Beauty Production at HERA – B

Performance of the Second Level Trigger and the $b\bar{b}$ Cross Section in 920 GeV pN Collisions

Thesis submitted for the degree of
Doctor of Philosophy in Physics
by

Cand. Scient
Brian Aagaard Petersen

May 2002

Niels Bohr Institute
University of Copenhagen

Abstract

HERA – B is a fixed target experiment, which uses the 920 GeV/c proton beam from the HERA accelerator at DESY to produce beauty hadrons. Beauty hadrons are only produced in about one in a million of the interactions and HERA – B is therefore a finely segmented spectrometer with large geometrical acceptance capable of running at high interaction rates. To select the events containing beauty hadrons out of the large background, a complex multi-level trigger system has been implemented which reduces the event rate by more than 10^5 in realtime.

The main subject of this thesis is the Second Level Trigger (SLT) system. The SLT is a software trigger running on a large farm of commodity PCs. The trigger searches for charged tracks (mainly leptons) in search windows defined by the First Level Trigger and follows these tracks through the complete spectrometer. The full algorithm of the SLT is described in detail and the performance in both the 2000 commissioning run and the upcoming 2002 physics run is evaluated.

Based on data from the 2000 run, the production cross section of beauty hadrons has been measured. The analysis uses inclusive beauty decays into J/ψ mesons. The J/ψ mesons are reconstructed in the $\mu^+\mu^-$ and e^+e^- decay channels and identified as beauty hadrons by their detached decay vertex with respect to the primary interaction. The analysis yields a total cross section of: $\sigma(b\bar{b}) = 32^{+14}_{-12\,stat} \pm 8_{sys}$ nb/nucleon.

Contents

1 Physics motivation **1**

1.1 Symmetries and CP violation . . . 1
1.2 CP violation in the Standard Model . . . 2
1.2.1 CKM matrix parameterizations . . . 3
1.2.2 Unitarity triangles . . . 4
1.3 Beauty hadrons and CP violation . . . 5
1.3.1 Notations . . . 5
1.3.2 The three types of CP violation . . . 8
1.4 The golden decay . . . 10
1.4.1 Measuring CP violation in the golden decay . . . 11
1.4.2 Current status of CP violation in the golden decay measurement . . . 12
1.5 Beauty production in hadronic collisions . . . 13

2 The HERA – B experiment **17**

2.1 Overview . . . 17
2.2 The HERA storage ring . . . 18
2.3 The target system . . . 19
2.4 The vertex detector . . . 21
2.5 The inner tracker . . . 22
2.6 Outer tracker . . . 23
2.7 Ring imaging Cherenkov detector . . . 24
2.8 Transition radiation detector . . . 25
2.9 Electromagnetic calorimeter . . . 25
2.10 Muon detector . . . 27
2.11 High-p_T chambers . . . 28

3 The HERA – B data acquisition and trigger system **29**

3.1 Trigger overview . . . 29
3.2 The pretriggers . . . 31
3.2.1 The MUON pretrigger . . . 32
3.2.2 The ECAL pretrigger . . . 32
3.2.3 The High-p_T pretrigger . . . 33
3.3 The first level trigger . . . 34
3.4 The second level trigger . . . 37
3.5 The third level trigger . . . 38

3.6 Fourth level trigger . . . 39
3.7 Data acquisition system . . . 39
3.7.1 Front end electronics . . . 40
3.7.2 Second level buffer and switch . . . 42
3.7.3 4LT and datalogging . . . 45
3.7.4 Deadtime . . . 45

4 SLT algorithm for the 2000 run 48
4.1 SLT configurations in year 2000 . . . 48
4.2 SLT Seeding . . . 49
4.2.1 Electron pretrigger seeding . . . 50
4.2.2 Muon pretrigger seeding . . . 52
4.2.3 FLT seeding . . . 54
4.3 Main tracker algorithm . . . 55
4.3.1 Slicer . . . 55
4.3.2 Refit . . . 57
4.4 Magnet propagation . . . 59
4.5 Vertex detector tracking . . . 63

5 SLT performance in the 2000 run 66
5.1 Detector performance . . . 67
5.2 ECAL+L2Sili trigger . . . 68
5.2.1 Trigger efficiency . . . 68
5.2.2 Tracking precision . . . 76
5.2.3 Trigger rates . . . 77
5.2.4 Timing . . . 79
5.2.5 Fake tracks . . . 81
5.2.6 Summary of the ECAL+L2Sili trigger performance . . . 84
5.3 ECAL+SLT trigger . . . 85
5.3.1 Trigger efficiency . . . 85
5.3.2 Tracking precision . . . 91
5.3.3 Trigger rates . . . 94
5.3.4 Timing . . . 95
5.3.5 Fake tracks . . . 96
5.3.6 Summary of the ECAL+SLT trigger performance . . . 98
5.4 MUON+SLT trigger . . . 98
5.4.1 Trigger Efficiency . . . 98
5.4.2 Tracking precision . . . 106
5.4.3 Trigger rates . . . 107
5.4.4 Timing . . . 108
5.4.5 Fake tracks . . . 109
5.4.6 Summary of the SLT+MUON trigger performance . . . 111
5.5 Summary of the 2000 SLT trigger performance . . . 112

6 The SLT algorithm for the 2002 run **113**
6.1 Detector upgrades/repairs . . . 114
6.2 Overview of the 2002 SLT trigger . . . 116
6.3 Seeding . . . 118
6.3.1 FLT seeding . . . 118
6.3.2 ECAL pretriggered seeds . . . 118
6.3.3 Muon pretriggered seeds . . . 118
6.3.4 Seed pairing . . . 119
6.4 Slicer/Refit . . . 120
6.5 L2Magnet . . . 121
6.6 L2Sili . . . 122
6.6.1 Comparison of the 2002 L2Sili version to the 2000 version . . . 124
6.7 L2Vertex . . . 125

7 Expected performance in 2002 **129**
7.1 Trigger rates . . . 130
7.1.1 FLT output rates . . . 130
7.1.2 SLT rates . . . 132
7.2 Timing . . . 134
7.3 Trigger efficiency . . . 136
7.4 Tracking precision . . . 140
7.5 Summary . . . 141

8 Measurement of $\sigma(b\bar{b})$ from the 2000 data **144**
8.1 The principle of the measurement . . . 144
8.1.1 Calculating $\sigma(b\bar{b})$ . . . 145
8.1.2 Results from other experiments used by the analysis . . . 147
8.2 Data sample . . . 148
8.3 Monte Carlo simulation . . . 149
8.3.1 The event generators . . . 149
8.3.2 Prompt produced J/ψ events . . . 149
8.3.3 Simulated $b \to J/\psi X$ events . . . 150
8.3.4 Monte Carlo samples . . . 151
8.4 Muon analysis . . . 152
8.4.1 Prompt $J/\psi \to \mu^+\mu^-$ selection . . . 152
8.4.2 Detached $J/\psi \to \mu^+\mu^-$ selection . . . 154
8.4.3 $b\bar{b}$ cross section determination in the $\mu^+\mu^-$ channel . . . 162
8.5 Electron analysis . . . 164
8.5.1 Prompt $J/\psi \to e^+e^-$ selection . . . 165
8.5.2 Detached $J/\psi \to e^+e^-$ selection . . . 169
8.5.3 $b\bar{b}$ cross section determination in the e^+e^- channel . . . 171
8.6 Results . . . 172
8.6.1 Systematic uncertainties . . . 172
8.7 Comparison with other results . . . 176

9 Summary **178**

A Track reconstruction with Kalman filters **179**
A.1 The filtering process 180
A.2 The smoothing process 182
A.3 Kalman filter in one dimension 182

Introduction

The Standard Model of particle physics has so far been a very successful theory. It is able to explain all measurements[1] and only one particle, the Higgs, remains to be found experimentally. Still most current experiments are in one way or another searching for phenomena, which can not be explained by the Standard Model. The hope is that discovery of "new physics" will give hints to what happens at higher energy scales where the Standard Model is known to break down.

The field of "CP violation" had few experimental constraints until recently as it was observed only in neutral kaon mesons. CP violation is the breaking in weak interactions of the combined symmetry of Charge-conjugation and Parity (coordinate inversion). In the Standard Model it is described by a single parameter and a wide range of effects in beauty mesons can be predicted. In particular the "golden decay" $B_d^0 \to J/\psi K_S^0$ is predicted to have a large CP violating asymmetry in its decay rate, which can be cleanly related to the Standard Model parameter.

In the beginning of the 1990's, three experiments were therefore proposed and approved to do measurements in beauty decays and in particular to search for CP violating effects. Two of the experiments, BaBar and Belle, use e^+e^- beams to produce $\Upsilon(4S)$ which decays into beauty mesons, while the third, HERA – B, uses hadronic collisions to produce the beauty mesons. In the e^+e^- experiments the major technological challenge is the accelerator, which has to provide high luminosities and collide beams of different energies, while in HERA – B the main challenge is to construct a detector capable of detecting the beauty decays in a dense hadronic environment.

HERA – B is a fixed target experiment at the HERA accelerator at DESY, Hamburg. Only one out of about a million inelastic interactions contains beauty mesons and of these only one in 10^4 to 10^5 has reconstructable decay modes where large CP violating effects are expected. HERA – B therefore has to have an interaction rate of about 40 MHz in order to be able to observe CP violation. As it has to be able to fully reconstruct beauty meson decays efficiently, the detector needs to be finely segmented with high hit efficiency and able to withstand very high radiation doses for several years. At the same time only a small fraction of the many interactions can be read out and stored for later physics analysis. Therefore a complex multi-level trigger system[2] has been designed to reduce the rate by more than five orders of magnitude. All of these requirements have proven more difficult to fulfill than anticipated and have caused both delays and reduction in the HERA – B physics program.

I joined HERA – B in the summer of 1999 and the first year I was stationed at the laboratory in Hamburg, working for the Second Level Trigger (SLT) group. During this time the last elements of the detector were being installed and the complete detector was being commissioned. My work was focused on getting the SLT to work online, adjusting it to the rapidly evolving

[1] One has to allow neutrinos to have mass, as neutrino oscillations now seem to be well established.

[2] Trigger systems are often also known as event filters. Throughout this thesis we will use the word "trigger".

detector and especially evaluating the performance of the trigger.

The accelerator was shutdown for a luminosity upgrade at the end of August 2000 for more than one year. During this time HERA – B has undergone a major upgrade and repair program and I have helped improve the SLT algorithm for the 2002 run based on the experience gained from the data taking in 2000. At the same time I have been heavily involved in the physics analysis of the data recorded during the last months of the 2000 run. In particular I worked on the measurement of beauty hadron production cross section, where beauty hadrons are reconstructed inclusively in decay modes into J/ψ mesons.

The thesis is structured as follows: Chapter 1 gives an introduction to CP violation and beauty mesons followed in Chapter 2 and 3 by a description of the experiment and the trigger system. Chapters 4–7 give a detailed description and performance evaluation of the SLT algorithm for both the 2000 and 2002 runs. In Chapter 8 the measurement of beauty production is presented.

Chapter 1

Physics motivation

HERA – B was proposed, designed and built to perform measurements in the field of beauty meson decays. In particular, it was meant to be the first experiment to observe CP violation in the so-called golden decay, $B_d^0 \rightarrow J/\psi K_S^0$. However HERA – B has been delayed for various reasons and in the meantime other experiments [1, 2] have observed CP violation in the golden decay. The physics goals of HERA – B have therefore been changed significantly. To understand the design of the HERA – B detector and trigger systems it is however useful with a short introduction to CP violation and the golden decay. This is the main purpose of this chapter. Most of the information presented in this chapter is based on Ref. [3]. A more thorough introduction to CP violation and beauty physics in general can be found also in Ref. [4].

At the end of the chapter an introduction to beauty production in hadronic collisions will also be given, as this is one of the main physics items in the upcoming HERA – B 2002 run and the subject of the analysis of HERA – B 2000 data presented in Chapter 8.

1.1 Symmetries and CP violation

Symmetry is one of the fundamental concepts of physics. A symmetry is a transformation applied to a physical system, which leaves the system itself unchanged. An example of symmetry is space translation. Symmetries are connected to conservation laws through the Noether Theorem [5], which states that to each symmetry a conserved quantity can be associated, for instance the symmetry under space translations leads to conservation of momentum. Symmetries can be very helpful to predict physical processes, for instance by forbidding certain processes to occur and forcing others to occur at equal rates. Some symmetries are exact, meaning that they hold under all circumstances, such as space translations, others can be violated and therefore only valid under certain circumstances. The CP transformation is an example of the latter. In strong and electromagnetic interactions CP has been observed to be a valid symmetry, while in weak interaction violations have been seen.

The CP transformation is in reality the combination of two discrete transformations: C and P. These are closely connected to a third discrete transformation, T. These three letters are short hand notation for the following transformations:

- **Charge conjugation (C):** Interchanging the particles with their anti-particles.
- **Parity transformation (P):** Inversion of spatial coordinates: $\vec{x} \rightarrow -\vec{x}$. This also inverts momentum, while axial vectors, such as angular momentum, are left unchanged.

- **Time reversal (T):** Inversion of time: $t \rightarrow -t$. This inverts both momentum and angular momentum.

None of the three transformations are symmetries in weak interactions. A fundamental result of field theory [6] however states that the combined transformation, CPT, is always an exact symmetry in any theory that can be described by a local, causal relativistic field theory. This is true for the Standard Model of particle physics and most extensions of it. Some string models exist where CPT is no longer an exact symmetry [7], but this possibility will not be considered further here. As a result of the CPT theorem, all fundamental particles have the same mass and lifetime as their anti-particle.

C and P transformation were predicted not to be valid symmetries in weak interactions in 1956 [8] and one year later violations of both were observed. Parity violation was the first to be observed using beta decays of spin-aligned ^{60}Co atoms [9]. Both "symmetries" are said to be maximally violated in weak interactions as both can transform commonly observed interactions into interactions never observed. In the combined transformation, CP, the symmetry breaking is much weaker.

After the discovery of C and P violation one started searching for CP violations that were found in 1964 in decays of neutral kaons [10]. The first CP violation observed was K_L^0 decaying into two charged pions. This decay mode would not be allowed if CP was a valid symmetry, but it has been seen with a branching ratio of a few permille, indicating a small violation of CP. Later, more CP violating effects have been observed in the neutral kaon system, for instance in semileptonic decays. If CP was a valid symmetry, K_L^0 should decay into e^+ and e^- equally often, but measurements give:

$$\delta = \frac{\Gamma(K_L^0 \rightarrow e^+\nu_e\pi^-) - \Gamma(K_L^0 \rightarrow e^-\bar{\nu}_e\pi^+)}{\Gamma(K_L^0 \rightarrow e^+\nu_e\pi^-) + \Gamma(K_L^0 \rightarrow e^-\bar{\nu}_e\pi^+)} = (3.27 \pm 0.12) \cdot 10^{-3} \tag{1.1}$$

All CP violating asymmetries observed in the kaon system are quite small, of order 10^{-3} or even smaller. If CPT is a valid symmetry, the observation of CP violation implies that the time transformation can also not be a symmetry in weak decays. The first T violating asymmetries have recently been observed in the neutral kaon system.

The observation of CP violation is interesting for several reasons. Several theories of CP violations were suggested after its discovery, but only one has so far been able to explain the observations. This was the proposal of a third quark family which gives rise to a complex phase in the weak interactions [11]. The third quark family was later discovered directly (b and t quarks). CP violation is also one of the three Sakharov conditions [12] needed to explain the baryon asymmetry of the universe. The other conditions are baryon number violation and absence from thermal equilibrium. It turns out that the CP violation given by the Standard Model is not enough to explain the observed baryons asymmetry. It is therefore a strong motivation for testing the predictions of CP violation from the Standard Model, for instance by measuring CP violation in beauty decays.

1.2 CP violation in the Standard Model

In the Standard Model weak interactions are mediated by three heavy vector bosons, $W^{\pm}$ and Z^0, but only charged boson exchange can give rise to CP violation. The vector bosons do not couple to the mass eigenstates of quarks, but rather to a linear combination of these. There

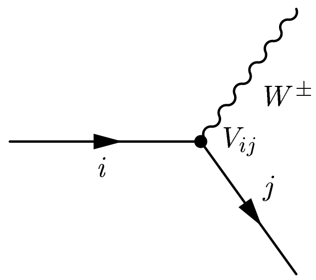

FIGURE 1.1: The weak interaction of two quarks i and j.

is quite some freedom in the choice of the linear combinations to use as basis for the weak eigenstates of the quarks, but, by convention, the up-type quarks are left unchanged and the quarks are sorted according to increasing mass. The weak eigenstates are then related to the mass eigenstates through the unitary Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{CKM} \begin{pmatrix} d \\ s \\ b \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{1.2}$$

The CKM matrix couples down-type quarks (d, s and b) to all up-type quarks (u, c and t) in the charged weak interactions with varying amplitudes as illustrated in Figure 1.1. This coupling for instance allows an s quark to decay into a u quark. If the CKM matrix contains complex phases a CP violating asymmetry can appear.

1.2.1 CKM matrix parameterizations

The CKM matrix represents a transformation between two complete bases and is therefore an unitary matrix. If a fourth quark generation exists, this is no longer true and deviations from unitarity could potentially be measured. A general unitary matrix can be described by three angles and six phases. One can freely choose the phases of the six quark fields and, by choosing the phase differences between the quarks correctly, one can remove five of the six phases from the CKM matrix. The CKM matrix can therefore be described by three angles and a single phase. The standard parameterization is [13]:

$$V_{CKM} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{13}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{13}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{13}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{13}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{13}} & c_{23}c_{13} \end{pmatrix}, \tag{1.3}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. The complex phase δ_{13} can be shown [3] to be the only cause of CP violation in the Standard Model. If δ_{13} was 0 or π, there would be no CP violation. If there were only two quark generations, the CKM matrix would be reduced to an unitary 2×2 matrix, which can be described by a single angle. The two quark generation case would not allow CP violation and this is what prompted Kobayashi and Maskawa to propose the third quark generation, before it was actually observed.

Most of the off-diagonal elements in the CKM matrix are small. It is often useful to use the so-called Wolfenstein parameterization, which expands the CKM matrix in terms of $\lambda = |V_{us}|$:

$$V_{CKM} = \begin{pmatrix} 1-\frac{\lambda^2}{2} & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1-\frac{\lambda^2}{2} & A\lambda^2 \\ A\lambda^3(1-\rho-i\eta) & -A\lambda^2 & 1 \end{pmatrix} + O(\lambda^4). \tag{1.4}$$

All four parameters, λ, A, ρ and η are real and current estimates based on many different measurements [14] are:

$$\begin{aligned} \lambda &= 0.2221 \pm 0.0021, && (1.5) \\ A &= 0.827 \pm 0.058, && (1.6) \\ \rho &= 0.23 \pm 0.11, && (1.7) \\ \eta &= 0.37 \pm 0.08. && (1.8) \\ & && (1.9) \end{aligned}$$

The values given for the four parameters are highly correlated.

1.2.2 Unitarity triangles

From the requirement of unitary of the CKM matrix we obtain nine equations:

$$\sum_{\alpha=u,c,t} V_{\alpha i} V^*_{\alpha i} = 1, \qquad i \in \{d,s,b\} \tag{1.10}$$

$$\sum_{\alpha=u,c,t} V_{\alpha i} V^*_{\alpha j} = 0, \qquad i,j \in \{d,s,b\} \text{ and } i \neq j \tag{1.11}$$

The left side of each equation is a sum of three complex numbers and the six equations contained in Equation 1.11, can be represented graphically as six triangles (unitarity triangles) in the complex plane. The advantage of this representation is that the angles in the triangles are independent of the parameterization used in the CKM matrix and thus represent real physical observables. The angles can be measured in CP violating processes and, if the Standard Model is correct, they should add up to 180°.

Most of these triangles have one side, which is much smaller (by λ^2 or more) than the two other sides and therefore one tiny angle, which is very difficult to measure. The triangle associated to the equation:

$$V_{ud}V^*_{ub} + V_{cd}V^*_{cb} + V_{td}V^*_{tb} = 0, \tag{1.12}$$

has sides of approximately the same length ($\sim \lambda^3$). This triangle is usually referred to as "The Unitarity Triangle", illustrated in Figure 1.2a. If all sides are divided with $|V_{cd}V^*_{cb}|$, the top of the Unitarity Triangle is (ρ, η) giving an easy relation of the different parameters. This relationship is illustrated in Figure 1.2b. The three angles of the triangle, defined in Figure 1.2, are related to the CKM matrix elements by:

$$\alpha = \arg\left(-\frac{V_{td}V^*_{tb}}{V_{ud}V^*_{ub}}\right), \quad \beta = \arg\left(-\frac{V_{cd}V^*_{cb}}{V_{td}V^*_{tb}}\right), \quad \gamma = \arg\left(-\frac{V_{ud}V^*_{ub}}{V_{cd}V^*_{cb}}\right). \tag{1.13}$$

All of angles contain factors relating to the b quark and α, β and γ can therefore be directly measured from CP violation in the beauty decays. The results of indirect measurements are

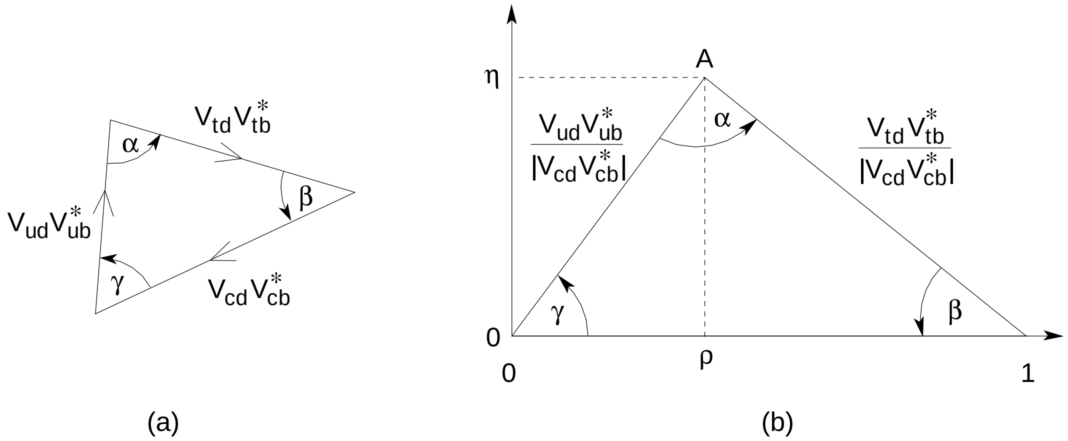

FIGURE 1.2: a) The Unitarity Triangle. b) The Unitarity Triangle rescaled to have two corners in $(0,0)$ respectively $(1,0)$ in the complex plane.

shown in the plane of the triangle in Figure 1.3. The goal of measuring CP violation is to over-constrain this triangle and if contradictions were to be found it would be a sign of new physics.

1.3 Beauty hadrons and CP violation

Beauty quarks (b quarks) are present in quite a few hadrons, both in mesons like B_d^0 and B^+ and in baryons like Λ_b. For measuring CP violation, the most important b hadrons are the two neutral b mesons, B_d^0 and B_s^0, which have a valence quark content of $(\bar{b}d)$ and $(\bar{b}s)$ respectively. For most of the discussion we shall not discriminate between the two mesons and simply denote them by B^0.

When studying these two mesons one has to distinguish between flavor eigenstates in which the meson has a specific quark content, such as $\bar{b}d$, and mass eigenstates, which have definite mass and lifetime. For the neutral b-mesons, the flavor and mass eigenstates are not the same, which means that a meson produced as B^0 can later decay as a $\overline{B}^0$. This B^0–$\overline{B}^0$ oscillation occurs mainly through double W exchange such as the box diagrams shown in Figure 1.4. The contributions from diagrams with top quark exchange dominate in the amplitudes. The same type of oscillation occurs also in kaon decays. Through the box diagrams, complex phase from the CKM matrix is introduced into the system, which makes neutral kaon and beauty mesons the simplest place to observe CP violation.

1.3.1 Notations

Given an arbitrary linear combination of B^0 flavor eigenstates:

$$|B(t)\rangle = a(t)|B^0\rangle + b(t)|\overline{B}^0\rangle, \tag{1.14}$$

the time evolution is given by a Schrödinger equation:

$$i\frac{d}{dt}\begin{pmatrix} a(t) \\ b(t) \end{pmatrix} = H\begin{pmatrix} a(t) \\ b(t) \end{pmatrix} = \left(M - \frac{i}{2}\Gamma\right)\begin{pmatrix} a(t) \\ b(t) \end{pmatrix}, \tag{1.15}$$

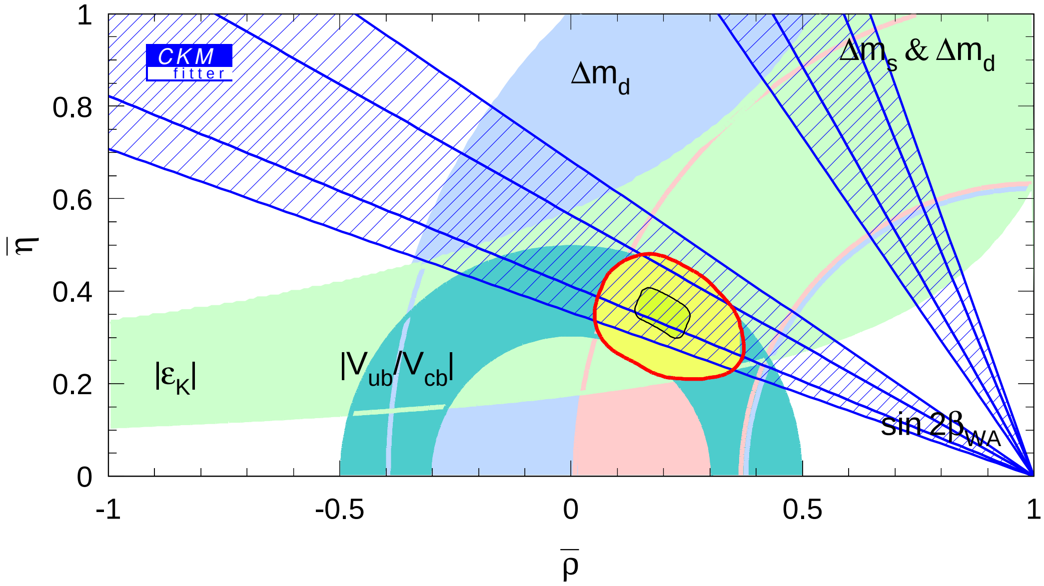

FIGURE 1.3: Constraints of the Unitarity triangle [14]. The fully colored bands represent 95% CL constraints from the different indirect measurements. The two circles represent the combined 68.3% and 95% CL regions for (ρ, η). The dashed areas are the 68.3% and 95% CL regions for the direct measurements of $\sin 2\beta$ from BaBar and Belle.

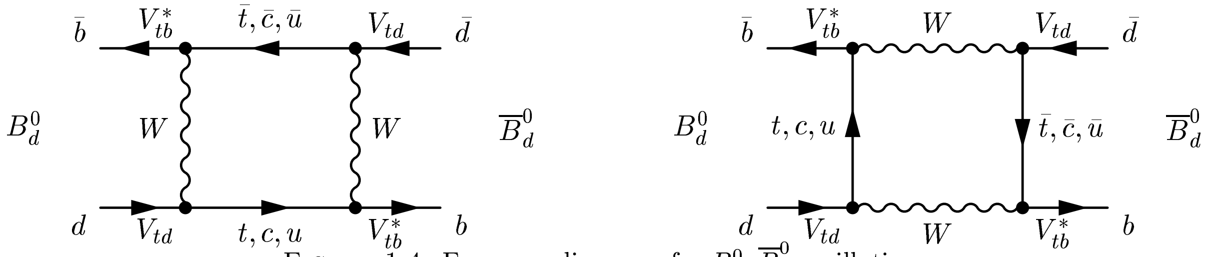

FIGURE 1.4: Feynman diagrams for B^0–$\overline{B}^0$ oscillations.

where M and Γ are 2×2 Hermitian matrices. The CPT invariance ensures that $H_{11} = H_{22}$. The off-diagonal elements M_{12} and Γ_{12} give the mixing between flavor states. They are respectively the dispersive, and absorptive part of the transition amplitude from B^0 to $\overline{B}^0$ and are in general both complex numbers.

The mass eigenstates, defined as the eigenstates of the Hamiltonian H, are normally parameterized as:

$$|B_L\rangle = p|B^0\rangle + q|\overline{B}^0\rangle, \tag{1.16}$$

$$|B_H\rangle = p|B^0\rangle - q|\overline{B}^0\rangle, \tag{1.17}$$

with $|q|^2 + |p|^2 = 1$. The indices refer to the light and heavy mass eigenstates. The mass eigenstate have definite masses $M_{H,L}$ and decay widths $\Gamma_{H,L}$. We define the mass and width difference as:

$$\Delta m = M_H - M_L, \qquad \Delta\Gamma = \Gamma_H - \Gamma_L, \tag{1.18}$$

where Δm by definition is positive while $\Delta\Gamma$ can be negative. The average mass and width are given by:

$$m_B = \frac{M_H + M_L}{2}, \qquad \Gamma_B = \frac{\Gamma_H + \Gamma_L}{2}. \tag{1.19}$$

By solving the eigenvalue equations we obtain the following relations:

$$\begin{aligned}
(\Delta m)^2 - \frac{1}{4}(\Delta\Gamma)^2 &= (4|M_{12}|^2 - |\Gamma_{12}|^2), &(1.20)\\
\Delta M \Delta\Gamma &= 4\mathrm{Re}(M_{12}\Gamma_{12}^*), &(1.21)\\
\frac{q}{p} &= -\frac{2M_{12}^* - i\Gamma_{12}^*}{\Delta m - \frac{i}{2}\Delta\Gamma}. &(1.22)
\end{aligned}$$

One can show that in the B system $|\Gamma_{12}| \ll |M_{12}|$ and the relations to first order simplify to:

$$\begin{aligned}
\Delta m &= 2|M_{12}|, &(1.23)\\
\Delta\Gamma &= 2\frac{\mathrm{Re}(M_{12}\Gamma_{12}^*)}{|M_{12}|}, &(1.24)\\
\frac{q}{p} &= -\frac{M_{12}^*}{|M_{12}|}\left(1 - \frac{1}{2}\mathrm{Im}\left(\frac{\Gamma_{12}}{M_{12}}\right)\right) \approx -\frac{M_{12}^*}{|M_{12}|}. &(1.25)
\end{aligned}$$

If we ignore $\Delta\Gamma$, which is small, the time evolution for a neutral b-meson produced as a $|B^0\rangle$ ($|\overline{B}^0\rangle$) state at $t = 0$ is given by:

$$\begin{aligned}
|B^0(t)\rangle &= e^{-\Gamma_B t/2 - i m_B t}\left[\cos\left(\frac{\Delta m}{2}t\right)|B^0\rangle + i\left(\frac{q}{p}\right)\sin\left(\frac{\Delta m}{2}t\right)|\overline{B}^0\rangle\right], &(1.26)\\
|\overline{B}^0(t)\rangle &= e^{-\Gamma_B t/2 - i m_B t}\left[i\left(\frac{q}{p}\right)\sin\left(\frac{\Delta m}{2}t\right)|B^0\rangle + \cos\left(\frac{\Delta m}{2}t\right)|\overline{B}^0\rangle\right]. &(1.27)
\end{aligned}$$

To measure CP violation one also has to consider the final state, $|f\rangle$, into which the B^0 decays. The final state can be either an exclusive or an inclusive state. Depending on which final state is being measured, one obtains different measures of CP violation parameters. For any final state we have a CP phase factor, ω_f, given by:

$$\mathrm{CP}|f\rangle = \omega_f|\bar{f}\rangle, \quad \mathrm{CP}|\bar{f}\rangle = \omega_f^*|f\rangle, \quad |\omega_f| = 1. \tag{1.28}$$

A final state is a CP eigenstate when $f = \bar{f}$ and the phase factor is then denoted $\eta_{f_{\mathrm{CP}}}$.

The decay amplitudes are defined according to:

$$A_f = \langle f|H_d|B^0\rangle, \quad \overline{A}_f = \langle f|H_d|\overline{B}^0\rangle, \tag{1.29}$$

where H_d is the decay Hamiltonian. A decay amplitude is given by a sum over the individual amplitude contributions given by the Feynman diagrams for the decay mode. Each amplitude can be written as:

$$A_{f,j} = |A|_j e^{i(\delta_j + \phi_j)}, \tag{1.30}$$

where δ_j comes from long distance rescattering effects in the decay and is known as the "strong phase" as the strong interactions dominate in the rescatterings. The ϕ_j phase is called the "weak phase" and is given by the complex phase in the CKM matrix. Applying the CP transformation

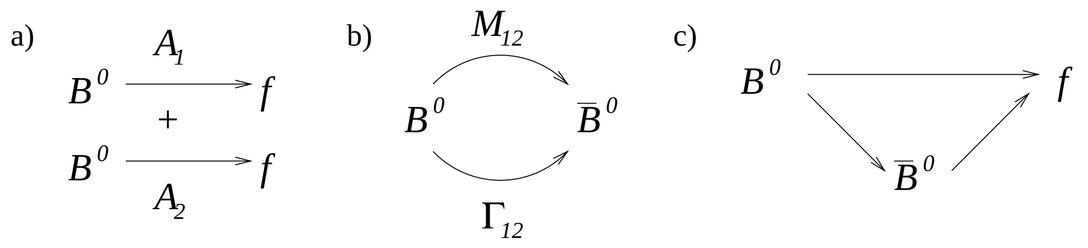

FIGURE 1.5: The three types of CP violations: a) CP violation in decay. b) CP violation in mixing. c) CP violation in the interference between decays with and without mixing.

will change the sign of ϕ_j as one has to use the complex conjugate of the CKM matrix elements in the calculations of the CP transformed Feynman diagrams. The strong interaction is CP invariant and the strong phase does not change sign. Generally one therefore has:

$$A_f = \sum_j |A|_j e^{i(\delta_j+\phi_j)}, \quad \overline{A}_{\bar{f}} = \sum_j |A|_j e^{i(\delta_j-\phi_j)}. \tag{1.31}$$

If one has more than one contribution to the amplitude, each of them with different, δ_j and ϕ_j phases, one will have CP violation as we shall see below.

1.3.2 The three types of CP violation

The weak interaction in general introduces a complex phase in the transition amplitude for a given process. When the CP transformation is applied to this process, this complex phase will have the opposite sign. We can however only measure transition probabilities, which are not sensitive to a change in phase. To get different transition probabilities for the CP transformed process, the transition amplitude should be a sum of several contributions with different phases. This will give rise to an interference term in the transition probability. Depending on how this interference occurs one classifies the CP violation into three different types described below and illustrated in Figure 1.5. Note that there are decay channels, where one can have contributions from several types of CP violation.

CP violation in decay: The decay rate into a given final state, f, is proportional to $|A_f|^2$. If there are several contributing diagrams to a decay mode, as indicated by Equation (1.31), and both the weak and strong phases are different, the rate of the decay and the rate of the CP transformed decay will not be equal as one has [4]:

$$|A_f|^2 - |\overline{A}_{\bar{f}}|^2 = -2\sum_{i,j} |A_i||A_j| \sin(\phi_i - \phi_j)\sin(\delta_i - \delta_j). \tag{1.32}$$

This can occur in decays of both neutral and charged b-hadrons and is the only CP violation possible in the charged b-hadrons. The CP violation in this case is measured by the asymmetry in decay rates:

$$a_f = \frac{\Gamma(B^+ \to f) - \Gamma(B^- \to \bar{f})}{\Gamma(B^+ \to f) + \Gamma(B^- \to \bar{f})}, \tag{1.33}$$

where the asymmetry will be given by:

$$a_f = \frac{1 - |\overline{A}_{\bar{f}}/A_f|^2}{1 + |\overline{A}_{\bar{f}}/A_f|^2}. \tag{1.34}$$

This type of CP violation is often called direct CP violation and has recently been observed in kaons [15, 16]. As it involves strong phases, which cannot be calculated directly from QCD, measurements are difficult to relate to the Unitarity Triangle.

CP violation in mixing: The mixing between B^0 and $\overline{B}^0$ is dominated by the dispersive part, M_{12}, but a phase difference between dispersive and absorptive part will introduce a CP violating effect. This effect manifests itself in the value of $\frac{q}{p}$ as can be seen from Equation (1.25). If $|q/p| \neq 1$ the probability of a produced B^0 to decay as a $\overline{B}^0$ will be different from the probability of a produced $\overline{B}^0$ to decay as a B^0. The effect is present in all B^0 decays, but is small as absorptive part is much smaller than the dispersive one and, to first order, they also have the same phase. To measure the CP violation in this case one would typically use inclusive decay modes, such as semileptonic b decays:

$$a_{sl} = \frac{\Gamma(\overline{B}^0(t) \to \ell^+\nu X) - \Gamma(B^0(t) \to \ell^-\nu X)}{\Gamma(\overline{B}^0(t) \to \ell^+\nu X) + \Gamma(B^0(t) \to \ell^-\nu X)}. \tag{1.35}$$

Using Equation (1.26) and ignoring CP violation in decay one obtains:

$$a_{sl} = \frac{1 - |q/p|^4}{1 + |q/p|^4}. \tag{1.36}$$

An average over current measurements [3] in the B_d^0 system gives:

$$a_{sl} = (0.2 \pm 1.4) \times 10^{-2}. \tag{1.37}$$

A measurement of a_{sl} is not easy to relate to CKM parameters as it involves calculating Γ_{12}, which is dominated by hadronic physics and has uncertainties of around 50%.

CP violation in the interference between decays with and without mixing: The third type of CP violation can be simpler to relate to CKM parameters. It uses decays into CP eigenstates, f_{CP}. Both B^0 and $\overline{B}^0$ can decay into these states and a produced B^0 meson can therefore either decay directly to f_{CP}, or first oscillate to $\overline{B}^0$, and then decay to f_{CP}. A phase difference between these two paths will introduce a CP violation effect. This effect can be measured as a rate difference between produced B^0 and $\overline{B}^0$ mesons decaying to f_{CP}:

$$a_{f_{\mathrm{CP}}}(t) = \frac{\Gamma(\overline{B}^0(t) \to f_{\mathrm{CP}}) - \Gamma(B^0(t) \to f_{\mathrm{CP}})}{\Gamma(\overline{B}^0(t) \to f_{\mathrm{CP}}) + \Gamma(B^0(t) \to f_{\mathrm{CP}})}. \tag{1.38}$$

This time dependent asymmetry is in general given by:

$$a_{f_{\mathrm{CP}}}(t) = -\frac{1 - |\lambda_{f_{\mathrm{CP}}}|^2}{1 + |\lambda_{f_{\mathrm{CP}}}|^2}\cos(\Delta m_B t) + \frac{2\mathrm{Im}\lambda_{f_{\mathrm{CP}}}}{1 + |\lambda_{f_{\mathrm{CP}}}|^2}\sin(\Delta m_B t), \tag{1.39}$$

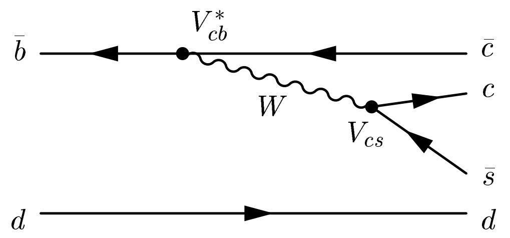

FIGURE 1.6: The tree level Feynman diagram for the golden decay, $B_d^0 \to J/\psi K_S^0$.

with:

$$\lambda_{f_{\rm CP}} = \eta_{f_{\rm CP}} \frac{q}{p} \frac{\overline{A}_f}{A_f}. \tag{1.40}$$

If the decay mode is dominated by a single diagram, there is little interference in the decay and we can ignore CP violation in decay ($|\overline{A}_f/A_f| = 1$). We have also approximately $|q/p| = 1$ and therefore $|\lambda_{f_{\rm CP}}| = 1$. This simplifies the asymmetry expression to:

$$a_{f_{\rm CP}}(t) = {\rm Im}\lambda_{f_{\rm CP}} \sin(\Delta m_B t), \tag{1.41}$$

where $\lambda_{f_{\rm CP}}$ has been reduced to a complex phase, which only contains the weak phases of the mixing and decay and no strong phases. $\lambda_{f_{\rm CP}}$ can therefore be directly related to CKM parameters. The prime example of such a decay is the golden decay, $B_d^0 \to J/\psi K_S^0$, described in the next section.

1.4 The golden decay

The $B_d^0 \to J/\psi K_S^0$ decay is dominated by the tree diagram shown in Figure 1.6. The decay also gets contributions from so-called penguin diagrams, which are lowest order loop diagrams. The dominant penguin diagram however has the same weak phase as the tree diagram and the other penguin diagrams are suppressed to less than 1%. To a very good approximation we therefore have no CP violation in decay and one can study the simple case of CP violation in the interference between decays with and without mixing as given by Equation (1.41).

The B_d^0 meson does not decay directly to $J/\psi\ K_S^0$, but rather we have $B_d^0 \to J/\psi K^0$ and $\overline{B}_d^0 \to J/\psi \overline{K}^0$. As with B^0 and $\overline{B}^0$ mesons the K^0 and $\overline{K}^0$ mesons mix and the mass eigenstates are known as K_S^0 and K_L^0. These are not quite CP eigenstates, but the CP violation is so small that we can ignore it here. The mixing of neutral kaons gives us the common final state and introduces an extra factor into $\overline{A}_f/A_f$:

$$\left(\frac{p}{q}\right)_K = \frac{V_{cs}V_{cd}^*}{V_{cs}^*V_{cd}}. \tag{1.42}$$

From the calculation of the Feynman diagram, shown in Figure 1.6, one therefore obtains:

$$\frac{\overline{A}_{J/\psi K_S^0}}{A_{J/\psi K_S^0}} = \eta_{J/\psi K_S^0} \left(\frac{V_{cb}V_{cs}^*}{V_{cb}^*V_{cs}}\right) \left(\frac{V_{cs}V_{cd}^*}{V_{cs}^*V_{cd}}\right), \tag{1.43}$$

where $\eta_{J/\psi K_S^0} = -1$ is the CP eigenvalue of the final state. $\left(\frac{p}{q}\right)_K$ is calculated from a box diagram similar to the one in Figure 1.4, but with s quarks instead of b quarks and the dominant contribution coming from the exchange of c quark. $\left(\frac{q}{p}\right)$ is calculated from Figure 1.4 with the t quark exchange dominating:

$$\left(\frac{q}{p}\right) = \frac{V_{td}V_{tb}^*}{V_{td}^*V_{tb}}. \tag{1.44}$$

We therefore find:

$$\lambda(B_d^0 \to J/\psi K_S^0) = -\left(\frac{V_{td}V_{tb}^*}{V_{td}^*V_{tb}}\right)\left(\frac{V_{cb}V_{cs}^*}{V_{cb}^*V_{cs}}\right)\left(\frac{V_{cs}V_{cd}^*}{V_{cs}^*V_{cd}}\right) = -e^{-i2\beta}, \tag{1.45}$$

where we have used Equation (1.13). Equation (1.41) therefore becomes:

$$a_{J/\psi K_S^0}(t) = \sin(2\beta)\sin(\Delta m_B t), \tag{1.46}$$

and the Unitarity Triangle angle, β, can be directly measured. Because of the simple and theoretically clean prediction of $a_{J/\psi K_S^0}(t)$, the decay is called the "golden decay". Other decays also exist which allows measurement of Unitarity Triangle parameters, but most of them contain strong phases or have very small branching ratios.

1.4.1 Measuring CP violation in the golden decay

The $B_d^0 \to J/\psi K_S^0$ decay is not "golden" only from a theoretical point of view, but also experimentally, as it has several signatures which make the asymmetry measurement relatively simple. HERA – B was designed to measure this asymmetry. A typical golden decay, as produced in HERA – B, is shown in Figure 1.7. The steps needed to do the asymmetry measurement are:

1. B_d^0 **production:** The golden decay having $\mathrm{Br}(B_d^0 \to J/\psi K_S^0) = 4.5\cdot 10^{-4}$ is not a common decay. Furthermore the final state is only reconstructed in specific modes (see next point) which reduces the rate by about a factor 20. To make a significant measurement one therefore needs to produce $10^7 - 10^8$ B_d^0 mesons. In HERA – B these are mesons produced in hadronic collisions, see Section 1.5.

2. $J/\psi K_S^0$ **reconstruction:** The $J/\psi K_S^0$ final state has to be picked out from millions of other events. Fortunately J/ψ decays into charged lepton pairs (e^+e^- or $\mu^+\mu^-$) with a branching ratio of about 6% for each type. The charged leptons are less common than other particles and can relatively easy be identified. The K_S^0 has a long lifetime ($c\tau = 2.68$ cm) and decays mainly to two pions, making K_S^0 mesons simple to identify and reconstruct.

3. **Decay time measurement:** The observable asymmetry varies with the time between the production and decay of B_d^0. The variation is characterized by the mass difference $\Delta m_B \approx 0.5$ ps and the time resolution should therefore be better than this. A time integrated measurement can also be done, but has less statistical power.

4. **Flavor tagging:** A key point in the asymmetry measurement is to determine whether the beauty meson was produced as a B_d^0 or $\overline{B}_d^0$. This is called flavor tagging. It cannot be determined from the decay itself as the final state can be reached by both meson decays. Instead one uses that b quarks are produced in pairs with the second b quark produced with the opposite flavor of the first. The second b quark will be part of a b-hadron and by reconstructing its decay, the flavor can usually be determined. The flavor tagging should

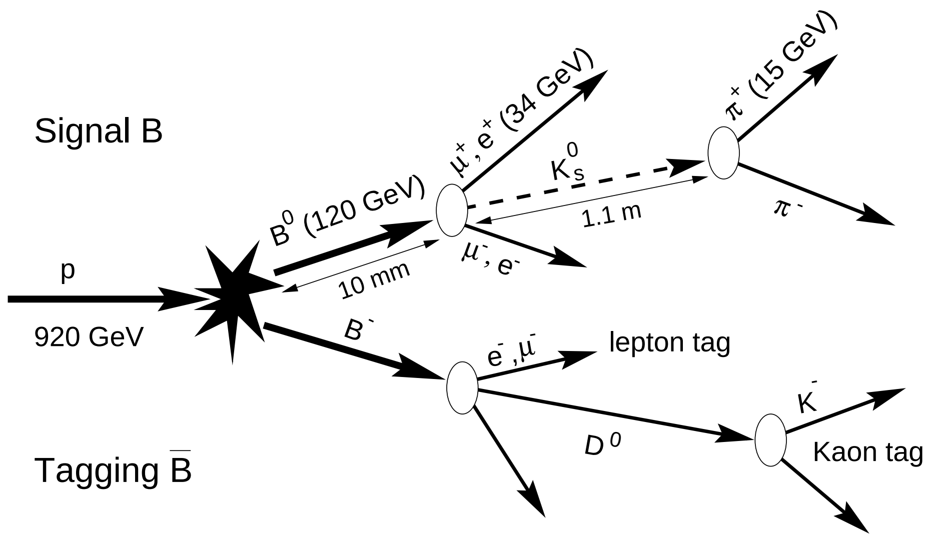

FIGURE 1.7: The topology of a golden decay with an accompanying possible flavor tags from the other b-hadron. The values are typically values for a golden decay produced at HERA – B.

be as efficient as possible and one therefore searches for inclusive decay modes such as semileptonic decays, where the lepton has the same charge-sign as the quark. Another possibility is to search for charged kaons originating from the decay chain $b \to c \to s$. Common for all tagging procedures is that they are not perfect and will tag the wrong flavor in some cases. This will dilute the asymmetry that has to be corrected by measuring the mistag probability in events in which both b hadrons are flavor tagged.

1.4.2 Current status of CP violation in the golden decay measurement

HERA – B was delayed by problems with the radiation hardness of the main tracker chambers and the first data taking with the complete detector was done only in spring 2000 and was mostly dedicated to commissioning studies. HERA – B has not yet reconstructed its first golden decay.

While HERA – B was built, two other experiments, BaBar and Belle, were also designed to measure CP violation in the B system. They both use an e^+e^- collider running at the $\Upsilon(4S)$ resonance to produce the B_d^0 mesons. They started data taking in 1999 and by winter 2002 they had collected in total more than 100 million B^0–$\overline{B}^0$ decays. From this sample they have reconstructed and tagged 1500 $B_d^0 \to J/\psi K_S^0$ decays. Using these reconstructed events along with decays in similar channels they obtained:

$$\sin 2\beta = 0.75 \pm 0.09_{stat} \pm 0.04_{syst}, \qquad \text{BaBar [1]} \tag{1.47}$$

$$\sin 2\beta = 0.99 \pm 0.14_{stat} \pm 0.06_{syst}. \qquad \text{Belle [2]} \tag{1.48}$$

Both measurement prove the existence of CP violation in the B system being consistent to each other and to the prediction from indirect measurements. In Figure 1.3 one can see the agreement

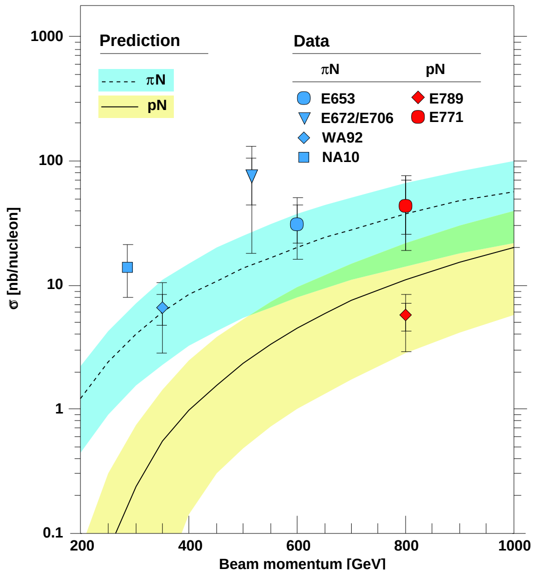

FIGURE 1.8: The experimental results on $b\bar{b}$ cross section in pN and πN data compared to QCD calculations [17] (continuous and dashed lines respectively). The shaded regions show the theoretical uncertainty.

between the indirect measurements and the average of the two $\sin 2\beta$ measurements. Over the next years, the two experiments will increase their statistics by about a factor of 10 allowing many interesting measurements in the B system beyond the golden decay mode to take place.

1.5 Beauty production in hadronic collisions

In HERA – B the beauty hadrons are produced by colliding protons on different nuclear materials, such as carbon and tungsten. Beauty production is a hard scattering process and it is therefore sufficient to consider the collisions between a proton and a single nucleon. The total beauty production cross section in both proton–nucleon and pion–nucleon is not very well-known experimentally or theoretically. Figure 1.8 shows the measurements/calculations available in 1998. The uncertainty in the cross section has always been a source of large uncertainty in HERA – B physics potential estimation as it depends strongly on the number of beauty hadrons produced per second. It has therefore always been one of the goals of HERA – B to provide a precise measurement of the beauty production cross section. Though HERA – B will not be able to do measurements of CP violation, it will still be able to measure the beauty production cross section. This will serve as input to the calculations of beauty quark production in hadronic collisions, which is also important for the understanding of the production of the heavier top quark.

The calculations of heavy quark[1] production are based on perturbative quantum chromodynamics (QCD). Perturbative QCD can be used as the large mass of the quarks introduces a lower bound on the energy scale, Q^2 involved in the scattering. This scale is much larger

[1]Heavy quarks are defined as quarks much heavier than QCD scale, $\Lambda \sim 200$ MeV. This is c, b and t quarks, though c quarks for some calculations are too light for perturbative QCD calculations.

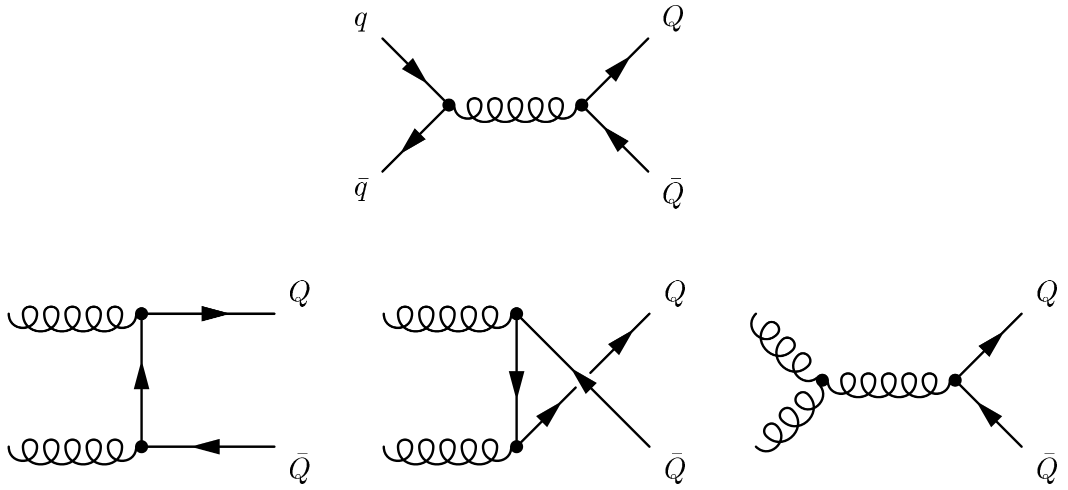

FIGURE 1.9: Leading order (LO) Feynman diagrams for heavy quark production.

than the QCD scale and the strong coupling $\alpha_S(Q^2)$ used in the calculation will therefore be much smaller than one, so that perturbative calculations are valid. Light quark production can be calculated within the framework of perturbative QCD as well, for instance the R-ratio. For hadronic production of light quarks, regulators (like cuts in the allowed transverse momentum) have to be introduced to secure the perturbative frame work.

In perturbative QCD the total hadronic cross section factorizes as[18]:

$$\sigma = \sum_{i,j} \int dx_1 dx_2 f_i(x_1, \mu^2) f_j(x_2, \mu^2) \hat{\sigma}_{ij}(\hat{s}, m^2, \mu^2), \qquad \text{with } \sqrt{\hat{s}} = s x_1 x_2 \tag{1.49}$$

where s is the center of mass energy in the hadron scattering. $f_i(x, \mu^2)$ are the parton distribution functions (PDFs) for the quarks, anti-quarks and gluons in the colliding hadrons. $\hat{\sigma}_{ij}(\hat{s}, m^2, \mu^2)$ is the short distance cross section which is calculated as a power series expansion in $\alpha_S(\mu^2)$. For heavy quark production the lowest order terms are $O(\alpha_S(\mu^2)^2)$ and are calculated from the Feynman diagrams in Figure 1.9. μ is the factorization scale, which is the energy scale that separates the hard scattering described by $\hat{\sigma}_{ij}(\hat{s}, m^2, \mu^2)$ and the soft QCD physics which is contained inside the PDFs. It should be of the order of the hard scale set by the heavy quark mass, m_Q. A second scale, the renormalization scale, is also present in perturbative QCD and is used to remove ultraviolet divergences. For simplicity this scale is usually set equal to the factorization scale as was done in the above equations. The physical cross section does not depend on the choice of the unphysical parameter μ, but some scale dependence of μ will be present in finite order calculations in α_S.

Figure 1.10 shows the scale dependence and the large variation of the leading order (LO) result and clearly indicates the need for higher order corrections. The calculation of the terms to α_S^3 has been done [19] and involves calculation of corrections from diagrams such as those shown in Figure 1.11.

Even after the NLO corrections, the beauty cross section at HERA – B still contains a large theoretical uncertainty from the scale dependence. One reason is that the production happens close to the kinematical threshold, which introduces a second energy scale [20], $m_{th} = (\hat{s} - 4m_Q^2)/2m_Q$. This affects the emission of soft gluons and gives rise to large logarithms in the perturbative coefficients giving rise to terms in the perturbative expansion of the form

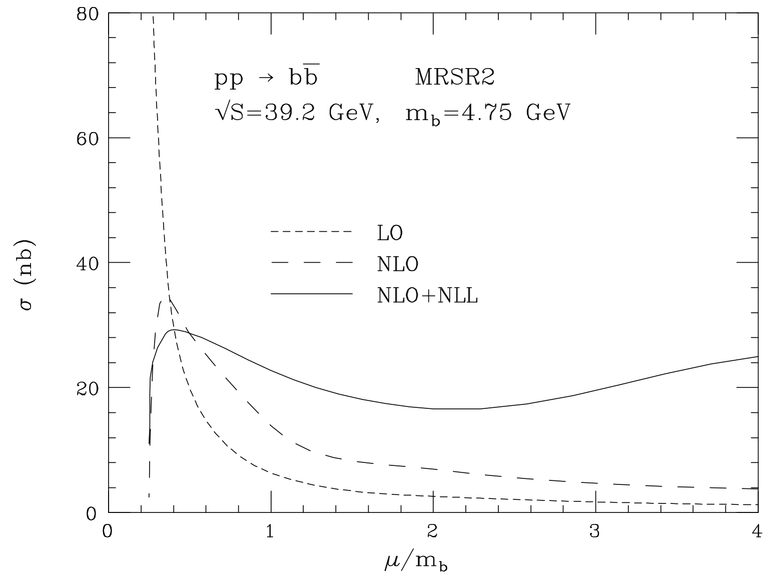

FIGURE 1.10: Scale dependence of the total beauty production cross section at HERA – B in LO, NLO and NLO+NLL calculations in perturbative QCD.

$(\alpha_S \log^2(1-4m_Q^2/\hat{s}))^n$ and $(\alpha_S \log(1-4m_Q^2/\hat{s}))^n$. As can be seen, these become important close to the threshold. The Next-to-Leading-Logarithmic (NLL) resummations[2] of this soft gluon emission have been done by several groups [21, 22] where the corrections have been summed to all orders in n and matched to the NLO fixed order result. This resummation significantly reduces the scale dependence, as can be seen in Figure 1.10.

Even with the reduction in scale dependence, the theoretical prediction still contains large uncertainties. A major uncertainty comes from the experimental uncertainty in mass of the b quark, which is usually considered to be between 4.5 and 5.0 GeV$/c^2$. Because the production happens so close to threshold this 10% uncertainty gives rise to large uncertainties in the cross section. The latest QCD calculations quote predictions of $\sigma(b\bar{b}) = 25^{+20}_{-13}$ nb/nucleon [22] and $\sigma(b\bar{b}) = 30 \pm 8 \pm 10$ nb/nucleon [21]. A precise measurement of the cross section would therefore serve as an important input to perturbative QCD calculations.

[2] The "Next" comes from summing not only the leading $\log^2()$ terms (LL) but also the sub-leading $\log()$ terms.

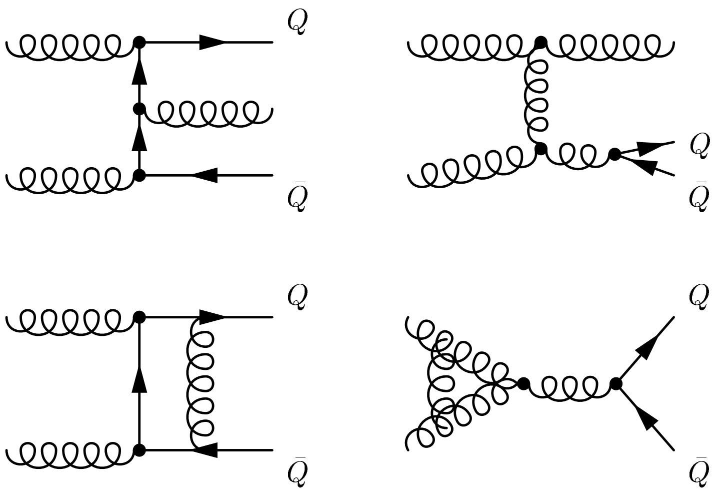

FIGURE 1.11: Examples of Next-to-Leading Order (NLO) Feynman diagrams for heavy quark production.

Chapter 2

The HERA – B experiment

2.1 Overview

In HERA – B the 920 GeV/c proton beam of the HERA[1] storage ring collides with fixed target wires giving a center-of-mass energy of 41.6 GeV. At this energy the beauty production cross section is around 25 nb/nucleon (see Section 1.5), while the inelastic cross section is 13 mb/nucleon for a titanium target. To produce the 10^8 B_d^0 mesons needed for the golden decay measurement, as described in Section 1.4.1, an inelastic interaction rate of about 40 MHz is needed. This rate together with the tracking, vertexing and particle identification requirements described in Section 1.4.1 are the basis for the design of HERA – B.

Figure 2.1 shows a schematic overview of HERA – B. It has the structure of a typical forward spectrometer with a large angular acceptance. Multiple target wires are inserted into the halo of the proton beam and off center protons collide with these, letting the main beam unharmed through the complete detector. The spectrometer is built around a warm dipole magnet with twenty tracking stations, called superlayers, in front, inside and behind the magnet. The magnetic field is oriented along the y axis[2] and has a maximal strength of 0.8 T and a field integral $\int B\,dl \approx 2.2$ Tm. Each of the twenty superlayers consists of several layers of active tracking devices in order to have a high tracking efficiency. Several detector technologies with varying segmentations are employed in order to have less than 20% hit occupancy, but at the same time keeping the number of channels and the cost down. Between the target system and the magnet, the decay products of the collision pass through the Vertex Detector (VDS) built from silicon micro-strip detectors and used to accurately reconstruct primary and secondary vertices. Each tracking station inside and behind the magnet is divided into two chambers: Inner and Outer tracking chambers. The Inner TRacker (ITR) is built from Micro Strip Gas Chambers (MSGC) and covers tracks with angles from 10 mrad to 100 mrad. The Outer TRacker (OTR) consists of honeycomb drift chambers and cover angles out to 250 mrad in x and 160 mrad in y. The tracking stations after the magnet (Main tracker) are used to find the charged tracks, while the stations inside the magnetic field are used to connect tracks from the Main tracker to tracks found in the VDS and to reconstruct long lived K_S^0 mesons.

Particle identification is provided by several detector systems. Before the last two tracking

[1]Short for the German "Hadron Elektron Ring Anlage".

[2]Throughout this thesis we will use a coordinate system with origin at the target. The z axis is along the direction of the proton beam, the y axis points upward and the x axis to the center of the HERA ring.

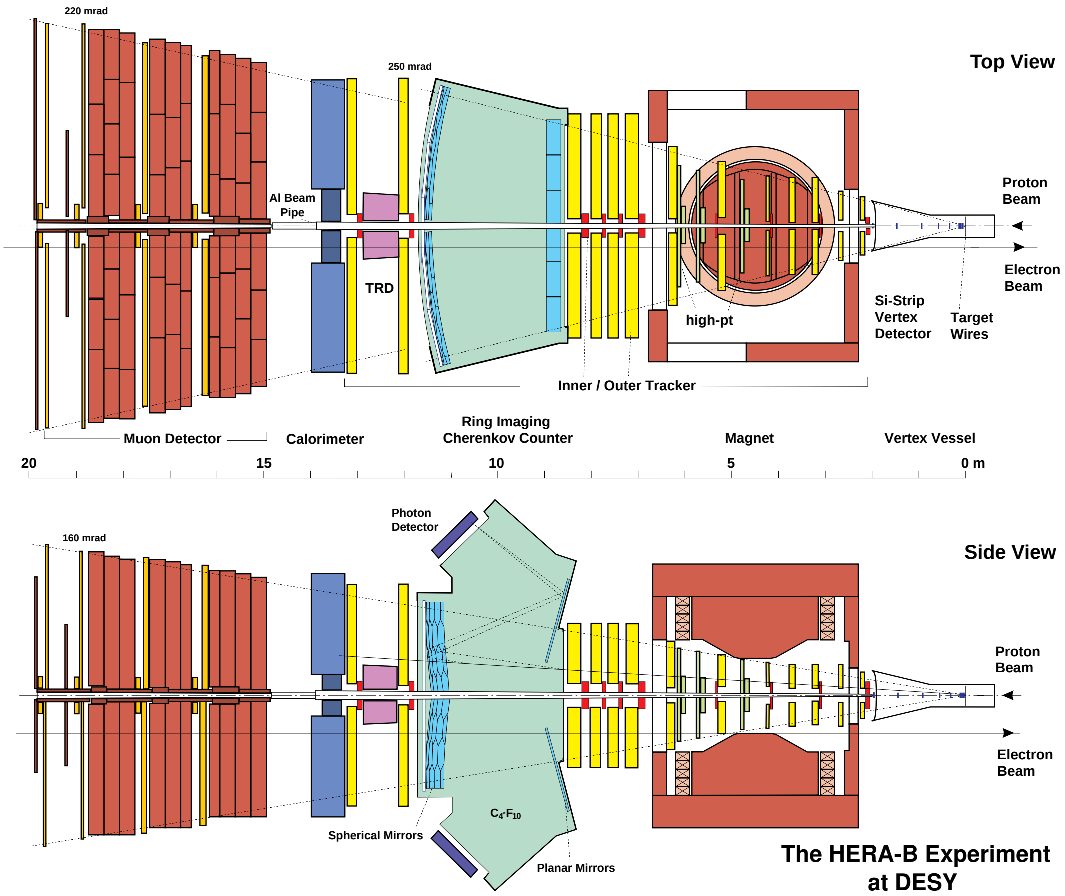

FIGURE 2.1: Schematic view of the HERA – B spectrometer from the top and side.

stations a large Ring Imaging Cherenkov Detector (RICH) provides kaon–pion separation over a large momentum range. Electrons[3] and photons are identified by the Electromagnetic Calorimeter (ECAL) placed after the tracking system. At small angles, where the track density is high, electron, pion separation is also provided by a Transition Radiation Detector (TRD). At the far end of the spectrometer, muons are identified by drift chambers shielded from the large flux of hadrons by large blocks of concrete. Finally, specially constructed tracking chambers with pad readout are installed inside the magnet to provide fast recognition of high p_T tracks for the trigger system.

2.2 The HERA storage ring

The HERA storage ring is located at the research laboratory DESY in Hamburg, Germany. Through a series of pre-accelerators, shown in Figure 2.2, protons and electrons[4] are filled into

[3]Electrons will unless otherwise noted always refer to both electrons and positrons.

[4]HERA is able to use either electrons or positrons for the lepton beam, but has mainly used positrons as they provide a more stable beam and are able to reach higher intensities.

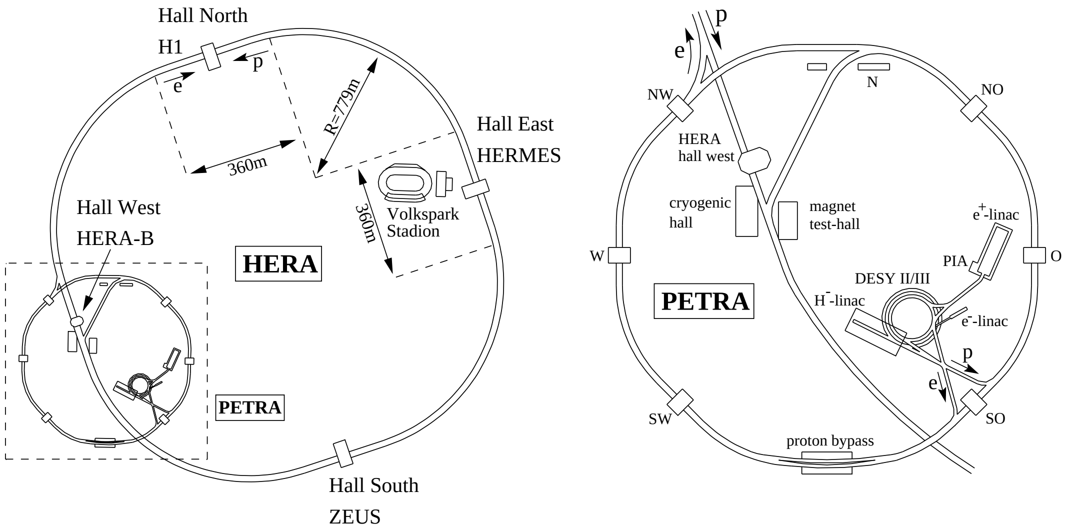

FIGURE 2.2: Left: The HERA storage ring and the four experiments. Right: The pre-accelerators used to fill HERA.

two separate storage rings located in the 6.3 km long HERA tunnel. The storage rings accelerate the protons to 920 GeV and the electrons to 27.5 GeV. The two beams are collided at two interaction points, where the experiments H1 [23] and ZEUS [24] are located. Both experiments use the deep inelastic *ep* collisions to study the proton structure functions. A third experiment, HERMES [25] uses only the longitudinally polarized electron beam, which is collided with a polarized gas target. This is used to study the spin structure of protons and neutrons. The last experiment is HERA – B, located in "Hall West" on the DESY site itself. In HERA – B only the proton beam is used.

The HERA proton beam is divided into bunches with approximately 10^{11} protons in each bunch. The time between bunches is 96 ns giving a bunch crossing rate at the HERA – B target of 10.42 MHz and the detector therefore has to be read out at that rate. It takes a bunch 21.12 μs to travel around the HERA ring and the storage ring can therefore theoretically contain 220 bunches. Due to the filling procedure only 180 of the 220 bunches are actually filled. The bunch structure is shown in Figure 2.3. Ten filled bunches follow right between each other in a so-called "group" and six groups, separated by single, empty bunches, are collected into a "train". A complete HERA fill contains three trains separated by five empty bunches (fifteen after the last). In HERA – B each bunch is identified with a bunch-crossing number, called BX.

2.3 The target system

The interactions in HERA – B are produced by inserting target wires into the halo of the HERA proton beam. The beam halo consists of protons, which are drifting away from the beam and would after some time have been removed by the collimators of the storage ring. By colliding the halo protons with target wires before they are removed, HERA – B is able to achieve high interaction rates without seriously affecting the central proton beam and therefore the other experiments at HERA. In particular the same proton fill can be used for several fills of electrons,

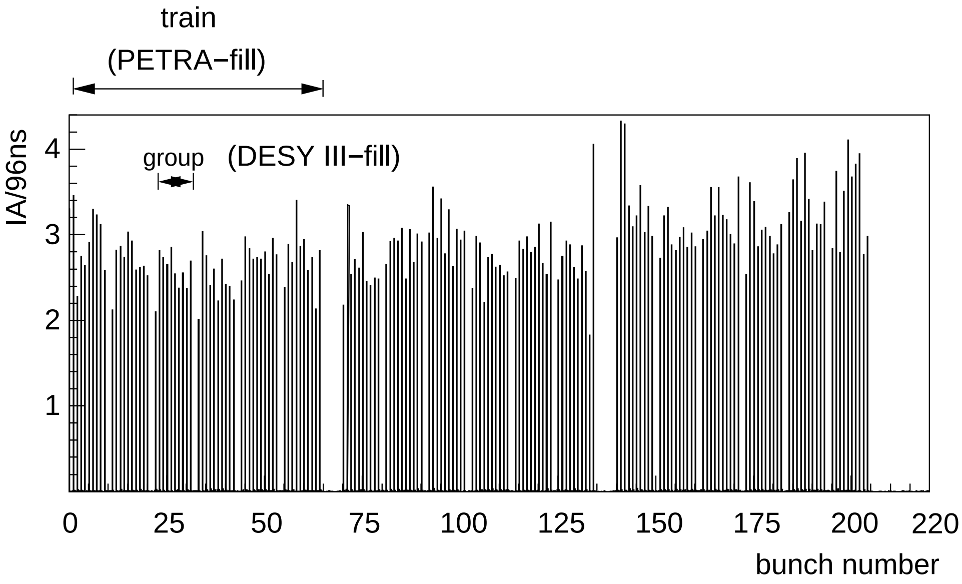

FIGURE 2.3: A typical proton bunch spectrum at HERA.

each lasting between eight and ten hours.

The design interaction rate at HERA – B is 40 MHz, but filled proton bunches only pass the target at 8.5 MHz ($180/220 \cdot 10.4$ MHz). An average of 4.7 interactions per bunch crossing is therefore needed. To achieve this and at the same time be able to distinguish between the different interactions, the HERA – B target system [26] uses eight target wires simultaneously, see Figure 2.4. The eight wires are mounted in two stations separated by approximately 4 cm. Each wire is attached to its own target fork and is able to move independently of the other target wires. Typically the wires are operated at a distance of 3–5 times the beam width from the center of the beam, where the beam width is typically around 400 μm.

Four scintillating hodoscope counters placed behind the RICH are used to monitor the interaction rate. If a too high rate is observed, the target wires are retracted slightly and at too low rates they are moved closer to the beam. The wire position adjustments are done ten times per second and the stepping motors are able to move a wire by as little as 50 nm. To distribute the interactions equally between the different target wires, each wire is equipped with a charge integrator which measures the charge produced by δ electron generation inside the target wires. In this way, only the wires producing too many/few interactions are adjusted and the number of interactions on each wire can be kept equal within 10%.

In the 2000 setup most of the target wires were ribbons with widths of 500 μm along the beam axis and 50 μm perpendicular to the beam. For the 2002 run most the wires have been replaced by round wires with a thickness of 50 μm as these provide better constraints for the vertex reconstruction. A round wire was successfully tested during the 2000 run. Different target materials are used, ranging from carbon over aluminum and titanium to tungsten which has the highest number of nucleons per nucleus. By using different target materials, HERA – B is able to measure nuclear effects in for instance prompt charmonium production.

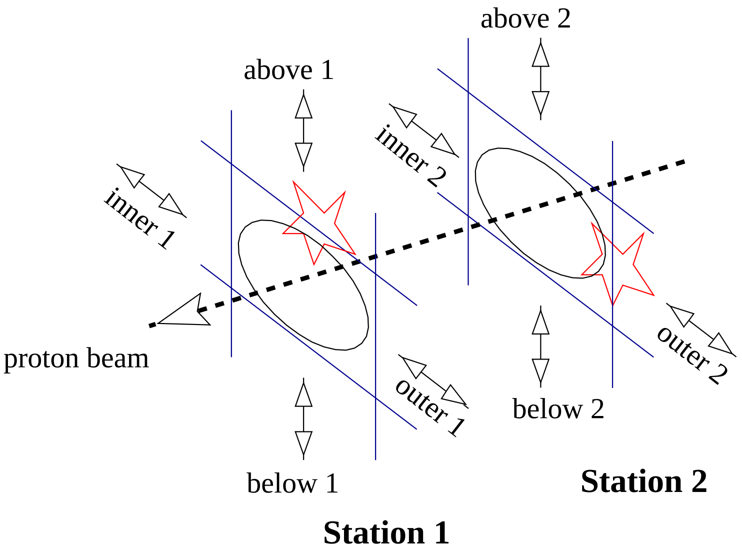

FIGURE 2.4: Overview and naming convention for the eight target wires.

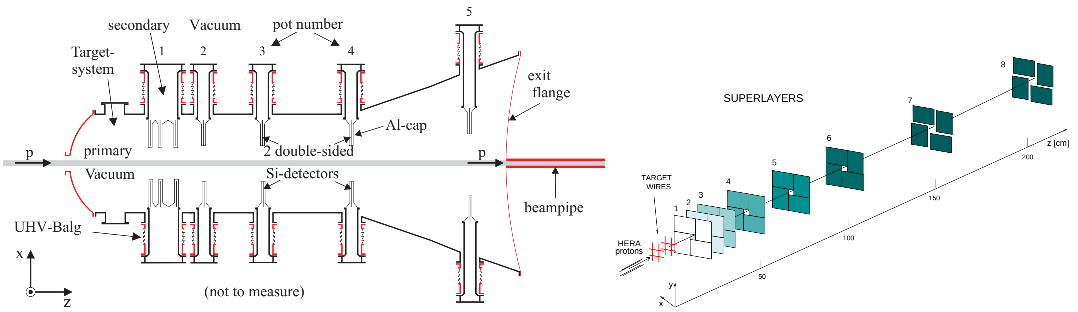

FIGURE 2.5: Left: Schematic overview of the VDS. The last superlayer is located after the exit flange and not shown. Right: The quadrant structure and superlayer numbering in the VDS.

2.4 The vertex detector

The VDS [27] is built from 64 double-sided silicon micro-strip detectors arranged in eight superlayers each divided into four quadrants as shown in Figure 2.5. In order to reconstruct the primary and secondary vertices with high precision the detectors need to be as close to the target as possible and to have the largest possible angular coverage. Furthermore, the amount of material should be kept as low as possible in order to minimize the amount of multiple scatterings. The first seven superlayers are therefore placed inside a large steel vessel which is part of the HERA storage ring. The vessel also contains the target system and is shown in Figure 2.5.

The detector module of each quadrant of a superlayer consists of two double sided silicon layers housed in a “Roman pot”. The first three superlayers are housed in the same Roman pot and due to space limitations one layer in each of the three superlayers only has single sided readout. The Roman pot system allows the silicon detectors to be retracted during injection, where the beam has larger spread. During data taking the detectors are moved to within 10 mm of the beam[5], giving an angular coverage from 10 to 250 mrad. The steering system allows the detector modules to be repeatedly placed at the same position with a precision of about a 1 μm and the VDS therefore does not need be re-aligned after each insertion. The seven superlayers have different angular acceptance and typically a charged track will pass through sensitive areas in three of the superlayers. The last superlayer (SI08) is not movable. It is located two meters from the target system just before the magnetic field. At the same position the ITR and OTR chambers provide measurements at larger angles and all charged tracks inside the HERA – B acceptance will therefore have a position measurement just before the magnet.

The silicon layers are made from 280 μm thick n-type material and have an active area of 50×70 mm^2. The strips on the two sides are orthogonal and, with a readout pitch of 50 μm, the two sides have 1024 and 1280 readout strips respectively. To facilitate better pattern recognition the strips are tilted by $\pm 2.5°$ relative to vertical/horizontal in the two layers of a superlayer. A single superlayer therefore provides measurements with stereo angles of $-2.5°$, $+2.5°$, $-87.5°$ and $92.5°$ relative to the y-axis. A charged track usually produces a signal in more than one strip in a layer and using a clusterisation algorithm a hit resolution of 12 μm has been achieved with a hit efficiency above 98%. With this performance, primary and secondary vertices can be reconstructed with a resolution along the z axis of about 500 μm.

2.5 The inner tracker

The ITR [28] covers angles from 10 mrad up to about 100 mrad. It has four superlayers inside the magnetic field labeled MS01–MS06[6], and six superlayers after the magnet labeled MS10–MS15. Each superlayer contains several layers in four quadrants as in the VDS. The strips of the layers are oriented along the vertical axis at angles of $0°$ or $\pm 5°$ and therefore only provide precise measurements of the horizontal coordinate, while the precision of the vertical coordinate is roughly $\frac{1}{\sin 5°} \approx 11.5$ times worse. This layout is used because the magnetic field bends charged particles in the horizontal plane and good precision in this plane is therefore needed for the momentum measurement. The four superlayers MS10, MS13, MS14 and MS15 are used by the First Level Trigger (FLT) and have two layers for each strip orientation as the FLT has strict requirements on chamber redundancy. The other superlayers have a varying number of layers optimized to provide high tracking efficiency and low ghost rate while at the same time keeping the amount of material as low as possible.

Each layer of the ITR is a so-called Micro Strip Gas Chamber (MSGC). The principle of a MSGC detector is shown in Figure 2.6. It consists of a 400 μm thick glass wafer with thin anode and cathode strips of gold placed in a gas filled volume (70% Ar and 30% CO_2). A large drift voltage is applied between the anode and the top of the chamber and the cathode strips. Passage of a charged particle will ionize the gas and the electrons will drift to the anode strips. The large electric field around the 10 μm wide anode strips provides a strong gas amplification of the primary electrons and a large signal can therefore be read out from the anode strip.

During the design phase it was found that if the field strength is adjusted to provide sufficient signals for minimum ionizing particles, heavy ionizing particles, such as α particles, would often lead to sparks in the chambers, and this would quickly destroy a chamber. Therefore

[5] During the 2000 run the detectors where limited to 15 mm from the beam and therefore the lower angular acceptance was about 15 mrad.

[6] MS01 is actually located in front of the magnetic field and there are no MS02 or MS04 superlayers in order to keep the amount of material down.

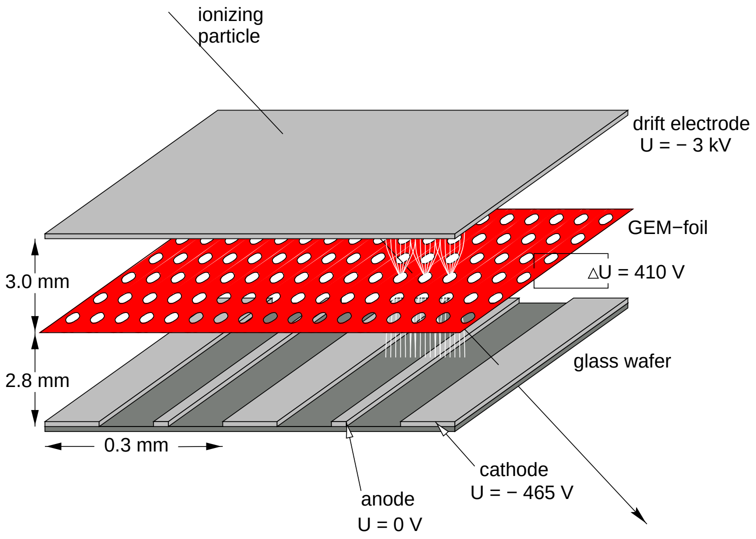

FIGURE 2.6: Schematic overview of a MSGC chamber with GEM technique.

an intermediate gas amplification is introduced by the Gas Electron Multiplier (GEM) layer, which allows one to reduce the voltage of the cathodes. The GEM layer is a 50 μm thick Kapton foil with thin copper layers on each side. The layer is filled with holes (typically with a 55 μm diameter and a 140 μm pitch) and by applying a voltage potential over the layer a gas amplification of a factor 15–100 takes place in the holes. At the MSGC layers the amplification is typically 200–1000.

The readout pitch is 300 μm (350 μm in MS14 and MS15) which gives a theoretical resolution of about 90 μm. During the 2000 run the resolution achieved was only 150 μm. The hit efficiency was measured to be between 91% and 97% depending on the quality of the GEM layer.

2.6 Outer tracker

The OTR [29] has angular acceptance overlapping slightly with the ITR and extending out to 250 mrad in the horizontal plane and 160 mrad in the vertical. In the magnet the OTR has three more superlayers than the ITR. They are labeled MC01 to MC08 starting from the front of the magnet and with no MC07 superlayer. After the magnet the OTR has six superlayers at approximately the same z positions as the ITR. The four before the RICH are labeled PC01–PC04 (Pattern Chambers) and the two after the RICH, TC01 and TC02 (Trigger Chambers). The names refer to the first four being the main chambers for pattern recognition, while the last two are mainly there to help the tracking in the FLT. Similar to the ITR, each superlayer in the OTR contains several layers with angles of 0° and ±5°. Due to the large size, the OTR (up to 6 × 4 meters in TC02) it is built from many modules. Chambers used by the FLT (PC01, PC04, TC01 and TC02) all have double layers to achieve high efficiency, while the other chambers use single layered modules.

The OTR layers are built from Honeycomb drift cells. Figure 2.7 shows the principle of both a single and double layered OTR module. A 25 μm gold-plated tungsten wire is used as anode wire in a honeycomb shaped drift cell. The drift cell (cathode) is made of 50 μm

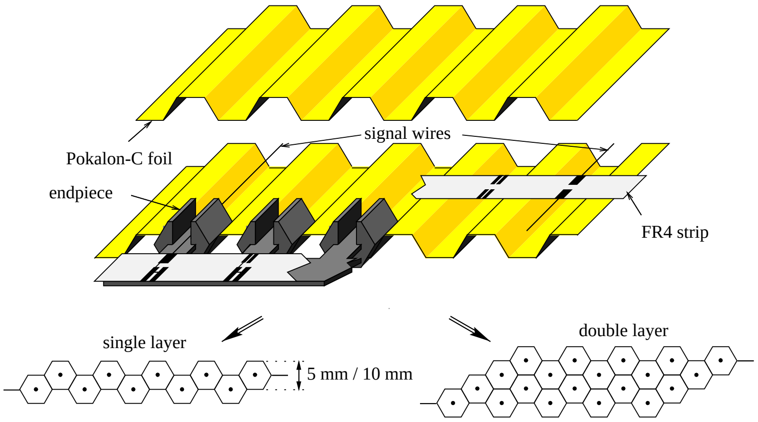

FIGURE 2.7: The honeycomb structure of OTR drift chambers.

polycarbonate (Pokalon-C) coated on the inside by a thin gold layer. The drift gas used is a mixture of $Ar/CF_4/CO_2$ which keeps the drift time below a single bunch crossing and minimizes aging effects[7]. The cell size is 5 mm in modules closest to the beam pipe and 10 mm further away, which keeps the occupancy at 40 MHz interaction rate below 20%. In order to have the best possible resolution, the readout of the OTR is equipped with Time-to-Digital Converters (TDCs) which measures at which time during a bunch crossing the drift electrons reach the signal wire. This drift time measures, with a left–right ambiguity, how far from the anode the charged particle passed. From this, a hit resolution of about 350 μm was achieved in the 2000 run for the 5 mm chambers. The hit efficiencies in 2000 were between 90% and 95% depending on the superlayer. This is well below the design value of 98% and was caused by noise in the readout, which forced rather higher signal threshold to be used. For the 2002 run the noise has been reduced significantly.

2.7 Ring imaging Cherenkov detector

The RICH [30] is used to identify charged kaons and protons. It uses the Cherenkov effect, where a charged particle traveling through a medium with a velocity higher than the velocity of light in the medium will emit photons in a cone around the charged particle. The angle θ_c of the emitted photons to the particle direction is given by:

$$\cos\theta_c = \frac{c}{v \cdot n}, \tag{2.1}$$

where n is the refraction index of the medium and v is the velocity of the particle. By measuring the Cherenkov angle, the velocity can be determined and by comparing it to the measured momentum, the mass and thus particle type can identified.

In HERA – B the RICH is a large vessel filled with C_4F_{10} as radiator gas. The photons are reflected to multi-channel photomultipliers, placed outside the detector acceptance, by a series of mirrors as shown in Figure 2.8. The threshold for Cherenkov light is 2.6 GeV/c for pions,

[7] Due to the high rate of charged particles aging was a major concern during the design of HERA – B and the OTR design had to be changed several times due to problems with aging in prototypes.

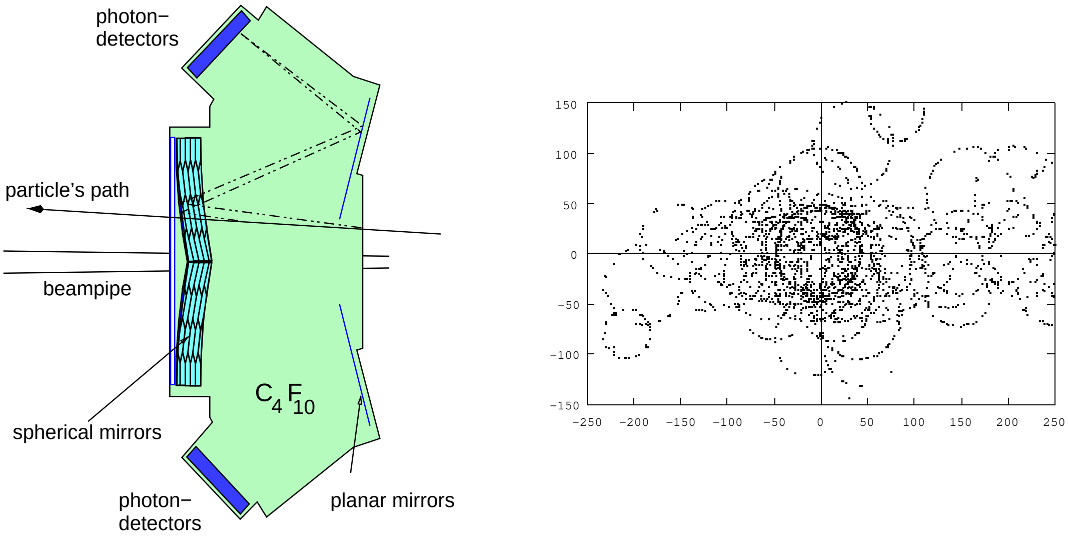

FIGURE 2.8: Left: Schematic of the RICH detector. Right: Example of the observed RICH hits in an inelastic event.

9.0 GeV/c for kaons and 17.2 GeV/c for protons. Pions and kaons can be separated up to momentum of about 60 GeV/c, protons up to 80 GeV/c and electrons up to 15 GeV/c.

2.8 Transition radiation detector

The TRD [31] provides electron-hadron separation in the inner part of the detector where the occupancy is very high. When a charged particle traverses media with different di-electricity constants it will emit transition radiation in the form of x-rays. The radiation emitted scales with the Lorentz-factor, $\gamma = \frac{E}{mc^2}$, and is therefore much more energetic for electrons. The TRD works by interleaving radiator material with straw tubes for detecting the x-rays.

In HERA – B the TRD is placed between TC01 and TC02 and has 32 layers of radiator material and straws. The radiator consists of polypropylene fibers with a 20 μm diameter and each layer has a thickness of 1.6 cm. The straws are made from thin multi-layered (Kapton + aluminum) cylindrical drift tubes with a diameter of 6 mm and are tilted by $\pm 30°$ relative to the horizontal axis. The gas used is 70% Xe, 20% CF_4 and 10% CO_2. The Xenon gas is used to detect the x-rays, while the other gases adjust the drift time so that a signal is collected within a single bunch crossing. The TRD gives good electron-pion separation in the energy range 1 to 100 GeV, but during the 2000 run it was not equipped with readout electronics for cost reasons and therefore not used.

2.9 Electromagnetic calorimeter

The ECAL [32] identifies electrons and photons by measuring their energy and comparing this to the charged tracks found in the tracking system. Photons do not leave a track in the tracker, and for electrons the measured momentum, p, should match the measured energy, E, as within

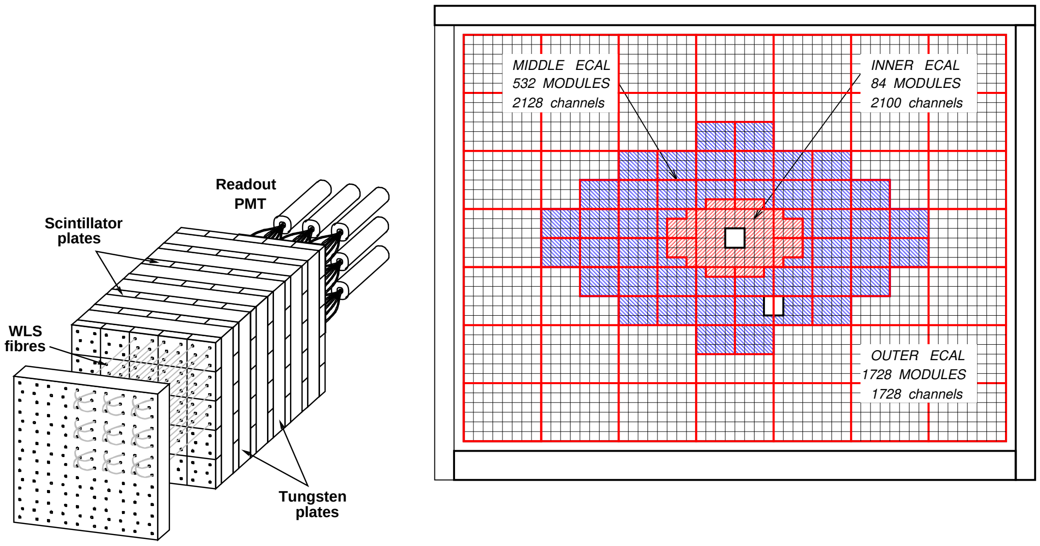

FIGURE 2.9: Left: Sandwich structure of an ECAL module. Right: The different regions of the ECAL.

measurement uncertainties one has $E/p \approx 1$. For hadrons, the calorimeter only measures a small part of the energy and therefore $E/p \ll 1$.

The ECAL is built as a sampling calorimeter with a sandwich structure as shown in Figure 2.9 where absorber plates are interleaved with plastic scintillator layers. Electrons and photons hitting the absorber will produce bremsstrahlung photons and electron-positron pairs respectively and a shower of secondary electrons, positrons and photons develops. The amount of particles in a shower is proportional to the total energy of the original electron or photon. The electron and positron are detected in the scintillator plates, whose light is collected by wavelength-shifting fibers inserted in the calorimeter parallel to the beam (so-called shashlik architecture). The amount of light is measured in photomultipliers located behind the calorimeter.

The calorimeter is divided into modules with a front face of 11.2×11.2 cm^2. In total it has 56 modules in the horizontal direction and 42 modules in the vertical one giving full angular acceptance. To keep occupancies below 10%, the modules are further divided into cells of varying size. In the inner most part ("inner ECAL") a module consists of 25 cells, in the "middle ECAL" of four cells and in the "outer ECAL" a module is a single cell. The division is shown in Figure 2.9. To keep the shower of an electron/photon confined laterally to a few cells, the inner ECAL uses a tungsten-alloy as absorber material, while the middle and outer ECAL use lead, which is cheaper.

A shower will typically deposit energy in several neighboring cells and the reconstruction therefore combines cells with energy depositions into clusters, which are used as candidates for electrons and photons. In the inner ECAL the expected energy resolution, $\sigma(E)$, is given by:

$$\frac{\sigma(E)}{E} = \frac{17\%}{\sqrt{E/\,\mathrm{GeV}}} \oplus 1.6\%, \tag{2.2}$$

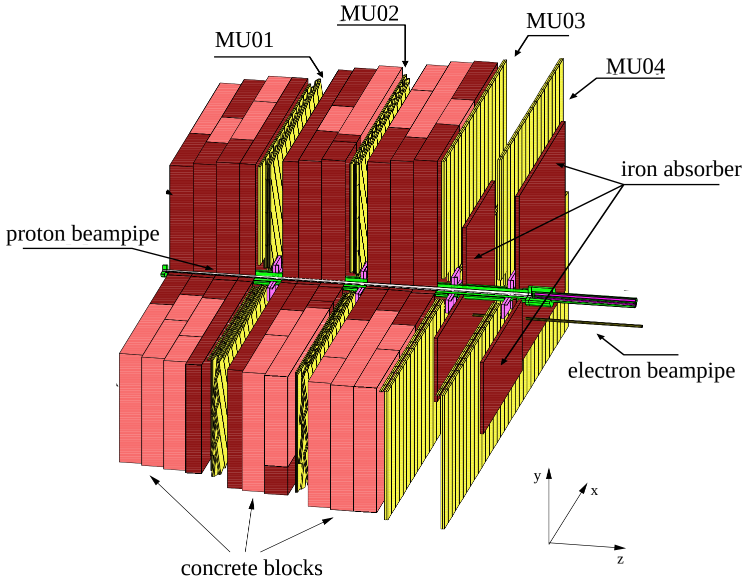

FIGURE 2.10: The four tracking stations of the MUON system and the concrete absorber separating them. The last two stations are only separated by about 5 cm iron.

while for the middle and outer ECAL it is given by:

$$\frac{\sigma(E)}{E} = \frac{9.5\%}{\sqrt{E/\,\mathrm{GeV}}} \oplus 1.0\%. \tag{2.3}$$

During the 2000 run the energy resolution was significantly worse because of the large amount of material in front of the calorimeter (mainly the ITR and OTR) and because of noise in the readout electronics.

2.10 Muon detector

The MUON detector [33] consists of four tracking stations (MU01–MU04) behind the calorimeter that are separated from each other by one meter thick walls of concrete and iron as shown in Figure 2.10. The absorber walls will stop all hadrons and only muons with momentum larger than 5 GeV/c are able to penetrate all walls and leave signals in all four tracking stations. They can therefore be identified quite reliably.

The inner part of the tracking stations (± 42 cm in x and ± 33 cm in y) is constructed from so-called gas-pixel chambers. The principle of such a chamber is illustrated in Figure 2.11. A square cell of 9×9 mm^2 is formed by four potential wires and a sense wire in the middle. The wires are 3 cm long, oriented parallel with the beam and surrounded by gas (80% Ar, 20% CF_4). Passage of a charged particle will ionize some of the gas and the electrons will drift to the anode wire and induce a signal. To limit the amount of readout channels the signals from 2×2 cells in MU01–MU02 and 1×4 cells in MU03–MU04 are combined together in the readout.

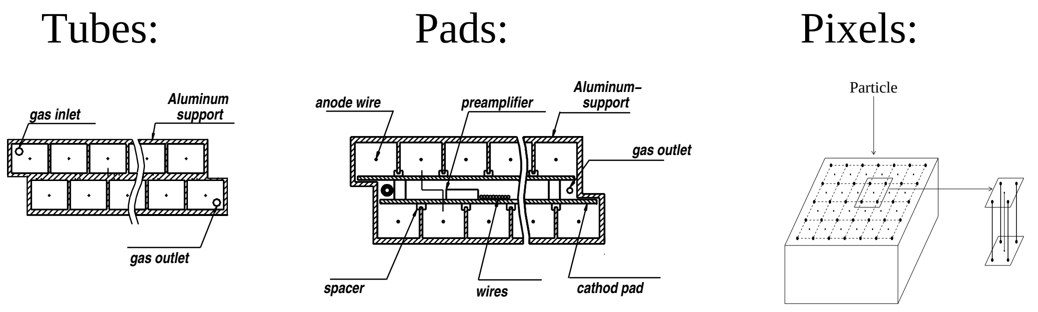

FIGURE 2.11: Left: Structure of a double layer in MU01 and MU02. Center: Structure of double layer in MU03 and MU04 with the cathode planes between the two layers. Right: Structure of a pixel chamber.

The outer part is made of proportional tube chambers. The tubes are 16 mm wide and in order to have high hit efficiency all layers are implemented as double layers with the second layer shifted half a cell size, as shown in Figure 2.11. The drift time of the cells is not used, only the presence of a signal or not. The tracking stations in MU01 and MU02 have three double layers with the tubes oriented at angles of 0° and $\pm 20^\circ$ relative to vertical. MU03 and MU04 only have a vertical oriented double layer, which instead is equipped with segmented cathodes (pads). The pads are 12×10 cm^2 large and provide the initial measurements of x and y coordinates of the muon for the trigger system. The signal from the pads is quite small and a pre-amplifier is therefore mounted directly on the backside of each pad and the signal from the two pads in a double layer is furthermore combined into a single measurement.

2.11 High-p_T chambers

To allow fast recognition of hadrons with high transverse momentum ($p_T > 2.0$ GeV$/c$) in the trigger system, three special tracking chambers called High-p_T chambers are installed in the magnetic field. These measure the position of charged tracks with a pad-type readout. The chambers have a projective segmentation, where high p_T particles, which have almost straight trajectories in the magnet, leave a characteristic pattern in the three stations. The pattern used by the trigger system is explained in the next chapter.

The inner part [34] of the High-p_T chambers is constructed from chambers similar to the gas-pixel chambers in the MUON detector. The outer part [35] uses straw tubes which are positioned against printed circuits with a pad pattern. A signal in the straw tube will induce an image charge on the pads, which can be read out. The straw tubes themselves are not read out. During the 2000 run both types of chambers were plagued by large amounts of electronics noise, which was only cured at the very end of the 2000 run. They were therefore not used during data taking.

Chapter 3

The HERA – B data acquisition and trigger system

For HERA – B to acquire enough beauty decays to measure CP violation, the detector needs to be read out for every bunch crossing. At the same time HERA – B is a finely segmented spectrometer with more than 500,000 channels to be read out for every event. The number of channels and amount of data per channel are given in Table 3.1. In total 470 kByte of data are produced per bunch crossing, which at the 10.4 MHz bunch crossing rate gives almost 5 TByte per second. No system in the world is able to read out and store that amount of data in a second. Even compressing the data by suppressing the empty channels only reduces the data rate by about a factor five. This can still not be stored and the data therefore have to be filtered.

The actual rate of interesting events is quite low. b-hadrons are only produced in about 10^{-6} of inelastic events and only $O(10^{-4})$ of these contain a golden decay. Other interesting b-hadron decays occur at the same or lower rates. The low rate (10^{-10}) is the reason in the first place why HERA – B is using such a high input rate, but this means that the event filtering has to maintain a high efficiency for signal events, reducing the event rate from 10.4 MHz to the about 20–50 Hz of events, which can be stored on tape. This event filtering is done using a multi-level trigger and buffer system, which uses an increasing amount of event data and processing power on a decreasing amount of events as events passes through the system. The trigger system will be described in this chapter. The implementation and performance of a part of the system (the Second Level Trigger) will be given in the following chapters.

The HERA – B trigger and data acquisition system description in this chapter corresponds more or less to the design. For the 2000 data taking only parts of the system were used and even in 2002 the full system will not be used. The rate and efficiencies estimated during design have been superseded and will only be mentioned, when needed to justify the design.

3.1 Trigger overview

The main physics channel for the HERA – B trigger design is the golden decay, $B_d^0 \rightarrow J/\psi K_S^0$. K_S^0 mesons are very common in hadronic collisions and therefore not well-suited to trigger on. The J/ψ meson though frequently decays into muon or electron pairs, particles which are less commonly produced, particularly with as high momentum and transverse momentum as the leptons produced in the J/ψ decay. Furthermore the lepton pair has an invariant mass

Detector:	Channels	Bits per channel
VDS	176,000	8
ITR	135,000	8
OTR	120,000	8
High-p_T	8,736	1
RICH	27,520	1
TRD	16,000	2
ECAL	5,800	16
MUON	30,814	1
Total	$\sim 520,000$	$\sim$ 470 kByte

TABLE 3.1: The number of channels and bits per channel for each detector component in HERA – B [36].

corresponding to the mass of the J/ψ (3.097 GeV/c^2), another good signal for the trigger system. b-hadrons also commonly decay semileptonically into high-p_T muons and electrons ($\mathrm{Br}(b \to \ell\nu X) \simeq 10\%$), thus providing a similar signature to the golden decay. In low multiplicity b-hadron decays, such as $B^0_d \to \pi^+\pi^-$, the hadrons from the decay, due to the high mass of the b-hadrons, will have a very high transverse momentum, which can be used to pick them out from the huge background of hadrons in inelastic decays. Finally b-hadrons have a relatively long lifetime ($\sim$ 1.5 ps) and, with the high boost of a fixed target experiment like HERA – B, they are able to fly up to several centimeters before they decay. In HERA – B the average distance between the production and decay point of a b hadron is 1.1 cm. This is a clear signature, since almost all other particles are produced in the primary interaction itself. The signature to trigger on is therefore muons, electrons and hadrons with high transverse momentum (0.5–3 GeV/c), high invariant mass of pairs of these ($>$ 2.0 GeV/c^2) and detached vertices (0.5–5.0 cm).

To filter events based on these signatures, a four-level trigger system, illustrated in Figure 3.1, has been designed. No trigger is able to filter the events in the 96 ns between two consecutive bunch crossings. The events under consideration therefore have to be buffered. The time available to each trigger level is therefore limited by the amount of events, which can be buffered at that level. To limit the time, the first trigger levels only considers parts of event data and all levels are highly parallelized, so they can treat multiple events at the same time.

The First Level Trigger (FLT) [37] looks at data from every event. It is designed to find and track leptons and hadrons with high transverse momentum through the tracking chambers behind the magnet and apply p_T cuts on individual tracks and mass cuts on pairs of tracks. It is impossible at this level to look for tracks in the complete tracker and the FLT tracking is therefore based on seeds found by so-called pretriggers. These use information from the MUON chambers, the calorimeter and the High-p_T chambers inside the magnet to define search windows for muons, electrons and hadrons respectively. The FLT is designed to reduce the event rate by a factor 200 by requiring at least two tracks to be found.

The input rate to the Second Level Trigger (SLT) [38] is therefore 50 kHz. This is still too high for a general track search. The SLT therefore refines the FLT tracks in the tracking chambers behind the magnet, prolonging them through the magnetic field and into the VDS. If the tracks are also found in the VDS, the SLT requires the tracks to come from a common vertex, well-separated from the target wires, or to have large impact parameters to the target wires. The latter requirement is mainly to keep double semileptonic beauty decays, where the

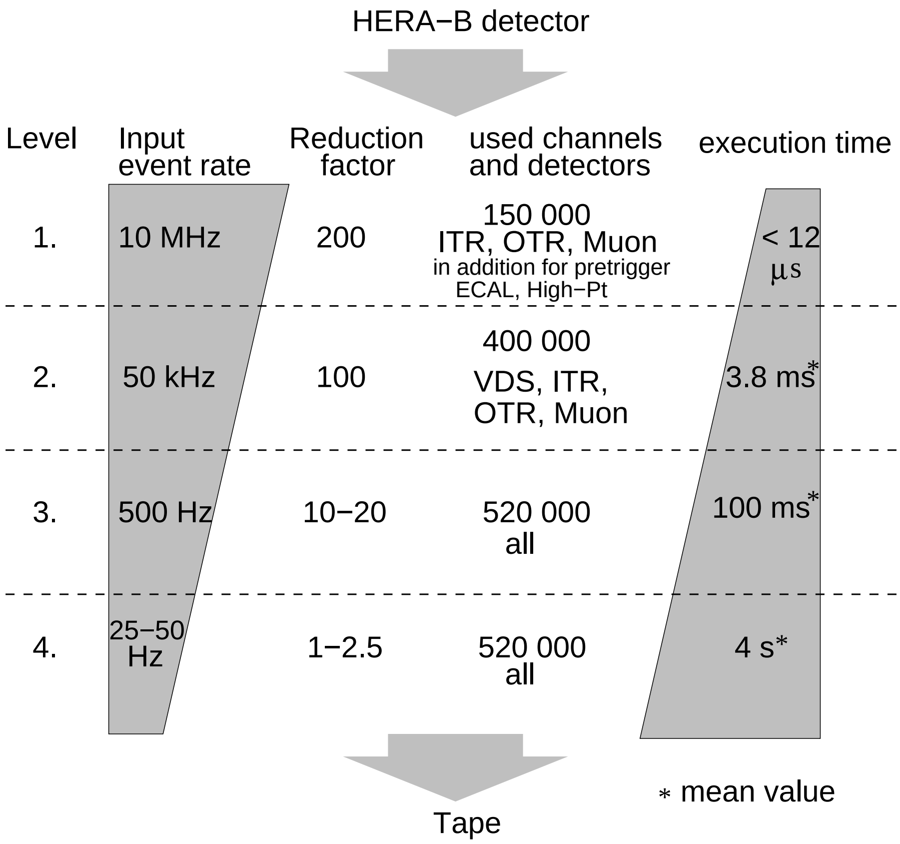

FIGURE 3.1: The four trigger levels in HERA − B with the design values for event rates, suppression and latency.

two leptons are coming from two separate detached vertices.

With a factor 100 suppression in the SLT, the Third Level Trigger (TLT) [39] has an input rate of 500 Hz. At this rate it can use data from the complete detector and is not restricted to search regions. The design calls for a complete reconstruction of the VDS to find primary and secondary vertices. By requiring the presence of secondary vertices the event rate should be reduced to 50 Hz. For both the 2000 and 2002 trigger configuration the TLT is not used.

The Fourth Level Trigger (4LT) [40] does a complete event reconstruction allowing immediate physics analysis and classification of the events into different physics classes for easy offline selection. In the design the 4LT should also reduce the event rate to tape to 20 Hz, however advances in network and storage technology mean that the full rate of 50 Hz can now be stored.

3.2 The pretriggers

The pretriggers do not reject or accept events. Instead they are designed to give the starting point for the FLT track search. The FLT searches for three types of tracks: muons, electrons and hadrons with high transverse momentum. Each type has its own pretrigger: MUON, ECAL and High-p_T pretriggers, which use information from the detector components of the same name.

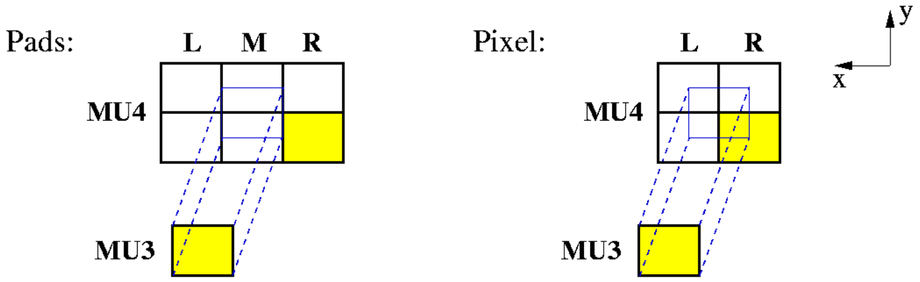

FIGURE 3.2: Coincidence schemes for the MUON pretrigger in the pad and pixel systems.

Each pretrigger searches for candidates in every bunch crossing and can give several candidates to the FLT in a single event.

3.2.1 The MUON pretrigger

The MUON pretrigger [41] uses hit information from the MU03 and MU04 superlayers at the very end of the spectrometer. Essentially only muons with momentum larger than 5 GeV/c are able to penetrate the thick absorber in front of these two layers. Between the two layers there is only a small amount of absorber and therefore little multiple scattering. A muon will therefore hit the two layers in approximately the same x, y position and this is the basis of the MUON pretrigger. From the outer part of the MUON superlayers it uses the hits in the 12×10 cm^2 large pads. In the inner part the pixels are combined into "pseudo pads" consisting of 6×4 pixel cells. If any of the pixels in pseudo pad fires, the pseudo pad is considered to have fired.

A pretrigger seed is generated if a (pseudo) pad in MU03 has fired and, in the same bunch crossing, one of the six (four) pads closest behind it in MU04 also fires. The possible pad coincidences are shown in Figure 3.2. The coincidence scheme allows the MUON pretrigger to find most muons with $p_T > 0.5$ GeV/c and forward the position of these to the FLT.

3.2.2 The ECAL pretrigger

The ECAL pretrigger [42] bases its seeds for the FLT on highly energetic clusters in the calorimeter. For each cell in the calorimeter an individual energy threshold, E_c, can be set, and if the energy in a cell is higher than E_c, a cluster search around that cell is done. The energy in the surrounding eight cells are added to the energy of the central cell and if the total energy exceeds a second threshold, E_{th}, a pretrigger seed is generated and sent to the FLT. To avoid generating several seeds for the same physical cluster, the energy in the central cell is required to be the largest of the nine cells. For the 2000 run a slightly different configuration was used, where only the cells in a cross around the central cell (see Figure 3.3) were used. This was done to avoid various problems with the pretrigger electronics, which was generating fake seeds.

The two thresholds for the pretrigger are position dependent and can be set differently for each cell. Normally the pretrigger uses $E_c = 0.5E_{th}$. The threshold should reflect the transverse momentum of the electron generating the cluster. This is complicated by the magnetic field in front of the calorimeter, which bends the electron trajectory either outward or inward in x depending on the sign of the charge. The simplest is simply to ignore the magnetic field and use

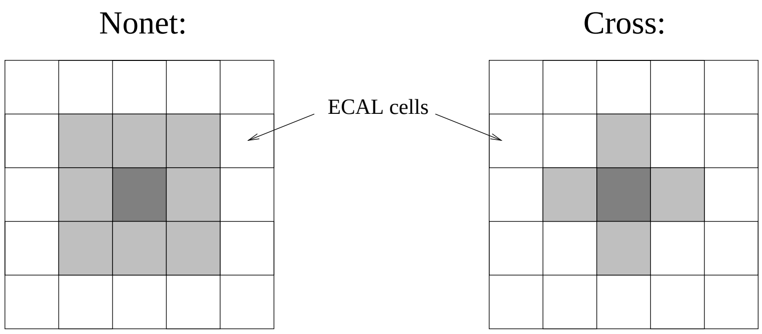

FIGURE 3.3: The two cluster types possible in the ECAL pretrigger. The center cell has to be the most energetic cell.

the "transverse energy", $E_T = E\frac{\sqrt{x^2+y^2}}{z}$, where E is the electron energy and x, y, z the position of the cluster in the calorimeter. A minimum transverse energy can then be used by setting the threshold to:

$$E_{th} = E_T^{\min}\frac{z}{\sqrt{x^2+y^2}}. \tag{3.1}$$

This definition favors electrons which are bend away from the center of the detector as their E_T will be overestimated. Empirically a slightly better definition of the threshold was found to be:

$$E_{th} = K_{trig}\left(\frac{1}{\sqrt{x^2+y^2}} + \frac{1}{\sqrt{x^2+|y|^3/\mathrm{cm}}}\right), \tag{3.2}$$

where the design called for $K_{trig} = 700$ GeV·cm. Until now only the simpler E_T threshold has been used.

An extra feature of the ECAL pretrigger is its ability to recover bremsstrahlung. When electrons pass through material they emit bremsstrahlung photons in more or less the same direction as the electron is traveling. These photons are not affected by the magnetic field and if they are emitted before the magnetic field, they will be visible as a cluster with the same y as the electron cluster, but separated in x. The separation depends on the momentum of the electron as it passes through the magnetic field and therefore on the measured electron cluster energy. The pretrigger is able to search for such bremsstrahlung clusters. If a bremsstrahlung cluster is found for a pretrigger seed, the energy in it is forwarded together with the seed cluster to the FLT. The FLT uses this to correct its estimate of the electron momentum in front of the magnet and therefore increases the p_T value and the invariant mass of possible electron pairs. As the spectrometer contains a significant amount of material, this correction can be important to achieve good trigger efficiency for electrons, depending on the p_T thresholds used in the FLT.

3.2.3 The High-p_T pretrigger

Pions with high transverse momentum are characterized by having only a small bending in the magnetic field. The High-p_T pretrigger uses this by having three chambers inside the magnetic field with a projective pad-like readout. The pretrigger uses a pad coincidence scheme similar to the MUON pretrigger, but requires hits in three layers. The scheme is illustrated in Figure 3.4,

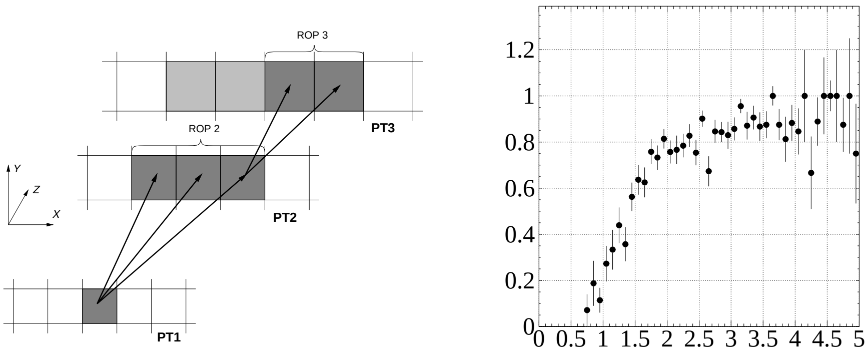

FIGURE 3.4: Left: Coincidence schemes for the High-p_T pretrigger. Right: Acceptance as function of p_T.

where a pad hit in PT1 is required to be in coincidence with one of the three pads behind it in PT2, which is then again required to have a hit in one of two pads in PT3. With the pad size used in the pad chambers this scheme gives an approximate p_T threshold of 1.5 GeV/c as seen in the efficiency plot in Figure 3.4. If a high-p_T coincidence is found, it is projected onto the TC02/MS15 superlayers, where the FLT will start its track search.

3.3 The first level trigger

The FLT uses the seeds from the pretriggers to search for tracks from muons, electrons and high p_T hadrons. It has to do this for practically every bunch crossing, where an average of five pretrigger seeds is expected per event at the design interaction rate of 40 MHz. While the FLT processes an event, the detector data is stored in a pipeline placed in the readout electronics, see Section 3.7.1. The pipeline is 128 events deep, which with 96 ns between bunch crossings means that the FLT has 12.3 μs to decide whether an event should be accepted or not. If the FLT is too late, the event data have already been overwritten by a new event. The 12.3 μs defines the maximum latency of the FLT and includes the latency of the pretriggers (between 1 and 4 μs depending on the type). In order to accomplish this, a highly parallel and pipelined architecture has been implemented.

The FLT is implemented as a simplified Kalman filter[1] as illustrated by Figure 3.5. The track search starts at the MU04 for the MUON pretrigger seeds and the TC02/MS15 superlayer for the ECAL and High-p_T pretrigger seeds. The seed defines a Region of Interest (RoI) in the superlayer, which is the area in which the track is supposed to have passed through and left hits. If a hit is found in the RoI, it is used to update the position and direction of the track and define a new RoI in the next superlayer based on the assumption that particles move in a straight line. For muon candidates this procedure, "search for hits in RoI and predict new RoI", starts with the pretrigger seed in MU04 and goes through MU03, MU01, TC02, TC01, PC04 and PC01, while for electrons and hadrons only the last four superlayers are used. If more than one hit is found in an RoI, several predictions are made for the next superlayer where the

[1]A general introduction to Kalman filters is given in Appendix A.

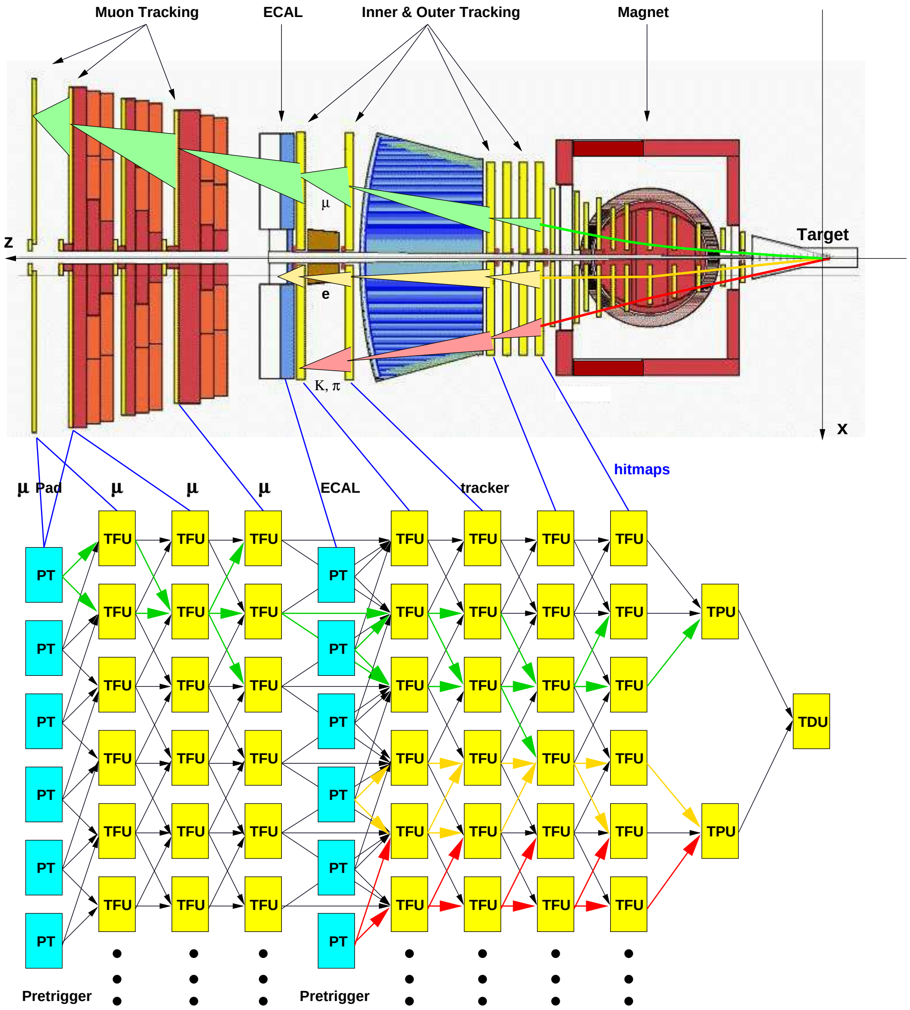

FIGURE 3.5: The FLT tracking concept.

FLT therefore have to search several RoIs. Only if hits are found in all of the superlayers is a track found. These tracks are then subjected to p_T cuts and if pairs of the same type, but with opposite signed charges and large invariant mass are found, the event is accepted by the FLT.

The FLT implements this system with three types of processing units: Track Finding Units (TFUs), Track Parameter Units (TPUs) and a Track Decision Unit (TDU). Each of the seven superlayers are split into between two and twelve regions and each region is serviced by its own TFU. A TFU receives track candidates from TFUs in the previous superlayer and searches for hits inside the RoI of these track candidates. If a hit is found, a new, refined track candidate is generated and sent to the TFUs of the next superlayer. Track candidates are represented as 80 bit messages which are forwarded through the TFU network, starting at the pretriggers and ending at the TFUs in PC01. If a TFU at PC01 accepts a track candidate it is send to one of the TPUs, which calculates the track momentum based on the track parameters found by the TFUs. Track candidates accepted by the TPU are forwarded to the TDU, which collects all the tracks and does the pairing of the tracks in an event. An event is accepted if the TDU finds a

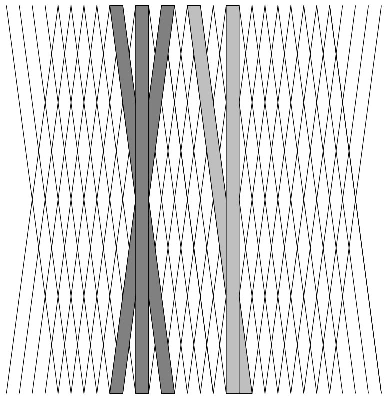

FIGURE 3.6: The TFU coincidence scheme. The lines delimit the $\pm 5°$ wires, while for clarity the empty $0°$ wires are not shown. The dark lines mark wires which fired. The three fired wires on the left constitute a coincidence and are used as a track hit by the TFU if it is inside a track RoI. The two fired wires on the right are missing a hit in the $+5°$ layer for it to be a TFU hit.

track pair which satisfy the trigger requirements. When that occurs a signal is sent to the Fast Control System (FCS) asking it to accept the event and send it to the SLT.

The whole FLT system is completely message driven and runs asynchronously from the detector. It is a highly parallel system as tracks from different events are propagating through a network of TFUs and TPUs at the same time and do not necessarily arrive time ordered at the TDU as a message from the MUON pretrigger has to pass more TFUs than ECAL and High-p_T pretrigger messages. The TFU, TPU and TDU boards have a pipeline architecture, where each message in a board has to go through a certain number of pipeline steps, each step taking one clock cycle and where other messages can be processed at the same time in other pipeline steps. A TFU for instance has 19 pipeline steps and runs with a clock frequency of 50 MHz. This means that a TFU can accept up to five messages at the input in a single bunch crossing, up to 19 track candidates can be processed at the same time and it takes 380 ns for a message to pass through a TFU[2]. At this speed, there is no time to do actual calculations and all calculations are therefore done with the help of Look Up Tables (LUTs). A calculation is done by combining the input variables, for instance a track parameter and a hit position, into a bit pattern which is used as an address for memory filled with the LUT. The value stored in the memory at that address is the result of the calculation and can be used in the next step in the pipeline. All LUTs are pre-calculated and thus remain fixed during runs.

The detector region covered by a TFU is internally represented as three layers ($0°$, $+5°$ and $-5°$) with up to 384 equidistantly spaced wires. So-called Trigger Linkboards (TLBs) placed on the detector have the task of mapping the cells of the real detector, which consists of many modules with different 5 and 10 mm cells, into this idealized geometry and for every bunch crossing transmit the hit information to the TFU. In the TFU, hits from the three layers are compared and if hits from all three layers cross each other in the same place as illustrated in Figure 3.6 it is called a “coincidence”. These coincidences are used by the TFU to update the

[2]This calculation is simplified as multiple hits can be found for a single track candidate and multiple messages have to be generated. This forces the pipelines further back to paused while this is being done and the effective rate is therefore lower.

track information. MU03 and MU04 do not have any stereo layers, so here only a single wire is required.

A track is required to have coincidences in all the superlayers used by the FLT. This puts rather strict requirements on the hit efficiency as a track is required to pass four superlayers in the main tracker (and three in the MUON system for muons). That means that the tracking efficiency scales with the hit efficiency, ε, as $\varepsilon^{3 \cdot 4}$ (4 superlayers with 3 hits each, for muons it is $\varepsilon^{3 \cdot 5+2}$), which with a hit efficiency of 98% would give a tracking efficiency of just 78% (70%) and with a hit efficiency of 95% it is reduced to 54% (41%) per track. In case of a high voltage failure in a detector layer, one would also have areas in the detector where no FLT tracks can be found. To avoid this, each layer used by the FLT is actually a double layer with separate high voltage connections. The hits from a double layer are in the TLB logical "ORed" into a single layer representations. The single wire efficiency therefore increases to $1-(1-\varepsilon)^2$ and the tracking efficiency is only reduced by a few percent, even if the single hit efficiency is only 95%. In case both layers in a double layer are dead or very inefficient in a region of the detector, both the TLB and the TFU have provisions for masking this by setting the corresponding wires in the wire memory to always contain hits. This reduces the hit requirement to just a double coincidence of hits in the two other layers and therefore increases the amount of ghost tracks. It is therefore clearly limited how much of the detector can be masked in this way.

During the 2000 run, the TFU and TPUs were also used in an alternative mode, where they just forwarded the track candidates from the pretriggers through the network to the TDU. The TDU would then trigger on events which had at least two candidates from the same type of pretrigger. This mode was used while the FLT was not ready to do the full tracking, but of course this did not allow one to have high interaction rates. For the 2002 run, a second TDU will be installed and pretriggers will send their candidates to both the TFUs and the second TDU. The second TDU can be read out from the SLT, which will therefore have access to both pretrigger messages and the FLT triggered tracks. For details on this mode, see Chapter 6.

3.4 The second level trigger

The SLT is a farm of 240 Pentium II PCs running at a clock frequency of 300 MHz (100 nodes) and 450 MHz (140 nodes). Each PC is equipped with 64 MB RAM and 100 Mbit Ethernet, but has no harddisk and runs only a minimal Linux system. The SLT runs as a process on each node and all processing of an event is carried out inside this process, called the Second Level Process (SLP). The event data are stored in the Second Level Buffer (SLB) and the SLP requests data from the SLB over a large switch. The switch and the SLB are made of DSPs and described below in Section 3.7.2.

Similar to the FLT, the SLT functions on the basis of RoIs. After the FLT accepts an event, all of the detectors sends the event data to the SLB and the event is assigned to an SLT node. This node first requests the FLT information (FLT tracks and mass of pairs), which have been sent from the TDU to the SLB. The accepted FLT tracks are used as RoIs for the SLT, which refines the tracks and follows them through the whole tracking system to the target region. The SLT algorithm for the 2000 run is described in detail in Chapter 4 and for 2002 in Chapter 6. It is split into several steps, where the complexity and the amount of event data needed increase with each step. At each step only the data from the detector inside the RoIs under investigation in that step are pulled from the SLB and processed in the SLT node. This is the opposite strategy to that of the FLT, where all event data are transfered to the processing power, but the processing there is spread over many processors. The SLT treats the complete event inside a single processor. The individual steps in the algorithm are:

- **Slicer**
  Removes FLT ghost tracks by requiring a certain number of hits in the tracking chambers behind the magnet to be lined up on an approximate straight line (inside a slice of the RoI). Only hits from $0°$ layers are used.
- **RefitX**
  Finds the track parameters in the xz plane by applying a Kalman filter to the hits used by Slicer. Tracks with too few hits or too poor χ^2 are rejected.
- **RefitY**
  Finds track parameters in the yz plane by applying a Kalman filter to the hits in the stereo layers behind the magnet.
- **L2Magnet**
  Propagates tracks from behind the magnet to the front of the magnet using a parameterization based on momentum and track slopes. Searches for hits in the tracking stations right in front of the magnet.
- **L2Sili**
  Follows tracks through the VDS using a Kalman filter approach. Track parameters are again fitted separately in the xz and yz planes.
- **L2Vertex**
  The two tracks, which have been found by the previous algorithms are required to come from a common vertex within errors. The vertex should be detached from all target wires. Alternatively, the tracks can have large impact parameter to the target wires, indicating a possible semileptonic b-hadron decay.

For the 2000 run not all steps were used. The FLT was only used for commissioning purposes and for most data taking the SLT had to base its RoIs on pretrigger information instead. This was done by introducing a seeding step in front of the rest of the SLT algorithm. The seeding step emulates part of the FLT functionality, but is less advanced. The procedure is also much slower than the FLT.

3.5 The third level trigger

The third level is also a software trigger and in fact runs on the same node and in the same process as the SLT. It is distinguished from the SLT by not limiting itself to RoIs, but instead works with data from complete subdetectors. It starts by reading out the complete event from the SLB (this is called “event building” as it is the first time the complete event is stored in the same place), which frees the memory in the SLB for storing a new FLT accepted event. It then does a full track reconstruction in the VDS and proceeds to reconstruct primary vertices from the reconstructed tracks and searches for secondary vertices. The presence of secondary vertices separated from the primary vertices by some millimeters would indicate a possible b-hadron decay and the event is triggered and send to the 4LT.

In the design, an average of 50 out of the 240 nodes in the SLT/TLT farm should be busy with TLT algorithm, which, with an estimated input rate of 500 Hz, would leave 100 ms for reading out the complete event and run the TLT algorithm. A tracking algorithm for the TLT has been written [39], but its performance in both precision and speed has since been surpassed by the algorithm used by full event reconstruction (CATS [43]). This would therefore be used instead, but the TLT was in fact not used at all for the 2000 run as there was no need for more suppression beyond that provided by the SLT. For the 2002 run prompt J/ψ events are a major part of the physics goals and a requirement of detached vertices cannot be used. It is therefore not planned to use the TLT in 2002 either, but instead have all rejection done in the FLT and SLT and after the event building, send the event directly to the 4LT.

3.6 Fourth level trigger

The 4LTs main purpose is to provide a complete event reconstruction and classification for the events accepted by the previous trigger level. This is done on a second farm of Linux PCs. For the 4LT, 100 PCs each with two 500 MHz Pentium III CPUs and 256 MB ran are used.

The event reconstruction is done inside a framework called ARTE, which mainly provides structured access in C++ and Fortran to event banks, which contain information such as raw data, hits and reconstructed tracks. The reconstruction itself is done by a long list of different packages. It starts by reconstructing partial tracks (called segments) independently in the different tracking detectors and electromagnetic clusters in the calorimeter. The segments from different detectors are then matched together into long tracks and a global refit of the tracks is done to estimate the optimal track parameters. The particle identification uses the tracks to search for rings in the RICH and to find hits in the MUON detector. The primary vertices are reconstructed from tracks which pass close to the target wires. All track and vertex information and some global event information are stored together with the raw data and no further reconstruction of the event is foreseen.

Based on the tracks and vertices found, the events are also classified into different groups, which allows quick selection later, so that not all events needs to be read for an analysis. Originally it was foreseen to have classes such as "golden decay"-like or "$b \to \pi^+\pi^-$-like. For the 2000 run more modest classes such "di-muon event" or "contains K^0_S candidate" were used. It was also planned to reject more than half of the events based on the event classification, but the event logging has proven itself capable of handling the full rate and the 4LT is no longer planned to be used as a trigger. The online reconstruction also allows an advanced data quality monitoring to be carried out in real time.

3.7 Data acquisition system

The Data AcQuisition system (DAQ) is intricately connected with the trigger system. An overview of the data paths from the detector to the different trigger levels and to the final event storage is provided in Figure 3.7. The complete detector is read out every 96 ns and the data is stored in 128 cell deep pipelines in the Front End Drivers (FEDs), most of which are either inside the detector or just next to it. The data for the pretriggers and the FLT are transmitted over a dedicated data path for every event. If the FLT accepts an event within 12.3 μs, the event is send from the FEDs to the SLB. In the SLB, data from different detectors are kept in different places, just as with the FEDs, and the SLT has to request data from several places in the SLB. An SLT accepted event is assembled in a single TLT node, from where it is send via an Ethernet network to first the 4LT and then to a logging machine, if all trigger levels accept the event. In the following sections the DAQ is described in more detail.

3.7.1 Front end electronics

The front end (or readout) electronics has the task of amplifying, shaping and digitizing the raw signals coming from the detector. It also stores the data while the data is being processed by the FLT and sends the data of accepted events to the SLB via the FEDs. The front end electronics in HERA – B is highly unified and only a few different readout chip types are used.

The VDS and the ITR use the HELIX chip [44] to read out the micro-strips. Each HELIX chip handles 128 input channels and for each channel consists of a charge integrating preamplifier followed by a shaping stage, which stores the measured charge in a pipeline of 141 capacitors. On an FLT accept, the 128 channels are read out sequentially via an analog multiplexer and transmitted via an optical fiber to the electronics trailer. Each fiber transmits data from two chips. In the trailer, the analog signals are digitized in the FEDs with a 10 bit ADC and the eight most significant bits are sent to the SLB for the 256 channels. The HELIX chip also contains a prompt digital output, which is set if the measured charge is over a certain threshold. This is used by the ITR to send hit information to the FLT for every bunch crossing.

The ECAL signals are transmitted over coaxial cables from each of the 8000 photomultipliers tubes to the readout boards in the electronics trailer 50 meters away. The signals are digitized with 12 bit ADCs and stored in dual ported memory, while waiting for a trigger accept from the FLT. The 12 bit ADC values are also transformed via a channel-specific LUT into an eight bit value to be used by the ECAL pretrigger. The first 7 bits encode the energy of the cell, while the 8th bit is set if the energy is higher than E_c, so that the ECAL pretrigger knows that this is a potential center of a cluster. Each channel has its own LUT in order to be able to take the calibration into account properly. The 8 bit values for all channels are transfered to the ECAL pretrigger in every bunch crossing.

The remaining subdetectors use the ASD8 chip [45] in the readout. Each chip provides preamplification, fast signal shaping with tail cancellation and a discriminator for eight channels. The signal threshold in each of the eight discriminators is set by external voltage source, which can be remotely controlled. The signals from the ASD8 are transmitted via shielded twisted pair cables to a Time-to-Digital Converter (TDC), which converts the drift time into an 8 bit value. Only the OTR uses the drift time, all other detectors use just one bit (signal or not). The TDC boards are located just outside the detector and the signals are stored in dual ported RAM as for the ECAL. On an FLT accept, the data is transfered to the SLB in the electronics trailer over up to 50 meter long cables.

For the OTR, ITR, MUON and High-p_T systems, one bit signals are also sent to the Trigger Link Boards (TLBs), which maps the hits into the FLT representation of the detector and does the ORing of hits from double layers. The data are sent from the TLBs to the pretriggers and the FLT via 900 Mbit/s optical links [46]. About 1500 links are used to send almost 1 TBit of data from the detector to the FLT each second.

The data are transfered from the FEDs to the SLB using twisted pair cables. The amount of data and rate of transfer over each cable depends on the subdetector and has been listed in Table 3.2. The transfer rate is one of the factors, which limits the maximal rate of accepted FLT

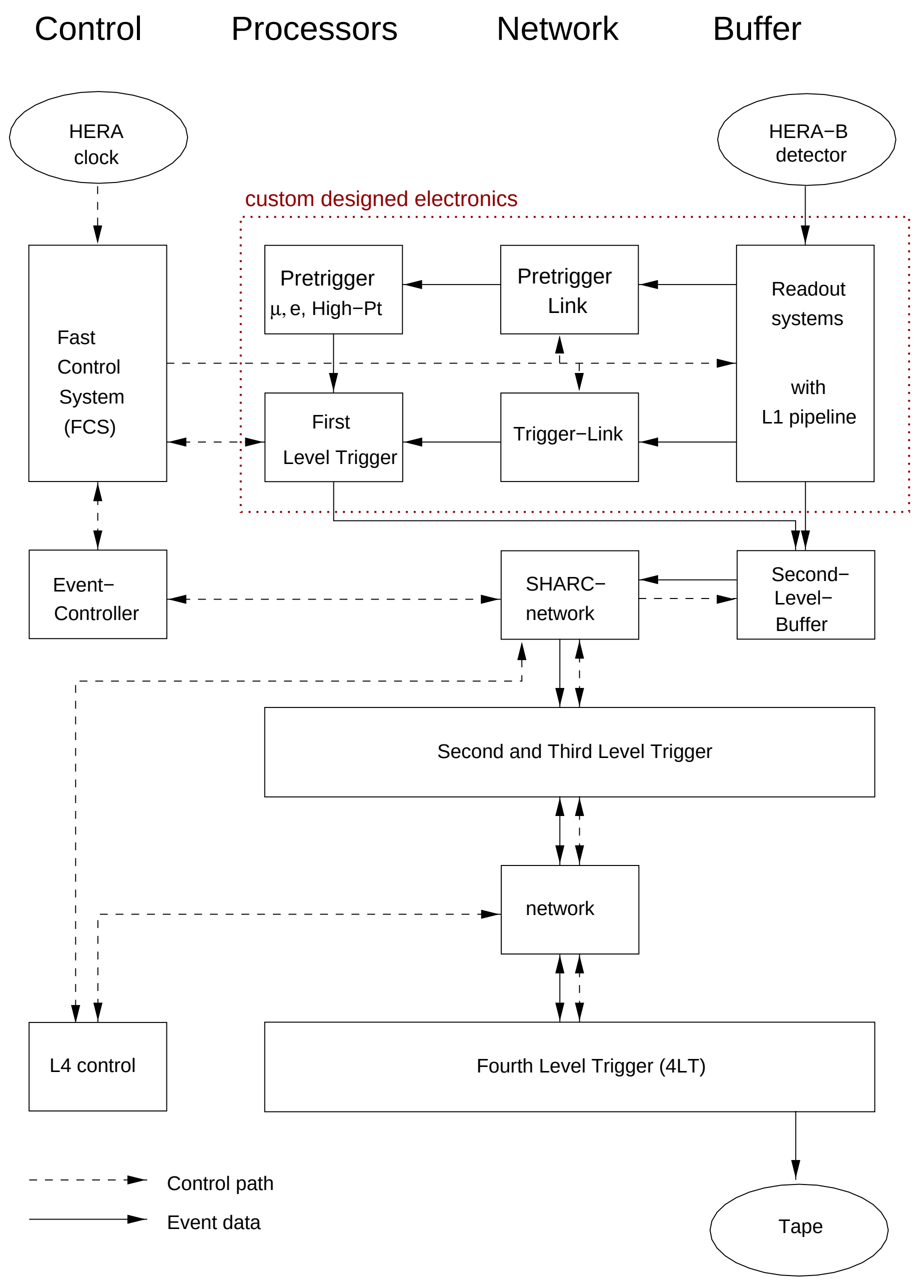

FIGURE 3.7: The data and control paths in the HERA – B DAQ system.

Detectors	Transfer rate	Bytes per link	Rate limit
VDS+ITR	20 (40) MByte/s	300 Byte	66.5 (133) kHz
ECAL	20 MByte/s	64 Byte	313 kHz
All other	15 MByte/s	144 Byte	104 kHz

TABLE 3.2: Transfer rates and data amounts for different FED→SLB connections [36].

trigger events as the system should have enough time to transfer the data from the FEDs into the SLB between two FLT triggered events. The rate is limited to 66.5 KHz due to VDS and ITR FEDs, which have to transfer 300 Byte per event at 20 MByte per second. This limit actually comes from the data transfer from the HELIX chip to the electronics trailer. The HELIX chip is designed to supply the data at up to 40 MByte per second, but during the 2000 run this mode was not able to run stable and data were transfered at just 10 MBytes per second.

The time between two FLT accepted events is randomly distributed, so even if the average FLT rate is below 66.5 kHz, one often gets an FLT accepted event, while the system is already busy transferring another FLT accepted event to the SLB. The system cannot wait until the transfer finishes as the data from the second event will then most likely have been overwritten in the FED pipeline. In order not to reject such events, all FEDs are constructed with so-called readout-FIFOs in which the data from up to 8 events are temporarily stored while waiting to be transmitted. This is the reasons the HELIX is equipped with 141 capacitors per channel as it needs 128 for the pipeline and 8 for the readout-FIFO. The five extra are needed because it takes 5 bunch crossings to transfer the signal to the analog multiplexer in the HELIX. The effect of the readout-FIFO on deadtime is discussed below in Section 3.7.4.

The system is controlled by the Fast Control System (FCS), which makes sure that the readout of all the detectors are synchronized to the same bunch crossing. It supplies a 10.4 MHz signal to all readout electronics triggering the readout of the detector into the pipeline. The signal has to be adjusted individually for the different subdetectors as there are different delays in the electronics and particles arrive later at the last layers compared to the first. Muons, for example, traveling at light speed take about 65 ns to pass through the complete detector. The FCS also supplies the bunch crossing number to all readout electronics for each event. This is used to identify the events and accompanies all data transfered to both the FLT and the SLB. The FCS is also responsible for receiving trigger signals from the FLT. If the trigger signal is not too late (within about 12 μs of the collision) and the event controller (see below) is not blocking triggers, the FCS signals the FEDs to start sending the event to the SLB.

3.7.2 Second level buffer and switch

The main component in both the SLB and the switch connecting the SLB to the SLT/TLT farm is the Digital Signal Processor (DSP), ADSP-21060, from Analog Devices [47], which runs at a 40 MHz clock cycle. It is commonly known as a SHARC, which stands for Super-Harvard-ARchitecture-Computer. The SHARC is a highly integrated chip, which is especially suitable for high rate data transfers. It has the following characteristics [36]:

- A main CPU capable of two memory accesses or calculations in a single clock cycle.
- A 512 kBytes internal memory usable for both program and data storage.

- A 48 bit wide external bus for connecting it to other SHARCs or external units. The bus can sustain transfer rates of up to 240 MBytes/s and allows a SHARC to access the memory of other SHARCs in a single clock cycle.
- Six data link-ports, 4 bit wide and running at up 80 MHz, thus allowing transfer rates of up to 40 MBytes/s per link.
- An independent I/O processor with DMA controllers for the link ports and external bus. This allows data transfers to be setup and carried out independently of the main CPU.

For the DAQ the most important feature is the six fast links and the I/O processor controlling them. These links are both used to receive the data from the FEDs and for the data transport in the switch.

The building block for the SLB and the switch is a SHARC module, which is a VME card containing six SHARCs connected to each other over a shared bus. The 6 times 6 link ports can be connected to different driver cards, depending on whether they receive data from the different FEDs or are used to communicate between different SHARC modules. In HERA – B all link ports always transfer in the same direction to simplify the programming.

140 of the SHARC modules are used to build the complete SLB. Each SHARC receives data from up to six FEDs over the data links (one FED per link, but not all links are always used). At least two of the links are used to communicate with the switch. The data received from the FEDs are stored in the internal memory of the SHARC and in total data from 280 events can be stored in the SLB at the same time. Unlike the FEDs where the data are kept in pipelines and therefore filled in a fixed, sequential order, the data in the SLB are kept for widely varying time periods. Some events are rejected after less than a millisecond, while other might have to be kept for several hundred milliseconds. Each SHARC therefore keeps a list of free "buffer indices", which are updated from the EVC. The buffer indices are used to setup the memory addresses for the DMA controllers for the links to the FEDs, so that the data can be received without interfering with the main CPU.

All communication between the SLT farm and the SHARCs in the SLB occurs through a message system called RPS. An RPS message can maximally be 1200 Bytes long and contains a header, which gives the destination and the origin of the message. When a SHARC receives a message over either the link ports or the shared bus, it first checks whether the message is destined for itself. If not, the message is forwarded to another SHARC according to an internal routing table. The routing of messages is completely DMA and interrupt controlled and therefore independent of the main process running on the SHARC, which only has to act on RPS messages destined for itself. The protocol also allows for multicast messages, where a received message is multiplied and sent to multiple destinations with each SHARC receiving it. This is, among other things, used by the EVC to distribute control messages.

To allow the SLT farm to communicate with the SLB, a special SHARC-to-PCI interface card has been constructed and installed in every SLT node. It gives each node access to two SHARC links, which are used to send and receive RPS messages from the SLT process via a specially constructed Linux kernel module. The 240 SLT nodes need to be able to request data from any of the 140 SHARC modules that are used in the SLB. This is handled by connecting the nodes and the SLB through "switch boards" constructed from SHARC modules as illustrated in Figure 3.8. The SLT nodes are divided into 12 processor blocks and the SLB modules into 10 buffer blocks. Each block has two switch boards, one for receiving messages and one for

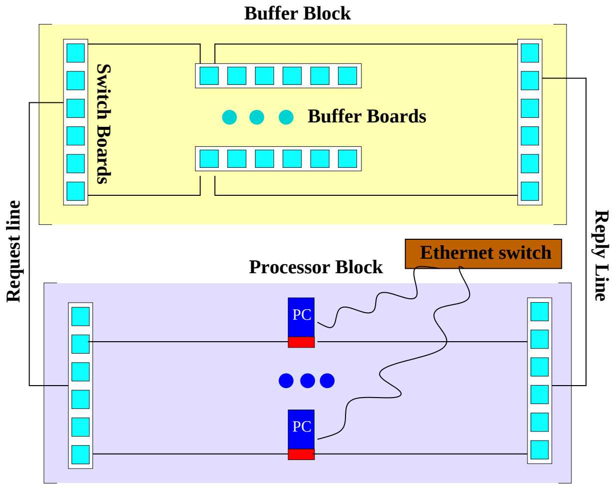

FIGURE 3.8: The connections between a processor and a buffer block.

transmitting messages. Both switch boards are connected to all modules/nodes in the block and the transmitter board of a processor block is connected to the receiver boards in all 10 buffer blocks and similarly for the transmitter boards in the buffer blocks. This way an SLT node can send a request message to any SLB SHARC by routing it through two layers of switch boards and the answer comes through two other layers of switch boards.

The SLT nodes send two kinds of requests to the SLB. For the tracking through RoIs it sends "range requests", where only a specific range of data from an FED is requested. For the event building for the TLT the messages request all data for a given event in a buffer module. To keep the number of messages down, the range requests can ask for data from several SHARCs in the same SHARC module. The first SHARC in the module to receive the message will fill in its part and forward the message to the next SHARC. The replies to range requests are typically only 140 Bytes long, while for event building the messages are about 800 Bytes. The performance of the switch is limited both by the capacity of a single link (40 MBytes/s) and the overhead of routing the messages (around 5.6 μs). For large messages the capacity dominates and a maximum of 32 MBytes/s has been obtained. Small messages are limited by overhead to about 80 kHz, giving a bandwidth of 11 MBytes/s for RoI replies. The bandwidth of the complete network has been measured to 1.2 GBytes/s with $2.6 \cdot 10^6$ messages going through the system every second. This is more than enough for the SLT, which has been estimated to need 500 MBytes/s and a message rate of about $1.8 \cdot 10^6$ [48].

The EVent Controller (EVC) [36] is the process responsible for controlling the SLT and SLB. It is implemented as a multiprocessor process on a SHARC module, where the different tasks are distributed over several SHARCs in the board. It receives the trigger signals from the FCS and assigns the processing of the event to a free SLT node. At the same time it has to keep track of which nodes and buffer indices are available. The latter has constantly to be distributed to the complete SLB as the data from the FEDs arrive at the SLB modules at approximately the same time as the FCS signal in the EVC and the buffer indices therefore have to be determined well in advance. In order not to "stress" the SLB SHARCs, the list of free buffer indices are updated in blocks of 20 indices using the multicast protocol. This has some latency and the EVC has to ensure there are always at least around 50 buffers free, so that one of the SHARCs does not suddenly run out of free buffers, because the message with new indices is still propagating through the switch. If the number of free buffers gets dangerously low or there are no free SLT nodes left, the EVC can block the FCS for accepting triggers in a few milliseconds. The current implementation of the EVC is able to handle trigger rates of up to 70 kHz.

3.7.3 4LT and datalogging

The bandwidth requirements from the SLT/TLT to the 4LT and beyond are much lower than in the first trigger levels. During event building the event size is compressed to about 100 kBytes and the addition of reconstructed track information in the 4LT roughly doubles the event size. With a 50 Hz event rate, about 5 MBytes/s is needed between the SLT and 4LT farms and about 10 MBytes/s for the logging. This can be handled by ordinary Ethernet connections.

The PCs in the two farms are connected via fast Ethernet (100 Mbit/s) to large Ethernet switches and the connections between the switches and to the logging machine in the DESY computer center are handled by gigabit Ethernet. The transfer protocol is based on TCP/IP and sustained rates in excess of 15 MBytes/s have been seen. The assignment of events in the SLT/TLT to a 4LT node is done by the L4 controller. This is a dedicated PC which keeps track of the available 4LT nodes and receives request for a 4LT node from the SLT nodes. The data transfer in the 4LT is not handled directly by the ARTE process, but rather by two separate processes, which store/read the events in shared memory accessible by the ARTE process. This gives better performance and stability.

The logging machine (in 2000 an SGI server, in 2002 a dedicated Linux PC) does not write the data directly to tape, but rather stores the events on a large disk in files of about 400 MBytes. These files are then copied to tape by a separate process. This buffering allows the logging to tape not to run at the full 10 MBytes/s, since it can also use the time between fills to copy to tape. This is sometimes necessary as the tape robot is shared between experiments and also used for offline access, which gives varying access time and bandwidth to tape.

3.7.4 Deadtime

Deadtime is defined as the time where the experiment is not able to treat an event, because somewhere the system is busy handling a previous event. The interactions occur every 96 ns and all detector components are able to sample at that rate. The trigger and DAQ system are designed to run with virtually no deadtime, but some deadtime can still occur if trigger rates are higher than expected. The rejection of events due to deadtime is done in the FCS, which keeps track of how many events were lost due to this.

Deadtime can occur for several reasons:

- **FLT latency:** The pretriggers and the FLT receive and process every event, but run asynchronously of the detector. The TFUs can only handle up to five track candidates per bunch crossing and in case of events with very high occupancy some congestion in the FLT network can occur. This can cause an event to be triggered so late that the pipelines in the FEDs containing the event data are already overwritten and the event has to be rejected.

- **FIFO overflow:** If an event is triggered, while a previous event is still being transfered from the FEDs to the SLB, it is stored in the readout-FIFO. This FIFO however has a limited size and when it is filled, events have to be rejected. With a depth of eight buffers this is only a problem when the trigger rate is close to the maximum transfer rate between the FEDs and the SLB (66 kHz for ITR and VDS FEDs). In the 2000 run, the system was however not able to use the FIFO due to a problem with the ITR. This gave some deadtime even at low trigger rates and was a major concern in the 2000 run. Figure 3.9 shows the deadtime as a function of trigger rate for different FIFO depths.

- **EVC blocking:** The EVC can block the FCS for accepting events. This can occur for several reasons. If the SLT is too slow to reject events either because of lack of CPU power or having to wait for data to arrive over the switch, the SLB fills up and events have to be rejected. It can also be that the 4LT is too slow and the SLT nodes will therefore be busy waiting for a 4LT to become free. Finally it can in principle be a lack of bandwidth between the SLT and 4LT or the 4LT and the logging machines, which again will keep the SLT nodes waiting. For parts of the 2000 run the 4LT was on average using more than four seconds to reconstruct an event and therefore not able to keep up with the rate.

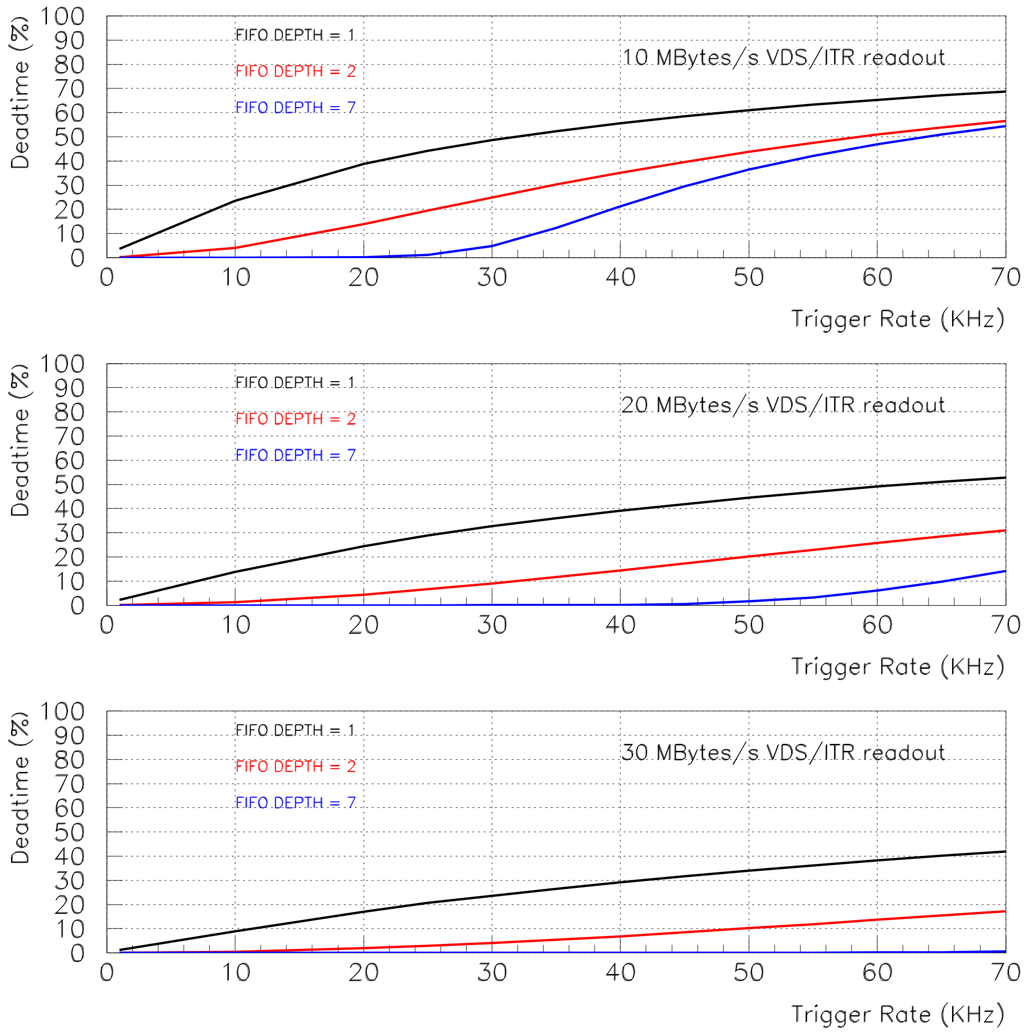

FIGURE 3.9: The deadtime due to overflow of the readout FIFO [49] at different trigger rates from the FLT. The three plots shows the deadtime at different readout speeds of the VDS/ITR. The top plot with a FIFO depth of 1 buffer corresponds to the 2000 run, while the middle plot with a FIFO depth of 7 buffer is the plan for the 2002 run.

Chapter 4

SLT algorithm for the 2000 run

The SLT is designed to refine the charged tracks found by the FLT by tracking them through the full spectrometer as described in Section 3.4. However for most of the 2000 run the FLT was still being commissioned and not able to provide tracks for the FLT. In order to get a sample of dilepton triggered events, the SLT algorithm therefore had to be extended and adjusted. The ECAL and MUON pretriggers were ready for data taking and, by routing the pretrigger candidates directly to the TDU, they could be used to trigger events. The SLT would use the pretrigger candidates to setup RoIs for the SLT tracking algorithms. Because the SLT is implemented completely in software it is flexible enough to allow quite complex seeding algorithm.

In this Chapter the seeding procedures and tracking algorithms used during the 2000 run will be described in full detail.

4.1 SLT configurations in year 2000

Although most of the detector was installed by the beginning of year 2000 there were constant changes and improvements throughout the run period up until the very end of data taking in August 2000. Consequently, not only the performance of the SLT, but also the configuration of the SLT was changed several times during the 2000 run.

The status of the detector components used by the SLT, during 2000 run, was as follows:

VDS: Fully installed and working for run 2000. However a problem with water condensation in the last superlayer (SI08) just in front of the magnet, made a large fraction of this superlayer unusable from early in the 2000 run.

OTR: Fully installed in January 2000. Commissioned and reasonably aligned in April 2000. The hit efficiency was low ($\sim 90\%$) due to many dead channels and noisy electronics.

ITR: Only four complete layers after the magnet were installed in the 2000 run. These chambers were working in July and August, but the SLT was not able to use them because the hit preparation routine for the SLT was not ready.

ECAL: The calorimeter was fully installed well before the 2000 run. Readout and pretrigger electronics for the inner and middle calorimeter was complete by May 2000.

MUON: Fully installed in November 1999, but the gas system was completed only in April 2000. The pretrigger system for Muon pad system was working from April, while the pretrigger system for the Muon pixel system was not installed for the 2000 run.

FLT: Ready for data taking in August 2000. Severely influenced by the low hit efficiencies in the OTR and had problems due to unstable data transmission. The system was able to forward pretrigger information to the SLT and to trigger on events with two or more messages from the same pretrigger already in 1999.

This situation gave rise to the following four main trigger configurations which, unless specified otherwise in the text, were used at a target interaction rate of about 5 MHz:

ECAL+L2Sili: This trigger configuration was used already in the fall of 1999. Its purpose is to trigger on electrons using only information given by ECAL and VDS. The FLT requires two clusters with high transverse energy to be found by the electron pretriggers in the calorimeter, but does not do any tracking. In the SLT these clusters are required to have a large invariant mass. Assuming the energy deposited in ECAL is equal to the particle momentum and the origin of the particle is at the target, an estimate of the trajectory through the detector is constructed. This is used to obtain an RoI for L2Sili, ignoring the main tracker completely. L2Sili is required to find a valid track for both clusters, but no vertex requirement was imposed.

ECAL+SLT: This configuration is similar to the previous except that OTR information is also used. The RoIs, calculated from the ECAL clusters, are processed by the Slicer and Refit algorithms and the output from Refit is propagated through the magnet to L2Sili.

MUON+SLT: Two MUON pretrigger candidates are required to be found and these are used as seeds for a tracking algorithm in the MUON chambers. The tracking algorithm is done in the SLT. The muon tracks are in turn used to setup RoIs for Slicer and the algorithm continues as for electrons. This trigger configuration is designed to run in parallel with the "ECAL+SLT" configuration described above. This is possible since only the seeding of Slicer is different.

FLT+SLT: This is the trigger configuration HERA – B was designed to use. It was used only during the last weeks of the 2000 run. The FLT finds tracks in the OTR seeded from the ECAL and MUON pretriggers. These tracks are send to the SLT, which use them to define RoIs for Slicer. Most of the data taking in this mode was done triggering on just a single very high p_T muon or electron. In this mode the trigger rejection is low and the limiting factor is the data logging rate. It was therefore used at a low interaction rate (2 MHz). The FLT was also used to trigger on a pair of leptons with an invariant mass cut, which were then send to the SLT. These runs have mainly been used to commission the FLT and the SLT performance for this configuration will not be evaluated in the next chapter.

4.2 SLT Seeding

The seeding of the SLT, which means setting up RoIs for Slicer, is the first part of the SLT. Originally it was designed to simply use the tracks found by the FLT as seeds, but the late arrival of the FLT caused a much more complicated seeding system to be constructed. Because of this it is usually not considered a part of the "SLT algorithm", even if it is done inside the SLT process.

In the following the different seeding procedures are described.

4.2.1 Electron pretrigger seeding

The first seeding used for the SLT is based on only calorimeter information. The SLT searches for $J/\psi \to e^+e^-$ decays, whose characteristic signature is two high-p_T electrons with large invariant mass.

Before the events reach the SLT they are filtered by the electron pretrigger in order to reduce the input rate to the SLT. The electron pretrigger, which is described in more details in Section 3.2.2, finds calorimeter clusters with high transverse energy. In the 2000 run, the threshold was set to 1.2 GeV for the ECAL-VDS trigger and to 1.0 GeV for the ECAL-SLT trigger. The transverse energy is calculated assuming no bending in the magnet, which is necessary since the charge of the particle cannot be determined from the cluster alone. The pretrigger information is sent to the TDU, which requires the presence of two high E_T clusters in order to forward the event to the SLT.

The fastest solution to seed the SLT algorithm would be to use the cluster information transmitted from the electron pretrigger to the SLT via the FLT hardware. Unfortunately during the 2000 run this transmission was not reliable. This problem is overcome by applying a full calorimeter reconstruction in the SLT. This implies reading out the complete calorimeter followed by the reconstruction of clusters using an adapted version of the offline calorimeter reconstruction code, CARE [50]. The SLT version of CARE uses less internal precision and start the clusters search only from calorimeter cells with higher energy than in the offline version. The changes have little impact on the ability to reconstruct the high E_T clusters, while the latency of the algorithm is significantly reduced.

The E_T of a cluster is calculated as in the pretrigger with no bending in the magnet assumed and a minimum transverse energy of $E_T \geq 1$ GeV$/c$. This allows us to find the same clusters as the electron pretrigger, while loosing a minimal amount of real clusters. The distribution of E_T in data and in simulated $J/\psi \to e^+e^-$ events can be seen in Figure 4.1.

The E_T cut ensures that there are at least two high-E_T clusters in an event. These clusters are combined into pairs with invariant mass larger than 2 GeV$/c^2$. To calculate the invariant mass of a pair we need to know the opening angle between the two tracks in front of the magnet and therefore to know their bending in the magnetic field. To calculate the bending in the magnet, a charge has to be assigned to each cluster. The most general option would be to use each cluster in two seeds with opposite charges, giving four possible pairs for each two clusters. Due to the high E_T cut already imposed on the clusters, most of the $J/\psi \to e^+e^-$ ($\sim 90\%$) are in one specific combination: the out-bender e^+e^- combination. An out-bender pair is a pair of opposite charged particles whose trajectories are "pushed" apart when they pass through the magnetic field. The E_T of the tracks in such a pair, calculated ignoring the magnetic field, will be larger than the true E_T. In the HERA – B magnetic field configuration the e^+ from an out-bender pair has a larger x-coordinate at calorimeter than that of the e^- and the charges can be assigned accordingly. There will be clusters which have either charge assigned: in the case of three clusters with x-coordinates, $x_1 < x_2 < x_3$, the second cluster is assigned two different charges, e^+ when it is combined with the first cluster and e^- when it is combined with the third.

The effect of the magnetic field is approximated with a p_T kick of 650 GeV$/c$ in x at the

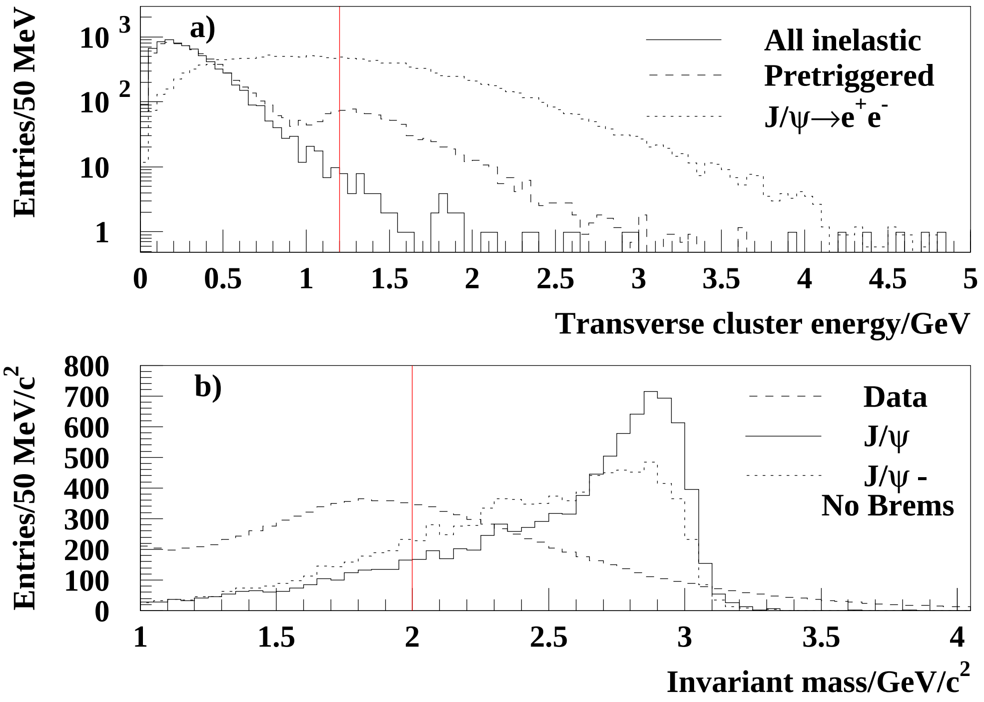

FIGURE 4.1: a) The E_T distribution in all inelastic events, electron pretriggered events and in simulated $J/\psi \to e^+e^-$ events. The vertical line shows the E_T cut used in the ECAL seeding of the SLT. b) The invariant mass distribution in electron pretriggered events and in simulated $J/\psi \to e^+e^-$ events. In this figure the degradation of the invariant mass spectrum in $J/\psi \to e^+e^-$ events is also shown when the correction for bremsstrahlung has not been included.

center of the magnet ($z = 450$ cm), allowing us to calculate the invariant mass of the e^+e^- pair. A complication with electrons compared to muons is the emission of bremsstrahlung photons during the passage through detector material in front of the calorimeter. The bremsstrahlung is emitted in three regions: before the magnet, inside the magnet and after the magnet. The radiation is emitted almost collinearly with the direction of the electron momentum, which allows us to measure with the calorimeter the emission in the first and third region. In the first region, the emitted photons hit the calorimeter at the place the electrons would have hit in the absence of a magnetic field, while radiation from after the magnet is detected at the same place as the electron and the energy is therefore automatically assigned to the electron. The radiation emitted in the magnet is smeared out over the calorimeter region spanned by the two other types and is thus almost impossible to recover.

Since we do a full reconstruction of the calorimeter with CARE we are able to search for bremsstrahlung clusters in the calorimeter. The search window is found by extrapolating the track direction before the magnet (which was already calculated using the p_T kick approximation) to the z position of the calorimeter and search for clusters within 1 cm in y and 10 cm in x. We can use a much narrower window in y as there is essentially no bending in the yz plane and the limitation in y is therefore given by the precision of the cluster position and not the track extrapolation. Before accepting a cluster as bremsstrahlung, we check that it is not identical to the other electron cluster in the e^+e^- pair. If a bremsstrahlung cluster is found, its energy is added to the electron/positron energy in the invariant mass calculation. This correction moves an extra 3% of $J/\psi \to e^+e^-$ events over the 2 GeV/c^2 threshold. The calculated invariant mass spectrum in data and $J/\psi \to e^+e^-$ events is shown in Figure 4.1 after the $E_T \geq 1.0\,\text{GeV}$ requirement has been applied.

When the high E_T clusters have been found and combined into e^+e^- pairs, RoIs are calculated for the following algorithm. Instead of the simple p_T kick approximation used for the invariant mass calculation, the field parameterization of L2Magnet is used to increase the precision. By keeping the two endpoints of the particle trajectory fixed to the target wire and the ECAL cluster, respectively, and assuming the momentum to be equal to the measured cluster energy, the L2Magnet routine is used iteratively to find track slopes before and after the magnet. The uncertainties on the target position[1] and the cluster position/energy are propagated through the calculation of track slopes and are used to define the size of the RoI. The RoIs are either given to Slicer (in the ECAL+SLT trigger) or propagated once more by L2Magnet into L2Sili (in the ECAL+L2Sili trigger).

4.2.2 Muon pretrigger seeding

The MUON pretrigger, described in more detail in Section 3.2.1, searches for coincidences between hits in the pad chambers, MU03 and MU04. Accepted events can are then send to the TDU, which triggers on an event if at least two coincidences are present. If a charged particle passes close to the edge of a MUON pad, the adjacent pad will often also record a hit, giving rise to two pretrigger coincidences for the same particle. This can be seen in in Figure 4.2, where the distance between two pretrigger coincidences is shown. These fake double coincidences can be removed with a simple distance cut without affecting the $J/\psi \to \mu^+\mu^-$ efficiency. Since the TDU is not able to impose a cut on the distance, the first step in the SLT is to reject all muon coincidences separated by less than 50 cm at MU03. This cut rejects approximately 50% of the muon pretriggered events.

The muon tracking uses hits from the 0° layers of MU01 and MU02. Due to the aluminum walls separating the tube cells, a cell in a single layer of the muon detector is only sensitive over approximately 85% of its width. To achieve $\sim 100\%$ hit efficiency, the muon chambers contain two layers shifted by half a cell size with respect to each other. If a particle passes through the insensitive area of the first layer it will pass through the sensitive area of the second layer, and vice versa. In MU02 we therefore use hits where only one layer fired. When cells in both layers fires, they are counted as a single hit. In MU01 the occupancy is more than four times higher than in MU02, as there is less absorber material in front. To have high purity, we therefore require that cells in both layers fires, which effectively doubles the insensitive area and the hit efficiency is reduced to 72%. Hadrons will shower in the ECAL and in the absorber in front of the muon chambers. If they are energetic enough, part of this shower can reach the muon layers. It will typically be spread over several neighboring cells and in order not to get a large number of combinatorics, we require that cells next to a hit are empty.

The principle of the muon tracking is to require hits in MU01 and MU02 to lie roughly on a straight line with the MUON pad coincidence. The x coordinate of the MU02 hit is required to lie within 20 cm of the MUON pad coincidence when projected back toward the target position.

[1]The uncertainties on the target position are described in Section 4.4

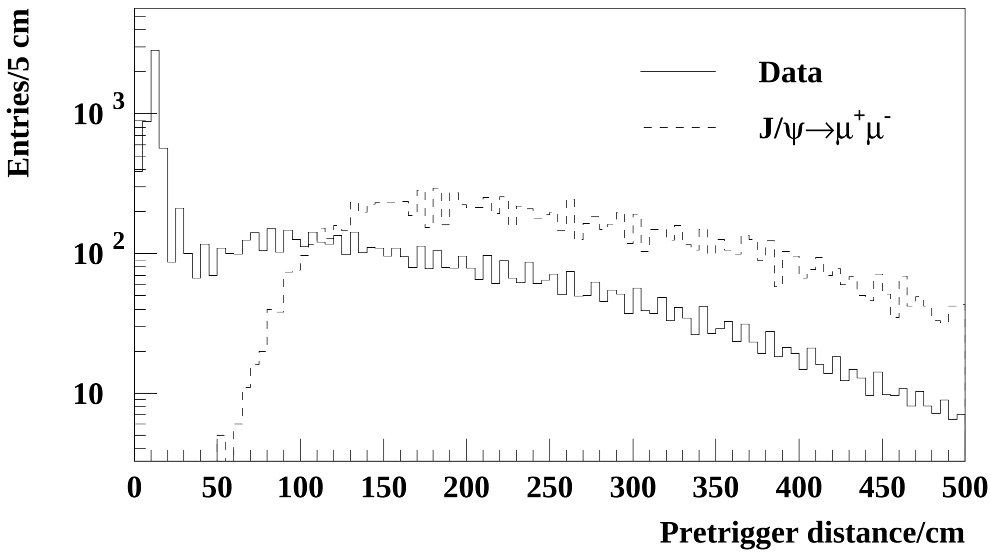

FIGURE 4.2: The distance between two MUON pretrigger coincidences. The peak at 10 cm is from coincidences in adjacent pads, which fired due to the same charged particle. Also shown is the distribution for muons from simulated $J/\psi \to \mu^+\mu^-$ decay. The vertical line shows the 50 cm cut applied by the SLT.

The hit found in MU02 is combined with all hits in MU01. The line connecting a hit in MU01 and MU02 is projected back to the pad in MU03 and it is required that the pad is within 20 cm of this line. The 20 cm cuts are adjusted according to the residuals observed in real data, shown in Figure 4.3. The residuals are dominated by the 12 cm wide pads. No cuts are done in the y-coordinate, except that all hits must reside in the same half-plane ($y > 0$ or $y < 0$) as the muon pad coincidence. To reduce combinations, only the first track found for each pad coincidence is used, which unfortunately makes the efficiency dependent on the hit occupancy.

When setting up the seeds for the SLT algorithm, we do not have an energy estimate for the muons like in the electron case. We have however an estimate of the track direction from hits in MU01 and MU02. It turns out that this track direction has very large uncertainties when extrapolated to PC01, due to the large amount of multiple scattering. Given an assumption for p_T and charge of the muon, the track position in MU01 and the fixed p_T kick in the magnet, one can calculate the energy the track must have had. Using this energy, the trajectory can be calculated in the same way as for electrons. Unlike in the electron case we make both a μ^- and μ^+ charge assumption for every muon candidate. By using an assumption of $p_T = 1.6\,\mathrm{GeV}/c$, the predicted trajectory is much better than when using the track direction. This can be seen from the comparison to the Monte Carlo in Figure 4.4. The RoI size is calculated by varying the p_T assumption between $0.7\,\mathrm{GeV}/c$ and $2.5\,\mathrm{GeV}/c$, which results in RoI sizes between 3 cm for the muon tracks closest to the beam pipe and up to 50 cm for the outer most tracks. The trigger efficiency as a function of muon p_T is shown in Figure 4.5. It does not have very sharp cutoffs, due to the large uncertainties present in the p_T calculation.

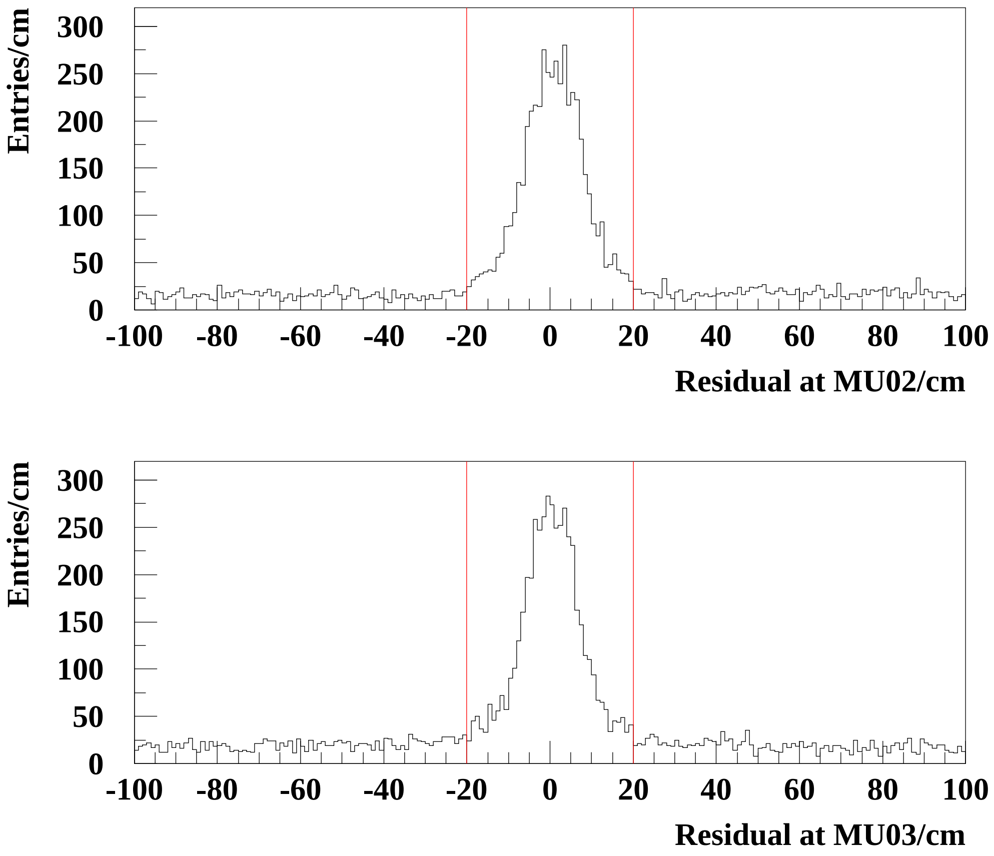

FIGURE 4.3: a) Residual between MUON pretrigger coincidences and hits in MU02. b) Residuals between MUON pretrigger coincidences and the line passing through the hits in MU01 and MU02.

4.2.3 FLT seeding

Using FLT tracks to setup the RoIs is more natural and less complex for the SLT compared to the previous two methods. As described in Section 3.3, the FLT follows the tracks through the chambers TC02 to PC01 (skipping PC02 and PC03) and provides both a good position and momentum measurement of the triggered lepton candidates. Most of the data triggered in the FLT+SLT configuration during the 2000 run are triggered in the single lepton mode, where the FLT requires the presence of a single, high-p_T muon or electron. In this mode the SLT simply uses the tracks found by the FLT as seeds.

When the FLT triggers on leptons pairs, the FLT sends to the SLT both the track information and the pairs of tracks with an invariant mass higher than the 2.0 GeV/c^2 threshold. In the 2000 run the FLT was however hampered by having many clones of the same track and we therefore let the SLT do its own pairing of FLT tracks. The SLT combines two muons or electron tracks with opposite charge into trigger pairs if the distance between the tracks at TC02 is more than 50 cm and if the invariant mass, recalculated more precisely by the SLT, is larger than 2.0 GeV/c^2.

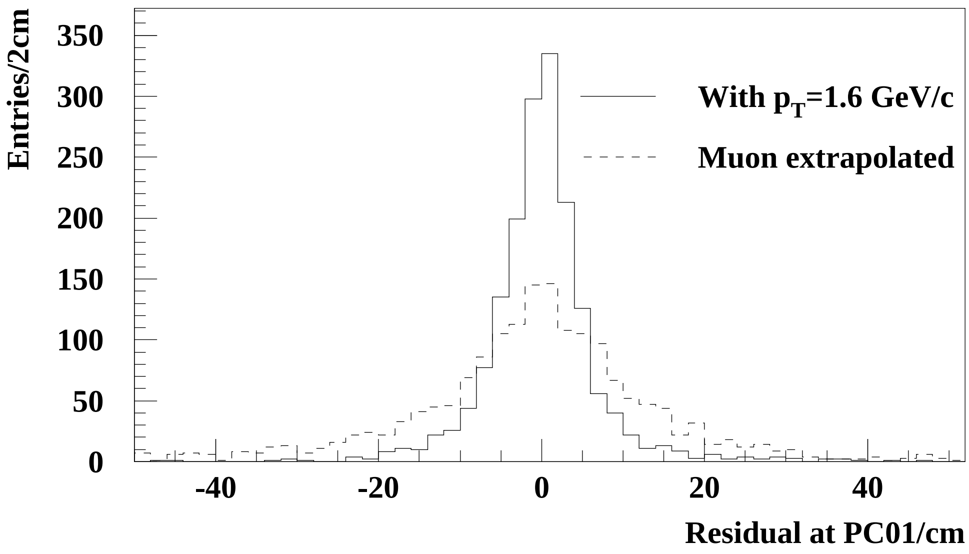

FIGURE 4.4: Residual in x at PC01 between Monte Carlo truth and RoIs calculated from either a fixed $p_T = 1.6\,\mathrm{GeV}/c$ assumption or from the muon track extrapolation. The fixed p_T assumption is observed to give the best prediction.

4.3 Main tracker algorithm

The main tracking system in HERA – B is located between the magnet and the calorimeter. It consists of six superlayers composed of drift chambers (OTR) and micro-strip gas chambers (ITR) as described in Chapter 2. Most muons and electrons pass all six superlayers before they go into the calorimeter/Muon system. Each superlayer consists of several individual layers, with the tubes/strips oriented either vertical or at $\pm 5^\circ$ angles to the vertical. During the 2000 run only the OTR chambers were used and the SLT did not use the drift time information. The lack of drift time reduces the precision, but avoids having to resolve the sign ambiguity of the drift time.

The SLT algorithm for the main tracker is divided into two parts: Slicer, which is designed to remove ghosts, and Refit, which assigns hits to the track and calculates the track parameters based on these hits. Each part is based on RoIs and all RoIs are processed before the SLT continues to the next algorithm.

4.3.1 Slicer

The purpose of Slicer is a fast ghost track rejection by checking that an RoI contains a minimum amount of hits located on a straight line. In order to reduce the latency and to limit the amount of data to be read out, only hits in the 0° layers are used.

The seeding algorithms define the RoIs, which are used to determine from which detector regions, the SLT should request data. These RoIs are however enlarged by requiring that they have a minimum width and height of 5 cm which is mostly the case for the MUON pretrigger

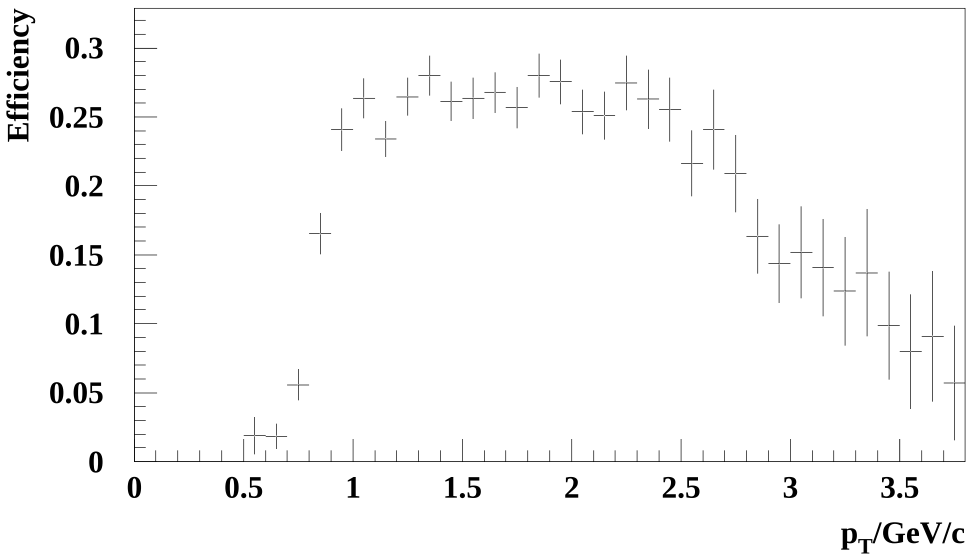

FIGURE 4.5: The trigger efficiency for muons from J/ψ decays as a function of the muon p_T.

seeding, see Section 4.2.2. An upper limit of 50 cm is also imposed, as larger regions are bound to contain more than one track. The enlargement is done in order to cope with non-perfect alignment and because the RoIs from the ECAL seeding are given with only 1σ boundaries[2].

After the hits have been read out, Slicer first checks that each superlayer contains at least one hit. In the 2000 run, one superlayer was allowed to be completely empty due to the unstable performance of the OTR[3]. The RoI is divided into "slices" by dividing each end of the RoI into 8 equally sized divisions and connecting all divisions from one end with all at the other end, as illustrated in Figure 4.6. For each of the 64 slices, the algorithm calculates how many hits are inside the slice and how many superlayers have no hits in the slice. A hit is defined to be inside a slice if it is within 1.5σ from the slice, where σ is the hit resolution, fixed to 1.5 mm for 5 mm OTR channels and 3 mm for 10 mm OTR channels. The 2000 algorithm required hits in a slice in at least 5 out of the possible 6 superlayers and a minimum number of hits in total in the slice. Unfortunately the 2000 run was triggered with a bug in the Slicer logic, which meant that it only considered slices with at least two empty superlayers. Consequently, the requirement of total hits in a slice was set to just nine hits as this was observed to have high trigger efficiency. Since the cut was so low, the bug did not affect the efficiency significantly, but it also meant that Slicer had low rejection power. This was not a major problem for the 2000 run, since there was enough CPU power available for the following SLT algorithms.

[2]The 1σ boundaries are necessary to be able to run without the main tracker algorithms.

[3]Allowing for an empty superlayer also helped increasing the 2000 acceptance, mainly in the case of tracks crossing MS10 instead of PC01.

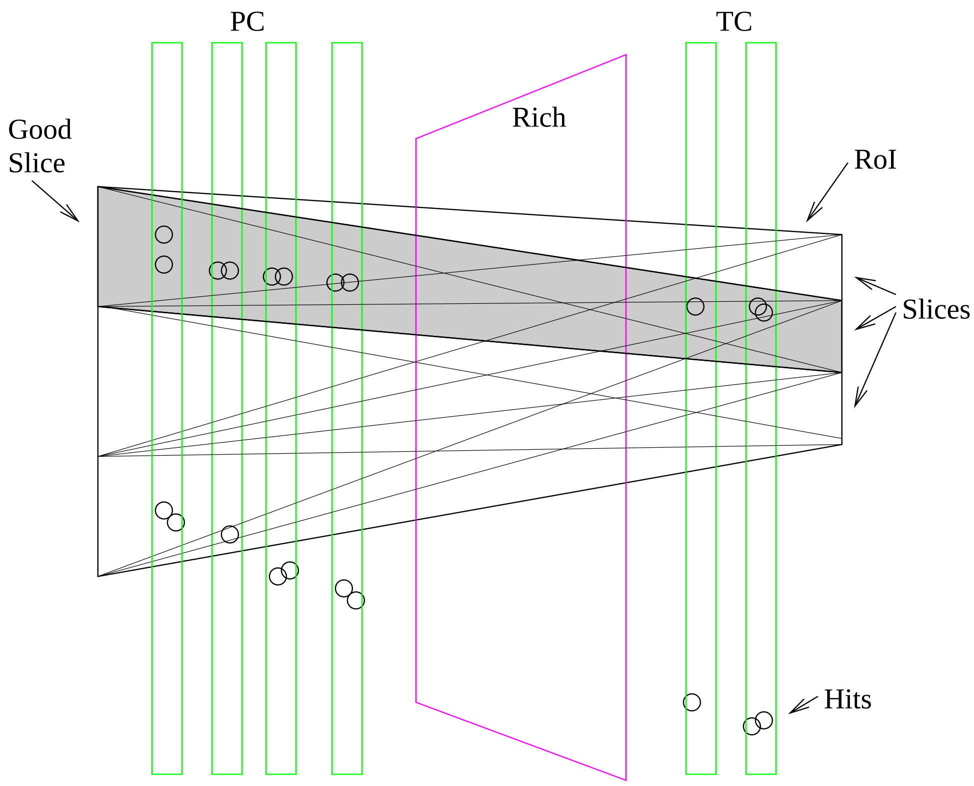

FIGURE 4.6: The subdivision of an RoI into slices by the Slicer algorithm. (The detectors are not to scale).

4.3.2 Refit

The Refit algorithm confirms and refines the charged track candidates in the Main tracker. It calculates the track parameters for a particle in the RoI and determines whether there is such a particle or not. At design, Refit was meant to work on small RoIs based on FLT tracks and not for general track finding. Hence the large RoIs from the MUON pretrigger seeding resulted in an efficiency drop as we shall see in Section 5.4.1. To simplify calculations and thus save CPU time, the tracking is done first in the horizontal plane using the hits already read out for Slicer. This step is called RefitX. Afterward hits in the stereo layers ($\pm 5°$) are read out and used to calculate the track parameters in y (RefitY).

The tracking is done with a simplified Kalman filter, see Appendix A. The main purpose of Refit is not to give the best possible track parameters, in particular covariance matrix for the track parameters. That is done by the 4LT/offline reconstruction. Therefore a significant saving in CPU time can be achieved by completely ignoring the multiple scattering in the detector. This reduces the parameter calculation to a least-squares fit of a straight line. The track parameters in a single coordinate (x or y) are parameterized by the track slope, t, and the intercept, b, at $z = 0$ (ignoring the magnetic field). The equations from Section A.3 for calculating the track parameters, $\vec{x}_k = (b, t)$, and covariance matrix $\mathbf{C}_k$, after tracking on k hits with positions m_k, are reduced to:

$$\mathbf{C}_k^{-1} = \mathbf{C}_{k-1}^{-1} + \frac{1}{\sigma_k^2}\begin{pmatrix} 1 & z_k \\ z_k & z_k^2 \end{pmatrix} = \sum_{i=1}^{k} \frac{1}{\sigma_i^2}\begin{pmatrix} 1 & z_i \\ z_i & z_i^2 \end{pmatrix} \tag{4.1}$$

$$= \begin{pmatrix} \sum_{i=1}^k \frac{1}{\sigma_i^2} & \sum_{i=1}^k \frac{z_i}{\sigma_i^2} \\ \sum_{i=1}^k \frac{z_i}{\sigma_i^2} & \sum_{i=1}^k \frac{z_i^2}{\sigma_i^2} \end{pmatrix}, \tag{4.2}$$

$$\vec{x}_k = \mathbf{C}_k \left(\mathbf{C}_{k-1}^{-1}\vec{x}_{k-1} + \begin{pmatrix} 1 \\ z_k \end{pmatrix} \frac{m_k}{\sigma_k^2} \right) = \mathbf{C}_k \left(\sum_{i=1}^{k} \begin{pmatrix} 1 \\ z_i \end{pmatrix} \frac{m_i}{\sigma_i^2} \right) \tag{4.3}$$

$$= \mathbf{C}_k \begin{pmatrix} \sum_{i=1}^k \frac{m_i}{\sigma_i^2} \\ \sum_{i=1}^k \frac{m_i z_i}{\sigma_i^2} \end{pmatrix}. \tag{4.4}$$

Because the multiple scattering is ignored it does not matter in which order the hits are added, unlike the situation of a full Kalman filter where a smoothing step has to be carried out to obtain track parameters at both ends of the detector. Since PC01 has a lower hit occupancy compared to TC02 the tracking is done from PC01 to TC02. This is not optimal for RoIs based on muon seeding, as the RoIs at PC01 can be up to 10 times larger in that case, compensating for the lower hit occupancy.

The pattern recognition in Refit works with hits in a single superlayer at a time. The track candidates are found by combining hits, which according to a χ^2 test lie on a straight line with respect to upstream track candidates. After each superlayer only the five track candidates of best quality are continued to the next superlayer. The track quality is determined by the χ^2 and the number of layers where a track candidate does not have hits. The initial track candidate is calculated from the RoI center and width.

More precisely the algorithm proceeds as follows:

- Track the candidates through a superlayer:
  - If there are no hits at all in the RoI in a given superlayer skip to the next superlayer.
  - Loop over layers within the superlayer:
    - Loop over track candidates:
      - For every hit in a layer calculate the χ^2 for adding the hit to the track candidate.
      - If $\chi^2 < 5 \cdot \#$"hits on track", a new track candidate including the hit is created.
      - If at least one new track candidate is created, the original candidate is removed.
      - A track candidate is also removed if this is the third consecutive empty layer, otherwise it continues to the next layer.
- Cleanup the list of track candidates:
  - Sort tracks according to their track quality, Q, given by Equation 4.6.
  - Remove all but the five tracks with the best Q.
  - If tracks are clones of each other, remove those with worst Q.

   – Tracks are defined as clones if their difference in slope is less than 0.3 mrad and their difference in intercept is less than 1 mm.

- Go to next superlayer if any candidates are left.

The χ^2 is calculated as:

$$\chi^2 = \sum_i \frac{(m_i - x_i)^2}{\sigma^2_{m_i} + \sigma^2_{x_i^{\mathrm{SL}}}}, \tag{4.5}$$

where x_i is the track position at the z of hit number i. The position is calculated from the first i hits, so that adding an extra hit to the track only requires the calculation of an extra term, not the full χ^2. The track position uncertainty $\sigma_{x_i^{\mathrm{SL}}}$ is only calculated once per superlayer (before any hits are added from the superlayer) in order to save CPU time. The track quality, Q, is defined as:

$$Q = 0.05 \cdot \chi^2 + \#\text{"empty layers"} \ , \tag{4.6}$$

and therefore weights both how “straight” a track is and how “complete” it is.

The above algorithm is given for RefitX. The algorithm for RefitY is identical with a few exceptions. The hits are measured in stereo layers with wires rotated $\pm 5°$ from vertical. A hit coordinate, $\tilde{m}$, is therefore related to the (x, y) position through the relation:

$$\tilde{m} = x\cos\theta - y\sin\theta, \qquad \theta = \pm 5° \tag{4.7}$$

These hits contain information on both the x and y position of the track. From RefitX we already have a reasonably good estimation of the x position and to save computing time we therefore transform the measurements points into a y coordinate using the track from RefitX:

$$m = y = \frac{x\cos\theta - \tilde{m}}{\sin\theta}, \quad \sigma_m = \frac{\sigma_{\tilde{m}}}{\sin\theta} \tag{4.8}$$

These measurements are used to calculate track parameters and covariance matrices for the y coordinates. We first check that the hit transformed according to Equation 4.8 lies inside the RoI in y. This removes approximately 85% of the hits that were read out. As the hit uncertainty is 12 times larger than in x, the clone removal cuts have been increased to 2.5 mrad in slope and 5 mm in intercept.

The OTR has 14 0° layers and 16 stereo layers, of which 4, respectively 8, are double layers. In Refit a double layer is treated as a single layer. If a particle produced hits in both layers, these hits will give rise to different candidates and will not be used in the same track. This gives a maximum of 14(16) hits that can be used to determine $x(y)$ coordinates. Figure 4.7 shows the number of hits used by Refit in data and simulation.

Only the track candidate with the best quality is passed on to L2Magnet.

4.4 Magnet propagation

L2Magnet propagates the Refit tracks through the magnet and into the VDS where they are used to set up the RoI for L2Sili. An early version of L2Magnet has been described in detail in Ref. [51]. The code was upgraded before the 2000 run to provide higher precision.

The magnetic field provided by the HERA – B magnet can be considered constant and oriented along the y-axis. In this approximation, the bending in the magnet is proportional to $\frac{Q}{p_{xz}}$,

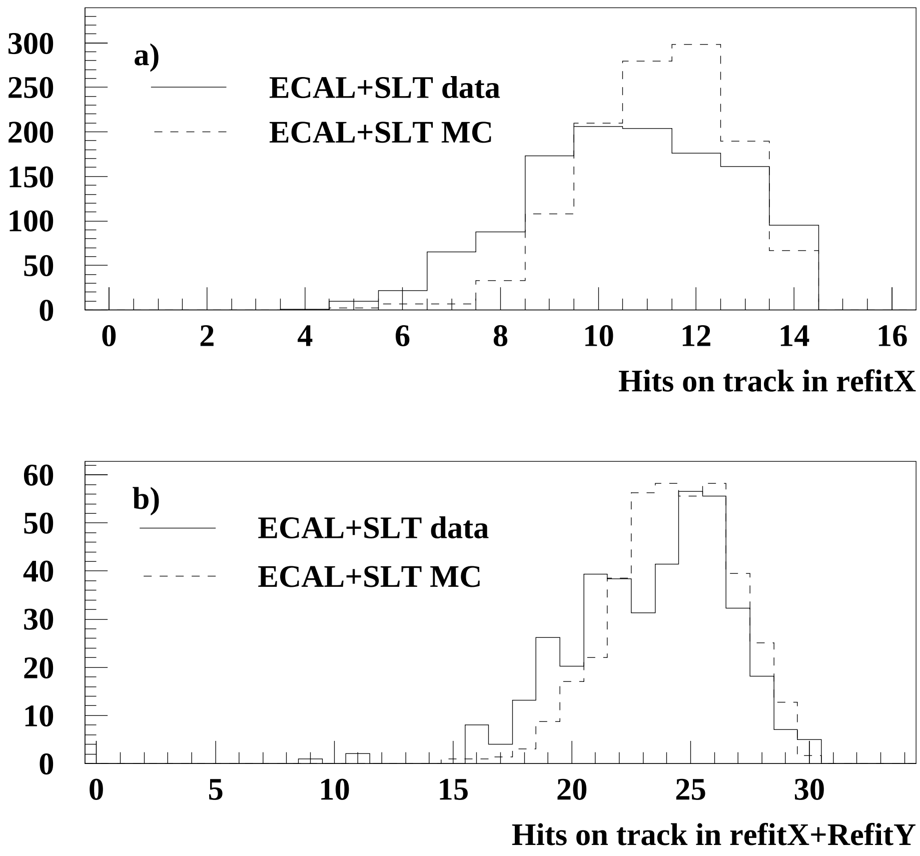

FIGURE 4.7: Number of hits on Refit tracks in data and $J/\psi \to \mu^+\mu^-$ Monte Carlo. a) Distribution after RefitX. b) Distribution after RefitY (with hits from RefitX counted also).

where Q is the particle charge and p_{xz} is the momentum in the $x-z$ plane. The approximation does not give precise enough track parameters and we therefore parameterize the bending as function of both the momentum and the entrance point of a particle in the magnet. The particles of interest come from the target region, and we can therefore use the track slopes before the magnet, $t_{x,b}$ and $t_{y,b}$, as parameters for the position dependence.

Figure 4.8 illustrates the algorithm used. The tracks are assumed to be straight lines up to the center of the magnet, where the positions are displaced by $(\Delta x, \Delta y)$ and the slopes by $(\Delta t_x, \Delta t_y)$. These changes are parameterized as:

$$\Delta t_x = \frac{Q}{p} F_{t_x}(t_{x,b}, t_{y,b}), \tag{4.9}$$

$$\Delta x = \frac{Q}{p} F_x(t_{x,b}, t_{y,b}), \tag{4.10}$$

$$\Delta t_y = \frac{Q}{p} F_{t_y}(t_{x,b}, t_{y,b}), \tag{4.11}$$

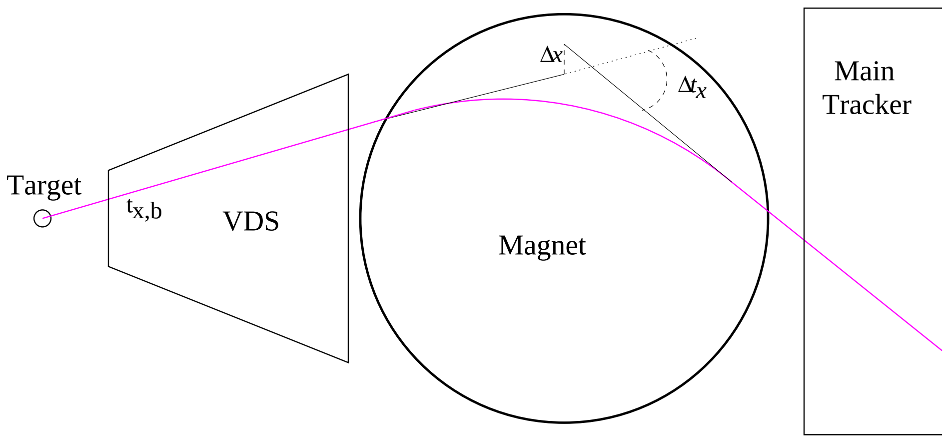

FIGURE 4.8: The approximation of the bending in x used by the magnet propagation.

$$\Delta y \quad = \quad \frac{Q}{p} F_y(t_{x,b}, t_{y,b}). \tag{4.12}$$

The four fieldmaps, F_{t_x}, F_x, F_{t_y} and F_y, are "measured" on a grid of 52×32 nodes within ± 260 mrad and ± 160 mrad in $t_{x,b}$ and $t_{y,b}$, respectively. The "measurement" is done using the full HERA – B detector simulation, which has a detailed map of the magnetic field. Monte Carlo muons of 100 GeV/c, one for each node in the grid, are tracked through the magnetic field and the deflections are measured. For the purpose of this measurement all multiple scattering effects were turned off in the simulation. The resulting fieldmaps are shown in Figure 4.9 and, with the exception of regions around the electron beam pipe and on the magnet edges, they are almost flat, and for the y coordinates close to zero. For a particular value of track slopes, the fieldmaps are calculated as a weighted average of the four closest nodes in the grid.

The target steering database sends the positions of the target wires every 10 seconds to all of the SLT nodes. The interaction region along the wire is set to 3 mm[4]. From this information a "target-box" is constructed, which is the smallest three dimensional box containing all the interaction regions of the inserted target wires. The center of this target-box is used by the magnet propagation to constrain the track parameters before the magnet. To allow for particles not coming from the primary interaction, in particular leptons from $b \to J/\psi X$ decays, the target-box is increased by an extra 2 mm in all directions.

The track slopes before the magnet and the momentum are estimated using an iterative procedure based on the center of the target box, the refit track parameters and the fieldmaps. The refit track is prolonged as a straight line to the center of magnet and this point is used together with the target box to calculate the first estimate of $(t_{x,b}, t_{y,b})$. From these, the fieldmaps can be looked up, followed by an estimation of $\frac{Q}{p}$ using Equation 4.9. The point in the magnet is corrected using Equation 4.10 and 4.12 and a new estimate of $(t_{x,b}, t_{y,b})$ is calculated. These are used to look up the fieldmaps again and the procedure converges for all three values, $t_{x,b}, t_{y,b}$ and $\frac{Q}{p}$, after a few iterations. Once the track parameters in front of the magnet have been calculated, the full covariance matrix for these is calculated based on the covariance matrix from the Refit track and the width of the target-box (divided by $\sqrt{12}$ to make it approximately

[4]The 3 mm correspond to approximately ± 3 times the Gaussian width of the proton beam.

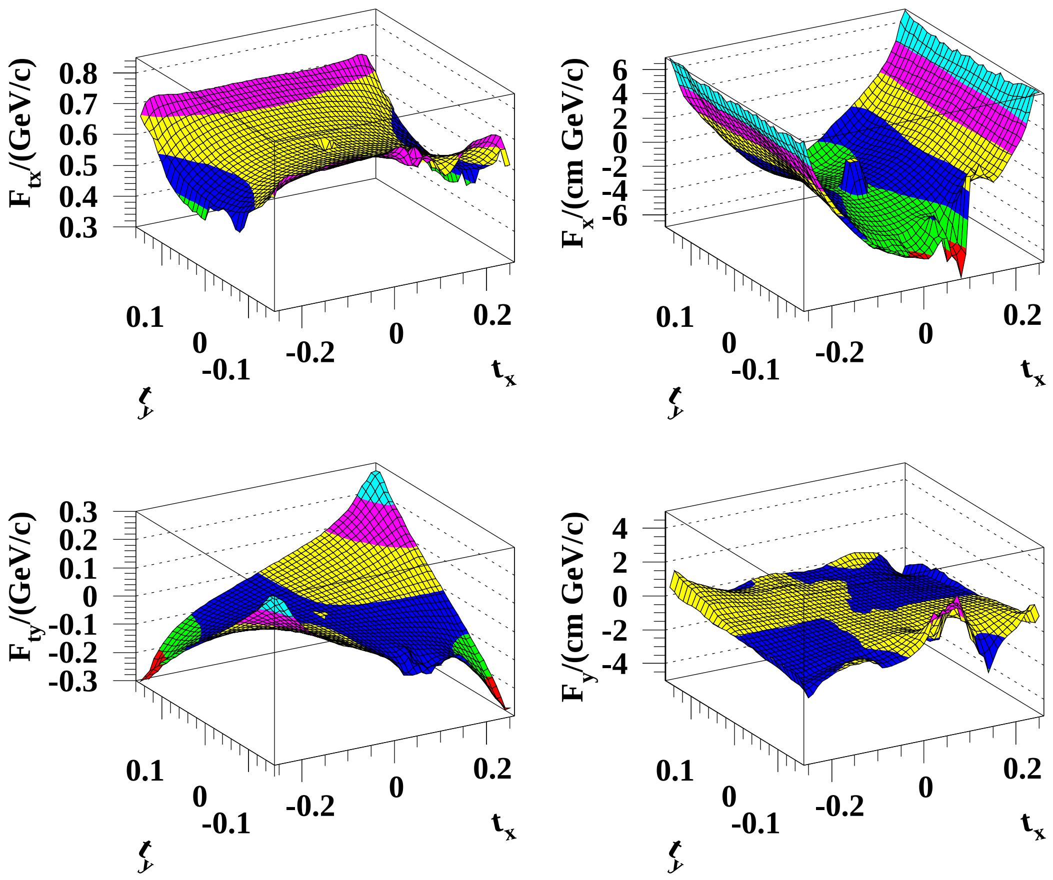

FIGURE 4.9: Fieldmaps used to propagate tracks through the magnet. The bumps at $(t_x, t_y) = (0.1, -0.1)$ are caused by the electron beam pipe.

Gaussian). The covariance matrix for the Refit track is corrected for multiple scattering effects in the magnet by adding a term proportional to $\frac{1}{p^2}$ where the proportionality factor has been determined from Monte Carlo.

L2Magnet was originally designed to use the first calculation of track parameters in front of the magnet for predicting an RoI in chambers right in front of the Magnet (SI08, MS01 and MC01). Using hits in these chambers, a space point is constructed and a second calculation of the track parameters is carried out. This reduces the dependence on the target constraint and gives smaller RoIs for the tracking in the VDS by L2Sili. However, in the 2000 run, the ITR detector (MS01) was not available and SI08 was heavily damaged. Therefore the hit finding capability of L2Magnet was not used and will be integrated only in the upgrade for 2002. The result of the first track parameter calculation was used directly to setup an RoI for L2Sili in the first seven VDS superlayers.

4.5 Vertex detector tracking

The tracking through the VDS detector in the SLT is performed by L2Sili [52]. The VDS is described in detail in Section 2.4. L2Sili only uses the seven superlayers inside the vacuum vessel as the last superlayer, SI08, is meant to be used by L2Magnet.

L2Sili starts the tracking at SI07 and follows the tracks toward the target. For speed reasons, the tracking is done independently in x and y coordinates. Afterward, the tracks are supposed to be fitted to a common vertex, though this feature was not used for the 2000 run. For vertexing, a precise estimate of not only the track parameters, but also of their covariance matrix is needed. To obtain this we need to take multiple scattering into consideration as this can be significant compared to the nominal hit resolution of 12 μm.

In Ref. [52] the calculation of track parameters is formulated as a least-squares fit of a straight line. Here it is treated in the frame work of a Kalman filter, which eases the comparison with the Refit algorithm and the 2002 version of L2Sili. In 2000, L2Sili was not a true Kalman filter, as the multiple scattering effects in L2Sili were treated in a different, non-optimal way. Instead of adding multiple scattering terms to the covariance matrix after each layer, as described in Appendix A, the cumulative amount of multiple scattering is calculated for each layer passed by the particle. This can be pre-calculated using the seed from L2Magnet for the momentum and to determine which layers are passed, and an internal database with the amount of material in each layer. The amount of multiple scattering increases as a particle moves downstream and L2Sili increases the measurement uncertainties in the different layers accordingly, based on the pre-calculated values. Such a multiple scattering correction gives approximately the correct weighting of the hits used in the calculation of track parameters, as the hits closest to the target have a smaller uncertainty and thus higher weight. This method has the advantage that it avoids a lot of CPU intensive covariance matrices inversion. However the hit uncertainties are increased, thus L2Sili is more likely to assign hits from different tracks to a single track.

The algorithm is further simplified by treating the x and y coordinates almost independently — almost, because the $\pm 2.5^\circ$ stereo angles introduce some correlations. In this case one works with two 2×2 covariance matrices instead of one 4×4 matrix in the calculation. The track coordinates are parameterized by track slopes (t_x, t_y) and intercepts, (b_x, b_y) at $z = 0$ as for Refit. A hit coordinate, $\tilde{m}$, is related to the x, y coordinates through:

$$\tilde{m} = x \cos\theta - y \sin\theta, \quad \theta \in \{2.5^\circ, -2.5^\circ, 87.5^\circ, 92.5^\circ\}. \tag{4.13}$$

For hits in the x view ($\theta = \pm 2.5^\circ$) the measurement, $\tilde{m}$, is modified to remove the y coordinate dependence to:

$$m = x \cos\theta = y \sin\theta + \tilde{m}. \tag{4.14}$$

The measurement uncertainty becomes:

$$\sigma_m^2 = \sigma_y^2 \sin^2\theta + \sigma_{\tilde{m}}^2 + \sigma_{\mathrm{MS}}^2, \tag{4.15}$$

where σ_y is the uncertainty in the track parameters in y at the z of the hit, $\sigma_{\tilde{m}}$ is the hit resolution and σ_{MS} is the pre-calculated increase due to multiple scattering. σ_y is calculated based on the hits already added to the track before the current measurement, while the update of the y is explained below. The hit resolution was increased to 30 μm in the 2000 run, as the alignment available online was not giving the nominal precision. This is however a minor correction, as the multiple scattering part, σ_{MS}, is typically 50–150 μm. Since there is no multiple scattering to be added to the covariance matrix separately, the covariance matrix and track parameters calculation (similar to the case of Refit) reduces to:

$$\mathbf{C}_k^{-1} = \sum_{i=1}^{k} \frac{1}{\sigma_i^2} \begin{pmatrix} \cos^2\theta_i & z_i\cos^2\theta_i \\ z_i\cos\theta_i & z_i^2\cos^2\theta_i \end{pmatrix}, \tag{4.16}$$

$$\vec{x}_k = \mathbf{C}_k \left(\sum_{i=1}^{k} \begin{pmatrix} \cos\theta_i \\ z_i\cos\theta_i \end{pmatrix} \frac{m_i}{\sigma_i^2} \right). \tag{4.17}$$

By substituting Equation 4.14 into Equation 4.17, the correction from y coordinates can be improved:

$$\vec{x}_k = \mathbf{C}_k \left(\sum_{i=1}^{k} \begin{pmatrix} \cos\theta_i \\ z_i\cos\theta_i \end{pmatrix} \frac{\tilde{m}_i + y\sin\theta_i}{\sigma_i^2} \right) \tag{4.18}$$

$$= \mathbf{C}_k \left(\sum_{i=1}^{k} \begin{pmatrix} \cos\theta_i \\ z_i\cos\theta_i \end{pmatrix} \left(\frac{\tilde{m}_i}{\sigma_i^2} + b_y \frac{\sin\theta_i}{\sigma_i^2} + t_y \frac{z_i\sin\theta_i}{\sigma_i^2} \right) \right). \tag{4.19}$$

Using the last equation, the correction from the y coordinates can be calculated from the best available information instead of using the instantaneous value of t_y and b_y for each hit, which could be poor, particularly at the first layers. The above equations also hold for (b_y, t_y) with y-view hits after a substitution of $(\cos\theta, \sin\theta) \to (-\sin\theta, -\cos\theta)$.

The RoI used by L2Sili is based on the track predicted by L2Magnet with the covariance matrix corrected for multiple scattering in the VDS. The center of the RoIs follow the L2Magnet track and the width/height is set to $\pm 3\sigma$, with σ given by the covariance matrix.

The pattern recognition in L2Sili has some similarities with the Refit philosophy. For each superlayer it combines track candidates using all possible hit combinations which have a reasonable χ^2. After each superlayer, candidates with few hits are removed, but, unlike Refit, there are no limitations on the number of empty layers. The algorithm starts at the superlayer closest to the magnet and, unlike Refit, it looks at x and y view simultaneously, updating the x and y alternatively according to the view of a hit.

The initial track candidate is given by the seed from L2Magnet and L2Sili loops over the seven superlayers according to following algorithm:

- Track the candidates through the superlayer:
  - Loop over layers in the superlayer:
    - Loop over the track candidates:
      - For every hit in the layer, check if the hit position is less than 5σ away from the predicted track position.
      - If this is the case, calculate the χ^2 (Equation 4.20) for adding the hit to the track candidate.
      - If the probability of the χ^2 is higher than 0.001, a new track candidate containing the hit is created.
      - Both the original and the possible new track candidates are kept for the tracking in the next layer.

- Cleanup the list of track candidates:
  - Sort the track candidates according to the total number of hits they contain.
  - Candidates which have at least two hits less in both the x and y views compared to the best candidate are removed.
  - Candidates whose hits are all contained in a different candidate are removed.
- Continue to the next superlayer.

The χ^2 is calculated as:

$$\chi^2 = \sum_i \frac{|\tilde{m}_i - ((b_x + z_i t_x)\cos\theta_i - (b_y + z_i t_y)\sin\theta_i)|^2}{\sigma_m^2}, \tag{4.20}$$

where the sum is made over x or y hits of the tracks. Unlike Refit where the χ^2 is calculated cumulatively, L2Sili recalculates the complete sum for every hit added, which can be computationally expensive. The check of whether hits of a candidate are a subset of another candidate can also be computationally expensive, as a double loop over all candidates to compare all of their hits is necessary. Finally, if the number of track candidates exceeds 1000 at any point, all but the candidate with most hits are removed. This typically occurs in about 3% of the RoIs.

At the end, a track candidate is accepted by L2Sili, if it has at least three hits in each view and if the hits come from at least two different superlayers. This is the minimum required for reliable track parameters to be calculated. In the design SLT configuration, L2Sili is followed by the vertexing in L2Vertex, but this was not used in the 2000 run as the suppression by the SLT was already large enough.

Chapter 5

SLT performance in the 2000 run

The performance of a trigger system is ultimately measured by the number of usable signal events written to tape per hour. This tells us how much statistics we will have for our physics analysis. The trigger algorithm should therefore be designed to maximize this number. Once data have been triggered and written to tape, the rejected events cannot be recovered by any means. It is therefore important to maximize the trigger performance as much as possible beforehand. For the 2000 run this was only partially done, mainly because the run period was limited to a few months, leaving very little time for optimizations. In this chapter we will evaluate the performance of the SLT for the 2000 run. A major goal of this evaluation is to locate the weak points in the SLT algorithm, which can be corrected before the start of the 2002 run.

We define signal events to be events in which a J/ψ decays into two muons or electrons. We use prompt produced J/ψ as signal, because the amount of $b \to J/\psi X$ events in the 2000 run was tiny, see Chapter 8. Production cross sections in proton-nucleus collisions is generally parameterized as:

$$\sigma = \sigma_0 A^{\alpha}, \tag{5.1}$$

where A is the number of nucleons in the target nucleus. This formula is used to calculate the cross section for both inelastic interactions ($\sigma_0 = 41.2 \pm 0.29$ mb, $\alpha = 0.7110 \pm 0.0016$ [53]) and prompt J/ψ ($\sigma_0 = 357 \pm 8 \pm 27$ nb, $\alpha = 0.955 \pm 0.005$ — see Chapter 8). With $\mathrm{Br}(J/\psi \to \ell^+\ell^-) = 5.93\%$, the signal to background is $0.94 \cdot 10^{-6}$ for a carbon target. For a mean interaction rate of 5 MHz, this gives a production rate of five $J/\psi \to \ell^+\ell^-$ per second or 17,000 per hour. Our estimation has a 10% uncertainty from the J/ψ cross section and a 10% uncertainty from the luminosity. A signal event is useful only if we are able to separate it from background and to reconstruct it. We will assume that if both leptons are found and reconstructed by the SLT, they can be reconstructed by the 4LT/offline reconstruction. We will ignore events where we, for instance, happen to trigger on just one of the leptons and the other triggered track comes from a pion decay.

The performance evaluation uses a combination of real data and simulated events. For real events we will mainly use SLT bypass events[1], as the SLT triggered events are a completely biased sample. The evaluation is divided into five points:

- **Efficiency:** The probability that an event passes the trigger. Ignoring deadtime in the DAQ system, the efficiency is equivalent to the measured number of signal events per

[1] The SLT allows a small fraction (1/400) of events, which passed the seeding step, to be accepted even in they are later rejected by the tracking algorithms. These events are called "SLT bypass" events.

Run no.	Run date	Runtime	No. of triggered events	Trigger type
15230	May 14	537 min	382,000	ECAL+VDS
16665	July 23	342 min	48,500	ECAL/MUON+SLT

TABLE 5.1: The runs which are the basis for the SLT performance evaluation.

hour. We will discriminate between inefficiencies caused by detector acceptance, detector performance and the trigger algorithm itself.

- **Tracking precision:** The tracks found by a given algorithm step are used to define the RoI for the following step. High precision is therefore important in order to ensure high efficiency and good background rejection.
- **Rejection:** In order to minimize the deadtime, and thus have maximal benefit of our trigger efficiency, we need to achieve sufficient rejection of background events at each step in the algorithm, in order to be able to handle the rate.
- **Timing:** Closely connected to rejection is the time each algorithm step takes to process an event. The faster the algorithm, the higher event rate it can handle.
- **Fake track rates:** The signal events are only a tiny part of the triggered events. Most of the triggered events contain either tracks which have similar properties to the signal, like a muon from pion decay, or fake tracks. Fake trigger tracks either consist of noise hits or are combinations of hits from different particles. Ideally these should not occur or at least must be rejected as much and as early as possible.

5.1 Detector performance

The SLT performance depends critically on the performance of the detector components used by the SLT. This can vary quite a lot from run to run and in particular from run period to run period. For the evaluation presented in this chapter we have chosen a two reference runs for the different trigger configuration used in the 2000 run. The selected runs are listed in Table 5.1. These runs are long and stable enough for a reasonable reliable detector performance measurement. The measured performance is used as input to the detector simulation for the masking of dead channels and to set realistic hit efficiencies. To understand the influence of the detector performance, we will compare the SLT performance with a realistic detector simulation with the performance under the ideal detector simulation.

The following describes the simulation of each detector system used by the SLT.

VDS: The performance of the VDS in the 2000 run was satisfactory. Inefficiencies were determined for each run and were normally limited to individual readout chips. The detector simulation however allows us only to set an average efficiency in a whole layer (8–10 readout chips). For good detector layers the hit efficiency was about 98%. Hit resolution was set to 12 μm, though alignment with this precision was only achieved after the data were triggered.

OTR: The large amount of dead and hot channels in the OTR (up to 15% and 6% respectively) were determined for each run. Both kinds of defective channels are masked out in the simulation. The hit efficiencies in working channels have been evaluated [54] to be 85%–95%, depending on the superlayer. For the simulation all layers are set to have the same hit efficiency, 90%, for non-masked channels. For an ideal detector we use 95%, corresponding to the best observed. The hit resolution is set to the measured one of 500 μm. This has no significance for the SLT as we do not use the drift time information, but only the cell hit, the resolution of which is given by the cell size.

ECAL: The calorimeter readout and pretrigger electronics were under constant development during the 2000 run. The simulation masks out both the readout and pretrigger boards which were not present together with individual dead or hot channels.

MUON: The muon readout allows the signal thresholds to be set on a channel by channel basis and many noisy channels were suppressed by increasing their threshold. The efficiency, as function of the threshold, is not very well known and all channels with high thresholds or which are observed to be dead have been masked out completely in the simulation. For working channels a hit efficiency of 98% is used. This is not correct for the muon pads due to the preamplifier mismatch (See Section 2.10), and the estimates for the pad efficiencies vary between 80% and 90%. However since a muon pretrigger coincidence requires two pads to be hit and we always require two muon pretrigger coincidences when triggering on two muons, we simply scale the trigger efficiencies with $(85\%)^4 = 52\%$.

SLT: The SLT offline analysis of both real and simulated data runs almost the same algorithm, as used online. The only differences are that the raw data are read from a file instead of being requested from the SLB and that the VDS data have been sparsified[2]. The simulated data are stored in the same format as real data and the SLT does not discriminate between the two types of input data.

5.2 ECAL+L2Sili trigger

The ECAL+L2Sili trigger, which was described in Section 4.1, is the simplest SLT trigger configuration used in the 2000 run. It was used as a stepping stone to the full SLT algorithm, while the OTR was not yet ready. It is therefore not planned to be used in the 2002 run. All data taking was done at 5 MHz interaction rate and using mainly the carbon target wire (Inner 2).

5.2.1 Trigger efficiency

For the efficiency evaluation we use a sample of 40,000 simulated $J/\psi \to e^+e^-$ events without any underlying event[3] and 20,000 simulated prompt $J/\psi \to e^+e^-$ events generated as described

[2] The sparsification of VDS data removes empty channels, does noise suppression and gathers neighboring hits into single clusters. This is done during event building and the raw VDS data is not stored.

[3] Underlying event refers to the other particles produced in the same proton–nucleus collision that produced the J/ψ. An event can also contain tracks from additional proton–nucleus interactions. These additional interactions are referred to as "overlapping events" and is simulated by adding the hits from inelastic Monte Carlo events to the J/ψ Monte Carlo events.

in Section 8.3. The 20,000 events have a variable amount of inelastic interactions added to estimate the effects of detector occupancies. The detector has been simulated according to both its ideal performance and to the performance in the run 15230.

To separate algorithm efficiency from effects caused by detector acceptance/efficiency we separate the total trigger efficiency into the following factors:

Geometrical acceptance: The probability that both leptons actually pass through the subdetectors used to trigger the J/ψ. For this trigger configuration the requirements are that the electrons pass through the VDS and deposit their energy in the calorimeter. The energy deposited has to be higher than the threshold $E_T > 1.2$ GeV, a kinematic constraint included in the geometrical acceptance.

The simulation stores every detector plane crossed by a particle (impact point) and it is therefore simple to check that each electron passes through at least two superlayers in the VDS, which is the minimum number of superlayers required to reconstruct a track. All VDS layers in the simulation are represented by the full silicon layer including non-sensitive edges. In the non-sensitive areas no hits will be generated, but we still get an impact point in the simulation, hence the acceptance is slightly overestimated by our method.

Almost all of the electrons pass through more than one radiation length of material before they reach the calorimeter and thus radiate a significant amount of bremsstrahlung photons. As described in Section 4.2.1, part of these photons are included in the same cluster as the electrons, while others are lost. To evaluate the E_T of an electron in the simulation, we can therefore not use the electron energy at the beginning of the track nor the energy it has when it reaches the calorimeter, as we would either overestimate or underestimate the effective E_T. One could in principle try to sum up all of the energy of the photons emitted after the magnet, but some of these photons will convert to e^+e^-. Furthermore, it is not simple to calculate from where in the magnetic field the photons, that are ending up in the electron cluster, start. Instead we define an electron to be within the acceptance with a given E_T cut if for ideal detector performance we can reconstruct a cluster with more than the required E_T and the cluster position is consistent with the impact point of the electron. The cluster reconstruction is done with the CARE version used by the SLT. Part of the reconstruction efficiency is therefore included into our definition of geometrical acceptance. Figure 5.1 shows the acceptance for $J/\psi \to e^+e^-$ events as a function of the E_T cut.

Calorimeter efficiency: The geometrical acceptance is calculated for an ideally working calorimeter. The calorimeter efficiency gives the amount of events which have both electrons depositing their energy in the working part of the realistic simulated calorimeter. The measured inefficiency is mainly due to missing/broken readout electronics.

Electron pretrigger efficiency: The electron pretrigger requires that at least two clusters have an $E_T > 1.2$ GeV. This is already taken into account by the geometrical acceptance cut and any further inefficiency in the pretrigger is caused by non-working channels in the pretrigger and differences in the clustering procedure which is significantly simpler for the ECAL pretrigger compared to CARE. The main inefficiency comes from the non-working channels. This can be seen in the simulation with an ideal detector where only 1.3% of the events in the acceptance are not triggered by the pretrigger, while with a realistic simulated detector we loose half of the events in the ECAL pretrigger.

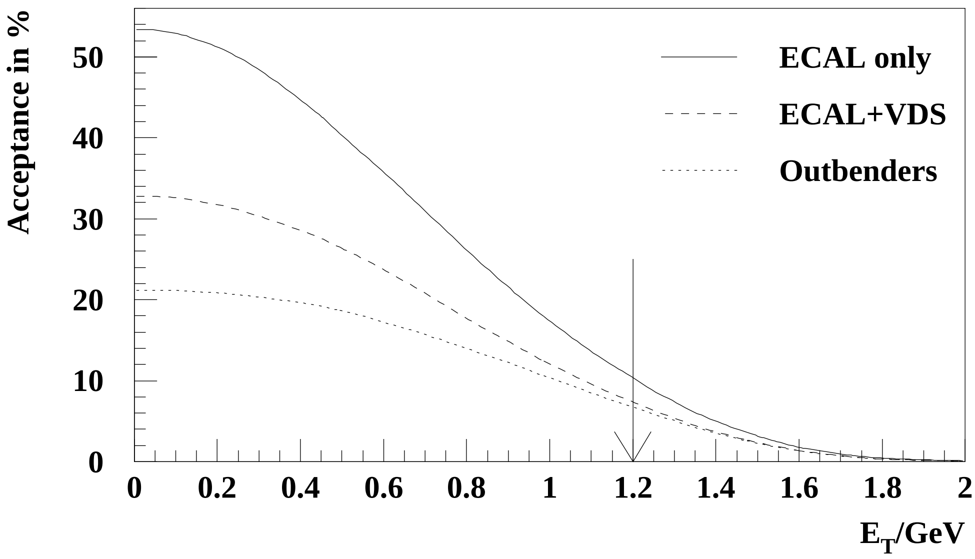

FIGURE 5.1: Geometrical acceptance of the ECAL+L2Sili (dashed line) and ECAL-only (full line) trigger as a function of the E_T cut applied to the two electrons. The dotted line shows the effect of requiring the two electrons to be out-bending. The arrow shows the cut used by the ECAL+L2Sili trigger.

Lepton charge assumption: A pair of ECAL clusters are assumed to come from an out-bending pair of e^+e^- particles. With the high E_T cut already applied, this charge assumption only removes 10% of the events as shown in Figure 5.1.

Invariant mass cut: The invariant mass is calculated using the energy of the two clusters and required to be larger than 2.0 GeV/c^2. The invariant mass spectrum is shown in Figure 4.1. This cut introduces some losses of $J/\psi \to e^+e^-$ events due to bremsstrahlung radiation either in or before the magnetic field, resulting in a reconstructed invariant mass of the J/ψ lower than the nominal 3.097 GeV/c^2.

L2Sili RoIs: The RoIs for L2Sili are predicted using the position of the interaction point and the cluster. These points are located more than 13 meters apart. We have to distinguish between the case of a single target wire and the case of multiple wires. If a single wire is used, the interaction point is known with a precision of a few hundred microns and thus gives a very good constraint for the RoI. For multiple target wires the interaction point is only known with a few millimeters precision. Figure 5.2 shows the deviation of the predicted track position from the true position for both cases. As can be seen, the positions are predicted to within a few millimeters of the truth, with the x position downstream being the worst due to the uncertainties of the bending in the magnetic field.

Figure 5.2 also shows the value of the deviations divided by the size of the RoIs. From this we see that most of the tracks are inside the RoI. Even if a track is not fully contained in an RoI, it can still have enough hits inside the RoI for the SLT to reconstruct it accurately. We therefore require that impact points of a track are contained inside the RoI on at least two VDS superlayers. This corresponds to the minimum requirement for a track to be in the geometrical acceptance.

VDS hits: L2Sili requires at least three hits in each view to accept a track and that these hits come from two different superlayers. The geometrical acceptance already requires that each electron passes at least two superlayers, which should ensure that the required hits are always present, if we have no inefficiencies in the VDS layers. We therefore ask that the lepton produce the required number of hits and that they are present inside the RoI, so that we can separate detector effects from the efficiency of the L2Sili algorithm. As can be seen from Table 5.2 even with an ideal detector we have a significant amount of events, which do not fulfill this requirement. It has been checked that more than 90% of this loss was due to tracks passing insensitive parts of the VDS and thus our definition of geometrical acceptance overestimate the real acceptance by about 5%.

L2Sili algorithm: The last step in the trigger is tracking in L2Sili. We have already required that the needed hits are present and any inefficiency will be due to the pattern recognition done by L2Sili. The main rejection in L2Sili is done by a cut on the χ^2 of a track. The χ^2 can become too large either due to multiple scattering or by combining hits from one track with hits belonging to other tracks. The first effect can be estimated from the sample without underlying events, and latter from adding overlapping inelastic events. Both effects are seen to give inefficiencies of less than a few percent.

Occupancy dependence

The trigger efficiency can depend on the occupancy in the detector, as hits belonging to other tracks can be added to lepton candidates and confuse the pattern recognition. We therefore study the trigger efficiency as a function of how many inelastic interactions we add on top of the J/ψ producing interactions. The extra inelastic interactions are added according to a Poisson distribution with variable mean.

For ECAL, an increase in occupancy will mainly give a slight increase in the cluster energy, which has little effect on the seeding of L2Sili. We can also get additional high E_T clusters. In the latter case the trigger efficiency will only be affected if we get more than ten clusters as this can overflow the buffer of track candidates in the SLT code, but this seems never to happen. We therefore have negligible influence on the performance of the calorimeter part of the trigger from an increased amount of particles.

In the VDS the additional hit occupancy come from two sources: tracks from the underlying events and detector noise. The underlying events include both tracks from the interaction that produced the J/ψ and other interactions in the same proton bunch. The number of interactions is Poisson distributed with a mean of 0.6 (5 MHz interaction rate with $\sim 8.3 \cdot 10^6$ filled bunches per second). We have measured that the track multiplicity in J/ψ interactions is underestimated by approximately 50% and the noise is not included at all in simulation. To evaluate how many additional hits are expected in an RoI, we used RoIs from J/ψ Monte Carlo to request hits in real data. For the data, we used a sample of random triggered events at 5 MHz with the empty events removed and a sample of SLT bypass events. The random triggered events will underestimate the occupancy as we have seen that J/ψ events have higher occupancy than ordinary inelastic events, while the SLT bypass events will overestimate the occupancy as the pretrigger is biased toward

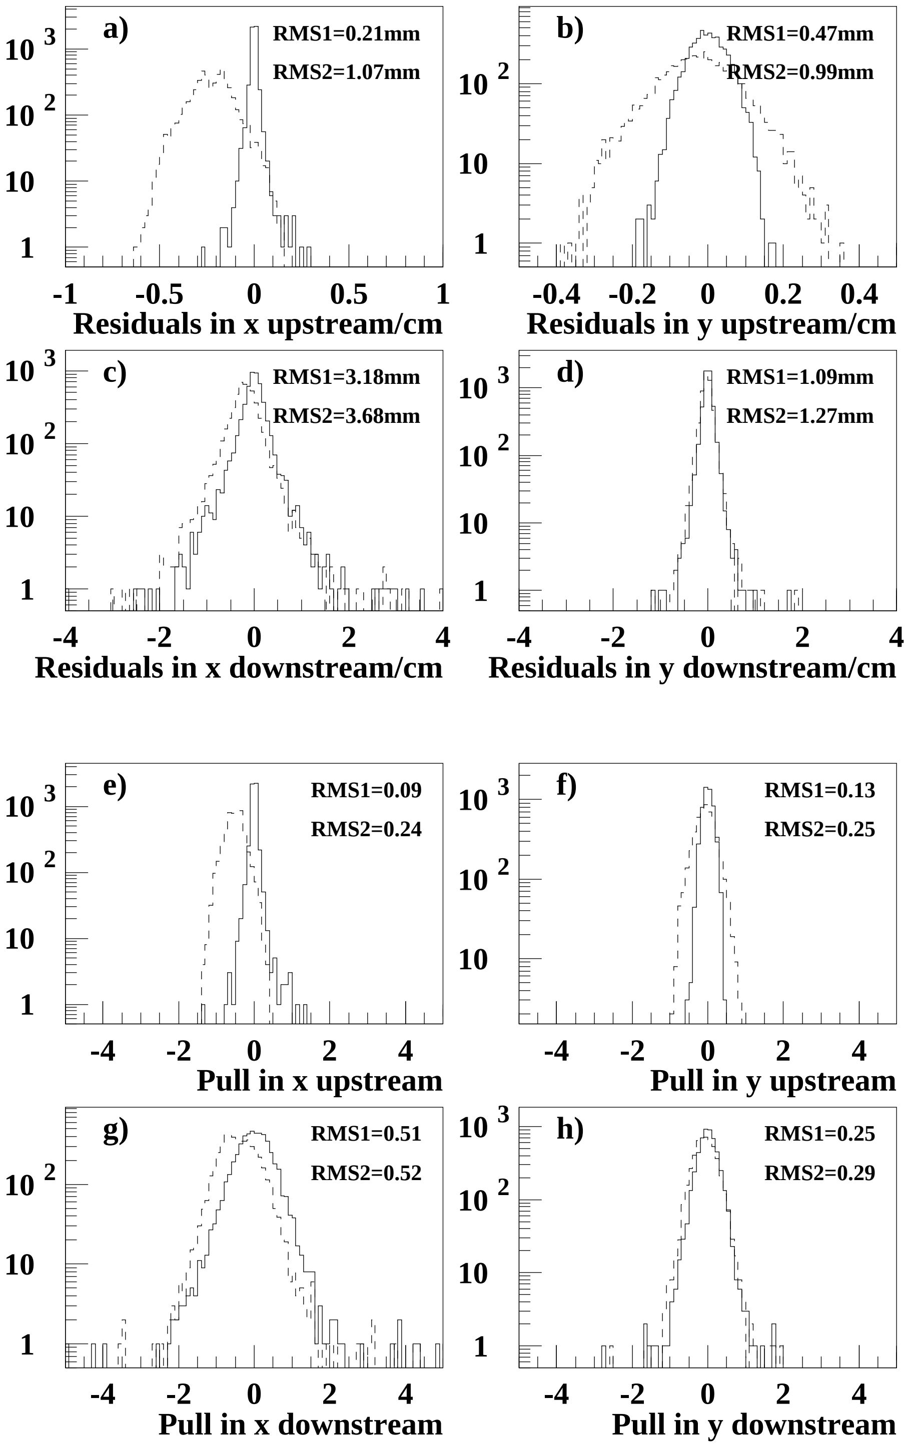

FIGURE 5.2: a-b) The difference between the predicted L2Sili RoI and the true position in x and y at $z = 6$ cm, for a single target wire (full line, RMS1) and two target wires (dashed line, RMS2). For the two target wires we use a target box including Inner 2 and Below 1 wires, but the J/ψ mesons are produced only at Inner 2, which gives rise to the shifted residuals. c-d) The same as a-b), but at $z = 155$ cm. e-h) The deviation of the RoI from a-d) divided by half the width of the RoI (pulls). Values outside $[-1, 1]$ indicate that the track is not fully contained in the RoI.

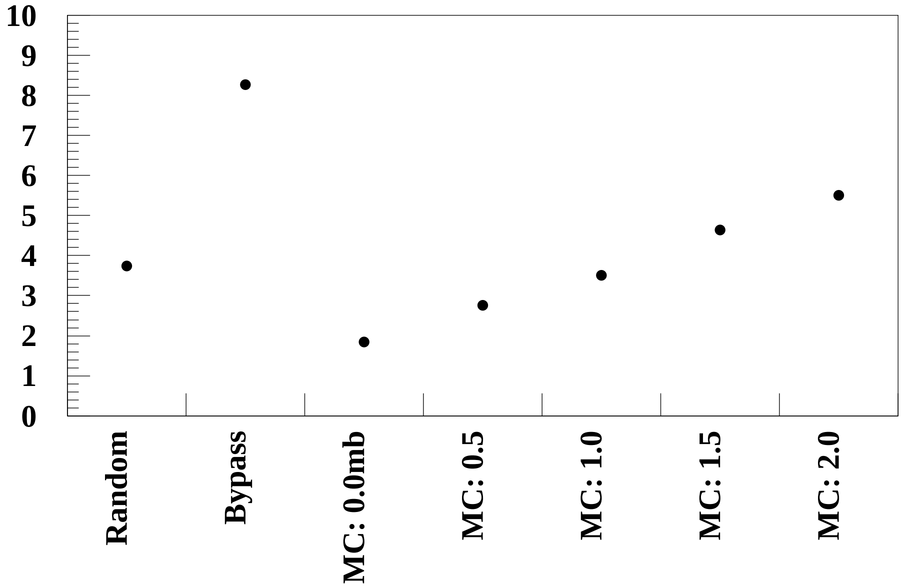

FIGURE 5.3: The mean number of hits in a J/ψ RoI in non-empty random trigger events, in SLT bypass events and in Monte Carlo events with an increasing amount of overlapping inelastic events (from a zero to two events in mean, Poisson distributed). The number of extra hits in an RoI is approximately given by an exponential distribution.

triggering on high multiplicity events. Figure 5.3 shows the mean number of hits measured in this way. It is compared to the observed number of extra hits in simulated $J/\psi \to e^+e^-$ events with different amounts of simulated inelastic events added (Poisson distributed). We see that adding between one and two inelastic events gives the observed RoI occupancy. However as can be seen in Table 5.2, the L2Sili efficiency does not depend strongly on the occupancy within the investigated range.

Monte Carlo results

The efficiency of the ECAL+L2Sili trigger separated in the different factors is listed in Table 5.2. The evaluation has been carried out on the following Monte Carlo samples:

1. Pure $J/\psi \to e^+e^-$ events, ideal detector performance.
2. Pure $J/\psi \to e^+e^-$ events, detector performance of run 15230.
3. $J/\psi \to e^+e^-$ events, no overlapping events (run 15230 performance).
4. $J/\psi \to e^+e^-$ events, 0.5 overlapping events (run 15230 performance).
5. $J/\psi \to e^+e^-$ events, 1.0 overlapping events (run 15230 performance).
6. $J/\psi \to e^+e^-$ events, 1.5 overlapping events (run 15230 performance).
7. $J/\psi \to e^+e^-$ events, 2.0 overlapping events (run 15230 performance).

Sample	1	2	3	4	5	6	7
Geo. acceptance	7.3%	7.3%	7.3%	7.3%	7.3%	7.3%	7.3%
Calorimeter	100.0%	72.9%	72.9%	72.9%	72.9%	72.9%	72.9%
Pretrigger	98.7%	51.3%	50.4%	50.4%	51.3%	50.5%	50.2%
Lepton charge	91.3%	91.6%	91.0%	91.2%	90.6%	90.3%	91.4%
Invariant mass	96.8%	96.1%	94.9%	95.5%	95.1%	95.7%	95.5%
L2Sili RoIs	98.4%	97.3%	97.0%	97.9%	97.7%	96.5%	96.3%
VDS hits	94.4%	94.3%	94.0%	93.6%	93.2%	93.7%	93.3%
L2Sili algorithm	98.9%	98.3%	97.6%	96.7%	98.1%	96.2%	96.2%
Total Efficiency	5.84%	2.16%	2.06%	2.07%	2.10%	2.02%	2.01%

TABLE 5.2: The $J/\psi \to e^+e^-$ trigger efficiencies with the ECAL+L2Sili trigger separated in individual factors. The statistical error on the total efficiency is about 5%. The numbers are correlated as the J/ψ kinematics is the same in the samples 1 and 2 and in 3-7. See the text for the definition of the different Monte Carlo samples.

Sample	1	2	3	4	5	6	7
L2Sili RoIs	94.2%	95.8%	94.2%	93.3%	94.5%	94.2%	94.0%
Silicon hits	93.4%	92.6%	92.9%	92.6%	93.0%	94.0%	92.9%
L2Sili algorithm	97.8%	97.3%	95.0%	97.3%	97.3%	96.6%	96.1%
Total Efficiency	5.46%	2.07%	1.92%	1.96%	2.01%	1.98%	1.95%

TABLE 5.3: The $J/\psi \to e^+e^-$ trigger efficiencies with the ECAL+L2Sili trigger as in Table 5.2, but with two target wires. Only values which differ are shown.

Overlapping events are the amount of inelastic events added on top of the $J/\psi \to e^+e^-$ interaction. Table 5.3 lists the results for the case of several target wires.

From the tables we see that the major inefficiencies come from the geometrical acceptance and the non-working areas in ECAL and the ECAL pretrigger. The geometrical acceptance is low mainly because of the high E_T threshold determined by rate limitations in the trigger chain. If a lower threshold had been used, we would have had worse S/B ratio in the triggered data. We would also have increased the deadtime in the DAQ and therefore not increased the $J/\psi \to e^+e^-$ rate. Similarly if the calorimeter or pretrigger acceptance had been higher we would also have more background events coming into the SLT, which leads to increased deadtime.

The SLT algorithm itself has an efficiency of about 80% if acceptance factors are excluded. Most of that is due to the out-bender assumption and bremsstrahlung tails in the J/ψ mass spectrum. The increase of the target area by adding a second wire lowers the total efficiency by about 5%.

Comparison to data

The 2% trigger efficiency from the Monte Carlo study above gives a predicted rate of 380 triggered J/ψ per hour. This efficiency estimate should be checked using real data. Unfortunately that is not completely straightforward as we must use the trigger to obtain a sample

Selection cut	$J/\psi \to e^+e^-$ efficiency	Expected $J/\psi \to e^+e^-$ events	# of events
Trigger	1.9%	2900	382,242
Matching	1.7%	2600	73,630
Vertex $\chi^2 < 4$	1.6%	2400	49,354
$\lvert E/p - 1\rvert < 3\sigma$	1.4%	2100	43,217
Bremsstrahlung	0.19%	290	443
Fitted $J/\psi \to e^+e^-$	0.18%	270	$230 \pm 21 \pm 23$

TABLE 5.4: The $J/\psi \to e^+e^-$ efficiency after applying different selection cuts. Also shown is the expected number of $J/\psi \to e^+e^-$ events in run 15230 compared to the total amount of events surviving the different cuts. Trigger efficiency has been corrected for a 10% deadtime compared to the value listed in Table 5.2.

of $J/\psi \to e^+e^-$ data. We therefore can only compare the observed number of J/ψ events in triggered data with the expected number, which relies on our estimate of the J/ψ production rate. It is not possible from the triggered data to precisely determine where $J/\psi \to e^+e^-$ events are lost compared to the simulation.

After the trigger, we still have a $S/B \sim O(10^{-2})$ and we therefore use the offline reconstruction to remove the background in the sample by applying particle identification cuts. The $J/\psi \to e^+e^-$ selection is a simplified version of the selection presented in Chapter 8. The cuts applied here are the following:

1. ECAL cluster and L2Sili tracks are confirmed by offline reconstructed ECAL+VDS tracks.
2. Dilepton vertex fit has a $\chi^2 < 4$.
3. $|E/p - 1| < 3\sigma$, where σ is the Gaussian width of the E/p distribution for electrons. Here E is the energy measured in the calorimeter and p is the momentum calculated from the bending of the track in the magnet.
4. One can associate a bremsstrahlung cluster to both electrons.

We do not use tracking information from the OTR and the momentum is therefore calculated using the VDS track and the ECAL cluster position.

Table 5.4 gives the efficiency of this selection on $J/\psi \to e^+e^-$ Monte Carlo and the amount of events after each selection step for run 15230. The main requirement for getting a clean sample is the bremsstrahlung association, although requiring the trigger tracks to be found offline removes a large amount of fake tracks. The bremsstrahlung association reduces the signal by more than a factor seven, but gives almost a factor hundred in background suppression and has a well understood efficiency. For further discussion on the bremsstrahlung requirement, please refer to Chapter 8.

The invariant mass spectrum, calculated from the offline reconstructed tracks, is shown in Figure 5.4. It has been fitted with a Gaussian for the signal and a second degree polynomial times an exponential function for the background. The fit estimates $230 \pm 21 \pm 23$ $J/\psi \to e^+e^-$ events, where the second uncertainty term is due to uncertainties in the background and signal shape. The observed number of J/ψ events is consistent with the expected the 270 events expected from Monte Carlo.

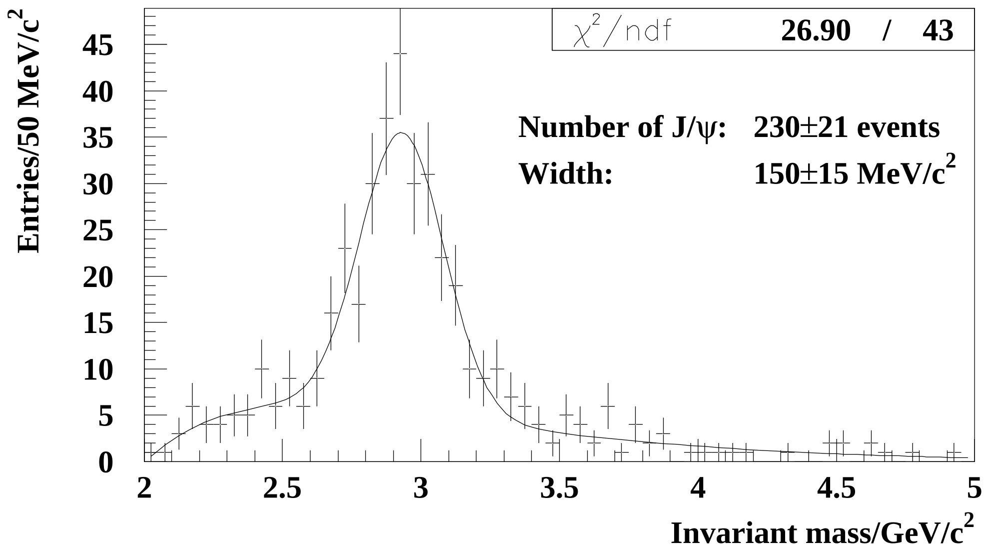

FIGURE 5.4: The invariant mass spectrum of e^+e^- pairs in run 15230 after requiring a bremsstrahlung cluster associated to both leptons. The full line shows the result of the fit.

5.2.2 Tracking precision

In the 2000 run L2Sili was the last step in the SLT trigger algorithm and the tracks found by it were used only to confirm that a particle passed within the RoI. The track precision achieved by L2Sili was therefore not important, but for the 2002 run we will use the tracks from L2Sili in L2Vertex. The L2Vertex algorithm requires precise track parameters as input and we will therefore study the precision of the code used for the 2000 run.

To measure the tracking precision, we use $J/\psi \to e^+e^-$ Monte Carlo with an average of one inelastic event added on top. We use the simulation with both one target wire and two target wires, but we observed no difference in the tracking precision. Figure 5.5 shows the residuals of both track slopes and positions to the true track. The track parameters are calculated at the z position of the hit on the L2Sili track which is closest to the target. Each residual distribution is fitted with a single Gaussian. The fitted widths are compared in Table 5.5 to the results obtained from a similar fit using tracks found by the offline VDS tracking (CATS) [43]. CATS reconstructs space points from the hits in the VDS and uses a full Kalman filter to calculate the track parameters. The track slopes found by L2Sili are 30% worse than what CATS finds. This is mainly due to the way multiple scattering is treated by L2Sili and because less correct hits are assigned to the track. The difference will be reduced for the 2002 version of L2Sili which has a more correct calculation of multiple scattering.

To check our estimate of the track uncertainty, calculated from the covariance matrices of L2Sili, we calculate the pull (normalized residuals) defined as:

$$\mathrm{Pull}(a) = \frac{a_{\mathrm{measured}} - a_{\mathrm{truth}}}{\sigma_a}, \qquad (5.2)$$

where a_{measured} is any of the measured track parameters, a_{truth} is the true value and σ_a is

	L2Sili	CATS
$\Delta x/\mu$m	11.8	10.3
$\Delta y/\mu$m	11.1	10.5
Δt_x/mrad	0.09	0.07
Δt_y/mrad	0.09	0.07
Pull(x)	0.44	1.08
Pull(y)	0.39	1.12
Pull(t_x)	0.69	1.28
Pull(t_y)	0.61	1.27

TABLE 5.5: Width of residuals and pulls of L2Sili track parameters extracted from Gaussian fits, see Figure 5.5. The rightmost column shows the equivalent results given by the offline tracking (CATS).

our uncertainty estimate. If the covariance matrices are calculated correctly we should have a Gaussian distribution for the pulls centered at zero and with a width of one. As can be seen in Figure 5.5 and Table 5.5, the slope and the position uncertainties are overestimated by L2Sili.

5.2.3 Trigger rates

The trigger rate is the number of events accepted by an SLT trigger algorithm step per second to be processed by the next step or, after the last step, to be written to tape. The trigger rates in the 2000 run were not always stable from run to run as the acceptance of the calorimeter was steadily increasing and the beam was not always equally distributed in the different bunches. The bunch distribution affects the trigger rate, since the probability to trigger on an event does not scale linearly with the number of interactions in that event. If some bunches have a higher interaction probability than others, the number of events with multiple interactions will increase even though the average interaction rate is kept fixed at 5MHz. Because the probability to trigger on an event does not increase linearly with the number of interactions, we will have higher trigger rates if the interactions are not distributed evenly over all bunches.

The trigger rates in run 15230, shown as function of time in Figure 5.6, are quite stable. In Table 5.6 the mean trigger rates from the run are shown together with the rejection factor. The 10% deadtime of the DAQ system are due to overflow of readout-FIFO as described in Section 3.7.4. The $E_T > 1.0$ GeV cluster cut in the SLT reduces the rate mainly by suppressing fake electron pretrigger messages, caused by noise in the pretrigger electronics. L2Sili is able to give a large suppression as most of the high E_T clusters found in the calorimeter are from photons and therefore do not give a track in the vertex detector. The output rate from the SLT consists of the 12 Hz of fully triggered events, 4 Hz$\cdot$0.9 = 3.6 Hz random triggers[4] and 550/200 = 2.75 Hz SLT bypass events[5]. As we shall see below, all parts of the SLT are able to handle rates that are this low.

The trigger rates cannot be predicted reliably by simulation for several reasons. As the total reduction is of $O(10^5 - 10^6)$, we need a large amount of Monte Carlo statistics, which has to be

[4]The factor 0.9 in the random trigger calculation is due the to the 10% deadtime, which the random trigger, nominally running at 4 Hz, also sees.

[5]For ECAL+L2Sili trigger configuration one in 200 of the ECAL accepted events were bypassing the VDS tracking.

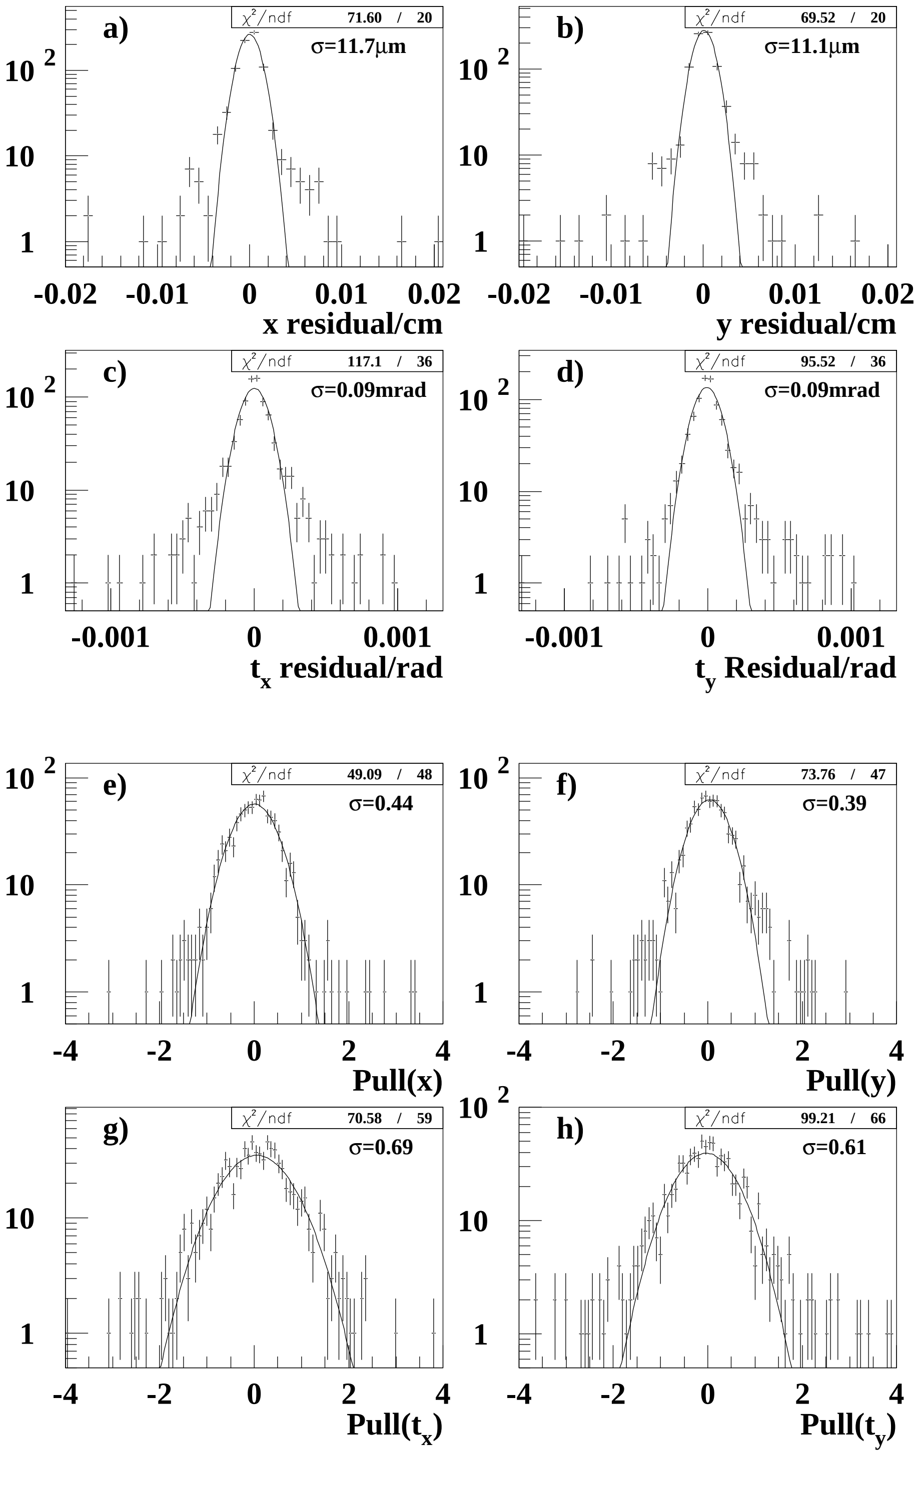

FIGURE 5.5: a-d) L2Sili tracks residuals for slopes and track position at the first at the VDS plane closest to the target. The full line shows the result of Gaussian fit. e-h) Pulls for the residuals shown in a-d) with Gaussian fit shown with the full line.

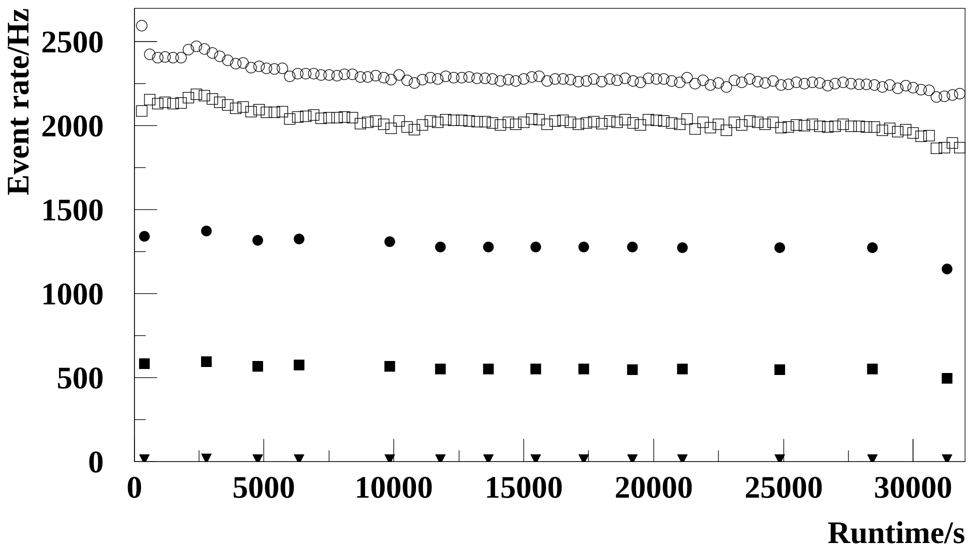

FIGURE 5.6: Event rates at different places in the trigger for run 15230 as a function of run time. Open circles: output rate of pretrigger system. Open squares: input rate to SLP (reduced by deadtime). Solid circles: after $E_T > 1.0$ GeV cut in SLP. Solid boxes: after mass cut (input to L2Sili). Triangles: output rate of the SLT.

accurate at that level. Furthermore many of the events are composed of fake tracks, where high E_T photons (from π^0 decays) in the calorimeter are matched by the SLT with charged hadron tracks in the VDS. Table 5.7 shows the result of running the trigger simulation on a sample of 600,000 inelastic events[6]. The number of interactions per event is Poisson distributed with a mean of 0.5 events[7]. This corresponds to 57 ms of data taking at 5 MHz interaction rate. We see that the Monte Carlo predicts a significantly lower trigger rate for the electron pretrigger, though part of the discrepancy can be explained by the noisy electronics, which is not simulated in the Monte Carlo. The invariant mass cut is reasonably well reproduced, while the statistics is too low to estimate the L2Sili reduction. Also shown in the table are rates expected for an ideal working detector. The rate from the pretrigger increases in this case by more than 600%. About 50% of that increase is due to the inclusion of the inner most part of the calorimeter, which was working in the 2000 run, but excluded from the pretrigger since electrons giving signal in this part of the calorimeter were not in the geometrical acceptance of the VDS.

5.2.4 Timing

To understand the rate limitation of the trigger system we need to know not only the trigger rates at different steps in the SLT trigger algorithm, but also the CPU time necessary for the calculations done at each step. Unfortunately this was not systematically recorded during

[6]This is the largest Monte Carlo sample of inelastic events available for the 2000 configuration.

[7]In real data, the mean number of interaction is 0.6 events per filled bunches as only 180 out 220 HERA bunches are filled.

Trigger step	Input rate/Hz	Output rate/Hz	Reduction factor
Pretrigger	$10.4 \cdot 10^6$	$2.25 \cdot 10^3$	4600
DAQ deadtime	$2.25 \cdot 10^3$	$2.0 \cdot 10^3$	1.12
Cluster $E_T > 1.0$ GeV	$2.0 \cdot 10^3$	$1.28 \cdot 10^3$	1.56
Mass> 2.0 GeV/c^2 cut	$1.28 \cdot 10^3$	550	2.31
L2Sili	550	12	45
4LT	12 (+6)	18	1

TABLE 5.6: Trigger rates measured for the run 15230.

Trigger step	Events (15230)	Red. factor	Events (ideal)	Red. factor
Pretrigger	40	15,000	292	2050
Cluster $E_T > 1.0$ GeV	36	1.1	235	1.2
Mass> 2.0 GeV/c^2	16	2.3	128	1.8
L2Sili	0	—	1	~ 128

TABLE 5.7: The number of inelastic Monte Carlo events which survive the different steps in the trigger algorithm out of an initial 600,000. Column two and three show the amount of events and the corresponding reduction factor for data simulated with run 15230 detector performance, while column four and five corresponds to an ideal detector.

running. Instead the timing has to be extracted by running the trigger simulation on already recorded data (using the SLT bypass events). The trigger simulation does not include the data readout via the SHARC switch or the sparsification of the VDS data. The time for these operations have been estimated from later runs, where the processing time was recorded online.

The steps in the SLT trigger algorithm which take a non-negligible time are:

- **ECAL readout** — Messages have to be sent via the SHARC switch to the SLB and back again. The time spent is independent of CPU speed.
- **Cluster reconstruction**
- **VDS hit request** — For each RoI, the SLT scans each VDS layer to check if it overlaps with the RoI, and sends a hit-request for the overlap region. The actual time for retrieving the data is small and it is partially done in parallel with the hit preparation, but the check of overlap with each layer is quite time consuming.
- **VDS hit preparation** — To extract the hits from from the raw data of the VDS, a sparsification is carried out, which involves a CPU intensive noise subtraction procedure.
- **L2Sili algorithm**
- **Event building**

Other steps such as the E_T and mass cuts on the electron clusters and the magnet propagation represent a negligible time. Table 5.8 shows the time spent by each step, every time it is called (per invocation) and the average time for all events.

Trigger step	Time per inv. (300 MHz)	Time per event (300 MHz)	Time per inv. (390 MHz)	Time per event (390 MHz)
ECAL readout	4.6 ms	4.6 ms	4.6 ms	4.6 ms
Cluster reconstruction	15 ms	15 ms	12 ms	12 ms
VDS hit req.	0.9 ms	0.3 ms	0.7 ms	0.2 ms
VDS hit prep.	7.3 ms	2.0 ms	5.6 ms	1.6 ms
L2Sili algorithm	7.8 ms	2.2 ms	6.0 ms	1.7 ms
Event building	270 ms	2.4 ms	220 ms	2.0 ms
Total	—	27 ms	—	22 ms

TABLE 5.8: Timing of the trigger steps in the ECAL+L2Sili trigger. The tables shows both the time needed for each invocation of a trigger step and the time spent per event after weighting with the reduction factors from Table 5.6. Column two and three show the timings for a 300 MHz CPU, while the last two columns show the timings for an "average CPU" of the full SLT farm.

Most of the ECAL+L2Sili runs were recorded using only a hundred SLT nodes[8], each with a 300 MHz CPU. In Table 5.8 we therefore list both the time for a 300 MHz CPU and for a "390 MHz CPU", corresponding to an average of 100 CPUs of 300 MHz and 140 CPUs of 450 MHz. All CPU related timings scale according to the CPU speed, while data readout is independent of CPU speed.

Table 5.8 shows that with 100 nodes we were able to handle input rates up to 3.7 kHz and with 240 nodes up to 11 kHz. During running we had only 1.5-3.0 kHz input rate, and we could therefore have lowered the E_T threshold on electron pretriggers. However the 4LT reconstruction was quite slow and only able to handle rates of about 30 Hz. This gave the limit for the trigger configuration, which led us to set the high E_T threshold.

5.2.5 Fake tracks

Only a small fraction of the tracks the SLT finds are electrons, the rest are hadrons interacting strongly in the calorimeter and fake tracks. The fake tracks can be divided into "fake matches" and "ghost tracks". Fake matches can occur when a charged particle (mainly hadrons) passes through the VDS RoI predicted by a cluster originating from a different particle (mainly photons). Ghost tracks are tracks in which either the cluster or the L2Sili track is not given by a real particle, but consists of hits from several particles and/or electronics noise. Since we have very high requirement for the cluster energy, there is a very small probability to manufacture a ghost cluster, except from merging of neighboring clusters. We therefore study ghost tracks only in L2Sili.

The amount of fake matches is studied using the technique of mixing hits from uncorrelated events ("event mixing"), while the ghost track rate is determined via a comparison with offline reconstructed VDS tracks. These studies are described in the following.

[8] 80 nodes were not cabled until beginning of June 2000, and other 60 were not running due to cooling problems

	Unmixed	Mixed e^-	Mixed e^+	Mixed e^+e^-	Real tracks
$P(e^-)$	11.7 ± 0.3	5.9 ± 0.2	—	—	5.8 ± 0.4
$P(e^+)$	13.8 ± 0.3	—	7.6 ± 0.3	—	6.2 ± 0.4
$P(e^+e^-)$	1.6 ± 0.1	0.71 ± 0.08	0.82 ± 0.08	0.62 ± 0.07	0.7 ± 0.1

TABLE 5.9: Probabilities in percent for L2Sili to find tracks in a e^+ or e^- RoI or both at the same time, when some of the hits come from a different event. The last column gives the calculated probability of finding tracks not due to fake matches.

Event mixing

The amount of fake matches can be studied using event mixing. In event mixing we use clusters from one event to setup L2Sili RoIs, but use the VDS hits from the previous event for the L2Sili algorithm. If a track is found in the event mixing, we know it is a fake match. We trigger on pairs of tracks and the probability to have a fake track is not independent of whether the other track is fake or not, as the fake rate increases faster with track multiplicity than the rate of real tracks. We therefore do the event mixing in three ways: only hits for the e^+ RoIs are taken from the previous event, only hits for the e^- RoIs are taken from the previous event and hits for both RoIs are taken from the previous event. From these samples we can extract the probability to trigger on a fake e^+ and any "e^- track", $P(a, f)$, on a fake e^- and any "e^+ track", $P(f, a)$, and on two fake matches, $P(f, f)$. Combined with the total probability of triggering on two tracks $P(a, a)$, we can calculate the probability to trigger on two real tracks, $P(r, r)$ as:

$$P(r, r) = P(a, a) - P(a, f) - P(f, a) + P(f, f). \tag{5.3}$$

The probabilities extracted from the event mixing are listed in Table 5.9. From this table we see that about 55% of the triggered events contain at least one fake match.

Table 5.9 also shows trigger probabilities for a single e^+ or e^- in mixed and unmixed events. Only 45% of the e^+ and 50% of the e^- are real tracks, which shows that there is a correlation between the fake e^+ and e^- as the pair probability is not given by the product, but rather that most fake matches come in pairs. The e^+ and e^- probabilities are different due to differences in acceptance between the negative-x and positive-x sides of the detector.

Track confirmation

To study the amount of ghost tracks in the trigger we compare the tracks found by L2Sili to the tracks found by the offline tracking (CATS) The ghost rate in CATS has been studied [43] using simulated events and found to be between 7% and 15% of the reconstructed tracks, depending on the number of interactions. For the SLT we cannot use the simulation since a huge amount of statistics would be necessary due to large suppression in the SLT. Instead we will measure the ghost rate of L2Sili by studying how often the L2Sili tracks, found in run 15230, are also found by CATS.

We will compare the tracks using three χ^2s defined as:

$$\chi_x^2 = \frac{|t_x(\mathrm{SLT}) - t_x(\mathrm{CATS})|^2}{\sigma_{t_x}^2(\mathrm{SLT})} + \frac{|x(\mathrm{SLT}) - x(\mathrm{CATS})|^2}{\sigma_x^2(\mathrm{SLT})}, \tag{5.4}$$

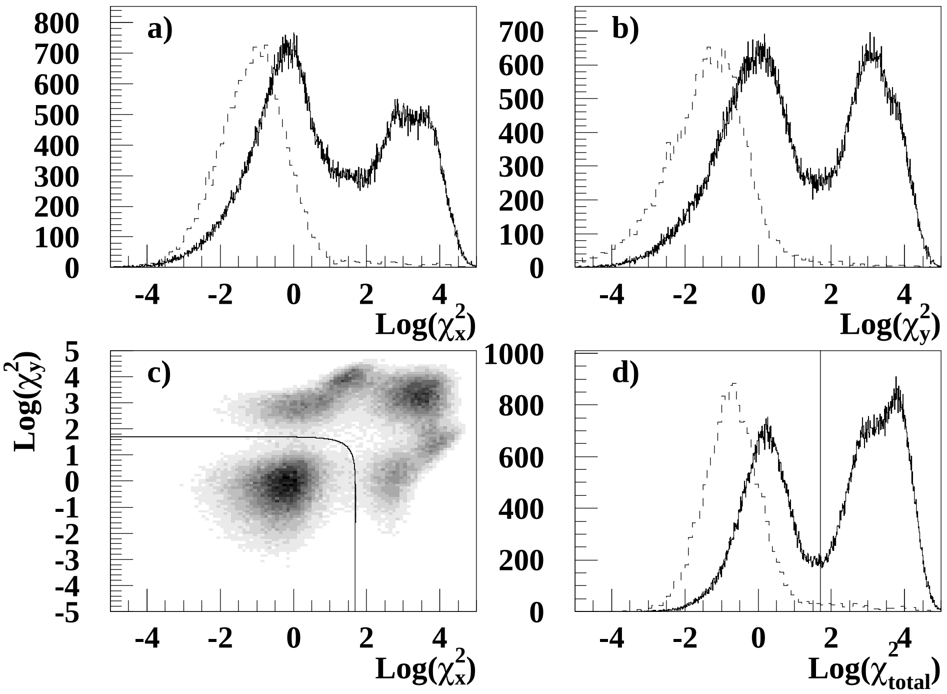

FIGURE 5.7: Log(χ^2) distributions from matching L2Sili tracks with offline reconstructed VDS tracks. a) χ^2 in x view. b) χ^2 in y view. c) χ^2 in y view versus. χ^2 in x view. d) Total χ^2. The dashed line in a), b) and d) is the result for $J/\psi \to e^+e^-$ Monte Carlo.

$$\chi^2_y = \frac{|t_y(\mathrm{SLT}) - t_y(\mathrm{CATS})|^2}{\sigma^2_{t_y}(\mathrm{SLT})} + \frac{|y(\mathrm{SLT}) - y(\mathrm{CATS})|^2}{\sigma^2_y(\mathrm{SLT})}, \tag{5.5}$$

$$\chi^2_{\mathrm{total}} = \chi^2_x + \chi^2_y. \tag{5.6}$$

The x and y coordinates and their uncertainties are calculated at the first hit on the CATS track. Given that a substantial amount of the hits are in common to the two tracks, $\chi^2_{x/y/\mathrm{total}}$ are not distributed according to a true χ^2 distribution, but they give a good separation of ghost from real tracks. This can be seen in Figure 5.7, which shows the χ^2s for triggered L2Sili tracks matched with CATS tracks. The χ^2s for a L2Sili track are only shown for the CATS track which gives the lowest χ^2_{total}.

The χ^2 distributions reflect three classes of L2Sili tracks: tracks with low χ^2s (real tracks), tracks with high χ^2s (ghosts) and tracks where χ^2_x is low and χ^2_y is high or the reverse. The last class consists of real tracks, which pass close to, but outside the L2Sili RoI. If a track outside the RoI passed through VDS micro-strips, which lie partially inside the RoI in one view, L2Sili can in some cases reconstruct the track almost correctly in that view and combine it with random hits in the other. As the tracks lies outside the RoIs the last class is also classified as ghost track.

Based on Figure 5.7 we define a L2Sili track as a ghost track if $\chi^2_{\mathrm{total}} > 50$. If CATS was inefficient we would overestimate the ghost rat. However, from Monte Carlo studies [43] CATS is expected to be around 97% efficient, consistent to the loss observed in triggered $J/\psi \to e^+e^-$ Monte Carlo. Using this definition of ghost tracks, 57% of the L2Sili tracks are classified as ghosts and 80% of the track pairs accepted by the SLT contain at least one ghost tracks.

Using CATS instead of L2Sili

We have seen that the tracks found by L2Sili are dominated by ghost tracks. We can use tracks found by CATS, which has much lower ghost rate, to estimate how many non-ghost tracks we should find and how many of those are due to fake matches. This is done by checking how often we find a track reconstructed by CATS inside the RoI given to L2Sili. We use SLT bypass events in order not to have L2Sili bias the sample. The fake matching rate is estimated by using event mixing as was done for L2Sili.

Figure 5.8 shows the normalized residuals of CATS tracks to the center of the RoI at the entrance point of the VDS ($z = 4$ cm) and at the exit point ($z = 160$ cm). The normalization is done with respect to half the width of the RoI and only residuals for tracks which are within the RoI at both endpoints are shown in Figure 5.8. In all residuals, except for the residual in x at the exit point of the VDS, we have clear peaks from tracks passing through the RoI, though the y residuals are shifted by about 1σ from zero. This shift is caused by the beam, which was shifted by several millimeters from the nominal position during the run. This shift was not taken into account in the online definition of the target box. However the tracks were still contained inside the RoI and it caused only little efficiency losses. By construction, almost all tracks are inside the RoI at the beginning of the RoI as they were defined to contain tracks originating from the target.

The residual distribution in x at the exit point of the VDS is shown in Figure 5.9 in a wider range. The peak position is seen to be charge dependent, which is not the case for the other residuals. This charge dependence is understandable if we assume that the ECAL cluster is caused by a hadron. In that case the $E/p = 1$ assumption used to setup the RoI is no longer true. Instead the energy measured in the calorimeter has a quite flat distribution between zero and the real momentum, p. We therefore overestimate the bending in the magnetic field, which depends on the charge, and we get the peak structure shown in Figure 5.8. L2Sili will only find real tracks inside the RoI, which are the tracks with $E/p \approx 1$, while part of the tracks with $E \ll p$ will be reconstructed as ghosts tracks in the y coordinate.

The fake matching rate is estimated from event mixing, similar to the procedure applied for L2Sili, and the results are listed in Table 5.10. We see that with CATS the amount of fake matches are reduced to about 30% of the single tracks and 45% of the RoI pairs accepted contain one or two fake matches. The fraction of fake matches is not much lower than for L2Sili, but there is almost a factor four more rejection from CATS. Given that with CATS we estimate that only 0.22% of the events into L2Sili contain two real tracks and we trigger on 1.6%, only 14% of the L2Sili triggered events contain two real tracks. This shows the potential gain we can have from additional suppression, for instance by using the OTR.

5.2.6 Summary of the ECAL+L2Sili trigger performance

Based on Monte Carlo, we have estimated the ECAL+L2Sili trigger to have an $J/\psi \to e^+e^-$ efficiency of about 2%. Our prediction was confirmed by the observed $J/\psi \to e^+e^-$ rate. The SLT algorithm itself is 80% efficient and the major inefficiencies in the trigger come from missing

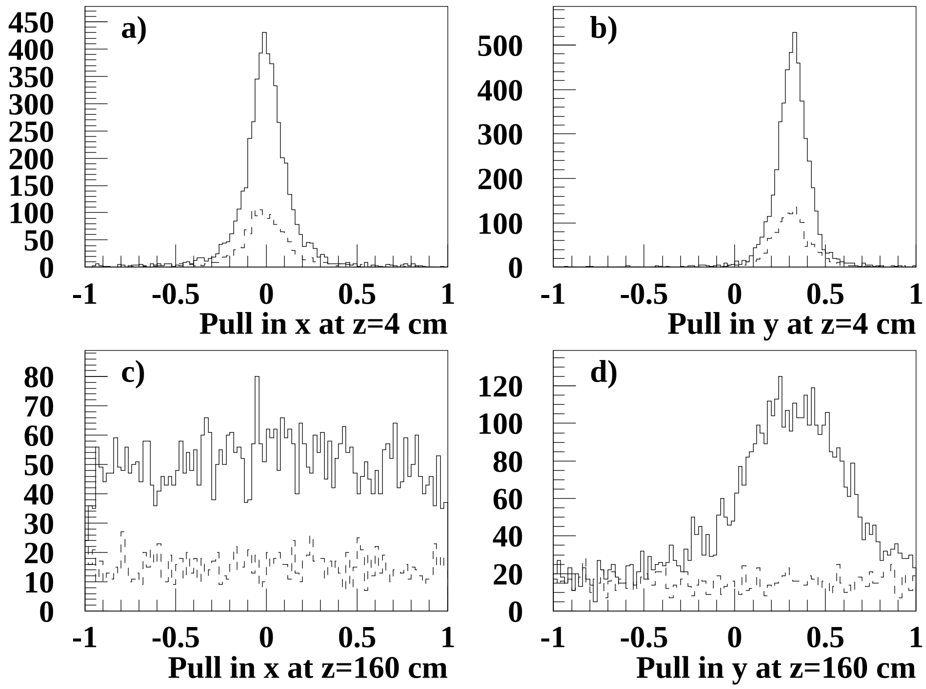

FIGURE 5.8: The normalized residuals of CATS tracks to the RoIs for L2Sili at the entrance point of the VDS, a)+b), and at the exit point of the VDS, c)+d). The dashed lines show the residuals for mixed events.

coverage in the ECAL and the high E_T threshold used by the electron pretriggers. This threshold is used to lower the trigger rates. The SLT had enough CPU power to handle up to five times higher input rate (with the full SLT farm), but the output rate was limited to around 30 Hz by the 4LT reconstruction. The SLT output is dominated by fake tracks, contributing to about 85% of the triggered events, mainly due to ghost tracks in L2Sili. This high ghost rate is one of the main motivations to introduce the OTR tracking in the SLT.

5.3 ECAL+SLT trigger

The ECAL+SLT trigger improves the ECAL+L2Sili trigger by tracking the electrons through the OTR before the VDS. This gives us faster and additional suppression of fake tracks, which allows us to run at higher input rates. Moreover the MUON+SLT trigger can be operated in parallel, since the trigger is not limited by the output rate, as is the case of the ECAL+L2Sili trigger.

5.3.1 Trigger efficiency

For the efficiency evaluation we use the same Monte Carlo samples as for the ECAL+L2Sili trigger, but with the detector simulated according to the conditions measured in run 16665.

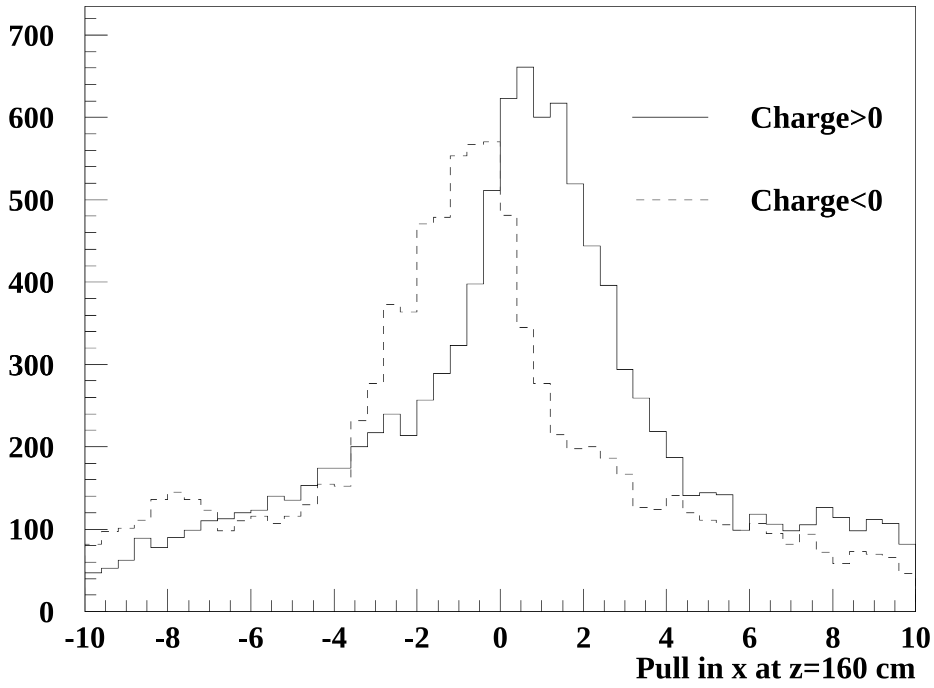

FIGURE 5.9: The residual between CATS tracks and L2Sili RoIs in x at the exit point of the VDS. The full line is for RoIs with a positive charge assumption and the dashed is for negative charge.

The trigger efficiency is separated into the same factors as for the ECAL+L2Sili trigger with a few extra factors describing hit and tracking efficiencies in the OTR. Below we describe the factors which are new or significantly changed compared to the ECAL+L2Sili trigger.

Geometrical acceptance: The electrons are required to pass through both the VDS and OTR and leave high E_T clusters in the calorimeter. For the VDS we require the electrons to pass through at least two super layers. In the OTR, the SLT tracking allows a single superlayer to be empty and we therefore require the electrons to pass through at least five of the six OTR superlayers after the magnet. The requirement for a track to pass through the OTR lowers the acceptance, since the OTR does not cover the inner most part of the calorimeter. This loss in acceptance is however compensated by lowering the E_T cut from 1.2 GeV to 1.0 GeV. The change in acceptance by requiring not only ECAL and VDS, but also the OTR to be present is shown in Figure 5.10 as a function of the E_T.

OTR RoIs: The OTR RoIs are predicted from the ECAL cluster position, the cluster energy and the target box. For an RoI to be acceptable we require that the electrons are contained inside the RoIs at least at five of the OTR superlayers, which matches the geometrical acceptance requirement. Figure 5.11 shows that the RoIs are well centered around the tracks and that the RoIs of ± 2.5 cm are large enough to contain almost all of the electrons.

	Unmixed	Mixed e^-	Mixed e^+	Mixed e^+e^-	Real track
$P(e^-)$	6.3 ± 0.1	2.0 ± 0.1	—	—	4.3 ± 0.2
$P(e^+)$	7.0 ± 0.1	—	2.1 ± 0.1	—	4.9 ± 0.2
$P(e^+e^-)$	0.42 ± 0.03	0.12 ± 0.02	0.13 ± 0.02	0.05 ± 0.01	0.22 ± 0.05

TABLE 5.10: Probabilities in percent for CATS to find tracks in a e^+ or e^- RoIs or both at the same time. We mix tracks from different events as was done in Table 5.9 for L2Sili.

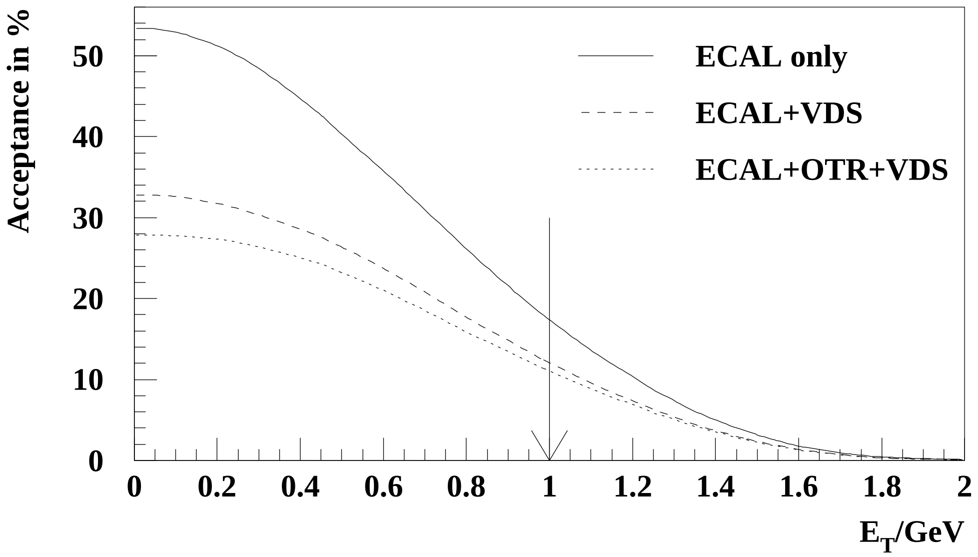

FIGURE 5.10: Geometrical acceptance for two electrons as a function of E_T and requiring them to pass through different detectors. The arrow shows the cut used by the ECAL+SLT trigger, which requires the tracks to be found in all three detectors.

Slicer hits: Slicer requires at least nine hits in the 0° layers and allows at most one of the super layers to be empty. We distinguish the algorithm efficiency (the hits lie in a “slice”) from detector efficiencies by requiring that each electron has at least nine hits in the OTR and that these hits are distributed over at least five of the six superlayers. With an ideal detector (Sample 1 in Table 5.11) with hit efficiencies of 95%, this requirement should always be fulfilled. However, a 11% loss is observed due to the electrons passing through insensitive regions of the OTR, mainly at the detector edges, but also in the overlapping region between different OTR modules.

Slicer: If the required hits are in the RoI, which is ensured by the previous factor, Slicer should only reject tracks, which scatter a lot in the OTR and thus do not lie in a single slice. However the 2000 version of Slicer had a logic bug, which effectively only counted hits in slices which had hits in a maximum of four superlayers. This introduced inefficiencies for tracks which were close to the limit in hits.

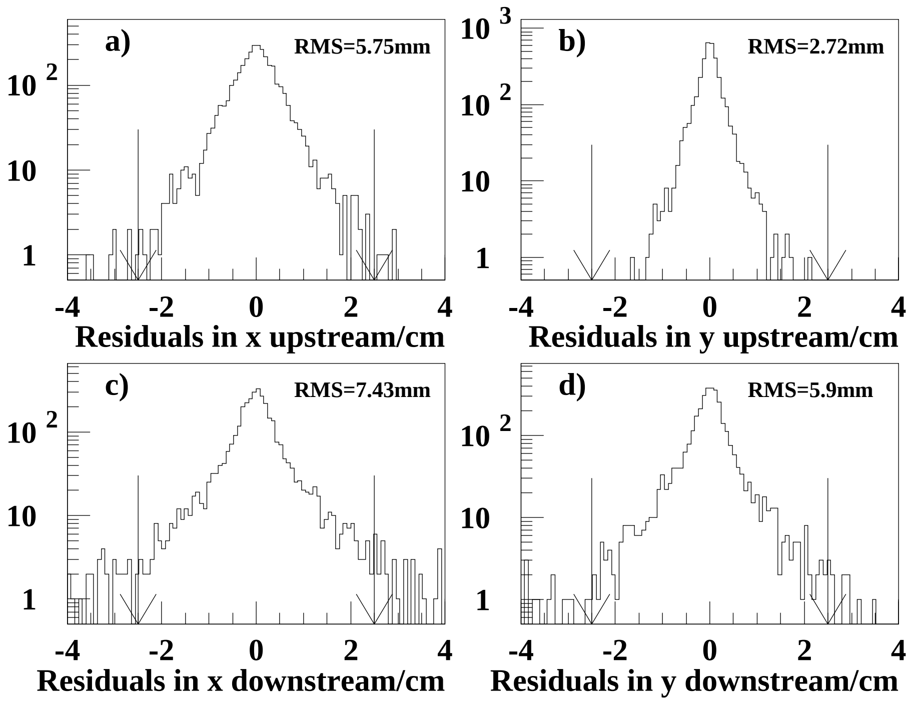

FIGURE 5.11: a-b) The difference between the predicted Slicer/Refit RoI and the true position in x and y right after the magnet ($z = 700$ cm). c-d) The same as a-b), but in the front of ECAL ($z = 1350$ cm). The arrows mark the limits of the RoI.

RefitX hits: RefitX does not require a specific number of hits, but instead allows for a maximum of two consecutive layers to be empty. This requirement is therefore factorized out of the RefitX algorithm efficiency. Only hits caused by the electron tracks are considered in the counting. The main efficiency loss is due to dead areas of the OTR and not to the hit efficiency.

RefitX: If the hits needed by RefitX are found, the algorithm itself only has minor efficiency losses, which are caused by hits rejected by the χ^2 cut or by the pattern recognition being confused by hits from other tracks.

RefitY hits: As for RefitX we require that the stereo layers have a maximum of two consecutive empty layers.

RefitY: The RefitY algorithm is observed to have a 5% inefficiency even when the required hits are present and there are no other tracks to confuse the pattern recognition. Most of this efficiency loss comes from the requirement that the y coordinate of a hit is inside the RoI, when the x coordinate is taken from RefitX. No uncertainty of the x coordinate is taken into account in this calculation, which leads to a loss of hits.

Sample	1	2	3	4	5	6	7
Geo. acceptance	11.0%	11.0%	11.0%	11.0%	11.0%	11.0%	11.0%
Calorimeter	100.0%	79.1%	79.1%	79.1%	79.1%	79.1%	79.1%
Pretrigger	94.2%	59.7%	58.0%	59.0%	60.1%	59.9%	60.1%
Lepton charge	87.7%	86.7%	85.9%	85.2%	85.3%	85.9%	85.8%
Invariant mass	95.1%	93.7%	92.9%	92.8%	92.7%	94.1%	93.0%
OTR RoIs	96.9%	98.8%	98.1%	98.1%	97.5%	97.5%	98.1%
Slicer hits	87.5%	71.5%	74.6%	75.4%	77.2%	77.6%	78.3%
Slicer algorithm	99.6%	94.9%	95.2%	97.7%	95.0%	97.0%	98.0%
RefitX hits	98.2%	84.6%	82.5%	82.2%	80.0%	77.9%	79.9%
RefitX algorithm	98.7%	98.0%	97.8%	97.3%	96.4%	97.5%	96.0%
RefitY hits	97.8%	90.6%	85.4%	86.8%	87.7%	88.0%	89.0%
RefitY algorithm	95.3%	92.2%	91.2%	88.7%	89.9%	91.3%	86.4%
L2Sili RoIs	99.7%	99.9%	99.1%	99.4%	98.9%	99.5%	99.5%
VDS hits	97.9%	97.2%	94.6%	96.1%	95.8%	93.6%	94.8%
L2Sili algorithm	99.6%	98.9%	98.8%	99.4%	99.1%	99.4%	98.9%
Total Efficiency	6.38%	1.88%	1.63%	1.71%	1.68%	1.74%	1.72%

TABLE 5.11: The $J/\psi \to e^+e^-$ trigger efficiencies with the ECAL+SLT trigger separated into individual factors. The statistical error on the total efficiency is about 5%, but the numbers are correlated as the J/ψ kinematics is the same in the samples 1 and 2 and in 3-7. See the text for the definition of the different Monte Carlo samples.

L2Sili RoIs: The tracking of the electrons through the OTR gives a much better prediction of the L2Sili RoIs. This can be seen by comparing the residuals in Figure 5.12 to the ones in Figure 5.2 from the ECAL+L2Sili trigger. The RoI position in x in front of the magnet is improved by more than a factor three, while the position in y is only marginally improved as the OTR has a limited resolution in y compared to the ECAL precision.

Monte Carlo results

The efficiencies measured for the ECAL+SLT trigger are listed in Table 5.11 for a single target wire and in Table 5.12 for two target wires. The Monte Carlo samples are numbered in the same way as in the study of the ECAL+L2Sili trigger in Section 5.2.1.

In the case of ideal detector performance (Sample 1) the efficiency loss caused by introducing the OTR in the SLT is compensated by lowering the E_T cut. In the realistic cases we gain an extra 25% efficiency over the ECAL+L2Sili runs from better coverage in the calorimeter, but we loose more from the large amount of dead channels in the OTR. The efficiency of the VDS tracking gets very close to 100% due to the more precise RoIs. The net result is a trigger which is about 20% less efficient than the ECAL+L2Sili trigger. Within the statistical error, we see no overall dependence of event occupancy. The slicer efficiency increases with the amount of underlying events as hits from other tracks are counted also, but the gain is lost again in Refit, where these hits are rejected after the χ^2 calculation.

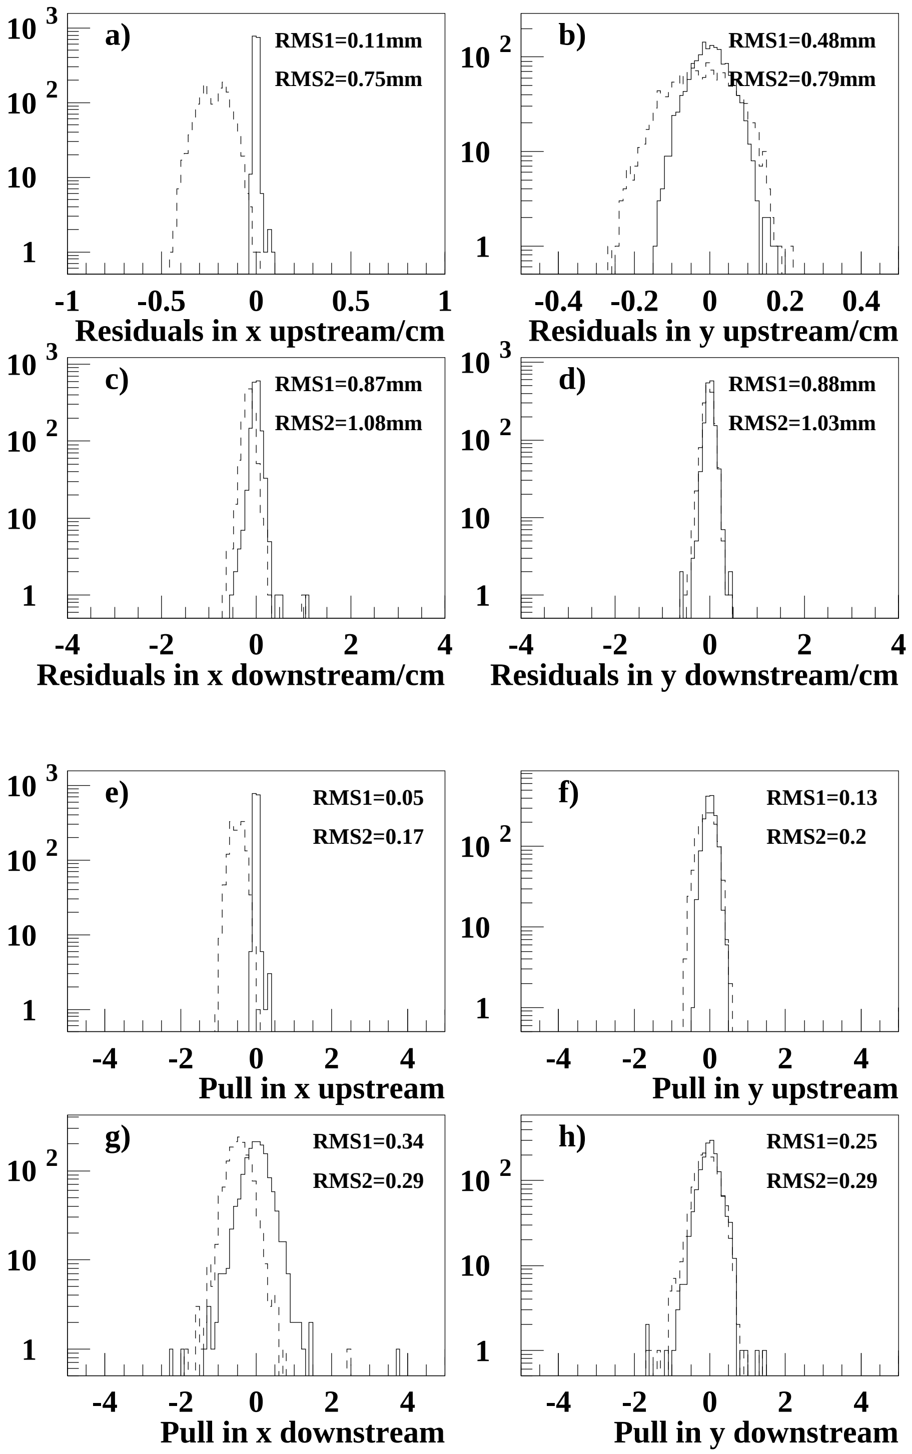

FIGURE 5.12: Residuals and normalized residuals of the L2Sili RoIs as in Figure 5.2, but for ECAL+SLT trigger configuration. The full lines are for a target box with a single target wire, while the dashed lines are for two target wires.

Sample	1	2	3	4	5	6	7
L2Sili RoIs	97.8%	99.6%	99.2%	99.4%	98.9%	98.9%	99.7%
Silicon hits	97.6%	95.9%	93.4%	95.5%	94.4%	92.5%	92.8%
L2Sili algorithm	99.2%	98.8%	99.1%	99.1%	98.8%	99.7%	98.2%
Total Efficiency	6.22%	1.85%	1.62%	1.67%	1.65%	1.73%	1.64%

TABLE 5.12: The $J/\psi \to e^+e^-$ trigger efficiencies with the ECAL+SLT trigger as in Table 5.11, but with two targets used to calculate RoIs.

Selection cut	$J/\psi \to e^+e^-$ efficiency	Expected $J/\psi \to e^+e^-$ events	# of events
Trigger	0.94%	910	27861
Matching	0.89%	860	17613
Vertex $\chi^2 < 4$	0.84%	810	13631
$\|E/P - 1\| < 3\sigma$	0.74%	720	9940
Bremsstrahlung	0.12%	115	128
Fitted $J/\psi \to e^+e^-$	0.10%	95	$64 \pm 9 \pm 6$

TABLE 5.13: The $J/\psi \to e^+e^-$ efficiency after applying different selection cuts. The expected number of $J/\psi \to e^+e^-$ events in run 16665 compared to total amount of events surviving the different cuts is also shown. Trigger efficiency has been corrected for the 45% deadtime compared to the value listed in Table 5.11 (Column 4).

Comparison to data

The trigger efficiency, extracted from Monte Carlo, is used to predict the amount of $J/\psi \to e^+e^-$ we should observe in run 16665. The $J/\psi \to e^+e^-$ selection is identical to the one used for run 15230 and the efficiencies and event rates are given in Table 5.13. The amount of expected $J/\psi \to e^+e^-$ events is almost a factor three lower than in run 15230, not only because of the lower trigger efficiency, but also a higher deadtime (45%) and shorter run time (35%). Figure 5.13 shows the resulting invariant mass spectrum and a fit to the spectrum. The uncertainty of the fit is large due to the low amount of statistics, which might partly explain why the observed value is 35% below the expected. However several other runs in the same period as run 16665 and other $J/\psi \to e^+e^-$ selection procedures have been tried and all give more or less same J/ψ deficit. Part of the deficit can be explained by an overestimation of the ECAL coverage in the Monte Carlo for the runs in this period, as a visual comparison of the ECAL occupancies in Monte Carlo and data shows clear differences in several regions.

5.3.2 Tracking precision

The tracks in the OTR reconstructed by RefitX/RefitY are compared to the true tracks in Figure 5.14. The figure shows the residuals of the track position and slopes at the first layer in the OTR and the result of a fit with a Gaussian. The Gaussian describes the residuals reasonably well and the fitted width is used as the track resolution listed in Table 5.14. Table 5.14 also lists the track resolutions obtained with three different reconstruction programs. OTR-CATS gives the best results and therefore is the current offline reconstruction program. The offline

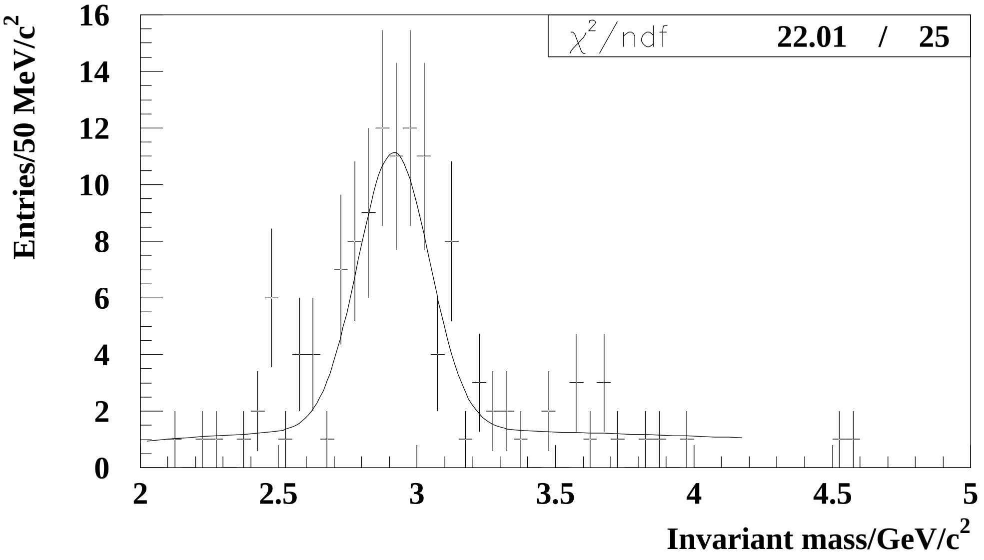

FIGURE 5.13: The invariant mass spectrum of e^+e^- pairs in run 16665 after requiring bremsstrahlung clusters associated to both the electron and positron. The full line shows the result of the fit with a Gaussian and a polynomial.

reconstruction is applied to any track produced in the primary interaction. Hence it includes many soft tracks and tracks not fully contained in the OTR, which intrinsically have worse tracking resolution than electrons from J/ψ decay. Even so, the Refit tracks have significantly better resolution than the offline reconstruction, particularly in y coordinates. Refit has the advantage over the offline reconstruction packages that it takes into account the target position and the ECAL cluster as measurements in the track. For y coordinates this completely dominates the resolution — the t_y resolution without any OTR hit is 0.31 mrad, consistent with the 0.29 mrad seen with the hits included. The offline reconstruction only gets better results for the track position in x, which is mainly determined by the hit resolution in the first superlayer. The offline reconstruction uses drift time resulting in an improved resolution over RefitX.

The Refit algorithm does not correct the covariance matrices for multiple scattering and therefore underestimates the errors on the track parameters. The pull distribution in Figure 5.14 shows that the errors are underestimated by up to a factor two. Figure 5.11 however shows that once the tracks have been propagated to the VDS area almost all of them are well located within the predicted area.

The L2Sili track resolution is measured in the same way as in Section 5.2.2 and Table 5.15 shows the results. The electrons triggered by the ECAL+SLT trigger have in average lower momentum than electrons triggered by the ECAL+L2Sili trigger, given the lower E_T threshold. The high momentum electrons are also mainly produced at low angles, outside the OTR coverage. The lower momentum tracks will have higher multiple scattering in the VDS planes and therefore worse track resolution. This is the reason the track resolution is about 25% worse than for the ECAL+L2Sili trigger (Table 5.5).

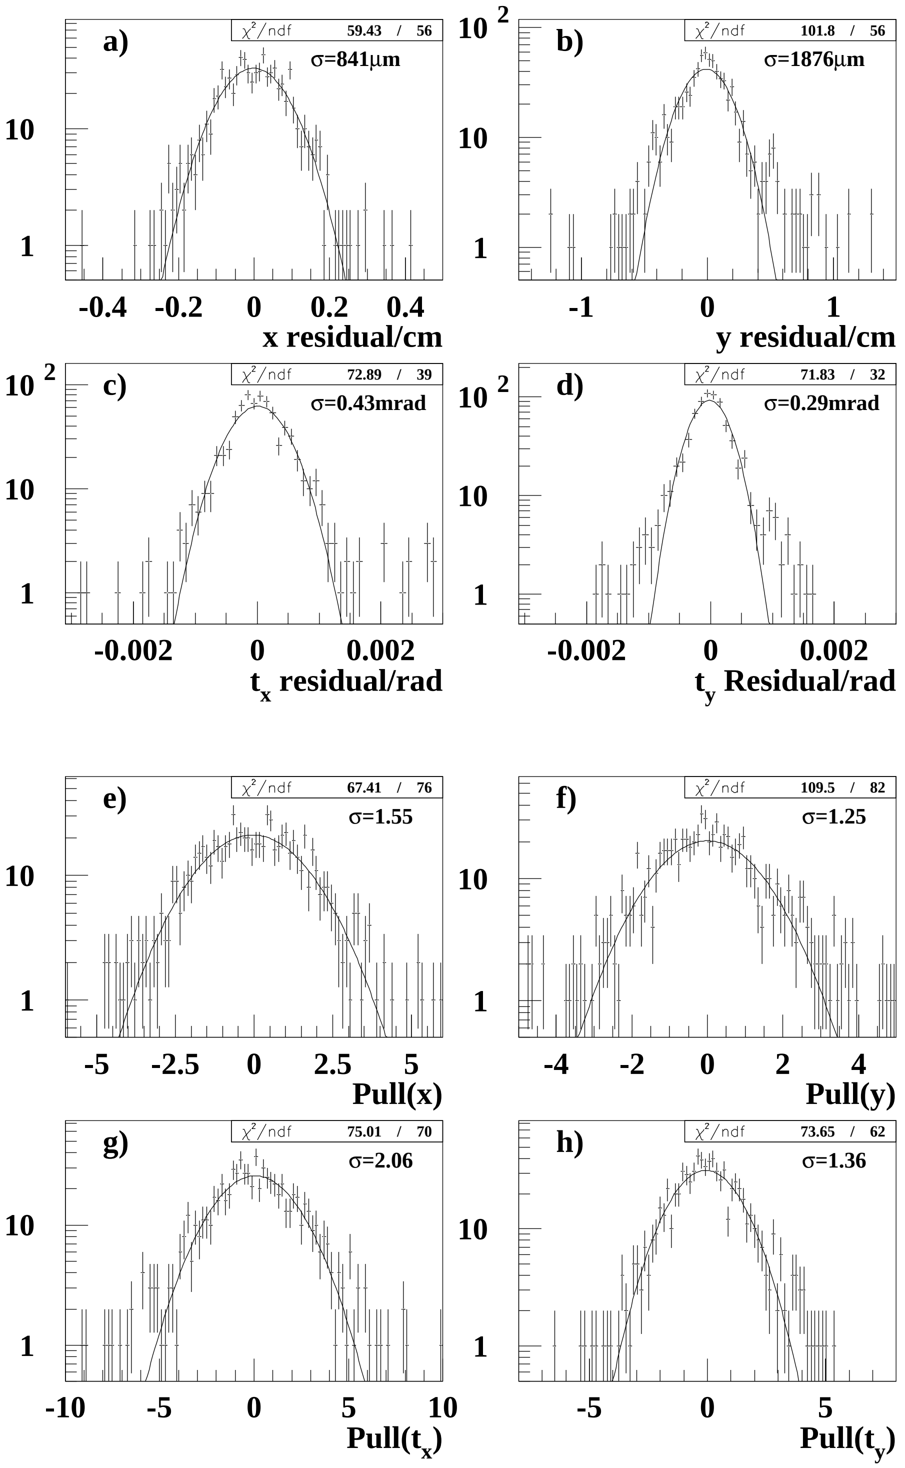

FIGURE 5.14: a-d) RefitX/RefitY track residuals for slopes and track position at the first OTR plane after the magnet. The tracks are seeded with ECAL clusters. The full line shows the result of Gaussian fit. e-h) Pulls for the residuals shown in a-d), with Gaussian fits shown in full line.

	Refit	OTR-CATS	RANGER	TEMA
$\Delta x/\mu$m	840	246	322	291
$\Delta y/$mm	1.9	3.7	5.0	4.1
$\Delta t_x/$mrad	0.43	0.62	0.71	0.76
$\Delta t_y/$mrad	0.29	4.73	6.96	5.39
Pull(x)	1.55	1.59	1.37	1.45
Pull(y)	1.25	1.52	1.25	1.81
Pull(t_x)	2.06	1.16	1.25	1.18
Pull(t_y)	1.36	1.53	1.39	1.92

TABLE 5.14: Width of residuals and pulls of Refit track parameters for the ECAL+SLT trigger. The other columns show the tracking resolution measured in different offline reconstructions programs [55].

	L2Sili	CATS
$\Delta x/\mu$m	14.0	10.3
$\Delta y/\mu$m	12.6	10.5
$\Delta t_x/$mrad	0.12	0.08
$\Delta t_y/$mrad	0.10	0.08
Pull(x)	0.48	1.08
Pull(y)	0.43	1.12
Pull(t_x)	1.02	1.34
Pull(t_y)	0.93	1.34

TABLE 5.15: Width of residuals and pulls of L2Sili track parameters for the ECAL+SLT trigger. The second column shows the equivalent results from the offline tracking.

5.3.3 Trigger rates

Table 5.16 shows the average rates of triggered events in run 16665 at different steps in the SLT algorithm. The rates listed are not the full event rates seen by the SLT, as both the ECAL and MUON pretriggers were used in this run. The rate of the MUON pretriggered events[9] into the SLT was much higher than the ECAL pretriggered event rate. This lead to a significant deadtime (45%) in the DAQ system due to overflow of the readout-FIFO, see Section 3.7.4. Besides this, the trigger/DAQ system was able to handle the rate. Table 5.17 shows the rejections predicted in the electron channel from the simulation of 600,000 bunch crossings.

The rate from the ECAL pretrigger is almost three times higher than in run 15230, mainly because of the lower E_T threshold, but also because the ECAL coverage was increased in the time between the two runs were recorded. The invariant mass cut rejects more events in run 16665 as the lower E_T cut lowers the average invariant mass in the background events. The rejection from E_T cut in CARE is almost unchanged compared to run 15230, but while the rejection in run 15230 was mainly due to events triggered by faulty ECAL pretrigger boards, the rejection in run 16665 was dominated by differences in the ECAL pretrigger and CARE cluster reconstructions.

[9] See Section 5.4.3.

Trigger step	Input rate/Hz	Output rate/Hz	Reduction factor
Pretrigger	$10.4 \cdot 10^6$	6200	1680
DAQ deadtime	6200	3400	1.8
Cluster $E_T > 1.0$ GeV	3400	2230	1.5
Mass> 2.0	2230	633	3.5
Slicer	633	126	5.0
RefitX	126	35	3.6
RefitY	35	8.4	4.2
L2Sili	8.4	1.5	5.6
4LT	1.5	1.5	1

TABLE 5.16: Trigger rates of ECAL pretriggered events measured for run 16665.

Trigger step	Events (16665)	Red. factor	Events (ideal)	Red. factor
Pretrigger	332	1800	2133	281
Cluster $E_T > 1.0$ GeV	245	1.4	930	2.3
Mass> 2.0	85	2.9	353	2.6
Slicer	3	28.3	13	27
RefitX	2	—	4	~ 3
RefitY	1	—	1	—
L2Sili	1	—	0	—

TABLE 5.17: The number of inelastic Monte Carlo events out of an initial 600,000 which survive the different steps in the ECAL+SLT trigger. Columns two and three show the amount of events and the reduction factor for data simulated with run 16665 detector performance, while columns four and five correspond to an ideal detector.

The SLT tracking algorithms reduce the rate by more than a factor 400, with more or less evenly distributed contributions from the different algorithms. The reduction is about nine times higher than for the ECAL+L2Sili trigger, partially because the OTR does not cover the inner region, but mainly because of a much better rejection of fake tracks. The high rejection makes the output rate to the 4LT low enough, that there are no problems with the reconstruction or data logging. For the full rates into the 4LT including the MUON pretriggered events, see Section 5.4.3.

The simulation does not have enough statistics to give rate estimates for the full algorithm. The rejection by Slicer is however overestimated substantially by the simulation. This could be caused by non-simulated hot channels.

5.3.4 Timing

The time spent in the different steps of the SLT trigger algorithm is listed in Table 5.18. The trigger steps have been divided into algorithmic, hit requesting and hit preparation parts, as we did in Table 5.8 for the ECAL+L2Sili trigger. The average time per event is still completely dominated by the ECAL readout and cluster reconstruction hiding the fact that hit request/preparation for the both the OTR and the VDS are rather slow and take as much time

Trigger step	Time per inv. (390 MHz)	Time per event (390 MHz)
ECAL readout	4.6 ms	4.6 ms
Cluster reconstruction	12 ms	12 ms
0^0 hit request	2.3 ms	0.4 ms
0^0 hit prep.	1.0 ms	0.2 ms
Slicer	0.5 ms	0.1 ms
RefitX	1.7 ms	0.06 ms
Stereo view hit req.	2.4 ms	0.02 ms
Stereo view hit prep.	0.6 ms	0.01 ms
RefitY	2.5 ms	0.03 ms
VDS hit req.	0.6 ms	$<$ 0.01 ms
VDS hit prep.	8.3 ms	0.02 ms
L2Sili algorithm	10.6 ms	0.03 ms
Event building	220 ms	0.1 ms
Total	—	17 ms

TABLE 5.18: Timing of the trigger steps in the ECAL+SLT trigger. The tables shows both the time needed for each invocation of a trigger step and the time spent per event after weighting with the reduction factors from Table 5.6. The timings are given for an "average CPU" of the full SLT farm.

as the tracking algorithms themselves. L2Sili is slower for this trigger configuration, because the candidates from photon clusters which were a major part of the candidates in ECAL+L2Sili trigger have already been rejected by the OTR tracking algorithms.

The total time spent per event (17 ms) would have allowed us to run at up to 14 kHz input rate, if the farm had not also been used to trigger on muons simultaneously.

5.3.5 Fake tracks

As the ECAL+SLT trigger follows the electron candidates through the two tracking systems, the amount of fake tracks is significantly reduced. Fake matches are less likely because the extrapolation distance from ECAL cluster to the OTR is quite short and the RoIs in the VDS are significantly smaller. The latter also reduces the amount of ghost tracks.

The amount of fake tracks is estimated in the same way as for the ECAL+L2Sili trigger, with the help of event mixing and by comparing the SLT tracks to tracks reconstructed by the offline reconstruction.

Event mixing

The event mixing becomes more complicated with the introduction of the OTR, as the fake matching can now happen between several detector systems. We will for the purpose of event mixing consider the OTR tracking algorithms, Slicer and the two Refits, as a single algorithm. The event mixing is done as for the ECAL+L2Sili trigger by replacing either the hits in just one of the RoIs by the hits from the previous event or by replacing the hits in both RoIs. Table 5.19 shows the results of the event mixing for track pairs. Three different results are shown:

	Unmixed	Mixed e^-	Mixed e^+	Mixed e^+e^-	"Real tracks"
OTR only	1.10 ± 0.06	0.14 ± 0.02	0.12 ± 0.02	0.02 ± 0.01	0.87 ± 0.07
OTR+VDS	0.22 ± 0.03	0.02 ± 0.01	0.02 ± 0.01	0.00 ± 0.00	0.20 ± 0.03
Only VDS mixed	0.22 ± 0.03	0.01 ± 0.01	0.01 ± 0.01	0.00 ± 0.00	0.18 ± 0.03

TABLE 5.19: Probabilities in percent to accept a track pair after applying event mixing and the derived fraction of real tracks. The definition of the different rows and the results are explained in the text.

- **OTR only:** The event mixing is done only for OTR hits and the track pair trigger probabilities are calculated after Slicer and Refit. This gives the rate of fake matches between the OTR tracks and the ECAL. About 20% of the events triggered by Slicer+Refit are seen to be fake matches.
- **OTR+VDS:** The event mixing is done for both OTR and VDS hits. From this we can calculate the rate of fake matches between tracks found in both detectors and ECAL clusters. The addition of the VDS reduces the amount of fake matches between ECAL and the tracking systems to $\frac{0.22\%-0.20\%}{0.22\%} \approx 10\%$ of the accepted tracks pairs.
- **Only VDS mixed:** We can also have a track in the OTR from the particle which deposited the energy in the ECAL, but the track is matched with a different track in the VDS. This can for instance occur if a photon converts into an e^+e^- pair inside one of the magnet chamber. The probability of such a fake match is estimated by doing the event mixing with just the VDS hits and using the OTR hits from the same event as the ECAL cluster. We see that even with the smaller RoIs we have $\frac{0.22\%-0.18\%}{0.22\%} \approx 20\%$ of triggered events with at least one fake matched track, to be compared with the 55% for the ECAL+L2Sili trigger configuration.

We have to count both types of fake matches — between ECAL and the OTR+VDS and between OTR and VDS tracks — in the total fake rate. There will be some correlation between the two types, but the statistics from which the two numbers are extracted is already limited to a few events. We will simply sum the two fake matching fractions and use $10\% + 20\% = 30\%$ as an upper limit on the amount of events which contain one or two wrongly matched tracks.

Track confirmation

The tracks found by L2Sili are compared to the CATS tracks in Figure 5.15 using the same χ^2 definition as in Section 5.2.5. The amount of ghost tracks is greatly reduced by the inclusion of the OTR in the trigger configuration. Table 5.20 shows the estimated ghost rate using the same χ^2 cut to define a ghost as in Section 5.2.5. Only 17% of the tracks are not found by CATS versus 57% in the ECAL+L2Sili trigger. These 17% ghost tracks, 75% are due to hits from real tracks located just outside the RoI (65% have a good y component, 35% a good x component) and only 25% are real ghost tracks.

The tracks reconstructed in the OTR by Refit are compared to tracks found by the offline reconstruction OTR-CATS [55] using a χ^2 definition similar to the L2Sili case described above. The parameters used in the χ^2 are the track slope and position at the first OTR layer after the magnet. The track uncertainties used in the χ^2 calculation are not given by the track covariance matrices. Instead fixed values are used for all track uncertainties. These values are

	Single track	Pair
Ghost in VDS	17%	30%
Ghost in OTR	14%	26%
Ghost in VDS or OTR	26%	43%

TABLE 5.20: The fraction of tracks and track pairs which from comparison with offline tracking is estimated to contain ghost tracks.

found by Gaussian fits to the difference of track parameters as estimated by Refit and OTR-CATS. Figure 5.16 shows χ^2 distributions separated in x and y components and the total χ^2. The ghost tracks are no so well separated from real tracks compared to the L2Sili tracks, but a small fraction of the Refit tracks are clearly not found by offline tracking. Defining tracks with $\chi^2 > 100$ to be ghost tracks we estimate that 14% of the Refit tracks are ghosts.

The amount of pairs with at least one ghost track in the OTR, the VDS or in either detector are shown in Table 5.20. From $J/\psi \to \mu^+\mu^-$ decays (Section 5.4.1) we find that approximately 10% of the lepton pairs triggered by the SLT are not reconstructed by OTR-CATS. The real amount of ghost tracks in the OTR is therefore lower than the value in the table. It is however not possible to simply correct by the 10%, as the hit occupancy in ordinary triggered events is higher than in J/ψ events and this affects the tracking efficiency of OTR-CATS. The 43% is therefore an upper limit of events with ghost tracks and is a large improvement over the 80% from the ECAL+L2Sili trigger.

5.3.6 Summary of the ECAL+SLT trigger performance

The introduction of the OTR tracking algorithms in the SLT increases the rejection in the SLT by almost a factor ten. This allows us to lower the E_T threshold applied in the ECAL pretrigger to 1.0 GeV. The total $J/\psi \to e^+e^-$ efficiency is however lower because of inefficiencies in the OTR and the loss over acceptance at small angles. The amount of ghost tracks has also been reduced by about a factor two. The algorithm is fast enough and has enough rejection to allow input rates of up to 14 kHz.

5.4 MUON+SLT trigger

The MUON+SLT trigger configuration triggers on $J/\psi \to \mu^+\mu^-$ events. This provides a much more pure sample of J/ψ events than the $J/\psi \to e^+e^-$ events since the muons do not have problems with bremsstrahlung emission and muons in general are easier to discriminate from background. The performance of the MUON+SLT trigger is evaluated in the same way as for the ECAL+SLT trigger.

5.4.1 Trigger Efficiency

The $J/\psi \to \mu^+\mu^-$ trigger efficiency is evaluated using a sample of 40,000 pure $J/\psi \to \mu^+\mu^-$ Monte Carlo events and 20,000 $J/\psi \to \mu^+\mu^-$ events, which have an variable amount of overlapping inelastic events. Below the efficiency factors related to the MUON tracking are described.

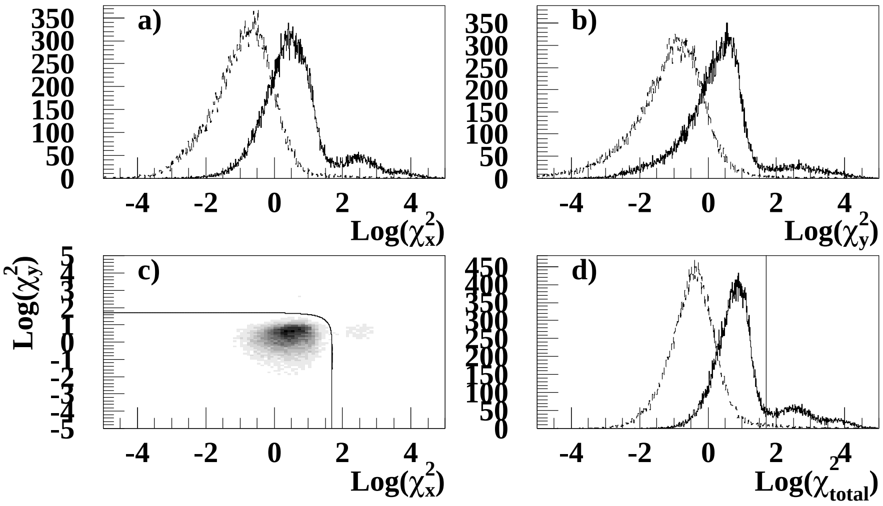

FIGURE 5.15: $\mathrm{Log}(\chi^2)$ distributions from matching ECAL pretrigger seeded L2Sili tracks with offline reconstructed CATS tracks. a) χ^2 in x view. b) χ^2 in y view. c) χ^2 in y view versus. χ^2 in x view. d) Total χ^2. The dashed lines in a), b) and d) are the results for $J/\psi \to e^+e^-$ Monte Carlo. The data has a higher average χ^2 than the Monte Carlo, due to worse alignment.

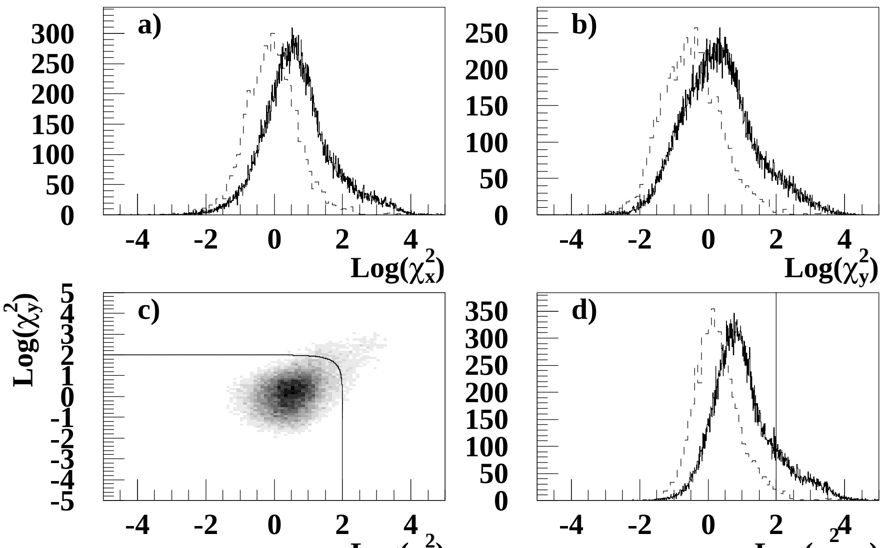

FIGURE 5.16: χ^2 distribution from matching ECAL pretrigger seeded Refit tracks with offline reconstructed OTR tracks. a) χ^2 in x view. b) χ^2 in y view. c) χ^2 in y view versus. χ^2 in x view. d) Total χ^2. The dashed lines in a), b) and d) are the results for $J/\psi \to e^+e^-$ Monte Carlo.

The SLT contributions unchanged compared to the ECAL+SLT trigger are not explicitly described here, but can be found in Section 5.3.1.

Geometrical acceptance: The geometrical acceptance for $J/\psi \to \mu^+\mu^-$ is much simpler to define than for $J/\psi \to e^+e^-$. We require each muon to pass through two superlayers in the VDS and five superlayers in the OTR as was done for electrons. Furthermore the muons have to cross the tube layers MU01 and MU02 and the pad layers MU03 and MU04.

Muon pad efficiency and pad coverage: The efficiency for having the required hits in MU03 and MU04 is separated into two factors: MUON pad efficiency and MUON pad coverage. The MUON pad efficiency is the hit efficiency for a working MUON pad. This was described in Section 5.1 and the expected efficiency is 52% for events triggered on two muons. The MUON pad coverage gives the efficiency after removing hits from the channels with higher readout thresholds or any other problem during data taking.

Muon pretrigger efficiency: The muon pretrigger had a few bad data links, which is reflected in the efficiency. Furthermore about 3% of the events are not triggered because the hits in the muon pad layers do not lie in the required pattern.

Coincidence distance cut: The SLT has a cut of 50 cm between the two MUON pretrigger coincidences. It is seen not to remove any $J/\psi \to \mu^+\mu^-$ events.

Muon hits: We require a hit in the 0° layers of MU01 and MU02. The hits are required to be isolated and the hit in MU01 has to be a double hit, see Section 4.2.2. The double hit requirement in MU01 introduces a large inefficiency, which for an ideal detector accounts for 85% of the loss. The remaining 15% are mainly due to overestimates of the geometrical acceptance. For a realistic detector we loose another 20% events due to non working channels. The overlapping events also cause a slight decrease of the efficiency due to the isolation requirement on the hits.

Muon tracking: The MUON tracking requires that the hits in MU01, MU02 and the pretrigger coincidence lie on a straight line within certain limits. With the chosen limits we have an efficiency of about 95%. The MUON tracking only considers the first candidate found for each pretrigger coincidence, which introduces an efficiency loss of up to 10% depending on the amount of overlapping events.

Muon track pairing: When the tracks found in the MUON system are combined into pairs, the SLT requires that the x and y coordinates of the two tracks at MU01 differ by more than 30 cm. This removes clones, but also 18% of the signal events, where the tracks are far apart while having almost identical x or y coordinates. In the realistic scenario the loss is slightly compensated by the missing coverage in the MUON pads.

OTR RoIs: The RoIs for Slicer and the Refits are much worse for the MUON+SLT trigger configuration than for the ECAL+SLT. This is because the y coordinates are calculated from muon pretrigger position, which has a resolution of about 10 cm. The x coordinates are calculated by assuming the muons all have $p_T = 1.6$ GeV/c. We therefore loose muons which have

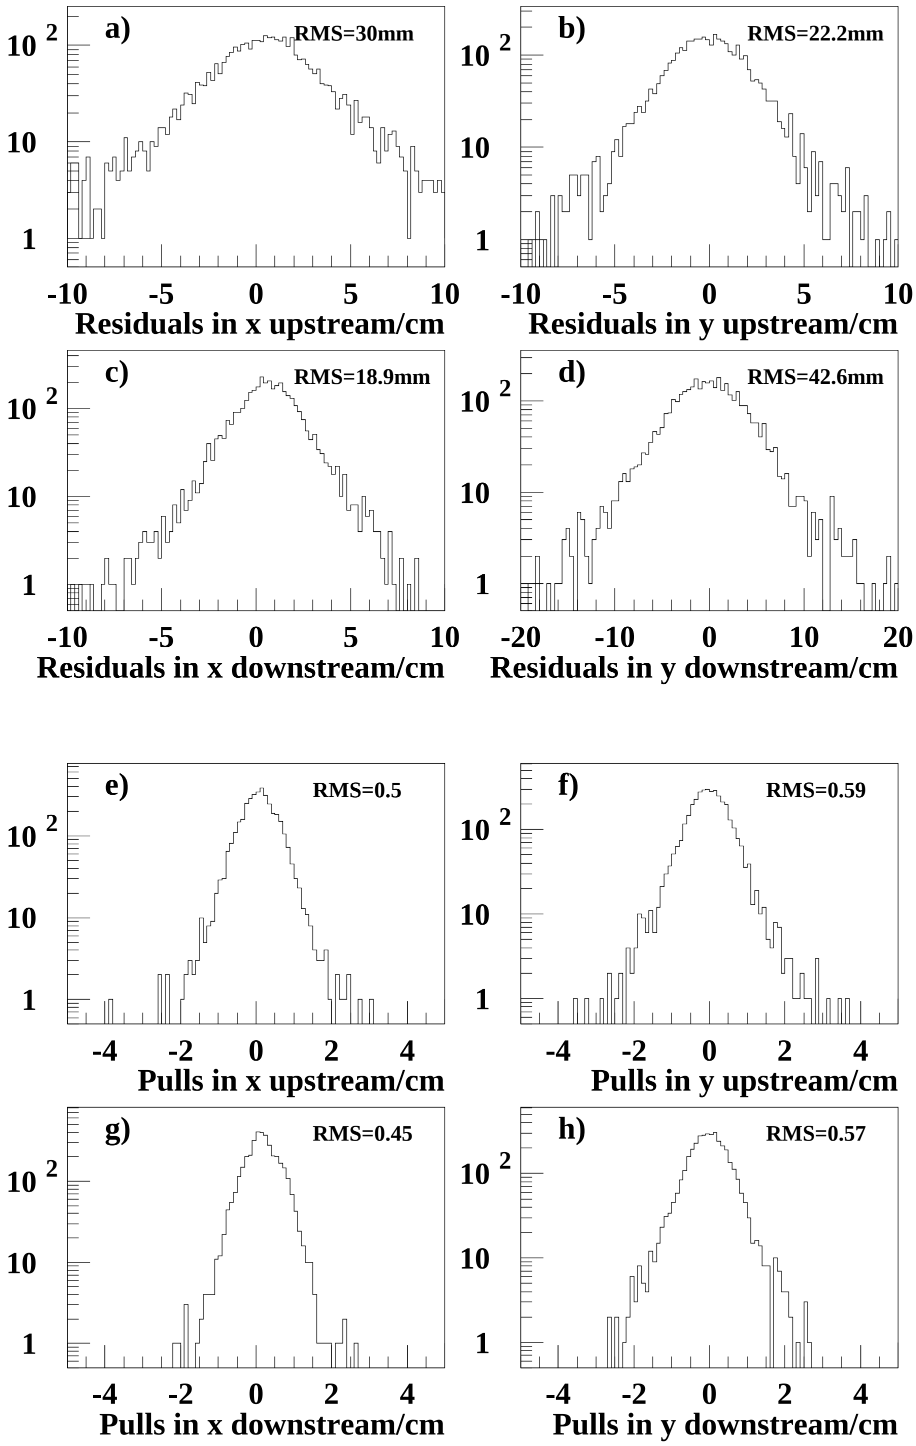

FIGURE 5.17: a-b) The difference between the predicted Slicer/Refit RoI and the true position in x and y right after the magnet ($z = 700$ cm). c-d) The same as a-b), but in front of the ECAL ($z = 1350$ cm). e-h) The residuals of a-d) divided by the width of the RoIs.

a p_T which deviates a lot from the p_T window ($0.7\,\mathrm{GeV}/c < p_T < 2.5\,\mathrm{GeV}/c$). The resulting precision of the RoIs is between 2.5 and 8 times worse than for the ECAL+SLT trigger as can be seen from the RoI residuals in Figure 5.17 compared to Figure 5.11. The figure also shows the "pulls" of the RoIs, where we can see that the RoI size corresponds on average to about 2σ. This is reflected in the 85% of the events, where both RoIs are correctly predicted.

L2Sili RoIs: The RoIs for L2Sili have a precision very similar to the ECAL+SLT case, as shown in Figure 5.18. The only exception is the position in y close to the magnet which is a factor two worse. The Refit track in y is mainly determined from RoI itself and not from the OTR hits and the muon pad has much worse resolution than the ECAL cluster. This is reflected in the worse y resolution for the L2Sili RoIs. The RoI are sized accordingly and no events are lost because of it. With increasing occupancy more wrong hits are added to the Refit tracks, resulting in a wrong prediction of the RoIs for a few percent of the events, which consequently are lost.

Monte Carlo results

The trigger efficiency factors are listed in Table 5.21 for a single target wire and in Table 5.22 for two target wires. The same sample numbering is used as for the ECAL+L2Sili in Section 5.2.1. Table 5.23 shows the total efficiency separated into the three detector components used in the trigger: MUON, OTR and VDS. The majority of the events are lost in the MUON system, because of inefficiencies in the pad system and the requirement on the hits in MU01. The hit requirement is a part of the trigger algorithm and will be improved significantly for the 2002 run. A substantial amount of events are also lost during the tracking through the OTR, because of the imprecise RoIs from the MUON and OTR performance . Within the statistical precision of the analyzed sample the efficiency does not change when using two target wires.

We also have a strong dependence on the detector occupancy. Figure 5.19 shows the efficiency for the tracking algorithms as a function of the average amount of overlapping events. The efficiencies have been normalized with respect to the tracking efficiency in event with just the two muons in the complete detector (Sample 2). The MUON tracking efficiency decreases with increasing occupancy because of the hit isolation cut and because only the first muon track candidate for each pretrigger coincidence is used. For the OTR tracking the inefficiencies comes from picking up wrong hits/tracks in the large RoIs given by the MUON tracking, while the VDS tracking efficiency is almost independent of the hit occupancy.

Comparison with data

The MUON-SLT trigger has a much higher suppression of background events in the J/ψ region than the ECAL+SLT trigger. $J/\psi \to \mu^+\mu^-$ events also do not suffer from bremsstrahlung tails and we are therefore able to observe the $J/\psi \to \mu^+\mu^-$ events based on the SLT tracks alone. The $J/\psi \to \mu^+\mu^-$ peak in the invariant di-muon mass spectrum is clearly seen in run 16665 as shown in Figure 5.20. The invariant mass spectrum is dominated by muons from pion decays and falls off exponentially for masses above 2.2 GeV/c^2. Below 2.2 GeV/c the mass spectrum falls off toward lower mass due to the p_T cutoff. The spectrum shows a small peak at the mass of ω meson. It has been confirmed to be real $\omega \to \mu^+\mu^-$ decays by using complete di-muon triggered data sample. With the full invariant mass spectrum one also observe $\phi \to \mu^+\mu^-$ and

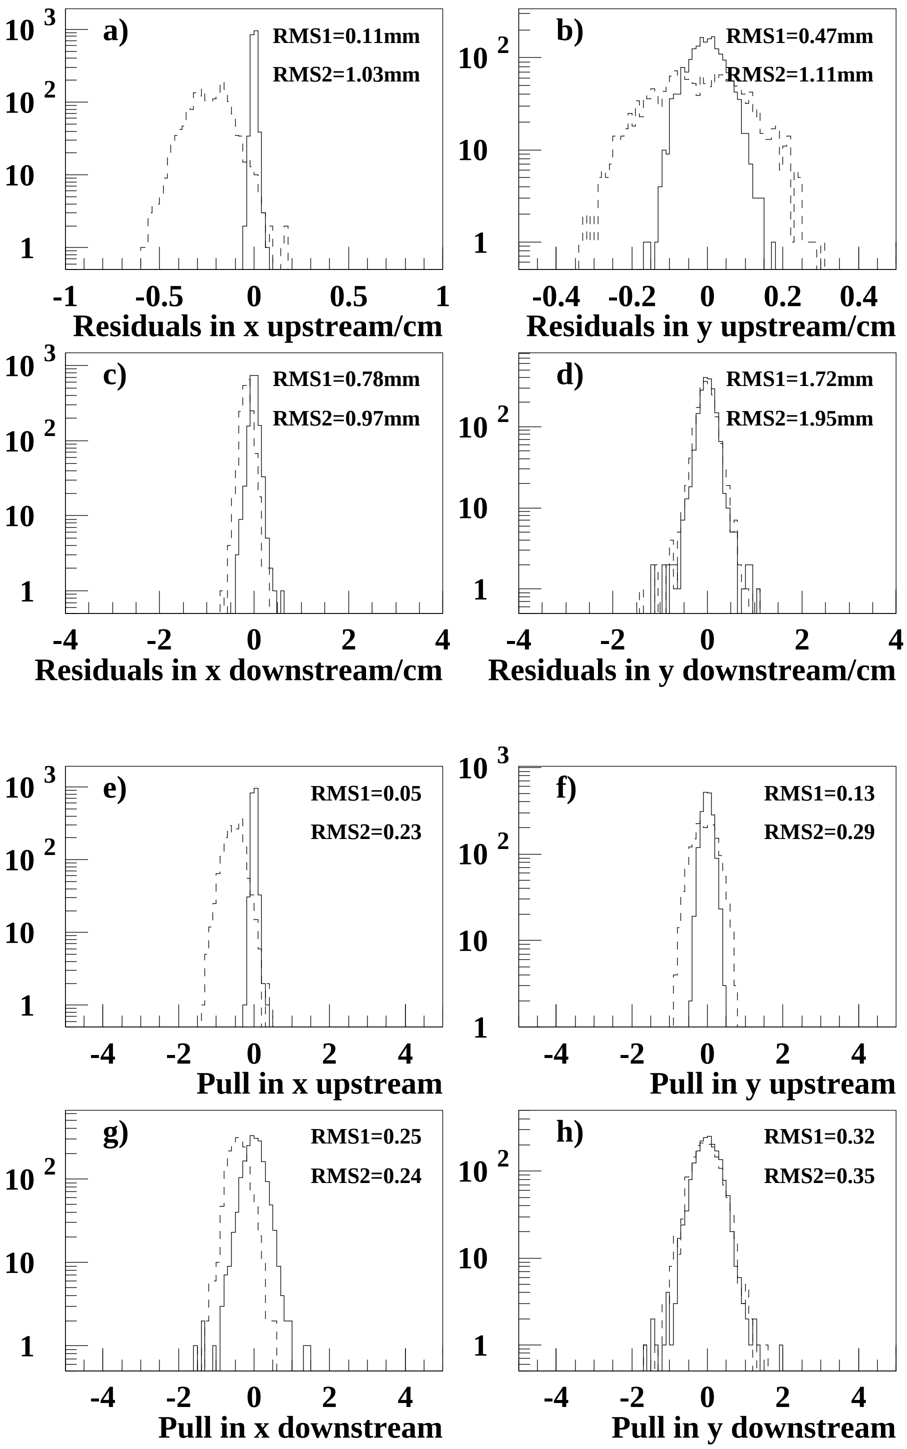

FIGURE 5.18: Residuals and normalized residuals of the L2Sili RoIs for the MUON+SLT trigger configuration. See Figure 5.2 for a comparison with the ECAL+SLT case.

Sample	1	2	3	4	5	6	7
Geo. acceptance	30.8%	30.8%	30.8%	30.8%	30.8%	30.8%	30.8%
Pad efficiency	52.2%	52.2%	52.2%	52.2%	52.2%	52.2%	52.2%
Pad coverage	96.5%	50.0%	50.1%	50.1%	50.1%	50.1%	50.1%
Pretrigger	97.2%	95.9%	95.3%	95.4%	95.4%	95.5%	95.5%
Distance cut	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Muon hits	44.4%	34.9%	33.5%	33.1%	32.8%	32.8%	32.1%
Muon tracking	95.6%	95.0%	92.0%	90.6%	89.0%	85.8%	87.7%
Muon pairing	81.9%	86.8%	86.6%	86.5%	87.1%	87.1%	87.1%
OTR RoIs	85.8%	86.5%	86.0%	86.1%	85.9%	86.0%	86.1%
Slicer hits	89.7%	80.6%	86.5%	87.5%	87.7%	88.0%	90.2%
Slicer algorithm	99.5%	88.3%	90.7%	92.7%	93.2%	94.7%	94.5%
RefitX hits	99.5%	90.5%	86.4%	87.1%	84.7%	85.6%	84.3%
RefitX algorithm	99.9%	99.7%	95.9%	92.2%	92.2%	89.7%	88.3%
RefitY hits	99.6%	94.1%	92.0%	90.9%	87.1%	86.9%	87.1%
RefitY algorithm	99.1%	96.8%	87.9%	85.2%	87.4%	81.6%	83.6%
L2Sili RoIs	99.9%	100.0%	99.1%	98.1%	98.7%	97.4%	97.8%
VDS hits	97.3%	94.5%	93.8%	94.9%	96.0%	94.3%	93.9%
L2Sili algorithm	99.5%	98.9%	97.8%	98.0%	96.6%	95.6%	97.2%
Total efficiency	3.82%	1.06%	0.84%	0.79%	0.74%	0.64%	0.67%

TABLE 5.21: The $J/\psi \to \mu^+\mu^-$ trigger efficiency with the MUON+SLT trigger separated into individual factors. The statistical error on the total efficiency is about 5%, but the numbers are correlated as the J/ψ kinematics is the same in the samples 1 and 2 and in 3-7. See the text for the definition of the different Monte Carlo samples.

Sample	1	2	3	4	5	6	7
L2Sili RoIs	96.4%	97.2%	96.8%	95.6%	95.7%	95.1%	97.1%
VDS hits	97.4%	95.8%	94.8%	95.4%	96.2%	94.1%	94.0%
L2Sili algorithm	99.5%	99.2%	98.1%	96.9%	96.1%	95.0%	98.0%
Total efficiency	3.74%	1.07%	0.83%	0.76%	0.72%	0.62%	0.67%

TABLE 5.22: The $J/\psi \to \mu^+\mu^-$ trigger efficiency with the MUON+SLT trigger as in Table 5.21, but with two targets used to calculate RoIs.

Sample	1	2	3	4	5	6	7
Geo	30.8%	30.8%	30.8%	30.8%	30.8%	30.8%	30.8%
MUON	17.1%	7.2%	6.7%	6.5%	6.3%	6.1%	6.1%
OTR	75.1%	50.7%	45.3%	43.4%	41.7%	39.0%	39.7%
VDS	96.8%	93.5%	91.0%	91.3%	91.5%	87.8%	89.3%
accept	3.82%	1.06%	0.84%	0.79%	0.74%	0.64%	0.67%

TABLE 5.23: The $J/\psi \to \mu^+\mu^-$ trigger efficiency with the MUON+SLT trigger separated into major components.

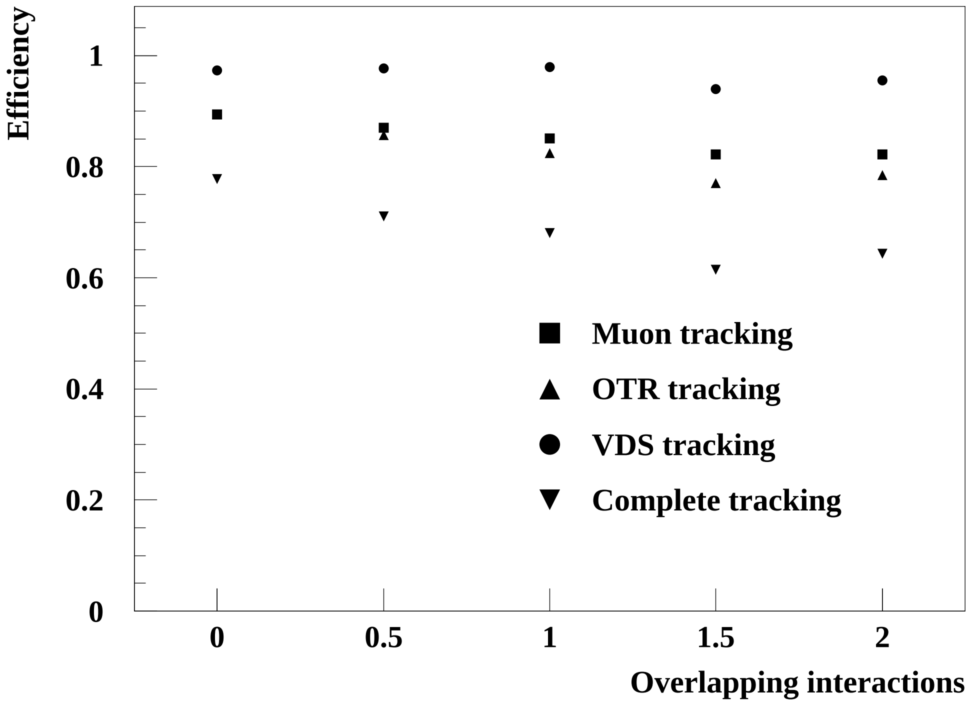

FIGURE 5.19: Tracking efficiencies in the different detectors as function of the mean number of overlapping interactions. The efficiencies have been normalized with respect to the measured efficiency with just two muons in the complete detector.

$\rho^0 \to \mu^+\mu^-$ decays. These mesons are produced much more copiously than J/ψ mesons, but have a much lower branching ratio to muon pairs ($O(10^{-5}$–$10^{-4})$) and only mesons produced with a high transverse momentum decay into muons that are accepted by the trigger.

The invariant mass spectrum around the J/ψ mass region has been fitted with a Gaussian for the signal and an exponential to describe the background. The fitted amount of $J/\psi \to \mu^+\mu^-$ events are compared in Table 5.24 with the prediction from the simulation after correcting for a down-scaling of four applied in FLT hardware and 45% deadtime in the DAQ system, see Section 5.4.3. The observed amount is almost a factor two higher than the expected one. Other runs have been investigated and the same mismatch between data and simulation is observed. The data available can not be used to determine the reason, but the most likely explanation is the masking in the MUON system which is very strict in the simulation. The uncertainty on the efficiency of working pads is also large.

Figure 5.20 shows the invariant spectrum after matching the SLT tracks with offline tracks in the VDS and the OTR. This reduces the background by a factor three, while the signal is only reduced by about 10%. The 10% loss comes from the matching of Refit tracks with OTR-CATS tracks, while there is no loss observed from the matching of tracks in the VDS. The observed amount of $J/\psi \to \mu^+\mu^-$ events still disagrees with the simulation by almost a factor two.

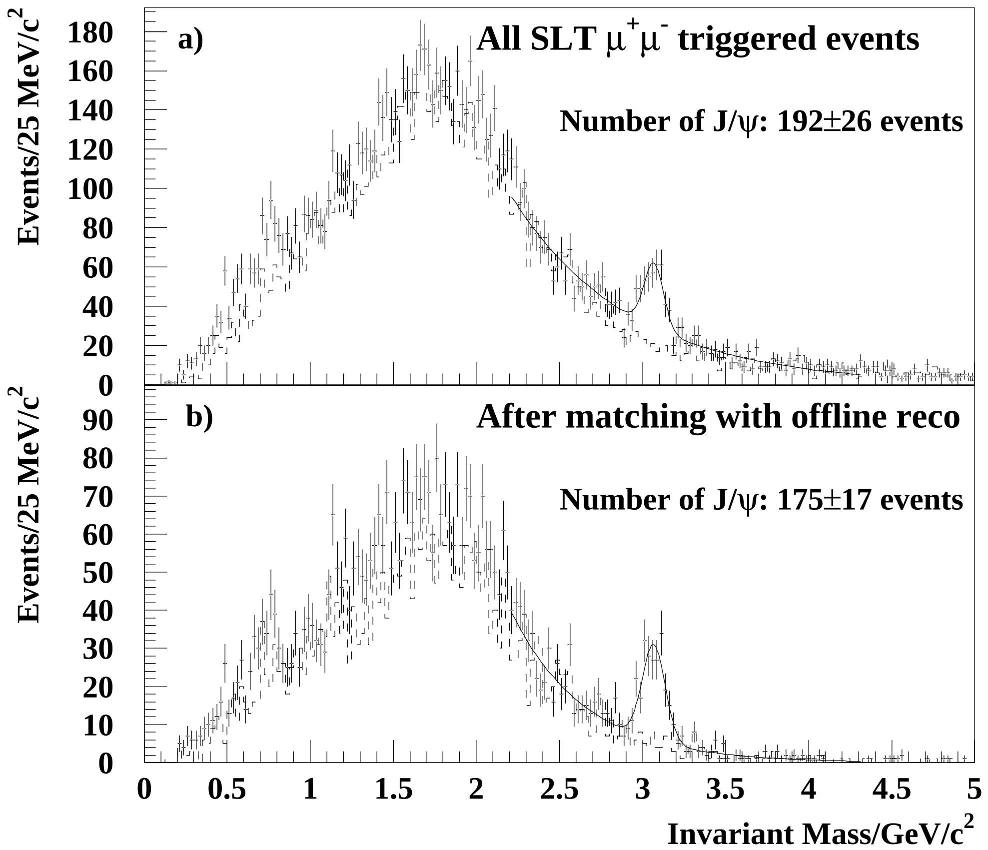

FIGURE 5.20: The invariant mass spectrum of $\mu^+\mu^-$ pairs in run 16665. The top spectrum shows the complete SLT triggered sample, while the lower spectrum only shows events, where both muons have been matched to offline VDS and OTR tracks. The dashed line shows the invariant mass spectrum for $\mu^+\mu^+$ and $\mu^-\mu^-$ events.

5.4.2 Tracking precision

The tracking precision of the Refit and L2Sili tracks is measured like it was done for the ECAL+SLT trigger in Section 5.3.2. The results are listed in Table 5.25 and 5.26. Compared to the ECAL+SLT trigger configuration, the precision of RefitY is almost three times worse for the MUON+SLT trigger configuration, but still far better than the resolution from the offline track reconstruction. The resolution is worse since it is completely dominated by the precision of the RoI seeds, which is significantly worse with MUON seeding.

The tracking resolution in L2Sili is also worse in the MUON+SLT case. This can be partially explained, because the triggered muons have 35% lower momentum than the triggered electrons ($p_T > 0.7$ GeV/c cut versus $E_T > 1.0$ GeV cut). The y resolution is also affected by the approximately twice larger RoIs, which allow for more wrong hits to be assigned to a track and thus worsen the resolution.

Selection cut	$J/\psi \to \mu^+\mu^-$ efficiency	Expected $J/\psi \to \mu^+\mu^-$ events	# of events	Fitted # of $J/\psi \to \mu^+\mu^-$
Trigger	0.10%	100	21147	—
Opposite charge	0.10%	100	11460	192 ± 26
Matching	0.09%	90	4411	175 ± 17

TABLE 5.24: The $J/\psi \to \mu^+\mu^-$ efficiency after applying different selection cuts. The trigger efficiency is taken from Table 5.23 (Sample 5, 0.74%) and has been corrected for the pretrigger down-scaling (factor four) and the DAQ deadtime (45%). Also shown is the expected number of $J/\psi \to \mu^+\mu^-$ events in run 16665, compared to the total amount of events surviving the different cuts.

	Refit	OTR-CATS	RANGER	TEMA
$\Delta x/\mu$m	720	246	322	291
Δy/mm	5.2	3.7	5.0	4.1
Δt_x/mrad	0.45	0.62	0.71	0.76
Δt_y/mrad	0.80	4.73	6.96	5.39
Pull(x)	1.24	1.59	1.37	1.45
Pull(y)	1.28	1.52	1.25	1.81
Pull(t_x)	1.57	1.16	1.25	1.18
Pull(t_y)	1.11	1.53	1.39	1.92

TABLE 5.25: Width of residuals and pulls of Refit track parameters for the MUON+SLT trigger. The other columns show the tracking resolution measured in different offline reconstructions programs[55].

5.4.3 Trigger rates

The trigger rates for MUON pretriggered events are shown in Table 5.27 for run 16665. The output rate from the MUON pretriggers, almost 100 kHz, is far too large for the SLT. We do not know the energy of the muons and can therefore not impose an E_T cut in the MUON pretrigger like we did for the ECAL pretrigger. Instead the MUON pretrigger rate is simply down-scaled by the FLT hardware (the TDU) by a factor four. Even with this down-scaling, the DAQ still has 45% deadtime due to overflow of the readout-FIFO and the event rate into the SLT is 13.2 kHz of MUON pretriggered events. In addition we have 3.4 kHz of ECAL pretriggered events. The 13.2 kHz is higher than the maximal rate we can read out the full MUON system via the SHARC switch, but the 50 cm distance cut between the MUON pretrigger coincidences reduces the rate by a factor two before the MUON system is read out. The readout therefore does not impose any limitation.

The SLT tracking in the MUON chambers gives a large suppression by eliminating events triggered on noise and on soft muons from meson decays. In the OTR tracking, Slicer has very little suppression because of the large RoIs and the majority of the rate reduction comes from RefitY.

The total reduction in the SLT is about 13000, a factor six higher than for the ECAL+SLT trigger. The total output rate in run 16665 was 7.7 Hz given by: 1.5 Hz ECAL+SLT trig-

	L2Sili	CATS
$\Delta x/\mu$m	16	10.4
$\Delta y/\mu$m	16	10.7
Δt_x/mrad	0.13	0.12
Δt_y/mrad	0.14	0.13
Pull(x)	0.57	1.10
Pull(y)	0.53	1.15
Pull(t_x)	1.08	1.34
Pull(t_y)	1.12	1.36

TABLE 5.26: Width of residuals and pulls of L2Sili track parameters for the MUON+SLT trigger. The last column shows the equivalent results from the offline tracking.

Trigger step	Input rate/Hz	Output rate/Hz	Reduction factor
Pretrigger	$10.4 \cdot 10^6$	96,000	108
Down-scaling	96,000	24,000	4
DAQ deadtime	24,000	13,200	1.8
Pretrigger distance cut	13,200	6,600	2.0
Muon tracking	6,600	550	11.9
Slicer	550	395	1.4
RefitX	395	100	3.9
RefitY	100	7.7	13
L2Sili	7.7	1.0	7.7
4LT	1.0	1.0	1.0

TABLE 5.27: Trigger rates measured for the run 16665.

gered events, 1.6 Hz ECAL+SLT bypass events[10], 1 Hz MUON+SLT triggered events, 1.4 Hz MUON+SLT bypass events and 2.2 Hz of random triggered events. Events recorded at this rate were easily reconstructed by the 4LT and logged to tape without any additional deadtime.

Table 5.28 shows the results of applying the MUON+SLT trigger simulation to 600,000 simulated inelastic events. The sample is not large enough to give rate predictions for the full trigger chain. The trigger simulation overestimates the reduction of the MUON pretrigger, but underestimates the reduction in the MUON tracking. This is most likely because the real MUON pretrigger also triggers on electronics noise and other effects not simulated. These events are rejected when hits in MU01 and MU02 are also required. The rejection by Slicer and RefitX is also overestimated by about factor two in the simulation.

5.4.4 Timing

The time spent in different parts of the trigger is listed in Table 5.29. The total time is dominated by the time used to read out the MUON system and the MUON tracking. The time spent in the rest of the SLT algorithm, in particular in the OTR, is higher per invocation than in the

[10]The SLT bypass events are 1 in 400 of the events which passed the seeding step of the SLT.

Trigger step	Events (16665)	Red. factor	Events (ideal)	Red. factor
Pretrigger	4082	147	6673	90
Pretrigger distance cut	1291	3.2	2420	2.75
Muon tracking	166	7.8	384	6.3
Slicer	51	3.3	153	2.5
RefitX	6	8	25	6
RefitY	0	—	4	~ 6
L2Sili	0	—	1	—

TABLE 5.28: The number of inelastic Monte Carlo events out of an initial 600,000 which survive the different steps in the MUON+SLT trigger. Column two and three show the amount of events and the reduction factor for data simulated with run 16665 detector performance, while column four and five corresponds to an ideal detector.

ECAL+SLT trigger, because of the larger RoIs. The large rejection in the MUON tracking lowers the load on the rest of the SLT algorithm, hence the timing is less critical here.

With an average processing time of just 2.9 ms per event, the full farm can in principle handle up to 80 kHz input rate. However the complete MUON system can only be read out at approximately 10 kHz and, because we have a factor two rejection from the 50 cm cut on the MUON pretrigger pairs, the maximally possible SLT input rate is around 20 kHz.

The MUON+SLT trigger configuration ran in parallel with the ECAL+SLT trigger during the 2000 run. The MUON+SLT trigger had four times higher input rate, but was six times faster than the ECAL+SLT trigger. The average time spent per event in run 16665 was therefore approximately 5.8 ms. With a total input rate of 16.6 kHz this meant that effectively only 100 SLT nodes out of 240 were busy during the running.

5.4.5 Fake tracks

The estimation of fake track rates follows closely the procedure used for the ECAL+SLT trigger, described in Sections 5.3.5. Since the RoIs given by the MUON seeding are larger than the RoIs from the ECAL seeding, the amount of fake tracks will increase .

Event mixing

The event mixing is done by mixing hits in the OTR and VDS as was done for the ECAL+SLT trigger. In principle one could also do the event mixing for the MUON tracking. This requires a data sample recorded with the MUON pretriggers, but without the MUON tracking. Unfortunately, during the 2000 run, no such data sample was recorded with enough statistics to study event mixing.

For simplicity we only consider $\mu^+\mu^-$ pairs in the event mixing. The result of the event mixing is given in Table 5.30. Compared to the ECAL+SLT trigger configuration, the amount of fake matches between the SLT seed (in this case the MUON system) and the OTR has increased. It has increased from 20% to $\frac{1.15\%-0.53\%}{1.15\%} \approx 55\%$, when measured before the tracking through the VDS. For tracks found by both Refit and L2Sili the ghost fraction has increased from 10% in ECAL+SLT trigger to $\frac{0.17\%-0.12\%}{0.17\%} \approx 30\%$ and for full tracks 30% of the events which contain at least one fake match. The fake matches between the OTR and the VDS have

Trigger step	Time per inv. (390 MHz)	Time per event (390 MHz)
Pretrigger readout	0.04 ms	0.04 ms
Muon readout	1.2 ms	0.6 ms
Muon tracking	3.0 ms	1.5 ms
0^0 hit request	3.6 ms	0.15 ms
0^0 hit prep.	4.0 ms	0.17 ms
Slicer	1.7 ms	0.07 ms
RefitX	8.4 ms	0.25 ms
Stereo view hit req.	3.3 ms	0.03 ms
Stereo view hit prep.	1.2 ms	0.01 ms
RefitY	3.4 ms	0.03 ms
VDS hit req.	0.7 ms	0.00 ms
VDS hit prep.	11.2 ms	0.01 ms
L2Sili algorithm	10.4 ms	0.01 ms
Event building	220 ms	0.02 ms
Total	—	2.9 ms

TABLE 5.29: Timing of the trigger steps in the MUON+SLT trigger. The tables shows both the time needed for each invocation of a trigger step and the time spent per event after weighting with the reduction factors from Table 5.27. The timings are given for an "average CPU" of the full SLT farm.

	Unmixed	Mixed μ^-	Mixed μ^+	Mixed $\mu^+\mu^-$	Real tracks
OTR only	1.15 ± 0.06	0.58 ± 0.05	0.56 ± 0.04	0.53 ± 0.04	0.53 ± 0.10
OTR+VDS	0.17 ± 0.02	0.04 ± 0.01	0.3 ± 0.01	0.01 ± 0.01	0.12 ± 0.03
Only VDS mixed	0.17 ± 0.02	0.03 ± 0.01	0.04 ± 0.01	0.01 ± 0.01	0.11 ± 0.03

TABLE 5.30: Probabilities in percent to accept a track pair after doing event mixing and the derived fraction of real tracks. The results in the different rows are explained in the text.

increased to 36% of the triggered events. Combining the two types of fake matches (assuming that they are uncorrelated) we estimate that 55% of the triggered $\mu^+\mu^-$ events contain at least one wrongly matched track.

Track confirmation

The Refit and L2Sili tracks are compared to the offline reconstructed tracks using the χ^2 definitions of Sections 5.2.5 and 5.3.5. The higher ghost fraction in the MUON+SLT trigger makes the separation between real tracks and ghost tracks in the OTR χ^2 more visible, as shown in Figure 5.21. The $\chi^2_{\mathrm{total}} < 100$ separation cut, also used for the ECAL+SLT trigger, is chosen based on this plot.

The estimated fractions of ghost tracks and events containing ghost tracks are listed in Table 5.31. The 63% events with ghost tracks listed in the table is an upper limit, since we have seen that the offline reconstruction is not able to reconstruct about 10% of the triggered $J/\psi \to \mu^+\mu^-$ events.

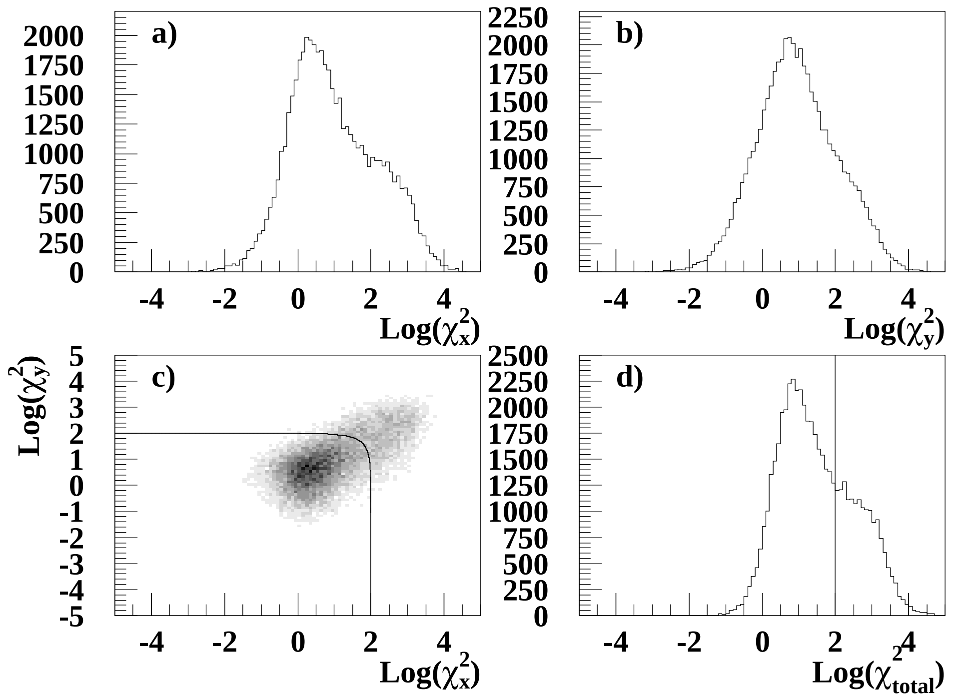

FIGURE 5.21: $\log(\chi^2)$ distributions from matching Muon pretrigger seeded Refit tracks with offline reconstructed OTR tracks. a) χ^2 in x view. b) χ^2 in y view. c) χ^2 in y view vs. χ^2 in x view. d) Total χ^2. The dashed line in a), b) and d) is the result for $J/\psi \to e^+e^-$ Monte Carlo.

5.4.6 Summary of the SLT+MUON trigger performance

The SLT+MUON trigger has been measured to be less efficient than the ECAL+SLT trigger. The inefficiencies comes mainly from detector inefficiencies, particularly in the MUON system, though the current simulation underestimates the actual detector performance. The trigger algorithm has some significant inefficiencies in the MUON and OTR tracking, the former due to the hit preparation and the latter because of large uncertainties in the RoIs given by the MUON seeding. Furthermore, the output rate from the pretrigger is far too high for the DAQ system to handle and had to be down-scaled in the 2000 run, thus reducing the total $J/\psi \to \mu^+\mu^-$ efficiency.

The SLT algorithm itself has a very large rejection and is fast enough to keep up with the rate. The only limitation comes from the readout of data from the MUON tube chambers. The output rate is only slightly lower than the rate from the ECAL+SLT trigger, but most of triggered events have low invariant mass and the J/ψ signal can be seen directly from the trigger tracks without further lepton identification from the offline reconstruction.

	Single track	Pair
Ghost VDS	33%	52%
Ghost OTR	30%	48%
Ghost VDS or OTR	43%	63%

TABLE 5.31: The fraction of tracks and track pairs containing a ghost track compared to the offline tracking.

5.5 Summary of the 2000 SLT trigger performance

The performance of the three main trigger configurations used in the 2000 run has been evaluated. All of them have efficiencies of 1–2%. Most of the efficiency loss is coming from the reduced detector performance and the missing coverage from the inner detectors (ITR and MUON pixel chambers). Because of the missing FLT, the trigger rates from the pretriggers also have to be kept low, causing large inefficiencies.

The MUON+SLT trigger configuration has some major inefficiencies in the SLT algorithm itself, mainly in the MUON hit preparation and in the seeding of the OTR tracking algorithms. It is also significantly dependent on the detector occupancy due to very large RoIs in the MUON and OTR tracking algorithms. For the configurations with OTR tracking, the SLT provides sufficient rejection and is fast enough. The rate limitation is given by the DAQ system and its ability to transfer the information from the detectors used in the SLT seeding at high rate to the SLT. Up to half of the events triggered by the SLT are seen to contain fake tracks, which can potentially be rejected by improvements in the SLT algorithm.

Chapter 6

The SLT algorithm for the 2002 run

The HERA accelerator was shut down for a luminosity upgrade in August 2000 and was planned to start the physics running in December 2001. However the startup has been delayed and the upgraded accelerator is still being recommissioned. The HERA – B collaboration has used the shutdown to upgrade and repair various subdetectors found to give problems during the 2000 run.

The physics program of HERA – B has been revised [56, 57] as we are not competitive in the field of b-hadron decays and CP violation, which has now been measured[1, 2]. The two main physics topics for HERA – B in the 2002 run are:

- **Nuclear dependence of charmonium production**
  The suppression of charmonium production in heavy ion collisions is one of the key signatures for quark-gluon plasma production. The interpretation of current data relies on models of charmonium production in hadron collisions and different models give very different predictions. HERA – B does not study heavy ion collisions, but is able to measure different aspects of charmonium production in proton-nucleus and thus provide constraints on the models. The goal is to measure the charmonium production cross section as a function of x_F. HERA – B has the ability to measure the dependence on the target nucleus over wide range ($-0.4 < x_F < 0.3$) unlike existing measurements [58], which only cover positive x_F. HERA – B will measure not only the production of J/ψ mesons, but also ψ', χ_c and, with lower statistics, different Υ states.

- **$b\bar{b}$ production cross section**
  As explained in Chapter 1, the beauty production cross section in hadronic collisions is not very well-known experimentally, nor theoretically at the center of mass energy of HERA – B. Using decays of beauty hadrons into J/ψ mesons, HERA – B will be able to measure the cross section with an uncertainty of about 15%. In Chapter 8 this method is demonstrated on the 2000 data sample.

The HERA – B detector and trigger system have undergone some changes to be able to accomplish this physics program, but the main requirement is still to trigger and reconstruct $J/\psi \to \ell^+\ell^-$ decays.

In this chapter we will describe the upgrade of SLT trigger algorithm carried out during the shutdown. The SLT has been prepared for the improved detector and the new physics goals. The trigger deficiencies described in the previous chapter has been fixed. As practically no data have been recorded with the upgraded detector, the algorithm has not yet been completely tuned for the 2002 detector configuration. Here the state of the SLT as of March 2002 will be described.

6.1 Detector upgrades/repairs

During the shutdown the HERA – B spectrometer components were repaired and some of them have also been upgraded. The details have been described elsewhere [59, 57]. We will give here a summary of the changes in different subsystems with focus on the changes important for the trigger performance:

- **Beam pipe:** The proton beam pipe between the VDS and the ECAL has been exchanged in order to reduce the amount of material in the detector. The new beam pipe has 0.3 mm thick aluminum walls. This upgrade reduces the background from photons conversions in the main tracker considerably, compared to the old beam pipe with 0.5 mm aluminum walls.

- **Target:** The target mechanics has been rebuild to replace the old, aging system. Similarly new hodoscope counters have been installed with increased acceptance. Together with improved software the result should be more stable interaction rates. For the 2002 run the proposed target configuration is two target wires, one of carbon and the second of titanium or tungsten. Running with two wires at the same time reduces a lot of the systematic uncertainties in the measurement of nuclear suppression of charmonium. This measurement also benefits from having the largest possible gap in the nuclear number between the two target materials[56]. The target wires have also been changed and are now 50 μm round wires instead of the thicker ribbons used during the 2000 run. This gives more precise determination of the z position of the primary interactions, which benefits the $b\bar{b}$ cross section analysis.

- **VDS:** The VDS was performing well during the 2000 run and has not suffered significant radiation damage. It has undergone only minor repairs of a few broken modules, mainly in the SI08 layer. It was planned that the VDS modules should be able to move closer to the beam in order to achieve 10 mrad coverage instead of the 15 mrad as in the 2000 run. This requires the RF shield placed between the VDS modules and the beam to be replaced with a smaller one. The replacement of the RF shield has not been approved by the HERA accelerator group and since the gains in our physics measurements from the increased acceptance do not look very significant, it has not pursued further.

- **Magnet chambers:** The magnet chambers were designed to allow matching of track segments in the VDS with track segments in the main tracker in a high multiplicity environment (running at 40 MHz interaction rate) and to reconstruct K_s^0 mesons which decay after the VDS detector. Given the changes in the physics program, the K_s^0 meson reconstruction have become less important. Also the interaction rate will be lower reducing the amount of fake matches between the VDS and main tracker track segments. The chambers contribute with a significant amount of material for tracks passing through the magnet. This affects mainly electron reconstruction due to the emission of unrecoverable bremsstrahlung, and photon reconstruction since they can convert in the material.

The bremsstrahlung broadens the $J/\psi \to e^+e^-$ peak, and photon conversion reduces the $\chi_c \to J/\psi\gamma$ reconstruction efficiency, both important for the 2002 physics program. All of the OTR and ITR chambers in the magnet have therefore been removed. Only the thin High-p_T chambers are left inside the magnet. The first superlayer of the magnet chambers is not located inside the magnetic field and has not been removed as it will be used by the SLT.

- **OTR:** More than 80% of the dead channels in the OTR were caused by short-circuits across capacitors on the high voltage boards inside the OTR modules. All 14000 capacitors have been replaced during the shutdown to avoid further short-circuits. The remaining 20% of the dead channels were mainly due to mechanical stability problems of some modules. These modules were stiffened by carbon rods and the worst of them were completely rebuilt. Tests have shown that the meantime between high voltage failures have improved by at least a factor of ten. Furthermore the OTR is able to run at slightly higher voltage leading to a higher gas amplification and thus hit efficiency. Most of the noisy channels in the OTR were caused by the readout electronics connecting the OTR to the FLT. The electronics has been modified and the noise has been significantly reduced. The lower noise allows the OTR to run with lower readout threshold. The net effect should be hit efficiencies above 95% [60].

- **ITR:** The ITR has been completely installed with four superlayers between the magnet and the RICH, two superlayers between the RICH and the ECAL and a superlayer in front of the magnet (MS01). The grounding scheme of the chambers has been improved to reduce noise and a new readout cheap (HELIX 3.1a) has been used. The reduction in noise and the new HELIX chip allows the use of the ITR by the FLT, though that option will most likely not be used due to manpower problems. A sparsification routine, which allows the SLT to reconstruct the hits in an RoI without reading out a complete ITR HELIX chip, has also been implemented.

- **ECAL:** The calorimeter has been fully equipped with readout and pretrigger electronics and all broken electronics boards repaired. The noise, which significantly reduced the precision in the 2000 run, have been drastically reduced by increasing the output impedance of the photomultiplier bases in the calorimeter. The ECAL pretrigger has been upgraded in order to implement bremsstrahlung recovery at the pretrigger level.

- **MUON:** The main problems in the MUON system, during the 2000 run, were high levels of noise and low intrinsic hit efficiency of the MUON pad system. Besides repairing broken electronics and wires, the noise situation was improved by using a different grounding scheme in the pad chambers. The low intrinsic hit efficiency has been understood [61] to be a mismatch in timing characteristic between the pre-amplifier mounted on each pad and the ASD8 readout chip. To solve this problem a complete disassemble of the MUON pad chambers is required, which was not possible during the shutdown. The pad efficiency for the 2002 run will therefore be around 80-90%. A procedure for measuring the efficiency precisely at the beginning of the run period has been prepared. The MUON pretrigger has been extended to cover both the pad and the pixels systems.

- **FLT:** The efforts for FLT upgrade have mainly been concentrated on making the system debugable and to ensure that the simulation gives an exact representation of the actual hardware. The data transmission from the OTR to the FLT was quite unstable during the 2000 run and work has been done to improve the cabling and electronics. The optical transmitters in the TLBs have been equipped with electronics which allows remote control of the light amplitude [62]. The integration of the ITR into the FLT has begun, but is not finished yet due to lack of manpower. The completion of this task will depend on the manpower needed for the commissioning of the OTR part of the FLT and the performance of the ITR in the 2002 run. Some algorithm improvements [63] have been implemented for the tracking efficiency, particularly in the MUON system. Finally a second TDU has been added to the FLT network, which allows simultaneous read out of the FLT tracks and the pretrigger candidates. This is useful for debugging the FLT and is required by the new 2002 default trigger configuration described below.

- **SLT:** The SLT farm has been upgraded with 1.2 GHz Celeron CPUs from Intel and 128 MByte RAM. The faster CPUs will allow the SLT farm to run at input rates of up to 30 kHz even with significant less rejection at the first triggers steps compared to the 2000 running. The CPU and memory upgrade also allows the SLT farm to be used for Monte Carlo and reprocessing in between fills, something the 4LT farm was already used for in the 2000 run. The algorithm upgrades in the SLT are described in the rest of this chapter.

- **4LT/DAQ:** The DAQ system should be able to run at up to 50 kHz input rate to the SLB without deadtime as the system can now run with a readout-FIFO depth of seven, see Section 3.7.4. The 4LT software has been improved in both stability and speed. A new dedicated logging server and an improved network connection have increased the rate capability and the 4LT should now be able to sustain a 50 Hz reconstruction and logging rate.

The HERA – B upgrade and repair program has been a major undertaking, necessary to ensure successful data taking in 2002.

6.2 Overview of the 2002 SLT trigger

Most of the trigger inefficiencies presented in Chapter 5 were found to be due to detector inefficiencies that should be much reduced for the 2002 run. Another problem was the trigger rate from the pretriggers, which had to be downscaled or reduced by using high p_T threshold. With the improvements in the ECAL and MUON systems the pretrigger rates will be even higher and it is clear that we need to use FLT to reduce the input rate to the SLT to a manageable level. The FLT tracking efficiency, ε, is estimated to be in the range of 27-54% for muons [57] depending strongly on the performance of the subdetectors. This is quite low compared to the SLT efficiency and if the FLT is required to find both tracks the trigger efficiency scales with a factor ε^2.

The efficiency of the design trigger configuration is therefore associated with large uncertainties. The use of the inner detectors (ITR and MUON pixel chambers) in the FLT also looks quite uncertain. Therefore a different trigger configuration has been chosen as the baseline for the 2002 run. Instead of requiring both leptons to be found by the FLT, the FLT will only require one lepton to be found. The trigger efficiency in this configuration scales as $1-(1-\varepsilon)^2 = 2\varepsilon-\varepsilon^2$, a large improvement. Requiring only one track gives of course much less rejection. The trigger can therefore not run at as high interaction rates as the design trigger. Even at the planned interaction rates of 5-10 MHz, a higher p_T threshold has to be applied to the FLT track.

Even with high p_T thresholds the rate of ECAL pretriggered FLT tracks will be very high. Furthermore the electrons in most analyses had to be identified using bremsstrahlung radiation. This method reduces the statistics significantly compared to the muon case. It was therefore decided to give the MUON pretriggered FLT tracks the major part of the bandwidth. Instead of simply downscaling the electron candidates or not using them at all, the FLT will require the electron candidates to have an associated bremsstrahlung cluster found by the ECAL pretrigger. This gives the needed rejection for the SLT to be able to handle the input rate.

The second lepton will have to be found by the SLT independently of the FLT. The SLT will use the pretrigger to seed the second track. Since both the pretriggers and the SLT use the inner detectors, the second lepton can be found also in the inner detector, giving us most of the forward acceptance, as only 3% of the J/ψ have both leptons in the inner detectors.

The improvements in the SLT algorithm are described in the following sections, but can be summarized as follows:

- **More efficient muon tracking:** The hit preparation of the MUON system has been modified. It no longer has a requirement of the hit isolation and or to have hits in both layers of MU01. The tracking has been changed to allow multiple candidates per pad coincidence.

- **Inclusion of the inner detectors:** The MUON pixel chambers and the ITR are now used by the muon tracking and Slicer+Refit in order to extend the acceptance in x_F.

- **Use of magnet chambers:** The track propagation through the magnet now uses hits from the first magnet chambers. This results in more precise RoIs for L2Sili.

- **VDS tracking:** L2Sili has been rewritten as a full Kalman filter to provide higher track precision and better ghost suppression. The rewrite also resulted in a much faster code.

- **Vertexing:** The two leptons are now required to come from a common vertex. This rejects track pairs coming from different interactions and many ghost tracks, which often do not point to target wires. There is no requirement on decay length since we need also to record the prompt J/ψ for the charmonium suppression measurement.

It was planned to restructure the hit request code as it was seen to be very slow. This has not been done yet and given the CPU upgrade it should no longer be necessary.

To improve the possibilities for understanding the trigger performance after the data has been recorded, a new type of "bypass" events have been implemented, called "SLP bypass". These are about 1 in 10,000 of the events triggered by the FLT. They will be accepted no matter in which point of the SLT algorithm the event was rejected. This is unlike the "SLT bypass" events, which were recorded only if they passed the seeding, where we can have quite large rejections already. The "SLT bypass" event will still be kept for better study of the algorithms later in the trigger chain. The downscaling of bypass events will be adjusted to use only a few Hz of the total output rate.

6.3 Seeding

The seeding procedure for the 2002 run has been changed significantly to allow combining seeds based on FLT tracks with tracks based on pretriggers. To do that the seeding is done in two parts: a seed finding part, which handles the readout of the FLT TDUs, muon tracking, etc., and a seed pairing part, which combines the tracks found into pairs for the real SLT algorithm.

The seeding has also been modified to run at higher input rates mainly reducing the amount of data read out per event to decrease the load on the SHARC switch and the SLB. The aim is to be able to run at input rates of up to 30 kHz.

6.3.1 FLT seeding

The FLT tracks are used as seeds directly as described in Section 4.2.3. The FLT uses the same pretrigger information to start its track search as the SLT uses for the ECAL/MUON pretriggered seeds. In order not to find the same track twice, the pretrigger clusters which seeded the tracks found by the FLT are removed from ECAL/MUON pretriggered seeding. The FLT track does not provide information on which pretrigger cluster seeded it and the removal is therefore done by extrapolating the FLT track to the ECAL or MU03 layer. ECAL pretrigger cluster within 5 cm and MUON pretrigger seeds within 20 cm of the extrapolated FLT track are removed. This might remove more than one pretrigger seed even if there is only one FLT track, but it has been found that efficiency loss from this is negligible. It also means that if the FLT found both the leptons from a $J/\psi \to \mu^+\mu^-$ decay, the SLT does not have to do any tracking in the MUON system.

6.3.2 ECAL pretriggered seeds

The ECAL pretrigger seeding of the SLT no longer uses CARE to find clusters but instead uses ECAL pretrigger information, read out from the 2nd TDU in the FLT network. This is possible because the problem with the data transmission has been solved. It gives a massive reduction in the CPU time used by the SLT and the data from the ECAL no longer needs to be read out, reducing the load on the SHARC switch. The RoIs for Slicer/Refit will be larger though, as the precision of the ECAL pretrigger is worse than the full calorimeter reconstruction in CARE.

The 2000 code assigned charges to ECAL seeds in pairs and was based on the assumption of out-benders, see Section 4.2.1. This is no longer done. Each ECAL pretrigger cluster will give two RoIs for the SLT, one with positive and one with negative charge. The RoIs are calculated in the same way as in 2000 assuming the momentum is equal to the measured energy and using the field parameterization from L2Magnet.

The ECAL pretrigger also looks for clusters, which from their position looks like bremsstrahlung radiation from the electron/positron which caused the main cluster. This information is used in the mass calculation when the track pairs are considered, see below. It could potentially be used also to assign a charge to the RoI, but this is currently not done.

6.3.3 Muon pretriggered seeds

The muon tracking in 2000 was reading out the complete data of MUON tube system, consisting of 18 FEDs each with a data size of 136 bytes. The FEDs were spread over three SHARCs placed in the same SHARC module. This limited the readout rate to 12-14 kHz due to the bandwidth of the SHARC links. We do however only use the data from 0° layers. For the 2002 run we have therefore implemented an RoI based MUON read out similar to the method used by the other tracking algorithms. Only data from the 0° layers in MU01 and MU02 are reads out and furthermore only data from the FEDs which overlap with RoIs will be read out. For simplicity we still read out full FEDs. From simulation we expect only an average of 0.7 RoIs in events triggered by a single muon track in the FLT. It is below one because 60% of the events have no second MUON pretrigger coincidence or it has already been used in a track by the FLT. This means that on average only 180 bytes of data are transfered per event. This allows us to run at more than 30 kHz and the actual limit will come from overhead of message passing in the MUON part of the SLB rather than from the actual amount of data transmitted.

The hit preparation of the MUON system was changed to be able to handle the FED based readout. At the same time the strict isolation and double layer requirement of the 2000 version, see Section 4.2.2, was dropped. To avoid getting too many track candidates from particles which generate showers in the muon system, the hits are clusterized. If the distance between two tube cell hits is less than 2.1 cm, they are joined into a single cluster. This distance corresponds to two neighboring cells or two cells located one behind the other in a double layer. Many cells can be joined into a single cluster. Cluster wider than 3 cells are broken into multiple clusters. For the pixels a clusterisation is also done. Here a pixel is added to a cluster if it lies inside or on the border of the rectangle spanned by the pixels already added to a cluster. The maximum size for a pixel cluster is 5.5 cm in both x and y corresponding to six pixels. Larger clusters are split into two.

The muon tracking itself is very similar to the 2000 version: a MUON pad coincidence is combined with a hit in the 0° layers of MU01 and MU02. The RoI for the hit search (and readout) is defined from the MUON pad coincidence and use the same window size as the FLT does. The size depends on the coincidence position and varies from 70 cm close to the beam pipe, where there is less multiple scattering, to 110 cm furthest away. If a hit in MU02 is found within the RoI, the position of the hit and the pretrigger coincidence are extrapolated in a straight line to MU01. Here a tube hit should lie within 30 cm of the extrapolation in x. For pixel hits, the x and y coordinate are required to be within 6 cm of the extrapolation. If several hits are found, all hit combinations are taken as valid muon tracks. The SLT algorithm limits the number of tracks to 24. The potential muon tracks are therefore sorted according to their quality, where the quality is defined as the difference between the MU01 hit position in x and the extrapolated track.

The RoIs for the OTR/ITR tracking algorithm are defined in exactly the same way as for the 2000 version (Section 4.2.2), assuming an average $p_T = 1.6$ GeV/c and creating RoIs for both positive and negative muon candidates.

6.3.4 Seed pairing

The RoIs found by the three seeding procedures described above are in principle combined into all possible pairs of RoIs with the only restriction that the RoIs should come from different pretrigger seeds. The pairs are sorted into electron pairs, muon pairs and electron+muon pair, where the last combination most likely will be rejected. For the others it is possible to require the charges to be opposite, the invariant mass of the pair should be higher than some threshold or the distance between the RoI centers at TC02 should be larger than some threshold. The distance cut is supposed to remove pairs made up of clone tracks which were not suppressed by the requirement of different pretrigger seeds. Currently a value of 40 cm is used. The two other cuts remove non-J/ψ events, which could be potentially useful, for example by using same sign events to describe the pion decay background under the J/ψ peak in the invariant mass spectrum. They can be used to reduce the rates into the next steps of the SLT algorithm if needed.

6.4 Slicer/Refit

The basic principles used for the tracking through the Main tracking system are unchanged from the 2000 run: Slicer determines that a reasonable amount of hits are available, followed by the Kalman filters of RefitX and RefitY to do further track rejection and calculate track parameters. The main improvement is the inclusion of the ITR into the tracking. The use of ITR was made possible by a sparsification routine provided by the ITR group. This routine performs the hit preparation even if only information from a part of a individual HELIX chip is read out. The old sparsification routine required the read out of a full HELIX chip in order to do noise suppression. The read out of a complete HELIX chip (144 bytes) would put too high load on the SHARC switch as the tracking requires hits from many ITR layers and consequently many chips.

For Slicer the inclusion of the ITR is straight forward. The algorithm must just take the better resolution of ITR hits into account. For the Refits it is slightly more complicated as a single track can pass through ITR chambers in some superlayers and OTR chambers in others. Usually it only passes through one type of chamber in a given superlayer, but in particular with the large Muon pretriggered RoIs, one might get hits in both the OTR and ITR at the same superlayer. When track candidates are propagated into a new superlayer they are therefore first combined with hits in the OTR and up to the five best tracks are selected. The original track candidates are then combined with hits in the ITR and the five best tracks found in either the ITR or OTR are propagated to the next superlayer. The best tracks are determined by the track quality, Q, as defined in Equation 4.6.

With the inclusion of the ITR and the improvements in detector performance we no longer need to allow one superlayer to be empty in Slicer. The number of hits required to be found by Slicer in a single slice for a track to be accepted still needs to be tuned when real data becomes available. The baseline requirement used in the evaluation in the next chapter is 12 hits per slice.

The direction of the track reconstruction in the Kalman filter of RefitX/RefitY has been changed to start at TC02/MS15 and move toward PC01/MS10. This change was made to reduce the number of possible hit combinations in MUON seeded RoIs, RoIs at TC02 being much smaller than at PC01. To further reduce the number of combinations in RefitX, a hit is only considered for addition to a track candidate if it lies within 5σ of the expected position. This is very similar to the strategy used in the new L2Sili, described below. The requirement makes practically all hits satisfy the χ^2 requirement. The calculated χ^2 is still used for the track quality evaluation.

It has also been tried [64] to use the OTR drift time in Refit. It turned out that the improvement in tracking precision is marginal. There is also no improvement in rejection, but the time spent in Refit went up significantly. The drift time is therefore not planned to be used in the 2002 run. There have also been attempts to change the definition of the χ^2 used in Refit to give better rejection of wrong tracks, but the conclusion was not reached yet. For this text we will use the old definition.

6.5 L2Magnet

The track propagation through the magnet uses the same magnetic field parameterization as in the 2000 version. Instead of using the propagated Refit track parameters to setup the RoI for L2Sili, they are used to define an RoI in the first magnet chambers placed between VDS vessel and the magnetic field. The hit information from these chambers are used to reconstruct spacepoints. If a spacepoint is reconstructed, the track parameters are recalculated using the target box, the spacepoint and the propagated refit track parameters. This method gives a more precise determination of the RoIs in the VDS. The smaller RoIs allow us to read out less of the VDS and spend less time for the time consuming sparsification. The smaller RoIs also reduced the probability to find a ghost track in the VDS.

The RoI in the magnet chamber is defined from the extrapolated Refit track and has a width of $\pm 3\sigma$, where σ is the calculated track uncertainty at the magnet chamber. The magnet chambers are constructed of three different types of chambers. Small angles (10–45 mrad) are covered by a VDS module, SI08, with two double layers similar to the ones in SI04-SI07. The majority of the angular space (20-120 mrad) is covered by an ITR module with four layers, two 0° and two $\pm 5^\circ$ layers. The outermost tracks (80–220 mrad) pass through three OTR layers, one 0° and two $\pm 5^\circ$ layers. If an RoI covers several modules, they are all read out and the spacepoints in each module are reconstructed independently.

The spacepoints are constructed slightly differently in the different detectors. In SI08 simple spacepoints are constructed by combining hits from both sides on a double-sided layer. These 2-hit spacepoints are joined with 2-hit spacepoints from the second double-sided layer to form a 4-hit spacepoint, if the two 2-hit spacepoint positions lie close each other. A 2-hit spacepoint which could not be combined with another spacepoint can be combined with single hits not used for a 4-hit spacepoint, giving a 3-hit spacepoint. Only 3- and 4-hit spacepoints are used in the track fit. In MC01 all possible 2-hit spacepoints are constructed. A 3-hit spacepoint is constructed if all three 2-hit spacepoints generated by the 3 hits lie close to each other. If no 3-hit spacepoints are found (happens for about 15% of tracks passing through the OTR) the 2-hit spacepoints are used in the track fit. In MS01 the hits from the two 0° layers are combined into "single" hits if they lie at the same x coordinate. 2-hit and 3-hit spacepoints are constructed as for the OTR. The 2-hit spacepoints are only kept if they are in fact constructed from two 0° layer hits and one stereo layer hit.

A spacepoint found by the above technique is used to calculate the RoIs for the VDS if the spacepoint itself lies inside of the RoI in the magnet chamber. The RoI is calculated using an iterative procedure as for the 2000 version. The refit track is prolonged to the center of the magnet and the initial track slope estimate is calculated by a fitting a straight line using this position, the spacepoint and the target box. The track slopes are used to look up the fieldmaps to correct the estimate for refit track position inside the magnet and the track momentum. This calculation is reiterated in the same way as is done without the extra spacepoint.

An RoI is created for each spacepoint. If no spacepoint is found we currently use the original prolongation of the refit track. L2Magnet therefore provides no event rejection. We could gain some rejection by requiring a spacepoint, but since there is only one redundant layer in the magnet chambers, low hit efficiencies can cause a large loss of trigger efficiency. This will be optimized once the actual chamber performances are known.

From the simulation of 2002 data, we expect 97% of the muons from $J/\psi \to \mu^+\mu^-$ decays to have a spacepoint assigned, while for general inelastic event only 65% of the RoIs are expected to have a spacepoint. The inclusion of spacepoints reduces the RoI size in the VDS by about 30% but because the data request in the VDS is limited[1] to complete HELIX chips, the amount of data transfered is reduced by less than 10%.

6.6 L2Sili

The L2Sili tracking algorithm has been completely rewritten for the 2002 run. The idea is to make it faster, give more precise track parameters and reduce the amount of ghosts tracks. The tracking efficiency is already very high (above 97% per track) for the 2000 version and can not be improved too much. The tracking precision and ghost rates are improved by turning the algorithm into a Kalman filter with correct treatment of multiple scattering. The latency is improved by using a much simpler pattern recognition and optimizing the code structure.

The basic concept of L2Sili has not changed. It uses the track seeds from L2Magnet with the RoI size given by three times the seed uncertainty. L2Magnet can now give several seeds for a single track in the main tracker. The RoIs of these are joined into a single large RoI for the hit request. The seeds are used as the initial track candidates for the Kalman filter, which closely follow the algorithm presented in Appendix A.3. The tracking is still done separately in x and y, with the hit measurement in the x view given by (see Equation (4.13–4.14)):

$$m = y \sin\theta + \tilde{m}. \tag{6.1}$$

We use then a measurement matrix: $\mathbf{H_k} = [\,\cos\theta_k \;\; z_k \cos\theta_k\,]$. The hit measurement uncertainty is given by:

$$\sigma_m^2 = \sigma_y^2 \sin^2\theta + \sigma_{\tilde{m}}^2, \tag{6.2}$$

where σ_y is the uncertainty of the y coordinate given by the tracking procedure performed up to the layer that contains the current hit. The equations for updating the state vector and the covariance matrix are:

$$\vec{x}_{k,proj} = \vec{x}_{k-1}, \tag{6.3}$$

$$\mathbf{C}_{k,proj} = \mathbf{C}_{k-1} + \sigma_{MS}^2 \begin{pmatrix} 1 & -z_k \\ -z_k & z_k^2 \end{pmatrix}, \tag{6.4}$$

$$\mathbf{C}_k = \left[(\mathbf{C}_{k,proj})^{-1} + \frac{1}{\sigma_k^2} \begin{pmatrix} z_k^2 \cos^2\theta_k & z_k \cos^2\theta_k \\ z_k \cos^2\theta_k & \cos^2\theta_k \end{pmatrix} \right]^{-1}, \tag{6.5}$$

$$\vec{x}_k = \mathbf{C}_k \left[(\mathbf{C}_{k,proj})^{-1} \vec{x}_{k,proj} + \frac{m}{\sigma_k^2} \begin{pmatrix} z_k \cos\theta_k \\ \cos\theta_k \end{pmatrix} \right]. \tag{6.6}$$

To reduce the amount of calculations and thus increase the speed, the multiple scattering is only added after processing the information for each superlayer. This approximation is reasonably well justified as the two silicon layers in a superlayer are only separated by around

[1] This is the same limitation as existed for the ITR. It is currently under investigation whether the new ITR sparsification routine can also be used for the VDS.

0.8 cm. The multiple scattering handling is further simplified by only adding multiple scattering terms to the covariance matrix for superlayers in which the track candidates had hits and the algorithm does not take into account the areas in the detector where two modules in the same superlayer overlap and has twice the material. σ_{MS} is calculated according to Equation A.17 assuming each superlayer to be a thin scatterer[65].

The pattern recognition is simplified by no longer using a χ^2. Instead a hit is added to a track candidate if it lies within 5σ of the predicted position. Both the uncertainty from the prediction and the hit uncertainty are included in σ. When processing hits in a superlayer all allowed track candidates are constructed including tracks without any new hits in the current superlayer. After each superlayer only the 10 candidates with the most hits assigned are kept. From these 10 we furthermore remove all candidates which are a subset of another. By imposing the 10 track candidate limit, the time spent comparing track candidates to find subsets has an upper limit, unlike in the case of the 2000 algorithm. The track candidate with no hits is kept up to the last superlayer in order not to loose a real track by picking up a wrong hit in a superlayer, where the real track does not have hits.

The algorithm therefore proceeds as follows (starting at the SI07 superlayer):

- Track the candidates through the superlayer:
  - Loop over the layers in the superlayer:
    - If the layer is empty, go to the next layer.
    - Loop over the track candidates:
      - Pre-calculate new covariance matrix to be used if a hit belonging to this layer will be added to the track candidate.
      - Calculate the 5σ boundaries of the predicted hit position.
      - For every hit in the layer, check if hit position is less than 5σ away from the predicted track position.
      - If this is the case, construct a new track candidate using the added hit.
- Cleanup the list of track candidates:
  - Find the 10 candidates with highest number of hits.
  - Remove all other candidates except the ones without any hits at all.
  - Candidates whose hits are subsets of the hits belonging to a different candidate are removed.
- Add multiple scattering to the covariance matrix for all candidates which had at least one hit in the current superlayer.
- Continue to next superlayer.

By pre-calculating various values as early as possible, a significant amount of CPU time is saved. Once the hits in all superlayers have been processed, only the candidate with the highest number of hits is kept. L2Sili requires this track to have at least 3 hits in each view and that its hits comes from two superlayers in each view. The cuts are identical to the ones used in the old version and represent the minimum requirement to ensure that the track parameters are reliably reconstructed.

Sample	1	2	3	4	5	6	7
Old L2Sili algorithm	99.5%	98.9%	97.8%	98.0%	96.6%	95.6%	97.2%
New L2Sili algorithm	99.7%	99.8%	99.7%	98.6%	99.3%	97.6%	99.2%

TABLE 6.1: The tracking efficiencies of the two L2Sili algorithms.

6.6.1 Comparison of the 2002 L2Sili version to the 2000 version

In the next chapter, the performance of the complete SLT trigger chain will be evaluated. As the 2002 detector is somewhat different from the 2000 one and the algorithms called before L2Sili in the SLT chain have changed, the results of the 2002 evaluation cannot be directly compared to the results of the 2000 evaluation in Chapter 5. For comparing the two L2Sili versions we have therefore applied the 2000 trigger algorithm to the same data as in Chapter 5 except that the new L2Sili has been used. The results of this analysis are given in the following and show that the new version is a significant improvement over the old version.

Efficiency:

The tracking efficiency of the old L2Sili algorithm was already very high. In Table 6.1 the efficiency of the old and new L2Sili algorithms are given for the same seven $J/\psi \to \mu^+\mu^-$ Monte Carlo samples as in Section 5.4.1. The efficiencies are given for events in which the required number of VDS hits are present. We see that the new algorithm is a few percent more efficient at reconstructing leptons from J/ψ decays than the old one.

Running the new version of L2Sili on SLT bypass events, we find 10% more real tracks than using the old version. Real tracks here are defined as tracks confirmed by the offline tracking code. The 10% additional tracks found are mainly at the limits of the RoI, which would explain why the increase in tracking efficiency is larger than in the J/ψ Monte Carlo, where most of the tracks lie in the center of the RoI.

Precision:

The tracking precision of L2Sili has been evaluated using $J/\psi \to \mu^+\mu^-$ Monte Carlo as in Section 5.4.2. The results obtained using the two different variants are compared in Table 6.2. We see that the new L2Sili has approximately as good tracking precision as CATS in both slope and position. The uncertainty in the track slope is underestimated as the multiple scattering in the Roman pots crossed by the tracks in their way between the target and the first layer with a track hit is not taken into account. If the measured track slope is instead compared to the true slope at the layer containing the first hit, the pulls are very close to one.

Timing:

Figure 6.1 shows CPU time spent per non-empty RoI in the two L2Sili versions. The data used are SLT bypass events from run 16665. The new L2Sili version is on average seven times faster and in particular the time distribution does not have the same long tails the old version has. The improvement in speed is mainly due to the new simplified pattern recognition, which keeps the number of track candidates down to a manageable level, and partially due to code optimizations.

	New L2Sili	Old L2Sili	CATS
$\Delta x/\mu$m	12	16	10.4
$\Delta y/\mu$m	12	16	10.7
Δt_x/mrad	0.12	0.13	0.12
Δt_y/mrad	0.11	0.14	0.13
Pull(x)	0.90	0.57	1.10
Pull(y)	0.91	0.53	1.15
Pull(t_x)	1.38	1.08	1.34
Pull(t_y)	1.47	1.12	1.36

TABLE 6.2: Width of residuals and pulls of L2Sili track parameters for the MUON+SLT trigger with the new and old L2Sili version . The CATS column shows the equivalent results from the offline tracking.

Ghosts fraction	Single track (electrons)	Single track (muons)
Old L2Sili	16%	32%
New L2Sili	17%	30%

TABLE 6.3: The fraction of tracks which, from comparison with offline tracking, is estimated to be ghost tracks.

Ghosts:

The ghost rate is investigated by running the trigger simulation on the SLT bypass events from run 16665 and comparing the found L2Sili tracks with the offline reconstructed VDS tracks. The comparison is done with the χ^2 definition of Section 5.2.5 and the estimated fractions of ghost tracks are given in Table 6.3 for both the old and new version of L2Sili. The table shows that the new version of L2Sili did not succeed in significantly reducing the amount of ghost tracks. Figure 6.2 however shows that the amount of hits associated to ghost tracks have been reduced. If needed, the ghost fraction can therefore be reduced by increasing the required number of hits on a track with little loss in efficiency.

6.7 L2Vertex

The L2Vertex algorithm is supposed to reduce the rate by requiring that the two leptons come from the same vertex. When using multiple target wires, tracks from interactions on different wires are easily rejected by the vertex requirement, while tracks from different interactions on the same wire are much harder to discriminate. The rejection factor of L2Vertex therefore increases with the number of target wires. Ghost tracks in general does not point back to the target wires either. A significant fraction of them is therefore rejected by L2Vertex. Originally L2Vertex was meant to select detached vertices from b hadron decays, where the vertex is separated from the target wire with a few millimeters. This can give a much larger rejection as most tracks come from the target wire. This selection cannot be used anymore as one of main physics topics is the study of prompt J/ψ production.

The L2Vertex algorithm has been described before in Ref. [52]. It finds the di-lepton vertex,

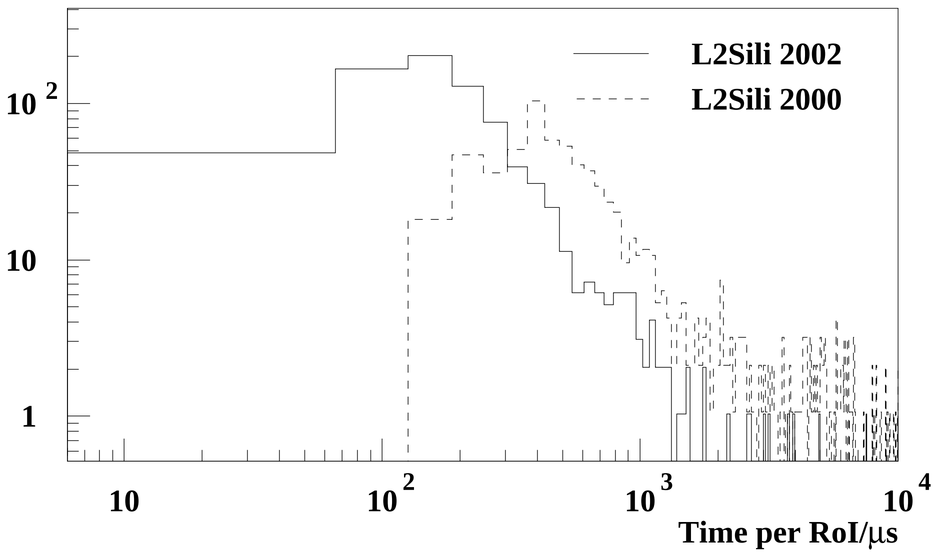

FIGURE 6.1: Time spent per non-empty RoI for the 2000 and 2002 L2Sili algorithm on a 933 MHz CPU.

(x_v, y_v, z_v) by minimizing a χ^2, which "measures" the minimum distance between the two tracks, when track uncertainties are taken into account. The χ^2 is defined as:

$$\chi^2 = \sum_{i=1,2} \left(\frac{(b_{x,i} + t_{x,i} z_v - x_v)^2}{\sigma_{x,i}^2} + \frac{(b_{y,i} + t_{y,i} z_v - y_v)^2}{\sigma_{y,i}^2} \right), \tag{6.7}$$

where the sum is done over the two tracks and σ_x is given by:

$$\sigma_x^2 = \sigma_{b_x}^2 + \sigma_{t_x}^2 z_v^2 + 2 z_v \mathrm{cov}\{b_x, t_x\}. \tag{6.8}$$

A similar formula is used for σ_y. The minimization of χ^2 is done by applying Newton-Rhapsons method [66] to the nonlinear equation system:

$$f_i(x_v, y_v, z_v) \quad = \quad 0, \qquad \text{for } i = 1 \ldots 3, \tag{6.9}$$

where:

$$f_1(x_v, y_v, z_v) \quad = \quad \frac{\partial \chi^2}{\partial x_v} = \sum_{i=1,2} \left(\frac{-2(b_{x,i} + t_{x,i} z_v - x_v)}{\sigma_{x,i}^2} \right), \tag{6.10}$$

$$f_2(x_v, y_v, z_v) \quad = \quad \frac{\partial \chi^2}{\partial y_v} = \sum_{i=1,2} \left(\frac{-2(b_{y,i} + t_{y,i} z_v - y_v)}{\sigma_{y,i}^2} \right), \tag{6.11}$$

$$f_3(x_v, y_v, z_v) \quad = \quad \frac{\partial \chi^2}{\partial z_v} = \sum_{i=1,2} \left(\frac{2t_{x,i}(b_{x,i} + t_{x,i} z_v - x_v)}{\sigma_{x,i}^2} + \frac{2t_{y,i}(b_{x,i} + t_{y,i} z_v - y_v)}{\sigma_{y,i}^2} \right). \tag{6.12}$$

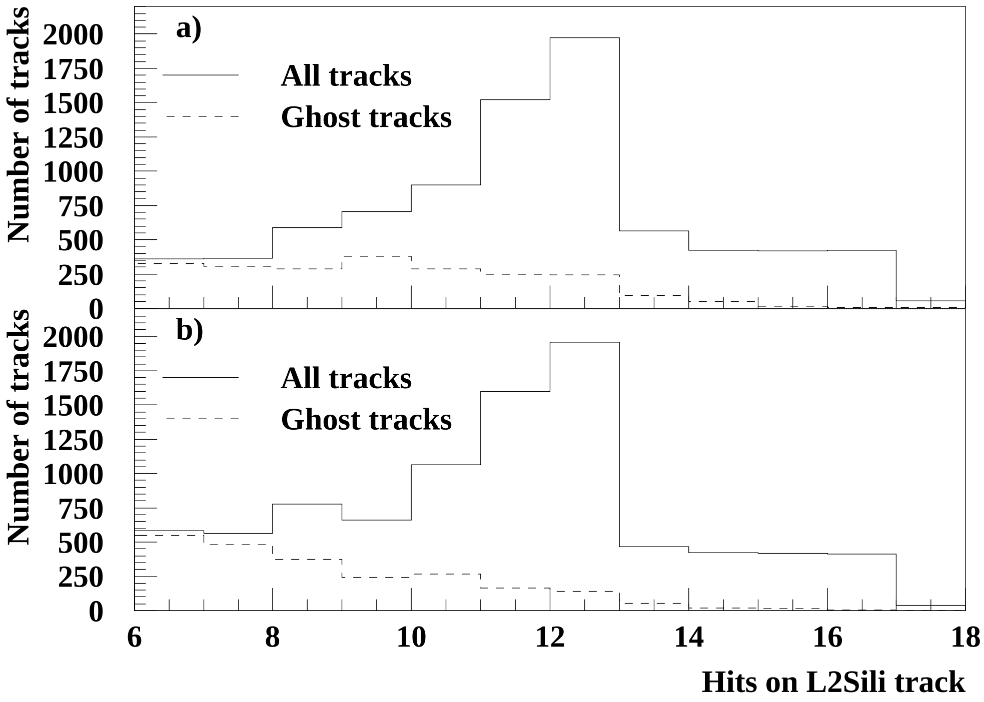

FIGURE 6.2: Hits per L2Sili track for all tracks (full line) and ghost tracks (dashed line). a) Results for the old version of L2Sili. b) Results for the new version of L2Sili.

The z_v dependence in $\sigma_{x/y}$ is ignored in the derivative to simplify the calculation. After each iteration in the fit, the uncertainties are updated with the new vertex position. The Newton-Rhapsons method iterates to the solution using the formula:

$$\begin{pmatrix} x_{v,n} \\ y_{v,n} \\ z_{v,n} \end{pmatrix} = \begin{pmatrix} x_{v,n-1} \\ y_{v,n-1} \\ z_{v,n-1} \end{pmatrix} - \mathcal{J}^{-1} \begin{pmatrix} f_1 \\ f_2 \\ f_3 \end{pmatrix}, \tag{6.13}$$

where $\mathcal{J}$ is the Jacobian given by:

$$\mathcal{J} = \begin{pmatrix} \sum\limits_{i=1,2} \frac{2}{\sigma_{x,i}^2} & 0 & \sum\limits_{i=1,2} \frac{-2t_{x,i}}{\sigma_{x,i}^2} \\ 0 & \sum\limits_{i=1,2} \frac{2}{\sigma_{y,i}^2} & \sum\limits_{i=1,2} \frac{-2t_{y,i}}{\sigma_{y,i}^2} \\ \sum\limits_{i=1,2} \frac{-2t_{x,i}}{\sigma_{x,i}^2} & \sum\limits_{i=1,2} \frac{-2t_{y,i}}{\sigma_{y,i}^2} & \sum\limits_{i=1,2} \left(\frac{-2t_{x,i}^2}{\sigma_{x,i}^2} + \frac{-2t_{y,i}^2}{\sigma_{y,i}^2} \right) \end{pmatrix}. \tag{6.14}$$

The vertexing starts at the point closest to the two tracks and iterates until the change is not numerical significant or until ten iterations have been performed.

L2Vertex accepts a vertex if the χ^2 is less than 20. This value reflects that the tracking errors are underestimated in L2Sili and therefore χ^2 can become quite large. With this cut, 73% of the events in run 16665 accepted by L2Sili are also accepted by L2Vertex, and 77% of the events rejected by L2Vertex contain at least one VDS ghost track. The fraction of VDS ghost pairs is thereby reduced by almost 50% by L2Vertex.

Chapter 7

Expected performance in 2002

In this chapter the performance of the HERA – B trigger for the 2002 run will be evaluated. Several options for the 2002 runs exist and here the following trigger configurations are studied:

- **Single FLT-Muon:**
  The FLT triggers on a single MUON pretrigger seeded track, which is then confirmed by the SLT. The SLT reconstructs a second muon, also seeded by the MUON pretriggers.

- **Single FLT-Electron:**
  The FLT triggers on a single ECAL pretrigger seeded track for which the pretrigger has also found an ECAL cluster, which could be from bremsstrahlung radiation. The SLT confirms the FLT track and searches for a second electron candidates seeded ECAL pretrigger, which is not required to have a bremsstrahlung cluster associated.

- **FLT-Muon pair:**
  The FLT triggers on two MUON pretrigger seeded tracks with an invariant mass above $2\,\mathrm{GeV}/c^2$. The track pair is then confirmed by the SLT.

- **FLT-Electron pair:**
  The FLT triggers on two ECAL pretrigger seeded tracks with an invariant mass above $2\,\mathrm{GeV}/c^2$. The track pair is then confirmed by the SLT. No bremsstrahlung clusters are required.

- **MUON pretriggers only:**
  The FLT does not do any tracking, but just requires two seeds from the MUON pretrigger and send them to the SLT. The SLT does the complete tracking based on the MUON pretrigger seeds.

- **ECAL pretriggers only:**
  The FLT does not do any tracking, but just requires two seeds from the ECAL pretrigger. The SLT does the complete tracking based on the ECAL pretrigger seeds.

The pretrigger-only configurations are only foreseen to be used at the beginning of the data taking and will be used to take data for the FLT commissioning. The evaluation in this chapter will be focused mainly on the baseline trigger configuration: Single FLT-Muon, which has the highest efficiencies. Some numbers will also be presented for the other configurations.

Due to the delayed startup of HERA, HERA – B has at the time of this writing (March 2002) not been able to take significant amounts of data. The performance of the different repaired subdetectors has therefore not been checked in realistic run conditions. We will therefore assume ideal detector performance (See Chapter 5) for the simulation. Also the evaluation is necessarily and exclusively based on simulated events. As seen in Chapter 5, the precision of the prediction, particularly for the trigger rates, can not be trusted at a very high degree. Finally the trigger groups are using the delay of the startup to improve the trigger algorithms and the latest changes, such as the improvement in the FLT tracking in Ref. [63], have not been included.

7.1 Trigger rates

The trigger rates from the FLT and SLT determine how loose we can set the cuts applied in the trigger and therefore influence the trigger efficiency. For the rate estimates, one million inelastic Monte Carlo events are used. These are equally distributed between the carbon wire, Inner 2, and the titanium wire, Below 2. The events have been used to provide two Monte Carlo samples where the number of interactions on each wire is simulated according to a Poisson distribution with a mean corresponding to 5 and 10 MHz total interaction rate, respectively. The one million Monte Carlo events corresponds to 0.2 and 0.1 seconds of data taking at 5 and 10 MHz interaction rate.

7.1.1 FLT output rates

The output rate of the FLT is limited by the rate which the SLT/DAQ system can handle. The SLT limit is given by the amount of time necessary to process each event. As we shall see below, the SLT has no problems to keep up at very high rates after the CPU upgrade. The DAQ has to be able to read out the FEDs on the detector to the SLB fast enough and supply the input data to the SLT. With the availability of the full readout-FIFO depth in the DAQ, the read out can run at 50 kHz without significant dead time, see Section 3.7.4.

The SLT needs to retrieve the information from the FLT at the full FLT trigger rate. The SLB containing the FLT information has been constructed to be able to provide the information at 50 kHz [36]. However the amount of information to be transfered has been doubled compared to the design as we now read out both the FLT and the pretriggers using two separate TDUs in the FLT network. By connecting the two TDUs two separate SHARC modules, this should however not pose any major problem as the 50 kHz limit comes from overhead in the SHARC and not from bandwidth limitations. Furthermore most of the triggered events will require the SLT to track through the MUON system. With the new RoI based muon tracking, the data from a given MUON FED are only requested in about 10% of the FLT triggered events. If the MUON FEDs are therefore spread over many SHARCs, the transfer of data for the MUON tracking should also be able to run at the 50 kHz. As a consequence, the limit on FLT output rate should be around 50 kHz. To be conservative we have used a limit of 30 kHz to allow for underestimations of the trigger rates in the Monte Carlo and unforseen problems given by the DAQ.

Figure 7.1 shows the expected FLT single track trigger rates as a function of the p_T requirement in the FLT. At low p_T thresholds (0.5–1.2 GeV/c), the rate falls off exponentially, while at higher p_T thresholds, it drops much slower. The latter happens because of the limited precision of the FLT for high momentum tracks and because ghost tracks in the FLT tend to

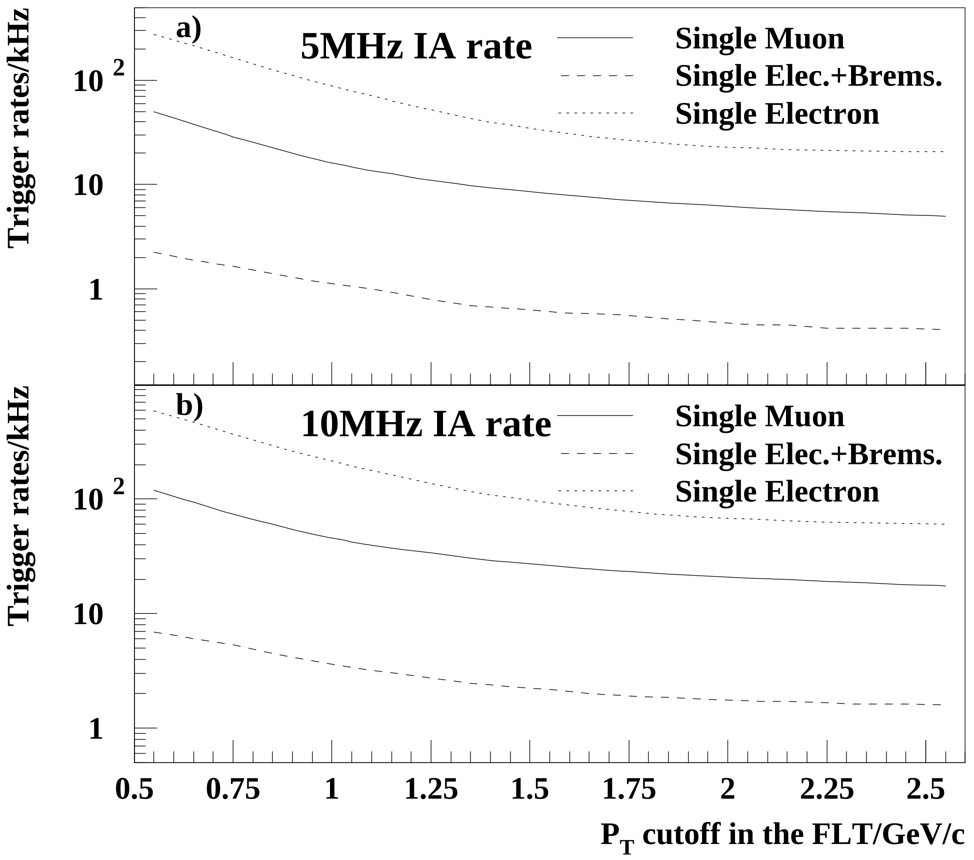

FIGURE 7.1: Expected output rate of the single FLT trigger at 5 MHz and 10 MHz interaction rate. The rate of FLT trigger events are shown as a function of the p_T cut applied to the muon and electron candidates.

point straight back to target (the center of the FLTs internal RoIs) and thus have "infinite" momentum.

The rate of MUON pretrigger seeded FLT tracks can be limited to around 30 kHz by requiring a p_T of at least 0.7 GeV$/c$ at 5 MHz interaction rate and 1.4 GeV$/c$ at 10 MHz. It is clear from the figure that the rate of ECAL pretrigger seeded FLT tracks is too high for the SLT/DAQ to handle, even with a very high p_T threshold. By requiring a bremsstrahlung cluster to be found in the ECAL pretrigger, the rate is easily reduced to a few kHz and most of the available rate can be dedicated to the muon trigger. For simplicity we will use the same p_T cut for the electrons as for the muons in this analysis. The cut will have to be adjusted to be as low as possible, once the true rates are measured in the detector.

In Table 7.1 FLT trigger rates using the mentioned p_T thresholds are listed. The trigger rates for the FLT pair triggers and for the two pretrigger only configurations are given also. The pair trigger rates are of course much lower than from the corresponding single track trigger configurations. The SLT/DAQ will be able to handle the rate from both pair triggers at the same time. A possibility would be to use a single FLT-Muon track trigger together with the

Trigger configuration	5 MHZ IA rate	10 MHZ IA rate
Single FLT muon	29 kHz	28 kHz
Single FLT electron	1.6 kHz	2.3 kHz
FLT muon pair	650 Hz	3.6 kHz
FLT electron pair	3.8 kHz	15 kHz
Muon pretriggers only	195 kHz	425 kHz
Electron pretriggers only	770 kHz	1.6 MHz

TABLE 7.1: Output rates from the FLT trigger at 5 and 10 MHZ interaction (IA) rate using different trigger configurations. The "Single FLT-track" configuration require $p_T > 0.7$ GeV/c at 5 MHz and $p_T > 1.4$ GeV/c at 10 MHz to keep the input rates acceptable for the SLT/DAQ.

FLT-Electron pair trigger as the FLT-Electron pair trigger has higher J/ψ efficiency than the single FLT-Electron trigger with bremsstrahlung requirement, as we shall see in Section 7.3. At 10 MHz interaction rate, one would however have to reduce the FLT-Muon trigger rate further or to impose a p_T cut in the FLT-Electron pair trigger, in order to keep the total trigger rate at around 30 kHz. The pretrigger only rates are of course much too high to be handled by the SLT/DAQ and must be downscaled. For the ECAL pretrigger one can also impose higher E_T threshold directly in the pretrigger (all the numbers are given for an $E_T > 0.7$ GeV threshold). The configurations are only meant to be used for the FLT commissioning and the rate as a function of E_T threshold has not been studied in detail.

7.1.2 SLT rates

After having determined the thresholds needed to keep the input rate to the SLT at a manageable level, the output rate of the SLT and the rates at intermediate steps have to be estimated. The intermediate rates are needed for the estimation of the CPU time. The output rate of the SLT is limited by the rate with which the 4LT can reconstruct and log the data to tape. For triggered events the rate limit given by the reconstruction time is expected to be around 50 Hz. The data logging is able to run at more than 100 Hz and therefore imposes no limitation. Besides the triggered events, about 10 Hz of monitoring events, such as random triggered and SLT bypass events, will be logged. The random triggered events have on average a much lower occupancy than the SLT triggered events and are therefore much faster to reconstruct. The bypass events will most likely not be reconstructed online as they are mainly meant for SLT related analysis. We therefore use the 50 Hz, as the expected limit for SLT triggered events.

Table 7.2 gives the simulated output rates of the single FLT-Muon trigger at 5 and 10 MHz interaction rate. The naming of the individual steps should be straight forward from the description of the trigger algorithm in Chapter 6. The "distance cut" is the 40 cm cut applied when the tracks found in the MUON detector and the ones reconstructed by the FLT are combined into pairs. The rejection of events is more or less equally done in the different steps in the trigger algorithm, though at 10 MHz interaction rate, where events have higher occupancies and a higher p_T threshold is used, more of the tracks accepted in the first steps are ghosts. The algorithms called later have therefore higher rejection power at 10 MHz interaction rate compared to 5 MHz. The SLT output rates are given by just a few events out of the initial million events and therefore have large statistical and systematic uncertainties. The rate estimates are close

Trigger step	5 MHZ IA rate	10 MHZ IA rate
FLT	29 kHz	28 kHz
Muon tracking	10 kHz	15 kHz
Distance cut	7.3 kHz	10 kHz
Slicer	2.4 kHz	4.8 kHz
RefitX	520 Hz	780 Hz
RefitY	250 Hz	380 Hz
L2Magnet	250 Hz	380 Hz
L2Sili	105 Hz	80 Hz
L2Vertex	60 Hz	30 Hz

TABLE 7.2: Trigger rates in the SLT for the single FLT-Muon trigger.

to our limit of 50 Hz, however most of the accepted events have low invariant mass. If, at the output of the trigger algorithm, a muon pair is required to have an invariant mass of at least 2.0 GeV$/c^2$, the output rate is reduced to about 20 Hz at both 5 and 10 MHz interaction rate, without removing any $J/\psi \to \mu^+\mu^-$ events. Alternatively one can increase the p_T threshold in the FLT, which also reduces the load on the SLT. Increasing the p_T cut reduces however the $J/\psi \to \mu^+\mu^-$ efficiency and the next section will show that the SLT has plenty of CPU power to keep up with the full 30 KHz of FLT triggered events.

Table 7.3 shows the corresponding trigger rates for the single FLT-Electron trigger. While the SLT input rate is 15–20 times lower than from the single FLT-Muon trigger, the SLT output rates from the two triggers are of the same size. This is because one gets much better rejection from requiring a second track in the MUON system instead of a second ECAL cluster with $E_T > 0.7$ GeV. The output rate for the single FLT-Electron trigger is on the high side, but as in the muon case a simple 2.0 GeV$/c^2$ invariant mass cut, applied at the output of the SLT algorithm, can reduce the SLT output by about a factor of four. For the electron candidates we estimate their energy already at the beginning of the SLT using the ECAL pretrigger. This can be used to apply an invariant mass cut very early in the algorithm as was done in the year 2000 or to increase the E_T threshold imposed on the clusters. The SLT however does not need this rejection before the events are to be send to the 4LT. One of the problems during the 2000 run was that the high E_T threshold and an imprecise mass cut were shaping the background in such a way that a background peak appeared to the left of the $J/\psi \to e^+e^-$ peak in the invariant mass spectrum. This made it rather difficult to precisely estimate the amount $J/\psi \to e^+e^-$ events and should therefore be avoided for the 2002 run. A mass cut at the end of the SLT is therefore preferable.

The output rates for all six trigger configurations are listed in Table 7.4 without taking into account that the pretrigger only triggers would need to be downscaled at the input of the SLT. The output rates from the MUON pretrigger based configurations are more or less of the same magnitude, though the output rate from the FLT-Muon pair trigger at 10 MHz interaction rate would be too high for the SLT. Here an invariant mass cut at the output of the SLT would not be of much help, as it is already used at the FLT, though with lower precision. Instead one would have to apply a p_T cut on at least one of the muons or maybe apply in addition a particle identification procedure at the output of the SLT chain, but this has not been investigated.

The ECAL pretrigger based configurations cannot be directly compared, as the single FLT-

Trigger step	5 MHZ IA rate	10 MHZ IA rate
FLT	1.6 kHz	2.3 kHz
Second electron cluster	1.4 kHz	2.2 kHz
Distance cut	1.3 kHz	2.0 kHz
Slicer	420 Hz	1.0 kHz
RefitX	240 Hz	430 Hz
RefitY	190 Hz	280 Hz
L2Magnet	190 Hz	280 Hz
L2Sili	65 Hz	70 Hz
L2Vertex	40 Hz	40 Hz

TABLE 7.3: Trigger rates in the SLT for the single FLT-Electron trigger.

Trigger configuration	5 MHZ IA rate	10 MHZ IA rate
Single FLT-Muon	90 Hz	40 Hz
Single FLT-Electron	40 Hz	60 Hz
FLT-Muon pair	20 Hz	140 Hz
FLT-Electron pair	400 Hz	1 kHz
MUON pretriggers only	30 Hz	100 Hz
ECAL pretriggers only	6500 Hz	15 kHz

TABLE 7.4: Predicted output rates for the SLT (ignoring bandwidth/CPU limitations) for the different trigger configurations.

Electron trigger requires a bremsstrahlung cluster to be found for one of the ECAL pretrigger clusters and the two other configurations are not imposing this requirement. The FLT-Electron pair trigger is interesting because it has a relatively low input rate to the SLT and, as we shall see in Section 7.3, three to five times higher $J/\psi \to e^+e^-$ efficiency. Unfortunately the SLT algorithm is in this case only able to reject about 90% of the FLT triggered events. The 4LT would therefore be completely overloaded. Further studies will be needed to see if other cuts, which reduce the rate significantly, can be implemented. One possibility would be to reconstruct the ECAL clusters with the full reconstruction package, CARE, as was done for the 2000 run. This could be used to confirm the ECAL pretrigger clusters, giving a small additional rejection, and to apply E/p cuts after the momentum has been reconstructed by the SLT.

7.2 Timing

The time necessary for the SLT to process an event in the different steps in the trigger algorithm has been evaluated using the same Monte Carlo samples as for the trigger rates and using the rejections obtained for the corresponding sample. As the algorithm is used offline, the simulation does not measure the time spent waiting for the requested detector data to arrive from the SLB. This is expected to contribute with less than 1 ms to the total time per event. Most of this latency is coming from the read out of the FLT information and the MUON detectors. The simulation also does not include the time spent for the sparsification and hit preparation of hits

Trigger step	Time per inv. 5 MHz IA rate	Time per event 5 MHz IA rate	Time per inv. 10 MHz IA rate	Time per event 10 MHz IA rate
FLT/pret. decoding	0.32 ms	0.32 ms	0.32 ms	0.32 ms
Muon Tracking	0.19 ms	0.19 ms	0.33 ms	0.33 ms
0^0 hit request	1.2 ms	0.30 ms	1.4 ms	0.53 ms
0^0 hit prep.	0.40 ms	0.10 ms	0.55 ms	0.21 ms
Slicer	0.47 ms	0.12 ms	0.84 ms	0.30 ms
RefitX	0.41 ms	0.033 ms	0.69 ms	0.12 ms
Stereo view hit req.	1.2 ms	0.021 ms	1.3 ms	0.035 ms
Stereo view hit prep.	0.42 ms	0.007 ms	0.50 ms	0.013 ms
RefitY	1.6 ms	0.029 ms	1.9 ms	0.050 ms
Magnet hit req.	0.22 ms	0.002 ms	0.21 ms	0.003 ms
Magnet hit prep.	0.080 ms	0.001 ms	0.069 ms	0.001 ms
L2Magnet	0.22 ms	0.002 ms	0.16 ms	0.002 ms
VDS hit req.	0.28 ms	0.003 ms	0.25 ms	0.003 ms
VDS hit prep.	7.7 ms	0.076 ms	6.2 ms	0.084 ms
L2Sili	1.4 ms	0.014 ms	0.97 ms	0.013 ms
L2Vertex	0.039 ms	0.000 ms	0.037 ms	0.000 ms
Total	—	1.22 ms	—	2.02 ms

TABLE 7.5: Timing of the SLT trigger steps in the single FLT muon trigger. Timings are given for 1200 MHz CPU and using rates as given by the Monte Carlo at 5 MHz and 10 MHz interaction rate.

in the ITR and VDS, which is known to be quite long. For the VDS, the hit preparation time has been added using the times measured online in 2000 (scaled to a 1200 MHz CPU). For the ITR no hit preparation times have so far been measured. The hit preparation routine is however expected to be much faster than one used for the VDS, simply because it only uses the part of the readout chip which is actually needed.

The time estimates for the single FLT-Muon and single FLT-Electron triggers are given in Table 7.5 and 7.6. The CPU upgrade alone gives a factor of three improvement compared to the 2000 data taking. The rejection in the SLT 2002 version is however much smaller for the early steps. The speed of the later steps therefore becomes more important. To increase the speed, major improvements were done in RefitX and L2Sili, which have been changed to keep the number of intermediate track candidates low during the pattern recognition. For the electron trigger, a major improvement is also the removal of the ECAL read out and reconstruction, which was completely dominating the time spent in the 2000 run.

The major contributions to the total SLT latency are coming from the decoding of the FLT information and the hit request protocol. The former should in principle be simple, but has so far not been optimized for speed. During the hit request procedure a lot of time is spent calculating which detector regions to read out. This was noticed already during the 2000 run and it was planned to be improved, but apart for a few trivial optimizations, the upgrade has not been done. With the SLT input rates given in Table 7.1, the single FLT-Muon would need 35 SLT nodes at 5 MHz interaction rate or 60 SLT nodes at 10 MHz to process the 29 kHz expected input rate. The single FLT-Electron trigger would only needed 4 to 6 SLT nodes for the approximate 2 kHz of input rate. These numbers underestimate the real number of needed

Trigger step	Time per inv. 5 MHz IA rate	Time per event 5 MHz IA rate	Time per inv. 10 MHz IA rate	Time per event 10 MHz IA rate
FLT/pret. decoding	0.31 ms	0.31 ms	0.31 ms	0.31 ms
0^0 hit request	1.0 ms	0.81 ms	1.1 ms	0.95 ms
0^0 hit prep.	0.29 ms	0.23 ms	0.35 ms	0.30 ms
Slicer	0.40 ms	0.29 ms	0.52 ms	0.39 ms
RefitX	0.46 ms	0.12 ms	0.42 ms	0.18 ms
Stereo view hit req.	0.89 ms	0.13 ms	0.85 ms	0.16 ms
Stereo view hit prep.	0.20 ms	0.031 ms	0.23 ms	0.044 ms
RefitY	0.98 ms	0.15 ms	1.0 ms	0.19 ms
Magnet hit req.	0.16 ms	0.017 ms	0.16 ms	0.023 ms
Magnet hit prep.	0.046 ms	0.005 ms	0.051 ms	0.007 ms
L2Magnet	0.13 ms	0.014 ms	0.16 ms	0.023 ms
VDS hit req.	0.21 ms	0.023 ms	0.20 ms	0.028 ms
VDS hit prep.	4.8 ms	0.51 ms	3.9 ms	0.55 ms
L2Sili	0.67 ms	0.072 ms	0.76 ms	0.11 ms
L2Vertex	0.016 ms	0.001 ms	0.025 ms	0.001 ms
Total	—	2.71 ms	—	3.27 ms

TABLE 7.6: Timing of the SLT trigger steps in the single FLT electron with bremsstrahlung trigger. Timings are given for 1200 MHz CPU and using rates as given by the Monte Carlo at 5 MHz and 10 MHz interaction rate.

nodes as the latency of the switch is not taken into account, neither is the time needed for the event building of accepted events. It does anyhow show that, with the upgraded 240 SLT nodes, there is plenty of CPU capacity in the SLT system to handle unforseen complications.

7.3 Trigger efficiency

The trigger efficiency has been evaluated using 10,000 simulated $J/\psi \to \mu^+\mu^-$ and 10,000 $J/\psi \to e^+e^-$ events. These events have been mixed with varying amounts of overlapping inelastic interactions. The results for the single FLT-Muon and FLT-Electron triggers are given in Tables 7.7 and 7.8. The efficiency is factorized in a similar way as in the procedure in Chapter 5. The samples refer to the amount of inelastic interactions as follows:

1. Pure $J/\psi \to \ell^+\ell^-$ events (no tracks from underlying interactions).
2. $J/\psi \to \ell^+\ell^-$ events, no overlapping inelastic interactions.
3. $J/\psi \to \ell^+\ell^-$ events, 0.5 overlapping inelastic interactions.
4. $J/\psi \to \ell^+\ell^-$ events, 1.0 overlapping inelastic interactions.
5. $J/\psi \to \ell^+\ell^-$ events, 1.5 overlapping inelastic interactions.
6. $J/\psi \to \ell^+\ell^-$ events, 2.0 overlapping inelastic interactions.

Most of the different factors are defined as in Chapter 5, but a few have been altered/added to reflect the changes in the algorithm and the addition of new detector components:

Geometrical acceptance: Both of the leptons are still required to pass through two superlayers in the VDS, while in the main tracker they have to pass all six superlayers, as the SLT no longer allows empty superlayers. Depending on the number FLT tracks required in the trigger configuration, zero, one or two leptons are required to pass through the six superlayers in the OTR, while the rest can completely or partially pass through superlayers in the ITR also. For the muon trigger, the muons are required to cross all four superlayers in the MUON detectors and again depending on the trigger configuration, some of these superlayers can be passed in the gas pixel detector part of the MUON detector. For the electron trigger, both electrons are required to have deposited a cluster which is reconstructable with the CARE package and have an E_T of at least 0.7 GeV, matching the requirement in the ECAL pretrigger.

FLT: For the trigger configuration with the FLT included, the required number of leptons naturally must be found by the FLT. A lepton is considered to be found if an FLT track passes within a few centimeters of the true position of the track at both PC01 and TC02. The FLT efficiency has not been factorized[1] and therefore include both detector and algorithm effects. In fact from the table we see that the FLT efficiency increases with the detector occupancy as the FLT can compensate for inefficiencies by adding hits from other nearby tracks. For the single FLT track trigger configurations a p_T threshold of 0.7 GeV/c has been applied in the FLT, while for the FLT pair triggers an invariant mass cut of 2.0 GeV/c^2 is used instead.

Bremsstrahlung: For the single FLT-Electron trigger we require that the FLT track is seeded by an ECAL pretrigger cluster, where the ECAL pretrigger also found a cluster consistent with bremsstrahlung emitted before the magnet. The bremsstrahlung cluster is required to lie on the correct side of the main cluster as given by the charge of the FLT track, but we do not require that the cluster was actually caused by a photon emitted from the electron. The efficiency of having a "bremsstrahlung cluster" therefore increases with the occupancy. The efficiency of the ECAL pretrigger to bremsstrahlung clusters is however still lower than the probability for the offline reconstruction to find bremsstrahlung which is of order 30%. This is mainly because the pretrigger has to predict the position of the bremsstrahlung cluster from the ECAL energy, while the offline reconstruction predicts it based on the much more precise VDS track.

MUON pads coverage: This factor gives the efficiency for leptons not already found by the FLT to have hits in either the two MUON pad chambers or in the two MUON gas pixel chambers. The simulation does not include the inefficiencies which are known to be present in the MUON readout and which could potentially reduce the efficiency by up to 50%, see discussion below.

Pretrigger: This factor only includes the pretrigger efficiency for tracks not found by the FLT. For tracks found by the FLT, the pretrigger efficiency is included into the FLT efficiency.

[1]See [57, 67] for a detailed evaluation of the FLT efficiency

Slicer hits: Slicer (and the two Refits) considers now also hits in the ITR. Hits are required in all six superlayers and at least 12 hits from each lepton are required. The inefficiency observed is mainly due to an overestimate of the geometrical acceptance for tracks passing close to the detector limits.

L2Vertex: The L2Vertex factor gives the probability for the vertex reconstructed by L2Vertex from the two L2Sili tracks, to have a χ^2 less than 20. From Table 7.7 one observes that efficiency is almost 100% for pure $J/\psi \to \mu^+\mu^-$ events, but drops when the detector occupancy increases. This inefficiency belongs actually to L2Sili as the inefficiencies are caused by the assignment of wrong hits to the L2Sili tracks. These wrong hits can cause one of the track slopes to be wrong by a few mrad, which is enough for L2Vertex to reject the event. However L2Sili still accepted the track and if the vertex requirement was not used there would be no trigger inefficiency from this, and the track could be reconstructed correctly by the offline reconstruction. One solution, currently under study, is to use all of the acceptable track candidates from L2Sili instead of only the one with the highest number of hits.

Simulation results

The Tables 7.7 and 7.8 give the detailed factorized results for the single FLT track triggers with different amounts of overlapping events. Table 7.9 gives the efficiency for J/ψ events for all of the six trigger configurations with an average of one overlapping inelastic interaction and the efficiency factorized into the major parts of the trigger.

The estimated efficiencies are quite satisfactory. The inefficiencies in the muon tracking, present during the 2000 data taking, have been almost completely eliminated. The dependence on hit occupancy for the Refit algorithms has also been reduced a lot. The total trigger efficiency drops less than 10% within the investigated range of overlapping inelastic events. The largest SLT related inefficiency comes from setting up incorrect RoIs in the main tracker in the single FLT-Muon trigger. This is caused by large multiple scattering effects in the MUON system, which prompted us to use an average p_T assumption for setting up the RoIs. This procedure is difficult to improve without simply moving the inefficiencies into the Slicer and Refit algorithms. All other SLT efficiency factors are estimated to be above 95% and in total the SLT efficiencies for single FLT-Muon and FLT-Electron triggers are 70% and 75%, respectively[2], which is similar to the FLT tracking efficiency (excluding the bremsstrahlung requirement).

Table 7.9 shows that the SLT efficiency improves with the number of tracks required to be found by the FLT. This is partially because the SLT gets better and smaller RoIs from the FLT compared to the ones it gets from its own seeding, but mainly because most of the detector inefficiencies gets folded into the FLT efficiency. The total efficiencies given in the table cannot be directly compared to each other as the pretrigger only configurations need to be downscaled. At interaction rates above 5 MHz the single FLT track triggers also need to apply a p_T cut higher than the 0.7 GeV/c cut already applied for the efficiencies given in the tables. Figure 7.2 shows the trigger efficiency for the single FLT track triggers as a function of the p_T cut applied in the FLT. For muons the p_T cut at 0.7 GeV/c contributes very little inefficiency by itself and the efficiency drops relatively slow as the p_T cut is increased. For electrons the efficiency drops faster with the increase in p_T cut, since the energy loss through bremsstrahlung emission

[2]The "SLT efficiencies" quoted include inefficiencies due to underestimated geometrical acceptance.

Sample	1	2	3	4	5	6
Geo. acceptance	44.5%	44.5%	44.5%	44.5%	44.5%	44.5%
FLT	61.0%	67.5%	67.3%	69.1%	69.4%	71.5%
Pad coverage	98.5%	98.6%	98.6%	98.8%	98.7%	98.8%
Pretrigger	97.2%	97.1%	96.8%	95.8%	95.3%	94.8%
Muon hits	95.4%	94.8%	95.2%	96.1%	95.8%	96.3%
Muon tracking	97.7%	97.1%	96.8%	97.2%	97.3%	97.0%
Muon pairing	100.0%	99.9%	99.9%	99.9%	99.9%	99.8%
Slicer RoIs	92.5%	93.4%	93.4%	91.8%	92.4%	91.4%
Slicer hits	94.4%	96.3%	96.5%	96.8%	97.6%	97.1%
Slicer algorithm	100.0%	100.0%	99.9%	100.0%	100.0%	100.0%
RefitX hits	99.1%	99.0%	99.0%	98.8%	98.8%	98.3%
RefitX algorithm	99.4%	97.1%	96.0%	96.0%	95.1%	94.7%
RefitY hits	98.3%	99.0%	98.7%	98.5%	98.2%	98.0%
RefitY algorithm	99.2%	99.2%	99.1%	98.5%	98.7%	98.5%
L2Sili RoIs	99.6%	99.0%	98.5%	98.8%	97.9%	97.2%
VDS hits	96.3%	96.0%	96.2%	95.8%	95.8%	96.2%
L2Sili algorithm	99.7%	99.2%	99.1%	98.8%	99.0%	98.8%
L2Vertex	99.6%	97.7%	96.9%	95.7%	94.7%	94.6%
Total efficiency	19.4%	20.7%	20.1%	19.9%	19.6%	19.4%

TABLE 7.7: Trigger efficiency factors for the single FLT-Muon trigger with different amounts of overlapping inelastic events.

lowers the p_T, even if part of this loss is regained by the bremsstrahlung recovery in the ECAL pretrigger.

To give a better comparison of the potential of the different trigger configurations, the expected number of triggered J/ψ events per hour (J/ψ yield) is given in Table 7.10 for all six trigger configurations at 5 and 10 MHz interaction rate. A production rate of 3400 $J/\psi \to \ell^+\ell^-$ per hour per MHz interaction rate has been assumed, as it was done in Chapter 5. The pretrigger seeded SLT configurations have been downscaled to have the SLT input rates of 30 kHz. This would not allow the two trigger configurations to run in parallel. The J/ψ yield in the ECAL pretrigger configuration can be improved significantly by increasing the E_T cut in the pretrigger instead of simply downscaling the rate. However for FLT commissioning purposes, which will be the main purpose of this trigger configuration, it is most likely more useful to have tracks which resemble as much as possible the tracks seen by the FLT. For the commissioning it is also not necessary to run the two triggers configurations in parallel.

For the single FLT-Muon trigger the doubling in the J/ψ production rate when going from 5 to 10 MHz interaction rate more than compensates for the loss given by the higher p_T cut used at 10 MHz. This is therefore seen as the preferred running mode, though the single FLT-electron mode does not gain in rate when going from 5 to 10 MHz interaction rate. The FLT pair triggers have not been downscaled, even though the rate study in Section 7.1 showed that the output rates of the SLT in these configurations are too high for the 4LT to handle. The FLT-Muon pair trigger does anyway not have a better yield than the single track trigger. The FLT-Electron pair trigger would be significant better than the single FLT-Electron trigger,

Sample	1	2	3	4	5	6
Geo. acceptance	27.5%	27.5%	27.5%	27.5%	27.5%	27.5%
FLT	71.9%	74.9%	74.6%	74.5%	75.2%	74.3%
Bremsstrahlung	16.7%	19.0%	19.4%	20.5%	20.2%	21.5%
Pretrigger	97.6%	98.2%	98.0%	97.4%	97.8%	96.1%
Electron pairing	97.8%	97.7%	97.9%	96.1%	97.3%	96.2%
Slicer RoIs	99.1%	99.2%	99.2%	99.2%	98.5%	98.5%
Slicer hits	94.9%	97.8%	97.1%	97.4%	97.7%	97.5%
Slicer algorithm	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
RefitX hits	99.7%	98.9%	99.2%	97.6%	98.4%	98.7%
RefitX algorithm	97.3%	95.3%	95.6%	95.7%	95.2%	95.0%
RefitY hits	99.0%	98.5%	98.3%	99.4%	98.0%	98.6%
RefitY algorithm	98.6%	98.8%	98.5%	98.9%	98.9%	98.0%
L2Sili RoIs	98.9%	99.1%	98.8%	98.3%	98.6%	98.3%
VDS hits	96.8%	96.4%	97.9%	96.5%	96.2%	96.5%
L2Sili algorithm	97.4%	98.1%	97.9%	98.5%	98.2%	98.2%
L2Vertex	98.9%	98.1%	96.9%	97.2%	96.0%	95.1%
Total efficiency	2.6%	3.1%	3.1%	3.2%	3.1%	3.1%

TABLE 7.8: Trigger efficiency factors for the single FLT-Electron trigger with different amounts of overlapping inelastic events.

which has the very inefficient bremsstrahlung requirement. The bremsstrahlung requirement has the advantage that the triggered data sample gets much more pure and for most studies of the 2000 data the bremsstrahlung was always required offline anyway. It is however hoped that bremsstrahlung requirement will not be necessary for the 2002 analyses, due to the improvements in the calorimeter and the removal of the magnet chambers. More studies of the rejection factor in the FLT-Electron pair trigger are therefore needed.

The $J/\psi \to \mu^+\mu^-$ yields have not been corrected for the known inefficiency in the MUON pad readout which affects all three muon trigger configurations equally. For the 2000 data the effect was estimated to be about an extra 50% efficiency loss, but this was not in agreement with the observed $J/\psi \to \mu^+\mu^-$ yield in the 2000 run, which was significantly higher than the prediction. The inefficiency is therefore probably smaller and the $J/\psi \to \mu^+\mu^-$ yield will still be of the order of several thousand $J/\psi \to \mu^+\mu^-$ events per hour, a vast improvement over the 30 per hour observed in the 2000 run.

7.4 Tracking precision

The track resolution in Refit and L2Sili is given in Table 7.11 for muons from $J/\psi \to \mu^+\mu^-$ decays triggered by the single FLT-Muon configuration. The Refit track resolution is listed separately for tracks in the OTR and ITR, which have different hit resolutions. The resolution in the OTR is somewhat worse in the y direction than the one measured for the 2000 run as this is heavily influenced by the target constraint. For the 2000 run the tracking resolution was measured with a single target wire, while for 2002 a target box with two target wires is used.

Trigger configuration	1 FLT-μ	1 FLT-e	2 FLT-μ	2 FLT-e	μ-pret.	e-pret.
Geo. acceptance	44.5%	27.5%	33.8%	23.0%	45.9%	27.8%
FLT/pretriggers	65.3%	14.9%	27.5%	39.4%	93.4%	96.0%
Seeding	93.4%	96.1%	99.9%	99.9%	85.2%	98.4%
Slicer/Refit	81.8%	88.8%	92.3%	94.9%	70.4%	81.8%
Vertexing	89.5%	90.8%	90.8%	89.8%	87.9%	86.8%
Total efficiency	19.9%	3.2%	7.8%	7.7%	22.6%	18.7%

TABLE 7.9: Trigger efficiencies separated into the main factors for different trigger configurations.

Trigger configuration	1 FLT-μ	1 FLT-e	2 FLT-μ	2 FLT-e	μ-pret.	e-pret.
5 MHz IA rate	3400	540	1350	1300	580	125
10 MHz IA rate	4900	520	2700	2600	540	115

TABLE 7.10: Expected J/ψ yields at 5 and 10 MHz interaction rate for the different trigger configurations.

The tracking resolution in the VDS is improved in the new SLT by about 20–30% as was already shown in Section 6.6.1. The slope uncertainties are underestimated as no correction for the multiple scattering in material between the target and the first active layer has been included.

Table 7.12 gives the vertex resolution achieved by L2Vertex. As the precision of the L2Sili is approximately the same as for offline reconstruction the vertex resolution is also very similar to the resolution achieved offline, see Section 8.4.2. The vertex uncertainties are underestimated as the track slopes uncertainties are underestimated. The χ^2 cut in L2Vertex is therefore kept relatively loose at 20.

7.5 Summary

The performance evaluation in this chapter shows that large increase in J/ψ yield are expected in the 2002 run thanks to the improvements in detector performance, trigger algorithm and the inclusion of the FLT. For $J/\psi \to \mu^+\mu^-$ events, the yield per hour is increased by approximately two orders of magnitude over the values obtained during the 2000 run by using the single FLT-Muon trigger, which has been shown to be the most effective. For $J/\psi \to e^+e^-$ events, the improvements are more modest, in particular for the single FLT-Electron trigger. The FLT-Electron pair trigger could increase the yield by an order of magnitude over the 2000 run, but more work needs to be done to achieve the needed rejection for the SLT to be able to run in this configuration. For the 2002 run HERA – B expects to take trigger data for about 2000 hours, which with the single FLT triggers at 10 MHz would give J/ψ samples of $O(10^7)$ in $J/\psi \to \mu^+\mu^-$ channel and $O(10^6)$ in the $J/\psi \to e^+e^-$ channel. This sample is sufficient to carry out the planned physics program.

The efficiency and timing of the upgraded SLT algorithm are very satisfactory, while more studies of the rejection will be necessary once real data are available.

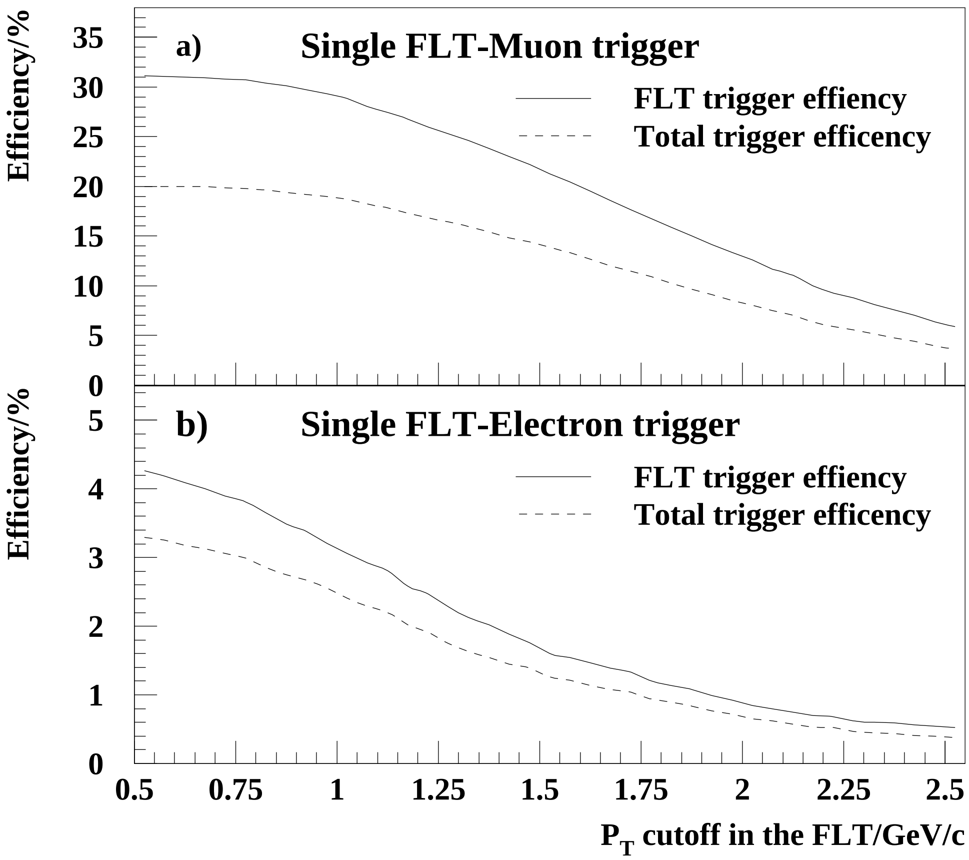

FIGURE 7.2: Trigger efficiency for the single FLT track triggers as a function of the p_T cut in the FLT. Both the FLT efficiency (including geometry) and the full (FLT+SLT) efficiency are shown.

	OTR	ITR	L2Sili
$\Delta x/\mu$m	700	100	12.5
$\Delta y/\mu$m	8400	1000	12.5
Δt_x/mrad	0.42	0.13	0.10
Δt_y/mrad	2.8	0.88	0.10
Pull(x)	1.3	1.1	0.99
Pull(y)	1.3	0.86	0.98
Pull(t_x)	1.7	1.7	1.31
Pull(t_y)	1.2	1.1	1.35

TABLE 7.11: Resolutions and pulls of Refit (in OTR and ITR) and L2Sili track parameters for the single FLT-Muon trigger.

	L2Vertex
$\Delta x/\mu$m	27
$\Delta y/\mu$m	28
$\Delta z/\mu$m	500
Pull(x)	1.35
Pull(y)	1.38
Pull(z)	1.42

TABLE 7.12: Resolution and pulls of the vertices found by L2Vertex for the single FLT-Muon trigger.

Chapter 8

Measurement of $\sigma(b\bar{b})$ from the 2000 data

With the J/ψ rates for 2002 estimated in the previous chapter, HERA – B will record about 10^7 J/ψ decays during the 2002 run. This sample will allow HERA – B to measure both the nuclear dependence of charmonium production and the $b\bar{b}$ production cross section, $\sigma(b\bar{b})$, with unprecedented precision. No data have been recorded yet this year and this chapter will therefore present a measurement of $\sigma(b\bar{b})$ based on the approximately 10^4 J/ψ events recorded by HERA – B between June and August 2000. The interest for measuring $\sigma(b\bar{b})$ has already been discussed in Chapter 1.

The analysis uses the known long lifetime of b-hadrons to separate $b \to J/\psi X$ decays from prompt J/ψ decays. To get the best possible precision, both the $J/\psi \to e^+e^-$ and $J/\psi \to \mu^+\mu^-$ decay channels are used to estimate the cross section. The background events in the two channels are quite different leading to significant differences in their analysis. I carried out the analysis for the muon channel, while the analysis for the electron channel was done by HERA – B members from NIKHEF and INFN, Bologna. For this reason the focus will here be on the muon channel even though it has the smallest statistical significance and therefore the smallest weight in the combined cross section measurement. A preliminary result of this analysis was presented in Ref. [68] and is about to be submitted for publication [69]. The full analysis of the muon channel will be presented in the following, together with a summary of the electron channel analysis and the combined result.

8.1 The principle of the measurement

As described in the previous chapters, HERA – B is designed to trigger on leptons from J/ψ decays. Almost all of the 2000 data were recorded with the di-lepton triggers described in Chapter 4. We use inclusive decays of b-hadrons to J/ψ ($b \to J/\psi X$) for measuring $\sigma(b\bar{b})$. The total trigger and reconstruction efficiencies are not known very well, neither is the luminosity[1]. The prompt produced J/ψ are kinematically very similar to J/ψ from b-hadron decays and are therefore detected with similar efficiency. The prompt cross section in proton–nucleus interactions is also rather well measured [71, 72]. By measuring the $b\bar{b}$ production relative to the prompt

[1] HERA – B was not designed for precise luminosity measurements. The best luminosity estimates [70] in HERA – B use random triggered event, which are only available in limited quantities for triggered runs.

J/ψ, the luminosity does not need to be measured and the systematic uncertainty due to the detection efficiency is greatly reduced.

The J/ψ from b-hadron decays are distinguished from prompt J/ψ by their long lifetime, which results in a decay vertex detached from the primary interaction vertex. The procedure for measuring $\sigma(b\bar{b})$ therefore consists of two steps:

1. The J/ψ events are selected and counted. These will be completely dominated by the prompt J/ψ.
2. The prompt J/ψ are removed by requiring a detached vertex and the number of J/ψ from b-hadron decays can be counted.

The calculations and external values needed for extracting $\sigma(b\bar{b})$ from the data are described in the following.

8.1.1 Calculating $\sigma(b\bar{b})$

The total $b\bar{b}$ production cross section, σ_B^A, depends on the target material, given by the number of nucleons A, and can be expressed as:

$$\sigma_B^A = \frac{n_B}{\mathcal{L} \cdot \varepsilon_B^{\text{tot}} \cdot \text{Br}(b\bar{b} \to J/\psi X) \cdot \text{Br}(J/\psi \to \ell^+\ell^-)}, \tag{8.1}$$

where n_B is the number of b events found, $\mathcal{L}$ the luminosity, $\varepsilon_B^{\text{tot}}$ the total efficiency of the b selection, including contributions from trigger, reconstruction and detached vertex selection. Similarly, the prompt J/ψ cross section is given by:

$$\sigma_P^A = \frac{n_P}{\mathcal{L} \cdot \varepsilon_P^{\text{tot}} \cdot \text{Br}(J/\psi \to \ell^+\ell^-)}, \tag{8.2}$$

where n_P is the number of prompt J/ψ events and $\varepsilon_P^{\text{tot}}$ is the total efficiency for this channel.

The analysis consists essentially of two parts: a prompt J/ψ selection in order to determine the number of J/ψ produced, followed by a selection of the detached $b \to J/\psi X$ events. Therefore the prompt J/ψ and b selection efficiencies are factorized as:

$$\varepsilon_B^{\text{tot}} = \varepsilon_B^{J/\psi} \cdot \varepsilon_B^{\Delta z}, \qquad \varepsilon_P^{\text{tot}} = \varepsilon_P^{J/\psi}, \tag{8.3}$$

where $\varepsilon_B^{J/\psi}$ and $\varepsilon_P^{J/\psi}$ are the J/ψ selection efficiency for $b \to J/\psi X$ and for the prompt J/ψ decays, while $\varepsilon_B^{\Delta z}$ is the efficiency of the detached vertex selection for the b-hadrons.

The cross section σ_B^A can be determined from the ratio of Equations 8.1 and 8.2:

$$\sigma_B^A = \sigma_P^A \cdot \frac{n_B}{n_P} \cdot \frac{1}{\varepsilon_R \cdot \varepsilon_B^{\Delta z} \cdot \text{Br}(b\bar{b} \to J/\psi X)}, \qquad \text{where } \varepsilon_R = \frac{\varepsilon_B^{J/\psi}}{\varepsilon_P^{J/\psi}}. \tag{8.4}$$

If the kinematics of the J/ψ from b-hadrons and prompt J/ψ were completely identical, ε_R would be equal to one and the normalization would not introduce systematic uncertainties. This is however not quite the case. The J/ψ kinematics is mainly described by the Feynman-x (x_F) and p_T of the J/ψ. The selection efficiency for J/ψ from either b-hadrons or prompt production is therefore given by:

$$\varepsilon_i = \int\int dx_F\, dp_T \, \frac{d\sigma_i}{\sigma_i dx_F dp_T} \varepsilon_i(x_F, p_T), \tag{8.5}$$

where $\frac{d\sigma_i}{\sigma_i dx_F dp_T}$ is the shape of the differential cross section. The efficiency $\varepsilon_i(x_F, p_T)$ depends on the production mechanism (b-hadron or prompt), because there are other things, such as the vertex position, which affects the efficiency. Even if we would have $\varepsilon_P(x_F, p_T) = \varepsilon_B(x_F, p_T)$ over the full kinematical range, the integrated efficiencies would still differ as the shape of the differential cross sections are somewhat different for the two sources of J/ψ mesons. If $\varepsilon_i(x_F, p_T)$ was constant we could could move it out of the integration and we would have simple cancellations of efficiencies in ε_R. This is approximately the case for the p_T dependence, see Figure 8.1, but the efficiency depends strongly on x_F as seen in the same figure. The best solution would be instead to measure the differential $b\overline{b}$ cross section by dividing the data sample in to small x_F bins where the efficiency is approximately constant. The low statistics available however does not allow this kind of study and we have rely on models to get the shape of the differential cross section. To reduce slightly the dependence on models, we will restrict the measurement of the $b\overline{b}$ production cross section to the x_F range where we actually have acceptance. The region we use is:

$$-0.25 < x_F < 0.15,$$

which covers the region where we have acceptance either for the $J/\psi \to e^+e^-$ or the $J/\psi \to \mu^+\mu^-$ channel. It would have been preferable to use a smaller region, where the efficiency is approximately constant as a function of x_F, but the loss of statistics with the 2000 sample when we use a more narrow x_F region would be significant.

For the cross section restricted to our x_F range, $\Delta\sigma_B^A$, we have the following expression:

$$\Delta\sigma_B^A = \Delta\sigma_P^A \cdot \frac{n_B}{n_P} \cdot \frac{1}{\varepsilon_R \cdot \varepsilon_B^{\Delta z} \cdot \mathrm{Br}(b\overline{b} \to J/\psi X)}, \tag{8.6}$$

where ε_R and $\varepsilon_B^{\Delta z}$ now are efficiencies for events in our x_F acceptance. Both efficiencies will be dependent on the models used to describe the differential spectra. The efficiencies will be estimated from simulated data based on Monte Carlo models of the differential spectra.

Cross sections in proton–nucleus interactions are usually parameterized as:

$$\sigma = \sigma_0 A^\alpha, \tag{8.7}$$

where σ is the cross section for an interaction channel in a proton–nucleus interaction, σ_0 is the cross section for the same channel in a proton–nucleon interaction (also denoted by $\sigma(pN \to \ldots)$) and α parameterize the influence of the type of nucleus on the reaction. For prompt J/ψ production α has been measured, see below, but for $b\overline{b}$ production no measurements exist. No nuclear suppression is expected in the open bottom or open charm hadron production and for D-meson production no such dependence has been observed [73]. We will therefore assume the formula $\sigma_B^A = \sigma(pN \to b\overline{b}) \cdot A^\alpha$, with $\alpha = 1$ to be valid. The final $b\overline{b}$ cross section, $\sigma(pN \to b\overline{b})$, can be evaluated as:

$$\Delta\sigma(b\overline{b}) \equiv \Delta\sigma(pN \to b\overline{b}) = \frac{\Delta\sigma_B^A}{A} = \Delta\sigma_r \cdot \frac{n_B}{n_P} \frac{1}{\varepsilon_R \cdot \varepsilon_B^{\Delta z} \cdot \mathrm{Br}(b\overline{b} \to J/\psi X)}, \tag{8.8}$$

with

$$\Delta\sigma_r = \Delta\varepsilon_P \cdot \sigma(pN \to J/\psi) \frac{A^\alpha}{A}, \tag{8.9}$$

where $\Delta\varepsilon_P$ is the fraction of J/ψ events inside our acceptance.

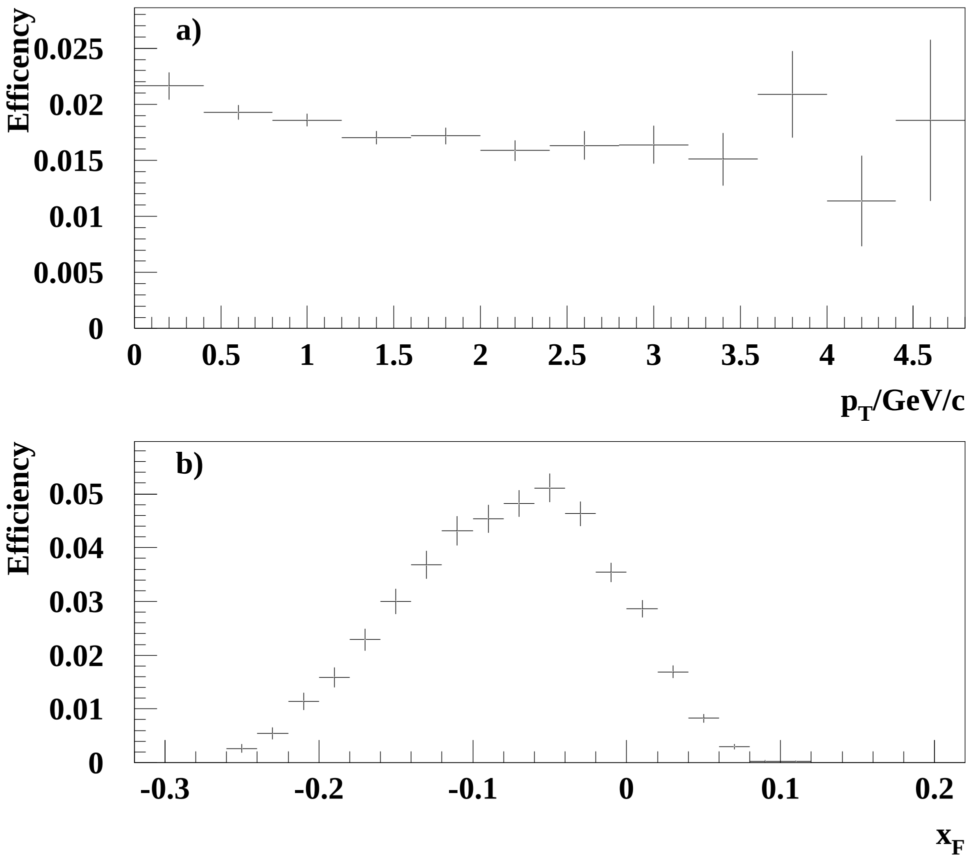

FIGURE 8.1: a) The $J/\psi \to \mu^+\mu^-$ selection efficiency as a function of p_T. b) The $J/\psi \to \mu^+\mu^-$ selection efficiency as a function of x_F. The $J/\psi \to e^+e^-$ channel (not shown) has acceptance up to $x_F \approx 0.15$.

8.1.2 Results from other experiments used by the analysis

The branching ratio $\mathrm{Br}(b\bar{b} \to J/\psi X)$ is the double of $\mathrm{Br}(b \to J/\psi X)$ since the two produced b quarks decay independently[2]. At the $\Upsilon(4S)$ resonance, CLEO and ARGUS measured $\mathrm{Br}(b \to J/\psi X) = (1.15\pm0.06)\%$ with a 50% B^0_d and 50% B^+ admixture [13], while LEP1 obtained $(1.16\pm 0.10)\%$ with 10% B^0_s and 10% b-baryons [13]. Since the b hadron admixture in hadronic collisions at HERA – B is expected to be similar the admixture to LEP, the value $\mathrm{Br}(b\bar{b} \to J/\psi X) = (2.32 \pm 0.20)\%$, obtained from the LEP1 result is used in this analysis.

The prompt J/ψ cross section (σ_P^A) has been measured at Fermilab in p-Au collisions by the E789 experiment [71] and in p-Si collisions by the E771 experiment [72], both fixed target experiments operating with a 800 GeV proton beam. Their measurements were done in different x_F ranges than ours (mainly positive x_F, see Table 8.1). The x_F dependence was measured and they extrapolated the results to the full range. We use the x_F dependence measured by E789 for

[2]In our measurement we only allow one $b \to J/\psi X$ candidate per event, so in principle we should correct for the events with two $b \to J/\psi X$ decays where both J/ψ decay into leptons. This effect is however completely negligible

	E789 [71]	E771 [72]
Target nucleus	Gold (A = 196.97)	Silicon (A = 28.09)
x_F acceptance	-0.035:1.0	-0.05:0.25
α value used	0.90 ± 0.02	0.920 ± 0.008
$\sigma(pN \to J/\psi)$ quoted (nb/nucleon)	$442 \pm 2_{stat} \pm 88_{sys}$	$375 \pm 4_{stat} \pm 30_{sys}$
σ_P^A (μb)	$51.3 \pm 0.1_{stat} \pm 10.2_{sys}$	$8.1 \pm 0.1_{stat} \pm 0.6_{sys}$
$\sigma(pN \to J/\psi)$ rescaled (nb/nucleon)	$331 \pm 12_{stat} \pm 67_{sys}$	$334 \pm 7_{stat} \pm 28_{sys}$

TABLE 8.1: Fermilab results on prompt J/ψ cross section in 800 GeV fixed target proton–nucleus collisions. In the last row, $\sigma(pN \to J/\psi)$ is calculated from the measured cross sections σ_P^A using the same value [58] $\alpha = (0.955 \pm 0.005)$ for both experiments.

our prompt J/ψ simulation and the extrapolations are therefore consistent. Both experiments expressed their results by parameterizing the nuclear dependence as in Equation 8.7, but they used different α values: E789 obtained $\sigma(pN \to J/\psi) = (442 \pm 2_{stat} \pm 88_{sys})$nb/nucleon with $\alpha = (0.90 \pm 0.02)$, while E771 measured $\sigma(pN \to J/\psi) = (375 \pm 4_{stat} \pm 30_{sys})$nb/nucleon with $\alpha = (0.920 \pm 0.008)$. More recently, the α value has been measured as a function of x_F and p_T by the E866 collaboration [58] using a very high statistics sample. The average value found in the x_F range of E771 and E789 is $\alpha = (0.955 \pm 0.005)$ [58]. After rescaling the two measurement results by the updated α value, they become more compatible, as shown in Table 8.1. The Fermilab cross section values is scaled to the HERA – B beam energy of 920 GeV according to the energy dependence given in Ref. [74], resulting in an upward shift of 7.3 %. The average value used as a reference for the prompt J/ψ cross section throughout the present analysis is therefore:

$$\sigma(pN \to J/\psi) = (357 \pm 8_{stat} \pm 27_{sys}) \text{ nb/nucleon}. \tag{8.10}$$

From this result, we can calculate the specific values of σ_P^A for both carbon and titanium target at HERA – B:

$$\begin{aligned} \sigma_P^C &= (3.84 \pm 0.08_{stat} \pm 0.27_{sys})\ \mu\text{b} \quad \text{and} \\ \sigma_P^{Ti} &= (14.4 \pm 0.3_{stat} \pm 1.1_{sys})\ \mu\text{b}. \end{aligned}$$

For $\Delta\sigma_r$ in Equation 8.8 an average of these are used, weighted according to the observed number of prompt J/ψ from the two target materials.

8.2 Data sample

Most of the HERA – B 2000 data taking was performed at a low interaction rate of 5 MHz, with a maximum of two target wires inserted in the beam halo in order to simplify the vertexing analysis. The wires used were "Inner 2" made of carbon and "Below 1" made of titanium. The present analysis uses events recorded using the "ECAL+SLT" and "MUON+SLT" trigger configurations described in Chapter 4. The events recorded with the "ECAL+L2Sili" trigger are not used, as the detector during runs taken with that configuration were quite unstable. A large fraction of the 'ECAL+L2Sili" runs for instance does not have any OTR information at all[3], and the OTR is essential for the background rejection.

[3] The OTR was turned off during those runs for investigations into high voltage failures

In total HERA – B collected ~450,000 di-muon and ~905,000 di-electron candidates with the two trigger configuration after runs with poor data quality have been removed. The data used for this analysis corresponds to 169 runs acquired in a time period of about 6 weeks. The accepted trigger events have been reconstructed several times using the 4LT farm to take advantage of improvements in the reconstruction code and in calibration and alignment constants[4].

8.3 Monte Carlo simulation

8.3.1 The event generators

The Monte Carlo generator in HERA – B is built using standard particle physics tools. Since no particular generator covers the full range of physics processes seen at HERA – B our Monte Carlo event generator is built using routines from several generators. For heavy flavor (b or c) quark production we use the PYTHIA package [75], which contains (at least partially) the hard scattering processes that result in heavy quarks and simulates the decay of mesons containing heavy quarks. PYTHIA is not able to simulate the effects given by the nuclear environment or the soft physics involved in the ordinary inelastic events. For the nuclear effects and inelastic events we use the FRITIOF package [76], which, like PYTHIA, is based on the Lund model of string fragmentations, but has been tuned to treat effects appearing in nuclear collisions, such as energy loss and multiple scattering of partons crossing the target nucleus.

The prompt J/ψ or the $b\bar{b}$ production is simulated, first, by generating the basic process $pN \rightarrow Q\bar{Q}X$ (where $Q\bar{Q}$ are heavy quarks) with PYTHIA and letting the heavy quarks hadronize. The energy left in the X part of the process is given as an input to FRITIOF to simulate further interactions taking place inside the nucleus. The generated particles are passed afterward to the GEANT 3 package [77] to simulate the behavior in the HERA – B detector using the geometry description of the detector.

The Monte Carlo generator produces exactly one interaction per event, while in reality the number of interactions per event is given by a Poisson distribution, whose mean depends on the interaction rate used. The simulations therefore allow the user to overlay several Monte Carlo events that have passed through GEANT, into a single event, which then is digitized and reconstructed. For this analysis, which uses data recorded at 5 MHz interaction rate, a simulated heavy flavor event is overlayed with simulated inelastic events, the number of which are distributed according to a Poisson distribution with a mean of 0.5 inelastic events.

8.3.2 Prompt produced J/ψ events

In the version of PYTHIA considered here (5.7), the physics processes leading to direct charmonium (and bottomonium) production are only partially taken into account. In order to have a prompt J/ψ simulation with the proper kinematic characteristics, the prompt J/ψ differential cross sections $d\sigma/d(p_T^2)$ and $d\sigma/dx_F$ measured at E789 [71] are used to correct the standard generator. To each event produced, one assigns a weight [78] which forces the x_F and p_T Monte Carlo distributions for prompt J/ψ to agree with the experimental data. In this analysis, all quoted Monte Carlo prompt J/ψ efficiencies are obtained using the E789 weighting of the events. Note that E789 only measured the cross section for x_F larger than -0.035, while we use events

[4]The final analysis is performed on the so-called "repro2" data sample.

with x_F as low as -0.25. The spectrum is expected [79] to be symmetric in x_F and that has been assumed in our simulation.

This reweighing procedure is equivalent, at the percent level, to the introduction of Color Octet Model (COM) processes [79] of quarkonium productions in the PYTHIA generator. The contribution to the systematic uncertainty given by the prompt J/ψ x_F and p_T spectra is estimated both by varying the spectra within the measurement errors of E789 and by using instead the result from the introduction of the COM in PYTHIA.

8.3.3 Simulated $b \to J/\psi X$ events

The PYTHIA version in the HERA – B generator is quite old and the $b\bar{b}$ production is only simulated to leading order. As our detection efficiency depends on the x_F and p_T of the J/ψ from b-hadron decays, our results are affected by the x_F and p_T spectrum of the initial b quark. To get a more precise estimate we therefore reweight the $b\bar{b}$ Monte Carlo in x_F and p_T similar to what was done for the prompt J/ψ. No measurements of x_F and p_T spectra exist for b-hadrons in proton–nucleus collision at our c.m.s. energy and instead we will use weights based on a theoretical calculation.

Our model is based on a NLO calculation [80] of the differential b quark cross section in x_F and p_T. The calculation has been updated compared to the Reference [80] to use the latest MRST parton distribution functions [81] for protons and neutrons. A b quark mass of $m_b = 4.75\ \mathrm{GeV}/c^2$ has been used and the QCD renormalization and factorization scale was set to $\mu = \sqrt{m_b^2 + p_T^2}$. The intrinsic transverse momenta of the interacting partons are distributed according to a Gaussian distribution set to give $\langle k_T^2 \rangle = 0.5(\mathrm{GeV}/c)^2$ [82]. The fragmentation of the b quark into a b-hadron is simulated according to a Peterson function [83] with $\epsilon = 0.006$. The resulting b-hadron x_F and p_T spectra are used to reweight the b hadron events generated by PYTHIA. The b-hadron flavor composition and decay chain is still given by PYTHIA with the average b-hadron lifetime updated to latest measurement $\tau_b = 1.564 \pm 0.014$.

To estimate the uncertainty introduced by using this model, several of the parameters have been varied within reasonable limits:

- The MRST parton distributions were replaced by the CTEQ5 [84].
- The b quark mass was varied between 4.5 and 5.0 GeV/c^2.
- The QCD renormalization and factorization scale is varied from: $\mu = 0.5\sqrt{m_b^2 + p_T^2}$ to $\mu = 2.0\sqrt{m_b^2 + p_T^2}$.
- The ϵ parameter of the Peterson function was varied between 0.002 and 0.008 corresponding to the different results of fits in Refs. [85, 86], measured in $Z \to b\bar{b}$ decays. The Peterson function was also replaced by the parameterization of Kartvelishvili [87] with $\alpha_\beta = 13.7 \pm 1.3$ [85].

The largest variations are seen when changing the parton distribution. The total systematic uncertainty is estimated to be 5%. The detection efficiencies are mainly affected by variations in x_F, as the acceptance in p_T is almost flat. The resulting x_F spectrum for J/ψ from b-hadrons with the default model and the model with changed parton distributions can be seen in Figure 8.2 together with prompt J/ψ spectrum, simulated as described in the previous section.

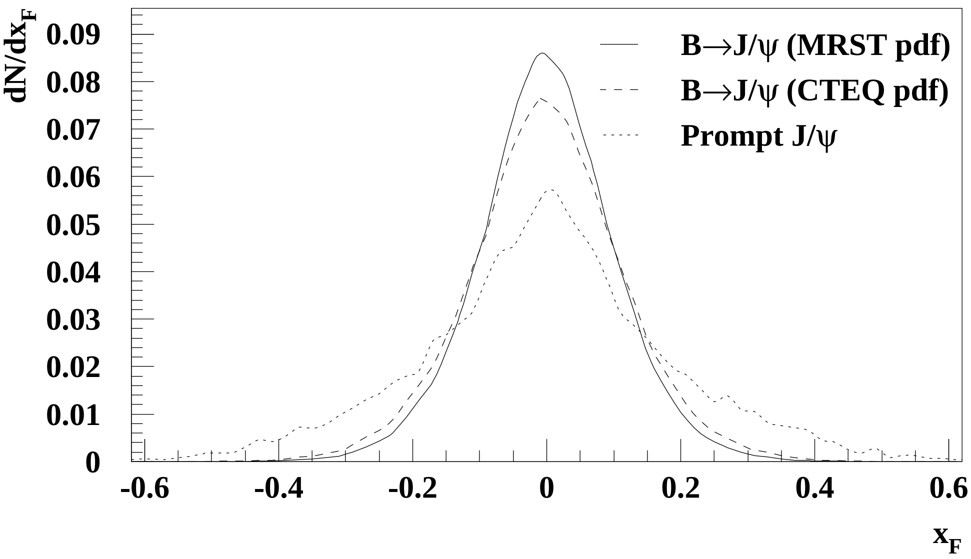

FIGURE 8.2: The J/ψ x_F spectra (normalized to unit integral) for our $b\bar{b}$ production model (full line), the $b\bar{b}$ production model with CTEQ5 parton distributions (dashed line) and for prompt produced J/ψ (dotted line).

8.3.4 Monte Carlo samples

We have used the Monte Carlo generator described above to generate a total of 600,000 prompt $J/\psi \to \mu^+\mu^-$ decays and 200,000 $b \to J/\psi X$ decays with the J/ψ decaying into muons. For the electron channel 5.4 million[5] $J/\psi \to e^+e^-$ and 150,000 $b \to J/\psi X$ decays were generated. 300,000 inelastic events have also been generated for overlaying heavy flavor events with inelastic interactions.

For studying the background given by semileptonic, open charm and beauty decays we furthermore generated about $100 \cdot 10^6$ open charm and $25 \cdot 10^6$ open beauty events, which at generator level were required to contain at least two muons(electrons). For most of these samples the leptons were furthermore required at the PYTHIA level to have p_T's higher than a certain threshold, in order to enhance the fraction of events accepted by the trigger simulation. For the heavy flavor background simulations no reweighting of heavy mesons was carried out.

All the events were subjected to the full detector simulation using the detector performances determined from run 16665 as described in Chapter 5. The "SLT+MUON" and "SLT+ECAL" triggers were simulated as was done in Section 5.4 and 5.3. To estimate the influence of the detector performance, the same events were also simulated with ideal detector performances. A variance of less than 3% in ε_R for the muon channel and 5% for the electron channel was observed, while $\varepsilon_B^{\Delta z}$ showed no change within the statistical uncertainty of the Monte Carlo generation. The variance of ε_R is used to evaluate the systematic uncertainty.

All Monte Carlo have been reconstructed and analyzed with the same reconstruction software

[5]On part of this sample kinematical cuts were imposed, before the GEANT simulation was applied, in order to speed up the generation process.

as used in the real data analysis.

8.4 Muon analysis

The measurement of $\sigma(b\bar{b})$ essentially consists of two steps: first the amount of prompt J/ψ events are determined, secondly the number of $b \to J/\psi X$ events is measured by requiring that the J/ψ vertex does not coincide with the primary interaction. From the ratio of these two numbers the $b\bar{b}$ production cross section can be calculated as described in Section 8.1. In the muon channel measuring the number of prompt J/ψ is quite straightforward as a clear J/ψ peak is seen already at the trigger level as for instance in Figure 5.20. The long lifetime of b-hadrons and the high relativistic boost of a fixed target experiment mean that a big fraction of the b-hadrons decay several millimeters or even a few centimeters from the primary interaction vertex. The main b-hadron cut to separate the few b events from the thousands of prompt J/ψ events is therefore to require a minimum decay length for the J/ψ. Some tracks from the primary interaction are invariably reconstructed slightly wrong, which can lead to a secondary vertex, which has a non-zero decay length. To suppress this background, we therefore impose cuts that ensure that both muons are incompatible with having been produced directly in the primary interaction. After these requirements, mainly background from other physics processes with long lived mesons such as semileptonic charm and beauty decays are left. These are of course not expected to produce peaks in the J/ψ invariant mass region. The semileptonic beauty decays are considered as background and not as an additional channel for measuring $\sigma(b\bar{b})$, since the kinematics is quite different from the J/ψ decay case and the normalization of efficiencies is therefore very problematic. Furthermore semileptonic beauty decays are difficult to distinguish from the semileptonic charm decays and other kind of background as they do not have any clear signature, such as a peak in the invariant mass spectrum, but are only characterized by having leptons not coming from primary interaction.

The details of the analysis are described in the following sections.

8.4.1 Prompt $J/\psi \to \mu^+\mu^-$ selection

The muon track selection is based on candidates found by the trigger system. The trigger system requires two muon tracks, but has no requirement on the invariant mass of the pair. Therefore it might sometimes trigger on a muon from a J/ψ and a muon from a decay in flight or, in the case of a $b \to J/\psi X$ decay, on one muon from the J/ψ and one muon from semileptonic decay of the other b hadron. The effect is significantly different for prompt J/ψ Monte Carlo, with less than 3% of events not triggered on both muons from the J/ψ, and for the $b \to J/\psi X$ Monte Carlo, where the effect is of 19%. The main part of the J/ψ selection efficiency is given by the trigger system and if one was to try to recover these events by introducing the muons not found by the trigger system, significant differences between the prompt J/ψ selection and the selection of J/ψ from b-hadrons would be seen. These differences introduce large uncertainties in ε_R and this is avoided by only using tracks which are matched with trigger tracks. The match is done using the χ^2 method described in Chapter 5 for both the VDS and OTR.

A fit of the di-muon invariant mass spectrum with a Gaussian signal and an exponential background gives 3840 ± 85 prompt J/ψ decays as shown in Figure 8.3. The invariant mass is shifted by 16 MeV/c^2 and the width of the J/ψ peak, 55 MeV/c^2, is about 35% larger than is expected from simulation. The differences are understood to be due to non-optimal alignment.

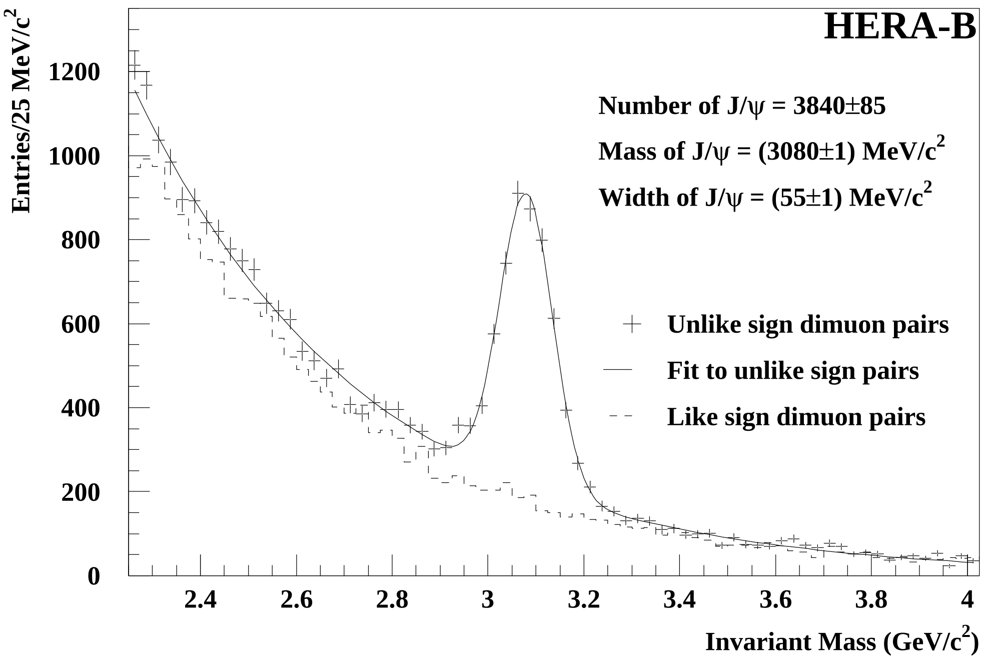

FIGURE 8.3: The di-muon invariant mass spectrum before applying the off-line particle identification.

The like sign spectrum, shown in Figure 8.3, is obtained using the same input events as the ones used for the unlike sign spectrum. The small discrepancy in the number of reconstructed events arises from the difference in trigger acceptance between the two cases and due to small physics contributions to the unlike sign spectrum (Drell-Yan, open charm production) not present in the like sign spectrum.

Monte Carlo studies show (see Section 8.4.2) that in the selection of detached J/ψ, the amount of badly reconstructed prompt J/ψ is insignificant compared to non-J/ψ background. The non-J/ψ background can be reduced by particle identification cuts and other kinds of cuts, besides the requirement of a detached vertex. To find the best set of cuts one should consider all the possible cuts in parallel, as the particle content in the background might be different in events with detached vertices. Given our low statistics we cannot do this. In the end it turns out that most of the background in the detached event selection most likely is coming from semileptonic charm and beauty decay. These are not suppressed by the particle identification cuts. The cut selection is therefore probably not the most optimal possible, but was kept unchanged to avoid biasing the final results by choosing cuts, which artificially enhance the significance of $b \to J/\psi X$ peak.

We choose the $J/\psi \to \mu^+\mu^-$ selection cuts based on the significance of the J/ψ signal using the $S/\sqrt{B}$ as measured from fits to a Gaussian signal and an exponential background in the invariant mass spectrum.

We investigated the following parameters:

1. Muon particle identification likelihood as given by the MUON system.

Parameter	Cut value	Eff. in data	Eff. in b MC	Eff. in Prompt MC
Trigger · Reconstruction	—	—	2.39%	2.20%
Truly triggered on J/ψ	—	—	80.7%	96.9%
MUON likelihood	> 0.4	88%	87.8%	88.3%
RICH likelihood	> 0.01	90%	95.5%	94.4%
$P(\chi^2)$	> 0.005	95%	97.5%	98.1%
Selection efficiency	—	75%	81.8%	81.7%
Total efficiency	—	—	1.59%	1.75%

TABLE 8.2: Cuts and efficiencies for the $J/\psi \to \mu^+\mu^-$ selection in data and Monte Carlo. The efficiencies are given for J/ψ events inside our x_F acceptance ($-0.25 < x_F < 0.15$) and no correction has been applied for the MUON pad efficiency or deadtime which affect the results in both channels equally.

2. Muon particle identification likelihood as given by the RICH system.
3. Probability of the di-muon vertex fit.
4. Kaon particle identification likelihood as given by the RICH system.
5. Hit occupancies in the VDS, RICH and OTR detectors.
6. Number of hits used by muon tracks in the OTR and/or the VDS.

After the first three cuts have been applied, no further significant improvements are seen and cuts 4–6 are not used. Figure 8.4 shows the significance of the J/ψ signal as a function of cut values. In Figure 8.5, the J/ψ selection efficiency in data is compared with both prompt J/ψ and $b \to J/\psi X$ Monte Carlo samples, as a function of the cut value. We observe that there is very little difference between the two Monte Carlo samples, but that they only reproduce the efficiency in data fairly well for the likelihood cut given by the muon system. The difference between data and Monte Carlo with the RICH likelihood cut can be at least partly explained by a higher hit occupancy in the real data and mis-alignment of the tracking chambers, both making the particle identification less efficient. The discrepancy in the vertex fit probability between data and Monte Carlo is related to the better resolution obtained in Monte Carlo (see comments in Section 8.4.2) and to underestimated errors in data.

The final selection, quoted in Table 8.2, has an efficiency of 75% with more than a factor three improvement in signal-to-background, which leaves $n_P = 2880 \pm 60$ J/ψ event. The distribution of the events in wire-configurations can be seen in Table 8.3 and the Monte Carlo samples have been weighted accordingly. The resulting invariant mass spectrum is shown in Figure 8.6. As can be seen from Figure 8.5 and Table 8.2, the selection introduces no new asymmetries between prompt and b efficiency, except the kinematic difference already seen by the trigger. We obtain $\varepsilon_R = 0.90 \pm 0.02$, where the error is due to Monte Carlo statistics.

8.4.2 Detached $J/\psi \to \mu^+\mu^-$ selection

Decay length determination

The decay length of J/ψ candidates is defined as the distance along the z-axis from the middle of the target wire to the J/ψ vertex. The distance flown in the transverse plane can be safely

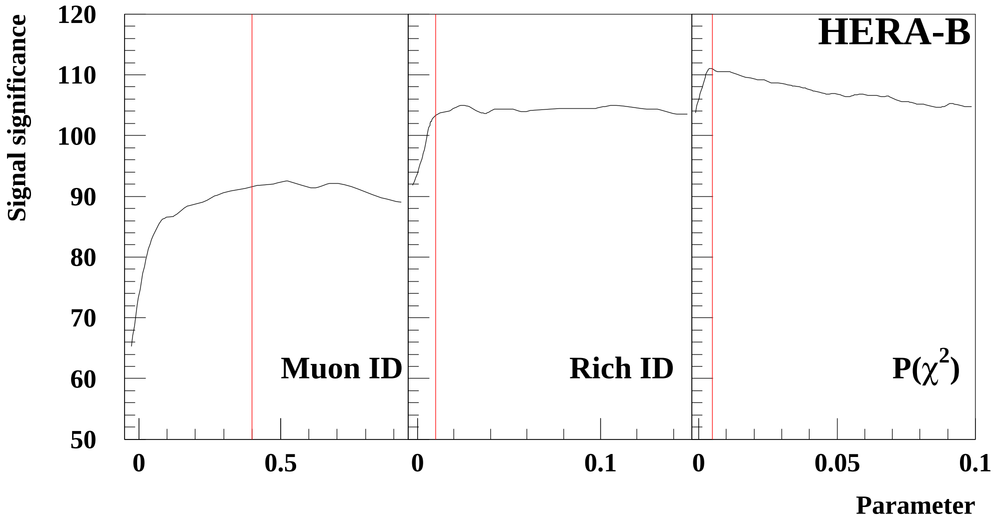

FIGURE 8.4: The significance ($S/\sqrt{B}$) of the J/ψ signal as a function of different cuts on muon identification likelihoods in the MUON and RICH system and the probability of J/ψ vertexing. Cuts have been applied sequentially from left to right. The vertical lines are drawn at the selected cut values. The parameters are slightly correlated, therefore only the last cut appears to be at the maximal significance.

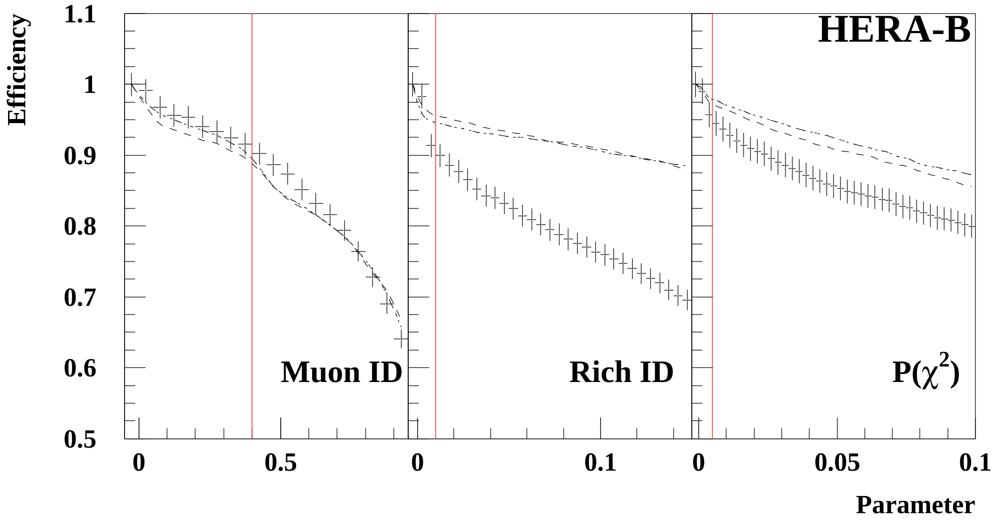

FIGURE 8.5: The J/ψ efficiency in data (error bars), prompt J/ψ Monte Carlo (dot-dashed line) and $b \to J/\psi X$ Monte Carlo (dashed line). The efficiencies are plotted as function of different cuts on muon identification likelihoods in the MUON and RICH system and the probability of J/ψ vertexing. The vertical lines are drawn at the selected cut values.

Target wire	Wire size ($\Delta x \times \Delta z$)	# of wires	# of $J/\psi \to \mu^+\mu^-$	Fraction
Carbon (Inner 2)	100μm$\times 1000\mu$m	1	1467 ± 42	51 %
Carbon (Inner 2)	100μm$\times 1000\mu$m	2	631 ± 27	22 %
Titanium (Below 1)	50μm$\times 500\mu$m	2	782 ± 31	27 %

TABLE 8.3: Number of selected $J/\psi \to \mu^+\mu^-$ in the different wire configurations used during the HERA – B 2000 data taking.

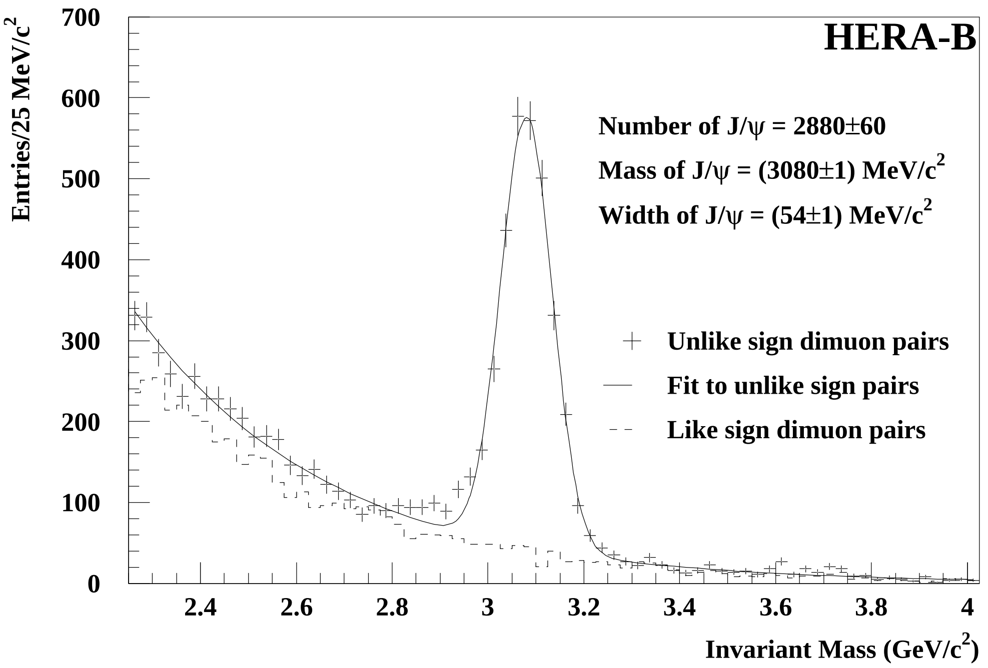

FIGURE 8.6: The di-muon invariant mass spectrum after J/ψ selection cuts.

ignored due to the much higher momentum along the beam axis. The decay length is calculated with respect to the wire position instead of the position of the reconstructed primary vertex, in order to give better precision and to avoid tails from badly reconstructed primary vertices.

The distance between the J/ψ vertex and primary vertices on the wire assigned to the J/ψ is calculated in the transverse plane. We require a primary vertex to be found within a maximum distance of 500 μm. From simulation we expect that 2% of the b events will have no primary vertex reconstructed and 10% will be lost because the distance is larger than 500 μm. The distance cut mainly affects the long lived b events and after applying the final detached selection cuts given in Section 8.4.2, the 500 μm distance cut removes 18% of the b events. At the same time the cut reduces the background by more than a factor of four.

In runs with multiple wires, one can assign the wrong wire or primary vertex to a primary J/ψ candidate and thus make it look detached. The above described method avoids wrong assignments completely with less than 0.1% wrongly assigned wires for b Monte Carlo and less than 0.01% for prompt Monte Carlo.

Detached vertex parameters

The following parameters are used for the selection of detached events:

1. Decay length divided by the uncertainty of the J/ψ vertex z position.
2. Impact parameter of muon tracks to the primary vertex.
3. Impact parameter of muon tracks to the assigned wire.

The prompt J/ψ which survive the detached vertex cut cannot be distinguished from $b \to J/\psi X$ events. The prompt J/ψ background must therefore be estimated from simulation and it is essential to check that the simulation describes the data well. To study the above parameters, we assume that the background under the J/ψ is described by the events in the sideband $2.1\,\mathrm{GeV}/c^2 < m < 2.8\,\mathrm{GeV}/c^2$, which we subtract from the signal region, defined as a 3σ region around the observed J/ψ mass.

Decay length The J/ψ decay length is shown in Figure 8.7 for the two target wires used during the HERA – B 2000 data taking for both data and Monte Carlo. Almost all of the J/ψ are produced at the wire and the spread in the decay length is therefore a convolution of the J/ψ vertex resolution and the wire thickness in the z direction. We observe that the vertex resolution is worse in data than Monte Carlo. After removing the wire contribution, the difference is about 20%. The wire position is known very well as it is an average over thousands of events and therefore the track resolution must be 20% worse in data. The same discrepancy is seen for the impact parameters to the wire in Figure 8.10 and we therefore smear the tracks in the Monte Carlo by 20%. We checked that the momentum spectrum and the number of hits per track are approximately the same in data and Monte Carlo, hence the worse resolution in data must be due to a worse hit resolution, either due to alignment or to the intrinsic hit resolution. The vertex resolution in the smeared Monte Carlo can be described by a sum of two Gaussians of weights 53% and 47% with widths 670 μm and 310 μm respectively. The weighted average vertex resolution in z for $J/\psi \to \mu^+\mu^-$ is therefore approximately 500 μm.

The estimated uncertainty on the secondary vertex position, σ_z, as given by the covariance matrix of the vertex fit, is only 5% larger in data than in the unsmeared Monte Carlo. For the unsmeared Monte Carlo, σ_z describes the vertex resolution very well, which means σ_z and hence the covariance matrices for the track parameters are underestimated in data. In the smeared Monte Carlo we therefore only increase the vertexing errors by 5% in order to get a reasonable simulation of the vertex separation, $\Delta z/\sigma_z$, which will be our main variable for selecting detached events. The resulting distribution can be seen in Figure 8.8. Note that the width of $\Delta z/\sigma_z$ would not be unity even if σ_z was correctly estimated, since the error does not include the finite size of the wire. Usually one would include a term of $l/\sqrt{12}$ in the uncertainty (the primary vertices are uniformly distributed in the wire), where l is the length of the wire in z. However we know that the primary interactions are confined to the wire, and including the width of the wire as an Gaussian distributed uncertainty would therefore weaken our separation power.

In order to study how well the simulation describes the tails in the decay length distribution of the prompt J/ψ and to see if there is any forward-backward asymmetry in the prompt J/ψ,

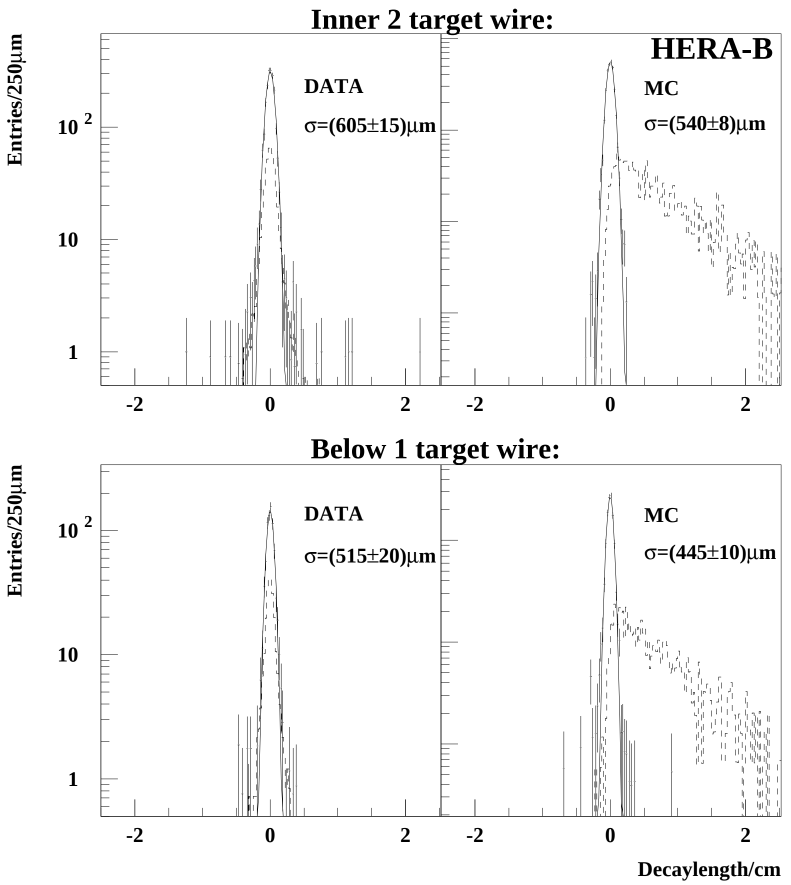

FIGURE 8.7: The decay length of J/ψ events in data and Monte Carlo. The background has been subtracted in the data plots. The Monte Carlo has to be smeared by 20% to match the observed resolution. The subtracted background is superimposed on the data plots (dashed line), while the spectrum for b hadrons is superimposed on the Monte Carlo plots (Not at the correct ratio!).

we studied the events which have an absolute decay length larger than three times σ_z. This leaves approximately 5% of the prompt J/ψ, which is enough to allow us to fit the invariant mass spectra and ensure that we are not affected by the presence of b events. In data we get 58 ± 11 J/ψ downstream of the target and 60 ± 10 J/ψ upstream, while from Monte Carlo, normalized to the total prompt signal, we expect 58 ± 3 downstream and 64 ± 3 upstream. No upstream–downstream asymmetry or deviation from simulation are therefore seen at the 5% level. Increasing the decay length cut prevents us from fitting the number of J/ψ in data, because of too small $\frac{\mathrm{S}}{\mathrm{B}}$ and due to the enhancing number of b events at positive decay length. Nevertheless, no significant forward-backward asymmetry is seen at any point in the prompt J/ψ Monte Carlo.

Muon impact parameters We consider both the impact parameter to the primary vertex and to the target wire. The impact parameter is here defined as the distance from the track

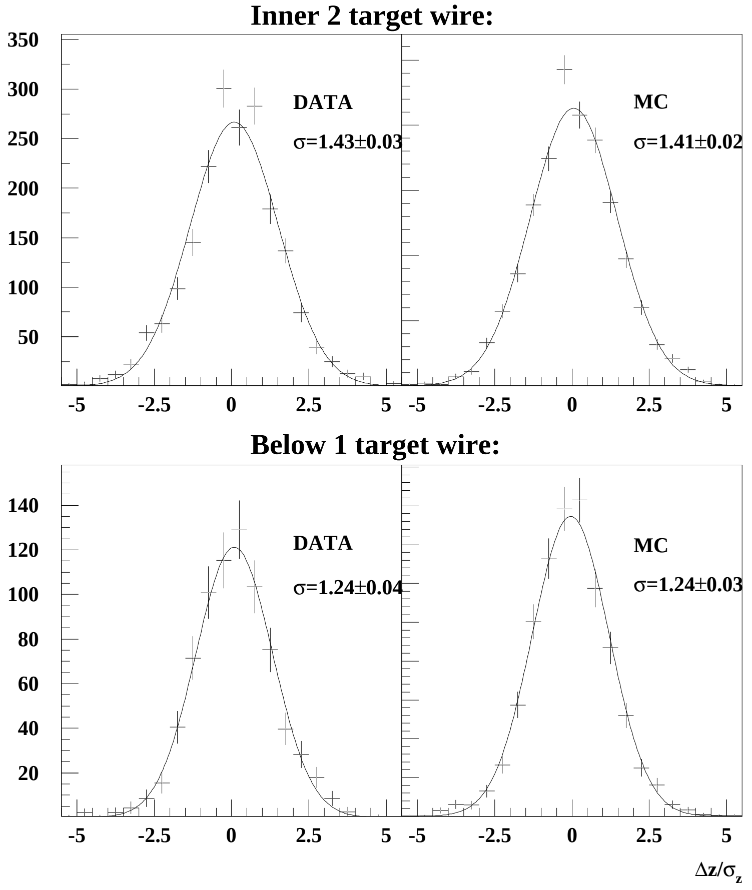

FIGURE 8.8: Vertex separation ($\Delta z/\sigma_z$) for prompt J/ψ in data and smeared Monte Carlo.

at the z of the primary vertex/wire to the vertex/wire. The impact parameter to the primary vertex gives the best separation between b signal and prompt J/ψ background as it is calculated in two dimensions. A soft cut on the impact parameter to the wire can remove background events with two primary interactions on the same wire, where one of the primary vertices was not reconstructed.

In Figure 8.9 the minimal impact parameter of the two muons to the assigned primary vertex is shown. Satisfactory agreement is observed comparing the shape in data and Monte Carlo with smeared tracks, as described above for the decay length.

The impact parameter cut to the target wire has to be analyzed separately for each wire since the Inner 2 wire is 100 μm thick while the Below 1 wire is only 50 μm thick. Also in this parameter we observe a 15-20% discrepancy between data and unsmeared Monte Carlo, as can be seen on Figure 8.10. After applying the smearing, Monte Carlo and data agree within errors. Since the Inner 2 wire distribution is 30% wider than the Below 1 wire and since we are cutting at just above one sigma, we are scaling the cut such that it is 30% larger for events assigned to the Inner 2 wire.

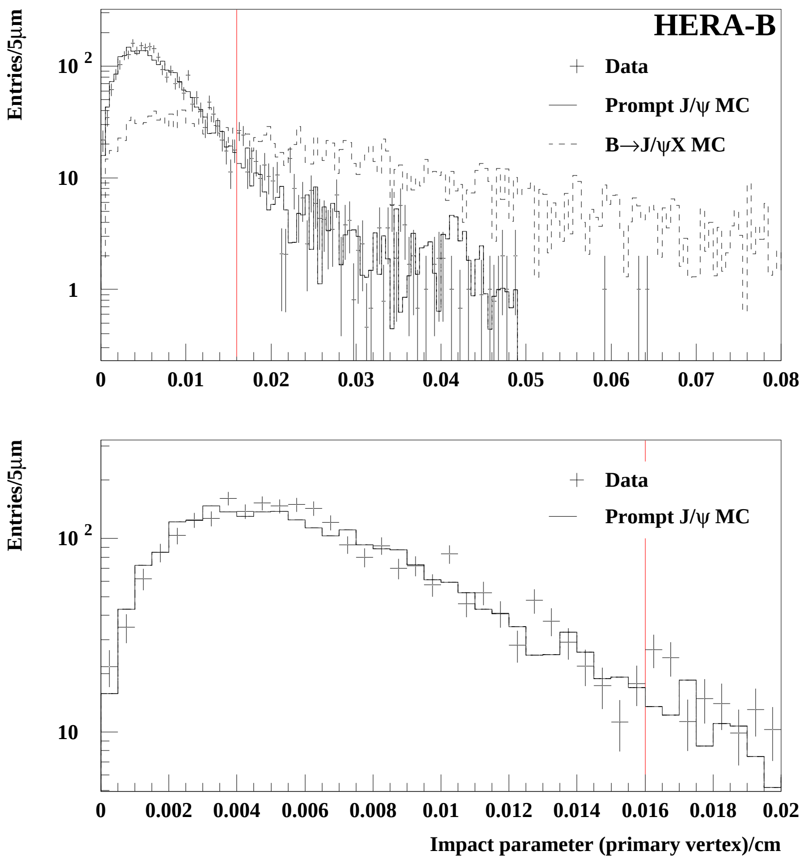

FIGURE 8.9: The minimal impact to the primary vertex of the two muons tracks in data and Monte Carlo with tracks smeared by 20%. The background has been subtracted in the data plots. The expected spectrum for b hadrons is superimposed (dashed line) and the vertical line is the cut applied in the detached event selection. The lower plot shows the details of the left part of the upper plot.

Cut selection procedure

We would like to select cuts, which give either the best precision on the $b\bar{b}$ cross section or the lowest possible upper limit. Because we expect very few signal events[6] and the expected background is less than an event, $S/\sqrt{B}$ is a bad parameter to optimize as it diverges for $B \to 0$. We cannot use $S/\sqrt{S+B}$ neither, since we don't know the size of the signal. Instead we optimize the sensitivity as defined in [88], which for an expected number of background events gives the expected upper limit on the number of $b\bar{b}$ events in the absence of a real signal. This in turn gives us the expected upper limit on $\sigma(b\bar{b})$ after correcting for the selection efficiency and thus the variable we need to optimize. Expected upper limit in this context is to be understood as the average of upper limits in an ensemble of similar experiments with a true $b\bar{b}$ cross section of 0 nb/nucleon.

[6]This can already be seen from Figure 8.7.

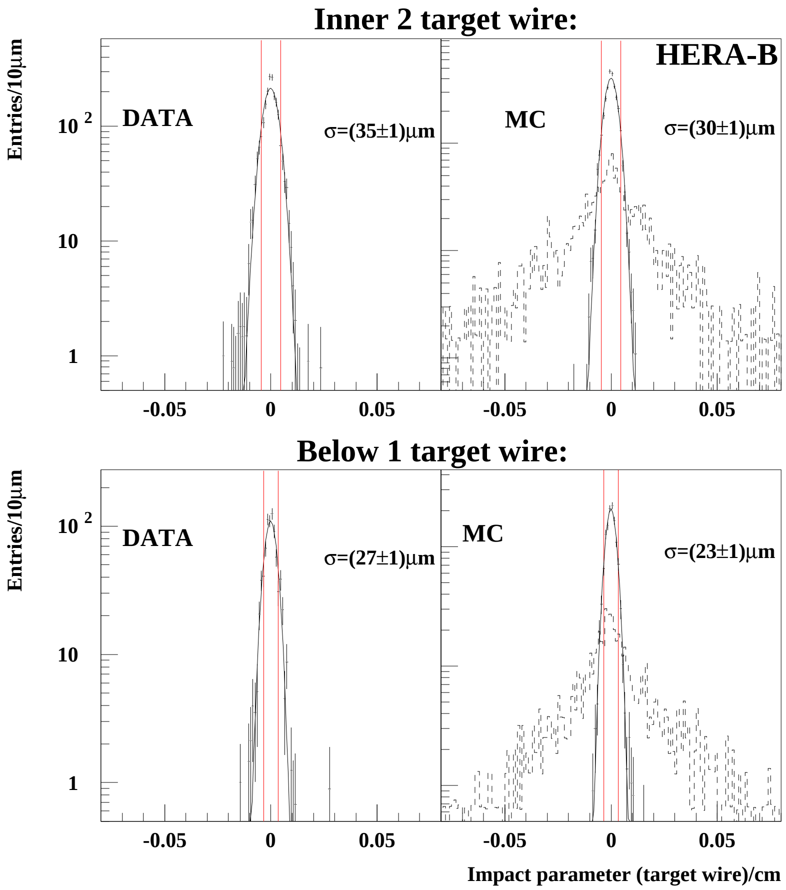

FIGURE 8.10: The minimal impact of the two muon tracks to the target wire, in data and unsmeared Monte Carlo, for each of the two wires. The background has been subtracted in the data plots. The expected spectrum for b hadrons is superimposed (dashed line). The vertical lines mark the cut applied in the detached event selection.

The b selection efficiency and the background of prompt J/ψ are estimated from simulation using the corrections mentioned in the previous section.

The non-J/ψ background contains several components:

1. **Combinatorics:** Pairs of muons from pion and kaon decays. This is the dominant background in the prompt spectrum.

2. **Semileptonic charm decays:** Two charm hadrons decay to muons close enough to each other that the muons are considered to come from a common, detached vertex.

3. **Semileptonic beauty decays:** Two beauty hadrons or cascade charm hadrons decay to muons that are joined into a common, detached vertex.

The first type of background cannot be estimated from simulation, but we assume that the shape of the invariant mass spectrum from combinatorics can be described by the same exponential shape as the background in the prompt J/ψ spectrum. The two other types of background can be predicted from simulation. Both invariant mass spectra can be described by a single exponential, for masses larger than 2.1 GeV/c^2, and both are less steep than in the case of combinatorics. The beauty background has the flattest shape. The amount of background from semileptonic charm and beauty decays is also predicted by simulation, but not very reliably, since the kinematics is quite different compared to J/ψ decays; in particular, due to the fact the muons are not coming from a common vertex. Furthermore, the beauty background scales with $\sigma(b\bar{b})$ [7]. In the selection optimization, we assume the background shape to be given by combinatorics and use the number of downstream events in the mass window $2.1 - 2.8$ GeV/c^2, scaled accordingly to a 3σ mass window around the J/ψ peak, as an estimate of the background count.

The cuts found after optimization procedure are:

1. The decay length larger than 7.5 times σ_z.
2. Impact parameter of muon tracks to the primary vertex larger than 160 μm.
3. Impact parameter of muon tracks to the assigned wire larger than 35 μm for the Below 1 wire and 45 μm for the Inner 2 wire.

These cuts give a b hadron selection efficiency of $\varepsilon_B^{\Delta z} = 45.4\%$, while the sideband estimate of the background is 0.6 events and none of the selected 8000 prompt J/ψ Monte Carlo events survive. The efficiency for each cut on b and prompt Monte Carlo is detailed in Table 8.4. Using Monte Carlo events without masking out dead detector channels, we have 25000 prompt J/ψ in the analysis and none of these survive the detached selection cuts. We will therefore ignore the possibility of having any prompt J/ψ left after our final detached event selection. The sensitivity of these cuts is $\sigma(b\bar{b}) < 16$ nb/nucleon at 68.3% confidence level in the absence of any true b signal. Assuming a $c\bar{c}$ cross section of 40 μb/nucleon [8] [19] and a D semileptonic average decay $\mathrm{Br}(D \to \mu^\pm X) = (11 \pm 2)\%$ [13], the estimated detached charm background with mass larger than 2.1 GeV/c^2 is 2.3 events. This corresponds to 0.2 events expected under the J/ψ peak. From semileptonic b decay (including cascade decays) we expect about three background events with mass larger than 2.1 GeV/c^2 for every $b \to J/\psi X$ event. This corresponds to 0.4 events expected under the J/ψ peak for every $b \to J/\psi X$ event. Both the charm and beauty numbers contain large systematic uncertainties from the simulation, as mentioned above.

8.4.3 $b\bar{b}$ cross section determination in the $\mu^+\mu^-$ channel

The cuts given in the previous section select twelve events with invariant mass above 2.1 GeV/c^2. One event is upstream of the target and lies in the sideband region, the rest of them are downstream with six in the sideband, three in the J/ψ region and the last two just outside the J/ψ region — one on each side. The mass spectrum of the downstream events can be seen Figure 8.11.

To take into account the full information of the invariant mass spectrum, we apply an unbinned likelihood fit of the invariant mass spectrum of the eleven downstream events. The probability distribution for the background, $P_b(m, a)$, is given by an exponential function where

[7] If we trusted the simulation we could also use these events to extract $\sigma(b\bar{b})$.

[8] This value has been rescaled to 920 GeV/c proton momentum from the 38 μb/nucleon in Ref. [19] measured in a 800 GeV/c fixed target experiment.

	Reaction			
	$b \to J/\psi$		Prompt J/ψ	
$\varepsilon_B^{J/\psi}(\varepsilon_P^{J/\psi})$	0.0134		0.0129	
Primary vertex assignment	0.896	0.896	0.820	0.820
$\Delta z/\sigma_z > 7.5\sigma$	0.649	0.559	$8.2 \cdot 10^{-4}$	$6.7 \cdot 10^{-4}$
Impact to primary$>$ 160 μm	0.619	0.499	0.094	$< 1 \cdot 10^{-4}$
Impact to wire$>$(35 or 45) μm	0.736	0.454	0.185	$< 1 \cdot 10^{-4}$
$\varepsilon_B^{\Delta z}$		0.454		$< 1 \cdot 10^{-4}$

TABLE 8.4: Selection efficiency for each of the detached selection cuts applied. The values are given with respect to $\varepsilon_B^{J/\psi}$ (second column) and $\varepsilon_P^{J/\psi}$ (fourth column) for each described cut. Cumulative efficiencies from top to bottom are given in the third and fifth columns. The total b selection efficiencies ($\varepsilon_B^{\Delta z}$) are also reported.

we fit not only the amount of background events, but also the slope, a, since we do not know the exact background composition. For the signal, $P_s(m)$, we use the Gaussian shape fitted to the prompt spectrum. The likelihood function is given by:

$$\mathcal{L}(n_s, n_b, a) = \left(\frac{(n_s + n_b)^N e^{-(n_s+n_b)}}{N!}\right) \cdot \prod_{i=1}^{N} \left(\frac{n_s P_s(m_i) + n_b P_b(m_i, a)}{n_s + n_b}\right), \tag{8.11}$$

where N is the observed number of events, n_s is the fitted number of b events and n_b the fitted numbers of background events. The first part of the likelihood function takes Poisson fluctuations into account.

The result of the likelihood fit is $1.9^{+2.2}_{-1.5}$ b events. For the $\sigma(b\bar{b})$ measurement, the σ_r reference cross section is evaluated from Equation 8.9, using the weighted average of our two target materials, according to the statistical size of the two samples (Table 8.3). We obtain: $\sigma_r = (314 \pm 7_{stat} \pm 31_{sys})$ nb/nucleon. Our chosen x_F range has 72% acceptance for prompt J/ψ, this gives $\Delta\sigma_r = (226 \pm 5_{stat} \pm 2_{sys})$ nb/nucleon. No additional systematic uncertainty is added as the reference cross section was measured inside our x_F acceptance and assuming the same shape of the differential cross section. Based on this reference cross section, the fitted number of b events corresponds to $\sigma(b\bar{b}) = 16^{+18}_{-12}$ nb/nucleon. The result of the fit has been superimposed on Figure 8.11 and the resulting likelihood function projected as function of the $b\bar{b}$ cross section is shown in Figure 8.12. It is clear that the result in itself is not very significant and the result is therefore combined with the cross section measured in the e^+e^- channel.

In the fit we describe the background by a single exponential with a free slope. The background should actually be described by several exponentials from the contributions described in Section 8.4.2, but unfortunately we do not have reliable estimates of the relative amount of each kind of background. It is also not possible to fit the data with several exponential functions with fixed slopes, but free amounts, due to the low statistics. Instead we tried to fit with slopes fixed to combinatorics, charm and beauty background separately. We observe variations in cross section of 20%, which we take as a systematic uncertainty due to the background shape. This uncertainty is most likely an overestimate as all the background shapes lie within the uncertainty of the fitted slope and we are therefore partially counting the statistical uncertainty also as systematical uncertainty.

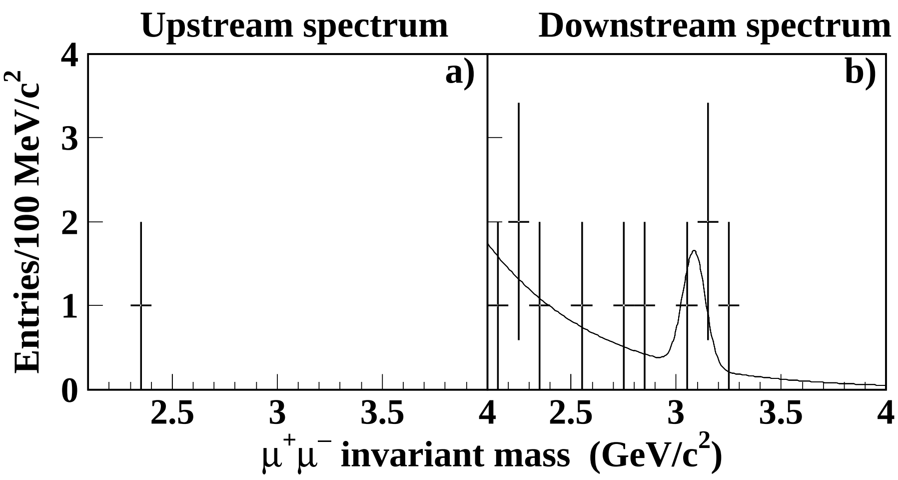

FIGURE 8.11: The invariant di-muon mass spectrum after the detached event selection. The curve shows the result of an unbinned likelihood fit.

If we assume that the amount of background is described correctly by the simulation and the combinatorics is the same both upstream and downstream, we would expect downstream 1 event from combinatorics, 2.3 semileptonic charm events, 5.7 semileptonic beauty events and 1.9 $b \to J/\psi X$ events. In total 10.9 events, fully consistent with the observed 11 events. Given the large uncertainties in the calculation, the good agreement should not be overvalued.

Because we use the same events to select our cuts and to estimate our background, the procedure will have a bias toward underestimating the background. In order to estimate this bias, we slightly soften all cuts, hence introducing slightly more background without changing the b efficiency significantly. We choose the following loosened detached event selection:

1. The decay length larger than 7.0 times σ_z.

2. Impact parameter of muon tracks to the primary vertex larger than 150 μm.

3. Impact parameter of muon tracks to the assigned wire larger than 30 μm for the Below 1 wire and 45 μm for the Inner 2 wire.

Applying these cuts the number of downstream background events increases by four and the upstream background by two events, while giving no new signal events. Due to the observed difference between the two cut selections, we assign a 23% systematic uncertainty contribution from a possible bias in the background estimation.

8.5 Electron analysis

The principle of the $\sigma(b\bar{b})$ measurement in the electron channel is the same as in the muon channel: First the number of prompt J/ψ is measured, followed by a measurement of the $b \to J/\psi X$ events performed by counting J/ψ with detached vertices. The cross section is calculated using the ratio of these two numbers. The analysis in the electron channel is however

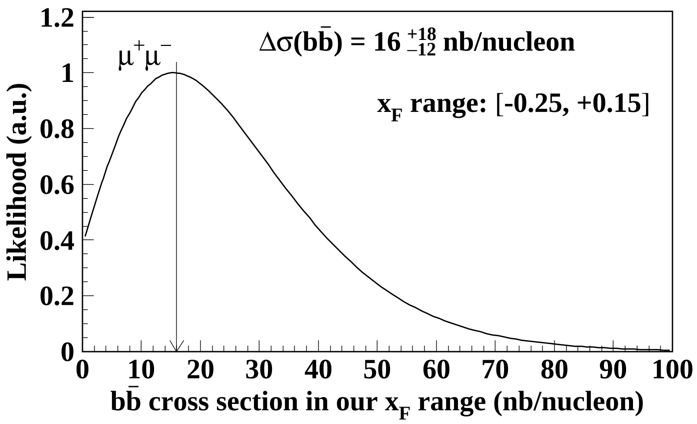

FIGURE 8.12: The maximum of the likelihood for the $J/\psi \to \mu^+\mu^-$ channel as function of $\sigma(b\bar{b})$.

significantly complicated by the much higher background in the prompt selection. In order to measure reliably the number of prompt J/ψ, hard particle identification cuts have to be applied, reducing significantly the statistics available for the detached vertex selection. It turns out that detached vertex selection in itself gives a good suppression of the background in the e^+e^- channel and the hard particle identification cuts for the prompt J/ψ selection are not needed for the b selection. The number of prompt J/ψ events before the particle identification cuts therefore has to be extrapolated from a more "clean" prompt J/ψ sample.

The following sections will describe the main points in the e^+e^- analysis, but the full details will not be presented as the work was not carried out by myself, though of course there was a large degree of cooperation and exchange of ideas between the two analysis groups.

8.5.1 Prompt $J/\psi \to e^+e^-$ selection

The prompt $J/\psi \to e^+e^-$ selection is based on the tracks found by the "ECAL+SLT" trigger and matched to offline reconstructed tracks in both the VDS and OTR. The offline reconstructed tracks are required to be associated to an ECAL cluster with transverse energy, E_T, larger than 1.0 GeV, similar to the trigger requirement. Furthermore the transverse momentum, p_T, of each track is required to be above 0.8 GeV$/c$. The two electron candidates are required to form a common a vertex by requiring the probability of the vertex fit to be higher than 2.5%.

The momentum of electrons has to be corrected for the energy loss due to ionization and bremsstrahlung emission in the detector material. As described in Chapter 4, the track momentum is measured by the bending in the magnet and therefore is giving the momentum of a particle as it passes through the magnet. Bremsstrahlung radiation emitted after the magnet therefore does not influence the momentum measurement. Bremsstrahlung radiation emitted before the magnet, emitted in the direction of the electron, is recovered by extrapolating the VDS track to the front side of the calorimeter. The energy of any nearby cluster is added to the measured momentum. This correction is of course not 100% efficient and the radiation loss inside the magnetic field is not recovered at all. Based on Monte Carlo studies the electron momentum is therefore additionally increased by 3.7% if no bremsstrahlung cluster was found, and by 1.3% if it was found.

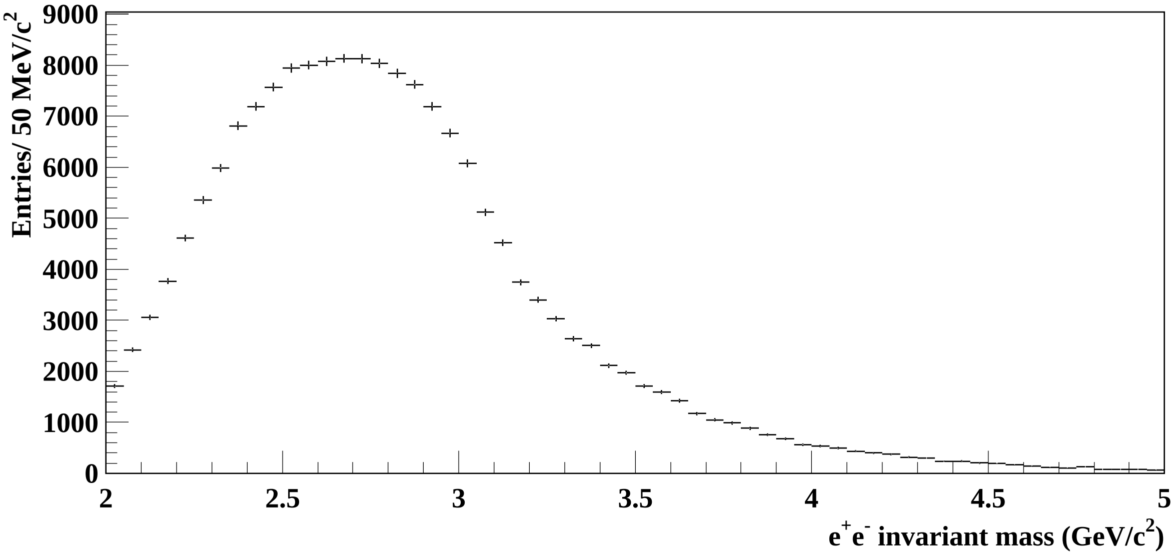

FIGURE 8.13: e^+e^- invariant mass spectrum without BR requirements. No J/ψ signal is seen due to the high level of background.

The invariant mass spectrum of the e^+e^- candidates, corrected by the above procedure, is shown in Figure 8.13. The spectrum is completely dominated by pions interacting in the ECAL and hadrons entering into energetic neutral showers. To see the $J/\psi \rightarrow e^+e^-$ peak one needs to apply particle identification cuts, such as the E/p ratio (energy over momentum), which is equal to one for electrons. The E/p is shown in Figure 8.14 for several classes of tracks. We see that for general trigger tracks it peaks at $E/p \approx 0.9$ and has a width (σ) of 0.16. This is an artifact of the trigger, which assumed $E/p = 1$ when searching for track candidates and the sample is therefore extremely biased. By requiring the tracks to have a bremsstrahlung cluster associated, a clean sample of electrons is found. From this sample we find that the E/p distribution is approximately Gaussian distributed with $E/p \approx 1.0$ and $\sigma = 0.09$. The width is dominated by the energy resolution of the calorimeter, still quite poor for particle identification purposes. By requiring E/p to be within 1σ of 1 for both of the J/ψ electrons, a $J/\psi \rightarrow e^+e^-$ peak can be seen in the invariant mass spectrum shown in Figure 8.15. The invariant mass spectrum is still dominated by the background.

HERA – B has also the RICH detector for particle identification, mainly designed to discriminate between pions and kaons and not electrons and pions. Using the RICH gives better suppression than the E/p requirement, but the electron selection efficiency is not higher than when the bremsstrahlung requirement is used. The RICH has therefore not been used further in the electron channel.

To truly reduce the hadronic background the most effective way is to require that one or both electrons have a bremsstrahlung cluster associated as also shown in Figure 8.15. Requiring a bremsstrahlung cluster for an electron reduces the background by an order of magnitude and the double bremsstrahlung sample is essentially free of background. The J/ψ selection efficiency naturally also suffers quite a lot as majority of the electrons will not emit bremsstrahlung radiation.

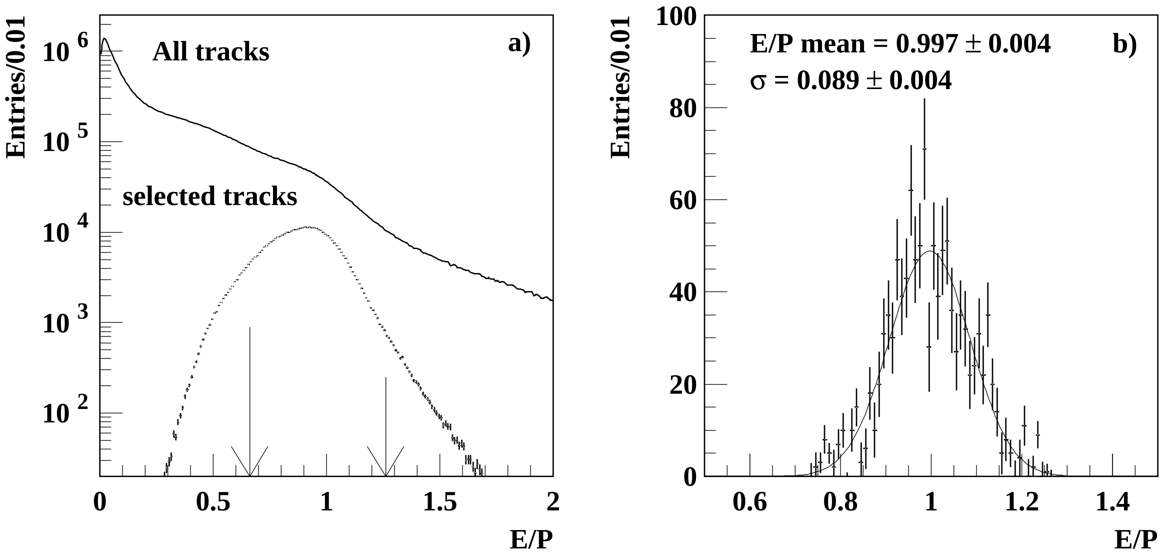

FIGURE 8.14: a) E/p distribution for all tracks in an unbiased sample (solid line) for selected trigger tracks in a trigger sample (data with error bars) The arrows define the maximum region of E/p we use ($\approx 3\sigma$). b) The E/p distribution for electron/positron tracks in a sample of e^+e^-, where both leptons are required to have a bremsstrahlung cluster associated.

Table 8.5 shows the fitted number of J/ψ with different requirements on the number of tracks with bremsstrahlung and spread in E/p. We assume that there is no correlation between the detection of bremsstrahlung emission for the two electrons, which is supported by Monte Carlo studies. The bremsstrahlung detection probability has also been seen to be independent of the charge. From the numbers of J/ψ in Table 8.5 with a minimum of 0, 1 or 2 tracks with bremsstrahlung, we can therefore calculate the probability, ϵ_{BR}, of detecting the bremsstrahlung radiation for an electron. ϵ_{BR} is given in Table 8.5 for the different E/p cuts. These values are correlated as the samples with hard E/p cuts are contained in the ones with softer cuts. The average, taking the correlations into account, is:

$$\epsilon_{BR} = 0.34 \pm 0.02_{stat} \pm 0.02_{sys}. \tag{8.12}$$

In the Monte Carlo the same bremsstrahlung detection probability is found.

We will use for the $b\bar{b}$ analysis a sample with a 3σ E/p cut and no requirement of bremsstrahlung clusters. We can use ϵ_{BR} to extrapolate the number of J/ψ from the sample with two bremsstrahlung clusters or the sample with at least one bremsstrahlung cluster to the sample without any bremsstrahlung requirement. Similarly we can use that E/p is Gaussian distributed and extrapolate for instance from the sample with a 1σ E/p cut and no bremsstrahlung requirement to the sample with a 3σ cut. All of the extrapolations are consistent within errors and we estimate the number of J/ψ to be:

$$n = 5710 \pm 380_{stat} \pm 280_{sys}. \tag{8.13}$$

A complication to this method is that fits used to extract the number of J/ψ in the different samples assumed a Gaussian shape for the J/ψ and various polynomials for the background.

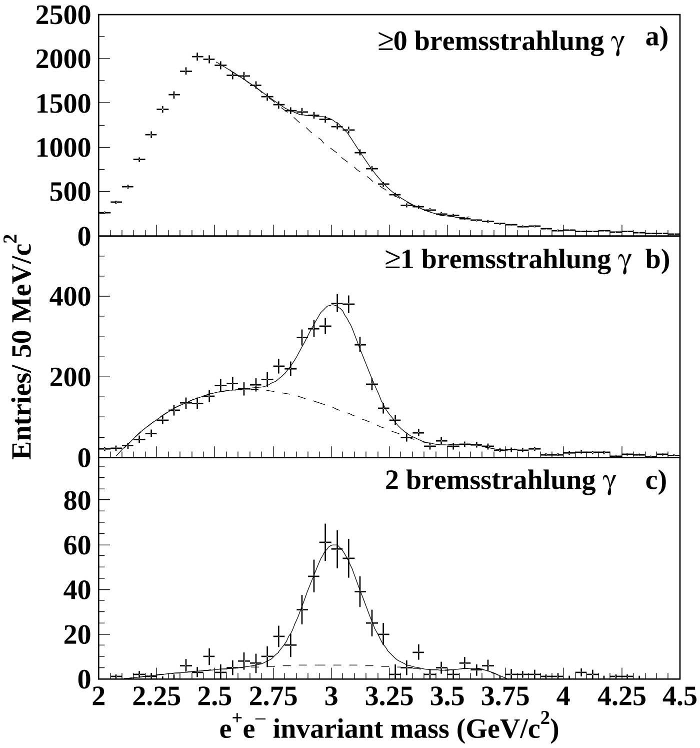

FIGURE 8.15: e^+e^- invariant mass spectra, with an E/p cut at 1σ on both e^+ and e^- tracks. a) no bremsstrahlung requirement. b) at least one associated bremsstrahlung cluster. c) two associated bremsstrahlung clusters. The fits (solid lines) assume a Gaussian signal and a polynomial background. The background shapes (dashed lines) differ from the di-muon case (Fig. 8.3), due to E_T invariant mass cuts in the trigger.

However from the simulation we know that $J/\psi \to e^+e^-$ events have a significant tail toward lower mass values due to unrecovered bremsstrahlung radiation, which in the fits will be counted as background. Because the background shape is not known, it is not possible to fit reliably the number of J/ψ events using the full $J/\psi \to e^+e^-$ spectrum. Instead we estimate the loss by adding "background events" distributed similar to the ones fitted in data to the invariant mass spectrum of simulated $J/\psi \to e^+e^-$ events. Using this method we estimate that the real number of $J/\psi \to e^+e^-$ events is 26% higher than the value calculated above[9]:

$$n_P = 7750 \pm 520_{stat} \pm 380_{sys}. \tag{8.14}$$

From simulation, the trigger and reconstruction efficiency for prompt $J/\psi \to e^+e^-$ events in our x_F acceptance has been estimated to be 2.53%, while it is 2.98% for J/ψ decays from b-hadron decays. The J/ψ mesons from b-hadron decays have higher reconstruction and trigger efficiency mainly due to the high E_T cut in the trigger system. We therefore have the following ratio of selection efficiencies:

$$\varepsilon_R = 1.18 \pm 0.05, \tag{8.15}$$

[9]In the Ref. [68, 69] the correction is instead done to ε_R, but in my opinion it is better to correct the J/ψ estimate. No systematic uncertainty for this correction has been assigned.

BR	E/P cut in units of σ				
req.	0.5σ	1σ	2σ	2.5σ	3σ
≥ 0	895 ± 75	2553 ± 184	5362 ± 292	-	-
≥ 1	$519 \pm 45 \pm 11$	$1420 \pm 70 \pm 48$	$2851 \pm 108 \pm 48$	$3100 \pm 126 \pm 170$	$3304 \pm 148 \pm 188$
≥ 2	$106 \pm 13 \pm 2$	$308 \pm 24 \pm 6$	$587 \pm 38 \pm 14$	$622 \pm 38 \pm 18$	$661 \pm 40 \pm 29$
ϵ_{BR}	$0.34 \pm 0.05 \pm 0.01$	$0.35 \pm 0.03 \pm 0.01$	$0.36 \pm 0.03 \pm 0.01$	$0.33 \pm 0.02 \pm 0.02$	$0.33 \pm 0.02 \pm 0.02$

TABLE 8.5: Number of prompt J/ψ with different bremsstrahlung requirements and different E/p cuts on both tracks. Above the 2σ cut on E/p, no signal is visible anymore in the spectrum without bremsstrahlung requirement due to the high background. The last row contains the calculated bremsstrahlung reconstruction efficiencies. In all numbers, the systematic uncertainty (second quoted error) contains only the effect of varying the fitted mass region.

where the uncertainty is due to acceptance uncertainties estimated by using ideal detector performance in the simulation.

8.5.2 Detached $J/\psi \to e^+e^-$ selection

With the much worse S/B ratio in the initial spectrum than in the $\mu^+\mu^-$ case it is clear that harder cuts need to be applied in the e^+e^- channel to suppress the background and to allow us to determine the amount of $J/\psi \to e^+e^-$ events from b-hadron decays. It turns out to be more efficient in the electron channel to rely almost exclusively on the wire position instead of the reconstructed primary vertex. A primary vertex is still required to be present to ensure that the correct wire has been assigned but no further use is made of the primary vertex position.

The detached $J/\psi \to e^+e^-$ selection cuts used are:

1. The impact parameters of both tracks to the wire should be larger than 200 μm. Failing that, a track can still be accepted if it is isolated. Here a track is considered to be isolated, if, at the z of the wire, no other track passes within 250 μm in the $X - Y$ plane.

2. Distance between the J/ψ flight direction and the wire position should be less than 300 μm.

3. The decay length, Δz, should be larger than 0.5 cm.

The second cut uses the fact that a J/ψ from a b-hadron decay is moving roughly in the same direction as the b-hadron and therefore points back to the interaction point, unlike tracks from fake vertices, which can generally point anywhere. The cut does not remove vertices produced in the primary interaction, that is the purpose of cuts 1 and 3.

The exact values of the cuts are chosen by "optimizing" the sensitivity, somewhat similarly to the procedure used for the muon channel. The signal is estimated from the $b \to J/\psi \to e^+e^-$ Monte Carlo and the background using the events outside the $J/\psi \to e^+e^-$ mass window between 2.8 and 3.3 GeV/c^2. Instead of optimizing all cuts at the same time, each cut is tuned individually. This means that correlations between cuts are not taken into account. As the background estimates are not taken from the final spectrum, any bias from the cut selection is smaller. Since during the tuning of a cut the other cuts are not applied, there are plenty of background events and the sensitivity is estimated simply by $S/\sqrt{B}$. The sensitivity and efficiency as a function of cuts are shown in Figure 8.16.

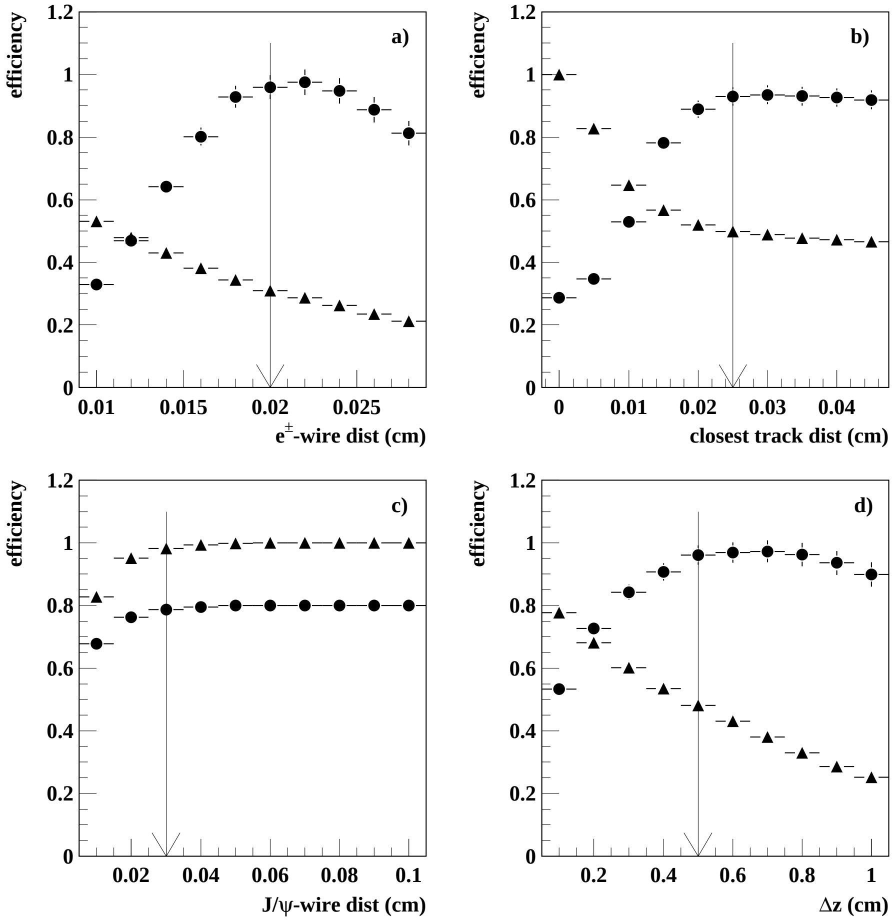

FIGURE 8.16: Sensitivity to the signal (filled circles, arbitrary units) and $b \to J/\psi \to e^+e^-$ efficiencies (triangles) after J/ψ reconstruction are shown for different cuts: a) lepton impact parameter to the wire, b) isolation cut, c) J/ψ impact parameter to the wire and d) Δz. The arrows show the cut positions used.

The efficiencies of the individual cuts on $b \to J/\psi \to e^+e^-$ and prompt $J/\psi \to e^+e^-$ events are listed in Table 8.6. No prompt $J/\psi \to e^+e^-$ events in the Monte Carlo survive the detached event selection. From simulation we also obtain an upper limit of 0.4 (90% C.L.) semileptonic charm decays in our final sample. These are much more suppressed in the electron channel compared to the muon channel because of the high E_T requirement in the trigger and the harder detached vertex requirements. Semileptonic beauty decays are also more suppressed than in the muon channel. We estimate to have 0.8 semileptonic beauty decays in our final sample for every $b \to J/\psi \to e^+e^-$ decay.

The cuts applied in the electron channel are quite different from the ones for the muon channel. Using the cuts from the electron channel in the muon channel, the background is more suppressed, but the efficiency drops a lot and the calculated sensitivity is worse than with the cuts used in the muon analysis. Almost all of the events in the $J/\psi \to \mu^+\mu^-$ mass region survive (no new ones appear) and, because of the lower efficiency estimate, the $\sigma(b\bar{b})$ value comes out

	Reaction			
	$b \to J/\psi$		Prompt J/ψ	
$\varepsilon_B^{J/\psi}(\varepsilon_P^{J/\psi})$	2.53		2.98	
Vertex assignment	0.916	0.916	0.862	0.862
Impact parameter/isolation cut	0.330	0.305	$8.0 \cdot 10^{-3}$	$7.8 \cdot 10^{-3}$
J/ψ direction	0.950	0.291	0.866	$7.8 \cdot 10^{-3}$
Decay length cut	0.523	0.280	$3.6 \cdot 10^{-4}$	$< 4 \cdot 10^{-5}$
$\varepsilon_B^{\Delta z}$		0.280 ± 0.010		$< 4 \cdot 10^{-5}$

TABLE 8.6: Trigger and reconstruction efficiencies for $b \to J/\psi$ and prompt J/ψ at different stages of the analysis chain. It shows the efficiencies for b selection given with respect to $\varepsilon_B^{J/\psi}$ (second) and $\varepsilon_P^{J/\psi}$ (fourth column) for each described cut. Cumulative efficiencies from top to bottom are given in the third and fifth columns. The b selection efficiencies ($\varepsilon_B^{\Delta z}$) are also reported. All numbers are evaluated as a weighted average on C and Ti Monte Carlo, with the weights chosen according to the J/ψ statistics in Table 8.3. All quoted errors include the Monte Carlo statistical contribution only.

higher. The relative uncertainty is however larger and therefore electron channel cuts are not used in the muon channel.

8.5.3 $b\bar{b}$ cross section determination in the e^+e^- channel

With the detached event selection presented in the previous section, we select 19 downstream events and 8 upstream events as shown in Figure 8.18. 10 of the events are in the J/ψ mass region and to fit the number of $b \to J/\psi \to e^+e^-$ events, an unbinned likelihood is used in the same way as in the muon channel. The background is expected to be given by both combinatorics, as witnessed by the upstream events, and by semileptonic beauty decays, as seen in the simulation. The shape of the semileptonic beauty decays has been extracted from simulation, while for the combinatorics a fit has been done to the invariant mass of the upstream events (with looser cuts in order to have more statistics). Both shapes are shown in Figure 8.17 and are quite similar since the drop-off toward lower mass is given by the trigger acceptance. For the likelihood fit, an average of the two shapes has been used and only the number of detached J/ψ events and the number of background events are fitted. By varying the shape used in the fit between the two background shapes a systematic uncertainty of 7% has been estimated.

From the likelihood fit the amount of $b \to J/\psi \to e^+e^-$ events after the detached event selection is estimated to be $8.6^{+3.9}_{-3.2}$ events. To estimate any systematics from the cut selection, the cuts have been varied and re-optimized using different selection procedures. The results obtained using two of the selection procedures are shown in Figure 8.19. The amount of detached J/ψ events vary naturally due to the low statistics, but all the selections result in an excess around the J/ψ mass in amounts consistent with our result. From the cut variation a systematic uncertainty of 11% has been assigned to the number of $b \to J/\psi \to e^+e^-$ decays.

Using $\varepsilon_R = 1.18$ and $\varepsilon_B^{\Delta z} = 0.28$ from the simulation, we convert the number of detached $b \to J/\psi \to e^+e^-$ events into a measurement of $\Delta\sigma(b\bar{b})$. Figure 8.20 shows the likelihood of the fit as a function of $\Delta\sigma(b\bar{b})$. We obtain:

$$\Delta\sigma(b\bar{b}) = 38^{+18}_{-15} \text{ nb/nucleon.} \tag{8.16}$$

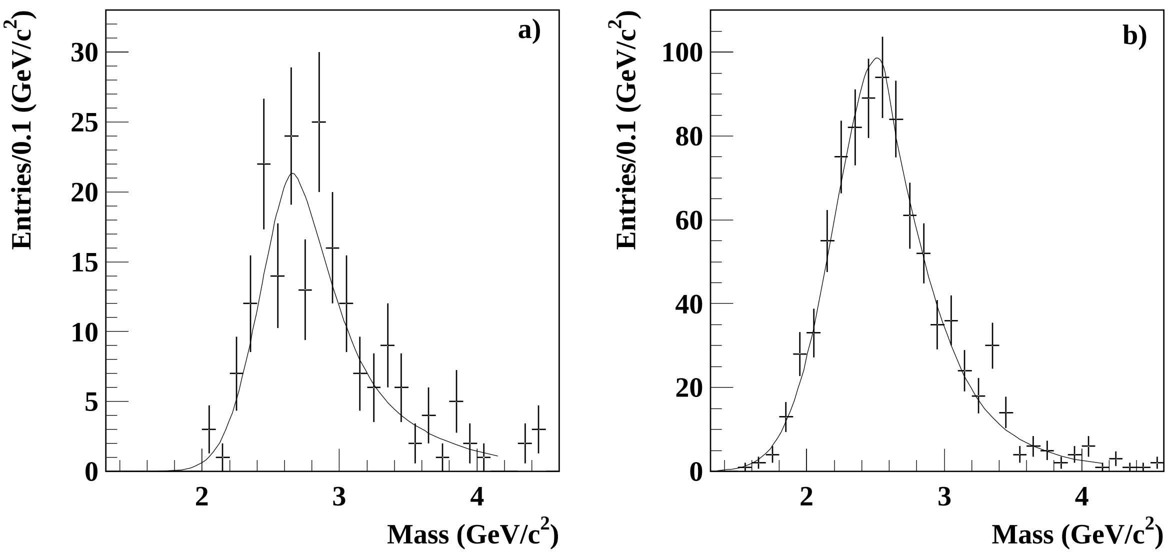

FIGURE 8.17: Background shapes used in the maximum likelihood fits (solid lines). On the left, a fit of $b\bar{b} \rightarrow e^+e^-$ Monte Carlo events was done, after removing the $b \rightarrow J/\psi \rightarrow e^+e^-$ contribution, while on the right the procedure was applied to real data upstream events with relaxed cuts.

The result is higher but consistent with the result obtained for the muon channel.

8.6 Results

The two measurements, $\sigma(b\bar{b}) = 16^{+18}_{-12}$ nb/nucleon and $\sigma(b\bar{b}) = 38^{+18}_{-15}$ nb/nucleon, obtained in the muon and electron channel respectively, are statically independent and compatible within the statistical errors.

In order to have a more precise measurement, we combine the $\mu^+\mu^-$ and e^+e^- likelihoods in a four parameter likelihood maximization — $\sigma(b\bar{b})$, $n_b(\mu^+\mu^-)$, $\mu^+\mu^-$ background slope and $n_b(e^+e^-)$ — of the detached candidates. The resulting likelihood is shown in Figure 8.21 and the fit provides our final result of the $b\bar{b}$ production cross section on the HERA – B 2000 data:

$$\Delta\sigma(b\bar{b}) \ = 30^{+13}_{-11\,stat} \text{ nb/nucleon}, \tag{8.17}$$

where the quoted uncertainty has been estimated directly by the fit.

8.6.1 Systematic uncertainties

The systematics can be divided into two groups: uncertainties coming from the normalization and uncertainties from the J/ψ counting. The former is mainly given by external parameters and difficult for HERA – B to improve upon, while the latter is mainly an effect of the low statistics available in the 2000 data.

Apart for small variations due the slightly different acceptance, the normalization systematic uncertainties are almost the same in the $\mu^+\mu^-$ and e^+e^- channels. We use therefore for each

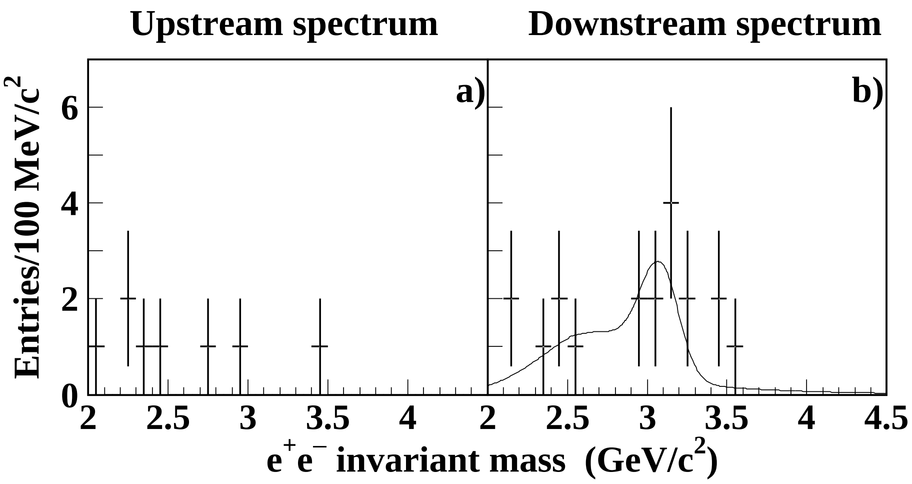

FIGURE 8.18: The upstream (left) and downstream (right) e^+e^- invariant mass spectra after detached event selection. The curve on the right shows the result of the unbinned likelihood fit.

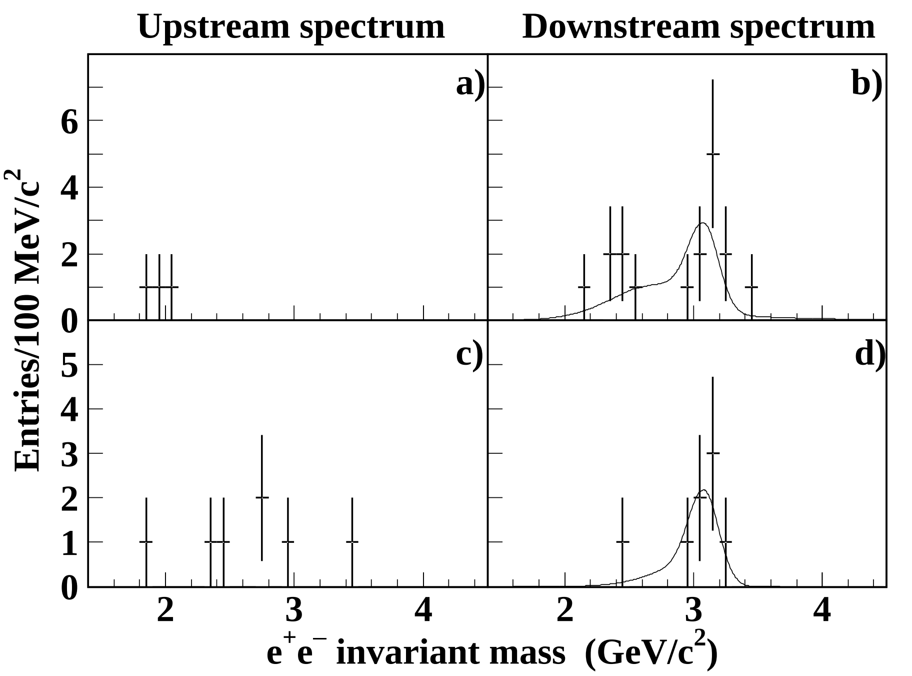

FIGURE 8.19: The upstream (left) and downstream (right) e^+e^- invariant mass spectra for different optimization requirements. a,b) simultaneous optimization of all cuts on upstream background and downstream Monte Carlo $b \to J/\psi \to e^+e^-$ events. c,d) simultaneous optimization of all cuts on downstream background regions, (2.0 GeV/c^2 $< m_{e^+e^-} <$ 2.8 GeV/c^2) or (3.3 GeV/c^2 $< m_{e^+e^-} <$ 6.0 GeV/c^2), and downstream Monte Carlo $b \to J/\psi \to e^+e^-$ events.

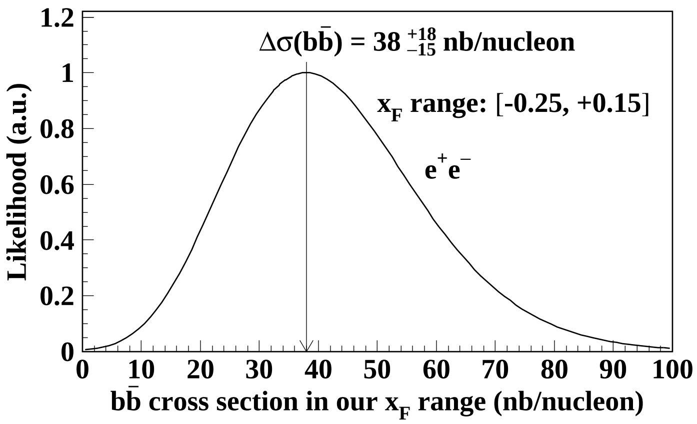

FIGURE 8.20: The maximum of the likelihood for the $J/\psi \rightarrow e^+e^-$ channel as function of $\Delta\sigma(b\bar{b})$. It has been maximized over the amount of background.

systematic contribution a common value for the two channels. We have considered uncertainties coming from the following sources:

- **Prompt J/ψ reference cross section:** The cross section calculated in Section 8.1.2 has a measurement uncertainty of 7.5%. Furthermore we have a 3.5% uncertainty from the A dependence on the cross section.

- **Br($b\bar{b} \rightarrow J/\psi X$):** The LEP measurement has a 9% uncertainty.

- **Prompt J/ψ model:** By varying the prompt J/ψ x_F and p_T spectrum within the uncertainty of its measurement [71] and by using the COM model instead, we see a 2.5% variation in the prompt J/ψ selection efficiency, and thus in ε_R.

- **$b\bar{b}$ production model:** The parton distribution, b lifetime and other parameters of the $b\bar{b}$ production model have been varied as described in Section 8.3.3. In total a 5% systematic uncertainty has been assigned.

- **Detector effects:** Several effects due to the detector have been investigated. While the assumed detector performance has a large influence on the prompt J/ψ and $b \rightarrow J/\psi X$ selection efficiencies, most of them cancel in the ratio ε_R as they apply almost identically on the two channels. The detector performance influence has been estimated to less than 5% by running the simulation with ideal detector performance instead of the performance measured for run 16665. We also estimate a 1.7% uncertainty due to the difference in the acceptance and multiplicity in events from the carbon and titanium wires. Finally a 1% uncertainty from the smearing of the tracking resolution is seen. The smearing has little influence on our selection as our tracking and vertex precision is much higher than the cuts used for the detached event selection.

In total, the systematic uncertainty due to normalization is estimated to be 14%. With more statistics the model dependence could be reduced by measuring the differential cross section

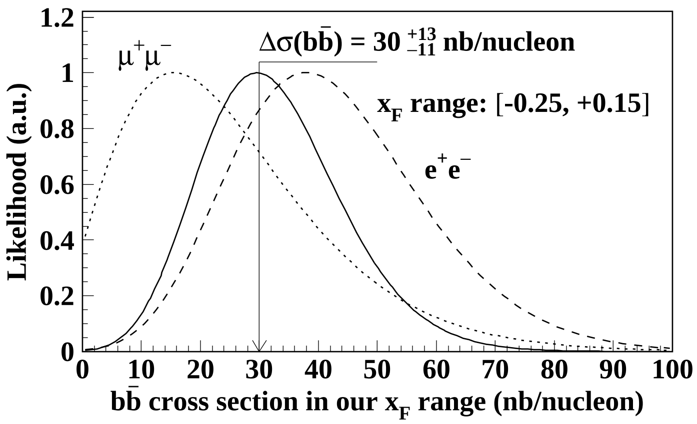

FIGURE 8.21: The combined likelihood function maximized with respect to the $\mu^+\mu^-$ and e^+e^- background yields and to the $\mu^+\mu^-$ background slope is shown as a function of the $b\bar{b}$ production cross section.

instead of the cross section in a single x_F interval. The uncertainty due to the detector acceptance should also be reducible with a better understanding of the detector performance. The improvements will however be marginal.

The systematic uncertainties from the J/ψ counting are given by the following contributions:

- **Prompt J/ψ counting:** In the $J/\psi \to e^+e^-$ channel we had to extrapolate the number of prompt J/ψ events from samples with harder particle identification cuts. A 5% systematic uncertainty has been assigned to this extrapolation. In the muon channel we have a clear Gaussian shaped J/ψ peak in the invariant mass spectrum and we furthermore use the same shape in the fit of the detached events. We therefore expect the systematic uncertainty to be negligible.

- **Background shape under $b \to J/\psi X$ peak:** In the muon channel an exponential function was used describe the background under the $b \to J/\psi X$ peak. The resulting slope was consistent with the shapes found from simulation for semileptonic beauty and charm decays and with the shape observed for the background in prompt decays. The true background shape is thought to be a combination of these three types and by varying the shape between these three variants, a 20% variation in the number of $b \to J/\psi$ candidates from the fitted result was seen. This is most likely a conservative estimate, as this uncertainty is partially contained in the statistical error. In the electron channel, the background is varied between the shape seen for simulated semileptonic beauty decays and the one for the upstream background events. A 7% variation is seen.

- **Cut selection procedure:** When choosing the cuts for the detached event selection the background was estimated from the sideband events in the invariant mass spectrum. Some of the same events are used also when the final mass spectrum is fitted to extract the number of signal events. This will have an tendency to underestimate the amount of background and thus to overestimate the amount of signal events. The effect has been estimated by varying the cuts and the selection procedure to be 11% in the electron channel and 23% in the muon channel.

The uncertainty from the J/ψ counting in muon channel adds up to 30%, while in the electron channel, with more statistics, it is only 14%. Weighted according to the weight of the two channels in the combined result, the systematic uncertainty from the J/ψ counting is estimated to be 20%. With more statistics this uncertainty could be drastically reduced. The uncertainty from the selection procedure could be removed by selecting the cuts on a small subsample of the statistics and measuring the cross section on the remaining sample. With enough statistics, a clear J/ψ in the detached sample should also be available and the amount of detached J/ψ events can be fitted as simple as was done for the prompt $J/\psi \to \mu^+\mu^-$. Some uncertainty in the electron channel will probably remain, due to bremsstrahlung tails in the spectrum.

With this, the final value for the $b\bar{b}$ cross section in our x_F acceptance including statistical and systematic uncertainty is:

$$\Delta\sigma(b\bar{b}) = 30^{+13}_{-11\,stat} \pm 7_{sys} \text{ nb/nucleon.} \tag{8.18}$$

Using our $b\bar{b}$ production model to extrapolate the result to the full x_F range we obtain the total $b\bar{b}$ production cross section:

$$\sigma(b\bar{b}) = 32^{+14}_{-12\,stat} \pm 8_{sys} \text{ nb/nucleon.} \tag{8.19}$$

We have not assigned any further systematics for the extrapolation.

8.7 Comparison with other results

Two other measurements of $\sigma(b\bar{b})$ in proton–nucleus collisions have been performed at energies close to the HERA – B energy. The measurements at E789 [89] and E771 [90] both used a 800 GeV proton beam at Fermilab and measured:

$$\sigma(b\bar{b}) = 5.7 \pm 1.5 \pm 1.3 \text{ nb/nucleon} \quad \text{(E789)}, \tag{8.20}$$

$$\sigma(b\bar{b}) = 43^{+27}_{-17} \pm 7 \text{ nb/nucleon} \quad \text{(E771)}. \tag{8.21}$$

E771 measured the cross section using double semileptonic beauty decays into muons, while the E789 measurement is quite similar to method used in the analysis presented here. The HERA – B and E789 measurements have overlapping x_F regions, the beam momentum differs only by 13% and the chosen b decay channel is $b \to J/\psi X$ for both measurements. In Table 8.7, we list all the parameters and values characteristic of the HERA – B and the E789 analysis. For a more precise comparison, the E789 result has been corrected in the third column using the updated $\mathrm{Br}(b\bar{b} \to J/\psi X)$ value (+12%) and for the predicted $\sqrt{s}$ dependence[21] of $\sigma(b\bar{b})$ (+35%).

The latest QCD calculations [22, 21] predict cross section values close to the one measured by HERA – B: $\sigma(b\bar{b}) = 25^{+20}_{-13}$ nb/nucleon[10] and $30 \pm 8 \pm 10$ nb/nucleon respectively. Figure 8.22 shows the two theoretical predictions together with the three measurements. Both theory and measurements have still large uncertainties and any definite conclusions cannot be made. Our measurement has somewhat higher precision than the E771 measurement, while the most precise measurement (E789) seems to be on low side compared to the theoretical prediction. The experimental situation will be significantly improved by the HERA – B measurement in 2002. With an increase in statistics of 10^3, the statistical uncertainty will be just a few percent, while based on systematic uncertainties estimated in the 2000 analysis presented here, the systematic uncertainties for the 2002 $\sigma(b\bar{b})$ measurement should be of the order 10–20%.

[10] The result of [22] has been updated to the latest MRST parton distribution functions [81].

	E789		HERA – B	
	Published ($\mu^+\mu^-$) [89, 71]	Extrapolated to HERA – B	$\mu^+\mu^-$ channel	e^+e^- channel
Target	Gold (A=197)	Gold (A=197)	77% C(A=12) +23% Ti(A=48)	
Interaction rate (MHz)	50	-	5	
p beam (GeV/c)	800	920	920	
$\sqrt{s}$ (GeV)	38.7	41.6	41.6	
α	0.90 ± 0.02	0.955 ± 0.005	0.955 ± 0.005	
$\sigma(J/\psi)$ (nb/nucleon)	442 ± 88	355 ± 28	357 ± 28	
Prompt J/ψ	71206	-	2880	7750
Detached J/ψ	19 ± 5	-	$1.9^{+2.2}_{-1.5}$	$8.6^{+3.9}_{-3.2}$
$\varepsilon_R \cdot \varepsilon_B^{\Delta z}$	0.47	0.47	0.41	0.32
x_F range	$0 < x_F < 0.1$	$0 < x_F < 0.1$	$-0.25 < x_F < 0.15$	
$BR(b\bar{b} \to J/\psi X)$ (%)	2.60 ± 0.34	2.32 ± 0.20	2.32 ± 0.20	
$\sigma(b\bar{b})$ (nb/nucleon)	$5.7 \pm 1.5 \pm 1.3$	$8.6 \pm 2.3 \pm 2.0$	$32^{+14}_{-12} \pm 8$	

TABLE 8.7: Comparison between the $\sigma(b\bar{b})$ measurement of E789 and of HERA – B. The E789 result has been corrected for the updated $BR(b\bar{b} \to J/\psi X)$ value (+12 %) and for the predicted dependence[21] of $\sigma(b\bar{b})$ on $\sqrt{s}$ (+35 %) in the third column.

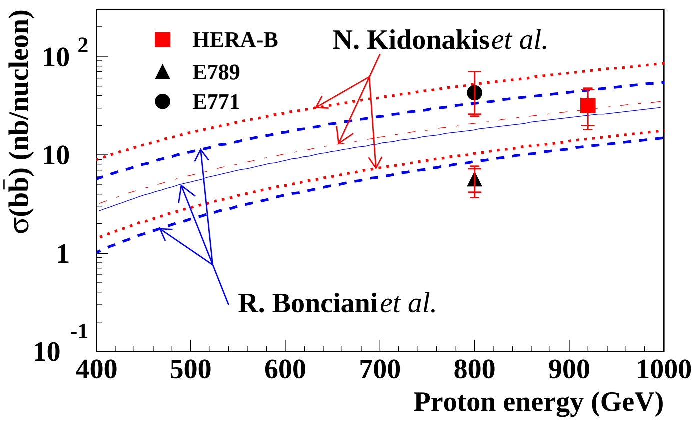

FIGURE 8.22: Comparison of the HERA – B (2000) $\sigma(b\bar{b})$ value with other experiments and with the theoretical predictions of R. Bonciani *et al.* [22] updated for the MRST parton distribution function in [81] (solid line: central value, dashed lines: upper and lower bounds) and N. Kidonakis *et al.* [21] (dot-dashed line: central value, dotted lines: upper and lower bounds).

Chapter 9

Summary

HERA – B has the difficult task of studying beauty and J/ψ mesons in an environment dominated by a large background of hadronic interactions. In order to do this, the HERA – B spectrometer has been equipped with a complex multi-level trigger system for selecting charged lepton pairs online. During the 2000 run many parts of the spectrometer were still being commissioned and most of the online selection was done by the Second Level Trigger system (SLT). The SLT is a software trigger running on a farm of 240 Pentium PCs. It does full track reconstruction inside limited regions of the detector.

In this thesis the SLT configuration as used during the 2000 run is presented and a detailed evaluation of the performance is given. Trigger efficiencies of $O(1\%)$ are found. The trigger inefficiencies is mainly related to the performance of the subdetectors used by the SLT and not to the trigger algorithm itself. However, some weak points of the 2000 trigger configuration are found. These have all been corrected for the upcoming 2002 run. The trigger performance with the improved trigger algorithm and improved detectors is evaluated. The study shows a much improved efficiency with sufficient background rejection to run with an interaction rate of 10 MHz. For the 2002 run $\sim 10^7$ $J/\psi \to \mu^+\mu^-$ and $\sim 10^6$ $J/\psi \to e^+e^-$ events are expected to be recorded.

One of the main physics topics for the 2002 run is a precise measurement of the beauty production cross section ($\sigma(b\bar{b})$). In this thesis a measurement of $\sigma(b\bar{b})$ has been presented based on a sample of $\sim 10^4$ J/ψ events recorded during the 2000 run. While the result, $\sigma(b\bar{b}) = 32^{+14}_{-12}{}_{stat} \pm 8_{sys}$ nb/nucleon, does not significantly improve the precision over previous measurements, it demonstrates that HERA – B is able to trigger and reconstruct beauty decays. With the huge improvement in statistics expected for the 2002 run, HERA – B will be able to make a precise measurement of $\sigma(b\bar{b})$.

Appendix A

Track reconstruction with Kalman filters

A Kalman filter is one of many methods, which can be used for track reconstruction. In the SLT algorithm this method is used for tracking in both the main tracker and the VDS. This appendix will present the basic equations of a Kalman filter for track reconstruction. It is based on the excellent introduction in Ref. [91] from which the notation, equations and figures are taken. For the derivation of the equations and proof that the method gives the optimal track parameters, see Ref. [92].

The track reconstruction in general consists of two tasks: track finding and track fitting. The track finding is the task of determining which hits should be assigned to a track. This is an exercise in pattern recognition and is in general not part of a Kalman filter, but, due to the iterative nature of a Kalman filter, it can be easily added to the Kalman filter algorithm. The track fitting is the task of calculating the track parameters from the hits assigned to the track. The fit should weight the contribution of each individual hit based on the hits resolution and the multiple scattering effects on the track. It can be shown that under certain assumptions[1] the track parameters found are the best which can be possible achieved.

A Kalman filter actually consists of two parts: a filtering step and a smoothing process. The filtering process is visualized in Figure A.1. It starts at one end of the detector with an initial estimate of the track parameters and uncertainty and proceeds sequentially through the detector updating the estimates by adding the information of the hits from each layer in the detector. The hits are added at each layer by first predicting the track parameters and uncertainties at the next layer, taking the multiple scattering in the current layer into account in the uncertainty. A new hit is added by weighting the predicted and measured position according to their uncertainties. The uncertainties are characterized with the help of covariance matrices.

After the filtering process, the track parameters are known at the last layer with the full information of all upstream layers included. In the intermediate layers, the track parameters does not include the information provided by the downstream layers and thus have not got the optimal precision. The track parameters in the intermediate layers is updated by the smoothing process, which is illustrated in Figure A.2. It starts at the last layer used in the filtering process and proceeds back to the first layer, refining the track parameters and uncertainties to include the downstream information. After the smoothing process the track parameters are optimally

[1] The hit uncertainty and the multiple scattering should be described by Gaussian distributions.

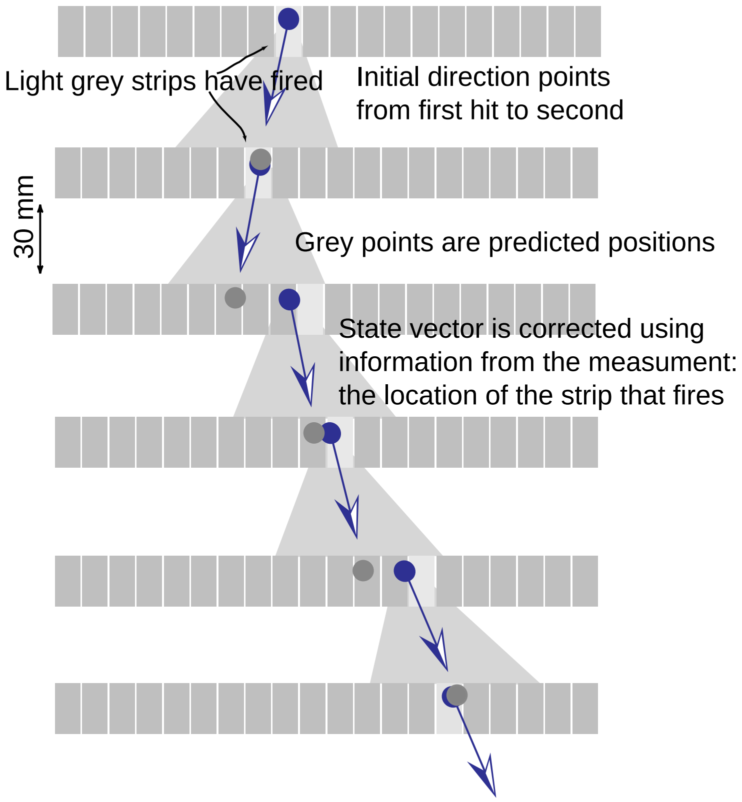

FIGURE A.1: The Kalman filtering process

determined everywhere.

A.1 The filtering process

In the Kalman filter, the track parameters are given at each plane, k, by a state vector, $\vec{x}_k$, which contains all relevant track information, such as position and slopes. The state vectors in each layer are related to the ones in the previous layer through the equation:

$$\vec{x}_k = \mathbf{F}_{k-1}\vec{x}_{k-1} + \vec{w}_{k-1}, \tag{A.1}$$

where $\mathbf{F}_{k-1}$ is the propagator of the state vector from the $k-1$ layer to the next one and $\vec{w}_{k-1}$ is the change in track parameters due to multiple scattering in layer $k-1$. The state vector in a layer is related to a hit measurement through the equation:

$$\vec{m}_k = \mathbf{H}_k\vec{x}_k + \vec{\epsilon}_k, \tag{A.2}$$

where $\vec{m}_k$ is the hit measurement in the layer k, $\mathbf{H}_k$ is the measurement matrix, which relates the state vector coordinates to the hit measurement and $\vec{\epsilon}_k$ is the measurement error. A hit measurement will typically be one-dimensional, for instance the x position of a vertical wire, and $\mathbf{H}_k$ will therefore often not be a square matrix.

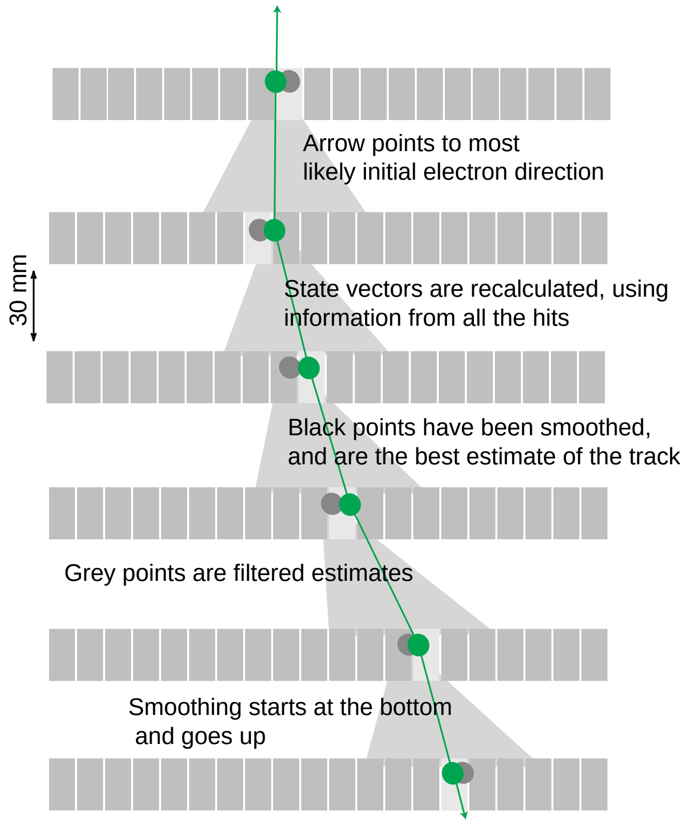

FIGURE A.2: The Kalman smoothing process

The multiple scattering and the measurement errors are given by random variables. We will assume that the state vectors are defined such that the expectation values of $\vec{w}_k$ and $\vec{\epsilon}_k$ are zero and we therefore only need to consider their covariance matrices:

$$\mathbf{Q}_k = \mathrm{cov}\{\vec{w}_k\}, \tag{A.3}$$

$$\mathbf{V}_k = \mathrm{cov}\{\vec{\epsilon}_k\}. \tag{A.4}$$

We similarly define the covariance matrix of the measured state vector, $\vec{x}_k$, to be:

$$\mathbf{C}_k = \mathrm{cov}\{\vec{x}_k - \vec{x}_{k,true}\}. \tag{A.5}$$

With these definitions the filtering process can be expressed with just a few equations. The state vector in layer k is predicted from the previous layer with the help of the propagator:

$$\vec{x}_{k,proj} = \mathbf{F}_{k-1}\vec{x}_{k-1}. \tag{A.6}$$

The covariance matrix is similarly propagated to the current layer, taking the multiple scattering in the previous one into account:

$$\mathbf{C}_{k,proj} = \mathbf{F}_{k-1}\mathbf{C}_{k-1}\mathbf{F}^T_{k-1} + \mathbf{Q}_{k-1}. \tag{A.7}$$

For layers of material without a hit measurement, either because of detector inefficiency or because the layer does not contain active material, the propagation simply continues until a layer with a hit measurement is reached. The hit measurement in the layer refine covariance matrix and track parameters through the equations:

$$\mathbf{C}_k = \left[(\mathbf{C}_{k,proj})^{-1} + \mathbf{H}_k^T \mathbf{V}_k^{-1} \mathbf{H}_k\right]^{-1}, \tag{A.8}$$

$$\vec{x}_k = \mathbf{C}_k \left[(\mathbf{C}_{k,proj})^{-1} \vec{x}_{k,proj} + \mathbf{H}_k^T \mathbf{V}_k^{-1} \vec{m}_k\right]. \tag{A.9}$$

These equations are most easily understood in the limit of scalar state vectors and measurements, where the covariance matrices are given by σ^2. Equation (A.8-A.9) then reduce to:

$$\sigma_k^2 = \left(\frac{1}{\sigma_{k,proj}^2} + \frac{1}{\sigma_{measured}^2}\right)^{-1}, \tag{A.10}$$

$$\frac{x_k}{\sigma_k^2} = \frac{x_{k,proj}}{\sigma_{k,proj}^2} + \frac{m}{\sigma_{measured}^2}. \tag{A.11}$$

From these it is seen that the predicted track parameters and the hit measurement are weighted according to their uncertainty.

A.2 The smoothing process

In the SLT algorithm the tracking starts with tracks seeds downstream of the interaction point and proceeds upstream. At the end only the track parameters at the interaction point are needed and the SLT does not need to perform the smoothing process. For completion we will just state the equations needed for the smoothing process.

To calculate the smoothed parameters in the layer k, we define an auxiliary matrix $\mathbf{A}_k$:

$$\mathbf{A}_k = \mathbf{C}_k \mathbf{F}_k^T (\mathbf{C}_{k+1,proj})^{-1}. \tag{A.12}$$

This matrix is used to propagate the refined track back to the previous layer through the equations:

$$\mathbf{C}_{k,smooth} = \mathbf{C}_k + \mathbf{A}_k(\mathbf{C}_{k+1,smooth} - \mathbf{C}_{k+1,proj})\mathbf{A}_k^T, \tag{A.13}$$

$$\vec{x}_{k,smooth} = \vec{x}_k + \mathbf{A}_k(\vec{x}_{k+1,smooth} - \vec{x}_{k+1,proj}). \tag{A.14}$$

A.3 Kalman filter in one dimension

The SLT algorithms use Kalman filters for fitting tracks in one dimension at the (x or y coordinates) and we will therefore reduce the general equations given above to this specific case. As an example we will use a detector with several layers distributed along the z axis with only vertical wires, as in Figure A.1. There is no measurement of y, the hit measurement can be represented by a scalar x coordinate. The state vector, $\vec{x}_k$, is given by the track slope t_k and the intercept, b_k at $z = 0$. For the multiple scattering estimation, we need an external estimate of the slope in y, denoted by t_y, as the path length in the material and thus the amount of multiple scattering increases with the track slope. We use the intercept instead of the coordinate at layer k as this reduces $\mathbf{F}_k$ to the unit matrix, while the measurement matrix becomes:

$$\mathbf{H}_k = [z_k \;\; 1] \text{ with the state vector } \vec{x}_k = \begin{pmatrix} t_k \\ b_k \end{pmatrix}. \tag{A.15}$$

The covariance matrix for the measurement is reduced to a scalar: $\mathbf{V}_k = \sigma_k^2$.

To describe multiple scattering we treat the layers as a so-called "thin scatterers"[65]. For such scatterers the slope variance is given by:

$$\mathrm{var}(t,t) = \frac{\sigma_{MS}^2}{p^2}(1+t^2)(1+t^2+t_y^2)\sqrt{1+t^2+t_y^2}, \tag{A.16}$$

where σ_{MS} is the variance of the multiple scattering angle for a 1 GeV particle impacting perpendicularly on the material and given by Molière theory[13]:

$$\sigma_{MS} = \frac{13.6\,\mathrm{MeV}}{\beta c}\sqrt{l/X_0}\left[1+0.038\ln(l/X_0)\right], \tag{A.17}$$

where βc is the velocity, X_0 is the radiation length for the detector material and l is the thickness of the layer. As we use the intercept as the second coordinate in the state vector, the intercept is correlated with a change in slope and we obtain the following covariance matrix for the multiple scattering:

$$\mathbf{Q}_k = \mathrm{var}\{t,t\} \cdot \begin{pmatrix} 1 & -z_k \\ -z_k & z_k^2 \end{pmatrix}. \tag{A.18}$$

The filtering equations are now reduced to:

$$\vec{x}_{k,proj} = \vec{x}_{k-1}, \tag{A.19}$$

$$\mathbf{C}_{k,proj} = \mathbf{C}_{k-1} + \sigma_{MS}^2 \begin{pmatrix} 1 & -z_k \\ -z_k & z_k^2 \end{pmatrix}, \tag{A.20}$$

$$\mathbf{C}_k = \left[(\mathbf{C}_{k,proj})^{-1} + \frac{1}{\sigma_k^2}\begin{pmatrix} z_k^2 & z_k \\ z_k & 1 \end{pmatrix}\right]^{-1}, \tag{A.21}$$

$$\vec{x}_k = \mathbf{C}_k \left[(\mathbf{C}_{k,proj})^{-1}\vec{x}_{k,proj} + \frac{m}{\sigma_k^2}\begin{pmatrix} z_k \\ 1 \end{pmatrix}\right]. \tag{A.22}$$

These are the equations used by the SLT[2]. The tracking algorithms get their initial values from the previous algorithm and then iterate sequentially over the layers in the subdetector used by the algorithm. At each layer, the position and covariance matrix are predicted followed by the refinement of adding a hit. Pattern recognition is easily included into this algorithm. For each hit in a layer one checks if that hit lies at the predicted track position within errors. Several hits might be consistent or a layer might be inefficient and there are no hits. Typically several track candidates are constructed and prolongated to the next layer. As the number of track candidates will grow exponentially with the number of layers passed, some of the track candidates will have to be removed according to some quality criteria, which could be a χ^2 for how close all the hits are to the prediction or how many empty layers a candidate had. The exact procedure varies from algorithm to algorithm as it depends on efficiency, occupancy and other characteristics of the subdetector used.

[2]For some algorithms they are further simplified by completely ignoring multiple scattering.

References

[1] B. Aubert (BaBar Collaboration), hep-ex/0203007.

[2] K. Abe (Belle Collaboration), hep-ex/0107061.

[3] Y. Nir, "CP violation — A New Era", Lecture notes, 55th Scottish Universities Summer School in Physics, 2001.

[4] "The BaBar Physics Book", SLAC-R-504, 1998.

[5] E. Noether, Mathe. Ann. **77** (1916) 89.

[6] G. Lüders, Kong. Dansk. Vid. Selskab, Mat.-fys. Medd. **28** (1954) No. 5.

[7] V. Kostelecky and R. Potting, Phys. Rev. **D51** (1995) 3923.

[8] T. D. Lee and C. N. Yang, Phys. Rev. **104** (1956) 254.

[9] C. S. Wu et al., Phys. Rev. **105** (1957) 1413.

[10] J. H. Christenson et al., Phys. Rev. Lett. **13** (1964) 138.

[11] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49** (1973) 652.

[12] A. D. Sakharov, JETP Lett. **5** (1967) 24.

[13] D. E. Groom et al., Eur. Phys. J. **C15** (2000) 1.

[14] A. Hocker et al., hep-ph/0104062 and updates on `http://www.slac.stanford.edu/~laplace/ckmfitter.html`.

[15] A. Lai et al., Eur. Phys. J. **C22** (2001) 231.

[16] J. Graham, "A New Measurement of ϵ'/ϵ from KTeV", `http://kpasa.fnal.gov:8080/public/epsprime/ktev2001.ps`.

[17] S. Frixione et al., Adv. Ser. Direct. High Energy Phys. **15** (1998), 609.

[18] R.K. Ellis et al., "QCD and Collider Physics", Cambridge University Press, 1996.

[19] S. Frixione *et al.*, Nucl. Phys. **B431** (1994) 453.

[20] P. Nason, QCD at High Energy, Proceedings of the 12^{th} Reconcentre de Physique de la Valle d'Aoste, La Thuile, hep-ph/9811468.

[21] N. Kidonakis et al., Phys. Rev. **D64** (2001) 114001.

[22] R. Bonciani et al., Nucl. Phys. **B529** (1998) 424.

[23] I. Abt et al.(H1 Collaboration), Nucl. Instr. Methods **A386** (1997) 310

[24] Zeus Collaboration, "The ZEUS Detector", DESY-PRC-93-05, (1993).

[25] K. Ackerstaff et al.(HERMES Collaboration), Nucl. Instr. Methods **A417** (1998) 230.

[26] K. Ehret, Nucl. Instr. Methods **A446** (2000) 190.

[27] C. Bauer et al., Nucl. Instr. Methods **A453** (2000) 103.

[28] T. Zeuner, Nucl. Instr. Methods **A446** (2000) 324;
Y. Bagaturia et al., in press on Nucl. Instr. Methods A.

[29] M. Capeans, Nucl. Instr. Methods **A446** (2000) 317.

[30] J. Pyrlik, Nucl. Instr. Methods **A446** (2000) 299;
I. Ariño et al., Nucl. Instr. Methods **A453** (2000) 289.

[31] V. Saveliev, Nucl. Instr. Methods **A408** (1998) 191.

[32] G. Avoni et al., Proc. of the IX Conference on Calorimetry in Particle Physics, Annecy, France, October 9-14, 2000, Calorimetry in High energy physics, (2001) 777;
A. Zoccoli, Nucl. Instr. Methods **A446** (2000) 246.

[33] M. Buchler et al., IEEE Trans. on Nucl. Sci. **46** (1999) 126;
A. Arefiev et al., IEEE Trans. on Nucl. Sci. **48** (2001) 1059.

[34] V. Balagura et al., Nucl. Instr. Methods **A379** (1996) 404.

[35] C. Leonidopoulos et al., Nucl. Instr. Methods **A427** (1999) 436.

[36] G. Wagner, "Aufbau und Test der mit Digitalen-Signal-Prozessoren realisierten Komponenten des Datennahmesystems von HERA – B ", Ph.D. Thesis, University of Hamburg, 2000.

[37] T. Fuljahn et al., IEEE Trans. Nucl. Sci. NS-45 (1998) 1782.

[38] P. Kreuzer, Nucl. Instr. Methods **A462** (2001) 212.

[39] U. Schwanke, "Trigger and Reconstruction Farms in the HERA – B Experiment and Algorithms for a third level trigger", Ph.D. Thesis, Humboldt-University, 2000.

[40] J. M. Hernández, "PC Farms for Triggering and Online Reconstruction at HERA-B", Proc. of the International Conference on Computing in High Energy Physics CHEP 2001, Beijing, September 2001.

[41] M. Böcker et al., IEEE Trans. Nucl. Sci. **48** (2001) 1270.

[42] V. Alberico et al., Il Nuovo Cimento, **110 A** (1997) 1453;
G. Avoni et al., Nucl. Instr. Methods **A461** (2001) 332.

[43] D. Emeliyanov and I. Kisel, "CATS track fitting algorithm based on the discrete Kalman filter", HERA – B note 00-032.

[44] S. Brenner et al., Nucl. Instr. Methods **A339** (1994) 477.

[45] K. Berkhan et al., "Large System Experience with the ASD8 Chip in the HERA – B Experiment", Proceedings of the 5th Workshop on Electronics for the LHC, Snowmass 1999.

[46] J. Gläßet al., "Terabit per Second Data Transfer for the HERA – B First Level Trigger", Proceedings og the IEEE Conference on Realtime Systems, Valencia 2001.

[47] Analog Devices, "ADSP-2106x SHARC Users's Manual", 1995.

[48] J. M. Hernández, V. Rybnikov, D. Ressing, F. Sánchez, G. Wagner, "HERA-B Data Acquisition System", Proc. of the Nuclear Science Symposium and Medical Imaging Conference (IEEE, NPSS), Lyon, October 2000.

[49] F. Sanchez, Private communication.

[50] M. Villa, "The reconstruction for the Electromagnetic Calorimeter of the HERA – B experiment", Proceedings of the 7th International Conference On Calorimeters in High Energy Physics, World Scient., 1997.

[51] S. Xella, "HERA – B experiment: preparation and first results", Ph.D. Thesis, University of Copenhagen, 1999.

[52] S. Schmidt, "Second Level Trigger Studies at the HERA – B experiment",Ph.D. Thesis, University of Copenhagen, 2000.

[53] J. Carvalho, "Cross-section Compilation for the Measurement of the Luminosity at HERA – B ", HERA – B note 01-144.

[54] W. Hulsbergen, "Hit Efficiency of the Outer Tracker", HERA – B note 01-077.

[55] D. Emeliyanov, "OTR/ITR-CATS: Tracking Based on Cellular Automaton and Kalman Filter", HERA – B note 01-137.

[56] The HERA – B collaboration, "HERA – B Physics in 2001/2", Letter to the DESY Directorate, December 2000.

[57] The HERA – B collaboration, "Update of status and plans for 2001/2", HERA – B note 01-064.

[58] M. J. Leitch et al., Phys. Rev. Lett. **84** (2000) 3256.

[59] The HERA – B Collaboration, HERA – B Status Report, DESY-PRC-00-04, (2000).

[60] A. Abyzovet al., "Study of the Hit Efficiency of the OTR PC Chambers", HERA – B note 02-034.

[61] Y. Gilitsky and T. Kvaratskheliia, "New Muon Chamber Readout Electronics — Summary of Test Beam Results", HERA – B note 01-058.

[62] B. Schwingenheuer, "Electronics Commisioning Experience at HERA – B ", Proceedings of the 7th Conference on Electronics for LHC Experiments, Stockholm, 2001.

[63] V. Balagura, "MU1 Efficiency at FLT", HERA – B note 02-035.

[64] T. Nunez, Private communication.

[65] R. Mankel and A. Spiridonov, "A Pattern Recognition Algorithm for the HERA – B Main Tracking System, Part IV: The Object-Oriented Track Fit", HERA – B note 01-137.

[66] D. Kincaid and W. Cheney, "Numerical Analysis", Brooks/Cole Publishing Company (1991).

[67] H. Fleckenstein, "Acceptance and Efficiency of the First Level Trigger", HERA – B note 01-142.

[68] The HERA – B Collaboration, "Measurement of the $b\bar{b}$ production cross section at HERA – B with 2000 data sample", HERA – B note 02-005

[69] The HERA – B Collaboration, "Measurement of the $b\bar{b}$ Production Cross Section in 920 GeV Fixed-Target Proton-Nucleus Collisions", to be submitted to The European Physical Journal C.

[70] J. Bastos, "A statistical method for monitoring the luminosity of HERA – B data", HERA – B note 02-039.

[71] M. H. Schub et al., Phys. Rev. **D52** (1995) 1307.

[72] T. Alexopoulos et al., Phys. Rev. **D55** (1997) 3927.

[73] M. J. Leitch et al., Phys. Rev. Lett. **72** (1994) 2542.

[74] T. Alexopoulos et al., Phys. Lett. **B374** (1996) 271.

[75] T. Sjöstrand, Comp. Phys. Comm. **82** (1994) 74.

[76] H. Pi, Comp. Phys. Comm. **71** (1992) 173.

[77] R. Brun et al., GEANT3, Internal Report CERN DD/EE/84-1, CERN, 1987.

[78] J. Ivarsson *et al.*, "PYTHIA and FRITIOF: Event Generators for HERA-B", HERA – B note 99-067.

[79] P. Cho and A. Leibovich, Phys. Rev. **D53** (1996) 150.

[80] M. Mangano, P. Nason and G. Ridolfi, Nucl. Phys. **B373** (1992) 295;
P. Nason, S. Dawson and R. K. Ellis, Nucl. Phys. **B327** (1988) 49.

[81] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, hep-ph/0201127;

[82] J. Chay, S. D. Ellis and W. J. Stirling, Phys. Rev. **D45** (1992) 46.

[83] C. Peterson et al., Phys. Rev. **D27** (1983) 105.

[84] H. L. Lai et al., Eur. Phys. J. **C12** (2000) 375.

[85] A. Heister et al., Phys. Lett. **B512** (2001) 30.

[86] J. Chrin, Z. Phys. **C36** (1987) 163;
D. Decamp et al., Phys. Lett. **B244** (1990) 551;
P. Nason and C. Oleari, Nucl. Phys. **B565** (2000) 245.

[87] V. G. Kartvelishvili, A. K. Likhoded and V. A. Petrov, Phys. Lett. **B78** (1978) 615.

[88] G. J. Feldman and R. D. Cousins, Phys. Rev. **D57** (1998) 3873.

[89] D. M. Jansen et al., Phys. Rev. Lett. **74** (1995) 3118.

[90] T. Alexopoulos et al., Phys. Rev. Lett. **82** (1999) 41.

[91] B. Jones and B. Tompkins, "A Physicist's Guide to Kalman Filters",
`http://egret0.stanford.edu/bbjones/explainkalman.ps` .

[92] R. Frühwirth, Nucl. Instr. Methods **A262** (1987) 444.

Acknowledgments

I would like to thank my supervisor Prof. Jørn Dines Hansen for giving me the opportunity to work in HERA − B and to stay a full year at DESY. He has always provided me with useful advice and guidance, while at the same time letting me develop my own ideas and made it possible for me to develop as a physicist. Thanks goes to the whole HEP group at the Niels Bohr Institute for providing a really comfortable working environment. It has always been a great pleasure to be a student in the HEP group. Special thanks goes to Søren Schmidt for introducing me to the SLT and life in Hamburg and to Anders Wäänänen and Björn Nilsson for showing me the wonders of Linux and keeping me away from the dark path of Windows.

I also wish to thanks the present and past members of the SLT group at DESY, Mike Medinnis, Peter Kreuzer, Teresa Núñez, Lars Sözüer and Kendall Reeves for a pleasant and fruitful collaboration, which not only has made the SLT one of the most successful components of HERA − B, but also my stay at DESY a pleasurable time. One should also not forget the $\sigma(b\bar{b})$ analysis group: Benedetto Giacobbe, Peter Kreuzer, Maaijke Mevius, Mauru Villa and especially our physics coordinator Antonio Zoccoli. We have had many useful (and long) discussions, which resulted in the first physics paper to come from HERA − B. I will also never forget the many excellent dinners, where I learned to appreciate the Italian food and wine. Many more people in HERA − B deserves to be mentioned for helping to make my time in HERA − B educating and fun, but I am afraid I shall forget someone. I will therefore thank the whole collaboration with a special thanks to the members of the "lunch" (mahlzeit) and "dinner" (pappe) groups.

Finally I would like to thank the people who helped by reading parts of my thesis and providing me with many suggestions for improvements: Jeppe Andersen, Peter Kreuzer, Raluca Mureşan, Jakob Nielsen and Teresa Núñez. I am in particular indebted to Raluca for helping to correct my written English.