

Development of a Monte Carlo Re-Weighting Method
for Data Fitting and Application to Measurement of
Neutral B Meson Oscillations

by

David J. Asgeirsson

B.Sc., Simon Fraser University, 1998

M.Sc., The University of British Columbia, 2005

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

August, 2011

© David J. Asgeirsson, 2011

Abstract

In experimental particle physics, researchers must often construct a mathematical model of the experiment that can be used in fits to extract parameter values. With very large data sets, the statistical precision of measurements improves, and the required level of detail of the model increases. It can be extremely difficult or impossible to write a sufficiently precise analytical model for modern particle physics experiments.

To avoid this problem, we have developed a new method for estimating parameter values from experimental data, using a Maximum Likelihood fit which compares the data distribution with a “Monte Carlo Template”, rather than an analytical model. In this technique, we keep a large number of simulated events in computer memory, and for each iteration of the fit, we use the stored true event and the current guess at the parameters to re-weight the event based on the probability functions of the underlying physical models. The re-weighted Monte-Carlo (MC) events are then used to recalculate the template histogram, and the process is repeated until convergence is achieved.

We use simple probability functions for the underlying physical processes, and the complicated experimental resolution is modelled by a highly detailed MC simulation, instead of trying to capture all the details in an analytical form.

We derive and explain in detail the “Monte-Carlo Re-Weighting” (MCRW) fit technique, and then apply it to the problem of measuring the neutral B meson mixing frequency. In this thesis, the method is applied to simulated data, to demonstrate the technique, and to indicate the results that could be expected when this analysis is performed on real data in the future.

Preface

The research presented in this thesis was performed by David J. Asgeirsson as a member of the *BABAR* Collaboration. This collaboration consists of 300–500 members at various times, who worked together to design, build, and operate the *BABAR* detector at the SLAC PEP-II accelerator. The experiment ran from 1999–2008. The author has been a student associate of the collaboration since 2002, and a full member since 2005. As only one of the many members of the collaboration, he is responsible for only a small part of the experimental operation of the detector, as described in Appendix D.

The data analysis described in this thesis has been carried out with the aid of many different pieces of software. The event generation software, EvtGen, is a product of members of the *BABAR* collaboration, and was used with an additional code module written by the author to save the specific information required for this analysis. The detailed detector simulation, based on GEANT4, was written by the members of the *BABAR* collaboration. The event reconstruction software was also written by members of the *BABAR* collaboration. A stand-alone program, operating within the *BABAR* software framework, was written by the author to study reconstructed events of interest, and extract and store separately the quantities needed for further analysis.

The Monte-Carlo Re-Weighting fit algorithm was conceived by the author's research supervisor, Dr. Thomas Mattison, and was encoded in two distinct versions by the author and by Dr. Mattison. The author used his own implementation of the MCRW algorithm to carry out the work described in this thesis. Within this software, pre-written pieces of code from Numerical Recipes were used to implement generic linear algebra and polynomial equation solving routines.

None of the work presented in this thesis has been previously published, although plans exist for a detailed publication of the MCRW algorithm, and hopefully a publication of the analysis of neutral *B* meson mixing with dileptons will be completed in the future.

Contents

Abstract	ii
Preface	iii
Contents	iv
List of Tables	vii
List of Figures	ix
Acknowledgements	xiii
1 Introduction	1
1.1 Particle Physics Data and Simulations	1
1.2 Neutral B Meson Mixing	2
1.3 The Monte Carlo Template Method	4
1.4 Barlow-Beeston Method	7
1.5 The Monte-Carlo Re-Weighting Method	7
1.6 Outline of This Thesis	9
2 Monte-Carlo Re-Weighting Fit Method	11
2.1 Likelihood Formulation	11
2.2 Likelihood Maximization	15
2.3 Binning Issues	18
2.4 Parameter Errors	19
2.5 Goodness of Fit	19
2.6 Example Application of MCRW Method	20
3 Dilepton B Mixing Analysis Issues	25
3.1 Boost Approximation	25
3.2 Background Processes	26
3.3 Monte-Carlo Models	30
3.4 Traditional Fitting Approach	31
3.5 Overview of MCRW Approach	36

4	Tracking Resolution	40
4.1	Contributions to Δz Resolution	40
4.2	<i>BABAR</i> Single Track Resolution in Simulated Events	42
4.3	Track Resolution Model	45
4.4	MCRW Fits to Simulated Track Resolution	48
5	Track and Event Selection Efficiency	55
5.1	Event Selection	55
5.2	Lepton Selection	58
5.3	Tracking Efficiency Tables	59
5.3.1	ChargedTracks Efficiency Table	60
5.3.2	Lepton Reconstruction Efficiency Tables	60
5.4	Lepton Identification Efficiency Tables	64
6	Monte-Carlo Template Preparation	69
6.1	EvtGen Template	69
6.2	Tertiary Lepton Template	72
6.3	Continuum Template	74
6.4	Template Event Counts	74
7	The MCRW Fit	75
7.1	Simulated Data Sample and Binning	75
7.2	Fit Parameters	77
7.3	Template Event Weights	78
7.4	Fit Results	80
8	Systematics Studies	87
8.1	B Lifetimes	87
8.2	Cascade Leptons	89
8.3	Track Resolution Model	90
8.4	Tertiary Leptons	90
8.5	Continuum Backgrounds	92
8.6	Optimizing Event Selection Cuts	94
8.7	Systematics Summary	97
9	Conclusions	99
9.1	Summary	99
9.2	Outlook - The MCRW Fit Technique	100
9.3	Outlook - Measuring Δm	101
	Bibliography	103
	Appendices	104
A	Abbreviations and Acronyms	105

B Standard Model and Physics of B Meson Oscillations	107
B.1 Experimental Particle Physics	107
B.2 The Standard Model and the CKM Matrix	109
B.3 Simplified Model of the Neutral B System	111
B.4 CP-Conserving B -mixing	112
B.5 $B^0 - \bar{B}^0$ Mixing and the CKM Matrix	113
B.6 Measuring Δm_d	115
B.7 Measuring B Meson Decay Times	117
B.8 Current Measurements	118
C The <i>BABAR</i> Detector	120
C.1 Overview	120
C.2 Silicon Vertex Tracker	121
C.3 Drift Chamber	122
C.4 Detector of Internally Reflected Cherenkov Light	123
C.5 Electromagnetic Calorimeter	124
C.6 Instrumented Flux Return	125
D Personal Contributions to the <i>BABAR</i> Collaboration	127

List of Tables

2.1	Comparison of example fit results using standard analytical likelihood (ROOT), and the MCRW method, with template parameters equal to the data parameters (V1) and 50% larger than the data parameters (V2) for the example problem in Section 2.6. We give extra digits in the errors in order to compare them more accurately.	21
3.1	Major systematic uncertainties, in ps^{-1} , in the unpublished <i>BABAR</i> result on mixing with dileptons [2].	34
4.1	Values for resolution parameters a, b, c , for both z_0 and d_0 in μm .	45
4.2	Results of a MCRW fit to single track resolution $\Delta z'$ for $\sim 15\text{M}$ fully simulated leptons.	51
6.1	Number of lepton pairs with at least one lepton in the given tertiary category. Note the total number of tertiary pairs is less than the sum of the categories, since a pair can contain more than one tertiary category lepton.	72
7.1	Results of a MCRW fit to 10M simulated lepton pairs in $\Delta t'$ with a template of $\sim 14\text{M}$ lepton pairs. See the text for more details. .	81
8.1	Preliminary estimates of the Δm systematic uncertainties due to fixing one of the two values parameterizing the B lifetimes in the fit. The percent error is based on an assumed value of $\Delta m = 0.489\text{ps}^{-1}$	88
8.2	Preliminary estimates of the systematic uncertainties due to differences in the relative amount of cascade leptons between simulated and real B decays.	89
8.3	Preliminary estimates of the systematic uncertainties due to the fixed parameters in our track resolution model.	90
8.4	Preliminary estimates of the systematic uncertainties due to the tertiary leptons, with the finalized cuts applied.	91
8.5	Preliminary estimates of the systematic uncertainty due to an assumed 1% uncertainty in the fixed continuum background scale factor.	93

8.6	Preliminary estimates of the systematic uncertainty due to the EvtGen model of reconstructed $NTracks$ and R_2	94
8.7	Preliminary estimates of the total systematic uncertainties on Δm . The total is the quadrature sum.	98
B.1	Systematic uncertainties in the published <i>BABAR</i> result on mixing with dileptons [1].	119

List of Figures

1.1	The time-dependent probability of a B^0 meson remaining a B^0 (red, solid line), or oscillating into a \bar{B}^0 (blue, dashed line). The right plot is on a logarithmic scale.	4
1.2	The probability distribution for decay-time difference for both Opposite-Sign (OS, red, solid line) and Same-Sign (SS, blue, dashed line) lepton pairs produced by the decay of $B^0\bar{B}^0$ pairs.	5
1.3	Reconstructed mass distribution for candidate top-quark events (solid). Also shown are the background shape (dotted) and the sum of background plus $t\bar{t}$ Monte Carlo for $M_{top} = 175 \text{ GeV}/c^2$ (dashed). The inset shows the likelihood fit used to determine the top mass. Figure taken from CDF [3].	6
2.1	ROOT Likelihood Fit to the smeared-exponential problem. Data is histogram (black), analytic fit is smooth curve (red). Upper left plot is linear scale, upper right plot is logarithmic scale. Lower left plot is residuals (Data-Function), and lower right plot is pulls (Residuals/Bin Error).	22
2.2	MCRW Fit to the smeared-exponential problem, using the template with longer lifetime (V2). Upper left, data histogram is black, scaled template (mq) histogram is red, weighted template (mwq) is blue. Upper right, average weight w is shown with error bars multiplied by 5. Lower left, residuals are Data-Fit. Lower right, pulls are Residuals/Errors.	23
3.1	Spectra, normalized to unit area, of MC truth electrons and muons, divided into direct from B decays (open histogram), and all other sources (shaded histogram) in fully simulated generic B decays.	27
3.2	Overlay of true Δt for the following categories: Direct-Direct $B^0\bar{B}^0$, Direct-Direct B^+B^- , Direct-Cascade $B^0\bar{B}^0$, etc. OS pairs only. Both leptons satisfy $p_{cm} > 0.5 \text{ GeV}/c$, and $0.5 < \theta < 2.5$	28
3.3	Overlay of true Δt for the following categories: Direct-Direct $B^0\bar{B}^0$, Direct-Direct B^+B^- , Direct-Cascade $B^0\bar{B}^0$, etc. SS pairs only. Both leptons satisfy $p_{cm} > 0.5 \text{ GeV}/c$, and $0.5 < \theta < 2.5$	29

3.4	Fitted values of Δm in MC, as a function of various lab θ cuts. The vertical line represents the MC generator value used for Δm_d . Taken from [2].	35
4.1	Schematic representation of the collision of an e^+e^- pair to create an $\Upsilon(4S)$, followed by the decay into a $B^0\bar{B}^0$ pair. Both B mesons decay, emitting direct leptons ($\ell^+\ell^-$), and the decay-point difference, Δz , is measured. The beam spot is actually much longer (~ 1.5 cm), than the flight length of the B mesons (~ 100 - 200 μm).	41
4.2	The difference between measured and true d_0 for true electrons in simulated $B\bar{B}$ events. The “truth-matching” cuts have a large effect on the electron d_0 tails.	43
4.3	The difference between measured and true d_0 (in cm) for true muons in simulated $B\bar{B}$ events. The “truth-matching” cuts have almost no effect on the muon d_0 tails.	43
4.4	The difference between measured and true z_0 (in cm) for true electrons in simulated $B\bar{B}$ events. The “truth-matching” cuts have a small effect on the electron z_0 tails.	44
4.5	The difference between measured and true z_0 (in cm) for true muons in simulated $B\bar{B}$ events. The “truth-matching” cuts have almost no effect on the muon z_0 tails.	44
4.6	Example of the application of the transformation from Δz (left) to $\Delta z'$ (right). Note the different horizontal and vertical scales.	50
4.7	Overlay of data and template in the MCRW fit to $\Delta z'$. Plots show three p_{cm} bins for electrons, and three p_{cm} bins for muons. The solid line represents the template, and the points with errors are the data.	52
4.8	Pulls ((Data-Template)/Error) of the MCRW fit to $\Delta z'$ (right). Plots show three p_{cm} bins for electrons, and three p_{cm} bins for muons.	53
5.1	The number of Charged Tracks $NTracks$ for a sample of offpeak data (shaded histogram) and $B\bar{B}$ MC events (open histogram). The offpeak data is scaled to match the integrated luminosity of the MC.	56
5.2	The quantity R_2 for a sample of offpeak data (shaded histogram) and $B\bar{B}$ MC events (open histogram). The offpeak data is scaled to match the integrated luminosity of the MC.	57
5.3	The measured ChargedTracks efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.	61
5.4	The measured ChargedTracks efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.	61
5.5	The measured ChargedTracks efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.	62

5.6	The measured muon reconstruction efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.	63
5.7	The measured muon reconstruction efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.	63
5.8	The measured muon reconstruction efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.	64
5.9	The measured electron PID efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.	65
5.10	The measured electron PID efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.	66
5.11	The measured electron PID efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.	66
5.12	The measured muon PID efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.	67
5.13	The measured muon PID efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.	67
5.14	The measured muon PID efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.	68
7.1	A sample transformation from the Δt (left) to $\Delta t'$ (middle) variables, for the same group of fully simulated $B\bar{B}$ muon pairs. The transformed histogram is plotted versus Δt on the right, with the variable bin widths resulting from the transformation. These plots show only the High-High bin, with both muons having $p_{cm} > 1.8$ GeV/c. The lower (red) histogram is the SS pairs, the upper (black) histogram is OS pairs.	76
7.2	The results of a fit to fully simulated ee pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.	83
7.3	The results of a fit to fully simulated $e\mu$ pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.	84
7.4	The results of a fit to fully simulated μe pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.	85
7.5	The results of a fit to fully simulated $\mu\mu$ pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.	86
8.1	Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Tertiary leptons, as a function of the lower limit on p_{cm} . Curves are shown for fits using both the Loose and Tight PID selectors.	95

8.2	Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Tertiary leptons, as a function of the lower limit on θ . Curves are shown for fits using both the Loose and Tight PID selectors.	96
8.3	Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Continuum background, as a function of the minimum allowed $NTracks$ value, for two different requirements on R_2	97
9.1	Measurements of Δm in neutral B meson mixing. Statistical error bars are black, systematic error bars are red. Our projected errors are centered on the PDG average value.	102
B.1	Unitarity triangles constructed from the CKM matrix in a) standard quark-mixing parameters and b) Wolfenstein parameterization.	111
B.2	Dominant box diagrams for $B^0 \rightarrow \bar{B}^0$ transitions. Similar diagrams exist where the t quark is replaced by u or c quarks, but they make much smaller contributions to the overall amplitude.	114
B.3	Published measurements of Δm_d from the B Factories, <i>BABAR</i> , and <i>BELLE</i> . From http://www.slac.stanford.edu/xorg/hfag/	116
C.1	Schematic view of the PEP-II accelerator at SLAC.	120
C.2	Cross-sectional view of the <i>BABAR</i> detector at the PEP-II accelerator. All dimensions are in mm.	121
C.3	Cross-sectional view of a fused silica radiator in the DIRC, showing the internal reflections of the Cherenkov photons.	124

Acknowledgements

I want to start by thanking all of my physics professors at both SFU and UBC who did such a wonderful job over the years of first teaching me the subject, and then welcoming me into the field. Without the passion and enthusiasm of those teachers, I would never have pursued a PhD in physics.

I wish to thank all of my colleagues in the *BABAR* Collaboration, who made many contributions, large and small, that allowed this project to be realized, along with all of the other analyses performed on the *BABAR* data set. Specifically I would like to thank the members of the UBC *BABAR* group: fellow students Doug Thiessen, Neil Knecht, Nasim Boustani, and Zhongxi Song for both teaching and learning with me the software tools needed for analyzing *BABAR* data. Special thanks to Doug Maas who set up a stable computing environment at UBC, and was always willing to help in any way he could. A warm thank you to Chris Hearty and Janis McKenna who were always ready to offer comments and advice over the years.

I also wish to thank Chih-hsiang Cheng, who has followed my work on this analysis for several years, and always had useful, supportive comments and suggestions.

Of course, this project would never have been realized without all of the time and energy that Tom Mattison, my supervisor, put into both my education and this work. It has been a privilege and a pleasure to work so closely, and so well with him during this time.

Finally, I need to give special thanks to all of my family, who have continued to encourage and support me, even if they didn't understand exactly what I have been working on, or why. And finally, I cannot express in words how much I love and appreciate my wife Elissa and daughter Emma, who have been a source of strength for me, and have been so patient while I worked on this thesis.

1

Introduction

1.1 Particle Physics Data and Simulations

In many fields of science, recorded data are simply related to the quantities which one wishes to study. For example, there may be a limited resolution in a measurement device, such as a balance, but the measurement readings are directly and simply related to the mass which is sought.

In experimental particle physics the relationship between the recorded data and the quantities which one wishes to study is much more complex. One of the fundamental units of particle physics data, especially in a collider experiment, is the event. An event can be thought of as a sort of electronic snapshot taken for an extremely brief moment in time. The items which appear in the snapshot are composed out of millions of electronic signals which have been extensively processed, filtered and combined before they can be interpreted as signs of the passage of particles. An electronic or software trigger is used to only fire the “camera” or the data recording system on events which look likely to contain phenomena of interest, in order to suppress the otherwise enormous amounts of background which would overwhelm the system.

Once an event has been recorded, the reconstruction process carefully assembles all of the millions of signals from individual electronic channels into larger composite objects. For example, a series of “hits” in the active regions of a segmented silicon detector can be interpreted as the trace of one portion of the path of a charged particle which traversed the detector. These composite objects are themselves combined into yet higher objects which can represent some of the known particles, such as electrons, protons, muons, etc.

Finally, these higher level versions of the events are analyzed in a systematic way to obtain a measured distribution for some quantity of interest. A decay lifetime, for example, can be determined by looking at the average length of time between the creation and destruction of a particular particle species, once that particular particle type has been identified and isolated in all of the recorded events. These data are often analyzed in the form of a histogram, to facilitate comparison with theoretical distributions.

In order to analyze these complicated, multi-level data samples, it is necessary to implement detailed simulations of the experiments, in order to compare the simulated outputs with the measured quantities. In these “Monte-Carlo” simulations, it is possible to generate an underlying physics event, and then simulate the passage of the particles through the detector material, the creation of the electronic signals, and their reconstruction into the higher-level objects

which are eventually interpreted as particle candidates. By using simulated events, the researcher can access the underlying “truth” of the event, as well as the observed quantities, in order to establish a detailed model of the experimental acceptance and resolution. The complex, highly abstract nature of particle physics data creates unique challenges in its analysis, and the methods in this Thesis were developed with these particular challenges in mind.

In general, when we try to measure a quantity in particle physics, (or any other physics experiment), we use an idealized model of the experiment to write down a tractable, analytic model for the results, that can then be used in either a χ^2 or Maximum Likelihood fit to extract parameter values. This analytical model almost always involves simplifications of the real world details and complications of the experiment. If the statistics are relatively poor, then these simplifications may not matter, compared to the statistical uncertainty.

For particle physics experiments in the current era of “Precision Flavour Physics”, with very high statistics, it becomes more and more difficult to write an adequately detailed analytical model including all aspects of the signal and backgrounds, and experimental resolution in energy, position or trajectory. Not only is it difficult to choose the right parameterization to accurately reflect the detailed underlying processes which contribute to the measured quantity, but it is also difficult to then choose, measure, or otherwise determine the numerical values for the parameters of the analytical model, since they typically cannot all be simultaneously floated in a fit.

In this thesis, we first explain in detail the “Monte-Carlo Re-Weighting” (MCRW) fit technique (Chapter 2), and then apply it to the problem of measuring the neutral B meson mixing frequency. The MCRW method is applied to fully simulated B meson decays to show how the technique works, and the results that could be expected by an analysis on real data, at a later date.

The rest of this Chapter further discusses the physics of neutral B meson mixing, and the motivation and development of the MCRW fit method.

1.2 Neutral B Meson Mixing

The physics topic that motivated the development of this technique is a measurement of neutral B meson mixing, using the data recorded by the *BABAR* experiment, from 1999–2007. We wish to use the so-called “dilepton” method, involving a very loose selection of events, and very high statistics. In the past, the dilepton analysis was published based on the first two years worth of *BABAR* data, and was statistically limited [1]. However, when a later attempt was made to use the larger dataset then available, the researchers found severe limitations in the analytical model that they were using to model the data [2]. The approximations inherent in the chosen PDFs (Probability Density Functions) resulted in systematic uncertainties as large, or larger, than the expected statistical errors. The present work represents an attempt to overcome these limitations and utilize the full statistical power available in the complete *BABAR* data set.

Neutral B meson mixing, or oscillations, is a quantum-mechanical phenomenon in which a matter particle (B^0) can transform itself into an antimatter particle (\bar{B}^0). For more details of the way in which this transformation occurs, and its relationship to the Standard Model of Particle Physics, please consult Appendix B.

The neutral B mesons, B^0 and \bar{B}^0 are too short-lived to observe directly, with lifetimes of only ~ 1.5 ps, so we must study them through their decay products. If a neutral B meson decays semileptonically, it emits a lepton (here we mean either e or μ), a neutrino, and a meson containing a charm quark. The B^0 emits a positively charged lepton (e^+ or μ^+), and the \bar{B}^0 emits a negatively charged lepton (e^- or μ^-). By observing the charge of the lepton, we are able to infer whether the original B meson was a B^0 or \bar{B}^0 at the moment of the decay.

A B^0 or \bar{B}^0 meson oscillates between matter and antimatter, or between the B^0 and \bar{B}^0 state, with a sinusoidal time-dependence. This oscillation has a frequency of Δm , when using “natural” units, such that $\hbar = c = 1$. Δm is the mass or energy difference between the two mass eigenstates of the neutral B system. The oscillation takes place within the envelope of exponential decay imposed by the neutral B meson lifetime, τ^0 . The time dependent probability of a B^0 meson to remain a B^0 , or to oscillate into a \bar{B}^0 is given by the following equations:

$$Prob(B^0 \rightarrow B^0) \propto e^{-t/\tau^0} (1 + \cos(\Delta m t)), \quad (1.1)$$

$$Prob(B^0 \rightarrow \bar{B}^0) \propto e^{-t/\tau^0} (1 - \cos(\Delta m t)). \quad (1.2)$$

Numerically, the mixing frequency Δm is approximately 0.5 ps^{-1} , and the lifetime is roughly 1.5 ps , giving the relationship depicted in Figure 1.1 for the time-dependent probability of a B^0 remaining a B^0 or turning into a \bar{B}^0 .

If we examine the decay-time difference of the two mesons, Δt , it contains the same information about the lifetime and mixing frequency of the $B^0\bar{B}^0$ mesons that would be available by studying individual B mesons. Figure 1.2 shows the probability of observing a same-sign (SS) or opposite-sign (OS) pair of leptons separated by decay-time difference Δt , after the creation of a $B^0\bar{B}^0$ pair. A same-sign pair would be either $\ell^+\ell^+$, or $\ell^-\ell^-$, and an opposite-sign pair would be $\ell^+\ell^-$ or $\ell^-\ell^+$, where the ℓ denotes either an electron (e) or muon (μ). The order of the pairs is chosen to be with the higher momentum lepton first.

We can write the probability of observing an OS or SS pair of direct leptons from the decay of a $B^0\bar{B}^0$ pair as:

$$Prob(OS, \Delta t) \propto e^{-|\Delta t|/\tau^0} (1 + \cos(\Delta m \Delta t)), \quad (1.3)$$

$$Prob(SS, \Delta t) \propto e^{-|\Delta t|/\tau^0} (1 - \cos(\Delta m \Delta t)). \quad (1.4)$$

The $B^0\bar{B}^0$ pair created in the decay of the $\Upsilon(4S)$ is in a coherent state, so that as time progresses, they are in an orthogonal superposition of states until one decays. Once the first B mesons decays, the second B meson becomes fixed

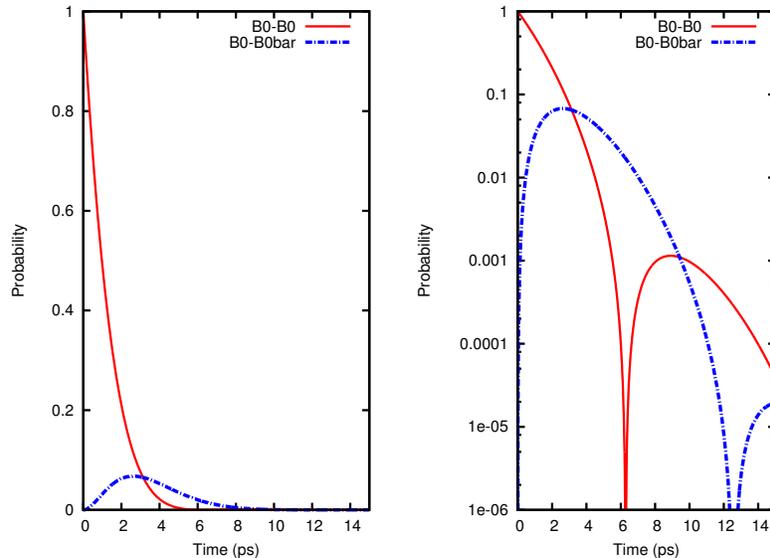


Figure 1.1: The time-dependent probability of a B^0 meson remaining a B^0 (red, solid line), or oscillating into a \bar{B}^0 (blue, dashed line). The right plot is on a logarithmic scale.

as the opposite flavour, and after that point in time it will evolve independently according to the probability shown in Figure 1.1.

If both B mesons decay at the same time so $\Delta t = 0$, or $\Delta t = 2n\pi$ we will always observe an Opposite-Sign (OS) lepton pair, and never a Same-Sign (SS) pair. If the decay-time difference satisfies $\Delta t = n\pi$, with n odd, then we will always observe an SS pair, and never an OS pair. In between these times, we will observe a mixture of OS and SS decays. The essence of a time-dependent mixing frequency measurement is to fit Δt histograms for OS and SS lepton pairs to extract the oscillation frequency.

1.3 The Monte Carlo Template Method

Conceptually, it should be possible to generate Monte-Carlo (MC) samples with different values for a given physical parameter of interest, and compare the agreement between data and MC distributions as a function of the generated parameter value. This method has been used occasionally in experimental particle physics, and is called the MC Template Method.

An illustration of the use of this method to measure the top quark mass, performed by the CDF collaboration, is shown in Figure 1.3 [3]. For each possible value of the mass of the top quark, ranging from 160 to 190 GeV/c^2 , the authors generated a series of evenly-spaced MC templates. They then calculated

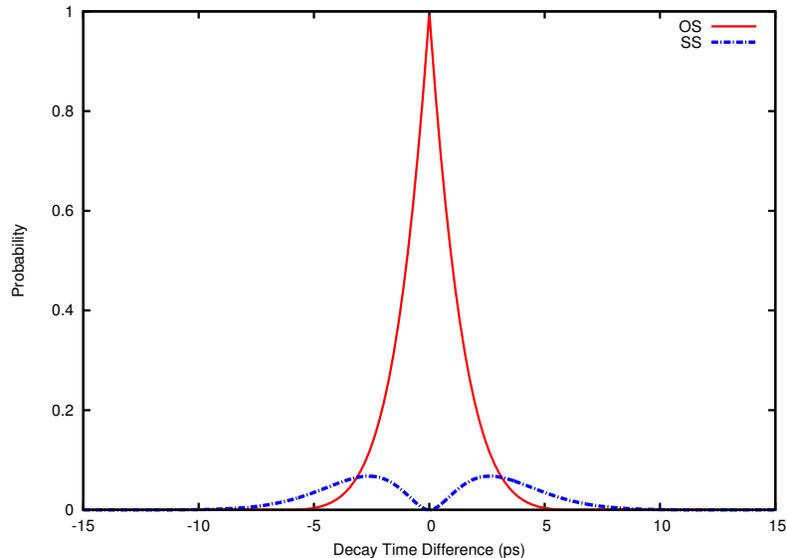


Figure 1.2: The probability distribution for decay-time difference for both Opposite-Sign (OS, red, solid line) and Same-Sign (SS, blue, dashed line) lepton pairs, produced by the decay of $B^0\bar{B}^0$ pairs.

the likelihood that their data sample and the MC template were drawn from the same distribution (*i.e.* they had the same mass) and plotted -2 times the log-likelihood ($-2LL$) versus the simulated mass, as shown in the inset of Figure 1.3. They then fit a parabolic curve to these points, and found the minimum point as their measurement of the top mass. The width of the parabola at a height of one unit in $-2LL$ above the minimum corresponds to the statistical uncertainty of the measurement.

There are several shortcomings to this MC template method. The first is that the user needs to generate many templates, each with many times more events than the data, in order to get the MC statistical error on the likelihood ($-2LL$) differences (treating data as fixed without error) small compared to 1. In other words, the size of the vertical error bars shown in the inset of Figure 1.3 depend on the number of MC events generated in each template, and in order to have sufficiently small errors, this must be significantly greater than the number of data events.

An obvious shortcoming is that the user would need an N -dimensional grid of templates in order to fit for N parameters. Hence the template method is quite impractical for fits with more than two parameters. In fact, later measurements of the top quark mass used much more sophisticated techniques.

The measurement of B mixing with dileptons is a problem that requires more than just one or two parameters. In addition to the leptons coming directly from

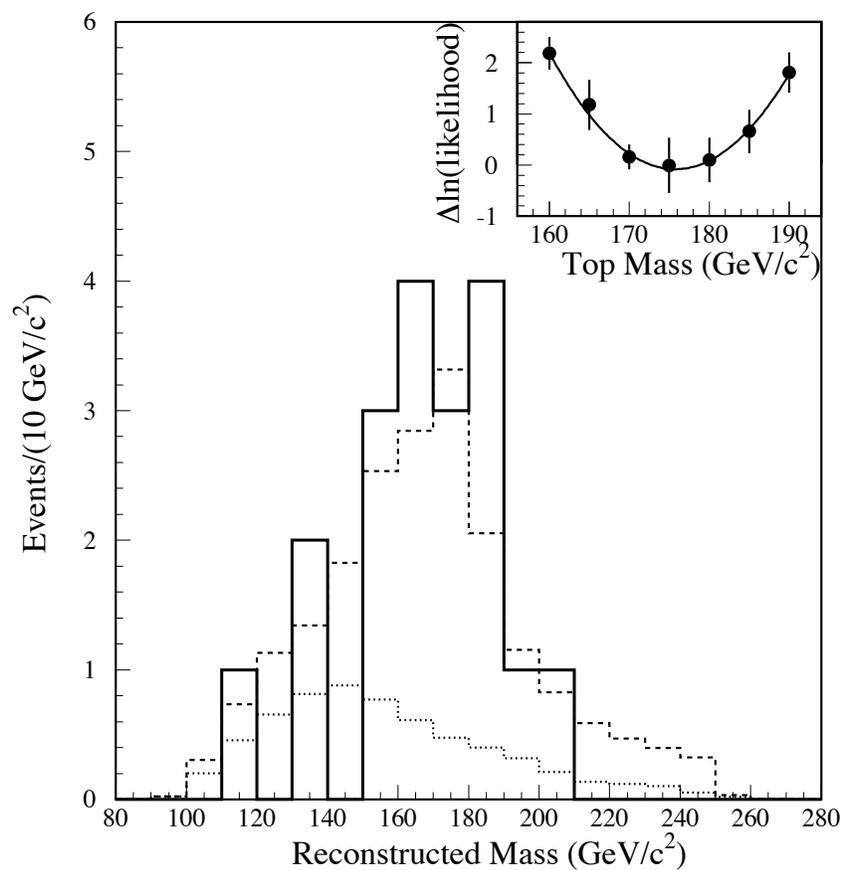


Figure 1.3: Reconstructed mass distribution for candidate top-quark events (solid). Also shown are the background shape (dotted) and the sum of background plus $t\bar{t}$ Monte Carlo for $M_{top} = 175$ GeV/c² (dashed). The inset shows the likelihood fit used to determine the top mass. Figure taken from CDF [3].

the decays of neutral B mesons, there are also charged B pairs that produce lepton pairs with no mixing and decay with the charged B^\pm lifetime, rather than the neutral B lifetime. In addition, the relative branching fractions of the $\Upsilon(4S)$ to neutral and charged B s aren't known precisely enough for the desired experimental sensitivity. There are also secondary leptons from decays such as: $B \rightarrow D \rightarrow \ell$, $B \rightarrow \tau \rightarrow \ell$, and other processes with poorly known branching fractions. We would like to have parameters available to measure the relative contributions of the direct and secondary leptons. The detector measures track positions with finite resolution so a large part of the smearing of the measured decay time difference, Δt , arises from the track resolution along the z axis of the detector. We would like to have parameters floating which allow the track resolution to be determined in the same fit. Fitting for this many parameters would be extremely impractical with the MC Template method.

1.4 Barlow-Beeston Method

A first step toward multiple parameter Monte-Carlo based fits is represented by the method of Barlow and Beeston [4]. This method is implemented as `TFractionFitter` inside the ROOT [5] data analysis software framework. This method essentially builds a template from an arbitrary number of MC distribution shapes, and then determines the fraction of the data contained in each MC category. The only parameters it can measure are the fraction for each component, but this method supports a large number of parameters, and requires only a single MC template per component. It uses Poisson statistics, so it works even with low-statistics or empty bins, in both the data and template. Support for Poisson statistics requires introducing an additional parameter for each bin of each template histogram, as well as the “physics” fraction parameters. This would be an impractical number of parameters for a conventional minimizer to deal with, but for fixed values of the “physics” parameters it is possible to solve bin by bin for these extra ones, adding only simple calculations to each iteration of the fit. The template statistics need only be comparable to the data statistics, although higher statistics lowers the fit errors.

However, the Barlow-Beeston method is limited to linear combinations of fixed component shapes, and cannot handle more general fitting problems, in which some of the fitted parameters determine the shape of the distribution, rather than just the normalization.

1.5 The Monte-Carlo Re-Weighting Method

Our approach has been to develop a new method for estimating physical parameters from experimental data, using a Maximum Likelihood fit that compares the data distribution to a “Monte Carlo Template”, rather than an analytical model. For simulated events it is possible to record both the true quantities of the particles generated in the event, and the reconstructed quantities that

represent the response of the detector hardware and software. In the method presented in this Thesis, we keep a large number of simulated events in memory, and for each iteration of the fit, we use the stored truth information to re-weight each simulated event based on the probability functions of the underlying physical models, and their relations to the parameters that are floating in the fit. The re-weighted Monte Carlo (MC) events are then used to recalculate the template distribution, by using the reconstructed quantities to determine in which bin of the template histogram a given event will appear, in the same way that data events are placed into a data histogram.

When we re-weight the template, we are essentially applying a probability ratio to each event in the template, which is the ratio of the probability of generating the event with the current estimate of the parameter values (\vec{p}_{fit}), to the probability of generating the event with the values of the parameters which were actually used to produce the template (\vec{p}_{gen}), as given in Equation 1.5.

$$w_{event}(\vec{p}_{fit}) = \frac{F(\vec{p}_{fit}, \vec{t}_{event})}{F(\vec{p}_{gen}, \vec{t}_{event})} \quad (1.5)$$

Here the function $F(\vec{p}, \vec{t})$ is a Probability Density Function (not necessarily normalized in our case) for obtaining the true MC quantities \vec{t} , given the set of parameters \vec{p} .

Re-weighting the template in this way allows us to compare the data to an arbitrary template distribution without having to spend the time needed to regenerate the simulated events for each iteration of a fit. The fit uses binned data and a binned template, so it works best with sufficient data and simulated events to populate all the bins in the fitted regions. The use of Poisson statistics allows the fit to treat empty, or low-statistics bins correctly.

By fitting to a template of simulated events, we allow the detailed calculations of the simulation to be included in the fit, without having to apply a simplifying analytical model that approximates the distribution of the observed output variable(s). We need to be able to write probability functions for each of the underlying physical mechanisms that we are interested in isolating or measuring in our analysis, but these should typically have relatively simple forms, such as exponentials, sinusoids, etc. for particle physics processes.

This method requires that the user has a very accurate model of the underlying physics and the experimental conditions established, but this is generally the case in all modern particle physics experiments, which tend to have large teams of researchers working together to develop these detailed simulations, and collect enormous amounts of data over periods of many years.

The MCRW method we have developed requires just a single MC template sample, but smoothly re-weights the events to different parameter values. The re-weighting requires that we store enough ‘‘MC truth’’ about each template event so that we can calculate the relative probability of such an event with different parameter values. This requires a large amount of memory and computation, which would have been impractical in the past, but is now possible with the computing resources that have become available in recent years. One

of the goals of this thesis is to determine if this method is practical for a high-statistics data sample. The dilepton B mixing analysis uses about 10 million lepton pairs, so it is an excellent stress-test for this algorithm.

A fitting method must also be able to handle complex physics models and complex detector models. In fact, this is another advantage of the MCRW method, because it is usually much easier to write a MC simulation of a complex process than to determine an analytic expression for the observables. Modern particle physics experiments have collaborations of hundreds of researchers, and often devote hundreds of person-years of effort to develop detailed MC simulations. These highly detailed MC simulations can be incorporated into a physics fit with minimal extra effort in the MCRW method. Another goal of this thesis is to determine if the full complexity of the B decay physics, background physics, and *BABAR* detector simulations required for the dilepton B mixing analysis can be introduced into the physics fit, in a simpler and less ad-hoc way than in an analysis using more conventional fitting methods.

1.6 Outline of This Thesis

We briefly mention the contents of the remaining Chapters of this Thesis.

Chapter 2 is a derivation of the algorithm used to implement the Monte-Carlo Re-Weighting fit procedure, as well as a discussion of its strengths and weaknesses. Chapter 3 is an overview of the “dilepton” neutral B meson mixing analysis, and the challenges involved in it, as well as a summary of the way in which we approach this analysis with the MCRW fitting method. Chapter 4 is a discussion of the model of tracking resolution that is used in our implementation of the MCRW fitting procedure. Chapter 5 is a discussion of the procedure used to select tracks and events of interest from the large dataset of real or simulated *BABAR* events in order to obtain a “dilepton” sample and carry out a measurement of neutral B meson mixing. It also contains a discussion of the measurement of various efficiencies in simulated data that are used in the creation of the MCRW template. Chapter 6 is a description of the procedure involved in assembling the three components of the template together in order to perform a MCRW fit to either real or simulated events from the *BABAR* experiment, to extract the neutral B mixing frequency. Chapter 7 describes how the Δt variable is transformed for more convenient plotting and fitting, the binning that is used to create the histograms for the MCRW fit, the weights that are applied to the template events for the fit, and the performance of a fit to a sample of 480 M simulated $B\bar{B}$ pairs. Chapter 8 describes the procedures used to calculate the sensitivity to systematic uncertainties that are expected in an analysis of neutral B meson mixing, using the dilepton sample. It also contains a discussion of the procedures used for optimizing the selection cuts for the MCRW B mixing with dileptons analysis. Chapter 9 summarizes this thesis, and discusses the outlook for two aspects of the work.

Appendix B discusses in more detail the theory of the Standard Model of particle physics, and explores the phenomenon of neutral B meson mixing within

the Standard Model, looking briefly at how B mixing can be used to search for new physics. Appendix C briefly discusses the design, operation and performance of the PEP-II accelerator at SLAC National Laboratory in the US, and the *BABAR* detector which was located there. Appendix D details the service work, or contributions that the author made to the operation of the *BABAR* detector, and data analysis efforts of the experimental collaboration.

2

Monte-Carlo Re-Weighting Fit Method

In this Chapter we first present the mathematical formulation of the likelihood that the data and template are drawn from the same distribution, followed by the mathematical calculations needed to obtain the set of equations that must be solved in order to maximize the likelihood (or minimize the negative log-likelihood) with respect to the parameters. We also present the algorithm of the software implementing this fitting method, which is used for the study detailed in later Chapters, as well as discussing its performance in comparison to more familiar methods. Finally, we present a simple example problem that can be treated exactly in the traditional analytical likelihood method, and compare it to the treatment with the Monte-Carlo Re-Weighting (MCRW) Method.

2.1 Likelihood Formulation

We begin by deriving the log likelihood function, $\log(\mathcal{L})$, describing a comparison of a Monte-Carlo template histogram to a data histogram, for each bin i of the N -bin histogram.

If a given bin in the data histogram contains n_i events, and the model predicts an expected value of x_i events in the bin, where x_i is a function of the parameters p_j or \vec{p} , then we can write the Poisson probability of observing n_i events when we expect to observe x_i events as:

$$P_i = e^{-x_i} \frac{x_i^{n_i}}{n_i!} \quad (2.1)$$

Taking the logarithm of the probability P_i , we obtain:

$$\log(P_i) = n_i \log(x_i) - x_i - \log(n_i!) \quad (2.2)$$

If we sum over all N bins of the histogram, the logarithm of the probability (or likelihood) is:

$$\log(\mathcal{L}) = \sum_{i=1}^N [n_i \log(x_i) - x_i - \log(n_i!)] \quad (2.3)$$

If the expected values x_i are trivial functions of a small number of parameters p_j , this log-likelihood can be minimized as a function of the parameters analytically. If, however, the x_i are moderately complicated or if there are many

parameters, it is still possible to perform the minimization using a standard numerical algorithm such as MINUIT [6]. The resulting parameter estimates have uncertainties which depend on the statistics of the data histogram, and decrease roughly as the square root of the total number of events in the data histogram.

In experimental particle physics, the model values x_i are typically the sum of a “signal” component (which can often be described relatively simply), and a “background” component (which can often be fairly complicated). Both of these components are also convoluted with the acceptance and resolution of the experimental apparatus, or detector, which is normally very complicated. If it is even possible to write down an exact analytical expression for the model expectations x_i , the expression will be very complicated. Since it is often not possible to write an analytic expression, approximate expressions are frequently used instead.

Particle physics collaborations put a large amount of effort into a shared Monte Carlo (MC) simulation of the signal physics, background physics, and the detector. These simulations are often used to derive approximate expressions for the model x_i values. The process of determining these approximations is itself a substantial effort, and may be different for each data analysis within the same experimental collaboration. The approximate model has both statistical errors from finite MC statistics, and systematic uncertainties from the approximate functional forms used, since the functional forms used must be simple enough for the final likelihood minimization to be practical on finite computing resources. A substantial effort is required to evaluate these statistical and systematic uncertainties for each approximate model.

If it were possible to incorporate the MC simulation directly in the fit, the user could avoid the systematic uncertainties caused by using simplifying approximate expressions, and also avoid the need to do intermediate fits to determine those approximations. The finite MC statistics will introduce an additional statistical uncertainty beyond that of the data, but it turns out that this complication can be handled in the fitting procedure without much difficulty.

The MC simulation must be run with particular choices made for the parameters \vec{p}_{gen} . Each simulated event will be described by a truth vector \vec{t} . This contains the true values which were randomly generated for the event, for example particle decay times and positions. The probability of generating an event with the given values of its truth vector for the chosen parameter values is some function $F(\vec{p}, \vec{t})$. It is important to note that this function is normally a product of a number of relatively simple functions describing the physics and the detector resolution, rather than a more complicated convolution. Parameters describing the detector resolution can be included as generator parameters. The given event is placed in bin i of the observable(s) histogram, and a total of m_i events appear in that bin. Obviously, the simulation must be binned in exactly the same way as the data, in order to be useful in determining the prediction x_i . The simulation may have more or fewer events in it than the actual data sample, so we introduce an overall normalization parameter q . Now we can write the predicted number of events in bin i for the chosen parameter values

as $x_i = m_i q$.

If we had chosen different parameter values to generate the MC, a given event would have been generated with a higher or lower probability. We define the weight of a given event as:

$$w_{event}(\vec{p}_{fit}) = \frac{F(\vec{p}_{fit}, \vec{t}_{event})}{F(\vec{p}_{gen}, \vec{t}_{event})} \quad (2.4)$$

If the current estimate of the parameters in the fit, \vec{p}_{fit} , are identical to the generator parameters \vec{p}_{gen} , then all events have a weight of exactly 1. The mean weight of the events in bin i is:

$$w_i(\vec{p}_{fit}) = \frac{1}{m_i} \sum_{event=1}^{m_i} \frac{F(\vec{p}_{fit}, \vec{t}_{event})}{F(\vec{p}_{gen}, \vec{t}_{event})} \quad (2.5)$$

where each event has a different truth vector \vec{t} . It is now possible to write the predicted number of data events as a function of the fit parameters:

$$x_i(\vec{p}_{fit}) = q m_i w_i(\vec{p}_{fit}) \quad (2.6)$$

This prediction now has two distinct sources of statistical uncertainty. The number of simulated MC events generated fluctuates with Poisson statistics, and the mean event weight in a bin also fluctuates due to variations in \vec{t}_{event} . We assume that the mean weight obeys Gaussian statistics with σ_i calculated from the distribution of event weights.

We can then introduce two new variables: y_i representing the unknown number of expected MC template events in bin i with the generator parameter values \vec{p}_{gen} , and v_i representing the unknown expected mean weight of events in bin i for parameter values \vec{p}_{fit} . Now we can write the predicted number of data events in bin i as function of y_i and v_i :

$$x_i(\vec{p}_{fit}) = q y_i v_i \quad (2.7)$$

The combined probability of observing n_i data events when we expect x_i , observing m_i MC events when we expect y_i , and measuring mean weight w_i when we expect v_i for a given bin i is given by:

$$P_i = e^{-x_i} \frac{x_i^{n_i}}{n_i!} \cdot e^{-y_i} \frac{y_i^{m_i}}{m_i!} \cdot \frac{e^{-(w_i - v_i)^2 / (2\sigma_i^2)}}{\sqrt{2\pi}\sigma_i} \quad (2.8)$$

We can now write the logarithm of the total likelihood of all the observations in all of the bins, ranging from $i = 1..N$ as:

$$\log(\mathcal{L}) = \sum_{i=1}^N \left[n_i \log(q y_i v_i) - q y_i v_i - \log(n_i!) + m_i \log(y_i) - y_i - \log(m_i!) - \frac{(w_i - v_i)^2}{2\sigma_i^2} - \frac{\log(2\pi)}{2} - \log \sigma_i \right] \quad (2.9)$$

This log-likelihood (LL) is a function of the parameters \vec{p} , since the mean weight w_i and dispersion σ_i depend explicitly on \vec{p} :

$$w_i(\vec{p}_{fit}) = \frac{1}{m_i} \sum_{events} w_{event}(\vec{p}_{fit}) = \frac{1}{m_i} \sum_{events} \frac{F(\vec{p}_{fit}, \vec{t}_{event})}{F(\vec{p}_{gen}, \vec{t}_{event})} \quad (2.10)$$

$$\sigma_i^2(\vec{p}_{fit}) = \frac{1}{m_i(m_i - 1)} \left(\sum_{events} w_{event}^2 \right) - w_i^2 / (m_i - 1) \quad (2.11)$$

By minimizing the negative log-likelihood (or -2LL), we obtain the parameter values which give the best agreement between the measured data and the MC simulation. However, we must minimize with respect to not only the parameters \vec{p} , but also the unknown v_i and y_i for each bin of the MC simulation.

At first glance the minimization problem seems more difficult than normal due to the presence of the two additional unknown parameters y_i and v_i for each bin. The problem is not, however, underdetermined, because we also have two extra ‘‘measurements’’ per bin: the observed m_i and w_i from the simulation. We can determine the y_i and v_i values by taking the partial derivatives of Equation 2.9 with respect to y_i and v_i , while treating w_i , σ_i and q as constants, and then equating them each to zero, resulting in the following equations:

$$\frac{d}{dy_i} \log(\mathcal{L}) = n_i/y_i - v_i q + m_i/y_i - 1 = 0 \quad (2.12)$$

$$\frac{d}{dv_i} \log(\mathcal{L}) = n_i/v_i - y_i q - (w_i - v_i)/\sigma_i^2 = 0 \quad (2.13)$$

Note that these equations are no longer summed over all the bins, so we can solve these two equations for y_i and v_i separately for each bin. We first solve Equation 2.12 in terms of y_i to obtain (dropping the subscript i for the moment, for brevity):

$$y = \frac{m + n}{1 + vq} \quad (2.14)$$

We then substitute this expression into Equation 2.13, and rearrange the terms to obtain:

$$\frac{n - mvq}{v(1 + vq)} - \frac{v - w}{\sigma^2} = 0 \quad (2.15)$$

which is a nonlinear equation in v . This equation always has a positive, real solution, given the allowed values of n , m , q , w and σ . Normally the solution will be close to $v \sim w$, within a few σ , and the exact value will depend on the level of agreement between n and mq . Once we have thus calculated v for each bin, we can use it to trivially determine the value of y for each bin, using Equation 2.14.

Given values for \vec{p} and q , we can solve for the v_i and y_i in each bin, and calculate $-\log(\mathcal{L})$ as a function strictly of \vec{p} and q .

2.2 Likelihood Maximization

This function can then be minimized by a conventional program like MINUIT [6] to find the best-fit parameter values and their errors. Calculating $-\log(\mathcal{L})$ for a given set of parameter values \vec{p}_{fit} requires the following steps:

- Loop over all the events in the MC sample, and calculate $F(\vec{p}_{fit}, \vec{t}_{event})$ from the stored event truth and the current \vec{p}_{fit} values, and accumulate the sums of Equation 2.10 and 2.11 for each bin in the histogram. To save computation time, the bin number in which the event falls, and $F(\vec{p}_{gen}, \vec{t}_{event})$ for each event can be calculated once and stored rather than being recalculated each time, since they do not depend on the changing parameter values of \vec{p}_{fit} .
- Loop over the bins, calculating w_i and σ_i for each, and then solving Equation 2.15 for v_i and calculating y_i . Then accumulate the contribution of each bin to the $-\log(\mathcal{L})$ sum.

The loop over all of the MC events normally dominates the computation time, since each event is treated individually, and there are typically as many or more MC events as data events, and one typically has many more data events than there are bins. The numerical solution for the v_i is normally fast in comparison, and the rest of the calculation time is quite negligible in comparison.

The MINUIT program has a long and distinguished record, but it is not particularly economical about the number of function evaluations it makes. For K parameters, it evaluates the function about K^2 times while initializing, then K plus a few times per iteration, for a number of iterations that ranges from K to a large number (depending on how non-linear the problem is, and the distance of the starting point from the minimum). In contrast, if the function being minimized is a sum of squares of quantities that are linear in the parameters (*e.g.*, χ^2 for a linear combination function), a program that exploits the linearity of the function only requires $K + 1$ function evaluations (basically to determine the derivatives) to find the minimum from an arbitrary starting point.

Near the solution, the $\log(\mathcal{L})$ should be nearly parabolic, and a linearized iterative solution should be possible and much faster than MINUIT. If, due to non-linear behaviour, the parameter update degrades $\log(\mathcal{L})$ rather than improving it, the Levenberg-Marquardt scheme [7] can be used find a shorter step in a steeper direction that is guaranteed to improve $\log(\mathcal{L})$.

We find the parameter derivatives of $-\log(\mathcal{L})$, treating q as an extra component of $\vec{p} = p_j$, and using the notation $D_i = w_i - v_i$ and $s_i = \sigma_i^2$ for compactness. The w_i and s_i values depend on \vec{p} , and because the solutions for v and y depend on w and s they also have \vec{p} derivatives, and therefore $x = qvy$ also has a \vec{p} derivative.

$$\frac{d}{dp_j} (-\log(\mathcal{L})) = \sum_{i=1}^N \left[\left(1 - \frac{n_i}{x_i}\right) \frac{dx_i}{dp_j} + \left(1 - \frac{m_i}{y_i}\right) \frac{dy_i}{dp_j} + \frac{D_i}{s_i} \frac{dD_i}{dp_j} - \frac{D_i^2}{2s_i^2} \frac{ds_i}{dp_j} \right] = R_j \quad (2.16)$$

For K parameters (including q), there are K such derivative values R_j . To calculate them requires the following derivatives: $dD_i/dp_j = dw_i/dp_j - dv_i/dp_j$, dy_i/dp_j , $dx_i/dp_j = d(qv_i y_i)/dp_j$, and ds_i/dp_j .

We calculate the derivatives of w_i and s_i directly from the Monte Carlo events

$$\frac{dw_i}{dp_j} = \frac{1}{m_i} \sum_{events} \frac{dw_{event}}{dp_j} = \frac{1}{m_i} \sum_{events} \frac{\frac{dF(\vec{p}_{fit}, \vec{t}_{event})}{dp_j}}{F(\vec{p}_{gen}, \vec{t}_{event})} \quad (2.17)$$

$$\frac{ds_i}{dp_j} = \frac{2}{m_i(m_i - 1)} \left(\sum_{events} w_{event} \frac{dw_{event}}{dp_j} - m_i w_i \frac{dw_i}{dp_j} \right) \quad (2.18)$$

If $F(\vec{p}_{fit}, \vec{t}_{event})$ is relatively simple, the derivatives can be done analytically, otherwise the derivatives can always be done numerically with K extra function evaluations.

We calculate the other derivatives by implicit differentiation of the equations we solve for v and y at the bin level (rather than the event level):

$$\frac{dv_i}{dp_j} = \frac{v_i(1 + qv_i)}{m_i q s_i + 2v_i + 3qv_i^2 - w_i - 2qv_i w_i} \left(\frac{dw_i}{dp_j} + \frac{ds_i}{dp_j} \frac{(v_i - w_i)}{s_i} \right) \quad (2.19)$$

$$\frac{dy_i}{dp_j} = \frac{-y_i q}{1 + qv_i} \frac{dv_i}{dp_j} \quad (2.20)$$

$$\frac{dx_i}{dp_j} = qv_i \frac{dy_i}{dp_j} + qy_i \frac{dv_i}{dp_j} \quad (2.21)$$

At the minimum of -2LL, all of the R_j are zero. If we are near the solution, the R_j are not zero, but we can linearize them to find the parameter steps that should make them all zero:

$$\begin{aligned} \frac{dR_j}{dp_k} = & \sum_{i=1}^N \left[+ \frac{dx_i}{dp_j} \frac{n_i}{x_i^2} \frac{dx_i}{dp_k} + \left(\frac{n_i}{x_i} - 1 \right) \frac{d^2 x_i}{dp_j dp_k} + \frac{dy_i}{dp_j} \frac{m_i}{y_i^2} \frac{dy_i}{dp_k} + \left(\frac{m_i}{y_i} - 1 \right) \frac{d^2 y_i}{dp_j dp_k} \right. \\ & + \frac{dD_i}{dp_j} \frac{1}{s_i} \frac{dD_i}{dp_k} - \frac{D_i}{s_i^2} \frac{dD_i}{dp_j} \frac{ds_i}{dp_k} + \frac{D_i}{s_i} \frac{d^2 D_i}{dp_j dp_k} \\ & \left. - \frac{D_i}{s_i} \frac{ds_i}{dp_j} \frac{dD_i}{dp_k} + \frac{ds_i}{dp_j} \frac{D_i^2}{s_i^3} \frac{ds_i}{dp_k} - \frac{D_i^2}{2s_i^2} \frac{d^2 s_i}{dp_j dp_k} \right] = W_{jk} \quad (2.22) \end{aligned}$$

This expression should be symmetric on exchange of j and k . Terms 6 and 8 are not, but their sum is, and the other terms are already symmetric. Most of the terms involve products of first derivatives rather than second derivatives. The factors $n_i/x_i - 1$ and $m_i/y_i - 1$ are close to zero when near the solution, and fluctuate between positive and negative so they tend to average to zero. The same is true for $D_i = w_i - v_i$. So most of the explicit second-derivative terms tend to cancel out in the sum.

Since we will solve the linearized equations iteratively, it is acceptable to approximate the dR_j/dp_k by dropping all the explicit second-derivative terms

(note that we do not approximate the R_j themselves). So we are left with

$$\begin{aligned} \frac{dR_j}{dp_k} = \frac{d^2(-\log(\mathcal{L}))}{dp_j dp_k} &= \sum_{i=1}^N \left[\frac{dx_i}{dp_j} \frac{n_i}{x_i^2} \frac{dx_i}{dp_k} + \frac{dy_i}{dp_j} \frac{m_i}{y_i^2} \frac{dy_i}{dp_k} + \frac{dD_i}{dp_j} \frac{1}{s_i} \frac{dD_i}{dp_k} \right. \\ &\quad \left. + \frac{ds_i}{dp_j} \frac{D_i^2}{s_i^3} \frac{ds_i}{dp_k} - \frac{D_i}{s_i} \frac{ds_i}{dp_j} \frac{dD_i}{dp_k} - \frac{D_i}{s_i} \frac{dD_i}{dp_j} \frac{ds_i}{dp_k} \right] \\ &= W_{jk} \end{aligned} \quad (2.23)$$

Now we solve the K linearized equations:

$$\frac{d(-\log(\mathcal{L}))}{dp_j} \approx R_j + \sum_k W_{jk} \Delta p_k = 0 \quad (2.24)$$

to obtain the the parameter steps:

$$\Delta p_k = - \sum_j (W_{jk})^{-1} R_j \quad (2.25)$$

For each iteration, we need to loop over events to evaluate $F(\vec{p}_{fit}, \vec{t}_{event})$ once for the w and s needed for the $-\log(\mathcal{L})$ value and evaluate $dF/d\vec{p}$ for the K parameter derivatives (or evaluate the F function K more times for numerical derivatives), accumulate sums for w_i and s_i and their parameter derivatives, then loop over bins to solve for v_i , y_i , calculate the additional derivatives, and accumulate the R_j vector and W_{jk} matrix. Then we invert the matrix and calculate the parameter steps. If the parameter steps and change in $-\log(\mathcal{L})$ are sufficiently small, we have found the solution.

Our approach to the minimization requires about K evaluations of F or its derivative per event per iteration, and one numerical solution for v_i per bin per iteration, and for starting points close to the solution typically requires only a few iterations. The MINUIT approach requires about K^2 evaluations of F per event and K^2 solutions for v_i per bin for initialization, about K evaluations of F per event per iteration and K solutions for v_i per bin, per iteration, and often a large number of iterations. So our algorithm is expected to be somewhat faster than MINUIT, particularly if there are a large number of parameters.

It may seem that MCRW fits are highly mathematical, but the math only needs to be done once and coded into a computer program, and conventional fitting programs contain a similar amount of formalism.

For a case where an analytical function describes the data adequately, a conventional fit to the (binned) data will be much faster than an MCRW fit that treats each MC template event individually, no matter what minimization algorithm is used. The MCRW execution time is expected to be similar to the execution time for an unbinned analytical likelihood fit to a data set the same size as the MCRW template.

2.3 Binning Issues

Computer minimization programs require smooth parameter dependence, so we must use the same set of MC events for all parameter values, and all events must remain in the same bins for all parameter values. The bin in which an event falls must be independent of the parameter values, and can only depend on the parameter-independent truth vector \vec{t} for the event. The parameter-dependence is only allowed through the smooth dependence of the mean event weights w_i and the σ_i on \vec{p} . The truth vector for each event must contain enough information to calculate the event weight from Equation 2.4, for arbitrary parameter values.

Note that MCRW fits in the present formalism require the data to be binned, which increases the statistical error somewhat compared to an unbinned fit. The degradation can be minimized by using finer binning, and since the formalism uses Poisson statistics there is no problem in principle with data bins containing only a few or even zero events. However, the MC template must use the same binning, and the formalism breaks down if there are empty template bins where the corresponding data bin is not empty (there is no problem if both data and template bins are empty). If there are only a handful of MC events in a bin, the w_i and σ_i may not be a very good description of the MC prediction for the bin.

The simplest way to avoid the low-population MC bin problem is to have many more MC events than data events. Then any bin that has a data event will typically have many MC events. This will also reduce the contribution to the statistical error from MC statistics, and is always recommended if the MC statistics are available and the fit computation time is acceptable.

If higher MC statistics is not available or practical to use, one can combine adjacent bins until enough MC events are included (and combine the corresponding data bins). Randomly varying bin widths would substantially complicate an analytical fit, but pose no problem at all for a MCRW fit. An interesting variation is to bin the events using a transformed variable, such that all bins have an adequate number of events, even if there are long tails to the underlying distribution. This would also be complicated for an analytical fit but trivial for an MCRW fit, provided both data and MC are binned in the same transformed variable.

Another approach to the low-population MC bin problem is to intentionally use different parameters when generating the template than are expected for the fit. For instance, use a wider gaussian width or longer exponential lifetime, to increase the template population in the tails. Inside the fitter, the mean w_i value in the tails will then be much less than 1, but be calculated from a substantial number of MC events. In fact, the formalism would allow the MC template to be generated with a completely flat distribution, with the entire job of making the MC look like the data being done by the F function.

2.4 Parameter Errors

The parameter errors are the range of parameter values where the $-\log(\mathcal{L})$ is within 1/2 unit of the value at the minimum. In most cases $-\log(\mathcal{L})$ is well approximated by a paraboloid near the minimum:

$$-\log(\mathcal{L}) = -\log(\mathcal{L})_{min} + \frac{1}{2} \sum_{j,k} \Delta p_j W_{jk} \Delta p_k \quad (2.26)$$

In this case, the parameter covariance matrix is the inverse of the W matrix, and the parameter errors are the square roots of the diagonal elements of the inverse. These are equivalent to the “parabolic” errors from MINUIT. In most cases these errors are sufficient, but in some cases the second-derivative terms that were dropped in our definition of the W matrix, or nonlinearities in the underlying problem, make the parabolic errors inaccurate. In such cases it is recommended to explicitly evaluate $-\log(\mathcal{L})$ to find the range where it is within 1/2 of the minimum.

The parabolic errors in Equation 2.26 reflect not just the statistical error of the data, but also of the Monte Carlo simulation. This is because the likelihood includes not just the Poisson probability of fluctuations of the data bins, but also the Poisson probability of fluctuations of the MC simulation bins, and the Gaussian fluctuations of the MC mean event weights. If the MC statistics are equal to the data statistics, and the MC was generated with close to the right parameter values, the parabolic errors are roughly $\sqrt{2}$ times what they would be when fitting the data with an analytic model. If the MC statistics are much larger than the data statistics, the parabolic errors are about equal to those of an analytic model fit.

In fact, if the MC statistics become infinite, then the MC errors become negligible, and the fit is precisely equivalent to a fit using a binned analytic probability density function. It will have exactly the same errors, and same bias or lack of bias.

2.5 Goodness of Fit

In a χ^2 fit, perfect agreement between the data and the model gives $\chi^2 = 0$, and the goodness of fit can be determined by the χ^2 value and the number of degrees of freedom. The χ^2 formalism can be derived from the maximum likelihood formalism, and the relation between them is $\chi^2 = -2 \log(\mathcal{L})$.

It is possible to add parameter-independent (but data-dependent) constants to our log likelihood such that its value is zero when the data and model agree perfectly. We also include the factor of -2 to agree with the normalization of χ^2 .

$$-2 \log(\mathcal{L}) = -2 \sum_{i=1}^N \left[n_i \log(x_i) - x_i + m_i \log(y_i) - y_i - \frac{(w_i - v_i)^2}{2\sigma_i^2} - n_i \log(n_i) + n_i + m_i \log(m_i) - m_i \right] \quad (2.27)$$

It is observed that the numerical value of this modified $-2\log(\mathcal{L})$ is very similar to the value of χ^2 [7] defined as:

$$\chi^2 = \sum_{i=1}^N \frac{(n_i - m_i w_i q)^2}{n_i + m_i w_i^2 q^2 + m_i^2 \sigma_i^2 q^2} \quad (2.28)$$

The numerator is the square of the difference between the observed data events and the prediction from the MC events, mean event weight, and MC normalization factor q . The denominator is the sum in quadrature of the $\sqrt{n_i}$ data error, the $\sqrt{m_i}$ MC error scaled by the weight and normalization, and the error of the mean MC event weight scaled by the MC events and normalization.

The number of degrees of freedom for the χ^2 expression is the number of data bins N minus K , the number of fit parameters (which include q), or $N - K$. The number of degrees of freedom for the $-2\log(\mathcal{L})$ expression is also $N - K$, but the accounting is different. There are not N measurements but $3N$, because we count m_i and w_i as measurements as well as n_i . In addition to fitting the K parameters, we calculate the N values of v_i and the N values of y_i to minimize $-2\log(\mathcal{L})$, so there are $2N + K$ fit parameters. Thus the degrees of freedom for a MCRW fit are $3N - (2N + K) = N - K$.

One could do a fit by minimizing the above χ^2 rather than $-2\log(\mathcal{L})$. This would still include both the bin-contents and mean-weight errors of the MC simulation as well as the bin-contents errors of the data. It would also avoid the need to numerically solve for v_i and y_i (although this is a relatively easy and fast task). The w_i and σ_i are still required, and the number of F and/or $dF/d\vec{p}$ evaluations required would be the same for a given minimization algorithm. We have found that the χ^2 formulation is somewhat better behaved when the starting point is far from the minimum, but essentially the same for a starting point near the minimum. The parameter values at the minimum $-2\log(\mathcal{L})$ and χ^2 are not identical, but that is not surprising, and they are much closer together than the parabolic errors.

χ^2 fits tend to be somewhat biased. When a data bin fluctuates upward its effective error grows and when it fluctuates downward its effective error shrinks, so the average effect of fluctuations is downward. Poisson likelihood fits do not have this problem because the effective error for a bin is determined not from the fluctuating data but from the prediction for the bin. In a MCRW fit based on χ^2 , there are biases not only from the data fluctuations, but also from the MC fluctuations. Amusingly, the MC fluctuations are included with the opposite sign, and if the MC has the same number of events as the data, the biases roughly cancel. If, however, the MC has many more events than the data, the cancellation is less perfect and the fit is typically biased.

2.6 Example Application of MCRW Method

This section describes the application of the MCRW method to the sample problem of an exponential decay smeared with a Gaussian resolution, and a

Parameter	ROOT	MCRW V1	MCRW V2
χ^2/DOF	105.4/117	114.9/117	93.4/117
τ Fit	0.9984	0.9991	0.9993
τ Error	$1.137 \cdot 10^{-3}$	$1.247 \cdot 10^{-3}$	$1.236 \cdot 10^{-3}$
τ Resid.	$-1.59 \cdot 10^{-3}$	$-8.93 \cdot 10^{-4}$	$7.19 \cdot 10^{-4}$
σ Fit	0.2496	0.2492	0.2495
σ Error	$5.09 \cdot 10^{-4}$	$5.54 \cdot 10^{-4}$	$5.45 \cdot 10^{-4}$
σ Resid.	$-3.92 \cdot 10^{-4}$	$-7.67 \cdot 10^{-4}$	$-2.50 \cdot 10^{-4}$

Table 2.1: Comparison of example fit results using standard analytical likelihood (ROOT), and the MCRW method, with template parameters equal to the data parameters (V1) and 50% larger than the data parameters (V2) for the example problem in Section 2.6. We give extra digits in the errors in order to compare them more accurately.

comparison of the results obtained via the MCRW method and the more traditional method using MINUIT [6] within the ROOT [5] software framework.

We generate 1 million simulated decays, with a lifetime, τ , of 1.0, and note the true decay time T . We smear the true decay times with a Gaussian resolution model, by generating a value δ , which is Gaussian-distributed with mean 0.0 and width $\sigma=0.25$, and adding it to the true decay time T , to obtain the observed measurement $x = T + \delta$.

In order to perform a traditional fit using the MINUIT minimizer in the ROOT software package, we first write down the analytical convolution of an exponential decay with a Gaussian resolution:

$$P(x) = A \cdot \frac{e^{-x/\tau}}{2\tau} \cdot e^{\sigma^2/2\tau^2} \cdot \operatorname{erfc} \left(\frac{1}{\sqrt{2}} \left(\frac{\sigma}{\tau} - \frac{x}{\sigma} \right) \right), \quad (2.29)$$

in which A is an amplitude whose value depends on both the number of events and the binning chosen, and erfc is the complementary error function. This is one of the few realistic convolutions for which it is possible to easily obtain an analytic solution, although the erfc function must be evaluated numerically.

We fill a histogram, consisting of 120 bins in the range $x=-1..5$, with the 1 million observed decay times generated above. This allows for potential negative measured values, to demonstrate the smearing effect more clearly, and removes the very far positive tail of the distribution which has sparse statistics. We perform a traditional fit using the PDF given in Equation 2.29 in the ROOT software package, employing the Likelihood formulation, and using the integral of the function across the bin, instead of the central value, to improve the accuracy of the fit. We obtain the results shown in Figure 2.1 and Table 2.1.

In order to perform a similar fit with the MCRW method, we first generate two sets of MC Template events. The first set has the same parameter values as the data above, with a size of 5 million decays. The second template is also generated with 5 million decays, but uses parameter values 50% larger than

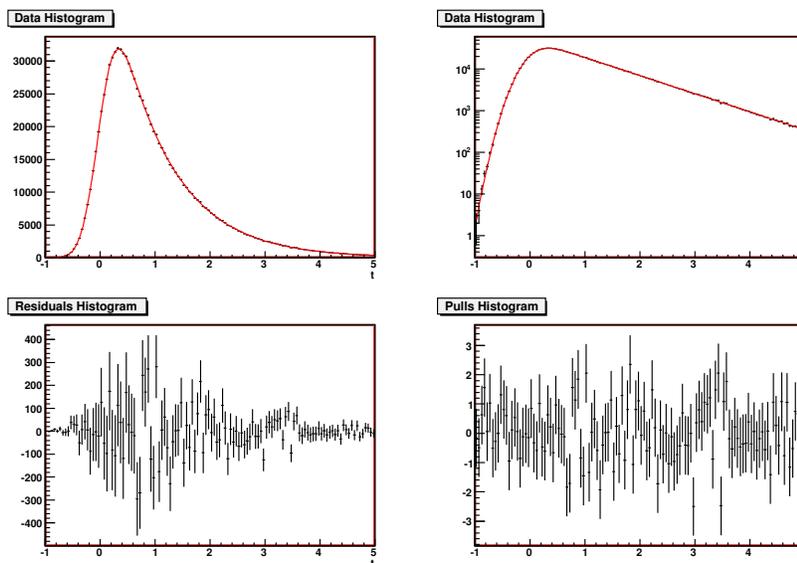


Figure 2.1: ROOT Likelihood Fit to the smeared-exponential problem. Data is histogram (black), analytic fit is smooth curve (red). Upper left plot is linear scale, upper right plot is logarithmic scale. Lower left plot is residuals (Data-Function), and lower right plot is pulls (Residuals/Bin Error).

those of the data, specifically $\tau=1.5$, and $\sigma=0.375$. The parameter vector for this problem is: $\vec{p} = (\tau, \sigma, q)$. The truth information stored for each event is $\vec{t} = (T, \delta)$, where T is the true, not observed decay time of the event, and δ is the difference between the true and observed decay times. We fill template histograms (which have the same number of bins, and the same limits as the data histogram) with the MC Template decays, using the same observed value $x = T + \delta$, that was used to fill the data histogram.

We use Equation 2.30 as the weighting function:

$$F(\vec{p}, \vec{t}) = \frac{e^{-T/\tau}}{\tau} \cdot \frac{e^{-\delta^2/(2\sigma^2)}}{\sqrt{2\pi}\sigma} \quad (2.30)$$

and obtain the results shown in Figure 2.2 and Table 2.1. All fits were performed on the same computer, a 2010 model desktop computer. The ROOT fit required 273 function calls, and took approximately 2 seconds to read the data from disk, and less than 2 seconds to perform the analytical fit. The MCRW fit needed 10 seconds in total to read both the data and the larger template from disk, and only used 5 iterations to converge with either template. The actual fit calculations required took 11s. While the MCRW fits for this problem were slower than analytical fits, the speed was still acceptable. Note that if we had

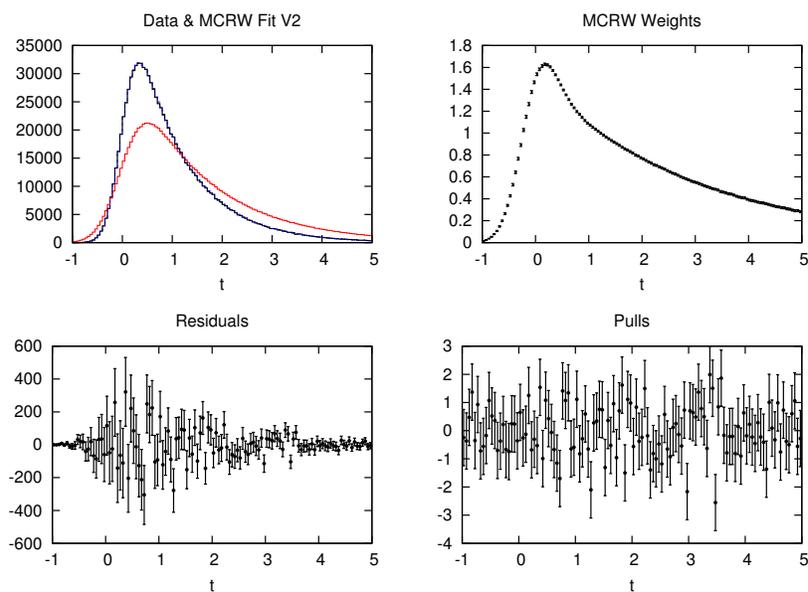


Figure 2.2: MCRW Fit to the smeared-exponential problem, using the template with longer lifetime (V2). Upper left, data histogram is black, scaled template (mq) histogram is red, weighted template (mwq) is blue. Upper right, average weight w is shown with error bars multiplied by 5. Lower left, residuals are Data-Fit. Lower right, pulls are Residuals/Errors.

used 1M template events (equal to number of data events) instead of 5M, the computation time would have been reduced to 2.2 seconds, although at a cost in statistical error.

Upon examination of the values in Table 2.1, we see no significant difference in the fitted values of the parameters. We do observe the expected difference in the size of the parameter errors, due to the contribution from the fluctuations in the MC template histogram and the MC template event weights. We expect the parameter errors in the MCRW fit to be equal to the conventional fit errors times $\sqrt{1 + 1/5} = 1.095$ when the template parameters are equal to the data parameters (V1 in the table) since there are 1M data events and 5M template events. The ratio of the errors agrees with the expected value to within a few parts per thousand. The fit errors are slightly smaller in the case when the template parameters are 50% larger than the data values (V2 in the table). This is likely because in this case the MC template has more events in the far negative and positive tails, which strongly influence the determination of σ and τ respectively.

To summarize these results, the well-established ROOT/MINUIT software, and our new method give perfectly compatible results when both are applied to an identical, tractable problem, once we account for the expected inflation of the parameter errors in the MCRW method.

We have conducted extensive tests of the fitting software, beginning with simple toy problems like the one presented here, and working our way slowly up to the most detailed example examined so far, the model of neutral B meson oscillations examined in subsequent chapters. In this case, we are able to obtain fits with the MCRW method which show no evidence of bias, where other researchers have been unable to devise analytical likelihoods that are precise enough to avoid significant biases. We expect to publish a journal article going into more details about the fit algorithm and its performance at a later date. We also plan to distribute the MCRW fitting software for wider use in the community of experimental particle physics.

3

Dilepton B Mixing Analysis Issues

In this Chapter we present an overview of the challenges faced in performing a time-dependent measurement of the neutral B meson mixing frequency in the “dilepton” sample. We discuss the detailed physics of the “Boost Approximation”, the sources of background, and the detector effects. We then compare the traditional method of performing this analysis with an analytic likelihood fit to the MCRW method of performing the analysis.

3.1 Boost Approximation

The experimental goal that has motivated the development of the analysis tools described in this thesis is the measurement of the neutral B meson oscillation frequency. A time-dependent measurement can be performed by measuring the flight distance of the B mesons if the relativistic velocity, or boost ($\beta\gamma$) is known for them. The cleanest environment for performing these measurements is to produce $B^0\bar{B}^0$ pairs in e^+e^- collisions tuned to the energy of the $\Upsilon(4S)$ resonance. The decay of the $\Upsilon(4S)$ produces pairs of B mesons which are nearly at rest with respect to the centre-of-mass. If the incoming beam energies are unequal, this sets sufficiently large values of the boost velocity, $\beta\gamma$, for the B mesons to allow the experimental measurement of their decay positions to be used to determine their decay times.

In these e^+e^- collisions the production point of the $\Upsilon(4S)$ is not known with sufficient precision (because the beam spot, or region in which the e^+e^- collision can occur, is relatively long), but we can use the difference of the two B meson decay points to measure the time difference between their decays. The large boost along the z axis leads one to the approximation that all the B meson flight is along the z axis. The spatial difference between the two B meson decay points along the z axis, Δz , can be converted into the decay time difference Δt , through the so-called “Boost Approximation”:

$$\Delta t \simeq \Delta z / (\beta\gamma c) \tag{3.1}$$

Tests based in generator-level Monte-Carlo showed that the boost approximation did not measurably bias the measurement of Δm in our method of MCRW fitting, compared to using the “true” B meson boosts for each of the two leptons.

The sign of Δt is determined by the lepton momenta, with the decay time of the higher momentum first, minus that of the second, lower momentum lepton. This expression is symmetric for the case where both leptons are produced directly from B decays.

The B meson decay point is estimated from the point-of-closest-approach (POCA) of the lepton track to the beam axis. This introduces uncertainty in the measurement of z for each lepton, and the difference Δz . The B mesons are also not exactly at rest in the $\Upsilon(4S)$ frame; they have a momentum of roughly 340 MeV/c in the center-of-mass frame of the $\Upsilon(4S)$. This means that the “Boost Approximation” has replaced the more precise calculation:

$$\Delta z = z_2^{lab} - z_1^{lab} = \gamma\gamma^{cm}\beta c(t_2 - t_1) + \gamma\gamma^{cm}\beta^{cm}c\cos(\theta^{cm})(t_2 + t_1), \quad (3.2)$$

where t_i is the time when each B decays in its own rest frame, $\gamma^{cm} \simeq 1.0022$ and $\beta^{cm} \simeq 0.065$ is the boost of the B meson in the $\Upsilon(4S)$ rest frame, γ and β define the boost of the $\Upsilon(4S)$ in the lab frame (in which the distances are measured), and θ^{cm} is the polar angle between the two decays in the $\Upsilon(4S)$ rest frame. The difference between Δz and $\beta\gamma c\Delta t$ can cause biases of order one percent in time-dependent measurements [8]. The POCAs of the leptons are also smeared by the transverse flight of the B mesons, in a non-trivial way.

While software tools have been designed that can fit the distribution in Figure 1.2 smeared by a resolution function represented by a sum of Gaussians (e.g. RooFit [9]), the boost smearings are not exactly Gaussian, and are time-dependent, so they can only be treated approximately. In addition, the analytic forms of these convoluted functions take a long time to calculate with adequate accuracy for a minimizer, like MINUIT [6] that works with small differences.

3.2 Background Processes

Within the dilepton sample there will be both signal and background leptons. Signal leptons are from neutral B meson decays, and specifically the case when each lepton has come directly from a B meson decay. Each pair of signal leptons is vertexed with the beam axis to obtain a measurement of Δz , the distance along the z axis between the origin points. This distance is converted into the decay-time difference Δt using the boost approximation, Equation 3.1. The signal lepton pairs are also divided into opposite-sign (OS or unmixed) and same-sign (SS or mixed) pairs.

Background leptons can originate from secondary or cascade decays of neutral B s, charged B events and from non $B\bar{B}$ continuum events. We can divide the background leptons into the following categories:

- B^+B^- direct-direct which has a different lifetime from the $B^0\bar{B}^0$ pairs. The ratio of charged to neutral B s is not known precisely, but there are no SS pairs in this category.
- Secondary, or cascade (e.g. $B \rightarrow D \rightarrow \ell$) leptons with (mostly) the “wrong” sign. There are multiple components here, from the same B

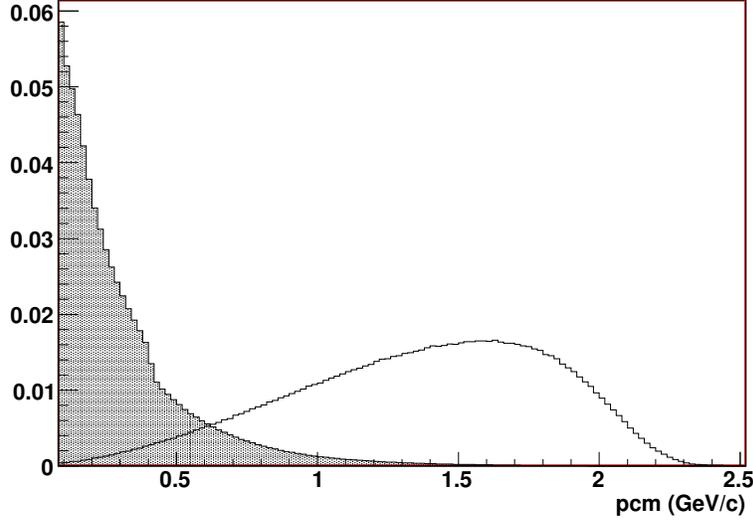


Figure 3.1: Spectra, normalized to unit area, of MC truth electrons and muons, divided into direct from B decays (open histogram), and all other sources (shaded histogram) in fully simulated generic B decays.

and the other B , with extra flight and different lifetimes compared to the direct leptons. They generally have lower momenta than the direct leptons. Their relative amount and composition are not known precisely.

- Tertiary leptons are real leptons from other sources in $B\bar{B}$ events besides those mentioned above, including hadrons appearing as fake leptons.
- Continuum leptons are from Non- $B\bar{B}$ events produced in the e^+e^- collisions, essentially divided into $u\bar{d}sc$ quark-antiquark pairs, τ pairs, and so-called QED e^+e^- and $\mu^+\mu^-$ pairs.

If we examine the centre-of-mass momentum of simulated electrons and muons, we can see a clear distinction between the signal leptons, which are direct from B decays, and all other sources of leptons, as shown in Figure 3.1. The p_{cm} variable is one of the most powerful discriminating tools available for distinguishing between signal and background leptons in the dilepton analysis of B mixing, in either the traditional fit or MCRW fit methods.

All of the above background components have the same analysis procedure applied to them to calculate Δz and Δt , and contribute in various amounts and shapes to the observed Δt distribution.

In Figure 3.2 and 3.3 we can see the shapes in Δt of various contributions to the dilepton sample, in both the OS and SS pairs. The signal or direct-

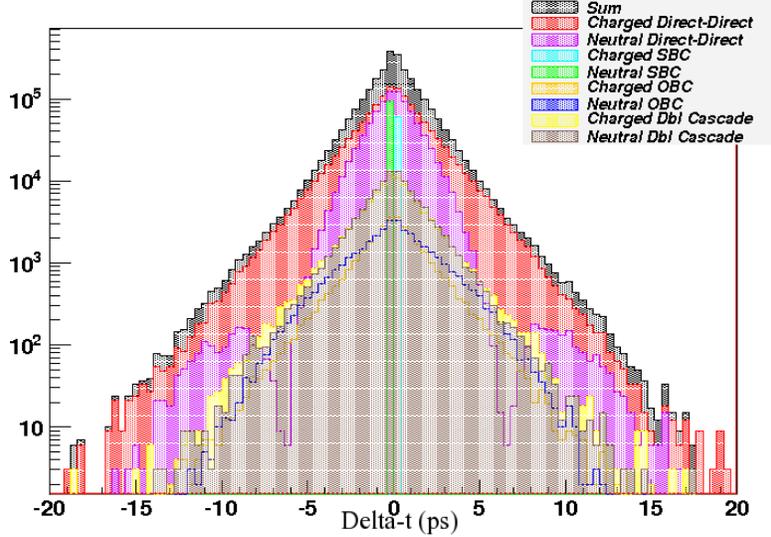


Figure 3.2: Overlay of true Δt for the following categories: Direct-Direct $B^0\bar{B}^0$, Direct-Direct B^+B^- , Direct-Cascade $B^0\bar{B}^0$, etc. OS pairs only. Both leptons satisfy $p_{cm} > 0.5$ GeV/c, and $0.5 < \theta < 2.5$.

direct pairs are symmetric in Δt , but the background pairs are biased towards positive Δt when one lepton is from a direct B decay, and the other lepton is from a charm decay. The lepton from direct B decay will most often have higher momentum than that of the charm decay, so we tend to have direct-charm more often than charm-direct. The sign of the lepton from the charm decay is normally opposite that of the B decay, which leads to an apparent increase in the number of same sign (SS) pairs.

All of these processes will be present in some amount in the final fit to data, and need to be modelled accurately in order to determine Δm precisely. Applying cuts to the data can help to reduce these contributions from background leptons (especially the p_{cm} cut, as mentioned above), but they cannot be removed completely. A dilepton analysis of neutral B mixing needs to be able to deal with a large number of background contributions in the final fit, in a precise manner. The Δt distributions in Figure 3.2 and 3.3 were obtained from generator-level Monte-Carlo dilepton generic B events, after imposing the requirements that both leptons have $p_{cm} > 0.5$ GeV/c, and $0.5 < \theta < 2.5$.

The measured lepton pairs are also affected by the usual experimental effects of detector efficiency, event selection efficiency, track reconstruction efficiency, track selection efficiency (we require a large number of hits in the tracking system), Particle-Identification (PID) efficiency, and tracking resolution. The efficiency for detecting the lepton pairs and the tracking resolution are often

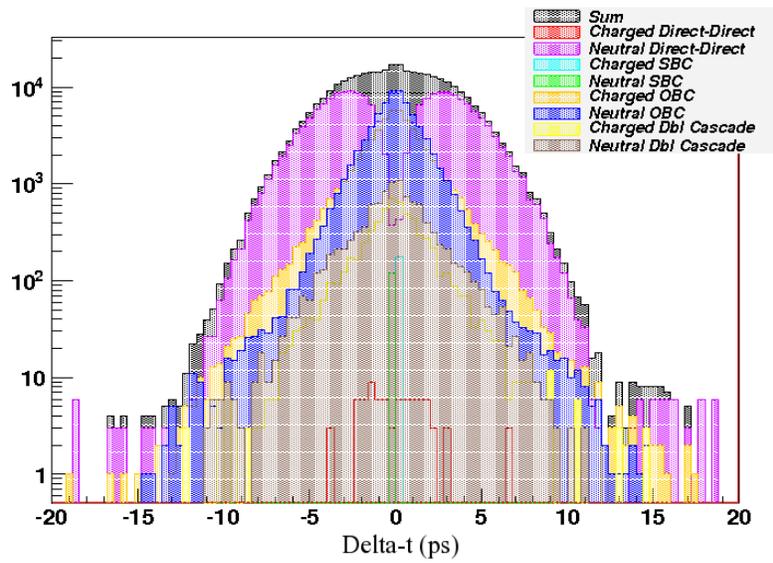


Figure 3.3: Overlay of true Δt for the following categories: Direct-Direct $B^0\bar{B}^0$, Direct-Direct B^+B^- , Direct-Cascade $B^0\bar{B}^0$, etc. SS pairs only. Both leptons satisfy $p_{cm} > 0.5$ GeV/c, and $0.5 < \theta < 2.5$.

measured in detailed Monte-Carlo simulations. They can be measured in data control samples, but these control samples are often only used as cross-checks of the simulation, because they generally have little similarity to the dilepton B events. Both the various efficiencies and the tracking resolution must be included in the model that is used to fit the data, when trying to extract Δm from a fit to the Δt distribution for OS and SS dilepton pairs.

3.3 Monte-Carlo Models

The *BABAR* Collaboration developed and utilizes the EvtGen [12] software package to generate Monte-Carlo decays used in studying B physics. This piece of software simulates the creation of $\Upsilon(4S)$ particles with the appropriate spread in energies and momenta, to match the operating conditions of the PEP-II collider. It then decays the $\Upsilon(4S)$ into $B^0\bar{B}^0$ and B^+B^- pairs, with equal probability. These B mesons decay with lifetimes and branching fractions specified in electronic data tables, up to the creation of particles with lifetimes $c\tau > 2$ cm, at which point the decays stop and the particles are declared stable. Particle decays which take place beyond 2 cm in radius are strongly affected by the detector materials, and must be dealt with in the fully detailed simulation, not the event generator. EvtGen can be configured to use the data tables assembled by the Particle Data Group (PDG) [13], which contain the most precise, up-to-date world averages of the properties of the B meson, and its decays. These include the probabilities for decays to various excited states, as well as ground states, and both inclusive and exclusive decay modes. In addition, there are a variety of decay models employed in EvtGen, which allow the program to simulate different form-factors in decays, angular distributions, and time-dependent CP-violation and mixing in B and charm decays. The EvtGen sample fully simulates the beam and B meson physics, but includes no model of the detector whatsoever.

There is another software package used within the *BABAR* Collaboration, referred to as **Moose**, which is an acronym of Monolithic-Simulation-Executable. This simulation is only described in *BABAR* Collaboration internal documents.

The software starts by running the EvtGen event generator to simulate B meson decays, but then passes the particles created with EvtGen into a detailed physical simulation of the detector built inside the GEANT4 [14] software framework. This allows the particles to interact with the various materials, both active and inactive, of the detector. This simulation contains an extremely detailed model of the geometry and materials of the *BABAR* detector, and accurately simulates the energy deposition, ionization, scattering and other material interactions of the particles with the detector. Finally, all of the information from the first two stages is passed into a third stage which runs the same reconstruction algorithms that are applied to real *BABAR* data on the simulated detector hits created by the previous stages.

An attempt is then made to “truth-match” the reconstructed particles with the true MC particles which led to their detection, a process which is quite

difficult and not always perfect. During simulation, the “hits” made by each true track are known, and the “hits” contributing to each “digitization” are also known, but a digitization may come from more than one track. A track is reconstructed from many digitizations. For simulated events, a list of all true tracks that made any contribution to a reconstructed (reco) track is computed. The single true track that made the largest contribution of hits to the “reco” track is considered to “match” the reconstructed track. For tracks with a reasonably large number of hits (certainly true for the leptons we use), this reconstructed-to-MC Truth matching is a quite unambiguous process. In some angle and momentum ranges, true tracks produce so few hits that they are not efficiently reconstructed. However, some of their hits may still be incorporated into (poorly) reconstructed tracks. One can still ask the question of which reconstructed track incorporates the most hits from a given true track, but this is much less useful than the reverse question.

At the end of the simulation process, the momentum vector and origin point of each true track in the decay chain is recorded, along with the truth-matching information. As well, all of the reconstructed candidates and their properties are recorded, just as they are when real experimental data is analyzed.

Throughout this Thesis we will refer at various points to either EvtGen or generator-level MC as events which have only been simulated at the level of the EvtGen program, not including any detector interactions or reconstruction. We will refer to fully-simulated MC events as those which have passed through the entire process contained within the Moose software program.

3.4 Traditional Fitting Approach

Most time-dependent B decay analyses have used an analytical PDF (Probability Density Function) constructed within the ROOT [5] software framework to perform a maximum likelihood fit to the measured data. Often, some of the parameters in these rather complicated PDFs are highly correlated, and cannot all be allowed to float at the same time in the fit. Fixed parameters in the fits are measured using either data or MC control samples, and they contribute to the systematic uncertainties in the measurement. We describe here the rather complicated method used to perform an analytical likelihood measurement of the mixing frequency in the dilepton sample of data from *BABAR* [2].

The likelihood in this analysis was composed of a sum of terms representing the contributions of $B^0\bar{B}^0$ lepton pairs, B^+B^- lepton pairs, and continuum lepton pairs, with a fraction parameter governing the total number of pairs in each of these three categories.

The PDFs for the $B^0\bar{B}^0$ and B^+B^- categories are very similar to each other, and they are composed of six contributions each, where the six components, roughly in order of their importance, are:

- Signal mode - Both leptons are directly from B meson decay, or primary-primary.

- Other B cascade (OBC) - Leptons are from opposite B mesons, one is direct or primary, and the other is cascade, or secondary.
- One direct lepton and one lepton from a τ decay.
- Same B cascade (SBC) - Both leptons are from the same B meson, one is direct, or primary, and the other is cascade, or secondary.
- Lepton pairs in which both leptons originate from the same decay (J/ψ or other vector mesons), with zero true decay-time difference.
- A grab bag of a small amount of other lepton pairs.

The PDF for the signal component contains an exponential lifetime and a sinusoidal mixing term convoluted with a signal resolution function that is the sum of three Gaussians. There is also a nearly equal contribution from charged B pairs with a different lifetime and no mixing, convoluted with the same sum of Gaussians.

The OBC and τ components have one lepton direct from one B and the other lepton from a finite-lifetime decay product of the other B , so the PDF has the B lifetime (with mixing for B^0 , but no mixing for B^+), and also a convolution with an additional exponential lifetime. Usually the lower momentum lepton is from the charm or τ decay, but in a fraction of the cases it is the higher momentum lepton, which flips the direction of the convolution. There are different effective lifetimes for τ , D^+ , and for D^0 and D_s (treated together), and tracking resolution sums of Gaussians (different from the signal because the tracking resolution is different since the cascade lepton momentum is lower).

The SBC component has one lepton direct from one B and the other lepton from a finite-lifetime decay product of the same B . The PDF has no B lifetime and no mixing terms, but different effective lifetimes for its D^+ and D^0/D_s components (with different relative sizes for B^0 and B^+), with a fraction flipped in direction, and yet more different tracking resolution sums of Gaussians.

The J/ψ component has both leptons from the same particle decay, so there is no lifetime or mixing term, just a sum of Gaussians for tracking resolution.

The final component contains every dilepton in the MC simulation not contained in one of the other categories. It is dominated by one lepton directly from B decay and one “lepton” that is actually a misidentified hadron or real lepton from some other process, from either B . These are modelled by an effective lifetime (different for B^+ and B^0), convoluted by another sum of Gaussians.

To add a final additional layer of complexity, the entire PDF structure described above was fitted separately to each of the 4 dilepton types: ee , $e\mu$, μe and $\mu\mu$. The fits were done separately because the PID efficiency as a function of momentum is different for electrons and muons, and therefore the effective tracking resolution is expected to be different as well. Again, it should be noted that there is a distinction between $e\mu$ and μe due to the momentum-ordering of the pair of leptons.

The last category to add to this model was that due to continuum lepton pairs. It was also relatively simple in comparison to the model for B events given

above. A single effective lifetime was used, and the signal resolution function was used. The continuum fraction was determined by counting the number of events in their selection for on-resonance and off-resonance data, and multiplying their ratio by the ratio of the total recorded luminosities of on-resonance and off-resonance data. In the final fit this continuum fraction was fixed.

The effective lifetimes, flip fractions, and Gaussian resolution parameters in the PDFs were determined by fitting Δt in reconstructed MC simulation for the various categories, selected using MC truth information. The flip fractions represent the portion of events in which the higher momentum lepton is directly from a B meson (unflipped) or in which it is from a D meson or other cascade decay (flipped). The relative amounts and “dilutions” (deviations from the expected same-sign vs opposite-sign behaviour) of the components were determined by MC event counting.

The mixing fit was first performed independently on the four different dilepton types, and the parameters which were allowed to float in each of these fits were the following:

- Δm , the mixing frequency,
- The width of the core Gaussian in the triple Gaussian signal resolution,
- The width of the second Gaussian in the triple Gaussian signal resolution,
- The fraction of the second Gaussian in the triple Gaussian signal resolution,
- The fraction of the third (outlier) Gaussian in the triple Gaussian signal resolution,
- The overall fraction of B^+B^- events,
- The fraction of OBC in the $B^0\bar{B}^0$ events. Note that the OBC fraction in B^+B^- events is fixed by a ratio to this fraction.
- The effective lifetime of continuum events.

When the fit was extended to fit all four dilepton types simultaneously, all the parameters in the above list were allowed to be different for each of the 4 categories, except for the mixing frequency Δm . In this case the high-degree of correlation between the parameters meant that some of the above parameters needed to be fixed and the 4 values for the width of the second Gaussian in the signal resolution function were fixed to those obtained in the individual fits to the 4 dilepton categories. In total there were 25 floating parameters in the final fit to the data, while the total PDF contained 520 numerical parameters describing 88 different shapes.

The authors of the analysis estimated the systematic uncertainties associated with their model in the standard way, by varying the values of the fixed parameters up and down by a “reasonable” amount, and observing the change

Bias on Δm	0.0062
z scale and SVT misalignment	0.0039
Signal resolution	0.0036
B lifetimes	0.0033
<i>Xtra</i> Parameterization	0.0025
Other B cascade (OBC) parameterization	0.0022
Total	0.0087

Table 3.1: Major systematic uncertainties, in ps^{-1} , in the unpublished *BABAR* result on mixing with dileptons [2].

in the fitted value of Δm . As in most B mixing measurements, they were unable to float both the B^0 and B^+ lifetimes simultaneously, and fixing the value of τ^+ to the world average value creates a noticeable systematic uncertainty. Table 3.1 contains a list of the dominant sources of systematic uncertainty (in ps^{-1}), in the above unpublished result. The dominant systematic, by far, is the “Bias on Δm ”, which was attributed primarily to “time-dependent resolution”.

The “Boost Approximation”, whereby the proper motion of the B mesons with respect to the beam spot is normally neglected, ignores different effective boost values ($\beta\gamma$) for each lepton, depending on the flight direction of the parent B meson. If the B meson is travelling along the z axis, then the lepton will have a larger boost in that direction, compared to when the B meson travels perpendicular to the z axis. The apparent decay time difference Δt will change, due to the difference in the boost component along the z axis. This can create smearing on the order of 1% of the central value [8], which is noticeable with our level of statistics. Furthermore, this smearing does not really follow a Gaussian distribution, and is therefore difficult to apply with an analytical convolution.

The authors of this past analysis attempted to use simple cuts to reduce the effect of the boost smearing on their data, mainly by using more restrictive cuts in the lab polar angle, since leptons which are travelling transverse to the beam should have less longitudinal boost smearing. Figure 3.4 shows the biases observed in fits to approximately 300 thousand fully simulated Monte-Carlo events, as a function of the cuts used on θ , the lepton lab-frame polar angle. Notice that as the cuts were tightened, the bias shrank very slowly, and the statistical error on the measurement grew, until the fit was within 1σ of the expected value. Note also that the larger samples fully include the smaller ones, so the samples are not completely statistically independent. This detail aside, it is clear that these cuts don’t remove the bias, and greatly lower the statistics of the sample as the cuts are tightened. In the end, they chose $1.17 < \theta < 1.97$, corrected their measurement for this bias, and quoted half of the size of the bias as a systematic uncertainty.

In the end, after attempting to minimize the size of the bias while keeping

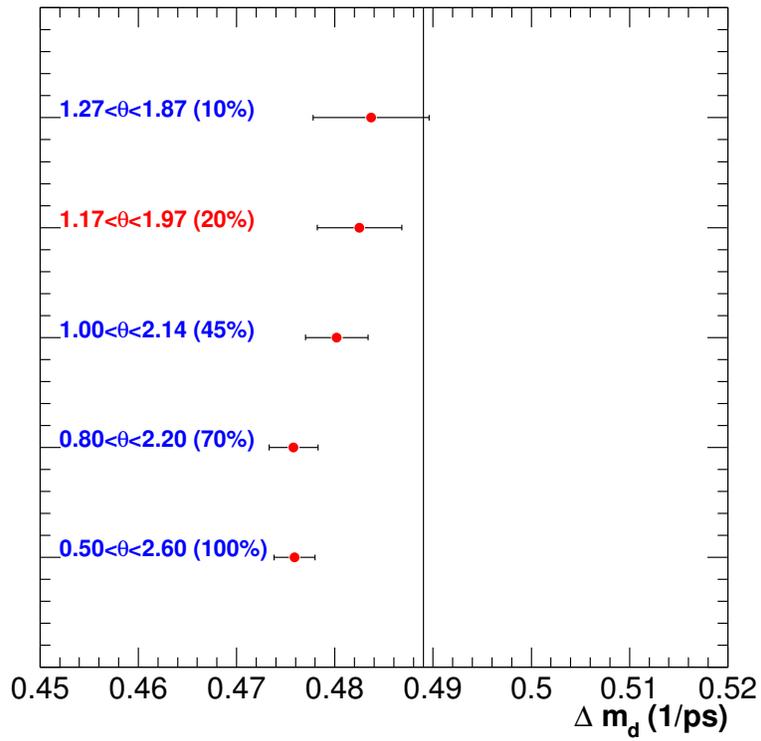


Figure 3.4: Fitted values of Δm in MC, as a function of various lab θ cuts. The vertical line represents the MC generator value used for Δm_d . Taken from [2].

as many events as possible, they obtained the (still blinded) result:

$$\Delta m = x.xxx \pm 0.009 \text{ (Stat.)} \pm 0.010 \text{ (Syst.) ps}^{-1} \quad (3.3)$$

with statistical and systematic uncertainties of roughly 2% each. This was based on the data in *BABAR* Runs 1–3. In comparison to the world average value:

$$\Delta m = 0.505 \pm 0.005 \text{ ps}^{-1} \quad (3.4)$$

which has a combined error of only 1%, we can see that this measurement was not able to improve our knowledge of Δm , and this led to the decision not to publish it.

We will see that in the MCRW approach we also rely heavily on the Monte-Carlo simulations, but in a more direct way, and without having to determine as many numerical parameters in fits to the Monte-Carlo. Our PDFs are generally simple and easy to implement, and the fit bias and associated systematic uncertainty that created so much trouble in the traditional analysis is cleanly dealt with in the MCRW approach.

3.5 Overview of MCRW Approach

For our analysis we do not attempt to use a highly complex analytical PDF to describe the distribution of Δt ; instead we use a large sample of simulated events to create what we call our Monte-Carlo (MC) Template. Our fit then maximizes the likelihood that the data and template histograms arise from the same underlying parameters, within statistical fluctuations. This general fitting technique was discussed in more detail in Chapter 2.

The template is composed of several parts, in order to allow us to have floating parameters in the fit corresponding to our major sources of uncertainty. The first main component is made up of generator-level events, produced with the EvtGen [12] software package, described earlier in this Chapter. Since the software is relatively simple compared to the full simulation involved in simulating the detector behaviour and the reconstruction procedure, we can generate as many of these events as we desire to use in our template, without being limited by available computing resources. This reduces the contribution to the overall measurement error which is due to finite statistics in the MC template.

We only include in the template true leptons which originate from the following sources in the EvtGen template: direct B to lepton decays, B to charm meson to lepton “cascade” decays, B to tau to lepton decays, and B to charm baryon to lepton decays. We refer to these as primary and secondary leptons.

It is important to note that: the EvtGen software includes a precise simulation of the effects of longitudinal and transverse boost smearing, which are ignored in the traditional fitting method. This key difference allows us to eliminate the large bias and associated systematic uncertainty that posed so much trouble for the dilepton analysis performed in the traditional way. We can successfully perform a MCRW fit to the EvtGen events that takes into account the

B^0 and B^+ lifetimes, the mixing frequency Δm , and the ratio of primary to secondary leptons, without any bias on the measured parameters.

We then apply our own, parameterized detector model to the EvtGen events, consisting of the following stages: first is bremsstrahlung emission by the electrons, followed by a track-scattering and momentum vector smearing model to simulate the processes of multiple-coulomb scattering and the geometric resolution of the detector in measuring charged particle trajectories. By modelling the track parameter resolution ourselves, we are able to introduce floating parameters into the fit to describe the tracking resolution and vary them because we know the true track parameters, as well as the amount by which they were smeared. As described in Chapter 4 we are able to measure these resolution parameters in a MCRW fit to the single track resolution in fully simulated events, and apply them to create our template. We are unable, however, to simultaneously fit all of the resolution parameters in a MCRW fit to the Δt distribution, because of the high degree of correlation among them.

The third step is an application of the reconstruction efficiency measured in fully-simulated B meson decays (using the *BABAR* MOOSE software package), or the efficiency to reconstruct and match to a MC-truth candidate a given charged particle. Finally another efficiency is applied representing the efficiency of particle-identification algorithms for a given true electron or muon, as determined in fully-simulated B meson decays. These tables are easy to measure using the fully-simulated events, and easy to apply to the generator-level events.

There are other sources of reconstructed leptons that will be present in either a fully-simulated sample of B meson decays, or in real *BABAR* data that are not possible to model with a template composed only of generator-level events. We have a second component of the template which is created directly from fully-simulated, reconstructed B meson decays. This includes all of the other reconstructed leptons in fully-simulated B meson decays which pass our data selection criteria, but don't originate from one of the specified primary or secondary decay paths which we use to create the EvtGen template. These "tertiary" leptons can be divided into the following categories:

- Large energy-loss bremsstrahlung electrons, electrons and positrons originating from photon conversions.
- "Other" electrons and muons, coming from the decays of relatively long-lived particles like pions, kaons and vector mesons.
- "Fake" electrons and muons coming from the mis-identification of pions, kaons and protons as electrons or muons.

These processes are difficult to model analytically, and are not critical in the measurement of the mixing frequency, so we control each of them with a single weight factor in our fits which describes how much of the process takes place in the data compared to our template. They depend heavily on the interaction of the leptons with the detector material and reconstruction software, and it is simpler to use the fully-detailed simulation to model them, rather than trying to

create our own ad-hoc model. We refer to these as “tertiary” leptons throughout this thesis.

Finally there is a third major component of the template which is composed of non $B\bar{B}$ events. There are many electron-positron collisions in the detector that produce other final states besides the $\Upsilon(4S)$. There are large cross-sections for the processes $e^+e^- \rightarrow q\bar{q}$, in which q is any of the quark species lighter than the b quark ($udsc$), and also for $e^+e^- \rightarrow \ell^+\ell^-$, where ℓ is any of three lepton families: e, μ, τ . There are roughly equal numbers of $q\bar{q}$ events, compared to the $B\bar{B}$ events, but there are many more e^+e^- and $\mu^+\mu^-$ pairs created. The *BABAR* trigger has a prescale factor intended to reduce the frequency of these QED events by a factor of 32 or 64 compared to the rate of $B\bar{B}$ events, but they still represent a large fraction of the recorded data. There are also so-called “two-photon” events in the data, which are not adequately simulated in the *BABAR* software, but they tend to produce mostly low p_T pairs, which are relatively easy to remove from the data selection.

In principle there are two ways to include these non $B\bar{B}$ events in the template. The first method is to use *BABAR* data recorded at Off-Peak energies, approximately 40 MeV below the $\Upsilon(4S)$ peak, which contains all the relevant processes except for $b\bar{b}$, in very nearly the same ratios as they occur at the peak energy (10.58 GeV). The difficulty with using the Off-Peak data in our template is that it was only recorded with one-tenth the statistics of the On-Peak data, and this will introduce statistical and systematic uncertainties into our fit when it is scaled up to match the luminosity of the On-Peak data. If the Off-Peak data is used in the template it is also not possible (or necessary when fitting data) to make any shape adjustments to the weights for the events, since they have no MC truth information stored.

The second method for including non $B\bar{B}$ events in the template is to use a cocktail of fully-simulated $udsc$ events, $ee, \mu\mu$ and $\tau\tau$ pairs, weighted according to their generated luminosities and their measured production cross-sections at the On-Peak energy. On the positive side, this approach allows us to have equally high statistics for the continuum events as we do for the simulated $B\bar{B}$ events in our template. It would also allow us to re-weight the events with shape parameters, if so desired, because of the availability of the MC truth information.

There are, however, several shortcomings with the use of this MC cocktail. Firstly, it is more complicated to compose the template and keep track of the relative fractions of the different categories, and the uncertainties of the measured production cross-sections for these processes become systematic uncertainties in our analysis. These cross-section measurements can have relatively large errors, on the order of 10% in several cases [15]. Additionally, the “two-photon” background events are not at all simulated adequately, and could not be included in this method. These are events produced through the process $e^+e^- \rightarrow e^+e^-\gamma\gamma \rightarrow e^+e^-X$, where the final state X is produced by the collision of two photons emitted by the incoming electron-positron pair, rather than their annihilation. In our case, we observe many events in the data in which the final state X is a pair of leptons, either e^+e^- or $\mu^+\mu^-$. The e^+e^- and $\mu^+\mu^-$

pairs were also not simulated in the same quantity and detail as the hadronic events, and would provide a poor model. We therefore choose to use the Off-Peak data for our template, not only because of the simplicity, but also because the contributions to the total error appear to be lower.

With these three major components in the template: $B\bar{B}$ primary and secondary leptons simulated with EvtGen, tertiary leptons from the full *BABAR* simulation and reconstruction, and continuum pairs from Off-Peak data, we have a template in place that can be used to fit the *BABAR* data, (either real experimental data, or detailed simulations) and extract a measurement of Δm . The rest of this Thesis will discuss the application of this analysis technique, using MCRW fits, to the measurement of Δm in fully-simulated B meson events.

4

Tracking Resolution

In this Chapter we examine the tracking resolution in the variables z_0 and d_0 in the *BABAR* experiment, through simulated data. We construct a physically motivated model that is simple to simulate in MC, and which describes the resolution well over many decades, and we use this model to construct a Monte-Carlo template that can be used to perform a MCRW fit to the single track resolution of simulated *BABAR* events. This tracking model could not be used in a conventional analytic likelihood fit, because a closed-form expression for the model could not be found.

4.1 Contributions to Δz Resolution

The *BABAR* detector (see Appendix C) has a nearly uniform magnetic field in the z direction, so a charged particle travels in a helix. Charged particles create hits in the 5 layer silicon vertex detector (SVT) and 40 layer drift chamber (DCH). Pattern recognition software organizes hits into tracks, and fitting software fits a helix to the hits of each track. The helix has a unique point of closest approach (POCA) to the z -axis of the *BABAR* coordinate system, which is within a few mm of the beam axis. The d_0 track parameter is the distance of the POCA from the z -axis, signed by the z -component of the angular momentum of the track around the z -axis. The z_0 track parameter is the z -coordinate of the helix at that point. Other track parameters are the ϕ and θ values of the tangent to the helix at that point, and the curvature of the helix.

The PEP-II beam envelope is a Gaussian ellipsoid, about 100 microns wide, less than 10 microns high, and 1 to 2 cm long. The centre and size of the beam envelope is determined from reconstructed tracks and updated continuously in recorded data. The $\Upsilon(4S)$ is produced within this Gaussian ellipsoid, as depicted schematically in Figure 4.1. The B and \bar{B} from the decay are produced nearly at rest in the $\Upsilon(4S)$ frame, but that frame has a boost of $\beta\gamma \simeq 0.56$ in the lab due to the energy asymmetry of the PEP-II beams. The B mesons both fly primarily in the boost direction and decay independently. The two B decay points nearly overlap in the plane perpendicular to the boost axis, and have essentially the Gaussian distribution of the beam envelope.

To find Δz for a dilepton pair, we first find the the x and y coordinates of the $\Upsilon(4S)$ production point. This is done by a 2D vertex finder using the centre and size of the $x - y$ beam spot, and the ϕ angles and d_0 coordinates of the two tracks, and the d_0 errors of the tracks, with an additional 25 microns added in

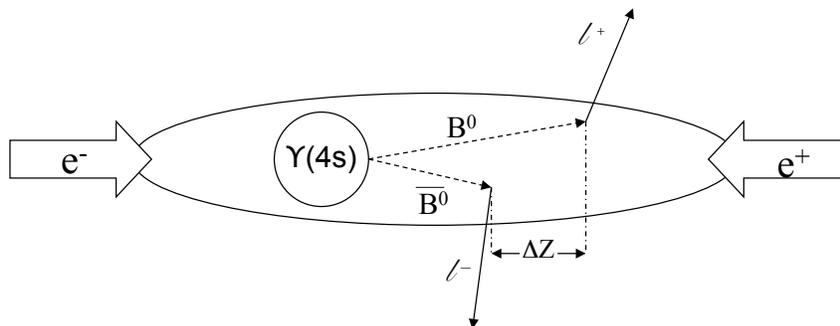


Figure 4.1: Schematic representation of the collision of an e^+e^- pair to create a $\Upsilon(4S)$, followed by the decay into a $B^0\bar{B}^0$ pair. Both B mesons decay, emitting direct leptons ($\ell^+\ell^-$), and the decay-point difference, Δz , is measured. The beam spot is actually much longer (~ 1.5 cm), than the flight length of the B mesons ($\sim 100\text{-}200$ μm).

quadrature to account approximately for the transverse flight of the B mesons from the $\Upsilon(4S)$ production point. Next, the POCA of each track helix to the axis defined by this $x - y$ point is determined. The Δz is this z of the POCA for the track with the higher p_{cm} minus the z of the POCA for the other track.

If a track is at $\theta = 90$ degrees, all points on the helix are at the same z coordinate, and it would not matter which point along the track helix was used. If both tracks were at 90 degrees, the only contribution to the resolution of Δz would be the z_0 resolutions of the individual tracks. As the track θ s deviate more from 90 degrees, the uncertainty of the $x - y$ vertex point adds to the uncertainty of which point along the helix to use for the z of the track, and thus the d_0 resolution contributes to the Δz resolution.

The z_0 and d_0 resolution for tracks have geometrical contributions from the strip width of the SVT layers and the spacing of the layers. Tracks typically deposit charge on several SVT strips and the hit coordinate is interpolated with a resolution considerably less than the width of a single strip, but there may be non-Gaussian tails on the SVT hit resolution.

For the tracks of interest in this analysis, the dominant contribution to the d_0 and z_0 resolution is from Coulomb scattering of tracks in the material of the beam pipe (and to a lesser degree the material of the SVT). It has a Gaussian core from a large number of small scatters, and a long power-law tail from rare large scatters. The angle scale for Coulomb scattering is inversely proportional to track momentum, and proportional to the square-root of the number of radiation lengths traversed, which depends on the track θ angle. The geometrical contributions to resolution are essentially independent of momentum but depend on the track θ and ϕ .

The track pattern recognition software is imperfect, so the wrong combina-

tion of hits may be fitted, resulting in long tails on d_0 and z_0 . Requiring that a track have hits in nearly all of the SVT layers greatly reduces the tails from imperfect pattern recognition, without rejecting very many correctly reconstructed tracks.

Electrons can lose energy due to bremsstrahlung in the beam pipe, SVT material and the inner wall of the drift chamber. This does not significantly change the θ or ϕ angle of the track, but causes an abrupt change in the track curvature. The track will be reconstructed with the post-brem curvature, and the d_0 parameter will be calculated with that curvature. The actual curvature before the brem was different, so the extrapolated d_0 will be wrong. Brems in the support tube and the drift chamber inner wall cause the track curvature to be different in the SVT and the DCH, which contributes to pattern recognition difficulties.

4.2 *BABAR* Single Track Resolution in Simulated Events

In simulated events, it is possible to measure the tracking resolution directly by comparing the true and observed track parameters for a lepton candidate. Figures 4.2-4.5 show the difference between true and reconstructed d_0 and z_0 parameters for leptons with $p_{cm} > 500$ MeV/ c in fully simulated *BABAR* MC events, for 4 cases. The first is with no additional cuts, the second is requiring hits in 2 of the inner 3 SVT layers, the third is also requiring either hits on both outer SVT layers or a hit in 1 outer layer and at least 10 hits in the drift chamber, and the fourth is also requiring that reconstructed p_T agree with the MC true p_T within 20% and that the reconstructed θ and ϕ agree with the true θ and ϕ within 0.1 radians.

It is clear that requiring two inner SVT hits substantially reduces the tails. Requiring the outer SVT hits or DCH hits makes a very small additional improvement with minimal loss of efficiency. The electron d_0 difference has large tails due to bremsstrahlung that are absent from the muon d_0 difference. These tails are substantially reduced by requiring the reconstructed momentum to match the true momentum within 20%. The electron z_0 difference has a smaller but still noticeable contribution from bremsstrahlung.

None of these distributions is very Gaussian, because they are sums of contributions that have Gaussian cores but with a continuous range of widths, and also tails from several sources. The previous *BABAR* dilepton mixing analysis treated the z_0 difference resolution as a sum of at most 3 Gaussian contributions, largely due to the limitations of fitting using analytical models, which is clearly a rather crude approximation.

Instead, we will use the power of the MCRW fitting method to construct a track by track model of the d_0 and z_0 smearing, based on geometrical contributions, core and tail Coulomb scattering contributions, and electron bremsstrahlung. We will fit some of the parameters of this model using the difference between

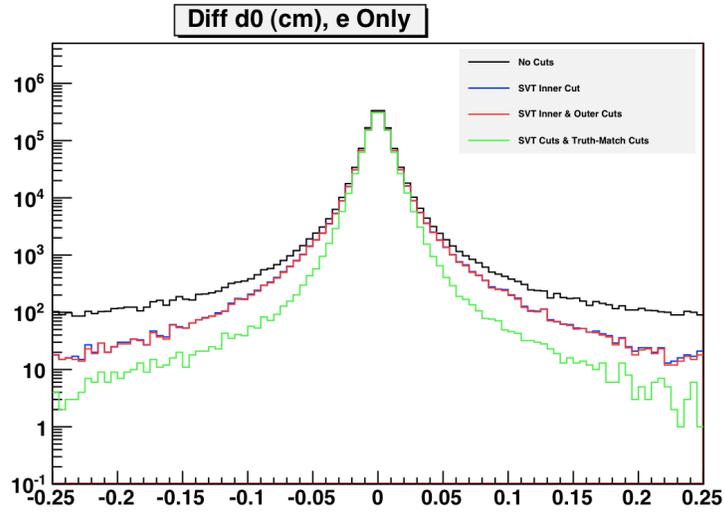


Figure 4.2: The difference between measured and true d_0 for true electrons in simulated $B\bar{B}$ events. The “truth-matching” cuts have a large effect on the electron d_0 tails.

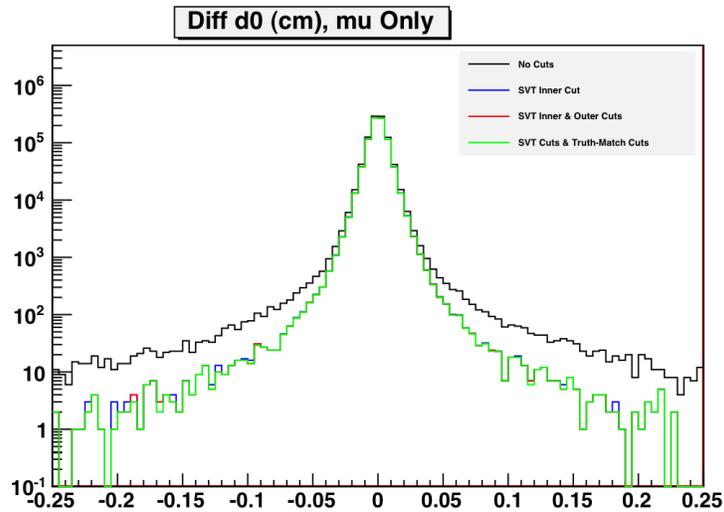


Figure 4.3: The difference between measured and true d_0 (in cm) for true muons in simulated $B\bar{B}$ events. The “truth-matching” cuts have almost no effect on the muon d_0 tails.

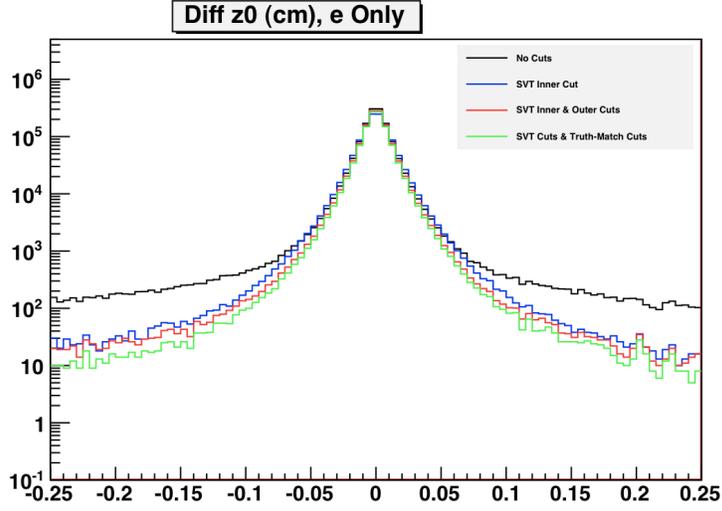


Figure 4.4: The difference between measured and true z_0 (in cm) for true electrons in simulated $B\bar{B}$ events. The “truth-matching” cuts have a small effect on the electron z_0 tails.

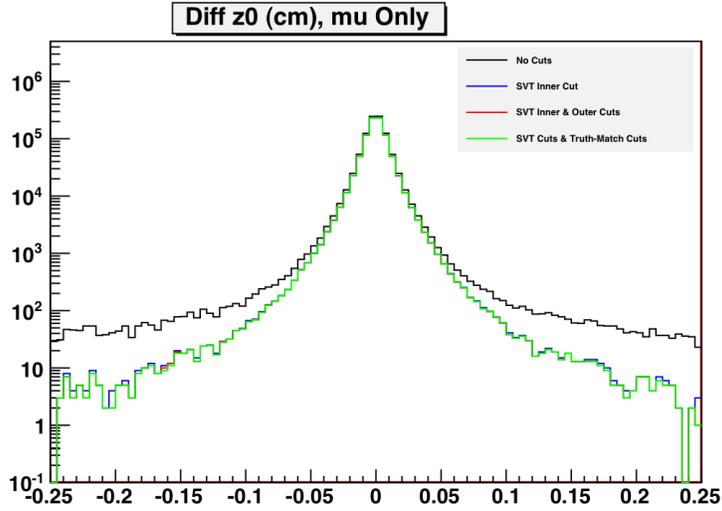


Figure 4.5: The difference between measured and true z_0 (in cm) for true muons in simulated $B\bar{B}$ events. The “truth-matching” cuts have almost no effect on the muon z_0 tails.

Resolution Variable	a	b	c
z_0	55	15	15
d_0	54	0.0	14

Table 4.1: Values for resolution parameters a , b , c , for both z_0 and d_0 in μm .

the true and reconstructed d_0 and z_0 values in simulated data, and others will be left floating in the final physics fit.

4.3 Track Resolution Model

We divide the track resolution model for both d_0 and z_0 into two main components; a core described by a sum of Gaussians, and a power-law tail. The core contains a dominant Gaussian, whose width is given by the following expression for the resolution in z_0 :

$$\sigma^2(z_0) = (c^2 + (b^2 \cot^2(\theta)) \cdot [1 + \sin^2(\phi_{SVT})]) + \left(a \frac{\sqrt{1 + \sin^2(\phi_{SVT})}}{\beta_\pi \cdot p_T^{0.9} \cdot \sin^{3/2}(\theta)} \right)^2 \quad (4.1)$$

and by the following slight variation (note the different power of $\sin(\theta)$) for the resolution in d_0 :

$$\sigma^2(d_0) = (c^2 + (b^2 \cot^2(\theta)) \cdot [1 + \sin^2(\phi_{SVT})]) + \left(a \frac{\sqrt{1 + \sin^2(\phi_{SVT})}}{\beta_\pi \cdot p_T^{0.9} \cdot \sin^{1/2}(\theta)} \right)^2 \quad (4.2)$$

Note that these are the quadrature sums of two Gaussians, the first depending entirely on geometrical factors, and the second containing a momentum dependence. The parameters a, b, c are determined separately for the resolutions of z_0 , and d_0 , and were obtained through a trial-and-error approach, involving repeated attempts to perform an analytical fit to the central 3 to 5 σ s of the muon resolution distributions shown in Figures 4.3 and 4.5. c is approximately 15 μm , and represents the spatial resolution of the SVT at normal incidence. b is also approximately 15 μm , and represents the polar angle (θ) dependence of the SVT hit resolution. a is approximately 55 μm , and represents the scale of the contribution to the resolution from Multiple-Coulomb Scattering (MCS) [10], [11]. This is consistent with the expected value given the density of beam pipe material, and its radius from the interaction point. Our final values for a , b and c for both z_0 and d_0 are given in Table 4.1.

The variable ϕ_{SVT} is the angle of incidence of the track on the SVT layers. The SVT has the shape of a hexagon in the $x - y$ plane. The factor of $\sqrt{1 + \sin^2(\phi_{SVT})}$ measures the effect of the radial distance from the track origin to a point on the SVT hexagon. The factor β_π is for the velocity dependence

of the width of the MCS Moliere distribution Gaussian core. The ‘ad-hoc’ exponent of -0.9 for the p_T dependence works noticeably better than using the expected value of -1 .

As part of the model for the MCS Moliere distribution, we also include a power-law tail, with a “Tail Fraction”, in the resolution for both d_0 and z_0 . This tail has a theoretical dependence proportional to r^{-3} , where r is the radial displacement from the original track trajectory. We use this same dependence in our model. The Tail Fraction or TF is allowed to assume different values for electrons and muons in the fits, and is normally 35% for electrons, and 33% for muons.

This model for the resolution is still somewhat idealized, and assumes that the SVT will function with perfect efficiency. Due to the possibility of missing hits in the SVT, and/or problems in the pattern-recognition software which associates hits to tracks, or in the fitting software which fits a helical trajectory to the hits, the core of the resolution distribution is actually somewhat wider. We can model this with a series of extra, wider Gaussians (normally one or two more is enough) in addition to the core Gaussian width given above. These additional terms generally are described by a fraction and a width, neither of which depend on any of the track kinematic variables. The physical meaning of the resolution model parameters is more clearly explained by describing exactly how the resolution model is applied to a simulated particle trajectory.

Before applying the main part of the tracking resolution model, we apply a simplified model of Bremsstrahlung emission to the electrons in our EvtGen sample. Bremsstrahlung emission has an energy spectrum proportional to $1/E$, and we generate the random fractional energy lost by the electrons, x , with the following:

$$x = e^{(r-1)/\lambda}, \quad (4.3)$$

where r is a random number between 0 and 1, and $\lambda = \lambda_0 / \sin(\theta)$ is the effective number of material lengths the electron traverses, dependent on the polar angle θ . In order to determine the appropriate value of λ_0 to use, we tune our model to match as closely as possible the energy loss spectrum for electrons in the fully-detailed *BABAR* simulation. From this we obtain a value of $\lambda_0 = 0.0475X_0$, or 4.75% of a radiation length. For those electrons whose energy loss would be greater than 10%, we remove them from the EvtGen sample, and allow them to be modelled in the full simulation, as Tertiary leptons, instead.

Once we have determined the energy loss of the electron, we calculate the apparent shift in the d_0 track parameter for the electron due to bremsstrahlung emission:

$$d_0\text{Shift} = \frac{xq}{(1-x)p_T} 141\mu\text{m}/\text{GeV} \quad (4.4)$$

where q is the charge of the particle (-1 for e^- and $+1$ for e^+). We also decrease the original leptons momentum and energy by the amount x , which assumes that the bremsstrahlung photon is emitted parallel to the electron path (a very good approximation). Finally we modify the original lepton origin point in space by the given d_0 shift in the x-y plane. We use the modified electron origin point

and momentum vectors as inputs to the following track resolution model, which applies to the muons as well.

The parameterizations for the d_0 and z_0 core resolution (Equations 4.1 and 4.2) are evaluated using the given track angles and $p_T=1000$ GeV/c to obtain the nominal geometrical resolution, and at the given track p_T to obtain the nominal resolution. Note that the track angles and momenta of the electrons have been previously modified by the bremsstrahlung smearing model above. The difference in quadrature of the nominal geometric resolution, and the nominal resolution, divided by the distance to the beam pipe, is the nominal scattering angle scale. The nominal scattering angle scale times the true track momentum is the momentum scattering scale. At this point we modify the momentum scattering scale by adding a small Gaussian random term, 15% of the width of itself. This additional noise term makes a noticeable improvement in the agreement between the model and the resolution in simulated *BABAR* data.

The shape of the scattering angle distribution is modelled by a radial Gaussian core, r_{core} with probability:

$$P_{\text{core}}(r_{\text{core}}) = r_{\text{core}} e^{-r_{\text{core}}^2/2} \quad (4.5)$$

and a radial power-law tail shape r_{tail} with probability:

$$P_{\text{tail}}(r_{\text{tail}}) = \begin{cases} r_{\text{tail}}^{-3} & \text{if } r_{\text{tail}} > 1, \text{ or} \\ r_{\text{tail}}^{+3} & \text{if } r_{\text{tail}} < 1. \end{cases} \quad (4.6)$$

We generate r_{core} in all cases, and in fraction TF of the cases we generate r_{tail} , otherwise we leave $r_{\text{tail}}=0$. This implementation of the tail generation means that the TF parameter actually describes a portion of the events lying within the central part of the of the distribution, when $r_{\text{tail}} < 1$, and thus the values of TF are much higher than one would naively expect when looking at only the far tails.

We generate a flat random ϕ_{core} and ϕ_{tail} in the range between 0 to 2π , and calculate:

$$\begin{aligned} x &= (r_{\text{core}} \cos(\phi_{\text{core}}) + r_{\text{tail}} \cos(\phi_{\text{tail}})) p_{\text{init}}, \\ y &= (r_{\text{core}} \sin(\phi_{\text{core}}) + r_{\text{tail}} \sin(\phi_{\text{tail}})) p_{\text{init}}, \\ z &= \sqrt{p_{\text{init}}^2 - x^2 - y^2}. \end{aligned} \quad (4.7)$$

This vector has magnitude p_{init} , which was the track momentum before beginning the scattering process, and is scattered from the z direction by a Gaussian core and a power law tail. The rotation that aligns the z axis with the true track direction is applied, and the rotated vector is used as the scattered momentum vector. The scattering is assumed to occur at the beam pipe radius, and the POCA of the track line to the beam axis gives the scattered d_0 and z_0 .

We then smear the d_0 and z_0 by random numbers Δ from a sum of Gaussians with the probability:

$$P_{\Delta}(\Delta, \vec{\sigma}) = \sum_j^n f_j \exp(-\Delta^2/2\sigma_j^2)/\sigma_j. \quad (4.8)$$

For d_0 we use only a single Gaussian term, for z_0 we use two terms.

At this point the track scattering and smearing process is complete, and we record r_{core} , r_{tail} , and the two Δ values and three σ values for re-weighting the track to different resolution parameters.

4.4 MCRW Fits to Simulated Track Resolution

Using the resolution model described above, we can perform a MCRW fit to the resolution in z_0 , or Δz for single tracks in the *BABAR* simulated data set. We separate the tracks into electrons and muons, and divide the sample into several bins of momentum, to allow for the possibility that the momentum dependence in the model is not completely accurate. For each true lepton in the simulated data, we record the true track co-ordinates, as well as the observed or reconstructed track coordinates, and histogram the quantity Δz , which is the difference between the true and reconstructed values.

We use the EvtGen simulation to produce a roughly comparable number of events simulated at the event-generator level, and apply the models of electron bremsstrahlung smearing, reconstruction efficiency, and PID efficiency, as discussed in Chapter 5 and 6. We apply the track smearing model described above (after the bremsstrahlung and before reconstruction efficiency), and use the difference between the generated and smeared values of z_0 as the quantity to histogram for the template. We store the true lepton momentum, and initial position in the template. We also must record the amount of scattering applied to each track, whether the scatter was in the Gaussian core or the power-law tail, and the amount of d_0 and z_0 smearing applied, in order to properly re-weight the template in the fit. We apply the same cuts to the lepton selection as those used for the event selection in Chapter 5. We require $p_{cm} > 0.8$ GeV/c, and $0.7 < \theta < 2.3$. After applying the cuts to all *BABAR* Run 6 fully simulated $B\bar{B}$ events, we obtain a data histogram containing ~ 15 million fully simulated leptons. After applying the same cuts, the template contained ~ 26 M EvtGen leptons, so the overall normalization was $q \simeq 0.61$.

We introduce an additional parameter, the Resolution Scale (RS) which scales the overall width of the track resolution as indicated in Equations 4.9-4.13. This parameter is one which is varied in the fit, and it has a separate value for electrons and muons. The parameter TF is the ‘‘Tail Fraction’’ described above, and it can also be varied in the fit, with separate values for electrons and muons. For each lepton in the template pair, the track resolution weight is given by the following:

$$w_{\text{track}} = w_{\text{core}} \cdot w_{\text{tail}} \cdot w_{d_0} \cdot w_{z_0} \quad (4.9)$$

$$w_{\text{core}} = \frac{P_{\text{core}}(r_{\text{core}}/RS)}{P_{\text{core}}(r_{\text{core}}) \cdot RS} \quad (4.10)$$

$$w_{\text{tail}} = \frac{P_{\text{tail}}(r_{\text{tail}}/RS)}{P_{\text{tail}}(r_{\text{tail}}) \cdot RS} \cdot \frac{TF_{\text{fit}}}{TF_{\text{gen}}} \quad \text{if } r_{\text{tail}} > 0$$

$$= \frac{1 - TF_{\text{fit}}}{1 - TF_{\text{gen}}} \quad \text{if } r_{\text{tail}} = 0 \quad (4.11)$$

$$w_{d_0} = \frac{P_{\Delta}(\Delta, RS \cdot \sigma_{d_0})}{P_{\Delta}(\Delta, \sigma_{d_0})} \quad (4.12)$$

$$w_{z_0} = \frac{P_{\Delta}(\Delta, RS \cdot \sigma_{z_{0,0}}, \sigma_{z_{0,1}})}{P_{\Delta}(\Delta, \sigma_{z_{0,0}}, \sigma_{z_{0,1}})} \quad (4.13)$$

Note that the second Gaussian in z_0 is not scaled by RS.

As long as the fit starts with values reasonably close to the final fitted ones, we can simultaneously float all of the following parameters in these fits:

- Resolution Scale for Electrons (RS-e)
- Resolution Scale for Muons (RS-mu)
- Tail Fraction for Electrons (TF-e)
- Tail Fraction for Muons (TF-mu)
- 2nd Core Gaussian Fraction (ResF1)
- 2nd Core Gaussian Width (Sigma1)
- Ratio of Electrons in Low, Medium, High p_{cm} Bins (Ne0, Ne1, Ne2)
- Ratio of Muons in Low, Medium, High p_{cm} Bins (Nmu0, Nmu1, Nmu2)

The Low, Medium, and High p_{cm} bins are defined by Low = $0.5 < p_{cm} < 1.0$, Medium = $1.0 < p_{cm} < 1.5$, and High = $p_{cm} > 1.5$ GeV/c. The ‘‘Ratio’’ parameters measure the ratio between the number of leptons in the bin for the data and the template, after the overall normalization has already been taken into account. In this way, they can be thought of as parameters which adjust the relative shape of the momentum spectrum of the template compared to the data.

As seen in Figures 4.5 and 4.4, the resolution quantities span many decades of scale, and it would be easier to fit and view the data if it were transformed in a way that flattened the histograms onto a narrower vertical scale. To do this, we apply the following transformation:

$$\Delta z' = N \begin{cases} \text{erf}(\Delta z/\sqrt{2}) & -K \leq \Delta z \leq +K \\ A + B(\Delta z^{1-P} - K^{1-P}) & \Delta z > +K \\ -A - B(\Delta z^{1-P} - K^{1-P}) & \Delta z < -K \end{cases} \quad (4.14)$$

which is essentially flattening the Gaussian core of the distribution with an error-function, and flattening the power-law tails with another power-law. The parameter K represents the number of σ at which the transition from the Gaussian flattening to the power-law flattening occurs, and we choose $K = 2.5$. The parameter P is chosen as the power by which to flatten the tail, and we use $P = 5$, which flattens the curve enough in the more central regions where it

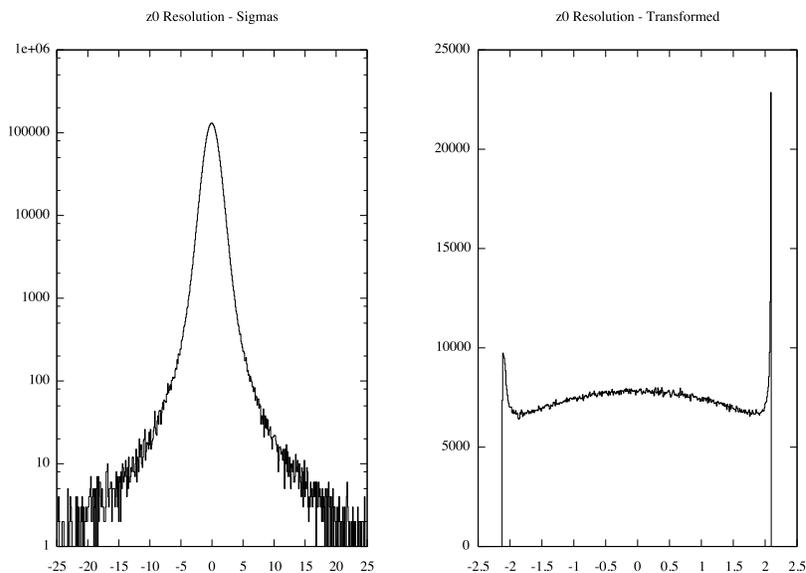


Figure 4.6: Example of the application of the transformation from Δz (left) to $\Delta z'$ (right). Note the different horizontal and vertical scales.

falls more steeply than x^{-3} , and too much in the very far tails, where the fall-off really does obey x^{-3} , so that the distribution falls off very rapidly past $\pm K$. The constants $A = \text{erf}(K/\sqrt{2})$ and $B = \frac{2K^P e^{-K^2/2}}{\sqrt{2\pi}(1-P)}$, are chosen to ensure that $\Delta z'$ is continuous in both the value and first derivative across the transition from the core to the tail at the points $\pm K$. $N = \frac{C}{\text{erf}(C/\sqrt{2})}$ is an overall normalization factor, so that $\Delta z = \Delta z'$ at the points $\pm C$. We have found $C = 2$ to give particularly informative and useful plots, although it can be varied as desired.

An example of the application of the transformation in Equation 4.14 is depicted in Figure 4.6. The same data histogram, containing the z_0 resolution for roughly 1M electrons in simulated *BABAR* events is presented in the normal Δz variable on the left (in units of σ), and in the transformed $\Delta z'$ variable on the right. We see that the distribution in $\Delta z'$ is much flatter, it does not require a logarithmic scale for display, that the core from 0 to 2 Gaussians is quite flat, and only the far tails near 2.5 are sharply shaped. Apparently small differences in the amount of tails are much clearer to see in the transformed variable, and the MCRW fits achieve slightly better results (lower χ^2/DOF) in the transformed variable.

Table 4.2 gives the numerical results of the fit. The fit was binned using 6 different p_{cm} bins, 3 each for electrons and muons. There were 500 bins in

Parameter	Template Value	Fit Value	Fit Error	Pull (σ)
RS-e	0.8401	0.8408	0.0012	+0.62
RS-mu	0.8549	0.8546	0.0013	-0.26
TF-e	0.3518	0.3468	0.0022	-2.23
TF-mu	0.3329	0.3320	0.0023	-0.37
ResF1	0.1865	0.1852	0.0069	-0.19
Sigma1	0.00399	0.00400	$5.9 \cdot 10^{-5}$	+0.26
Ne0	1.000	1.002	0.0010	+2.50
Ne1	1.000	1.002	0.0007	+2.35
Ne2	1.000	0.996	0.0007	-6.14
Nmu0	1.000	1.004	0.0012	+3.68
Nmu1	1.000	1.002	0.0007	+2.71
Nmu2	1.000	0.998	0.0007	-2.83

Table 4.2: Results of a MCRW fit to single track resolution $\Delta z'$ for ~ 15 M fully simulated leptons.

the transformed variable $\Delta z'$, ranging between -2.5 to +2.5 for each of the p_{cm} bins, for a total of 3000 bins. The outermost bins were empty for these fits, giving a total number of Degrees of Freedom (DOF) = 2520. The fit had a $\chi^2 / \text{DOF} = 2763/2520$. We present a graphical version of these fit results in Figures 4.7 and 4.8. In Figure 4.7, the error bars are the quadrature sum of the statistical error on the data, the statistical error on the template, and the error in the average weight of the template. In Figure 4.8, the pulls are calculated as (Data-Template)/Error, and are plotted with error bars of unit size, to aid in visualization. In each of these plots, only every third point is plotted, to improve legibility. We make an exception, however, for the three outermost points in the tails of each histogram.

By looking at the overlays, we cannot clearly see any significant discrepancies between the data and the template. In the pulls plot, we can see individual bins which are 2 or 3 σ s from zero, but they are generally isolated, and a few of these can be expected when we examine 500 bins in total.

We iterated the process of generating and fitting the resolution template several times, in order to minimize the differences between the data and the template. This was because when we move on to fitting Δt for lepton pairs, with additional physics model details, it is very difficult to float all of the resolution parameters, because they are highly correlated, and not resolved as easily in a lepton pair Δt fit, as they are here in a single lepton Δz fit. So before fitting in Δt , we fix several resolution parameters to the values we obtain here, in the single track fits, and we want them fixed to the values which best describe the resolution in simulated *BABAR* data, so that they don't introduce significant systematic uncertainties in the Δt fit.

This fit had statistically significant pulls in the normalization parameters, indicating that the high p_{cm} bins contained fewer leptons in the data, than

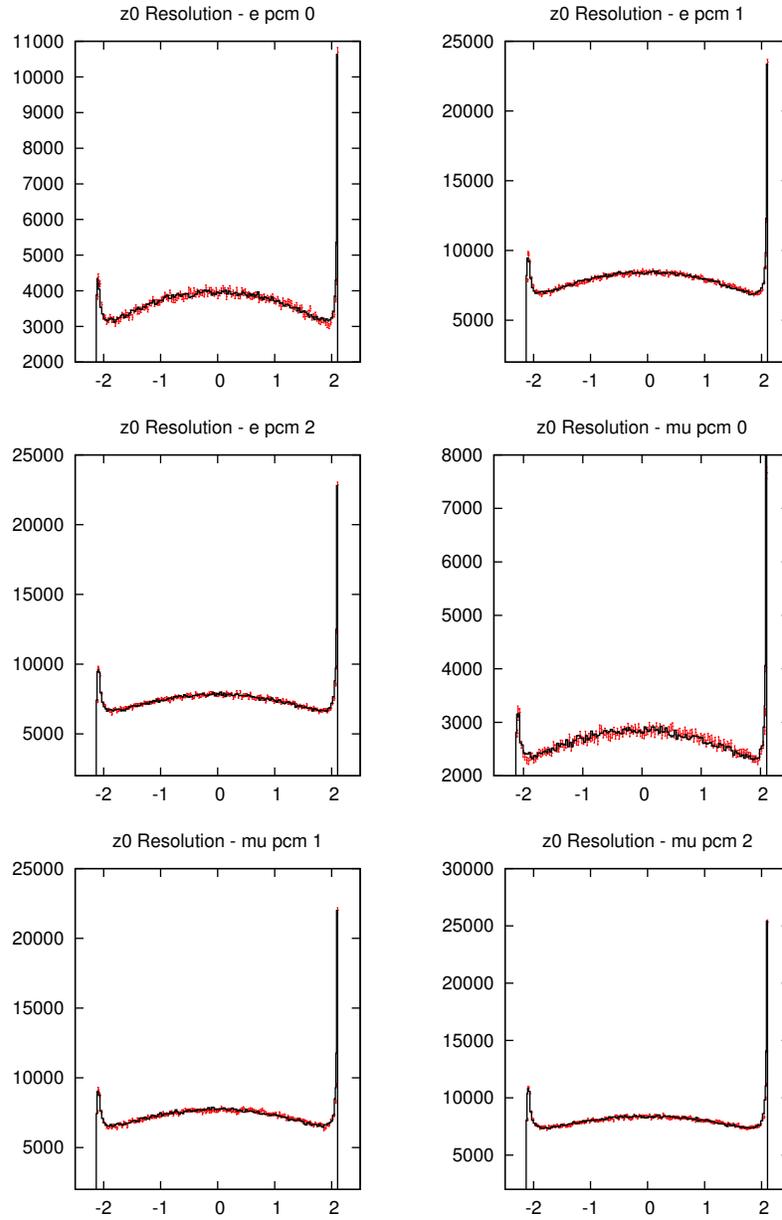


Figure 4.7: Overlay of data and template in the MCRW fit to $\Delta z'$. Plots show three p_{cm} bins for electrons, and three p_{cm} bins for muons. The solid line represents the template, and the points with errors are the data.

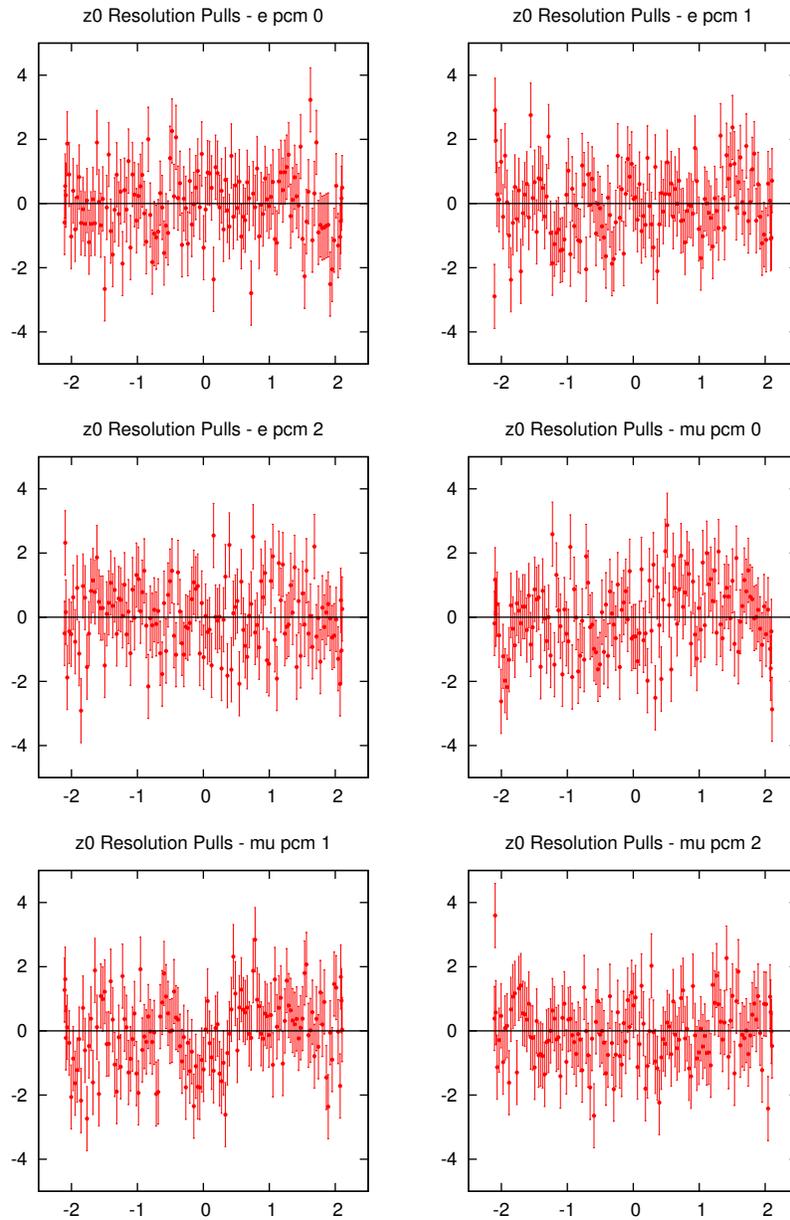


Figure 4.8: Pulls $((\text{Data}-\text{Template})/\text{Error})$ of the MCRW fit to $\Delta z'$ (right). Plots show three p_{cm} bins for electrons, and three p_{cm} bins for muons.

in the template, while the lower p_{cm} bins were more heavily populated in the data than in the template. This indicates some inadequacy in our efficiency modelling for the lepton spectra, used to produce the EvtGen template, at the level of a few parts per thousand. This disagreement didn't appear to affect the results when we fit for the track resolution parameters, and we don't believe that it will affect the results when we fit for Δm in a noticeable way.

In these single-track resolution fits, we notice the fact that RS came out to about 0.85, rather than the value of 1.0 one might expect. That means that the core of the fitted data looks like a narrower Gaussian than described by the resolution model in Equations 4.1 and 4.2. This is because the PDG convention for the width of MCS (which the *BABAR* and GEANT software follow, and which is the basis of the parameterizations) is to include some of the non-Gaussian tails when determining the effective σ . Our more detailed model has additional, physics-motivated, contributions to the tail, rather than widening the core to account for part of the tail.

To summarize, in this Chapter we examined the tracking resolution in the variables z_0 and d_0 in the *BABAR* experiment, in both simulations and data. We constructed a physically motivated model that is simple to simulate, and which describes the resolution well over many decades, and we used this model to construct a Monte-Carlo template that can be used to perform a MCRW fit to the single track resolution of simulated *BABAR* events. This resolution model works well with the MCRW fit software, but could not be used in a traditional analytic likelihood fit because it doesn't have a closed-form expression. The single-track resolution fit serves as a successful test of the MCRW fit software, and the results of that fit will be used when we prepare a template to fit the Δt distribution, in later Chapters.

5

Track and Event Selection Efficiency

This Chapter details the procedure that is used to select dilepton events from the *BABAR* data set for analysis in a MCRW fit, designed to extract the neutral B meson mixing frequency, Δm . In this Thesis we apply this selection procedure to fully-simulated $B\bar{B}$ events, and Off-Peak data. We also apply the same selection requirements to the generator-level and fully-simulated events that are used to make up the template for the MCRW fit. To apply these selection criteria to the generator-level simulated events requires the use of efficiency tables which are measured in fully-simulated events.

5.1 Event Selection

The dilepton method of measuring neutral B meson oscillations requires events with two leptons (electrons and/or muons), and doesn't really depend on any other information about the event. At the same time, however, it is important to try and minimize the backgrounds that make an accurate measurement more difficult. The analysis should reduce background from non $B\bar{B}$ processes (QED production of e^+e^- and $\mu^+\mu^-$ pairs, $\tau^+\tau^-$ pairs, and $udsc$ pair creation), which are dominated by the leptons produced in the $e^+e^- \rightarrow e^+e^-$ and $e^+e^- \rightarrow \mu^+\mu^-$ reactions. We also need to reduce the background from pion, kaon and muon decays-in-flight, photon conversions, and fake leptons, since these will not carry any useful information for the fit. Finally we want to have good resolution on the measurement of the lepton z_0 (and d_0 for vertexing) track parameters, to obtain good resolution on the measurement of Δt .

In order to ensure that the EvtGen based template is an accurate model of the B decays in the data sample, it is necessary to use relatively simple quantities to select our events from the full data set, so that the efficiency of the selection criteria applied to the data can be accurately modelled in the generator-level simulation, where detector related information is not available.

We allow for more than one lepton pair per event. If we tried to choose the "best" pair, we would need to simulate the masking effect of fake leptons, leptons from decays-in-flight, and other sources which blocked one of our signal leptons. This simulation would be extremely difficult, so instead we allow multiple pairs per event, and we find that our reconstructed MC sample has an average of 1.1 pairs per event passing all the selection requirements.

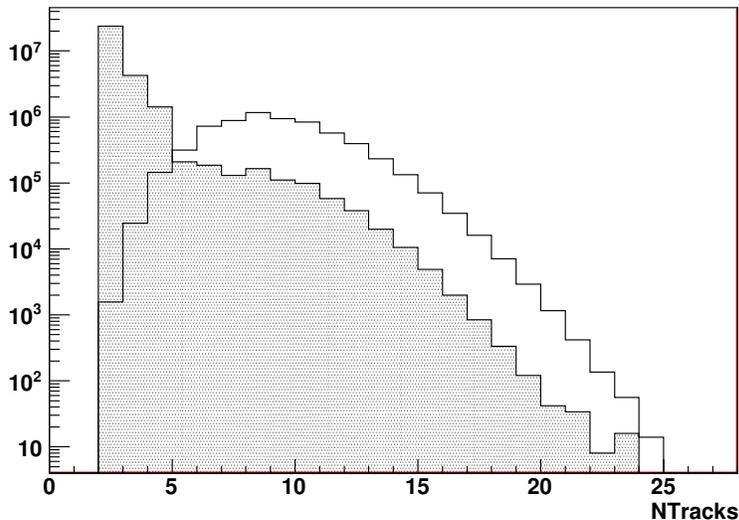


Figure 5.1: The number of Charged Tracks $NTracks$ for a sample of offpeak data (shaded histogram) and $B\bar{B}$ MC events (open histogram). The offpeak data is scaled to match the integrated luminosity of the MC.

We apply several requirements designed to reduce the number of non- $B\bar{B}$ events in the data sample. There are very large amounts of QED process lepton pairs in the data initially (even after the *BABAR* trigger system removes the majority of them), which tend to have roughly half the total CM energy for each lepton, or roughly 5 GeV per lepton. By imposing a maximum of 2.8 GeV/c on p_{cm} , we remove most of these leptons, and leave all leptons resulting from B meson decay in the sample.

Additionally, B meson events tend to have an average charged track multiplicity of around 10, and these QED processes have multiplicity generally in the range of 1-4 visible charged tracks, as seen in Figure 5.1. The continuum events with higher multiplicities are from $udsc$ processes. This Figure was created by plotting the $NTracks$ variable for all Run 6 $B\bar{B}$ MC events with at least two of our selected leptons (described below), and the corresponding distribution for the Run 6 offpeak data was prepared with the same conditions imposed, but the number of off-peak events was scaled up to match the integrated luminosity of the Run 6 $B\bar{B}$ MC in the plot. The ChargedTracks list is the fundamental list of all charged particle candidates that is available to *BABAR* users after the reconstruction process, and $NTracks$ is simply the length of this unrefined list.

Based on the relative $NTracks$ distributions for $B\bar{B}$ events and offpeak data, we can see that by imposing a requirement of $NTracks \geq 3$, we can remove a great deal of the QED background, while removing a relatively small amount

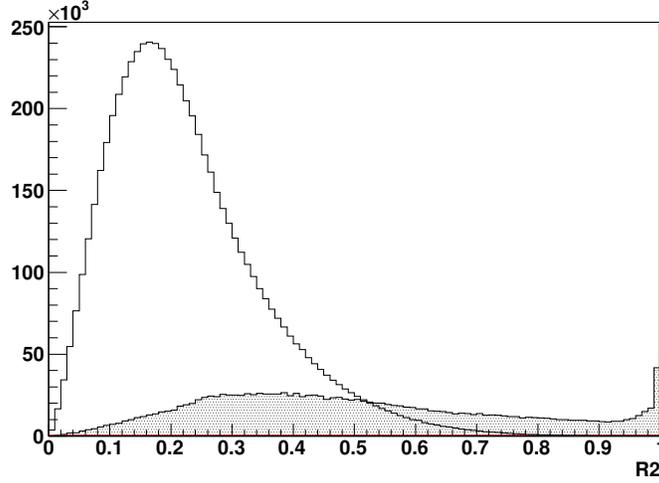


Figure 5.2: The quantity R_2 for a sample of offpeak data (shaded histogram) and $B\bar{B}$ MC events (open histogram). The offpeak data is scaled to match the integrated luminosity of the MC.

of $B\bar{B}$ events. As will be shown in Chapter 8, the optimal requirement for the Δm measurement appears to be $NTracks \geq 5$.

We can use the event shape as another method of discriminating between the lepton pair, or $q\bar{q}$ events, and the $B\bar{B}$ events. In general, the $q\bar{q}$ or lepton pair events will have two tracks or “jets”, in roughly opposite directions in the x-y plane. B pair events are much more spherically distributed in the CM frame, and in the x-y plane. The ratio of the second and zeroth Fox-Wolfram moments (R_2) [16] is a commonly used shape variable for discriminating between these “jetty” and spherical events. It is calculated according to the formula given in Equation 5.1,

$$R_2 = \sum_{i,j} \frac{|p_i||p_j|}{E_{vis}^2} P_2(\cos \theta_{ij}), \quad (5.1)$$

in which $p_{i,j}$ are the momenta of the particles, E_{vis} is the total visible energy in the event, θ_{ij} is the angle between particles i and j , and P_2 is the second Legendre polynomial. As shown in Figure 5.2, “jetty” events create a small peak towards $R_2=1$, while the majority of B pair events have R_2 between 0.0 and 0.5. The R_2 distributions shown in this Figure were prepared with the same selection requirements imposed on the events as in Figure 5.1 above, and the additional requirement that $NTracks \geq 3$.

The R_2 criterion is not as powerful in our final analysis, especially after the p_{cm} and $NTracks$ cuts have already been applied, but we optimize it to be $R_2 \leq 0.8$ in Chapter 8 in order to minimize the total error in our Δm

measurement, both statistical and systematic.

5.2 Lepton Selection

The main criteria for selecting an event for inclusion in our fits are the presence of two high-energy leptons. We begin by searching the list of GoodTracksLoose, which contains those tracks from ChargedTracks which meet the additional requirements of:

- $|d_0| < 1.5$ cm
- $|z_0| < 2.5$ cm
- $p_T > 0.05$ GeV/c
- $|p| < 10$ GeV/c

GoodTracksLoose also has so-called “V0 daughters” removed when it is derived from ChargedTracks. These are the daughters of long-lived decays like K_s and Λ , which are found by a dedicated code module in the *BABAR* reconstruction sequence. We then remove all electrons and positrons from the GoodTracksLoose list that are identified by the “Gamma Conversion Finder” standard *BABAR* software as the daughters of a converted photon.

We require that the electrons and muons lie within the polar angle region $0.7 < \theta < 2.3$ to ensure that they traverse the vertex detector and the drift chamber, and that the tracks should be well reconstructed and have a good chance of being correctly identified. We discuss the optimization of these polar angle requirements to the above values in Chapter 8.

We require that the electrons and muons have a centre-of-mass momentum (p_{cm}) of at least 0.8 GeV/c, in order to ensure they have enough energy to cross the drift chamber in a relatively straight line, and their momentum resolution will be good. In addition, leptons from B decays tend to have higher p_{cm} values, while those originating from charm or other processes will normally have lower p_{cm} values. Raising the lower limit on p_{cm} selectively removes these leptons, and this criterion is optimized in Chapter 8.

We use the most recent and sophisticated particle identification (PID) software available within the *BABAR* collaboration for identifying electrons and muons.

For electrons we use the “eKM” selector, which is based on “ECOC” or error-correcting output codes, and uses information from the SVT, DCH, EMC and DIRC to probabilistically categorize particles as electrons. We record the so-called “PID bitmap”, for each candidate, so that we may apply different selections, if necessary, in the analysis. We require that all electron candidates pass the eKM “Tight” criteria. The choice between the “Loose” and “Tight” electron selectors is optimized in Chapter 8. The eKM selectors are documented in an internal *BABAR* document [17]. We also use a “Gamma Conversion Finder”, which attempts to fit electron-positron pairs to the trajectory associated with

measured photon clusters in the calorimeter, and removes those e^+e^- pairs from the list we use. After applying this algorithm to simulated events, we find an extremely small number of electrons and positrons in our sample resulting from converted photons, and so we model this process with fully simulated events. We don't make any attempt to correct for, or remove electrons which have emitted bremsstrahlung photon(s) at this point in the selection process.

For muon identification, we use a “Bagged Decision Tree”, which incorporates information from all of the detector subsystems in order to probabilistically assign a likelihood of being a muon to candidate charged tracks. We record the “PID bitmap” for each candidate, so that we may apply tighter selections, if necessary, later in the analysis. We require that all muon candidates pass the muonBDT Low-Fake-Rate “Tight” criteria. The choice between “Loose” and “Tight” muon selectors is optimized in Chapter 8. The muonBDT selectors are documented in an internal *BABAR* document [18]. It is worth noting that there was a substantial improvement in the efficiency of the muon particle selectors for Runs 5 and 6 data, due to the replacement of many parts of the *BABAR* IFR hardware at the end of Runs 4 and 5.

These PID selectors offer the best ratio of true/fake leptons for the momentum and angular ranges we consider in this analysis, and in the context of events containing pairs of B mesons. To determine whether or not the “Loose” selectors are adequate, or whether it would be better to use tighter selectors, we must take into account the effects of background leptons, and other sources of error. We examine the uncertainty of the Δm measurement as a function of the choice of selectors in Chapter 8.

With this initial selection of our lepton pairs in place, we further refine and “clean” the sample by examining the tracking measurements in detail. Specifically, we wish to reduce the tails of the Δz distribution for leptons, where Δz is the difference in z -coordinate between the true and observed track (only available in simulations). We place requirements on the number of SVT and DCH hits of the tracks, as discussed in detail in Chapter 4.

We now prepare to simulate the effect of these selection requirements on EvtGen, or generator-level MC events, which do not contain any reconstructed quantities. In order to carry out this, we must measure the efficiency of these selections on the simulated events.

5.3 Tracking Efficiency Tables

As part of the preparation of an EvtGen template, we require the calculation of quantities simulating the various reconstructed variables that are used to apply selection criteria to real or fully simulated *BABAR* data. Consider, first, the variables $NTracks$ and R_2 . In reconstructed events, these quantities are based on the list of track candidates called “ChargedTracks”. This is the loosest user-level list of candidate charged particles available in the *BABAR* framework. $NTracks$ is the length of the list, and R_2 is calculated according to Equation 5.1 using the list of ChargedTracks. This is essentially the list of all charged candidates

reconstructed in an event.

5.3.1 ChargedTracks Efficiency Table

To simulate selection criteria on $NTracks$ and R_2 in EvtGen events, we need to measure the efficiency with which a given EvtGen particle appears on the ChargedTracks list, and apply it to our generator level events, in order to determine effective values of $NTracks$ and R_2 . To do this we first construct a list of all the Monte-Carlo charged particles that could appear on the reconstructed ChargedTracks list. This includes all positive and negative electrons, muons, kaons, pions and protons in the event, but excludes the daughters of photon conversions ($\gamma \rightarrow e^+e^-$), and requires that the particles were created within the radius of the beam pipe, and within 5 cm in the z direction of the centre of the detector. We exclude the photon conversion electrons here because we remove them from the reconstructed events with a dedicated code module. We then check the entire ChargedTracks list and examine each reconstructed track to see if it can be “truth-matched” to one of these selected MC particles. If so, we consider the MC particle as “found”, and if not, as “missed”. In this sense, our measurement of “ChargedTracks” efficiency also includes implicitly the efficiency of the imperfect MC truth-matching algorithm. We measure the efficiency as a function of both the polar angle θ and transverse momentum p_T of the tracks, and the results are shown in Figures 5.3, 5.4, and 5.5. These efficiency tables are only used in the calculation of the values of $NTracks$ and R_2 for EvtGen events.

5.3.2 Lepton Reconstruction Efficiency Tables

We divide the efficiency to detect real electrons and muons into two parts for modelling with EvtGen leptons, “Reconstruction” and particle identification (PID). The procedure for each is similar to that described above for the measurement of the ChargedTracks efficiency. We note, however, that the use of the name “Reconstruction” is not identical to the basic procedure of reconstructing a charged track, and we include a number of other effects. We construct a separate lepton reconstruction efficiency table for a subtle reason. If an electron undergoes bremsstrahlung, the reconstructed momentum can be quite different than the original momentum, and it is the reconstructed momentum that controls the efficiency of particle identification. So the electron PID efficiency table discussed in the next section is binned by reconstructed momentum in simulated events. But then the table measures the PID efficiency, given that the track was reconstructed at all. So we need to apply both a reconstruction efficiency and a PID efficiency to tracks simulated by EvtGen.

To measure the “Reconstruction” efficiency, we begin with fully simulated $B\bar{B}$ events, and select all true MC leptons in the event. Then we examine the list of GoodTracksLoose, which contains those tracks from ChargedTracks that meet the additional requirements stated above. We impose additional requirements before declaring a lepton to have been definitely detected and

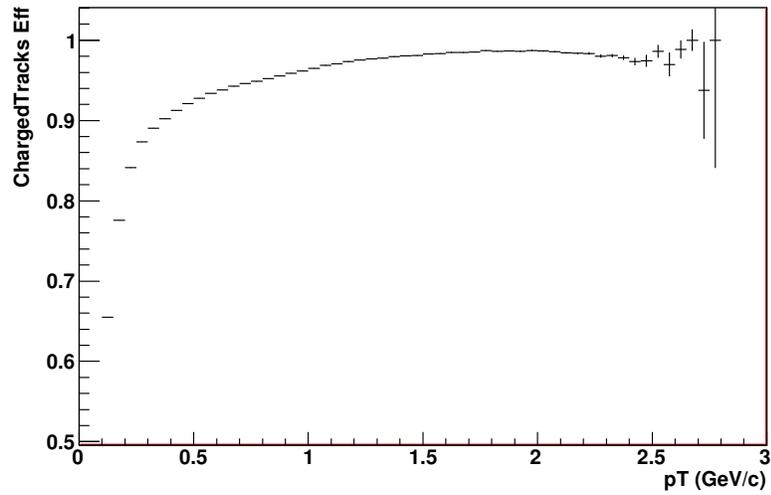


Figure 5.3: The measured ChargedTracks efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.

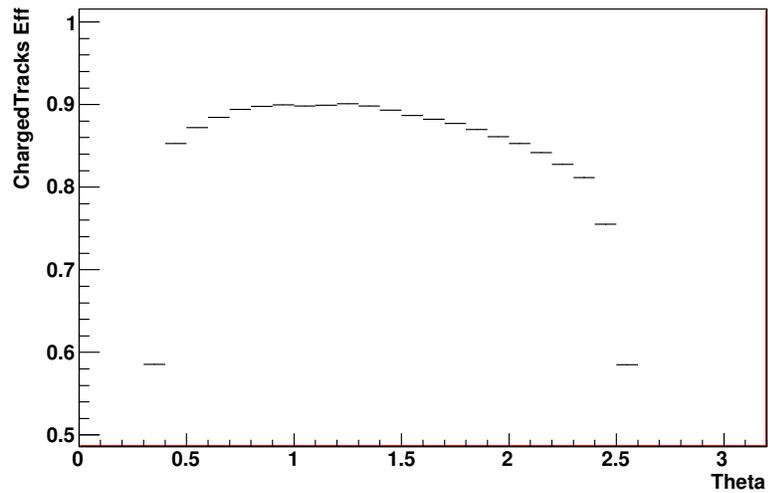


Figure 5.4: The measured ChargedTracks efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.

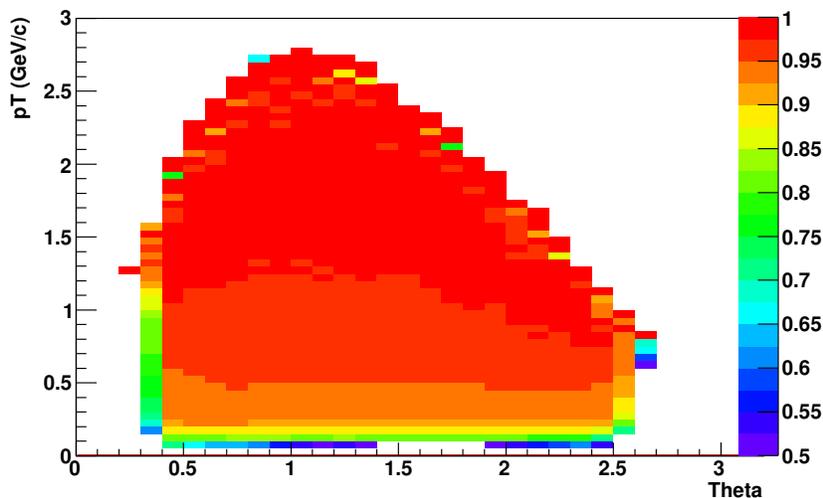


Figure 5.5: The measured ChargedTracks efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.

reconstructed to our satisfaction. We require that the true lepton originates within 2 cm in the x-y plane and 5 cm in z of the beam spot. We remove all true leptons which originate from long-lived kaon or pion decays. We remove all true leptons which are the daughters of a bremsstrahlung electron. We require that the reconstructed tracks meet the criteria specified above, intended to improve the d_0 and z_0 resolutions, by placing requirements on the number of SVT and DCH hits. Finally, we require an improved truth-match, designed to remove large energy-loss bremsstrahlung electrons from the calculations. This requires that the reconstructed and true tracks agree on θ within 0.1, ϕ within 0.1, and p_T within 5%. By imposing all of the additional criteria described here, the “Reconstruction” efficiency is substantially lower than the “ChargedTracks” efficiency measured previously.

The next step is to perform a search for “truth-matched” pairs on this list and the MC truth list, noting whether or not each true MC electron and positron is truth-matched to a reconstructed electron or positron from our refined list, or not. At this point, we declare each remaining true lepton to either be reconstructed, or not, and measure the reconstruction efficiency as a 2D function of both θ and p_T .

Figures 5.6, 5.7, and 5.8 present the results of these measurements of reconstruction efficiency for fully simulated $B\bar{B}$ events generated with conditions matching *BABAR* Run 6.

We measure the reconstruction efficiency separately for MC electrons and muons, finding only a very small difference between them, which we ascribe

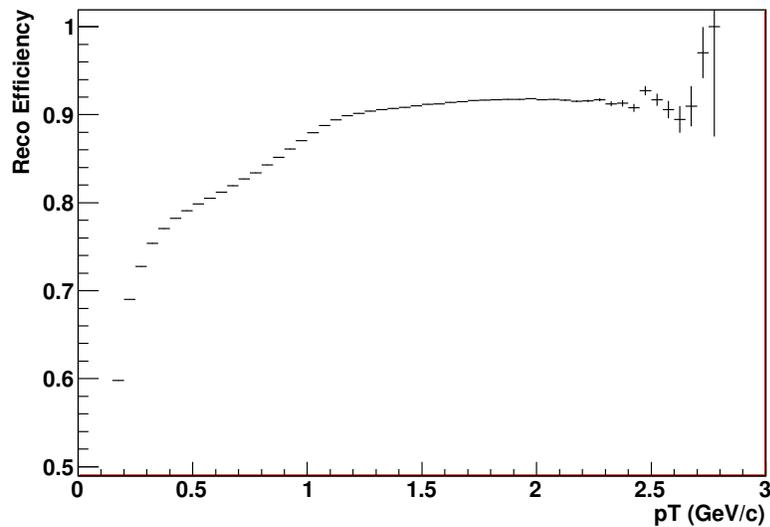


Figure 5.6: The measured muon reconstruction efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.

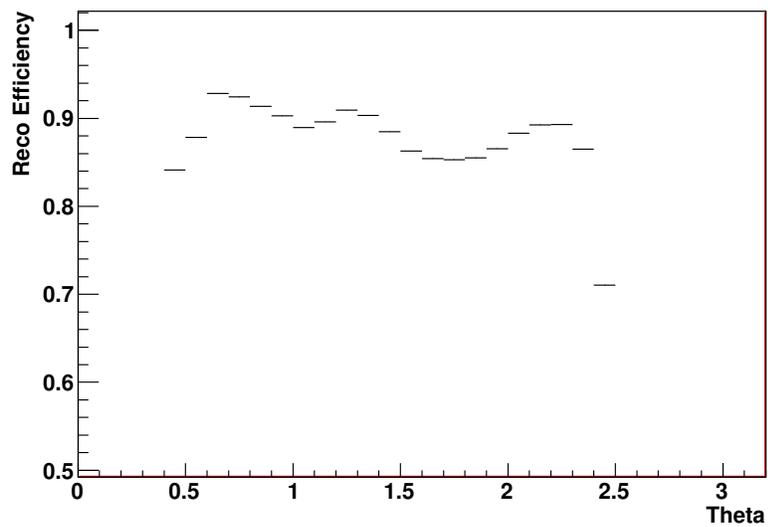


Figure 5.7: The measured muon reconstruction efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.

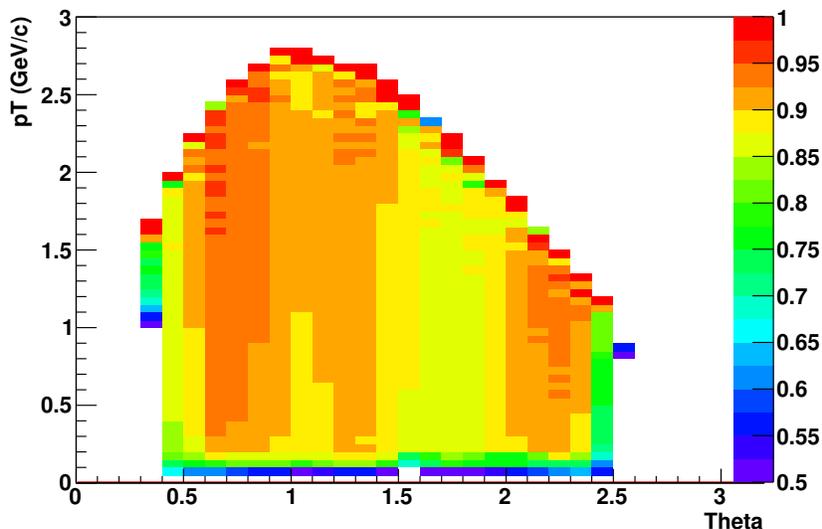


Figure 5.8: The measured muon reconstruction efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.

to the effect of bremsstrahlung emission on the electrons, and the “improved truth-matching” requirements designed to remove these from the samples. This process is repeated for each of the six major Run periods of the *BABAR* experiment, to allow for changes in the detector performance as a function of time.

5.4 Lepton Identification Efficiency Tables

In order to measure the PID efficiency, we repeat the process described above for the determination of the “Reconstruction” efficiency, but we require additionally that the electrons are passed by our chosen “eKMTight” selection routine, and the muons are passed by our chosen “muBDTConstFakeTight” selection routine. We then measure the efficiency as the probability of a true, reconstructed electron or muon to pass each selector, as a 2D function of particle θ and p_T . In other words, we are measuring the conditional probability of identifying a track as an electron or a muon, once it has already been reconstructed and passed all of the other selection criteria. Figures 5.9, 5.10, and 5.11 present the results of these measurements of PID efficiency for electrons in fully simulated $B\bar{B}$ events generated with conditions matching Run 6, while Figures 5.12, 5.13, and 5.14 present the same measurements for muons. We note again, that large energy-loss bremsstrahlung electrons are removed from these calculations, and do not appear as a source of inefficiency. We only include simulated electrons

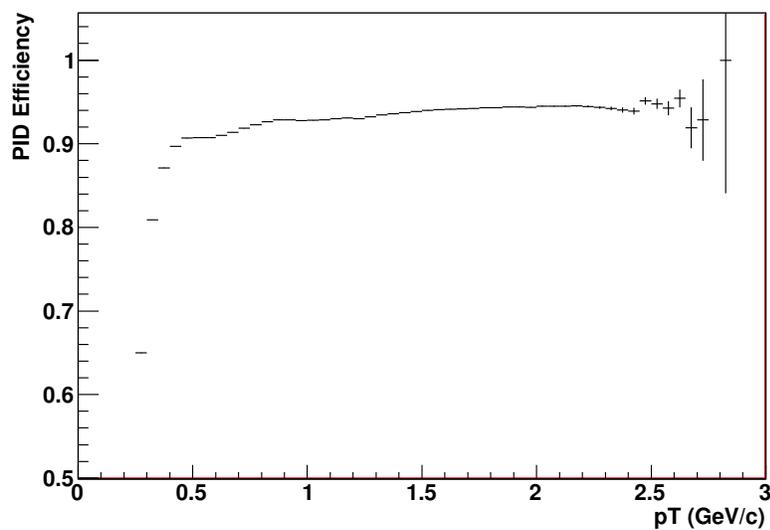


Figure 5.9: The measured electron PID efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.

which have lost less than 5% of their energy to bremsstrahlung in the efficiency calculations, both for “Reconstruction” and PID.

Once these various efficiencies have been measured in fully-simulated B events, we are ready to use them along with the tracking resolution model discussed in Chapter 4 to create the EvtGen portion of the template that will be used in the B mixing dilepton fit, as described in Chapter 6.

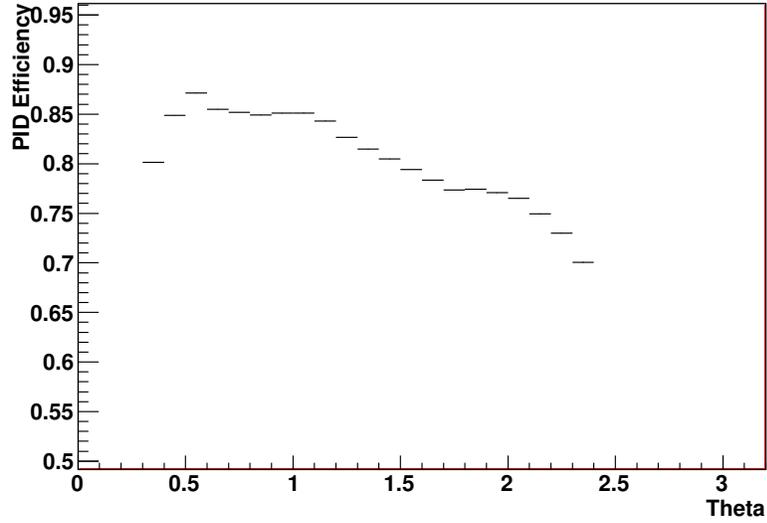


Figure 5.10: The measured electron PID efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.

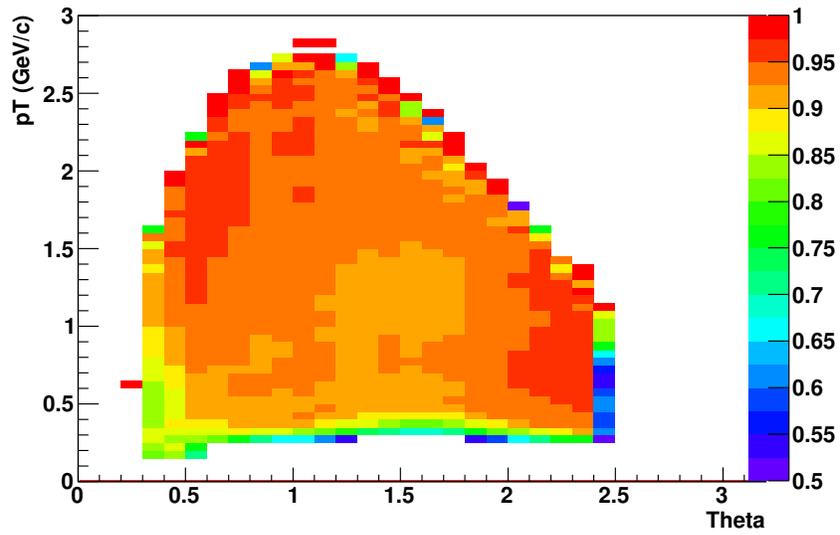


Figure 5.11: The measured electron PID efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.

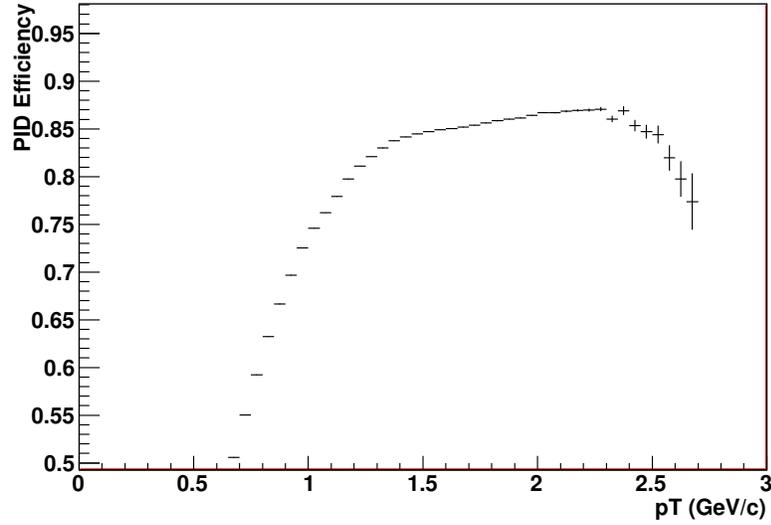


Figure 5.12: The measured muon PID efficiency as a function of track p_T in fully simulated $B\bar{B}$ events.

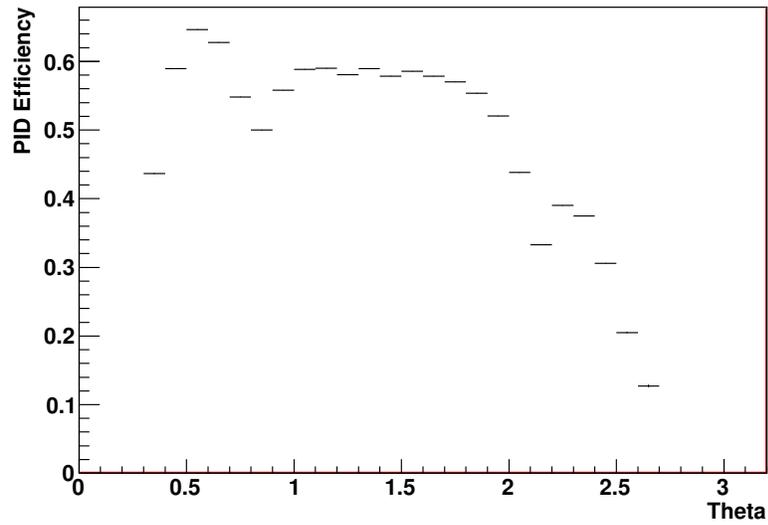


Figure 5.13: The measured muon PID efficiency as a function of track θ (in radians) in fully simulated $B\bar{B}$ events.

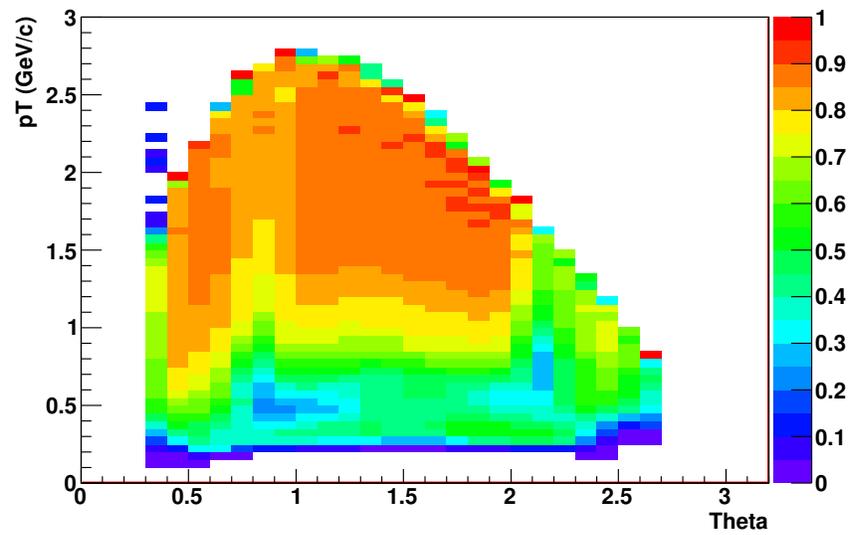


Figure 5.14: The measured muon PID efficiency as a 2D function of both track p_T and θ (in radians) in fully simulated $B\bar{B}$ events.

6

Monte-Carlo Template Preparation

This chapter describes the process of assembling the template used to perform a MCRW fit to *BABAR* data, selected according to the procedures described in Chapter 5. The template is composed of both simulated and real data events, which have enough information stored for each lepton pair to allow us to calculate the binning variables Δt , p_{cm} and lepton species, apply all the selection cuts, or simulate their effect, and re-weight the lepton pairs as the fit parameters are varied. The information needed for re-weighting the lepton pairs includes, for example, the true decay times of simulated B decays, to allow the fitter to calculate the weights as the B lifetime parameters are varied, and the true mixing state of neutral B mesons, to allow re-weighting as the mixing frequency Δm is varied. All of the necessary pieces of information stored for each event in the template are described in the section detailing the event weights used in the fit, in Chapter 7.

There are three major components needed for the template in order to fit either real or simulated *BABAR* data and extract the neutral B meson mixing frequency from dilepton events: the so-called “EvtGen”, “Tertiary Lepton” and “Continuum” parts.

6.1 EvtGen Template

The first and largest part of the template is made up of simulated B meson decays from the EvtGen event generator program, and they are not run through any additional *BABAR* simulation software. This means that the simulations are very fast, and we can produce a very large quantity of events (several times as much as the entire recorded data set) with relatively modest amounts of computing resources, roughly a few hundred CPU hours. In contrast, running a complete simulation of an equal number of events to the data set takes thousands of computers many months to complete. This is the portion of the template that is primarily responsible for measuring the B physics related parameters in the MCRW fit: the mixing frequency Δm , the lifetimes τ^0 and τ^+ , the ratio of charged to neutral B s, and the ratio of cascade to direct leptons. For this reason it is important to have as large a sample as practically possible, in order to decrease the size of the errors due to the template statistics.

We analyze the MC truth information for each lepton from EvtGen, and

only keep those which originate from the following parents in this portion of the template:

- Direct from B decay.
- From longer-lived D meson decay (D^\pm)
- From shorter-lived D meson decay (D^0, D_s)
- From J/ψ decay.
- From τ decay.
- From charmed baryon decay.

leaving the rest of the leptons to be simulated by the Tertiary and Continuum portions of the template. In order to be included as part of the EvtGen portion, we want the leptons to originate within the beam pipe, and to carry some useful information for the fit. We refer to the leptons from the sources listed above as “primary” or “secondary” leptons. Lepton daughters from very long-lived decays, originating outside the beam pipe, don’t meet these criteria, and are included instead in the Tertiary sample, described later in this Chapter.

We use our own simplified, multiple step simulation process to treat the EvtGen template events in a way similar to the much more detailed full *BABAR* simulation, capturing the details which are important for measuring the mixing frequency in dilepton events.

The first selection cuts which need to be applied to the EvtGen events are the requirements on $NTracks$ and R_2 . In reconstructed events, these quantities are based on the list of track candidates called “ChargedTracks”. For each true charged track, we generate a flat random number between 0 and 1, and reject the track if this is less than the value in the ChargedTrack efficiency table at the theta and pt of the track. Because EvtGen does not produce tracks from kaon or lambda decay, photon conversions, or material interactions, the mean value of $NTracks$ in the full simulation is at this point 3.6 tracks larger than the mean value for the efficiency-corrected EvtGen events, so we simply add a Poisson random integer with mean 3.6 to the value of $NTracks$ determined for the EvtGen events. At this point, the means agree quite well, and the widths of the distribution agree to within a few per cent.

The calculation of R_2 is based on this list of tracks, without any modification. The agreement between our EvtGen events smeared with the ChargedTracks efficiency and the fully simulated $B\bar{B}$ events appears to be adequate in this variable, even though the average number of ChargedTracks is significantly different. We store the modified $NTracks$ and R_2 values for the EvtGen template events, and will apply cuts based on them later. For estimates of the systematic uncertainties due to the modelling of $NTracks$ and R_2 in the EvtGen part of the template, please see the discussion of systematic uncertainties in Chapter 8.

We apply a simplified model of Bremsstrahlung emission to the electrons in the EvtGen sample. We smear their momentum and track parameters as

described in Chapter 4, and feed the modified electron tracks into the next step of the tracking resolution model. Each track is independently scattered, and has resolution smearing in both z_0 and d_0 applied according to the procedure described in Chapter 4. The true track origin and direction, and the final smeared track z_0 , d_0 , and direction are recorded, along with enough of the smearing information to re-weight the track to different parameter values of RS and TF.

We divide the efficiency to detect real electrons and muons into two parts for modelling with EvtGen leptons, “Reconstruction” and particle identification (PID). We measured the “Reconstruction” efficiency for both electrons and muons in fully simulated B events, following the procedure described in Chapter 5. In order to apply this “Reconstruction” efficiency to the EvtGen portion of the template, we simply look up the recorded efficiency from a table for the values of θ and p_T of each MC lepton. We then throw a random number between 0 and 1, and if the number is greater than the measured efficiency, we reject the lepton from our EvtGen sample.

We measured the PID efficiency for both the electron and muon selectors in fully simulated B events, following the procedure described in Chapter 5. In order to apply this PID efficiency to the EvtGen portion of the template, we simply look up the recorded efficiency from a table for the values of θ and p_T of each MC lepton. We then throw a random number between 0 and 1, and if the number is greater than the measured efficiency, we reject the lepton from our EvtGen sample.

If we allow all of the $B^0\bar{B}^0$ pairs in the EvtGen sample to mix normally, then we will have definite zeroes in the true proper-time distributions for both the OS and SS distributions. The problem with these zeroes is that events which fall too close to the zero due to resolution smearing will have enormous, or even singular weights due to the mixing part of the weight. To avoid this problem, we simply “dilute” the mixing of the $B^0\bar{B}^0$ pairs in the EvtGen sample, by throwing a random number between 0 and 1 for each $B^0\bar{B}^0$ pair, and if the random number is between 0.0 and the dilution factor (we have chosen to use 2%), then we switch the sign and mixing state of half of the lepton pairs, again selected randomly. In this way, the apparent amplitude of the mixing signal in the PDF is 0.98, instead of 1.00, and 1% of the $B^0\bar{B}^0$ lepton pairs switch signs. We record the true mixing state for $B^0\bar{B}^0$ events in the template, for use in the fit.

We record the true proper decay times of the two B mesons, whether it was a $B^0\bar{B}^0$ or B^+B^- event, the true parents of each lepton, and the other pieces of information noted above, for use in the MCRW fit. After applying this procedure to the EvtGen leptons, all remaining leptons in the EvtGen template at this point are ready for use in the MCRW fit.

Category	N Pairs
Large Energy Loss Bremsstrahlung e	625795
Photon Conversion e	36184
Other Electrons	34506
Other Muons	100426
Fake Electrons	43534
Fake Muons	1061804
Total	1842034

Table 6.1: Number of lepton pairs with at least one lepton in the given tertiary category. Note the total number of tertiary pairs is less than the sum of the categories, since a pair can contain more than one tertiary category lepton.

6.2 Tertiary Lepton Template

There are many other sources of lepton candidates: hadrons misidentified as leptons, unrecognized photon conversions, and real leptons from decays of particles that EvtGen treats as stable. Also, the details of the d_0 shift and the loss of SVT hits correlated to momentum shift from hard Bremsstrahlung are difficult to model inside EvtGen. All of these processes are, however, modelled in detail in the full Moose simulation of *BABAR*. Hence, we create a tertiary template from the full simulation to model these processes. This limits the statistics of this template component to about three times the full *BABAR* data statistics, but tertiary leptons are not a dominant component of the data set. Also, the tracking resolution of the tertiary template is fixed to the simulation production. Note that for tertiary pairs we still have available true proper times, whether the event was B^+B^- or $B^0\bar{B}^0$ and whether a neutral B event is mixed or not, so these events can also be re-weighted to different mixing and lifetime values. It is also possible to re-weight different components of the tertiary lepton background.

The Tertiary template includes the contributions listed in Table 6.1. To produce the tertiary template we process all of the fully-simulated generic B events available, and require at least two leptons in an event which pass all of our selection cuts. We record the lepton pair if either reconstructed lepton candidate doesn't match back to a true primary or secondary lepton (so in fact most "tertiary" pairs have one true primary or secondary lepton).

The large energy-loss Bremsstrahlung electrons are those which have lost more than 10% of their initial energy (known through MC truth) to the emission of a Bremsstrahlung photon. We apply a cut to the ratio of p_T Reconstructed/True, requiring that this ratio lie between 0.9 and 1.0. The process of Bremsstrahlung is highly-material dependent, and the tracks of the electrons can end up pointing in quite a different direction after the photon emission, so we leave this relatively rare process to be modelled by the fully-detailed simulation

software.

Like Bremsstrahlung, photon conversions are also a highly material-dependent process. In fact, one can produce a map of the detector material structure by reconstructing the vertices of converted photons. The daughter electrons and positrons from these conversions therefore provide very little useful information for the dilepton B mixing analysis, and we also leave their modelling to be performed by the fully-detailed simulation software.

Other electrons and muons are those which do not fall into either of the above two Tertiary categories, and don't fall under the list of acceptable parents defined in the section detailing the EvtGen portion of the template above. They are predominantly electrons and muons originating from the decays of pions, kaons, and neutral vector mesons. These long-lived particles travel a large distance from the beam-spot before decaying, so again they don't contribute much useful information about B physics, and we leave them as a Tertiary background to be modelled by the full simulation software.

Finally, a small fraction of identified electrons and muons are actually pions, kaons and protons which have been mistakenly identified (misID'ed) by the detector hardware and software as leptons. This process is relatively rare, as *BABAR* measurements on data control samples indicate the Fake rate for our chosen muon selector to range between 0–3%, for each of π , K , and p , as a function of lab momentum and angle. The same calibration studies show a Fake rate for our chosen electron selector of less than 1% for a wide range of momenta. Simulating this misidentification process requires a very detailed model of the detector elements and the particle ID software, so we leave this Tertiary background to be handled in the full event simulation.

We select all lepton pairs containing one or more of these Tertiary leptons, and passing our other selection criteria, from the centrally produced generic B meson simulation samples with Run 5 and 6 conditions, and store them in an ntuple containing all of the necessary MC truth information used for the EvtGen portion of the template, as well as the reconstructed information used when analyzing real data. We use an equal number of events from Run 5 and 6, 240 M simulated $B\bar{B}$ events from each simulated run period. This is approximately the same number of $B\bar{B}$ events as the final *BABAR* data set. After applying all the selection cuts, we have a total of 1.84 M Tertiary lepton pairs from the MC sample with Run 5 and Run 6 conditions.

We record the tertiary lepton category for re-weighting. Each of these categories of Tertiary leptons will be multiplied by a single, adjustable scaling factor in the MCRW fit. We record both the reconstructed values of $NTracks$ and R_2 , and values calculated in the same manner as done for the EvtGen template above. We study the difference between these two values as a source of systematic uncertainty in Chapter 8. We save the true proper times, whether the event was B^+B^- or $B^0\bar{B}^0$ and whether a neutral B event is mixed or not, so these events can also be re-weighted with different mixing and lifetime values.

6.3 Continuum Template

We make this portion of the template out of off-peak data lepton pairs selected according to the same criteria as applied to the on-peak data, or fully-simulated events. We analyzed the off-peak data, applying the same requirements which would be applied to real *BABAR* data. We required at least two reconstructed, PID leptons. We analyzed all of the off-peak data, from all run periods, and selected 145 K lepton pairs. We recorded all of the reconstructed quantities necessary to place these lepton pairs in the Δt histograms: the leptons p_{cm} , θ , Δt , charges, and species. We added these lepton pairs into our template as a model for the continuum (no $\Upsilon(4S)$) events present in *BABAR* data.

By using the off-peak data for this part of the template, we are guaranteed to have the same mixture of different processes in the background as we would have in the real data. While there is a sample of fully-simulated MC events for *udsc* and τ pairs, in equal quantities to the generic *B* MC, there isn't enough QED MC, especially for two-photon processes, which are notoriously difficult to simulate efficiently. Even if there was an adequate MC sample of these processes, the exact normalization relative to $B\bar{B}$ is problematic because the $\Upsilon(4S)$ production rate depends on the exact beam energy and spread, and not just the luminosity. The *BABAR* experiment ran about 10% of the integrated luminosity below the $\Upsilon(4S)$ threshold, and the relative luminosity is known quite accurately, and all processes are automatically included in this data. We don't need to worry about the relative cross-sections of the *udsc*, e^+e^- , $\mu^+\mu^-$ and $\tau^+\tau^-$ processes, as we would if we used simulated continuum events. The main source of uncertainty is simply related to the ratio of on-peak and off-peak data luminosities, which has been measured quite well by my *BABAR* colleagues. We therefore use a correspondingly simple template model for the off-peak component, with just a single adjustable scaling factor in the MCRW fit.

6.4 Template Event Counts

We used EvtGen to generate 476M *B*-pairs in EvtGen with the Run 5 beam conditions and efficiency tables, and 357M *B*-pairs with the Run 6 beam conditions and efficiency tables, resulting in 11.79 M lepton pairs in the EvtGen portion of the template. We used 240M fully simulated *B*-pairs simulated with the Run 5 beam and detector conditions and 240M fully simulated *B*-pairs with the Run 6 beam and detector conditions, resulting in 1.84M lepton pairs in the tertiary portion of the template. We used all of the *BABAR* off-peak real data to add about 145,000 lepton pairs to the continuum portion of the template. The total size of the template is 13.75 M lepton pairs. The relative normalizations of the components are handled by re-weighting the template, at the same time as the fitting for physics and resolution parameters.

7

The MCRW Fit

In this Chapter we describe the particular details used to perform a MCRW fit to fully simulated *BABAR* events, using a template constructed according to the procedure described in Chapter 6. The first step is to calculate the decay-time difference Δt for the lepton pairs, then transform it into a variable which is more convenient for use in fitting and plotting. After this we create a binned histogram, separating the lepton pairs into bins in the transformed decay-time variable, bins of momentum, splitting the pairs into Opposite-Sign (OS) or Same-Sign (SS), and separating the pairs by lepton species.

We then recap briefly the weighting function which is applied to the template events in order to allow the fit to modify and determine the desired parameters, and conclude by presenting the results of a fit to $\sim 480\text{M}$ fully simulated $B\bar{B}$ events.

7.1 Simulated Data Sample and Binning

We perform a full-scale test of our fitting scheme with simulated $B\bar{B}$ events, and off-peak real data in the place of the continuum component. Our data sample is drawn from roughly 480M fully-simulated B -pairs, representing a large fraction of the simulated $B\bar{B}$ events with Run 5 and 6 conditions. We use equal amounts of simulated events, (240M each) from the Run 5 and Run 6 collections. This total sample size is similar to the total *BABAR* recorded sample of approximately 471M B -pairs [20]. We require the events to contain at least five charged tracks, to have charged track $R_2 < 0.8$, and two identified leptons satisfying: $0.8 < p_{cm} < 2.8 \text{ GeV}/c$, and $0.7 < \theta < 2.3$, along with the other track selection criteria described in Chapter 5. After imposing these requirements, we have 8.87 M lepton pairs resulting from the fully simulated generic B decays. The off-peak real data events are required to pass the same cuts, and to roughly scale to the same luminosity, each off-peak lepton pair is included ten times, for a total of 1.45 M “continuum” pairs. We end up with approximately 10.29 M lepton pairs in the simulated data sample.

We use the measured decay time difference Δt for each selected lepton pair as determined by the procedure detailed in Section 4.1. Before binning the lepton pairs from the data and template histograms for the MCRW fit, we first transform the decay time variable from Δt to a new variable, $\Delta t'$ by the following:

$$\Delta t' = \pm\tau' \times (1 \mp e^{-\Delta t/\tau'}) \quad (7.1)$$

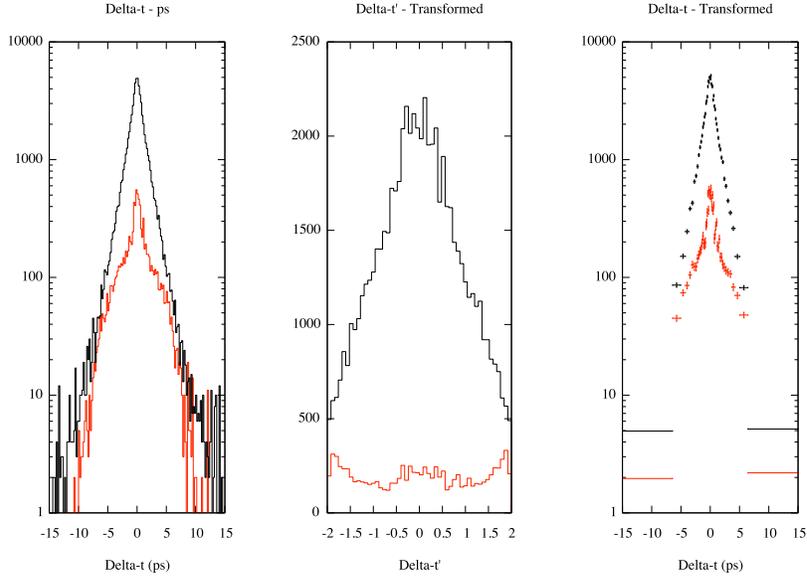


Figure 7.1: A sample transformation from the Δt (left) to $\Delta t'$ (middle) variables, for the same group of fully simulated $B\bar{B}$ muon pairs. The transformed histogram is plotted versus Δt on the right, with the variable bin widths resulting from the transformation. These plots show only the High-High bin, with both muons having $p_{cm} > 1.8$ GeV/c. The lower (red) histogram is the SS pairs, the upper (black) histogram is OS pairs.

The transformation uses \pm to treat positive and negative Δt values in a symmetric way. This transformation increases the emphasis on the core of the distribution, at low values of Δt , and decreases the importance of the far tails, by transforming the range from 0 to ∞ in Δt into the range from 0 to τ' in $\Delta t'$. We choose the value $\tau' = 2.0$ ps, which is slightly larger than the B lifetimes, so that the exponential decay is mostly flattened out, but not entirely. This transformation is applied to both the data and the template, and example of its application is shown in Figure 7.1.

The MCRW fit is binned in several different dimensions in order to improve the precision of the Δm measurement, and to help the fit discriminate between lepton pairs from different sources. We use the following binning scheme:

- We bin the variable $\Delta t'$ in 50 bins between -2.0 and +2.0, for a bin width of 0.08 in $\Delta t'$. These bins have variable widths in the original Δt .
- We split the pairs into OS and SS categories. In the case of both leptons directly from B meson decay, and from cascade decays, we expect

distinctly different Δt shapes for OS and SS pairs.

- We use six p_{cm} bins, low-low, medium-low, medium-medium, high-low, high-medium, and high-high. The low p_{cm} bin ranges from 0.8 to 1.3 GeV/c, the medium bin ranges from 1.3 to 1.8 GeV/c and the high bin is $p_{cm} > 1.8$ GeV/c. This binning helps distinguish between the direct leptons, which are mostly in the highest bin, and cascade leptons, which are mostly in the lowest bin.
- We split the pairs into four dilepton categories: ee , $e\mu$, μe and $\mu\mu$. This split helps to separate efficiency issues for the two types of leptons and their particle identification algorithms.

In total there are $N = 50 \times 2 \times 6 \times 4 = 2400$ bins in the histograms. We have separate weight factors for normalizing the momentum spectra for electrons and muons, separately in the three p_{cm} bins, Ne_0 , Ne_1 , Ne_2 , $N\mu_0$, $N\mu_1$, and $N\mu_2$. The same binning scheme must be used for both the data and template.

7.2 Fit Parameters

The ratio of the Δt (or $\Delta t'$) distribution for same-sign lepton pairs to the distribution for the sum of same-sign and opposite-sign pairs contains the information about Δm , the frequency of neutral B mixing, which is our primary interest. The data set and template were generated with $\Delta m = 0.489 \text{ ps}^{-1}$.

The distribution for the sum of same-sign and opposite-sign pairs obviously also gives information about the B^0 and B^+ lifetimes. The two lifetimes are nearly equal, and B^0 and B^+ mesons are produced in nearly equal numbers. If we fit for the two lifetimes independently, they have highly correlated errors, and the errors are highly correlated to Δm .

The following identity is instructive:

$$e^{-t/\tau_1} + e^{-t/\tau_2} = e^{-t/2[1/\tau_1+1/\tau_2]} \cdot \left[e^{-t/2[1/\tau_1-1/\tau_2]} + e^{+t/2[1/\tau_1-1/\tau_2]} \right] \quad (7.2)$$

This implies that an equal mix of exponentials with two slightly different lifetimes is equivalent to an exponential with the average of the inverse lifetimes, modulated by a function of the difference of the inverse lifetimes. The modulation function is nearly flat if the lifetimes are nearly equal. If we change the fit parameters, we can determine the sum of inverse lifetimes with a small error, and the difference of inverse lifetimes with a large error, and these errors are nearly uncorrelated.

Another instructive identity is:

$$\frac{e^{-t/\tau_1}(1 \pm \cos(\Delta mt))}{e^{-t/\tau_1} + e^{-t/\tau_2}} = \frac{1 \pm \cos(\Delta mt)}{1 + e^{+t[1/\tau_1-1/\tau_2]}} \quad (7.3)$$

This implies that the ratio of same-sign dileptons from B^0 to total dileptons from an equal mix of B^+ and B^0 , which is essentially our signal for Δm , depends

on the B^+ and B^0 lifetimes principally through the difference of the inverse lifetimes, and is essentially independent of the sum of the inverse lifetimes.

We cannot measure the difference of inverse lifetimes very well, so we fix it to the EvtGen default value of 0.051703 ps^{-1} . Changing this value does make a significant systematic shift in our result for Δm , which is discussed in Chapter 8. We do float the sum of inverse lifetimes in the fit.

The precise ratio of B^+ to B^0 in $\Upsilon(4S)$ decay is not known with high precision, so we include a ChargeFactor parameter that multiplies the B^+B^- event weights in the template. Our test data set and template assume the same, equal production rate, so we expect the value of the ChargeFactor = 1.0.

The number and spectrum of secondary or cascade leptons relative to primary or direct leptons is not precisely known, so we include a CascadeFactor parameter that multiplies the weights of secondary leptons in the EvtGen template. A double-cascade lepton pair would have CascadeFactor applied twice. Since the test data and template have precisely matching lepton spectra, we expect the value of the CascadeFactor = 1.0.

We expect real data to have tracking resolution modestly different from the parameterization used to generate the template. The template can be re-weighted to a different resolution by the fit. The multiple scattering core and tail, and the geometrical resolution, are scaled by a common factor RS, separately for electrons and muons. We expect these parameters from fitting $\Delta t'$ to be the values found when fitting reconstructed minus true z_0 difference and d_0 difference distributions in the full BABAR Monte Carlo simulation.

Finally, the electron and muon spectrum and identification efficiency vs momentum may be modestly different in real data from those in the template, so we have a normalization factor for each of the three electron p_{cm} bins and three muon p_{cm} bins. We expect these parameters to come out to 1.0 since the test data set should match the template.

7.3 Template Event Weights

The weight of an event in the EvtGen portion of the template is given by

$$w_{\text{EvtGen}} = w_{\text{nEvtGen}} \cdot w_{\text{lifetime}} \cdot w_{\text{mix}} \cdot w_{\text{chg}} \cdot w_{\text{casc1}} \cdot w_{\text{casc2}} \cdot w_{\text{track1}} \cdot w_{\text{track2}} \quad (7.4)$$

The w_{nEvtGen} value was fixed at 0.596, to make the normalization of primary and secondary lepton pairs consistent with the template. The weight is assigned by counting the number of direct and secondary lepton pairs in the MC-as-data, and dividing by the number in the template. This weight decreases as the size of the EvtGen template increases in comparison to the data. The next term in the event weight, w_{lifetime} is used to encode the probabilities of the charged and neutral B meson lifetimes, τ^+ and τ^0 . We calculate τ^+ and τ^0 from the actual fitted parameters SumInvTau and (fixed) DiffInvTau. For each B meson, the proper decay time t , is used in the following way:

$$P_{\text{lifetime}}(t, \tau) = e^{-t/\tau} / \tau$$

$$w_{\text{lifetime}} = \frac{P_{\text{lifetime}}(t_1, \tau_{\text{fit}})P_{\text{lifetime}}(t_2, \tau_{\text{fit}})}{P_{\text{lifetime}}(t_1, \tau_{\text{gen}})P_{\text{lifetime}}(t_2, \tau_{\text{gen}})} \quad (7.5)$$

where P is the probability function for exponential decay, and t_1 and t_2 are the true decay times of the two B mesons in the event. We use the appropriate value of τ for either B^+B^- or $B^0\bar{B}^0$ events.

The next term in the event weight is for the mixing frequency Δm , and depends on the difference of the true B meson proper decay times, as well as the mixing amplitude D . The mixing weight is given by:

$$\begin{aligned} P_{\text{mix}}(t_1, t_2, \Delta m, D) &= 1.0 + gD \cos[\Delta m(t_1 - t_2)] \\ w_{\text{mix}} &= \frac{P_{\text{mix}}(t_1, t_2, \Delta m_{\text{fit}}, D = 1)}{P_{\text{mix}}(t_1, t_2, \Delta m_{\text{gen}}, D = 0.98)} \end{aligned} \quad (7.6)$$

where g is equal to $+1.0$ for unmixed $B^0\bar{B}^0$ events (*i.e.* $B^0\bar{B}^0$ or \bar{B}^0B^0), -1.0 for mixed $B^0\bar{B}^0$ events (*ie* B^0B^0 or $\bar{B}^0\bar{B}^0$), and 0.0 for all other types of events in the template. We consider the amplitude to be 1.0 in data, and dilute the mixing with an amplitude of $D = 0.98$ in the template. We apply an amplitude of less than 1.0 to the generated template mixing signal, in order to avoid the possibility of having any weight precisely equal to zero, which would lead to numerical problems, as the ratio of the initial and final weights is then undefined.

We also apply the following weight for the ratio of charged and neutral B mesons in the sample:

$$\begin{aligned} w_{\text{chg}} &= \text{ChargeFactor} && \text{if event is } \Upsilon(4S) \rightarrow B^+B^- \\ &= 1 && \text{if event is } \Upsilon(4S) \rightarrow B^0\bar{B}^0 \end{aligned} \quad (7.7)$$

and the following for the ratio of direct to cascade leptons:

$$\begin{aligned} w_{\text{casc1,2}} &= \text{CascadeFactor} && \text{if lepton is not directly from } B \text{ decay} \\ &= 1 && \text{if lepton is directly from } B \text{ decay} \end{aligned} \quad (7.8)$$

The tracking weight, w_{track} is calculated as described in Equations 4.9 – 4.13, independently for each track in the EvtGen template.

We only float two of the track resolution parameters in the fits to the Δt distribution, a Resolution-Scale (RS) for electrons, and one for muons. The other parameters (TailFraction, 2nd gaussian Fraction, 2nd Gaussian width) are fixed to the values obtained in the single track fits to the z_0 resolution in generic B fully simulated events.

The weight of an event in the tertiary portion of the template is given by

$$w_{\text{tertiary}} = w_{\text{lifetime}} \cdot w_{\text{mix}} \cdot w_{\text{chg}} \cdot w_{\text{category}} \quad (7.9)$$

where the first three factors are the same as for the EvtGen portion, and w_{category} allows different sources of the tertiary leptons to be varied. The categories are:

- Large Energy Loss Bremsstrahlung

- Gamma Conversion Daughters
- Other Electrons
- Other Muons
- Fake Electrons
- Fake Muons

As long as the systematics associated with each of these categories is small, we don't need any more detailed models, than just weighting their contribution. We rely on GEANT to correctly model their contribution to the Δt distribution, while we can vary the overall amount of each category in the fit. We fix the weight factor for each of these categories to be equal to 1.0 in our current fits, because the same tertiary leptons are included in both the data and template. We vary the weights of each category to establish systematic uncertainties in Chapter 8.

We have several parameters which we use to adjust the relative p_{cm} spectrum of the template in comparison with the data. We apply the factors: Ne0, Ne1, Ne2 to EvtGen or Tertiary electrons in the low, medium and high p_{cm} bins. We apply Nmu0, Nmu1, Nmu2 to EvtGen or Tertiary muons in the same way. Floating these parameters allows the fit to account for differences in the particle spectra (in our case, these must be due to the efficiency modelling process) between the data and template. We expect these parameters to all be equal to 1.0 in our fit.

The weight of an event in the continuum portion of the template is simply w_{nonBB} . This was fixed to 10.0, consistent with the fact that each off-peak event in the template was included exactly ten times in the “data” and only once in the template.

7.4 Fit Results

We float the following parameters in the fit: SumInvTau, Δm , ChargeFactor, CascadeFactor, RS-e, RS-mu, Ne0,1,2 and Nmu0,1,2. The fit takes approximately 50 minutes to run on a single CPU machine, including the time spent loading the ntuples over the network. This process of reading in the ntuples takes approximately 3–4 minutes, and the rest of the time is spent performing five iterations of the fit. The final fit results are given in Table 7.1, showing both the fitted and MC Truth values for each parameter. The pulls are generally within the expected range, except perhaps for Ne2, which may indicate a small problem in our template description of high-energy electrons, or may be just a random fluctuation.

We present overlay plots showing the data (red points) and fitted template (black histogram) for all the bins in the fit. The error bars shown on the data points are the quadrature sum of the statistical error of the data, the statistical

Parameter	MC Truth	Fit Value	Fit Error	Pull (σ)
Δm	0.4890	0.4886	0.0026	-0.16
SumInvTau	1.2460	1.2435	0.0014	-1.88
ChargeFactor	1.0000	0.9879	0.0104	-1.16
CascadeFactor	1.0000	1.0077	0.0042	+1.83
RS-e	0.8401	0.8328	0.0054	-1.34
RS-mu	0.8549	0.8542	0.0063	-0.11
Ne0	1.0000	1.0000	0.0033	+0.0034
Ne1	1.0000	0.9945	0.0030	-1.82
Ne2	1.0000	0.9885	0.0033	-3.51
Nmu0	1.0000	1.0080	0.0032	+2.52
Nmu1	1.0000	1.0063	0.0031	+2.05
Nmu2	1.0000	1.0053	0.0033	+1.61

Table 7.1: Results of a MCRW fit to 10M simulated lepton pairs in $\Delta t'$ with a template of ~ 14 M lepton pairs. See the text for more details.

error of the template, and the error of the weight in each bin. Figures 7.2 through 7.5 show the bins containing ee , $e\mu$, μe and $\mu\mu$ pairs, respectively.

The statistical error of Δm in this fit is 0.0026, or 0.53% for the generator true value of 0.489 ps^{-1} . This is noticeably smaller than the statistical error achieved in published measurements of Δm , with an error of 1% on the world-average value (see Appendix B). Our final fit contains 12 floating parameters, and approximately 10 fixed parameters. The previous *BABAR* measurement using more traditional methods floated 25 parameters in the final fit, with 488 fixed.

It is straightforward to increase the size of the EvtGen portion of the template in order to decrease the associated component of the statistical error, up to the limits imposed by the available RAM of the computer which performs the fit. If we were able to increase the statistics of the template to an infinite number of events, we would obtain a statistical error of 0.0021, rather than 0.0026, based only the size of the data. In the near future it should be possible to access a sample of fully-simulated $B\bar{B}$ events which is three times larger, to increase the size of the Tertiary template, which would reduce the total statistical error somewhat towards this limit. It is not possible to increase the size of the off-peak data set, (used for continuum Template) now that the *BABAR* experiment has ended. We don't believe, however, that this is a large contribution to the statistical error.

The final fit χ^2 is 1065 for 2388 Degrees-of-Freedom. This value is misleadingly low, because the exact same tertiary pairs and off-peak data pairs are used in both the data and template, so there are far less statistical fluctuations than one would expect.

At first glance one might naively assume that the presence of the same Tertiary and continuum lepton pairs in both the data and the template will affect

the parameter estimates or their errors. On close inspection of the Equations in Chapter 2, however, we note that the residuals between data and template don't enter directly into the matrix W_{jk} which is used to estimate the parameter values and their errors. As a simple way of verifying that the replicated events only modify the fit χ^2 and don't affect the parameter estimates or their errors, we perform the following test. We generate Poisson random numbers with mean 1.0 when reading each of the tertiary and continuum lepton pairs into the fitter, and add the lepton pair to either the data or template the corresponding random number of times. This procedure effectively erases the correlations between the data and template samples.

After this modification, we find that the fit χ^2/DOF increases from 1065/2388 to 2475/2388. For a χ^2 statistic with $\text{DOF} = 2388$, the standard deviation is $\sigma = \sqrt{2 \times \text{DOF}} = 69$. Our modified fit is therefore within two σ of a perfect fit. The fit error on Δm changes from 0.002613 to 0.002612, which is a negligible difference. The estimated value of Δm changes by an amount that is much less than one standard deviation (approximately 0.01σ). This test demonstrates that, as expected, the presence of the same template and continuum lepton pairs in both the data and the template does not affect the parameter estimates or errors. Finally, we note that this issue of replicated events would no longer be an issue when a fit was performed to real experimental data, and the template could be used directly as we constructed it.

Equally important to the statistical error in an analysis of real data is the sum of systematic uncertainties that must be considered, and they are examined in the following Chapter.

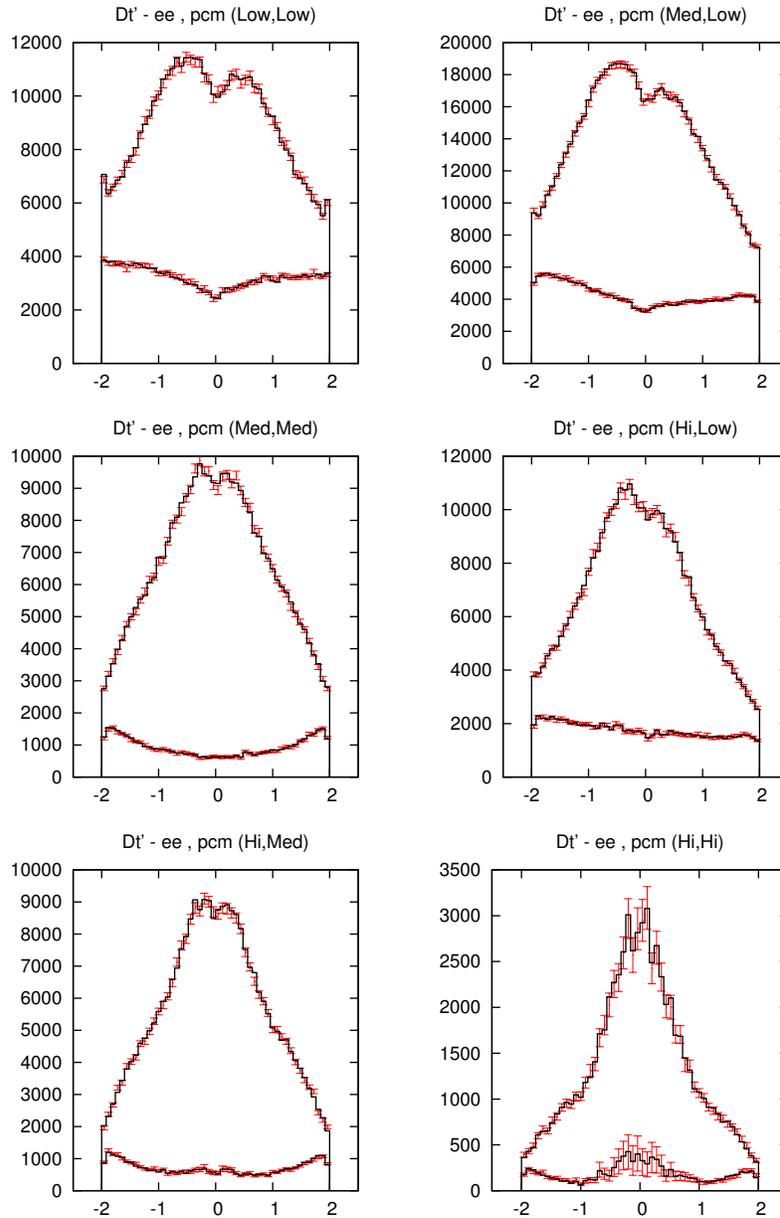


Figure 7.2: The results of a fit to fully simulated ee pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.

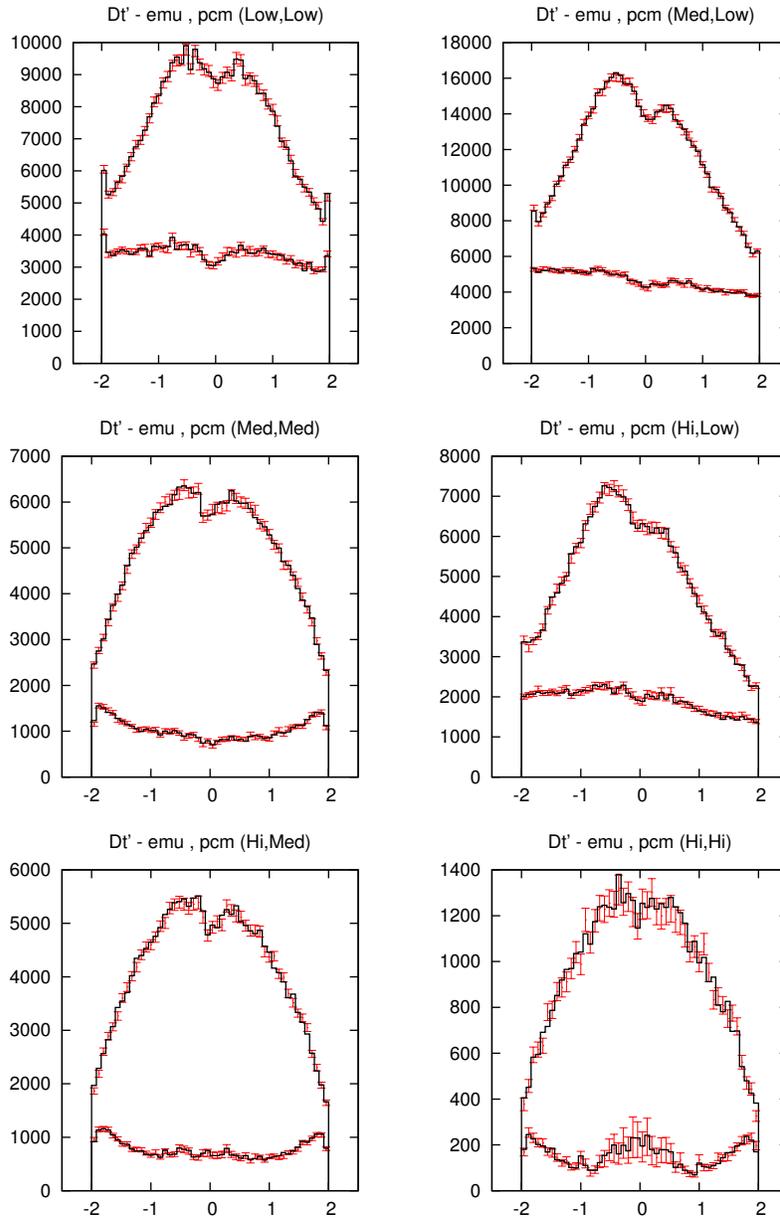


Figure 7.3: The results of a fit to fully simulated $e\mu$ pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.

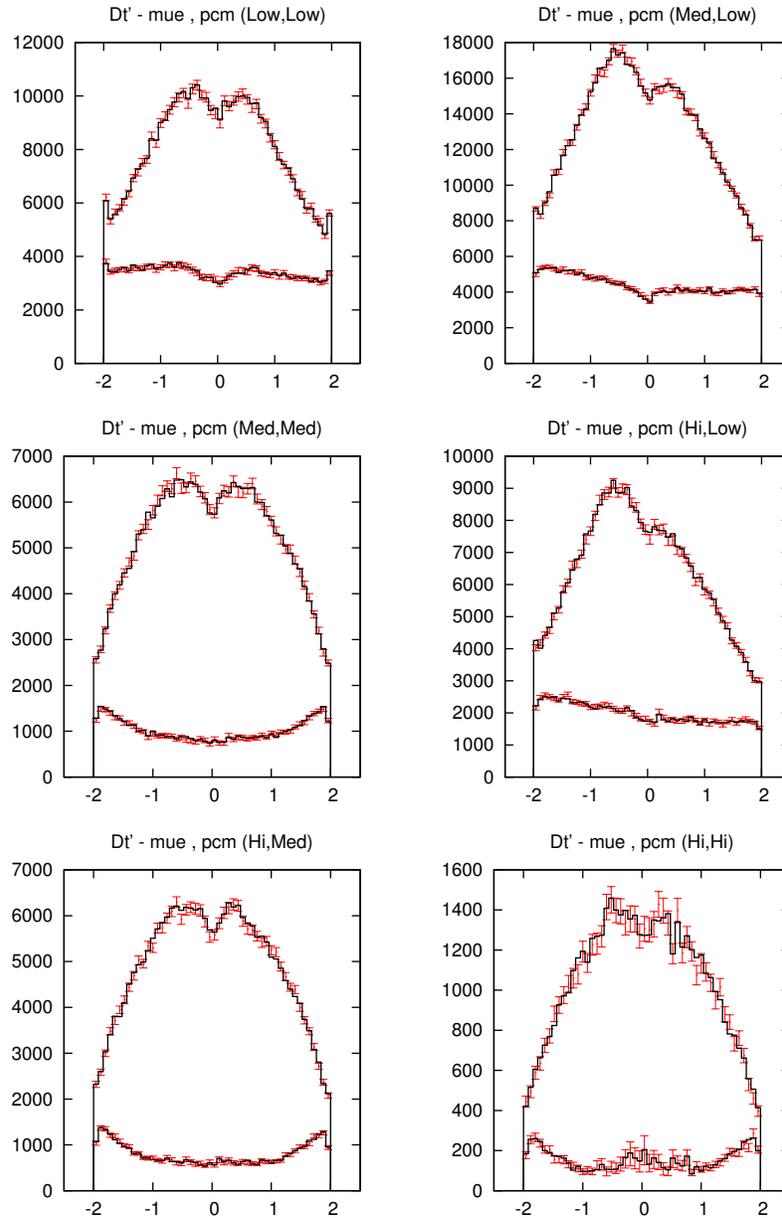


Figure 7.4: The results of a fit to fully simulated μe pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.

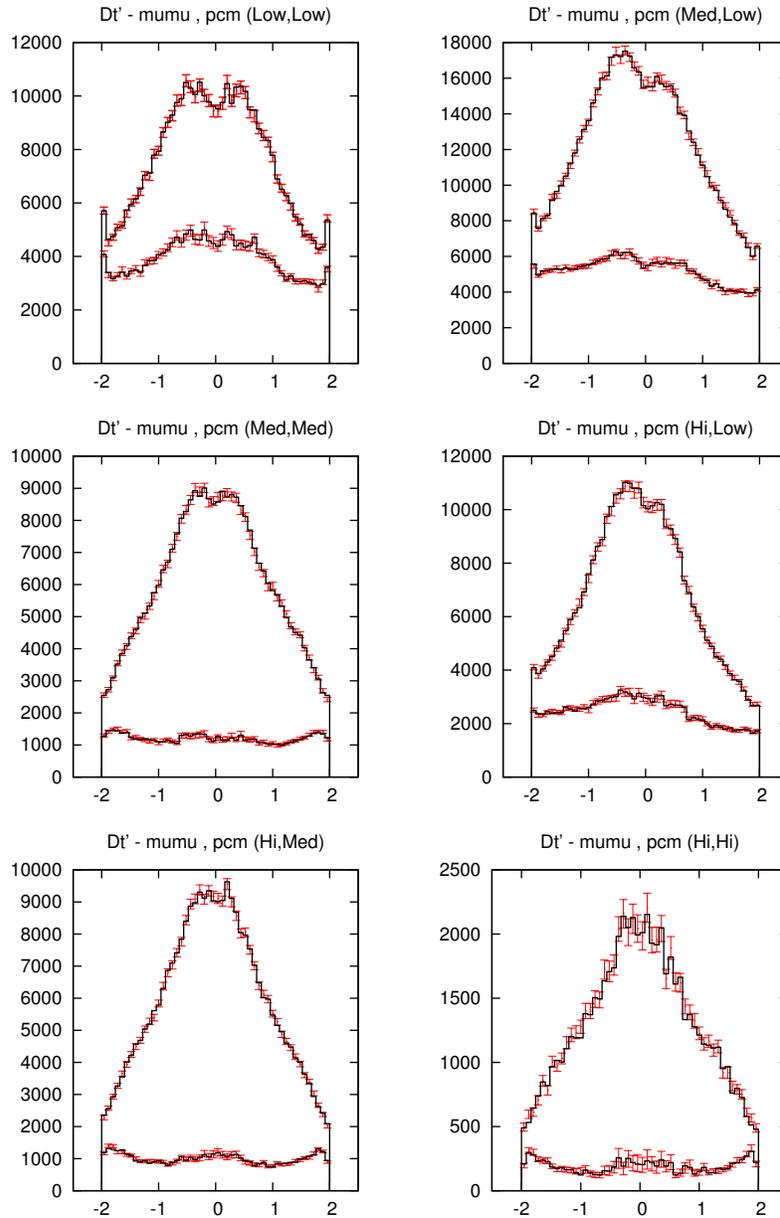


Figure 7.5: The results of a fit to fully simulated $\mu\mu$ pairs in $\Delta t'$. OS pairs are the upper curve in all plots, SS pairs are lower. Data is red points with errors, black histogram is weighted template.

8

Systematics Studies

This chapter details our estimations of the sensitivity to systematic uncertainties in our measurement method for Δm in neutral B mixing. Essentially, these correspond to differences between the fitted data distribution and the template that cannot be removed by adjusting the fitted parameters. In our case, the fitted data distribution is composed of detailed *BABAR* simulation events, so the systematics will be estimated by looking at differences between our composite template and the detailed simulation. We use large changes in the fixed parameters (generally $\pm 10\%$, except for the B lifetimes) to examine how sensitive the fitted result will be. These numbers do not represent a precise estimate of systematic uncertainties in a completed analysis. We have examined what we believe to be the dominant sources of systematic uncertainty that could be expected in a fit to real *BABAR* data, and they are discussed below.

8.1 B Lifetimes

We cannot easily distinguish the lifetimes of the B^0 and the B^+ in our fit to the Δt distribution, and so we must fix one lifetime parameter to the accepted world-average value [13], while we float the other in our fit. We examined several different parameterizations of the two lifetimes, in order to find the best one for minimizing the fit errors, and the contribution of the uncertainty due to a fixed lifetime parameter on the Δm measurement. The current (as of April, 2011) world-average lifetime ratio (τ^0/τ^+) is given with 0.9% uncertainty. The current single best measurement of τ^0 was performed by the BELLE collaboration [19], with a value of $\tau^0 = 1.534 \pm 0.008 \pm 0.010$ ps. The current best measurement of τ^+ is a result in the same paper, with a value of $\tau^+ = 1.635 \pm 0.011 \pm 0.011$ ps. These experimental uncertainties translate into possible shifts in the fixed lifetime parameter, with the other parameter left floating.

We considered four different methods of parameterizing the two B meson lifetimes, the lifetime of the B^0 (τ^0) and the lifetime of the B^+ (τ^+). The first parameterization is the simplest, just using the two lifetimes directly in the fit. We find that unless the fit starts extremely close to the final values, it has a hard time distinguishing between the two lifetimes, and it is difficult to converge reliably. This is generically true, for any fitting algorithm, not just MCRW. So we fix one of the two lifetime parameters, and allow the other to float. We vary the fixed parameter by an amount equivalent to varying τ^0 or τ^+ by approximately 1%, to estimate a potential systematic uncertainty, and

Fixed Parameter	Shift Amount	Δm Shift	Δm Shift %
SumInvTau	$\pm 1\%$	2.19×10^{-3}	0.45
DiffInvTau	$\pm 1\%$	2.56×10^{-3}	0.5
TauAvg	$\pm 0.5\%$	7.17×10^{-3}	1.5
TauDiff	$\pm 1\%$	3.41×10^{-3}	0.7
TauGeom	$\pm 1\%$	1.28×10^{-2}	2.6
TauRatio	$\pm 1\%$	3.69×10^{-3}	0.8
τ^0	$\pm 1\%$	6.50×10^{-3}	1.3
τ^+	$\pm 1\%$	6.36×10^{-3}	1.3

Table 8.1: Preliminary estimates of the Δm systematic uncertainties due to fixing one of the two values parameterizing the B lifetimes in the fit. The percent error is based on an assumed value of $\Delta m = 0.489\text{ps}^{-1}$.

take half the difference between the two fit results.

The second parameterization is in terms of the arithmetic average and difference of the two lifetimes, TauAvg, and TauDiff. When TauAvg is fixed, we vary the difference by a conservative 1% of the average lifetime, both up and down. When the difference is fixed, we vary the average lifetime up and down by 0.5%.

The third parameterization is in terms of the geometric average TauGeom ($\sqrt{\tau^0\tau^+}$), and the ratio of the lifetimes TauRatio. When the geometric mean is fixed, we vary the ratio up and down by 1%, and when the ratio is fixed, we vary the geometric mean up and down by 1%.

The fourth parameterization is based upon the form of the weight factors for the two lifetimes, and the parameters used are SumInvTau = $1/\tau^0 + 1/\tau^+$, and DiffInvTau = $1/\tau^0 - 1/\tau^+$. This parameterization should have nearly orthogonal errors in the two parameters, because the construction matches that of the weight factors in the fit. When SumInvTau is fixed, we vary DiffInvTau up and down by an amount equivalent to changing the lifetimes by 1%. When DiffInvTau is fixed, we vary SumInvTau by 1%.

The four different estimated systematic uncertainties on Δm due to knowledge of B lifetimes are given in Table 8.1. Based on the information in Table 8.1, we conclude that the best parameterization for us to use is the one with SumInvTau and DiffInvTau, keeping DiffInvTau fixed in the fit, and floating SumInvTau. In this case, we expect a systematic uncertainty on Δm of roughly 0.5%, or $\pm 0.0024 \text{ps}^{-1}$, based on the MC generator value of 0.489ps^{-1} . This is comparable to the expected statistical error, seen at the end of Chapter 7. It is possible that in the near future, results from the LHC experiments may provide more precise measurements of the average or ratio of B meson lifetimes, allowing for a decrease in this contribution to the systematic uncertainty.

Category	Shift Amount	Δm Shift	% Δm Shift
D^\pm	$\pm 10\%$	1.05×10^{-3}	0.215
J/ψ	$\pm 10\%$	6.6×10^{-4}	0.135
D^0, D_s	$\pm 10\%$	5.8×10^{-4}	0.119
τ	$\pm 10\%$	4.0×10^{-4}	0.082
A_c	$\pm 10\%$	1.7×10^{-4}	0.035
Quadrature Sum		1.44×10^{-3}	0.294

Table 8.2: Preliminary estimates of the systematic uncertainties due to differences in the relative amount of cascade leptons between simulated and real B decays.

8.2 Cascade Leptons

This is another of the systematics associated with the EvtGen part of the template, and it corresponds roughly to differences between the MC branching fractions and the spectra for B meson decays in the data and template of the fit. When fitting real data, these would be differences between reality and the *BABAR* simulations, but in our fit of simulation to simulation, we must manually adjust these categories to create a mismatch.

Our fit has floating p_{cm} bin normalization variables ($N_{e0,1,2}$ and $N_{m0,1,2}$), that will absorb some of the spectral differences. The influence of changes in the sources of the leptons can have a limited effect on the Δt distributions, as compared to a mass measurement or branching fraction measurement, which would be more sensitive. We apply weight factors to various categories of non-direct leptons in the fit, and see how much Δm and the average lifetime shifts, in order to estimate the systematic uncertainties. These fits were performed with the lifetime model using the sum and difference of inverse lifetimes (SumInvTau and DiffInvTau). The statistics corresponded to roughly half of the full data statistics (200M simulated B pairs), and the event selection criteria in Chapter 5 were used, although the continuum suppression cuts were relaxed slightly to $NTracks \geq 4$, and $R_2 \leq 0.8$, and there were no off-peak data or continuum events included in the fit tests. It should also be mentioned here that the variance of 10% used to estimate the systematics was not carefully considered, and is likely to be an overly conservative estimate of the possible discrepancies between the *BABAR* simulations and reality, at least for some of the decays in the list.

We take the absolute value of the largest of the two shifts as our systematic. Based on the results in Table 8.2, it appears that our Δm measurement is most sensitive to relative differences between the data and template in the proportion of A_c , D^0 , and D_s , in the cascade decays. If this systematic is large enough to cause concern, it would be straightforward to introduce additional parameters to the fit to allow one or both of these contributions to float, although the stability and reliability of introducing such a parameter has not been tested.

Parameter	Shift Amount	Δm Shift	% Δm Shift
Geometric Resolution	$\times 0.5, \times 2.0$	9.4e-4	0.19
TailFraction	$\pm 10\%$	5.1×10^{-4}	0.10
Core σ Noise	$\pm 10\%$	2.6×10^{-4}	0.05
2nd Gaussian σ	$\pm 10\%$	2.0×10^{-4}	0.04
2nd Gaussian Fraction	$\pm 10\%$	7.0×10^{-5}	0.01
Quadrature Sum		1.1×10^{-3}	0.23

Table 8.3: Preliminary estimates of the systematic uncertainties due to the fixed parameters in our track resolution model.

8.3 Track Resolution Model

In addition to the Resolution-Scale, or ResScale parameters which are floated separately for electrons and muons, there are several resolution parameters which are fixed, and were set based on studying MC events, by fitting to single-track z_0 differences between reconstructed and truth values in Chapter 4. We varied those fixed parameters by a reasonable amount ($\pm 10\%$ in most cases, except for the geometric resolution, which was multiplied by 0.5 and by 2.0) and determined how much the fitted values of Δm shifted, in order to estimate the systematic uncertainties. We took the difference between the +10% and -10% cases, and divided in half to obtain the shifts in Δm . To evaluate these systematic uncertainties, the fits were done with flattened histograms, using the transformed variable: $\Delta t'$. The fits included the normal amount of Tertiary leptons in the both the data and template ntuples, and ten copies of the off-peak data events as continuum in the data histogram, and one copy of the off-peak data, with a weight of 10.0, in the template.

We used 480 million fully-simulated B pairs as data, roughly equivalent to the statistics expected for the final fit to real data. We used roughly twice as many EvtGen events in the template, compared to the data, for these fits. We applied all of the selection criteria as described in Chapters 5 and 7. The size of the systematics for each fixed resolution parameter are given in Table 8.3. We see that the largest resolution model systematic is due to the variation of the geometric resolution by a factor of one half or two. This large range was chosen in a very conservative manner, and is likely an overestimate of the uncertainty in the core geometric resolution. In the future, a careful comparison of tracking resolution between data control samples and similar simulated control samples could be used to estimate these uncertainties more carefully.

8.4 Tertiary Leptons

The systematic uncertainties due to the Tertiary leptons correspond to differences between the MC detector material model and simulated performance and reality, when one fits real data. In our case where we are fitting simulated events,

Category	Shift Amount	Δm Shift	% Δm Shift
Fake Muons	10%	9.8×10^{-4}	0.200
Large Brem Electrons	10%	6.6×10^{-4}	0.135
Other Muons	10%	2.6×10^{-4}	0.053
γ Conversions	10%	1.7×10^{-4}	0.035
Other Electrons	10%	1.0×10^{-4}	0.020
Fake Electrons	10%	4.0×10^{-5}	0.008
Quadrature Sum		0.0012	0.251

Table 8.4: Preliminary estimates of the systematic uncertainties due to the tertiary leptons, with the finalized cuts applied.

we can vary them, and see how much the fitted Δm shifts, in order to estimate systematics, similar to the procedure used for the Cascade leptons. The differences in our fit correspond to differences between the GEANT simulation and real data, coming from the fully-simulated B event part of the template.

We used 480 million fully-simulated B -pairs as data, roughly equivalent to the statistics expected for the final fit to real data. We used roughly twice as many EvtGen events in the template for these fits. Initially the value for the weight factor applied to each category in Table 8.4 was fixed to 1.0, representing the ratio of data/template for these tertiary leptons. We adjusted the fixed weight of each of these categories individually, and used the exact same lepton pairs for both the data and template side. This causes double-counting of lepton pairs, and artificially lowers the fit χ^2 , but should not affect the shifts in the parameter values in which we are interested. We apply the full set of event selection cuts as described in Chapter 5 to both the data and template. The data includes ten copies of the off-peak data acting like continuum background, and the template contains one copy of the off-peak data, with a weight of 10.0.

We performed a series of MCRW fits, and varied the fixed weight factor for the given category down by 10%, in order to create a reasonably sized mismatch between the data and template in each fit. We used the difference between the fitted parameter value with the weight factor fixed at 1.0 and 0.9, to estimate the systematic uncertainty, and divided by the expected true value of the parameter Δm , to obtain the percentage shift, given in Table 8.4.

From this table we see that the largest contribution to the Tertiary lepton systematic uncertainty is due to the Fake Muons, followed by the Large Energy Loss Bremsstrahlung Electrons. It is perhaps easiest to reduce the Fake muon contribution by tightening the PID requirements for muons, from the Tight, to the “Very Tight” selector. This may not help overall, however, due to the reduced data statistics available in the fit. On the other hand, it is not immediately clear how to improve the modelling of the Large-Energy-Loss Bremsstrahlung Electrons in our template, as would be necessary if this is one of the dominant systematic uncertainties. If the tertiary lepton systematic uncertainty is a limiting factor when trying to analyze real data, further tightening of the p_{cm} and

θ cuts may be necessary to achieve an acceptable size of the systematics due to these tertiary leptons, even though the statistical error would suffer.

8.5 Continuum Backgrounds

There are several systematic uncertainties associated with the treatment of the continuum events that are present in data, and must be simulated in the template. We apply one overall scaling factor to the off-peak data which we use in our template, and we must simulate the requirements on the reconstructed quantities $NTracks$ and R_2 on the EvtGen portion of the template. These selection requirements would not have been used if we didn't have to remove continuum background. To study both of these issues, we will vary the fixed parameters, and estimate systematic uncertainties by measuring shifts in the measured value of Δm when fitting reconstructed Monte-Carlo events.

The treatment of continuum background events creates three sources of systematic uncertainty associated with our fits. The first, and most important, is the uncertainty in the amount of continuum background present in our data sample. Because we are fixing this amount in the fit, based on the measured relative luminosities of on-peak and off-peak data, the uncertainty in this continuum scale factor can potentially contribute to the uncertainty in the Δm measurement.

To estimate this systematic uncertainty, we used off-peak data in the template and data samples, scaled by the appropriate ratio of luminosities to match the expected composition of real *BABAR* data. We then varied the weighting of the continuum (off-peak data) component in the template, to see how much the measured Δm value shifted when the amount of continuum in the data and template don't agree. To remove other sources of disagreement, we used our EvtGen-based template lepton pairs in both the data and template parts of the fit, so there was no difference between the data and template. The correct weighting of the continuum background gave a perfect fit, and changing this weight caused a change in the fitted value of Δm and the other floating parameters. These shifts depended on the requirements chosen for $NTracks$ and R_2 , which determined the amount of continuum background present in the fitted samples. The amount of continuum and parameter shifts depend more strongly on $NTracks$ which is a more powerful discriminating variable than R_2 .

We performed all the fits described in this section with the standard event and track selection requirements of Chapter 5 applied, except for the changing requirements on $NTracks$ and R_2 . We observed the shifts in Δm given in Table 8.5 for values of the required $NTracks$ ranging from 3 to 6.

In order to determine the actual systematic uncertainty from the numbers in Table 8.5, we need to know the uncertainty on the amount of continuum background in our selected data sample. When analyzing real *BABAR* data, this uncertainty will be based on the uncertainties in the relative luminosities of the on-peak and off-peak data, as well as the statistical error in the total number of events (which is negligible in comparison).

NTracks Cut	Shift Amount	Δm Shift	% Δm Shift
3	1%	8.4×10^{-4}	0.17
4	1%	4.8×10^{-4}	0.10
5	1%	3.5×10^{-4}	0.07
6	1%	3.3×10^{-4}	0.07

Table 8.5: Preliminary estimates of the systematic uncertainty due to an assumed 1% uncertainty in the fixed continuum background scale factor.

Our colleagues in the *BABAR* collaboration have recently performed an updated measurement of the total integrated luminosity, both on and off the $\Upsilon(4S)$ peak for the entire data set of the experiment, and arrived at the following result [21]:

$$L_{\text{on}}/L_{\text{off}} = 9.658 \pm 0.003 \pm 0.006 \quad (8.1)$$

where the first error is statistical, and the second is systematic. If we combine the statistical and systematic uncertainties in quadrature, we obtain ± 0.007 , or $\pm 0.07\%$. This remarkably small error is possible because many of the errors involved in determining the integrated luminosity cancel out in the ratio. Our normalization depends on the ratio:

$$\frac{N_{\text{Continuum}}(\text{OnPeak})}{N_{\text{Continuum}}(\text{OffPeak})} = \frac{L_{\text{on}} \sigma_{\text{On}} \text{Eff}_{\text{On}}}{L_{\text{off}} \sigma_{\text{Off}} \text{Eff}_{\text{Off}}} \quad (8.2)$$

where the ratio $\sigma_{\text{On}}/\sigma_{\text{Off}} \simeq 1.007$, and we safely assume that the efficiency ratio $\text{Eff}_{\text{On}}/\text{Eff}_{\text{Off}} = 1.00$. We conservatively assign an uncertainty of $\pm 1\%$ to the ratio $N_{\text{On}}/N_{\text{Off}}$ for the continuum background, which is used in the calculations of Table 8.5. From these numbers, we assign a systematic uncertainty in Δm to our present analysis (requiring $N\text{Tracks} \geq 5$ and $R_2 \leq 0.8$) of 0.07%.

The second and third sources of uncertainty are related to our ability to model the continuum suppression requirements on the $N\text{Tracks}$ and R_2 variables in our template, as a reflection of the procedure in the data. In the EvtGen portion of the template, these variables are not normally available for use, because the reconstruction software has not been run. Therefore we use a measured ChargedTracks efficiency, and a Poisson random correction, with the same mean as the observed difference between the average number of charged tracks in real $B\bar{B}$ events, and our simulated-calculated number for the EvtGen only events. This procedure is described in detail in Chapter 5. Differences in the observed distributions for $N\text{Tracks}$ and R_2 can then contribute to small shifts in the measured value of Δm .

To estimate the systematics, we fit reconstructed Monte-Carlo events as the data in two different ways. First we use the reconstructed values $N\text{Tracks}$ and R_2 to perform the selections, and then we used the EvtGen-equivalent values (based on MC truth in the events) to perform the selections. The difference

$NTracks$ Cut	R_2 Cut	Δm Error	% Δm Error
4	0.8	3.0×10^{-5}	0.005
5	0.6	0.0002	0.04
5	0.7	0.0006	0.12
5	0.8	0.0008	0.15
5	0.9	0.0008	0.17
6	0.8	0.0010	0.20

Table 8.6: Preliminary estimates of the systematic uncertainty due to the EvtGen model of reconstructed $NTracks$ and R_2 .

between the fitted value of Δm in these two cases is taken as the estimate of the systematic uncertainty.

We document the shifts observed in Δm for a range of requirements on both $NTracks$ and R_2 in Table 8.6.

We choose to use requirements of $NTracks \geq 5$ and $R_2 \leq 0.8$, as explained in Section 8.6. Furthermore, we follow the conventions of the *BABAR* Collaboration and assign the size of the systematic uncertainty to be equal to one half the size of the observed shift. From these numbers, we assign a systematic uncertainty to our present analysis of 0.15%, due to the model for $NTracks$, and 0.12%, due to the model for R_2 in EvtGen template events.

Assuming that the three systematic uncertainties described in this section are uncorrelated, we combine the three of them (0.07, 0.15, 0.12) in quadrature, to arrive at a total systematic uncertainty related to the treatment of continuum backgrounds for the present analysis of 0.20%. This means that the continuum background represents a noticeably smaller contribution to the systematic uncertainties than either the EvtGen or Tertiary parts of the template model.

8.6 Optimizing Event Selection Cuts

We have made a basic attempt to optimize the p_{cm} , θ , $NTracks$ and R_2 cuts, based on fits to simulated events, in order to minimize the quadrature sum of statistical and systematic uncertainties, hereafter referred to as the “Total” error. We also chose between the “Loose” and “Tight” PID selectors based on looking at the Δm error resulting in fits to simulated events.

We did not perform a “global” analysis of the errors. Instead, we examined the categories of Cascades, Tertiary Leptons, and Continuum errors individually. For each of these categories, we systematically varied the cuts on the variables: p_{cm} and θ , and recorded the Δm fit errors, as well as the estimated systematic uncertainties. We performed this analysis for $NTracks$ and R_2 , taking into account only the Continuum systematic uncertainties. These curves then contain information about the impact on the statistics of the signal as the cut is varied, as well as the dilution of the signal due to background, and the

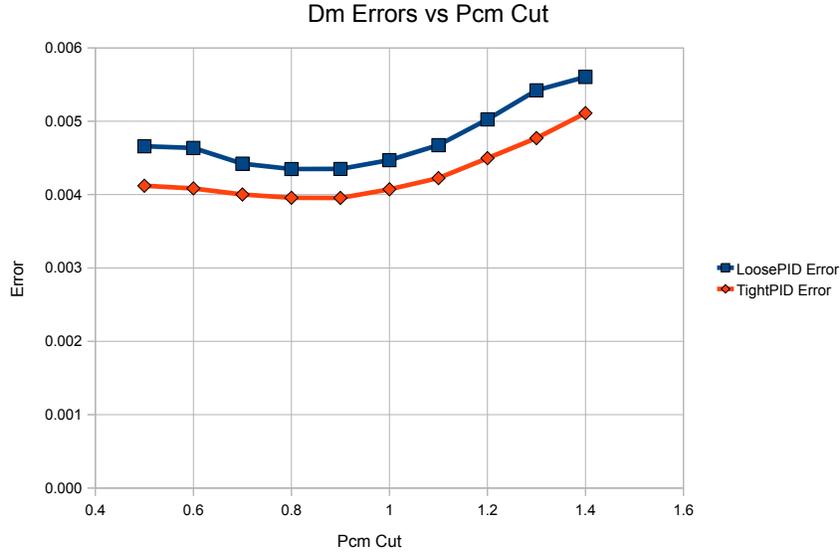


Figure 8.1: Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Tertiary leptons, as a function of the lower limit on p_{cm} . Curves are shown for fits using both the Loose and Tight PID selectors.

size of the systematic uncertainty which generally scales with the amount of background present in the sample.

Figures 8.1 and 8.2 show plots of the quadrature sum of the Δm statistical error and the estimated systematic uncertainty due to Tertiary lepton backgrounds, as functions of the cuts on p_{cm} and θ respectively. Similar curves were obtained for the systematic uncertainty due to Cascade lepton backgrounds. Inspection of the curves in Figure 8.1 shows a fairly mild dependence on the lower p_{cm} cut, presumably due to a near-balance between the removal of background leptons and signal leptons as the cut value is raised. The preferred cut value appears to be located near $p_{cm} \geq 0.8$, and it appears that Tight PID selectors, despite the lower efficiency for signal leptons, are better overall. This is likely due to the removal of more fake leptons, than signal leptons.

Inspection of the curves in Figure 8.2 shows a fairly mild dependence on the lower θ cut, again presumably due to a near-balance between the removal of background leptons and signal leptons, until a threshold near $\theta \geq 0.8$ is reached. The preferred cut value appears to be located anywhere between 0.5 and 0.8 in this plot. The results for the Cascade lepton systematics indicate a slight preference for a cut near 0.7. We see a small preference here for the Tight PID selectors, compared to the Loose ones.

In Figure 8.3 we varied the $NTracks$ and R_2 cuts, while determining the



Figure 8.2: Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Tertiary leptons, as a function of the lower limit on θ . Curves are shown for fits using both the Loose and Tight PID selectors.

quadrature sum of the Δm statistical error and the systematic error due to our fixed contribution from continuum leptons. From these curves it is evident that a looser R_2 cut of 0.8 yields lower total errors, and that a cut at $NTracks \geq 5$ appears to be preferred. The curves have a clear minimum that is presumably due to the removal of more QED and $udsc$ events, and keeping the maximum number of B events.

Based on these simple studies, we chose the following set of selection criteria as “optimal” for this analysis, and used them to perform the fit in Chapter 7, and evaluate the systematic uncertainties above:

- $0.8 < p_{cm} < 2.8 \text{ GeV}/c$
- $0.7 < \theta < 2.3$
- $NTracks \geq 5$
- $R_2 \leq 0.8$
- Tight PID Selectors, rather than Loose.

A more thorough, global analysis including the simultaneous effect of all the selection criteria on all systematic uncertainties should be performed before this fitting technique can be used to publish a measurement on real data.

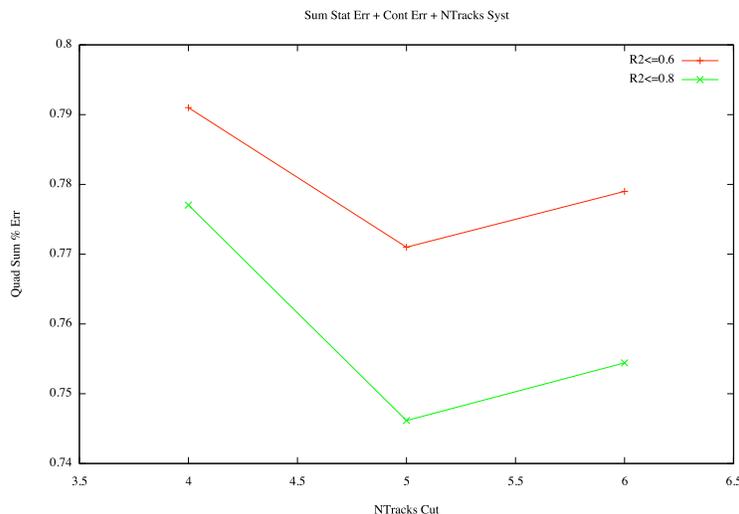


Figure 8.3: Quadrature sum of the statistical error on Δm and the systematic uncertainty due to Continuum background, as a function of the minimum allowed $NTracks$ value, for two different requirements on R_2 .

8.7 Systematics Summary

In this section we present Table 8.7, which contains a summary of the different categories of systematic uncertainties described and estimated above. We see that currently the B lifetimes are the dominant systematic uncertainty, followed by the contribution to uncertainties in the model of Cascade leptons. The tracking resolution model, Tertiary leptons, and Continuum leptons create smaller, and roughly equal systematic uncertainties. In order to make substantial improvements in the total systematic uncertainty, those due to B lifetimes and Cascade leptons would need to be reduced.

It is possible that a future, more careful evaluation of the systematic uncertainties due to the Cascade leptons could decrease these estimates somewhat. This would require a detailed comparison of inclusive rates of leptons from the various charm decays and τ decays in *BABAR* data and simulations, to estimate the size of possible differences. The systematic uncertainty due to the uncertainty in the B lifetimes can only improve if the particle physics community is able to better measure the average or ratio of the B lifetimes in the extremely high statistics of the LHC experiments, which may be possible. If this effort is made, the decrease in uncertainty on the measured B lifetimes would shrink the associated uncertainty in a roughly linear fashion.

Description	% Δm Shift
<i>B</i> Lifetimes	0.50
Cascade Leptons	0.29
Track Resolution	0.23
Tertiary Leptons	0.25
Continuum Leptons	0.20
Total (Excluding Lifetimes)	0.49

Table 8.7: Preliminary estimates of the total systematic uncertainties on Δm .
The total is the quadrature sum.

9

Conclusions

This Chapter contains a brief summary of the work that is presented in this thesis, as well as an outlook for the future of our new fitting technique, Monte-Carlo-Re-Weighting (MCRW fits). We close with a look at the prospects for publishing a new, more precise measurement of Δm for neutral B mesons.

9.1 Summary

In this thesis, we began by discussing the difficulties associated with creating detailed analytical models including experimental resolution and acceptance for use in standard likelihood fits. We then developed a novel method of analyzing experimental data to extract model parameters, which we call the Monte-Carlo Re-Weighting Method (MCRW). This method relies upon the existence of a detailed Monte-Carlo simulation of the experiment, and uses the simulated events directly in determining the likelihood, rather than using an analytical expression which is often only approximate. We presented a simple example problem to demonstrate the operation of the MCRW software, and compared its results to those obtained using the traditional software tools of experimental particle physics.

We then introduced the experimental issues involved in measurements of the neutral B meson mixing frequency, Δm , in the dilepton sample of the *BABAR* experiment. Attempts to use the full sample of data recorded by the experiment were thwarted by the shortcomings of the analytical model previously used to fit the data and associated biases which appeared in these fits. Boost smearing was identified as one possible origin of these biases, and the associated systematic uncertainties. By using the MCRW method to analyze the measured data from the dilepton sample to extract the value of Δm , we are able to incorporate a fully detailed model of boost smearing into the fit, and completely remove this potential source of bias and associated systematic uncertainty from the results, thus overcoming one of the obstacles in the path to an update of the *BABAR* measurement.

The remainder of the thesis describes the application of the MCRW method to the problem of measuring Δm in fully simulated B meson events. We incorporated a detailed tracking resolution model, event and track selection criteria, detector efficiencies and particle ID in our template, as well as included all of the major sources of background. Estimates presented in Chapters 7 and 8 indicate that the MCRW method should obtain an excellent measurement of Δm .

9.2 Outlook - The MCRW Fit Technique

This thesis has presented a new method of fitting experimental data, which we call the Monte-Carlo-Re-Weighting (MCRW) method. This method was designed to allow the user to analyze complicated experimental data in the case where a detailed MC simulation of the experiment is possible, but it is extremely difficult or impossible to write a concise analytical PDF describing the distribution of the measured variable(s). The technique uses a large sample of simulated events in memory to recompute a template histogram for comparison to the data histogram on each iteration of the fit. As the fitted parameters change, the events in the template are re-weighted according to simple PDFs which describe the underlying physical processes relating to the parameters of interest, and then the template histogram is re-computed.

Rather than being CPU instruction intensive, this technique requires a relatively large amount of memory to store the simulated events of the template (up to 2 GB in our example application of fitting for Δm), but this is easily achievable on modern computers.

The MCRW fit method is completely general-purpose, and could be useful for other data analysis projects, especially in experiments in which a detailed Monte-Carlo simulation of the experimental apparatus exists, and there are ample data statistics for a detailed measurement. In situations in which the researchers are searching through a large amount of data for extremely rare events, or when the statistical errors dominate, a reduction of systematic uncertainties may be unimportant, but the convenience of implementing a model directly using MC events, instead of analytical approximations, will remain.

Any scientific field in which there is a detailed Monte-Carlo simulation of the experimental apparatus, and for which this Monte-Carlo simulation is a more precise description of the measured data than any simple analytical function, could benefit from the use the MCRW method. This method is intended to estimate the values of numerical model parameters, not to test the null hypothesis in experiments which seek to determine whether or not any relationship exists between variables. A large amount of recorded and simulated data is needed for this technique to be applicable, which also limits the range of analyses where it may be applied. There are likely applications within computational finance where it is necessary to estimate model parameters by comparing historical financial data and the output of detailed financial simulations. There also likely applications in computational particle physics where computer models of fragmentation and hadronization must be calibrated against experimental data. It is also possible that the method could be used in genetics research, where researchers use computer models to analyze enormous amounts of data in a probabilistic manner.

The author and his supervisor plan to continue developing the MCRW technique, and sharing it with other members of the experimental particle physics community, both inside and beyond the *BABAR* collaboration. They also plan to publish a journal paper describing the fit method with several simple, detailed examples to show that it is capable of performing as well as a traditional fit in

those cases where one can write an exact analytical fit model, and is capable of out-performing the traditional methods when an exact analytical model cannot be found. Only time will tell whether or not this method will see further use in the scientific community.

9.3 Outlook - Measuring Δm

The motivation for the development of the MCRW software was the systematics-limited analysis of neutral B mixing in the dilepton sample, carried out at *BABAR*. The experiment has now finished, after accumulating roughly 470 M $B\bar{B}$ pairs, and this increased data set, compared to that which was used to publish earlier measurements (23 M $B\bar{B}$ pairs), presents an opportunity to greatly improve the precision of measurements of Δm . We saw in Chapter 3 that this analysis reached a point at which it was limited by systematic uncertainties in the past, and had not been updated for several years. The limiting systematic was associated with an observed fit bias which we believe may have been caused in part by the lack of a detailed model for boost smearing in the fit.

Using Monte-Carlo B decays as our template allows us to include a fully detailed description of boost smearing in our template, and eliminate this potential source of bias. Based on the fit results shown in Chapter 7, and on earlier, preliminary studies which only examined generator-level Monte-Carlo events, we no longer see any evidence for a fit bias due to boost smearing. Those same fit results suggest that a MCRW fit could obtain a statistical error of 0.5% for a measurement of Δm , in neutral B mesons, with the dilepton sample, based on the complete *BABAR* data set.

The results presented in Chapter 8 indicate that the methods of this Thesis can achieve a total systematic uncertainty smaller than 1%. It is possible that a future, more careful evaluation of the systematic uncertainties due to the tracking resolution, cascade leptons, and tertiary leptons could decrease these estimates somewhat. The dominant systematic uncertainty due to uncertainty in B lifetimes is only going to improve if the particle physics community is able to better measure the average or ratio of the B lifetimes in the extremely high statistics of the LHC experiments.

We expect that an analysis performed on the complete *BABAR* data set following the methods detailed in this Thesis would observe:

$$\Delta m = \text{x.xxxx} \pm 0.0026 \text{ (Stat.)} \pm 0.0024 \text{ (Lifetimes)} \pm 0.0024 \text{ (Syst.) ps}^{-1} \quad (9.1)$$

where x.xxxx represents the central value from the fit. If we assume that all three sources of error are completely uncorrelated, and add them in quadrature we would obtain:

$$\Delta m = \text{x.xxxx} \pm 0.0043 \text{ (Combined) ps}^{-1} \quad (9.2)$$

These statistical and systematic uncertainties would be significantly better than any other published measurement, and would improve the errors on the world

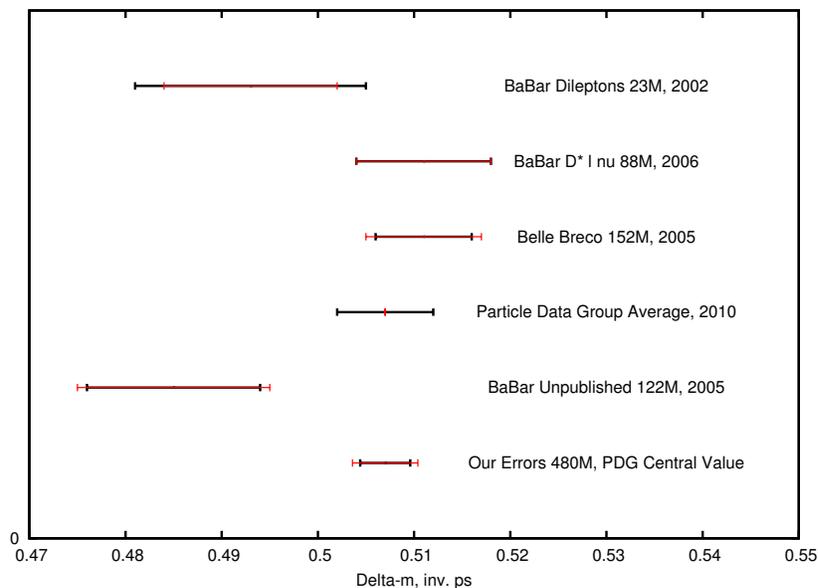


Figure 9.1: Measurements of Δm in neutral B meson mixing. Statistical error bars are black, systematic error bars are red. Our projected errors are centered on the PDG average value.

average value, as shown in Figure 9.1. This figure is adapted from Figure B.3, and our error bars are shown centered at the current world average value of Δm .

Of course, there is always the possibility that an increase in the precision of the measurement may actually point out a disagreement between the measurement of Δm and the other experimental values which are tested in the CKM picture of electroweak physics. Even if no such disagreement were found, improved precision of measurements in the B meson sector may be useful in constraining the possible scenarios for new physics at the Energy Frontier experiments such as at the LHC.

Bibliography

- [1] B. Aubert *et al.* [*BABAR* collaboration], Phys. Rev. Lett. **88**, 221803 (2002)
- [2] M. Chao and D. Best, Measurement of the $B^0\bar{B}^0$ Mixing Frequency with Inclusive Dilepton Events, *BABAR* Analysis Documents 691, 1304. (2005)
- [3] F. Abe *et al.* [CDF Collaboration], Phys. Rev. Lett. **74**, 2626 (1995) [arXiv:hep-ex/9503002].
- [4] R. J. Barlow and C. Beeston, Comput. Phys. Commun. **77**, 219 (1993).
- [5] I. Antcheva *et al.*, Comput. Phys. Commun. **180**, 2499 (2009).
- [6] F. James and M. Roos, Comput. Phys. Commun. **10**, 343 (1975).
- [7] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C, 2nd Ed. (1992)
- [8] F. Le Diberder, *BABAR* Note 42. (1990)
- [9] W. Verkerke and D. P. Kirkby, *In the Proceedings of 2003 Conference for Computing in High-Energy and Nuclear Physics (CHEP 03), La Jolla, California, 24-28 Mar 2003, pp MOLT007* [arXiv:physics/0306116].
- [10] G. Moliere, Z. Naturforsch. **3a**, 78 (1948)
- [11] H. A. Bethe, Phys. Rev. **89**, 1256 (1953)
- [12] Lange, D. J. 2001, Nucl. Instrum. Meth. A, **462**, 152 (2001).
- [13] K. Nakamura *et al.* [Particle Data Group], J. Phys. G **37**, 075021 (2010).
- [14] S. Agostinelli *et al.* [GEANT4 Collaboration], Nucl. Instrum. Meth. A **506**, 250 (2003).
- [15] P. F. Harrison, H. R. Quinn [*BABAR* Collaboration], “The BaBar physics book: Physics at an asymmetric B factory,” SLAC-R-0504, (1998).
- [16] G. C. Fox and S. Wolfram, Phys. Rev. Lett. **41**, 1581 (1978).
- [17] P. Ongmongkolkul, A. Gaz, K. Mishra, A.V. Telnov, Particle Identification Using Error Correcting Output Code Multiclass Classifier, *BABAR* Analysis Document 2199, (2009).

-
- [18] C.O. Vuosalo, K. Flood, A.V. Telnov, Muon Identification Using Decision Trees, *BABAR* Analysis Document 1853, (2010).
- [19] K. Abe *et al.* [BELLE Collaboration], Phys. Rev. D **71**, 072003 (2005) [Erratum-ibid. D **71**, 079903 (2005)] [arXiv:hep-ex/0408111].
- [20] G.D. MacGregor, B Counting at Babar, MSc. Thesis, UBC, (2008).
- [21] Abi Soffer, J.P. Burke, Eisner, A., Guttman, N., Hearty, C., Long, O., Neal, H., Peimer, D., Roney, M. The Integrated Luminosity of the *BABAR* Experiment, *BABAR* Analysis Document 2365, (2011).
- [22] K. Abe *et al.* [Belle Collaboration], Phys. Rev. Lett. **87**, 091802 (2001).
- [23] E. Noether, Nachr. v.d. Ges.d.Wiss.zu Gottingen, 235 (1918).
- [24] T. D. Lee, C. N. Yang, Phys. Rev. **104**, 254 (1956).
- [25] C. S. Wu, *et al.*, Phys. Rev. **105**, 1413 (1957).
- [26] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, Phys. Rev. Lett. **13**, 138 (1964).
- [27] A. D. Sakharov, Sov. Phys. Usp. **34**, 417 (1991).
- [28] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).
- [29] L. Wolfenstein, Phys. Rev. Lett. **51**, 1945 (1983).
- [30] V. A. Kostelecky, Phys. Rev. D **64**, 076001 (2001).
- [31] A. J. Buras, W. Slominski and H. Steger, Nucl. Phys. B **245**, 369 (1984).
- [32] A. J. Buras, W. Slominski and H. Steger, Nucl. Phys. B **238**, 529 (1984).
- [33] B. Aubert *et al.* [*BABAR* Collaboration], Nucl. Instrum. Meth. A **479**, 1 (2002).
- [34] R. Santonico and R. Cardarelli, Nucl. Instrum. Meth. **187**, 377 (1981).
- [35] G. Battistoni *et al.* Nucl. Instrum. Meth. **164**, 57 (1979).

Appendix A

Abbreviations and Acronyms

- MC = Monte-Carlo - A term denoting simulated events, generated using pseudo-random numbers in a computer program.
- MCRW = Monte-Carlo Re-Weighting, the type of maximum likelihood fit developed in this thesis.
- SLAC = Stanford Linear Accelerator Center, now known as SLAC National Laboratory, where the letters SLAC are no longer an acronym.
- SVT = Silicon Vertex Tracker - The innermost component of the *BABAR* detector. Responsible for geometrical vertexing of charged tracks.
- DCH = Drift Chamber
- EMC = Electromagnetic Calorimeter
- DIRC = Detector of Internally Reflected Cerenkov Light.
- IFR = Instrumented Flux Return
- LL = Log-likelihood - Maximum Likelihood fits are almost always performed in practice by minimizing the quantity $-LL$ or $-2LL$.
- MCS = Multiple-Coulomb-Scattering - The description of charged particle deflections resulting from the passage through matter. The distribution has a complex shape, described by the theory of Moliere [10].
- OBC = Other *B* Cascade - An event in which one lepton is directly from *B* decay, and the other is from a cascade decay of the other *B*.
- OS = Opposite Sign - A pair of leptons, each with opposite charge.
- PDF = Probability Density Function - The normalized function describing the probability of x , given some parameters \vec{p} .
- PID = Particle IDentification - The combination of hardware and software used to distinguish different particle species in the *BABAR* detector.
- POCA = Point Of Closest Approach - The point on a track which is closest in distance to the measured beam spot.

- SBC = Same B Cascade - An event with one lepton directly from B decay, and the other from a cascade decay of the same B .
- SS = Same Sign - A pair of leptons, each with the same charge.

Appendix B

Standard Model and Physics of B Meson Oscillations

B.1 Experimental Particle Physics

Experimental particle physics, broadly speaking, is the pursuit of the understanding of the fundamental building blocks and laws of nature, on the smallest experimental scales. Equivalently, it can also be described as the effort to understand the laws of nature at higher and higher energy scales. Historically, experiments have moved rapidly over the last one hundred years from the discovery and verification of the existence of atoms, to the discovery of electrons and the nucleus, to the discovery of the components of the nucleus: protons and neutrons, to the discovery of the components of protons and neutrons: quarks and gluons. We currently believe that quarks, gluons and leptons are indivisible, fundamental particles, but are actively engaged in the search for new particles and laws of nature.

There have been additional discoveries along the way of a range of other particles besides quarks and gluons, including the mysterious neutrinos, the muon and tau, heavier cousins of the electron, the gauge bosons W and Z , which mediate the weak force, and of course antimatter partners for all of the known particles (except perhaps for neutrinos, where the distinction between matter and antimatter is not yet clear). All these discoveries have been made with a combination of theoretical predictions leading the way, or unexpected experimental findings pointing the field in new directions.

Current particle physics experiments may be divided into three broad categories, based upon their method and object of investigation:

- The Energy Frontier
- The Cosmic Frontier
- The Intensity Frontier

The famous Large Hadron Collider (LHC), which began operations in 2010, is currently the leading particle accelerator in the world dedicated to exploring the “Energy Frontier”, taking over from the Tevatron at Fermilab, which is

due to be retired before the end of 2011. Many of the most famous discoveries in the history of particle physics have taken place at the Energy frontier, and the scientists working there hope that the LHC will continue in that tradition. By probing to higher and higher energies in particle collisions, experiments are able to search for heavier short-lived particles. The quest for higher energy collisions can also be seen as an effort to examine shorter and shorter distance scales; to search for structure within what seem to be point-like particles that we see today. Evidence for new types of particles, or new laws of physics, would hopefully take the form of a direct observation at the LHC. This machine has taken decades of effort, billions of dollars, and the work of thousands of physicists to be designed, built and commissioned. There are also at least two proposals in place for future e^+e^- colliders, one known as the ILC, or International Linear Collider, and another known as CLIC, or the Compact Linear Collider. These would operate at similar energies to the LHC, but producing much cleaner events through e^+e^- annihilation, and allowing for higher precision measurements of new particles which may be discovered at the LHC. There also early discussions and plans for constructing a muon collider, as a way to reach higher energies beyond those of the LHC, ILC or CLIC. It is unclear which proposed future accelerator, if any, can surpass the LHC and push into even higher energies in the future, or whether the Energy Frontier will eventually be abandoned due to the enormous costs and difficulties involved.

As the name implies, the “Cosmic Frontier”, describes a series of experiments, (or observations), combining traditional particle physics and astrophysics and astronomy. The synergy works in both directions: using the technology and techniques of particle physics to learn more about astrophysics and astronomical objects, and studying different types of particles that arrive at the earth from outer space, to look for new particles beyond the ones we already know and understand. This is an area of research that has a long history, especially involving the study of cosmic rays, and neutrinos emanating from our sun. Efforts are under way to look for ultra-high energy neutrinos, dark matter particles, high energy gamma rays, and others. Currently this field appears to have a bright future, with new facilities under construction, and in operation, as well as experiments being launched aboard satellites and the International Space Station. The study of neutrinos, in particular, is a field which bridges both the Cosmic Frontier (through the study of solar neutrinos, and high-energy neutrinos from outer space) and the Intensity Frontier, at accelerator-based experiments.

The *BABAR* Experiment at SLAC, and the *BELLE* Experiment at KEK in Japan [22], fall under the third of these categories, described as exploring the “Intensity Frontier”. Rather than building a larger, higher-energy, and more expensive accelerator, like the LHC, these experiments generally upgrade the beam currents at existing, lower-energy facilities in order to acquire enormous data samples of particle interactions that have already been observed. By studying millions or billions of these interactions in great detail, and comparing the experimental results to precise predictions from the Standard Model of particle physics, scientists hope to either confirm and support the current theories, or ideally to discover an unexplained discrepancy between the observations and

the theoretical predictions. This type of discrepancy would generally be an “indirect” signal of new particles or new forces, in which the unknown new particles subtly influenced the behaviour of the known and observed species. *BABAR* was an experiment conceived and designed to carry out these types of studies in the field of “Flavour” physics, which means primarily studying the Weak force interactions between quarks and leptons. Among the particles studied at *BABAR*, were the B mesons, charm mesons, kaons, and τ leptons. The data analysis method developed in this thesis was created in order to improve the study of quantum mechanical mixing between neutral B mesons. The *BABAR* Experiment has finished taking new data, but is still publishing analyses of their data. The BELLE experiment and the KEK collider are performing upgrades to continue running at higher luminosity in the near future. There is also a new project called SuperB, which is soon to begin construction in Italy, and it will continue to push forward the Intensity Frontier.

B.2 The Standard Model and the CKM Matrix

One of the primary motivations of the *BABAR* Experiment was the measurement of CP-violation phenomena in B -mesons – mesons containing at least one b (bottom or beauty) quark [15]. CP is a quantum-mechanical operator, corresponding to changing the sign of the charge, and applying the 3-dimensional parity or mirror-reversal operator. We know that if a symmetry is respected in nature, there must be a corresponding quantity which is conserved; a result proved by Emmy Noether in 1918 [23]. Translation symmetry leads to conservation of energy, rotational symmetry leads to conservation of angular momentum. CPT symmetry is associated with gauge invariance. It was believed that both C and P should be good symmetries of nature, until the possibility of CP violation was suggested by Lee and Yang in 1956 [24]. This prediction was verified when Madame Wu [25] experimentally observed parity violation in 1957. After this discovery, it was known that the nuclear weak force is maximally parity violating, or that parity was never conserved in a purely weak interaction. The subsequently developed V-A theory of the weak interaction restored CP as a good (unbroken) symmetry, but CP-violation was observed experimentally at a level of 2 parts in 1000 in Kaon decay by Cronin and Fitch in 1964 [26]. Since that time, CP-violation has been one of the most heavily studied topics in particle physics, particularly due to Sakharov’s realization that it is a necessary ingredient in any explanation of the matter-antimatter asymmetry of the universe [27]. The CPT symmetry is always preserved in any quantum field theory which is local, Lorentz invariant and has a Hermitian Hamiltonian.

It was shown by the theorists Kobayashi and Maskawa that CP violation is allowed, or in fact predicted, in the Standard Model of particle physics if there are at least three generations of quarks that are mixed through the weak interaction [28]. If there are exactly three generations of quarks, then the weak mixing between quark flavours can be described mathematically by a 3×3

matrix normally written as:

$$V = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \quad (\text{B.1})$$

where the subscripts u, d, s, c, t, b stand for the “up”, “down”, “strange”, “charm”, “top” and “bottom” species of quarks, respectively. In principle each element of the so-called CKM (Cabibbo, Kobayashi, Moskawa) matrix can be complex, but the matrix must be unitary if there are only three generations of quarks in nature. Requiring V to be unitary reduces the number of freely adjustable parameters from 18 (9 real and 9 imaginary numbers) to only 9 real and imaginary numbers. These nine parameters can be described generically as three real rotation angles, and six imaginary phases. With six flavours of quark in the Standard Model, there are five relative phases between the wavefunctions of these quarks. These five arbitrary phases can be multiplied away to reduce the number of free parameters in V from nine to four. A 3×3 matrix containing four parameters can then be written as the product of three rotations (real numbers) and one irreducible phase (an imaginary number). It is this single imaginary phase, that cannot be rotated away through a change of basis, that is ultimately responsible for the observed phenomena of CP-violation, in the Standard Model CKM picture.

The location of the imaginary phase within the elements of the CKM matrix is an arbitrary convention, and several forms of the matrix are frequently used in particle physics literature. The parameterization of Wolfenstein [29], which is based on experimental measurements, is frequently used when discussing the latest results of tests on the unitarity of the CKM matrix. This parametrization is given by:

$$V = \begin{pmatrix} 1 - \lambda^2/2 & \lambda & \lambda^3 A(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & \lambda^2 A \\ \lambda^3 A(1 - \rho - i\eta) & -\lambda^2 A & 1 \end{pmatrix} \quad (\text{B.2})$$

where $\lambda = \sin(\theta_c)$ with θ_c the Cabibbo angle and A, ρ and η can be experimentally measured. In this parameterization, all Standard Model CP-violation is related to the imaginary phase η in the bottom-left and top-right elements of the CKM matrix. It is also clear that as written, this form of the CKM matrix (derived by expanding in powers of λ), is not unitary at the level of λ^4 , but it still provides a convenient method of comparison for experimental results, given the measured value of $\lambda \simeq 0.22$.

By experimentally verifying whether or not the CKM matrix is truly unitary, we can test for the possible existence of new particles beyond the three generations of quarks thus far observed, or new interactions beyond those of the Standard Model of particle physics. One convenient way of testing unitarity is by multiplying one column by a row of the conjugate matrix to obtain an equation with 6 elements:

$$V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* = 0 \quad (\text{B.3})$$

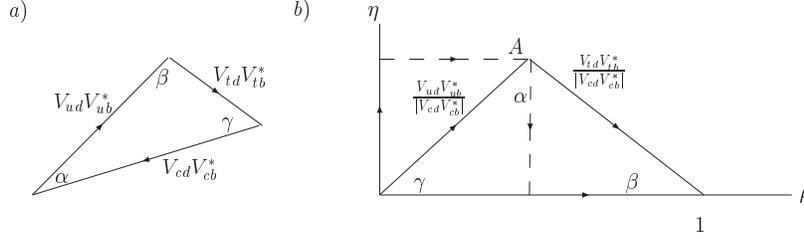


Figure B.1: Unitarity triangles constructed from the CKM matrix in a) standard quark-mixing parameters and b) Wolfenstein parameterization.

which can also be interpreted as the requirement that the sum of three complex quantities vanishes. This can be represented in the Argand plane as the so-called Unitarity Triangle, shown in Figure B.1a. If the three sides are divided by $V_{cd}V_{cb}^*$, the triangle can also be expressed in terms of the Wolfenstein parameters, as shown in Figure B.1b. The experimental programs known as “Precision Flavour Physics” consist of making many measurements related to the angles and sides of this Unitarity Triangle in order to test the Standard Model, and hopefully uncover evidence for new physics beyond the Standard Model. The Kaon system has been thoroughly explored, and in more recent years particle physicists have focussed on examinations of the B meson system to study “heavy” flavour physics in high precision.

B.3 Simplified Model of the Neutral B System

Mixing in the neutral B system is a consequence of the fact that the mass eigenstates for B mesons are not identical to the flavor eigenstates. The mathematical development which follows is similar to that used for the kaon system. We begin by writing an effective Hamiltonian for the system as a 2×2 matrix:

$$i\hbar \frac{d}{dt} \begin{pmatrix} B^0 \\ \bar{B}^0 \end{pmatrix} = H \begin{pmatrix} B^0 \\ \bar{B}^0 \end{pmatrix} = \begin{bmatrix} M_{11} - \frac{i}{2}\Gamma_{11} & M_{12} - \frac{i}{2}\Gamma_{12} \\ M_{21} - \frac{i}{2}\Gamma_{21} & M_{22} - \frac{i}{2}\Gamma_{22} \end{bmatrix} \begin{pmatrix} B^0 \\ \bar{B}^0 \end{pmatrix} \quad (\text{B.4})$$

where the mass and decay matrices, M_{ij} and Γ_{ij} are real. The diagonal elements of M and Γ are associated with the flavor-conserving transitions between $B^0 \rightarrow B^0$ and $\bar{B}^0 \rightarrow \bar{B}^0$. The off-diagonal elements represent the flavor-changing transitions $B^0 \rightarrow \bar{B}^0$ and $\bar{B}^0 \rightarrow B^0$.

We can infer a great deal about the behaviour of the system based on the character of the matrix elements. If $H_{11} \neq H_{22}$, then CPT symmetry is violated. Mixing between the neutral B states will occur in general unless $H_{12} = H_{21} = 0$. If there is mixing, and $|H_{12}| \neq |H_{21}|$, then either the CP or T symmetry is violated. This is generally referred to as “indirect” CP-violation, since it takes place as part of the mixing process, and not as part of the tree-level weak

decays. This form of CP-violation means that the mass eigenstates $B_{H,L}$ are not exact CP eigenstates. This type of CP-violation is the predominant form in the neutral kaon system, but is expected to occur only at the level of several parts per thousand in the neutral B system [30], which has been beyond the reach of measurements so far.

In the simplest case, with $H_{11} = H_{22}$ (CPT conserving) and $|H_{12}| = |H_{21}|$ (CP-conserving), the mass eigenstates are just the symmetric and antisymmetric combinations of B^0 and \bar{B}^0 , and will have two different lifetimes and masses: τ_H, τ_L and m_H, m_L . For the purposes of this thesis, we assume that CP-violation in B mixing is negligible.

B.4 CP-Conserving B -mixing

Once we neglect any CP-violation in B mixing, then diagonalization of the effective Hamiltonian allows us to write the two mass eigenstates of the system as linear combinations (not normalized) of the flavor eigenstates:

$$|B_{L,H}\rangle = |B^0\rangle \pm |\bar{B}^0\rangle \quad (\text{B.5})$$

where the two have a mass difference of $\Delta m = m_H - m_L > 0$. In addition they have a small but finite difference in lifetimes, $\Delta\Gamma = \Gamma_H - \Gamma_L$. The mass difference is given by: $\Delta m = 2|M_{12}|$ and $\Delta\Gamma = 2|\Gamma_{12}|$, where M_{12} and Γ_{12} are the off-diagonal elements of the mass and decay matrices. This mass difference, Δm is also referred to as the B mixing frequency, for reasons which will be explained shortly. Throughout this thesis, all references to Δm will refer to the quantity associated with B_d mesons, consisting of b, d quark-antiquark pairs. Similarly, we can define Δm_s for the B_s mesons, which consist of b, s quark-antiquark pairs. Where the possibility for confusion exists, we will specify either Δm_d or Δm_s .

If we consider the time evolution of these neutral B meson states, we start with pure flavor states created by the $\Upsilon(4s)$ decay at time $t = 0$:

$$|B^0(t)\rangle = g_+(t)|B^0\rangle + \frac{q}{p}g_-(t)|\bar{B}^0\rangle \quad (\text{B.6})$$

$$|\bar{B}^0(t)\rangle = g_+(t)|\bar{B}^0\rangle + \frac{p}{q}g_-(t)|B^0\rangle \quad (\text{B.7})$$

and allow them to evolve into either the same (+) or opposite flavor (-). The ratio q/p is always 1 unless CP violation takes place in the mixing process. The time-dependent coefficients are combinations of both decay and mixing terms:

$$|g_{\pm}(t)|^2 = \frac{e^{-\Gamma t}}{2} [\cosh(\Delta\Gamma t/2) \pm \cos(\Delta m t)] \quad (\text{B.8})$$

where $\Gamma = (\Gamma_H + \Gamma_L)/2$. So after a period of time, a B^0 meson is said to “mix” into a superposition of both B^0 and \bar{B}^0 mesons, governed by the mass difference, or frequency Δm .

Historically, B mixing was first measured using a less-sensitive time-integrated measurement. We can write the time-integrated mixing probability as:

$$\chi = \frac{\int |g_-(t)|^2 dt}{(\int |g_-(t)|^2 dt + \int |g_+(t)|^2 dt)} \quad (\text{B.9})$$

and once the integrals are performed we obtain:

$$\chi = \frac{x^2 + y^2}{2(x^2 + 1)} \quad (\text{B.10})$$

where $x = \Delta m/\Gamma$ and $y = \Delta\Gamma/2\Gamma$. Time-integrated mixing measurements are sensitive to a combination of these parameters. We can measure χ by calculating the time-independent asymmetry:

$$\chi = A = \frac{N_{SS}}{N_{OS} + N_{SS}} \quad (\text{B.11})$$

In the above, N_{OS} are events in which the two neutral B 's are of opposite signs, therefore they are unmixed, and if they are the same sign (N_{SS}) then the B 's have mixed. By ‘‘sign’’, we mean the distinction between B^0 and \overline{B}^0 , and not the charged B 's, B^+ and B^- , which are ignored throughout this discussion.

Rather than study the integrated asymmetry, we can improve on these measurements by looking at time-dependent mixing. If we are able to precisely measure the times of the B decays, we can more precisely determine Δm . We can form the following time-dependent asymmetry:

$$A(\Delta t) = \frac{N_{OS}(\Delta t) - N_{SS}(\Delta t)}{N_{OS}(\Delta t) + N_{SS}(\Delta t)} \quad (\text{B.12})$$

In the simplest case, taking the limit of $\Delta\Gamma \rightarrow 0$, and assuming no CP-violation, we can write the following expressions for the mixed and unmixed rates:

$$N_{SS} \sim e^{-|\Delta t|/\tau} [1 - \cos(\Delta m \Delta t)] \quad (\text{B.13})$$

$$N_{OS} \sim e^{-|\Delta t|/\tau} [1 + \cos(\Delta m \Delta t)] \quad (\text{B.14})$$

We will use these formulae for the time dependence of the opposite sign and same sign pairs in the current work. In addition, substituting these expressions into the time-dependent asymmetry $A(\Delta t)$ yields the simple relation:

$$A(\Delta t) = \cos(\Delta m \Delta t) \quad (\text{B.15})$$

which emphasizes the mixing contribution only, and factors out the B lifetimes and exponential decay.

B.5 $B^0 - \overline{B}^0$ Mixing and the CKM Matrix

The Standard Model processes dominating neutral B -mixing are depicted in Figure B.2. These Feynman diagrams depict the weak W -box diagrams which

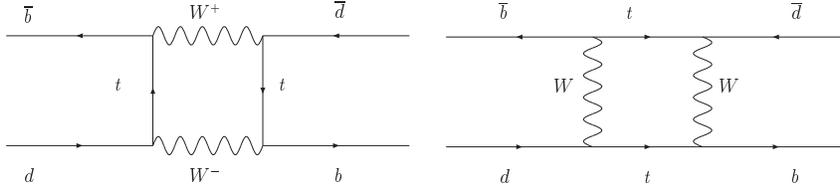


Figure B.2: Dominant box diagrams for $B^0 \rightarrow \bar{B}^0$ transitions. Similar diagrams exist where the t quark is replaced by u or c quarks, but they make much smaller contributions to the overall amplitude.

are the leading, 2nd-order contributions to $B\bar{B}$ -mixing. Changing from a b to \bar{b} quark is a $\Delta F = 2$ transition, (where F is flavour) which is forbidden at the tree-level in the Standard Model. Through the box diagrams, particles heavier than the top quark can theoretically contribute to the amplitudes, allowing a window into possible new physics through precision measurements of neutral B mixing.

Calculation of the amplitudes for the box diagrams yields the following predictions [31] for the off-diagonal elements of the mass and decay matrices:

$$M_{12} = -\frac{G_F^2 m_W^2 \eta_B m_B B_B f_B^2}{12\pi^2} S_0(m_t^2/m_W^2) (V_{td}^* V_{tb})^2 \quad (\text{B.16})$$

$$\Gamma_{12} = \frac{G_F^2 m_b^2 \eta'_B m_B B_B f_B^2}{8\pi} \left[(V_{td}^* V_{tb})^2 + V_{td}^* V_{tb} V_{cd}^* V_{cb} \mathcal{O}\left(\frac{m_c^2}{m_b^2}\right) + (V_{cd}^* V_{cb})^2 \mathcal{O}\left(\frac{m_c^4}{m_b^4}\right) \right] \quad (\text{B.17})$$

where G_F is the Fermi coupling constant of the electroweak force, m_b is the mass of the b quark, m_W is the mass of the W boson, m_t is the mass of the top quark, and m_B is the mass of the B meson. η_B and η'_B are QCD numerical corrections of order 1, B_B is a “bag factor” which can only be calculated through Lattice QCD, f_B is the B^0 decay constant and V_{ij} are elements of the CKM matrix describing weak transitions between quark flavors. The function $S_0(x)$ is approximated quite well by $0.784x^{0.76}$, as calculated in Reference [32].

There are also contributions to neutral B mixing from similar diagrams in which the t quark is replaced by u or c quarks. Due to the $(m_t/m_W)^2$ dependence in Equation B.16, however, these contributions are greatly suppressed, by powers of $(m_u/m_t)^2$ and $(m_c/m_t)^2$, respectively. In Equation B.16, terms involving $S_0(m_c^2/m_W^2)$ and $S_0(m_u^2/m_W^2)$ have been dropped.

We note that M_{12} contains a factor of m_W^2 , while Γ_{12} contains m_b^2 in its place. Since all of the other factors in the two expressions are roughly comparable, we therefore expect M_{12} to be much larger than Γ_{12} . This justifies the earlier assumption of taking the limit for $\Delta\Gamma \rightarrow 0$, or assuming negligible CP violation in mixing.

The neutral B meson system is similar to the familiar neutral kaon system, with the following important difference: in the kaon system the two mass eigen-

states have dramatically different lifetimes and nearly identical masses, while in the neutral B system the two mass eigenstates have nearly identical lifetimes and noticeably different masses. In the kaon system, one mass-eigenstate (which is approximately a CP-eigenstate) decays to two pions, while the other decays to three pions, so the wavefunction and phase-space differences result in quite different lifetimes, known as the “long” and “short” kaons, K_L and K_S . In the neutral B system, the two mass-eigenstates have essentially the same number of correct CP final states available to each of them, so they have a very small lifetime difference, and we refer to either the mass eigenstates, “heavy” and “light”, B_H and B_L , or the flavour eigenstates B^0 and \bar{B}^0 .

By carefully measuring the time-dependent distribution of neutral B decays, physicists can determine these values of Δm and M_{12} and thereby test the CKM matrix and the Standard Model. Measuring the ratio of Δm_d and Δm_s can remove some of the systematic uncertainties associated with the hadronic parameters, and provide tighter constraints on the consistency of the Standard Model, CKM Matrix picture of particle physics.

B.6 Measuring Δm_d

The mixing frequency of neutral B mesons can be expressed simply as Δm_d , in natural units (where $\hbar = c = 1$). This is the mass difference between the two mass eigenstates of the neutral B_d mesons, B_H and B_L , where the subscripts denote the heavier and lighter eigenstates, respectively. The subscript d denotes the oscillation frequency for mesons containing a b or \bar{b} quark, and a d or \bar{d} quark. Another current topic of interest in experimental particle physics is the measurement of the mixing frequency for B_s mesons, which contain s , or “strange” quarks or antiquarks, in place of the d quarks or antiquarks.

Figure B.3 lists the published measurements of Δm_d obtained from time-dependent measurements at the asymmetric energy B Factories, *BABAR* and *BELLE* [22]. All the published measurements are based on much less than the currently recorded data sets, due to the limiting systematic uncertainties (and manpower). These published results are discussed in more detail below. In order to substantially improve the precision of the world average value, an updated analysis requires better than 1% precision, in both the statistical and systematic uncertainties.

The *BELLE* experiment, which is in many ways very similar to *BABAR* has provided an important source of competition as well as an important cross-check for the consistency of physics quantities which are being measured for the first time in these B -factory experiments. The remainder of this thesis will only reference the technology and measurements of the *BABAR* experiment, which ran from 1999 until 2008.

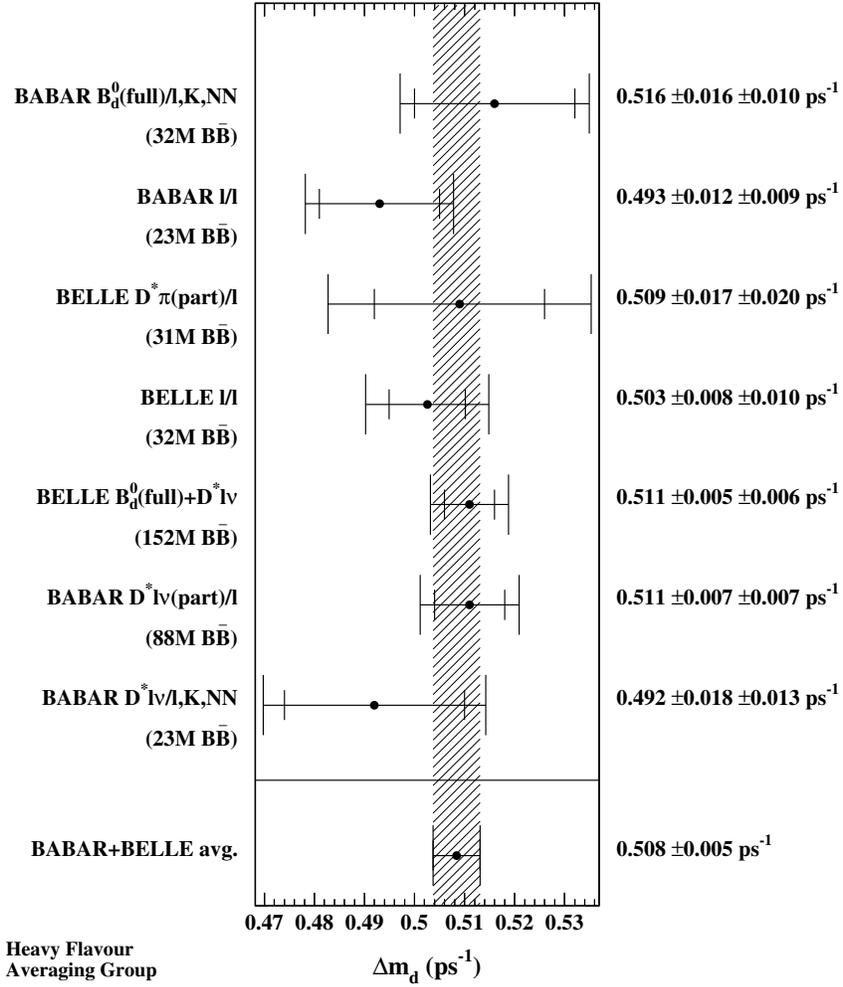


Figure B.3: Published measurements of Δm_d from the B Factories, $BABAR$, and $BELLE$. From <http://www.slac.stanford.edu/xorg/hfag/>

B.7 Measuring B Meson Decay Times

The cleanest possible process for producing pairs of B_d mesons is to collide electron-positron pairs at the resonant energy to produce the $\Upsilon(4S)$ meson. The $\Upsilon(4S)$ decays into a pair of B mesons, either $B^0\bar{B}^0$ or B^+B^- in nearly equal amounts. In the B Factory experiments, electrons and positrons are annihilated to produce the $\Upsilon(4S)$ meson, which decays almost instantly into a pair of B -mesons. The two B mesons are very nearly at rest in the centre-of-mass frame of the $\Upsilon(4S)$, with energies of roughly 340 MeV each. If the beam energies are initially equal, the $\Upsilon(4S)$ would be produced at rest in the lab frame, and the resulting B mesons would also be very nearly at rest in the lab. In this case, time-dependent analyses are not feasible because the B lifetime is too short, approximately 10^{-12} seconds. This decay time would only allow the B mesons to travel extremely short distances before decaying, and it would not be possible with current detector technology to resolve the two separate decays. At the PEP-II and KEKB colliders, however, electrons and positrons are collided with asymmetric or unequal beam energies. At PEP-II electrons were accelerated to an energy of 9 GeV, and positrons were accelerated to an energy of 3.1 GeV, before colliding. These collisions created a system with a CM (centre-of-mass) energy of $\sqrt{s} = 10.58$ GeV, corresponding to the mass of the $\Upsilon(4S)$ meson, and a Lorentz boost in the lab frame of $\beta\gamma \simeq 0.56$ because of the asymmetric beam energies. This means that the $\Upsilon(4S)$ and the B mesons are moving in directions very close to the beam axis, in the direction of the e^- beam, with a velocity of roughly one-half the speed of light.

This Lorentz boost converts timescales of order 10^{-12} seconds into distances on the order of 10^{-4} metres, or roughly 100 μm . This allows the two B meson decay vertices to be resolved with a suitably precise silicon vertex detector (SVT), located close enough to the collision point. Clearly, excellent position resolution along the beam axis of the detector is crucial to any time-dependent analysis. Furthermore, knowledge of this resolution often constitutes one of the dominant sources of systematic uncertainty in time-dependent *BABAR* results.

The decay time difference, Δt , is measured with the help of the asymmetric beam energies at PEP-II. The boost of the CM frame with respect to the lab frame allows us to convert the decay position difference, Δz , into a decay time difference, Δt , since the boost is very nearly along the z -axis of the detector. The so-called ‘‘boost approximation’’ calculation is given by:

$$\Delta t \simeq \Delta z / (\beta\gamma c) \tag{B.18}$$

There is a small systematic uncertainty introduced because the actual boost directions of the B mesons are not generally exactly along the z axis of the detector. This ‘‘boost smearing’’, and its dependence on the lepton momenta, is an important component of the experimental Δt resolution. In time-dependent asymmetry measurements the sign of Δt is somewhat arbitrary, and is chosen in our case to be the decay time of the higher-momentum lepton, minus the decay time of the lower-momentum lepton.

B.8 Current Measurements

There are three main techniques that have been used to make time-dependent B -mixing measurements at the B -Factory experiments, as shown in Figure B.3. The first method is using the very inclusive dilepton sample (denoted ℓ/ℓ in Figure B.3), which simply requires the presence of two high-energy leptons (electrons or muons) in the event. A large fraction of these leptons will originate from B meson decay, but there are substantial backgrounds due to leptons from charm decay and other processes. These measurements have the highest available statistics, but tend to be dominated by systematic uncertainties involved in describing the many contributions to the resolution of the lepton vertex measurements, and the different indirect sources of leptons.

The second technique is somewhat less inclusive, in which a high energy lepton or kaon is used to “tag”, or identify the weak charge of one B in the event, while another high energy lepton or pion, and a soft pion are used to form the vertex of a B meson decaying to a lepton, lepton anti-neutrino and D^* meson. The D^* is assumed to decay into a $D\pi$, giving the soft pion, which is used as a constraint on the vertex of the B decay. This method has fairly high statistics because there are relatively few requirements made on the final states of the events, but the use of the soft pion improves the resolution on the vertex, and greatly decreases the contributions of many background categories. This method is often limited by systematic uncertainties as well.

The third technique is exclusive, in that fully reconstructed B mesons are used. One B is tagged by a high energy kaon or lepton, to determine its flavour, and the other B is fully reconstructed in one of many decay modes. In this case, the events are very clean, with a small amount of background, and very precise determination of the decay vertices of the B mesons, but because of the difficulty involved with fully reconstructing a B meson, especially if there are neutrinos resulting from semileptonic decays, the statistics are rather low. These measurements are ultimately statistically limited.

The *BABAR* Collaboration has published papers using all three of the above methods. A paper was published in 2002 on the measurement of Δm for neutral B 's using the dilepton technique [1]. The data sample used at the time (consisted of 23M $B\bar{B}$ pairs), with no reconstruction of the B 's yielded the result:

$$\Delta m = 0.493 \pm 0.012 \text{ (stat)} \pm 0.009 \text{ (syst)} \text{ ps}^{-1} \quad (\text{B.19})$$

Table B.1 contains a list of the dominant sources of systematic uncertainty (in ps^{-1}), in the above published result. The systematic uncertainty due to B lifetimes is present because the final fit could not float both the neutral and charged B lifetimes. As the uncertainty on the world-average values for the lifetimes is decreased, this systematic should decrease as well. The error due to the “ Δt dependence of resolution” is essentially a case in which the analytical PDF describing the experimental resolution is not detailed enough at the level of precision achievable by the statistics. The “Other B Cascade” components represent one of the major sources of background in which one lepton is directly from the decay of a B meson, but the other lepton results from the cascade

Description	Uncertainty (ps^{-1})
B lifetimes	0.0064
Δt dependence of resolution	0.0043
Other B cascade resolution/lifetimes	0.0026
Other B cascade fractions, wrong tags	0.0020
z scale and SVT misalignment	0.0020
Total	0.0087

Table B.1: Systematic uncertainties in the published $BABAR$ result on mixing with dileptons [1].

decay of the other B , for example a charm meson, charm baryon, or τ resulting from the other B . Because of the many possible sources of these OBC leptons it is quite difficult to create an analytical model describing them all accurately, and experimentally determining the inclusive rates is difficult. The “ z scale and SVT misalignment” systematic uncertainty is due to the limits of how well the vertex detector is calibrated. Firstly, the distance between electronic sensors on the detector must be measured precisely through metrology of the detector. Secondly, the alignment of the multiple layers must be calibrated *in situ*, after the SVT is mounted in place. These z scale calibration uncertainties have since been reduced by more careful analysis within the $BABAR$ Collaboration.

For a discussion of the attempts to update the dilepton analysis of neutral B mixing, please refer to Chapter 3. This thesis describes the development of a new analysis tool to measure the B mixing frequency in the dilepton mode, intended for use on the full $BABAR$ data set.

Appendix C

The *BABAR* Detector

C.1 Overview

The *BABAR* detector was located at the PEP-II accelerator at the SLAC National Laboratory, in Stanford, California, shown schematically in Figure C.1. Beginning shortly after the experiment was shutdown in 2008, the long process of decommissioning and dismantling the detector began. This Chapter briefly describes the detector in its operational state, from 1999 to 2008.

The PEP-II accelerator accelerated electrons to an energy of 9 GeV, and positrons to an energy of 3.1 GeV, before colliding them in the interaction region IR2. These energies create a CM energy of 10.58 GeV, and a Lorentz boost in the lab frame of $\beta\gamma \simeq 0.58$. The *BABAR* detector, shown in Figure C.2, was

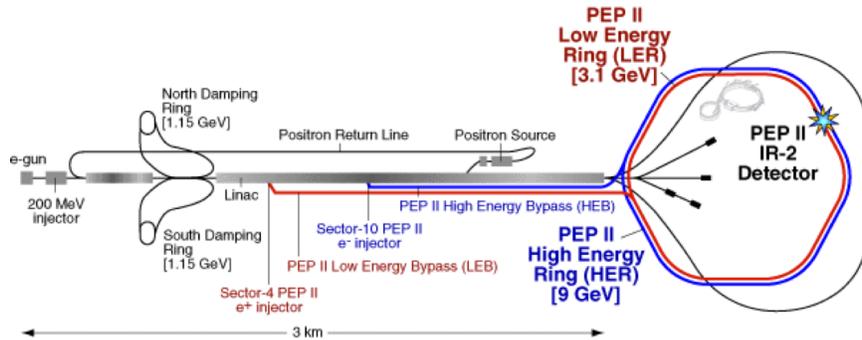


Figure C.1: Schematic view of the PEP-II accelerator at SLAC.

comprised of five subsystems. The innermost was the Silicon Vertex Tracker, based on silicon micro-strip detectors. Next, proceeding outward radially, was a multi-wire drift chamber filled with a Helium-Isobutane gas mixture which was used for charged particle tracking and dE/dx measurements, followed by a detector of internally-reflected Cherenkov light used for particle identification. An electromagnetic calorimeter, composed of CsI(Tl) crystals, measured the photon energies and angles. A 1.5T superconducting magnet enclosed the inner four detector layers and the axial magnetic field forced charged particles to follow curved trajectories in the detector. The outer part of the detector, the instrumented magnetic flux return was made of interleaved iron absorber plates

and detector elements used to detect neutral hadrons and identify muons by penetration. Measurement of the track curvature determined the track momentum and combining the momentum with the velocity measured by the DIRC allowed the determination of the particle mass and species. Full details of the *BABAR* detector design and its performance can be found in [33], and an updated publication which is forthcoming from the *BABAR* Collaboration.

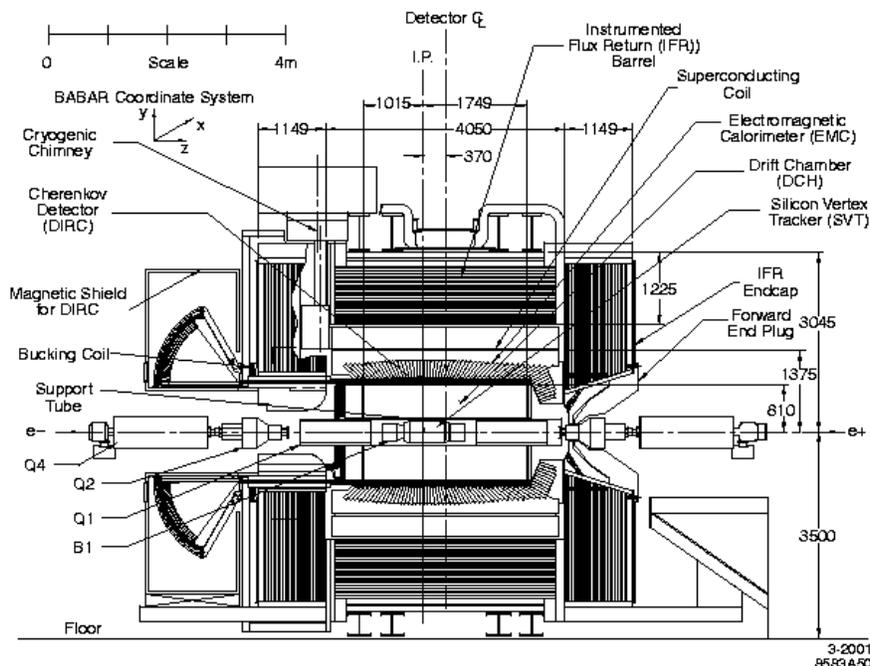


Figure C.2: Cross-sectional view of the *BABAR* detector at the PEP-II accelerator. All dimensions are in mm.

C.2 Silicon Vertex Tracker

The micro-strip silicon vertex tracker (SVT) is composed of five double-sided layers. The layers are segmented into strips used to determine a passing particle's location along both the z and ϕ axes. In total, there are roughly 150,000 individual strips read out from the SVT. When a charged particle passes through the active regions of silicon, it deposits energy which leads to the creation of electron-hole pairs. These pairs are swept apart by an applied voltage and the resulting current pulse is detected as a "hit" in one of the strips. The locations of hits are combined to infer the location of the track at each layer of the SVT.

When attempting to measure the vertex of dilepton events, the resolution is primarily determined by the SVT, since the hits in the silicon are the clos-

est points along the track to the beamspot and their spatial resolution is the best. For tracks with high transverse momentum (p_T), the SVT can resolve the point of closest approach to the beamspot to within $50\ \mu\text{m}$, and sometimes as accurately as $10\ \mu\text{m}$. This resolution is primarily determined by the following factors: the strip size in the SVT, the angle at which the track traverses the SVT layers and the track momentum. Values are calculated for each track as they are reconstructed in the software, providing resolution estimates on a track by track basis. In addition, the alignment of the SVT with respect to the beams is calibrated on a continual basis, and contributes another systematic uncertainty to the vertex measurements.

C.3 Drift Chamber

The drift chamber (DCH) is filled with a mixture of helium and isobutane (C_4H_{10}) gases kept at 4 mbar above atmospheric pressure. The sense wires are kept at high-voltages ranging between 1900–2000V. There are a total of 28,768 wires, and they are arranged into forty different concentric layers. Twenty-four of the forty layers are placed at a small angle relative to the z axis to measure the θ angle of passing tracks.

When a charged particle passes through the DCH, some of the gas is ionized along the path of the track and the ionized gas molecules are attracted to the ground wires while the liberated electrons are drawn toward the sense wires. The electrons collide frequently with gas molecules on their way to the sense wire, resulting in a constant drift velocity. When the electrons are very close to the wires, the increased electric field gives them enough energy to knock loose additional electrons from each gas atom. This creates a large gain in the signal picked up by the sense wire, on the order of 50,000 for this particular gas mixture and voltage setting. The resulting voltage pulses on the sense wires are read out and their arrival times are converted into distances from the wire, using the constant drift velocity. The position resolution ranges between 0.1 and 0.4 mm, depending on the distance from the wire the original track was located. This resolution, and the effects of Multiple Coulomb Scattering (MCS) in the detector are ultimately the limiting factors in determining how well the track curvature parameter ω , and hence the momentum, can be measured.

The other important measurement made by the DCH is the energy loss per unit distance, or dE/dx . The total charge deposited in each cell of the chamber is used to calculate the total energy loss over the entire chamber width. There are a number of corrections applied to improve the measurement accuracy, including continuous updating of the gas pressure, temperature and gain, wire-by-wire geometrical constants and the variation of the energy loss with the angle of the incident particle. The accuracy of dE/dx measurements is roughly 7%. The value of dE/dx for a track depends only on the speed of the particle, not its momentum, and it can therefore be used with measurements of the momentum to determine the particle mass. The measurement of dE/dx within the DCH is used for particle identification in conjunction with the Detector of Internally

Reflected Cherenkov Light (DIRC), or by itself for tracks which are within the acceptance of the DCH but miss the DIRC.

C.4 Detector of Internally Reflected Cherenkov Light

A charged particle travelling with velocity $\beta = v/c$ in a medium of refractive index n produces Cherenkov light if $n\beta \geq 1$. For relativistic charged particles, a cone of Cherenkov radiation is emitted with a characteristic angle given by:

$$\cos(\theta_c) = 1/(n\beta) \quad (\text{C.1})$$

Measurement of the Cherenkov angle, in conjunction with knowing the particle momentum from the drift chamber, allows a determination of the particle mass and hence, the particle type.

The DIRC sub-detector is a novel device used for the first time in the *BABAR* experiment. It is a ring imaging Cherenkov detector based on total internal reflection and uses fused silica bars as both the radiator and the light guide. The DIRC Cherenkov radiators are 4.9 m long rectangular fused silica bars oriented parallel to the z axis of the detector. The fused silica has an index of refraction $n=1.473$. Through internal reflections, the Cherenkov light from the passage of a particle through the DIRC is carried to the ends of the bar, as shown in Figure C.3. This radiation is captured by internal reflection in the bar and transmitted to the photon detector array located at the backward end of the detector (forward-going light is reflected by a mirror located at the end of the bar). The high optical quality of the fused silica preserves the angle of the emitted Cherenkov light during total internal reflection.

An advantage of the DIRC for an asymmetric collider, like PEP-II, is that the high momentum tracks are boosted forward, causing a much higher light yield than for particles at normal incidence. This is due to two effects: the longer path length in the fused silica and a larger fraction of the produced light being internally reflected in the bar.

The photon detector array consists of about 11,000 conventional 2.5 cm-diameter photomultiplier tubes. They are organized in a close-packed array at a distance of about 120 cm from the end of the radiator bars. The photo-tubes, together with modular bases, are located in a gas-tight volume as protection against helium leaks from the drift chamber.

To maintain good photon transmission for all track dip angles, the standoff region is filled with water to minimize the change in the index of refraction. The water seal occurs at a fused silica window that is glued to the fused silica wedges. The standoff box is surrounded by a steel box which provides adequate magnetic shielding for the photo-tubes.

The main design goal for the DIRC was to be able to distinguish pions and kaons at momenta greater than 2 GeV/c. The statistical separation between the two species ranges from about 10σ at 2 GeV/c to approximately 3σ at 4 GeV/c, easily meeting the specified design requirements.

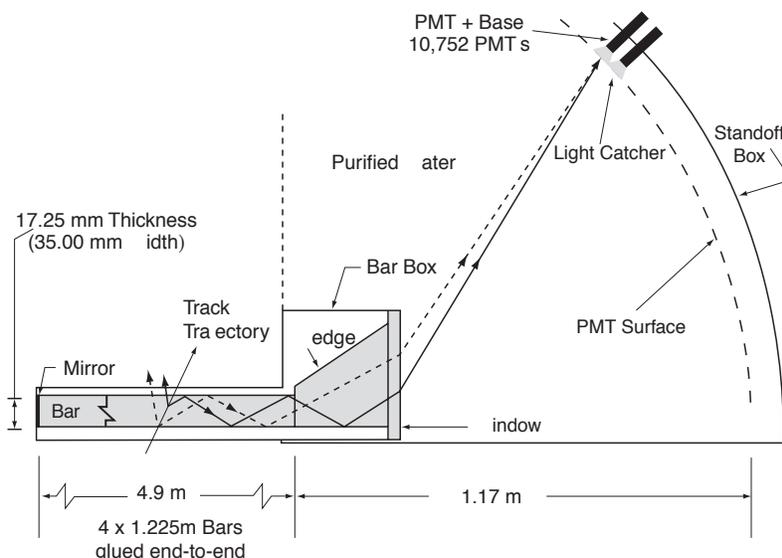


Figure C.3: Cross-sectional view of a fused silica radiator in the DIRC, showing the internal reflections of the Cherenkov photons.

C.5 Electromagnetic Calorimeter

The Electromagnetic Calorimeter (EMC) was designed to measure electromagnetic showers with high efficiency, high angular resolution, and high energy resolution over the range from 20 MeV to 9 GeV. The lower energy limit is determined by the need to reconstruct π^0 's resulting from B decays, and the higher energy limit is necessary to measure Bhabha-scattered electrons and muons. The EMC is also helpful in identifying electrons.

This sub-detector is divided into two parts, a cylindrical barrel surrounding the inner detector systems, and a conical forward end-cap. The barrel is composed of 5,880 Thallium-doped Caesium-Iodide crystals, arranged into 280 separate modules, while the end-cap contains 20 modules of 41 crystals each. CsI(Tl) was chosen for the crystals primarily because of its high light yield and small Moliere radius which allow for excellent energy and angular resolution.

High-energy photons and electrons travelling through the material of the calorimeter create electromagnetic showers and the resulting photons are absorbed by the crystals and re-emitted as visible light. These scintillation photons are then detected by high-speed silicon photo-diodes mounted on the outer surfaces of the crystals. Silicon photo-diodes were chosen because of their ability to function properly inside the 1.5T magnetic field of the *BABAR* solenoid and because it is possible to match their absorption spectrum to that of the crystal light output.

The measured energy resolution of the EMC ranges from 5% with a 6.13 MeV radioactive source, down to 1.9% for Bhabha-scattered electrons at 7.5 GeV. The angular resolution varies from 12 mrad at low energies to 3 mrad at high energies. Finally, the energy measured in the EMC is used with the momentum determined by the drift chamber in order to identify electrons with an efficiency of between 88–94%, and a pion misidentification probability ranging from 0.15–0.3%.

C.6 Instrumented Flux Return

The Instrumented Flux Return (IFR) was designed to identify muons with high efficiency and low contamination, as well as to detect long-lived neutral hadrons over a wide range of angles and momenta. The main objectives for the IFR are large coverage angles, good efficiency and good rejection of muon background down to momenta lower than 1 GeV/c. For the neutral hadrons, high efficiency of detection and good angular resolution are the most important criteria.

The IFR is divided into three separate sections: a cylindrical barrel and forward and backward end-caps. Each section is composed of layers of iron absorber material interspersed with active detector layers. The iron absorbs energy from highly-penetrating muons and long-lived neutrals, as well as acting as the flux return for the solenoidal superconducting magnet. The number of layers of steel, and their thickness, were chosen using simulations to optimize the detection of muons and neutral hadrons. The steel absorber consists of 18 steel plates with thicknesses ranging from 2–10 cm. The active detector layers were originally Resistive Plate Chambers (RPCs) [34], that provide two-dimensional coordinates and precise timing information for passing particles. The barrel originally contained 19 RPC layers and the end-caps had 18 active layers.

An RPC consists of two bakelite plastic plates separated by a 2 mm gap containing a gas mixture of argon, freon and isobutane. The bakelite has a volume resistivity of approximately $10^{11} - 10^{12} \Omega\text{cm}$. The inner surfaces of the gas-filled gap are coated with linseed oil in order to reduce high-voltage breakdown. A voltage of roughly 8 kV is applied between the two plates and the passage of a particle leads to the creation of a limited streamer between the plates. The plates are read capacitively with strips running along both the z axis and in the ϕ direction. There are a total of approximately 50,000 channels.

The efficiency of individual RPC modules can be determined using cosmic rays. Initially more than 75% of the modules had an efficiency greater than 90%, but at least 50 modules had already failed by the summer of 1999. Between 1999 and 2004 the overall efficiency of this detector system continued to degrade from 98% to roughly 75%. The angular resolution of the system continued to meet the design criteria.

Due to the continuing problems with RPCs, the *BABAR* collaboration decided in 2004 to undertake their replacement with another technology known as Limited Streamer Tubes (LSTs) [35]. These use extruded plastic tubes with wires inside, instead of the flat gaps. They have a long successful history, and

are manufactured with a high degree of quality and reliability. The first two sextants of the IFR were disassembled and the LSTs were installed during the summer shutdown of 2004. The rest of the RPC modules in the barrel were replaced during the summer shutdown of 2006. The LST modules performed well until the end of data-taking.

The IFR is primarily used for particle identification of high energy muons and K_L^0 s. The efficiency of muon ID continually decreased over the first five years due to the degradation of the RPCs, but a large amount of data was taken after the LST installations, in data Runs 5 and 6. The efficiency of muon identification affects the total statistics available for the measurement in this thesis, but does not directly affect the mixing measurement itself, as it would for a branching fraction measurement, or in a search for a rare state or decay including muons.

Appendix D

Personal Contributions to the *BABAR* Collaboration

This appendix describes the various forms of service work that I did over the years for the *BABAR* Collaboration.

I joined the *BABAR* Collaboration as an “Associate Member” in September, 2002, when I began my MSc. in the UBC Department of Physics and Astronomy. Associate members are able to use the computing resources and access the data of the experiment, but do not appear on the published author list.

My first task as a member of the collaboration was the reviewing of many *BABAR* journal article drafts before their submission for publication. This began in 2003, and has continued until today. Each draft journal article is sent out for review to approximately 10–12 member institutions, who are encouraged to comment on both the analysis and results contained in the article, as well as spelling and grammar. As part of this process, I have read and commented on approximately 70 journal articles, which have appeared primarily in Physical Review Letters and Physical Review D. This allowed me to gain exposure to a wide range of particle physics topics which are studied using the data recorded by *BABAR*, learn a great deal about how to summarize a complicated analysis in a concise manner, and also to perform a helpful service for my colleagues.

The second task that I undertook began in 2004 and continued until mid-2006. I was in charge of the production of simulated events (Monte-Carlo Simulation) on the WestGrid and UVic computing clusters. This work consisted of daily monitoring, job submission, checking for success or failure, running scripts for archiving and transferring results from the local computer systems I used to SLAC. Several times over this period I was asked to upgrade all of the installed *BABAR* software necessary to run the simulations, which was an intensive process requiring nearly two days. Finally I also worked with several experts to try and track down system performance issues, and maximize the throughput of the computing systems I was using.

The third major task on which I worked was a detailed study of the efficiency of the *BABAR* Drift Chamber tracking system, using a carefully selected data control sample. By selecting nearly ideal, two-track events, I could compare the number of hits recorded and associated to the track, with the theoretical maximum number with perfect efficiency. I studied the different sources of “missing” hits, at both early and late stages in the reconstruction process, and performed measurements of these various sources of inefficiency as functions

of the collider beam currents and instantaneous luminosity. A question under consideration at the time, was whether or not planned increases in machine luminosity up to 10^{34} cm²/s would have a significant impact on the performance of the drift chamber. The answer turned out to be no, my measured trends versus luminosity indicated that the loss of hits in the tracking system would continue to be tolerable, even with a doubling of the instantaneous luminosity, and the associated increase in background.

The fourth task that I performed was a performance analysis (both of the quality of output, as well as computing resources) of modified pattern recognition algorithms for the reconstruction of tracks in the drift chamber. The pattern recognition software was originally transliterated from FORTRAN code that had been used on an earlier experiment, and it was discovered after several years of use that the pattern recognition code had several long-standing mistakes in it, which represented potential track hit patterns that had not been included in the hard-coded list of search patterns. The programmer who worked on this code also tried to make other performance improvements, by allowing the hits that were dropped from one track to be re-used for subsequent tracks, among other changes to the algorithm. I built the *BABAR* data analysis executable software with all of the different code changes included, and used run-time configuration flags to turn on and off the different modifications of the tracking software. I then ran a number of control data samples, and carefully compared the output data, as well as the computing resources used on a standard type of machine for the different configurations. From these studies we were able to select a subset of the code changes which provided noticeable improvements in the tracking performance, for a modest increase in computing requirements.

Finally, as with all other members of the *BABAR* Collaboration, I performed numerous Detector Control Room shifts as both the real-time Data-Quality Expert (navigator in *BABAR* parlance), and the Shift Leader (pilot), in 2006 and 2007. One of the highlights of these shifts was a flood which occurred during a rainstorm in a room full of high-voltage racks powering accelerator components near the interaction region. Another was an unexpected trip of the solenoid magnet supplies, causing a quench of the magnet. The experiment finished data-taking in 2008, approximately 4 months after my final shifts, and has since been decommissioned and disassembled.