CDF/DOC/BOTTOM/CDFR/7042
Version 1.2
June 4, 2004

Report of the B_s Task Force

F. Bedeschi, M. Herndon, J. Kroll, T. Liss, K. Pitts, M. Shochet, R. Snider, W. Trischuk

Contents

1 Introduction

The B_s task force was formed in December 2003 to address technical issues that are limiting the progress on the B_s^0 mixing analysis. The charge to the task force is as follows:

`Form a task group that includes broad expertise from the collaboration in the following areas`

- `Silicon tracking`
- `Layer 00 software and hardware`
- `Time-of-flight hardware and software`
- `Trigger`
- `Flavor tagging`
- B^0 `and` B_s^0 `mixing analysis`

`Your goal is to expedite and optimize the elements of the` B_s^0 `mixing analysis so that the collaboration can report a meaningful measurement on as short a timescale as possible. You should plan to optimize and if possible increase the` B_s^0 `yield, improve the proper time resolution and increase tagging power. Involvement of the` B `group conveners is necessary as well as people working on the analysis.`

`We request that you present regular updates at the weekly CDF meeting, and a document that contains the achieved proper time resolution, the` B_s^0 `yield, and the achieved tagging` ϵD^2 `values. This report should be completed by May 1, 2004`

The task force membership consists of the authors of this document plus (ex officio) the B conveners Christoph Paus and Stefano Giagu and the conveners of the B Physics Analysis Kernel (BPAK) group Ivan Furic and Steve Nahn.

The committee met weekly on Fridays for invited presentations and discussion as well as following relevant presentations in the B group and its subgroups. This document is the committee's final report.

1.1 Overview of Measuring B_s^0 Flavor Oscillations

To set the stage for what follows, we begin with a very short introduction to B_s^0 mixing.

B_s^0 mixing occurs, like B^0 mixing, through the standard model box diagrams shown in Figure 1. The amplitude is proportional to the vertex CKM element V_{ts} (the light quarks also contribute, but the matrix element is proportional to the quark mass so it is dominated by the top quark).

The measurement of B_s^0 flavor oscillations requires the following:

1. Reconstruction of the proper decay time of the B_s^0. Layer 00 provides an important improvement here.

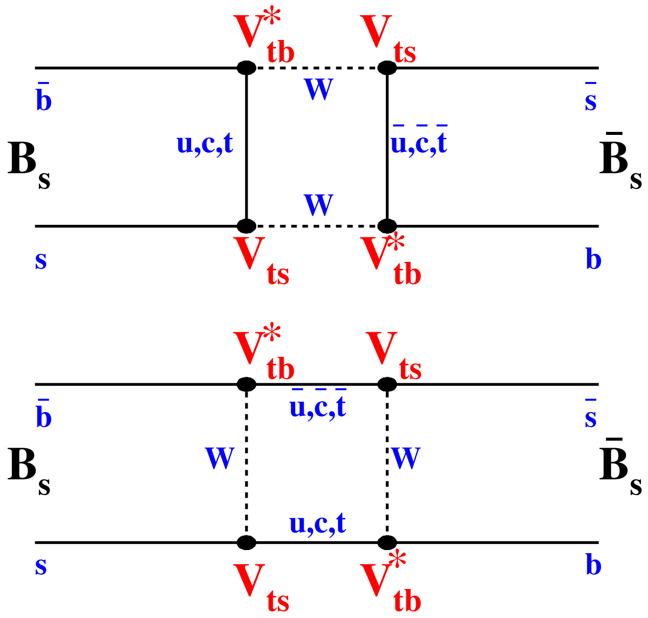

Figure 1: Feynman diagrams for B_s^0 mixing.

2. Determining the flavor of the B_s^0 at production, *i.e.*, did it contain a b or a $\bar{b}$ quark. TOF plays a crucial role here.

3. Determining the flavor of the B_s^0 at decay. For the decay modes considered in this Report, the final state identifies the flavor at decay unambiguously.

The signature for flavor oscillations is a time-dependent asymmetry given by

$$\frac{N_{\text{unmixed}}(t) - N_{\text{mix}}(t)}{N_{\text{unmixed}}(t) + N_{\text{mix}}(t)} = D \cdot \cos\left(\Delta m_s t\right), \tag{1}$$

where $N_{\text{unmixed}}(t)$ ($N_{\text{mix}}(t)$) are the number of decays at proper decay time t in which the flavor at production is the same (opposite) the flavor at decay, Δm_s is the mass difference between the two mass eigenstates (Δm_s is related to the mixing parameter x_s by $x_s = \Delta m_s \tau(B_s^0)$), and D is the dilution of the flavor tag, which is related to the probability $\mathcal{P}$ that the flavor tag at production is correct: $D = 2\mathcal{P} - 1$. The proper decay time t is not measured with infinite precision, and this introduces further dampening of the amplitude of the cosine. Of course measuring the flavor oscillations is predicated on the ability to extract a B_s^0 signal with adequate statistics and signal-to-noise. In Run II we have extracted clean signals in both semileptonic decay modes such as $B_s^0 \rightarrow D_s^- \ell^+ \nu_\ell X$ and in fully hadronic decay modes such as $B_s^0 \rightarrow D_s^- \pi^+$.

The proper time t is determined from the measured decay length L and the momentum p:

$$t = m(B) \cdot \frac{L}{p}, \tag{2}$$

where $m(B)$ is the mass of the B hadron. The uncertainty δt on t is

$$(\delta t)^2 = \left(\frac{m(B)}{p}\delta L\right)^2 \oplus \left(t \cdot \frac{\delta p}{p}\right)^2, \tag{3}$$

where δL and δp are the uncertainties on the decay length and momentum, respectively. The semileptonic decays are only partially reconstructed. This means not all of the B momentum is measured, and this degrades the proper time resolution at larger t according to the second term in Equation 3. This limits the sensitivity to B_s^0 flavor oscillations in the semileptonic modes to smaller values of Δm_s. The fully reconstructed hadronic modes do not suffer from this particular limitation as the uncertainty in the proper decay time due to the momentum reconstruction is unimportant.

The current world limit for B_s^0 mixing is 14.9 ps^{-1} [1] (for comparison, the measured frequency in B^0 mixing is about 0.5 ps^{-1}). In Run I CDF published a B_s^0 mixing limit of $\Delta m_s >$5.8 ps^{-1} based on 110 pb^{-1} of semileptonic $B_s^0 \rightarrow \phi \ell X \nu$ decays. The world average limit of 14.9 ps^{-1} includes the limits of many experiments, including CDF. The best single experiment limit is currently $\Delta m_s >$10.9 ps^{-1} [2]. The standard model expected range for Δm_s is between the limit and 24 ps^{-1}, with a favored value of 17-18 ps^{-1}.

The factors that determine the Δm_s reach for B_s^0 mixing are the proper time resolution σ_{ct}, the tagging power ϵD^2 where ϵ is the flavor tagging efficiency and D is the dilution, the yield of reconstructed B_s^0 events and the signal to background (S/B). The significance of a mixing measurement can be written as

$$\sigma(\Delta m_s) = \sqrt{\frac{S\epsilon D^2}{2}} e^{-(\Delta m_s \sigma_{ct})^2/2} \sqrt{\frac{S}{S+B}} \tag{4}$$

In the "Yellow Book" [3], CDF claimed to be able to exceed the world limit with 200 pb^{-1} of Run II data. Reality has proved otherwise. Of the factors appearing in the equation above, only σ_{ct} is relatively close to what was claimed in the Yellow Book (67 fs has been achieved vs. 45 fs in the Yellow Book, but the 67 fs does not yet include improvements from Layer 00). The achieved tagging power is, as of this writing, *significantly* below what it was claimed we would achieve, and the achieved yields are lower as well. There is some understanding of the lower yields (see [4]), but there is no understanding of how to achieve the Yellow Book projected tagging power. The committee decided not to invest time in a detailed understanding of Yellow Book claims vs. reality, but rather to focus on the current capabilities and the paths to improvements.

This document is organized as follows: In Sections 2 and 4 we describe the current performance of the trigger and understanding of B_s^0 yields. Section 2 also includes a description of planned improvements, their timescales and expected impact on the yields. Section 3 describes the current performance of silicon tracking, including layer 00 (L00). Here we discuss recent improvements in the silicon stand-alone tracking efficiency, the impact parameter resolution improvements that are achieved with layer 00 hits attached and the efforts to improve the efficiency of attaching layer 00 hits to silicon tracks. Section 5 describes the current performance of the TOF detector and the COT dE/dX. We describe the algorithm improvements in TOF reconstruction, issues related to the calibration of both the TOF and the COT dE/dX, and occupancy issues for the TOF. In Section 6, we describe the current performance of the lepton and jet charge taggers. We also discuss the status of the work on same-side and opposite-side kaon tagging. Finally in Section 7 we address the question of our current reach for B_s^0 mixing and try to project what our future, but near term, reach can be. We end this report with the recommendations of the committee.

2 Trigger Issues

As described in Section 4, one of the primary issues in maximizing B_s^0 yield is the trigger. In this section, we describe the trigger strategies utilized for the search for B_s^0 mixing, discuss the limitations, recent improvements and prospects.

There are two main classes of trigger paths: the "all hadronic" path and the "semileptonic" path. Dating back prior to Run II, the primary plan of the B group has been to utilize the semileptonic sample for flavor tagging calibration, while the hadronic decay modes would yield the best measurement of B_s^0 mixing. For low values of Δm_s, however, the semileptonic sample may indeed be useful for the mixing search.

2.1 Hadronic Trigger Path

In the B_s^0 mixing search, the goal of the hadronic trigger path is to use the SVT to isolate a large sample of fully reconstructed B_s^0 decay modes, such as $B_s^0 \to D_s^- \pi^+$, with $D_s \to \phi\pi^+$ and $\phi \to K^+K^-$. The primary issues with the hadronic trigger path are bandwidth limitations and XFT/SVT efficiency. The strategy at Level 1 (L1) is to use track-only topologies to reduce the rate to a level that is acceptable for Level 2 (L2), where SVT information can be utilized to preferentially select long-lived heavy flavor particles.

In Level 1, the two-track trigger is used to select track topologies consistent with hadronic B decays. Since the triggers are track-only, the cross sections are large, but still reflect a drastic reduction over the inelastic $p\overline{p}$ cross section. For all Level 1 two-track triggers, only XFT tracks with $p_T > 2\,\mathrm{GeV}$ are considered. The Level 1 hadronic trigger paths are:

- **"Scenario A"**: L1_TWO_TRK2_OPPQ_DPHI135_SUMPT5.5_DPS
  - two XFT tracks, $p_T > 2\,\mathrm{GeV}$
  - opposite charge
  - $\delta\phi_{\mathrm{SL6}} < 135°$
  - $p_T(1) + p_T(2) > 5.5\,\mathrm{GeV}$ (scalar p_T sum).
- **"Scenario C"**: L1_TWO_TRK2.5_OPPQ_DPHI135_SUMPT6.5
  - two XFT tracks, $p_T > 2.5\,\mathrm{GeV}$
  - opposite charge
  - $\delta\phi_{\mathrm{SL6}} < 135°$
  - $p_T(1) + p_T(2) > 6.5\,\mathrm{GeV}$ (scalar p_T sum).
- **"Scenario Low"**: L1_TWO_TRK2_DPHI90_DPS
  - two XFT tracks, $p_T > 2\,\mathrm{GeV}$
  - $\delta\phi_{\mathrm{SL6}} < 90°$.
- **"Auto accept path"**: L1_SEVEN_TRK2 (> 6 XFT tracks with $p_T > 2\,\mathrm{GeV}$)

The cuts listed in the trigger name are fully descriptive. The term "DPS" refers to dynamic prescaling, where the prescale factor changes to maximally utilize bandwidth. The Scenario C trigger is the only path which is not dynamically prescaled. The "auto-accept path" is a consequence of the L1 Two-Track Trigger hardware. For events with more than 6 tracks, it is not possible to evaluate all track pairs. Since the cross section (although it has a very large growth term) for this type of event is small, they are simply accepted regardless of track-topology.[1]

At Level 2, XFT tracking information is supplemented by silicon (SVT) information. Furthermore, the L2 processor has the ability to loop over the full track list, so there are no multiplicity limitations (no auto-accept paths) as there are in L1. For the Level 2 triggers, the following cuts are used:

- prerequisite on L1 trigger path

[1] Beginning with PHYSICS_2_04_v1, the auto-accept path is merged with standard two-track paths at Level 1.

- two SVT tracks[2] with $\chi^2 < 15$
- tracks must satisfy L1 cuts and have $120\mu\text{m} < d_0 < 1\text{mm}$
- track pair: $L_{xy} > 200\mu\text{m}$, $2^\circ < \delta\phi_0 < 90^\circ$,

where L_{xy} is defined as the two-dimensional distance from the primary vertex to the intersection of the two tracks, and $\delta\phi_0$ is the difference in ϕ_0 for the two tracks. The L2 hadronic triggers have the name "B_CHARM" along with a descriptor denoting which Level 1 path was required: "B_CHARM_HIGHPT" is driven from Level 1 Scenario C, "B_CHARM" is driven from Scenario A, and "B_CHARM_LOWPT" is driven from Scenario Low. For efficiency studies, there are heavily prescaled backup paths without cuts, *e.g.* SVT, applied.

The Level 3 triggers use COT and silicon tracking to verify the Level 2 criteria.

Algorithm issues

The Level 1 trigger paths have been studied for years [5, 6, 7]. We know that Scenario A is about as good as we can do in terms of efficiency versus purity. Scenario C provides a lower rate at the cost of a factor of ~ 2 in efficiency. At higher luminosity, when the Scenario C cross section becomes a concern, we have implemented a "two track + muon" trigger that attempts to preserve flavor tagged events. This is implemented in PHYSICS_2_04_v1 (April 17, 2004). The cross section for a "two-track plus 2 GeV electron" is too high to be of use, although we could implement a two-track plus 4 GeV electron trigger if that is deemed useful.

For each of the different L1 paths, the L2 trigger topology cuts (p_T, charge, $\delta\phi$) are set to match the L1 path. The SVT impact parameter and L_{xy} cuts are the same for all of the paths listed here.

Efficiency concerns

As described later in this document (see Figure 10), the hadronic D yield was seen to drop over the summer 2003 and this drop was ultimately attributed to reduced COT single-hit efficiency coming from chamber aging. In addition to a reduction in plateau efficiency, the XFT turn-on curve has softened, which provides an additional inefficiency for this sample, which is dominated by tracks with $p_T \sim 2\,\text{GeV}$. Some efficiency was regained by lowering the COT thresholds in December 2003, although further aging mitigates any improvement. It is important to note that there are no further handles to improve XFT efficiency if the COT hit efficiency continues to fall.

With the COT fully operational, the hadronic trigger paths use a large portion of our total L1 bandwidth, but they do not dominate the total bandwidth at L2 or L3. At $\mathcal{L} = 5 \times 10^{31}\text{cm}^{-2}\text{s}^{-1}$, the approximate Scenario A L1A/L2A/L3A (L1A = "Level 1 accept") rates are 14kHz/60Hz/10Hz out of a total budget of 25kHz/350Hz/75Hz. In the CDF trigger system, there are efficiency and bandwidth considerations at all levels of the trigger.

[2]Prior to PHYSICS_2_01_v4, the cut was SVT $\chi^2 < 25$.

In the following subsections, we outline some of the relevant issues associated with the hadronic trigger paths and steps to optimize the yields to tape.

2.1.1 Issues at L1

Fundamentally, there are two L1 trigger issues for the all-hadronic modes:

- available trigger bandwidth (limited to $\sim$25 kHz)
- L1 tracking (XFT) efficiency.

Prior to Run II, several studies showed that Scenario A would be manageable up to an instantaneous luminosity of $\mathcal{L} = 1 \times 10^{32}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [5, 6, 7]. The initial Run II observed two-track cross sections were significantly higher than expected. This was partially due to the XFT running in "2-miss" mode. Since Oct 2002, we have been running the XFT in "1-miss" mode, and the observed two-track cross sections are still higher than expected. Since the predicted cross sections came from Run I track data, we expected them to be very accurate. The high Run II cross sections are believed to arise from a higher than expected COT occupancy due to more material inside the Run II COT bore than in the Run I CTC bore. The cabling and readout infrastructure of the silicon detector at higher η within the COT bore produces secondaries as forward tracks scatter off the material.

This occupancy leads to fake (and real) segments in the XFT which,in turn, translate into an XFT track rate that is larger than expected. Table 1 shows the observed and predicted L1 cross sections for the two-track triggers.

The additional occupancy coupled with a loss in single-hit efficiency from chamber aging leave the XFT and track-based triggers in a difficult position. Although the XFT device itself is performing as expected, there are few handles to further improve efficiency or reduce fake rate in the current system. The XFT upgrade is expected to allow for improved fake rejection, but the upgrade will not be installed until 2005. The large growth rates seen in the two-track triggers is a major concern for higher luminosity running.

The amount of available L1 bandwidth is less than expected. The specification for a peak L1A rate of 45 kHz will likely never be achieved. Improvements to the system in early 2004 permit L1 accept rates as high as 25 kHz. The primary limitation comes from a combination of SVX readout and L2 processing. There is an additional limitation arising from the need to prevent the SVX readout from running at a resonant frequency that will excite and break wire bonds.

With 25 kHz L1A rate, the system as a whole is limited by L2 processing (SVT + L2) and data rates into the event builder. Because of these limitations, the Scenario A trigger must be prescaled at luminosities above $5 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The Scenario C trigger with higher momentum cuts remains unprescaled up to $7 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the highest luminosity seen so far.

To recover additional hadronic decay modes, the Scenario Low trigger was introduced with a dynamic prescale in November 2002. The dynamic prescaling allows for

Luminosity ($\times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$)	Scenario A σ (μb) predicted	Scenario A σ (μb) observed [extrapolated]	Scenario C σ (μb) predicted	Scenario C σ (μb) observed [extrapolated]
3	170	280		115
5	205	330		130
7	240	410		167
15		[880]	163	[500]

Table 1: The predicted cross sections from [5] for the Scenario A & C triggers as a function of instantaneous luminosity. The observed and extrapolated cross sections come from data taken after the Fall 2003 shutdown. XFT efficiencies were seen to fall over the Summer 2003.. After the Fall 2003 shutdown, COT discriminator thresholds were lowered, improving efficiency and leading to larger cross sections.

unused bandwidth to be filled at lower luminosities. The Scenario Low trigger has lower purity than Scenario A, but studies have shown that B_s^0 signals can be reconstructed in the Scenario Low path with good signal-to-noise [8]. The Scenario Low trigger has no opposite charge requirement on the two tracks, allowing for more trigger combinations in real B_s^0 decays.

A new additional implementation of dynamic prescaling called Über-prescaling (ÜPS) is currently being tested as a way to improve the L1 throughput [9]. This can be thought of as dynamic prescaling on the microsecond timescale. Two concerns with Über-prescaling are the SVX resonance detector and the L2A throughput. The mechanics of Über-prescaling have been tested and are working, however to optimally utilize the trigger bandwidth, it is necessary to prioritize the triggers. In other words, if limited bandwidth is available, it is optimal to preferentially choose a Scenario A trigger over a Scenario Low trigger. Work is ongoing to fold this additional layer of logic into our prescaling system.

In addition, we can further optimize our higher purity DPS triggers by implementing fractional prescaling (FDPS). The time spent between a prescale factor of 2 and 1 is lengthy, while the system awaits enough available bandwidth to double the DPS trigger rate. With fractional prescaling, we can reduce the prescale factor in smaller steps and further optimize our bandwidth utilization. Fractional prescaling has been tested successfully in the two-track trigger board and will likely be implemented in conjunction with the new Über prescaling prioritization. We anticipate operating with ÜPS and FDPS in June 2004.

The new enhanced two-track board will allow for additional combinations at L1 [10]. The advantage of this is that the auto accept path (> 6 tracks with $p_T > 2\,\mathrm{GeV}$ in the current system) becomes > 9 tracks in the new system. This saves L1 bandwidth, and further maintains the option of running with lower p_T cuts. The enhanced two-track board will also have the option to cut on transverse mass at L1. Although this is immediately useful for the $B \to h^+h^{-\prime}$ path, it is currently being studied as a discriminator that may allow for better purity with lower track momentum cuts. The timescale for the enhanced two-track board (formerly known as the three-track board) is July 2004.

At lower luminosity ($\mathcal{L} < 2.5 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$) there is additional bandwidth at L1 that is currently unused. A Scenario A+ trigger has been proposed that drops the track p_T cuts to 1.7 GeV. This trigger again has lower purity, but MC studies show that B_s^0 signals can still be reconstructed with good S/N. The Scenario A+ cross section is approximately 1 mb, so this is the ultimate bandwidth-filler. This trigger can be implemented in conjunction with ÜPS, and can be maintained (or extended to even lower p_T) with the enhanced two-track board.

2.1.2 Issues at L2

Many of the issues at L1 feed directly into L2 as well, because higher L1A rates mean higher L2A rates.

The SVT efficiency is lower than anticipated due to dead SVX ladders, lower SVX hit efficiency, and broad SVT clusters, which lead to mismeasured tracks. In June 2003, the SVT went from `4/4` logic to `4/5` logic. This significantly improved the SVT efficiency and doubled the cross section for SVT triggers.

The rejection factor on the hadronic triggers is lower than anticipated. Based upon [5, 6], we anticipated a L1/L2 rejection factor of roughly 500. Instead, we see a rejection factor of roughly 220 at $3 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and 160 at $5 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. (These numbers are from `4/5` running.) The lower rejection factors are likely arising from fake XFT tracks, tails in the SVT d_0 resolution, and beamline slope.

Level 2 bandwidth limitations coming from L3 processing power, the event builder and TDC readout have all been addressed in early 2004. The L3 processing power was significantly improved with the inclusion of new L3 nodes. The event builder throughput was improved by re-cabling the input so that the data volume is more balanced. The TDC readout time was reduced by switching over to a new version of the front end processing (DSP) code. Combined, these improvements allow the system to run with L2A rates at or above 350Hz.

In addition, a number of improvements have been incorporated to further reduce system deadtime and improve trigger purity. Beginning with trigger table PHYSICS _2_04_v1, we are cutting on muons at L2. This reduces the total L2 cross section by $\sim 1\mu$b out of a total L2 cross section of $\sim 6\mu$b, which significantly improves available bandwidth. Other improvements, such as the SVT fast-reject and streamlining of L2 processing code have resulted in further improvements in system performance.

Rolf Oldeman has recently shown that we can further reduce the L2 cross section for SVT triggers by requiring barrel information consistent with the two tracks originating from the same vertex [11]. This selection is especially powerful at higher luminosities, where multiple interactions become problematic. Studies have shown that this additional criteria reduces the two-track cross section by as much as 40% at high luminosity with no loss in efficiency. This is currently being tested and is expected become part of the default selection for all SVT two-track triggers in the very near future.

Although the system is meeting or exceeding the Run IIA specifications, the unexpected growth in trigger cross sections require further throughput improvements. One of the primary L2A limitations is the event builder. An upgrade is in progress which, combined with readout improvements, should allow a L2A rate of $\sim 1kHz$ [12]. The timescale for these improvements is the Summer 2005.

2.1.3 Issues at L3

In the summer of 2003, we began using silicon tracking in L3. This improves the hadronic trigger path by tightening the confirmation on the SVT tracks. Since the L3 processor upgrade in January 2004, L3 processing power is not a limitation.

With all of the improvements outlined in the previous section, it is possible that we will begin to reach the limits of the CSL at very high luminosity. This has not yet been explored due to compromised COT running. The CDF proposal to upgrade the CSL has been accepted by the lab and is crucial for running at higher rates [13].

2.1.4 Current Operating conditions

Since many of the improvements outlined in the previous section have come while we have been operating with the COT compromised, we do not yet have a clear picture of exactly where the current trigger table begins to incur significant deadtime. With all of the recent improvements, we estimate the following operating conditions:

- $\mathcal{L} > 8 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$**.** Scenario C requires prescale. Two-track + CMU/CMX triggers unprescaled.
- $8 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1} > \mathcal{L} > 5 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, Scenario C unprescaled, Scenario A DPS (with 1 <prescale< 3).
- $5 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1} > \mathcal{L} > 4 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, Scenario C unprescaled, Scenario A DPS (prescale=1), Scenario Low with DPS to fill in bandwidth.
- $\mathcal{L} < 4 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, All paths have prescale =1.

These operating points reflect major improvements in the system between January and May 2004. We congratulate members of the L2, SVT, L3, trigger and operations groups who implemented these improvements. As discussed below, more improvements are forthcoming, and these improvements will be needed to keep these triggers live at higher luminosity.

2.1.5 Future Improvements

In the near future, we anticipate implementing a prioritized DPS/ÜPS trigger along with fractional prescaling. These elements combined, along with the performance improvements outlined above, will fully saturate and optimize our L1 bandwidth usage.

Supplementing the suite of L2 triggers with a L2 DPS $\phi \to K^+K^-$ trigger [14] or a Scenario A trigger with relaxed d_0 cuts will insure that the L2 bandwidth is fully saturated at all luminosities. Beyond these final steps to improve our bandwidth usage, further work on the trigger system will primarily involve continued efforts to maintain good livetime at high luminosity. All further steps taken to improve performance at high luminosity will translate into additional bandwidth at all luminosities.

One improvement that is in the pipeline is the SVT Road Warrior board, which will suppress redundant fits in the SVT, reducing the total number of fits performed by a factor of ~ 2.5. This should reduce the SVT timing (one of the contributors to L2 processing time) by $\sim 25\%$. Another near term improvement to L2 processing is additional pipelining of the L2 decision: non-SVT triggers will be processed in parallel with SVT, followed by processing of SVT triggers.

Longer term improvements that should further reduce L2 processing time include the SVT Associative Memory Upgrade and the L2 Pulsar system. To further improve the L2A rate, the event builder upgrade is the crucial piece, followed by further improvements to the TDC readout. Finally, the CSL upgrade will allow us to further increase the L3 accept rate. All of these upgrades are necessary to preserve the track-trigger program to luminosities beyond $1 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

2.2 Semileptonic Trigger Path

As stated above, the primary motivation for the semileptonic trigger path is to accumulate a large sample of semileptonic B decays for flavor tagging studies. As described in the Yellow Book [3], the sensitivity for B_s^0 mixing using semileptonic B_s^0 decays is limited by the resolution on the $\beta\gamma$ (boost) factor due to the missing momentum. For low values of Δm_s, the semileptonic sample can contribute to the mixing sensitivity.

The strategy to accumulate large semileptonic samples is to require a 4 GeV electron or muon at L1, along with a single SVT track at L2 that is near the lepton.

The Level 1 trigger paths utilized are:

- **Electron**: L1_CEM4_PT4_&_TRK2_DPHI100
- **Muon**: L1_CMUP6_PT4

In the case of the electron + SVT path, we require the track be found at L1. Since the SVT tracks are a subset of the XFT tracks, this causes no inefficiency.

The Level 2 trigger paths are:

- **Electron**: L2_TRK2_D120_&_CEM4_PT4_CES3_2DPHI90
  - one SVT track with $\chi^2 < 15$,
  - CES confirmation on the electron,
  - tracks must satisfy L1 cuts and have $120\mu\mathrm{m} < d_0 < 1\mathrm{mm}$,
  - track pair: $2^\circ < \delta\phi < 90^\circ$.

- **Muon**: L2_CMUP6_PT4_&_TRK2_D120_DPHI90
  - one SVT track with $\chi^2 < 15$
  - tracks must satisfy L1 cuts and have $120\mu\text{m} < d_0 < 1\text{mm}$,
  - track pair: $2° < \delta\phi < 90°$.

For efficiency studies, there are heavily prescaled backup paths without cuts, *e.g.* SVT, applied. This version of the μ+SVT trigger was implemented in April 2004 with table PHYSICS_2_04_v1. Prior to that, muon information was not available at L2, so it was not possible to prevent the muon from being the SVT track. These events were rejected at L3, but we can now reject them at L2.

For the semileptonic paths, there are really no major trigger issues at L1. The electron path has 10μb cross section with big overlap with Scenario A/C, while the muon cross section is only about 1μb. These can be easily maintained even at high luminosity.

At L2, there are bandwidth concerns associated with these triggers. The electron path is 700nb at $5\times10^{31}\text{cm}^{-2}\text{s}^{-1}$ with a growth that appears linear. For the muon path the trigger cross section is 300nb at $5\times10^{31}\text{cm}^{-2}\text{s}^{-1}$ with a growth term that appears linear. Combined, the semileptonic trigger path is 1μb at $5\times10^{31}\text{cm}^{-2}\text{s}^{-1}$. To reduce the electron cross section, we are currently testing a two-track L_{xy} cut. For efficiency studies, we will keep a prescaled trigger that accepts events without an L_{xy} cut.

The L3 cross sections at $5\times10^{31}\text{cm}^{-2}\text{s}^{-1}$ are 50nb (37nb) for the electron (muon) path. This is less than 10% of the total L3 bandwidth and is not a problem. Silicon tracking is not used due to concern about the problem with semileptonic lifetimes. Instead, the COT/SVT hybrid tracks are used to implement a mass(lepton+track) $< 5\,\text{GeV}$ cut.

Comments on Semileptonic Samples

The semileptonic triggers were specifically designed to accumulate large, clean samples of semileptonic B decays for flavor tagging studies and calibrations. These samples provide a number of interesting physics measurements, but are not ideal for a B_s^0 mixing search in this channel.

A number of people have suggested additional semileptonic triggers that would be better suited for accumulating samples for a B_s^0 mixing search in the semileptonic channel [15, 16]. The proposals vary in the details, but the general idea is to trigger on a topology where there are trigger tracks on the side opposite the trigger lepton. The idea is that both b hadrons are needed to tag the flavor of the B_s^0 at production and at decay. The lepton would tag the flavor at the time of decay, and the away-side information would be used for tagging flavor at production. Topologies of this type may allow the impact parameter cuts on the trigger (B_s^0) side to be relaxed or removed completely. An example of this would be a 4 GeV lepton plus track with no impact parameter cut in addition to an away-side SVT track. For mixing, it is best to have the B_s^0 unbiased by any impact parameter cuts, since the best proper time resolution is at low L_{xy}. This trigger will also supplement the fully reconstructed B_s^0 sample, although studies are needed to determine whether it is the most efficient use of bandwidth.

We have the ability to implement triggers of this type. However, they would utilize a non-negligible amount of L2 bandwidth and would be a significant departure from the current B_s^0 mixing strategy. If the collaboration decided to place more emphasis on the semileptonic B_s^0 search, the bandwidth for these new triggers would have to come from the existing hadronic triggers. We feel the existing trigger strategy should be maintained. However, if the collaboration chose to take a more near-term view and fully emphasize a measurement or limit of low values of Δm_s, enhancing the semileptonic trigger paths would help this effort.

2.3 Summary of Trigger Issues

Recent improvements to the CDF trigger and DAQ system allow us to keep the B triggers live at high luminosities and also free up additional bandwidth that can be filled at lower luminosities. Further improvements to prescaling will result in slight, not drastic, improvements in yields.

The bottom line is that, at current Tevatron luminosities, we are doing just about as well as we can with the trigger. There are no smoking guns or major sources of correctable inefficiency. Modest improvements will continue to come, but no large factors remain to be gained.

The two major (related) concerns over the longer term are the track trigger behavior at higher luminosities and the state of the COT. There are growth terms in the two-track cross sections that are quadratic in luminosity, which indicates that we will ultimately not be able to keep the Scenario C trigger unprescaled at $\mathcal{L} \sim 10^{32}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Since the XFT efficiency is quite sensitive to COT single hit efficiency, it is not possible to keep efficiency high and fake rates under control with a COT that is significantly degraded. We have seen this already when running in the "COT Compromised" mode. If the COT aging effect is not arrested, it is unlikely that CDF will be able to measure B_s^0 mixing.

As described above, further improvements such as the Road Warrior, L2 Pulsar, event builder and CSL upgrades are all crucial to improve bandwidth for future, higher luminosity running. The system is running close to its limits in a number of places, so further improvement is needed at all levels of the trigger to truly improve throughput. It is crucial that all of these upgrades are in place as soon as possible.

In early 2006, the XFT upgrade will help reduce fake rates and improve L1 two-track trigger cross sections. It has been shown that a 25% reduction in two-track cross sections can be achieved by raising the p_T threshold on the finder masks [17]. This is a firmware change to the XFT and could be done at any time. We would sacrifice 1.5 GeV tracks, which would affect other samples (like $J/\psi \to \mu^+\mu^-$) but should be considered if the luminosity exceeds $8 - 9 \times 10^{31}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ before the XFT upgrade is commissioned.

As a final point, more studies are needed as soon as possible to optimize trigger purity for higher luminosity running. The Scenario C trigger must be prescaled before the Tevatron luminosity gets to $\sim 1\times10^{32}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. It will be crucial to accumulate the type of events (by flavor tag or by topology) that are of most use to the B_s^0 mixing search. This type of optimization has not yet been performed, and will be very important once the current inclusive program must be abandoned.

3 Tracking

Tracking issues can significantly impact the B_s^0 mixing measurement in several ways. The sensitivity of the measurement depends exponentially on the square of the proper lifetime resolution. In addition, the measurement is sensitive to a number of other tracking related issues including tracking efficiency, mass resolution and the rate of mis-measured tracks. Each of these issues are discussed below and, where possible, numerical estimates of the impact on the B_s^0 mixing measurement are given. Finally some suggestions on where future efforts on tracking issues would best be concentrated in the context of the B_s^0 mixing measurement are given.

The studies described in this section are based on ProductionExe versions 4.8.4, 5.1.3 and 5.3.1. ProductionExe is our primary reconstruction program and all tracking reconstruction, except for adding L00 hits, is done in ProductionExe. The ProductionExe versions are usually divided into two groups: the 4 series version, 4.8.4, and the 5 series versions, 5.1.3 and 5.3.1. Between the 4 series version and the 5 series versions the COT and silicon alignment was substantially improved. Between versions 5.1.3 and 5.3.1 the silicon only tracking was reprogrammed increasing the efficiency of the algorithm. In general, where improvements in tracking performance are discussed we give either absolute efficiencies and resolutions or relative improvements compared to the 4 series release performance. Detailing the relative improvement compared to the 4 series release performance is useful since most B physics yield and resolution studies where done using the data processed by the 4 series tracking and the relative improvement numbers given in this report can be used to multiply those measurements. The 4 series data processing was only performed for the data through September 2003 and the tracking performance for the data taken after September 2003 is currently being evaluated.

Discussions on the six tracking algorithms that will help in fully understanding this section can be found in the CDF note catalog: COT segment linking (SL) [18] and histogram linking (HL) [19], combined COT silicon outside in (OI) [20] and Karlruhe OI [21], silicon standalone (SiSA) [22] and silicon seeded combined silicon COT inside out (IO) [23].

3.1 Tracking Efficiencies and Mass Resolution

Improvements in tracking efficiency can directly improve the B_s^0 mixing measurement by increasing the yield of reconstructed B states or increasing the number of tracks used for tagging purposes. Specifically, improvements in low momentum tracking efficiency are useful for reconstructing high multiplicity decays and soft fragmentation tracks, while improvements in general tracking efficiency are useful for reconstructing leptons and other higher momentum tracks. In addition, improvements in silicon tracking efficiency increases the number of tracks with significant lifetime information. Improvements have been made in both the COT and silicon low momentum tracking efficiency and in the SiSA tracking performance in the forward region.

The efficiency and resolution of the COT tracking code has been improved relative to the 4 series releases. These improvements are a result of better alignment and a new, aggressive, hit dropping scheme. In addition, the default quality cuts on the number of axial and stereo COT hits where redefined to be compatible with the new hit dropping scheme and improve the efficiency for finding tracks at low momentum and higher η. The performance of low momentum tracking efficiency can be be evaluated by looking at the yield of states that decay to one or more slow moving particles. For instance, the D^* can be reconstructed by adding a slow pion to the D^0 state. The yield of slow pions, reconstructed in the this mode, indicates a 9% relative improvement in central tracking for low momentum tracks above 400 MeV [24]. This improvement can be used to estimate the increase in yield in decay chains involving low momentum particles. The COT and IO algorithms also have sensitivity for tracks below 400 MeVand the yield of D^* should be studied for those candidates with soft pions in that transverse momentum region. The hit dropping scheme primarily effects mass resolution. The width of the mass peaks for two track states such as the J/ψ is improved from 27MeV to 23MeV [25]. The better mass resolution improves the signal to background of the B_s^0 mixing samples, which will enhance the mixing fit.

The refined COT alignment and resolution also have allowed us to develop an improved silicon alignment enhancing the Outside-In (OI) silicon tracking efficiency. The OI tracking efficiency, measured using J/psi di-muon candidates, rises from 88.4 ± 0.1(stat)% to 90.3 ± 0.1(stat)% [26]. This is the efficiency to add axial silicon hits using a only one missed hit allowed definition for the numerator. The 4 and 5 series data used for this comparison were nearly identical. The only difference was in the number of denominator COT tracks found in the two releases which is negligible for tracks above 1.5GeV. Another advantage of the new alignment is that now we find that when using a slightly looser definition of the efficiency numerator (requiring only three layers of SVXII with silicon hits) that the efficiency rises to 93.6 ± 0.1(stat)% [27] with a negligible change to the fake rate. Also the OI efficiency is found to be flat with p_T where previously it dropped at high momentum. In addition, the efficiency for adding both silicon stereo and axial information to COT tracks is 84.8 ± 0.1(stat)% [28]. This will be useful if we decide to use 3D vertexing for the mixing measurement.

The overall resolution for mass states that include tracks with silicon information is improved, as well. The width for two track states such as the J/ψ or the D^0 are improved from 15MeV to 14MeV and 11MeV to 9MeV , respectively (Figure 2) [24]. The mass resolution of a three track state such as the $B \to J/\psi K$ improves from 16MeV to 14MeV. This result includes a mass constraint on the J/ψ [25].

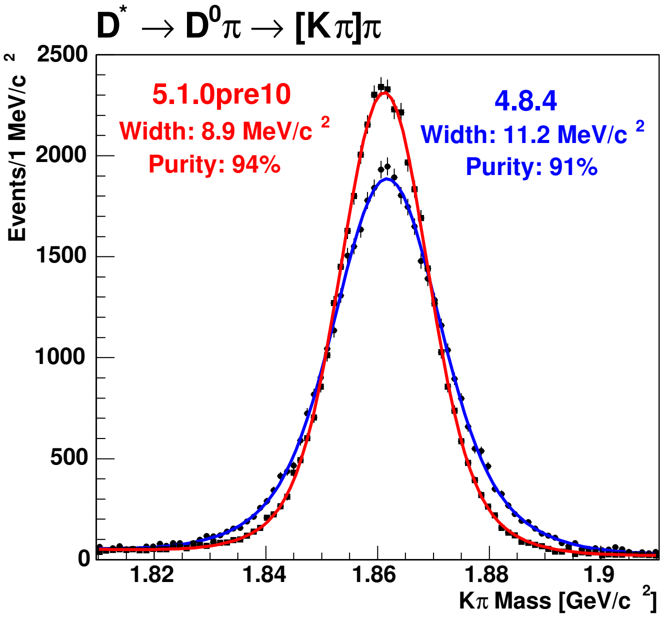

Figure 2: The invariant mass distribution of $D^0 \to K\pi$ candidates in 4.8.4(blue) and 5.1.0(red). Also given are the width and purity of the sample.

Considerable advances have been made in the performance of the silicon only tracking algorithm. This algorithm finds forward tracks and low momentum tracks in the central region that are missed by the COT or are too low momentum to reach the COT. For higher p_T leptons in the forward region, we are able to measure the absolute efficiency using forward calorimeter energy deposits that are consistent with coming from $Z \to ee$ candidates as the denominator. The efficiency in the forward direction for high p_T isolated leptons increases by 5% for tracks with $|\eta|$ near one and by 15% at $|\eta|$ near two (Figure 3) [29] [30]. In addition, the yield of low momentum tracks is increased by a factor of two [31]. These improvements in lepton efficiency portend an increase in yield of opposite side forward leptons for flavor tagging. The increased efficiencies should also improve the yields of multi-body decays where one or more of the tracks may be forward going and for flavor tagging using jet charge techniques. The SiSA tracking algorithm or a similar algorithm should also have efficiency for tracks below the minimum transverse momentum threshold for tracking in the COT and in the region of transverse momentum below 400 MeVwhere the COT efficiency is sharply falling. Such an algorithm would also increase the yield of slow pions reconstructed in D^* decays by adding candidates where the slow pion has transverse momentum below 400 MeV. Studies of algorithms of this type are being conducted by the MIT and Duke groups. The MIT algorithm which is based on the standard SiSA algorithm included in release 5.3.1 and can be tested by users.

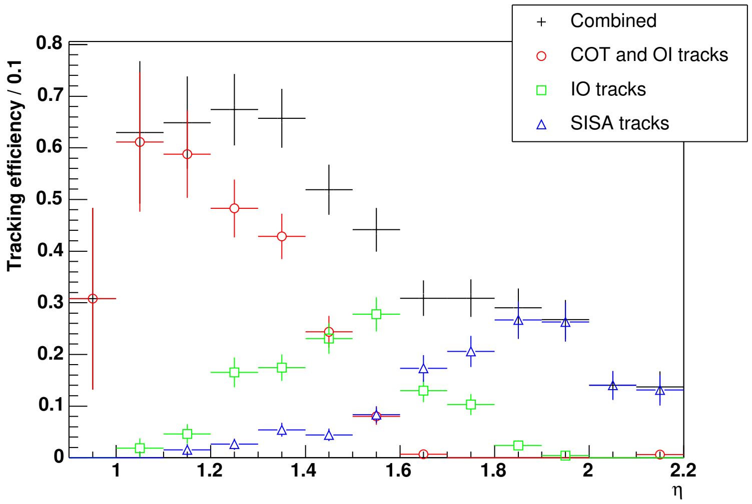

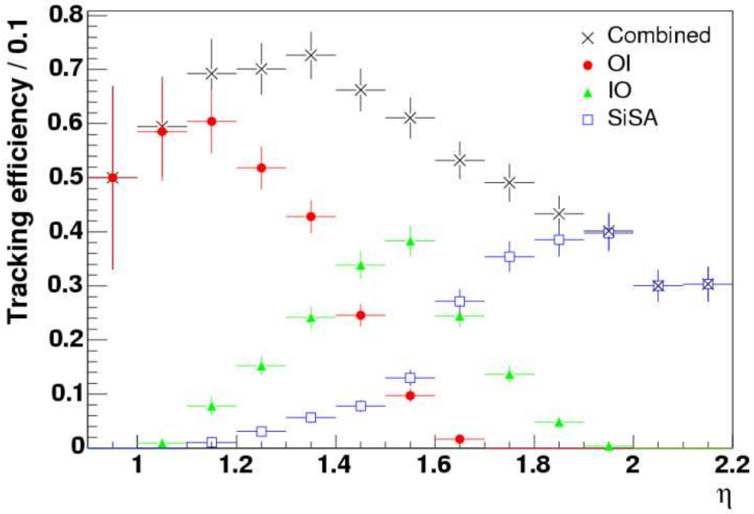

Figure 3: Forward isolated high pt lepton efficiency for 5.1.3(upper plot) and 5.3.1(lower plot).

3.2 Vertexing and Lifetime resolution

Improvements in vertexing can affect the B_s^0 mixing analysis in two ways. An improvement in vertexing resolution directly enhances the sensitivity of the mixing measurement. Further, reducing the rate of mis-measured tracks can improve the fits for the mixing parameters. Better alignment is the primary contributor to improved vertexing in recent releases. The effect of the COT alignment improvements is clearly seen in the reduction in systematic offsets in the impact parameter (d_0) vs. ϕ distribution (Figure 4) [32] [25]. The deviations are reduced by a factor of 3 and the distribution no longer exhibits an average negative signed impact parameter. The COT alignment also, indirectly, improves the silicon efficiency. A better silicon alignment has reduced systematic effects in the distribution of d_0 vs. ϕ for silicon tracks and improved the z positions of the barrels. The systematic improvements, though significant, were far smaller than the per track resolution on d_0. Therefore, the average vertexing resolution is not noticeably improved. However, the rate of mis-measured tracks is improved and the number of candidates occurring in the tails of lifetime distributions is reduced, leading to improved fits to lifetimes distributions. In addition, improvements in the z alignment have fixed biases that were observed in lifetime distributions when using silicon stereo and z hits [33]. Using the stereo information further reduces the number of candidates seen in the tails of the lifetime distributions. This may be because the small angle stereo strips add axial as well as stereo information. A comparison of the $B^+ \to J/\psi K^+$ lifetime in release 4.8.4 and release 5.3.1 using stereo information is given in figure 5 [34]. These studies show that in release 5.3.1 the silicon stereo information can be used and is beneficial.

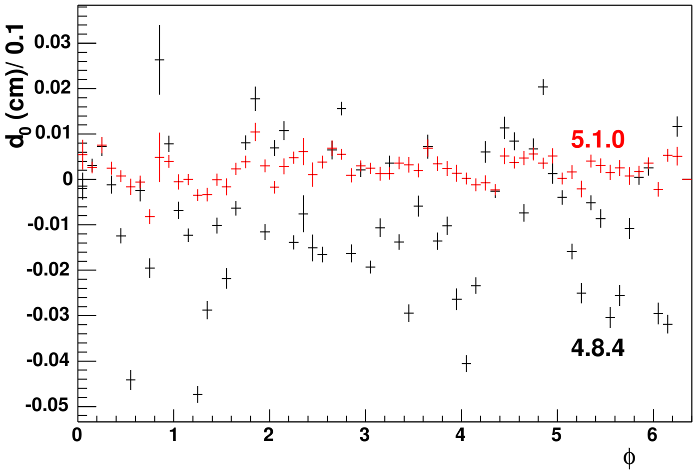

Figure 4: d_0 vs. ϕ profile distribution for COT tracks from 4.8.4 (black) and 5.1.0 (red).

Another method of improving vertex resolution for lifetimes is to use an event-by-event vertex. Currently the use of an event-by-event vertex has been attempted in both

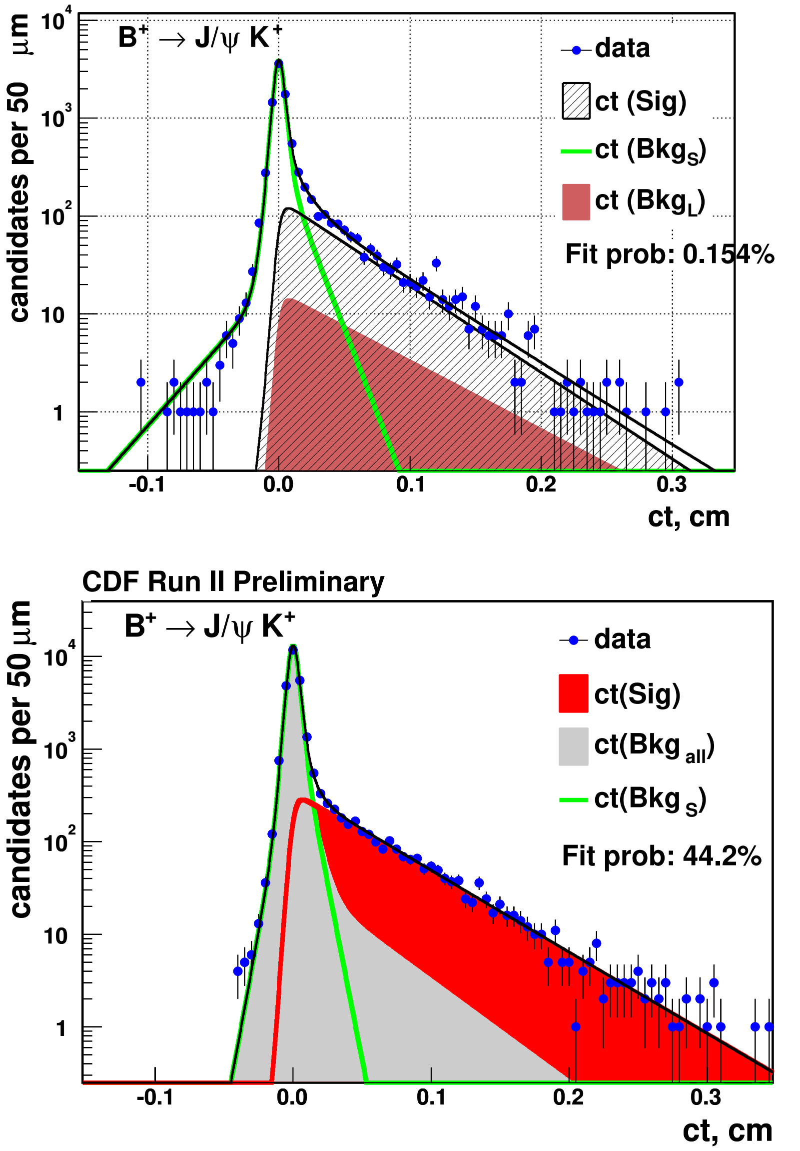

Figure 5: The $c\tau$ distribution and fit from $B \to J/\psi K$ data for 4.8.4 (upper plot) and 5.3.1 including stereo information (lower plot). Note the lower plot also includes additional data from the spring of 2004 and still has less events in the tails as indicated by the number of entries at negative lifetime.

data and MC and only a minimal improvement has been seen in the r-ϕ plane though a large improvement is seen in the z dimension [35]. We expect a larger improvement from using an event by event vertex. One possible source of the degradation of the vertexing resolution is from events where the primary vertex that is reconstructed is contaminated with tracks from the opposite side B. Another possible source of degradation is that primary vertex finder used was a naive implementation following the work of the Top physics group, which may not be optimum for low multiplicity B physics events. Further studies and tuning of the performance of the primary vertex finding algorithm is in progress. In addition, using L00 in connection with primary vertex finding will be beneficial.

3.3 Using L00

The use of L00 is expected to directly improve the resolution for lifetime measurements. Preliminary attempts indicate that L00 considerably improves the impact parameter resolution for low momentum tracks (Figure 6) [36]. However, the efficiency for attaching L00 hits is low: 59% in J-stream data. Even with this low efficiency, adding L00 information to tracks where possible in $B \to J/\psi K$ decays results in a 20% average improvement in the lifetime resolution [37]. This improvement was estimated by observing the reduction in the width of the prompt background component of the $J/\psi K$ candidate lifetime distribution.

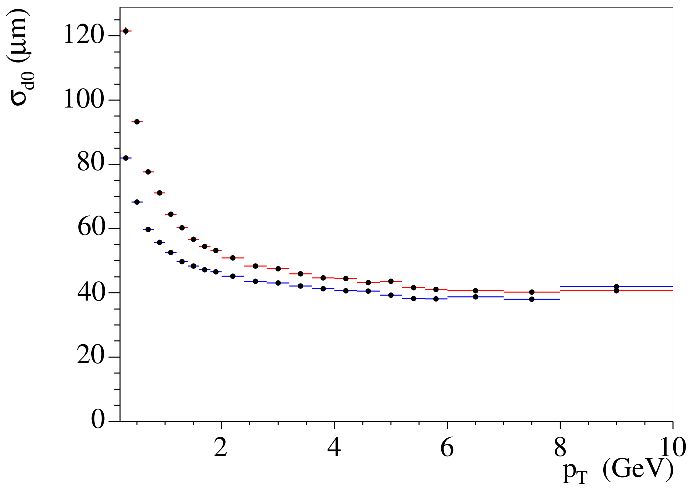

Figure 6: The width of the impact parameter distribution for tracks with (blue) and without (red) l00 hits as a function of pt.

Continuing work on L00 is focused on improving the efficiency. Currently low level issues such as dead strips and the performance of the L00 pedestal fit have been studied and the 5% inefficiency due to these sources is irreducible [38]. In the course of these studies it was also found that several ladders, representing 5% of the acceptance of the detector, were underbiased and this problem was corrected. Studies on improving the L00 pattern recognition are underway and current results indicate that 5-10% improvements in efficiency may be possible [39]. The remaining source of inefficiencies are from the quality cuts applied to L00 data which could possibly be optimized to improve efficiency without sacrificing performance. These quality cuts include cuts on the maximum amount of charge and number of strips in a cluster which are designed to eliminate clusters from large η and looping tracks. Further work on improving the L00 alignment is also necessary.

3.4 Tracking Conclusions

Improvements to the tracking efficiency in the 5 series releases have resulted in relative improvements of 9% for low momentum COT tracks, 5.5% for OI tracks and an absolute improvement of 5-10% up to a factor of two for isolated forward tracks and low momentum forward tracks respectively. These numbers can be used to multiply the 4 series release yields to estimate the expected yield of reconstructed B states and improvements in the performance of lepton taggers, same side taggers and jet charge taggers in a 5 series analysis for the data taken up to September 2003. Also there have been improvements in the mass resolution, and therefore the signal to background, of approximately 10% depending on the decay mode. Finally, the tails of the proper time distributions have been reduced, which will considerably improve the mixing fits. Finally the use of L00 will directly improve lifetime resolutions by 20%.

Many further improvements could be made in tracking to improve the sensitivity of B_s^0 mixing analysis. The sensitivity of a B_s^0 mixing measurement is highly dependent on the proper time resolution and this is where effort should be concentrated. Several areas of tracking could be improved which would impact proper time resolution. The largest gain will come from full utilization of L00. The current gain in proper time resolution is on order 20%, where we expect to be able to improve the resolution by 25%. Work to improve L00 should be concentrated in the areas of efficiency of the clustering cuts and alignment. The next largest gain should come from using an event by event primary vertex. Current work indicates a modest gain in decay length resolution from using an event by event primary vertex. The gain from using an event by event primary vertex is expected to be larger and continuing efforts should be concentrated on understating the resolution of primary vertices from B physics events in data and whether there are unexpected problems reducing the effectiveness of using this technique. The final source of improvement to proper time resolution should come from improvements in the global alignment. The tracking group is concentrating on another round of improving the COT alignment by investigating whether corrections to wire by wire positions can be improved. This should improve systematic effects in the COT alignment and will improve the silicon alignment, which uses COT tracks as a starting point. In addition, B physics data should be used to study the silicon alignment. High statistics J/ψ events could be used to understand whether there are remaining systematic effects in the silicon alignment. The Top group conducted similar studies using high multiplicity primary vertices and demonstrated that physics probes could be used to make substantial improvements in the silicon alignment.

In addition, other tracking issues effect the B_s^0 mixing measurement. The primary remaining issue is the efficiency of the forward tracking. The tracking group is conducting an ongoing program to improve the forward tracking. However, input from the B physics group on the current performance of the tracking in 5.3.1 would be very useful. For instance, measurements of the improvements in yields of forward tagging leptons and high multiplicity B states would useful. Particularly evaluations of the tracking performance as a function of η, ϕ, momentum, z_0 position and density of the track environment would give valuable feedback for the tracking algorithm authors. Such input from the Electroweak and Exotic groups was critical in developing the current large improvements in the standalone tracking performance. A second tracking efficiency issue is the performance of the tracking at low momentum. An evaluation of performance as described above should also be performed for low the momentum tracking based on the COT and the silicon in both the forward and central regions. This would include testing the new low momentum tracking algorithm that is now available in release 5.3.1 and evaluating the performance of the COT tracking below 400 MeV.

The B_s^0 mixing analysis is highly dependent on the tracking performance. While the B group has made valuable contributions in areas such as calibrating the tracking momentum scale, calibrating the magnetic field strength and validating the physics performance of new releases, the group has not in general has made substantial contributions in areas that would benefit the B_s^0 mixing analysis. Concentrating only on physics related tracking issues that effect the simpler analyses that have been performed to date has delayed reaching the longer term goal of measuring B_s^0 mixing. Though MOU responsibilities such as pattern recognition algorithms and beamline production are being vigorously pursued by the responsible parties all the issues discussed above need additional manpower in order to perform a world's best B_s^0 mixing measurement.

4 Yields

B_s^0 decay modes of interest for the mixing measurement all have D_s or D_s^* in the final state. CDF triggers on both semileptonic modes, where the D_s is accompanied by an electron or muon, and the fully hadronic modes. The hadronic modes come through the two-track trigger and semileptonic modes through the lepton+SVT trigger. Measurements of the yields in all modes have been complicated by changing conditions. Some of these are by design, such as the change in the SVT trigger to using 4/5 hits from 4/4, others not, such as the COT aging which has resulted in a yield that decreases with time.

The B_s^0 yields depend crucially on the trigger configuration and trigger efficiency.

Trigger issues were discussed in Section 2 and in this section we describe studies of event yields, which come about from trigger and reconstruction efficiency and acceptance.

4.1 Hadronic Modes

The best studied B_s^0 hadronic decay mode is $B_s^0 \to D_s\pi$ with $D_s \to \phi\pi$, $\phi \to$KK. Measured yields for this mode are documented in [40]-[41]. These measurements (and all others cited here) are done with 4.8.4 reconstruction and therefore do not take advantage of the 5 series tracking improvements discussed in Section 3. We note at this point that all studies of B_s^0 yields have involved some kind of offline trigger confirmation (requiring that the reconstructed B_s^0 tracks include the SVT trigger tracks). Such confirmation, while of obvious utility for a cross section measurement, may not be needed for mixing. The task force requested of the various groups making yield measurements that they also measure them without the trigger confirmation. A preliminary study indicates that the gain may be only a few percent [42]. We encourage the completion of this study although we recognize that it may be non-trivial to include "volunteer" events and only a more substantial gain would make it worthwhile. The yields have been studied in three separate datasets, `hbot0h`, comprising 70 pb^{-1}, `hbot1i`, comprising 54 pb^{-1} and SVT `4/5` data, comprising 68 pb^{-1}. The luminosity numbers are after good run selection. The `4/5` result is not yet documented in a CDF note. These datasets comprise the complete pre-September shutdown data. We do not yet have results with post shutdown data and we comment on this below.

The datasets consist of three trigger paths: B_CHARM, B_CHARM_LOWPT and B_CHARM_HIGHPT, each of which is described in Section 2. As described in detail in the notes, reconstruction of the mode $B_s^0 \to$D$_s\pi$ with $D_s \to \phi\pi$, $\phi \to K^+K^-$ begins with reconstruction of $\phi \to K^+K^-$, then ϕ mesons are grouped with candidate charged-π tracks to make $D_s^{\pm}$ candidates which are then grouped with opposite-sign π candidates to make candidate $B_s^0 \to D_s\pi$ decays. A number of selection requirements are made in the reconstruction process, such as vertex and mass window cuts for the associated tracks. We refer the reader to the original notes for the details of these cuts. The selection was optimized for $S/\sqrt{S+B}$.

Signal yields are extracted from fits to the invariant mass distribution around the B_s^0 mass. In addition to a peak from $B_s^0 \to D_s\pi$, these distributions also include a slightly lower mass peak from $B_s^0 \to D_s^*\pi$ and $D_s\rho$ (where one of the final-state particles is not reconstructed). The signal is fit to a Gaussian, the low-mass peak (both $D_s^*\pi$ and $D\rho$) to a smeared exponential plus a Gaussian, the background from other D_sX decays as a linear function with a cutoff and combinatorics to an exponential. The fit procedure is documented in [43]. The blessed results, based on 119 pb^{-1}, are documented in [44]. The yields have been studied [45], though not yet blessed, for the full pre-shutdown dataset (192 pb^{-1}). In Figure 7 we show the fit for the combined `hbot0h`, `hbot1i` and `4/5` datasets using all three trigger paths. The fit gives a combined yield of 157$\pm$14 events. The three dataset split and yields are shown in Table 2. The `hbot0h` and `hbot1i` numbers are from [44], and the `4/5` number is from [45].

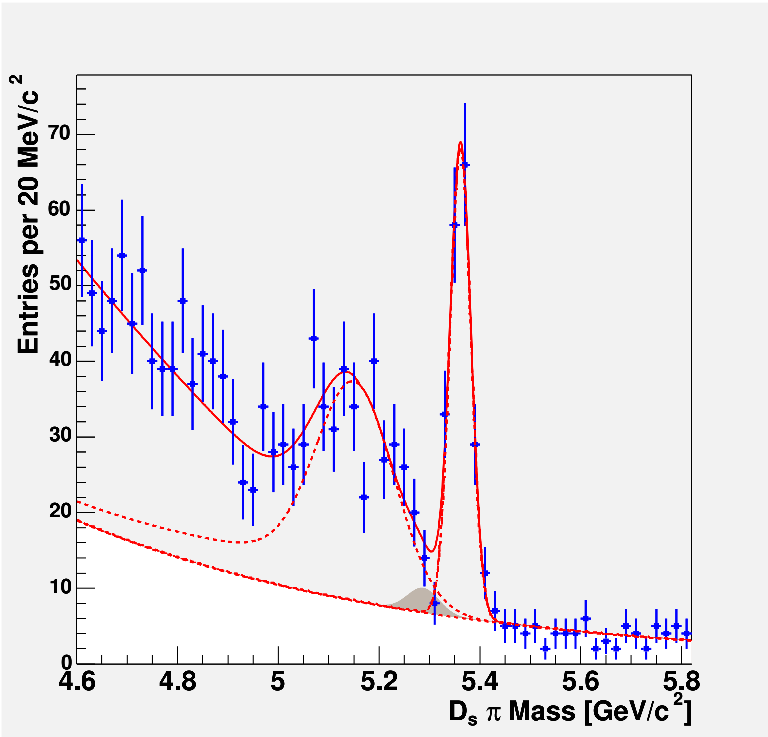

Figure 7: Fit to the B_s^0 mass spectrum for $B_s^0 \to D_s\pi$ with $D_s \to \phi\pi$ (from [45]).

Dataset	Integrated Lum	Number of events	Yield (fb)
`hbot0h`	70 pb^{-1}	45±8	643±114
`hbot1i`	54 pb^{-1}	38±7	704±130
`4/5`	68 pb^{-1}	70±10	1029±147

Table 2: Measured yields for $B_s^0 \to D_s\pi$ with $D_s \to \phi\pi$, $\phi \to K^+K^-$.

The total yield of 157 events is relevant for evaluating our reach with the current dataset (modulo the inclusion of the post-shutdown data). However, it is not a useful number for extrapolating to the future because it integrates over the change from SVT "`4/4`" to `4/5`. For that, the `4/5` number alone is more useful, but as discussed below cannot be directly used either because of the COT aging.

There are other useful modes of B_s^0 decays. The modes with D_s^* in the final state are of interest because the mass fit in the D_s mode indicates a competitive number in D_s^*, albeit with smaller S/B. It may be possible to reconstruct the low-energy photon in these events to regain full reconstruction, or use them with a K-factor. To our knowledge, no serious attempt to understand if these events are usable has been made. In addition there is the $B_s^0 \rightarrow D_s\pi\pi\pi$. Yellow Book expectations were that this mode would be comparable in yield to $D_s\pi$. Initial searches for this mode [46] have failed to find more than a hint of a signal and raise more questions than answers regarding the expectations.

Next we come to modes of D_s decay other than $\phi\pi$. We expect two other modes to have significant yields: K^*K and 3π.

In $B_s^0 \rightarrow D_s\pi$ with $D_s \rightarrow 3\pi$ a small signal of about 30 events was observed with `hbot0h+hbot1i` [47]. This is in line with Yellow Book expectations that it would be about 1/3 of the $\phi\pi$ yield. The 30 events come from a selection that was not yet fully optimized. A more recent analysis in the same dataset, including optimization of selection cuts [48], fits a signal of 57$\pm$11 events in 124 pb^{-1}. The mass spectrum and fit are shown in Figure 8.

$B_s^0 \rightarrow D_s\pi$ with $D_s \rightarrow K^*K$ has also been measured with `hbot0h+hbot1i` [49]. There is a substantial signal peak as shown in Figure 9, but the yield estimate is complicated by a B$_d$ reflection. The reflection occurs when a π^- in $D^- \rightarrow K^+\pi^-\pi^-$ is reconstructed as a K^-. To extract the B_s^0 yield the Monte Carlo is used to measure the ratio of the number of B^0 mesons that are mis-reconstructed as B_s^0 to the number that are correctly reconstructed. This ratio is then used in the data to fix the normalization of the B^0 Gaussian template used in the fit. The width and peak position of the B^0 template are also fixed from the Monte Carlo. The fit, shown in Figure 9, gives a B_s^0 yield of 176$\pm$32 events. The B^0 contribution is fixed at 146 events. The S/B is of order 1/1. The yield here exceeds the measurements in the $\phi\pi$ channel, but it is difficult to know how reliable the measurement is. The fit has little sensitivity to varying the normalization of the reflection, as the reflected peak lies nearly on top of the B_s^0 peak (ΔM=26 MeV/c^2) and the widths are very similar. In a measurement of B_s^0 mixing, this mode could certainly be used, but in a limit a significantly better understanding is needed in order to gain confidence in the fit technique and results.

The measured S/B ratios for these modes are not well documented but are approximately 3:1 for the $B_s^0 \rightarrow$D$_s\pi$ with D$_s \rightarrow \phi\pi$ mode and 1:1 or worse for the other modes.

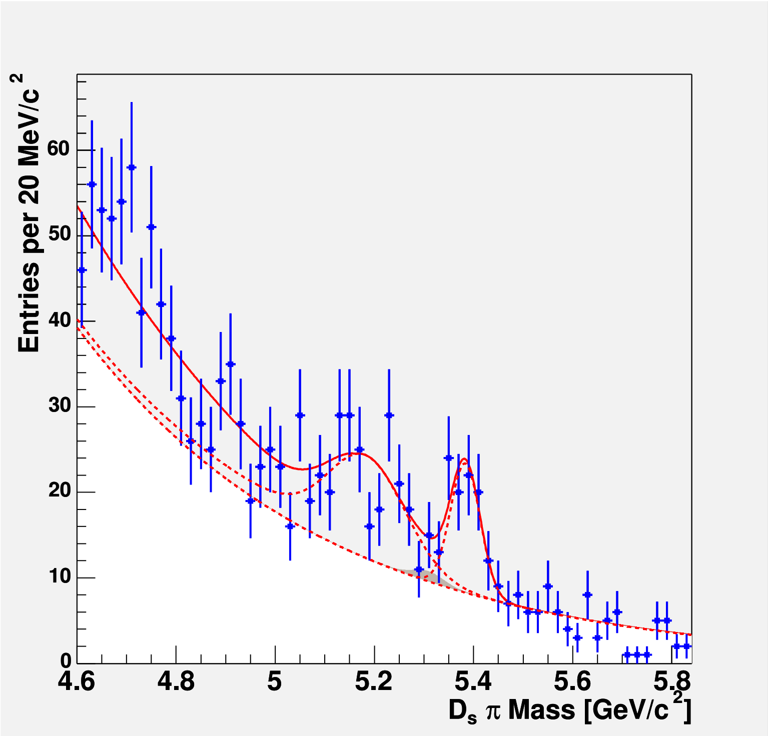

Figure 8: Fit to the B_s^0 mass spectrum for $B_s^0 \to D_s\pi$ with $D_s \to \pi\pi\pi$. From [48].

D_s Mode	`hbot0h` yield (fb)	`hbot1i` yield (fb)	`4/5` yield (fb)
$\phi\pi$	643±114	704±130	1029±147
K*K	1544±281		-
3π	460±89		-

Table 3: Summary of measured, pre-shutdown, yields in $B_s^0 \to D_s\pi$

.

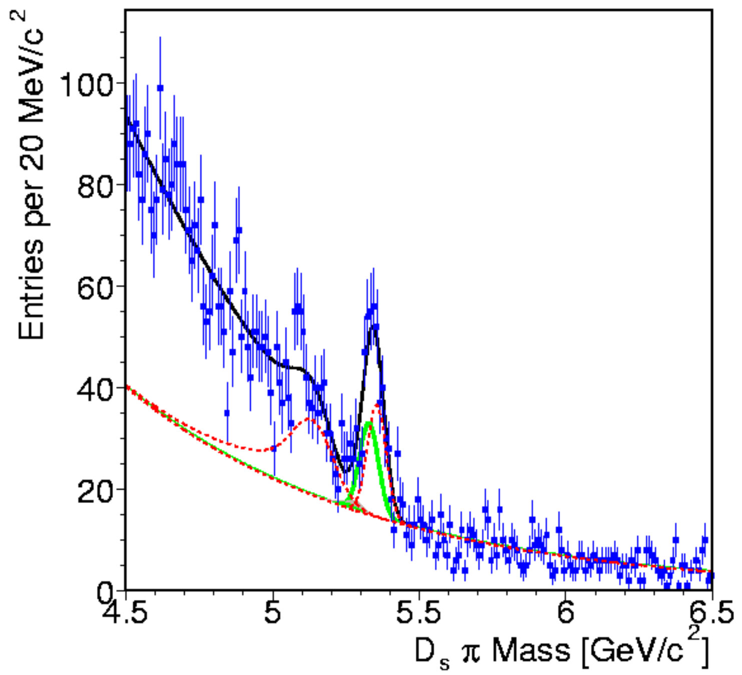

Figure 9: Fit to the B_s^0 mass spectrum for $B_s^0 \to D_s\pi$ with $\mathrm{D}_s \to K^*K$. From [49]

4.2 Time Dependence

Because of the aging of the COT, extrapolation of these yields into the future is non-trivial. Measurements of the yields in the post shutdown data, which would provide useful information, are not yet complete. There are two issues: One is what is the total yield and S/B in all useful B_s^0 modes in the current dataset through February 13, 2004 when the COT running conditions were modified. This will allow us to understand the B_s^0 mixing reach in the current data. The second is how we extrapolate our current yields into the future to predict the mixing reach as a function of integrated luminosity.

The first question is in principle relatively straightforward, though there remain significant uncertainties. We simply need a measurement of the hadronic yields in the data through February 13. What we have are measurements in the data up to either the September, 2003 shutdown, or the switch to `4/5`. We can extrapolate these to February 13 with varying degrees of uncertainty. Furthermore, the only well understood measurement is for the channel $B_s^0 \to D_s\pi$ with $D_s \to \phi\pi$. Other D_s decay modes are not well understood at this time. Finally, we expect improvements to these measured numbers from the improved reconstruction in 5.3. The hadronic datasets are not yet available in 5.3 and we will not attempt to guess at the level of improvement.

The second question is even more difficult to answer. We know that the yields have been decreasing as a function of time, but we do not know the slope very well, and more importantly we do not know the second derivative. We can only make assumptions.

Figure 10 shows the yield vs. time for $D^0 \to K\pi$ for pre-shutdown data. One notes the increase in yield due to the switch to SVT `4/5` at around day 160. After that a linear decrease in the yield is observed, attributed to COT aging, and by the shutdown the yield is about where it was pre-`4/5`.

In December 2003 the COT discriminator thresholds were lowered to recover some of the lost gain. This is expected to have brought the yields back to where they were when the switch to `4/5` was made. However, we expect significant additional yield loss due to aging in the post-shutdown, pre-February 13 data. There is therefore substantial uncertainty when trying to use these results to project into the future. We use pre-`4/5` yields as a well defined point from which to extrapolate. We assume that, future yields will be at this level. On the one hand we know there has been further deterioration of the COT, on the other hand we expect that there will be some improvements from 5.3 reconstruction. We also assume that the data with a compromised COT will not be useful for a mixing measurement.

4.3 Semileptonic Modes

Semileptonic B_s^0 decay modes do not have the same Δm reach as the hadronic modes because they are not fully reconstructed. As a result there is a "K-factor" that leads to a term in the proper time resolution that is proportional to decay time. For large Δm, corresponding to very rapid oscillations, the semileptonics do not contribute to our sensitivity in measuring Δm_s. However, the semileptonic yield is significant and

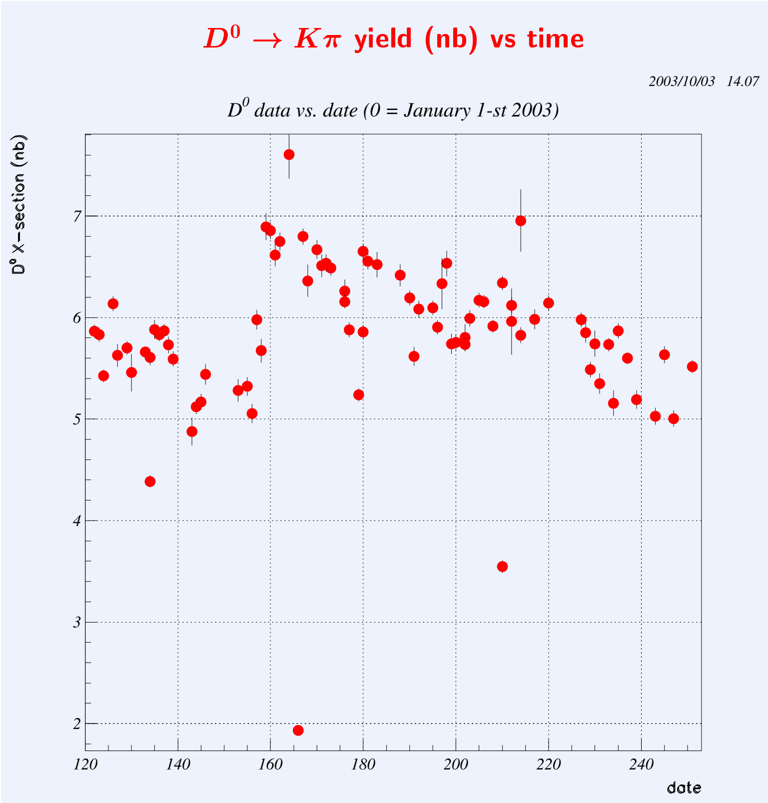

Figure 10: $D^0 \to K\pi$ yield vs. time from [50]. Day 0 is January 1, 2003.

D_s Mode	Fit Event Yield	Yield (fb)
$\phi\pi$	1390±50	7514±270
K*K	1423±271	7692±1465

Table 4: Summary of measured, pre-shutdown, yields in $B_s^0 \rightarrow \ell D_s X$ from [51]

.

the semileptonic modes *should* have an impact on the first CDF mixing measurement.

The selection of a semileptonic sample has been documented [51]. The selection requires an electron or muon and a reconstructed D meson from the lepton+SVT trigger path. Trigger confirmation is required, thus eliminating volunteers. As for the hadronic modes, this may not be necessary for the mixing measurement. The selection is optimized by maximizing the mixing significance shown in Equation 4, with respect to a variety of decay kinematic and candidate quality variables. For this optimization ϵD^2 is assumed constant and Δm=15 ps^{-1}. The *ct* resolution is taken as the quadrature sum of the L_{xy} and K-factor resolutions multiplied by the proper time itself. The details are described in [51].

Figure 11 shows the reconstructed $D_s \rightarrow \phi\pi$ peak in $B_s^0 \rightarrow \ell D_s X$ candidates. The signal is also reconstructed in other D_s decay modes, of which only K^*K has a significant yield. The semi-leptonic yields are summarized in Table 4.

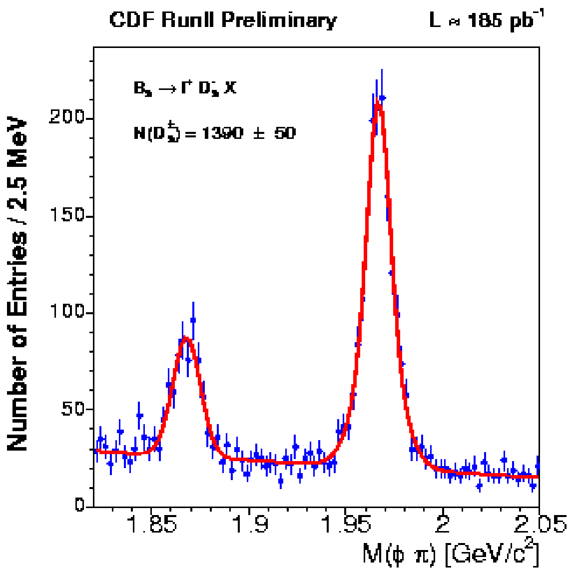

Figure 11: Reconstructed $D_s \rightarrow \phi\pi$ mass spectrum in ℓD_s events from [51].

A recent preliminary study [52] of the mixing sensitivity of the semileptonic sample indicates that with about 250 pb^{-1} CDF could approach the world's best single-experiment limit of 10.9 ps^{-1}. The study uses the measured semileptonic yields, but assumes tagging power in excess of what we have achieved. Nevertheless, it does show that while the Δm reach is not great, and a mixing measurement is not yet at hand, a meaningful result from CDF might be achievable with the current dataset.

The semileptonic reach might be improved by considering semileptonic D_s decays, which provided a significant yield enhancement (with worse K-factor resolution) in Run 1. One could also consider these decays in the hadronic analysis.

5 Particle identification

Same-side and opposite side flavor tagging relies heavily upon identifying charged kaons using the time of flight from the interaction point to the TOF detector, and/or the measured dE/dx within the COT. The contributions to ϵD^2 from these techniques arise from the efficiency with which the TOF and dE/dx measurements yield a particle ID, and the rate at which particles (principally pions) are mis-identified as kaons. We expect the latter effect to be relatively more important than the efficiency since the dilution enters as the square.

In the next sections, we will discuss the current state of these measurements as they pertain to flavor tagging, the limitations of the techniques and the ongoing work to improve them.

5.1 Time-of-flight

The TOF detector is designed to measure the time-of-flight of particles with a resolution of 100 ps over along a trajectory that is typically about 150 cm long. Figure 12 shows the time differences between various particle species as a function of momentum for a typical trajectory. Assuming the design resolution, we expect at least two sigma $K-\pi$ separation up to a momentum of about 1.6 GeV/c, where the difference in the TOF of the two species is 200 ps.

The time resolution obtained in practice is a function of the available calibrations, reconstruction algorithms and detector occupancy. The legacy reconstruction, that which is available in release 5.3.0 and earlier, achieves a TOF resolution of about 110 ps for a large fraction of TOF measurements. Non-Gaussian tails in the time resolution function, however, degrade the effective 2σ time to about 470 ps [53] from the design value of 200 ps. From Figure 12, we see that this effective time resolution amounts to a reduction in the maximum momentum at which to expect 2σ $K-\pi$ separation to about 1.0 GeV/c. Efforts are now under way within the TOF reconstruction group to improve the effective time resolution.

A second issue that adversely affects the performance of the TOF is a larger than expected occupancy. The observed average channel occupancy is in the range of 25%

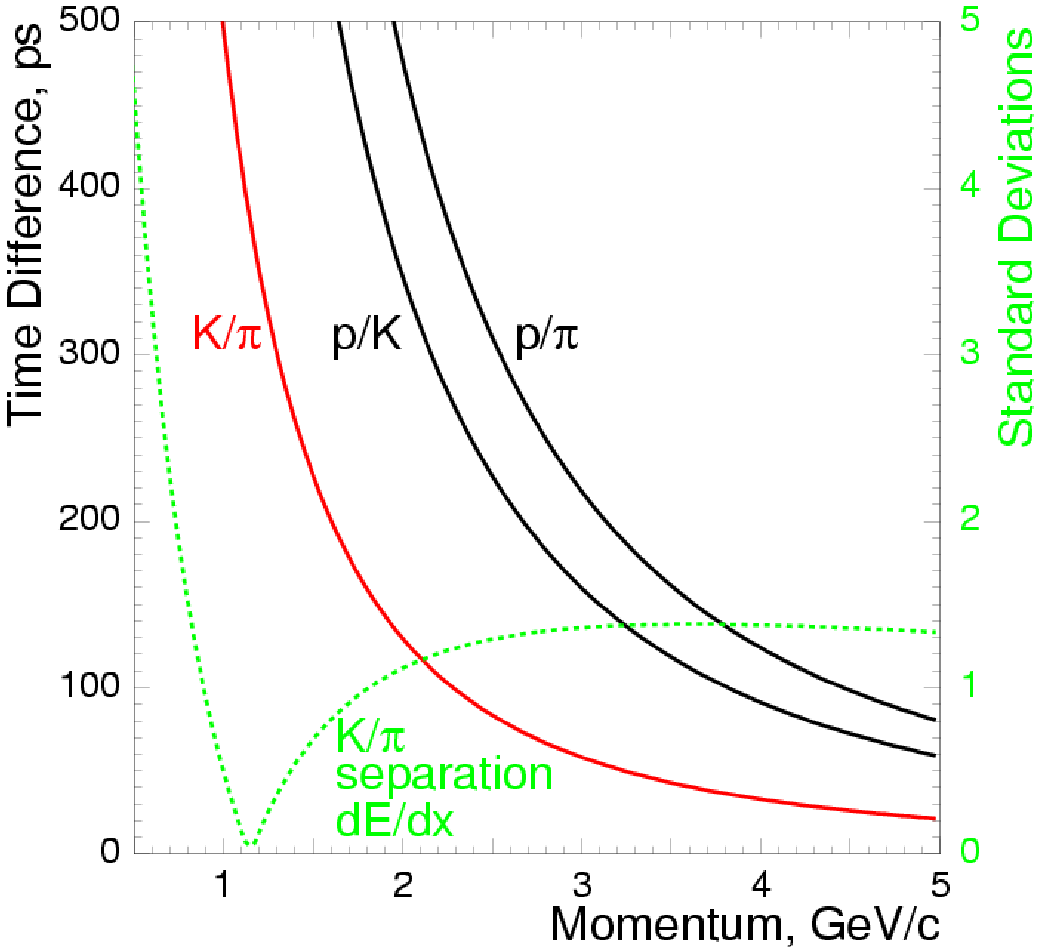

Figure 12: The time difference (ps) between the TOF for various particle species as a function of the particle momentum. The difference is for particles traveling in a plan perpendicular to the beam. The scale on the right corresponds to the number of standard deviations represented by the time difference assuming a time resolution of 100 ps. For comparison, the green dashed curve shows the design $K - \pi$ separation using the COT dE/dx in units of the number of standard deviations.

to 30% (see discussion in Sect. 5.1.3), compared to an expected value of about 12% in a typical B sample at luminosities of 5×10^{31} cm^{-2}s^{-1}[54]. As the occupancy becomes larger, the incidence of multiple hits in a readout channel increases, an effect which compromises the timing information for these channels.

In the following sections, we will summarize the salient features of the TOF detector, the reconstruction, performance metrics and the status of efforts to improve the performance of reconstructed TOF measurements.

5.1.1 TOF detector and reconstruction

Details of the TOF detector and electronics can be found in Refs. [55] and [56]. Briefly, the TOF detector consists of 216 scintillator bars, each about 2.8 m long and 4×4 cm in cross section, located at the outer diameter of the COT. Both ends of each bar are instrumented with photo-multiplier tubes (PMTs). Signals from each PMT are split between a leading edge discriminator, set to a few percent of the pulse height from a minimum ionizing particle, and a voltage-to-current circuit. Charge from the pulses is integrated within a fixed-width gate about 10 ns wide that is triggered by the discriminator. This gated charge is then digitized. The discriminator also triggers a Time-to-Amplitude-Converter (TAC). Times are measured by digitizing the TAC output voltage at a common stop that is synchronized to the beam crossing.

The raw times from the detector are corrected for non-linearities in the TAC response, and channel-to-channel differences arising from variations in the electronics, variations in cable lengths, etc. The pulse heights are corrected for a time-dependent pedestal. Once the position of a hit along the bar is known, the time can be further corrected for the propagation time in the bar and slewing effects in the discriminator. The latter correction can be as large as several nanoseconds, and is therefore one of the most critical calibrations to be applied. The pulse height can be corrected for attenuation length in the bar, as well as light collection and reflection effects near each end. Differences in channel-to-channel gain, which can be as large as a factor of four, are also measured and used in the reconstruction.

The function of the TOF reconstruction is to extract the arrival times and event t_0 values necessary to calculate the TOF for reconstructed tracks. To obtain an arrival time, the reconstruction extrapolates a track to determine the scintillator bar (or bars) it hits, and then determines whether the bar contains accurate timing information. Typically, accurate times can be extracted only from bars that are struck by exactly one particle; any other signal deposition in the bar can lead to potentially large errors in the reconstructed time. Defining matching criteria and understanding the resulting trade-off between matching efficiency and timing purity (*i.e.*, dilution) is one of the central tasks of the reconstruction.

The event t_0 is calculated from a sub-set of the tracks determined to have good arrival times. The legacy reconstruction uses all tracks that pass a loose set of quality criteria that include a minimum number of axial and stereo hits, a minimum number of hits in the outer super layer and impact parameter cuts. Notably, no vertex requirement

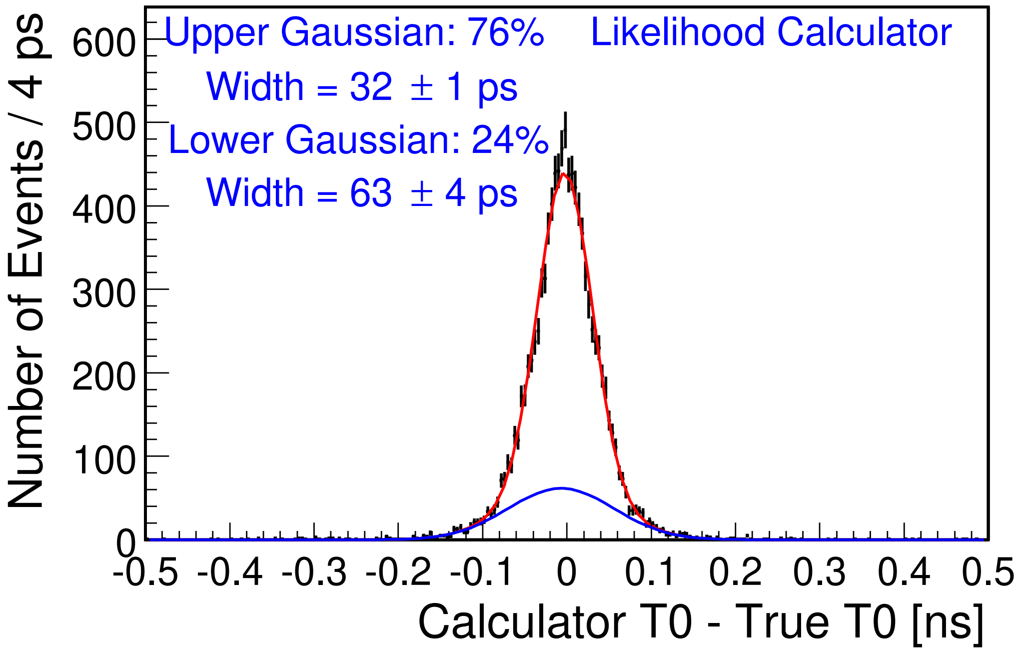

Figure 13: The difference in the fitted and the true t_0 for MC tracks with Gaussian distributed measurement errors with 100 ps time resolution (from Ref. [57]). A double Gaussian fit is applied in order to better account for the tails. The effective 2σ time is about 80 ps.

is imposed. The t_0 is calculated using a negative log likelihood fit that simultaneously takes into account pion, kaon and proton particle ID's for each track. These particles are assumed to occur with probabilities of 80%, 10% and 10%, respectively. Figure 13 shows the difference between the actual and fitted t_0 value for simulated data. Such studies show that the fitter performs well when provided input data with Gaussian distributed errors, and in fact produces a nearly efficient estimator of the event t_0 [57].

5.1.2 TOF performance measures

The performance of the arrival time reconstruction is quantified in part by the time resolution function, which can be measured by comparing the t_0 for well-identified tracks from a common vertex:

$$\begin{aligned}\Delta t &= t_0(1) - t_0(2) \\ &= (t_1 - \frac{s_1}{c}\sqrt{1 + \frac{m_1^2}{p_1^2}}) - (t_2 - \frac{s_2}{c}\sqrt{1 + \frac{m_2^2}{p_2^2}}) \qquad (5)\end{aligned}$$

where $t_0(i)$ is the estimated t_0, t_i is the arrival time, p_i the momentum and s_i the path-length from the common vertex to the TOF bar for track i. Typical methods for extracting this function include using the electron tracks in $Z \to ee$ events, or muon tracks in $J/\psi \to \mu\mu$ events.

Figure 14 shows the Δt distribution for a sample of J/ψ events from the legacy reconstruction. The central core of this distribution is a Gaussian with a width equal to $\sqrt{2}\sigma$, where σ is the intrinsic resolution of the system. Times obtained from bars with signals from more than one particle populate the tails of the distribution.

We can use the Δt distribution to define both the efficiency and purity of the reconstruction algorithm. In one such method, we define the track pair efficiency as the number of events in which both tracks have matches divided by the total number of events in which both muons are fiducial to the TOF. The purity can be defined as the number of events in the central Gaussian of the Δt distribution divided by the total number of events with matches. Figure 15 shows the efficiency versus purity for the legacy algorithm. Improvements to the reconstruction by definition move the curve to the upper right.

A similar time resolution function can be defined for the event t_0. For each event, we calculate the time difference between the calculated t_0 and that estimated from a single well identified probe track:

$$\begin{aligned}\Delta t_0 &= t_0(\text{event}) - t_0(\text{track}) \\ &= t_0(\text{event}) - (t_a - \frac{s}{c}\sqrt{1 + \frac{m^2}{p^2}}) \end{aligned} \tag{6}$$

where t_a is the measured arrival time, m the mass and p the momentum of the probe track, and s is the path length from the event vertex to the point at which the probe track enters the TOF. The probe track is not used in the t_0 fit. Samples in which a clean measurement of the Δt_0 distribution can be measured include include $W \to e\nu$ and $J/\psi \to \mu\mu$ decays. Figure 16 shows Δt_0 distribution measured in J/ψ events. The central core of the distribution is dominated by the arrival time resolution with a much smaller contribution from the t_0 resolution. The tails arise from large errors in the arrival time or t_0 estimates, but are dominated by t_0 errors. (The arrival time contribution can be significantly reduced by requiring agreement between the two muon tracks.)

The Δt_0 distribution can be used to define an efficiency and purity in much the same way as for the arrival time distribution. No such plot is yet available, although it will be the fundamental tool by which the TOF reconstruction group assesses improvements to the t_0 algorithm.

The J/ψ sample provides other diagnostics beyond those mentioned above. By using the two muon tracks and the event t_0 in combination, for instance, it is possible to determine which of the two muon tracks was mis-measured, and therefore whether the track was measured too early or too late.

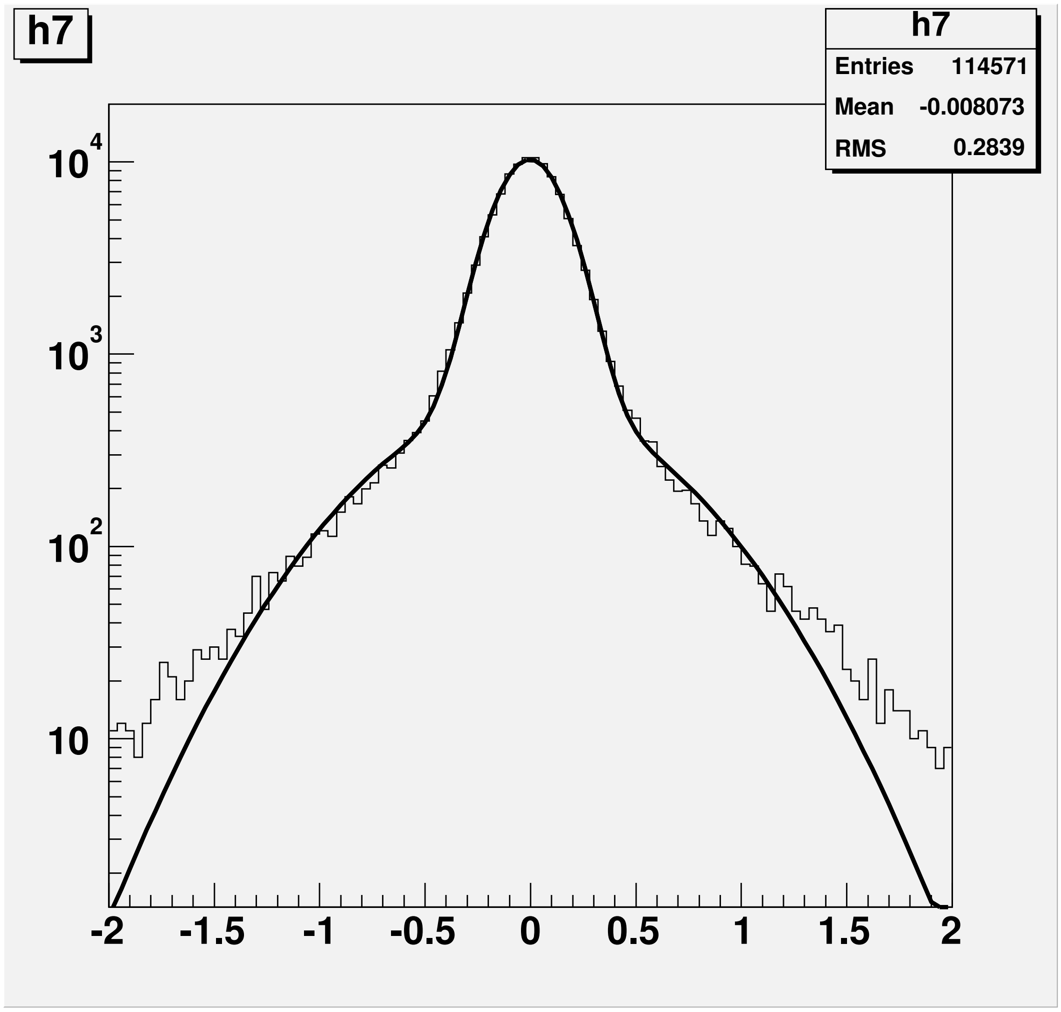

Figure 14: Δt (ns) from Eq. 5 for the two muon tracks in a sample of J/ψ events obtained from the legacy reconstruction. The curve is a double Gaussian fit where the narrow Gaussian has a width of 160 ps, and the broad one a width of 550 ps.

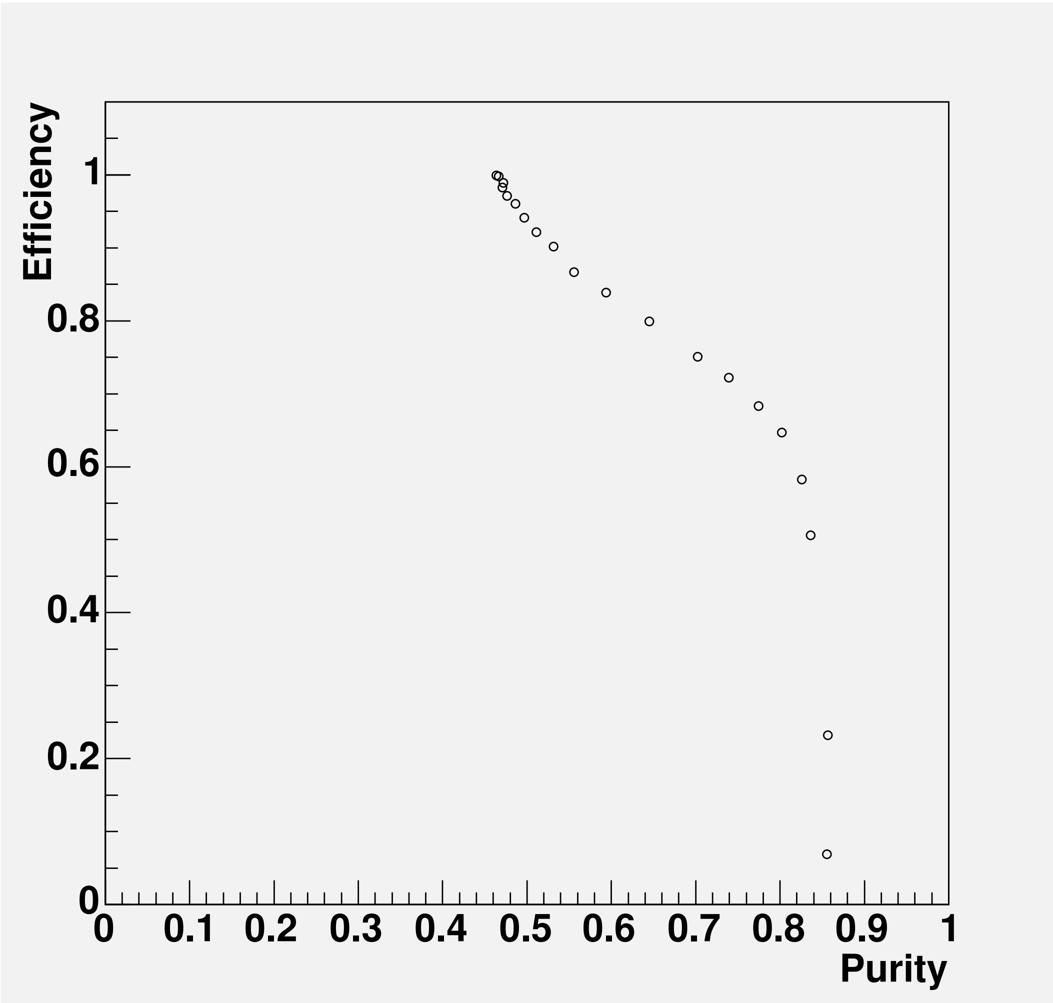

Figure 15: Efficiency versus purity for muons in J/ψ events obtained using the legacy reconstruction.

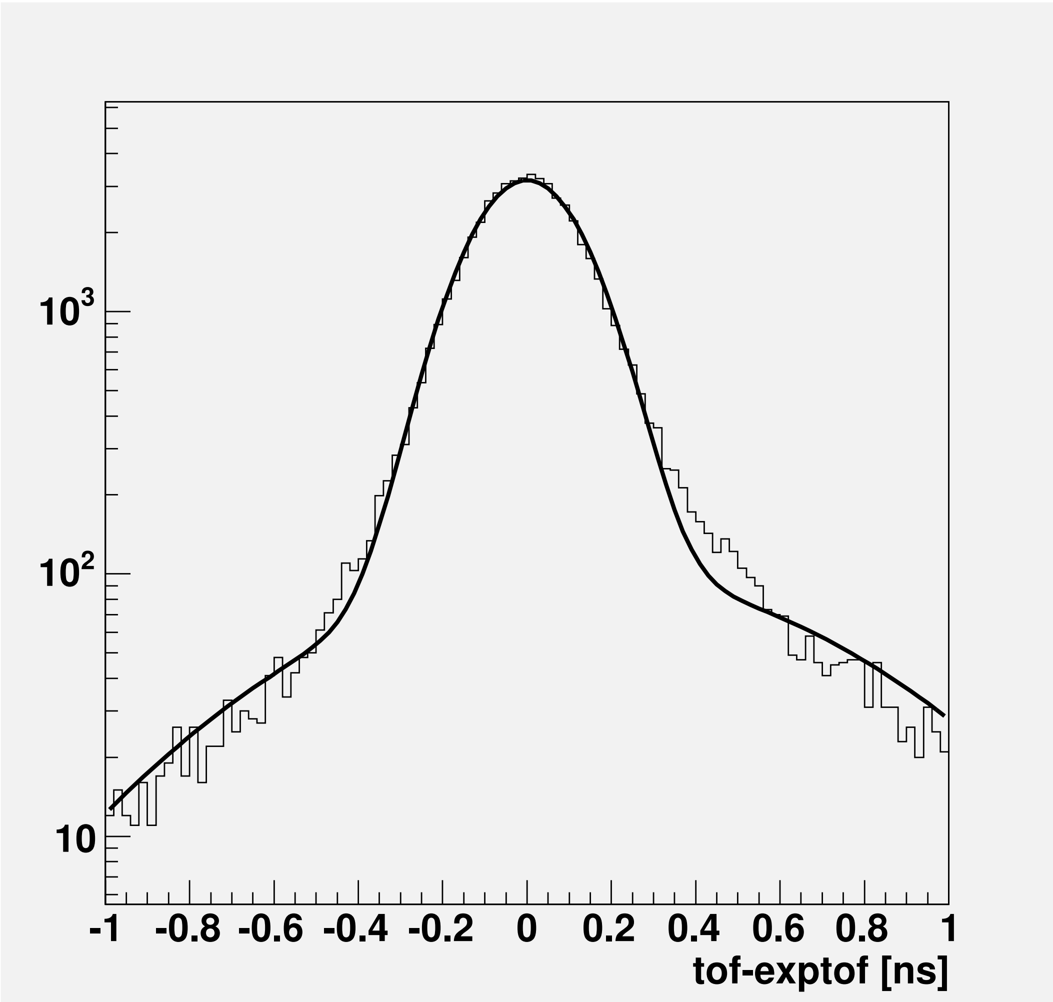

Figure 16: The Δt_0 (ns) from Equation 6 observed in J/ψ events using the legacy reconstruction. The asymmetry arises because arrival time errors more often produce late times than early times due to slewing effects from multiple particles in a bar, and the t_0 fit more often mis-identifies pions as protons than visa-versa.

5.1.3 Status of TOF reconstruction

The major limitations to the performance of the TOF system are the tails in the arrival time and event t_0 resolution functions. Results from MC suggest that the tails in the t_0 resolution function arise in large measure from the tails in arrival time resolution function. The top priority for the TOF reconstruction group is to understand the source of the tails in the arrival time resolution function and improve the discriminating power of the matching algorithm.

Besides explicitly requiring that no more than one track point to a bar, the legacy reconstruction uses the time difference between the two PMTs on a bar to establish a good match. Figure 17 shows the efficiency versus purity for several different matching criteria in addition to the legacy algorithm. The right-most curve, representing the best matching algorithm among those shown, is obtained using a chi-squared cut, where the chi-squared is defined by the time difference normalized by the uncertainty in the difference, and the normalized pulse height difference.[3] The improvement in purity results from almost a factor of two decrease in the tails of the arrival time resolution function relative to that from the legacy reconstruction for some values of the chi-squared cut. Additional discriminants are under investigation, such as the isolation of a bar in the COT (eg., the number of hits in the outer layer of the COT that are associated with the bar beyond those belonging to the extrapolated track).

Studies using simulated data demonstrate that high purity tracks are critical to the performance of the t_0 fitter. Eliminating tracks that do not originate from the primary interaction vertex or tracks that may be poorly measured is therefore the most important task for calculating the t_0. The legacy reconstruction associated neither tracks nor t_0 values with vertices, which was a serious structural design flaw. This defect of the legacy reconstruction has been corrected. The tails in the t_0 resolution function drop dramatically after this change alone, as seen in Figure 18. Another important structural improvement recently introduced is a user-level interface to the t_0 fitter. This feature provides complete control over the tracks used in the t_0 calculation within the context of a given analysis, allowing, for instance, a user to calculate unbiased TOF values for an arbitrary set of tracks by excluding the desired tracks from the t_0 fit. A user can also use this feature to choose, if necessary, the precise efficiency vs. purity trade-off based upon the needs of a particular analysis.

Poorly measured tracks will enter the t_0 calculation despite the addition of tight quality and vertex association cuts. The event contains additional information, however, that can be brought to bear on the problem. First, the dE/dx of tracks can be used to modify the assumed particle ID probabilities assumed in the likelihood fit. Second, the CLC detector can also be used to extract an event t_0 with a precision of the same order of magnitude of the TOF t_0. The COT can also be used to calculate the event t_0, although with a precision considerably worse that that from the TOF.

[3] Note that the current calibration scheme for the TOF detector does not include calibration of the pulse height. Use of the chi-squared matching scheme therefore requires a considerable extension to the calibration database. The TOF calibration group is actively pursuing adding this functionality.

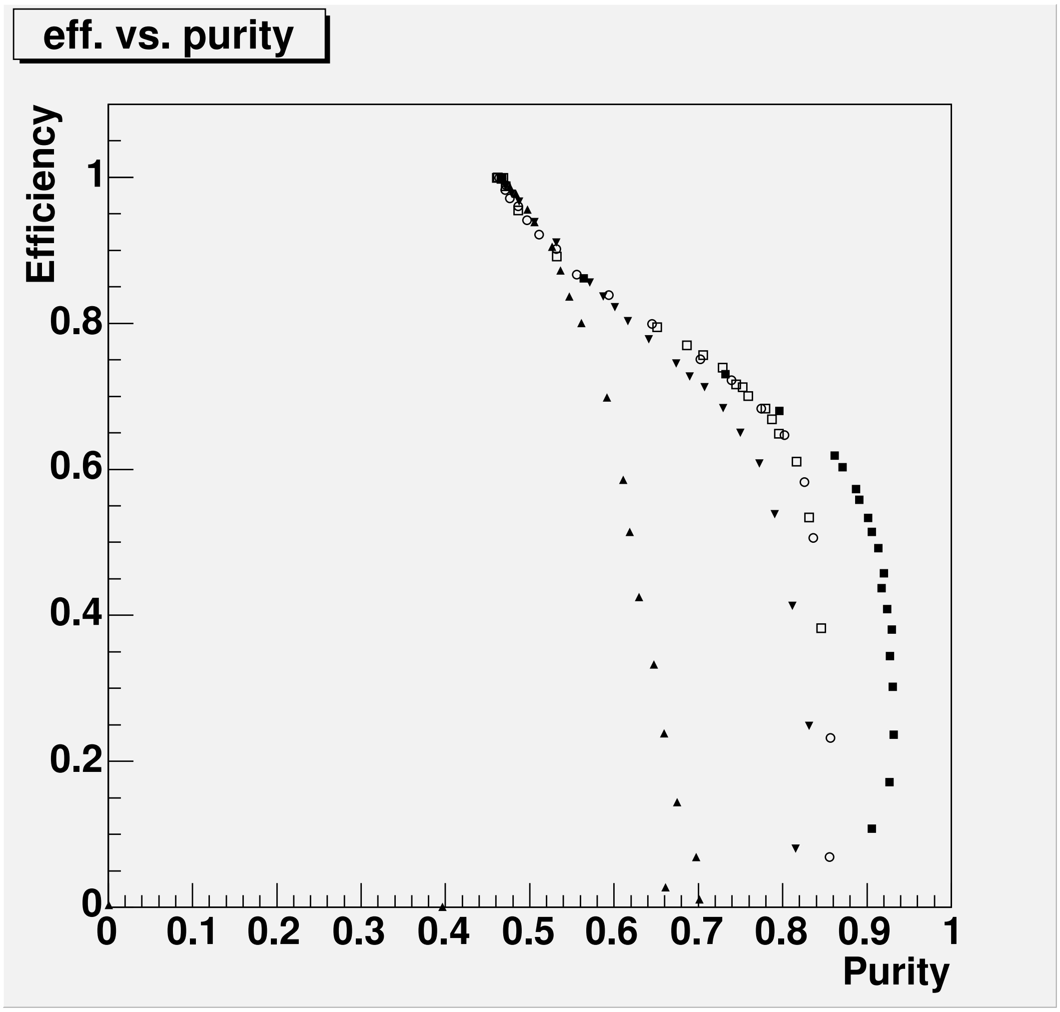

Figure 17: The efficiency vs. purity for four different matching criteria. The left-most curve is the result using the charge difference as the matching criterion; the next two curves to the right, the normalized charge and normalized time differences, respectively (the latter being almost the same as the legacy reconstruction); and the right-most, the chi-squared cut. The apparent decreases in purity as cuts are tightened in some cases are likely caused by the low statistics of these particular data points.

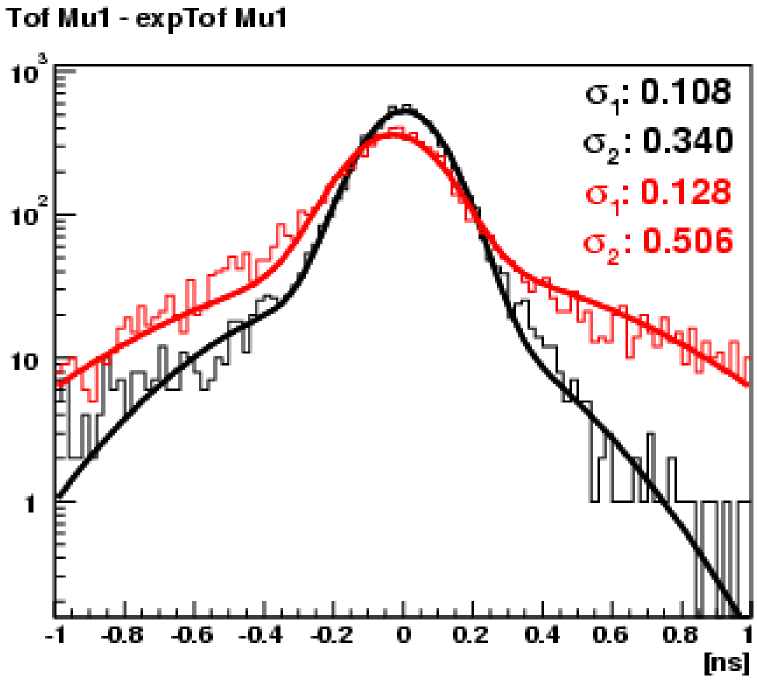

Figure 18: The t_0 resolution function (ns) from Eq. 6 for J/ψ events using the legacy reconstruction (broad, red curve) and for the new implementation (narrower, black curve) that calculates the t_0 for each vertex in the event using only tracks that belong to the particular vertex. A substantial reduction in the tails is observed from this change alone.

It can nonetheless be used to constrain or cross check the value obtained by the TOF algorithm. Finally, the event t_0 itself provides a constraint on the individual tracks that contribute to the t_0 calculation. Outliers in the unbiased residual, for instance, can be used to identify and eliminate tracks that are likely mis-measured. All of these avenues are under investigation by the TOF reconstruction group.

The final area of concern to the TOF performance is the high occupancy and the effect on the efficiency and purity of TOF measurements. Figure 19 shows the fraction of the occupancy that is associated with tracks hitting TOF bars, with segments in the outer SL of the COT that are not associated with tracks but extrapolate to TOF bars, and bars that have neither tracks nor segments pointing to them. Almost one half of the occupancy cannot be associated with tracks or segments in the COT. The TOF reconstruction group is attempting to understand the origin of this activity in the detector. Ultimately, losses to high occupancy may be irreducible unless a reliable scheme for using a single PMT can be developed. No effort has yet been directed at this task.

5.2 dE/dx measurements

The energy loss of high velocity charged particles traversing a material is dominated by momentum transfers to the electrons in the material. The average energy loss per unit length under such circumstances is given by the well known Bethe-Block formula:

$$\frac{dE}{dx} = \frac{4\pi N e^4}{mc^2\beta^2} z^2 (\ln \frac{2mc^2\beta^2\gamma^2}{I} - \beta^2 - \frac{\delta(\beta)}{2}),$$

where N is the number density of electrons, m is the electron mass, z is the incident particle charge, I is the mean excitation energy of the atoms in the medium, β is the velocity divided by c, $\gamma = 1/(1-\beta^2)$ and δ is a correction for high values of β within dense materials (not relevant for the gas in drift chambers). In principle, the charge collected by the COT along the trajectory of a reconstructed particle track can be used to measure dE/dx. Given the particle momentum, we can then use the measured dE/dx to estimate the value of $\beta\gamma$ and thereby obtain an estimate of the particle mass.

In practice, the ionization in a drift chamber is subject to large fluctuations due to a logarithmic divergence in the electron scattering cross section [58]. These fluctuations lead to large errors when the estimated average dE/dx is based upon the measurement of total ionization in a drift chamber. The conventional method used to reduce the measurement errors is to calculate a truncated mean in which some fraction of the measurements, typically around 20%, for a given trajectory are ignored. A number of detector effects also contribute significant systematic errors to the measured energy deposition. These effects will be described in the next section.

The performance of the dE/dx reconstruction is quantified primarily using two measures. The first is the statistical separation power between two particle species, such as the $K-\pi$ separation noted above. Improvements in the dE/dx appear as an increase in the separation power. The second is the dE/dx correlation between particles

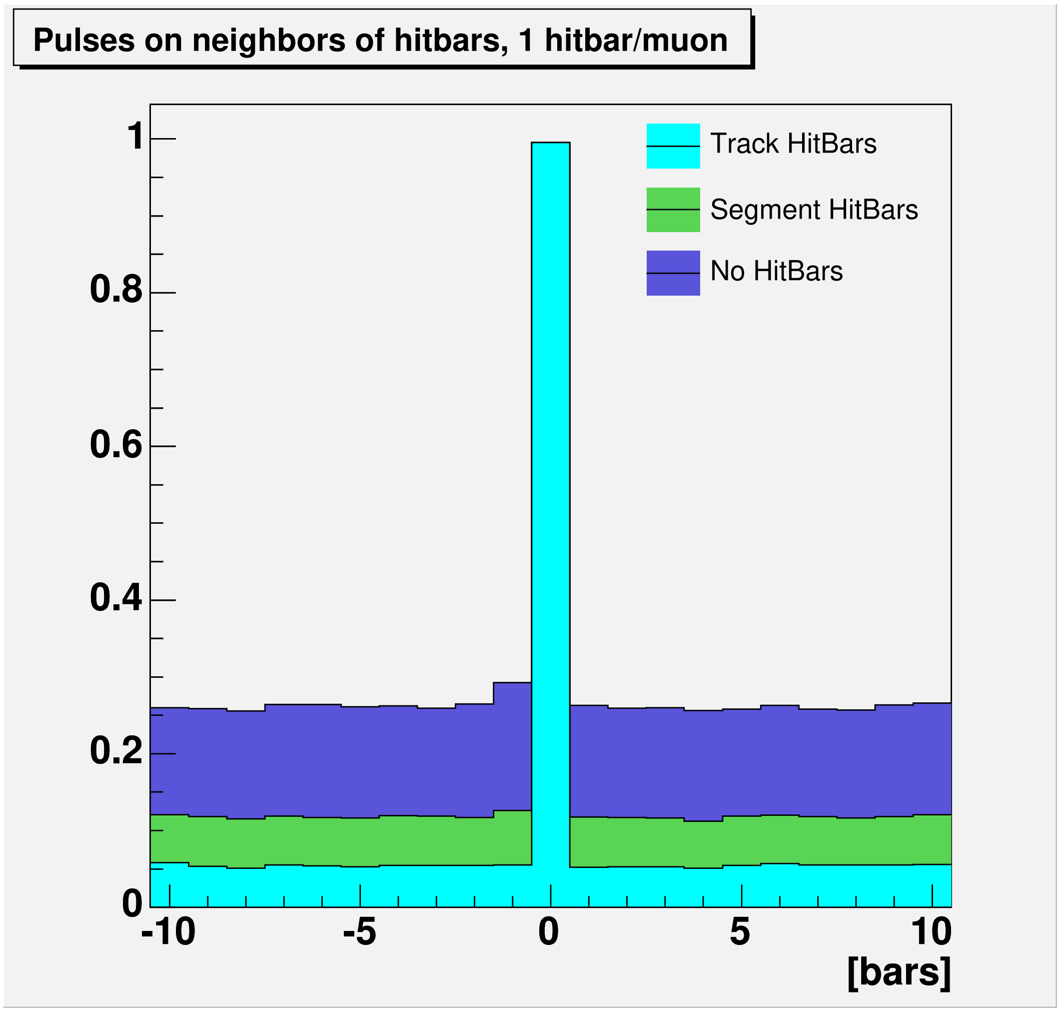

Figure 19: The TOF occupancy within ± 10 bars of the bar hit by a J/ψ muon in cases where the muon hits exactly one bar. For each case where a neighbor bar has two pulses, the occupancy is classified according to whether there is another track pointing to the bar, a segment in the outer layer of the COT that points to the bar but does not belong to a track, or neither segments nor tracks point to the bar. Cross-talk is a negligible contribution to the neighbor occupancy.

in the same event. Ideally, there are no correlations between the dE/dx measured for different tracks in the same event. Non-trivial correlations indicate the presence of significant uncorrected detector or physics effects.

5.2.1 The COT detector and dE/dx reconstruction

Details of the dE/dx measurement in the COT are discussed in [59]. Briefly, the charge deposited by dE/dx energy losses in the COT gas is collected on 96 layers of sense wires. Analog pulses from the front-end electronics on each sense wire are converted to a digital pulse by a leading-edge discriminator, where the time over threshold is related to the amount of charge collected. The digital pulse width is approximately linear in the amount of charge collected in a given pulse. TDC's digitize the leading edge and pulse width times.

To convert the individual raw pulse widths into charge measurements requires that several systematic effects be removed during several stages of calibrations. Some of these calibrations depend upon properties of the drift chamber and electronics: variations in high voltage, wire location within a drift cell, super-layer location and channel-to-channel electronic gain variations, etc. Other effects depend upon the properties of the particle trajectory as well: z location of the hit along the wire and various details of the trajectory through the cell such as path-length, drift distance and aspect angle. Finally, a global correction for the gas pressure is applied to all measurements.

The corrections currently available are entirely hit-level and are based upon generation 4 tracking and the first 65 pb^{-1} of data. These corrections are described in Ref. [59]. Additional track-based corrections developed using samples of XFT tracks from 200 pb^{-1} of data are described in [60]. The $K - \pi$ separation obtained using the combined hit-level and track-based calibrations is about 1σ for $p_T > 2$ GeV/c.

5.2.2 Status of dE/dx measurements in the COT

There are several significant issues and known deficiencies with the dE/dx corrections described in the previous section. First, the fully corrected dE/dx shows a charge asymmetry (Fig. 20). The origin of this effect is not understood. Second, the dE/dx resolution function (obtained from high-purity samples of exclusive particle decays) shows a linear dependence on the number of hits rather than the expected $1/\sqrt{N}$ dependence, suggesting the presence of residual effects that are not properly taken into account. Finally, and most significantly, changes in the performance of the COT over time have introduced time-dependent effects in azimuth (ϕ) (Figure 21), polar angle and hit multiplicity, among others, that are taken into account. Correcting for these effects requires that certain aspects of the original calibration framework to be re-visited.

To address these issues, the dE/dx group has undertaken and recently concluded a detailed study of track-based dE/dx corrections using the first 200 pb^{-1} of data and Generation 4 tracking [61]. These studies have not only confirmed effects observed in previous studies, but also allowed the measurement of effects and correlations not

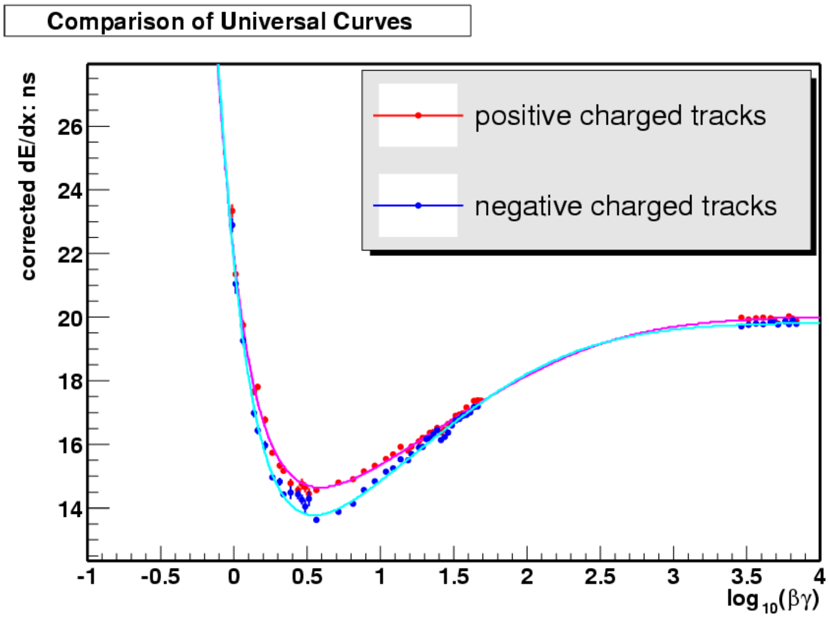

Figure 20: Comparison of universal curve for measured dE/dx for positive and negative tracks. The curves are measured from from samples of $\Lambda \to p\pi$, $K_s \to \pi\pi$ and $J/\psi \to \mu\mu$ decays and $\gamma \to ee$ tracks [59].

previously observed in the earlier, smaller data samples, such as a significant drift-time dependent correction. Figure 22 shows the residuals to the expected average dE/dx as a function of several variables for which corrections are applied. New correlations with the hit multiplicity have been discovered in the dE/dx resolution function. Correcting for these effects has improved the the $K-\pi$ separation to 1.4σ for $p_T > 2$ GeV/c. Track-to-track correlations within an event have been successfully exploited as a powerful tool to understand the extent of uncorrected effects that remain in the data. The new corrections reduce these track-to-track correlations by about a factor of two.

Despite these advances, there remain a number of significant issues that need to be resolved in order to make optimal use of dE/dx. All the existing track-level corrections, for instance, are based upon samples of XFT tracks [59][60]. As a result, the dE/dx measured for tracks with p_T below 2 GeV/c exhibit significant departures from the universal curves based upon the new calibrations (see Figure 23) [62]. Work is now under way to remedy this problem.

A second, more pervasive problem stems from the nature of track-based corrections themselves. These corrections necessarily introduce sample dependence of the correction values due to such issues as background subtraction and sample bias. The magnitude of all track-based corrections to date are large relative to the track independent, hit-level corrections. Many time-dependent effects, for instance, are treated at the track-level while in principle they could be handled at the hit-level. This observation has led to speculation that the remaining residual effects and sample dependence

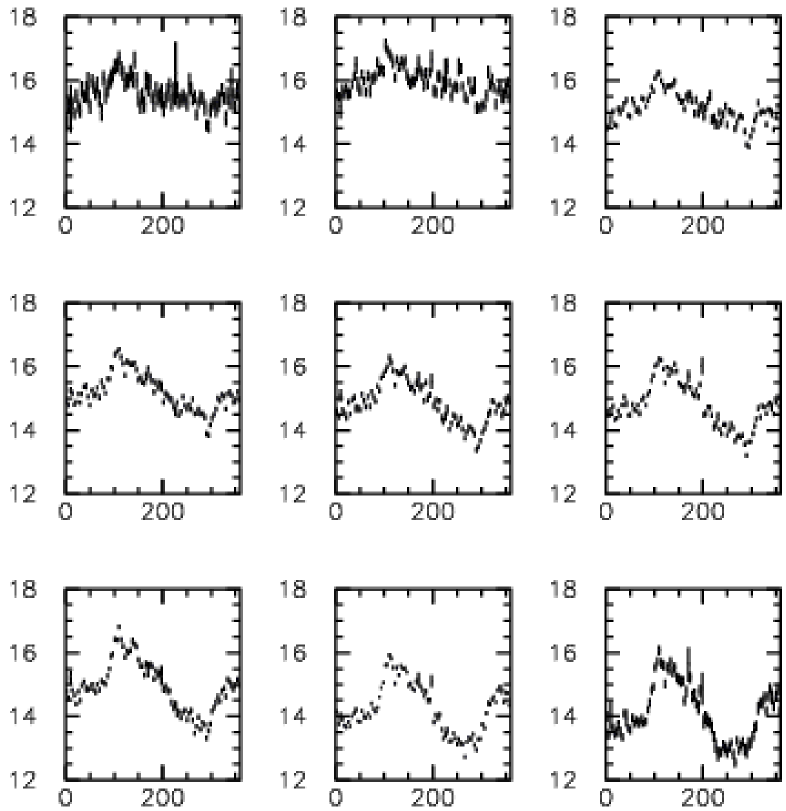

Figure 21: Average dE/dx as a function of ϕ (deg) of the track for positive kaons and pions. Each plot corresponds to a specific run range. The plots are ordered in time from left to right, then top to bottom. A strong time dependence in azimuthal response of the detector is evident. (Figure from [61].)

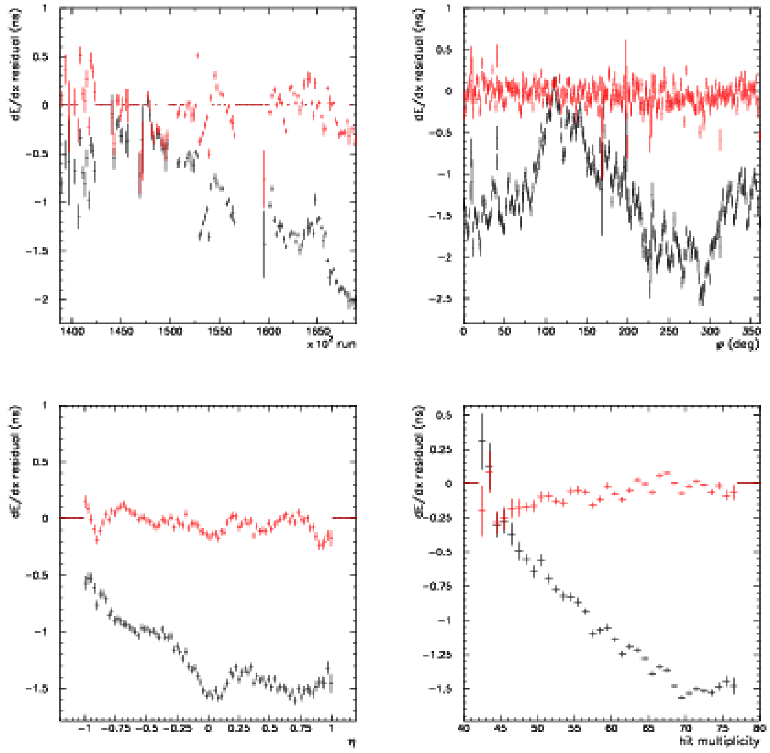

Figure 22: Residuals between the measured and expected average dE/dx as a function of run number (top left), ϕ (top right), pseudo-rapidity (η, bottom left) and hit multiplicity (bottom left). Black (red) dots are the residuals before (after) corrections are applied. (From [61].)

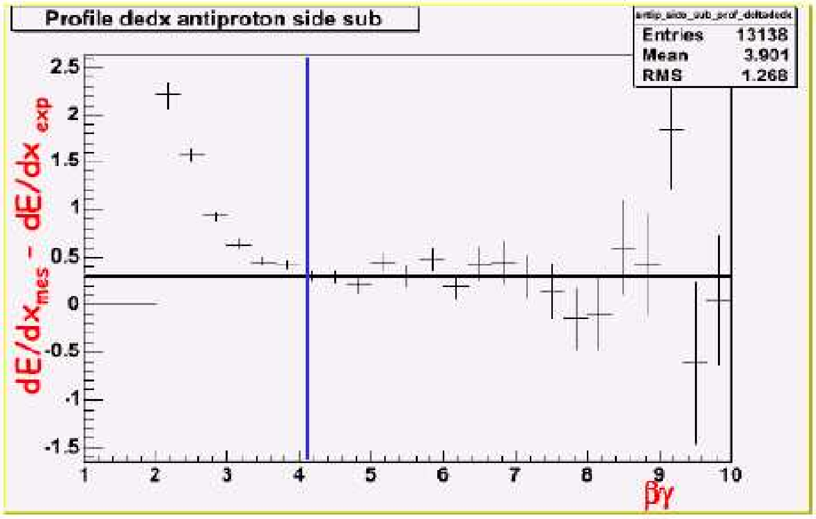

Figure 23: The difference between the measured and expected average dE/dx for a sample of anti-protons. The vertical blue line corresponds to the lower bound of the momentum range covered by the XFT-based corrections.

may be greatly reduced by implementing an improved framework of time-dependent hit-level corrections. Many of the time-dependent effects in this scheme may be handled online using Stage 0. The resulting track-based corrections are expected to be significantly smaller than the current set. Track-to-track correlations should continue to be useful in probing the presence of residual effects, and will be used to evaluate the success of this approach. The dE/dx group considers this issue to have a high priority, and hopes to have results ready in time for summer of 2004.

The task of identifying and prioritizing issues related to dE/dx measurement in the COT is being coordinated by the dE/dx group. Several individuals have committed significant effort toward advancing the agenda of the the group.

The committee notes that the silicon detectors can in principle make a significant measurement of dE/dx in the $1/\beta^2$ region before the density effect eliminates all dependence of the ionization on $\beta\gamma$. Although no presentations on this topic were made to the committee, we note that there are efforts in progress that are attempting to use the silicon detectors to measure dE/dx. Although significant obstacles remain before this work demonstrates that the silicon detectors can effectively contribute to particle ID, the committee encourages continued work in this direction.

6 B Hadron Flavor Identification

One of the components of measuring neutral B meson flavor oscillations is identifying whether the B meson was produced as a B, which contains $\bar{b}$ antiquark, or a $\bar{B}$, which contains a b quark. We refer to this B hadron flavor identification as "b flavor tagging."

At a hadron collider or a continuum e^+e^- collider operating above the $\Upsilon(4S)$ resonance the methods of b flavor tagging may be classified into two categories: (1) opposite-side flavor tags (OST) and (2) same-side flavor tags (SST). The dominant production mechanisms of b quarks in hadron collisions produce $b\bar{b}$ pairs. Opposite-side flavor tags exploit this fact: to identify the production flavor of the B hadron of interest (*e.g.*, a $B_s^0(\bar{b}s)$ that eventually decays to $D_s^-\pi^+$), we identify the flavor of the second B hadron in the event and *infer* the flavor at production of the first B. Same-side flavor tags are based on the charge of particles produced in association with the production of the B hadron of interest. These two approaches to b flavor identification are described below.

The performance of the b flavor tags may be quantified conveniently by their efficiency ϵ and their dilution D. The efficiency is the fraction of B candidates to which the flavor tag can be applied ($0 < \epsilon < 1$). The dilution is related to the probability $\mathcal{P}$ that the tag is correct: $D = 2\mathcal{P} - 1$, so a perfect tag has $D = 1$, and a random tag has $D = 0$.

The measurement of B flavor oscillations involves examining an asymmetry:

$$\mathcal{A}_{\text{meas}} = \frac{N_{\text{unmixed}} - N_{\text{mixed}}}{N_{\text{unmixed}} + N_{\text{mixed}}}, \tag{7}$$

where N_{mixed} (N_{unmixed}) is the number of (background subtracted) B hadrons that are classified as having the opposite (same) flavor at decay as the flavor they had at production. This asymmetry can be measured as a function of the decay time t in the rest frame of the B hadron, and the true asymmetry is equal to $\cos(\Delta mt)$, where Δm is the mass difference between the two weak eigenstates (usually referred to as the "heavy" and the "light" states) of the neutral B hadron. Since the lifetime of B hadrons is on the order of picoseconds, we usually quote the mass difference in units of inverse picoseconds. The misclassification of the B meson production flavor reduces the true asymmetry by a factor of D, the dilution:

$$\mathcal{A}_{\text{meas}}(t) = D\cos(\Delta mt). \tag{8}$$

Measured CP asymmetries of neutral B mesons to CP eigenstates are reduced by this same factor due to misclassification of production flavor. The statistical error on the true asymmetry $\mathcal{A}$ is

$$\delta\mathcal{A} = \sqrt{\frac{1 - D^2\mathcal{A}^2}{\epsilon D^2 N}} \approx \sqrt{\frac{1}{\epsilon D^2 N}} \qquad (\text{for} D^2\mathcal{A} \ll 1), \tag{9}$$

where N is the total signal before applying the flavor tag, and the efficiency is the fraction of these N signal to which the flavor tag can be applied. The statistical power of the data sample scales with ϵD^2. At a hadron collider, a flavor tag with $\epsilon D^2 \approx 1\%$ is respectable. In comparison, the B factory experiments have achieved overall tagging efficiencies of $\epsilon D^2 \approx 30\%$.

6.1 Performance of Opposite Side Flavor Tags in Run I

In Run I, CDF published physics papers using two types of opposite-side flavor tags: (1) the soft lepton tag (SLT) and (2) the jet charge tag (JQT). The soft lepton tag is based on semileptonic B hadron decays. The charge of the lepton signals the flavor of the b quark: negatively (positively) charged leptons come from primary b quark ($\bar{b}$ antiquark) decay. Misidentification of the B flavor can arise from detecting a lepton from the sequential decay of charm hadrons produced in the B hadron decay. About 70% to 80% of these sequential leptons have the "wrong" charge, *i.e.,* opposite the expected charge from primary B decay. In addition, lepton misidentification contributes to soft lepton flavor mistags. The jet charge tag is based on the momentum weighted sum of the charges of tracks associated with a jet. The jet charge $Q_{\rm jet}$ in [63] is calculated using a p_L weighting scheme:

$$Q_{\rm jet} = \frac{\sum_{\rm tracks} q_i(\vec{p}_i \cdot \vec{p}_{\rm jet})}{\sum_{\rm tracks}(\vec{p}_i \cdot \vec{p}_{\rm jet})}.$$

On average $Q_{\rm jet}$ from the charged particles associated with a jet produced by a b quark ($\bar{b}$ antiquark) is negative (positive). Flavor misidentification occurs because the jet charge only works on average and because the jet selected may not be the other b jet produced in the collision. In fact, the efficiency of opposite side flavor tagging techniques is limited by the acceptance of this other b. In the central acceptance of $|\eta| < 1$, a generator level Monte Carlo study [64] finds that the maximum efficiency of any opposite side tag varies from 10% to 30%. This studied was based on the `HERWIG` and `PYTHIA` Monte Carlo generators and included higher order contributions (gluon splitting and flavor excitation) to $b\bar{b}$ production in addition to the lowest order tree level process (flavor creation). The variation of the acceptance in the opposite side b is due to the possible variation in the p_T requirements of the trigger side b and the opposite side b. Some of the projections of b flavor tag performance for Run II described below were partially based on the `PYTHIA` Monte Carlo with only the flavor creation $b\bar{b}$ process, which overestimated the opposite side b acceptance by as much as a factor of two.

In Run I the performance of the SLT and JQT were measured [63, 65, 66] using B hadron mixing in an inclusive lepton sample. The trigger lepton was associated with a jet containing a secondary vertex that was reconstructed with the `BVTX` [67, 68] algorithm, a modified version of the high-p_T `SECVTX` algorithm that was tuned to have better efficiency for the lower p_T b's in these inclusive lepton events. The sign of the charge of the trigger lepton identified the B hadron flavor at decay, and the secondary vertex position and the measured B hadron momentum were used to reconstruct the proper decay time. The SLT and JQT flavor tagging algorithms were used to determine the flavor of the B hadron at production. The measured performance [63] of these flavor tags was

$$\epsilon D^2 = 0.91 \pm 0.10\,({\rm stat.}) \pm 0.11\,({\rm syst.})\,\%$$

for the soft lepton tag (electrons and muons combined) and

$$\epsilon D^2 = 0.78 \pm 0.12\,(\text{stat.}) \pm 0.08\,(\text{syst.})\,\%$$

for the jet charge tag.

The dilution of these two opposite side flavor tags was found to depend on various physics related quantities, and these quantities were used to increase the statistical separation power of these two flavor tags. This technique is usually referred to as "binning," and we say that the flavor tag has been "binned" in ϵD^2. For the soft leptons, the binning quantity was p_T^{rel}, the transverse momentum of the soft lepton with respect to the tracking jet that contained the lepton. Leptons from sequential charm decay ($b \to c \to \ell^+$) and fake leptons tend to have low values of p_T^{rel}, and leptons from primary B hadron decay ($b \to \ell^-$) tend to have high values of p_T^{rel}. The variation of the lepton dilution as a function of p_T^{rel} was measured in the data, and this variation was used as part of the likelihood description in the extraction of the overall soft lepton dilution and Δm_d. Isolated soft leptons that are not associated with a jet were found to have high dilution, and these leptons were included separately in the likelihood. In the case of the jet charge tag, the dilution of this tag was found to increase linearly as a function of the jet charge Q_{jet}. The slope of this linear dependence was determined from the data, and the overall dilution was extracted as part of the likelihood fit for mixing. Jet charge tags with $Q_{jet} = \pm 1$ were treated separately, just like the isolated soft leptons. Further separation power for the jet charge was achieved by separating the opposite side jets that were identified as having a vertex tag using the `BVTX` algorithm. The dilution of these jets was found to be significantly larger on average than the dilution of the generic opposite-side jet selection for the jet charge. This increase in dilution presumably came from (1) that it was more likely that the `BVTX` identified jet was from the opposite side B hadron and (2) the `BVTX` tag increased the average p_T of the jet increasing the jet charge separation power. Attempts to reproduce the performance of the SLT and JQT observed in the data using Monte Carlo were not successful. In particular the performance of the jet charge was several times better in the Monte Carlo than in the data [69, 70] Even the dilution achieved with `BVTX` identified jets in the data and the Monte Carlo disagreed [70], although the disparity in performance was reduced in this particular comparison.

Opposite side flavor tags were also used in the measurement [71] of $\sin(2\beta)$, the CP asymmetry in the production and decay chain $B^0 \to J/\psi K^0_S$ and the charge conjugate process. For the measurement of $\sin(2\beta)$, the dilution of the flavor tag must be determined in an independent sample. The same two opposite side flavor tags were used in this measurement, however, there was no binning of the soft lepton in p_T^{rel} or the jet charge in Q_{jet}. The jet charge tags were not classified into jets with and without `BVTX` vertex tags. The jet charge tag was modified, however, to increase efficiency for the lower p_T B hadrons in the dimuon trigger sample used for this CP measurement compared to the inclusive lepton sample used in the mixing measurement described above. The most important changes were (1) the jet finding algorithm based on clustering tracks in $\eta - \phi$ space was replaced with an algorithm based on invariant mass, and (2) tracks from a larger cone in $\eta-\phi$ space were included in the determination of the jet charge. The calculation of the jet charge also used the impact parameter significance of the tracks as part of the weight used in the jet charge calculation. The dilution of the SLT and the modified JQT were measured using a sample of 998 ± 51 fully reconstructed $B^\pm \to J/\psi K^\pm$ decays. (The efficiencies of these two tags were measured from the $B^0 \to J/\psi K^0_\mathrm{S}$ sample.) The measured[4] tagging performances were

$$\epsilon D^2 = (2.2 \pm 1.2)\,\%$$

for the soft lepton tag (e and μ combined) and

$$\epsilon D^2 = (2.2 \pm 1.3)\,\%$$

for the jet charge (the errors are the statistical errors only). These values agreed within statistics with the values determined in the inclusive lepton mixing analysis. A study comparing the performance of the jet charge flavor tag in events with fully reconstructed $B^\pm \to J/\psi K^\pm$ and the inclusive lepton+BVTX data sample may be found in reference [72].

6.2 Performance of Same Side Flavor Tags in Run I

Same side flavor tagging is based on the correlation between particles produced in the fragmentation of b quarks to B hadrons or in the decay of P wave B mesons (B^{**}). The principle of same side tagging is illustrated in Fig. 24. In the most simple picture, where only pseudoscalar mesons are produced directly by the fragmentation process, the following charged stable mesons are expected: a $\bar{B}^0$ is expected to be produced along with π^-, a B^- is expected to be produced with a π^+ or a K^+, and a $\bar{B}^0_s$ is expected to be produced with a K^-. Similarly, a B^0 is expected to be produced along with π^+, a B^+ is expected to be produced with a π^- or a K^-, and a B^0_s is expected to be produced with a K^+. The charge of the particles produced nearby a B meson can be used to identify the production flavor of that meson.

One of the potential advantages of same side tagging compared to opposite side tagging is the acceptance: if the B meson we want to tag is in the acceptance of the experiment, then probably the same side tag track or tracks are in the acceptance too. In contrast, as mentioned above, the opposite side B hadron is often not in the experimental acceptance. However, same side tags have a disadvantage as well: the expected charge correlation depends on which species of meson we are tagging, and we can expect that the efficiency and dilution of the same side tag will be different for $\bar{B}^0$, B^-, and B^0_s . Since the fragmentation of the b and the $\bar{b}$ quark are expected to be independent, opposite side tags are expected to have the same dilution for all B hadrons. We can calibrate an opposite side tag for measuring B^0_s flavor oscillations using a variety of B hadron decays. In this calibration we only need to know if the

[4] These total tagging efficiencies were determined from the efficiencies and dilutions reported in Table I of [71].

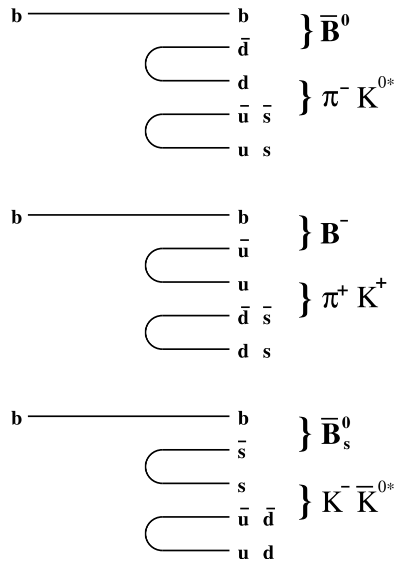

Figure 24: Illustration of some of the possible species of particles produced in the fragmentation of a b quark to a B meson.

B hadron contains a b quark or a $\bar{b}$ antiquark. We cannot use the measured efficiency and dilution for same side tagging for $\bar{B}^0$ and B^- decays and apply these to B_s^0 mixing. The situation is further complicated by the existence of P wave B mesons, the so-called B^{**}. These P wave mesons can decay down to the light pseudoscalar mesons, and these decays produce the same charge b-flavor correlation illustrated in Fig. 24[5]. The B_s^{**} mesons, however, decay to B^- and $\bar{B}^0$ (and a kaon) and not to B_s^0, so these P wave mesons do not contribute to same side B_s^0 flavor tagging.

In Run I same side tagging was studied using partially reconstructed semileptonic $\bar{B}^0$ and B^- decays such as $\bar{B}^0 \to D^{*+}\ell^-\bar{\nu}_\ell X$. A time-dependent measurement [73, 75, 76] of Δm_d was made using the SST, and the SST was used in the first measurement [74, 77] of $\sin(2\beta)$. The efficiency and dilution of the SST depends on $p_T(B)$, the transverse momentum of the B hadron. There was insufficient statistics to measure the dilution of the SST for $\bar{B}^0$ mesons using the Run I sample of $\bar{B}^0 \to J/\psi \bar{K}^{*0}$, so the dilution had to be extrapolated from the dilution measured with the partially reconstructed semileptonic samples. This extrapolation required Monte Carlo simulation. An extensive comparison between the data and the Monte Carlo revealed that the `PYTHIA` Monte Carlo, after it had been tuned to reproduce the multiplicity distributions measured in the data, reproduced well the efficiency and dilution of the same-side tag. It was found using the Monte Carlo that production of P wave B mesons caused a substantial increase in the SST efficiency and dilution. When applied to $\bar{B}^0 \to J/\psi K_S^0$

[5]For example, $\bar{B}^{0**} \to B^-\pi^+$ producing a π^+ with a B^-. The fragmentation process that produced the $\bar{B}^{0**}$, however, has a π^- produced in association in the most naive picture, which complicates the same side tag.

decays, the SST was found[6] to have

$$\epsilon D^2 = (2.1 \pm 0.6)\,\%.$$

The dilution of the SST for B^- was found to be larger than the dilution of the SST for $\bar{B}^0$. This difference was anticipated from the simple fragmentation picture in Fig. 24: in the case of the charged B meson, a charged pion or a charged kaon may be produced; in the fragmentation of a b quark to a $\bar{B}^0$ only a charged pion is produced. Furthermore, if a neutral K^{*0} meson is produced with the $\bar{B}^0$, then the charged kaon from the decay of this particle has the "wrong" charge, *i.e.,* a positive charge. The difference in the SST dilution between charged and neutral B mesons was confirmed in the Monte Carlo; only allowing charged pions to be used as tagging tracks in the Monte Carlo brought the dilution of the B^- and $\bar{B}^0$ into agreement. This is illustrated in Fig. 25.

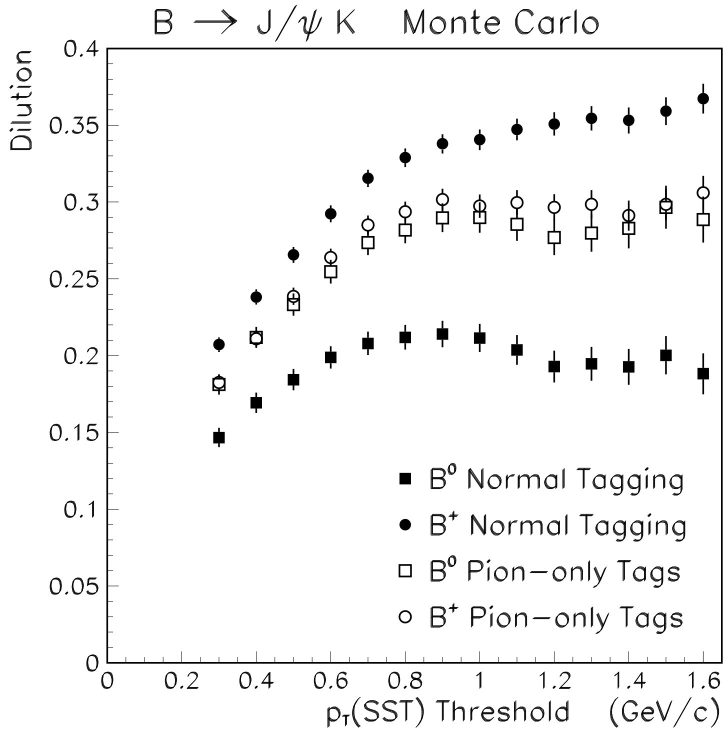

Figure 25: The dilution of the same side flavor tag for charged and neutral B mesons as a function of the minimum required transverse momentum for the tag track from the `PYTHIA` Monte Carlo. The open points show the dilution when only pions are allowed to be used as tag tracks. The dilution for charged B mesons and neutral B mesons agree under this restriction, as expected. This figure is taken from reference [73].

[6]This total tagging efficiency is the sum of the ϵD^2 determined from the first two rows in Table I of [71].

For the final Run I measurement of $\sin(2\beta)$ [71], the same side tag, the soft lepton tag, and the jet charge tag were combined, and the overall efficiency of these three tags combined was

$$\epsilon D^2 = (6.3 \pm 1.7)\,\%.$$

6.3 Expected Improvements in Flavor Tagging in Run II

With the upgrade of the CDF detector to CDF II, many improvements were expected in the capability for B hadron flavor identification. The detector upgrades that had particularly important impact on flavor tagging included the following:

1. The installation of a Time of Flight detector (TOF) outside the main drift chamber, the Central Outer Tracker (COT). The main purpose of the TOF was to add particle identification capabilities to CDF, in particular, to make it possible to identify low momentum kaons. As described below, these kaons would be used for an opposite side kaon tag, which could be applied to both $\bar{B}^0$ and $\bar{B}_s^0$ flavor oscillation measurements and CP asymmetry measurements, and a same side kaon tag, which could be used for $\bar{B}_s^0$ mixing and CP measurements. The same side kaon tag was predicted to be the most powerful of any of the CDF II b flavor tags.

2. The extension of the coverage of the silicon tracking system from $|\eta| < 1$ (Run I) to $|\eta| < 2$ (Run II). This potentially extends the angular coverage for opposite side flavor tags, increasing the efficiency for soft lepton tagging and jet charge tagging.

3. The improved plug calorimeter and the intermediate muon system (IMU) were expected to increase the coverage for soft lepton tags from $|\eta| < 1$ (Run I) to $|\eta| < 2$ (Run II).

4. The introduction of an additional layer of silicon (Layer 00) located at a radius of 1.4 cm to 1.6 cm from the beam was expected to greatly improve our impact parameter resolution in the central region of pseudorapidity ($|\eta| < 1$). The main impact of Layer 00 is on lifetime resolution, however, it is important for tagging as well. It can allow further improvement of vertexing for jet charge tagging methods, and for the opposite side kaon tag, it can be used to help discriminate between prompt kaons from the $p\bar{p}$ underlying event and kaons originating from B hadron decay, which are the basis of opposite side kaon tagging.

5. The improved trigger capabilities of CDF II. Again the main improvement in the trigger is important for obtaining a fully reconstructed $\bar{B}_s^0$ signal, which is crucial for obtaining sufficient lifetime resolution for $\bar{B}_s^0$ mixing measurements. The increased capabilities of the trigger system, however, can allow potential improvements in mixing sensitivity by making it possible to trigger on opposite

side leptons for flavor tagging. The flavor tags discussed so far in this Section are for flavor tags that are unbiased by the trigger.

The most significant expected improvement from CDF to CDF II for $\bar{B}_s^0$ flavor identification is based on the kaon identification provided by the TOF. With 100 ps resolution, the TOF achieves 2σ separation between kaons and pions up to a *momentum* of $p = 1.6\,\mathrm{GeV}/c$. This momentum range covers a significant fraction of opposite side kaons and same side kaons associated with $\bar{B}_s^0$ production. The opposite side kaon tag (OSKT) is based on kaons produced in B hadron decay. The decay chain $b \to c \to s$ produces a K^- (and $\bar{b} \to K^+$). This method of flavor tagging is the most effective method at the B factory experiments[7]. Excellent impact parameter resolution (provided by the CDF II silicon system, in particular, Layer 00) is crucial to separate kaons from B hadron decay from kaons coming from the primary $p\bar{p}$ interaction. The same side kaon tag (SSKT) relies on identifying the kaon produced in the fragmentation of a b quark to a $\bar{B}_s^0$ meson. This process is illustrated in Fig. 24; a K^- (K^+) is produced in association with a $\bar{B}_s^0$ (B_s^0). If a heavier strange meson, such as a vector meson, is produced in association with the fragmentation of a b quark to a $\bar{B}_s^0$, then this may decay down to a K^- too. An example of this process is illustrated in Fig. 24 as well: a $\bar{K}^{*0}$ decays to $K^-\pi^+$. Kaon identification has been found to be very important for separating these fragmentation kaons from pions from the underlying event and fragmentation.

The expected performance for the OSKT and the SSKT were reported in the proposal [78] for the beyond the baseline upgrades to CDF II. The overall tagging efficiency of the OSKT using TOF for kaon identification was $\epsilon D^2 = 2.4\%$ (1.9%) with (without) the the improved impact parameter resolution provided by Layer 00. The tagging efficiency of the SSKT was found to depend on the average $p_T(B)$ of the sample of $\bar{B}_s^0$ used for tagging. The results were $\epsilon D^2 = 2.6\%$ for a sample of $\bar{B}_s^0 \to J/\psi\,\phi$ decays collected with the dimuon trigger; and $\epsilon D^2 = 4.2\%$ for a sample of $\bar{B}_s^0 \to D_s^+\pi^-$ decays collected with the hadronic trigger (Scenario A). The prediction for the overall flavor tagging efficiency after all opposite side tags and same side tags were included for this golden mode for $\bar{B}_s^0$ flavor oscillations was 11.3%, which was the sum[8] of 2.4% for the OSKT, 4.2% for the SSKT, 1.0% for the SLT(μ), 0.7% for the SLT(e), and 3.0% for the JQT. The increase of the SLT ϵD^2 over that observed in the inclusive lepton analysis was based on the expectation that we would extend our low p_T lepton detection out to larger values in pseudorapidity. The increase of the JQT ϵD^2 over that observed in the inclusive lepton analysis was based on a study reported in [79], a generator level study of jet charge tagging. The predicted increase is based on using lifetime information with the extended coverage provided by the CDF II silicon system.

---

[7] The lepton tag is cleaner (higher dilution), but less efficient, leading to lower overall ϵD^2.

[8] This sum does not include any correlations between the various flavor tags.

6.4 Current Performance of Flavor Tags in Run II Data

The performance of several of the Run I flavor tags has been evaluated using Run II data. For the opposite side tags, both the SLT and the JQT have been investigated. Preliminary values of ϵD^2 for the SLT(μ) and the JQT have been approved (blessed) for public presentation. A preliminary value of the performance of the SLT(e) is nearly ready for blessing as well. The opposite side kaon tag has been evaluated too, however, the performance of the TOF must be improved substantially (efficiency and tails in the time of flight resolution) before this tag will have useful ϵD^2. Several groups (Cantabria, MIT, Penn, Purdue) have observed statistically significant dilutions based on opposite side kaon tagging.

The same side tag performance (not using kaon identification) has been studied using partially reconstructed semileptonic B meson decays and fully reconstructed B meson decays collected with the dimuon trigger and the hadronic trigger. Preliminary values of ϵD^2 for both the B^- and the $\bar{B}^0$ mesons have been blessed, and the dilution for the SST for the neutral B meson was determined in a simultaneous extraction of the $\bar{B}^0$ mixing parameter Δm_d. No performance evaluations of the same side kaon tag for $\bar{B}^0_s$ based on the data have been presented so far.

6.4.1 Performance of Opposite-side Tags in Run II

The performance of the opposite side tags has been studied in a very high statistics sample of semileptonic B hadron decays collected with the lepton plus displaced (SVT) track trigger. This trigger requires a CEM electron (CMUP muon) with $E_T > 4\,\mathrm{GeV}$ ($p_T > 4\,\mathrm{GeV}/c$) and a SVT track with $p_T > 2\,\mathrm{GeV}/c$ and $120\,\mu\mathrm{m} < d_0 < 1\,\mathrm{mm}$. Furthermore, the lepton and the SVT track must satisfy $2° < \Delta\phi(\ell, \mathrm{SVT}) < 90°$ and $m(\ell + \mathrm{SVT}) < 5\,\mathrm{GeV}/c^2$. Studies of the composition of this data sample led to a technique [80] based on the sign of the impact parameter of the SVT trigger track that yields background subtracted measurements of the efficiencies and dilutions of the opposite side flavor tags.

Whether or not the opposite side flavor tag made the correct inference about the trigger side production b flavor is determined on an event by event basis. The sign of the charge of the trigger side lepton determines whether the trigger side B hadron contained a b quark or $\bar{b}$ antiquark when it decayed. An average correction must be made to the measured *flavor tag dilution* D_{raw} to account for the the fraction of events in which the flavor assignment of the trigger side B hadron is incorrect because the trigger side lepton originates from a sequential leptonic decay ($b \to c \to \ell$) or the trigger side B hadron has mixed to the flavor opposite its flavor at production. This average correction is called the "trigger side" dilution D_{trig}, and the value is found to be $D_{\mathrm{trig}} = 0.6412 \pm 0.0012(\mathrm{stat.})$ [80]. The true dilution of the opposite side flavor tag is given by $D = D_{\mathrm{raw}}/D_{\mathrm{trig}}$. This value of D_{trig} applies only to flavor tag dilutions determined following the prescription defined in reference [80].

The first b flavor tag based on Run II data that was approved for public presentation was the soft muon tag or SLT(μ). Using approximately 100,000 electron+SVT events and 100,000 muon+SVT events (corresponding to $64 \pm 4\,\mathrm{pb}^{-1}$), the total tagging efficiency for the SLT(μ) was measured [81] to be

$$\epsilon D^2 = (0.66 \pm 0.19)\,\%.$$

This value of ϵD^2 has been corrected for the trigger side dilution and is the sum of the values of ϵD^2 obtained in six independent bins of p_T^{rel}. The variation of the *raw* dilution with p_T^{rel} is shown in Fig. 26. This value of ϵD^2 includes all the central muon

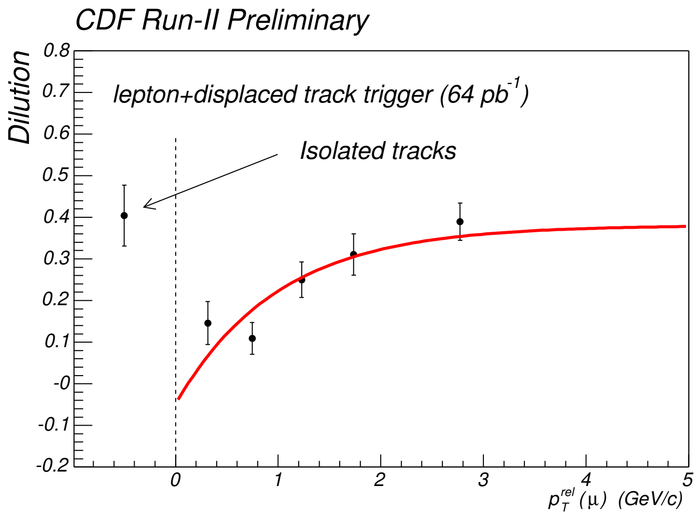

Figure 26: Variation of the raw dilution of soft muon tags with p_T^{rel}. These values are the average of the dilutions measured in the electron+SVT sample and the muon+SVT sample. The isolated muons correspond to opposite side muons that are not associated with a tracking jet. The curve is a parameterization [81] that is fit to the data to provide D_{raw} as a function of p_T^{rel} for the likelihood function employed in a mixing measurement.

subsystems (CMU, CMUP, CMP, CMX) and the BMU. Including the BMU increases ϵD^2 by 15%. Although further improvements in reconstruction (in particular track reconstruction at larger values of η) should enhance the contribution of the BMU, this increase is less than the size anticipated in our projections[9].

The performance of the jet charge flavor tag in Run II has been evaluated [82] using $190\,\mathrm{pb}^{-1}$ of lepton+SVT data, corresponding to about 321,000 electron+SVT events and 402,000 muon+SVT events. The jet charge algorithm follows the algorithm described in reference [63], that is, it uses track-based jets found using an $\eta - \phi$ algorithm, and it identifies opposite side b jets using a vertexing algorithm, if possible. Since a Run II equivalent of `BVTX` does not exists, the high-p_T b tag algorithm, `SECVTX`, was used

[9]The raw dilution for the muons in the BMU is $(31.4 \pm 5.8)\%$ [81], which is good. The efficiency is $(0.30 \pm 0.02)\%$, which is low.

as a substitute[10]. The jet charge $Q_{\rm jet}$ is calculated using a p_L weighting scheme, and the overall ϵD^2 is the sum of the ϵD^2 measured in eleven bins of $Q_{\rm jet}$. The measured values are $\epsilon D^2 = (0.318 \pm 0.020)\%$ in the muon+SVT sample and $\epsilon D^2 = (0.306 \pm 0.023)\%$ in the electron+SVT sample. These values of ϵD^2 have been corrected by $1/D^2_{\rm trig}$. The performance of the jet charge tag in Run II was compared with the performance found in Run I and found to be comparable, if an additional requirement of $p_T(\ell) > 8\,{\rm GeV}/c$ was applied to the Run II data to match more closely the Run I trigger requirements.

A CPU-intensive optimization process was carried out to improve the performance of the jet charge tag. The optimization procedure involved the evaluation of the clustering parameters in the jet algorithm, the jet selection parameters, and the parameters used in `SECVTX`. Notable modifications to the parameters included increasing the ΔR of the cone associated with finding track-based jets (originally set at $\Delta R = 0.7$) or including tracks in the jet charge calculation (originally set at $\Delta R = 0.8$) to $\Delta R = 1.5$. A similar increase was found to be effective when the jet charge tag was applied to the dimuon sample in the Run I measurement of $\sin(2\beta)$. The optimization of the parameters was carried out using the muon+SVT sample and then applied to the electron+SVT sample to evaluate the improvement. In the biased muon+SVT sample, ϵD^2 increased from 0.32% to 0.53%; in the unbiased electron+SVT sample, ϵD^2 increased from 0.31% to 0.42%. The jet-charge distribution measured in the electron+SVT sample is shown in Fig. 27; $Q_{\rm jet}$ is expected to be opposite the sign of the trigger lepton charge on average. Figure 28 shows the measured variation of dilution with $Q_{\rm jet}$. The variation of D with $Q_{\rm jet}$ is fit to a line, and the slope of this line may be used in the likelihood function used in a measurement of B mixing. Just as in the Run I mixing analysis, the dilution corresponding to $Q_{\rm jet} = 1$ is not included in the linear model of $D_{\rm JQT}(Q_{\rm jet})$.

The present performance of the jet charge tag is

$$\epsilon D^2 = (0.419 \pm 0.024\,({\rm stat.})^{+0.020}_{-0.034}\,({\rm syst.}))\,\%.$$

This is significantly below the value measured in Run I, however, this corresponds to a data sample with lower $p_T(B)$ on average. Since the tree level production of $b\bar{b}$ pairs produces back-to-back topologies balanced in transverse momentum, there should be a correlation between the trigger-lepton transverse momentum and the transverse momentum of the tag jet. The dependence of the jet charge dilution with the transverse momentum of the tag jet is shown in Fig. 29 for jets with and without a `SECVTX` vertex. The data suggest a possible increase of dilution with increasing transverse momentum of the tag jet, however, the data are also compatible with no significant dependence. There is a clear difference between the dilution of jets that have a `SECVTX` vertex compared to jets with no `SECVTX` vertex.

Evaluations of a soft electron tag and several variations of an opposite side kaon tag have been shown in various meetings. A recent presentation of the soft electron

[10] Track-based jets, instead of calorimeter jets, were used to drive the `SECVTX` secondary vertex finding algorithm.

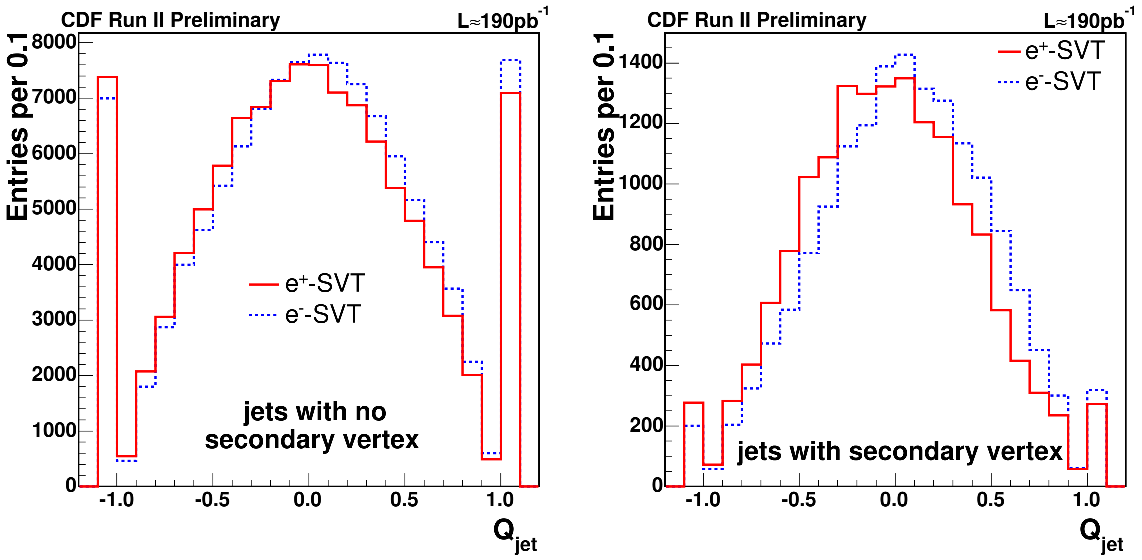

Figure 27: The jet charge distribution measured in the electron+SVT data sample for positrons (red) and electrons (blue). The distribution on the left (right) is for jets without (with) a secondary vertex found by SECVTX.

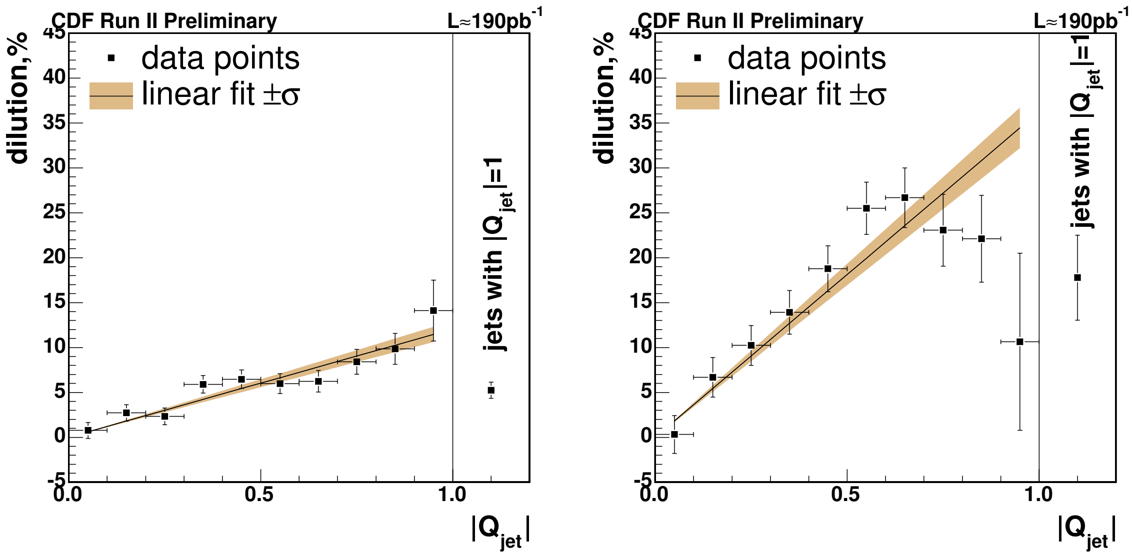

Figure 28: The variation of the dilution of the jet charge tag with $Q_{\rm jet}$ measured in the electron+SVT data sample. The distribution on the left (right) is for jets without (with) a secondary vertex found by SECVTX.

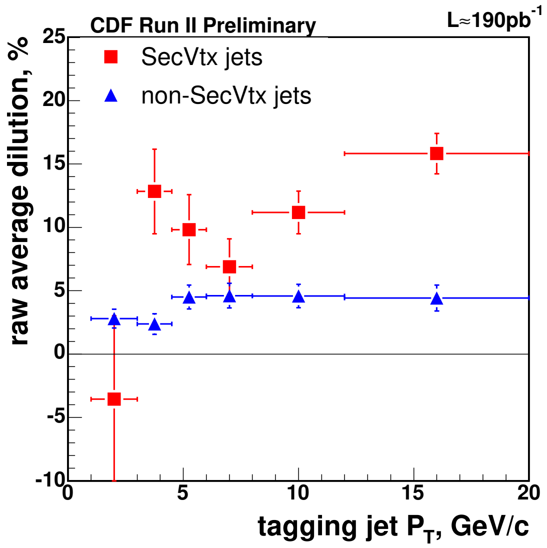

Figure 29: The variation of the dilution of the jet charge tag with the p_T of the tag jet for both jets with a `SECVTX` vertex (red points) and jets with no identified secondary vertex (blue points).

tag at the Semileptonic B subgroup meeting (18 May 2004) reported an overall tagging efficiency that is about 50% of the ϵD^2 of the soft muon tag (*i.e.*, half as effective). A statistically significant dilution using opposite side kaons has been demonstrated in the lepton+SVT data, however, the overall ϵD^2 using these kaons is too small (less than 0.1%) to be of practical use.

The present overall performance of the opposite side flavor tags is difficult to estimate since such an estimate requires including the correlations between the various opposite side tagging algorithms. A best-case estimate comes from assuming there are no correlations, and this estimate yields approximately 2%, well below the 7.1% reported in [78].

6.4.2 Performance of Same-side Tags in Run II

The performance of the same side tags has been evaluated [83] with fully reconstructed samples of $B^- \to J/\psi K^-$ (signal 3011 ± 79), $B^- \to D^0\pi^-$ (signal 4809 ± 88), $\bar{B}^0 \to J/\psi \bar{K}^{*0}$ (signal 1040 ± 40), and $\bar{B}^0 \to D^+\pi^-$ (signal 4865 ± 111) in 245 pb^{-1} of Run II data. The same side tagging algorithm is the same algorithm [73, 76] used in Run I: tracks that are tag candidates are required to fall in a cone of $R = 0.7$ around the B candidate momentum vector; the candidate tracks must have $p_T > 0.35\,\mathrm{GeV}/c$ and must be consistent with originating from the primary vertex. If there is more than one track satisfying these criteria, then the track with the smallest magnitude of the transverse momentum relative to the vector sum of the momentum of the B candidate and the candidate track is selected (the "minimum p_T^{rel} algorithm").

The variation of the asymmetry between "right-sign" and "wrong-sign" correlations is measured as a function of proper decay time. This variation is expected to be flat for charged B mesons, and the asymmetry is fit to a line; the height of this line gives the dilution D_+ for same side tagging applied to $B^\pm$. For neutral B^0 mesons, this asymmetry is fit to $D_0 \cos(\Delta m_d t)$, where t is the proper decay time of the neutral B meson. The resolution on this proper time is taken into account in this fit, although it is unimportant, since the oscillation frequency times proper lifetime is small[11] for B^0 mesons: $x_d = \Delta m_d \tau = 0.76$. The resulting flavor tagging performances are

$$\epsilon D^2 = (2.8 \pm 0.5)\,\%$$

for the B^- and

$$\epsilon D^2 = (1.0 \pm 0.5 \pm 0.1)\,\%$$

for the $\bar{B}^0$. The difference between charged and neutral B mesons comes from the dilution, not the efficiency: $D_+ = 0.209 \pm 0.020$ and $D_0 = 0.124 \pm 0.033$. A similar difference was observed in Run I, and is expected, as discussed previously, due to the presence of charged kaons produced in association with B^-. This measurement of the same side tag performance is being repeated [84] using partially reconstructed semileptonic decays such as $B^- \to D^0 \ell^- \bar{\nu}_\ell$ and $\bar{B}^0 \to D^{*+} \ell^- \bar{\nu}_\ell$. Although there are much higher statistics in the semileptonic modes, the analysis is more complicated due to cross-talk between the charged and neutral B mesons and the unresolved systematic problem in the semileptonic lifetime measurement. The measured efficiencies and dilutions agree well with the values quoted above.

The performance of the same-side tag for B^- and $\bar{B}^0$ as measured in the data does not tell us the performance for B_s^0. This performance will be measured only when we observe B_s^0 oscillations.

6.5 Recommendations for Further Progress in Flavor Tagging

The performance of the b flavor identification algorithms is significantly below the performance anticipated based on Run I data and studies carried out for the CDF II upgrade. Although there are several areas where we could direct effort on flavor tags, there are two particularly important areas:

1. understanding the performance of the opposite side flavor tags with Monte Carlo in order to determine if we can improve these tag algorithms in a significant way, and

2. making a systematic comparison between data and Monte Carlo of the identity and kinematic properties of particles produced in association with B^-, B^0, and B_s^0 mesons.

[11] This resolution is crucial for B_s^0 oscillations, and the method of extraction of Δm_d used here is not applicable to measuring, or placing a limit on, Δm_s.

This second area is necessary if we want a defensible prediction of the performance of the same side kaon tag. If our first result on B_s^0 mixing is a limit on Δm_s, then we must predict the dilution of the SSKT from Monte Carlo.

The performance of the opposite side flavor tags has been evaluated with the letpon+SVT data sample. At present, the lifetimes of the B^0 and B^- mesons measured with this data sample are about 10% low. The cause of this discrepancy is not firmly established; the current favored hypothesis is that there is a background in this sample, *e.g.*, from $c\bar{c}$ production or from fake leptons, that has not been taken into account with the lifetime fit. If this background is the cause of this discrepancy, then the affect on the evaluation of the flavor tags is probably small. We must make a quantitative estimate of this affect, however.

We should also invest effort on improving the current flavor tags and implementing new flavor tags. For example, inclusive $D^{*\pm}$ reconstruction from the decay of the opposite side B hadron could be used as a flavor tag: the charge of the bachelor pion in the D^* decay indicates whether the D meson originated from a hadron with a b quark (π^+) or a $\bar{b}$ antiquark (π^-). Further discussion of how to proceed with the flavor tags may be found in [85].

Finally, a relatively new trigger strategy has been implemented that incorporates the flavor tag as part of the trigger. In this trigger, a 1.5 GeV/c CMU muon or a 2.0 GeV/c CMX muon is combined with an opposite side SVT pair. All of the tags discussed in this section are not biased by the trigger. This new trigger potentially provides a tagged sample of hadronic B hadron decays and may ultimately be the most viable path to measuring B_s^0 mixing. Effort to study the performance of this flavor tag should be carried out as soon as possible to make sure we are using the most optimal trigger strategy for getting to B_s^0 mixing.

7 Conclusions and Recommendations

We discuss the current B_s^0 mixing reach and summarize the basic components that contribute to the present situation. We shall see that this is an extremely complex measurement which requires a good understanding of many CDF subsystems and their reconstruction as well as the development of several dedicated analysis tools. No single outstanding problem emerges that, if solved, could improve drastically the current situation. We rather observe that progress can be made only by improving every single element of this complex analysis.

In the last subsection we summarize our recommendations.

7.1 Current Reach in B_s^0 mixing

Our reach in the measurement of the mixing parameter Δm_s is described analytically by Equation 4 presented in the introduction. The key ingredients are the signal yield and the associated background level, the proper lifetime resolution and the flavor tagging power of our flavor tag combination. Summarizing the results from the previous sections we can optimistically expect a signal yield of 3000 events/fb^{-1} with an average S/N ~ 1 after including the contributions from the $D_s \to K^*K$, which roughly doubles the sample; the average proper lifetime resolution is not well understood, but lower bounds are estimated to be about 50 fs with L00 and 67 fs without L00 information; finally the currently achieved flavor tagging power, ϵD^2, is 1.7% from opposite side tagging only if we assume that the electron tag (still not clearly characterized) eventually gives a performance similar to the muon tag. Unfortunately while we have demonstrated a same side tagging performance of about 1% for B^0 we have yet no demonstration of a non-zero tagging power using a kaon signature either in the opposite side or the same side. If we assume that we need a measurement at the 5σ level to discover B_s^0 mixing we can see by application of Equation 4 that we have no sensitivity to any Δm_s, no matter how small, until we integrate almost 2 fb^{-1}. If we rescale our goals to just improve the current world limit ($\Delta m_s = 14.9$) at the 95% confidence level we can see that the requirement is between 350 and 500 pb^{-1} depending on the final proper time resolution. However this ignores the additional systematics arising from the use of the $D_s \to K^*K$ channel, which in the end will make us use more statistics to achieve the same result. As an example if only the $\phi\pi$ and 3π channels are used we may need as much as 1fb^{-1} to set such a limit.

This situation is clearly unsatisfactory as well as embarrassing for the CDF collaboration. We shall now examine each ingredient of our B_s^0 mixing reach and summarize the possibilities for improvement relative to the current situation.

7.2 Signal Yield

Current yields are based on $B_s \to D_s\pi$ with the D_s decaying to either $\phi\pi$, 3π or K^*K. Adding other decay channels, such as $B_s \to D_s 3\pi$, is clearly to be pursued further even if the currently obtained yields are much lower than originally expected and the expectation is to have more background in these higher multiplicity modes. It should be noted that improvements in the tracking efficiency, in particular at low p_t, may significantly affect our efficiency for these higher multiplicity final states.

Similar modes where the D_s is replaced by a D_s^* have been observed at a rate similar to the D_s modes although with a worse S/N. While the resolution is not as good due to the loss of the soft photon, we should explore in detail how much can be gained by the inclusion of these final state and, possibly, attempt to reconstruct and associate the missing γ.

Finally extending the signal search to include regions not covered by the COT could also give significant yield gains. This would require however significant improvements of the quality of tracking in the forward regions.

It appears from section 2 that the current trigger is already highly optimized to give the maximum possible yield. On the other hand such trigger yields may soon be degraded if we do not aggressively continue to improve and update our trigger system, in particular if the Tevatron luminosity continues the positive improvement trend that we have observed over the past months. We are very concerned that we will not be able to maintain the current B_s^0 yields and recommend that every effort be made to complete and install all trigger and DAQ upgrades currently undertaken or planned by the CDF collaboration as soon as possible.

7.3 Lifetime Resolution

The lifetime resolution depends critically on the quality of the tracking code. Lifetime resolution can be degraded by many effects: inefficiencies in the association of hits to tracks, association of wrong hits, incomplete understanding of tracking resolution and the associated material effects, and last, but certainly not least, the quality of the alignment.

While significant improvements have been made in tracking from the 4 to 5 offline release series, we think it is very important to continue studying and optimizing the tracking performance in all detector regions.

In particular the use of L00 has demonstrated a significant improvement in the track parameter resolution, but work still needs to be done to improve its efficiency and alignment. Extending our reach to large Δm_s may depend critically on the final lifetime resolution achieved. A proper lifetime resolution difference at the 20 - 30% level may look small, but it can be equivalent in certain cases to doubling the required integrated luminosity. We strongly encourage more work on the optimization of the use of L00 in our standard tracking algorithms.

More work is needed to improve the determination of the primary vertex on an event by event basis. While this approach has been highly succesful for the analysis of top data no substantial improvements on the transverse beam position determination have been obtained so far. Further optimization of the primary vertex finding for low p_t B samples could lead significant improvements.

In addition we need to have all systematic effects under control. For instance the lifetime measurements in our semileptonic samples need to be fully understood. These samples are of great importance for the calibration of our mixing dilutions and it is hard to feel comfortable about them until the lifetime issue is fully understood.

7.4 Flavor Tagging

The flavor tagging power reached so far with Run II data is clearly one of the greatest concerns with regards to B_s^0 mixing analyses. Indeed our results so far are significantly worse than both our predictions or extrapolations and the previous Run I performance. We believe this discrepancy is not yet fully understood and should be studied and modeled in detail in the hope of improving the performance or at least understanding the source of the problem.

In addition we notice that no flavor tagging power has been demonstrated yet for the kaon tagging both on the opposite side and on the same side. This is where we would expect most of our tagging power to come from and therefore much effort should be placed in this area. It should be remarked here that the improvement of the dE/dx in the COT and of the time of flight reconstruction are very critical. Only accurate work on the calibration and reconstruction of these systems will allow their full potential to be reflected in the tagging power. We note that this tagging technique is completely new at CDF so a major effort should be invested in its detailed understanding and optimization.

7.5 Reconstruction Infrastructure

In the previous sub-sections we have discussed the basic elements of the mixing analyses, but there are additional analysis tools that also need to be developed. In the following we mention just a few:

- the fitting technique and the methods to handle multiple flavor tags
- the use of Fourier transform space for the determination of the mixing frequency
- the handling of backgrounds in the fits
- the study of analysis-specific systematics
- possibly the combination of several flavor taggers to obtain a more powerful one

It is very important that all these tools be well understood before application to the most difficult and challenging problem, B_s^0 mixing. Natural areas to gain full confidence in all of this are the measurements of B lifetimes using the hadronic final states and those of B^0 mixing, which use many of the same techniques. While CDF is not competitive on the measurement of B^0 mixing it is still very important that we demonstrate our mixing analysis capabilities there and use it to tune all the basic algorithms and develop a common infrastructure for this type of measurement across the B group.

7.6 Organizational Aspects

B_s^0 mixing is a complex analysis that requires optimization of every one of its many components. Even if individually each is pursued aggressively, there is no guarantee that the entire enterprise will progress as rapidly as it could. To accomplish that, the overall organization must be strengthened with special attention paid to communications among the subgroups. With that in mind, we have the following organizational recommendations.

- Improve coordination and commonality of the mixing analyses among the various B sub-groups.
- Strengthen the relevant detector reconstruction and operations groups, in particular the tracking groups, including L00, dE/dx and TOF.

- Improve coordination of mixing analysis groups with relevant detector reconstruction groups.

- Recruit more people to carry out this extremely important analysis.

- Assign two senior experienced CDF people the task of coordinating the mixing analyses.

7.7 Final Words

It is clear that the B_s^0 mixing measurement is very challenging, and it may seem that a result is far in the future. Certainly it is true that if Δm_s is at the upper end of the standard model expected range, it will take a long time before we have the sensitivity to reach it. But we do not know the value of $\Delta\ m_s$, and it may, in fact, lie near the current limit of 14.9 ps^{-1}. In this case a vigorous analysis effort could have the needed sensitivity in the not so distant future. Sooner still should be the first CDF result on B_s^0 mixing in the form of a limit, which becomes interesting when we exceed the best single-experiment limit of 10.9 ps^{-1}. None of these results will come easily. Rather than being discouraged by this, we hope that the collaboration will be excited by the challenge and contribute to the effort. There is no shortage of things on which to work.

References

[1] Review of Particle Properties, Physical Review **D66** 670 (2002).

[2] ALEPH Collaboration, Eur. Phys. J. **C29**, 143 (2003).

[3] B Physics at the Tevatron Run II and Beyond ; Fermilab-Pub-01/197

[4] A. Belloni, I.K. Furic & C. Paus, Update of B_s^0 event yields using `GenTrig`, CDF note 6816.

[5] J. Lewis and I. Yu, CDF Note 4095.

[6] S. Donati and G. Punzi, CDF Note 3780.

[7] J. Mueller and P. Wilson, CDF Note 2665.

[8] M. Rescigno, Presentation at the CSL Review, June 23, 2003.

[9] D. Saltzberg, CDF Note 6434.

[10] K. Pitts, Presentation to the B Physics Trigger Workshop, October 9, 2003.

[11] R. Oldeman, CDF Note 6926.

[12] The CDFIIB Collaboration, CDF Note 6261.

[13] N. Lockyer *et al.* , CDF Note 6792.

[14] R. Oldeman, CDF Note 6028.

[15] F. Ukegawa, Presentation at B Reconstruction Subgroup, June 13, 2003.

[16] M. Tanaka, Presentation at TDWG, Oct. 31, 2003.

[17] Kevin Lannon, Presentation to the Run 2b XFT Review Committe, May 25, 2004.

[18] A. Mukherjee, CTC and VTX Tracking, CDF note 5490, note The Run 2 SL algorithm is an adaptation of the Run 1 CTC algorythm.
C. Hays, P. Tamburello, A. Kotwal, P. Wittich, R. Snider, The COT Pattern Recognition Algorithm and Offline Code, CDF note 6992

[19] P. Azzi, G. Busetto, P. Gatti and A. Ribon, Histogram Tracking in the COT, CDF note 5562

[20] W. Yao and K. Bloom "Outside-In" Silicon Tracking at CDF, CDF Note 5991

[21] S. Menzemer, TrackingKal, CDF Note 5968

[22] S. Menzemer, TrackingKal, CDF Note 5968

[23] Y. Huang, C. Hays, and A. Kotwal, Inside-Out Tracking, CDF note 6707

[24] S. Giago, 5.1.0 Validation with D*, Presentation at the September 17, 2003 Tracking Meeting.

[25] M. Herndon, 5.1.0 Performance and Validation, Presentation at the September 17, 2003 Tracking Meeting

[26] M. Herndon, 5.1.0 Validation,Presentation at the August 3, 2003 Tracking Meeting.

[27] M. Herndon, Tracking: Winter Conference Development and Validation, Presentation at the September 3, 2003 Tracking Meeting.

[28] M. Herndon, Tracking Status and Stereo Studies,Presentation at the October 15, 2003 Tracking Meeting.

[29] T. Scheidle, Studies of Silicon Forward Tracking Efficiency, Presentation at the December 10, 2003 Tracking Meeting.

[30] Y. Huang, Tracking Validation in 5.3.1pre2, Presentation at the February 11, 2004 Tracking Meeting.

[31] M. Herndon, 5.3.1 Tracking Validation, Presentation at the February 11, 2004 Tracking Meeting.

[32] C. Hays, A. Kotwal, and R. Culbertson, Alignment, Presentation at the July 16, 2003 Tracking Meeting.

[33] D. Sherman, J. Guimaraes and K. Burket, Update on Prime Vtx Status, Presentation at the December 3, 2003 Tracking Meeting.

[34] K. Anikeev, New Measurements of Lifetimes in B+ to JPsi K+. B0 to JPsi K*0 and Bs to JPsi Phi, Presentation at the April 28, 2004 B meeting.

[35] M. Shapiro, A. Ceri, S. Da Ronco, D. Lucchesi, S. Dececco and G. Salamanna, Primany Vertex Finding an Update, Presentation at the April 9, 2004 BPAK meeting.

[36] B. Brau, D. Stuart and M. Herndon, L00 Studies, Presentation at the September 24, 2003 Tracking Meeting, and A Study of L00 Performance in Release 5.1.0, CDF Note 6720

[37] M. Herndon, L00 Studies, Presentation at the December 2, 2003 J/ψ meeting.

[38] R. Mumford, Understanding L00 Efficiency, Presentation at the April 9, 2004 BPAK meeting.

[39] J. Miles and C. Paus, L00 Effects in B+ Lifetime Fitting, Presentation at the March 19, 2004 BPAK meeting.

[40] I. Furic & C. Paus, Event Yield Estimates for Hadronic B_s^0 Decays in Run II, CDF note 5349.

[41] D. Lucchesi, et al., An Update to the Study of B_s^0 Yields in the Hadronic Trigger, CDF note 6692.

[42] Donatella Lucchesi, Private communication.

[43] A. Belloni, I. Furic & C. Paus, Study of B_s^0 , B^+ and B^0 Yields in the Hadronic B samples, CDF note 6614.

[44] A. Belloni, et al., A Combined Measurement of Br($B_s^0 \to D_s^- \pi^+$)/Br($B^0 \to D^- \pi^+$), CDF note 6708.

[45] D. Lucchesi, Presentation at the February 13, 2004 B_s^0 Expediters Meeting.

[46] B. Iyutin, Presentation at the October 13, 2003 B Hadronic Meeting.

[47] S. Pagan Griso, Presentation at the September 15, 2003 B Hadronic Meeting.

[48] P. Squillacioti, Presentation at the April 5, 2004 B Hadr onic Meeting.

[49] A. Belloni, Presentation at the November 3, 2003 B Hadronic Meeting.

[50] R. Carosi, Presentation at the October 31, 2003 Joint Physics Meeting.

[51] G. Bauer, et al., Selection of Semileptonic B Decays Optimized for B Mixing Measurements, CDF Note 6780.

[52] N. Leonardo et al., Presentation at the January 13, 2004 Semileptonic B Meeting.

[53] R. Snider, et al., TOF reconstruction group presentation to the Bs Task Force, Jan. 9, 2004.

[54] CDF-II Collaboration, "Proposal for Enhancement of the CDFII Detector: An Inner Silicon Layer and a Time of Flight Detector", CDF-5096, Oct., 1998.

[55] K. Anikeev, *et al.*, "Construction and Installation of the CDF Time-of-Flight Counters for Run II", CDF-5818, Dec., 2001.

[56] The electronics for the TOF detector are documented in CDF-5361, CDF-5362, CDF-5612 and CDF-5632.

[57] I.K. Furic, *et al.*, "Event t_0 Calculation in Time-of-Flight Reconstruction", CDF-5689, Aug., 2001.

[58] W. Blum and L. Rolandi, "*Particle Detection with Drift Chambers*", Springer-Verlag.

[59] S-S. Yu, *et al.*, "COT dE/dx Measurement and Corrections", CDF-6361, 2001.

[60] S. d'Auria, *et al.*, "dE/dx Performance Studies Using $D*^{\pm}$", CDF-6369, March, 2003.

[61] S. d'Auria, *et al.*, "Track-based Calibration of the COT Specific Ionization", CDF-6932, Apr., 2004.

[62] S. Giagu, et al., dE/dx working group presentation to the Bs Task Force, March 2, 2004.

[63] The CDF Collaboration, F. Abe *et al.*, Phys. Rev. **D60**, 072003 (1999).

[64] K. Lannon, CDFNOTE 6254.

[65] Owen Long, Ph. D. dissertation, 1997.

[66] Marc D. Peters, Ph. D. dissertation, 1997.

[67] C. Gay, G. Michail, R. Strohmer, J. de Troconiz, CDFNOTE 3751

[68] O. Long, M. D. Peters, CDFNOTE 4400.

[69] O. Long, J. Kroll, CDFNOTE 3069.

[70] O. Long *et al.*, CDFNOTE 4315.

[71] The CDF Collaboration, T. Affolder *et al.*, Phys. Rev. **D61**, 072005 (2000).

[72] J. Kroll, O. Long, K. Pitts, CDFNOTE 4847.

[73] The CDF Collaboration, F. Abe *et al.*, Phys. Rev. **D58**, 032001 (1999).

[74] The CDF Collaboration, F. Abe *et al.*, Phys. Rev. Lett. **81**, 5513 (1998).

[75] The CDF Collaboration, F. Abe *et al.*, Phys. Rev. Lett. **80**, 2057 (1998).

[76] Petar Maksimović, Ph. D. dissertation, 1998.

[77] Kenneth Kelley, Ph. D. dissertation, 1999.

[78] The CDF II Collaboration, *Proposal for Enhancement of the CDF II Detector an Inner Silicon Layer and Time of Flight Detector* (P-909), this proposal may be obtained from the following web URL: http://www-cdf.fnal.gov/upgrades/upgrades.html; the direct link to the postscript file is http://www-cdf.fnal.gov/upgrades/btb_proposal.ps

[79] F. Dejongh, CDFNOTE 2629.

[80] M. Jones, D. Usynin, J. Kroll, B. Wicklund, CDFNOTE 6480.

[81] M. Jones, D. Usynin, J. Kroll, B. Wicklund, CDFNOTE 6483.

[82] G. Bauer *et al.*, CDFNOTE 6951.

[83] G. Bauer *et al.*, CDFNOTE 6964.

[84] G. Bauer *et al.*, CDFNOTE 7011.

[85] see J. Kroll and M. Shapiro, http://surf.hep.upenn.edu/ kroll/tagging.txt.