DD

cea

DAPNIA

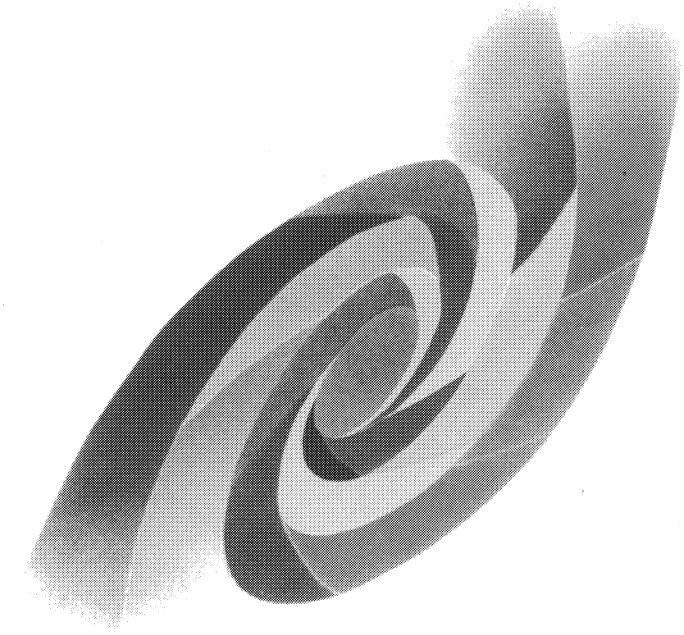

CERN LIBRARIES, GENEVA

CM-P00048628

DAPNIA-03-134 JUNE 2003

The run IIb trigger upgrade for the D0 experiment

M. Abolins, D. Calvet, P. Le Dû, E. Perez, G. Tarte et Al.

Presented at the 13th IEEE Real Time Conference (RT-03),
Montréal (Canada), 18 May-23 May, 2003

Département d'**A**strophysique, de Physique des **P**articules, de Physique **N**ucléaire et de l'**I**nstrumentation **A**ssociée

DSM/DAPNIA, CEA/Saclay F - 91191 Gif-sur-Yvette Cédex
Tél : (1) 69 08 24 02 Fax : (1) 69 08 99 89
http : //www-dapnia.cea.fr

Couverture : Sylvie Cabirit et Marc Sauvage

The Run IIb Trigger Upgrade for the DØ Experiment

M. Abolins, J. Ban, J. Bystricky, D. Calvet, D. Chapin, D. Edmunds, H. Evans, L. Han, U. Heintz, M. Hildreth, R. Hirosky, K. Johns, M. Johnson, P. Laurens, P. Le Dû, J. Mitrevski, M. Narain, J. Olsen, P. Padley, J. Parsons, R. Partridge, E. Perez, W. Sippach, J. Steinberg, G. Tarte, G. Wilson, D. Wood, S. Wu

***Abstract*—The increase in instantaneous luminosity anticipated in Run IIb of the Tevatron collider at Fermilab requires increased background rejection capabilities from the trigger system of the DØ detector. A set of upgrades is under way to improve triggering at level 1 in the calorimeter and tracker, and at level 2 in the silicon track trigger and software triggers. Reductions of up to a factor of ten on the rates of high transverse momentum triggers are anticipated with the upgrades described.**

Manuscript received May 30, 2003.

This work was supported in part by the U.S. Department of Energy, by the U.S. National Science Foundation (under Grant PHY-0216360), and by CEA in France.

J. Bystricky (telephone: +169083345, e-mail: bystri@hep.saclay.cea.fr) is with the Centre National de la Recherche Scientifique at CEA, 91191 Gif-sur-Yvette, France .

D. Calvet (telephone: +169086909, e-mail: calvet@hep.saclay.cea.fr), P. Le Dû (telephone: +169084073, e-mail: ledu@hep.saclay.cea.fr), E. Perez (telephone: +169085228, e-mail: eperez@hep.saclay.cea.fr), and G. Tarte (telephone: +169088465, e-mail: tarte@hep.saclay.cea.fr) are with the Commissariat à l'Energie Atomique, 91191 Gif-sur-Yvette, France .

J. Ban (telephone: 914-591-2866, e-mail: ban@nevis.columbia.edu), H. Evans (telephone: 914-591-2815, e-mail: evans@nevis.columbia.edu), J. Mitrevski (telephone: 212-854-6386, e-mail: jmitrevs@nevis.columbia.edu), J. Parsons (telephone: 914-591-2820, e-mail: parsons@nevis.columbia.edu), and W. Sippach (telephone: 914-591-2853, e-mail: sippach@nevis.columbia.edu) are with Columbia University, Nevis Labs, Irvington, NY, USA .

M. Abolins (telephone: 517-35-9200 ext. 2121, e-mail: abolins@pa.msu.edu), D. Edmunds (telephone: 517-355-9200 ext. 2521, e-mail: edmunds@pa.msu.edu), and P. Laurens (telephone: 517-355-9200, ext. 2522, e-mail: laurens@pa.msu.edu) are with Michigan State University, East Lansing, MI, USA.

D. Chapin (telephone: 630-840-2338, e-mail: dpchapin@fnal.gov), and R. Partridge (telephone: 401-863-2334, e-mail: partridge@hep.brown.edu) are with Brown University, Providence, RI, USA

U. Heintz (telephone 617-353-9057, e-mail: heintz@bu.edu), M. Narain (telephone 617-353-6027, e-mail: narain@bu.edu), and S. Wu (telephone 617-353-5102, e-mail: wusx@bu.edu) are with Boston University, Boston, MA, USA.

M. Hildreth (telephone 574-631-6458, e-mail: mikeh@fnal.gov) is with University of Notre Dame, South Bend, IN, USA.

R. Hirosky (telephone 434-982-5721 e-mail: hirosky@virginia.edu) is with University of Virginia, Charlottesville, VA, USA

K. Johns (telephone 520-621-6791, e-mail: johns@fnal.gov), and J. Steinberg (telephone 520-621-9740, e-mail: steinberg@physics.arizona.edu) are with University of Arizona, Tucson, AZ, USA.

M. Johnson (telephone 630-840-3168, e-mail: mjohnson@fnal.gov), and J. Olsen (telephone 630-840-2779, jamieson@fnal.gov), are with Fermi National Accelerator Laboratory, Batavia, IL, USA.

P. Padley (telephone 713-348-4703, e-mail: padley@phsyics.rice.edu) is with Rice University, Houston, TX, USA.

D. Wood (telephone 617-373-2958, e-mail: darien@neu.edu) is with Northeastern University, Boston, MA, USA.

L. Han (telephone 630-840-8321, e-mail: hanl@fnal.gov) is with University of Manchester, Manchester, UK.

G. Wilson (telephone 785-864-5231, e-mail: graham@fnal.gov) is with University of Kansas, Lawrence, KS, USA.

I. Introduction

The current run (Run IIa) of Tevatron proton-antiproton collider at Fermilab is expected to accumulate at total integrated luminosity of ~2 fb^{-1} by ~2005. The run which follows, Run IIb, will be more demanding for the DØ detector with increased instantaneous and integrated luminosity. In particular, it is essential to have a powerful and flexible trigger, because it is the trigger which dictates what physics processes can be studied properly and what is ultimately left unexplored. 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.

The DØ Run IIb Trigger Upgrade [1] is designed to meet these goals within the context of the physics program and the challenges of triggering at the high instantaneous luminosities that will be present in Run IIb. The upgrade will allow DØ to select with high efficiency the wide variety of data samples required for the Higgs search and the high-p_T physics program, while providing sufficient background rejection to meet constraints imposed by the readout electronics and DAQ system.

The DØ Run IIb Trigger Upgrade is designed for operation at a peak luminosity of $2x10^{32}$ $cm^{-2}s^{-1}$ with either 396 or 132 ns bunch spacing. The 396 ns mode of operation yields an average of 5 minimum bias interactions accompanying the high-p_T interaction and either mode requires a factor of ~2.5 increase in trigger rejection over the Run IIa design.

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

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. An important part of the L2 trigger is the Silicon Track Trigger (STT), which performs track fits using fiber and silicon tracker hits and allows triggering on displaced vertices. The Level 3 (L3) trigger is based on high-performance processors and is completely asynchronous. The L1 and L2 triggers rely on dedicated trigger data paths, while the L3 trigger utilizes the DAQ readout to collect all event data in a L3 processing node.

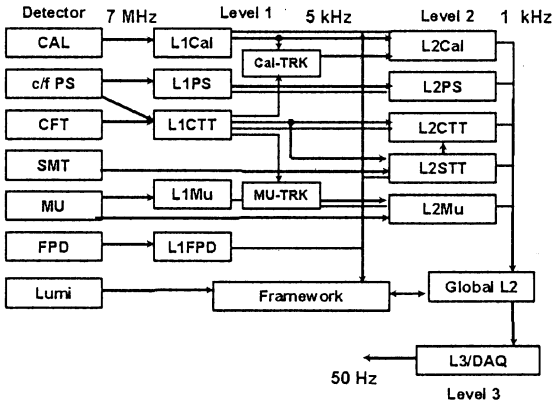

Fig. 1. The architecture of the DØ trigger system for Run IIb. The detector elements which feed the triggers are the uranium-liquid argon calorimeter (CAL), the central and forward preshower detectors (c/f PS), the central fiber tracker (CFT), the silicon microstrip detector (SMT), the muon detection system (MU), the forward proton detector (FPD) and the luminosity counters (Lumi).

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.

Given the limited time available to implement upgrades, and the desire to preserve as much as possible the working systems of the Run IIa detector, we have select five specific subsystems for targeted upgrades:

1. the L1 calorimeter trigger
2. the L1 central track trigger (CTT)
3. a new L1 system to match energy clusters in the calorimeter with tracks (Cal-TRK)
4. the L2 processors (L2 Beta)
5. the L2 silicon track trigger (STT)

II. Level 1 Calorimeter Trigger

The upgrade of the Level 1 calorimeter trigger is described in more detail in [2].

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 pseudorapidity. The calorimeter tower segmentation in $\eta x \phi$ is 0.1 x 0.1, which results in towers whose transverse size is larger than the expected sizes of electromagnetic 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 (TTs) with a segmentation of 0.2 x 0.2 in $\eta x \phi$. This yields an array that is 40 in η and 32 in ϕ or a total of 1,280 Electromagnetic (EM) and 1,280 Hadronic (H) tower energies as inputs to the L1 calorimeter trigger.

The first major element of the Level 1 Calorimeter trigger upgrade for Run IIb is the introduction of digital filters on the signals from the trigger towers. The digital filtering improves rejection of pileup effects and reduces the probability of triggering on the incorrect beam crossing.

The second major element of the upgrade is the use of algorithms relying on "sliding" windows of TTs. These can significantly improve the trigger performance, compared to the current 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 ATLAS Level-1 Trigger Technical Design Report [3]. The sliding window 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.

A block diagram of the Run IIb L1 calorimeter trigger is shown in Fig. 2. The system is divided in three major parts: 80 6u VME cards with ADCs and digital filters (ADF) followed by a system of eight Trigger algorithm boards (TAB) which execute the sliding windows algorithm and a single Global Algorithm Board (GAB) which gathers the results and computes global quantities such as missing transverse energy.

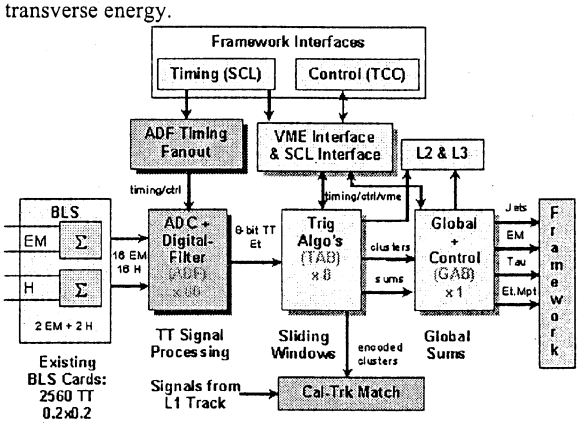

Fig. 2. Block diagram of the L1 calorimeter trigger, with the baseline subtractor (BLS) cards shown at left.

Each ADF board receives 32 channels (16 TT's of EM and H) and digitizes the signals with 10-bit ADCs (AD9218) at 30.28 MHz. The digital filter algorithm is executed in FPGAs (Xilinx XC2V500). The filtered data are transmitted

to the TAB boards over 3 LVDS channel links per ADF card. The total bandwidth for data flow from the ADF system to the TAB system is 480 Gbit/s.

In each TAB card, a set of 10 Altera Stratix FPGAs are used to execute the sliding windows algorithm. Local maxima are found by requiring the transverse energy sums in windows of 2x2 TTs to exceed those in all other 2x2 windows in a 5x5 region about the candidate tower. The 2x2 EM sum is used directly for EM objects, while for jets, the total energy is computed as the sum in the 4x4 window surrounding the maximum window. The EM algorithm also requires that there be little energy in the H cells directly behind the EM cells of the local maximum, and that there be little energy deposited in the surrounding cells. Hadronic tau decays are also found by requiring a large ratio of the transverse energy in the 2x2 window to that in the 4x4 window.

The results from each TAB card are sent to the GAB card. Here, the total counts of EM objects, jets, and tau candidates are accumulated and sent to the trigger framework, along with globally summed quantities such as the missing transverse energy and the scalar sum of transverse energy. In addition, the GAB can be programmed to require particular topological configurations of the objects found in the TABs.

III. Level 1 Central Track Trigger

The central track trigger (CTT) uses patterns of hits from the central fiber tracker (CFT) to select tracks with momenta above four pre-set thresholds. The CFT is made of scintillating fibers mounted on eight low-mass cylinders. Each of these cylinders supports four layers of fibers arranged in 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. These are referred to as Stereo Doublet layers. Only the Axial Doublet layers are incorporated into the L1 CTT. Each fiber is connected to a visible light photon counter (VLPC) that converts the light pulse to an electrical signal.

The tracking trigger hardware has three main functional elements. The first element is 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. The second hardware element is the Mixer System (MS) that receives the signals from the AFE boards and sorts them for the following stage. 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, 5 DFE Stereo (DFES) and 16 DFEF boards that handle the FPS signals. It is in the DFEA boards that the track finding algorithms are executed.

Track finding is the most difficult and expensive function of the DFEA daughterboard. To identify tracks down to 1.5 GeV, relatively large FPGAs must be used. These FPGAs match the raw data to a predefined list of track equations and serialize the found tracks to be read out at 53 MHz. The present daughter board houses five Xilinx Virtex-I chips. They are housed in ball grid array packages. The PC board requires 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 tracking trigger algorithm currently implemented is based on hits constructed from pairs of neighboring fibers, referred to as a "doublet". Fibers in doublet layers are arranged on each cylinder. In the first stage of the track finding, doublet layer hits are formed from the individual axial fiber hits. A 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.

The idea behind singlet equations is illustrated Fig. 3, which shows a fragment of three CFT doublet layers. Improving the resolution of the L1 CTT by treating CFT axial layers as singlets rather than doublet layers in the L1 trigger significantly improves the background rejection. Simulation studies show a factor of ten improvement in fake rejection rate at high-p_T by treating the hits from fibers on all axial layers as singlets.

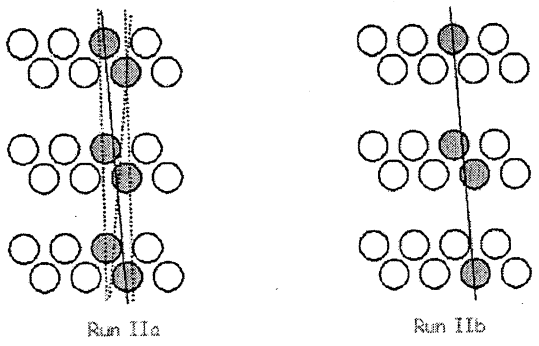

Figure 3: The doublet (left) and singlet (right) schemes for CCT triggers. The circles represent three doublet layers of axial fibers in the CTT (not to scale). The solid line represents the actual track trajectory and the dark circles show which fibers are validated in each scheme. The dashed lines show the range of possible trajectories consistent with the validated fibers for the doublet scheme.

The Run IIb upgrade of the CTT requires the replacement of the 80 DFEA daughter boards. The implementation of singlet equations requires more FPGA resources than are available in the Run IIa DFEAs. The Virtex-II series FPGAs have 8 to 10 times larger logic cells than the largest chips that we are currently using. The new design calls for four XC2V6000 chips on each daughter board. These chips come

in ball grid array packages similar in size to the existing parts. Thus, we will be able to fit four 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.

The proposed FPGA upgrade provides a major increase in the number of equations and number of terms per equation that can be handled, and provides increased flexibility in the track finding algorithms that may be implemented. Depending on the pT range, either mixtures of doublet and singlet layers or full singlet layers are proposed. We have demonstrated the technical feasibility of the upgrade by implementing the proposed algorithm in currently available FPGAs (e.g. with 4 Xilinx Virtex II XC2V6000, about 40% of the available resources would accommodate the prototype algorithm.).

IV. Level 1 Calorimeter-Track Matching

The goal of the L1CalTrack trigger is to exploit matches in the azimuthal 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 azimuth 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 azimuth 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. Specifically, we will use minimally modified muon trigger (MTCxx) cards with a new flavor board (MTFB) that performs the calorimeter-track match algorithms. The new flavor board is 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.

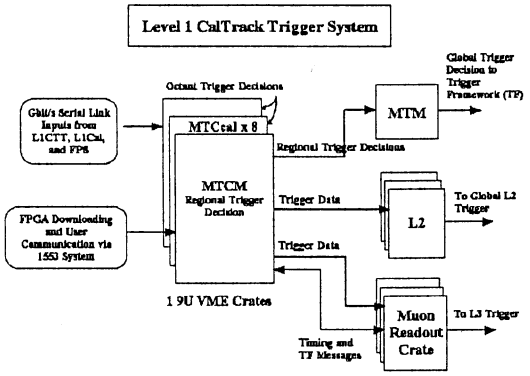

Figure 3. Block diagram of the L1 calorimeter-track matching system.

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. The four FPGA's on the present MTC05 MTFB can be replaced with one larger FPGA. In addition, the MTCcal MTFB will contain a Gbit/s serial link. This will allow the MTCM to be bypassed in forming the trigger decision that is sent to the Trigger Framework.

V. Level 2 Beta Processors

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 are included (L2Muon, L2STT). 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 L2β processors 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 currently available "off the shelf" from several vendors including Advantech-nc, VMIC, Diversified Technology Inc., and Teknor. The remaining functionality of the board is implemented in a large FPGA and Universe II VME interface mounted on a 6U-to-9U VME adapter card.

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.

For Run IIb, we plan a partial upgrade of the Level 2β system that replaces 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.

VI. Level 2 Silicon Track Trigger Upgrade

The DØ Level 2 Silicon Track Trigger (L2STT) receives the raw data from the silicon microstrip tracker (SMT) on every level 1 accept. It processes the data from the axial strips in the barrel detectors to find hits in the SMT that match tracks found by the level 1 track trigger in the CFT. It then fits a trajectory to the CFT and SMT hits. This improves the resolution in momentum and impact parameter, and the rejection of fake tracks, compared to the central track trigger alone.

The modules comprising the STT are the Fiber Road Card (FRC), which receives the data from L1CTT and fans them out to all other cards the process hits from the same sector, the Silicon Trigger Card (STC), which receives raw data from the SMT front ends and filters the hits to associate them with the L1CTT tracks, and the Track Fit Card (TFC) which fits trajectories to the L1CTT tracks and SMT hits. These modules are housed in six VME crates corresponding to six azimuthal sectors of the detector.

The Run IIa SMT consists of four concentric layers of detectors, while the detector being built for Run IIb will contain six layers. In order to take advantage of the additional layers, the silicon track trigger (STT) must be expanded to accommodate more inputs. Monte Carlo studies showed that the best results in efficiency vs. rejection were obtained with all six layers instrumented, but that instrumenting only layers 0, 1, 2, 3, and 5 could give reasonable performance at a reduced cost.

Instrumenting these five SMT layers requires one additional STC per crate. Increasing the CPU power for fitting requires two additional TFCs per crate. All of these will require motherboards. The designs for all of these already exist in the Run IIa STT. Additional optical splitters and fibers must also be purchased for the larger number of silicon input channels.

VII. Simulation of Rates

In order to evaluate the performance of the trigger upgrade, background events were processed with simulated versions of the Run IIa trigger and with the upgraded Run IIb trigger, including all three upgrades at Level 1. The background events included an average of 7.5 simulated minimum bias events to mimic the conditions at a luminosity of 2×10^{32} $cm^{-2}s^{-1}$. A list of ten core triggers was used, which was intended to trigger efficiently on the high-p_T physics processes of interest in Run IIb: $W \to \ell\nu$, $Z \to \ell\bar{\ell}$, $t\bar{t} \to \ell + \text{jets}$, $WH \to \ell\nu b\bar{b}$, $ZH \to \ell\bar{\ell}b\bar{b}$, $ZH \to \nu\bar{\nu}b\bar{b}$, $H \to \tau\bar{\tau}$, etc. With the Run IIa trigger conditions, the Level 1 accept rate for this list was about 30 kHz, which far exceeds the available 5kHz bandwidth. With the upgrades included, the simulated rate was about 3.2 kHz. The simulation did not include any special topological terms in the level 1 calorimeter trigger, so we expect the rate can be reduced further.

VIII. References

[1] DØ Collaboration, "Run IIb upgrade technical design report", Fermilab Pub-02-327-E, September, 2002.

[2] D. Calvet, "Algorithms and Architecture for the L1 Calorimeter Trigger at D0 Run IIb", these proceedings.

[3] ATLAS Trigger Performance Status Report, CERN/LHCC 98-15.