many of which have their own embedded process control computers) over two types of field bus: (a) the VME parallel backplane and (b) the MIL/STD1553B serial bus.

An IOC processor is either a Motorola 68K or PowerPC single-board computer. The operating system for these processors is VxWorks and the controls-specific software consists of EPICS (Experimental Physics and Industrial Control System) plus a number of EPICS extensions developed by the DØ controls group. The node-specific EPICS configuration, which is loaded into an IOC at boot time, is maintained in an ORACLE database from which the IOC configuration files are extracted. The processors have both Ethernet connections to the Online network and serial line connections from their console ports to network terminal servers. The console connections allow users to communicate directly with the VxWorks shell and the EPICS local user interface.

The IOC nodes perform downloading, monitoring, calibration, and other time-critical functions essential to the operation of the detector. The host-level nodes execute applications such as operator monitoring and control GUI's, the detector configuration manager, the central significant event (alarm) system, and other, less time-critical tasks.

Multiple MIL/STD1553B busses provide the communication link to all the electronic components on the platform and many of the components in the movable counting house. Controller cards that are located in VME crates containing an IOC processor drive these busses. The MIL/STD1553B bus was selected for its robustness and for the low electrical noise environment required for devices located on the detector platform.

A significant fraction of the sensors that are managed by the slow controls system are connected via a generic analog/digital interface unit called a rack monitor (RM) that communicates with an IOC via a MIL/STD1553B bus. The environmental conditions in electronics racks on the platform and in the movable counting house are monitored by a Rack Monitor Interface (RMI) unit that, in turn, is connected to a RM.

8.2.9.2 Hardware Upgrade Options

The operation of a number of IOC's is being compromised by insufficient memory and a number of these must be replaced in the immediate future to meet Run 2a needs. The current inventory of MIL/STD1553B controllers, RM's, and RMI's is insufficient to meet the requirements of the Run 2b detector. This includes units for the detector expansion plus an adequate number of spares to maintain the detector.

- *PowerPC Processors*

Both of the Motorola 68K and PowerPC processor families have limited lifetimes. By the beginning of Run 2b, repairs or replacements for the 68K processor boards will no longer be available and, by the end of the run, the same situation may also exist for the PowerPC boards. Without a change in single-

board computer architecture (for example, moving to Intel processors with a significant accompanying software effort) DØ must be able to sustain operation with the existing systems through Run 2b. At a minimum, an adequate number of spare PowerPC boards - the number based on operational experience - must be purchased and the existing 68K boards in the slow controls system must be replaced.

The functionality of the 68K processors in the Muon detector read-out crates is now severely limited by their available memory (4 Mbytes). Due to the age of these processors, memory upgrades are no longer available. Monitoring, control, and calibration functionality would be greatly improved by a complete replacement of these aging processors.

- *MIL/STD1553B Controllers*

There are seven MIL/STD1553B controller spares held by the controls group and an additional five to ten controllers are located in various test benches at Fermilab and at several of the collaborating universities. Ten controllers are being built but parts are only available for five.

Controller failures occur approximately once a month and the parts to repair most of the failures are obsolete and in short supply. The controllers are located in the movable counting house and not on the platform. When repair or replacement is needed they can be replaced without requiring an access. The board could be redesigned and new ones built at Fermilab. This requires financing for board engineering and construction.

Commercially available replacements can be purchased that fit into VME crates, as the current ones do, or into PMC (PCI mezzanine card) slots in a PowerPC processor. Unfortunately, they are all military grade and very expensive. For either case, some research and testing must be done to insure that the commercial hardware meets our specifications. Additional effort will be required from the controls group to develop a device driver to support the new controller.

There are too many custom designed hardware components in the detector that communicate via the MIL/STD1553B bus to consider replacing it with a contemporary communication link.

- *Rack Monitors*

These devices were designed during the early 1980's and have many components, e.g. the MIL/STD1553B protocol chips, which have not been manufactured since the early 1990's. DØ has purchased the available stocks of these components from known suppliers and cannibalized old equipment in order to maintain our existing stock of approximately 175 RM's.

Although the RM's used during Run I were very reliable, there is an increasing failure trend for these modules that was observed during the enhancement of the Run 1 slow control system in preparation for Run 2a. Considering that: (1) these modules (or an equivalent) are critical for the

operation of the detector, (2) some fraction of these modules will fail before the end of Run 2b, (3) the number of working spare modules is ~3%, and (4) the Run 2b upgrade will require additional RM's, a plan for the production of additional units and the repair and replacement of existing units is clearly required.

Several options exist for supporting or replacing the rack monitors:

- Additional RM's could be manufactured. This would require functional replacements for the components that are no longer available. For example, the critical protocol chip might be replaced by a piggyback FPGA that reproduced the external behavior of the original protocol chip. This might appear to be the least-cost solution; however, there are other components in the RM that will become unavailable in the next few years and a plan that relies upon building functional replacements for these has risks.
- The RM itself could be redesigned using contemporary components to provide a functionally equivalent module with the same form factor. This would allow the existing modules to be replaced, as they fail, without requiring moving other modules mounted in the same rack. This solution requires an investment in engineering design and testing in addition to the direct manufacturing and these have not yet been estimated.
- An existing commercial product might be found; however there will probably be significant software development costs to adapt it to the existing slow controls system. There is also the probability that the form factor will be larger than the present module (1U height) and that would require rearrangement many of the electronics racks, which are already filled.

One possible commercial replacement would be the Internet Rack Monitor (IRM), manufactured by the Bi-Ra Corporation. This device is already supported by our slow controls system; unfortunately, however, it has a 4U form factor -- four times the size of the existing module.

- There is an ongoing project, the HOTlink rack monitor (HRM), in the Beams Division, where the original RM was designed. This project is in the early design and testing phase and could provide a staged replacement solution.

Although the individual RM replacement modules could have the same 1U form factor, the organization has significant differences. Individual RM-like modules, which have the same external connector format, are linked in a star network of high-speed serial links to a central processor that could be an existing or additional IOC in our slow controls system. The communication link exists as a PMC card and, with VME extender cards, a single IOC could service six PMC cards and, therefore, six RM replacements. This solution is particularly attractive in the moving counting house and the platform where a single IOC could service up to six adjacent racks.

The HRM has a much higher maximum data acquisition rate (~100 KHz) than the RM and, while these rates are not needed for existing monitoring connections, the availability of higher rates could be very useful in studying transient phenomena.

Some software development time would be required to integrate the VxWorks driver for the communication links into the DØ slow controls system.

The first version of HRM should be available for testing later this year and, by committing to the use of the HRM, DØ could influence some of the features of the design to minimize the impact of RM replacement. The cost of a single RM-equivalent with its communication link is estimated to be ~\$2K which, at current prices, would provide a replacement for six RM's at a cost of ~\$16K.

- *Rack Monitor Interfaces*

The Rack Monitor Interface (RMI) monitors the environmental conditions (air temperature, air flow, cooling water flow, smoke, water drips) in most of the electronics racks on the detector platform and in the movable counting house. The RMI is constructed from currently available components.

An unknown number of these units will be required for the additional electronics racks required by the Run 2b detector upgrades.

8.2.9.3 Operating System Upgrade Options

We are currently using version 5.3.1 of the VxWorks operating system. The Fermilab Computing Division provides support for VxWorks 5.3.1 and is proposing to support the "Tornado" product from Wind River Systems. It appears that we can delay until the end of Run2b without changing the operating system, assuming we can get legacy support from the Computing Division.

Tornado provides an updated version of VxWorks (version 5.4) and integrated development tools. A move to Tornado would offer developers a suite of tools to help develop their programs. In particular, Tornado provides a very useful debugging tool that operates at the source language level. EPICS is in use at other sites in conjunction with Tornado. DØ might be required to pay for Tornado licenses - instead of having it provided by the Computing Division - and there are issues related to the restricted number of hosts on which the Tornado development tools will run.

It is also conceivable to use a different real-time operating system. EPICS is currently in the beta test phase for a version that is operating system (OS) independent. The current version R3.13 is closely tied to VxWorks. Conversion will require the controls group to port the DØ specific extensions to EPICS to work on the new OS. Extensive testing of the new configuration is also required.

8.2.9.4 Run 2b Upgrade Plan

The 68K processors currently installed in the muon readout crates and scheduled to be installed in the luminosity and forward proton readout crates do

not have sufficient memory to support the operating system, the EPICS system, and the readout tasks. An immediate upgrade of 19 processors from the MV162 model (68K) to MV2301 models (PowerPC) is essential. A total of 21 MV2301 processors should be purchased, 19 for installation in readout crates plus 2 spares.

The design and production of new 1553 controllers at Fermilab will be an expensive option. The least costly option is to locate all the existing controllers and, if sufficient parts can be located, to construct 20 additional controllers. If this cannot be accomplished, studies should begin immediately to select a commercial supplier.

For the RM, two options should be pursued. If sufficient parts can be located, an additional 50 units should be constructed in the near future. In addition, the HRM development project in the Beams Division should be closely monitored and, if possible, a prototype replacement for the RM should be built and tested.

In anticipation of the additional electronics racks that will be installed for Run 2b, 25 additional RMI's should be built in the next few years.

New development tools would aid in the debugging of EPICS and other VxWorks applications at DØ. Effort should be assigned to investigate whether Tornado should be used. Switching to a new OS, such as Real-Time Linux, is possibly expensive and will require extensive development time. More study on this issue is required before a decision is possible.

8.3 Online Computing Cost Summary

Table 44 Run 2b Online Computing Cost Summary, thousands of \$
Funds type is indicated in last column

Equipment	M&S Total	Contingency %	Cost	Total Cost	Funds Type
Network	40			50	
16-port 100baseT modules	22	25	6	28	OP
16-port 1000baseTx module	11	25	3	14	OP
patch panels, cables	6	25	2	8	OP
Level 3 filter nodes	150			165	
Nodes, racks, power supplies, cables	150	10	15	165	EQ
Fibre Channel SAN	23			29	
R&D Hubs, cables	3	25	1	4	OP
Hubs, cables	20	25	5	25	EQ
RAID storage array	52			65	
FC RAID controllers	10	25	3	13	EQ
Disk crates, rack	15	25	4	19	EQ
Hot-swappable disks	18	25	5	23	EQ
Backup system	9	25	2	11	OP
DAQ HOST system	118			148	
R&D mid-range server	20	25	5	25	OP
Primary high-end server	32	25	8	40	EQ
Secondary mid-range server	20	25	5	25	EQ
R&D Fibre Channel JBOD disk	6	25	2	8	OP
Fibre Channel JBOD buffer disk	40	25	10	50	EQ
Database System	52			65	
Primary high-end server	32	25	8	40	EQ
Secondary mid-range server	20	25	5	25	EQ
File Server system	40			50	
Primary mid-range server	20	25	5	25	EQ
Secondary mid-range server	20	25	5	25	EQ
Control System	317			449	
Muon processor replacements	44	25	11	55	OP
PPC crate adapter card	6	25	2	8	OP
CAL, MUO 1553 PPC replacement	17	25	4	21	OP
CAL download 68K replacement	5	25	1	6	OP
VXWORKS target licences	10	25	3	13	OP
Rack monitors	125	50	63	188	OP
Rack monitor interfaces	50	50	25	75	OP
1553 controllers	40	50	20	60	OP
PowerPC spares	20	25	5	25	OP
Control Room systems	50			63	
(5) years annually @ \$10K	50	25	13	63	OP
Monitoring systems	50			63	
(5) years annually @ \$10K	50	25	13	63	OP
Total	891			1145	

8.4 Online Computing Summary

The need to update and replace DØ Online computing equipment is based mainly on the problems associated with the rapid aging and obsolescence of computing hardware. Maintenance costs, particularly 7x24 costs for high availability systems, rapidly approach replacement costs by systems with much greater functionality. Additionally, software support for operating systems and critical applications (ORACLE) is potentially problematic for the platforms currently in use. There is a need for higher bandwidth data logging made possible by improved Offline capabilities. There are very real benefits to be accrued from more complex trigger filters and data monitoring software. For these reasons, we plan to update and replace the Online systems.

Replacement systems, wherever possible, will be based on commodity Linux solutions. This is expected to provide the best performance at the lowest cost. The Fermilab Computing Division is expected to support Linux as a primary operating system, with full support of local products and commercial applications. We plan to follow a “one machine, one function” philosophy in organizing the structure of the Online system. In this way, less costly commodity processors can replace costly large machines.

The maintenance of the Control System will be difficult in the Run 2b era as the component parts become difficult to acquire. The plan is to build the items needed for Run 2b expansion and acquire a sufficient number of spares.

9 Summary and Conclusions

The DØ experiment has an extraordinary opportunity for discovering new physics, either through direct detection or precision measurement of SM parameters. An essential ingredient in exploiting this opportunity is a powerful and flexible trigger that will enable us to efficiently record the data samples required to perform this physics. Some of these samples, such as $p\bar{p} \rightarrow ZH \rightarrow b\bar{b} \nu\bar{\nu}$, are quite challenging to trigger on. Furthermore, the increased luminosity and higher occupancy expected in Run 2b require substantial increases in trigger rejection, since hardware constraints prevent us from increasing our L1 and L2 trigger rates. Upgrades to the present trigger are essential if we are to have confidence in our ability to meet the Run 2b physics goals.

To determine how best to meet our Run 2b trigger goals, a Run 2b Trigger Task Force was formed to study the performance of the current trigger and investigate options for upgrading the trigger. Based on the task force recommendations, we have adopted the following plan for the trigger upgrade:

1. Replacement of the Level 1 Central Track Trigger (CTT) DFEA daughter boards. The CTT is very sensitive to occupancy in the fiber tracker, leading to a large increase in the rate for fake high- p_T tracks in the Run 2b environment. The new daughter board will utilize more powerful FPGAs to implement individual fiber “singlets” in the trigger, rather than the “doublets” currently used. Preliminary studies show significant reductions in the rate of fake tracks can be achieved with this upgrade.
2. Replacement of the Level 1 calorimeter trigger. The calorimeter trigger is an essential ingredient for the majority of DØ triggers, and limitations in the current calorimeter trigger, which is essentially unchanged from the Run 1, pose a serious threat to the Run 2b physics program. The two most serious issues are the long pulse width of the trigger pickoff signals and the absence of clustering in the jet trigger. The trigger pickoff signals are significantly longer than 132 ns, jeopardizing our ability to trigger on the correct beam crossing. The lack of clustering in the jet trigger makes the trigger very sensitive to jet fluctuations, leading to a large loss in rejection for a given trigger efficiency and a very slow turn-on. Other limitations include exclusion of ICD energies, inability to impose isolation or HAD/EM requirements on EM triggers, and very limited capabilities for matching tracking and calorimeter information. The new L1 calorimeter trigger would provide:
 - A digital filter that utilizes several samplings of the trigger pickoff signals to properly assign energy deposits to the correct beam crossing.
 - Jet triggers that utilize a sliding window algorithm to cluster calorimeter energies and significantly sharpen jet energy thresholds.
 - Inclusion of ICD energy in the global energy sums to improve missing E_T resolution.

- Electron/photon triggers with options to impose isolation and/or HAD/EM requirements to improve jet rejection.
 - Topological triggers that aid in specific event topologies, such as acoplanar jets.
3. A new calorimeter-track match system. Significant improvements in rates have been demonstrated for both EM and track-based τ triggers from correlating calorimeter and tracking information. The cal-track match system utilizes boards that have already been developed for the muon-track matching system.
 4. No major changes are foreseen for the Level 1 Muon trigger. Since the muon trigger matches muon and tracking information, it will benefit indirectly from the track trigger upgrade.
 5. Some of the L2 β processors will be replaced to provide additional processing power.
 6. The L2 Silicon Track Trigger (STT) requires additional cards to accommodate the increased number of inputs coming from the Run 2b silicon tracker.
 7. Maintaining Level 3 trigger rejection as the luminosity increases will require increasing the processing power of the L3 processor farm as part of the upgrade to the online system.
 8. The online computing systems require upgrades in a number of different areas. These upgrades are largely needed to address the rapid aging and obsolescence of computing hardware. We anticipate upgrading our networking infrastructure, L3 farm processors, the online host system, control and monitoring systems, database and file servers, and the slow control system.

Simulation studies indicate that the above upgrades will provide the required rejection for the Run 2b physics program. In particular, the expected trigger rate for the primary Higgs channels, WH and ZH, is compatible with our trigger rate limitations.

The technical designs for these systems are making rapid progress. The designs are based on existing technologies, such as FPGAs and commercial processors, minimizing the technical risk. The current status of these designs is presented in this document, along with summaries of the cost and schedule.

Detailed cost and schedule documents have been prepared separately from this document. The total M&S cost for the trigger and online upgrade is given in Table 45 below. The schedules for installation of these upgrades are consistent with the expected shutdown in 2005 for the silicon tracker replacement.

Table 45. M&S cost summary for the Run 2b trigger upgrade. University labor costs that require new funding are included. Fermilab labor, as well as university labor supported by the university or base grants, is not included.

Trigger System	Cost (\$k)	Contingency	Total Cost (\$k)
Track Trigger	780	43%	1117
Calorimeter Trigger	1344	26%	1690
Calorimeter-Track Match	176	28%	225
L2βeta Processors	72	36%	98
L2 Silicon Track Trigger	402	40%	564
Online Computing (Equipment only)	397	20%	475
<i>Total</i>	<i>3171</i>	<i>32%</i>	<i>4169</i>