



The DØ Run 2b Upgrade

The DØ Collaboration

April 12, 2001

Prepared for the April 2001 meeting of the Fermilab PAC

Introduction

The case for maximizing the discovery potential of the Tevatron collider has never been stronger. Indirect measurements continue to point towards a light Higgs, with direct searches showing an excess for $M_H \sim 115$ GeV where the Tevatron is able to make a decisive discovery. Meanwhile, precision measurements may be showing hints of new physics beyond the Standard Model.

For over a year, an effort has been underway in DØ to identify the additional upgrades needed to maximize our sensitivity to new physics in the “Run 2b” era. We believe our most critical needs are to upgrade the silicon tracker and improve trigger rejection. For the silicon tracker, our goals are two-fold: develop a device capable of handling the 15-30 fb^{-1} expected for Run 2b, and maximize our sensitivity to Higgs decays and other high- p_T physics processes. For the trigger, we seek to increase our trigger rejection power to maintain the present trigger rates in the face of higher luminosity, and ensure efficient triggering for key discovery channels such as $p\bar{p} \rightarrow ZH \rightarrow b\bar{b}nn$.

Considerable progress has been made in developing our Run 2b upgrade plans since the fall PAC meeting. We have selected a baseline design for the silicon tracker upgrade, begun a number of detailed design studies, and developed a preliminary cost estimate and schedule. An NSF Major Research Instrumentation (MRI) proposal was submitted in early February by eight university groups seeking partial funding for the silicon tracker upgrade. In the area of trigger upgrades, we are developing plans for replacing the Level 1 calorimeter trigger and are studying various options for upgrading the tracking trigger. In the remainder of this document, we summarize the progress that has been made and describe the conceptual design of the DØ Run 2b upgrade.

Silicon Tracker for Run 2b

The current DØ silicon tracker was built to withstand the $2 - 4 \text{ fb}^{-1}$ of integrated luminosity originally projected for Run 2. The higher integrated luminosity expected in Run 2b will render the inner layers of the present tracker inoperable due to radiation damage. Of particular importance to being able to collect the data needed for the Higgs discovery in a timely manner is completion of the replacement detector in approximately three years with minimal Tevatron down time. The collaboration has carefully studied two options for a Run 2b silicon tracker upgrade: "Partial Replacement" and "Full Replacement." In the Partial Replacement option, the present tracker design is retained and the inner two silicon layers are replaced with new radiation tolerant detectors. In the Full Replacement option, the Run 2a tracker is replaced with a new device. An internal review of these two options identified significant risks with the Partial Replacement option, including the risk of damage to the components not being replaced, the long down-time required to re-fit the existing detector, an inadequate supply of the SVX2 readout chips, difficulties in adequately cooling the inner layers, and marginal radiation hardness in the layers not being replaced. Furthermore, it is very difficult in the Partial Replacement option to re-optimize the detector for the Run 2b physics program. For these reasons, DØ has decided to proceed with the Full Replacement option and build a new silicon tracker that is optimized for the Higgs search and other high- p_T physics processes.

Silicon Tracker Overview

The proposed silicon tracker will provide 6 axial and 4 stereo measurements over the region $|\eta| < 2$ with the inner layers covering at least 2 standard deviations of the interaction point distribution. The innermost layer will have a significantly reduced radius relative to the current tracker to improve the impact parameter resolution. A new outer layer at larger radius will reduce occupancy and improve pattern recognition in the higher rate environment of Run 2b. The outer layer is also important for providing stand-alone silicon tracking with acceptable momentum resolution in the region $1.7 < |\eta| < 2.0$ where DØ has good muon and electron coverage but lacks coverage in the fiber tracker.

The minimum tracker radius is determined by the Tevatron optics since one wants the detector to be in the shadow of the low-beta quadrupoles, both during injection and under normal colliding beam conditions. Even with this protection, the inner layer is expected to be exposed to $\sim 1 \text{ MRad/fb}^{-1}$. The aperture of the fiber tracker limits the radius of the outermost layer to 160 mm.

The completed and installed Run 2a tracker includes silicon disk detectors that provide forward tracking coverage and extend the effective length of the silicon tracker to accommodate the large size of the luminous region ($\sigma_z \sim 25 \text{ cm}$) during the start of Run 2a. We expect the Tevatron to operate with a crossing angle during Run 2b, reducing the size of the luminous region to $\sigma_z \sim 12\text{-}15 \text{ cm}$, obviating the need for disk detectors.

Of paramount importance to the successful construction of the new detector in the ~ 3 years available is a simple, modular design with a minimum number of part types. Figures 1 – 2 depict one possible layout of the proposed tracker. This design has 1656 silicon sensors with 7512 SVX4 readout chips for a total of 962K readout channels. For comparison, the present SMT has 793K readout channels. Table 1 shows a breakdown of the components needed per layer. Detailed engineering studies are underway to determine the optimum layout, mechanical structure, and cooling.

Table 1. Parameters of the proposed baseline design

Layer	Radius (mm) (approximate)	Axial/ Stereo	Silicon Sensor (#SVX4, length)	# Phi	#Z	# Sensors	# Chips
0	16	A	1 or 2chip, 8cm	12	8	96	144
1	35	A	2 chip, 8cm	18	8	144	288
2	62	A+S	5 chip, 12cm	12	6	144	720
3	95	A+S	5 chip, 12cm	18	8	288	1440
4	127	A+S	5 chip, 12cm	24	8	384	1920
5	160	A+S	5 chip, 12cm	30	10	600	3000
Total		6A, 4S	3 types			1656	7512

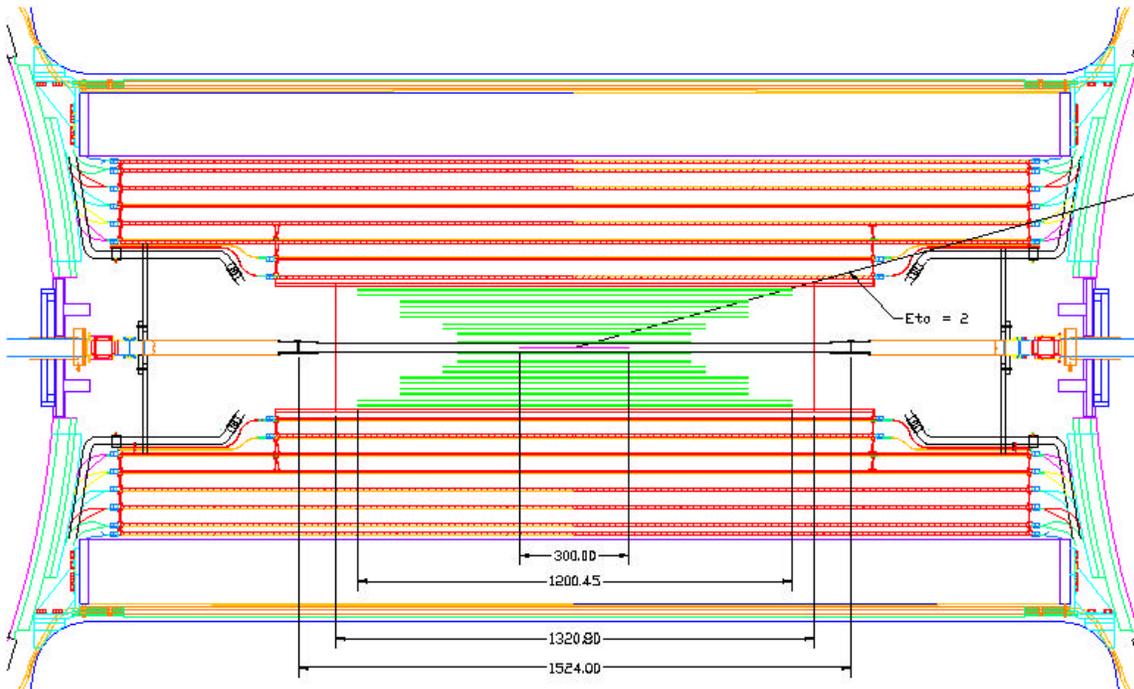


Figure 1: Side view of the proposed tracker upgrade. The 6 axial and 4 stereo silicon layers are shown in green. The $\pm 1\sigma$ size of the luminous region (± 150 mm) is shown in purple. The length of the tracker is limited to 52" (1321 mm), and the beam pipe to 60" (1524 mm), by the maximum opening between cryostats obtainable in the collision hall.

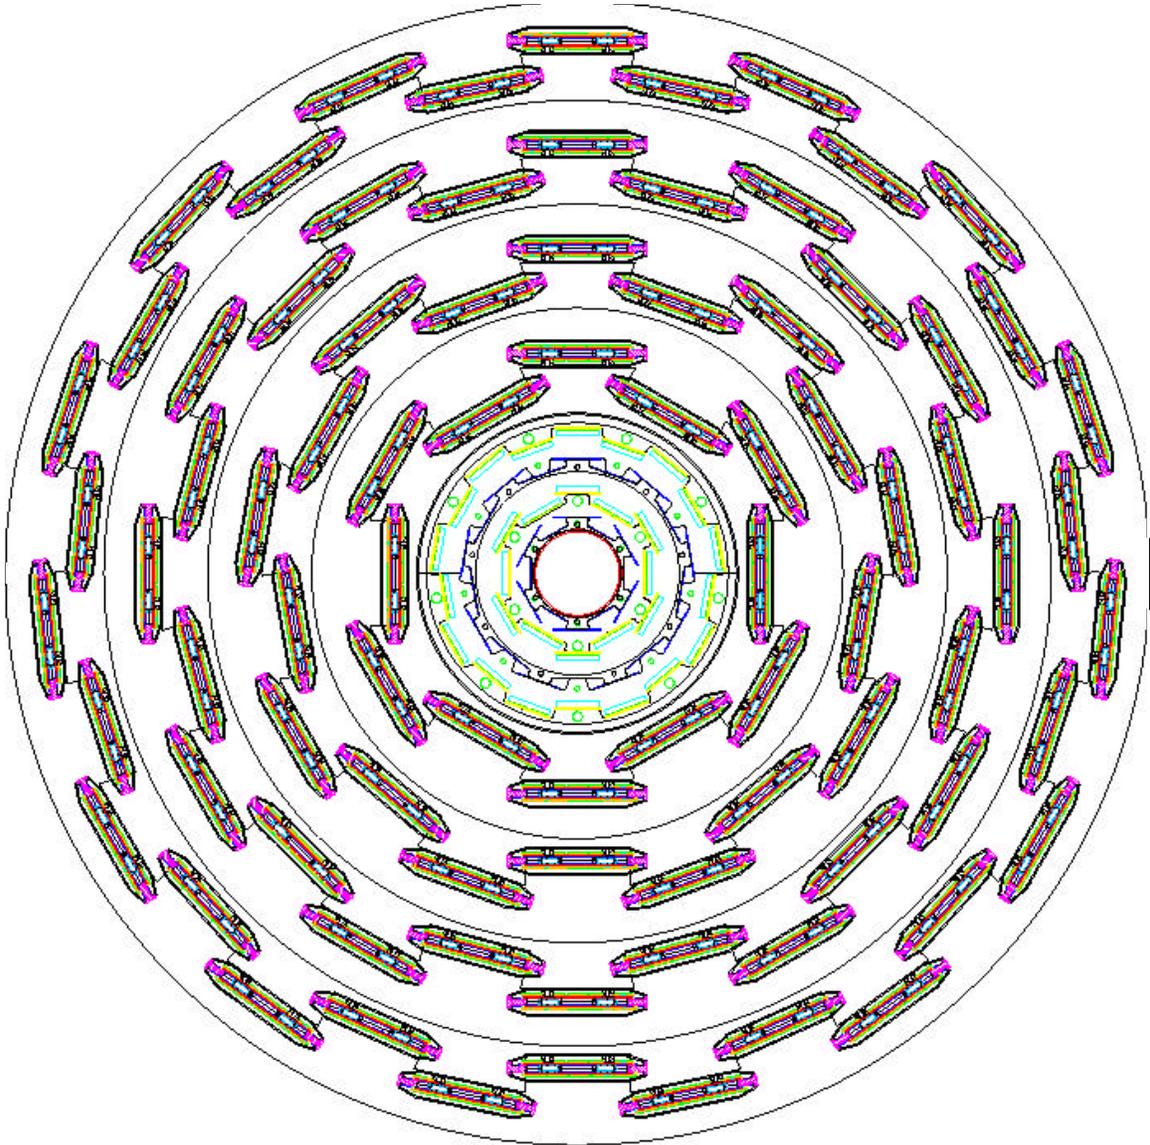


Figure 2: Axial view of the proposed tracker upgrade. The outer four layers provide both axial and stereo track measurements, while the inner two layers only provide axial track measurements.

Silicon Microstrip Sensors

The proposed design uses single-sided silicon sensors. There are several reasons for this choice. The primary reason is that double-sided sensors are significantly less radiation tolerant. In addition, the number of available vendors for double-sided sensors is small, the cost is high, and they are much more susceptible to damage during handling. The principle drawback of using single-sided sensors is the introduction of an additional $0.32\% X_0$ of material per stereo measurement; this amounts to 12-15% of the total mass per layer. There are currently two radiation hard sensor technologies being considered: low-resistivity silicon¹ and oxygenated silicon². Both technologies have been extensively tested for use in LHC detectors. These results, as well as studies of details such as guard ring structures, will be used to guide our final sensor design. Pre-

production sensors will be procured and tested in the next several months, followed shortly by production orders. As mentioned previously, our goal is to minimize the number of different part types and maximize interchangeability in order to meet the very aggressive schedule. Two narrow sensor types are required for the innermost layer to wrap closely around the beam tube. The wider of these will be used in the first layer as well. The outer layers will use a common sensor and hybrid design. We are considering two options for the stereo sensors: single metal stereo sensors and rotated axial sensors. Stereo sensors have the advantages of simpler fixturing and easier handling of subassemblies. Rotated axial sensors have the advantage of greater interchangeability of parts, but this gain may not be a great one given that the parts likely become non-interchangeable after the first production steps.

We have explored using the actual mask designs available from other experiments. We believe that we can use the present mask designs for the CDF Run 2a "layer 00" sensors that have a 50 micron readout pitch with 25 micron pitch intermediate strips. The intermediate strips are coupled to the readout strips via inter-strip capacitance and therefore the charge from them is efficiently (~90%) collected. The use of intermediate floating strips improves the single hit resolution by almost a factor of two, with only a modest increase in noise from the additional capacitance. There is some degradation of two-track resolution, but it is outweighed by the improvement in single-hit resolution. These detectors are 8 cm long, with both 1-chip and 2-chip wide sensor designs manufactured in 4" wafer technology.

The 6-inch wafer technology is now readily available for manufacturing the larger detectors needed in the outer layers. The readout pitch will be 50-70 microns, and is driven by the desired single-hit resolution, the geometrical constraints imposed by the available space, and the 6-fold symmetry required by the silicon track trigger. These sensors will be 12 cm long by ~3.7 cm wide and may or may not have intermediate strips. Two of these sensors can be fabricated on a single 6" wafer.

The sensors will undergo various levels of radiation exposure based on their radial positions. As a result, the evolution of their depletion voltages will be quite different. This necessitates a flexible high voltage system. We estimate that a bias voltage of 500V will be needed to fully deplete the sensors of the innermost two layers at the end of the run. To verify the performance of irradiated inner sensors, DØ is conducting radiation tests of sensors from Hamamatsu, SGS-Thomson, Micron, and ELMA. Due to uncertainties in the particle flux, we are designing a system that can deliver twice the voltage that we believe we will need. A separate high voltage cable will be required to deliver the bias voltage directly to the sensor. The bias voltage will be applied to the backside of the sensor, which must be capacitively coupled to the readout ground at the sensor to suppress noise. While the existing power supplies are capable of providing this voltage, the present high voltage distribution system is limited to ~200V. Therefore, a new distribution system, including cables and interface cards, will be required for the inner two layers. The current system is sufficient for the outer four layers.

SVX 4 Readout Chips

To read out the silicon detector, custom readout chips capable of handling 128 silicon strips are wired to the sensors, with wire bonds making the electrical connections to the strips and the readout chip. The fabrication process used to produce the SVX2 chips in the current detector is no longer available. In addition, the SVX2 are not able to withstand the radiation dose of ~15 Mrad expected in the inner layer. Fortunately, commercially available deep sub-micron (0.25 – 0.35 micron) processes are available

that are intrinsically radiation tolerant. Both the APV25 chip developed for CMS and the VA1 chip developed for Belle have been produced in deep sub-micron processes and tested to >20Mrad with no sign of radiation damage.

During the fall, DØ examined the various options for reading out the silicon detector, and made a decision in early December to adopt the SVX4 readout chip. The SVX4 is a joint effort by chip designers at Fermilab, LBL, and Padova to develop a new readout chip for use by both CDF and DØ in Run 2b. It is based on converting the SVX3 chip used by CDF to 0.25 micron technology. To utilize this chip in the DØ data acquisition system, the SVX4 must emulate an SVX2 chip. By building a small adapter board that performs a relatively simple re-mapping of control signals, we have demonstrated that it is possible to readout an SVX3 chip using the existing DØ DAQ system. This allows DØ and CDF to work together on a common design for the chip core, with only small differences in the metalization layer that provides the bonding pads needed to reproduce the SVX2/SVX3 footprint.

Modern chip design and simulation tools are expected to significantly aid the conversion of the SVX3 design to the 0.25 micron process. While the development time for a new readout chip is a concern, the turn-around time should be considerably faster than for the previous rad-hard submissions since we will be using an industry standard process. The 0.25 micron process requires the SVX4 be operated with a 2.5V power rail, rather than the 5V rail of the SVX2. To maintain compatibility with the existing readout electronics, level-shifting circuitry must be incorporated into the design. We anticipate using custom radiation hard transceiver chips that were previously purchased by DØ to perform the necessary level shifting of the SVX4 output signals. We anticipate that the first prototype SVX4 chips will be available for testing by late 2001 or early 2002.

Hybrid Assemblies

In the outer layers, the SVX4 readout chips will be mounted on hybrids mounted directly on, or in very close proximity, to the silicon sensors. This allows for wire bonding directly from the chips to the sensors. We are currently investigating "double-ended" hybrid designs where two sensors are read out on each readout cable (see Fig. 3). The hybrid technology is not yet finalized; options include ceramic thick-film technology, such as that used for the CDF L00 hybrids, or Kapton circuits mounted on either ceramic or beryllium substrates. Radiation length, thermal management, fabrication tolerances, cost, assembly time, and layout rules are among the issues that must be weighed in this decision.

The innermost layers require a substantially different support design due to their very small radius and close proximity to the beam tube. For these layers a low-mass analog readout cable will couple the silicon and hybrids. This allows the hybrids and silicon to be mounted independently, moving the mass and heat load of the hybrids out of the active detector volume. While the added capacitance from the flex cable degrades the signal-to-noise (S/N) ratio, we expect to achieve $S/N > 10$ for the SVX4 with flex cable readout.

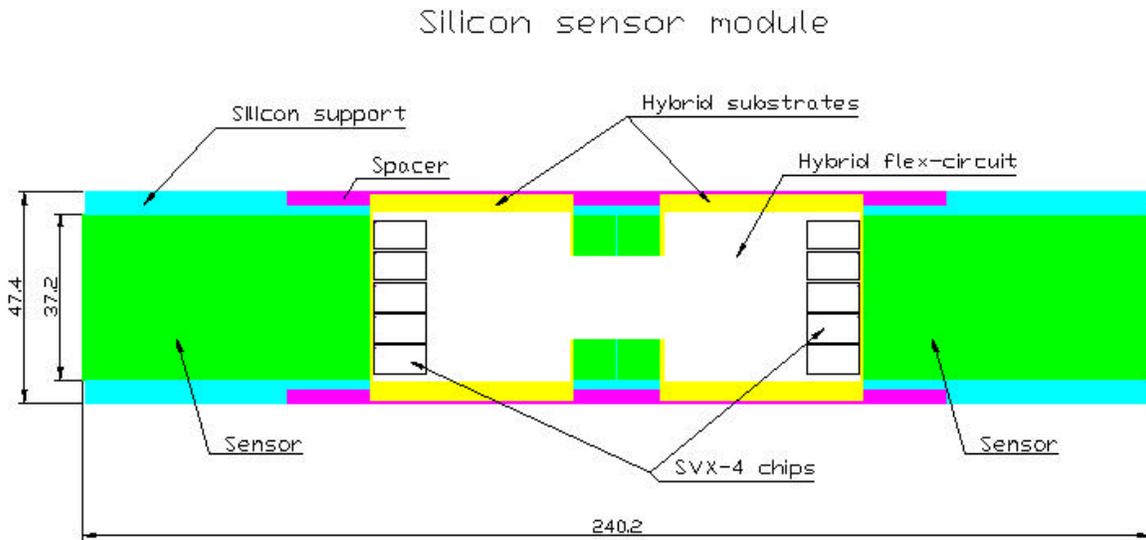


Figure 3: Double-ended hybrid design with two sensors sharing a common readout cable.

Mechanical Assemblies

The sensor/hybrid modules are mounted on support beams with integrated cooling channels. Each beam structure will have a number of sensor/hybrid modules mounted on it forming a “stave.” Readout cables are brought out the ends of the stave, and cooling also enters and exits at the end. Figure 4 shows an example of a stave design for Layer 2.

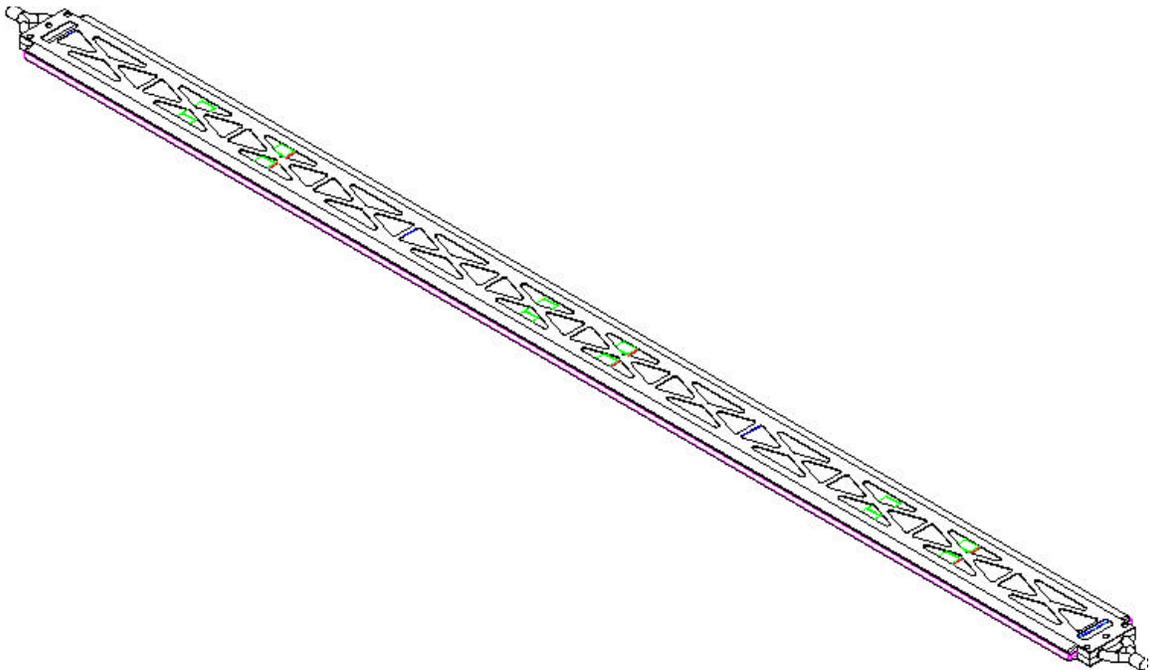
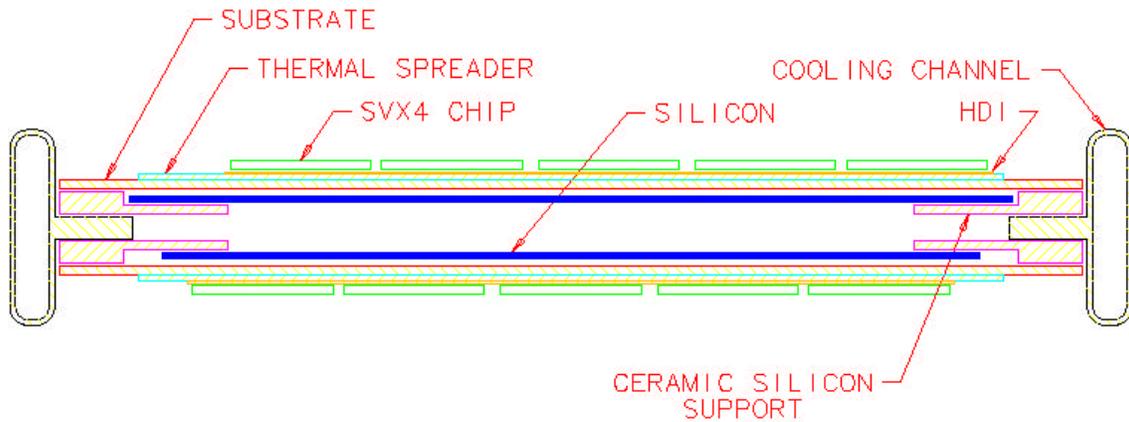


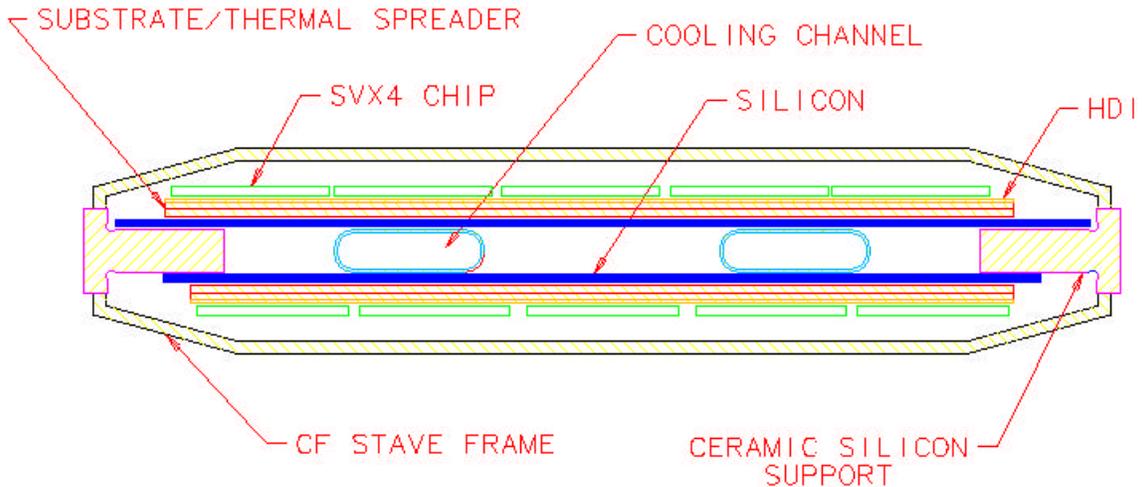
Figure 4: Example of a stave design for Layer 2 with 6 axial and 6 stereo sensors.

Several options are being considered for supporting and cooling the sensor/hybrid modules. One design alternative uses edge cooling along the length of the sensor and may allow for separate cooling of the sensors and hybrids (Fig. 5). A second alternative would place the cooling tubes between the axial and stereo sensors providing a large area of surface contact between the cooling lines and sensors (Fig. 6). The hybrids would be mounted directly to the silicon and their heat load transferred through the silicon to the cooling.



DO RUN 2B SILICON PROPOSAL "A"

Figure 5: Module design with edge-mounted cooling tubes.



DO RUN 2B SILICON PROPOSAL "B"

Figure 6: Module design with interior cooling tubes.

Both of these designs have been modeled using ANSYS finite element analysis and found to have satisfactory heat transfer characteristics, namely they each maintain the silicon sensor temperature to within 57 °C of the coolant bulk temperature. Cross sections of these structures and ANSYS results are shown in Figures 7-10.

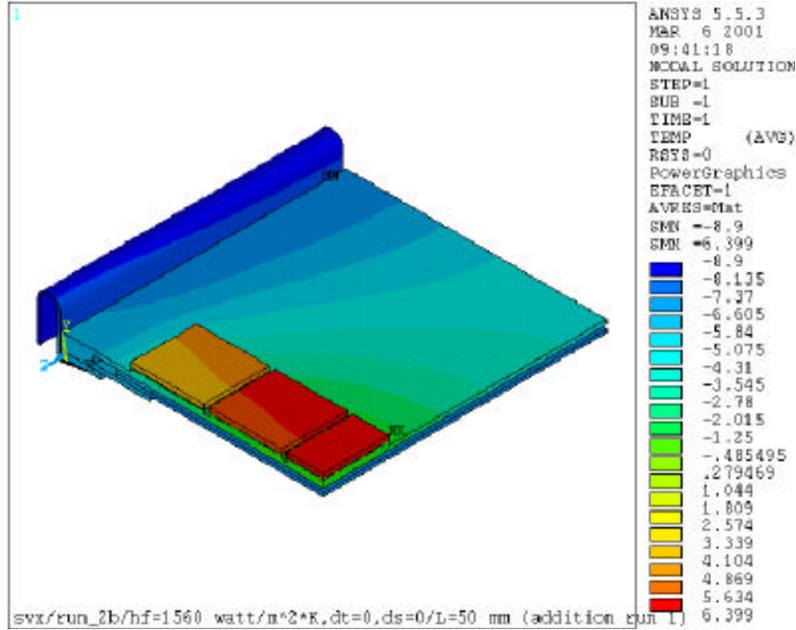


Figure 7: Finite Element calculation of thermal properties for a module with edge-mounted cooling tubes. The temperature profile is symmetric about the chip and module centerline, with only the upper left quarter of the module shown.

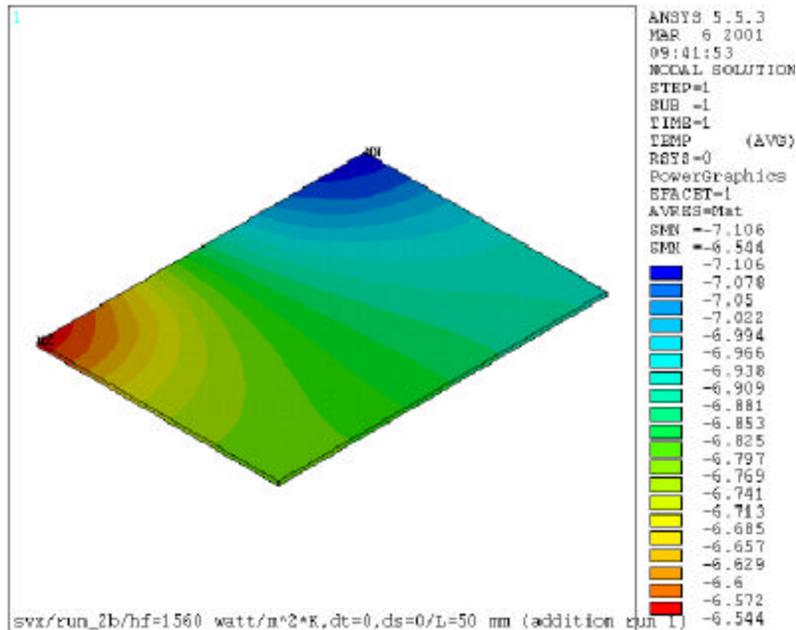


Figure 8: Same as Fig. 7 except only silicon sensor is shown.

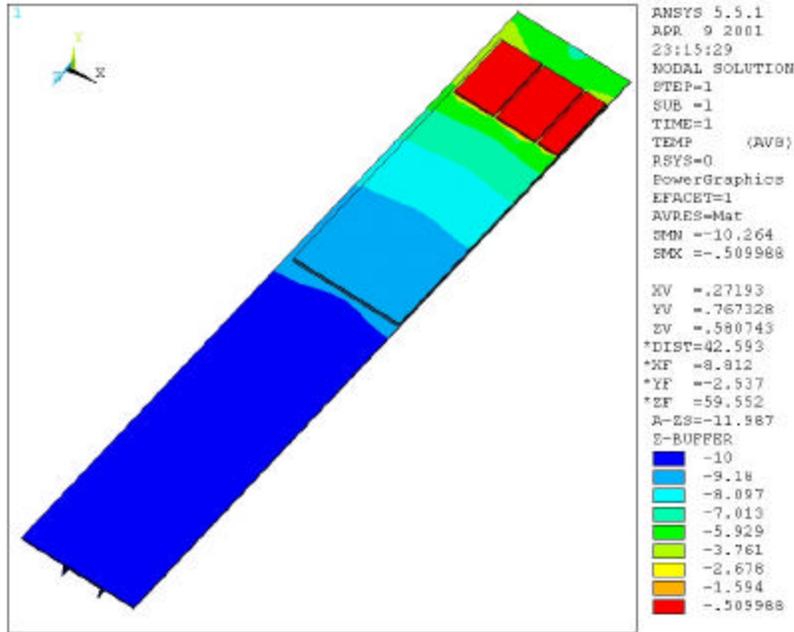


Figure 9: Finite Element calculation of thermal properties for a module with interior cooling tubes. The temperature profile is symmetric about the module centerline, with only the left half of the module shown.

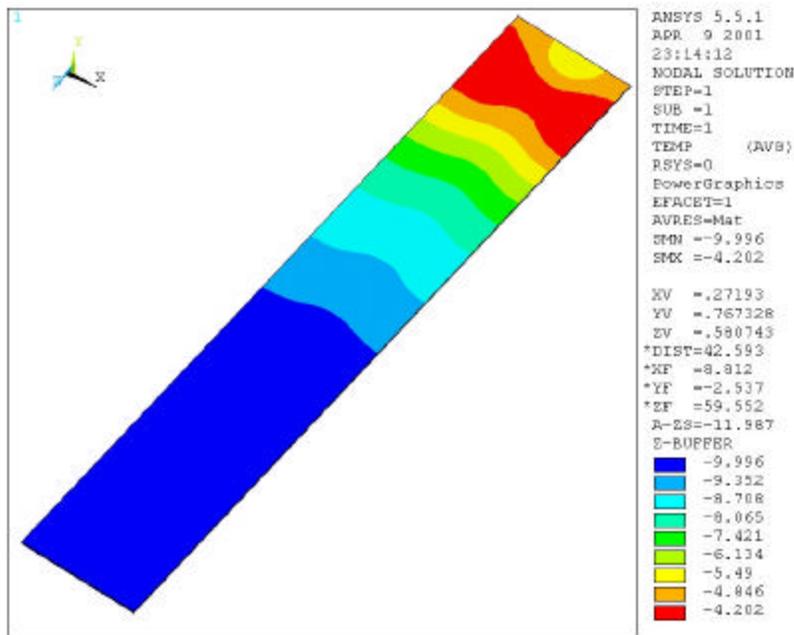


Figure 10: Same as Fig. 9 except that only the silicon sensor is shown.

A final module design choice requires a detailed understanding of available materials and fabrication methods, thermal management for the sensors and hybrids, cost, and schedule.

Staves will be mounted on the outer surfaces of thin carbon fiber reinforced epoxy (CFRE) cylinders to form barrel subassemblies. The stave mountings will need to accommodate thermal contraction of the stave cooling tubes, which may be ~ 0.8 mm over the length of the longest staves. An outer double-walled cylinder will be connected to end rings of other cylinders in order to provide the required mechanical support and stiffness. Thin CFRE membranes are likely to be needed at intermediate locations to limit deflections. Removable extensions of the outermost cylinder will support the silicon tracker from the ends of the inner barrel of the fiber tracker. Mountings of the extensions will need to be reproducible and may include provisions for remote adjustment.

In order to locate the innermost silicon layer at as small a radius as practical, we plan to follow the CDF Layer 00 design. A new beryllium beam tube will be needed. Layers 0-1 will be mounted on the beam tube in advance and then inserted through the outer silicon assembly. End connections will tie Layers 2-5 with Layers 0-1 and locate the beam tube.

Additional structures at the ends of the CFRE cylinders will interconnect cooling tubes and support low-mass cable connections. If possible, the existing low-mass cables will be used between the ends of staves and the existing adapter boards. The effect of material in this region is under study.

Installation

Each of the end calorimeters of the present detector can be rolled outward slightly more than 39" for tracker access while the end calorimeters remain on the platform; that limits the length of a replacement silicon tracker. We believe that a 52" long, 14" diameter detector can be installed in the available gap. A full-length silicon tracker with $|\eta| < 2$ coverage is a minimum of 48" long. The additional time required to de-cable an end calorimeter, remove it from the platform, replace it, and re-cable it leads to an unacceptably long down time for the installation of the new tracker. Careful design of the cable and cooling plant should allow a single-piece tracker to be installed in the available space.

Tracker Performance

A number of factors affect the tracker performance, and consequently the physics performance, of the detector. Among these factors are tracker acceptance, amount of material, resolution, and pattern recognition capabilities.

Figure 11 shows the tracker acceptance as a function of pseudorapidity, taking into account the size of the luminous region. The proposed tracker has $>50\%$ acceptance for hitting all 6 layers out to $|\eta| = 2$, significantly extending the coverage of the fiber tracker. We have also calculated the momentum resolution for silicon-only tracking, which will be most useful in the forward region where the fiber tracker doesn't provide full coverage. We find that silicon-only tracking should provide $>3\sigma$ determination of the charge for $p_T < 60$ GeV.

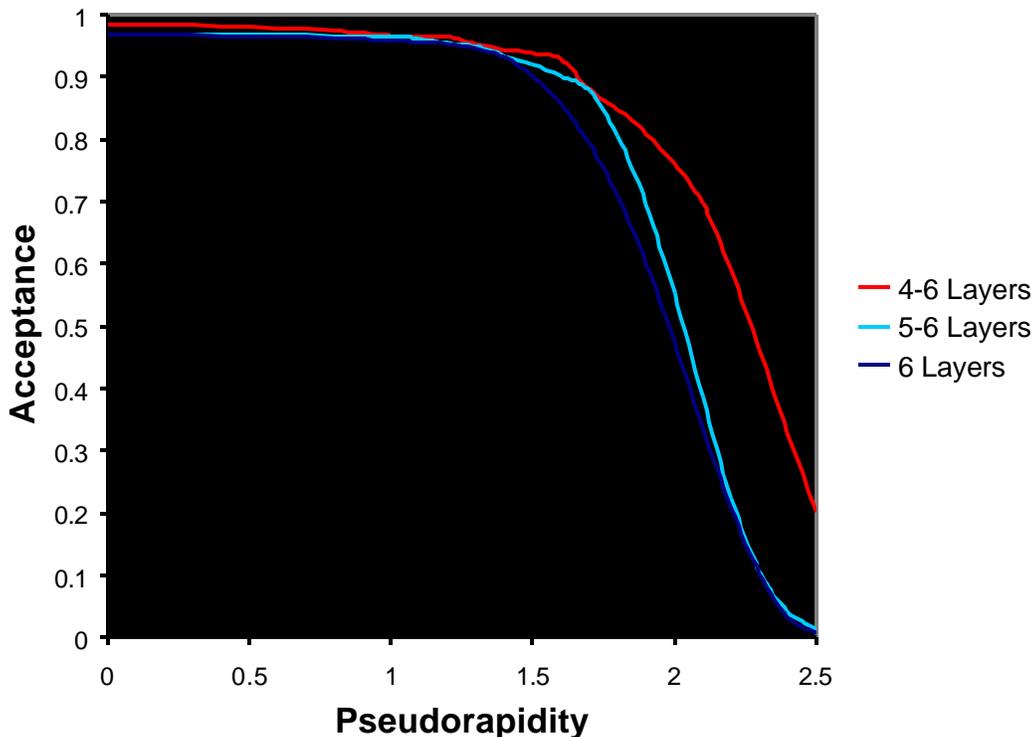


Figure 11: Tracker acceptance as a function of pseudorapidity and the minimum number of hit layers required. The size of the luminous region is taken to be $\sigma_z=15$ cm.

Rough material estimates have been made to determine the impact of multiple scattering on the tracker resolution and pattern recognition. These estimates include the silicon sensors, a model of the hybrids using ceramic substrates, support beams, coolant, and cables. The current estimate is 12-14% X_0 (averaged over the sensor area) where 3.4% X_0 is due to the silicon sensors. Additional material beyond the fiducial region for cooling manifolds, readout cables, and mechanical support will be minimized to ensure that they do not contribute significantly to occupancies nor degrade the performance of other detector subsystems.

Run 2b Simulation Studies

Full GEANT-based Monte Carlo simulations are underway to thoroughly understand the capabilities of this new Run 2b silicon tracker. These studies investigate the full pattern recognition performance of the detector with respect to tracking and the efficiency for b-tagging. The ultimate goal is a full simulation and reconstruction of Higgs decays and W+jet backgrounds with the Run 2b tracker.

Initial studies have been performed by adding an inner Layer 0 (L0) and an outer Layer 5 (L5) to the DØ GEANT-based detector simulation package. Tracks have been reconstructed using a full-fledged pattern recognition code. We have investigated track reconstruction efficiency and quality in very busy events, e.g. top pair production with a number of minimum bias events overlaid.

The addition of an inner L0 significantly improves the impact parameter resolution of the tracker, which is a crucial factor in b-tagging efficiency. This is shown in Fig. 12, which compares the impact parameter distribution of primary tracks with and without L0.

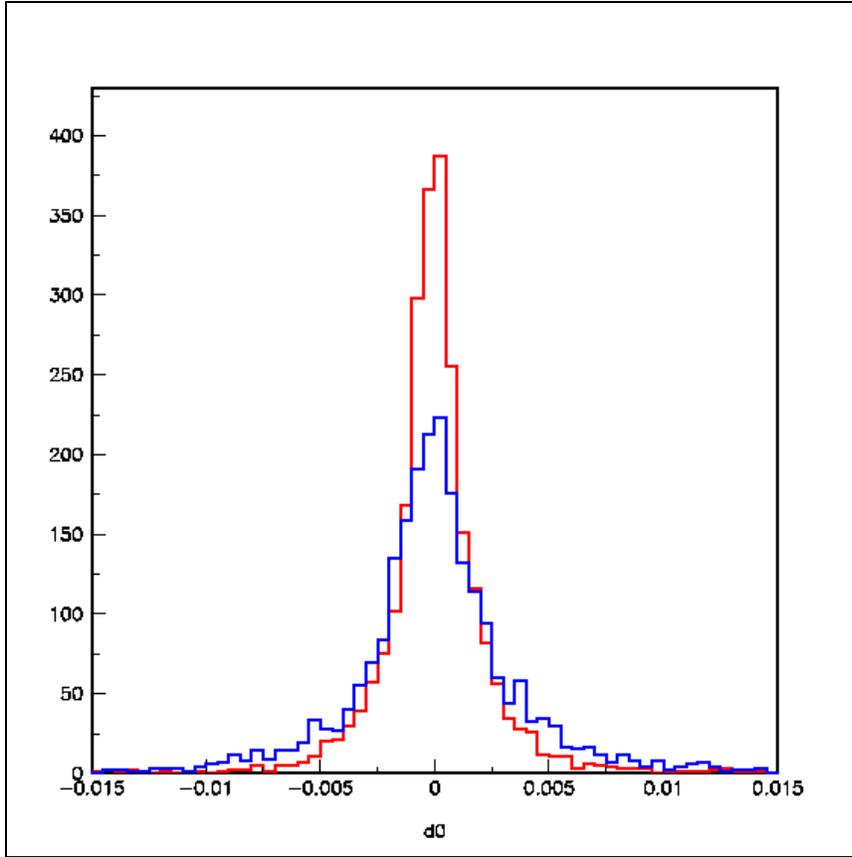


Figure 12: Impact parameter resolution (in cm) for tracks in jets with (red) and without (blue) the Layer 0 detector.

To compare the b-tagging efficiency of different tracker options, we used a very simple impact-parameter tagging algorithm that allowed us to keep the b-quark purity fixed. The results are summarized in Table 2.

Table 2: Impact parameter resolution (σ) and b-tagging efficiency (ε_b) measured for 3 tracker configurations: SMT, SMT+L0, and SMT+L5.

Option	$\sigma(d_0)$, μm	ε_b (%) $t\bar{t} + 0$ Minbias	ε_b (%) $t\bar{t} + 6$ Minbias
SMT	25	40	30
SMT+L0	15	47	35
SMT+L0+L5	15	48	40

The addition of L5 to the tracker has very little effect on the impact parameter resolution, but it plays a crucial role in improving the pattern recognition, and thus b -tagging efficiency, in very busy events. Note that the current simulation does not include the extended z -coverage of the Run 2b design and is based on a simple b -tagging algorithm. Thus, it is primarily useful for comparing the general features of the Run 2b design options, and the tagging rates which it yields should not be taken as a good estimate of the expected performance of the device.

Run 2b Trigger Upgrades

This section presents a brief overview of the DØ trigger system and the impact of the high instantaneous luminosities expected during Run 2b on the present trigger. We are currently reviewing a number of possible trigger upgrades for Run 2b, which we briefly describe below.

The DØ 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 Fig. 13). 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 μ s to complete, the equivalent of 27 beam crossings. Since ~ 10 μ s of deadtime for SVX readout is incurred following a L1 trigger, we have set a maximum L1 trigger rate of 10 kHz and would like to run well below this rate.

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 1.8 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.

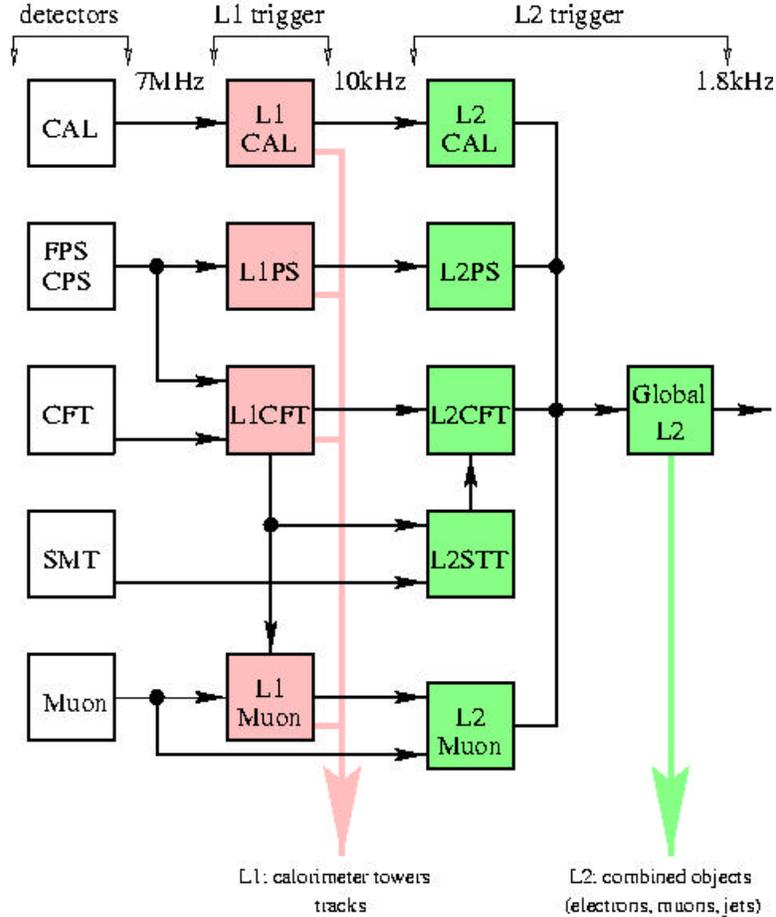


Figure 13: Block diagram of Level 1 and Level 2 triggers.

Simulation of Trigger Behavior at High Luminosity

Simulations have been performed to study the behavior of the Run 2a trigger system. Table 3 shows the expected L1 and L2 rates for a representative trigger menu that includes both high- p_T and low- p_T triggers³. Given that the maximum Level 1 rate is limited to 10 kHz by the trigger architecture and the Level 2 rate is limited to 1.8 kHz by the calorimeter readout electronics, it is clear from Table 3 that the current trigger will be compromised for Run 2b luminosities. Eliminating low- p_T triggers helps meet the rate requirements, but is not sufficient. Reductions by a factor of three or four in the Level 1 rate is needed in order to exploit the expected luminosity for Run 2b.

Table 3: Luminosity parameters and trigger rates for present Run 2 trigger, where $\langle n \rangle$ is the mean number of interactions per beam crossing. The L1 and L2 rates are in kHz.

Luminosity (cm ² /sec)	Crossing Time (ns)	$\langle n \rangle$	L1 High p_T Rate	L1 Total Rate	L2 High p_T Rate	L2 Total Rate
2×10^{32}	396	6.6	4.3	9.5	1.4	3.2
5×10^{32}	132	5.5	10.8	23.8	3.5	8

Solutions to these problems are currently being pursued along several paths. For Level 1, hardware upgrades can utilize advances in DSP and FPGA technology to increase the functionality and logic density at modest cost. A more flexible Level 1 trigger with sharper trigger thresholds will enable both a reduction in overall rate and an increase in the ability to select events of interest with less background. Since events which do not pass the Level 1 trigger are forever lost to the experiment, it is desirable to make the most informed decision possible at this time, including correlations between detectors whose signals have been reconstructed with the most sophisticated algorithms possible under the severe time constraints. The following sections describe in more detail the two major targets for improvement, the Level 1 calorimeter trigger and the Level 1 Charged Track Trigger.

The Level 1 Calorimeter Trigger

The Level 1 Calorimeter trigger (L1CAL) is used to select those events with large energy or large missing energy transverse to the beam axis. It takes as its input the raw energies from the individual calorimeter elements and computes both scalar and vector sums of the visible energy.

The L1CAL divides the calorimeter up into 1280 projective non-overlapping trigger towers for each of the electromagnetic (EM) and Hadronic (H) longitudinal layers. A trigger tower is comprised of four readout towers, each of which subtends a solid angle of 0.1×0.1 in $\Delta\phi \times \Delta\eta$. The electromagnetic trigger towers include a sum over 4 longitudinal readout segments in depth; the hadronic trigger towers are summed over 3 longitudinal segments. The input for each element in the trigger tower sum is provided by a fast trigger pick-off in the preamplifier shaping circuit of each readout tower, and the sums of the energies in each of the trigger towers is performed locally.

The inputs to the Level 1 trigger decision are formed by vector or scalar sums over the energies of the individual trigger towers. For example, the total transverse electromagnetic energy EM_T is given by $EM_T = \sum EM_T^i$, where the sum runs over all towers above threshold. Hadronic and Electromagnetic sums are combined to form total transverse energy and total missing transverse energy; each of the individual sums is also available for comparison to thresholds in order to form a Level 1 trigger decision. No clustering or jet-finding is currently performed at Level 1.

The Level 1 Jet Trigger essentially requires that the transverse energy deposited in a 0.2×0.2 electromagnetic and/or hadronic calorimeter trigger tower exceed one of several programmable thresholds. Typically, however, the majority of the energy in a jet is deposited over several contiguous trigger towers. The result of using a count of individual towers-over-threshold is a very slow turn-on in trigger efficiency as a function of the jet transverse energy. Thus, jet thresholds must be set much lower than the desired threshold in order to achieve full efficiency.

Simulations of the Run 2 environment show that 80% of the Level 1 trigger rate is due to the Level 1 Calorimeter Jet Trigger; electron and di-electron candidates have a relatively low rate due to the required CFT-CPS-Calorimeter matching. The main cause of such high rates is the poor energy resolution and lack of clustering in the current Level 1 calorimeter trigger. Given that L1CAL dominates the overall Level 1 trigger rate, upgrading the calorimeter trigger is likely to be an essential component of the Run 2b trigger upgrade.

One method to refine and sharpen the jet turn-on curve and decrease the rate from low energy jets is to group adjacent trigger towers in order to maximize the energy

collected and assigned to a single object. The technique of clustering towers using the method of *sliding windows* was widely studied and developed for the Level 1 Calorimeter triggers for the LHC experiments⁴. Even a modest increase in the energy threshold will have a huge effect on the trigger rate due to the steeply-falling rate of jet production as a function of transverse energy. The same clustering technique can be used to select electrons and hadronic τ 's with additional isolation criteria and track matching with the Level 1 track trigger.

Work has begun to design a replacement for the current calorimeter trigger. To achieve the greatest possible improvement in L1CAL performance, the size of the electromagnetic trigger tower will be reduced to 0.1×0.2 in $\Delta\phi \times \Delta\eta$, and the entire trigger logic will be replaced by an FPGA-based system to perform the clustering and process the other algorithms.

Each Trigger tower signal will be digitized using a dedicated 8 bit ADC followed by a Digital Signal Processor performing a Digital Filter to extract the transverse energy for each beam crossing. The result is a set of 1280 0.2×0.2 basic trigger towers (TT) with two electromagnetic energies and the corresponding hadronic energy. The full treatment of these 1280 TT can be achieved using 160 standard size (6U) Euro VME cards, each one treating 8 TTs. A schematic diagram of the proposed system is shown in Fig. 14.

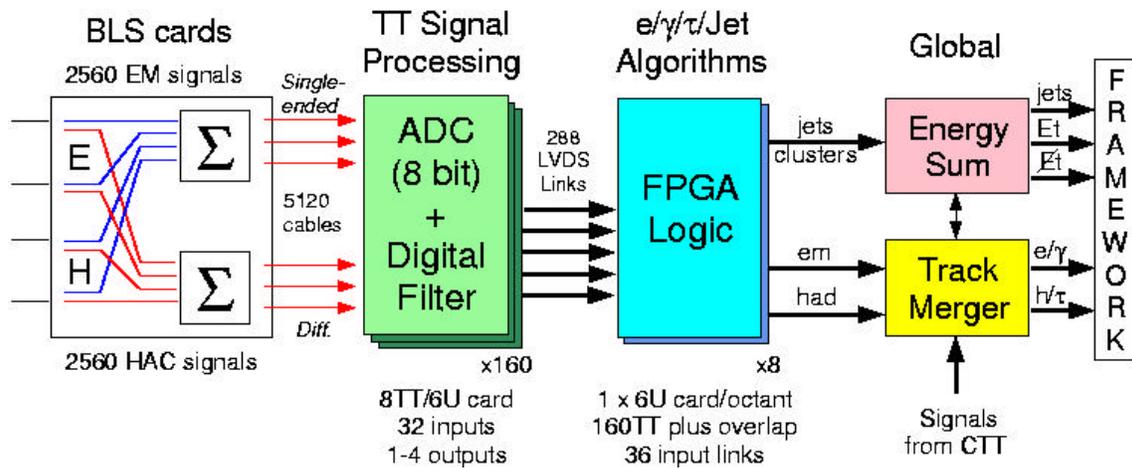


Figure 14: A schematic overview of the proposed L1CAL upgrade, showing the signal paths and logical elements. The current $D\phi$ BLS cards provide sums over 0.2×0.2 trigger towers; they will be modified so that the EM sums are done over a smaller tower size.

The subsequent calorimeter trigger is fully digital and is divided into two logical parts. One part performs a search for high p_T electrons/photons/taus using the increased granularity of the electromagnetic calorimeter. The other part searches for high- E_T jets and calculates the missing- E_T and total scalar- E_T values.

For the electron/photon and the hadron/tau trigger, there are eight sets of thresholds that can be programmed independently. Each set consists of a threshold of the E_T of the cluster, an isolation threshold on the surrounding E_T in the electromagnetic calorimeter and a hadron-veto threshold on the E_T in the associated hadron calorimeter tower. For both algorithms, the optimization of the window size, isolation area, and sliding step is underway using simulation tools.

For the jet trigger, there are eight programmable thresholds that are compared to the E_T in a "jet window." The jet window size is programmable from 0.4×0.4 to 0.8×0.8 in steps of 0.2. Communication between neighboring boards is done using LVDS links.

Finally, a summation is performed over the trigger towers to calculate the missing- E_T vector and the total scalar energy for the event. Comparisons are made with eight threshold values for the missing E_T and four for the scalar E_T . In addition, matching with the CTT track trigger can be done in order to improve the purity of electron/tau clusters.

The Level 1 Central Track Trigger

The purpose of the Level 1 Central Track Trigger (L1CTT) is to find all charged tracks above a pre-defined transverse momentum (p_T) threshold. These tracks are passed to the L1 Muon system for muon identification, and are matched to the Central Preshower (CPS) detector for electron identification. The L1CTT takes as its input the discriminator signals from the fiber readout front-end boards; the track-finding is done in re-programmable FPGAs.

The discriminator signals from each of the axial Analog Front End (AFE) boards are sorted into 80 azimuthal trigger sectors, each of which subtends 4.5 degrees. The hits from a single sector and the adjacent sectors are processed, looking for hit patterns that are consistent with charged tracks originating from the interaction point. The track finding is based on pre-calculated patterns or "roads" in the axial tracker layers. Each road corresponds to a unique azimuthal (ϕ) position and p_T range. Each of the 80 trigger sectors contains approximately 11,000 roads, which are processed by FPGAs. Track finding takes place in less than 500 ns, at which time the six highest p_T tracks per sector are transmitted to the L1 Muon Trigger to be matched to muon detector segments. A list of all found tracks is also organized for later transmission to the L1 Trigger Manager (L1CTTM) and the STT. An additional step in the L1CTT processing is to match the found tracks to axial clusters in the CPS, forming embryonic electron candidates. The ensemble of found tracks, their cluster matches, and a measure of their isolation is processed by the L1CTTM and trigger terms are passed to the Level 1 trigger framework where the global Level 1 trigger decision is made.

We have performed simulations to examine the behavior of the L1CTT rate as the number of interactions per beam crossing (and hence the CFT occupancy) increases. It is estimated⁵ that the rate of found tracks in the L1CTT reaches 35 kHz at a luminosity of 2×10^{32} cm⁻²/sec with 132 ns bunch crossings, and that it increases nearly exponentially with occupancy after this point. At 5×10^{32} cm⁻²/sec, the L1CTT rate will exceed 100 kHz. The overall L1 rate is lower due to the required coincidence of calorimeter and track information.

Concerns about the sensitivity of the trigger performance to radiation damage have also been raised. As the CFT fibers are irradiated, their light output drops. While it is not expected that the light yield will become sufficiently low to require replacement of the CFT, it may be the case that fluctuations in photon statistics will sometimes result in a signal too small to pass the discriminator threshold. If this occurs, relaxing the trigger requirement from 8 out of 8 axial layers on a found track to 7 out of 8 will be required in order to maintain full efficiency. Simulations show that if this change is enacted, the number of fake tracks increases by nearly another factor of ten⁵.

Fake tracks also pose problems for the Silicon Track Trigger (STT). All L1CTT tracks must be processed through the STT. Firstly, many fake tracks make this step very time-consuming, potentially leading to deadtime if the processors cannot keep up.

Secondly, the purpose of the STT is to select b -quark events on the basis of lifetime information. Large numbers of fake tracks can result in high mistag rates.

Currently, the individual hit fibers in the CFT are converted within the track-finding logic into hit "doublets." A doublet is a pair of fibers spanning the inner and outer layer of fibers in each CFT super-layer. A doublet is formed whether or not the second in the pair of fibers is hit. This method was chosen in order to be less sensitive to fiber inefficiencies and to lower the number of possible track roads in each trigger sector. Unfortunately, the formation of doublets also degrades the position information inherent in the geometry of the CFT construction. By combining two fibers into a hit doublet, whether or not both fibers are hit, the granularity of the track trigger is reduced by almost a third, leading to more fake tracks formed from doublet hits which are not from actual tracks. At a luminosity of 5×10^{32} cm⁻²/sec, the single-fiber occupancy of the inner layers will be approximately 20%. Increasing the effective width of each fiber in this environment is much more likely to render these layers as "hit" as far as the trigger is concerned, drastically reducing the rejection power against fake tracks.

A simple way to circumvent this problem is to build a "drop-in" replacement for the current AFE that allows track-finding equations for single-fiber track roads. Such an upgrade is made possible because the combination of hits into doublets is currently done on the track-finding boards themselves; all of the individual channel signals are available for additional processing. Switching to single-fiber roads improves the local position resolution by up to a factor of two and also reduces the effective occupancy from background hits, leading to fewer fake tracks. Calculations have shown a substantial decrease in the rate of fake tracks found when compared with the standard trigger, as can be seen in Figure 15.

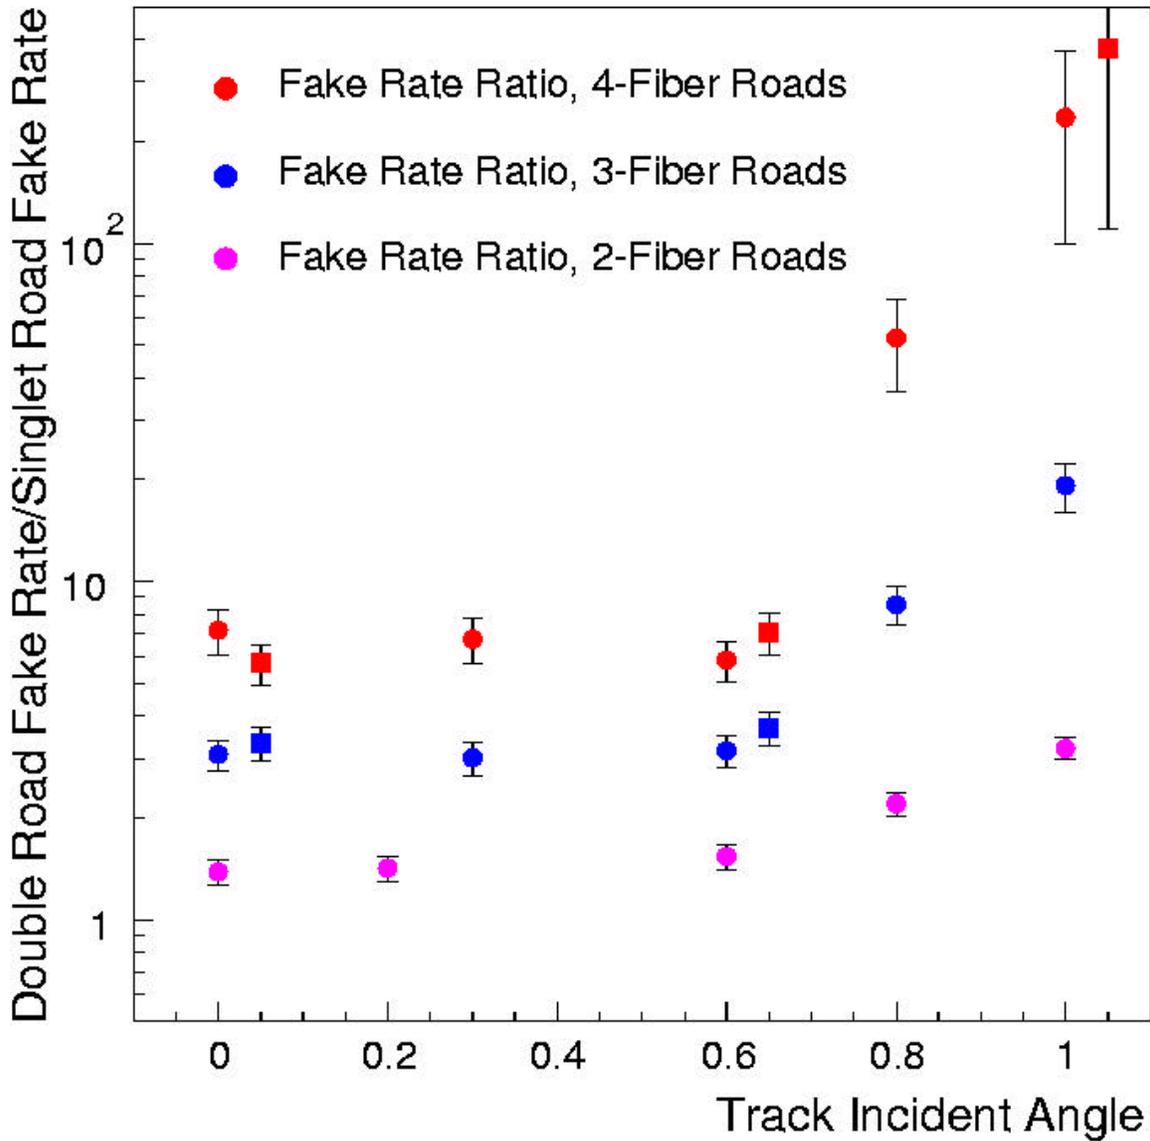


Figure 15: Results of a calculation comparing the fake rates for found track segments in two layers of the CFT. Each track type is defined by how many fibers should be hit in each of the two superlayers, for a maximum of four fibers. Tracks formed from single fiber hits require an exact geometric match between hit fibers and the track road. Tracks formed from doublet hits (the L1CTT default) only require that a track trigger doublet (one or two hit fibers) lies within the road. A large reduction in fake rate is seen when the single-fiber granularity is enforced.

Another strategy to reduce the L1CTT fake rate is to implement a second fast digital processor which can add the hits in the stereo layers of the CFT to the L1CTT. This has the additional advantage that three-dimensional tracks will be available at Level 1 for matching to calorimeter hits. We will refer to this processor as the Level 1 Stereo Track Processor, or L1STP.

The overall strategy for adding stereo hit information to the axial tracks is predicated on having the axial track parameters (ϕ and p_T) in order to make the problem tractable.

Then, a pattern of 7 or 8 stereo hits can be matched to the pre-determined trajectory of the axial track.

Due to the stereo angle, each axial trigger sector crosses 6 stereo sectors, making the number of possible track-hit combinations very large. In addition, since such a large fraction of the CFT stereo hits will need to be compared to an arbitrary axial track, much of the stereo hit information will need to be available on each L1STP track-finding board. The data-transfer rate requirements are severe, but tractable, especially since the stereo information can be transferred while waiting for the axial track parameters to arrive from the current axial track-finding boards. Large buffers and substantial on-board data buses will be required.

This upgrade would require small modifications to the transition boards at the back end of the current L1CTT system so that the axial track candidates can also be shipped to the new stereo track processor. The inputs to the L2STT would also be modified so that the L2 can take full advantage of the refined track information. Algorithms for rapid pattern matching between a given axial track and its complement of stereo hits are under development.

To estimate the advantages of the L1STP system, simulations of the stereo track finding were run on generic high-occupancy events including minimum bias and di-jet samples. The first studies of stereo trigger performance⁶ showed factors of 10-12 reduction in the L1CTT rate from adding the stereo information to minimum bias events containing 6 interactions. The additional rejection of di-jet events was approximately a factor of 5 over the axial-only track trigger. Recent simulations have achieved similar results. Since there are uncertainties in the simulation of the underlying event and minimum-bias interactions, it is important to verify these results with algorithms applied to collider data.

Level 2 and Level 3 Upgrades

While we can envision substantial changes to the architecture of the Level 1 trigger system, modifications to Levels 2 and 3 are more likely to be evolutionary. We will be able to take advantage of future increases in processor speed in order to upgrade the capacity and sophistication of the event-processing done at the higher trigger levels. A general trend in these upgrades will be the migration of more sophisticated algorithms downward in the trigger hierarchy: many of the calculations formerly done at Level 2 will be done in hardware at Level 1. The one exception is the processing of silicon tracker hits by the STT, which will play a critical role in Run 2b. As processing power increases, the Level 3 nodes should be able to apply better and better approximations of the full offline algorithms to the reconstructed events. For these reasons, processor upgrades for both Level 2 and Level 3 should be considered an integral part of the Run 2b trigger upgrade.

References

- ¹ RD20 Collaboration Technical Reports, http://albert.hep.ph.ic.ac.uk/RD20/RD20_Home.html.
- ² RD48 Collaboration Technical Reports, <http://rd48.web.cern.ch/RD48/>.
- ³ J. Blazey, "Triggering at High Luminosity," DØ Run 2b Workshop, FNAL, February 28, 2000.
- ⁴ The ATLAS Collaboration, "Level 1 Trigger Technical Design Report," ATLAS TDR-12 (1998).
- ⁵ F. Borcharding, "Tracking Triggers at High Luminosity," DØ Run 2b Workshop, FNAL, February 28, 2000.
- ⁶ P. Bloom, S. Glenn, C. Klopfenstein, R. Lander, S. Mani, "UV Trigger Processor for the DÆ Upgrade," DØ Note 2970 (1996); C. Klopfenstein and S. Mani, "Deadtimeless Level 1 UV Trigger for the DÆ Upgrade," DØ Note 3054 (1996).